

A Project Report on

Automated Failover Disaster Recovery of Virtual Machines in Geocluster Environment over Layer 3

University Of Alberta

Submitted by: Faisal Shaikh
8/8/2011

Automated Failover Disaster Recovery of Virtual Machines in Geocluster Environment over Layer 3

A PROJECT REPORT

Submitted by
Faisal Shaikh

*In Partial fulfillment for the award of the degree
Of*

MASTER OF SCIENCE

IN

INTERNETWORKING

University Of Alberta
Edmonton, Alberta T6G 2R3

August 2011

Abstract

Technology changes rapidly and become affordable so does the way of doing things for the growing business to protect data availability specially when the business involve services that general public need access to and companies generate revenue by providing such services.

Clustering of the servers is a method when minimum downtime and data protection is required. Geocustering or WAN clustering is the use of multiple redundant computing resources located in deferent geographical location to form what appears to be a single high available system. WAN clustering can be used for just about any computing resource, including mainframes, file servers and application stacks.

The biggest challenge in Geocustering is to make sure that system states and their related data are synchronized at different locations. Two developments have helped to meet this challenge, faster WAN connection and single virtualized master identity.

Traditional disaster recovery for physical machines are often costly, complex and keeping two sets of hardware in sync is nearly impossible and downtime of critical business application can cost lots of revenue to the business. Automated disaster recovery in Geocuster virtual environment will eliminate the need to perform numerous manual steps, which are prone to human error, dramatically reduce the downtime and improve reliability and availability of critical application.

Virtual Machine recovery on Geocuster VM host over L2 have lots of benefits over L3 but it comes with the cost, but what if we can achieve nearly similar benefits of L2 on L3. So my goal will be to design Geocuster virtual environment and find the way to do block level change transfer of data for minimum traffic over L3 and script the trigger recovery.

In this study I have created primary site in virtualized environment and replicate it to secondary site using VMware virtualized product and Veeam data replication product and trigger the disaster recovery process using self made script and monitor the downtime.

After the testing using different scenarios I have come to conclusion that downtime and data loss of highly critical machine in Geocuster environment is depend on replication cycle time between two sites, how much data is need to be transmit and system boot time which is hardware depended. In Lab environment I have started 50 MB of data transfer over 10Mb mimic WAN link and trigger the recovery process after 5 minutes, secondary site system came online in less than 3 minute without any data loss.

TABLE OF CONTENTS

CHAPTER 1	#
1.1 INTRODUCTION	1
1.2 OBJECTIVE	1
1.3 ISSUES	2
CHAPTER 2	#
2.1 WHAT IS VIRTULIZATION?	3
2.2 PHYSICAL vs VIRTUAL	3
2.3 BENEFITS OF VIRTULIZATION	4
2.4 BENEFITS OF VIRTUAL MACHINES	4
2.5 INTRODUCTION TO VMWARE	5
2.6 VMWARE ESX VS ESXI	7
2.7 CREATING NEW VIRTUAL MACHINE	10
CHAPTER 3	#
3.1 WHAT IS VEEAM BACKUP AND REPLICATION	11
3.2 VEEAM SERVER REQUIREMENT	11
3.3 REQUIREMENT PERMISSION	14
3.4 VEEAM USED PORTS	15
3.5 SYNCHRONOUS VS ASHYNCHRONOUSE REPLICATION	16
3.6 INSTALLING VEEAM BACKUP & REPLICATION	18
3.7 ADDING VIRTUAL CENTER	22
3.8 ADDING ESXI/ESX SERVER	24
3.9 REPLICATING VIRTUAL MACHINE	27

CHAPTER 4.....	#
4.1 NETWORK DESIGN.....	41
4.2 SITE CREATION.....	42
4.3 NETWORK CONFIGURATION.....	43
4.4 REPLICATION SERVICE.....	45
4.5 PERFORMING RELICATION FAILOVER USING VEEAM SERVICE	47
4.6 FAILING OVER REPLICA	48
4.7 UNDOING FAILOVER	50
4.8 MONITORING SERVICE	52
4.9 AUTOMATIC FAILOVER USING SCRIPT.....	52
CHAPTER 5.....	#
5.1 TESTING	53
5.2 CONCLUSION	55
REFERENCES	56

Chapter 1

1.1 Introduction:

In the past few years explosive growth in Internet based services and applications arise a question of reliability of such application and make us think what measures should we take to minimize the downtime and increase the high availability of online products while keeping the data intact.

Increasing use of Virtualization have bring down the hardware cost significantly for the organization while increasing the efficiency, utilization, and flexibility of the assets and has open new ways of doing rapid disaster recovery onsite. On the other hand remote site disaster recovery will add up the cost of direct connection such as dark fiber, L2TP etc.

WAN clustering also known as Geocustering is one option to create a remote disaster recovery site without the cost of direct link but the biggest challenge in this case is the data integrity and the system states over the internet although faster WAN connection and single virtualized master identity have helped to meet this challenge.

1.2 Objective:

Traditional disaster recovery process of physical machines involve similar sets of hardware and it's pretty hectic to keep them in sync and engage manual work which could easily prone to human error. Virtual Machine recovery on Virtual machine host in Geocluster environment over a LAN link have lots of benefits over WAN network such as high speed and no latency, but what if we can achieve virtually comparable result of LAN over WAN.

Goal here is to design Geocluster virtual environment and find a way to do block level data transfer of virtual machine for minimum traffic over WAN after the initial replication and then trigger the automated recovery process using a script to power on the remote Virtual machine after initiating different failover scenarios and monitor the amount of time of node left without service.

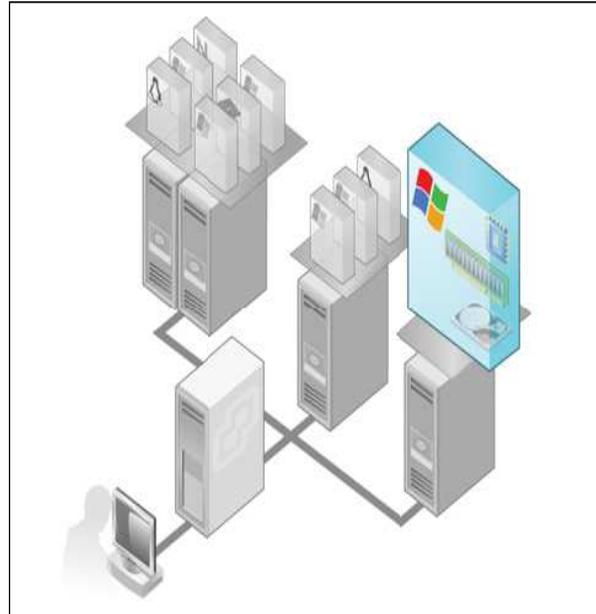
1.3 ISSUES:

Main issue with the Geocluster setup over WAN is the link latency and link speed between two sites, as rapid change in technology and lower cost of high speed data link from ISP's has somewhat help achieving this goal another big issue is the data replication, transferring full copy of virtual machine data for synchronization every time is not an option as it take numerous time to complete and failure of virtual machine will end up losing data. Using a block level data transfer technology enables us to track and copy only changes. Other issues are the monitoring and automated site recovery which I have achieved using simple scripts.

Chapter 2

2.1 What is Virtualization?

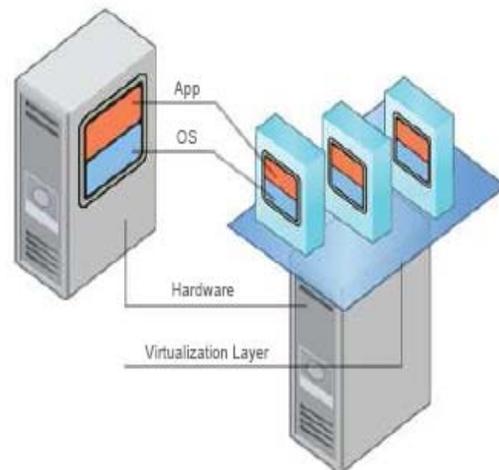
A virtual machine is a software computer that, like a physical computer, runs an operating system and applications. An operating system installed on a virtual machine is called a guest operating system. Every virtual machine has virtual devices that provide the same functionality as physical hardware. virtual machines get CPU and memory, access to storage, and network connectivity from the hosts they run on. In vSphere, virtual machines run on hosts or clusters. Multiple virtual machines can run on the same host or cluster at the same time.



2.2 Physical vs. Virtual

A virtual machine is a software computer that, like a physical computer, runs an operating system and applications. It has its own set of virtual hardware on which a guest operating system and its applications run. The operating system sees a consistent set of hardware regardless of the actual physical hardware components.

Virtual machines are not emulators or simulators. They are real machines that can do the same things physical computers can do and more. Because of the flexibility of virtual machines, physical computers become less a way to provide services (applications, databases, and so on) and more a way to house the virtual machines that provide those services.



2.3 Benefits of Virtualization:¹

1. **Get more out of your existing resources:** Pool common infrastructure resources and break the legacy “one application to one server” model with server consolidation.
2. **Reduce datacenter costs by reducing your physical infrastructure and improving your server to admin ratio:** Fewer servers and related IT hardware means reduced real estate and reduced power and cooling requirements. Better management tools let you improve your server to admin ratio so personnel requirements are reduced as well.
3. **Increase availability of hardware and applications for improved business continuity:** Securely backup and migrate entire virtual environments with no interruption in service. Eliminate planned downtime and recover immediately from unplanned issues.
4. **Gain operational flexibility:** Respond to market changes with dynamic resource management, faster server provisioning and improved desktop and application deployment.
5. **Improve desktop manageability and security:** Deploy, manage and monitor secure desktop environments that users can access locally or remotely, with or without a network connection, on almost any standard desktop, laptop or tablet PC.

2.4 Benefits of Virtual Machines

A physical machine is hard to move, difficult to copy, and bound to a specific set of hardware. Virtual machines are easy to copy and move because they are independent of physical hardware. Also, they are easy to manage because they are isolated from other virtual machines running on the same physical hardware and insulated from physical-hardware changes. Virtual machines enhance infrastructure by providing:

_ **Freedom from physical hardware constraints**

Virtual machines allow the operation of heterogeneous operating systems running across heterogeneous hardware.

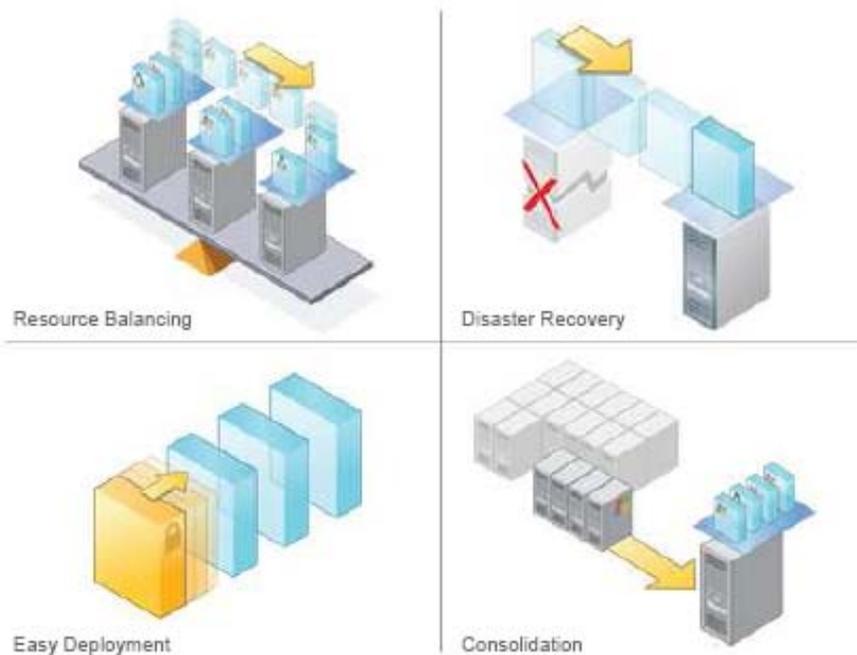
_ **Backup and recovery with little or no down-time**

You can configure virtual machines with operating systems and applications once and then clone them many times.

Backing up a virtual machine is as easy as backing up a few files. In this way, virtual machines ensure fast deployment and reliability.

