

# NAS Concepts with VMware Infrastructure 3

An overview of using Network Attached Storage in a VMware World

IIPG Group, Storage

Sun Microsystems

December 12, 2006



© 2005 Sun Microsystems, Inc., 4150 Network Circle, Santa Clara, CA 95054 USA

All rights reserved.

This product or document is protected by copyright and distributed under licenses restricting its use, copying, distribution, and decompilation. No part of this product or document may be reproduced in any form by any means without prior written authorization of Sun and its licensors, if any. Third-party software, including font technology, is copyrighted and licensed from Sun suppliers.

Parts of the product may be derived from Berkeley BSD appliances, licensed from the University of California.

Sun, Sun Microsystems, Sun StorEdge, the Sun logo, are trademarks, registered trademarks, or service marks of Sun Microsystems, Inc. in the U.S. and other countries.

UNIX is a registered trademark in the United States and other countries, exclusively licensed through X/Open Company, Ltd.

Windows is a registered trademark of Microsoft Corporation in the United States and other countries.

All SPARC trademarks are used under license and are trademarks or registered trademarks of SPARC International, Inc. in the U.S. and other countries. Products bearing SPARC trademarks are based upon an architecture developed by Sun Microsystems, Inc.

The OPEN LOOK and Sun's Graphical User Interface was developed by Sun Microsystems, Inc. for its users and licensees. Sun acknowledges the pioneering efforts of Xerox in researching and developing the concept of visual or graphical user interfaces for the computer industry. Sun holds a non-exclusive license from Xerox to the Xerox Graphical User Interface, which license also covers Sun's licensees who implement OPEN LOOK GUIs and otherwise comply with Sun's written license agreements.

RESTRICTED RIGHTS: Use, duplication, or disclosure by the U.S. Government is subject to restrictions of FAR 52.227-14(g)(2)(6/87) and FAR 52.227-1987), or DFAR 252.227-7015(b)(6/95) and DFAR 227.7202-3(a). DOCUMENTATION IS PROVIDED AS IS AND ALL EXPRESS OR IMPLIED CONDITIONS, REPRESENTATIONS AND WARRANTIES, INCLUDING ANY IMPLIED WARRANTY OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE OR NON-INFRINGEMENT, ARE DISCLAIMED, EXCEPT TO THE EXTENT THAT SUCH DISCLAIMERS HELD TO BE LEGALLY

# Table of Contents

Introduction.....	1
Background.....	1
NAS Configurations.....	3
ISCSI Software Initiator.....	3
ISCSI Hardware Initiator.....	4
General NFSv3.....	5
Direct Virtual Machine Connection.....	6
Conclusion.....	7
References.....	8

