

PHD Virtual Backup

for VMware vSphere™

version 5.1
User Guide

Software Release Date: December 2010

Document Release Date: May 04, 2011

www.phdvirtual.com



Legal Notices

Copyright © 2010-2011 PHD Virtual Technologies Inc. All rights reserved. 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 trademarks 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 or 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 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

Date	Chapter	Description
2010-12-27	1	"Changed Block Tracking" (on page 11). Added new section with additional details on VMware's CBT.
2011-02-01	1	"Frequently Asked Questions" (on page 17). Updated information about automatic snapshot removal.
2011-02-01	1	"How PHD Virtual Backup Works" (on page 12). Added notes sections for backups and restores.
2011-02-01	2	"The PHD Virtual Backup Appliance" (on page 21). Updated information about automatic snapshot removal. (5.1.2)
2011-02-01	3	" The PHD Virtual Backup Console" (on page 24). Added information for using non-standard ports when accessing the console as a stand-alone application. (5.1.2)
2010-12-27	3	"Jobs" (on page 37). Added new sub-section with additional details about PHD Virtual Backup job speeds and deduplication.
2011-02-01	3	"Jobs" (on page 37). Added information about the Snap Hunt system job. (5.1.2)
2011-02-01	3	"Job History" (on page 40). Job History tab was enhanced to include icons in the result column. (5.1.2)
2011-03-25	3	"Retention" (on page 50). Updated information about when the Delete trim job runs. (5.1.4)
2011-03-25	6	"Limiting the PHD Console to a Single PHD VBA" (on page 81). Added section that describes how to limit the PHD Console to show one PHD VBA, only. (5.1.4)

Contents

- Chapter 1 - Welcome** 7
 - What's New 8
 - About This Guide 9
 - Benefits of PHD Virtual Backup 10
 - Changed Block Tracking 11
 - How PHD Virtual Backup Works 12
 - Backups 13
 - Restores 14
 - PHD Virtual Backup Components 16
 - Frequently Asked Questions 17
 - Best Practices 19
 - Getting Help 20
- Chapter 2 - The PHD Virtual Backup Appliance** 21
 - The PHD VBA Console 23
- Chapter 3 - The PHD Virtual Backup Console** 24
 - Dashboard 26
 - Backup Catalog 29
 - File Recovery 32
 - Restoring Files 32
 - Restoring Files from a Linux VM on Windows 34
 - Mounting iSCSI Targets on Other Devices 36
 - Deleting iSCSI targets 36
 - Jobs 37
 - Job Details 38
 - Job Speeds, Deduplication, and Data Written 39
 - Job Types 39
 - Job History 40
 - Configuration 41
 - General 42
 - Storage 44
 - Network 46

Using DHCP.....	47
Using Static IP Addresses.....	47
Email.....	48
Retention.....	50
Connector.....	53
Support.....	55
Chapter 4 - The Backup Wizard.....	56
Chapter 5 - The Restore Wizard.....	62
Chapter 6 - Using PHD Virtual Backup.....	65
Creating Backup Jobs.....	66
Running a Backup Now.....	67
Scheduling Backups.....	69
Viewing Jobs.....	71
Restoring Backups.....	73
Restoring Files.....	74
Configuring Email Alerts.....	76
Verifying Backups and Restores with TrueRestore™.....	77
Backup Retention and Archiving.....	78
Excluding VMs and Disks.....	79
Sending Backup Files to Tape.....	80
Limiting the PHD Console to a Single PHD VBA.....	81
Increasing Backup Storage (Attached Disk).....	83
Updating PHD Virtual Backup.....	84
Appendix A - Troubleshooting.....	86
Support Files.....	87
What To Do If a PHD VBA Crashes.....	88
Resetting PHD VBA Network Settings.....	89
Problems Accessing the BDC Share.....	90
Cannot Power on a VBA.....	91
Appendix B - Errors and Warnings.....	93
Index.....	95

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 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.....	8
About This Guide.....	9
Benefits of PHD Virtual Backup.....	10
Changed Block Tracking.....	11
How PHD Virtual Backup Works.....	12
Frequently Asked Questions.....	17
Best Practices.....	19
Getting Help.....	20

What's New

- PHD Virtual Backup now includes support for VMware vSphere environments (ESX and ESXi):
 - Backup and restore VMs using integrated menus within vSphere Client.
 - Create backup jobs for vCenter Pools, Clusters, Datacenters, Folders, and hosts.
- File Level Restore - mount a backup as an iSCSI share and restore individual files. See ["Restoring Files" \(on page 74\)](#) for details.
- Flexible backup storage options - send your backups to attached local storage or to external locations, including NFS or SMB/CIFS shares. Refer to the appliance deployment instructions for details and ["Storage" \(on page 44\)](#).
- Use the new Archive Backup feature to preserve backups outside of normal retention policy settings. See ["Backup Catalog" \(on page 29\)](#) for details.
- Enhanced backup retention settings let you select from pre-configured policies or customize your own. For details, see ["Retention" \(on page 50\)](#).
- Define custom levels to warn you when your backup storage runs low. See ["Storage" \(on page 44\)](#)
- Backup Data Connector - save your backups to tape or archive them to disk by accessing them directly from the PHD Virtual Backup Appliance via an SMB share. For step-by-step instructions, see ["Connector" \(on page 53\)](#)
- Additional improvements, including:
 - New, faster compression for backups.
 - An improved snapshot model for more efficient use of space during backups.
 - Improved virtual machine disk restore speeds.

About This Guide

This guide is designed to introduce you to PHD Virtual Backup for 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 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
PHD Virtual Backup Plug-in	The integrated component of PHD Virtual Backup found within vSphere Client and installed via the PHD Virtual Backup MSI.
PHD Virtual Backup Console	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.
VBA™	Virtual Backup Appliance. A small virtual machine used to backup and restore other VMs. The PHD Virtual Backup Appliance is a VBA.
PHD Virtual Backup Appliance	The VBA that is deployed and used to perform backups and restores of virtual machines.
PHDVB	PHD Virtual Backup
PHD VBA	The PHD Virtual Backup Appliance.
PHD Console	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 SMB/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 "[Changed Block Tracking](#)" (on page 11).

Note: VMware vSphere Hypervisor (the free version of ESXi) is not supported.

Changed Block Tracking

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 "[The Backup Wizard](#)" (on page 56).
- 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.

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" \(on page 13\)](#)
- ["Restores" \(on page 14\)](#)
- ["PHD Virtual Backup Components" \(on page 16\)](#)

Backups

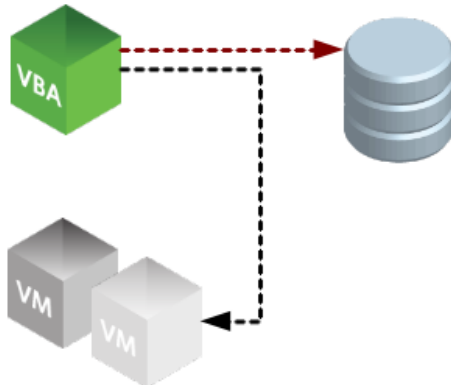
When a backup is run, the PHD Virtual Backup Appliance interacts with the vStorage API to create a snapshot of the virtual machine targeted for backup




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" \(on page 74\)](#) 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)
 - Disk Properties, including max IOPS
 - Additional default devices
 - Storage vMotion parameters dMotion.enabled and *.DMotionParent

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" \(on page 21\)](#)
- **PHD Virtual Backup Console** - Installed with the Plugin, 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 launched from within the vSphere Client or from the Windows Start Menu. For more information, see ["The PHD Virtual Backup Console" \(on page 24\)](#).
- **Backup Wizard** - The wizard which guides you through the steps of creating and editing backup jobs. See ["The Backup Wizard" \(on page 56\)](#) 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 ["The Restore Wizard" \(on page 62\)](#) 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" \(on page 50\)](#).

What happens if my appliance is rebooted 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 automatically start again.

Can I edit a job while it is running?

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

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 53).

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 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

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

The PHD VBA requires direct write access to the NFS export. During backup, the 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 anti-virus software scans, including the network shares and directories used for backup targets (CIFS and NFS).


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 also 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) that demonstrate how to install and use PHD Virtual Backup.

Chapter 2 - The PHD Virtual Backup Appliance

The PHD Virtual Backup Appliance (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, SMB/CIFS share, or NFS share).

Figure 1 - The PHD Virtual Backup Appliance in vSphere Client



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

Note: If the appliance 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 appliance 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 24) 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 23).

How many VBAs do I need?

- The number of appliance 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 an appliance 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 appliances within each.

If you need to deploy additional appliances, 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 appliance 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 the appliance is started again, it may become unresponsive and a reboot will be required to reestablish the appliance. Avoid suspending the appliance 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.



Note: Powered off PHD Virtual Backup Appliances are not available within the Console. 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 launch the PHD Virtual Backup Console" (on page 71)
- The Console can also be launched 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 81)

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

- "Dashboard" (on page 26)
- "Backup Catalog" (on page 29)
- "File Recovery" (on page 32)
- "Jobs" (on page 37)
- "Configuration" (on page 41)

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	192.168.40.251	6.7 GB	13.1 GB	3.2 GB	4:1
PHDVBA_cifs	192.168.40.252	24.6 GB	266 GB	25.4 GB	N/A

Storage

- Used space: 25.4 GB (50.8%)
- Free space: 24.6 GB (49.2%)

Deduplication (Post-compression)

- Used Storage
- Duplicate

System Alerts

Displays appliance messages and alerts.

Appliance	Message	Recommended Action
PHDVBA	Appliance has debug enabled. This will degrade performance.	
PHDVBA_cifs	Appliance has debug enabled. This will degrade performance.	

Backup Appliances

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 descriptions

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

Note on 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 appliance.

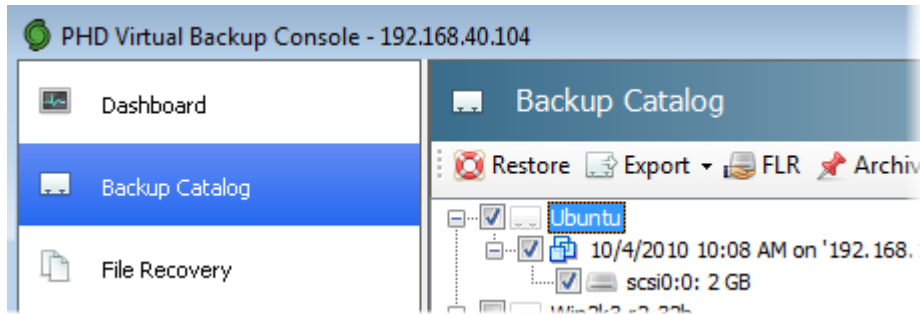
The following table provides additional information about some of the system alerts you may encounter.

