

Optimizing VMware protection with HP and Veeam

Technical white paper

Table of contents

Executive summary.....	2
Key points.....	2
Introduction.....	2
Technology overview.....	2
VMware vStorage APIs and Change Block Tracking.....	3
Veeam Backup source-side deduplication.....	3
HP StoreOnce D2D appliance target-side deduplication.....	3
Combining Veeam Backup inline and HP D2D target-based deduplication.....	3
HP StoreOnce D2D appliances and Veeam Backup.....	4
An appliance vs. general purpose disk for data protection.....	4
Important considerations when sizing your HP StoreOnce D2D appliance.....	5
Performance considerations.....	5
Configuring a NAS share on the HP StoreOnce D2D appliance.....	5
Defining a new backup job.....	7
Processing Mode.....	7
Backup Destination.....	8
Advanced features: Backup tab.....	8
Backup mode.....	9
Advanced Settings: Storage tab.....	9
Deduplication and HP StoreOnce D2D appliance replication.....	9
Compression.....	10
Advanced Settings: vSphere tab.....	10
Change Block Tracking.....	10
Advanced Settings: Advanced tab.....	11
Integrity checks.....	11
Post-job activity.....	11
Integration with physical tape backup.....	11
Migrating backups to physical tape.....	12
Conclusion.....	14
For more information.....	15



Executive summary

VMware administrators know how important a reliable and flexible data protection strategy is. Traditional backup methods are not optimized for VMware environments. These traditional methods are adequate when making the transition to a VMware infrastructure, but as more workloads and critical workloads are moved to virtual machines, VMware optimized backup and recovery solutions are needed.

Data protection software and hardware need to provide an effective solution that protects VMware files, images, templates, and applications running on virtual machines by integrating directly into VMware's rich feature set while offering features such as backup data deduplication and replication.

Hewlett Packard and Veeam Backup software offer VMware optimized protection solutions for virtual data center administrators.

Key points

- For VMware-centric environments, Veeam Backup is an excellent application to manage backup and recovery to HP StoreOnce D2D appliances.
- Using Veeam Backup Incremental backup mode provides for faster incremental backups and allows for leveraging synthetic full backups.
- Both the Veeam Backup and the HP StoreOnce D2D appliances offer backup data deduplication. Highest backup performance is achieved by enabling both the software and hardware deduplication.
- Veeam Backup has no native backup to tape support. For longer term protection and archiving, backup data can be copied to tape using a simple scripted integration with other backup applications.
- Backup to HP StoreOnce D2D appliances can be configured and run in less than 20 minutes with the Veeam Backup Job Wizard.

Introduction

This technical brief provides an overview of the latest VMware data protection technologies available and their recommended best practices when configuring the HP StoreOnce D2D appliance with Veeam Backup.

This paper focuses on the recommended settings to be used when configuring an HP StoreOnce D2D appliance with Veeam Backup Incremental backup mode with Synthetic Full. By following these recommendations, the VMware data protection solution can achieve optimum results.

The technical information in this brief has been tested and certified by both HP and Veeam Backup.

Technology overview

Technologies covered include VMware Change Block Tracking, VMware vStorage APIs, Veeam Backup inline source-side deduplication, HP StoreOnce D2D appliance target-side deduplication and HP StoreOnce replication. When used together, this delivers a data protection solution that can be implemented in any VMware environment.

VMware vStorage APIs and Change Block Tracking

Change Block Tracking (CBT) is a feature that was introduced in VMware vSphere 4.0 that keeps track of blocks of data stored on a virtual disk that have been changed since a certain point in time. CBT is a VMkernel feature. Access is via the VMware vStorage APIs. By default, CBT is not enabled due to the slight (1–2%) overhead incurred when using this feature.

The vStorage APIs are a set of interfaces exposed for third party hardware and software integration with the VMware storage features. A subset of these is specifically for backup and recovery software vendors to optimize VMware data protection. These are known as the vStorage APIs for Data Protection (VADP). The VMware recommended method of backup and recovery is to use products that support VADP. Veeam Backup fully supports VADP.

