



esXpress 4.0
DeDupe Appliance
Users Guide

Software Release Date: March 2010

Document Date: July 30, 2010

www.phdvirtual.com



Legal Notices

Copyright © 2005-2010 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.

Microsoft®, Windows® and MS-DOS are registered trademarks of Microsoft Corporation.

VMware®, VMware® ESX Server, VMware® VMotion™, VMware® VirtualCenter, vCenter™, and vSphere™ are registered trademarks of VMware Corporation.

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

Support, Sales, Renewals, and Licensing

For information on new sales, licensing and support renewals you can email sales@phdvirtual.com or info@phdvirtual.com.

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

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

For customers and partners with an active support agreement, you can use the esXpress 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 esXpress release and software updates.

Documentation Updates

The following table contains updates made to this document. To verify that you are using the latest edition of a document, check the PHD Virtual website.

Chapter	Version	Description
All	4.0	Formatting and editing.
1	4.0	"DeDupe Appliance Passwords" (on page 9): Added new section. 4/12/10
2	3.6	Added note about adding multiple DNS servers.
2	3.6	"DeDupe Appliance System Requirements" (on page 11): added note about DeDupe configuration for optimal performance.
2	3.6	"VMDK Attached Storage" (on page 20): added caution about /.data directory.
2	4.0	"Manually Updating the DeDupe Appliance" (on page 18): added new section. 5/27/10
3	4.0	Added "System Page" (on page 27) information (was previously the Summary Page).
3	3.6-4	"Stats Page" (on page 34): updated to reflect new functionality.
3	3.6-8	"System Options" (on page 41): added options for automatic check for updates process.
3	3.6-10	" Auto Delete - DeDupe Backups" (on page 37): added THIN_SPECIAL.
3	3.6-12	" Auto Delete - DeDupe Backups" (on page 37): added KEEP_MINIMUM_UNIQUE, and KEEP_MAXIMUM_WEEKS. Removed old purge options.
5	4.0	"Troubleshooting" (on page 53): added new chapter for troubleshooting information. 4/12/10

Contents

Chapter 1 - The esXpress DeDupe Appliance	7
About This Guide.....	7
DeDupe Backup Speeds.....	7
Date Smart Dynamic Export.....	8
DeDupe Appliance Passwords.....	9
Chapter 2 - Installing and Configuring the DeDupe Appliance	10
DeDupe Appliance System Requirements.....	11
Pre-Installation Steps.....	12
Installing the DeDupe Appliance.....	13
Assigning Storage Space for VMFS Backups.....	14
Configuring the DeDupe Appliance.....	15
Checking for DeDupe Appliance Updates.....	17
Manually Updating the DeDupe Appliance.....	18
Configuring Storage for the DeDupe Appliance.....	19
NFS External Storage.....	19
SMB External Storage.....	19
VMDK Attached Storage.....	20
Remaking Block Folders.....	20
esXpress Backup Target Configuration.....	21
Appliance Maintenance.....	22
Extending Appliance Disks.....	22
Appliance Logs.....	22
Backing Up the Attached VMDK.....	23
DeDupe Appliance Text Menu Options.....	24
Chapter 3 - DeDupe Appliance Administration	25
Accessing the DeDupe Appliance Web Page.....	26
System Page.....	27
Maintenance Mode.....	27
DeDupe Appliance Updates.....	28
Storage.....	28
VMs Page.....	29
Managing Backups.....	30

FLR (File Level Restores) Page	31
Stats Page	34
Trim Report	34
DeDupe Stats	35
Backup Counts	36
Other Counts	36
Options Page	37
Auto Delete - DeDupe Backups	37
Planning for Auto Deletes/Trim	40
Statistics Options	40
System Options	41
Users Page	43
Logs Page	44
Exporting Backups	45
Chapter 4 - Implementation Scenarios	47
Local VMFS Storage	48
Remote NFS Storage	50
Chapter 5 - Troubleshooting	53
What Happens If My DeDupe Appliance Crashes?	54
Index	55

Chapter 1 - The esXpress DeDupe Appliance

DeDuplication technology (DeDupe) enables the creation of Virtual Full Backups, which capture unique data blocks across all virtual machines. Data DeDuplication will eventually replace the creation of separate full and delta backups for virtual machines.

PHD Virtual's esXpress DeDupe Appliance aims to change the prevalent Virtual Machine backup paradigm by implementing cutting edge DeDuplication technology to provide immediate and visible cost savings to IT budgets by dramatically reducing disk space requirements. Reduction of network traffic will also result in an improvement of backup speeds and reduce stress on the existing network infrastructure.

esXpress DeDupe technology creates backups of VMDKs (virtual disks) using the Virtual Backup Appliance, but, different than traditional Full and Delta backups, a Virtual Full backup is created.

During the first backup, the entire virtual disk is compressed and stored on the DeDupe Appliance. Subsequent backups of any VM whose targets are the same pre-existing Virtual Backups managed by the DeDupe Appliance provide significant storage savings. Data that already exists on the DeDupe Appliance target that matches data on the VM being backed up, is not written to the appliance. This ensures only unique data blocks are stored, creating Virtual Full backups.

Additional storage can easily be added to the DeDupe Appliance to accommodate varied usage scenarios. Storage can be in the form of an attached VMDK, an SMB share or an NFS share.

PHD Virtual's DeDupe feature requires the installation of the newest version of the DeDupe Appliance. The appliance will run on the existing ESX infrastructure and become a backup target for esXpress.

About This Guide

This guide describes how to install and configure the esXpress DeDupe Appliance including how to configure the storage the appliance will use to store your backups.

In addition to installation and configuration information, this guide describes how to use the features available from the DeDupe Appliance web page, including how to do DeDupe File Level Restores and how to view the statistics and summary information about all DeDupe backups.

DeDupe Backup Speeds

The initial DeDupe backups of your Virtual Machines may be slow as the DeDupe Appliance is seeded with the new backup blocks. In addition, each block is verified after writing to ensure they are valid. This validation takes a little extra time but is important in ensuring proper backups.

In testing, the initial backup of a VM may only get around 5 to 8 MB/sec. However, the more backups that are run, the DeDupe technology will start to show benefits as much less blocks are written to the appliance. Speeds can range from 25 to 40 MB/sec to upwards of 75 MB/sec depending on the VM size, the number of blocks being written, and the underlying infrastructure.

The following example is an actual test case of a DeDupe backup.

In the initial DeDupe backup for a 55GB VMDK, almost the entire disk was written to the DeDupe Appliance, 53 GB, which is 97% of the data. With compression, 25 GB of data was sent to the target. This backup took 3 hours to complete and averaged 5 MB/sec.

Initial Backup

2009-05-22 23:28:18.387a PHDD Blocks: 56331 Total, **54761 Wrote**, 14 Skipped, 1556 Zero for: test1, 01-test1_1.vmdk

2009-05-22 23:28:18.391a PHDD Stats : **25866 sent meg (47%/45.92%) / 54761 written meg (97.21%)** / 56331 total meg for: test1, 01-test1_1.vmdk

2009-05-22 23:28:22.469a OK PHDD 'test1' - '01-test1_1.vmdk' - 2/2 disks, (97%) 53g/55g (97% Data), **3:00:10s, 5mb/s** (17gb/hr) NET1

Next, let's look at a DeDupe backup of this same VM run a few days later. Because of the DeDupe, only 294 blocks (81 MB compressed) were written to the DeDupe Appliance. The backup of the 55 GB VM completed in just 12 minutes and 28 seconds and averaged 75 MB/sec.

Subsequent Backup

2009-05-26 19:59:53.129d PHDD Blocks: 56331 Total, 294 Wrote, 54481 Skipped, 1556 Zero for: test1, 01-test1_1.vmdk

2009-05-26 19:59:53.136d PHDD Stats : 81 sent meg (28%/0.14%) / 294 written meg (0.52%) / 56331 total meg for: test1, 01-test1_1.vmdk

2009-05-26 19:59:59.193d OK PHDD 'test1' - '01-test1_1.vmdk' - 2/2 disks, (0.5%) 294m/55g (97% Data), 12:28s, 75mb/s (263gb/hr) NET1

The above example shows the considerable time and storage savings you can expect from esXpress DeDupe backups - the more you backup, the greater the benefit.

Date Smart Dynamic Export

esXpress DeDupe backups can be exported to save to tape or archive on disk. Date Smart Dynamic Export allows you to select the exact backups you want to export as uncompressed VMDKs, when you want. Data is selected using a dynamically generated folder structure on the DeDupe Appliance that allows you to select backups by day, week, month, or even the latest backup stored for every VM available, and more.

Selecting only the specific data you want allows you to save time transferring data and to save money on storage and tapes, which helps to increase the overall ROI for DR and data protection.

For additional details, see ["Exporting Backups" \(on page 45\)](#).

DeDupe Appliance Passwords

The DeDupe Appliance web page and text menu interfaces are shipped with the following default passwords:

DeDupe Appliance web page

User	Default Password
admin	admin

To change the DeDupe Appliance web page default password, use the Users page. See ["Users Page"](#) (on page 43).

DeDupe Appliance console

User	Default Password
root	esxpress

To change the console login default password, use the console text menu. See ["DeDupe Appliance Text Menu Options"](#) (on page 24).

Chapter 2 - Installing and Configuring the DeDupe Appliance

Review the system requirements and pre-installation information in this section before installing and configuring the esXpress DeDupe Appliance.

Note You can deploy the esXpress appliances using the esXpress deployment wizard (recommended). If you deployed the appliance using this method, you do not need to follow the configuration steps in this section.

- "DeDupe Appliance System Requirements" (on page 11)
- "Pre-Installation Steps" (on page 12)
- "Installing the DeDupe Appliance" (on page 13)
- "Configuring the DeDupe Appliance" (on page 15)
- "Configuring Storage for the DeDupe Appliance" (on page 19)
- "Appliance Maintenance" (on page 22)
- "DeDupe Appliance Text Menu Options" (on page 24)

DeDupe Appliance System Requirements

The following table shows the minimum system requirements for running the DeDupe Appliance.

Table 1 - System Requirements

VMware version	ESX Server Version 3.5x, vSphere (ESX4)
VM Memory	1GB of memory allocated to each VM
Storage Requirements	<p>A Minimum of 8GB of disk space is required on the VMFS partition of the ESX host to import the compressed DeDupe VMDK file. It is recommended to expand the system disk to 16GB when you initially set up the appliance.</p> <p>Storage space for backups must be available either on VMFS partition, an NFS share, or an SMB share.</p> <p>The space requirement depends on the actual data that will be sent to the appliance.</p> <p>A good rule of thumb is to assume you will get a 10:1 compression ratio. So calculate storage by dividing the total size of your VMDKs by 10, and then multiply by the number of backups you plan to keep stored on disk.</p>
Supported Browsers	Internet Explorer 7, Firefox 3.0.7
Other Requirements	VMware Virtual Infrastructure Client 3 (VI3), or vSphere Client.

Note: The DeDupe Appliance is designed to work optimally on your best hardware and best storage equipment. If you attach older or slower hardware to the DeDup appliance you may experience performance degradation on backups.

