

PHD Virtual Backup

for **CITRIX** XenServer®

version 5.1

User Guide

Software Release Date: December 2010

Document Release Date: March 23, 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.

Citrix, Xen, XenServer, XenDesktop, XenMotion, and XenCenter are either trademarks or registered trademarks of Citrix Systems, Inc.

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
2011-02-01	1	"How PHD Virtual Backup Works" (on page 11). Added a notes section for restores.
2011-02-24	1	"Backups" (on page 12). Updated the conceptual overview.
2010-12-27	3	"Jobs" (on page 34). Added new sub-section with additional details about PHD Virtual Backup job speeds and deduplication.
2011-02-01	3	"Job History" (on page 37). Job History tab was enhanced to include icons in the result column. (5.1.2)

Contents

Chapter 1 - Welcome	7
What's New.....	8
About This Guide.....	9
Benefits of PHD Virtual Backup.....	10
How PHD Virtual Backup Works.....	11
Backups.....	12
Restores.....	13
PHD Virtual Backup Components.....	14
Frequently Asked Questions.....	15
Best Practices.....	17
Getting Help.....	18
Chapter 2 - The PHD Virtual Backup Appliance	19
The PHD VBA Console.....	21
Chapter 3 - The PHD Virtual Backup Console	22
Dashboard.....	23
Backup Catalog.....	26
File Recovery.....	29
Restoring Files.....	29
Restoring Files from a Linux VM on Windows.....	31
Mounting iSCSI Targets on Other Devices.....	33
Deleting iSCSI targets.....	33
Jobs.....	34
Job Details.....	35
Job Speeds, Deduplication, and Data Written.....	36
Job Types.....	36
Job History.....	37
Configuration.....	38
General.....	39
Storage.....	41
Network.....	43
Using DHCP.....	44

Using Static IP Addresses	44
Email	45
Retention	47
Connector	50
Support	52
Chapter 4 - The Backup Wizard	53
Chapter 5 - The Restore Wizard	59
Chapter 6 - Using PHD Virtual Backup	63
Creating Backup Jobs	64
Running a Backup Now	65
Scheduling Backups	67
Viewing Jobs	69
Restoring Backups	71
Restoring Files	72
Configuring Email Alerts	74
Verifying Backups and Restores with TrueRestore™	75
Backup Retention and Archiving	76
Excluding VMs and Disks	77
Skipping VMs	78
Sending Backup Files to Tape	80
Using Tags to Backup VMs	81
Increasing Backup Storage (Attached Disk)	82
XenServer System Logging	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
Appendix B - Errors and Warnings	91
Index	93

Chapter 1 - Welcome

PHD Virtual™ Backup for Citrix XenServer® provides reliable backup and recovery for all of the virtual machines (VMs) in your XenServer environment. With PHD Virtual Backup, you can manage backup and recovery right from within XenCenter using simple, integrated menus. Using the PHD Virtual Backup Console and wizards, you can you create and manage custom backup and restore jobs to meet all of your data protection requirements - without interrupting workflow and without directly accessing Xen Domain Zero (Dom0).

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 XenServer 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
How PHD Virtual Backup Works.....	11
Frequently Asked Questions.....	15
Best Practices.....	17
Getting Help.....	18

What's New

- File Level Restore - mount a backup as an iSCSI share and restore individual files. See ["Restoring Files" \(on page 72\)](#) 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" \(on page 41\)](#).
- Use the new Archive Backup feature to preserve backups outside of normal retention policy settings. See ["Backup Catalog" \(on page 26\)](#) for details.
- Enhanced backup retention settings let you select from pre-configured policies or customize your own. For details, see ["Retention" \(on page 47\)](#).
- Define custom levels to warn you when your backup storage runs low. See ["Storage" \(on page 41\)](#)
- 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" \(on page 50\)](#)
- 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 Citrix XenServer 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 Citrix XenServer - 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 XenCenter 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

The First Citrix Ready Solution for Virtual Backup and Recovery

PHD Virtual Backup is built upon the next generation of PHD Virtual's VBA architecture and version 5.1 extends the Citrix XenServer virtualization backup and recovery solution with new capabilities. PHD Virtual Backup provides:

- XenCenter management integration. With the plug-in for XenCenter, PHD Virtual Backup provides "single pane of glass" management of your virtual machine backup and restore right from the XenCenter 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 (hosts, pools, folders, tags) 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.

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 XenServer resource pool, the PHD Virtual Backup Appliance performs the backup and restore processing for all of the VMs within that pool (as long as they are using shared storage).

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


- ["Backups" \(on page 12\)](#)
- ["Restores" \(on page 13\)](#)
- ["PHD Virtual Backup Components" \(on page 14\)](#)

Backups


When a backup is run, the PHD VBA first reads the target VM metadata and creates a snapshot.




Next, the virtual disks created with the snapshot are attached to the PHD VBA as new virtual disks and the snapshot is removed.



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




Finally, the virtual disks are detached from the PHD VBA and removed.



Restores

When a virtual machine restore job is created, the appliance searches the storage location for the matching VM metadata, networking information, 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" \(on page 72\)](#) 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 VM backup selected is restored using the PHD Virtual Backup Appliance used to perform the backup (the VBA that has access to the storage location on which the backup resides).
- Restoring VMs to the same location where the original VM is located will result in VMs with duplicate names in XenCenter.

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" \(on page 19\)](#)
- **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 XenCenter or from the Windows Start Menu. For more information, see ["The PHD Virtual Backup Console" \(on page 22\)](#).
- **Backup Wizard** - The wizard which guides you through the steps of creating and editing backup jobs. See ["The Backup Wizard" \(on page 53\)](#) 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" \(on page 59\)](#) 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 XVA - XenServer's virtual appliance format. 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" \(on page 47\)](#).

What happens if my appliance is rebooted during a backup?

The running backup job will be canceled and when the appliance restarts, any leftover snapshots or data will be cleaned up and removed. 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" \(on page 50\)](#).

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).

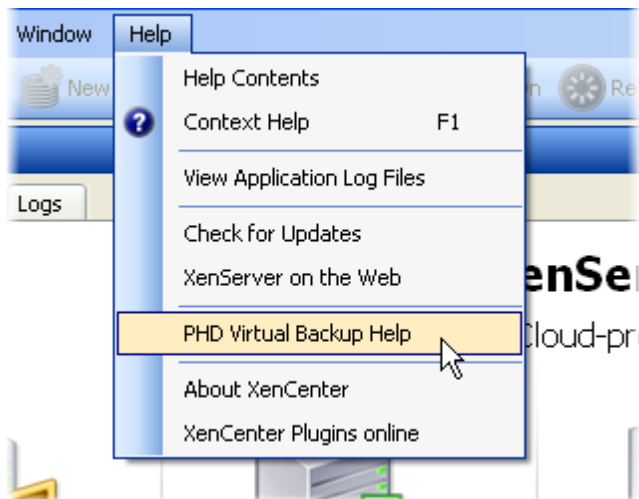
Disk Defragmenter

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

Figure 1 - The PHD Virtual Backup Appliance in XenCenter



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.

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 22) for details.


PHD VBA status and log information can also be seen by selecting the PHD VBA virtual machine within XenCenter then clicking the Console tab. See "[The PHD VBA Console](#)" (on page 21).

How many VBAs do I need?

- You will need to deploy at least one PHD Virtual Backup Appliance per resource pool. Each appliance can perform backups and restores for VMs within the same resource pool. If you have more than one resource pool, you will need to deploy an additional appliance to each pool. Depending on your environment, you may choose to use multiple appliances within each pool, though only one per pool is required.

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

Figure 2 - PHD Virtual Backup VBA in a XenServer Resource Pool 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. When the appliance starts up again, a system job runs and cleans up any snapshots leftover from the job that was in progress. 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.

The PHD VBA Console

Viewing the PHD VBA virtual machine console within XenCenter (in XenCenter, 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 XenCenter

```

PHD Virtual Backup for Citrix XenServer v5.0.0.2282 09:29:20 2
Worker Queue Depth: 0           Utility Queue Depth: 0
Worker Threads: 4              Utility Threads: 3
Store: 1 GB used, 9 GB free    Deduplication Ratio: 57:1
PHDVB Appliance Log:
09:28:59 Worker-2 r Debian Etch 4.0: Collected metadata
09:28:59 Worker-2 r Debian Etch 4.0: 2 disk(s) to restore
09:28:59 Worker-2 r Debian Etch 4.0: Recreated VM as r Debian Etch 4.0
09:28:59 Worker-2 r Debian Etch 4.0: Recreated network(s)
09:28:59 Worker-2 r Debian Etch 4.0: Recreated optical drive(s)
09:28:59 Worker-1 r Debian Etch 4.0: Allocated new disk 1
09:29:00 Worker-1 r Debian Etch 4.0: Linked new disk 1 to PHDVB
09:29:02 Worker-1 r Debian Etch 4.0: Attached new disk 1 to PHDVB
09:29:02 Worker-1 r Debian Etch 4.0: Restoring disk 1: 0% of 549 MB
09:29:05 Worker-4 r Debian Etch 4.0: Allocated new disk 0
09:29:06 Worker-4 r Debian Etch 4.0: Linked new disk 0 to PHDVB
09:29:08 Worker-4 r Debian Etch 4.0: Attached new disk 0 to PHDVB
09:29:08 Worker-4 r Debian Etch 4.0: Restoring disk 0: 0% of 4 GB
PHDVB Worker Thread Status:
Worker-1: r Debian Etch 4.0: Restoring disk 1: 70% of 549 MB
Worker-2: (idle)
Worker-3: Windows Server 2003: Backing up disk 0: 7% of 8 GB @ inf:1
Worker-4: r Debian Etch 4.0: Restoring disk 0: 5% of 4 GB


```

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 XenCenter, see "To launch the PHD Virtual Backup Console" (on page 69)
- The Console can also be launched as a stand-alone application from the Windows Start Menu.

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

- "Dashboard" (on page 23)
- "Backup Catalog" (on page 26)
- "File Recovery" (on page 29)
- "Jobs" (on page 34)
- "Configuration" (on page 38)

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.

