

Administration Guide - Virtual Server iDataAgent (VMware)

TABLE OF CONTENTS

OVERVIEW

Introduction

Key Features

- Complete Virtual Machine Protection
- Granular Recovery of Virtual Machine Data
- Minimal Load on Host Resources
- Automatic Protection of Virtual Machines
- SnapProtect Backup
- Block Level Deduplication

Terminology

SYSTEM REQUIREMENTS - VIRTUAL SERVER /DATAAGENT (VMWARE)

GETTING STARTED - VMWARE DEPLOYMENT

Installation

Method 1: Install Software from the CommCell Console

Method 2: Interactive Install

GETTING STARTED - VMWARE CONFIGURATION

GETTING STARTED - VMWARE BACKUP

Perform a Backup

- What Gets Backed Up
- What Does Not Get Backed Up

GETTING STARTED - VMWARE RESTORE

Perform a Restore

ADVANCED - VMWARE CONFIGURATION

Understanding the CommCell Console

Optimizing Performance

Configuring Backups for vSphere VADP Environments

Configuring Backups for Combined vSphere VADP and VCB Environments

Configuring Backups for VCB Environments

Configuring Backups of Virtual Centers

Configuring Backups of ESX Servers

Configuring Backups of Volumes

Configuring Backups of Files and Folders

Configuring Automatic Discovery of Virtual Machines

- Discover Non-Associated Virtual Machines to Default Subclient
- Discover By Regular Expression
- Discover By Single-ESX Server Affinity
- Discover By Multiple-ESX Server Affinity
- Discover By DataStore Affinity
- Discover By Resource Pool Affinity
- Discover By OS Group

Configuring Transport Modes

Accessing Virtual Clients in the CommCell Browser

- Show or Hide Discovered Virtual Clients
- Locate Virtual Clients

Modifying an Agent, Instance, or Subclient

Deleting an Instance, Backup Set, or Subclient

- Delete an Instance

Delete a Backup Set
Delete a Subclient

ADVANCED - VMWARE BACKUP

Choosing Backup Types

Full Backup

Backing Up Individual Files/Folders

Incremental Backups

Differential Backups

Scheduling A Backup

Managing Jobs

Restarting Jobs
Controlling Jobs

Additional Options

ADVANCED - VMWARE RESTORE

Choosing Restore Types

Restoring a Virtual Machine to a Different Virtual Center or ESX Server

Role of a Proxy Computer
Automatically Power On Restored Virtual Machines

Restoring VMDK Files

Attaching Restored VMDK Files to an Alternate Virtual Machine
Mounting VMDK Files as a Complete Volume
Mounting VMDK Files Backed up from VSphere Environments

Restoring Volumes

Restoring Files and Folders

Restore Destinations

Scheduling a Restore

Restores from Unix MediaAgents

Creating Virtual Machines from Restored VMX Files

Browse Data

Browsing Data From Before the Most Recent Full Backup

List Media

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

Restore by Jobs

Manage Restore Jobs

Restart Jobs
Resubmit Jobs
Control Jobs

Additional Restore Options

DATA AGING - VMWARE

DISASTER RECOVERY - VMWARE

Planning for a Disaster Recovery

Restoring Virtual Machine Data

ADDITIONAL OPERATIONS

Audit Trail

Storage Policy

Subclient Policy

Schedule Policy

Auxiliary Copy

Operation Window

Operating System and Application Upgrades

Uninstalling Components

Online Help Links

FREQUENTLY ASKED QUESTIONS

BEST PRACTICES

TROUBLESHOOTING - VMWARE BACKUP

Backup Failure

Incorrect Virtual Machine Is Backed Up During Incremental Backup

Version 4.0 vStorage option, backs up full disk rather than data only copy

TROUBLESHOOTING - VMWARE RESTORE

Restore Failure

Virtual Machine not powering on after restore of Independent Disk/Physical RDM

SAN Mode restore slower than NBD Transport Mode restore

While Restoring Virtual Machine as Template, registering Template Virtual Machine Fails

Recovering Data Associated with Deleted Clients and Storage Policies

VM LIFE CYCLE MANAGEMENT

VM Lifecycle Management - Administrator (Overview)

VM Lifecycle Management - Administrator (Install Virtual Server iDataAgent)

VM Lifecycle Management - Administrator (Install Web Search Server and Client)

VM Lifecycle Management - Administrator (Setup)

VM Lifecycle Management - Administrator (User Information)

VM Lifecycle Management - Administrator (Advanced)

VM Lifecycle Management - End-User (Getting Started)

VM Lifecycle Management - End-User (Advanced)

VIRTUALIZE ME

Virtualizing a Windows Computer - Getting Started

Virtualizing a Windows Computer - Advanced

Virtualizing a Windows Computer - Best Practices

Virtualizing a Windows Computer - FAQ

Virtualizing a Windows Computer - Troubleshooting

Virtualizing a Linux Computer - Getting Started

Virtualizing a Linux Computer - Advanced

Virtualizing a Linux Computer - Best Practices

Virtualizing a Linux Computer - FAQ

Virtualizing a Linux Computer - Troubleshooting

System Requirements - Virtual Server iDataAgent (VMware)

[System Requirements](#) | [Supported Features](#)

The following requirements are for the Virtual Server iDataAgent:

VIRTUAL SERVER APPLICATION SOFTWARE

VMWARE ESX SERVER

VMware ESX/ ESXi 4.x, 5.0, 5.1

vSphere 4.x, 5.0, or 5.1

Before configuring Virtual Server iDataAgent for ESXi servers, consider the following:

- VADP is not available in the free version of ESXi and requires the Essentials licensing level or higher.
- Proxy computers must be running Windows Server 2008 or higher for ESX Server version 5.0 or 5.1.
- VDDK 5.0 is required for this version of ESX server. VDDK 5.0 update 1 is installed automatically with Service pack 6A or higher.

vSphere 5.5 and VDDK 5.5 are supported only in Calypso Version 10. If you need this support, CommVault recommends upgrading to Version 10. You can begin by upgrading the CommServe system, Virtual Server Agent, and MediaAgent, and upgrade the rest of the CommCell environment later.

VMWARE VCENTER SERVER APPLIANCE

VMware vCenter Server Appliance 5.0.0

VIRTUAL MACHINE HARDWARE

Version 4.0, 7.0, 8.0, 9.0

VIRTUAL MACHINE OPERATING SYSTEMS

All Guest Operating Systems supported by VADP and VCB

PROXY COMPUTER OPERATING SYSTEMS

Microsoft Windows Server 2012 R2 is not supported for Calypso Version 9.

Microsoft Windows Server 2012 x64 Editions

Microsoft Windows Server 2008 32-bit and x64 Editions

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

Microsoft Windows 7 32-bit and x64 Editions

Microsoft Windows Vista 32-bit and x64 Editions

Microsoft Windows XP Professional 32-bit and x64 Editions

VM LIFECYCLE MANAGEMENT

The following conditions must be considered when utilizing VM Lifecycle Management:

- VMware vCenter 4.1 is required
- VMware Tools (version 8288) must be installed on templates
- Windows Server 2008, Windows 7, and Linux (RedHat5) are supported

HARD DRIVE

100 GB recommended.

If performing backups with the granular recovery option enabled, the location of the Job Results folder should contain additional space to accommodate at least 2 percent of the total amount of data being backed up.

MEMORY

1 GB RAM minimum required; 4 GB RAM recommended.

PROCESSOR

All Windows-compatible processors supported

PERIPHERALS

DVD-ROM drive

Network Interface Card

MISCELLANEOUS

ALLOCATION UNIT SIZE OF THE NTFS VOLUMES

The cluster size or the allocation unit size of an NTFS volume in a virtual machine must be multiple of 1024 bytes. You can set the cluster size before formatting a volume. The default cluster size is 4096 bytes.

NET FRAMEWORK

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

MICROSOFT VISUAL C++

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

NETWORK

TCP/IP Services configured on the computer.

NOTES ON VIRTUAL SERVER iDATAAGENT INSTALLATIONS FOR VMWARE

For VMware, the Virtual Server iDataAgent is installed on a proxy computer that can communicate with the host computer.

Ensure that the VDDK 5.1 is not installed on any host.

SUPPORT FOR VCB ENVIRONMENT

The following utilities must be installed prior to performing any backup operations in the VCB Environment using the Virtual Server iDataAgent (these packages can be obtained from the VMware web site):

- VMware Consolidated Backup (refer to the VMware customer support web site for support information for the VMware Consolidated Backup utility.)
 - The Virtual Server iDataAgent supports all modes supported by the VMware Consolidated Backup utility.
 - Refer to the VMware customer support web site for support information for the VMware Consolidated Backup utility.
- VMware Converter 4.0.1 (refer to the VMware customer support web site for support information for the VMware Converter utility.)
 - Ensure that the guest operating systems are using the supported version of the VMware Converter.
 - The vCenter Converter Standalone 4.0.1 package must be installed for Windows Vista and Windows Server 2008 environments.

Ensure that the latest version of VMware Consolidated Backup and that VMware Converter 4.0.1 are installed.

- Restores of entire virtual machines from ESX Server 4.0 to a prior version of ESX Server are not supported.
- Virtual machines from ESX Server 3.5 can be backed up through vcbmounter, but they are restored through vStorage.

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

Getting Started - VMware Deployment



The Virtual Server iDataAgent is used to protect VMware virtual machine data. Follow the steps given below to install the Virtual Server iDataAgent.

INSTALLATION

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

METHOD 1: INTERACTIVE INSTALL

Use this procedure to install directly on proxy computer that can communicate with ESX Server.

METHOD 2: INSTALL SOFTWARE FROM COMMCELL CONSOLE

Use this procedure to install remotely on a proxy computer that can communicate with the ESX Server.

METHOD 1: INTERACTIVE INSTALL

1. Run **Setup.exe** from the **Software Installation Disc** on a proxy computer that can communicate with the ESX Server.
2. Select the required language.
Click **Next**.

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

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

5. Expand **Client Modules | Backup & Recovery | File System** and select **Virtual Server Agent**.

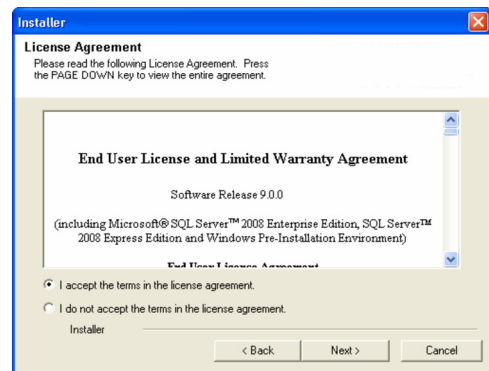

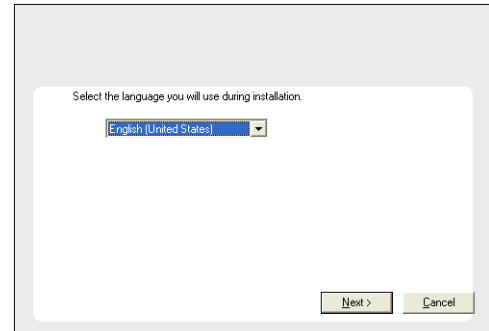
BEFORE YOU BEGIN

Download Software Packages

Download the latest software package to perform the install.

System Requirements

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



Click **Next**.

6. Click **YES** to install Microsoft .NET Framework package.
 - This prompt is displayed only when Microsoft .NET Framework is not installed.
 - Once the Microsoft .NET Framework is installed, the software automatically installs the Microsoft Visual J# 2.0 and Visual C++ redistributable packages.
7. If this computer and the CommServe is separated by a firewall, select the **Configure firewall services** option and then click **Next**.
 For firewall options and configuration instructions, see Firewall Configuration and continue with the installation.
 If firewall configuration is not required, click **Next**.

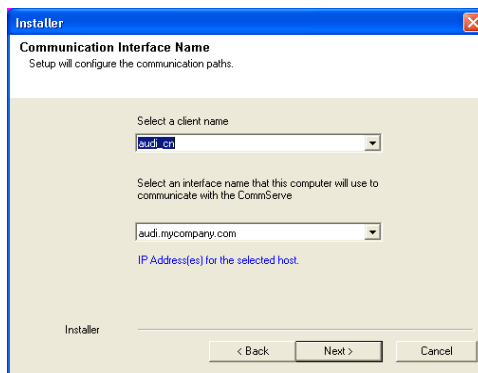
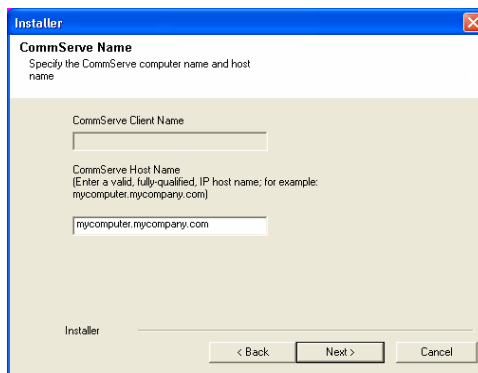
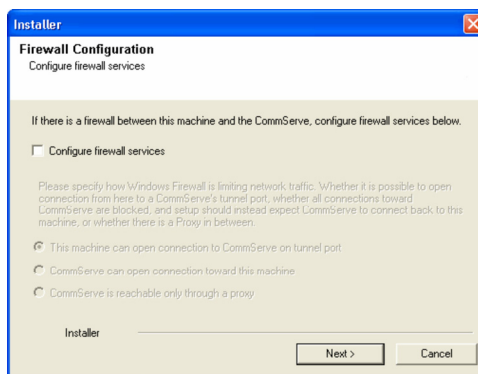

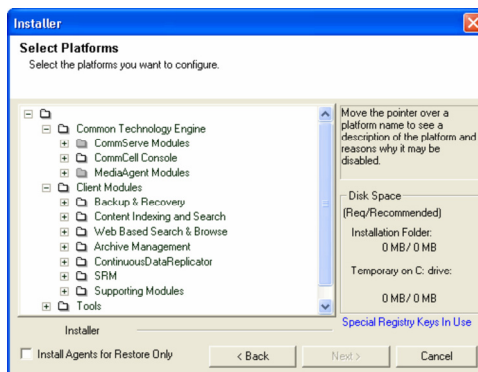
8. Enter the fully qualified domain name of the **CommServe Host Name**.
 Click **Next**.

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

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

9. Click **Next**.

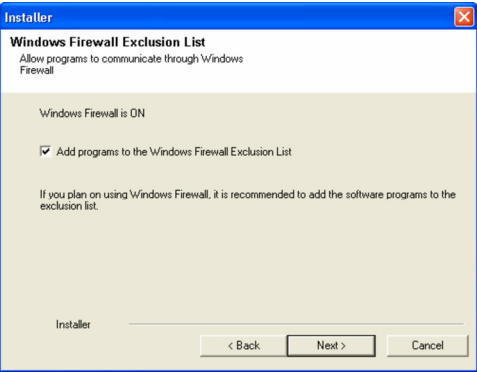
10. Select **Add programs to the Windows Firewall Exclusion List**, to add CommCell programs and services to the Windows Firewall Exclusion List.



Click **Next**.

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

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



11. Verify the default location for software installation.

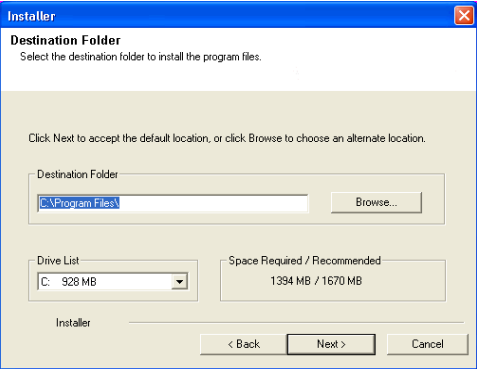
Click **Browse** to change the default location.

Click **Next**.

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

/ : * ? " < > | #

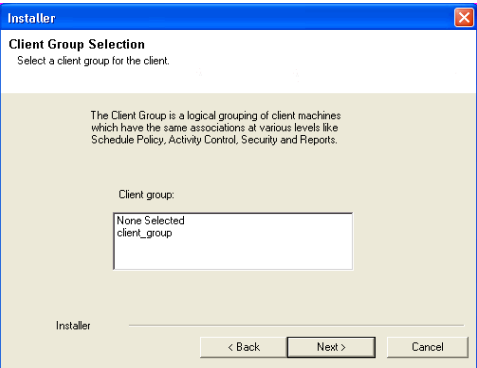
It is recommended that you use alphanumeric characters only.



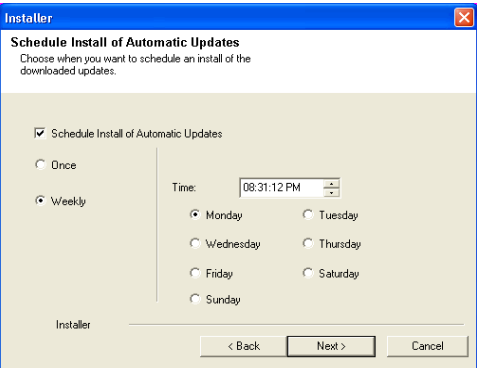
12. Select a Client Group from the list.

Click **Next**.

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



13. Click **Next**.



14. Click **Next**.

15. Select a **Storage Policy**.
Click **Next**.





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

Click **Next**.

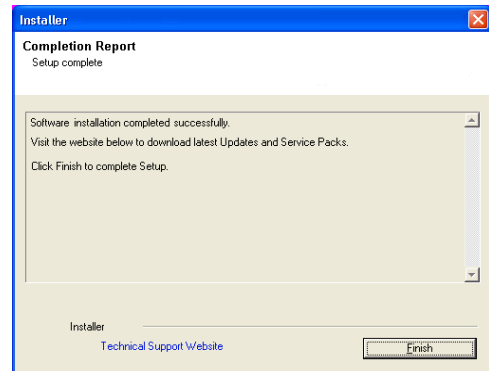
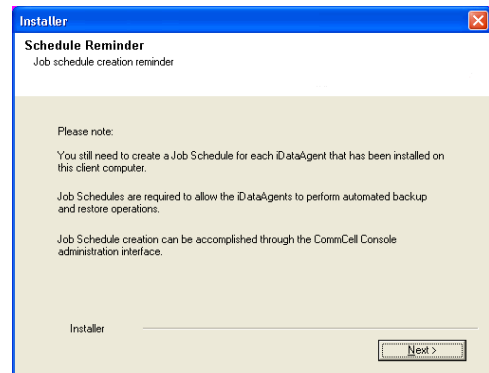
You can create the Storage Policy later in step 19.

16. Click **Next**.

17. Click **Next**.



18. Click **Finish**.



If you already have a storage policy selected in step 15, click the **Next** button available on the bottom of the page to proceed to the **Configuration** section.

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

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

DISK LIBRARY CREATION:

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

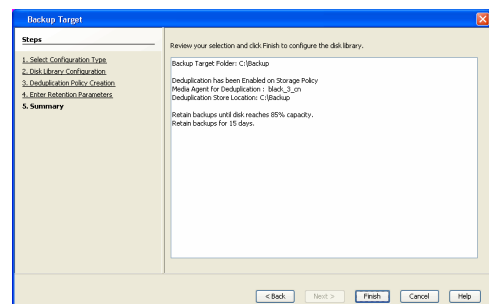
Type the name of the folder in which the disc library must be located in the **Enter backup destination folder** box or click the **Browse** button to select the folder.

Click **Next**.

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

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

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



STORAGE POLICY CREATION:

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

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

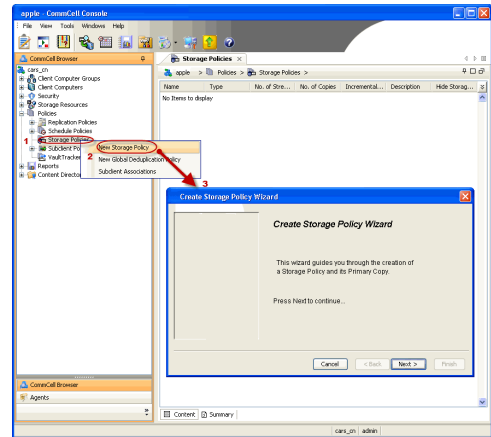
current release.

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

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


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

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

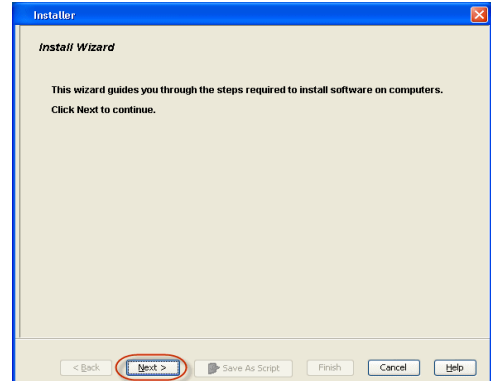


METHOD 2: INSTALL SOFTWARE FROM COMMCELL CONSOLE


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



2. Click **Next**.



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



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

5. Enter the fully qualified domain name of the proxy computer that can communicate with the ESX Server.

For example: apple.domain.company.com

The Virtual Server iDataAgent will be installed on this proxy computer.

Click **Next**.





6. Click **Next**.

7. Specify **UserName** and **Password** of the proxy computer.


Click **Next**.

8. Select **Virtual Server Agent**.


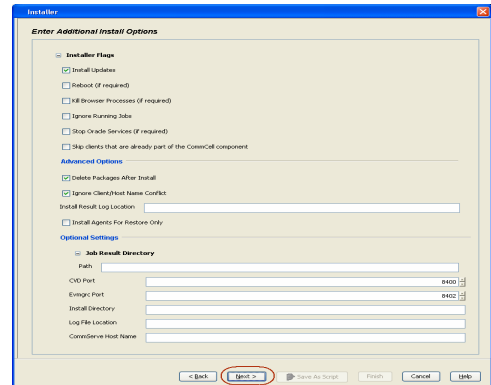
Click **Next**.




9. Click **Next**.




10. Select **Immediate**.
Click **Next**.



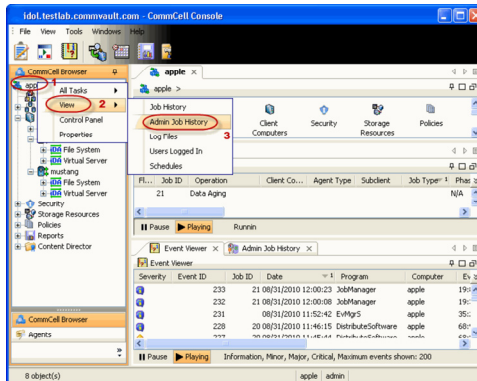
11. Click **Finish**.



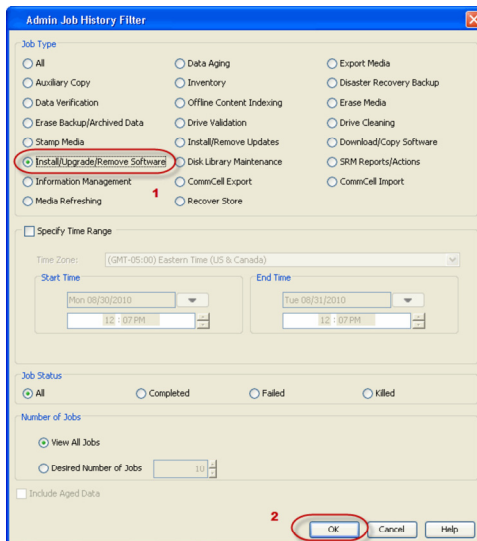
12. You can track the progress of the job from the **Job Controller** or **Event Viewer** window.



13. Once the job is complete, right-click the **CommServe** computer, click **View** and then click **Admin Job History**.




14. Select **Install/Upgrade/Remove Software**.
Click **OK**.



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

- Items that succeeded during the job
- Items that failed during the job
- Details of the job
- Events of the job
- Log files of the job




Getting Started - VMware Configuration

◀ Previous Next ▶

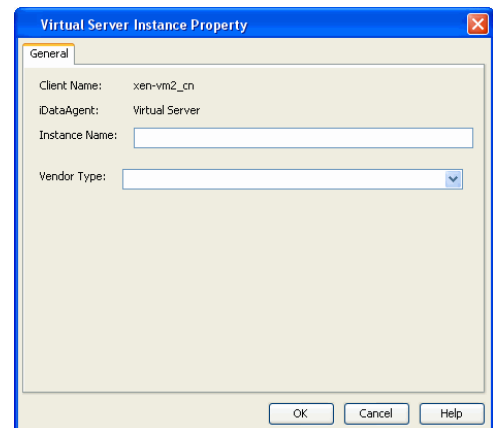
CONFIGURATION

Once the Virtual Server iDataAgent has been installed, configure an Instance, a Backup Set and a Subclient to facilitate backups. The following sections provide the necessary steps required to create and configure these components for a first backup of a Virtual Center.

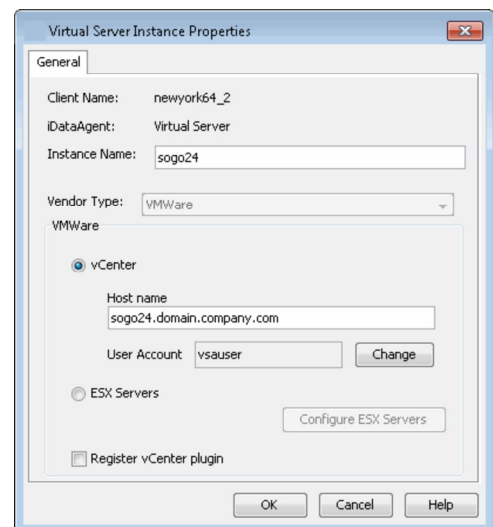
- From the CommCell Browser, navigate to **<Client> | Virtual Server**.
 - Right-click the **Virtual Server | All Tasks** and click **Create New Instance**.



- Enter the **Instance Name**.
 - Select **VMware** from **Vendor Type** menu.





- Click **Virtual Center**.
 - Click **Configure Password**.
 - Enter the username and password associated with the Virtual Center.
 - Ensure that the **Host name** is a fully qualified domain name.




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

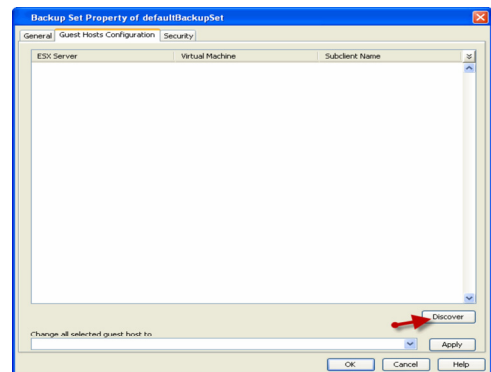
5.
 - From the CommCell Browser, right-click the **Default Backup Set**.
 - Click **Properties**.



6.
 - Ensure the **Automatically add new VMs that do not qualify for membership in any of the subclients** is selected.
 - Click **OK**.



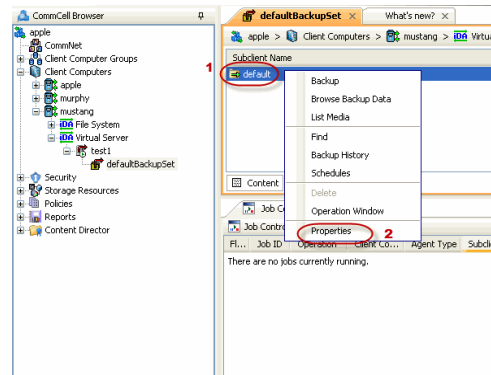
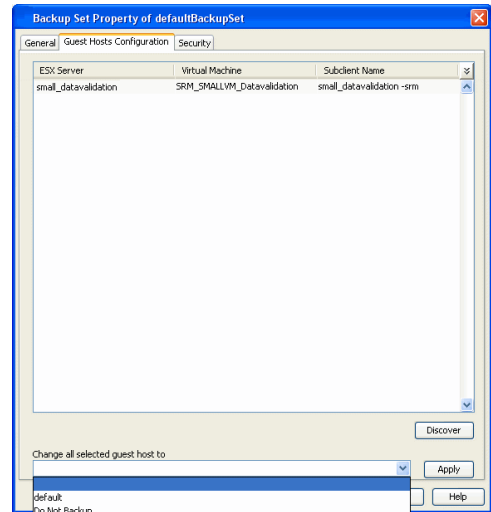
7. Click **Discover** on the **Guest Hosts Configuration** tab.
The discovery process might take several minutes to complete.




8.
 - Select a virtual machine to back up from the **Virtual Machine** column.
 - Select the **Default Subclient** from the **Subclient** column for the virtual machine you want to back up.

- Click **Apply**.
- Click **OK**.

- 9.
- From the CommCell Browser, navigate to the **Default Subclient**.
 - Click **Properties**.



- 10.
- Select **Disk-Level** from the **Backup Type** space.
 - Select a proxy computer from the **Use Proxy** menu.



- 11.
- Click the **Storage Device** tab. Ensure that a Storage Policy is selected.
 - Click **OK**.

If you do not have Storage Policy created, continue with **Step 12**.

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

DISK LIBRARY CREATION:

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

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

Click **Next**.

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

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

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

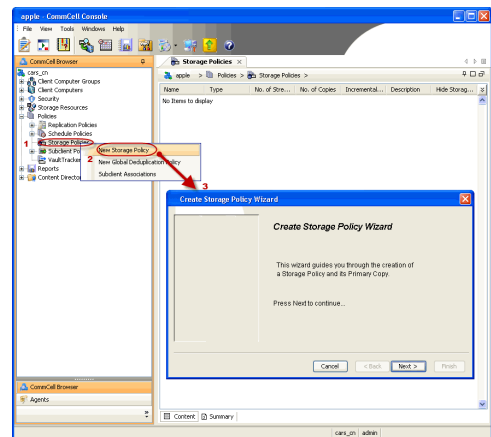
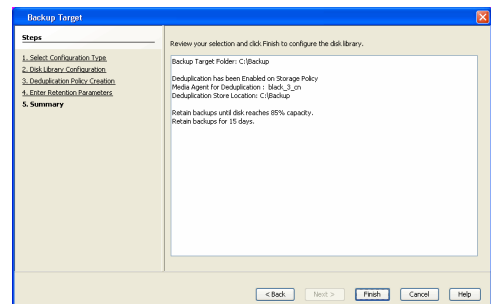
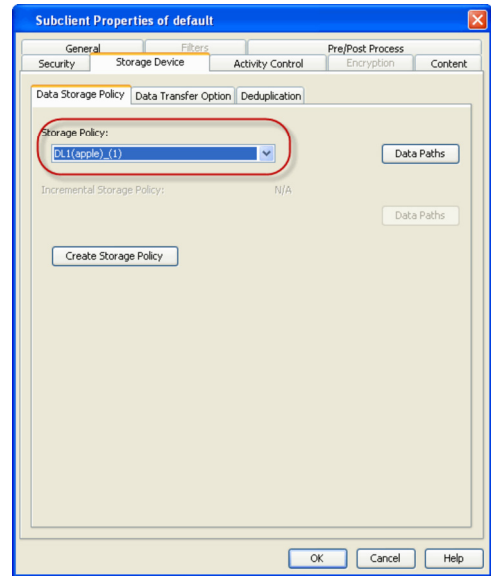
STORAGE POLICY CREATION:

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

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

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

Type the name of the folder in which the deduplication database must be located in the Deduplication Store Location or click the Browse button to



select the folder and then click **Next**.

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

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



Getting Started - VMware Backup

PERFORM A BACKUP

After configuring your Instance, Backup Set and Subclient, you are ready to perform your first backup.

WHAT GETS BACKED UP

The Virtual Server iDataAgent backs up the following:

- Windows Virtual Machines
- Linux Virtual Machines
- Stand-alone ESXI with standard license level
- Volumes
- Files/folders
- Virtual RDMs


WHAT DOES NOT GET BACKED UP

The Virtual Server iDataAgent does not back up the following:

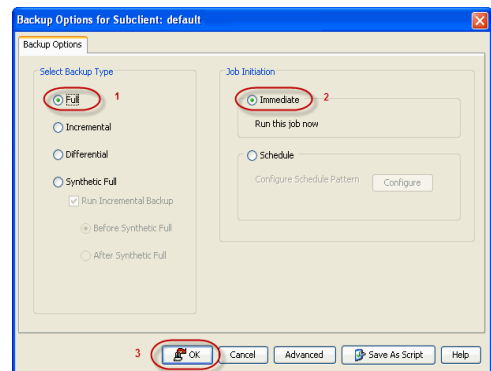
- Virtual machines containing:
 - SCSI bus sharing set to `virtual`
 - Fault Tolerance
- GPT or dynamic disk volumes (only disk and file level backups are supported, volume level backups are not supported).
- Hard links
- VMDKs in the same location with same names, where only case is different.
- Physical RDMs
- Independent Disks

The following section provides step-by-step instructions for running your first full backup of a single virtual machine immediately.


1.
 - From the CommCell Console, navigate to **Client Computers | Virtual Server**.
 - Right-click the **Subclient** and click **Backup**.



2.
 - Select **Full** as backup type and **Immediate** to run the job immediately.
 - Click **OK**.



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






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

- Click **OK**.

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

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



Getting Started - VMware Restore

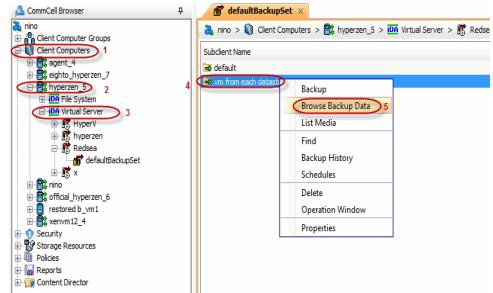


PERFORM A RESTORE

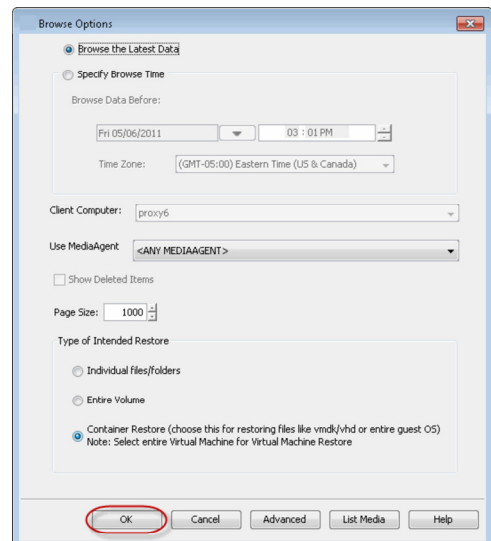
It is recommended that you perform a restore operation immediately after your first full backup to understand the process.

The following sections describe the steps involved in restoring a virtual machine to a different Virtual Center/ESX Server.


- From the CommCell Console, navigate to the <Client> | **Virtual Server**.
 - Right-click the subclient that contains the data you want to restore, point to **Browse Backup Data**.



- Click **OK**.




- Select the virtual machine under the backup set. Its entire contents will be automatically selected in the right pane.
 - Click **Recover All Selected**.



- Select the **Destination ESX Server** to which the virtual machine will be restored.

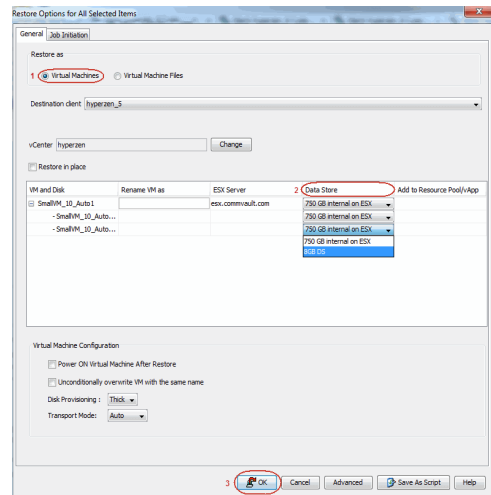
5. Select the **Datastore** to which the disk will be restored.

If the selected datastore does not meet the minimum requirements needed to restore the virtual machine, you can repeat this step until an acceptable datastore is found.




6. Enter the **VM Name** for the virtual machine.


Ensure that you provide a fully qualified name for the virtual machine. Entering an IP address will cause the restore operation to fail.



7.
 - Click the **Job Initiation** tab.
 - Select **Immediate** to run the job immediately.
 - Click **OK**.

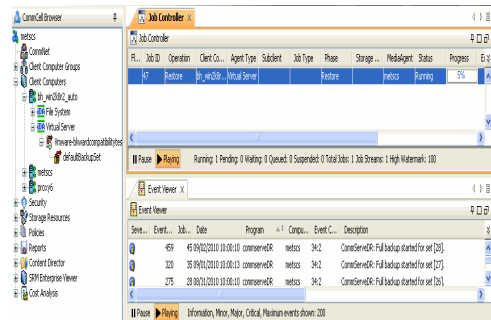


8. You can monitor the progress of the restore job in the **Job Controller** or **Event Viewer** window of the CommCell Console.

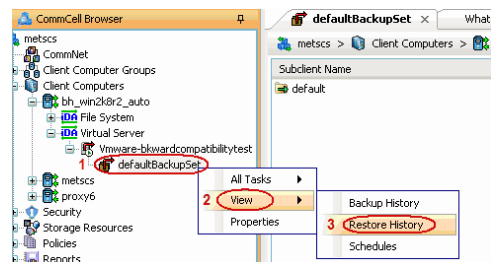


9. Once the restore job has completed, right-click the entity (e.g. agent, instance) and click **View | Restore History**.

If the entity chosen is the client computer, click **View | Job History**.




10. Click **OK**.



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

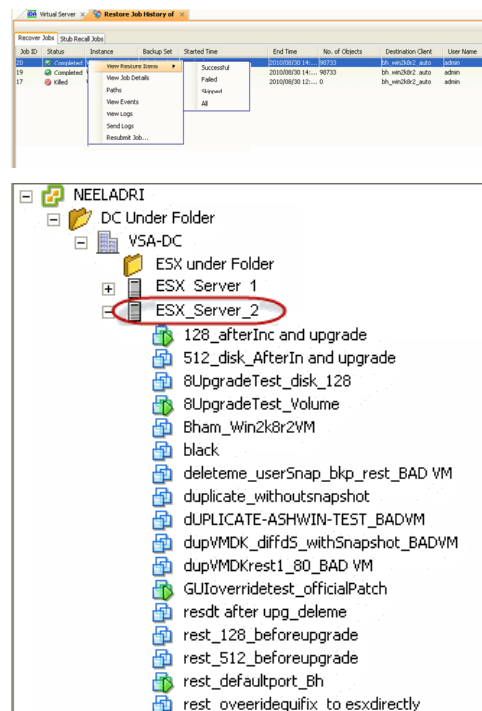
- View Restore Items



You can view them as **Successful, Failed, Skipped** or **All**.

- View Job Details
- View Events of the restore job.
- View Log files of the restore job.

- Once the virtual machine is restored, it is automatically mounted to the virtual center/ESX Server you selected.



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

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

Advanced - VMware Configuration

TABLE OF CONTENTS

Understanding the CommCell Console

Creating User-Defined Instances

Creating User-Defined Backup Sets

Creating User-Defined Subclients

Configuring User Accounts

- Virtual Center Credentials
- Permissions for Custom User Accounts

Configuring Backups for vSphere VADP Environments

Configuring Backups for Combined vSphere VADP and VCB Environments

Configuring Backups for VCB Environments

Configuring Backups of ESX Servers

Configuring Backups of Volumes

Configuring Backups of Files and Folders

Configuring Automatic Discovery of Virtual Machines

- Discover Non-Associated Virtual Machines to Default Subclient
- Discover By Regular Expression
- Discover By Single-ESX Server Affinity
- Discover By Multiple-ESX Server Affinity
- Discover By Cluster Affinity
- Discover by Data Center Affinity
- Discover By DataStore Affinity
- Discover By Resource Pool Affinity
- Discover By OS Group
- Discover By vApp Affinity

Configuring Transport Modes

- Forcing Transport Modes in vSphere VADP Environments

Pre/Post Processing using VMware Tools

Modifying an Agent, Instance, or Subclient

Deleting an Instance, Backup Set, or Subclient

- Delete an Instance
- Delete a Backup Set
- Delete a Subclient


UNDERSTANDING THE COMMCELL CONSOLE

The Virtual Server iDataAgent uses three main components to perform backup and restore operations from the CommCell Console.

Instance - defines the Virtual Center to be backed up.

Backup Set - collective information of all virtual machines contained in the virtual center.

Subclient - defines the virtual machines to be backed up.



CREATING USER-DEFINED INSTANCES

While the Virtual Server iDataAgent is capable of backing up the entirety of your VMware environment using a single instance, backup set, and subclient configuration, it is often not practical to do so. For example, a Virtual Center may contain as few as two ESX Servers, each containing no more than 10 virtual machines. In this scenario, a single instance configured for the entire Virtual Center may be sufficient. However, if a Virtual Center contains upwards of 50 ESX Servers, each containing at least 20 virtual machines, a backup of a single Virtual Center instance may take a significant amount of time to complete. As such, it is important to consider how much data must be backed up at a given time.

As a best practice, consider the following recommendations:

- If creating an instance for an entire Virtual Center containing a large number of ESX Servers and virtual machines, consider performing backups at the backup set or subclient level. These backups can be staggered throughout the day to ensure minimal resources are consumed.
- If backups for a Virtual Center instance must be performed at the instance level, consider utilizing the Incremental and Differential backup options described

in the **Getting Started** documentation.

- If it is desirable to be able to perform backups at the instance level for both the Virtual Center and individual ESX Servers, consider creating one instance for the Virtual Center and additional instances for each ESX Server. This provides the flexibility to perform comprehensive backups for both the Virtual Center and ESX Servers as needed.
1. From the CommCell Browser, right-click the agent and select **Create New Instance**.
 2. Follow the procedure described in the **Getting Started** configuration documentation.

CREATING USER-DEFINED BACKUP SETS

Because backup sets dictate the data assigned to each subclient, it is recommended that backup sets be configured with the number of anticipated subclients in mind. For example, if an ESX Server has 100 virtual machines, all of which need to be backed up, then consider creating two backup sets. Each backup set can then contain five subclients with 10 virtual machines assigned to each. Such a scenario provides the flexibility of performing narrower backups staggered at different times, saving time and resources.

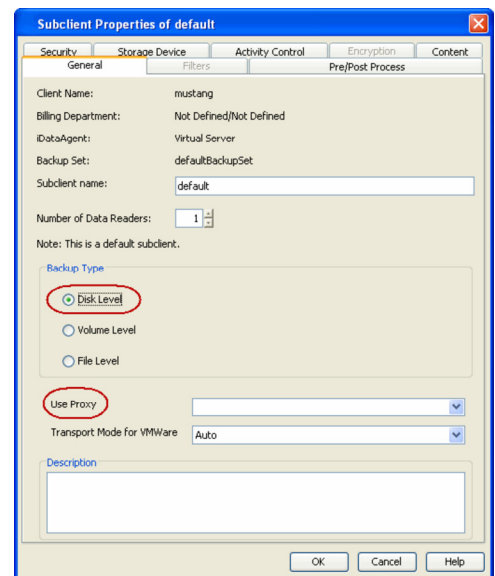
1. From the CommCell Browser, right-click the instance for which you want to create a new backup set and select **New Backup Set**.
2. Follow the procedure described in the **Getting Started** configuration documentation.

CREATING USER-DEFINED SUBCLIENTS

For Virtual Centers or ESX Servers that have a large number of virtual machines, it is recommended that the virtual machines be distributed across more than one subclient. This helps to enhance backup and recovery performance for each subclient, as well as preserves network resources. Additionally, user-defined subclients also provide the facility to group virtual machine backups of a specific entity, such as an ESX Server or datastore, together.

To further enhance performance in environments where a proxy server is used, consider assigning a different proxy server to each subclient. This ensures no single computer is given a substantially higher load.

1. From the CommCell Browser, right-click the backup set for which you want to create a new subclient and select **New Subclient**.
2. Enter the name for the subclient in the **Subclient Name** text box.
3. Click **Content** tab.
4. Select the virtual machines to be added to the subclient.
5. Click **OK**.



CONFIGURING USER ACCOUNTS

The Virtual Server iDataAgent requires user accounts that have sufficient privileges for the software to:

- Access the Virtual Center and ESX Servers
- Access virtual machines
- Access volumes, files, and folders within virtual machines

An administrative account configured with the **VCB Role** and the following additional privileges can be used:

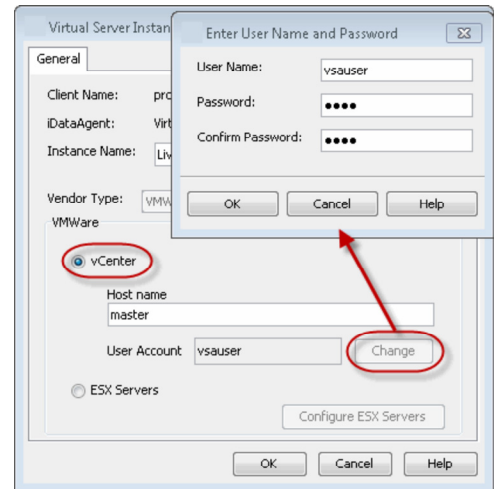
- Virtual Machine
- Resource
- Datastore

User-defined roles can also be created provided the above-mentioned privileges are included in those roles.

The following sections provide the steps necessary to configure user accounts for Virtual Centers and ESX Servers.

VIRTUAL CENTER CREDENTIALS

1. Navigate to **Client Computers | <Client> | Virtual Server**.
2. Right-click the instance and click **Properties**.
3. In the **VMware | vCenter** area, click the **Change** button.
4. Enter the username and password.
The password must not contain single-quote (') or double-quote (") characters.
5. Click **OK** to save your changes.
6. Click **OK**.



PERMISSIONS FOR CUSTOM USER ACCOUNTS

You can create a separate account for backup and restore operations. If you are creating a user account other than administrator, ensure that the following permissions are assigned to the user account:

BACKUP PERMISSION REQUIREMENTS

Category	Available Permissions
Datastore	Allocate space Browse datastore Low level file operations
Global	Disable methods Enable methods Licenses
Virtual machine - Configuration	Add existing disk Add new disk Add or remove device Change resource Disk change tracking Disk lease Remove disk Settings
Virtual machine - Provisioning	Allow read-only disk access Allow virtual machine download Clone virtual machine
Virtual machine - Snapshot management ("Virtual machine - State" in vSphere 4.1)	Create snapshot Remove Snapshot

RESTORE PERMISSION REQUIREMENTS

Category	Available Permissions
Datastore	Allocate space
Network	Assign network
Resource	Assign vApp to resource pool Assign virtual machine to resource pool
Virtual machine - Configuration	Add existing disk Add new disk Add or remove device Advanced Change CPU Count

	Change resource Disk change tracking Disk lease Host USB device Memory Modify device settings Raw device Reload from path Remove disk Rename Reset guest information Settings Swapfile placement Upgrade virtual machine compatibility ("Upgrade virtual hardware" in vSphere 4.1)
Virtual machine - Interaction	Power Off Power On
Virtual machine - Inventory	Create new Register Remove Unregister
Virtual machine - Provisioning	Allow disk access Allow read-only disk access Allow virtual machine download Mark as template (to restore VM template)
Virtual machine - Snapshot Management ("Virtual machine - State" in vSphere 4.1)	Create snapshot Remove Snapshot Revert to snapshot

VM FILE RECOVERY PLUGIN PERMISSIONS

Category	Available Permissions
Extension	Register extension Unregister extension Update extension
Global	Log event Set Custom attribute Diagnostics

VM LIFECYCLE MANAGEMENT PROVISIONING REQUIREMENTS

Category	Available Permissions
Datastore	Allocate space Browse datastore Low level file operations Remove file Update virtual machine files
Resource	Assign virtual machine to resource pool
Virtual machine - Configuration	Add existing disk Add new disk Add or remove device Advanced Change CPU count Change resource Display connection settings (not in vSphere 4.1) Extend virtual disk

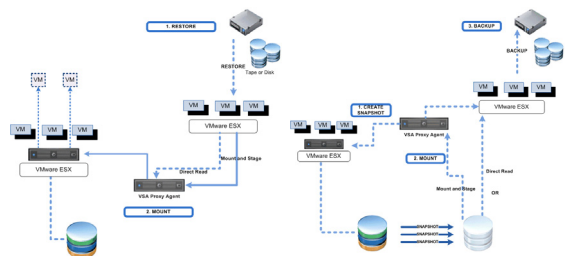
	Memory Modify device settings Raw device Reload from path Remove disk Rename Reset guest information Set annotation (not in vSphere 4.1) Settings Swapfile placement Unlock virtual machine Upgrade virtual machine compatibility ("Upgrade virtual hardware" in vSphere 4.1)
Virtual machine - Guest Operations (not in vSphere 4.1)	Guest Operation Modifications Guest Operation Program Execution Guest Operation Queries
Virtual machine - Interaction	Power Off Power On Reset Suspend
Virtual machine - Inventory	Create from existing Move Remove
Virtual machine - Provisioning	Allow disk access Allow read-only disk access Allow virtual machine download Clone template Clone virtual machine Customize Deploy template Mark as virtual machine Modify customization specification Promote disks Read customization specifications
Virtual machine - Snapshot management ("Virtual machine - State" in vSphere 4.1)	Create snapshot Remove Snapshot Rename Snapshot Revert to snapshot

CONFIGURING BACKUPS FOR VSPHERE VADP ENVIRONMENTS


Backups using vSphere VADP environment provide significant performance improvements over VCB.

During backups, the Virtual Server iDataAgent requests a snapshot of the virtual machine hosted on the VMFS datastore. The snapshot VMDK and virtual machine configuration files are then copied to the storage media without requiring any dedicated disk cache on the proxy server. In the case of incremental backups, Change Block Tracking (CBT) helps quickly identify the data blocks on the virtual machine that have changed since the last backup.

Similarly, during restores, virtual machines are restored directly to the appropriate ESX Server and datastore without the need for staging on the proxy server with VMware Converter. This capability provides for much faster restores.



1. From the CommCell Browser, right-click the backup set and click **Properties**.
2. Select **vStorage**.
3. Click **OK**.
4. From the CommCell Browser, right-click the subclient click **Properties**.
5. Select the proxy computer from the **Use Proxy** menu.



- Click **OK**.

CONFIGURING BACKUPS FOR COMBINED vSPHERE VADP AND VCB ENVIRONMENTS

You can also perform backup operations from an environment wherein both technologies exist. To accomplish this, simply ensure the following components are installed on the proxy computer:

- VMware Consolidated Backup (VCB)
- vSphere VADP
- the Virtual Server iDataAgent

Once these components are prepared, backups will leverage between VCB and vSphere as appropriate.

- From the CommCell Browser, right-click the backup set and click **Properties**.
- Select **Automatic**.
- Click **OK**.
- From the CommCell Browser, right-click the subclient and click **Properties**.
- Select the proxy computer from the **Use Proxy** menu.
- Click **OK**.

The vStorage API can be used to restore backups performed with VCB. The backup data will first be staged to the proxy and then vStorage will perform the restore.

CONFIGURING BACKUPS FOR VCB ENVIRONMENTS

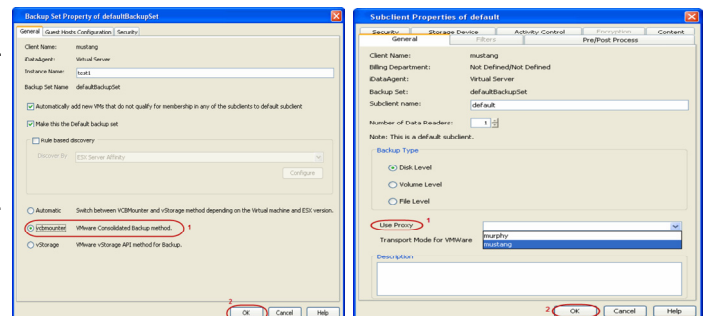
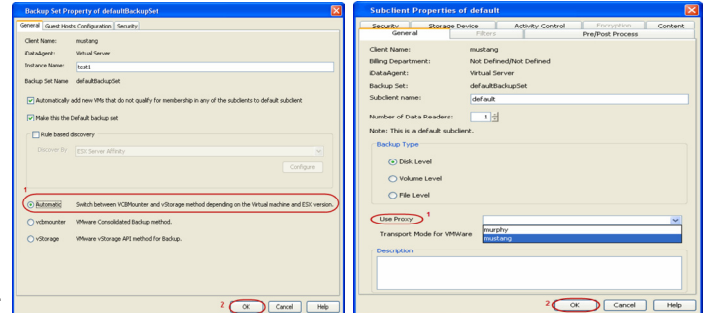
For ESX Servers prior to version 4.x, backups using VCB environment requires a proxy server to copy the data to be processed before being backed up or restored. In addition to ample disk space for copying the virtual machine data, the proxy server also requires the Virtual Server iDataAgent, VMware Consolidated Backup, and VMware Converter to be installed to process all the data prior to moving to the backup media or restore destination.