Veeam Backup queries the VMkernel through the VADP, to identify the blocks of data that have changed on the virtual disk since the last backup, which greatly improves the performance of incremental backups.

The first full backup of any virtual machine results in all the data being transferred to the HP StoreOnce D2D appliance, but all incremental data consists of only the changed blocks as identified by CBT.

Note:

The use of CBT requires VMware virtual hardware 7 for virtual machines.

Veeam Backup source-side deduplication

Inline deduplication performs deduplication processing during the backup process before writing data to the HP StoreOnce D2D appliance. This reduces the amount of data packets sent across a network while increasing storage utilization. The scope of the deduplication is the data in the backup job. Deduplication is a configurable option for the backup job.

Inline deduplication works by cataloging and comparing the processed data to data previously backed up data. Duplicate data is skipped while unique data is sent to the HP StoreOnce D2D appliance.

Veeam Backup inline deduplication uses resources on the client machine or the Veeam Backup server. The hardware needs to be sized to provide the required resources for the deduplication process and not impact other applications running. Veeam Backup inline deduplication block sizes are configurable from 256 KB to 1 MB. These block sizes are optimized for the distance to the backup storage from the virtual machine. Setting the storage target to a WAN target results in a block size of 256 KB, 512 KB for LAN targets and 1 MB for locally attached targets.

HP StoreOnce D2D appliance target-side deduplication

Target-side deduplication is performed on the HP StoreOnce D2D appliance that has been configured as either a Virtual Tape Library (VTL) or NAS share and is presented to the backup application. A major benefit of target-side deduplication is that the backup storage appliance removes the processing I/O from the backup server or clients enabling the dedicated hardware to use finer deduplication chunking, which improves deduplication ratios and storage utilization.

Combining Veeam Backup inline and HP D2D target-based deduplication

The HP and Veeam recommendation is to enable Veeam Backup inline deduplication along with HP StoreOnce D2D appliance target-side deduplication. Both deduplication techniques provide complementary benefits.

Veeam Backup's inline deduplication reduces the amount of backup data sent over the network, but works only within single backup jobs. The HP StoreOnce D2D appliance uses smaller blocks to identify repeated data across single and multiple backups produced by different Veeam backup jobs. This further reduces the amount of data to be physically stored, thus greatly improving storage utilization.

HP StoreOnce D2D appliances and Veeam Backup

This paper focuses on best practices and techniques for using the Veeam Backup Incremental backup mode in conjunction with the HP StoreOnce D2D appliances. Veeam Backup's other backup mode—Reversed Incremental is not covered in this paper since this mode is designed for backups to general purpose disk storage. Testing has confirmed that this backup mode may impact performance when used with deduplicating backup appliances such as the HP StoreOnce D2D appliance.

When Veeam Backup Incremental backup mode is selected, Veeam Backup creates a full backup file (.vbk) at the first run of the backup job. At subsequent backups, it only saves changes that have taken place since the last performed backup (whether full or incremental) and saves them as incremental backup files (.vib).

As well as being optimized for deduplicating backup appliances, the Veeam Backup Incremental backup mode is the best choice if regulation and policies require you to regularly move a backup file to tape. With the Incremental backup mode, only incremental changes, not the full backup file, are moved, which takes less time and requires less tape. Tape copy is not supported directly by Veeam Backup, but it can be scripted as a post-process activity. More detail on this feature is provided in a later section.

An appliance vs. general purpose disk for data protection

There are many reasons why the HP StoreOnce D2D disk appliance more robust and easier to manage for data protection than general purpose disk:

- The HP StoreOnce D2D appliance is a single purpose machine, optimized for backup. It is robust, and access is limited to prevent unwanted deletions and data corruption. There is a dedicated management interface to control and report on all the data protection features. General purpose disk is cheaper to buy but much more expensive to manage and less robust.
- The HP StoreOnce D2D appliance has the added benefit of reliable performance with fine-grain chunking and Sparse Indexing (Intellectual property of HP that reduces memory and disk I/O requirements, enabling efficient deduplication for all workloads) resulting in higher deduplication ratios thus achieving better utilization of storage capacity.
- Consolidating backups onto an appliance enables deduplication across backup jobs resulting in increased backup data deduplication and improved storage utilization.
- With smaller deduplication blocks than available in the Veeam Backup, deduplication ratios are higher and replication from D2D appliance to D2D appliance is more efficient by sending smaller data packets, thus utilizing a smaller amount of bandwidth.

