

Features - MySQL iDataAgent

TABLE OF CONTENTS

OVERVIEW

SYSTEM REQUIREMENTS - MYSQL /DATAAGENT

INSTALLATION

- Install the MySQL iDataAgent -Unix

BACKUP - MYSQL

RESTORE DATA - MYSQL

CONFIGURATION

- Instances - MySQL
- Subclients - MySQL

MANAGEMENT

- Backup Job History
- Restore Job History

FREQUENTLY ASKED QUESTIONS

Overview - MySQL iDataAgent

Choose from the following topics:

- Introduction
- Supported Data Types
- Tree Levels in the MySQL iDataAgent
- License Requirement
- User Privileges
- Client Terminology
- Instances, Databases, and Subclients
- Securing MySQL Application Data Using the File System iDataAgent
- Transaction Logs
- Disaster Recovery Considerations

INTRODUCTION

The MySQL iDataAgent is responsible for backing up and restoring MySQL servers that comprises of user created MySQL databases. The MySQL Server system databases (information_schema and mysql database) contain information that the MySQL Server needs to operate, such as user database characteristics, permissions, grants, and other information. The system generated mysql database is backed up and restored using the File System iDataAgent. For more information on the types of data in system databases, refer to your MySQL documentation. MySQL Server user databases contain your enterprises data, related administrative files and transaction logs. The MySQL iDataAgent supports two types of MySQL Servers: servers containing few databases of sizes larger than 200GB and servers containing large number of databases of smaller sizes (ranging between 3GB to 10GB).

SUPPORTED DATA TYPES

Data Protection Operations for the following data types are supported by the MySQL iDataAgent:

- MySQL system databases
- MySQL user databases
- Transaction logs

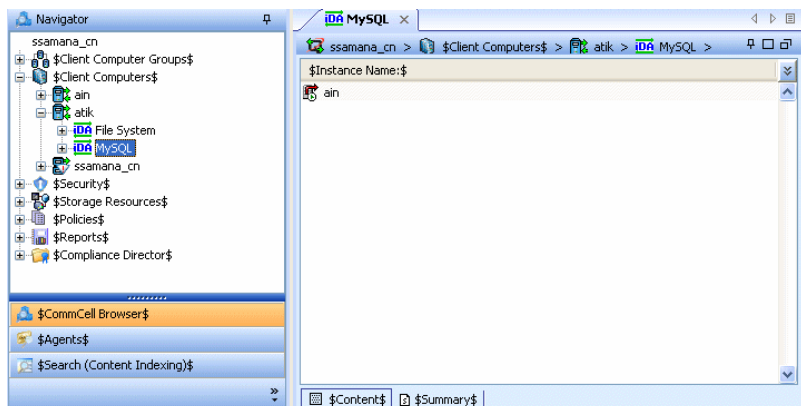
Data Protection Operations for all other data types not mentioned in the above list are not supported by the MySQL iDataAgent.

- MySQL application files
- Operating system data
- System generated MySQL database

The above data types can be secured by the appropriate File System iDataAgent.

TREE LEVELS IN THE MYSQL /DATAAGENT

When the iDataAgent is installed, the following levels are automatically created in the CommCell Browser.



atik: Client

MySQL: Agent

<user_defined_instance>: Instance



In order to run the first backup, instances must be added. Each instance creates a default subclient.

[Back to Top](#)

LICENSE REQUIREMENT

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

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

USER PRIVILEGES

MySQL users should have special privileges to perform backups and recovery.

For Backup:

As a MySQL user, you should have the following permissions to perform backups:

SHOW DATABASES, SUPER, RELOAD,SELECT, LOCK

Example Query to grant the required permissions for backup is as follows:

```
mysql> GRANT SHOW DATABASES, SELECT, LOCK TABLES, RELOAD,SUPER ON *.* to <backup_agent username>@localhost IDENTIFIED BY '<backup_agent password>';
```

However, it is recommended to grant all Database Administrator privileges to perform backup operations too.

Example:

```
mysql> GRANT ALL PRIVILEGES ON *.* TO <backup_agent username>@localhost IDENTIFIED BY '<backup_agent password>';
```

For Restore:

You should have full Database Administrator privileges to perform restore operations:

Example Query to grant the required permissions for restore is as follows:

```
mysql> GRANT ALL PRIVILEGES ON *.* TO <restore_agent username>@localhost IDENTIFIED BY '<restore_agent password>' WITH GRANT OPTION;
```

CLIENT TERMINOLOGY

The following definitions are offered to clarify the term 'client':

- MySQL clients are computers that access MySQL Servers but are not themselves MySQL servers. Both MySQL application binary files and file system data on these computers can be backed up using the File System iDataAgent.
 - A client is a computer whose data is backed up via an iDataAgent. This may include MySQL, File System, etc.
-

INSTANCES, DATABASES, AND SUBCLIENTS

For the MySQL iDataAgent, each MySQL instance corresponds to a MySQL server. The MySQL server may contain many databases. Multiple instances can be created for the MySQL iDataAgent, where each instance represents a single MySQL server. It is recommended to not assign a MySQL server to more than one instance.

MySQL instances must be created manually. However, for each instance that you create, a default subclient is automatically created. Each instance can support multiple subclients. Initially, each default subclient contains all of the MySQL databases for the instance. Thereafter, you can create one or more user-defined subclients to suit your backup needs. The content of a subclient can contain one or more MySQL databases.

Instances have various configuration options, and these options affect the associated subclients. For example, if **Enable Auto Discovery** is selected, newly created databases will be automatically added to the default subclient. This and other options are accessed using the Instance Properties dialog box.

SECURING MYSQL APPLICATION DATA USING THE FILE SYSTEM iDATAAGENT

In addition to the MySQL Server, there may be MySQL application files stored on the computer (that is, the client, as defined above). Such data is not backed up by the MySQL iDataAgent. To secure such data from a MySQL client, you must back it up using the appropriate iDataAgent for the computer's file system.

TRANSACTION LOGS

A transaction log is a file containing a record of changes to a database.

Truncation of transaction logs follows these rules for the MySQL iDataAgent:

- Transaction logs will be truncated at the completion of a Transaction Log backup.
 - Full backups truncate transaction logs.
-

DISASTER RECOVERY CONSIDERATIONS

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

[Back to Top](#)

System Requirements - MySQL iDataAgent

System Requirements | Supported Features

The following requirements are for the MySQL iDataAgent:

APPLICATION/OPERATING SYSTEM		PROCESSOR
MYSQL COMMUNITY SERVER EDITION OR ENTERPRISE, STANDARD OR CLASSIC EDITION - 5.0.26 AND HIGHER, 5.1.X, 5.5.X AND 5.6.X ON:		
AIX	AIX 5.3 with technology level 6 (or higher) and runtime library xIc.rte 8.0.0.0 or higher	Power PC (Includes IBM System p)
FREEBSD	FreeBSD 8.x	Intel Pentium, x64 or compatible processors
HP-UX	HP-UX 11i v3 (11.31)	Itanium
	HP-UX 11i v3 (11.31)	PA-RISC
	HP-UX 11i v2 (11.23)	PA-RISC
LINUX	DEBIAN	
	Debian 6.x with glibc 2.11.x	Intel Pentium, x64 or compatible processors
	Debian 5.x	Intel Pentium, x64 or compatible processors
	GENTOO	
	Gentoo release 10.1 with glibc 2.6.x	Itanium, x64 or compatible processors
	OPENSUSE	
	OpenSuSE 12.1 with glibc 2.14.x	x64
	OpenSuSE 11.1 with glibc 2.9.x	Itanium, x64 or compatible processors
	RED HAT ENTERPRISE LINUX/CENTOS	
	Red Hat Enterprise Linux/CentOS 6.x with glibc 2.12.x	x64
	Red Hat Enterprise Linux/CentOS 5.x with glibc 2.5.x	Intel Pentium, x64 or compatible processors
	Red Hat Enterprise Linux/CentOS 4.x with a minimum of glibc 2.3.4	Intel Pentium, x64 or compatible processors
SOURCE MAGE LINUX		
Source Mage Linux	Itanium, x64 or compatible processors	
SUSE LINUX (SLES)		
SuSE Linux 11.x with glibc 2.9.x and above	Intel Pentium, x64 or compatible processors	
SuSE Linux 10.x with glibc 2.4.x	Intel Pentium, x64 or compatible processors	
UBUNTU		
Ubuntu 9.x	Intel Pentium, x64 or compatible processors	
Ubuntu 8.04	Intel Pentium, x64 or compatible processors	
Ubuntu 12.04 LTS	Intel Pentium, x64 or compatible processors	
Ubuntu 10.04 LTS	Intel Pentium, x64 or compatible processors	
SOLARIS	Solaris 9	Sparc5 or higher recommended
	Solaris 11.x	x64, Sparc5 (or higher recommended)
	Solaris 10.x with a minimum of SunOS (Sparc) Patch 119963-14	Sparc5 or higher recommended
	Solaris 10.x	x64
WINDOWS	WINDOWS 2008	

Microsoft Windows Server 2008 32-bit and x64 Editions*

All Windows-compatible processors supported

*Core Editions not supported

WINDOWS 2003

Microsoft Windows Server 2003 32-bit and x64 Editions with a minimum of Service Pack 1

All Windows-compatible processors supported

CLUSTER - SUPPORT

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

For information on supported cluster types, see Clustering - Support.

HARD DRIVE

220 MB minimum of hard disk space for software

50 MB of additional hard disk space for log file growth

725 MB of temp space required for install (where the temp folder resides).

MEMORY

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

Swap space = 2*RAM size

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.

PERIPHERALS

DVD-ROM drive

Network Interface Card

MISCELLANEOUS

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

On Ubuntu 9.x 64-bit computers, it is required that you install the libstdc++5 compatibility package. To install the package, use the following command:

```
[root]# sudo apt-get install libstdc++5
```

On Ubuntu and Debian computers, it is recommended that you install the latest packages for the operating system before installing the iDataAgent. To install the latest packages, use the following command:

```
[root]# apt-get upgrade
```

NETWORK

TCP/IP Services configured on the computer.

