Veeam Backup & Replication 5.0

User Guide

April 2011

Important! Please read the End User Software License Agreement before using the accompanying software program(s). Using any part of the software indicates that you accept the terms of the End User Software License Agreement.

CONTENTS

ABOUT THIS GUIDE

Overview

This user guide provides information about main features, installation and use of Veeam Backup & Replication 5.0. The document is intended for network administrators, consultants, analysts and any other IT professionals using the product.

Conventions

In order to help you get the most out of this guide, we have used the following formatting conventions, terms and abbreviations in the document:

Convention	Description
Emphasis	This type of formatting is used to designate user interface elements (names of dialog windows, buttons and so on).
Italics	This type of formatting is used to designate names of files, file paths, cross-references and options of choice (for example, in a drop-down list).
Notes	This type of formatting is used for tips, notes and important information the user should pay attention to.

Term/Abbreviation	Description
Datastore	Storage for a managed ESX server
NAS	Network attached storage
SAN	Storage area network
VCB	VMware Consolidated Backup
VM	Virtual machine
VMFS	Virtual machine file system
VSS	Windows Volume Shadow Copy Service
Source host	A host on which a VM to be backed up or replicated resides or where a restored VM should be started.
Target host	A host where a created backup should be stored or replica should be started, or from which VM data to be restored is retrieved.

ABOUT VEEAM SOFTWARE

Contacting Veeam Software

At Veeam Software we pay close attention to comments from our customers. It is important to us not only to quickly help you with your technical support issues — we make it our mission to listen to your input, and to build our products with your suggestions in mind.

Should you have a Customer Support issue or question, please feel free to contact us. We have qualified English speaking technical and customer support staff in the USA and Europe who will help you with any inquiry that you may have.

Phone Support

United States	+1 (614) 339 8252
UK	+44 (845) 508 70 05
Germany	+49-2241-955-89-99
France	+33 (1) 70 61 83 74
Spain	+34 (91) 182 97 60
New Zealand	+64 (9) 974-9594
Australia	+61 (26) 108-4305
Netherlands	+31 (858) 88 06 55
Norway	+47 (85) 40 43 85
Denmark	+45 (78) 77 54 76
Belgium	+32 (78) 48 02 54
Sweden	+46 (10) 199 25 77

Company Contacts

Office	**U.S. Headquarters**	**EMEA Headquarters**	**APAC Headquarters**	**Veeam Software Benelux**
Address	6479 Reflections Drive, Suite 200 Columbus, Ohio 43017 USA	400 Thames Valley Park Thames Valley Park Drive Reading, Berkshire RG6 1PT UNITED KINGDOM	Level 21 & 22, 201 Miller Street North Sydney NSW 2060, AUSTRALIA	Evert van de Beekstraat 310, 1118 CX Schiphol Centrum THE NETHERLANDS
Phone	+1-614-339-8200	+44 (0) 1276-804-501	+61 2 8014 7476	+31(0)20 654 18 05
Fax	+1-614-675-9494	+44 (0) 1276-804-676	+61 2 8088 6899	+31(0)20 654 1801

Contacting Veeam Support

We offer e–mail and phone technical support for customers on maintenance and assistance during the evaluation period. For better experience please provide the following when contacting our technical support:

- Information about the operating system and database you are using.
- Error message and/or accurate description of the problem.
- Log files. To browse to the log files, select **Help > Support Information...** from the main menu.

To submit your support ticket or obtain additional information, please visit http://www.veeam.com/support.html.

Note: Before contacting technical support, you may be able to find a resolution to your issue at Veeam Technical Support Forum at: http://www.veeam.com/forums/

OVERVIEW

Veeam Backup & Replication 5.0 is a disaster recovery solution for VMware infrastructure that combines backup and replication, as well as the fastest file-level restore, in a single product. Enabling these options from one interface, it serves to solve most critical problems of the VMware infrastructure management and protects mission-critical virtual machines from both hardware and software failure.

Veeam Backup & Replication 5.0 shares a common interface with Veeam FastSCP, file management freeware, allowing you to manage backup, replication and file copying jobs from a single console.

Veeam Backup & Replication 5.0 provides the following features and functionality:

VMware ESX / ESXi Support

Veeam Backup & Replication 5.0 provides full support for VMware ESX Server and ESXi for backup with or without the VMware Consolidated Backup (VCB) proxy, restore and failover processes.

Native vSphere and vStorage Support

Veeam Backup & Replication 5.0 features native support for VMware vSphere, including all vSphere and vStorage functionality: support for thin-provisioned disks, ESX4 changed block tracking, new vStorage APIs for Data Protection, vApps and virtual backup appliances using the vSphere HotAdd technology.

Veeam vPower

Veeam Backup & Replication 5.0 offers vPower — a new patent-pending technology enabling you to start a VM directly from a compressed and deduplicated backup file. vPower eliminates the need to extract a VM from the backup file and allows you to:

- Immediately recover a failed VM, thus reducing downtime of production VMs to the minimum.
- Verify recoverability of every backup by starting and testing VMs directly from backups in the isolated environment (SureBackup).
- Restore items from any virtualized applications with U-AIR (Universal Application Item-Level Restore).

Veeam Backup Enterprise Manager

Veeam Backup & Replication 5.0 comes with Veeam Backup Enterprise Manager — a management and reporting component that allows you to manage multiple Veeam Backup & Replication installations from a single web console. In case of distributed backup infrastructure, Veeam Backup Enterprise Manager acts as a single management point, allowing you to perform backup and replication jobs across the entire backup infrastructure, and providing enhanced notification and reporting options.

Veeam Backup Enterprise Manager performs the role of a coordinator in U-AIR procedures, allowing you to monitor and delegate recovery processes. It is also responsible for replicating and consolidating index files from backup servers to enable file browsing and search functionality, and acts as a license center, allowing you to centrally update licenses and get statistics on their usage.

Guest OS Files and VM Files Recovery

Veeam Backup & Replication 5.0 provides a possibility to perform granular VM guest OS file- or folder-level recovery for FAT16, FAT32 and NTFS file systems without extracting a full VM image to the local drive. The file-level restore for VMs running other file systems can be performed with the Veeam File Level Restore wizard.

Along with VM OS files recovery, it allows restoring specific VM files (VMDK, VMX, etc). Individual files or folders can be restored to their latest state or to a specific point in time.

Windows Volume Shadow Copy Service (VSS) Support

Veeam Backup & Replication 5.0 supports Windows Volume Shadow Copy Service (VSS) enabling backup and replication on live and open systems running Windows applications or working with databases (for example, Domain Controller, Exchange Server, SQL Server) without shutting them down.

Veeam Backup & Replication 5.0 provides advanced options to control truncating of transaction logs so that you can ensure correct backup of applications that use transaction logs, and meet requirements of any backup scenario. You can select to truncate transaction logs after each backup job, each successful backup or not to truncate logs at all.

File Indexing and Search

Veeam Backup & Replication 5.0 indexes guest OS files in Windows-based VMs, allowing you to perform quick and accurate search for files within backed up VM images without the need to restore them. Using Veeam Backup Enterprise Manager, you can browse and search for files in a single selected VM backup or use the advanced search option to find necessary files in all VM backups within your backup infrastructure.

Incremental and Reversed Incremental Backup

Depending on the type of backup storage you use, you can choose between two backup methods — incremental and reversed incremental. Incremental backup is recommended for disk-to-disk-to-tape and remote site backups — it reduces the time spent to move backups to tape or a remote site, and the amount of tape required. Reversed incremental backup is recommended for disk-to-disk backup, allowing you to keep the latest image of a VM in a ready-to-restore state on disk. With advanced options from Veeam Backup & Replication 5.0, you can select to perform incremental backup and schedule creation of synthetic full backups on specific days, which will let you combine advantages of incremental backup with those of reversed incremental.

Granular Backup Options

Along with backing up a VM as a whole, Veeam Backup & Replication 5.0 lets you back up specific VMDK disks — for example, only system disks or disks with application data. Additionally, you may choose to include VM templates into backup (either in both full and incremental backups or in the full backup only).

Data De-Duplication and Compression

In order to decrease the size of created backups, Veeam Backup & Replication 5.0 de-duplicates identical blocks inside a backup file. Higher de-duplication rates are achieved when backing up multiple VMs created from a single template or VMs with gigabytes of free space within. You can also decrease the backup file size by using compression.

Reporting

Veeam Backup & Replication 5.0 features comprehensive real-time job statistics (start/end time, performance metrics), as well as the current job activity description. With the reporting option, you can generate HTML reports with statistics for the performed job.

ARCHITECTURE

This chapter provides a high-level overview of the Veeam Backup & Replication 5.0 architecture and functionality.

Veeam Backup & Replication Components

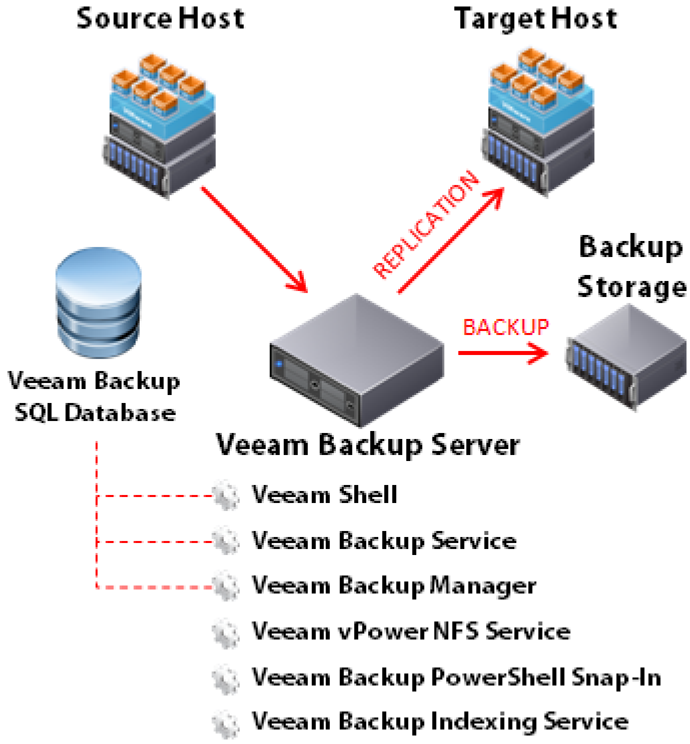

Veeam Backup & Replication 5.0 comprises a set of components responsible for performing the following functions.

Veeam Backup Server:

- *Veeam Shell* provides the application user interface enabling control over Veeam Backup & Replication 5.0 and access to its functionality.
- *Veeam Backup Service* is a Windows service running on the Veeam Backup console responsible for scheduling and coordinating backup, replication and copying jobs. Veeam Backup Service runs under the administrator account with the *Log on as service right* granted.
- *Veeam Manager* is a Windows process running on the Veeam Backup & Replication console that is activated by Veeam Backup Service at startup of every job. Veeam Manager controls Veeam Agents on the source and target hosts to perform jobs according to the set job parameters.
- *Veeam vPower NFS service* is a Windows service that enables the Veeam Backup server to act as an NFS server. Veeam vPower NFS service provides ESX servers with transparent access to backed up VM images and saves changes that take place when a VM is up and running. This service is used for recovery verification, instant VM recovery and U-AIR procedures.
- *Veeam Indexing service* is a Windows service that manages VM guest file indexing and replicates system index data files to enable search through VM guest OS files. Veeam Indexing service running on the Veeam Backup server works in conjunction with search

components installed on Veeam Backup Enterprise Manager and a dedicated search server.

- *Veeam Backup PowerShell snap-in* allows users to automate backup and replication tasks by running single cmdlets or custom automation scripts via the command-line interface.
- *U-AIR wizards* allow users to perform item-level restore from any applications running on VMs. U-AIR wizards can be installed on Veeam Backup & Replication console, or any machine in the production environment.

ESX hosts/Veeam Backup server:

- *Veeam agents* are deployed on target and source hosts, or on the Veeam Backup & Replication console (depending on the selected backup mode) and are controlled by the Veeam Manager that initiates work of all agents on hosts simultaneously. Veeam agents are responsible for general activities performed within the frames of a specific job: scanning virtual machine file systems, communicating with VMware Tools utilities, copying VM files, performing data de-duplication and compression and so on.

SQL Server:

- *Veeam SQL Database* with which Veeam Shell, Veeam Backup Service and Veeam Manager communicate in the process of work is used to store program activities and job relevant data (job options, performance metrics and statistics). You can use SQL Server 2005 or SQL Server 2008 installed locally or remotely. During installation, the Veeam Backup setup installs a new SQL Server 2005 Express Edition instance, creates a new *VeeamBackup* database on the existing SQL Server, or connects to the *VeeamBackup* database installed by the previous version of Veeam Backup & Replication.

Veeam Backup Enterprise Manager:

- *Veeam Backup Enterprise Manager* is a management and reporting component that aggregates data from multiple Veeam Backup & Replication servers and provides centralized control over these servers from a single web console.
- *Veeam Indexing service* running on the Veeam Backup Enterprise Manager server performs catalog replication and consolidation, and manages work of search servers (if any are deployed).

Search server:

- *MOSS Integration Service* installed on a dedicated pre-configured Microsoft Search Server initiates updates of index databases on Microsoft Search Server. It also filters and initiates execution of search queries that should be performed by the search server when requested by the Indexing service on Veeam Backup Enterprise Manager.

Backup Methods

Veeam Backup & Replication 5.0 offers two backup methods to back up virtual machines:

- Reversed incremental, or synthetic backup
- Incremental backup

Reversed incremental backup is recommended for disk-based backup. Reversed incremental backup implies that during the first run of a backup job a full backup of a VM is created. VM data is copied block by block, compressed using an appropriate compression level, and stored in a resulting full backup file (.vbk). All subsequent backups are incremental — Veeam Backup & Replication 5.0 backs up only those data blocks that have changed since the last job run. During incremental backup, Veeam Backup & Replication "injects" changes into the created .vbk file to rebuild it to the most recent state of a VM. It also creates a reversed incremental backup file (.vrb) containing data blocks that were replaced when the full backup file was rebuilt. Therefore, the most recent point-in-time is always a full backup which gets updated after each backup cycle.

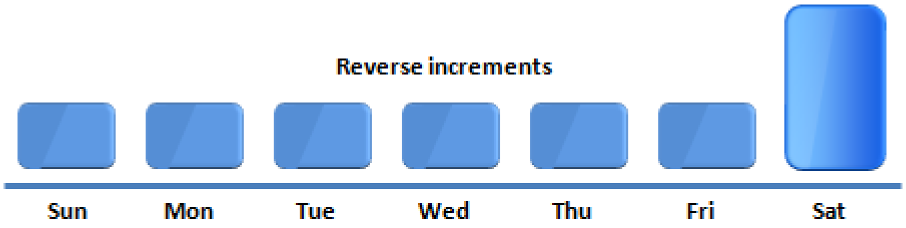

This backup method enables you to perform forever-incremental backup and save disk space as you have to store only one full backup. With reversed incremental backup, you do not have to perform periodic full backups to keep up with the specified retention policy. If the number of backups allowed by the retention policy is exceeded, Veeam Backup & Replication 5.0 will simply delete the oldest reversed increments.

Reversed incremental backup also enables you to immediately restore a VM to the most recent state without extra processing. If you need to restore a VM to a particular point in time, Veeam Backup & Replication 5.0 will apply related .vrb files to the .vbk file in the reversed order to get you to that point in time.

Incremental backup is recommended for disk-to-disk-to-tape and remote site backups. If this method is selected, Veeam Backup & Replication 5.0 creates a full backup file (.vbk) at the first run of a 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) next to a full backup.

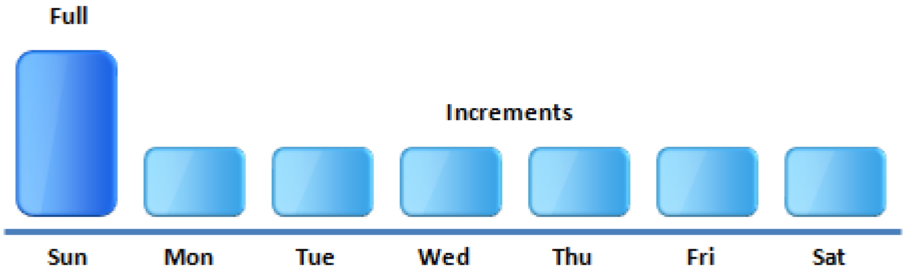

Incremental backup is the best choice if company regulation and policies require you to regularly move a created backup file to tape or a remote site. With incremental backup, you move only incremental changes, not the full backup file, which takes less time and requires less tape. Writing backups to tape or a remote site can be initiated through Veeam Backup & Replication 5.0 — to learn more, see the Integration with Traditional Backup section.

To let you get the most out of backup, Veeam Backup & Replication allows you to run incremental backups and schedule creation of **full synthetic backup** on specific days. This option lets you combine advantages of tape-friendly incremental backup with those of reversed incremental — you can write small incremental changes to tape and at the same time have the latest VM image in a ready-to-restore state on disk.

For example, if you select an incremental backup with synthetic fulls scheduled on Thursday, Veeam Backup & Replication 5.0 will perform incremental backup through Sunday to Wednesday as usual. On Thursday, however, it will first create an increment, and then at the end of the backup job use the previous full backup and a chain of increments from Monday to Thursday to build a full VM backup so that you can have the latest state of the VM in a ready-to-restore state on disk. After that, it will delete the increment created on Thursday. Such mechanism will work only once a day on which it is scheduled — if you run the backup job once again on Thursday, Veeam Backup & Replication 5.0 will perform an incremental backup.

When creating a full synthetic backup, Veeam Backup & Replication 5.0 does not address VI to retrieve VM data – it uses the chain of full and incremental backups that are already kept on backup storage. After a full synthetic backup is created, Veeam Backup & Replication will use it as a starting point to create increments.

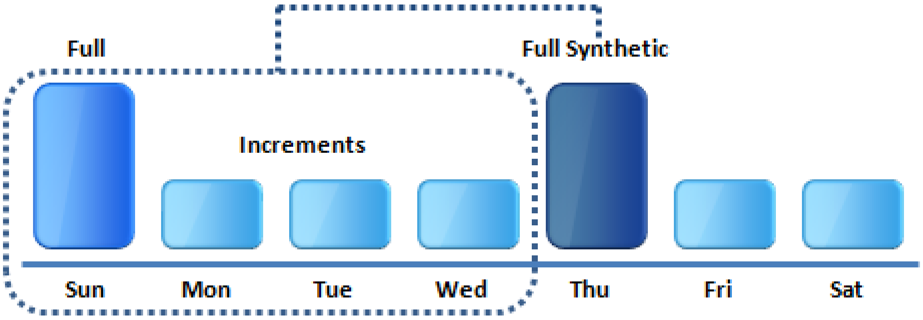

If you select to create a full synthetic backup, you can additionally choose to transform all previous full backup chains to the reversed incremental backup sequence. That is, all incremental .vib files, as well as a full backup created on Sunday will be transformed to reversed increments (.vrb), and you will only have a full backup created on Thursday. This option allows you to keep only one full backup image on disk and so reduce the amount of space required to store backups. However, such transformation takes more time than simply creating a full synthetic backup.

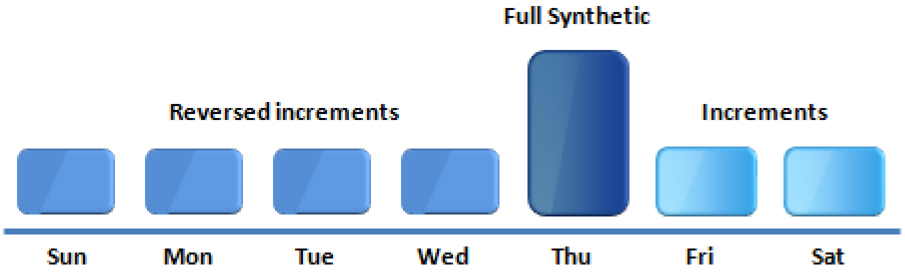

Note: If you choose the forward incremental backup method, you must select to perform either full synthetic or full active backups regularly for safety purposes and to keep up with specified retention policy. When full active backup is performed, Veeam Backup & Replication 5.0 simply creates full backups by the defined schedule. In contrast to the synthetic full backup, during full active backup it addresses VI to retrieve VM data from.

With Veeam Backup & Replication 5.0, you can easily switch the selected backup mode to the other one at any moment of time. In this case, Veeam Backup & Replication 5.0 will not transform the previously created chain — it will create a new chain next to the existing one. For example, if you switch from the reversed incremental backup mode to incremental one, it will create a new full backup next to the reversed incremental chain and will further use it to create increments.

Some companies have to obey regulations and policies requiring that a full backup is performed every time, or with certain periodicity. To conform to these requirements, Veeam Backup & Replication 5.0 offers an ability to configure a job for performing active full backups instead of forever-incremental backup. You can schedule full backups on specific week days or specific day of month, or instruct Veeam Backup & Replication 5.0 to create a full backup manually using the job's shortcut menu. Creation of a new full backup resets the chain of rollbacks or increments, so all subsequent backups processing will use the new full backup. A previously used full backup file will remain on disk until it is automatically deleted by backup retention policy, just like rollback files.

Backup Retention Policy

Backup retention policy controls for how long VM backups should be retained. In Veeam Backup & Replication 5.0, retention policy is defined by the number of VM restore points that should be kept. Once the specified number is exceeded, the earliest restore points are automatically removed.

To maintain retention policy, Veeam Backup & Replication 5.0 deletes not backup files on the whole, but restore points for separate VMs from the backup.

Let's imagine a backup job contains two virtual machines and retention policy is set to three restore points. During the first run of a job, both VMs were backed up. During the next job runs, one VM was first skipped for some reason (for instance, due to VSS failure), and at the third job run two VMs were successfully backed up again. As a result, we will have three restore points for one VM, and two restore points for the other one. During the next backup cycle, Veeam Backup & Replication 5.0 will remove one restore point for the first VM from the earliest backup file but will keep the backup file as it still contains a restore point for the second VM.

In case of reversed incremental backup, Veeam Backup & Replication 5.0 deletes the earliest reverse increments. In case of incremental backup, it takes into consideration a full backup and a chain of increments created after it. For example, if the retention policy is set to three restore points and there are a full backup and three increments, Veeam Backup & Replication 5.0 will not simply remove a full backup because that would make the whole chain inoperable. Instead, Veeam Backup & Replication 5.0 will wait for the next full backup to be created. Once the number of increments created starting from this new full backup exceeds two, Veeam Backup & Replication 5.0 will delete the whole previous chain of full backup and increments.

Automatic Retry of Backup Jobs

With Veeam backup & Replication 5.0, you can select to retry a backup job for several times if the initial backup of VMs fails.

When Veeam Backup & Replication 5.0 re-runs a job with several VMs during retry, it does not create a new backup file for failed VMs – it updates the backup file that has already been created. If some VM still fails to be backed up, Veeam Backup & Replication 5.0 will perform backup of a corresponding type for this VM during the next job run. For example, if a full backup of a VM failed, during the next job run Veeam Backup & Replication 5.0 will create a full backup for this failed VM, incremental backup for other VMs in the job, and write data to the same backup file.

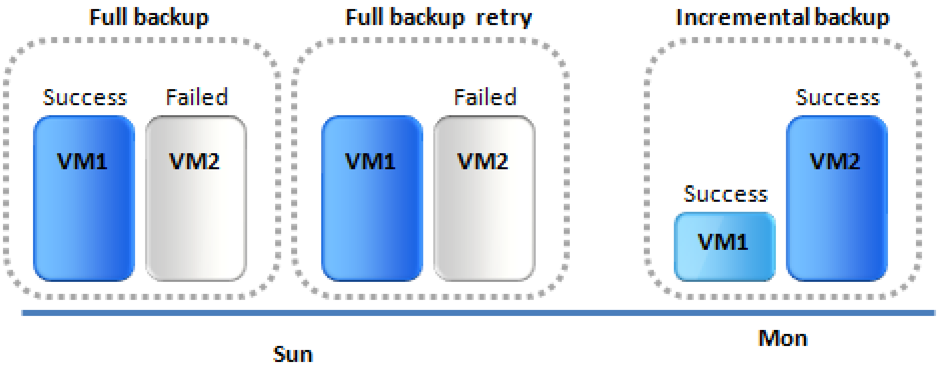

De-Duplication and Compression

To decrease disk space required for backup files, Veeam Backup & Replication 5.0 provides mechanisms of de-duplication and compression.

De-duplication is applied when backing up multiple virtual machines that have similar blocks within (for example, if virtual machines were created on the basis of the same template), or in case virtual machines with great amount of free space on their logical disks are backed up. Veeam Backup & Replication 5.0 does not store 0 byte blocks or space that has been pre-allocated but not used. With de-duplication, identical blocks or blocks of free space are eliminated, which decreases the size of the created backup file.

Depending on the type of storage you select as a backup target, Veeam Backup & Replication uses data blocks of different size to process VMs, which optimizes the size of a backup file and job performance. You can choose one of the following options:

- The **Local target** option is recommended if you use SAN, DAS or local storage as a target. SAN identifies larger blocks of data and therefore can process larger quantities of data at a time. This option provides the fastest backup job performance but reduces the de-duplication ratio — the larger a data block is, the lower is the chance to find an identical block.
- The **LAN target** option is recommended for backup to NAS and on-site replication. It provides a better de-duplication ratio and reduces the size of an incremental backup or replica file.
- The **WAN target** option is recommended if you are planning to use WAN for offsite backup. Veeam Backup & Replication 5.0 will use small data blocks, which will result in the maximum de-duplication ratio and the smallest size of a backup file, allowing you to reduce the amount of traffic over the WAN link.

Another means of reducing the size of a backup file is **compression**. Use of compression decreases the size of created backups but affects duration of the backup procedure. Veeam Backup & Replication 5.0 allows you to select one of the following compression levels:

- **No compression** is recommended if you use storage devices with hardware compression and de-duplication tools to store created backups.
- **Low compression** is an optimized compression level for very low CPU usage. It is recommended if you are backing up VMs to another ESX server and do not want to load it heavily.
- **Optimal compression** is the recommended compression level providing the best ratio between the size of a result file and time of the backup procedure.
- **Best compression** provides the smallest size of a backup file but may reduce backup performance. We recommend that you install Veeam Backup & Replication 5.0 on computers with modern multi-core CPU (8 cores recommended) if you intend to use best compression.

In case of virtual machine backup, both a full backup and subsequent increments are compressed. In case of virtual machine replication, compression is not performed for full replicas — it affects subsequent increments only.

Note: Please note that changing compression of an existing job will not have any effect on already created backup files — it will affect only those backups that were created after you set the new compression level.

Backup Content

When creating a backup or replication job, you can select to process separate VMs or VM containers. Veeam Backup & Replication 5.0 also enables you to add the whole datastore as a VM container to the job. This option can be useful if you are planning to back up or replicate VMs residing on the same datastore. Instead of performing several jobs targeted at the same datastore, which typically slows down the backup speed and can potentially cause the datastore to hang out, you can back up the whole datastore as one instance and so improve the backup job performance.

Alongside with a general case of backing up a VM or VM container as a whole, Veeam Backup & Replication 5.0 allows you to determine the content of the created backup by including or excluding specific elements from it: VM disks and VM templates.

In some situations it may be necessary to back up only **specific VM disks**. For example, you may want to back up only the system disk instead of creating a full backup which would take much more space than you actually require. Veeam Backup & Replication 5.0 provides the following options for disks selection: you may choose to back up all VM disks (selected by default), the 0:0 disks (which are commonly the system disks of a VM) or select custom disks at your discretion. Disk processing settings are specified granularly for each VM in the backup job.

When creating a job, you can select to include **VM templates** into the created backup. Backing up VM templates warranties supplementary safety of your production environment, though demands additional space. As a concession, Veeam Backup & Replication 5.0 allows you to include a VM template only in the full backup and omit it in all subsequent increments.

Backup Process

This chapter provides information on modes of backup that can be performed with Veeam Backup & Replication 5.0.

Note: You can use custom pre-freeze and post-thaw scripts before creating a snapshot of the virtual machine that is being backed up or replicated; this is done by means of VMware Tools. For more information about creating pre-freeze and post-thaw scripts please refer to VMware's documentation.

VMware vStorage API Backup

Veeam Backup & Replication 5.0 provides full support for vSphere. It uses VMware vStorage APIs for Data Protection to access VMware virtual disk storage and copy virtual disk data directly through SAN, thus enabling LAN-free backup of VMs.

vStorage APIs is the VMware recommended method for efficient off-host backup of vSphere virtual machines.

The VMware vStorage API backup can be used for VMware vSphere 4 (including ESX/ESXi), vCenter Server 4, VMware Virtual Infrastructure (ESX/ESXi 3.5 and VirtualCenter 2.5). Please note that ESX 3.0 is not supported.

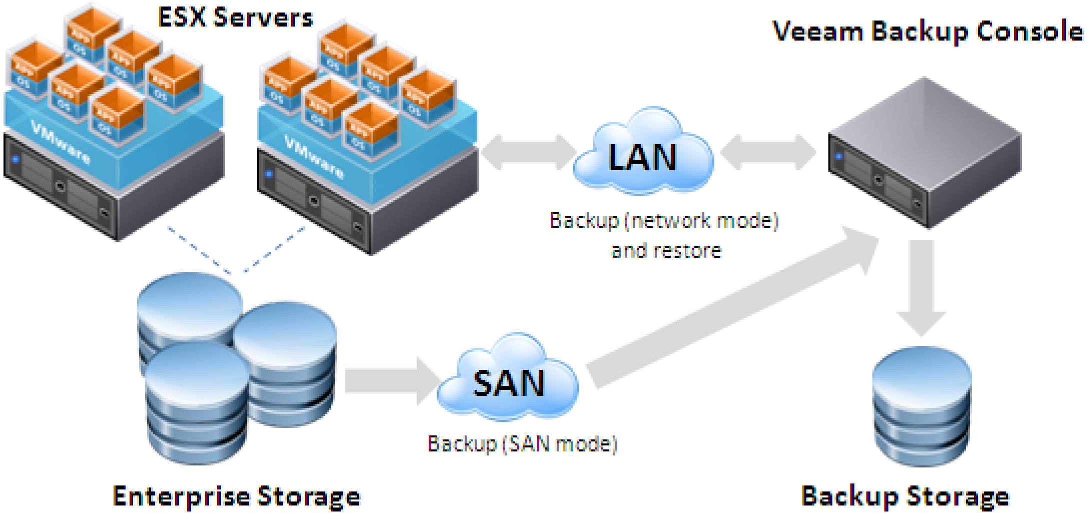

You can select to process your VMs in one of the three modes: **Direct SAN access**, **Virtual Appliance** and **Network** mode.

Note: The VMware vStorage API backup mode with ESX4 changed block tracking enabled is a recommended backup method and is used by default for created backup jobs. To learn more about ESX4 changed block tracking, see the Native vSphere and vStorage Support section.

Direct SAN Access

This mode is recommended if your ESX hosts are using shared storage. In this mode, VM data is retrieved directly from FC/iSCSI shared storage (Storage Area Network, or SAN) using the VMware vStorage API for Data Protection. The SAN mode uses metadata about layout of virtual disks on SAN to directly read data blocks off SAN LUN, providing, therefore, LAN-free transfer of VM data. Keep in mind that the Veeam Backup server must be connected directly into the SAN fabric for backup to work in this mode. VM processing will fail if direct SAN connection is not configured, or not available when the job starts.

With Veeam, you can also fail over to the network mode and retrieve VM data through the ESX host over LAN if SAN becomes inaccessible. The **Failover to network mode if primary backup mode fails** option is selected by default; it allows your backup jobs to still complete successfully. However, it puts additional load on your local area network and thus may potentially affect production environment if you are performing backup and replication during business hours.

Virtual Appliance

This mode is recommended and can only be used if Veeam Backup & Replication 5.0 is installed on a VM running on ESX/ESXi host. The Virtual appliance mode uses SCSI hot-add capability of ESX to attach disks of a backed up VM to the Veeam Backup & Replication VM, or to the helper VM (depending on vCenter version you are using). In this mode, VM data is retrieved directly from storage through the ESX I/O stack, instead of going through the network stack, which improves performance. Please note that disks of a backed up VM must be located on storage accessible by the ESX host on which the VM with Veeam Backup & Replication 5.0 is running.

If you are using vCenter Server earlier than version 4.0, a helper VM named *VeeamBackupVMName(VCB-HELPER)* must also be created on the same ESX server where Veeam Backup VM is running. For example, if your Veeam Backup VM name is *vbsrv01*, the helper appliance name must be *vbsrv01(VCB-HELPER)*. The helper VM is a blank dummy VM without virtual disks or OS installed. This VM is only used to temporarily hot-add disks of backed up VMs to.

Network

This mode is recommended when your ESX host uses local storage, which makes direct storage access not possible. In this mode, VM data is retrieved via the ESX host over the network using NBD (Network Block Device) protocol. You can also choose to transfer disks data over encrypted SSL connection. Use of encryption affects backup performance and CPU usage of the ESX server slightly, but provides secure data transfer.

Note: Veeam Backup & Replication 5.0 processes VM disks one by one. If VM disks are located on different storages (for example, on SAN and local storage subsystem), Veeam Backup & Replication 5.0 will use different transport modes to process VM disks. In such scenario, using the **Failover to network mode if primary backup mode fails** option is strongly recommended.

Legacy Backup Modes

Along with VMware vStorage API backup, Veeam Backup & Replication 5.0 provides support for legacy processing modes — **VCB-enabled backup** and **Network backup**. These modes are left for compatibility with previous versions.

Important! Legacy modes are disabled by default. To enable them, select **Tools > Options…** from the main menu of Veeam Backup & Replication 5.0. On the **Advanced** tab, select the **Enable legacy processing modes** check box.

VCB-Enabled Backup

Veeam Backup & Replication 5.0 integrates with VMware Consolidated Backup — a solution that provides a fast and efficient way of backing up virtual machines.

In case the VCB-enabled backup method is used, Veeam Agents are deployed on the VCB proxy server — a physical server running Windows 2003. As a result, backup processes are moved to the VCB proxy, offloading the ESX server. Due to its proprietary "VCB on-the-fly" technology, Veeam Backup & Replication 5.0 doesn't require extra space on the VCB proxy for VM images.

Note: VCB 1.5 Update 1 offers official support for 32-bit and 64-bit versions of Microsoft Windows Server 2008 proxy server. Previous VCB versions provide only experimental support for Windows Server 2008.

The Veeam Agent started on the VCB proxy works in conjunction with the VMware Consolidated Backup that comprises a set of scripts and utilities responsible for performing main backup activities.

Veeam Backup & Replication 5.0 offers two transport modes for VCB-enabled backup:

- **SAN mode** is used in case VM disks reside on the Fibre Channel SAN or iSCSI SAN. In this mode, Consolidated Backup reads disk data directly from the storage devices. This operation mode is LAN-free: disk data is accessed via the Fibre Channel adapter. The SAN operation mode provides the best performance and the least impact on the production environment.
- **Network mode** is used in case VM disks reside on the local storage devices or NAS. In this mode, Consolidated Backup uses network connection to the ESX server to send unencrypted disks data over the Network Block Device (NBD) protocol. This processing mode is applicable for ESX Server 3.5 or ESX Server 3i version 3.5; VirtualCenter version 2.5, ESX/ESXi version 4 and virtual disks not larger than 1TB each.
  You can select to transfer disks data over networked encrypted SSL connection. Use of encryption puts more stress on the CPU of the ESX server, providing, however, secure data transfer.

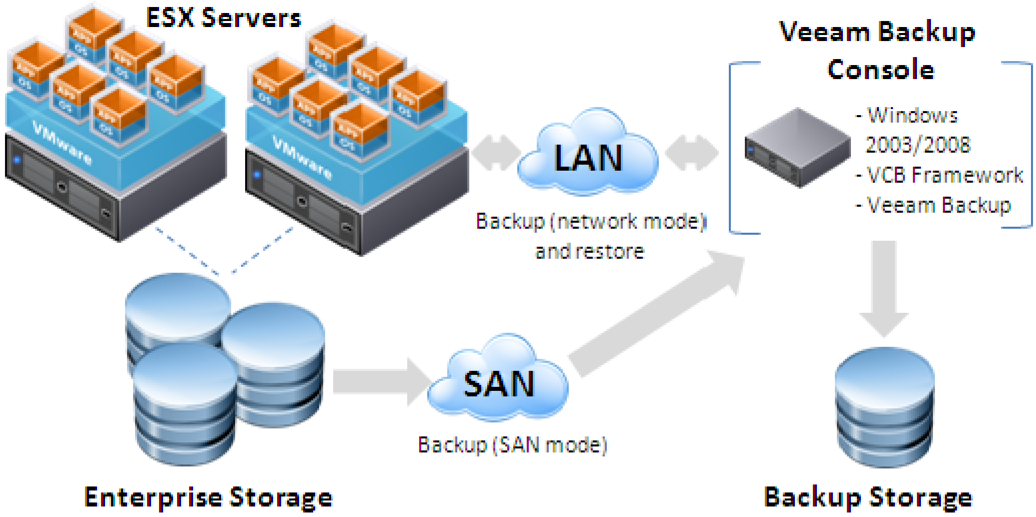

Note: For VCB version 1.0, only SAN operation mode can be used.

To perform VCB-enabled backup, Veeam Backup & Replication 5.0 must be installed locally on a properly functioning VCB proxy that should be directly connected to the SAN fabric. Remote VCB proxy connection is not supported — you will have to install a separate instance of Veeam Backup & Replication 5.0 on every VCB proxy. To learn more about the VCB proxy configuration, see the Installing VCB Proxy section.

Network Backup

During network backup, the data of a VM is retrieved directly from the ESX host through the local area network. Veeam Manager initiated by the Veeam Backup service when the job is launched starts Veeam Agents deployed on the source and target hosts. The ESX server acts as a source, and the localhost or a Linux-based server acts as a target. In case a shared network drive is selected as a target, the Veeam Agent is deployed on the localhost.

By default, Veeam Backup & Replication 5.0 uses a service console agent to achieve the best performance. If the service console is not available (in case of ESXi), the agentless mode is used. In this case, Veeam Backup & Replication 5.0 uses VI API to enable backup to ESX servers and restore from ESX servers. The agentless mode can be used for all types of ESX servers; however, it may provide lower performance if compared with using a service console agent.

You can force service agent mode for data transfer, or select to use agentless mode for all ESX servers using the **Properties** window for an added ESX server (right-click a server and select **Properties**).

The Veeam Agent started on the source host is responsible for performing main job activities: scanning virtual disks data, performing de-duplication if necessary and transferring backup data over the network to the target host. The Veeam Agent started on the target host is responsible for storing the created backup file to its destination.

When a backup job is being performed, the job statistics and program activities are written by Veeam Manager to the *VeeamBackup* database and can be viewed from Veeam Backup & Replication 5.0.

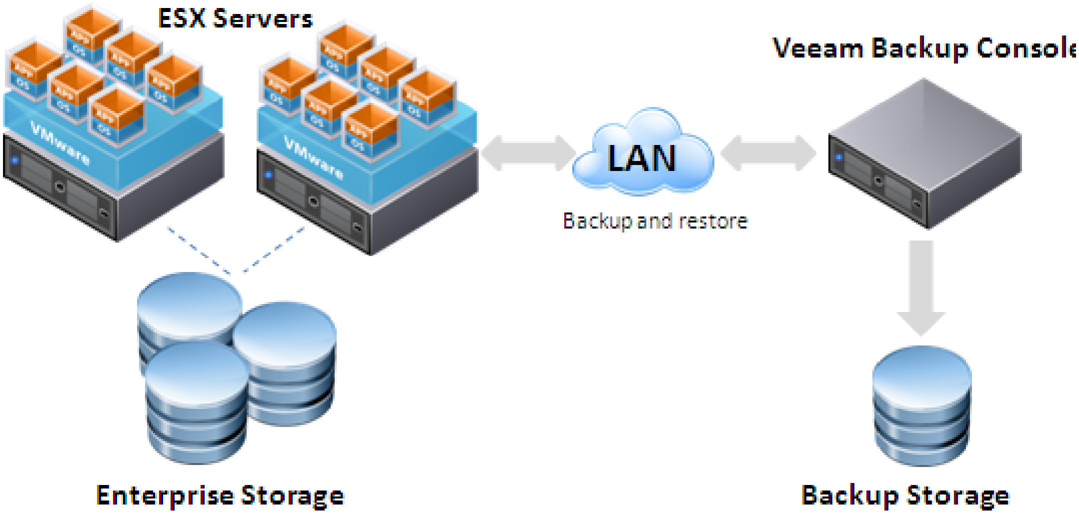

As the main job activities are performed on the source ESX server, the speed of the backup job is reduced and the work efficiency of VMs running on the ESX server is decreased. An alternative to the network backup method is VMware vStorage API and VCB backup modes that allow shifting backup workloads from the source host and accessing storage devices directly through SAN.

Restore Process

Veeam Backup & Replication 5.0 allows you to perform both image-level and file-level restore of backups and replicas — you can restore a virtual machine as a whole to start it on the target ESX server, recover VM files (.vmdk. .vmx and so on) or VM guest OS files and folders and save them on your local machine. At that, VMs or files can be restored at any of the available restore points.

The restore process is always performed via the network.

- **For instant VM recovery,** Veeam Backup & Replication uses vPower engine that mounts a VM image to an ESX host directly from a compressed and deduplicated backup file. To finalize VM recovery, you can vMotion it or replicate to the production datastore and then fail over to this replica during the next maintenance window.
- **For image-level restore**, Veeam Agents are deployed on the target localhost or Linux-based server, and the source ESX server where a restored VM should be started. You can restore a VM to the ESX server of the same or later version than the server on which the backup was created. For instance, if you created a backup of a VM running on ESX 3.0, you can restore this VM to ESX 3.0, 3.0.1, 3.0.2 and 3.5 or ESXi.

  If a VM is restored to an ESXi server or an ESX server for which service console credentials are not provided, Veeam Backup & Replication 5.0 uses the agentless restore mode, and the restore procedure can take much time to perform. For a faster restore process, it is recommended to use ESX server and set credentials for it.
- **For restore of VM guest OS files**, Veeam Agents are deployed on the target localhost/Linux-based server (in case files are restored from a backup) or the ESX server (in case files are restored from a replica), and the source host — a local machine running Windows OS.
The Veeam Agent running on the target host mounts the VM file system without extracting the full virtual machine image to the local drive. Once the restore job is completed, the virtual machine file system is displayed in the Backup Browser. You can copy necessary files and folders to your local machine drive or save them anywhere within the network.
When restoring is performed, ACL is ignored. The restored files get a default ACL set for the folder in which restored files are stored.
- **For VM files restore, Veeam Agents are deployed on the target localhost/Linux-based server and the source** host — the ESX host or a local machine running Windows OS.

Note: For VM guest OS file restoring, Veeam Backup & Replication 5.0 supports FAT and NTFS guest file systems. To restore files from VMs running other guest files systems, use the Veeam File Level Restore wizard. To learn more about the wizard, see the Using Veeam File Level Restore Wizard section.

Replication Process

As well as for backing up virtual machines, for VM replication Veeam Backup & Replication 5.0 offers the following modes: **Direct SAN access**, **Virtual Appliance**, **Network**, as well as two legacy modes: **VCB-enabled replication** and **Network replication** .

Note: For replication, the target ESX server must be of the same or later version than the source ESX server.

VMware vStorage API Replication

Veeam Backup & Replication 5.0 uses VMware vStorage API for Data Protection to access VMware virtual disk storage and copy virtual disk data directly through SAN, thus enabling LAN-free replication of VMs.

The VMware vStorage API mode can be used for VMware vSphere 4 (including ESX/ESXi), vCenter Server 4, VMware Virtual Infrastructure (ESX/ESXi 3.5 and VirtualCenter 2.5). Please note that ESX 3.0 is not supported.

You can select to process your VMs in one of the three modes: **Direct SAN access**, **Virtual Appliance** and **Network** mode. To learn more about modes, see the VMware vStorage API Backup section.

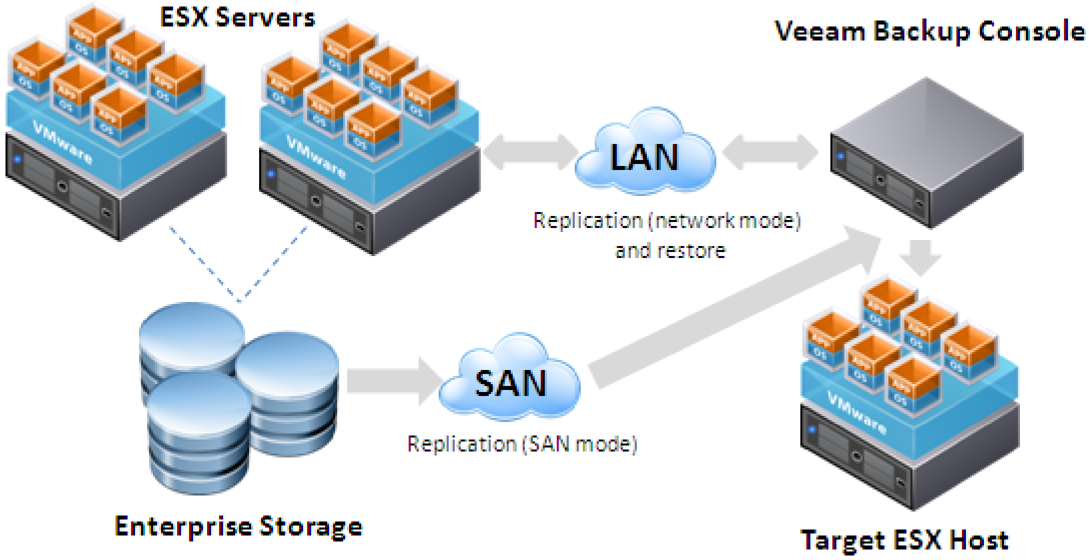

Note: The VMware vStorage API backup mode with Changed Block Tracking enabled is a recommended backup method and is used by default for created replication jobs. To learn more about Changed Block Tracking, see the Native vSphere and vStorage Support section.

Legacy Replication Modes

As well as for backup, for replication Veeam Backup & Replication provides support for legacy processing modes — **VCB-enabled replication** and **Network replication**. These modes are left for compatibility with previous versions. To learn more about legacy modes, see the Legacy Backup Modes section.

Failover

In order to diminish the risk of failure and make the work process seamless, Veeam Backup & Replication 5.0 provides a possibility to failover a virtual machine to its replicated version. In case of software or hardware malfunction, you can recover a corrupted virtual machine by failing over to its replica or its last known good point-in-time incremental.

At replica failover, the Veeam Agent is started on the target ESX server where a replicated virtual machine resides. A snapshot of a replica is created to protect a replicated VM from user's changes, and the replica is started on the target host.

In case the undo failover operation is performed, the replica reverts to the created snapshot. Any changes made to a replicated VM will not be committed to the original VM when undo failover operation is performed.

At performing failover, the original VM should be stopped.

Important! If possible, avoid powering on a replica manually in case its original has failed. Use the **Perform failover** option in the **Restore** wizard instead. Otherwise, the subsequent replication sessions will be failing.
However, for testing your replicas and disaster recovery plan, you can use the following procedure: http://www.virtualizationteam.com/veeam/veeam-backup-replication/testing-veeam-backup-replication-replica-testing-procedure.html.

Transaction-Consistent Backup

Veeam Backup & Replication 5.0 provides two techniques for creating transaction consistent backup images — the **Enable VMware tools quiescence** option and **Enable application-aware image processing** option that utilized Windows VSS. In contrast to restoring a crash-consistent backup, which is essentially equivalent to rebooting a server after a hard reset, restoring transaction consistent backups ensures safety of data of applications running on VMs.

Please note that when you select both options for a job at the same time, the VSS module will only be used for processing backed up and replicated VMs. However, if you use options and select the **Ignore application processing failures** option for backup or replication jobs, all your VMs will be processed with VSS first, and in case of VSS failure (e.g., in case of Linux VMs), VMs will be processed with the VMware tools quiescence option enabled.

This can be very useful when you have both Windows- and Linux-based VMs in one job, so all VMs will be processed in a transaction consistent manner by using VSS or VMware tools quiescence option.

VMware Tools Quiescence

Backing up and replicating a running virtual machine without quiescencing may result in inconsistent backup or replication. To avoid this, the **Enable VMware tools quiescence** option should be used.

The **Enable VMware tools quiescence** option enables freezing of the file-system for proper snapshot creation. With this option enabled, creation of a snapshot is performed with the help of the sync driver responsible for holding incoming I/O and flushing all dirty data to a disk, thus making the file systems consistent.

Note: The **Enable VMware tools quiescence** option is disabled by default. It is strongly recommended to leave it disabled if you are backing up or replicating Windows systems that support Windows VSS — for these systems, it is recommended to use the **Enable application-aware image** processing option.

Application-Aware Image Processing

With the **Enable application-aware image processing** option selected, Veeam Backup & Replication 5.0 utilizes the Windows Volume Shadow Copy Service (VSS) that ensures consistent backup of VSS-aware application running within your virtual machines (domain controllers, databases and other applications) without shutting them down. The **Enable application-aware image processing** option allows creating a transaction-consistent backup image of a VM, which, in contrast to a crash-consistent backup image, ensures successful VM recovery, as well as proper recovery of all applications installed on the VM without any data loss.

In the process of its work, VSS freezes all I/O at a specific point-in-time by interfacing with all VSS-aware applications and the Windows operating system. Consequently, there remain no unfinished database transactions or incomplete application files. Such backups, when restored correctly, result in fully functional applications.

Microsoft Windows VSS integration is supported for Windows XP (64 bit only), Windows 2003, Windows Vista, Windows 2008, Windows 2008 R2, and Windows 7. Microsoft Windows VSS backup option requires that your guest OS has VMware Tools, and all the latest service packs and patches installed.

Please note that administrator credentials are required to access the guest OS. Veeam Backup & Replication 5.0 allows you to provide administrator credentials for each VM in the job separately.

Transaction Logs Truncation

If you are performing backup or replication of database systems that use transaction logs (for example, Microsoft Exchange or Microsoft SQL), you can select to truncate transaction logs after the job so that they don't overflow the storage space. Veeam Backup & Replication provides advanced options of transaction logs handling for different backup scenarios.

- You can choose to truncate transaction logs after any VM backup to save disk on storage.
- You can choose to truncate logs after successful VM backup only. With this option selected, if backup of a VM fails, you will be able to restore the database to any point in time between the last successfully performed backup and a failed backup job. To do so, you will have to restore the database from a successful backup, get transaction logs from the VM that failed to back up, and apply them to a restored database.
- You can choose not to truncate transaction logs at all. This option is recommended if, together with Veeam Backup & Replication, you are using another backup tool to perform guest-level backup, and this tool maintains consistency of the database state. In this case, truncation of logs with Veeam Backup & Replication will break the guest-level backup chain and cause it to fall out of sync.

ESXi Support

Many organizations deploying virtual infrastructure incline to use ESXi servers in their production environment. However, use of ESXi may cause a problem when taking into consideration the disaster recovery strategy. Owing to the absence of the service console inherent to the standard ESX server, common methods of backup and replication are not applicable to ESXi.

Veeam Backup & Replication 5.0 offers full support for ESXi. It uses VMware APIs to access ESXi remotely and enable backup and restore of VMs running on ESXi servers over the network. You may also select to perform backup of VMs running on ESXi with the VCB proxy and using VMware vStorage API. However, at this time you cannot select ESXi as a backup target — as it is possible to do with "full" ESX servers.

Important! Please note that since June 2009 Veeam has discontinued support for ESXi Free in Veeam Backup & Replication in order to comply with VMware`s updated licensing policy.

Native vSphere and vStorage Support

Veeam Backup & Replication 5.0 offers full and native support for VMware vSphere, including all new vSphere and vStorage functionality:

- **Thin-provisioned disks**. With VMware vSphere thin provisioning, you do not lose space that is provisioned to a VM, but is not actually used — the space is allocated and committed by a VM on demand, and becomes available to the rest of your system if free. Veeam Backup & Replication 5.0 allows you to back up and replicate VMs using thin-provisioned disks. When restoring a backed up VM, you can select to restore it in its initial state, or force all VM disks thin or thick at your option.
- **ESX4 changed block tracking**. Veeam Backup & Replication 5.0 leverages a new VMware block copy mechanism, Changed Block Tracking, which minimizes CPU and memory resource consumption on the ESX host up to several times.

  Changed block tracking (or CBT) is vSphere functionality that keeps track of blocks of a virtual disk that have changed since the last backup or replication cycle. With CBT, Veeam Backup & Replication 5.0 does not have to scan VM virtual disk for changes – they can simply query this information using an API call to the VMkernel. This enables much faster incremental backup and replication cycles and greatly reduces resource usage. For example, if a VM only had 5 percent change since the last backup, the incremental backup time will be 20 times faster.

CBT requires ESX version 4 and virtual machines with hardware version 7. As mentioned before, you can use the vStorage APIs for Data Protection mode for ESX(i) 3.5 servers as well. However, they will not be processed with Changed Block Tracking.

- **New vStorage APIs**. Veeam Backup & Replication 5.0 provides native support for the new vStorage APIs. The vStorage APIs provides a possibility to perform off-host backup and replication and thus reduce workload on LAN.
- **Support for virtual applications (vApp)**. Veeam Backup & Replication 5.0 supports virtual application that are now presented as a part of VI hierarchy and treated as VM containers that can be easily backed up or replicated.

Integration with Traditional Backup

One of the major questions concerned by organizations at choosing a backup solution is a question of its deployment within the frames of the existing data protection strategy. Implementing a new solution and introducing changes in the established scheme may seem to be a risky point.

Veeam Backup & Replication 5.0 provides a possibility of integration with the functioning backup scheme, offering a flexible approach to protect your VI environment data and ensuring its flawless operation. This may be very useful, for example, for organizations using a traditional backup technology — tape. With the option of **performing post-backup activities**, you may choose to execute necessary actions once the backup procedure is completed. One of the most common scenarios in such case is to run a custom script that will write a ready backup file to the tape as soon as the backup process is finished. At that, a desired post-backup activity may be performed once after a set of backup job runs, which will protect against redundant loading of the VI environment.

PLANNING AND PREPARATION

This chapter describes the planning and preparation steps that you should take before the Veeam Backup & Replication 5.0 deployment.

Prerequisites

- Veeam Backup & Replication 5.0 requires .NET Framework 2.0 SP1. If it is not available, the Veeam Backup & Replication setup will install it on your computer.
- Veeam Backup & Replication 5.0 uses SQL Server instance installed either locally or remotely. In case it is not installed, the Veeam Backup & Replication setup will install SQL Server 2005 Express SP3 on your computer. If an SQL Server instance has already been installed by the previous version, Veeam Backup & Replication 5.0 will connect to the existing database, upgrade it (if necessary) and use it for work.
- In case you are planning to perform VCB-enabled backup or replication, you must install Veeam Backup & Replication 5.0 locally on a properly configured VCB proxy server. Depending on the versions of ESX servers you have, you may need to install a specific VCB version. To learn more, refer to the compatibility matrix.

Requirements

The present chapter describes the list of system requirements to the VMware Infrastructure, Veeam Backup & Replication console, virtual machines and backup target hosts, Veeam Backup Enterprise Manager, search server, necessary rights and permissions, as well provides information on ports used by Veeam Backup & Replication 5.0.

System Requirements

To ensure successful usage of Veeam Backup & Replication 5.0, the following system requirements should be met:

Virtual Infrastructure	
Platforms	VMware vSphere 4. VMware Infrastructure 3 (VI3).
Hosts	ESX(i) 4.x ESX(i) 3.x Free ESXi is not supported
VMware Infrastructure	vCenter Server 4.x (optional) Virtual Center 2.x (optional)
Virtual Machines	
Hardware	All types and versions of virtual hardware are supported, except physical RDM (raw device mapping) and Independent disks. You can use disk exclusion functionality to exclude some of the unsupported disks from backup. MBR disk partition table is required for file-level restore, GPT disks are not supported.
OS	Any operating system supported by VMware. Application-aware image-level processing option is supported on Windows XP x86, Windows 2003, Windows

	Vista, Windows 2008, Windows 2008 R2 and Windows 7. Windows file-level restore option is supported on NTFS, FAT and FAT32 file systems. GPT disks are not supported. To restore files from non-Windows guests (Linux, Solaris, BSD) use the Multi-OS File Level Restore wizard.
Software	VMware Tools (optional, recommended). Application-aware image-level processing option requires that your guest has VMware Tools and all latest service packs and patches.
Veeam Backup & Replication Console	
Hardware	*CPU*: modern x86/x64 processor (minimum 4 cores recommended for optimal backup performance). Using faster processors generally improves backup performance. *Memory*: 1024MB RAM (2048MB RAM when using local SQL Express installation). Using faster memory (DDR3) generally improves backup performance. *Hard disk space*: 100 MB. *Network*: 1Gbit/sec recommended due to backup performance considerations.
OS	Both 32-bit and 64-bit versions of the following operating systems are supported: • Microsoft Windows XP SP3 • Microsoft Windows 2003 SP2 • Microsoft Windows Vista SP2 • Microsoft Windows 2008 SP2 • Microsoft Windows 2008 R2 SP1 • Microsoft Windows 7 SP1
Software	Microsoft .NET Framework 2.0 SP1 (included in the setup) Microsoft PowerShell 2.0 or later
Backup Target	
Hardware	*CPU*: modern x86/x64 processor. Using faster processors generally improves backup performance when using Linux targets and **Best** compression option. *Memory*: 256 MB RAM. *Hard disk*: Using faster storage (fast high-RPM hard drives, RAID0 configurations) and optimal storage controller settings generally improves backup performance. *Hard disk space*: Sufficient disk space required to store backup files. *Network*: 1Gb/sec recommended due to backup performance considerations.
OS	Microsoft Windows. All major Linux distributions. ESX 3.x or later (ESXi is not supported).
Replication Target	
Hosts	ESX(i) 3.x or later.
SQL Database	
Database	Microsoft SQL Server 2005 Express, Microsoft SQL Server 2005 or Microsoft SQL Server 2008. If you do not have

	one, the Veeam Backup & Replication setup will install Microsoft SQL Server 2005 Express SP3.
Veeam Backup Enterprise Manager	
Hardware	*CPU*: x86/x64 processor *Memory*: 1024MB RAM (2048MB RAM when using local SQL Express installation). *Hard disk space*: 25MB. *Network*: 1Gbit/sec recommended due to backup performance considerations.
OS	Both 32-bit and 64-bit versions of the following operating systems are supported: • Microsoft Windows XP SP3 • Microsoft Windows 2003 SP2 • Microsoft Windows Vista SP2 • Microsoft Windows 2008 SP2 • Microsoft Windows 2008 R2 SP1 • Microsoft Windows 7 SP1
SQL	Microsoft SQL Server 2005 Express, Microsoft SQL Server 2005 or Microsoft SQL Server 2008. If you do not have one, the Veeam Backup Enterprise Manager setup will install Microsoft SQL Server 2005 Express SP3.
Software	Microsoft .NET Framework 2.0 SP1 or later. Microsoft Internet Information Services 5.1 or later (IIS 6 Management Compatibility and Windows Authentication components for IIS 7.0). If not installed, the MS Windows installation disk to set up IIS. *Browser*: Internet Explorer 6.0 or later, Mozilla Firefox 3.0 or later. Microsoft Excel 2003 or later (to view report data exported from Veeam Backup Enterprise Manager).
Veeam Backup Search Server	
Hardware	Refer to corresponding Microsoft Search Server version system requirements.
OS	Both 32-bit and 64-bit versions of the following operating systems: • Microsoft Windows Server 2003. • Microsoft Windows Server 2008. • Microsoft Windows Server 2008 R2. All the latest service packs and security updates should be installed.
Software	Microsoft Search Server 2008 (including Express edition) Microsoft Search Server 2010 (including Express edition)

Required Permissions

The accounts used for installing and using Veeam Backup & Replication 5.0 should have the following permissions:

Account	Required Permission
Setup Account	Local Administrator permissions on the Veeam Backup & Replication console to install Veeam Backup & Replication 5.0.
Target/Source **Host Permissions**	Root permissions on the source ESX/ESXi server. Root (or equivalent) permissions on the target Linux host. Write permission on the target folder and share. If vCenter is used, administrator credentials are required.
SQL Server	The user account must have database owner rights for the *VeeamBackup* database on the SQL Server instance.
Veeam Backup Enterprise Manager	Local Administrator permissions on the Veeam Backup Enterprise Manager server to install Veeam Backup Enterprise Manager. To be able to work with Veeam Backup Enterprise Manager, users should be members of the Portal Administrators or Portal Viewers group.
Veeam Backup Search Server	Local Administrator permissions on the Veeam Backup Search Server console to install Microsoft Search Server and the Veeam Backup Search component

Hardware Recommendations

- At least 1Gbit/s network is required. We do not recommend running Veeam Backup & Replication 5.0 on slower connections due to performance considerations.
- Using faster processors configurations on the Veeam Backup & Replication console generally improves the backup performance. We recommend installing Veeam Backup & Replication 5.0 on powerful computers with multi-core processors (Intel Core Duo/Quad, AMD Phenom X2/X4).
- You can additionally improve the backup speed by ensuring that a backup file is saved to the fast storage (high-RPM hard drives, RAID0 configurations).
- The amount of RAM installed on the Veeam Backup & Replication console does not affect the backup performance significantly.

Used Ports

From	To	Protocol	Port	Notes
Veeam Backup server	vCenter Server	HTTPS	443	Default VMware web service port that can be customized in vCenter settings.
	ESX(i) Server	HTTPS	443	Default VMware web service port that can be customized in ESX host settings. Not required if vCenter connection is used.
		TCP	902	VMware data mover port.
		TCP	22	Default SSH port used as a control channel, only for jobs with full ESX target with service console agent enabled.
		TCP	2500-5000	Used as transmission channels, only for jobs with full ESX target with service console agent enabled. For one job, one port from this range is used. You can open only a small range of ports for the concurrent jobs, depending on your environment. For example, you need to open 2500–2510 to be able to perform 10 concurrent jobs.
	Linux Server	TCP	22	Port used as a control channel from the console to the target ESX/Linux host.
		TCP	2500-5000	Used as transmission channels for jobs with Linux target. For one job, one port from this range is used. You can open only a small range of ports for the concurrent jobs, depending on your environment. For example, you need to open 2500–2510 to be able to perform 10 concurrent jobs.
	Veeam Backup Enterprise Manager	TCP	9392	Default port used for Enterprise Manager interaction; can be changed during Veeam Backup & Replication installation.
		TCP	2500 9393	Default port used by the Veeam Backup Catalog interaction; can be changed during Veeam Backup & Replication installation.
Veeam vPower NFS service	ESX host	UDP	111 1058 2049	Standard NFS ports.
Veeam Backup Enterprise Manager	Veeam Backup server	TCP	9394	Default port used for interaction with backup servers; can be changed during Enterprise Manager installation.

From	To	Protocol	Port	Notes
	Microsoft Search Server	TCP	9395	Default port used by Veeam Backup Search service integration component; can be changed during Veeam Backup Search installation
IIS extension	Veeam Backup Enterprise Manager	HTTP HTTPS	9080 9443	Default ports used by Enterprise Manager service to communicate with the web site; can be changed during Enterprise Manager installation.
	Enterprise Calalog Service	HTTP	9393	Port used to enable advanced file search

Job Planning

To ensure sufficient use of resources and faster jobs performance, it is important to plan your backup, replication and copying jobs in a proper way. The present chapter contains a number of recommendations that may be helpful in organizing and scheduling jobs.

Organizing Jobs Content

- Consolidate VMs created on the basis of one or similar template in the same job for the best de-duplication space savings. Balance this amount with the manageability of your backup job (the more VMs you include in the job, the longer the job will take in case you need to restart it).
- Veeam Backup & Replication 5.0 does not support backing up VMs using RDM in physical mode. You can use disk exclusion functionality to exclude some of the unsupported disks from backup/replication. VMs using virtual RDM may be backed up by means of VMware vStorage API and VCB-enabled modes.
- Veeam Backup & Replication 5.0 supports VMs using independent disks only for powered off VMs. Backup of such VMs can be performed only in the network legacy mode via the service console of ESX host.

Configuring Jobs Settings

- Use of compression decreases the size of created backup files, but may affect the duration of the backup procedure. You may try backing up VMs using different compression levels to define the best ratio between the size of the result backup file and the time of the backup procedure.
- The Direct SAN access option is the best when your ESX servers are using SAN storage device because it provides the best performance and the least impact on your production environment.
- If you are using a local storage on your ESX servers, use the Network option of the VMware vStorage APIs backup mode.

Running Multiple Jobs at Once

- Avoid having parallel backup jobs sharing the same sources and targets.
- If your backup window is not enough to back up all your VMs, install multiple Veeam Backup & Replication consoles and spread jobs across them.
- For VCB-enabled Backup. The maximum number of parallel VCB-enabled jobs (backup or replication) is limited to 8 due to VCB limitations.

- If you need to perform backup of VMs residing on one datastore, instead of creating several backup jobs targeted at this datastore, you can create a single backup job and add the datastore as a VM container to it.

DEPLOYMENT

The Veeam Backup & Replication 5.0 setup comprises the following components:

- **Veeam Backup & Replication 5.0** itself.
- **Veeam Backup Enterprise Manager**, allowing you to manage multiple Veeam Backup & Replication installations from a single web console.
- **Veeam Backup Search**, enabling advanced search functionality in Veeam Backup Enterprise Manager.
- **U-AIR wizards** allowing you to restore individual items and objects from applications installed in VMs.

All components can be installed on the same machine, either physical or virtual, or can be set up separately.

Before you begin the installation process, take the following steps to prepare for deployment:

- Check system requirements. Make sure the computers on which Veeam Backup & Replication 5.0, Veeam Backup Enterprise Manager and Veeam Backup Search are to be installed meet the system requirements (see the System Requirements section).
- Check account permissions. Make sure all accounts you will be using have sufficient permissions defined in the Required Permissions section. You will not be able to use Veeam Backup & Replication 5.0, Veeam Backup Enterprise Manager and Veeam Backup Search successfully if the accounts do not have required permissions.
- Verify that VCB is working (if VCB-enabled backup should be performed).

Installing Veeam Backup & Replication

This section will guide you through the Veeam Backup & Replication 5.0 installation process.

Note: If you are planning to perform VCB–enabled backup, install Veeam Backup & Replication 5.0 on the VCB server. To learn more about VCB configuring, see the Configuring VCB Proxy (VMware Reference) section.

Step 1. Download and Run Veeam Backup & Replication Setup

Download the latest version of Veeam Backup & Replication 5.0 from: http://www.veeam.com/downloads/. Unpack the downloaded archive and run the setup file (*Veeam_Backup_Setup_x64.exe* or *Veeam_Backup_Setup_x86.exe*).

Step 2. Accept the License Agreement

Read, then accept or decline the License Agreement. If you select **I do not accept the terms in the license agreement**, the installation process will be terminated.

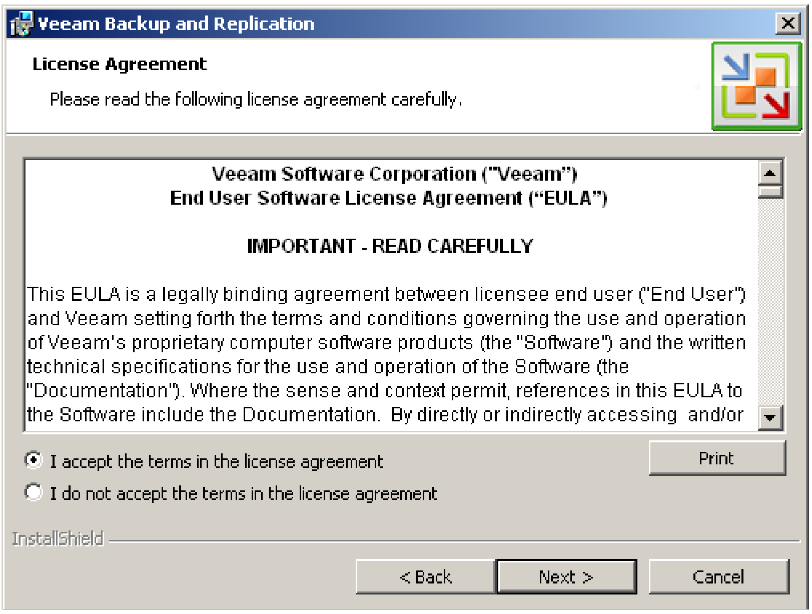

Veeam Backup and Replication

License Agreement

Please read the following license agreement carefully.

Veeam Software Corporation ("Veeam")
End User Software License Agreement ("EULA")

IMPORTANT - READ CAREFULLY

This EULA is a legally binding agreement between licensee end user ("End User") and Veeam setting forth the terms and conditions governing the use and operation of Veeam's proprietary computer software products (the "Software") and the written technical specifications for the use and operation of the Software (the "Documentation"). Where the sense and context permit, references in this EULA to the Software include the Documentation. By directly or indirectly accessing and/or

● I accept the terms in the license agreement

○ I do not accept the terms in the license agreement

Step 3. Install a License

At this step, you should install a license that was sent to you after registration. Click the **Browse…** button and select a necessary .lic file.

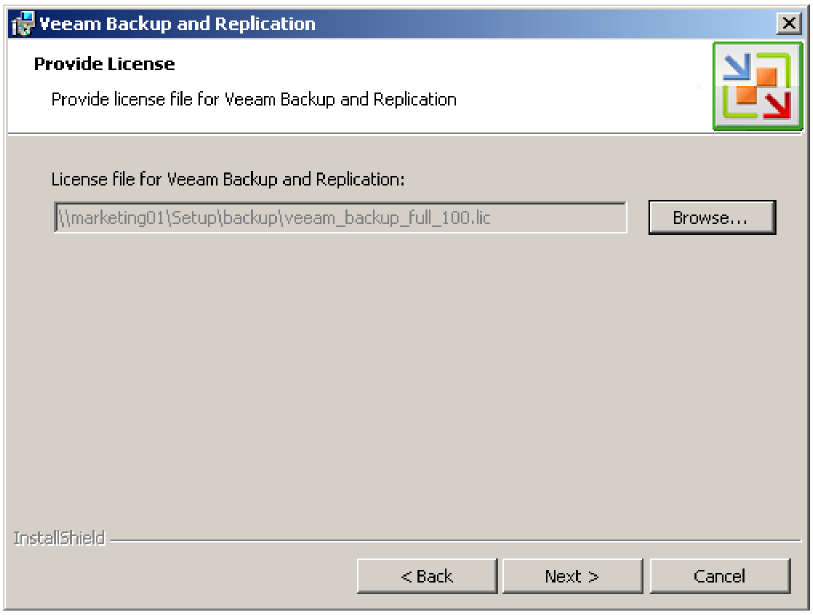

Veeam Backup and Replication

Provide License

Provide license file for Veeam Backup and Replication

License file for Veeam Backup and Replication: \\marketing01\Setup\backup\veeam_backup_full_100.lic

Step 4. Choose Destination for Installation

During installation, the setup installs Veeam Backup & Replication itself, Veeam Backup Catalog component responsible for indexing VM guest OS files, and Veeam Backup PowerShell snap-in for automating backup and replication activities via scripts. Note that the Veeam Backup PowerShell component is disabled by default.

Specify the installation folder for each component. Note that at least 150 MB is required to install Veeam Backup & Replication 5.0, at least 55 Mb to install Veeam Backup Catalog, and at least 400 Kb to install Veeam Backup PowerShell snap-in.

Use the **Space** button to estimate how much free space is available on your local drives.

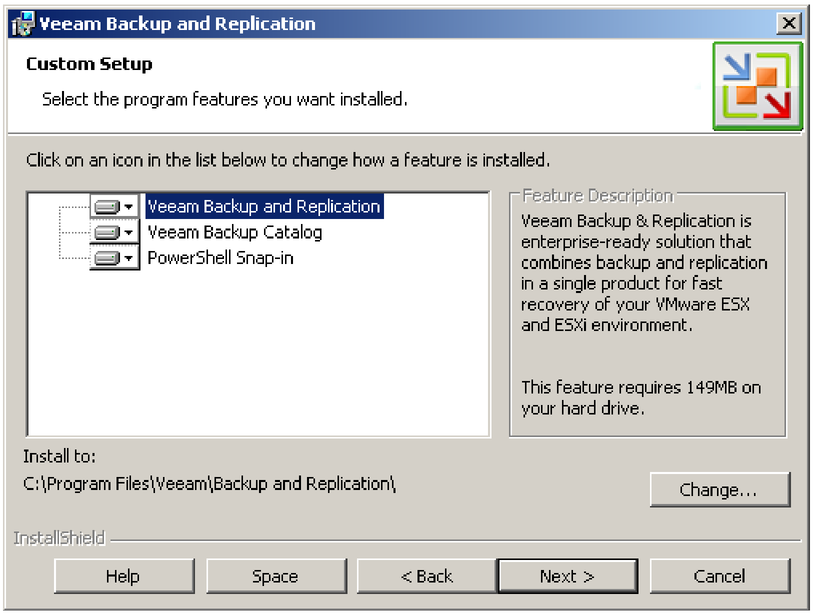

Veeam Backup and Replication

Custom Setup

Select the program features you want installed.

Click on an icon in the list below to change how a feature is installed.

- Veeam Backup and Replication
- Veeam Backup Catalog
- PowerShell Snap-in

Feature Description

Veeam Backup & Replication is enterprise-ready solution that combines backup and replication in a single product for fast recovery of your VMware ESX and ESXi environment.

This feature requires 149MB on your hard drive.

Install to:
C:\Program Files\Veeam\Backup and Replication\

Step 5. Choose or Install SQL Server

At this step, you should select an SQL Server instance on which the *VeeamBackup* database should be created or choose to install a new SQL Server instance.

If the SQL Server is already installed, select the **Use existing instance of SQL Server** option and enter the instance name in the *HOSTNAME\INSTANCE* format and specify the name of the database to be used in the **Database** field.

If the SQL Server is not installed, select the **Install new instance of SQL Server** option.

The user account under which the installation is being performed should have sufficient rights to log on to the selected SQL Server instance using Windows integrated authentication and create a database on the selected instance.