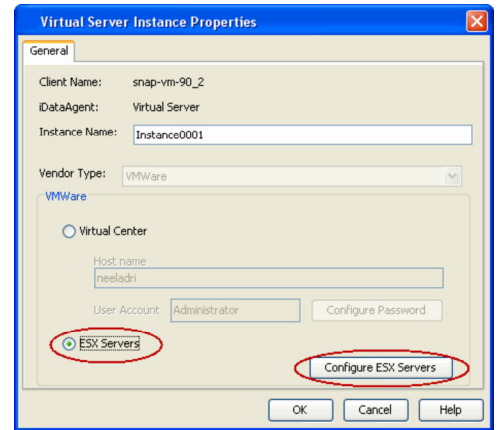
- From the CommCell Browser, right-click the backup set and click **Properties**.
- Select **vcbmounter**
- Click **OK**.
- From the CommCell Browser, right-click the subclient and click **Properties**.
- Select the proxy computer from the **Use Proxy** menu.
- Click **OK**.

CONFIGURING BACKUPS OF ESX SERVERS

Backups of entire ESX Servers consist of backing up all virtual machines that reside on the ESX Server. This option is useful in scenarios where each ESX Server contains a number of virtual machines that must be protected together.

- From the CommCell Browser, right-click the Virtual Server **Instance** and click **Properties**.
- Select **ESX Servers**.
- Click **Configure ESX Server**.
- Click **Add**.
- Enter the following:
 - ESX Server name**
 - ESX User Name** and **Password**. Re-confirm the password.
- Click **OK**.

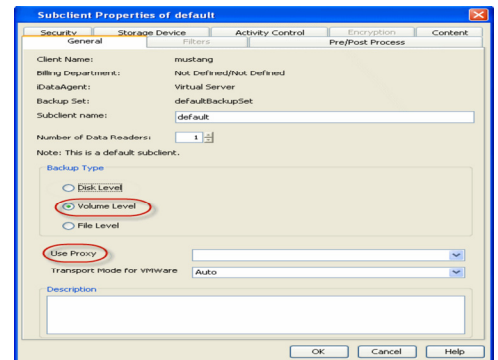




CONFIGURING BACKUPS OF VOLUMES

You can create separate subclients to backup specific volumes of the virtual machine data. Use this option if you do not want to backup an entire virtual machine.

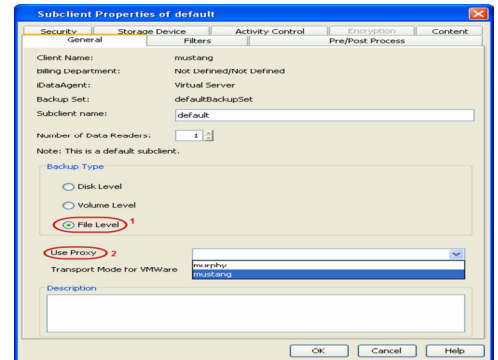
1. From the CommCell Browser, navigate to the **Subclient** and click **Properties**.
2. Select **Volume-Level** from the Backup Types.
3. Select the proxy computer from the **Use Proxy** menu.
4. Click **Storage Device** tab. Ensure that a Storage Policy is selected.
5. Click **Ok**.



CONFIGURING BACKUPS OF FILES/FOLDERS

In addition to backing up virtual server disks and volumes, you can also create separate subclients to backup specific files/folders in the virtual machine. Use this option if you only need to back up a select number of important files rather than an entire virtual machine or volume.

1. From the CommCell Browser, navigate to the **Subclient** and click **Properties**.
2. Click **Properties**.
3. Select **File-Level** from the Backup Types.
4. Select the proxy computer from the **Use Proxy** menu.
5. Click **Storage Device** tab. Ensure that a Storage Policy is selected.
6. Click **Ok**.



CONFIGURING AUTOMATIC DISCOVERY OF VIRTUAL MACHINES

While virtual machine discovery can always be performed manually, it may be useful to discover virtual machines based on a specific set of criteria, particularly in environments where virtual machines are added, removed, or relocated on a routine basis. The following sections describe several methods of discovering virtual machines based on specific criteria, as well as the configuration steps required.

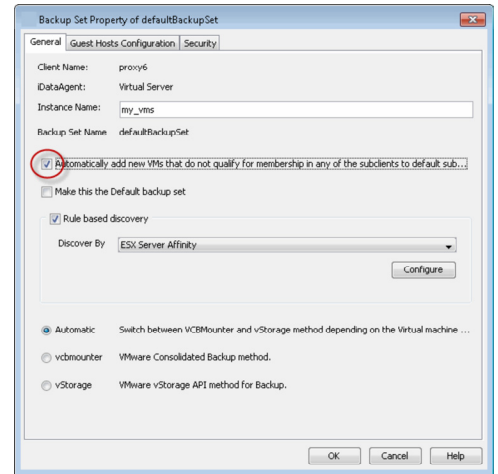
Virtual machines in HotAdd configurations should always be excluded from subclient content.

DISCOVER NON-ASSOCIATED VIRTUAL MACHINES TO DEFAULT SUBCLIENT

While rule-based discovery provides certain levels of granularity in populating virtual machines to back up, there may be scenarios where an entire ESX Server or Virtual Center requires complete protection on a routine basis. The **Automatically add new VMs that do not qualify for membership in any of the subclients** option provides the facility to assign any virtual machines found during discovery to the default subclient provided the virtual machine is not already configured in another subclient.

Note, however, that selecting this option overrides any previously configured regular expressions, ESX Server affinity, or the selection of **Do Not Backup** for any virtual machines.

1. From the CommCell Browser, right-click the backup set and select **Properties**.
2. Select **Automatically add new VMs that do not qualify for membership in any of the subclients**.
3. Click **OK**.



DISCOVER BY REGULAR EXPRESSION

When configured, virtual machines matching the wildcard expressions are automatically either assigned to or filtered from the associated subclient.

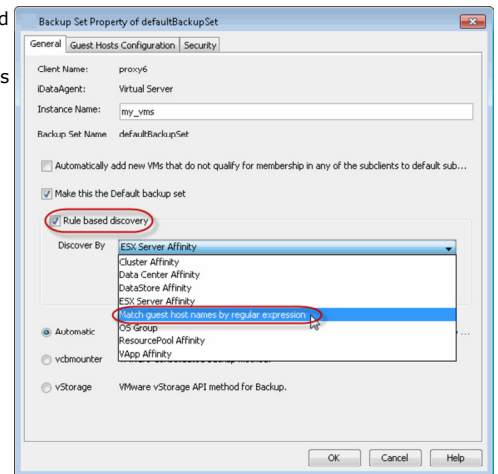
1. From the CommCell Browser, right-click the backup set for which you want to discover guest hosts and click **Properties**.
2. Select **Rule-Based Discovery**.
3. Click **Discover By** drop-down box and select **Match Guest Host Names By Regular Expression**.
4. Click **Configure**.
5. Highlight the desired subclient name.
6. Click **Configure**.
7. Type the name of the host name in the **Guest host name** text box.
8. Click **Add** to add the name to the **Add Guest host name** list.

Some examples of wildcard filters are:

- o *exchangeVM* for a virtual machine named **exchangeVM**
- o *vm* for virtual machines containing names that begin with **vm**

9. Click **OK**. The name is added to the Wildcard names for Virtual Machines column.
10. Click **OK** to save your changes.
11. Click **OK**.

Add a filter for the desired operating system to the **Wildcard names for OS** column for the subclient to which you want to apply the filter.



DISCOVER BY SINGLE-ESX SERVER AFFINITY

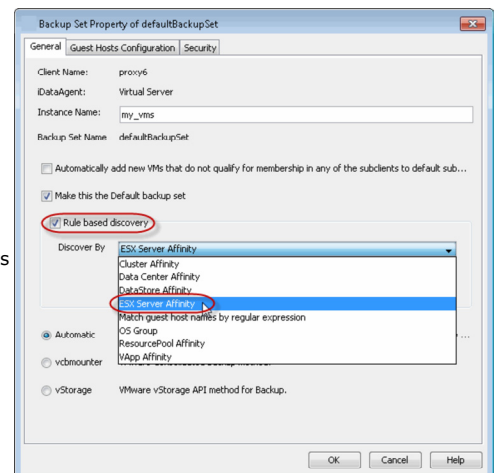
When configured, new virtual machines found on the ESX Server are automatically assigned to the desired subclient and backed up when backup operations on the subclient are performed. This capability is useful in scenarios where:

- a large number of virtual machines exist, rendering the manual process of adding them to a subclient unfeasible.
- virtual machines are added to an ESX server on a frequent basis.

Optionally, you can choose to exclude specific virtual machines from being added without deselecting this option. However, note that these virtual machines will not be backed up unless each is reassigned to the desired subclient.

1. From the CommCell Browser, right-click the backup set for which you want to discover guest hosts and then click **Properties**.
2. Select **Rule-Based Discovery**.
3. Click **Discover By** drop-down box and select **ESX Server Affinity**.
4. Click **Configure**.
5. Select the ESX Server you wish to configure for auto discovery and select the subclient you wish to associate with the ESX Server.

To exclude a select number of virtual machines from being backed up, choose the **Do Not Backup** option in the **Subclient** column.



6. Click **OK** to save your changes.
7. Click **OK**.

DISCOVER BY MULTIPLE-ESX SERVER AFFINITY

In some cases, having access to other ESX Servers containing virtual machines you want to include in your backup. To accomplish this:

1. Creating a dedicated instance for multiple-ESX Server affinity with the instance type defined for the **Virtual Center**.
2. Configure ESX Server Affinity as discussed in By Single-ESX Server Affinity.
3. Click **Configure ESX Server**.
4. Manually add the additional virtual machines from the second ESX Server.

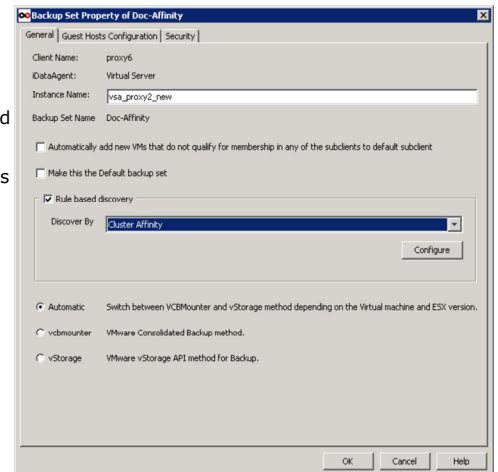
Once configured, all virtual machines from the pre-configured ESX Server, as well as the virtual machines selected from the second ESX Server, will be backed up during the next backup operation.

DISCOVER BY CLUSTER AFFINITY

When configured, new virtual machines found on the cluster are automatically assigned to the desired subclient and backed up when backup operations on the subclient are performed. This capability is useful if your virtual machines are strategically arranged among clusters in your environment.

Optionally, you can choose to exclude specific clusters from being added without deselecting this option. However, note that these clusters will not be backed up unless each is reassigned to the desired subclient.

1. From the CommCell Browser, right-click the backup set for which you want to discover guest hosts and then click **Properties**.
2. Select **Rule-Based Discovery**.
3. Click **Discover By** drop-down box and select **Cluster Affinity**.
4. Click **Configure**.
5. Click **Discover**.
6. Select the cluster you wish to configure for auto discovery and select the subclient you wish to associate with it.
To exclude a select number of clusters from being backed up, choose the **Do Not Backup** option in the **Subclient** column.
7. Click **OK** to save your changes.
8. Click **OK**.

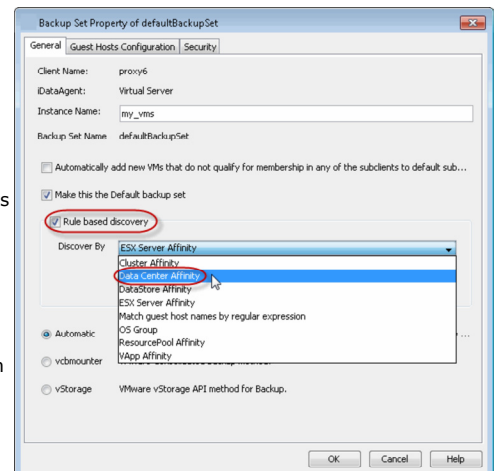


DISCOVER BY DATA CENTER AFFINITY

When configured, new virtual machines found on a Data Center are automatically assigned to the desired subclient and backed up when backup operations on the subclient are performed. This capability is useful if your virtual machines are strategically arranged among Data Centers in your environment.

Optionally, you can choose to exclude specific Data Centers from being added without deselecting this option. However, note that these Data Centers will not be backed up unless each is reassigned to the desired subclient.

1. From the CommCell Browser, right-click the backup set for which you want to discover guest hosts and then click **Properties**.
2. Select **Rule-Based Discovery**.
3. Click **Discover By** drop-down box and select **Data Center Affinity**.
4. Click **Configure**.
5. Select the Data Centers you wish to configure for auto discovery and select the subclient you wish to associate with it.
To exclude a select number of Data Centers from being backed up, choose the **Do Not Backup** option in the **Subclient** column.
6. Click **OK** to save your changes.
7. Click **OK**.



DISCOVER BY DATASTORE AFFINITY

When configured, new virtual machines found in the datastore are automatically assigned to the desired subclient and backed up when backup operations on the subclient are performed. This capability is useful if you want to limit discovery of virtual machines to a specific datastore in scenarios where:

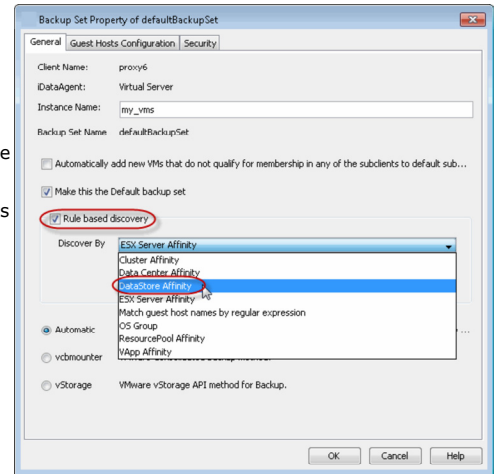
- a particular datastore houses virtual machines of vital importance or that are prone to considerable change.
- you have a large number of datastores and intend to limit your backup and restore operations to a select group of datastores.

Optionally, you can choose to exclude specific virtual machines from being added without deselecting this option. However, that these virtual machines will not be backed up unless each is reassigned to the desired subclient.

1. From the CommCell Browser, right-click the backup set for which you want to discover guest hosts and then click **Properties**.
2. Select **Rule-Based Discovery**.
3. Click **Discover By** drop-down box and select **DataStore Affinity**.
4. Click **Configure**.
5. From the **DataStore Configuration for Auto Discovery** dialog box, locate the datastores you wish to configure for auto discovery and select the subclient you wish to associate each with.

To exclude a select number of datastores from being backed up, choose the **Do Not Backup** option in the **Subclient** column.

6. If necessary, click **Configure** to add, edit, or remove ESX Servers.
7. Click **OK** to save your changes.
8. Click **OK**.



DISCOVER BY RESOURCE POOL AFFINITY

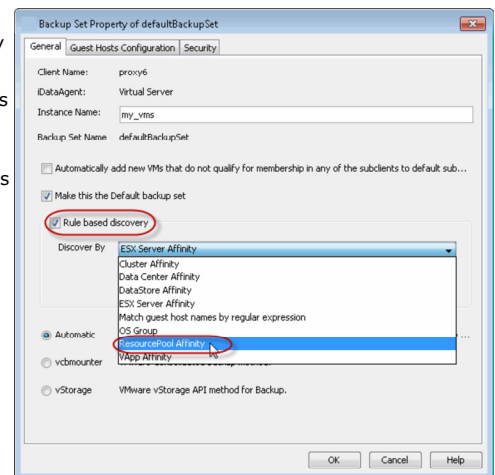
When configured, new virtual machines found in the resource pool are automatically assigned to the desired subclient and backed up with backup operations on the subclient are performed. This capability is useful in scenarios where virtual machines are added to a datastore server on a frequent basis.

Optionally, you can choose to exclude specific resource pools from being added without deselecting this option. However, that these virtual machines will not be backed up unless each is reassigned to the desired subclient.

1. From the CommCell Browser, right-click the backup set for which you want to discover guest hosts and then click **Properties**.
2. Select **Rule-Based Discovery**.
3. Click **Discover By** drop-down box and select **Resource Pool Affinity**.
4. Click **Configure**.
5. From the **Resource Pool Configuration for Auto Discovery** dialog box, locate the resource pools you wish to configure for auto discovery and select the subclient you wish to associate each with.

To exclude a select number of resource pools from being backed up, choose the **Do Not Backup** option in the **Subclient** column.

6. Select the ESX Server that should be used to discover the resource pools.
7. If necessary, click **Configure** to add, edit, or remove ESX Servers.
8. Click **OK** to save your changes.
9. Click **OK**.



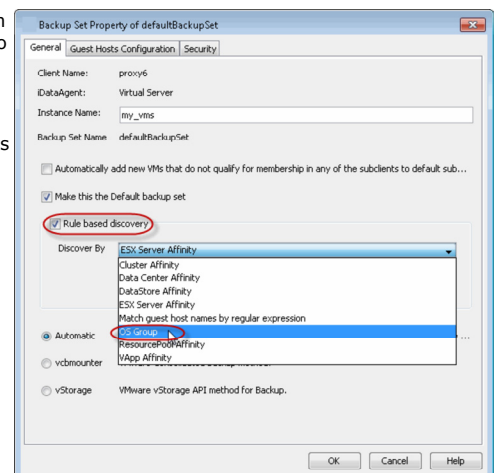
DISCOVER BY OS GROUP

If your virtual machines consist of many different operating systems, it may be useful to discover them based on the operating systems they are running. The Virtual Server iDataAgent provides the facility to select the virtual machine operating systems that should be included in the subclient. When discovery operations are later performed, only virtual machines running the operating systems selected will be incorporated into the subclient content for backups.

1. From the CommCell Browser, right-click the backup set for which you want to discover guest hosts and then click **Properties**.
2. Select **Rule-Based Discovery**.
3. Click **Discover By** drop-down box and select **OS Group**.
4. Click **Configure**.
5. Add a filter for the desired operating system to the **Wildcard names for OS** column for the subclient to which you want to apply the filter.

Some examples of wildcard filters are:

- o *win* for all Windows operating systems
- o *sus* for all SuSE Linux operating systems
- o *lin* for all Linux operating systems



- Click **OK**.

DISCOVER BY VAPP AFFINITY

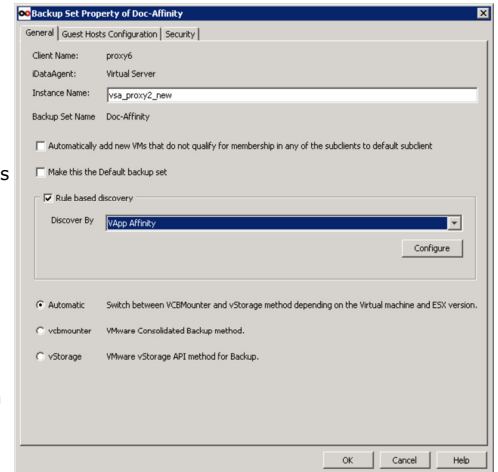
When configured, new virtual machines associated with the selected vApps are automatically assigned to the desired subclient and backed up when backup operations on the subclient are performed. This capability is useful if you want to back up virtual machines sharing similar vApps at the same time.

Optionally, you can choose to exclude specific vApps from being added without deselecting this option. However, note that these vApps will not be backed up unless each is reassigned to the desired subclient.

- From the CommCell Browser, right-click the backup set for which you want to discover guest hosts and then click **Properties**.
- Select **Rule-Based Discovery**.
- Click **Discover By** drop-down box and select **vApp Affinity**.
- Click **Configure**.
- Select the vApps you wish to configure for auto discovery and select the subclient you wish to associate with it.

To exclude a select number of vApps from being backed up, choose the **Do Not Backup** option in the **Subclient** column.

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



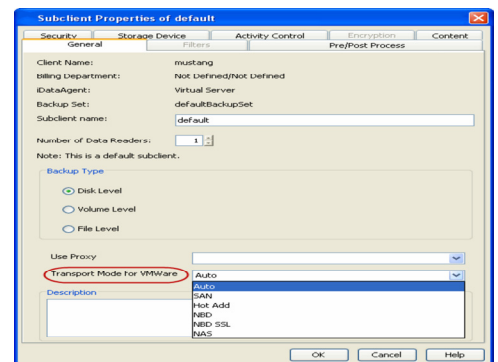
CONFIGURING TRANSPORT MODES

By default, the software automatically determines the transport mode currently being used in your environment prior to each backup operation. This process can be eliminated by configuring the subclient to use a specific transport mode for all backups of that subclient. Doing so negates the need for the software to search your environment for the appropriate mode, thereby increasing backup performance.

- From the CommCell Browser, right-click the subclient and click **Properties**.
- Click the **Transport Mode for VMware** drop-down box and select any of the following transport modes:
 - o **SAN**
 - o **HotAdd**
 - o **NAS**
 - o **NBD SSL**
 - o **NBD**

If you want the software to automatically set the transport mode based on your setup, select **Auto**.

- Click **OK**.
 - In VCB environments, the **NAS** transport mode is not attempted during auto discovery operations. However, you can force the NAS mode by selecting it in the CommCell Console.
 - If you are performing backup of VMware templates, the **NBD** mode is used during the backup operation irrespective of any transport mode that you have selected.



FORCING TRANSPORT MODES IN VSPHERE VADP ENVIRONMENTS

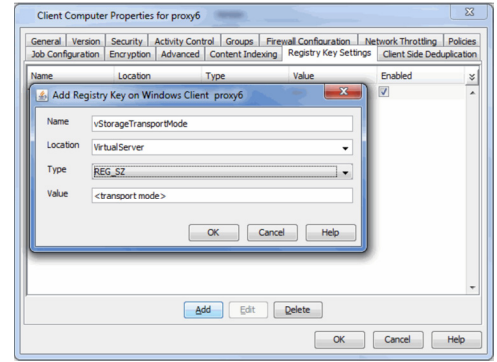
In vSphere VADP environments, you can force one of the following transport modes to be used for all virtual machine backups by configuring the vStorageTransportMode registry key. The following transport modes can be forced using this registry key:

- SAN
- HotAdd
- NBD

Follow the steps below to configure this key:

- From the CommCell Browser, navigate to **Client Computers**.
- Right-click the **<Client>**, and then click **Properties**.
- Click the **Registry Key Settings** tab.
- Click **Add**.
- In the **Name** field, type vStorageTransportMode.
- In the **Location** field, type VirtualServer.

7. In the **Type** list, select **REG_SZ**.
8. In the **Value** field, type the transport mode you want to use in lower-case letters (i.e., san, hotadd, nbd).
9. Click **OK**.



PRE/POST PROCESSING USING VMWARE TOOLS

You can use the VMware tools to perform the any operations before or after the backup. For example: A virtual machine hosts a oracle database and you want to enable the Hot backup of the database before performing the backup, you can run a script using the VMware Tools.

For more information about running scripts on Windows virtual machines, refer to <http://kb.vmware.com/kb/1006671>

On the Linux virtual machines, the script /usr/sbin/pre-freeze-script will be executed when the software snapshot is created and /user/sbin/post-freeze-script will be executed when the software snapshot is removed. Ensure that these scripts are executable by the VMware tools user.

The following scripts are available to perform the Pre/Post processes using VMware tools:

OPERATION	PROCEDURE	SCRIPTS
The scripts can be used to take snapshot of a VM with DB2 Application. This script allows VMWARE to suspend I/O updates to DB2 database till VMWARE LINUX guest snap shot is completed.	<ol style="list-style-type: none"> 1. On a Linux virtual machine, copy the pre-freeze-script and post-thaw-script to /usr/bin directory. 2. Copy write_suspend and write_resume scripts to a location where DB2 database can execute it. (Preferably to a directory under DB2 home) 	pre-freeze-script.sh post-thaw-script.sh write_resume.sh write_suspend.sh
The scripts can be used to take snapshot of a VM with MAXDB Application. This script allows VMWARE to suspend I/O updates to MAXDB database till VMWARE LINUX guest snap shot is completed.	<ol style="list-style-type: none"> 1. On a Linux virtual machine, copy the pre-freeze-script and post-thaw-script to /usr/bin directory. 2. Copy suspend_logwriter and resume_logwriter scripts to a location where MAXDB database can execute it. (Preferably to a directory under MAXDB home) 	pre-freeze-script.sh post-thaw-script.sh resume_logwriter.sh suspend_logwriter.sh
The scripts can be used to take snapshot of a VM with Oracle Application. This script allows VMWARE to suspend I/O updates to Oracle database till VMWARE LINUX guest snap shot is completed.	<ol style="list-style-type: none"> 1. On a Linux virtual machine, copy the pre-freeze-script and post-thaw-script to /usr/bin directory. 2. Copy pre-freeze-script.sql and post-thaw-script.sql scripts to a location where Oracle database can execute it. (Preferably to a directory under Oracle home) <p>These scripts change the entire database to backup mode. If the virtual machine has Oracle iDataAgent installed on it, you can also use the consistent-archivelog-backup.rman script as a sample to run archivelog backup. This will get a consistent snap and also all the latest archived logs and current control file.</p>	pre-freeze-script.sh post-thaw-script.sh pre-freeze-script.sql post-thaw-script.sql consistent-archivelog-backup.rman
The scripts can be used to take snapshot of a VM with Sybase Application. This script allows VMWARE to suspend I/O updates to Sybase database till VMWARE LINUX guest snap shot is completed.	<ol style="list-style-type: none"> 1. On a Linux virtual machine, copy the pre-freeze-script and post-thaw-script to /usr/bin directory. 2. Copy pre-freeze-script.sql and post-thaw-script.sql scripts to a location where Sybase database can execute it. (Preferably to a directory under Sybase home) <p>These scripts execute sybase quiesce commands which will stop updates to databases.</p>	pre-freeze-script.sh post-thaw-script.sh pre-freeze-script.sql post-thaw-script.sql

MODIFYING AN AGENT, INSTANCE, OR SUBCLIENT

Certain properties of Agents, Instances, and Subclients can be modified to accommodate changes to your configuration, data, or desired backup behavior.

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

The following table describes the properties that can be configured from these levels.

OPTION	DESCRIPTION	RELATED TOPICS
--------	-------------	----------------

Job Results Directory	<p>Ensure that the full path name for the Job Results Directory, combined with the VMDK file name, is no greater than 255 characters in length.</p> <ol style="list-style-type: none"> 1. From the CommCell Browser, right-click the icon of the client computer whose job results path you want to change, then click Properties. 2. From the Job Configuration tab of the Client Computer Properties dialog box, click User Name/Password to establish or change the Impersonate User account to access the Job Results Directory. Click OK once you have administered the account. 3. From the Job Configuration tab, type a new job results path in the Job results path field. You can also click Browse to browse to a new job results path from the Browse for Job Result Path dialog box. Click OK. 4. Click OK to save your changes. 	Refer to Job Management.
Change the User Account Details	<p>In order to perform discovery, backup, and restore operations within a single instance, you need administrator-level global, virtual machine, and resource privileges.</p> <p>You can configure the user account associated with the Virtual Center.</p> <ol style="list-style-type: none"> 1. From the CommCell browser, right-click the virtual server instance. 2. Click Properties. 3. Click Configure Password. 4. Type the user name. 5. Type password and confirm password. <p>You can configure the user account associated with the ESX Server.</p> <ol style="list-style-type: none"> 1. Click Configure ESX Server. 2. Select the ESX server. 3. Click Edit. 4. Type the username. 5. Type the password and confirm password. 	
Change Storage Policies	<p>You can modify the storage policies in any of the following situations:</p> <ul style="list-style-type: none"> • To include a different media for the backup operation. • To use a storage policy with a different retention criteria. <p>You can change the storage policies from the subclient level.</p> <ol style="list-style-type: none"> 1. From the CommCell Browser, right-click the subclient. 2. Click Properties. 3. Click Storage Device. 4. Select the Storage policy from the drop-down menu. 5. Click OK. 	Refer to Storage Policies.
Rename a Backup Set or Subclient	<p>You can rename a subclient:</p> <ol style="list-style-type: none"> 1. From the CommCell Browser, right-click the subclient. 2. Click Properties. 3. Type the new name in the Subclient name field. 4. Click OK. 	
Data Transfer Options	<p>You can efficiently configure the available resources for transferring data secured by data protection operations from the subclient level. This includes the following:</p> <ul style="list-style-type: none"> • Enable or disable Data Compression either on the client or the MediaAgent. • Configure the transfer of data in the network using the options for Network Bandwidth Throttling and Network Agents. <p>You can configure the data transfer options.</p> <ol style="list-style-type: none"> 1. From the CommCell Browser, right-click the subclient. 2. Click Properties. 3. Click Storage Device. 4. Click Data Transfer Option tab. 5. Choose the appropriate software compression option for this subclient. 6. Select Throttle Network Bandwidth and set the required bandwidth. 	Refer to Data Compression and Network Bandwidth Throttling.

	7. Click OK .	
View Data Paths	<p>You can view the data paths associated with the primary storage policy copy of the selected storage policy or incremental storage policy. You can also modify the data paths including their priority from the subclient level.</p> <ol style="list-style-type: none"> From the CommCell browser, right-click the subclient. Click Properties. Click Storage Device. Select Storage Policy from the drop-down menu. Click Data Paths. 	
Number of Data Readers	<p>Use this option to specify the number of simultaneous backup data streams allowed for this subclient.</p> <ol style="list-style-type: none"> From the CommCell browser, right-click the subclient. Click Properties Type the Number of Data Readers. Click OK. 	Refer to Streams.
Configure a Subclient for Pre/Post Processing of Data Protection	<p>You can add, modify or view Pre/Post processes for the subclient. These are batch files or shell scripts that you can run before or after certain job phases.</p> <ol style="list-style-type: none"> From the CommCell browser, right-click the subclient. Click Properties. Click Pre/Post Process. Click one of the following phases and type the full path of the process that you want to execute during that phase. Alternatively, click Browse to locate the process (applicable only for paths that do not contain any spaces). <ul style="list-style-type: none"> PreBackup Process PostBackup Process Click OK. Select Run Post Scan Process for all attempts to run a post scan process for all attempts . Select Run Post Backup Process for all attempts to run a post backup process for all attempts. For subclients on Windows platforms, Run As displays Not Selected. If you want to change the account that has permission to run these commands, click Change. <ol style="list-style-type: none"> In the User Account dialog box, select Use Local System Account, or select Impersonate User and enter the user name and password. Click OK. If you selected Local System Account, click OK to the message advising you that commands using this account have rights to access all data on the client computer. 	Refer to Pre/Post Processes.
Configure Activity Control	<p>You can enable backup and restore operations from the agent and subclient level. However, you can enable restore operations only from the agent level.</p> <ol style="list-style-type: none"> From the CommCell browser, right-click the subclient. Click Properties. Click Activity Control, select or clear option(s) as desired. Click OK. 	Refer to Activity Control.
Configure User Security	<p>You can configure user security from the agent or subclient level.</p> <p>You can perform the following functions:</p> <ul style="list-style-type: none"> Identify the user groups to which this CommCell object is associated. Associate this object with a user group. Disassociate this object from a user group. <ol style="list-style-type: none"> From the CommCell browser, right-click the subclient. Click Properties. Click Security. Select the appropriate user groups to which you want to associate to the CommCell object from the Available Groups pane, and then move the user group to the Associated Groups pane. Click OK. 	Refer to User Administration and Security.
Enable/Disable Data Encryption	<p>You can enable data encryption from the subclient level. Encryption must be enabled at the client level prior to configuring any instances residing on that client.</p> <ol style="list-style-type: none"> From the CommCell browser, right-click the subclient. 	Refer to Data Encryption.

	<ol style="list-style-type: none"> 2. Click Properties. 3. Click Encryption. 4. Select the desired encryption. 5. Click OK. 	
View Software Version and Installed Updates	<p>The Version tab, at the Agent level displays the software version and post-release service packs and updates installed for the component.</p> <ol style="list-style-type: none"> 1. From the CommCell browser, right-click the agent. 2. Click Properties. 3. Click Version. 4. Click OK. 	
CommCell Configuration Report	<p>The CommCell Configuration Report provides the properties of the CommServe, MediaAgents, clients, agents, subclients, and storage policies within the CommCell based on the selected filter criteria.</p> <ol style="list-style-type: none"> 1. From the CommCell browser, click Reports icon. 2. Select CommCell Configuration. 3. Click Run. 	Refer to CommCell Configuration.

DELETING AN INSTANCE, BACKUP SET, OR SUBCLIENT

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

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

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

DELETING AN INSTANCE

Consider the following before deleting an instance:

- When you delete a specific instance all job schedules and job histories that pertain to any of the levels within the deleted instance are deleted.
 - You cannot delete an instance if it is being backed up. Attempts to delete an instance under such conditions cause the deletion to fail. If a backup is in progress, either wait for the backup to complete or kill the backup job using the Job Manager. Once the backup is no longer in progress, you can delete the instance level.
 - You cannot delete an instance if there is only one instance present for an agent. To delete the final instance, you must remove the agent software from the client computer.
1. From the CommCell Browser, right-click the instance that you want to delete, click **All Tasks** and then click **Delete**.
 2. Click **Yes** to confirm the deletion. (Clicking **No** cancels the deletion and retains the node.)
 3. Type the requested phrase in the **Enter Confirmation Text** dialog box and click **OK**. This should delete the instance.

DELETING A BACKUP SET

Consider the following before deleting a backup set:

- You cannot delete a default backup set.
 - Schedules associated with the backup set are also automatically deleted.
1. From the CommCell Browser, right-click the user-defined backup set 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 backup set.
- Click **No** to cancel the deletion and retain the backup set, or click **Yes** to continue the deletion.

DELETING A SUBCLIENT

Consider the following before deleting a subclient:

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

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

[Back to Top](#)

Advanced – VMware Backup

TABLE OF CONTENTS

Choosing Backup Types

Full Backup

Incremental Backups

Differential Backups

Synthetic Full Backups

- Running Synthetic Full Backups
- Running an Incremental Backup Before or After a Synthetic Full Backup
- Verifying Synthetic Full Backups
- Ignoring Errors in Synthetic Full Backups
- Accelerated Synthetic Full Backups (DASH Full)

Scheduling A Backup

Verifying the Virtual Machine Protection

Managing Jobs

- Restarting Jobs
- Controlling Jobs

Additional Options

CHOOSING BACKUP TYPES

BACKUP LEVEL	RESTORE LEVEL	NOTES
Disk-Level	Disk-Level (as virtual machine) Disk Level (as VMDK file)	
Disk-Level with Enable Granular Recovery enabled	Disk-Level (as virtual machine) Disk-Level (as VMDK file) File-Level	For File-level restores, the following apply: <ul style="list-style-type: none"> • The MediaAgent must be a Windows-based computer when restoring from Windows servers. • The Index Cache must be on an NTFS disk. • To perform Linux file-level restores, the LinuxMetadataSupport registry key must first be configured on the computer on which the Virtual Server iDataAgent is installed prior to performing the backup. You can restore files only from ext2, ext3 and NTFS file systems.
Volume-Level	Volume-Level (as physical volume) Volume-Level (as VHD) Volume-Level (as VMDK)	Supported only with volumes formatted with the NTFS file system.
Volume-Level with Enable Granular Recovery enabled	Volume-Level (as physical volume) Volume-Level (as VHD) Volume-Level (as VMDK) File-Level	For File-level restores, the following apply: <ul style="list-style-type: none"> • The MediaAgent must be a Windows-based computer when restoring from Windows servers. • The Index Cache must be on an NTFS disk.
File-Level	File-Level	Supported only with volumes formatted with the NTFS file system.

FULL BACKUP

Full backups provide the most comprehensive protection of data. However, full backups also consume the most amount of time and resources. To streamline the backup process, several additional backup types are available. The sections below describe the additional backup types that are available.

1. From the CommCell Console, navigate to **<Client> | Virtual Server**.
2. Right-click the **BackupSet** and click **All Tasks | Backup All Subclients**.
Click **Yes**.
3. Select **Full** as backup type and **Immediate** to run the job immediately.
Click **OK**.
4. You can track the progress of the job from the **Job Controller or Event Viewer** window.
5. Once job is complete, view the details of job from the **Job History**.

RELATED TOPICS

Scheduling

Provides comprehensive information on scheduling jobs.

Job Management

Provides comprehensive information on managing jobs.

Right-click the client computer, click **View** and then click **View Job History**.


6. Click **OK**.
7. You can view the following details about the job by right-clicking the job:
 - o Items that succeeded during the job
 - o Items that failed during the job
 - o Details of the job
 - o Media associated with the job
 - o Events of the job
 - o Log files of the job.

INCREMENTAL BACKUPS

An incremental backup contains only data that is new or has changed since the last backup, regardless of the type. On average, incremental backups consume far less media and place less of a burden on resources than full backups.

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

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




1. From the CommCell Console, navigate to **<Client> | Virtual Server | defaultBackupSet**.
2. Right click the subclient and click **Backup**.
3. From the **Backup Options** dialog box, select from the following options:
 Backup Type:
 - o Select **Incremental**.
 Backup Schedule:
 - o Select **Run Immediately**
4. Click **OK**.
5. You can track the progress of the backup job from the **Job Controller** window.
6. When the backup has completed, the Job Controller displays **Completed**.

DIFFERENTIAL BACKUPS

A differential backup contains only the data that is new or has changed since the last full backup. Like incremental backups, differential backups, on average, consume less media and place less of a burden on resources than full backups. Differential backups are cumulative. This means that each differential backup contains all changes accumulated since the last full backup. Each successive differential backup contains all the changes from the previous differential backup.

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

Backup #1 is a full backup and therefore writes all the data to the backup media. Backups #2 through #n-1 are differential backups and only back up those files that changed since the time of the last full backup. For example, files A, B, and E changed after the full backup and were therefore backed up in Backup #2 as well as all subsequent differential backups. File C changed sometime after Backup #2 and was consequently backed up in Backup #3 and all subsequent differential backups. File F did not change; hence, it was not backed up in any of the differential backups, but it was included in both full backups, which, by definition, back up everything.



1. From the CommCell Console, navigate to **<Client> | Virtual Server | defaultBackupSet**.
2. Right click the subclient and click **Backup**.
3. From the **Backup Options** dialog box, select from the following options:
 Backup Type:

- o Select **Differential**.

Backup Schedule:

- o Select **Run Immediately**
4. Click **OK**.
 5. You can track the progress of the backup job from the **Job Controller** window.
 6. When the backup has completed, the Job Controller displays **Completed**.

SYNTHETIC FULL BACKUPS

A synthetic full backup is a synthesized backup, created from the most recent full backup and subsequent incremental and/or differential backups. The resulting synthetic full backup is identical to a full backup for the subclient.

Unlike full, incremental, and differential backups, a synthetic full backup does not actually transfer data from a client computer to the backup media. Therefore, they do not use any resources on the client computer.

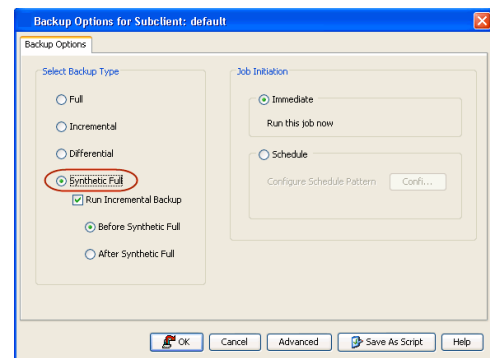
Synthetic full backups are media-based; they read backup data from one media and simultaneously write (restore) the result to the new active media within the same media group. As a result, synthetic full backups require that at least two media drives for the same storage policy be available at the time the job is started. Synthetic full backups cannot be performed on subclients where the storage policy is associated with a stand-alone drive.

RUNNING SYNTHETIC FULL BACKUPS

Synthetic full backups can either be started manually or scheduled to occur at regular intervals. Do not concurrently run more than one synthetic full backup, especially synthetic full backups with multiple streams.

1. From the CommCell Console, navigate to **<Client> | Virtual Server | defaultBackupSet**.
2. Right click the **<Subclient>** and click **Backup**.
3. From the **Backup Options** window, select **Synthetic Full** for **Backup Type**.
4. Select **Immediate** for **Job Initiation**.
5. Click **OK**.

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



RUNNING AN INCREMENTAL BACKUP BEFORE OR AFTER A SYNTHETIC FULL BACKUP

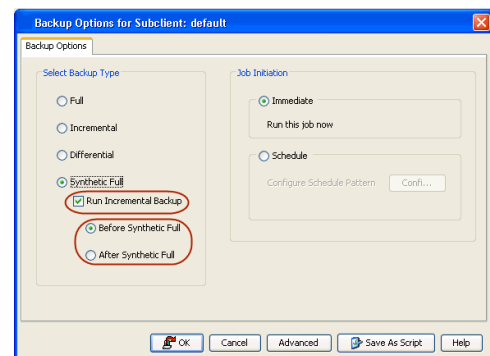
An incremental backup can be run either before or after a synthetic full backup.

- Running an incremental backup before the synthetic full ensures that any new or recently changed data is included in the synthetic full.
- Running an incremental backup after the synthetic full ensures that any new or recently changed data since the backup that occurred prior to the synthetic full, but was not included in the synthetic full, is backed up by the incremental.

The sections below describe the steps to run an incremental backup before or after a synthetic full backup:

1. From the CommCell Console, navigate to **<Client> | Virtual Server | defaultBackupSet**.
2. Right click the **<Subclient>** and click **Backup**.
3. From the **Backup Options** window, select **Synthetic Full** for **Backup Type**.
4. Select **Run Incremental Backup** and then select **Before Synthetic Full** or **After Synthetic Full** options.
5. Select **Immediate** for **Job Initiation**.
6. Click **OK**.

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



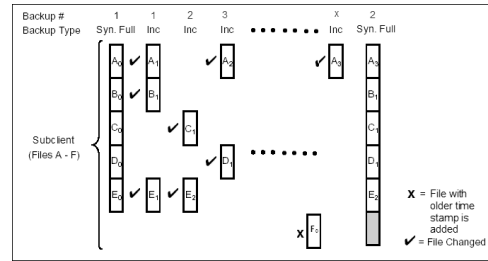
VERIFYING SYNTHETIC FULL BACKUPS

This option ensures that files that have not backed up by incremental jobs since the last full backup are backed up by the next incremental.


For example: If you perform the conventional full backup only once for a given subclient and then perform incremental (or differential) backups and periodic synthetic full backups. After the conventional full backup, a file is added to the subclient and the file has a time stamp older than the time of full backup. In such scenario, the file will never get backed up. Eventually, such files may be pruned and no backup of such files will be available.

The illustration on the right explains the necessity of verifying synthetic full backups.

In order to overcome this problem, it is recommended to verify the synthetic full backups job periodically as follows:




1. From the CommCell Console, navigate to **<Client> | Virtual Server | defaultBackupSet**.
2. Right click the **<Subclient>** and click **Backup**.
3. From the **Backup Options** dialog box, select **Synthetic Full**.
4. Click **Advanced**.
5. In the **Advanced Backup Options** dialog box, click the **Data** tab.
6. Select **Verify Synthetic Full**. Click **OK** to continue.
7. Click **OK**.



IGNORING ERRORS IN SYNTHETIC FULL BACKUPS

Errors encountered in synthetic full backups can be ignored, even if backup media is partially corrupted. Errors such as inability to read data (bad tape, files missing on disk media, etc.) will be ignored when this option is enabled.

1. From the CommCell Console, navigate to **<Client> | Virtual Server | defaultBackupSet**.
2. Right click the subclient and click **Backup**.
3. From the **Backup Options** dialog box, select **Synthetic Full**.
4. Click **Advanced**.
5. In the **Advanced Backup Options** dialog box, click the **Data** tab.
6. Select **Verify Synthetic Full** and then select **Ignore Error in Synthetic Full**.
7. Click **OK** to close the **Advanced Backup Options** dialog box.
8. Click **OK**.



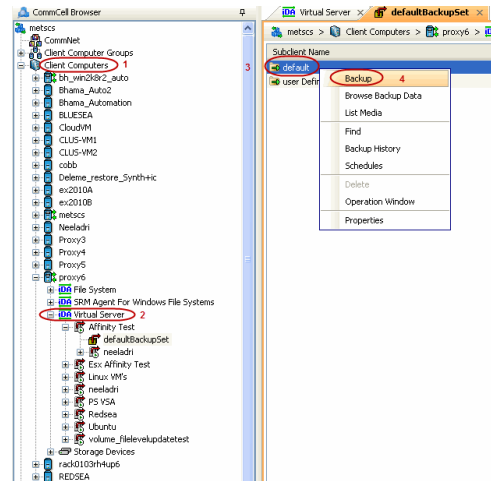
ACCELERATED SYNTHETIC FULL BACKUPS (DASH FULL)

When a storage policy copy is deduplicated, synthetic full backups can be created in an accelerated mode to significantly reduce the copy duration. This is done by identifying and transferring the data signatures (instead of the data itself) to the target wherever possible.

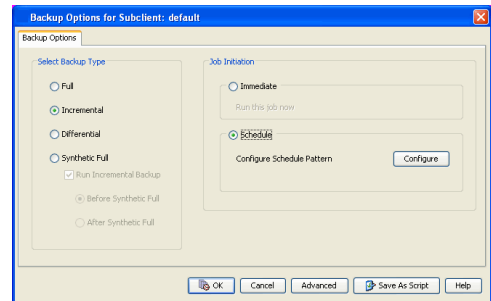
Learn more...

SCHEDULING A BACKUP

1.
 - From the CommCell Console, navigate to **Client Computers | Virtual Server**
 - Right-click the **Subclient** and click **Backup**.

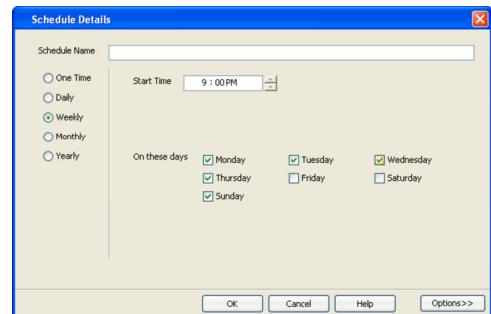


2.
 - Click **Schedule** to schedule the backup for a specific time.
 - Click **Configure** to set the schedule for the backup job. The Schedule Details dialog displays.



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

The backup job will execute as per the schedule.



RELATED TOPICS

Scheduling: Provides comprehensive information on scheduling.

Schedule Policy: Provides comprehensive information on creating schedule policies.

VERIFYING THE VIRTUAL MACHINE PROTECTION COVERAGE

You can get information about the number of virtual machines protected at any point in time. This information is useful to verify that all the virtual machines are getting protected as specified in the backup schedule. You can run a SQL query on the CommServe database to verify the status of backups performed for each virtual machine in the last *<n>* days. When you schedule backups for a subclient, you should run this query periodically to check whether all the virtual machines are getting backed up according to the schedule.

Use the following steps to run the query:

1. Log on to the CommServe computer. Ensure that the user account used for logging in has administrative privileges.
2. Open the Microsoft SQL Server Management Studio and run the following query on the CommServe database:

```
select * from VMPProtectionCoverage (<n>, '<backup type>')
```

where *<n>* is the number of days and *<backup type>* is the type of backup.

For example:

```
select * from VMPProtectionCoverage (30, 'INCR')
```

This query provides the status of incremental backups performed for each virtual machine protected in the last 30 days. The query is executed for all virtual machines protected by any Virtual Server iDataAgent within the CommCell.

To verify the status of full and incremental backups performed for each virtual machine in the last <n> days, run the query as follows:

VMProtectionCoverage (<n>, ")

The VMProtectionCoverage query does not provide information about synthetic full backups.

The results of the query contains the status of the backups performed for each virtual machine. The status is one of the following:

STATUS	DESCRIPTION
Currently Protected	All backups for the virtual machine were successful in the last <n> days.
Discovered, Not Protected	The virtual machine was discovered but never backed up.
Manually Excluded	The virtual machine was excluded from the subclient and was not backed up in last <n> days.
Not Protected in the time range	No backup was performed for the virtual machine in the last <n> days, but a backup was performed in the past.

MANAGING JOBS

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

RESTARTING JOBS

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

It is also possible to override the default CommServe configuration for individual jobs by configuring retry settings when initiating the job. This configuration, however, will apply only to the specific job.

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

CONFIGURE JOB RESTARTABILITY AT THE COMMSERVE LEVEL

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

CONFIGURE JOB RESTARTABILITY FOR AN INDIVIDUAL JOB

1. From the **Backup Options** dialog box, click **Advanced**, then select the **Job Retry** tab and specify the following as desired:
 - o **Total Running Time** - The maximum elapsed time, in hours and minutes, from the time that the job is created.
 - o **Number of Retries** - The number of times that Job Manager will attempt to restart the job.
 - o **Kill Running Jobs When Total Running Time Expires** - Option to kill the job when the specified Total Running Time has elapsed, even if its state is "Running".
2. Click **OK**.

CONTROLLING JOBS

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

SUSPEND	Temporarily stops a job. A suspended job is not terminated; it can be restarted at a later time. Only preemptible jobs can be suspended.
RESUME	Resumes a job and returns the status to Waiting, Pending, Queued, or Running depending on the availability of resources or the state of the operation windows and activity control settings. Backup jobs for this Agent are resumed from the point-of-failure.
KILL	Terminates a job.

These controls can be applied to:

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

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

- When killing a job:

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

- When suspending a job:

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

- When resuming a job:

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

ADDITIONAL OPTIONS

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

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

OPTION	DESCRIPTION	RELATED TOPICS
Catalog	<p>The Catalog section helps you to select the index cache sharing and granular restartability options for the job.</p> <ol style="list-style-type: none"> Right-click the Subclient and select Backup. From the Backup Options for Subclient window, click Advanced. From the Data tab, select the Use shared profile if present with transaction logging option. This option helps you to use the shared index cache profile of the MediaAgent (if configured) to save a copy of the job's index cache for index cache sharing with transaction logging feature for granular job restartability. Shared profile with transaction logging provides job restartability in failover situations. Select the Use shared profile if present without transaction logging option. This option helps you to use the shared index cache profile of the MediaAgent (if configured) to save a copy of the index cache for index cache sharing without transaction logging feature for granular job restartability. This is not a recommended configuration. Select the Use Transaction Logging option. This option helps you to use transaction logging feature for granular job restartability (without using shared profile for index cache sharing). Select None for not using shared profile and transaction logging. Select the Enable Granular Recovery. This is to perform Granular Recovery. It allows the recovery of Individual files and folders. Click OK. 	Refer Index Cache Server for more information.
Startup Options	<p>The Job Manager will use the startup priority setting when allocating the required resources. This is useful if you have jobs that are very important and must complete, or jobs that can be moved to a lower priority.</p> <ol style="list-style-type: none"> Right-click the Subclient and select Backup. From the Backup Options for Subclient window, click Advanced. From Startup tab, select the Default Priority option. Select Change Priority option to change the priority for a job, between 0 (highest priority) and 999 (lowest priority). Select Start up in Suspended State option to specify the selected job to start in the Job Controller in a suspended state. Click OK. 	Refer Job Priority and Priority Precedence for more information.
Job Retry Options	<p>The Job Retry option helps in configuring the retry behavior of the backup jobs. You can specify the maximum elapsed time before a job can be restarted or killed and the maximum number of restart attempts.</p> <ol style="list-style-type: none"> Right-click the Subclient and select Backup. From the Backup Options for Subclient window, click Advanced. From the Job Retry tab, select the Enable Total Running Time check box and specify the maximum elapsed time in hours and minutes. Select the Number Of Retries check box and specify the number of retries in the Number Of Retries box. Based on this number, the Job Manager attempts to restart the job. 	Refer Job Management for more information.