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.
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 updates from PHD Virtual. Update your license file using the General tab.
Appliance has debug enabled. This will degrade performance.	On the Support tab, Debug mode can be enabled to provide expanded logs when working with PHD Virtual Support. Enabling Debug will impact backup and restore performance and should only be enabled if instructed to do so by PHD Virtual Support.

Backup Catalog









The Backup Catalog displays all available backups in an expandable tree-view. From here, you can 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, and if the VM was powered on during the backup, they additionally display the host on which the VM was running during the backup. For backups taken while a VM was powered off, only the date and time is displayed.

If your backup catalog contains VMs with identical names, the UUID of the VM will be appended to one of the VM names in the backup catalog.


Table 4 - Backup Catalog Toolbar Buttons

Button icon	Description
 Restore	Launches the Restore Wizard. For details, see "The Restore Wizard" (on page 62) .
 Export	Opens the Export dialog from which you can export the selected disks as VMDK, VHD, or Raw formatted files.
 FLR	Launches the File Recovery wizard. For details, see "File Recovery" (on page 32) .
 Archive	Lets you set selected backup files as archived, which means they cannot be deleted by the trim process or manual deletes. For details, see "Backup Retention and Archiving" (on page 78) .
 Delete	Deletes the selected backup files. Note that backups marked as archived will not be deleted.
 Refresh	Refreshes the catalog.
 View by	Changes the catalog view to display backups by Virtual Machine, Date, or Appliance.
 Expand All	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 catalog tree view. 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 ["The Restore Wizard" \(on page 62\)](#) for details.


Deleting backups

1. Find the VM backup you want to delete using the catalog tree view. 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" \(on page 37\)](#) 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 delete.
3. 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 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 and click **Save**.

Note: 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, within the Backup Catalog you would find an entry only for NewVM1. Under the NewVM1 backup tree, you would then find each backup, including the backup that was taken when the VM was named TestVM1. This backup would be noted under the NewVM1 tree 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 to restore individual files. By creating an iSCSI target from a backup, 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 a Linux machine you must install an iSCSI Software Initiator for your Linux 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
 Create	Launch the File Recovery wizard to guide 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.
 Mount	Mount an existing iSCSI target locally.
 Copy	Copy an existing iSCSI target's credentials to the Windows clipboard.
 Delete	Delete an iSCSI target. Note that the target must not be connected in order to be removed - you can disconnect targets using the iSCSI initiator.
 Refresh	Refresh the list of iSCSI targets.
 Collapse / Expand	Collapse or expand the list of iSCSI targets.
 Open iSCSI Initiator	Launch the iSCSI Initiator.
 Open Computer Management	Open 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 32).
- "Restoring Files from a Linux VM on Windows " (on page 34).
- "Mounting iSCSI Targets on Other Devices" (on page 36).
- "Deleting iSCSI targets" (on page 36).

Restoring Files

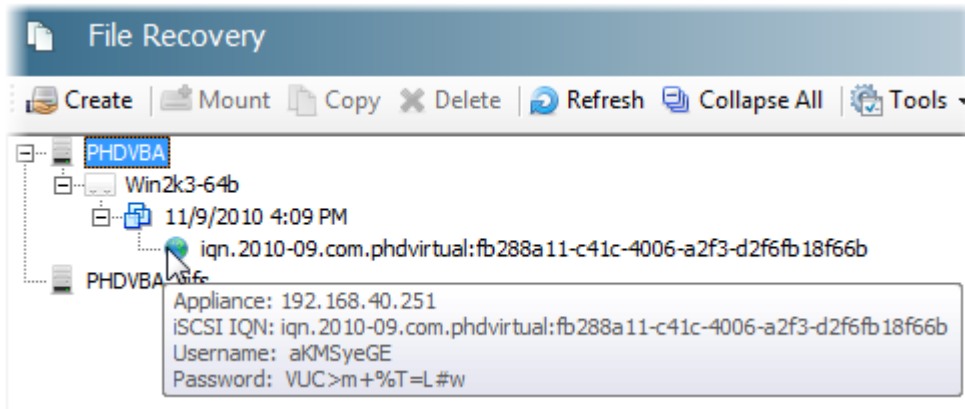
Restoring files and folders from your backups is as simple as creating and mounting an iSCSI target. Follow the steps below to create, mount, and browse files on an iSCSI target created from an existing backup.

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 automatically mount the target locally after the wizard completes (to mount iSCSI targets the Microsoft iSCSI Software Initiator 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 file.



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

When mounted, the target 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.

- 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 "File Recovery" (on page 32) for details.

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

If you need to mount an iSCSI target created from a Linux VM, see "Restoring Files from a Linux VM on Windows" (on page 34).

To mount an iSCSI target on another device, see "Mounting iSCSI Targets on Other Devices" (on page 36).

To delete an iSCSI target, see "Deleting iSCSI targets" (on page 36).

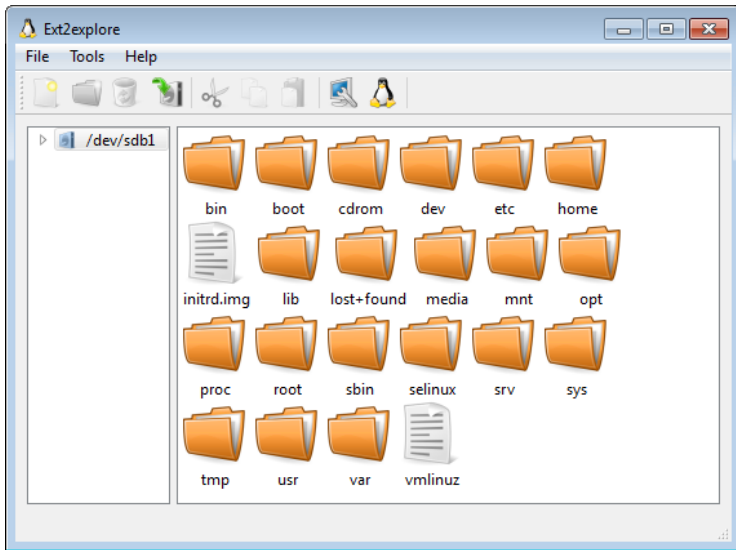
Restoring Files from a Linux VM on Windows

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

To restore files from a Linux VM backup on a Windows machine

In order to restore files from an iSCSI target created from a Linux backup you will need to use a third-party tool, for example Ext2explore, 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 Linux file system explorer tool, for example, Ext2explore, to view the contents of the mounted Linux disk.



To mount an iSCSI target on a Windows machine, see ["Restoring Files" \(on page 32\)](#).

To mount an iSCSI target on another device, see ["Mounting iSCSI Targets on Other Devices" \(on page 36\)](#).

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 machine you must install an iSCSI Software Initiator for your Linux operating system, for example, on an Ubuntu machine, you can install the Linux Open-iSCSI Initiator.

Note: that the following steps use Windows 7; your specific steps may vary based on your operating system.

1. Open the Windows iSCSI Software Initiator (Click **Start** > **Run** and type: **iSCSI Initiator**, then select it from the list of programs)
2. If the service is not running, click **Yes** to start it.
3. In the Targets tab, enter the IP address associated with the iSCSI target you created. This will be the IP address of the PHDVB appliance where the target was created.
4. Select the target from the list and click **Connect**. The Connect to Target dialog opens.
5. Click **Advanced** and select **Enable CHAP log on**.
6. Enter the username and password of the iSCSI target and click **OK**.
7. Click **OK** again. The target is mounted and available from within Windows Explorer as a new drive.

If you need to mount an iSCSI target created from a Linux VM, see ["Restoring Files from a Linux VM on Windows " \(on page 34\)](#).

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 Virtual Backup Console to the File Recovery page and select the iSCSI target.
3. Click **✗ Delete**.

Jobs

The Jobs area displays the status of running and scheduled jobs as well as maintaining a history of all jobs run and the result of each. 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 **Inactive** status.

Jobs

Backup Restore Edit Start Pause Cancel Delete Hide Details View Log Options

Current History

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











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

Tasks

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

The Jobs toolbar can be used to launch the Backup Wizard and the Restore Wizard or to control job status. The Jobs toolbar buttons are described in the following table.


Table 6 - Jobs Toolbar Buttons

Button Icon	Description
 Backup	Launches the Backup Wizard which guides you through the process of creating backup jobs. See " The Backup Wizard " (on page 56) for details.
 Restore	Launches the Restore Wizard which guides you through the process of restoring stored backups. See " The Restore Wizard " (on page 62) for details.
 Edit	Edit the selected job. The Backup Wizard launches allowing you to edit the Job settings.
 Start	Start an Inactive job or resume a paused job.
 Pause	Pause a job that is currently running. Note that average speed is not adjusted for paused jobs.
 Cancel	Cancels a job that is currently running. A cleanup process removes any unneeded snapshots or partial backup files.
 Delete	Deletes a current job.
 Show Details	Opens the Details pane which displays additional information about the selected job.
 View Log	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.
 Options	Select Show system jobs 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.

Job Details Parameter	Description
Recurrence	When the job is set to recur. For details on recurrence, see "Scheduling Backups" (on page 69) .
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 on 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 ($20,000 \text{ MB} / 240 \text{ seconds} = 83.3333 \text{ MB/s}$).

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 ["Changed Block Tracking"](#) (on page 11).