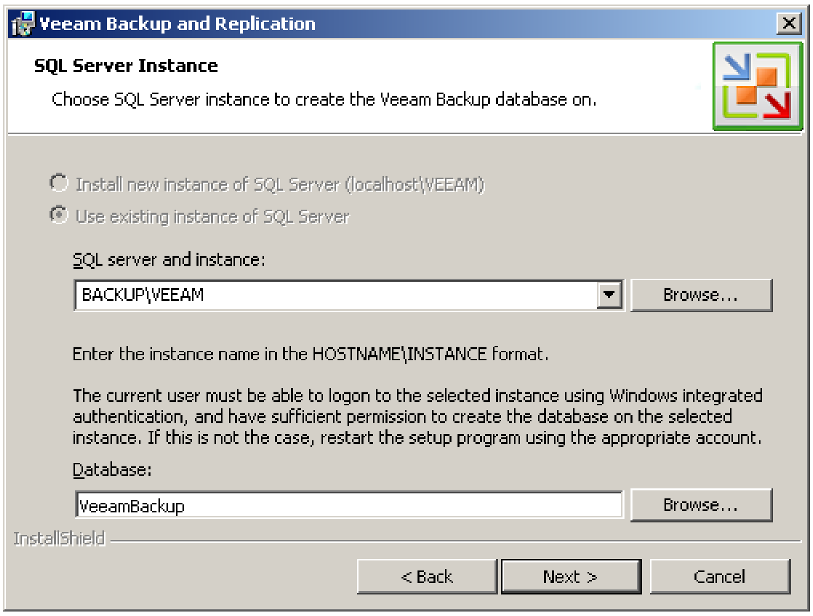

Veeam Backup and Replication

SQL Server Instance

Choose SQL Server instance to create the Veeam Backup database on.

○ Install new instance of SQL Server (localhost\VEEAM)

● Use existing instance of SQL Server

SQL server and instance: BACKUP\VEEAM

Enter the instance name in the HOSTNAME\INSTANCE format.

The current user must be able to logon to the selected instance using Windows integrated authentication, and have sufficient permission to create the database on the selected instance. If this is not the case, restart the setup program using the appropriate account.

Database: VeeamBackup

Note: In case the *VeeamBackup* database already exists on the SQL Server instance (that is, it was created by the previous installations of Veeam Backup & Replication), a warning message notifying about it will be displayed. Click the **Use Existing** button to connect to the detected database. If necessary, the existing database will be upgraded to the latest version.

Step 6. Specify Service Credentials

Enter the administrative credentials of the account under which you want to run the Veeam Backup Service. The user name should be specified in the *DOMAIN\USERNAME* format.

The user account must have database owner rights for the *VeeamBackup* database on the SQL Server instance and full control NTFS permissions on the *VBRCatalog* folder where index files are stored. The *Log on as service* right will be automatically granted to the specified user account.

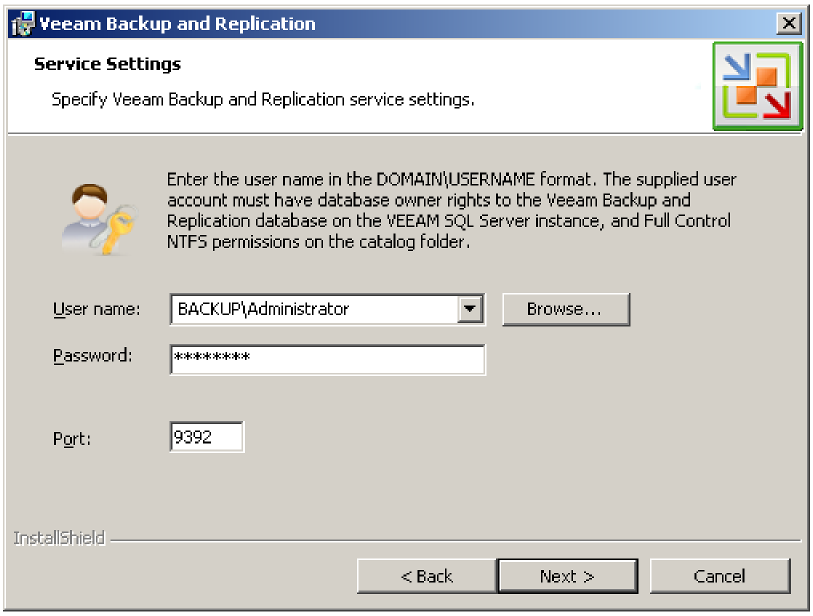

Veeam Backup and Replication

Service Settings

Specify Veeam Backup and Replication service settings.

Enter the user name in the DOMAIN\USERNAME format. The supplied user account must have database owner rights to the Veeam Backup and Replication database on the VEEAM SQL Server instance, and Full Control NTFS permissions on the catalog folder.

User name: BACKUP\Administrator

Password: ********

Port: 9392

If necessary, change the number of TCP port. By default, Veeam Backup & Replication services use port 9392.

Step 7. Specify Catalog Options

Specify the name and destination for catalog folder where index files should be stored. By default, catalog is located at: *C:\VBRCatalog*. If necessary, change the number of port to be used by Veeam Backup Catalog components. By default, port 9393 is used.

In the **vPower NFS** section, specify the folder where instant VM recovery write cache will be stored. Please note that the selected volume should have at least 100 Gb of free disk space.

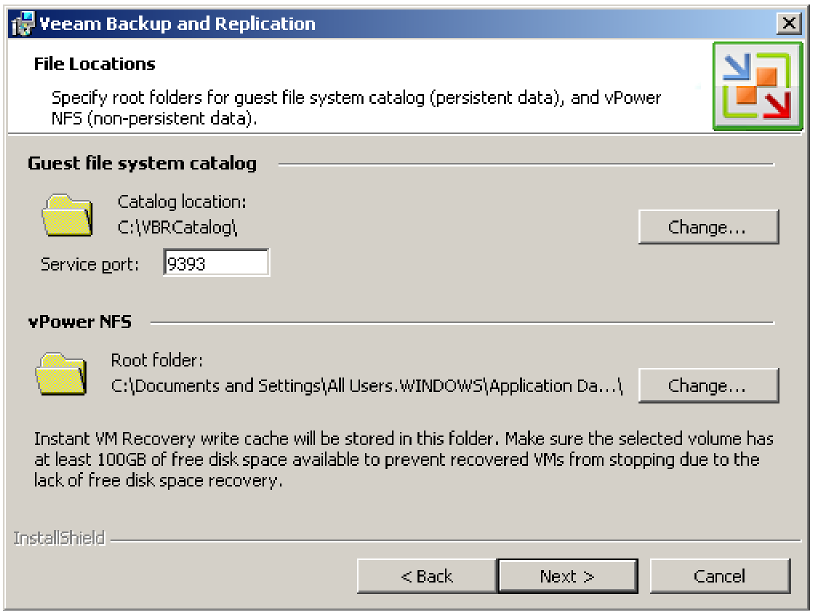
Veeam Backup and Replication

File Locations

Specify root folders for guest file system catalog (persistent data), and vPower NFS (non-persistent data).

Guest file system catalog

Catalog location:
C:\VBRCatalog\

Service port: 9393

vPower NFS

Root folder:
C:\Documents and Settings\All Users.WINDOWS\Application Da...\

Instant VM Recovery write cache will be stored in this folder. Make sure the selected volume has at least 100GB of free disk space available to prevent recovered VMs from stopping due to the lack of free disk space recovery.

Step 8. Install Veeam Backup & Replication 5.0

Click **Next**, then click **Install** to start the installation. Once the installation is complete, launch Veeam Backup & Replication 5.0 by clicking the Veeam Backup & Replication icon on your desktop.

Installing Veeam Backup Enterprise Manager

This section will guide you through the installation process of Veeam Backup Enterprise Manager.

Important! Before installing Veeam Backup Enterprise Manager, make sure you have IIS installed on your computer. When installing IIS version 5.1 to 6.0, make sure that the IIS component is selected. When installing IIS starting from version 7, the following components should be selected: **Web Management Tools**; ASP and ASP.NET under **Application Development Features**; Default Document, Directory Browsing, HTTP Errors and Static Content under **Common HTTP Features**, Static Content Compression under **Performance Features**; and **Security** components.

Step 1. Run Enterprise Manager Setup

After you have downloaded the latest version of Veeam Backup & Replication, run the setup file (*Veeam_Backup_Enterprise_Manager_Setup_x64.exe* or *Veeam_Backup_Enterprise_Manager_Setup_x86.exe*) from the downloaded archive.

Step 2. Accept License Agreement

Read, then accept or decline the License Agreement. If you select **I do not accept the terms in the license agreement**, the installation process will be terminated.

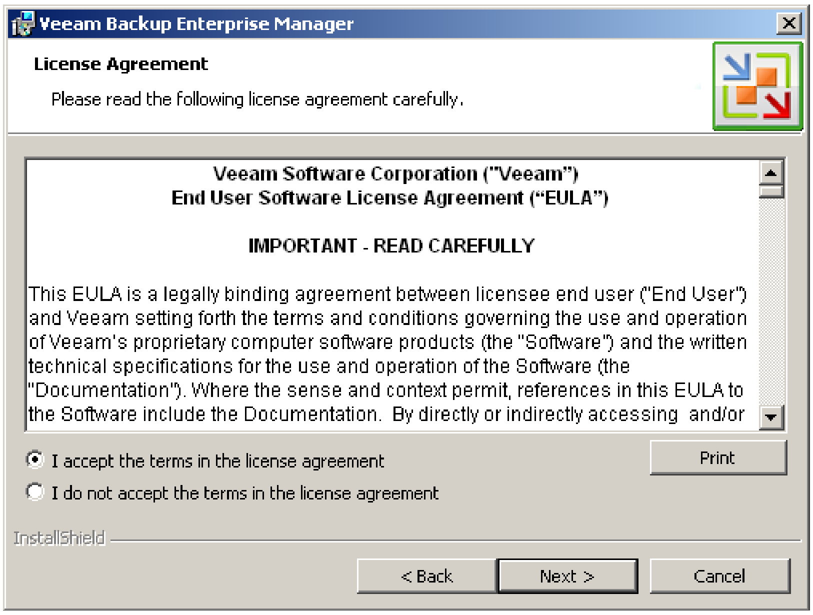

Veeam Backup Enterprise Manager

License Agreement

Please read the following license agreement carefully.

Veeam Software Corporation ("Veeam")
End User Software License Agreement ("EULA")

IMPORTANT - READ CAREFULLY

This EULA is a legally binding agreement between licensee end user ("End User") and Veeam setting forth the terms and conditions governing the use and operation of Veeam's proprietary computer software products (the "Software") and the written technical specifications for the use and operation of the Software (the "Documentation"). Where the sense and context permit, references in this EULA to the Software include the Documentation. By directly or indirectly accessing and/or

● I accept the terms in the license agreement

○ I do not accept the terms in the license agreement

Step 3. Install a License

At this step, you should install an enterprise license that was sent to you after registration. Click the **Browse…** button and select a necessary .lic file.

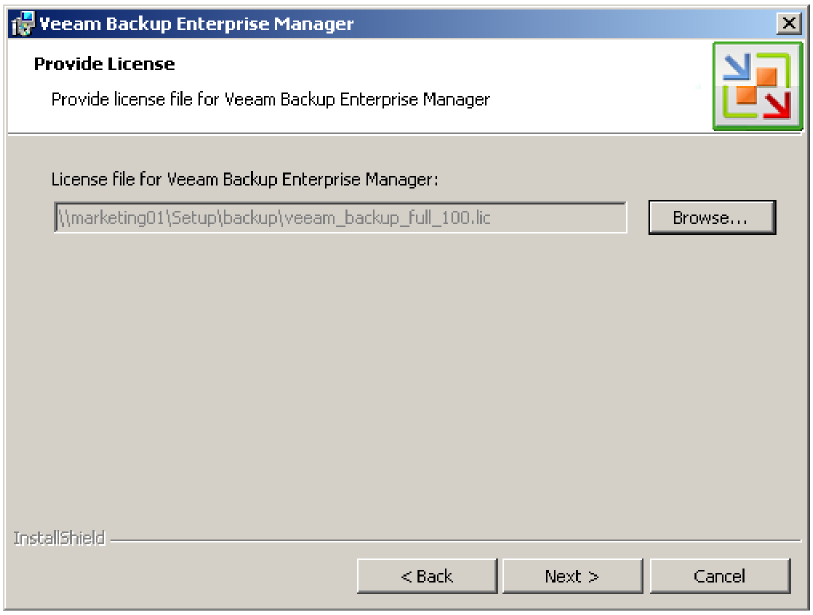

Veeam Backup Enterprise Manager

Provide License

Provide license file for Veeam Backup Enterprise Manager

License file for Veeam Backup Enterprise Manager: \\marketing01\Setup\backup\veeam_backup_full_100.lic

Step 4. Confirm Component Installation and Choose Destination

At this step, you should select the destination folder for installation and confirm that the Veeam Backup Enterprise Manager components should be installed on your computer.

The Enterprise Manager setup installs two components: *Veeam Backup Enterprise Manager Web Site*, which requires at least 20 Mb, and *Veeam Backup Enterprise Manager Server*, which requires at least 26 Mb on the local hard drive, and Veeam Backup Catalog used for indexing and search activities. Use the **Space** button to estimate how much free space is available on your disks.

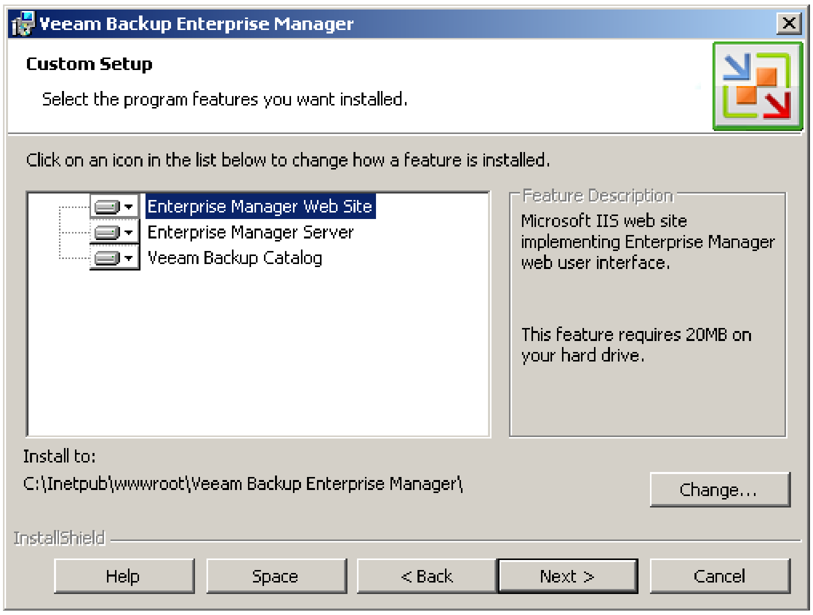

Veeam Backup Enterprise Manager

Custom Setup

Select the program features you want installed.

Click on an icon in the list below to change how a feature is installed.

- Enterprise Manager Web Site
- Enterprise Manager Server
- Veeam Backup Catalog

Feature Description

Microsoft IIS web site implementing Enterprise Manager web user interface.

This feature requires 20MB on your hard drive.

Install to:

C:\Inetpub\wwwroot\Veeam Backup Enterprise Manager\

Note: If you are installing Veeam Backup Enterprise Manager on the Veeam Backup & Replication machine where Veeam Backup Catalog is already installed, Veeam Backup Catalog will be excluded from this list of components.

Step 5. Set up a Database to Be Used

Select to install a new SQL server or use the existing one. If the second option is selected, enter the SQL Server instance name in the *HOSTNAME\INSTANCE* format.

Please make sure the account under which Veeam Backup Enterprise Manager is installed is granted administrative privileges on both the specified SQL Server instance and on the local machine, as Windows authentication for SQL Server is required.

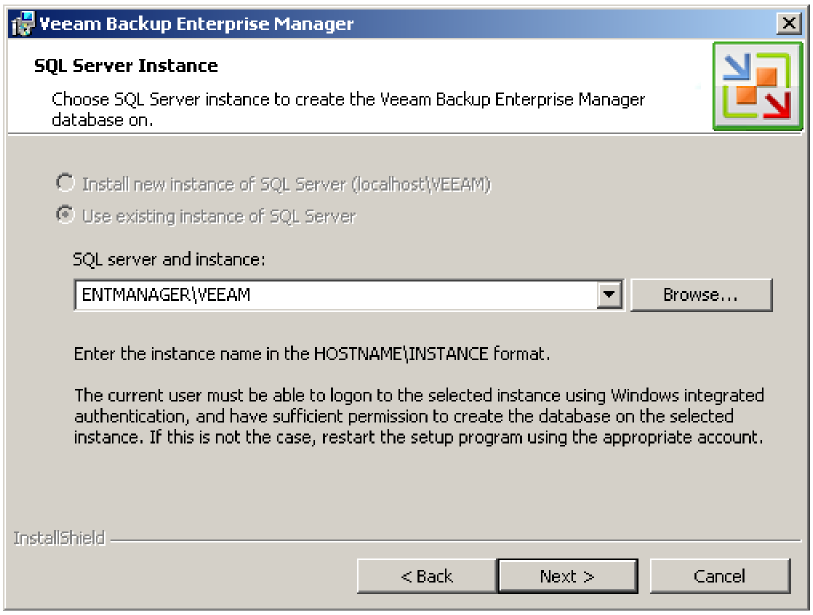

Veeam Backup Enterprise Manager

SQL Server Instance

Choose SQL Server instance to create the Veeam Backup Enterprise Manager database on.

○ Install new instance of SQL Server (localhost\VEEAM)

● Use existing instance of SQL Server

SQL server and instance: ENTMANAGER\VEEAM

Enter the instance name in the HOSTNAME\INSTANCE format.

The current user must be able to logon to the selected instance using Windows integrated authentication, and have sufficient permission to create the database on the selected instance. If this is not the case, restart the setup program using the appropriate account.

Tip: You can use the same SQL server for both Veeam Backup & Replication 5.0 and Veeam Backup Enterprise Manager, or different SQL servers.

Step 6. Specify Service Credentials

Specify the user name and password to be used by the Veeam Backup Enterprise Manager Service. Please note that the user should have the database owner rights to the Enterprise

Manager database on the Veeam SQL Server instance and full control NTFS permissions on the *VBRCatalog* folder where index files are stored.

If necessary, change the number of TCP port. By default, port 9394 is used.

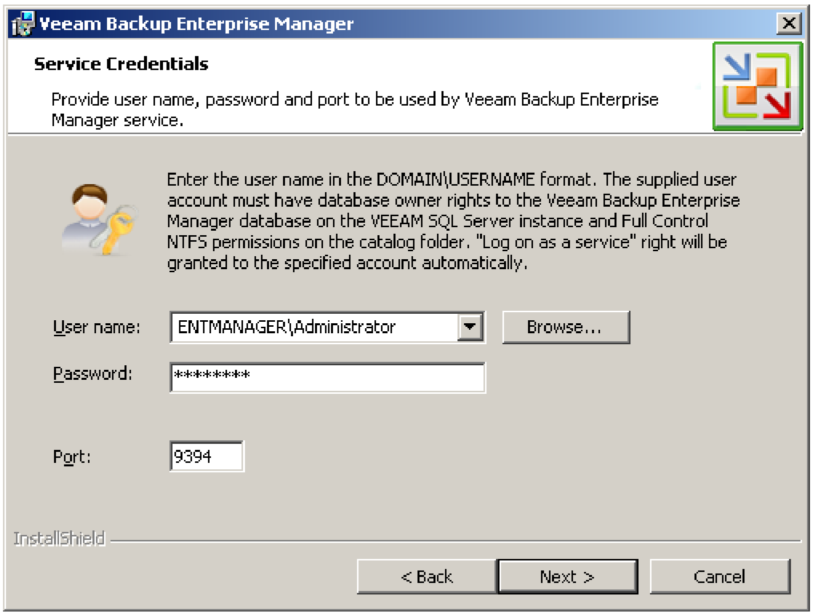
Veeam Backup Enterprise Manager

Service Credentials

Provide user name, password and port to be used by Veeam Backup Enterprise Manager service.

Enter the user name in the DOMAIN\USERNAME format. The supplied user account must have database owner rights to the Veeam Backup Enterprise Manager database on the VEEAM SQL Server instance and Full Control NTFS permissions on the catalog folder. "Log on as a service" right will be granted to the specified account automatically.

User name: ENTMANAGER\Administrator

Password: *********

Port: 9394

Step 7. Specify Catalog Options

Specify the name and destination for catalog folder where content index files should be stored. By default, catalog is located at: *C:\VBRCatalog*. If necessary, change the number of port to be used by Veeam Backup Catalog components. By default, port 9393 is used.

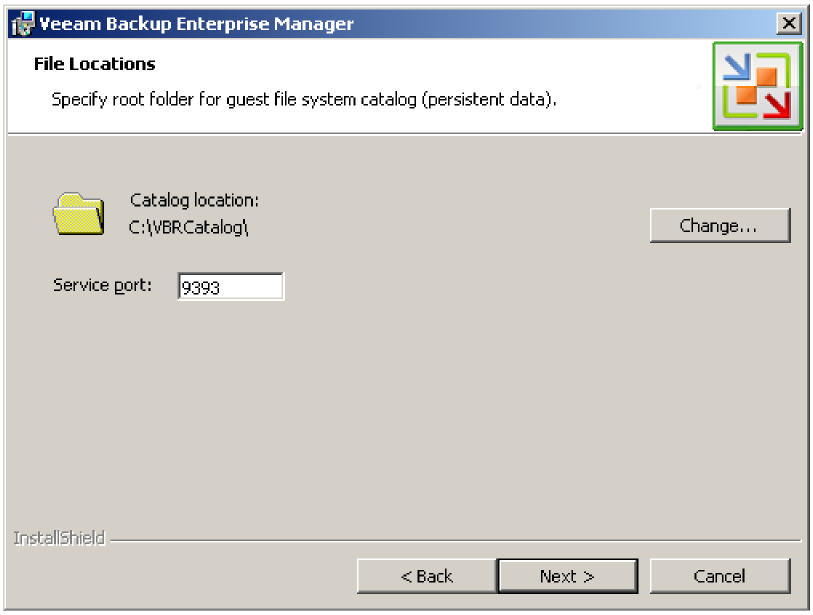
Veeam Backup Enterprise Manager

File Locations

Specify root folder for guest file system catalog (persistent data).

Catalog location:
C:\VBRCatalog\

Service port: 9393

Step 8. Specify TCP Ports

Review and, if necessary, change HTTP and HTTPS ports and the certificate that Veeam Backup Enterprise Manager Web site will use. If the setup does not find an appropriate certificate on the machine where Veeam Backup Enterprise Manager is installed, it will generate a self-signed certificate.

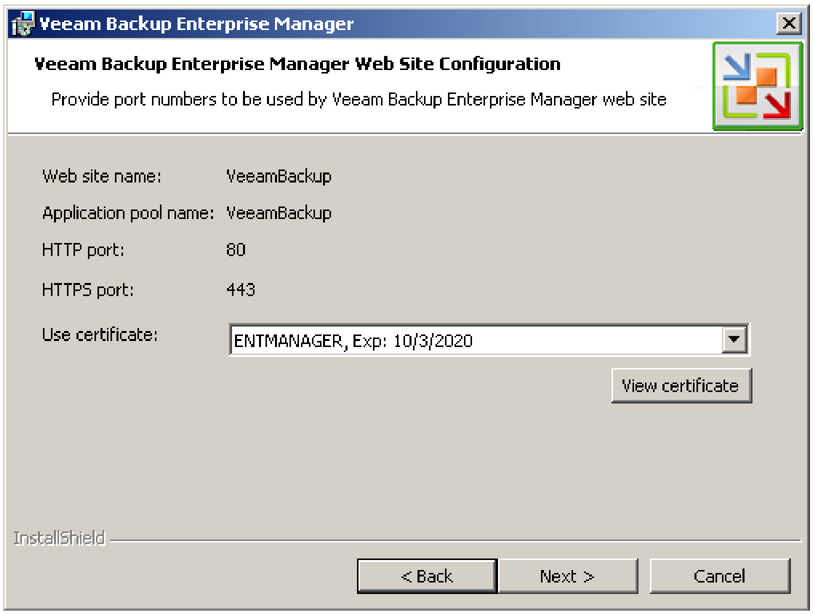

Veeam Backup Enterprise Manager Web Site Configuration

Provide port numbers to be used by Veeam Backup Enterprise Manager web site

Web site name: VeeamBackup

Application pool name: VeeamBackup

HTTP port: 80

HTTPS port: 443

Use certificate: ENTMANAGER, Exp: 10/3/2020

Step 9. Install Veeam Backup Enterprise Manager

Click **Next**, then click **Install**. The Veeam Backup Enterprise Manager will be installed on your computer. Once installation is complete, click **Finish** to finish working with the setup wizard.

Important! If Veeam Backup Enterprise Manager is installed on Windows XP OS, you should enable access to Veeam Backup Enterprise Manager Web site after installation. Select **Tools > Options** from the main Windows Explorer menu, then click **View** and clear the **Use simple file sharing** check box.

Installing Veeam Backup Search

Veeam Backup Search allows you to perform catalog replication and indexing which is required for the file search feature to work. It must be installed on a dedicated Microsoft Search Server. This section will guide you through the installation process of Veeam Backup Search.

Step 1. Install Microsoft Search Server

Select a machine which will function as a search server and install Microsoft Search Server on this machine. Keep in mind that Microsoft Search Server can be installed on Windows Server machine only.

Microsoft Search Server must be installed by a user who has administrator permissions on the computer.

To learn more about hardware and software requirements to Microsoft Search Server, see http://technet.microsoft.com/en-gb/library/bb905370(office.12).aspx (for Microsoft Search Server 2008) and http://technet.microsoft.com/en-gb/library/bb905370.aspx (for Microsoft Search Server 2010).

Step 2. Run Veeam Backup Search Setup

Run the setup file (*Veeam_Backup_Search_Setup_x64.exe* or *Veeam_Backup_Search_Setup_x86.exe*) from the downloaded archive on the machine where Microsoft Search Server is installed.

Step 3. Accept License Agreement

Read through the License Agreement and accept it to continue. If you select **I do not accept the terms in the license agreement**, the installation process will be terminated.

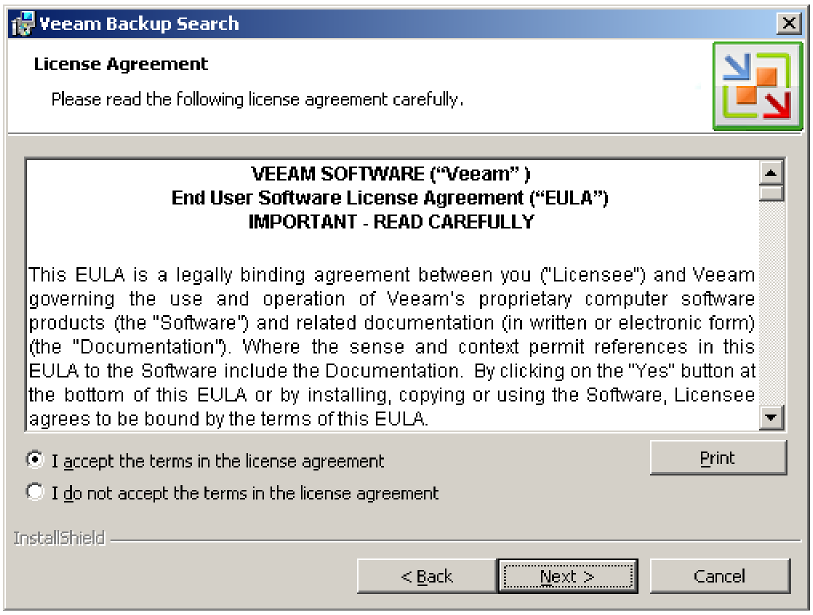

Veeam Backup Search

License Agreement

Please read the following license agreement carefully.

VEEAM SOFTWARE ("Veeam")
End User Software License Agreement ("EULA")
IMPORTANT - READ CAREFULLY

This EULA is a legally binding agreement between you ("Licensee") and Veeam governing the use and operation of Veeam's proprietary computer software products (the "Software") and related documentation (in written or electronic form) (the "Documentation"). Where the sense and context permit references in this EULA to the Software include the Documentation. By clicking on the "Yes" button at the bottom of this EULA or by installing, copying or using the Software, Licensee agrees to be bound by the terms of this EULA.

● I accept the terms in the license agreement

○ I do not accept the terms in the license agreement

Step 4. Choose Destination for Installation

Specify the installation folder. Use the **Change…** button to select a necessary installation folder. Please note that at least 51 MB is required to install the Veeam Backup Search component.

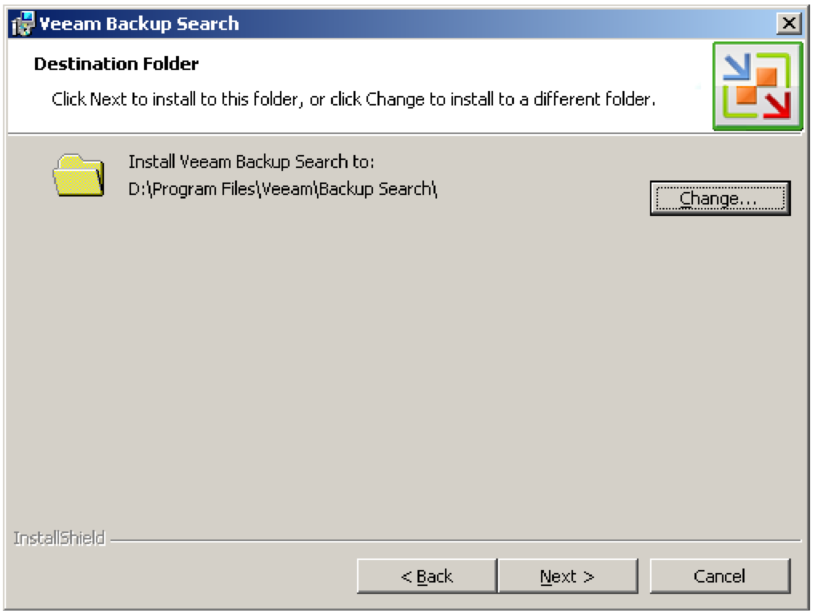

Veeam Backup Search

Destination Folder

Click Next to install to this folder, or click Change to install to a different folder.

Install Veeam Backup Search to:
D:\Program Files\Veeam\Backup Search\

Step 5. Specify Service Credentials

Specify the user name and password to be used by the Veeam Backup Search service. Change the number of TCP port if it is necessary. By default, Veeam Backup Search uses port number 9395.

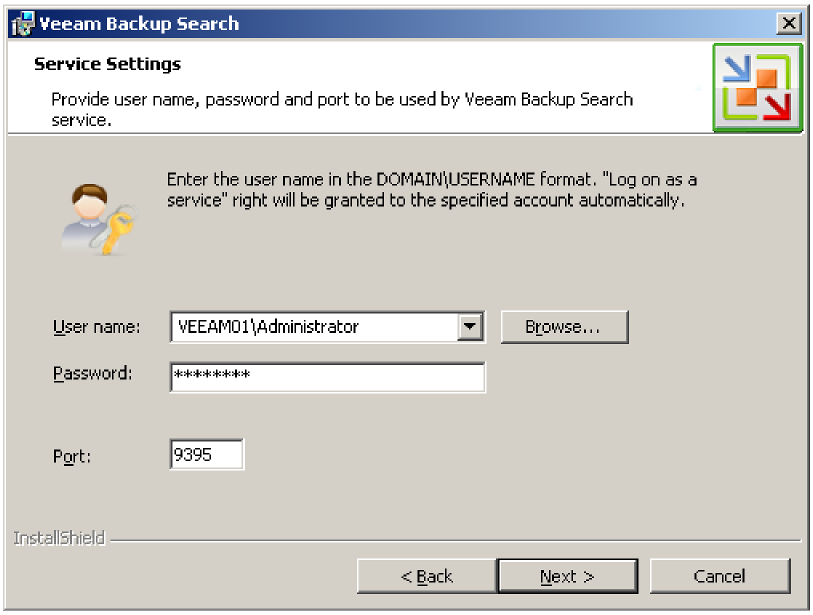

Veeam Backup Search

Service Settings

Provide user name, password and port to be used by Veeam Backup Search service.

Enter the user name in the DOMAIN\USERNAME format. "Log on as a service" right will be granted to the specified account automatically.

User name: VEEAM01\Administrator

Password: *********

Port: 9395

Step 6. Install Veeam Backup Search

Click **Install** to begin installation. Veeam Backup Search will be installed on your computer. Once installation is complete click **Finish** to exit the setup wizard.

Step 7. Specify the Default Account for Crawling Content

Specify the default account that will be used by Microsoft Search Server for crawling indexing content. This account must have read access to the shared *VBRCatalog* folder on the Veeam Backup Enterprise Manager server.

For Microsoft Search Server 2008:

1. Select **Programs > Search Server Administration** from the **Start** menu on the search server.
2. Click **Crawling** on the left.
3. On the **Crawling** page, click **Default content access account** and enter the required account data – user name and password.

For Microsoft Search Server 2010:

1. Select **Programs >Microsoft SharePoint 2010 Products > SharePoint 2010 Central Administration** from the **Start** menu on the search server.
2. Click **Application management** on the left. In the **Application Management** section, click **Manage service applications.**
3. On the **Manage Service Applications** page, click the **Search service application**.
4. In the **System Status** section on the **Search Administration** page, locate the default content access account, which is in the form *Domain\Username*. Click the default content access account name and enter the required account data — user name and password —in the displayed **Default Content Access Account** window.

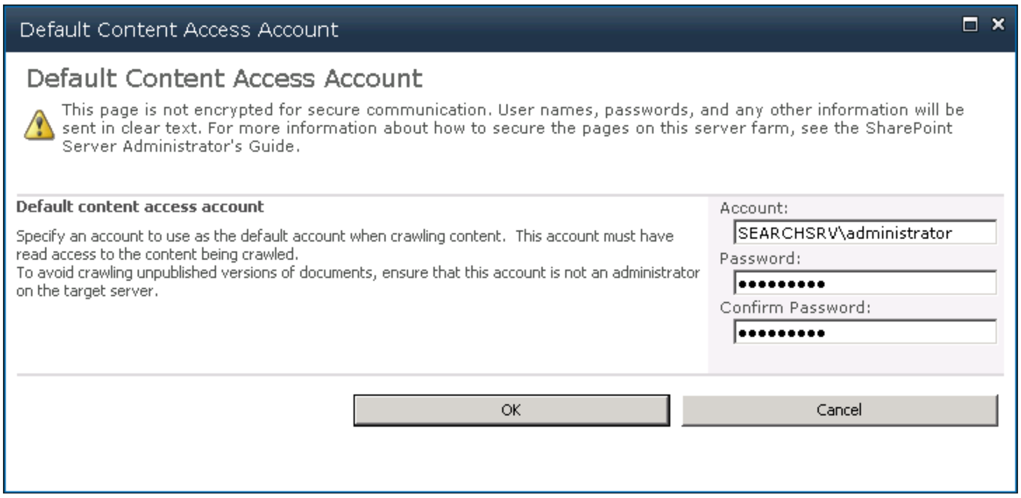

Default Content Access Account

This page is not encrypted for secure communication. User names, passwords, and any other information will be sent in clear text. For more information about how to secure the pages on this server farm, see the SharePoint Server Administrator's Guide.

Default content access account

Specify an account to use as the default account when crawling content. This account must have read access to the content being crawled.
To avoid crawling unpublished versions of documents, ensure that this account is not an administrator on the target server.

Account: SEARCHSRV\administrator

Password: ●●●●●●●●●

Confirm Password: ●●●●●●●●●

Enterprise and Standard Editions of Veeam Backup & Replication 5.0

Veeam Backup & Replication 5.0 is available in two editions — Standard and Enterprise. To get a desired edition of Veeam Backup & Replication, you should run the setup file and install a license of the proper type during the product setup. You can also re-install the license later at any moment of time: to do so, select **Help > License information…** from the main menu of Veeam Backup & Replication 5.0, click the **Install license** button and browse to a necessary license file.

The differences between the two editions are listed in the table below.

Feature	Standard Edition	Enterprise Edition
SureBackup recovery verification	Manual You can verify the recoverability of the latest backup by mounting a VM from the backup file with the help of instant VM recovery and manually testing it.	Automatic The recovery verification process is automated. You can select any restore point, not only the latest one.
Universal Application Item-Level Recovery	Not available.	Available Includes Active Directory restore wizard and the wizard for user-directed recovery.
Browsing and searching for VM guest OS files within indexed backups	You can browse and search for files in current backups only.	You can search and browse for files in both current and archived backups.

Veeam Backup & Replication Licensing

Veeam Backup & Replication 5.0 is licensed per CPU Socket ("CPU Sockets") for each Managed Server. Managed Server is defined as VMware ESX server that is backed up, recovered, collected data from or otherwise managed by the software. "CPU Sockets" means a single, physical chip that houses not more than six (6) processor cores on the managed server.

The trial license key is sent to you after registering the product with Veeam Software at: http://www.veeam.com/vmware-esx-backup/download.html. After registering the product you will receive a trial license key. The trial license is valid for 30 days from the moment of registration.

To obtain a full license key for the desired number of sockets, refer to http://www.veeam.com/buy-end-user.html.

The full license includes a one-year maintenance plan. To renew your maintenance plan, please contact Veeam Customer Support at: support@veeam.com.

Installing Veeam Backup & Replication License

When installing Veeam Backup & Replication 5.0, you will be asked to specify the license file that was sent to you after registration. If you do not have a license, you will not be able to install Veeam Backup & Replication 5.0.

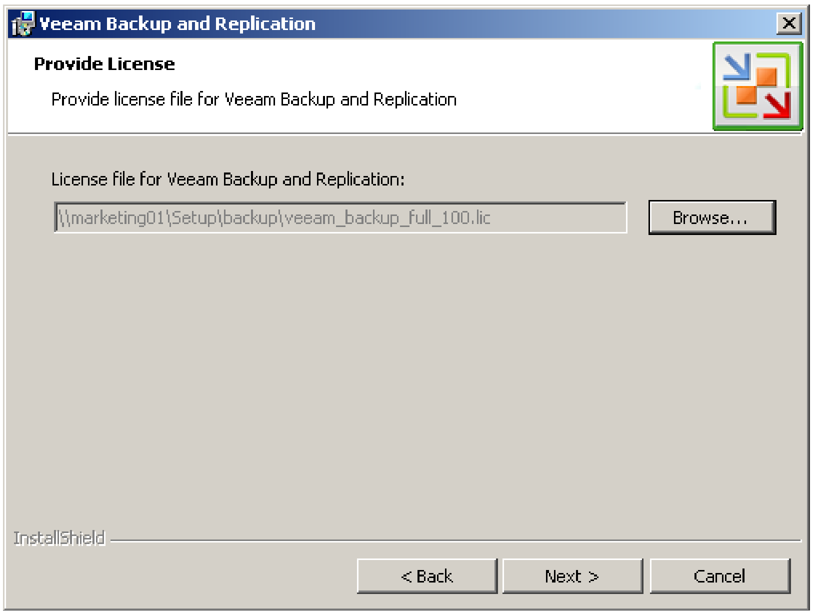

Provide License

Provide license file for Veeam Backup and Replication

License file for Veeam Backup and Replication:

\\marketing01\Setup\backup\veeam_backup_full_100.lic

To view information on the currently installed license, select **Help > License Information…** from the main menu. To change a license, click the **Install license** button and browse to a necessary .lic file.

To learn about license handling for backup servers added to Veeam Backup Enterprise Manager, see the Managing Licenses from Veeam Backup Enterprise Manager section.

Revoking ESX Servers from the License

Veeam Backup & Replication 5.0 offers you a possibility to revoke unused ESX servers from the license — that is, to re-use the license applied to one ESX server to another ESX server. This may be required if the ESX server to which the license is applied does not need backup or replication anymore (for example, in case it is no longer used).

To revoke an ESX server, select **Help > License Information…** from the main menu. In the displayed window, click the **Licensed hosts** button — as a result, the list of hosts using the license will be displayed.

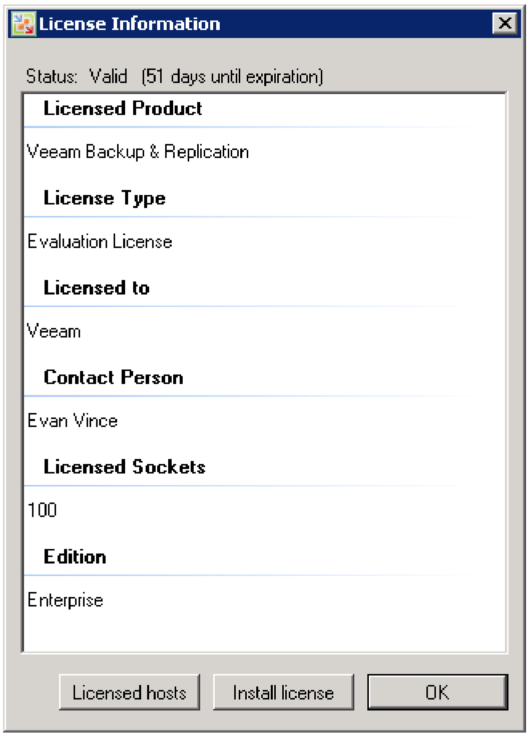

Status: Valid (51 days until expiration)

Licensed Product

Veeam Backup & Replication

License Type

Evaluation License

Licensed to

Veeam

Contact Person

Evan Vince

Licensed Sockets

100

Edition

Enterprise

The **Licensed hosts** list displays all ESX hosts to which the license is applied. When you start Veeam Backup & Replication 5.0 for the first time, the list will be empty. After you run a backup or replication job targeted at some VI objects, this section will display the list of ESX servers that were engaged in the job, with the number of sockets per each.

To revoke a specific ESX server, select it in the list and click the **Revoke** button. Licensed sockets used by it will be freed and will become available for use by other ESX servers.

Upgrading Veeam Backup & Replication

During the installation process, the Veeam Backup & Replication setup scans the system for previous versions and, if any is detected, upgrades it to a newer version.

The SQL database instance installed and used by the previous version of Veeam Backup is not removed at the uninstall process. All jobs data stored in it remains as well. At new installation, the database gets upgraded and becomes available for usage with the newly installed version.

Repairing and Uninstalling

To uninstall Veeam Backup & Replication 5.0, from the Start menu, select **Control Panel** > **Add or Remove Programs** > **Veeam Backup & Replication 5.0** and click the **Remove** button.

Then, repeat the procedure for the Veeam Backup Enterprise Manager component (if it is installed).

Configuring VCB Proxy (VMware Reference)

In case you are planning to take advantage of VCB-enabled backup, you should install Veeam Backup & Replication 5.0 on the VCB proxy server.

The VCB proxy server represents a dedicated physical server on which VMware Consolidated Backup and Veeam Backup & Replication 5.0 are installed. The server should run Microsoft Windows 2003 or Windows 2008 and have direct access to the ESX VMFS LUNs.

This section provides a description of the VCB proxy configuring, as well as main requirements to the VCB proxy.

Note: The present section contains a general description of the VCB proxy configuration. To learn more about the VMware Consolidated Backup and aspects of the VCB proxy configuring, see http://www.vmware.com/pdf/vi3_301_201_vm_backup.pdf.

Prerequisites

Before installing VMware Consolidated Backup on the VCB proxy server, make sure your VI environment meets the following requirements:

ESX Server

ESX Server should be set up to use VMware File System (VMFS) or virtual compatibility raw device mappings (RDMs). VCB does not support RDMs in physical compatibility mode.

SAN

- If the VCB proxy does not have access to storage LUNs managed by ESX Server systems, VCB-enabled backup will not be performed. Therefore, the VCB proxy should be added to the same fabric zones where your ESX servers reside. On the other hand, all VMs running on the local ESX storage should be moved to a SAN LUN for you to be able to back them up.
- For every LUN containing VMFS or RDM data, the LUN ID on the VCB proxy must match the UN ID as seen by the ESX Server.

VCB Proxy

- The VCB proxy server should run Microsoft Windows 2003. Besides, VCB 1.5 Update 1 offers official support for 32-bit and 64-bit versions of Microsoft Windows Server 2008 proxy server. Previous VCB versions provide only experimental support for Windows Server 2008.
- The VCB proxy requires the following hardware components: Network adapter (NIC) and Fibre Channel host bus adapter (HBA).
- Networking on the backup proxy should be configured so that the proxy can establish a connection to VirtualCenter. If there is a firewall between the backup proxy and the VirtualCenter, the firewall must permit TCP/IP connections to VirtualCenter. By default, VirtualCenter expects incoming connections at TCP/IP port 902.

Installing VCB Proxy

The present section describes the procedure of the VCB installation.

Step 1. Disable Automatic Drive-Letter Assignment

Microsoft Windows automatically assigns drive letters to each new visible NTFS or FAT volume. It is necessary to disable this feature so that volumes are not automatically mounted on the proxy.

Shut down the Windows proxy and disconnect it from the SAN. Boot the proxy and log on to the account with the administrator privileges. Then run the next commands using the command-line interface:

- *diskpart* — to start the diskpart utility
- *automount disable* — to disable automatic drive-letter assignment to newly seen volumes
- *automount scrub* — to clean out entries of previously mounted volumes in the registry
- *exit* — to exit the disport utility

Shut down Windows and reconnect the proxy to the SAN. Then boot the proxy.

Step 2. Install VMware Consolidated Backup

Log on to the proxy using an account with administrative privileges and install VMware Consolidated Backup by running the setup file from the VCB distribution.

Note: Ensure that you are using a correct VCB version. Depending on versions of ESX servers you have, you may need to install a specific VCB version. For more information, refer to the following compatibility matrix.

Step 3. Install Veeam Backup & Replication 5.0

Install Veeam Backup & Replication 5.0 on the configured VCB proxy. To learn more about the installation process, see the Installing Veeam Backup & Replication section.

ADMINISTRATION

This section provides description of main activities performed with Veeam Backup & Replication 5.0: creating backup, replication and VM copy jobs, performing file- and image-level restore, accomplishing replica failover, importing backups, reporting and logging.

Adding Servers

Prior to performing backup or replication processes, it is necessary to add servers you want to work with to Veeam Backup & Replication 5.0. You may add an ESX/ESXi server, VirtualCenter or Linux server. If you are planning to use an ESX server being a part of the VirtualCenter hierarchy, we recommend adding a corresponding VirtualCenter instead of a single ESX server to ensure more flexibility and convenience at work with servers.

To add a server, do one of the following:

- Right–click the **Servers** node in the management tree and select **Add Server**.
- Click the **Add Server** button on the toolbar.
- Select the **Add Server...** command from the **File** menu.
- Press **Alt+A** on the keyboard.

Then follow the **Add Server** wizard steps.

Adding VirtualCenter

Step 1. Specify Server Type and Name

Enter a full DNS name or IP address of the server and select the server type: **vCenter server**.

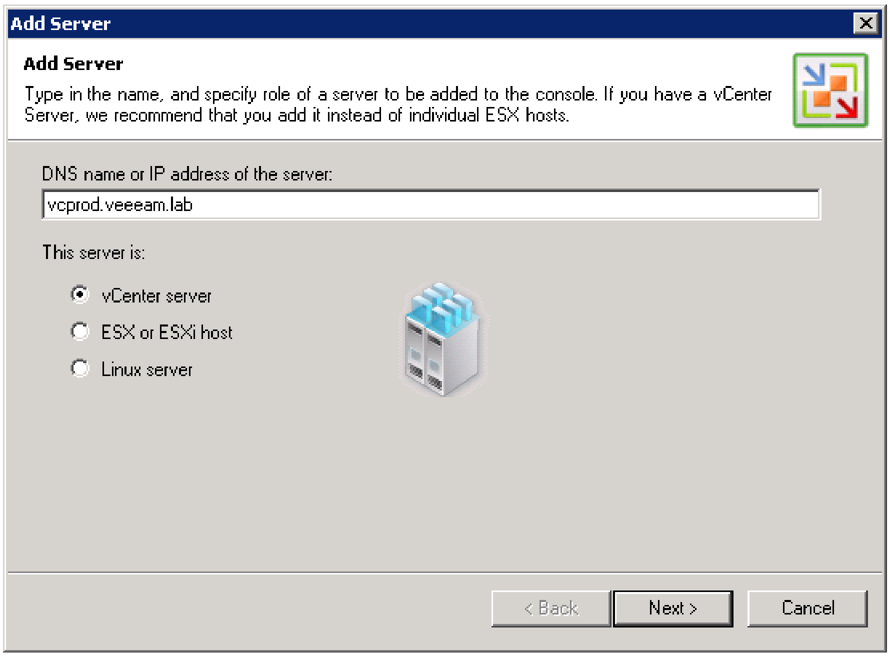

Add Server

Add Server

Type in the name, and specify role of a server to be added to the console. If you have a vCenter Server, we recommend that you add it instead of individual ESX hosts.

DNS name or IP address of the server:

vcprod.veeeam.lab

This server is:

● vCenter server

○ ESX or ESXi host

○ Linux server

Step 2. Specify Server Connection Settings

At this step, you should enter administrator's credentials to connect to the VirtualCenter server: user name and password. To avoid problems, we recommend specifying the user name in the *DOMAIN\USERNAME* format.

Select the **Save password** check box. Otherwise the entered credentials will be used for one work session of Veeam Backup & Replication 5.0. When Veeam Backup & Replication 5.0 is closed and started anew, you will have to enter credentials again as soon as the server is addressed.

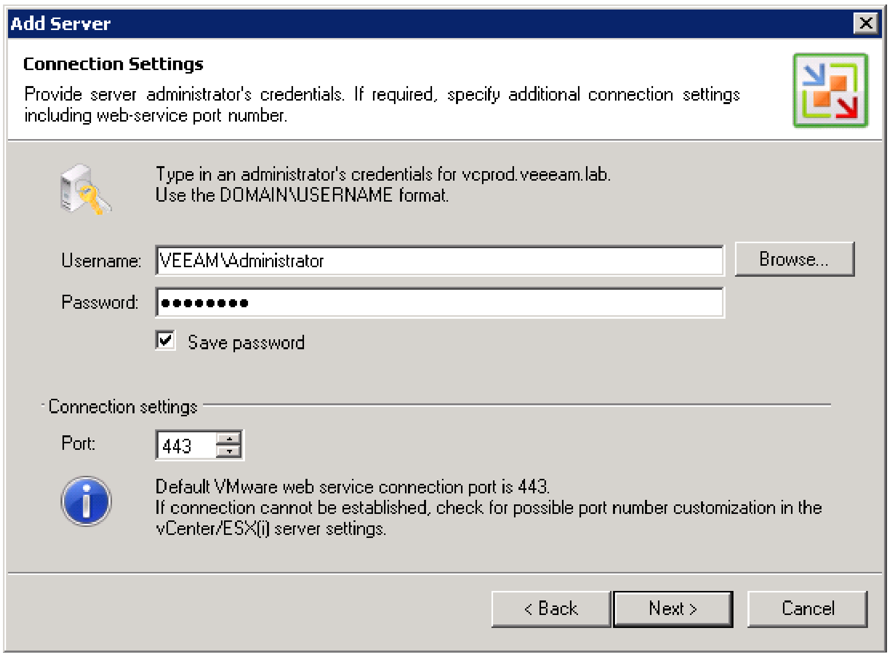

Add Server

Connection Settings

Provide server administrator's credentials. If required, specify additional connection settings including web-service port number.

Type in an administrator's credentials for vcprod.veeeam.lab.
Use the DOMAIN\USERNAME format.

Username: VEEAM\Administrator

Password: ●●●●●●●●●

☑ Save password

Connection settings

Port: 443

Default VMware web service connection port is 443.
If connection cannot be established, check for possible port number customization in the vCenter/ESX(i) server settings.

Change the web service port if necessary. By default, port 443 is used for VMware vCenter and VMware ESX.

Step 3. Finish Working with the Wizard

If you want to connect to the added VirtualCenter server on finishing work with the wizard, select the **Connect when I click Finish** check box. Then click **Finish**.

Adding ESX/ESXi Server

Step 1. Specify Server Type and Name

Enter a full DNS name or IP address of the server and select the server type: **ESX or ESXi host**.

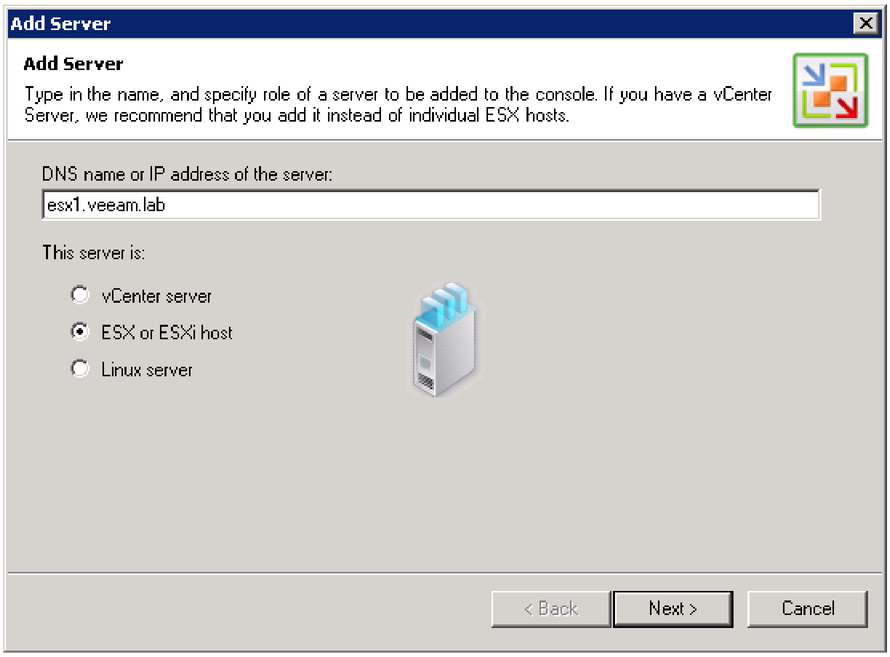

Add Server

Add Server

Type in the name, and specify role of a server to be added to the console. If you have a vCenter Server, we recommend that you add it instead of individual ESX hosts.

DNS name or IP address of the server: esx1.veeam.lab

This server is:

○ vCenter server

● ESX or ESXi host

○ Linux server

Step 2. Specify Server Connection Settings

At this step, you should enter administrator's credentials to connect to the ESX/ESXi server: user name and password. Select the **Save password** check box. Otherwise the entered credentials will be used for one work session of Veeam Backup & Replication 5.0. When Veeam Backup & Replication 5.0 is closed and started anew, you will have to enter credentials again as soon as the server is addressed.

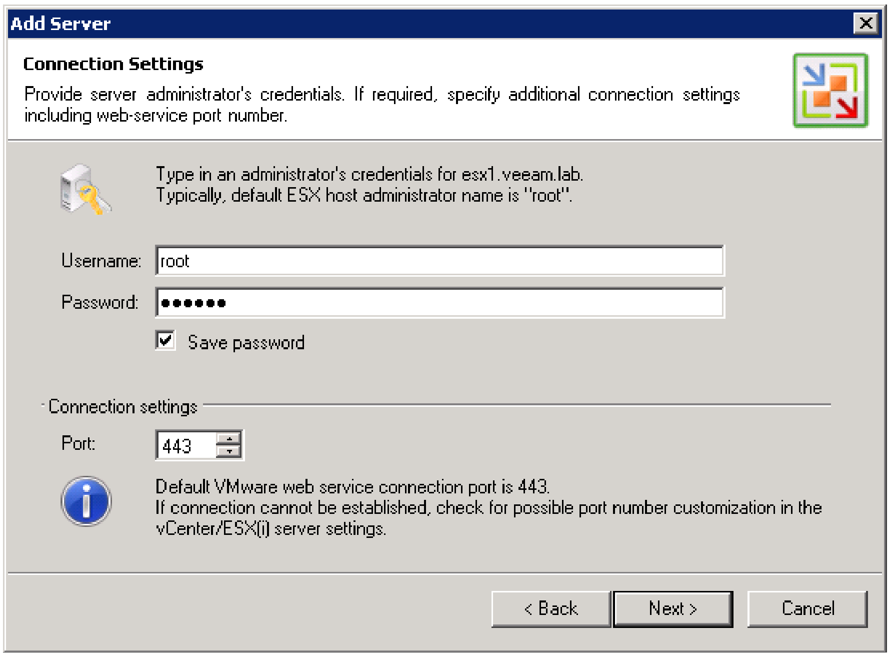

Add Server

Connection Settings

Provide server administrator's credentials. If required, specify additional connection settings including web-service port number.

Type in an administrator's credentials for esx1.veeam.lab.
Typically, default ESX host administrator name is "root".

Username: root

Password: ••••••

☑ Save password

Connection settings

Port: 443

Default VMware web service connection port is 443.
If connection cannot be established, check for possible port number customization in the vCenter/ESX(i) server settings.

Change the web service port if necessary. By default, port 443 is used for VMware vCenter and VMware ESX.

Step 3. Specify Service Console Connection Settings

This step is available if you are adding the ESX server only; when adding the ESXi server, you will pass immediately to step 4.

At this step, you should specify service console connection settings and adjust SSH port number if necessary. Specifying console connection settings is optional. If you do not want to use the service console, clear the **Use service console connection to this server** check box and click **Next**. In this case, Veeam Backup & Replication 5.0 will work with the server in the agentless mode. The agentless mode may be used to work with ESX server 3.5 and higher; for ESX server 3.0, the agentless mode is not supported. However, we recommend that you work with ESX servers using the service console.

By default, the **Use service console connection to this server** check box is selected. Enter the user name and password to connect to the service console of the server. Select the **Save password** check box. Otherwise the entered credentials will be used for one work session of Veeam Backup & Replication 5.0. When Veeam Backup & Replication 5.0 is closed and started anew, you will have to enter credentials again as soon as the server is addressed.

If you choose to use a non-root account that does not have sudo permissions on the ESX server, you can use the **Non-root account** options section to grant sudo rights to this account. Select the **Elevate account to root** check box to provide a non-root user with access to the added server. You can add the account to sudoers file automatically by selecting the **Add account to the sudoers file automatically** check box. If you do not select this option, you will have to manually add the user to the sudoers file.

Note: Make sure that in the sudoers file the *NOPASSWD:ALL* option is enabled for the user account you want to elevate to root to prevent the user from entering a password.

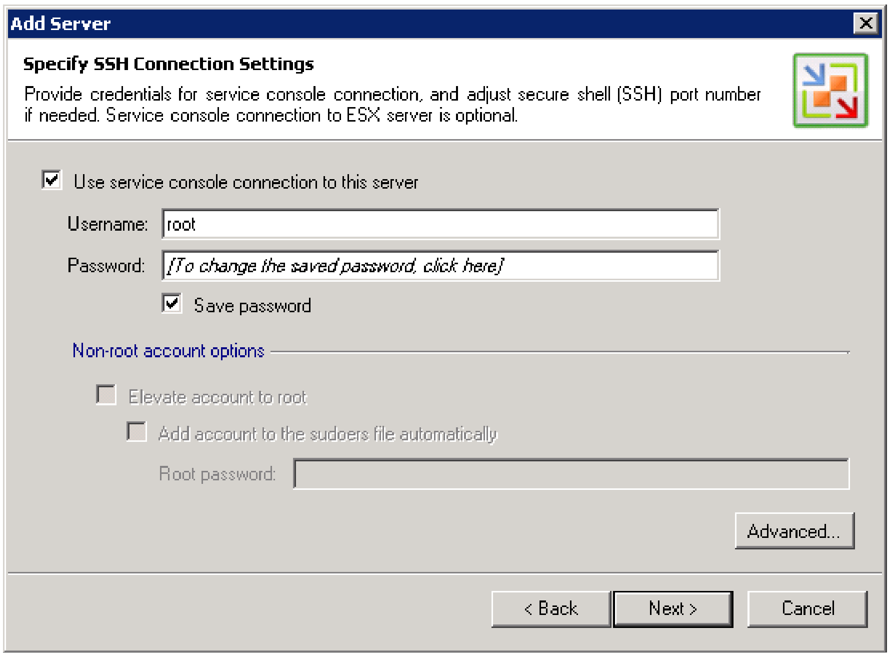
Add Server

Specify SSH Connection Settings

Provide credentials for service console connection, and adjust secure shell (SSH) port number if needed. Service console connection to ESX server is optional.

☑ Use service console connection to this server

Username: root

Password: [To change the saved password, click here]

☑ Save password

Non-root account options

□ Elevate account to root

□ Add account to the sudoers file automatically

Root password:

Click the **Advanced...** button to change advanced SSH settings: SSH port and SSH timeout.

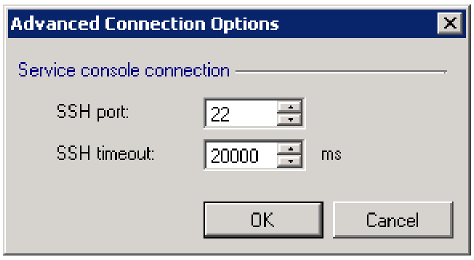
Advanced Connection Options

Service console connection

SSH port: 22

SSH timeout: 20000 ms

Tip: You can start PuTTY, a popular SSH client enabling safe logging on to a remote server, directly from Veeam Backup & Replication 5.0. Select **Tools > PuTTY...** from the main menu. To learn about PuTTY, see http://www.chiark.greenend.org.uk/~sgtatham/putty/docs.html.

Step 4. Finish Working with the Wizard

If you want to connect to the added ESX server on finishing work with the wizard, select the **Connect when I click Finish** check box. Then click **Finish**.

Adding Linux Server

Step 1. Specify Server Type and Name

Enter the full DNS or IP address of the server and select the server type: **Linux server**.

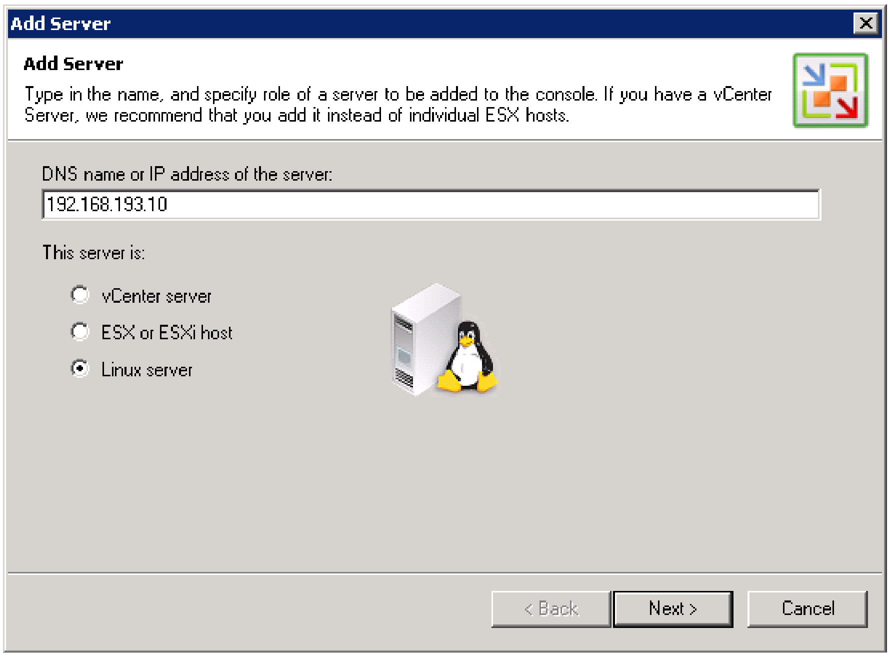

Add Server

Add Server

Type in the name, and specify role of a server to be added to the console. If you have a vCenter Server, we recommend that you add it instead of individual ESX hosts.

DNS name or IP address of the server:

192.168.193.10

This server is:

○ vCenter server

○ ESX or ESXi host

● Linux server

Step 2. Specify SSH Connection Settings

At this step, you should enter administrator's credentials to connect to the Linux server: user name and password. Select the **Save password** check box. Otherwise the entered credentials will be used for one application session. When Veeam Backup & Replication 5.0 is started anew, you will have to enter credentials again as soon as the server is addressed.

If you choose to use non-root account that does not have sudo permissions on the Linux server, you can use the **Non-root account** options section to grant sudo rights to this account. Select the **Elevate account to root** check box to provide a non-root user with access to the added server. You can add the account to sudoers file automatically by selecting the **Add account to the sudoers file automatically** check box. If you do not select this option, you will have to manually add the user to the sudoers file.

Note: Make sure that in the sudoers file the *NOPASSWD:ALL* option is enabled for the user account you want to elevate to root to prevent the user from entering a password.

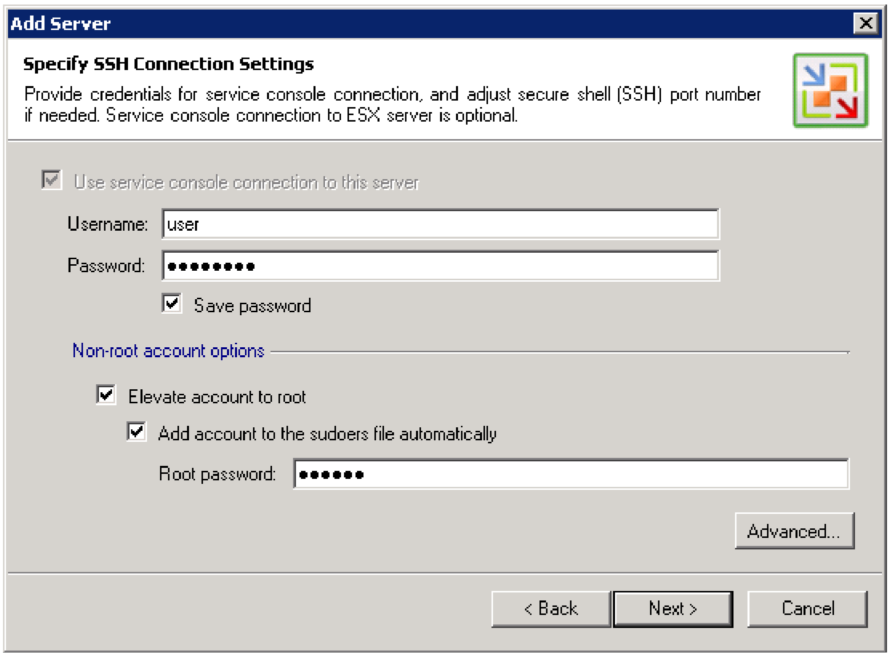

Add Server

Specify SSH Connection Settings

Provide credentials for service console connection, and adjust secure shell (SSH) port number if needed. Service console connection to ESX server is optional.

☑ Use service console connection to this server

Username: user

Password: ●●●●●●●●

☑ Save password

Non-root account options

☑ Elevate account to root

☑ Add account to the sudoers file automatically

Root password: ●●●●●●

Click the **Advanced...** button to change advanced SSH settings: SSH port and SSH timeout.

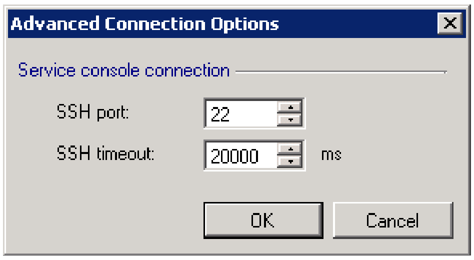

Advanced Connection Options

Service console connection

SSH port: 22

SSH timeout: 20000 ms

Step 3. Finish Working with the Wizard

If you want to connect to the added Linux server on finishing work with the wizard, select the **Connect when I click Finish** check box. Then click **Finish**.

Disconnecting and Removing Servers

Along with adding servers (ESX hosts, VirtualCenter servers and Linux servers), you have an ability to use the following options: **Disconnect server** and **Remove Servers** from your Veeam Backup & Replication console. All options are available from the shortcut menu.

- The **Disconnect** option is used to close connection to any host, VirtualCenter, Linux server previously added to the console.
- If you choose the **Remove Servers** option, all objects (such as jobs, backups, replicas) that have references to the ESX host or VirtualCenter server being removed will be deleted from the configuration SQL database and the Veeam Backup & Replication console.
  Please note that all backup files (VBK and VRB) will stay intact, so you can easily import these files later to the Veeam Backup & Replication console for restore operations if needed.
  As for replication jobs, all references to replicas will be removed from the Veeam Backup & Replication console. However, all replicated VMs will still reside on your target hosts, so you can start them manually after Remove Servers option is performed. Also you can power on replicated VMs to any available restore point in time using the Veeam Backup & Replication console before performing **Remove Servers** option. To learn more, see the Failing Over VM Replicas section.

Please note that import option for replicas is not supported.

Managing Backup Jobs

Any backup, replication or VM copy operation performed with Veeam Backup & Replication 5.0 is run within the frames of a separate job. A job is a specific task that can be accomplished immediately after its creation, saved or scheduled for specific time. Every job is marked with one of the following types: *Backup*, *Replication*, *VM Copy* and *File Copy*. To create a job, the user should run a corresponding wizard and complete all wizard steps.

All created jobs are listed in the **Jobs** section under the **Backup** node in the management tree. You can edit job properties, start and stop jobs, re-start failed jobs, view job statistics data and delete unnecessary jobs. Commands for any of the listed operations are available from the shortcut menu.

Tip: When a job is being run, Veeam Backup & Replication 5.0 checks disk space on the destination storage. If the disk space is below a specific value, a warning will be displayed. To specify the disk space threshold, select **Tools > Options…** from the main menu. On the **Global Warnings** tab, specify the amount of free disk space required in percent.

Creating a Backup Job

To perform backup of a VM, you should create a backup job by means of the **New Backup Job** wizard. You can perform the created job immediately, schedule or save it. This section will guide you through all steps of the wizard and provide explanation on offered options.

Before You Begin

- Prior to creating a backup job, make sure you have enough free space on the destination disk. When a backup job runs for the first time, full backup is performed: the disk space required is equal to the actual size of a virtual machine adjusted to the selected compression level. At all subsequent backups, only incremental data will be saved.
  To learn how much disk space is available on storage devices used by a specific server, right–click a corresponding server in the management tree, select the **Properties** command from the shortcut menu and click the **Populate** button. You will also be able to check disk space resources right from the wizard.
- Make sure all servers you want to work with are available in the management tree: you will not be able to add them once the **New Backup Job** wizard is launched.

Step 1. Launch the New Backup Job Wizard

To run the **New Backup Job** wizard, do one of the following:

- Click the **Backup** button on the toolbar.
- Select **Backup > Backup…** from the main menu.
- Click **Jobs** under the **Backup and Replication** node in the management tree, right-click anywhere on the blank area of the informational pane and select **Backup…**.
- Right-click the **Backups** node under **Backup and Replication** in the management tree and select **Backup...** from the shortcut menu.

Step 2. Specify Job Name and Description

At the first step of the wizard, enter a name and description of the created job. By default, the following description is initially provided for the created job: time at which the job was created and user who created the job.

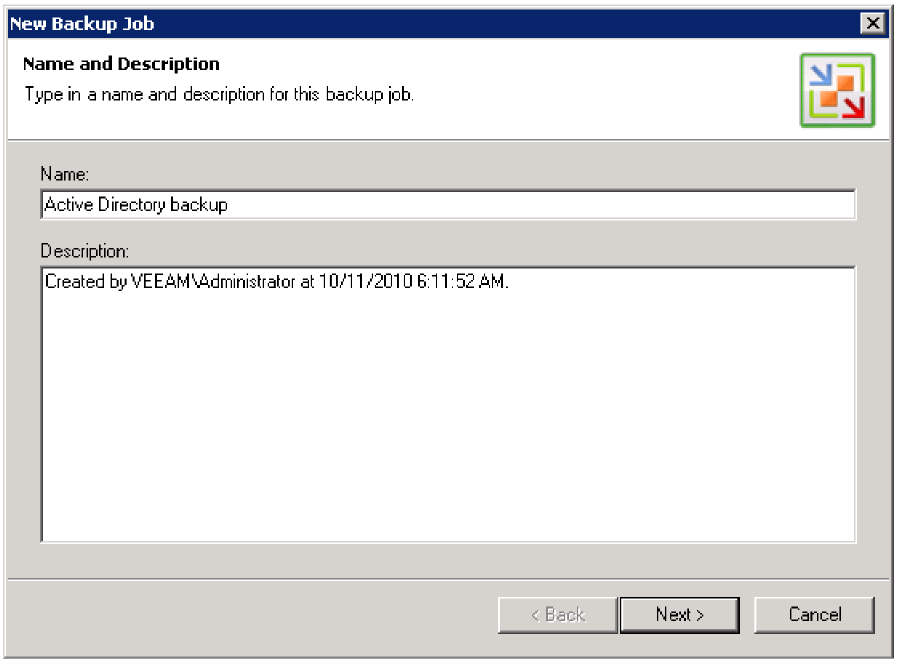

Active Directory backup

Created by VEEAM\Administrator at 10/11/2010 6:11:52 AM.

Step 3. Select Backup Mode

You can back up VMs in one of the three modes using VMware vStorage APIs— **Direct SAN access**, **Virtual Appliance** and **Network** mode.

- By default, if the **Direct SAN access** or **Virtual Appliance** mode is selected, Veeam Backup & Replication will automatically fail over to network data transfer in case the primary selected backup mode fails during the job run. To disable failover, click the **Advanced…** button and clear the **Failover to network mode if primary backup mode fails** check box.
- If the **Network** VMware vStorage APIs mode is selected, you can choose to transfer disks data over encrypted SSL connection. Click the **Advanced…** button and select the **Encrypt LAN traffic** check box. Use of encryption puts more stress on CPU of an ESX server, providing, however, secure data transfer.

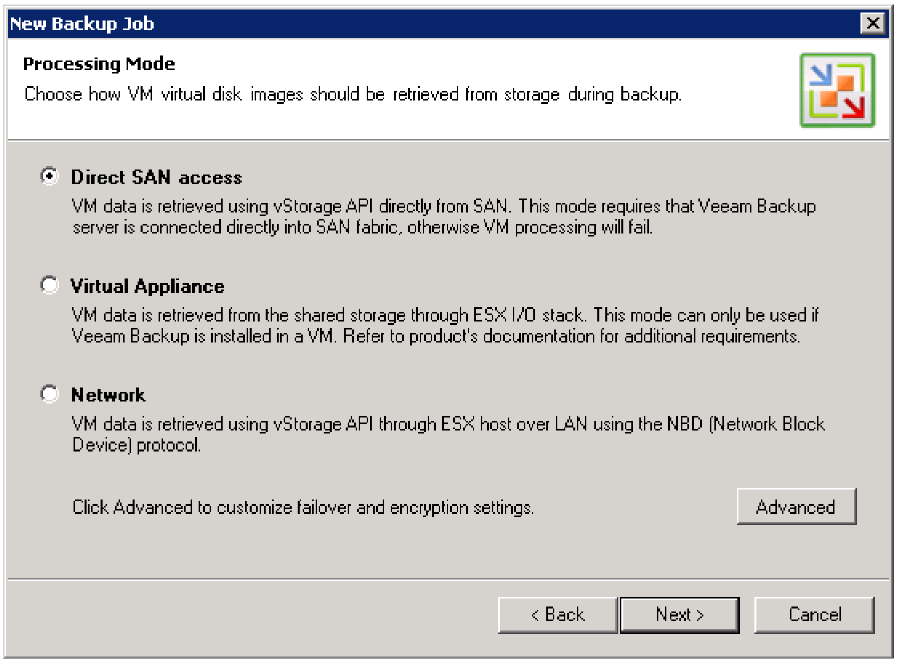
New Backup Job
Processing Mode
Choose how VM virtual disk images should be retrieved from storage during backup.

