# VMware vSphere® Data Protection Advanced™ 5.5

## EMC® Data Domain® Integration

TECHNICAL WHITE PAPER

vmware®

## Table of Contents

Data Domain Configuration . . . . . 4
Architecture Overview . . . . . 4
vSphere Data Protection Advanced Client Support . . . . . 5
Best Practices . . . . . 5
Data Domain Limitations . . . . . 5
Backup . . . . . 6
Restore . . . . . 6
Security . . . . . 6
Encryption . . . . . 6
User Access . . . . . 6
Preintegration Requirements . . . . . 7
Network Throughput . . . . . 7
Network Configuration . . . . . 7
NTP Configuration . . . . . 8
Licensing . . . . . 8
Port Usage and Firewall Requirements . . . . . 8
Capacity . . . . . 8
Data Domain System Streams . . . . . 8
Existing Backup Products in Use with Data Domain . . . . . 9
Preparing the Data Domain System for vSphere Data Protection Advanced Integration . . . . . 9
Adding a Data Domain System . . . . . 10
Prerequisites . . . . . 10
Procedure . . . . . 10
Changing the Maximum Streams Value . . . . . 12
Editing Data Domain . . . . . 12
Deleting the Data Domain System from the vSphere Data Protection Advanced Appliance . . . . . 13
Procedure . . . . . 13
Backups with vSphere Data Protection Advanced and Data Domain . . . . . 15
How Backups Work with vSphere Data Protection Advanced and Data Domain . . . 15
Supported Backup Types . . . . . 16
Canceling and Deleting Backups . . . . . 16
Selecting a Data Domain Target for Backups . . . . . 16

**Table of Contents (continued)**

Replication Control . . . . . 17
Replication Data Flow . . . . . 17
Replication Schedule . . . . . 17
Replication Configuration . . . . . 17
Replication Monitoring with vSphere Data Protection Advanced . . . . . 17
Server Maintenance Activity Monitoring . . . . . 18
Monitoring Data Domain from the vSphere Data Protection Advanced Appliance . . . 18
Monitoring Using the vSphere Web Client . . . . . 18
Monitoring Using the vSphere Data Protection Configuration UI . . . . . 18
Data Domain Capacity Monitoring . . . . . 19
Reclaiming Storage on a Full Data Domain System . . . . . 19
Common Problems and Solutions . . . . . 22
Backup Fails if the Data Domain System Is Offline . . . . . 22
Rolling Back After Deleting a Data Domain System . . . . . 22

# Data Domain Configuration

VMware vSphere® Data Protection Advanced™ typically is implemented to create image-level backups of virtual machines, including virtual servers, databases with low change rates, and others. In addition, it can utilize application plug-ins to back up Microsoft Exchange Server, SQL Server, and SharePoint Server.

vSphere Data Protection and EMC® Data Domain® system integration enable the following:

- Data Domain systems to be a backup target for vSphere Data Protection Advanced backups
- The target destination of backup data, which is set during the creation of a backup job
- Transparent user interaction to the backup destination (vSphere Data Protection Advanced or Data Domain)

# Architecture Overview

A Data Domain system performs deduplication through Data Domain Operating System (DD OS) software. vSphere Data Protection Advanced source-based deduplication to a Data Domain system is facilitated through the use of the Data Domain Boost library.

vSphere Data Protection Advanced uses the Data Domain Boost library through API-based integration to access and manipulate directories, files, and so on, contained on the Data Domain file system. The Data Domain Boost API provides vSphere Data Protection Advanced with visibility into some of the properties and capabilities of the Data Domain system. This enables vSphere Data Protection Advanced to control backup images stored on Data Domain systems. This also enables it to manage maintenance activities and to control replication to remote Data Domain systems. Data Domain Boost is installed on the vSphere Data Protection Advanced appliance during the addition of a Data Domain system.

Figure 1 depicts a high-level architecture of the combined vSphere Data Protection Advanced and Data Domain solution. With integration of these two components, users can specify whether a particular backup policy targets a vSphere Data Protection Advanced appliance or a Data Domain system.

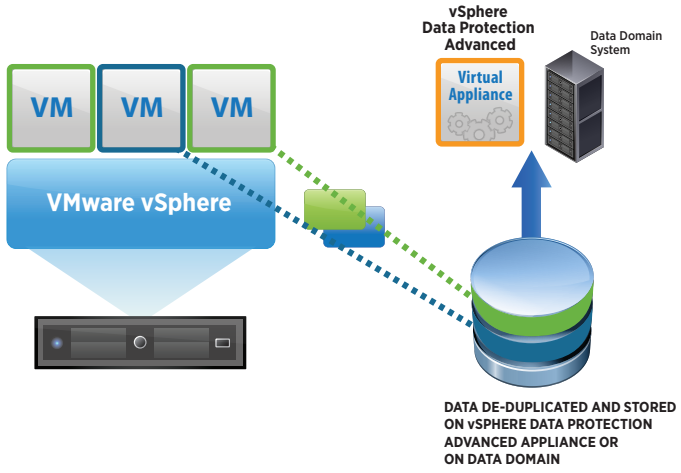


**Figure 1.** vSphere Data Protection Advanced and Data Domain Solution

When the vSphere Data Protection Advanced appliance is selected as the target for backup storage, it performs deduplication segment processing.

When a Data Domain system is selected as the backup target, backup data is transferred to the Data Domain system. The related metadata that is generated is simultaneously sent to the vSphere Data Protection Advanced appliance for storage. The metadata enables the vSphere Data Protection Advanced appliance to perform restore operations directly from the Data Domain system.