SELINUX

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

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

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

File Name	Content of the File
<code><directory>/<file_name>.te</code>	The content of the file should be as follows:
where:	<code>policy_module(<name>,<version>)</code>

<pre><directory> is /usr/share/selinux/devel</pre> <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>	<pre>#####</pre> <p>where:</p> <p><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.</p> <p><version> is the version of the policy module. It can be any number, such as 1.0.0.</p> <p>For Example: While creating a policy module for the backup_IDA application, you can use the following content.</p> <pre>policy_module(backup_IDA,1.0.0)</pre>
<pre><directory>/<file_name>.fc</pre> <p>where:</p> <pre><directory> is /usr/share/selinux/devel</pre> <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

Install the MySQL iDataAgent - Unix

TABLE OF CONTENTS

Where to Install

Install Requirements

Before You Begin

Install Procedure

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

Post-Install Considerations

WHERE TO INSTALL

INSTALLING ON A GLOBAL ZONE OR UNIX SERVER

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

INSTALLING ON A NON-GLOBAL ZONE

Install the software on each of the non-global zones where you have application data.

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

- System Requirements - MySQL iDataAgent
- 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

It is recommended to install the software on the global zone to protect non-changing or static data on non-global zones. If the data is dynamic or contains application data, install the software on the non-global zone.

INSTALL REQUIREMENTS

The following procedure describes the steps involved in installing the MySQL iDataAgent and the Unix File System iDataAgent. The MySQL iDataAgent is installed on the computer on which the MySQL application resides. (This computer is referred to as the *Client* computer in this install procedure.)

Review the following Install Requirements before installing the software:

GENERAL

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

PACKAGE MANAGEMENT SYSTEMS

You can use the Red Hat Package Manager (RPM) package management system to install all Linux versions of the agent. Also, you can use the Advanced Packaging Tool (APT) package management system to install all Debian Linux versions of the agent. For step-by-step instructions, see Base Software Install for RPM Package Manager and Base Software Install for APT Package Manager.

RED HAT LINUX

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

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


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

```
192.168.1.111 bluesky localhost
```

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

```
127.0.0.1 localhost
```

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

```
192.168.1.111 bluesky
```

BEFORE YOU BEGIN

- The MySQL application has been installed on the client. Ensure that the `expire_log_days` parameter is set to 0 on the MySQL Server.
- Log on to the client as `root`.
- The install package requires `super-user` permissions to execute.

INSTALL PROCEDURE

GETTING STARTED

1. If you are installing the software on a global zone or Unix Server, do the following:
Place the software installation disc for the Unix platform into the disc drive.

You can also install the product using a disc drive mounted on another computer on the network.

- On Solaris, double-click the **cvpkgadd** program from the File Manager window.
- On other Unix platforms, open the Terminal window, navigate to the software installation disc and then enter **./cvpkgadd**.

If you are installing the software on a non-global zone, do the following:

- Mount the installation disc on the non-global zone.

```
mkdir <Non-Global Zone root location>/<Non-Global Zone local directory>
```

```
mount -F lofs <Global zone software Install Disc mount point> <Non-Global Zone root location>/<Non-Global Zone local directory>
```

Connect to Non-Global Zone terminal

- Run the following command from the Software Installation Disc:

```
./cvpkgadd
```

2. Read the license agreement. Type **y** and press **Enter** to continue.
3. Enter the number corresponding to the setup task you want to perform.

NOTES

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

To create a custom package and for record and play setup, follow the steps described in Custom Package - Unix.

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

4. The product banner and other information is displayed.

Press **Enter** to continue.

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

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

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

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

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

Your choice: [1]

4. The product banner and other information is displayed.

Press **Enter** to continue.

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

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

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

- 1) All platforms (32 bit)
- 2) FS and MA only (64 bit)

Your choice: [1]

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

Press **Enter** to continue

NOTES

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

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

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

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

NOTES

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

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

NOTES

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

Verify the default network interface name.

Press **Enter** to accept the default network interface name and continue, or Type the default network interface name, and then press **Enter** to continue.

NOTES

- This prompt will be displayed only when you have multiple network interfaces for the computer.

8. Specify the client name for the computer.

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

and activating the IP address/services on the other server.

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

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

- 1) Install on a physical machine
- 2) Install on a virtual machine
- 3) Exit

Your choice: [1]

Network interface with the following IPs have been found available on your system. One of these interfaces should be associated with the physical machine being installed. It will also be used by the CommServe to initiate connections to the physical machine. Note that you will be able to additionally customize Datapipe Interface Pairs used for the backup data traffic later in the Calypso Java GUI.

Please select the correct network interface below.
 1) client (201.42.33.598)
 2) hk97:489:9glg:hk8d:9490
 3) client.company.com (hr90:8842:2:78:013:8ghh:hg8k:9x54)

Interface number: [1]

Network interfaces with the following IPs have been found configured on your system. Please select one to be used by the CommServe to contact the client.

- 1) dunk.company.com (480.19.65.333)
- 2) 480.19.66.326

Interface number: [1]

Please verify the physical machine interface name below. Make it as complete (with fully qualified domain name) as possible:

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

Please specify the client name for this machine.

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

Physical Machine Client name: [angel]

SELECT COMPONENTS FOR INSTALLATION

9. Enter the number corresponding to the **CVGxMySQL** module.

NOTES

- Your screen may look different from the example shown.
- Components that either have already been installed, or which cannot be installed, will not be shown.

Press **Enter** to continue.

Install Calypso on physical machine client.company.com

Select the Calypso module that you would like to install

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

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

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

**BASE SOFTWARE INSTALLATION**

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

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

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

**NOTES**

- Select **Download from the software provider website** to download the latest software scripts from your software provider website.  
 Make sure you have internet connectivity when you are using this option.
- Select **Use the one in the installation media**, to install the software scripts from the disc or share from which the installation is performed.
- Select **Use the copy I already have by entering its unix path**, to specify the

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

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

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

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

path if you have the software script in an alternate location.

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

**NOTES**

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

Press **Enter** to continue.

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

**NOTES**

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

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

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

14. Specify the location for the log files.

**NOTES**

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

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

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

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

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

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

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

Press **Enter** again to continue.

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

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

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

**NOTES**

- For more information about Network TCP Ports, see Network TCP Port Requirements.
- For more information about these services, see Services.
- If the port number you entered already exists, a message will be displayed Port

Your choice: [1] 2

Keep Your Install Up to Date - Latest Service Pack

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

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

Do you want to download the latest service pack now? [no]

Press <ENTER> to continue ...

Please specify where you want us to install Calypso binaries.

It must be a local directory and there should be at least 209MB of free space available. All files will be installed in a "calypso" subdirectory, so if you enter "/opt", the files will actually be placed into "/opt/calypso".

Installation Directory: [/opt]

..

Calypso will be installed in /opt/calypso.

Press ENTER to continue ...

Please specify where you want to keep Calypso log files.

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

Log Directory: [/var/log]

..

Calypso log files will be created

in /var/log/calypso/Log\_Files.

Press ENTER to continue ...

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

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

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

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

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

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

Group name: dba

REMINDER

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

Press <ENTER> to continue ...

Every instance of Calypso should use a unique set of network ports to avoid interfering with other instances running on the same machine.

The port numbers selected must be from the reserved port number range and have not been registered by another application on this machine.

Please enter the port numbers.

Port Number for CVD : [8600]

Port Number for EvMgrC: [8602]

#### is already reserved in /etc/services. To work around this issue, enter different port number.

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

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

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

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

#### NOTES

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

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

#### NOTES

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

Click **Enter** to continue.

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

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

CommServe Host Name:

Enter your CommCell user name and password:

User Name :

Password :

Press <ENTER> to continue ...

## KERNEL PARAMETERS

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

#### NOTES

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

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

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

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

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

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

Number of streams: [10]

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

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

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

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

Changes saved into /etc/system.gal.1744

Press <ENTER> to continue.

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

#### NOTES

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

Press **Enter** to continue.

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

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

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

## ENABLE GLOBAL FILTERS

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

### NOTES

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

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

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

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

## CLIENT GROUP SELECTION

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

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

### NOTES

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

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

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

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

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

26. Press **Enter** to continue.

### NOTES

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

+-----+

IMPORTANT:

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

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

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

+-----+

## STORAGE POLICY SELECTION

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

### NOTES

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

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

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

Storage Policy: [3]

28. If you have multiple instances of the database, ensure that the items listed are satisfied. Press **Enter**.

IMPORTANT

If you are planning to use Calypso <package> with several instances of the database, each running on behalf of different UNIX users, you must make sure that:

- 1) All database users have been added to each of the database groups involved. For example, if there is a "db2inst1" user from group "dba1" and "db2inst2" user from group "dba2", then both "db2inst1" and "db2inst2" should be included into "dba1" and "dba2" groups. On most UNIXes this can be accomplished by editing the /etc/group file.

- 2) The .profile of each database user must contains a "umask 002" command at the end, which will guarantee that

- all files created by one database user can be open for writing by other database users.
- 3) All database users should be added to the "dba" group that you have chosen for Calypso at the time of CVGxBase install.
- If all of the above requirements are not fulfilled, Calypso will run into various "Permission denied" errors, and backups will not work.
- Press <ENTER> to continue ...
29. The install program now starts copying the software to the computer. The progress of the operation is displayed.
30. Press **Enter** to continue.
31. This prompt is displayed only when you are installing on HP-UX, Linux, or Solaris computers. Enter the number corresponding to the **Exit** option and then press **Enter** to continue.
- The installation is now complete.
- ```

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

Successfully copied xx files

.....
.....

Successfully installed <package_name>.

Press ENTER to continue ...

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

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

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

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

Your choice: [1]

```

POST-INSTALL CONSIDERATIONS

GENERAL

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

Backup - MySQL

Topics | How To | Related Topics

Overview

Supported Backup Types

MySQL Backup Jobs

Backup Considerations

Advanced Backup Options

OVERVIEW

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

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

SUPPORTED BACKUP TYPES

This agent supports the following backup types:

- Full Backups
- Transaction Log Backups
- Do Not Truncate Log

The state of the databases must be online to ensure successful data protection operations.

MYSQL BACKUP JOBS

This agent has the following unique functionality. The MySQL iDataAgent backs up very specific sets of data depending upon what type of operation you choose. Review the following information to understand exactly what will be backed up when you perform a backup job at the indicated level.