● Direct SAN access

VM data is retrieved using vStorage API directly from SAN. This mode requires that Veeam Backup server is connected directly into SAN fabric, otherwise VM processing will fail.

○ Virtual Appliance

VM data is retrieved from the shared storage through ESX I/O stack. This mode can only be used if Veeam Backup is installed in a VM. Refer to product's documentation for additional requirements.

○ Network

VM data is retrieved using vStorage API through ESX host over LAN using the NBD (Network Block Device) protocol.

Click Advanced to customize failover and encryption settings.

You can also choose one of the legacy modes — **VCB-enabled backup** or **Network backup**. To enable legacy modes, select **Tools > Options…** from the main menu of Veeam Backup &

Replication, click the **Advanced** tab and select the **Enable legacy processing modes** check box. Legacy modes may be used for ESX/ESXi servers earlier than 3.5; for ESX/ESXi servers 3.5 and higher, it is recommended to use vStorage API backup modes.

Step 4. Select Virtual Machines to Back Up

At this step, you should select an individual VM or a VM container you want to back up. Jobs with VM containers are dynamic in their nature: if a new VM is added to the container after a backup job is created, the job will be automatically updated to include the added VM.

Click the **Add…** button to browse to VMs and VM containers that should be backed up. In the displayed VI tree, select a necessary object and click the **Add** button.

To facilitate objects selection, you can:

- Switch between VI views: click the **Hosts and Clusters**, **VMs and Templates** or **Datastores and VMs** buttons at the top of the tree.
- Use a search field at the bottom of the window: click the button on the left of the field to select a necessary type of object that should be searched for (*Everything*, *Folder*, *Cluster*, *Host*, *Resource Pool*, *Virtual Application* or *VM*), enter an object's name or a part of it and click the **Start search** button on the right.

Note: Depending on the view you select, some VI objects may be not available: for example, if you select the VMs and Templates view, you will not be able to see and find resource pools.

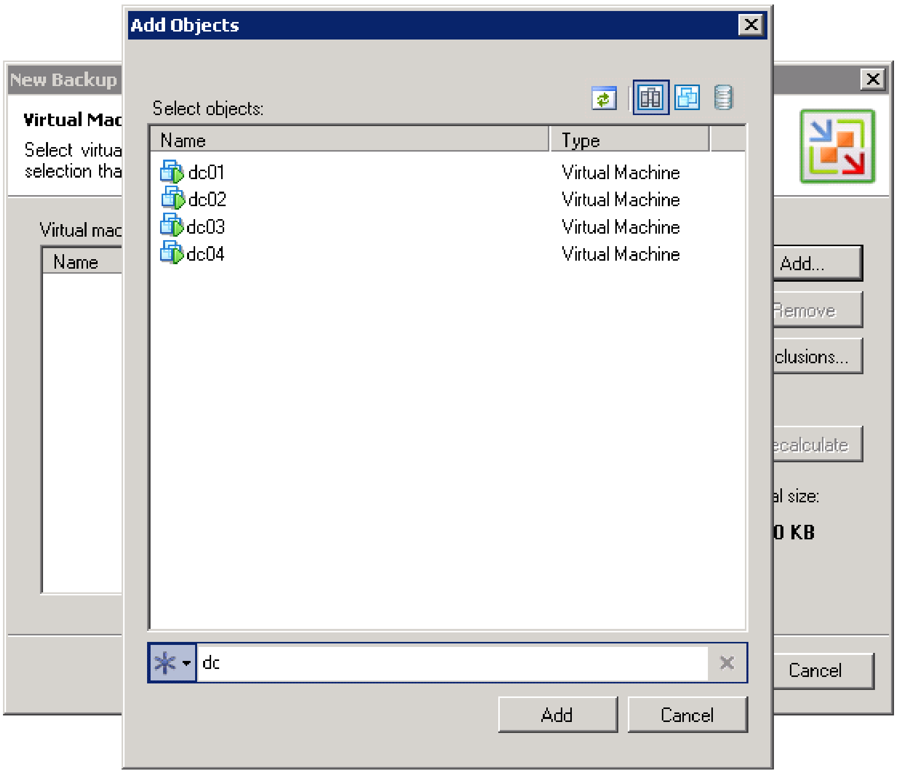

Name	Type
dc01	Virtual Machine
dc02	Virtual Machine
dc03	Virtual Machine
dc04	Virtual Machine

To remove an object from the list, select it and click the **Remove** button on the right.

The initial size of VMs and VM containers added to a backup job is displayed in the **Size** column in the list. The total size of backed up objects is displayed in the **Total size** field. Use the **Refresh** button to refresh the total size value after you add a new object to the job.

Step 5. Exclude Objects from Backup Job

After you have added VMs and VM containers to the list, you can specify which objects should be excluded from backup. Veeam Backup & Replication 5.0 allows excluding the following types of objects: VMs and VM templates from VM containers, as well as specific VM disks.

To select which objects should be excluded, click the **Exclusions…** button on the right.

- To exclude VMs from a VM container (for example, if you need to back up the whole ESX server excluding several VMs running on this server), click the **VMs** tab. Click the **Add…** button on the right and select VMs that should be excluded.
  To display all hosts added to Veeam Backup & Replication 5.0, select the **Show full hierarchy** check box.
  To facilitate objects selection, you can switch between the **Hosts and Clusters**, **VMs and Templates** and **Datastores and VMs** views, and use the search field just as in the main window of the wizard.
- To select what VM disks you want to back up, click the **Disks** tab, select a necessary VM in the list and click the **Edit…** button. If a VM is not in the list, you can add it by clicking the **Add...** button. You can choose to process all disks, 0:0 disks (typically, the system disks) or select custom disks.
  If you select the **Remove excluded disks from VM configuration** check box, Veeam Backup & Replication 5.0 will modify VMX file to remove disks you want to skip from VM configuration. If this option is used, you will be able to restore, replicate or copy VM to a location where excluded disks are not accessible with the original paths. If you do not use this option, you will have to manually edit VM configuration file to be able to power on a VM.
- If you select to use the **Network** backup mode, you can back up VM templates together with VMs. Click the **VM Templates** tab. By default, the **Backup VM templates** check box is selected. Clear it if you do not want to include VM templates into the backup. The **Exclude templates from incremental backup** option allows you to include VM templates into a full backup only.

Note: Veeam Backup & Replication 5.0 automatically excludes VM log files from backup to make backup process faster and reduce the size of a backup file.

Step 6. Specify Backup Destination

At this step of the wizard, you should select destination for the created backup.

From the **Destination** list, select a host where the created backup should be stored. The list contains servers that were added to Veeam Backup & Replication 5.0. You can store a backup to a local host, network shared folders, and hosts added to the Veeam Backup & Replication 5.0. Use the **Host Properties...** button to view available disk resources, specify SSH and SOAP connection and data transfer information for a selected host.

Important! ESXi cannot be used to accommodate backup (VMware limitation). For this reason, ESXi servers are not displayed in the **Destination** list.

In the **Path to folder** field, specify a folder where the created backup should be stored. Make sure you have enough free space on your storage device. Use the **Check Space** button to check how much free space is available on the backup destination, and how much space you will require to store a full backup and its increments according to specified retention policy settings.

In the **File name** field, enter a name for the created backup file.

Specify the number of restore points that should be kept on the disk. If this number is exceeded, the earliest restore point will be deleted. The number of restore points is a relative value and doesn't correspond to the number of days to store them.

Please keep in mind that such retention policy mechanism works for reversed incremental backup. To learn about the retention policy for incremental backup, see the Backup Retention Policy section.

Specify a retention period for deleted VMs. When a VM is deleted and new backups are not created for it, the number of restore points specified previously stays on the disk. By specifying a retention period you make sure that these unnecessary restore points are removed from the target storage after the specified number of days.

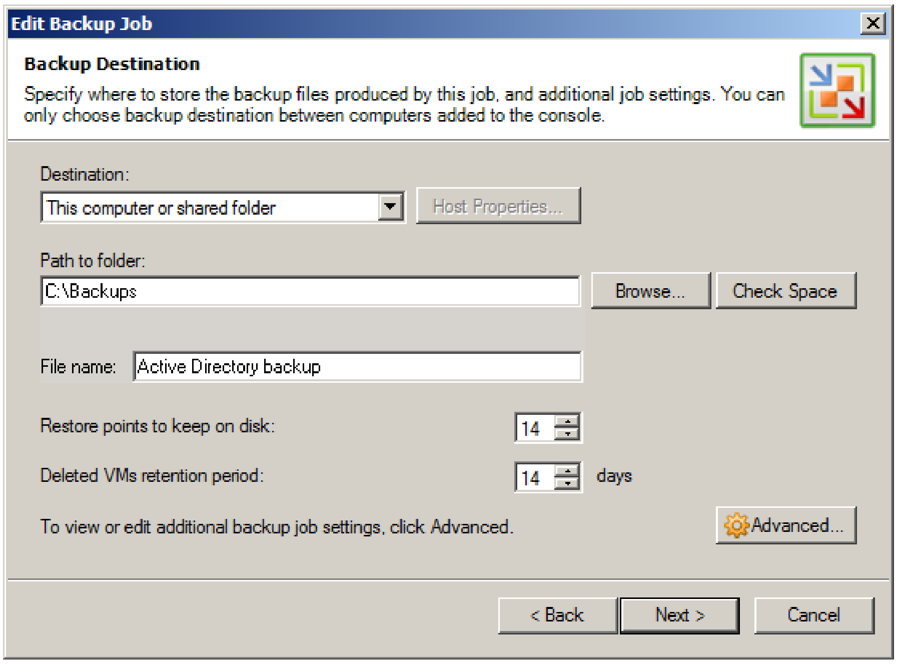

Edit Backup Job

Backup Destination

Specify where to store the backup files produced by this job, and additional job settings. You can only choose backup destination between computers added to the console.

Destination: This computer or shared folder [Host Properties...]

Path to folder: C:\Backups [Browse...] [Check Space]

File name: Active Directory backup

Restore points to keep on disk: 14

Deleted VMs retention period: 14 days

To view or edit additional backup job settings, click Advanced. [Advanced...]

[< Back] [Next >] [Cancel]

Step 7. Specify Advanced Backup Settings

Click the **Advanced…** button to specify advanced options for the created backup job:

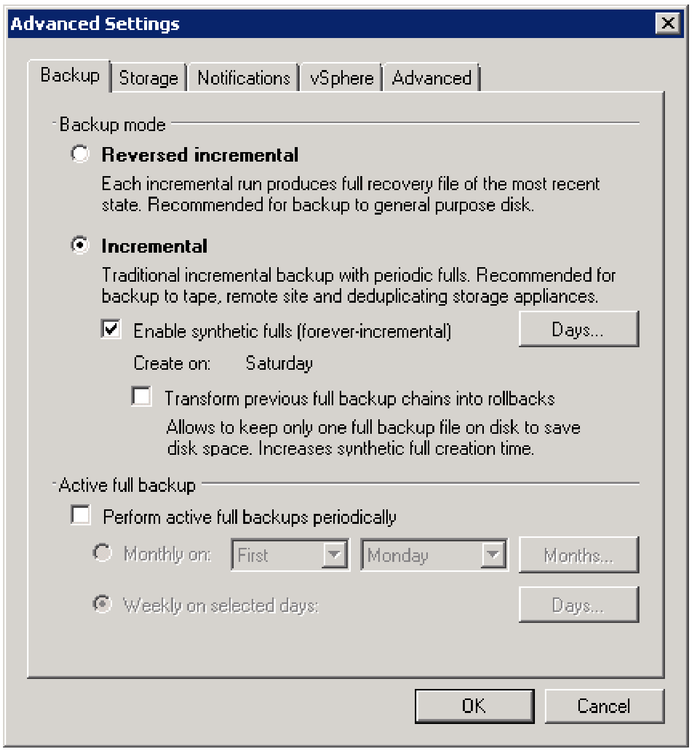

Advanced Settings

Backup | Storage | Notifications | vSphere | Advanced

Backup mode

○ **Reversed incremental**

Each incremental run produces full recovery file of the most recent state. Recommended for backup to general purpose disk.

● **Incremental**

Traditional incremental backup with periodic fulls. Recommended for backup to tape, remote site and deduplicating storage appliances.

☑ Enable synthetic fulls (forever-incremental) [Days...]

Create on: Saturday

□ Transform previous full backup chains into rollbacks

Allows to keep only one full backup file on disk to save disk space. Increases synthetic full creation time.

Active full backup

□ Perform active full backups periodically

○ Monthly on: First Monday [Months...]

● Weekly on selected days: [Days...]

[OK] [Cancel]

Backup settings

Select the method you want to use to back up VMs — **Reversed incremental** or **Incremental**. To learn about incremental and reversed incremental backup methods, see the Backup Methods section.

If you choose the **Incremental** method, you must select to periodically create a full synthetic backup or perform active full backups regularly.

- To create a full synthetic backup, select the **Enable synthetic fulls (forever incremental)** check box and click the **Days…** button to select necessary days. You can additionally choose to transform all previous full backup chains to the reversed incremental backup sequence. To do so, select the **Transform previous full backup chains into rollbacks** check box. Veeam Backup & Replication will leave only one full backup on the disk, and all .vib files will be transformed to .vrb files. Keep in mind that .vbk files that were created manually via the shortcut menu or using the **Perform active full backups periodically** option will not be transformed or deleted.
  The transformation option allows you to keep only one full backup on disk and so reduce the amount of space required to store backups. At the same time, it takes more time than simply creating a full synthetic backup.
- To perform full backups regularly, select the **Perform active full backups periodically** check box and define scheduling settings. The created full backup will be used as an initial point for subsequent increments.
  Note that if the active full backup and transformation to the reversed incremental backup are scheduled on the same day, only full backup will be performed – transformation will be skipped.

Tip: Before you select to perform periodic full backup, make sure you have enough free space on the backup destination. As an alternative, you can perform full backup manually: right-click a created backup job in the list and select **Perform Full Backup** from the shortcut menu.

Storage settings

In the **Storage** section, select the type of backup target you are planning to use. Depending on the chosen option, Veeam Backup & Replication will use data blocks of different size to optimize the size of backups and job performance:

- **Local target**. This option is recommended if you are planning to use SAN, DAS or local storage as a target. SAN identifies larger blocks of data and therefore can process larger quantities of data at a time. This option provides the fastest backup job performance but reduces the de-duplication ratio — the larger a data block is, the lower is the chance to find an identical block.
- **LAN target**. This option is recommended for NAS and on-site replication. It provides a better de-duplication ratio and reduces the size of an incremental backup file.
- **WAN target**. This option is recommended if you are planning to use WAN for offsite backup. Veeam Backup & Replication uses small data blocks, which results in the maximum de-duplication ratio and the smallest size of a backup file, allowing you to reduce the amount of traffic over the WAN link.

You can disable de-duplication at all by clearing the **Enable inline deduplication** check box. De-duplication provides a smaller size of a resulting backup file but may reduce backup performance.

In the **Compression** section, specify a compression level for the created backup: *None*, *Low*, *Optimal* or *Best*. To learn more about compression, see the De-duplication and Compression section.

Notifications settings

- Select the **Send e-mail notifications to the following recipients** check box if you want to receive notifications by e–mail in case of job failure or success. In the field below, specify a recipient's e-mail address. You can enter several addresses separated by a semicolon.

  E–mail notifications will be sent only if you have selected the **Enable email notification** check box in the **Options** window and specified e–mail notification settings (select **Tools > Options...** from the main menu). To learn more, see the Specifying Notification Settings section.
- Select the **Enable SNMP notification for this job** check box if you want to receive SNMP traps when a job is completed and a backup is created. SNMP traps will be sent if

you configure SNMP settings in Veeam Backup & Replication and on the recipient's computer. To learn more, see the Specifying SNMP Settings section.

- In the **VM notes** section, select the **Set successful backup details to this VM attribute** check box to write to a VM custom attribute information about successfully performed backup and data on backup results (backup date and time, backup console name, and path to the backup file). In the field below, enter the name of a necessary attribute. If the specified attribute does not exist, Veeam Backup & Replication 5.0 will create it.

vSphere settings

Use the **vSphere tab** to specify if vSphere changed block tracking should be used. By default, this option is selected. If you want to force using changed block tracking for VMs for which changed block tracking is disabled on the ESX server, select the **Enable changed block tracking for all processed VMs** check box. Please note that you can use this option only for VMs using virtual hardware version 7 or later.

Advanced settings

- The **Enable VMware tools quiescence** option enables freezing of the file-system for proper snapshot creation. With this option enabled, creation of a snapshot is performed with the help of the sync driver responsible for holding incoming I/O and flushing all dirty data to a disk, thus making the file systems consistent.
  It is strongly recommended to leave this option disabled if you are backing up Windows systems that support Windows VSS — for these systems, it is recommended to use the **Enable application-aware image** processing option. To learn more about VMware tools quiescence, see the Transaction-Consistent Backup section.
- If you are running pre-ESX 3.5 Update 2 hosts, consider enabling the **Safe snapshot removal** option. Because full image-level backup can take long time depending on the VM size, the VM snapshot can grow very large. When a large snapshot is removed on a VM with heavy disk I/O, a consolidation helper snapshot may grow large too, and will then require long time to be committed. While a helper snapshot is being committed into VM virtual disk files, VM remains completely "frozen", and depending on the consolidation helper snapshot size, the freeze time may be so long that some applications running on a VM would time out. To prevent such situation, Veeam Backup & Replication 5.0 offers a procedure of safe snapshot removal which includes creating an additional snapshot in cases when the "main" snapshot size is above the specified threshold. An additional snapshot is used to host writes while the "main" snapshot is being deleted. This ensures that a consolidation helper snapshot does not grow large. To use this option, select the **Safe removal for snapshots larger than … Mb** check box and specify a threshold for the size of a snapshot that should not be exceeded.
- Select the **Enable automatic backup integrity checks** check box if you want Veeam Backup & Replication 5.0 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 increments corrupted, too.
  A backup check is performed every time a job is started, and a full backup file is re-built to include new incremental changes. If the check determines a full backup file to be corrupted, a notification message will be displayed, prompting you to perform full backup anew. During such full backup, no integrity check will be performed.
- Select the **Run the following command** check box if you want to execute post-backup actions, for example, to launch a script recording the resulting backup file to tape. Use the **Browse…** button to select an executable file.
  You can select to execute post-backup actions after a number of backup cycles or on specific week days. If you select the **Run every… backup cycle** option, specify the number of a backup cycle after which the file should be executed. If you select the **Run on** selected days only option, click the **Days…** button and specify week days when actions should be performed.

Step 8. Enable Application-Aware Image Processing and Indexing

At the **Guest Processing** step of the wizard, you can enable guest file indexing and select to create a transactionally consistent backup.

If you want to create a transactionally consistent backup ensuring successful recovery of VM applications without any data loss, select the **Enable application-aware image processing (recommended)** check box.

If you want to index guest files in a VM you back up, select the **Enable guest file system indexing** check box. Veeam Backup & Replication will perform file indexing and enable you to perform fast and accurate search for VM guest OS files. To learn more about indexing, see the Searching for VM Guest Files section.

To coordinate proper indexing and VSS activities, Veeam Backup & Replication deploys a small agent inside a VM. The agent is installed only during VSS quiescence and indexing procedure and removed immediately after the processing is finished (depending on the selected option, during the backup job or after it is finished), thus producing low impact on VM performance and stability. In the **Guest OS credentials** section, specify guest operating system credentials for a target VM that are required to deploy the agent. Please note that the user name must be supplied in the *DOMAIN\USERNAME* format.

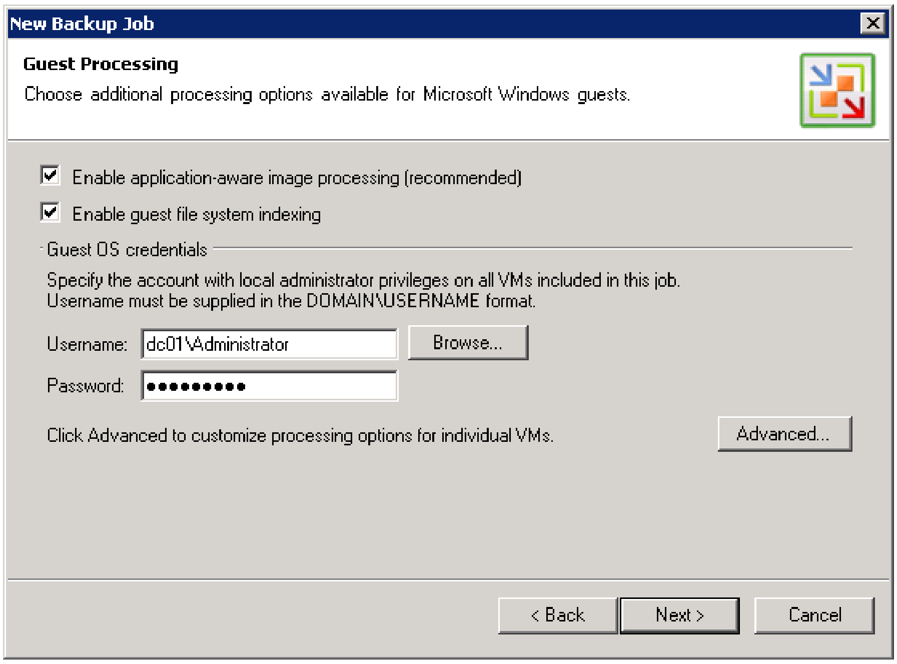

New Backup Job

Guest Processing

Choose additional processing options available for Microsoft Windows guests.

☑ Enable application-aware image processing (recommended)

☑ Enable guest file system indexing

Guest OS credentials

Specify the account with local administrator privileges on all VMs included in this job.
Username must be supplied in the DOMAIN\USERNAME format.

Username: dc01\Administrator

Password: ●●●●●●●●●●

Click Advanced to customize processing options for individual VMs.

Click the **Advanced…** button to specify advanced options for Veeam VSS and indexing processing.

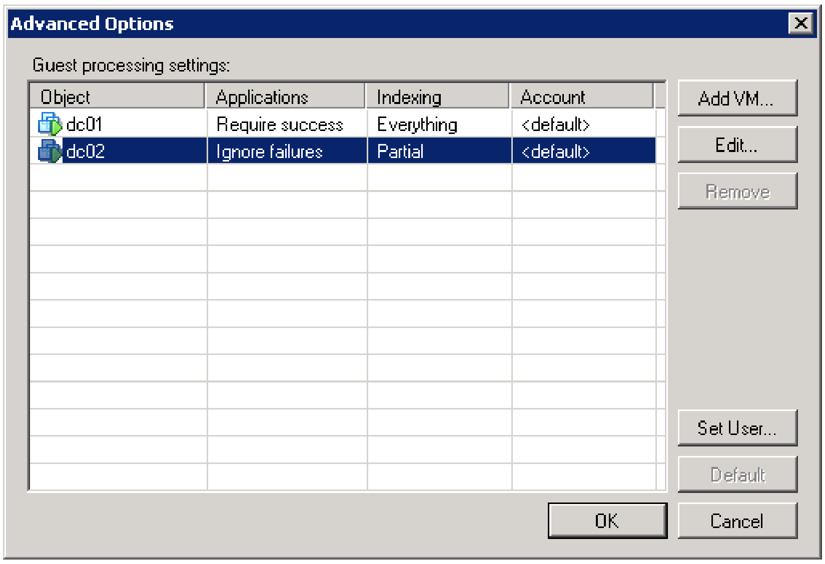

Advanced Options

Guest processing settings:

Object	Applications	Indexing	Account
dc01	Require success	Everything	<default>
dc02	Ignore failures	Partial	<default>

The **Advanced Options** window contains a list of VMs that will be processed with Veeam VSS and indexing tools. You can exclude specific VMs from processing or add them:

- To add a VM, click the **Add VM…** button and select a VM you want to process. The **Add objects** list contains only those VMs that you added to the backup job. To display all VMs in the virtual infrastructure hierarchy, select the **Show full hierarchy** check box.
- To exclude a VM, select it in the list and click the **Remove** button.

To provide granular quiesencing and indexing options for a VM, select it in the list and click the **Edit…** button.

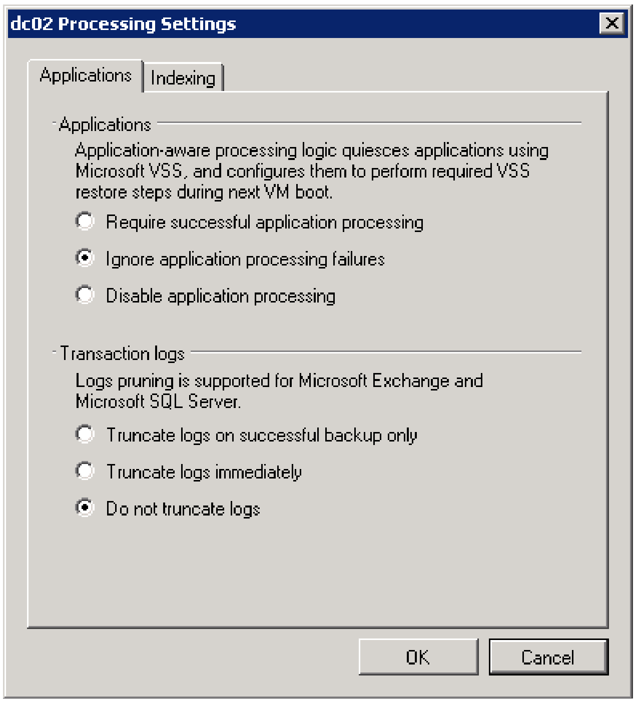

dc02 Processing Settings

Applications

Application-aware processing logic quiesces applications using Microsoft VSS, and configures them to perform required VSS restore steps during next VM boot.

○ Require successful application processing

● Ignore application processing failures

○ Disable application processing

Transaction logs

Logs pruning is supported for Microsoft Exchange and Microsoft SQL Server.

○ Truncate logs on successful backup only

○ Truncate logs immediately

● Do not truncate logs

In the **Applications** section on the **Applications** tab, specify the VSS behavior scenario:

- Select the **Require successful application processing** option if you want Veeam Backup & Replication to stop backup up a VM if any VSS errors occur.
- Select the **Ignore application processing failures** option if you want to continue backing up a VM even if VSS errors occur. This option is recommended to guarantee completion of the job. The created backup image will be not transactionally consistent, but crash consistent.

- Select the **Disable application processing** option if you do not want enable quiescencing for a VM.

Use the **Truncation logs** section to define the scenario of transaction log handing:

- Select the **Truncate logs on successful backup only** option if you want Veeam Backup & Replication to truncate logs only after the job is finished successfully. In this case, Veeam agent will wait for the backup to complete, and then truncate transaction logs. If the agent will not manage to truncate transaction logs for some reason, it will be remain in the VM guest OS till the next start of Veeam VSS.
- Select the **Truncate logs immediately** option if you want Veeam Backup & Replication to truncate logs in any case, no matter whether the job finishes successfully or fails.
- Select the **Do not truncate logs** option if you do not want Veeam Backup & Replication to truncate logs at all. This option is recommended if, together with Veeam Backup & Replication, you are using another backup tool to perform guest-level backup, and this tool maintains consistency of the database state. In such scenario, truncation of logs with Veeam Backup & Replication will break the guest-level backup chain and cause it to fall out of sync.

Click the **Indexing** tab to specify the indexing option for a VM. Please keep in mind that file indexing is supported for Windows-based VMs only.

- Select the **Disable indexing** option if you do not want to index guest OS files of a VM and enable the search option.
- Select the **Index everything** option if you want to index all guest OS files inside a VM.
- Select the **Index everything except** option if you want to index all guest OS files except those defined in the list. By default, system folders are excluded from indexing. You can add or delete folders to exclude using the **Add…** and **Remove** buttons on the right. You can use any system environment variables, for example: *%windir%*, *%Program Files%* and *%Temp%*.
- Use the **Index only following folders** option to select specific folders that you want to index. To form a list of folders, use the **Add…** and **Remove** buttons.

Step 9. Define Job Schedule

The **Job Schedule** step of the wizard allows you to choose to manually run the created job or schedule performing the backup job for a specific period of time — for example, the least busy hours to reduce impact on the VI environment.

To specify the job schedule, select the **Run the job automatically** check box. If this check box is not selected, the job is supposed to be run manually.

You can choose to perform the job at specific time on defined week days, monthly and with specific periodicity.

You can also select to back up a VM continuously. In this case, the next run of a backup job will be started once the previous one is complete, maintaining your backup always in the most recent state.

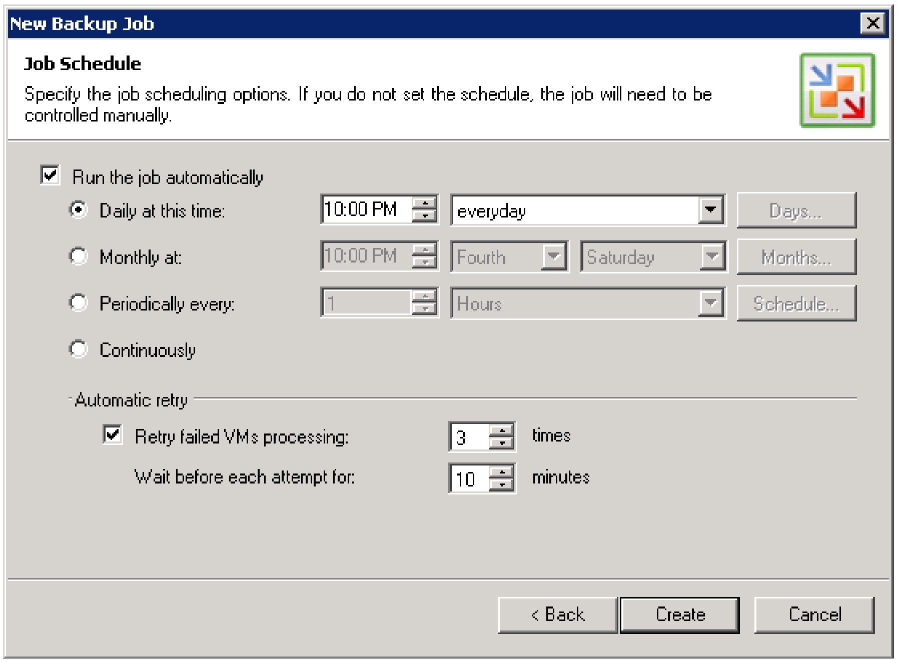

Job Schedule

Specify the job scheduling options. If you do not set the schedule, the job will need to be controlled manually.

☑ Run the job automatically

● Daily at this time: 10:00 PM everyday Days...

○ Monthly at: 10:00 PM Fourth Saturday Months...

○ Periodically every: 1 Hours Schedule...

○ Continuously

Automatic retry

☑ Retry failed VMs processing: 3 times

Wait before each attempt for: 10 minutes

In the **Automatic retry** section, select to repeat an attempt to run a backup job in case it fails for some reason. A repeatedly run job will include failed VMs only. Enter the number of attempts to run the job and define time spans between them. If you select continuous backup, Veeam Backup & Replication 5.0 will retry the job for the defined number of times without any time intervals between the job runs.

To learn more about job retries, see the Automatic Retry of Backup Jobs section.

Note: After you have created a scheduled job, you can temporarily disable it — hold it for some time without changing the set time schedule. Right-click a job in the list and select **Disable Job** from the shortcut menu. To enable the job schedule, right-click the job and deselect **Disable Job** in the shortcut menu.

Step 10. Finish Working with Wizard

After you have specified schedule settings, click **Create**. Select the **Run the job when I click Finish** check box if you want to start the created job right after you complete working with the wizard; then click **Finish**.

Creating a Replication Job

To replicate a VM, you should create a replication job by means of the **New Replication Job** wizard. You can perform the created job immediately, schedule or save it. This section will guide you through all steps of the wizard and provide explanation on offered options.

Before You Begin

- Prior to creating a VM replica, make sure you have enough free space on the destination disk. When a replication job runs for the first time, a full replica is created: the disk space required is equal to the actual size of the virtual machine. At all subsequent runs of the replication job, only incremental data will be saved.
  To learn how much disk space is available on storage devices used by a specific server, right–click a necessary server in the management tree, select the **Properties** command from the shortcut menu and click the **Populate** button. You will also be able to check disk space resources right from the wizard.
- Make sure all servers you want to work with are available in the management tree: you will not be able to add them once the **New Replication Job** wizard is launched.

Step 1. Launch the New Replication Job Wizard

To run the **New Replication Job** wizard:

- Click the **Replication** button on the toolbar.
- Select **Backup** > **Replication…** from the main menu.
- Click **Jobs** under the **Backup and Replication** node in the management tree, right-click anywhere on the blank area of the informational panel and select **Replication….**Click **Jobs** under the **Backup** node in the management tree, right-click anywhere on the blank area of the informational pane and select **Backup….**
- Right-click the **Replicas** node under **Backup and Replication** in the management tree and select **Replication...** from the shortcut menu.

Step 2. Specify Job Name and Description

At the first step of the wizard, enter a name and description of the created job. By default, the following description is initially provided for the created job: time at which the job was created and user who created the job.

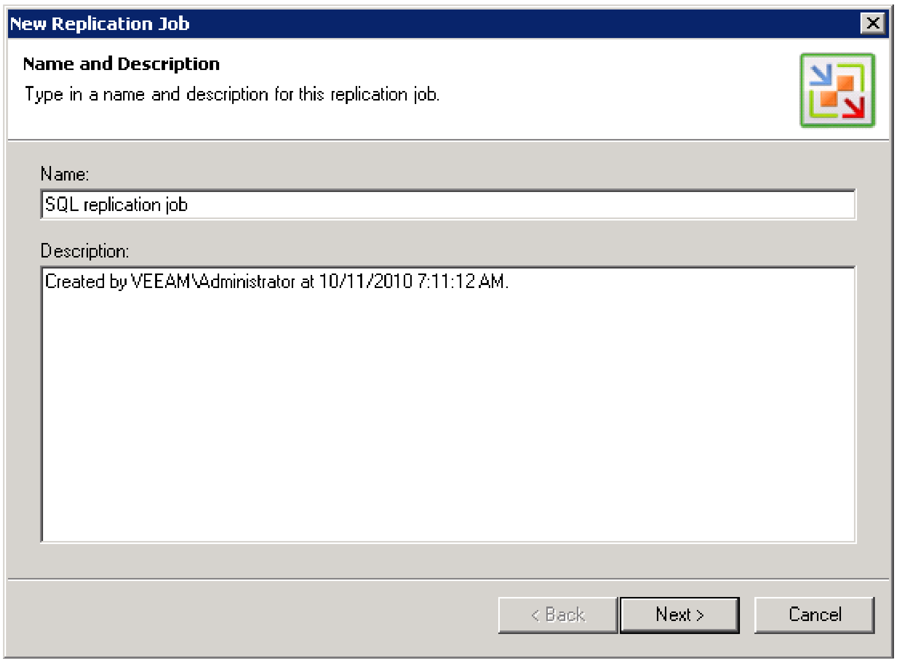

Name and Description

Type in a name and description for this replication job.

Name:
SQL replication job

Description:
Created by VEEAM\Administrator at 10/11/2010 7:11:12 AM.

Step 3. Select Replication Mode

You can replicate VMs in one of the three modes using VMware vStorage APIs — **Direct SAN access**, **Virtual Appliance** and **Network mode**.

- By default, if the **Direct SAN access** or **Virtual Appliance** mode is selected, Veeam Backup & Replication will automatically fail over to network data transfer in case the primary selected replication mode fails during the job run. To disable failover, click the **Advanced…** button and clear the **Failover to network mode if primary backup mode fails** check box.
- If the **Network** VMware vStorage APIs mode is selected, you can choose to transfer disks data over encrypted SSL connection. Click the **Advanced…** button and select the **Encrypt LAN traffic** check box. Use of encryption puts more stress on CPU of an ESX server, providing, however, secure data transfer.

You can also choose one of the legacy modes — **VCB-enabled backup** or **Network backup**. To enable legacy modes, select **Tools** > **Options…** from the main menu of Veeam Backup & Replication, click the **Advanced** tab and select the **Enable legacy processing modes** check box. Legacy modes may be used for ESX/ESXi servers earlier than 3.5; for ESX/ESXi servers 3.5 and higher it is recommended to use vStorage API replication modes.

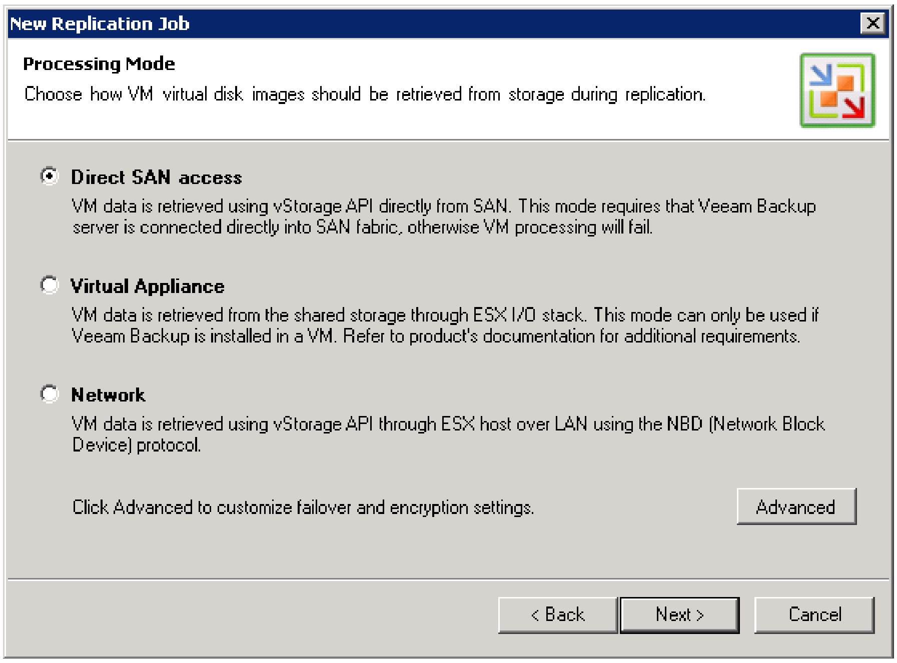
New Replication Job

Processing Mode

Choose how VM virtual disk images should be retrieved from storage during replication.

● Direct SAN access
VM data is retrieved using vStorage API directly from SAN. This mode requires that Veeam Backup server is connected directly into SAN fabric, otherwise VM processing will fail.

○ Virtual Appliance
VM data is retrieved from the shared storage through ESX I/O stack. This mode can only be used if Veeam Backup is installed in a VM. Refer to product's documentation for additional requirements.

○ Network
VM data is retrieved using vStorage API through ESX host over LAN using the NBD (Network Block Device) protocol.

Click Advanced to customize failover and encryption settings.

Step 4. Select Virtual Machines to Replicate

At this step, you should select an individual VM or a VM container you want to replicate. Jobs with VM containers are dynamic in their nature: if a new VM is added to the container after a replication job is created, the job will be automatically updated to include the added VM.

Click the **Add…** button to browse to VMs and VM containers that should be replicated. In the displayed VI tree, select a necessary object and click the **Add** button.

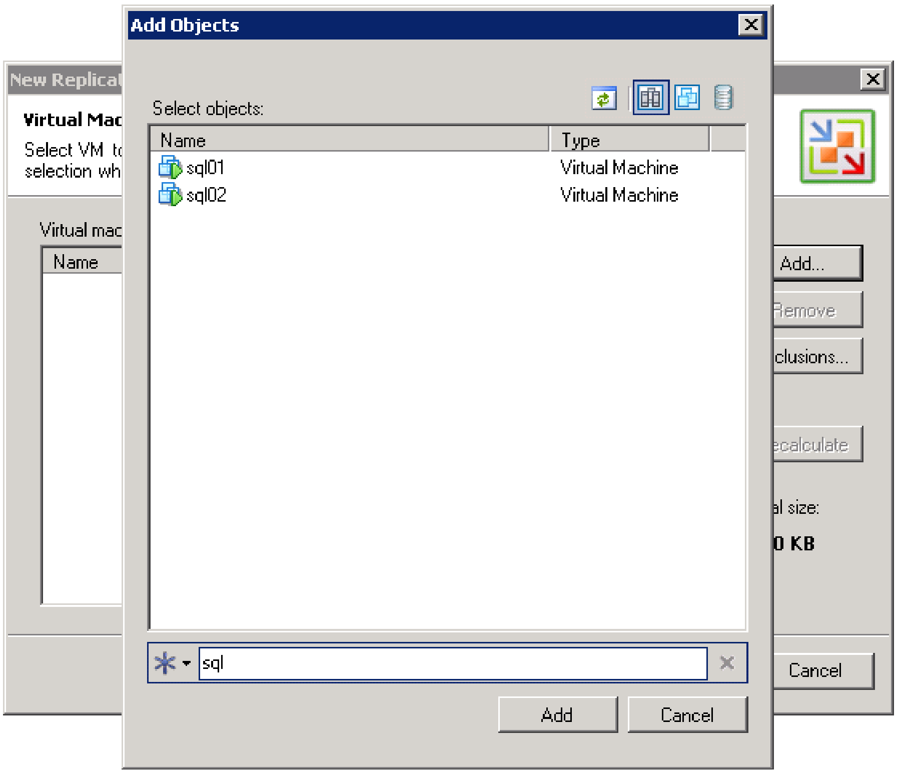
Add Objects

Select objects:

Name	Type
sql01	Virtual Machine
sql02	Virtual Machine

sql

To facilitate objects selection, you can:

- Switch between VI views: click the **Hosts and Clusters**, **VMs and Templates** or **Datastores and VMs** buttons at the top of the tree.

- Use a search field at the bottom of the window: click the button on the left of the field to select a necessary type of object that should be searched for (*Everything*, *Folder*, *Cluster*, *Host*, *Resource Pool*, *Virtual Application* or *VM*), enter an object's name or a part of it and click the **Start search** button on the right.

Note: Depending on the view you select, some VI objects may be not available: for example, if you select the VMs and Templates view, you will not be able to see and find resource pools.

To remove an object from the list, select it and click the **Remove** button on the right.

The initial size of VMs and VM containers added to a replication job is displayed in the **Size** column in the list. The total size of objects is displayed in the **Total size** field. Use the **Refresh** button to refresh the total size value after you add a new object to the job.

Step 5. Exclude Objects from Replication Job

After you have added VMs and VM containers to the list, you can specify which objects should be excluded from the replication job. Veeam Backup & Replication 5.0 allows excluding the following types of objects: VMs and VM templates from VM containers, as well as specific VM disks.
To select which objects should be excluded, click the **Exclusions…** button on the right.

- To exclude VMs from a VM container (for example, if you need to replicate the whole ESX server excluding several VMs running on this server), click the **VMs** tab. Click the **Add…** button on the right and select VMs that should be excluded. To display all hosts added to Veeam Backup & Replication 5.0, select the **Show full hierarchy** check box. To facilitate objects selection, you can switch between the **Hosts and Clusters**, **VMs and Templates** and **Datastores and VMs** views, and use the search field just as in the main window of the wizard.
- To select what VM disks you want to replicate, click the **Disks** tab, select a necessary VM in the list and click the **Edit…** button. If a VM is not in the list, you can add it by clicking the **Add...** button. You can choose to process all disks, 0:0 disks (typically, the system disks) or select custom disks.
  By default, the **Remove excluded disks from VM configuration** check box is selected, which means that Veeam Backup & Replication 5.0 will modify VMX file to remove disks you want to skip from VM configuration. When this option is used, you will be able to restore, replicate or copy VM to a location where excluded disks are not accessible with the original paths. If you do not use this option, you will have to manually edit VM configuration file to be able to power on a VM.

Note: Veeam Backup & Replication 5.0 automatically excludes VM log files from replicas to make replication process faster and reduce the size of the replica.

Step 6. Specify Replica Destination

At this step of the wizard, you should select destination for the created replica.

In the **Replica destination** section, select where the created replica should be located. Click the **Choose...** button to select a necessary host and storage. The displayed list will contain hosts that were added to Veeam Backup & Replication 5.0. The **Summary** section at the bottom of the window will display general information on a selected datastore.

Use the **Check Space** button to check how much free space is available on destination storage, and how much space you will require to store a full replica and its increments according to specified retention policy settings.

Beside storing a replica to a host, you can select to store an initial replica to a removable physical storage. Storing an initial replica to a removable storage may be useful if you want to replicate a VM to a remote site (for example, from one company affiliate to another) and need to minimize traffic over WAN.

Select the **Perform initial replication over this removable storage** check box and choose a necessary device from the list. Veeam Backup & Replication 5.0 will save a replica to the selected device and along with it will create a *README.txt* file with a path on the target host where a replica should be transferred (path you specified in the **Replica destination** section). When you transfer a replica to the specified location and run a replication job again, Veeam Backup & Replication 5.0 will store incremental changes next to this imported replica.

If you select a removable storage as a replica destination, make sure you have enough free space on your storage device.

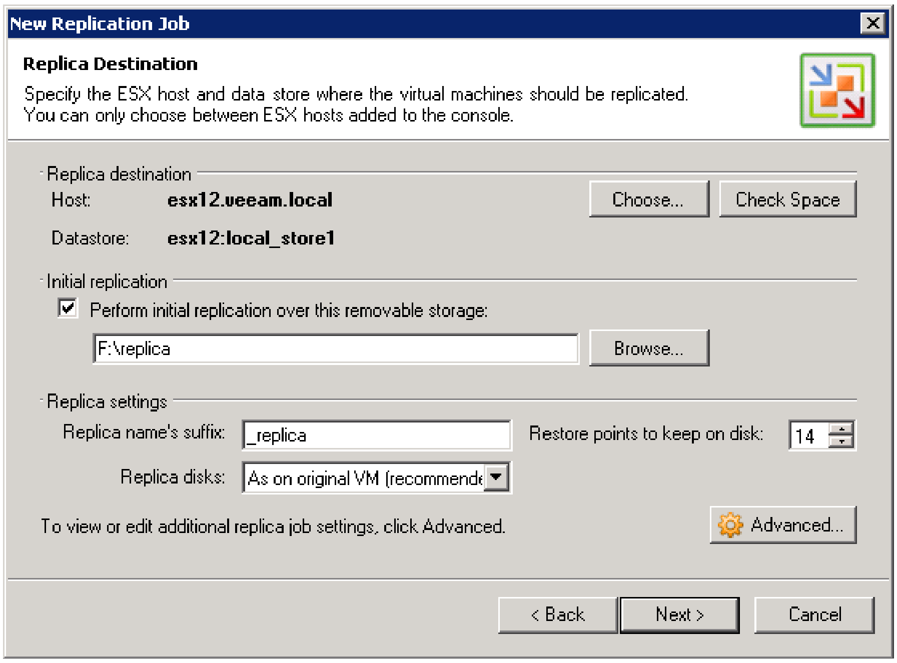

New Replication Job

Replica Destination

Specify the ESX host and data store where the virtual machines should be replicated.
You can only choose between ESX hosts added to the console.

Replica destination

Host: esx12.veeam.local

Datastore: esx12:local_store1

Initial replication

☑ Perform initial replication over this removable storage:

F:\replica

Replica settings

Replica name's suffix: _replica

Restore points to keep on disk: 14

Replica disks: As on original VM (recommende

To view or edit additional replica job settings, click Advanced.

In the **Replica's name suffix** field, enter a suffix that will be appended to a name of the virtual machine you are replicating. This name, with the suffix added, will be used to register the replicated virtual machine on the target server. Files of a replicated VM will be placed to the selected datastore in the */VeeamBackup/VMname(vm-ID)* folder.

From the **Replica disks** list, select the type of disks for a replicated VM. You can select to replicate a VM in its original state (recommended), or force all VM disks thick or thin. Please note that this option is available only for VMs using virtual hardware version 7 or later.

In the **Restore points to keep on disk** field, specify the number of restore points that should be maintained by the replication job. If this number is exceeded, the earliest restore point will be deleted. The number of restore points is a relative value and doesn't correspond to the number of days to store them.

Step 7. Specify Advanced Replica Settings

Click the **Advanced…** button to specify advanced options for the created replication job:

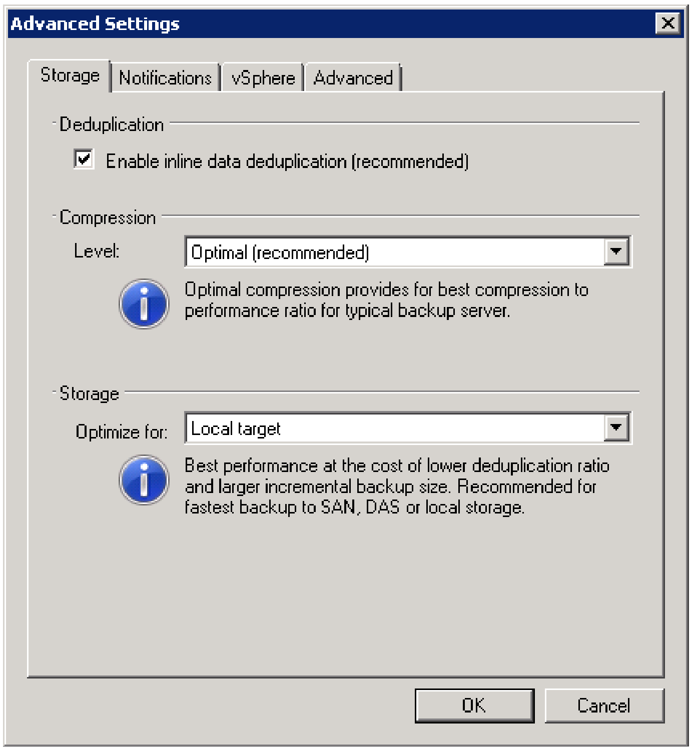

Deduplication

☑ Enable inline data deduplication (recommended)

Compression

Level: Optimal (recommended)

Optimal compression provides for best compression to performance ratio for typical backup server.

Storage

Optimize for: Local target

Best performance at the cost of lower deduplication ratio and larger incremental backup size. Recommended for fastest backup to SAN, DAS or local storage.

Storage settings

In the **Storage** section, select the type of replication target you are planning to use. Depending on the chosen option, Veeam Backup & Replication will use data blocks of different sizes to optimize the job performance:

- **Local target**. This option is recommended if you are planning to use SAN, DAS or local storage as a target. SAN identifies larger blocks of data and therefore can process larger quantities of data at a time. This option provides the fastest replication job performance but reduces the de-duplication ratio — the larger a data block is, the lower is the chance to find an identical block.
- **LAN target**. This option is recommended for NAS and on-site replication. It provides a better de-duplication ratio and reduces the size of an incremental replication file.
- **WAN target**. This option is recommended if you are planning to use WAN for offsite replication. Veeam Backup & Replication uses small data blocks, which results in the maximum de-duplication ratio and the smallest size of a replica file, allowing you to reduce the amount of traffic over the WAN link.

You can disable de-duplication by clearing the **Enable inline deduplication** check box. De-duplication provides a smaller size of a resulting replica file but may reduce the job performance.

Use the **Compression** tab to specify a compression level for the created replica: *None*, *Low*, *Optimal* or *Best*. To learn more about compression, see the De-duplication and Compression section.

Notification settings

- Select the **Send e-mail notifications to the following recipients** check box if you want to receive notifications by e–mail in case of job failure or success. In the field below, specify a recipient's e-mail address. You can enter several addresses separated by a semicolon.
  E–mail notifications will be sent only if you have selected the **Enable email notification** check box in the **Options** window and specified e–mail notification settings (select **Tools > Options...** from the main menu). To learn more, see the Specifying Notification Settings section.

- Select the **Enable SNMP notification for this job** check box if you want to receive SNMP traps when a job is completed and a backup is created. SNMP traps will be sent if you configure SNMP settings in Veeam Backup & Replication and on the recipient's computer. To learn more, see the Specifying SNMP Settings section.

vSphere settings

Use the **vSphere tab** to specify if vSphere changed block tracking should be used. By default, this option is selected. If you want force using changed block tracking for VMs for which changed block tracking is disabled on the ESX server, select the **Enable changed block tracking for all processed VMs** check box. Please note that you can use this option only for VMs using virtual hardware version 7 or later.

Advanced settings

- The **Enable VMware tools quiescence** option enables freezing of the file-system for proper snapshot creation. With this option enabled, creation of a snapshot is performed with the help of the sync driver responsible for holding incoming I/O and flushing all dirty data to a disk, thus making the file systems consistent.
  It is strongly recommended to leave this option disabled if you are replicating Windows systems that support Windows VSS — for these systems, it is recommended to use the **Enable application-aware image** processing option. To learn more about VMware tools quiescence, see the Transaction-Consistent Backup section.
- If you are running pre-ESX 3.5 Update 2 hosts, consider enabling the **Safe snapshot removal** option. Because full image-level replication can take long time depending on the VM size, the VM snapshot can grow very large. When a large snapshot is removed on a VM with heavy disk I/O, a consolidation helper snapshot may grow large too, and will then require long time to be committed. While a helper snapshot is being committed into VM virtual disk files, VM remains completely "frozen", and depending on the consolidation helper snapshot size, the freeze time may be so long that some applications running on a VM would time out. To prevent such situation, Veeam Backup & Replication 5.0 offers a procedure of safe snapshot removal which includes creating an additional snapshot in cases when the "main" snapshot size is above the specified threshold. An additional snapshot is used to host writes while the "main" snapshot is being deleted. This ensures that a consolidation helper snapshot does not grow large.
  To use this option, select the **Safe removal for snapshots larger than ... Mb** check box and specify a threshold for the size of a snapshot that should not be exceeded.
- Select the **Enable automatic replication integrity checks** check box if you want Veeam Backup & Replication 5.0 to periodically check a full replica. An automatic replication check allows you to verify integrity of a replica and avoid a situation when a replica is corrupted, making all further increments corrupted, too.
  A replication check is performed every time a job is started and a replica is re-built to include new incremental changes. If the check determines a replica to be corrupted, a notification message will be displayed, prompting you to perform full replication anew. During such full replication, no integrity check will be performed.
- Select the **Run the following command** check box if you want to execute post-replication actions, for example, to launch a script recording the resulting replica to tape. Use the **Browse…** button to select an executable file.
  You can select to execute post-replication actions after a number of replication cycles or on specific week days. If you select the **Run every... replication cycle** option, specify the number of a replication cycle after which the file should be executed. If you select the **Run on** selected days only option, click the **Days…** button and specify week days when actions should be performed.

Step 8. Enable Application-Aware Image Processing

If you want to create a transactionally consistent replica ensuring successful recovery of VM applications without any data loss, select the **Enable application-aware image processing (recommended)** check box.

To coordinate proper VSS activities, Veeam Backup & Replication installs a small agent inside a VM. The agent is installed only during VSS quiescence procedure and removed immediately after the processing is finished (depending on the selected option, during the replication job or after it is finished), thus producing low impact on VM performance and stability.

In the **Guest OS credentials** section, specify guest operating system credentials for a target VM. Please note that the user name must be supplied in the *DOMAIN\USERNAME* format.

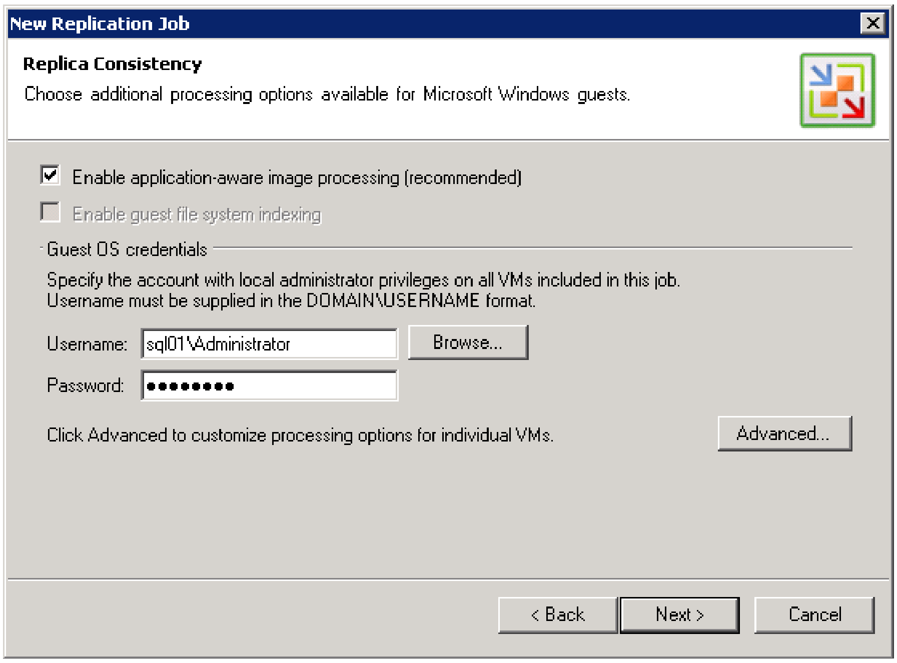
New Replication Job

Replica Consistency

Choose additional processing options available for Microsoft Windows guests.

☑ Enable application-aware image processing (recommended)

☐ Enable guest file system indexing

Guest OS credentials

Specify the account with local administrator privileges on all VMs included in this job.
Username must be supplied in the DOMAIN\USERNAME format.

Username: sql01\Administrator

Password: ••••••••

Click Advanced to customize processing options for individual VMs.

Click the **Advanced…** button to specify advanced option for Veeam VSS processing.

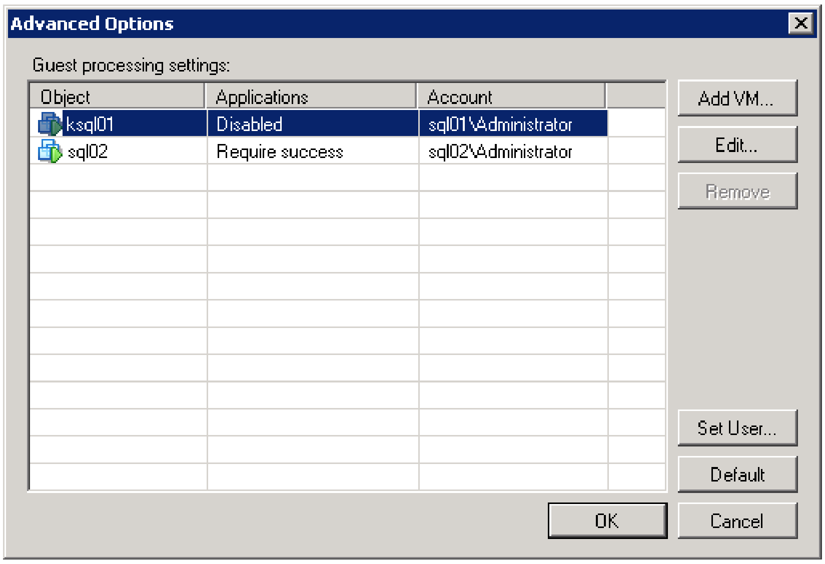
Advanced Options

Guest processing settings:

Object	Applications	Account
ksql01	Disabled	sql01\Administrator
sql02	Require success	sql02\Administrator

The **Advanced Options** window contains a list of VMs that will be processed with Veeam VSS. You can exclude specific VMs from processing or add them:

- To exclude a VM, select it in the list and click the **Remove** button.
- To add a VM, click the **Add VM…** button and select a VM you want to process. The **Add objects** list contains only those VMs that you added to the replication job. To display all VMs in the virtual infrastructure hierarchy, select the **Show full hierarchy** check box.

To provide granular quiesencing options for a VM, select it in the list and click the **Edit…** button.

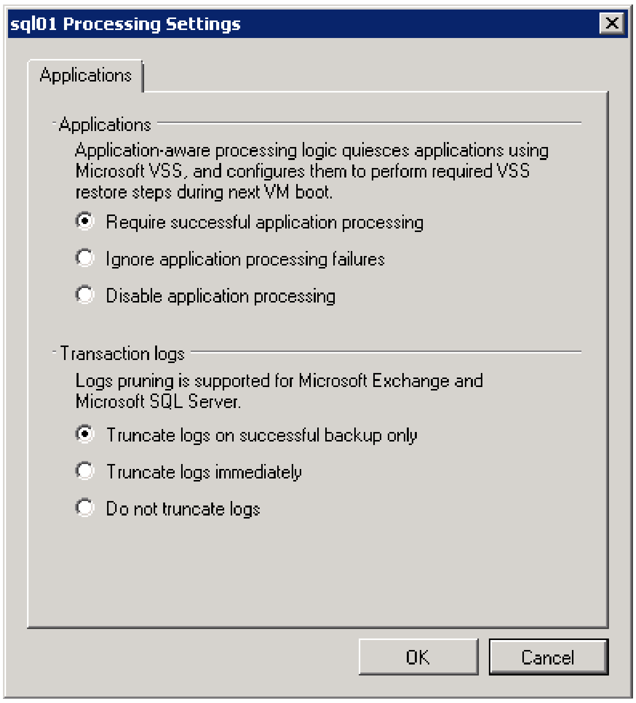

sql01 Processing Settings

Applications

Applications

Application-aware processing logic quiesces applications using Microsoft VSS, and configures them to perform required VSS restore steps during next VM boot.

● Require successful application processing

○ Ignore application processing failures

○ Disable application processing

Transaction logs

Logs pruning is supported for Microsoft Exchange and Microsoft SQL Server.

● Truncate logs on successful backup only

○ Truncate logs immediately

○ Do not truncate logs

In the **Applications** section on the **Applications** tab, specify the VSS behavior scenario:

- Select the **Require successful application processing** option if you want Veeam Backup & Replication to stop replicating a VM if any VSS errors occur.
- Select the **Ignore application processing failures** option if you want to continue replicating a VM even if VSS errors occur. This option is recommended to guarantee completion of the job. The created replica will be not transactionally consistent, but crash consistent.
- Select the **Disable application processing** option if you do not want enable quiescencing for a VM.

Use the **Truncation logs** section to define the scenario of transaction log handing:

- Select the **Truncate logs on successful backup only** option if you want Veeam Backup & Replication to truncate logs only after the job is finished successfully. In this case, Veeam agent will wait for the replication job to complete, and then truncate transaction logs. If the agent will not manage to truncate transaction logs for some reason, it will be remain in the VM guest OS till the next start of Veeam VSS.
- Select the **Truncate logs immediately** option if you want Veeam Backup & Replication to truncate logs in any case, no matter whether the job finishes successfully or fails.
- Select the **Do not truncate logs** option if you do not want Veeam Backup & Replication to truncate logs at all. This option is recommended if, together with Veeam Backup & Replication, you are using another tool to perform guest-level replication, and this tool maintains consistency of the database state. In such scenario, truncation of logs with Veeam Backup & Replication will break the guest-level replication chain and cause it to fall out of sync.

Step 9. Define Job Schedule

The **Job Schedule** step of the wizard allows you to choose to manually run the created job or schedule performing the replication job for a specific period of time — for example, the least busy hours to reduce impact on the VI environment.

To specify the job schedule, select the **Run the job automatically** check box. If this check box is not selected, the job is supposed to be run manually.

You can choose to perform the job at specific time on defined week days, monthly and with specific periodicity.

You can also select to replicate a VM continuously. In this case, the next run of a replication job will be started once the previous one is complete, maintaining your replica always in the most recent state.

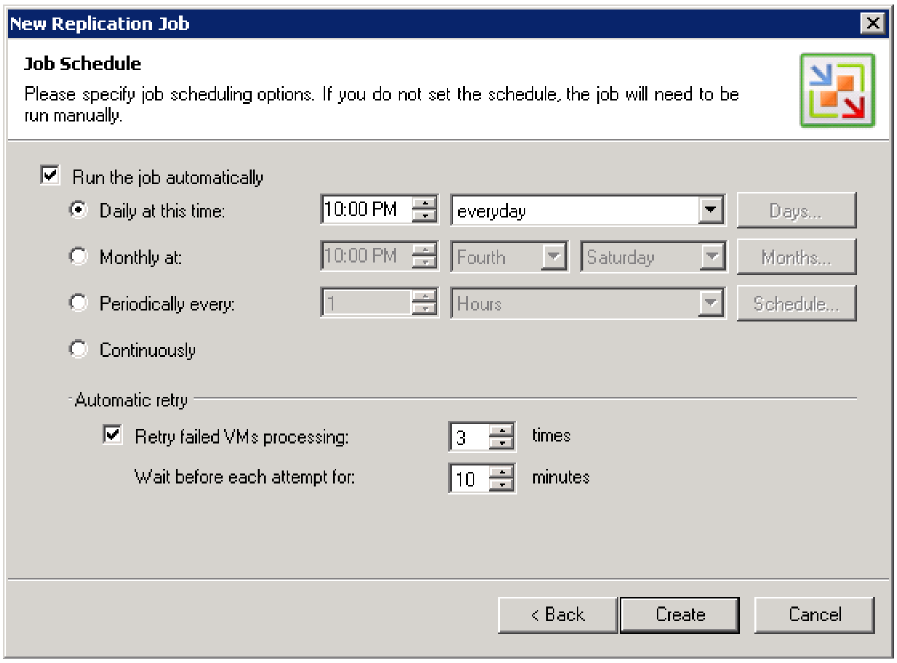

New Replication Job

Job Schedule

Please specify job scheduling options. If you do not set the schedule, the job will need to be run manually.

☑ Run the job automatically

● Daily at this time: 10:00 PM everyday

○ Monthly at: 10:00 PM Fourth Saturday

○ Periodically every: 1 Hours

○ Continuously

Automatic retry

☑ Retry failed VMs processing: 3 times

Wait before each attempt for: 10 minutes

In the **Automatic retry** section, select to repeat an attempt to run a replication job in case it fails for some reason. A repeatedly run job will include failed VMs only. Enter the number of attempts to run the job and define time spans between them. If you select continuous replication, Veeam Backup & Replication 5.0 will retry the job for the defined number of times without any time intervals between the job runs.

To learn more about job retries, see the Automatic Retry of Backup Jobs section.

Note: After you have created a scheduled job, you can temporarily disable it — hold it for some time without changing the set time schedule. Right-click a job in the list and select **Disable Job** from the shortcut menu. To enable the job schedule, right-click the job and deselect **Disable Job** in the shortcut menu.

Step 10. Finish Working with Wizard

After you have specified schedule settings, click **Create**. Select the **Run the job when I click Finish** check box if you want to start the created job right after you complete working with the wizard; then click **Finish**.

Creating a VM Copy Job

With a VM copy job, you can create a fully-functioning copy of a VM (both stopped and running) that will require no manual editing and adjustments. VM copying can be helpful if you want to move your datacenter, mirror your production environment to test lab storage and so on. Just as backup and replication jobs, a VM copy job can be performed in the VMware vStorage API and Network modes, supports VSS options, and can be run on demand or scheduled.

This section will guide you through all steps of the VM Copy wizard and provide explanation on offered options.

Before You Begin

- Prior to creating a VM copy, make sure you have enough free space on the destination disk. To learn how much disk space is available on storage devices used by a specific server, right–click a necessary server in the management tree, select the **Properties** command from the shortcut menu and click the **Populate** button. You will also be able to check disk space resources right from the wizard.

- Make sure all servers you want to work with are available in the management tree: you will not be able to add them once the **New Virtual Machine Copy Job** wizard is launched.

Step 1. Launch the New Virtual Machine Copy Job Wizard

To run the **New Virtual Machine Copy Job** wizard:

- Click the **VM Copy** button on the toolbar.
- Select **Backup > VM Copy…** from the main menu.
- Click **Jobs** under the **Backup & Replication** node in the management tree, right-click anywhere on the blank area of the informational pane and select **VM Copy...**.

Step 2. Specify Job Name and Description

At the first step of the wizard, enter a name and description of the created job. By default, the following description is initially provided for the created job: time at which the job was created and user who created the job.

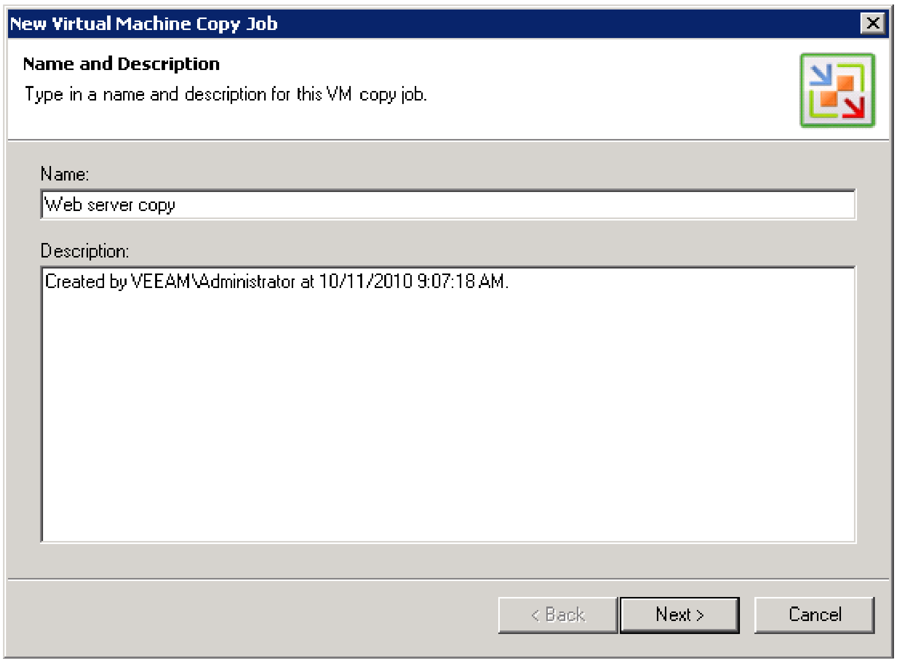

New Virtual Machine Copy Job

Name and Description

Type in a name and description for this VM copy job.

Name: Web server copy

Description: Created by VEEAM\Administrator at 10/11/2010 9:07:18 AM.

Step 3. Select VM Copy Mode

You can copy VMs in one of the three modes using VMware vStorage APIs — **Direct SAN access**, **Virtual Appliance** and **Network mode**.

- By default, if the **Direct SAN access** or **Virtual Appliance** mode is selected, Veeam Backup & Replication will automatically fail over to network data transfer in case the primary selected copy mode fails during the job run. To disable failover, click the **Advanced…** button and clear the **Failover to network mode if primary backup mode fails** check box.
- If the **Network** VMware vStorage APIs mode is selected, you can choose to transfer disks data over encrypted SSL connection. Click the **Advanced…** button and select the **Encrypt LAN traffic** check box. Use of encryption puts more stress on CPU of an ESX server, providing, however, secure data transfer.

You can also choose one of the legacy modes — **VCB-enabled backup** or **Network backup**. To enable legacy modes, select **Tools > Options…** from the main menu of Veeam Backup &

Replication, click the **Advanced** tab and select the **Enable legacy processing modes** check box. Legacy modes may be used for ESX/ESXi servers earlier than 3.5; for ESX/ESXi servers 3.5 and higher it is recommended to use VMware vStorage API copy modes.

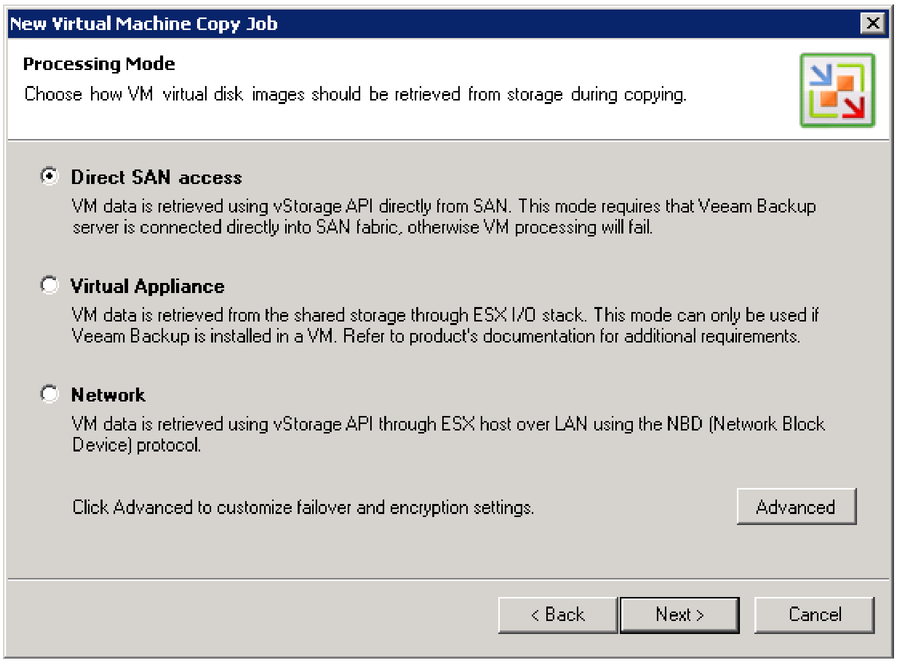

New Virtual Machine Copy Job

Processing Mode

Choose how VM virtual disk images should be retrieved from storage during copying.

● **Direct SAN access**
VM data is retrieved using vStorage API directly from SAN. This mode requires that Veeam Backup server is connected directly into SAN fabric, otherwise VM processing will fail.

○ **Virtual Appliance**
VM data is retrieved from the shared storage through ESX I/O stack. This mode can only be used if Veeam Backup is installed in a VM. Refer to product's documentation for additional requirements.

○ **Network**
VM data is retrieved using vStorage API through ESX host over LAN using the NBD (Network Block Device) protocol.

Click Advanced to customize failover and encryption settings.

Step 4. Select Virtual Machines to Copy

At this step, you should select an individual VM or a VM container you want to copy. Jobs with VM containers are dynamic in their nature: if a new VM is added to the container after a copy job is created, the job will be automatically updated to include the added VM.