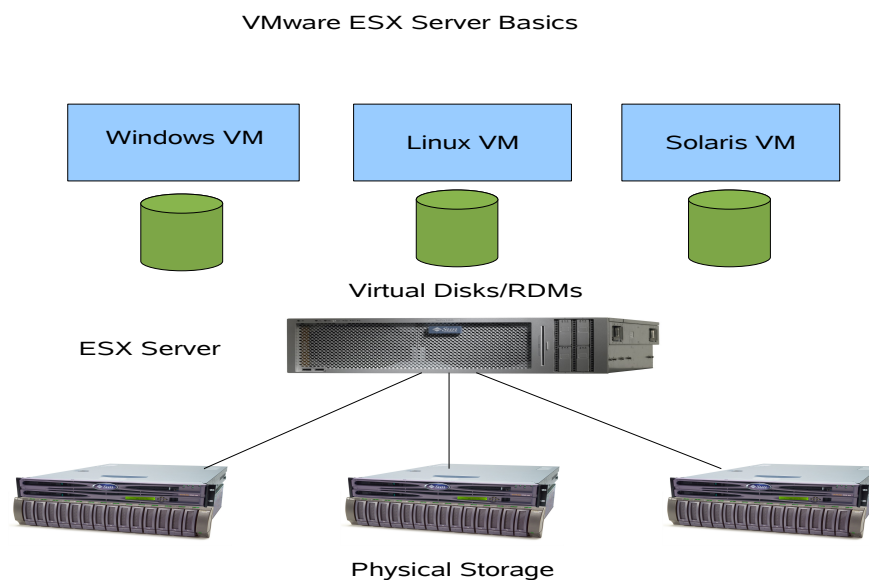
## Introduction

This short paper details a few of the possible combinations when using NAS devices, such as the Sun StorageTek 5000 Family, with VMware Infrastructure 3. It will detail a few of the combinations that are possible in an abstract sense, how they differ from one another and perhaps most importantly, what is currently possible with current Sun products and what is coming in the future.

## Background

VMware Infrastructure 3 is a bundling of VMware's ESX server, Virtual Infrastructure Client and various other add-on services. ESX Server serves as the foundation for virtualization and allows IT administrators to abstract physical hardware resources such as CPUs, Memory, Network and Storage and build multiple virtual machines, allowing for greater flexibility and higher utilization of those resources. ESX Server should be thought of as its own Operating System and various configurations in this paper deal with how to connect those NAS resources directly to the ESX server. Other configurations deal with how those NAS resources are connected to the Virtual Machines themselves.

The basic concept of storage for Virtual Machines is shown in the diagram below.



Internal Only

In the above diagram, the ESX server can be seen as presenting the physical storage to the Virtual Machines as Virtual Disks (vdisks) or RDMs (Raw Device Mappings).

In the vdisk type of configuration, luns and/or volumes of various sizes are created on the storage device and then connected via FibreChannel or Network connections to the ESX server. The ESX server then is able to create filesystems (in the case of Block FC or iSCSI storage) or use the NFSv3 filesystem (in the case of NAS devices) on which large files, called Virtual Disks are created. These vdisks are then associated with the Virtual Machines via a virtual SCSI adapter. The Virtual Machines see the vdisks as SCSI luns and are able to install an Operating System and the needed applications onto those SCSI disks.

In the RDM type of configuration, luns/volumes are still attached to the ESX server, but instead of creating a filesystem to hold various vdisks, as in the Virtual Disk model, the entire lun is mapped directly to the Virtual Machine, bypassing the ESX filesystem (and associated caching). This is useful for Virtual Machines needing higher levels of performance and those requiring Microsoft Cluster installations. RDMs also present the opportunity to do some interesting things with array level snapshots, mirrors and remote replication. These topics will be discussed in follow-on whitepapers.

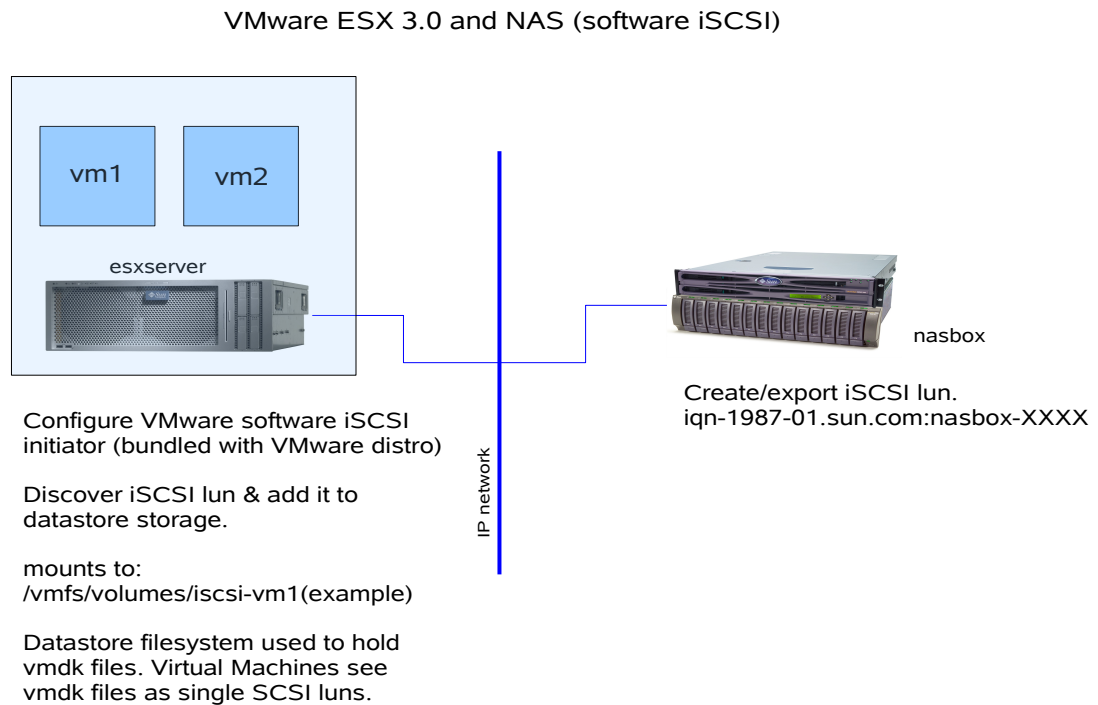
These are the configurations that are tested and certified in order to be added to VMware's Hardware Compatibility List. This testing verifies general connectivity, boot-from-san capabilities and multi-pathing/failover support.