Job Types


PHD Virtual Backup Job types include:

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

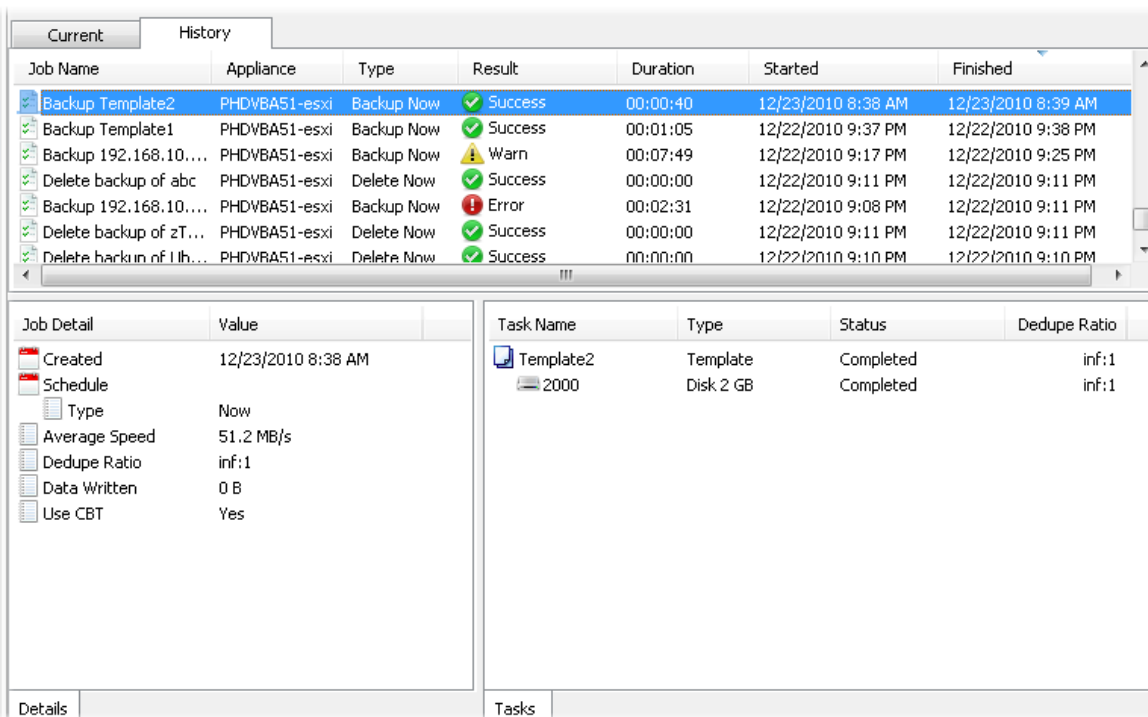
System Jobs include:

- **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 Saturday at 9 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 50) 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 contains a table with the following columns: Job Name, Appliance, Type, Result, Duration, Started, and Finished. The table lists several backup and delete jobs with their respective statuses (Success, Warn, Error) and completion times.

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 1h...	PHDVBA51-esxi	Delete Now	Success	00:00:00	12/22/2010 9:10 PM	12/22/2010 9:10 PM

Below the table, there are two detailed views for the selected job 'Backup Template2':

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 appliance must be configured separately; the drop-down menu at the top of the Configuration page indicates which appliance's settings are displayed.

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 select appliance drop-down 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 42)
- "Storage" (on page 44)
- "Network" (on page 46)
- "Email" (on page 48)
- "Retention" (on page 50)
- "Connector" (on page 53)
- "Support" (on page 55)

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 Appliance. 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 "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")
 - 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 next to the expiration dates.

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 each 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 Host Name text box. Unless you are using a non-standard port to communicate with your server, leave the default port set to 443.

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

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 drop-down menu at the top of the Configuration page 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 and also apply product updates.
- 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/SMB share, or to an NFS share.

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

The screenshot shows the configuration interface for the PHDVBA appliance. At the top, a dropdown menu is set to 'PHDVBA'. Below this are several tabs: 'General', 'Storage' (which is selected), 'Network', 'Email', 'Retention', 'Connector', and 'Support'. The 'Storage' tab contains two main sections:

- Backup storage:** A dropdown menu for 'Storage Type' is set to 'Attached Virtual Disk'. Below it, a green checkmark indicates 'Using attached disk 10 GB'.
- Advanced options:**
 - A checked checkbox for 'Enable compression for new backups'.
 - A 'Warning level % free' spinner set to '10.00', with a note 'Warns at 1 GB of free storage'.
 - A 'Stop level % free' spinner set to '3.00', with a note 'Stops at 307.2 MB of free storage'.
 - A 'Reset to Defaults' button.

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

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 on 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.

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** drop-down menu, select the type of storage to use. If you select to use an NFS or CIFS share you will be prompted to enter the share location and credentials the appliance should use.
4. Click **Save** and restart the appliance.

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 the speed at which you backups are taking place, 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 receive a warning that your backup storage is running low on available free space.
- **Stop level % free** - use this option to cause PHD Virtual Backup to stop running backups when 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

General Storage **Network** Email Retention Connector Support

Adapter

MAC Address: 00:50:56:9e:00:0e

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

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.

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** drop-down box at the top of the page.
3. Click the **Network** tab.
4. Select **Obtain an IP address automatically**.

When you've selected to automatically obtain an IP (using DHCP), you have the option to obtain DNS information automatically by selecting **Obtain DNS address automatically**, or you can specify your DNS settings.

5. Click **Save**.

To configure a PHD VBA to use a static IP address, see ["Using Static IP Addresses" \(on page 47\)](#)

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** drop-down box.
3. Click the **Network** tab.
4. and select **Use the following IP address**.
5. Enter your IP address, Subnet mask, and Gateway.
6. When manually assigning networking information, you must also define your DNS settings. Enter a preferred and alternate DNS address.
7. Click **Save**.

To configure a PHD VBA to use DHCP to automatically obtain an IP address, see ["Using DHCP" \(on page 47\)](#).

Email

Use the Email tab if you want to receive email alerts from PHD Virtual Backup. You can enter your email server options.

You can select to send email alerts for Critical errors, Errors, or All, which includes backup and restore completions, system alerts, and errors. Warnings are not sent as 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** drop-down box.
3. Click the **Email** tab and 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 drop-down list.
 - **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 reports will come from).
8. Select the **Alert Level**
 - **All** - include all alerts, including backup and restore job results and all system level alerts, in the emailed alert report. Warnings are not sent as email alerts though they are included in the backup and restore reports.
 - **Errors** - include all errors (Error and Critical Error) in the emailed alert report.
 - **Critical**- include critical errors only in the email alert report.
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. 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** drop-down box.
3. Click the **Email** tab and select **Do not email alerts from the appliance**
4. 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 use pre-defined settings selected from the drop down menu, 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. Click the **Retention** tab and use the **Retention setting** dropdown menu to select your retention policy.
3. Click **Save**.

To keep only a certain number of backups per VM

1. Open the PHD Virtual Backup Console and click **Configuration**.
2. Click the **Retention** tab and use the **Retention setting** dropdown menu to select **Custom**.
3. 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.
4. Set the **Days, Weeks, Months, and Years** values to 0.
5. Click **Save**. Now, only the five last 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	2
Total Backups Retained			20

* 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 'Connector' tab in 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 'Connector' tab is active, displaying the 'Backup Data Connector' configuration section. This section includes a checked checkbox for 'Enable share at \\192.168.40.52\backups'. Below this are three input fields: 'User name' with the value 'phd', 'Set Password' with a masked password of 12 dots, and 'Confirm Password' with a masked password of 12 dots. A 'Save' button is located at the bottom right of the configuration area.

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. Click the **Connector** tab.
3. Select **Enable Share at...** This will display your appliance IP address and the share name, for example, \\192.168.1.100\backups.
4. Set a password to access the share and confirm the password entered. The default username *phd* cannot be changed.

5. Click **Save** and restart the appliance (any backups or restores in progress will be canceled).
1. Open the PHD Virtual Backup Console and click **Configuration**.
2. Click the **Connector** tab.
3. Select **Enable Share at...** This will display your appliance IP address and the share name, for example, \\192.168.1.100\backups.
4. Set a password to access the share and confirm the password entered. The default username *phd* cannot be changed.
5. Click **Save** and restart the appliance (any backups or restores in progress will be canceled).

When created, 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.

In addition to accessing the files through the share, you can manually export individual backups using the Export backup feature. See "Backup Catalog" (on page 29) 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 "Problems Accessing the BDC Share" (on page 90).

Support

Use the Support tab to enable debugging mode, download support files, apply updates to the PHD Virtual Backup Appliances, and find the installed version information.

Select the appliance to configure: PHDVBA

General Storage Network Email Retention Connector Support

Enable debug logging on appliance

Debug logging provides additional diagnostic information. Enable this option only if instructed to by support as this will degrade appliance performance.

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: 5.1.2.6156 (for VMware vSphere)

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](#)

Save

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.

Before downloading and sending support files, you may also be asked to enable debugging mode. This is also accomplished using the Support tab by selecting **Enable debug logging on appliance**.

Caution: Enabling debug mode will negatively impact the performance of the PHD Virtual Backup Appliance. Only enable this option if instructed to do so by support.

Uploading Appliance Patches

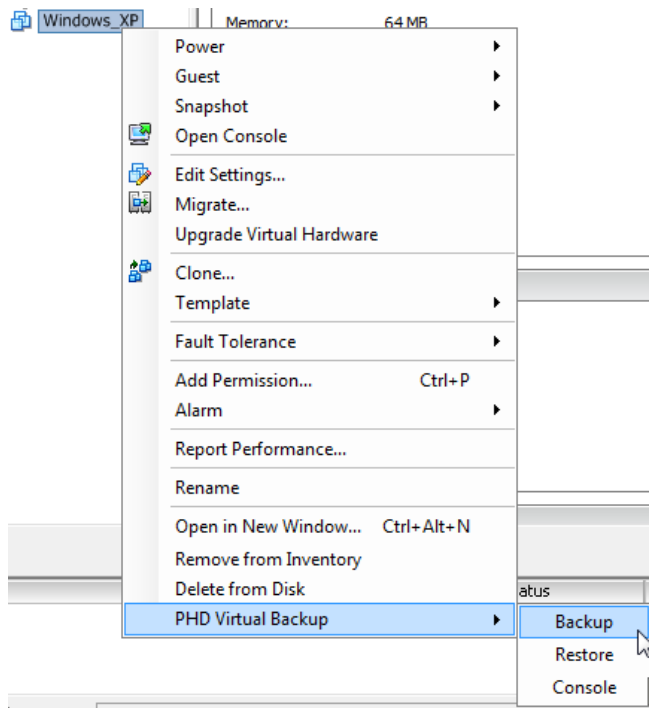
Periodically, update patches for the PHD Virtual Appliance will be available for download from the PHD Virtual Web site. When downloaded to your local computer, they can be uploaded through the PHD Virtual Backup console using the **Upload Appliance Patch** link. Clicking this link will allow you to select the downloaded appliance patch file. For additional information, see "Updating PHD Virtual Backup" (on page 84).