Click the **Add…** button to browse to VMs and VM containers that should be copied. In the displayed VI tree, select a necessary object and click the **Add** button.

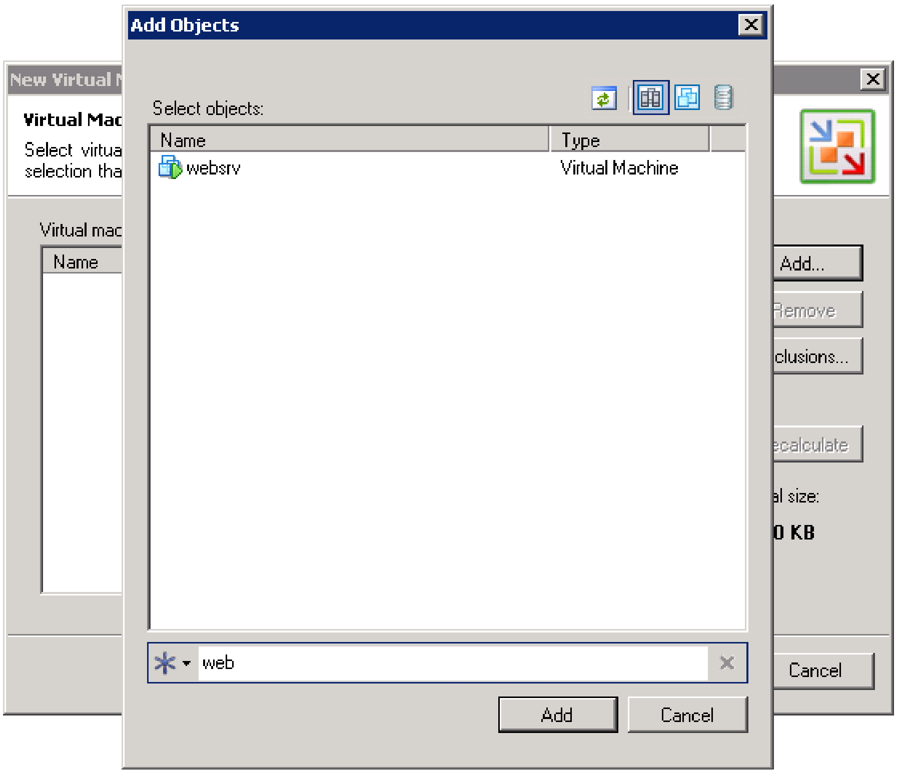

Add Objects

Select objects:

Name	Type
websrv	Virtual Machine

web

To facilitate objects selection, you can:

- Switch between VI views: click the **Hosts and Clusters**, **VMs and Templates** or **Datastores and VMs** buttons at the top of the tree.
- Use a search field at the bottom of the window: click the button on the left of the field to select a necessary type of object that should be searched for (*Everything*, *Folder*, *Cluster*, *Host*, *Resource Pool*, *Virtual Application* or *VM*), enter an object's name or a part of it and click the **Start search** button on the right.

Note: Depending on the view you select, some VI objects may be not available: for example, if you select the VMs and Templates view, you will not be able to see and find resource pools.

To remove an object from the list, select it and click the **Remove** button on the right.

The initial size of VMs and VM containers added to a copy job is displayed in the **Size** column in the list. The total size of objects is displayed in the **Total size** field. Use the **Refresh** button to refresh the total size value after you add a new object to the job.

Step 5. Exclude Objects from VM Copy Job

After you have added VMs and VM containers to the list, you can specify which objects should be excluded from the VM copy job. Veeam Backup & Replication 5.0 allows excluding the following types of objects: VMs and VM templates from VM containers, as well as specific VM disks.

To select which objects should be excluded, click the **Exclusions…** button on the right.

- To exclude VMs from a VM container (for example, if you need to copy the whole ESX server excluding several VMs running on this server), click the **VMs** tab. Click the **Add…** button on the right and select VMs that should be excluded. To display all hosts added to Veeam Backup & Replication 5.0, select the **Show full hierarchy** check box. To facilitate objects selection, you can switch between the **Hosts and Clusters**, **VMs and Templates** and **Datastores and VMs** views, and use the search field just as in the main window of the wizard.
- To select what VM disks you want to copy, click the **Disks** tab, select a necessary VM in the list and click the **Edit…** button. If a VM is not in the list, you can add it by clicking the **Add...** button. You can choose to process all disks, 0:0 disks (typically, the system disks) or select custom disks.
  If you select the **Remove excluded disks from VM configuration** check box, Veeam Backup & Replication 5.0 will modify VMX file to remove disks you want to skip from VM configuration. If this option is used, you will be able to restore, replicate or copy VM to a location where excluded disks are not accessible with the original paths. If you do not use this option, you will have to manually edit VM configuration file to be able to power on a VM.
- If you select to use the **Network** backup mode, you can copy VM templates together with VMs. Click the **VM Templates** tab. By default, the **Backup VM templates** check box is selected. Clear it if you do not want to copy VM templates. The **Exclude templates from incremental backup** option allows you to process VM templates into a with a full copy job only.

Note: Veeam Backup & Replication automatically excludes VM log files from a copy to make copying process faster and reduce the size of the VM copy.

Step 6. Specify Copy Destination

At this step of the wizard, you should select destination for the created VM copy.

From the **Destination** list, select a host where the created copy should be stored. The list contains hosts that were added to Veeam Backup & Replication 5.0. You can store a copy to a local host, network shared folders, and hosts added to the Veeam Backup & Replication 5.0.

Use the **Host Properties...** button to view available disk resources, specify SSH and SOAP connection and data transfer information.

In the **Path to folder** field, specify a folder where the created copy should be stored. Use the **Check Space** button to check how much free space is available on copy destination.

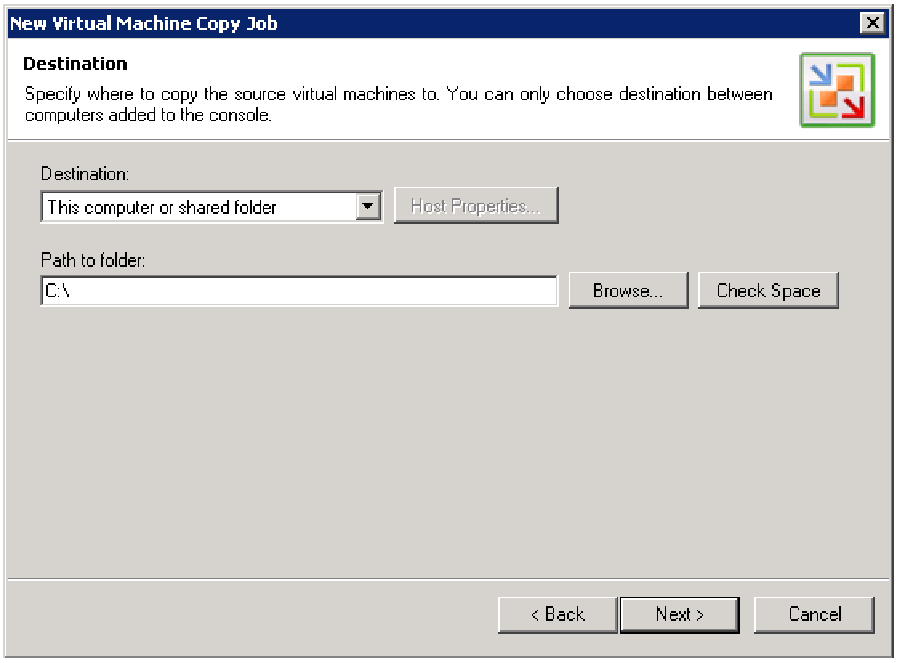

New Virtual Machine Copy Job

Destination

Specify where to copy the source virtual machines to. You can only choose destination between computers added to the console.

Destination: This computer or shared folder

Path to folder: C:\

Step 7. Enable Application-Aware Image Processing

If you want to create a transactionally consistent VM copy ensuring successful recovery of VM applications without any data loss, select the **Enable application-aware image processing (recommended)** check box.

To coordinate proper VSS activities, Veeam Backup & Replication installs a small agent inside a VM. The agent is installed only during VSS quiescence procedure and removed immediately after the processing is finished (depending on the selected option, during the copy job or after it is finished), thus producing low impact on VM performance and stability.

In the **Guest OS credentials** section, specify guest operating system credentials for a target VM. Please note that the user name must be supplied in the *DOMAIN\USERNAME* format.

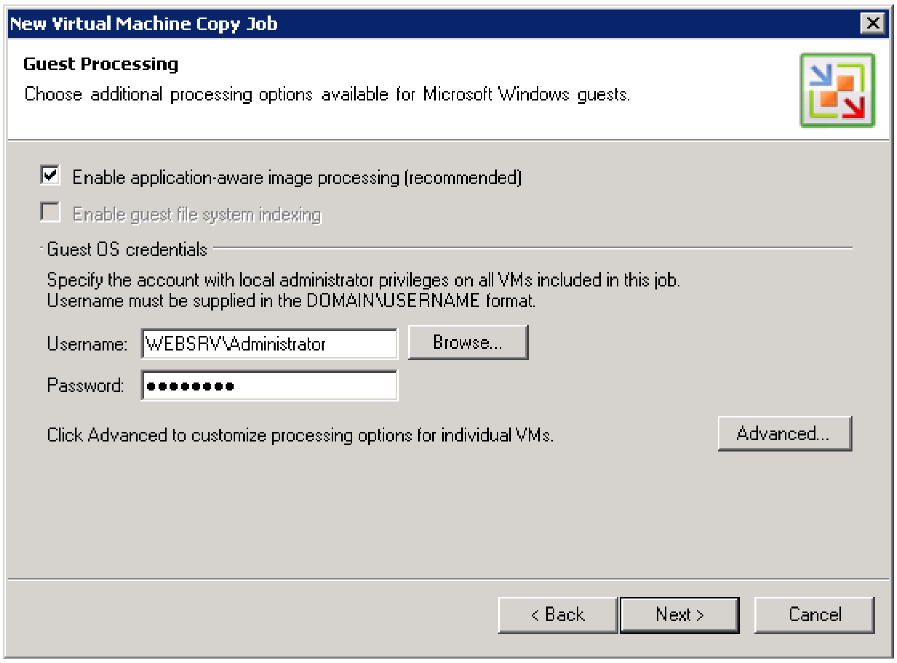

New Virtual Machine Copy Job

Guest Processing

Choose additional processing options available for Microsoft Windows guests.

☑ Enable application-aware image processing (recommended)

☐ Enable guest file system indexing

Guest OS credentials

Specify the account with local administrator privileges on all VMs included in this job. Username must be supplied in the DOMAIN\USERNAME format.

Username: WEBSRV\Administrator

Password: ●●●●●●●●●

Click Advanced to customize processing options for individual VMs.

Click the **Advanced…** button to specify advanced option for Veeam VSS processing.

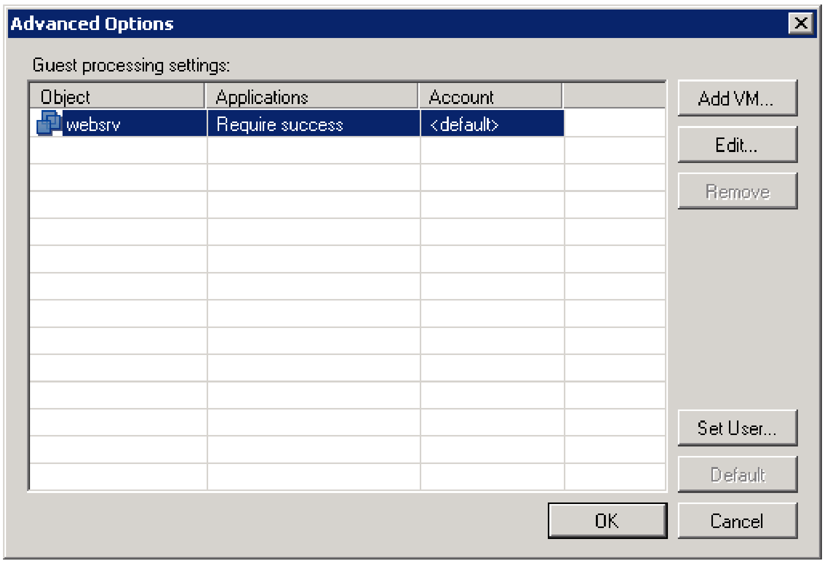

Advanced Options

Guest processing settings:

Object	Applications	Account
websrv	Require success	<default>

The **Advanced Options** window contains a list of VMs that will be processed with Veeam VSS. You can exclude specific VMs from processing or add them:

- To exclude a VM, select it in the list and click the **Remove** button.
- To add a VM, click the **Add VM…** button and select a VM you want to process. The **Add objects** list contains only those VMs that you added to the replication job. To display all VMs, select the **Show full hierarchy** check box.

To provide granular quiesencing options for a VM, select it in the list and click the **Edit…** button.

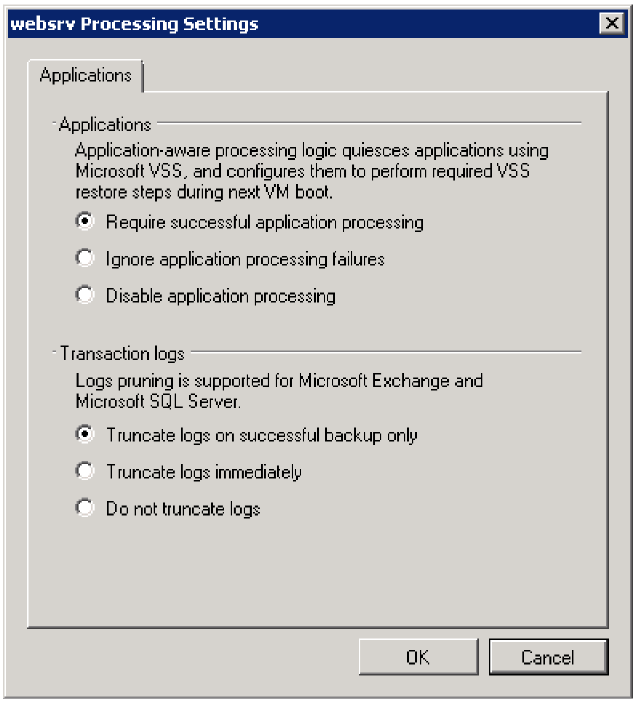

websrv Processing Settings

Applications

Application-aware processing logic quiesces applications using Microsoft VSS, and configures them to perform required VSS restore steps during next VM boot.

● Require successful application processing

○ Ignore application processing failures

○ Disable application processing

Transaction logs

Logs pruning is supported for Microsoft Exchange and Microsoft SQL Server.

● Truncate logs on successful backup only

○ Truncate logs immediately

○ Do not truncate logs

In the **Applications** section on the **Applications** tab, specify the VSS behavior scenario:

- Select the **Require successful application processing** option if you want Veeam Backup & Replication to stop copying a VM if any VSS errors occur. In this case, Veeam agent will wait for the copy job to complete, and then truncate transaction logs. If the agent will not manage to truncate transaction logs for some reason, it will be remain in the VM guest OS till the next start of Veeam VSS.
- Select the **Require successful application processing** option if you want to continue copying a VM even if VSS errors occur. This option is recommended to guarantee

completion of the job. The created copy will be not transactionally consistent, but crash consistent.

- Select the **Disable application processing** option if you do not want enable quiescencing for a VM.

If you are copying VMs running database systems that use transaction logs, you can select to truncate transaction logs after the job so that they don't overflow the storage space. Use the **Truncation logs** section to define the scenario of transaction log handing:

- Select the **Truncate logs on successful backup only** option if you want Veeam Backup & Replication to truncate logs only after the job is finished successfully.
- Select the **Truncate logs immediately** option if you want Veeam Backup & Replication to truncate logs in any case, no matter whether the job finishes successfully or fails.
- Select the **Do not truncate logs** option if you do not want Veeam Backup & Replication to truncate logs at all. This option is recommended if, together with Veeam Backup & Replication, you are using another tool to perform guest-level copy, and this tool maintains consistency of the database state. In such scenario, truncation of logs with Veeam Backup & Replication will break the guest-level copy chain and cause it to fall out of sync.

Step 8. Define Job Schedule

The **Job Schedule** step of the wizard allows you to choose to manually run the created job or schedule performing the copy job for a specific period of time — for example, the least busy hours to reduce impact on the VI environment.

To specify the job schedule, select the **Run the job automatically** check box. If this check box is not selected, the job is supposed to be run manually.

You can choose to perform the job at specific time on defined week days, monthly and with specific periodicity.

You can also select to copy VM data continuously. In this case, the next run of a copy job will be started once the previous one is complete, maintaining your VM always in the most recent state.

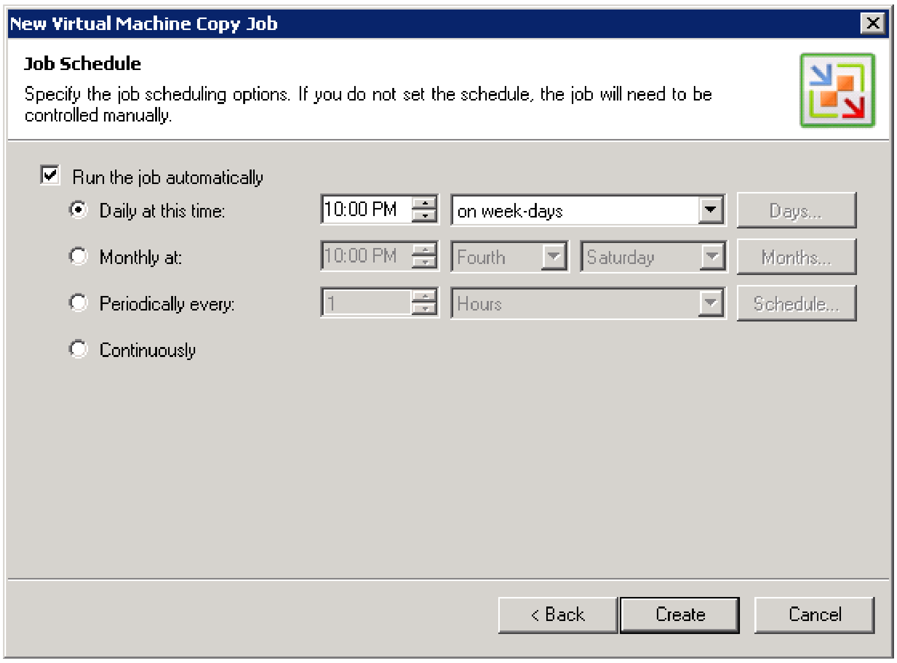

New Virtual Machine Copy Job

Job Schedule

Specify the job scheduling options. If you do not set the schedule, the job will need to be controlled manually.

☑ Run the job automatically

● Daily at this time: 10:00 PM, on week-days

○ Monthly at: 10:00 PM, Fourth, Saturday

○ Periodically every: 1 Hours

○ Continuously

Note: After you have created a scheduled job, you can temporarily disable it — hold it for some time without changing the set time schedule. Right-click a job in the list and select **Disable Job** from the shortcut menu. To enable the job schedule, right-click the job and deselect **Disable Job** in the shortcut menu.

Step 9. Finish Working with Wizard

After you have specified schedule settings, click **Create**. Select the **Run the job when I click Finish** check box if you want to start the created job right after you complete working with the wizard; then click **Finish**.

Performing Restore

Veeam Backup & Replication 5.0 allows you to instantly restore entire virtual machines and specific VM files from backups, and restore individual VM guest OS files and folders from backups and replicas that have been successfully created.

Performing Instant VM Recovery

With Veeam Backup & Replication, you can immediately recover an entire VM from the backup. Instant VM recovery accelerates VM restore, allowing you to improve recovery time objectives and decrease downtime of production VMs.

When instant VM recovery is performed, Veeam Backup & Replication 5.0 runs a VM from a backup file that resides on a regular backup storage, so you do not have to extract the VM from the backup and move it to the production storage. Similar to the SureBackup recovery verification technology, when instant VM recovery is performed, Veeam Backup & Replication mounts a VM directly from a compressed backup file on a selected ESX host. The archived image of a VM remains in a read-only state to avoid unexpected modifications. All changes to a virtual disk that take place while a VM is running are logged to an auxiliary file that resides on the Veeam Backup server or on a datastore. These changes are discarded as soon as a restored VM is removed.

Once the VM is started from the backup, you can move it to your production storage using Storage vMotion and cold migration to finalize recovery. Alternatively, you can replicate a restored VM with Veeam Backup & Replication and then fail over to the created replica during the next maintenance window. Beside disaster recovery matters, instant VM recovery can also be used for testing purposes to make sure VM guest OS and applications are functioning properly.

Instant VM recovery is wizard-driven. To launch the wizard, click the **Restore** button on the toolbar or select **Backup** > **Restore…** from the main menu. Then, select **Instant VM recovery.**

Alternatively, you can right-click **Instant recovery** in the management tree and select **Run VM from backup…**

Step 1. Select a Virtual Machine

Select a necessary VM in the list of available backup jobs. You can instantly recover a VM that has been successfully created at least once.

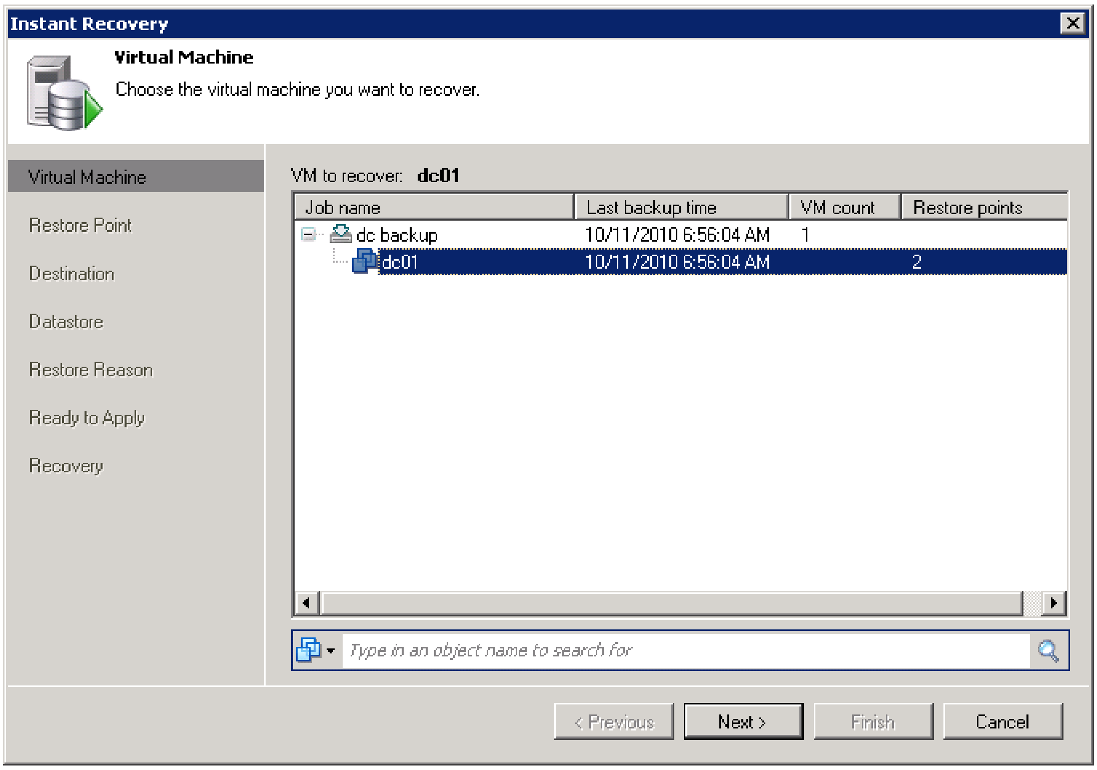

Virtual Machine

Choose the virtual machine you want to recover.

VM to recover: **dc01**

Job name	Last backup time	VM count	Restore points
dc backup	10/11/2010 6:56:04 AM	1	
dc01	10/11/2010 6:56:04 AM		2

Type in an object name to search for

Tip: To quickly find VMs in jobs, use the search field at the bottom of the window.

Step 2. Select a Restore Point

Select a necessary restore point for the virtual machine.

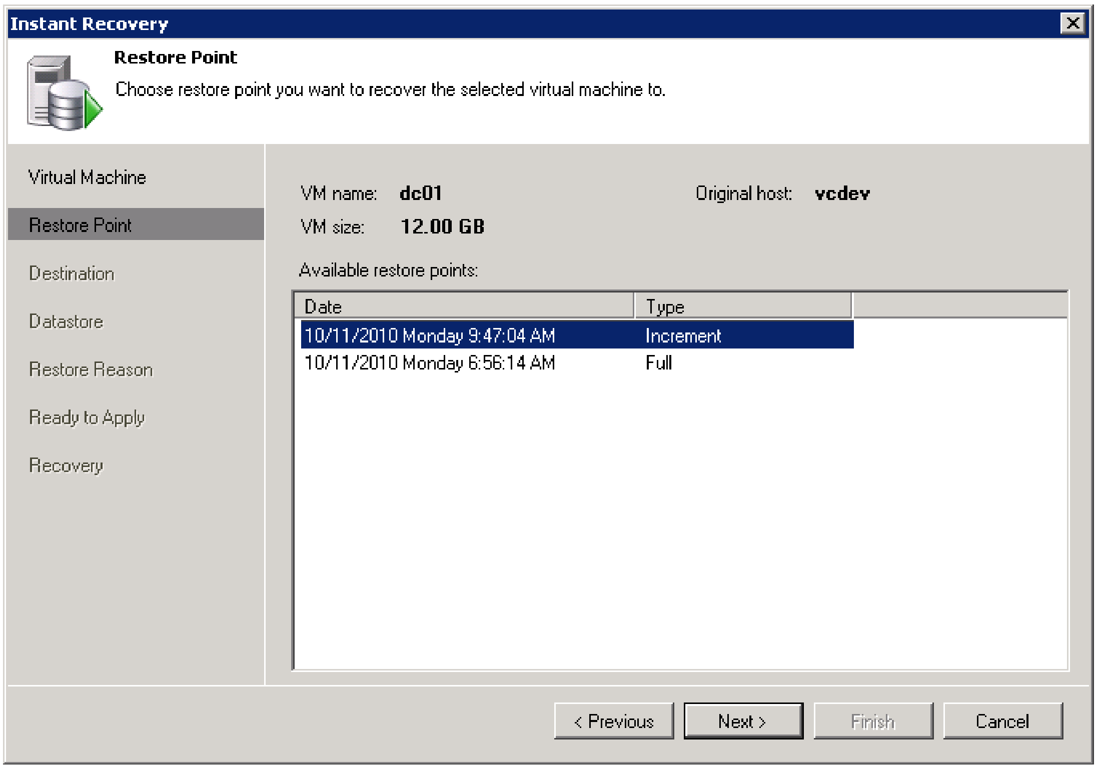

Restore Point

Choose restore point you want to recover the selected virtual machine to.

VM name: **dc01** Original host: **vcdev**

VM size: **12.00 GB**

Available restore points:

Date	Type
10/11/2010 Monday 9:47:04 AM	Increment
10/11/2010 Monday 6:56:14 AM	Full

Step 3. Select Destination for Recovered VM

Select the destination where the recovered VM should reside. In the **Host** field, specify a host on which a VM should run.

In the **Restored VM name** field, enter a name under which the VM should be restored and registered (if necessary). By default, the original name of the virtual machine is used. If you are restoring a VM to the same ESX host or the same datacenter where the original VM is registered, and the original VM still resides there, change the name to avoid conflicts.

In the **Resource pool** list, select a resource pool to which the virtual machine should be recovered.

If you are recovering a production VM that has failed and want restore with initial network settings, select the **Connect VM to network** check box. If you are recovering a VM for testing disaster recovery while the initial VM is still running, leave this check box not selected. Before you power on a VM, you will have to manually change VM network configuration from being connected to the production network, and re-connect it to an isolated non-production network to avoid conflicts.

To start a VM immediately after recovery, select the **Power on VM automatically** check box.

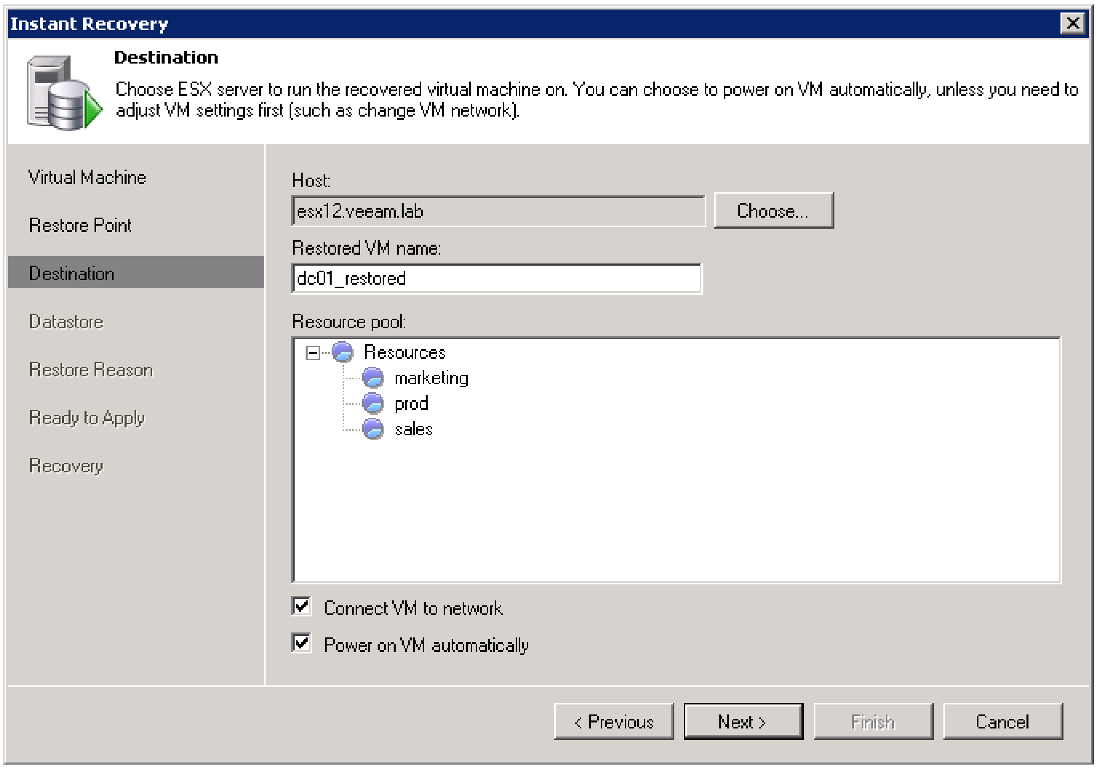

Destination

Choose ESX server to run the recovered virtual machine on. You can choose to power on VM automatically, unless you need to adjust VM settings first (such as change VM network).

Host: esx12.veeam.lab

Restored VM name: dc01_restored

Resource pool:

- Resources
  - marketing
  - prod
  - sales

☑ Connect VM to network

☑ Power on VM automatically

Step 4. Select Destination for Virtual Disk Updates

Select where disk changes should be written when a VM is restored. By default, disk changes are stored directly on the Veeam Backup Server. However, you can store disk changes to any datastore in your VMware environment. Select the **Redirect virtual disk changes** and choose a necessary datastore. Redirecting disk changes improves recovery performance but makes Storage vMotion not possible.

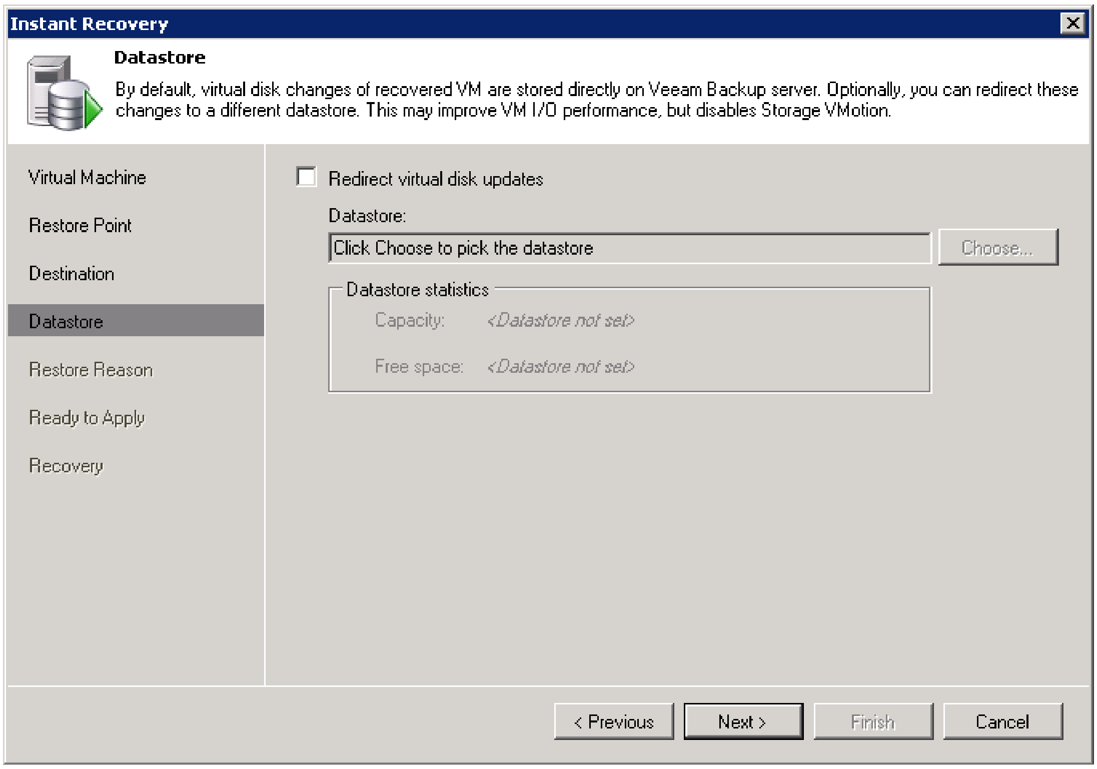

Instant Recovery

Datastore

By default, virtual disk changes of recovered VM are stored directly on Veeam Backup server. Optionally, you can redirect these changes to a different datastore. This may improve VM I/O performance, but disables Storage VMotion.

□ Redirect virtual disk updates

Datastore: Click Choose to pick the datastore

Datastore statistics

Capacity: <Datastore not set>

Free space: <Datastore not set>

Step 5. Specify Restore Reason

If necessary, enter the reason for performing instant restore of a VM. The information you provide will be saved in the session history so that you can reference it later.

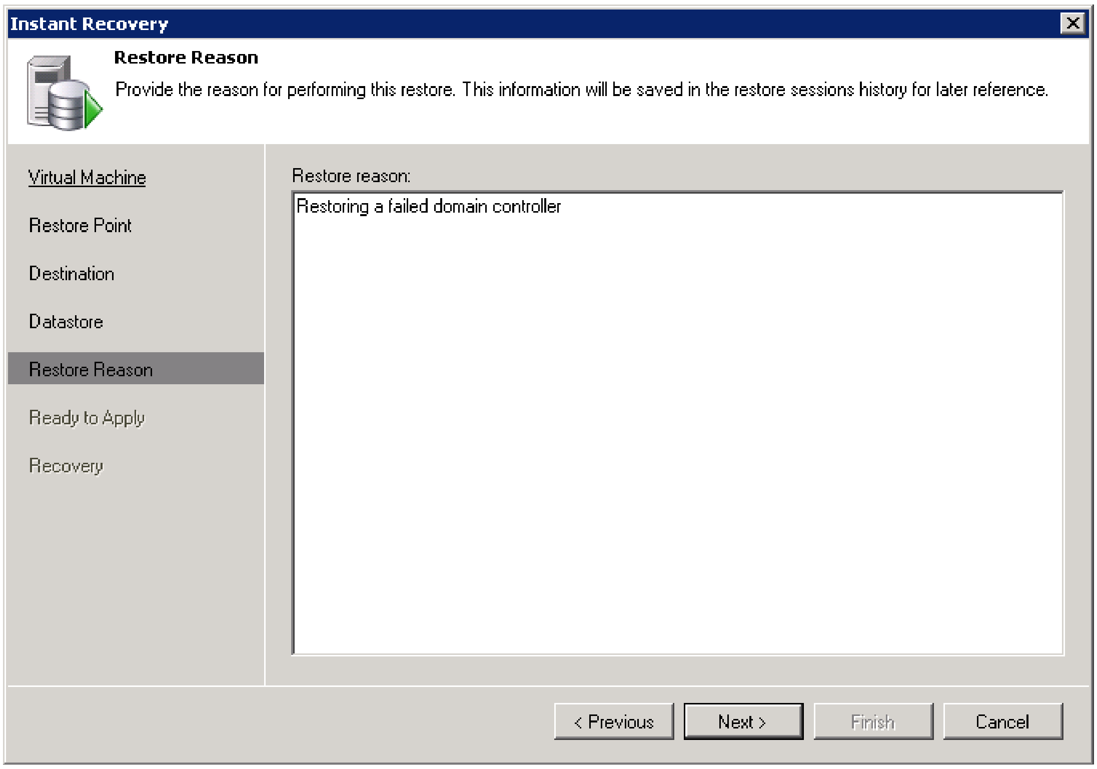

Instant Recovery

Restore Reason

Provide the reason for performing this restore. This information will be saved in the restore sessions history for later reference.

Restore reason: Restoring a failed domain controller

Step 6. Verify Instant Recovery Settings

Check specified settings for instant recovery of a VM and click **Next**. Veeam Backup & Replication will restore a VM on the selected ESX host.

Tip: To check the progress of instant VM recovery and view session details, click **Sessions** under **Restore** in the inventory tree and double-click a necessary instant VM restore session.

Restoring Full VM

With the **Restore** wizard, you can restore the entire VM and start it on the target host if necessary. This section will guide you through all steps of the wizard and provide explanation on offered options.

Step 1. Launch the Restore Wizard

To launch the Restore wizard, click the **Restore** button on the toolbar. Alternatively, you can select **Backup > Restore…** from the main menu.

You can also click the **Backups** node in the management tree, right–click a necessary VM in the corresponding backup job and select the **Restore entire VM...** command from the shortcut menu. In this case, you will immediately pass to the step 4 of the wizard.

Step 2. Select a Task

At the first step of the wizard, select **Entire VM (including registration)**.

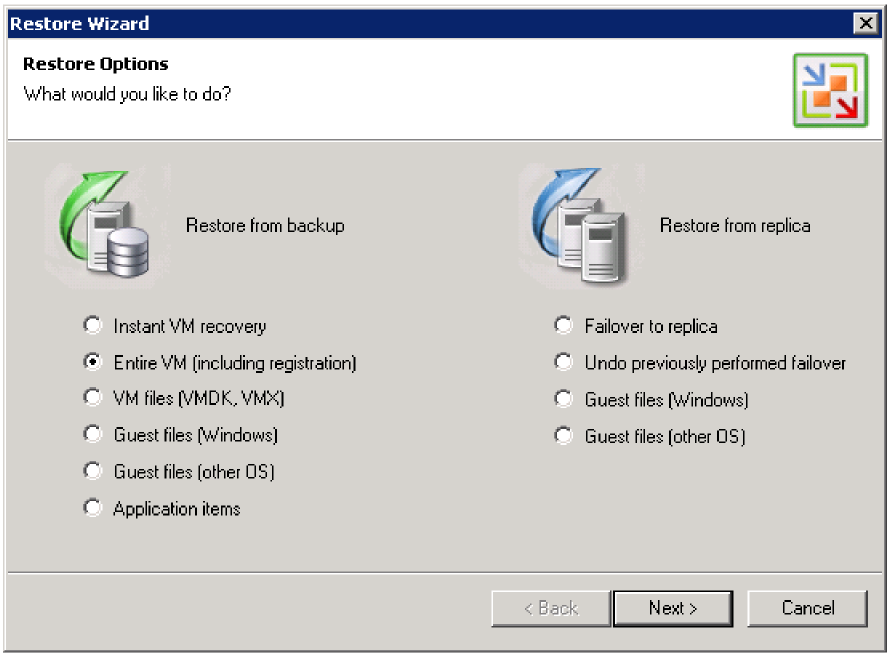

Restore Wizard

Restore Options
What would you like to do?

Restore from backup

○ Instant VM recovery
● Entire VM (including registration)
○ VM files (VMDK, VMX)
○ Guest files (Windows)
○ Guest files (other OS)
○ Application items

Restore from replica

○ Failover to replica
○ Undo previously performed failover
○ Guest files (Windows)
○ Guest files (other OS)

Step 3. Select a Virtual Machine

Select a necessary virtual machine in the list of available jobs.

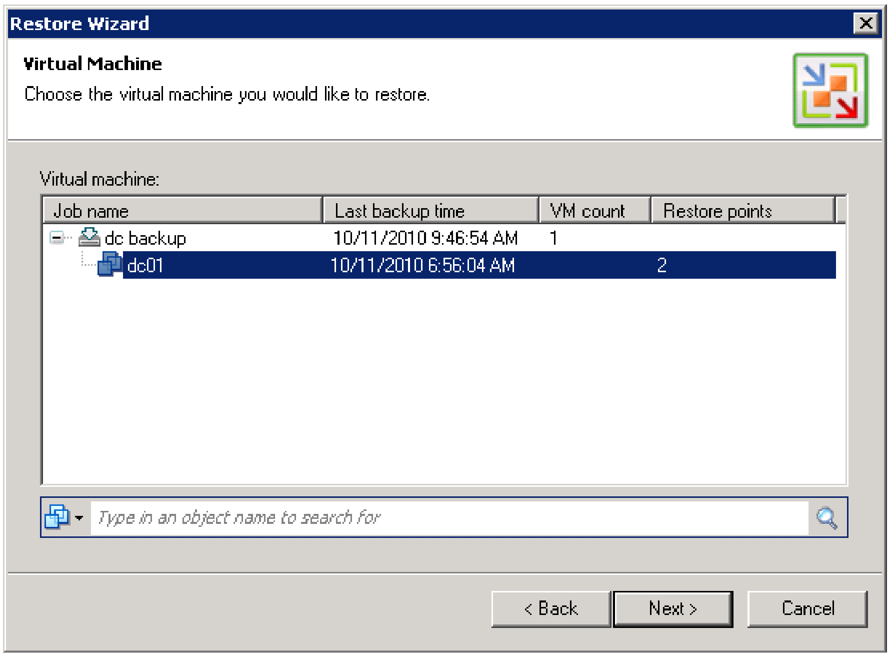

Virtual Machine

Choose the virtual machine you would like to restore.

Virtual machine:

Job name	Last backup time	VM count	Restore points
dc backup	10/11/2010 9:46:54 AM	1	
dc01	10/11/2010 6:56:04 AM		2

Type in an object name to search for

Tip: To quickly find VMs in jobs, use the search field at the bottom of the window.

Step 4. Select a Restore Point

Select a necessary restore point for the virtual machine. If you want to start the virtual machine after the work with the wizard is complete, select the **Power on VM after restoring** check box under the list of restore points.

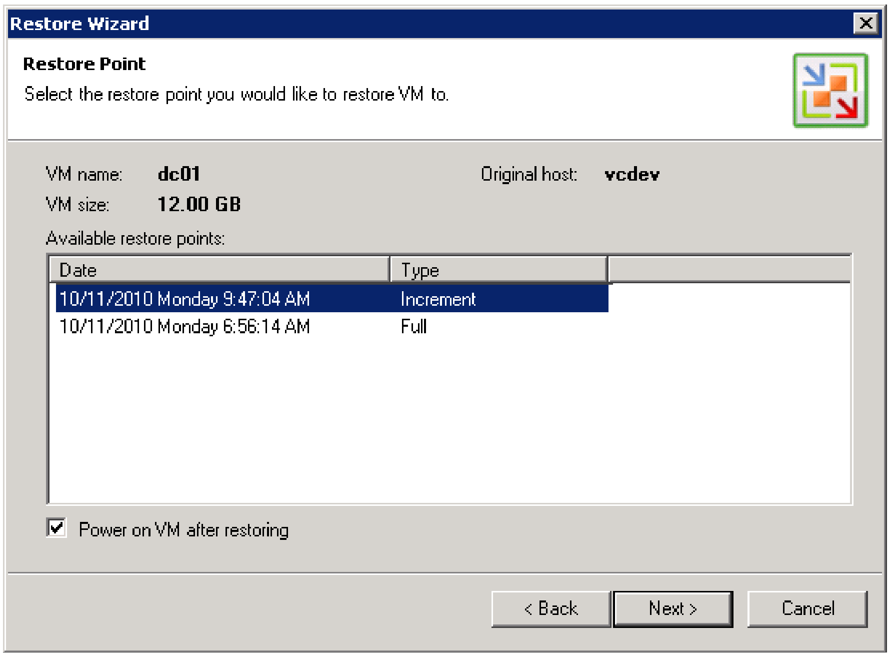

Restore Point

Select the restore point you would like to restore VM to.

VM name: **dc01** Original host: **vcdev**

VM size: **12.00 GB**

Available restore points:

Date	Type
10/11/2010 Monday 9:47:04 AM	Increment
10/11/2010 Monday 6:56:14 AM	Full

☑ Power on VM after restoring

Step 5. Select Destination for Restored VM

At this step of the wizard, you should select the destination where the restored VM should reside. From the **Host** list, select a host. Use the **Host Summary…** button to view information on storage resources.

In the **Virtual machine name** field, enter a name under which the VM should be restored and registered (if necessary). By default, the original name of the virtual machine is used.

Select a datastore and resource pool to which the virtual machine should be restored.

From the **Restore disks** list, select the type of disks for a restored VM. You can select to restore a VM in its original state, or force all VM disks thick or thin. Please note that this option is available only for VMs using virtual hardware version 7.

If you want to locate virtual disks of the restored VM on different datastores, click the **Choose a separate datastore for each virtual disk** link and select corresponding datastores.

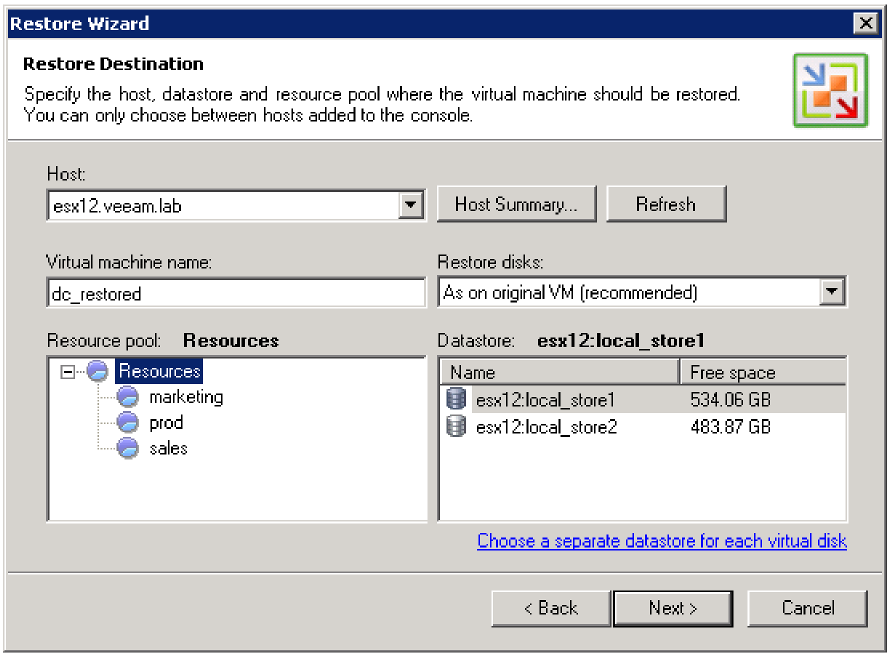

Specify the host, datastore and resource pool where the virtual machine should be restored. You can only choose between hosts added to the console.

Host: esx12.veeam.lab

Virtual machine name: dc_restored

Restore disks: As on original VM (recommended)

Resource pool: **Resources**

Datastore: **esx12:local_store1**

Name	Free space
esx12:local_store1	534.06 GB
esx12:local_store2	483.87 GB

Step 6. Complete the Work with the Wizard

Click **Finish** to start restoring the virtual machine.

Restoring VM Files: VMX, VMDK, etc

The **Restore** wizard allows you to restore specific VM files — .vmdk, .vmx, .vmsd, vmsn, .nvram files and so on. This section will guide you through all steps of the wizard and provide explanation on offered options.

Step 1. Launch the Restore Wizard

To launch the Restore wizard, click the **Restore** button on the toolbar. Alternatively, you can select **Backup > Restore…** from the main menu.

You can also click the **Backups** node in the management tree, right–click a necessary VM in the corresponding backup job and select the **Restore VM files...** command from the shortcut menu. In this case, you will pass to the step 4 of the wizard.

Step 2. Select a Task

At the first step of the wizard, select **VM files (VMX, VMDK)**.

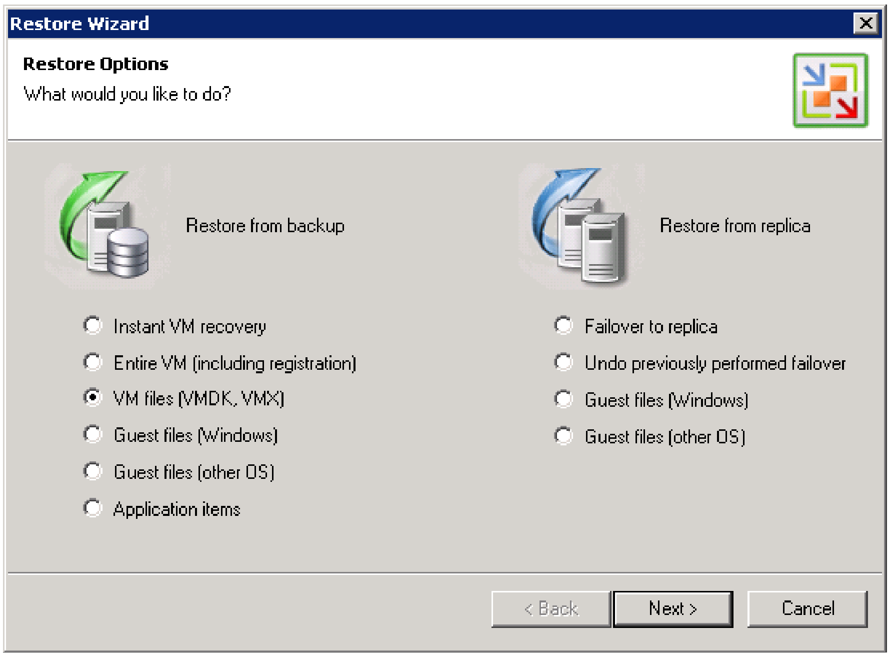

Restore Wizard

Restore Options

What would you like to do?

Restore from backup

- ○ Instant VM recovery
- ○ Entire VM (including registration)
- ● VM files (VMDK, VMX)
- ○ Guest files (Windows)
- ○ Guest files (other OS)
- ○ Application items

Restore from replica

- ○ Failover to replica
- ○ Undo previously performed failover
- ○ Guest files (Windows)
- ○ Guest files (other OS)

Step 3. Select a Virtual Machine

Select a necessary virtual machine in the list of available jobs.

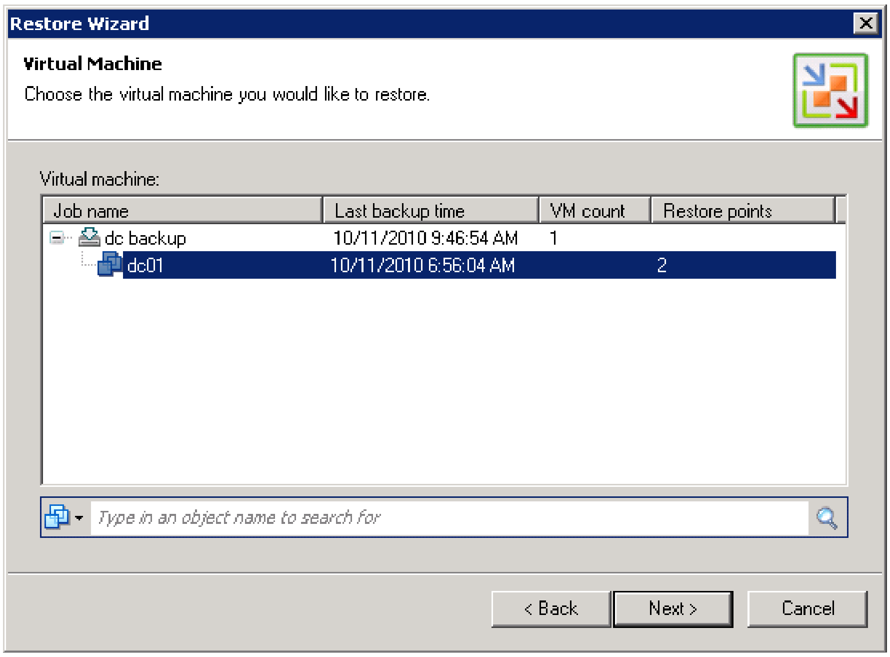

Restore Wizard

Virtual Machine

Choose the virtual machine you would like to restore.

Virtual machine:

Job name	Last backup time	VM count	Restore points
dc backup	10/11/2010 9:46:54 AM	1	
dc01	10/11/2010 6:56:04 AM		2

Type in an object name to search for

Tip: To quickly find VMs in jobs, use the search field at the bottom of the window.

Step 4. Select a Restore Point

Select a necessary restore point for the virtual machine.

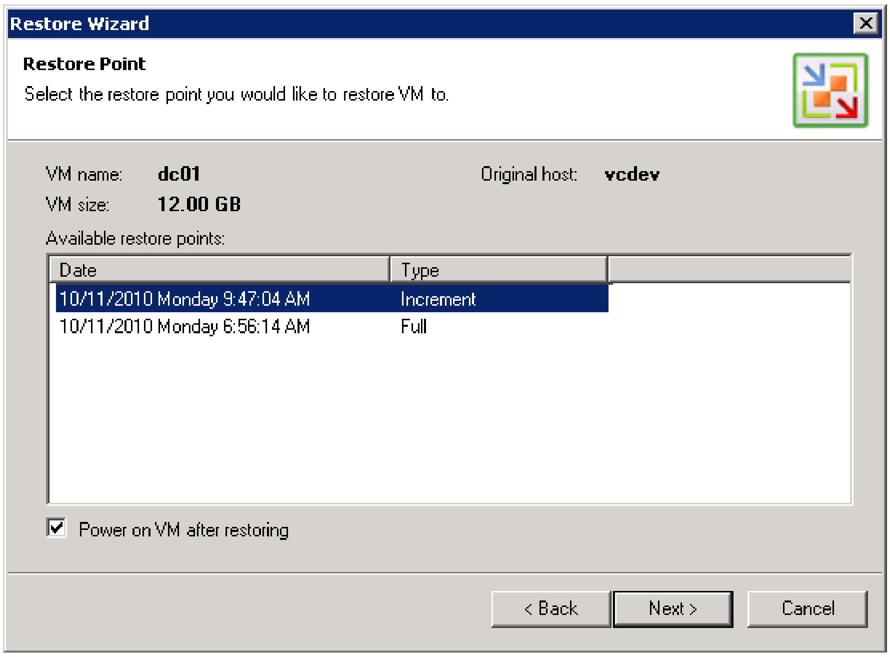

Restore Wizard

Restore Point

Select the restore point you would like to restore VM to.

VM name: **dc01** Original host: **vcdev**

VM size: **12.00 GB**

Available restore points:

Date	Type
10/11/2010 Monday 9:47:04 AM	Increment
10/11/2010 Monday 6:56:14 AM	Full

☑ Power on VM after restoring

Step 5. Select VM Files and Destination

At this step of the wizard, you should select the VM files you want to restore and the destination where the restored files should be stored. From the **Destination** list, select where to store VM files: to an ESX host or the local machine. Use the **Host Summary…** button to view information on storage resources. Specify the path to the folder on the selected host where files should be restored.

In the **VM files to restore** section, select check boxes next to files that should be restored. By default, all VM files are selected.

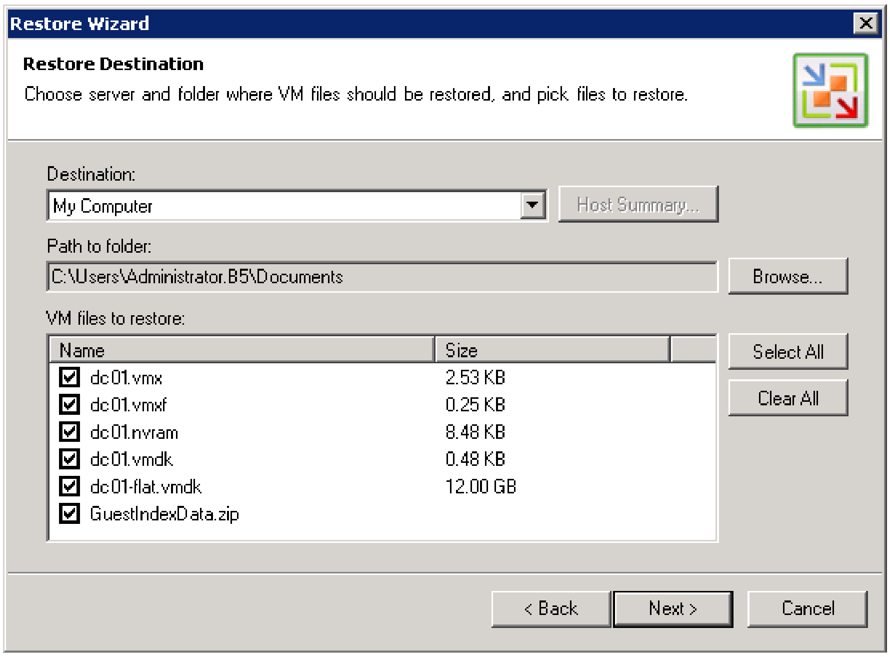

Restore Wizard

Restore Destination

Choose server and folder where VM files should be restored, and pick files to restore.

Destination: My Computer

Path to folder: C:\Users\Administrator.B5\Documents

VM files to restore:

Name	Size
☑ dc01.vmx	2.53 KB
☑ dc01.vmxf	0.25 KB
☑ dc01.nvram	8.48 KB
☑ dc01.vmdk	0.48 KB
☑ dc01-flat.vmdk	12.00 GB
☑ GuestIndexData.zip	

Step 6. Complete the Work with the Wizard

Click **Finish** to start restoring the VM files.

Restoring VM Guest Files

With the **Restore** wizard, you can restore individual VM guest OS files from any successfully created backup or replica. This section will guide you through all steps of the wizard and provide explanation on offered options.

Important! You cannot restore files from a replica that is currently running, or in case the replication or backup job with the VM from which you want to restore files is being performed.

Step 1. Launch the Restore Wizard

To launch the Restore wizard, click the **Restore** button on the toolbar. Alternatively, you can select **Backup > Restore…** from the main menu.

You can also click the **Backups** or **Replicas** node in the management tree, right–click a necessary backup or replica and select the **Restore guest files...** command from the shortcut menu. In this case, you will pass to the step 4 of the wizard.

Step 2. Select a Task

In the **Restore from backup** or **Restore from replica** section, select the **Guest files (Windows)**.

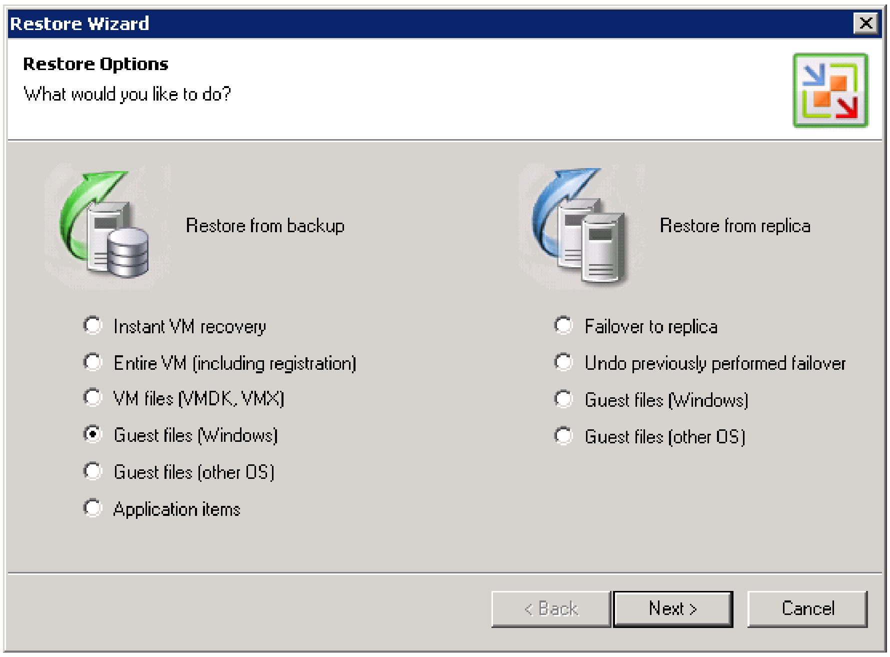

Restore Wizard

Restore Options

What would you like to do?

Restore from backup

○ Instant VM recovery
○ Entire VM (including registration)
○ VM files (VMDK, VMX)
● Guest files (Windows)
○ Guest files (other OS)
○ Application items

Restore from replica

○ Failover to replica
○ Undo previously performed failover
○ Guest files (Windows)
○ Guest files (other OS)

Step 3. Select a Virtual Machine

In the list of available jobs, select a necessary virtual machine.

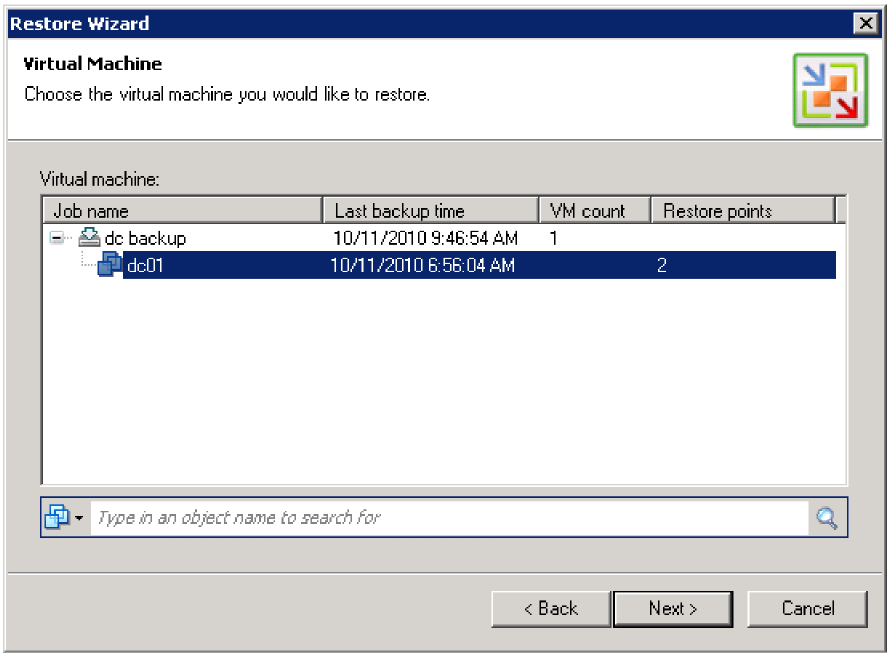

Restore Wizard

Virtual Machine

Choose the virtual machine you would like to restore.

Virtual machine:

Job name	Last backup time	VM count	Restore points
dc backup	10/11/2010 9:46:54 AM	1	
dc01	10/11/2010 6:56:04 AM		2

Tip: To quickly find VMs in jobs, use the search field at the bottom of the window.

Step 4. Select a Restore Point

Select a necessary restore point for the virtual machine.

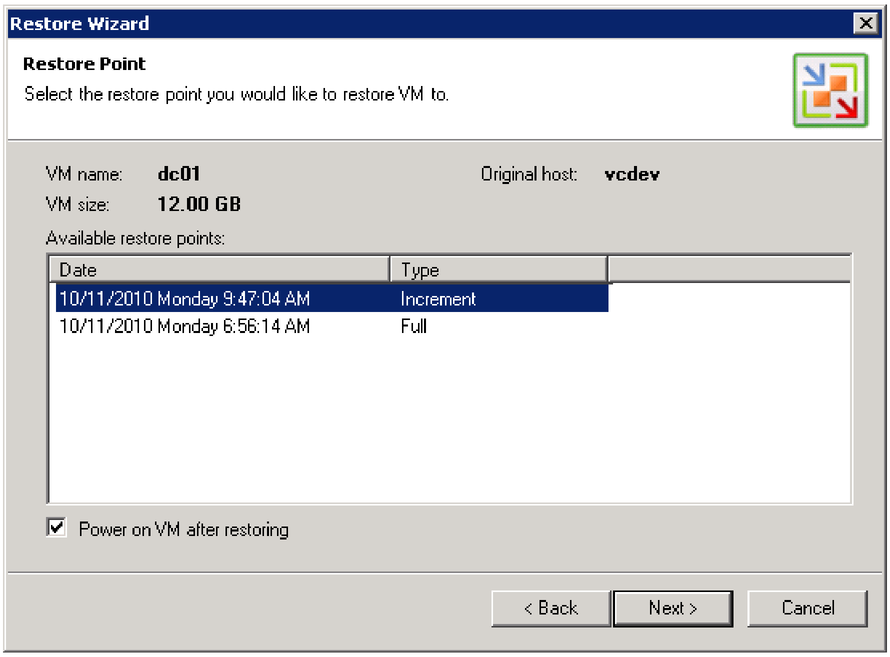

Restore Wizard

Restore Point

Select the restore point you would like to restore VM to.

VM name: **dc01** Original host: **vcdev**

VM size: **12.00 GB**

Available restore points:

Date	Type
10/11/2010 Monday 9:47:04 AM	Increment
10/11/2010 Monday 6:56:14 AM	Full

☑ Power on VM after restoring

Step 5. Specify Restore Reason

If necessary, enter the reason for performing VM guest file restore. The information you provide will be saved in the session history so that you can reference it later.

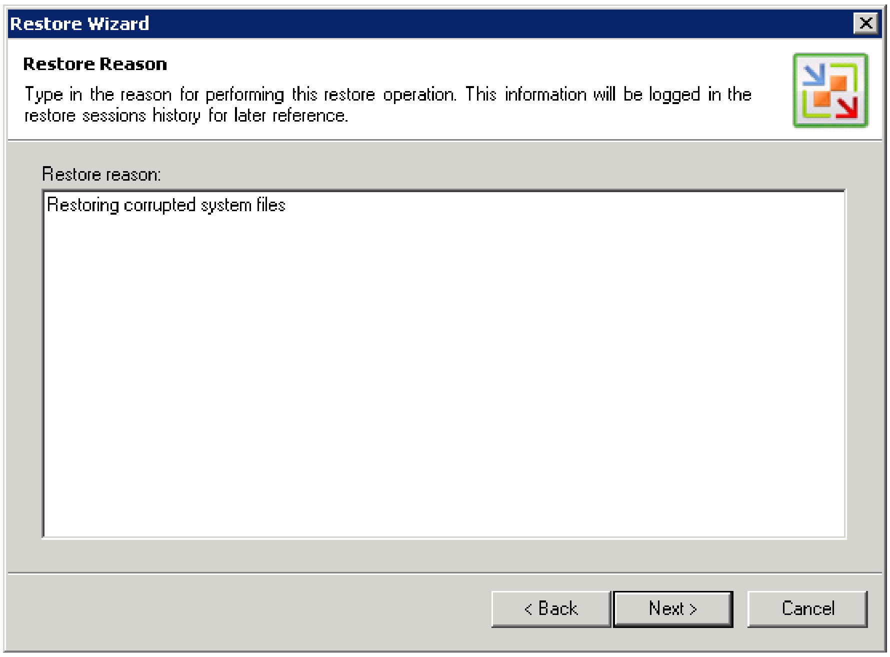

Restore Wizard

Restore Reason

Type in the reason for performing this restore operation. This information will be logged in the restore sessions history for later reference.

Restore reason: Restoring corrupted system files

Step 6. Complete the Work with the Wizard

Click **Finish** to start restoring files from a backup or replica. Once restoring is completed, a file browser displaying the file system tree of the restored virtual machine will be opened. Please note that the names of the restored machine drives may differ from the original ones.

Step 7. Save Restored Files

To save restored files or folders on the local machine or within the network, right–click a necessary node in the file system tree and select the **Copy To...** command from the shortcut menu.

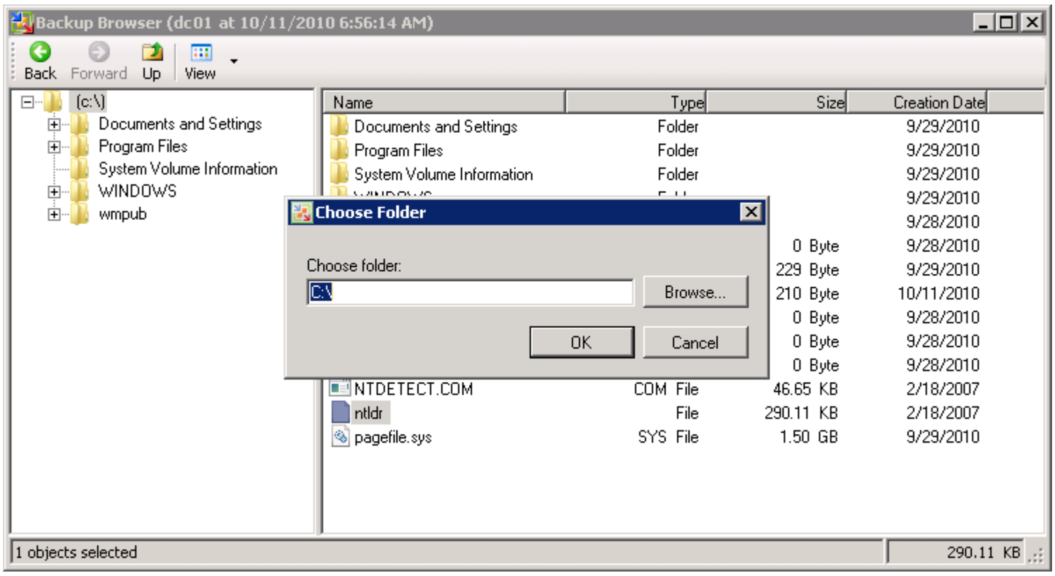

Backup Browser (dc01 at 10/11/2010 6:56:14 AM)

- (c:\)
  - Documents and Settings
  - Program Files
  - System Volume Information
  - WINDOWS
  - wmpub

Name	Type	Size	Creation Date
Documents and Settings	Folder		9/29/2010
Program Files	Folder		9/29/2010
System Volume Information	Folder		9/29/2010
WINDOWS	Folder		9/29/2010
			9/28/2010
		0 Byte	9/28/2010
		229 Byte	9/29/2010
		210 Byte	10/11/2010
		0 Byte	9/28/2010
		0 Byte	9/28/2010
		0 Byte	9/28/2010
NTDETECT.COM	COM File	46.65 KB	2/18/2007
ntldr	File	290.11 KB	2/18/2007
pagefile.sys	SYS File	1.50 GB	9/29/2010

Choose Folder

Choose folder: C:\

1 objects selected | 290.11 KB

Using Veeam File Level Restore Wizard

The Veeam File Level Restore wizard is intended for performing instant file-level restore of VM guest files and folders directly from image-level backups created with Veeam Backup & Replication 5.0, without requirement to extract actual VM disks files.

The key difference between the Veeam File Level Restore wizard and built-in file-level restore functionality that has been in the product since version 1.0, is additional file systems support. Built-in file level restore only supports Windows file system (NTFS and FAT). The Veeam File Level Restore wizard supports the following Linux, Unix, BSD, Solaris (ZFS) and Mac file systems:

Guest OS	Supported File System
Linux	ext2 ext3 ext4 ReiserFS (Reiser3 only) JFS XFS
Unix	JFS XFS UFS
BSD	UFS UFS2
Solaris	UFS ZFS
Mac	HFS HFS+
Windows	NTFS FAT FAT32

Tip: Recovery for file systems mentioned in the table is wizard-driven; however, with the new vPower engine in Veeam Backup & Replication, you can recover files from any file system. To restore individual files from file systems other than those specified in the table, you should leverage Instant VM Recovery functionality to publish VMDK from backup on vPower NFS datastore (without actually starting the VM). With the VMDK files readily available, you can mount these VMDKs to any VM that can read the corresponding file system (including the original VM), and restore the required files using native OS file management tools. Alternatively, you can mount the VMDK to a Windows VM, and use tools such as Portlock Explorer.

Starting from version 5 of Veeam Backup & Replication, the File Level Restore wizard can be used if Veeam Backup & Replication is installed both on a physical or virtual machine.

When restoring guest OS files, the File Level Restore wizard does not extract the VM image from the backup. Veeam Backup & Replication presents a compressed and deduplicated backup file as an NFS datastore to an ESX server so it is seen as a standard VM.

To perform file-level restore, the wizard uses a proxy appliance. The proxy appliance is a small VM that is created in the same network where a restored VM is located. Whenever you perform file-level restore, the File Level Restore wizard automatically starts the appliance in the necessary network and mounts to it disks of a backed up VM you want to restore files from. The wizard then displays the file browser window providing you with direct access to VM file system. You can copy the individual files and folders from VM disks to your local machine drive, network share or to a remote host.

The **File Level Restore** wizard provides support for Windows Logical Disk Manager (dynamic disks). Spanned/striped/raid-5 volumes and GPT disks are not supported.

The File Level Restore wizard provides support for Linux Logical Volume Manager. If LVM volumes are detected, the LVM node will be added to the file tree in the file browser.

If ZFS pools are detected, the ZFS node will be added to the file tree in the file browser.

Encrypted LVM volumes are not supported.

Restoring VM Guest Files

This section will guide you through all steps of the wizard and provide explanation on offered options.

Step 1. Launch the Veeam File Level Restore Wizard

To launch the Veeam File Level Restore wizard, select **Programs** > **Veeam** > **Veeam File Level Restore** from the **Start** menu or select **Tools** > **File Level Restore** > **Other OS** from the main menu of Veeam Backup & Replication 5.0. You can also start it from the general Restore wizard by selecting the **Guest files (other OS)** option.