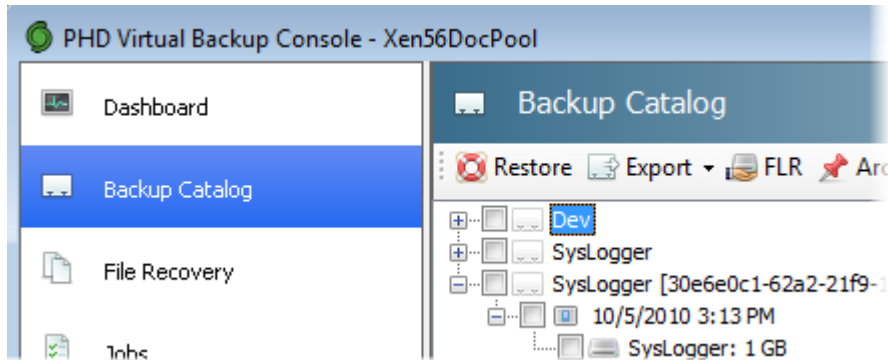
The following table provides additional information about some of the system alerts you may encounter.

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 XenCenter 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, as seen in the following image.





Table 4 - Backup Catalog Toolbar Buttons


Button icon	Description
Restore	Launches the Restore Wizard. For details, see "The Restore Wizard" (on page 59) .
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 29) .
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 76) .
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" \(on page 59\)](#) 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" \(on page 34\)](#) 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 29).
- "Restoring Files from a Linux VM on Windows " (on page 31).
- "Mounting iSCSI Targets on Other Devices" (on page 33).
- "Deleting iSCSI targets" (on page 33).

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 29) 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 31).

To mount an iSCSI target on another device, see "Mounting iSCSI Targets on Other Devices" (on page 33).

To delete an iSCSI target, see "Deleting iSCSI targets" (on page 33).

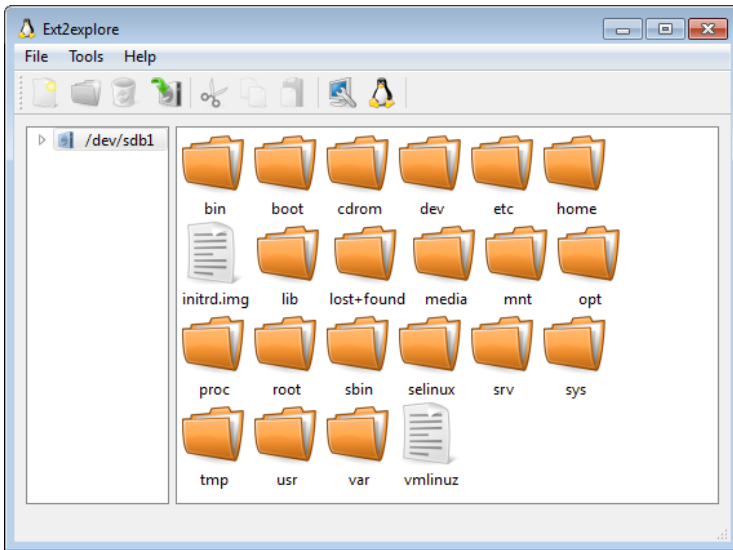
Restoring Files from a Linux VM on Windows

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" \(on page 29\)](#).

To mount an iSCSI target on another device, see ["Mounting iSCSI Targets on Other Devices" \(on page 33\)](#).

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 " \(on page 31\)](#).

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.

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



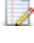







Job Name	Appliance	Type	Status	Progress	Current Speed	Time Remaining
Backup Exchange Server	PHDVBA	Backup Now	Running	100%		
Daily Backup	PHDVBA	Backup Daily	Inactive			
Weekly Template Backup	PHDVBA	Backup Weekly	Inactive			

Job Detail	Value
Schedule	
Type	Now
Other	
Created	11/10/2010 2:53 PM
Next Run	
Started	11/10/2010 2:53 PM
Duration	00:00:29
Average Speed	35.3 MB/s
Dedupe Ratio	211:1
Data Written	4.9 MB

Task Name	Type	Status	Dedupe Ratio
Exchange_Server	Virtual Machine	Completed	211:1
SysLogger	Disk 1 GB	Completed	211:1
Windows Server 2...	Virtual Machine	0%	inf: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 53) for details.
 Restore	Launches the Restore Wizard which guides you through the process of restoring stored backups. See " The Restore Wizard " (on page 59) 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.

Table 7 - Job Details

Job Details Parameter	Description
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 67).
Created	The date and time the job was created.
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.

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.

Job Types

PHD Virtual Backup Job types include:


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

System Jobs include:

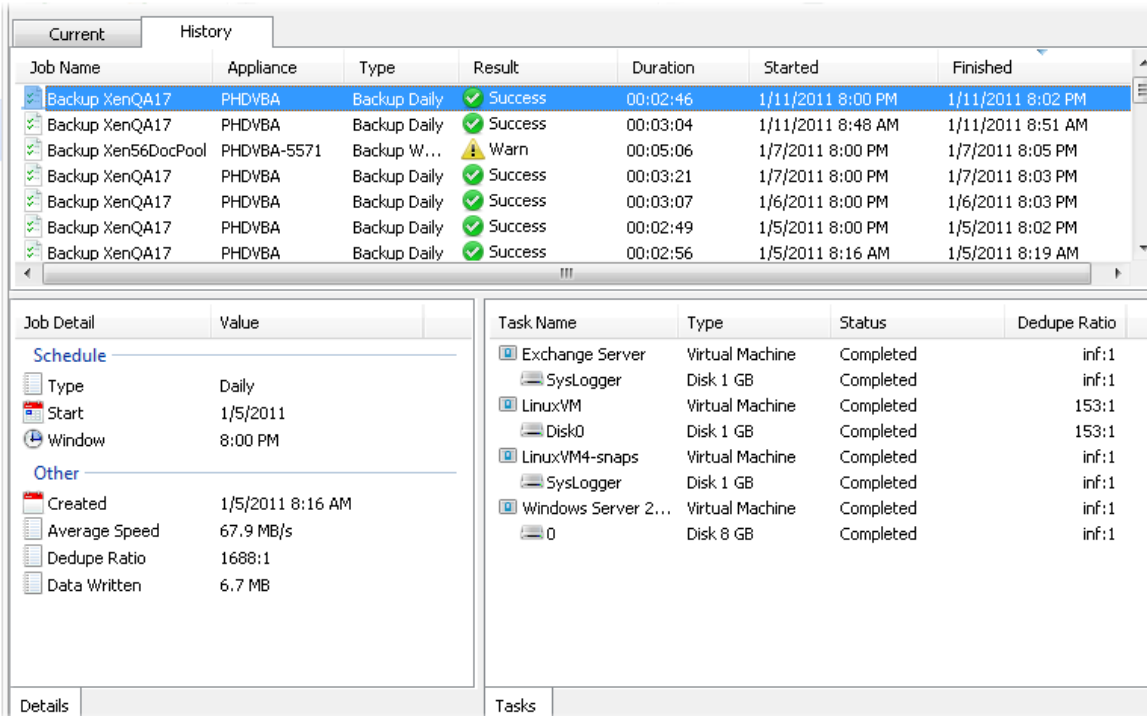
- **Startup** - The job that runs when the appliance first starts. This job cleans up any unfinished processes as well as synchronizes the backup catalog with the backup storage.

- **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 47) for details on setting your retention policy.

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).



Job Name	Appliance	Type	Result	Duration	Started	Finished
Backup XenQA17	PHDVBA	Backup Daily	Success	00:02:46	1/11/2011 8:00 PM	1/11/2011 8:02 PM
Backup XenQA17	PHDVBA	Backup Daily	Success	00:03:04	1/11/2011 8:48 AM	1/11/2011 8:51 AM
Backup Xen56DocPool	PHDVBA-5571	Backup W...	Warn	00:05:06	1/7/2011 8:00 PM	1/7/2011 8:05 PM
Backup XenQA17	PHDVBA	Backup Daily	Success	00:03:21	1/7/2011 8:00 PM	1/7/2011 8:03 PM
Backup XenQA17	PHDVBA	Backup Daily	Success	00:03:07	1/6/2011 8:00 PM	1/6/2011 8:03 PM
Backup XenQA17	PHDVBA	Backup Daily	Success	00:02:49	1/5/2011 8:00 PM	1/5/2011 8:02 PM
Backup XenQA17	PHDVBA	Backup Daily	Success	00:02:56	1/5/2011 8:16 AM	1/5/2011 8:19 AM

Job Detail	Value
Schedule	
Type	Daily
Start	1/5/2011
Window	8:00 PM
Other	
Created	1/5/2011 8:16 AM
Average Speed	67.9 MB/s
Dedupe Ratio	1688:1
Data Written	6.7 MB


Task Name	Type	Status	Dedupe Ratio
Exchange Server	Virtual Machine	Completed	inf:1
SysLogger	Disk 1 GB	Completed	inf:1
LinuxVM	Virtual Machine	Completed	153:1
Disk0	Disk 1 GB	Completed	153:1
LinuxVM4-snaps	Virtual Machine	Completed	inf:1
SysLogger	Disk 1 GB	Completed	inf:1
Windows Server 2...	Virtual Machine	Completed	inf:1
0	Disk 8 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 39)
- "Storage" (on page 41)
- "Network" (on page 43)
- "Email" (on page 45)
- "Retention" (on page 47)
- "Connector" (on page 50)
- "Support" (on page 52)

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.

Select the appliance to configure: PHDVBA

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

Appliance options

- Select time zone: America
- Select region: New_York
- NTP Server 1: ntp.ubuntu.com
- NTP Server 2:
- Data Streams: 4

Hypervisor credentials

- Pool Master: 192.168.10.117
e.g., server.example.com or IP address
- User Name: root
- Password: ●●●●●●

Professional License: PHD Virtual

- Product Expiration: Friday, November 18, 2011
- Support Expiration: Friday, November 18, 2011
- [Update](#)

Save

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.

Enter the Management IP address of the Pool Master XenServer along with the root account credentials when configuring each appliance. If you are using a single host, only, use the IP address or host name of that host, instead.

Note: If the XenServer Pool Master changes, each PHD Virtual Backup Appliance must be updated with the new Pool Master credentials in order to run backups and restores.