Chapter 4 - The Backup Wizard

The Backup Wizard lets you create backup jobs to protect the virtual machines in your environment.

To launch the Backup Wizard

- There are multiple ways to launch the wizard using the integrated menus. All options allow you to select Backup from the integrated **PHD Virtual Backup** menus.

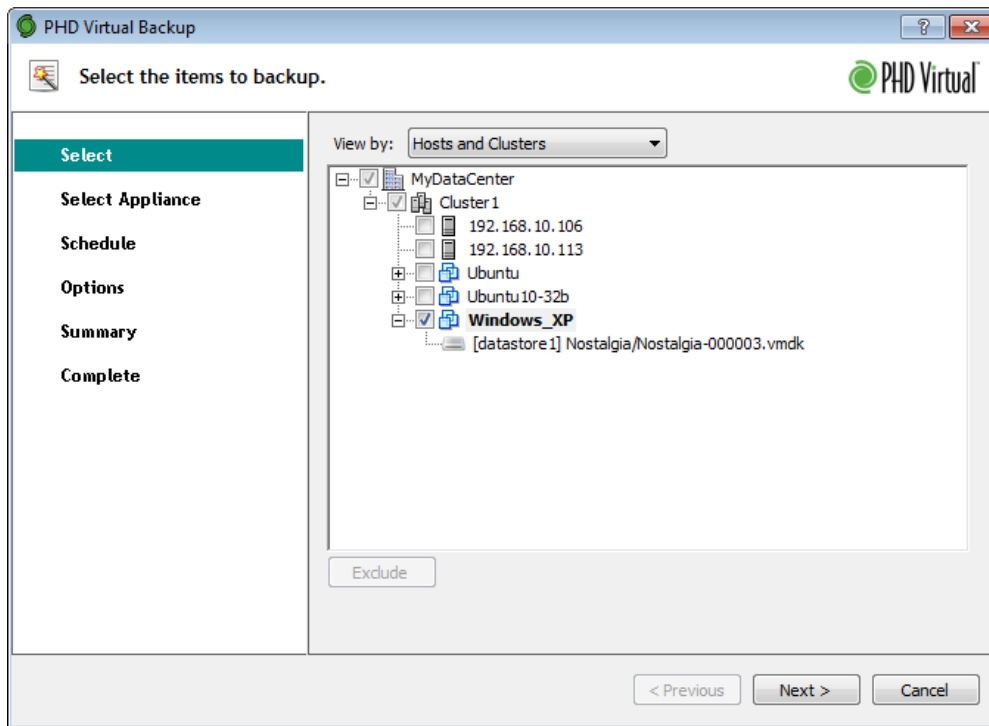


Using the Backup Wizard

1. When the wizard opens, you are presented with the **Select** step. Here you can use the **View by:** drop-down options to change how the VMs 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.



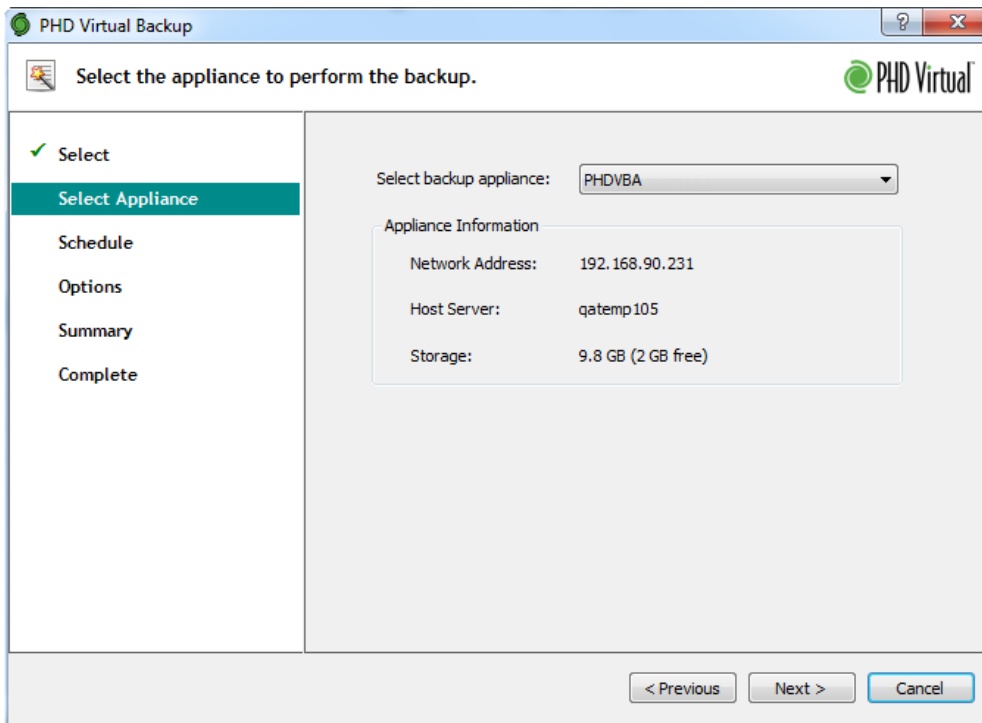
If you select the top container in any view (for example, Cluster1 in the image above) all VMs in the container will be included in the backup job. Also, after the job is created, any VMs added to or removed from the selected container will also be included or excluded from the backup job, respectively.

- **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 the Exclude button. VMs can be included again by editing the job.

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, use the drop-down box to select the appliance 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.

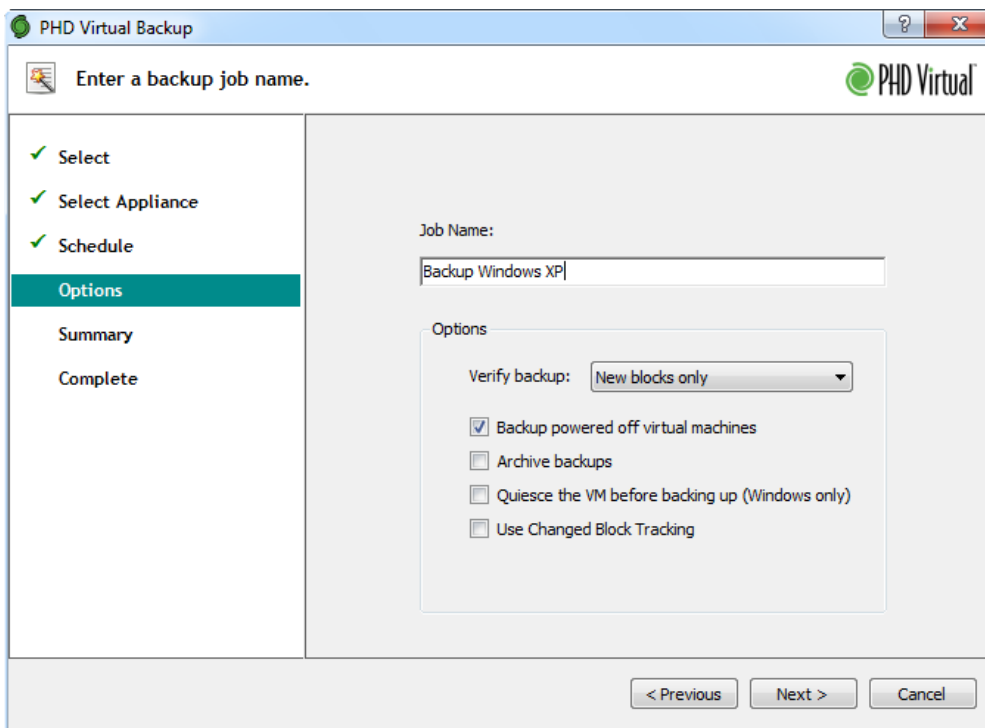
Note: If you will be backing up a VM located on local storage, you must select an appliance that is located on the same host as the VM or else the backup will fail. Virtual disks for any VMs that are unreachable by an appliance (on different local or shared storage, for example) will be displayed after 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. 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 69).



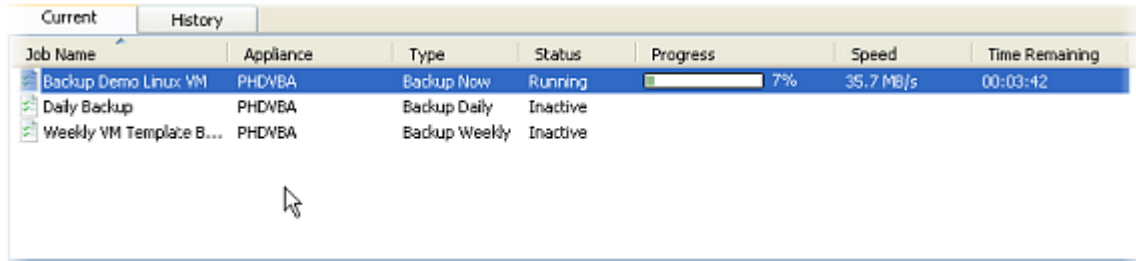
- **Start Date**- The date the scheduled job will begin.
- **Start Time**- The time the job should start.
- **Do not start after**- The time after which the job should not start. In a situation where many backup jobs or very large jobs are running and this time passes before the job can begin, it will 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).

6. Select the type of backup to create and click **Next**.
7. The **Options** step lets you name the backup job and define options specific to the backup.



- **Verify backup** - These options tell PHD Virtual Backup how the backups should be verified. By default, this is set to **New blocks only**, which verifies only information that has changed since the last backup. **All blocks** verifies every block of data each time the backup runs and **None** does not verify any data. PHD Virtual recommends selecting either **New blocks only** or **All blocks** for your backup jobs. For detailed information on the verify options, see "[Verifying Backups and Restores with TrueRestore™](#)" (on page 77).
 - **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 in a VM between backups. Instead of reading the entire VMDK each time to discover what has changed, 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. Selecting this option enables Changed Block Tracking 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 finished adding a job name and selecting job options, click **Next**.
 9. Review the Summary information and click **Submit**. The backup job is submitted for processing.
 10. Click **Finish** to close the wizard.

11. 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></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 37)

Chapter 5 - The Restore Wizard

The Restore Wizard lets you restore the virtual machines you backed up with PHD Virtual Backup.

