

ESX Server 3.0 Tips and Tricks

Mostafa Khalil, VCP

VMware Product Support Engineering



VMWORLD 2006

Agenda

- Storage Area Network
- VMFS-3
- iSCSI Storage
- NFS

Storage Area Network (SAN)


- Components
- Topology
- Identification

SAN Components

- Initiators (host HBAs)
 - QLogic 1/2/4Gb
 - Emulex 1/2/4Gb
 - See HCL for supportability
 - Set HBA's BIOS to "Point-to-Point" or similar to that
- Targets (Storage Processors' Ports)
 - Active/Active Arrays
 - Active/Passive Arrays
 - Check the HCL for supportability
- Fabrics (Switches and Fibre Connections)
 - No specific listing on HCL
 - 1/2/4Gb switches
 - May need to set switch port to "Fx" or equivalent (e.g. F-Port)

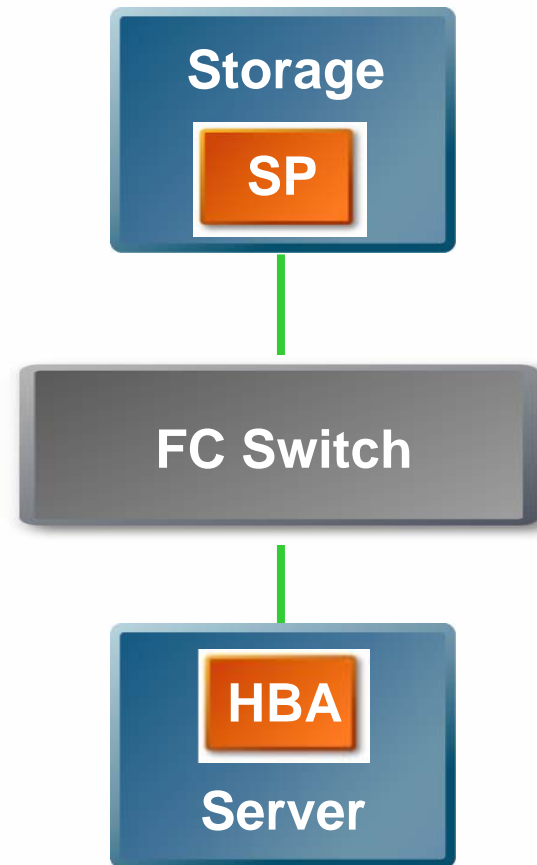
SAN Topology – Point-to-Point

- AKA Direct Connect
- Not supported
- Except CX100




SAN Topology – Absolute Minimum

- Switch allows additional connections of more storage or more servers
- No redundancy
 - Many single points of failure
- Supported by VMware for ESX Server as a bare minimum configuration
 - If the Storage Arrays are not listed on the hardware compatibility list (HCL), VMware only supports them with this simple configuration




SAN Topology – Arbitrated Loop

- Not Supported




SAN Topology – Multipath Fabric



SAN Topology – Fully Redundant Fabrics

- Recommended

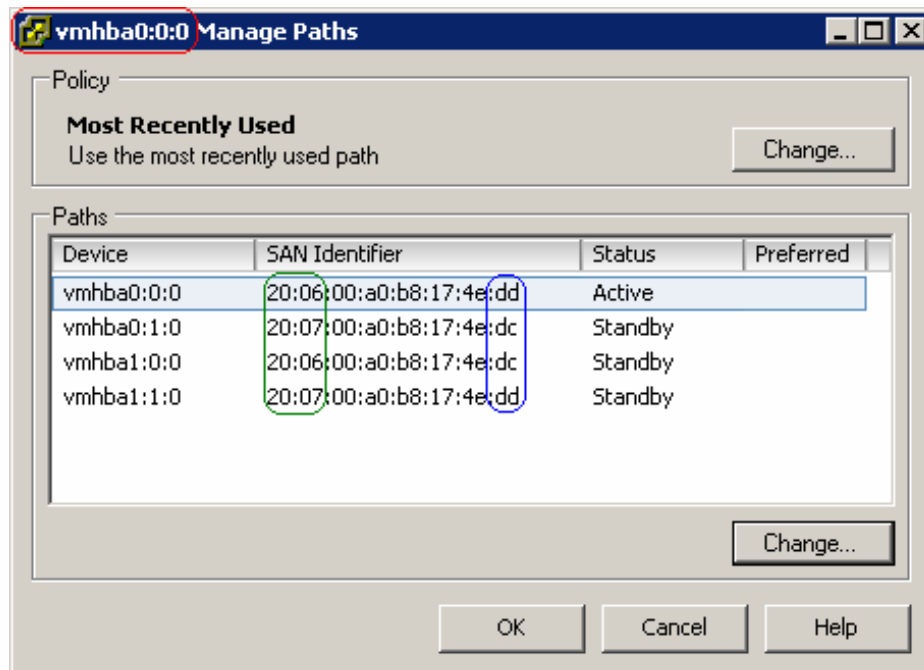


Identifying SAN Configuration

- VI Client
- Console

Multipath Analysis

- Canonical Name
 - vmhbaC:T:L
- In this example
 - Vmhba0:0:0
- Policy: MRU or Fixed
- In this Example: MRU
- Path States:
 - On
 - Off
 - Standby
 - Dead




```
root@giza:~  
[root@giza root]# esxcfg-mpath -l  
Disk vmhba0:0:0 /dev/sda (512MB) has 4 paths and policy of Most Recently Used  
FC 2:2.0 210000112593a86a<->200600a0b8174edd vmhba0:0:0 On active preferred  
FC 2:2.0 210000112593a86a<->200700a0b8174edc vmhba0:1:0 Standby  
FC 2:2.1 210000112593a86b<->200600a0b8174edc vmhba1:0:0 On  
FC 2:2.1 210000112593a86b<->200700a0b8174edd vmhba1:1:0 Standby
```

Multipath Analysis - Identifying Targets

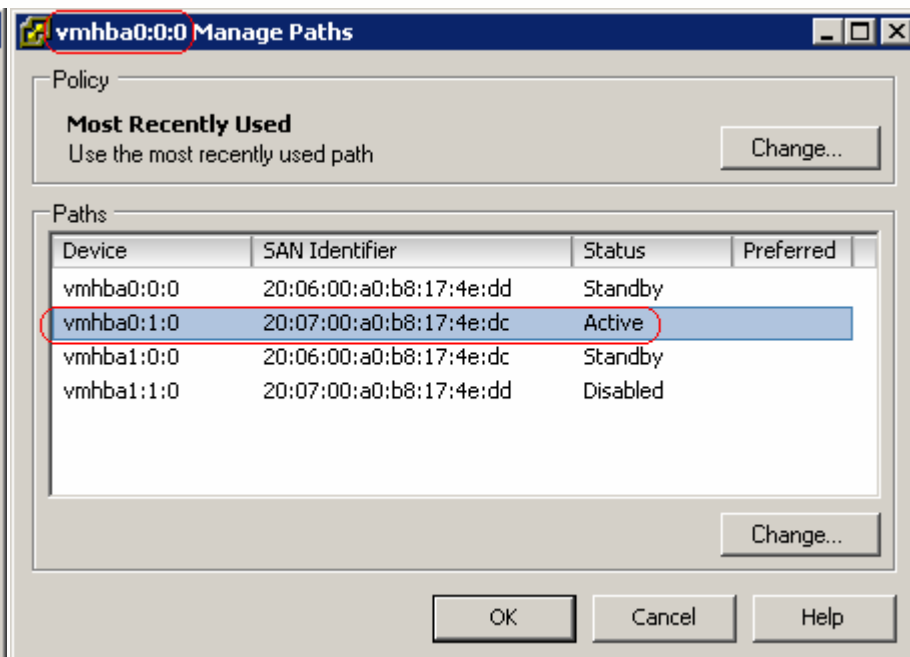
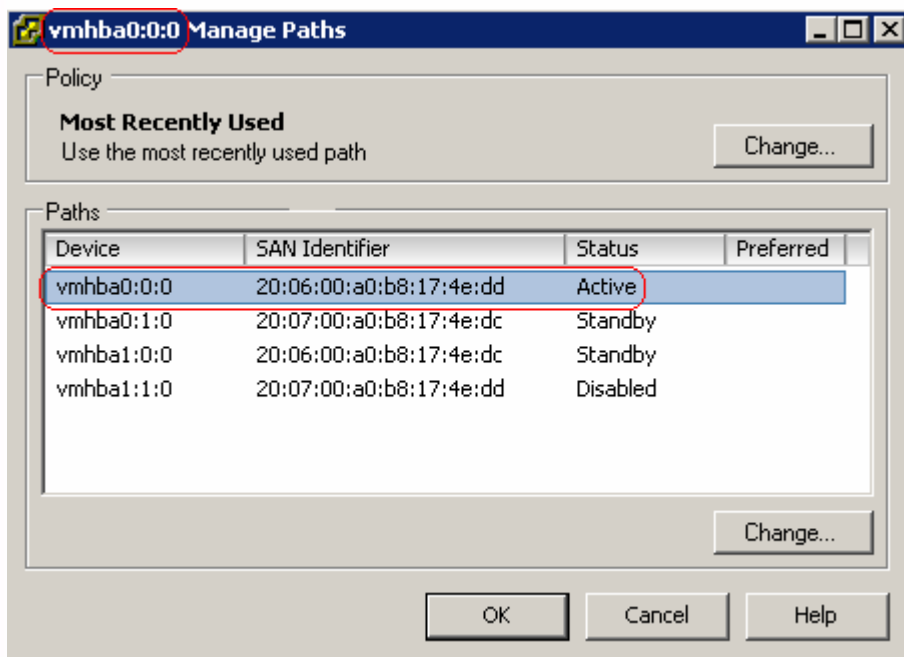
SAN Type	Port ID	WWN
EMC Clariion	SPA0	xx:xx:xx:60:xx:xx:xx:xx
	SPA1	xx:xx:xx:61:xx:xx:xx:xx
	SPA2	xx:xx:xx:62:xx:xx:xx:xx
	SPA3	xx:xx:xx:63:xx:xx:xx:xx
	SPB0	xx:xx:xx:68:xx:xx:xx:xx
	SPB1	xx:xx:xx:69:xx:xx:xx:xx
	SPB2	xx:xx:xx:6A:xx:xx:xx:xx
	SPB3	xx:xx:xx:6B:xx:xx:xx:xx
HP EVA	SPA1	xx:xx:xx:xx:xx:xx:xx:x9
	SPA2	xx:xx:xx:xx:xx:xx:xx:x8
	SPB1	xx:xx:xx:xx:xx:xx:xx:xD
	SPB2	xx:xx:xx:xx:xx:xx:xx:xC
IBM FASTT	N/A	20:0X:00:00:00:00:xx
	N/A	20:0Z:00:00:00:00:zz

For IBM FASTT:
 Compare X and Z:
 Higher number is the
 secondary processor.
 Compare xx and zz:
 Higher number is the
 higher number port.



