

# PHD Virtual Backup

---

## for VMware vSphere™

version 5.2  
User Guide

Document Release Date: August 15, 2011  
[www.phdvirtual.com](http://www.phdvirtual.com)



## Legal Notices

PHD Virtual Backup for VMware vSphere User Guide

Copyright © 2010-2011 PHD Virtual Technologies Inc. All rights reserved. [www.phdvirtual.com](http://www.phdvirtual.com)

PHD Virtual believes the information in this publication is accurate as of its publication date. The information is subject to change without notice.

THE INFORMATION IN THIS PUBLICATION IS PROVIDED “AS IS.” PHD VIRTUAL TECHNOLOGIES MAKES NO REPRESENTATIONS OR WARRANTIES OF ANY KIND WITH RESPECT TO THE INFORMATION IN THIS PUBLICATION, AND SPECIFICALLY DISCLAIMS IMPLIED WARRANTIES OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE.

Use, copying, and distribution of any PHD Virtual software described in this publication requires an applicable software license.

Linux is a registered trademark of Linus Torvalds.

Windows is a registered trademark of Microsoft Corporation.

VMware, VMotion, vCenter, and vSphere are either trademarks or registered trademarks of VMware Corporation.

All other trademarks and copyrights referred to are the property of their respective owners.

## Support, Sales, Renewals, and Licensing

For information on new sales, licensing and support renewals you can email [sales@phdvirtual.com](mailto:sales@phdvirtual.com) or [info@phdvirtual.com](mailto:info@phdvirtual.com).

For additional information about PHD Virtual's products and services, go to: <http://www.phdvirtual.com>.

To license and register this product, go to: <http://www.phdvirtual.com>.

For customers and partners with an active support agreement, you can use the support web board or <http://phdvirtual.com> or email [support@phdvirtual.com](mailto:support@phdvirtual.com) for information about software patches, technical documentation, and support programs.

Note: A valid support agreement is necessary to receive new release and software updates.

# Documentation Updates

Chapter	Version	Description
All	5.2	Changing some configuration options (Email, Retention, Debug, BDC, NTP) no longer requires a PHD VBA restart. Updates to these steps were made where appropriate.
1	5.2	"What's New" on page 9: Updated for v5.2.
1	5.2	"PHD Virtual Backup and Changed Block Tracking" on page 12: Added additional notes about using CBT.
1	5.2	"Restores" on page 15: Updated the lists of what is and is not included with restored VMs.
3	5.2	"Dashboard" on page 27: Added information for additional System Alerts.
3	5.2	"Storage" on page 47: Added note about using storage with multiple PHD VBAs. (7/20/2011)
3	5.2	"Mounting iSCSI Targets on Other Devices" on page 37: Added details for mounting iSCSI targets on multiple Windows platforms.
3	5.2	"Network" on page 49: Added information for using multiple network adapters with a PHD VBA (to access backup storage on another network).
4	5.2	"Using the Backup Wizard" on page 61: 'Verify backup' is now set to None by default.
6	5.2	"Using Multiple Network Adapters" on page 90: Added new section.
6	5.2	"Updating PHD Virtual Backup" on page 93: Added information for updating incompatible PHD VBAs.
A	5.2.1	"Backup Alerts" on page 102: Updated information about alerts.
A	5.2	"TCP/IP Ports" on page 104: Added new section describing the ports used by PHD Virtual Backup.
B	5.2.1	"Errors and Warnings" on page 105: Updated section with additional information.

# Contents

---

<b>Chapter 1 - Welcome</b> .....	<b>8</b>
What's New.....	9
About This Guide.....	10
Benefits of PHD Virtual Backup.....	11
PHD Virtual Backup and Changed Block Tracking.....	12
How PHD Virtual Backup Works.....	13
Backups.....	14
Restores.....	15
PHD Virtual Backup Components.....	17
Frequently Asked Questions.....	18
Best Practices.....	20
Getting Help.....	21
<b>Chapter 2 - The PHD Virtual Backup Appliance</b> .....	<b>22</b>
The PHD VBA Console.....	24
<b>Chapter 3 - The PHD Virtual Backup Console</b> .....	<b>25</b>
Dashboard.....	27
Backup Appliances List.....	28
System Alerts.....	28
Backup Catalog.....	30
File Recovery.....	33
Restoring Files.....	34
Restoring Files from a Linux or Unix VM on Windows.....	36
Mounting iSCSI Targets on Other Devices.....	37
Deleting iSCSI targets.....	38
Jobs.....	39
Job Details.....	41
Job Speeds, Deduplication, and Data Written.....	42
Job Types.....	42
Job History.....	43
Configuration.....	44
General.....	45

---

---

Storage.....	47
Network.....	49
Using DHCP.....	50
Using Static IP Addresses.....	50
Email.....	51
Retention.....	53
Connector.....	56
Support.....	58
<b>Chapter 4 - The Backup Wizard.....</b>	<b>59</b>
Accessing the Backup Wizard.....	60
Using the Backup Wizard.....	61
<b>Chapter 5 - The Restore Wizard.....</b>	<b>66</b>
Accessing the Restore Wizard.....	67
Using the Restore Wizard.....	68
<b>Chapter 6 - Using PHD Virtual Backup.....</b>	<b>71</b>
Creating Backup Jobs.....	72
Running a Backup Now.....	73
Scheduling Backups.....	75
Viewing Jobs.....	77
Restoring Backups.....	79
Restoring Files.....	80
Configuring Email Alerts.....	82
Verifying Backups and Restores with TrueRestore™.....	83
Backup Retention and Archiving.....	84
Excluding VMs and Disks.....	85
Sending Backup Files to Tape.....	86
Limiting the PHD Console to a Single PHD VBA.....	87
Increasing Backup Storage (Attached Disk).....	89
Using Multiple Network Adapters.....	90
Updating PHD Virtual Backup.....	93
<b>Appendix A - Troubleshooting.....</b>	<b>95</b>
Downloading Support Files.....	96
Recovering Backups from an Unavailable PHD VBA.....	97
Resetting PHD VBA Network Settings.....	98
BDC Share and Local Security Policies.....	99

---

---

PHD VBA will not Power On.....	100
Backup Alerts.....	102
TCP/IP Ports.....	104
<b>Appendix B - Errors and Warnings.....</b>	<b>105</b>
<b>Index.....</b>	<b>107</b>

# Chapter 1 - Welcome

PHD Virtual™ Backup for VMware vSphere™ provides reliable backup and recovery for all of the virtual machines (VMs) in your VMware environment. With PHD Virtual Backup, you can manage backup and recovery right from within vSphere Client using simple, integrated menus. Using the PHD Virtual Backup Console and wizards, you can you create and manage custom backup and restore jobs to meet all of your data protection requirements.

PHD Virtual Backup is built on the next generation of PHD's award winning VBA™ (Virtual Backup Appliance) architecture. Purpose-built for virtualization, the PHD VBA architecture enables backup and recovery to be deployed as a virtualized workload directly on the VMware platform. This approach enables high-performance data protection that seamlessly scales for large and distributed deployments. With PHD Virtual Backup, there is no need to deploy and manage separate physical servers, additional software, scripts, or agents. After you've deployed and configured the PHD Virtual Backup Appliance and plug-in, you're ready to begin protecting your virtual environment, right away.

Topics in this chapter include:

What's New.....	9
About This Guide.....	10
Benefits of PHD Virtual Backup.....	11
PHD Virtual Backup and Changed Block Tracking.....	12
How PHD Virtual Backup Works.....	13
Frequently Asked Questions.....	18
Best Practices.....	20
Getting Help.....	21

## What's New

### PHD Virtual Backup v5.2

- An additional network interface can be used for each PHD VBA, allowing added flexibility when accessing backup storage locations on other networks.
- A PHD VBA restart is no longer required for some configuration changes, streamlining the configuration process and the overall PHD Console experience.
- The delete process was improved with version 5.2. Manual and automatic (trim) deletes take much less time to complete, especially when a large number of files must be processed.
- Earlier versions of the PHD VBA can now be updated from the latest version of the PHD Console right from the PHD Console's Dashboard.
- Restores now have priority over backups - allowing you to recover your data without the need to wait for backups in progress to finish.
- Improvements were made to improve the overall stability and speed of backup and restore processing, including enhancements to Change Block Tracking backup performance.

*Refer to the Release Notes for additional details about each PHD Virtual Backup update.*

## About This Guide

This guide is designed to introduce you to PHD Virtual Backup for VMware vSphere and to:

- Illustrate the steps necessary to perform the available product functions, including virtual machine backups and restores.
- Describe the PHD Virtual Backup Appliance configuration options.
- Explain what to do when troubleshooting certain scenarios.

**Note:** This guide contains information tailored to using PHD Virtual Backup for VMware vSphere - if you are using PHD Virtual Backup on another hypervisor, refer to the specific User Guide for that hypervisor.

In addition to this guide, an Installation Guide is available that can assist you with the installation of the product, including the PHD Console and Plug-in and the deployment of the PHD Virtual Backup Appliance. The Installation Guide is available on the [PHD Virtual Web site](#) as well as in the installation package.

**Table 1 - Terms used in this guide**

Term or acronym	Definition
<b>PHD Virtual Backup Plug-in</b>	The integrated component of PHD Virtual Backup found within vSphere Client and installed via the PHD Virtual Backup MSI.
<b>PHD Virtual Backup Console</b>	The graphical interface used to configure PHD VBA settings and to configure and run backups and restores. Installed via the PHD Virtual Backup MSI along with the plug-in.
<b>VBA™</b>	Virtual Backup Appliance. A small virtual machine used to backup and restore other VMs. The PHD Virtual Backup Appliance is a VBA.
<b>PHD Virtual Backup Appliance</b>	The VBA that is deployed and used to perform backups and restores of virtual machines.
<b>PHDVB</b>	PHD Virtual Backup
<b>PHD VBA</b>	The PHD Virtual Backup Appliance (also, sometimes referred to as 'the appliance').
<b>PHD Console</b>	The PHD Virtual Backup Console.

## Benefits of PHD Virtual Backup

PHD Virtual Backup is built upon the next generation of PHD Virtual's VBA architecture and supports vSphere deployments using ESX and ESXi hypervisors. PHD Virtual Backup provides:

- vSphere Client management integration. With the plug-in for vSphere Client, PHD Virtual Backup provides "single pane of glass" management of your virtual machine backup and restore right from the vSphere Client management console.
- Reduced storage requirements and optimized network backup with TrueDedupe™. Source-side deduplication and compression occur before the data leaves the host, reducing the network impact and providing an ideal solution for backup over distributed networks and WAN environments.
- TrueRestore™ allows you to restore VM backups with confidence. Data integrity is checked during both the backup and restore processes, ensuring the restored data matches the original.
- Flexible backup storage options. You can send your backup data to locally attached storage or external storage locations such as NFS or CIFS shares.
- Job scheduling and container backups. Create backup jobs based on containers (datacenters, hosts, clusters, folders) so that any VM added to that container later will automatically be backed up based on the job settings. Also, VMs within each container can be excluded from the job, if needed.
- File Level Restore for any operating system. Restore individual files and folders without the need to restore the entire VM.
- Support for tape backup solutions via the Backup Data Connector. Quick and easy integration with tape backup solutions, providing the ability to sweep VM backups to tape.
- Scalable and fault-tolerant deployment. Distributed architecture minimizes a single point of failure. Multiple VBAs can be configured to support backup across large and distributed environments.
- Backup retention and archiving. Define and configure flexible retention policies for storing VM backups. Trim options can automatically remove old backups based on customizable policies. Archiving provides the ability to mark specific backups for archive to exclude them from being deleted by the retention policy.
- Take advantage of VMware's vStorage API and enable Changed Block Tracking to increase the speed of your backups. For additional details, see "[PHD Virtual Backup and Changed Block Tracking](#)" on page 12.

## PHD Virtual Backup and Changed Block Tracking

VMware's Changed Block Tracking (CBT) reduces both the backup window and storage requirements for your backup jobs. With PHD Virtual Backup, you can take advantage of VMware's vStorage API and enable Changed Block Tracking at the job level to reduce the time your backups take and the amount of data sent to your backup storage location.

When CBT is enabled for a backup job, each VM in the job is checked to see if it is hardware version 7. If the VM meets this requirement, the CBT configuration parameter is enabled for that VM (`ctkEnabled = true`) the first time the job runs. The initial backup that takes place with CBT enabled reads all blocks of the VM's virtual disks to create a change ID for the VM. The next time the backup job is run, the change ID is used to determine only the blocks that have changed since the last backup for each VM. Only the changed blocks are then included in the backup.

### CBT Notes

- CBT can be enabled or disabled when you create a backup job using the Backup Wizard. For details, see "[Using the Backup Wizard](#)" on page 61.
- VM Hardware Version 7 is required to run CBT. If a VM is hardware version 4 and included in a job with CBT enabled, a WARN message is included in the logs and a regular backup takes place. VMs can be upgraded from version 4 using vSphere Client.
- The initial backup with CBT enabled will take the same amount of time as a regular (non-CBT) backup, as all blocks must be read. Each backup thereafter will take much less time as only the changed blocks are read and sent.
- For restored, cloned, moved, or copied virtual machines, VMware cannot enable CBT until the new VM is powered on for the first time. Any backups run before the VM is powered on will be non-CBT backups.

## How PHD Virtual Backup Works

PHD Virtual Backup uses jobs to perform backups, restores, and backup storage maintenance (manual and automatic deletes). When a job is created, the PHD Virtual Backup Appliance (VBA) performs the requested action right away or based on a defined schedule.

When deployed to a vCenter Server or individual ESX or ESXi host, the PHD Virtual Backup Appliance performs the backup and restore processing for the VMs within that vCenter or Host environment.

The next few sections present a conceptual overview of how PHD Virtual Backup works and the components used.


- ["Backups" on page 14](#)
- ["Restores" on page 15](#)
- ["PHD Virtual Backup Components" on page 17](#)

## Backups


When a backup is run, the PHD Virtual Backup Appliance interacts with the vStorage API to create a snapshot of the virtual machine targeted for backup




Next, it attaches that snapshot to itself as a new virtual disk.



The data is then deduplicated, verified, and compressed and then sent to the defined backup storage location.




Finally, the virtual disk is detached from the appliance and the snapshot is destroyed.



## Restores

When a virtual machine restore job is created, the appliance searches the storage location for the matching VM metadata and data blocks. All of the data is then uncompressed, verified, and written to the restore location.



PHD Virtual Backup can be used to restore entire VMs or you can restore individual files with an iSCSI connection. See ["Restoring Files" on page 80](#) for details. Individual backups can also be restored from exported backup files either manually or using the Backup Data Connector.

### Restore Notes

- When a restore job is created, the PHD Virtual Backup Appliance that performed the backup is used to perform the restore (the VBA that has access to the storage location on which the backup resides).
- When a restore job completes, a new VM is created with the restored virtual disks and a recreated VMX file.
- Default VM hardware devices that were explicitly removed from a backed up VM will be included again with the restored VM. This is because restored VMs are created using a default virtual machine as a base, to which the backed up metadata is then added during the restore. If necessary, any additional devices that were not part of the originally backed up VM can be removed manually (Edit Settings...) after the restore is complete. The list of default VM hardware devices includes:
  - Memory: 256 MB
  - CPU: 1
  - IDE Controller: 2
  - PS2 Controller: 1
  - PCI Controller: 1
  - SIO Controller: 1
  - Keyboard: 1
  - Pointing Device: 1
  - Video Card: 1
  - VMCI device: 1
  - Floppy Drive: 1

- The following virtual machine configuration items are **not** included with a restored VM:
  - Attached images to CD, DVD, or floppy drives
  - CPU feature mask, CPU affinity/allocation
  - Memory affinity/allocation
  - Network Shaper
  - Network attached to NIC (if network is a distributed virtual switch)
  - Network Adapter Types VMXNET 2 or 3 are restored as type VMXNET
  - Disk Properties, including max IOPS
  - Additional default devices
  - Storage vMotion parameters dMotion.enabled and \*.DMotionParent
  - vApp Options (EULA, IP Allocation Policy, etc.)
  - Serial Port
  - Parallel Port
  - USB Controller
  - SCSI device

## PHD Virtual Backup Components

- **PHD Virtual Backup Appliance** - The Virtual Backup Appliance (VBA) which performs the backup and restore processing and presents the target for backup storage. The appliance VM can be configured to use locally attached storage or an external data store. For more information, see ["The PHD Virtual Backup Appliance" on page 22](#).
- **PHD Virtual Backup Console** - Installed with the Plug-in, the PHD Virtual Backup Console displays the status of running jobs, maintains a job history, and is used to create and manage backup and restore jobs. The console can be opened from within the vSphere Client or from the Windows Start Menu. For more information, see ["The PHD Virtual Backup Console" on page 25](#).
- **PHD Virtual Backup Plug-in** - Installed with the Console, the plug-in provides access to PHD Virtual Backup right from within vSphere Client, through simple, integrated menus.
- **Backup Wizard** - The wizard which guides you through the steps of creating and editing backup jobs. See ["Using the Backup Wizard" on page 61](#) for a detailed description of each step of the wizard.
- **Restore Wizard** - The wizard which guides you through the process of restoring a VM. See ["Using the Restore Wizard" on page 68](#) for detailed information about each step of the wizard.

## Frequently Asked Questions

This section contains frequently asked questions about PHD Virtual Backup.

### How many appliances do I need?

The number of PHD Virtual Backup Appliances you will need is determined by how your virtual machine environment is configured. Appliances must be able to access the storage where virtual machine disks are located in order to perform the backup. If you have some VMs on local storage and others on shared, you will need to deploy at least one appliance that can access the local storage on the individual host. For more information, refer to the Installation Guide.

### How many backups can I store per appliance?

The number of backups you can store per appliance depends on the size of the target storage you are using. Due to deduplication and compression, typically, to store one month of backups per VM, you need to allocate only enough backup storage equal to the total size of your VM data. For example, if you have 500 GB of VMs, allocate 500 GB of space to store one month of backups for each VM. Visit the PHD Virtual web site for additional information, including a whitepaper on planning for deduplicated backup storage.

### How is the PHD Virtual Backup Appliance deployed?

The appliance is deployed via an OVF. Refer to the Installation Guide for details.

### Why does my deduplication ratio display as inf:1?

When a deduplicated backup is performed, only new blocks of data are written to the storage location for each backup. Since this ratio is calculated while the backup is in progress, before any new data is written, the deduplication ratio is essentially infinite for the current virtual disk backup and is therefore displayed as a ratio of inf (infinite) to 1. When new data is encountered and written to disk, the deduplication ratio is updated.

### How do I configure my backup retention policy?

The retention policy (how long to keep backups for each virtual machine) is configured using the PHD Virtual Backup Console, Configuration page. For details, see ["Retention" on page 53](#).

### What happens if my appliance is restarted during a backup?

The running backup job will be canceled and any leftover snapshots will be removed the next time a backup is run. In addition, a job runs on startup and once daily to find and remove any leftover snapshots. If snapshots cannot be removed automatically, they can be removed manually using vSphere Client. If the job was a scheduled backup job, and the appliance restarts within one hour of the job's start time, the job will start again, automatically, when the PHD VBA finishes restarting.

### Can I edit a job while it is running?

Yes – scheduled jobs can be edited while in progress but any changes will not take place until the next time the job runs.

### Can I restore Exchange mailboxes or database objects?

PHD Virtual Backup is application-aware - using the File Recovery feature, you can mount an individual virtual disk where an Exchange mailbox or database was stored then access that data using your existing software. For example, to recover a database, you could create an iSCSI target from the backed up disk that contained the database then mount that target on a machine where SQL Server was installed. Then you could use SQL Server to attach the backed up database by simply

browsing the attached disk.

**Can I back up the same VM multiple times per day?**

Because PHD Virtual Backup uses backup jobs, you can create any number of customized jobs to protect your virtual machines. For example, you could create a job that backs up all of your VMs each night, then create another job that runs in the afternoon for specific VMs that have shorter RPO requirements.

**Can I replicate VMs from one host to another?**

You can restore individual VM backups to any host that the appliance performing the restore has access to. A specific replication feature is planned for a future release.

**How do I export my backups to tape?**

Using the Backup Data Connector, you can enable an SMB/CIFS share on the appliance to access all of your backup data in uncompressed format. For details, see ["Connector" on page 56](#).

**Can I order my backups?**

Using Backup Jobs, you can define a schedule for specific VMs that should run first each night. For example, create a job that backs up critical VMs beginning at 8 PM. You could then create a second backup job that includes the next tier of VMs to begin at 10 PM, and so on. In this way, you can ensure that your most critical machines have priority and are protected each night.

## Best Practices

To help ensure optimal performance when running PHD Virtual Backup in your environment, review the best practices included in this section.

### CIFS/SMB Shares

When using a CIFS share as backup storage, the CIFS service account must have full permissions (read/write/delete) for the share used as the backup target. Also, antivirus software should not be configured to analyze or scan the PHD VBA CIFS storage repository.