Important! Before starting the wizard, make sure no file-level restore operation is performed in Veeam Backup & Replication 5.0. In the opposite case, the work of the wizard will be aborted.

The welcome screen of the wizard will be displayed. If you do not want to see the welcome screen at subsequent launches of the wizard, select the **Don't show this step again** check box at the bottom of the screen.

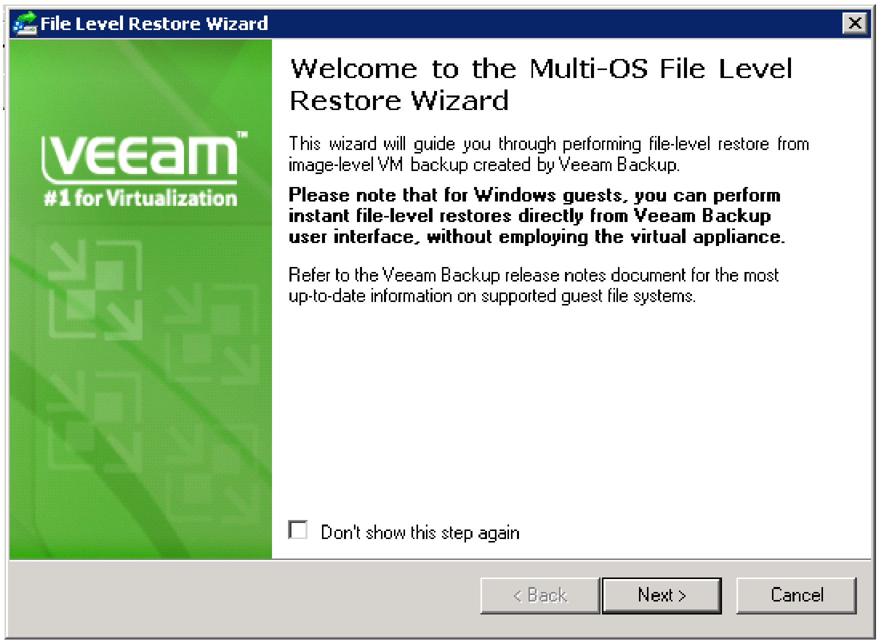

File Level Restore Wizard

Welcome to the Multi-OS File Level Restore Wizard

This wizard will guide you through performing file-level restore from image-level VM backup created by Veeam Backup.

Please note that for Windows guests, you can perform instant file-level restores directly from Veeam Backup user interface, without employing the virtual appliance.

Refer to the Veeam Backup release notes document for the most up-to-date information on supported guest file systems.

☐ Don't show this step again

Step 2. Select a Virtual Machine

At this step of the wizard, you will see a list of backups created with Veeam Backup & Replication 5.0. In the list of available backups, select a necessary virtual machine.

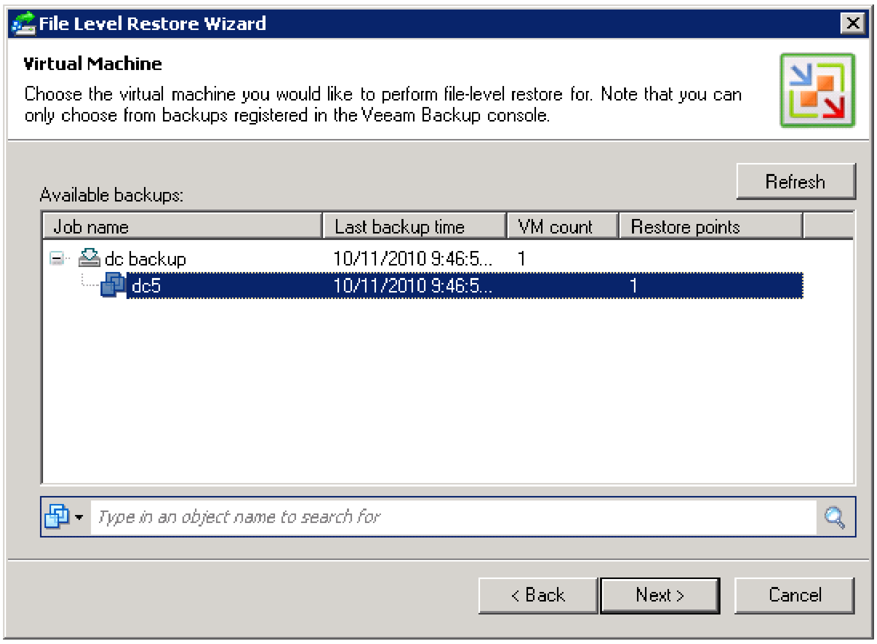
File Level Restore Wizard
Refresh
Type in an object name to search for
< Back
Next >
Cancel

Virtual Machine

Choose the virtual machine you would like to perform file-level restore for. Note that you can only choose from backups registered in the Veeam Backup console.

Available backups:

Job name	Last backup time	VM count	Restore points
dc backup	10/11/2010 9:46:5...	1	
dc5	10/11/2010 9:46:5...		1

Step 3. Select the Restore Point

Select a necessary restore point for the virtual machine.

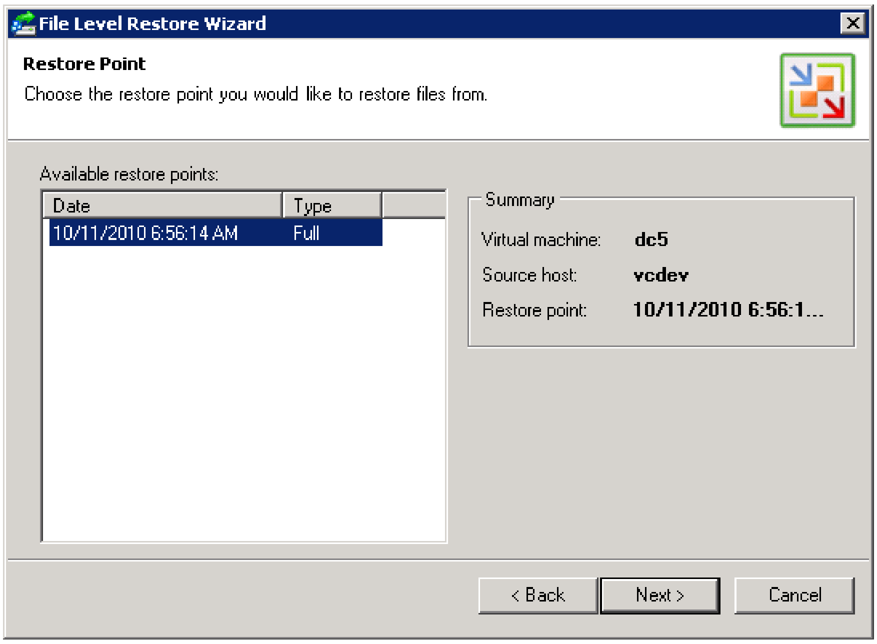

Restore Point

Choose the restore point you would like to restore files from.

Available restore points:

Date	Type
10/11/2010 6:56:14 AM	Full

Summary

Virtual machine: **dc5**

Source host: **vcdev**

Restore point: **10/11/2010 6:56:1...**

Step 4. Specify Restore Reason

If necessary, enter the reason for performing restore of VM guest OS files. The information you provide will be saved in the session history so that you can reference it later.

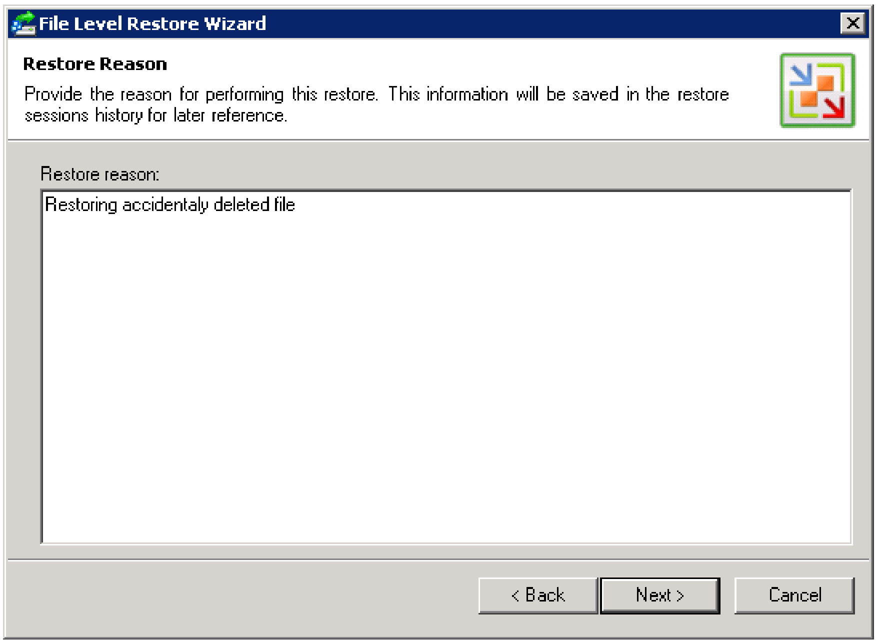

Restore Reason

Provide the reason for performing this restore. This information will be saved in the restore sessions history for later reference.

Restore reason: Restoring accidentaly deleted file

Step 5. Select Location for Proxy Appliance

At this step of the wizard, you should select where a proxy appliance used to mount VM file system should be run.

To locate the appliance, at the last step of the wizard, click the **Choose** button.

In the **FLR Appliance Settings** window, select the ESX host, resource pool and the network on which the proxy appliance will be run. Keep in mind that you should locate the proxy appliance in the same network where a VM from which you want to restore files resides.

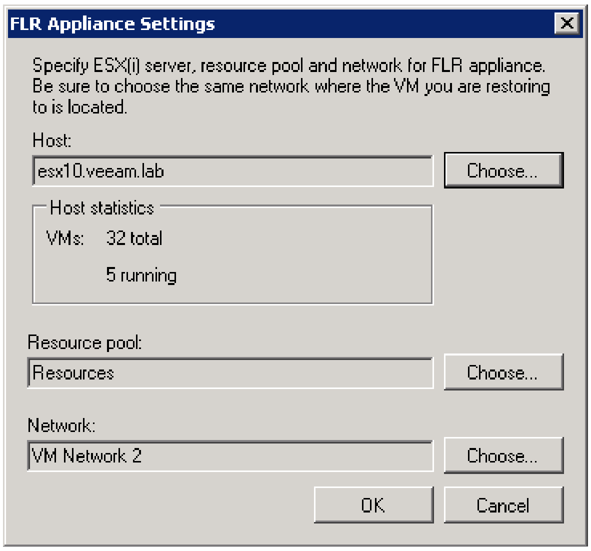

Specify ESX(i) server, resource pool and network for FLR appliance. Be sure to choose the same network where the VM you are restoring to is located.

Host: esx10.veeam.lab

Host statistics

VMs: 32 total

5 running

Resource pool: Resources

Network: VM Network 2

Step 6. Complete the Work with the Wizard

Click **Finish** to start restoring files from a backup or replica. Please note that the file-level restore appliance may take about 30-40 seconds to boot.

Step 7. Save Restored Files

Once the restore process is completed, a file browser displaying the file system tree of the restored virtual machine will be opened.

To save restored files or folders on the local machine or within the network, right–click a necessary file or folder and select the **Copy to...** command from the shortcut menu and select a necessary destination and folder on the local or remote host. The file or folder will be saved at the specified folder on the host.

Note: If you are recovering files to a remote Linux host, you can select the **Preserve permissions and ownership** check box to keep original permission settings for recovered files. Ownership settings are restored only if you have privileges to change the owner at the remote Linux host where files are restored.

Restoring VMs with the Extract Utility

Veeam Backup & Replication comes with an extract utility that can be used to recover VMs from a backup file to the latest restore point. The utility can be used as an independent tool both on Linux and Windows computers as it does not require any interaction with Veeam Backup & Replication.

The extract utility can be helpful, for example, if it is written to the tape next to backup files — in this case, you get a possibility to recover VMs from backups at any moment of time even if backups are removed from Veeam Backup & Replication or the application is uninstalled at all. Please keep in mind that the extract utility always restores a VM to the latest restore point.

The installation folder of Veeam Backup & Replication contains two executable files named *extract* for Windows and Linux. You can run the executable file and work with the utility in the interactive mode, or start the utility from the command line.

Using the Extract Utility in the Interactive Mode

To start the extract utility in the interactive mode, run the *Extract.exe* file from the installation folder of Veeam Backup & Replication (if Veeam Backup & Replication is installed on a Linux machine, run the *Extract* file).

You will have to sequentially enter the following arguments:

1. Path to the backup file from which VMs should be restored. After you enter the path, the restore utility will display a list of all VMs included in the backup and their description.
2. Name of a VM(s) you want to restore. If there are more than one VM with the specified name in the backup, you will be asked to specify the host on which the backed up VM resides. If you want to restore all VMs from the backup, press **Enter** on the keyboard.
3. Output directory to which VMs should be restored. If you want to restore VM(s) to the current directory, press **Enter**.
4. Press **Y** on the keyboard to restore a VM to the directory you selected. If you want to abort the operation, press **Enter**.

Using the Extract Utility from the Command Line

If you run the extract utility from the command line, you can perform the following actions:

- Run the extract utility in the interactive mode
- Display help information for the utility usage
- Display the list of all VMs in the backup file
- Restore all or selected VMs from the backup

Running the Extract Utility in the Interactive Mode

This command runs the extract utility in the interactive mode.

Syntax

extract.exe [pathtovbk]

Parameters

Parameter	Description	Required/Optional
pathtovbk	Path to the backup file from which VM(s) should be restored	Optional

Displaying Help Information for the Utility Usage

This command prints all variants of the extract utility usage along with required and optional parameters.

Syntax

extract.exe -help

Displaying the List of VMs in the Backup

This command displays the list of all VMs in the backup file from which you want to perform restore.

Syntax

extract.exe -dir [-vm vmname] [-host hostname] pathtovbk

Parameters

Parameter	Description	Required/Optional
vm	Name of a VM that you want to display. Use this parameter to filter VMs in the backup job.	Optional
host	Name of the host on which a backed up VM resides. This parameter is used if the vm parameter is specified. Use this parameter to filter VMs that have the same name but reside on different hosts.	Optional
pathtovbk	Path to the backup file from which VM(s) should be restored.	Required

Restoring VMs from Backup

This command restores all or selected VM files from the backup file.

Syntax

extract.exe -restore [-vm vmname] [-host hostname] pathtovbk [outputdir]

Parameters

Parameter	Description	Required/Optional
vm	Name of a VM that you want to restore. If you do not specify this parameter, all VMs from the backup will be restored.	Optional
host	Name of the host on which a backed up VM resides. This parameter is used if the vm parameter is specified. Use this parameter to filter VMs that have the same name but reside on different hosts.	Optional
pathtovbk	Path to the backup file from which VM(s) should be restored.	Required
outputdir	Path to the directory to which VM(s) should	Optional

	be restored. If this parameter is not specified, VM(s) will be restored to the current directory.	

Performing Replica Failover

With the virtual machine replica failover option, you can recover a corrupted virtual machine in case of software or hardware malfunction. The failover option can be used for any virtual machine replica that was successfully created at least once.

Failing Over VM Replicas

Failing over replicas is performed by means of the **Restore** wizard. This section will guide you through all steps of the wizard and provide explanation on offered options.

Note: Remember to power off the original virtual machine on the source host before starting failover. To avoid unwanted interference with the replica files, stop the corresponding replication job, too.

Step 1. Launch the Restore Wizard

To launch the Restore wizard, click the **Restore** button on the toolbar. Alternatively, you can select **Backup > Restore…** from the main menu.

You can also click the **Replicas** node in the management tree, right–click a necessary virtual machine in the information pane and select the **Failover to a Particular Version...** command from the shortcut menu. In this case, you will pass to the step 4 of the wizard.

Step 2. Select a Task

Select the **Failover to replica** option.

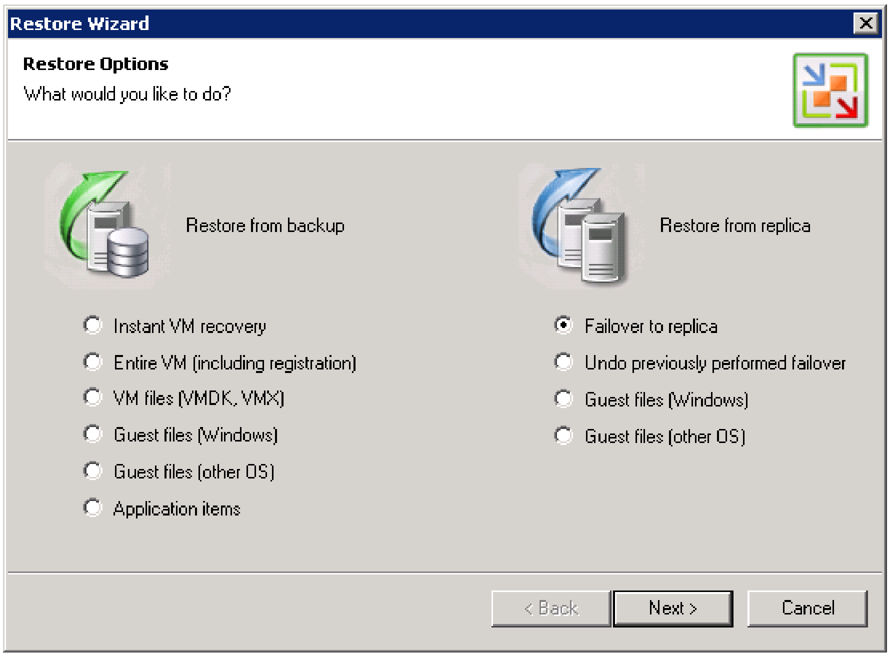

Restore Wizard

Restore Options

What would you like to do?

Restore from backup

○ Instant VM recovery
○ Entire VM (including registration)
○ VM files (VMDK, VMX)
○ Guest files (Windows)
○ Guest files (other OS)
○ Application items

Restore from replica

● Failover to replica
○ Undo previously performed failover
○ Guest files (Windows)
○ Guest files (other OS)

Step 3. Select a Virtual Machine

Select a necessary virtual machine in the list of available jobs.

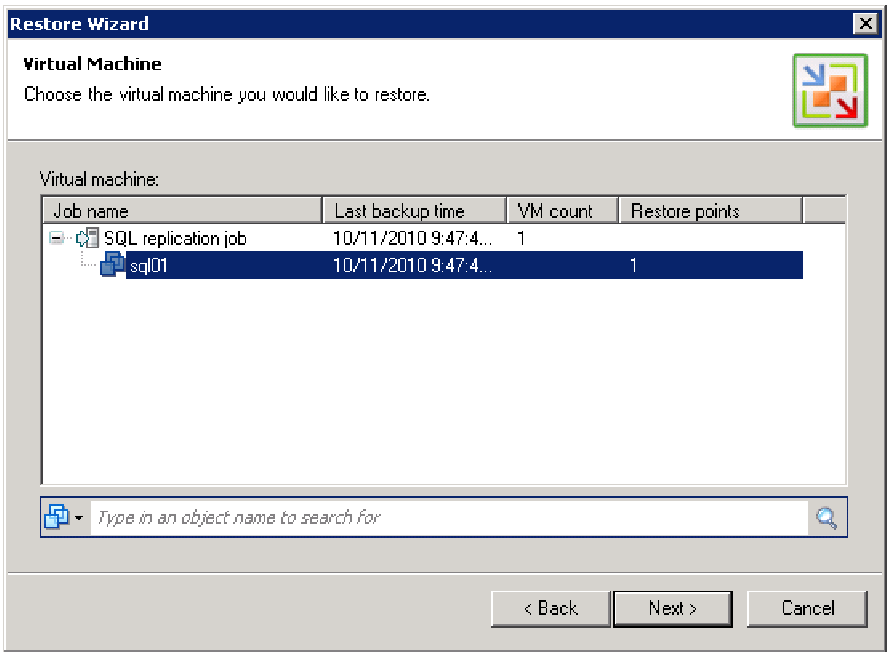

Restore Wizard

Virtual Machine

Choose the virtual machine you would like to restore.

Virtual machine:

Job name	Last backup time	VM count	Restore points
SQL replication job	10/11/2010 9:47:4...	1	
sql01	10/11/2010 9:47:4...		1

Step 4. Select a Restore Point

Select a necessary restore point for the virtual machine.

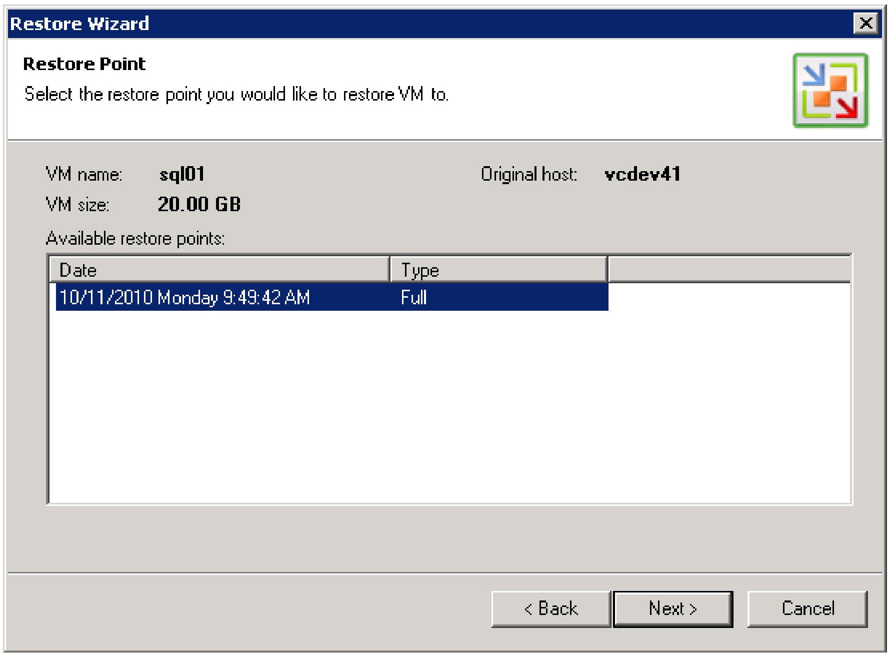

Restore Wizard

Restore Point

Select the restore point you would like to restore VM to.

VM name: **sql01** Original host: **vcdev41**

VM size: **20.00 GB**

Available restore points:

Date	Type
10/11/2010 Monday 9:49:42 AM	Full

Step 5. Complete the Work with the Wizard

Click **Finish** to start failing over the selected restore point. The virtual machine will be powered on on the target host.

Undoing Failover

The Undo failover option allows powering off failed over virtual machines on the target host and rolling back to their initial state. Undoing failover is performed by means of the Restore wizard. This section will guide you through all steps of the wizard and provide explanation on offered options.

Step 1. Launch the Restore Wizard

To launch the **Restore** wizard, click the **Restore** button on the toolbar. Alternatively, you can select **Backup > Restore…** from the main menu.

You can also click the **Replicas** node in the management tree, right–click a necessary virtual machine in the information pane and select the **Undo Failover** command from the shortcut menu. In this case, the undo failover operation will be immediately performed for the selected VM.

Step 2. Select a Task

Select the **Undo previously performed failover** option.

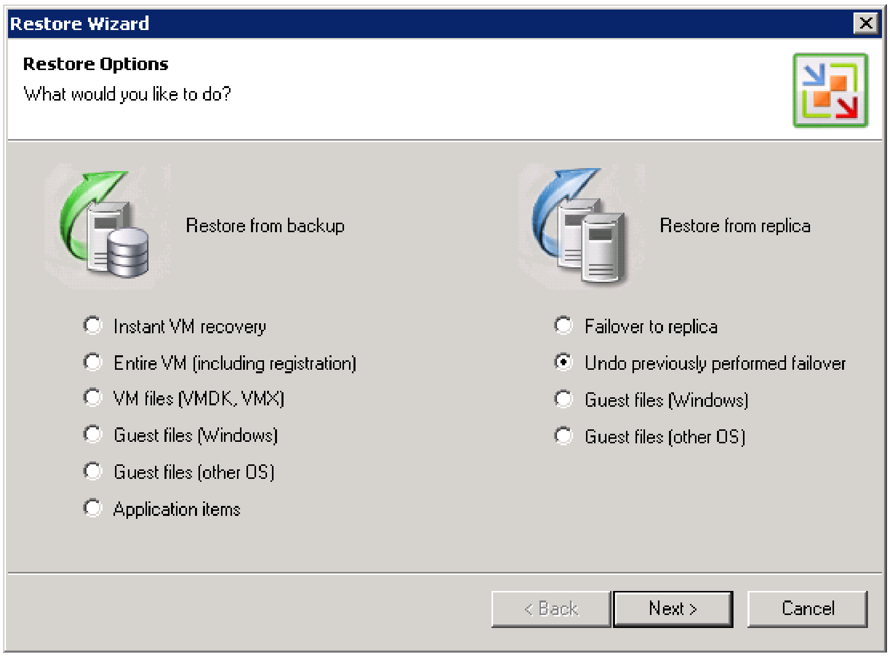

Restore Wizard

Restore Options

What would you like to do?

Restore from backup

- ○ Instant VM recovery
- ○ Entire VM (including registration)
- ○ VM files (VMDK, VMX)
- ○ Guest files (Windows)
- ○ Guest files (other OS)
- ○ Application items

Restore from replica

- ○ Failover to replica
- ● Undo previously performed failover
- ○ Guest files (Windows)
- ○ Guest files (other OS)

Step 3. Select a Virtual Machine

Select a necessary virtual machine in the list of available jobs.

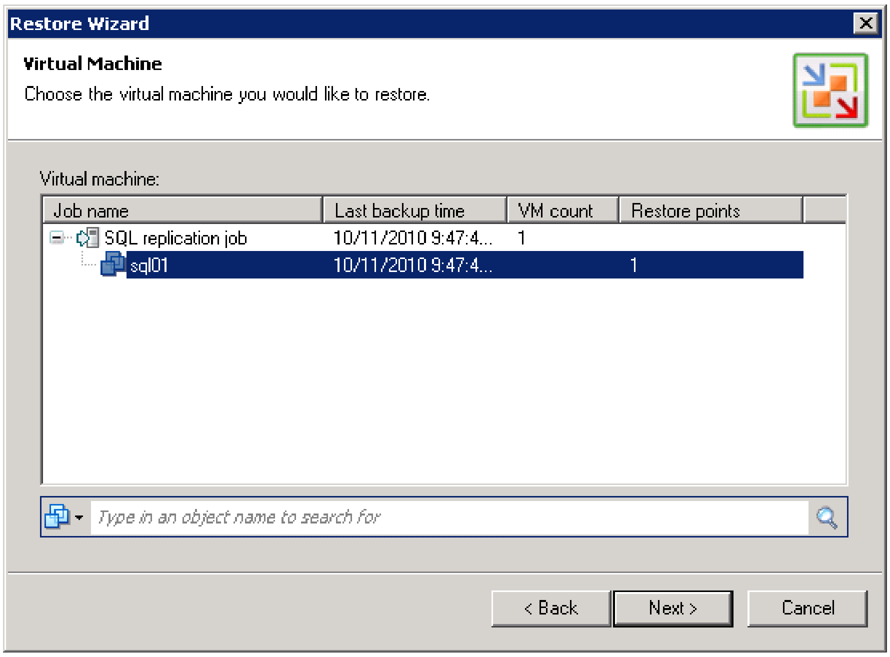

Restore Wizard

Virtual Machine

Choose the virtual machine you would like to restore.

Virtual machine:

Job name	Last backup time	VM count	Restore points
SQL replication job	10/11/2010 9:47:4...	1	
sql01	10/11/2010 9:47:4...		1

Type in an object name to search for

Step 4. Complete the Work with the Wizard

Click **Finish** to undo failover for the selected machine.

Note: When undoing failover, you will lose all changes that you made to the replicated virtual machine since it was powered on.

Managing Backups & Replicas

Veeam Backup & Replication 5.0 offers the following management options for your backups and replicas: removing from backups/replicas, deleting from disks and viewing properties. All options are available from the shortcut menu.

- The **Remove from Backups or Replicas** option is used when you want to remove records about backup and replica files from the Veeam Backup configuration database. Please note that all backup files (VBK and VRB) will stay safe on the destination backup storage, so you can easily import these files later to the Veeam Backup & Replication console for restore operations if needed.
  As for replicas, all references will be removed from the Veeam Backup & Replication console; however, all your replicated VMs will still reside on your target hosts, so you can start them manually after the **Remove from Replicas** option is performed.
- In addition to removing records about backup and replica files from the Veeam Backup configuration database, the **Delete from Disk** option also removes actual backups and replicas from the destination storage. Note that you should avoid deleting backup files manually from your destination storage, otherwise all subsequent job sessions will be failing.
  You can use this option for all VMs in the backup or replication job or for each VM separately. Granular deletion of VMs can be useful, for example, if some replica has been started directly from vSphere Client instead of using the failover option in Veeam Backup & Replication 5.0. In this case, further replication jobs would be failing. Previously you would have to delete all replicas created with such job. Now, Veeam Backup & Replication 5.0 creates a separate storage of configuration files and replica data checksum files for each replicated VM. This storage is kept next to the replicated VM on the replication target. When a separate replicated VM is deleted, Veeam Backup & Replication 5.0 deletes the replica itself and files created for it.
- The **Properties** option for backups is used to view summary information on backups you made. It contains information on compression and de-duplication ratios, available restore points for a particular backup, as well as date, data size and backup size.

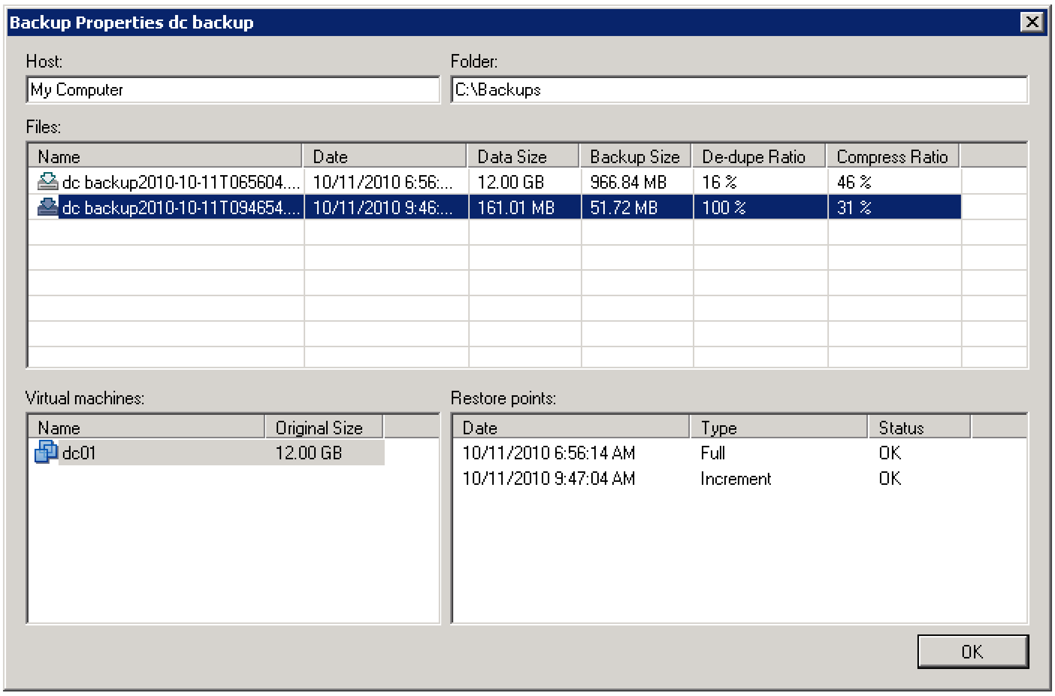
Backup Properties dc backup

Host: My Computer

Folder: C:\Backups

Files:

Name	Date	Data Size	Backup Size	De-dupe Ratio	Compress Ratio
dc backup2010-10-11T065604....	10/11/2010 6:56:...	12.00 GB	966.84 MB	16 %	46 %
dc backup2010-10-11T094654....	10/11/2010 9:46:...	161.01 MB	51.72 MB	100 %	31 %

Virtual machines:

Name	Original Size
dc01	12.00 GB

Restore points:

Date	Type	Status
10/11/2010 6:56:14 AM	Full	OK
10/11/2010 9:47:04 AM	Increment	OK

Importing Backups

Importing backups can be useful if you need to restore backups from tape or from .vbk files of other Veeam Backup & Replication versions or instances, if you happened to delete the server with which the backup was associated from the management tree, or in case the application has been uninstalled.

To import backups to Veeam Backup & Replication 5.0, click the **Import Backup** button on the toolbar or select **Backup > Import Backup…** from the main menu. From the **Computer** list, select the host on which a backup you want to import is stored. Then browse to a necessary .vbk file by clicking the **Browse…** button.

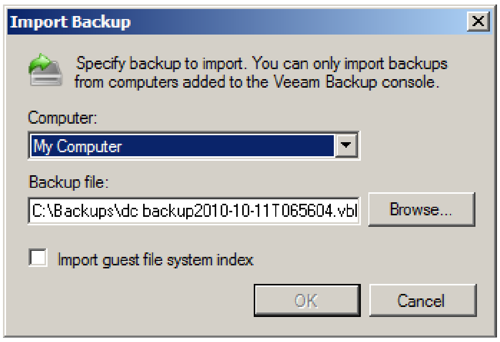

Specify backup to import. You can only import backups from computers added to the Veeam Backup console.

Computer: My Computer

Backup file: C:\Backups\dc backup2010-10-11T065604.vbl

□ Import guest file system index

Note: By default, index data of the guest OS file system is not imported with the backup file to speed up the import process. However, if it is necessary, select the **Import guest file system index** check box.

Click **OK** to import the selected backup. The imported backup data will be stored under the original backup job name with the *_imported* suffix appended.

You can also import a backup that was deleted from the list of backups (to display the list, click the **Backups** node under **Backup** in the management tree). In this case, the backup will be imported under its original name.

Note: To be able to perform any restore operation for previous points in time (rollbacks) for your backed up VM, before importing a full backup file to the Veeam Backup & Replication console, make sure that you have all required increments (either forward or reverse) in the same folder.

Performing Recovery Verification with SureBackup

SureBackup is a new technology in Veeam Backup & Replication 5.0 developed to automate and simplify backup verification process – one of the most crucial parts of data management and protection. It is a collection of features that allows you to start VMs directly from VM backups in a fenced-off environment and perform backup reliability and availability testing as a routine part of the backup process.

Note: Recovery verification functionality is available in Veeam Backup & Replication 5.0 Enterprise Edition only.

vPower NFS

For instant VM recovery, SureBackup recovery verification, U-AIR and multi-OS file-level restore operations, Veeam Backup & Replication 5.0 uses vPower NFS service. vPower NFS service is a Windows service that runs on the Veeam Backup server and enables it to act as an NFS server. The vPower technology lets you publish a compressed and deduplicated backup file as a regular VMDK file directly to the ESX server via NFS, so ESX servers get transparent access to backed up VM images.

To be able to successfully connect an ESX host to the NFS server, you should make sure that the ESX host has a proper network interface configuration and can access the Veeam Backup & Replication server on which vPower NFS service is running.

When connecting to the NFS server, the ESX(i) host uses a VMkernel interface. That is why ESX(i) host you are using must have a VMKernel interface – otherwise mounting vPower NFS on ESX(i) host will fail.

By default, VMKernel interfaces are not available for non-ESXi versions — so you will have to add them on the ESX host to be able to connect to the NFS server.

- If the Veeam Backup server and ESX host are located in the same subnet, the ESX server should have a VMkernel interface in the same IP-network as the Veeam Backup server has.
- If the Veeam Backup server and the ESX host are located in different subnets and use a router for network access, in addition to creating a new VMkernel interface, you will have to manually specify routing settings in the IP routing table on the ESX server.

Tip: To check whether an ESX host can access the Veeam Backup server or not, you can use the *vmkping* utility on the ESX server. The *vmkping* utility is similar to the ping tool – the only difference is that ICMP packets are sent via the VMkernel interface, not the service console interface.

SureBackup Overview

When performing recovery verification of VM backups, Veeam Backup & Replication runs VMs directly from backup files without restoring them to a production datastore. This is achieved by utilizing the vPower NFS service running on the backup server. This service presents VM images on the backup storage as an NFS datastore to an ESX server so a VM backup is seen as a regular VM image.

During verification, a backed up VM image remains in the read-only state – all changes that take place when a VM is running are written to redo log files that are stored on a selected datastore in the production environment. Once the recovery verification process is complete, the redo logs are removed.

To perform recovery verification testing, you need to create an application group required to verify full functionality of backed up VMs, an isolated virtual lab where VMs should be tested, and a recovery verification job.

Application Group

An application group contains VMs running production applications on which VMs to be verified are dependent. That is, it includes all components and services that should be started to enable fully functional work of VMs you want to test.

For instance, to verify an Exchange Server, it is necessary to test its functionality in cooperation with other components. Thus, other VMs such as Active Directory Domain Controller and DNS server need to be started in the same test environment to make verification of the Exchange Server possible.

The application group specifies roles for each VM, boot priority and boot delays, so all services requested for VM recovery verification are started in the required order. Typically, an application group contains at least a Domain Controller, DNS server and DHCP server.

To learn more, see the Creating an Application Group section.

Virtual Lab

A virtual lab is an isolated virtual test environment where verified VMs with all components required for their proper operation are started and tested. A virtual lab is created using existing resources in your VI environment and ensures secure integrity and functionality testing for backed up VMs.

When a new virtual lab is created, Veeam Backup & Replication adds a new resource pool, VM folder and vSwitch on the host where the virtual lab is registered. The network configuration in the virtual lab reproduces the configuration of the production network. For example, if a tested VM and its dependencies are located in two logical networks in your production environment, these two networks will be recreated in the virtual lab and mapped to corresponding production networks.

To enable communication between an outer world and VMs and the virtual lab, Veeam Backup & Replication uses a proxy appliance that is created and registered in the folder and resource pool of the virtual lab. This proxy appliance is a VM that acts as a gateway routing requests from the production network to the isolated network.

To connect to isolated networks, Veeam Backup & Replication adds to the proxy appliance a vNIC adapter for each network. Each vNIC adapter gets an IP address from the network to which it is connected, which is typically the same as the IP address of a default gateway in the corresponding production network.

If the application group to be started in the virtual lab does not have a DHCP server and some applications in this group as well as verified applications require DHCP, you can enable the DHCP service on a vNIC adapter for each isolated network. You can also select specific DNS servers from the production network that should be started in the isolated network. Keep in mind that to be able to add a DNS server, you should have it virtualized in your production environment, and you should also have its backup.

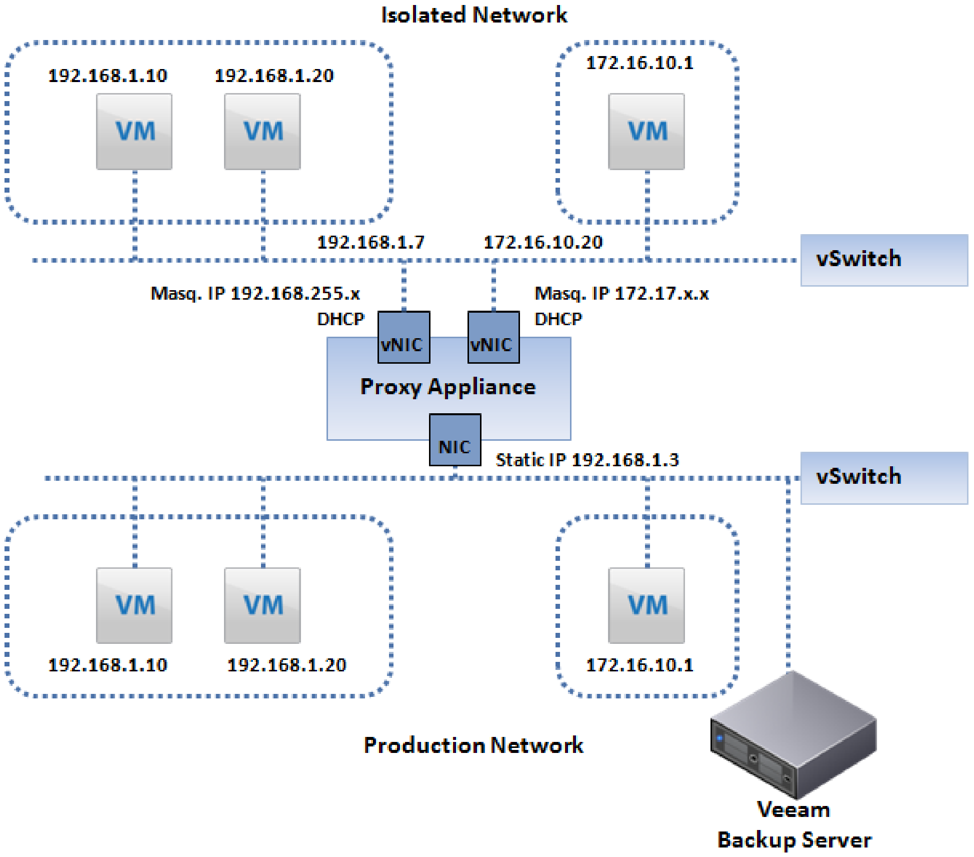

To ensure correct work of applications, VMs in isolated networks are run with the same IP addresses as in the production network. To avoid IP address conflicts between VMs in production and isolated networks, Veeam Backup & Replication 5.0 uses IP masquerading.

For each isolated logical network, Veeam Backup & Replication 5.0 assigns a masquerade IP address, and adds a new route to the IP routing table on the Veeam Backup & Replication console, where a proxy appliance is specified as a gateway to access VMs in this network.

For example, when trying to access a VM with IP address 172.16.10.1 in the isolated network, Veeam Backup & Replication 5.0 will send a request to the masquerade IP address 172.17.10.1. According to the routing rule added to the IP routing table, all requests will first be sent to the next hop — the proxy appliance, which will then resolve the IP address and forward the request to a necessary VM in the isolated network – in our case, 172.16.10.1.

Sometimes it is necessary to provide many clients with access to a restored VM, which is especially the case for user-directed U-AIR restores. For example, you may want to provide access to a backup copy of the Exchange server for employees using Web based access (Outlook Web Access). In this situation, it is impossible to update the routing table on each client machine. Veeam Backup & Replication 5.0 allows you to get access to a VM in the isolated network directly from a production environment.

To get access to a VM in the isolated network, you should reserve a static IP address in the pool of production IP addresses and specify which IP address of the VM powered on in the isolated environment it matches. This static IP address will be assigned to the appliance NIC connected to the production network. IP traffic directed to the specified static IP address will be routed by the appliance to the VM powered on in isolated network.

For example, to access a VM that has IP address 192.168.1.20 in the isolated network, you can reserve IP address 192.168.1.3 (in production) for it. You should also register an alias record in the production DNS server for the reserved IP address. For the example mentioned above, you can register alias *backup.exchange.local* for the IP address 192.168.1.3.

To learn more, see the Creating a Virtual Lab section.

Note: Veeam Backup & Replication 5.0 automatically adds a new route to the routing table on the server where it is installed. To be able to access VMs in isolated networks from other servers, you will have to manually modify the routing table using the *route add* command.

Recovery Verification Job

A recovery verification job aggregates all settings and policies of a recovery verification task, such as required application group, virtual lab to be used and backups of VMs that should be verified in the created environment.

When a recovery verification job runs, VMs from the application group are started from backups in the required order and remain running while VMs from verified backups are booted and tested. If Veeam Backup & Replication does not find a successful backup for any of VMs from the application group, the recovery verification job will fail.

Veeam Backup & Replication 5.0 then takes the latest backed up VMs from the selected backups, and depending on the specified job settings, starts, tests and verifies them one by one, or creates several streams and tests a number of VMs simultaneously. If Veeam Backup & Replication does not find a successful backup for any of verified VMs, verification of this VM will fail, but the job will continue to run.

Note: By default, you can start and test 3 VMs at the same time. You can also increase the number of VMs to be started and tested simultaneously; however if these VMs are resource-demanding, performance of the SureBackup job may decrease.

Once verification is complete, VMs from the application group are powered off. Optionally, you can leave the VMs from the application group running to perform manual testing or enable user-directed application item-level recovery.

To learn more, see the Creating a Recovery Verification Job section.

Creating an Application Group

To create a new application group:

- Right-click **Application Group** in the management tree and select **Create Application Group** from the context menu.
- Click **Application Group** under **SureBackup** in the management tree, right-click anywhere on blank area in the information pane and select **Create Application Group** from the context menu.

Step 1. Specify Application Group Name and Description

Enter a name and description for the new application group. The default description contains time at which the group was created and user who created it.

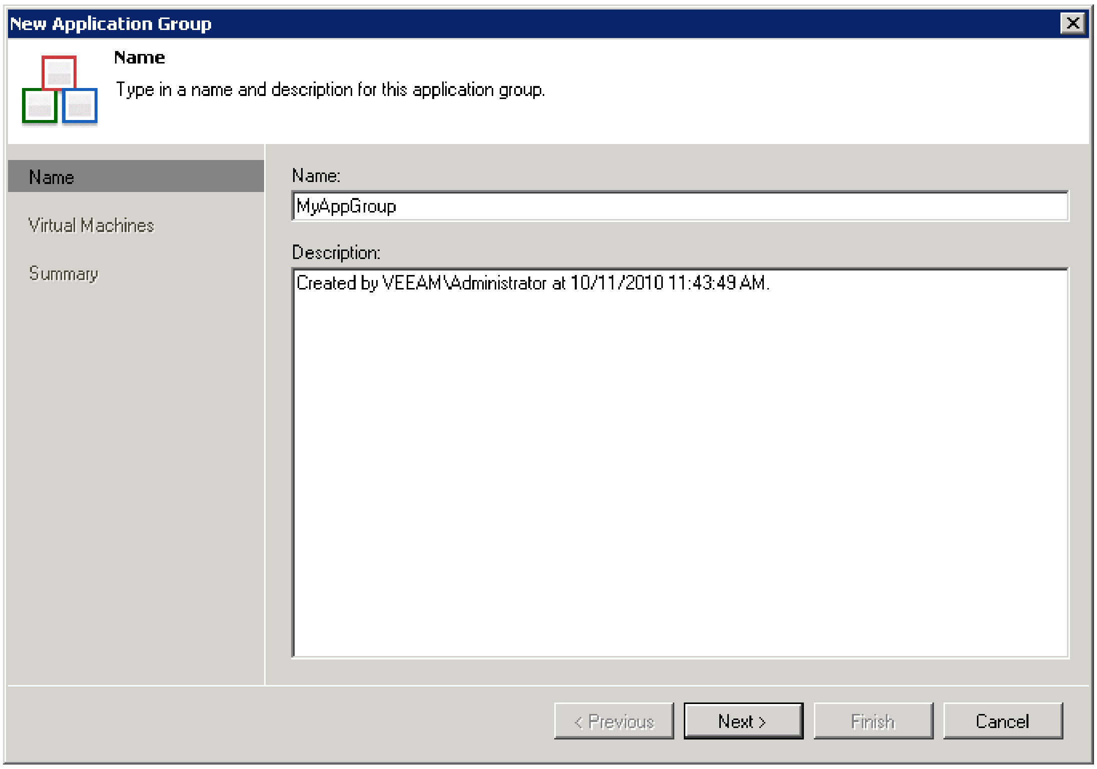

Name: MyAppGroup

Description: Created by VEEAM\Administrator at 10/11/2010 11:43:49 AM.

Step 2. Select Virtual Machines

To add a VM to the group, click **Add VM** and select where to browse for the machines:

- **From VI** — browse the VI environment. As VMs from the application group are started from backups, by the time you are planning to run a SureBackup job you need to make sure that VMs you selected have been successfully backed up at least once.
- **From Backups** – browse existing backups

VMs in the list are specified in the order of their boot priority. To move a VM up or down in the list, select it and click the **Move Up** or **Move Down** button.

To remove a VM from the list, select it and click the **Remove** button.

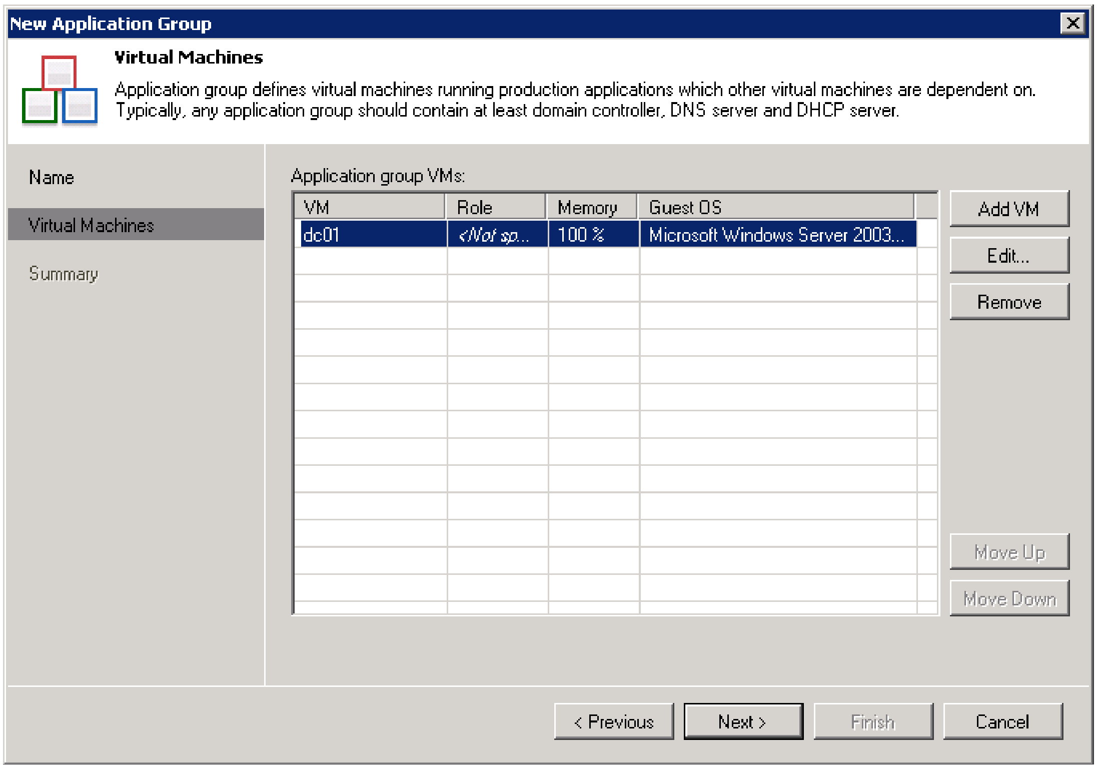

New Application Group

Virtual Machines

Application group defines virtual machines running production applications which other virtual machines are dependent on. Typically, any application group should contain at least domain controller, DNS server and DHCP server.

VM	Role	Memory	Guest OS
dc01	<Not sp...	100 %	Microsoft Windows Server 2003...

Step 3. Specify Recovery Verification Options and Tests

After you have added all VMs to the application group, you should specify a role, VM startup options and select tests to be performed for each VM.

Note: To be able to perform tests, Veeam Backup & Replication 5.0 requires VMware tools to be installed in a verified VM. If VMware tools are not installed, the VM will be started, but tests will not be performed.
VMs without VMware tools can still be used as auxiliary VMs that should be started to enable proper work of other VMs. In this case, you may leave check boxes tests not selected in role settings.

Select a necessary VM in the list and click the **Edit…** button on the right.

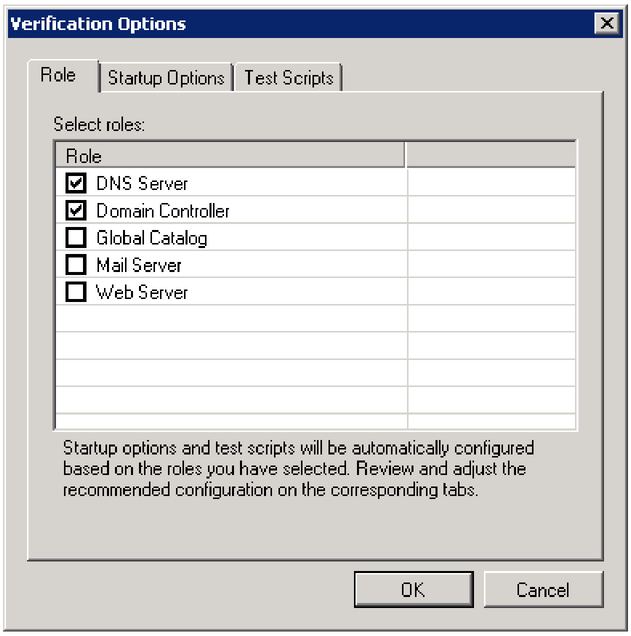

Verification Options

Role	
☑ DNS Server	
☑ Domain Controller	
□ Global Catalog	
□ Mail Server	
□ Web Server	

Startup options and test scripts will be automatically configured based on the roles you have selected. Review and adjust the recommended configuration on the corresponding tabs.

Role settings

On the **Role** tab, select the role that a VM performs. Veeam Backup & Replication offers the following predefined roles for VMs:

- DNS Server
- Domain Controller
- Global Catalog
- Mail Server
- Web Server

VM roles are described in .xml files stored in the *SbRoles* subfolder in the product installation folder. You can add your own roles by creating new .xml files and specifying roles settings and test scripts to be performed in them. To learn more, see the Creating XML files with VM Roles Description section.

Once you select a necessary role, Veeam Backup & Replication will automatically configure startup options and provide predefined test scripts applicable for the chosen role. You can use these settings or specify custom ones using the **Startup Options** and **Test Scripts** tabs.

To verify VMs that perform roles other than those specified in the list, you will have to manually configure startup options and specify test scripts to be used.

Startup Options

On the **Startup Options** tab, specify settings that should be used when a VM is started.

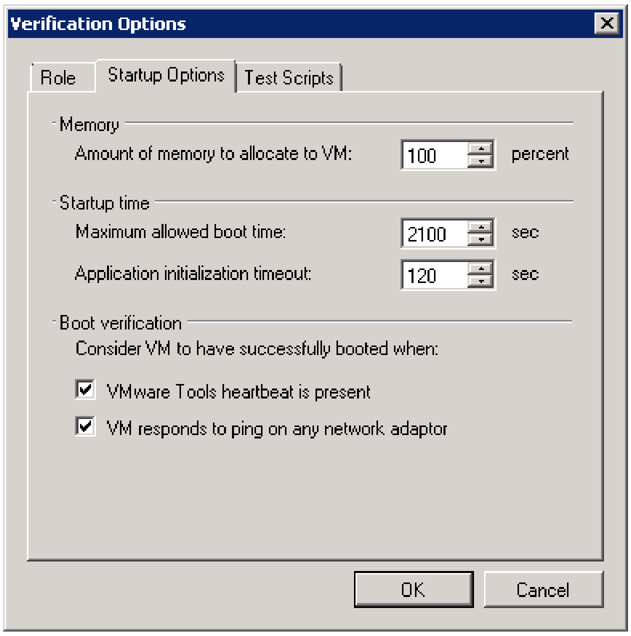

Verification Options

Role | Startup Options | Test Scripts

Memory

Amount of memory to allocate to VM: 100 percent

Startup time

Maximum allowed boot time: 2100 sec

Application initialization timeout: 120 sec

Boot verification

Consider VM to have successfully booted when:

☑ VMware Tools heartbeat is present

☑ VM responds to ping on any network adaptor

- In the **Memory** section, specify the amount of memory you want to pre-allocate to VM on the system boot (in percent). This will be the percentage of the memory level set for the corresponding production VM.
- In the **Startup time** section, specify an allowable timeout to initialize applications that will be started in a VM.
- In the **Boot verification** section, specify when a VM should be regarded to have been booted successfully: **VMware tools heartbeat is present** and **VM responds to ping on any network adapter**.

Note: Make sure that firewall on the tested VM allows ping requests.

Test Scripts

On the **Test Scripts** tab, click the **Add** button to specify what tests should be performed for a VM during verification.

- To use a predefined script to verify VMs of common roles, choose the **Use predefined test scripts** option and select a necessary script from the list.
- To provide a custom script, select the **Use the following script** option. Enter the name of the script to be run, path to an executable script file and arguments that should be passed to the script. You can use the following variables as arguments: *%vm_ip%* - IP address of a virtual lab VM and *%vm_fqdn%* — a fully qualified domain name of a virtual lab VM.

To edit a script, select it in the list and click the **Edit** button. To delete a script, select it in the list and click the **Delete** button.

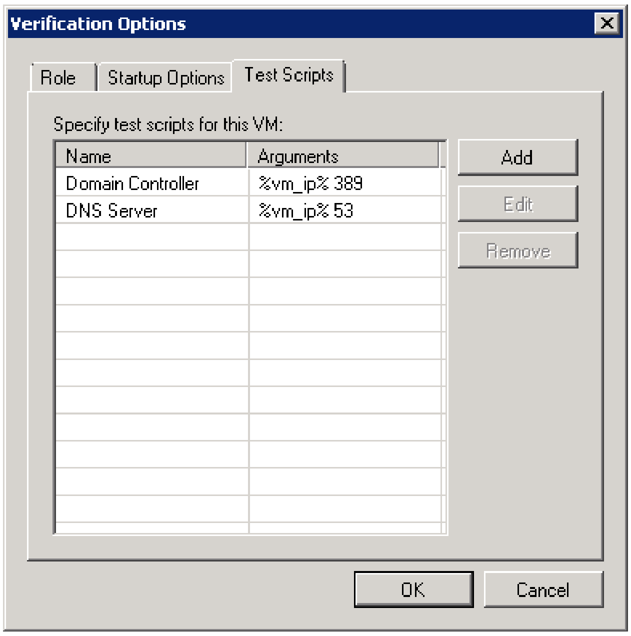

Specify test scripts for this VM:

Name	Arguments
Domain Controller	%vm_ip% 389
DNS Server	%vm_ip% 53

Note: If a VM performs several roles running a number of applications at once, you can verify their work by adding several verification scripts. For such VMs, it is recommended to specify maximum startup timeouts and allocate the greatest amount of memory.

Step 4. Review the Application Group Summary and Finish Working with Wizard

After the group is created, review the application group summary and click **Finish** to exit the wizard.

Creating a Virtual Lab

When setting up a virtual lab, you should select an ESX host on which it should be created, a datastore to hold redo logs and files of the proxy appliance, and specify settings for a proxy appliance and isolated networks.

To create a new virtual lab, you have to start the **New Virtual Lab** wizard:

- Right-click the **Virtual Lab** item in the menu on the left and select **Create Virtual Lab** in the context menu.

Step 1. Specify Name and Description

Enter a name and description for the new virtual lab. The default description contains time at which the lab was created and user who created it.

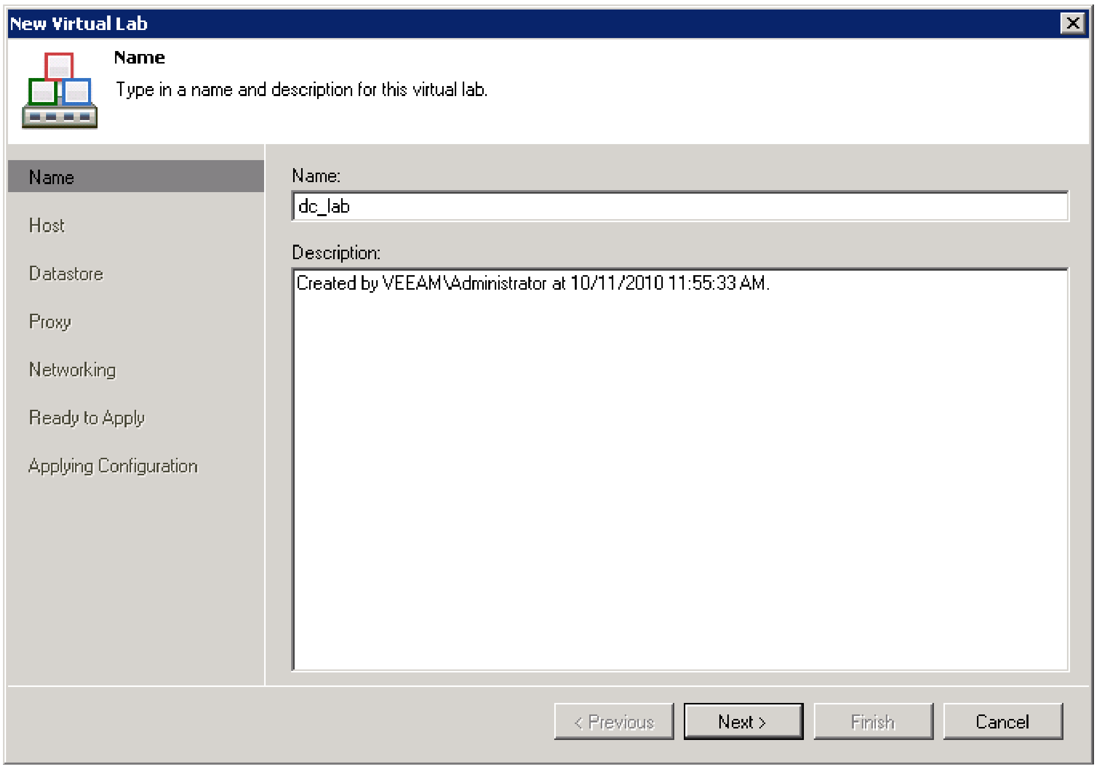

Name: dc_lab

Description: Created by VEEAM\Administrator at 10/11/2010 11:55:33 AM.

Step 2. Select a Host

Click the **Choose…** button to select an ESX/ESXi host on which the new virtual lab will be created. You can select a standalone ESX/ESXi host or the one being a part of a cluster.

Note: If you want to create a virtual lab on the ESX/ESXi server being a part of the vCenter hierarchy, make sure that this vCenter server is added to the Veeam Backup & Replication console. If such ESX/ESXi server is added as a standalone host, a virtual lab will not be created on it.

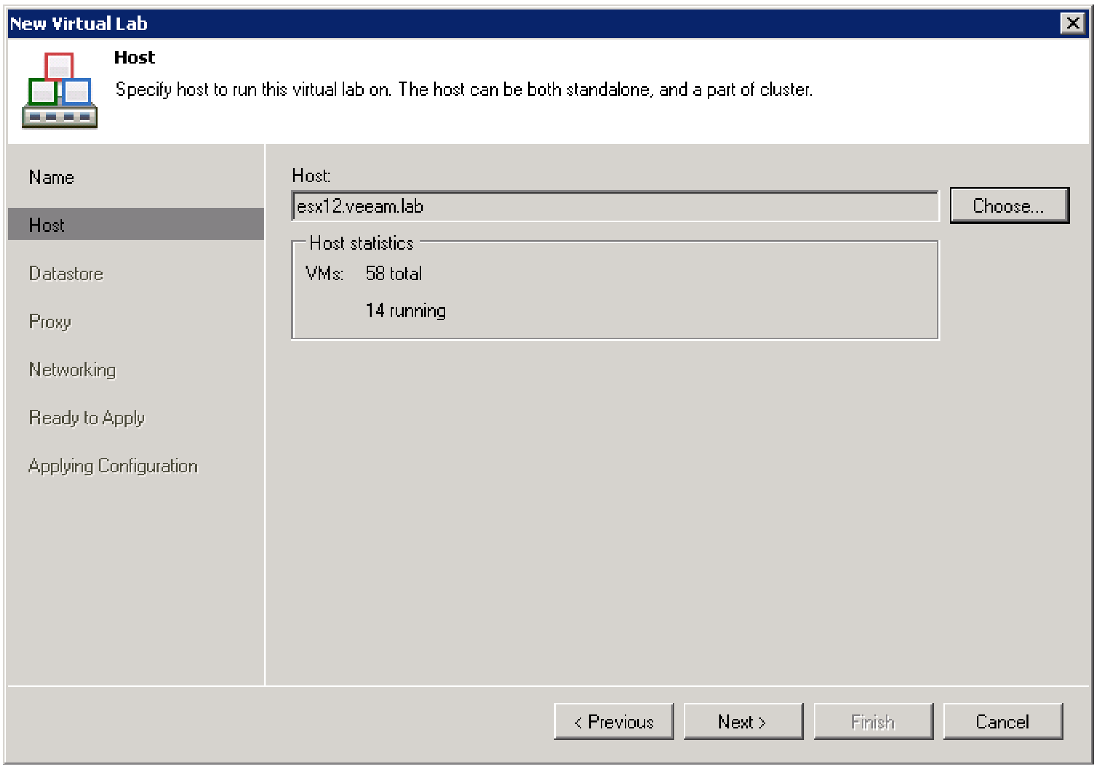

Host: esx12.veeam.lab

Host statistics
VMs: 58 total
14 running

Step 3. Select Datastore

Click **Choose** to select a datastore on which redo logs for tested VMs should be stored. Redo logs are auxiliary files used to store all changes that take place when a VM is run from a read-only backup. As soon as a recovery verification jobs completes, redo logs are deleted.

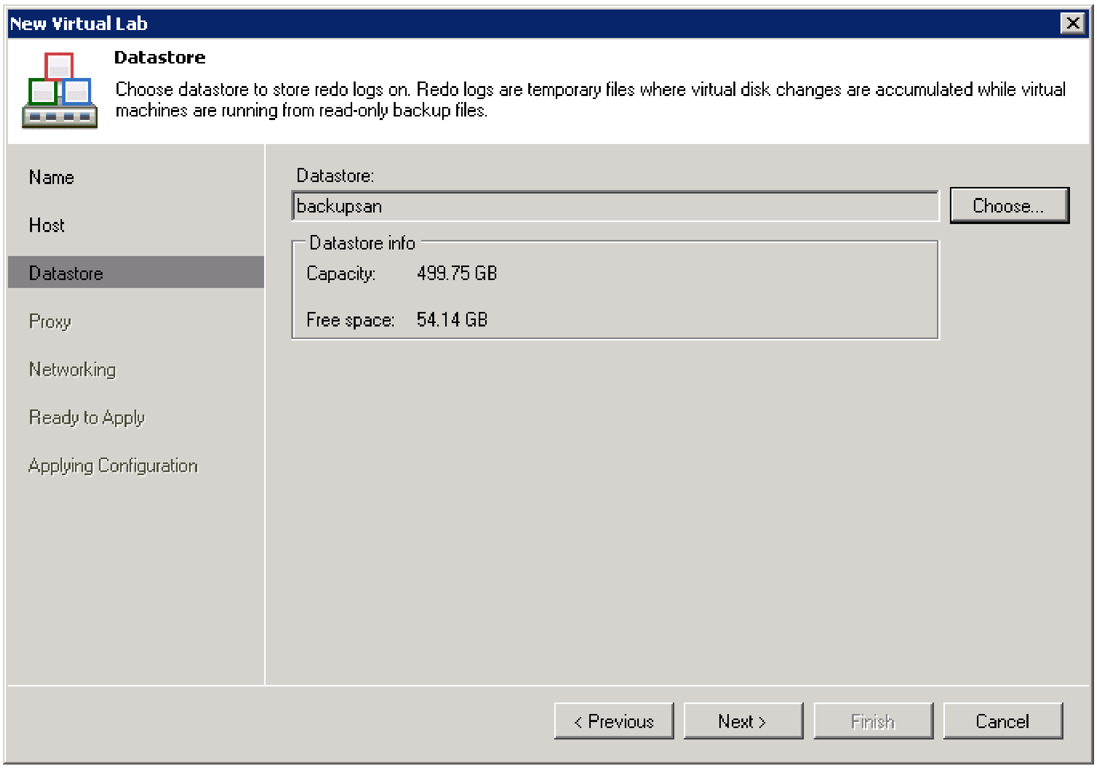

New Virtual Lab

Datastore

Choose datastore to store redo logs on. Redo logs are temporary files where virtual disk changes are accumulated while virtual machines are running from read-only backup files.

Datastore: backupsan

Datastore info

Capacity: 499.75 GB

Free space: 54.14 GB

Step 4. Set Up Proxy Appliance

To enable automatic recovery verification of VMs, select the **Use proxy appliance in this virtual lab** check box. The proxy appliance acts as a gateway that provides access from Veeam Backup Server to VMs running in the isolated virtual lab. If you do not select this check box, you will only be able to verify VMs and perform item-level restore using built-in temporary VM console in Veeam Backup & Replication, or using vSphere Client, and perform heartbeat tests.

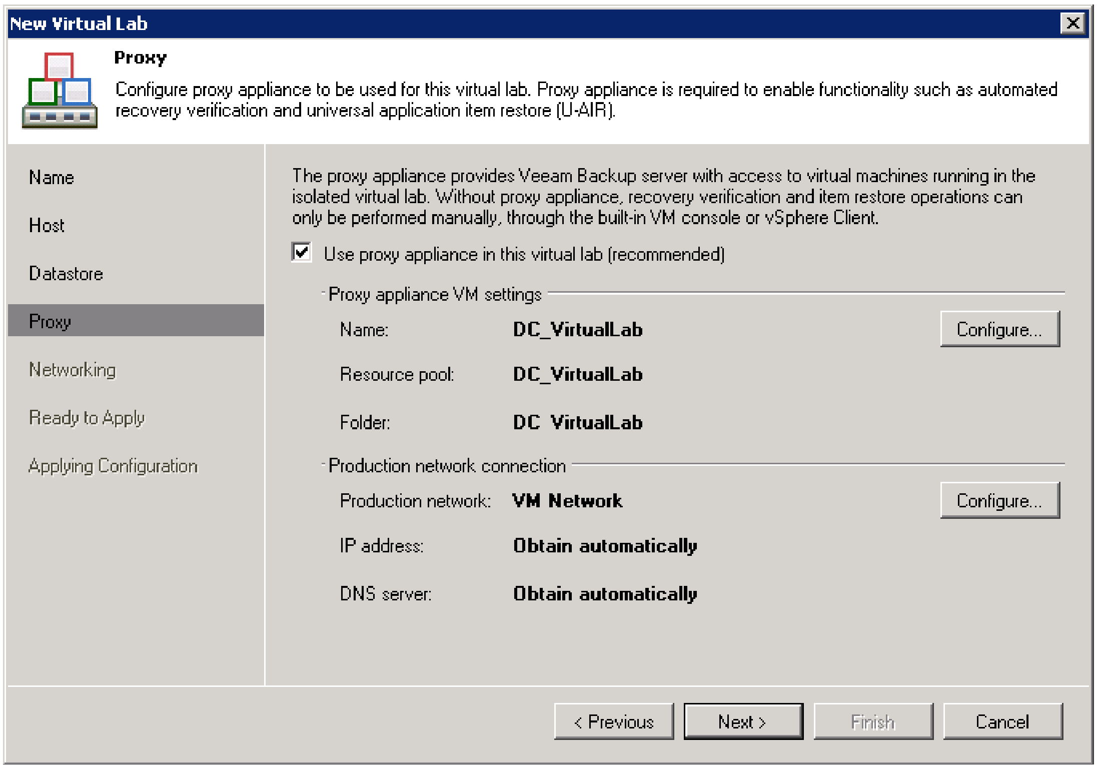

New Virtual Lab

Proxy

Configure proxy appliance to be used for this virtual lab. Proxy appliance is required to enable functionality such as automated recovery verification and universal application item restore (U-AIR).

The proxy appliance provides Veeam Backup server with access to virtual machines running in the isolated virtual lab. Without proxy appliance, recovery verification and item restore operations can only be performed manually, through the built-in VM console or vSphere Client.

☑ Use proxy appliance in this virtual lab (recommended)

Proxy appliance VM settings

Name: **DC_VirtualLab**

Resource pool: **DC_VirtualLab**

Folder: **DC_VirtualLab**

Production network connection

Production network: **VM Network**

IP address: **Obtain automatically**

DNS server: **Obtain automatically**

The virtual proxy is created and started in a new resource pool where all tested VMs are run during recovery verification process. Click the **Configure…** button in the **Proxy appliance VM settings** section and specify the name of the created virtual appliance, resource pool and folder where it should be created. By default, the name of a virtual lab is used.

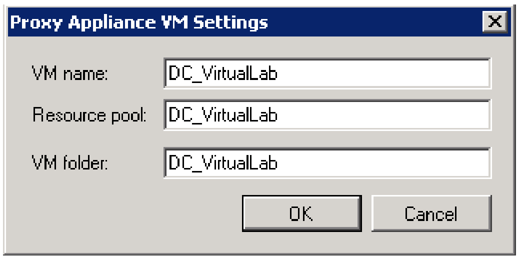

Proxy Appliance VM Settings

VM name: DC_VirtualLab

Resource pool: DC_VirtualLab

VM folder: DC_VirtualLab

Click the **Configure…** button in the **Production network connection** section to select a network where the proxy appliance should be created, specify its IP address and settings of DNS server to be used. You can choose to automatically obtain IP address for the proxy appliance and DNS server, or set them manually.

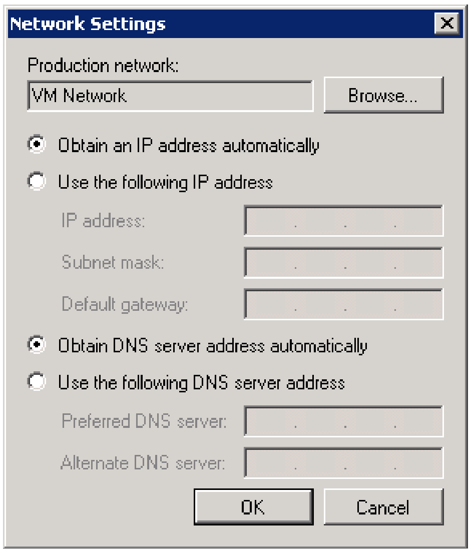

Network Settings

Production network: VM Network

● Obtain an IP address automatically

○ Use the following IP address

IP address: . . .

Subnet mask: . . .

Default gateway: . . .

● Obtain DNS server address automatically

○ Use the following DNS server address

Preferred DNS server: . . .

Alternate DNS server: . . .

Important! If you assign a proxy appliance an IP address from the same network where the Veeam Backup & Replication server is located, Veeam Backup & Replication will automatically add a new route to the routing table on the Veeam Backup & Replication server. If you assign a proxy appliance an IP address from the network other than that where the Veeam Backup & Replication server is located, you will have to manually add a new route to the routing table on the router in the production network. Otherwise you will not be able to access virtual machines in isolated networks.

