VMware vSphere® Data Protection™ Evaluation Guide

REVISED APRIL 2015

vmware®

Table of Contents

Introduction

VMware vSphere® Data Protection™ is a backup and recovery solution from VMware. It is fully integrated with VMware vCenter Server™ and VMware vSphere Web Client, providing disk-based backup of virtual machines and applications. vSphere Data Protection is based on the industry-leading EMC® Avamar® data protection solution.

Features and Benefits of vSphere Data Protection

- Wizard-driven setup and management to quickly and easily implement a data protection solution for a VMware vSphere virtual machine environment
- Significantly reduced backup data disk space requirements, with the patented, variable-length Avamar deduplication technology
- Use of VMware vSphere Storage APIs – Data Protection as well as Changed Block Tracking (CBT) to reduce load on the vSphere host infrastructure and minimize backup window requirements
- Agentless virtual machine backup and restore that reduces complexity and deployment time
- Integration with EMC Data Domain® for additional scale, efficiency, and reliability
- Microsoft Exchange Server agent for application-consistent backup and restore of databases and mailboxes, including those protected by a database availability group (DAG)
- Microsoft SQL Server agent that leverages the Virtual Device Interface feature for proper backup and restore of databases in standalone configurations and clustered environments
- Microsoft SharePoint Server agent that enables application-consistent backup and restore
- Reliable, efficient replication of backup data between vSphere Data Protection appliances for redundancy and offsite data protection
- Flexibility to restore replicated backup data at both the source and target locations
- Automated backup verification that provides the highest level of confidence in backup data integrity
- Secure, efficient backup data replication to Avamar for offsite data protection
- Direct-to-host emergency restore operation that enables virtual machine recovery even when vCenter Server and vSphere Web Client are offline
- File Level Restore (FLR), which enables granular file and folder restoration without the need for an agent in Microsoft Windows and Linux virtual machines
- Simple Web browser–based administration through vSphere Web Client
- Appliance and backup data protection via a checkpoint-and-rollback mechanism
- Deployment of external proxies enabling as many as 24 parallel backup operations

This document focuses on the evaluation of vSphere Data Protection.

Requirements

In this document, the assumption is made that the following items are already properly installed and configured:

- Domain Name System (DNS) server with forward and reverse lookup enabled
- Static IP addresses available for vCenter Server, each vSphere host, and each vSphere Data Protection virtual appliance
- Host (A) records in DNS for vCenter Server, each vSphere host, and each vSphere Data Protection virtual appliance
- vCenter Server 5.5 or higher, including VMware vCenter™ Single Sign-On™, vCenter Inventory Service, and vSphere Web Client—vCenter Server 6.0 recommended
- At least one host running vSphere 5.1 or higher—vSphere 6.0 recommended
- Three to five virtual machines running a supported guest operating system (OS) with VMware Tools™ installed
- Storage capacity for vSphere Data Protection virtual appliances

  NOTE: vSphere Data Protection can be deployed using thin-provisioned storage for evaluation purposes. vSphere Data Protection virtual appliances are deployed to VMware Virtual SAN™, VMware vSphere VMFS, and NFS datastores. See "vSphere Data Protection Sizing" in the vSphere Data Protection Administration Guide *for details.*

- Client machine—for example, laptop—with a supported Web browser and the latest version of Flash installed for running vSphere Web Client and vSphere Data Protection Restore Client

Evaluating the application backup and restore capability of vSphere Data Protection requires one or more virtual machines running Exchange Server, SQL Server, or SharePoint Server.

To evaluate the email-reporting feature of vSphere Data Protection, the virtual appliance must have an email account on and access to a functional email system.

Backup data replication requires at least two vSphere Data Protection virtual appliances.

This document does not contain specific, step-by-step instructions for deploying and administering vCenter Server, vSphere, and vSphere Data Protection. See the following documentation for these details, as needed.

vCenter Server and vSphere documentation

vSphere Data Protection administration documentation

Recommendation: *Use default settings for all components—for example, installation paths, TCP port settings, and so on—wherever possible, to minimize installation and configuration complexity in the evaluation environment. Write down all usernames and passwords used during evaluation environment deployment.*

Evaluation Workflow

Overview

The following exercises are discussed at a high level in this document:

1. Deploying and configuring a vSphere Data Protection virtual appliance
2. Creating an image backup job
3. Running a backup job manually
4. Restoring a virtual machine
5. Performing an FLR
6. Creating an application backup job
7. Restoring an application
8. Configuring email reporting
9. Replicating backup data
10. Creating a backup verification job

The following checklist can be used to track the progress of the evaluation. The sections after the checklist provide more details on each exercise, including recommendations, documentation references, VMware Knowledge Base articles, and other resources. This document does not contain detailed, step-by-step instructions for completing the tasks in each exercise. They are documented in items such as the *vSphere Data Protection Administration Guide*, the VMware Knowledge Base, and so on. Links and references to these items are provided in each exercise. In some cases, one exercise is dependent on another one. For example, an application restore cannot be performed until an application backup has successfully been run—even if an image backup has been completed. Perform the exercises in the order documented in this guide.

Evaluation Checklist

SUCCESS CRITERIA	STATUS
Deploy vSphere Data Protection virtual appliance	
Create image backup job	
Run image backup job	
Restore a virtual machine from image backup	
Restore individual file from image backup	
Install vSphere Data Protection application agent	
Create application backup job	
Restore application	
Configure vSphere Data Protection email reporting, report received	
Replicate backup data from one vSphere Data Protection virtual appliance to another	
Restore a virtual machine from replicated backup data	
Successfully validate via backup verification job that backup data can be restored	

Exercise 1: Install and Configure vSphere Data Protection

vSphere Data Protection Virtual Appliance Deployment

Prior to deploying a vSphere Data Protection virtual appliance, verify that network connectivity and DNS are properly configured for the environment. All vSphere hosts and the vCenter Server system should have static IP addresses and be able to ping each other using IP address, host name, and fully qualified domain name (FQDN). There must also be a static IP address and DNS host (A) record for each vSphere Data Protection appliance to be deployed. For details on DNS, see this article. Also see "DNS Configuration" in the *vSphere Data Protection Administration Guide*.

Recommendations: Create a DNS host (A) record prior to deploying a vSphere Data Protection virtual appliance. Failure to do so will likely lead to vSphere Data Protection deployment issues. DNS forward and reverse lookup must be enabled.