### NFS Shares

When using an NFS share as backup storage, the PHD VBA requires direct write access to the NFS export. During backup, the PHD VBA will directly mount and copy files to the NFS share. It is important to configure the export to allow this behavior.

### Antivirus software


Running antivirus software on a backup target can result in file locking or deletions and may cause additional issues with writing and deleting backups. PHD Virtual recommends excluding backup targets from antivirus software scans, including the network shares and directories used for backup targets.


### Disk Defragmenter

Defragmenting virtual disks can impede the overall performance of PHD Virtual Backup, resulting in lower deduplication rates, which in turn produces larger backup files written to storage and longer backup durations. To ensure consistent backup performance, PHD recommends running disk defragmentation programs only when necessary.

Running defragmentation on any disks used as backup storage is not recommended.

## Getting Help

In addition to the Release Notes, Installation Guide, and Users Guide, PHD Virtual Backup includes context-sensitive, online help which can be accessed by clicking the help button  within any of the wizards or the PHD Console or by selecting **PHD Virtual Backup Help** from within the vSphere Client Help menu.



The PHD Virtual Web site also contains additional information about PHD Virtual Backup and its benefits.


### Video Tutorials

Along with product guides and a searchable HTML library, video tutorials are available on the PHD Virtual Web site ([www.phdvirtual.com](http://www.phdvirtual.com)) that demonstrate how to install and use PHD Virtual Backup.

## Chapter 2 - The PHD Virtual Backup Appliance

The PHD Virtual Backup Appliance (PHD VBA) performs all of the backup and restore processing including source-side deduplication and compression. After it is deployed, the appliance must be configured to use a backup storage location (an attached virtual disk, CIFS share or NFS share).

**Figure 1 - The PHD Virtual Backup Appliance in vSphere Client**



When creating backup jobs, you select which PHD Virtual Backup Appliance to use to perform the job. The PHD VBA you select also determines where the backup data is stored, based on the configured storage location.

**Note:** If the PHD VBA is stored on a VMFS volume with the default formatting of 1 MB block sizes, you will be able to backup VMDK files up to 256 GB, only. If you need to backup VMDK files larger than 256 GB, you will need to store the PHD VBA on a volume formatted with larger block sizes.

- 1 MB block size = 256 GB max file size
- 2 MB block size = 512 GB max file size
- 4 MB block size = 1024 GB max file size
- 8 MB block size = 2048 GB max file size

### Configuring the PHD VBA

All configuration for the PHD VBA is done using the PHD Virtual Backup Console. See "[The PHD Virtual Backup Console](#)" on [page 25](#) for details.


PHD VBA status and log information can also be seen by selecting the PHD VBA virtual machine within vSphere Client then clicking the Console tab. See "[The PHD VBA Console](#)" on [page 24](#).

### How many VBAs do I need?

- The number of PHD VBAs you will need to deploy should be determined by how your VMware environment is configured. Each appliance can perform backups and restores for the VMs with the same shared resources. If you have configured your environment with multiple clusters or pools or other container using different shared resources, you will need to deploy a PHD VBA within each container to allow the VMs within to be backed up. Depending on the number of VMs and available resources within each pool, cluster, or Datacenter, you may choose to deploy multiple PHD VBAs within each.

If you need to deploy additional PHD VBAs, refer to the Installation Guide.

Figure 2 - PHD Virtual Backup VBA in a vCenter Cluster with shared storage



**Note:** If a PHD Virtual Backup Appliance is restarted while a backup or restore job is in progress, the job will be canceled and any leftover snapshots will be removed the next time a backup is run. In addition, a job runs each time the PHD Virtual Backup Appliance starts up, as well as once daily, to locate and remove any leftover snapshots. Any snapshots that cannot be removed automatically can be removed manually using vSphere Client.

If the job in progress was a scheduled daily or weekly backup **and the appliance is started within one hour of the scheduled start time**, the job will automatically start again.

If the job in progress was a backup Now or backup Once job, or if the PHD VBA is started more than one hour after the scheduled start time, then the job will need to be started manually.

**Caution:** If a PHD Virtual Backup Appliance is suspended during a job, when it is started again, it may become unresponsive and a restart will be required. Avoid suspending the PHD VBA during backup and restore jobs, when possible.

## The PHD VBA Console

Viewing the PHD VBA virtual machine console within vSphere Client (select the appliance, then click the Console tab) displays the number of licensed worker threads (each worker thread can perform a single backup or restore process for a virtual disk image), the available free space on the backup storage location, the latest log information, and thread status. The number of threads used during each backup and restore job can be adjusted using the Configuration area of PHD Virtual Backup Console.

The following figure shows a sample appliance console as it begins a new backup and simultaneously restores another VM.

Figure 3 - The PHD Virtual Backup Appliance console in vSphere Client

```

PHD Virtual Backup for VMware vSphere v5.1.0.4203 14:54:10
Worker Queue Depth: 0           Utility Queue Depth: 0
Worker Threads: 4              Utility Threads: 3
Store: 7.5 GB used, 2.3 GB free Deduplication Ratio:
PHDVB Appliance Log:
14:51:06 Worker-1 Archiving job into history
14:51:07 Worker-1 Restore TestVM: Job is complete
14:51:26 Worker-3 Windows 7: Backing up disk scsi0:0: 12% of 24 GB @ 2:1
14:52:00 Worker-3 Windows 7: Backing up disk scsi0:0: 16% of 24 GB @ 2:1
14:52:27 Worker-3 Windows 7: Backing up disk scsi0:0: 20% of 24 GB @ 3:1
14:52:45 Worker-3 Windows 7: Backing up disk scsi0:0: 24% of 24 GB @ 3:1
14:53:02 Worker-3 Windows 7: Backing up disk scsi0:0: 28% of 24 GB @ 3:1
14:53:41 Worker-3 Windows 7: Backing up disk scsi0:0: 32% of 24 GB @ 3:1
14:53:48 Worker-2 Restore Windows 7: Expanding job ...
14:53:48 Worker-2 Restore Windows 7: Expansion complete
14:53:48 Worker-2 Restore Windows 7: Queueing 1 VM job ...
14:53:48 Worker-2 Restore Windows 7: Queueing VM Windows 7 32 bit ...
14:53:48 Worker-1 Windows 7 32 bit: Collected metadata
14:53:50 Worker-1 Windows 7 32 bit: 1 disk(s) to restore
14:53:50 Worker-1 Windows 7 32 bit: Recreated VM as Windows 7 32 bit
14:53:53 Worker-4 Windows 7 32 bit: Allocated new disk 6000C291-20ae-e89c-c30d
14:53:56 Worker-4 Source hash unavailable, will not compute restore hash for c
14:53:56 Worker-4 Windows 7 32 bit: Restoring disk 6000C291-20ae-e89c-c30d-f9f
14:53:59 Worker-3 Windows 7: Backing up disk scsi0:0: 36% of 24 GB @ 3:1
PHDVB Worker Thread Status:
Worker-1: (idle)
Worker-2: (idle)
Worker-3: Windows 7: Backing up disk scsi0:0: 39% of 24 GB @ 4:1
Worker-4: Windows 7 32 bit: Restoring disk 6000C291-20ae-e89c-c30d-f9fb3808028


```

**Tip:** You can type Ctrl-N within the console to access appliance networking options.

## Chapter 3 - The PHD Virtual Backup Console

The PHD Virtual Backup Console allows you to manage all of your backup and restore jobs and configure your PHD Virtual Backup Appliances.

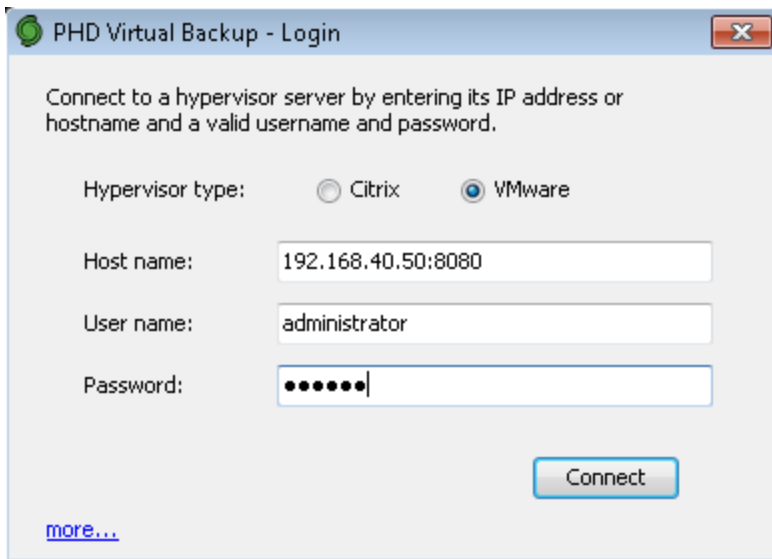
When the Console is opened, the Dashboard displays all of the available appliances (for information on deploying additional appliances, refer to the Installation Guide or online help).



**Note:** Powered off PHD Virtual Backup Appliances are not displayed. To view or manage all of your deployed appliances, make sure they are powered on.

### To access the PHD Virtual Backup Console

- The Console opens automatically after creating a job with the Backup Wizard or Restore Wizard or it can be accessed from the PHD Virtual Backup menu within vSphere Client, see ["To start the PHD Virtual Backup Console" on page 77](#)
- The Console can also be opened as a stand-alone application from the Windows Start Menu. If you are using a non-standard port to access the console, enter the port number after the server IP address you are connecting to, for example, 192.168.40.50:8081, as seen in the following image.



**Tip:** If you have multiple PHD VBAs deployed but would like to view information for a single PHD VBA only, click **more...** and enter the PHD VBA's display name. For details, see "Limiting the PHD Console to a Single PHD VBA" on page 87

The PHD Virtual Backup Console areas are described in the following sections:

- "Dashboard" on page 27
- "Backup Catalog" on page 30
- "File Recovery" on page 33
- "Jobs" on page 39
- "Configuration" on page 44

## Dashboard

The PHD Virtual Backup Console's Dashboard shows all of the deployed PHD Virtual Backup Appliances. Selecting any appliance displays multiple pie charts which represent the available storage and deduplication information. The System Alerts area displays all of the messages and alerts for each appliance.

Dashboard
?

### Backup Appliances

Lists your currently available appliances. Select an appliance to display more information.

Appliance	IP Address	Free Storage	Total Backup Data	Used Storage	Dedupe Ratio
PHDVBA 1	192.168.42.46	86.8 GB	466.9 GB	9.3 GB	50:1
PHDVBA	192.168.42.11	15.3 GB	2 GB	483.8 MB	4:1

#### Storage

- Used space: 483.8 MB (3%)
- Free space: 15.3 GB (97%)

#### Deduplication

- Used space: 483.8 MB
- Saved space: 1.5 GB

### System Alerts

Displays appliance messages and alerts.

Appliance	Message	Recommended Action

## Backup Appliances List

This area of the Dashboard displays all available appliances as well as each appliance's IP address and storage information. Pie charts display a graphical representation of the available free space and deduplication. The following table describes each column in the Backup Appliances area of the console.

**Table 2 - Backup Appliances list column descriptions**

Column	Description
Appliance	PHD Virtual Backup Appliance name.
IP Address	IP address of the PHD Virtual Backup Appliance.
Free Storage	Amount of free storage space available.
Total Backup Data	The total amount of source data backed up by the PHD Virtual Backup Appliance (before deduplication and compression).
Used Storage	The amount of actual storage space consumed by the backup data in the storage repository after deduplication and compression (if enabled). In addition to the backups, this value includes a small amount of PHD Virtual Backup system data.
Dedupe Ratio	Ratio of total backup data to used storage.


**Note: CIFS shares:** Since Windows Explorer is not aware of hard links (used with deduplication) CIFS share directories will not display used space accurately when directory properties are viewed. To see the actual disk usage for a CIFS share directory you can download the [Disk Usage](#) utility from the Microsoft Web site and use the -u option when displaying disk details.

## System Alerts

The System Alerts area provides informational messages and alerts about each available PHD VBA. The following table provides additional information about some of the alerts and messages you may encounter.

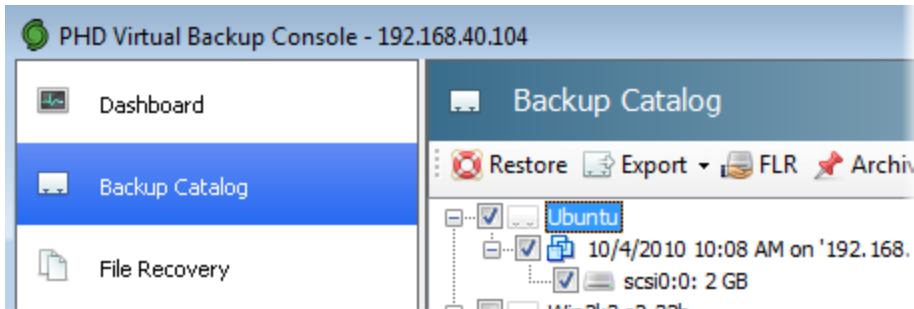
**Table 3 - System Alert descriptions**

Alert Message	Description
Appliance has no network address.	The PHD Virtual Backup Appliance does not have an IP address configured. You can manually change the network settings by opening the appliance VM's console in vSphere Client and typing CTRL-N.
Appliance has no backup storage currently mounted.	No backup storage is mounted for the appliance. Click the Storage tab to configure the storage target.
Hypervisor credentials have not been configured.	Use the General tab to configure the Hypervisor credentials for the appliance.
Appliance does not have enough free backup storage.	The storage location used to store backups is running out of free space and no new backup files can be stored. Increase the amount of space allocated to your target storage location.


Alert Message	Description
Appliance is running low on free backup storage.	The storage location used to store backups is running out of free space. Increase the amount of space allocated to your target storage location.
The product license on the appliance has expired.	PHD Virtual Backup requires a valid license to perform backups. Update your license file using the General tab.
The support license on the appliance has expired.	A valid Support license is required to receive support and upgrades from PHD Virtual. Update your license file using the General tab.
Appliance not compatible: [version]	The PHD VBA found by the PHD Console is an older version and is not compatible with the current PHD Console version. Click the link "Upgrade this appliance" to apply the latest PHD VBA update file (.phd) from the installation or update package.
The post-backup process is encountering a storage error.	When a job completes, additional tasks take place after the job processing has finished. For example, after a backup is finished writing to storage, a file-linking process takes place. If an error is encountered during post-processing, an alert is logged and an email is sent (if configured). The System Alerts Viewer contains additional information about the errors encountered as well as possible solutions, click <b>View alerts</b> to open the viewer window. The Backup Catalog also displays an alert icon  next to any backups that encountered a post-processing error.

## Backup Catalog

The Backup Catalog displays a consolidated view of all backups from each PHD VBA you have configured. From here, you can view and sort backups, select backups to restore, export backups to a file, archive, or manually delete backups by VM, Date, or the PHD Virtual Backup Appliance used.




Backups displayed in the catalog show the date and time of the backup, as seen in the image above. If a VM was powered on during the backup, the host on which the VM was running during the backup is also displayed (VMs that were powered off at the time the backup was taken display only the date and time).









**Note:** A backup marked with an alert icon  indicates an error was encountered. See ["Backup Alerts"](#) on page 102 for details.

### Duplicate VM names

If the Backup Catalog contains VMs with identical names, a UUID will be appended to one of the VM names, as seen in the following image.




**Table 4 - Backup Catalog Toolbar Buttons**


Button icon	Description
 <b>Restore</b>	Opens the Restore Wizard. For details, see <a href="#">"Using the Restore Wizard" on page 68</a> .
 <b>Export</b>	Opens the Export dialog from which you can export the selected disks as VMDK, VHD, or Raw formatted files.
 <b>FLR</b>	Opens the File Recovery wizard. For details, see <a href="#">"File Recovery" on page 33</a> .
 <b>Archive</b>	Lets you set selected backup files as archived, which means they cannot be deleted by the trim process or manual deletes. For details, see <a href="#">"Backup Retention and Archiving" on page 84</a> .
 <b>Delete</b>	Deletes the selected backup files. Note that backups marked as archived will not be deleted.
 <b>Refresh</b>	Refreshes the catalog.
 <b>View by</b>	Changes the catalog view to display backups by Virtual Machine, Date, or Appliance.
 <b>Expand All</b>	Expands or collapses the entire backup catalog tree view.

The next few sections describe some of the functions that can be performed from the Backup Catalog area of the Console with links to additional details and steps.


### Restoring Virtual Machines


1. Find the VM backup you want to restore using the Backup Catalog. You can sort the backups by VM name, Date, or PHD Virtual Backup Appliance.
2. Select the Backup file, then click  **Restore**.
3. The Restore Wizard opens. Follow the steps in the wizard to complete the restore. See ["Using the Restore Wizard" on page 68](#) for details.

### Deleting backups

1. Using the Backup Catalog, find the VM backup you want to delete. You can sort the backups by VM name, Date, or PHD Virtual Backup Appliance.
2. Select the Backup file, then click  **Delete**.
3. A Delete job is created and the backup is removed from the catalog. View the Jobs page to see the progress of the job. See ["Jobs" on page 39](#) for details.

### Deleting all backups for a specific date

1. Within the Backup Catalog, click  **View by** and select **Date**.



2. Find and select the date that contains the backups you want to remove and click **Delete**.

### Exporting Backups

Individual virtual disk backups can be exported as VMDK, Virtual Hard Disks (VHD) or Raw files. This may be useful when saving backup files to tape or when creating new VMs on different hosts. Additionally, Windows 7 and Windows Server 2008 R2 machines have the ability to mount .vhd files as native disks or boot off of these disk images. For more information about using VHD files with Windows, refer to Microsoft's knowledge base, online.

When using the VMDK export option, both the descriptor file and the data file (the flat file) are created for each VM disk you export. For example, an exported disk for the virtual machine *examplevm* will require the descriptor file, *examplevm.vmdk* and the data file, *examplevm-flat.vmdk*. The files can be renamed after export, if necessary.

1. To export a backup to a file, select the backup in the Backup Catalog and click **Export**.
2. Select the type of file to export to (VMDK, VHD, or Raw) and the virtual disk to export and click **OK**.
3. Enter a name and location for the file then click **Save**.

**Tip:** You can also right-click an individual disk in the Backup Catalog and select **Export**.

### Backup Catalog Notes

- If you renamed a VM after backing it up, all of the future backups for that VM will be included under the new VM name in the Backup Catalog. Any backups that were taken with the VM's original name will be noted in the catalog. For example, if you backed up *TestVM1*, changed the name to *NewVM1*, then ran another backup, you would find an entry only for *NewVM1* in the Backup Catalog. Under the *NewVM1* backup tree, you would then find each backup, including the backup that was taken with the VM's previous name. This backup would be noted as:  
1/24/2011 2:30 PM on 'Server1' as 'TestVM1'









## File Recovery

Instead of restoring an entire backup, you can use PHD Virtual Backup's File Recovery feature to restore individual files, folders, and application objects. By creating an iSCSI target from a backup file, you can mount and browse the backed up virtual machine disks to find the files you want to recover. File Recovery can be performed on any operating system that has an iSCSI initiator available.

**Note:** To mount iSCSI targets on a Windows machine you will need the Microsoft iSCSI Software Initiator, which is installed, by default with Windows Vista, Windows 7, and Windows 2008 Server. For earlier versions of Windows, the Initiator can be downloaded from the Microsoft web site. To mount iSCSI targets on Linux or Unix machines you must install an iSCSI Software Initiator for your specific operating system, for example, on an Ubuntu machine, you can install the Linux Open-iSCSI Initiator.

The File Recovery area of the PHD Console displays all of the iSCSI targets that have been created. From here, you can create new iSCSI targets, mount existing targets, or find the credentials needed to mount a target on another device.

**Table 5 - File Recovery Toolbar Buttons**

Button Icon	Description
 <b>Create</b>	Opens the File Recovery wizard which guides you through the process of creating a new iSCSI target from an existing backup. When created, you can mount the iSCSI target to recover files and folders.
 <b>Mount</b>	Mounts an existing iSCSI target locally.
 <b>Copy</b>	Copies an existing iSCSI target's credentials to the Windows clipboard.
 <b>Delete</b>	Deletes a selected iSCSI target. Note that the target must not be connected in order to be removed - you can disconnect targets using the iSCSI initiator.
 <b>Refresh</b>	Refreshes the list of iSCSI targets.
 <b>Collapse / Expand</b>	Collapses or expands the list of iSCSI targets.
 <b>Open iSCSI Initiator</b>	Opens the iSCSI Initiator.
 <b>Open Computer Management</b>	Opens the Windows Computer Management dialog.

The next few sections describe how to use the PHD Virtual Backup File Recovery feature in detail.


- ["Restoring Files" on page 34.](#)
- ["Restoring Files from a Linux or Unix VM on Windows " on page 36.](#)
- ["Mounting iSCSI Targets on Other Devices" on page 37.](#)
- ["Deleting iSCSI targets" on page 38.](#)

## Restoring Files


Restoring files and folders from your backups is as simple as creating and mounting an iSCSI target. Follow the steps below to create, mount, and browse files on an iSCSI target created from an existing backup.

