VAssure
Value Assured

Vensoft India

Technical Case Study:

VMware Virtual Infrastructure
For
Backup, Restoration and Disaster Recovery

This is subsequent to Summary of Projects –

VMware - ESX Server to Facilitate:
IMS, Server Consolidation, Storage & Testing with Production Server

PRIMA - Panel for Remote Infrastructure Management Applications

VAssure | Virtualization Labs | trRIMS | Offshore-QA | BI | Portals

http://www.vassure.com

26 Nagarjuna Hills Panjagutta Hyderabad - 500 082 AP India

Infrastructure Management Services

In this project, we are providing Infrastructure Management Services to clients through using ESX Server and Virtualization Tools. It has production server which is an ESX server and it is attached to SAN (Storage Area Network) or RAID to store the data and other useful information as shown in Figure 1. Client is connected to the production server through internet. ESX Server does not require any guest operating system hence it needs minimum two processors to maintain and perform the required tasks. It may have many operating systems running simultaneously without interfering to data on different virtual machine.

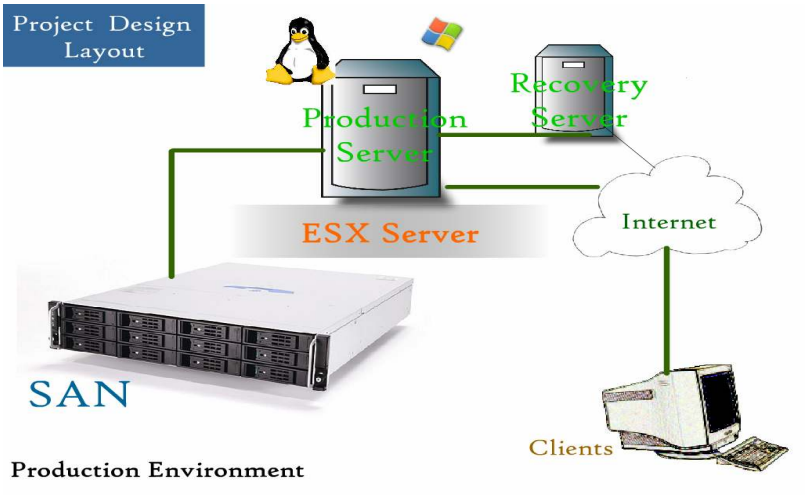

To include any patch or module at a production server at run time (i.e. when it is running and giving services to its clients), it is not advisable and feasible because of down time consideration, as server may halt which may disturb the working of organization as well as could suffer financial loss.

To avoid such a situation, we take a snapshot of the ESX Server (Production Server) and work on that snapshot in lab, which have same environment as the production server site has. Ssnapshot is a copy of a set of files and directories as they were at a particular point in the past. With the help of virtual machine, we can take snapshot of production server in different ways like cold, warm or hot snapshot as per requirement and situation.

R&D team first check the requirements and enhanced the patches or create a patch and send that patch to testing team, where the testing department persons attach it to kernel and check it by using automated tools. They perform functional test, regression test, performance Test. Once all are satisfied that there is no bug in the added patch, it is sent to the implementation team who attached this module to the client site.