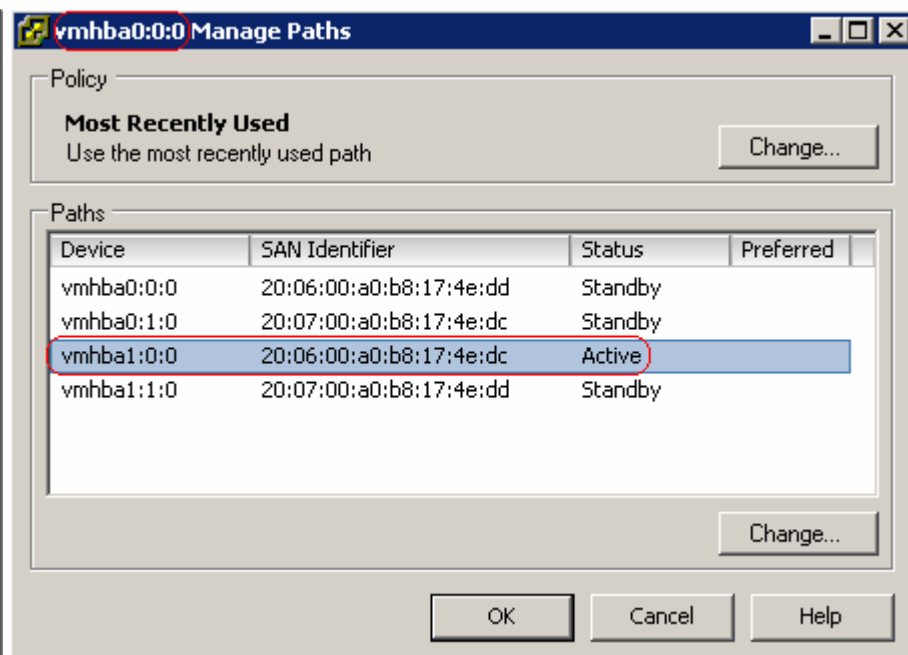
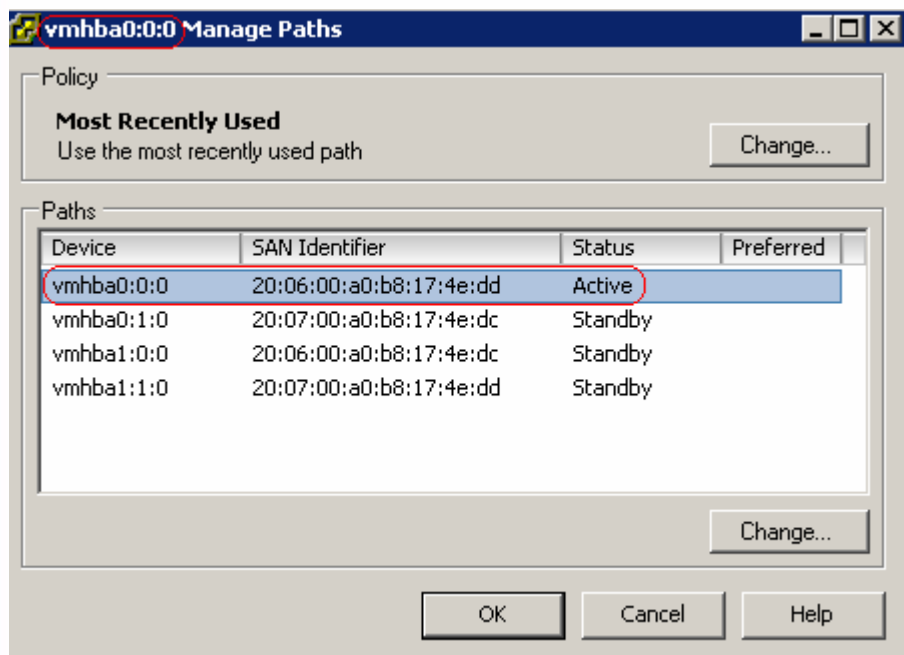
VI Client – Multipath Analysis (Target Failover)

- Canonical name not equal to active path → Failover
- Target Failover example



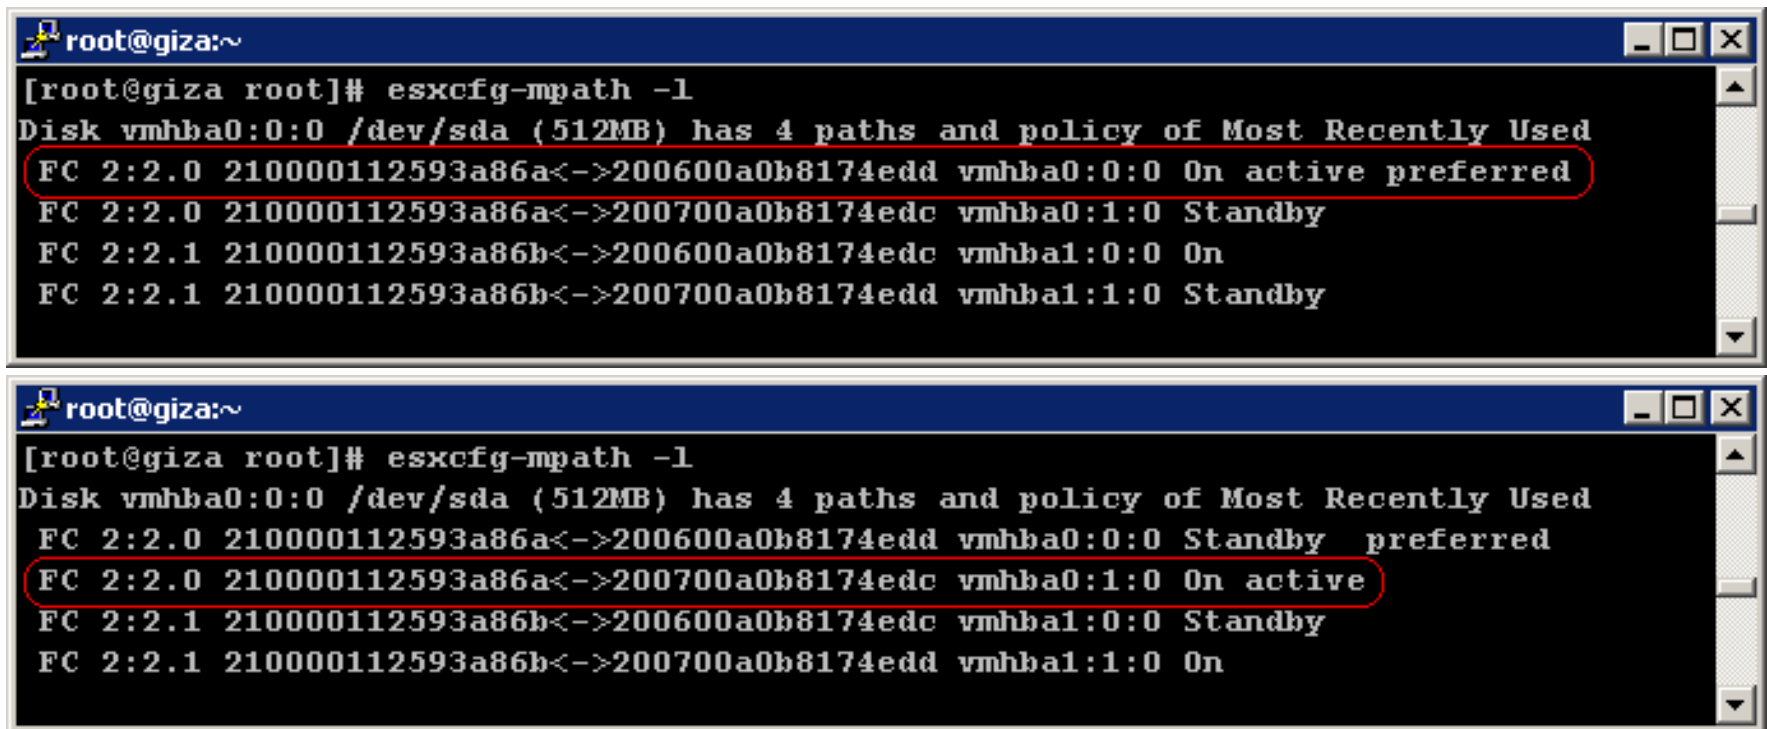
VI Client – Multipath Analysis (HBA Failover)

- Canonical name not equal to active path -> Failover
- HBA Failover example



Console – Multipath Analysis (Target Failover)

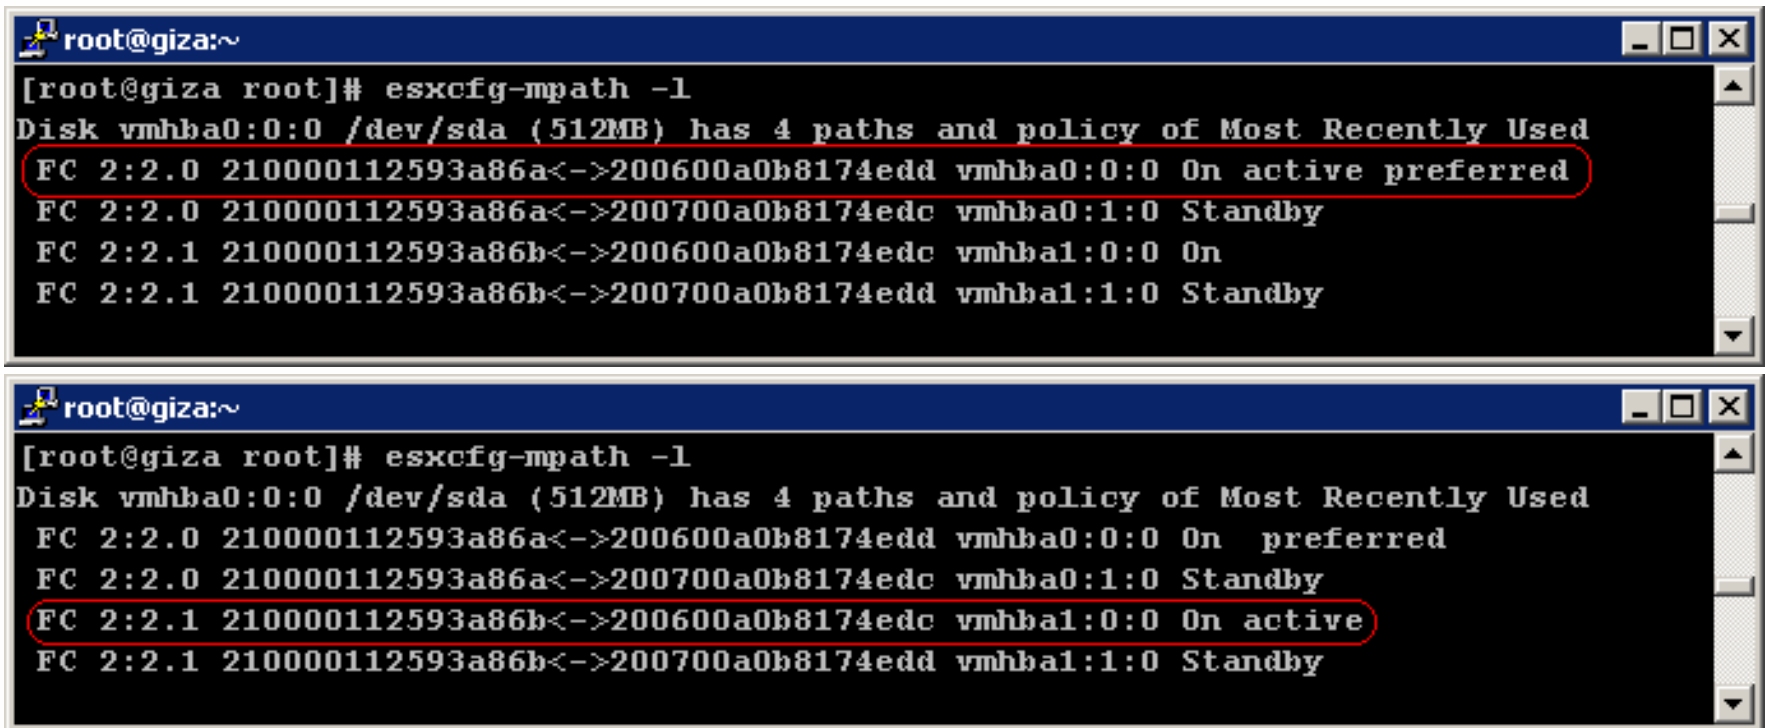
- esxcfg-mpath -l
- Console View of the target failover event



```
root@giza:~  
[root@giza root]# esxcfg-mpath -l  
Disk vmhba0:0:0 /dev/sda (512MB) has 4 paths and policy of Most Recently Used  
FC 2:2.0 210000112593a86a<->200600a0b8174edd vmhba0:0:0 On active preferred  
FC 2:2.0 210000112593a86a<->200700a0b8174edc vmhba0:1:0 Standby  
FC 2:2.1 210000112593a86b<->200600a0b8174edc vmhba1:0:0 On  
FC 2:2.1 210000112593a86b<->200700a0b8174edd vmhba1:1:0 Standby  
  
root@giza:~  
[root@giza root]# esxcfg-mpath -l  
Disk vmhba0:0:0 /dev/sda (512MB) has 4 paths and policy of Most Recently Used  
FC 2:2.0 210000112593a86a<->200600a0b8174edd vmhba0:0:0 Standby preferred  
FC 2:2.0 210000112593a86a<->200700a0b8174edc vmhba0:1:0 On active  
FC 2:2.1 210000112593a86b<->200600a0b8174edc vmhba1:0:0 Standby  
FC 2:2.1 210000112593a86b<->200700a0b8174edd vmhba1:1:0 On
```

Console – Multipath Analysis (HBA Failover)