### To restore individual files

1. Open the PHD Virtual Backup Console and click **Backup Catalog**.
2. Select the checkbox for the backup that contains the file or files you would like to recover.




3. Click **FLR**. The File Recovery wizard opens.
4. Follow the steps in the wizard to create an iSCSI target for the selected backup. You can use the wizard to create custom target credentials and to mount the target locally after the wizard completes (to mount iSCSI targets, an iSCSI Software Initiator for your operating system must be installed).
5. When the wizard completes, the target is available within the File Recovery area. The following image displays an iSCSI target created from a backup of a Windows VM.



6. If you selected to mount the target locally (if it is a Windows VM backup) the target is added as a new drive on your local computer. Open Windows Explorer to view the newly added drive. Note that mounting may take a few moments - you can open the iSCSI Software Initiator to make sure the target is connected. Additionally, you can view Computer Management, Storage, Disk Management to make sure it is mounted.

When mounted, a target created from a Windows VM backup should appear in Windows Explorer as a new hard drive.



**Note:** If the target disk does not appear in Windows Explorer, open **Computer Management > Disk Management** and find the newly mounted disk. Make sure it is set to **Online**. Additionally, you may need to import the disk if it displays as "foreign." This may happen if it is a dynamic disk created with a version of Windows different than the version running on the computer you are using to mount the target. Use the right-click menu options to import or configure the disks as necessary. Changing the attributes of a mounted disk will not affect your backup data.

- If you did not select to mount the target during the wizard, you can still mount it locally by clicking  **Mount**.

**Note:** If the iSCSI Service is not running, you will encounter an error when attempting to mount the backup. Make sure the service is running before attempting to mount any targets.

- To mount the target on another device, use the iSCSI Software Initiator and the target credentials. See "[Mounting iSCSI Targets on Other Devices](#)" on page 37 for details.

7. Using Windows Explorer, you can now browse the new drive to find the files to restore.

#### Related topics

- If you need to mount an iSCSI target created from a Linux or Unix VM on Windows, see "[Restoring Files from a Linux or Unix VM on Windows](#)" on page 36.
- To mount an iSCSI target on another device, see "[Mounting iSCSI Targets on Other Devices](#)" on page 37.
- To delete an iSCSI target, see "[Deleting iSCSI targets](#)" on page 38.

## Restoring Files from a Linux or Unix VM on Windows

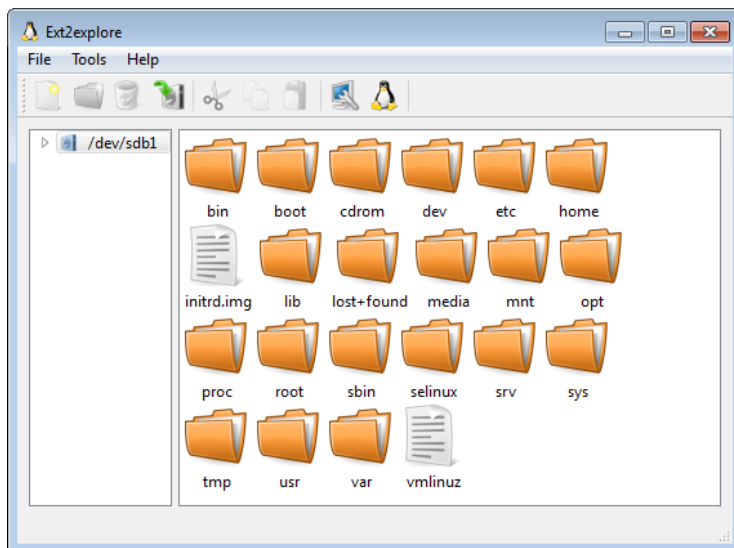
If you need to restore files from a Linux or Unix VM, but you only have access to a Windows machine to do the restore, you can use a third-party tool to view the mounted iSCSI target and browse the disk.

**Note:** In order to view the contents of the disk, the third-party tool must support the filesystem used by the Linux or Unix operating system. View your operating system documentation for filesystem details.

### To restore files from a Linux or Unix VM backup on a Windows machine

In order to restore files from an iSCSI target created from a Linux or Unix backup you will need to use a third-party tool, for example Ext2explore or explore2fs, to view the mounted disks from a Windows computer.

1. Follow the steps above to create the iSCSI target and mount the disk, making sure it is available and online within the Disk Management interface.
2. Use a file system explorer tool to view the contents of the mounted disk. The following image shows one example of a third-party tool, Ext2explore, used to browse a mounted Linux disk on Windows.



### Related topics

- To mount an iSCSI target on a Windows machine, see "Restoring Files" on page 34.
- To mount an iSCSI target on another device, see "Mounting iSCSI Targets on Other Devices" on page 37.

## Mounting iSCSI Targets on Other Devices

After creating an iSCSI target, you can either mount the target locally from the machine where the PHD Console is installed, or you can copy the target's credentials and mount the target on another device.

### To mount an iSCSI target on another device

Mount the iSCSI target using its credentials found in the File Recovery area. You can mount the target on any Windows machine that has the Microsoft iSCSI Software Initiator installed. To mount iSCSI targets on a Linux or Unix machine you must install an iSCSI Software Initiator for your operating system, for example, on an Ubuntu machine, you can install the Linux Open-iSCSI Initiator.

1. Open the Windows iSCSI Software Initiator.
2. If the service is not running, click **Yes** to start it.
3. Follow the specific steps for your operating system, below.

#### Windows 7, Windows Vista, and Windows Server 2008 R2:

- a. Use the **Targets** tab and in the **Target** dialog, enter the IP address of the PHD VBA where the iSCSI target was created.
- b. Select the IQN from the list of **Discovered targets** (click **Refresh** if needed) and click **Connect**.
- c. In the dialog that opens, click **Advanced** and select **Enable CHAP log on**.
- d. Enter the username and password of the iSCSI target in the **Name** and **Target secret** text boxes and click **OK**.
- e. Click **OK** again.

#### Windows 2003, Windows XP, Windows Server 2008:

- a. Use the **Discovery** tab and in the Target Portals area, click **Add**.
  - b. Enter the IP address of the PHD VBA where the iSCSI target was created.
  - c. Click the **Targets** tab and select the IQN of the iSCSI target from the list, and click **Log On....**
  - d. In the dialog that opens, click **Advanced...** and select **CHAP logon information**.
  - e. Enter the username and password of the iSCSI target in the **User name** and **Target secret** text boxes and click **OK**.
  - f. Click **OK** again.
4. The target is mounted and available from within Windows Explorer as a new drive.

### Related topics


- If you need to mount an iSCSI target created from a Linux or Unix VM, see ["Restoring Files from a Linux or Unix VM on Windows"](#) on page 36.

## Deleting iSCSI targets

If you need to delete an iSCSI target, you must first disconnect or log off the target using the iSCSI Initiator.

### To delete iSCSI targets

**Note:** To delete iSCSI targets, they must first be disconnected/logged off and not in use on any device (there must be no open files or directories).

1. To disconnect/log off a target:
  - a. **Windows 7 and Windows Vista:** To disconnect a target, open the Microsoft iSCSI Software Initiator, select the target and click disconnect.
  - b. **Windows 2003, Windows XP, and Windows 2008:** To log off a target, open the Microsoft iSCSI Software Initiator, click the Targets tab and select the target you want to delete. Click Details, then select the target identifier and click Log Off.
2. Open the PHD Console to the File Recovery page, select the iSCSI target, then Click  **Delete**.

## Jobs

The Jobs area lets you create new jobs, manage existing jobs, and monitor jobs in progress with two main tabs, **Current** and **History**. The **Current** tab displays scheduled and running jobs. When a running job is complete, it is moved to the **History** tab for archiving. Scheduled jobs remain in the Current tab with an **Inactive** status.

The screenshot shows the 'Jobs' window with the following data:

Job Name	Appliance	Type	Status	Progress	Current Speed	Time Remaining
Backup Daily	PHDVBA	Backup Daily	Inactive			
Backup WinXP-SP3	PHDVBA	Backup Now	Running	4%	23.3 MB/s	00:14:00
Weekly Template Backup	PHDVBA	Backup Weekly	Inactive			











Job Detail	Value
Created	11/10/2010 2:36 PM
Schedule	
Type	Now
Next Run	
Started	11/10/2010 2:36 PM
Duration	00:01:05
Average Speed	13.4 MB/s
Dedupe Ratio	2:1
Data Written	376.4 MB
Use CBT	No

Task Name	Type	Status	Dedupe Ratio
WinXP-SP3	Virtual Machine	4%	2:1
2000	Disk 20 GB	4%	2:1

The Jobs toolbar contains the options you will use to create, edit, and manage backup and restore jobs. Jobs toolbar options are described in the following table.


**Table 6 - Jobs Toolbar Buttons**

Button Icon	Description
 <b>Backup</b>	Opens the Backup Wizard which guides you through the process of creating backup jobs. See <a href="#">"Using the Backup Wizard" on page 61</a> for details.
 <b>Restore</b>	Opens the Restore Wizard which guides you through the process of restoring stored backups. See <a href="#">"Using the Restore Wizard" on page 68</a> for details.
 <b>Edit</b>	Edits the selected job.
 <b>Start</b>	Start an Inactive job or resume a paused job.
 <b>Pause</b>	Pause a job that is currently running. Note that average speed is not adjusted for paused jobs.
 <b>Cancel</b>	Cancel a job that is currently running. A cleanup process removes any unneeded snapshots or partial backup files.
 <b>Delete</b>	Deletes a current job.
 <b>Show Details</b>	Opens the Details pane which displays additional information about the selected job.
 <b>View Log</b>	Open the Log Viewer for the selected job. The Log Viewer contains the detailed log messages for the job in progress and when the job is complete.
 <b>Options</b>	Select <b>Show system jobs</b> to show or hide PHD Virtual Backup System jobs (Appliance Startup, Trim, and Orphan jobs).

## Job Details

The Job Details windows in both the Current and History tabs display additional information about each job. Detail information is based on the type of job and the options selected during the backup wizard. Details can be displayed for a job by either double-clicking the job or using the Jobs toolbar.

### To display Job Details

1. Within the Current or History tab, click to highlight a job, then click  **Show Details**.
2. The Details pane opens, displaying the information about the selected job.

Job Details Parameter	Description
Created	The date and time the job was created.
Type	The type of job. See Job Types, below, for details about each job type.
Start	The start date for the job.
Window	The window in which the job is scheduled to run, for example, 8:00 PM to 5:00 AM each night.
Recurrence	When the job is set to recur. For details on recurrence, see " <a href="#">Scheduling Backups</a> " on page 75.
Next Run	When the scheduled job will run next.
Started	The date and time the job was queued.
Duration	The total time the job took to run.
Average Speed	The total data processed by the job divided by the job duration.
Dedupe Ratio	The ratio of the total job data (all VMs, etc) to the actual data written to the backup store.
Data Written	The size of the actual data written to the backup store.
Use CBT	Indicates whether or not Changed Block Tracking is enabled for the job.

**Note: CIFS shares and displayed storage:** Since Windows Explorer is not aware of hard links (used with deduplication) CIFS share directories will not display used space accurately when viewing folder properties. To see the actual disk usage for a CIFS share directory you can download the [Disk Usage](#) utility from the Microsoft web site and use the -u option when displaying disk details.

## Job Speeds, Deduplication, and Data Written

The average job speed displayed in the console is calculated by dividing the total time the job ran by the total data processed. Therefore, if you had a single backup job for a 20 GB Windows XP VM that took 4 minutes to run, you would see an average speed of about 83 MB per second (20,000 MB / 240 seconds = 83.3333 MB/s).

The DeDupe (or deduplication) ratio for each job is determined by calculating the ratio of the total job data for all VMs in the backup job to the actual data written to the backup storage. For example, our Windows XP example backup job included 20 GB of total data. After deduplication and compression, only 100 MB of data was written to the backup store when the backup ran, resulting in a ratio of 200:1. The 100 MB is then reported as the Data Written in the Job Details for our example job. Note that Data Written reports only the actual amount of data written to the backup store - it does not include the total data of all VMs in the job.

With CBT enabled for a backup job, backup speeds will be much faster, as only the changed data is read and written for each VM. In our Windows XP job example, with CBT enabled, the initial backup took 4 minutes to process the entire 20 GB disk. The next time the backup job ran with CBT enabled, since only a small amount of data had changed on the VM, the job took only 30 seconds, and with deduplication and compression enabled, only 6.5 MB of data was actually written to the backup store. For additional details on CBT see "[PHD Virtual Backup and Changed Block Tracking](#)" on page 12.

## Job Types


PHD Virtual Backup Job types include:

- Backup Now
- Backup Daily
- Backup Once
- Backup Weekly
- Restore Now
- Delete Now

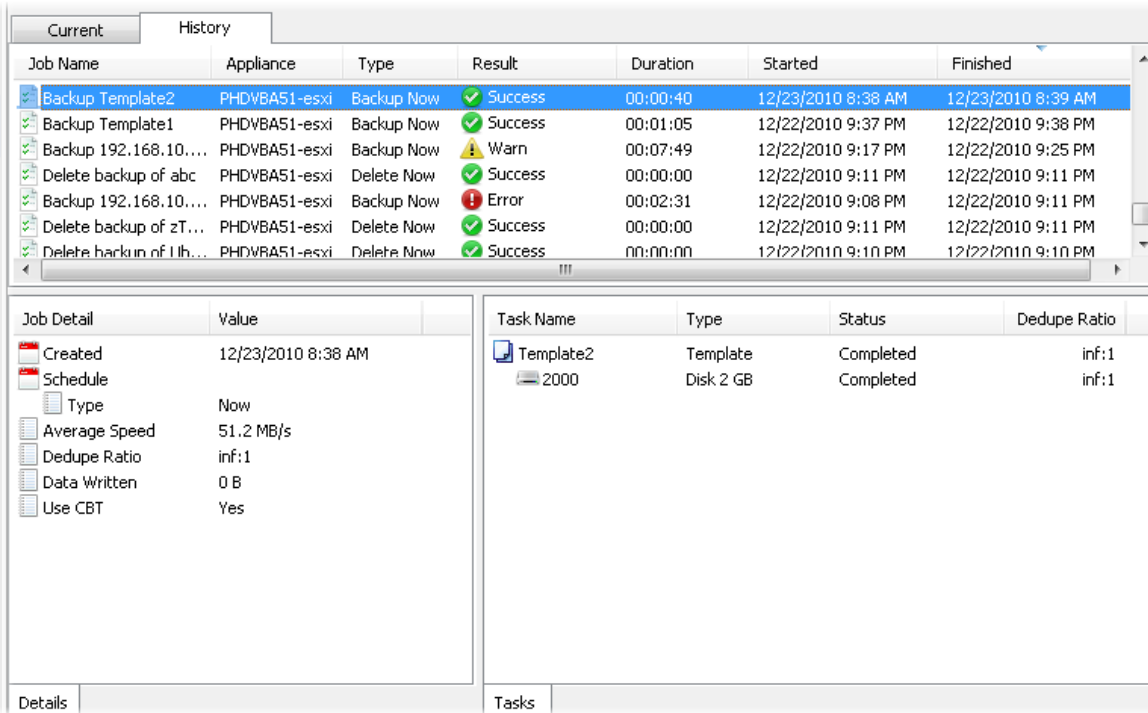
System Jobs include:

- **Startup** - The job that runs when the appliance first starts. This job cleans up any unfinished processes as well as synchronizes the backup catalog with the backup storage.
- **Orphan Weekly** - A weekly job that runs each Sunday at 7 AM to reclaim storage space used by unique and unreferenced blocks created during a backup that did not complete (failed backup, canceled backup, appliance shutdown, etc.).
- **Delete trim** - The system job that removes older backups based on your archive retention policy settings. See "[Retention](#)" on page 53 for details on setting your retention policy.
- **Snap Hunt** - A system job that runs once on PHD VBA start up and also once daily to remove any snapshots that may have been left behind by any PHD VBA.

## Job History

The Jobs page also contains a History tab that lets you see all of the jobs that have completed. Clicking **Show Details**  will display the detailed information about the completed jobs.

History information is retained for 90 days (it may be available for up to 120 days).



The screenshot shows the 'History' tab in the software interface. It features a table of job history and two detail panels below it.

Job Name	Appliance	Type	Result	Duration	Started	Finished
Backup Template2	PHDVBA51-esxi	Backup Now	Success	00:00:40	12/23/2010 8:38 AM	12/23/2010 8:39 AM
Backup Template1	PHDVBA51-esxi	Backup Now	Success	00:01:05	12/22/2010 9:37 PM	12/22/2010 9:38 PM
Backup 192.168.10...	PHDVBA51-esxi	Backup Now	Warn	00:07:49	12/22/2010 9:17 PM	12/22/2010 9:25 PM
Delete backup of abc	PHDVBA51-esxi	Delete Now	Success	00:00:00	12/22/2010 9:11 PM	12/22/2010 9:11 PM
Backup 192.168.10...	PHDVBA51-esxi	Backup Now	Error	00:02:31	12/22/2010 9:08 PM	12/22/2010 9:11 PM
Delete backup of zT...	PHDVBA51-esxi	Delete Now	Success	00:00:00	12/22/2010 9:11 PM	12/22/2010 9:11 PM
Delete backup of 1lh...	PHDVBA51-esxi	Delete Now	Success	00:00:00	12/22/2010 9:10 PM	12/22/2010 9:10 PM

Job Detail	Value
Created	12/23/2010 8:38 AM
Schedule	
Type	Now
Average Speed	51.2 MB/s
Dedupe Ratio	inf:1
Data Written	0 B
Use CBT	Yes


Task Name	Type	Status	Dedupe Ratio
Template2	Template	Completed	inf:1
2000	Disk 2 GB	Completed	inf:1


## Configuration

The Configuration page of the PHD Virtual Backup Console contains all of the options to configure your PHD Virtual Backup Appliances.

**Tip:** To access the console, you can right click any VM and select **PHD Virtual Backup > Console**.

Each PHD VBA must be configured separately; the menu at the top of the Configuration page indicates which PHD VBA's settings are displayed, as seen in the following image.

Select the appliance to configure: PHDVBA 

You can reload the values for any changed configuration area before saving them by clicking the refresh button  to the right of the appliance selection menu.

**Note:** The **Hypervisor Credentials** on the General tab and the **Backup storage** selection on the Storage tab are the only configuration options that are required to run backups. All of the additional settings are optional.

The Configuration page contains multiple tabs, described in the following sections:

- "General" on page 45
- "Storage" on page 47
- "Network" on page 49
- "Email" on page 51
- "Retention" on page 53
- "Connector" on page 56
- "Support" on page 58

## General

The General tab contains appliance options including the time zone, Data Streams, Hypervisor Credentials, and License information for the currently selected PHD Virtual Appliance.

The screenshot shows the configuration interface for a PHD Virtual Backup Appliance (PHDVBA). At the top, a dropdown menu is set to 'PHDVBA'. Below this are several tabs: 'General', 'Storage', 'Network', 'Email', 'Retention', 'Connector', and 'Support'. The 'General' tab is active and contains three main sections:

- Appliance options:**
  - Select time zone: America
  - Select region: New\_York
  - NTP Server 1: ntp.ubuntu.com
  - NTP Server 2: (empty)
  - Data Streams: A slider set to 4.
- Hypervisor credentials:**
  - vCenter Server: 192.168.40.50 (with a note: e.g., server.example.com or IP address) and Port: 443
  - User Name: administrator
  - Password: (masked with dots)
- Professional License: PHD Virtual:**
  - Product Expiration: Friday, November 18, 2011
  - Support Expiration: Friday, November 18, 2011
  - An [Update](#) link is present.

A 'Save' button is located at the bottom right of the configuration window.

### Appliance options

- The **time zone** and **region** defined here affect when each job will run. Scheduled jobs will run according to the time in the configured time zone, which may not be the same time zone as your desktop or host server.
- **NTP servers** are used to synchronize the time on multiple computers. You can configure up to two NTP servers here to synchronize each PHD Virtual Backup Appliance.
- **Data Streams** perform the individual job processes on the appliance. The Data Streams slider lets you set the number of processes that will operate concurrently while a job is in progress. For example, when set to four, up to four virtual disks can be processed at once during a backup job. In some cases, with older or slower hardware, you may need to reduce the number of threads to avoid saturating host server resources. If you are experiencing performance issues, you can reduce the number of streams used by the appliance at one time by moving this slider to the left.

### Hypervisor Credentials

Hypervisor Credentials are used by the PHD Virtual Backup Appliance to perform the steps required to backup and restore virtual machines.

If you are using vCenter to manage your environment, enter your vCenter Server name or IP address in the **vCenter Server** text box. Unless you are using a non-standard port to communicate with your server, leave the default port set to 443.

For **User name** and **Password**, enter your vCenter administrator credentials.