Verify that Network Time Protocol (NTP) is properly configured throughout the environment—time must be synchronized across all machines. See "NTP Configuration" in the vSphere Data Protection Administration Guide.

Verify that the latest version of VMware Tools is installed in all virtual machines. VMware Tools is required for FLR. Virtual machines typically get the current time from the host on which they are running via VMware Tools. The vCenter Server virtual appliance and vSphere Data Protection virtual appliance are deployed with VMware Tools preinstalled. Do not update or reinstall VMware Tools in these virtual appliances.

vSphere Data Protection requires a user account with administrator permissions in vCenter Server. See Figure 1. This user should be explicitly added as an administrator to the vCenter Server root node. Users who inherit permissions from group roles are not valid. See "User Account Configuration" in the *vSphere Data Protection Administration Guide*.

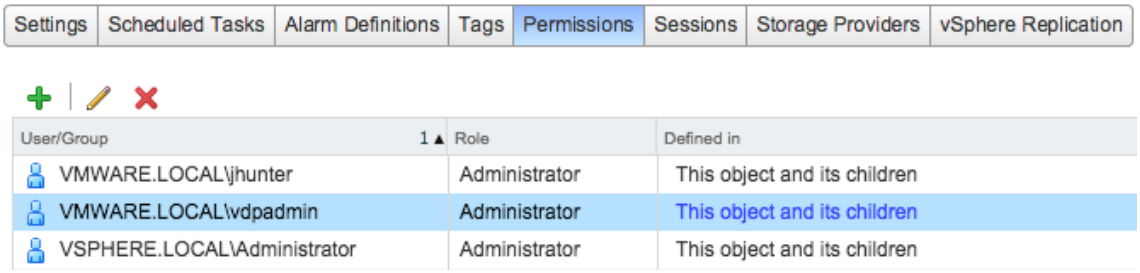

Figure 1. User Account with vCenter Server Administrator Permissions

Recommendations: Prior to vSphere Data Protection deployment, review and follow "VDP Appliance Best Practices" in the vSphere Data Protection Administration Guide.

Deploy a vSphere Data Protection virtual appliance with enough backup data capacity for the environment, taking into consideration items such as number of protected virtual machines, size of protected virtual machines (storage consumed), data change rates, and retention policy requirements. For most evaluation environments, the smallest vSphere Data Protection appliance is sufficient (.5TB). To minimize storage consumption, thin provisioning can be used for the protected virtual machines and the vSphere Data Protection virtual appliance in the evaluation environment. However, thick provisioning is recommended for production environments.

A vSphere Data Protection appliance is deployed from an OVA file, which is a "zipped" OVF template. vSphere Web Client is used to deploy vSphere Data Protection. See "Deploy the OVF Template" in the *vSphere Data Protection Administration Guide*. Deployment of the vSphere Data Protection virtual appliance might take several minutes, depending on the performance characteristics of the evaluation environment.

Configuring vSphere Data Protection

After the vSphere Data Protection virtual appliance has been deployed and powered on, the vSphere Data Protection configuration user interface (UI) is used to perform initial configuration. See the "VDP Installation and Configuration" section in the *vSphere Data Protection Administration Guide* for detailed steps on completing the configuration of vSphere Data Protection.

Recommendations: If issues are encountered during vSphere Data Protection virtual appliance deployment—for example, forgetting to create a DNS host record for the vSphere Data Protection appliance prior to deployment—the easiest resolution might be to deploy a new vSphere Data Protection virtual appliance rather than troubleshooting the existing installation.

vSphere Data Protection features a storage performance analysis to verify that the storage that vSphere Data Protection is deployed on meets or exceeds the minimum performance requirements for backing up virtual machines, running integrity checks, and other disk I/O–intensive tasks. To verify that these requirements are met, this analysis should be run during deployment.

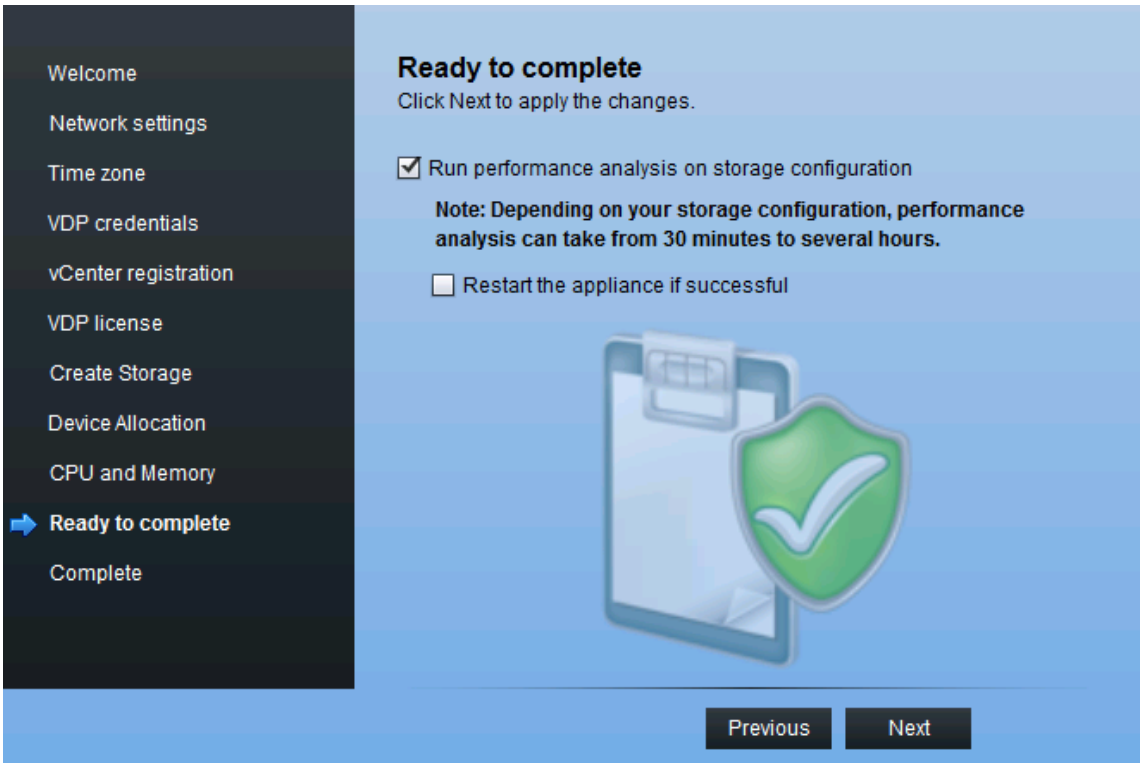

Figure 2. Run Performance Analysis on Storage Configuration

After successful completion of the initial vSphere Data Protection configuration wizard, the vSphere Data Protection virtual appliance must be rebooted. During the reboot, vSphere Data Protection completes several automated configuration steps, which can take as long as 30 minutes or more. Registering the vSphere Data Protection plug-in with vCenter Server is one of these steps. vSphere Data Protection will appear in vSphere Web Client after this step has been completed, as is shown in Figure 3.