# vSphere Data Protection Advanced Client Support

vSphere Data Protection Advanced and Data Domain system integration supports the following plug-ins:

- vSphere Data Protection Advanced plug-in for Exchange Server Volume Shadow Copy Service (VSS)
- vSphere Data Protection Advanced plug-in for SharePoint Server VSS
- vSphere Data Protection Advanced plug-in for SQL Server VSS

# Best Practices

**What are the limitations on vSphere Data Protection Advanced with a Data Domain system attached?**

We suggest protecting as many as 25 virtual machines per terabyte of capacity on a vSphere Data Protection Advanced appliance. This variable is dependent upon the size of the virtual machines, the typical change rate, and the amount of data on each virtual machine. With these considerations, we suggest protecting as many as 200–400 virtual machines per vSphere Data Protection Advanced appliance backing up to a Data Domain system.

- 1 vSphere Data Protection Advanced per DD160 and DD620
- 2 vSphere Data Protection Advanced per DD2500 and DD4xxx
- 3 vSphere Data Protection Advanced per DD7200 and DD990

**What size vSphere Data Protection Advanced appliance do I need if I want all my backups to go to Data Domain?**

With a Data Domain back end, the vSphere Data Protection Advanced appliance is storing only the metadata. It has been determined that a 16TB Data Domain system requires only a 256GB disk on the front end. Therefore, a 0.5TB vSphere Data Protection Advanced appliance should handle a standard Data Domain system; to be sure, use a 1TB vSphere Data Protection Advanced appliance with a 64TB Data Domain system.

**I have many images and pictures in my virtual machines. Should I set the vSphere Data Protection Advanced appliance or the Data Domain system as the destination for these backups?**

The vSphere Data Protection Advanced appliance does not deduplicate images and pictures very well, so vSphere Data Protection Advanced with a Data Domain back end is a better option.

# Data Domain Limitations

The following are the current defined limitations for the use of a Data Domain system with a vSphere Data Protection Advanced appliance.

- If a vSphere Data Protection Advanced appliance has a Data Domain system attached to it, its **Import existing storage** functionality cannot be used for its VMDKs.
- Only one Data Domain system at a time can be attached to a vSphere Data Protection Advanced appliance.
- The Data Domain system cannot be deleted from the vSphere Data Protection Advanced configuration user interface (UI). Use the manual steps defined in “Deleting the Data Domain System from the vSphere Data Protection Advanced Appliance” on page 13 to delete a Data Domain system.

- If the Data Domain and vSphere Data Protection Advanced connection is broken, the vSphere Data Protection Advanced appliance does not monitor the Data Domain system. A failure of the integrity check, hfscheck, or backups is among the indications that the connection between the appliances might have been broken.
- The Data Domain system and vSphere Data Protection Advanced appliance cannot be upgraded if the connection between them is broken.

# Backup

During a backup, the vSphere Data Protection Advanced appliance generates a request for the backup destination. If this request includes the option to use a Data Domain system as the destination, backup data is stored on the Data Domain system. Metadata is stored on the vSphere Data Protection Advanced appliance.

# Restore

The process of data recovery is transparent to the backup administrator. The backup administrator uses the same vSphere Data Protection Advanced recovery processes that are native to current vSphere Data Protection Advanced implementations.

# Security

The following topics provide details on security for encryption and user access in a vSphere Data Protection Advanced environment with an attached Data Domain system.

## Encryption

When using a vSphere Data Protection Advanced appliance with a Data Domain system attached, there are two potential backup data streams. If the backup data is being written to the vSphere Data Protection Advanced appliance, it is always compressed and encrypted. If the backup data is being routed to the Data Domain system, the Data Domain Boost utility does not encrypt it during transmission over the network to the Data Domain system.

## User Access

Use caution when granting users access to the Data Domain system. A user should not be able to directly access the Data Domain system and manually delete data.

Backup data cannot be migrated directly from the vSphere Data Protection Advanced appliance to the Data Domain system. To start using the Data Domain system instead of the vSphere Data Protection Advanced appliance as the target for backing up a virtual machine or appliance, edit the backup job and define the destination as the Data Domain system; start performing backups to the Data Domain system. After changing the backup target to the Data Domain system, the next backup run will be a full backup.

After successfully performing a backup to the Data Domain system, users can delete the earlier backups from the vSphere Data Protection Advanced appliance. Refer to “Deleting a Backup Job” in the *vSphere Data Protection Administration Guide* for information on how to delete backups.

# Preintegration Requirements

Before integrating a Data Domain system with a vSphere Data Protection Advanced appliance, review the following topics:


- "Network Throughput" on page 7
- "Network Configuration" on page 7
- "NTP Configuration" on page 8
- "Licensing" on page 8
- "Port Usage and Firewall Requirements" on page 8
- "Capacity" on page 8
- "Data Domain System Streams" on page 8
- "Existing Backup Products in Use with Data Domain" on page 9


*NOTE: This section assumes that the vSphere Data Protection Advanced appliance and the Data Domain system have been installed and configured.*

## Network Throughput

The vSphere Data Protection Advanced appliance and all Data Domain systems must be on the same local network. Do not connect the vSphere Data Protection Advanced appliance and Data Domain systems over a wide area network (WAN). Configurations that use a WAN are not supported.

Use vSphere Data Protection Advanced appliance replication over a WAN to replicate data from source vSphere Data Protection Advanced appliance and Data Domain systems to target vSphere Data Protection Advanced appliances, provided they also have a Data Domain system attached.

Before integrating a Data Domain system with a vSphere Data Protection Advanced appliance, ensure that enough network bandwidth is available. To obtain the maximum throughput available on a Data Domain system—for restores, level-zero backups, and subsequent incremental backups after a level-zero backup—verify that the network infrastructure provides more bandwidth than that required by the maximum throughput of the Data Domain system.