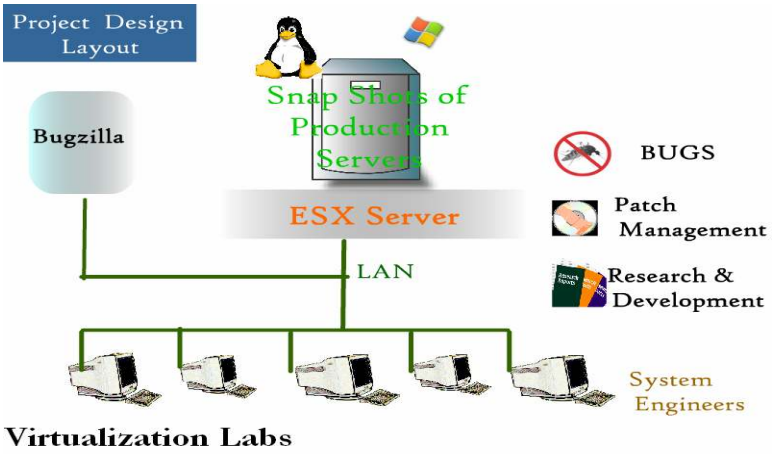

Introduction of ESX Server

VMware ESX Server is data-center class virtual machine software for consolidating and partitioning servers in high-performance environments. It is suited for corporate IT and service provider data centers, VMware ESX Server is a secure, cost-effective, highly scalable virtual machine platform. With advanced resource management capabilities it is also VMware's highest performance platform for building Virtual Infrastructure. VMware recognizes that in order for its customers to entrust valuable data to VMware ESX Server virtual machines, the platform must provide strong security. VMware provides for security in the ESX Server environment in several different ways,
Including:

i. ESX Server Architecture and the design of Virtual Machines
ii. Network security and VLAN's
iii. Strong Corporate Security Response Policy
iv. Independent Security Audit
v. Best practices for running an ESX Server securely

ESX Server Architecture and the Design of Virtual Machines

VMware ESX Server consists of the three major components as shown in Figure 1:

i. VMware virtualization layer – the vmkernel
ii. Service Console or Console OS
iii. Virtual Machines

Figure 1 – ESX Server Architecture

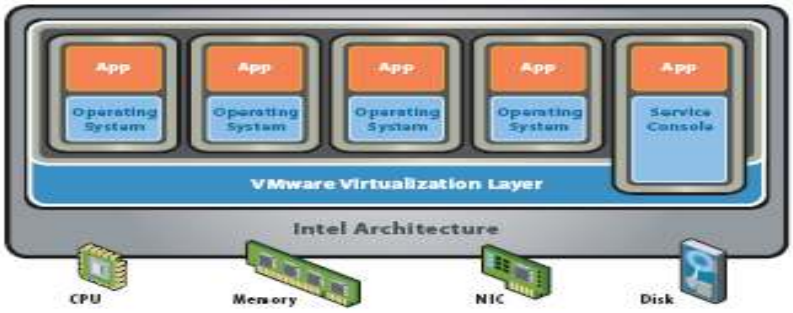

VMware Virtualization Layer – vmkernel

The vmkernel is a proprietary microkernel1 developed by VMware specially for running virtual machines. It is optimized for running virtual machines in the high performance ESX Server environment. The vmkernel controls the hardware and schedules the allocation of these resources between the virtual machines and the service console. The vmkernel has no public interfaces, and cannot execute a “process” in the traditional operating system sense. Hence, it is highly secure – there are no public interfaces to connect.

Virtual Machines

Virtual machines are the containers inside which guest operating systems are run. All VMware virtual machines have been designed to be completely isolated from each other. This isolation of virtual machines is key to enabling multiple virtual machines to run securely while sharing hardware and was a key factor in the design of virtual machines. This isolation of virtual machines applies to both their ability to access hardware and their performance characteristics.

Isolation - hardware

An executing virtual machine is isolated from other virtual machines running on the same hardware. This isolation is complete - while they share physical resources such as CPU, memory and I/O devices, they cannot see any device other than the virtual devices made available to it by the virtual machine monitor. Each guest operating system running inside a virtual machine behaves just as it was running on a separate machine and has no knowledge of other virtual machines running on that hardware

