



VMware vSphere® Data Protection

TECHNICAL WHITE PAPER
JUNE 2012

Table of Contents

Introduction 3

Architectural Overview 3

Deployment and Configuration 4

Management 5

 Virtual Machine Backup.....5

 Virtual Machine Restore.....6

Reporting7

Preventing Backup-Data Corruption 9

Conclusion 10

About the Author 10

Introduction

With VMware vSphere® 5.1 (“vSphere 5.1”), VMware® is releasing a new backup and recovery solution for virtual machines called vSphere Data Protection (VDP). This solution is fully integrated with VMware vCenter Server™ (vCenter Server) and provides agentless, disk-based backup of virtual machines to deduplicated storage.

Benefits of VDP include the following:

- It ensures fast, efficient protection for virtual machines even if they are powered off.
- It uses patented deduplication technology across all backup jobs, significantly reducing disk space consumption.
- VMware vSphere® APIs – Data Protection (VADP) and Changed Block Tracking (CBT) are utilized to reduce load on the vSphere hosts and minimize backup window requirements.
- It performs full virtual machine and File-Level Restore (FLR) without installing an agent in every virtual machine.
- Installation and configuration is simplified using an appliance form factor.
- Management is performed utilizing the VMware vSphere® Web Client (vSphere Web Client).
- The VDP appliance and its backups are protected using a checkpoint and rollback mechanism.
- Windows and Linux files can easily be restored by the end user with a Web browser.

This paper presents an overview of the architecture, deployment, configuration, and management of VDP.

Architectural Overview

VDP requires VMware vCenter Server 5.1 or higher. vCenter Server can be the traditional Windows implementation—or the new Linux-based VMware vCenter Server Virtual Appliance (VCVA), first released in VMware vSphere® 5.0. VDP supports the backing up of virtual machines on VMware vSphere® (“vSphere”) versions 4.0 and higher. The following Web browsers are supported for configuration and management of VDP:

- Microsoft Internet Explorer versions 7 and 8
- Mozilla Firefox versions 3.6 and higher

VDP is deployed as a preconfigured Linux-based appliance. Each appliance supports as many as 100 virtual machines, and as many as 10 VDP appliances can be deployed per vCenter Server instance. The Windows-based VMware vSphere® Client™ is used to deploy VDP. After the appliance has been deployed, management is performed using the vSphere Web Client with any supported Web browser. Adobe Flash must be installed in the Web browser.