	<ol style="list-style-type: none"> Select the Kill Running Jobs When Total Running Time Expires check box. The job will be killed when the job is still in running status even after reaching the maximum elapsed time. Click OK. 	
Start New Media	<p>The Start New Media option helps in starting the backup/archive operation on a new media. This media management feature provides a degree of control over where the data physically resides.</p> <ol style="list-style-type: none"> Right-click the Subclient and select Backup. From the Backup Options for Subclient window, click Advanced. From the Media tab, select the Start New Media check box. Click OK. 	<p>Refer the Creating an Exportable Media Set section of the Export Media for more information.</p> <p>Another form of the Start New Media option is available from the library properties. Refer the Start New Media section of the Library Properties for more information.</p>
Mark Media Full	<p>The Mark Media Full on Success option marks the Media as Full, 2 minutes after the successful completion of the backup/archive. This feature prevents any other data being written to the same media.</p> <ol style="list-style-type: none"> Right-click the Subclient and select Backup. From the Backup Options for Subclient window, click Advanced. From the Media tab, select the Mark Media Full on Success check box. Click OK. 	<p>Refer the Create an Exportable Media Set section of the Export Media documentation for more information.</p> <p>Refer Export Media for more information.</p>
Allow other Schedules to use Media Set	<p>The Allow Other Schedules to use Media Set option allows jobs that are part of the schedule or schedule policy and using the specific storage policy to start a new media. It also prevents other jobs from writing to the same set of media.</p> <ol style="list-style-type: none"> Right-click the Subclient and select Backup. From the Backup Options for Subclient window, click Advanced. From the Media tab, select the Allow Other Schedules to Use Media Set check box. Click OK. 	<p>Refer the Creating an Exportable Media Set section of the Export Media for more information.</p>
Data Path Options	<p>Data Protection operations use a specific data path (Library, MediaAgent, Drive Pool, and Drive) to perform the backup operations as configured in the CommCell. By default, the system automatically identifies the data path for the backup operations. The following are some of the important situations where you may need to change the data path:</p> <ul style="list-style-type: none"> When the preferred MediaAgent, library or drive pool is not available during the backup When you need to use a different MediaAgent, library or drive pool for load balancing purposes. When you need to direct that backup to a particular location for compliance purposes. <ol style="list-style-type: none"> Right-click the Subclient and select Backup. From the Backup Options for Subclient window, click Advanced. From the Data Path tab, select Use MediaAgent and Use Library or Use Drive Pool and Use Drive from the respective drop-down combo boxes. Click OK. 	<p>Refer Change Data Path for more information.</p>
Vault Tracker	<p>The VaultTracker feature provides the facility to manage media that are removed from a library and stored in offsite locations. The VaultTracker function provides the following capabilities in your day-to-day operations:</p> <ul style="list-style-type: none"> Identifies media that must be sent off-site for storage or brought back from off site locations. Automatically moves the media in sequence in the library and provides a pick-up list for the operators. Identifies and tracks the media during transit. Records and tracks the movement of media not used by all MediaAgents or Foreign Media. <ol style="list-style-type: none"> Right-click the Subclient and select Backup. From the Backup Options for Subclient window, click Advanced. From the VaultTracking tab, select the necessary options. Click Ok. 	<p>Refer to the following documentation for a comprehensive overview prior to using this feature:</p> <ul style="list-style-type: none"> VaultTracker if a standard CommServe license is available. VaultTracker Enterprise if the VaultTracker Enterprise license is available.
Alerts	<p>The Alert option is used for setting up the criteria to raise notifications/alerts for job statuses such as failure, success, or any other conditions triggered by the backup job. Adding alerts helps the user or the user group to get the notification automatically about the status of the backup job.</p>	<p>Refer Alerts for more information.</p>

	<ol style="list-style-type: none"> 1. Right-click the Subclient and select Backup. 2. On the Backup Options for Subclient window, click Advanced. 3. On the Alert tab, click Add Alert and it opens Add Alert Wizard window. 4. From the Add Alert Wizard window, select the Threshold and Notification Criteria. Click Next. 5. Select the Notification Types. Click Next. 6. Select the Users and User Groups. Click Next. 7. Click Finish. 	
Command Line Backups	Command Line Interface enables you to perform backups from the command line. The commands can be executed from the command line or can be integrated into your own scripts or scheduling programs. In addition, you can also generate scripts for specific operations from the CommCell Console using the Save As Script option. These scripts can later be executed using the commands from the command line interface.	Refer Command Line Interface for more information.
CommCell Readiness Report	The CommCell Readiness Report provides you with vital information about the condition of your CommCell.	Refer CommCell Readiness Report for more information.
Backup Job Summary Report	The Backup Job Summary Report provides the details of all the backup jobs of clients.	Refer Backup Job Summary Report for more information.

Back to Top

Advanced - VMware Restore

TABLE OF CONTENTS

Choosing Restore Types

Restoring Multiple Virtual Machines

Restoring Virtual Machines to Same Destination (In-Place Restore)

Restoring Virtual Machines to Different Destination (Out-of-place Restore)

- Different Destination Client
- Different Virtual Center
- Different ESX Server
- Different Data Store
- Different Resource Pool/vApp
- Different Name

Restoring Virtual Machines with Disk Provisioning

Restoring Virtual Machines with Transport Mode

Overwrite Virtual Machines with Same Name

Automatically Power On Restored Virtual Machines

Restoring Virtual Machine Disk (VMDK) Files

- Restoring VMDK Files
- Mounting VMDK Files Backed up from VStorage Environments
- Mounting VMDK Files as a Complete Volume

Restoring Volumes

Restoring Files and Folders

- Windows Clients
- Linux Clients

Finding and Restoring

- Prerequisites
- Finding a File or Folder
- Finding a File or Folder On Multiple Virtual Machines
- Finding A File or Folder on Specific Volume
- Finding a File Backed Up in Specific Time Range
- Quick Search Using Sampling
- Adding Search Results to a List
- Using Find From a Browse Window
- Restoring Files or Folders from the Search Results

Scheduling a Restore

Restoring Guest Files from Unix Media Agent

Creating Virtual Machines from Restored VMX Files

Browse Data

- Browsing Data From Before the Most Recent Full Backup

List Media

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

Restore by Jobs

Manage Restore Jobs

- Restart Jobs
- Resubmit Jobs
- Control Jobs

Additional Restore Options

CHOOSING RESTORE TYPES

When restoring VMware data, it is important to consider the backup-level that was originally performed. The following table illustrates the types of restores available for each backup level:

BACKUP LEVEL	RESTORE LEVEL	NOTES
Disk-Level	Disk-Level (as virtual machine) Disk Level (as VMDK file)	
Disk-Level with Enable Granular Recovery	Disk-Level (as virtual machine)	For File-level restores, the following apply:

RELATED TOPICS

Scheduling

Provides comprehensive information on scheduling jobs.

Job Management

Provides comprehensive information on managing jobs.

List Media

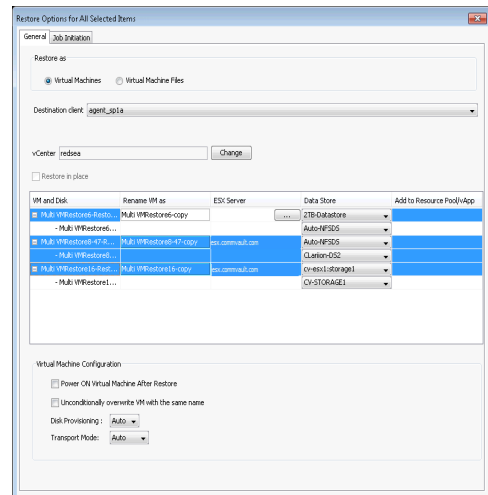
Provides comprehensive information on media management and prediction.

enabled	Disk-Level (as VMDK file) File-Level	<ul style="list-style-type: none"> The MediaAgent must be a Windows-based computer when restoring from Windows servers. The Index Cache must be on an NTFS disk. To perform Linux file-level restores, the LinuxMetadataSupport registry key must first be configured on the computer on which the Virtual Server iDataAgent is installed prior to performing the backup. <p>You can restore files only from ext2, ext3 and NTFS file systems.</p>
Volume-Level	Volume-Level (as physical volume) Volume-Level (as VHD) Volume-Level (as VMDK)	Supported only with volumes formatted with the NTFS file system.
Volume-Level with Enable Granular Recovery enabled	Volume-Level (as physical volume) Volume-Level (as VHD) Volume-Level (as VMDK) File-Level	For File-level restores, the following apply: <ul style="list-style-type: none"> The MediaAgent must be a Windows-based computer when restoring from Windows servers. The Index Cache must be on an NTFS disk.
File-Level	File-Level	Supported only with volumes formatted with the NTFS file system.

RESTORING MULTIPLE VIRTUAL MACHINES

To fasten and ease the restore operation, you can restore multiple virtual machines in a single job. Follow the steps given below, to restore multiple virtual machines:

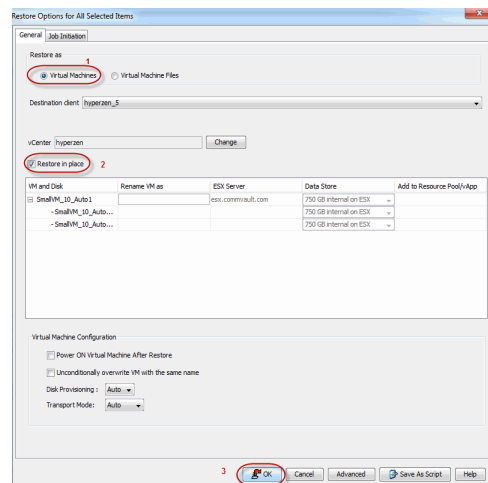
- From the CommCell Console, navigate to the <Client> | **Virtual Server**.
- Right-click the subclient that contains the data you want to restore, point to **Browse Backup Data**.
- In the **Browse Options** window, click **OK**.
- From the **Browse** window, select the virtual machine in the left pane. The configuration and VMDK files of the virtual machine are displayed in the right pane.
- Click **Recover All Selected** in the bottom of the window.
- From the **General** tab of the **Restore Options for All Selected Items** dialog box:
 - Click **Virtual Machines**.
 - All the selected virtual machines are displayed in the **VM and Disk** column.
 - Enter corresponding values for each virtual machine you want to restore.
 - To select multiple virtual machines hold down the SHIFT key and select the virtual machines you want to restore.
- Click **OK**. All the virtual machines that were selected are restored in a single job.



RESTORING VIRTUAL MACHINES TO SAME DESTINATION (IN-PLACE RESTORE)

When restoring data, the Virtual Server iDataAgent provides the facility to restore data to the same destination host, and placing all of the disks back to their original datastores. Follow the steps given below to perform an in-place restore:

- From the CommCell Console, navigate to the <Client> | **Virtual Server**.
- Right-click the subclient that contains the data you want to restore, point to **Browse Backup Data**.
- In the **Browse Options** window, click **OK**.
- From the **Browse** window, select the virtual machine(s) in the left pane. The configuration and VMDK files of the virtual machines are displayed in the right pane.
- Click **Recover All Selected** in the bottom of the window.
- From the **General** tab of the **Restore Options for All Selected Items** dialog box:
 - Click **Restore in place**.
- Click **OK**. When a virtual machine is restored, it will be restored to the same destination per the time of backup.



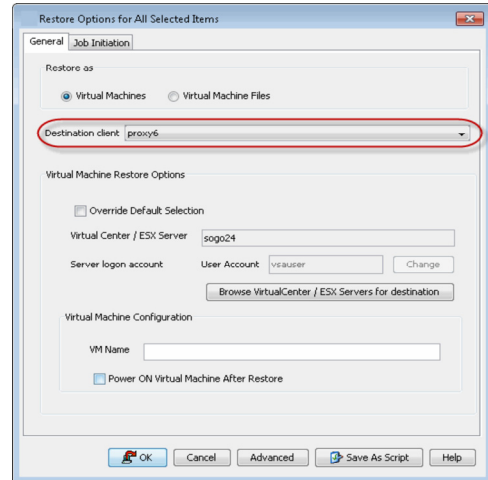
RESTORING VIRTUAL MACHINES TO DIFFERENT DESTINATION (OUT-OF-PLACE RESTORE)

Virtual machines can be restored to a different destination host from which the data originated; this is referred to as an out-of-place restore. The following sections provide the options available for an out-of-place restore.

RESTORING A VIRTUAL MACHINE TO A DIFFERENT DESTINATION CLIENT

By default, the virtual machine is associated to the client machine, as it was at the time of backup. However, an entire virtual machine can be restored to a different client, by following the steps given below:

1. From the CommCell Console, navigate to the <Client> | **Virtual Server**.
2. Right-click the subclient that contains the data you want to restore, point to **Browse Backup Data**.
3. In the **Browse Options** window, click **OK**.
4. From the **Browse** window, select the virtual machine in the left pane. The configuration and VMDK files of the virtual machine are displayed in the right pane.
5. Click **Recover All Selected** in the bottom of the window.
6. From the **General** tab of the **Restore Options for All Selected Items** dialog box:
 - o Click **Virtual Machines**.
 - o Choose a different **Destination Client** from the dropdown list.
7. Click **OK**.

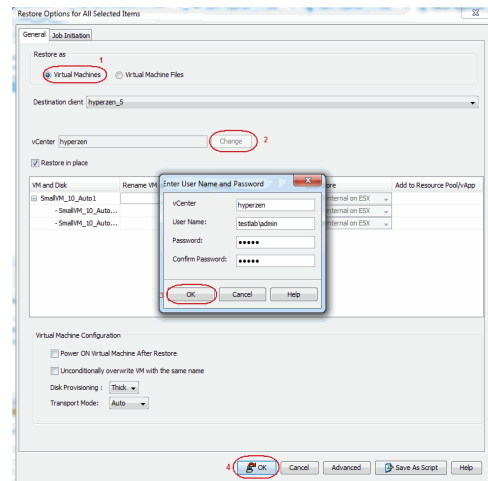


RESTORING TO A DIFFERENT VIRTUAL CENTER

An entire virtual machine can be restored to a datastore of the original Virtual Center or to a different one using the credentials of the Virtual Center. Specifying a different location provides the flexibility to distribute restored virtual machines to a location with greater space and resource availability, should the virtual machine's original location prove not optimal.

When restoring a virtual machine, a Disk level backup must have been used to back it up. Follow the steps given below to restore the virtual machine to a different vCenter:

1. From the CommCell Console, navigate to the <Client> | **Virtual Server**.
2. Right-click the subclient that contains the data you want to restore, point to **Browse Backup Data**.
3. In the **Browse Options** window, click **OK**.
4. From the **Browse** window, select the virtual machine in the left pane. The configuration and VMDK files of the virtual machine are displayed in the right pane.
5. Click **Recover All Selected** in the bottom of the window.
6. From the **General** tab of the **Restore Options for All Selected Items** dialog box:
 - o Click **Virtual Machines**.
 - o Choose a different **Virtual Center** by clicking on the **Change** button.
 - o Enter the name of the Virtual Center, that you want your virtual machine to be restored into.
 - o Enter the User credentials.
 - o Click **OK**.
7. Click **OK**.



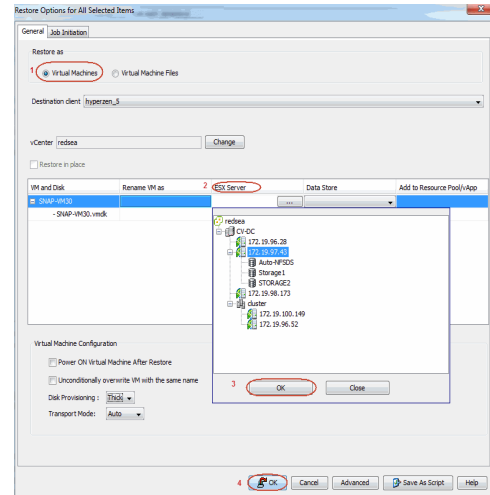
RESTORING TO A DIFFERENT ESX SERVER

By default, the virtual machine is restored to the ESX server, as it was at the time of backup. However, an entire virtual machine can be restored to a different ESX server, by following the steps given below. Specifying a different location provides the flexibility to distribute restored virtual machines to a location with greater space and resource availability should the virtual machine's original location prove not optimal.

Before restoring a virtual machine ensure the following:

- A Disk Level backup is performed for the virtual machine.
 - The ESX server supports the hardware version of the virtual machine. If the hardware version of the virtual machine is not supported by the ESX server, the restore job will complete successfully but the virtual machine will be inaccessible.
1. From the CommCell Console, navigate to the <Client> | **Virtual Server**.
 2. Right-click the subclient that contains the data you want to restore, point to **Browse Backup Data**.
 3. In the **Browse Options** window, click **OK**.

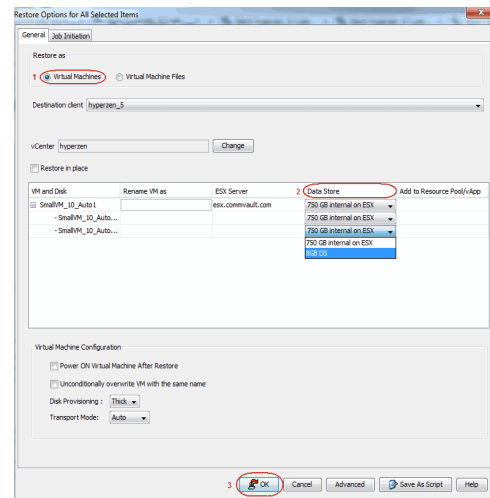
- From the **Browse** window, select the virtual machine in the left pane. The configuration and VMDK files of the virtual machine are displayed in the right pane.
- Click **Recover All Selected** in the bottom of the window.
- From the **General** tab of the **Restore Options for All Selected Items** dialog box:
 - Click **Virtual Machines**.
 - Choose a different **ESX Server** by clicking on the Browse button.
 - Select a different ESX Server from the **Browse** window.
 - Click **OK**.
- Click **OK**.



RESTORING TO A DIFFERENT DATA STORE

You can specify a different datastore on the destination host by following the steps given below. When restoring a virtual machine, a Disk Level backup must have been used to back it up. After restore, the virtual machine is automatically associated to the new data store selected.

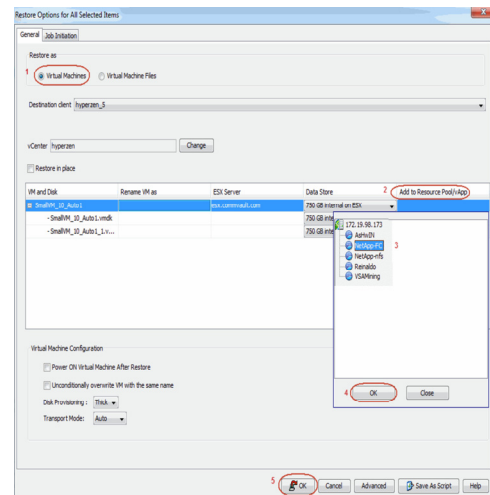
- From the CommCell Console, navigate to the **<Client> | Virtual Server**.
- Right-click the subclient that contains the data you want to restore, point to **Browse Backup Data**.
- In the **Browse Options** window, click **OK**.
- From the **Browse** window, select the virtual machine(s) in the left pane. The configuration and VMDK files of the virtual machines are displayed in the right pane.
- Click **Recover All Selected** in the bottom of the window.
- From the **General** tab of the **Restore Options for All Selected Items** dialog box:
 - Click **Virtual Machines**.
 - Choose a different **Data Store** from the dropdown list.
 - If a virtual machine has multiple disks a destination datastore can be selected for each disk.
- Click **OK**.



RESTORING TO A DIFFERENT RESOURCE POOL/VAPP

By default, when a virtual machine is restored it will be restored to its original resource pool. You can select a specific resource pool/vApp on the host that the virtual machine can be restored into. vApp is a customized form of a resource pool. Follow the steps given below to add/change a virtual machine to a Resource Pool/vApp:

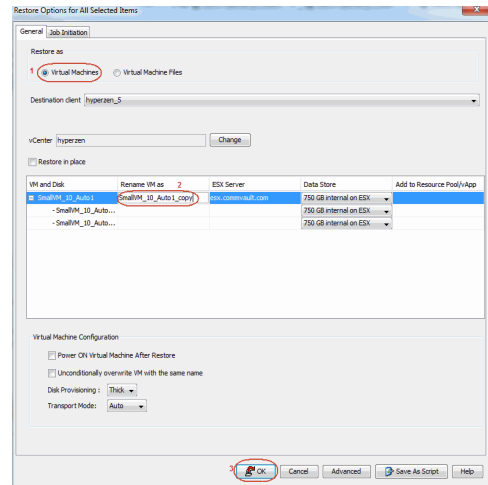
- From the CommCell Console, navigate to the **<Client> | Virtual Server**.
- Right-click the subclient that contains the data you want to restore, point to **Browse Backup Data**.
- In the **Browse Options** window, click **OK**.
- From the **Browse** window, select the virtual machine in the left pane. The configuration and VMDK files of the virtual machine are displayed in the right pane.
- Click **Recover All Selected** in the bottom of the window.
- From the **General** tab of the **Restore Options for All Selected Items** dialog box:
 - Click **Virtual Machines**.
 - Add the virtual machine to a **Resource Pool/vApp** by clicking on the Browse button.
 - Select a different Resource Pool/vApp from the **Browse** window.
 - Click **OK**.
- Click **OK**.



RESTORING VIRTUAL MACHINES WITH A DIFFERENT NAME

By default, the Virtual Machine is restored with the original name, exactly as it was at the time of backup. This can be changed to a more appropriate name by following the steps given below:

1. From the CommCell Console, navigate to the <Client> | **Virtual Server**.
2. Right-click the subclient that contains the data you want to restore, point to **Browse Backup Data**.
3. In the **Browse Options** window, click **OK**.
4. From the **Browse** window, select the virtual machine in the left pane. The configuration and VMDK files of the virtual machine are displayed in the right pane.
5. Click **Recover All Selected** in the bottom of the window.
6. From the **General** tab of the **Restore Options for All Selected Items** dialog box:
 - o Click **Virtual Machines**.
 - o Type in a different name to the virtual machine by double clicking the **Restore as** text box.
7. Click **OK**. When the virtual machine is restored, it is restored with the new name.



RESTORING VIRTUAL MACHINES WITH DISK PROVISIONING

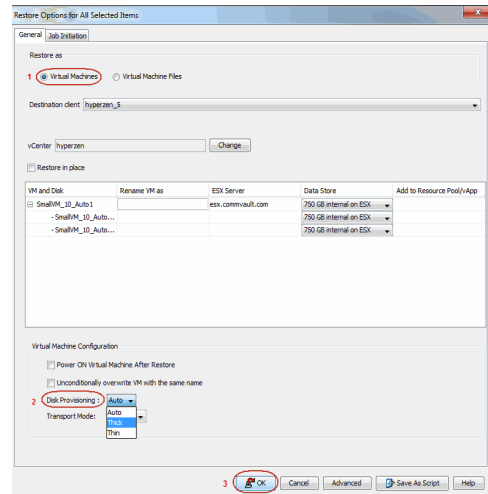
Ensure that Service Pack 2 has been installed on the CommServe and Client computers prior to using this feature.

By default the disk is restored with its original provisioning type, as it was at the time of backup. The disk provisioning types available are:

- **Auto**: same as the source virtual machine, chosen during backup.
- **Thick**: Thick provisioning allocates and commits full size to the data store.
- **Thin**: Thin provisioning allocates full size and commits only as much storage space as the disk needs for its initial operation.

Follow the steps given below to change the provisioning type on restore:

1. From the CommCell Console, navigate to the <Client> | **Virtual Server**.
2. Right-click the subclient that contains the data you want to restore, point to **Browse Backup Data**.
3. In the **Browse Options** window, click **OK**.
4. From the **Browse** window, select the virtual machine in the left pane. The configuration and VMDK files of the virtual machine are displayed in the right pane.
5. Click **Recover All Selected** in the bottom of the window.
6. From the **General** tab of the **Restore Options for All Selected Items** dialog box:
 - o Click **Virtual Machines**.
 - o Choose an appropriate **Disk provisioning** option from the drop down list.
7. Click **OK**. When the virtual machine is restored, it is restored with the disk provisioning type selected.



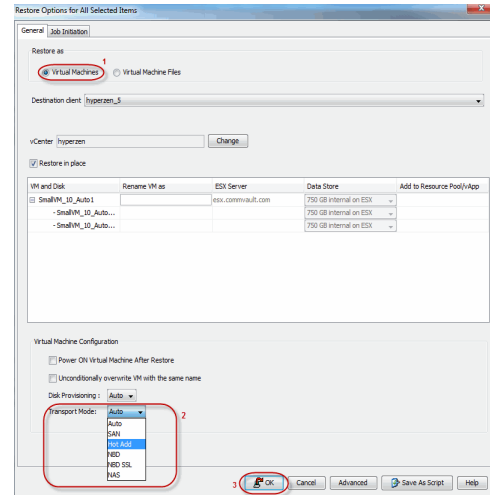
RESTORING VIRTUAL MACHINES WITH TRANSPORT MODE

Ensure that Service Pack 2 has been installed on the CommServe and Client computers prior to using this feature.

By default, the software automatically defaults to the **auto** transport mode. Choosing Auto, software will automatically select the best suitable transport mode based on the setup. Follow the steps given below to restore virtual machines with a different Transport Mode option:

1. From the CommCell Console, navigate to the <Client> | **Virtual Server**.
2. Right-click the subclient that contains the data you want to restore, point to **Browse Backup Data**.
3. In the **Browse Options** window, click **OK**.
4. From the **Browse** window, select the virtual machine in the left pane. The configuration and VMDK files of the virtual machine are displayed in the right pane.
5. Click **Recover All Selected** in the bottom of the window.
6. From the **General** tab of the **Restore Options for All Selected Items** dialog box:
 - o Click **Virtual Machines**.
 - o Choose an appropriate **Transport Mode** option from drop down list.

- Click **OK**. When the virtual machine is restored, it is restored with the transport mode you selected.



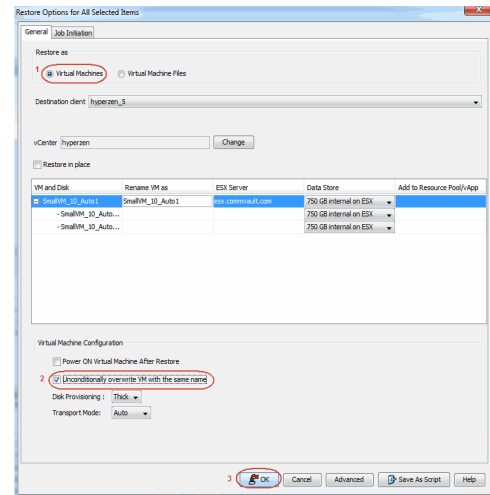
OVERWRITE VIRTUAL MACHINES WITH SAME NAME

Ensure that Service Pack 2 has been installed on the CommServe and Client computers prior to using this feature.

By default, this option is deselected. When Virtual machines with the same name exist on the destination host, you can choose the overwrite an existing virtual machine option, to free up space on the datastore. By doing this, the existing virtual machines are deleted and the newly restored virtual machines replace them.

Follow the steps given below to overwrite virtual machines with same name:

- From the CommCell Console, navigate to the **<Client> | Virtual Server**.
- Right-click the subclient that contains the data you want to restore, point to **Browse Backup Data**.
- In the **Browse Options** window, click **OK**.
- From the **Browse** window, select the virtual machine in the left pane. The configuration and VMDK files of the virtual machine are displayed in the right pane.
- Click **Recover All Selected** in the bottom of the window.
- From the **General** tab of the **Restore Options for All Selected Items** dialog box:
 - Click **Virtual Machines**.
 - Select **Unconditionally overwrite VM with same name** option to overwrite an existing virtual machine name.
- Click **OK**.

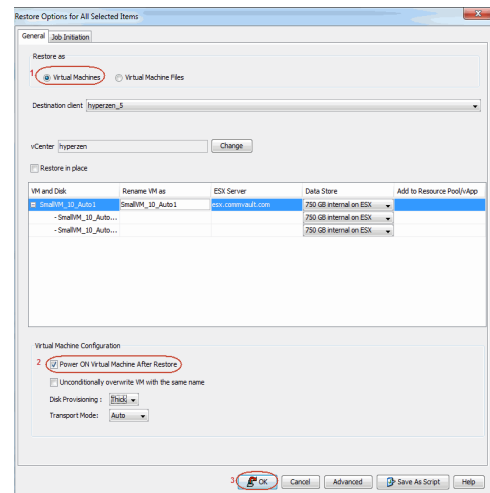


AUTOMATICALLY POWER ON RESTORED VIRTUAL MACHINES

Ensure that Service Pack 2 has been installed on the CommServe and Client computers prior to using this feature.

Once a virtual machine has been restored, it can be automatically powered on by following the steps given below:

- From the CommCell Console, navigate to the **<Client> | Virtual Server**.
- Right-click the subclient that contains the data you want to restore, point to **Browse Backup Data**.
- In the **Browse Options** window, click **OK**.
- From the **Browse** window, select the virtual machine in the left pane. The configuration and VMDK files of the virtual machine are displayed in the right pane.
- Click **Recover All Selected** in the bottom of the window.
- From the **General** tab of the **Restore Options for All Selected Items** dialog box:
 - Click **Virtual Machines**.
 - Select **Power On Virtual Machine After Restore** to power on the virtual machine immediately following restore completion.
- Click **OK**.



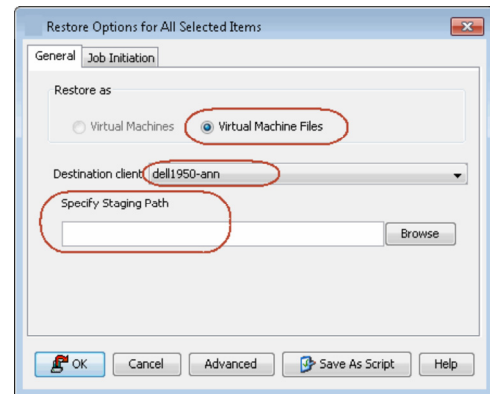
RESTORING VIRTUAL MACHINE DISK (VMDK) FILES

RESTORING VMDK FILES

You can restore single or multiple VMDK files from a virtual machine.

1. From the CommCell Console, navigate to the **Client Computers | <Client> | Virtual Server | <Instance> | <BackupSet>**.
2. Right-click the subclient that contains the VMDK file you want to restore and select **Browse Backup Data**.
3. In the **Browse Options** window, click **OK**.
4. The **Browse** window displays list of backed up virtual machines in the left pane. The configuration and VMDK files of the virtual machine are displayed in the right pane.
5. Select the required VMDK file from the right pane,
6. Click **Recover All Selected** in the bottom of the window.
7. From the **General** tab, select a **Destination client** (e.g. a proxy computer).
8. Click **Browse** and specify the **Staging Path**.
9. Click **OK**.

The VMDK file will be restored to the location specified in the staging path.



MOUNTING VMDK FILES BACKED UP USING VSTORAGE ENVIRONMENTS

To mount VMDK files if a backup was done using the vStorage method, execute the following at the command line:

```
<software install path>\Base\VMWARE\VDDK\bin>vmware-mount.exe <DriveLetter> <Path to disk>
```

For example: C:\Program Files\Calyпсо\Base\VMWARE\VDDK\bin>vmware-mount.exe Z: "G:\restoreasVMDK\W2K8_X64_1.vmdk"

To unmount VMDK files if backup was done using the vStorage method, execute the following at the command line:

```
<software install path>\Base\VMWARE\VDDK\bin>vmware-mount.exe <DriveLetter> /d
```

For example: C:\Program Files\Calyпсо\Base\VMWARE\VDDK\bin>vmware-mount.exe Z: /d

When restoring the VMDK files, do not include the configuration files like vmsd. It is recommended that you select and restore the .VMDK files to mount as Volume.

MOUNTING VMDK FILES AS A COMPLETE VOLUME

In some cases, it may be desirable to mount virtual machine VMDK files directly using the VMware Disk Mount utility following a restore operation. Doing so mounts all the VMDK file data as a complete volume, providing easy access to specific data residing on the original VMDK file.

To mount VMDK files from disk-level backups execute the following at the command line:

```
vmware-mount.exe <DriveLetter> <Path to disk>
```

For example: C:\Program Files\VMware\VMware DiskMount Utility>vmware-mount.exe Q: "G:\DiskFiles\scsi0-0-0-EXCH-VM2.vmdk"

To unmount VMDK files from disk-level backups execute the following at the command line:

```
VMware-mount.exe /d <DriveLetter>
```

For example: C:\Program Files\VMware\VMware DiskMount Utility>vmware-mount.exe /d Q: "G:\DiskFiles\scsi0-0-0-EXCH-VM2.vmdk"

To mount VMDK files from volume-level backups execute the following at the command line:

```
vmware-mount.exe <DriveLetter> <Path to disk> E
```

For example: C:\Program Files\VMware\VMware DiskMount Utility>vmware-mount.exe Q: "G:\DiskFiles\C.vmdk"

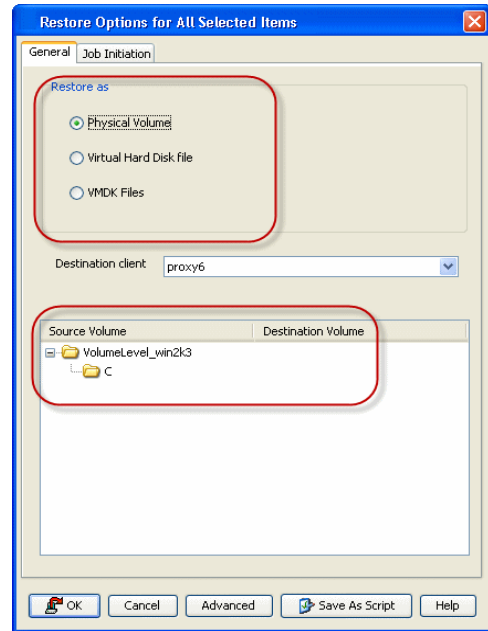
RESTORING VOLUMES

Volumes can be restored to a physical volume, a VMDK file, or a VHD file format. The volumes that are supported are for Windows-based VMware Virtual Servers and must have been formatted with the NTFS file system.

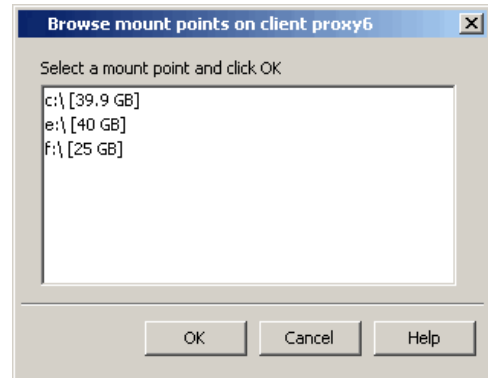
When restoring a volume, a Volume Level backup must have been used to back it up.

1. From the CommCell Console, perform a **Browse and Restore** operation.
2. In the **Browse Options** window, click **Entire Volume** and click **OK**.

3. In the **Browse** window, click a volume in the right pane and click **Recover All Selected**.
4. From the **General** tab, click either:
 - o **Physical Volume**, which automatically mounts the volume on the destination VMware virtual machine.
 - o **Virtual Hard Disk file**, which provides the facility to manually mount the volume on another virtual machine of any virtualization vendor type (such as Hyper-V) to mine specific data.
 - o **VMDK Files**, which provides the facility to manually mount the volume on another virtual machine to mine specific data.
5. Choose the **Destination Client** (e.g., a proxy computer).



6. If restoring a physical volume and if you want to restore the volume on a different volume from the source, double-click **Destination Volume** and click ... to browse and select a volume.
7. Click **OK**.



RESTORING FILES AND FOLDERS

WINDOWS CLIENTS

To restore files and folders to a Windows client, the client must have one of the following components installed:

- Windows File System iDataAgent
- Restore Only Agent for Windows File System.


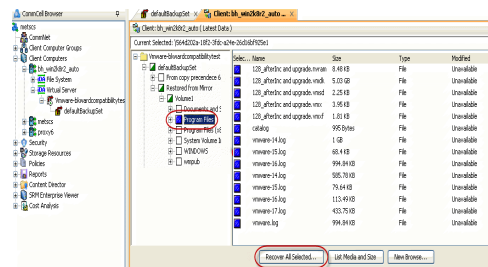
Consider that the following before restoring files and folders from a virtual machine:

- The virtual machine has the MBR partition. If the virtual machine has GPT partition, you can restore files from a backup copy. For more information, refer to Restore Data from a Backup Copy.
- You must perform the backup using the **VMware Storage API method**. For more information, refer to Configuring Backups for vSphere VADP Environments.
- You cannot restore any archived files and folders.
- It is recommended to perform the file-level restores from disk or volume-level backups only when you are restoring small files. For example, restoring a 2GB file from a disk-level backup is not recommended.

You can restore files from NTFS file systems with the following limitations:

- The formatted cluster size (allocation unit) must be 1024 or greater.
- You cannot restore any archived files and folders.
- You cannot restore any file that has zero bytes, symbolic links, or hard links associated with it.
- You cannot restore files that have been dehydrated by Windows 2012 dedup.
- You cannot restore files from volumes created on Windows Storage Spaces.

If you cannot restore a file, you can restore the complete virtual machine or the disk that



contains the file.

When restoring a file or folder on a virtual machine, the **Enable Granular Recovery** must have been selected when the backup was performed.

1. From the CommCell Console, perform a **Browse and Restore** operation.
2. In the **Browse Options** window, click **Individual files/folders**.
3. Click **OK**.
4. In the **Browse** window, click a file or folder in the right pane, and then click **Recover All Selected**.
5. From the **General** tab, specify a **Staging Path**. This is the destination path to which the file will be restored.
6. Click **OK**.

The **Preserve Source Path/Remove Source Path** feature is supported for File level restores from File level backup jobs. It is not relevant for Volume Level and Disk Level backups and restores.

LINUX CLIENTS


You can restore files and folders from a Linux virtual machine. Before restoring the files and folders, ensure that the virtual machine satisfies the following prerequisites:

- The virtual machine has the ext2 or ext3 file system.
- Image Level iDataAgent or Restore Only Agent for Image Level iDataAgent for Unix is installed on the Linux virtual machine.
- A Disk-Level backup has been performed and the **Enable Granular Recovery** has been selected when the backup was performed.

Do not restore Windows backup data to a Linux client.

You cannot restore any file that has zero bytes, symbolic links, or hard links associated with it. You must restore the complete virtual machine.

1. From the CommCell Console, perform a **Browse and Restore** operation.
2. In the **Browse Options** window, click **Individual files/folders**.
3. Click **OK**.
4. In the **Browse** window, select a file or folder in the right pane and click **Recover All Selected**.
5. From the **General** tab:
 - Choose the **Destination Client** from the dropdown list.
The destination client must have ext2 or ext3 file system.
 - Enter the **Staging Path** of the Linux client or click **Browse** to locate the path.
 - The **Restore ACLs** option is not supported for this operation.
 - The **Preserve Source Paths** option is not supported when you are restoring files or folders from a virtual machine.
 - When performing a file-level restore from a Linux virtual machine to a Linux client, select the **Unconditional Overwrite** option to restore folder permissions.
6. Click **OK**.



FINDING AND RESTORING

Use Find to search and locate files or folders from a data protection backup. Find is available from a backup set, subclient, or by right-clicking a virtual machine, volume or a folder from a Browse and Restore operation. If you perform the Find operation at the backupset level, the last backed up subclient will be searched. It is recommended to perform the Find operation at a subclient level to get accurate results.

PREREQUISITES

Ensure that the following prerequisites are satisfied before using the Find feature:

- Ensure that you are not using Linux MediaAgent to backup the data. You cannot find files or folders from the data that is backed up using a Linux MediaAgent.
- If you want to find files or folders on a Linux virtual machine, ensure that the virtual machine has ext2 or ext3 file systems.
- You can find files or folders from a SnapProtect Backup only for Windows virtual machines.

FINDING A FILE OR FOLDER

If you want to restore a specific file or folder from a backup set or subclient, you can search the file or folder in the backup set or subclient.


Before searching the subclient, ensure that all the files and folders in the subclient are included in the index. The files and folders are included in the index when you perform the **File Level** backup or **Enable Granular Recovery** before performing the **Volume Level** or **Disk Level** backup.

1. From the CommCell Browser, navigate to **Client Computers | <Client> | Virtual Server | <Instance>**.
2. Right-click **Backup Set** or **Subclient**, click **All Tasks | Find**.
3. Enter the file name in the **Name/Pattern** box. You can use wildcards to specify the pattern of file names.

These wildcards, or a combination of these wildcards, are supported in the **Subject** field:

- o * for any number of characters e.g., Accounting*.
 - o ? for any one character e.g., user_group?.
 - o [] for a range of characters e.g., [a-k]Lee.
 - o [!] for a negation of a range of characters e.g., ![a-k] Lee.
4. Enter the path of the location where you want to search the file. Do not type the drive letter while specifying the path.
 5. Click **Search**.

The files that match the criteria are displayed in Search Results.




FINDING A FILE OR FOLDER ON MULTIPLE VIRTUAL MACHINES

If a file or folder exists on multiple virtual machines and you want to restore files from specific machines, follow the steps given below to search the file:

1. From the CommCell Browser, navigate to **Client Computers | <Client> | Virtual Server | <Instance>**.
2. Right-click **Backup Set** or **Subclient**, click **All Tasks | Find**.
3. Enter the file name in the **Name/Pattern** box. You can use wildcards to specify the pattern of file names.
4. Enter the path of the location where you want to search the file. Do not type the drive letter while specifying the path.
5. Click **Advanced Options**.
6. Enter the name of Virtual Machine in the **VM Name Pattern** box. You can also use wildcards to specify the pattern of virtual machine names.
7. Click **Search**.

The files that match the criteria are displayed in Search Results.




FINDING A FILE OR FOLDER ON SPECIFIC VOLUME

If a file or folder exists on multiple volumes and you want to restore files from specific volumes, follow the steps given below to search the file:

1. From the CommCell Browser, navigate to **Client Computers | <Client> | Virtual Server | <Instance>**.
2. Right-click **Backup Set** or **Subclient**, click **All Tasks | Find**.
3. Enter the file name in the **Name/Pattern** box. You can use wildcards to specify the pattern of file names.
4. Enter the path of the location where you want to search the file. Do not type the drive letter while specifying the path.
5. Click **Advanced Options**.
6. Enter the name of Virtual Machine in the **VM Name Pattern** box. You can also use wildcards to specify the pattern of virtual machine names.
7. Enter the drive letter of the volume in the **Volume Name Pattern** box. You can also use wildcards to specify the pattern of volumes.
8. Click **Search**.

The files that match the criteria are displayed in Search Results.




FINDING A FILE BACKED UP IN SPECIFIC TIME RANGE

When you are using the Find option for a subclient or a backupset you cannot specify time range. You can find files and folders from the latest backup cycle. If you want to find data backed up during a specific time period, follow the steps given below:

1. From the CommCell Browser, navigate to **Client Computers | <Client> | Virtual Server | <Instance>**.
2. From the CommCell Browser, navigate to **Client Computers | <Client> | Virtual Server | <Instance>**.
3. Right-click **Backup Set** or **Subclient**, click **Browse Backup Data**.
4. Select **Specify Browse Time**.
5. Select a date and time to **Browse Data Before** it.
6. Click **OK**.
7. Right-click a volume, virtual machine or a folder and select **Find**.
5. Enter search criteria (e.g., File name or Folder name) as explained in the previous sections.
6. Click **Search**.

The files that match the criteria are displayed in Search Results.



QUICK SEARCH USING SAMPLING

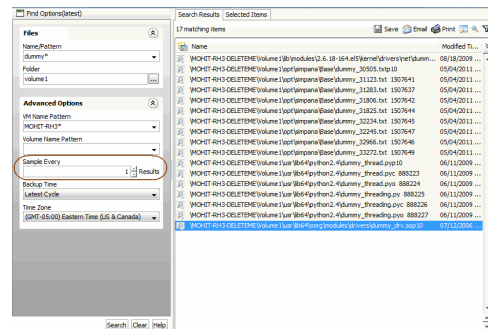
1. From the CommCell Browser, navigate to **Client Computers | <Client> | Virtual Server | <Instance>**.
2. Right-click **Backup Set** or **Subclient**, click **All Tasks | Find**.
3. Click **Advanced Options**.
4. In **Sample Every**, specify the rate at which files are sampled to determine whether they match the search criteria. For example: If 20 files on a virtual machine are matching the search criteria and specified sampling rate is 5, then every 5th file is displayed in the search result.

Keep in mind that the order in which files are sampled is based on the order in which the files were backed up.

The default value of 1 indicates that every file will be sampled.

5. Click **Search**.

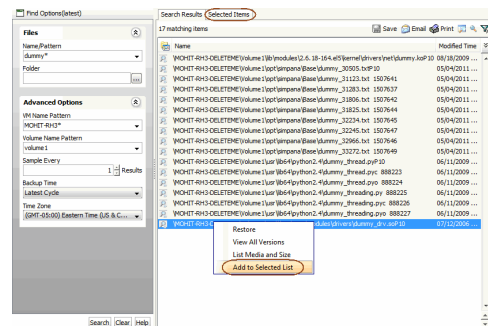
The files that match the criteria are displayed in Search Results.



ADDING SEARCH RESULTS TO A LIST

You can control the the items you need to restore by moving items to the Selected List. You can narrow down the list from the Search Results so that you are able to recover only the items from the Selected List.

1. From the CommCell Browser, navigate to **Client Computers | <Client> | Virtual Server | <Instance>**.
2. Right-click **Backup Set** or **Subclient**, click **All Tasks | Find**.
3. Enter search criteria (e.g., File name, Folder, VM Name, etc.) as explained in the previous sections.
4. Click **Search**.
- The files that match the criteria are displayed in Search Results.
5. Right-click a file and select **Add to Selected List**.
6. Click the **Selected Items** tab and the file will be displayed.




USING FIND FROM A BROWSE WINDOW

When you access Find as part of a Browse operation, keep in mind that the time range values are not pre-populated by any previously selected time range settings entered on the Browse Options or Advanced Browse Options dialog.

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


2. Right-click **Backup Set** or **Subclient**, click **Browse Backup Data**.
3. Right-click a volume, virtual machine or a folder and select **Find**.
5. Enter search criteria (e.g., File name or Folder name) as explained in the previous sections.
6. Click **Search**.

The search results are displayed in the right pane.



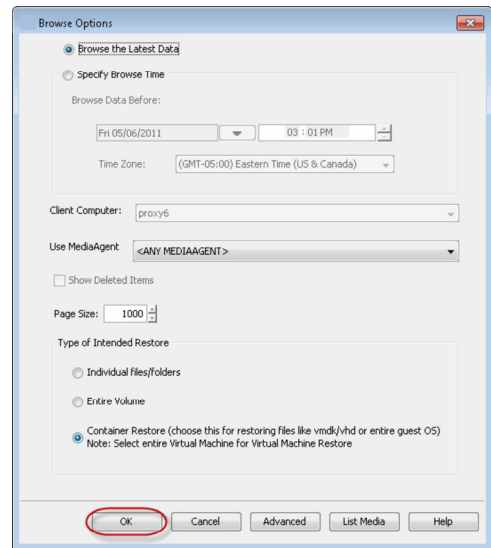
RESTORING FILES OR FOLDERS FROM THE SEARCH RESULTS

1. From the CommCell Browser, navigate to **Client Computers | <Client> | Virtual Server | <Instance>**.
2. Right-click **Backup Set** or **Subclient**, click **All Tasks | Find**.
3. Enter search criteria and click **Search**.
4. Right-click a file or folder from the search results and then select **Restore**.
5. Specify the Staging Path.
6. Click **OK**.

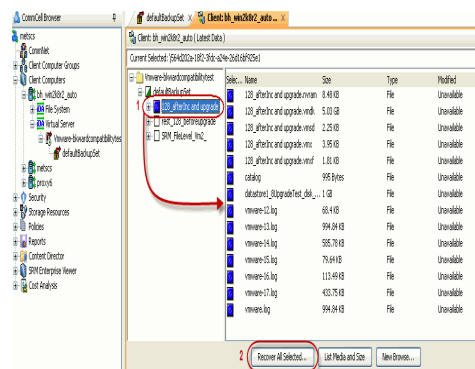


SCHEDULING A RESTORE

1.
 - From the CommCell Console, navigate to **Client Computers | Virtual Server**
 - Right-click an instance and click **All Tasks | Browse**.
 - Select a restore type and click **OK**.



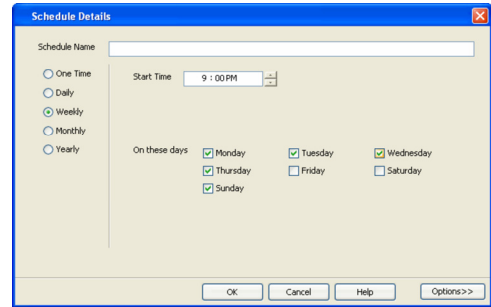
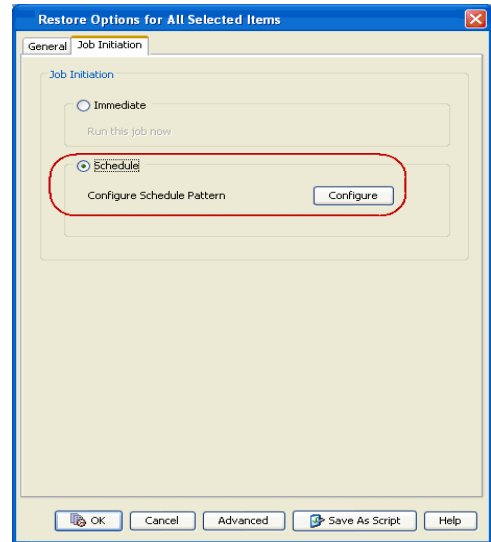
2.
 - Select a virtual machine, volume, or file to restore.
 - Click **Recover All Selected**.



3.
 - Click **Job Initiation** tab and click **Schedule**.
 - Click **Configure**.

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

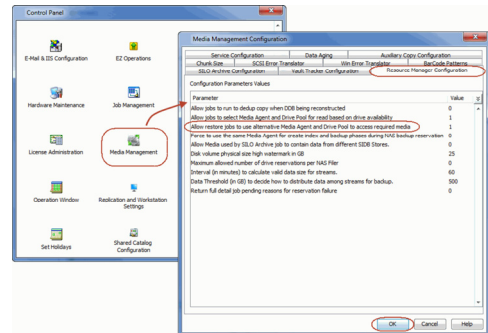
The restore job will execute as per the schedule.



RESTORING GUEST FILES FROM UNIX MEDIAAGENT

For container, volume, and file-level restores from disk and volume-level backups using any Unix (AIX, HP-UX, or Solaris) MediaAgent, perform the following steps:

1. From the CommCell Console, click **Control Panel** and then double-click **Media Management**.
2. Click the **Resource Manager Configuration** tab.
3. Enter the **Value 1** for the **Allow restore jobs to use alternative Media Agent and Drive Pool to access required media** parameter.
4. Click **OK**.
5. From the CommCell Console, navigate to **Client Computers | Virtual Server**
6. Right-click an instance and click **All Tasks | Browse**.
7. Select a Windows MediaAgent.
8. Select a restore type and click **OK**.
9. Select the virtual machine from which you want to restore the files and folders. The **List Media** dialog box appears.
10. Click **Select All**.
11. Click **Restore Index**. The index of the Unix MediaAgent will be restored to the selected Windows MediaAgent. All the volumes, files and folders on the selected virtual machine will be displayed.



The index restore operation can be avoided if the index cache is shared between the Unix and Windows MediaAgents. For more information about sharing the index cache, refer to Configure a MediaAgent for Index Cache Sharing.

12. Select the guest files and folder for restore and click **Recover All Selected**. You cannot select an entire volume for restore.
13. Click **OK**.

CREATING VIRTUAL MACHINES FROM RESTORED VMX FILES

New virtual machines can be created with select volumes or disks from **Container** restores of disk-level backups. The VMware Converter utility is used to accomplish this. To do so:

1. Perform a **Browse and Restore** operation.
2. From the **Browse** window, select the container that was restored from the virtual machine to which it was restored.
3. Select the .vmx file from the location containing the container.
4. Click **OK**.
5. Provide the credentials for the Virtual Center or ESX Server to which you want to restore the .vmx file.

If the credentials provided are for the Virtual Center, you must also select the desired ESX Server to which you want to restore the .vmx file.

6. Select the **Select volume copy from Data copy** option.
7. Deselect any volumes you do not want to restore.

The virtual machine will be restored with the selected volumes.