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 PHD Virtual Backup Console. At the top, there is a dropdown menu labeled "Select the appliance to configure:" with "PHDVBA" selected. Below this are several tabs: "General", "Storage", "Network", "Email", "Retention", "Connector", and "Support". The "Storage" tab is active. The "Backup storage" section contains a "Storage Type" dropdown menu set to "Attached Virtual Disk" and a status indicator: a green checkmark followed by "Using attached disk 10 GB". The "Advanced options" section includes a checked checkbox for "Enable compression for new backups", a "Warning level % free" spinner set to "10.00" with the text "Warns at 1 GB of free storage", and a "Stop level % free" spinner set to "3.00" with the text "Stops at 307.2 MB of free storage". A "Reset to Defaults" button is located below these options. A "Save" button is positioned 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.

The screenshot shows the configuration interface for a PHD Virtual Backup Appliance (VBA). 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" (which is currently selected), "Email", "Retention", "Connector", and "Support". The "Network" tab is divided into two sections: "Adapter" and "Name Servers".

Adapter Section:

- 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 Section:

- Obtain DNS address automatically
- Use the following DNS addresses
- Preferred DNS: [. . .]
- Alternate DNS: [. . .]

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

Note: if you are experiencing network problems you can manually assign network settings by selecting the VBA within XenCenter 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" \(on page 44\)](#)

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" \(on page 44\)](#).

Email

Use the Email tab if you want to receive email alerts from PHD Virtual Backup. You can enter your email server options. If you choose to not send alerts from the appliance, you can still use XenCenter to receive alerts (see your XenServer documentation for instructions on setting up email alerting). By default, warning and error alerts are sent to XenCenter by the appliance and will be displayed in the System Alerts dialog.

You can select to send email alerts for Critical errors, Errors, or All, which includes backup and restore completions, system alerts, and errors. Warnings, though logged within XenCenter, are not sent as email alerts.

The screenshot shows the configuration window for the PHD Virtual Backup Console. At the top, there is a dropdown menu labeled "Select the appliance to configure:" with "PHDVBA" selected. Below this are several tabs: "General", "Storage", "Network", "Email", "Retention", "Connector", and "Support". The "Email" tab is active. The configuration options are as follows:

- 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:

At the bottom right of the window is a button.

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.

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

- "Retention setting" is a dropdown menu currently set to "Typical".
- "Recent backups to keep" is a text input field containing the number "5".
- Below this is the text "And keep the most recent backup from each of the last:" followed by four rows of settings:
 - "Days" is a text input field containing "7".
 - "Weeks" is a text input field containing "4".
 - "Months" is a text input field containing "12".
 - "Years" is a text input field containing "5".

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

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 and after every backup completes.

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

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

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.

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 26) 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: 5.1.0.4203 (for Citrix XenServer)

PHD Console Version: 5.1.0.4211

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

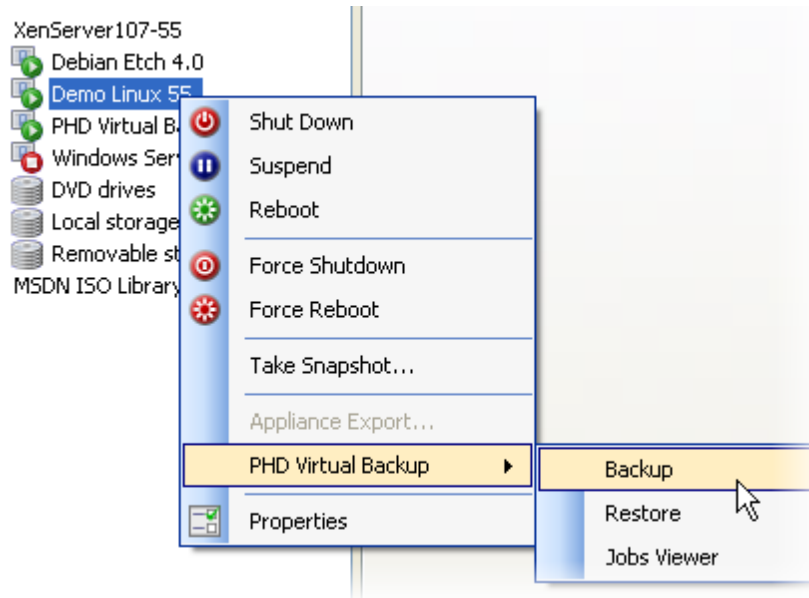
Periodically, update patches for the PHD Virtual Appliance will be available for download from the PHD Virtual Web site. When downloaded to your local computer, they can be uploaded through the PHD Virtual Backup console using the **Upload Appliance Patch** link. Clicking this link will allow you to 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.



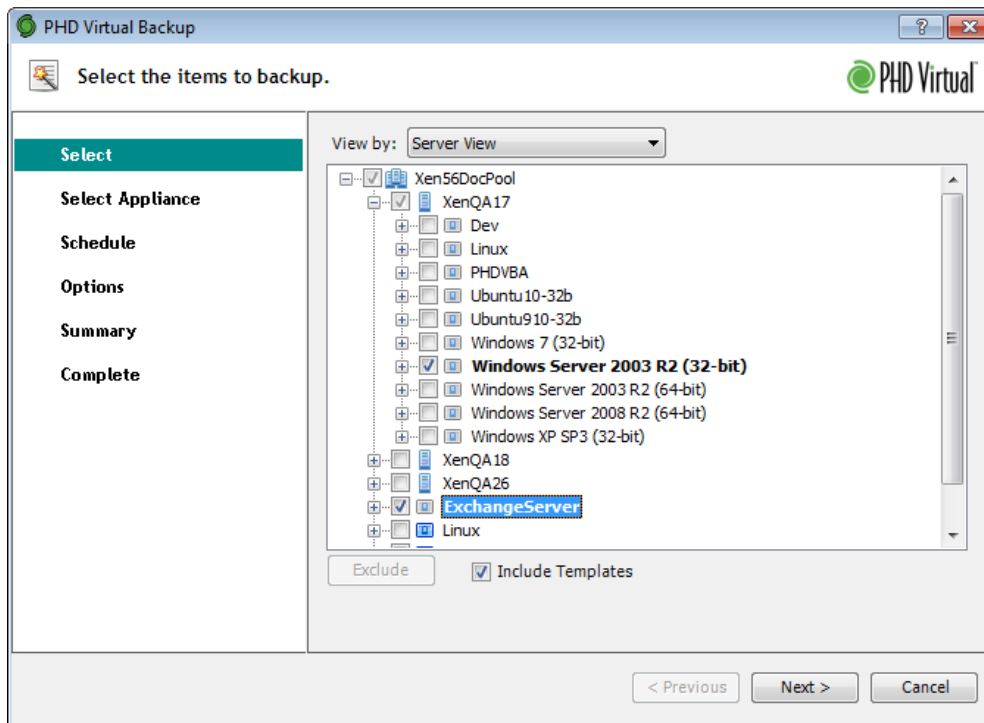
Using the Backup Wizard

- 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.

Server View - Display all VMs within the XenServer Pool, by XenServer host.

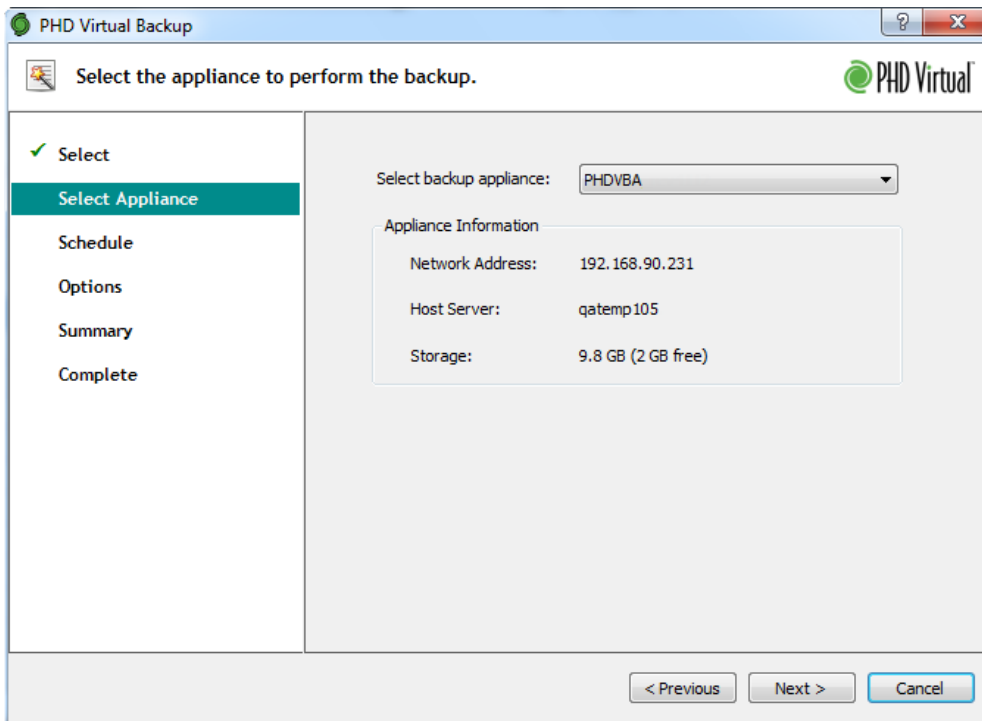
Folder - Display only VMs assigned to folders.

Tag - Display only VMs that contain tags.



If you select the top container in any view (for example, Xen56DocPool in the image above) all VMs in that pool or folder or with that tag applied will be included in the backup job. Also, any VMs added to a selected pool or folder or with the tag applied in the future will also be included in the backup job. Likewise, any VMs removed from the selected pool, folder, or tag will no longer be included in the job.