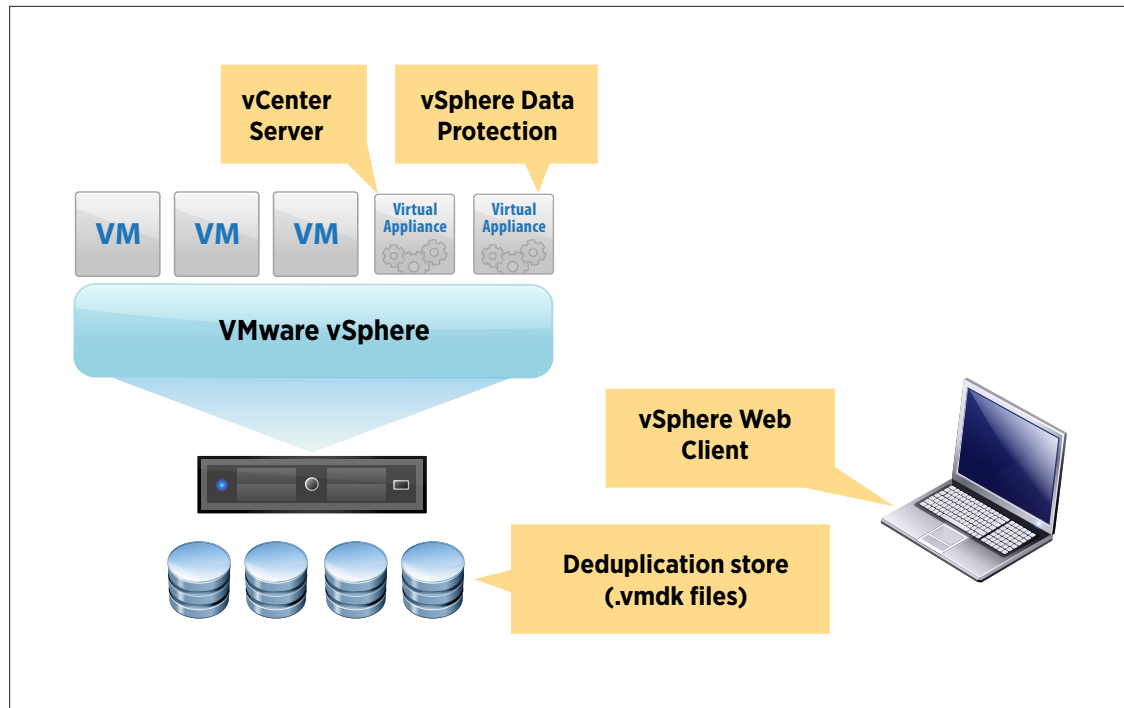


Figure 1. VDP Components

The VDP appliance is deployed with four processors (vCPUs) and 4GB of RAM. Three configurations of usable backup storage capacity are available: .5TB, 1TB and 2TB, which respectively consume 850GB, 1,300GB and 3,100GB of actual storage capacity. Proper planning should be performed to help ensure that proper sizing and additional storage capacity cannot be added after the appliance is deployed. Storage capacity requirements are based on the number of virtual machines being backed up, amount of data, retention periods and typical data change rates. Recommendations for VDP sizing can be found in the *VMware vSphere Data Protection Administration Guide*.

Deployment and Configuration

VDP is deployed using the vSphere Client from a prepackaged Open Virtualization Archive (.ova) file. The .ova files are labeled to easily identify the amount of backup storage capacity included with the appliance.

After the appliance is deployed and powered on, a Web browser is used to access the VDP-configure user interface (UI) and perform the initial configuration. The first time the user connects to the VDP-configure UI, it will be running in installation mode. With the installation mode wizard, items such as IP address, host name, DNS, time zone and vCenter Server connection information are configured. Upon successful completion of the installation mode wizard, the appliance must be rebooted. This reboot can take up to 30 minutes to complete as the appliance finishes initial configuration.

After the initial configuration, the VDP-configure utility runs in maintenance mode. In this mode, the VDP-configure UI is utilized to perform functions such as starting and stopping services on the appliance, collecting logs and rolling back the appliance to a previous valid configuration state (discussed later in this document).

The vSphere Web Client is used to create and maintain backup jobs and perform entire virtual machine restores, as well as for reporting and configuration of VDP. Figure 2 shows VDP in the vSphere Web Client.



Figure 2. vSphere Data Protection in the vSphere Web Client

The initial backup of a virtual machine takes comparatively more time, because all of the data for that virtual machine is being backed up. Subsequent backups of the same virtual machine take less time, because VDP utilizes CBT and deduplication.

Management

Virtual Machine Backup

Creating and editing a backup job is accomplished using the **Backup** tab of the VDP UI in the vSphere Web Client. Individual virtual machines can be selected for backup. Containers such as datacenters, clusters, hosts, resource pools and folders also can be selected for backup. All virtual machines in the container at the time the backup job runs will be backed up. New virtual machines added to the container will be included when the next backup job runs. Similarly, any virtual machines removed from the container no longer will be backed up.

Backup jobs can be scheduled daily, weekly or monthly. Each job runs once on the day it is scheduled and begins when the backup window opens (default is 8:00 p.m. local time). As many as eight backup jobs can run simultaneously on each VDP appliance.

The retention policy can be defined in the following several ways:

- Forever
- For x number of days, weeks, months or years
- Until a specific date
- Custom schedule: for example, daily for 15 days, weekly for 4 weeks, monthly for 6 months, and yearly for 5 years

The retention policy determines how long backups are retained. After this time period expires, they are deleted from the system.

Keep: ☐ Forever

☒ for 60 day(s)

☐ until 07/28/2012

☐ this Schedule:

Daily for: 60 day(s)

Weekly for: 0 week(s)

Monthly for: 0 month(s)

Yearly for: 0 year(s)

Figure 3. VDP Retention Policy Configuration

Virtual Machine Restore

The restore of an entire virtual machine is performed using the **Restore tab** of the VDP UI in the vSphere Web Client. The administrator can browse the list of virtual machines backed up by VDP and then select one or more restore points. By leveraging CBT during a restore of a virtual machine to its original location, VDP offers fast and efficient recovery. During the restore process, VDP queries VADP to determine which blocks have changed since the selected restore point, and it recovers only those blocks. This reduces data transfer within the vSphere environment during a recovery operation and decreases recovery time. VDP automatically compares and evaluates the workload of the two restore methods (full-image restore and restore leveraging CBT) and utilizes the method resulting in the fastest restore time. This is useful in scenarios where the change rate since the selected restore point is very high and the overhead of a CBT analysis operation would be more costly than that of a full-image recovery. VDP intelligently determines which deployment method will result in the fastest recovery time. A new virtual machine name and destination datastore also can be specified to prevent overwriting an existing virtual machine. Choosing a restore location other than the original will result in a full-image restore (CBT is not leveraged).