If you are using a standalone ESX/ESXi host, enter that server's fully qualified name or IP address and the administrator credentials for that host.

### License

PHD Virtual Backup is installed with a trial license. To avoid any interruption in your ability to run backups, you will need to upload a new license before the trial period expires.

**To update your PHD Virtual Backup license**, click **Update** in the **License** area to apply the new license file. New licenses must be applied to each PHD Virtual Backup Appliance you have deployed. Use the menu at the top of the Configuration window to select each appliance to update.

- The **Product expiration** date displays when PHD Virtual Backup expires. After the product expiration date, you can no longer run backups, but you can still restore your backed up files.
- The **Support expiration** date determines when your support license expires. A valid support license is required to install product upgrades.

## Storage

The storage tab is used to define where your backups are sent. Backups can be sent to an attached virtual disk, a CIFS share, or an NFS share.

The storage currently in use is shown in the **Backup storage** area.

Select the appliance to configure: PHDVBA

General | **Storage** | Network | Email | Retention | Connector | Support

**Backup storage**

Storage Type: Attached Virtual Disk

✓ Using attached disk 10 GB

**Advanced options**

Enable compression for new backups

Warning level % free: 10.00 Warns at 1 GB of free storage

Stop level % free: 3.00 Stops at 307.2 MB of free storage

Reset to Defaults

Save

To run backups, storage must be defined when the appliance is first deployed and configured. If you need to change your storage location later, you can do so using the Storage tab.

**Note:** Backup storage cannot be shared by multiple PHD VBAs - each VBA must use its own unique storage location.

### To change the backup storage location

1. Open the PHD Virtual Backup Console and click **Configuration**.
2. Click the **Storage** tab.
3. From the **Storage Type** menu, select the type of storage to use. If you select to use an NFS or CIFS share you must also enter the share location and credentials the appliance should use.
4. Click **Save**.

### Advanced storage options

Advanced options include compression and settings for storage level warnings.

- **Enable compression for new backups** - enabled by default, this option instructs PHD Virtual Backup to use compression when creating backups. If you have a reason to store backup data uncompressed, you can disable this option. For example, if you have a large amount of storage available and need to increase backup speeds, you can disable this option to skip the compression.
- **Warning level % free** - use this option to set the threshold at which you would like to display a warning that your backup storage is running low on available free space.
- **Stop level % free** - use this option to force PHD Virtual Backup to stop running backups when the free storage capacity reaches this threshold.

**Note:** CIFS and NFS shares may have additional free space thresholds defined that, when exceeded, could potentially prevent new backups from completing. Check with your local administrator for details.

## Network

Use the Network tab to define a PHD Virtual Backup Appliance's network settings. By default, the appliance will attempt to obtain an IP address automatically after it is deployed.

Select the appliance to configure: PHDVBA-514-ovf

General | Storage | **Network** | Email | Retention | Connector | Support

Adapter | **Storage Adapter**

MAC Address: 00:50:56:94:00:40

Obtain an IP address automatically

Use the following IP address

IP address: . . .

Subnet mask: . . .

Gateway: . . .

Name Servers

Obtain DNS address automatically

Use the following DNS addresses

Preferred DNS: . . .

Alternate DNS: . . .

Save

If necessary, you can configure a second network adapter to be used by the PHD VBA. This may be useful if your storage location exists on a network that is unreachable by your VM management network (where the PHD Virtual Backup Appliance resides). To use multiple network adapters you first need to add a second adapter to the PHD VBA virtual machine then use the PHD Console to configure the second adapter. See ["Using Multiple Network Adapters" on page 90](#) for details.

**Note:** If you are experiencing network problems you can manually assign network settings by selecting the VBA within vSphere Client then clicking the Console tab and typing Ctrl-N. Note that this method allows you to configure settings for the first network adapter, only. If configured, the settings for the second adapter will be reset to obtain an IP address automatically.

The next few sections describe how to use the Network tab to configure the network settings for your PHD VBAs.

## Using DHCP

By default, each PHD VBA will attempt to acquire an IP address automatically using DHCP. If you had set a PHD VBA to use a static address, but would like to switch to using DHCP, follow the steps below.

### To obtain the appliance IP address automatically

1. Open the PHD Virtual Backup Console and click **Configuration**.
2. Select the appliance you want to configure from the **Select the appliance to configure** menu at the top of the page.
3. Click the **Network** tab and select **Obtain an IP address automatically**.
4. When obtaining an address automatically, you also have the option to obtain DNS information automatically by selecting **Obtain DNS address automatically**, or you can specify your DNS settings. When complete, click **Save**.

## Using Static IP Addresses

PHD VBAs can be configured to use static IP addresses using the Network tab of the PHD Console's Configuration area.

### To assign static appliance network settings

1. Open the PHD Virtual Backup Console and click **Configuration**.
2. Select the appliance you want to configure from the **Select the appliance to configure** menu.
3. Click the **Network** tab and select **Use the following IP address**.
4. Enter your IP address, Subnet mask, and Gateway.
5. When manually assigning networking information, you must also define your DNS settings. Enter a preferred and alternate DNS address, then click **Save**.

## Email

Use the Email tab to configure PHD Virtual Backup to send email alerts and reports.

You can select to send email alerts for Critical errors, Errors, or All, which includes backup and restore job results, system alerts, and errors. Warnings are not sent as separate email alerts.

Select the appliance to configure: PHDVBA

General | Storage | Network | **Email** | Retention | Connector | Support

Do not email alerts from the appliance  
 Email alerts using the following information

Server Name:  Port:   
 Security:   
 Server requires credentials  
 User name:   
 Password:   
 From Email Address:   
 Alert Level:   
 Recipients:

### To enable alerts

1. Open the PHD Virtual Backup Console and click **Configuration**.
2. Select the appliance you want to configure from the **Select the appliance to configure** menu.
3. Click the **Email** tab then select **Email alerts using the following information**:
4. Enter the IP address or FQDN of the email server you would like to use to send email alerts.
5. If your email server requires security, select the type from the **Security** menu.
  - **None** - do not use security.
  - **STARTTLS** - use STARTTLS security when sending email alerts.
  - **SMTP over SSL** - use SMTP over SSL when sending email alerts.

6. If the server requires authentication, select the checkbox and enter a username and password.
7. Enter a **From Email Address** (this is the address the PHD Virtual Backup emails will come from).
8. Select the **Alert Level**
  - **All** - send all email alerts, including backup and restore job results and all system level errors (warnings are not sent as email alerts though they are included in the backup and restore reports).
  - **Errors** - send an email alert only for errors (Error and Critical Error).
  - **Critical** - send an email alert only for critical errors.
9. Click **Add** to add the email addresses that will receive the email alerts. When added, the addresses will be displayed within the **Recipients** dialog box. To remove any email addresses, select the address in the **Recipients** dialog and click **Remove**.
10. When you are finished configuring email alerts, click **Save**.

#### To disable email alerts

1. Open the PHD Virtual Backup Console and click **Configuration**.
2. Select the appliance you want to configure from the **Select the appliance to configure** menu.
3. Click the **Email** tab and select **Do not email alerts from the appliance** then click **Save**.

## Retention

Use the Retention tab to define your backup retention policy.

Select the appliance to configure: PHDVBA

General Storage Network Email **Retention** Connector Support

Retention

Retention setting: Typical

Recent backups to keep: 5

And keep the most recent backup from each of the last:

Days: 7

Weeks: 4

Months: 12

Years: 5

Save

By default, PHD Virtual Backup will keep all backups for each VM. Using the Retention options, you can select how many backups you want to keep for each virtual machine to meet your individual compliance and storage requirements. When a retention policy is set, a job runs (Delete trim) and performs the retention processing at the top of each hour.

You can select to use pre-defined settings, or you can set specific values for each setting. The available **Retention Settings** are:

- **Keep All** - Retain all backups for all VMs. This is the default setting.
- **Typical** - Retain the 5 most recent backups as well as the most recent backup from each of the last 7 days, 4 weeks, 12 months, and 5 years.
- **Custom** - You define the values for each retention setting.

### Retention Notes

- **Days** start at 00:00:00 and include the current day.
- **Weeks** start on Monday and include the current week.
- **Months** are based on the calendar month and include the current month.
- **Years** are based on the calendar year and include the current year.
- Retention adjusts for Daylight Savings Time.
- Backup files marked as Archive will never be deleted.

### To define backup retention settings

1. Open the PHD Virtual Backup Console and click **Configuration**.
2. Select the appliance you want to configure from the **Select the appliance to configure** menu.
3. Click the **Retention** tab then use the **Retention setting** menu to select your retention policy.
4. When finished, click **Save**.

### To keep only a certain number of backups per VM

1. Open the PHD Virtual Backup Console and click **Configuration**.
2. Select the appliance you want to configure from the **Select the appliance to configure** menu.
3. Click the **Retention** tab then use the **Retention setting** menu to select **Custom**.
4. Set the **Recent backups to keep** to the number of backups you would like to keep for each VM. For example, to keep only 5 backups for each VM, set this value to 5.
5. Set the **Days**, **Weeks**, **Months**, and **Years** values to 0.
6. When finished, click **Save**.

Now, only the last five backups will be kept for each VM.

### Advanced Retention Scenario

The following example scenario describes how backups are retained when using advanced retention settings. We will assume the following:

- Today is 10/29/2010
- Backup Frequency is set to Daily (and the daily backup has run today)
- Backups have been collected for the last 5 years
- Retention Settings set to Custom with Recent backups set to 3, Days set to 0, Weeks to 5, Months to 13, and Years to 3. The following image illustrates the current settings.

**Retention**

Retention setting Custom ▼

Recent backups to keep

And keep the most recent backup from each of the last:

Days

Weeks

Months

Years

The following table describes the backups that will be retained based on this scenario.

Backup Period	Retention Setting	Backups Retained (by date)	Unique Backups
Most Recent	3	10/29, 10/28, 10/27	3
Days	0		0
Weeks	5	10/29*, 10/24, 10/17, 10/10, 10/3	4
Months	13	10/29*, 9/30, 8/31, 7/31, 6/30, 5/31, 4/30, 3/31, 2/28, 1/31, 12/31/09, 11/30/09, 10/31/09	12
Years	3	10/29/2010*, 12/31/2009*, 12/31/2008	1
<b>Total Backups Retained</b>			<b>20</b>

\* Backup already retained; not unique.

## Connector

Use the Connector tab to enable and configure the Backup Data Connector (BDC) to export backups.

The screenshot shows the configuration interface for the Backup Data Connector (BDC) in the PHD Virtual Backup Console. At the top, there is a dropdown menu labeled "Select the appliance to configure:" with "PHDVBA" selected. Below this are several tabs: "General", "Storage", "Network", "Email", "Retention", "Connector", and "Support". The "Connector" tab is currently active. Inside the "Connector" tab, there is a section titled "Backup Data Connector" with the following options:

- Enable share at \\192.168.40.52\backups
- User name:
- Set Password:
- Confirm Password:

At the bottom right of the configuration area, there is a "Save" button.

The Backup Data Connector lets you access backups in an uncompressed format which can be useful if you need to save backups to tape or archive backups to disk. With the connector, you enable an SMB/CIFS share that allows access to your backup files in a simple folder structure. You can then use third-party tools or your own scripting to compress, select and move these files to tape or to other disk locations, as necessary.


**Note:** VMDK files available from the Backup Data Connector share will always be hardware version 7 (regardless of the version at time of backup).

### To access backups using the Backup Data Connector

1. Open the PHD Virtual Backup Console and click **Configuration**.
2. Select the appliance you want to configure from the **Select the appliance to configure** menu.
3. Click the **Connector** tab.

4. Select **Enable Share at....** This will display your appliance IP address and the share name, for example, \\192.168.1.100\backups.
5. Enter and confirm a password. The default username *phd* cannot be changed.
6. When finished, click **Save**.

When enabled, you can access the share to view the uncompressed backups, as seen in the example image, below.



The folders in the share organize backups into categories by when each backup was taken.

- **Last7Days** - All backups taken within the last seven days, not including today.
- **Last32Days** - All backups taken within the last 32 days, not including today.
- **LatestofEach** - The latest backup file for each VM available.
- **Today** - All backups taken today.
- **Yesterday** - All backups taken yesterday.

**Note:** After a backup has finished, some larger backup files may not be visible in the BDC share, right away. If you do not see a backup that has recently finished, wait a few moments then refresh the share.

In addition to accessing backup files through the Backup Data Connector share, you can manually export individual backups using the Export backup feature. See "[Backup Catalog](#)" on page 30 for details.

**Note:** If you experience problems connecting to the Backup Data Connector share, you may need to adjust the local security policy on your Windows computer. See "[BDC Share and Local Security Policies](#)" on page 99.

## Support

Use the Support tab to configure debugging logs, to download support files, apply updates to the PHD Virtual Backup Appliances, and to find current version information.

The screenshot shows the 'Support' tab of the PHD Virtual Backup Console configuration interface. At the top, there is a dropdown menu labeled 'Select the appliance to configure:' with 'PHDVBA' selected. Below this are several tabs: 'General', 'Storage', 'Network', 'Email', 'Retention', 'Connector', and 'Support'. The 'Support' tab is active and contains the following content:

- Enable debug logging on appliance  
Debug logging provides additional diagnostic information.
- Diagnostics**
  - [Download Support File](#)  
The support file contains information useful when diagnosing appliance problems.
  - [Download Console Logs](#)  
The console logs contain useful information when diagnosing console problems.
- Version Information**
  - PHD VBA Version: 5.2.0
  - PHD Console Version: 5.2.0
  - Patches are bundles downloaded from the PHD Virtual support website that contain updates for your appliance.
  - [Upload Appliance Patch](#)

A 'Save' button is located at the bottom right of the configuration area.

### Debug Logging

Enabled by default, **Enable debug logging on appliance** includes additional diagnostic log messages in the PHD Virtual Backup logs. These messages are useful when troubleshooting product issues.

### Support Files

When communicating with PHD Virtual Support, you may be asked to download and send support files to help resolve any issues. Use the links in the **Diagnostics** area to do this. A compressed package will be downloaded and can then be sent to PHD Virtual, if requested.

### Uploading Appliance Patches


Periodically, update patches for the PHD Virtual Appliance will be available for download from the PHD Virtual Web site. When downloaded to your local computer, they can be uploaded through the PHD Virtual Backup console using the **Upload Appliance Patch** link. Clicking this link will allow you to select the downloaded appliance patch file. For additional information, see "Updating PHD Virtual Backup" on page 93.

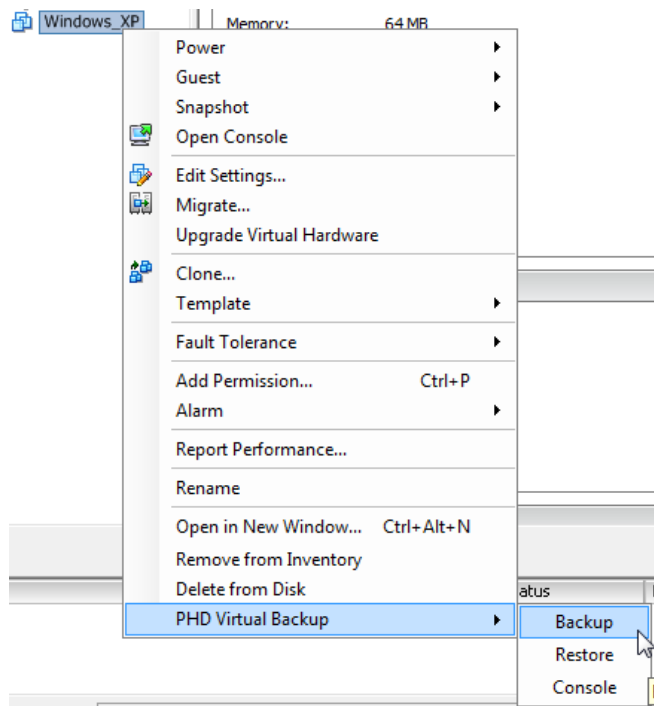
## Chapter 4 - The Backup Wizard

The Backup Wizard guides you through the process of creating backup jobs to protect the virtual machines in your environment. The following sections describe how to access and use the wizard.

Accessing the Backup Wizard.....	60
Using the Backup Wizard.....	61

## Accessing the Backup Wizard

There are multiple ways to start the wizard, using either the PHD Console or the integrated menus within vSphere Client. From the Console, **Jobs** area, you can click  **Backup**. Within vSphere Client, you can right-click a VM name then select **PHD Virtual Backup > Backup** from the integrated menu, as shown in the following image.



When opened, the wizard presents the steps required for backing up your virtual machines. See "Using the Backup Wizard" on page 61 for details.

## Using the Backup Wizard

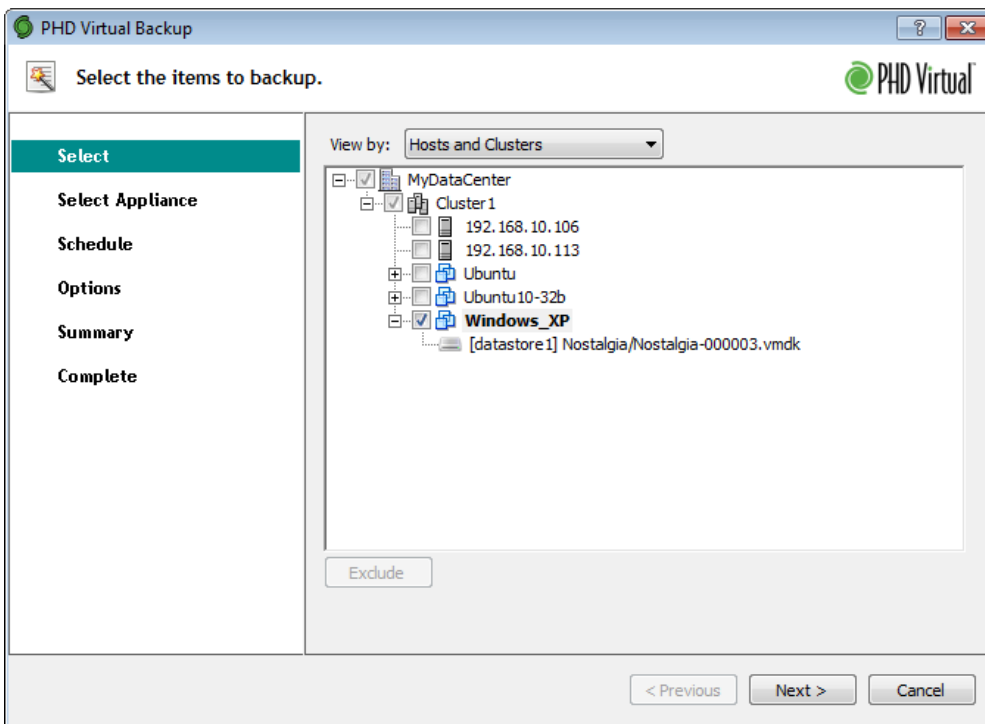
The following procedure provides detailed information about each step of the Backup Wizard.

### Using the Backup Wizard

1. When the wizard opens, you are presented with the **Select** step. Here you can use the **View by:** menu to change how the virtual machines available for backup are displayed.

**Hosts and Clusters** - Display all VMs based on the Hosts and Clusters (containers) they belong to.

**VMs and Templates** - Display only VMs and Templates.




If you select the top container in any view (for example, Cluster1 in the image above) all VMs in the container will be included in the backup job. Also, after the job is created, any VMs added to or removed from the selected container will also be included or excluded from the backup job, respectively.

- **Exclude/Include** - When backing up groups of VMs, an entire folder, for example, you can choose to exclude specific VMs or individual disks from the backup job by selecting the VM or disk and clicking **Exclude**. Excluded VMs can be included again select the VM and clicking **Include**.

2. Select the VMs you want to backup and click **Next**.

**Note:** By design, the PHD Virtual Backup Appliance is not included in the list of VMs you can select for backup.

3. At the **Select Appliance** step, select the PHD VBA you want to use to perform the backup.



The backup wizard searches for all available appliances within the current resource pool. The appliance you select will perform the backup processing and store the backup file on its configured storage location.