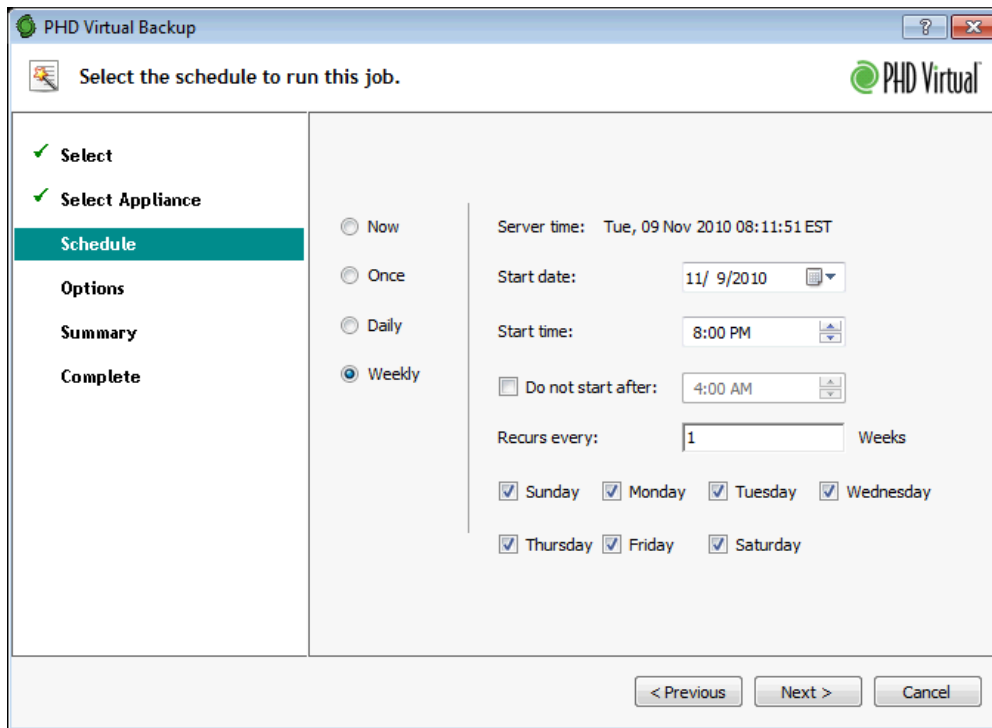
- **Include Templates** - Show or hide Templates in the list of displayed objects.
 - **Exclude/Include** - When backing up groups of VMs, an entire folder, for example, you can choose to exclude specific VMs or individual disks from the backup job by selecting the VM or disk and clicking the Exclude button. VMs can be included again by editing the job.
2. Select the VMs you want to backup and click **Next**.
 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.

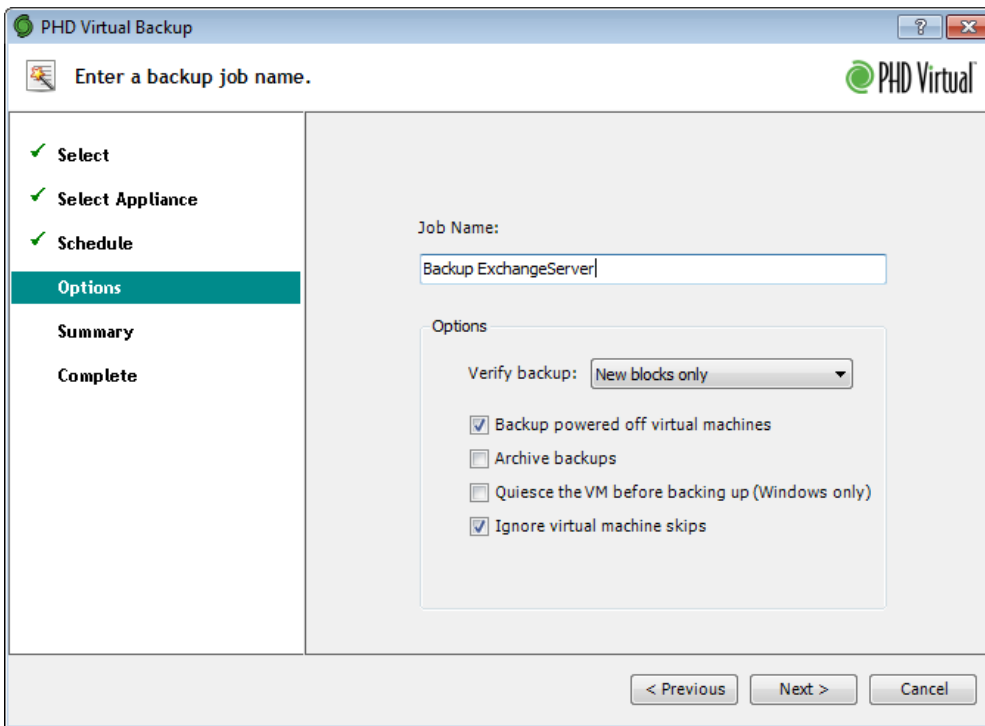
Note: If you will be backing up a VM located on local storage, you must select an appliance that is located on the same host as the VM or else the backup will fail. Virtual disks for any VMs that are unreachable by an appliance (on different local or shared storage, for example) will be displayed after 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 67).



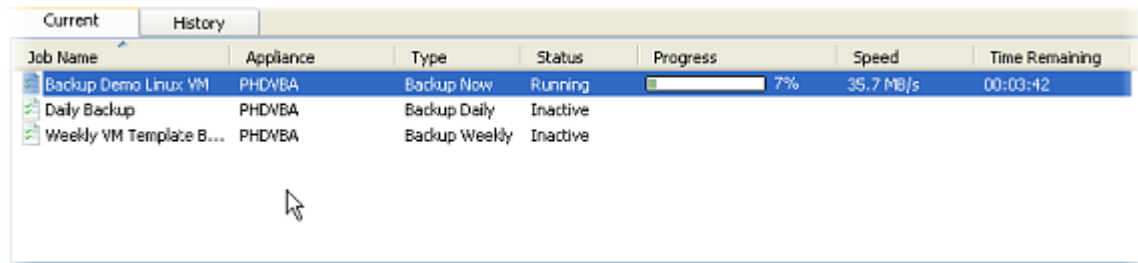
- **Start Date**- The date the scheduled job will begin.
- **Start Time**- The time the job should start.
- **Do not start after**- The time after which the job should not start. In a situation where many backup jobs or very large jobs are running and this time passes before the job can begin, it will 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 75).
 - **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 XenServer tools are installed, you can choose to quiesce the VM before backing it up, to take advantage of Microsoft's Volume Shadow Copy Services.
 - **Ignore virtual machine skips** - Select this option to ignore any PHDVB:skip tags added to any of the VMs included in the backup job. For additional information on skipping VMs, see ["Skipping VMs"](#) (on page 78).
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.

Chapter 4 - The Backup Wizard



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

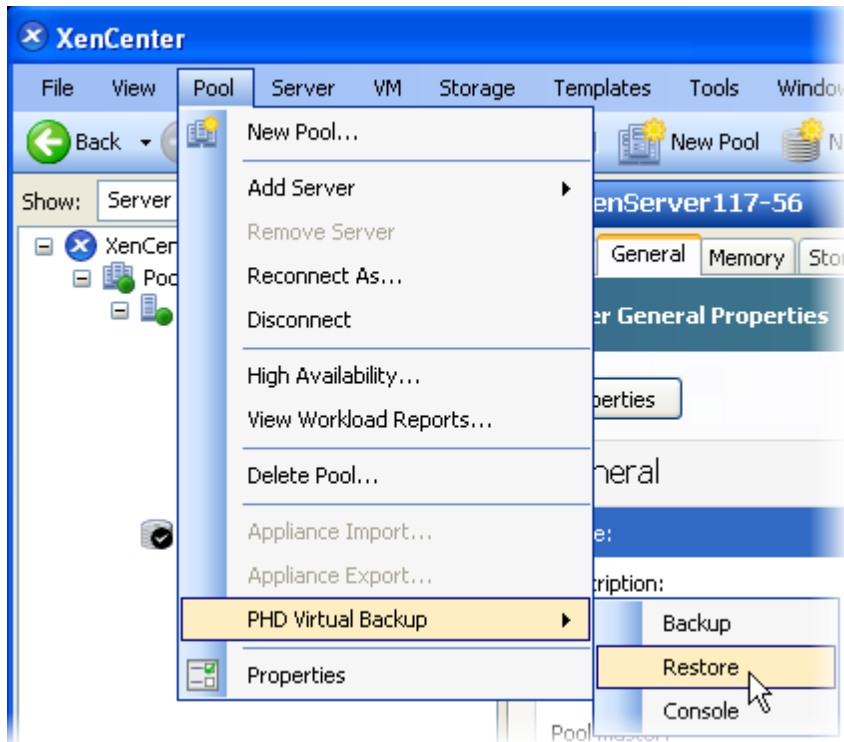
For more information on using the Jobs area of the console, see "Jobs" (on page 34)