The suggested optimal performance configuration for DeDupe is to use an NFS or SMB share with by-pass mode.

Using the attached VMDK storage or using NFS or SMB without by-pass mode is not preferred and may result in performance degradation.

Note: To secure connections between the VBA and the DeDupe datastore, an automatically generated password is used. These passwords are used once and cannot be reused.

Pre-Installation Steps

Please ensure the following pre-requisites are met prior to installing the DeDupe Appliance.

1. Download the latest esXpress zip file from PHD Virtual's web site (www.phdvirtual.com).
2. The following files within the zip will be imported to the ESX server:
 - a. phd-x.x-x.ovf
 - b. phd-x.x-x-disk1.vmdk

(where x.x-x represents the latest version of the file)
These files can be imported by downloading them to a local machine and then imported to the ESX Server using the VI3 or vSphere client.
3. The client must be logged on to your ESX host.
4. If an NFS share is to be used for storing backup data, please verify that NFS service is running on the NFS server and there is sufficient storage space available.
5. If VMFS storage is to be used for storing backup data, after the installation of the appliance please create a new virtual hard drive and assign it to the imported Virtual Machine. Also, make sure sufficient storage space is available.
6. Before beginning with the installation, you must have the following network information ready:
 - a. an IP Address for the DeDupe Appliance
 - b. DNS
 - c. Default Gateway
 - d. Proxy Server information, if required
7. If the Host is a Dual CPU host, make sure to allocate only 1 virtual CPU to the appliance VM.
8. Ensure that esXpress is configured to use PHDD as a backup target.

After you've reviewed the system requirements and pre-installation information, install and configure the DeDupe Appliance. The following sections explain the required steps:

- "Installing the DeDupe Appliance" (on page 13)
- "Assigning Storage Space for VMFS Backups" (on page 14)
- "Configuring the DeDupe Appliance" (on page 15)

Installing the DeDupe Appliance

The following sections describe how to install the DeDupe Appliance using a vSphere or VI3 client. Be sure to follow the steps for the VMware client you are using.

To install the DeDupe Appliance

1. Unzip the esXpress zip file.
2. Using a VMware client, log in to your ESX host.
3. Import the appliance OVF:

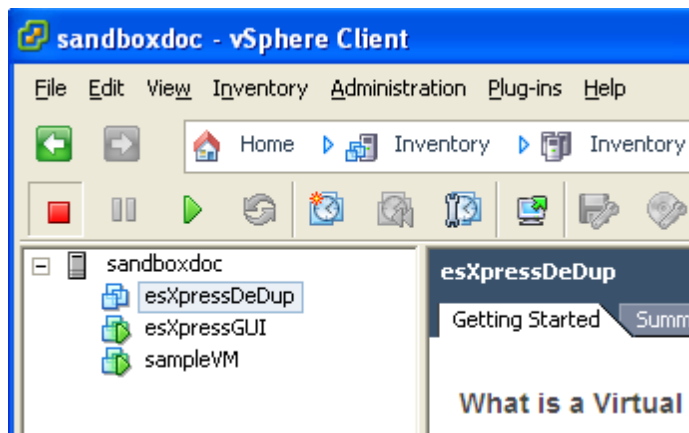
If you are using a VI3 client, select **File > Virtual Appliance > Import**

If you are using a vSphere client, select **File > Deploy OVF Template...**

4. Browse to the location you stored the extracted OVF, select the file, and click **Next**.
5. Follow the remaining steps to import the file and accept the End User License Agreement.

When you are finished, a progress indicator opens and the appliance is added to your ESX host.

6. When the process completes, click **Close** to exit the progress indicator window.
7. The new appliance is displayed in the client window.



8. Next, define the following recommended settings for the DeDupe Appliance VM.
 - Minimum Memory: 2 GB.
 - Hard Disk Capacity: 16 GB.

To adjust these values, right-click the new DeDupe VM and select **Edit Settings**. Note that if you are planning to use VMDK storage for backup data, a second hard disk must be added to the VM. For details, see "Assigning Storage Space for VMFS Backups" (on page 14).

When you are finished installing the appliance, follow the steps in the next sections to complete the required post-installation configuration tasks:

- "Configuring the DeDupe Appliance" (on page 15)
- "Configuring Storage for the DeDupe Appliance" (on page 19)

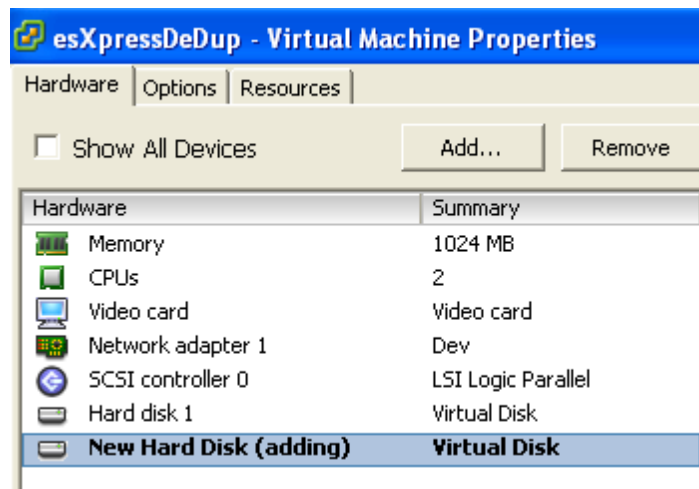
If you will be using an attached VMDK as the storage location, see, "[Assigning Storage Space for VMFS Backups](#)" (on page 14).

Assigning Storage Space for VMFS Backups

Only one type of backup storage can be assigned per DeDup Appliance. If you are going to use VMFS storage for backup data, you must create a new virtual hard drive for the imported DeDup Appliance VM. This is not required if you are using an external NFS or SMB share. Follow the steps below to add an additional virtual drive to the appliance.

To assign storage space

1. In the VMware client, right-click the newly added VM and select **Edit Settings**. Note: the VM must be powered off before you can change hardware settings.
2. Click **Add**.
3. Select **Hard Disk** and click **Next**.
4. Select **Create a new virtual disk** and click **Next**.
5. Enter the **Disk Capacity** and choose a location then click **Next**. The disk capacity is the space required for subsequent backups. Recommended capacity is a minimum of 16 GB.
6. Select the **SCSI ID**. This can be either left as default or customized. Click **Next**.
7. Review the summary information and click **Finish**. The new hard disk is added.



8. Click **OK** to complete the hard disk installation.

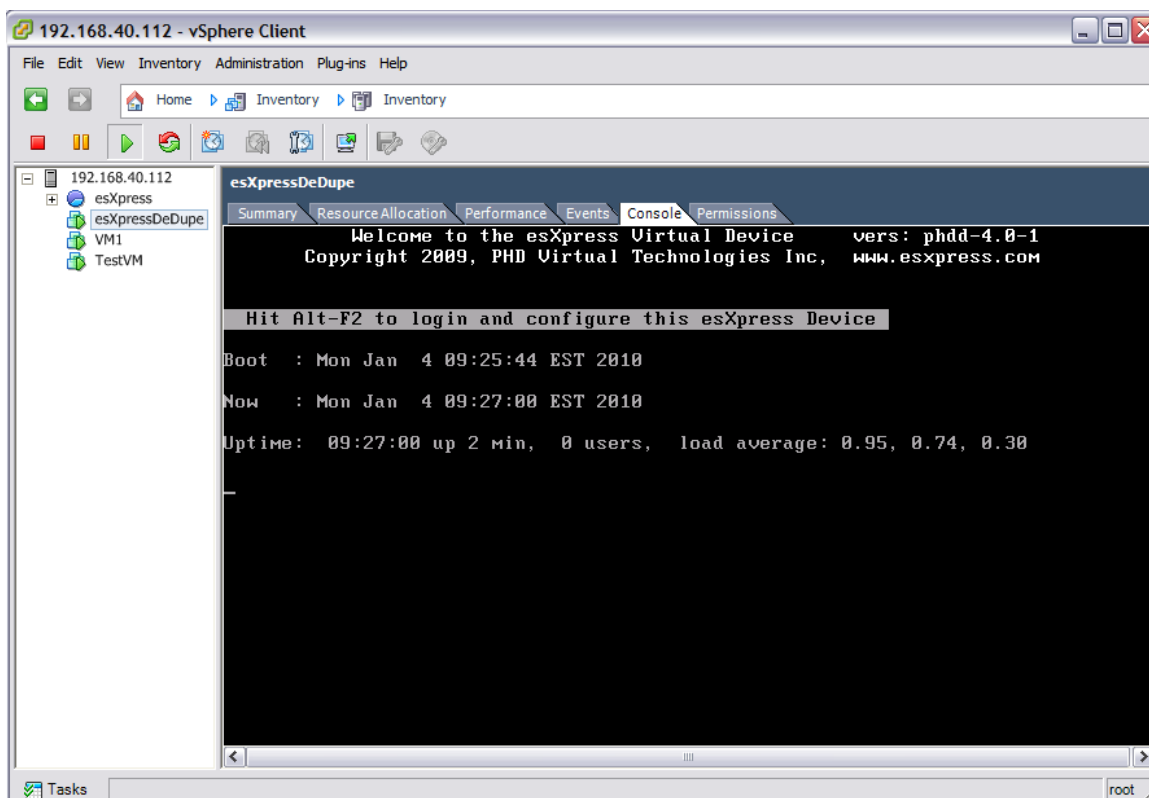
When added, continue with the appliance configuration in the section, "[Configuring the DeDup Appliance](#)" (on page 15). and finally, configure the appliance to use the newly added attached VMDK. See "[VMDK Attached Storage](#)" (on page 20).

Configuring the DeDupe Appliance

After it is deployed, the DeDupe Appliance must be configured. This section details the steps required to complete the DeDupe Appliance configuration.