Note:

With Veeam Backup Incremental backup mode, you must regularly perform either Synthetic (build from incremental data) or Active full backups to maintain a specified retention policy. Multiple full backups will not consume significant storage space, because all similar blocks will only be stored once by the HP StoreOnce D2D appliance.

Important considerations when sizing your HP StoreOnce D2D appliance

When sizing a D2D backup appliance for a VMware environment it is important that the correct HP StoreOnce D2D products are selected. The key factors to consider are:

- Number of backup jobs
- Backup frequency
- Retention periods
- Source data change rate
- Estimation of the achievable deduplication ratio

These variables are set to meet the Restore Point Objectives (RPOs) and Restore Time Objectives (RTOs) as determined by Service Level Agreements (SLAs).

Correct sizing will avoid delayed backups and/or not meeting SLAs. Information and links regarding the HP Sizing tools can be found at the end of this document.

Performance considerations

Listed are three considerations regarding performance when sizing an HP StoreOnce D2D appliance for use with Veeam Backup Incremental backup mode with Synthetic Full Backups:

1. Backup storage opening—When starting, Veeam Backup opens existing backup files and reads small portions of metadata from each. This operation happens each incremental backup cycle, and involves opening previous full backup files and all subsequent incremental backup files. As a result, the further you go from a full backup, the more backup files need to be opened resulting in higher performance requirements.
2. Synthetic full backup creation (creating a new synthetic full backup using a previous full backup file and incremental backup files)—In addition to (1), this requires reading and writing a lot of other data, and requires the backup storage to provide enough IOPS for both random and sequential reads, and mostly sequential writes.
3. Booting up a stored virtual machine from a backup file—On a high level, from a storage device perspective, the workload of booting a stored virtual machine from an HP StoreOnce D2D appliance has similar requirements to (1) and (2), with added read randomness resulting in some loss of performance. Although the testing of this feature resulted in acceptable performance, deduplicating backup appliances are not designed as primary storage. When combined with Veeam Backup, booting from a backup VM provides an excellent method to restore service quickly, on a temporary basis, until primary storage is operational.

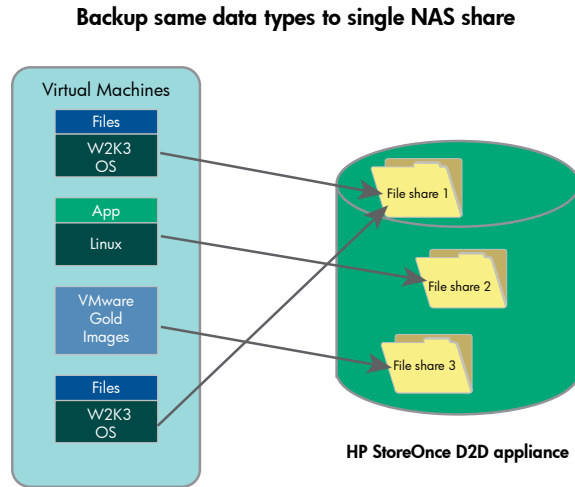
Configuring a NAS share on the HP StoreOnce D2D appliance

The HP StoreOnce D2D appliance is architected as a deduplicating backup appliance that scales in capacity and performance. The HP StoreOnce D2D appliance is also a data “agnostic” device that efficiently integrates into Windows®, Linux, and virtual environments.

Although the HP StoreOnce D2D appliance supports NAS (NFS and CIFS) and VTL emulation, best practices for VMware protection with Veeam Backup is to utilize the CIFS protocol. Veeam Backup does not natively support tape and therefore not the D2D appliance’s VTL emulation.