FULL BACKUPS

MySQL full backups are database backups and does not include the traditional logs. For detailed information on full backups, see Full Backups.

TRANSACTION LOG BACKUPS (INCREMENTAL BACKUPS)

A transaction log backup is a backup of the transaction log of the database. This backup captures the transaction log which contains a record of all transactions whether the transaction was committed or not. Transaction log backups are consistent to the start time of the backup. Binary log must be enabled to perform transaction log backup.

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

ENABLING THE BINARY LOG AT MYSQL SERVER

Add the following under [mysqld] section in MySQL Configuration file (my.cnf for UNIX and my.ini for WINDOWS) to enable the Binary Logs:

```
[mysqld]
log-bin=mysql-bin.log
```

Restart the MySQL server for the changes to take effect. After restarting the MySQL Server, you can verify binary logging using the following query:

```
mysql > show variables like "%log_bin%".
```

If it is enabled then the value of 'log_bin' will be **ON** otherwise the value will be **OFF**.

IMPORTANCE OF BINARY/TRANSACTION LOGS

The binary/transaction logs are important for MySQL iDataAgent. However, they are not mandatory. The use of transaction log backups makes point in time

recovery possible. You can restore to any point in time within the transaction log. If you use only full and backups, you will be able to restore to the time of the backup, but not to a point in time between backups. It is your choice to enable (ON) or disable (OFF) the transaction logs.

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

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

BACKING UP THE TRANSACTION LOG OF A DAMAGED DATABASE (DO NOT TRUNCATE LOG)

In the scenario of a database failure where it is unacceptable to lose any data, you will want to restore to the point of failure. To do so means you will need to capture all transaction log events that occurred after the last backup job was run.

It is possible to log performance counters for successful backups about time taken and bytes read/write for each backup and the time taken for buffer allocation [MySQLBackup.log (Windows) and MySQLBackupChild.log (Unix)]. The counters can be enabled by setting the following Registry Key:

- nPerformanceLog

The Performance Counters appear as follows:

```
3900 4b0 03/22 11:53:42 48 ::MySQLBackupDb::DoBackupMySQL() - Performance counter enabled from Registry
3900 4b0 03/22 12:02:05 48 ::MySQLBackupDb::DoBackupMySQL() - Performace for Database auto_4 is Bytes= 9406567384, Time= 190.661203
Sec(s), Average Throughput= 165.413699 GB/h
```

MYSQL DATABASE BACKUP

The MySQL database is a system database that gets created when the MySQL Server application is installed. It contains all the meta data information, such as database access permissions, user information etc., that is used to manage the MySQL Server. Hence it is important that you take frequent backups of this database.

It is mandatory that you backup the MySQL database every time a new user is created or whenever this database is modified. These backups can later be used during disaster recovery situations when you need to restore the MySQL database.

MySQL dump uses the default value of max_allowed_packet variable while creating dump from MySQL server.

BACKUP CONSIDERATIONS

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

- Make sure to provide correct Log data directory path (path of the Transaction logs location, when MySQL Server Log_bin is ON) while configuring an Instance for MySQL iDataAgent.
- When you add new databases or add new content to the user defined subclients that are already created for MySQL iDataAgent and run a log backup, the system will detect the newly added databases or content, updates the Index, and performs a data backup after the log backup.
- For default Subclient, you must ensure that the Enable Auto Discovery option in the Instance General Properties dialog box is selected. This option is required to update the default subclient with the newly added database(s).
- Make sure to perform a full backup before running an incremental backup or restore.
- When you backup a default subclient, if some of the auto discovered databases are removed/missing from the MySQL server, the backup job for the default subclient still gets completed successfully. However, if manually discovered databases of the default or user-defined subclients are removed/missing from the MySQL server, then the specific backup job will complete with errors.
- If you perform a backup of the default subclient without changing its default properties, the backup would include all the databases for the MySQL server instance, including any new objects that may have been created. However, if you have created subclients for database, then the default subclient will exclude those databases in other subclients during the backup operation.

If you have created subclients for a database, you can create a complete backup of the database by backing up these subclients. If you have created a group of selected subclients for the database, you must back up all of these subclients to create a complete backup. If you have created separate subclients for the log files, and if these subclients are not included with the subclients defined for databases or subsets, then the subclients for these files must also be backed up to create a complete backup.

- Multiple streams are supported for data protection and recovery operations with this agent; however, before utilizing multiple streams, you should understand how their use affects operations. For more information, see Streams.
- Backup and restore operations support utf-8 based multilingual databases (including db names).
- When you backup two subclients with different data storage policy but same log storage policy, the data backup phase will run simultaneously for both the subclients. However, during the log backup phase, you can run only one log backup at a time. The log backup for the second subclient will wait till the first log backup operation is completed.

- Prior to running a log backup of MySQL databases, ensure that the logs are not pruned by the MySQL Server automatically. To do this, set the `expire_log_days` parameter to 0 on the MySQL server.
- The subclient which has "mysql" system database in its content will perform data phase in Log Backups to protect the latest system database. Enable `bExcludeSysDBFromLogJob` Registry key to disable this behavior.
- When you perform an incremental log backup on a default subclient which contains system database, the browse window for such log backup should display the MySQL system database only. If a new database is auto discovered or added after a full backup, then the browse window for such log backup will display this newly added database along with the system database.

[Back to Top](#)

Backup - MySQL - How To

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

[Start a Backup](#)

[Schedule Backups](#)

[Start a Transaction Log Backup](#)

[Start a Backup in the Suspended State](#)

[Start a Backup on New Media](#)

[Start a Backup that Marks Media Full on Completion](#)

[Start a Backup with a Set Job Priority](#)

[Start a Backup with Vault Tracking enabled](#)

START A BACKUP

Before You Begin

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

Required Capability: See Capabilities and Permitted Actions

▶ To start an immediate backup job:

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

In certain circumstances a non-full backup may automatically be converted to a full backup. For a listing of these circumstances, see [When a Non-Full Backup is Automatically Converted to a Full Backup](#).

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

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

- You can also run backups of the following:
 - For a user-defined backup set or instance, right-click the backup set you want to back up, click **All Tasks**, and click **Backup All Subclients**.
 - For the Lotus Notes Document iDataAgent, to back up a partition, right-click the partition you want to back up, click **All Tasks**, and click **Backup Default Backup Set**.

For the Lotus Notes Database iDataAgent, to back up a partition, right-click the partition you want to back up, click **All Tasks**, and click **Backup All Subclients**.

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

SCHEDULE BACKUPS

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

Before You Begin

- **All Agents**

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

Required Capability: See Capabilities and Permitted Actions

▶ To schedule a backup operation:

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

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

START A TRANSACTION LOG BACKUP

Use this procedure to backup Transaction logs.

SQL Server Transaction Log Backups can be automatically performed immediately after a full or differential backup by selecting the **Start Log Backup after Successful Backup** advanced backup option.

Before You Begin

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

Required Capability: See Capabilities and Permitted Actions

▶ To start an immediate backup job:

1. From the CommCell Browser, select one of the following:
 - To backup a subclient, right-click the subclient to want to backup and click **Backup**.
 - To backup a user-defined backup set or instance, right-click the backup set you want to backup, click **All Tasks**, and click **Backup All Subclients**.
 - To backup the default backup set, right-click the agent or instance node, click **All Tasks**, and click **Backup Default Backup Set**.
2. If you chose a level higher than subclient (i.e., backup set, etc.), you are prompted to confirm that you want to back up all the subclients below that level/node. Click **Yes**.
3. From the Backup Options dialog box, select **Run Immediately**.
4. Select **Transaction Log** backup.
In certain circumstances a non-full backup may automatically be converted to a full backup. For a listing of these circumstances, see [When a Non-Full Backup is Automatically Converted to a Full Backup](#).
5. Click **OK**. You can track the progress of the backup job from the Job Controller window.
6. 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.
7. When the backup has completed, Job Controller displays *Completed*.

NOTES

• All Agents

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

• Microsoft SQL iDataAgents

- When File/File Group subclient jobs are scheduled and the instance property **Use VSS** is selected, backups for File and File Group subclients run successfully as scheduled.
- In order to perform a Transaction Log backup the following is required:
 - SQL Server database must be in full or bulk-logged recovery mode.
- When running a transaction log backup through the CommCell Console, the system default is to truncate the log. To override the default, select **Do not truncate log** from the Backup Options dialog box.
When running a transaction log backup through the command line, the system default is to not truncate the log. To override the default, run the backup job with an `-af` input file containing the argument file [options] QB_TRUNC_LOG.

START A BACKUP IN THE SUSPENDED STATE

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

Before You Begin

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

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

▶ To start an immediate backup job with advanced backup options:

1. From the CommCell Browser, select one of the following:
 - To backup a subclient, right-click the subclient to want to backup and click **Backup**.
 - To backup a user-defined backup set or instance, right-click the backup set you want to backup, click **All Tasks**, and click **Backup All Subclients**.
 - To backup the default backup set, right-click the agent or instance node, click **All Tasks**, and click **Backup Default Backup Set**.
 - For Lotus Notes iDataAgent, to backup a partition, right-click the partition you want to backup, click **All Tasks**, and click **Backup Default Backup Set**.
 - For Agents that do not have backup set or instance levels, to backup all subclients, right-click the agent icon, click **All Tasks**, and click **Backup All Subclients**.
2. If you chose a level higher than subclient (i.e., backup set, etc.), you are prompted to confirm that you want to back up all the subclients below that level/node. Click **Yes**.
3. From the Backup Options dialog box, select **Run Immediately**.
4. Select the type of backup that you want to initiate.

In certain circumstances a non-full backup may automatically be converted to a full backup. For a listing of these circumstances, see [When a Non-Full Backup is Automatically Converted to a Full Backup](#).