- esxcfg-mpath -l
- Console View of the HBA failover event



```
root@giza:~  
[root@giza root]# esxcfg-mpath -l  
Disk vmhba0:0:0 /dev/sda (512MB) has 4 paths and policy of Most Recently Used  
FC 2:2.0 210000112593a86a<->200600a0b8174edd vmhba0:0:0 On active preferred  
FC 2:2.0 210000112593a86a<->200700a0b8174edc vmhba0:1:0 Standby  
FC 2:2.1 210000112593a86b<->200600a0b8174edc vmhba1:0:0 On  
FC 2:2.1 210000112593a86b<->200700a0b8174edd vmhba1:1:0 Standby  
  
root@giza:~  
[root@giza root]# esxcfg-mpath -l  
Disk vmhba0:0:0 /dev/sda (512MB) has 4 paths and policy of Most Recently Used  
FC 2:2.0 210000112593a86a<->200600a0b8174edd vmhba0:0:0 On preferred  
FC 2:2.0 210000112593a86a<->200700a0b8174edc vmhba0:1:0 Standby  
FC 2:2.1 210000112593a86b<->200600a0b8174edc vmhba1:0:0 On active  
FC 2:2.1 210000112593a86b<->200700a0b8174edd vmhba1:1:0 Standby
```


VMFS-3

- Journaling
- Hierarchical
- Metadata
- LVM
- Snapshot handling
- VMotion with Storage Migration
- New Virtual Disk Types

Journaling

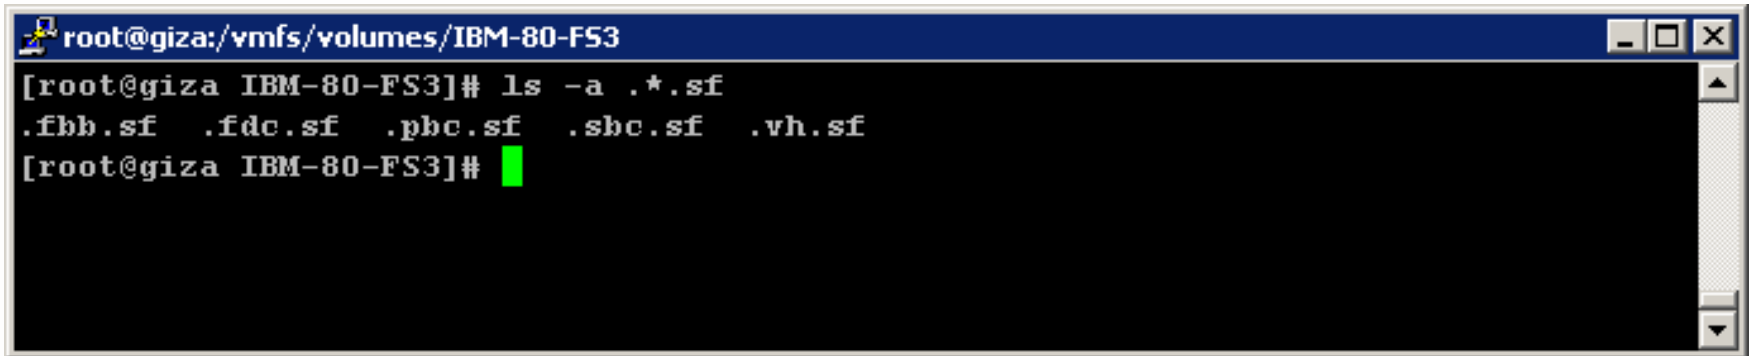
- VMFS-3 is now a journal based file system
- Better recovery from crashes
- Online File System Check

Hierarchical Structure

- VMFS-3 volumes support directory structure
- Virtual Machine files stored in the same directory
 - Configuration file (VMX)
 - VM Swap File
 - Virtual Disks
 - Logs
 - Monitor dump

VMFS-3 Metadata

- Copies of metadata stored at the root level of each volume
- Backup these files on a regular basis
- Run vm-support on a regular basis
 - Collects Partition Table info

A terminal window with a blue title bar containing the text 'root@giza:/vmfs/volumes/IBM-80-F53'. The terminal shows a command prompt '[root@giza IBM-80-F53]# ls -a *.sf' followed by the output '.fbh.sf .fdc.sf .phc.sf .shc.sf .vh.sf'. The prompt is followed by a green cursor.

```
root@giza:/vmfs/volumes/IBM-80-F53
[root@giza IBM-80-F53]# ls -a *.sf
.fbh.sf .fdc.sf .phc.sf .shc.sf .vh.sf
[root@giza IBM-80-F53]#
```

LVM - VMware Logical Volume Manager

- Spanned VMFS volumes
 - Volume still valid with missing extent(s)
- Snapshot LUN handling
- VMFS file system modules
 - vmfs2
 - vmfs3
 - fsaux

VMFS3 Volume UUID

- A volume UUID looks like this:
42263200-74382e04-b9bf-009c06010000
- Also Known As “Signature”
- Created as a response to non-persistent vmhba names
- Volume names are supposed to be used instead of UUIDs
- Not to be confused with host UUIDs

Snapshot LUN Handling

- Basic rules:
 - A LUN shared by multiple hosts MUST be presented with the SAME LUN ID to all hosts
 - A VMFS volume signature is associated with the LUN ID and the UUID among a few other elements (this may change at a later release)
- AX100 and IBM ESS (Shark) may not meet these rules.
- Symptoms
 - Volume hidden
 - Volume identified as a snapshot

Snapshot LUN Handling

■ Sample log entries

Jul 18 10:58:31 ServerB vmkernel: 0:14:17:59.787 cpu13:1046)LVM: 5670: Device vmhba1:0:5:1 is a snapshot:
Jul 18 10:58:31 ServerB vmkernel: 0:14:17:59.787 cpu13:1046)LVM: 5676: disk ID: <type 2, len 22, lun 5,
devType 0, scsi 5, h(id) 10179760818951437974>
Jul 18 10:58:31 ServerB vmkernel: 0:14:17:59.787 cpu13:1046)LVM: 5678: m/d disk ID: <type 2, len 22, lun 1,
devType 0, scsi 5, h(id) 10179760818951437974>
Jul 18 10:58:31 ServerB vmkernel: 0:14:17:59.790 cpu13:1046)LVM: 5670: Device vmhba1:0:6:1 is a snapshot:
Jul 18 10:58:31 ServerB vmkernel: 0:14:17:59.790 cpu13:1046)LVM: 5676: disk ID: <type 2, len 22, lun 6,
devType 0, scsi 5, h(id) 11552037668126695191>
Jul 18 10:58:31 ServerB vmkernel: 0:14:17:59.790 cpu13:1046)LVM: 5678: m/d disk ID: <type 2, len 22, lun 2,
devType 0, scsi 5, h(id) 11552037668126695191>
Jul 18 10:58:31 ServerB vmkernel: 0:14:17:59.794 cpu13:1046)LVM: 5670: Device vmhba1:0:7:1 is a snapshot:
Jul 18 10:58:31 ServerB vmkernel: 0:14:17:59.794 cpu13:1046)LVM: 5676: disk ID: <type 2, len 22, lun 7,
devType 0, scsi 5, h(id) 13372428508588014685>
Jul 18 10:58:31 ServerB vmkernel: 0:14:17:59.794 cpu13:1046)LVM: 5678: m/d disk ID: <type 2, len 22, lun 3,
devType 0, scsi 5, h(id) 13372428508588014685>

Correcting Hidden Volumes

- See KB 6482648 at: <http://kb.vmware.com/kb/6482648>
- Correct LUN ID so that all hosts “see it” with the same LUN number
 - Symmetrix: Assign the LUNs to the FAs using the same LUN number
 - Clariion: Assign the LUNs to the same Storage Group
 - IBM DS4000/FAStT family: Assign the LUNs to the same Host Group (Logical Partitioning option enabled)
 - Other Arrays: use equivalent features to present the LUN with the same Host LUN ID to all hosts sharing it
- If the above is not possible:
 - Set advanced option LVM.DisallowSnapshotLUN to “0” then rescan
 - **DO NOT** present actual snapshot LUNs to that server with this option setting in place. Data Corruption may result

AutoResignature

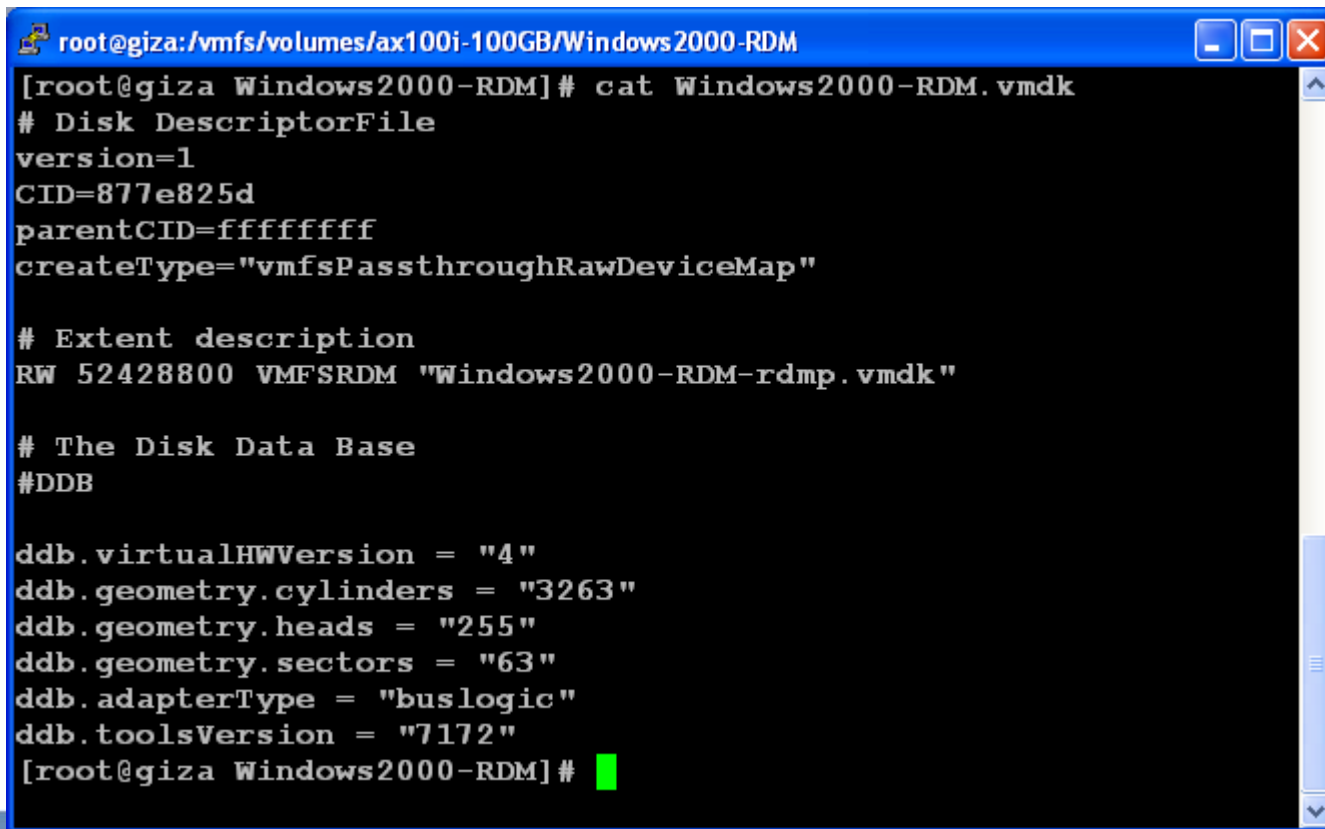
- Auto Resignature feature is disabled by default
- To enable it set advanced option LVM.EnableResignature to “1” then rescan
- This will apply to ALL volumes visible to this host.
- This will affect ALL hosts sharing this volume
- All VMs configurations need to be edited to reflect the new volume ID
- This will be addressed in a future release

VMFS Kernel Modules

- VMFS2
 - Loaded for “Read-Only” by default for all VMFS2 volumes
 - Can be unloaded after all VMFS2 volumes have been upgraded
- VMFS3
 - All operations related to VMFS3 volumes
- FSAUX
 - Auxillary File System functions
 - Upgrading VMFS2 to VMFS3
 - Other functions internal fo VMware

VMFS3 – Virtual Disk Structure Changes

- Virtual machine virtual disk now has a descriptor file which is in ASCII format
- ESX 2.x contained this information in last 512 bytes of vmdk