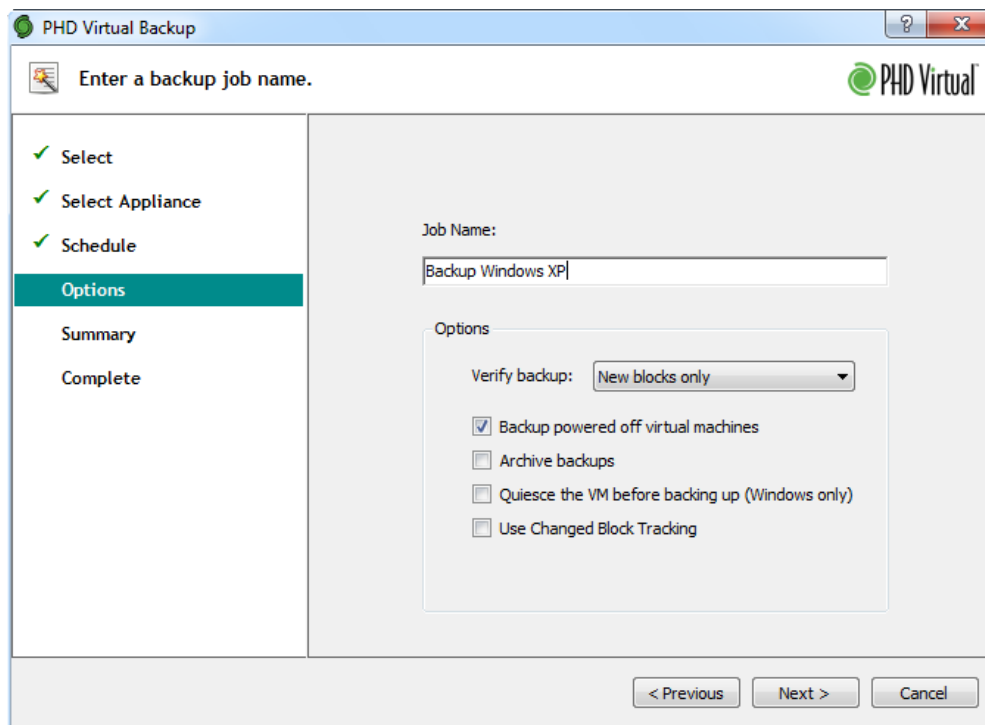
**Note:** If you will be backing up a VM located on local storage, you must select an appliance that is located on the same host as the VM or else the backup will fail. Virtual disks for any VMs that are unreachable by an appliance (on different local or shared storage, for example) will be displayed after the appliance is selected. You can then choose to click Previous and exclude those VMs or disks or select another appliance with access to those disks.

4. When you've selected the PHD VBA, click **Next**.
5. The **Schedule** step lets you run a backup **Now**, schedule a backup **Once** for later, create a **Daily** backup or a **Weekly** backup. Select the type of backup to create and define any required options and click **Next**. For additional details on scheduling backup jobs, see "[Scheduling Backups](#)" on page 75.

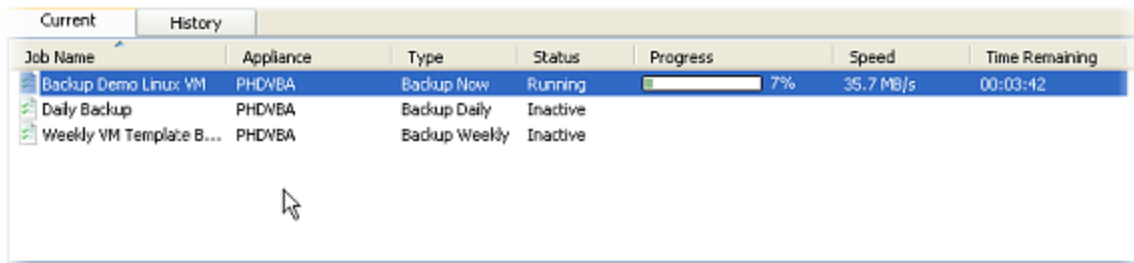
- **Start Date**- The date the scheduled job will begin.
- **Start Time**- The time the job should start.
- **Do not start after**- The time after which the job should not start. In a situation where many or very large backup jobs are running, and this time passes before the job can begin, it will not start until the next scheduled start time. Jobs already in progress after this time will not stop - they will complete as normal.
- **Recurs every *n* Days/Weeks**- How often the job will run. A daily job, by default, will run once per day. If you'd like a job to run every other day, set this to 2, for example. Weekly jobs will run once per week, by default. To create a job that runs only once every two weeks, select a Weekly job then set this value to 2. Recurring jobs begin based on the first day of each month. For instance, if you create a daily job that recurs every 10 days, it will run on the first of the month, the eleventh, the twenty-first and the thirty-first, if available. This schedule is reflected in the **Next Run** date within the Job Details. Therefore, if on August 19th you created a daily job that recurs every 10 days, the Next Run date will be August 21st. Though this may appear to be only two days from the day the job was created, it represents the third recurrence date of the job for that month (1st, 11th, 21st, and 31st).

6. Select the type of backup to create, then click **Next**.

7. The **Options** step lets you name the backup job and define options specific to the backup.



- **Verify backup** - This option determines how the backup should be verified. By default, this is set to **None** which means data is written, but not verified. **All blocks** instructs PHD Virtual Backup to verify every block of data (for a scheduled job, this would happen every time the backup job runs). **New blocks only** verifies only the information that has changed since the last backup. For additional information on the verify options, see "[Verifying Backups and Restores with TrueRestore™](#)" on page 83.
  - **Backup powered off virtual machines** - Select this check box to backup VMs included in the backup job even if they are powered off.
  - **Archive backups** - Select this option to flag backups created with this job as archived backups. This means the backups will never be deleted by the automatic retention policy. Archived backups also cannot be manually deleted. To remove an archive flag, or to archive existing backups, see the Backup Catalog in the console.
  - **Quiesce the VM before backing up (Windows only)** - When backing up a Windows VM, if VMware tools are installed, you can choose to quiesce the VM before backing it up, to take advantage of Microsoft's Volume Shadow Copy Services.
  - **Use Changed Block Tracking** - Enabled by default, this option lets you take advantage of VMware's Changed Block Tracking when performing backups. Enabling this feature instructs PHD Virtual Backup to use the VMware vSphere API to keep track of only the disk sectors that have changed between backups. Instead of reading the entire VMDK each time to discover changes, only the changed blocks are read to establish backups, saving significant time in the process. **Note:** VMs must be hardware version 7 to support Changed Block Tracking - VMs that are not hardware version 7 will still be backed up but a warning will be logged. When this option is selected, Changed Block Tracking is enabled on each individual VM in the job, if not enabled already. If Change Block Tracking is enabled on a VM but this option is not selected, a regular backup is performed (no Changed Block Tracking). VM templates do not support changed block tracking.
8. When you've finished adding a job name and selecting job options, click **Next**.
  9. Review the **Summary** information, then click **Submit**. The backup job is then submitted for processing. Click **Finish** to close the wizard.
  10. The PHD Virtual Backup Console opens and displays the status of the backup job.



Job Name	Appliance	Type	Status	Progress	Speed	Time Remaining
Backup Demo Linux: VM	PHDVBA	Backup Now	Running	<div style="width: 7%;"><div style="width: 7%;"></div></div> 7%	35.7 MB/s	00:03:42
Daily Backup	PHDVBA	Backup Daily	Inactive			
Weekly VM Template B...	PHDVBA	Backup Weekly	Inactive			

For more information on using the Jobs area of the console, see "Jobs" on page 39.

## Chapter 5 - The Restore Wizard

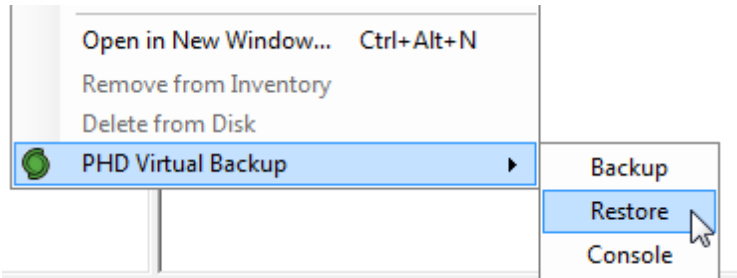
The Restore Wizard lets you restore the virtual machines you backed up with PHD Virtual Backup. Restored VMs include the metadata and all of the virtual disks associated with the VM. For additional details about what attributes are included with each restored virtual machine, see "Restores" on page 15.

The following sections describe how to access and use the wizard.

<a href="#">Accessing the Restore Wizard</a> .....	67
<a href="#">Using the Restore Wizard</a> .....	68

## Accessing the Restore Wizard

The wizard can be started by right-clicking an object within vSphere Client or by using the File menu and selecting **Restore** from the integrated **PHD Virtual Backup** menu, as shown in the following image.



The wizard can also be accessed using the PHD Console, either from the **Backup Catalog** or from the **Jobs** area.

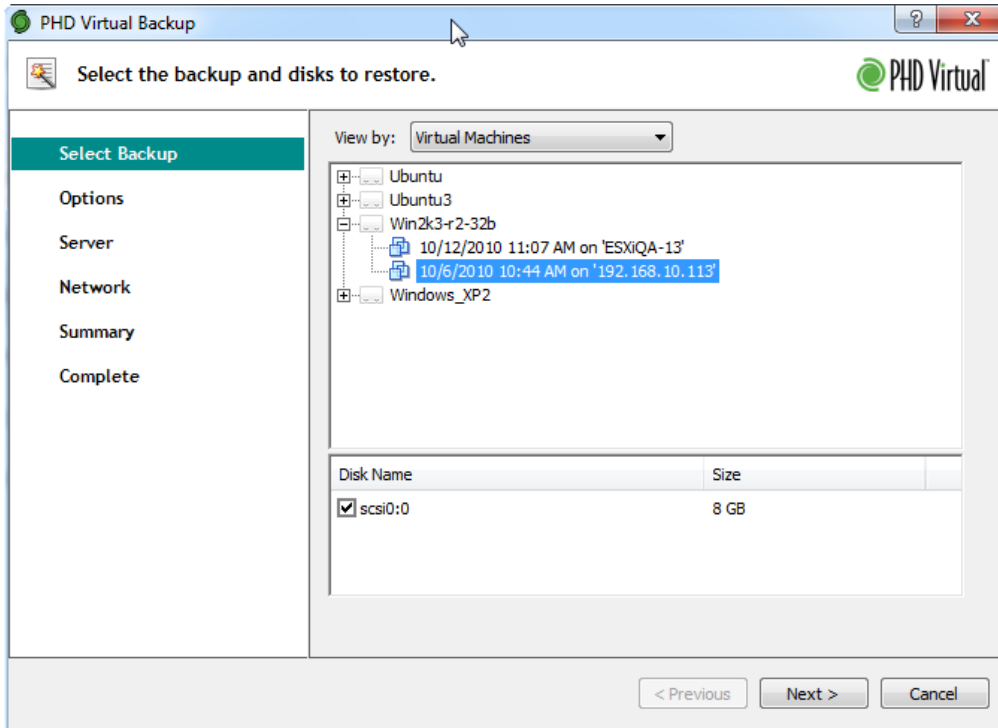
When opened, the wizard presents the steps for restoring your virtual machine backups. See ["Using the Restore Wizard"](#) on [page 68](#) for details.

## Using the Restore Wizard

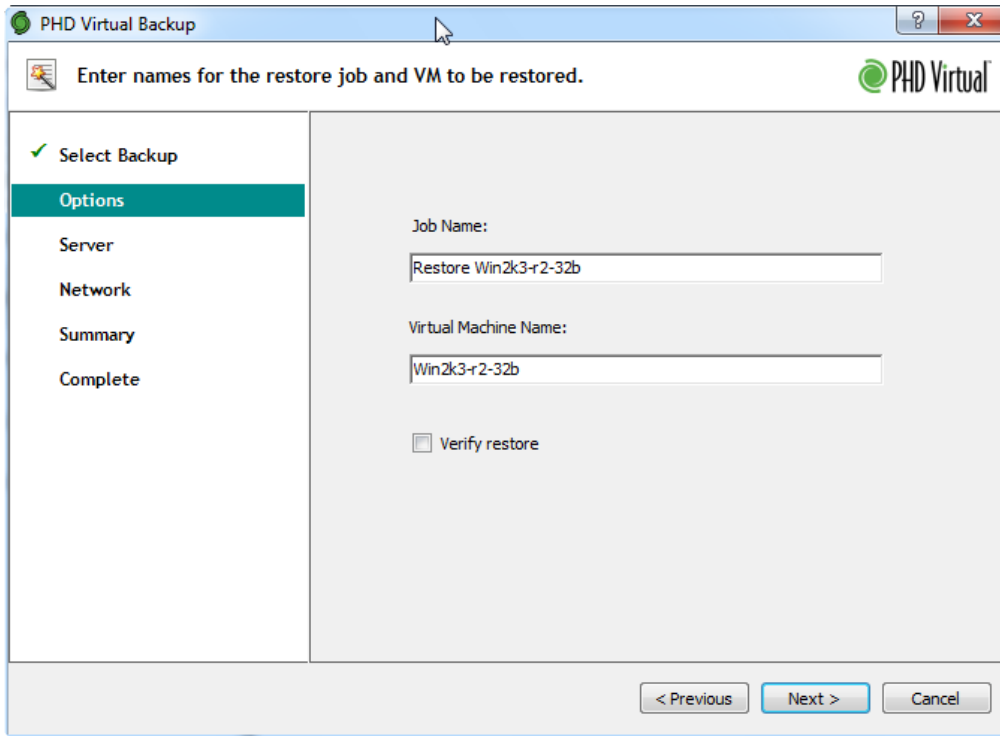
The following procedure provides detailed information about each step of the Restore Wizard.

### Using the Restore Wizard

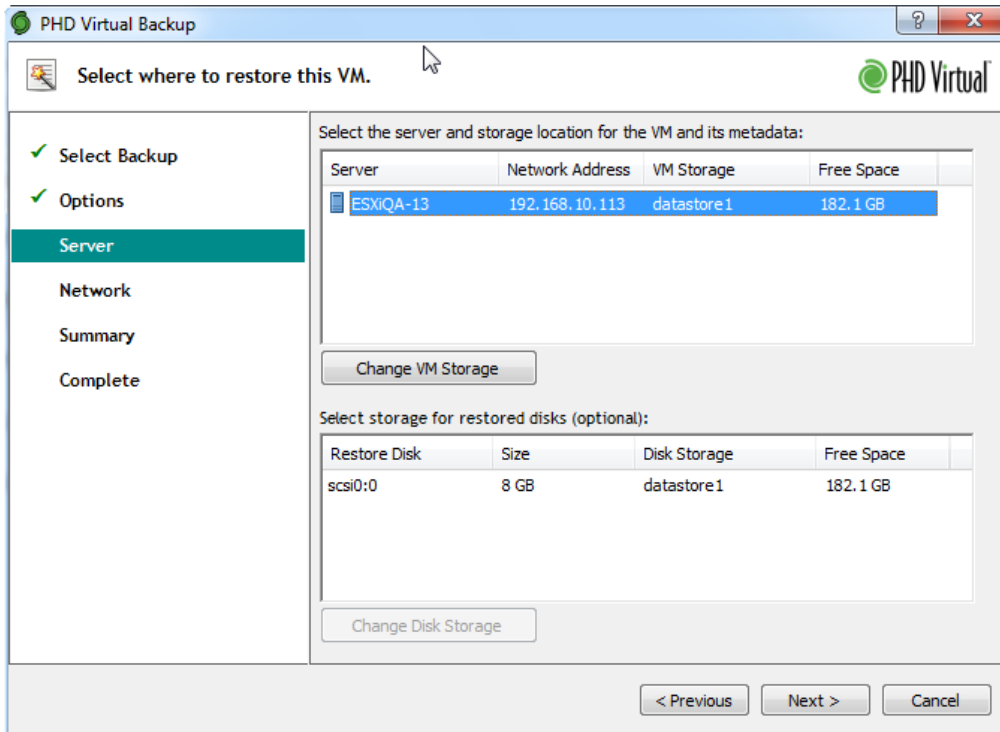
1. When the Restore Wizard opens, the **Select Backup** step presents you with all of the backups available for restore. Use the **View by** menu and the navigation tree to locate the backup you'd like to restore.



2. When you select the backup, the available disks are also displayed, as seen in the image above. All available disks are selected for restore by default. To exclude a particular disk from the restore, clear the check box in the **Disk Name** column. After selecting the VM and disks to restore, click **Next**.
3. At the **Options** step, enter a name for the restore job and a name for the VM to be restored.



4. If you want to add additional verification during the restore process, select **Verify Restore**. For more information on verifying backups and restores, see ["Verifying Backups and Restores with TrueRestore™"](#) on page 83. When ready, click **Next**.
5. The next step lets you select where the VM should be restored. Select the server from the available list (be sure to select a location with enough free space).



6. If you need to send an individual disk somewhere other than the selected server's default storage, select the disk and click **Change Disk Storage**. When you've selected where you will restore your VM and disks, click **Next**.
7. At the **Network** step, select the network device to use for the restored VM. If you need to change any settings, select a network interface, then click **Edit**. The Edit dialog lets you change the Network and MAC address used by the network interface. Click **Next** when complete.
8. Review the **Summary** information for the restore job, then click **Submit**. Click **Finish** to close the wizard.

Use the Jobs area of the PHD Virtual Backup Console to view the progress of the restore job. When the restore is complete, the VM is available within vSphere Client.

## Chapter 6 - Using PHD Virtual Backup

The topics in this chapter include short reference topics as well as step-by-step instructions for using PHD Virtual Backup features.

Creating Backup Jobs.....	72
Running a Backup Now.....	73
Scheduling Backups.....	75
Viewing Jobs.....	77
Restoring Backups.....	79
Restoring Files.....	80
Configuring Email Alerts.....	82
Verifying Backups and Restores with TrueRestore™.....	83
Backup Retention and Archiving.....	84
Excluding VMs and Disks.....	85
Sending Backup Files to Tape.....	86
Limiting the PHD Console to a Single PHD VBA.....	87
Increasing Backup Storage (Attached Disk).....	89
Using Multiple Network Adapters.....	90
Updating PHD Virtual Backup.....	93

## Creating Backup Jobs

PHD Virtual Backup protects your virtual machines using Backup Jobs that you create and customize. Jobs can be run immediately or they can be created with a schedule to backup VMs every night, for example.

You can create backup jobs to protect individual virtual machines or you can create jobs to backup an existing VMware container (Resource Pool, Cluster, Folders, or Datacenter). When you create a job using a container, a Cluster, for example, VMs added to the cluster will be included in the job automatically, the next time the job runs. Likewise, if you remove a VM from that cluster, it will not be backed up the next time that job runs.

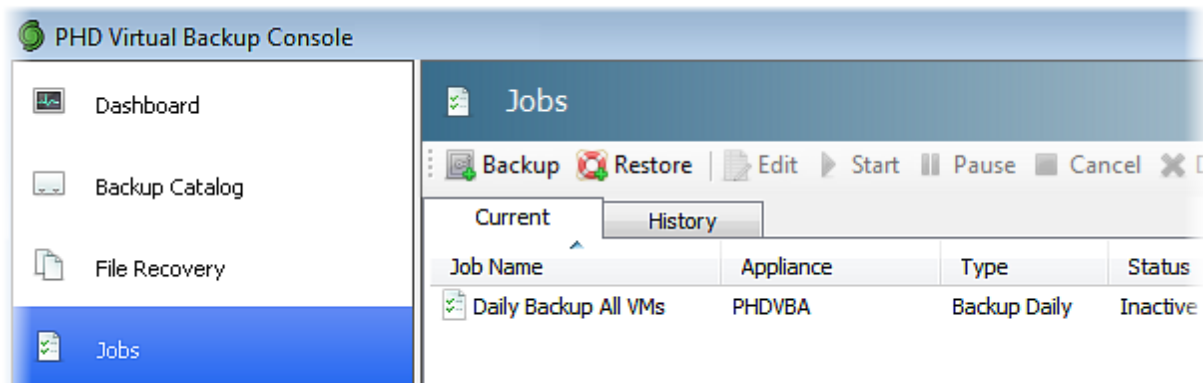
Backup jobs are created using the Backup Wizard. The wizard is accessed by selecting **PHD Virtual Backup > Backup** from the integrated PHD Virtual Backup menus in vSphere Client, or when you click **Backup** within the PHD Virtual Backup Console.

### To create a Backup Job

1. Start the Backup Wizard by right-clicking a VM and selecting **PHD Virtual Backup > Backup**.
2. Follow the steps in the wizard to select VMs for backup and define a backup schedule. For detailed information about each step in the wizard, see ["Using the Backup Wizard" on page 61](#).

### To edit a job

1. Start the PHD Virtual Backup Console, then click **Jobs**. The Current tab displays all jobs in progress as well as any scheduled jobs.