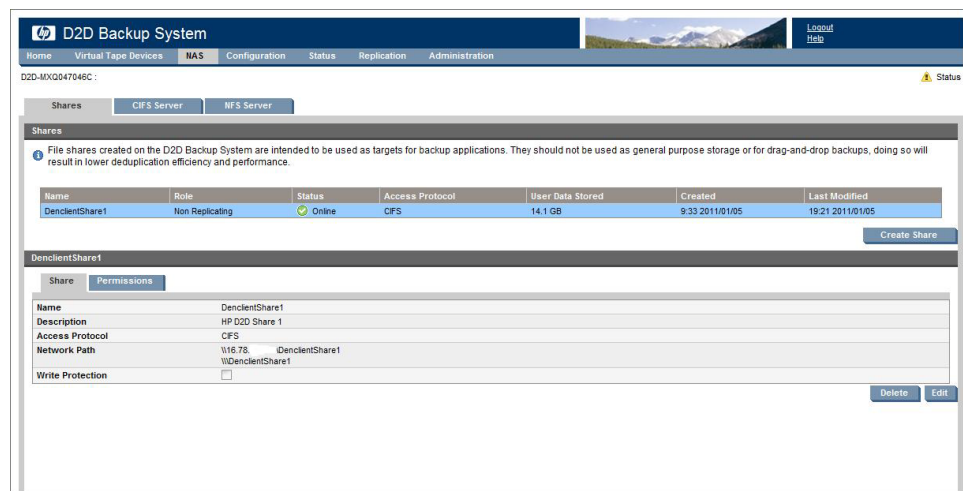
When setting up a NAS share to store backup data, it is advisable to configure multiple shares, each as the target for similar data types. For example, file system data should be stored in one NAS share, while database data should be stored in a separate share. This is recommended in order to improve the likelihood that data will deduplicate against similar data in order to maximize deduplication ratios while minimizing space usage. The scope of the D2D deduplication is across a single share.

Figure 1: Backing up the same data types to individual NAS shares



For example, as noted in figure 1, a virtual machine running a Windows Server 2003 operating system storing file data will likely contain the same data (at a block level) as another virtual machine running Windows Server 2003 storing file data. The same will hold true for databases with like fields (such as SQL Server) or when backing up virtual machine golden images or templates.

Screenshot 1: Creating an HP StoreOnce D2D appliance NAS CIFS share using the appliance graphical interface



Defining an HP StoreOnce D2D appliance NAS share is straightforward and can be completed in a few steps. A few things to consider are the share name, permissions, and whether it will connect to an Active Directory server or not. There isn't a need to define the share type such as file system or database share, as with hash based chunking, the deduplication algorithm will identify data blocks and store accordingly.

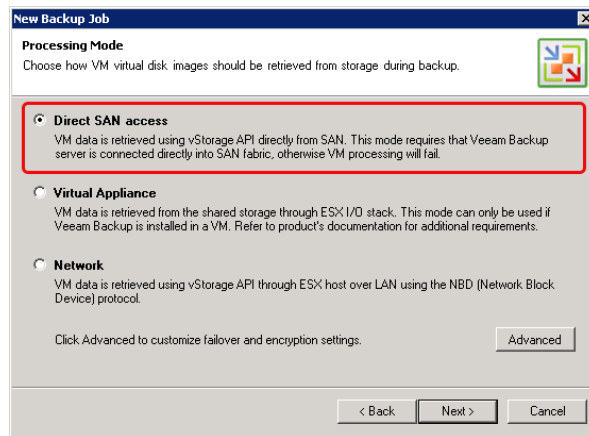
An additional benefit to the HP StoreOnce D2D appliance is that the defined NAS share will automatically adjust its capacity when storing data. It is not necessary to define a NAS share capacity limit; however, you can set a capacity water mark in Veeam Backup to warn you when you reach a capacity threshold.

Once the NAS share has been created it can be accessed by mounting it directly to the Veeam Backup server using the HP StoreOnce D2D appliance network address path provided. It is then configured and accessed directly as a backup storage location.

When configuring an Incremental Backup with Synthetic full with the HP StoreOnce D2D appliance, HP and Veeam recommend the default backup settings be used when configuring your virtual machine backup(s). Basic steps of setting up and managing your backup job can be found in the [Veeam Backup User Guide](#). However, to ensure the HP StoreOnce D2D appliance is able to take advantage of the extra features provided by HP, Veeam, and VMware, it is recommended to use the following best practice methods. Implementing the following recommendations will result in quicker backups, improved network utilization, deduplication ratios and efficient D2D replication (if utilized).

