

---

# A VMware vCenter Administrator's Guide to IBM InfoSphere Copy Data Management



# Contents

<b>Chapter 1 - Introduction</b> .....	<b>1</b>
IBM InfoSphere Data Virtualization .....	1
Capture Mechanisms.....	2
Capturing Virtual Server Data.....	3
Replicating Captured VMware Data .....	5
Datastore Space Monitoring.....	5
Accessing Data .....	6
Workflows to Automate Access to SQL Server Data .....	7
<b>Chapter 2 - Discovering and Protecting VMware VMs</b> .....	<b>9</b>
Discovering VMs .....	9
Deleting VMs.....	11
Discovering Applications on a VM .....	12
<b>Chapter 3 - Mounting a VMware VM Image</b> .....	<b>15</b>
Mounting a VMware VM (Standard Mount) .....	15
Recovering a Mounted VMware VM to Production Storage.....	17
<b>Chapter 4 - Replicating VMware Data to a Datastore</b> .....	<b>19</b>
<b>Chapter 5 - Restoring Virtual Machines</b> .....	<b>21</b>
<b>Chapter 6 - IBM InfoSphere VDP and VMware HA</b> .....	<b>23</b>
Shutting Down a Managed IBM InfoSphere VDP Appliance .....	24
Use of vMotion, DRS, DPM.....	24
VMware FT (Fault Tolerance) Configurations.....	24
Use of Resource Pools with IBM InfoSphere VDP Appliance Instances .....	24
Licensing Considerations .....	25
Networking Considerations.....	25
Best Practices for IBM InfoSphere VDP Appliance in an HA Failover Configuration.....	26
System Recovery Steps After an HA failover .....	26
Migrating IBM InfoSphere VDP Appliances.....	27
<b>Chapter 7 - VMware Permissions</b> .....	<b>29</b>
Creating the VDPReadOnly vCenter Role.....	30

Creating the VDPOperations vCenter Role .....31  
The vCenter Permissions List, vCenter 6.0.....32  
The vCenter Permissions List, vCenter 6.5.....33  
The vCenter Permissions List, vCenter 6.7.....34  
Assigning Minimum Permissions .....35

# 1 Introduction


This chapter provides a high-level overview of basic IBM InfoSphere concepts and procedures used to capture, manage, and access virtual machines (VMs). Specifically, this chapter describes:

- [IBM InfoSphere Data Virtualization on page 1](#)
- [Capture Mechanisms on page 2](#)
- [Capturing Virtual Server Data on page 3](#)
- [Accessing Data on page 6](#)
- [Workflows to Automate Access to SQL Server Data on page 7](#)

## IBM InfoSphere Data Virtualization

An InfoSphere VDP Appliance is a highly scalable copy data management platform that virtualizes application data to improve the resiliency, agility, and cloud mobility of your business. It works by virtualizing data in much the same way other technologies have virtualized servers and networks.

VDP Appliances enable you to capture data from production systems, manage it in the most efficient way possible, and access virtual or physical copies of the data whenever and wherever they are needed.



### Capture, Manage and Use Application Data

Application data is captured at the block level, in application native format, according to a specified SLA. A Golden copy of that data is created and stored once, and is then updated incrementally with only the changed blocks of data in an “incremental forever” model. Unlimited virtual copies of the data can then be made available instantly for use, without proliferating physical copies and taking up additional storage infrastructure.

## Capture Mechanisms

An InfoSphere VDP Appliance captures an initial full copy of an application's data or a VM. From then on, only the changed data is captured. To track changes, the InfoSphere VDP Appliance uses either VMware API calls or the VDP Connector.

### VMware API Calls for Entire VMs

An InfoSphere VDP Appliance can take advantage of VMware API for data protection (VADP) calls to capture an entire virtual server. Specifically, the API calls can:

**Perform change block tracking:** Makes an initial full snapshot of a database, then going forward only snapshots the changes to the database thereby enabling IBM InfoSphere's incremental forever capture strategy.

**Quiesce applications:** Ensures application consistency during capture.

When an entire VM is captured, a fully functional VM (operating system, applications, and its data) is captured. Having a copy of the entire VM guarantees that the data can be accessed fast and without issues. Because the image presented is a fully functional VM, if needed, it can be started and run from an InfoSphere VDP Appliance directly and then optionally migrated to a new, permanent location.

Virtual servers and their applications can be grouped and captured with a single SLA.

### The VDP Connector, for Individual Applications, Databases, and File Systems

The VDP Connector is used to capture applications. The VDP Connector is a small-footprint, lightweight service that can be installed on either virtual or physical servers.

The VDP Connector provides a more granular capability than what is provided by VMware API calls. It allows you to capture applications that cannot be snapped by VMware. In addition, it also allows you to capture Microsoft SQL Server clusters and offers options for handling individual database transaction logs.

Specifically, the VDP Connector:

- Discovers applications
- Quiesces applications
- Where applicable, takes advantage of Microsoft VSS Writer for discovery, capture, and access operations.
- Identifies changes to application data for IBM InfoSphere's incremental forever capture strategy.
- Captures databases in clustered application deployments.
- Captures database transaction logs:
  - o Captures database(s) and logs with one Policy Template
  - o Truncates database transaction logs as needed
  - o Rolls logs forward for point-in-time recovery
- Allows you to apply a single Policy Template to multiple VMs and/or applications.
- For VMware VMs:
  - o Captures databases that use pRDMs and vRDMs
  - o Avoids virtual server "stun" issues.

## Capturing Virtual Server Data

When capturing virtual machine data, you can capture:

- Applications on a VM
- Applications in a Consistency Group
- Application(s) along with the VM's boot volume
- Entire VMs individually or in groups

Capturing VMs consists of four steps in the VM Onboarding Wizard in the IVGM Application Manager:

1. Choose the Server for VM Discovery
2. Select Virtual Machines
3. Manage VMs
4. Approve the Onboarding Summary

---

**Note:** You can create Production Direct-to-Dedup policies VMware VMs without keeping a snapshot in the Snapshot Pool. Capturing VMware VMs directly to a Dedup Backup Pool is meant for long term retention when instant access from a Snapshot Pool is not required. See the IVGM online help for details.

---


---

**Note:** If you capture an entire VM with one policy and also capture individual applications on that VM with another policy, ensure that one capture operation completes before the other capture operation completes.

---

### Capturing Applications on a VM


Installing the VDP Connector on a VM allows you to capture applications on that VM. Multiple applications can be captured with a single policy template, or multiple policies can be used to capture individual applications.



**Connectors on Multiple Virtual Machines**

## Capturing VMs Individually or in Groups

To capture entire VMware VMs, the InfoSphere VDP Appliance takes advantage of VMware APIs.