BROWSE DATA

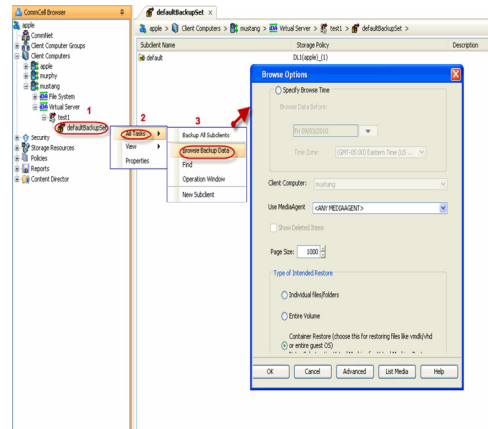
The option to browse the backup data provides the facility to view and selectively restore/recover the data that was backed up. The browse option can be invoked from the client, agent, instance, backup set, or subclient level depending on the functionality of a given agent.

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

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

Use the following steps to browse data:

1. From the CommCell Browser, navigate to **<Client>**.
2. Right-click **Instance** or **Backup Set** or **Subclient**, click **All Tasks | Browse/Browse Backup Data**.
3. Click **OK**.



BROWSING DATA FROM BEFORE THE MOST RECENT FULL BACKUP


To browse data that is older than the most recent full backup specify a **Browse Data Before** date that pre-dates the full backup. Remember, the **Browse Data Before** date establishes the ending point of the search. Consequently, using a **Browse Data Before** date that pre-dates the most recent full backup starts the search in the previous full backup cycle. This is only valid if the data in that full backup cycle has not expired.

BROWSE DATA BEFORE A SPECIFIED TIME

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

1. From the CommCell Browser, navigate to **<Client>**.
2. Right-click **Instance** or **Backup Set** or **Subclient**, click **All Tasks | Browse/Browse Backup Data**.
3. Select **Specify Browse Time** and the **Time Zone**.

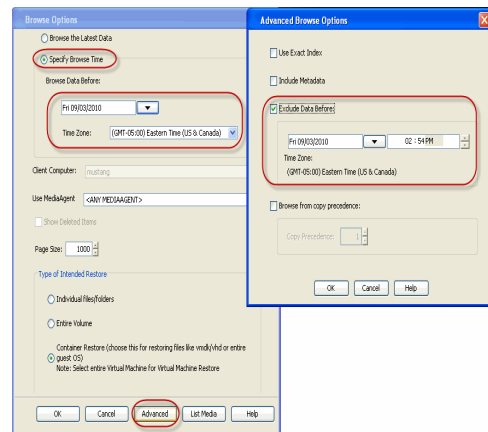
4. Click **OK**.



BROWSE DATA BETWEEN A SPECIFIED TIME

Follow the steps given below to browse data backed up during a specified time interval:

1. From the CommCell Browser, navigate to **<Client>**.
2. Right-click **Instance** or **Backup Set** or **Subclient**, click **All Tasks | Browse/Browse Backup Data**.
3. Select **Specify Browse Time** and the **Time Zone**.
4. Click **Advanced**.
5. Select **Exclude Data Before**, select the date and time from which you wish to view the data in the browse window.
6. Click **OK**.



LIST MEDIA

Use List Media to predict media required for the following operations:


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

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

LISTING MEDIA FOR A SUBCLIENT

Do the following:

1. From the CommCell Browser, navigate to the subclient and right-click **List Media**.
2. Select one of the following options:
 - o Click **Media For the Latest Data** to list media associated with the most recent data protection cycle.
 - o Click **Specify Time Range** to list media associated with data protection operations up to the specified date and time range.
Specify the end date and time in the **Data Before** box.
3. Select the **Type of Intended Restore**.
4. Click **OK**.




LISTING MEDIA FOR A BACKUP SET OR INSTANCE

Do the following:

- From the CommCell Browser, navigate to the backup set or instance, right-click **All Tasks | Browse**.
- Select one of the following options:
 - Click **Browse the Latest Data** to list media associated with the most recent data protection cycle.
 - Click **Specify Browse Time** to list media associated with data protection operations up to the specified date and time range.

Use the **Browse Data Before** box to specify the end date and time.
- Click **List Media**.
- Select one of the following options:
 - Click **List Media for restore within specified time range**
 - Click **List Media containing index required for browse**
- Click **OK**.



RESTORE BY JOBS

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

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

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

Consider the following when restoring by jobs:

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

Use the following steps to restore data by jobs.

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

MANAGE RESTORE JOBS

Once you initiate the restore operation, a restore job is generated in the Job Controller. Jobs can be managed in a number of ways. The following sections provide information on the different job management options available:

RESTART JOBS

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

It is also possible to override the default CommServe configuration for individual jobs by configuring retry settings when initiating the job. This configuration, however, will apply only to the specific job.

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

CONFIGURE JOB RESTARTABILITY AT THE COMMSERVE LEVEL

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

CONFIGURE JOB RESTARTABILITY FOR AN INDIVIDUAL JOB

1. From the **Restore Options** dialog box, click **Advanced**, then select the **Job Retry** tab and specify the following as desired:
 - o **Total Running Time** - The maximum elapsed time, in hours and minutes, from the time that the job is created.
 - o **Number of Retries** - The number of times that Job Manager will attempt to restart the job.
 - o **Kill Running Jobs When Total Running Time Expires** - Option to kill the job when the specified Total Running Time has elapsed, even if its state is "Running".
2. Click **OK**.

RESUBMIT JOBS

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

RESUBMIT A RESTORE JOB

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

CONTROL JOBS

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

SUSPEND	Temporarily stops a job. A suspended job is not terminated; it can be restarted at a later time. Only preemptible jobs can be suspended.
RESUME	Resumes a job and returns the status to Waiting, Pending, Queued, or Running depending on the availability of resources or the state of the operation windows and activity control settings. Restore jobs for this Agent are resumed from the beginning.
KILL	Terminates a job.

These controls can be applied to:

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

CONTROL A JOB

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

Click **Yes** when the confirmation prompt appears if you are sure you want to kill the job. The job status may change to **Kill Pending** for a few moments while the operation completes. Once completed, the job status will change to **Killed** and it will be removed from the Job Controller window

after five minutes.

- When suspending a job:

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

- When resuming a job:

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

ADDITIONAL RESTORE OPTIONS

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

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

OPTION	DESCRIPTION	RELATED TOPICS
Use hardware revert capability if available	<p>This option allow you to revert the data to the time when the snapshot was created. Selecting this option brings back the entire virtual machine to the point when the snapshot was created, overwriting all the modifications to data since the snapshot creation.</p> <p>Revert operations for a virtual machine are supported on NetApp File Servers but not from SnapVault or SnapMirror snapshots. Other file servers are not supported. To perform a revert operation, the SnapRestore license is required on the NetApp file server.</p> <ol style="list-style-type: none"> 1. From the CommCell Browser, navigate to Client Computers Virtual Server <Instance> 2. Right-click <backup set> and click All Tasks Browse Backup Data. 3. From the Browse Options dialog box, select Container Restore and click OK. 4. Select the virtual machine that you want to revert and click Recover All Selected. 5. From the Restore Options dialog box, click Advanced. 6. Select the Use hardware revert capability if available option. 7. Click OK to confirm the revert operation. 8. Click OK from the Advanced Restore Options dialog box. 9. Click OK to start the revert. 	SnapProtect Backup
Automatically Detect Regular Expressions	<p>This option allows you to automatically detect regular expressions (wildcard characters) in the virtual machines assigned to a subclient.</p> <p>This allows you to restore the files and folders in virtual machines which contain regular expressions in their name (e.g., c:\[test1]).</p> <ol style="list-style-type: none"> 1. From the CommCell Browser, click Client Computers Virtual Server <instance name> BackupSet. The default and other subclients (if available) are displayed on the right-hand windowpane. 2. Right-click the backupset that contains the data you want to restore. Click All Tasks and then click Browse Backup Data. 3. From the Browse Options window, select the Type of Intended Restore and click OK. 4. From the browse window, select the data to be restored and click Recover all Selected. 5. From the Restore Options for All Selected Items dialog box, click Advanced. 6. From the Advanced Restore Options (General) tab, select Automatically Detect Regular Expressions option. 7. Click OK. 	
Impersonate User	<p>This information is needed only if you intend to restore data to a mapped network drive, shared network drive, or directory to which you do not have access privileges.</p> <p>You can impersonate a windows logon information with an account that has access rights to restore data.</p> <ol style="list-style-type: none"> 1. From the CommCell Browser, click Client Computers Virtual Server <instance name> BackupSet. The default and other subclients (if available) are displayed on the right-hand windowpane. 2. From the CommCell Browser, right-click the backupset that contains the data you want to restore. Click All Tasks and then click Browse Backup Data. 3. From the Browse Options window, select the Type of Intended Restore and click OK. 4. From the browse window, select the data to be restored and click Recover all Selected. 5. From the Restore Options for All Selected Items dialog box, click Advanced. 6. From the Advanced Restore Options (General) tab, select Impersonate User option. 7. In User Name and Password boxes, enter a user name and password which has access 	

	<p>privileges. In Confirm Password box, type the password again.</p> <p>8. Click OK.</p>	
Startup Options	<p>The Job Manager will use the startup priority setting when allocating the required resources. This is useful if you have jobs that are very important and must complete, or jobs that can be moved to a lower priority.</p> <ol style="list-style-type: none"> 1. From the CommCell Browser, click Client Computers Virtual Server <instance name> BackupSet. The default and other subclients (if available) are displayed on the right-hand windowpane. 2. From the CommCell Browser, right-click the backupset that contains the data you want to restore. Click All Tasks and then click Browse Backup Data. 3. From the Browse Options window, select the Type of Intended Restore and click OK. 4. From the browse window, select the data to be restored and click Recover all Selected. 5. From the Restore Options for All Selected Items dialog box, click Advanced. 6. From the Advanced Restore Options (Startup) tab, select Use Default Priority option. 7. Select Change Priority option to change the priority for a job, between 0 (highest priority) and 999 (lowest priority). 8. Select Start up in Suspended State option to specify the selected job to start in the Job Controller in a suspended state. 9. Click OK. 	Refer to Job Priority and Priority Precedence.
Copy Precedence	<p>When you select Restore from copy precedence option, the system recovers or retrieves data from the selected storage policy copy (Synchronous Copy or Selective Copy). If data does not exist in the specified copy, the data recovery or retrieve operation fails even if the data exists in another copy of the same storage policy.</p> <p>When cleared, (or by default) the system recovers or retrieves data from the storage policy copy with the lowest copy precedence. If the data was pruned from the primary copy, the system automatically recovers or retrieves from the other copies of the storage policy in the lowest copy precedence to highest copy precedence order. Once the data is found, it is recovered or retrieved, and no further copies are checked.</p> <ol style="list-style-type: none"> 1. From the CommCell Browser, click Client Computers Virtual Server <instance name> BackupSet. The default and other subclients (if available) are displayed on the right-hand windowpane. 2. From the CommCell Browser, right-click the backupset that contains the data you want to restore. Click All Tasks and then click Browse Backup Data. 3. From the Browse Options window, select the Type of Intended Restore and click OK. 4. From the browse window, select the data to be restored and click Recover all Selected. 5. From the Restore Options for All Selected Items dialog box, click Advanced. 6. From the Advanced Restore Options (Copy Precedence) tab, select Restore from copy precedence option. 7. Select the number to change the selected copy's precedence. The numbers will move a copy to a higher or lower precedence in increments of 1. 8. Click OK. 	Refer to Recovering Data From Copies.
Data Path Options	<p>Data Recovery operations use a specific data path (Library, MediaAgent, Drive Pool, and Drive) to perform the restore operations as configured in the CommCell. By default, the system automatically identifies the data path for the restore operations. You can change the data path at the restore level by selecting a specific MediaAgent, library, drive pool or drive in the Advanced Restore Options dialog.</p> <p>The following are some of the important situations where you may need to change the data path:</p> <ul style="list-style-type: none"> • To free heavily loaded resources such as MediaAgents and libraries. • A library suffers severe hardware damage and you need to move its media to a different library in order to access data secured by data protection operations. • A drive pool within a library is being used much more frequently than another, and a storage policy copy from the more active drive pool can be assigned to the less active one. <ol style="list-style-type: none"> 1. From the CommCell Browser, click Client Computers Virtual Server <instance name> BackupSet. The default and other subclients (if available) are displayed on the right-hand windowpane. 2. From the CommCell Browser, right-click the backupset that contains the data you want to restore. Click All Tasks and then click Browse Backup Data. 3. From the Browse Options window, select the Type of Intended Restore and click OK. 4. From the browse window, select the data to be restored and click Recover all Selected. 5. From the Restore Options for All Selected Items dialog box, click Advanced. 	Refer to Change Data Path.

	<ol style="list-style-type: none"> 6. From the Advanced Restore Options (Data Paths) tab, select the MediaAgent and Library that you wish to perform the restore. 7. Select the Drive Pool and Drive for optical and tape libraries. 8. Select the name of the Proxy server if you wish to restore using a proxy server. 9. Click OK. 	
Encryption	<p>You can configure encryption for the specific subclient during the data recovery operation. You must configure encryption at the client level first and then at the subclient level to encrypt data during data recovery operations using the CommCell Console.</p> <p>If you configure data encryption with a Pass-Phrase and do not elect to export the pass-phrase to destination clients, you will be required to enter the pass-phrase during immediate data recovery operations. Scheduled data recovery operations require you to export the pass-phrase.</p> <ol style="list-style-type: none"> 1. From the CommCell Browser, click Client Computers Virtual Server <instance name> BackupSet. The default and other subclients (if available) are displayed on the right-hand windowpane. 2. From the CommCell Browser, right-click the backupset that contains the data you want to restore. Click All Tasks and then click Browse Backup Data. 3. From the Browse Options window, select the Type of Intended Restore and click OK. 4. From the browse window, select the data to be restored and click Recover all Selected. 5. From the Restore Options for All Selected Items dialog box, click Advanced. 6. From the Advanced Restore Options (Encryption) tab, enter a pass-phrase in the Pass Phrase dialog box. 7. Enter the pass-phrase again in the Re-enter Pass Phrase dialog box. 8. Click OK. 	Refer to data encryption.
Job Retry Options	<p>The restore jobs can be restarted either by a user or automatically by the Job Manager. You can specify the maximum elapsed time before a job can be restarted or killed and the maximum number of restart attempts.</p> <ol style="list-style-type: none"> 1. From the CommCell Browser, click Client Computers Virtual Server <instance name> BackupSet. The default and other subclients (if available) are displayed on the right-hand windowpane. 2. From the CommCell Browser, right-click the backupset that contains the data you want to restore. Click All Tasks and then click Browse Backup Data. 3. From the Browse Options window, select the Type of Intended Restore and click OK. 4. From the browse window, select the data to be restored and click Recover all Selected. 5. From the Restore Options for All Selected Items dialog box, click Advanced. 6. From the Advanced Restore Options (Job Retry) tab, select Enable Total Running Time option. 7. Select Enable Number of Retries option. 8. Select Kill Running Jobs When Total Running Time Expires option. 9. Click OK. 	Refer to Job Management.
ESX Browse selection	<p>Virtual machines can be restored to any datastore on the ESX Server using the credentials of the default ESX Server or the Virtual Center.</p> <p>You can Override the default selection and select the desired location to which the virtual machine will be restored. These options provide the flexibility to distribute restored virtual machines to a location with greater space and resource availability.</p> <ol style="list-style-type: none"> 1. From the CommCell Browser, click Client Computers Virtual Server <instance name> BackupSet. The default and other subclients (if available) are displayed on the right-hand windowpane. 2. From the CommCell Browser, right-click the backupset that contains the data you want to restore. Click All Tasks and then click Browse Backup Data. 3. From the Browse Options window, select the Type of Intended Restore and click OK. 4. From the browse window, select the data to be restored and click Recover all Selected. 5. From the Restore Options for All Selected Items dialog box, click Advanced. 6. From the Advanced Restore Options (ESX Browse) tab, select Default Selection option. Select the ESX Server from the drop down list. 7. Select the Override Default Selection to change the default selection. 8. Enter the desired Virtual Center/ESX Server name. 9. Select Server Logon Account and enter the User Account. 10. Click Configure Password to create a password. 	

	<ol style="list-style-type: none"> 11. Click Browse Virtual Center/ESX Servers for destination option to select the desired location to which the virtual machine will be restored. 12. Click OK. 	
Alert	<p>The Alert option is used for setting up the criteria to raise notifications/alerts for job statuses such as failure, success, or any other conditions triggered by the restore job. Adding alerts helps the user or the user group to get the notification automatically about the status of the restore job.</p> <ol style="list-style-type: none"> 1. From the CommCell Browser, click Client Computers Virtual Server <instance name> BackupSet. The default and other subclients (if available) are displayed on the right-hand windowpane. 2. From the CommCell Browser, right-click the backupset that contains the data you want to restore. Click All Tasks and then click Browse Backup Data. 3. From the Browse Options window, select the Type of Intended Restore and click OK. 4. From the browse window, select the data to be restored and click Recover all Selected. 5. From the Restore Options for All Selected Items dialog box, click Advanced. 6. From the Advanced Restore Options (Alerts) tab, Click Add Alert. 7. From the Add Alert Wizard window, select the Threshold and Notification Criteria. Click Next. 8. Select the Notification Types. Click Next. 9. Select the Users and User Groups. Click Next. 10. Verify the Summary. 11. Click Finish. 	Refer to Alerts.
Command Line Restores	<p>Command Line Interface enables you to perform restore operations from the command line. The qcommands can be executed from the command line or can be integrated into your own scripts or scheduling programs.</p> <p>In addition, you can also generate scripts for specific operations from the CommCell Console using the Save As Script option. These scripts can later be executed using the commands from the command line interface.</p>	Refer to Command Line Interface.
CommCell Readiness Report	CommCell Readiness Report provides you with important information about the potential problems that can impact your restore operations before they occur.	Refer to CommCell Readiness Report.
Restore Job Summary Report	Restore Job Summary Report provides you the summary of all restore jobs during a specified time period for each client. Select Restore option from the Job Summary Report Selection dialog box, to view the Restore Job Summary Report.	Refer to Restore Job Summary Report.

[Back to Top](#)

Data Aging - VMware

TABLE OF CONTENTS

Getting Started


Extended Retention Rules

GETTING STARTED


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

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


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




2. Highlight the **Storage Policy**.



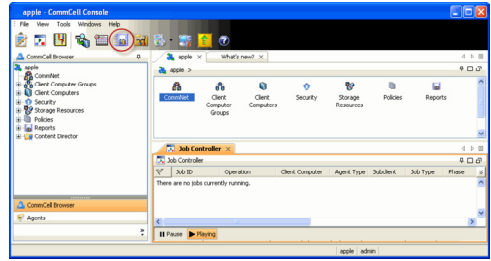
3. From the right pane, right-click the **Storage Policy Copy** and click the **Properties**.




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




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



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



7. Click **Run**.

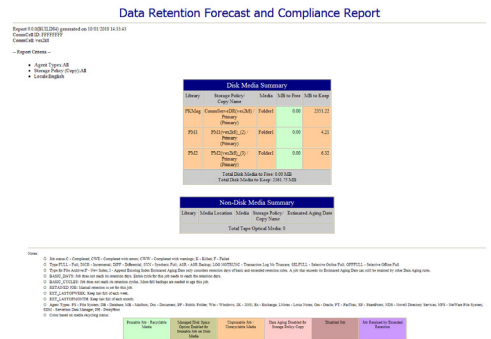


8. The report will display the data to be pruned when a data aging job is run.


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

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


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



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




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



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

Make sure that the job completes successfully. If the job did not complete successfully, re-run the job.



EXTENDED RETENTION RULES

Extended retention rules allow you to keep specific full (or synthetic full) backups for longer periods of time.

Extended retention rules can be used in the following circumstances:

- If you have a single drive tape library
- If you want to create a hierarchical retention scheme (grandfather-father-son tape rotation)

In all other cases, it is recommended that the Auxiliary Copy feature be used for extended storage as it actually creates another physical copy of the data, thereby reducing the risk of data loss due to media failure.

UNDERSTANDING EXTENDED RETENTION RULES

Extended retention allows you to retain a specific full (or synthetic full) backup for an additional period of time. For example, you may want to retain your monthly full backups for 90 days.

Extended retention rules allow you to define three additional "extended" retention periods for full (or synthetic full) backups. For example:


- You may want to retain your weekly full backups for 30 days.
- You may want to retain your monthly full backup for 90 days.
- You may want to retain your yearly full backup for 365 days.

A backup job will be selected for extended retention based on its start time. For example: If a backup job starts at 11:55 pm on August 31st and ends at 1 am on September 1st, then it will be selected as the last full backup for the month of August and will be picked up for extended retention.

SETTING UP EXTENDED RETENTION RULES

Use the following steps for setting up the extended retention rules:

1. Right-click the storage policy copy and click **Properties**.
2. Click the **Retention** tab.
3. Set the basic retention rules by clicking **Retain for** and entering the number of days and cycles appropriate for your organization.
4. Set the extended retention rules as follows:
 - Click the **For** button.
 - Enter the number of **Days Total** to retain the backup.
 - Click the **Keep** drop-down list, and select the desired backup criteria (e.g., Monthly Full).
 - Click the **Grace Days** drop-down list and select the number of days (e.g., 2).
5. Repeat Step 4 to configure additional extended retention.
6. Click **OK**.



ADVANCED TOPICS

Data Aging - Advanced

Provides comprehensive information on additional Data Aging capabilities.

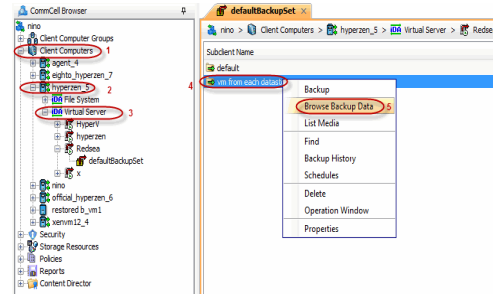
Disaster Recovery - VMware

PLANNING FOR A DISASTER RECOVERY

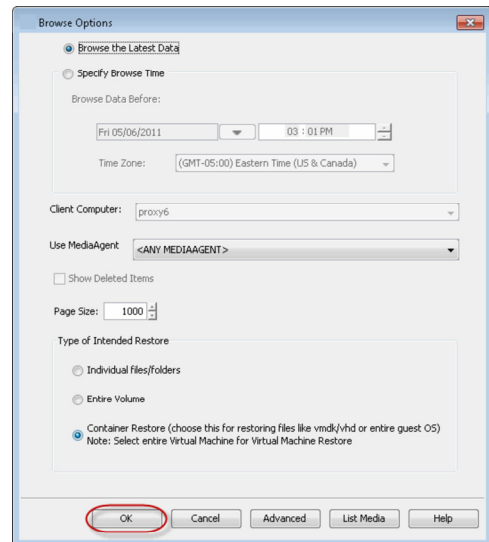
While multiple virtual machines can always be restored simultaneously at any time, disaster scenarios may require them to be readily available at the moment they are needed.

To achieve this, the recommended disaster recovery plan consists of creating a Stand-by Server to which virtual machines are restored on a routine basis. In the event of a disaster, the Stand-by server is readily available with the virtual machines.


1. Create a **Stand-by Server** consisting of the following:
 - The VMware application
 - Sufficient memory and disk space to hold all the virtual machines you may need to recover in a disaster scenario.
2.
 - From the CommCell Console, navigate to the <Client> | **Virtual Server**.
 - Right-click the subclient that contains the data you want to restore, point to **Browse Backup Data**.



3. Click **OK**.



4.
 - Select the backup set in the left pane. Its entire contents will be automatically selected in the right pane.
 - Click **Recover All Selected**.

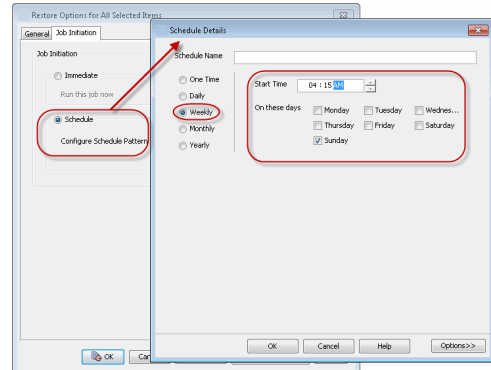
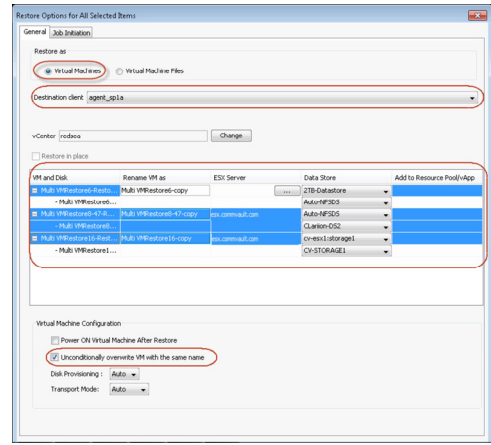


5.
 - Click **Virtual Machines**.
All the selected virtual machines are displayed in the **VM and Disk** column.
 - Enter corresponding values for each virtual machine you want to restore.
 - Choose the Stand-by Server in the **Destination Client** dropdown list.
 - Select **Unconditionally overwrite VM with same name** option.

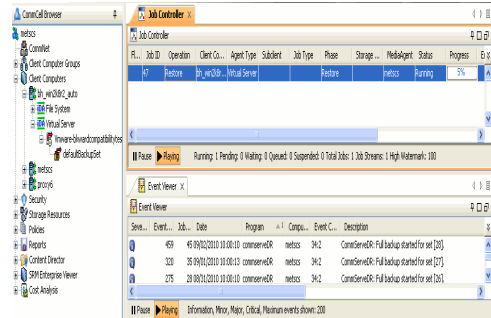
6.
 - Click the **Job Initiation** tab.
 - Select the **Schedule** radio button and then click **Configure**.
 - Enter a name for the schedule and configure the desired frequency for which the restore will occur.

To ensure the latest virtual machine data is available for the restore operations, create a schedule that corresponds to any backup schedules you have configured. For example, if you back up your virtual machines weekly on a Saturday morning, you may consider scheduling your restores to occur weekly on a Sunday morning.

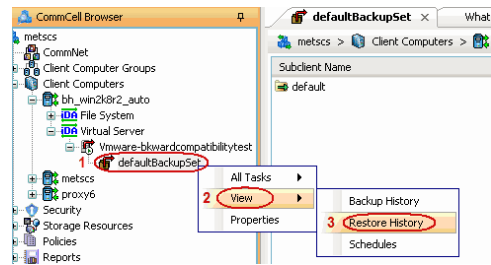
 - Click **OK** to save your changes.
 - Click **OK**.



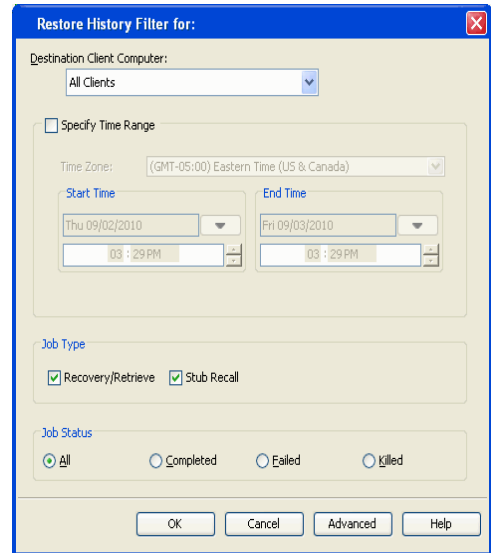
7. You can monitor the progress of the restore job in the **Job Controller** or **Event Viewer** window of the CommCell Console.



8. Once the restore job has completed, right-click the entity (e.g. agent, instance) and click **View | Restore History**.
If the entity chosen is the client computer, click **View | Job History**.



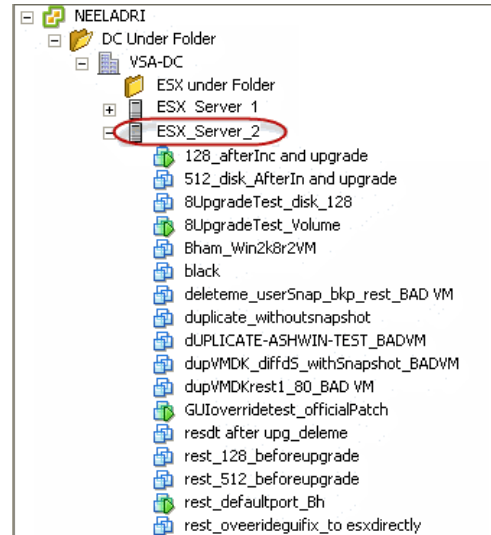

9. Click **OK**.



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

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

11. Once the virtual machines are restored, they are automatically mounted to the Stand-by Server's virtual center/ESX Server.



Back to Top

Additional Operations - VMware

TABLE OF CONTENTS

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

AUDIT TRAIL

The Audit Trail feature allows you to track the operations of users who have access to the CommCell. This capability is useful if a detrimental operation was performed in the CommCell and the source of that operation needs to be determined.

Audit Trail tracks operations according to four severity levels:

- **Critical:** This level records operations that will result in imminent loss of data.
- **High:** This level records operations that may result in loss of data.
- **Medium:** This level records changes to the general configuration of one or more entities. Such changes may produce unintended results when operations are performed.
- **Low:** This level records changes to status, addition of entities, and other operations that have minimal impact on existing CommCell functions.

To set Audit Trail retention periods:

1. From the **Tools** menu in the CommCell Console, click **Control Panel**, and then select **Audit Trail**.
2. From the **Audit Trail** dialog box, select the desired retention time (in days) for each severity level.
3. Click **OK**.

STORAGE POLICY

A Storage policy defines the data lifecycle management rules for protected data. Storage policies map data from its original location to a physical storage media and determine its retention period.

1. Expand the **Policies** node, right-click **Storage Policies**, and select **New Storage Policy**.
2. Click **Next**.
3. Select **Data Protection and Archiving** to create a regular storage policy or **CommServe Disaster Recovery Backup** to backup the CommServe database and click **Next**.
4. Click **Next**.
5. Enter the name of storage policy and click **Next**.
6. Enter the name of the primary copy and click **Next**.
7. From the drop down box, select the default library for the primary copy and click **Next**.
8. From the drop down box, select the MediaAgent and click **Next**.
9. Enter number of data streams and set the retention period for the policy and click **Next**.
10. Click **Next**.
11. Click **Browse**, browse to your designated deduplication store location and click **Next**.
12. Confirm your selections and click **Finish**.

SUBCLIENT POLICY

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

RELATED TOPICS

Additional Operations - Concepts

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

- [Audit Trail](#)
- [Storage Policies](#)
- [Storage Policy Copies](#)
- [Subclient Policies](#)
- [Schedule Policy](#)
- [Auxiliary Copy](#)
- [Operation Window](#)
- [Cost Analysis](#)
- [Uninstalling Components](#)

1. Expand the **Policies** node, right-click **Subclient Policies** and click **New Subclient Policy**.
2. Enter policy name in the **Name** entry field, then select the appropriate file system type from the **iDataAgent** drop-down list.
3. In the Subclient pane, assign a storage policy for each subclient template by clicking in the drop down box under the **Storage Policy** column.
4. Click **OK**.

SCHEDULE POLICY

A schedule policy is a defined schedule or group of schedules for specific operations to be performed on associated objects within the CommCell. When the schedules from a policy are run, the specified operations, (e.g., auxiliary copy, backup, etc.) will be performed on the associated CommCell objects.

1. Expand the **Policies** node, right-click **Schedule Policies** and click **Add**.
2. Type the **Name** of the schedule policy.
3. Select the **Type** of schedule policy.
4. Select the **Agent Type**.
5. Type a description of the schedule policy.
6. Click **Add**.
7. Enter a Schedule Name in the **Schedule Pattern** tab.
8. Click **OK**.
9. On the **Associations** tab, select the objects to be associated with the schedule policy.
10. Click **OK**.

AUXILIARY COPY

An auxiliary copy operation allows you to create secondary copies of data associated with data protection operations, independent of the original copy.

1. Expand the **Policies** node, right-click storage policy for which you want to perform an auxiliary copy, click **All Tasks**, and then click **Run Auxiliary Copy**.
2. If you are starting the auxiliary copy operation from the CommServe level, select the storage policy for which you wish to perform the auxiliary copy.

If you are starting the auxiliary copy operation from the Storage Policy level, the Storage Policy field is already populated with the name of the Storage Policy you selected.

3. If the source copy is configured with a shared library, select the **Source MediaAgent** for the auxiliary copy.
4. Click **OK** to start the auxiliary copy operation. A progress bar displays the progress of the operation.

OPERATION WINDOW

By default, all operations in the CommCell can run for 24 hours. To prevent certain operations from running during certain periods of the day, you can define operation rules so that these operations are disabled during those times.

When operation rules are configured, operations that are started within the time window specified will go to a queued (as opposed to pending) state. Once the time window specified in the operation rule has elapsed, these queued or running operations will resume automatically.

1. In the CommCell Browser, right-click the appropriate entity, click **All Tasks**, and then click **Operation Window**.
2. Click **Add**.
3. From the **Operation Window** dialog box:
 - o Enter the name of the rule in the **Name** field.
 - o Select either an administration, data protection (either full or non-full), and/or a data recovery operation from the **Operations** pane.
4. Click **OK**.

OPERATING SYSTEM AND APPLICATION UPGRADES

Operating system upgrades are only supported when upgrading from one version of an OS to a different version of the same OS (e.g., Win2003 to Win2008). The two methods of upgrading are:

- **Seamless Upgrade** - This involves uninstalling the Agent software, upgrading the operating system, and then re-installing the Agent software.
- **Full OS Replacement** - This involves performing a clean install of a new version of the OS, re-installing any application software, then re-installing the CommServe, MediaAgent, and/or Agent software.

For Full OS Replacement, the client computer must be configured to have the CommServe, MediaAgent, and/or Client software re-installed to the same location, the same Fully Qualified Domain Name or short domain name, the same partitions, disk drive format (FAT, NTFS, et. al.), and IP configuration as previously.

If it is necessary to remove Agent software to facilitate an operating system or application upgrade, do not delete the icon for the Agent from the CommCell Console, or all associated backed up data will be lost.

Use the following strategy to upgrade the operating system software:

- Identify the computers you want to upgrade and the CommCell components installed on each of these computers.
- Choose the type of upgrade procedure you want to use on each computer: seamless or full replacement.
- CommServe, MediaAgent, and Client computers can be upgraded in any order.

UNINSTALLING COMPONENTS

You can uninstall the components using one of the following method:

Method 1: Uninstall Components Using the CommCell Console

Method 2: Uninstall Components from Add or Remove Programs

METHOD 1: UNINSTALL COMPONENTS USING THE COMMCELL CONSOLE

1. From the CommCell Browser, right-click the desired Client Computer and click **All Tasks -> Add/Remove Programs** and click **Uninstall Software**.
2. **Uninstall Software Options** dialog will appear.
3. In the **Uninstall Software** tab, select **Uninstall All** to uninstall all the software packages.
4. In the **Job Initiation** tab, select **Immediate** to run the job immediately.

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

METHOD 2: UNINSTALL COMPONENTS FROM ADD OR REMOVE PROGRAMS

1. Click the **Start** button on the **Windows** task bar and then click **Control Panel**.
2. Double-click **Add/Remove Programs**.
For Windows Vista/Windows 2008, click Uninstall a Program in the **Control Panel**.
3. Click **Remove** to uninstall the components in the following sequence:
 1. <Agent>
 2. File System iDataAgent
 3. Base Software

ONLINE HELP LINKS

Use the following links to view the online help for the corresponding tabs in the CommCell Console:

OPERATIONS	ENTITY	ONLINE HELP LINKS	SUB LINKS
CONFIGURATION	Agent	Virtual Server Agent Properties (General) Virtual Server Agent Properties (Version) Virtual Server Agent Properties (Index) Virtual Server Agent Properties (Security) Virtual Server Agent Properties (Activity Control)	
	Instance	VMware Instance Properties (General) VMware Instance Properties (Security)	
	Subclient	VMware Subclient Properties (General) VMware Subclient Properties (Filters) VMware Subclient Properties (Pre/Post Process) VMware Subclient Properties (Security) VMware Subclient Properties (Encryption) VMware Subclient Properties (Storage Device)	

		VMware Subclient Properties (Auto Discover) VMware Subclient Properties (Activity Control) VMware Subclient Properties (Content) VMware Subclient Properties (SnapProtect Operations)	
	BackupSet	VMware BackupSet (General) VMware BackupSet (Guest Hosts Configuration) VMware BackupSet (Security)	
BACKUP	Backup Options	VMware Backup Options	Advanced Backup Options Save As Script Command Line XML Options
	Advanced Backup Options	VMware Advanced Backup Options (Data) VMware Advanced Backup Options (Startup) VMware Advanced Backup Options (Media) VMware Advanced Backup Options (Data Path) VMware Advanced Backup Options (VaultTracking) VMware Advanced Backup Options (Alert)	
BROWSE		VMware Browse Options VMware Advanced Browse Options VMware List Media	
RESTORE	Restore Options	VMware Restore Options (General)	Save As Script Command Line XML Options
	Advanced Restore Options	VMware Advanced Restore Options (General) VMware Advanced Restore Options (Copy Precedence) VMware Advanced Restore Options (Data Path) VMware Advanced Restore Options (Encryption) VMware Advanced Restore Options (Startup) VMware Advanced Restore Options (Alert) VMware Advanced Restore Options (Options) VMware Advanced Restore Options (Browse)	
Operation Window	Operation Window	VMware Operation Window	

[Back to Top](#)

VM Lifecycle Management - Administrator

[Overview](#)
[Install the Virtual Server iDataAgent](#)
[Install the Web Search Server and Client](#)
[Setup](#)
[User Information](#)
[< Previous](#)
[Next >](#)

The VM Lifecycle Management feature provides virtual machine users with the ability to create, manage, and recover their own virtual machines from a user-friendly interface. Users can perform following tasks:

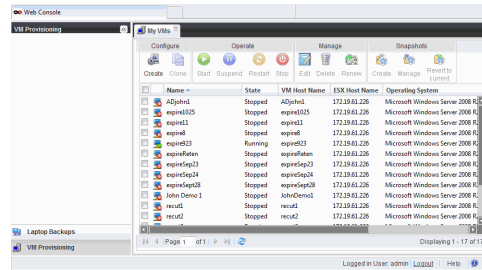
CREATE AND MANAGE VIRTUAL MACHINES

- Options are available to create, customize, and clone virtual machines to suit their needs.
- Administrators can use pre-defined templates and configuration thresholds, such as maximum allowable disk space, memory allocation, and virtual machine expiration to effectively manage resources.

RECOVER ARCHIVED VIRTUAL MACHINES

- When virtual machines are backed up, expired virtual machines are automatically powered off and marked as archived, thereby conserving essential resources.
- The capacity-based licenses for virtual machines are automatically released when virtual machines are archived. As a result, only actively used virtual machines will consume license capacity.
- Users can reactivate and power on virtual machines that have been archived directly from the web console interface. Once reactivated, the virtual machines will be included in Virtual Server iDataAgent backups and count against the capacity-based license.

Click [Next >](#) to setup the VM Life Cycle Management.


[< Previous](#)
[Next >](#)

VM Lifecycle Management - Administrator

[Overview](#)
[Install the Virtual Server iDataAgent](#)
[Install the Web Search Server and Client](#)
[Setup](#)
[User Information](#)
[◀ Previous](#)
[Next ▶](#)

The first step in setting up VM Lifecycle Management is to install the Virtual Server iDataAgent as described below.

PREREQUISITES

The following conditions must be considered before installing Virtual Server iDataAgent:

- VMware vCenter 4.1 is required
- VMware Tools (version 8288) must be installed on templates
- A default password must be set in the template for Linux Virtual machines. For more information, refer to Setting the Password for a Linux VM
- Windows Server 2008, Windows 7, and Linux (RedHat5) are supported

WHERE TO INSTALL

Install the Virtual Server iDataAgent on a proxy computer that can directly communicate with vSphere.

INSTALLATION

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

METHOD 1: INTERACTIVE INSTALL

Use this procedure to install directly on a proxy computer that can communicate with vCenter.

METHOD 2: INSTALL SOFTWARE FROM COMMCELL CONSOLE

Use this procedure to install remotely on a proxy computer that can communicate with vCenter.

BEFORE YOU BEGIN

Download Software Packages


Download the latest software package to perform the install.

System Requirements

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

METHOD 1: INTERACTIVE INSTALL

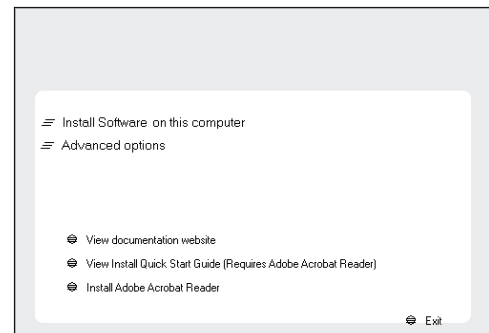
1. Run **Setup.exe** from the **Software Installation Disc** on a proxy computer that can communicate with vCenter.
2. Select the required language.
Click **Next**.



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

NOTES

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



4. Click **Next**.

5. Click **OK**.

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

7. Expand **Client Modules | Backup & Recovery | File System** and select **Virtual Server Agent**.
Click **Next**.


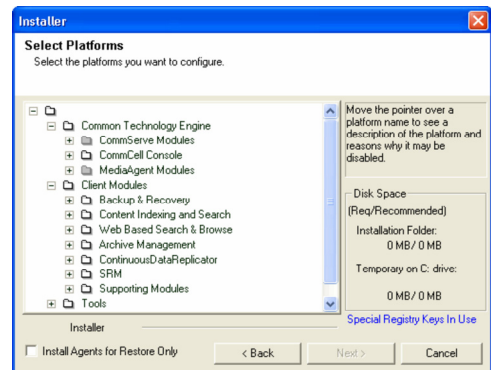

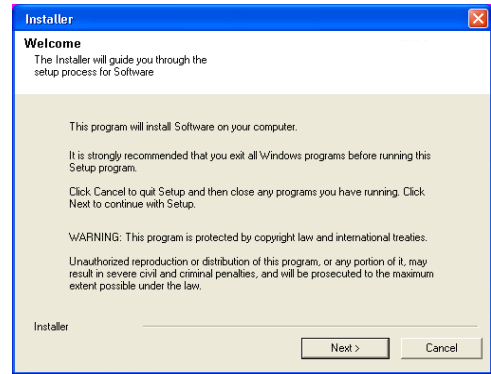
8. Click **YES** to install Microsoft .NET Framework package.

- This prompt is displayed only when Microsoft .NET Framework is not installed.
- Once the Microsoft .NET Framework is installed, the software automatically installs the Microsoft Visual J# 2.0 and Visual C++ redistributable packages.

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

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

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



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

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

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

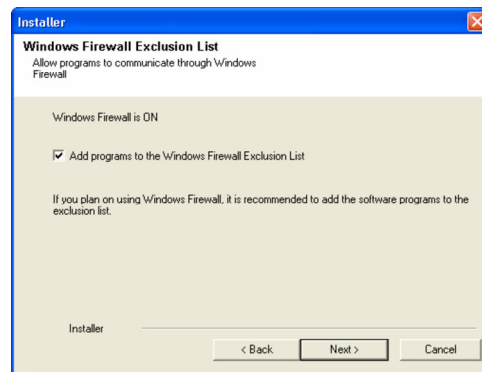



- Click **Next**.

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

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

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

- Click **Next**.



14. Verify the default location for software installation.

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

Click **Next**.

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

/ : * ? " < > | #

It is recommended that you use alphanumeric characters only.

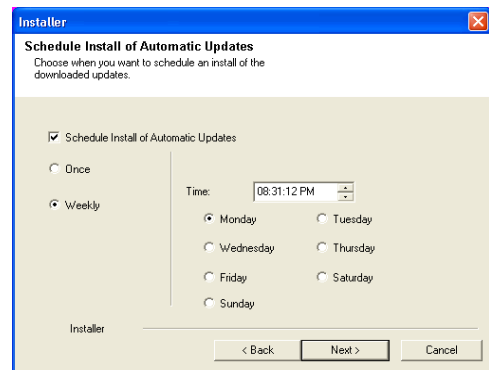
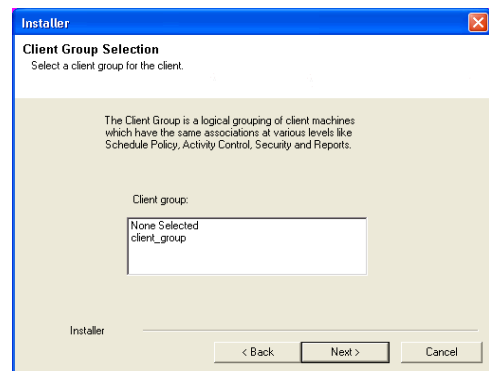
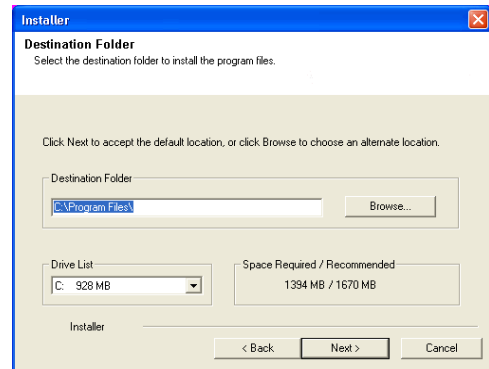
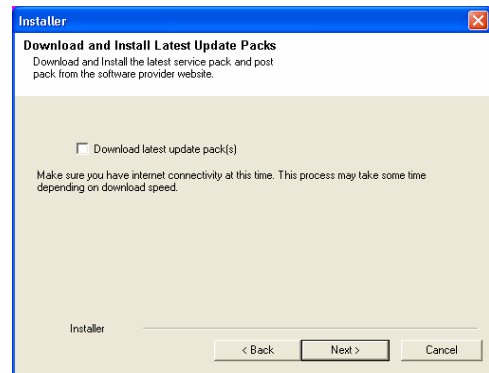
15. Select a Client Group from the list.

Click **Next**.

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

16. Click **Next**.

17. Click **Next**.




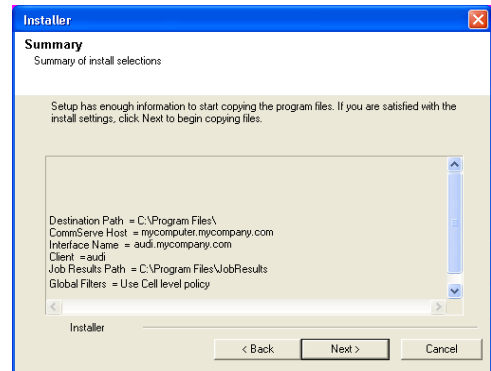



18. Select a **Storage Policy**.
Click **Next**.


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

19. Click **Next**.

20. Click **Next**.

21. Click **Finish**.






THIS COMPLETES THE INSTALLATION OF THE VIRTUAL SERVER IDATAAGENT

Click **Next** to continue.

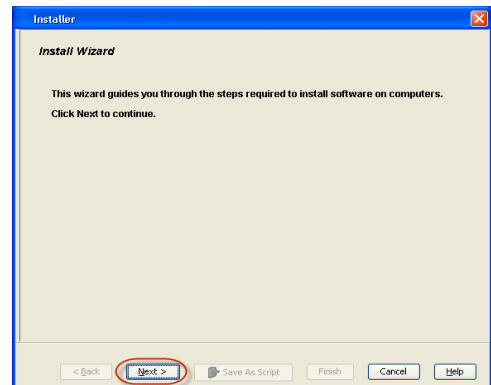


METHOD 2: INSTALL SOFTWARE FROM COMMCELL CONSOLE


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



2. Click **Next**.



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



4. Select **Manually Select Computers**.

Click **Next**.

5. Enter the fully qualified domain name of the proxy computer that can communicate with the ESX Server.

For example: apple.domain.company.com

The Virtual Server iDataAgent will be installed on this proxy computer.

Click **Next**.

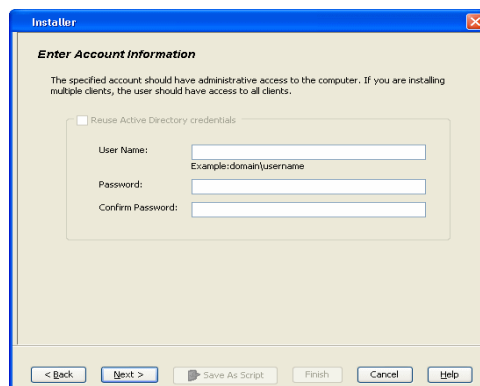
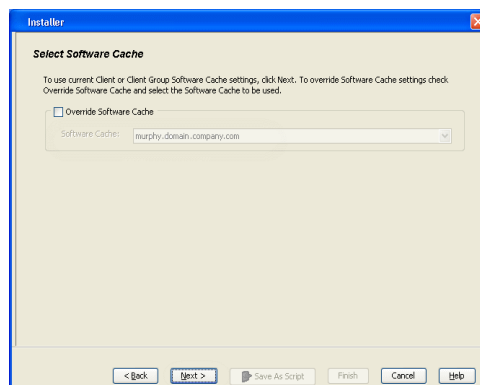
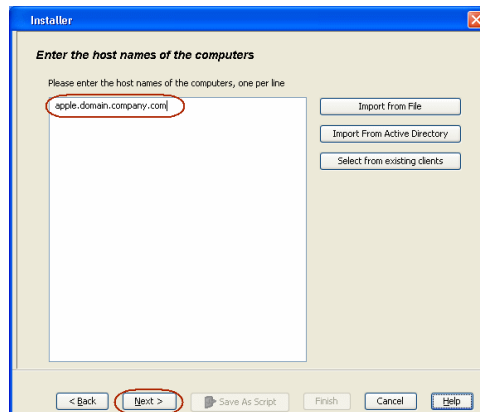
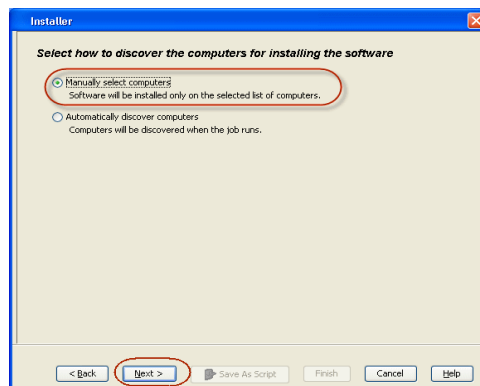
6. Click **Next**.

7. Specify **UserName** and **Password** of the proxy computer.

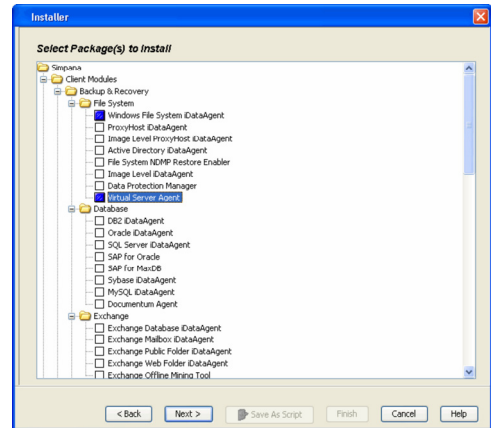
Click **Next**.

8. Select **Virtual Server Agent**.

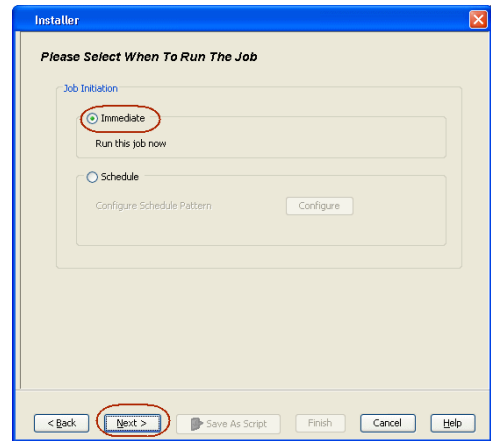
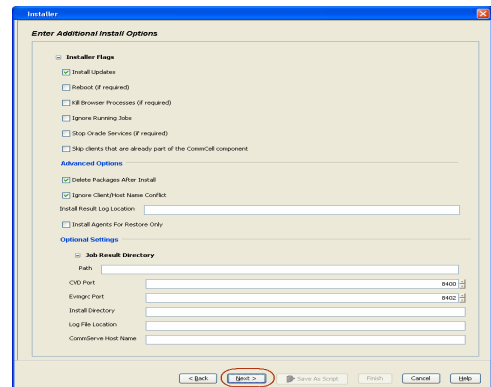
Click **Next**.