The rest of this paper will delve further into the various storage configurations with NAS devices.

## NAS Configurations

### iSCSI Software Initiator

In this configuration, NAS volumes are exported as iSCSI volumes and connected to the ESX server by using VMware's bundled software iSCSI initiator. No iSCSI hardware is needed.



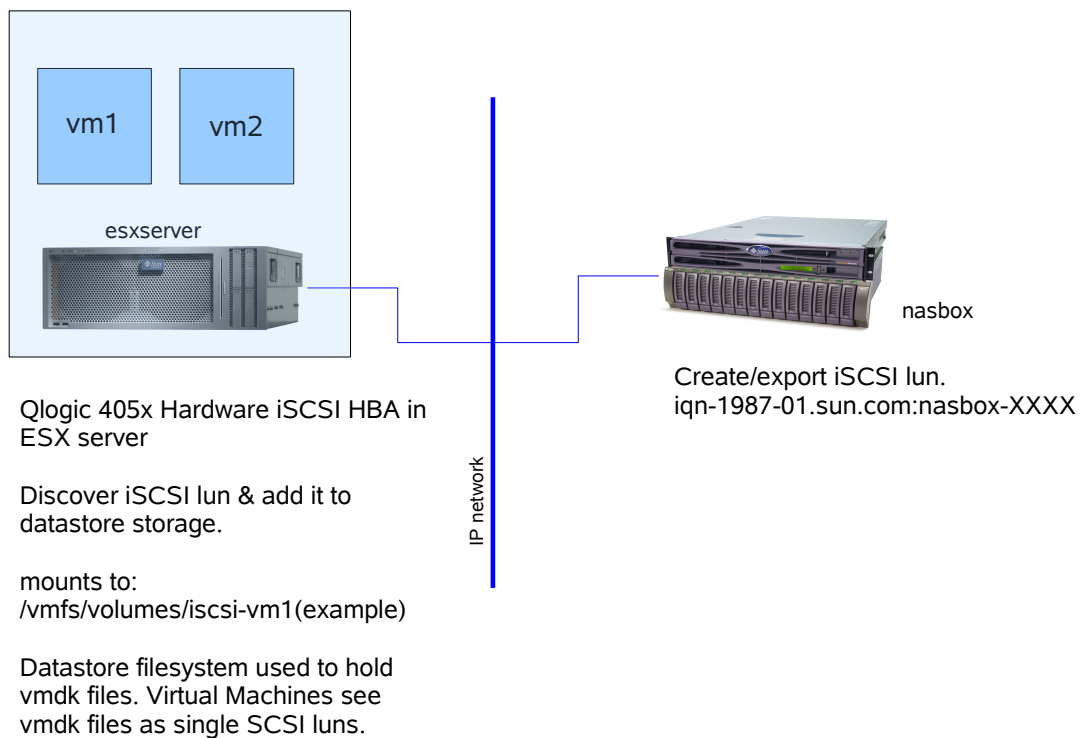
As depicted in the above diagram, an iSCSI lun is created on the NAS device and presented to the ESX Server's software iSCSI initiator. The lun then shows up as a SCSI device to ESX and can be either configured with a vmfs3 filesystem for storing vdisks or can be presented as a RDM lun directly to the Virtual Machine.