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

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

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

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

START A BACKUP ON NEW MEDIA

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

Before You Begin

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

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

▶ To start an immediate backup job with advanced backup options:

1. From the CommCell Browser, select one of the following:
 - To backup a subclient, right-click the subclient to want to backup and click **Backup**.
 - To backup a user-defined backup set or instance, right-click the backup set you want to backup, click **All Tasks**, and click **Backup All Subclients**.
 - To backup the default backup set, right-click the agent or instance node, click **All Tasks**, and click **Backup Default Backup Set**.
 - For Lotus Notes iDataAgent, to backup a partition, right-click the partition you want to backup, click **All Tasks**, and click **Backup Default Backup Set**.
 - For Agents that do not have backup set or instance levels, to backup all subclients, right-click the agent icon, click **All Tasks**, and click **Backup All Subclients**.
2. If you chose a level higher than subclient (i.e., backup set, etc.), you are prompted to confirm that you want to back up all the subclients below that level/node. Click **Yes**.
3. From the Backup Options dialog box, select **Run Immediately**.
4. Select the type of backup that you want to initiate.

In certain circumstances a non-full backup may automatically be converted to a full backup. For a listing of these circumstances, see [When a Non-Full Backup is Automatically Converted to a Full Backup](#).
5. Click the **Advanced** button to open the **Advanced Backup Options** dialog box.
6. Click on the Advanced Backup Options (Media) tab, and select **Start new media** and click **OK**.

If you would like jobs with other Job IDs to use this new media, also select the **Allow other schedule to use media set** option.
7. From the **Backup Options** dialog box, click **OK**. You can track the progress of the backup job from the Job Controller window.
8. If you are using a stand-alone drive, you are prompted to load a specific cartridge into the drive. If you are using a library, you will not receive this prompt. The system loads the tapes automatically.

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

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

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

START A BACKUP THAT MARKS MEDIA FULL ON COMPLETION

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

Before You Begin

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

Required Capability: See Capabilities and Permitted Actions

▶ To start an immediate backup job with advanced backup options:

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

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

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

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

START A BACKUP WITH A SET JOB PRIORITY

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

Before You Begin

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

Required Capability: See Capabilities and Permitted Actions

▶ To start an immediate backup job with advanced backup options:

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

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

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

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

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

START A BACKUP WITH VAULT TRACKING ENABLED

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

For additional information, see the following:

- VaultTracker
- VaultTracker Enterprise

Before You Begin

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

Required Capability: See *Capabilities and Permitted Actions*

▶ To start an immediate backup job with advanced backup options:

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

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

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

Restore Data - MySQL

Topics | How To | Full System Restore | Related Topics

Overview

Restore Considerations for this Agent

Restore Destinations

- In-Place Restore
- Out-of-Place Restore
- Cross-Platform Restores

Restoring Database to a Point in Time

OVERVIEW

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

The MySQL iDataAgent supports the following types of restores:

- Restoring and recovering entire MySQL server
- Restoring and recovering individual MySQL databases
- Restoring transaction logs

MySQL databases can be restored to the same instance or to different instances on the current machine. The databases can also be saved to the same instance on a destination machine (cross-machine restore).

You can restore MySQL system database. Selecting MySQL database from browse and restore converts restore into full instance restore. MySQL dump uses the default value of `max_allowed_packet` variable while creating dump from MySQL server. If you want to specify the value of `max_allowed_packet` explicitly for dump/restores, you can set the value in MySQL iDataAgent Registry keys.

Additionally, the MySQL iDataAgent supports:

- Restoring and recovering the MySQL database online. (See Recover Data with the Database Online for more information)
- Restoring and recovering the transaction logs online. (See Recover Logs with the Database Online for more information)
- Restoring and recovering the transaction logs offline. (See Recover Logs with the Database Offline for more information)
- Recurring restores of Log files, which allows you to replicate data from one server to another to be used in Fail over cases. (See Recurring Restores of Logs for more information)
- Point-in-time restores, which allow you to restore data from the full backup created at or before the point in time based on the date and time that you specify. (See Restore Data to a Point in Time for more information)

To enable the non-FULL Recurring Restore job, use `bFreshRecRestore` Key. This would clear the earlier history. If you do not enable this Registry Key, the earlier history will be cleared only for the instances with FULL restores.

For the MySQL iDataAgent, restore operations can be performed from the client, agent and instance levels in the CommCell Browser.

When there is a problem with the MySQL database or the client system (operating system, hardware, hard drives, etc.), full system restores may be required. See Restore Data - MySQL - Full System Restore for more information.

RESTORE CONSIDERATIONS FOR THIS AGENT

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

- Review the general restore requirements described in What You Need to Know Before Performing a Restore.
- Backup and restore operations do not support non-English database names.

It is recommended not to drop MySQL database before performing a restore operation. MySQL database should be running when you perform restore operations.

RESTORE DESTINATIONS

By default, the MySQL 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 the MySQL iDataAgent. See Restore/Recover/Retrieve Destinations - Support for a list of Agents supporting each restore destination type.

IN-PLACE RESTORE

- Same path/destination

OUT-OF-PLACE RESTORE

Keep in mind that out-of-place restore operations are also subject to the conditions described in Cross-Platform Restores.

- Same path/destination

CROSS-PLATFORM RESTORES

The following types of cross platform restores are supported:

Same Operating System - Different Version

Restoring data associated with different versions of the same operating system is supported. You can also have different versions of MySQL in one machine. You can restore the data across these same/ different versions too. For example, you can restore data secured from a Solaris 9 to Solaris 10.

Different Operating System - Same/Different Version

You can restore data across different platforms too. You can restore data across different machines with different operating systems running same/different versions of MySQL server. However, you should be aware of how the data is restored under certain circumstances.

Same Machine - Multiple servers/instances

You can have several instances of MySQL databases in one server or multiple servers. Restoring data across these instance is also supported.

See Cross-Platform Restores from CommCell Console for a step-by-step procedure to perform a cross platform restores.

RESTORING DATABASE TO A POINT IN TIME

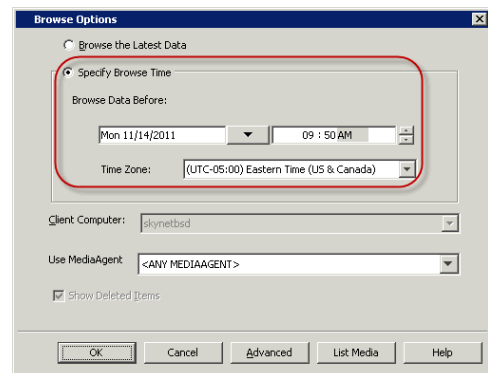
The point-in-time database restore is useful in the following scenarios:

- If any undesired transaction occurs in the database, you can revert the database to a state just before the transaction.
- If a database fails, you can restore to the state just before the point of failure.

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

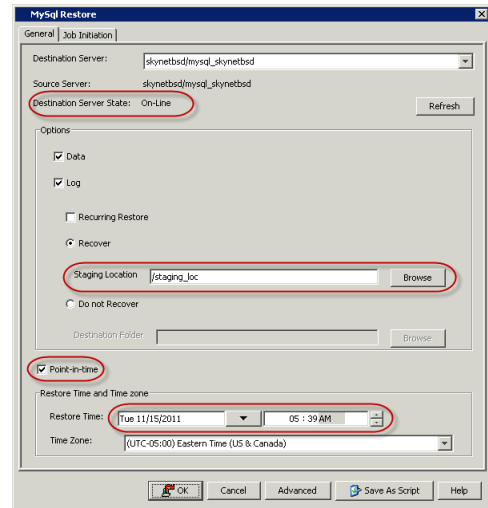
Use the following steps to restore a database to a previous point-in-time:

1. From the CommCell Browser, navigate to **Client Computers | <Client> | MySQL**.
2. Right-click the **<Instance>** and then click **All Tasks | Browse Backup Data**.
3. Click **Specify Browse Time** and select the date and time of the data backed up to be displayed in the Browse window.
4. Click **OK**.
5. In the right pane of the Browse window, select one or more databases and click **Recover All Selected**.
6. In the **MySQL Restore** dialog box, select the name of the client computer with instance to which you want to restore the database to in the **Destination Server**



from the list.

7. Verify that the status of the database is displayed as **On-Line**; if necessary click the **Refresh** button to refresh the status.
8. Type or click **Browse** to select the staging path for data in the **Staging Location** box.
9. Click **Point-In-Time** and select the date and time later than the date and time selected in the Browse Options dialog box.
10. Click **OK** to start the restore.



[Back to Top](#)

Restore Data - MySQL - How To

[Topics](#) | [How To](#) | [Full System Restore](#) | [Related Topics](#)

[Restore Data to a Point in Time](#)

[Recover Data with the Database Online](#)

[Recover Logs with the Database Online](#)

[Recover Logs with the Database Offline](#)

[Recurring Restores of Logs](#)

[Cross-Platform Restores from CommCell Console](#)

RESTORE DATA TO A POINT IN TIME

Before You Begin:

- Review the general and agent-specific restore requirements accessed from Restore Backup Data prior to performing any restore operation.
- When you browse and restore your data to a point-in-time, make sure that the point-in-time is later than the browse time selected in the browse options dialog box.

Required Capability: See Capabilities and Permitted Actions

▶ To restore to a point in time:

1. From the CommCell Browser, right-click the instance, click **All Tasks**, and click **Restore Server**.
2. From the Restore Options (General) dialog box, select **Point in Time** and select a **Restore Time and Time Zone**.
3. Continue your restore.

RECOVER DATA AND BRING THE DATABASE ONLINE

Before You Begin:

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

Required Capability: See Capabilities and Permitted Actions

▶ To recover data and bring the database online:

1. From the CommCell Browser, right-click the instance, click **All Tasks**, and click **Restore Server**.
2. From the Restore (Options) dialog box, click **Recover**.
3. Continue with your restore.

RECOVER LOGS WITH THE DATABASE ONLINE

Before You Begin:

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

Required Capability: See Capabilities and Permitted Actions

▶ To recover data and logs with the database online:

1. From the CommCell Browser, right-click the instance, click **All Tasks**, and click **Restore Server**.
2. From the MySQL Restore (General) dialog box, select **Recover**. The database will be online for recovery. When you select **Recover** option, the logs will be applied to the server.
3. Enter the path or click **Browse** to specify the destination folder path.
4. Continue with your restore.