A terminal window with a blue title bar showing the command prompt and the output of the 'cat' command. The output is an ASCII descriptor file for a VMFS3 virtual disk.

```
root@giza:/vmfs/volumes/ax100i-100GB/Windows2000-RDM
[root@giza Windows2000-RDM]# cat Windows2000-RDM.vmdk
# Disk DescriptorFile
version=1
CID=877e825d
parentCID=ffffffff
createType="vmfsPassthroughRawDeviceMap"

# Extent description
RW 52428800 VMFSRDM "Windows2000-RDM-rdmp.vmdk"

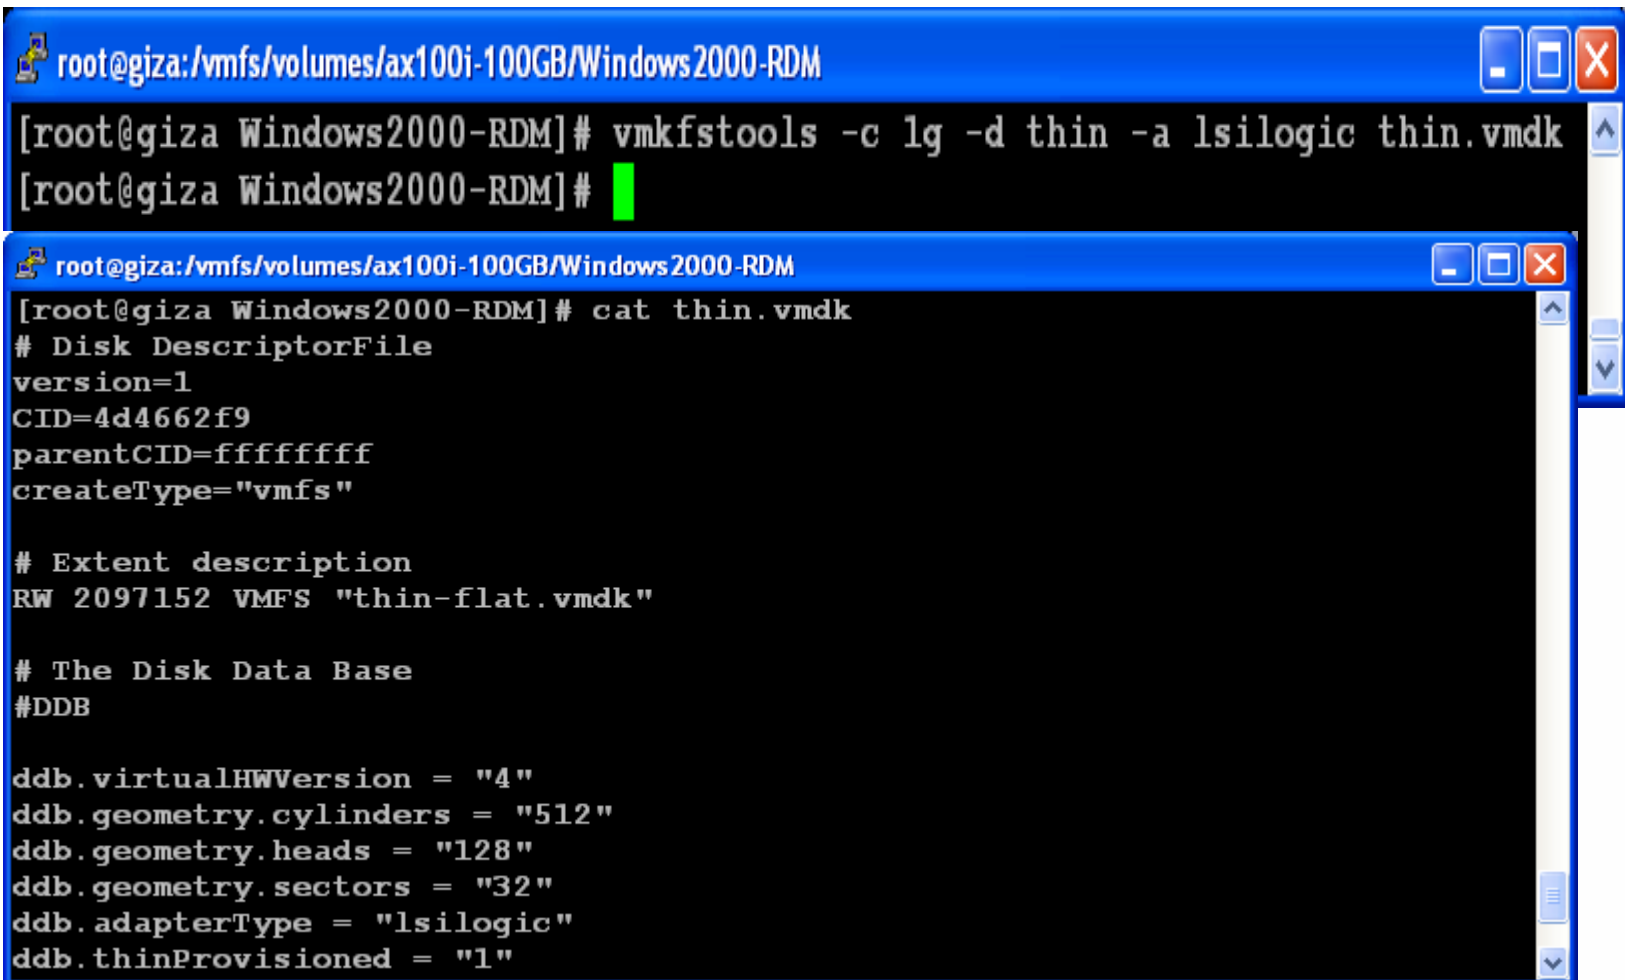
# The Disk Data Base
#DDB

ddb.virtualHWVersion = "4"
ddb.geometry.cylinders = "3263"
ddb.geometry.heads = "255"
ddb.geometry.sectors = "63"
ddb.adapterType = "buslogic"
ddb.toolsVersion = "7172"
[root@giza Windows2000-RDM]#
```

VMFS3 – Virtual Disk Structure Changes

Format	Example
<code>createType="<type> # Extent description RW <size in 512bytes blocks> <extent>"</code>	<code>createType="vmfsPassthroughRawDeviceMap" # Extent description RW 52428800 VMFSRDM "Win2K-rdmp.vmdk"</code>
Virtual Disk Type	Extent
<code>vmfsPassthroughRawDeviceMap</code>	VMFSRDM "<name>-rdmp.vmdk"
<code>vmfsRawDeviceMap</code>	VMFSRDM "<name>-rdm.vmdk"
<code>vmfs (Thin / Thick)</code>	VMFS "<name>-flat.vmdk"

Sample Virtual Disk



The image displays two terminal windows from a VMware environment. The top window shows the command `vmkfstools -c 1g -d thin -a lsilogic thin.vmdk` being executed to create a 1GB thin-provisioned virtual disk using the lsilogic adapter. The bottom window shows the command `cat thin.vmdk` being used to inspect the disk's descriptor file, which contains metadata such as version, CID, parent CID, and geometry.

```
root@giza:/vmfs/volumes/ax100i-100GB/Windows2000-RDM
[root@giza Windows2000-RDM]# vmkfstools -c 1g -d thin -a lsilogic thin.vmdk
[root@giza Windows2000-RDM]#

root@giza:/vmfs/volumes/ax100i-100GB/Windows2000-RDM
[root@giza Windows2000-RDM]# cat thin.vmdk
# Disk DescriptorFile
version=1
CID=4d4662f9
parentCID=ffffffff
createType="vmfs"

# Extent description
RW 2097152 VMFS "thin-flat.vmdk"

# The Disk Data Base
#DDB

ddb.virtualHWVersion = "4"
ddb.geometry.cylinders = "512"
ddb.geometry.heads = "128"
ddb.geometry.sectors = "32"
ddb.adapterType = "lsilogic"
ddb.thinProvisioned = "1"
```

Question?

iSCSI Storage

- Block-level I/O over TCP/IP using SCSI-3 protocols
- Pros
 - Direct access to random blocks on disks or LUNs for performance
 - Block and file traffic on single Ethernet infrastructure
 - Enabled by adoption of Gigabit Ethernet
 - Can use normal IP authentication, encryption, routing features
- Cons
 - Protocol overhead added by TCP/IP
 - Small packets, connection-oriented chatter
 - Gigabit traffic can saturate a 2.4-GHz CPU
 - Slower than Fibre Channel (for now)
- TCP Offload Engines (TOE)
 - Reduce CPU overhead, but expensive
 - Plain Gigabit NICs catching up with better drivers

iSCSI Requirements


- Install a VMotion and IP Storage licenses on VC
- Host and storage ports on the same subnet
- A supported Ethernet NIC OR
- A [QLogic 4050](#) or other card on the HCL (Experimental)
- A supported iSCSI Storage

iSCSI Configuration

- Authentication
 - CHAP
 - Per HBA or per target
 - None
- Target discovery
 - Static targets
 - SendTargets discovery (Dynamic)
 - No SLP discovery
 - No iSNS discovery
- Administration
 - SNIA-based IMA library
 - Configured through VMware Virtual Infrastructure Client
 - Common approach for all initiators
 - No need for vendor-specific tools

iSCSI – Configuring HW Initiator

- Select: Configuration – Storage Adapters
- Select the HBA and click Properties (in details pane)
- Click “Configure”



iSCSI – Configuring SW Initiator


- COS port group (for Authentication)
- VMKernel Port Group (for Data Traffic)
- Initiator IQN
- Security
- Discovery
- Multipathing
 - Only with both ports on the same switch and subnet

iSCSI – SW Initiator

Configuring Network

Virtual Switch: vSwitch1

[Remove...](#) [Properties..](#)



iSCSI – SW Initiator

- Enable software iSCSI client in firewall

The screenshot shows the VMware ESX Server 3.0.0 configuration interface. The top navigation bar includes tabs for Summary, Virtual Machines, Resource Allocation, Performance, Configuration (highlighted with a red box), Users & Groups, Events, Permissions, and Maps. The left sidebar has two main sections: Hardware and Software. Under Hardware, there are links for Processors, Memory, Storage (SCSI, SAN, and NFS), Networking, Storage Adapters, and Network Adapters. Under Software, there are links for Licensed Features, DNS and Routing, Virtual Machine Startup/Shutdown, Security Profile (highlighted with a red box), System Resource Allocation, and Advanced Settings. The main content area is titled 'Security Profile' and contains a 'Firewall' section. The Firewall section has a 'Properties...' link (highlighted with a red box) and lists incoming and outgoing connections. The incoming connections list includes CIM SLP (427 (UDP,TCP)), CIM Secure Server (5989 (TCP)), SSH Server (22 (TCP)), EMC AAM Client (2050-5000,8042-8045 (TCP,UDP)), and CIM Server (5988 (TCP)). The outgoing connections list includes CIM SLP (427 (UDP,TCP)), VMware License Client (27000,27010 (TCP)), VMware VirtualCenter Agent (902 (UDP)), and EMC AAM Client (2050-5000,8042-8045 (TCP,UDP)). Below the Firewall section is the 'Virtual Machine Delegate' section, which includes an 'Edit...' link and a description: 'Read and write to virtual machine files using these credentials. The host must be in maintenance mode to change this setting.' The 'User Name:' field is set to 'root'.

giza.vmware.com VMware ESX Server, 3.0.0, 23447

Summary Virtual Machines Resource Allocation Performance **Configuration** Users & Groups Events Permissions Maps

Hardware

- Processors
- Memory
- Storage (SCSI, SAN, and NFS)
- Networking
- Storage Adapters
- Network Adapters

Software

- Licensed Features
- DNS and Routing
- Virtual Machine Startup/Shutdown
- Security Profile**
- System Resource Allocation
- Advanced Settings

Security Profile

Firewall [Properties...](#)

Incoming Connections

CIM SLP	427 (UDP,TCP)
CIM Secure Server	5989 (TCP)
SSH Server	22 (TCP)
EMC AAM Client	2050-5000,8042-8045 (TCP,UDP)
CIM Server	5988 (TCP)

Outgoing Connections

CIM SLP	427 (UDP,TCP)
VMware License Client	27000,27010 (TCP)
VMware VirtualCenter Agent	902 (UDP)
EMC AAM Client	2050-5000,8042-8045 (TCP,UDP)