Set the restore options for each backup that you are restoring.

Client: vm1

Backup: 05/24/2012 08:07 PM

☐ Restore to Original Location

New Name: vm1_5_24_2012

Destination: /Datacenters/datacenter1/host3.vmware.local Choose

Datastore: vms (461.8 GiB free) ▼

Figure 4. Specifying a New Location for the Restored Virtual Machine

It also is possible to restore individual files and folders/directories within a virtual machine. An FLR is performed using a Web-based tool called vSphere Data Protection Restore Client. The process enables end users to perform restores on their own without the assistance of a VDP administrator. The end user can select a restore point and then browse the file system as it looked at the time that the backup was performed. After the end user locates the item(s) to be restored, a destination (on the local machine) is selected and the job is started. The progress of the restore job can also be monitored in the tool.

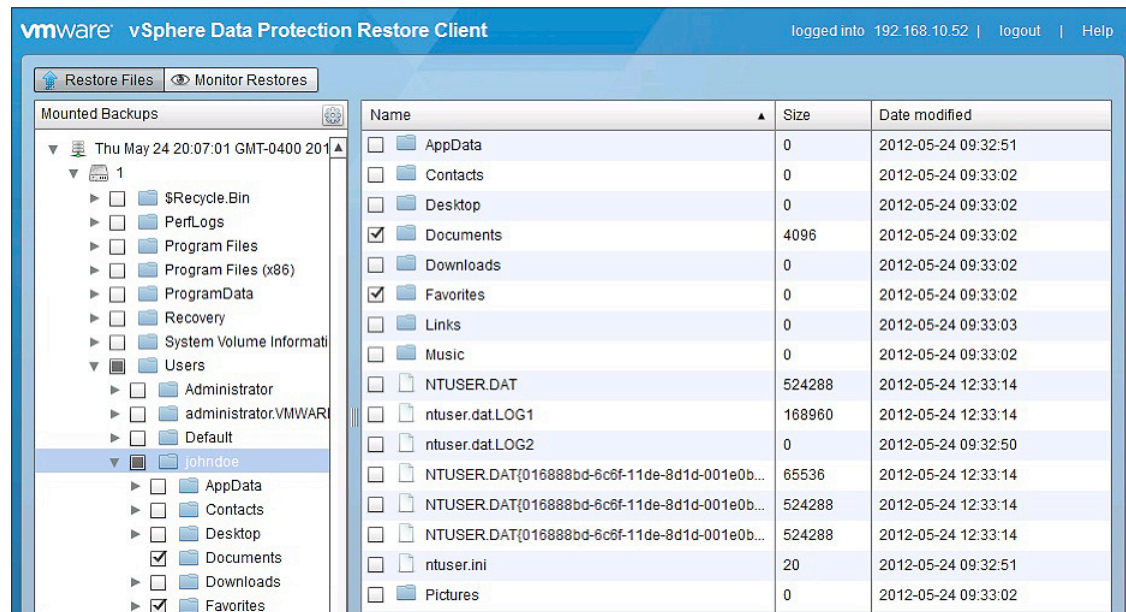





Figure 5. vSphere Data Protection Restore Client



Reporting


The Reports tab displays the following information: VDP Appliance Status, Used Capacity, backup job information, virtual machine backup details, and so on. There are links to the Event Console and Task Console for additional information and troubleshooting purposes. Users can filter the list of virtual machines by means of several criteria, including Virtual Machine Name, Backup Jobs and Last Successful Backup date. The virtual machine information details section displays the Virtual Machine Name, guest operating system, backup status, backup date and other useful items.

vSphere Data Protection

vdp1 (192.168.10.50) Switch Appliance: vdp1   All Actions 

Getting Started Backup Restore **Reports** Configuration

 Refresh 

 Appliance Status: Normal
 Used Capacity: 1.54%
 Integrity Check Status: Normal
 Recent Successful Backups: 1
 Recent Failed Backups: 0

[Event Console](#)
[Task Console](#)

Virtual Machines [Reset Filter](#)

Filter: Show All ▼

Virtual Machine Information				Last Backup Job		
Virtual Machine Name	State	Backup Jobs 1 ▲	Last Successful Backup	Status	Date	Backup
vm2	poweredOff	External Web	05/29/2012 08:57 AM	SUCCESS	05/29/2012 08:57 AM	External Web
vc1	poweredOn		Never		Never	
vdp1	poweredOn		Never		Never	
dc1	poweredOn		Never		Never	
vm1	poweredOn		05/23/2012 08:07 PM		Never	
vdr1	poweredOn		Never		Never	