_ **Greater resource utilization**

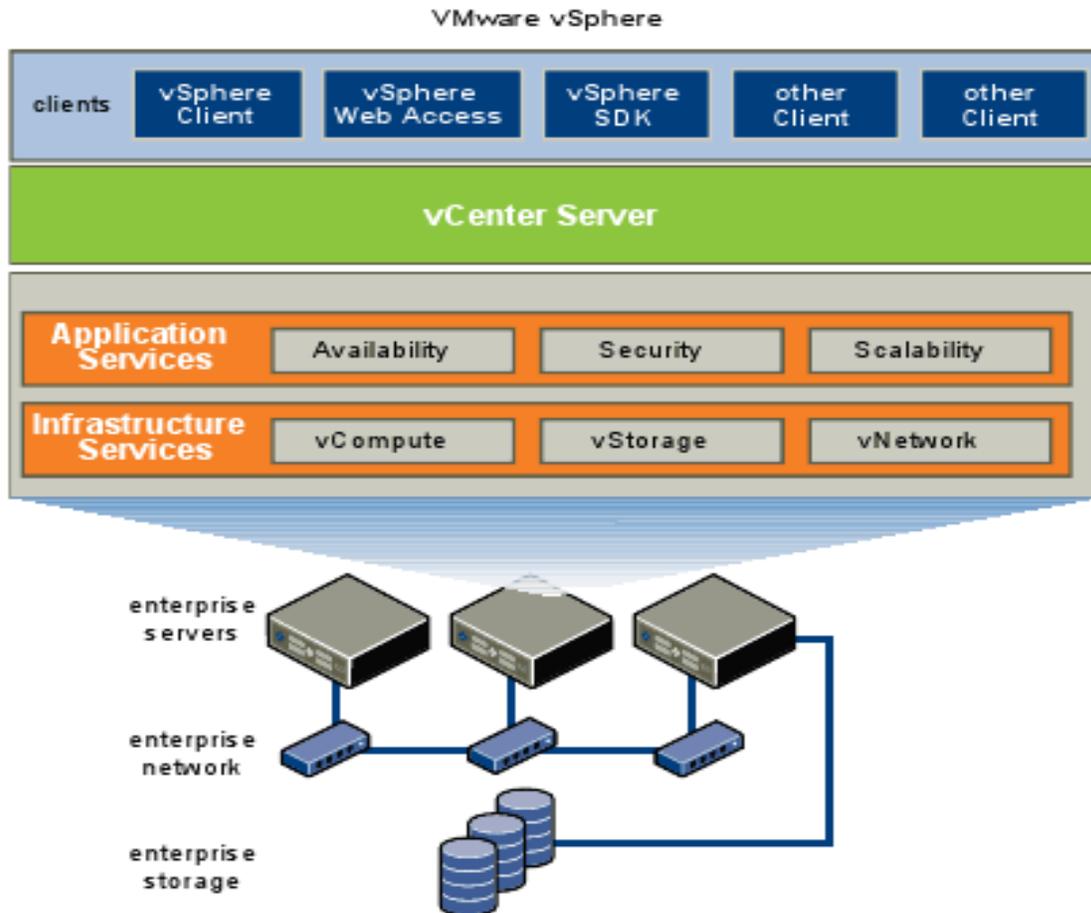
Multiple virtual machines can run on the same physical server. In addition, consolidating computing power to fewer physical computers can substantially increase power savings in your enterprise.



2.5 Introduction to VMWare vSphere:¹

In a simple term VMware is a software program that emulates PC system. VMware Sphere leverages the power of virtualization to transform datacenters into simplified cloud computing infrastructures and enables IT organizations to deliver flexible and reliable IT services. VMware vSphere virtualizes and aggregates the underlying physical hardware resources across multiple systems and provides pools of virtual resources to the datacenter.

As a cloud operating system, VMware vSphere manages large collections of infrastructure (such as CPUs, storage, and networking) as a seamless and dynamic operating environment, and also manages the complexity of a datacenter. The following component layers make up VMware vSphere.



- **Infrastructure Services** Infrastructure Services are the set of services provided to abstract, aggregate, and allocate hardware or infrastructure resources. Infrastructure Services can be categorized into:
 - ❖ VMware vCompute—the VMware capabilities that abstract away from underlying disparate server resources. vCompute services aggregate these resources across many discrete servers and assign them to applications.
 - ❖ VMware vStorage—the set of technologies that enables the most efficient use and management of storage in virtual environments.
 - ❖ VMware vNetwork—the set of technologies that simplify and enhance networking in virtual environments.

- **Application Services** Application Services are the set of services provided to ensure availability, security, and scalability for applications. Examples include HA and Fault Tolerance.

- **VMware vCenter Server** VMware vCenter Server provides a single point of control of the datacenter. It provides essential datacenter services such as access control, performance monitoring, and configuration.
- **Clients** Users can access the VMware vSphere datacenter through clients such as the vSphere Client or Web Access through a Web browser.

2.6 VMware ESX 4.1 vs ESXi 4.1 Comparison:¹

There are some limitations that you need to put into consideration before planning your disaster recovery site and chose the product according to your need, following is the comparison between licensed ESX vs free ESXi.

Capability	VMware ESX	VMware ESXi
Service Console	Service Console is a standard Linux environment through which a user has privileged access to the VMware ESX kernel. This Linux-based privileged access allows you to manage your environment by installing agents and drivers and executing scripts and other Linux-environment code.	VMware ESXi is designed to make the server a computing appliance. Accordingly, VMware ESXi behaves more like firmware than traditional software. VMware has created APIs through which monitoring and management tasks – traditionally done through Service Console agents – can be performed. VMware has provided remote scripting environments such as vCLI and PowerCLI to allow the remote execution of scripts and commands. Tech Support Mode (TSM) provides a command-line interface that can be used by the administrator to troubleshoot and correct abnormal conditions on VMware ESXi hosts.
CLI-Based Configuration	VMware ESX Service Console has a host CLI through which VMware ESX can be configured. VMware ESX can also be configured using vSphere CLI (vCLI) or vSphere PowerCLI.	The vSphere CLI (vCLI) is a remote scripting environment that interacts with VMware ESXi hosts to enable host configuration through scripts or specific commands. It replicates nearly all the equivalent COS commands for configuring ESX. VMware vSphere PowerCLI is a robust command-line tool for

		automating all aspect of vSphere management, including host, network, storage, virtual machine, guest operating system, and more.
Scriptable Installation	VMware ESX supports scriptable installations through utilities like KickStart.	VMware ESXi supports scriptable installations using a mechanism similar to Kickstart, and includes the ability to run pre- and post-installation scripts. VMware ESXi also provides support for post installation configuration using PowerCLI- and vCLI-based configuration scripts.
Boot from SAN	VMware ESX supports boot from SAN. Booting from SAN requires one dedicated LUN per server.	VMware ESXi may be booted from SAN. This is supported for Fibre Channel SAN, as well as iSCSI and FCoE for certain storage adapters that have been qualified for this capability. Please check the Hardware Compatibility List for supported storage adapters.
Serial Cable Connectivity	VMware ESX supports interaction through direct-attached serial cable to the VMware ESX host.	VMware ESXi does not support interaction through direct-attached serial cable to the VMware ESXi host at this time.
SNMP	VMware ESX supports SNMP.	VMware ESXi supports SNMP when licensed with vSphere Essentials, vSphere Essential Plus, vSphere Standard, vSphere Advanced, vSphere Enterprise, or vSphere Enterprise Plus. The free vSphere Hypervisor edition does not support SNMP.
Active Directory Integration	VMware ESX provides native support for Active Directory integration.	VMware ESXi provides native support for Active Directory integration.
HW Instrumentation	Service Console agents provide a range of HW instrumentation on VMware ESX.	VMware ESXi provides HW instrumentation through CIM Providers. Standards-based CIM Providers are distributed with all versions of VMware ESXi. VMware partners include their own proprietary CIM Providers in customized versions of VMware ESXi. These customized versions are available either from

		VMware’s web site or the partner’s web site, depending on the partner. Remote console applications like Dell DRAC, HP iLO, IBM RSA, and FSC iRMC S2 are supported with ESXi.
Software Patches and Updates	VMware ESX software patches and upgrades behave like traditional Linux based patches and upgrades. The installation of a software patch or upgrade may require multiple system boots as the patch or upgrade may have dependencies on previous patches or upgrades.	VMware ESXi patches and updates behave like firmware patches and updates. Any given patch or update is all-inclusive of previous patches and updates. That is, installing patch version “n” includes all updates included in patch versions n-1, n-2, and so forth. Furthermore, third party components such as OEM CIM providers can be updated independently of the base ESXi component, and vice versa.
vSphere Web Access	vSphere Web Access is only experimentally supported in VMware ESX.	VMware ESXi does not support web access at this time.
Licensing	For licensing information, see the VMwareSphere Editions Comparison.	For licensing information, see the VMwareSphere Editions Comparison.
Diagnostics and Troubleshooting	VMware ESX Service Console can be used to issue command that can help diagnose and repair support issues with the server.	VMware ESXi has several ways to enable support of the product: <ul style="list-style-type: none"> • Remote command sets such as the vCLI include diagnostic commands such as vmkfstools, resxtop, and vmware-cmd. • The console interface of VMware ESXi (known as the DCUI or Direct Console User Interface) has functionality to help repair the system, including restarting of all management agents. • Tech Support Mode, which allows low-level access to the system so that advanced diagnostic commands can be issues.
Jumbo Frames	VMware ESX 4.1 fully supports Jumbo Frames.	VMware ESXi 4.1 fully supports Jumbo Frames.

2.7 Creating a New Virtual Machine

Creating a new virtual machine allows you to customize options such as number of processors, memory, network connections, and storage. You can create new virtual machines on hosts or clusters. Before you create a virtual machine, decide which host or cluster the new virtual machine should reside on, the type of guest operating system you will install on the new virtual machine, and the location of the CD or image files for the installation. You also need appropriate vCenter Server permissions to create virtual machines.

To create a new virtual machine, select the host or cluster in the inventory you want to run the new virtual machine on and choose

Inventory > Host > New Virtual Machine or
Inventory > Cluster > New Virtual Machine.

A new virtual machine is like a physical computer with a blank hard disk. After you create the new virtual machine, you need to install a guest operating system on it. You can also change the settings of the virtual machine at any time.

Installing a guest operating system on a virtual machine is essentially the same as installing it on a physical computer. You must have a CD-ROM or ISO image containing the installation files from an operating system vendor. As with physical computers, a separate operating system license is required for each installation. After you install a guest operating system, the vCenter Server allows you to use the virtual machine, just as you would a physical computer, for tasks such as installing applications or managing power operations.



You can install the guest operating system from several different locations:



Chapter 3

3.1 What is VEEAM Backup and Replication?

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.

3.2 Veeam Server 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

VMware Infrastructure	
Platforms	VMware vSphere 4. VMware Infrastructure 3 (VI3).
Hosts	ESX(i) 4.x ESX(i) 3.x Free ESXi is not supported
Software	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.

	<p>Application-aware image-level processing option is supported on Windows XP x86, Windows 2003, Windows Vista, Windows 2008, Windows 2008 R2 and Windows 7.</p> <p>Windows file-level restore option is supported on NTFS, FAT and FAT32 file systems. GPT disks are not supported.</p> <p>To restore files from non-Windows guests (Linux, Solaris, BSD) use the Multi-OS File Level Restore wizard.</p>
Software	<p>VMware Tools (optional, recommended).</p> <p>Application-aware image-level processing option requires that your guest has VMware Tools and all latest service packs and patches.</p>
Veeam Backup & Replication Console	
Hardware	<p><i>CPU:</i> modern x86/x64 processor (minimum 4 cores recommended for optimal backup performance). Using faster processors generally improves backup performance.</p> <p><i>Memory:</i> 1024MB RAM (2048MB RAM when using local SQL Express installation). Using faster memory (DDR3) generally improves backup performance.</p> <p><i>Hard disk space:</i> 100 MB.</p> <p><i>Network:</i> 1Gbit/sec recommended due to backup performance considerations.</p>
OS	<p>Both 32-bit and 64-bit versions of the following operating systems are supported:</p> <ul style="list-style-type: none"> • Microsoft Windows XP SP3. • Microsoft Windows 2003 SP2. • Microsoft Windows Vista SP2. • Microsoft Windows 2008 SP2. • Microsoft Windows 2008 R2. • Microsoft Windows 7.
Software	<p>Microsoft .NET Framework 2.0 SP1 (included in the setup)</p> <p>Microsoft PowerShell 2.0 or later</p>
Backup Target	
Hardware	<p><i>CPU:</i> modern x86/x64 processor. Using faster processors generally improves backup performance when using Linux targets and Best</p>

	<p>compression option.</p> <p><i>Memory:</i> 256 MB RAM.</p> <p><i>Hard disk:</i> Using faster storage (fast high-RPM hard drives, RAID0 configurations) and optimal storage controller settings generally improves backup performance.</p> <p><i>Hard disk space:</i> Sufficient disk space required to store backup files.</p> <p><i>Network:</i> 1Gb/sec recommended due to backup performance considerations.</p>
OS	<p>Microsoft Windows. All major Linux distributions. ESX 3.x or later (ESXi is not supported).</p>
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	<p><i>CPU:</i> x86/x64 processor</p> <p><i>Memory:</i> 1024MB RAM (2048MB RAM when using local SQL Express installation).</p> <p><i>Hard disk space:</i> 25MB.</p> <p><i>Network:</i> 1Gbit/sec recommended due to backup performance considerations.</p>
OS	<p>Both 32-bit and 64-bit versions of the following operating systems are supported:</p> <ul style="list-style-type: none"> • Microsoft Windows XP SP3. • Microsoft Windows 2003 SP2. • Microsoft Windows Vista SP2. • Microsoft Windows 2008 SP2. • Microsoft Windows 2008 R2. • Microsoft Windows 7.