RECOVER LOGS WITH THE DATABASE OFFLINE

Before You Begin:

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

Required Capability: See Capabilities and Permitted Actions

▶ To recover data and logs with the database offline:

1. From the CommCell Browser, right-click the instance, click **All Tasks**, and click **Restore Server**.
2. From the MySQL Restore (General) dialog box, select **Do Not Recover**. The database will be offline for recovery. When you select **Do Not Recover** option, you need to specify the path of the Destination Folder to store the logs without applying to the server.
3. Enter the path or click **Browse** to specify the destination folder path.
4. Continue with your restore.

RECURRING RESTORES OF LOGS

Before You Begin:

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

Required Capability: See Capabilities and Permitted Actions

▶ To recover data and bring the database online:

1. From the CommCell Browser, right-click the instance, click **All Tasks**, and click **Restore Server**.
2. From the MySQL Restore (General) dialog box, select **Recurring Restore** and then select **Recover** option for replicating online data from one server to another server.
3. From the MySQL Restore (General) dialog box, select **Recurring Restore** and then select **Do Not Recover** option for replicating offline data from one server to another server. Enter the path or click **Browse** to specify the destination folder path.
4. Continue with your restore.

The Recurring Restore option is applicable only to Log Restore. Make sure to execute a data restore operation prior to performing a recurring log restore.

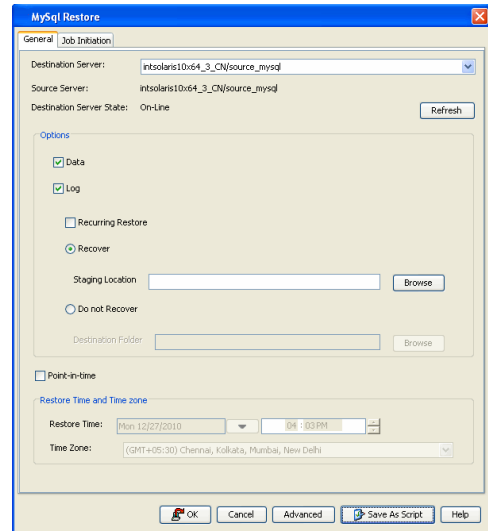
CROSS-PLATFORM RESTORES FROM COMMCELL CONSOLE

Perform the following step-by-step procedure to restore the data across same/different versions or platforms from CommCell Console:

1. From the CommCell Browser, right-click the MySQL server instance on the source client.
2. Select **All Tasks** and click **Restore MySQL Server**.
3. Select the name of the client computer with the database you want to restore in the **Destination Server** from the list.
4. Verify that the Status of the database is displayed as **On-Line**; if necessary click on

the **Refresh** button to refresh the status.

5. Select the following options to restore and recover the database. Click **Browse** and select the staging location.
 - Data**
 - Logs**
 - Recover**
6. Click **OK**.



[Back to Top](#)

Basic Restore

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

[Overview](#)

[Time Range Options](#)

OVERVIEW

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

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

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

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

TIME RANGE OPTIONS

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

Basic Restore - How To

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

[Basic Restore](#)

[Time Range Options for Basic Restore](#)

BASIC RESTORE

Before You Begin:

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

Required Capability: See Capabilities and Permitted Actions

▶ To restore data without using browse:

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

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

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

TIME RANGE OPTIONS FOR BASIC RESTORE

Before You Begin:

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

Required Capability: See Capabilities and Permitted Actions

▶ To provide time range options for basic restore:

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

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

[Back To Top](#)

Browse and Restore

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

OVERVIEW

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

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

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

Browse and Restore

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

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

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

Browse and Restore - How To


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

BROWSE AND RESTORE

Before You Begin:

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

Required Capability: See Capabilities and Permitted Actions

 To browse and restore data:

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

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

[Back To Top](#)

Instances - MySQL

Topics | How To | Related Topics

Overview

Configurable Properties

Configuring User Accounts

OVERVIEW

As MySQL instances reference a MySQL server, you must create a MySQL instance for each of the MySQL servers that you want to back up. In order to create a MySQL instance, the MySQL User should have DBA and SYSDBA privileges.

See Instances for an overview of Instances. Given below is a list of what you need to provide before creating a MySQL instance.

- A name for the MySQL server instance
 - The MySQL User account which will be used to perform all system operations
 - The path to the MySQL server Binary directory. For example, `/usr/bin`.
 - The path to the Log Data directory. For example, `/var/lib/mysql/mysqllogs`.
 - The path of the configuration file for the instance. For example, `/etc/my.cnf`.
 - The path to the socket file For example, `/var/lib/mysql/mysql.sock`.
-

CONFIGURABLE PROPERTIES

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

You can view or change the instance configuration from the **Instance Properties** dialog box. The following information can be configured.

ENABLE AUTOMATIC DISCOVERY OF DATABASES

For the MySQL iDataAgent, you can assign the system to search the instance for unconfigured databases and add them to the system. See Modify an Instance for step-by-step instructions.

CHANGE THE DESIGNATED HOME DIRECTORY

For the Oracle, Oracle RAC, and SAP for Oracle iDataAgents, you can change the Oracle Home directory. Similarly, you can change the Informix (Home) directory for the Informix iDataAgent, the Sybase Home directory for the Sybase iDataAgent, the MySQL Home directory for the MySQL iDataAgent, the PostgreSQL home directory for the PostgreSQL iDataAgent, and the Instance Home directory for the DB2 DPF pseudo-client.

See Change the Home Directory for step-by-step instructions.

CHANGE THE LOG DATA DIRECTORY

When you enable binary log, an index file (`.index` extension) is generated and reside in data directory. Archive log Index file is a text file which comprises of all the running list of transaction logs for MySQL database. Typically, this file is used to backup the logs. For the MySQL iDataAgent, you can change the path of the archive log index file for the instance. See Change the Archive Log Index File for step-by-step instructions. If binary log is disabled on the MySQL server, then the path of the data directory is displayed.

CHANGE THE CONFIGURATION FILE

You can change the path of the configuration file for the instance. See Change the Configuration File for step-by-step instructions.

If you encounter any compatibility mismatch between Calypso and Linux/Linux flavored operating system, you should rename and map it to common library that you are using for both OS and Calypso.

For Example: If a log reports an error: MySQL iDataAgent unable to configure instance - `/opt/<instance_name>/Base/libgcc_s.so.1: version `GCC_4.2.0' not found (required by /usr/lib64/libstdc++.so.6)`

Do not delete `"/opt/<instance_name>/Base/libgcc_s.so.1"`. Instead, rename it as `"/opt/<instance_name>/Base/libgcc_s.so.1.bkp"`.

CHANGE THE SOCKET FILE

For the MySQL iDataAgent, you can change the path of the socket file for the instance. See [Change the Socket File](#) for step-by-step instructions.

USER ACCOUNTS

Use the Accounts tab to identify the account that the system uses to access the associated instance.

See MySQL iDataAgent section in **User Accounts and Passwords** for more information.

For step-by-step instructions on changing the user accounts for the instance, see [Change Account for Accessing Instances](#).

VIEW AND CHANGE INSTANCE-LEVEL STORAGE POLICIES

From the Log Backups tab, you can change the storage policy for backing up the logs on the selected instance. (See [Storage Policies](#) for an overview.)

For step-by-step instructions, see [Select a Storage Policy for Log Backups](#).

USER SECURITY

You can perform the following functions:

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

For more information, see [User Administration and Security](#).

CONFIGURING USER ACCOUNTS

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

- Perform backups and restores
- Access the MySQL Server application
- Stop or start the MySQL Server services


The following table illustrates the necessary privileges the user account should have to perform backup and restore operations:

OPERATIONS	PRIVILEGES MYSQL USER SHOULD HAVE:	EXAMPLE QUERY TO GRANT THE PERMISSION
Backup	<ul style="list-style-type: none"> • SHOW DATABASES • SUPER • RELOAD • SELECT • LOCK 	<ol style="list-style-type: none"> 1. <code>mysql> GRANT SHOW DATABASES, SELECT, LOCK TABLES, RELOAD,SUPER ON *.* to <backup_agent username>@localhost IDENTIFIED BY '<backup_agent password>';</code> 2. However, it is recommended to grant all Database Administrator privileges to perform backup operations too. <code>mysql> GRANT ALL PRIVILEGES ON *.* TO <backup_agent username>@localhost IDENTIFIED BY '<backup_agent password>';</code>
Restore	Full Database Administrator privileges	<code>mysql> GRANT ALL PRIVILEGES ON *.* TO <restore_agent username>@localhost IDENTIFIED BY '<restore_agent password>' WITH GRANT OPTION;</code>

AT THE INSTANCE LEVEL

Use the following steps to change the MySQL Server user account for an Instance:

1. From the CommCell Browser, navigate to **Client Computers | <Client> | MySQL**.
2. Right-click the **<Instance>** and click **Properties**.
3. Click the **Accounts** tab.
4. Type the SA user name in the **SA User Name** box.
5. Type the password in the **SA Account Password** box and retype it in the **SA Confirm Password** box.
6. Click **OK**.



[Back to Top](#)

Instances - MySQL - How To

Topics | How To | Related Topics

Create an Instance

Modify an Instance

Delete an Instance

Change the Home Directory

Change the Archive Log Index File

Change the Configuration File

Change the Socket File

Change Account for Accessing Instances

Select a Storage Policy for Log Backups

CREATE AN INSTANCE

You must create a MySQL instance for each MySQL server that you want to back up. Ensure that you have a user account with full DBA and SYSDBA privileges to create a new instance.

Use the following steps to create a new instance on an Unix client:

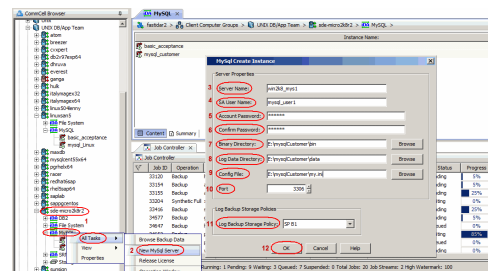

1. From the CommCell Browser, navigate to **Client Computers | <Client> | MySQL Server**.
2. Right-click **MySQL Server**, point to **All Tasks** and then click **New MySql Server**.
3. In the **Server Name** box, type the name of the MYSQL Server instance that you want to assign to this instance.
4. In the **SA User Name** box, type the user name that you used to create an instance.
5. In the **SA Account Password** box, type the password that you want to assign to this instance.
6. In the **SA Confirm Password** box, type the password again to confirm the assigned password to this instance.
7. In the **Binary Directory** box, enter the path or click **Browse** to select the directory where the mysql binaries are stored.
8. In the **Log Data Directory** box, enter the path or click **Browse** to select the directory where the log files are stored.
9. In the **Config File** box, enter the path or click **Browse** to select the directory where the configuration (.cnf) file exist.
10. In the **Socket File** box, enter the path or click **Browse** to select the directory where the socket file exist.