Step 5. Select the Networking Mode

Select the type of network settings configuration. Veeam Backup & Replication offers two types of networking for the created virtual lab:

- **Basic** — this type of networking is recommended if you have only one production network, and the Veeam Backup & Replication server is located in that network. Veeam Backup & Replication will use parameters of this network to automatically configure an isolated network to verify tested VMs.
- **Advanced** — this type of networking is recommended if you are planning to verify VMs that have dependencies on other VMs located in different networks. In this case, you will have to configure network parameters these isolated networks manually.

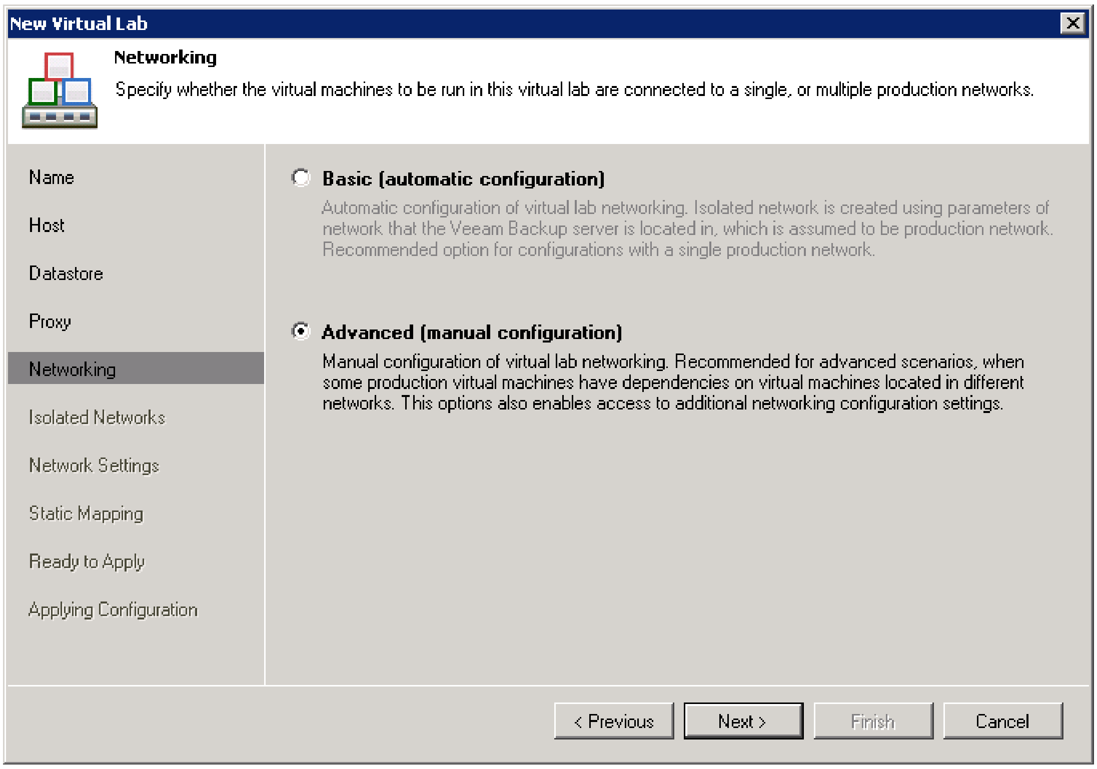

New Virtual Lab

Networking

Specify whether the virtual machines to be run in this virtual lab are connected to a single, or multiple production networks.

○ **Basic (automatic configuration)**

Automatic configuration of virtual lab networking. Isolated network is created using parameters of network that the Veeam Backup server is located in, which is assumed to be production network. Recommended option for configurations with a single production network.

● **Advanced (manual configuration)**

Manual configuration of virtual lab networking. Recommended for advanced scenarios, when some production virtual machines have dependencies on virtual machines located in different networks. This options also enables access to additional networking configuration settings.

Step 6. Specify Isolated Networks

This step is available if you have selected the **Advanced networking** option at the **Networking** step of the wizard.

At this step of the wizard, you should create isolated networks where verified VMs should be started, and map them to production networks where these VM are located.

To add a network, click the **Add..** button and select a production network in which a VM from the application group or a verified VM resides. Then, specify a name for an isolated network that should be mapped to this production network, and enter an identifier for the created virtual network.

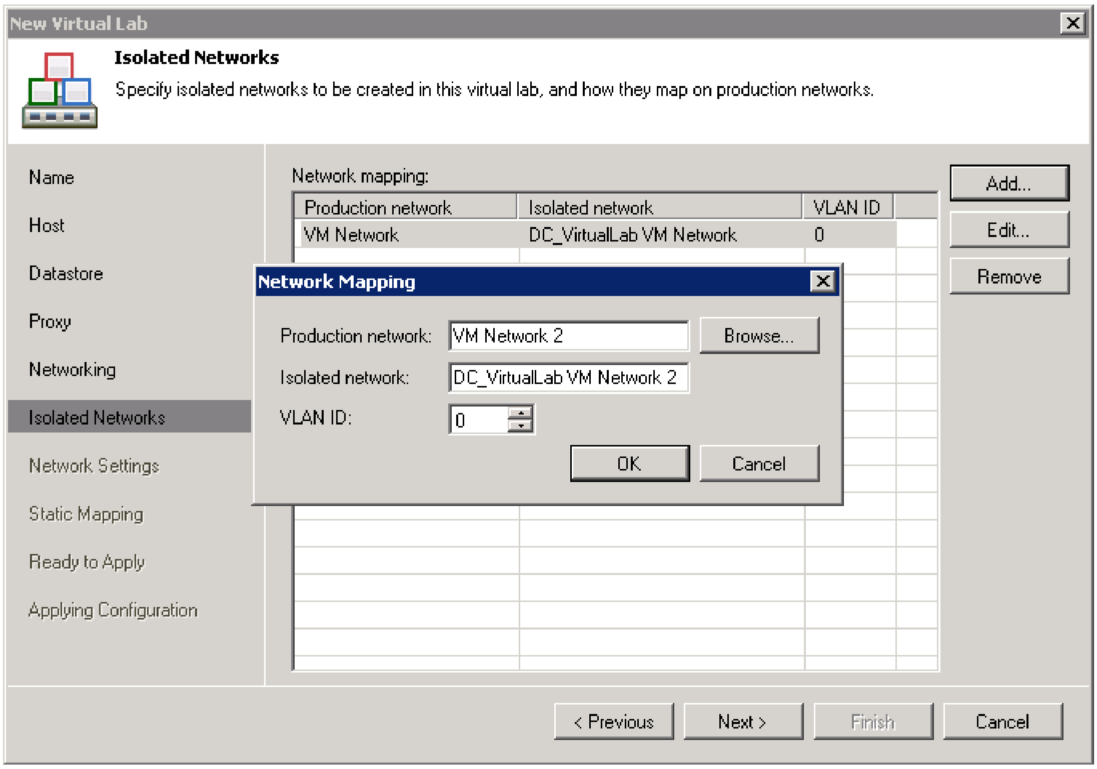

New Virtual Lab

Isolated Networks

Specify isolated networks to be created in this virtual lab, and how they map on production networks.

Network mapping:

Production network	Isolated network	VLAN ID
VM Network	DC_VirtualLab VM Network	0

Network Mapping

Production network: VM Network 2

Isolated network: DC_VirtualLab VM Network 2

VLAN ID: 0

Step 7. Specify Network Settings

This step is available if you have selected the **Advanced networking** option at the **Networking** step of the wizard.

At this step of the wizard, you should specify settings for every created isolated networks and how a proxy appliance should connect the production network to these networks.

Communication between the production network and an isolated network is carried out through the vNIC adapter that is added to the proxy appliance. A vNIC adapter is added for each isolated network.

To add an adapter, click the **Add…** button and specify its connection settings.

Select the network to which you want this adapter to be connected. Specify the IP address that the proxy appliance should have in this isolated network, and the subnet mask. Typically, the IP address should coincide with the gateway IP address in the production network.

Note: Network addresses for different adapters should be different. For example, if the first adapter has address 192.168.0.1 with mask 255.255.255.0, and the second one – 192.168.0.2 with mask 255.255.255.0, such configuration will not be supported.

Once you specify the IP address, Veeam Backup & Replication will automatically configure a masquerade IP address for accessing VMs running in the virtual lab through the production network.

Select the **Enable DHCP service on this adaptor** check box and specify settings of a virtualized DNS server if necessary.

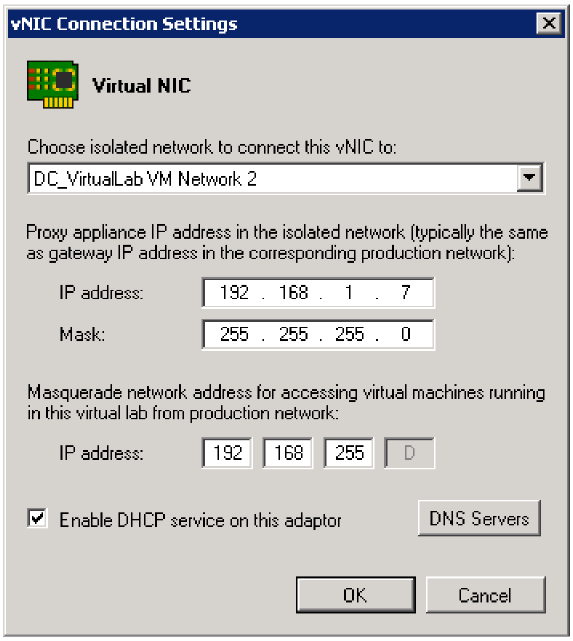

vNIC Connection Settings

Virtual NIC

Choose isolated network to connect this vNIC to: DC_VirtualLab VM Network 2

Proxy appliance IP address in the isolated network (typically the same as gateway IP address in the corresponding production network):

IP address: 192 . 168 . 1 . 7

Mask: 255 . 255 . 255 . 0

Masquerade network address for accessing virtual machines running in this virtual lab from production network:

IP address: 192 168 255 D

☑ Enable DHCP service on this adaptor

Click the **Route network traffic between vNICs** check box to enable communication between isolated networks. When you select this option, make sure that the IP address of the proxy appliance in the isolated network matches the IP address of a proxy appliance in the production network.

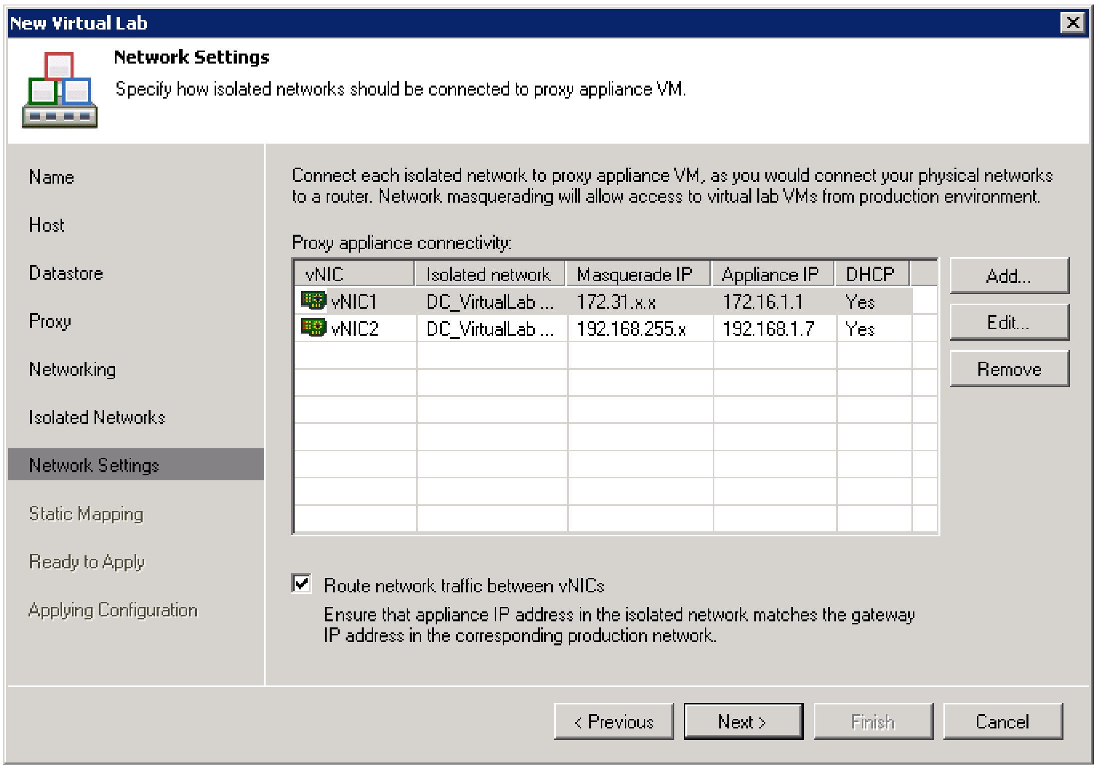

New Virtual Lab

Network Settings

Specify how isolated networks should be connected to proxy appliance VM.

Connect each isolated network to proxy appliance VM, as you would connect your physical networks to a router. Network masquerading will allow access to virtual lab VMs from production environment.

Proxy appliance connectivity:

vNIC	Isolated network	Masquerade IP	Appliance IP	DHCP
vNIC1	DC_VirtualLab ...	172.31.x.x	172.16.1.1	Yes
vNIC2	DC_VirtualLab ...	192.168.255.x	192.168.1.7	Yes

☑ Route network traffic between vNICs

Ensure that appliance IP address in the isolated network matches the gateway IP address in the corresponding production network.

Step 8. Specify Static IP Mapping

At this step of the wizard, you can specify static IP address mapping rules to make VMs in the virtual lab accessible from any computer in the production network.

To add a new static IP relation, click the **Add** button. In the IP relation window, specify an IP address of a VM in the production network, and its masquerade IP – a free IP address from the production network that will be used to access it in the isolated network from the production environment.

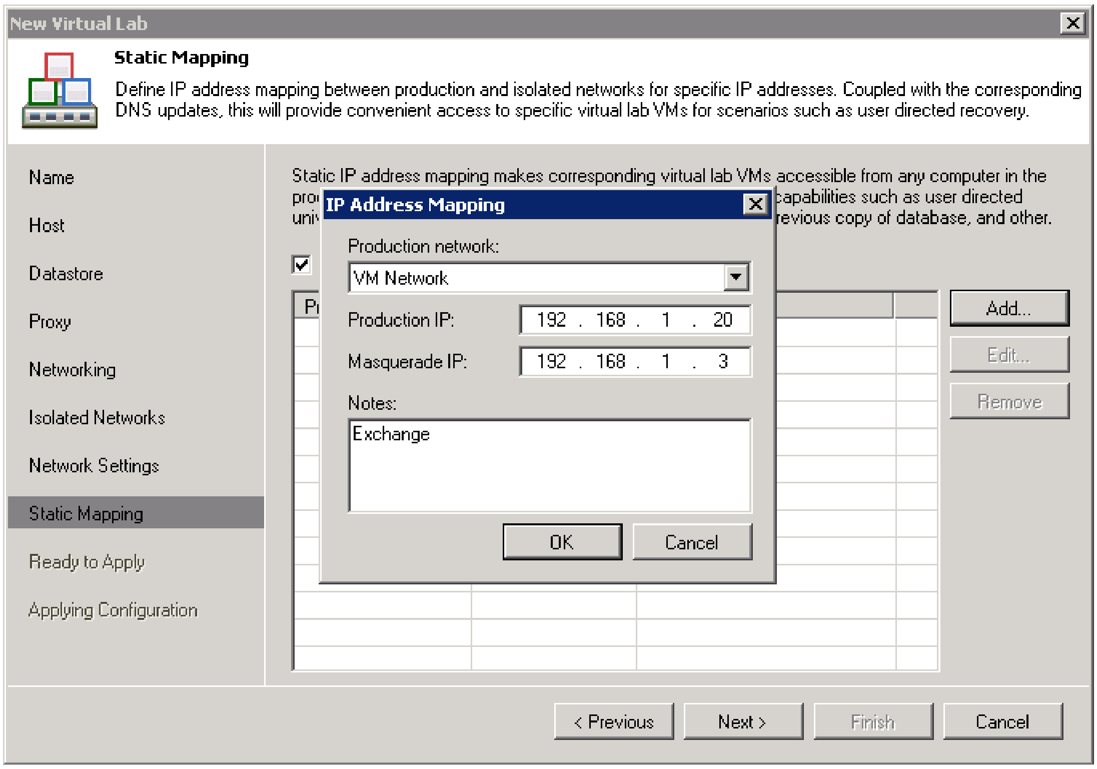

New Virtual Lab

Static Mapping

Define IP address mapping between production and isolated networks for specific IP addresses. Coupled with the corresponding DNS updates, this will provide convenient access to specific virtual lab VMs for scenarios such as user directed recovery.

Static IP address mapping makes corresponding virtual lab VMs accessible from any computer in the pro… capabilities such as user directed uni… revious copy of database, and other.

IP Address Mapping

Production network: VM Network

Production IP: 192 . 168 . 1 . 20

Masquerade IP: 192 . 168 . 1 . 3

Notes: Exchange

Step 9. Apply Parameters

Review the parameters of the virtual lab which will be created. You can go back to any previous step to adjust the parameters. If everything is fine, click **Next** to create the virtual lab.

Important! Use Veeam Backup & Replication 5.0 to modify or delete a virtual lab. If you change lab settings or delete any of its components from outside (for example, using vSphere Client), the lab will be corrupted and its component such as created vSwitch, resource pool and so on will remain in the virtual infrastructure.

Creating a Recovery Verification Job

To create a new recovery verification job:

- Right-click **Jobs** under **SureBackup** in the management tree and select **Create SureBackup Job** from the context menu.
- Click **Jobs** under **SureBackup** in the management tree, right-click anywhere on blank area in the working area and select **Create SureBackup Job** from the context menu.

Step 1. Specify Name and Description

Enter a name and description for the new recovery verification job. The default description contains time at which the job was created and user who created it.

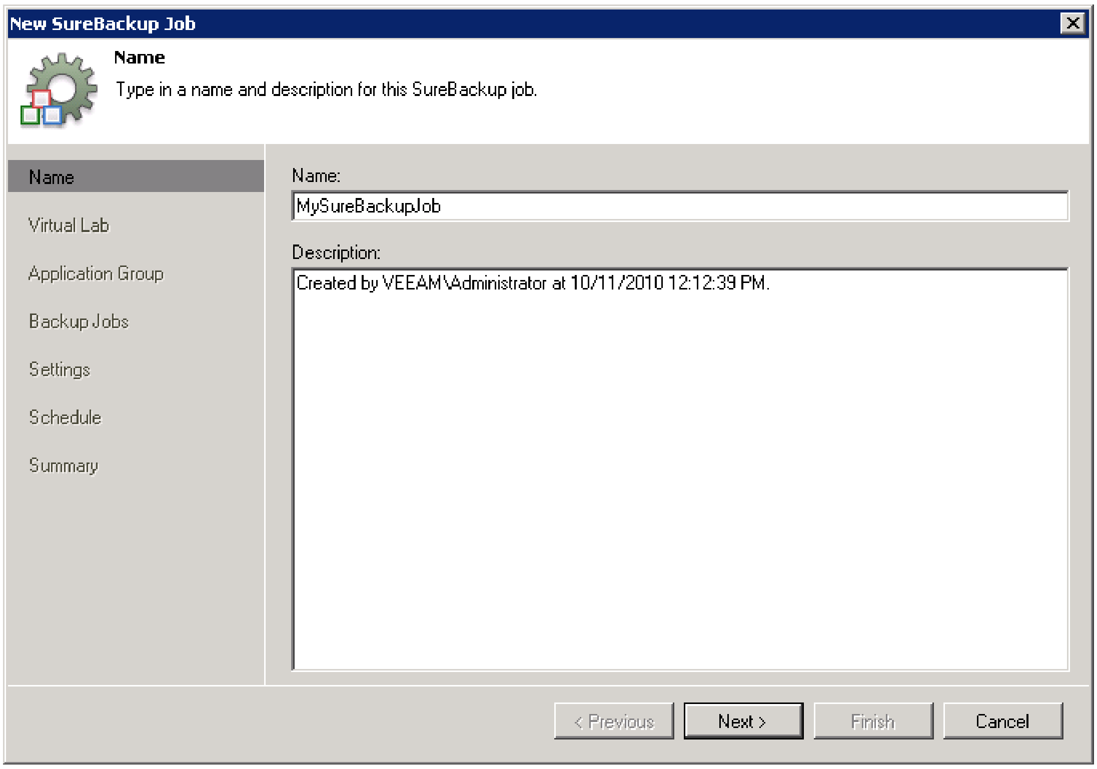

New SureBackup Job

Name

Type in a name and description for this SureBackup job.

Name:

MySureBackupJob

Description:

Created by VEEAM\Administrator at 10/11/2010 12:12:39 PM.

Step 2. Select a Virtual Lab

From the **Virtual lab** list, select one of existing virtual labs in which recovery verification should be performed. Information about the selected virtual lab will be displayed in the **Virtual lab info** section.

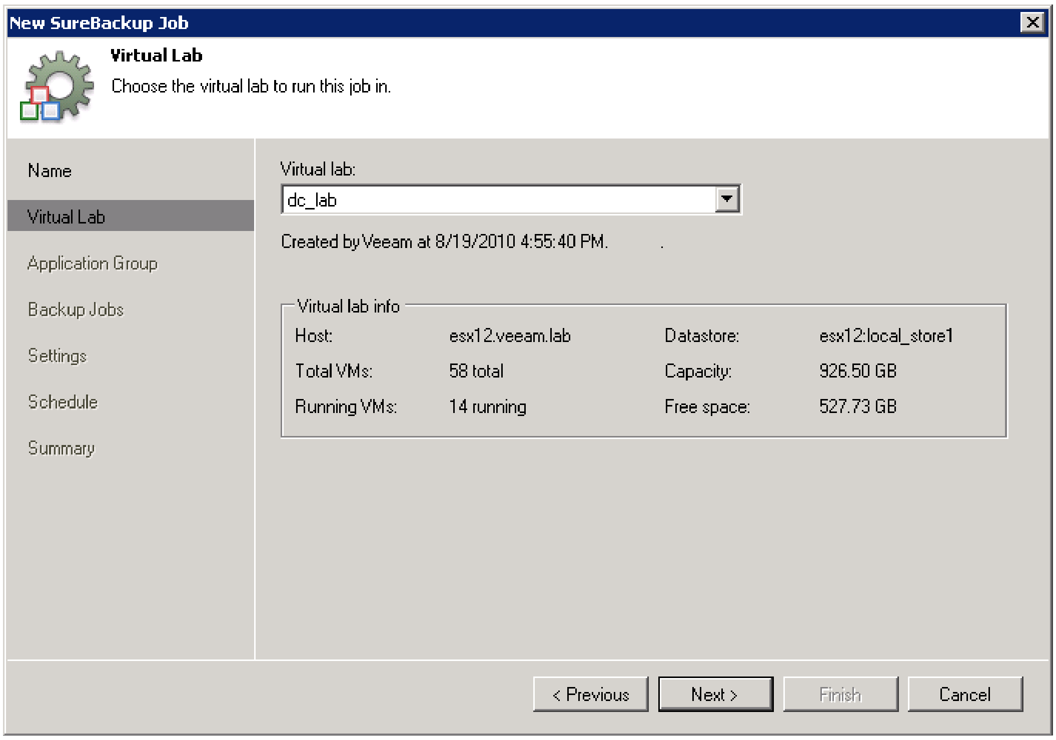

New SureBackup Job

Virtual Lab

Choose the virtual lab to run this job in.

Virtual lab: dc_lab

Created by Veeam at 8/19/2010 4:55:40 PM.

Virtual lab info

Host:	esx12.veeam.lab	Datastore:	esx12:local_store1
Total VMs:	58 total	Capacity:	926.50 GB
Running VMs:	14 running	Free space:	527.73 GB

Step 3. Select an Application Group

From the **Application group** list, select the application group containing all components and services required to perform recovery verification of VMs you want to test. Refer to the **Backup Status** column in the **Application group info** list to make sure that the backups of VMs in this group are available.

You can either select an application group or skip this step. If the application group is not selected, you must link a backup job to the created SureBackup job at the next step. In this case, Veeam Backup & Replication 5.0 will only start and verify VMs from the linked backup job when the SureBackup job is run.

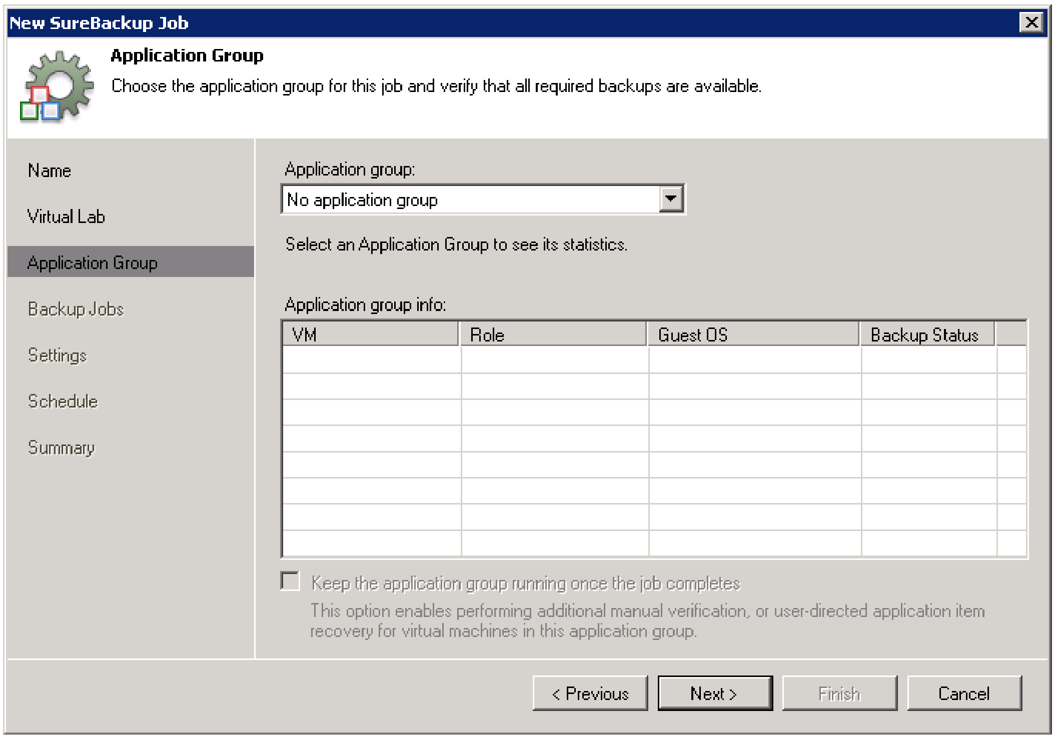

New SureBackup Job

Application Group

Choose the application group for this job and verify that all required backups are available.

Application group: No application group

Select an Application Group to see its statistics.

Application group info:

VM	Role	Guest OS	Backup Status

□ Keep the application group running once the job completes

This option enables performing additional manual verification, or user-directed application item recovery for virtual machines in this application group.

Select the **Keep the application group running once the job completes** check box to leave VMs from the application group running after the recovery verification job is finished. This option lets you manually test these VMs and perform item-level restore with U-AIR. If you select this check box, the lab will not be powered off when the SureBackup job completes, and Veeam Backup & Replication 5.0 will not have to start it again to perform U-AIR procedures.

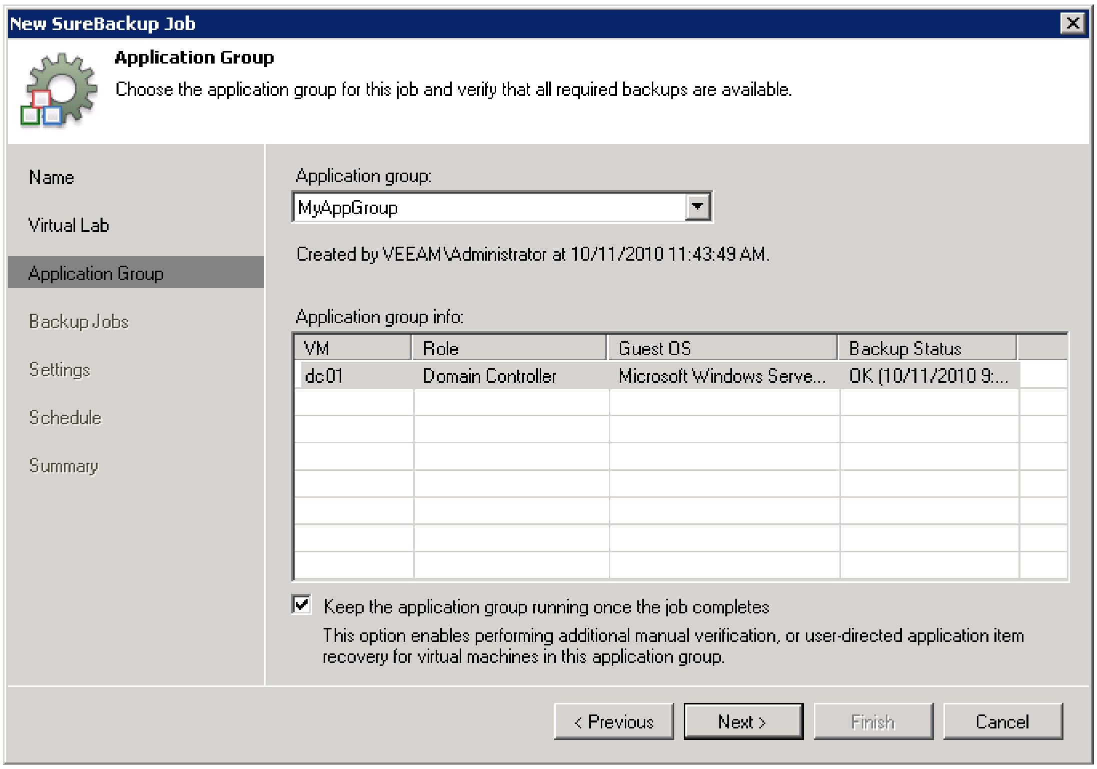

Application Group

Choose the application group for this job and verify that all required backups are available.

Application group: MyAppGroup

Created by VEEAM\Administrator at 10/11/2010 11:43:49 AM.

Application group info:

VM	Role	Guest OS	Backup Status
dc01	Domain Controller	Microsoft Windows Serve...	OK (10/11/2010 9:...

☑ Keep the application group running once the job completes

This option enables performing additional manual verification, or user-directed application item recovery for virtual machines in this application group.

Step 4. Link a SureBackup Job to Backup Job(s)

At this step of the wizard, you should select VM backups that you want to verify with the created recovery verification job. Once you run a recovery verification job, Veeam Backup & Replication will start VMs from the application group in the required order, and then boot and process VMs from the selected backup one by one.

Select the **Link backup jobs** check box. Click **Add** and select necessary backup jobs in the **Job Browser** window. If this check box is not selected, Veeam Backup & Replication 5.0 will only start and verify VMs from the selected application group.

In the **Process simultaneously up to … VMs** field, specify the maximum number of VMs that can be started at the same time. For example, if you select to start three VMs at the same time, Veeam Backup & Replication 5.0 will create three streams in which each VM will be started. Once a VM is verified and powered off, the next VM will be started in the stream. After all VMs are verified, the application group will also be powered off or keep running if corresponding settings are specified in job options.

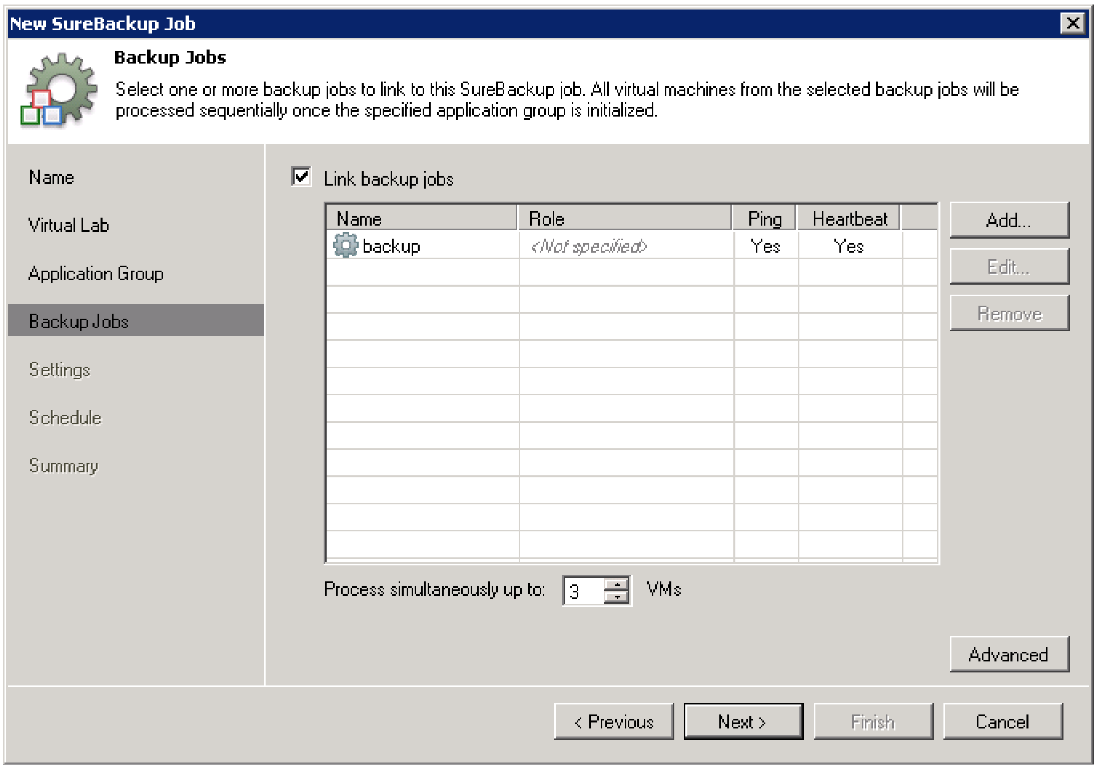

Select one or more backup jobs to link to this SureBackup job. All virtual machines from the selected backup jobs will be processed sequentially once the specified application group is initialized.

☑ Link backup jobs

Name	Role	Ping	Heartbeat
backup	<Not specified>	Yes	Yes

Process simultaneously up to: 3 VMs

To remove a backup job from the list, select it and click **Remove**.

Step 5. Specify Recovery Verification Options and Tests

After you add a backup to be verified, you should define roles, specify startup options and select tests to be performed for VMs in the selected backup. If a backup you want to verify contains VMs performing one role, you can specify recovery verification settings for the whole VM backup in bulk. Alternatively, you can granularly set recovery verification options and select tests to be used for each VM in the backup.

- To specify recovery verification options for the whole VM backup, select a backup in the list and click the **Edit…** button on the right.
- To specify recovery verification options for each VM in the backup separately, select a backup in the list and click the **Advanced…** button on the right. Then click **Add** and select a necessary VM in the **Add Object** window. Select the added VM in the list, click the **Edit** button and specify recovery verification settings as described below.

Note: To be able to perform tests, Veeam Backup & Replication 5.0 requires VMware tools to be installed in a verified VM. If VMware tools are not installed, the VM will be started, but tests will not be performed.
VMs without VMware tools can still be used as auxiliary VMs that should be started to enable proper work of other VMs. In this case, you may leave check boxes tests not selected in role settings.

Role settings

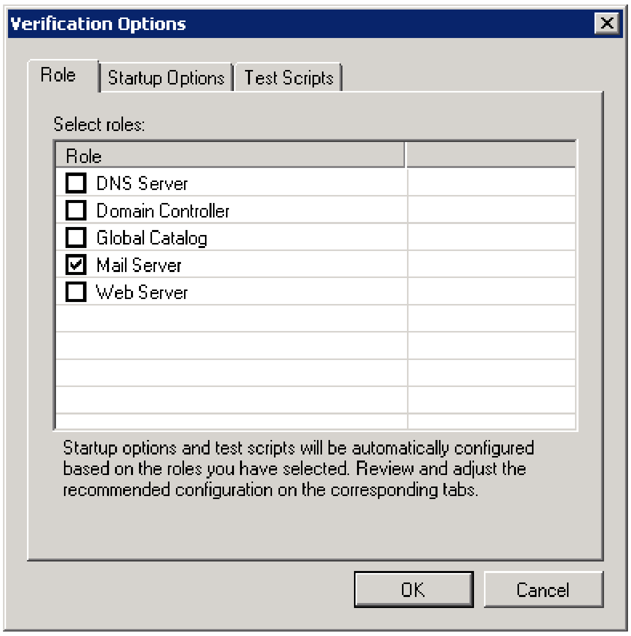

Select roles:

Role	
□ DNS Server	
□ Domain Controller	
□ Global Catalog	
☑ Mail Server	
□ Web Server	

Startup options and test scripts will be automatically configured based on the roles you have selected. Review and adjust the recommended configuration on the corresponding tabs.

On the **Role** tab, select the role that VMs in the backup perform. Veeam Backup & Replication offers the following predefined roles for VMs:

- DNS Server
- Domain Controller
- Global Catalog
- Mail Server
- Web Server

VM roles are described in .xml files stored in the *SbRoles* subfolder of the product installation folder. You can add your own roles by creating new .xml files and specifying roles settings and test scripts to be performed in them.

Once you select a necessary role, Veeam Backup & Replication will automatically configure startup options and provide predefined test scripts applicable for the chosen role. You can use these settings or specify custom ones using the **Startup Options** and **Test Scripts** tabs.

To verify backups with VMs that perform roles other than those specified in the list, you will have to manually configure startup options and specify test scripts to be used.

Startup Options

On the **Startup Options** tab, specify settings that should be used when a VM is started.

- In the **Memory** section, specify the amount of memory you want to pre-allocate to VM on the system boot (in percent). This will be the percentage of the memory level set for the corresponding production VM.
- In the **Startup time** section, specify an allowable timeout to initialize application that will be started in a VM.
- In the **Boot verification** section, specify when a VM should be regarded to have been booted successfully: **VMware tools heartbeat is present** and **VM responds to ping on any network adapter**.

Note: Make sure that firewall on the tested VM allows ping requests.

Test Scripts

On the **Test Scripts** tab, click the **Add** button to specify what tests should be performed for a VM during verification.

- To use a predefined script to verify VMs of common roles, choose the **Use predefined test scripts** option and select a necessary script from the list.
- To provide a custom script, select the **Use the following script** option. Enter the name of the script to be run, path to an executable script file and arguments that should be passed to the script. You can use the following variables as arguments: *%vm_ip%* - IP address of a virtual lab VM and *%vm_fqdn%* — a fully qualified domain name of a virtual lab VM.

To edit scripts, select it in the list and click the **Edit** button. To delete a script, select it in the list and click the **Delete** button.

Note: If a VM performs several roles running a number of applications at once, you can verify their work by adding several verification scripts. For such VMs, it is recommended to specify maximum startup timeouts and allocate the greatest amount of memory.

Step 6. Specify Additional Job Settings

Using the Notifications section, you can select to send notifications about the job result when a recovery verification job completes.

- Select the **Send SNMP trap** check box if you want to receive SNMP traps. SNMP traps will be sent if you configure SNMP settings in Veeam Backup & Replication and on the recipient's computer. To learn more, see the Specifying SNMP Settings section.
- Select the **Send email notifications to the following recipients** check box if you want to receive notifications by e–mail. In the field below, specify a recipient's e-mail address. You can enter several addresses separated by a semicolon.

  E–mail notifications will be sent only if you configure general e-mail notification settings in Veeam Backup & Replication. To learn more, see the Specifying E-Mail Notification Settings section.

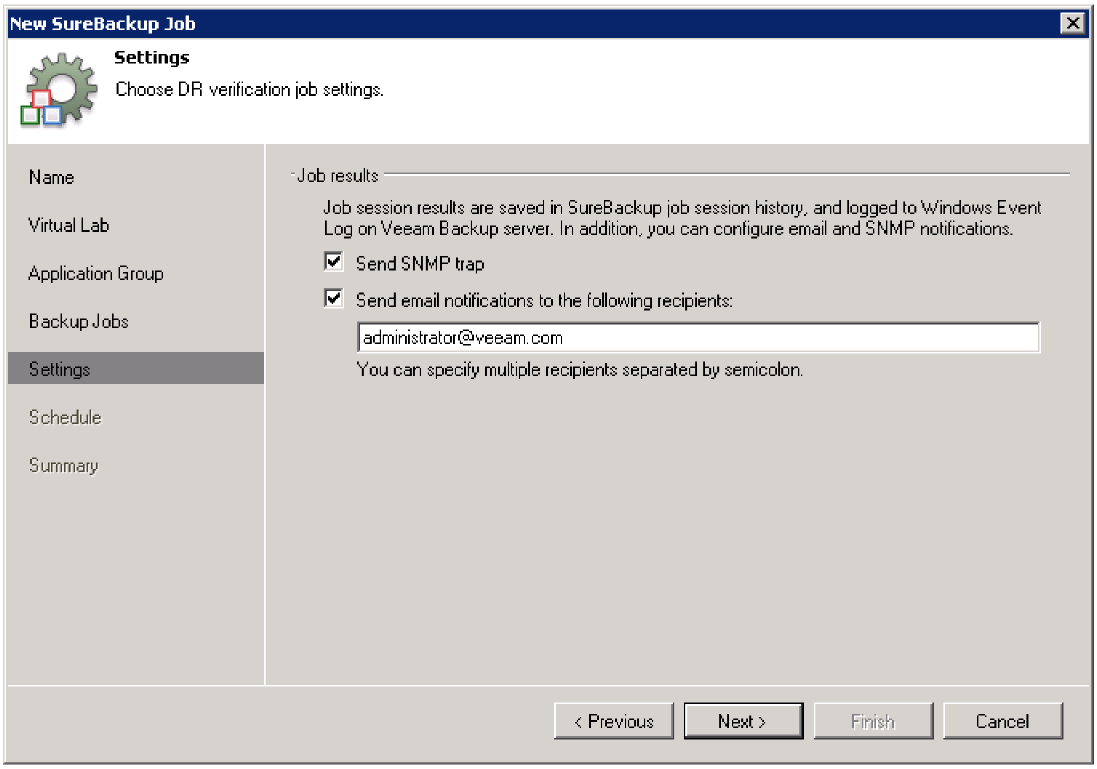

Job results

Job session results are saved in SureBackup job session history, and logged to Windows Event Log on Veeam Backup server. In addition, you can configure email and SNMP notifications.

☑ Send SNMP trap

☑ Send email notifications to the following recipients:

administrator@veeam.com

You can specify multiple recipients separated by semicolon.

Step 7. Specify Job Schedule

The **Schedule** step of the wizard allows you to choose to manually run the created job or schedule the recovery verification job for specific time — for example, after the corresponding backup job completes.

To specify the job schedule, select the **Run the job automatically** check box. If this check box is not selected, the job is supposed to be run manually.

You can choose to perform the job at specific time on defined week days or monthly.

To avoid interference with the corresponding backup job, select the **If some linked backup jobs are still running, wait up to … minutes** and define a necessary timeout.

Note: If a backup job linked to a SureBackup job is started while this SureBackup job is still running, the SureBackup job will be automatically stopped.

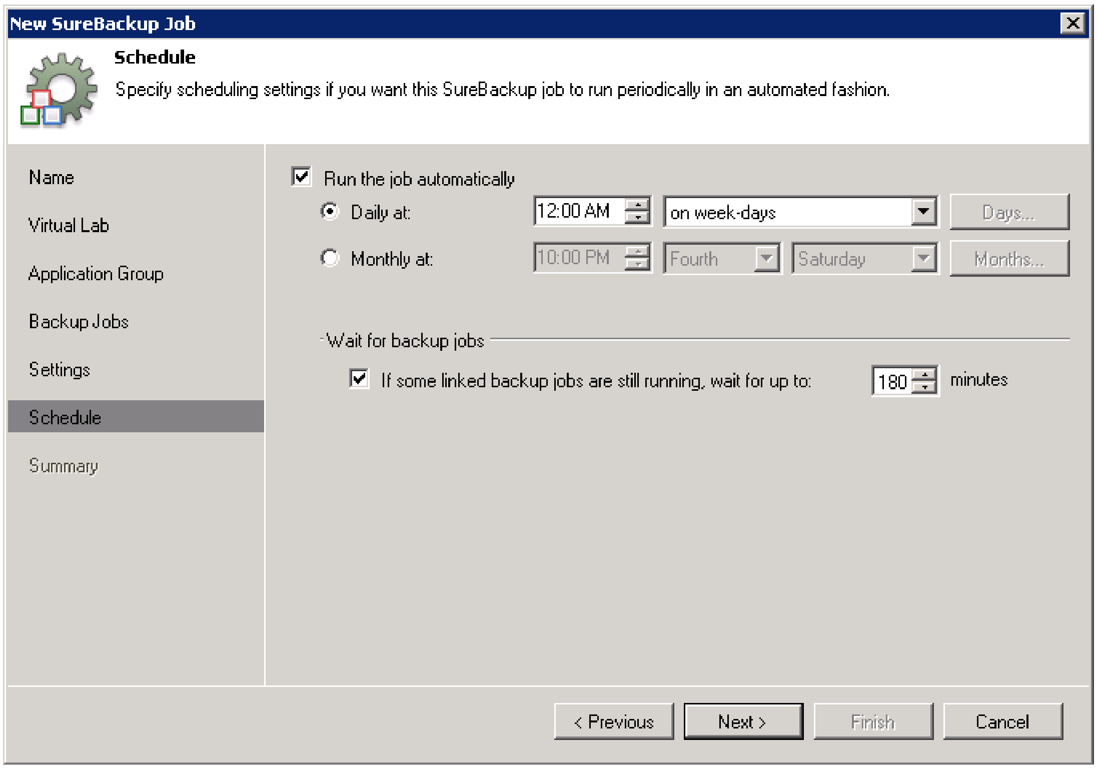

New SureBackup Job

Schedule

Specify scheduling settings if you want this SureBackup job to run periodically in an automated fashion.

☑ Run the job automatically

● Daily at: 12:00 AM on week-days Days...

○ Monthly at: 10:00 PM Fourth Saturday Months...

Wait for backup jobs

☑ If some linked backup jobs are still running, wait for up to: 180 minutes

Step 8. Review Job Summary and Finish Working with Wizard

Review the summary of the created recovery verification job. Select the **Run the job when I click Finish** check box to start the created job right after you finish working with the wizard; then click **Finish**.

Viewing Recovery Verification Job Statistics

When a recovery verification job is running, you can monitor how tests for verified VMs are performed and see their results in the real-time mode. To see the status of VM tests, right-click a necessary recovery verification job and select **Realtime Statistics** from the shortcut menu.

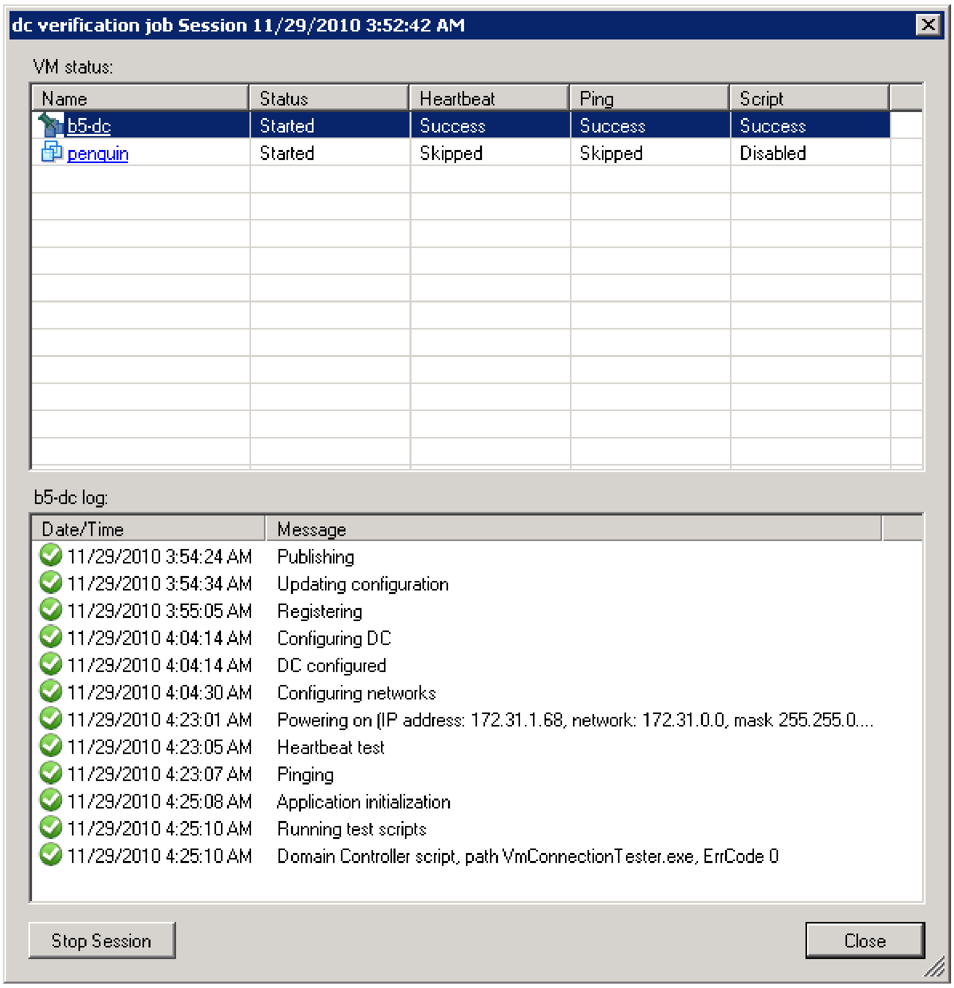

dc verification job Session 11/29/2010 3:52:42 AM

VM status:

Name	Status	Heartbeat	Ping	Script
b5-dc	Started	Success	Success	Success
penguin	Started	Skipped	Skipped	Disabled

b5-dc log:

Date/Time	Message
✓ 11/29/2010 3:54:24 AM	Publishing
✓ 11/29/2010 3:54:34 AM	Updating configuration
✓ 11/29/2010 3:55:05 AM	Registering
✓ 11/29/2010 4:04:14 AM	Configuring DC
✓ 11/29/2010 4:04:14 AM	DC configured
✓ 11/29/2010 4:04:30 AM	Configuring networks
✓ 11/29/2010 4:23:01 AM	Powering on (IP address: 172.31.1.68, network: 172.31.0.0, mask 255.255.0....
✓ 11/29/2010 4:23:05 AM	Heartbeat test
✓ 11/29/2010 4:23:07 AM	Pinging
✓ 11/29/2010 4:25:08 AM	Application initialization
✓ 11/29/2010 4:25:10 AM	Running test scripts
✓ 11/29/2010 4:25:10 AM	Domain Controller script, path VmConnectionTester.exe, ErrCode 0

The **Verification Job Session** window displays statistics for all VMs that are started during the recovery verification job – VMs from the application group in the specified order, and VMs from the linked backup job. For your convenience, these VMs are marked with different icons.

The recovery verification process includes the following steps:

1. **Getting virtual lab configuration.** Veeam Backup & Replication gets information about configuration of the virtual lab where verified VMs should be started.
2. **Starting virtual lab routing engine.** Veeam Backup & Replication starts a proxy appliance used as a gateway to provide access to the virtual lab.
3. **Publishing**. Veeam Backup & Replication creates an NFS-datastore with a VM backup and registers it on the selected ESX server. Veeam Backup & Replication does not deploy the whole VM from the backup file, it deploys VM configuration files only. Virtual disks are deployed per force and per required data blocks.
4. **Updating configuration**. Veeam Backup & Replication updates configuration files for VMs that should be run in the isolated network.
5. **Registering**. Veeam Backup & Replication registers the verified VM on the selected ESX host.
6. **Configuring DC**. If a verified VM has the Domain Controller or Global Catalog role, the VM is re-configured.

7. **Powering on**. Veeam Backup & Replication powers on the verified VM in the isolated network.

   To be able to perform tests for a verified VM without errors, Veeam Backup & Replication needs to know that the VM is ready for testing. To determine this, Veeam Backup & Replication waits for the VM to reach a "stabilization point"— that is, waits for the VM to boot completely and report it is ready for tests. After the stabilization point has been established, Veeam Backup & Replication can start performing heartbeat tests, ping tests and running test scripts against the VM.
   Veeam Backup & Replication establishes the stabilization point with the help of VMware parameters that it gets from the VM. Depending on the VM configuration, it uses one of the three algorithms to do that:

   - *Stabilization by IP*. This algorithm is used if the VM has VMware Tools installed, there are NIC(s) and mapped network(s) for these NIC(s). In this case, Veeam Backup & Replication waits for an IP address of the VM for mapped networks, which is sent by VMware Tools running in the VM. The sent IP address should be valid and should not change for a specific period of time.
   - *Stabilization by heartbeat*. This algorithm is used if the VM has VMware Tools installed but there are no NIC(s) and mapped networks for them. In this case Veeam Backup & Replication waits for a corresponding heartbeat signal (*Green* or *Yellow*) to come from the VM. As well as in the first case, the signal is sent by VMware Tools running in the VM.
   - *Stabilization by Maximum allowed boot time*. This algorithm is used if the VM has neither VMware Tools installed, nor NIC(s) and mapped networks for them. In this case, Veeam Backup & Replication will simply wait for the time specified in the **Maximum allowed boot time** field, which is considered to be a stabilization period for the VM. Once this time interval is exceeded, Veeam Backup & Replication will consider that the VM is successfully booted and is ready for testing.

Note: The stabilization process cannot exceed the value specified in the **Maximum allowed boot time** field. If the stabilization point cannot be determined within the **Maximum allowed boot time**, the recovery verification process will be finished with the timeout error. For this reason, you should be careful when specifying this value — typically, the VM started within the frames of a SureBackup job requires more time to boot if compared to a regular VM startup. When such an error situation occurs, you will need to increase the **Maximum allowed boot time** value and start the job once again.

Once the stabilization point has been established, Veeam Backup & Replication runs ping, heartbeat tests and performs test scripts against the verified VM.

8. **Pinging**. Veeam Backup & Replication checks if the VM responds to the ping requests or not. If the VM has no NIC(s) and mapped networks for them and/or has no VMware tools installed, the ping test will not be performed, and a notification will be written to the session details.
9. **Performing heartbeat test**. Veeam Backup & Replication checks whether the VMware Tools heartbeat signal (*Green* or *Yellow*) is coming from the VM or not. If the VM has no VMware Tools, the test will not be performed, and a notification will be written to the session details.
10. **Application initialization**. Veeam Backup & Replication waits for the applications installed in the VM (for example, SQL Server, web server, mail server) to start. The application initialization period is defined in the corresponding properties of a SureBackup job, and by default equals to 120 sec. However, depending on the software installed in a VM, the application initialization process may require more time than specified in the SureBackup job settings. If applications installed in a VM are not initialized within the specified period of time, test scripts can be completed with errors. If such an error situation occurs, you will need to increase the **Application initialization timeout** value and start the job once again.

11. **Running test scripts**. Veeam Backup & Replication runs scripts to test whether the application installed in the VM is working correctly or not. If the VM has no VMware Tools installed and/or there are no NIC(s) and mapped networks for them, Veeam Backup & Replication will skip tests that use variables such as *%vm_ip%* and so on, as the IP address of the VM cannot be determined.
    Test results are written to the session details. To define whether the script has completed successfully or not, Veeam Backup & Replication 5.0 uses return codes. If the return code is equal to 0, the script is considered to complete successfully. Other values in the return code mean that the script has failed.
12. **Powering off**. After all tests have been performed, SureBackup powers off the verified VM.
13. **Unregistering**. Veeam Backup & Replication unregisters the verified VM on the selected ESX host.
14. **Clearing redo logs**. Veeam Backup & Replication deletes redo logs that were created to store changes made to the VM while it is running from the backup file.
15. **Unpublishing**. Veeam Backup & Replication unpublishes the content of the backup file on the ESX host.

Once the verified VM is powered on, its name is displayed as a hyperlink. You can click the link to open the VM console (just like in the vSphere Client) to see what is happening inside a VM, or perform manual testing. To open the VM console, click the VM name link in the list of verified VMs.

After the verified VM is started and the application running there is initialized, you can start U-AIR wizards right from the **Realtime statistics** window to perform granular application item-level recovery. To do so, right-click the verified VM and select a corresponding command from the shortcut menu. Depending on the type of a running VM, you can start the **Active Directory item recovery** wizard, **Exchange item recovery** wizard, or **SQL item recovery** wizard.

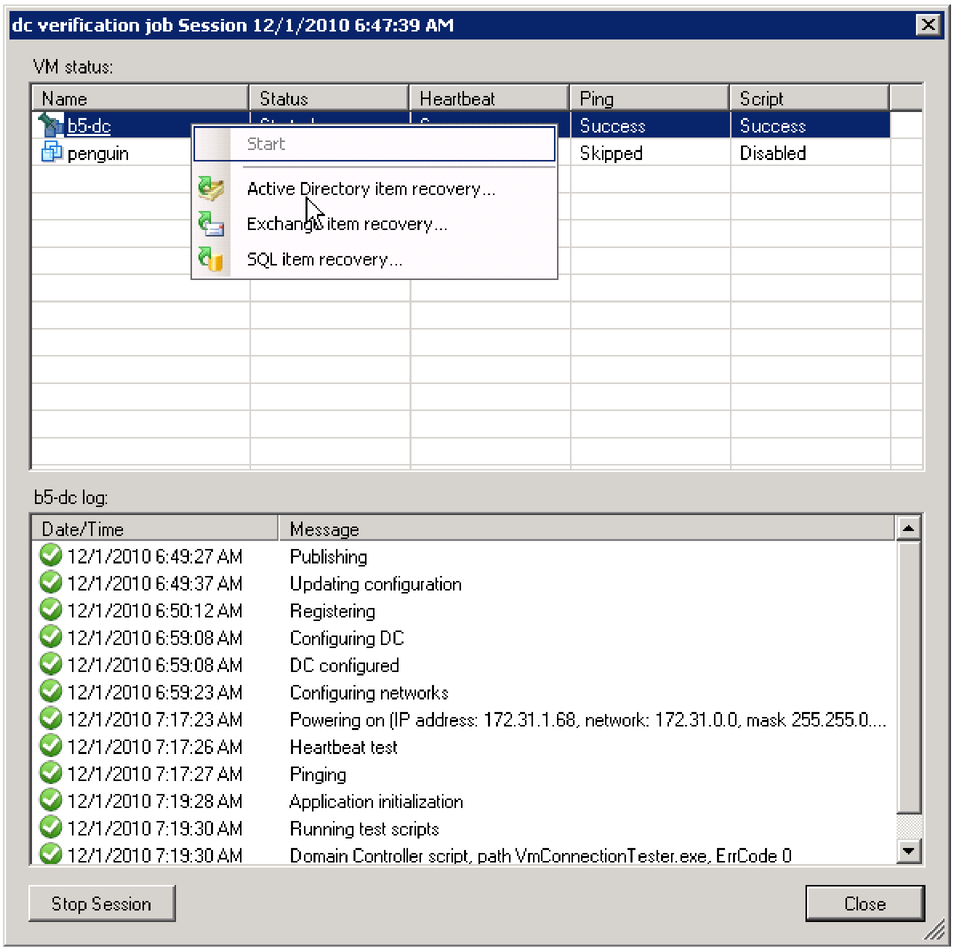
dc verification job Session 12/1/2010 6:47:39 AM
Start
Active Directory item recovery...
Exchange item recovery...
SQL item recovery...
Stop Session
Close

VM status:

Name	Status	Heartbeat	Ping	Script
b5-dc			Success	Success
penguin			Skipped	Disabled

b5-dc log:

Date/Time	Message
12/1/2010 6:49:27 AM	Publishing
12/1/2010 6:49:37 AM	Updating configuration
12/1/2010 6:50:12 AM	Registering
12/1/2010 6:59:08 AM	Configuring DC
12/1/2010 6:59:08 AM	DC configured
12/1/2010 6:59:23 AM	Configuring networks
12/1/2010 7:17:23 AM	Powering on (IP address: 172.31.1.68, network: 172.31.0.0, mask 255.255.0....
12/1/2010 7:17:26 AM	Heartbeat test
12/1/2010 7:17:27 AM	Pinging
12/1/2010 7:19:28 AM	Application initialization
12/1/2010 7:19:30 AM	Running test scripts
12/1/2010 7:19:30 AM	Domain Controller script, path VmConnectionTester.exe, ErrCode 0

If some VM fails to be verified automatically, once it is powered off, you can start it by right-clicking it in the list and selecting the **Start** command from the shortcut menu. If the application group has already been powered off by that time, it will be started again. After that, you can open the VM console and perform verification and testing manually.

Note: To define whether the script has completed successfully or not, Veeam Backup & Replication 5.0 uses return codes. If the return code is equal to 0, the script is considered to complete successfully. Other values in the return code mean that the script has failed.

Creating SureBackup Session Reports

Veeam Backup & Replication 5.0 allows you to generate HTML reports with statistics on a performed SureBackup job, a separate job session and multiple jobs sessions.

A report generated for a job contains detailed data on job sessions: job status, start and end time and details of the session performance, as well as the status of verified VMs and test results.

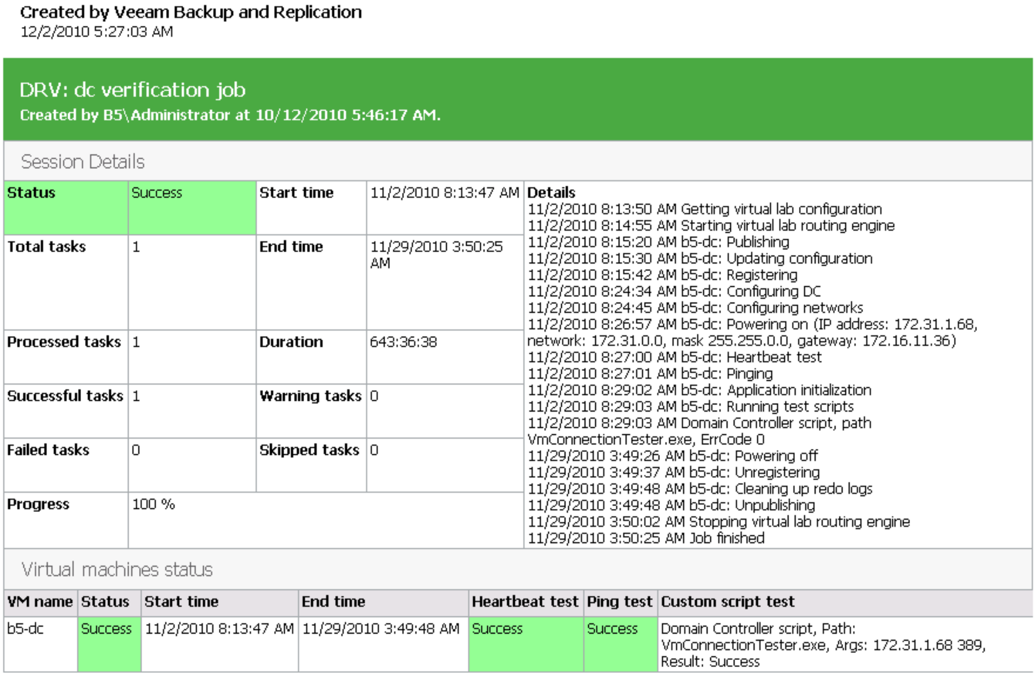

Created by Veeam Backup and Replication
12/2/2010 5:27:03 AM

DRV: dc verification job
Created by B5\Administrator at 10/12/2010 5:46:17 AM.

Session Details

Status	Success	**Start time**	11/2/2010 8:13:47 AM	**Details** 11/2/2010 8:13:50 AM Getting virtual lab configuration 11/2/2010 8:14:55 AM Starting virtual lab routing engine 11/2/2010 8:15:20 AM b5-dc: Publishing 11/2/2010 8:15:30 AM b5-dc: Updating configuration 11/2/2010 8:15:42 AM b5-dc: Registering 11/2/2010 8:24:34 AM b5-dc: Configuring DC 11/2/2010 8:24:45 AM b5-dc: Configuring networks 11/2/2010 8:26:57 AM b5-dc: Powering on (IP address: 172.31.1.68, network: 172.31.0.0, mask 255.255.0.0, gateway: 172.16.11.36) 11/2/2010 8:27:00 AM b5-dc: Heartbeat test 11/2/2010 8:27:01 AM b5-dc: Pinging 11/2/2010 8:29:02 AM b5-dc: Application initialization 11/2/2010 8:29:03 AM b5-dc: Running test scripts 11/2/2010 8:29:03 AM Domain Controller script, path VmConnectionTester.exe, ErrCode 0 11/29/2010 3:49:26 AM b5-dc: Powering off 11/29/2010 3:49:37 AM b5-dc: Unregistering 11/29/2010 3:49:48 AM b5-dc: Cleaning up redo logs 11/29/2010 3:49:48 AM b5-dc: Unpublishing 11/29/2010 3:50:02 AM Stopping virtual lab routing engine 11/29/2010 3:50:25 AM Job finished
Total tasks	1	**End time**	11/29/2010 3:50:25 AM	
Processed tasks	1	**Duration**	643:36:38	
Successful tasks	1	**Warning tasks**	0	
Failed tasks	0	**Skipped tasks**	0	
Progress	100 %			

Virtual machines status

VM name	Status	Start time	End time	Heartbeat test	Ping test	Custom script test
b5-dc	Success	11/2/2010 8:13:47 AM	11/29/2010 3:49:48 AM	Success	Success	Domain Controller script, Path: VmConnectionTester.exe, Args: 172.31.1.68 389, Result: Success

You can generate the following reports:

- *Job report*. This type of report contains data on all sessions initiated for a specific SureBackup job. To make up a job report, right-click a necessary job in the list and select **HTML Report**.
- *Session report*. This type of report contains data on a single job session. To make up a session report, click **Sessions** under the **SureBackup** node, right-click a necessary session in the list and select **HTML Report**.
- *Multiple session report*. You can also generate a report for a number of random job sessions. To make up a multiple sessions report, click **Sessions** under the **SureBackup** node and select necessary sessions (use the **Ctrl** and **Shift** keys to select a number of sessions). Then select the **HTML Report** command from the shortcut menu. The generated report will contain data on the sessions that have been selected.

Creating XML Files with VM Roles Description

VM roles are described in .xml files stored in the *SbRoles* subfolder in the product installation folder. To add a new role, you should create a new .xml file and save it to the *SbRoles* subfolder.

.xml files describing VM roles have the following structure:

```
<SbRoleOptions>
  <Role>
    <SbRole>
      <Id>4CDC7CC4-A906-4de2-979B-E5F74C44832F</Id>
      <Name>Web Server</Name>
    </SbRole>
  </Role>
  <Options>
    <SbVerificationOptions>
      <ActualMemoryPercent>100</ActualMemoryPercent>
      <MaxBootTimeoutSec>300</MaxBootTimeoutSec>
      <AppInitDelaySec>120</AppInitDelaySec>
      <TestScripts>
        <TestScripts>
          <TestScript>
            <Name>Web Server</Name>
            <Type>Predefined</Type>
            <TestScriptFilePath>VmConnectionTester.exe</TestScriptFilePath>
            <Arguments>%vm_ip% 80</Arguments>
          </TestScript>
        </TestScripts>
      </TestScripts>
      <HeartbeatEnabled>True</HeartbeatEnabled>
      <PingEnabled>True</PingEnabled>
    </SbVerificationOptions>
  </Options>
</SbRoleOptions>
```

Available XML tags are described in the table below.

Tag	Required/ Optional	Description
<SbRoleOptions>	Required	Encapsulates the VM role file.
<Role>	Required	Parent tag for a role assigned to a VM. <SbRole>, <Id> and <Name> are children of this tag.
<SbRole>	Required	Encapsulates basic information for a VM role – ID and name.
<Id>	Required	A unique identifier of a VM role.

<Name>	Required	Name of a VM role that is displayed in the roles list on the **Role** tab.
<Options>	Required	Parent tag for startup and test script options to be used for the defined role. <SbVerificationOptions>, <ActualMemoryPercent>, <MaxBootTimeoutSec>, <AppInitDelaySec>, <TestScripts>, <Name>, <Type>, <TestScriptFilePath>, <Arguments>,<HeartbeatEnabled>, <PingEnabled> are children of this tag.
<SbVerificationOptions>	Required	Encapsulates options data for a VM role.
<ActualMemoryPercent>	Optional	Percent of the original memory level set for a production VM that should be pre-allocated to a verified VM on the system boot.
<MaxBootTimeoutSec>	Optional	Maximum allowed time to boot a VM.
<AppInitDelaySec>	Optional	Maximum allowed time to initialize an application inside the VM.
<TestScripts>	Optional	Encapsulates test script data for a VM role.
<Name>	Optional	Name of a VM role to be displayed on the **Test Scripts** tab.
<Type>	Optional	Type of the test script – *Predefined* or *Custom*
<TestScriptFilePath>	Optional	Path to an executable file with a test script to be performed. Can be absolute or relative.
<Arguments>	Optional	Arguments to be passed to the script. You can use two variables here: - *%vm_ip%* - IP address of a virtual lab VM - *%vm_fqdn%* — a fully qualified domain name of a virtual lab VM
<HeartbeatEnabled>	Required	Should the heartbeat test be enabled for this VM role: *True* or *False*.
<PingEnabled>	Required	Should the ping test be enabled for this VM role: *True* or *False*.

Performing Universal Application Item-Level Restore

Veeam Backup & Replication 5.0 offers a new technology — U-AIR, or Universal Application Item-Level Restore, that allows you to restore individual items from any virtualized application – Active Directory, Microsoft SQL, Microsoft Exchange and many others. U-AIR does not require any special backups or additional tools – it starts an application and all components required for its proper work in the isolated virtual lab, and connects to the application using its native tools so that users can restore items they need.

For such applications as Active Directory, Microsoft SQL and Microsoft Exchange, U-AIR is a wizard-driven process — that is, you can restore necessary items from applications using Veeam's wizards. For other applications, U-AIR is user-driven — that is, Veeam Backup & Replication 5.0 starts the application and all required components in the virtual lab so that users can connect to that application and restore items themselves.

U-AIR wizards are not tied to Veeam Backup & Replication 5.0 — these are standalone components that can be downloaded, installed and updated independent of the product release. You can install U-AIR wizards on any machine in the production environment.

As a restore procedure requires specific knowledge and is commonly performed by application administrators or users working with these applications, the U-AIR process is distributed between two roles:

- Application administrators or users submit requests for virtual labs in which the required application should run, wait for the request to be approved, and perform the application item-level restore itself when the lab is ready.
- All submitted requests are registered at the Veeam Backup Enterprise Manager server. Portal administrators working with Veeam Backup Enterprise Manager make sure that users who submitted requests are eligible to access the corresponding application's data backup. After that, they approve or reject lab requests, and select necessary backups and virtual labs where restore should be performed.

To help users who requested virtual labs monitor the state of their request, Veeam Backup & Replication 5.0 offers a special tool — Virtual Lab Manager. Virtual Lab Manager runs on the machine from which the request has been sent and connects to Veeam Backup Enterprise Manager to notify users about the state of their requests. When the request is approved or rejected, a virtual lab is ready or its time elapses, Virtual Lab Manager displays a message hovering over its icon in the system tray.

Virtual Lab Manager is launched once the virtual lab request is submitted and continues running in the background even when the New Virtual Lab Request wizard is closed. Beside monitoring the state of your requests, you can Virtual Lab Manager it to create new virtual lab requests, open ready virtual labs and dismiss unnecessary requests.

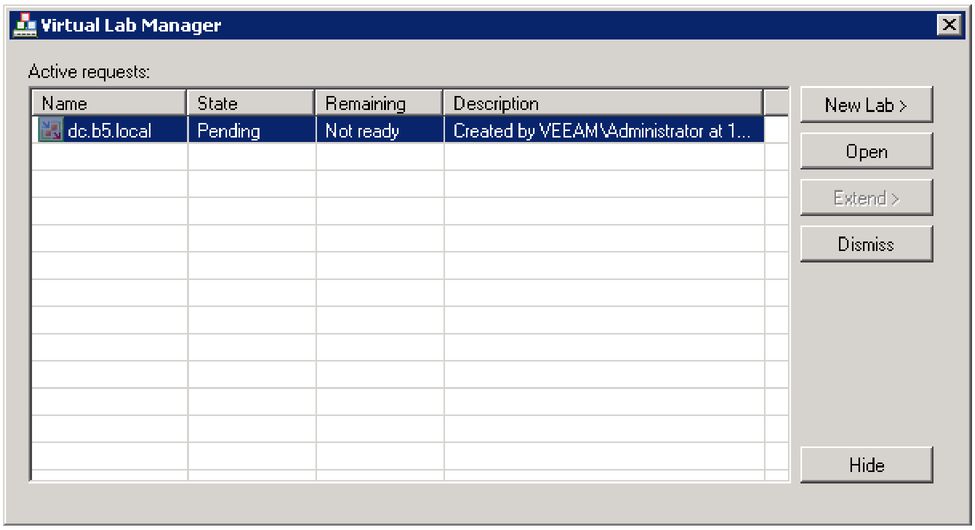

Name	State	Remaining	Description
dc.b5.local	Pending	Not ready	Created by VEEAM\Administrator at 1...

Recovering Objects from Active Directory Backups

Using the Active Directory Restore wizard, you can quickly restore deleted or modified Active Directory objects such as user accounts and groups, or recover individual attributes from AD backups to your production Active Directory.

To enable object restore from Active Directory, Veeam Backup & Replication 5.0 starts an isolated virtual lab from the selected SureBackup job and runs there a Domain Controller from the selected backup. The virtual lab should be configured beforehand and should contain all necessary VMs on which Active Directory is dependent, such as DNS server (if the Domain Controller itself does not perform the role of the DNS server). The Active Directory restore wizard connects to the production Domain Controller and the Domain Controller in the virtual lab, compares the current state of Active Directory with that from the backup, and lets you see what data has been changed or deleted so that you can recover it in its initial state.

Veeam Active Directory wizard works with objects in Domain Partition and can restore the following types of deleted or moved objects: *User*, *InetOrgPerson*, *Group*, *Computer*, *Contact*, *Printer*, *Organizational Unit*, *Container*. You can also restore attribute values of these class objects: *User*, *InetOrgPerson*, *Group*, *Computer*, *Contact*, *Printer*, *Organizational Unit*, *Container*, *Shared Folder*.

Note: When the Domain Controller starts in the virtual lab, it is first booted in the Directory Services Restore mode, and then automatically rebooted into a normal mode.