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	<p>Microsoft .NET Framework 2.0 SP1 or later.</p> <p>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.</p> <p><i>Browser:</i> Internet Explorer 6.0 or later, Mozilla Firefox 3.0 or later.</p> <p>Microsoft Excel 2003 or later (to view report data exported from Veeam Backup Enterprise Manager).</p>
Veeam Backup Search Server	
Hardware	Refer to corresponding Microsoft Search Server version system requirements.
OS	<p>Both 32-bit and 64-bit versions of the following operating systems:</p> <ul style="list-style-type: none"> • Microsoft Windows Server 2003. • Microsoft Windows Server 2008. • Microsoft Windows Server 2008 R2. <p>All the latest service packs and security updates should be installed.</p>
Software	<p>Microsoft Search Server 2008 (including Express edition)</p> <p>Microsoft Search Server 2010 (including Express edition)</p>

3.3 Required Permissions

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

Account	Required Permissions
Setup Account	Local Administrator permissions on the Veeam Backup & Replication console to install Veeam Backup & Replication 5.0.
Target/Source Permissions	<p>Host</p> <p>Root permissions on the source ESX/ESXi server.</p> <p>Root (or equivalent) permissions on the target Linux host.</p>

	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 <i>VeeamBackup</i> 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.

3.4 Veeam Used Ports

- Veeam Backup & Replication 5.0 uses SSH (TCP port 22) and/or HTTPS (TCP port 443) protocols as control channels (from the console to the source and/or target) and the range of ports from 2500 to 5000 as transmission channels (between the source and the target). For one job (copy/backup/replication), 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.
- If you are using vStorage APIs or VCB backup mode, make sure that port 902 is open on the ESX host to establish NFC connection.
- In case of using VSS, Veeam Backup & Replication 5.0 uses NetBIOS (TCP port 139, UDP ports 137, 138), SMB (TCP port 445) and RPC (TCP port 135) protocol connections from the console to guest VM additionally.
- Veeam NFS service uses ports 135, 1058 and 2049.
- The Veeam Backup Catalog component installed on the Veeam Backup & Replication server and responsible for indexing guest OS files inside Windows-based VMs communicates with Veeam Backup Enterprise Manager through port 9393. Veeam Backup Enterprise Manager Service uses port 9394 to collect data from backup servers added to Veeam Backup Enterprise Manager.
- Veeam Backup Service uses port 9392 (for calls are coming from Enterprise Manager Service).
- Veeam Backup Enterprise Manager Web site uses HTTP (TCP port 9080) and HTTPS (TCP port 9443) protocols.
- Veeam Backup Search service running on a dedicated Microsoft Search Server uses TCP port 9395.

Veeam Backup and Replication uses the synthetic backup method. Synthetic backup presupposes that a full backup is performed only once. All subsequent backups are incremental: only data changed relative to the most recent version is backed up. In contrast to repeatedly performing full backups, this approach ensures a faster and less space-intensive backup.

When full backup is performed, the resulting backup .vbk file is written to the target host. At each incremental backup, Veeam Backup and Replication rebuilds the most recent state of a VM and uses historical data to calculate a reverse increment. Obtained changes are backed up and saved as a service .vrb file next to the full backup .vbk file. The most recent point-in-time version is always a full backup, which allows the user to perform restore in the shortest period of time.

3.5 Synchronous vs. Asynchronous Replication:²

Replication is the process of copying data from one host to another host in a block-level, incremental fashion. So as replication is typically done at either a file level or a volume level, as things change on that volume, the blocks that are changed on the source are then immediately replicated to the target.

Synchronous replication technology does not acknowledge the write from the primary application until the block has been replicated to the target site. Asynchronous replication then acknowledges the write and then replicates that block over time.

Each has advantages and disadvantages in different scenarios. Synchronous replication has the advantage of being continually up to date at the target site. You always know that the data at the target site is as current as the data at the source site. The challenge is that since it won't acknowledge the write until it knows that the block has been replicated, the length of time it takes to get that block to the target can change the performance of the front-end application. So typically, synchronization is only done within a data center or at a very short distance -- less than 50 miles, or even 20 miles away. There are some technologies that are allowing people to go out further than that, but they are newer technologies.

The advantage of asynchronous replication is that no matter what the bandwidth or latency is, it's not going to impact the performance of the primary application. The downside of asynchronous is that it can get out of synch with the primary application and can actually get so out of synch that it can never catch up. Some products have the ability to go into special modes to try and catch up, but if you don't have enough bandwidth or have too much latency, you can actually get

so far behind that you won't meet your recovery point objective (RPO), which is the whole point of replication.

So one is always up to date, but can impact your performance and the other never impacts your performance, but can become out of date pretty quickly.

How does asynchronous replication differ from "point-in-time replication"?

Technically, point-in-time replication is a subset of one of the ways to do asynchronous replication, in that since asynchronous just means that you're not forcing the write to be acknowledged, before you acknowledge the write back to the primary application. What point-in-time replication means is that you take a snapshot at a certain time, typically once an hour then your replication product looks at the bytes that have been changed in between the last snapshot and the current snapshot and then replicates those bytes necessary to create those points in time at the replication destination so some of them can continually replicate and then will take a snapshot at the source site. Then it will just replicate that status to the other side. But the big difference is that with a point-in-time replication system, you have one or many points-in-time to go back to if one you have with corruption.

With asynchronous replication, depending on how to up to date you are, you're continually copying over everything, including the corruption. If you were to do something like drop a table, you could potentially overwrite the target with that corruption.

3.6 Installing Veeam Backup & Replication 5.0

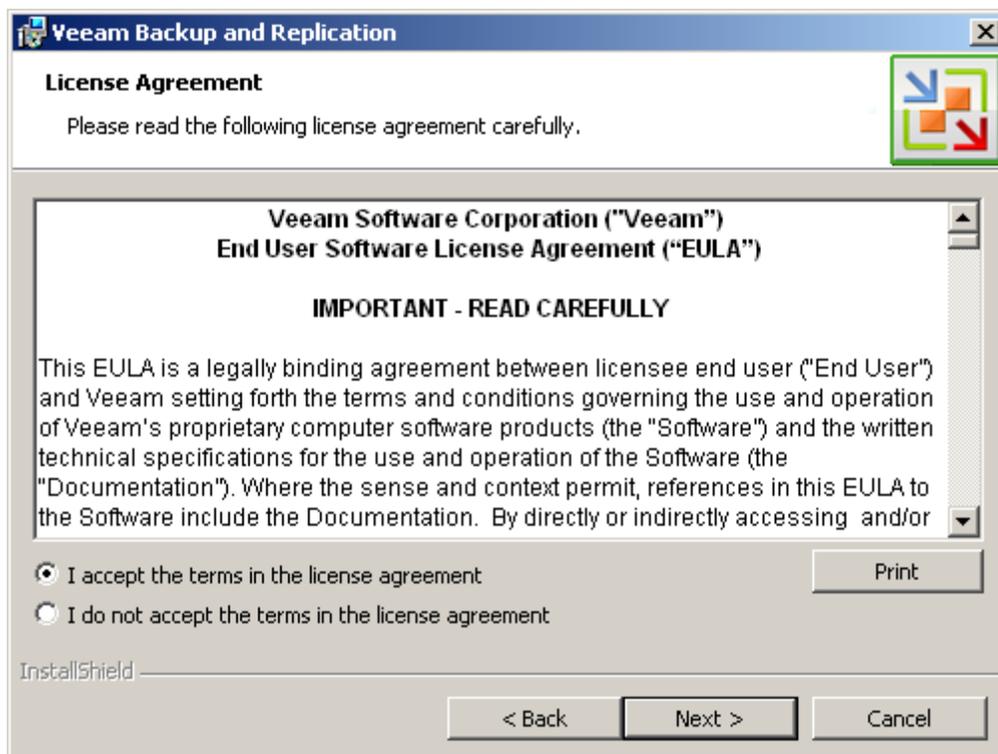
This section will guide you through the Veeam Backup & Replication 5.0 installation process.

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 *VeeamBackup.exe* setup file.

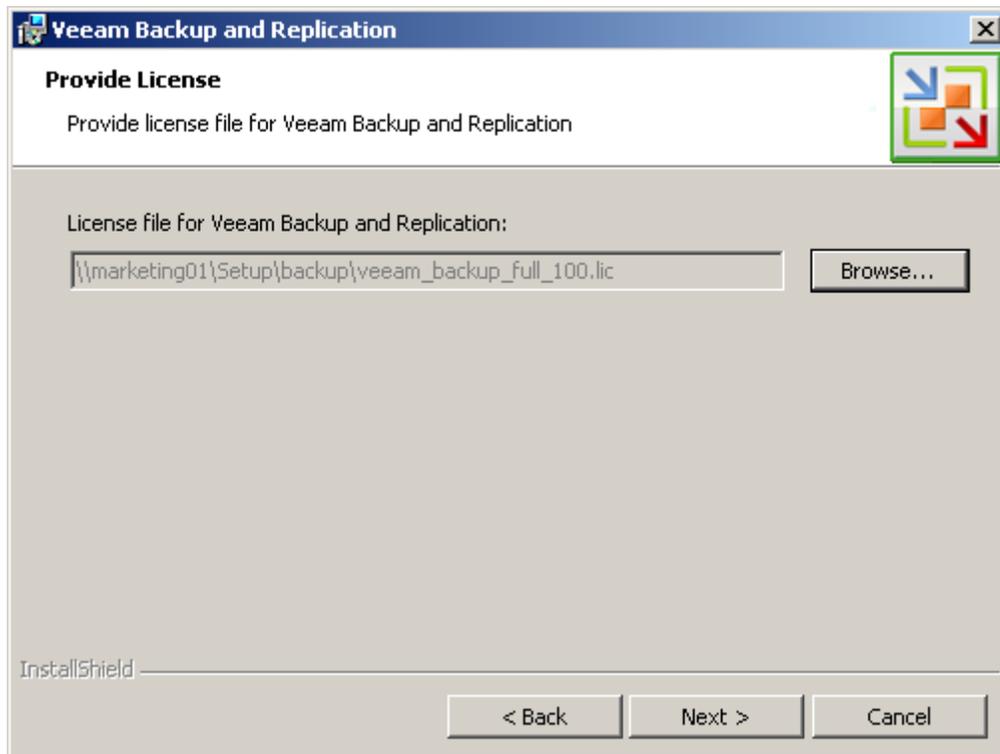
Step 2. Accept the License Agreement

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.



Step 3. Install 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.



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.

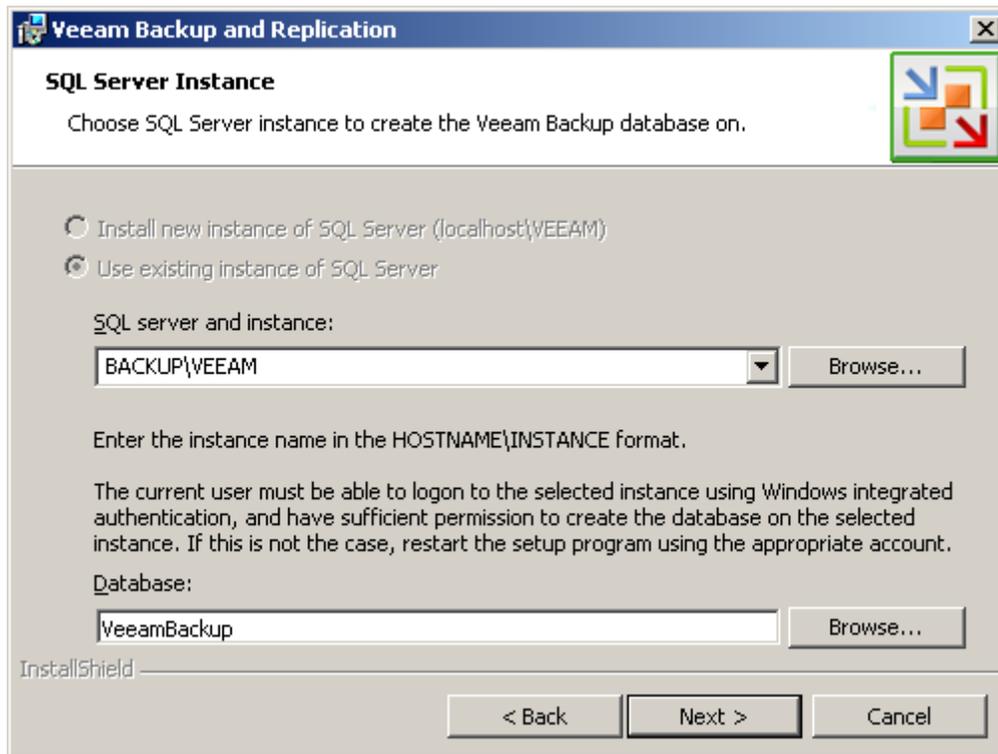
Use the **Space** button to estimate how much free space is available on your local drives.

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.