Chapter 5 - The Restore Wizard

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

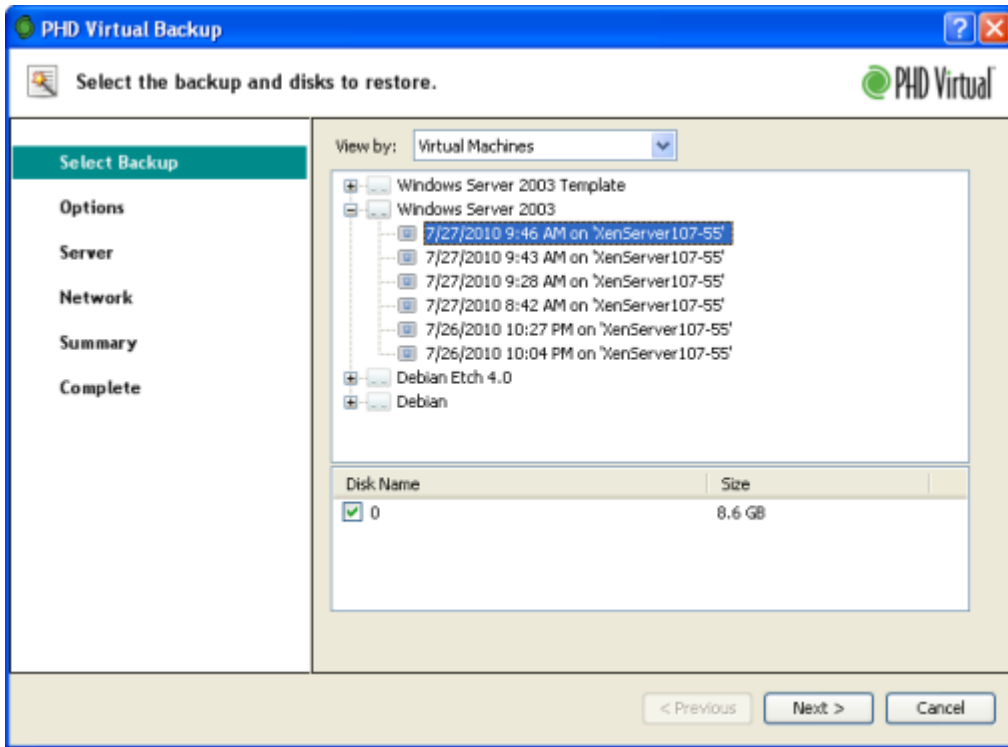
To launch the Restore Wizard

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



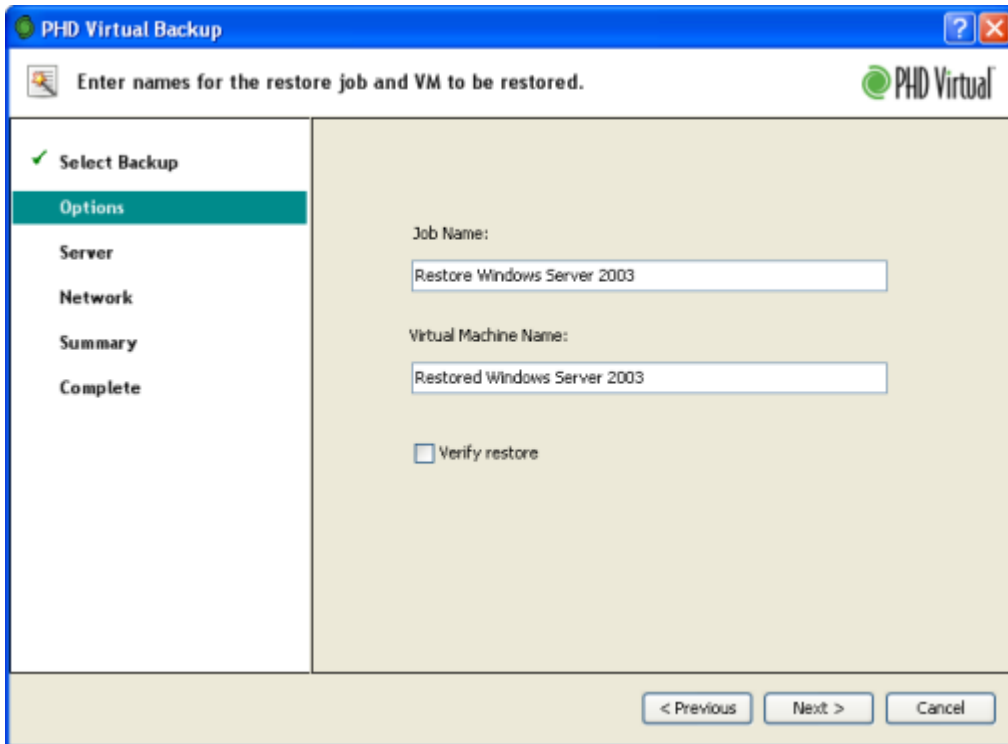
Using the Restore Wizard

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



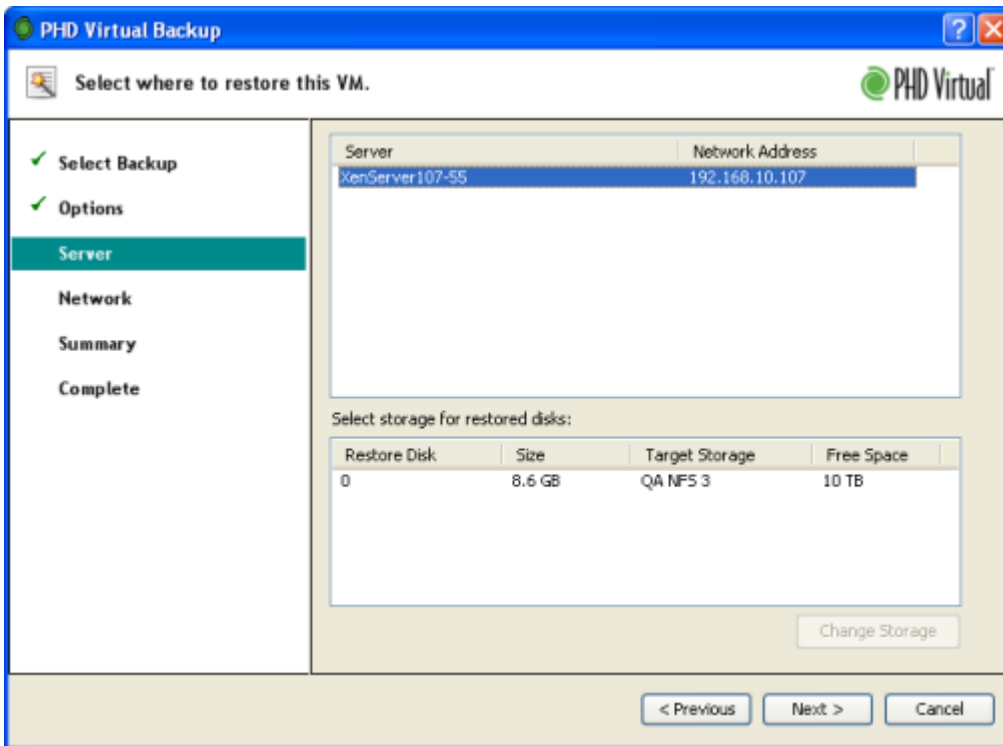
When you select the backup, the available disks are also displayed, as seen in the image above. All available disks are selected for restore by default. To exclude a particular disk from the restore, clear the check box in the **Disk Name** column.

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 75) .
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**.

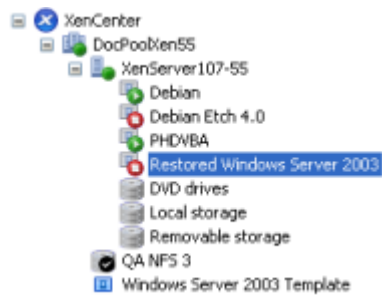
7. Select the network device to use for the restored VM. If you need to change any of the settings, click **Edit**.
8. Click **Next**.
9. Review the summary information for the restore job and click **Submit**.
10. Click **Finish** to close the wizard.

Use the Jobs area of the PHD Virtual Backup Console to view the progress of the restore job. While in progress, the VM being restored is displayed with an ellipsis before its name.

Chapter 5 - The Restore Wizard



When the restore is complete, the VM is available within XenCenter .



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.....	64
Running a Backup Now.....	65
Scheduling Backups.....	67
Viewing Jobs.....	69
Restoring Backups.....	71
Restoring Files.....	72
Configuring Email Alerts.....	74
Verifying Backups and Restores with TrueRestore™.....	75
Backup Retention and Archiving.....	76
Excluding VMs and Disks.....	77
Skipping VMs.....	78
Sending Backup Files to Tape.....	80
Using Tags to Backup VMs.....	81
Increasing Backup Storage (Attached Disk).....	82
XenServer System Logging.....	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 Pool, Folders, or Tags. When you create a job using a Pool, Folder, or Tag, VMs added to the pool or folder or assigned the tag will be included in the job automatically. Likewise, if you remove a VM from a pool or folder or remove the tag from a VM, it will not be included in the backup job. For instance, if you create a scheduled daily job that backs up all of the VMs in your production pool, any new VMs added to that pool in the future will also be included in the back up the next time the job runs.

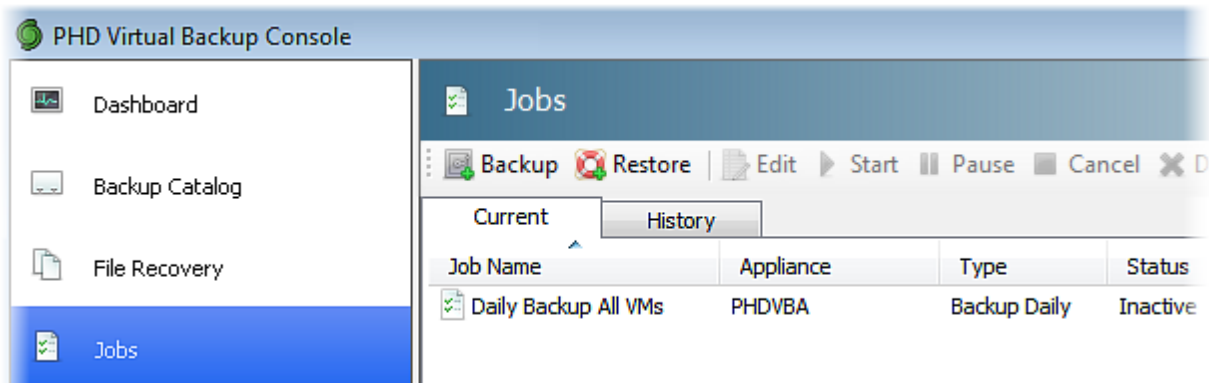
Backup jobs are created using the Backup Wizard, which is launched when you select **PHD Virtual Backup > Backup** from the integrated PHD Virtual Backup menus in XenCenter, or when you click  **Backup** within the PHD Virtual Backup Console.


To create a Backup Job

1. Launch the Backup Wizard by right-clicking a VM and selecting **PHD Virtual Backup > Backup**.
2. Follow the steps in the wizard to select VMs for backup and define a backup schedule. For detailed information about each step in the wizard, see ["The Backup Wizard"](#) (on page 53).

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.