To recover individual objects or attributes, you have to perform the following actions:

- Create a request for a new virtual lab
- Wait for the virtual lab request to be approved
- Perform item-level restore

Creating Lab Request

To create a new Active Directory lab request, use one of the following options:

- Select **Tools** > **Application Item Restore** > **Microsoft Active Directory** from the main menu of Veeam Backup & Replication 5.0.
- Click the **Restore** button on the toolbar, select **Application Item** and click **Next**. At the **Select Application** step of the wizard, choose **Microsoft Active Directory** and click **Next**.

As a result, the **New AD Virtual Lab Request** wizard will be started. Follow the wizard steps.

Note: The Active Directory wizard is not installed by default with Veeam Backup & Replication 5.0. When you start it for the first time using one of the options mentioned above, Veeam Backup & Replication 5.0 will offer you to download and install it. Alternatively, you can download it from www.veeam.com web site and install it on any machine in your production environment.

Step 1. Specify Lab Description and Time Span to Run a Lab

Specify a description for the created lab request. By default, the following description is initially provided for the created job: time at which the lab request was created and user who created the request.

At the bottom of the window, specify the time span for which you want the created lab to run. Use time links to set a necessary period of time. By default, the virtual lab will be up for 30 minutes.

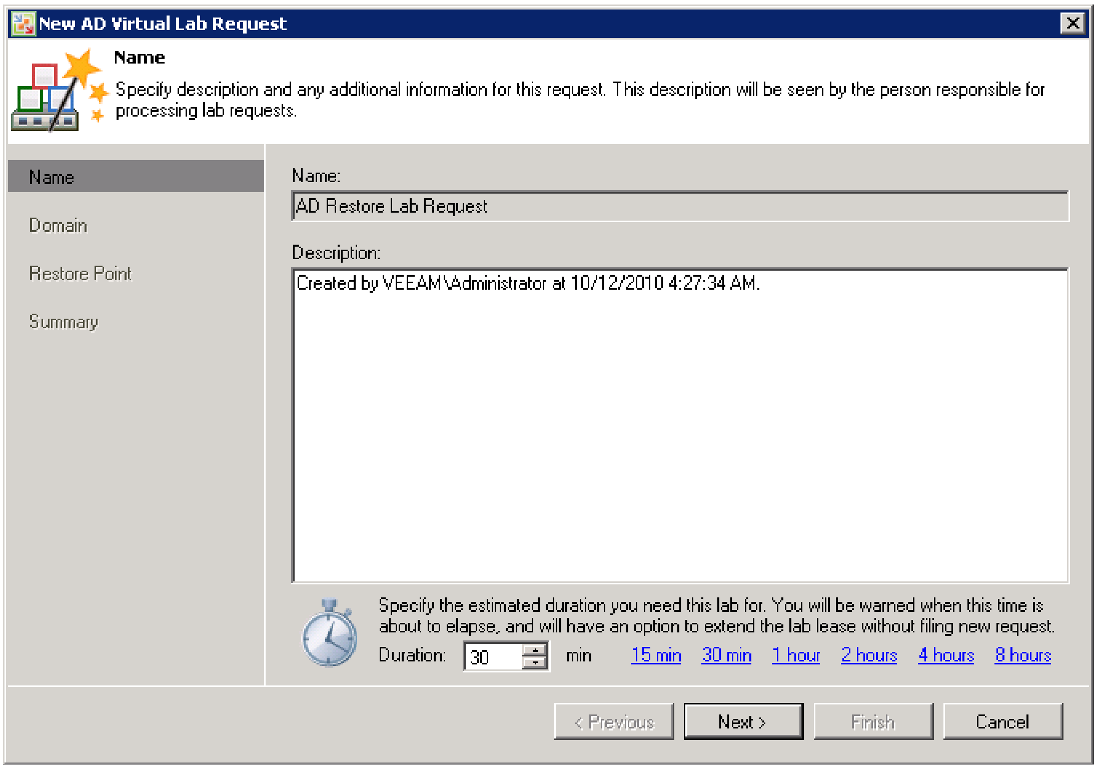

New AD Virtual Lab Request

Name

Specify description and any additional information for this request. This description will be seen by the person responsible for processing lab requests.

Name: AD Restore Lab Request

Description: Created by VEEAM\Administrator at 10/12/2010 4:27:34 AM.

Specify the estimated duration you need this lab for. You will be warned when this time is about to elapse, and will have an option to extend the lab lease without filing new request.

Duration: 30 min

Note: Once the time span set for the virtual lab elapses, you will be notified about that via Virtual Lab Manager. You can extend the time span for lab running by clicking the **Extend** button in Virtual Lab Manager without having to issue a new lab request.

Step 2. Specify Domain Controller

Enter a DNS name or IP address of a virtualized Domain Controller you are using in your production environment.

Enter a user name and password of the account you are planning to use to connect to Active Directory. You can use the account under which you are currently logged on to Active Directory, or specify another account.

Click the **Connect** button. Veeam Backup & Replication 5.0 will resolve the Domain Controller name and connect to it.

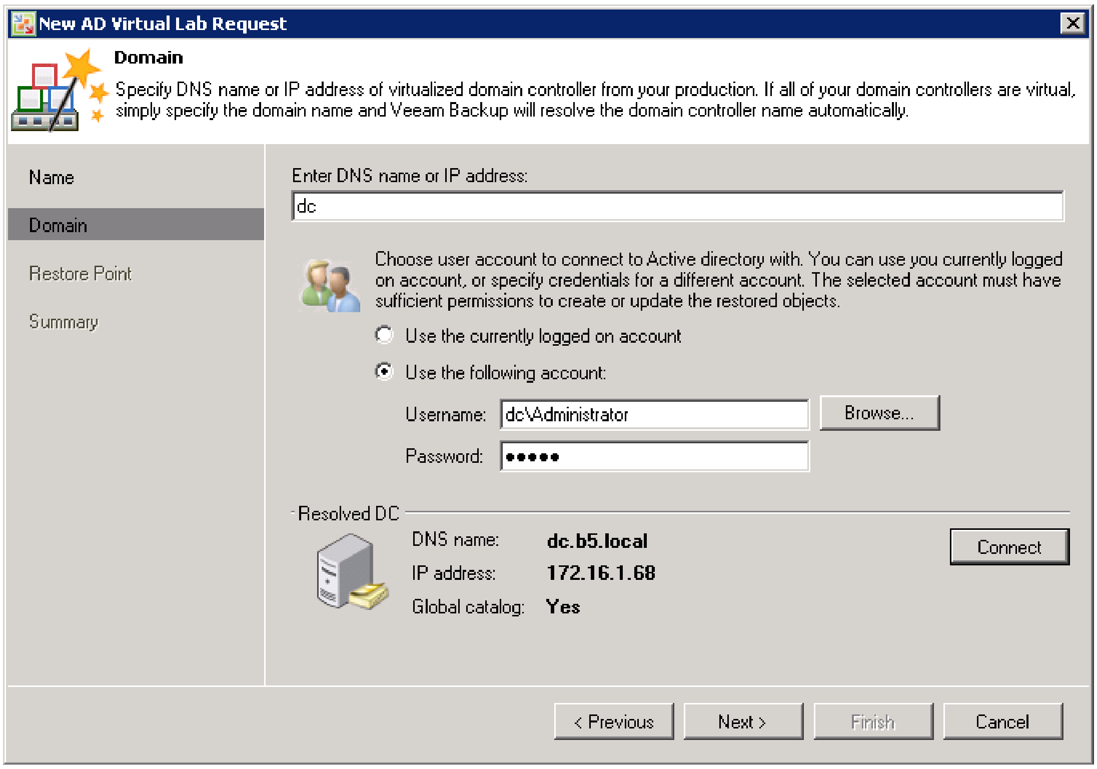

New AD Virtual Lab Request

Domain

Specify DNS name or IP address of virtualized domain controller from your production. If all of your domain controllers are virtual, simply specify the domain name and Veeam Backup will resolve the domain controller name automatically.

Enter DNS name or IP address: dc

Choose user account to connect to Active directory with. You can use you currently logged on account, or specify credentials for a different account. The selected account must have sufficient permissions to create or update the restored objects.

○ Use the currently logged on account

● Use the following account:

Username: dc\Administrator

Password: •••••

Resolved DC

DNS name: **dc.b5.local**

IP address: **172.16.1.68**

Global catalog: **Yes**

Step 3. Select a Restore Point for Active Directory Backup

From the list of available restore points, select the one when Active Directory was in the desired state, for example, before some objects were deleted. You can select the latest performed backup, last Friday night backup or a backup preceding a specific date.

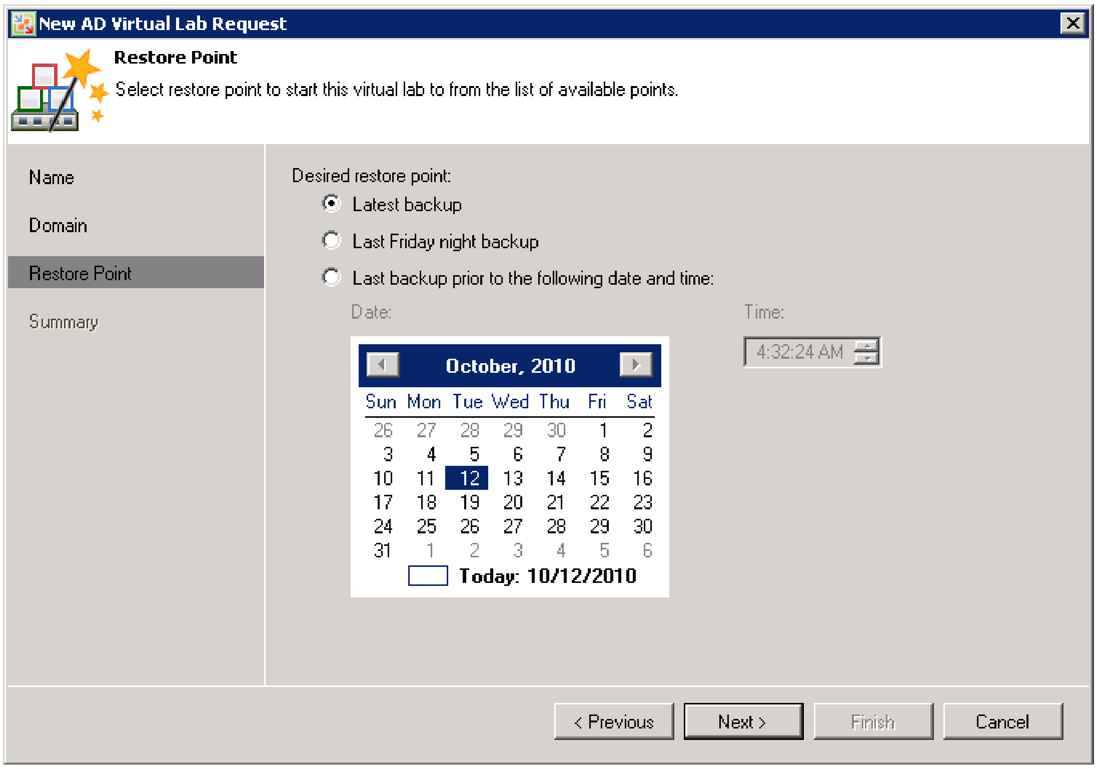

New AD Virtual Lab Request

Restore Point

Select restore point to start this virtual lab to from the list of available points.

Desired restore point:

● Latest backup

○ Last Friday night backup

○ Last backup prior to the following date and time:

Date:

Sun	Mon	Tue	Wed	Thu	Fri	Sat
26	27	28	29	30	1	2
3	4	5	6	7	8	9
10	11	12	13	14	15	16
17	18	19	20	21	22	23
24	25	26	27	28	29	30
31	1	2	3	4	5	6

Time: 4:32:24 AM

Step 4. Submit the Lab Request

Review the virtual lab request settings and submit the request. As soon as you click the **Finish** button, Veeam Backup & Replication 5.0 will register the request at the Veeam Backup Enterprise Manager server, and Virtual Lab Manager will start monitoring this request.

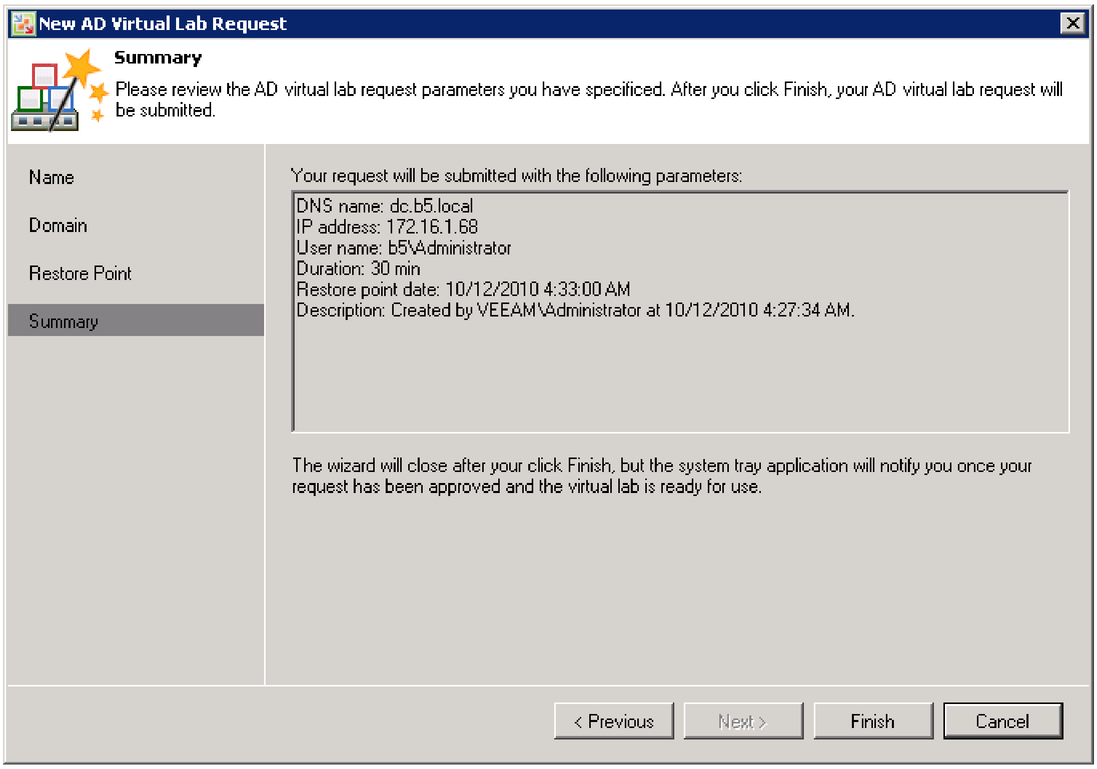

Please review the AD virtual lab request parameters you have specificed. After you click Finish, your AD virtual lab request will be submitted.

Your request will be submitted with the following parameters:

```
DNS name: dc.b5.local
IP address: 172.16.1.68
User name: b5\Administrator
Duration: 30 min
Restore point date: 10/12/2010 4:33:00 AM
Description: Created by VEEAM\Administrator at 10/12/2010 4:27:34 AM.
```

The wizard will close after your click Finish, but the system tray application will notify you once your request has been approved and the virtual lab is ready for use.

Approving Virtual Lab Requests

All virtual lab requests submitted by different users are listed on the **Lab Requests** tab in Veeam Backup Enterprise Manager. Administrators working with Veeam Backup Enterprise Manager can approve submitted lab requests, reject them or prolong the time for which a requested virtual lab should run.

Please keep in mind that you should have Portal Administrator rights in Veeam Backup Enterprise Manager to be able to work with lab requests.

To approve a lab request, select a necessary request in the list and click the **Approve…** button. Then follow the **Edit Lab Request** wizard steps.

Step 1. Review Lab Request Settings

At this step of the wizard, you can review and, if necessary, edit a virtual lab request — for example, change the time span for which the lab should run. To edit virtual lab request data, click **Edit request** link at the bottom of the window.

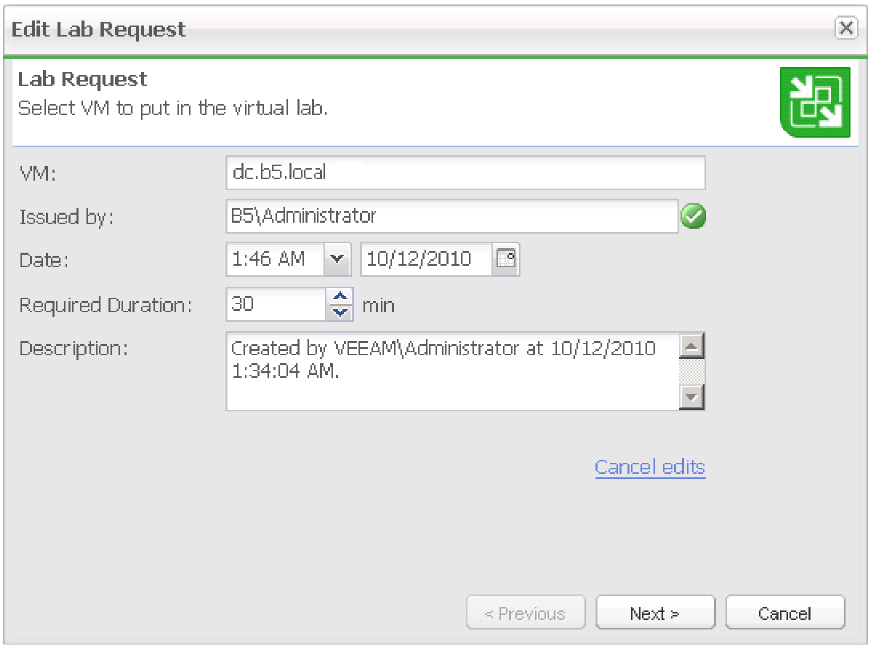

Edit Lab Request

Lab Request

Select VM to put in the virtual lab.

VM: dc.b5.local

Issued by: B5\Administrator

Date: 1:46 AM 10/12/2010

Required Duration: 30 min

Description: Created by VEEAM\Administrator at 10/12/2010 1:34:04 AM.

Step 2. Select a VM from the Backup

Select a backed up VM from which you want to restore Active Directory objects. The list of backed up VMs is formed automatically depending on the name of a VM specified at the previous step of the wizard — in our case, *dc*.

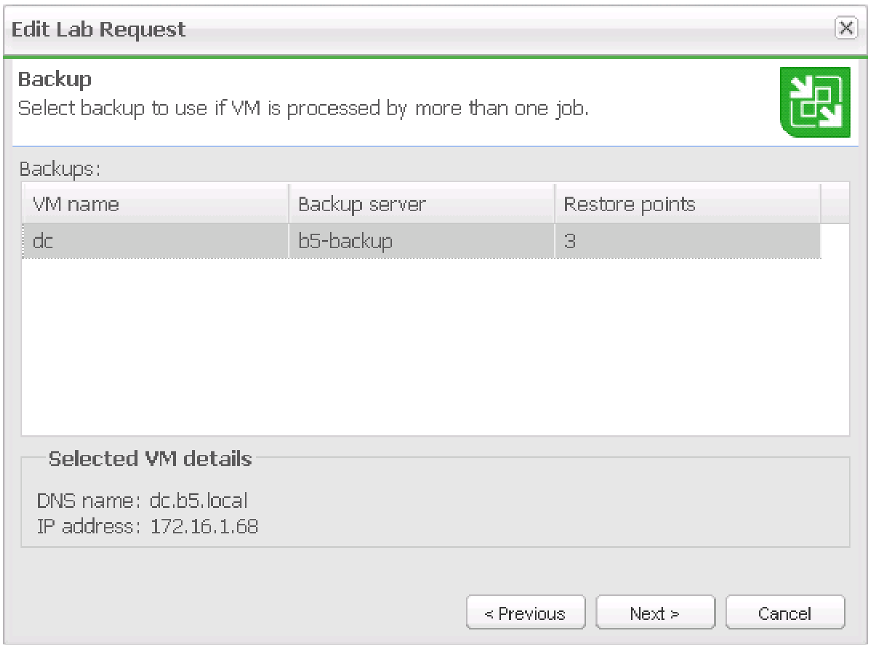

Edit Lab Request

Backup

Select backup to use if VM is processed by more than one job.

Backups:

VM name	Backup server	Restore points
dc	b5-backup	3

Selected VM details

DNS name: dc.b5.local
IP address: 172.16.1.68

Step 3. Select a Restore Point

Select the restore point when Active Directory was in the desired state. The list of restore points is formed depending on the choice of the user who submitted the virtual lab request. For example, if the user selected the **Last Friday night backup**, all restore points created on the last Friday night will be displayed.

If you want to display all restore points that were created for this VM, select the **Show all available restore points** check box.

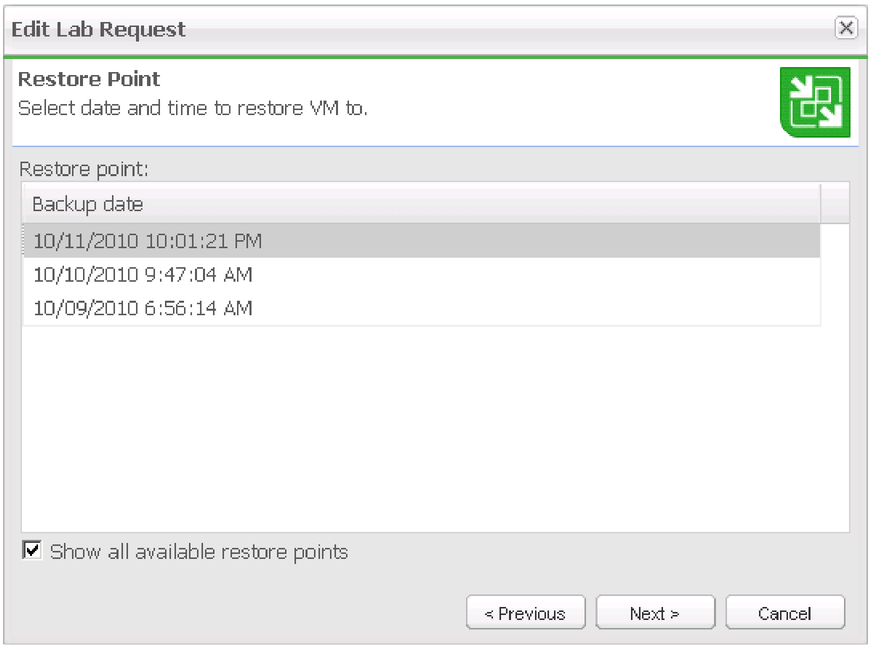

Edit Lab Request

Restore Point

Select date and time to restore VM to.

Restore point:

Backup date
10/11/2010 10:01:21 PM
10/10/2010 9:47:04 AM
10/09/2010 6:56:14 AM

☑ Show all available restore points

Step 4. Select the SureBackup Job to Be Used

Select one of the existing SureBackup job that you want to run to create an isolated sandbox in which Domain Controller(s) from the selected backup should be started. The application group and virtual lab used by this Surebackup job will be displayed in the **Selected Job details** section.

By default, the list of jobs displays only those jobs that contain the selected virtual Domain Controller. If you want to display all SureBackup jobs that were created, select the **Show all available SureBackup jobs** check box.

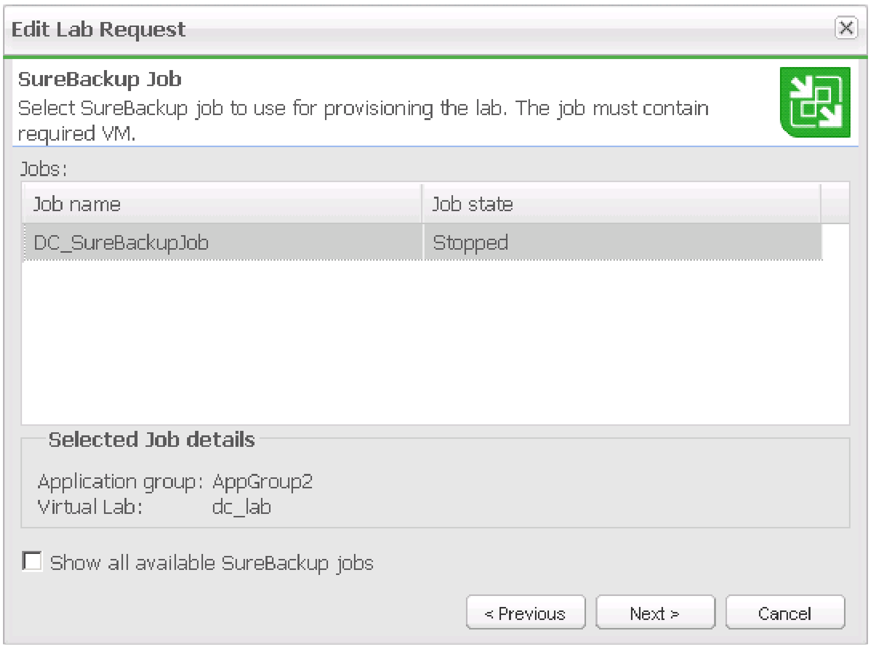

Edit Lab Request

SureBackup Job

Select SureBackup job to use for provisioning the lab. The job must contain required VM.

Jobs:

Job name	Job state
DC_SureBackupJob	Stopped

Selected Job details

Application group: AppGroup2
Virtual Lab: dc_lab

☐ Show all available SureBackup jobs

Step 5. Approve the Lab Request

Review the settings you have configured for the virtual lab and click **Finish**. Veeam Backup & Replication 5.0 will initiate the selected SureBackup job, start the virtual lab and restore Domain Controller into it. Once the virtual lab is ready, it will send a notification to Virtual Lab Manager so that the user who requested the virtual lab can start restoring Active Directory objects.

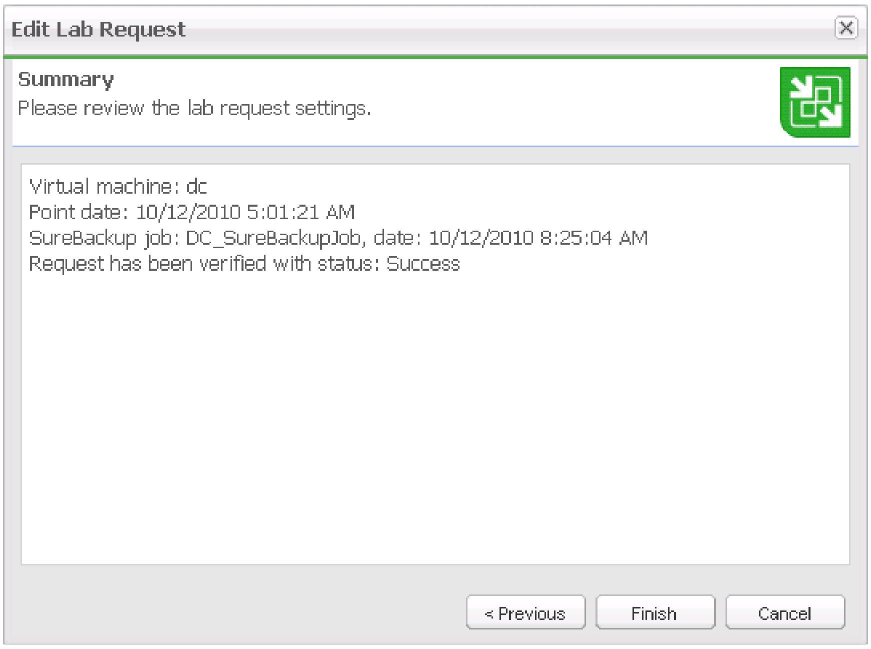

Virtual machine: dc
Point date: 10/12/2010 5:01:21 AM
SureBackup job: DC_SureBackupJob, date: 10/12/2010 8:25:04 AM
Request has been verified with status: Success

Declining Virtual Lab Requests and Prolonging Virtual Lab Existence

Using the **Lab Requests** tab in Veeam Backup Enterprise Manager, you can also decline virtual lab requests and prolong the time for which a created virtual lab should run.

- To decline a lab request, select it in the list, click the **Reject…** button and enter the reason for declining the request. The user who submitted the request will be notified about it via Virtual Lab Manager.
- To prolong the time span for which the lab should run, select it in the list and click the **Prolong...** button.

Performing Item-Level Restore

Once the virtual lab has been created and is running, you will be notified about it via Virtual Lab Manager and can start restoring necessary Active Directory objects.

Using the Active Directory wizard, you can either restore objects that have been deleted, or recover object attributes that have been changed.

Step 1. Open the Virtual Lab Request Wizard

Click the **Open** button in Virtual Lab Manager to bring up the **New Virtual Lab Request** wizard and make sure the virtual lab is ready.

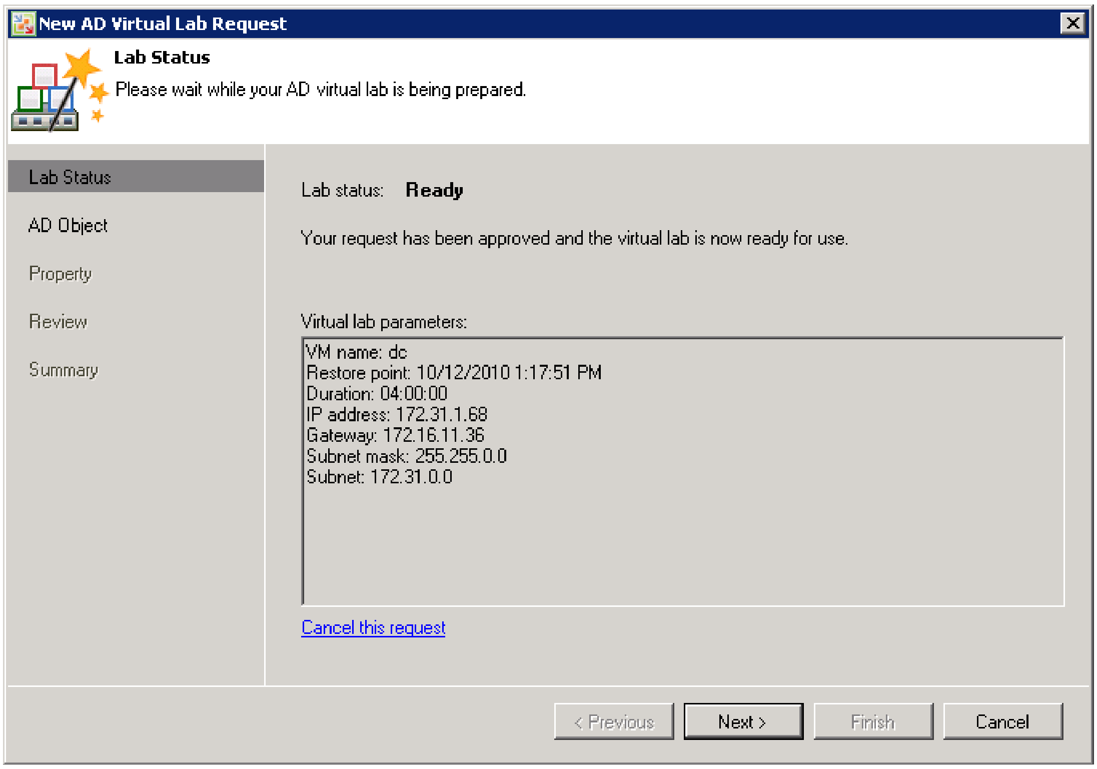

Lab status: **Ready**

Your request has been approved and the virtual lab is now ready for use.

Virtual lab parameters:

```
VM name: dc
Restore point: 10/12/2010 1:17:51 PM
Duration: 04:00:00
IP address: 172.31.1.68
Gateway: 172.16.11.36
Subnet mask: 255.255.0.0
Subnet: 172.31.0.0
```

Step 2. Select Objects to Restore

Select the object you want to restore or which attributes you want to recover.

Tip: To quickly find a necessary object, use the search field at the bottom of the window.

- Objects that have been deleted and no longer exist in the production Active Directory are marked with the *(DELETED)* word appended to them. Objects that have been moved are marked with the *(MOVED)* word appended to it. To restore such objects, select them in the tree and click **Next**. You can select several objects using **SHIFT** and **CTRL** keys on the keyboard.
  To learn about object recovery, see Scenario 1 below.
- To recover attributes of objects that still exist in the production Active Directory, select them in the tree and click **Next**. To learn about attributes recovery, see Scenario 2 below.

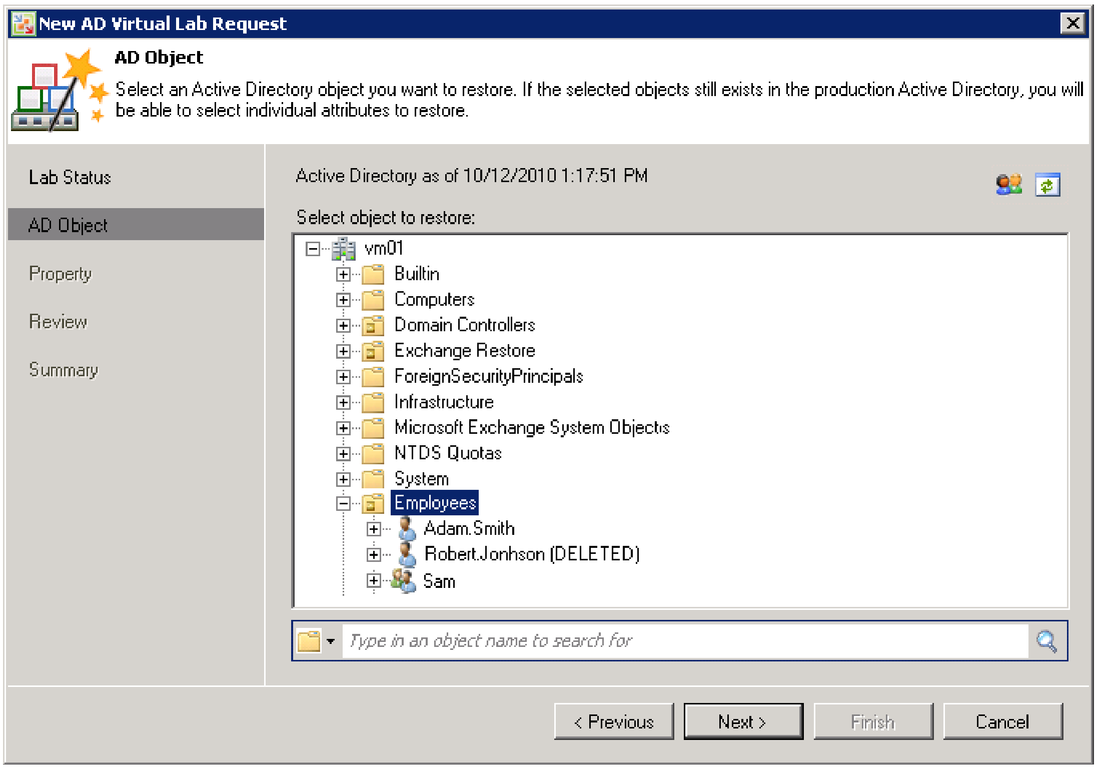

New AD Virtual Lab Request

AD Object

Select an Active Directory object you want to restore. If the selected objects still exists in the production Active Directory, you will be able to select individual attributes to restore.

Active Directory as of 10/12/2010 1:17:51 PM

Select object to restore:

- vm01
  - Builtin
  - Computers
  - Domain Controllers
  - Exchange Restore
  - ForeignSecurityPrincipals
  - Infrastructure
  - Microsoft Exchange System Objects
  - NTDS Quotas
  - System
  - Employees
    - Adam.Smith
    - Robert.Jonhson (DELETED)
    - Sam

Scenario 1. Restoring a Deleted or Moved Object

This scenario describes the procedure of restoring Active Directory objects that have been deleted or moved.

Veeam Backup & Replication 5.0 does not restore passwords for user that have been deleted and restores user accounts in the disabled state. To enable user accounts, select the **Enable users** check box. In this case, after you restore users to the production Active Directory, user accounts will be enabled right after the restore procedure is complete.

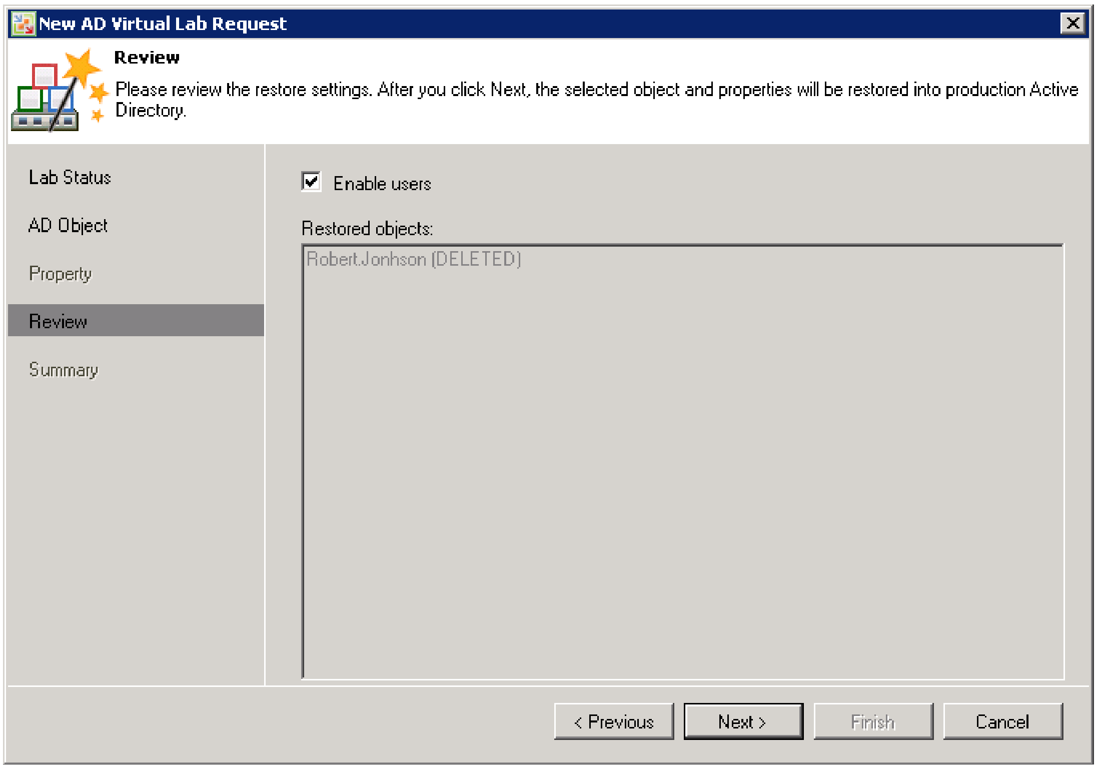

New AD Virtual Lab Request

Review

Please review the restore settings. After you click Next, the selected object and properties will be restored into production Active Directory.

☑ Enable users

Restored objects:

Robert.Jonhson (DELETED)

Once you click **Next**, you will be asked to provide a new password. If you are restoring several users at once, you can provide the same password for all of them by selecting the **Use for all restored users** check box.

Note: Passwords are requested only if you have selected the **Enable users** option.

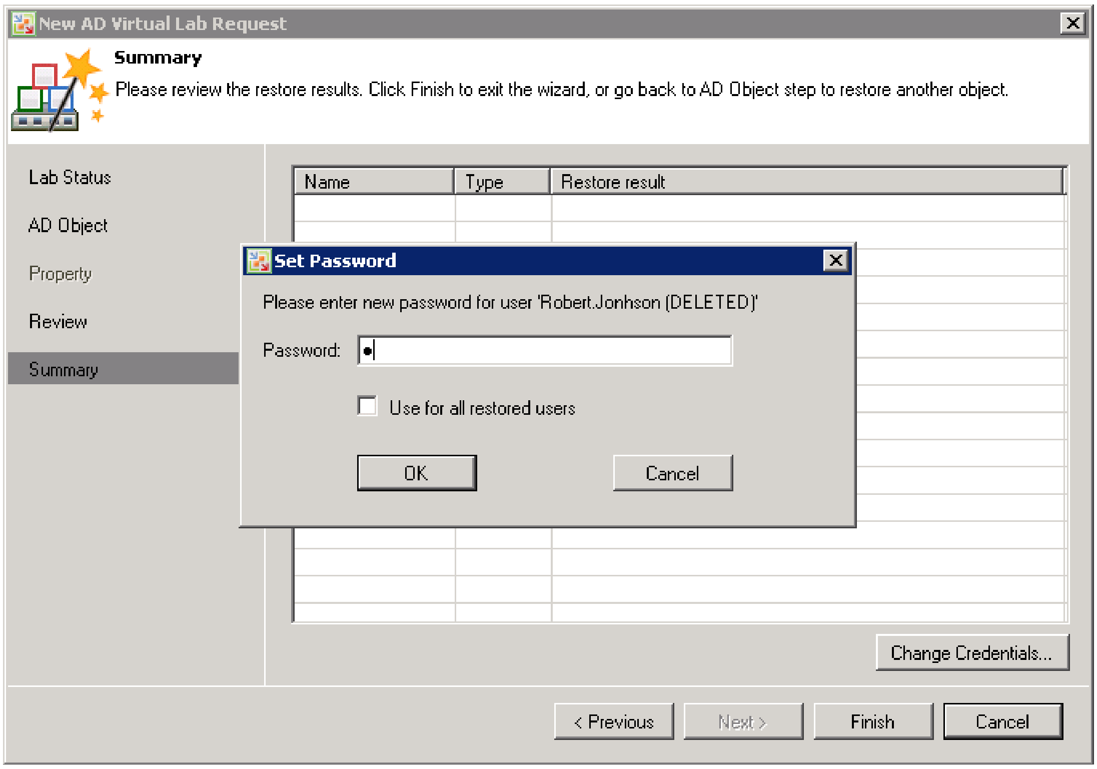
New AD Virtual Lab Request
Summary
Please review the restore results. Click Finish to exit the wizard, or go back to AD Object step to restore another object.
Lab Status
AD Object
Property
Review
Summary
Name
Type
Restore result
Set Password
Please enter new password for user 'Robert.Jonhson (DELETED)'
Password:
□ Use for all restored users
OK
Cancel
Change Credentials...
< Previous
Next >
Finish
Cancel

Click **Next**, then review the settings of a restored object. Once you click **OK**, the object will be restored to the production Active Directory.

Note: Veeam Backup & Replication 5.0 does not restore passwords for computer accounts you recover. To finalize computer account recovery once work with the Active Directory wizard is completed, you will have to enable restored computer accounts and re-join to domain computers whose accounts were restored.

Scenario 2. Restoring Attributes of Active Directory Objects

This scenario describes the procedure of restoring attributes of Active Directory objects that have been modified.

Once you have selected a necessary object and clicked **Next**, the Active Directory wizard will display attribute values for the selected object from the Active Directory backup, and the object existing in the production Active Directory. To display only those attributes that have changed, select the **Show differences only** check box.

Select check boxes next to attributes you want to restore.

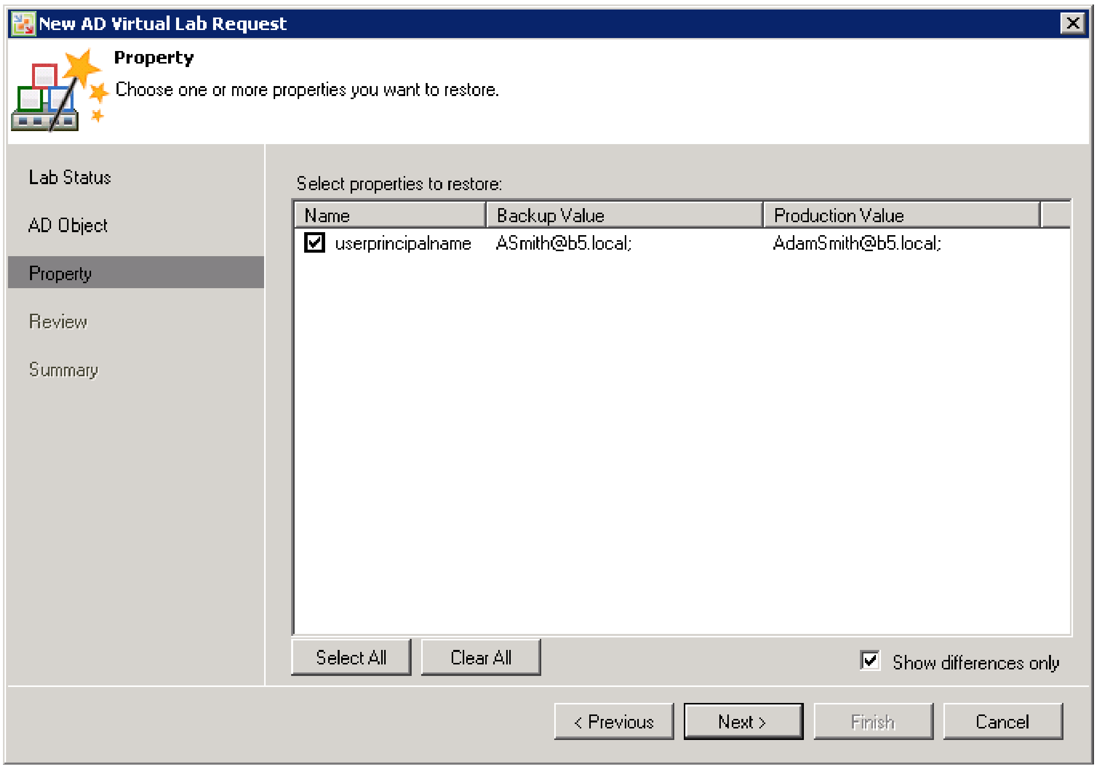

Property

Choose one or more properties you want to restore.

Select properties to restore:

Name	Backup Value	Production Value
☑ userprincipalname	ASmith@b5.local;	AdamSmith@b5.local;

☑ Show differences only

Click **Next** to review the attributes you selected. Once you click **Next**, attributes in the production Active Directory will be replaced with those from the Active Directory backup.

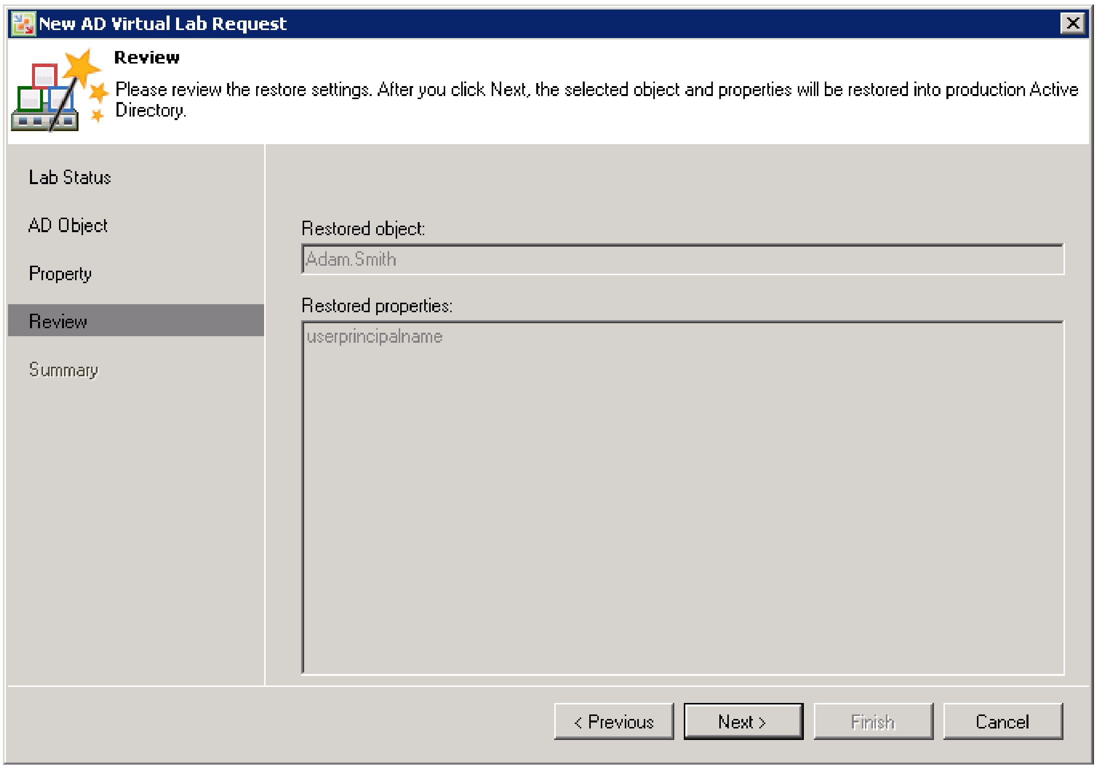

Review

Please review the restore settings. After you click Next, the selected object and properties will be restored into production Active Directory.

Restored object: Adam.Smith

Restored properties: userprincipalname

Specifying Veeam Backup & Replication Options

This section provides a detailed description about general Veeam Backup & Replication options.

Specifying E-Mail Notification Settings

With Veeam Backup & Replication 5.0, you can select to receive e–mail messages in case of success or failure of a created backup or replication job. To be able to receive e-mail notifications, you should configure general e-mail notification settings and select to receive a notification when creating a corresponding job.

Tip: To be able to receive e-mail notification about all performed jobs at once, use Veeam Backup Enterprise Manager. To learn more, see the Veeam Backup Enterprise Manager section.

Configuring General E-Mail Notification Settings

To configure general e-mail notification settings, select **Tools > Options…** from the main menu. Select the **Enable e–mail notification** check box and specify e–mail notification settings:

1. In the **SMTP Server** field, enter the DNS name or IP address of the SMTP server that will be used for sending e–mail messages.
2. Use the **Advanced…** button to specify user credentials and connection options — port number and connection timeout.
3. In the **From** field, specify the e–mail from which e-mail notifications should be sent.
4. In the **To** field, specify the recipient address(es). Use semicolon to enter multiple addresses.
5. In the **Subject** field, specify the subject for a sent message. You can use two variables in the subject: %Job Name% and %Job Result%.
6. Select the **Notify on success**, **Notify on warning** and/or **Notify on failure** check boxes to receive e–mail notification in case a job is run successfully, not successfully or with a warning.

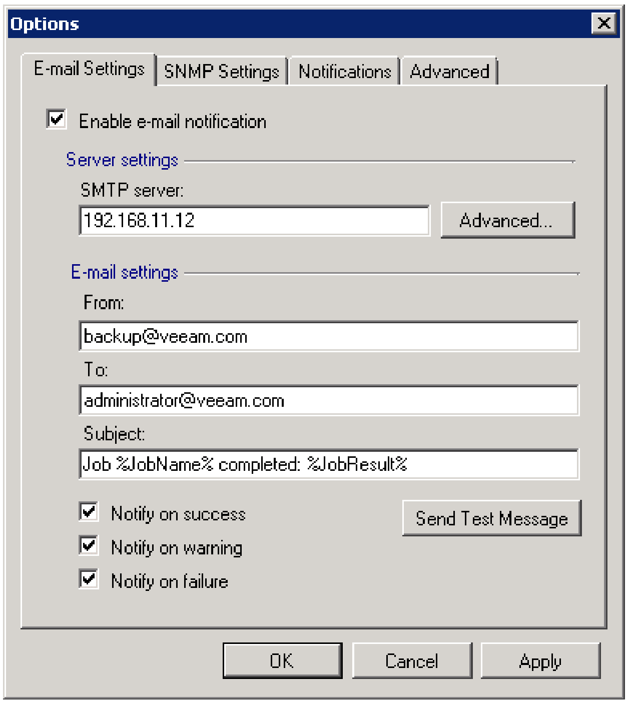

☑ Enable e-mail notification

SMTP server: 192.168.11.12

From: backup@veeam.com

To: administrator@veeam.com

Subject: Job %JobName% completed: %JobResult%

☑ Notify on success

☑ Notify on warning

☑ Notify on failure

Veeam Backup & Replication 5.0 allows sending a test e–mail to check if all settings have been configured correctly: click the **Send Test Message** button to receive a test e–mail.

Configuring Job Notification Settings

To configure job notification settings:

1. At the step of specifying destination for the created job, click the **Advanced…** button.
2. On the **Notifications** tab, select the **Send email notifications to the following recipients** check box.
3. In the field below enter an e-mail to which a notification should be sent. To enter several e-mails, use semicolon.

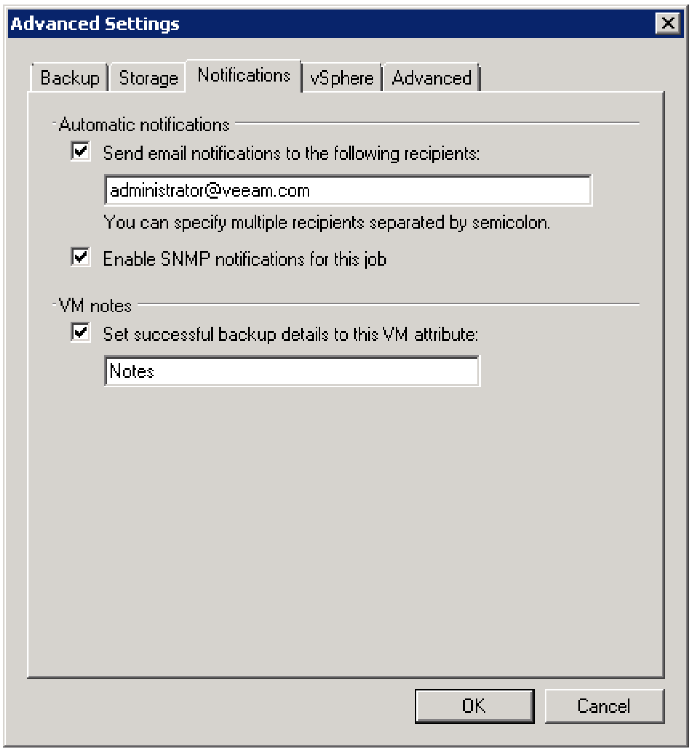

Automatic notifications

☑ Send email notifications to the following recipients:

administrator@veeam.com

You can specify multiple recipients separated by semicolon.

☑ Enable SNMP notifications for this job

VM notes

☑ Set successful backup details to this VM attribute:

Notes

Specifying SNMP Settings

Veeam Backup & Replication provides a possibility to monitor execution of backup and replication jobs using SNMP traps. You can select receive SNMP notifications once each job is completed and backup or replica is created. SNMP traps can be used to feed data into other popular system monitors, such as CA Unicenter, BMC Patrol, IBM Tivoli or HP OpenView.

To be able to receive SNMP traps, you should:

- Configure general SNMP settings in Veeam Backup & Replication
- Configure SNMP service properties on the trap recipients' computers
- Select to receive SNMP settings for a specific job

Configuring General SNMP Settings

To configure general SNMP settings:

1. Select **Tools > Options...** from the main menu of Veeam Backup & Replication.
2. Click the **SNMP Settings** tab.
3. In the **Receiver** field, specify an IP address of the SNMP recipient.

4. In the field on the right, enter the port number to be used.
5. In the **Community String** field, enter the community identifier.

Trap notifications can be sent to 5 different destinations.

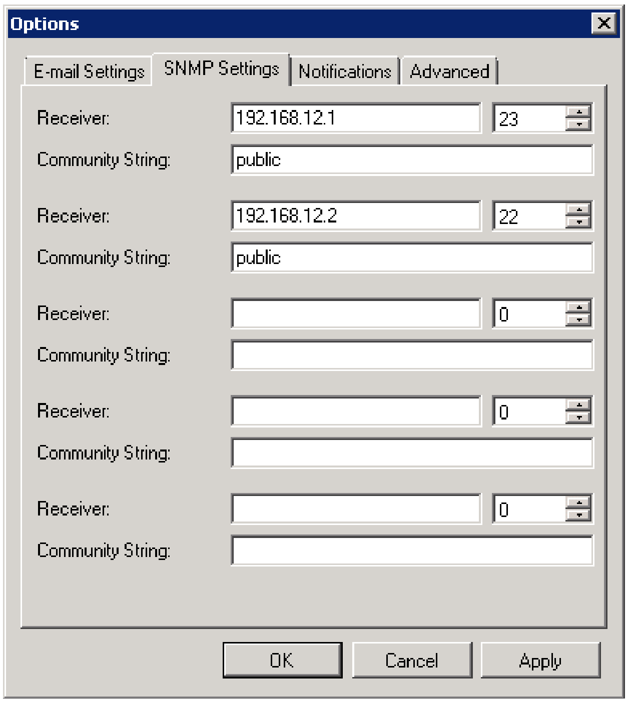

Receiver: 192.168.12.1 23
Community String: public

Receiver: 192.168.12.2 22
Community String: public

Receiver: 0
Community String:

Receiver: 0
Community String:

Receiver: 0
Community String:

Configuring SNMP Service Properties

To configure SNMP service properties on the trap recipients' computers:

1. Install standard Microsoft SNMP agent from the Windows distribution.
2. From the **Start** menu, select **Control Panel > Administrative Tools > Services**.
3. Double-click SNMP Service to open the **SNMP Service Properties** window.
4. Click the **Traps** tab.
5. Add the public string to the **Community name** list and a necessary host name — to the **Trap destinations** list.
6. Click the **Security** tab.
7. Make sure the **Send authentication trap** option is selected.
8. Add the public string to the **Accepted community names** list.
9. Select the **Accept SNMP packets from any hosts** option.
10. Click **Apply** and then **OK** to accept changes.

Specifying SNMP Settings for Jobs

To be able to receive SNMP traps with results for a specific job:

1. At the step of specifying destination for the created job, click the **Advanced…** button.
2. On the **Notifications** tab, select the **Enable SNMP notifications for this job** check box.

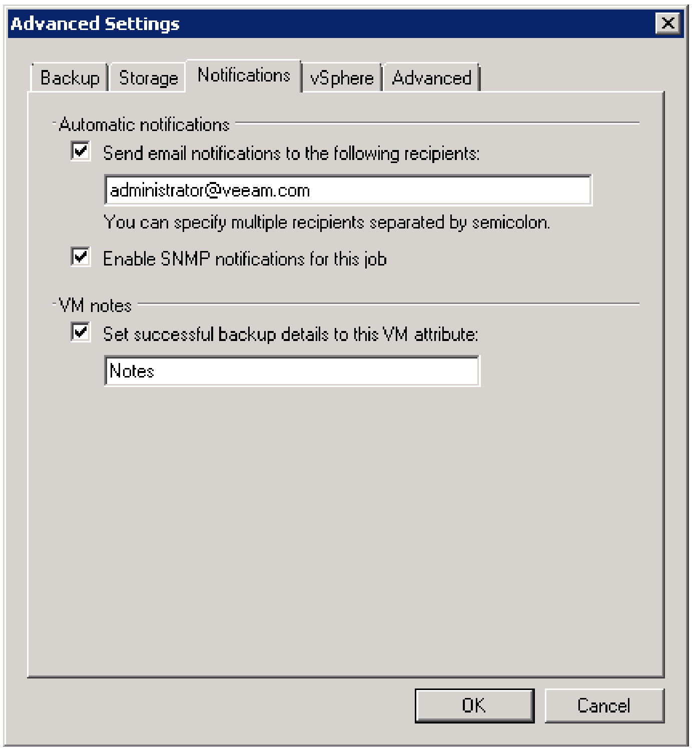

Automatic notifications

☑ Send email notifications to the following recipients:

administrator@veeam.com

You can specify multiple recipients separated by semicolon.

☑ Enable SNMP notifications for this job

VM notes

☑ Set successful backup details to this VM attribute:

Notes

Specifying Global Notification Settings

When a job is run, Veeam Backup & Replication 5.0 checks disk space on the backup storage and on production datastores. If the disk space is below a specific value, a warning will be displayed. To specify the disk space threshold:

1. Select **Tools** > **Options…** from the main menu.
2. Click the **Notifications** tab.
3. In **Backup storage** and **Production datastores** sections, select the **Warn me if free disk space is below N percent** options and specify a desired disk space threshold.

Specifying Advanced Settings

Using advanced settings of Veeam Backup & Replication, you can enable legacy processing backup modes and specify session history settings.

1. Select **Tools** > **Options...** from the main menu.
2. Click the **Advanced** tab.
3. To be able to use legacy backup modes such as **VCB-enabled backup** and **Network** modes, select the **Enable legacy processing modes** check box. Legacy modes are left for compatibility with previous versions and can be enabled if you are using ESX hosts earlier than 3.5.
4. In the **Sessions** section, specify the number of sessions to display in the **Sessions** list and the number of sessions to keep in the database.

Reporting

When a job is being run, jobs statistics and operation data is written to the *VeeamBackup* database. Veeam Backup & Replication 5.0 allows viewing real-time statistics on a performed job and generating HTML reports with statistics on a job, a separate job session and multiple jobs sessions.

To view real-time statistics for a job being run, right-click the corresponding job in the information area and select **Realtime Statistics** from the shortcut menu.

A report generated for a job contains detailed data on job sessions: job status, start and end time, total number of processed and failed objects, size, performance rate and details of the session performance (for example, errors that have occurred in the process of operation). Additionally, it contains detailed data on each object processed within the frames of a job (that is, a virtual machine).

Depending on the type of data to be covered, you can generate the following reports:

- **Job report**. This type of report contains data on all sessions initiated for a specific job. To make up a job report, right-click a necessary job in the list and select **HTML Report**.
- **Session report**. This type of report contains data on a single job session. To make up a session report, right-click a necessary session in the list and select **HTML Report**.
  A session report can also be generated from the **Statistics** window. To view job statistics, right-click a necessary job and select **Realtime Statistics**. Select a necessary job session by pressing the **Next Session** and **Previous Session** buttons and click the **HTML Report** button.
- **Multiple session reports**. Veeam Backup & Replication 5.0 provides a possibility to generate a report for a number of random job sessions. To make up a multiple sessions report, select necessary sessions in the Sessions list (click **Backup** > **Sessions** in the management tree). Use the **Ctrl** and **Shift** keys to select a number of sessions. Then select the **HTML Report** command from the shortcut menu. The generated report will contain data on the sessions that have been selected.

Users and Roles

There are three levels of security that can be granted to users who work with Veeam Backup & Replication 5.0:

- Veeam Restore Operators
- Veeam Backup Viewers
- Veeam Backup Operators
- Veeam Backup Administrators

A security scheme in Veeam Backup & Replication 5.0 is mainly used for work with Veeam Backup Enterprise Manager. To learn more about security settings in Veeam Backup Enterprise Manager, see the Specifying Security Settings section of Veeam Backup Enterprise Manager documentation.

In Veeam Backup & Replication 5.0, security settings are checked for managing (starting and stopping) jobs and performing the restore operations.

Role	Operations
Veeam Restore Operator	Can perform restore operations using existing backups and replicas.
Veeam Backup Viewer	Has the “read-only” access to Veeam Backup & Replication – can view existing and performed jobs and review the job session details.
Veeam Backup Operator	Can start and stop existing jobs and perform restore operations.
Veeam Backup Administrator	Can perform all administrative activities in Veeam Backup & Replication.

To specify user security settings:

1. Select **Tools > Users and Roles...** from the main menu.
2. Click the **Add** button.
3. In the **User name** field, enter the name of a user or group in the *DOMAIN\Username* format.
4. From the **Role** list, select a necessary role to be assigned: *Veeam Backup Administrator* or *Veeam Backup Operator*.

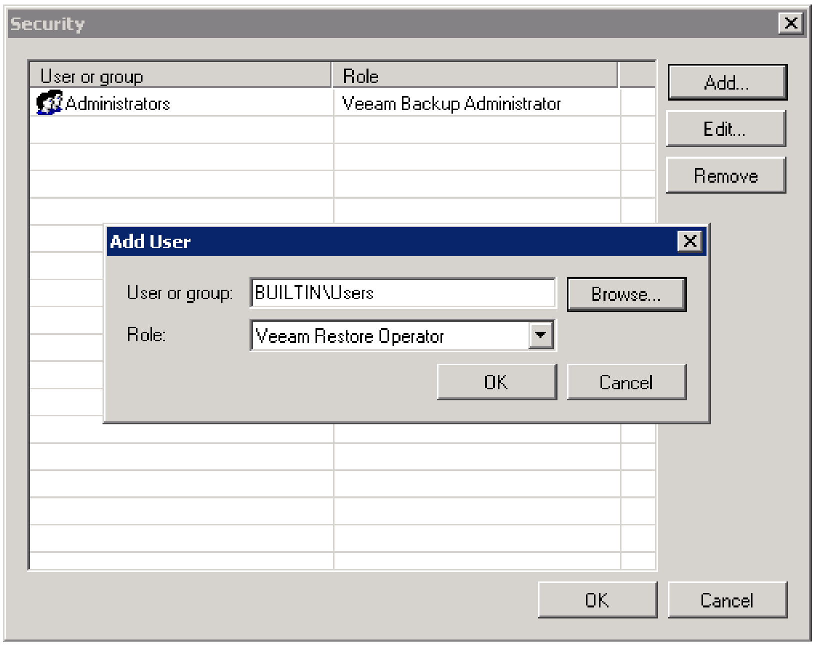

Security

User or group	Role
Administrators	Veeam Backup Administrator

Add User

User or group: BUILTIN\Users

Role: Veeam Restore Operator

Tip: By default, during installation the Veeam Backup Administrator role is assigned to users listed in the local Administrators group.

Logging

Veeam Backup & Replication 5.0 provides detailed logging of performed activities, initiated jobs, Backup Agent work and so on. Log files are stored at: *%userprofile%\Local Settings\Application Data\Veeam\Backup* (for Windows Vista and higher, log files are stored at: *%userprofile%\AppData\Local\Veeam\Backup*).

Veeam Backup & Replication 5.0 keeps a separate log file for each of its components: *Veeam Shell*, *Veeam Backup Service*, *Veeam Indexing Service, Veeam vPower NFS Service, Veeam Agents*, *Veeam Manager* and performed jobs. Please note that logs for Veeam Backup services are stored under the account that was used to run the service, and logs for Veeam Shell are stored under the account that was used to start Veeam Backup shell.

To facilitate browsing to the log files, select **Help > Support Information…** from the main menu. As a result, a folder with log files will be opened.

Beside the Veeam Backup & Replication console, log files are also stored on ESX servers in folder */var/log/VeeamBackup/*.

Use log files to submit a support ticket. It is recommended that you send the whole content of the logs folder to ensure that overall and comprehensive information is provided to the support team.

VEEAM BACKUP ENTERPRISE MANAGER

Veeam Backup & Replication 5.0 comes with Veeam Backup Enterprise Manager — a management and reporting component that allows you to manage multiple Veeam Backup & Replication installations from a single web console.

In case of an enterprise with distributed architecture when a number of Veeam Backup & Replication instances are installed on different servers, Veeam Backup Enterprise Manager acts as a single management point, allowing you to perform backup and replication jobs across the entire VMware backup infrastructure, and providing enhanced reporting options.

With Veeam Backup Enterprise Manager, you can:

- Manage jobs across a number of Veeam Backup & Replication servers
- View on-going reporting data for all jobs
- Receive e-mail notifications about the status of all jobs
- Search for Windows guest files in current and archived backups
- Centrally monitor license usage and update them

A new distributed architecture of Veeam Backup & Replication 5.0 provides you with a possibility to create custom backup infrastructure meeting your company needs and manage backup and replication according to your administrative, business and security requirements and restrictions. While Veeam Backup Enterprise Manager provides centralized backup and reporting options, Veeam Backup & Replication servers still make it possible to perform decentralized backup and recovery. For example, Exchange recoveries can be handled by the Exchange administrators group, while domain controller recovery requires another skill set, and is best performed by Active Directory administrators.

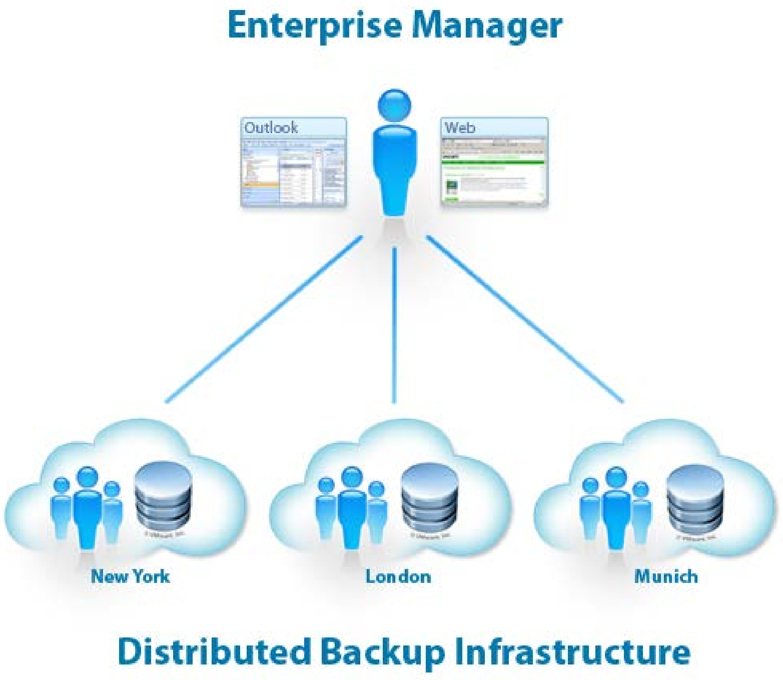

Configuring Veeam Backup Enterprise Manager

As soon as you start Veeam Backup Enterprise Manager for the first time, you should configure it to start working with backup servers. This section provides a detailed description of main configuration settings you should specify.

First Steps

To start working with Veeam Backup Enterprise Manager, follow the next steps:

1. Install Veeam Backup Enterprise Manager. Veeam Backup Enterprise Manager is installed as a separate component, either on the Veeam Backup & Replication console or apart from Veeam Backup & Replication 5.0. To learn more about the Veeam Backup Enterprise Manager installation, see the Installing Veeam Backup Enterprise Manager section.
2. Start Veeam Backup Enterprise Manager Web UI. If you are starting Enterprise Manager Web UI from the console on which it is installed, double-click the **Veeam Backup Enterprise Manager** icon on the desktop or select **Programs > Veeam > Veeam Backup Enterprise Manager** from the **Start** menu. If you are starting Enterprise Manager Web UI remotely, use HTTPS address *https://host-name/VeeamBackup* (for site on Windows XP 32) and *https://host-name:9443* (for site on other operating systems). As soon as you start Veeam Backup Enterprise Manager Web UI, you will be prompted to log on. Enter credentials of a user with local administrator rights or the user who installed Veeam Backup Enterprise Manager and click **Login**.
3. Configure backup server settings. Click the **Configuration** link at the top right corner of the main view. Then, click **Backup Servers** on the left and add all backup servers you want to manage. To learn more, see the Specifying Backup Server Settings section.
4. Collect job data for added backup servers. Click the **Start Collecting** button at the top of the Backup Servers view to collect data about all backup and replication jobs from added backup servers. You can also schedule data collection. To learn more, see the Collecting Data from Backup Servers section.
5. Configure security settings for the Veeam Backup Enterprise Manager. To be able to work with Veeam Backup Enterprise Manager, the user should be a member of Portal Administrators or Portal Viewers groups. To configure security settings for Veeam Backup Enterprise Manager, click the **Configuration** link at the top right corner of the window. Then, click **Roles** and configure groups as required. To learn more, see the Specifying Security Settings section.
6. Configure e-mail notification settings. To be able to receive e-mail notifications about the status and details of jobs, click **Notification** on the left of the Configuration view and specify e-mail notification settings. To learn more, see the Specifying Notification Settings section.

Once you have configured these settings, you can start working with managed backup servers.

Specifying Backup Server Settings

To start working with backup servers, you should add all servers you want to manage to Veeam Backup Enterprise Manager.

1. Click the **Configuration** link at the top of the main Veeam Backup Enterprise Manager view.
2. Click **Backup Servers** on the left of the Configuration view.

3. Click the **Add...** button at the top of the Backup Servers view.
4. In the **Backup Server Settings** window, enter a full DNS name or IP address of the server you want to add, and provide a server description.
5. By default, an account under which Veeam Backup Enterprise Manager Service is run is used for an added server. If this account does not have administrative rights on the server you want to add, select the **Use these credentials** check box and provide name and password of the user with administrative rights on the added server.
6. Specify the port used by Veeam Backup Service. By default, port 9392 is used.
7. Click the **OK** button to add the server.

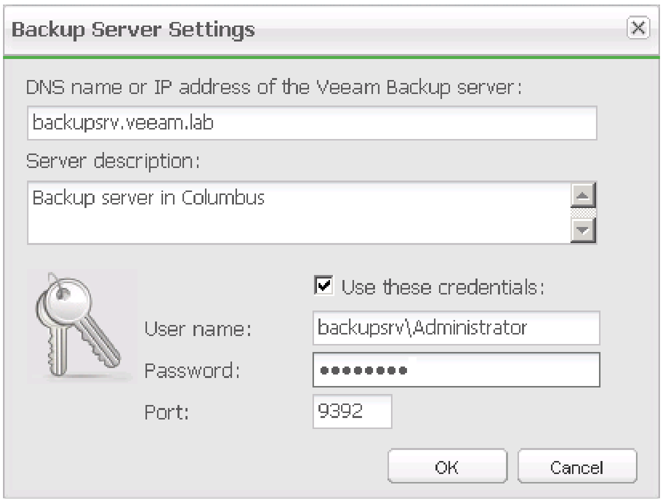

DNS name or IP address of the Veeam Backup server: backupsrv.veeam.lab

Server description: Backup server in Columbus

☑ Use these credentials:

User name: backupsrv\Administrator

Password: ••••••••

Port: 9392

To edit settings of an added server, select it in the Backup Servers list and click the **Edit...** button on the toolbar. Then, edit server connection settings as required.

To delete an added backup server, select it in the Backup Servers list and click the **Remove** button on the toolbar.

Collecting Data from Backup Servers

To retrieve data from added backup servers, Veeam Backup Enterprise Manager uses data a collection job. A data collection job represents a task for collecting information about backup and replication jobs from backup servers. Collected data is stored to the SQL backend and can be accessed by multiple users via the web browser.

Veeam Backup Enterprise Manager provides two options for running a data collection job:

- Running data collection job manually. To run a data collection job manually, click **Backup Servers** on the left of the **Configuration** view and click the **Start Collecting** button on the toolbar.
- Scheduling data collection job. To schedule a data collection job, click **Backup Servers** on the left of the **Configuration** view and click the **Schedule...** button on the toolbar. In the displayed window, select the **Periodically every...** option and specify an interval at which a data collection job should be run.