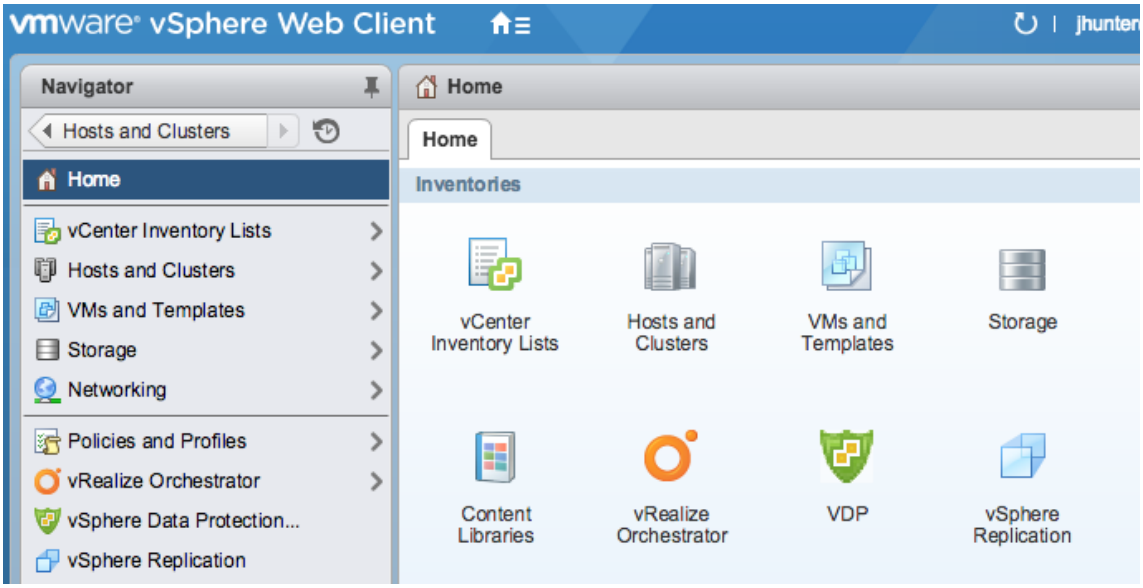

Figure 3. vSphere Data Protection in vSphere Web Client

Recommendations: If the vSphere Data Protection menu option does not appear after the vSphere Data Protection virtual appliance has rebooted and been fully configured—again, this might take as long as 30 minutes or more—log out of vSphere Web Client and log back in.

Review "Using vSphere Data Protection" in the vSphere Data Protection Administration Guide *before proceeding with the remaining exercises.*

NOTE: vSphere Data Protection has a set of alarms that are added to vCenter Server and enabled by default. Details on these alarms can be found in the vSphere Data Protection Administration Guide. *One alarm in particular can be triggered soon after vSphere Data Protection has been deployed: "Maintenance services are not running." This is expected because vSphere Data Protection maintenance services are disabled for the first 24 to 48 hours. This provides additional time—a larger backup window—that might be needed for the backup jobs to perform the initial full backups.*

Exercise 2: Create an Image Backup Job

This type of backup protects all components of a virtual machine—not only the information contained in the virtual machine (guest OS, applications, and data) but also the virtual machine configuration (number of virtual processors, amount of memory, and so on). This enables recovery of the entire virtual machine.

To create a new backup job or edit an existing backup job, see "Managing Backups" in the *vSphere Data Protection Administration Guide*.

Recommendations: The initial backup of multiple virtual machines will most likely take longer than the initial backup of one virtual machine. To verify functionality for this evaluation, configure the first backup job to back up a single virtual machine. After the backup job completes successfully, proceed with adding more virtual machines to the backup job.

After a backup job has been created, it will run as scheduled. Backup jobs should be scheduled to start and finish during the backup window. See Figure 4. For more details, see "Configuring VDP Details" in the vSphere Data Protection Administration Guide.

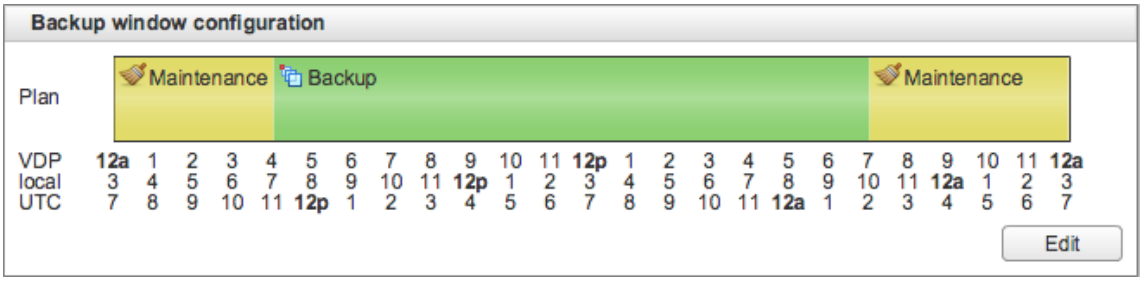

Figure 4. vSphere Data Protection Backup Window

After a backup job has successfully been completed at least once, it is possible to perform an image (entire virtual machine) restore and an FLR. Subsequent backup jobs should take much less time due to the Changed Block Tracking (CBT) feature of vSphere. More details on CBT can be found here.

Exercise 3: Run a Backup Job Manually

Recommendation: Running a backup job should be avoided during the maintenance window. During this time, vSphere Data Protection performs maintenance activities such as integrity check and garbage collection.

There are two methods for manually initiating a backup job:

1. On the **Backup** tab, click **Backup now** and choose one of the two options from the menu. See Figure 5. **Backup all sources** backs up all virtual machines or applications in the backup job regardless of the last time they were backed up. **Backup only out of date sources** backs up the virtual machines or applications that failed to back up the last time the backup job ran.