Figure 2 Virtual Devices

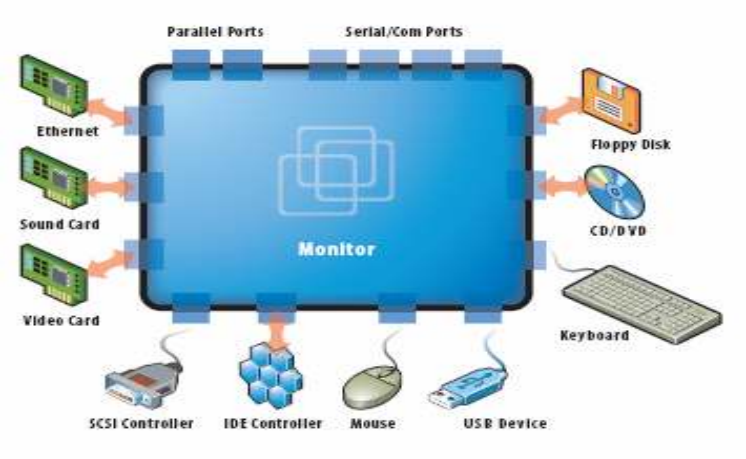

The only way a virtual machine can communicate with another virtual or physical machine is through the network, and its virtual network card (VMNIC) just like on a physical machine as shown below

Figure 3

If a virtual network card is not configured then the machine is completely isolated. Hence, any virtual machine that follows security procedures to protect itself from a network that a physical machine would, such as a firewall, antivirus etc. will be fully protected. *Figure 3 – Virtual Networking through Virtual Switches* It is important to note that the isolation of virtual machines operates at the virtual hardware level. Hence, it is operating below the guest operating system and so even a user with system administrator privileges for a guest operating system running inside a virtual machine cannot "reach out" and access another virtual machine without explicit access from the ESX Server system administrator.

Isolation – performance

The performance of a virtual machine can also be isolated completely from other virtual machines. Through the fine-grained resource controls available in ESX Server, we can configure a virtual machine to always get, for example, a minimum of 10% of a CPU. This protects a virtual machine from a loss of performance if another virtual machine was to consume too many resources on shared hardware. This can also prevent a malicious user in another virtual machine from affecting the performance of a virtual machine through a denial of service type attack. Since the vmkernel mediates the physical resources, and all physical hardware access is through the vmkernel, a virtual machine has no way to side step this performance isolation.

Service Console or Console OS

The ESX Service Console, also sometimes called the Console OS, is a limited distribution of Linux based on the Red Hat 7.2 distribution. The service console provides an execution environment to monitor and administer the entire ESX Server. In some cases, if the service console is compromised it can result in the virtual machines also being compromised

Backup, Restoration and Disaster Recovery

Backup, restoration, and disaster recovery are among the most crucial elements of datacenter management. VMware ESX Server System provides many different capabilities to improve these processes. The flexibility provided by Virtual Infrastructure, allows ESX Server backup and restores procedures into existing methodologies and procedures. VMware technology gives new and more advantageous ways of approaching these critical tasks. This technical note describes what resources should be backed up on an ESX server and explains all the options available for that backup, including advantages and disadvantages of each option.

Disk Structure of ESX Server

ESX Server uses VMware ESX Server File System (VMFS) for storage of virtual machines. VMFS is a high-performance file system on physical SCSI disks or IDE and partitions, optimized for storing large files such as virtual disk images and the memory images of suspended virtual machines. ESX Server 2 uses VMFS-2, earlier versions of the product used VMFS-1, which was more limited in its capabilities. VMFS-2 volumes can span multiple partitions, across the same or multiple (up to 32) LUNs or physical disks. A VMFS-2 volume is a logical grouping of physical disk partitions, which may also be called “physical extents”. Because the files stored on the VMFS may exceed 2GB in size, they cannot always be accessed using the same tools as files on a standard ext2, ext3, FAT, or NTFS file system.