## Network Configuration

Configure, or verify, the following network configuration:

- Assign a fully qualified domain name (FQDN) to the Data Domain system.
- Do not use IP addresses in place of host names when registering a Data Domain system. This can limit the ability to route optimized duplication traffic exclusively through a registered interface.
- Ensure that the Domain Name System (DNS) is properly configured on the Data Domain system.
- Ensure that forward and reverse DNS lookups work between the following systems:
  - vSphere Data Protection Advanced appliance
  - Data Domain system
  - Backup and restore clients
  - VMware® vCenter Server™
  - vSphere hosts
- Use host files to resolve host names to nonroutable IP addresses.
- Do not create secondary host names to associate with alternate or local IP interfaces.

## NTP Configuration

Configure the vSphere Data Protection Advanced appliance, Data Domain system, vCenter Server, and vSphere hosts to use the same Network Time Protocol (NTP) server.

## Licensing

Ensure that the environment meets the licensing requirements in the following table.

| PRODUCT | LICENSING REQUIREMENT |
|---|---|
| vSphere Data Protection Advanced appliance | The vSphere Data Protection appliance requires a valid VMware ESXi™ host license (minimum of VMware vSphere Essentials Plus Kit). In addition a valid vSphere Data Protection Advanced license must be applied through the vSphere Data Protection configuration UI to enable the advanced feature set. |
| Data Domain system | The Data Domain Boost license must be installed on the Data Domain system. |

## Port Usage and Firewall Requirements

To enable communication between the vSphere Data Protection Advanced appliance and the Data Domain system, review and implement the port usage and firewall requirements in the following documents:

- "vSphere Data Protection Port Usage" in the *vSphere Data Protection Administration Guide*
- "Port Requirements for Allowing Access to Data Domain System Through a Firewall," available on the Data Domain support portal at https://my.datadomain.com

## Capacity

Carefully assess backup storage needs when evaluating how much data to store on the Data Domain system and the vSphere Data Protection Advanced appliance. Include estimates of data that is sent to the Data Domain system from any other servers.

When the Data Domain system reaches its maximum storage capacity, no further backups to the Data Domain system occur until additional capacity is added or old backups are deleted.

"Data Domain Capacity Monitoring" on page 19 provides details on how to monitor capacity.

## Data Domain System Streams

Each Data Domain system has a soft limit to the maximum number of connection and data streams that can be sustained simultaneously while maintaining performance. The number of streams varies depending on the Data Domain system model. For example, the EMC Data Domain DD990 can support 540 backup streams; the EMC Data Domain DD620 can support 20 backup streams.

As a default, the vSphere Data Protection Advanced appliance is configured to use a maximum of 16 streams. To override the maximum number of streams, on the vSphere Data Protection Advanced appliance, change the `/usr/local/vdr/etc/vdp-options.properties` file. Add the `property com.vmware.vdp.option.datadomain.maxstreamsoverride=XX` (where `XX` is the new maximum number of streams) and save the file. Consult the *Data Domain Operating System Administration Guide* for your Data Domain system to review current recommended stream settings. These changes will not take effect until the Data Domain is edited or a new Data Domain is added to the vSphere Data Protection Advanced appliance. These changes will be reflected in the `/usr/local/avamar/var/ddr_info` file or can be seen by running the `ddrmaint read-ddr-info` command.

Setting this value forcibly overrides the maximum-streams value for every add or edit operation of a Data Domain for a given vSphere Data Protection Advanced appliance.

These changes will not take effect until the existing Data Domain is edited or a new one is added to the vSphere Data Protection Advanced appliance. The changes will be reflected in the `ddrmaint read-ddr-info` file.

### Existing Backup Products in Use with Data Domain

Data Domain systems can use other third-party backup-and-archiving software. The vSphere Data Protection Advanced appliance does not have exclusive access to the Data Domain system. Ensure that proper sizing is evaluated if the system is shared with other software products. The vSphere Data Protection Advanced appliance makes no use of the native Data Domain system snapshot and replication features.

Replication occurs through the Data Domain Boost SDK library by using copying and cloning. However, other third-party products can use the native Data Domain system snapshot and replication features. In this case, a snapshot of an entire Data Domain system or a replication of an entire Data Domain system includes the vSphere Data Protection Advanced appliance data.

## Preparing the Data Domain System for vSphere Data Protection Advanced Integration

To support vSphere Data Protection Advanced and Data Domain integration, ensure that the environment meets the Data Domain system requirements listed in the following table.

| DATA DOMAIN FEATURE OR SPECIFICATION | REQUIREMENT FOR USE WITH THE vSPHERE DATA PROTECTION ADVANCED APPLIANCE |
|---|---|
| Data Domain Operating System (DD OS) | vSphere Data Protection Advanced integration requires DD OS 5.3 with DD OS 5.4.x. |
| Data Domain Boost | vSphere Data Protection Advanced integration requires Data Domain Boost 2.6.x. Data Domain Boost software enables backup servers to communicate with storage systems without the need for Data Domain systems to emulate tape.<br>There are two components to Data Domain Boost: one runs on the backup server; the other runs on the Data Domain system.<br>The component that runs on the backup server, Data Domain Boost libraries, is integrated into the vSphere Data Protection Advanced appliance. Data Domain Boost software is an optional product that requires a license to operate on the Data Domain system. |
| Data Domain device type | The vSphere Data Protection Advanced appliance supports any Data Domain system that supports the execution of the required DD OS version. |
| Data Domain Boost user account | The Data Domain Boost library uses a unique login account name created on the Data Domain system. This account name is known as the Data Domain Boost account. Only one Data Domain Boost account exists per Data Domain system.<br>If the account is renamed or the password is changed, these changes must be immediately updated on the vSphere Data Protection Advanced appliance by editing the Data Domain configuration options. Failure to update the Data Domain Boost account information can potentially yield integrity check errors or backup-and-restore problems. The Data Domain Boost account must have administrator privileges. |

Before adding a Data Domain system to the vSphere Data Protection Advanced configuration, prepare the Data Domain system by enabling Data Domain Boost and creating a Data Domain Boost user account for the vSphere Data Protection Advanced appliance to use to access the Data Domain system for backups and restores, as well as for replication if applicable.