This configuration is currently supported by VMware as of VI3 3.0 (ESX 3.0). Sun is working towards certifying the 5320/5310 with this configuration in the FY07 Q2/Q3 timeframe.

## iSCSI Hardware Initiator

In this configuration, NAS volumes are exported as iSCSI volumes and connected to the ESX server by using an add-on Hardware iSCSI initiator card. As the diagram shows, it is very similar to the software iSCSI configuration, with only the initiator technology (hardware vs. software) being different.

### VMware ESX 3.0 and NAS (Hardware iSCSI)



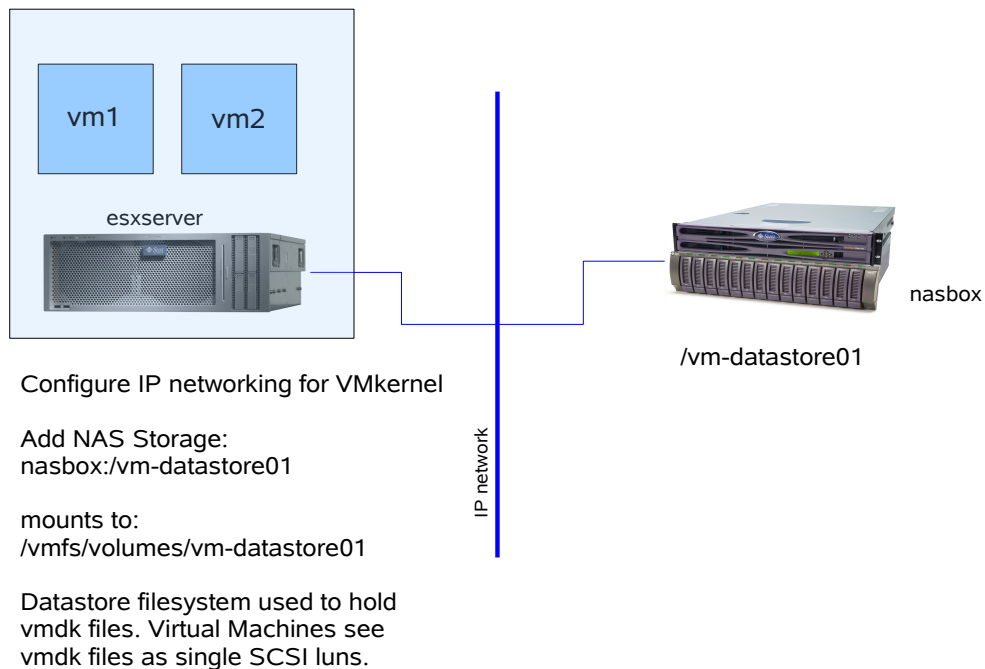
VMware has indicated that this support will be available late 2007 in VI3 3.0.1 (ESX 3.0.1) with the Qlogic 4040/4052 iSCSI cards only. Sun will pursue certification for this configuration using Hardware iSCSI initiator connections.

## General NFSv3

In this configuration, the NAS volume is exported as a general lun that connects to the ESX server via NFSv3 protocols, very similar to any other Solaris/Linux NFS mounted device. The ESX server must be given an additional IP address for the Vmkernel network configuration.