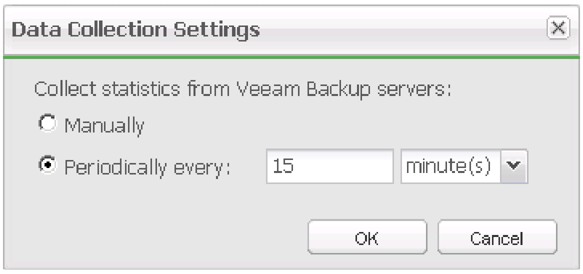

Collect statistics from Veeam Backup servers:

○ Manually

● Periodically every: 15 minute(s)

Note: When a data collection job is run, data from all added backup servers is collected at once.

Every run of a data collection job initiates a new job session. To view details on job sessions, click **Sessions** on the left of the Configuration view. In the list of sessions, select the one you need and follow the **click here** link in the **Log** column.

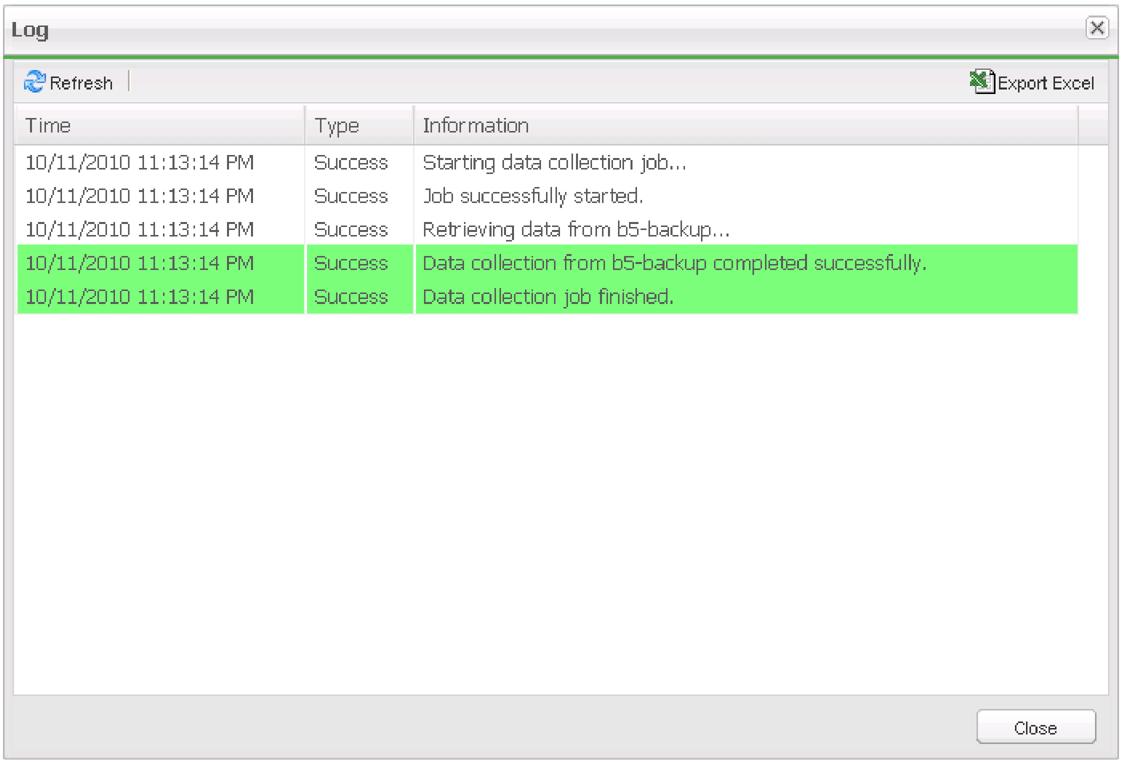

Time	Type	Information
10/11/2010 11:13:14 PM	Success	Starting data collection job...
10/11/2010 11:13:14 PM	Success	Job successfully started.
10/11/2010 11:13:14 PM	Success	Retrieving data from b5-backup...
10/11/2010 11:13:14 PM	Success	Data collection from b5-backup completed successfully.
10/11/2010 11:13:14 PM	Success	Data collection job finished.

Specifying Notification Settings

To be able to receive e-mail notification about the status of performed backup and replication jobs, you should configure e-mail notification settings.

1. Click the **Configuration** link at the top of the main Veeam Backup Enterprise Manager view.
2. Click **Notifications** on the left of the **Configuration** view.
3. In the **Email server settings** section, specify a full DNS name or IP address of the SMTP server that will be used for sending e-mail messages. Change port over which you want to communicate with the mail server if necessary. By default, port 25 is used.
4. (Optional) If your SMTP server requires SMTP authentication, select the **Requires authentication** check box and specify authentication credentials: login and password.
5. If you want to receive daily e-mail notifications, In the **Email notifications** section select the **Send daily notifications at** check box and specify the time at which a notification e-mail should be sent.
6. In the **From** field, enter an e-mail address of the notification sender.
7. In the **To** field, enter an e-mail address of the notification recipient. To specify multiple addresses, use a comma.
8. Enter a subject of e-mail notifications. You can use the following variables in the subject:
   %1 — number of jobs that ended with an error for the last 24 hours
   %2 — number of jobs that ended with a warning for the last 24 hours

%3 — number of jobs that ended successfully for the last 24 hours
%4 — number of jobs that ended with an errors for the last session
%5 — number of jobs that ended with a warning for the last session
%6 — number of jobs that ended successfully for the last session

After you configure e-mail settings, you will be able to receive e-mail notifications. A notification e-mail will contain a report about the number of jobs performed with the *Error*, *Warning* and *Success* statuses, and provide a link to Veeam Backup Enterprise Manager Web site so that you can see jobs statistics in detail.

9. If you want notifications to be sent when lab requests obtain some status, select the **Send notifications** check box in the **Lab request notifications** section. Specify addresses for the e-mail notification sender, recipient, the subject of the e-mail message and statuses on which the notification should be sent.

Note: To verify if you have configured e-mail settings correctly, use the **Test** button. Veeam Backup Enterprise Manager will send a test e-mail to specified e-mail addresses.

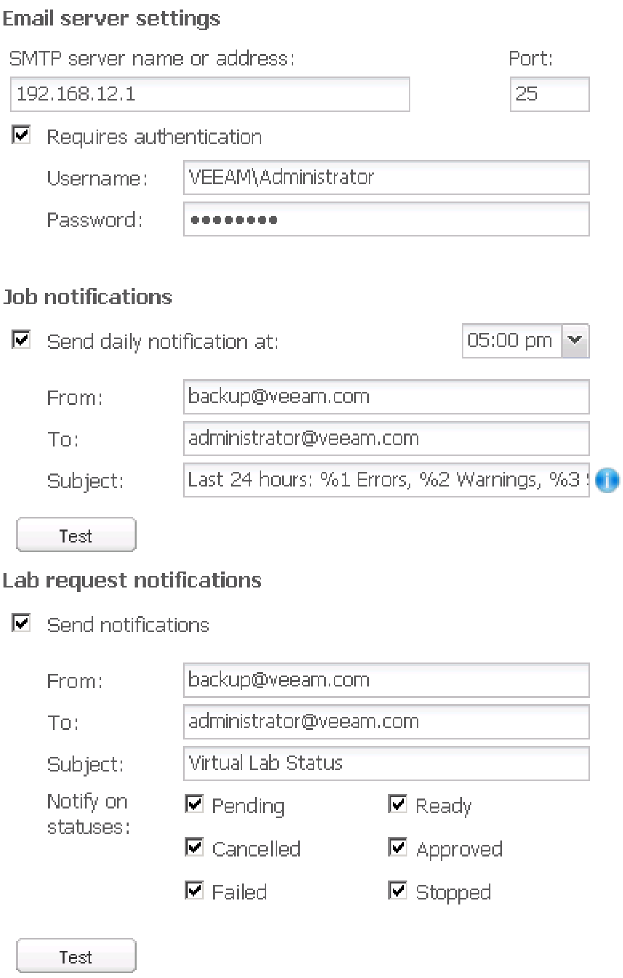

Email server settings

SMTP server name or address: 192.168.12.1

Port: 25

☑ Requires authentication

Username: VEEAM\Administrator

Password: ••••••••

Job notifications

☑ Send daily notification at: 05:00 pm

From: backup@veeam.com

To: administrator@veeam.com

Subject: Last 24 hours: %1 Errors, %2 Warnings, %3 !

Lab request notifications

☑ Send notifications

From: backup@veeam.com

To: administrator@veeam.com

Subject: Virtual Lab Status

Notify on statuses: ☑ Pending ☑ Ready ☑ Cancelled ☑ Approved ☑ Failed ☑ Stopped

Specifying Dashboard Settings

The **Dashboard Settings** view allows you to customize appearance of **Backup Servers** graphs.

To customize graph settings:

1. Click **Dashboard Settings** on the left of the **Configuration** view.

2. Use the **Activity graph scale** option to switch between graph types: *Linear* and *Logarithmic*.
3. By default, the **Backup Servers** graph on the **Last 24 hours** tab highlights a time interval for a planned backup window. You can change the width of a backup window and shift it backward or forward.
   By default, the **Show backup window** check box is selected. If you do not want to display a backup window on the graph, clear the check box.
4. Select necessary time values in the **Backup start time** and **Backup stop time** fields.
5. Click **Save** to save your changes.

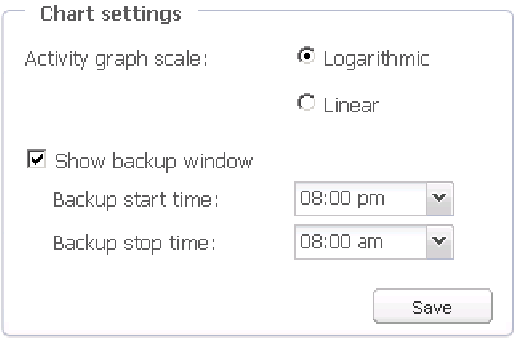

Chart settings

Activity graph scale: ● Logarithmic ○ Linear

☑ Show backup window

Backup start time: 08:00 pm

Backup stop time: 08:00 am

Specifying Security Settings

To configure a security scheme for distributed backup infrastructure with Veeam Backup Enterprise Manager, you should assign roles to users working with the Veeam Backup Enterprise Manager portal and Veeam Backup & Replication servers.

Veeam Backup Enterprise Manager Roles

Security settings of Veeam Backup Enterprise Manager are used to authenticate administrative rights for users working with Enterprise Manager Web UI. To be able to log on to the Veeam Backup Enterprise Manager Web UI, the user must be a member of either the Portal Administrators or the Portal Viewers group.

- Users listed in the *Portal Administrators* group can work with the main view of the portal (view on-going reporting data, manage jobs and so on), as well as configure the web portal settings.
- Users listed in the *Portal Viewers* group can work with the main view of the portal only, and cannot specify configuration settings for Veeam Backup Enterprise Manager.

Note: By default, the Portal Administrator role is assigned to users listed in the local Administrator groups and the user who installs Veeam Backup Enterprise Manager.

To specify security settings for a user or a group of users who should work with the Enterprise Manager Web UI:

1. Log on to Veeam Backup Enterprise Manager using an account with local administrator rights.
2. Click the **Configuration** link at the top of the main Veeam Backup Enterprise Manager view.
3. Click **Roles** on the left of the **Configuration** view.
4. Click the **Add...** button on the toolbar.
5. In the **Account type** field, select to which type of account you want the role to be assigned — *User* or *Group*.

6. In the **Account** field, specify a user account in the *DOMAIN\Username* format.
7. From the **Role** list, select a necessary portal role to be assigned — *Portal Administrator* or *Portal Viewer*.

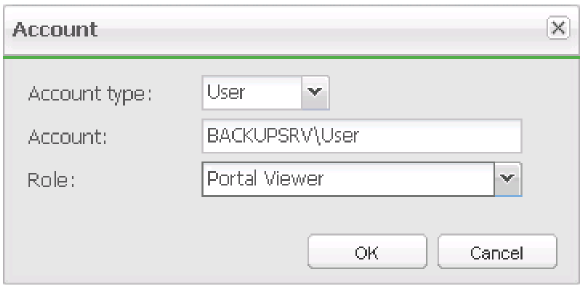

Account type: User
Account: BACKUPSRV\User
Role: Portal Viewer

To edit settings of an added user or group, select it in the list of roles and click the **Edit...** button on the toolbar. Then, edit user or group settings as required.

To delete an added user or group, select it in the list and click the **Remove** button on the toolbar.

Veeam Backup & Replication Roles

Security settings for the Veeam Backup & Replication server are used to authenticate user administrative rights for two operations: collecting data from backup servers and managing backup and replication jobs. Both operations are performed by Veeam Backup Service that verifies beforehand if the user has rights to accomplish these actions or not.

- To be able to retrieve data from backup servers, the user should be a member of *Veeam Backup Viewers* or *Veeam Backup Administrators* group on the backup server. Administrative rights should be granted to the account under which the *Veeam Backup Service* runs. Or, in case a backup server was added to Veeam Backup Enterprise Manager with other user credentials, administrative rights should be granted to the account that was used for adding the backup server.
- To be able to manage backup and replication jobs, the user should be listed in the *Veeam Backup Operators* or *Administrators* group on the backup server.

Note: By default, the Veeam Backup Administrator role is assigned to users listed in the local Administrators group.

To specify security settings on the backup server:

1. Select **Tools > Users and Roles...** from the main menu of Veeam Backup & Replication 5.0.
2. Click the **Add** button.
3. In the **User name** field, enter the name of a user or group in the *DOMAIN\Username* format.
4. From the **Role** list, select a necessary role to be assigned —*Veeam Backup Administrator*, *Veeam Backup Viewer* or *Veeam Backup Operator*.

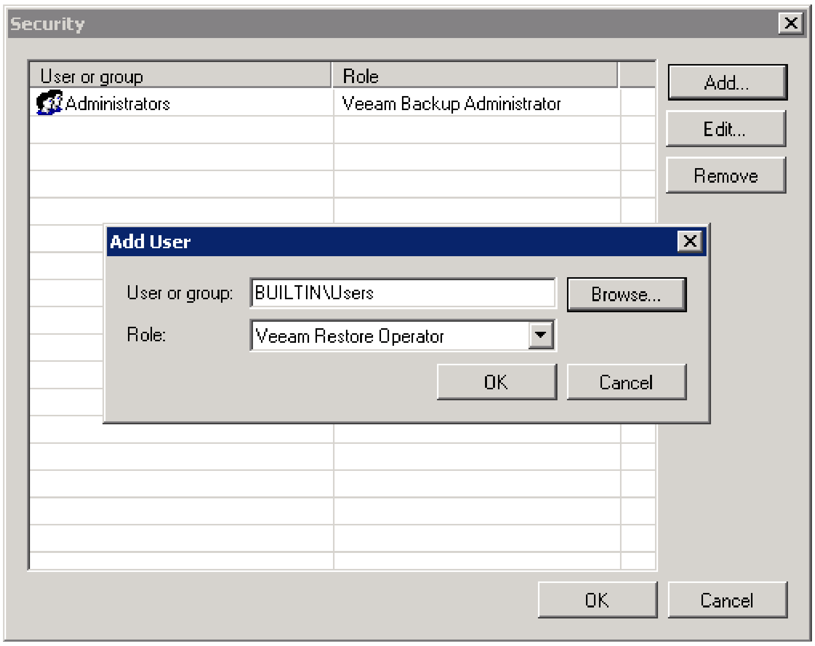

User or group	Role
Administrators	Veeam Backup Administrator

User or group: BUILTIN\Users

Role: Veeam Restore Operator

To edit settings of an added user or group, select it in the list of roles and click the **Edit** button on the right. Then, edit user or group settings as required.

To delete an added user or group, select it in the list and click the **Remove** button on the right.

Managing Licenses from Veeam Backup Enterprise Manager

Veeam Backup Enterprise Manager collects information about all licenses installed on backup servers that are connected to it. When Veeam Enterprise Manager replicates databases from backup servers, it also synchronizes license data — that is, checks if the license installed on the backup server coincides with the license installed on the Veeam Backup Enterprise Manager server. If the licenses do not coincide, the license on the backup server will be automatically updated with that on Veeam Backup Enterprise Manager.

Using the **Licensing** section of Veeam Backup Enterprise Manager, you can manage and activate licenses for the whole of the backup infrastructure from a single web console and thus reduce administration overhead.

To work with licenses for backup servers added to Veeam Backup Enterprise Manager:

1. Click the **Configuration** link at the upper right corner of the window.
2. Click **Licensing** on the left.

The upper pane of the **Licensing** section displays information on each installed license and used sockets. To update a license, click the **Update license** button and select a necessary .lic file.

The lower pane of the **Licensing** section displays information on all CPU sockets engaged in backup and replication jobs. You can revoke unused ESX servers from the license — that is, to re-use the license applied to one ESX server to another ESX server. This may be required if the ESX server to which the license is applied does not need backup or replication anymore (for example, in case it is no longer used).

To revoke a server, select it in the list and click the **Revoke…** button at the top of the section.

Managing Backup and Replication Jobs

Veeam Backup Enterprise Manager acts as a single point for managing backup and replication jobs from all added backup servers. To view a list of jobs, click the **Jobs** tab in the main view of Veeam Backup Enterprise Manager.

You can centrally run jobs from all added backup servers.

- To run a job, select it in the list and click the **Start job** button on the toolbar.
- To stop a job, select it in the list and click the **Stop job** button on the toolbar.
- To re-run a job that has failed, select it in the list and click the **Retry job** button.

Tip: To facilitate job search, use the filter at the top of the list. You can filter jobs by one or a number of filtering criteria: backup server, last job result and job name. Once you have selected necessary filter criteria, click the **Find** button to apply a filter to the list.

Viewing Jobs Reporting Data

Along with a possibility to manage backup and replication jobs, Veeam Backup Enterprise Manager provides a convenient way to work with job data collected from a number of sources — backup servers. Veeam Backup Enterprise Manager offers a wide range of reporting options, presenting information about performed jobs in various profiles. Being a common business requirement to IT infrastructure, reports allow you to get granular information about jobs created on managed backup servers, and data related to jobs, namely:

- Jobs performed for the last 24 hours
- Jobs performed for the last 7 days
- Data for all performed jobs
- Data for all VMs engaged in jobs
- Data about specific job sessions, and so on

You can both view reporting data on the on-going basis using the web browser, and export it to files of the Excel format which can be saved for documenting and archiving purposes.

On-Going Jobs Data

Veeam Backup Enterprise Manager displays on-going data for two time periods — data collected for the last day and data collected for the last week. To see on-going jobs data, click the **Last 24 hours** or **Last 7 days** tab, correspondingly.

To visualize jobs data, Veeam Backup Enterprise Manager uses graphs, informing about time and date when jobs were performed, and the network throughput rate.

Jobs relating to one backup server are marked with a separate color on the graph. The legend on the right interprets the color scheme used for all managed backup servers.

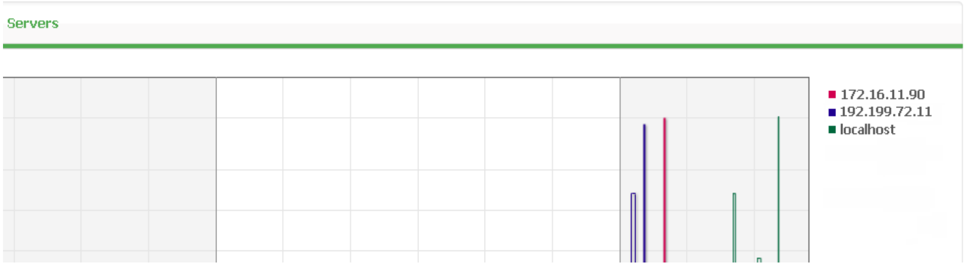

Beside a graph, the Last 24 hours and Last 7 days views provide information for the following data related to performed jobs:

- The **Summary** block reports on the total number of managed backup servers, jobs, processed VMs and VM templates
- The **Data block** reports on the average processing speed, total size of processed VMs, size of backups and the average of compression ratio
- The **Last 24 hours/7 days** block reports on the total number of jobs, and jobs completed with different statuses
- The **Status block** reports on the health status of managed backup servers and Veeam Backup Enterprise Manager

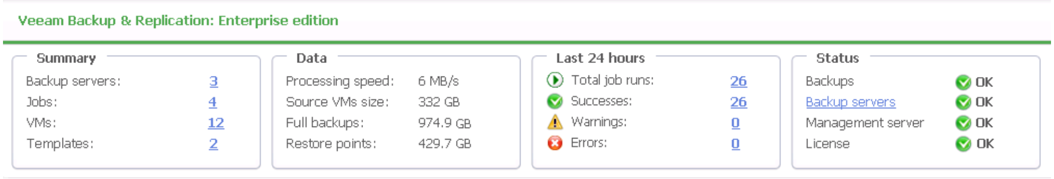

Veeam Backup & Replication: Enterprise edition

Summary	
Backup servers:	3
Jobs:	4
VMs:	12
Templates:	2

Data	
Processing speed:	6 MB/s
Source VMs size:	332 GB
Full backups:	974.9 GB
Restore points:	429.7 GB

Last 24 hours	
Total job runs:	26
Successes:	26
Warnings:	0
Errors:	0

Status	
Backups	✓ OK
Backup servers	✓ OK
Management server	✓ OK
License	✓ OK

Job Data

To view information about all jobs from managed backup servers, click the **Jobs** tab in the main view of Veeam Backup Enterprise Manager.

Every job in the list is described with the following data: job name, type, backup server on which a job was created, current job state, date of the latest run, date of the next run (if the job is scheduled) and description.

Beside information presented in the list of jobs, the Jobs tab allows you to view advanced job data: number of sessions for each job and detailed statistics for a job session.

- To see a list of job sessions, click the job name link in the *Name* column.
- To see detailed statistics on the last job run, click the state link in the *Current State* column.

VM Data

The VM tab provides information about all VMs engaged in performed jobs: VM name, path to a backup file, number of restore points, backup server to which the job relates, job name and status of the last job run.

Tip: To display detailed information about a VM, click its name in the VM column. To display detailed information about VM restore points, click a link in the **Restore Points** column.

Reports Data

The **Reports** tab of Veeam Backup Enterprise Manager allows you to granulate information for managed backup servers in the following succession: servers > jobs > job sessions > session details.

To get the information you need, click the **Reports** tab. Then, click a necessary link in a corresponding column of the displayed view.

Tip: You can export displayed information to a file of Excel format at any moment of time — to do so, click the **Export Excel** button on the toolbar.

Searching for VM Guest Files

Veeam Backup Enterprise Manager enables you to perform quick and accurate search for guest OS files in a backed up VM without the need to restore it. This can be useful, for example, if a file you need has been deleted on the VM and you want to restore it from a backup. Once you find a necessary file, you can use Veeam's file-level restore to recover the file from the VM backup.

At present, the search functionality is supported for Windows-based VMs only; however, it will be expanded to other file systems in future releases.

To be able to perform search within VM image backup, you need to enable file indexing in properties of a corresponding backup job. When such a backup job is run, Veeam Backup & Replication creates a catalog, or index, of the VM guest OS files and stores index files on the Veeam Backup server in the *C:/VBR Catalog/Index/Machines/[vm_name]* folder. Creation of index is extremely fast — the search engine works in the background outside the backup window and has minimal impact on network and VMware environment.

Once the index is created and stored on backup servers, the indexing service on Veeam Backup Enterprise Manager performs index replication — it aggregates index data for all VM image backups from multiple backup servers. This consolidated index is stored on the Veeam Backup Enterprise Manager server in the *C:/VBR Catalog/Index/* catalog and is used for search queries.

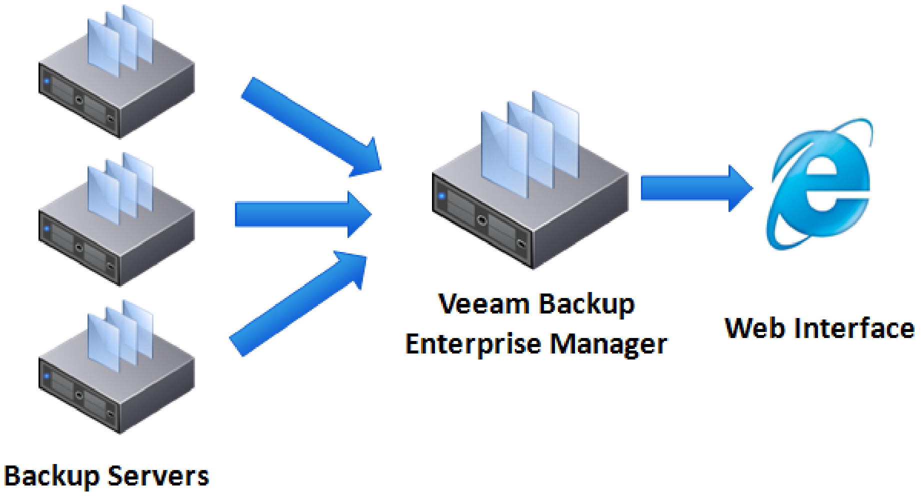

Veeam Backup Enterprise Manager offers two options of search for guest OS files in indexed VM backups:

- **Browsing through VM guest file system** — this option lets you browse inside the guest file system of a selected VM and perform quick file search through its guest OS files.
- **Performing advanced search** — this option lets you search across all backup servers in your backup infrastructure and quickly find a necessary file in any created VM backup.
  To enable advanced search, you need to configure a search server and add it to Veeam Backup Enterprise Manager. When performing advanced search queries, Veeam utilizes Veeam Backup Search — a special tool that installed on a dedicated Microsoft Search Server. Veeam Backup Search uses Microsoft Search Server functionality to crawl aggregated index files on Veeam Backup Enterprise Manager and create a content index on the search server that is used to serve search queries.

Note: Even if VM backups were moved to an external storage device or tape, indexing data for such VMs still remain in the catalog and will be displayed in search results. You can use the **Import** feature in Veeam Backup & Replication to import the backup to the backup server, and then recover the file.

Preparing for File Browsing and Advanced Searching

This section describes what settings and system components you should configure to be able to use file browsing and advanced search functionality.

First Steps

If you already have Veeam Backup & Replication and Veeam Backup Enterprise Manager installed, you need to perform the following steps to use the file browsing feature.

1. **Enable file indexing in backup job properties**. At the **Backup Consistency** step of the backup job, enable file indexing for a VM. To learn more, see the Creating a Backup Job section.
2. **Run the backup job** with file indexing enabled.
3. **Perform catalog replication**. Open the **Search Servers** view in Veeam Backup Enterprise Manager and click the **Sync Catalog Now** button on the toolbar. To learn more, see the Performing Catalog Replication and Indexing section.

To enable advanced search, you will have to additionally perform the following steps:

1. **Install Microsoft Search Server** on a dedicated machine. Veeam Backup & Replication can work with Microsoft Search Server 2008/Microsoft Search Server Express 2008 and Microsoft Search Server 2010/Microsoft Search Server Express 2010. Keep in mind that Microsoft Search Server can be installed on machines running Windows Server only. To learn more about hardware and software requirements, see http://technet.microsoft.com/en-gb/library/bb905370(office.12).aspx (for Microsoft Search Server 2008) and http://technet.microsoft.com/en-gb/library/bb905370.aspx (for Microsoft Search Server 2010).
2. **Install Veeam Backup Search** on the machine with Microsoft Search Server. To learn more, see the Installing Veeam Backup Search section.
3. **Add a search server** to Veeam Backup Enterprise Manager. Click **Configuration** link at the top of the main Veeam Backup Enterprise Manager view. Click **Search Servers** on the left and click the **Add** button. To learn more, see the Adding a Search Server section.

Performing Catalog Replication and Indexing

Once you have run backup jobs with file indexing enabled, you need to perform catalog replication to consolidate index files from multiple backup servers. During this operation, Veeam Backup Enterprise Manager aggregates index data from multiple backup servers and stores them on the Veeam Backup Enterprise Manager server to enable file browsing and advanced search.

Note: Catalog replication is performed for VM images with indexed guest OS file systems on all managed backup servers.

Veeam Backup Enterprise Manager provides two options to perform catalog replication:

- To perform manual catalog replication, select **Search Servers** on the left of the **Configuration** view and click **Sync Catalog Now** on the toolbar.
- To automatically run catalog replication after every backup job, select **Search Servers** on the left of the **Configuration** view and click **Schedule** on the toolbar. In the displayed window, select **Automatically after every backup job** and specify other options as necessary.

Every run of a catalog replication job initiates a new job session which can be tracked under **Sessions** in the **Configuration** view. To view detailed information for a specific session, find it in the **Sessions** view and click the corresponding **click here** link in the **Log** column.

Adding a Search Server

If you are planning to use advanced file search across all backup servers in your backup infrastructure, you should configure at least one search server and add it to Veeam Backup Enterprise Manager. Search servers are not required if you are planning to perform file browsing operations only.

The capacity of a search server is limited and depends on the type of search server you are planning to use. If you have a great number of backup servers and/or require storing index documents for a long period of time, you may want to deploy a number of search servers. In this case, the query processing and indexing load will be automatically spread across all deployed search servers.

To add a search server:

1. Click the **Configuration** link at the top of the main Veeam Backup Enterprise Manager view.
2. Click **Search Servers** in the left pane of the **Configuration** view.
3. Click **Add** at the top of the **Search Servers** view.
4. In the **Search Server Settings** window, enter a full DNS name or IP address of the server you want to add (provide a description if necessary).
5. By default, the account under which Veeam Backup Service is running will be used for the added server. To specify a different account, select **Use these credentials** and provide a name and a password of a user with administrative privileges on the search server.
6. Specify the port used by Veeam Backup Service. By default, port number 9395 is used.
7. In the **Capacity** section, select the type of database server you are planning to use to set a recommended capacity value for the indexing server. Specify the limit of index documents to be created. By default, the limit is set to 300,000 documents.
8. Click **OK** to add the server.

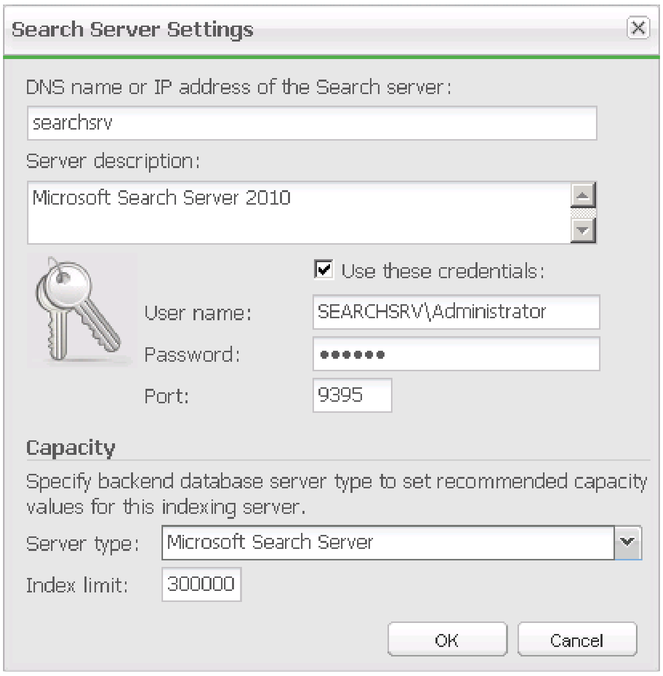

Search Server Settings

DNS name or IP address of the Search server: searchsrv

Server description: Microsoft Search Server 2010

☑ Use these credentials:

User name: SEARCHSRV\Administrator

Password: ••••••

Port: 9395

Capacity

Specify backend database server type to set recommended capacity values for this indexing server.

Server type: Microsoft Search Server

Index limit: 300000

Note: Microsoft Search Server crawls content in the shared *VBRCatalog* folder on the Veeam Backup Enterprise Manager server. Therefore, the *VBRCatalog* folder should be seen to the search server, and the search server should be granted access to the folder. When you click **OK** to add a search server, Veeam Backup Enterprise Manager will verify if required permissions are granted, and display a warning message if these permissions are not enough.

To edit the settings of an existing server, select it in the **Search Servers** list and click **Edit** on the toolbar. To remove a search server, select it in the **Search Servers** list and click **Remove** on the toolbar.

Browsing Guest OS Files in VM Backups

After you add a search server and perform catalog replication, you can browse VM backups for OS guest files. File browsing does not require you to configure and add a search server. However, in contrast to the advanced search functionality, it allows you to browse and search for files in the selected VM backup at a specific restore point only.

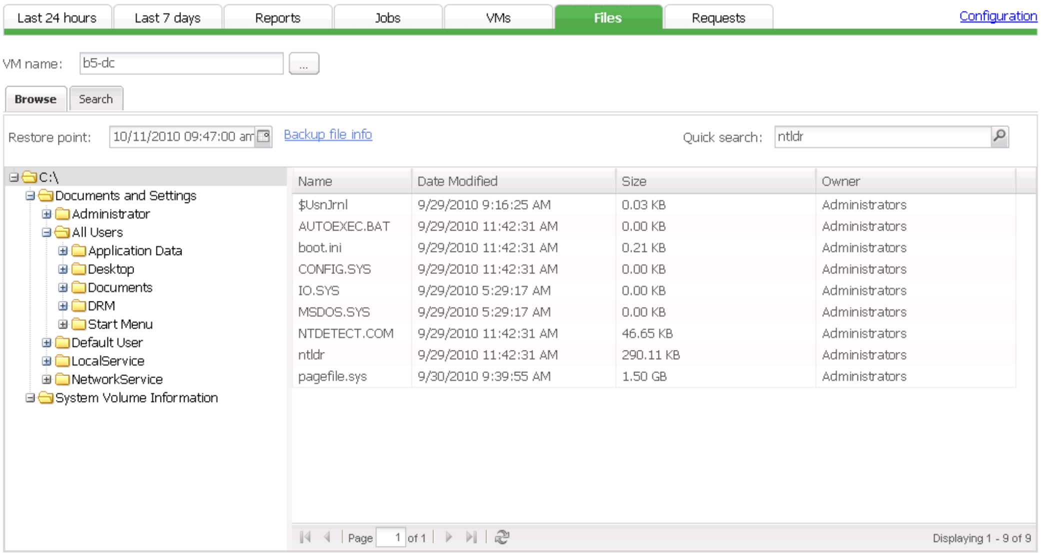

VM name: b5-dc

Restore point: 10/11/2010 09:47:00 am

Quick search: ntldr

- C:\
  - Documents and Settings
    - Administrator
    - All Users
      - Application Data
      - Desktop
      - Documents
      - DRM
      - Start Menu
    - Default User
    - LocalService
    - NetworkService
  - System Volume Information

Name	Date Modified	Size	Owner
$UsnJrnl	9/29/2010 9:16:25 AM	0.03 KB	Administrators
AUTOEXEC.BAT	9/29/2010 11:42:31 AM	0.00 KB	Administrators
boot.ini	9/29/2010 11:42:31 AM	0.21 KB	Administrators
CONFIG.SYS	9/29/2010 11:42:31 AM	0.00 KB	Administrators
IO.SYS	9/29/2010 5:29:17 AM	0.00 KB	Administrators
MSDOS.SYS	9/29/2010 5:29:17 AM	0.00 KB	Administrators
NTDETECT.COM	9/29/2010 11:42:31 AM	46.65 KB	Administrators
ntldr	9/29/2010 11:42:31 AM	290.11 KB	Administrators
pagefile.sys	9/30/2010 9:39:55 AM	1.50 GB	Administrators

To browse for guest OS files in a VM backup:

1. Click the **Files** tab in Veeam Backup Enterprise Manager.
2. Click the **Browse** tab.
3. In the **VM name** field, select the VM that you want to browse.
4. In the **Restore point** field, select a necessary date of backup and a restore point. Note that dates when backup of the selected VM was performed are marked green in the calendar.

As a result, the file tree of the VM as of the selected backup and restore point date will be displayed. You can manually browse the file tree to find a necessary file, or use the **Quick search** field and the top left corner.

Depending on the number of files on the VM, the search may take some time. Results are presented as a list with entries in the following format:

C:\WINDOWS\Folder\file.exe

VM: Virtual Machine

Owner: domain\user

Last modified: 7/8/2010 11:23:49 pm

Size: 1.68 MB

To get detailed data about a file, click the file name link. The file properties window will display information about the VM containing the file, backup and restore point you can use to recover the file. You can print out this information by clicking the **Print** button at the bottom of the window.

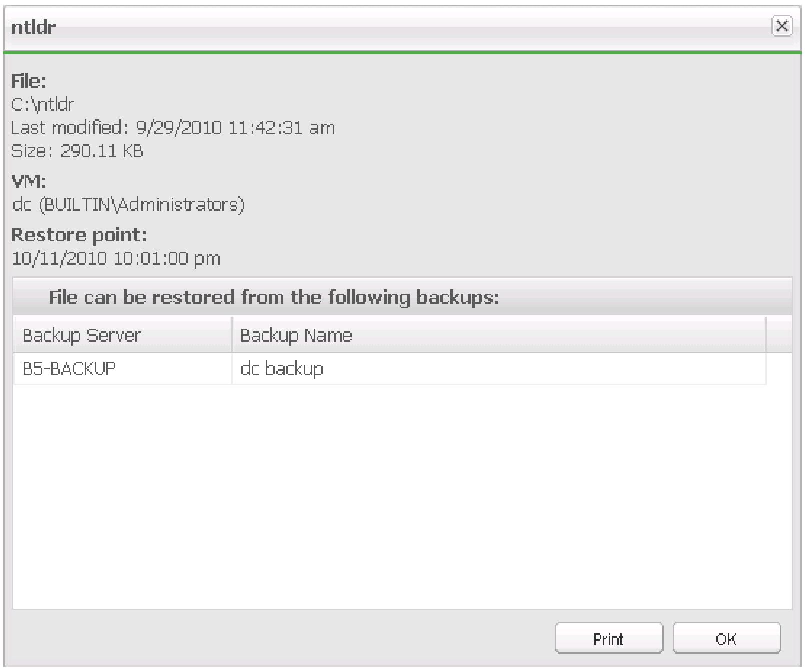

ntldr

File:
C:\ntldr
Last modified: 9/29/2010 11:42:31 am
Size: 290.11 KB

VM:
dc (BUILTIN\Administrators)

Restore point:
10/11/2010 10:01:00 pm

File can be restored from the following backups:

Backup Server	Backup Name
BS-BACKUP	dc backup

Using Advanced Search

Using the advanced search in Veeam Backup Enterprise Manager, you can search for guest OS files in all VM backups created by different backup servers in your backup infrastructure. The advances search functionality can only be used in you configure a search server and add it to Veeam Backup Enterprise Manager.

To perform advanced search, click the **Files** tab in Veeam Backup Enterprise Manager. Then, click the **Search** tab.

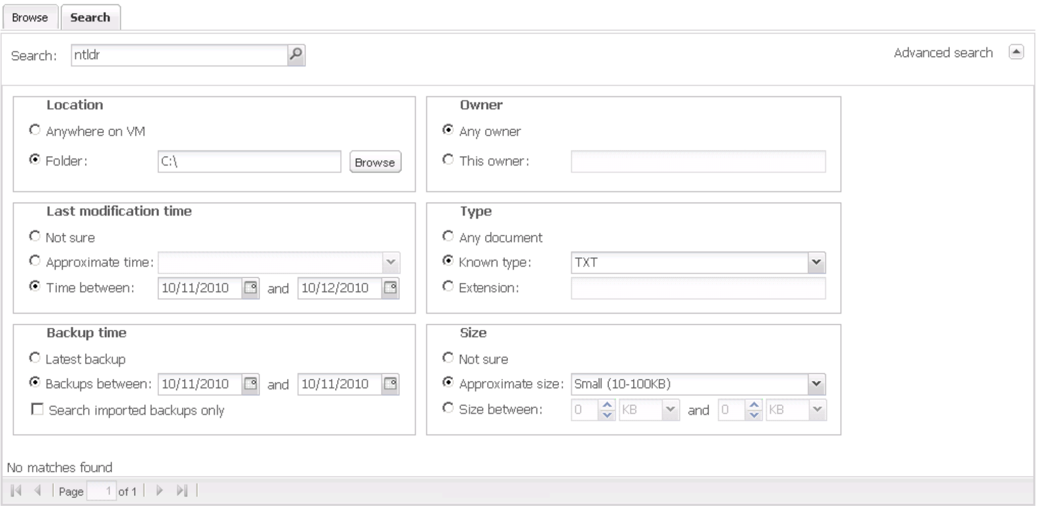

Search: ntldr

Location
○ Anywhere on VM
● Folder: C:\ Browse

Owner
● Any owner
○ This owner:

Last modification time
○ Not sure
○ Approximate time:
● Time between: 10/11/2010 and 10/12/2010

Type
○ Any document
● Known type: TXT
○ Extension:

Backup time
○ Latest backup
● Backups between: 10/11/2010 and 10/11/2010
□ Search imported backups only

Size
○ Not sure
● Approximate size: Small (10-100KB)
○ Size between: 0 KB and 0 KB

No matches found

The **Search** menu offers the following options:

- **Location** – select a specific folder on the VM to search in.
- **Last modification time** – specify approximate time when the file was last modified or set a time interval.
- **Backup time** – choose to search through the latest backup of the specified VM or all backups of the VM created within a certain time interval.
- **Owner** – select to search for files with a specific owner.
- **Type** – select to search for files of specific type or with a certain extension.
- **Size** – specify approximate size of file or set a size range.

Working with Virtual Lab Requests

The **Requests** tab allows you to approve and reject virtual lab requests, as well as prolong the time of virtual lab existence as a part of the U-AIR process. To learn more, see the Approving Virtual Lab Requests section.

POWERSHELL ACCESS

Veeam Backup & Replication 5.0 comes with PowerShell extension — a snap-in to Microsoft Windows PowerShell 2.0. Windows PowerShell is a powerful command-line tool that allows administrators to automate some Veeam Backup & Replication activities. Veeam extends functionality of Windows PowerShell 2.0, and now administrators may use PowerShell to automate Veeam backup, replication and copy job creation and editing, VMs restores, replica failover and other operations.

Before installing Veeam PowerShell snap-in, make sure that Microsoft Windows PowerShell 2.0 is installed on the Veeam Backup & Replication console. To download Microsoft Windows PowerShell, use the following link: http://support.microsoft.com/kb/968929.

PowerShell uses cmdlets — simple single-function commands that can be run in the command-line shell. Cmdlets are specialized .NET classes that implement specific actions. Veeam PowerShell provides a set of its own cmdlets which correspond to actions you can perform via Veeam Backup & Replication UI. Please keep in mind that actions performed with PowerShell have the same force as actions performed via Veeam Backup & Replication 5.0 — for example, if you delete some job with PowerShell scripts, the job will be removed from the *VeeamBackup* database, and you will not be able undo changes.

Work with Veeam PowerShell cmdlets and scripts in many respects depends on your imagination, skills and expertise in Windows PowerShell 2.0. To learn more about Windows PowerShell 2.0 and its basics, use the book by Dr. Tobias Weltner: Microsoft Master-PowerShell.

Important! Please keep in mind that Veeam support team does not write PowerShell scripts on demand.

Cmdlets operate with objects which they can accept and return. Every object has properties describing it, and methods that can be performed on it.

For example, the *Get-VBRJob* cmdlet has the following output:

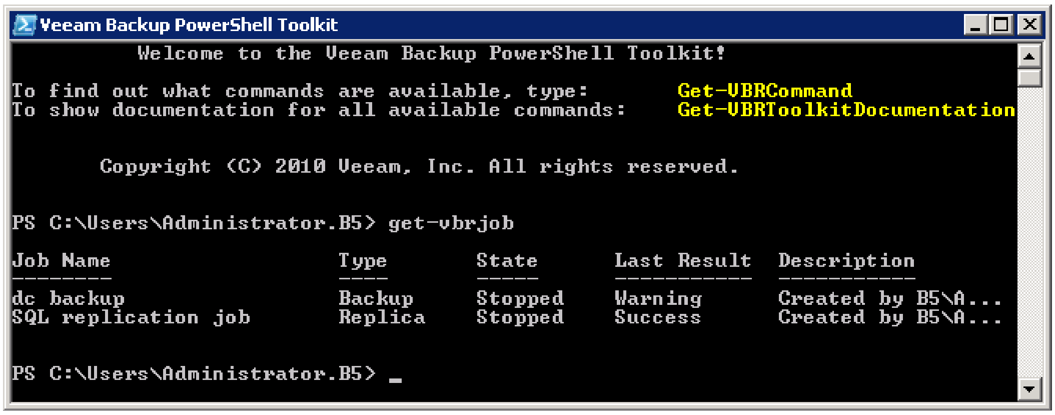

```
          Welcome to the Veeam Backup PowerShell Toolkit!

To find out what commands are available, type:       Get-VBRCommand
To show documentation for all available commands:    Get-VBRToolkitDocumentation


       Copyright (C) 2010 Veeam, Inc. All rights reserved.


PS C:\Users\Administrator.B5> get-vbrjob

Job Name                  Type       State      Last Result  Description
--------                  ----       -----      -----------  -----------
dc backup                 Backup     Stopped    Warning      Created by B5\A...
SQL replication job       Replica    Stopped    Success      Created by B5\A...


PS C:\Users\Administrator.B5> _
```

You can filter and sort results at your discretion using the (*where*) and (*sort*) commands (to learn more about these commands, please refer to Windows PowerShell documentation).

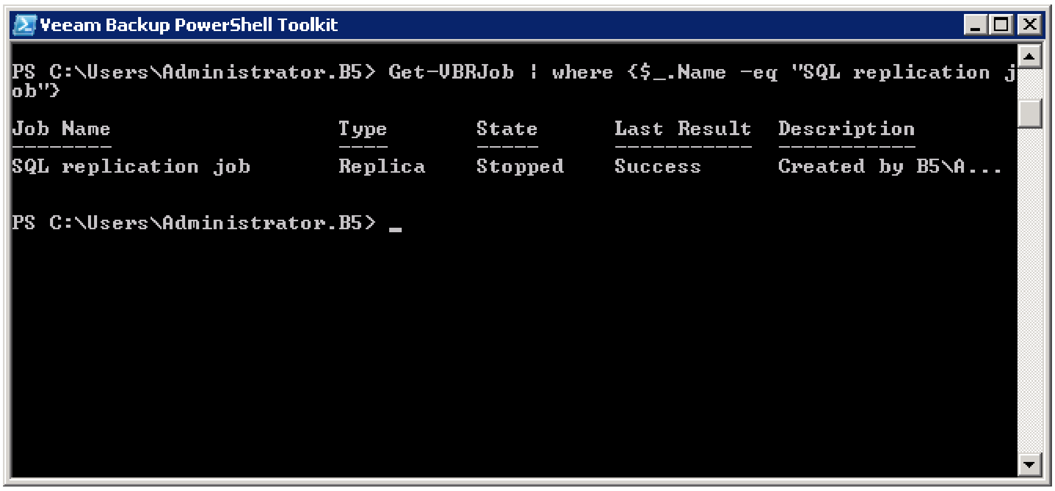

```
PS C:\Users\Administrator.B5> Get-VBRJob | where {$_.Name -eq "SQL replication j
ob"}

Job Name                  Type        State       Last Result  Description
--------                  ----        -----       -----------  -----------
SQL replication job       Replica     Stopped     Success      Created by B5\A...


PS C:\Users\Administrator.B5> _
```

The result of filtering and sorting can be saved to a variable:

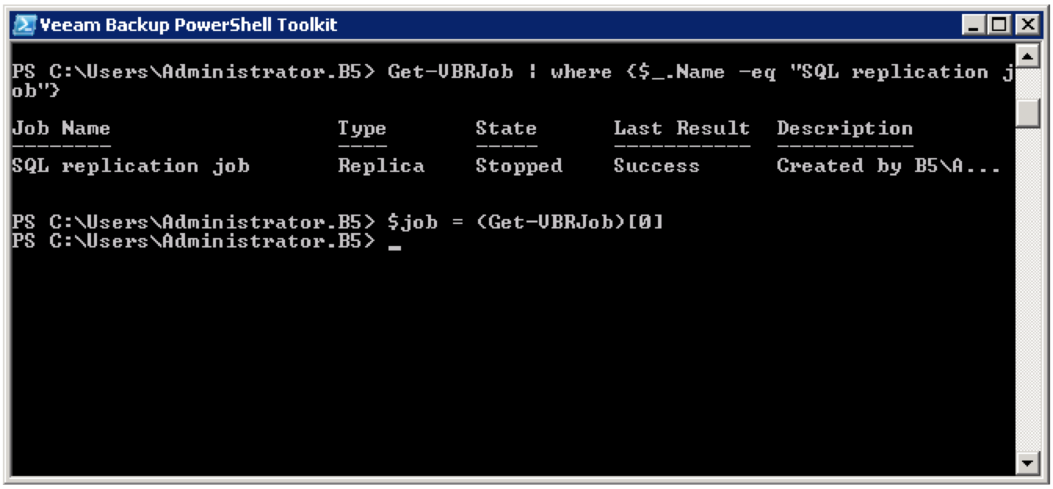

```
PS C:\Users\Administrator.B5> Get-VBRJob | where {$_.Name -eq "SQL replication j
ob"}

Job Name                  Type        State       Last Result  Description
--------                  ----        -----       -----------  -----------
SQL replication job       Replica     Stopped     Success      Created by B5\A...


PS C:\Users\Administrator.B5> $job = (Get-VBRJob)[0]
PS C:\Users\Administrator.B5> _
```

You can get detailed information about the object:

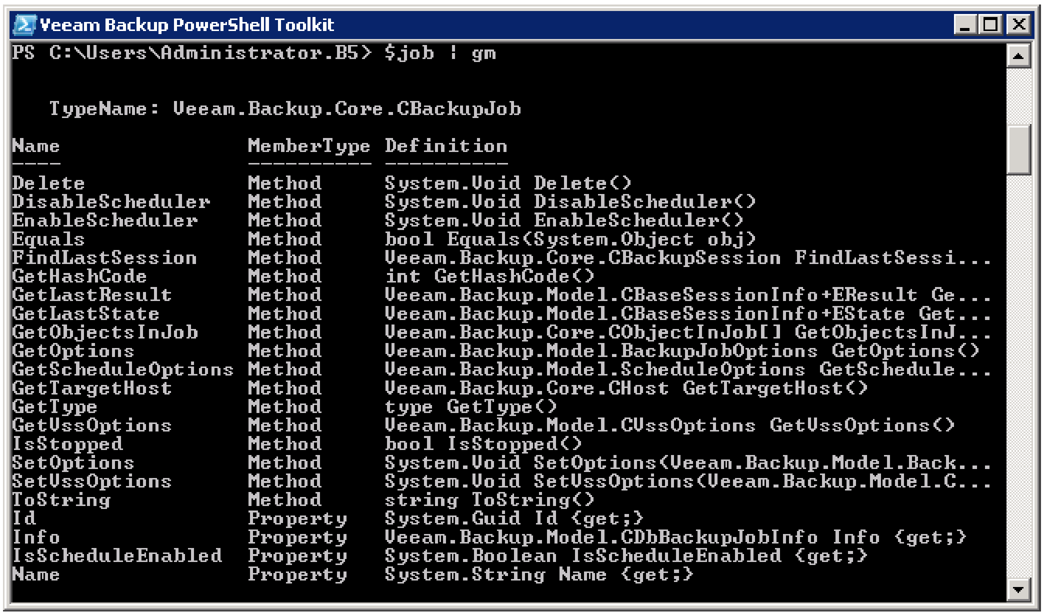

```
PS C:\Users\Administrator.B5> $job | gm


   TypeName: Veeam.Backup.Core.CBackupJob

Name                 MemberType Definition
----                 ---------- ----------
Delete               Method     System.Void Delete()
DisableScheduler     Method     System.Void DisableScheduler()
EnableScheduler      Method     System.Void EnableScheduler()
Equals               Method     bool Equals(System.Object obj)
FindLastSession      Method     Veeam.Backup.Core.CBackupSession FindLastSessi...
GetHashCode          Method     int GetHashCode()
GetLastResult        Method     Veeam.Backup.Model.CBaseSessionInfo+EResult Ge...
GetLastState         Method     Veeam.Backup.Model.CBaseSessionInfo+EState Get...
GetObjectsInJob      Method     Veeam.Backup.Core.CObjectInJob[] GetObjectsInJ...
GetOptions           Method     Veeam.Backup.Model.BackupJobOptions GetOptions()
GetScheduleOptions   Method     Veeam.Backup.Model.ScheduleOptions GetSchedule...
GetTargetHost        Method     Veeam.Backup.Core.CHost GetTargetHost()
GetType              Method     type GetType()
GetVssOptions        Method     Veeam.Backup.Model.CVssOptions GetVssOptions()
IsStopped            Method     bool IsStopped()
SetOptions           Method     System.Void SetOptions(Veeam.Backup.Model.Back...
SetVssOptions        Method     System.Void SetVssOptions(Veeam.Backup.Model.C...
ToString             Method     string ToString()
Id                   Property   System.Guid Id {get;}
Info                 Property   Veeam.Backup.Model.CDbBackupJobInfo Info {get;}
IsScheduleEnabled    Property   System.Boolean IsScheduleEnabled {get;}
Name                 Property   System.String Name {get;}
```

And call a method or a property of the object, for example, view last session results:

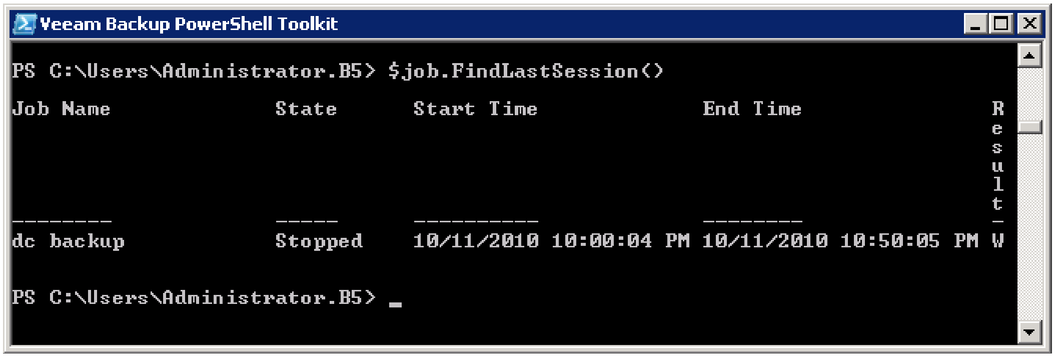

Veeam Backup PowerShell Toolkit

```
PS C:\Users\Administrator.B5> $job.FindLastSession()

Job Name                 State      Start Time              End Time                R
                                                                                    e
                                                                                    s
                                                                                    u
                                                                                    l
                                                                                    t
--------                 -----      ----------              --------                -
dc backup                Stopped    10/11/2010 10:00:04 PM 10/11/2010 10:50:05 PM W

PS C:\Users\Administrator.B5> _
```

Since Veeam Backup & Replication 5.0 uses Windows PowerShell 2.0, scripts that you have created with previous versions of Veeam Backup & Replication which used Windows PowerShell 1.0 may be not working.

Important! In Windows Vista and later, you must run PowerShell with elevated permissions if you are already an administrator and User Account Control (UAC) is enabled. To run PowerShell under elevated permissions in Windows Vista and later, right-click its shortcut and choose **Run as administrator**.

If you're logged on to Windows XP or Windows 2003 as a standard user, you can right-click the PowerShell shortcut, choose **Run as…**, and enter administrator account credentials.

Interactive PowerShell Help

Veeam PowerShell includes an interactive console-based help for each cmdlet. To get a cmdlet description, you can use the following commands:

Get-VBRcommand – displays a list of all Veeam PoweShell cmdlets that can be used.

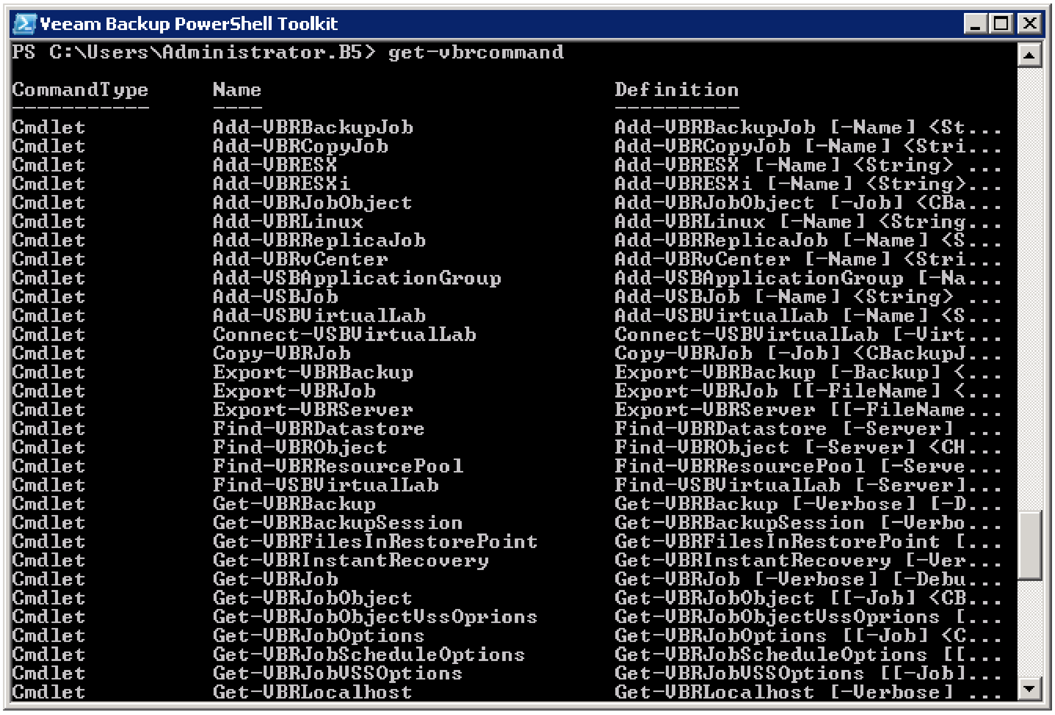

Veeam Backup PowerShell Toolkit

```
PS C:\Users\Administrator.B5> get-vbrcommand

CommandType     Name                             Definition
-----------     ----                             ----------
Cmdlet          Add-VBRBackupJob                 Add-VBRBackupJob [-Name] <St...
Cmdlet          Add-VBRCopyJob                   Add-VBRCopyJob [-Name] <Stri...
Cmdlet          Add-VBRESX                       Add-VBRESX [-Name] <String> ...
Cmdlet          Add-VBRESXi                      Add-VBRESXi [-Name] <String>...
Cmdlet          Add-VBRJobObject                 Add-VBRJobObject [-Job] <CBa...
Cmdlet          Add-VBRLinux                     Add-VBRLinux [-Name] <String...
Cmdlet          Add-VBRReplicaJob                Add-VBRReplicaJob [-Name] <S...
Cmdlet          Add-VBRvCenter                   Add-VBRvCenter [-Name] <Stri...
Cmdlet          Add-VSBApplicationGroup          Add-VSBApplicationGroup [-Na...
Cmdlet          Add-VSBJob                       Add-VSBJob [-Name] <String> ...
Cmdlet          Add-VSBVirtualLab                Add-VSBVirtualLab [-Name] <S...
Cmdlet          Connect-VSBVirtualLab            Connect-VSBVirtualLab [-Virt...
Cmdlet          Copy-VBRJob                      Copy-VBRJob [-Job] <CBackupJ...
Cmdlet          Export-VBRBackup                 Export-VBRBackup [-Backup] <...
Cmdlet          Export-VBRJob                    Export-VBRJob [[-FileName] <...
Cmdlet          Export-VBRServer                 Export-VBRServer [[-FileName...
Cmdlet          Find-VBRDatastore                Find-VBRDatastore [-Server] ...
Cmdlet          Find-VBRObject                   Find-VBRObject [-Server] <CH...
Cmdlet          Find-VBRResourcePool             Find-VBRResourcePool [-Serve...
Cmdlet          Find-VSBVirtualLab               Find-VSBVirtualLab [-Server]...
Cmdlet          Get-VBRBackup                    Get-VBRBackup [-Verbose] [-D...
Cmdlet          Get-VBRBackupSession             Get-VBRBackupSession [-Verbo...
Cmdlet          Get-VBRFilesInRestorePoint       Get-VBRFilesInRestorePoint [...
Cmdlet          Get-VBRInstantRecovery           Get-VBRInstantRecovery [-Ver...
Cmdlet          Get-VBRJob                       Get-VBRJob [-Verbose] [-Debu...
Cmdlet          Get-VBRJobObject                 Get-VBRJobObject [[-Job] <CB...
Cmdlet          Get-VBRJobObjectVssOprions       Get-VBRJobObjectVssOprions [...
Cmdlet          Get-VBRJobOptions                Get-VBRJobOptions [[-Job] <C...
Cmdlet          Get-VBRJobScheduleOptions        Get-VBRJobScheduleOptions [[...
Cmdlet          Get-VBRJobVSSOptions             Get-VBRJobVSSOptions [[-Job]...
Cmdlet          Get-VBRLocalhost                 Get-VBRLocalhost [-Verbose] ...
```

Get-help <*cmdletname*> - displays a full description of the specified cmdlet along with its syntax and full parameter description.

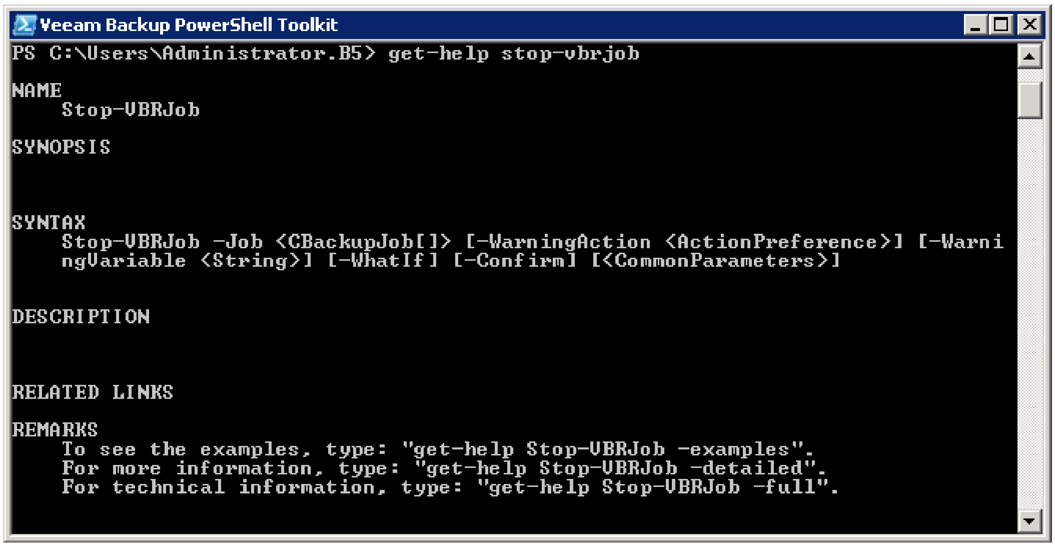

```
PS C:\Users\Administrator.B5> get-help stop-vbrjob

NAME
    Stop-VBRJob

SYNOPSIS


SYNTAX
    Stop-VBRJob -Job <CBackupJob[]> [-WarningAction <ActionPreference>] [-Warni
    ngVariable <String>] [-WhatIf] [-Confirm] [<CommonParameters>]


DESCRIPTION


RELATED LINKS

REMARKS
    To see the examples, type: "get-help Stop-VBRJob -examples".
    For more information, type: "get-help Stop-VBRJob -detailed".
    For technical information, type: "get-help Stop-VBRJob -full".
```

PowerShell Remoting

Veeam PowerShell supports remote execution of cmdlets and scripts. That is, you can run cmdlets and scripts directly on the Veeam Backup & Replication console, or against remote computers. A remote session can be started on one remote computer, or a number of remote computers at a time.

To enable PowerShell remoting:

1. Install the WinRM Service.
2. Enable PowerShell remoting. Start Windows PowerShell 2.0 as an administrator and type in: *Enable-PSRemoting*.

To learn more about PowerShell remoting and its configuration, see http://www.computerperformance.co.uk/powershell/powershell_remote.htm.

Example of Use

This section provides several examples of operations performed with the help of PowerShell scripts.

Performing Full VM Restore

In this example, we will review the full VM restore process performed by means of Veeam PowerShell script. To let you get most out of this example, each command will be illustrated by the action from the Veeam Backup & Replication UI that provides the result similar to execution of the PowerShell script.

First, we get a list of all available backup jobs with the cmdlet:

```
Get-VBRBackup
```

The similar action is performed by the **Restore** wizard in Veeam Backup & Replication 5.0:

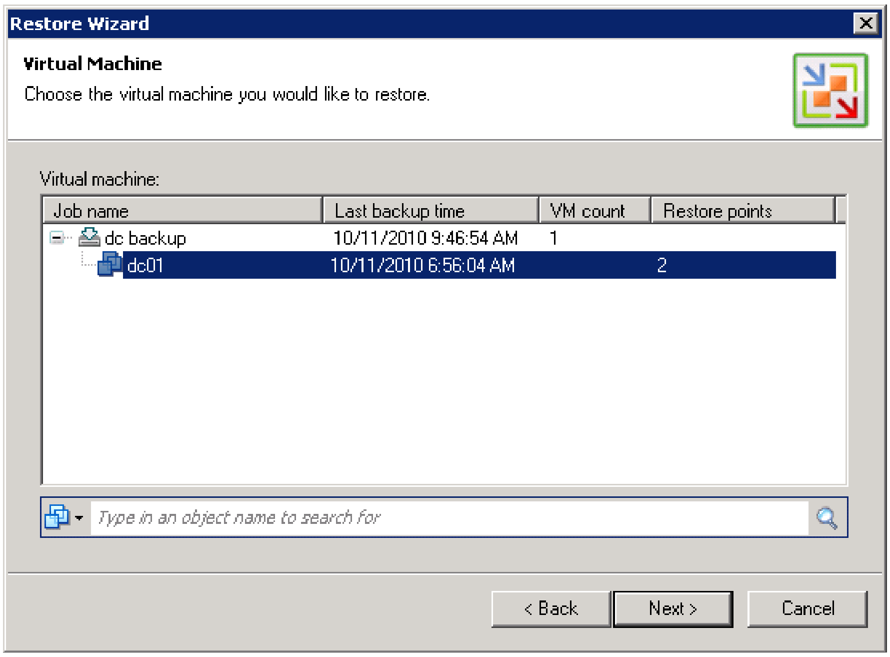

Restore Wizard

Virtual Machine

Choose the virtual machine you would like to restore.

Virtual machine:

Job name	Last backup time	VM count	Restore points
dc backup	10/11/2010 9:46:54 AM	1	
dc01	10/11/2010 6:56:04 AM		2

Then, we get a list of available restore points, select the last restore point and save it into a variable:

```
$rp = Get-VBRRestorePoint -Backup ((Get-VBRBackup)[2])
```

A similar action is performed by the Restore wizard in Veeam Backup & Replication 5.0:

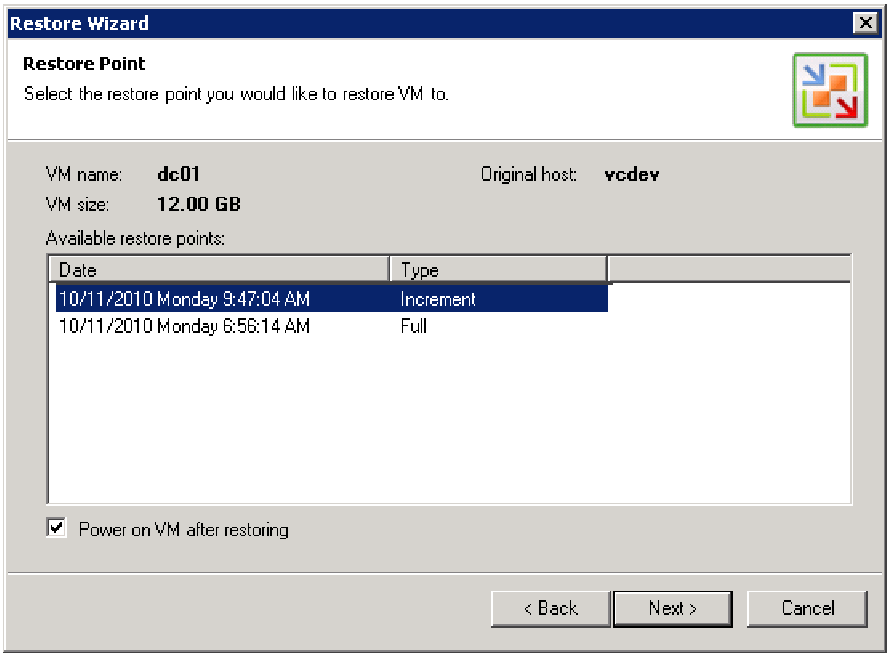

Restore Wizard

Restore Point

Select the restore point you would like to restore VM to.

VM name: **dc01**　Original host: **vcdev**

VM size: **12.00 GB**

Available restore points:

Date	Type
10/11/2010 Monday 9:47:04 AM	Increment
10/11/2010 Monday 6:56:14 AM	Full

☑ Power on VM after restoring

To restore a VM, you will need:

- ESX host to which the VM should be restored
- Resource pool
- Datastore on which the VM will reside

Let's get them and save into variables:

```
$server = Get-VBRServer | where {$_.Name -eq "esx12.veeam.lab"}
$server
```

Resource pool:

```
$res = Find-VBRResourcePool -Server ($server) | where {$_.Name -eq "Resources"}
$res
```

And datastore:

```
$datast = Find-VBRDatastore -Server ($server) | where {$_.Name -eq "localstore1"}
$datast
```

A similar action is performed by the **Restore** wizard in Veeam Backup & Replication 5.0:

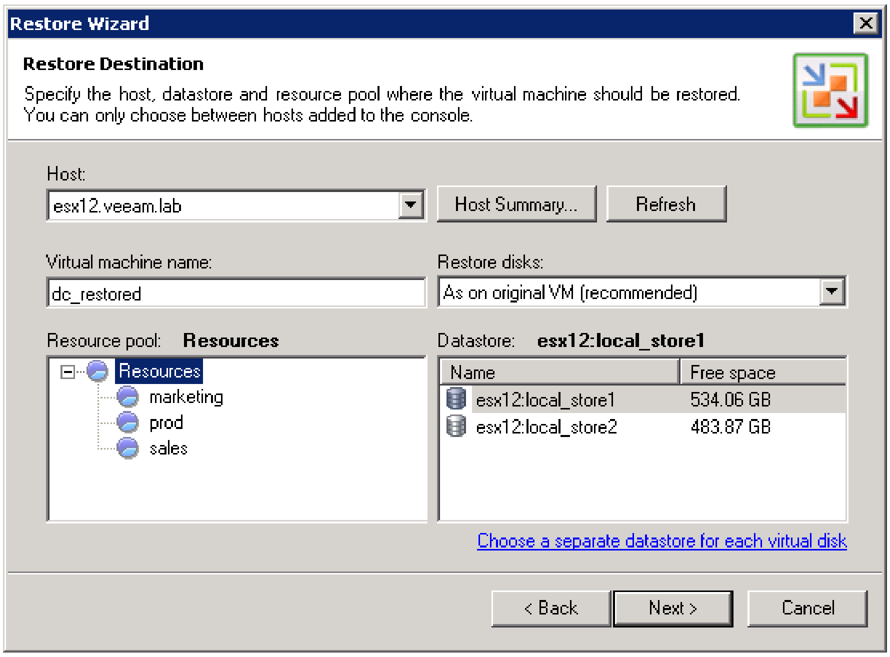

Restore Wizard

Restore Destination

Specify the host, datastore and resource pool where the virtual machine should be restored. You can only choose between hosts added to the console.

Host: esx12.veeam.lab

Virtual machine name: dc_restored

Restore disks: As on original VM (recommended)

Resource pool: **Resources**

Datastore: **esx12:local_store1**

Name	Free space
esx12:local_store1	534.06 GB
esx12:local_store2	483.87 GB

Now, let's perform restore of the VM:

```
Start-VBRRestoreVM -RestorePoint ($rp) -Server ($server) -ResourcePool ($res) -
Datastore ($datast) -VMName "dc01"
```

Parameters such as Virtual machine name (-*VMName*) and Restore reason (-*Reason*) are optional for the restore.

Performing VM Files Recovery Using Pipeline Input

In this example, we will review the process of restoring VM files (such as VMX, VMDK and so on) performed by means of Veeam PowerShell script.

The present example illustrates use of objects and pipeline input in PowerShell. Cmdlets operate with objects which they can accept and return. Every object has properties describing it, and methods that can be performed on it. Use of objects enables you to pipeline commands – that is, pass the output of one command as the input to another one. So, you can work with PoweShell interactively, entering commands and getting output step by step, or combine more complex commands using pipeline.

The example below describes step by step how such complex command is formed.

First, we will get an object of the backup named "*Backup Job 1*":

```
Get-VBRBackup | where {$_.JobName -eq "Backup Job 1"}
```

Then, we will get all restore points:

```
Get-VBRBackup | where {$_.JobName -eq "Backup Job 1"} | Get-VBRRestorePoint
```

Now we sort the restore points in the descending order and get the most recent one:

```
Get-VBRBackup | where {$_.JobName -eq "Backup Job 1"} | Get-VBRRestorePoint |
sort CreationTime -Descending | select -First 1
```

Once we got the most recent restore point, we can start the restore operation and save VM files locally to the *C:\restore\Documents* folder:

```
Get-VBRBackup | where {$_.JobName -eq "Backup Job 1"} | Get-VBRRestorePoint |
sort CreationTime -Descending | select -First 1 | Start-VBRRestoreVMFiles -Server
(Get-VBRLocalhost) -Path "C:\restore\Documents"
```

Changing Job Scheduling Settings

In this example, we will change the scheduling settings for an existing job — schedule it to be run every day.

First, we will get an object of the required job and save it to a variable:

```
$job = Get-VBRJob | where {$_.name -match "Backup Job 1"}
```

Now, we will get scheduling settings of the job and save them to a variable:

```
$sh = $job | Get-VBRJobScheduleOptions
```

At lastly, will change the scheduling settings:

```
$sh.OptionsDaily.Enabled = $true
$sh.OptionsDaily.Kind = "Everyday"
$sh.OptionsDaily.Time = "17:20:00"
Set-VBRJobScheduleOptions -Job ($job) -Options ($sh)
```

You can also change scheduling settings for a job in a different way:

```
$opt = $job | Get-VBRJobOptions
$opt.RunManually = $false
Set-VBRJobOptions -Job ($job) -Options ($opt)
```