To prepare the Data Domain system:

1. Disable Data Domain Boost on the Data Domain system by logging into the Data Domain command-line interface (CLI) as an administrative user and typing the following:

   ```
   ddboost disable
   ```

2. Create a Data Domain Boost account and password:

   a. Create the user account with administrative privileges by typing the following:

      ```
      user add USER role admin
      ```

      where **USER** is the username for the new account

   b. Set the new account as the Data Domain Boost user by typing the following:

      ```
      ddboost set user-name USER
      ```

      where **USER** is the username for the account

   c. Enable Data Domain Boost to allow the changes to take effect by typing the following command:

      ```
      ddboost enable
      ```

*NOTE: If the Data Domain Boost account name or password changes, make sure to edit the Data Domain system configuration in the vSphere Data Protection configuration UI. Otherwise, all backups, restores, and maintenance activities will fail.*

# Adding a Data Domain System

A Data Domain system performs deduplication through DD OS software. When a Data Domain system is selected as the backup target, backup data is transferred to it. Backing up to a Data Domain system is supported in vSphere Data Protection Advanced mode only, and only one Data Domain system can be configured.

## Prerequisites

- The vSphere Data Protection storage disks are distributed across the available datastore locations and the disks are validated.
- Data Domain version 5.3 is the minimum supported version.

## Procedure

1. To access the vSphere Data Protection configuration utility, open a Web browser and type the following:

   ```
   http:<vSphere Data Protection appliance IP>:8543/vdp-configure
   ```

2. Click the **Storage** tab.

   The storage summary displays statistics about the total usable storage and available capacity for the Data Domain and for each datastore.

3. From the **Action** list, select **Add Data Domain**.

   The **Host Configuration** dialog box appears.

4. Specify Data Domain system information:

   a In the Data Domain FQDN or IP box, enter the FQDN or IP address of the Data Domain system to add.

*NOTE: Do not use an IP address or a secondary host name that associates with alternative or local IP interfaces. It can limit the ability of the vSphere Data Protection Advanced appliance to route optimized deduplication traffic.*

   b. In the **Data Domain Boost User Name** box, type the name of the Data Domain Boost account for vSphere Data Protection Advanced to use to access the Data Domain system for backups, restores, and replication.

   c. In the **Password** box, type the password for the account for vSphere Data Protection Advanced to use to access the Data Domain system for backups, restores, and replication.

   d. In the **Confirm Password** box, type the password again to verify it.

5. To configure SNMP, click **Next**.

   The **SNMP** dialog box appears. The following configurations are among the SNMP options for vSphere Data Protection Advanced and Data Domain system integration:

   - The **Getter/Setter Port Number** text box lists the port on the Data Domain system from which to receive and on which to set SNMP objects. The default value is 161.
   - The **SNMP Community String** text box lists the community string that vSphere Data Protection Advanced uses for read-only access to the Data Domain system.
   - The **Trap Port Number** text box lists the trap port. The default value is 163.

6. Click **Next**.

   The **Ready to Complete** dialog box appears.

7. Click **Add** to save your Data Domain system configuration.

   A successful **Add Data Domain** operation causes the following UI changes to occur:

   The system creates a new checkpoint, which takes approximately 10 minutes.

   Data Domain system information appears on the vSphere Data Protection appliance in the following locations:

   - **Backup** tab – Data Domain is available as the storage target in the **Create a new backup** job wizard.
   - **Restore** tab – This displays Data Domain in the **Name** column of the **Restore backup wizard**.
   - **Reports** tab – This provides backup status reports for the Data Domain system.
   - Storage summary – This displays statistics on total usable storage and available capacity for the Data Domain system. Refer to "Viewing the Storage Configuration" on page 27 of the *vSphere Data Protection Administration Guide* for details.
   - Email reporting – This displays a summary of the Data Domain configuration.

   *NOTE: When a Data Domain system is added to the vSphere Data Protection Advanced configuration, the vSphere Data Protection Advanced appliance creates an MTree on the Data Domain system for itself. The MTree refers to the directory created within the Data Domain Boost path. Data Domain systems support a maximum of 100 MTrees. If this limit is reached, the Data Domain system cannot be added to the vSphere Data Protection Advanced configuration.*

# Changing the Maximum Streams Value

As a default, the vSphere Data Protection Advanced appliance is configured to use a maximum streams value of 16.

Perform the following steps to modify the number of streams associated with a Data Domain system from the vSphere Data Protection Advanced appliance. The changes applied using these steps will take effect only during subsequent edits of, or the addition of a Data Domain system to, the vSphere Data Protection Advanced appliance.

1. Go to the command line of the vSphere Data Protection Advanced appliance (either with SSH/PuTTY or terminal of the appliance).

   `cd /usr/local/vdr/etc/`

2. Use your preferred file editor to edit the `vdp-options.properties` file.
3. Insert the `com.vmware.vdp.option.datadomain.maxstreamsoverride=XX` field, where `XX` is the maximum number of streams for the Data Domain system.
4. Save the modified file.
5. Add or edit a Data Domain system. Allow five minutes for the appropriate process to run.