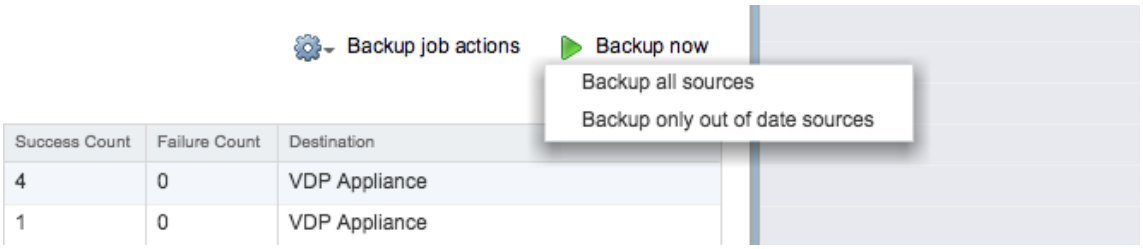

Figure 5. "Backup Now" Menu Options

2. In vSphere Web Client, select a virtual machine and click **Actions** or right-click a virtual machine. Then click **All VDP Actions** > **Backup Now**. See Figure 6. **Backup Now** is available for any virtual machine in the environment. For this operation to succeed, however, the virtual machine must be part of an existing backup job.

Figure 6. "All VDP Actions > Backup Now"

Backing up multiple virtual machines concurrently can generate a considerable amount of load on the storage system. In environments containing storage with lower performance characteristics, it might be necessary to limit the number of simultaneous backup and restore clients. This can be accomplished by managing the proxy throughput of vSphere Data Protection. See "Proxies" in the v*Sphere Data Protection Administration Guide* for more information.

Exercise 4: Restore a Virtual Machine

It is very easy to restore a virtual machine with vSphere Data Protection. There are several options available during the restore process, such as renaming the virtual machine (as registered in vCenter Server—not the guest operating host name). It is also possible to restore a virtual machine to an alternate location in the vCenter Server hierarchy or to a different storage location.

When restoring a virtual machine to its original location when the original virtual machine is still intact, vSphere Data Protection can leverage CBT during the restore process. As a result, restore times when "rolling back" a virtual machine to a restore point can be significantly reduced. For more information on the use of CBT by vSphere Data Protection during a restore, see "Image-Level Backup and Restore" in the *vSphere Data Protection Administration Guide*.

Recommendation: Perform a virtual machine restoration to its original location ("rolling back" an existing virtual machine to a restore point) and restore a virtual machine to a new location to observe the benefit of vSphere Data Protection using CBT.

Performing an image-level restore basically consists of two steps:

1. Locate and select the desired restore point.
2. Specify the location where the virtual machine will be restored.

Detailed steps and an explanation of the options available during an image-level restore can be found in the "Managing Restores" section of the *vSphere Data Protection Administration Guide*.

Restore Rehearsals

A restore rehearsal can easily be initiated by right-clicking a protected virtual machine in vSphere Web Client and choosing **All VDP 6.0 Actions** and **Restore Rehearsal**. See Figure 7.

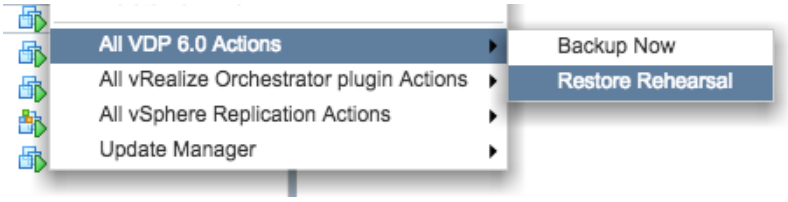

Figure 7. "Restore Rehearsal"

There are options to power on the virtual machine and reconnect the virtual network interface card (vNIC) of the virtual machine after it has been restored. These are useful when performing a restore rehearsal to verify that a functional virtual machine can be recovered.

Recommendation: Perform restore rehearsals often. vSphere Data Protection is a reliable solution based on the proven, mature Avamar product. With any backup solution, it is a best practice to routinely verify that virtual machines can be restored.

Exercise 5: Restore an Individual File

Using the FLR function, vSphere Data Protection can restore individual files and folders within a virtual machine running Windows or Linux. There is no requirement to install an additional backup software client to perform an FLR. A supported Flash-enabled Web browser is used to access vSphere Data Protection Restore Client. See Figure 8. This enables end users such as guest OS administrators and application owners to perform self-service individual file restores without the installation or complexity of backup client software.

There are some requirements and limitations to the FLR functionality of vSphere Data Protection. For example, the credentials that are used to perform an FLR must have administrator or root permissions in the guest OS of the local virtual machine. Requirements and limitations are fully documented in the "Using File Level Restore" section of the *vSphere Data Protection Administration Guide*.

Recommendation: vSphere Data Protection does not support FLRs for all logical partitions when multiple logical partitions are configured on the same VMDK file—for example, a Windows server with the C: and D: drives on one VMDK file. To work around this limitation, configure separate VMDK files for each logical partition. This approach also makes it easy to back up data in a more granular fashion. Consider this scenario: An administrator wants to create a backup job that backs up only the data belonging to an application. The OS and application are installed on the C: drive. The application data resides on the D: drive. A backup job can then be configured to back up only the VMDK file that contains the D: drive.

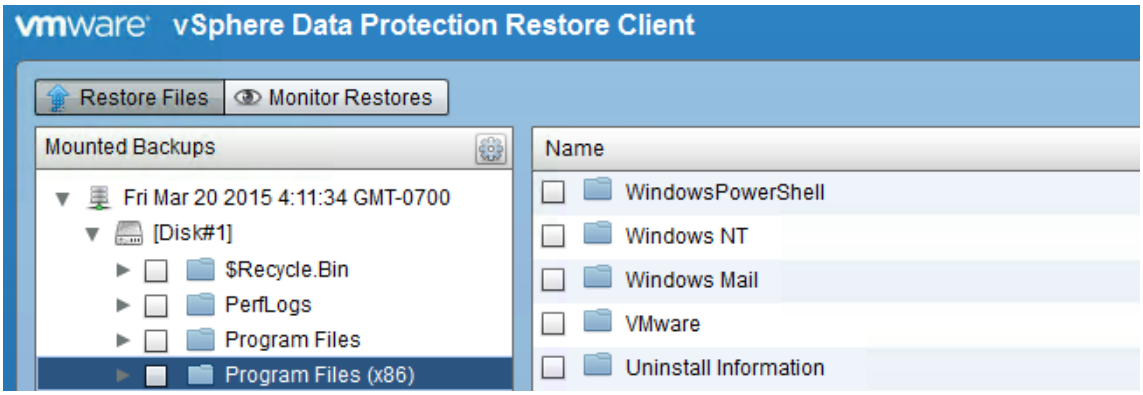

Figure 8. vSphere Data Protection Restore Client

When restoring files, folders, or directories, users can select the original location as the destination for the restored item, which will overwrite an existing item with the same name, or they can select an alternate location.