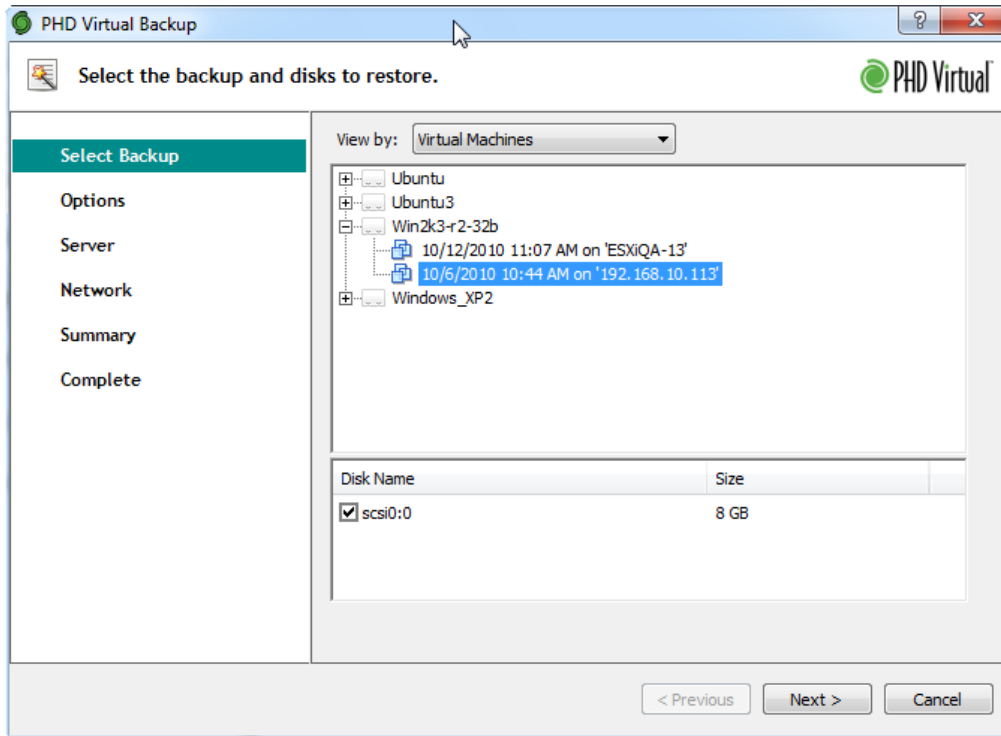
To launch the Restore Wizard

- The wizard can be launched by right-clicking an object within vSphere Client or by using the File menu and selecting **Restore** from the integrated **PHD Virtual Backup** menu.




Using the Restore Wizard

1. Use the **View by** drop-down and the navigation tree to locate the backup you'd like to restore.



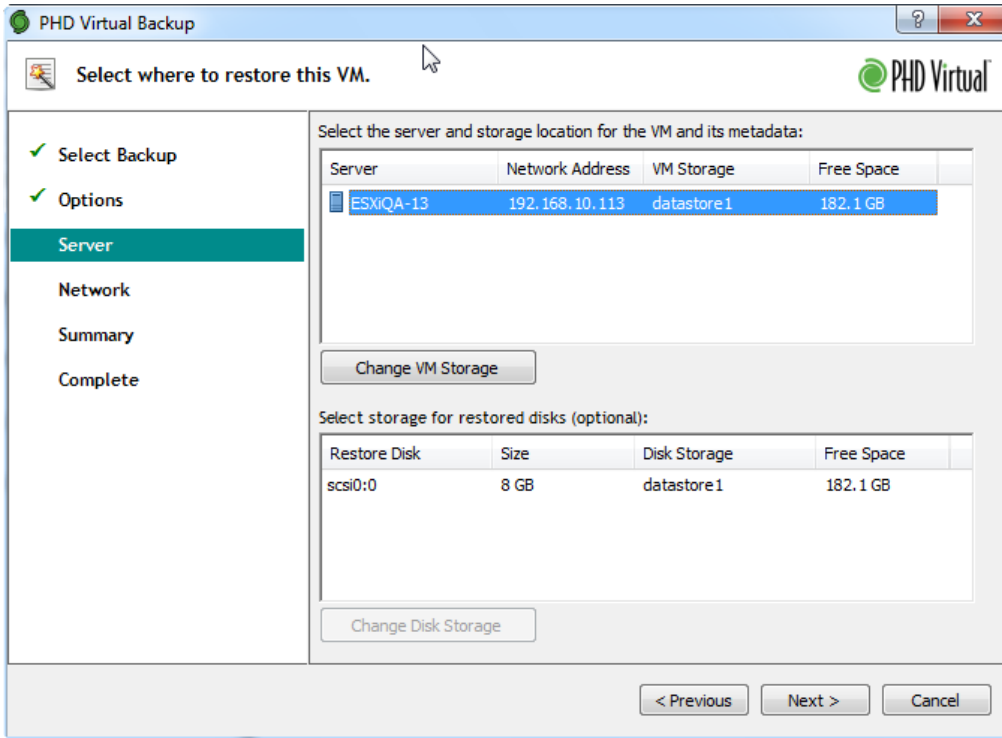
When you select the backup, the available disks are also displayed, as seen in the image above. All available disks are selected for restore by default. To exclude a particular disk from the restore, clear the check box in the **Disk Name** column.

2. Select the disks to restore and click **Next**. The Options step opens.



3. Enter a name for the restore job and a name for the Virtual Machine to be restored.

4. If you want to add additional verification during the restore process, select **Verify Restore**. PHD Virtual recommends selecting this option for your restore jobs. For more information on verifying backups and restores, see "Verifying Backups and Restores with TrueRestore™" (on page 77).
5. Click **Next**.
6. The next step lets you select where the VM should be restored.



Select the target from the available list. You must select a target location with a sufficient amount of free space.

If you need to send an individual disk somewhere other than the selected target, select the disk and click **Change Storage**.

When you've selected where you will restore your VM and disks, click **Next**.

7. Select the network device to use for the restored VM. If you need to change any of the settings, click **Edit**.
8. Click **Next**.
9. Review the summary information for the restore job and click **Submit**.
10. 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 quick reference and step-by-step instructions for using PHD Virtual Backup features.

Creating Backup Jobs.....	66
Running a Backup Now.....	67
Scheduling Backups.....	69
Viewing Jobs.....	71
Restoring Backups.....	73
Restoring Files.....	74
Configuring Email Alerts.....	76
Verifying Backups and Restores with TrueRestore™.....	77
Backup Retention and Archiving.....	78
Excluding VMs and Disks.....	79
Sending Backup Files to Tape.....	80
Limiting the PHD Console to a Single PHD VBA.....	81
Increasing Backup Storage (Attached Disk).....	83
Updating PHD Virtual Backup.....	84

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 by 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.

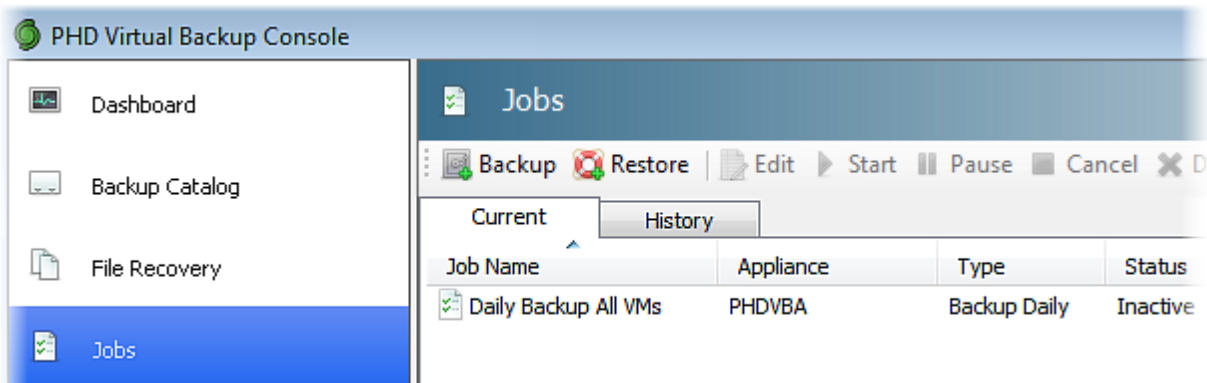
Backup jobs are created using the Backup Wizard, which is launched when you select **PHD Virtual Backup > Backup** from the integrated PHD Virtual Backup menus in vSphere Client, or when you click  **Backup** within the PHD Virtual Backup Console.

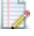
To create a Backup Job

1. Launch the Backup Wizard by right-clicking an object within your vSphere Client 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 ["The Backup Wizard"](#) (on page 56).

To edit a job

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



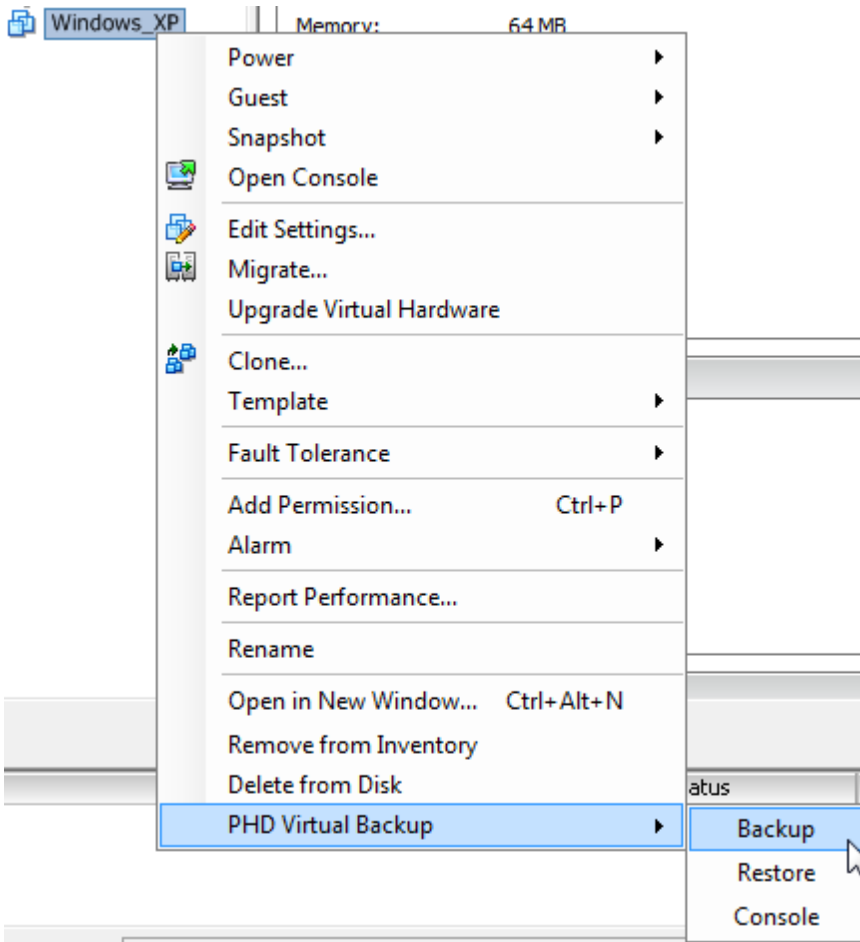
2. Select the job you would like to edit and click  **Edit**.
3. The Backup Wizard opens with the settings you originally defined for the job. Use the wizard to make any edits and submit the job again. For details on each step of the wizard, see ["The Backup Wizard"](#) (on page 56).