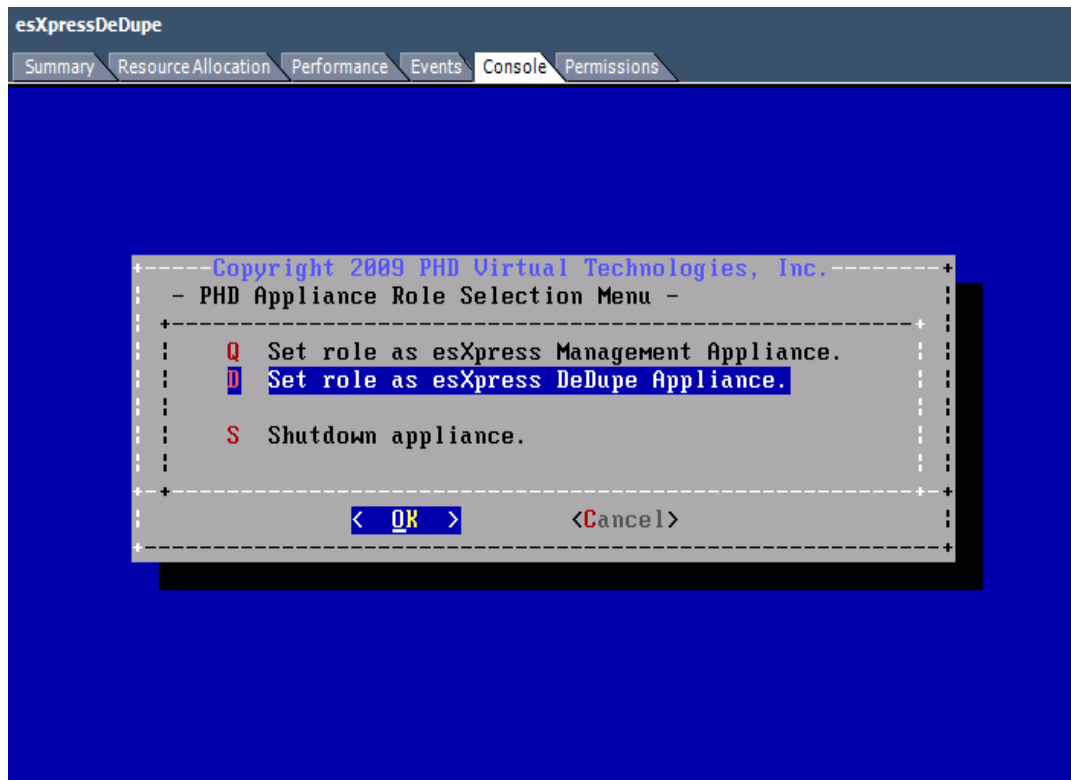
To configure the DeDupe Appliance

1. From the VMware client, start the newly installed DeDupe Appliance virtual machine.
2. Open a console for the VM (click the **Console** tab or right-click the VM name and select **Open Console**). To display status messages during the loading process, hit Alt+F2. When the machine is finished booting, you will see the main menu for the appliance.



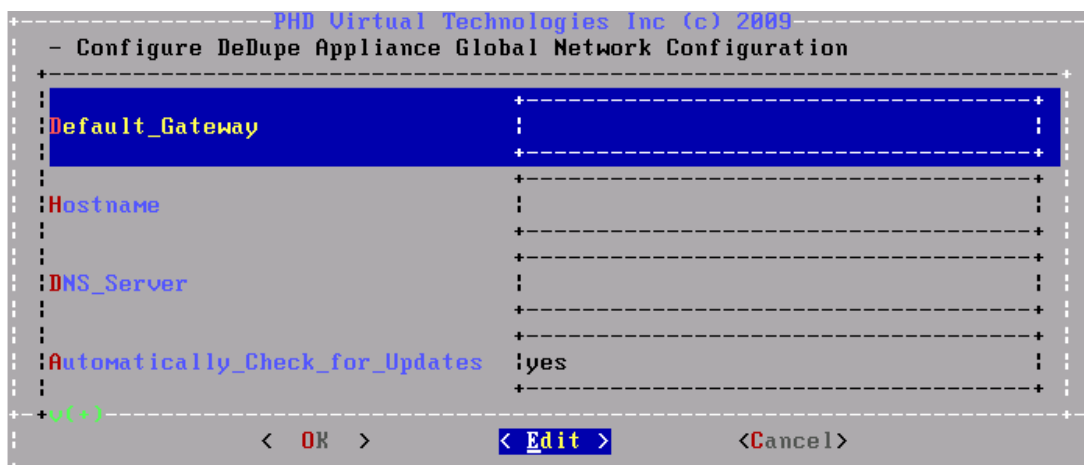
3. Hit **Alt+F2** to enter the console login prompt.
4. Login as **root** with password **esxpress**.
5. At the command prompt, type **esxpress** and hit **Enter**.

The PHD Appliance Role Selection menu opens.



6. Select **D, Set role as esXpress DeDupe Appliance**.
7. Select **Yes** to confirm.
The appliance is set as a DeDupe Appliance.
8. Select **OK** to open the appliance menu.
9. Select **C, Configure DeDupe Appliance settings**. For details on each text menu option, see "DeDupe Appliance Text Menu Options" (on page 24).
10. At the next menu, select **C, Configure DeDupe Appliance Networking**.
11. Select a Network Adapter number to configure the IP address for each network adapter assigned to the appliance. Additional Network Adapters are displayed only if you have added them to the DeDupe Appliance from the VMware client.

Note: Assign the appliance a unique IP Address - do not use the IP Address of the ESX server, Management Appliance, or the NFS share (if configured). This IP will be used to access the DeDupe Appliance web interface.
12. Select **OK**.
13. Select **Yes** to confirm the changes.
14. A window opens alerting you that you must define a hostname and default gateway. Select **OK**.
The DeDupe Target Global Network Configuration menu opens.



The following information is required:

- **Hostname** (must be a fully qualified domain name, do not use localhost)
 - **Default Gateway**
 - **DNS Server** (If required, you can enter multiple DNS Servers separated by a comma)
15. Set **Automatically_Check_for_Updates** to **yes** to allow the DeDupe Appliance to periodically check for updated software. After installation, this setting can be configured within the DeDupe Appliance Options Page, SYSTEM options. See "System Options" (on page 41).
 16. If required, also define the value for your proxy server. The appliance must be able to connect to the Internet to receive updates. You must define the server name and port for the proxy server, for example, `http://proxy.example.com:8080/`.

If your proxy server requires authentication, use the following syntax:
`http://username:password@proxy.example.com:8080/`

17. When finished defining network settings, select **OK**.
18. Select **Yes** to confirm the changes.
19. Select **OK** to configure the network settings.
20. When complete, select **OK** again.
21. Select **Quit** to return to the main menu.

Now that the appliance has been configured, check for any software updates.

22. Select **U, Check for DeDupe Appliance Updates** and follow the steps to download the latest software. When the updates are complete, you are prompted to restart the appliance.

The DeDupe Appliance is configured. Next, configure the storage the appliance should use for storing backups. See "Configuring Storage for the DeDupe Appliance" (on page 19).

Checking for DeDupe Appliance Updates

As new updates and versions are released, you can check for appliance updates using the **Check for Updates** option on the DeDupe Appliance web page or the **Check for DeDupe Appliance Updates** option on the DeDupe Appliance Text Menu. The updates will be downloaded from an esXpress server using HTTP, so the appliance must be able to access the Internet. Follow the instructions to apply the updates to the appliance.

If you are using a proxy server make sure you have configured the HTTP_PROXY option in the Network Settings as described earlier.

In addition to checking for updates manually, you can set a schedule to check for DeDupe Appliance updates. See "System Page" (on page 27) and "System Options" (on page 41) for more information.

If your appliance does not have access to the Internet, you can upgrade the appliance manually, see [Manually Updating the DeDupe Appliance](#).

Manually Updating the DeDupe Appliance

If your DeDupe Appliance is located behind a firewall or the appliance cannot access the Internet for another reason you will not be able to use the Check for Updates process to upgrade the appliance.

To manually upgrade the appliance, you will need to download the upgrade package using a machine with Internet access then copy the package to the appliance, manually.

To manually upgrade the DeDupe Appliance

1. Use a web browser and navigate to:

http://www.esxpress.com/Documents/phdd_updates/PHDD_Server_Update.tgz

2. Download the update package.
3. Copy the files using WinSCP or another program to the **/home/phd/updates** directory on the DeDupe Appliance.
4. Using PuTTY or another program, open an SSH window to the appliance.
5. Uncompress the copied update package using the following command:

```
tar -xzvf PHDD_Server_Update.tgz
```

6. Change to the directory that was uncompressed (phdd-20100520-4.0-3 for example).
7. Run the update:

```
./run_update
```

Configuring Storage for the DeDupe Appliance

Only one storage type can be used per appliance, VMDK, SMB, or NFS. The following sections describe how to configure each of these storage types for the DeDupe Appliance:

- "NFS External Storage" (on page 19)
- "SMB External Storage" (on page 19)
- "VMDK Attached Storage" (on page 20)

Caution: Running anti-virus software on a backup target can result in file locking or deletions and may additionally cause issues with writing and deleting backups. PHD Virtual recommends excluding backup targets from anti-virus software scans, including the DeDupe Appliance, network shares, and directories used for backup targets (SMB, NFS, or FTP).

NFS External Storage

The following steps describe how to configure the DeDupe Appliance to use an external NFS server to store backups.

To configure NFS storage

1. From the DeDupe Appliance Text Menu, select **C Configure DeDupe Appliance Settings**.
2. Select **N Connect to external NFS share**.
3. Set NFS Export Path, NFS server's IP and share path, for example xxx.xxx.xxx.xxx:/mnt/Backup.
4. Set **Direct VBA Connection** to **Yes**.

Note: Set Direct VBA Connection allows the VBA to connect directly to the NFS share and avoids sending the data through the DeDupe Appliance. This lowers the I/O load on the appliance. If you wish to not allow this, **Set Direct VBA Connection = No**

5. When finished defining settings, select **OK**.
6. Select **Yes** to confirm the changes.

After configuring your NFS storage for the DeDupe Appliance, verify if the daemon has started successfully by viewing the PHDD node log. Additionally, you can view the mounted share using the main DeDupe Appliance menu. From the text menu, press ALT+F1 to open the main menu.

SMB External Storage

The following steps describe how to configure the DeDupe Appliance to use an external SMB share to store backups.

To configure SMB storage

1. From the DeDupe Appliance Text Menu, select **C Configure DeDupe Appliance Settings**.
2. Select **S Connect to external SMB share**.
3. Set SMB Export Path, SMB server's IP and share path, for example //xxx.xxx.xxx.xxx/Share.
4. Set the SMB username and password.
5. When finished defining settings, select **OK**.

6. Select **Yes** to confirm the changes.

After configuring your SMB storage for the DeDupe Appliance, verify if the Daemon has started successfully by viewing the PHDD node log. Additionally, you can view the mounted share using the main DeDupe Appliance menu. From the text menu, press ALT+F1 to open the main menu.

VMDK Attached Storage

The following steps describe how to configure the DeDupe Appliance to use an attached VMDK to store backups.

To configure VMDK storage

1. From the DeDupe Appliance Text Menu, select **C Configure DeDupe Appliance Settings**.
2. Select **V Use attached VMDK**.
3. Select **Yes** to begin the VMDK configuration, which will format the VMDK volume and create its file system. This process may take several minutes. When the disk is configured, a **Configuration Successful** message displays.
4. Start the PHDD service.
5. Before proceeding with backups, check if the three directories **Backups**, **Blocks** and **Config** have been created in **/.data** folder. If not, refer to "[Remaking Block Folders](#)" (on page 20)

Caution: Do not change, delete, or move any files or directories within **/.data** without the assistance of support as this may result in corruption or loss of backup data.

Note: Please refer to the esXpress Management Appliance Guide to configure esXpress to use the DeDupe Appliance as a backup target. You must configure esXpress prior to taking backups.

Remaking Block Folders

After the storage has been mounted, three directories **Backups**, **Blocks** and **Config**, are created in **/.data**. Also, in the **Blocks** folder, 0-255 blocks are created.