Details

VM Information	Last Backup Job
Name: vm2	Status: SUCCESS
Guest OS: SUSE Linux Enterprise 11 (32-bit)	Date: 05/29/2012 08:57 AM
Host: host2.vmware.local	Backup Job: External Web
Ip Address: 192.168.10.131, 00:50:56:ad:27:af	

Figure 6. Reports Tab

In addition to the reporting capabilities of its UI, VDP can be configured to send email reports, which can be scheduled at a specific time once per day on any or every day of the week. Similar to the UI, these email messages contain details on the VDP appliance, backup jobs and the virtual machines that are backed up.

vdp1 - (192.168.10.50)	
Report Date:	May 29, 2012 - 12:00
Last Report Date:	May 25, 2012 - 07:00
Appliance Status:	Normal
Byte Capacity:	498.945 GiB
Bytes Free:	491.261 GiB
Used Capacity:	1.54%
Bytes Protected:	60.043 GiB
Bytes Deduped:	7.684 GiB
Integrity Check Status:	Normal
Successful Backups (Last 72 hours):	1
Failed Backups (Last 72 hours):	0
Backup Jobs Summary	
Backup Job: External Web	
Backup Sources:	External Web, vm2
Last Start Time:	May 29, 2012 - 08:48
Next Run Time:	May 29, 2012 - 20:00
Last Successful Backups:	1
Last Failed Backups:	0
Virtual Machines Summary	
Virtual Machine: dc1	
State:	poweredOn

Figure 7. VDP Email Report

Preventing Backup-Data Corruption

VDP features a checkpoint-and-rollback mechanism. A checkpoint is a system-wide backup of the VDP appliance that is performed to handle events that might cause data corruption, for example, an unexpected power-off of the VDP appliance. In this case, the VDP appliance would roll back to the last validated checkpoint. Any backup jobs performed after that checkpoint would be lost, but data corruption—that is, loss of all backup information—would be prevented.

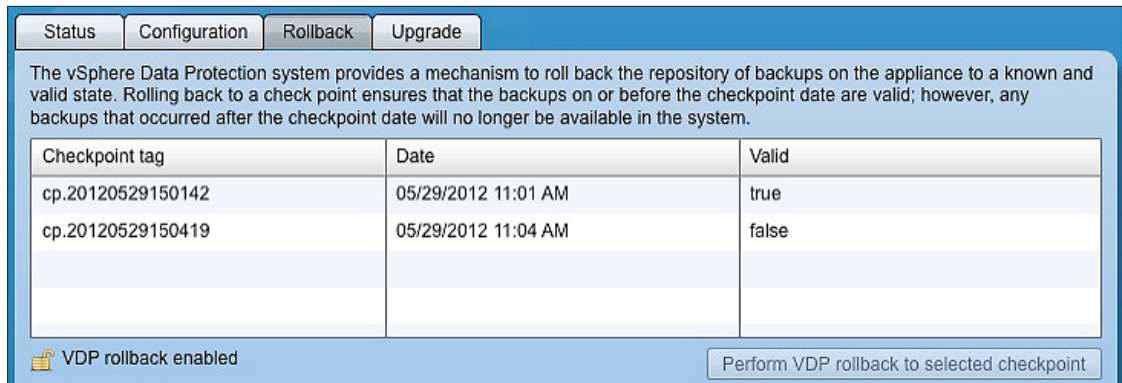


Figure 8. VDP Rollback

Conclusion

Data protection is a key component of any business continuity plan. VMware vSphere Data Protection (VDP) offers an efficient solution for protecting a VMware vSphere virtual machine infrastructure. It can be deployed quickly and provides a Web-based graphical user interface for management. End users can restore files without the need for assistance from a backup administrator. VDP also features a checkpoint-and-rollback protection system to help ensure that backup data is available for restoration when data loss occurs or disaster strikes.

About the Author

Jeff Hunter is Senior Technical Marketing Manager with a focus on vSphere data protection and other business continuity solutions at VMware. He has been employed at VMware for more than 5 years. Prior to Jeff's tenure at VMware, he spent several years assisting with the build-out and administration of VMware virtual infrastructures at two Fortune 500 companies.



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

Copyright © 2012 VMware, Inc. All rights reserved. This product is protected by U.S. and international copyright and intellectual property laws. VMware products are covered by one or more patents listed at <http://www.vmware.com/go/patents>. VMware is a registered trademark or trademark of VMware, Inc. in the United States and/or other jurisdictions. All other marks and names mentioned herein may be trademarks of their respective companies. Item No: VMW-WP-WHATS-NEW-vSPHR-DATA-PRO-USLET-101

Docsource: OIC-12VM007.08