NOTE:

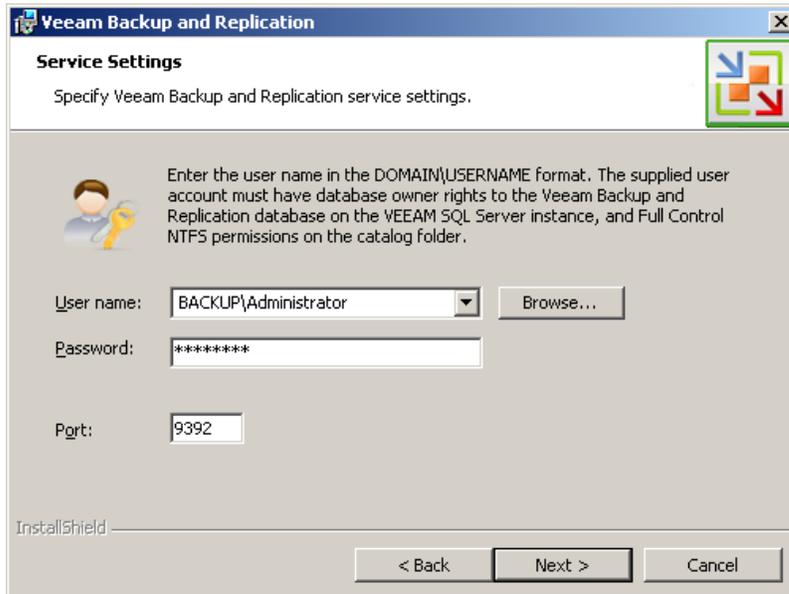
In case the *VeeamBackup* database already exists on the SQL Server instance (that is, it was created by the previous installations of Veeam Backup & Replication), a warning message notifying about it will be displayed. Click the **Use Existing** button to connect to the detected database. If necessary,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

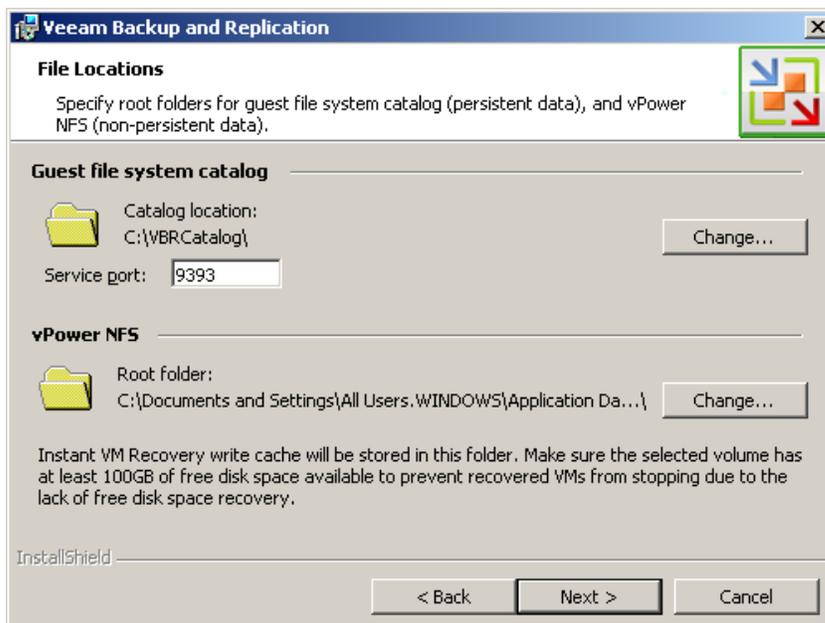
If necessary, change the number of TCP port. By default, Veeam Backup & Replication services use port 9392.



Step 7. Specify Catalog and vPower NFS 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.

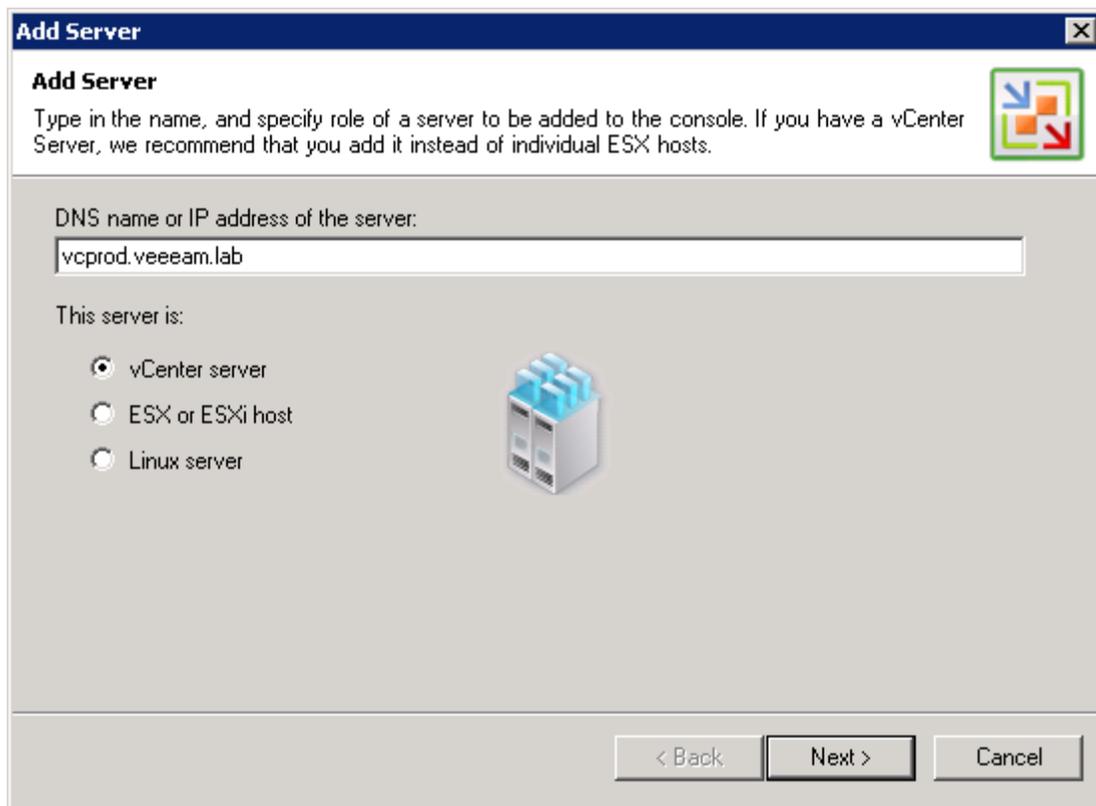


3.7 Adding VirtualCenter

To add a VirtualCenter server, follow the next steps.

Step 1. Specify Server Type and Name

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

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

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 vcpod.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.

< Back Next > Cancel

Step 3. Finish Working with the Wizard

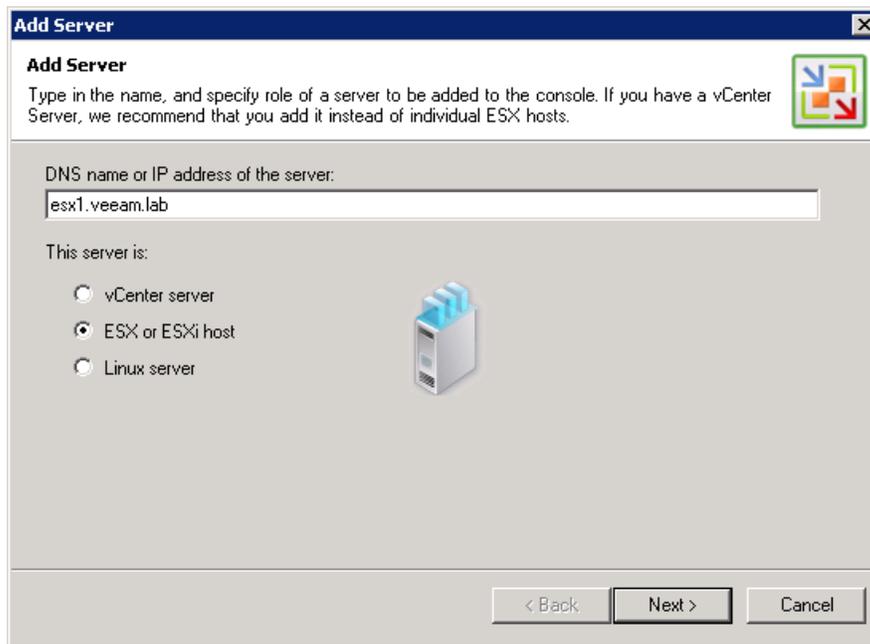
If you want to connect to the added VirtualCenter server on finishing work with the wizard, select the **Connect when I click Finish** check box. Then click **Finish**. If you do not select this option, you will have to manually connect to the added server.

3.8 Adding ESX/ESXi Server

To add an ESX/ESXi server, follow the next steps.

Step 1. Specify Server Type and Name

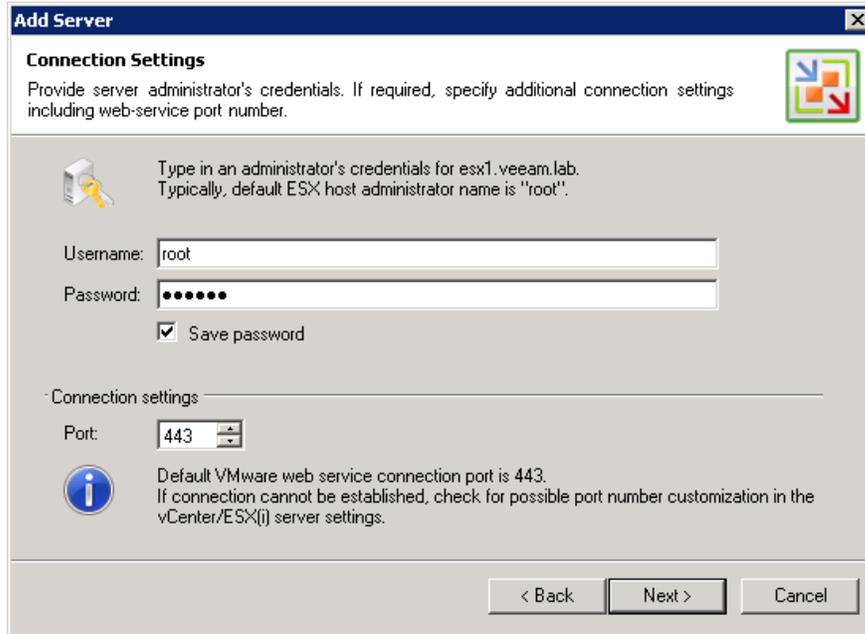
Enter a full DNS name or IP address of the server and select the server type: **ESX or ESXi host**.



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 a new, you will have to enter credentials again as soon as the server is addressed.

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



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:

Password:

Save password

Connection settings

Port:

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.

< Back Next > Cancel

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.

The screenshot shows a dialog box titled "Add Server" with a close button (X) in the top right corner. The main heading is "Specify SSH Connection Settings". Below the heading is a descriptive text: "Provide credentials for service console connection, and adjust secure shell (SSH) port number if needed. Service console connection to ESX server is optional." To the right of this text is a small icon with four arrows pointing outwards from a central square. The dialog contains several input fields and checkboxes:

- A checked checkbox labeled "Use service console connection to this server".
- A "Username:" label followed by a text input field containing the text "root".
- A "Password:" label followed by a text input field containing the text "[To change the saved password, click here]".
- A checked checkbox labeled "Save password".
- A section header "Non-root account options" followed by a horizontal line.
- An unchecked checkbox labeled "Elevate account to root".
- An unchecked checkbox labeled "Add account to the sudoers file automatically".
- A "Root password:" label followed by an empty text input field.
- An "Advanced..." button located at the bottom right of the main content area.
- At the bottom of the dialog are three buttons: "< Back", "Next >", and "Cancel".

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

The screenshot shows a dialog box titled "Advanced Connection Options" with a close button (X) in the top right corner. The main heading is "Service console connection" followed by a horizontal line. Below the heading are two input fields:

- "SSH port:" followed by a spin box containing the number "22".
- "SSH timeout:" followed by a spin box containing the number "20000" and the text "ms".

At the bottom of the dialog are two buttons: "OK" and "Cancel".

3.9 Replicating Virtual Machines

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.

NOTE:

You can use custom pre-freeze and post-thaw scripts before creating a snapshot of a virtual machine that is being 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.

Before You Begin Replication

- Prior to creating a VM replica, make sure you have enough free space on the destination disk. When the 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 Replication 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 the 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 Replication Job

Name and Description
Type in a name and description for this replication job.