If these directories are not being created and the blocks are not present in the **Blocks** directory, go to the DeDupe target configuration menu and select **Remake Block Folders**.

esXpress Backup Target Configuration

When the DeDupe Appliance is configured, it will become a new Backup Target for esXpress and is configured like any other esXpress backup target. The DeDupe Appliance is a PHDD target type.

- This configuration is done outside of the DeDupe Appliance and performed when configuring esXpress, from either the PHD text menu or the esXpress Management Appliance. Refer to the Management Appliance Guide for details.

Appliance Maintenance

The following sections explain DeDupe Appliance maintenance tasks including information on extending the appliance disk and where to find the DeDupe Appliance logs.

- ["Extending Appliance Disks" \(on page 22\)](#)
- ["Appliance Logs" \(on page 22\)](#)
- ["Backing Up the Attached VMDK" \(on page 23\)](#)

Extending Appliance Disks

You can extend the disks for the DeDupe Appliance, both the system VMDK and your data disk. Initially, the system VMDK ships as an 8 GB disk (if you used the deployment wizard, the appliance disk will be 16 GB). This can be extended using the VMware client as you would for any other VM. After extending the disk, reboot the DeDupe Appliance. The reboot will take care of what is needed in the appliance for extending the disk. It will reboot a second time automatically.

When the appliance boots the second time, the additional space should be visible to the appliance.

Note: The DeDupe Appliance allows only one additional disk to be added for DeDupe backups. We recommend extending the system VMDK to at least 16 GB before deploying the DeDupe Appliance. Over time, as backups are written to the appliance database, the size of the database and logs will grow. If space becomes an issue, the recommended solution is to extend the system VMDK.

Appliance Logs

There are a number of log files written on the appliance to view the status and activity of the esXpress daemon processing. All of the logs can be found in `/var/log/messages` on the DeDupe Appliance.

- **phd_daemon.log** - the main DeDupe Appliance daemon log.
- **phd_delete.log** - appliance auto delete processing log (Trim the Herd).
- **phd_smb.log** - samba logs used when exporting backups.
- **phd_update.log** - Check for Updates and `run_update` logs.
- **phd_verify.log** - daemon process log that checks the appliance blocks and verifies there are not any existing bad blocks.

The logs can be viewed using the DeDupe Appliance web page **Logs** tab. The logs can be viewed as a static page which shows the logs as they existed when you launched the page or they can be viewed in real time by checking the **Live Log View**

esXpress DeDupe Appliance [admin@doc_dedup_vmdk.espress.local | Logout]

Disk Space is OK, Var: 16%, Data: 8% System VMs FLR Stats Options Users Logs

Logs

Summary Daemon Deletes Verify Updates SMB

Mon Jan 4 17:06:58 EST 2010

2010-01-04 17:06:08 - /var/log/phdd_daemon.log

Last 50 lines from: /var/log/phdd_daemon.log

```
2010-01-04 17:06:08 import: Starting Import Daemon
2010-01-04 17:05:59 daemon: Starting phdd_import program
2010-01-04 17:05:59 daemon:
2010-01-04 17:05:53 import: Running Daily Stats for Imported Backups
2010-01-04 17:05:53 import: Complete
2010-01-04 17:05:53 import: Checking for daily stats run
2010-01-04 17:05:38 import: Running Daily Stats for Imported Backups
2010-01-04 17:05:38 import: Complete
2010-01-04 17:05:38 import: Checking for daily stats run
2010-01-04 17:05:33 import: Re-Running Stats for Deleted Backups
2010-01-04 17:05:33 import: Complete
2010-01-04 17:05:33 daemon: Running Daily Stats
2010-01-04 17:05:15 daemon: Disk Space is OK, Var: 16%, Data: 8%
2010-01-04 17:05:15 daemon: Calling delete process
2010-01-04 17:05:14 import: Looking for and deleting unsuccessful imports
2010-01-04 17:05:14 import: Complete
```

Backing Up the Attached VMDK

You can backup the DeDupe Appliance's attached VMDK as a Full/Delta backup, only. To do this you will need to create a Delta/Full target (FTP server, SMB target, SSH target) and follow your normal procedures for creating and running Delta or Full backups.

Alternatively, you can exclude the entire DeDupe Appliance from the esXpress automatic backups by adding the [x0] command to the VM Notes or by using a skip file.

DeDupe Appliance Text Menu Options

The following table describes the DeDupe Appliance text menu options. After initially configuring the appliance network settings, you do not need to use the text menu as all configuration can be done from the DeDupe Appliance web page. See "[DeDupe Appliance Administration](#)" (on page 25) for details. You will only need to use the text menu if you must change the appliance networking configuration.

Table 2 - DeDupe Appliance Text Menu Options

Menu Option	Description
Daemon Status (D)	Shows the status of the daemon running in the background. This must be running to accept backups.
Lock Status (A)	Set = Reset the lock status Clear = The appliance is free to take backups or in progress
Show Storage Status (S)	View the Free Space in the Backup Storage
Show status from TTY1, TTY10, TTY11 (T,0,1)	TTY1 – shows appliance main screen TTY10 – shows outbound daemon messages TTY11 – shows inbound daemon messages
Configure DeDupe Appliance settings (C)	Configure the DeDupe VM Networking and the target storage.
System Config Settings (O)	Maintain the DeDupe Appliance settings, including stats and trim options.
View Appliance Logs (L)	View the event logs on the machine.
Check DeDupe Appliance Updates (U)	Check for and install DeDupe Appliance updates.
Verify Repair Appliance Database	Use this function if you are experiencing Database corruption or crashed errors.
Restart Web Services	Restart the DeDupe Appliance web services.
Please change the ROOT Password (R)	Change the ROOT Password.
Quit (Q)	Quit the menu

Chapter 3 - DeDupe Appliance Administration

The DeDupe Appliance web page contains configuration options and statistics you can use to manage the appliance as well as an interface for restoring DeDupe backup files and directories.

The following sections describe the DeDupe appliance web interface:

- "Accessing the DeDupe Appliance Web Page" (on page 26)
- "System Page" (on page 27)
- "VMs Page" (on page 29)
- "FLR (File Level Restores) Page" (on page 31)
- "Stats Page" (on page 34)
- "Options Page" (on page 37)
- "Users Page" (on page 43)
- "Logs Page" (on page 44)

Accessing the DeDupe Appliance Web Page

The DeDupe Appliance web page can be accessed from any supported browser. See "[DeDupe Appliance System Requirements](#)" (on page 11) for details.

To access the DeDupe appliance web page

1. Open a web browser and enter the IP address you assigned to the DeDupe Appliance.
2. Enter the log in credentials, by default these are user **admin** and password **admin**.

System Page

When you first log in to the DeDupe appliance you are presented with a System page that displays general DeDupe Appliance information including the appliance hostname, the esXpress version, and disk usage. You can use the controls on the page to change the System Status, or to restart or shutdown the appliance.

Table 3 - DeDupe Appliance System page items

System page item	Description
Hostname	Displays the appliance hostname
System Uptime	Displays the time the appliance has been online.
System Status	Displays the current DeDupe Appliance status. Operational - The appliance is working and the DeDupe database will import backups, trim backups, and verify data as normal. Maintenance Mode - The appliance will not import backups, deletes will not occur, and data will not be verified. See " Maintenance Mode " (on page 27)
Version	The current installed version of esXpress.
Check for Updates	Displays available appliance updates. See " DeDupe Appliance Updates " (on page 28)
Network IP address	Displays all IP addresses configured for the appliance.
Storage	Displays available space for the backup, database and logs, and system volumes for the appliance. Backup displays the space available to store backups. See " Storage " (on page 28)

Maintenance Mode

When in Maintenance Mode, backups will still run, but the DeDupe Appliance will not import backups to the database, deletes will not occur and data will not be verified.

Note: Entering Maintenance Mode is also known as **Locking** the DeDupe Appliance, and may be referenced as setting a lock or **Locks** in other areas of the product.

Maintenance Mode should be used only if you are performing maintenance on your DeDupe storage (NFS, or SMB, or attached VMDK) or if instructed by support to prevent the appliance from attempting to manage backup data.

- To enter Maintenance Mode, on the DeDupe Appliance **System** page, click **Enter Maintenance Mode** in the **System Status** row. The System Status will change from Operational to **Maintenance Mode**.

System Status: Maintenance Mode  [Exit Maintenance Mode](#)

DeDupe Appliance Updates

You can use the System page to check for DeDupe Appliance updates.

- To check for updates, click **Refresh** on the **Check for Updates** row. When updates are found, click **Install Update** to apply the updated appliance software.

Updates can be checked for and installed automatically using the **System** options, on the Option page. See "[System Options](#)" (on page 41) for details.

Storage

The Storage displayed on the System page contains a visual representation of the used and available space for the backup, database and logs, and system volumes for your DeDupe Appliance.

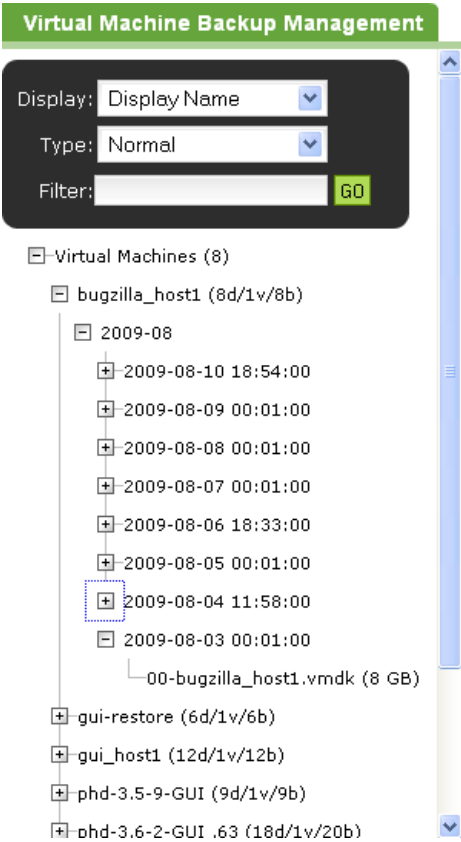
- **Backup**- represents the available storage for the DeDupe storage device. This can be an attached VMDK, NFS, or SMB (CIFS) share. If you begin to run out of space for this volume, either contact PHD Virtual Support, delete older backups manually, or configure trim settings to be more aggressive by adjusting the purge time to 12 and setting the deletes to run from 0 to 23 on the Options page (see "[Auto Delete - DeDupe Backups](#)" (on page 37)). Alternatively, you can increase the amount of space available on the backup target share.
- **Database & Logs** - displays the available space for the DeDupe operating system volume. The DeDue database is stored on this volume, so there must be available space for the database to grow. If you begin to run out of space for this volume, either contact PHD Virtual Support or shutdown the DeDupe Appliance VM and use vCenter to grow the size of the disk. The recommended size for this disk is 16 GB.
- **System** - shows the available space on the DeDupe Appliance root volume. If you begin to run out of space for this volume, contact PHD Virtual Support.