Users can monitor the progress and results of the restore job by clicking **Monitor Restores** in vSphere Data Protection Restore Client.

For detailed steps on using vSphere Data Protection Restore Client to perform an FLR, see "Using File Level Restore" in the *vSphere Data Protection Administration Guide*.

Exercise 6: Create an Application Backup

vSphere Data Protection can perform application-consistent database backups and restores of Exchange Server, SQL Server, and SharePoint Server.

Recommendations: For evaluation purposes, it is easiest to clone an existing virtual machine with the supported application already installed and to use this clone in the evaluation environment. The evaluation environment must be isolated from the production environment to prevent conflicts with production systems.

Review "VDP Application Support" in the vSphere Data Protection Administration Guide *for the list of supported applications, versions, and prerequisites that must be in place prior to installing the application agent and creating an application backup job.*

Agent Installation

vSphere Data Protection includes client agents for specific applications. The agents are downloaded from the **Configuration** tab in the vSphere Data Protection UI. See Figure 9.

Figure 9. Client Agent Downloads in vSphere Data Protection

Agents are installed in the guest OS of the virtual machines where the applications are installed. The name of the vSphere Data Protection appliance is entered during the agent installation. Use the FQDN when entering the name of the vSphere Data Protection appliance. Agent support for SQL Server clusters and Exchange Server DAG requires additional configuration beyond installing the agent. See "Installing VDP for SQL Server Client" and "Installing VDP for Exchange Server Client" in the *vSphere Data Protection Administration Guide* for details. After agent installation and configuration have been completed, an application backup job can be created. See Figure 10.

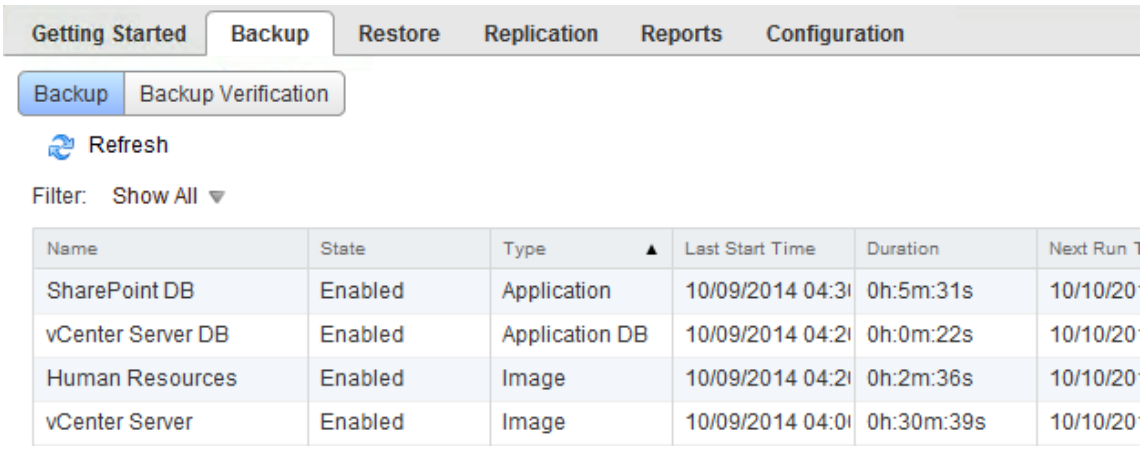

Figure 10. Application Backup Job in vSphere Data Protection

When protecting Exchange Server with vSphere Data Protection, an account must be created in Exchange Server for vSphere Data Protection as part of the agent installation. By default, this user is named **VDPBackupUser.** The agent installation makes the creation of this account easy, including the configuration of group memberships required for this user account. For details on the required Microsoft Active Directory and Exchange Server group memberships, see "Backing Up and Restoring Microsoft Exchange Servers" in the *vSphere Data Protection Administration Guide*.

Backup Job Creation

With the client agent(s) deployed to the evaluation environment, it is possible to create and edit application backup jobs. The backup-job wizard provides administrators with multiple options such as selecting the type of backup, enabling multiple backup streams, and managing database logs. Figure 11 shows the first step in creating an application backup job.

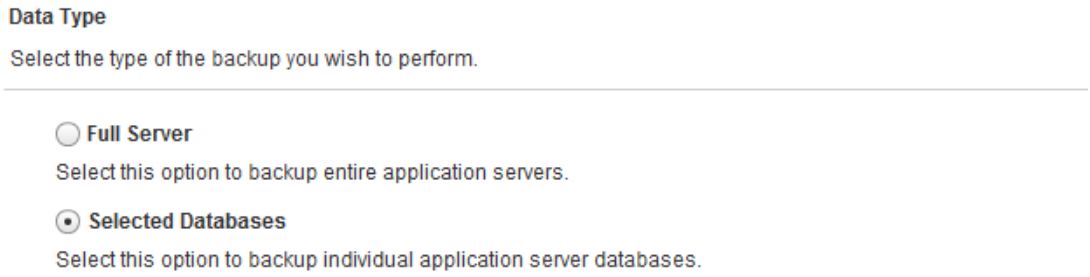

Figure 11. Full Server (All Databases) or Selected Databases Backup Job Option

Selecting **Full Server** backs up all application databases; **Selected Databases** provides the flexibility to select individual databases for backup, as is shown in Figure 12.