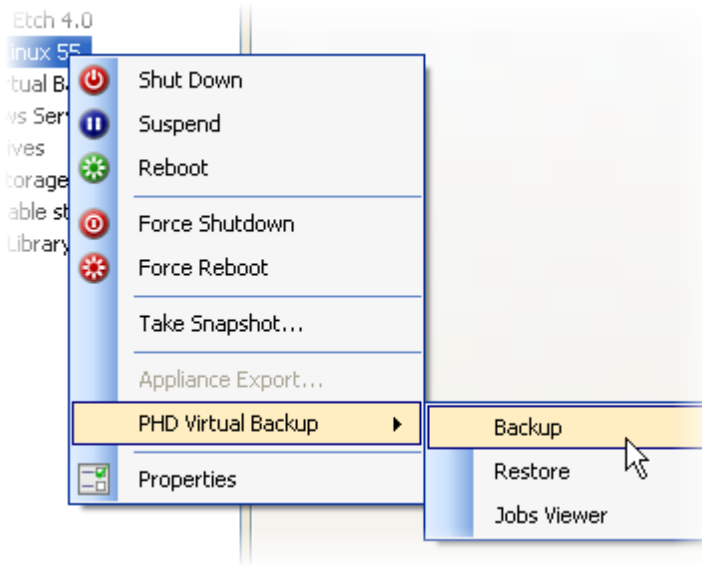
2. Select the job you would like to edit and click  **Edit**.
3. The Backup Wizard opens with the settings you originally defined for the job. Use the wizard to make any edits and submit the job again. For details on each step of the wizard, see ["The Backup Wizard"](#) (on page 53).

Running a Backup Now

There are multiple ways to run a backup with PHD Virtual Backup - the easiest is to right-click the VM you want to back up 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 XenCenter, 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 53)

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 XenCenter, 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 53).
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 34) for more information.

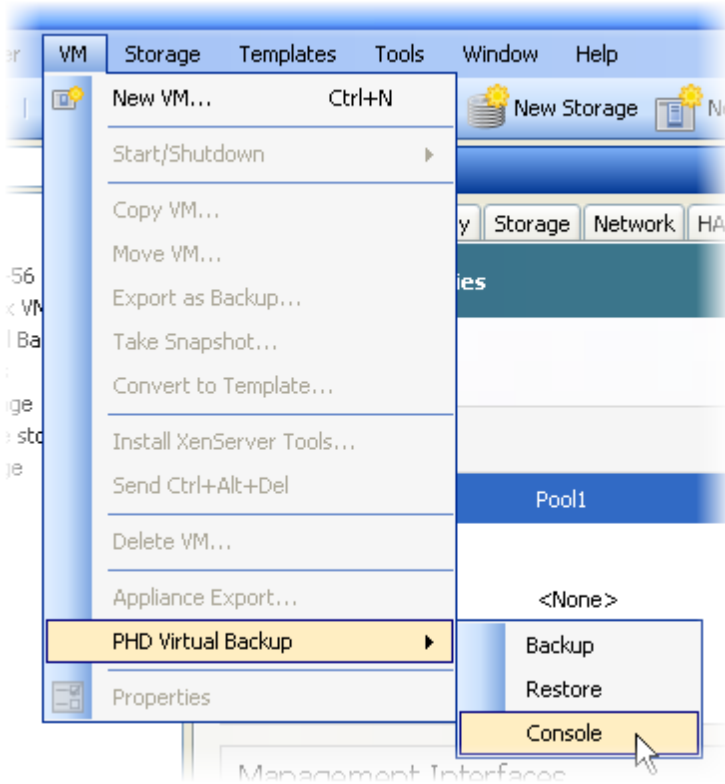
Note: If the PHD Virtual Backup Appliance is restarted within one hour of a scheduled Daily or scheduled Weekly job's start time, the scheduled job will be run again.

Viewing Jobs

To view the backup and restore jobs that are in progress or that have been created, 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 XenCenter.

To launch the PHD Virtual Backup Console


1. From XenCenter, expand the **Pool**, **Server** or **VM** menu and select **PHD Virtual Backup > Console**.



Alternatively, from within XenCenter you can right-click any VM, Server, or Pool object 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%

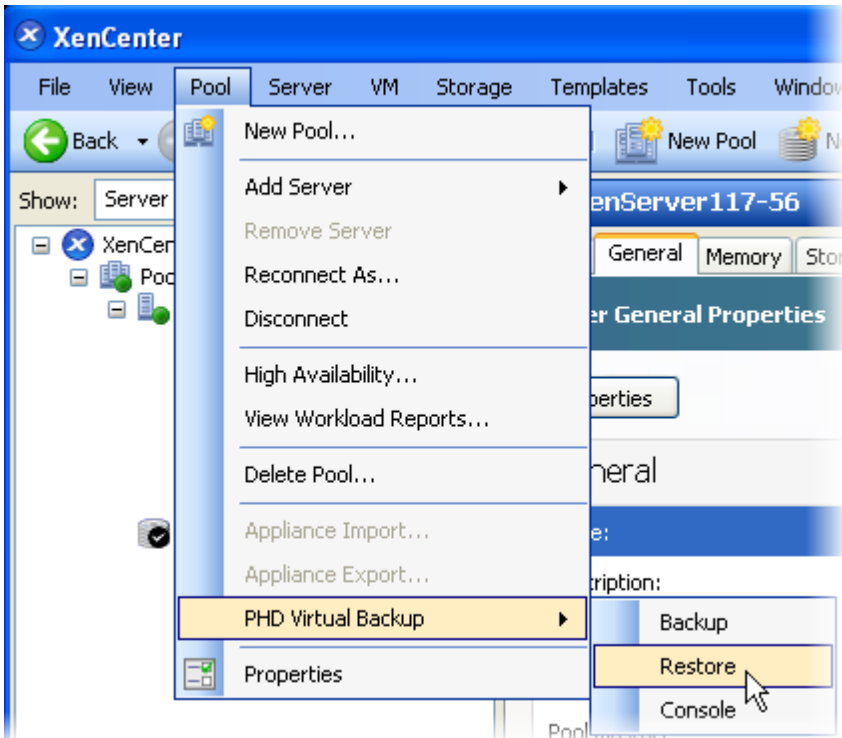
Details Tasks

Restoring Backups

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

To restore a Virtual Machine

1. From within XenCenter, select either the Pool, Server, or VM menu item and select **PHD Virtual Backup > Restore**.



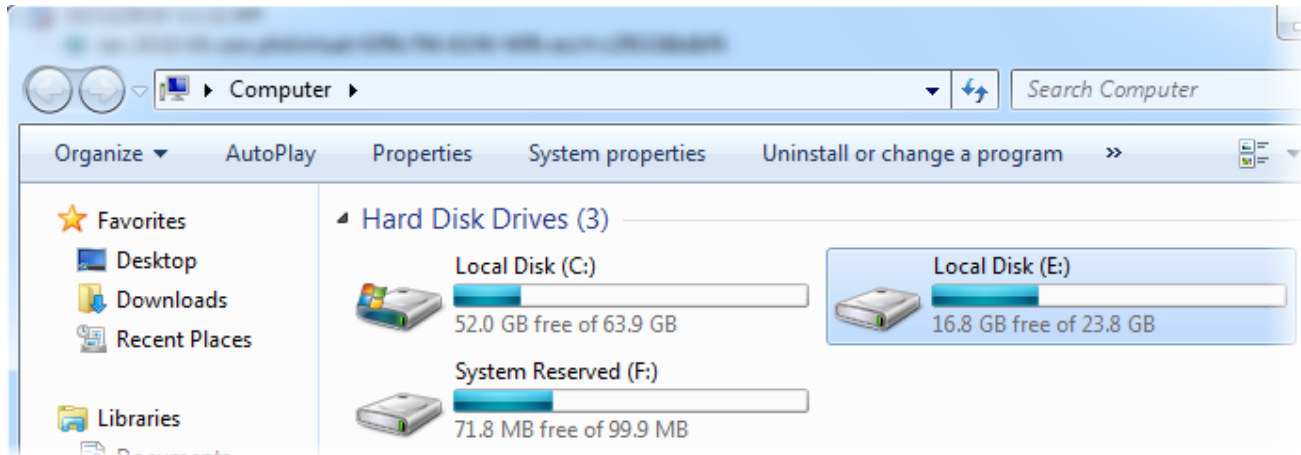
Alternatively, you can right-click directly on a VM, Pool, or Server name and select **Restore** from the PHD Virtual Backup menu. If a backup for that VM is available, the VM is pre-selected within the Restore Wizard catalog.

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 59)


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" \(on page 29\)](#).

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, you can rely on XenServer, if you are licensed for and have enabled email alerting, or if you're using the free version of XenServer and would like to receive alerts directly from PHD Virtual Backup, you can use the PHD Virtual Backup Console's Configuration page to configure email alerting.

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 are several tabs: "General", "Storage", "Network", "Email", "Retention", "Connector", and "Support". The "Email" tab is active. The configuration options are as follows:

- Radio buttons for "Do not email alerts from the appliance" (unselected) and "Email alerts using the following information" (selected).
- Text input for "Server Name" containing "smtp.example.com" and a "Port" input containing "587".
- Dropdown menu for "Security" set to "None".
- Checked checkbox for "Server requires credentials".
- Text input for "User name" containing "PHD".
- Text input for "Password" containing masked characters "••••••••".
- Text input for "From Email Address" containing "phdvh@phdvh.com".
- Dropdown menu for "Alert Level" set to "All".
- A "Recipients:" section with a text input containing "phd@phdvh.com" and "Add" and "Remove" buttons.
- A "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 configuration option, see "Email" (on page 45).

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 and after each backup is complete, 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 47)


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 34) for details.