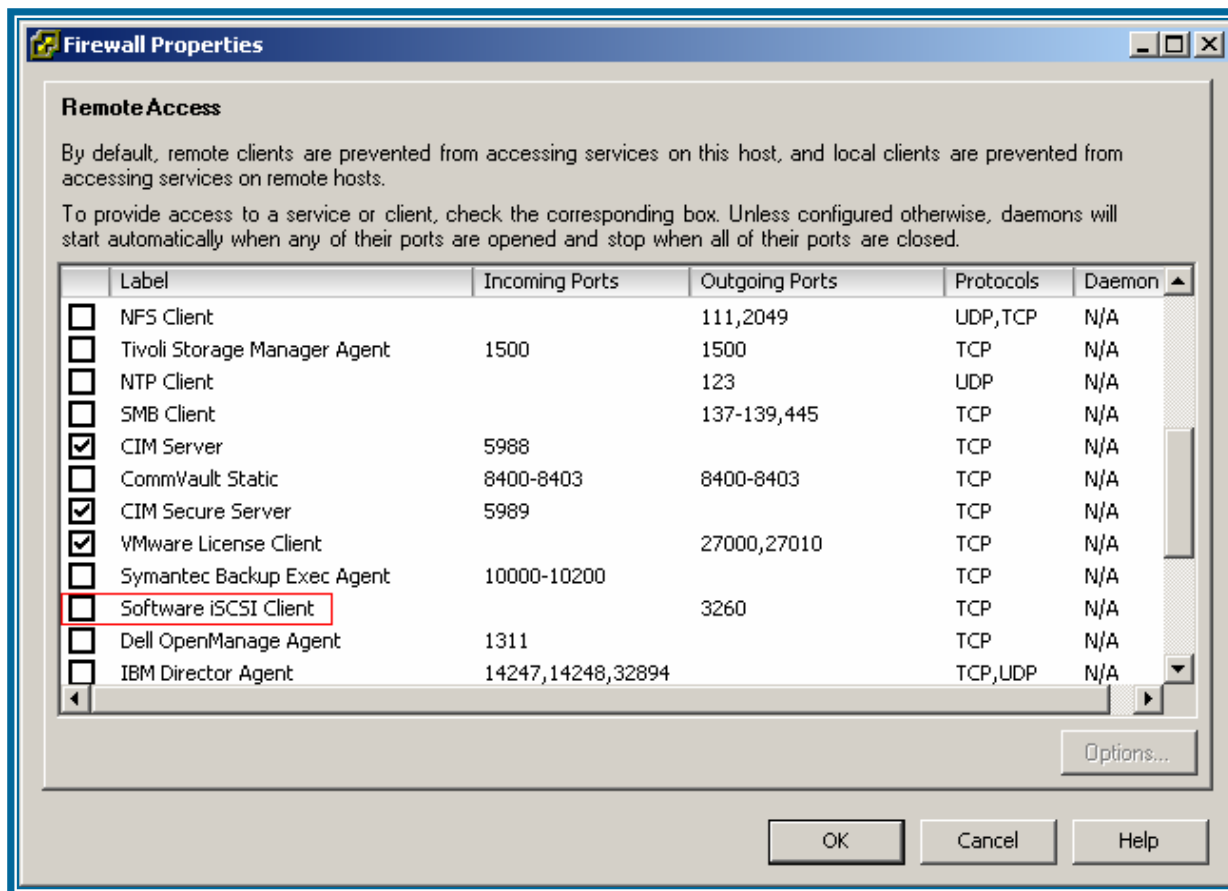
Virtual Machine Delegate [Edit...](#)

Read and write to virtual machine files using these credentials. The host must be in maintenance mode to change this setting.

User Name: root

iSCSI – SW Initiator

- Enable “Software iSCSI Client” in firewall



iSCSI – SW Initiator

- Enable the Software Initiator
 - Storage Adapters – Select “iSCSI” Software Adapter

The screenshot shows the VMware ESX Server 3.0.0 configuration interface. The top navigation bar includes tabs for Summary, Virtual Machines, Resource Allocation, Performance, Configuration, Tasks & Events, Alarms, Permissions, and Maps. The left sidebar contains a tree view with categories: Hardware (Processors, Memory, Storage (SCSI, SAN, and NFS), Networking, Storage Adapters, Network Adapters) and Software (Licensed Features, DNS and Routing, Virtual Machine Startup/Shutdown, Security Profile, System Resource Allocation, Advanced Settings). The main content area is titled 'Storage Adapters' and includes a 'Rescan' button. A table lists the storage adapters:

Device	Type	SAN Identifier
vmhba1	Fibre Channel SCSI	21:00:00:11:2
QLA231x/2340		
vmhba0	Fibre Channel SCSI	21:00:00:11:2
iSCSI Software Adapter		
iSCSI Software Adapter	iSCSI	

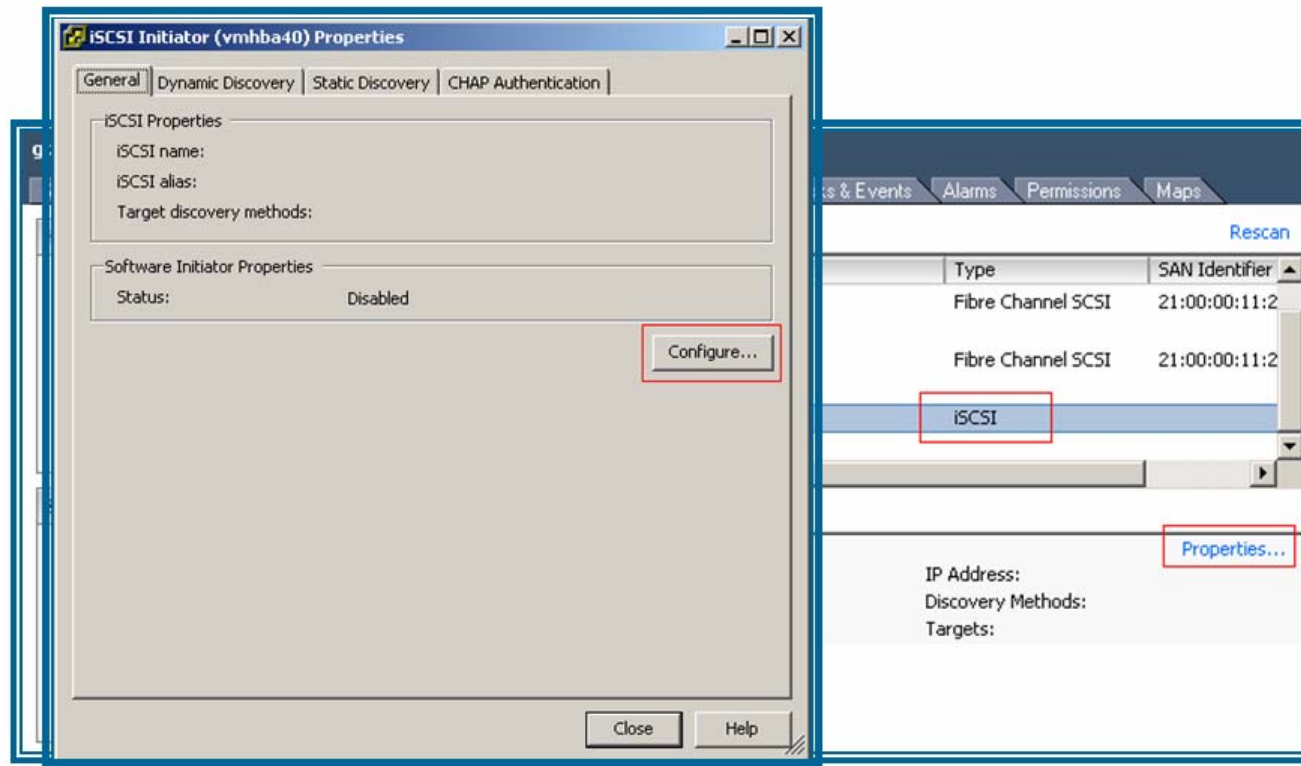
Below the table is a 'Details' section with a 'Properties...' button. The details are organized into two columns:

Details	
Model:	IP Address:
iSCSI Name:	Discovery Methods:
iSCSI Alias:	Targets:

iSCSI – SW Initiator

- Enable the Software Initiator
 - > Storage Adapters – Select “iSCSI” Software Adapter

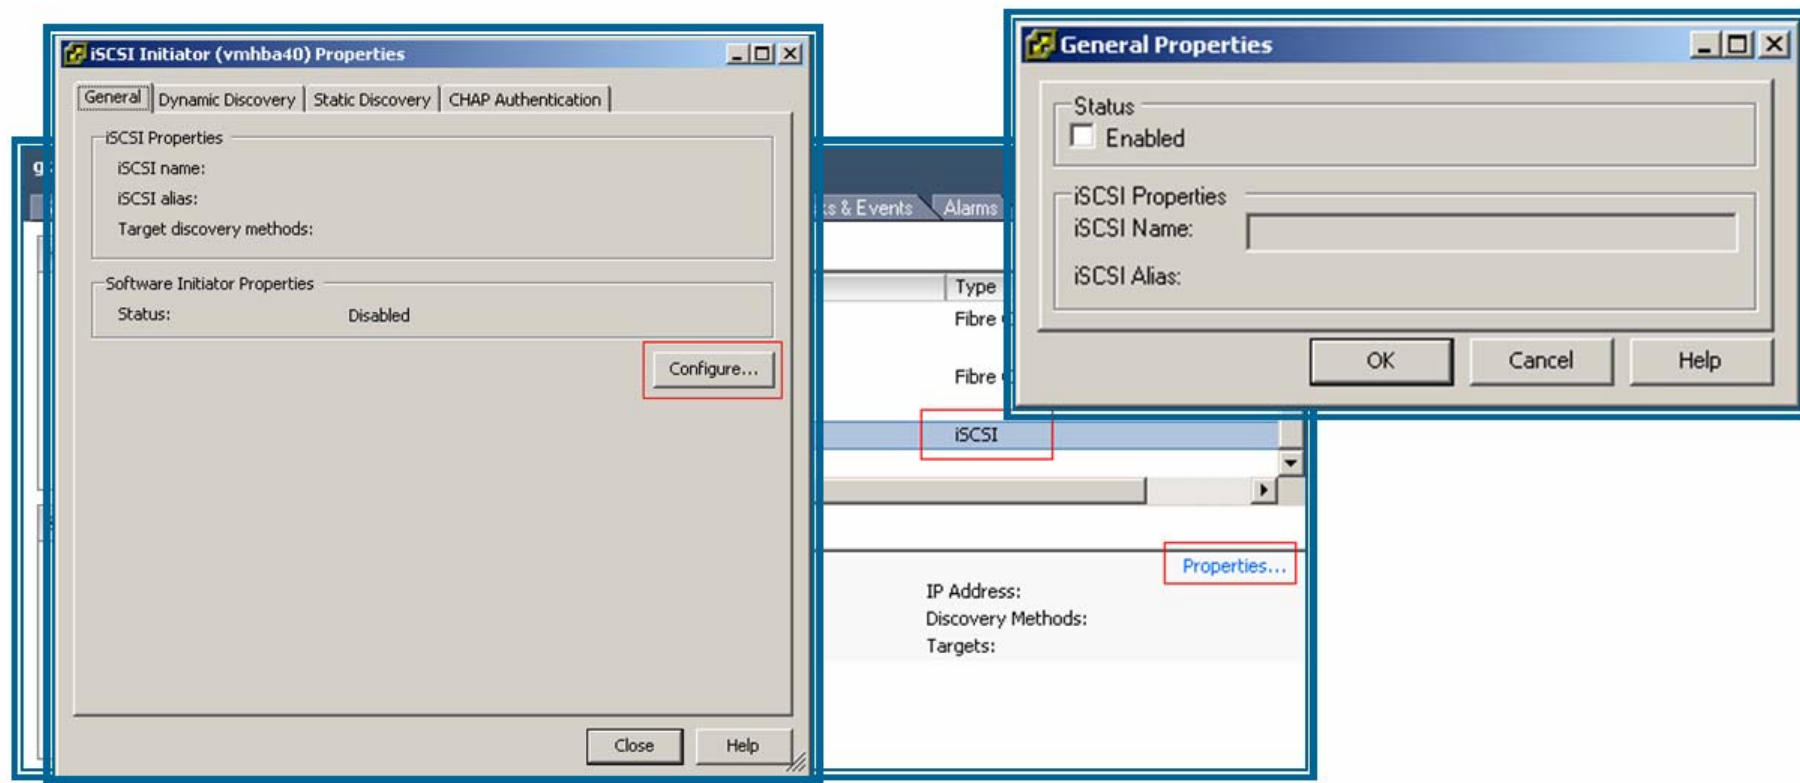
> Select “properties”