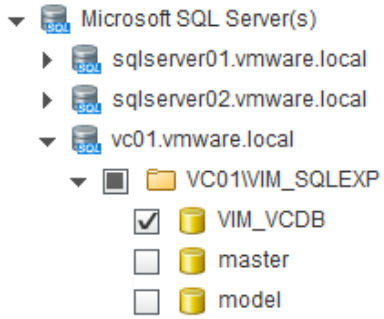

Figure 12. Selecting Individual Databases for Backup

A full list and descriptions of application backup options are available in the "VDP Application Support" section of the *vSphere Data Protection Administration Guide*.

Clicking **Enable multi-stream backup**, as is shown in Figure 13, will likely reduce the amount of time needed for backup. However, it might also increase resource utilization on the client. This feature should be used with caution so as not to impact performance if the application is in use.

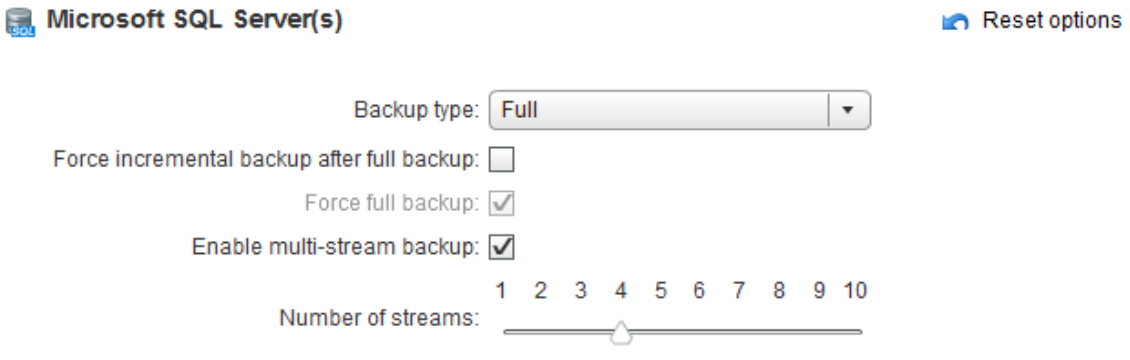

Figure 13. "Enable Multi-Stream Backup" for Microsoft SQL Server

After selection of the items to be backed up, the remaining steps (schedule, retention policy, and so on) are very similar to those for creating an image-level backup job.

Exercise 7: Restore an Application

With vSphere Data Protection, administrators can choose to restore an entire application (all databases) or individual databases. To begin the process, click the **Restore** tab. Backup types can be seen in Figure 14.

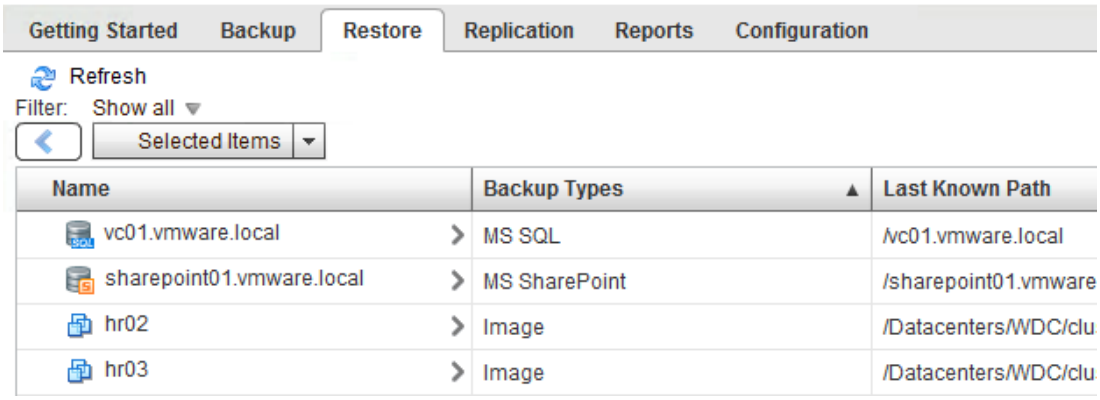

Figure 14. Backup Types on the Restore Tab

Click the application to reveal the list of restore points. Additional clicks reveal more-granular selection options such as individual databases in SQL Server and mailboxes in Exchange Server (Figure 15).

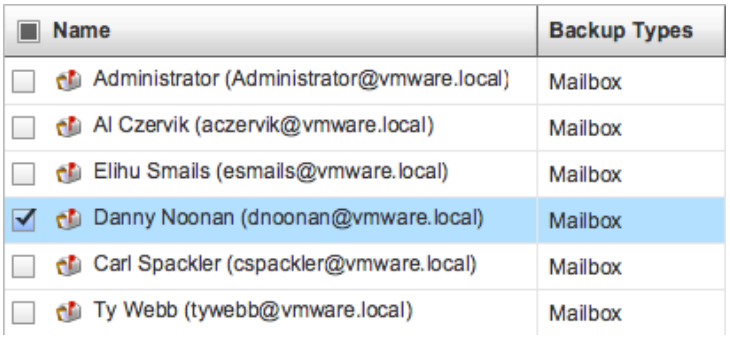

Figure 15. Restoring an Exchange Server Mailbox

As with guest-level and application backups, there are many options for restoring these protected applications. **Restore to original location** is a default setting, but it is possible to restore to an alternate location. Figure 16 shows additional options available when **Restore to original location** is not checked.

Figure 16. "Restore to Original Location" Not Checked

The options for application restore are detailed in the "VDP Application Support" section of the *vSphere Data Protection Administration Guide*.

Exercise 8: Configure Email Reporting

vSphere Data Protection features a **Reports** tab in the UI that provides information on the health of vSphere Data Protection, capacity status, and success or failure of recent backup jobs. The tab also provides details on each backup job, replication job, and backup verification job.

vSphere Data Protection can also be configured to send email reports with these details. The day(s) of the week and the time of day when the reports are sent are configurable—for example, Monday through Friday at 7:00 a.m. Figure 17 shows a screenshot of the email reports configuration UI. More information on configuring this feature can be found in the "Configuring VDP Details" section of the *vSphere Data Protection Administration Guide*. This exercise is optional and does not impact the primary backup and recovery functionality of vSphere Data Protection.