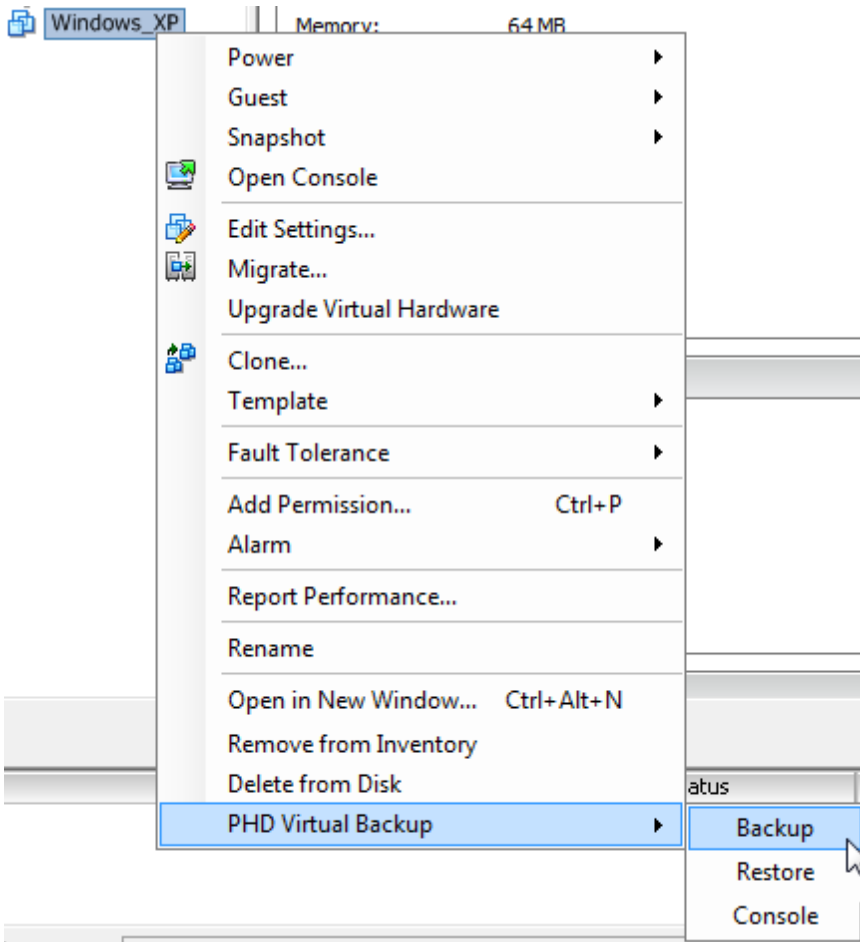
2. Select the job you would like to edit , then click **Edit**.
3. The Backup Wizard opens with the settings you originally defined for the job. Use the wizard to make any edits and submit the job again. For details on each step of the wizard, see ["Using the Backup Wizard" on page 61](#).

## Running a Backup Now

There are multiple ways to run a backup with PHD Virtual Backup - the easiest is to right-click a VM name within vSphere Client and select **Backup** from the PHD Virtual Backup context menu. This will open the Backup Wizard which guides you through the process of creating your Backup Job.

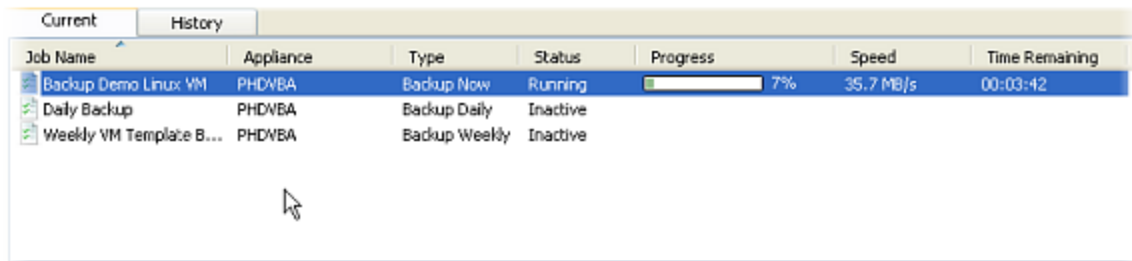
### To run a single backup

1. Within vSphere Client, right-click the name of the VM you want to backup.
2. From the context menu, select **Backup** from the PHD Virtual Backup menu.




The Backup wizard opens and guides you through the process of creating the Backup Job that will back up your selected VM. For detailed information about each step of the wizard, see ["Using the Backup Wizard" on page 61](#)


When the wizard completes, the PHD Virtual Backup Console opens and displays the progress of your backup job.



Job Name	Appliance	Type	Status	Progress	Speed	Time Remaining
Backup Demo Linux VM	PHDVBA	Backup Now	Running	<div style="width: 7%;"></div> 7%	35.7 MB/s	00:03:42
Daily Backup	PHDVBA	Backup Daily	Inactive			
Weekly VM Template B...	PHDVBA	Backup Weekly	Inactive			

**Tip:** Another way to run a backup right away is to force a scheduled backup to run now by selecting the job and clicking  **Start**.

### To run a scheduled backup now

1. Open the PHD Virtual Backup Console and click **Jobs**.
2. Click the scheduled job you want to run and then click  **Start**.
3. The job status changes from **Inactive** to **Running** and the backup begins.

When complete, the job remains in the Current tab and the status returns to Inactive, but the History tab will contain a record of the job you just ran.

## Scheduling Backups

Backups can be scheduled to run Once, Daily, or Weekly, using the PHD Virtual Backup Wizard.

### To create a scheduled backup job

1. From within vSphere Client, start the PHD Virtual Backup Wizard using the Pool, Server, or VM menu item: **PHD Virtual Backup > Backup**.
2. Use the check boxes to select the VMs to include in the scheduled backup job then click **Next**.
3. Select the appliance to use for the backup then click **Next**.
4. At the **Schedule** step, use the option buttons to set your schedule.

For example, to create a weekly backup schedule, select **Weekly**, then set the date to start the backups, the time the backups should be allowed to run, and the day of the week.

- **Start Date**- The date the scheduled job will begin.
- **Start Time**- The time the job should start.
- **Do not start after**- The time after which the job should not start. In a situation where many backup jobs or very large jobs are running and this time passes before the job can begin, it will not start until the next scheduled start time. Jobs already in progress after this time will not stop - they will complete as normal.
- **Rekurs every *n* Days/Weeks**- How often the job will run. A daily job, by default, will run once per day. If you'd like a job to run every other day, set this to 2, for example. Weekly jobs will run once per week, by default. To create a job that runs only once every two weeks, select a Weekly job then set this value to 2. Recurring jobs begin based on the first day of each month. For instance, if you create a daily job that recurs every 10 days, it will run on the first of the

month, the eleventh, the twenty-first and the thirty-first, if available. This schedule is reflected in the **Next Run** date within the Job Details. Therefore, if on August 19th you created a daily job that recurs every 10 days, the Next Run date will be August 21st. Though this may appear to be only two days from the day the job was created, it represents the third recurrence date of the job for that month (1st, 11th, 21st, and 31st).

5. When the schedule is set, click **Next**.
6. Enter a name for the job, for example, **Nightly Backup - Production VMs**.
7. Configure any job options. For more information, see "[Using the Backup Wizard](#)" on page 61. Then click **Next**.
8. Review the summary information then click **Submit**. Click **Finish** to close the wizard.

The selected VMs will be backed up based on the schedule you defined.

Use the Console, Jobs page to manage the existing scheduled backup jobs. From there you can run the job immediately to test your settings or edit the job details. See "[Jobs](#)" on page 39 for more information.

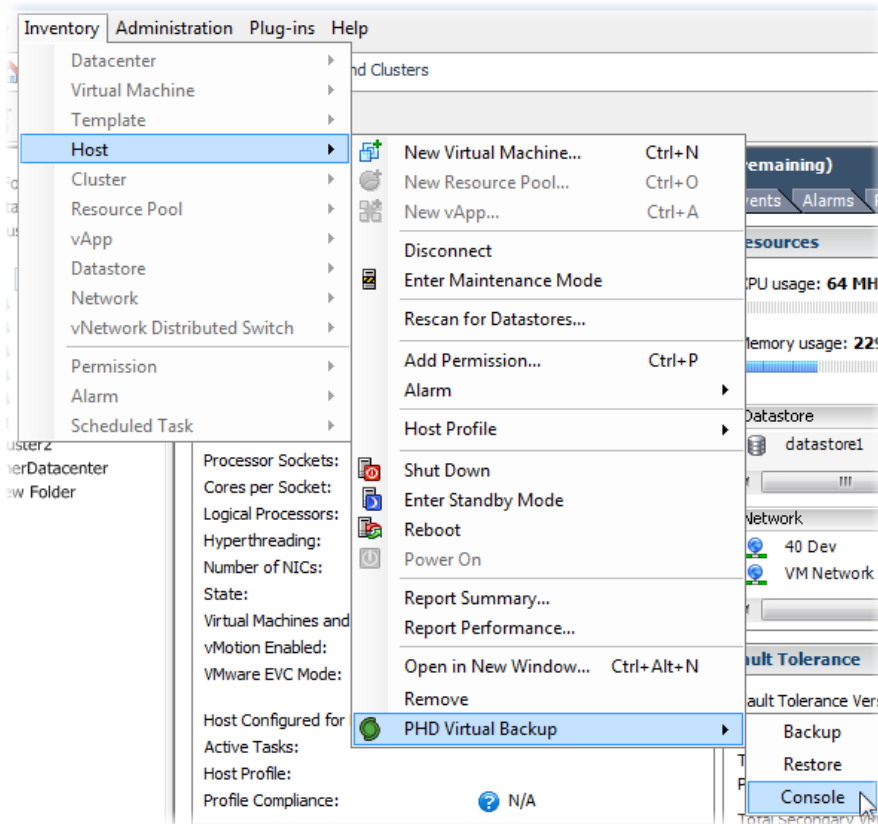
**Note:** If the PHD Virtual Backup Appliance is restarted within one hour of a scheduled Daily or scheduled Weekly job's start time, the scheduled job will be run again.

## Viewing Jobs

To view all of the jobs in progress or scheduled, use the PHD Virtual Backup Console, **Jobs** area. The PHD Console opens to the Jobs area automatically after creating a job with either the Backup Wizard or Restore Wizard.

### To start the PHD Virtual Backup Console


1. From vSphere Client, expand the Inventory menu and from the selected object menu, select **PHD Virtual Backup > Console**.



Alternatively, you can right-click an object within vSphere Client and select **PHD Virtual Backup > Console** from the context menu.

The Console opens and displays any jobs scheduled or currently in progress.

Job Name	Appliance	Type	Status	Progress	Speed	Time Remaining
Backup Demo Linux: VM	PHD/VBA	Backup Now	Running	<div style="width: 7%;"><div style="width: 7%;"></div></div> 7%	35.7 MB/s	00:03:42
Daily Backup	PHD/VBA	Backup Daily	Inactive			
Weekly VM Template B...	PHD/VBA	Backup Weekly	Inactive			

To see additional details about any job, first select the job and click  **Show Details**.

Job Name	Appliance	Type	Status	Progress	Speed	Time Remaining
Backup Windows Serv...	PHDVBA	Backup Now	Running	<div style="width: 12%;"><div style="width: 12%;"></div></div> 12%	47.8 MB/s	00:02:37
Daily Backup	PHDVBA	Backup Daily	Inactive			
Weekly VM Template B...	PHDVBA	Backup Weekly	Inactive			

Job Detail	Value
Created	7/27/2010 9:46 AM
Schedule	
Type	Now
Start	N/A
Window	N/A
Recurrence	N/A
Next Run	
Started	7/27/2010 9:46 AM
Duration	00:00:27
Message	
Dedupe Ratio	inf:1

Task Name	Type	Status
Windows Serve...	Virtual Machine	<div style="width: 12%;"><div style="width: 12%;"></div></div> 12%
0	Disk 8.6 GB	<div style="width: 12%;"><div style="width: 12%;"></div></div> 12%

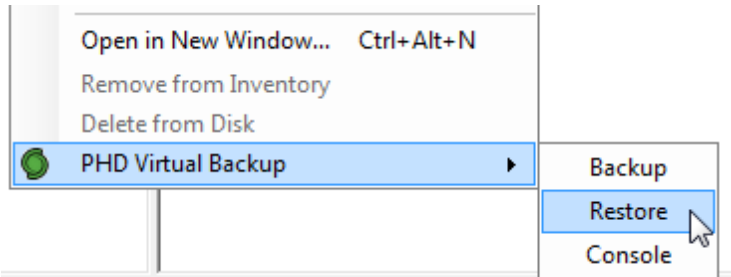
Details | Tasks

## Restoring Backups

Virtual Machine backups can be restored in the same way they were backed up, using the PHD Virtual Backup menu options within vSphere Client. By right-clicking an existing VM name, you can restore previous versions of that VM, or you can search through all existing backups to find the VM to restore.

### To restore a Virtual Machine

1. From within vSphere Client, select **PHD Virtual Backup > Restore** from any of the integrated menus.



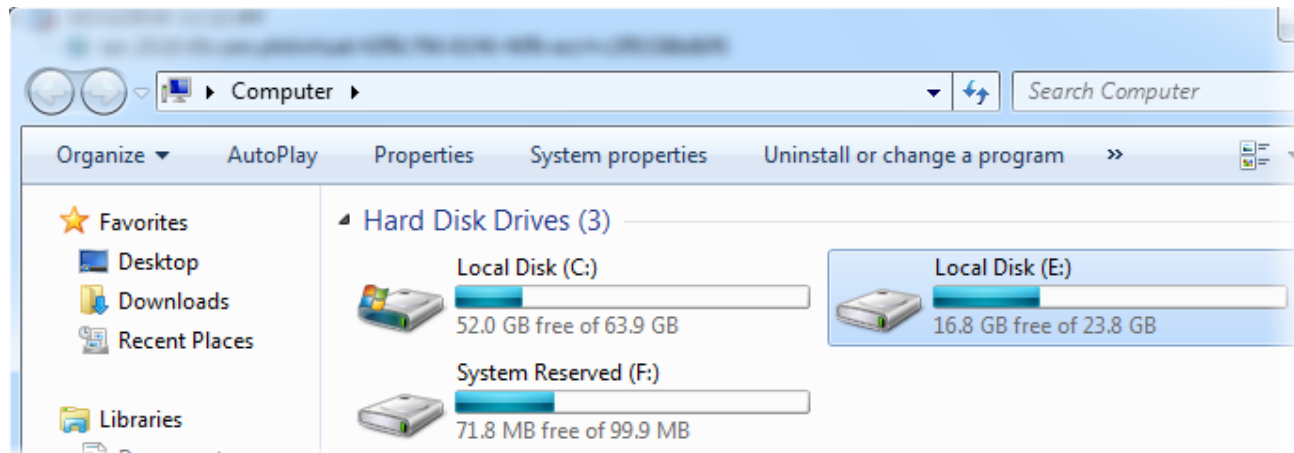
If you had right-clicked a VM Name when you opened the wizard, if a backup exists for that VM, it will be pre-selected within the Restore Wizard catalog.

The Restore Wizard guides you through the process of restoring your selected VM. For detailed information about each step of the wizard, see ["Using the Restore Wizard" on page 68](#)


When the wizard completes, the PHD Virtual Backup Console opens and displays the progress of your job.

## Restoring Files


With PHD Virtual Backup, you can restore an entire VM or you can restore individual files from a VM backup. By creating iSCSI targets from your backup files, you can mount your backed up virtual disks and browse them using Windows Explorer.



You can use the Backup Catalog to locate the backup that contains the files you want to restore then open the File Recovery wizard,



or you can start the File Recovery wizard right from the File Recovery page and browse the available backup files there.



When the wizard completes, an iSCSI target is created and available in the File Recovery area.

**File Recovery Notes**

- When running Windows, you can use the Microsoft iSCSI Software Initiator to mount the target locally or from another device. When mounted, you can browse the virtual disk using Windows Explorer to find the individual files you want.
- When running Windows, to restore files from a Linux backup, you will need to install and use a third-party Linux file system browser, for example, Ext2explore, to view the contents of the Linux disk.
- When running Linux, to mount iSCSI targets you must install an iSCSI Software Initiator for your Linux operating system, for example, on Ubuntu, you can install the Linux Open-iSCSI Initiator.

For detailed instructions on restoring individual files, see ["File Recovery"](#) on page 33.

**Note:** In order to mount iSCSI shares, the iSCSI Software Initiator must be installed on your Windows computer. The initiator is installed with Windows Vista, Windows 7, and Windows 2008 Server, by default. For earlier versions of Windows, download and installed the initiator from Microsoft's web site.

## Configuring Email Alerts

To receive email alerts from PHD Virtual Backup, use the PHD Virtual Backup Console's Configuration page.

### To enable email alerting

1. Open the PHD Virtual Backup Console.
2. From the menu on the left, click **Configuration**.
3. Click the **Email** tab.

The screenshot shows the configuration interface for email alerts. At the top, there is a dropdown menu labeled "Select the appliance to configure:" with "PHDVBA" selected. Below this is a horizontal navigation bar with tabs: "General", "Storage", "Network", "Email", "Retention", "Connector", and "Support". The "Email" tab is active. The main configuration area contains the following options:

- Do not email alerts from the appliance
- Email alerts using the following information
- Server Name:  Port:
- Security:
- Server requires credentials
- User name:
- Password:
- From Email Address:
- Alert Level:
- Recipients:

A "Save" button is located at the bottom right of the configuration area.

4. Use the options available to configure the mail server to use and any required security settings or authentication credentials, then click **Save**.

You will begin receiving email alerts from the PHD VBA you configured. If you are using multiple PHD VBAs and would like to receive alerts from each, you will need to configure each PHD VBA, separately.

For additional information about each available email configuration option, see "Email" on page 51.


## Verifying Backups and Restores with TrueRestore™

With data verification and self-healing, PHD Virtual Backup's TrueRestore technology ensures the data you backup is the data you can restore. Available verification options can be enabled during the backup and restore processes. For backups, you can additionally set the level of verification to use to None, New blocks only, or All blocks.

In addition to verify options, TrueRestore includes backup data self-healing, which means when a bad block is identified, it is flagged, and the PHD Virtual Backup Appliance will then attempt to repair that bad block, further ensuring the integrity of your data.

### To verify backups


1. At the **Options** step of the Backup Wizard, use the **Verify** menu to select the type of verify to use.



- **None** - Data is written but not checked. If a bad copy occurs or the target storage has a defective sector, valid restoration will not be possible.
- **New blocks only** - Verify only new data. Because deduplication allows for the reuse of data blocks, using this option lets you verify only the new blocks of data written to the data store. This ensures that all blocks written to the data store have been verified once after being written. Note that this option is useful only if **None** is never used. If both **None** and **New blocks only** are used, then some blocks for the VM being backed up, even with **New blocks only** selected, may never be verified. Selecting this option will impact performance.
- **All blocks** - Verify every data block needed for a restore after a backup. This includes blocks that are common to multiple backups and will result in the same blocks being verified multiple times. This option will impact backup performance.

### To verify restores

1. During the Restore Wizard, at the **Options** step, select the check box **Verify virtual machine after restore**.



This option instructs PHD Virtual Backup to verify the restored VM. What that means is, during the restore process, each block that is written is immediately read back and verified against the backup file.

## Backup Retention and Archiving

By default, PHD Virtual Backup will keep all backups for each VM. You can adjust the number of backups retained in the backup catalog using the PHD Virtual Backup Console's Retention tab in the Configuration area. After defining a retention policy, if you'd like to retain some backups indefinitely, you can use the Backup Catalog to set Archive flags for individual or groups of backup files.

### Backup Retention

Every hour, a trim job runs and removes older backups based on the defined policy. By default, no backup files are removed (Retention is set to Keep All). For details about the available settings (Keep All, Typical, and Custom), see "Retention" on page 53


Individual backups can also be deleted using the Backup Catalog. Select the backups to delete and click **Delete** in the Jobs area toolbar.




To delete all backups for a specific VM, within the Backup Catalog, select the VM name and click **Delete**.

### Archiving Backups

If you'd like to retain certain backup files indefinitely, for example if you needed to keep a master copy available on demand, you can use the Backup Catalog to set an Archive flag by selecting the backup, then clicking **Archive**.



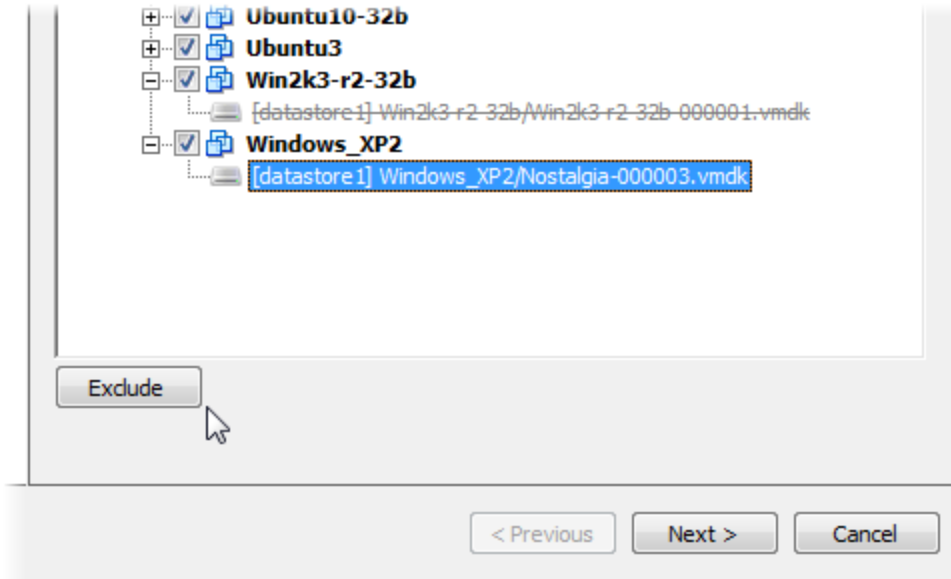
Backups flagged for archive display an archive icon  in the backup catalog, as seen in the image above. To remove the archive flag, select the backup and click **Archive** again.