Defining a new backup job

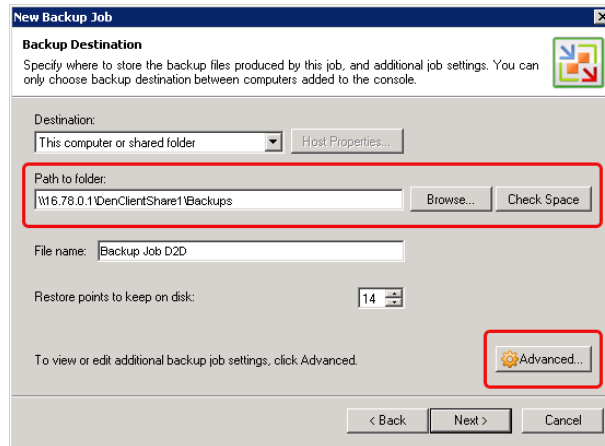
Screenshot 2: Selecting Processing Mode in Veeam Backup



Processing Mode

When determining the Processing Mode, it is recommended to use the Default "Direct SAN access" mode when running the Veeam Backup server on a physical server and "Virtual Appliance" when running the Veeam Backup server on a virtual machine, and allow the network to be used as a failover if necessary.

Screenshot 3: Backup Destination



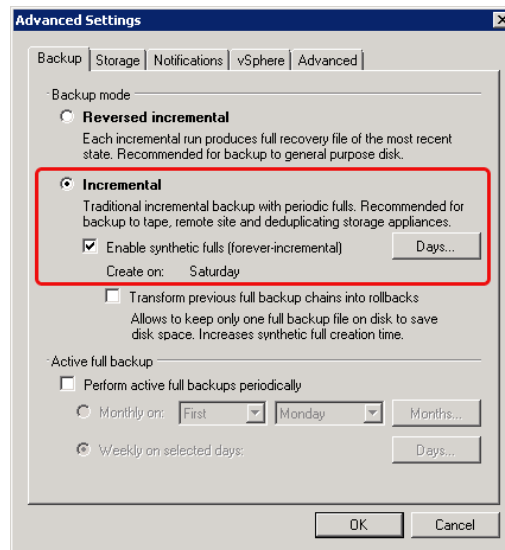
Backup Destination

When defining the Backup Destination, the UNC path to the HP StoreOnce D2D appliance's configured CIFS share must be specified in the job settings. It is at this step where adding advanced backup settings also need to be set.

Advanced features such as backup mode, inline deduplication, and CBT need to be enabled to maximize backup performance, network efficiency, and D2D capacity.

Advanced features: Backup tab

Screenshot 4: Advanced Settings



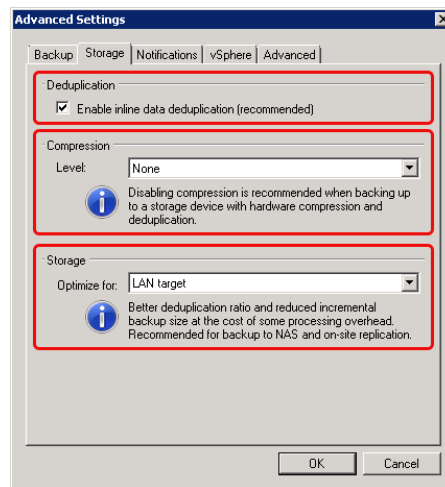
Backup mode

HP and Veeam recommend specifying Incremental backup mode with Synthetic full backups. When choosing the Incremental mode, you must select to periodically create a full synthetic backup or perform active full backups regularly. Not performing this important step will result in degraded restore performance as data will have to be built from generations of incremental backups instead of a single consolidated full and perhaps a handful of incremental backups.

Selecting “Active full backup” as the way to create a weekly full backup results in losing the benefit of forever incremental backups.