### Capturing Entire VMs

---

**Note:** An VDP Appliance is a VMware VM. It can be on the same ESX server as the VMs it manages.

---

When an entire virtual server is captured, a fully functional virtual server (operating system, applications, and its data) is captured. Having a copy of the entire virtual server guarantees that the data can be accessed quickly and without issues. Because the image presented is a fully functional virtual server, it can be migrated to a new, permanent location.

Capturing whole virtual servers allows groups of virtual servers and their applications to be captured with a single SLA Policy Template.

## Capturing Applications in IBM InfoSphere Consistency Groups

A consistency group is enabled by the VDP Connector. As the name implies, consistency groups ensure consistent point-in-time capture and recovery across multiple applications on the same host.

To achieve application consistency, members of a consistency group are quiesced and captured together via a single policy.

If the InfoSphere VDP Appliance captures database logs along with the associated database (Microsoft SQL Server and Oracle only), then all databases in that group can be recovered to the same point-in-time. Recovery and rolling forward of the logs (for databases) in a group is performed via the IBM InfoSphere user interface with a single action.

In addition to making capture and recovery operations easy and fast, consistency groups consume fewer system resources (VDisks).

## Capturing Applications and Boot Volumes

When capturing application data on VMs you have the option of also capturing the VM's boot volume.

When a VM's boot volume is captured along with its application data, if needed, an image can be presented that is a fully functional VM and its applications. The image can then be migrated to a new, permanent location.

## Replicating Captured VMware Data

To replicate data, at least two InfoSphere VDP Appliances must be joined and have exchanged certificates. Details on joining InfoSphere VDP Appliances can be found in IVGM Online Help.

Once InfoSphere VDP Appliances are joined, IBM InfoSphere Resource Profiles are used to control where data is replicated.

IBM InfoSphere Resource Policies can specify replication of data from either:

- A local Snapshot Pool to a remote Snapshot Pool
- A local Dedup Backup Pool to a remote Dedup Backup Pool
- A local Snapshot Pool to a remote OnVault pool
- A local Snapshot Pool to a remote VMware datastore. See [Replicating VMware Data to a Datastore](#) on page 19 for details.

## Datastore Space Monitoring

Datastore space utilization is checked before creating the snapshot and also monitored throughout the data movement process while data is copied from VMware snapshot to IBM InfoSphere staging disks/direct dedup objects.

In the case where the data is being replicated to a remote VMware datastore, the local datastore and the remote datastore space usage are monitored during data movement. If a critical threshold is crossed in any of the data movement jobs, all subjobs are canceled and the job fails.

Critical threshold and the frequency at which datastore usage is monitored are defined in: **Domain Manager > Configure Appliance (Appliance Configuration) > Storage Pools.**

## Accessing Data

### Role-based Access Control

IBM InfoSphere administrators can control which users have access to data, IBM InfoSphere features, processes, and resources. In addition, captured data can be defined as sensitive or non-sensitive. IBM InfoSphere users can be granted permission to access sensitive or non-sensitive data.

Captured VMs, VM data, or applications on a VM can be accessed in these three ways:

### Mounts

The IBM InfoSphere mount function provides instant access to data without moving data. Captured copies of databases can be rolled forward via the IBM InfoSphere user interface and mounted on any database server.

An InfoSphere VDP Appliance provides two ways to mount data:

- **The standard mount** presents and makes the captured data available to a target server as a file system, not as a VM or application. This is useful if a VM or application is corrupt, lost, or if a server is being replaced. In such cases you cannot use a restore operation to recover the application or VM. Instead, you can mount an image and copy the data from the mounted image to their original location on a server.
- **The Application Aware mount** presents and makes a captured database (Microsoft SQL Server or Oracle) available to a target server as a database. This allows you to address the unique challenges associated with creating and managing copies of production databases for non-production use. Application Aware mounts are performed from the InfoSphere VDP Appliance and do not require manual intervention by database, server, or storage administrators. Application Aware mounts can be used for such things as database reporting, analytics, integrity testing, and test and development.

### LiveClones

The LiveClone is an independent copy of a VM or an application on a VM that can be refreshed when the source data changes. The advantage of a LiveClone is that they are independent copies of data that can be incrementally refreshed and masked before being made available to users. This allows teams such as development and test to ensure they are working on the latest set of data without having to manually manage the data or access the production environment.

### Restores

The restore function reverts production data to a specified point in time. Restore operations actually move data. Typically restore operations are performed to restore a VM or application to a valid state after a massive data corruption or storage array failure. The amount of time required to complete a restore operation depends on the amount of data involved.

---

**Note:** IBM InfoSphere provides the flexibility to restore to the original server or to an alternate server. To restore to an alternate server, the VDP Connector must be installed on the alternate server before initiating the restore operation.

---

To restore a database and then apply logs, the restored database must be in Restoring Mode. IBM InfoSphere's log capture and restore functionality allows you to, from the IBM InfoSphere user interface, restore the database in Restoring Mode and then roll the logs forward to a specific point in time.

If you restore a database through the IBM InfoSphere user interface without specifying Restore with no Recovery, the database will be restored and brought on line without applying logs.

## Workflows to Automate Access to SQL Server Data

While SLA Policy Templates govern the automated capture of production data, Workflows automate access to the captured data.

Workflows are built with captured data. Workflows can present data as either a direct mount or as a LiveClone:

- Direct mounts (standard or application aware) work well for application data that does not need to be masked prior to being presented. A mounted copy of data can be refreshed manually or automatically on a schedule. Direct mounts allow you to instantly access captured Microsoft SQL Server data without actually moving the data.
- A LiveClone is a copy of your production Microsoft SQL Server data that can be updated manually or on a scheduled basis. You can mask sensitive Microsoft SQL Server data in a LiveClone prior to making it available to users.


Combining IBM InfoSphere’s automated data capture and access control with Workflows and their optional data masking capabilities allows you to create self-provisioning environments. Now, instead of having to wait weeks for DBAs to update test and development environments, users can provision their own environments almost instantly.

For example, an IBM InfoSphere administrator can create an SLA Template Policy that captures data according to a specified schedule. Optionally, the administrator can mark the captured production data as sensitive and only accessible by users with the proper access rights.

After access rights have been defined and data has been captured, the administrator can create a Workflow that:

- Makes the captured data available as a LiveClone or a direct mount
- Updates the LiveClone or mountable data on a scheduled or on demand basis
- Optionally automatically applies scripts to the LiveClone’s data after each update. This is useful for masking sensitive data.

Once the Workflow completes, users with proper access can, via the IBM InfoSphere user interface, provision their environments with the LiveClone or mountable data.