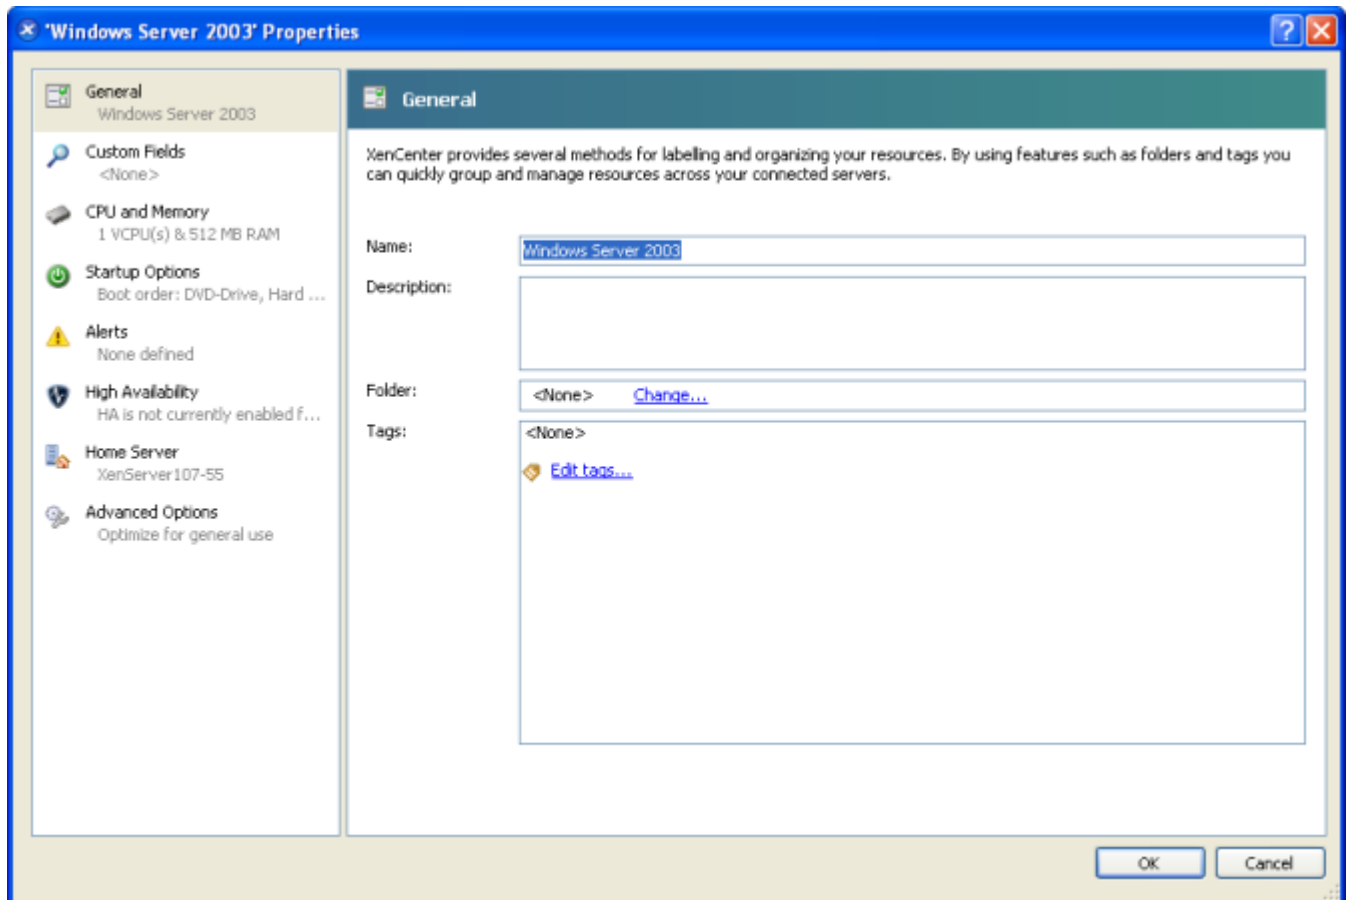
Skipping VMs

If you create a backup job that includes multiple VMs but would like to exclude some of those VMs from the scheduled backups, rather than editing the job, you can use a Tag within XenCenter, PHDVB:skip.

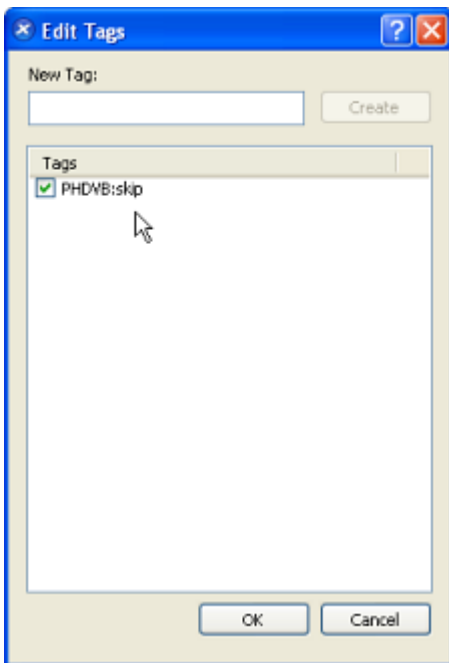
When you would like to begin backing up the VM again, you can simply remove the tag the same way it was applied.

To skip a VM with a tag in XenCenter

1. Within XenCenter, right-click the VM you would like to skip and select **Properties**. The properties window for that VM opens.



2. In the Tags area, click **Edit tags...** The Edit Tags dialog opens.
3. In the New Tag dialog box, type PHDVB:skip and click **Create**. The tag is added to the Tags dialog.



4. Click **OK**.

The next time you run a backup that includes the VM, it will be skipped, unless **ignore virtual machine skips** is selected during the backup wizard.

Note: In addition to the VM, the skip tag can be added to individual virtual disks. When applied at the virtual disk level, though, the disks are always skipped regardless if **ignore virtual machine skips** is selected during the backup wizard.


Skipping Swap File Disks

If you place your swap file data for each virtual machine on a separate virtual disk, you can use the PHDVB:skip tag to skip this disk each time the VM is backed up. Skipping the virtual disks that contain swap file data is a good idea because you do not need to save this data. Also, backup compression and deduplication ratios for your VMs will improve because swap file disks contain effectively random data that does not compress or deduplicate well.

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 50).


Using Tags to Backup VMs

Using the Backup Wizard, you can select VMs to backup individually, or you can select entire Pools, Folders, or Tags. When **View by: Objects by Tag** is selected, tagged VMs are displayed, grouped by the tag applied.



Tags can be applied to the Properties of each VM individually, within XenCenter.

When a backup job is created for a specific tag, any VMs that have that tag applied will be included in the backup job. Also, any VMs that have the tag applied in the future will be also included in the job, automatically. There is no need to edit the job each time a new VM is added - you can simply apply the tag using XenCenter when the VM is deployed and it will be backed up the next time the scheduled backup job runs.



This same concept, dynamically backing up containers, applies to XenServer Pools and Folders. Creating a scheduled backup job for a specific pool or folder will backup all of the VMs that belong to that folder or pool. Adding or removing VMs from a pool or folder will cause them to be backed up or excluded during the backup, respectively.

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 XenCenter, right-click the PHD Virtual Backup Appliance and select Shut Down.
2. When the appliance is powered off, click the **Storage** tab, then select the virtual disk used for storage, PHD Virtual Backup Store, for example.
3. Click **Properties**.
4. Click **Size and Location** and enter a new size for your backup storage.
5. Click **OK** to close the Properties dialog and start the appliance. The new size will be reflected in the PHD Virtual Backup Console's Dashboard.

XenServer System Logging

The PHD Virtual Appliance, by default, sends all of its log messages to the system log server configured on the Pool Master XenServer. If you already have remote system logging configured, once the appliance is deployed, you will begin to receive log messages.

If you enable system logging after the appliance has been deployed, you will need to restart the appliance before the log messages will be sent to the remote server. The PHD Virtual Backup Appliance uses the pool master settings to determine where to send the system logs

Updating PHD Virtual Backup

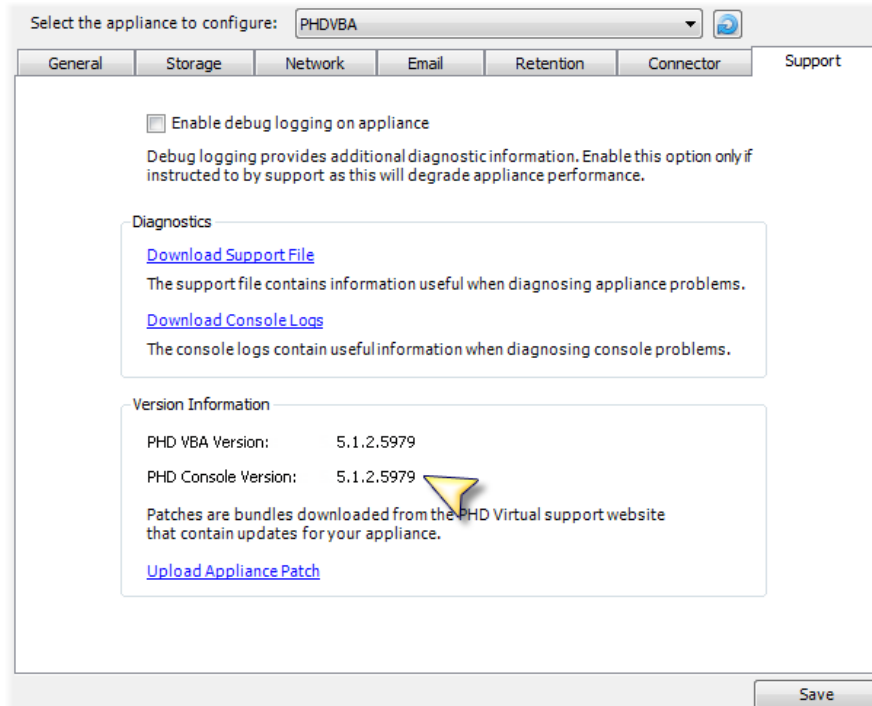
When available, updates to PHD Virtual Backup can be downloaded from the PHD Virtual Web site or obtained from Support. Console and Plug-in updates are made available via an updated MSI file and PHD Virtual Backup Appliance updates are available as update files (.phd files).

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.