Advanced Settings: Storage tab

Screenshot 5: Storage Settings



Deduplication and HP StoreOnce D2D appliance replication

It is recommended to enable Veeam Backup’s inline data deduplication feature and set Storage Optimization to “LAN Target” when writing backups to the HP StoreOnce D2D appliance. Setting the correct optimization level is important as this determines the most efficient storage block size for inline deduplication. The slower the link between the backup server and the backup storage, the smaller the dedupe block size should be. This will result in using smaller blocks for incremental backup packets, resulting in less data being transmitted across the network to the storage device.

Setting the optimization level to “LAN Target” will use a constant 512 KB block size for incremental backups. When the data is received by the HP StoreOnce D2D appliance, the StoreOnce algorithm will further deduplicate the data using smaller variable boundary 4KB blocks for optimum storage utilization and more efficient data replication.

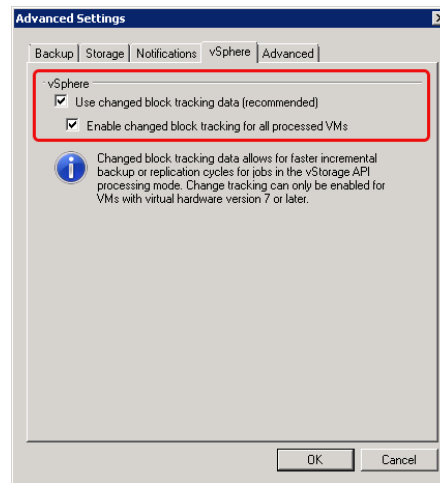
Additional consolidation benefits can be achieved by implementing a many to one fan in which multiple virtualized environments are backed up to a local HP StoreOnce D2D appliance that can replicate similar data to a single centralized HP StoreOnce D2D appliance.

Compression

Disabling Veeam Backup's compression (or setting the compression level to "None") is recommended when backing up to the HP StoreOnce D2D appliance. Compression is an integrated part of the HP StoreOnce D2D appliance duplication process where deduplicated data is compressed before being written to disk. This is another benefit of the HP StoreOnce D2D appliance as it moves compression processing from server to target. Sending compressed data to the HP StoreOnce D2D appliance reduces the appliance's deduplication ratio.

Advanced Settings: vSphere tab

Screenshot 6: vSphere

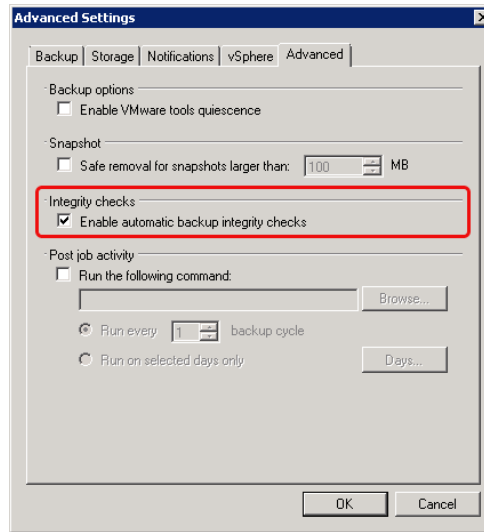


Change Block Tracking

Hewlett Packard and Veeam recommend enabling VMware CBT for all processed virtual machines. Once selected, Veeam Backup will automatically enable the feature for all processed virtual machines during the first backup run, which is when full backups are created.

Advanced Settings: Advanced tab

Screenshot 7: Advanced Settings



Integrity checks

It is recommended to select “Enable automatic backup integrity checks” if you want Veeam Backup to periodically check a full backup file. An automatic backup check allows you to verify integrity of a backup file and avoid a situation when a full backup is corrupted, making all further incremental backups corrupted, too. A backup check is performed every time a full backup file is rebuilt to include incremental changes. If the check determines a full backup file to be corrupted, a notification message will be displayed, prompting you to perform a new full backup. During such full backup, no integrity check will be performed.

Post-job activity

Select the “Run the following command” check box if you want to execute post-backup actions, such as launching a script to copy the backup file to tape. This is covered in greater detail in the next section.