VMs Page

The VMs page allows you to manage the backups stored on the DeDupe Appliance. Use the tree view to display all of the stored backups and to view the individual VMDK files. As of version 4.0, VMs with IDE disks are supported and any IDE disks that were backed up to the DeDupe Appliance are displayed. IDE disks are identified with the letter (A or B) that corresponds to their device ID.

Note: If VM names contain special characters other than A-Z, a-z, 0-9, <space>, #, .<dot> or , <comma> they will be replaced with _<underscore>. For example, a VM with the name **sample-vm** will be displayed in the DeDupe Appliance as **sample_vm**.

Expand each tree to view the DeDupe backups stored on the DeDupe Appliance. The top element displays the total number of VMs with backups. For example, the screenshot below shows 8 total VMs with backups, as designated by **Virtual Machines (8)**.



Each VM tree also displays a summary of the total number of days, VMDKs, and backups in the parentheses after each name. In the screenshot example above, the bugzilla_host1 VM has a total of 8 days with 1 VMDK file for a total of 8 backups as designated by **bugzilla_host1 (8d/1v/8b)**.

To modify the backup tree view

- Use the view options, **Display**, **Type**, and **Filter**, to change the backups that are displayed in the tree view.
- After selecting your options, click **Go** to update the view. Only the backup files that meet the criteria you entered are displayed.

- To display all files again, remove any **Filter** values and set the **Type** dropdown to **All**, then click **Go**.

Managing Backups

By drilling in to a particular backup file, the individual VMDK files that are backed up and stored on the DeDupe Appliance are displayed on the right side of the page along with the available management options.

The screenshot below shows the options available for a sample backup.

The screenshot displays the esXpress PHDD Appliance interface. At the top, the navigation bar includes the esXpress logo, the title 'esXpress PHDD Appliance', and user information '[admin@DEDUP-64 | Logout]'. Below this, a status bar indicates 'Disk Space is OK, Var: 6%, Data: 61%' and provides navigation tabs for 'Summary', 'VMs', 'FLR', 'Stats', 'Options', and 'Users'. The main section is titled 'Virtual Machine Backup Management'. On the left, a sidebar shows a tree view of virtual machines, with 'bugzilla_host1 (8d/1v/8b)' expanded to show a list of backup dates from 2009-08-10 to 2009-08-03. The selected backup file is '00-bugzilla_host1.vmdk (8 GB)'. The right pane shows a table with the following data:

Archive	Special	Delete	Bad	Date/Time	Display Name	VMDK
<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	2009-08-03 00:01:00	bugzilla_host1	00-bugzilla_host1.vmdk
Select All						

Below the table, it shows 'Row 1 - 1 of 1' and a 'Page Records' button. A 'Save' button is located at the top right of the table area.

The following flags can be applied to each VMDK file, individually:

- **Archive** – Setting the Archive flag will preserve the backup file indefinitely—it will not be deleted according to the Auto Delete/trim rules you’ve defined.
- **Special** – Use the Special flag to indicate a backup that was not part of the normal daily backups.
- **Delete** – Flags the backup for deletion the next time the purge process runs.
- **Bad** – Designates a bad backup. Set automatically by esXpress and cannot be changed.

After applying a flag to any backup file, click **Save** to commit your changes. When saved, you can modify the tree view to display only the files you’ve flagged using the available filters.

FLR (File Level Restores) Page

The esXpress DeDupe Appliance allows you to perform File Level Restores (FLR) through the web Interface. This permits a view into Windows and Linux backups, allowing immediate retrieval of individual files or entire directories.

FLR supported file systems:

- NTFS
- EXT2
- FAT16
- XFS
- FAT32
- ReiserFS (experimental support only)
- EXT3
- JFS

Note: only MBR partition tables are supported at this time. GPT is not currently supported.

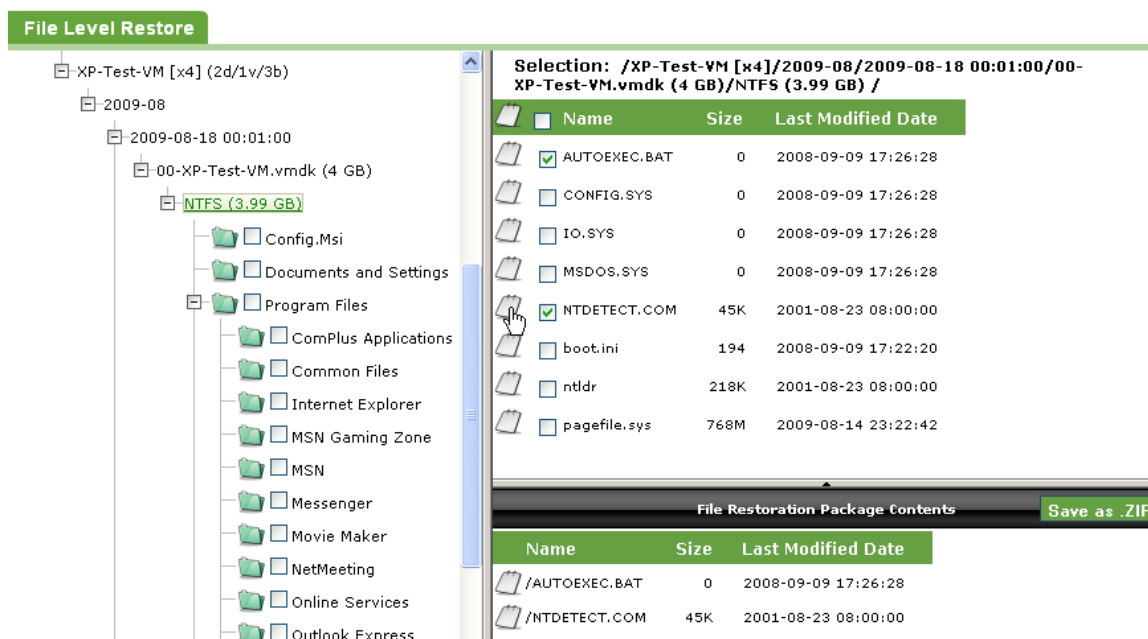
FLR Limitations:

- File-level restore does not support NTFS encrypted files/folders.
- File-level restore does not support multi-drive dynamic disks; however, single drive dynamic disks (simple dynamic disks) ARE supported.
- File-level restore does not support multi-drive Linux LVM; however, single drive LVM (such as Red Hat's installation defaults) IS supported.
- File-level restore does not allow you to create a single restore archive from multiple VMs. You must restore from each VM individually.
- ReiserFS is experimentally supported, only. FLR on Reiserfs volumes may work in some cases but may also cause issues with the DeDupe Appliance (DeDupe Appliance may crash). This is especially prevalent in SLES10 VMs.

Note: Creating a file-level restore zip archive that is larger than 4 GB will fail (though restoring single files larger than 4 GB is fine). This is a limitation of the zip file format.

To restore individual files

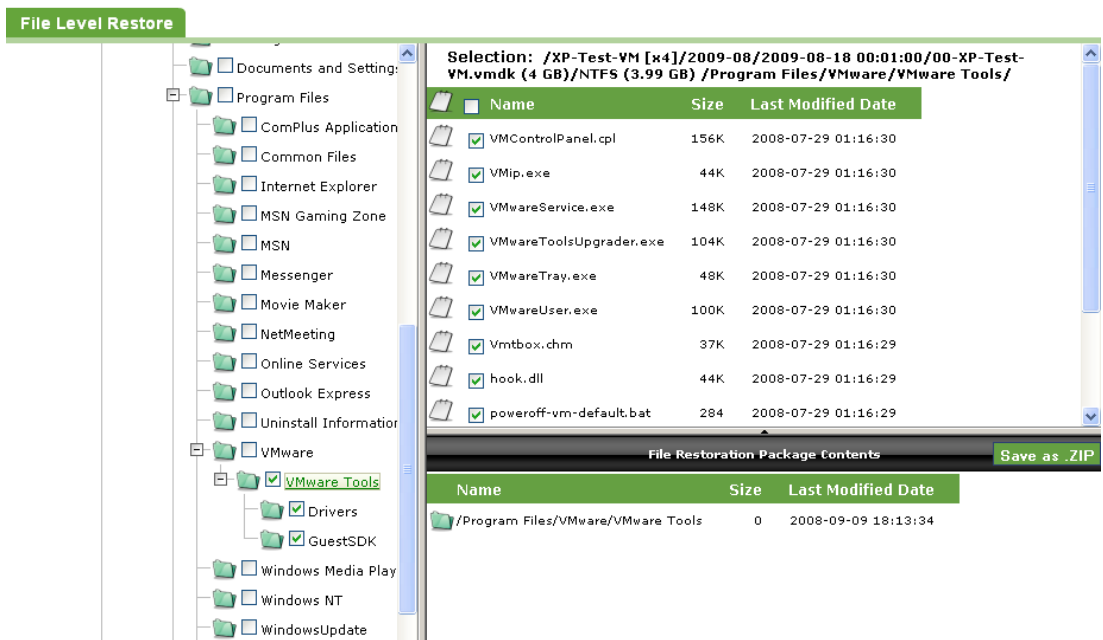
1. Using a web browser, log in to the DeDupe Appliance.
2. At the top of the page, click **FLR**.
3. From the list of backups, expand the tree view and select the VM that contains the file you want to restore (Identified by the time stamp when the backup was taken).



4. Click the VMDK file to expand the directory structure for the disk.
5. Continue to click each directory until you find the folder that contains the file you want to restore. Individual files are displayed in the file list to the right.
6. From the file list, locate the file that you wish to restore and click to the left of the file.
7. Save the file.

To restore a directory with FLR

1. Using a web browser, log in to the DeDupe Appliance.
2. At the top of the page, click **FLR**.
3. From the list of backups, expand the tree view and select the VM that contains the directory you want to restore (Identified by the time stamp when the backup was taken).
4. Click the VMDK file to expand the directory structure for the disk.
5. Continue to click each directory until you find the folder that contains the files you want to restore.
6. Select the directory. The selected directory is displayed in the window to the right.



7. Click **Save as ZIP**.
8. Specify the location to save the file and when the file is downloaded, double-click on the restore.zip file to view the contents.

Stats Page

To view a summary of the DeDupe backup data, click **Stats** at the top of the DeDupe Appliance web page.

The Stats page of the DeDupe Appliance web interface shows the summary of the backup data. Statistics are not dynamically updated each time you view this page—the frequency of updates is determined by the STATS options settings (see "Statistics Options" (on page 40)). By default, detailed statistics are run daily.

Stats can also be recalculated using the **ReCalc** button on this page.

The screenshot shows the esXpress PHDD interface. At the top, there's a navigation bar with 'Summary', 'VMs', 'FLR', 'Stats' (selected), 'Options', 'Users', and 'Logs'. A warning message states 'WARNING: Disk Space Low, 100% Used on /data'. Below the navigation bar, the 'Backup Statistics' section is active. It includes tabs for 'Trim Report', 'By Date', 'By VM Name', 'By Display Name', and 'By VMDK'. A 'ReCalc' button is present, with the text 'Last Updated: Pending Update'. The main content area is divided into three columns: 'esXpress v3.6, PHDD DeDupe Stats', 'Backup Counts', and 'Other Counts'. Each column contains a table of statistics. Below these tables are buttons for 'Re-Check Bad Backups' and 'Delete Bad Backups'. At the bottom, there's a calendar view titled 'Months of the Year 2009 and 2008' and 'Count of Backups by each Week for'. The calendar shows data for 'GUI-TEST1' and '00-GUI-TEST1.vmdk' with color-coded cells (green, orange, red) indicating backup status.