If you have a multiple instances installed on the client computer, make sure that the **Binary Directory**, **Log Data Directory**, **Config File** and **Socket File** locations are different for each instance that you want to backup.

11. Select a storage policy from the **Log Backup Storage Policy** drop down list.
12. Click **OK**.

Use the following steps to create a new instance on a Windows client:

1. From the CommCell Browser, navigate to **Client Computers | <Client> | MySQL Server**.
2. Right-click **MySQL Server**, point to **All Tasks** and then click **New MySql Server**.
3. In the **Server Name** box, type the name of the MySQL Server instance that you want to assign to this instance.
4. In the **SA User Name** box, type the user name that you used to create an instance.
5. In the **SA Account Password** box, type the password that you want to assign to this instance.




6. In the **SA Confirm Password** box, type the password again to confirm the assigned password to this instance.
7. In the **Binary Directory** box, enter the path or click **Browse** to select the directory where the mysql binaries are stored.
8. In the **Log Data Directory** box, enter the path or click **Browse** to select the directory where the log files are stored.
9. In the **Config File** box, enter the path or click **Browse** to select the directory where the configuration (.cnf) file exist.
10. In the **Port** box, type or select the TCP/IP port number to open the communication between MySQL server and clients.

If you have a multiple instances installed on the client computer, make sure that the **Binary Directory**, **Log Data Directory**, **Config File** locations and **Port** number are different for each instance that you want to backup.

11. Select a storage policy from the **Log Backup Storage Policy** drop down list.
12. Click **OK**.

MODIFY AN INSTANCE

Required Capability: See Capabilities and Required Actions

 To modify an Instance

1. From the CommCell Browser, right-click the instance you want to modify, and then click **Properties** from the shortcut menu.
2. Click the General tab of the Instance properties box to change the required access information for this instance.
3. For MySQL iDataAgent, use this Accounts tab to specify the MySQL Server user account.
4. If your agent uses the Restore View Names tab, use this tab to enter or modify a list of text patterns for the view through which the system displays a database's data documents when you browse.
5. If your agent has the Security tab, you can view CommCell resource associations to enable or disable user groups that can perform system operations on associated resources.
6. Click the Details tab, and type or select values for the fields at the bottom of the dialog box as appropriate.
7. From the Storage Device tab(s), select the storage policy for the log backups from the Log Backups tab.
8. If your agent uses the Software Compression tab, and you want to enable software compression for the selected backups, click the tab and select the appropriate option.
9. If you want to enable data encryption, from the Encryption tab select encryption settings for third-party or RMAN command line backups.
10. Click **OK** to save the configuration changes.
11. A message advises you to schedule backups for your new instance. It is recommended you elect to set a schedule now. If you click **No**, this procedure is now complete. If you click **Yes**, a dialog box prompts you to back up all subclients of this instance.
12. Click **Yes**. The **Backup Options** dialog box appears, which allows you to create the necessary schedules.
13. From the Backup Options dialog box, select the type of backup that you want to schedule.
14. If you want to access **Advanced Backup Options**, click **Advanced**.
15. After selecting the backup type and any advanced options, click **OK**. The **Schedule Details** dialog box appears.
16. From the Schedule Details tab of the **Schedule Details** dialog box, select the scheduling options that you want to apply, then click **OK**.

DELETE AN INSTANCE

Before You Begin

- For most agents except Oracle and SQL Server, you cannot delete an instance if there is only one instance present for an agent. To delete the final instance for most agents except the ones identified, you must first remove the agent software from the client computer. For Oracle RAC, which does not involve agent software, you can delete the instance under the RAC pseudo-client.
- When you delete a specific instance positioned in the level below the corresponding agent or pseudo-client in the CommCell Console, the following actions occur:
 - The corresponding backup data for that instance becomes immediately unavailable, thus making it inaccessible for restore purposes. In effect, the backup data is deleted.
 - All child levels contained within the level are deleted in the CommCell Browser.

- All job schedules and job histories that pertain to any of the levels within the deleted instance are deleted.
- If you delete an Oracle iDataAgent instance included within an Oracle RAC instance (node), the backup data is not deleted, and it is available for restore by another eligible client.
- 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.

Required Capability: See Capabilities and Required Actions

▶ To delete an instance positioned in the level below the corresponding agent or pseudo-client

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.

▶ To delete an Oracle iDataAgent instance included within an Oracle RAC instance (node)

1. Right-click the Oracle RAC instance, click **Properties**, and click **Details**.
2. Click the instance that you want to delete within **Client/Instance/Status** and click **Delete**.
3. Repeat the previous step for each Oracle iDataAgent instance (except the last one remaining) within the Oracle RAC instance that you want to delete.
4. Click **OK**. This should delete the Oracle iDataAgent instance(s).

With the exceptions noted above, you can no longer access the corresponding backup data for restore purposes.

CHANGE THE HOME DIRECTORY

Required Capability: See Capabilities and Required Actions

▶ To change the Home Directory:

1. From the CommCell Browser, right-click the icon for any supported instance and then click **Properties**.
2. For agents other than Oracle RAC, from the General tab of the Properties dialog box, type the path in the **Home** or **Dir** (Directory) box or use **Browse** to locate the path. Then go to Step 6.
3. For the MySQL iDataAgent, from the General tab of the Properties dialog box, type the path in the **Binary Directory** box or use **Browse** to locate the path. Also, type the path to the MySQL data directory in the **Data Directory** box or use **Browse** to locate the path. Then go to Step 6.
4. For Oracle RAC, from the Details tab of the Properties dialog box, click the appropriate client/instance/status and then click **Modify**.
5. In the Update Instance dialog box, enter the path in the **Home** box or use **Browse** to locate the path.
6. Click **OK** to save your settings.
7. Continue to make additional settings as desired.

CHANGE THE ARCHIVE LOG INDEX FILE

Required Capability: See Capabilities and Required Actions

▶ To change the Archive Log Index file:

1. From the CommCell Browser, right-click the icon for any supported instance and then click **Properties**.
2. From the General tab of the Properties dialog box, enter the path for the archive index file in the **Archive Log Index File** field or use **Browse** to locate the path.
3. Click **OK** to save your settings.

CHANGE THE CONFIGURATION FILE

Required Capability: See Capabilities and Required Actions

▶ To change the configuration file:

1. From the CommCell Browser, right-click the icon for any supported instance and then click **Properties**.

2. From the General tab of the Properties dialog box, enter the path for the configuration file in the **Config File** field or use **Browse** to locate the path.
 3. Click **OK** to save your settings.
-

CHANGE THE SOCKET FILE

Required Capability: See Capabilities and Required Actions

▶ To change the Archive Log Index file:

1. From the CommCell Browser, right-click the icon for any supported instance and then click **Properties**.
 2. From the General tab of the Properties dialog box, enter the path for the archive index file in the **Socket File** field or use **Browse** to locate the path.
 3. Click **OK** to save your settings.
-

CHANGE ACCOUNT FOR ACCESSING INSTANCES

Required Capability: See Capabilities and Permitted Actions

▶ To change a user account for accessing an instance:

1. From the CommCell Browser, expand the tree to view the instance for the affected agent. Then right-click the instance, click **Properties**, and click **Accounts**.
 2. In the Instance Properties (Accounts) tab, enter the MySQL Server user account information.
 3. Click **OK** to save the settings.
-

SELECT A STORAGE POLICY FOR LOG BACKUPS

Required Capability: See Capabilities and Required Actions

▶ To select a Storage Policy for Log Backups:

1. From the CommCell Browser, right-click the icon for any supported instance and then click **Properties**.
 2. From the Log Backups tab of the Properties dialog box, select from the list of available storage policies.
 3. Click **OK** to save your settings.
-

[Back to Top](#)

Subclients - MySQL

Topics | How To | Related Topics

Overview

Subclient Content

Configurable Properties

Things to Consider when Creating and Configuring MySQL Subclients

OVERVIEW

The following table shows subclient creation and configuration details specific to the MySQL iDataAgent.

AGENT	Type of Data	Default Subclient created during install of the Agent	Supports Default Subclient	Supports User Defined Subclient	Contents of the default subclient when user-defined subclient is present	Other Types of subclients supported by the Agent	Notes
MySQL iDataAgent	MySQL Server databases, transaction logs	No	Yes	Yes	database objects not assigned to other subclients, unless otherwise configured*	None	Default subclient is created when the instance is configured *See Caution Against Re-configuring Default Subclient Content.

SUBCLIENT CONTENT

The MySQL iDataAgent is designed to back up and restore MySQL Server databases along with the associated administration files and transaction logs. After you install a MySQL iDataAgent, a default subclient is created by the system only after you first configure a MySQL instance from the CommCell Console. Since a MySQL instance references a MySQL Server, the default subclient includes all the databases for the MySQL Server associated with that instance. However, you can use the CommCell Console to create user-defined subclients to distribute some of the database content by configuring the **Subclient Properties (Content)** tab. Any content that you do not distribute to the user-defined subclients will remain in the default subclient.

When you configure subclient content, you click the **Subclient Properties (Content)** tab for the subclient and then **Configure** to view the eligible databases for inclusion into the subclient content. You then include the desired databases into your subclient content by associating the databases with the user-defined subclient. Once you do this, the subclient contents may be displayed as follows:

Contents of subclient:

```
/db2
/db3
/db4
/db5
/db6
/db7
```

In this example, db2, db3, db4, db5, db6, and db7 are the databases that are included in the subclient content.

A useful way to distribute subclient content for this iDataAgent is to create separate user-defined subclients for frequently used databases. For example, if the db2 and db3 databases shown above undergo frequent changes, you can create a separate subclient for each database. For example:

- User-defined subclient: Temp1

Content:

```
db2
```

- User-defined subclient: Temp2

Content:

```
db3
```

Distributing the client data using subclients in this way can help improve backup performance by organizing the workload on the client into logical groupings.