The `ddrmaint read-ddr-info` file should now contain a "max-streams" attribute with the previously configured value.

# Editing Data Domain

1. To access the vSphere Data Protection configuration utility, open a Web browser and type the following:

   `http:<vSphere Data Protection appliance IP>:8543/vdp-configure`

2. Click the **Storage** tab.

   The storage summary displays statistics on the total usable storage and available capacity for the Data Domain system and each datastore.

3. From the **Action** list, select **Edit Data Domain**.

   The **Host Configuration** dialog box appears.

4. Edit the settings for the Data Domain system as necessary. "Adding a Data Domain System" on page 10 provides details on each setting in the dialog box.
5. Click **Next**.
6. After the edits are completed, click **Finish**.

*NOTE: If the Data Domain host name, the Data Domain Boost username, or the Data Domain Boost password is edited, the system automatically creates a new checkpoint, which takes approximately 10 minutes. For instructions, refer to the "Rolling Back an Appliance" section of the* vSphere Data Protection Administration Guide. *A rollback will fail if it is performed to a checkpoint with the outdated Data Domain system name or Data Domain Boost information.*

# Deleting the Data Domain System from the vSphere Data Protection Advanced Appliance

Before deleting the Data Domain system from the vSphere Data Protection Advanced appliance, note the following:

- Prior to deleting the Data Domain system, all restore points stored on it must be deleted via the vSphere Web Client.
- No backup jobs with the Data Domain system as the backup can exist. If any such jobs exist with the Data Domain configured as the destination, either edit them to set a new destination or delete them.
- After the restore points have been checked and the backup jobs have been verified, best practice is to run an integrity check from the **Configuration** tab of the vSphere Data Protection Advanced appliance.
- Remove the Data Domain system from vSphere Data Protection Advanced using the CLI. See the detailed instructions that follow.
- After the Data Domain system has been deleted, run another integrity check from the vSphere Data Protection Advanced UI to verify that the Data Domain system is invalid.

*NOTE: VMware Knowledge Base article 2063806 provides detailed information about deleting a Data Domain system. This is an internal article, so contact technical support for assistance.*

## Procedure

1. Prior to deleting the Data Domain system, all restore points stored on it must be deleted via the vSphere Web Client. To delete restore points:
   a. Navigate to the **Restore** tab of the vSphere Data Protection Advanced plug-in.
   b. Select the **Manual restore** tab on the navigation bar.
   c. For clients that have been backed up to the Data Domain system, remove all restore points where the location shows that they are stored on the Data Domain server.
2. In addition, no backup jobs can have the Data Domain system as a destination. If any such jobs exist, either edit them to set a new destination or delete them.
3. After the restore points have been checked and the backup jobs have been verified, best practice is to run an integrity check from the **Configuration** tab of the vSphere Data Protection plug-in.
4. After the integrity check and validation of the integrity check have been completed, attempt to remove the Data Domain system from the vSphere Data Protection Advanced appliance. This is accomplished from the command line with the following steps.
   a. SSH or PuTTY into your vSphere Data Protection Advanced appliance.
   b. Run the status.dpn command and verify that the **Last checkpoint** and **Last hfscheck** indicate that they have completed. If they have not, repeat this step until they indicate that they have completed.

c. Run the **mccli server show-prop** command. This will return results similar to the following:

Attribute Value

---

| | |
|---|---|
| State | Full Access |
| Active sessions | 0 |
| Total capacity | 575.9 GB |
| Capacity used | 0 bytes |
| Server utilization | 0.0% |
| Bytes protected | 0 bytes |
| Bytes protected quota | Not configured |
| License expiration | Never |
| Time since server initialization | 1 days 20h:58m |
| Last checkpoint | 2013-10-10 09:03:48 MDT |
| Last validated checkpoint | 2013-10-09 09:02:16 MDT |
| System name | gs-pod187.test.domain |
| **System ID** | **1381255529@00:50:56:86:46:10** |
| HFSAddr | gs-pod187.test.domain |
| HFSPort | 27000 |

The **System ID** contains a number, then the @ icon, and then the MAC address of the vSphere Data Protection Advanced appliance. Capture the number that precedes the @ icon. The Data Domain system refers to this as the DPN ID.

5. Run the **ddrmaint has-backups –dpnid=XXXX –ddr-server=DDRSERVER | grep hasbackups** command, where **XXXX** is the DPN ID captured in step 4c and **DDRSERVER** is either the host name or the IP address of the DDR server. There is a space in the grep between the single quote and the word "hasbackups."

   This should return one of the two following results:

   **hasbackups="true" or hasbackups="false"**

6. If the information returned is **hasbackups="true",** check whether step 1 and step 2 must be repeated. After verifying that step 1 and step 2 have been repeated, repeat step 5.

7. If step 5 still shows **hasbackups="true"**, proceed to step 7a. Otherwise, run step 16. If you have attempted to remove backups from the Data Domain system using the vSphere Data Protection UI and the Data Domain system still indicates that the backup data is present, you must mount the Data Domain data partition to a Linux virtual machine to clear out the data directory. If no Linux virtual machine exists within, use the vSphere Data Protection Advanced appliance for the next steps.

   By default, all of the data for backups on a Data Domain system is stored under a single logical storage unit (LSU) on the system. The LSU for vSphere Data Protection Advanced is named Avamar-<DPNID> and is located under **/data/col1**.

If you are unable to access the file system from the DD OS interface, you must grant remote access to the LSU. To do so, you must access the Data Domain system remotely with the following steps.

a. PuTTY or SSH to your Data Domain system.

b. Run the `nfs add /data/col1 <IP of Linux VM>` command.