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 launch the Backup Wizard which guides you through the process of creating your Backup Job.

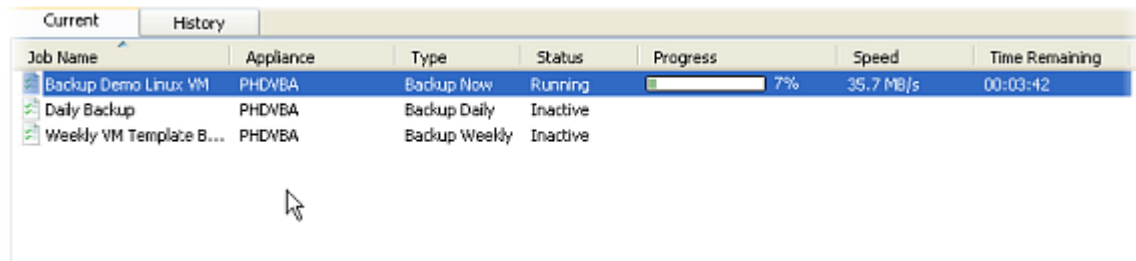
To run a single backup

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 ["The Backup Wizard"](#) (on page 56)

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 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			

Tip: Another way to run a backup right away is to force a scheduled backup to run now.

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, launch 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 and click **Next**.
3. Select the appliance to use for the backup and 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 "[The Backup Wizard](#)" (on page 56).
8. Click **Next**.
9. Review the summary information and click **Submit**.
10. 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 37) for more information.

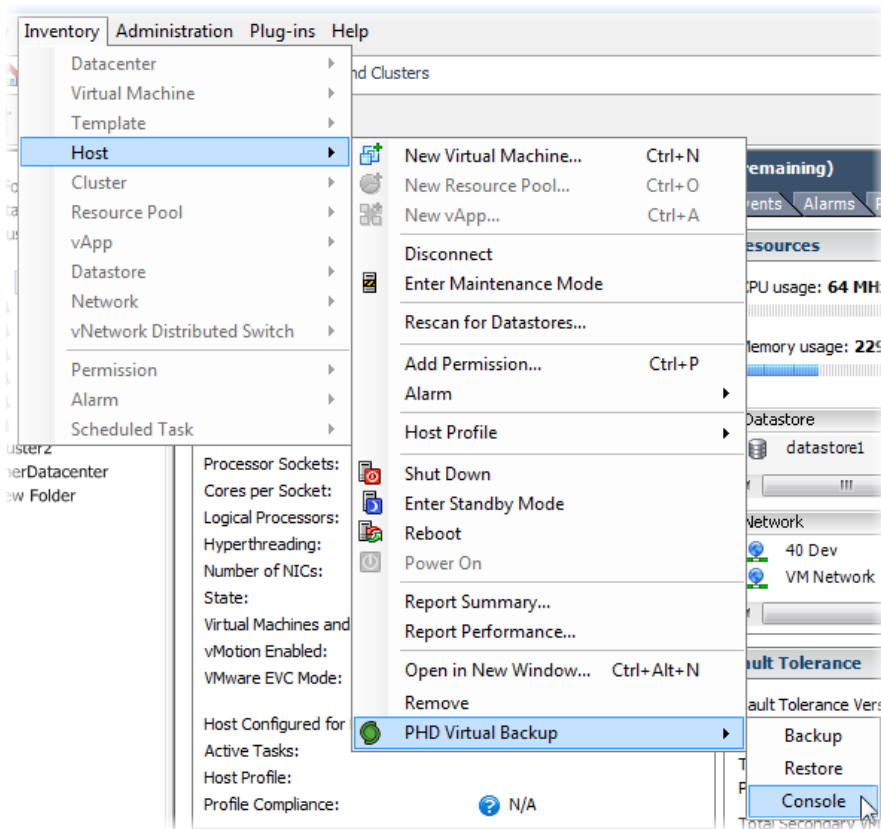
Note: If the PHD Virtual Backup Appliance is restarted within one hour of a scheduled Daily or scheduled Weekly job's start time, the scheduled job will be run again.

Viewing Jobs

To view all of the jobs in progress or scheduled, use the PHD Virtual Backup Console, Jobs page. The console opens automatically after creating a job with either the Backup Wizard or Restore Wizard or it can be launched from within vSphere Client.

To launch 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 currently in progress.

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			

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> 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> 12%
0	Disk 8.6 GB	<div style="width: 12%;"></div> 12%

Details Tasks

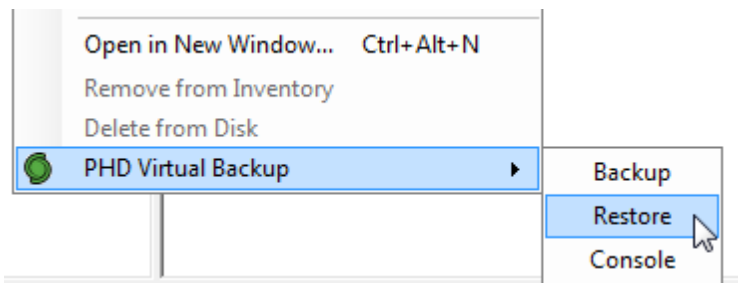
Restoring Backups

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.

Additionally, you can browse the Backup Catalog within the PHD Virtual Backup Console to find the backup you want to restore.

To restore a Virtual Machine

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




Alternatively, you can right-click directly on a VM, and select **Restore** from the PHD Virtual Backup menu. If a backup for that VM is available, the VM is pre-selected within the Restore Wizard catalog.

The Restore Wizard opens and guides you through the process of restoring your selected VM. For detailed information about each step of the wizard, see ["The Restore Wizard" \(on page 62\)](#)


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 launch the File Recovery wizard,



or you can launch 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" \(on page 32\)](#).

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 tab bar with tabs for "General", "Storage", "Network", "Email", "Retention", "Connector", and "Support". The "Email" tab is active. The configuration options include:

- Radio buttons for "Do not email alerts from the appliance" (unselected) and "Email alerts using the following information" (selected).
- Text input for "Server Name" (smtp.example.com) and "Port" (587).
- Dropdown menu for "Security" (None).
- Checked checkbox for "Server requires credentials".
- Text input for "User name" (PHD).
- Text input for "Password" (masked with dots).
- Text input for "From Email Address" (phd@phd@phd.com).
- Dropdown menu for "Alert Level" (All).
- Section "Recipients:" with a text input containing "phd@phd@phd.com" and "Add" and "Remove" buttons.
- "Save" button at the bottom right.

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

The appliance will restart and you will begin receiving PHD Virtual Backup alerts from the appliance you configured. If you are using multiple appliances, you will need to configure alerts for each appliance, individually.

For additional information about each available Email configuration option, see "Email" (on page 48).

Verifying Backups and Restores with TrueRestore™


PHD Virtual Backup's TrueRestore technology ensures the data you backup is the data you can restore.

During the backup and restore processes, PHD Virtual recommends you take advantage of the available verification options. 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. When a bad block is identified, it is flagged, and PHD Virtual Backup Appliance will then attempt to repair the bad block, further ensuring the integrity of your data.

To verify backups


1. At the Options step of the Backup Wizard, use the Verify drop-down box 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" \(on page 50\)](#)


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 on demand, you can use the Backup Catalog to set an Archive flag by clicking **Archive**.



Backups flagged for archive display an archive icon  in the backup catalog, as seen in the image above.

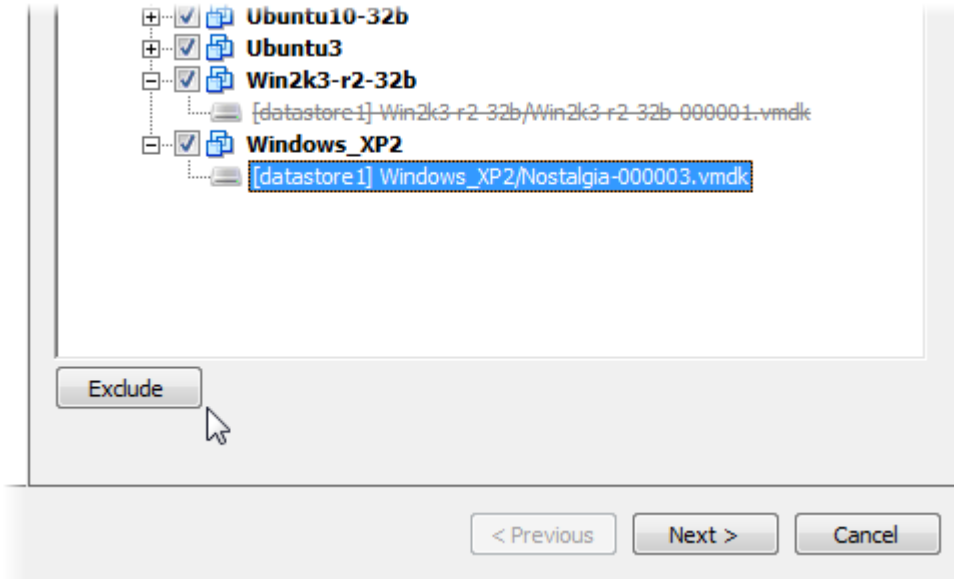
To remove the archive flag, select the backup and click **Archive** again.


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 37) 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 53).

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 may want to limit the PHD Console to display only information for a single PHD VBA. 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.

PHD Virtual Backup - Login

Connect to a hypervisor server by entering its IP address or hostname and a valid username and password.