You can also set the archive flag during the Backup Wizard. At the **Options** step, select **Archive Backups**. When the backup job runs, all backups created will be flagged for archive.

## Excluding VMs and Disks

Using the Backup Wizard, you can exclude VMs or individual virtual disks from a backup job. For instance, if you wanted to backup all VMs within a Folder with the exception of one, you could select the Folder within the Backup Wizard, select the VM you wanted to skip, and click **Exclude**. Then, when the backup job runs, all VMs within the folder will be backed up with the exception of the VM you chose to exclude.

When excluded, the virtual disk name is displayed with a strikethrough.




Later, if you decide you want to include the disks in the backup job, you can select the job within the Console's Job page and click  **Edit**. See "Jobs" on page 39 for details.

## Sending Backup Files to Tape

With the Backup Data Connector, you can allow access to all of your backup files via an SMB/CIFS share. Then you can use third-party tools or your own scripting to copy and move these uncompressed files to tape or to another disk location.

The Backup Data Connector is enabled using the Connector tab in the Configuration area of the Console. When enabled, you can access the share using the appliance's IP address and browse all of the available backups.



For more information about using the Backup Data Connector to allow access to your backups, see "Connector" on page 56.

## Limiting the PHD Console to a Single PHD VBA

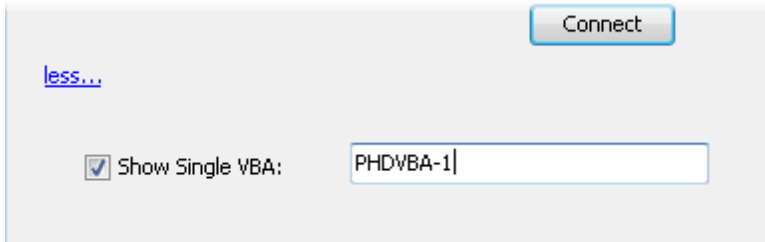
The PHD Console displays backup and configuration information for all available PHD Virtual Backup Appliances. If you have a larger environment with multiple PHD VBAs deployed, there may be situations where you want to limit the PHD Console to display only information for a single PHD VBA. For example, this may be useful if you need to make a configuration change immediately for one PHD VBA and you do not want to wait for the backup and configuration information for all other running PHD VBAs to load.

To limit the PHD Console to a single PHD VBA, use the PHD Virtual Backup login dialog which is accessed via the Windows Start menu.

### To access a single PHD VBA from the PHD Console

1. If open, close the PHD Console.
2. From the Windows Start menu, select **PHD Virtual Backup**.  
The login dialog opens.
3. Enter the credentials for your vCenter or ESX/ESXi server.
4. At the bottom of the dialog, click [more...](#) to expand the additional option.

5. Select **Show Single VBA** and enter the display name of the PHD VBA you want to access.



The screenshot shows a light gray rectangular window. In the top right corner, there is a blue button labeled "Connect". In the top left corner, there is a blue link labeled "less...". Below the link, there is a checkbox labeled "Show Single VBA:" which is checked. To the right of the checkbox is a white text input field with a blue border, containing the text "PHDVBA-1".

6. Click **Connect**.

The PHD Console opens and only information for the PHD VBA you specified is displayed.

## Increasing Backup Storage (Attached Disk)

If you are using an attached virtual disk to store your backups and you are beginning to run out of space, you can grow the storage by shutting down the PHD Virtual Appliance and adjusting the size of the storage disk, manually.

### To increase the size of your backup storage


1. Within vSphere Client, right-click the PHD Virtual Backup Appliance and select **Power > Power Off**.
2. When the appliance is powered off, click **Edit Settings**, then select the virtual disk used for storage.
3. In the Disk Provisioning area on the right, increase the provisioned size to the desired amount.
4. Click **OK** to close the window and then restart the appliance. The new size will be reflected in the PHD Virtual Backup Console's Dashboard.

## Using Multiple Network Adapters

In environments that have storage resources located on a separate network, virtual machines will require a separate network adapter to reach those resources. If a PHD VBA's backup storage is located on a network other than where the VM resides, a second network adapter must be configured.

You can add a second adapter to the PHD VBA virtual machine, then use the PHD Console to configure its settings. By default, a newly added network adapter will use DHCP to obtain an IP address. The following figure illustrates an environment with two networks, one used for VM management where the VMs and PHD VBA reside, and another for storage resources.

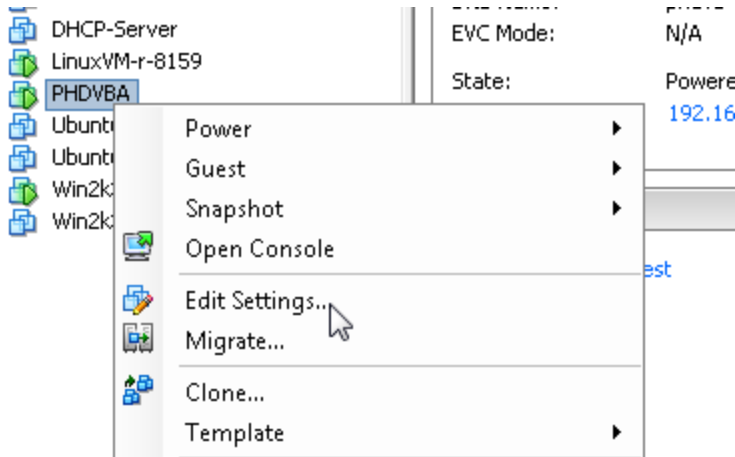
**Figure 1** - PHD VBA with a second network adapter (Storage Adapter) configured.



**Caution:** Configuring two adapters on the same network segment may result in unwanted behavior. You should only configure two adapters if you need to reach storage on a second network location.

**To add a second network adapter to a PHD VBA virtual machine**

1. Within vSphere Client, right-click the PHD VBA virtual machine and select **Edit Settings**.

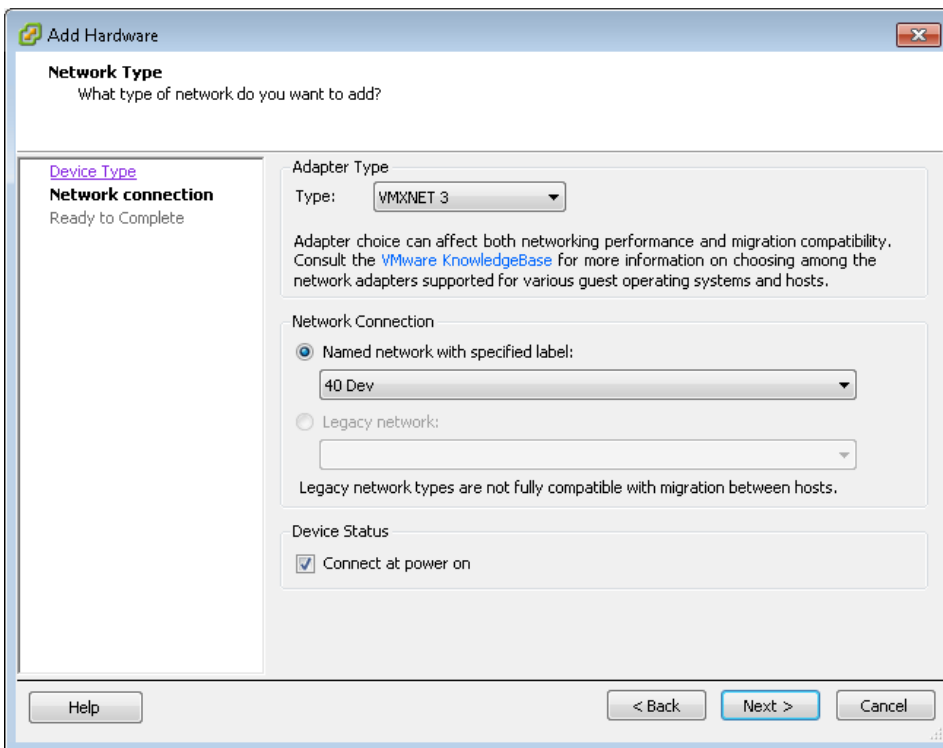



2. In the Virtual Machine Properties window that opens, click **Add...**
3. Select **Ethernet Adapter** then click **Next**.
4. Use the wizard to finish adding the new adapter, with the following attributes.

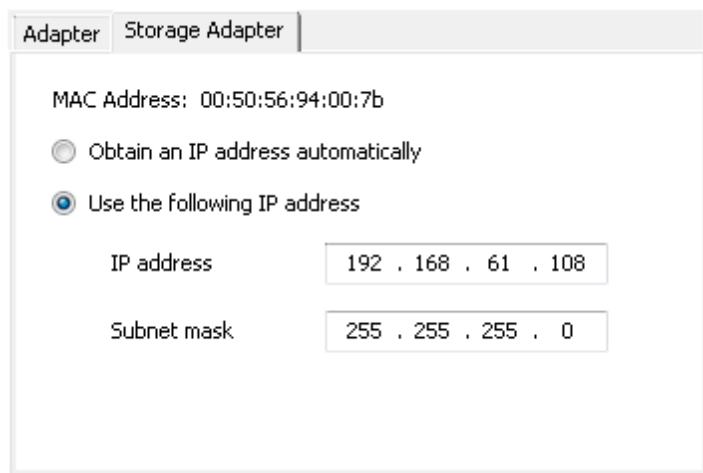
**Adapter Type:** VMXNET 3 is recommended.

**Network Connection:** Select the network that has access to your storage.

**Device Status:** Leave **Connect at power on** selected.



5. Click **Next** to review your changes, click **Finish**, then close the properties window.
6. Open the PHD Virtual Backup Console and click **Configuration** then click the **Storage** tab. Make sure the PHD VBA with the newly added adapter is selected in the appliance menu at the top of the window, then enter the information for your network storage.
7. On the **Network** tab, click the **Storage Adapter** tab - this is where you will configure the second network adapter you added, above. If the new adapter is not recognized, click . If prompted, save your changes but do not restart.



Adapter | Storage Adapter

MAC Address: 00:50:56:94:00:7b

Obtain an IP address automatically

Use the following IP address

IP address      192 . 168 . 61 . 108

Subnet mask     255 . 255 . 255 . 0

8. Enter the storage adapter settings, click **Save**, then click **OK** to restart the PHD VBA.

When the PHD VBA restarts, the network storage will be configured and you can begin running backups.

**Note:** You cannot configure a second adapter when **Storage Type** is set to **Attached Virtual Disk** - you must select a network storage option.

## Updating PHD Virtual Backup

When available, updates to PHD Virtual Backup can be downloaded from the PHD Virtual Web site or obtained from Support. Within the update packages, you will find the files necessary to update the PHD Virtual Backup components. PHD Console and Plug-in updates are installed using an updated MSI file. PHD Virtual Backup Appliances are updated by applying (.phd files).

**Note:** If you need to update your license, use the **General** tab of the Configuration page.


### What you will need:

- An existing PHD Virtual Backup installation (Console and VBA).
- The update package downloaded from the PHD Virtual Web site or obtained from PHD Virtual Support.
- Extract and save the contents of the update package (.MSI and .phd files).

### To update the PHD Virtual Backup Console and Plug-In


1. Use the Windows Control Panel, **Add Remove Programs**, to remove the current version of PHD Virtual Backup.
2. When removed, double-click the MSI from the update package and follow the steps to install the new Console and plug-in.

The new PHD Virtual Backup Console version is displayed in the **Version Information** area of the Support tab.



**To update the PHD Virtual Backup Appliance**

1. Open the PHD Virtual Backup Console, click **Configuration**, then click the **Support** tab.
2. Use the menu at the top of the window to select the PHD VBA to update.



3. In the **Version Information** area, click **Upload Appliance Patch**.


Version Information

PHD VBA Version: 5.1.2.5979

PHD Console Version: 5.1.2.5979

Patches are bundles downloaded from the PHD Virtual support website that contain updates for your appliance.

[Upload Appliance Patch](#)



4. Select the PHD VBA update file (for example, PHDVB\_1234.phd) from the update package and click **Open**.
5. After the update is applied, the PHD VBA must be restarted. Click **Yes** to restart the appliance.


The new PHD Virtual Backup Appliance version is displayed in the **Version Information** area of the Support tab.

**Note:** If the PHD VBA is an earlier version that is not supported by an updated PHD Console, you can update the PHD VBA from the Dashboard, as displayed in the image below.

System Alerts

Displays appliance messages and alerts.

Appliance	Message	Recommended Action
PHDVBA	Appliance version (v5.1) is not compatible with PHDVB Console (v5.2)	<a href="#">Upgrade this appliance</a>



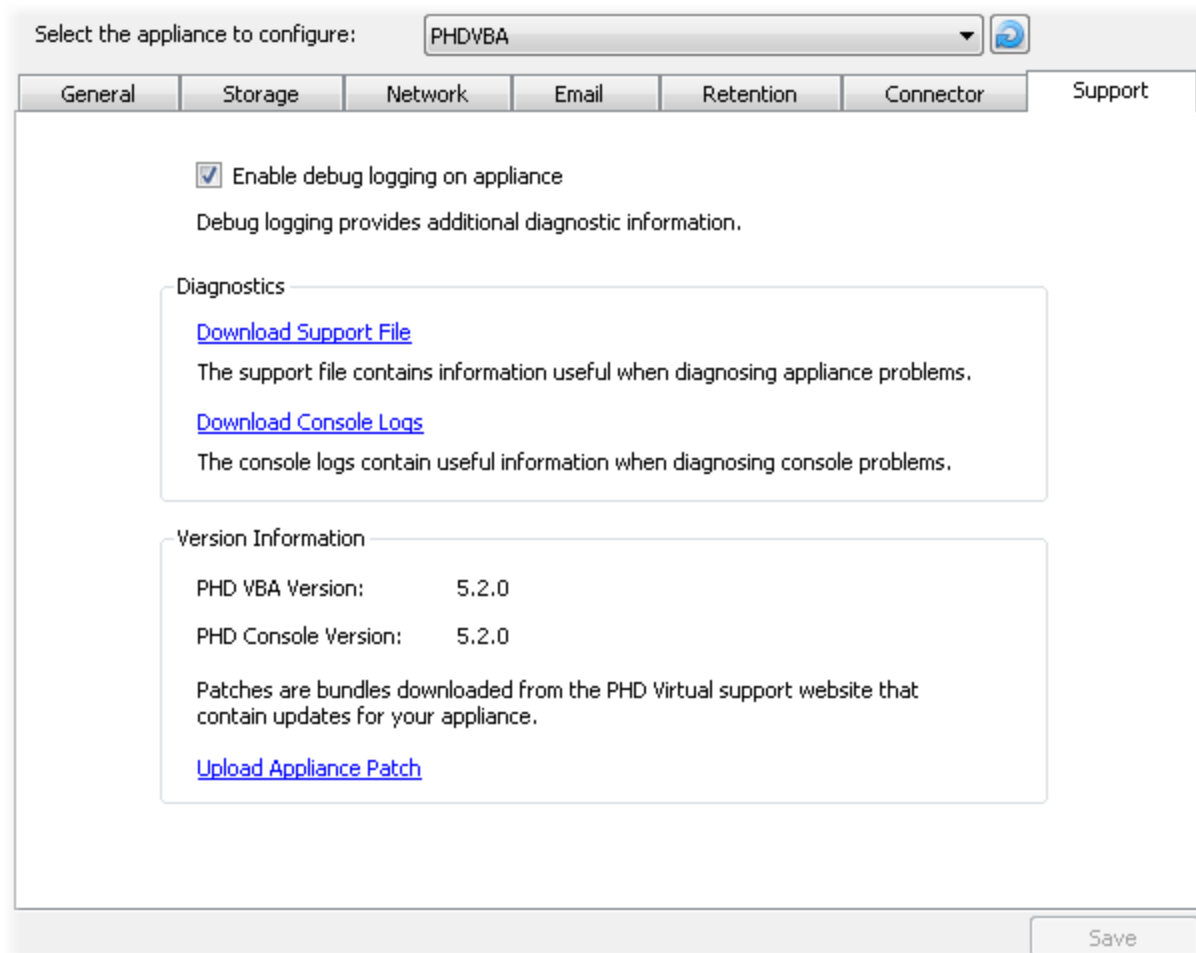
## Appendix A - Troubleshooting

The following topics contain information to help resolve issues encountered when using PHD Virtual Backup.

Downloading Support Files.....	96
Recovering Backups from an Unavailable PHD VBA.....	97
Resetting PHD VBA Network Settings.....	98
BDC Share and Local Security Policies.....	99
PHD VBA will not Power On.....	100
Backup Alerts.....	102
TCP/IP Ports.....	104

## Downloading Support Files

If you need to contact PHD Virtual Support, you may be asked to submit Support Files. These can be downloaded from the Support tab in the Configuration area of the PHD Virtual Backup Console.



For additional details about the Support tab, see "Support" on page 58.

## Recovering Backups from an Unavailable PHD VBA

If your PHD Virtual Backup Appliance becomes unavailable for some reason, you can still access your backups by deploying a new appliance and pointing to the previously used storage repository or attaching the existing virtual disk used to store backups.

### To recover backups if using an attached disk

1. Open vSphere and select the problematic PHD Virtual Backup Appliance. If running, power off the appliance (right-click and select **Power > Power Off**).
2. Right-click the appliance again and select **Edit Settings....**
3. Select the virtual hard disk used to store the backups and click **Remove**.
4. In the **Removal Options**, select **Remove from virtual machine** and click **OK**.
5. Deploy a new appliance. Use the OVF that came with your installation package. Follow the steps in the installation guide for details, making sure to select **Attached Virtual Disk** as the storage type.
6. Within vSphere Client, right-click the new appliance and select **Edit Settings....**
7. Click **Add...**
8. Select **Hard Disk** and click **Next**.
9. Select **Use and existing virtual disk** and click **Next**.
10. Browse to the location where the previous attached disk was created and select it, then click **Next**.
11. Click **Next** again, then click **Finish**
12. Power on the appliance. The appliance recreates the backup catalog automatically and you can begin backing up and restoring VMs using the new storage location.

### To recover backups if using CIFS or NFS

1. Power off the problematic appliance within vSphere Client.
2. Deploy a new appliance using the OVF that came with your installation package.
3. Follow the steps in the installation guide for details, making sure to select CIFS or NFS as the storage type.
4. Click **Save**, then restart the appliance.
5. Power on the appliance. The appliance recreates the backup catalog automatically and you can begin backing up and restoring VMs using the new storage location.


## Resetting PHD VBA Network Settings

If you are experiencing networking issues with a PHD Virtual Backup Appliance that cannot be resolved using the PHD Virtual Backup Console, or if you are deploying a new appliance and do not have DHCP enabled in your environment, you can use the VBA's virtual machine console within vSphere Client to configure the network settings.

**Note:** This method will allow you to configure the settings for the first network adapter on the PHD VBA, only. If configured, the settings for the second adapter (Storage Adapter) will be reset to use DHCP.

### To reset the PHD Virtual Backup Appliance's Network settings

1. Open vSphere Client and select the PHD Virtual Backup Appliance virtual machine.
2. Click the **Console** tab then click inside the console window.
3. Type CTRL-N to open the **Network Configuration** menu.



4. Use the Arrow keys on your keyboard to select either **DHCP** or **Static** and enter the new network settings.
5. When complete, select **OK** and hit **Enter**.
6. Restart the appliance to confirm the updated network settings.