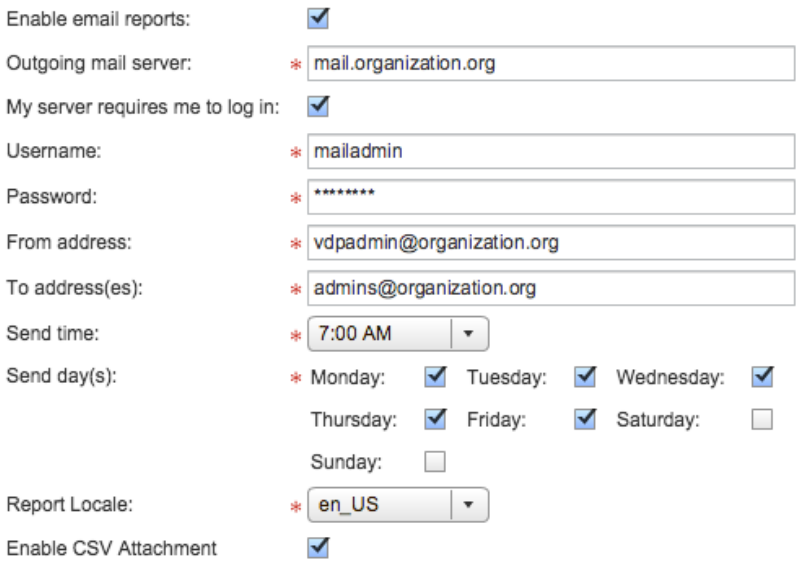

Figure 17. Email Reports Configuration User Interface

Exercise 9: Replicate Backup Data

vSphere Data Protection can replicate backup data to another vSphere Data Protection virtual appliance and to EMC Avamar Data Store. vSphere Data Protection replication provides a reliable, network-efficient, secure means to move backup data to another location within the same site or offsite. Backup data is deduplicated and compressed in each vSphere Data Protection virtual appliance. Only unique data segments are replicated from the source to the target. This approach minimizes the amount of data that must be replicated across the network. The replicated backup data is also encrypted for security.

Replicated backup data can be replicated again either back to the source vSphere Data Protection virtual appliance or to another vSphere Data Protection virtual appliance. This enables considerable flexibility in where the backup data can be restored and how many copies of the backup data are maintained. Here are a few examples to further illustrate this concept:

Example 1: A vSphere Data Protection virtual appliance performs backups at site A. This backup data is replicated to a vSphere Data Protection virtual appliance at site B. Because a copy of the backup data resides at both site A and site B, restores can be performed locally at either site.

Example 2 (building on Example 1): If disaster strikes site A and all virtual machines are lost, a new vSphere Data Protection virtual appliance is deployed when site A is back online. Backup data at site B can be replicated to the new vSphere Data Protection virtual appliance at site A, and the virtual machines can be recovered.

Recommendation: Before proceeding with this exercise, read the "Replication" subsection of the vSphere Data Protection Administration Guide *to better understand the replication capabilities of vSphere Data Protection.*

Creating a Replication Job

The first step in creating a replication job is selecting the backup data to be replicated. It is possible to select specific virtual machine or application backup data for replication. For example, an administrator can create a replication job that replicates backup data for two virtual machines (Figure 18).

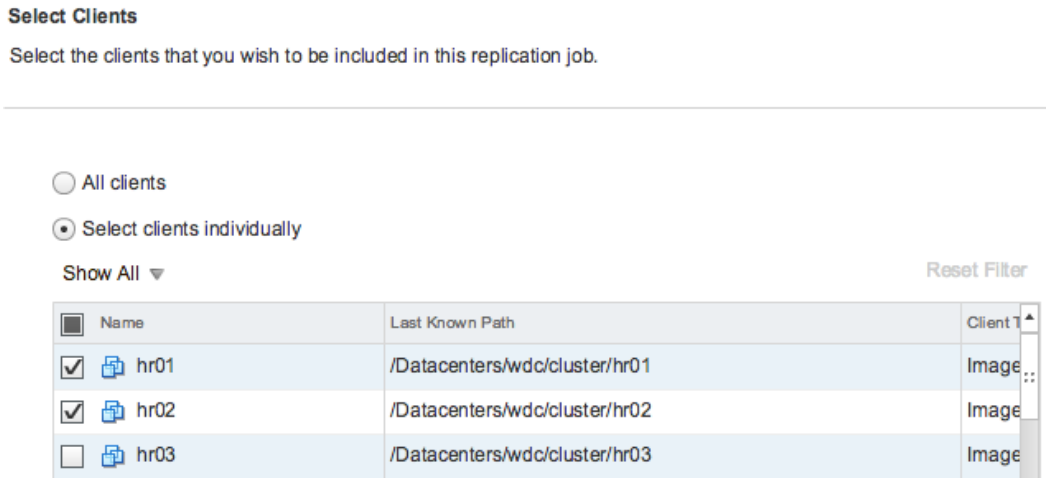

Figure 18. Virtual Machines (Backup Data) Selected for Replication

The next step in creating a replication job is defining options to limit the types and number of backups replicated. For example, an administrator might want to limit the job to three daily backups that occurred in the past 7 days. Details on the available options are found in the "Creating a Replication Job" subsection of the *vSphere Data Protection Administration Guide*.

By default, the retention policy assigned to the replicated backup data is the same as that assigned by the backup job. It is possible to set a different retention policy for replicated backup data. Consider this scenario: The administrator creates a backup job with a retention policy of 30 days. There is a requirement to keep weekly backups offsite for 12 months. The administrator creates a replication job and sets the retention policy accordingly (Figure 19).

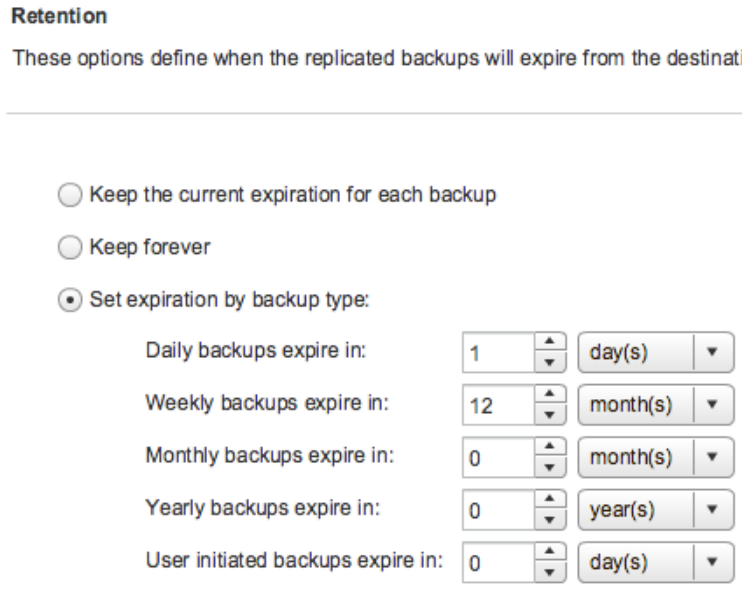

Figure 19. Replication Job Retention Policy

After the replication job has been created, click **Replicate now** to manually start the job. A replication job must be completed at least once to perform the next part of the exercise, "Restoring a Virtual Machine from Replicated Backup Data."