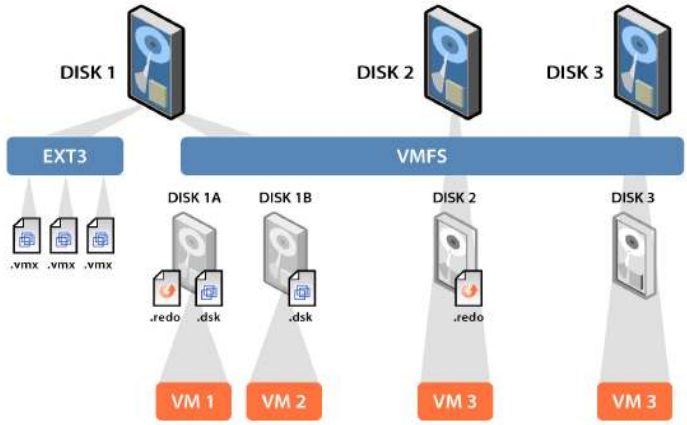

Virtual Machine Disks

As mentioned above, the disk files of virtual machines are stored on the VMFS file system. These files are in a special format and generally use .dsk or .vmdk file extensions. The disk files can comprise all the information the virtual machine stores on the virtual disk or be a symbolic link from a VMFS to a raw LUN when raw device mappings (RDM) are used.

Disk Files and Redo Logs

Redo logs make performing backups much easier. In the default state, a virtual machine disk is simply a single file. All changes to that disk are written directly and immediately to that .vmdk file. When a redo is added to a .vmdk file, that base disk file becomes static and unchanging. All writes are “trapped” in the redo log. This state is represented by file name: if the base disk is called disk.vmdk, the redo log will be called base.vmdk.REDO. A disk file may have a maximum of two redo logs, which is three files: base. vmdk, base.vmdk.REDO, and base.vmdk.REDO.REDO. The format of the redo log is a bitmap record of changes to the disk. Redo logs are, therefore, useful for disk snapshots. When a disk is represented by the two files disk.vmdk and

base.vmdk.REDO, disk.vmdk reflects the state of the drive at the time the disk snapshot was taken while base.vmdk.REDO is a bit-by-bit map of changes to the hard drive since that time.

Raw Device Mappings

With Raw Device Manager, VMware now has the most flexible storage capability for virtual machines in the industry. This mapping allows all of the features of VMware Virtual Infrastructure to be used in conjunction with raw SAN LUNs. The mapping file—the file that is used to connect the raw LUN to the virtual machine—is what is referenced in the virtual machine's configuration. The information about the raw LUN is stored within the mapping file, providing a consistent location for the virtual machine to find its disk across ESX servers even if the LUN is presented differently to each. There are two modes for RDMs: virtual compatibility and physical compatibility. Virtual compatibility mode allows a mapping to act exactly like a virtual disk file, including the use of redo logs. Physical compatibility mode allows direct SCSI access to the device being mapped for those applications that need lower level disk access and control. In both cases, data is stored on the LUN or SCSI device, not on the disk file. In both modes, an RDM file in a VMFS volume manages Metadata for its mapped device. There is a one-to-one mapping between mapping files and mapped devices. The mapping file is presented to the VMware Service Console as an ordinary disk file, available for file system operations.

To the virtual machine, the ESX Server presents the mapped device as a locally attached SCSI device. In physical compatibility mode, RDM provides minimal SCSI virtualization of the mapped device. In this mode, the VMkernel passes all SCSI commands to the device with one exception: the Report LUNs command is virtualized so that the VMkernel can isolate the LUN to the virtual machine that owns it. Otherwise, all physical characteristics of the underlying hardware are exposed. Physical compatibility mode is useful when there is a need to run SAN management agents or other SCSI target-based software in the virtual machine. This mode is also useful for virtual-to-physical clustering for cost-effective high availability.

Accessing and Managing Virtual Disk Files

Virtual disk files on the VMFS are accessible through ESX Server Service Console, ESX Server MUI, Virtual Center, and VMware SDK. From the service console, files can be viewed and manipulated on VMFS volumes under the /vmfs directory with ordinary file commands, such as ls and cp. Although mounted VMFS volumes may appear similar to any other file system, such as ext3, VMFS is primarily intended to store large files, such as disk images. The FTP, SCP, and cp commands can be used for copying files to and from a VMFS volume as long as the host file system supports these large files. The nfs protocol is subject to a 2G file size limitation and should not be used. Additional file operations are enabled through the vmkfstools command. This command supports the creation of a VMware ESX Server file system (VMFS) on a SCSI disk and

can be used to create, manipulate, and manage files stored in VMFS volumes. This command is also used to list files on the VMFS volume, add a redo log, commit a redo log, and export .dsk files into other formats. The vmsnap and vmres scripts that automate many of the common backup and restore tasks are run from the service console. This will be discussed in a later section of this document. Importing and exporting disk files can be done through the ESX Server MUI by copying the files from VMFS mount and pasting them to a partition running ext3 file system.