Name:
SQL replication job

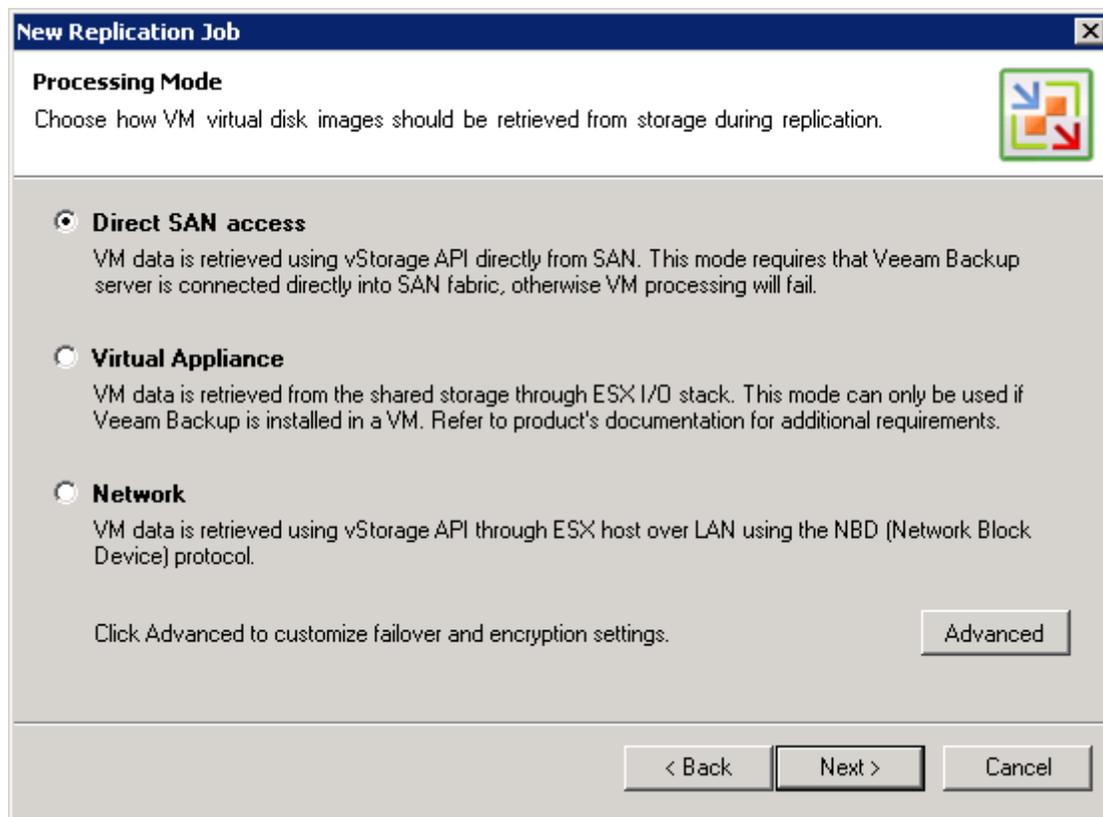
Description:
Created by VEEAM\Administrator at 10/11/2010 7:11:12 AM.

< Back Next > Cancel

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.

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.

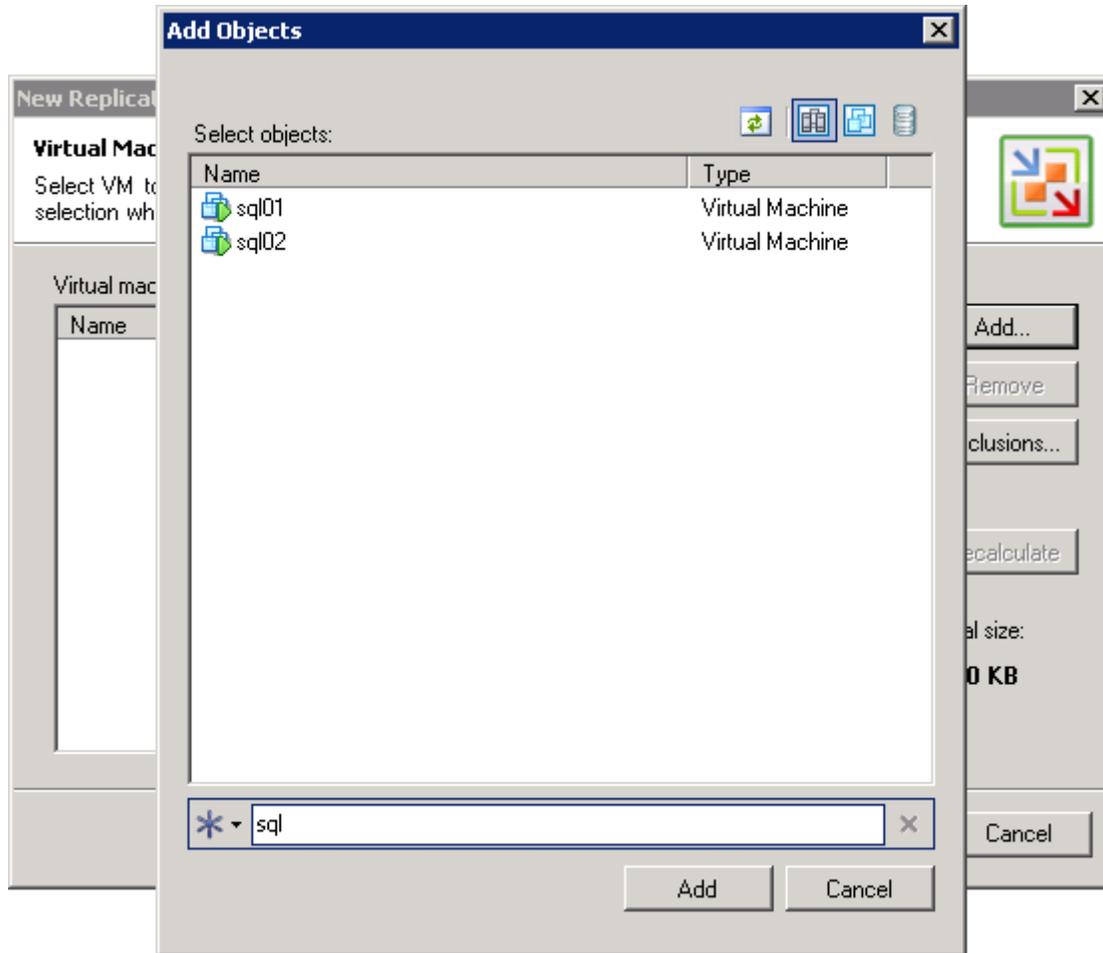
Click the **Add...** button to browse to VMs and VM containers that should be replicated. In the displayed VI tree, select a necessary object and click the **Add** button.

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.
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.

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** Choose... Check Space
Datastore: **esx12:local_store1**

Initial replication
 Perform initial replication over this removable storage:
F:\replica Browse...

Replica settings
Replica name's suffix: Restore points to keep on disk:
Replica disks: Advanced...

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

< Back Next > Cancel

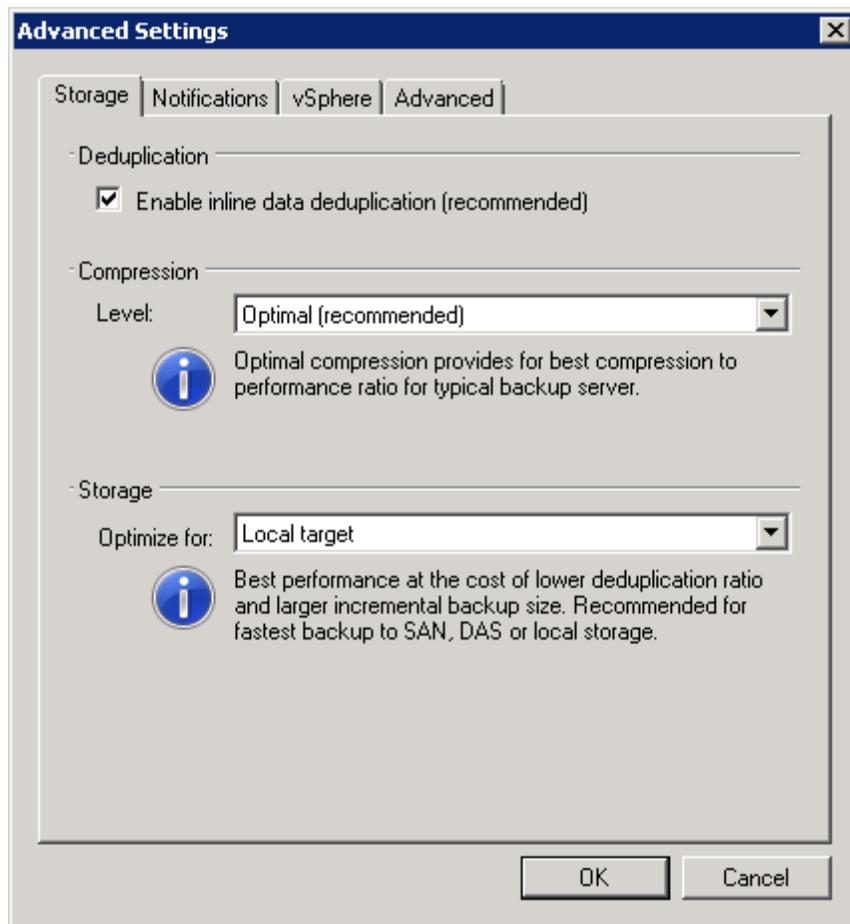
In the **Replica's name suffix** field, enter a suffix that will be appended to a name of the virtual machine you are replicating. This name, with the suffix added, will be used to register the replicated virtual machine on the target server. Files of a replicated VM will be placed to the selected datastore 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:



Storage settings

You can disable de-duplication by clearing the **Enable inline data 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*.

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.

Notification settings

- Select the **Send email 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).
- 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.

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. To learn more about VMware tools quiescence, see Veeam Backup & Replication 5.0 User Guide at www.veeam.com.
- If you are running pre-ESX 3.5 Update 2 hosts, consider enabling the **Safe removal for snapshots larger than ... Mb** 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

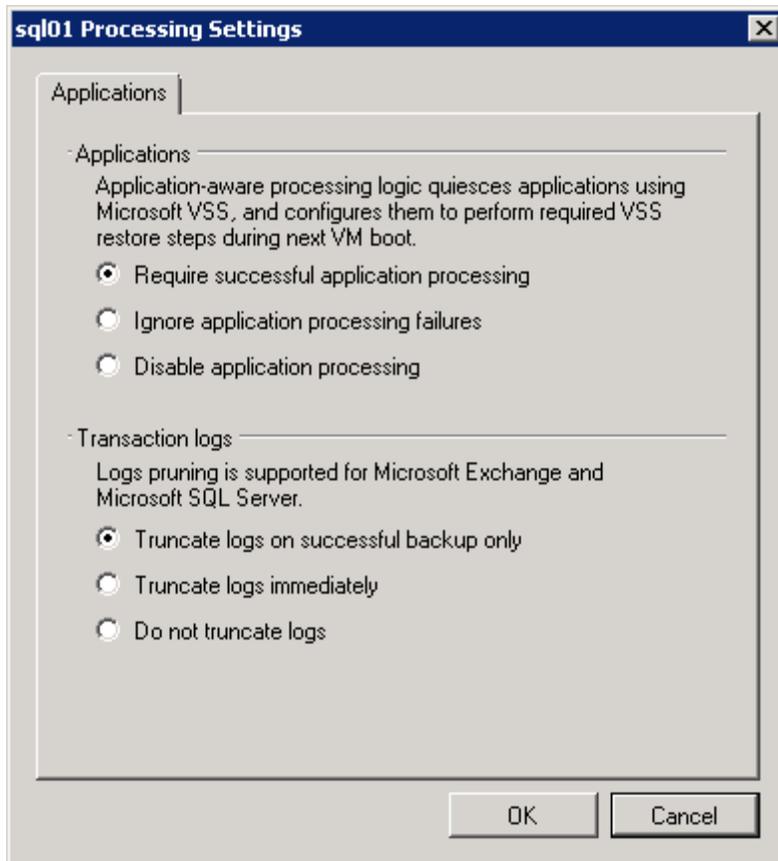
If you want to create a transactionally consistent replica ensuring successful recovery of VM applications without any data loss, select the **Enable application-aware image processing (recommended)** check box.

To coordinate proper VSS activities, Veeam Backup & Replication installs a small agent inside a VM. The agent is installed only during VSS quiescence and indexing 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.

- 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 quiescing options for a VM, select it in the list and click the **Edit...** button.



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 quiescing 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 the 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 [X]

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 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

< Back Create Cancel

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.