This should return the results of NFS export for `/data/col1` added. If this does not return the expected results, type `nfs help` for a *man page* on the command. If this returns the expected results, you can exit the SSH or PuTTY session.

8. PuTTY or SSH to the Linux virtual machine used in step 4b as the root user.
9. Run the `mkdir /mnt/DataDomain01` command.
10. Run the `mount <IP of DD>:/data/col1 /mnt/DataDomain01` command.
11. Run the `ls –ltr /mnt/DataDomain01/avamar-<DPNID>` command, where `DPNID` is the value captured in step 4c. This should show subdirectories where the vSphere Data Protection Advanced backups are stored.
12. Run the `rm –rf /mnt/DataDomain01/avamar-<DPNID>/*` command, where `DPNID` is the value captured in step 4c. This will remove all data from the vSphere Data Protection Advanced backups.
13. Repeat step 11 to verify that all data has been removed.
14. Exit the Linux virtual machine.
15. PuTTY or SSH to the vSphere Data Protection Advanced appliance.
16. Run the `mccli dd delete --name=<DD IP or hostname> --force=true` command.
17. After the Data Domain system has been deleted, run an integrity check again from the vSphere Data Protection UI, because the old checkpoints with the Data Domain information will be invalid.

# Backups with vSphere Data Protection Advanced and Data Domain

The following topics describe vSphere Data Protection Advanced and Data Domain system backups:

- How backups work with vSphere Data Protection Advanced and Data Domain
- Selecting a Data Domain target for backups

## How Backups Work with vSphere Data Protection Advanced and Data Domain

During a backup, the vSphere Data Protection Advanced appliance sends a backup request to the vCenter Server. If the backup request includes the option to use a Data Domain system as the target, backup data is stored on the Data Domain system and metadata is stored on the vSphere Data Protection Advanced appliance.

The following topics provide additional details on how backups work with vSphere Data Protection Advanced and Data Domain.

### Where Backup Data Is Stored

All data for a backup is stored under a single dedicated MTree on a single Data Domain system.

### How vSphere Data Protection Advanced Appliance Manages Backup Data

During a backup, the vSphere Data Protection Advanced appliance sends the metadata for the backup from the client to the vSphere Data Protection Advanced data partitions. This process enables the vSphere Data Protection Advanced appliance to manage the backup even though the data is stored on a Data Domain system.

The vSphere Data Protection Advanced appliance does not store the original path and filename for a file on the Data Domain system. Instead, it uses unique filenames on the Data Domain system.

### Supported Backup Types

Users can perform full backups, incremental backups, and differential backups, as well as VMware backups with Changed Block Tracking enabled.

Store the full backup for a client and all subsequent incremental and differential backups on either the vSphere Data Protection Advanced appliance or a single Data Domain system.

The vSphere Data Protection Advanced appliance does not support the following:

- Full backup on a Data Domain system and incremental or differential backups on the vSphere Data Protection Advanced appliance
- Full backup on the vSphere Data Protection Advanced appliance and incremental or differential backups on a Data Domain system
- Full backup on one Data Domain system and incremental or differential backups on another Data Domain system

If the device on which backups for a client are stored is changed, a full backup must be performed before any further incremental or differential backups are done.

### Canceling and Deleting Backups

If a backup is canceled while it is in progress, the vSphere Data Protection Advanced appliance will delete the backup data that was written to the Data Domain system. The data deletion will take place during the next cycle of the vSphere Data Protection Advanced appliance garbage collection process.

If a backup is deleted in vSphere Data Protection Advanced, it will be deleted from the Data Domain system during the next cycle of the vSphere Data Protection Advanced appliance garbage collection process.

"Deleting a Backup Job" in the *vSphere Data Protection Administration Guide* provides instructions on how to cancel or delete a backup.

# Selecting a Data Domain Target for Backups

After the vSphere Data Protection Advanced appliance and the Data Domain system have been integrated, any backup target for the vSphere Data Protection Advanced appliance can use the Data Domain storage as the **Destination** in the **Create a new backup job** work flow, as is shown in Figure 2.

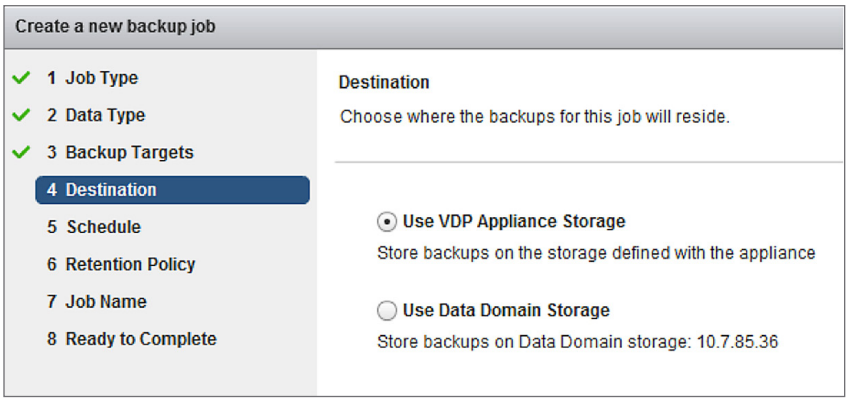


**Figure 2.** Create a New Backup Job Wizard – Destination Page

Use the **Edit a backup job** wizard to change the destination for a backup job. See "Editing Data Domain" on page 12 for more information.

*NOTE: If the destination of a backup job is modified, the next backup performed will be a full backup because the new destination will not have the previous full backup data stored on it.*

# Replication Control

When a vSphere Data Protection Advanced appliance with a Data Domain system attached replicates backups, the Data Domain system replicates the data between systems. The vSphere Data Protection Advanced replication feature is not used in this instance.