Hypervisor type: Citrix VMware

Host name: 192.168.42.5

User name: administrator

Password: ●●●●●●●●

Connect

[less...](#)

Show Single VBA: PHDVBA-1

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.

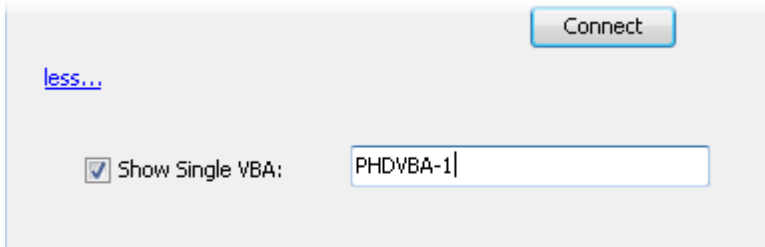
User name: administrator

Password: ●●●●●●●●

Connect

[more...](#)

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 checked checkbox followed by the text "Show Single VBA:". To the right of this 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.

Updating PHD Virtual Backup

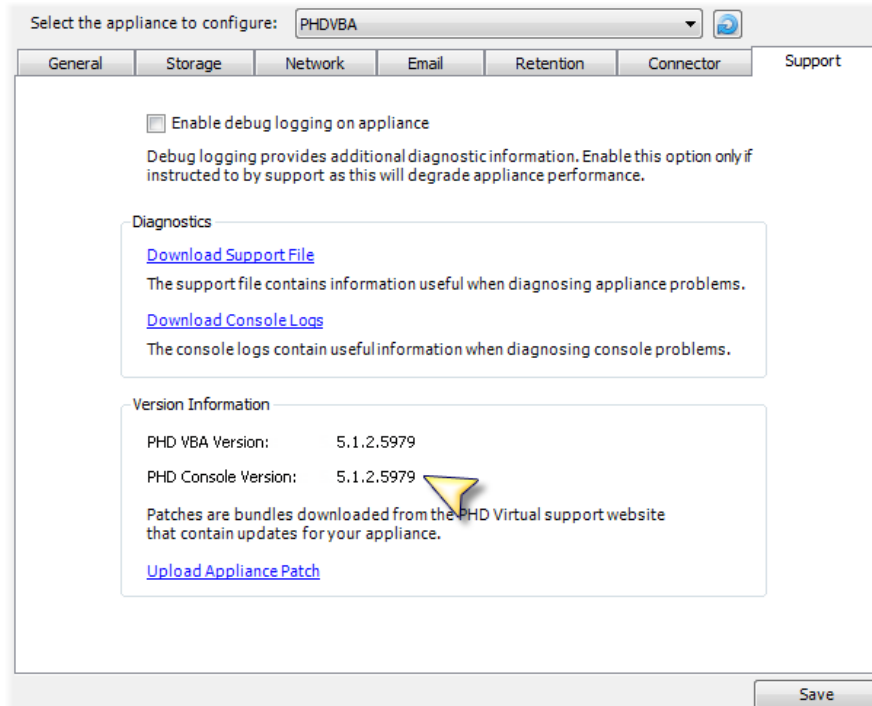
When available, updates to PHD Virtual Backup can be downloaded from the PHD Virtual Web site or obtained from Support. Console and Plug-in updates are made available via an updated MSI file and PHD Virtual Backup Appliance updates are available as update files (.phd files).

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

To update the PHD Virtual Backup Console and Plug-In


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

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

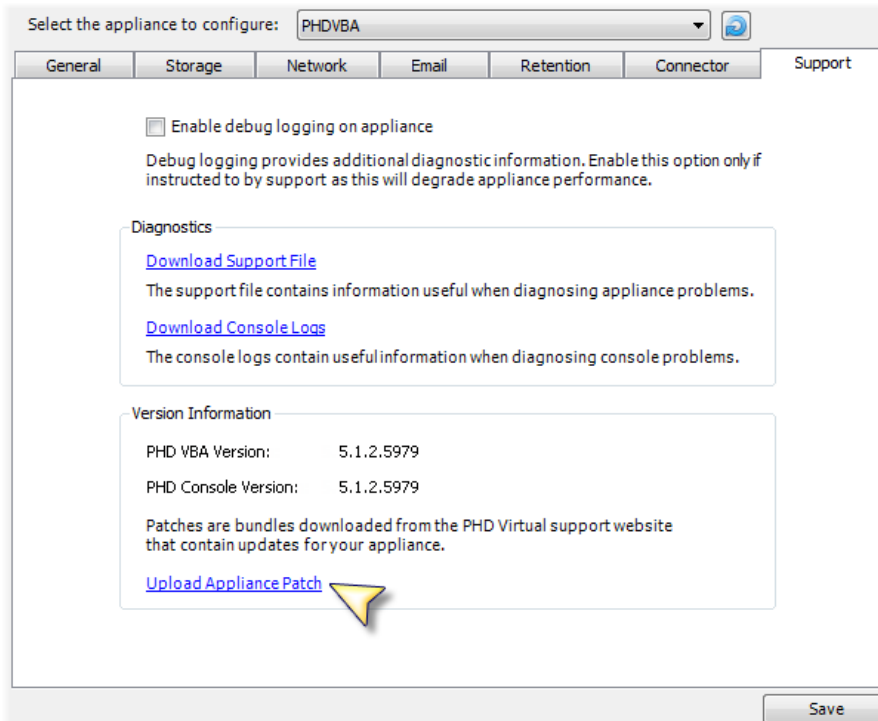


To update the PHD Virtual Backup Appliance

1. Extract the contents of update package.
2. Open the PHD Virtual Backup Console, and click **Configuration** then click the **Support** tab.
3. Use the drop-down menu at the top of the page to select the PHD VBA to update.



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



5. Select the VBA update file (for example, PHDVBA_1234.phd) from the update package and click **Open**.
6. After the update is applied, the appliance 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.

Appendix A - Troubleshooting

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

Support Files.....	87
What To Do If a PHD VBA Crashes.....	88
Resetting PHD VBA Network Settings.....	89
Problems Accessing the BDC Share.....	90
Cannot Power on a VBA.....	91

Support Files

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

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

- An unchecked checkbox labeled "Enable debug logging on appliance".
- A paragraph: "Debug logging provides additional diagnostic information. Enable this option only if instructed to by support as this will degrade appliance performance."
- A section titled "Diagnostics" containing two links:
 - [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."
- A section titled "Version Information" containing:
 - PHD VBA Version: 5.1.2.6156 (for VMware vSphere)
 - PHD Console Version: 5.1.2.5979
 - A paragraph: "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 window.

For additional details about the Support tab, see "Support" (on page 55).

What To Do If a PHD VBA Crashes

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 share

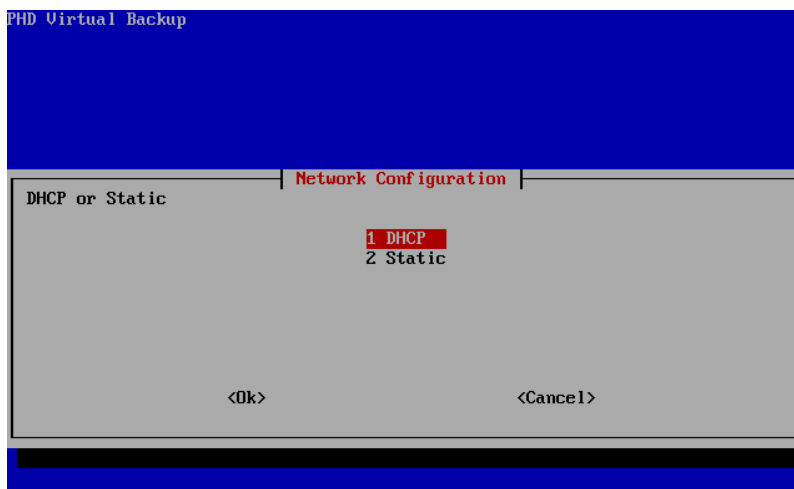
1. Power off the problematic appliance within vSphere Client.
2. Deploy a new appliance. Use the OVF that came with your installation package. Follow the steps in the installation guide for details, making sure to select **CIFS** or **NFS** as the storage type.
3. Click **Save**, then restart the appliance.
4. 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.

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. Reboot the appliance to confirm the updated network settings.

Problems Accessing the BDC Share

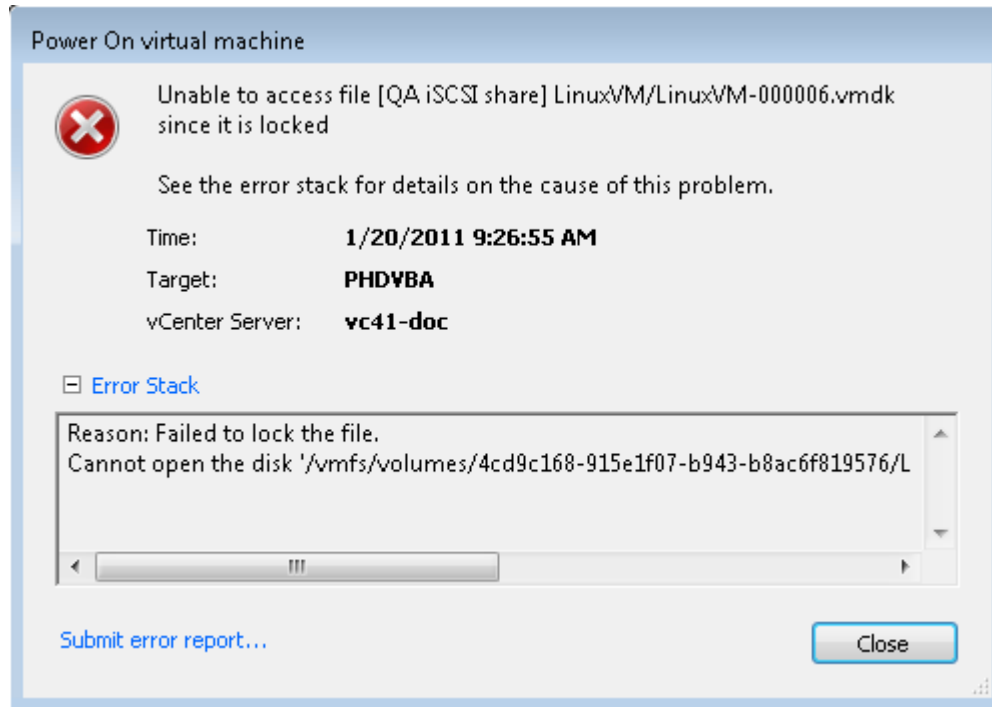
If you cannot access the Backup Data Connector (BDC) share from Windows Vista, Windows 7, or Windows 2008, you may need to adjust your local security policy, LAN Manager authentication level to "LM and NTLM - use NTLMv2 session security if negotiated."