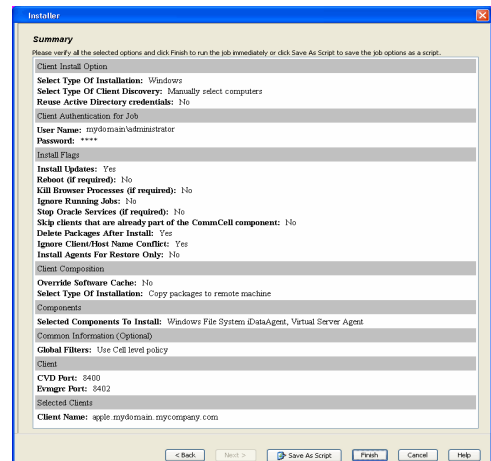
9. Click **Next**.




10. Select **Immediate**.
Click **Next**.



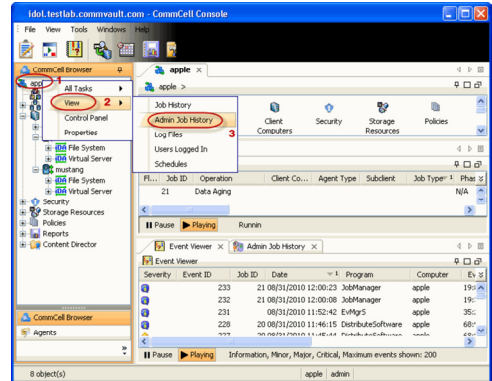
11. Click **Finish**.



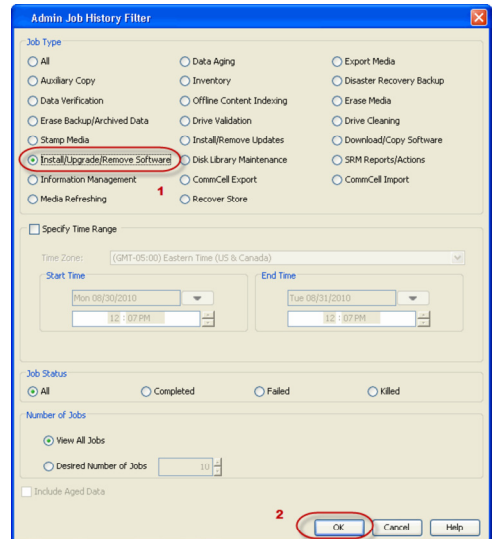
12. You can track the progress of the job from the **Job Controller** or **Event Viewer** window.



13. Once the job is complete, right-click the **CommServe** computer, click **View** and then click **Admin Job History**.




14. Select **Install/Upgrade/Remove Software**.
Click **OK**.



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

- Items that succeeded during the job
- Items that failed during the job
- Details of the job
- Events of the job
- Log files of the job



THIS COMPLETES THE INSTALLATION OF THE VIRTUAL SERVER IDATAAGENT

VM Lifecycle Management - Administrator

[Overview](#)
[Install the Virtual Server iDataAgent](#)
[Install the Web Search Server and Client](#)
[Setup](#)
[User Information](#)
[◀ Previous](#)
[Next ▶](#)

Once the Virtual Server iDataAgent is installed, the Web Search Server and Web Search Client must be installed. These components can be installed on the same computer or on different computers.

WHERE TO INSTALL

Install the Web Search Server and Client on any computer which satisfies the system requirements.

BEFORE YOU BEGIN

Download Software Packages

Download the latest software package to perform the install.

System Requirements

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

INSTALLATION

1. Ensure that the IIS is installed on the CommServe computer before installing the Web Search Server and Web Search Client.
2. Run **Setup.exe** from Software Installation Discs on the **CommServe** computer.
3. Select the required language.
Click **Next**.

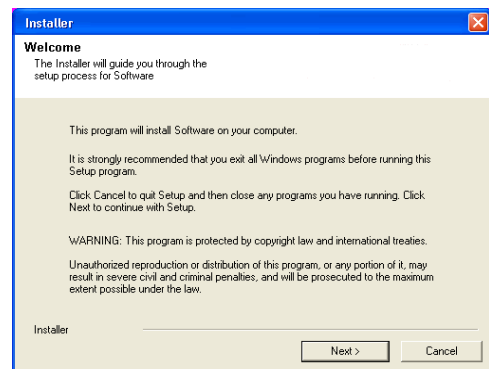


4. Select the option to install software on this computer.

NOTES


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

5. Click **Next**.


6. Click **OK**.



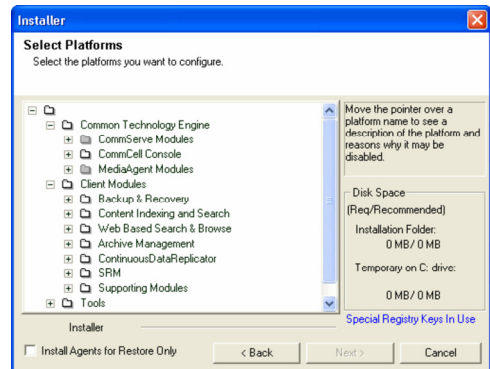
7. Select **I accept the terms in the license agreement.**
Click **Next.**




8. Expand **Common Technology Engine | Web Console Modules** and select **Web Server** and **Web Client.**
Click **Next.**



9. Click **Yes.**



10. Click **Yes.**




11. Specify the SQL Server System Administrator password.

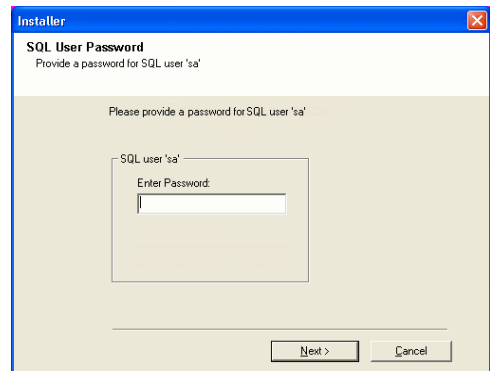
NOTES

- This is the password for the administrator's account created by SQL during the installation.

Click **Next.**



12. Modify **Apache Tomcat Server port number** to **82.**



Click **Next**.

NOTES:

- Ensure that these port numbers are different and are not already used by any other services or application.

13. Click **Next**.

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

Click **Next**.

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

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

15. Enter the Web Search Server database installation path.

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

Click **Next**.





16. Click **Next**.

NOTES

- The **Compliance site**, **End-user site**, and **Admin site** details are populated automatically.

Note down the **End-user** site name.

This is needed later for users to access the VM Provisioning web interface.



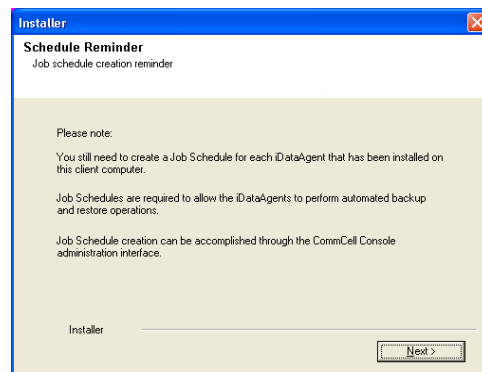
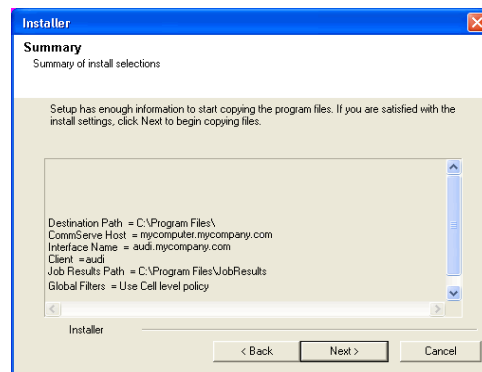


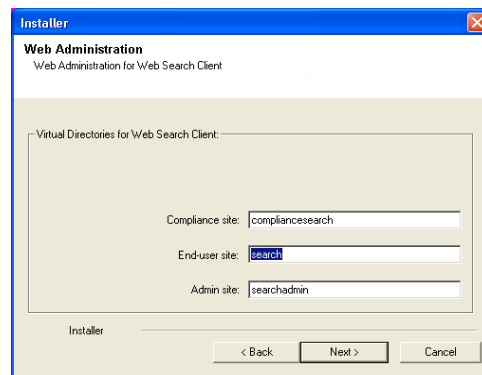
17. Click **Yes**.


18. Click **OK**.

19. Click **Next**.

20. Click **Next**.

21. Click **Reboot** to finish the installation.





VM Lifecycle Management - Administrator

[Overview](#)
[Install the Virtual Server iDataAgent](#)
[Install the Web Search Server and Client](#)
[Setup](#)
[User Information](#)
[< Previous](#)
[Next >](#)

Once the Virtual Server iDataAgent, Web Search Server, and Web Search Client are installed, perform the following configurations:

1. Register Users in the CommCell Console
2. Configure Backup for Virtual Machines
3. Populate Datacenter Information
4. Create a VM Lifecycle Policy
5. Schedule a Backup to Archive Virtual Machines

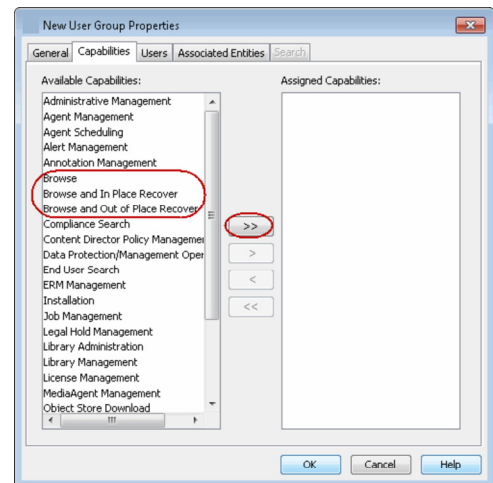
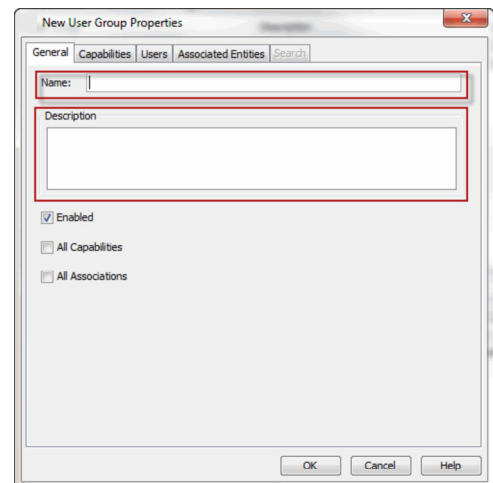

REGISTER USERS IN THE COMMCELL CONSOLE

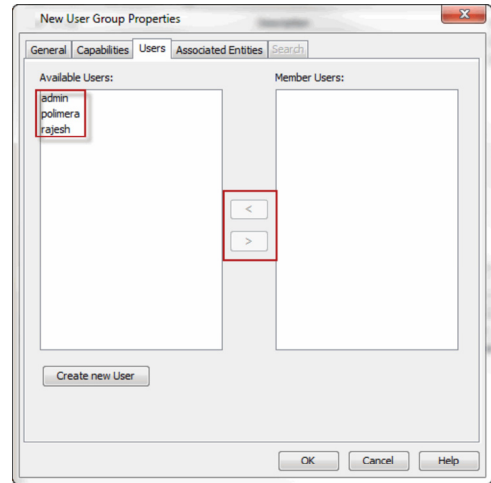
Follow the steps given below to create a group of all the users who will use the web console to create and manage the virtual machines:

1.
 - From the CommCell Browser, click **Security**.
 - Right-click **CommCell User Groups** and click **New User Group**.
2.
 - Type the **Name** of the user group.
The name can have a maximum of 32 characters and should not include trailing spaces.
 - Type a **Description** for the user group.
3. Click the **Capabilities** tab.

Select the following capabilities from the list of **Available Capabilities** and click **>>** to add them to the list of **Assigned Capabilities**:

 - **Browse**
 - **Browse and In Place Recover**
 - **Browse and Out of Place Recover**
4.
 - Click the **Users** tab.
 - Select the appropriate user from the list of **Available Users** and click **>>** to add them to the list of **Member Users**.
 - Click **OK**.

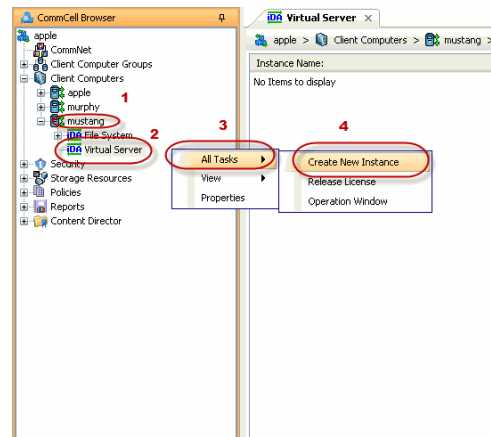




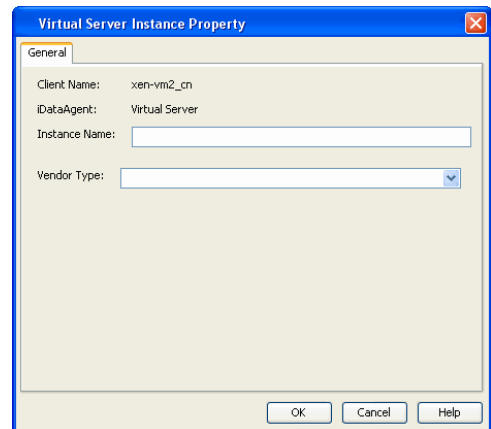
CONFIGURE BACKUP FOR VIRTUAL MACHINES

You must perform regular backups for all the virtual machines created by the users. Follow the steps given below to configure a backupset for such virtual machines:

1.
 - From the CommCell Browser, navigate to **<Client> | Virtual Server**.
 - Right-click the **Virtual Server | All Tasks** and click **Create New Instance**.

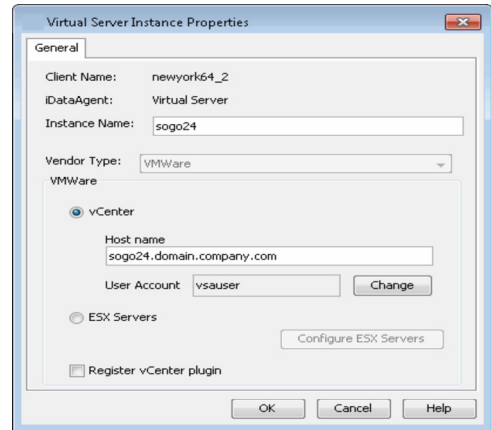


2.
 - Enter the **Instance Name**.
 - Select **VMware** from **Vendor Type** menu.



3.
 - Click **vCenter**.
 - Click **Change**.
 - Enter the username and password associated with the Virtual Center.

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

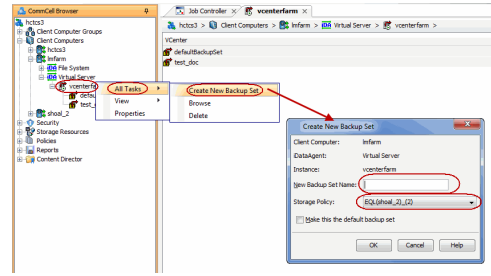
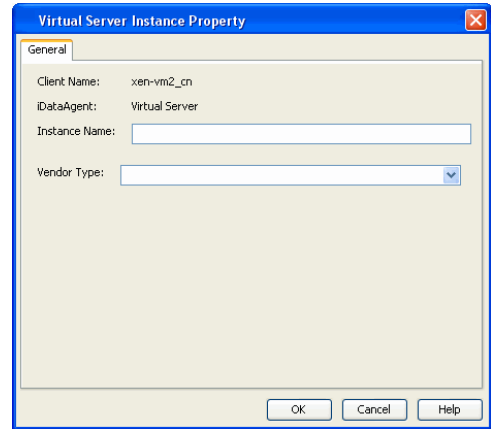


- 5.
- From the CommCell Browser, navigate to **<Client> | Virtual Server**.
 - Right-click the instance and select **All Tasks | Create New Backup Set**.
 - Enter the **New Backup Set Name**.

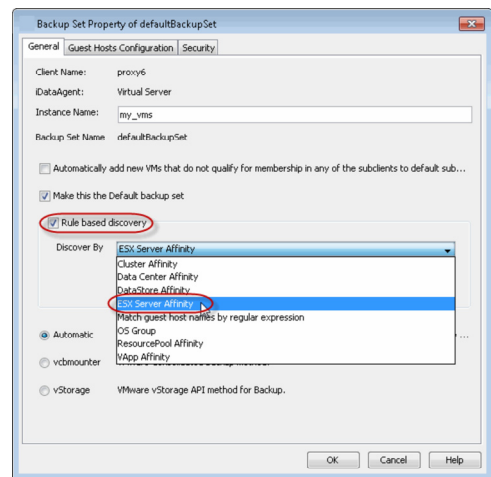
Note down the name of the backup set. You will need it when you will schedule a backup of virtual machines.

This backup set should be dedicated for VM Lifecycle management.

- Select a **Storage Policy**.
It is recommended to select a storage policy which is configured to retain the data for longer time.
- Click **OK**.



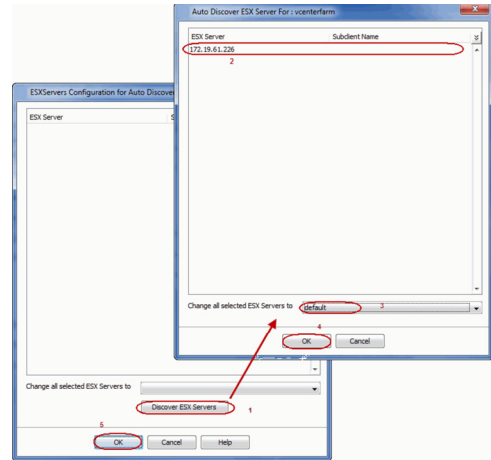
- 6.
- From the CommCell Browser, navigate to **<Client> | Virtual Server | <Instance>**.
 - Right click the **<Backup Set>** and select **Properties**.
 - Select **Rule-Based Discovery**.
 - Click **Discover By** drop-down box and select **ESX Server Affinity**.
 - Click **Configure**.



- 7.
- Click **Discover ESX Servers**.
 - Select an ESX Server in which users are going to create the virtual machines.

Note down the hostname of the ESX server. You will need it when you will create the **VM Lifecycle Policy**.


- Select the **default** subclient from **Change all selected ESX Server** to list.
- Click **OK**.
- Click **OK** on the **ESX Server Configuration for Auto Discover** dialog box.
- Click **OK** on the **Backup Set Properties** dialog box.




POPULATE DATACENTER INFORMATION

Follow the steps give below to collect information from the vCenter’s datacenters, such as templates and virtual machine operating systems. This information will be used later to create a VM Lifecycle Policy.

1.
 - From the CommCell Browser, navigate to **<Client> | Virtual Server**.
 - Right-click **Virtual Server**, select **All Tasks** and click **Refresh Datacenters**.
This process may take several moments to complete.




2. Click **OK**.



CREATE A VM LIFECYCLE POLICY

VM Lifecycle Policy establishes the basic configurations and thresholds users will need to create virtual machines, such as available operating systems, allowable disk space, expiration date of virtual machines. Follow the steps given below to create a VM Lifecycle policy:

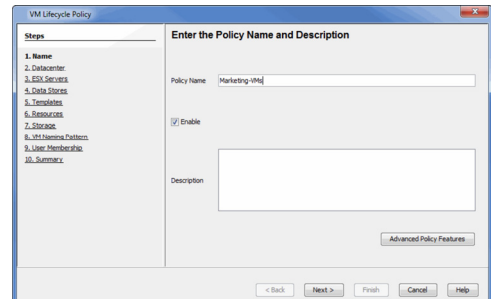
1.
 - From the CommCell Browser, expand **Policies**.
 - Right-click **VM Lifecycle Policies** and click **New Lifecycle Policy**.



2.
 - Enter a **Policy Name**.
 - Enter a **Description** for the policy.

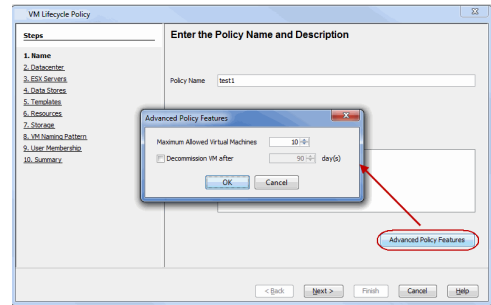
Note down the **Policy** name.

This name should be provided to the users to create virtual machines.

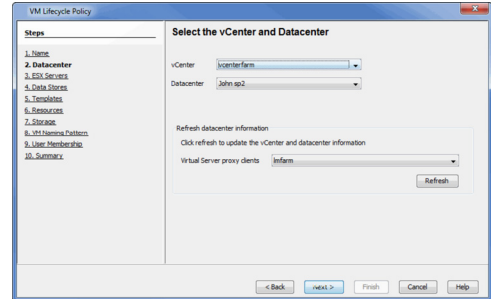


3.
 - Click **Advanced Policy Features**.
 - Enter the maximum number of virtual machines, which can be created using the VM Lifecycle Policy, in the **Maximum Allowed Virtual Machines** box.

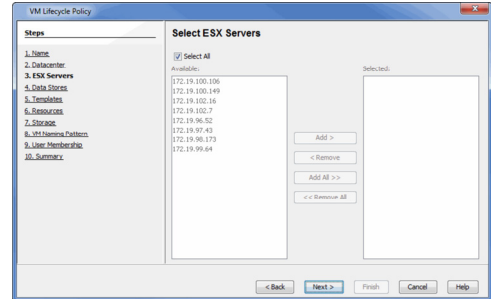
- Select the **Decommission VM after** check box.
 - Enter the number of days after which virtual machines will be decommissioned.
- After the specified time period, the virtual machines will be automatically turned off. When you perform the next backup, all such VMs will be flagged as **Archived**.



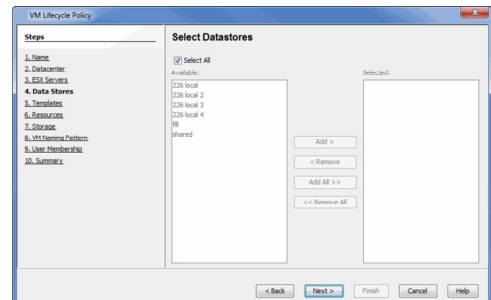
- 4.
- Select a **vCenter**.
 - Select a **Data Center**.
 - Click **Refresh** to update vCenter and Datacenter information.
 - Select a **Virtual Server Proxy Client**.



5. Select the **ESX Server(s)** from the **Available** list and click **Add**.
The users will create the the virtual machines on the selected **ESX Server(s)**.



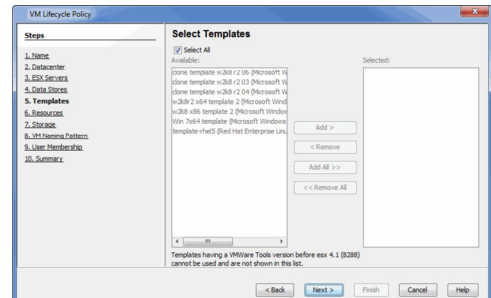
6. Select the **Data Stores** from the **Available** list and click **Add**.
The selected **Data Stores** will be used to create the virtual machines.



7. Select the **Template** from the **Available** list and click **Add**.
The selected **Template** will be used to create the virtual machines.

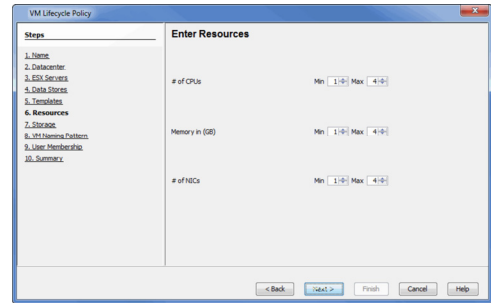
User would need to have one template per Operating System type that is planned to be provisioned.

User would need to have one template per Operating System type that is planned to be provisioned.



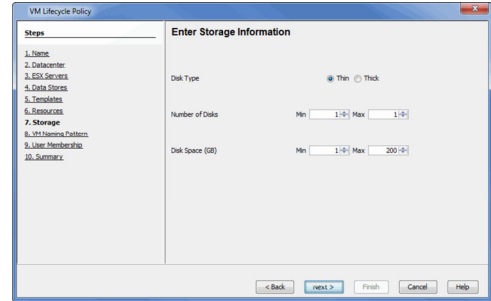
8. Enter the following minimum and maximum values allowed for each virtual machine:
- **# of CPUs**
 - **Memory (in GB)**
 - **# of NICs**

9.
 - Select the **Disk Type**.
The **Disk Type** specifies how the disk space will be allocated to a virtual machine.
 - Select the minimum and maximum **Number of Disks** allowed in a VM.
 - Enter the minimum and maximum amount of **Disk Space** allowed for each disk (in GB).

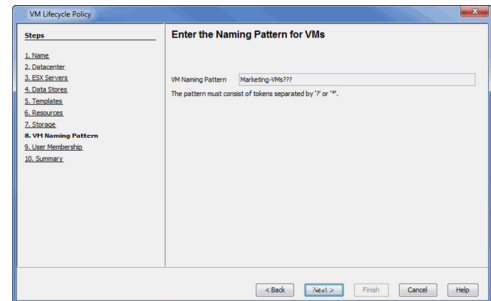


10. Create a **Virtual Machine pattern**. This pattern will be used when naming a new virtual machine.
 - Use '?' to require a single character in a portion of each virtual machine name.
 - Use '*' to require multiple characters in a portion of each virtual machine's name.

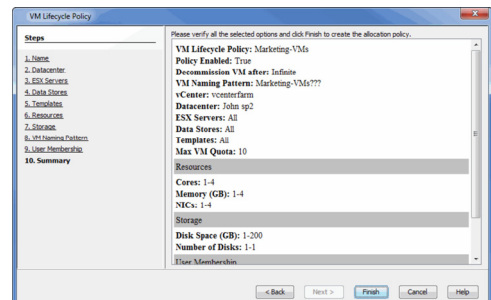
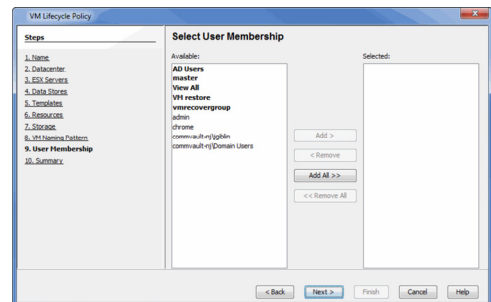
For example, the default value Policy??? requires each virtual machine name to contain three additional characters following the word **Policy** (e.g., Policy123, Policy-VM, etc.)



11. Select the user groups that will have access to VM Lifecycle Management feature. The **Available** list displays all the user groups which you have registered in the CommCell.



12. Click **Finish**.



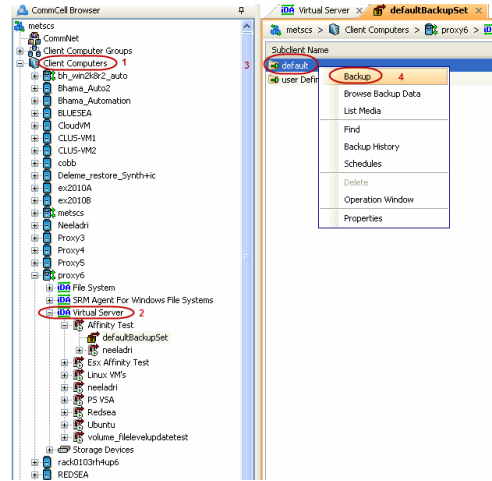
SCHEDULE A BACKUP TO ARCHIVE VIRTUAL MACHINES

Any virtual machines which have reached the expiration date are automatically turned off. When you perform a backup of such virtual machines, they will be flagged as Archived. It is recommended to schedule regular backups for virtual machines. You must schedule at least one backup after the expiration date of the virtual machine. Follow the steps given below to schedule a backup for virtual machines:

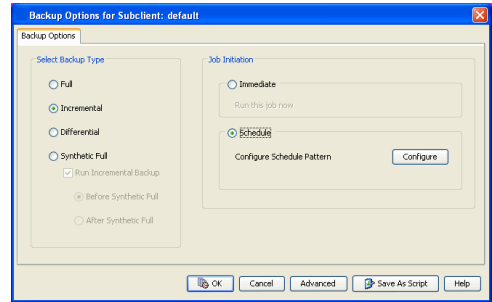
1.
 - From the CommCell Console, navigate to **Client Computers | Virtual Server |**

<Instance> | <Backup Set>.

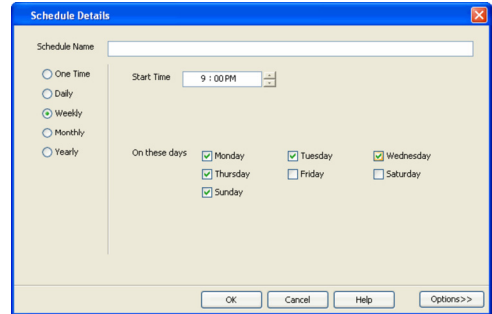
- Right-click the default subclient and click **Backup**.



2.
 - Click **Schedule** to schedule the backup for a specific time.
 - Click **Configure**.



3.
 - Enter the **Schedule Name** and select the appropriate scheduling options. You must schedule at least one backup after the time period specified in the step3 of Creating VM Lifecycle Policy.
 - Click **OK**.



CONGRATULATIONS - YOU HAVE SUCCESSFULLY CONFIGURED THE VM LIFECYCLE MANAGEMENT.

Before users can begin using this feature, provide them with the information described in User Information.



VM Lifecycle Management - Administrator

Overview	Install the Virtual Server iDataAgent	Install the Web Search Server and Client	Setup	User Information
----------	---------------------------------------	--	-------	------------------

[◀ Previous](#) [Next ▶](#)

DISTRIBUTE THE FOLLOWING INFORMATION TO VM LIFECYCLE MANAGEMENT USERS:

1. Web Console URL:

This URL is configured in Step 16 of the Web Search Server installation.
2. User Name:

This user account is the Active Directory user account for the user.
3. Password:

This password is associated with the Active Directory user account for the user.
4. Login Domain:

This domain is the domain associated with the Active Directory user account for the user.
5. Name of the VM Lifecycle Policy

This name is specified in Step 2 of the Create a VM Lifecycle Policy.
6. Naming Pattern for naming the virtual machines

This pattern is specified in Step 10 of the Create a VM Lifecycle Policy.

[◀ Previous](#) [Next ▶](#)

VM Lifecycle Management - Administrator (Advanced)

SETTING THE PASSWORD FOR A LINUX VM

For creating a Linux VM, you must set the password of the machine as `cvx1_<hostname>`. Follow the steps given below to set the password for a Linux VM:

1. Use the following script to create the `CronSetPwd.sh` file:

```
echo "hostnametemp=`/bin/hostname`" >> /var/tmp/cronSetPwd.sh
echo "echo hostname is \${hostnametemp} >> /var/tmp/vmlifecycle.log" >> /var/tmp/cronSetPwd.sh
echo "echo PATH is \${PATH} >> /var/tmp/vmlifecycle.log" >> /var/tmp/cronSetPwd.sh
echo "echo `root:cvX1__\${hostnametemp}` | /usr/sbin/chpasswd" >> /var/tmp/cronSetPwd.sh
echo "echo error code is \${?} >> /var/tmp/vmlifecycle.log" >> /var/tmp/cronSetPwd.sh
echo "/usr/bin/crontab -r" >> /var/tmp/cronSetPwd.sh
```

2. Open the command prompt and enter the following command to assign the `Execute` permissions for the `CronSetPwd.sh` file:

```
chmod 744 /var/tmp/cronSetPwd.sh
```

3. Enter the following command to execute the `CronSetPwd.sh` file:

```
echo "@reboot /var/tmp/cronSetPwd.sh" > /var/tmp/cronfile.txt
crontab /var/tmp/cronfile.txt
rm -f /var/tmp/cronfile.txt
```

The password of any Linux virtual machine, created using the VM Lifecycle Policy, will be set as `cvx1_<hostname>`. Here the `<hostname>` is the hostname of the virtual machine. The user can change this password later. If the system is unable to set the password specified by the user while creating the Linux virtual machine, the user can use the password `cvx1_<hostname>` to access the newly created Linux VM.

VM Lifecycle Management - End-User (Getting Started)

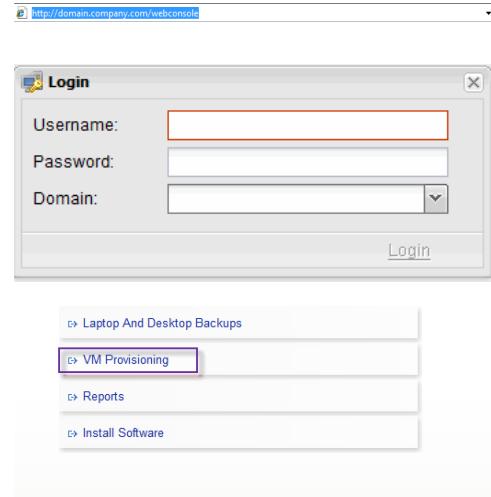
Getting Started | **Advanced**

◀ Previous Next ▶

LOG ON TO THE WEB CONSOLE

- Obtain the following information from your administrator:
 - Web Alias** name
 - Policy** name
 - Naming Pattern** for naming virtual machines

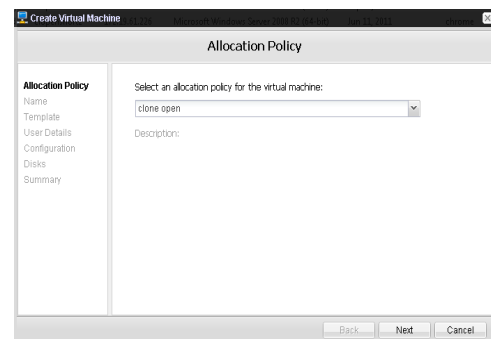
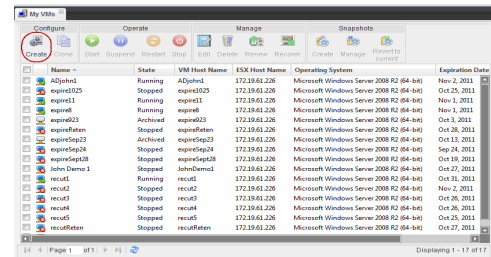
You will need this information as you proceed.
- Open a supported Web Browser and type the Web Console URL in the address line. For example, `http://amber.domain.company.com/<web_alias_name>`
- Type the login credentials and click **Login**.



- Click **VM Provisioning**.

CREATE A VIRTUAL MACHINE

- Click **Create** from the **Configure** menu.
- Select an allocation policy provided to you by your administrator.
 - Click **Next**.
- Enter a name for the virtual machine. The name must follow the pattern displayed. For example, if the pattern displayed is **Marketing?VMs***, then your virtual machine name must consist of the following:
`Marketing<single_character>VMs<any_additional_characters>`
 - Enter a **Description**.
 - Enter an **Expiration Date**.
 The virtual machine will not be available past the expiration date.



- Click **Next**.

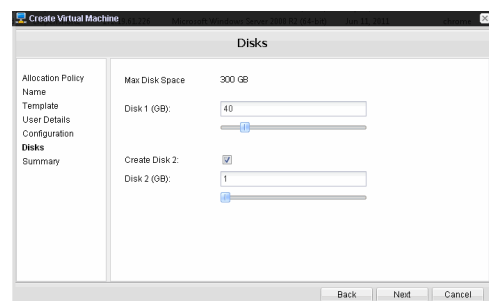
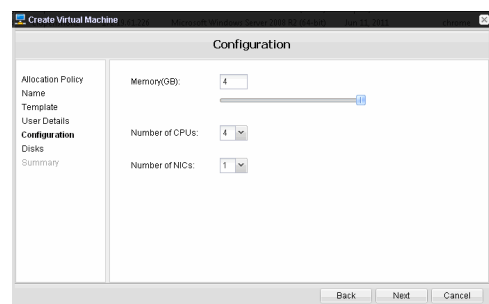
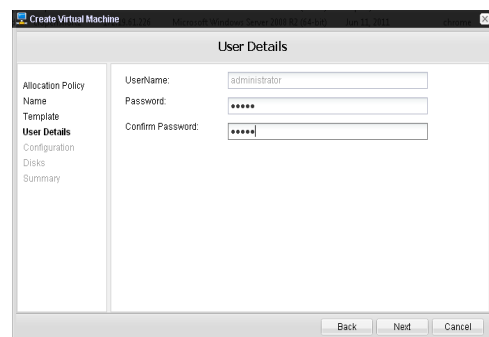
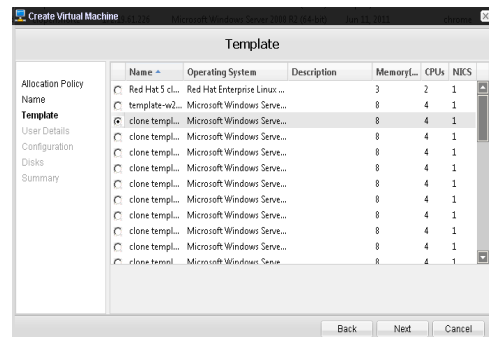
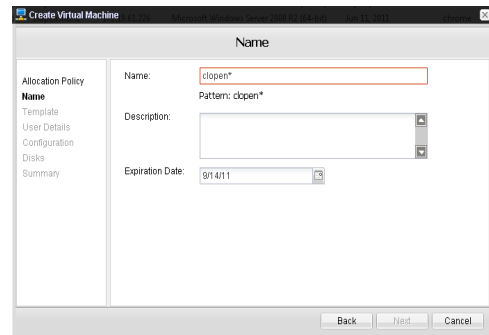
- 4.
- Select a **Template**.
 - Click **Next**.

- 5.
- Enter a password for the virtual machine.
 - Confirm the password.

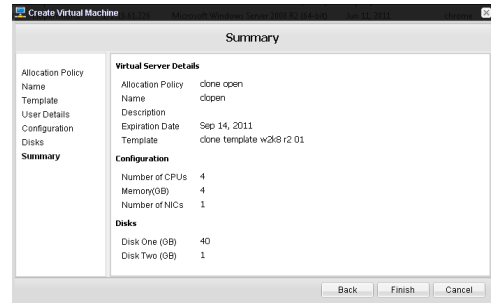
- 6.
- Enter the amount of **Memory** (in GB) for the virtual machine.
Drag the customization bar to refine the memory amount in MB.
 - Select the **Number of CPUs**.
 - Select the **Number of NICs**.
 - Click **Next**.

- 7.
- Enter the amount of disk space for the virtual machine.
Drag the customization bar to refine the memory amount in MB.
 - Click **Next**.

8. Click **Finish**.
- Once the virtual machine is created, it will appear in the **My VMs** list. The creation process may take several moments to complete.

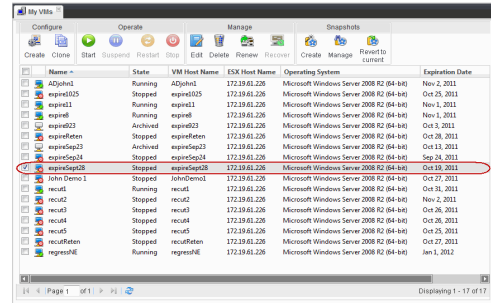


An email will be automatically sent to you confirming the creation of your virtual machine.

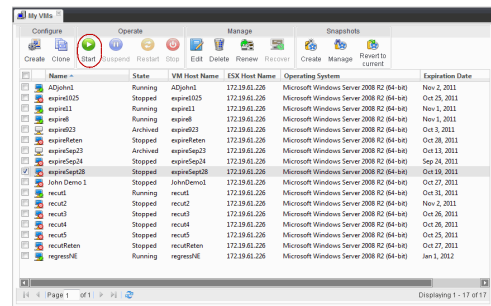


START A VIRTUAL MACHINE

1. Select the virtual machine you want to start.



2.
 - Click the **Start** button in the **Operate** menu.
 - You can view the progress of this operation in the **State** column of the **My VMs** list.



CONGRATULATIONS - YOU HAVE SUCCESSFULLY CREATED AND STARTED YOUR FIRST VIRTUAL MACHINE.

If you want to further explore this feature's capabilities, read the **Advanced** section of this documentation by clicking **Next**.



VM Lifecycle Management - End-User (Advanced)

Getting Started | **Advanced**

◀ Previous | Next ▶

TABLE OF CONTENTS

Virtual Machine Operations

Modify a Virtual Machine

- Rename a Virtual Machine
- Edit Memory, CPU, or NIC Configurations
- Add or Edit Available Disk Space

Delete a Virtual Machine

Renew a Virtual Machine

Recover an Archived Virtual Machine

Manage Snapshots

- Create a Snapshot
- Revert to a Snapshot
- Revert from a Snapshot


Clone a Virtual Machine

VIRTUAL MACHINE OPERATIONS

The **Operate** menu provides the facility to perform basic functions for your virtual machine.

- Click **Start** to turn a virtual machine on.
- Click **Suspend** to temporarily disable use of the virtual machine.
- Click **Restart** to restart the virtual machine.
- Click **Stop** to turn a virtual machine off.

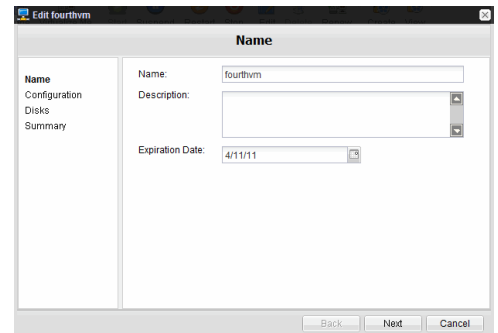
Suspending, restarting, or stopping a virtual machine will disrupt access to the virtual machine for other users. Ensure no other users are currently using the virtual machine prior to performing these functions.



MODIFY A VIRTUAL MACHINE

RENAME A VIRTUAL MACHINE

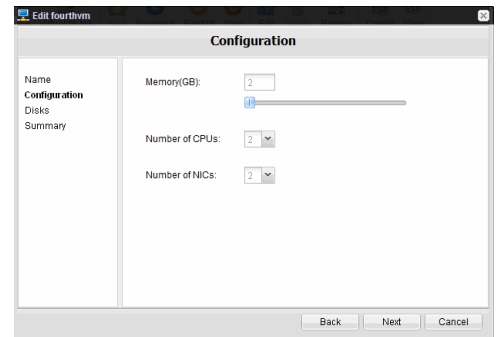
- Select the virtual machine you want to modify.
- Click **Edit** from the **Manage** menu.
- Navigate to the **Name** screen.
- Enter a name for the virtual machine. The name must follow the pattern displayed.
For example, if the pattern displayed is **Marketing?VMs***, then your virtual machine name must consist of the following:
`Marketing<single_character>VMs<any_additional_characters>`
- Enter a new **Description**.
- Enter a new **Expiration Date**.
The virtual machine will not be available past the expiration date.
- Navigate to the **Summary** screen and click **Finish**.



EDIT MEMORY, CPU, OR NIC CONFIGURATIONS

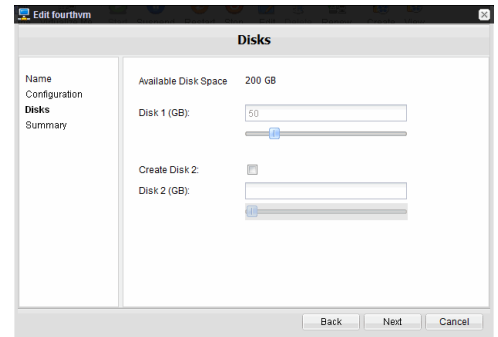
- Select the virtual machine you want to modify.
- Click **Edit** from the **Manage** menu.
- Navigate to the **Configuration** screen.
- Enter the amount of **Memory** (in GB) for the virtual machine.
Drag the customization bar to refine the memory amount in MB.
- Select the **Number of CPUs**.

6. Select the **Number of NICs**.
7. Navigate to the **Summary** screen and click **Finish**.



ADD OR EDIT AVAILABLE DISK SPACE


1. Select the virtual machine you want to modify.
2. Click **Edit** from the **Manage** menu.
3. Navigate to the **Disks** screen.
4. Enter the amount of disk space for the virtual machine.
Drag the customization bar to refine the memory amount in MB.
5. Navigate to the **Summary** screen and click **Finish**.



DELETE A VIRTUAL MACHINE


If you delete a virtual machine, it will no longer be available for use and you can no longer access any data or applications residing on it.

1. Select the virtual machine you want to delete.
2. Click **Delete** from the **Manage** menu.
3. Click **OK**.



RENEW A VIRTUAL MACHINE

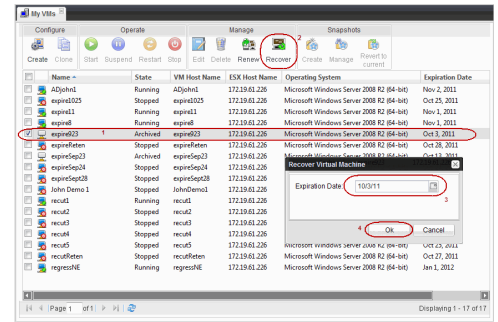
1. Select the virtual machine you want to modify.
2. Click **Renew** from the **Manage** menu.
3. Enter a new **Expiration Date**.
The virtual machine will not be available past the expiration date.
4. Navigate to the **Summary** screen and click **Finish**.



RECOVER AN ARCHIVED VIRTUAL MACHINE

After the expiration date a virtual machine is turned off automatically and its **State** is displayed as **Archived**. Follow the steps given below to recover an **Archived** virtual machine:

1. Select the archived virtual machine.
2. Click **Recover** from the **Manage** menu.
3. Enter a new **Expiration Date**.
By default an expiration date is set after the number of days specified in the **allocation policy**.
4. Click **OK**.

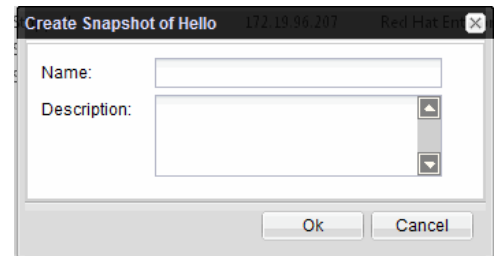


MANAGE SNAPSHOTS

The **Snapshot** feature allows you to create an exact image of a virtual machine as it exists at the time the snapshot is taken. If the virtual machine encounters an irreversible problem (i.e., important files are permanently deleted, the operating system becomes corrupted, etc.), you can revert a virtual machine to its snapshot, thereby bringing the virtual machine back to the state it was in when the snapshot was taken. You can also revert from a snapshot if you no longer wish to use the virtual machine in its snapshot state.

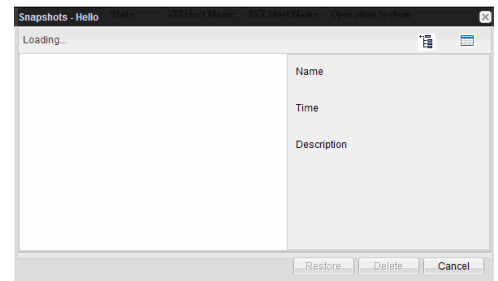
CREATE A SNAPSHOT

1. Select the virtual machine for which you want to create a snapshot.
2. Click **Create** from the **Snapshots** menu.
3. Enter a **Name** for the snapshot.
4. Enter a **Description** of the snapshot (such as a date description of files existing on the virtual machine at the point the snapshot is taken).
5. Click **OK**.



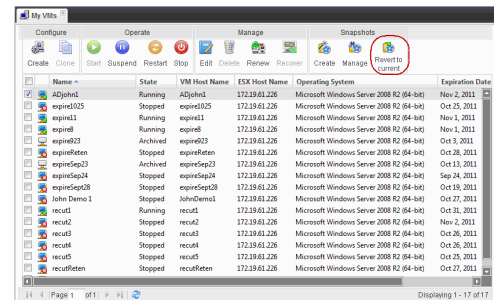
REVERT TO A SNAPSHOT

1. Select the virtual machine for which you want to revert to a snapshot.
2. Click **View** from the **Snapshots** menu.
3. Navigate to the desired snapshot.
4. Click **Revert**.



REVERT FROM A SNAPSHOT

1. Select the virtual machine for which you want to revert from a snapshot.
2. Click **Revert to Current**.



CLONE A VIRTUAL MACHINE

The **Clone** option creates a copy of an existing virtual machine. The new virtual machine retains the same characteristics of the virtual machine from which it was cloned (such as the operating system). This option may be especially useful if you want to create a virtual machine from a template that is no longer available.

1. Shut down the virtual machine you want to clone.

2. Select the virtual machine you want to clone.
3. Enter the following details:
 - **Name** of the virtual machine.
 - **Username** of the virtual machine owner.
 - **Password** of the virtual machine owner.
 - **Description** of what the virtual machine will be used for.
 - **Expiration Date** after which the virtual machine will no longer be available.

The screenshot shows a dialog box titled "Clone clopenshort12". It contains the following fields and values:

- Name: clopen
- Pattern: clopen*
- UserName: administrator
- Password: [masked with dots]
- Confirm Password: [masked with dots]
- Description: Clone of my VM
- Expiration Date: 6/17/11

Buttons: OK, Cancel, Previous, Next

Virtualizing a Windows Computer - Getting Started

Getting Started | **Advanced** | Best Practices | FAQ | Troubleshooting

◀ Previous Next ▶

You can convert a client computer into a virtual machine (VM). Perform the following to convert a client computer into a VM:

1. Verify the Prerequisites
2. Setup 1-Touch Server
3. Setup the Client
4. Create the Boot CD Image
5. Setup the ESX Server
6. Create an Instance in Virtual Server iDataAgent
7. Virtualize the Client

VERIFY THE PREREQUISITES

WHAT CAN BE VIRTUALIZED	WHAT CANNOT BE VIRTUALIZED
Any client that satisfies the System Requirements for 1-Touch client.	Any client that has Windows XP operating system. Any client which has firewall setup between the client computer and the CommServe. Any client which has Windows operating system installed on a disk other than the first disk. Any client which has Windows operating system installed on a drive other than the C drive. Any client computer that can boot from SAN or SAN attached disks.

INSTALL AND SETUP VIRTUAL SERVER iDATAAGENT

Ensure that the Virtual Server iDataAgent is installed on atleast one client computer in the CommCell. For more information, refer to Getting Started - VMware Deployment.

SETUP 1-TOUCH SERVER

Follow the steps given below to setup a 1-Touch server:

1. Install the 1-Touch Server software using the Calypso Installer Version released with Service Pack 4 (or higher). For more information, refer to Deployment - 1-Touch for Windows.
 - If the 1-Touch server has an older version of the 1-Touch server software, perform the following:
 - Uninstall the 1-Touch server software using the **Add/Remove Programs** utility in the Windows **Control Panel**.
 - Re-install the 1-Touch Server software using the Calypso Installer Version released with Service Pack 4 (or higher).
 - 1-Touch Server software is available on the Disc 2. For more information, refer to Software Installation Discs.
2. Configure remote server cache on the 1-Touch server. For more information, refer Configure a Remote Cache for 1-Touch Server.
3. Download all the software updates for 1-Touch Server software, Windows File System iDataAgent and WinPE. For more information, refer Download and Install Updates on 1-Touch Server.
4. Apply Updates to Live CD image. For more information, refer Apply updates in 1-Touch Server Cache to 1-Touch Live CD image.