Configure and monitor replication on the vSphere Data Protection Advanced appliance. The replication activity can be monitored through the Data Domain system by checking the Data Domain Boost activity. Refer to the *Data Domain OS Administration Guide* for instructions on how to monitor this activity.

Do not use Data Domain replication functionality to initiate replication of data to another Data Domain system that is configured for use with vSphere Data Protection Advanced. When Data Domain replication is utilized, the replicated data will not refer to the associated vSphere Data Protection Advanced appliance, because the metadata stored on the vSphere Data Protection Advanced appliance will not have been replicated.

## Replication Data Flow

vSphere Data Protection Advanced replicates the data directly from one Data Domain system to another. The replication process examines each backup to be replicated; if it determines that the backup data is stored on a Data Domain system, it will issue a request to replicate the data from the source Data Domain system to the target Data Domain system via Data Domain Boost. In this instance, the Data Domain systems are responsible for the replication of the data. This is analyzed for each backup being replicated.

## Replication Schedule

The replication of vSphere Data Protection Advanced data on a Data Domain system occurs on the vSphere Data Protection Advanced replication schedule. Replication of data on the Data Domain system cannot be scheduled separately from that on the vSphere Data Protection Advanced appliance.

## Replication Configuration

Configure replication through the vSphere Web Client when using a Data Domain system as a backup target for vSphere Data Protection Advanced. Refer to the *vSphere Data Protection Administration Guide* for more information on configuring vSphere Data Protection Advanced replication.

NOTE: When replicating from a vSphere Data Protection Advanced appliance with a Data Domain system attached, only a vSphere Data Protection Advanced appliance with a Data Domain system attached can be the target for replication.

## Replication Monitoring with vSphere Data Protection Advanced

To monitor replication activity with the vSphere Data Protection Advanced appliance, including replication activities associated with a Data Domain system, perform the following steps:

1. In the vSphere Web Client, log in to the vSphere Data Protection plug-in.
2. Click the **Replication** tab.

- The **Replication** tab displays all replication jobs, the last runtime, and the next scheduled runtime.

- Highlighting a replication job displays the destination server and the clients included in the **Replication job details** frame.
- When the **Replication** column is checked, the **Reports** tab displays the replication job and the last replication runtime for each protected client.

# Server Maintenance Activity Monitoring

The vSphere Data Protection Advanced appliance performs the system maintenance operations for backup data on the Data Domain system, including HFS checks, checkpoints, rollbacks, garbage collection, and secure backup deletion.

The `ddrmaint` utility implements all required operations on the Data Domain system for the vSphere Data Protection Advanced appliance. It is installed on the vSphere Data Protection Advanced appliance during the addition of the Data Domain system to the appliance.

The `ddrmaint` utility logs all maintenance activities on the vSphere Data Protection Advanced appliance in the `ddrmaint.log` file. This log file can be located in the `/usr/local/avamar/var/ddrmaintlogs` directory. The `ddrmaint.log` file is rotated when it reaches 25MB.

# Monitoring Data Domain from the vSphere Data Protection Advanced Appliance

To review high-level information about the Data Domain system attached to a vSphere Data Protection Advanced appliance, look in either the vSphere Web Client or the vSphere Data Protection configuration UI.

## Monitoring Using the vSphere Web Client

1. In the vSphere Web Client, open the vSphere Data Protection plug-in.
2. Navigate to the **Configuration** tab.

   In the Data Domain storage summary, the following information is displayed:

   - Data Domain system FQDN or IP address
   - **Capacity** of the Data Domain system
   - **Free Space** on the Data Domain system
   - **Used Capacity** on the Data Domain system

## Monitoring Using the vSphere Data Protection Configuration UI

1. In the vSphere Data Protection configuration UI, navigate to the **Storage** tab.
2. In the **Storage Summary** section, the following information is displayed:

   - Data Domain host name or IP address
   - **Total Usable Storage**
   - **Storage Available**
   - **Capacity Consumed** (in percentage)

### Data Domain Capacity Monitoring

Check the capacity of the Data Domain system by monitoring the vSphere Web Client or the vSphere Data Protection Advanced configuration UI.

1. In the vSphere Web Client, open the vSphere Data Protection plug-in and navigate to the **Configuration** tab to view a capacity summary for the Data Domain system.
2. In the vSphere Data Protection Advanced configuration UI, open the **Storage** tab to view a storage summary for the Data Domain system.

When the Data Domain system reaches its capacity limit, reclaim space on the device by using the instructions in "Reclaiming Storage on a Full Data Domain System."

*NOTE: When the Data Domain system reaches 99 percent capacity, maintenance operations will fail. Best practice is to limit Data Domain capacity usage to 80.*

# Reclaiming Storage on a Full Data Domain System

If all storage space on a Data Domain system is used, the following issues might occur:

- Backups do not succeed and might not start.
- Operations that change information on the Data Domain system fail, including the deletion of checkpoints, active backups, and expired backups during garbage collection. These operations can fail because they involve directory renames, which are not allowed on a full Data Domain system.