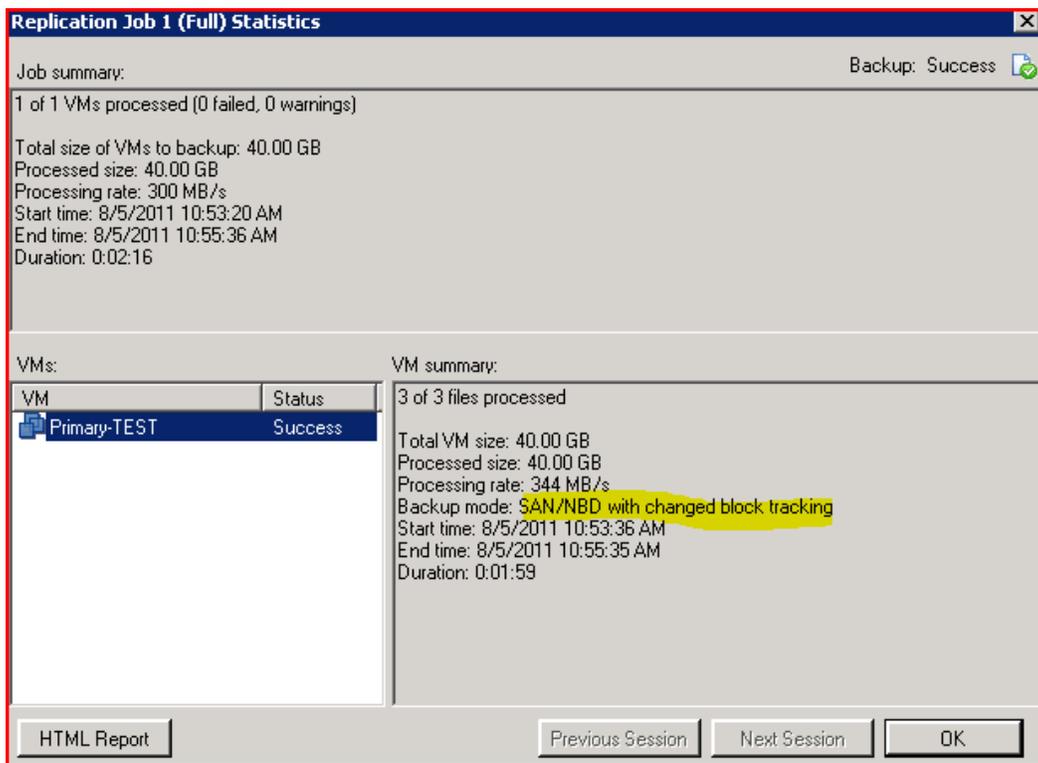
NOTE:

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.

Note: ESXi Server required licensing to enable SAN replication. In the LAB I have used ESXi Full featured Trail Version.

Once the initial replication done which approximately took 8 Hours over 10 Mbps WAN link, then every incremental with changed block tracking will take 3 to 6 minutes for every 20 to 50 MB of data changed.

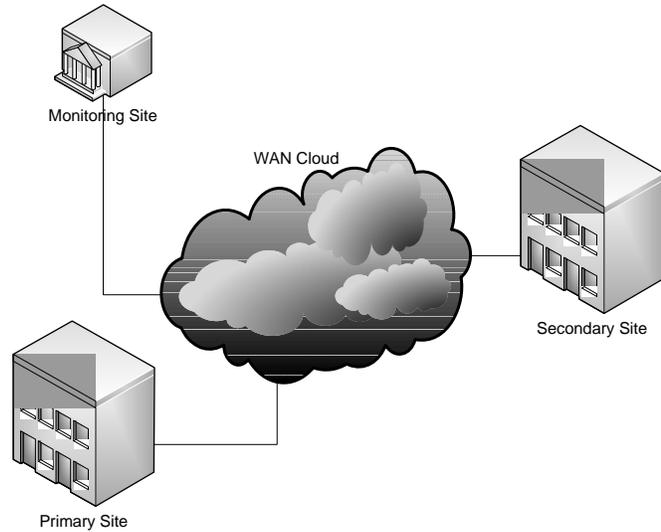
You can setup automatic replication occurrence depending on your need. In my Lab test I have setup replication every 10 min and add, and remove 20 to 50 MB of data in separate instances and monitor the time and data loss and found average time it took to transfer 50 MB of data is 6 Minutes.



Chapter 4:

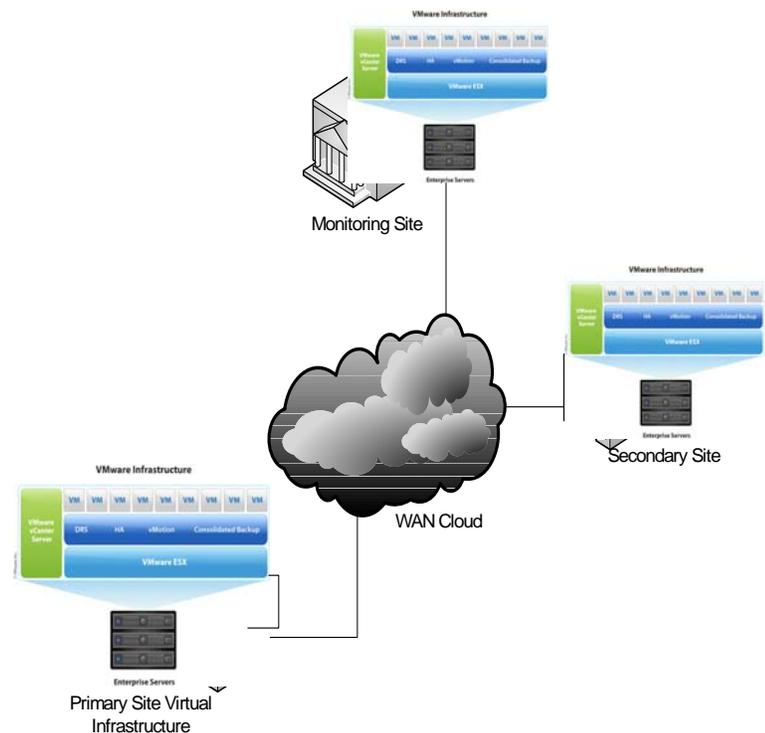
4.1 Lab Testing Network Design

In the LAB Virtual environment I have setup One primary site which will run One Primary virtual machine, One secondary site that holds the replication data and become active in failover scenario and I have setup the Monitoring site which will perform replication between primary site to secondary site and keep monitoring the primary site heartbeat and trigger the automated disaster recovery process in the failover scenario.



In this LAB environment I have created all sites on a single ESXi host with different subnet and mimic the WAN environment on the switches with the limitation of 10Mbps/sec transfer rate between subnets. In a real world scenario it's a best practice that you build your Primary site on VMware ESX in a cluster of at least 3 physical hosts for high availability.

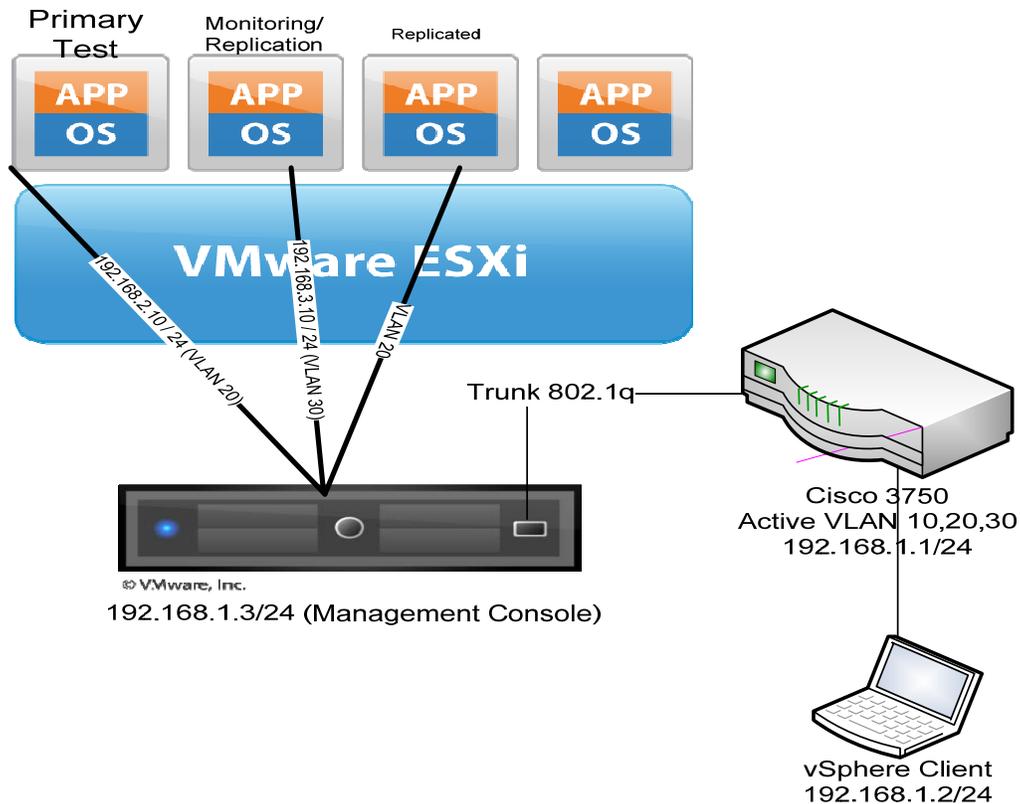
It is also important to create a monitoring site in a separate location other than the Primary and secondary because in case of node failure of either site, separate monitoring site will have no effect and it keep monitoring site to take decision and act accordingly.



4.2 Site Creation:

In a LAB test environment I have installed ESXi 4.1 on a single processor server with the single network card and 4GB of RAM. Network card is attached to a Cisco 3750 Switch with a trunk dot1q encapsulation and assigned 192.168.1.3/24 IP to Management console with the VLAN tag of 10.

In a LAB test environment I have created one Virtual machine (Windows XP) called it “Primary-TEST” on a subnet of 192.168.2.0/24 (VLAN 20), another Virtual machine (Windows XP) for monitoring and replication on 192.168.3.0/24 (VLAN 30) subnet as shown in the figure. In a real world situation create the same virtual network with same VLAN tag and IP scheme on a remote recovery site as the replicated VM will carry the same network setting and require the same subnet IP’s.



4.3 Network Configuration:

Network configuration on the router is simple, I have created 3 VLAN and enable inter vlan routing with the rate limit of 10Mbits/s to mimic the real world WAN link with the 10Mbits/s up and down stream transfer rate with the minimum latency. Following is the copy of router configuration.

```
Current configuration : 10915 bytes
!
! Last configuration change at 10:32:50 MDT Fri Jul 5 2011 by faisals
! NVRAM config last updated at 10:32:35 MDT Fri Jul 5 2011 by faisals
!
version 12.2
no service pad
service timestamps debug uptime
service timestamps log datetime localtime show-timezone
no service password-encryption
service sequence-numbers
!
hostname SWMAIN
!
boot-start-marker
boot-end-marker
!
logging console informational
clock timezone MST -7
clock summer-time MDT recurring
system mtu routing 1500
ip subnet-zero
ip routing
ip multicast-routing distributed
!
port-channel load-balance src-dst-mac
!
spanning-tree mode rapid-pvst
spanning-tree etherchannel guard misconfig
spanning-tree extend system-id
!
vlan internal allocation policy ascending
!
```

```
interface GigabitEthernet0/0/1
Description "Connected to VMHost"
switchport trunk encapsulation dot1q
switchport mode trunk
spanning-tree portfast trunk
!
interface GigabitEthernet0/0/48
description "Connected to vSphere Client"
switchport access vlan 10
switchport mode access
spanning-tree portfast
!
interface GigabitEthernet0/0/49
!
interface GigabitEthernet0/0/50
!
interface GigabitEthernet0/0/51

interface Vlan1
no ip address
!
interface Vlan10
ip address 192.168.1.1 255.255.255.0
rate-limit input 10000000 10000000 10000000 conform-action transmit exceed-action drop
no ip redirects

interface Vlan20
ip address 192.168.2.1 255.255.255.0
rate-limit input 10000000 10000000 10000000 conform-action transmit exceed-action drop
no ip redirects

interface Vlan30
ip address 192.168.3.1 255.255.255.0
rate-limit input 10000000 10000000 10000000 conform-action transmit exceed-action drop
no ip redirects
!
vlan internal allocation policy ascending
!
ip classless
end
```

4.4 Replication Service Installation:³

Mirroring or replicating data essentially involves writing two copies of the data, one to each of a mirrored half. There are a few ways to implement mirroring. Many OS's have the ability to create disk mirrors, or "Software RAID-1". There are also third party software products that can implement asynch mirroring on behalf of the host. In general, it is preferable to have an independent, shared storage server implement the data mirroring, as it can consolidate the activity for many hosts and offload those hosts from the performance penalty of the extra writes.

As for Synchronous vs Asynchronous, as the terms would indicate, it's all about time. In particular, it's about the time that data written to disk is "committed". In Synchronous Data Mirroring, the write acknowledgement isn't returned to the requesting server until the data has been written to both storage arrays. This is akin to "RAID-1": the two mirror halves stay in "lockstep" — you know that at any moment you have two identical copies of the data on two devices.

Because the acknowledgement isn't returned until both storage devices have a copy of the write, this implies that the bandwidth of the channels used for the writes needs to be sufficient so as not to cause undue latency which would affect the responsiveness of the production servers.

As this latency directly affects the performance, Synchronous Data Mirroring is often not a viable solution for DR where distance and limited bandwidth would cause roundtrip latency to go way up. In such cases, Async Replication is the method of choice.

While not a hard and fast rule, most people associate Synchronous Mirroring with Business Continuity, whereas Asynch Mirroring is typically associated with Disaster Recovery.

In Asynchronous Mirroring, we still have the concept of mirror halves, but for the purpose of clarity we will refer to them as a "source" volume and a "destination" or "replicated" volume. The source volume is the local production volume a host (such as your SQL server) writes to and reads from. Some agent then copies any of the writes over to the destination volume at the distant DR or remote site.