iSCSI – SW Initiator

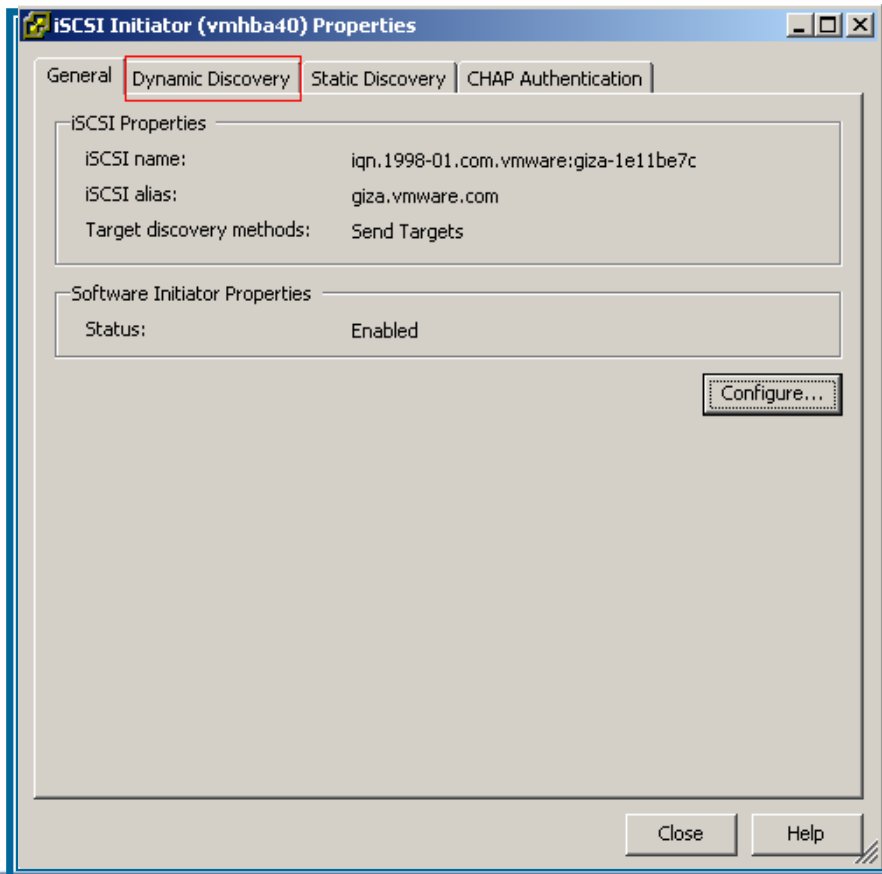
- Enable the Software Initiator
 - > Storage Adapters – Select “iSCSI” Software Adapter

- > Select “properties”
- > Select “Configure”
- > Click “Enabled” then Click “OK”



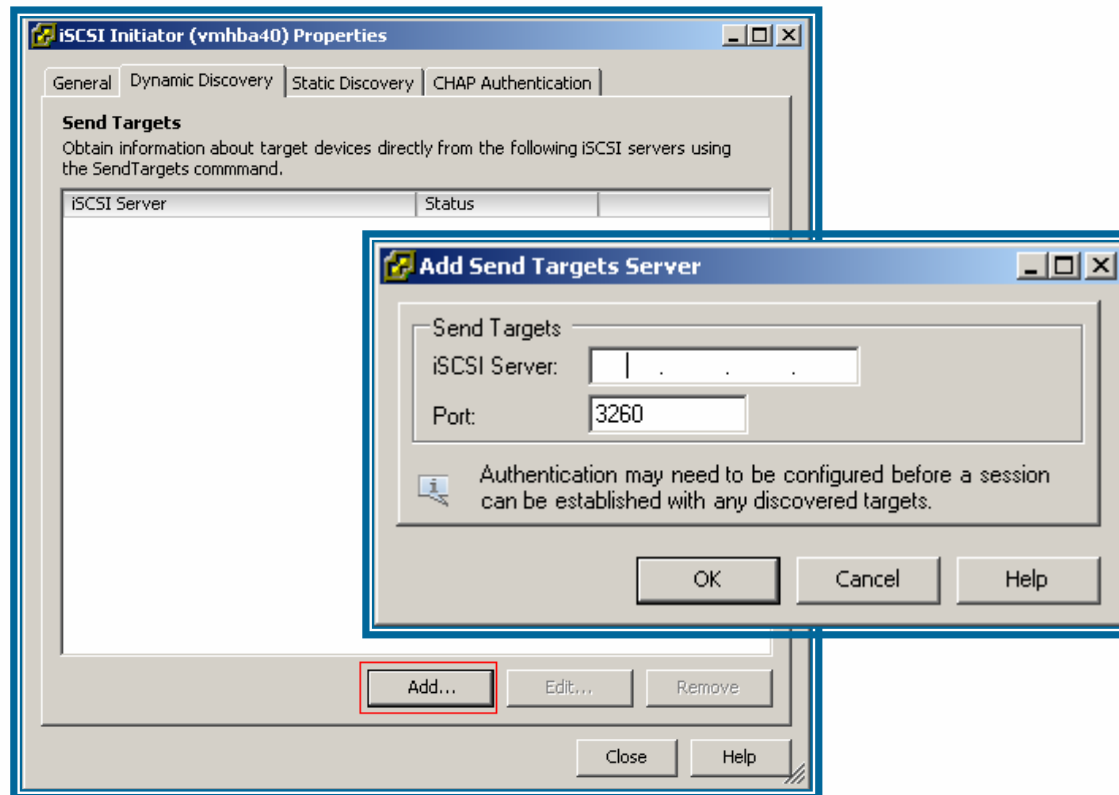
iSCSI – SW Initiator

- iSCSI properties get filled automatically
- Select “Dynamic Discovery”




iSCSI – Dynamic Discovery

- Select “Add” then input the iSCSI Server’s IP address
- Repeat for each port on the iSCSI storage



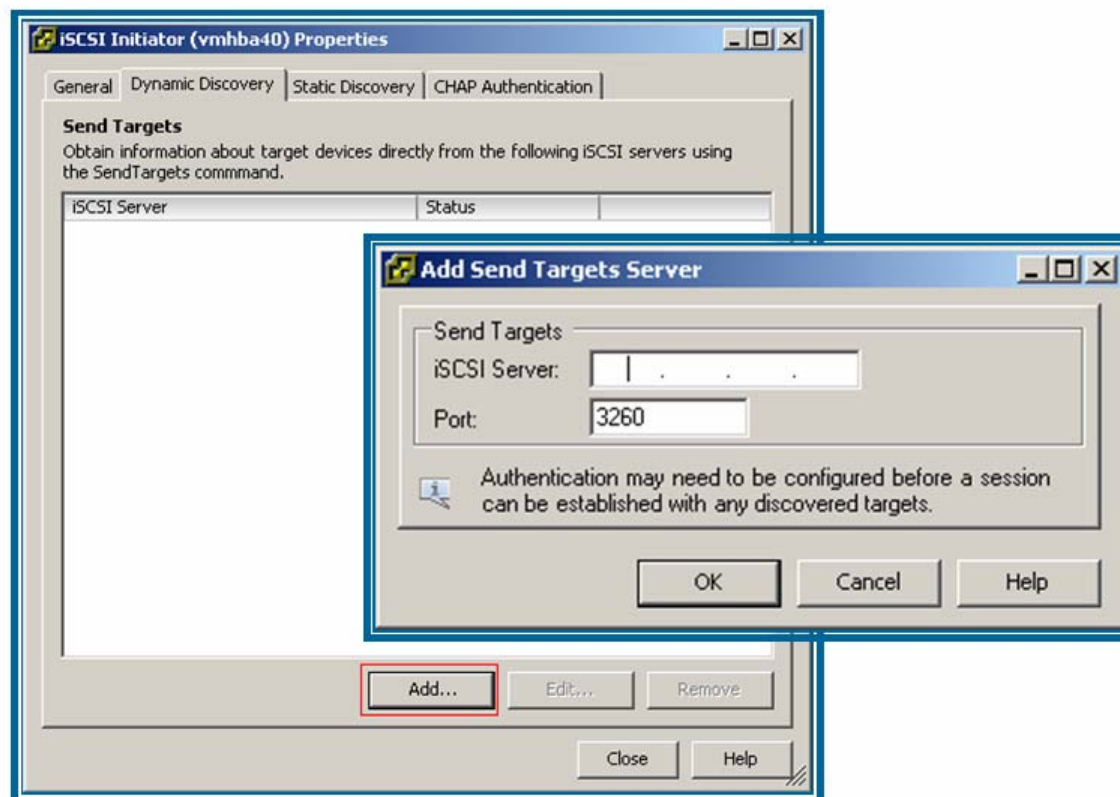
iSCSI – Dynamic Discovery

- Select “Add” then input the iSCSI Server’s IP address



iSCSI – Dynamic Discovery

- Select “Add” then input the iSCSI Server’s IP address
- Repeat for each port on the iSCSI storage

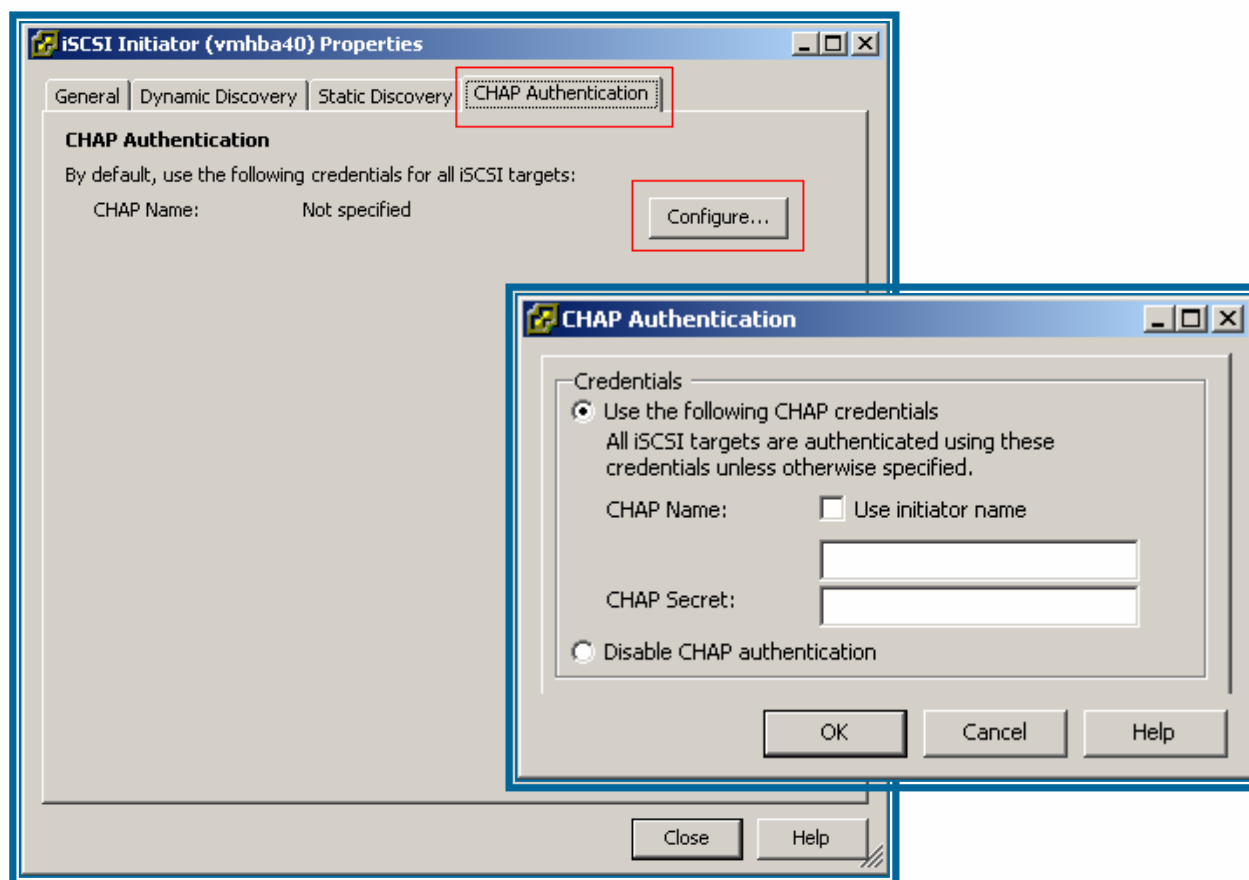


iSCSI – Static Discovery

- Same as Dynamic but select the “Static” tab instead
- Works with HW initiators ONLY