Once connected, the lun appears as a datastore just like hardware/software iSCSI connected luns. However, a NFSv3 lun cannot be used as a RDM device.

### Vmware ESX (VI3) & NAS (General NFS)

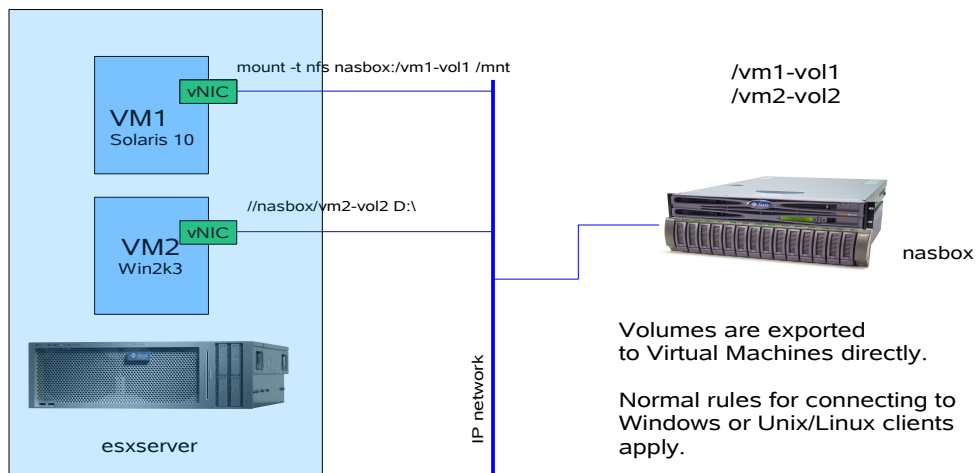


VMware has yet to announce a formal program for certifying this configuration so at this time it would not be supported as a Production level configuration. Sun is working with VMware on the details and is investigating any roadmaps for future certifications.

## Direct Virtual Machine Connection

The last configuration goes away from connecting the NAS luns/volumes directly to the ESX server and instead configures the Virtual Machines to have direct access themselves. Each Virtual Machine has a Virtual NIC that access the IP network via the ESX Server Vmkernel and the ESX Server's physical network card. The Virtual Machines have their own IP addresses and participate on the network as any other Operating System would. It therefore has access to NAS luns/volumes as any other Operating System would. (CIFS for Windows, NFSv3 for Solaris/Linux, or perhaps iSCSI). The fact that the machine is virtual is transparent to both the Operating System and the NAS device. Any access should be treated as if the machine was running on its own dedicated hardware.

VMware ESX 3.0 and Direct VM NAS connect



Virtual Machines use Virtual NICs (vNIC) to participate on the network. Have separate IP addresses from ESX Server.

NAS volumes configured and connected to VMs as normal NFS shares or Windows shares.

The expectation is that the Operating Systems and applications will behave as they would in a purely physical hardware world. There is no additional VMware certification or testing that is needed to run in this configuration. The normal rules for NAS device and OS connectivity would apply.

## Conclusion

This paper looked at various configurations that are possible in a VMware VI3 (with ESX server) and NAS array deployment.

- NAS connectivity to ESX via software iSCSI initiator
- NAS connectivity to ESX via hardware iSCSI initiator
- NAS connectivity to ESX via general NFSv3
- NAS connectivity to Virtual Machines

VMware certification status is only relevant when the connection is to the ESX server. Normal OS Interoperability rules apply when connecting to Virtual Machines.

It is expected that as the certification status changes for the various configurations, this paper will be updated to reflect the most current status in conjunction with the normal channels of communication. (Onestop, Interop Tool).

## References

- i. <https://extranet.storitek.com/interop/interop> – Interop Tool Website
- ii. <https://onestop/vmware> – Onestop VMware site
- iii. <http://www.vmware.com/support/pubs/> - VMware documentation and HCL's
- iv. <http://www.sun.com/storagetek/nas/5320/> - Sun StorageTek 5320 page