The agent is almost always software, implemented either on the host itself, on a SAN-based or DAS-based Storage Array, or on an intermediary engine, perhaps a Storage Virtualization platform.

If the agent is running on the host itself or on a file server (a.k.a. "NAS" storage such as a CIFS or NFS service), the replication can be implemented at the file level. When the agent is run on a Storage Array or Virtualization platform, the replication is invariably implemented at a raw disk or block level, where only the low-level writes or disk blocks that have been modified are copied to the destination volume. The scope of this article is limited to disk- or block-level replication.

With Asynchronous Replication it is understood that the bandwidth between the two sites is limited, and so writes to the local source volume are acknowledged immediately, and the agent subsequently replicates those changes to the destination volume. The destination volume is usually always in a "catch up" mode, anywhere from a few writes to a few gigabytes worth of writes behind the current state of the source. How far the destination falls out of synch (or lags behind) the source volume is determined by the quantity of changes occurring on the source volume, the bandwidth of the intersite link used for the replication, network QoS policies and/or replication schedules or throttles. There are many other factors affecting overall replication system performance, some of which we will discuss later. Finally, network outages that break the replication link can cause the two volumes to fall further out of synch.

Clearly, provisions must be made to keep track of changes to the source volume in order to replicate those changes and bring the source and destination as close to synch as time and bandwidth permit. There are a number of ways to do this, ranging from marking changed blocks that need to be replicated, to buffering the changes in a reserved storage area.

Some vendors offer de-duplicating technologies, which attempt to reduce the replication traffic by deleting buffered writes for blocks that have been written (and buffered) more than once; the idea is that they will only replicate the latest change to the block. There are pros and cons to the different technologies used; I personally prefer at least having the option to buffer every change to every block, as that facilitates snapshotting and newer technologies such as CDP (Continuous Data Protection) at the DR site.

There are lots of products available in the market which can perform replication of raw disk or block level, where only disk blocks that have been modified are copied to the destination location volume.

In the event a major incident occurs at the primary site, a DR site provides some level of recoverability that could be measured in minutes, hours, or days depending on how much planning and resources have gone into the DR implementation.

Remember I am try to build a DR site with automated recovery not a HA (High Availability) or BC (Business Continuity) Site, it is important at this point to clear up the differences.

DR compared to HA or BC

To begin with, let's delineate between a few of the common terms often employed when discussing matters of data protection and availability of service: H/A, BC and DR. The concepts are different, but they are not mutually exclusive.

High Availability or H/A refers to systems and components designed to withstand a variety of non-catastrophic local failures. The vendors implement H/A in servers and storage arrays via redundancy: redundant power, cooling, cabling, switches, RAID groups, dual processors, etc. The idea is that a fault of a component (such as an GBIC on an HBA) or a pulled cable shouldn't stop the show. An alternate path or component can take over without missing a beat. With H/A, users shouldn't notice any disruption in service when such a failure occurs.

Business Continuity / Continuance or BC takes this one step further. It is the idea of adding some additional level of redundancy to the architecture so that it can withstand the failure of entire systems without stopping production. Often when storage vendors talk about Business Continuity, they are implying the use of Synchronous Mirroring between two of their high-end storage arrays, perhaps separated over some short distance, such as between two buildings on a campus.

Whereas Business Continuity is usually implemented on a local campus or metro basis, DR is understood to imply geographical separation. Like BC, that separation can be "across the parking lot" or "across campus", but more typically it's "across the state" or "across the country".

As I said, H/A, BC and DR are not mutually exclusive. Many shops have requirements to be both Highly Available and have a DR plan in place; in some cases the two may be deployed using the same mechanisms, such as a stretch cluster across a campus or metropolitan area, where the two ends are both active production sites, synchronous mirrors of each other.

However, in the majority of cases, DR sites are not just BC extensions of the main site. Many shops with critical availability requirements will implement BC locally, and have a contingency DR site in another state or part of the country.

Bringing the DR site into production is a decision taken only when the primary site is deemed unavailable because of a disaster. How quickly that site is brought online depends on how much effort and resources have gone into the planning. Bringing the site online typically requires network changes (DNS, etc.), and an understanding that once you've flipped the switch, some planning will be required to flip it back once the primary site is restored to service.

4.5 Performing Replica Failover Using Veeam Service

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.

NOTE:

Always use the Veeam Backup & Replication 5.0 to work with replicated virtual machines. Starting or stopping a virtual machine replica outside Veeam Backup & Replication 5.0 (e.g., in Virtual Infrastructure client) will make the restore points not valid. If the machine hosting Veeam Backup & Replication 5.0 is not available, the virtual machine replica can be started on the target host using Virtual Infrastructure client. In this case, all restore points will become invalid.

4.6 Failing Over 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 the 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.

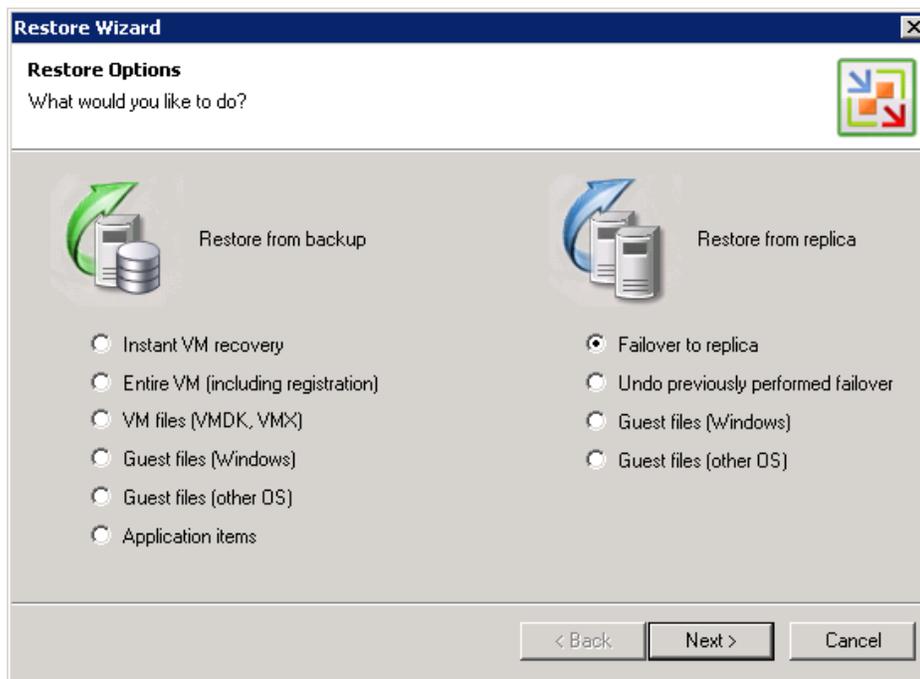
Step 1. Launch the Restore Wizard

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.

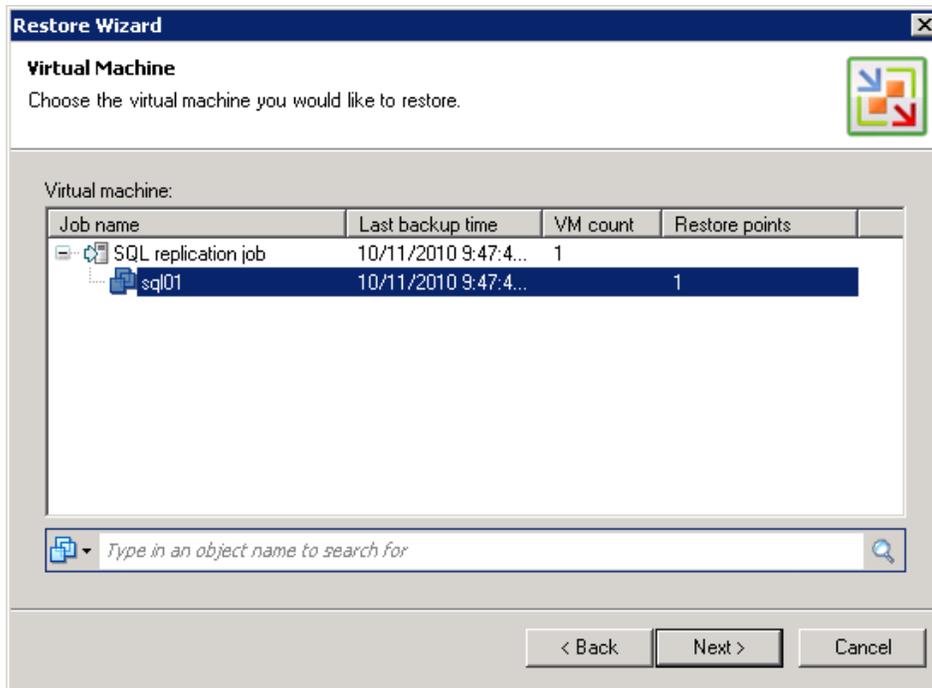
Step 2. Select a Task

Select the **Failover to replica** option.



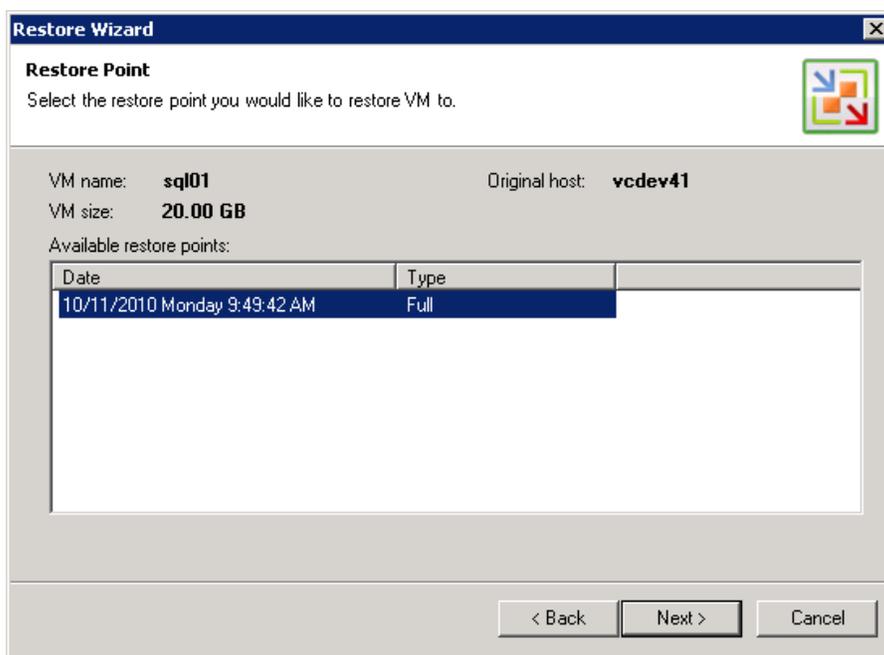
Step 3. Select a Virtual Machine

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



Step 4. Select the Restore Point

Select a necessary restore point for the virtual machine.



Step 5. Finish Working with the Wizard

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

3.7 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 the offered options.

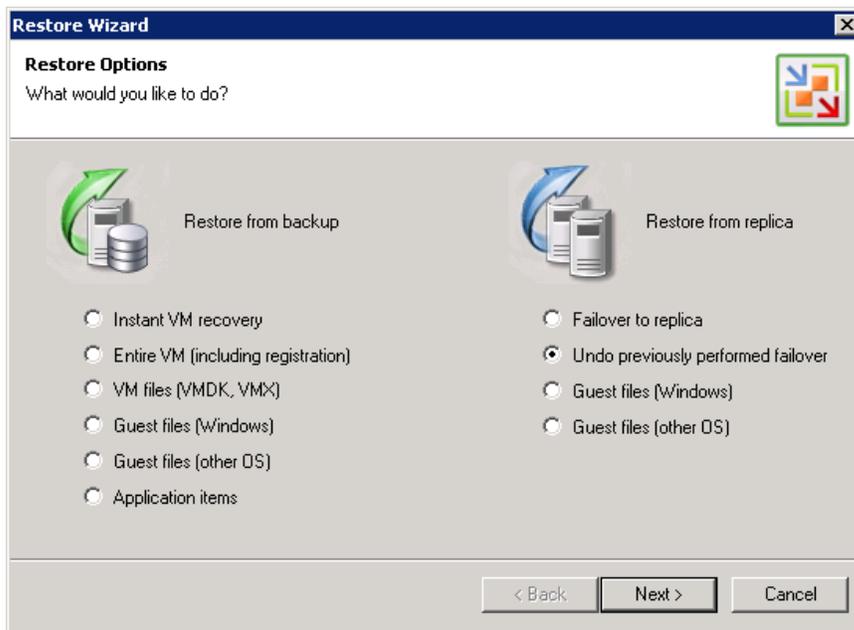
Step 1. Launch the Restore Wizard

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.