To reclaim the used storage on a full Data Domain system, perform the following steps:

1. Determine the source of the data that is using the storage. The data might be from a specific client, a group of clients associated with a specific vSphere Data Protection Advanced appliance, or a different backup product that stores data on the Data Domain system.
2. Cancel any backups that are in progress. You must do this from the command line of the vSphere Data Protection Advanced appliance.
3. Open an SSH or PuTTY session to the vSphere Data Protection Advanced appliance.
   a. Enter the `su – admin` command.
   b. Enter the `ssh-agent bash` command.
   c. Enter the `ssh-add .ssh/dpnid` command.
   d. Enter the `mccli activity show` command. This should return results similar to the following:

```
admin@gs-pod192:~/>: mccli activity show
0,23000,CLI command completed successfully.

ID              9138660744236309
Status          Running
Error Code      0
Start Time      2013-12-09 09:44 MST
Elapsed         00h:27m:25s
End Time        2013-12-10 09:44 MST
Type            On-Demand Backup
```

**Progress Bytes** 54.3 GB

**New Bytes** 4.2%

**Client Domain** Win2008R2-GSClone /10.7.242.175/VirtualMachines

**ID** 9138660744234709

**Status** Completed

**Error Code** 0

**Start Time** 2013-12-09 09:44 MST

**Elapsed** 00h:02m:51s

**End Time** 2013-12-09 09:47 MST

**Type** On-Demand Backup

**Progress Bytes** 40.0 GB

**New Bytes** <0.05%

**Client Domain** GermanExchange /10.7.242.175/VirtualMachines

**ID** 9138660718256909

**Status** Completed

**Error Code** 0

**Start Time** 2013-12-09 09:39 MST

**Elapsed** 00h:01m:06s

**End Time** 2013-12-09 09:40 MST

**Type** On-Demand Backup

**Progress Bytes** 40.0 GB

**New Bytes** <0.05%

**Client Domain** GermanExchange /10.7.242.175/VirtualMachines

**ID** 9138660744235609

**Status** Completed

**Error Code** 0

**Start Time** 2013-12-09 09:44 MST

**Elapsed** 00h:20m:37s

**End Time** 2013-12-09 10:04 MST

**Type** On-Demand Backup

**Progress Bytes** 40.0 GB

**New Bytes** 2.6%

**Client Domain** ActiveDirectory /10.7.242.175/VirtualMachines \

To run the command that cancels the backup jobs in progress, the appliance password (referred to as the `AppliancePassword` value in the following) is needed. Also note the ID of any running jobs.

e. Enter the `mccli activity cancel --mcsuserid=MCUser --mcpasswd=AppliancePassword --id=XXXXX` command, where ***`AppliancePassword`*** is the appliance password and ***`XXXX`*** is the ID of the running job that is chosen for cancellation. This should return results similar to the following:

```
admin@gs-pod192:~/>: mccli activity cancel --mcsuserid=MCUser
--mcspasswd=Test12345 --id=9138660744236309

0,22205,Backup cancelled via console

Attribute     Value

-----------   ----------------

activity-id   9138660744236309
```

f. Repeat step e for all jobs in the running state.

4. Suspend backups and restores. On the vSphere Data Protection Advanced appliance, this can be done by disabling the proxies from the command line. Prior to running these commands, verify with end users that there are no critical backups or restores that must be performed.
   a. Open an SSH or PuTTY session to the vSphere Data Protection Advanced appliance.
   b. Enter the `service avagent-vmware stop` command.
   c. Suspend server maintenance operations on the vSphere Data Protection Advanced appliance.
5. Open the vSphere Data Protection configuration UI by opening a Web browser and navigating to `https://<vSphere Data Protection_IP_Address>:8543/vdp-configure`.
6. If the **Maintenance services** show as **Running**, click **Stop**.
7. On the Data Domain system, manually delete the existing STAGING, DELETED, or /DELETED directories for the vSphere Data Protection Advanced appliance.
8. Use the Data Domain Enterprise Manager to initiate the Data Domain file system cleaning operation. This process should free enough space to enable vSphere Data Protection Advanced appliance service maintenance operations to complete successfully.
9. Restart server maintenance operations on the vSphere Data Protection Advanced appliance.
   a. Open the vSphere Data Protection configuration UI by opening a Web browser and navigating to `https://<vSphere Data Protection_IP_Address>:8543/vdp-configure`.
   b. If the **Maintenance services** show as **Stopped**, click **Start**.
10. Restart the proxies on the vSphere Data Protection Advanced appliance so that backups and restores can run.
    a. Open an SSH or PuTTY session to the vSphere Data Protection Advanced appliance.
    b. Enter the `service avagent-vmware start` command.

# Common Problems and Solutions

This section lists common problems and solutions for vSphere Data Protection appliance backups stored on a Data Domain system.

## Backup Fails if the Data Domain System Is Offline

If the Data Domain system is offline when a backup starts, the backup might take five minutes or more before it fails. The failure occurs because there is a minimum timeout period of five minutes for almost all Data Domain Boost operations.

To resolve the failed backup, set the Data Domain system online and then retry the backup.

## Rolling Back After Deleting a Data Domain System

Rolling back to a checkpoint after following the procedure for "Deleting the Data Domain System from the vSphere Data Protection Advanced Appliance" should result in a state where the Data Domain system is removed from the vSphere Data Protection Advanced appliance.

To restore the Data Domain system to the vSphere Data Protection Advanced appliance, use the vSphere Data Protection configuration UI. Refer to "Adding a Data Domain System."

If rollback to a checkpoint was done prior to the deletion of the Data Domain system, the system should still be attached and properly configured. To remove the Data Domain system, follow the procedure for "Deleting the Data Domain System from the vSphere Data Protection Advanced Appliance."

If a rollback of the appliance results in a different outcome, contact support to define the proper resolution.

**VMware, Inc.** 3401 Hillview Avenue Palo Alto CA 94304 USA Tel 877-486-9273 Fax 650-427-5001 www.vmware.com

 Item No: VMW-TWP-vSPHR-VDP-Adv-Dta-Dmn-Int-USLET-101 Docsource: OIC - 14-FP-1152