**Workflow With Masked Social Security Data**



## 2 Discovering and Protecting VMware VMs

This chapter details:

[Discovering VMs on page 9](#)

[Deleting VMs on page 11](#)

[Discovering Applications on a VM on page 12](#)

### Discovering VMs

Use the VM Onboarding Wizard from the Application Manager to discover virtual machines (VMs) managed by a vCenter or by an individual ESXi server. Once you have discovered one or more VMs, you can protect them all at once by applying an SLA Template and Profile or you can simply add them to the Applications list as unmanaged or ignored VMs.

During IVGM configuration, you may have already added the ESXi servers and vCenters as hosts to InfoSphere VDP Appliances. See [Adding a Host](#) for more information. The VM Onboarding Wizard also allows you to add a server in case it was not added before.

---

**Note:** When you discover a VMware vCenter, all ESXi hosts are automatically discovered.

---


---

**Note:** Virtual machine discovery on a hypervisor requires an IBM InfoSphere user with 'Host Manage' IBM InfoSphere rights.

---


To discover a VM:

1. Open the **Application Manager**.
2. Click **Add Virtual Machines** in the upper right corner.




**Opening the VM Onboarding Wizard**

- The Onboarding Wizard opens.



**Page 1 of the Onboarding Wizard**

- Follow the Wizard to the Onboarding Summary at the end and click Finish. After discovery, the virtual machines and hypervisors are added as hosts in the Domain Manager.



**The Onboarding Summary for a Simple Job**

**Note:** The InfoSphere VDP Appliance relies on synchronicity between an InfoSphere VDP Appliance and its discovered hosts. Hosts that are not connected to an NTP server can drift, resulting in differences between the host's record and the InfoSphere VDP Appliance's record of the time snapshots taken or other actions performed by the InfoSphere VDP Appliance.

## Deleting VMs

You can delete unprotected VMs. To delete protected VMs, first unprotect them by disassociating all SLAs.


---

**Note:** Remote VMs that appear in the Application Manager under the Remote category should be deleted from the remote InfoSphere VDP Appliance.

---

To delete a VM:

1. Open the Application Manager.
2. Click the VM filter tab from the filters on the left side.
3. Select the VM to delete.
4. Right-click it to open the service menu, and click **Delete Application**.



**The Application Manager Navigation Pane List Filter**

5. Click **Yes** in the confirmation dialog.

Images from a deleted application appear as orphans in the navigation pane under Orphan. You can see an application in the orphan section only if there are any images of that application.

You can delete a resource profile or a policy template only when the resource profile or policy template is not used to protect an application.

---

**Note:** Deleting a VM or removing its protection cleans up all Dedup-Async or StreamSnap related images (if replication is configured for that VM). If any stale images are left on an InfoSphere VDP Appliance (usually due to a remote appliance unavailability), in the left bottom menu list you will see an operation called **Cleanup Dedup Async** or **Cleanup StreamSnap**. For more information, see the IVGM Online Help.


---

## Discovering Applications on a VM

You can discover applications on VMs that are known to the InfoSphere VDP Appliance. You must have the 'Host Manage' or 'Application Manage' rights to discover applications and the VDP Connector must be installed and configured on the host.


To discover an application:

1. Open the VDP Desktop to the **Application Manager**. In the upper right corner, click **+ Add Application**.



### Adding Applications from the Application Manager

2. The Add Applications page appears. Select **Discover connector supported Applications** and **Using Existing Host**.
3. From the Appliance drop-down list, select the InfoSphere VDP Appliance that connects to the host(s) that includes the applications you would want IVGM to discover. You can choose All for multiple appliances or select a specific appliance.
4. Select the host that includes the application you would like to protect. If you have many hosts, then you can use the filters to make the list more manageable.



### Discovering Applications on a VM with a Connector

5. Select one or more hosts from the Available Hosts table, then click **Add Applications**.

---

**Note:** *The application discovery process can take some minutes, depending on the number of hosts and the applications associated with the selected hosts.*

---

6. Verify the application discovery status in the Notification Center.
7. Add additional applications or proceed to the application list in the Application Manager. The discovered applications are added to the list of applications in the Application List window of the Application Manager.

---

**Note:** *Instructions for installing and configuring the VDP Connectors are in **Connecting Hosts to IBM InfoSphere VDP Appliances**.*

---



---

# 3 Mounting a VMware VM Image

---

This chapter provides instruction for mounting VMware VMs:

- [Mounting a VMware VM \(Standard Mount\)](#) on page 15
- [Recovering a Mounted VMware VM to Production Storage](#) on page 17

When mounting to an existing VM, no settings are preserved. Captured VMDK volumes are presented as vRDM/pRDM disks to the target VM.

The InfoSphere VDP Appliance can present 45 LUNs to a VM, and a VM can have a maximum of 60 LUNs (including existing ones and mounted volumes).

## Mounting a VMware VM (Standard Mount)

To mount an active image to a VMware VM:

1. From the IVGM left-hand navigation, click the Application Manager. The Applications page opens.
2. Select the VM image that you want to mount, then choose **Access** from the drop-down list at the bottom of the Applications page.

The Timeline ramp view is a time-based visualization of 7 days of captured images for the selected application. You can use your mouse scroll wheel or the up and down arrows in the bottom left corner of the page to move the timeline through the captured images and make a selection.

3. Select an image, then select **Mount** from the list of access operations. The Mount page opens.
4. Select a host to mount to: **Existing Host** or **New Virtual Machine**.

If you select Existing Host, select a physical or virtual host from Host drop-down list. You can select any known host from the drop-down list, grouped into Physical Machines and Virtual Machines. If you need a host that has not yet been added, add it from the Domain Manager.

If you select New Virtual Machine, make the following selections specific to the virtual machine:

- o **VM Name:** Enter a name for the new VM that you want to mount,
- o **VCENTER:** Select a vCenter from the drop-down list for the new VM you want to mount.
- o **ESX HOST:** Select an ESX Host from the drop-down list for the new VM you want to mount.
- o **DATASTORE:** Select a datastore that has the required storage available from the drop-down list for the new VM you want to mount.

---

**Note:** The target InfoSphere VDP Appliance must write configuration data to the selected datastore to point to the mounted volumes, but no storage will be consumed by the image virtual copy.

---

When mounting as a new VM the VM version, Guest Id, Number of CPUs, Memory, Hardware details are preserved. Backed up VMDK volumes are presented as vRDM/pRDM disks to the new VM.