FILTERING

Data protection filtering is not applicable for this agent.

WILDCARDS AND REGULAR EXPRESSIONS

Wildcards and regular expressions are not supported for this agent when defining subclient content.

CONSIDERATIONS

For more information, see Things to Consider when Creating and Configuring MySQL Subclients.

CONFIGURABLE PROPERTIES

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

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

ACTIVITY CONTROL

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

CONTENT/DATABASES

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

DATA TRANSFER OPTIONS

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

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

DATA ENCRYPTION

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

DATA PATHS

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

PRE/POST PROCESSES

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

STORAGE POLICIES

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

SUBCLIENT NAME

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

THINGS TO CONSIDER WHEN CREATING AND CONFIGURING MYSQL SUBCLIENTS

When creating and configuring subclients for the MySQL iDataAgent, keep in mind the following considerations:

- **Caution Against Re-configuring Default Subclient Content**

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

[Back to Top](#)

Subclients - MySQL - How To

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

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

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

[Configure Subclient Content](#)

[Configure MySQL Databases](#)

[Configure the Subclient for Data Encryption](#)

[Create a New Subclient](#)

[Delete a User-Defined Subclient](#)

[Enable or Disable Operations](#)

[Enable Software Compression for a Subclient](#)

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

[Rename a Subclient](#)

[View Data Paths Associated with a Subclient](#)

ASSOCIATE A SUBCLIENT TO A STORAGE POLICY

Required Capability: See Capabilities and Permitted Actions

▶ To associate a subclient to a storage policy:

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

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

Before You Begin

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

Required Capability: Capabilities and Permitted Actions

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

1. From the CommCell Browser, right-click the subclient for which you want to configure a pre/post process, and then click **Properties** from the shortcut menu.
2. Click the Pre/Post Process tab of the Properties dialog box.
3. For an agent other than the Oracle RAC iDataAgent, click inside the space that corresponds to one of the following phases and type the full path of the process that you want executed during that phase. Alternatively, click **Browse** to locate the process (applicable only for paths that do not contain any

spaces). For the Oracle RAC iDataAgent, click **Browse** for the corresponding process, click the name of the control node client in the Select Client for Browse dialog box, and click **OK**. Then browse for and click the process.

- PreBackup
- PreScan
- PreArchive
- PreCopy
- PreSnap
- PostBackup
- PostScan
- PostArchive
- PostCopy
- PostSnap

Click **OK**.

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

CONFIGURE SUBCLIENT CONTENT

Before You Begin

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

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

Required Capability: See Capabilities and Permitted Actions

▶ To configure subclient content:

1. From the CommCell Browser, right-click the subclient for which you want to configure content, click **All Tasks** (if applicable) and then click **Properties**.
2. Follow the procedure below that is applicable for your agent:
 - For File System, Active Directory, File Archiver, Exchange Public Folder iDataAgents, NDS, and SharePoint Server iDataAgents click the Subclient Properties (Content) tab and configure content for the subclient as described below for your agent:
 - For File System, Active Directory, File Archiver, NDS, and SharePoint Server iDataAgents: Type the full path of the data that you want to include as subclient content in the **Enter New Content** field, then click **Add**. Optionally, click **Browse** to enter the content. When browsing content while configuring SharePoint subclients, you can add content via multiple selections with the CTRL or SHIFT keys. For Windows, when specifying a UNC Path, click **As User**, and enter the user account information for the domain user with permissions for that path. For NetWare/DNS, see the Notes

section below for content path examples. For Unix File Systems, you can enter the mount point of an NFS-mounted file system, see the Notes section below for examples.

- For Exchange Public Folder iDataAgents: Click **Browse**, select folders to include as content, then click **Add**.
- For the Unix File System iDataAgents, to facilitate the management of resource fork data in Apple double-encoded Macintosh files, click **Enable Apple Double Support**.
- For the Unix File System iDataAgents, to view the actual data path for any symbolic link in the subclient content, click **Expand symbolic links of subclient content** and then click **Discover**.
- For NAS NDMP iDataAgents, configure the **Backup Content Path** field(s) as described below, then click **Add**:
 - Click the drop-down list arrow to display the root volumes on the file server. To change the root volume, click one in the list. If you want to refine the content path further, use the space to the right of (or below) the root volume list to enter additional path information. Note the following:
 - For NetApp, the root volume is the mount path of each volume.
Example: for volume FS1 the root volume will be /vol/FS1.
 - For EMC Celerra, the root volume is the mount point created for a volume.
Example: for volume FS1 with mount point /FS1 the root volume will be /FS1.
 - For Hitachi, no root volumes are shown in the drop down list. Type the full path of the root volume.
Example: for volume FS1 with mount point /mnt/FS1 the root volume will be /mnt/FS1.
 - For BlueArc, the root volume is a combination of a descriptor of the path and the volume name.
Example: for volume FS1 with a mount point of / the root volume will be /__VOLUME__/_/FS1.
 - Optionally, for NetApp NAS NDMP, click **Browse** to enter the content.
- For Exchange Mailbox and Exchange Mailbox/Public Folder Archiver Agents follow the procedure to Discover and Assign New Mailboxes or Assign Mailboxes to Another Subclient.
- For Lotus Notes Database and Document iDataAgents follow the procedure to Discover and Assign New Databases or Assign Databases to a Subclient.
- For DB2, DB2 DPF, Exchange Database, Novell GroupWise, SharePoint Server, SQL Server Database, Sybase, and MySQL iDataAgents, click the Subclient Properties (Content) tab and configure content for the subclient as described below for your agent:
 - For the DB2 iDataAgent, specify whether you want to include the entire database or a subset of this data as content for the subclient. For the DB2 DPF iDataAgent, specify whether you want to include all the affected database partitions or a subset of this data as content for the subclient.
 - For Exchange and GroupWise iDataAgents: Click **Configure**. From the Add/Modify Subclients dialog box click the subclient entry for the database element/Storage Group that you want to add to the new subclient and select the name of the destination subclient from the list that appears. Alternatively, you can select and assign a range of databases/storage groups using the **Change all selected databases/storage groups to list**. Note that you must have at least one database element/Storage Group assigned to this subclient in order to save the configuration.

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

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

For more information, see When to Use Volume Explorer. For a step-by-step procedure, see Configure a CXBF Device in Volume Explorer.
 - Image Level and Image Level ProxyHost on Windows iDataAgents: Click **Add**. Then in the **Add Content** dialog box, type the full path of the volume or mount point that you want to include as subclient content, then click **Add**. Optionally, click **Browse** to select the content. Click **OK**. The volume or mount point is added to the **Contents of subclient** list. Add additional content by repeating this step.
 - ProxyHost iDataAgent: Select a backup host from the **Backup Host** list. This is the computer to which the BCV is connected. Click **Add**. In the **Content** field of the Add/Edit Content for Subclient dialog box, type the primary host path of the content that you want to back up, or click **Browse**

to find and select this data. In the **Backup Host BCV Path** field of the Add/Edit Content for Subclient dialog box, type the path through which the backup host accesses this data on the BCV, or click **Browse** to find and select this path. Click **OK**. The primary host data path and corresponding backup host BCV path are added as a single entry in the **Contents of subclient** list. To add additional entries, repeat these steps. Refer to Notes below for more information.

- o For Quick Recovery Agents, click the Subclient Properties (Content) tab and configure the following options:
 - Click **Add Volume**. From the Adding Volume dialog box, select volume(s) that you want to add to the subclient content (use CTRL + click to select multiple volumes). You can add/edit additional advanced options for the selected volume by select **Advanced** on the Adding Volume dialog box. Click **OK**.
 - Click Add App to select an application and associated volumes. Click **OK**.

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

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

See also Configure Subclients for Overlapping Content.

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

NOTES

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

<primary_host_path> --> <backup_host_path>

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

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

CONFIGURE MYSQL DATABASES

Related Topics:

- See Subclients

Required Capability: See Capabilities and Permitted Actions

▶ To configure databases:

1. From the CommCell Browser, if you are creating a new subclient, start to create the subclient and then click the Content tab from the Subclient dialog box. Otherwise, right-click the subclient whose content you want to configure, click **Properties**, and then click the Content tab from the Subclient dialog box.
2. For an existing subclient, check the current subclient content in the **Content** tab. If you want to delete any database(s) from the subclient content, select the database(s) and then click **Delete**.
3. Click **Configure**.
4. If the Database Configuration dialog box is blank, or to discover new database elements, click **Discover**.
5. In the Database configuration dialog box, check whether each database is associated to the desired subclient. If not, to change the subclient association

for the database, click the current associated subclient and then select the appropriate subclient from the list. Optionally, to associate more than one database to the same subclient, select the appropriate databases and then use the **Change all selected storage groups to** list to select the common subclient.

6. Click **OK**. The Subclient dialog box should display the desired subclient contents.
7. Make any other desired or required changes and then click **OK**.

CONFIGURE THE SUBCLIENT FOR DATA ENCRYPTION

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

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

Before You Begin

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

Required Capability: Capabilities and Permitted Actions

▶ To configure the subclient for data encryption:

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

CREATE A NEW SUBCLIENT

Before You Begin

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

Required Capability: See Capabilities and Permitted Actions

▶ To create a new subclient:

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

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

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

This task is now complete.

DELETE A USER-DEFINED SUBCLIENT

Related Topics:

- Command Line Interface - qdelete subclient
- Subclients

Required Capability: See Capabilities and Permitted Actions

▶ To delete a user-defined subclient:

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

ENABLE OR DISABLE OPERATIONS

Required Capability: See Capabilities and Permitted Actions

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

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

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



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

ENABLE OR DISABLE SOFTWARE COMPRESSION FOR A SUBCLIENT

Before you Begin

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

Required Capability: Capabilities and Permitted Actions

▶ To enable software compression for a subclient:

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

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

This task is now complete.

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

Before You Begin

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

Required Capability: Capabilities and Permitted Actions

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

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

RENAME A SUBCLIENT

Before You Begin

- You can rename a subclient at any time. However, we recommend that you don't rename a subclient while a data protection or archive operation is running on that subclient.
- In cases where a subclient is renamed using the same name as a deleted subclient, the system will append a Unix time stamp to the deleted subclient's name in data protection job history reports and views to distinguish the two subclients. For example, `subclientname_1104257351`.