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 drop-down menu and select **LM and NTLM - use NTLMv2 session security if negotiated**.
5. Click **OK**.
6. Now try accessing the Backup Data Connector share again.

Cannot Power on a VBA

If you're experiencing issues 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.



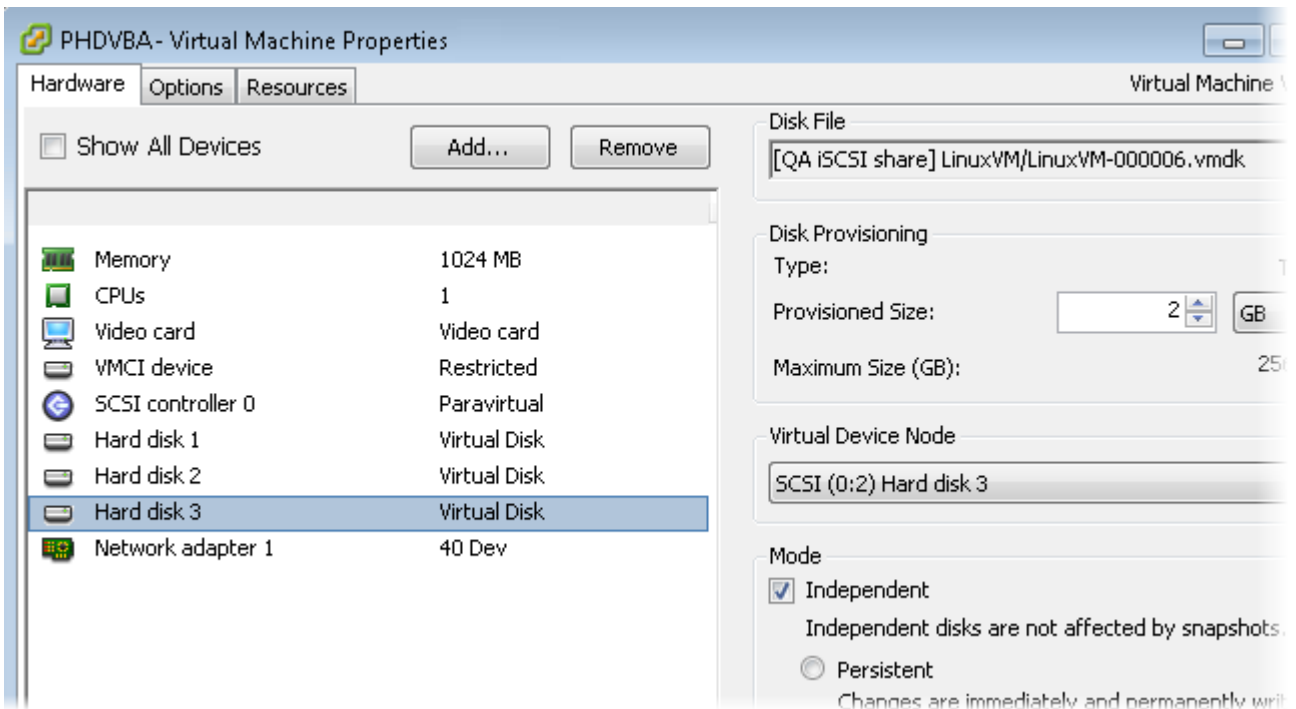
In some instances, when a PHD Virtual Backup Appliance is shutdown in the middle of a backup, it can leave behind a snapshot on the VM it was backing up. Also, that snapshot will remain attached to the powered off PHD Virtual Backup Appliance as an attached virtual disk. When the Appliance starts up,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

Appendix B - Errors and Warnings

Review the following section for information about errors and warnings encountered when using PHD Virtual Backup.

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 rebooting, 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 rebooting 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 after

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

Index

.		Backup is stopped	94
.phd files	84	Backup Jobs	
A		creating	66
Advanced storage options	45	Backup Now	67
Alert Level	49	Backup powered off virtual machines	60
All	48	Backup Retention	78
All blocks	60	Backup storage	44
Another PHDVB VBA has a snapshot	93	Backup Storage	
Antivirus	19	increasing	83
Antivirus Software	19	Backup Virtual Machine	38
Appliance	27	Backup Wizard	56, 66-67
Appliance Crash	88	using	56
Appliance options	42	Backups	13
Appliance updates	84	verify	77
archive backups	53	BDC	53
Archive backups	60	Best Practices	19
Archiving Backups	78	C	
assign static appliance network settings	47	Cancel	38
Average Speed	39	Change Storage	64
B		Changed Block Tracking	60
Backup Appliances	27	CIFS/SMB Shares	19
Backup Catalog	29	Collection of metadata failed	93
deleting backups	78	Configuration	41
Backup Catalog Notes	30	reload values	41
backup data		Configuration page	41
self-healing	77	Connector	53
Backup Data Connector	53, 80, 90	Connector tab	53
troubleshooting	90	Console and Plug-in updates	84
		Could not attach	93
		Could not write and close backup block	93

Index

create a Backup Job	66	edit a job	66
create a scheduled backup job	69	Email	48
Creating Backup Jobs	66	Email Alerts	76
Critical	48	enable alerts	48
Critical errors	48	Enable compression for new backups	45
CTRL-N	89	Enable debug logging on appliance	55
Custom		Errors	48, 93
retention setting	50	Exclude	57
D		Excluding VMs	79
Daily	69	export backups	53
Dashboard	26	Exporting Backups	30
Backup Appliances list columns	27	F	
Data Streams	42	Failed to add snapshot	94
Data Written	39	Failed to save changes	94
debug mode	55	File Recovery	32
debugging mode	55	Folder	
Dedupe Ratio	27	View by option	56
DeDupe Ratio	39	Free Storage	27
defragmentation	19	Frequently Asked Questions	17
Defragmenting	19	G	
Delete	38	General	42
delete iSCSI target	36	General tab	41-42
Delete trim	40, 50	H	
Deleting backups	30	Help	20
Deleting iSCSI targets	36	How many appliances do I need?	21
disable email alerts	49	How PHD Virtual Backup Works	12
Disk Defragmenter	19	Hypervisor Credentials	41-43
display Job Details	38	I	
Do not start after	59, 69	Individual backups	
Documentation Updates	4	deleting	78
does not have a UUID	94	inf	17
E			
Edit	38		

IP address		Network Settings	
appliance	46	reset	89
obtain automatically for appliance	47	Network tab	46
IP Address	27	New blocks only,	60
iSCSI Software Initiator	75	Next Run	59, 70
J		NFS Shares	19
Job Details	38	NTP servers	42
Jobs	37	O	
Jobs History	40	Once	69
K		Options	
Keep All		backup wizard	59
retention setting	50	Orphan Weekly	40
L		P	
Last32Days	54	Pause	38
Last7Days	54	PHD Console	9, 81
LatestofEach	54	limiting to a single PHD VBA	81
launch the PHD Virtual Backup Console	71	PHD VBA	9
launch the Restore Wizard	62	PHD Virtual Backup	
license		benefits	10
update	43	receiving alerts	76
upload new	43	updating	84
Licensing	43	PHD Virtual Backup Appliance	9
Limiting the PHD Console	81	reset network settings	89
lock error	91	unavailable	88
M		updating	55, 85
Mount iSCSI target on this computer	34	PHD Virtual Backup Components	16
Mounting iSCSI Targets on Other Devices	36	PHD Virtual Backup Console	9, 24, 41
MSI	84	accessing	24
N		updating	84
Network	46	PHD Virtual Backup Plug-in	9
		PHD Virtual Support	55
		PHDVB	9

Index

Plug-In		Retention tab	50
updating	84	run a scheduled backup now	68
Pool		run a single backup	67
View by option	56	Running a Backup Now	67
port	43		
port 443	43	S	
Power on a VBA	91	Schedule	
Product expiration	43	backup wizard	58
		Scheduling Backups	69
Q		self-healing	77
Quiesce the VM before backing up	60	Sending Backup Files to Tape	80
		Show Details	38, 72
R		Show system jobs	38
Raw		Show/Hide Details	38
exporting backups as	30	Single PHD VBA	81
Recent backups to keep	51	Snap Hunt	40
Recurrence	39	snapshot	91
Recurring jobs	59, 69	Speed	39
Recurs every	59, 69	Start Date	59, 69
Renaming VMs	30	Start Time	59, 69
Resetting VBA Network Settings	89	Start/Resume	38
restore a Virtual Machine	73	Startup	40
restore file		Static IP Address	47
Linux	34	Stop level % free	45
Restore Notes	14	Storage	44
Restore Virtual Machine	38	Support	3, 55
Restore Wizard	62, 73	Support expiration	43
Restores	14	Support Files	87
verify	77	submitting	87
Restoring Backups	73	Support tab	55
Restoring Files	32, 74	System Alert descriptions	28
Restoring Files from a Linux VM	34	System Alerts	26, 28
Restoring Virtual Machines	29	System Jobs	40
Retention	78		
Retention Settings	50		

T			
tape	53	View by	56
Terms	9	View Log	38
The PHD Virtual Backup Appliance	21	Viewing Jobs	71
The Restore Wizard	62	Virtual Backup Appliance	7
Total Backup Data	27	Virtual Hard Disks	
Trim	78	exporting backups as	30
Troubleshooting	86	VMDK	30
TrueRestore	77	VMDK export	30
Typical		VMware vSphere	7
retention setting	50	Volume Shadow Copy Services	60
U		W	
update packages	84	Warning level % free	45
Updating PHD Virtual Backup	84	Warnings	48, 93
Upload Appliance Patch	55	Weekly	69
Uploading Appliance Patches	55	What's New	8
Use Changed Block Tracking	60		
Used Storage	27		
Using PHD Virtual Backup	65		
Using the Restore Wizard	62		
UUID	29		
V			
VBA	7, 9, 16		
VBA already has a snapshot	94		
VBA Console	23		
Verify backup	60		
verify backups	77		
Verify Restore	64		
verify restores	77		
Verifying Backups and Restores	77		
VHD	30		
Video Tutorials	20		