SET UP THE CLIENT

INSTALL WINDOWS FILE SYSTEM iDATAAGENT

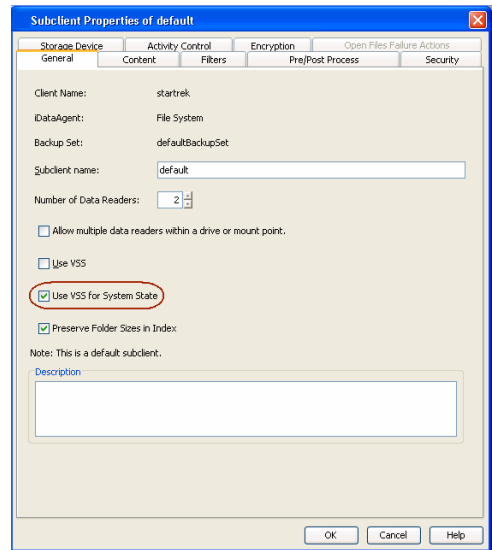
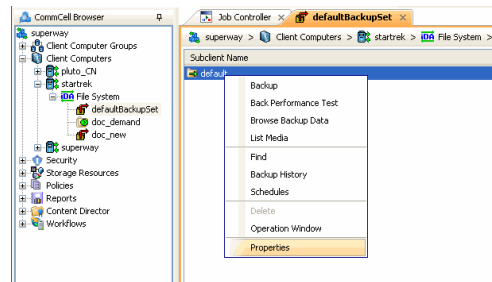
Install the Windows File System iDataAgent using the Calypso Installer Version released with Service Pack 4 (or higher). For more information, refer to Getting Started - Windows File System Deployment.

ENABLE SYSTEM STATE BACKUP

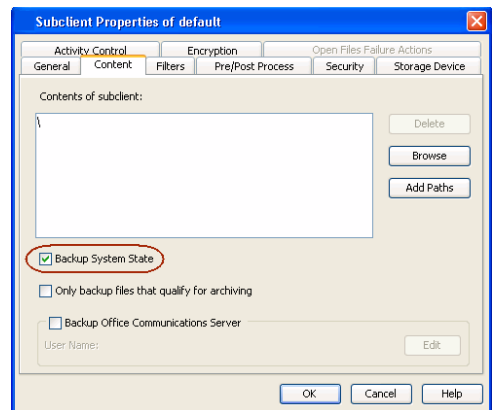
1.
 - From the CommCell Console, navigate to <Client> | **File System iDataAgent** | **defaultBackupSet**.
 - Right-click the default subclient, and click **Properties**.

You must enable the system state backup for the default subclient in the default backupset.

2.
 - Click the **General** tab.
 - Ensure that **Use VSS for System State** check box is selected.




3.
 - Click the **Content** tab.
 - Ensure that the **Backup System State** check box is selected.



PERFORM A BACKUP

If you are performing backup of a Windows 7 client, ensure that the Administrator account is enabled. You cannot perform the virtualization of the client if administrator account is not enabled while performing the backup.

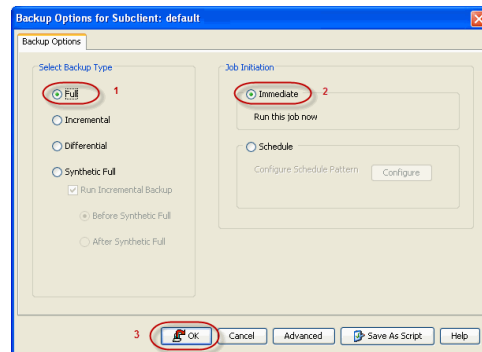
1.
 - From the CommCell Console, navigate to **<Client> | File System iDataAgent | defaultBackupSet**.
 - Right-click the default subclient, and click **Backup**.



2.
 - Select **Full**.
 - Select **Immediate**.

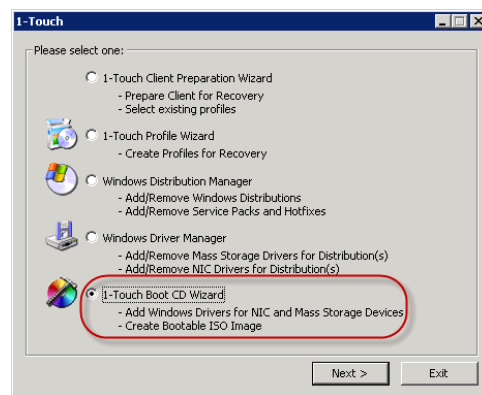
- Click **OK**.

Before virtualizing a client computer, ensure that you have successfully completed at least one full backup of the default subclient.

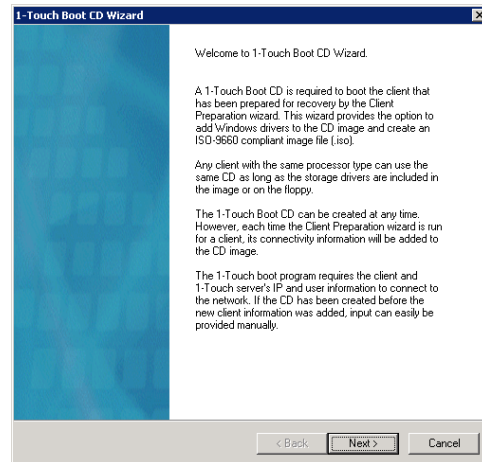


CREATE THE BOOT CD IMAGE

1. Login to computer which has 1-Touch Server software installed on it.
2. Click **Start | Programs | Bull Calypso | Calypso | 1-Touch**.
3. Click **1-Touch Boot CD Wizard**.

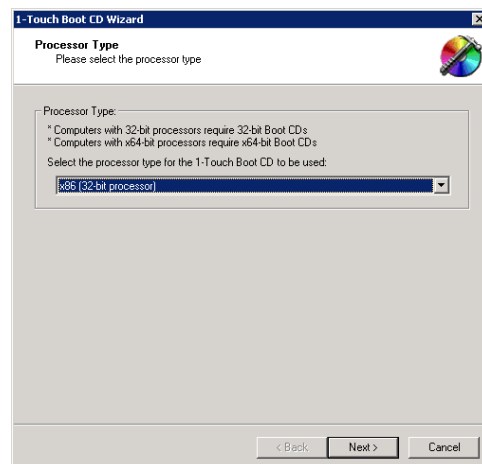


4. Click **Next**.

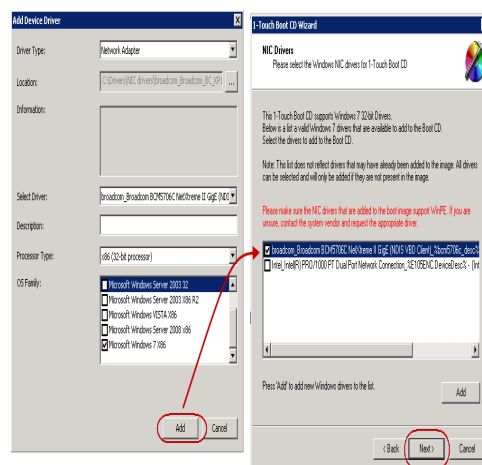
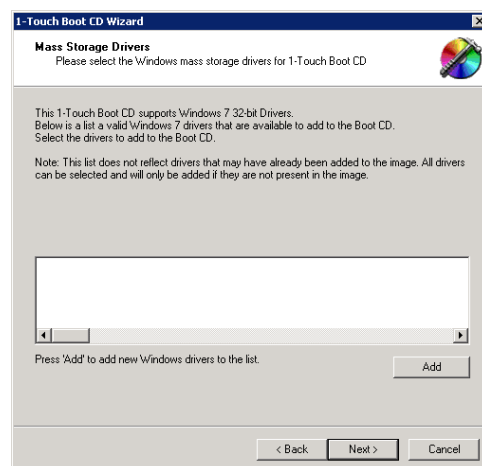


5. Select the processor type from the list.
 The processor type must match with the processor of the computer which you want to convert into a VM.
 For example: If you want to virtualize a windows computer which has x64 version of operating system, you must select x64 as the processor type.

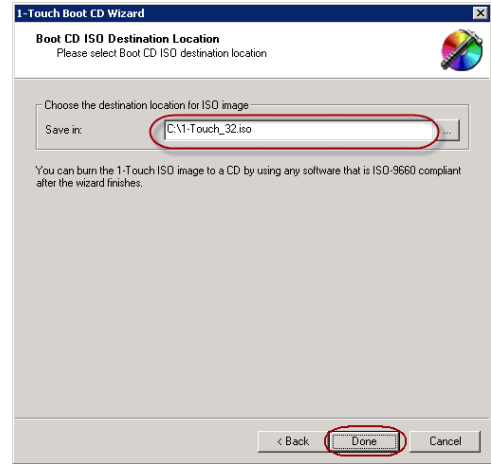
6. Click **Next**.



7. Click **Next**.



- 8.
- Enter the location to save the Boot CD image as an `.iso` file.
 - Click **Done**.
 - Verify that the Boot CD image is created at the specified location.



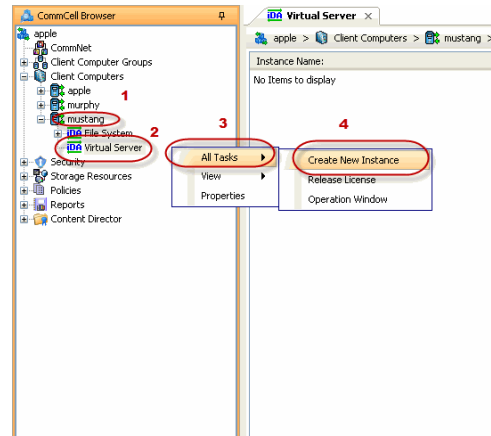
SETUP THE ESX SERVER

1. Verify the connectivity between the CommServe and the vCenter in which you want to create the VM
2. Copy the boot CD Image to any datastore on the ESX server in which you want to create the VM.

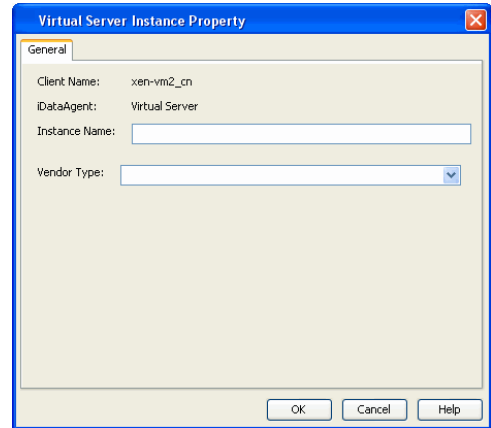
CREATE AN INSTANCE IN THE VIRTUAL SERVER IDATAAGENT

Follow the steps given below to create an instance for a vCenter in which you want to create the VM:

1.
 - From the CommCell Browser, navigate to **<Client> | Virtual Server.**
 - Right-click the **Virtual Server | All Tasks** and click **Create New Instance.**



2.
 - Enter the **Instance Name.**
 - Select **VMware** from **Vendor Type** menu.

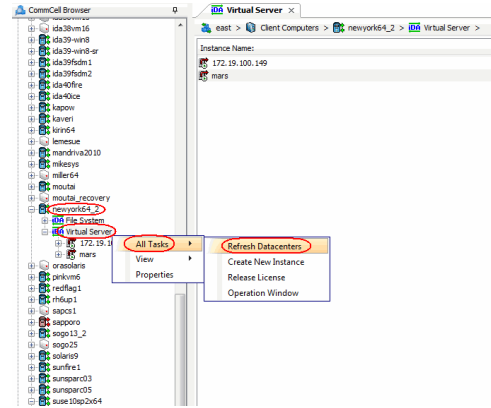



3.
 - Click **Virtual Center.**
 - Enter the host name of the vCenter.
 - Click **Configure Password.**
 - Enter the username and password associated with the Virtual Center.
 - Click **OK** to save the instance.

4.
 - From the CommCell Browser, navigate to **Client Computers** | **<Client>**.
 - Right-click the **Virtual Server** and select **All Tasks** | **Refresh Datacenters**.

It could take several minutes for a newly created instance to appear in the list of instances.

Once you create the instance, the information about the ESX sever, vCenter and the Data Store is pre-populated in the **Virtualize Me** dialog box.



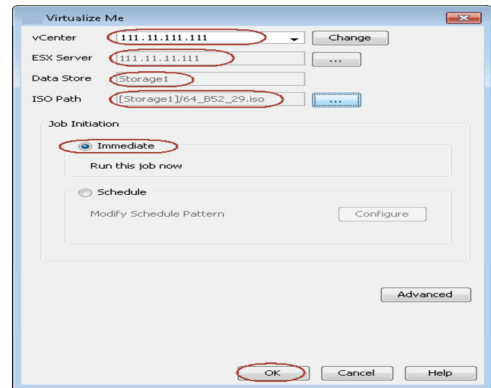
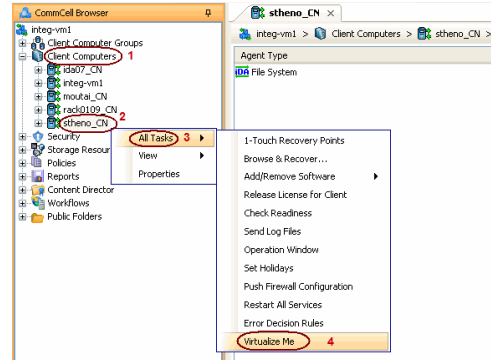
VIRTUALIZE THE CLIENT

Follow the steps given below to convert a client computer into a VM:

1. Turn off the client computer.
2.
 - From the CommCell Console, navigate to **Client Computers**.
 - Right-click the **<Client>** and select **All Tasks** | **Virtualize Me**.
3.
 - Select a vCenter from the **vCenter** list.
 - Click the **...** button and browse to select an **ESX Server** and a **Data Store**.
 - Click the **...** button for the **ISO Path** box and browse to the location of the Boot CD image on the ESX Server.

If you want to change the hardware configuration, click **Advanced**. For more information, refer to Advanced section of this documentation.

 - Click **OK**.




The virtual machine will be created in the specified ESX server.

4. You can track the progress of the Virtualization job from the **Job Controller** window of the CommCell Console. The following three jobs are launched during the virtualization operation:

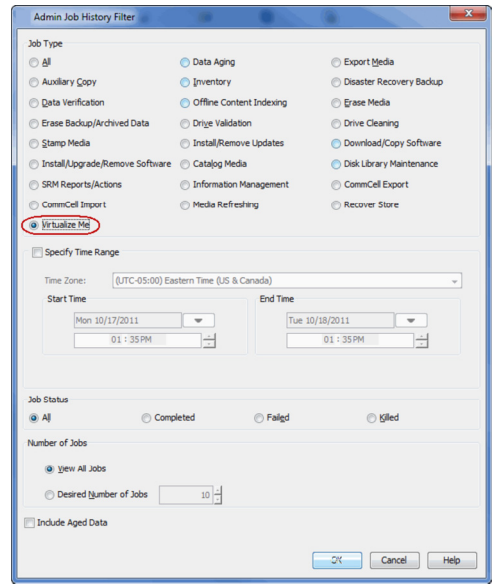

- **Virtualize Machine**
- **1-Touch Configuration Restore**
- **Full System Restore**

You can not suspend or kill the **Virtualize Machine** job when the other two jobs are running.

5. Once the job is complete, view the job details from the **Admin Job History**. From the CommCell Browser, right-click **<CommServe>** and select **View | Admin Job History**.



- 6.
- Select **Virtualize Me** under **Job Type**.
 - Click **OK**.



7. You can view the following options about the job by right-clicking the job:

- Details of the job
- Events of the job
- Log files of the job
- Resubmit the job to run it again

Job Type	Job ID	Status	Start Time	End Time	User Name	Error Code	Storage Policy
Virtualize Machine	19790	Completed	10/10/2011 01:36:03	10/10/2011 15:36:36	admin		
Virtualize Machine	19144	Completed	10/10/2011 01:35:07	10/10/2011 14:34:00	admin		
Virtualize Machine	19143	Failed	10/10/2011 01:31:54	10/10/2011 01:35:12	admin	[19381]	

CONGRATULATIONS - YOU HAVE SUCCESSFULLY CONVERTED A CLIENT COMPUTER INTO A VIRTUAL MACHINE.

If you want to further explore additional features read the Advanced section of this documentation.



Virtualizing a Windows Computer - Advanced Options

Getting Started | **Advanced** | Best Practices | FAQ | Troubleshooting

◀ Previous Next ▶

TABLE OF CONTENTS

Creating a Virtual Machine Using a Recovery Point

Comparing the Client Configurations Associated with Various Recovery Points
Creating a Virtual Machine

Setting an Expiration Date for the Virtual Machine

Modifying the Network Configuration for a Virtual Machine

Configuring WINS or DNS for a Windows Computer
Assigning a Network Label

Modifying the Hardware Configuration for a Virtual Machine

Configuring Thin Provisioning for a Virtual Machine

Creating a Virtual Machine Using a Backup Copy

Scheduling the Virtualization Operation

Related Reports

Job Summary Report

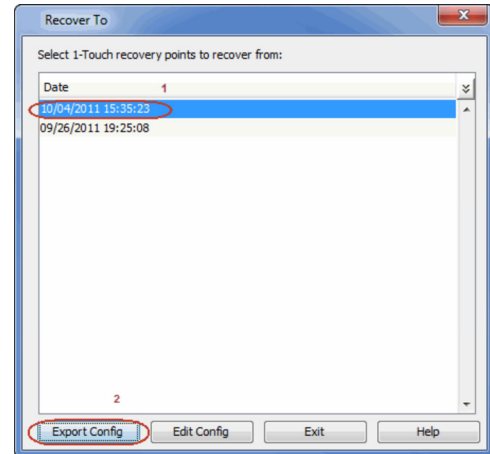
CREATING A VIRTUAL MACHINE USING A RECOVERY POINT

A recovery point is generated whenever you change the hardware or network configuration of the client and perform a backup. You can use recovery points to create a virtual machine with an old hardware or network configuration. By default, the most recent recovery point is used to create the virtual machine.

COMPARING THE CLIENT CONFIGURATIONS ASSOCIATED WITH VARIOUS RECOVERY POINTS

Before using a recovery point, you can compare the client configurations associated with various recovery points and then choose a recovery point for creating the VM. You can generate an xml file for each recovery point. The xml file contains the client configuration details. Follow the steps given below to generate the xml file for a recovery point:

1. From the CommCell Console, navigate to **Client Computers**.
2. Right-click the **<Client>** and select **All Tasks | 1-Touch Recovery Point**.
3. Select a recovery point and click **Export Config**.
4. Specify the location to save the xml file and click **Save**.
By default, the the xml file is saved in the *My documents* folder on the computer that is running the CommCell Console.
5. Open the xml file using any text editor and verify the client configuration.

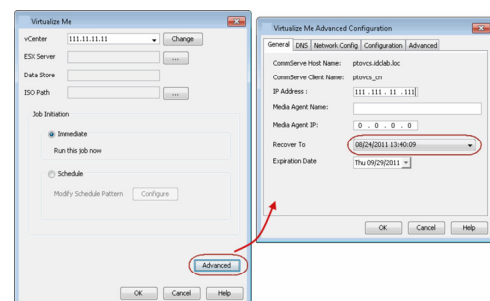


CREATING A VIRTUAL MACHINE

Once you choose a recovery point for creating the VM, follow the steps given below to create a VM using the recovery point:

1. From the CommCell Console, navigate to **Client Computers**.
2. Right-click the **<Client>** and select **All Tasks | Virtualize Me**.
3. Select a **vCenter**, **ESX Server** and a **Data Store**. Specify the **ISO Path** on the ESX server.
4. Click **Advanced**.
5. Select a date and time of the recovery point from the **Recover to** list.

A virtual machine will get created using the hardware configuration which is backed up on the selected date and time. However, the data on the virtual machine will be the data backed up during the most recent backup.



- Click **OK**.
- Click **OK** on the **Virtualize Me** dialog box.

SETTING AN EXPIRATION DATE FOR THE VIRTUAL MACHINE

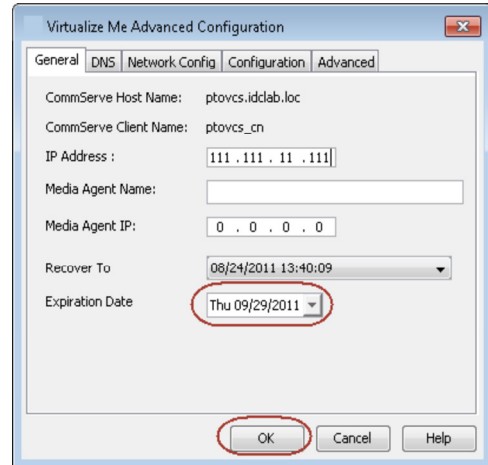
You can set an expiration date for the virtual machine. Follow the steps given below to set the expiration date of the virtual machine:

- From the CommCell Console, navigate to **Client Computers**.
- Right-click the **<Client>** and select **All Tasks | Virtualize Me**.
- Select a **vCenter**, **ESX Server** and a **Data Store**. Specify the **ISO Path** on the ESX server.
- Click **Advanced**.
- Clear the **Never** check box and select the **Expiration Date**.
- Click **OK**.

At the end of expiration date, the virtual machine is powered off and decommissioned. However, if you have performed regular backups of the virtual machine using Virtual Server iDataAgent, you can recover a decommissioned virtual machine at any time. It is recommended to perform regular backups for the virtual machine using Virtual Server iDataAgent.

You can change the expiration date later using the VM Lifecycle management feature. For more information about how to setup this feature, refer to Deployment - VM Lifecycle Management (Administrator).


Once you setup the VM Lifecycle management feature, refer to Renew a Virtual Machine.



MODIFYING THE NETWORK CONFIGURATION FOR A VIRTUAL MACHINE

When you convert a client computer into a virtual machine, you can modify it's network configuration such as IP address, default gateway etc.

- From the CommCell Console, navigate to **Client Computers**.
- Right-click the **<Client>** and select **All Tasks | Virtualize Me**.
- Select a **vCenter**, **ESX Server** and a **Data Store**. Specify the **ISO Path** on the ESX server.
- Click **Advanced**.
- Click the **Network Config** tab.
- Select the network interface and click **Edit**.
- Modify the IP address or any other network configuration details.
Click **Use DHCP** if you want to use DHCP instead of static IP address.
- Click **OK**.
- Click **OK** on the **Virtualize Me Advanced Configuration** dialog box.



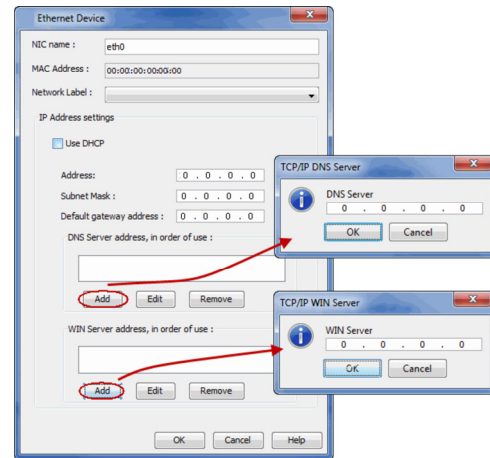
CONFIGURING WINS OR DNS FOR A WINDOWS MACHINE

When the client computer and CommServe are in the same domain, DNS will be used for name resolution between client computer and CommServe.

WINS will be used for name resolution in the following scenarios:

- CommServe and Client computer are in different networks.
 - CommServe is registered in a workgroup
- From the CommCell Console, navigate to **Client Computers**.
 - Right-click the **<Client>** and select **All Tasks | Virtualize Me**.
 - Select a **vCenter**, **ESX Server** and a **Data Store**. Specify the **ISO Path** on the ESX server.
 - Click **Advanced**.
 - Click the **Network Config** tab.

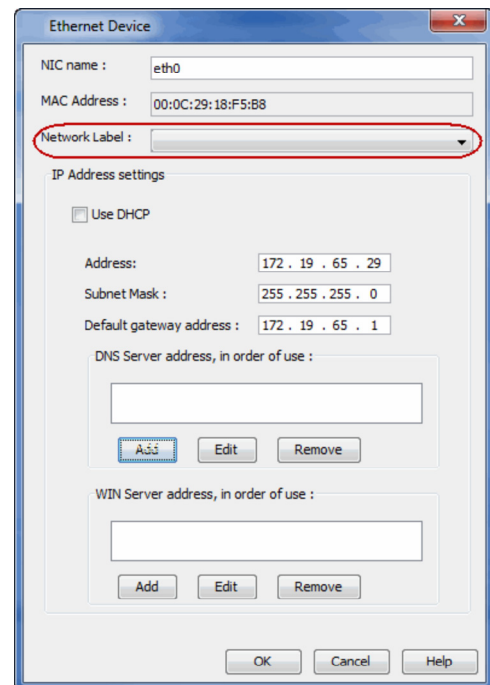
6. Select the network interface and click **Edit**.
7. If you want to use **DNS**, click **Add** and enter the IP address of the DNS Server.
If you want to use **WINS**, click **Add** and enter IP address of the WINS Server.
8. Click **OK**.
9. Click **OK** on the **Virtualize Me Advanced Configuration** dialog box.



ASSIGNING A NETWORK LABEL

Virtual Center uses labels to identify the virtual network adapter is associated with a physical network. If you want to associate the network adapter of the VM with a specific virtual server network, follow the steps given below:

1. From the CommCell Console, navigate to **Client Computers**.
2. Right-click the <Client> and select **All Tasks | Virtualize Me**.
3. Select a **vCenter**, **ESX Server** and a **Data Store**. Specify the **ISO Path** on the ESX server.
4. Click **Advanced**.
5. Click the **Network Config** tab.
6. Select the network interface and click **Edit**.
7. Select a network label from the Network Label list. This list displays all the network labels available on the ESX server or the vCenter.
8. Click **OK**.
9. Click **OK** on the **Virtualize Me Advanced Configuration** dialog box.



MODIFYING THE HARDWARE CONFIGURATION FOR A VIRTUAL MACHINE

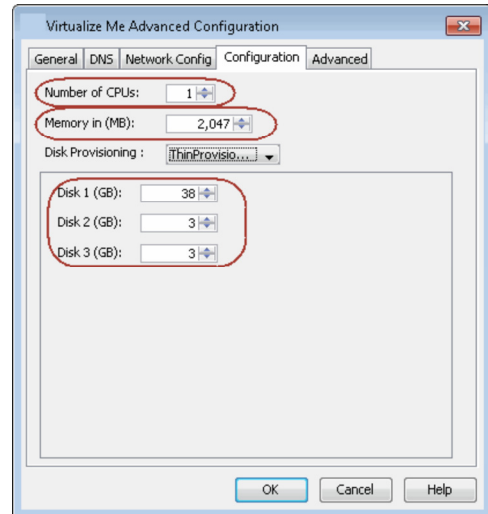
When you convert the client computer into a virtual machine, you can modify it's hardware configuration such as memory, number of CPUs, disk space etc.

1. From the CommCell Console, navigate to **Client Computers**.
2. Right-click the <Client> and select **All Tasks | Virtualize Me**.
3. Select a **vCenter**, **ESX Server** and a **Data Store**. Specify the **ISO Path** on the ESX server.
4. Click **Advanced**.
5. Click the **Configuration** tab.
6. Modify the required hardware configuration as follows:
 - o **Number of CPUs** - The existing number of CPUs in the virtual machine are displayed. You can increase the Number of CPUs.
 - o **Memory in (MB)** - The existing memory of the virtual machine is displayed. You can increase or decrease the memory.

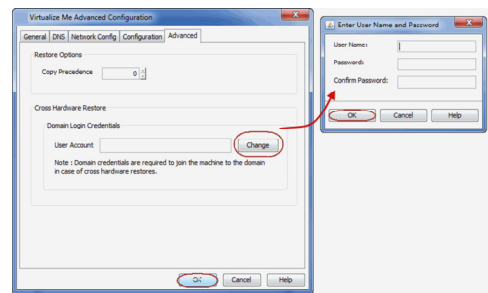
- o **Disk <n> (GB)** - The disk size of existing disks is displayed. You can increase the disk space for each disk.

The displayed disk size is always larger than the actual disk size.

For example: If the actual disk size of a disk on the physical machine is 50GB, the disk size appears as 54GB.



7. If the client computer is member of a domain, click **Advanced** tab.
8. Click **Change** and enter the credentials of the domain in which you want to add the virtual machine.
9. Click **OK**.
10. Click **OK** on the **Virtualize Me Advanced Configuration** dialog box.



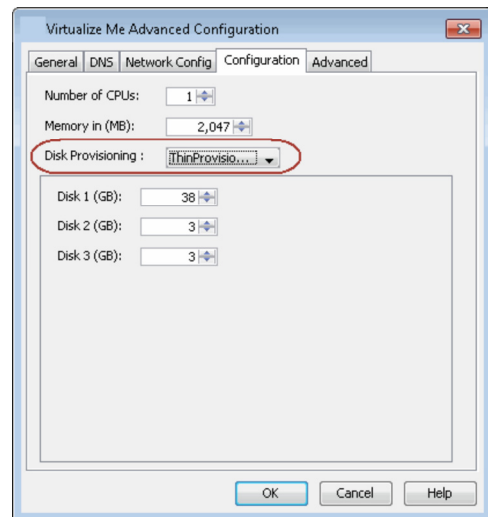
CONFIGURING THIN PROVISIONING FOR A VIRTUAL MACHINE

Thin Provisioning is a technique of allocating disk storage space in a flexible manner among multiple users, based on the minimum space required by each user at any given time. Follow the steps given below to use thin provisioning for disk space allocation:

1. From the CommCell Console, navigate to **Client Computers**.
2. Right-click the **<Client>** and select **All Tasks | Virtualize Me**.
3. Select a **vCenter, ESX Server** and a **Data Store**. Specify the **ISO Path** on the ESX server.
4. Click **Advanced**.
5. Click the **Configuration** tab.
6. Ensure that **Thin Provisioning** is selected from the **Disk Provisioning** list.

You cannot pre-allocate disk space for a virtual machine. Thin provisioning is always used to allocate disk space even if you select Thick Provisioning from the **Disk Provisioning** list.

7. Click **OK**.

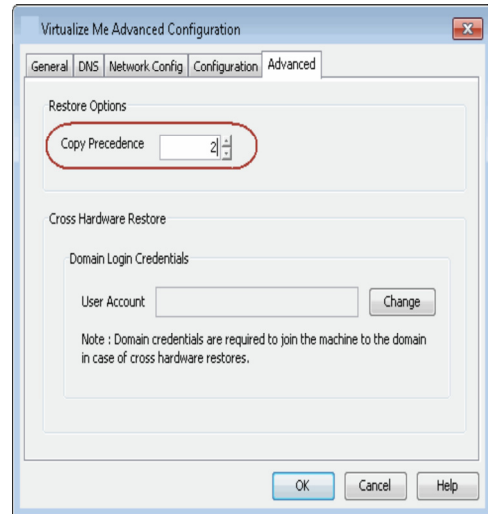


CREATING A VIRTUAL MACHINE USING A BACKUP COPY

You can use a specific copy of the backup data to create the virtual machine. By default the virtual machine is created using the backup data from the storage policy copy which has the lowest copy precedence. Follow the steps given below to specify the copy precedence of the backup copy:

1. From the CommCell Console, navigate to **Client Computers**.
2. Right-click the **<Client>** and select **All Tasks | Virtualize Me**.
3. Select a **vCenter, ESX Server** and a **Data Store**. Specify the **ISO Path** on the ESX server.

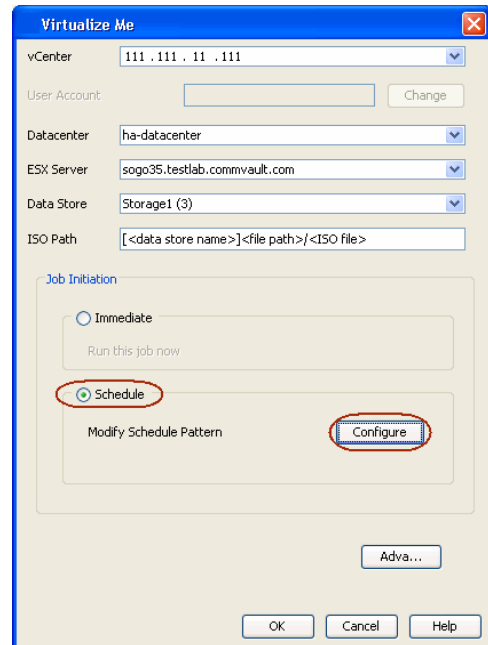
4. Click **Advanced**.
5. Click the **Advanced** tab.
6. Enter the Copy Precedence of the copy which you want to use to create the virtual machine.
7. Click **OK**.



SCHEDULING THE VIRTUALIZATION OPERATION

Follow the steps given below to schedule the virtualization of a physical machine:

1. From the CommCell Console, navigate to **Client Computers**.
2. Right-click the **<Client>** and select **All Tasks | Virtualize Me**.
3. Select a vCenter from the **vCenter** list.
4. Click **...** and browse to select an **ESX Server** and a **Data Store**.
5. Click **...** and browse to the location of the ISO file on the ESX Server.
6. Click **Schedule** and then click **Configure**.
7. Select the appropriate scheduling options.
8. Click **OK** to close the **Schedule Details** dialog.



RELATED REPORTS

JOB SUMMARY REPORT

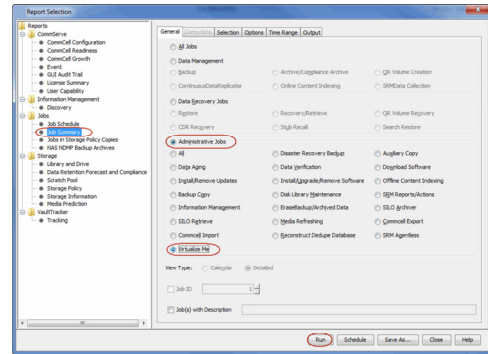
You can use the Virtualize Me Job Summary Report to review information related to all Virtualize Me jobs that run in the CommCell during the specified time period. This report is useful if you need to know:

- All Virtualize Me jobs with a specified status during a specified time period
- The user who set up the job
- The client for which the job run
- The failure reasons, if applicable

1. On the CommCell Console toolbar, click the **Reports** button.
2. From the Reports pane, select **Job Summary**.
3. On the General tab, select **Administrative Jobs**, and then select **Virtualize Me**.

4. Click **Run**.

The report appears in your default Web browser.



Virtualizing a Windows Computer - Best Practices

[Getting Started](#) | [Advanced](#) | [Best Practices](#) | [FAQ](#) | [Troubleshooting](#)

[◀ Previous](#) [Next ▶](#)

DO NOT CONTROL THE VIRTUALIZATION OPERATION USING VSPHERE CLIENT

If you have access to a vSphere client, you can monitor the progress of virtualization on the client console. However, do not use client console for any modifications in the VM when the virtualization job is running. The job may fail or suspend.

[◀ Previous](#) [Next ▶](#)

Virtualizing a Windows Computer - Frequently Asked Questions

[Getting Started](#)
[Advanced](#)
[Best Practices](#)
[FAQ](#)
[Troubleshooting](#)
[◀ Previous](#)
[Next ▶](#)

CAN I PERFORM 1-TOUCH RECOVERY IF THE CLIENT AND THE COMMSERVE ARE SEPARATED BY A FIREWALL?

No. You cannot perform 1-Touch Recovery (Interactive, Non-Interactive, or Online Restore) when firewall is configured between the client and the CommServe.

CAN I CONVERT A VIRTUAL MACHINE INTO A PHYSICAL MACHINE?

Yes. Follow the steps given below to convert a virtual machine into a physical machine:

1. Ensure that you have enabled 1-Touch backup on the virtual machine and you have performed atleast one full backup.
2. Create a response file for the virtual machine and perform the Non-Interactive Restore using the response file.

For more information about creating response file for a Windows machine, refer to Non-Interactive Restore.

For more information about creating response file for a Linux machine, refer to Non-Interactive Restore.

During Non interactive restore, copy the response file on a USB drive and connect the USB drive to a physical machine.

WHICH SUBCLIENT CAN BE USED TO PERFORM THE BACKUP OF THE MACHINE?

You must use default subclient in the default backupset to perform the 1-Touch or system state backup on the machine. Any other subclient or backup set is not supported for Virtualize Me.

DO I NEED TO PERFORM FULL BACKUP OF THE CLIENT BEFORE VIRTUALIZING IT?

If you have performed the full backup of the client computer before installing the SP4 (or higher) version of File System iDataAgent, you can perform incremental backup of the client and then start the virtualization operation.

IS IT MANDATORY TO CREATE AN INSTANCE FOR THE VCENTER IN WHICH I WANT TO CREATE THE VM?

No. It is not mandatory to create an instance for the vCenter. You must install Virtual Server iDataAgent in the CommCell. If you do not create an instance of the vCenter, follow the steps given below:

1. Click **Change** on the **Virtualize Me** dialog box and enter the hostname of the **vCenter**.
2. Enter the **User Name** and **Password** for the vCenter.
3. Enter the **ESX Server** name and **Data Store** name.
4. Enter the path of the Boot CD image in the **ISO Path** box.

In such scenario, the the information about the ESX servers and Datastores on the vCenter is not pre-populated in the **Virtualize Me** dialog box. Therefore, you cannot browse and select ESX server and ISO path on the vCenter. You can enter the ESX server and ISO path manually in the **Virtualize Me** dialog box.

WHICH DISK TYPES ARE SUPPORTED?

All the disks on the client computer can be created in a VM. However, following disk types are not supported:

- USB
- iSCSI Disks

DO I NEED TO CREATE A VM LIFECYCLE POLICY WHEN I VIRTUALIZE A COMPUTER?

No. There is no need to create a VM Lifecycle Policy during the process of virtualizing a computer.

[◀ Previous](#)
[Next ▶](#)

Virtualizing a Linux Computer - Getting Started

Getting Started | Advanced | Best Practices | FAQ | Troubleshooting

Previous Next

You can convert a client computer into a virtual machine (VM). Perform the following to convert a client computer into a VM:

1. Verify the Prerequisites
2. Setup the Client
3. Download the Boot CD Image
4. Setup the ESX Server
5. Create an Instance
6. Virtualize the Client

VERIFY THE PREREQUISITES

WHAT CAN BE VIRTUALIZED	WHAT CANNOT BE VIRTUALIZED
Any client that satisfies the System Requirements for 1-Touch client.	

INSTALL AND SETUP VIRTUAL SERVER iDATAAGENT

Ensure that the Virtual Server iDataAgent is installed on atleast one client computer in the CommCell. For more information, refer to Getting Started - VMware Deployment.

SETUP THE CLIENT


INSTALL THE LINUX FILE SYSTEM iDATAAGENT

Install the Linux File System iDataAgent using the Calypso Installer Version released with Service Pack 4 (or higher). For more information, refer to Getting Started - Linux File System Deployment.

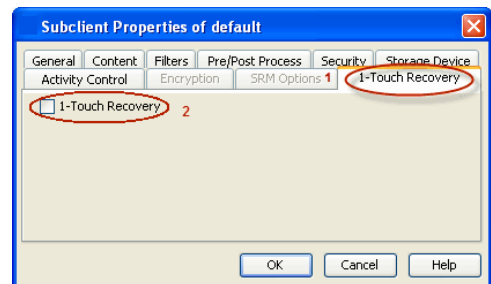
ENABLE 1-TOUCH

1.
 - From the CommCell Console, navigate to <Client> | **File System iDataAgent | defaultBackupSet**.
 - Right-click the default subclient and click **Properties**.

You must use default subclient to perform the 1-Touch backup.



2.
 - Click **1-Touch Recovery** tab.
 - Select **1-Touch Recovery**.
 - Click **OK**.

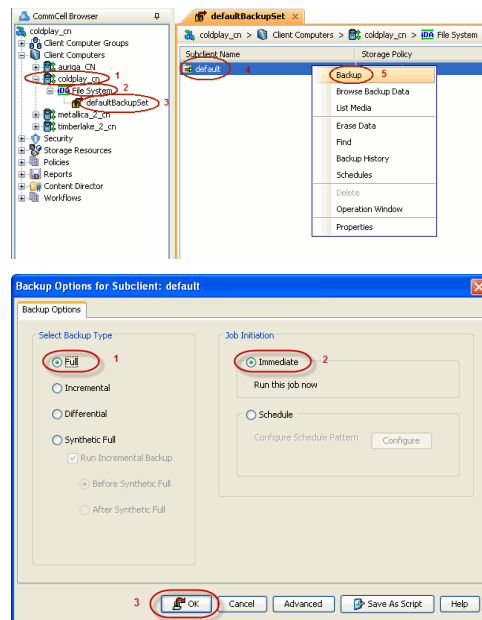


PERFORM A BACKUP

1.
 - From the CommCell Console, navigate to <Client> | **File System iDataAgent | defaultBackupSet**.
 - Right-click the subclient, and click **Backup**.

2.
 - Select **Full**.
 - Select **Immediate**.
 - Click **OK**.

Before virtualizing a client computer, ensure that you have successfully completed at least one full backup of the default subclient.



DOWNLOAD THE BOOT CD IMAGE

DVD4.iso is the Boot CD image for a Linux client. It will be available in the Linux Live CD directory of the Calypso Installer package. You can also download the DVD4.iso. Ensure that you have SP4 or higher version of DVD4.iso.

1. Log on to the Maintenance Advantage web site.
2. Click the **Downloads & Packages** tab to download the DVD4.iso.

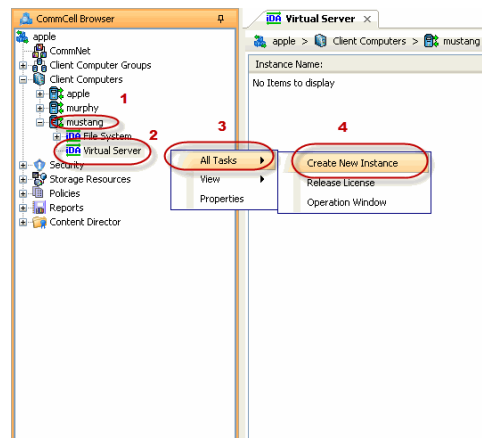
SETUP THE ESX SERVER

1. Verify the connectivity between the CommServe and the vCenter in which you want to create the VM
2. Copy the boot CD Image to any datastore on the ESX server in which you want to create the VM.

CREATE AN INSTANCE IN THE VIRTUAL SERVER IDATAAGENT

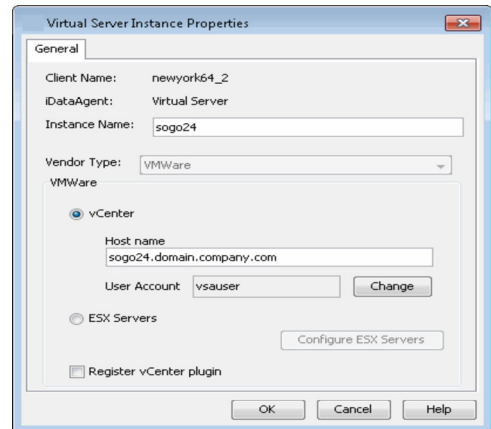

Follow the steps given below to create an instance for a vCenter in which you want to create the VM:

1.
 - From the CommCell Browser, navigate to **<Client> | Virtual Server**.
 - Right-click the **Virtual Server | All Tasks** and click **Create New Instance**.



2.
 - Enter the **Instance Name**.
 - Select **VMware** from **Vendor Type** menu.

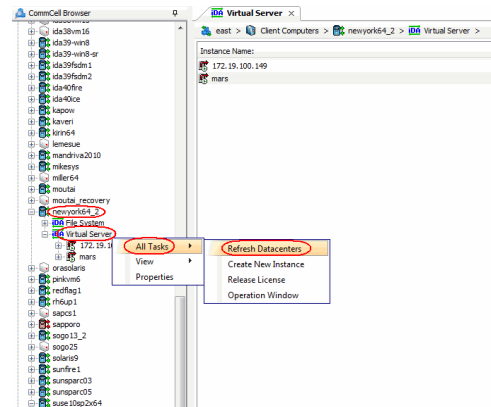
3.
 - Click **Virtual Center**.
 - Enter the host name of the vCenter.
 - Click **Configure Password**.
 - Enter the username and password associated with the Virtual Center.
 - Click **OK** to save the instance.



4.
 - From the CommCell Browser, navigate to **Client Computers | <Client>**.
 - Right-click the **Virtual Server** and select **All Tasks | Refresh Datacenters**.

It could take several minutes for a newly created instance to appear in the list of instances.

Once you create the instance, the information about the ESX sever, vCenter and the Data Store is pre-populated in the **Virtualize Me** dialog box.



VIRTUALIZE THE CLIENT

Follow the steps given below to convert a client computer into a VM:

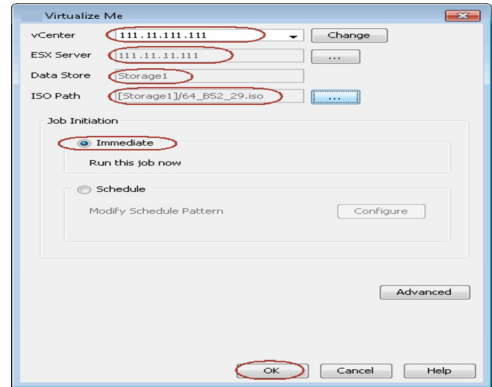
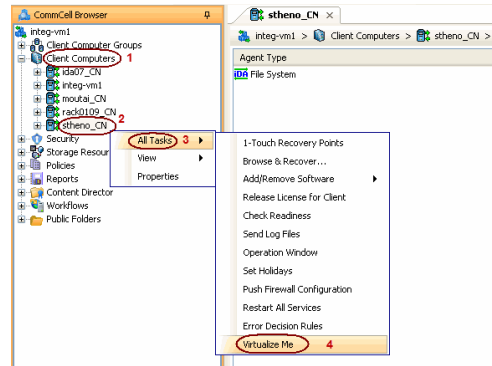
1. Turn off the client computer.
2.
 - From the CommCell Console, navigate to **Client Computers**.
 - Right-click the **<Client>** and select **All Tasks | Virtualize Me**.

3.
 - Select a vCenter from the **vCenter** list.
 - Click the ... button and browse to select an **ESX Server** and a **Data Store**.
 - Click the ... button for the **ISO Path** box and browse to the location of the Boot CD image on the ESX Server.

If you want to change the hardware configuration, click **Advanced**. For more information, refer to Advanced section of this documentation.


 - Click **OK**.

The virtual machine will be created in the specified ESX server.




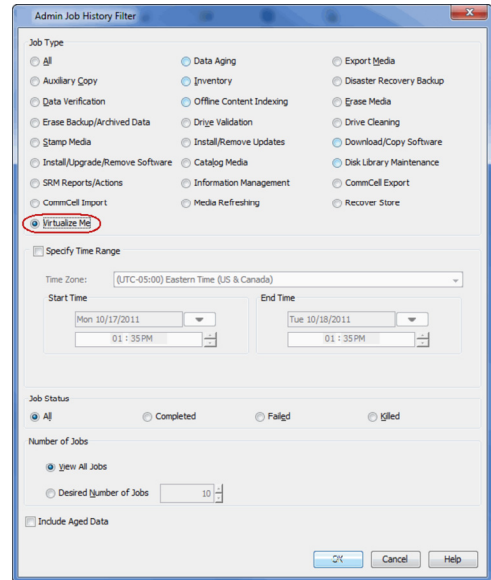
4. You can track the progress of the Virtualization job from the **Job Controller** window of the CommCell Console. The following three jobs are launched during the virtualization operation:
 - **Virtualize Machine**
 - **1-Touch Configuration Restore**
 - **Full System Restore**

You can not suspend or kill the **Virtualize Machine** job when the other two jobs are running.



5. Once the job is complete, view the job details from the **Admin Job History**. From the CommCell Browser, right-click **<CommServe>** and select **View | Admin Job History**.
6.
 - Select **Virtualize Me** under **Job Type**.
 - Click **OK**.





7. You can view the following options about the job by right-clicking the job:

- Details of the job
- Events of the job
- Log files of the job
- Resubmit the job to run it again

Job Type	Job ID	Status	Start Time	End Time	User Name	Error Code	Storage Policy
Virtualize Machine	18790	Completed	10/12/2011 20:36:03	10/12/2011 18:38:36	admin		
Virtualize Machine	18144	Completed	10/30/2011 15:38:17	10/31/2011 14:34:40	admin		
Virtualize Machine	18143	Killed	10/30/2011 15:11:54	10/30/2011 15:55:12	admin	[18861]	

CONGRATULATIONS - YOU HAVE SUCCESSFULLY CONVERTED A CLIENT COMPUTER INTO A VIRTUAL MACHINE.

If you want to further explore additional features read the Advanced section of this documentation.



Virtualizing a Linux Computer - Advanced Options

Getting Started | **Advanced** | Best Practices | FAQ | Troubleshooting

◀ Previous Next ▶

TABLE OF CONTENTS

Creating a Virtual Machine Using a Recovery Point

Verifying the Client Configuration Associated with a Recovery Point
Creating a Virtual Machine

Setting an Expiration Date for the Virtual Machine

Modifying the Network Configuration for a Virtual Machine

Assigning a Network Label

Modifying the Hardware Configuration for a Virtual Machine

Configuring Thin Provisioning for a Virtual Machine

Creating a Virtual Machine Using a Backup Copy

Scheduling the Virtualization Operation

Related Reports

Job Summary Report

CREATING A VIRTUAL MACHINE USING A RECOVERY POINT

A recovery point is generated whenever you change the hardware or network configuration of the client and perform a backup. You can use recovery points to create a virtual machine with an old hardware or network configuration. By default, the most recent recovery point is used to create the virtual machine.

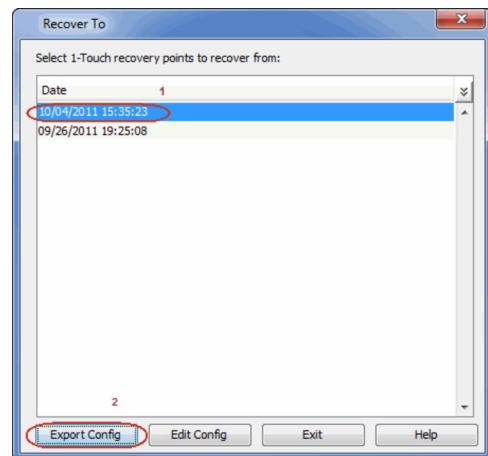
VERIFYING THE CLIENT CONFIGURATION ASSOCIATED WITH A RECOVERY POINT

Before using a recovery point, you can compare the client configurations associated with various recovery points and then choose a recovery point for creating the VM. You can generate an xml file for each recovery point. The xml file contains the client configuration details. Follow the steps given below to generate the xml file for a recovery point:

1. From the CommCell Console, navigate to **Client Computers**.
2. Right-click the <Client> and select **All Tasks | 1-Touch Recovery Point**.
3. Select a recovery point and click **Export Config**.
4. Specify the location to save the xml file and click **Save**.

By default, the the xml file is saved in the `My documents` folder on the computer that is running the CommCell Console.

5. Open the xml file using any text editor and verify the client configuration.



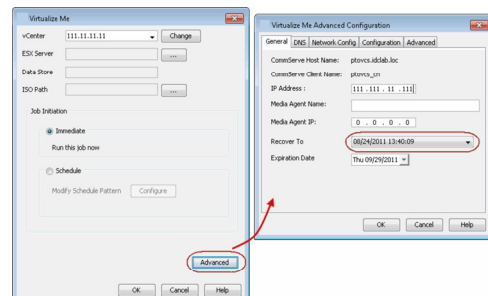
CREATING A VIRTUAL MACHINE

Once you choose a recovery point for creating the VM, follow the steps given below to create a VM using the recovery point:

1. From the CommCell Console, navigate to **Client Computers**.
2. Right-click the <Client> and select **All Tasks | Virtualize Me**.
3. Select a **vCenter**, **ESX Server** and a **Data Store**. Specify the **ISO Path** on the ESX server.
4. Click **Advanced**.
5. Select a date and time of the recovery point from the **Recover To** list.

A virtual machine will get created using the hardware configuration which is backed up on the selected date and time. However, the data on the virtual machine will be the data backed up during the most recent backup.

6. Click **OK**.



7. Click **OK** on the **Virtualize Me** dialog box.

SETTING AN EXPIRATION DATE FOR THE VIRTUAL MACHINE

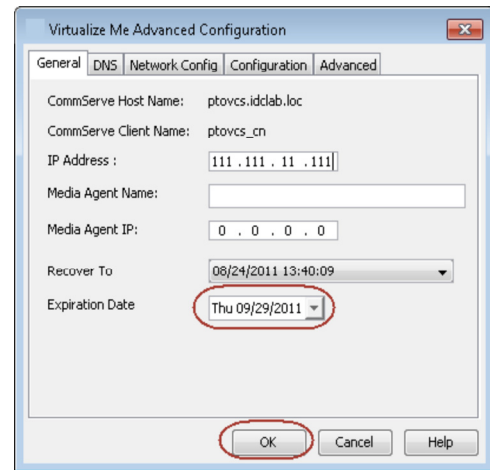
You can set an expiration date for the virtual machine. Follow the steps given below to set the expiration date of the virtual machine:

1. From the CommCell Console, navigate to **Client Computers**.
2. Right-click the **<Client>** and select **All Tasks | Virtualize Me**.
3. Select a **vCenter**, **ESX Server** and a **Data Store**. Specify the **ISO Path** on the ESX server.
4. Click **Advanced**.
5. Clear the **Never** check box and select the **Expiration Date**.
6. Click **OK**.