Restoring a Virtual Machine from Replicated Backup Data

Performing restores from replicated backup data is very similar to performing a restore from local backup data. Figure 20 shows restore points for replicated backup data in a vSphere Data Protection replication.

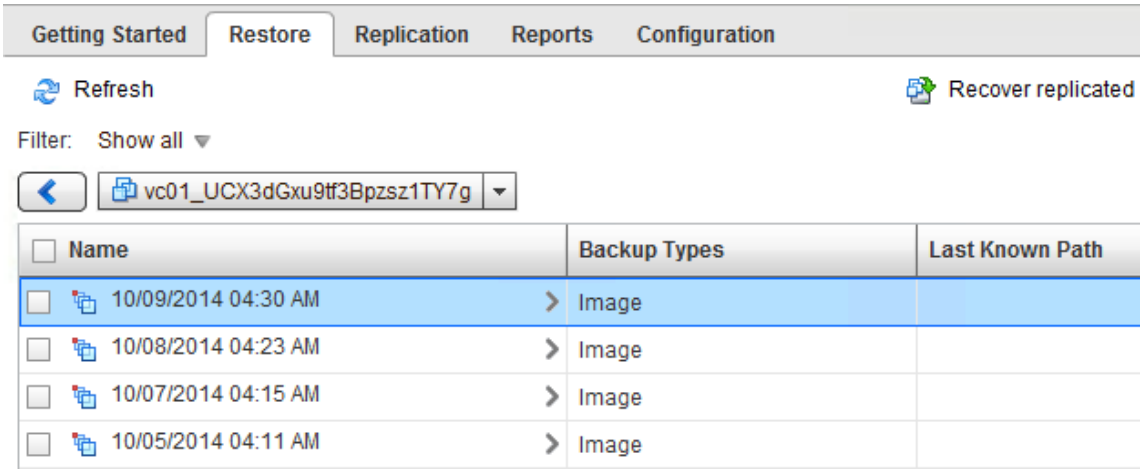

Figure 20. Replicated Backup Data

Nearly all of the same restore options are available for replicated backup data. One exception is application-level restores such as an individual SQL Server database. Application data can be restored only to virtual machines that have the vSphere Data Protection agent installed and registered to the appliance from which the restore is being performed. Consider this scenario: An SQL Server virtual machine named SQL1 contains a database that is backed up by a vSphere Data Protection virtual appliance named VDP1. VDP1 replicates this backup data to VDP2. The replicated backup data cannot be restored directly to SQL1 unless the SQL Server agent is reinstalled in SQL1 and associated with VDP2 rather than VDP1.

Exercise 10: Create a Backup Verification Job

The best way to verify that backup data can actually be restored is to perform "practice restores" on a regular basis. vSphere Data Protection automates this process to provide the highest level of confidence in backup data integrity. A backup verification job is created and scheduled to run daily, weekly, or monthly. At the scheduled time, a virtual machine is restored, disconnected from the network, and powered on. vSphere Data Protection validates the successful restore of the virtual machine by detecting VMware Tools heartbeats. Optionally, an administrator can define a script that runs in the guest OS of the restored virtual machine to provide further verification—for example, a script that verifies that an application service has started. See Figure 21. When the verification process completes, vSphere Data Protection deletes the restored virtual machine. Details about the success or failure of a backup verification job are found in the **Reports** tab and in email reports sent by vSphere Data Protection.

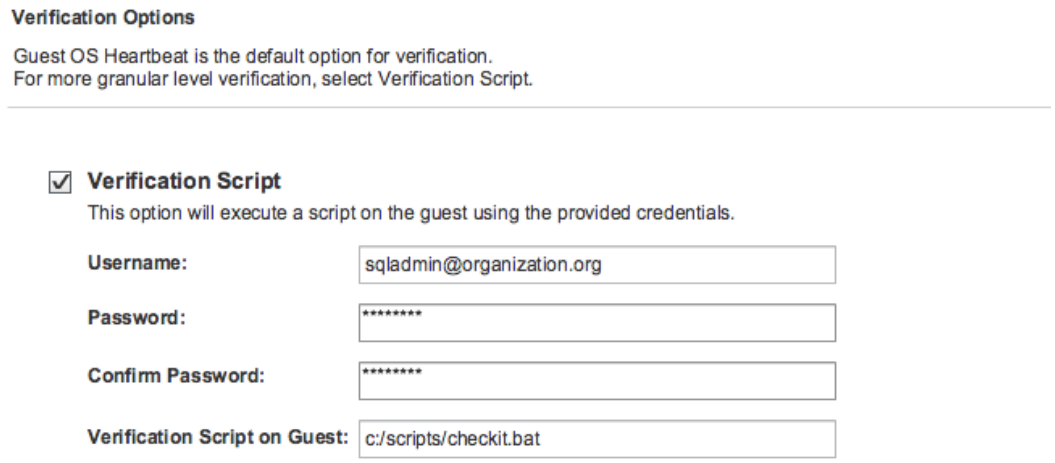

Figure 21. Backup Verification Script Configuration

Creating a backup verification job is very similar to performing a virtual machine restore. Detailed steps can be found in the "Automatic Backup Verification" section of the *vSphere Data Protection Administration Guide*.

Conclusion

VMware vSphere Data Protection is an effective data protection solution designed primarily for small and midsized environments. It enables quick, simple, and complete data protection for virtual machines and important workloads such as Microsoft Exchange Server and Microsoft SQL Server. vSphere Data Protection incorporates EMC Avamar deduplication backup and recovery software to help ensure reliable backups and restores while minimizing backup data storage consumption. It features a robust backup data replication engine to easily move backup data offsite in an efficient, secure, and reliable manner. vSphere Data Protection also utilizes VMware vSphere Storage APIs – Data Protection, including Changed Block Tracking, which facilitates image backup and recovery without the need for an in-guest client agent.

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