Integration with physical tape backup

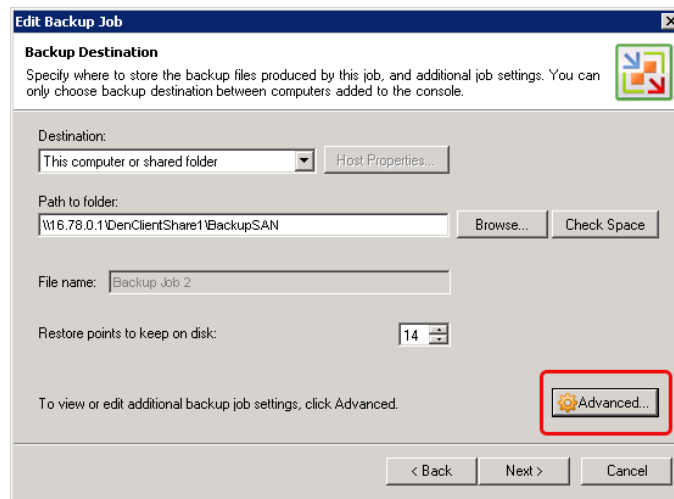
Veeam Backup can be integrated with an existing backup scheme, offering a flexible approach to protect VMware environments. This is useful for organizations using tape as a backup target. With the option of performing post-backup activities, actions can be executed once the backup procedure is completed. A custom script can initiate the write of a Veeam backup file to tape. The desired post-backup activity may be performed once after a set of backup jobs runs, which will protect against redundant loading of the virtualized environment.

Migrating backups to physical tape

The best protection strategies use disk and tape. Veeam Backup does not support tape directly; therefore, another backup application is required to move data from an HP StoreOnce D2D appliance to tape. This is done by backing up the D2D NAS share that contains the backup data required to be moved to tape using a backup and recovery application that supports tape. This can be automated by running a script or post-backup job as an option in the backup job. The script depends upon the backup and recovery software used. An example of specifying this is shown in Figures 8 and 9. Note that this backup method requires a two-stage recovery from tape. Stage 1 is to restore the backup from tape to the D2D device using the backup and recovery software that wrote the backup to tape. Stage 2 is a recovery from the D2D device using Veeam Backup.

Veeam Backup provides a simple way to script post-job activities by including it as an option in the backup job.

Screenshot 8: Advanced settings in Veeam Backup



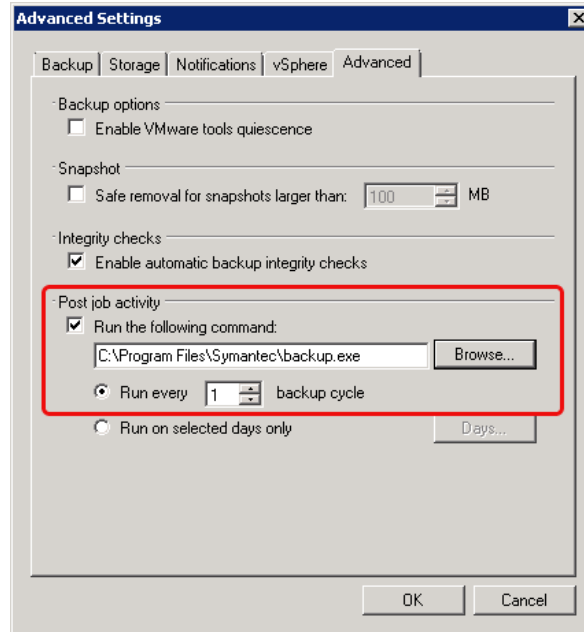
To add this option to a predefined backup job you will need to edit the job policy. Select the backup job and right click "Properties". Click through the configuration to the "Backup Destination" step and click the "Advanced" button to access the Advanced Settings. Click the "Advanced" tab and check the "Post job activity" box.

Once a Veeam backup job is completed, if a post job is scheduled, the requested command or script will be performed.

Depending on the policy and regulation guidelines, tape copy jobs may need to occur at different times. Copying daily incremental backups will use less tape but will likely result in longer restore times. Copying weekly full backups will use more tape, but will have quicker restore times.