Required Capability: See Capabilities and Permitted Actions

▶ To rename a subclient:

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

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

VIEW DATA PATHS ASSOCIATED WITH A SUBCLIENT

Required Capability: See Capabilities and Permitted Actions

▶ To view data paths:

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

Backup Job History

Topics | How To | Related Topics

Overview

Items That Were Backed Up

Items That Failed

Pruning Backup History Information

Supported Features

Content Indexing History Information

OVERVIEW

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

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

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

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

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

ITEMS THAT WERE BACKED UP

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

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



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

See [View the Items That Were Protected During a Data Protection Operation](#) for step-by-step instructions.

ITEMS THAT FAILED

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

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

application files.

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




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

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

See [View the Items That Failed For a Data Protection Operation](#) for step-by-step instructions.

PRUNING BACKUP HISTORY INFORMATION

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



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

SUPPORTED FEATURES

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

CONTENT INDEXING HISTORY INFORMATION

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

ITEMS THAT WERE SUCCESSFULLY CONTENT INDEXED

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

CONTENT INDEXING FAILURES

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

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

[Back to Top](#)

Backup Job History - How To

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

[View Backup Job History](#)

[View the Items That Were Protected During a Data Protection Operation](#)

[View the Items That Failed For a Data Protection Operation](#)

[View Job History Details](#)

[View the Media or Mount Paths of a Job History](#)

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

[View the Items that were Moved to Media during SnapProtect Backup](#)

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

View the Items that Were Not Indexed During Content Indexing

View the Items that Were Successfully Content Indexed

Resubmit a Backup Job

VIEW BACKUP JOB HISTORY

▶ To view backup history:

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. The system displays the Backup Job History window.
3. Click **OK**.

VIEW THE ITEMS THAT WERE PROTECTED DURING A DATA PROTECTION OPERATION



This option is available for File System-like agents.

Required Capability: none required

▶ To view the list of items that were protected during a data protection operation.

1. From the CommCell Browser, right-click the entity whose history of data protection operations you want to view, click **View**, and then click the necessary options to view a job history.
2. From the Job History Filter dialog box, select the filter options, if any, that you want to apply, and then click **OK**.
3. From the Job History window, right-click the operation whose list of protected items you want to view, and then select **View backup file list/View Backed Up Messages**. The **Backup file List** window displays a list of the backed up files/messages that were included in the backup job. You can use the **Search** option to find items in the window.
4. Click **File -> Exit**.
5. Click **Close** from the **Job History** window.

VIEW THE ITEMS THAT FAILED FOR A DATA PROTECTION OPERATION



A listing of files and folders that failed is not available for the Quick Recovery Agent, nor the Image Level and Image Level ProxyHost iDataAgents. These agents do not perform a file level backup/copy.

▶ To view the list of items that failed for a data protection operation:

1. From the CommCell Browser, right-click the entity whose history of data protection operations you want to view, click **View**, and then click to view a job history.
2. From the Job History Filter dialog box, select the filter options, if any, that you want to apply, and then click **OK**.
3. From the Job History window, right-click the operation whose list of failed items you want to view, and then select **View Failed Items**. The **Unsuccessful Backup Files** window (for DataArchiver Agents, **Items On Which Archive Failed**) displays those items that failed. If no items failed, a message to that effect is displayed.
4. Click **Close**.

VIEW JOB HISTORY DETAILS

Required Capability: See Capabilities and Permitted Actions

▶ To view the details of a job history:

1. From the CommCell Browser, right-click the entity (e.g., subclient, client computer, etc.) whose job history you want to view, click **View**, and then click job history.
2. From the Job History Filter dialog box, select the filter options that you want to apply and click **OK**.

3. From the Data Management Job History window, right-click the job whose job details you want to view, and then click **View Job Details**.
4. The Job Details dialog box appears, displaying detailed job history in General, Details, Phase Details and Attempts tabs for the selected job.
5. Click **OK**.



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

VIEW THE MEDIA OR MOUNT PATHS OF A JOB HISTORY

▶ To view media or mount paths associated with a job history:

1. From the CommCell Browser, right-click the entity (e.g., subclient, client computer, etc.) whose job history you want to view, click **View**, and then select the appropriate history.
2. From the Job History window select the filter options, if any, that you want to apply, and then click **OK**.
3. From the job history window, right-click the backup whose media or mount paths you want to view, and then click **View Media**.
4. The Media Used By Job ID window displays a list of media or mount paths used by the operation.
5. Click **OK**.

VIEW THE EVENTS OF A JOB HISTORY

Required Capability: See Capabilities and Permitted Actions

▶ To view the events associated with a job:

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

VIEW THE ITEMS THAT WERE MOVED TO MEDIA DURING SNAPPROTECT BACKUP



This option is available for the SnapProtect Backup.

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

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



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

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

VIEW THE LOG FILES OF A JOB HISTORY

Required Capability: See Capabilities and Permitted Actions

► To view the log files of a Job History:

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

VIEW THE ITEMS THAT WERE SUCCESSFULLY CONTENT INDEXED



This option is available for operations that performed content indexing.

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

1. From the CommCell Browser, right-click the entity whose operations you want to view, click **View**, and then click the necessary options to view a job history.
2. From the Job History Filter dialog box, select the filter options, if any, that you want to apply, and then click **OK**.
3. From the Job History window, right-click the job for which you want to view the successfully content indexed items, select **View Content Index**, and click **Successful Items**.
4. Click **Close**.
5. Click **Close** from the **Job History** window.

VIEW THE ITEMS THAT FAILED TO CONTENT INDEX



This option is available for operations that performed content indexing.

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

1. From the CommCell Browser, right-click the entity whose operations you want to view, click **View**, and then click the necessary options to view a job history.
2. From the Job History Filter dialog box, select the filter options, if any, that you want to apply, and then click **OK**.
3. From the Job History window, right-click the job for which you want to view the list of items failed to content index, select **View Content Index**, and click **Failed Items**.
4. Click **Close**.
5. Click **Close** from the **Job History** window.


RESUBMIT A BACKUP JOB

► To resubmit a backup job:

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

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


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




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

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


3. The Backup Job History window displays with the specified filter options.
4. Right-click on any job, and select **Resubmit Job**.



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




6. If you need to run the backup operation immediately, select **Immediate** from the **Job Initiation** tab. Go to step 11.
7. If you need to schedule the restore operation, select **Schedule** from the Job Initiation tab and click **Configure**.
8. From the **Schedule Details** dialog box that appears, select the appropriate scheduling options.




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

9. From the **Advanced Schedule Options** dialog box:
 - **Range of recurrence:** Specify the date on which you want this schedule to take effect.



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

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



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

[Back to Top](#)

Restore Job History

Topics | How To | Related Topics

Overview

Items That Restored

Supported Features

OVERVIEW

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

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

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

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

ITEMS THAT ARE RESTORED

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

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

SUPPORTED FEATURES

Consider the following.

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

[Back to Top](#)

Restore History - How To

Topics | How To | Related Topics

[View Restore Job History](#)

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

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

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

VIEW RESTORE JOB HISTORY

▶ To view the restored items associated with a job:

1. From the CommCell Browser, right-click the entity (e.g., subclient, client computer, etc.) whose job restore history you want to view, click **View**, and then click **Restore History**.

2. From the Job History filter window, select the filter options, if any, that you want to apply, and then click **OK**.
 3. From the Job History window, right-click the job whose restored items you want to view; click **View Restore Items**, and select from the type of items to view: **Successful**, **Failed**, **Skipped** or **All**.
 4. The **Restored Files** window will display the selected type of restored items for the job.
 5. Click **OK**.
-

VIEW THE EVENTS OF A JOB HISTORY

Required Capability: See Capabilities and Permitted Actions

▶ To view the events associated with a job:

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

VIEW THE MEDIA OR MOUNT PATHS OF A JOB HISTORY

▶ To view media or mount paths associated with a job history:

1. From the CommCell Browser, right-click the entity (e.g., subclient, client computer, etc.) whose job history you want to view, click **View**, and then select the appropriate history.
 2. From the Job History window select the filter options, if any, that you want to apply, and then click **OK**.
 3. From the job history window, right-click the backup whose media or mount paths you want to view, and then click **View Media**.
 4. The Media Used By Job ID window displays a list of media or mount paths used by the operation.
 5. Click **OK**.
-

VIEW THE LOG FILES OF A JOB HISTORY

Required Capability: See Capabilities and Permitted Actions

▶ To view the log files of a Job History:

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

[Back to Top](#)

Frequently Asked Questions - MySQL

CAN WE PROTECT DATA ON NON-GLOBAL ZONES IF THE FILE SYSTEM iDATAAGENT 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 MYSQL DATA ON A NON-GLOBAL ZONE?

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

HOW DO WE BACKUP MYSQL DATABASES?

We use mysqldump utility to perform backups on MySQL databases. Various storage engines such as InnoDB, MyISAM etc in MySQL server store the tables in different ways.

By default, mysqldump utility will use the `-opt` flag for databases containing **MyISAM** tables. The `-opt` flag uses the `-lock tables` command that will lock all the tables to be dumped before dumping them for backup. Also, `-lock-tables` uses READ LOCAL lock which permits concurrent inserts (inserts at the end of the table) by other sessions while dumping tables.

In the case of transactional tables such as InnoDB, the mysqldump utility uses the `--single-transaction` flag, which is a much better option than `--lock-tables` because it does not need to lock the tables at all.

HOW TO ATTAIN CONSISTENT BACKUP OF DATABASES HAVING MYISAM AND MYISAM + INNODB MIX TABLES?

We have a way to accept user options to get consistent backup of databases having MYISAM or MYISAM/INNODB mix tables. You can use `--skip-single-transaction` `--lock-all-tables` options in the `sAdvancedBackupOptions` registry key under Options can be supplied to mysqldump utility to get consistent backup.

DO WE SUPPORT NDB CLUSTER CONFIGURATION FOR MYSQL iDATAAGENT?

No, we do not support NDB cluster configuration for MySQL iDataAgent.