Statistics can be displayed by date, by VM name, display name, or by VMDK using the tabs on the page. The **Trim Report** tab also contains a calendar table showing the total backups for each VM by month and by week.

This screenshot shows the 'Trim Report' section of the interface. It features the same navigation tabs as the previous screenshot: 'Trim Report', 'By Date', 'By VM Name', 'By Display Name', and 'By VMDK'. Below the tabs is a 'ReCalc' button and the text 'Last Updated: 2009-09-21 10:05:12'.

Trim Report

In the Trim Report calendar table, the orange and red highlighting gives you a quick view of which backups will be deleted based on the settings you have configured. Delete settings are configured on the Options page. See "Options Page" (on page 37) for details.

Orange highlighting indicates the backup set is partially within the delete zone—some of these backups will be deleted. Red highlighting indicates all backups in the set will be deleted during the next delete window. The following image shows an example backup set that contains backup files to be deleted, highlighted accordingly.



You can drill in to the report by clicking any backup link in the calendar table, as shown in the following image.

phd-3.5-8-GUI .215	8	3	1
00-phd-3.5-8-GUI .215.vmd	8	3	1
quantum-esxpress-rosco	8	3	1
00-quantum-esxpress-rosco	8	3	1
Test	.	1
00-Test.vmdk	.	1

When you click a backup link, a new table is displayed which shows statistics for that set of backups. The following image shows an example table.

Showing Backups where month = '200907' and scsi = '00' and vmdk = 'GUI-TEST1.vmdk'

All Numbers in Blocks / Megabytes

Hostname	Backup ID	Week	Date	Time	VM Name	SCSI ID	VMDK Name	Size in Meg	Real Data	NULL Blocks	Sent Meg	OS Type
phddev0.esxpress.local	463	200928	2009-07-18	01:01	GUI-TEST1	00	GUI-TEST1.vmdk	8,192	4,625	3,567	83	rhel5
phddev0.esxpress.local	590	200929	2009-07-25	01:01	GUI-TEST1	00	GUI-TEST1.vmdk	8,192	4,809	3,383	244	rhel5
phddev0.esxpress.local	602	200930	2009-07-29	01:01	GUI-TEST1	00	GUI-TEST1.vmdk	8,192	4,832	3,360	271	rhel5
phddev0.esxpress.local	319	200927	2009-07-11	00:01	GUI-TEST1	00	GUI-TEST1.vmdk	8,192	4,452	3,740	34	rhel5
phddev1.esxpress.local	407	200928	2009-07-15	00:01	GUI-TEST1	00	GUI-TEST1.vmdk	8,192	4,356	3,836	0	rhel5

You can continue to drill in to the table by clicking any of the links available within each column, filtering the backup information based on the information you click. The table header displays what data is currently represented in the table.

DeDupe Stats

The first table on the Stats page shows the main overall DeDupe stat values for all of your DeDupe backups. Each value is listed along with the total size in MB and the overall ratio.

esXpress v3.6, PHDD DeDupe Stats

Stat Value	in Megabytes	Ratio
All Blocks of All VMs	606,576	34:1
Null Blocks	296,141	49%
Real Data Blocks	310,435	17:1
Total Unique Blocks	65,127	4:1
Total Unique Blocks Compressed	17,956	1:1

- **All Blocks of all VMS** – This is the total number of blocks for the backups sent to the DeDupe Appliance. In the figure above, the total was 606,576. The ratio is the overall deduplication ratio. Note that the ratio 34:1 is a test host in a lab with many similar Virtual Machines. Normal user environments usually return about 25:1
- **Null Blocks** – This is the number of Null Blocks. In the example above, 49% of all blocks from the backups are Nulls.
- **Real Data Blocks** – This is the total amount of actual data blocks in the backups.
- **Total Unique Blocks** – The total amount of actual unique blocks sent to the DeDupe Appliance for your backups. The Ratio is the unique blocks against the total data.
- **Total Unique Blocks Compressed** – This is the total unique blocks compressed.

To learn more about the Stats page, view the [Knowledge Center web site](#).

Backup Counts

The Backup Counts table shows totals for all backups and virtual disks for the DeDupe Appliance.

Backup Counts

Stat Value	Total Count
Blocks of All VMs	17,311
Total Backup Count	11
Total Unique Dates	7
Total Unique VMDKs	3
Total Virtual Machines	3

Other Counts

The **Other Counts** table indicates if there are bad blocks, the number of deleted backups and blocks waiting to purge. The **Size of Purged Blocks** displays the space waiting to be reclaimed from the Trim process.

Click **Re-Check Bad Backups** to update the Other Counts table. You can initiate a delete of bad backups directly from this page by clicking **Delete Bad Backups**.

Other Counts

Stat Value	Total Count
Bad Backup Count	76
Bad Block Count	2,911
Backups to be Deleted	0
Deleted Backups	564
Blocks to be Purged	0
Size of Purged Blocks	0

Re-Check Bad Backups

Delete Bad Backups

Bad blocks can occur for a number of reasons. The following are the most common:

- The block doesn't physically exist on disk.
- The md5sum is not correct.
- gzip fails.
- The block has become corrupt.

Options Page

Use the Options page to edit the attributes associated with your backups.

DeDupe options are separated into the following groups:

- DELETE: see "Auto Delete - DeDupe Backups" (on page 37)
- STATS: see "Statistics Options" (on page 40)
- SYSTEM: see "System Options" (on page 41)


Available options can be filtered using the **Group** drop-down filter box at the top of the options table. The following screenshot shows the Group filter options.

Group	Sub
DELETE	DEFAULT
DELETE	WINDOW
DELETE	WINDOW

To edit DeDup options

1. Use a web browser to log in to the DeDupe Appliance.
2. At the top of the page, click **Options**.
3. View the list of DeDupe options and find the attribute you would like to change. When you change an attribute value, the Save and Cancel buttons are displayed at the end of the attribute row.

Group	Sub	Name	Value
DELETE	DEFAULT	ENABLE	<input type="checkbox"/>
DELETE	WINDOW	START	10 [0 - 23]

4. Make your changes and click .

Auto Delete - DeDupe Backups

As of esXpress 3.6, the auto delete functionality for DeDupe backups now occurs at the DeDupe Appliance and not at the host level. The auto delete approach for the DeDupe Appliance is called Trim the Herd. Over time, you can trim down the number of backups kept on the appliance by determining how many backups to keep per week, per month, and per year, keeping only the number of backups you define and deleting the others (trimming them). For auto delete planning information, see "Planning for Auto Deletes/Trim" (on page 40).

DELETE options are configured on the Options page. See the table below for details.

Minimum Unique and Maximum Weeks

Based on the auto delete/trim settings, over time, backups older than a certain date should be purged from disk. The esXpress KEEP_MAXIMUM_WEEKS setting allows you to set the maximum age of a stored backup. When a backup exceeds this maximum age, esXpress first checks if this backup is set as Archived (do no delete), and then checks the KEEP_MINIMUM_UNIQUE setting. This setting ensures that esXpress

will not delete all backups for a VM. If, for example, you stop backing up a specific VM, given enough time, the MAX_WEEKS setting will allow all backups of this VM to be deleted, unless MIN_UNIQUE is set to greater than 0. In this event, esXpress will always keep a minimum number of unique backups of a VM, overriding the Maximum Weeks setting.

Table 4 - Delete Options

Group	Delete Option Name	Sub	Default Value	Description
DELETE	ENABLE (Yes/No)	DEFAULT	<i>(selected)</i>	When enabled, deletion of the backup from the appliance database (Soft Delete) will occur. When not enabled, auto delete will not run on the DeDupe Appliance.
DELETE	KEEP_MINIMUM_UNIQUE	MIN_UNIQ	2	The minimum number of unique backups to keep for a VM. This setting will ensure the trim process never deletes all backups of a VM. This can be set to 0, if desired.
DELETE	KEEP_MAXIMUM_WEEKS	MAX_WEEKS	12	The maximum number of weeks of backups to keep for each VM. Note that backups set to Archive are not deleted. Also, if the total number of remaining backups is less than the MIN_UNIQUE value, no backups will be deleted.
DELETE	THIN_SPECIAL	DEFAULT	<i>(not selected)</i>	Select this option to also include Special backups in the 'Thin the Herd' logic. By default, Special backups are not removed.
DELETE	THIN_PER_DAY (1 - 99999)	DAY	1	This indicates how many backups you wish to keep per day.
DELETE	THIN_AFTER_DAYS (1 - 99999)	DAY	5	This indicates in days when to start the 'Thin the Herd' logic. Before this time no deletes will occur.

Group	Delete Option Name	Sub	Default Value	Description
DELETE	THIN_KEEP_WHICH_DAY (first/last)	DAY	Last	If you have more than one daily backup, this indicates which one to keep, either the first or the last backup taken for the day.
DELETE	THIN_PER_WEEK (1 to 99999)	WEEK	1	Indicates how many backups you want to keep per week.
DELETE	THIN_AFTER_WEEKS (1 to 99999)	WEEK	1	This indicates the number of weeks to wait before starting the weekly 'Trim the Herd' logic. Until this time is reached, the daily 'Trim the Herd' logic will run.
DELETE	THIN_KEEP_WHICH_WEEK (first/last)	WEEK	Last	When the weekly 'Trim the Herd' logic runs, this indicates which backup to keep, either the first or the last backup for the week.
DELETE	THIN_PER_MONTH (1 - 99999)	MONTH	1	Indicates how many backups you want to keep per month. These are sequential days from either the last or the first of the month depending on your setting.
DELETE	THIN_AFTER_MONTHS (1 to 99999)	MONTH	1	This indicates the number of months to wait before starting the monthly 'Trim the Herd' logic. Until this time is reached, the weekly 'Trim the Herd' logic will run.
DELETE	THIN_KEEP_WHICH_MONTH (first/last)	MONTH	Last	When the monthly 'Trim the Herd' logic runs, this indicates which backup to keep, either the first or the last backup of the month.

Planning for Auto Deletes/Trim

When planning delete/trim settings, it is a good idea to look at it from the inside out; starting with your daily delete rules, then weekly rules if required, and then finally, the monthly rules. With this approach, you can slowly trim backups over time while retaining robust archive requirements. When planning how to set these values, consider the following:

Global Trim Settings

1. In order to ensure older backups are always removed and backup storage is available, define the value for KEEP_MAXIMUM_WEEKS to set the age per VM backup. Backups older than this setting will be trimmed.
2. If you would like to ensure at least one backup is kept for every VM (especially for VMs that may not be backed up any longer) use the KEEP_MINIMUM_UNIQUE setting. This will ensure a certain number of backups are retained for each VM, regardless of your other trim settings.