iSCSI – Authentication

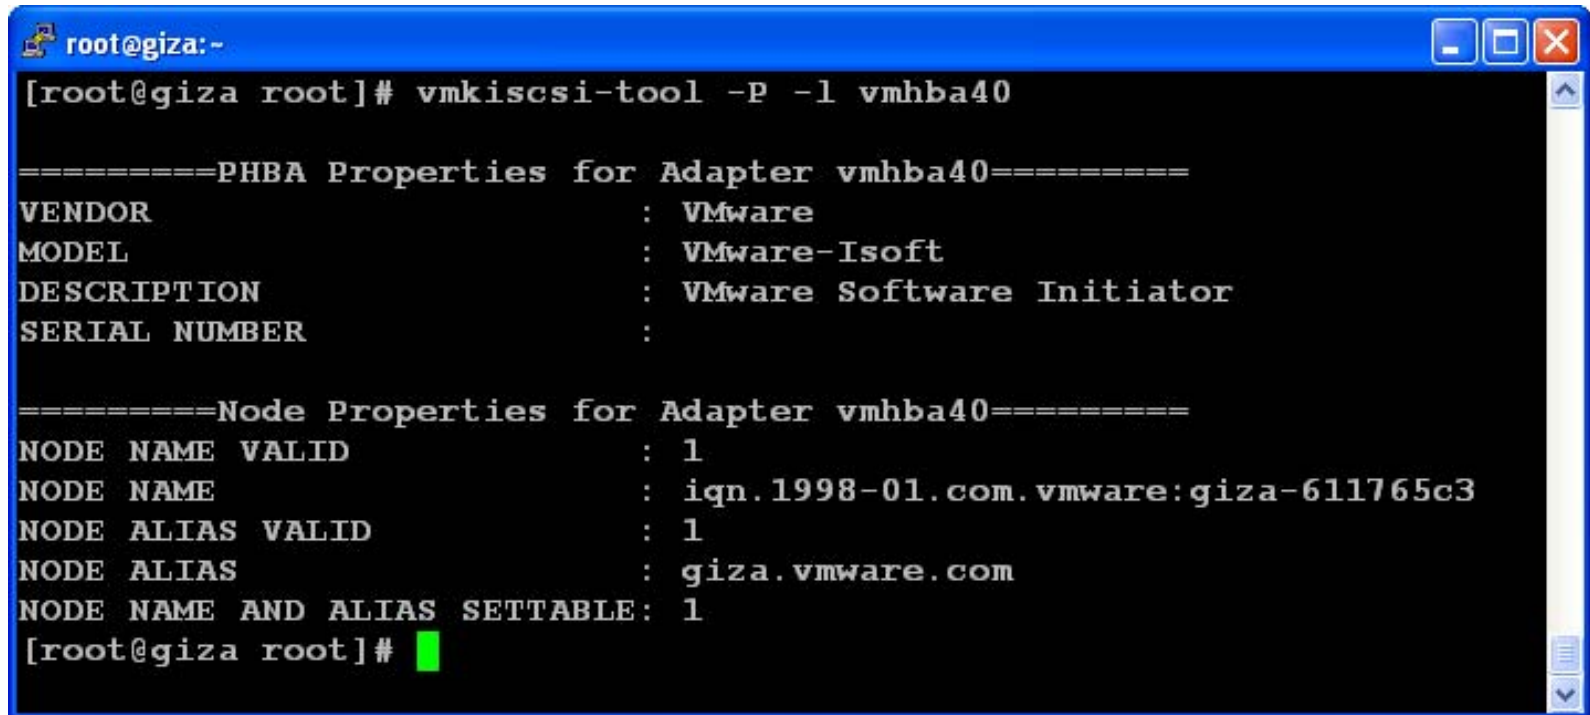
- Select the “CHAP Authentication” Tab then “Configure”



iSCSI – Troubleshooting

- vmkiscsi-tool

- Usage: vmkiscsi-tool <command> <subcommand> adapterName

A terminal window titled 'root@giza:-' with standard window controls. The command '[root@giza root]# vmkiscsi-tool -P -l vmhba40' has been executed. The output displays two sections of properties for the adapter 'vmhba40'. The first section, 'PHBA Properties', lists Vendor as 'VMware', Model as 'VMware-Issoft', Description as 'VMware Software Initiator', and Serial Number as empty. The second section, 'Node Properties', lists Node Name Valid as 1, Node Name as 'iqn.1998-01.com.vmware:giza-611765c3', Node Alias Valid as 1, Node Alias as 'giza.vmware.com', and Node Name AND Alias Settable as 1. The prompt '[root@giza root]#' is followed by a green cursor.

```
root@giza:-  
[root@giza root]# vmkiscsi-tool -P -l vmhba40  
  
=====PHBA Properties for Adapter vmhba40=====  
VENDOR                : VMware  
MODEL                 : VMware-Issoft  
DESCRIPTION           : VMware Software Initiator  
SERIAL NUMBER         :  
  
=====Node Properties for Adapter vmhba40=====  
NODE NAME VALID       : 1  
NODE NAME             : iqn.1998-01.com.vmware:giza-611765c3  
NODE ALIAS VALID      : 1  
NODE ALIAS            : giza.vmware.com  
NODE NAME AND ALIAS SETTABLE: 1  
[root@giza root]#
```

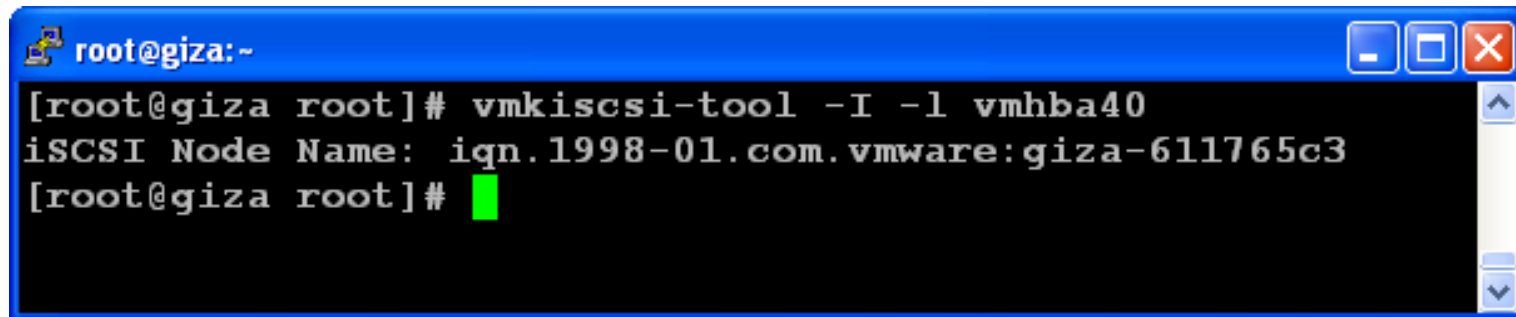
shows the Initiators properties

iSCSI – Troubleshooting

```
root@giza:~  
[root@giza root]# vmkiscsi-tool -D -l vmhba40  
  
=====Discovery Properties for Adapter vmhba40=====  
iSnsDiscoverySettable      : 0  
iSnsDiscoveryEnabled       : 0  
staticDiscoverySettable    : 0  
staticDiscoveryEnabled     : 0  
sendTargetsDiscoverySettable : 0  
sendTargetsDiscoveryEnabled : 1  
slpDiscoverySettable       : 0  
  
Discovery Status: Timed out. Displayed information may be incomplete.  
DISCOVERY ADDRESS      : 10.16.92.36:3260  
DISCOVERY ADDRESS      : 10.16.112.76:3260  
DISCOVERY ADDRESS      : 10.16.92.40:3260  
  
Static Discovery not supported for this adapter  
[root@giza root]#
```

- Shows Discovery Properties
- If this were a Hardware Initiator, Static Targets would have been listed

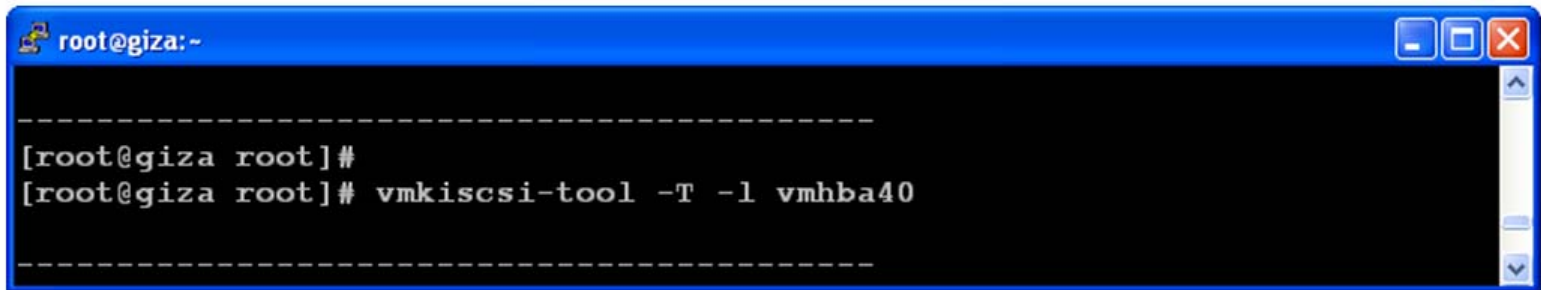
iSCSI – Troubleshooting

A terminal window with a blue title bar and standard window controls. The title bar shows 'root@giza:~'. The terminal text shows the command 'vmkiscsi-tool -I -l vmhba40' being executed, followed by the output 'iSCSI Node Name: iqn.1998-01.com.vmware:giza-611765c3'. A green cursor is visible on the line following the output.

```
root@giza:~  
[root@giza root]# vmkiscsi-tool -I -l vmhba40  
iSCSI Node Name: iqn.1998-01.com.vmware:giza-611765c3  
[root@giza root]#
```

- Shows iSCSI Node Name (IQN)

iSCSI – Troubleshooting



A terminal window titled 'root@giza:~' with standard window controls (minimize, maximize, close) and a scrollbar on the right. The terminal content shows a root shell prompt, followed by a command to list iSCSI targets using the 'vmkiscsi-tool' utility on the 'vmhba40' controller.

```
-----  
[root@giza root]#  
[root@giza root]# vmkiscsi-tool -T -l vmhba40  
-----
```

- Lists Targets

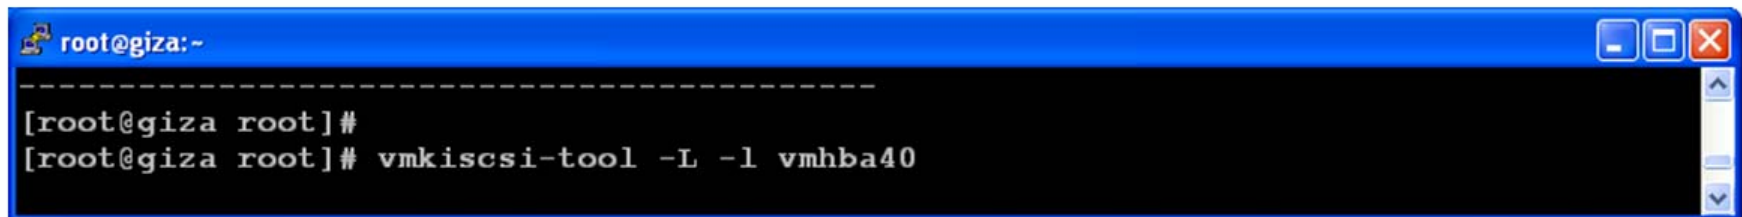
iSCSI – Troubleshooting



```
root@giza: ~  
root@giza: ~  
-----  
NAME                : iqn.1992-04.com.emc:ax.apm00054207419.b0  
ALIAS                : 7419.b0  
DISCOVERY METHOD FLAGS : 0  
SEND TARGETS DISCOVERY SETTABLE : 0  
SEND TARGETS DISCOVERY ENABLED : 0  
Portal 0             : 10.16.92.37:3260  
  
-----  
  
-----  
NAME                : iqn.1992-04.com.emc:ax.apm00054207419.a0  
ALIAS                : 7419.a0  
DISCOVERY METHOD FLAGS : 0  
SEND TARGETS DISCOVERY SETTABLE : 0  
SEND TARGETS DISCOVERY ENABLED : 0  
Portal 0             : 10.16.92.36:3260  
  
-----  
[root@giza root]#
```

- Lists Targets

iSCSI – Troubleshooting



```
root@giza:~  
-----  
[root@giza root]#  
[root@giza root]# vmkiscsi-tool -L -l vmhba40
```