5. For Mount Mode, select one of the following:
  - o **vRDM** (virtual raw device mapping) if you need the ability to move the mounted image with VMware VMotion without taking down the VM. The maximum vRDM size for ESXi 5.0 and 5.1 is 2 TB minus 512 B. In ESXi 5.5, the size was increased to 62 TB. By default vRDM mode is selected. VMware snapshots treat mounted vRDMs as Independent and are not included in snapshots. Because of this, by default, IBM InfoSphere does not include vRDMs when protecting a mounted VM. IBM InfoSphere does provide an option where you can mark vRDMs as Dependent. Although rarely used, when this option is enabled, vRDMs will be included in VMware snapshots. IBM InfoSphere SLA templates will capture vRDMs marked as Dependent and the captured data will be counted in MDL usage.
  - o **pRDM** (physical raw device mapping) is used for file level restore operations, and if you want to share the mounted image. pRDMs can be up to 64 TB. In many cases, pRDM is your best choice.
6. If necessary, change the default storage pool from the Storage Pool drop-down list. The default storage pool is act\_per\_pool (the Snapshot Pool).
7. If desired, enter a unique name associated with the mount in **Label**.
8. Specify the following mount selections:
  - o **Mount Drive:** (Windows only). Specifies a drive letter to be assigned to the volume. If the drive letter is not available, the job fails. If multiple volumes are found, then it assigns subsequent drive letters. If no Mount Drive is specified, the VDP Connector chooses a drive letter itself, if available.
  - o **Mount Point:** The full path at which you want to mount the volume. If the path exists as an empty folder, the VDP Connector will use it. If it does not exist, the VDP Connector will create it. If it exist as a file or as a folder that is not empty, then the job will fail. If there are multiple volumes to be mounted, the VDP Connector chooses the user specified for one of the volumes and for the remaining it appends an underscore ( \_ ) followed by a number (for example, <user\_specified>\_#).
9. Click **Map to All ESX Hosts** to instruct IVGM to map the mount settings to all ESX hosts.
10. Select a single volume or multiple volumes to mount from the **Select Volumes To Mount** area. By default, all the volumes are selected, and the first volume cannot be deselected. If you mount to an existing VM, the VM disks will show up as a new filesystem drive on the VM.

---

**Note:** IVGM assumes that the first volume of the VM is the boot volume. If the selected first volume is not the boot volume, contact IBM InfoSphere support for further assistance.

---

Specify the Storage Pool, Mount Drive, and Mount Point for each volume you want to mount.

11. If you are mounting an image from the dedup pool that was captured by an InfoSphere VDP Appliance, you have the option of:
  - o Mounting to a host directly from the Dedup Pool. With this option, the mount operation is almost instant, however, because the data in the dedup pool is deduplicated, access performance is impacted because data must be rehydrated before it can be accessed.
  - o Rehydrating the data to the IBM InfoSphere Snapshot Pool before mounting. With this option, the mount operation is delayed until rehydration is complete. Access performance is fast because the data is already rehydrated and ready for access.
12. Click **Submit**. A job is submitted to mount the image to the selected host. Once completed, the image becomes active and is available in the Active Images view (Manage Active Images) of the Application Manager.

## Recovering a Mounted VMware VM to Production Storage

To mount a VMware VM and then migrate the VM data via VMware ESX vMotion:

1. Choose the datastore for the RDM files and VM configuration files. Do not choose the datastore that will be the permanent home of the VM. For example, if you want the VM to be on DS1 after the vMotion operation, then choose DS2 as the datastore for the mount.
2. Use vMotion to migrate the VM to the production array:
  - o Change the datastore. The new datastore cannot be the same as the source datastore.
  - o Change the format of the virtual disk. Leaving it as Same format as source will not work.

---

**Note:** *If both of the above criteria are not met, you will have a very fast vMotion that only moves RDM files or has no action because the source datastore and destination datastore are the same.*

---

3. Unmount & Delete the mount:
  - a. From the Domain Manager, right-click on the appliance and select **Configure Appliance**.
  - b. Click **Images > Mounted/Unmounted** from the navigation pane. All mounted and unmounted images are listed in the right panel.
  - c. Select the image.
  - d. Click **Unmount & Delete**.
  - e. Click **Yes** in the confirmation dialog.

---

**Note:** *VMware snapshots treat mounted vRDMs as Independent and are not included in snapshots. Because of this, by default, IBM InfoSphere does not include vRDMs when protecting a mounted VM. IBM InfoSphere does provide an option where you can mark vRDMs as Dependent. Although rarely used, when this option is enabled, vRDMs will be included in VMware snapshots. IBM InfoSphere SLA templates will capture vRDMs marked as Dependent and the captured data will be counted in MDL usage.*

---



# 4 Replicating VMware Data to a Datastore

A VMware VM can be replicated to a datastore. To replicate data, at least two InfoSphere VDP Appliances must be joined and have exchanged certificates. Details on joining InfoSphere VDP Appliances can be found in the IVGM online help.


Resource Profiles define where to replicate data. By default, Resource Profiles replicate data to either a Snapshot Pool or a Dedup Backup Pool on a remote site. When coupled with a Production to Mirror Policy, a Resource Profile can replicate VMware data directly to a datastore. To use this option:

- The datastore must be part of an ESX server/vCenter added/discovered by the remote InfoSphere VDP Appliance to which the local InfoSphere VDP Appliance is joined. See the IVGM online help for details.
- Data must be replicated via a Production to Mirror Policy that uses either Dedup-Async or StreamSnap replication. See the IVGM online help for details.

**Note:** Once defined, the local InfoSphere VDP Appliance's Resource Profiles that include the remote InfoSphere VDP Appliance will replicate VMware data to the specified datastore.

To replicate VMware data directly to a datastore:

1. Open the **Domain Manager**. Right-click the appliance and select **Configure Appliance**.
2. In the Appliance Configuration page, go to **System > Configuration > Appliance Settings > Storage** tab:



**Storage Options**

If the local InfoSphere VDP Appliance is joined with multiple remote InfoSphere VDP Appliances, select the remote InfoSphere VDP Appliance needed.

3. Click the **VM override** check box.
4. From the dropdown menus, select a vCenter host/ESX host

5. Click the green plus sign next to the required datastore name. When selecting datastores:
  - o Select as many datastores as needed. When multiple datastores are selected, VMDKs will be written to the datastores in round robin fashion.
  - o Ensure the datastore(s) free space equals the amount of data that will be replicated plus enough extra space to accommodate future growth
6. Click **Save Settings**. Resource Profiles on this InfoSphere VDP Appliance that include the remote InfoSphere VDP Appliance set up with the VM override will replicate VMware data to the selected datastore.

If you exceed the capacity of the selected datastore(s), you can add more. Replicated VMDKs will be written to the new datastore(s). Data will not be balanced across datastores when new datastores are added.

In the case where the data is being replicated to a remote VMware datastore, the local datastore and the remote datastore space usage are monitored during data movement. If a critical threshold is crossed in any of the data movement jobs, all subjobs are canceled and the job fails.