## BDC Share and Local Security Policies

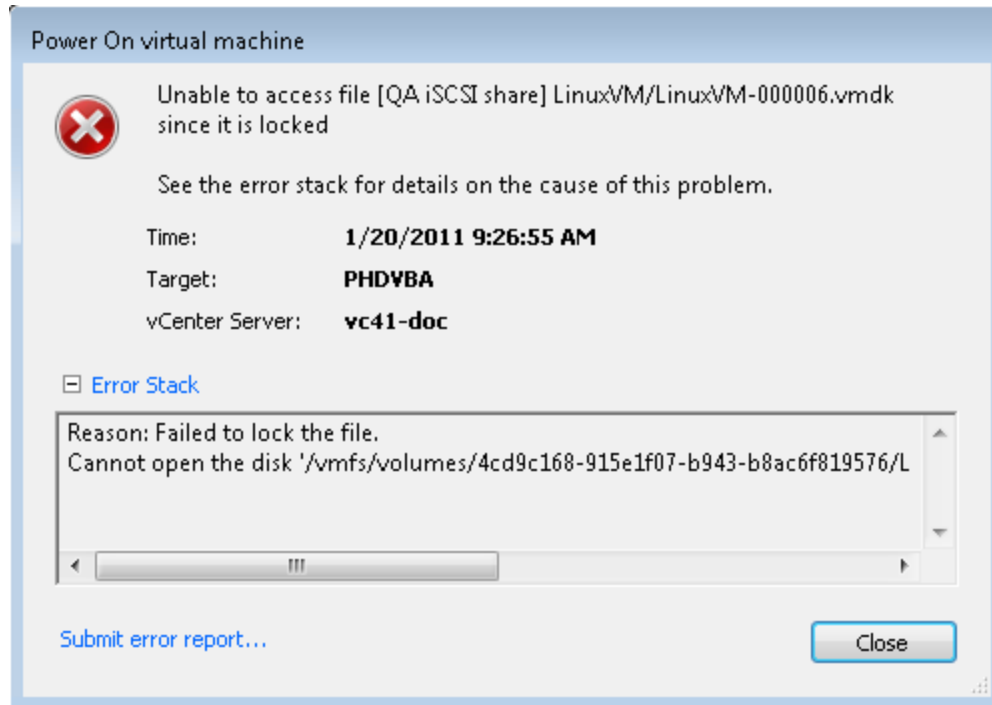
If you cannot access the Backup Data Connector (BDC) share from Windows Vista, Windows 7, or Windows 2008, you may need to adjust your local security policy, LAN Manager authentication level to "LM and NTLM - use NTLMv2 session security if negotiated."

### To adjust your LAN Manager authentication level

1. On your Windows machine, click **Start > Run** then type **secpol.msc** and hit Enter.
2. Click **Local Policies** then click **Security Options**
3. Next, navigate to and double-click **Network Security: LAN Manager authentication level**.
4. Use the menu to select **LM and NTLM - use NTLMv2 session security if negotiated**.
5. Click **OK**.
6. Now try accessing the Backup Data Connector share again.

## PHD VBA will not Power On

If you encounter an error when powering on a PHD Virtual Backup Appliance, for example, a lock error as seen in the image below, you may need to remove any attached snapshots that are no longer valid.



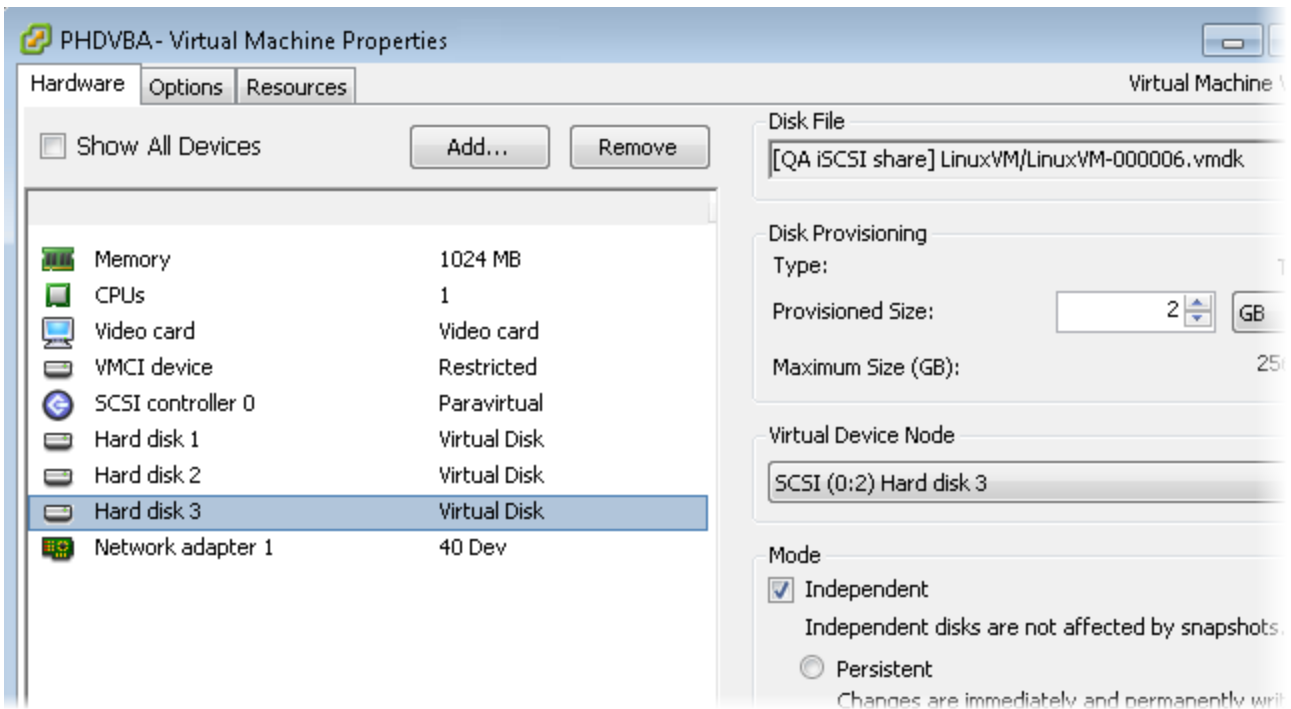
In some instances, when a PHD Virtual Backup Appliance is shutdown in the middle of a backup, it can leave behind a snapshot on the VM it was backing up. Also, that snapshot will remain attached to the powered off PHD Virtual Backup Appliance as an attached virtual disk. When the Appliance starts up, any leftover snapshots are removed by a snapshot cleanup process.

In environments using multiple PHD Virtual Backup Appliances, you may encounter a situation where a snapshot is left behind by a powered down PHD Virtual Backup Appliance but then that snapshot is removed by a second PHD Virtual Backup Appliance running the snapshot cleanup process. When you attempt to power on the original Appliance, you may encounter an error within vSphere Client, since the snapshot it had attached is no longer available. To start the Appliance, the snapshot must be removed. Follow the steps below to remove an attached snapshot from a PHD Virtual Backup Appliance.

### To remove a leftover snapshot from a VBA


1. Use vSphere Client and edit the PHD Virtual Backup Appliance's VM settings (right-click the PHD Virtual Backup Appliance VM and select **Edit Settings**).
2. Select the attached snapshot disk.

**(Warning: do not remove an attached disk that is being used as backup storage or the VBA's operating system disk).**




3. Click **Remove**. The attached snapshot disk is deleted and the PHD Virtual Backup Appliance should power on successfully.

## Backup Alerts

In the Backup Catalog, you may encounter a backup that has been flagged with an alert . This indicates that an error occurred while the backup was created and the backup should be run again (if the original backup was run with CBT on, run the next backup without CBT). Backups flagged with alerts may not restore correctly, depending on where in the backup file the error occurred. As an alternative, File Recovery can be used to access the backed up disks and restore individual files.


Backups flagged with an alert will remain in the catalog until they are deleted, either manually or based on the configured retention policy.

**Figure 1 - Backup with an alert in the Backup Catalog**




The **System Alerts** area of the Dashboard also displays an alert that indicates an error has occurred, as illustrated in the following image.

**Figure 2 - System Alerts area of the Dashboard displaying a backup alert.**

System Alerts		
Displays appliance messages and alerts.		
Appliance	Message	Recommended Action
 PHDVBA v5.2.1	The backup process encountered an error.	<a href="#">View alerts</a>

Click **View alerts** to display the System Alerts Viewer. Here you will find additional details about which VM backups were affected and the type of error encountered, as seen in the following image.

**Figure 3 - System Alerts Viewer**



To clear the alert from the Dashboard, click **Delete All Alerts** in the System Alert Viewer. Note that deleting these alerts removes only the alert messages, not the actual backup files.

**Alert Types**

The Alert Type column in the System Alerts Viewer displays the type of error encountered. The following table includes additional details about these alerts.

**Table 7 - Alert Type descriptions**

Type	Description
<b>Backup post-processing error</b>	Indicates an error occurred during the processing that takes place after a backup is completed. For example, if a problem is encountered during the linking process on the backup storage after the backup data has been written, the backup is flagged with this alert.
<b>Backup processing error</b>	Indicates an error was encountered while the backup was in progress. For example, if an individual block could not be written to the backup storage, the backup is flagged with this alert.

## TCP/IP Ports

The following table lists the ports that PHD Virtual Backup uses for communication.

**Table 8 - TCP/IP ports**

Port	Function
22	Required for SSH access to the PHD Virtual Backup Appliance (when debug logging is enabled).
139 and 445	Used by the Backup Data Connector when accessing backups over a CIFS/SMB share.
443	HTTPS access is required for communication between the PHD VBA and Console and each host.
3260	Required when accessing iSCSI targets created using File Recovery.

**Caution:** If any of the ports above are restricted (using a firewall or any other method), PHD Virtual Backup may fail to function correctly. If you require additional assistance with port configurations or firewall settings, contact your network administrator.

## Appendix B - Errors and Warnings

Review the following section for information about errors and warnings encountered when using PHD Virtual Backup. Typically, errors and warnings can be found in the log files available in the PHD Console's Jobs area.

### Could not attach...

When attempting to backup a VM that is on local storage with a PHD Virtual Backup Appliance on a different host, the appliance can not attach the VM's virtual disks to create a snapshot for backup. For example, when backing up VM1 which was deployed to local storage on Host1 with a PHD Virtual Backup Appliance that is located on Host2, you would see an error similar to:

```
VM1: Could not attach 1728279c-025a-4472-0987-0ca0f376839c to VBA
```

The message contains the name of the virtual machine (VM1) the error occurred on and the UUID of the virtual disk that could not be attached.

### Collection of metadata failed, backup aborted

When a backup job is run that includes a VM that no longer exists or was moved, PHD Virtual Backup cannot access the VM metadata to begin the backup. For example, if you scheduled a Job that backs up three VMs: VM1, VM2, VM3, then deleted VM3 before the backup job ran, you would see an error similar to:

```
VM 'Unknown': Collection of metadata failed, backup aborted
```

### Dedupe store has less than hard stop limit of 104857600 bytes free space, aborting backup job

This warning indicates the virtual disk used for storing backups has exceeded the stop level. Use the PHD Virtual Backup Console Dashboard to verify the amount of free space left. The stop level can be configured in the PHD Virtual Backup Console, Configuration page, Storage tab. Note that if you are using an attached disk to store your backups, the size of the disk can be increased by shutting down the PHD Virtual Backup Appliance and then growing the disk.

### Could not write and close backup block

When the backup datastore has run out of free space, no additional blocks of data can be written. The backup that was in progress will be aborted and the data that was partially backed up will be removed.

### Dedupe has less than 10.0% free space

This warning indicates the backup storage has exceeded the warning level configured within the PHD Virtual Backup Console, Configuration page, Storage tab.

### Another PHDVB VBA has a snapshot for this VM, backup aborted

A snapshot created by a different PHD Virtual Backup Appliance already exists for the VM being backed up. Wait for any currently running jobs to complete then try the backup again. If you still encounter this error, try restarting the other appliance that was backing up the VM. If this does not solve the problem, the snapshot can be deleted from the VM

manually using vSphere Client.

#### **The PHDVB VBA already has a snapshot for this VM, backup aborted**

The PHD Virtual Backup Appliance created a snapshot for this VM already. Additional backups cannot complete until the snapshot is removed by the appliance or manually using vSphere Client. After any currently running jobs complete, try the backup again. If you encounter the same error, restart the appliance. If you still encounter the error after restarting, you can manually delete the snapshot from the VM using vSphere Client.


#### **...does not have a UUID**

If a VM was migrated or upgraded from an older version of VMware, it may not contain a UUID. To create a UUID for the VM, use vSphere Client to edit the VM settings and add a UUID. Refer to the VMware documentation for additional details.

#### **Failed to add snapshot**

This error may be encountered if a VM has a configuration setting that prevents snapshots from occurring, for example, if Bus Sharing is enabled for the SCSI controller. The snapshot operation must be allowed to complete successfully to take backups.

#### **Failed to save changes: System unavailable due to restart**

You may encounter this error in the PHD Console if the PHD VBA takes too long to finish restarting after making a configuration change. If the timeout limit expires and you see this message, you can refresh the connection by clicking refresh  on the Configuration page.

#### **Backup is stopped**

If a backup encounters a critical error, any VM backups that were in progress will be stopped and they will be logged with this error. For example:

Windows Server: Backup is stopped

#### **CBT base compression type does not match current compression type, continuing with normal backup**

If a CBT backup is run after compression was either enabled or disabled, a regular, non-CBT backup must be run again. After the initial backup is run, subsequent backups will be CBT backups, capturing only the changed data, once again.

#### **Could not get VM metadata, retrying without EULA and annotation fields**

The EULA or Annotation field of the VM may contain non-ASCII characters that are not supported. The backup will continue, but will exclude both the EULA and Annotation fields.

#### **Excluding Independent disk...**

PHD Virtual Backup utilizes snapshots to take backups of VMs. Because VMware does not allow snapshots to be taken of disks in independent mode, these disks cannot be backed up. When an independent disk is encountered, a warning is written to the log, and the disk is excluded from the backup.

#### **...contains a Consolidate Helper Snap, backup aborted**

To avoid possible issues, VMs that have an existing Consolidate Helper snapshot will not be backed up until the Consolidate Helper snapshot is removed.

# Index

.		Backup Catalog	30, 102
.phd files	93	deleting backups	84
<b>A</b>		Backup Catalog Notes	32
Accessing the Backup Wizard	60	backup data	
Accessing the Restore Wizard	67	self-healing	83
Advanced storage options	48	Backup Data Connector	56, 86, 99
alert	102	troubleshooting	99
Alert Level	52	Backup is stopped	106
Alert Types	103	Backup Jobs	
All		creating	72
Alert Level	51	Backup Now	73
All blocks	64	Backup post-processing error	103
Another PHDVB VBA has a snapshot	105	Backup powered off virtual machines	64
Antivirus	20	Backup processing error	103
Appliance	28	Backup Retention	84
Appliance not compatible	29	Backup storage	47
Appliance options	45	Backup Storage	
appliance selection menu	44	increasing	89
Appliance updates	93	Backup Wizard	59, 72-73
Archive backups		using	61
Backup Wizard option	64	Backups	14
archive backups to disk	56	verify	83
Archiving backups	84	Backups with Alerts	102
Archiving Backups	84	BDC	56
assign static appliance network settings	50	BDC Share and Local Security Policies	99
Average Speed	41	Best Practices	20
<b>B</b>		<b>C</b>	
Backup		Cancel a backup	40
Toolbar button	40	CBT	12, 106
Backup Alerts	102	CBT base compression	106
Backup Appliances	28	CBT Notes	12

## Index

Change Disk Storage	70	Debug logging	58
Changed Block Tracking	12, 64	Dedupe Ratio	28, 41
CIFS/SMB Shares	20	Delete a backup job	40
Collection of metadata failed	105	Delete All Alerts	103
Configuration	44	delete iSCSI target	38
reload values	44	Delete trim	42, 53
Configuration page	44	Deleting backups	31
Connector	56	Deleting iSCSI targets	38
Connector tab	56	disable email alerts	52
Console and Plug-in updates	93	Disk Defragmenter	20
Consolidate Helper	106	Disk Properties	16
Could not attach	105	display Job Details	41
Could not get VM metadata	106	Do not start after	63, 75
Could not write and close backup block	105	Documentation Updates	4
CPU	15	does not have a UUID	106
CPU affinity	16	<b>E</b>	
CPU feature mask	16	Edit	
create a Backup Job	72	Toolbar button	40
create a scheduled backup job	75	edit a job	72
Creating Backup Jobs	72	Email	51
Critical		Email Alerts	82
Alert Level	51	enable alerts	51
Critical errors		Enable compression for new backups	48
email	51	Enable debug logging on appliance	58
ctkEnabled	12	Errors	105
CTRL-N	98	Alert Level	51
Custom		EULA	16
retention setting	53	Exclude	61
<b>D</b>		Excluding Independent disk	106
Daily		Excluding VMs	85
Backup schedule option	75	export backups	56
Dashboard	27	Exporting Backups	32
Backup Appliances list columns	28	<b>F</b>	
Data Streams	45	Failed to add snapshot	106
Data Written	41	Failed to save changes	106

File Recovery	33	<b>L</b>	
Floppy Drive	15	Last32Days	57
Folder		Last7Days	57
View by option	61	LatestofEach	57
Free Storage	28	license	
Frequently Asked Questions	18	update	46
<b>G</b>		upload new	46
General	45	Licensing	46
General tab	44-45	Limiting the PHD Console	87
<b>H</b>		lock error	100
Help	21	<b>M</b>	
How many PHD VBAs do I need?	22	MAC address	70
How PHD Virtual Backup Works	13	max IOPS	16
Hypervisor Credentials	44-46	Memory	15
<b>I</b>		Memory affinity	16
IDE Controller	15	Mount iSCSI target on this computer	35
Include	61	Mounting iSCSI Targets on Other Devices	37
Independent disk	106	MSI	93
Individual backups		<b>N</b>	
deleting	84	Network	49
inf	18	network adapters	49
IP address		Network Settings	
appliance	49	reset	98
obtain automatically for appliance	50	Network Shaper	16
IP Address	28	Network tab	49
iSCSI Software Initiator	81	New blocks only	64
<b>J</b>		Next Run	63, 76
Job Details	41	NFS Shares	20
Jobs	39	NTP servers	45
Jobs History	43	<b>O</b>	
<b>K</b>		Once	
Keep All		Backup schedule option	75
retention setting	53	Options	
Keyboard	15	backup wizard	63
		Orphan Weekly	42

## Index

### P

Pause a backup	40
PCI Controller	15
PHD Console	10, 87
limiting to a single PHD VBA	87
PHD VBA	10
PHD VBA will not Power On	100
PHD Virtual Backup	
benefits	11
receiving alerts	82
updating	93
PHD Virtual Backup Appliance	10
reset network settings	98
unavailable	97
updating	58, 94
PHD Virtual Backup Components	17
PHD Virtual Backup Console	10, 17, 25, 44
accessing	25
updating	93
PHD Virtual Backup Plug-in	10, 17
PHD Virtual Support	58
PHDVB	10
Plug-In	
updating	93
Pointing Device	15
Pool	
View by option	61
port	46
port 443	46
Product expiration	46
PS2 Controller	15

### Q

Quiesce the VM before backing up	64
----------------------------------	----

### R

Raw	
exporting backups as	32
Recent backups to keep	54
Recovering Backups	97
Recurrence	41
Recurring jobs	63, 75
Rekurs every	63, 75
Renaming VMs	32
Resetting VBA Network Settings	98
Restore	
Toolbar button	40
restore a Virtual Machine	79
restore file	
Linux	36
Restore Notes	15
Restore Wizard	66, 79
Restores	15
verify	83
Restoring Backups	79
Restoring Files	34, 80
Restoring Files from a Linux VM	36
Restoring Virtual Machines	31
Retention	84
Retention Settings	53
Retention tab	53
run a scheduled backup now	74
run a single backup	73
Running a Backup Now	73
<b>S</b>	
Schedule	
backup wizard	62
Scheduling Backups	75
self-healing	83
Sending Backup Files to Tape	86

Show Details	41, 78	The Restore Wizard	66
Show system jobs		Total Backup Data	28
Toolbar button	40	Trim	84
Show/Hide Details		Troubleshooting	95
Toolbar button	40	TrueRestore	83
Single PHD VBA	87	Typical	
SIO Controller	15	retention setting	53
Snap Hunt	42	<b>U</b>	
snapshot	100	Unix	36
Speed	41	Updating PHD Virtual Backup	93
Start Date	63, 75	Upload Appliance Patch	58
start the PHD Virtual Backup Console	77	Uploading Appliance Patches	58
Start Time	63, 75	Use Changed Block Tracking	64
Start/Resume		Used Storage	28
Toolbar button	40	Using PHD Virtual Backup	71
Startup	42	Using the Backup Wizard	61
Static IP Address	50	Using the Restore Wizard	68
Stop level % free	48	UUID	30
Storage	47	<b>V</b>	
Storage Adapter	92	vApp	16
Support	3, 58	vApp Options	16
Support expiration	46	VBA	8, 10, 17
Support Files	58, 96	VBA already has a snapshot	106
submitting	96	VBA Console	24
Support tab	58	Verify backup	64
System Alert descriptions	28	verify backups	83
System Alerts	27-28	Verify Restore	69
System Alerts Viewer	29, 102	verify restores	83
System Jobs	42	Verifying Backups and Restores	83
<b>T</b>		VHD	32
tape	56	Video Card	15
TCP/IP Ports	104	Video Tutorials	21
Terms	10	View by	61
The PHD Virtual Backup Appliance	22	View Log	
The post-backup process is encountering a storage erro29		Toolbar button	40

## Index

Viewing Jobs	77
Virtual Backup Appliance	8
Virtual Hard Disks	
exporting backups as	32
VM Hardware Version 7	12
VMCI device	15
VMDK	32
VMDK export	32
VMware vSphere	8
Volume Shadow Copy Services	64
<b>W</b>	
Warning level % free	48
Warnings	51, 105
Weekly	
Backup schedule option	75
What's New	9