Back Up on ESX Server

With ESX Server, there are three major components, which may need to back up:
- Virtual disks
- Virtual machine configuration files
- The configuration of the ESX Server itself

All the information normally backed up in the enterprise infrastructure, including the operating system, applications, and data, is included in the virtual disks. Because a virtual machine is just like a physical machine, one possible approach is to back it up in the same manner as a physical machine, using backup software running inside a virtual machine. With only a few files encapsulating an entire virtual machine, it is very simple to get back to a previous known state at a known time. Another possible approach is to back up all the files on the ESX Server that make up a virtual machine. This approach also allows for easy check pointing without any additional third-party software or hardware. Two levels of redo logs can be created, allowing maintenance of multiple snapshots. The use of redo logs allows for hardware-independent by taking snapshot of virtual machines, yielding true point-in-time copies without the use of SAN features. Yet another possibility is off-line backups where files encapsulating virtual machines are accessed and backed up without loading the ESX Server that the virtual machines normally run on. Storage replication must be used to assist in this process

Virtual Machines as Physical Machines

Virtual machines provide complete guest operating systems (OS) on virtualized hardware. These OS installations can be backed up in the same ways as their physical counterparts. One may attach backup hardware to the virtual machines and install backup servers in them. Alternately, one may install a backup agent within each VM, and back up data over the network to other backup servers. Finally, one might copy data manually or with a script to another machine. Backing up a virtual machine in this way is precisely like backing up a physical machine.

Treating Virtual Machines as Files on ESX Server

This approach takes advantage of the service console's ability to see each virtual machine's virtual disk as a file. ESX Server creates one file per virtual machine disk device, with redo logs stored separately. These .dsk or .vmdk files can be backed up, effectively backing up an entire virtual hard drive in a single pass. This approach is not possible if we are using RDM disks for virtual machines. To take advantage of this method, we must add and commit REDO logs to the base disk of each VM as the base disk is backed up.

Treating Virtual Machines as Files on Shared Storage

When virtual machine files reside on shared storage, it is possible to use SAN-based imaging or an independent backup server to back up virtual machine files without creating an additional load on the ESX Server where the virtual machines normally run.

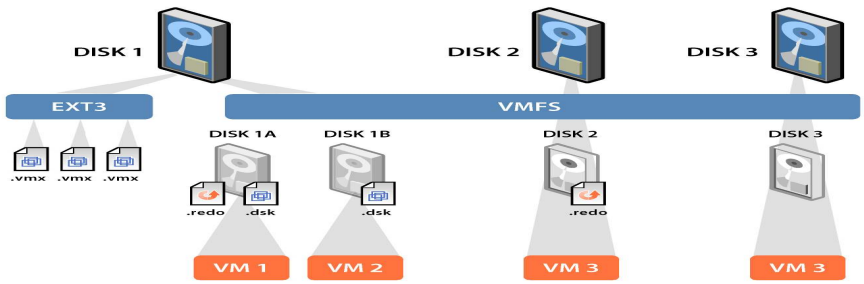

Backing up the Service Console

The service console is a specialized virtual machine running Linux, which provides a management interface for both command prompt and Web-based, to the ESX Server and the virtual machines running on it. The service console is nearly stateless; however, having it safely backed up may save time in a disaster recovery scenario.

Backing Up Virtual Machines as Physical Machines

Compatible with the following virtual driver formats:
• Virtual disks (all formats)
• RDM disks (all formats)

Issues to Consider

The following issues need to be considered before deciding on a backup method:
• Whether the backup server is on a physical or virtual machine
• The network configuration
• Type of long-term storage

Implementation Steps

1) Setup.

   a) Install the backup agent of your choice on each virtual machine to be backed up.
   b) Schedule the backups and manage the tapes as directed by the documentation for your backup program.
   c) Configure your backup server or node. If the backup server is installed in a virtual machine on the ESX Server, the following steps are required.
      i) Attach the tape drive or library to a SCSI port on the ESX Server.
      ii) Through the ESX Server Web management interface, assign the tape drive to a virtual machine.
   d) Configure the virtual machine to use the tape drive or library, installing the appropriate drivers moreover, backup server of choice.