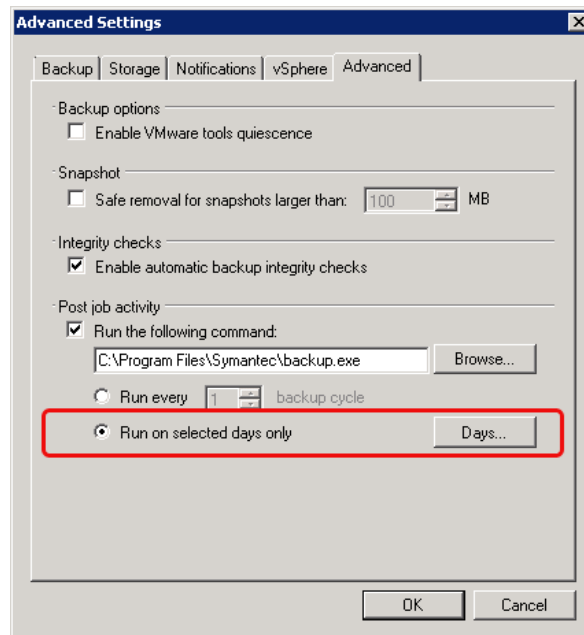
To perform a daily tape copy job (e.g., after each incremental) set the number of backup cycle days to 1.

Screenshot 9: Daily post-job activity



To set a weekly tape copy job, for example, after each Synthetic full, select the same day as the full backup.

Screenshot 10: Weekly post-job activity



Below is a list of commands that can be used to provide the tape offload functionality in conjunction with leading backup and recovery applications that support tape:

HP Data Protector

```
C:\Program Files\OmniBack\bin\Omnib -datalist <backupjobName> -mode <full|incremental>
```

Symantec NetBackup

```
C:\Program Files\Veritas\NetBackup\bin>bpbackup.exe" -p <backupJobName> -i
```

NetBackup Options:

- p Policy Name
- i immediate manual backup

Symantec Backup Exec

```
C:\Program Files\Symantec\Backup Exec\bemcmd.exe" -o01 -j<BackupjobName>
```

Backup Exec Options:

- o1 Runs backup job
- j Job Name

Conclusion

Server virtualization with VMware enables new efficiencies to be achieved. These efficiencies can be extended to backup and recovery using HP and Veeam solutions. The HP StoreOnce D2D appliance is a high-performance and reliable backup target that addresses data growth challenges through HP StoreOnce data deduplication technology. Together, the HP StoreOnce D2D appliance and Veeam Backup & Replication provide a comprehensive and complete data protection solution for VMware environments.

For more information

HP StoreOnce D2D Appliance

<http://www.hp.com/go/d2d> and <http://www.hp.com/go/storeonce>

HP StoreOnce D2D Disk Backup Application HP D2D 4312 user guide

<http://bizsupport1.austin.hp.com/bc/docs/support/SupportManual/c02295179/c02295179.pdf>

HP StoreOnce D2D Disk Backup Appliance HP D2D 4112 user guide

<http://bizsupport1.austin.hp.com/bc/docs/support/SupportManual/c02295156/c02295156.pdf>

HP StorageWorks Sizing Tool

<http://www.hp.com/go/nssizer>


HP StoreOnce D2D Disk Backup Appliance best practices guide

<http://bizsupport2.austin.hp.com/bc/docs/support/SupportManual/c02511912/c02511912.pdf>

Veeam Backup & Replication <http://www.veeam.com/backup>

Veeam Backup & Replication user guide & evaluator's guide (require registration)

<http://www.veeam.com/vmware-esx-backup/resources.html>



© Copyright 2011 Hewlett-Packard Development Company, L.P. The information contained herein is subject to change without notice. The only warranties for HP products and services are set forth in the express warranty statements accompanying such products and services. Nothing herein should be construed as constituting an additional warranty. HP shall not be liable for technical or editorial errors or omissions contained herein.

UNIX is a registered trademark of The Open Group. Windows is a U.S. registered trademark of Microsoft Corporation.

4AA3-3435ENW, Created April 2011