Daily, Weekly, and Monthly Trim Settings

1. Set the value for the option THIN_AFTER_DAYS to control how many days of backups to retain before backups are trimmed. For example, if you wanted to keep all backups for 14 days, set this value to 14; the trim process would then begin after 14 days had passed.
2. If you are running multiple backups per day, use the option TRIM_PER_DAY to determine how many you want to keep each day. If you run backups only once per day, this value will not apply.
3. Use the THIN_AFTER_WEEKS option to decide when to start trimming backups on a weekly basis. For example, if you decide that after one month (4 weeks) you would like to start trimming backups into weekly buckets, you would set this value to 4. With this value, the daily trim options would apply for the first four weeks.
4. Next, use the option THIN_PER_WEEK to determine how many backups to keep each week. In most situations, one backup per week is sufficient, but if you do decide to keep multiple backups per week, they will be sequential days. Use the option THIN_KEEP_WHICH_DAYS to determine which day of the week to keep.
5. Use the option THIN_AFTER_MONTHS to determine after how many months backups should be trimmed into monthly buckets. For example, if you determine that after 6 months you will start monthly trimming, then set this value to 6.
6. There are two additional considerations for the monthly trim. The first is how many monthly backups to keep. In most situations, one backup per month is sufficient, but if you wish to keep additional backups each month, then adjust the THIN_PER_MONTH value. When set to greater than one, this option refers to sequential days of backups. Also decide which backup to keep each month by setting THIN_KEEP_WHICH_MONTH to either last for the last of first.

Statistics Options

On the Options Page, there are a number of values which control when and what type of esXpress statistics are generated. Refer to the following table for details about each STATS attribute.

Table 5 - Stats Options

Group	Stat Option	Sub	Default Value	Description
STATS	RERUN_AFTER_X_HOURS (1 to 99 hours)	GLOBAL	12	This option indicates how often the global DeDupe statistics will run. Set this value to a number indicating

Group	Stat Option	Sub	Default Value	Description
				after how many hours the global stats will be generated. For example, if this is set to 12 the statistics will run every 12 hours.
STATS	RUN_STATS (DAILY / ALWAYS)	DETAILED	DAILY	Indicates when detailed statistics are run. Daily – run the detailed statistics daily. Always – run the detailed statistics after every backup is imported.
STATS	START_TIME (0 to 23 hours)	DETAILED	10	Defines the start time for the window when detailed statistics are allowed to run.
STATS	END_TIME (0 to 23 hours)	DETAILED	20	Defines the end time for the window when detailed statistics are allowed to run.

System Options

The Options page includes SYSTEM settings which can be used to modify the appearance of the DeDupe Appliance web page as well as configure SAMBA settings.

Table 6 - System Options

Group	System Option	Sub	Default Value	Description
SYSTEM	MOTD	LOGIN	esXpress DeDupe Appliance	Define the message displayed at the web login screen for the DeDupe Appliance.
SYSTEM	TITLE	DEFAULT	DeDupe Appliance	Use this option to change the DeDupe Appliance web page title.
SYSTEM	ENABLED	CHECK_FOR_UPDATE	<i>(selected)</i>	Enable automatic check for updates process.
SYSTEM	START_TIME	CHECK_FOR_UPDATE	4	Enter the start time (0-23 in 24 hour notation) for the automatic check for updates process to run. The default is 4 (4 AM).
SYSTEM	ENABLED	SAMBA	<i>(selected)</i>	Use this option to enable backup exports (Date Smart Dynamic Export) using SAMBA. For details, see "Exporting Backups" (on

Group	System Option	Sub	Default Value	Description
				page 45).
SYSTEM	WORKGROUP	SAMBA	ESXPRESS	Defines the workgroup the server will indicate it is a member of when servicing client requests. the default is ESXPRESS .
SYSTEM	SERVER_STRING	SAMBA	esXpress_Backups	The text string displayed with the machine's NetBIOS name when viewed over the network. The default is esXpress_Backups .
SYSTEM	UNCOMPRESSED_SHARE_NAME	SAMBA	espress_uncompressed	Samba share name. The default is espress_uncompressed .
SYSTEM	UNCOMPRESSED_COMMENT	SAMBA	esXpress_Backups	Comment field for the SMB share. The default is esXpress_Backups .
SYSTEM	UNCOMPRESSED_BROWSEABLE	SAMBA	<i>(selected)</i>	Enable or disable the SAMBA share for browsing.

Users Page

Use this page to add users with admin or read-only rights.

Note: The default admin account cannot be used as the SMB authentication user

To add a new user





1. On the DeDupe web page, click **Users**.
2. Click **New User**.
3. Enter the user credentials and click **OK**. Note that currently, only the Admin role is supported.

The user is added to the list.

If you are using the Export Share, you will need to indicate at least one user as the SMB user. The admin account cannot be used for SMB.


User Management

[New User](#) Filter: [Clear Filter](#)

Enabled		UserName	Role	Email	Full Name	SMB
<input checked="" type="checkbox"/>	 	admin	Admin		PHD Admin	No
<input checked="" type="checkbox"/>	 	smbuser	Admin	admin	SMB User	Yes

2 Rows [Export to CSV](#)

To change a user password

1. On the DeDupe web page, click **Users**.
2. In the list of users, click  at the beginning of the line.
3. Use the **Password** and **Confirm Password** text boxes to enter the new password.
4. Click **Ok**.

Logs Page

The Logs page displays DeDupe logging information separated into tabbed categories. Click any tab to display the latest log information.

With esXpress DeDupe backups you will see some additional informational messages written to your backup logs.

Below is a sample portion of a DeDupe backup log.

```
2009-04-29 07:42:59.393b PHDD: 67%, 5496 MB of 8192 MB at 36 (50) MB/Sec. Dur:
02m:32s, 267M / 40M (2)
2009-04-29 07:43:15.989b PHDD: 81%, 6638 MB of 8192 MB at 39 (71) MB/Sec. Dur:
02m:48s, 287M / 42M (2)
2009-04-29 07:43:32.620b PHDD: 92%, 7577 MB of 8192 MB at 41 (62) MB/Sec. Dur:
03m:03s, 295M / 42M (1)
2009-04-29 07:43:49.795b
2009-04-29 07:43:49.803b PHDD Blocks: 8192 Total, 295 Wrote, 3546 Skipped, 4351
Zero for: TEST1, 00-TEST1.vmdk
2009-04-29 07:43:49.813b PHDD Stats: 42 sent meg (14%/0.51%) / 295 written meg
(3.60%) / 8192 total meg for: TEST1, 00-TEST1.vmdk
```

The first group of lines from the log is the running status entries of the DeDupe backup. For example, the line showing PHDD: 67% indicates that the backup is 67% complete, is getting 36 MB/sec (50 MB/sec with both threads), and has been running for 2 minutes and 32 seconds.

The next key line is the PHDD Blocks entry. This shows your detailed DeDupe starts for this backup.

In the example above:

- 8192 blocks were the total for this VMDK (an 8 GB disk).
- Only 295 Blocks were actually written to the DeDupe Appliance, these were the unique blocks in this backup.
- 3546 Blocks were skipped. These are the blocks that already existed in the DeDupe Appliance so are not needed to be saved again as esXpress DeDupe will only write unique blocks once.
- 4351 Blocks in the VMDK were Null (Zero).

The last line is the PHDD stats log entry. This shows how much data was actually sent to the DeDupe Appliance for this backup.

In this example, only 42 MB of compressed data out of the original 8 GB VMDK was sent to the DeDupe Appliance for this backup (the original 295 MB of unique blocks were compressed to 42 MB)

Exporting Backups

You can export backups from the esXpress DeDupe Appliance with Date Smart Dynamic Export (DSDE). With DSDE, backups are exported as uncompressed VMDKs in order to save them to tape or archive them to disk. These exported backups are no longer a data-duplication backup, but simply a compressed VMDK which you can then access.

The export process uses an SMB(CIFS) share on the DeDupe Appliance. It is not enabled by default—you must set the Option (SYSTEM-SAMBA-ENABLED) to **Yes**. This can be maintained from the Options page in the DeDup web page or from the PHD text menu under System Options.

The global configuration options for this feature are also maintained on the Options page and have a sub-category of **SAMBA**. These are:

- **WORKGROUP** - defines the workgroup the server indicates when servicing client requests.
- **ENABLED (yes / no)** - set this option to 'Yes' to enable the SMB share on the DeDupe Appliance, allowing the export share to be accessible. This option is set to **No** by default.
- **SERVER_STRING** - the text string displayed with the machine's NetBIOS name when viewing over the network. The default is **esXpress_Backups**
- **UNCOMPRESSED_SHARE_NAME** - sets the samba share name. The default is **esxpress_uncompressed**.
- **UNCOMPRESSED_COMMENT** - the comment field for the SMB share. The default is **esXpress_Backups**.
- **UNCOMPRESSED_BROWSEABLE (yes / no)** - set this option to **Yes** to allow browsing for the share.

To access the share, map the following: **//DeDupe Appliance IP/esxpress_uncompressed**

- The authentication for the share is defined by setting up an SMB user on the Users page. When you define the user, select the **SMB User** check box.
 - The default admin account cannot be used as the SMB authentication user
 - Special characters are not allowed as part of the user name for SMB users.
- When the share is mapped you will see the following directories:
 - **Archived**- Backups indicated as Archived Backups. These will not be deleted according to the auto delete options.
 - **Normal** - Normal Backups.
 - **Special**- Backups indicated as Special.
- Under each folder the following dynamic folder structure exists:
 - **ByDay** - all backups sorted by day of the week. Contains the most recent backup taken each day for the last eight days.
 - **ByName** - all backups sorted by VM name.
 - **ByWeek** - backups sorted by week, 1-52. Contains the most recent backup taken each week.
 - **CurrentMonth** - all backups taken within the current month, sorted by day.
 - **Last32** - all backups stored within the last 32 days.
 - **Last7** - backups from the last seven days (from the current day).

- **LastMonth** - all backups taken last month.
- **OneOf** - the single, most recent backup taken for every VM available.
- **Today** - all backups that completed today.
- **Yesterday** - all backups that completed yesterday.
- Drilling deeper into the folders you will find the backup data, including the VMX file, the VMDK file and the –flat file. All files are uncompressed.
- The following is a sample directory structure drilled into the Archived folder:
 - |-- Archived
 - |-- ByWeek
 - |-- 47
 - |-- 20081124215800--VIMA--500c6b4c-8451-8306-fb82-39ae9607c4d7
 - |-- 00-VIMA-flat.vmdk
 - |-- 00-VIMA.vmdk
 - |--00-VIMA.vmx
- esXpress makes the exported backup share available. You would then need third-party tools or your own scripting to select and move these files to tape or other disk locations, as well as compression, if necessary.