2) Ensure that networking is configured for access between the backup server and virtual machines to be backed up. If both virtual machines to be backed up and the backup server are on the same ESX Server, we may use a private virtual network switch to connect them to each other.

3) Install the backup agent on all virtual machines to be backed up.

4) Backup and restore.
   a) Follow the instructions for the backup software we installed.

Architecture for the Setup:

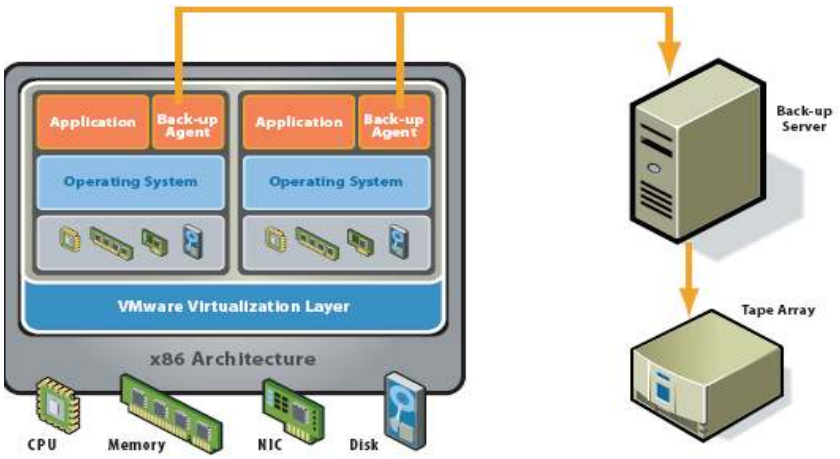

Creating Backup Copies of Virtual Machine Files on ESX Server

Up to three virtual disk files that live on the VMFS and the virtual machine configuration file, which resides on the file system managed by the ESX Server Service Console, represent each virtual machine. Disk files are frequently larger than 2GB, and not all backup programs can directly access its files. By default, Virtual disk files cannot be backed up while the virtual machine is powered on because, during this time, the disk file is open and being written to by a virtual machine. Powering off or suspending the virtual machine closes the virtual disk file and makes it safe to back up. This may not a tenable solution for most enterprise situations. With the REDO log capability, one can add a REDO log to the VM. All new writes are trapped in the REDO log, making the base virtual machine disk file static and available for back up. Alternatively, we can take a snapshot the virtual machine. This captures all new writes to a redo log, making the base virtual machine disk file .dsk available for back up. Doing so gives us a safe and consistent snapshot of the virtual disk to back up.

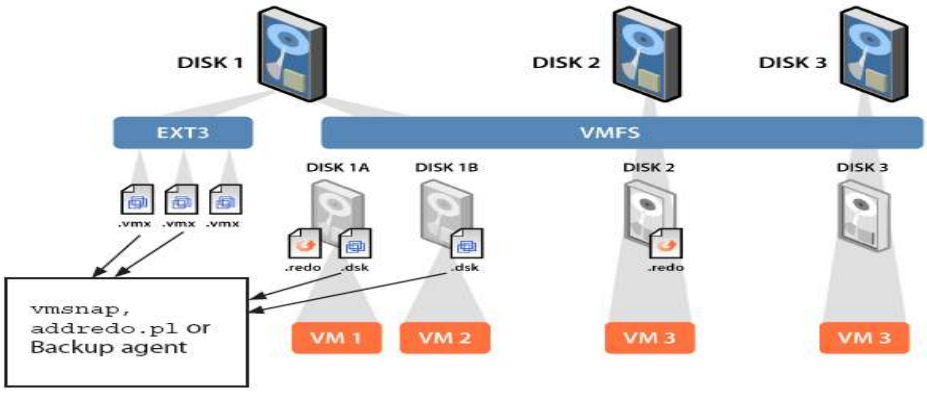

Built-in VMware File Operations

Compatible with the following virtual driver formats:
• Virtual disks only (all formats)
Virtual disks and other virtual machine files, such as the configuration file, logs, and memory, can be manipulated from ESX Server Service Console command line through scripting API commands and VMware Virtual Infrastructure SDK.