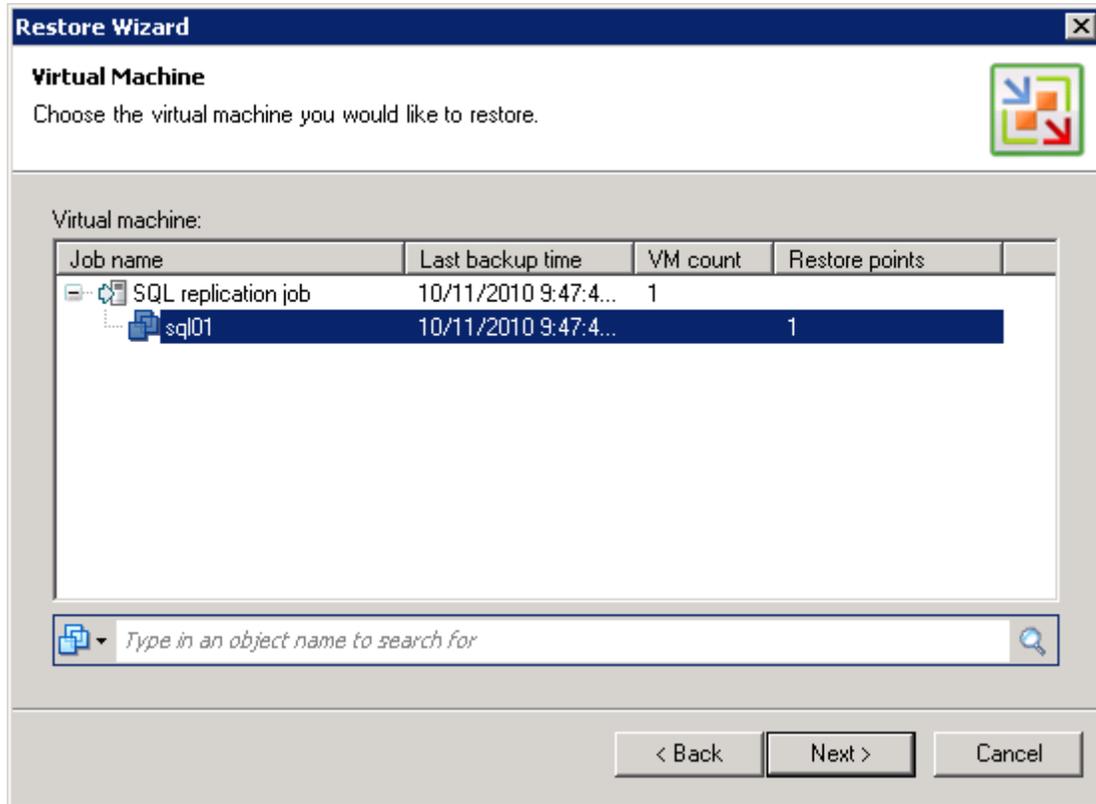
Step 2. Select a Task

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



Step 3. Select a Virtual Machine

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



Step 4. Finish Working with the Wizard

Click **Next**, then click **Finish** to finish working with the wizard.

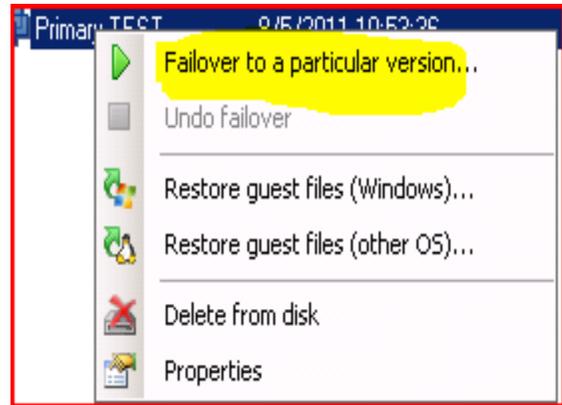
NOTE:

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

3.8 Monitoring Service:

It is very important to have some kind of network monitoring system in place which can send alerts to you in case of any failure.

Typically DR software for virtual environment have the capability to keep monitoring the virtual machine for heartbeat or some time for the application status, In the Veeam Software It do have the ability to trigger the recovery process but manually without any monitoring option.



3.9 Automatic Failover using Shell Script:

So, I have created simple shell script which will monitor the heartbeat of the Primary virtual machine every 10 seconds and trigger the recovery process if it did not receive ping response. You can also use http-ping utility and script it to monitor Http port.

```
@ECHO OFF
```

```
:loop
```

```
ping 192.168.2.10 -i 10 -w 500 > nul
```

```
if ERRORLEVEL 1 GOTO NOPING
```

```
Echo "Ping Successful"
```

```
GOTO loop
```

```
:NOPING
```

```
plink.exe -ssh -P 22 192.168.1.3 -l root -pw passwordhere -m cmds.txt
```

```
:END
```

I used Putty plink tool to connect to the Virtual host over SSH and put command line in a separate text file called cmds.txt

```
vim-cmd vmsvc/poweron 112
```

```
vim-cmd vmsvc/poweroff 64
```

Note: Script batch file and text file must be saved under same putty folder to make it work. Once script execute the SSH commands it will power on the secondary site and turn off the primary site virtual machine.

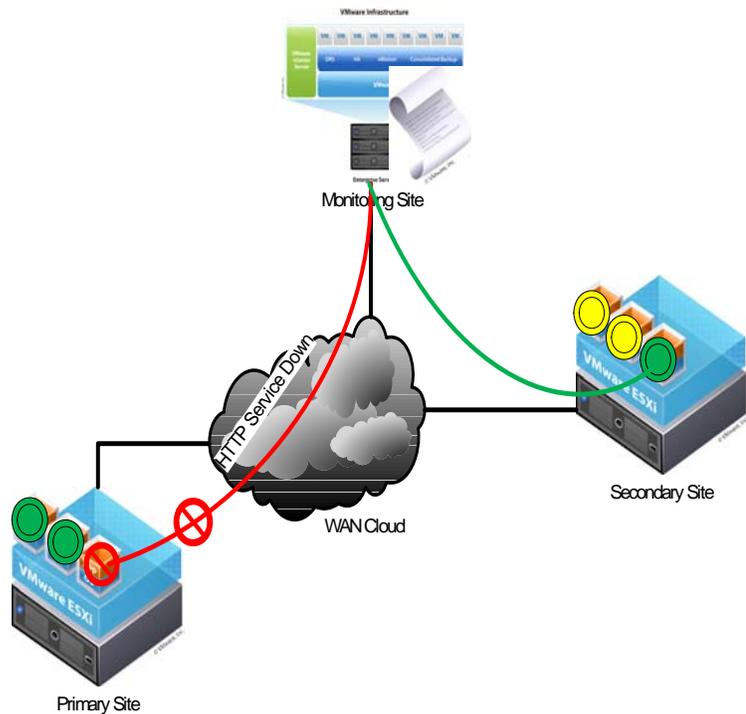
Chapter 5:

5.1 Testing:

Scenario 1.

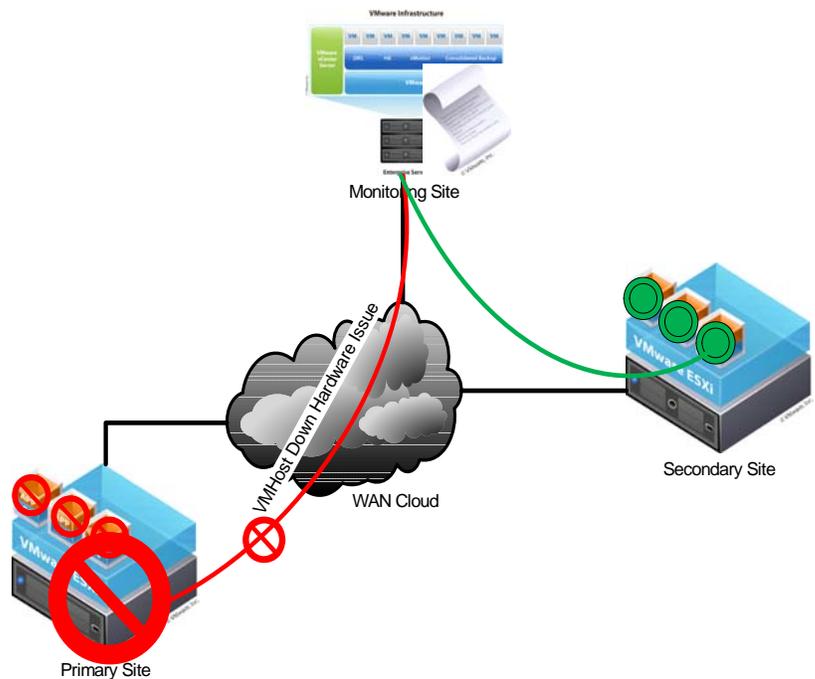
If Virtual machine is running http service, script will monitor the http response if using http-ping utility and will trigger the recovery process depending on what ever delay time you have set in the script. You can also setup another loop in the script if you want to check the server response on ping as well.

It is important to keep the same IP subnet and mask on the remote site and attached to the replicated virtual machine network and add additional Global DNS CNAME entry for remote site external IP to your DNS service provider.



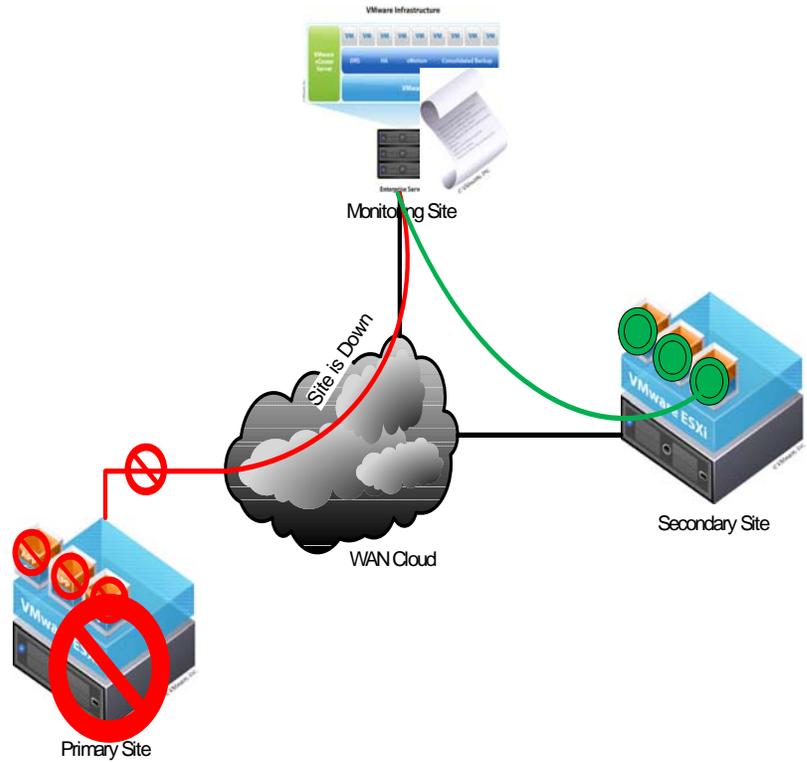
Scenario 2:

If Virtual host server fried up and unable to respond to the heartbeat request from the monitoring site, again the script will execute the recovery process and power on the remote site. One thing to mind here is that you need to make sure that Primary host will not come back online while the secondary site was up.



Scenario 3:

If the whole site for any reason go down and monitoring site is not getting response from the server or servers and also from the network, for example script will keep checking for heartbeat sending the ICMP request every 10 min in lab test environment to server and network devices and if the script ICMP request timed out then script will execute in the same manner as scenario 2 and bring up the remote site and shutdown the primary site if reachable by the monitoring system.



5.2 Conclusion:

When establishing a relationship between a replication source volume and its partner destination volume at the DR-site, a baseline mirror will need to be created. The table below is a rough estimate of how long it would take to initialize a volume given different network technologies. Keep in mind that these measurements are based on initializing the volumes one at a time, and presuming that each has dedicated use of the bandwidth.

Avoid routines that generate unnecessary write traffic. The first and most obvious that comes to mind is Disk Defragmentation. Presumably you will be implementing block-level replication in a SAN-based or shared storage environment where disk defragmentation has less value.

A relationship is created between a source volume and snapshot volume or replicated volume that present an image of the source at some particular fixed point of time. First replication or snapshot taken, we will call that time zero T0, with each replication Tn, we are copying most recent changes only you need to assure that the snapshots updates that need to be replicated from source volume to remote volume can take place within the defined replication latency window.

You will want to periodically test your DR environment to assure it is ready to deploy in case of failover.

Technology	Est. Hours To Replicate Capacity in GB					
	20	80	120	200	300	730
T1	42.33	169.31	253.97	423.28	634.92	1544.97
10Base-T LAN	6.50	26.01	39.01	65.02	97.52	237.31
DS3 / T3	1.50	6.02	9.03	15.05	22.57	54.93
100Base-T LAN	0.65	2.60	3.90	6.50	9.75	23.73
OC3	0.42	1.68	2.52	4.19	6.29	15.31
OC12	0.10	0.42	0.63	1.05	1.57	3.82

References:-

¹VMware.com

²Searchdatabackup.com

³Las.Solanas

Veeam.com

Cisco.com

VM knowledge base

Wiki references