At the end of expiration date, the virtual machine is powered off and decommissioned. However, if you have performed regular backups of the virtual machine using Virtual Server iDataAgent, you can recover a decommissioned virtual machine at any time. It is recommended to perform regular backups for the virtual machine using Virtual Server iDataAgent.

You can change the expiration date later using the VM Lifecycle management feature. For more information about how to setup this feature, refer to Deployment - VM Lifecycle Management (Administrator).


Once you setup the VM Lifecycle management feature, refer to Renew a Virtual Machine.



MODIFYING THE NETWORK CONFIGURATION FOR A VIRTUAL MACHINE

When you convert a client computer into a virtual machine, you can modify its network configuration such as IP address, default gateway etc.

1. From the CommCell Console, navigate to **Client Computers**.
2. Right-click the **<Client>** and select **All Tasks | Virtualize Me**.
3. Select a **vCenter**, **ESX Server** and a **Data Store**. Specify the **ISO Path** on the ESX server.
4. Click **Advanced**.
5. Click the **Network Config** tab.
6. Select the network interface and click **Edit**.
7. Modify the IP address or any other network configuration details.
Click **Use DHCP** if you want to use DHCP instead of static IP address.
8. Click **OK**.
9. Click **OK** on the **Virtualize Me Advanced Configuration** dialog box.

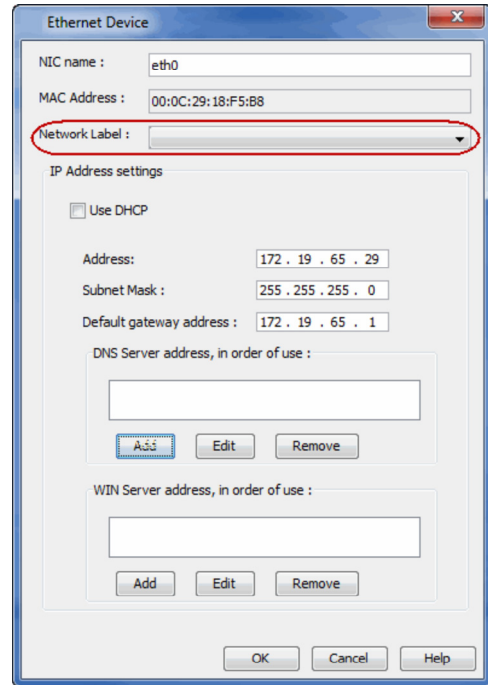


The modified network configuration is used only during virtualization operation. Once the virtualization operation completes and the virtual machine reboots, you must perform the network configuration again.

ASSIGNING A NETWORK LABEL

Virtual Center uses labels to identify the virtual network adapter is associated with a physical network. If you want to associate the network adapter of the VM with a specific virtual server network, follow the steps given below:

1. From the CommCell Console, navigate to **Client Computers**.
2. Right-click the **<Client>** and select **All Tasks | Virtualize Me**.
3. Select a **vCenter**, **ESX Server** and a **Data Store**. Specify the **ISO Path** on the ESX server.
4. Click **Advanced**.
5. Click the **Network Config** tab.
6. Select the network interface and click **Edit**.
7. Select a network label from the Network Label list. This list displays all the network labels available on the ESX server or the vCenter.
8. Click **OK**.
9. Click **OK** on the **Virtualize Me Advanced Configuration** dialog box.



MODIFYING THE HARDWARE CONFIGURATION FOR A VIRTUAL MACHINE

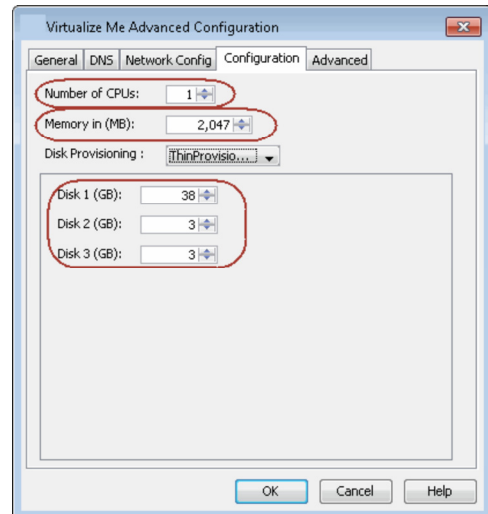
When you convert the client computer into a virtual machine, you can modify its hardware configuration such as memory, number of CPUs, disk space etc.

1. From the CommCell Console, navigate to **Client Computers**.
2. Right-click the **<Client>** and select **All Tasks | Virtualize Me**.
3. Select a **vCenter, ESX Server** and a **Data Store**. Specify the **ISO Path** on the ESX server.
4. Click **Advanced**.
5. Click the **Configuration** tab.
6. Modify the required hardware configuration as follows:
 - o **Number of CPUs** - The existing number of CPUs in the virtual machine are displayed. You can increase the Number of CPUs.
 - o **Memory in (MB)** - The existing memory of the virtual machine is displayed. You can increase or decrease the memory.
 - o **Disk <n> (GB)** - The disk size of existing disks is displayed. You can increase the disk space for each disk.

The displayed disk size is always larger than the actual disk size.

For example: If the actual disk size of a disk on the physical machine is 50GB, the disk size appears as 54GB.

7. Click **OK**.



CONFIGURING THIN PROVISIONING FOR A VIRTUAL MACHINE

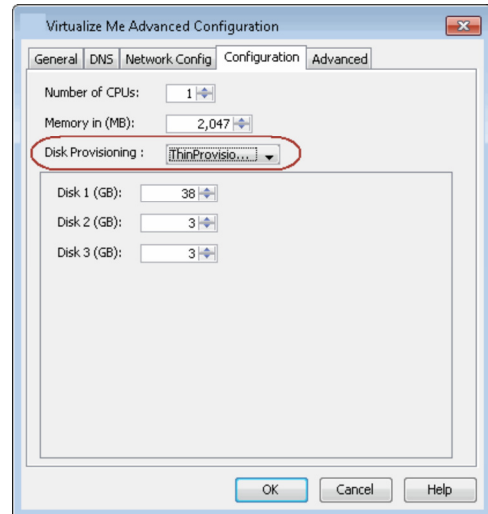
Thin Provisioning is a technique of allocating disk storage space in a flexible manner among multiple users, based on the minimum space required by each user at any given time. Follow the steps given below to use thin provisioning for disk space allocation:

1. From the CommCell Console, navigate to **Client Computers**.
2. Right-click the **<Client>** and select **All Tasks | Virtualize Me**.
3. Select a **vCenter, ESX Server** and a **Data Store**. Specify the **ISO Path** on the ESX server.
4. Click **Advanced**.
5. Click the **Configuration** tab.
6. Ensure that **Thin Provisioning** is selected from the **Disk Provisioning** list.

You cannot pre-allocate disk space for a virtual machine. Thin provisioning is always

used to allocate disk space even if you select Thick Provisioning from the **Disk Provisioning** list.

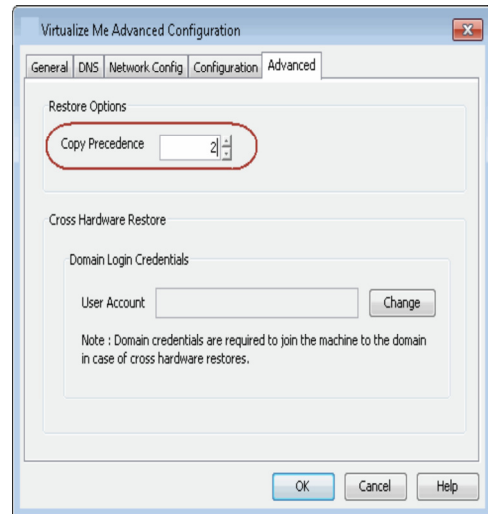
7. Click **OK**.



CREATING A VIRTUAL MACHINE USING A BACKUP COPY

You can use a specific copy of the backup data to create the virtual machine. By default the virtual machine is created using the backup data from the storage policy copy which has the lowest copy precedence. Follow the steps given below to specify the copy precedence of the backup copy:

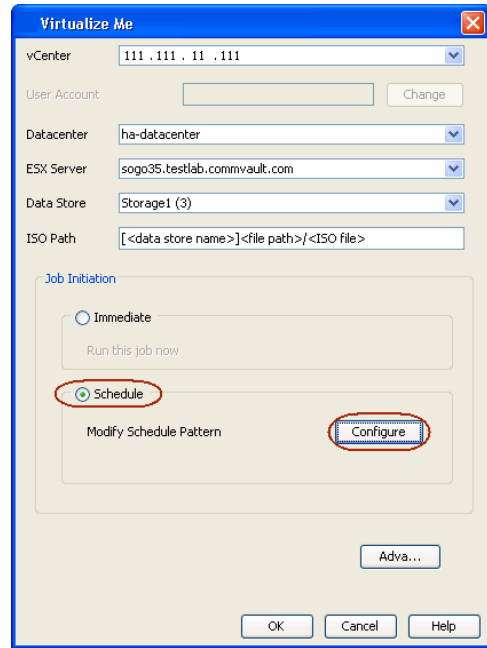
1. From the CommCell Console, navigate to **Client Computers**.
2. Right-click the **<Client>** and select **All Tasks | Virtualize Me**.
3. Select a **vCenter**, **ESX Server** and a **Data Store**. Specify the **ISO Path** on the ESX server.
4. Click **Advanced**.
5. Click the **Advanced** tab.
6. Enter the Copy Precedence of the copy which you want to use to create the virtual machine.
7. Click **OK**.



SCHEDULING THE VIRTUALIZATION OPERATION

Follow the steps given below to schedule the virtualization of a physical machine:

1. From the CommCell Console, navigate to **Client Computers**.
2. Right-click the **<Client>** and select **All Tasks | Virtualize Me**.
3. Select a vCenter from the **vCenter** list.
4. Click **...** and browse to select an **ESX Server** and a **Data Store**.
5. Click **...** and browse to the location of the ISO file on the ESX Server.
6. Click **Schedule** and then click **Configure**.
7. Select the appropriate scheduling options.
8. Click **OK** to close the **Schedule Details** dialog.



RELATED REPORTS

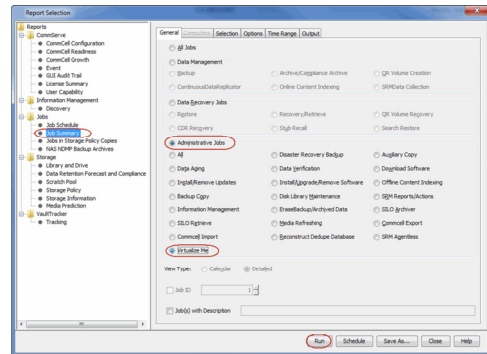
JOB SUMMARY REPORT

You can use the Virtualize Me Job Summary Report to review information related to all Virtualize Me jobs that run in the CommCell during the specified time period. This report is useful if you need to know:

- All Virtualize Me jobs with a specified status during a specified time period
- The user who set up the job
- The client for which the job run
- The failure reasons, if applicable

1. On the CommCell Console toolbar, click the **Reports** button.
2. From the Reports pane, select **Job Summary**.
3. On the General tab, select **Administrative Jobs**, and then select **Virtualize Me**.
4. Click **Run**.

The report appears in your default Web browser.



Virtualizing a Linux Computer - Best Practices

[Getting Started](#) | [Advanced](#) | [Best Practices](#) | [FAQ](#) | [Troubleshooting](#)

[◀ Previous](#) [Next ▶](#)

DO NOT CONTROL THE VIRTUALIZATION OPERATION USING VSPHERE CLIENT

If you have access to a vSphere client, you can monitor the progress of virtualization on the client console. However, do not use client console for any modifications in the VM when the virtualization job is running. The job may fail or suspend.

[◀ Previous](#) [Next ▶](#)

Virtualizing a Linux Computer - Frequently Asked Questions

[Getting Started](#)
[Advanced](#)
[Best Practices](#)
[FAQ](#)
[Troubleshooting](#)
[◀ Previous](#)
[Next ▶](#)

CAN I CONVERT A VIRTUAL MACHINE INTO A PHYSICAL MACHINE?

Yes. Follow the steps given below to convert a virtual machine into a physical machine:

1. Ensure that you have enabled 1-Touch backup on the virtual machine and you have performed atleast one full backup.
2. Create a response file for the virtual machine and perform the Non-Interactive Restore using the response file.

For more information about creating response file for a Windows machine, refer to Non-Interactive Restore.

For more information about creating response file for a Linux machine, refer to Non-Interactive Restore.

During Non interactive restore, copy the response file on a USB drive and connect the USB drive to a physical machine.

WHICH SUBCLIENT CAN BE USED TO PERFORM THE BACKUP OF THE MACHINE?

You must use default subclient in the default backupset to perform the 1-Touch or system state backup on the machine. Any other subclient or backup set is not supported for Virtualize Me.

DO I NEED TO PERFORM FULL BACKUP OF THE CLIENT BEFORE VIRTUALIZING IT?

If you have performed the full backup of the client computer before installing the SP4 (or higher) version of File System iDataAgent, you can perform incremental backup of the client and then start the virtualization operation.

IS IT MANDATORY TO CREATE AN INSTANCE FOR THE VCENTER IN WHICH I WANT TO CREATE THE VM?

No. It is not mandatory to create an instance for the vCenter. You must install Virtual Server iDataAgent in the CommCell. If you do not create an instance of the vCenter, follow the steps given below:

1. Click **Change** on the **Virtualize Me** dialog box and enter the hostname of the **vCenter**.
2. Enter the **User Name** and **Password** for the vCenter.
3. Enter the **ESX Server** name and **Data Store** name.
4. Enter the path of the Boot CD image in the **ISO Path** box.

In such scenario, the the information about the ESX servers and Datastores on the vCenter is not pre-populated in the **Virtualize Me** dialog box. Therefore, you cannot browse and select ESX server and ISO path on the vCenter. You can enter the ESX server and ISO path manually in the **Virtualize Me** dialog box.

DO I NEED TO CREATE A VM LIFECYCLE POLICY WHEN I VIRTUALIZE A COMPUTER?

No. There is no need to create a VM Lifecycle Policy during the process of virtualizing a computer.

[◀ Previous](#)
[Next ▶](#)

Frequently Asked Questions - VMware

HOW CAN I BACK UP AND RESTORE ALL VIRTUAL MACHINES WITHIN A SPECIFIC DATASTORE?

All virtual machines in a specific datastore can be backed up and restored together as follows:

1. Create a new subclient that will be dedicated to the datastore containing the virtual machines you want to back up. Consider giving this subclient the name of the datastore you are backing up for easy identification in the CommCell Console. Refer to [Creating User-Defined Subclients](#) for complete step-by-step instructions.
2. From the backup set in which the new subclient is created, configure automatic discovery of virtual machines based on datastore affinity. You can then assign the desired datastore to the new subclient. Refer to [Discover by Datastore Affinity](#) for complete step-by-step instructions.
3. Schedule routine backups of the subclient. Refer to [Scheduling a Backup](#) for complete step-by-step instructions.
4. When a restore is needed, you can restore all virtual machines within the datastore by performing a browse operation on the subclient, selecting all virtual machines displayed in the **Browse** window, and selecting either the original datastore or a new datastore as the destination from the **Restore Options** dialog box.

Refer to the following for complete step-by-step instructions:

- o Restoring Virtual Machines to Same Destination (In-Place Restore) - This procedure describes the steps in restoring the virtual machines to the exact location from which they were backed up.
- o Restoring Virtual Machines to Different Destination (Out-of-place Restore) - Different Datastore - This procedure describes the steps in restoring the virtual machines to a different datastore of your choice.

The above-mentioned steps can also be customized to group backups and restores of other entities, such as ESX Server, Resource Pools, etc.

THE BACKUP THROUGHPUT IS VERY LOW DURING THE FULL BACKUP.

If you are performing a full backup of virtual machines which have thin provisioned disks on NFS datastore, the backup throughput may be less than a VMFS datastore. VMware does not support the retrieving allocated blocks on NFS volume. Therefore, during the full backup, the software reads the complete disk and writes only valid data blocks on the media and ignores the white spaces. This reduces the backup throughput during the full backup. In case of incremental backups, software uses Change Block Tracking (CBT) and thus reads and backs up only the changed data.

For more information, refer to "Changed Block tracking on Virtual disks" section in the following document:

http://pubs.vmware.com/vsphere-51/topic/com.vmware.ICbase/PDF/vddk51_programming.pdf

If the Thin provisioned disk is on VMFS volume, the software reads and backs up only the allocated part of the disk.

HOW DO I ENSURE MIGRATED VIRTUAL MACHINES ARE BACKED UP?

If you change a virtual machine's GUID when you migrate the virtual machine from one storage location to another, auto discover does not recognize the migrated virtual machine, and the virtual machine is no longer automatically backed up.

- When you migrate a virtual machine from one storage location to another, maintain the virtual machine's GUID.
- If you change the virtual machine's GUID during migration, manually add the virtual machine to the subclient content after the migration is complete.

HOW DO I EXCLUDE VIRTUAL MACHINES FROM BACKUPS?

Virtual machines can be excluded from backups by using regular expressions to filter them out during the discovery process. There are two ways to do this:

- Assign expressions for the virtual machines to be excluded (such as ***exch*** for virtual machine names starting with **exch**) to the **Do Not Back Up** subclient.
- Assign expressions for only the virtual machines you wish to back up to a user-defined subclient.

For example, if you only want to back up virtual machines with names beginning with **win** and **sql**, then you can create a user-defined subclient with the expressions ***win*** and ***sql*** assigned to it. Thus, only virtual machines containing these expressions in their names will be backed up.

Note, however, that the default subclient does not support regular expressions for excluding virtual machines from backups.

For more information on discovery options, refer to [Advanced - VMware Configuration](#).

HOW DO I EXCLUDE CERTAIN FOLDERS FROM FILE-LEVEL BACKUPS?

Folders residing in specific drives cannot be filtered during vSphere VADP file-level backups. You can, however, replace the drive letter in the filter with an asterisk (*). Note, however, that this option will filter out the selected folders and files from all drives.

For Example:

- To filter [C:**\Windows***], the following filter can be used:

[*:**\Windows***]

- To filter [C:**\Program Files***], the following filter can be used:

[*:**\Program Files***]

Also note that global filters are not supported with volume-level backups.

I HAVE VIRTUAL MACHINES USING THE SAME GUID ACROSS VCENTER/ESX SERVERS. HOW CAN I BACK THESE UP?

If there are multiple virtual machines with same GUID:

1. Configure the instance with **Virtual Center** option
2. Right-click the BackupSet and select **Properties**
3. Select the **Automatically add new VMs that do not qualify for membership in any of the subclients** option.
4. Run a Full backup. This will automatically fix the virtual machines with duplicate GUID's and update the subclient.

CAN I RESTORE MULTIPLE VIRTUAL MACHINES FROM THE BACKUP SET LEVEL?

You can restore multiple virtual machines from the backup set level as long as you are restoring the virtual machines to a different destination. If you want to perform in-place restore of multiple virtual machines, you must perform the browse and restore operation at the subclient level and not at the backup set level.

THE GUIDS OF SOME VIRTUAL MACHINES HAVE CHANGED SINCE THEY WERE DISCOVERED. HOW DO I ENSURE THESE ARE INCLUDED IN BACKUPS?

If the GUIDs of the virtual machines associated to a subclient are modified:

1. Run a manual discovery. Follow the steps given below to run a manual discovery:
 - Right-click the Subclient and select **Properties**
 - Select the **Content** tab and click **Configure**
 - From the **Guest Host Configuration** dialog box, click **Discover**
2. The GUIDs of the virtual machines that were modified, will be marked for deletion
3. Re-associate the virtual machines in order for the backup to continue.

ARE BACKUPS OF VCLLOUD DIRECTOR 1.0 PROVISIONED VIRTUAL MACHINES AND VAPPS SUPPORTED?

Yes, virtual machines created with vCloud Director are supported for backups, including SnapProtect. No special configuration is required. In addition to the standard protection and recovery capabilities, affinity-based automatic discovery can leverage Organization and Virtual Data Centers as discovery criteria as those entities appear as resource pools. Full virtual machine and vApp restores are supported, including single-file recovery.

WHY IS AN INCREMENTAL BACKUP CONVERTED TO A FULL BACKUP, AFTER AN IN-PLACE RESTORE?

After performing an In-place restore, the first run backup job will always be converted to a Full backup. The system assumes that it is a newly created virtual machine and hence defaults to a full backup.

HOW DO I INSTALL A SEPARATE INSTANCE OF THE VDDK ON A COMPUTER?

Version 5.0 update 2 of the VDDK is supported for backups if it is installed manually in Virtual Server iDataAgent proxy machines, and is automatically installed with Service Pack 6A or higher.

To install a newer version of VDDK on a 32-bit computer, perform the following steps:

1. VDDK must first be installed on the proxy server to its default location.
For example, **C:\Program Files\VMware\VMware Virtual Disk Development Kit**

To install a newer version of VDDK on a 64-bit computer, perform the following steps:

1. VDDK must first be installed on the proxy server to its default location.
For example, **C:\Program Files (x86)\VMware\VMware Virtual Disk Development Kit**
2. Create a new folder and point it to the path below on the proxy computer.
For example, **C:\Program Files\VMware\VMware Virtual Disk Development Kit**
3. From the VDDK x64 installation directory **C:\Program Files (x86)\VMware\VMware Virtual Disk Development Kit\bin**, unzip the **vddk64** folder to a temporary location. Once unzipped, copy the **bin**, **lib**, and **plugins** folders to the new folder created in the previous step.

4. Create the following registry key:

HKEY_LOCAL_MACHINE\SOFTWARE\VMware, Inc.\VMware Virtual Disk Development Kit\InstallPath

with the following string value: (this should be the path created in Step 2)

C:\Program Files\VMware\VMware Virtual Disk Development Kit

HOW DO BACKUP AND RESTORE OPERATIONS HANDLE INDEPENDENT/RDM DISKS?

If a virtual machine undergoing a backup job includes independent disks, physical or virtual RDMs, these disks will be skipped. During a full VM restore the independent disk/Physical or virtual RDMs will get restored as a regular disk with 0MB data.

If a subclient contains virtual machines with independent disks/physical or virtual RDMs, the backup job will always complete with the status "Completed w/ one or more errors". However, if you create the IgnoreUnsupportedDisks registry key on the proxy computer, the backup job will complete successfully.

HOW DO BACKUP AND RESTORE OPERATIONS HANDLE VIRTUAL RDM DISKS?

Virtual RDMs are protected by the backup job (but not during IntelliSnap backup). However at the time of restore, the data is restored as a regular VMDK on a datastore. A virtual RDM is not re-created and the data is not restored to the virtual RDM's device.

HOW CAN THE VIRTUAL SERVER /DATAAGENT BE SET UP TO SEE THE STORAGE NETWORK?

In an isolated network, add an additional network connection to the proxy computer.

USING THE VIRTUAL SERVER IDATAAGENT IN HOTADD MODE CONFIGURATIONS

- While deploying the Virtual Server iDataAgent, install the software on a datastore with the largest VMFS block size. This is necessary to ensure that the Virtual Server iDataAgent can mount and back up virtual machines residing on all datastores.
- Helper virtual machines are not required for Hot-Add Virtual Server iDataAgent servers utilizing VADP.

UPGRADING THE VIRTUAL SERVER IDATAAGENT FROM VERSION 8.0.0

VDDK 1.2 is installed to the Base folder with the Virtual Server iDataAgent. Software version 8.0 Virtual Server iDataAgent machines that were upgraded to v9 will need to have the Virtual Disk Development Kit registry key that was manually created during the installation of VDDK 1.1 deleted. Failure to do so will result in the VDDK 1.1 to still be utilized or fail if it has been uninstalled. (VDDK 1.1 does not support HotAdd mode on ESX 4.1)

USING VSTORAGE API

- The vStorage API can be used to restore backups performed with VCB. The backup data will first be staged to the proxy then vStorage will perform the restore.
- Backups performed using the vStorage API cannot be restored through the VMware converter. Converter does not understand the format of vStorage.

CHANGE BLOCK TRACKING

When Change Block Tracking is unavailable backups will revert to CRC to determine changed blocks. Since the Virtual Server iDataAgent needs to read the entire virtual machine disk, CRC incremental backups may take almost as long as full backups. However, the amount of data transferred and stored by an incremental backup is limited to the changed blocks within the virtual disk. Correcting CBT on the problematic virtual machine is recommended as quickly as possible to take full advantage of VADP based incremental backups. Tools are available to check the correct installation of the VDDK and for verifying and correcting CBT issues. Contact your Software Provider for information on these tools.

SUPPORTED VCENTER CONVERTER VERSION

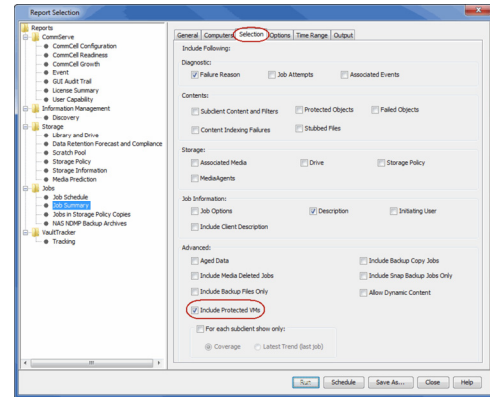
- VMware vCenter Converter Standalone 4.3 no longer supports VCB sources for restores. Version 4.0.1 must be used.
- Fault Tolerant VM's are not supported for snapshot operations and cannot be backed up.

HOW DO I INCLUDE THE STATUS OF PROTECTED VIRTUAL MACHINES IN BACKUP JOB SUMMARY REPORTS?

The Backup Job Summary report can show a detailed status of each virtual machine included in the backup. Follow the steps given below to include the status of protected virtual machines in the Job Summary Report:

1. On the CommCell Console toolbar, click the **Reports** button.
The **Report Selection** dialog box appears.
2. From the **Reports** pane, select **Job Summary**.
3. On the **General** tab, select **Data Management**, and then select **Backup**.
4. Click the **Selection** tab, and then select the **Include Protected VMs** check box.

- Click **Run**.



The report appears in your default Web browser. It contains a detailed status of all the virtual machines included in the backup.

- For snap backups, VM details appear in all columns, except for the columns that are not applicable, such as Backup Type, Transport Mode, and CBT Status.
- For jobs that ran before the Service Pack 5B update, Machine Name and Status appear, but all other columns display N/A.
- File Level and Volume Level standard backups display N/A in the VM Size and CBT Status columns.

90z2v9sgv4d65	Virtual Server	defaultBackupSet (V)	2107	DNCR	N/A	02/28/2012 18:45:58	02/28/2012 18:52:24	1.66	1.68	0%	1.67	0:04:24	22.676	8.833	0	0
Protected VMs:																
• Total Clients: 3																
• Successful Clients: 3																
• Failed Clients: 0																
• Client Priority: Successful: 0																
Machine Name	Status	VM Size	Start Time	End Time	Hardware Version	Operating System	Backup Type	Transport Mode	CBT Status	Owner						
Opst1Localvm1	Completed	2.00 GB	02/28/2012 18:47:27	02/28/2012 18:48:00	vmx-08	Microsoft Windows XP Professional (32-bit)	VADP	sb-d	Used							
Opst1Localvm1	Completed	2.00 GB	02/28/2012 18:47:33	02/28/2012 18:49:11	vmx-08	Microsoft Windows XP Professional (32-bit)	VADP	sb-d	Used							
Opst1Localvm1	Completed	2.00 GB	02/28/2012 18:48:00	02/28/2012 18:51:45	vmx-08	Microsoft Windows XP Professional (32-bit)	VADP	sb-d	Used							

WHAT ARE THE REQUIREMENTS TO ENABLE THIN PROVISIONED DISK RESTORES?

- VM Hardware version 7 or higher
- Source VM Disk backup on a VMFS datastore

Disk level backups can use allocated block tracking, which is part of VMware Changed Block Tracking (CBT). When enabled, allocated block tracking identifies portions of the virtual machine disk that have not been used, so those portions can be skipped during backup or restore.

Thin provisioned disk restores are supported only if allocated block tracking was used for the backup. Only source disks using lazy zero provisioning can be restored from thick provisioning to thin provisioning; disks using eager zero are always restored using thick provisioning.

If allocated block tracking was not active for the backup, the restore operation includes all of the blocks of the disk that were read during backup. This can result in thick provisioning for the restore, even if thin provisioning was requested.

The disk restore operation recovers all space that is reported in use (not zeroed out), so the VMDK can be larger than the size that was reported by the guest OS.

HOW CAN I RESTORE LINUX EXT4 FILES?

To restore ext4 files, or other Linux files that need to retain ACLs and permissions, restore the virtual machine disk (VMDK) to a new destination server and datastore, then ask your vCenter administrator to add the required files to the original VM or a test VM.

- In the CommCell Browser, go to **Client Computers** | **<Client>** | **Virtual Server** | **<Proxy>** | **<BackupSet>**.
- Under **<BackupSet>**, right-click the subclient and select **Browse Backup Data**.
- On the **Browse Options** dialog, ensure that **Container Restore** is selected and click **OK**.
- On the resulting content selection tab, expand the backup set and select the virtual machine disk that contains the files to be restored.
- Click **Recover All Selected**.
- On the **Restore Options for All Selected Items** dialog: Select the destination client, specify the folder where the VMDK file is to be restored, and click **OK**.
- The vCenter administrator can add the VMDK files to the original VM or to a test VM using the **Upload File** option from the Datastore Browser in the vSphere client.
- When the files are added to the VM, the Linux administrator can import the disks, mount the disks, and copy the restored files.

The original ACLs and permissions for the files are retained.

HOW IS VM BACKUP SIZE CALCULATED FOR CAPACITY LICENSING?

For VM backups, capacity licensing is based on the total backup size, calculated as the sum of backup sizes for all VM backup jobs after white spaces (blocks of zeros) are removed. The license counts the backup size of all configured subclients; virtual machines that are included in multiple subclients will be counted multiple times. The backup size is measured for usage tracking and shown on the Backup Job Summary Report.

The backup size can be different from the guest host size or used space value shown for the VM in the disk properties dialog by Microsoft Windows.

The following factors can affect the backup size calculation:

- The presence of virtual machine snapshots.
- The presence of sparse files or deleted files in the guest can cause the backup size to vary.
- The backup size reported for VMDKs in a VM can vary depending on white space and change allocation tracking for that VM.

For example:

- A Windows VM with a single volume of 80 GB has 30 GB occupied and 50 GB free.
- The guest size would be 30 GB.
- The backup size is the amount of data transferred and written for that backup, which can be up to 80 GB.

The backup size reflects the size after eliminating white spaces; but data that was written and deleted still counts as reserved (allocated) space. The layering effects of multiple virtual file systems can cause differences between the size reported by the guest host running within the VM and the reported backup size. Frequent deletion of large files can easily cause these numbers to be out of sync.

Version 9.0 reports on all allocated blocks in the VM. The amount reported for allocated blocks can be the same size or larger than what is actually in use and can contain reserved space for deleted items. For each VMware instance, Version 9.0 has an additional reporting column of the actual size of VMs.

The following measures can help reduce backup size:

- Delete or move unnecessary data before virtualizing physical machines. This saves resources and time, and ensures that new VMs only contain used blocks.
- If you already have a large number of VMs with significant reserved and unused blocks, use a tool such as the Windows SDelete utility to release reserved space on VMs, as described in SDelete v1.61.
- Ensure that VM templates used for provisioning do not contain unreferenced blocks in the VMDK.
- Wherever possible, configure virtual machine disks with the VMware Thin Provisioned disk option enabled. Thin provisioning ensures only valid blocks occupy space in the VMDK file. With thin provisioned VMs, VMware APIs only return occupied blocks.

[Back to Top](#)

Best Practices - VMware

CREATING MULTIPLE BACKUPSETS WHEN USING VIRTUAL SERVER IDATAAGENT WITH SNAPPROTECT

It is recommended to create multiple backupsets for the same ESX server or vCenter when using Virtual Server iDataAgent with hardware snapshots. You can configure multiple backupsets as follows:

BackupSet 1 - Use this backupset only for snapshots and not for backup copies. This backupset will provide multiple recovery points and fast recovery from any point-in-time.

BackupSet 2 - Use this backupset for regular backups. After performing the first full backup, you can keep on performing incremental backups. You can enable the DASH copy option while creating the secondary copy. The VMware's Change Block Tracking feature is used internally during regular backups. The Change Block Tracking and DASH copy enable faster backups.

This configuration has following advantages:

- A dedicated ESX proxy is not required for creating the backup copy.
- ESX proxy is required only for restoring data from the snapshots.
- You can perform the regular incremental backups and still you will get all benefits of a SnapProtect backup.
- No additional impact on production server or storage.

RELEASING VIRTUAL SERVER /DATAAGENT LICENSE

In case the Virtual Server iDataAgent is no longer required to run data protection operations, it is recommended to release the Virtual Server iDataAgent's license instead of uninstalling it. If the iDataAgent is uninstalled then you will not be able to do the following:

- Perform Live-Browse.
- Unmount any Virtual Server iDataAgent mounted snapshots.

USING THE DEFAULT SUBCLIENT

Use and run the default subclient regularly since it performs the VM discovery operations.

While the default subclient will be used regularly in most deployments, the following exceptions apply:

- Multiple Virtual Server iDataAgents are deployed in the environment and the 'Master' model will be in use.
- Multiple Storage Policies are required for different virtual machines.
- Rule-based discovery is in use and different subclients will be used to target a particular affinity.

CREATE A VIRTUAL MACHINE FOR MULTIPLE SUBCLIENTS

Creation of a Virtual Machine should be done for multiple subclients only. Creating one virtual machine per subclient can cause scheduling problems, and can become unmanageable in larger environments.

HOT-ADD DEPLOYMENTS

When deploying the Virtual Server iDataAgent for Hot-Add Mode Backups, choose the datastore with the largest VMFS block size to ensure backups can mount and back up virtual machines residing on all datastores.

ACTIVATING CHANGE BLOCK TRACKING FOR BACKUPS

When Change Block Tracking is unavailable, backups revert to Cyclic Redundancy Check (CRC) to determine changed blocks. Since the Virtual Server iDataAgent needs to read the entire virtual machine disk, CRC incremental backups may take almost as long as full backups even though the amount of data transferred and stored by an incremental backup is limited to the changed blocks within the virtual disk. Therefore, correcting CBT on the problematic virtual machine is recommended as quickly as possible to take full advantage of VADP-based incremental backups.

Tools are available to check the correct installation of the VDDK and for verifying and correcting CBT issues. Contact your software provider for information on obtaining these tools.

CONFIGURATION CONSIDERATIONS

- Regular backups should be performed on the default subclient. This will ensure new virtual machines are discovered on a regular basis along with the backup operation.
- Creating one virtual machine per subclient is not recommended. This creates issues around scheduling and becomes unmanageable in larger environments.
- Typically only the default Subclient will be used in most deployments except in the following cases:
 - Multiple Agents are deployed in the environment and the 'Master' model will be in use. (See the Master Agent Configurations section.)

- Multiple Storage Policies are required for different virtual machines.
- Rule-based discovery is in use and different subclients will be used to target a particular affinity.

LOCATION OF JOB RESULTS DIRECTORY

The Job Results directory stores the job results files from backup and restore operations of a client.

The user impersonation account specified for the job results directory will take precedence and will be used to backup and restore data from a virtual machine. This may result into file access related issues during the backup. Therefore, it is recommended to use a local folder on the client computer as the job results directory.

NAME OF NFS DATASTORE

It is recommended to use a short name for an NFS datastore. When you perform the SnapProtect backup of an NFS datastore, Calypso can append up to 20 characters with the name of NFS datastore to create the volume label. The ESX server supports a volume label of 42 characters.

SETTING THRESHOLD FOR FREE SPACE ON A DATASTORE

During the backup of any virtual machine, a snapshot of the virtual machine gets created on its datastore. If the datastore of any virtual machine doesn't have sufficient space for snapshots, the backup job completes with one or more errors. Before performing the backup, you can set a threshold of required free space by creating the nDatastorePercentageFreeSpace registry key. If a datastore doesn't have the required free space, the software will not perform backup of all the virtual machines residing in that datastore.

When you perform the backup of a subclient, a virtual machine will be skipped from the backup if its datastore doesn't have the required free space.

PROXY ESX SERVERS SHOULD RESIDE IN A STANDALONE (NON CLUSTERED) STORAGE GROUP /INITIATOR GROUP /MASKING VIEW

This is due to the fact that the Initiator group or the transport mechanism has cluster grouping. The clones are exposed to the existing transport mechanism having the cluster grouping where the rescan is done at the cluster level. Orphaned data stores with (_Gxbackup) will be exposed on all ESX servers and unmounts thereby not initiating a cleanup at the cluster level. Cleanup will be done only for the proxy ESX Server leaving behind stale data stores on other ESX Servers in a cluster.

Troubleshooting - VMware Backup

[Backup](#) | [Restore](#)

The following section provides information on some of the troubleshooting scenarios related to backup:

BACKUP FAILURE

During create snapshot phase	When the virtual machine version is upgraded, make sure all the VMware tools are upgraded.
After Upgrading Proxy Server from 8.0 to 9.0	Uninstall VDDK 1.1 on the proxy computer and re-install VDDK 1.2 before performing backup. *For 64-bit operating systems: From the VDDK x64 installation directory, unzip the vddk64 folder to a temporary location. Once unzipped, copy the bin, lib, and plugins folders to: <code>C:\Program Files\VMware\VMware Virtual Disk Development Kit</code>
Backup may fail with an error that a driver cannot be found	Cause 1: You may be using an older version of VDDK. To resolve this issue, upgrade your VDDK software to version 1.2.1 or higher. Cause 2: The driver may be in a stopping state. To resolve this issue, ensure the VMware driver service is in a healthy state. This may require a reboot.
With "Received empty disk extents Mount VM_VCB4 failed" error	Use the utility Change tracking to reset the Change tracking. Follow the commands given below: <code>CheckChangeTracking /host x-vcenter /user branzini /password Brown010 /vm iks-dcl.iks.it /disable</code> <code>CheckChangeTracking /host x-vcenter /user branzini /password Brown010 /vm iks-dcl.iks.it /enable</code> <code>CheckChangeTracking /host x-vcenter /user branzini /password Brown010 /vm iks-dcl.iks.it</code> It may be necessary to run this recycling of change tracking with the virtual machine powered down. It may also be necessary to power cycle the virtual machine after disabling the change tracking and again after enabling it.
With "Disk query change failures" error	Power cycle the virtual machine under backup.
With "Duplicate index entry" error	Check to see if there are virtual machines with duplicate disk names and, if so, rename one.
For a VADP backup	Changed Block Tracking (CBT) must be enabled at the time of the backup. While the Virtual Server iDataAgent enables this automatically, it may not be able to do so for cloned and migrated VMs. Check to ensure CBT is enabled for the VM for which a VADP backup failed as described below. To check if a virtual disk has CBT enabled, open the vSphere Client, select a powered-off VM, and click Edit...Settings>Options>Advanced/General>Configuration Parameters . Ensure that: The VM's (.vmx) file contains the entry: <code>ctkEnabled = "TRUE"</code> For each virtual disk, the .vmx file contains the entry: <code>scsix:x.ctkEnabled = "TRUE"</code> For each virtual disk and snapshot disk there is a .ctk file. For example: <code>vmname.vmdk</code> <code>vmname-flat.vmdk</code> <code>vmname-ctk.vmdk</code> <code>vmname-000001.vmdk</code> <code>vmname-000001-delta.vmdk</code> <code>vmname-000001-ctk.vmdk</code>
HotAdd backups that have a hard failure may result in mounted disks on the proxy that are not automatically cleaned up	To resolve this issue, manually remove the disks from the proxy virtual machine and delete the \tmp\vmware directory created in the current user's Documents and Settings directory, typically found at C:\Documents and Settings\username\Local Settings\Temp\vmware-username\. The next time a snapshot is created or deleted for the virtual machine the redo logs will get cleaned up.
File-Level and Volume-Level Backup Failures	If the virtual machine that is backed up contains disks that have not been initialized, then the volume-level and file-level backup operations will fail for that virtual machine.
with following error: FAILED to download config file [Some path].vmx, [1][12002] [Unknown native error code 12002	This error appears if WinHTTP proxy is set on the client computer where Virtual Server iDataAgent is installed. Before performing backup, ensure that WinHTTP proxy is not set on the client computer. Follow the steps given below to verify WinHTTP proxy configuration: 1. On the client computer, open the command prompt and enter following command: <code>C:\>netsh WinHTTP Show Proxy</code> If you get following response, it indicates that the WinHTTP proxy is set on the client computer. <code>Current WinHTTP proxy settings:</code> <code>Direct access (no proxy server)</code> 2. Use the following command to reset the WinHTTP proxy: <code>C:\>netsh WinHTTP Reset Proxy</code>

INCORRECT VIRTUAL MACHINE IS BACKED UP DURING INCREMENTAL BACKUP

During incremental backup, if there are virtual machines with duplicate GUIDs, rename the virtual machine GUIDS to ensure they are unique.

VERSION 4.0 VSTORAGE OPTION, BACKS UP FULL DISK RATHER THAN DATA ONLY COPY

To back up virtual machine as a data only copy using vStorage option, upgrade the virtual machine to Version 7.0.

If a virtual machine has been cloned or migrated from another ESX Server, ensure that the virtual machine is rebooted prior to running your first backup.

HOW TO FIX VIRTUAL MACHINES THAT HAVE THE SAME GUID ACROSS VCENTRE/ESX SERVER

If there are multiple virtual machines with same GUID:

1. Configure the instance with **Virtual Center** option
2. Right-click the BackupSet and select **Properties**
3. Select **Automatically add new VMs that do not qualify for membership in any of the subclients** option.
4. Run a Full backup. This will automatically fix the virtual machines with duplicate GUID's and update the subclient.

If the GUIDs of the virtual machines associated to a subclient are modified:

1. Run a manual discovery. Follow the steps given below to run a manual discovery:
 - o Right-click the Subclient and select **Properties**
 - o Select the **Content** tab and click **Configure**
 - o From the **Guest Host Configuration** dialog box, click **Discover**
2. The GUIDs of the virtual machines that were modified, will be marked for deletion.
3. Re-associate the virtual machines in order for the backup to continue.

DISK LEVEL BACKUP OF WINDOWS VIRTUAL MACHINE COMPLETES WITH ERROR

If the disk level backup of a Windows virtual machine completes with the following error: "Metadata collection failed for Virtual machine [B_VM1_disk_Full], File and Volume browse will not be available for this VM. [Initialization from Disk failed, <6001><1><MBR signature does not match>]":

The virtual machines may contain disks that are in uninitialized state. Initialize the disks before running a backup, in order to collect metadata enabling file level restores.

CONSOLIDATE HELPER-0 SNAPSHOT ON VCENTER SNAPSHOT MANAGER

SYMPTOM

The consolidate helper-0 snapshots are created on the vCenter Snapshot Manager and they cannot be deleted after the backup job completes. The following error message appears if you delete the snapshot from vCenter Snapshot Manager:

```
Unable to access file <unspecified filename> since it is locked.
```

CAUSE

This scenario can arise when the consolidate helper-0 snapshot is not committed to VMDK. This issue can be VMware I/O issue.

SOLUTION

Adjust the virtual environment to resolve high I/O. Restart the virtual machine to give the quiet time that is required commit the snapshot.

If you are unable to delete the snapshot, contact VMware Support for assistance.

VMWARE DRIVER NOT FOUND

SYMPTOM

After installing any service pack, the following error message appears when you start the backup operation:

```
CVMDiskInfo::_MountVolumes() - VixMntapi_MountVolume failed Err=5ef3 Driver not found
```

SOLUTION

1. Check the status of the driver services using the following commands:

```
sc query vstor2-mntapi10-shared
```

```
sc start vstor2-mntapi10-shared
```

2. If the driver service is in stop or pending state, reboot the machine where you have installed Virtual Server iDataAgent.

COMPLETED WITH ONE OR MORE ERRORS

Backup jobs from Virtual Server iDataAgent will be displayed as "Completed w/ one or more errors" in the Job History in the following cases:

- If the virtual machine, virtual machine disk, or virtual machine configuration file fails to back up.
- If one or more virtual machines in a backup job fail to back up.
- If communication fails with vCenter.
- If a disk included in a backup is not supported (i.e., independent disks or physical RDM).
- The meta data collection operation fails during a snap backup job, the job will be displayed as "Completed" in the **Backup Job History** of the subclient. You can create the bCWEJobMDataFails registry key if you want to display the status as "Completed w/ one or more errors" in such scenario.
- If the datastore of any virtual machine doesn't have sufficient space for snapshots, the backup job completes with one or more errors. Before performing the backup, you can set a threshold of required free space by creating the nDatastorePercentageFreeSpace registry key.

BACKUP FAILS AFTER MODIFICATIONS IN vCENTER CREDENTIALS AT THE INSTANCE LEVEL

SYMPTOM

If an ESX server moves from one vCenter to another, you have to modify the instance created for the vCenter. You must provide the Host name and login credentials for the new vCenter. In such scenario, the the subsequent backup of any virtual machine on the ESX server may fail.

CAUSE

In the CommSeve database the vCenter hostname and login credentials are updated only for the instance and not updated for all the subclients in the instance.

SOLUTION

If you update the vCenter hostname and login credentials for any instance, perform the following:

1. From the CommCell Browser, navigate to **<Client> | Virtual Server | <Instance> | defaultBackupset**.
2. Right- click any subclient and click **Properties**.
3. Click **OK**.
4. Repeat above steps for all the subclients in the instance. You can now perform backup of any subclient in the instance.

BACKUP FAILS TO DOWNLOAD VMX FILES

SYMPTOM

The backup operation fails with a error: Unable to download VMX files.

CAUSE

The name of the datastore contains + or @.

SOLUTION

1. Check if the datastore of any virtual machine included in the backup contains + or @ characters in its name.
2. Rename the datastore and remove + and @ characters from the name.
3. Perform the backup operation.

VIRTUAL MACHINE CLIENT NAMES ARE GETTING CREATED WITH '_1' APPENDED TO THE ORIGINAL CLIENT NAME

SYMPTOM

When viewing virtual machines in the Client Computers list, you may see duplicated client names (for example, *<ClientName>* and *<ClientName>_1*).

CAUSE

Virtual machine information is added to the database during the discovery phase of a backup. If information changes, a new client is discovered during the discovery phase.

Making changes to information such as the virtual machine fully qualified domain name (FQDN), the GUID, or the host name can cause this issue.

RESOLUTION**TO AVOID DUPLICATE CLIENT ENTRIES CAUSED BY DIFFERENT CLIENT OR HOST NAME:**

Use the install software option for the client level in the CommCell Console (rather than interactively installing software to the virtual machine Guest Operating System).

TO CORRECT EXISTING CLIENTS:

If duplicate clients are created with '_1' appended to the original client name, you can perform the following steps to merge the duplicate client back into the original:

1. At a command prompt, navigate to the software installation path, log in to the CommServer, and run the following script:

```
qoperation execscript -sn QS_SetVMClient -si @sourceClient='<ClientName>_1' -si @destClient='<ClientName>'
```

where <ClientName> is the original client name and <ClientName>_1 is the duplicate.

This script reassigns all backup history from <ClientName>_1 to <ClientName>. This enables you to view backup history, and to generate Job Summary Reports with the **Include Protected VMs** option enabled.

2. Remove the duplicate clients:
 - a. In the CommCell Console, go to **Control Panel | User Preferences**.
 - b. Click the **Client Computer Filter** tab.
 - c. Select the **Show Virtual Server Discovered Clients** option.
 - d. Delete the duplicate clients from the CommCell Browser.

Troubleshooting - VMware Restore

[Backup](#) | [Restore](#)

The following section provides information on some of the troubleshooting scenarios related to restore:

RESTORE FAILURE

After Upgrading Proxy Server from 8.0 to 9.0	<p>Uninstall VDDK 1.1 on the proxy computer and re-install VDDK 1.2 before performing restore.</p> <p>*For 64-bit operating systems: From the VDDK x64 installation directory, unzip the vddk64 folder to a temporary location. Once unzipped, copy the bin, lib, and plugins folders to: C:\Program Files\VMware\VMware Virtual Disk Development Kit</p>
Browse operation yields a message stating "There are no items to show in this view"	The MediaAgent selected in the Browse Options dialog box may be installed on a Unix computer. Unix-based MediaAgents are not supported for direct restores. To resolve this issue, select a Windows-based MediaAgent.
When using IP Address of the Virtual Machine	Provide the name of the virtual machine you are restoring instead of the IP address.
File level restores cannot be performed	If the cluster size (allocation unit) on the disk that you backed up was less than 1024 bytes, the backup job completes with the error: "Metadata collection failed for Virtual machine [B_VM1_disk_Full], File and Volume browse will not be available for this VM. [Failed to open volume for metadata processing, Cluster size is not supported]". Increase the disk size to perform a file level restore on this disk.
File level restores complete with errors	Do not perform a restore of a SymLink file.
When using SAN transport mode if proxy is on Windows Server 2008	<p>Clear the read-only mode, by following the steps given below:</p> <ol style="list-style-type: none"> 1. Open Command Prompt 2. Run diskpart 3. List and select the SAN shared disks 4. Type attribute disk clear readonly <p>For Example:</p> <pre>C:\Users\Administrator> diskpart Microsoft DiskPart version 6.0.6002 Copyright (C) 1999-2007 Microsoft Corporation. On computer: COMPUTERNAME DISKPART> list disk Disk ### Status Size Free Dyn Gpt ----- Disk 1 Offline 78 GB 2048 KB <--- The share disk is usually reported as Offline DISKPART> select disk 1 Disk 1 is now the selected disk. DISKPART> detail disk ROCKET IMAGEFILE SCSI Disk Device Disk ID: 02872505 Type : iSCSI Bus : 0 Target : 0 LUN ID : 0 Read- only : Yes <--- Make sure of read-only mode Boot Disk : No Pagefile Disk : No Hibernation File Disk : No Crashdump Disk : No There are no volumes. DISKPART> attribute disk clear readonly Disk attributes cleared successfully.</pre>
VMX and VMDK files are on different datastores	<p>Ensure the following:</p> <ol style="list-style-type: none"> 1. Before performing the backup, ensure that VMX and VMDK files are on the same datastore. Create at least one VMX file on the datastore that contains the VMDK files. 2. If configured for ESX Server backups, ensure a folder exists on the secondary datastore(s) and that the folder is given the exact name of the virtual machine.
With "Duplicate GUID" error	Virtual machines with the same GUID may fail on restore.
VMware converter restores of backups done via the vStorage API	Backups performed via the vStorage API cannot be restored through the VMware converter. The converter does not understand the format of vStorage.

RESTORING INDEPENDENT DISK/PHYSICAL RDM RESULTS IN VM NOT POWERING ON

PROBLEM

If the virtual machine under backup includes independent disks/physical RDMs, these disks will be skipped. During a full restore the independent disk/Physical RDM also gets restored with OMB, and the virtual machine cannot be powered on.

SOLUTION

Prior to restore, manually delete the independent disks/Physical RDMs and then power on the virtual machine.

FILE-LEVEL RESTORES INTRODUCE A DIFFERENT FILE SIZE AND CHECKSUM

PROBLEM

If a file-level restore is performed from a backup that used VCB mounter, the restored files may have a different size and a checksum.

SOLUTION

1. From the **Backup Set Properties** dialog box, switch the backup mode to **vStorage**.
2. Re-run the backup.
3. Restore the files again.

SAN MODE RESTORES MAY BE SLOWER THAN NBD TRANSPORT MODE

PROBLEM

When attempting restoration of a VMware virtual machine from disk backup in SAN environments, the restore job slows down with the events in VMware client GUI displaying "clear lazy zero". This results in the performance of restores with SAN transport mode to be slower than in NBD transport mode.

SOLUTION

To resolve this issue, perform restores directly to ESX server instead of vCenter.

WHILE RESTORING VM AS TEMPLATE, REGISTERING TEMPLATE VM FAILS

PROBLEM

When attempting restoration of a VMware virtual machine as a Template virtual machine, registering the Template virtual machine might fail.

SOLUTION

Manually register the template virtual machine if the host is present under the folder.