SAN Image

It is compatible with the following virtual driver formats:
• Virtual disks (all formats)
• RDM disks (all formats)

If our virtual disk files are stored on a SAN, we can use features supplied by our SAN vendor to create a copy of his production LUN, containing all virtual disks. These copies can then be sent to our backup media. With this method, we do not have to use virtual machine to take snapshot functionality during the backup process because the SAN snapshot guarantees consistency. Taking Snapshot and replicating SAN volumes requires the use of layered applications with SAN. ESX Server is compatible with some SAN layered applications but not all as standards are still evolving.

Server-Based Replication

Unlike SAN-based replication, this approach requires a replication agent in each of the virtual machines. Because of this, the replication agent has visibility into the file system and the application on the virtual disk and is capable of incremental replication where only the files that have been modified since the time of the previous replication are copied. On the other hand, we cannot take advantage of the encapsulation of virtual machines into few files. A replication server and replication targets are required.
Compatible with the following virtual driver formats:
• Virtual disks (all formats)
• RDM disks (all formats)

Using Specialized ESX Server for Backup

ESX Server's built-in snapshot function can also be used in the backup off-load scenario. We can create a snapshot of a virtual disk located on shared storage by adding a redo log. Thereafter, ESX Servers can access the virtual disk snapshot. A designated backup ESX Server can then access the disk snapshot and back it up using a backup agent installed in the backup ESX Server.

Implementation Steps

1. Using the SAN management software, schedule snapshots on the disk backend.
2. Schedule the backups of the SAN snapshots and manage the tapes as directed by the documentation for our backup program

Architecture: using SDK snapshot interface and ESX Server

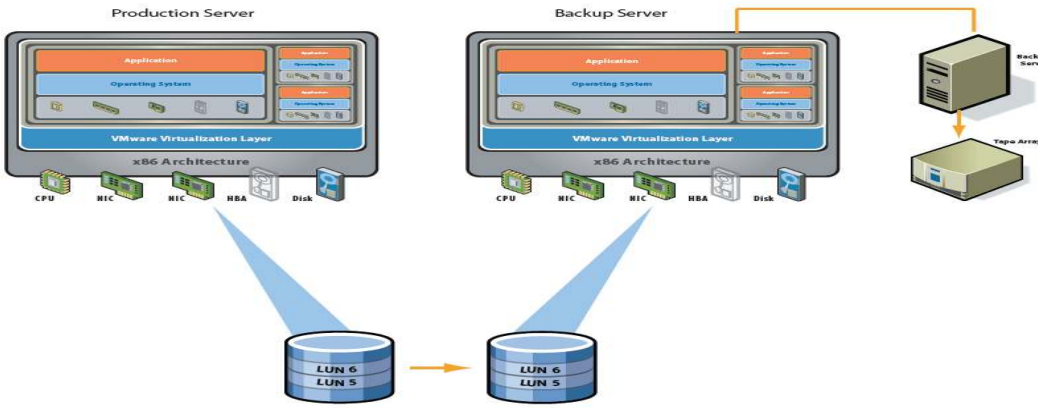

Compiled by: Dushyant Sharma
dushyant.sharma@vassure.com

This paper is not intended to be a definitive implementation guide. Many factors are not addressed in this document. Expertise may be required to solve logistical problems when the system is designed and built. VAssure team has not tested this procedure with all the combinations of hardware and software options available on all VMware products or guest OS variants. There may be significant differences in your configuration that will alter the procedures necessary to accomplish the objectives outlined in this paper.