---

# 5 Restoring Virtual Machines

---

You can restore a VM to its original host or to a replacement host at the same IP address, overwriting the existing VM. Restoring always involves some data loss: any data that came in between the last snapshot job and the application failure is lost. The InfoSphere VDP Appliance offers other options for recovering data.

Applications on VMs that are protected through the VDP Connector can be protected and restored as individual applications. See the IVGM online help for details.

When you restore an image, the SLA options (Run Schedule, Expire Data) of the protected VM are turned off.

## Cloning VMware VMs

VMware VMs can also be restored by cloning a VM to a new VM. In most cases, that is the better way to restore the VM. Cloning is detailed in the IVGM online help.

## Restoring VMware VMs

Restoring from a dedup image requires space in your snapshot pool for rehydration. The space required is equal to the full Backup Size of your application as shown in the application information in the VDP Desktop. If you need to add a new snapshot pool see the IVGM online help for details.

To restore a point-in-time image from a managed VM:

1. From the IVGM left-hand navigation, click the Application Manager. The Applications page opens.
2. Select the managed VM that you want to restore, then choose **Access** from the drop-down list at the bottom right corner of the Applications page. The Access page opens listing captured images in the Timeline ramp view.  
  
The Timeline ramp view is a time-based visualization of 7 days of captured images for the selected application. You can use your mouse scroll wheel or the up and down arrows in the bottom left corner of the page to move the timeline through the captured images and make a selection.
3. Select the image type by clicking the corresponding Snapshot, Dedup, Remote Dedup, or Remote Snapshot (Dedup Async or StreamSnap) image in the Access page.
4. Select **Restore** from the list of operations. The Restore page opens.
5. Check the **Power On Virtual Machine After Restore** check box if you want the restored VM to be powered on after the restore operation is complete.
6. Select a single volume or multiple volumes to restore. By default all the volumes are selected.
7. Click **Submit**. A warning dialog appears. Read it and then enter **DATA LOSS** to confirm. A second warning appears. Enter **OVERWRITE OTHER APPS** to confirm the restore operation.

The restore job starts. You can verify that the restore operation is successful by viewing the job status in System Monitor. When the image is restored, the InfoSphere VDP Appliance creates new VMs populated with data copied from the selected point-in-time image.

---

**Note:** The Restore operation cannot be performed from a remote appliance. However, you can use a remote-dedup image for a restore operation on the source appliance itself.

---



---

# 6 IBM InfoSphere VDP and VMware HA

---

This chapter describes:

- [Shutting Down a Managed IBM InfoSphere VDP Appliance on page 24](#)
- [Use of vMotion, DRS, DPM on page 24](#)
- [VMware FT \(Fault Tolerance\) Configurations on page 24](#)
- [Use of Resource Pools with IBM InfoSphere VDP Appliance Instances on page 24](#)
- [Networking Considerations on page 25](#)
- [Licensing Considerations on page 25](#)
- [Best Practices for IBM InfoSphere VDP Appliance in an HA Failover Configuration on page 26](#)
- [System Recovery Steps After an HA failover on page 26](#)
- [Migrating IBM InfoSphere VDP Appliances on page 27](#)

IBM InfoSphere VDP Appliance is VMware HA and DRS/DPM friendly, and is supported in environments configured with these VMware services. There are some special handling considerations for HA failover with VDP Appliance and potential functional/performance degradation while the dedup store goes through its recovery and integrity checks, post VDP Appliance VM restart.

Once an IBM InfoSphere VDP Appliance instance has failed over and is running on a new ESX host in the cluster due to an ESX HA failover, snapshot operations should resume within several minutes.

Recovery and integrity checks on the dedup store can take a significant amount of time to perform. The dedup engine keeps a large index in memory which is lost in the event the VDP Appliance VM experiences an HA Failover where the VDP Appliance VM is shutdown/powered off without flushing the cache.

So while VMware HA failover can be used to restart a VDP Appliance VM instance in the event of a ESX host failure, it is best to configure the ESX environment and IBM InfoSphere VDP Appliance instance to minimize the need for HA failover that requires the VDP Appliance VM to shut down without the proper quiesce and shutdown procedure.

## Shutting Down a Managed IBM InfoSphere VDP Appliance

It is of the utmost importance that the IBM InfoSphere VDP Appliance instance does not get powered down without flushing its memory cache to disk!

It is critical that you always shut down the IBM InfoSphere VDP Appliance instance from within the IBM InfoSphere VDP Appliance itself:

1. Right-click the appliance to shut down and select **Configure Appliance** to open the Appliance Set up page.
2. Go to Configuration > **Appliance Settings**.
3. Click the **Control Panel** tab. The control panel opens.
4. Click **Shutdown**. Complete shutdown can take 10 to 20 minutes.

IBM InfoSphere VDP Appliance does **not** necessarily shut down the dedup engine cleanly from the vSphere interface (Shutdown Guest/Restart Guest) via VMware tools. While IBM InfoSphere VDP Appliance VMs are hooked into the vSphere VM guest OS shutdown process, the dedup engine sometimes takes longer to quiesce and shutdown than VMware allows, which will result in a kill -9 being executed after a period of time. The procedure above allows time for dedup to shutdown.

5. Shut the IBM InfoSphere application down inside the VDP Appliance guest, and then when the application has completed its shutdown procedure, you can power the VM off via the vSphere interface.

## Use of vMotion, DRS, DPM

Movement of a VDP Appliance VM from one ESX host or storage system to another via online vMotion and/or DRS/DPM is supported and should result in a functioning instance on the target ESX host. Performance may be degraded during the VDP Appliance VM migration resulting in SLA misses, and best practice is to perform the migration during relatively low workloads on the IBM InfoSphere VDP Appliance instance. It is recommended that VDP Appliance disk devices are located on storage that is accessible to all hosts in the ESX cluster as a host migration will be less impactful than a storage migration. But both VM and Storage migrations should be expected to complete successfully assuming required resource levels.

Host vMotion between hosts shows very little performance impact at reasonable loads. Storage vMotion appears to have greater impact on the IBM InfoSphere VDP Appliance instance. For instance, CPU utilization can trigger CPU alarms when running a Storage vMotion, GC, and 3-4 snapshot jobs in parallel. It is therefore recommended to not perform a Storage vMotion while the VDP Appliance is powered on. VM vMotion should be limited to when few active jobs (including GC tasks) are running to reduce the risk of SLA misses.

## VMware FT (Fault Tolerance) Configurations

The VMware Fault Tolerance feature is not currently supported with IBM InfoSphere VDP Appliance instances.

## Use of Resource Pools with IBM InfoSphere VDP Appliance Instances

IBM InfoSphere VDP Appliance instances can be installed and have their resources managed via resource pools.