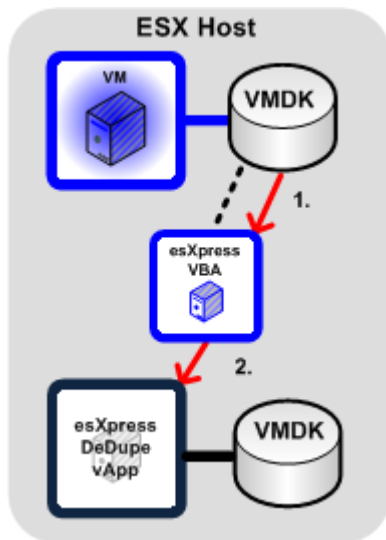
Chapter 4 - Implementation Scenarios

The DeDupe Appliance can be configured to receive backups in a host of scenarios. The following sections describe some of the possible implementation scenarios.

- ["Local VMFS Storage" \(on page 48\)](#)
- ["Remote NFS Storage" \(on page 50\)](#)

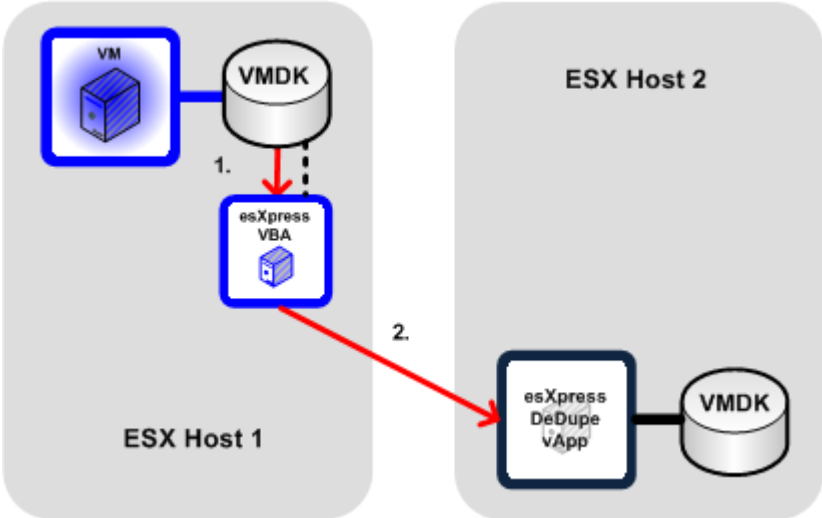
Local VMFS Storage

DeDupe appliance on the same host as the VM it is backing up. Local VMFS storage



1. VBA backs up the data from the VM Guest's disk.
2. VBA sends data directly to the DeDupe Appliance local disk (VMDK can be on Local or SAN storage)


DeDupe Appliance on a remote host from the VMs being backed up. Local VMFS storage



- 1. VBA backs up data from the VM Guest's disk.
- 2. VBA send data directly to the DeDupe Appliance on a remote host.
- 3. Data goes through the DeDupe Appliance to a local disk (VMDK can be on Local or SAN storage).

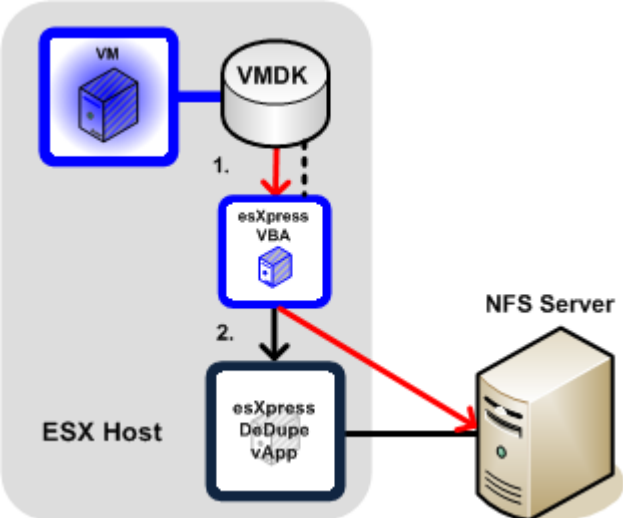
Remote NFS Storage

DeDupe appliance on the same host as the VMs it is backing up. Remote NFS storage




1. VBA backs up data from the VM Guest's disk.
2. VBA sends data directly to the DeDupe Appliance.
3. Data goes through the DeDupe Appliance to a remote mounted NFS Server.

Same as above but the VBA can bypass the DeDupe Appliance and write data directly to NFS storage




- 1. VBA backs up data from the VM Guest's disk.
- 2. VBA writes the backup data blocks directly to the NFS server. Other backup related data is still sent to the DeDupe Appliance (This lowers the I/O load on the appliance). Note that in order to bypass the appliance and write backup data directly to the storage location, the VBA must be on the same subnet as the storage location and it must be routable.

DeDupe Appliance on a remote host from the VMs being backed up. Remote NFS storage



- 1. VBA backs up data from the VM Guest's disk.
- 2. VBA sends data directly to the DeDupe Appliance on a remote host.
- 3. Data goes through the DeDupe Appliance to a remote mounted NFS server.

Same as above but the VBA can bypass the DeDupe Appliance and write directly to NFS storage



1. VBA backs up data from the VM Guest's disk.
2. VBA writes the backup data blocks directly to the NFS server. Other backup related data is still sent to the DeDup Appliance (This lowers the I/O load on the DeDupe Appliance). Note that in order to bypass the appliance and write backup data directly to the storage location, the VBA must be on the same subnet as the storage location and it must be routable.

Chapter 5 - Troubleshooting

This section contains the following DeDupe Appliance troubleshooting topics:

- ["What Happens If My DeDupe Appliance Crashes?"](#) (on page 54)

What Happens If My DeDupe Appliance Crashes?

If your DeDupe Appliance is deleted, crashes, or becomes unrecoverable due to another issue, in most cases, your stored backup data will be unaffected. The DeDupe Appliance contains the database that is used for the export directory, File Level Restores, and for deletes - all of this information can be re-imported to a new database from your stored backup data.

If you cannot recover your DeDupe Appliance, contact PHD Virtual Support. You will be asked to deploy a new DeDupe Appliance (do not recreate the blocks on the DeDupe storage) and then you will be guided through the process of importing the data from your stored backups to recreate the DeDupe database on the new appliance.

Note that you will only be able to recover this information if your backup data is available, for instance:

- If your backup data is stored on an attached VMDK, the storage disk must still exist and be accessible and stable.
- If you're using SMB or NFS storage, the Blocks and Data folders and subfolders must still exist and be accessible and stable.

Index

A		Config	20
About This Guide	7	Configure DeDupe Appliance settings	24
All	30	Configuring Storage	19
All Blocks of all VMS	35	D	
anti-virus	19	Daemon Status	24
Appliance Logs	22	Data DeDuplication	7
Appliance Passwords	9	Date Smart Dynamic Export	8, 41, 45
Archive	30	DeDupe	7
assigning storage space	14	DeDupe Appliance	
Attached VMDK		configuring	15
backing up	23	installing	13
auto delete	37	installing and configuring	10
Auto Deletes		maintenance tasks	22
planning for	40	manually upgrading	18
Automatically_Check_for_Updates	17	optimal configuration	11
B		preinstall steps	12
Backup Counts	36	System Requirements	11
Backup Target Configuration	21	DeDupe Appliance maintenance tasks	22
Backups	20	DeDupe Appliance Text Menu	24
Bad	30	DeDupe appliance web page	25
bad blocks	36	DeDupe Appliance web page	
Block Folders	20	accessing	26
Blocks	20	DeDupe backup	
C		example	7
Check for DeDupe Appliance Updates	17, 24	DeDupe Backup Speeds	7
Check for Updates	27-28	DeDupe logs	44
CHECK_FOR_UPDATE	41	DeDupe options	37
Checking for Updates	17	editing	37
		DeDupe stats	35
		Default Passwords	9

Index

Delete	30		
DELETE	37	L	
Delete Bad Backups	36	Local VMFS Storage	
Display	29	scenario	48
Documentation Updates	4	Lock Status	24
E		Locking	27
ENABLE (Yes/No)	38	Locks	27
ENABLED	41, 45	M	
END_TIME	41	Maintenance Mode	27
Enter Maintenance Mode	27	Managing Backups	30
export backups	45	Manually Updating the DeDupe Appliance	18
Extending Appliance Disks	22	MOTD	41
F		N	
File Level Restores	31	Network IP address	27
limitations	31	NFS	7, 19
Filter	29	NFS External Storage	19
FLR	31	NFS storage	
supported file systems	31	configuring	19
G		Null Blocks	35
Group drop-down	37	O	
H		Operational	27
Hostname	27	Options page	37
I		Other Counts	36
implementation scenarios	47	Other Requirements	11
Initial Backup	8	P	
K		Passwords	9
KEEP_MAXIMUM_WEEKS	38	PHD Virtual	7
KEEP_MINIMUM_UNIQUE	38	phd_daemon.log	22
Knowledge Center	35	phd_delete.log	22
		phd_smb.log	22
		phd_update.log	22
		phd_verify.log	22

Please change the ROOT Password	24	STATS	37
		Stats Page	34
Q		Storage	27-28
Quit (Q)	24	Backup	27-28
		Database & Logs	28
R		System	28
Re-Check Bad Backups	36	Storage Requirements	11
Real Data Blocks	35	Summary Page	27
ReCalc	34	Support	3
Remake Block Folders	20	Supported Browsers	11
Remaking Block Folders	20	SYSTEM	37, 41
Remote NFS Storage		SYSTEM-SAMBA-ENABLED	45
scenario	50	System Config Settings	24
RERUN_AFTER_X_HOURS	40	System Options	41
Restart Web Services	24	System Page	27
restoring directories	32	System Status	27
restoring individual files	31		
RUN_STATS	41	T	
		THIN_AFTER_DAYS	38, 40
S		THIN_AFTER_MONTHS	39-40
SAMBA	41	THIN_AFTER_WEEKS	39-40
SERVER_STRING	42, 45	THIN_KEEP_WHICH_DAY	39
setting a lock	27	THIN_KEEP_WHICH_MONTH	39-40
Show status from TTY1, TTY10, TTY11	24	THIN_KEEP_WHICH_WEEK	39
Show Storage Status	24	THIN_PER_DAY	38
Size of Purged Blocks	36	THIN_PER_MONTH	39
SMB	7, 19	THIN_PER_WEEK	39
SMB External Storage	19	THIN_SPECIAL	38
SMB storage		TITLE	41
configuring	19	Total Unique Blocks	35
Special	30	Total Unique Blocks Compressed	35
START_TIME	41	Trim Report	34
Stat Values	35	Trim the Herd	37, 40
Statistics Options	40	TRIM_PER_DAY	40
Stats	34	Type	29

U

UNCOMPRESSED_BROWSEABLE	42, 45
UNCOMPRESSED_COMMENT	42
UNCOMPRESSED_SHARE_NAME	42, 45
UNCOMPRESSED_COMMENT	45
Users	
add new	43
change password	43
Users Page	43

V

Verify Repair Appliance Database	24
Version	27
View Appliance Logs	24
Virtual Full Backups	7
VM Memory	11
VMDK	19
VMDK Attached Storage	20
VMDK storage	
configuring	20
VMDKs	7
VMs	29
VMs Page	29
VMware version	11

W

WORKGROUP	42, 45
-----------	--------