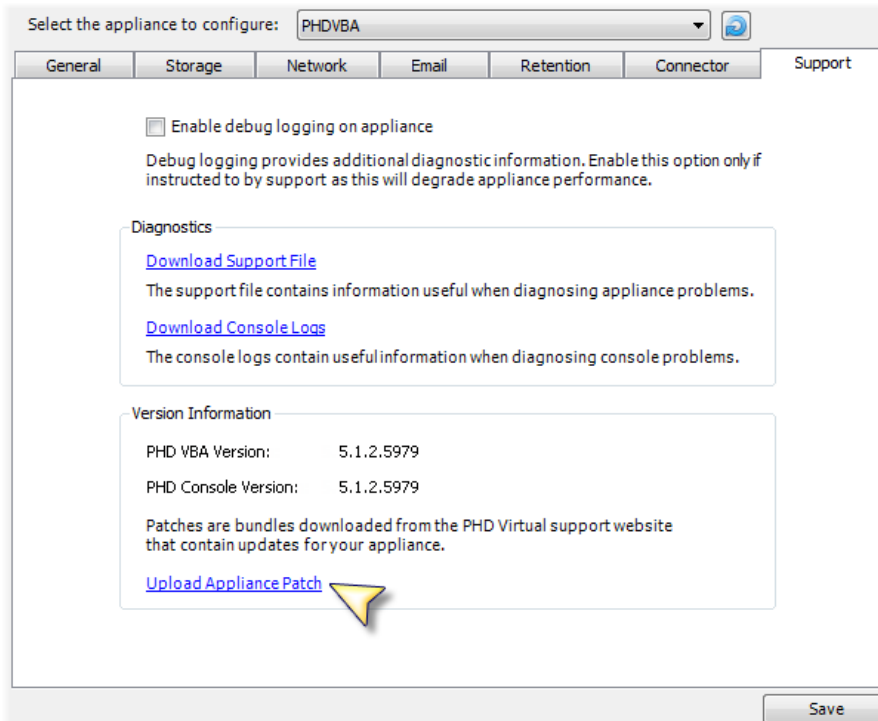


To update the PHD Virtual Backup Appliance

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

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 the configuration interface 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 sections:

- 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: 5.1.0.4203 (for Citrix XenServer)
 - PHD Console Version: 5.1.0.4211
 - Patches are bundles downloaded from the PHD Virtual support website that contain updates for your appliance.
 - [Upload Appliance Patch](#)

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

For additional details about the Support tab, see "Support" (on page 52).

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 XenCenter and select the problematic PHD Virtual Backup Appliance. If running, power off the appliance (right-click and select **Shut Down**).
2. Click the **Storage** tab, then select the virtual disk used to store your backups, for example, PHD Virtual Backup Store.
3. With the disk selected, click **Detach**.
4. Deploy a new appliance. Use the XVA that came with your initial installation package. Follow the steps in the installation guide for details.
5. Before powering on the new appliance, instead of creating a new virtual disk for backup storage, use the virtual disk you detached earlier. With the new appliance selected, click the **Storage** tab.
6. Click **Attach...**
7. From the list of available disks, select the PHD Virtual Backup Store you detached earlier.
8. Power on the appliance. You can begin backing up and restoring VMs once again and all backups stored previously should be available in the backup catalog.

To recover backups if using CIFS or NFS share

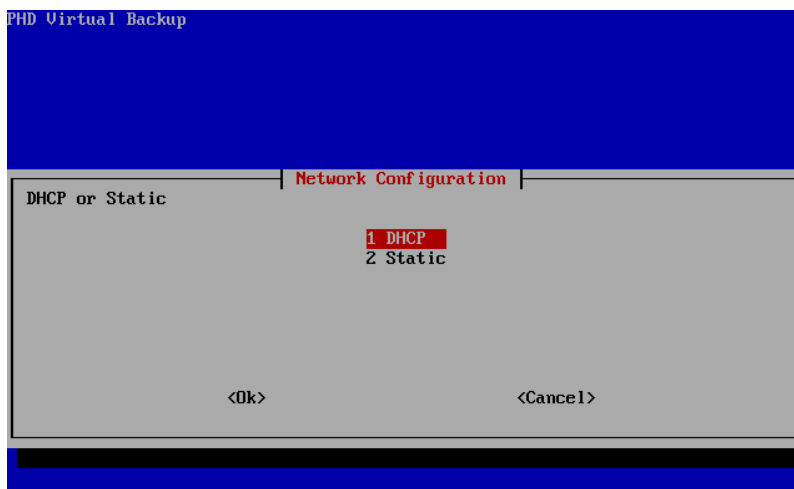
1. Power off the problematic appliance within XenCenter.
2. Deploy a new appliance. Use the XVA that came with your installation package. Follow the steps in the installation guide for details, making sure to select **CIFS** or **NFS** as the storage type.
3. Click **Save**, then restart the appliance.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

To reset the PHD Virtual Backup Appliance's Network settings

1. Open XenCenter and select the PHD Virtual Backup Appliance virtual machine.
2. Click the **Console** tab.
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.

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.

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

.			
.phd files	84		
A			
Advanced storage options	42		
Alert Level	46		
All	45		
All blocks	57		
Antivirus	17		
Antivirus Software	17		
Appliance	24		
Appliance Crash	88		
Appliance options	39		
Appliance updates	84		
archive backups	50		
Archive backups	57		
Archiving Backups	76		
assign static appliance network settings	44		
Average Speed	36		
B			
Backup Appliances	24		
Backup Catalog	26		
deleting backups	76		
Backup Catalog Notes	28		
backup data			
self-healing	75		
Backup Data Connector	50, 80, 90		
troubleshooting	90		
Backup is stopped	92		
		Backup Jobs	
		creating	64
		Backup Now	65
		Backup powered off virtual machines	57
		Backup Retention	76
		Backup storage	41
		Backup Virtual Machine	35
		Backup Wizard	53, 64-65
		using	53
		Backups	12
		verify	75
		BDC	50
		Best Practices	17
		C	
		Cancel	35
		Change Storage	61
		CIFS/SMB Shares	17
		Citrix XenServer®	7
		Collection of metadata failed	91
		Configuration	38
		reload values	38
		Configuration page	38
		Connector	50
		Connector tab	50
		Console and Plug-in updates	84
		Could not attach	91
		Could not write and close backup block	91
		create a Backup Job	64
		create a scheduled backup job	67
		Creating Backup Jobs	64

Index

Critical	45	enable alerts	45
Critical errors	45	Enable compression for new backups	42
CTRL-N	89	Enable debug logging on appliance	52
Custom		Errors	45, 91
retention setting	47	Exclude	54
D		Excluding VMs	77
Daily	67	export backups	50
Dashboard	23	Exporting Backups	27
Backup Appliances list columns	24	F	
Data Streams	39	Failed to save changes	91
Data Written	36	File Recovery	29
debug mode	52	Folder	
debugging mode	52	View by option	53
Dedupe Ratio	24	Free Storage	24
DeDupe Ratio	36	Frequently Asked Questions	15
defragmentation	17	G	
Defragmenting	17	General	39
Delete	35	General tab	38-39
delete iSCSI target	33	H	
Delete trim	37, 47	Help	18
Deleting backups	27	How many appliances do I need?	19
Deleting iSCSI targets	33	How PHD Virtual Backup Works	11
disable email alerts	46	Hypervisor Credentials	38-40
Disk Defragmenter	17	I	
display Job Details	35	Ignore virtual machine skips	57
Do not start after	56, 67	Include Templates	54
Documentation Updates	4	Increasing Backup Storage	82
Dom0	7	Individual backups	
E		deleting	76
Edit	35	inf	15
edit a job	64	IP address	
Email	45	appliance	43
Email Alerts	74		

obtain automatically for appliance	44	Next Run	56, 68
IP Address	24	NFS Shares	17
iSCSI Software Initiator	73	NTP servers	39
J		O	
Job Details	35	Once	67
Jobs	34	Options	
Jobs History	37	backup wizard	56
		Orphan Weekly	37
K		P	
Keep All		Pause	35
retention setting	47	PHD Console	9
L		PHD VBA	9
Last32Days	51	PHD Virtual Backup	
Last7Days	51	benefits	10
LatestofEach	51	receiving alerts	74
launch the PHD Virtual Backup Console	69	updating	84
launch the Restore Wizard	59	PHD Virtual Backup Appliance	9
license		crash	88
update	40	reset network settings	89
upload new	40	updating	52, 85
Licensing	40	PHD Virtual Backup Components	14
M		PHD Virtual Backup Console	9, 22, 38
Management IP	40	accessing	22
Mount iSCSI target on this computer	31	updating	84
Mounting iSCSI Targets on Other Devices	33	PHD Virtual Backup Plug-in	9
MSI	84	PHD Virtual Support	52
N		PHDVB	9, 57, 78
Network	43	Plug-In	
Network Settings		updating	84
reset	89	Pool Master change	40
Network tab	43	Pool Master XenServer	40
New blocks only,	57	Product expiration	40

Index

Q

Quiesce the VM before backing up 57

R

Raw

 exporting backups as 27

Recent backups to keep 48

recover backups 88

Recurrence 36

Recurring jobs 56, 67

Rekurs every 56, 67

Renaming VMs 28

Resetting VBA Network Settings 89

restore a Virtual Machine 71

restore file

 Linux 31

Restore Virtual Machine 35

Restore Wizard 59, 71

Restores 13

 verify 75

Restoring Backups 71

Restoring Files 29, 72

Restoring Files from a Linux VM 31

Restoring Virtual Machines 27

Retention 76

Retention Settings 47

Retention tab 47

run a scheduled backup now 65

run a single backup 65

Running a Backup Now 65

S

Schedule

 backup wizard 55

Scheduling Backups 67

self-healing 75

Sending Backup Files to Tape 80

Server View

 View by option 53

Show Details 35, 70

Show system jobs 35

Show/Hide Details 35

skip 57, 78

Skipping Swap File 79

Skipping VMs 78

Start Date 56, 67

Start Time 56, 67

Start/Resume 35

Startup 36

Static IP Address 44

Stop level % free 42

Storage 41

Support 3, 52

Support expiration 40

Support Files 87

 submitting 87

Support tab 52

Swap File 79

System Alert descriptions 25

System Alerts 23, 25

System Jobs 36

System Logging 83

T

Tag

 View by option 53

Tags

 backing up VMs by 81

tape 50

Terms	9	Virtual Hard Disks	
The PHD Virtual Backup Appliance	19	exporting backups as	27
The Restore Wizard	59	VMDK	27
Total Backup Data	24	VMDK export	27
Trim	76	Volume Shadow Copy Services	57
Troubleshooting	86		
TrueRestore	75	W	
Typical		Warning level % free	42
retention setting	47	Warnings	45, 91
U		Weekly	67
update packages	84	What's New	8
Updating PHD Virtual Backup	84	X	
Upload Appliance Patch	52	Xen domain zero	7
Uploading Appliance Patches	52		
Used Storage	24		
Using PHD Virtual Backup	63		
Using the Restore Wizard	59		
UUID	26		
V			
VBA	7, 9, 14		
VBA Console	21		
Verify backup	57		
verify backups	75		
Verify Restore	61		
verify restores	75		
Verifying Backups and Restores	75		
VHD	27		
Video Tutorials	18		
View by	53		
View Log	35		
Viewing Jobs	69		
Virtual Backup Appliance	7		