It is important not to memory starve the IBM InfoSphere VDP Appliance VM for dedup processing, which can be time sensitive, and through which any deduplicated data must pass for remote data movement. The memory limits defined in ***Installing and Upgrading InfoSphere Global Manager*** represent the minimum resources required to run the IBM InfoSphere VDP Appliance instance at expected levels.

While it is important not to vCPU starve the IBM InfoSphere VDP Appliance instance as well, memory reservation is more important than vCPU reservation for a IBM InfoSphere VDP Appliance instance.

For production deployments, use CPU and Memory reservations where possible, or ensure resources are not shared on the ESX host to ensure reliable performance is achieved.

For the latest specifications on IBM InfoSphere VDP Appliance CPU and memory requirements, see ***Installing and Upgrading InfoSphere Global Manager***.


## Licensing Considerations

IBM InfoSphere VDP Appliance licensing uses some portion of the MAC address of eth0 (or the first available NIC if eth0 is replaced) to generate the unique ID - which is then used for creating licenses. Make sure the MAC addresses remain the same for the VDP Appliance VM regardless of where it is running in the HA cluster.

Most HA and vMotion operations (including VM and Storage movement) should not cause the VDP Appliance VM change its MAC address. There are operations and conditions that can cause the MAC address to change, including Cloning of the VM and MAC address conflicts occurring because ESX does not check MAC addresses on powered down VMs for conflicts with running or suspended VMs, when MAC addresses are assigned.

There are several options for managing MAC address assignment at both the vCenter and ESX level. These are discussed in detail in the VMware Networking documentation.

Set a static MAC address for VDP Appliance VMs, to ensure that you retain the MAC address regardless of conditions within the cluster. Enable this for all VDP Appliance VMs. Without this, if a MAC address change occurs, the VDP Appliance VM will run in an evaluation mode (without a license) for 15 days and will require re-applying that license. Make a note of all VDP Appliance VM MAC addresses for reference. If a MAC address conflict occurs, resolve it by changing the MAC address assigned to the conflicting VM, not the VDP Appliance VM.



## Networking Considerations

Ensure that the network implementation for the HA cluster allows for seamless failover of the IBM InfoSphere VDP Appliance instances. If you have a dedicated backup network to all hosts being protected with a VDP Appliance for data movement, make sure all eligible hosts in the HA cluster have access to this network. If you have multiple networks through which a VDP Appliance is protecting VMs or physical hosts, again, make sure that infrastructure is available and named consistently across all ESX hosts in the HA cluster. Ensure all VDP Appliance required network infrastructure is accessible to IBM InfoSphere VDP Appliance instances during failover, for instance, DNS and NTP.

## Best Practices for IBM InfoSphere VDP Appliance in an HA Failover Configuration

- When adding an IBM InfoSphere VDP Appliance to a Resource Pool, do not over-commit the pool resources. Consider configuring a dedicated resource pool for the IBM InfoSphere VDP Appliance instance.
- Ensure that the VMware HA cluster nodes have sufficient resources to handle all recovered IBM InfoSphere VDP Appliance instances. IBM InfoSphere VDP Appliance instances consume a fair amount of CPU and Memory depending on License size. It is important that VMware HA slot calculations for the HA cluster VDP Appliance is running take these values into consideration. For more on HA slot calculations, see this excellent article by Duncan Epping on the Yellow-Bricks Blog. <http://www.yellow-bricks.com/VMware-high-availability-deepdiv/>
- When installing multiple IBM InfoSphere VDP Appliance instances into the same VMware HA cluster:
  - Balance multiple IBM InfoSphere VDP Appliance instances between the hosts in an HA cluster rather than installing them all on the same host. This will minimize downtime and SLA impact from a single ESX host failure.
  - Use DRS affinity rules with IBM InfoSphere VDP Appliance instances to help ensure they are always on different hosts (“VM/VM anti-affinity”).

## System Recovery Steps After an HA failover

There are two primary failover use cases with IBM InfoSphere VDP Appliance:

**Planned Failover:** This includes DRS, DPM, and vMotion migrations of the VDP Appliance VM to other clusters due to operational requirements, maintenance windows, and so on.

- These operations should be expected to succeed and running jobs will continue and complete during the VDP Appliance VM migration. The VDP Appliance should continue to operate normally during this operation though performance may be impacted.


**ESX Host failure:** This assumes that the host VDP Appliance VM’s ESX has failed and the VDP Appliance instance was not shut down cleanly. VMware HA can restart the VDP Appliance instance on another host in the HA cluster.

- When the VDP Appliance VM restarts, backup jobs that were in-flight during the failure will be failed and discarded. Restore points for these failed jobs will not be available.
- New snapshot backups start based on the normal SLA schedule. These accumulate until the dedup engine is available to ingest the snapshot images. If more than one snapshot is taken for the same application during this time only the latest will be ingested into dedup. If the snapshot expires before the dedup engine is available to process it, those restore points are lost.
- Dedup can take from two to several hours to recover on a large, very busy IBM InfoSphere VDP Appliance instance. As stated previously, when a VDP Appliance VM goes down without following the proper shutdown procedure, the in-memory cache that the dedup engine uses is lost and a recovery and integrity check is initiated before restarting.
- When the dedup engine comes back online, snapshots accumulated between the VDP Appliance instance starting up and dedup coming online will begin their dedup ingest process. Subsequent remote dedup jobs will process if configured.

## Migrating IBM InfoSphere VDP Appliances

Migrate IBM InfoSphere VDP Appliance instances as infrequently as possible. Only move IBM InfoSphere VDP Appliance instances between cluster hosts when necessary for maintenance operations or long term migrations. Configure the instance to ensure DRS and DPM are not moving VDP Appliance around the cluster on a frequent basis.

- DRS/DPM initiated movement of IBM InfoSphere VDP Appliance VMs will succeed assuming resource availability. Allowing for these automated IBM InfoSphere VDP Appliance instance migrations is supported and should perform well on clusters migrated during periods of minimal load. It is not a recommended configuration for heavily loaded IBM InfoSphere VDP Appliance instances. Performance impact on the IBM InfoSphere application will vary significantly during these migrations based on the vMotion speed (1GB, 10GB or multi-NIC) and cluster resource management strategies in place.
- The recommended configuration for heavily utilized IBM InfoSphere VDP Appliance instances requires setting the VDP Appliance VM Automation level to **Manual**. In the cluster settings, under vSphere DRS -> Virtual Machine Options, change the Virtual Machine Automation Level to **Manual** from the default of **Fully Automated**. When a migration event is generated for the IBM InfoSphere VDP Appliance instance, placement and migration recommendations are displayed, but do not run until you manually apply the recommendation. This will allow the VDP Appliance VM to be vMotioned and restarted by VMware HA services, but will not allow the instance to be automatically moved by DRS or DPM when the ESX Host is under performance load or relatively idle. This should result in DRS choosing to move other VMs rather than the VDP Appliance to relieve over subscription of resources. It is important to maintain a high level of performance in order for IBM InfoSphere to meet SLA targets.