RECOVERING DATA ASSOCIATED WITH DELETED CLIENTS AND STORAGE POLICIES

The following procedure describes the steps involved in recovering data associated with the following entities:

- Deleted Storage Policy
- Deleted Client, Agent, Backup Set or Instance

BEFORE YOU BEGIN

This procedure can be performed when the following are available:

- You have a Disaster Recovery Backup which contains information on the entity that you are trying to restore. For example, if you wish to recover a storage policy (and the data associated with the storage policy) that was accidentally deleted, you must have a copy of the disaster recovery backup which was performed before deleting the storage policy.
- Media containing the data you wish to recover is available and not overwritten.
- If a CommCell Migration license was available in the CommServe when the disaster recovery backup was performed, no additional licenses are required. If not, obtain the following licenses:
 - IP Address Change license
 - CommCell Migration licenseSee License Administration for more details.
- A standby computer, which will be used temporarily to build a CommServe.

RECOVERING DELETED DATA

1. Locate the latest Disaster Recovery Backup which contains the information on the entity (Storage Policy, Client, Agent, Backup Set or Instance) that you are trying to restore.
 - You can check the Phase 1 destination for the DR Set or use Restore by Jobs for CommServe DR Data to restore the data.

- If the job was pruned and you know the media containing the Disaster Recovery Backup, you can move the media in the **Overwrite Protect Media** Pool. See *Accessing Aged Data* for more information. You can then restore the appropriate DR Set associated with the job as described in *Restore by Jobs for CommServe DR Data*.
 - If the job is pruned and you do not know the media containing the Disaster Recovery Backup, you can do one of the following:
 - If you regularly run and have copies of the Data on Media and Aging Forecast report you can check them to see if the appropriate media is available.
2. On a standby computer, install the CommServe software. For more information on installing the CommServe, see *CommServe Deployment*.
 3. Restore the CommServe database using the CommServe Disaster Recovery Tool from the Disaster Recovery Backup described in Step 1. (See *Restore a Disaster Recovery Backup* for step-by-step instructions.)
 4. Verify and ensure that the **Bull Calypso Client Event Manager Bull Calypso Communications Service (EvMgrS)** is running.
 5. If you did not have a CommCell Migration license available in the CommServe when the disaster recovery backup was performed, apply the IP Address Change license and the CommCell Migration license on the standby CommServe. See *Activate Licenses* for step-by-step instructions.
 6. Export the data associated with the affected clients from the standby CommServe as described in *Export Data from the Source CommCell*.

When you start the Command Line Interface to capture data, use the name of the standby CommServe in the `commcell` argument.

7. Import the exported data to the main CommServe as described in *Import Data on the Destination CommCell*.

This will bring back the entity in the CommServe database and the entity will now be visible in the CommCell Browser. (Press F5 to refresh the CommCell Browser if the entity is not displayed after a successful merge.)
8. If you have additional data that was backed up after the disaster recovery backup and before the deletion of the entity, use the procedure described in *Import Metadata from a Tape or Optical Media* to obtain the necessary information.
9. You can now browse and restore the data from the appropriate entity.

As a precaution, mark media (tape and optical media) associated with the source CommCell as READ ONLY before performing a data recovery operation in the destination CommCell.

RESTORES FROM SNAP FAILS IN CASE OF READ-ONLY DISKS IN SAN MODE

In case of restores for SAN storage with any vendor type, you need to assign write permission to the disk on the Virtual Server proxy.

For Example:

If disk 12 is the disk you are trying to restore:

```
Z:\Henry\vmware>diskpart
select disk 12
DISKPART> disk attributes
attributes disk clear readonly
```

REVERT OPERATION IN PENDING STATE

SYMPTOM

The Revert job may go to Pending state with the following error:

```
File level revert is not supported for a Virtual Machine [%VMNAME%] which is on VMFS datastore. [Datastore(s) [%DatastoreName%] does not support file level revert.
```

CAUSE

The virtual machine may have NFS and iSCSI disks. Ensure that all the disks reside on the NFS data store. If the data does not reside on the NFS data store, you cannot perform the revert operation.

SOLUTION

Perform the conventional **Container Restore** to bring the data back to the point-in-time.

REVERT OPERATION FAILS

Before starting the revert operation, ensure that the Snapshot which you are using to perform the revert operation is not mounted.

FILE LEVEL RESTORE FAILS

SYMPTOM

While restoring files and folders from a virtual machine, you may get following error:

Volume name not found in the restore source path <source_path>, this is a volume restore.

CAUSE

This error appears when you have selected all files and folders in a virtual machine for the restore.

The file level restore of entire contents of the VM is not supported.

SOLUTION

If you want to restore entire contents of a VM, perform a full VM restore.

VIRTUAL MACHINE IS RESTORED SUCCESSFULLY BUT NOT ADDED TO THE CORRECT RESOURCE POOL

SYMPTOM

The virtual machine is not added to the correct resource pool after the restore.

CAUSE

This error appears in the following scenarios:

- You have performed the backup of the virtual machine using an ESX Server and the ESX Server is managed by a vCenter.
- You are restoring virtual machine to an ESX Server and the ESX Server is managed by a vCenter.

SOLUTION

If the ESX Server is managed by the a vCenter, do not configure the instance for the ESX server. Configure an instance for the vCenter and then perform the backup.

Similarly, before performing the restore, ensure that the instance is configured for the vCenter and not for the ESX server.

VIRTUAL MACHINE CLIENT NAMES ARE GETTING CREATED WITH '_1' APPENDED TO THE ORIGINAL CLIENT NAME

SYMPTOM

When viewing virtual machines in the Client Computers list, you may see duplicated client names (for example, <ClientName> and <ClientName>_1).

CAUSE

Virtual machine information is added to the database during the discovery phase of a backup. If information changes, a new client is discovered during the discovery phase.

Making changes to information such as the virtual machine fully qualified domain name (FQDN), the GUID, or the host name can cause this issue.

RESOLUTION

TO AVOID DUPLICATE CLIENT ENTRIES CAUSED BY DIFFERENT CLIENT OR HOST NAME:

Use the install software option for the client level in the CommCell Console (rather than interactively installing software to the virtual machine Guest Operating System).

TO CORRECT EXISTING CLIENTS:

If duplicate clients are created with '_1' appended to the original client name, you can perform the following steps to merge the duplicate client back into the original:

1. At a command prompt, navigate to the software installation path, log in to the CommServer, and run the following script:

```
operation execscript -sn QS_SetVMClient -si @sourceClient='<ClientName>_1' -si @destClient='<ClientName>'
```

where <ClientName> is the original client name and <ClientName>_1 is the duplicate.

This script reassigns all backup history from *<ClientName>_1* to *<ClientName>*. This enables you to view backup history, and to generate Job Summary Reports with the **Include Protected VMs** option enabled.

2. Remove the duplicate clients:
 - a. In the CommCell Console, go to **Control Panel | User Preferences**.
 - b. Click the **Client Computer Filter** tab.
 - c. Select the **Show Virtual Server Discovered Clients** option.
 - d. Delete the duplicate clients from the CommCell Browser.

FAILED TO START THE VIRTUAL MACHINE

SYMPTOM

When a virtual machine has been replicated in vSphere and backed up, and a full VM restore is performed from the backup, the following status message might be displayed in vCenter:

```
Failed to start the virtual machine.
```

CAUSE

If the **Power ON Virtual Machine After Restore** option is selected when performing a full VM restore, VMware attempts to power on the restored VM before disabling replication.

RESOLUTION

Clear the **Power ON Virtual Machine After Restore** option when initiating the restore. VMware automatically disables replication when the restore is completed, and you can power on the virtual machine manually.

CANNOT RESTORE FILES FROM A WINDOWS 2012 VIRTUAL MACHINE USING DEDUPLICATION

SYMPTOM

When restoring from a backup of a Windows 2012 virtual machine that has deduplication enabled, a file-level restore completes successfully but only creates stub files.

CAUSE

File-level restores are not supported for deduplicated volumes from a Windows 2012 VM.

RESOLUTION

To retrieve files from a backup for a Windows 2012 VM using deduplication, restore the disk that contains the file.

Alternatively, you can install a local file system agent on the Windows 2012 VM to enable file-level restores.

Online Help Topics

Agent Properties (General)

Use this dialog box to obtain or provide general and/or user account information for the agent.

Client Name

Displays the name of the client computer on which this Agent is installed.

Billing Department

Displays the name of the billing department, when the Agent is associated with a billing department.

iDataAgent

Displays the identity of the Agent that is installed on the client computer.

Installed date

Displays the date on which the Agent was installed or upgraded on the client computer.

Description

Use this field to enter a description about the entity. This description can include information about the entity's content, cautionary notes, etc.

Agent Properties (Version)

Use this dialog box to view the software version.

Version Information

Displays the Agent software version installed on the computer.

Agent Properties (Index)

Use this tab to establish the parameters for creating new index during backups by the agent.

Create new index on full backup

Specifies whether the system will create a new index with each full backup. If cleared, all backups build upon the index from the previous backup.

Convert to full backup on indexing failures

Specifies whether the system will create a new index if the index is irretrievable and whether, for agents that support non-full backup types, the system will convert the backup to a full backup. If cleared, the backup will fail if the index is irretrievable.

Security

Use this dialog box to:

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

Available Groups

Displays the names of the user groups that are not associated with this CommCell object.

Associated Groups

Displays the names of user groups that are associated with this CommCell object.

Activity Control

Use this dialog box to enable or disable data management and data recovery operations on a selected client or client computer group.

Enable Backup

Specifies whether data management operations will occur from this agent.

If cleared:

- The button **Enable after a Delay** appears to the right side of this operation activity. It allows to pre-determine the date and time at which data management operations will begin on the selected agent.
- Data management operations from this agent cannot be started and are skipped.
- Running and waiting data management operations for this agent run to completion.
- Stopped data management operations for this agent cannot be resumed until this option is enabled.
- Pending data management operations for this agent do not run until this option is enabled.

(Job states are indicated in the Job Controller.)

Enable Restore

Specifies whether data recovery operations will occur from this agent.

If cleared:

- The button **Enable after a Delay** appears to the right side of this operation activity. It allows to pre-determine the date and time at which data recovery operations will begin on the selected agent.
- Data recovery operations for this agent cannot be started and are skipped.
- Running data recovery operations for this agent run to completion.
- Pending data recovery operations for this agent do not run until data recovery operations are enabled.

(Job states are indicated in the Job Controller.)

Virtual Server Instance Properties (General)

VMware | Microsoft Virtual Servers/Hyper-V | Xen Server

Use this dialog box to define a new instance or view the properties of an existing instance.

Client Name

Displays the name of the client computer on which this instance is found.

iDataAgent

Displays the name of the iDataAgent to which this instance belongs.

Instance Name

- When you create a new instance, type the name of the instance.
- For an existing instance, the name of the instance is displayed.

Vendor Type

- When you create a new instance, select VMWare.
- For an existing instance, the vendor type is displayed.

VMware

When you create a new instance, the following options must be configured.

For an existing instance, the following options are displayed. You can modify them if required.

- **Virtual Center**

Select this option if this instance will be used for Virtual Center backups.

Host Name

Enter the name of the host computer in which the Virtual Center resides.

Configure Password

Click this button to configure the user name and password associated with the Virtual Center.

User Account

When you create a new instance, click the **Configure Password** button to establish the user name and password associated with the Virtual Center. Once

the credentials are configured, the user account will be displayed.

For an existing instance, the user account associated with the configured credentials is displayed.

- **ESX Servers**

Select this option if this instance will be used for ESX Server backups.

- **Configure ESX Servers**

Click this button to include the specific ESX Server to this instance.

NOTES ON USER CREDENTIALS

Users performing backup and restore operations must have credentials with sufficient permissions granted on the Virtual Center or ESX Server. **Datacenter Administrator** is an example of an account with acceptable privileges for these operations.

RECOMMENDED ROLES

For VCB environments, the **VCB Role** is sufficient for both backup and restore operations.

For vSphere VADP environments, the **VCB Role** can be used for backups, but require the following additional privileges for restore operations:

- Virtual Machine
- Resource
- Datastore

In both VCB and vSphere VADP environments, user-defined roles can be created provided the above-mentioned privileges are included in those roles.

Passwords for ESX Server or Virtual Center instances cannot have single-quote (') or double-quote (") characters.

CONFIGURING USER CREDENTIALS

To configure the user account associated with the **Virtual Center**, click the **Configure Password** button. Enter the credentials in the dialog displayed.

To configure the user account associated with the **ESX Server**:

1. Click the **Configure ESX Server** button. The ESX Server Configuration dialog box is displayed.
2. Click **Add/Edit/Remove** as desired.
3. If adding a new ESX server, the **Add ESX Server** dialog box appears. Enter the ESX Server name as well as the user name and password associated with the ESX Server (the password must be confirmed in the space provided). Click **OK**.
4. If editing an existing ESX Server, the **Add ESX Server** dialog box appears. Change the ESX Server name, user name, or password as desired (the password must be confirmed in the space provided). Click **OK**.

Security

Use this dialog box to:

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

Available Groups

Displays the names of the user groups that are not associated with this CommCell object.

Associated Groups

Displays the names of user groups that are associated with this CommCell object.

Automatically add New Users to the "View All" Group (CommServe Level Only)

Specifies that all new users will automatically be added to the `view All` user group. This group allows users to view all CommCell objects, regardless of the associations of the other user groups they belong to.

Require Authentication for Agent Installation (CommServe Level Only)

Enable this option to prompt users for authentication for any agent installation on the CommCell.

Client Owners

Specify a user to allow CommCell user access to this client computer through CommCell Console and allow permission to run Web Console browse and restore operations on this client. Each user must be listed in a separate line.

Subclient Properties (General)

VMware | Microsoft Virtual Servers/Hyper-V | Xen Server

Use this tab to add, modify or view general information on the selected subclient.

Client Name

Displays the name of the Client computer to which this subclient belongs.

iDataAgent

Displays the name of the Agent to which this subclient belongs.

Backup Set

Displays the name of the Backup Set to which this subclient belongs.

Subclient Name

Displays the name of this Subclient. You can use this space to enter or modify the name of the subclient.

Number of Data Readers

Use this space to specify the number of simultaneous backup data streams allowed for this subclient. For best performance, this should be set no higher than the number of physical drives that hold this subclient's data, except for specialized hardware such as RAID. For more information, see *Automatic File System Multi-Streaming* in Books Online.

Backup Type

Select the backup type for this subclient.

- **Disk Level**

Select this option if this subclient will be used to back up virtual server disks.

- **Volume Level**

Select this option if this subclient will be used to back up individual volumes.

- **File Level**

Select this option if this subclient will be used to back up individual files and/or folders.

Use Proxy

When selected, additional load balancing will be provided for operations occurring between multiple proxy computers.

Transport Mode for VMware

Select the mode on which you desire to backup this subclient. Each available mode represents a different system environment scenario.

Description

Use this field to enter a description about the entity. This description can include information about the entity's content, cautionary notes, etc.

Subclient Properties (Filter)

Volume Level | File Level

Use this dialog box to add or edit data paths to be filtered out of the subclient content and to specify exceptions to the exclusion list.

Include Global Filters

Displays a list of options for enabling or disabling Global Filters which can be used in addition to the filters set at the subclient level to exclude data from data protection operations for this subclient. To change the setting, select one of the following values from the list:

- **OFF** - Disables Global Filters for this subclient.
- **ON** - Enables Global Filters for this subclient.
- **Use Cell Level Policy** - Enables or disables Global Filters for this subclient depending on whether the **Use Global Filters on All Subclients** option is enabled or disabled on the Control Panel (Global Filters) dialog. If the **Use Global Filters on All Subclients** option is enabled, then the **Use Cell Level Policy** option will enable Global Filters on this subclient. If the **Use Global Filters on All Subclients** option is disabled, then the **Use Cell Level Policy** option will disable Global Filters on this subclient.

Exclude these files/folders/directories/patterns

Displays entries that will be excluded from data protection operations for this subclient.

Pre/Post Process

VMware | Microsoft Virtual Servers/Hyper-V | Xen Server

Use the Pre/Post Process tab to add, modify or view Pre/Post processes for the selected subclient.

Run Post Scan Process for all attempts

Specifies whether this process will execute for all attempts to run the phase. Selecting this option will execute the post scan for all attempts to run the phase, including situations where the job phase is interrupted, suspended, or fails. Otherwise, when the checkbox is cleared the post scan process will only execute for successful, killed, or failed jobs.

Pre Backup process

Displays the name/path of the backup that you want to run before the respective phase. You can use this space to enter a path for the backup process that will execute before this phase on the subclient, or you can use the **Browse** button to search for and select the name/path of the process. The system allows the use of spaces in the name/path, provided they begin with an opening quotation mark and end with a closing quotation mark.



If you are using the Linux File System iDataAgent to back up VMware virtual machines, you must create a PreBackup process script that will allow the system to create the required REDO files. See "Create Pre/Post Backup Process Scripts to Back up VMware Virtual Machines Using the Linux File System iDataAgent" in Books Online for step-by-step instructions.

Post Backup process

Displays the name/path of the backup process that you want to run after the respective phase. You can use this space to enter a path for the backup process that will execute after this phase on the subclient, or you can use the **Browse** button to search for and select the name/path of the process. The system allows the use of spaces in the name/path, provided they begin with an opening quotation mark and end with a closing quotation mark.



If you are using the Linux File System iDataAgent to back up VMware virtual machines, you must create a PostBackup process script that will allow the system to post data back to the virtual machines. See "Create Pre/Post Backup Process Scripts to Back up VMware Virtual Machines Using the Linux File System iDataAgent" in Books Online for step-by-step instructions.

Run Post Backup Process for all attempts

Specifies whether this process will execute for all attempts to run the phase. Selecting this option will execute the post backup process for all attempts to run the phase, including situations where the job phase is interrupted, suspended, or fails. Otherwise, when the checkbox is cleared the specified process will only execute for successful, killed, or failed jobs.

Run As

Displays either the Local System Account, or for added security, another account as having permission to run these commands.

Change

Click to add or modify the account that has permission to run these commands.

Browse

Click to search for and select the Pre/Post Process.

[Back to Top](#)

Security

Use this dialog box to:

- Associate this object with a user group.

- Disassociate this object from a user group.

Available Groups

Displays the names of the user groups that are not associated with this CommCell object.

Associated Groups

Displays the names of user groups that are associated with this CommCell object.

Encryption

Use this dialog box to select the data encryption options for the selected content. When accessing this dialog box from the Subclient Properties Encryption tab, this setting applies only to the selected subclient content for operations run from the CommCell Console. When accessing this dialog box from the Instance Properties Encryption tab, this setting applies only to third-party Command Line operations. The functionality is not propagated to the Subclient Properties Encryption tabs.

None

When selected, no encryption will take place during a data protection operations.

Media Only (MediaAgent Side)

When selected, for data protection operations, data is transmitted without encryption and then encrypted prior to storage. During data recovery operations, data is decrypted by the client.

When using this setting in conjunction with the client property **With a Pass-Phrase**, you will be required to provide a pass-phrase for data recovery operations unless you export the client pass-phrase to the destination client(s). When using pass-phrase security for third-party Command Line operations or DataArchiver Agents stub recovery operations, you must export the pass- phrase to the destination client.

Network and Media (Agent Side)

When selected, for data protection operations, data is encrypted before transmission and is stored encrypted on the media. During data recovery operations, data is decrypted by the client.

When using this setting in conjunction with the client property **With a Pass-Phrase**, you will be required to provide a pass-phrase for data recovery operations unless you export the client pass-phrase to the destination clients.

Network Only (Agent Encrypts, MediaAgent Decrypts)

When selected, for data protection operations, data is encrypted for transmission and then decrypted prior to storage on the media. During data recovery operations, data is encrypted by the MediaAgent and then decrypted in the client.

When using this setting in conjunction with the client property **With a Pass-Phrase**, you will not be required to provide a pass-phrase for data recovery operations.

Script Preview

Click to display the backup script, based on the current subclient configuration, that will be submitted to RMAN when backups are performed for the selected Oracle subclient.

Storage Device

Use this dialog box to establish the storage device related settings on the selected subclient.

The following tabs are displayed:

- (Data) Storage Policy
- Data Transfer Option
- Deduplication

STORAGE POLICY

Use this tab to select or view storage policy settings on the selected subclient.

Data/Database/Transaction Log Storage Policy

Displays the storage policy to which this subclient is currently associated. To associate a storage policy to a new subclient or to change the storage policy associated with an existing subclient, click one in the list.

Incremental Storage Policy

Displays the name of the Incremental Storage Policy associated with this subclient, if the storage policy has the Incremental Storage Policy option enabled.

Data Paths

Click to view or modify the data paths associated with the primary storage policy copy of the selected storage policy.

Create Storage Policy

Click to launch the Create a Storage Policy wizard. Once the storage policy has been created, it will be displayed in the list of storage policies to which the selected subclient can be associated.

DATA TRANSFER OPTION

Use this tab to establish the options for data transfer.

Software Compression

Indicates whether software compression for the subclient or instance is enabled or disabled, and if enabled whether it is performed on the client or MediaAgent computer.

- **On Client**

Click to enable software compression on the client computer.

- **On MediaAgent**

Click to enable software compression on the MediaAgent computer.

- **Off**

Click to disable software compression.

Note that hardware compression has priority over the software compression. Hardware compression is established in the Data Path Properties dialog box. The above software compression option will take effect when the data path is associated with a disk library, or when hardware compression is disabled in the data path associated with tape libraries.



If the subclient is associated with a storage policy copy that is deduplicated, then the compression settings on the storage policy copy takes precedence. See Copy Properties (Deduplication) - Advanced tab for compression settings on deduplicated storage policy copy.

Resource Tuning

Indicates the processes used by the client to transfer data based and whether bandwidth throttling is enabled or not.

- **Network Agents**

Specifies the number of data pipes/processes that the client uses to transfer data over a network. Increasing this value may provide better throughput if the network and the network configuration in your environment can support it. On non-UNIX computers, the default value is 2 and a maximum of 4 can be established if necessary. On UNIX computers the default value is 1 and a maximum of 2 can be established if necessary.

- **Throttle Network Bandwidth (MB/HR)**

Specifies whether the backup throughput is controlled or not. (By default this option is not selected and therefore the throughput is not controlled). When selected, use the space to specify a value for the throughput. By default, this is set to 500. The minimum value is 1 and there is no limit to the maximum value. (In this case the backup throughput will be restricted to the maximum bandwidth on the network.)

Use this option to set the backup throughput, based on the network bandwidth in your environment. Use this option to reduce the backup throughput, so that the entire network bandwidth is not consumed, especially in slow links. Increasing this value will end up consuming the bandwidth with the maximum throughput limited to the network bandwidth capability.

Note that throttling is done on a per Network Agent basis.

DEDUPLICATION

Use this tab to establish the options for deduplication on the subclient. It indicates whether deduplication for the subclient is enabled or disabled, and if enabled whether the signature generation (a component of deduplication) is performed on the client or MediaAgent computer.

Note that deduplication is supported on disk storage devices. So the deduplication options are applicable only if the subclient is associated with a Storage Policy containing disk storage.

On Client

Click to enable signature generation on the client computer.

On MediaAgent

Click to enable signature generation on the MediaAgent computer.

[Back to Top](#)

Subclient Properties (Auto Discover)

Use this tab to view and establish automatically discovered virtual servers.

Add Guest Host Name

Lists the names of the guest hosts currently discovered. Click **Remove** to delete a guest host.

Guest Host Name

Use this space to specify the name of a guest host. Click **Add** to add the guest host to the list of discovered guest hosts.

Activity Control

Use this dialog box to enable or disable backups on a selected subclient.

If data management and data recovery operations are disabled at the client computer group or client level, then these operations below these levels will be disabled. If data management / data recovery operations are enabled at the client computer group or client level, then these operations below these levels will be enabled.

Enable Backups

Specifies whether Backups will occur from this subclient.

If cleared:

- The button **Enable after a Delay** appears to the right side of this operation activity. It allows to pre-determine the date and time at which data management operations will begin on the selected subclient.
 - Backup operations from this subclient cannot be started and are skipped.
 - Running and waiting data management operations for this subclient run to completion.
 - Stopped data management operations for this subclient cannot be resumed until this option is enabled.
 - Pending data management operations for this subclient do not run until this option is enabled.
-

Subclient Properties (Content)

Use this tab to define the contents of a new subclient or to change the content of an existing subclient. Items included in subclient content are backed up by the subclient.

Virtual Machine

Displays a list of virtual machines included as content for this subclient.

Configure

Click this button to configure existing or additional virtual machines for this subclient.

Subclient Properties (SnapProtect Operations)

Use this tab to enable SnapProtect backup and to establish the options for SnapProtect backup operation, for the selected subclient.

SnapProtect

Specifies whether SnapProtect backup for the selected subclient is enabled or disabled.

Available Snap Engines

Lists the available snapshot engine vendors. To select a snapshot engine vendor, click the pull-down arrow and click the one that you want to use.

• Enter Array Credentials

Click to access the Array Management dialog box, which enables you to add or modify access information for a array.

Virtual Center / ESX Server

Specify the name of the ESX Server in the Virtual Center.

• Select ESX server for snap mount

Click to access the **Browse for ESX server** dialog box, which enables you to select a different ESX server.

Host

Enter an IP address of the Host computer.

Application aware backup for granular recovery

Select to enable snap mining for this subclient.

Truncate ExDB Logs

Select to prune Exchange Database logs that are no longer needed and to prevent them from growing to an unmanageable size.

Use source if proxy is unreachable

Select to use the source machine for SnapProtect backup if the proxy client is not available.

Allow multiple data readers within a drive or mount point

Specifies whether multiple data reads are allowed for a single physical drive or mount point during backups on this subclient. It is recommended that this option is selected for specialized hardware such as RAID, or possibly in the case of spanned volumes. For more information, see Automatic File System Multi-Streaming in Books Online.

Number of Data Readers for Backup Copy

Use this space to specify the number of simultaneous backup data streams allowed for this subclient, during the movement to media operation. For best performance, this should be set no higher than the number of physical drives that hold this subclient's data, except for specialized hardware such as RAID.

[Back to Top](#)

Backup Set Properties (General)

VMware | Microsoft Virtual Servers/Hyper-V | Xen Server

Use this dialog box to define a New Backup Set or view and modify the properties of an existing Backup Set.

Client Computer

Displays the name of the Client computer to which this Backup Set belongs.

iDataAgent

Displays the name of the iDataAgent to which this Backup Set belongs.

Instance Name

Displays the name of the instance to which this Backup Set belongs.

The instance name can be modified for an existing Backup Set from the properties dialog box.

Backup Set / New Backup Set Name

When creating a new Backup Set, type the name of the Backup Set.

For an existing Backup Set, the name of the Backup Set is displayed.

Automatically add new VMs that don't qualify for membership in any of the subclients

Select this option to automatically add virtual machines that don't qualify for membership in any of the existing subclients to this backup set.

Make this the default Backup Set

Specifies whether the selected backup set is the default or a user-defined backup. Selecting this option asserts this backup set as the default. The previously designated default backup set is automatically reallocated as a user-defined backup set. Clearing this option, is accomplished by opening the Backup Set Properties dialog box of some other backup set and clicking this option there, thereby designating that backup set as the default.

Rule based discovery

Select this option to configure rule based discovery.

- **Match guest host names by regular expression**

Select this option to configure discovery rules based on guest host names for the existing subclients of this backup set. Click **Configure** to setup the configuration for this selection.

- **ESX Server Affinity**

Select this option to configure discovery rules based on ESX Servers for existing subclients of this backup set. Click **Configure** to setup the configuration for this selection.

- **Data Store Affinity**

Select this option to configure discovery rules based on data stores affinity for existing subclients of this backup set. Click **Configure** to setup the configuration for this selection.

- **Resource Pool Affinity**

Select this option to configure discovery rules based on resource pool affinity for existing subclients of this backup set. Click **Configure** to setup the configuration for this selection.

- **OS Group**

Select this option to discover virtual machines based on the operating systems they are running on.

Select the appropriate operating systems from the displayed list of supported Windows and Unix operating systems.

vcbmounter

Select this option when using older version of vmware. A proxy computer will be used when this option is selected.

vStorage

Select this option when using newer version of vmware. No proxy computer is needed.

Backup Set Properties (Guest Hosts Configuration)

VMware | Microsoft Virtual Servers/Hyper-V | Xen Server

Use this dialog to configure one or more guest hosts for the selected backup set.

Server Name/Virtual Center Name

This column displays the name of the ESX Server/Virtual Center configured for the selected backup set.

Virtual Machine

This column displays the virtual machine associated with the selected backup set.

Subclient Name

This column displays the subclient associated with the selected Virtual Machine.

Discover

Click this button to discover new guest hosts.

Change all selected Guest Host to

Use the drop-down menu to change selected server's subclient association. Select **Do Not Backup** if you wish to exclude the selected servers from data protection operations.

Apply

Click this button to apply the changes made to the guest host configuration.

Security

Use this dialog box to:

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

Available Groups

Displays the names of the user groups that are not associated with this CommCell object.

Associated Groups

Displays the names of user groups that are associated with this CommCell object.

Automatically add New Users to the "View All" Group (CommServe Level Only)

Specifies that all new users will automatically be added to the `view All` user group. This group allows users to view all CommCell objects, regardless of the associations of the other user groups they belong to.

Require Authentication for Agent Installation (CommServe Level Only)

Enable this option to prompt users for authentication for any agent installation on the CommCell.

Client Owners

Specify a user to allow CommCell user access to this client computer through CommCell Console and allow permission to run Web Console browse and restore operations on this client. Each user must be listed in a separate line.

Advanced Backup Options (Data)

You can select advanced backup/archive data options for the operation. Note that all the options described in this help may not be available and only the options displayed in the dialog box are applicable to the agent or enabler for which the information is being displayed.

Catalog

Use this section to select index cache sharing and granular restartability options for the job. Note that these options are not applicable for the disaster recovery data protection.

- **Use shared profile if present with transaction logging**

Select this option if you wish to use the shared index cache profile of the MediaAgent (if configured) to save a copy of the job's index cache for index cache sharing *with* transaction logging feature for granular job restartability.

Shared profile with transaction logging provides job restartability in failover situations. This is a recommended configuration.

- **Use shared profile if present without transaction logging**

Select this option if you wish to use the shared index cache profile of the MediaAgent (if configured) to save a copy of the index cache for index cache sharing *without* transaction logging feature for granular job restartability. This is not a recommended configuration.

- **Use transaction logging**

Select this option to use transaction logging feature to provide granular job restartability (without using shared profile for index cache sharing).

- **None**

Select this option to use neither shared profile nor transaction logging.

Create Backup Copy immediately

Select this option to perform inline backup copy operation.

- **Enable Granular Recovery**

This option is selected by default. It enables recovery of files and folders from the backup copy.

Enable Granular Recovery

Click to perform Granular Recovery. When clicked, allows recovery of Individual files and folders from the primary copy.

Advanced Backup Options (Startup)

You can specify advanced backup/archive/content indexing startup options for the operation. For more information, select one of the following:

Set Priority

- **Use Default Priority**

Select this option to use default priority for the job.

- **Change Priority**

Use this option to modify the priority of a job between 0 (highest priority) and 999 (lowest priority). The Job Manager will use the priority setting when allocating the required resources.

Start job in suspended state

Specifies that this job will start in the Job Controller in a suspended state and cannot run until the job is manually resumed using the **Resume** option.

Description

Use this field to enter a description about the entity. This description can include information about the entity's content, cautionary notes, etc.

Advanced Backup Options (Media)

You can select advanced backup/archive media options for the operation. Note that all the options described in this help may not be available and only the options displayed in the dialog box are applicable to the agent or enabler for which the information is being displayed.

Start New Media

This option starts the backup/migration/archive operation on a new media, which causes the following to occur:

- If removable media is used, the current active media is marked as Appendable and a new media is used for the backup/migration/archive.
- If disk media is used, a new volume folder is created for the backup/migration/archive.

If cleared, the operation automatically uses the current active media or volume.



When using the Informix iDataAgent, if you are using tape and do all of the following, the tape will be marked full for both the data storage policy and the logical logs storage policy, and new media will be made available for both the logs and data even though you are backing up just the logical logs:

- Define a subclient with different storage policies for the logical logs and the data
- Prepare to back up just the logical logs
- Select the **Start new media** option before you start the backup

Mark Media Full after Successful Operation

This option marks media full, 2 minutes after the completion of the backup/archive operation. If any jobs are initiated within the 2 minutes, they are allowed to write to the media. If the job was associated with the prior media, new media (such as a new tape) will be used for subsequent jobs. (Applies to all backup/archive types.)

Allow other Schedule to use Media Set

This option allows jobs that are part of a schedule policy or schedule and using a specific storage policy to start a new media and also prevent other jobs from writing to the set of media. It is available only when the **Start New Media** and **Mark Media Full** options are enabled, and can be used in the following situations:

- When one scheduled job initiates several jobs and you only want to start new media on the first job.
- When you want to target specific backups to a media, or a set of media if multiple streams are used.

Override Job Retention

- **Infinite:** Select this option to retain this job indefinitely.
- **Number of Days:** Select this option to prune this job after the number of days specified.
- **Storage Policy Default:** Select this option to apply the retention rules of the associated storage policy, which is the default option.

Advanced Backup Options (Data Path)

Select the data path to be used for the backup/archive operation.

Ensure that the Library, MediaAgent, Drive Pool, and the Drive selected for this operation is available online and is a part of the associated Storage Policy.

Use Library

Specifies the name of the library that will be used to perform the backup operation. Use this option when you wish to backup to a specific library.

Use MediaAgent

Specifies the name of the MediaAgent that will be used to perform the backup operation. If necessary, you can change the name of the MediaAgent.

For example, if the library is shared and you wish to use a specific MediaAgent (instead of the system selected MediaAgent, or a MediaAgent which may be idle, or less critical) or if you know that the library attached to the specified MediaAgent.

Use Drive Pool

Specifies the name of the Drive Pool that will be used to perform the backup operation. Use this option when you wish to backup using a specific Drive Pool.

Use Drive

Specifies the name of the Drive that will be used to perform the backup operation. Use this option when you wish to backup using a specific Drive from the selected Drive Pool.

Vault Tracking

Select options to export and track media, using Vault Tracker.



Vault Tracking Options will be displayed only when a Vault Tracker license is available in the CommServe.

Vault Tracking options are only applicable for data protection operations using a storage policy associated with a library containing removable media (e.g., tape, optical or stand-alone.)

Export media after the job finishes

Specifies the media used by the data protection operation and media with the specific Media Status (if specified) will be exported and tracked by Vault Tracker.

Exclude Media Not Copied

When selected, allows you to exclude media with jobs that have to be copied.

Media Status

- **All**
Click to select all media. Clear this option to select media with a specific status.
- **Active**
Click to select media with its status marked as active.
- **Full**
Click to select media with its status marked as full.
- **Overwrite Protected**
Click to select media with its status marked as read-only .
- **Bad**
Click to select media with its status marked as bad.

Export Location

Specifies the destination location and lists the stationary locations entered using the **Export Location Details** dialog box.

Track Transit

Specifies that transit information must be tracked, and lists the transit locations entered using the **Export Location Details** dialog box.

Use Virtual Mail Slots

Specifies the exported media is stored within the library in the virtual mail slots defined in the **Library Properties (Media)** dialog box.

Filter Media By Retention

Specifies that the system must automatically filter media based on whether the media has extended retention jobs or not.

- **Media with extended retention job(s)**
Specifies that media with at least one extended retention job will be exported.
- **Media with no extended retention job(s)**
Specifies that media with no extended retention jobs will be exported.

Alert

Use this tab to configure an alert for a schedule policy.

Configure Alert

- **Alert**
The currently configured Alert.
- **Add/Modify Alert**
When clicked, opens the Alert Wizard to configure alerts for this operation.
- **Delete Alert**
When clicked, deletes any existing alerts that are already configured.

Browse Options

Use this dialog box to select the basic options under which you want to browse the data secured by a data protection operation in the selected client, agent, instance/partition, or backup set level depending on the functionality of a given agent. The system provides several browse options. These options influence the behavior of the browse operation in the subsequent Browse window.

Browse the Latest Data

Specifies whether the browse operation must display data from the most recent data protection operation (from now back to the last full backup) in the Browse window. (This option is selected by default.)

Specify Browse Time

Specifies whether the browse operation must display data, up to the specified date and time in the Browse window. Data secured after the specified date and time will be omitted in this window.

- **Browse Data Before/Browse to Point-In-Time**

Specifies the date and time that must be used to display the data in the Browse window.

- **Time Zone**

Lists the time zones. To change the time zone, click one in the list.

Client Computer

Lists the client computers available for the browse operation. This option is available only when the browse operation is performed either from the client node in the CommCell Browser, **Browse & Recover** option from the **Tools** menu in the CommCell Console or when the **New Browse** button is clicked in the Browse window.

Use MediaAgent

Specifies the name of MediaAgent that will be used to perform the browse (and restore) operation. If necessary, you can change the name of the MediaAgent. For example, if the library is shared and you wish to use a specific MediaAgent (instead of the system selected MediaAgent, or a MediaAgent which may be idle, or less critical) or if you know that the media containing the data you wish to restore is available in the library attached to the specified MediaAgent.

If the media containing the data is not available in the tape/optical library attached to the MediaAgent, the system will automatically prompt you to insert the appropriate media. In the case of a disk library, the operation will fail if the requested data is not available in the disk library attached to the specified MediaAgent.

Show Deleted Items

Specifies whether the operation will browse all the data (including deleted items) secured by all data protection operations for the selected backup set as of the specified browse items.

Page Size

Use the space to enter the default number of objects to be displayed in the Browse window. This field reverts back to the default setting after use. In the Browse window, you can browse through each page list by selecting the appropriate page number.

Type of Intended Restore (Virtual Server iDataAgent)

This section provides the facility to select the type of restore you want to perform following the Browse operation.

- **Individual Files/Folders**

Select this option if you intend on performing a file/folder-level restore.

- **Entire Volume**

Select this option if you intend on performing a volume-level restore.

- **Container Restore**

Select this option if you intend on performing a disk-level restore.

[Back to Top](#)

Advanced Browse Options

Use this dialog to access additional browse capabilities for the browse window.

Include Metadata

Specifies that the browse operation results must also include data containing metadata information.

Exclude Data Before

Specifies that the browse operation must return data that has been backed up after the specified date. Also provides a space for you to type the date and time.

- **Time Zone**

Lists the time zones. To change the time zone, click one in the list.

Browse from copy precedence

When selected, the system retrieves the data from the storage policy copy with the specified copy precedence number. If data does not exist in the specified copy, the operation fails even if the data exists in another copy of the same storage policy.

When cleared, (or by default) the system retrieves data from the storage policy copy with the lowest copy precedence. If the data was pruned from the primary copy, the system automatically retrieves the data from the other copies of the storage policy starting with the copy with the lowest copy precedence and proceeding through the copies with higher copy precedence. Once the data is found, it is retrieved, and no further copies are checked.

List Media

Use this dialog box to select the options that must be used to predict the required media.

Media For the Latest Data

Specifies whether the list media operation must display media associated with the most recent data protection operation (from now back to the last full backup). This option is selected by default.

Specify Time Range

Specifies whether the whether the list media operation must display media associated with the specified date and time.

- **Data Before/Time Range**

Specifies the date and time that must be used to display media.

- **Time Zone**

Lists the time zones. To change the time zone, click one in the list.

Advanced

Click to access additional options.

Restore Options (General)

VMware | Microsoft Virtual Servers/Hyper-V | Xen Server

Use this dialog box to control how the specified data is to be restored. For VMWare, Virtual Machine Level, Volume and File Level restores are available.

VIRTUAL MACHINE RESTORE OPTIONS

- **Virtual Machine**

The name of the virtual machine to be restored. Selecting this option also allows selecting additional **Virtual Machine Restore Options**.

- **Virtual Machine Files**

Select this option to restore the selected content as an entire virtual machine disk file.

- **Destination Client/Computer/Server**

Displays the name of the client computer to which the selected data will be restored. To change the destination computer, select one from the list. The list includes clients:

- Which are established as clients within the CommCell.
- Clients with Operating Systems that support the cross-platform restore operation from this client. (Refer to the topic titled **Cross Platform Restore Operations** in Books Online for additional information.)

By default, data is restored to the same computer from which it was backed up.

- **vCenter**

The name of the vCenter in which the data will be restored. Click **Change** to select a vCenter.

- **Restore in place**

Select this option to restore the virtual machine or virtual machine files to the same destination from which it was backed up.

- **VM and Disk**

The name of the virtual machines and disks to be restored.

- **Rename VM as**

Use this space to enter a new name for the virtual machine if you do not wish to continue using the virtual machine's existing name.

- **ESX Server**

Select the ESX Server to which the data will be restored.

- **Data Store**

Select the data store to which the data will be restored.

- **Add to Resource Pool/vApp**

Select a Resource Pool or vApp if you wish to include the restored virtual machine within a specific Resource Pool or vApp.

- **Virtual Machine Configuration**

- **Power ON Virtual Machine After Restore**

Click to power ON Virtual Machine after restore completes.

- **Unconditionally Overwrite VM with the same name**

Select this option to overwrite an existing virtual machine with the same name.

- **Disk Provisioning**

Select a provisioning type. By default, **Auto** is selected.

Auto: same as the source virtual machine, chosen during backup.

Thick: Thick provisioning allocates and commits full size to the data store.

Thin: Thin provisioning allocates full size and commits only as much storage space as the disk needs for its initial operation.

- **Transport Mode**

Select a transport mode. By default, **Auto** is selected.

SAN

Hot Add

NAS

NBD SSL

NBD

VOLUME-LEVEL RESTORE OPTIONS

Restore As (Volume-Level Restore)

- **Physical Volume**

Select this option to restore the selected content as a complete physical volume.

- **Virtual Hard Disk file**

Select this option to restore the selected content as a virtual hard disk file.

- **VMDK Files**

Select this option to restore the selected content as an entire virtual machine disk file.

- **Destination Client/Computer/Server**

Displays the name of the client computer to which the selected data will be restored. To change the destination computer, select one from the list. The list includes clients:

- Which are established as clients within the CommCell.

- Clients with Operating Systems that support the cross-platform restore operation from this client. (Refer to the topic titled **Cross Platform Restore Operations** in Books Online for additional information.)

By default, data is restored to the same computer from which it was backed up.

FILE-LEVEL RESTORE OPTIONS

Restore ACLs (File-Level Restore)

Specifies that the backed up Access Control Lists (ACLs) will be applied to files that have been restored without their ACLs, or to the existing files at the restore destination. Clear this option to restore without ACLs and consequently the data assumes the ACLs of the parent directory.

Unconditional Overwrite

Specifies whether the restored data is unconditionally written to the specified location.

If cleared:

- Files/directories whose names are the same as those in the restore path and where the backed up data is as old or older are not restored.
- files/directories whose names are the same as those in the restore path and where the backed up data is newer are restored.
- Files/directories whose names are different from those in the restore path are restored.



Avoid using this option whenever you run a Restore by Jobs. Otherwise, an older version of the files might be retrieved.

- **Destination Client/Computer/Server**

Displays the name of the client computer to which the selected data will be restored. To change the destination computer, select one from the list. The list includes clients:

- Which are established as clients within the CommCell.
- Clients with Operating Systems that support the cross-platform restore operation from this client. (Refer to the topic titled **Cross Platform Restore Operations** in Books Online for additional information.)

By default, data is restored to the same computer from which it was backed up.

- **Destination folder**

Use this space to enter the path on the destination computer to which the data will be restored. If the specified path does not exist, it will be created during the restore process.

On Windows computers, UNC path of a shared drive can be specified. (e.g., \\servername\sharename\directory). Click the **Advanced** button, and enter the Windows logon information for this shared drive in the **Advanced Restore Options** dialog box.

- **Browse**

Click to select a path for the **Destination folder**.

Preserve Source Paths

Specifies the action taken when the data is restored to another destination. The options are available only when the **Restore to Same Paths** option is cleared.

- **Preserve n levels from end of source path**

Specifies whether the restore operation must keep the specified number of levels at the end of source path of the selected data. For example, if the selected data is Documents\Memos\Sales, this option is set to 1 and the **Destination folder** is temp, the data will be restored as \temp\Sales.



However, if this option is used to restore a file in a path that contains any valid wildcard expression (e.g., Documents**\Sales), the default value of 1 for the option cannot be changed. Therefore, only the last found occurrence of the file is restored, and the restores for all the other occurrences of the file are overwritten.

- **Remove n levels from beginning of source path**

Specifies whether the restore operation must remove the specified number of levels from the beginning of source path of the selected data. For example, if the selected data is Documents\Memos\Sales, this option is set to 1 and the **Destination folder** is \temp, the data will be restored as \temp\Sales. (This option becomes available only when the **Restore to Same Paths** option is cleared.)

Save As Script

Click to open the Save As Script dialog, which allows you to save this operation and the selected options as a script file (in XML format). The script can later be executed from the Command Line Interface using `qoperation execute` command.

When you save an operation as a script, each option in the dialog will have a corresponding xml parameter in the script file. When executing the script, you can modify the value for any of these XML parameters as per need.

To view the XML values for each of the options in the dialog, see the following:

Command Line XML Options

[Back to Top](#)

Advanced Restore/Recover/Retrieve Options (General)

Use this dialog box to access additional restore/recover/retrieve options. Note that all the options described in this help may not be available and only the options displayed in the dialog box are applicable to the component installed on the client.

Automatically Detect Regular Expressions

Specifies whether to use regular expressions in defining the **Source Paths**. This does not apply to **Filter Paths**. When selected, the system will recognize supported regular expressions (wildcards) in the specified **Source Path**. When cleared, the system will read the specified **Source Path** as literal.

Skip Errors and Continue

For index-based agents, this advanced restore/recover/retrieve option enables a restore/recover/retrieve job to continue despite media errors. This option also provides an output file that lists the full path names of the files that failed to restore/recover/retrieve.

Impersonate (Windows, NetWare) User

Specifies whether to submit the Windows or NetWare logon information for the current restore/recover operation. This information is needed only if you intend to restore/recover data to a mapped network drive, shared network drive, or directory to which you have no write, create, or change privileges, or for restoring to applications.

Depending on the agent, there are as many as three separate User Impersonation accounts: this one for the function described, one for Pre and Post backup/migrate processes, and one for Pre and Post restore/recover processes. Each account operates independently of the other(s).

Advanced Restore Options (Copy Precedence)

Choose the copy from which you wish to recover or retrieve. Select from the following options:

Restore/recover from copy precedence

When selected, the system retrieves the data from the storage policy copy with the specified copy precedence number. If data does not exist in the specified copy, the operation fails even if the data exists in another copy of the same storage policy.

When cleared, (or by default) the system retrieves data from the storage policy copy with the lowest copy precedence. If the data was pruned from the primary copy, the system automatically retrieves the data from the other copies of the storage policy starting with the copy with the lowest copy precedence and proceeding through the copies with higher copy precedence. Once the data is found, it is retrieved, and no further copies are checked.

Advanced Restore Options (Data Path)

Select the data path for the restore/recovery operation. You can specify the MediaAgent, Library, Drive Pool, and Drive from which the restore operation must be performed.

Use MediaAgent

Specifies the name of the MediaAgent that will be used to perform the restore operation. If necessary, you can change the name of the MediaAgent.

For example, if the library is shared and you wish to use a specific MediaAgent (instead of the system selected MediaAgent, or a MediaAgent which may be idle, or less critical) or if you know that the media containing the data you wish to restore is available in the library attached to the specified MediaAgent.

If the media containing the data is not available in the tape/optical library attached to the MediaAgent, the system will automatically prompt you to insert the appropriate media. In the case of a disk library, the operation will fail if the requested data is not available in the disk library attached to the specified MediaAgent.

Use Library

Specifies the name of the library that will be used to perform the restore operation. Use this option when you wish to restore using a specific library.

For example, if you know that the media containing the data you wish to restore is available in a specific library.

Use Drive Pool

Specifies the name of the Drive Pool that will be used to perform the restore operation. Use this option when you wish to restore using a specific Drive Pool.

To restore NAS data, select the drive pool type that was used to perform the backup, *i.e.*, if a drive pool associated with an NDMP Remote Server was used to perform the backup, select a drive pool associated with an NDMP Remote Server. Similarly, if an NDMP drive pool was used, specify an NDMP drive pool.

Use Drive

Specifies the name of the drive in the drive pool that will be used to perform the restore operation. Use this option when you wish to restore using a specific Drive in the Drive Pool.

Use Proxy

Specifies the name of the proxy server that will be used to perform the restore operation. Use this option when you wish to restore using a proxy server.

Advanced Restore Options (Encryption)

Pass-Phrase

Enter the pass-phrase that is currently assigned to the client, whose data you are restoring. Note that if you have changed the pass-phrase since you secured the client data, you need to provide the new pass-phrase here, not the old one.

Re-enter Pass-Phrase

Re-enter the pass-phrase for confirmation.

If you attempt an immediate restore of encrypted data that was pass-phrase protected without entering the pass-phrase here, the restore operation will fail.

If you have an exported pass-phrase set up, and you enter the pass-phrase under **Decryption**, you over-ride (not overwrite) the client properties pass-phrase. Thus, if you enter the pass-phrase incorrectly, the restore does not complete successfully.

Startup

Select from the following options. Note that all the options described in this help may not be available and only the options displayed in the dialog box are applicable to the operation for which the information is being displayed.

Priority

- **Use Default Priority**

If selected, the default priority for this type of job will be used in determining how the Job Manager will allocate resources for this job.

- **Change Priority**

Use this option to manually specify the priority for the job, between 0 (highest priority) and 999 (lowest priority). The Job Manager will use the priority setting when allocating the required resources. This is useful if you have jobs that are very important and must complete, or jobs that can be moved to a lower priority.

Start up in suspended state

Specifies that this job will start in the Job Controller in a suspended state and cannot run until the job is manually resumed using the **Resume** option. This option can be used to add a level of manual control when a job is started. For example, you could schedule jobs to start in the suspended state and then choose which scheduled jobs complete by resuming the operation started in the suspended state.

Description

Use this field to enter a description about the entity. This description can include information about the entity's content, cautionary notes, etc.

Alert

Use this tab to configure an alert for a schedule policy.

Configure Alert

- **Alert**

The currently configured Alert.

- **Add/Modify Alert**

When clicked, opens the Alert Wizard to configure alerts for this operation.

- **Delete Alert**

When clicked, deletes any existing alerts that are already configured.

Advanced Restore Options (Options)

Use this dialog box to make modifications in the names and locations of databases for restore.

Database

Displays the name of the database. You can edit the name by clicking within the box. Changes cause the data to be restored to a new database on the selected server.

For a detailed description on how to change the name of a SQL database, see the procedure titled "Restore a Database with a Different Name" in *Books Online*.

Device Name

Displays the device name of the database.

File Name

Displays the name of the database file to be restored.

Physical Path

Displays the complete path to the selected database file. You can edit the path by clicking within the box. Changes cause the database files to be restored to a new location. If the specified path does not exist, it will be created during the restore process.

For SQL, if the database file path is changed without changing the name of the database, then the existing database will be overwritten and will point to the new location. For a detailed description, see the procedure titled "Restore with a Different Data File Path" in *Books Online*.

Size (MB)

Displays the size of the database file in megabytes.

Find & Replace

For SQL, click to change multiple Database Name or Database File Path.

Advanced Recover Options (Browse Options)

Use this dialog box to specify advanced restore options.

Restore Deleted Items

Specifies that the restore/recovery operation will include data that was secured during all backup/archive operations, including deleted items.

Specify Time Range

Specifies to restore/recover data within a specified time zone and time threshold. Select the checkbox to enable this option. Clearing the checkbox disables this option. If this option is enabled, set the following parameters:

- **Time Zone** - Restores/recovers only those jobs within a specified time zone. Click the pull-down list and click the desired time zone.
- **Start Time** - Displays the start time for the restore/recovery operation. Click the pull-down arrow and click a date; use the up and down arrows to specify a time.
- **End Time** - Displays the end time for the restore/recovery operation. Click the pull-down arrow and click a date; use the up and down arrows to specify a time.

Operation Window

Use this dialog box to add, modify, or delete operation rules for system operations.

Name

The friendly name assigned to the operation window.

Operations

The operations included in the rule.

Days

Displays the days in which the operations in this rule should not be run.

Do not run intervals

Displays the times in which the operations in this rule should not be run.

Add

Click this button to add an operation rule from the Operation Rule Details dialog box.

Modify

Click this button to modify an operation rule from the Operation Rule Details dialog box.

Delete

Click this button to delete an operation rule.

Ignore Operation Rules at Higher Levels (not applicable at the CommServe and Client group level)

Select this option to disable all operation rules set at levels above the currently selected entity.