---

# 7 VMware Permissions

---

VMware sometimes combines, separates, renames, and adds permissions with new releases of vCenter Server.

This section includes:

- [Creating the VDPReadOnly vCenter Role on page 30](#)
- [Creating the VDPOperations vCenter Role on page 31](#)
- [The vCenter Permissions List, vCenter 6.0 on page 32](#)
- [The vCenter Permissions List, vCenter 6.5 on page 33](#)
- [The vCenter Permissions List, vCenter 6.7 on page 34](#)
- [Assigning Minimum Permissions on page 35](#)

## Before You Begin

In order for IBM InfoSphere to back up and recover VMware virtual machines, the InfoSphere VDP Appliance must authenticate to the VMware vCenter Server with a user id that has sufficient privileges to perform the required operations. Create a custom IBM InfoSphere user account assigned a custom VDPReadOnly role and a custom VDPOperations role with a reduced set of privileges. A custom user also enables traceability within VMware logs to find commands used by the InfoSphere VDP Appliance. In this document, the custom user is referred to as **IBM InfoSphereUser**.

This document provides the minimum set of privileges needed to have the InfoSphere VDP Appliance perform all backup and recovery operations.

---


**Note:** Consider setting the password for this user to never expire. If the password expires then your InfoSphere VDP Appliances will be unable to work with vCenter until the password is updated, which would be a manual process.

---


## Creating the VDPReadOnly vCenter Role

You will create two vCenter roles. The first one is an VDPReadOnly role to assign the licenses permission and no other permissions:

1. Log into vSphere as a user with Administrator privileges.
2. On the vSphere Client Home page, under Administration, click **Roles**.
3. Right-click the **Read-Only** role and click **Clone**. A new *Clone of Read-Only* role appears in the list of roles.



4. Right-click **Clone of Read-Only** and click **Edit**.
5. Rename the new role **VDPReadOnly**.
6. Under **Global**, check:
  - o **Disable methods**
  - o **Enable methods**
  - o **Licenses**
7. Assign no other privileges; you will add privileges as needed for the VM, cluster, etc. Click **OK**.



---


**Note:** These examples show the vSphere client application running on a Windows host. Your screens will look a little different if you use the VMware web interface.

---

## Creating the VDPOperations vCenter Role

After the VDPReadOnly role exists, create a new vCenter role for IBM InfoSphere operations:

1. Log into vSphere as a user with Administrator privileges.
2. On the vSphere Client Home page, under Administration, click **Roles**.
3. Create a new role called **VDPOperations**.
4. Check the checkboxes for each of the privileges listed in:
  - o [The vCenter Permissions List, vCenter 6.0 on page 32](#)
  - o [The vCenter Permissions List, vCenter 6.5 on page 33](#)
  - o [The vCenter Permissions List, vCenter 6.7 on page 34](#)
5. Click **OK** to save the role.



### Set the Permissions by Checking their Checkboxes

---

**Note:** An IBM InfoSphere VM capture operation does not require capturing of .vmx files. To capture .vmx metadata files, the role must include the **All Privileges > Datastore > Update Virtual Machine Metadata** option.

---

## The vCenter Permissions List, vCenter 6.0

The IBM InfoSphere vCenter Server user must have the following permissions:

- Datastore: Allocate Space
- Datastore: Browse Datastore
- Datastore: Low Level File Operations
- Datastore: Remove file
- Datastore: Update virtual machine files
- Global: Cancel task
- Global: Disable methods
- Global: Enable methods
- Global: Licenses
- Global: Log Event
- Host: Configuration: Storage partition configuration
- Host: Local Operations: Create virtual machine
- Host: Local Operations: Delete virtual machine
- Host: Local Operations: Reconfigure virtual machine
- Network: Assign network
- Resource: Assign virtual machine to resource pool
- Tasks: Create task
- Tasks: Update task
- vApp: Export
- vApp: View OVF Environment
- vApp: vApp application configuration
- Virtual machine: Configuration: Add existing disk
- Virtual machine: Configuration: Add new disk
- Virtual machine: Configuration: Add or remove device
- Virtual machine: Configuration: Advanced
- Virtual machine: Configuration: Change resource
- Virtual machine: Configuration: Settings
- Virtual machine: Configuration: Disk change tracking
- Virtual machine: Configuration: Disk lease
- Virtual machine: Configuration: Modify device settings
- Virtual machine: Configuration: Query unowned files
- Virtual machine: Configuration: Raw device
- Virtual machine: Configuration: Remove disk
- Virtual machine: Guest Operations: Execute
- Virtual machine: Guest Operations: Modify
- Virtual machine: Guest Operations: Query
- Virtual machine: Interaction: Power off
- Virtual machine: Interaction: Power on
- Virtual machine: Interaction: Suspend
- Virtual machine: Inventory: Create from existing
- Virtual machine: Inventory: Create new
- Virtual machine: Inventory: Remove
- Virtual machine: Provisioning: Allow disk access
- Virtual machine: Provisioning: Allow read-only disk access
- Virtual machine: Provisioning: Allow virtual machine download
- Virtual machine: Provisioning: Clone virtual machine
- Virtual machine: Snapshot management: Create snapshot
- Virtual machine: Snapshot management: Remove snapshot
- Virtual machine: Snapshot management: Rename snapshot
- Virtual machine: Snapshot management: Revert to snapshot (*In vCenter 5.0 was Virtual machine: State.*)

## The vCenter Permissions List, vCenter 6.5

The IBM InfoSphere vCenter Server user must have the following permissions:

- Datastore: Allocate Space
- Datastore: Browse Datastore
- Datastore: Low Level File Operations
- Datastore: Remove file
- Datastore: Update virtual machine files
- Global: Cancel task
- Global: Disable methods
- Global: Enable methods
- Global: Licenses
- Global: Log Event
- Host: Configuration: Storage partition configuration
- Host: Local Operations: Create virtual machine
- Host: Local Operations: Delete virtual machine
- Host: Local Operations: Reconfigure virtual machine
- Network: Assign network
- Resource: Assign virtual machine to resource pool
- Tasks: Create task
- Tasks: Update task
- vApp: Export
- vApp: View OVF Environment
- vApp: vApp application configuration
- Virtual machine: Configuration: Add existing disk
- Virtual machine: Configuration: Add new disk
- Virtual machine: Configuration: Add or remove device
- Virtual machine: Configuration: Advanced
- Virtual machine: Configuration: Change resource
- Virtual machine: Configuration: Disk change tracking
- Virtual machine: Configuration: Disk lease
- Virtual machine: Configuration: Modify device settings
- Virtual machine: Configuration: Query unowned files
- Virtual machine: Configuration: Raw device
- Virtual machine: Configuration: Remove disk
- Virtual machine: Configuration: Settings
- Virtual machine: Guest Operations: Guest Operation Modifications
- Virtual machine: Guest Operations: Guest Operation Program Execution
- Virtual machine: Guest Operations: Guest Operation Queries
- Virtual machine: Interaction: Power off
- Virtual machine: Interaction: Power on
- Virtual machine: Interaction: Suspend
- Virtual machine: Inventory: Create from existing
- Virtual machine: Inventory: Create new
- Virtual machine: Inventory: Remove
- Virtual machine: Provisioning: Allow disk access
- Virtual machine: Provisioning: Allow read-only disk access
- Virtual machine: Provisioning: Allow virtual machine download
- Virtual machine: Provisioning: Clone virtual machine
- Virtual machine: Snapshot management: Create snapshot
- Virtual machine: Snapshot management: Remove snapshot
- Virtual machine: Snapshot management: Rename snapshot
- Virtual machine: Snapshot management: Revert to snapshot (*In vCenter 5.0 was Virtual machine: State.*)

## The vCenter Permissions List, vCenter 6.7

The IBM InfoSphere vCenter Server user must have the following permissions:

- Datastore: Allocate space
- Datastore: Browse datastore
- Datastore: Low level file operations
- Datastore: Remove file
- Datastore: Update virtual machine files
- Global: Cancel task
- Global: Disable methods
- Global: Enable methods
- Global: Licenses
- Global: Log Event
- Host: Configuration: Storage partition configuration
- Host: Local Operations: Create virtual machine
- Host: Local Operations: Delete virtual machine
- Host: Local Operations: Reconfigure virtual machine
- Network: Assign network
- Network: Configure
- Resource: Assign virtual machine to resource pool
- Tasks: Create task
- Tasks: Update task
- Virtual machine: Configuration: Acquire disk lease
- Virtual machine: Configuration: Add existing disk
- Virtual machine: Configuration: Add new disk
- Virtual machine: Configuration: Add or remove device
- Virtual machine: Configuration: Advanced configuration
- Virtual machine: Configuration: Change settings
- Virtual machine: Configuration: Change resource
- Virtual machine: Configuration: Configure raw device
- Virtual machine: Configuration: Modify device settings
- Virtual machine: Configuration: Query unowned files
- Virtual machine: Configuration: Remove disk
- Virtual machine: Configuration: Toggle disk change tracking
- Virtual machine: Edit Inventory: Create from existing
- Virtual machine: Edit Inventory: Create new
- Virtual machine: Edit Inventory: Remove
- Virtual machine: Guest Operations: Guest Operation Modifications
- Virtual machine: Guest Operations: Guest Operation Program Execution
- Virtual machine: Guest Operations: Guest Operation Queries
- Virtual machine: Interaction: Power off
- Virtual machine: Interaction: Power on
- Virtual machine: Interaction: Suspend
- Virtual machine: Provisioning: Allow disk access
- Virtual machine: Provisioning: Allow read-only disk access
- Virtual machine: Provisioning: Allow virtual machine download
- Virtual machine: Provisioning: Clone virtual machine
- Virtual machine: Snapshot management: Create snapshot
- Virtual machine: Snapshot management: Remove snapshot
- Virtual machine: Snapshot management: Rename snapshot
- Virtual machine: Snapshot management: Revert to snapshot (*In vCenter 5.0 was Virtual machine: State.*)
- vApp: Export
- vApp: View OVF environment
- vApp: vApp application configuration

## Assigning Minimum Permissions

To limit access of IBM InfoSphereUser, assign the VDPReadOnly role to IBM InfoSphereUser at the vCenter level and the VDPOperations role to IBM InfoSphereUser at the Datacenter level, then set NoAccess at the highest level necessary to restrict IBM InfoSphereUser from all VMs and ESXi servers that will never be mounted to or backed up by the InfoSphere VDP Appliance.

To assign to IBM InfoSphereUser the minimum permissions necessary to perform all required functions:

1. Log into vSphere as a user with Administrator privileges. On the vSphere Client Home page, click **Hosts and Clusters**.
2. Select the vCenter to ensure that permissions are propagated correctly. On the **Permissions** tab or under the Actions dropdown, select **Add Permission**.
3. Select **VDPReadOnly** from the Assigned Role drop-down menu.
4. Check the **Propagate to Children** check box at the bottom of the window.
5. Click **Add** to open the Select Users or Groups dialog box.
6. Select the domain where **IBM InfoSphereUser** is located from the Domain drop-down menu and type **IBM InfoSphereUser** in the Search box. Click **Add**. IBM InfoSphereUser is added to the Users list. Click **OK**.
7. Select the Datacenter to ensure that permissions are propagated correctly.
8. On the **Permissions** tab or under the Actions dropdown, select **Add Permission**.
9. Select **VDPOperations** from the Assigned Role drop-down menu.
10. Check the **Propagate to Children** check box at the bottom of the window.
11. Click **Add** to open the Select Users or Groups dialog box.
12. Select the domain where **IBM InfoSphereUser** is located from the Domain drop-down menu and type **IBM InfoSphereUser** in the Search box. Click **Add**. IBM InfoSphereUser is added to the Users list. Click **OK** and then click **OK** again.
13. Go back to Inventory > Hosts and Clusters. Right-click each branch that will have no IBM InfoSphere jobs, select IBM InfoSphereUser, and assign the **No Access** role to IBM InfoSphereUser. Click **OK** to finish.

**In this example vCenter hierarchy, if you want to:**

(Assume IBM InfoSphereOperations role was assigned at Datacenter 2).

**Protect a Single VM**

If you want IBM InfoSphere to back up VM2111, then assign the No Access role to IBM InfoSphereUser at VM2111 and at ESXi Cluster 22 in [Step 13](#) above.

**Protect Multiple VMs, Access within Cluster**

If you want IBM InfoSphere to back up some or all VMs in ESXi Cluster 21, and if you expect to mount or restore the images within the same cluster, then select ESXi Cluster 22 in [Step 13](#) above.

**Protect Multiple VMs, Access to Different DataCenter**

If you want IBM InfoSphere to back up some or all VMs in ESXi Server 211, and if you expect to mount or restore the images to an ESXi server in DataCenter 1, then select vCenter in [Step 13](#) above.