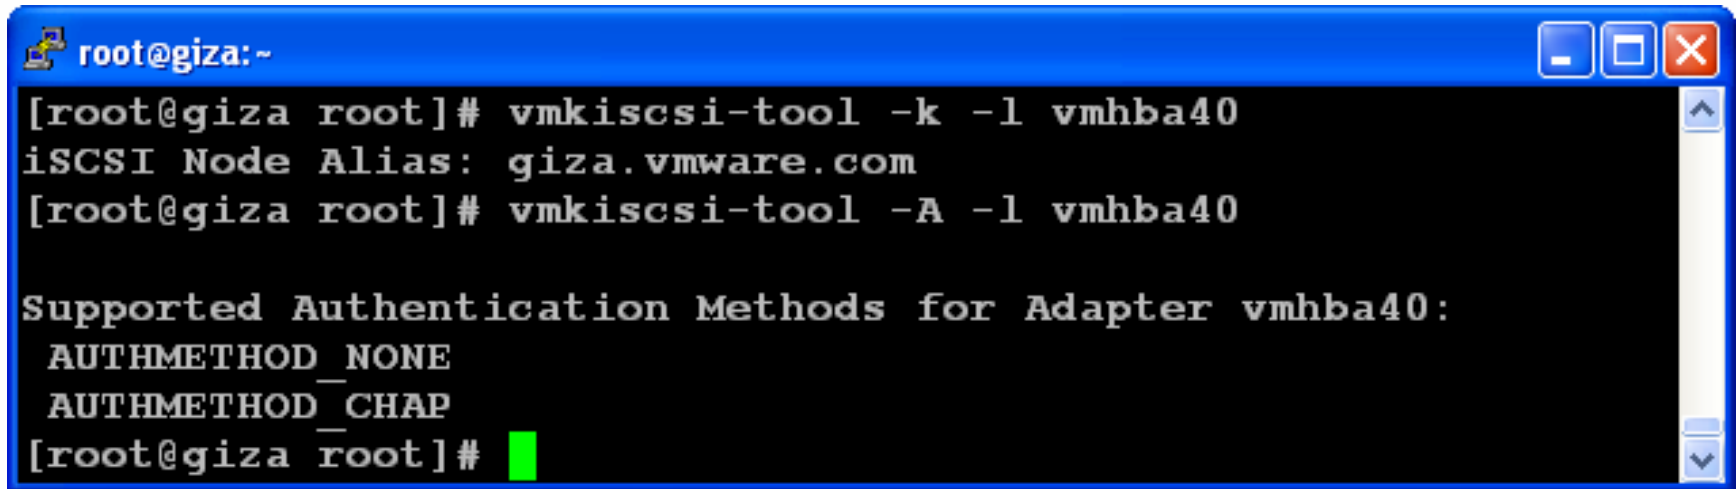
- Lists LUNs

iSCSI – Troubleshooting

```
root@giza:~  
-----  
Target iqn.1992-04.com.emc:ax.apm00054207419.b0:  
-----  
OS DEVICE NAME      : vmhba40:3:0  
BUS NUMBER          : 0  
TARGET ID           : 3  
LUN ID              : 0  
-----  
  
Target iqn.1992-04.com.emc:ax.apm00054207419.a0:  
-----  
OS DEVICE NAME      : vmhba40:4:0  
BUS NUMBER          : 0  
TARGET ID           : 4  
LUN ID              : 0  
-----  
  
OS DEVICE NAME      : vmhba40:4:1  
BUS NUMBER          : 0  
TARGET ID           : 4  
LUN ID              : 1
```

- Lists LUNs

iSCSI – Troubleshooting



```
root@giza:~  
[root@giza root]# vmkiscsi-tool -k -l vmhba40  
iSCSI Node Alias: giza.vmware.com  
[root@giza root]# vmkiscsi-tool -A -l vmhba40  
  
Supported Authentication Methods for Adapter vmhba40:  
AUTHMETHOD_NONE  
AUTHMETHOD_CHAP  
[root@giza root]#
```

- Shows Node Alias then shows Authentication Method

Q&A

NAS and NFS – Overview

- Use Network FileSystem Protocol
- NFS 3 TCP only (no UDP)
- Supported NAS filers only but also works with Linux NFS 3.x
- No CIFS/SMB yet
- NFS Naming Convention
 - nfs.remote.com:remote/filesystem
- Locking Handled by VMkernel
- Lease-based locks
- 8 NFS mounts by default. Can be increased to 32

NFS

Configuring datastore

- Configure the network

Add Network Wizard

VMkernel - Network Access
Use network labels to identify VMkernel connections while managing your hosts and datacenters.

[Connection Type](#)
Connection Settings
Summary

Port Group Properties

Network Label:

VLAN ID (Optional):

☐ Use this port group for VMotion

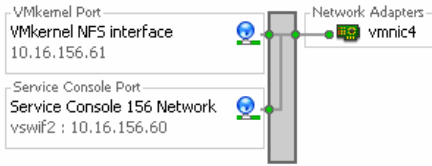
IP Settings

IP Address:

Subnet Mask:

VMkernel Default Gateway:

Preview:

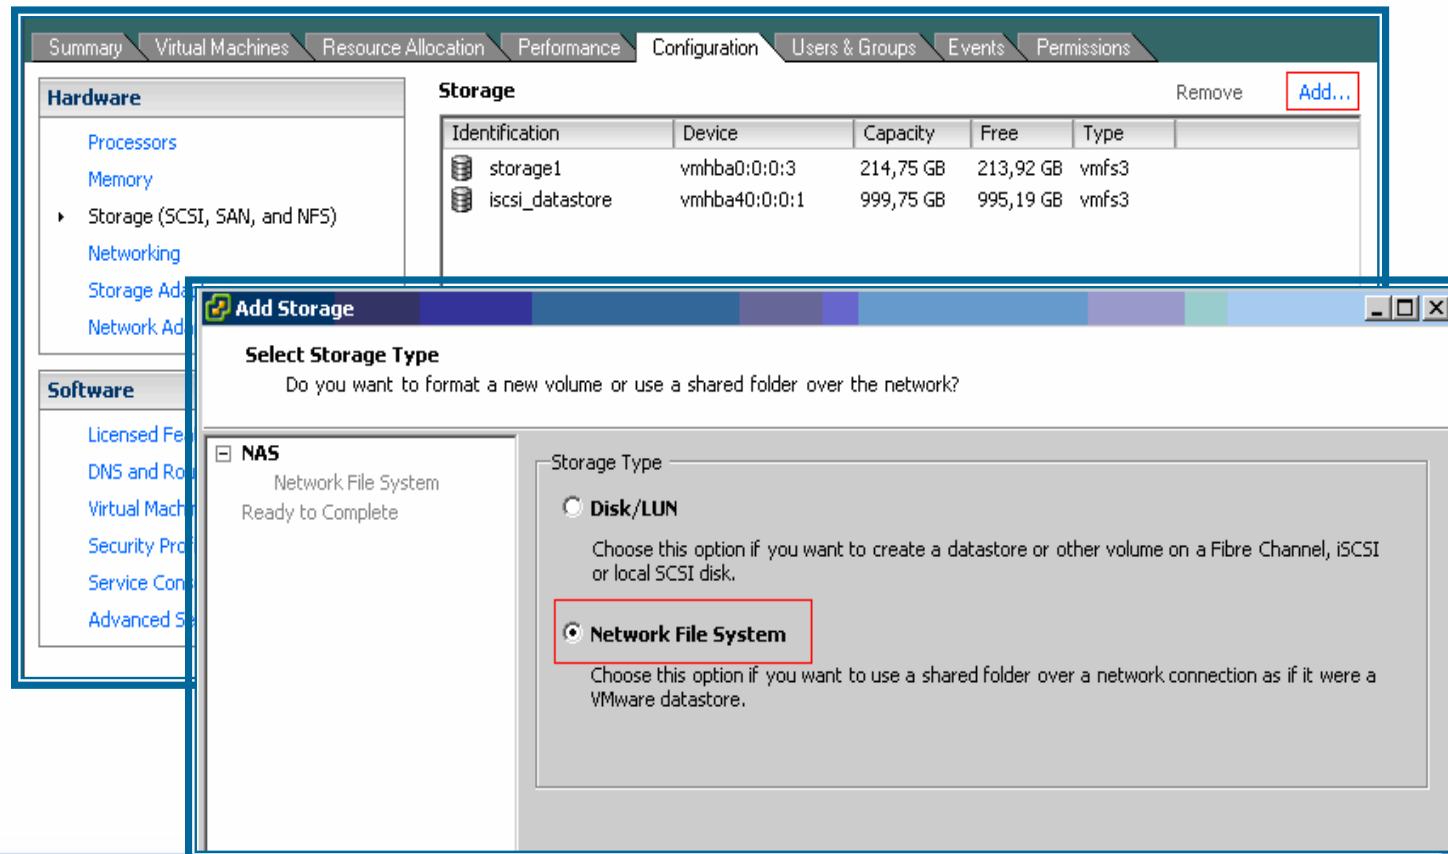


For more information about this wizard, see the [online documentation](#).

NFS

Configuring datastore (cont.)

- In “Storage” pane select “add” then “Network File System”



NFS

Configuring datastore (cont.)

- Fill in the NFS info

Add Storage

Locate Network File System
Which shared folder will be used as a VMware datastore?

NAS
Network File System
Ready to Complete

Properties

Server:
Examples: nas, nas.it.com or 192.168.0.1

Folder:
Example: /vols/vol0/datastore-001

Datastore Name

User interfaces display this name for human readers.

NFS

Configuring datastore (cont.)

- Now the storage shows the new DataStore

The screenshot shows the VMware vSphere Configuration tab for a virtual machine. The left sidebar contains a tree view with 'Hardware' and 'Software' sections. The 'Storage' option under 'Hardware' is selected. The main area displays the 'Storage' configuration, which includes a table of storage devices. The 'NFS datastore' is highlighted with a red box. Below the table, the 'Details' section for the 'NFS datastore' is visible, showing the server and folder paths.

Storage [Remove](#) [Add...](#)

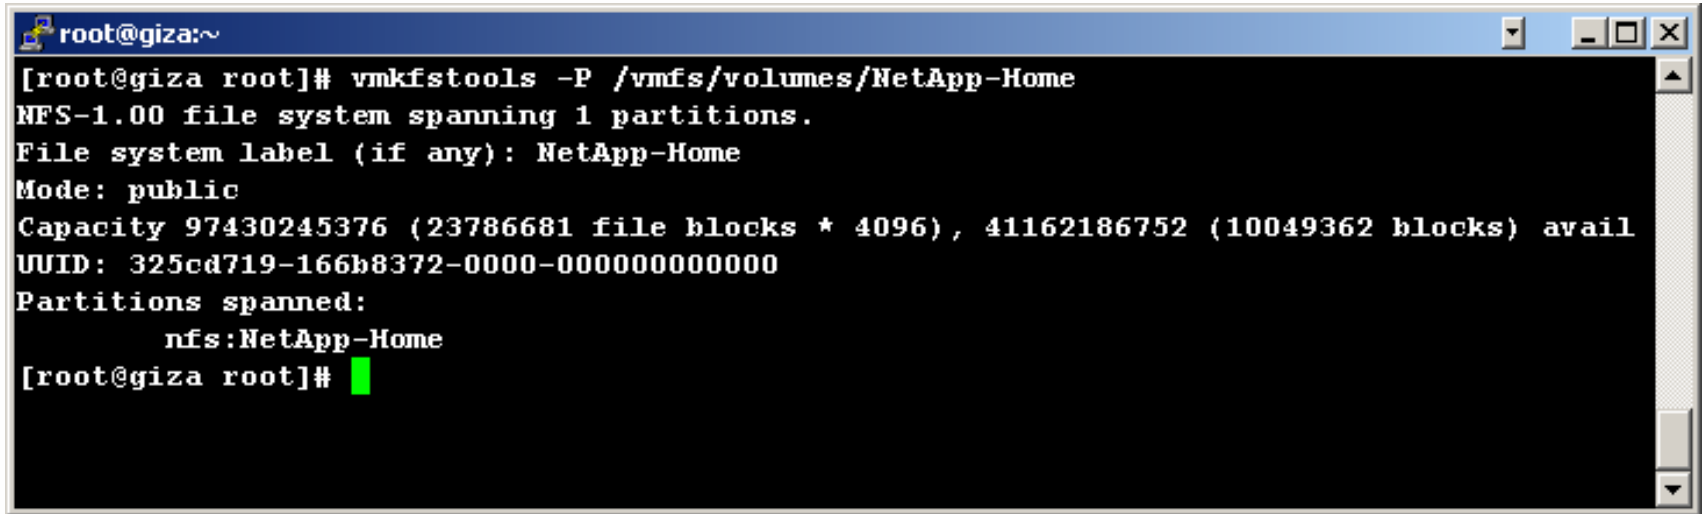
Identification	Device	Capacity	Free	Type
storage1	vmhba0:0:0:3	214,75 GB	213,92 GB	vmfs3
iscsi_datastore	vmhba40:0:0:1	999,75 GB	995,19 GB	vmfs3
NFS datastore	10.16.156.7;/med...	150,27 GB	9,05 GB	nfs

Details [Properties...](#)

NFS datastore

Server: 10.16.156.7
Folder: /media/usbdisk1/ISOS_REPO

NFS – Console View

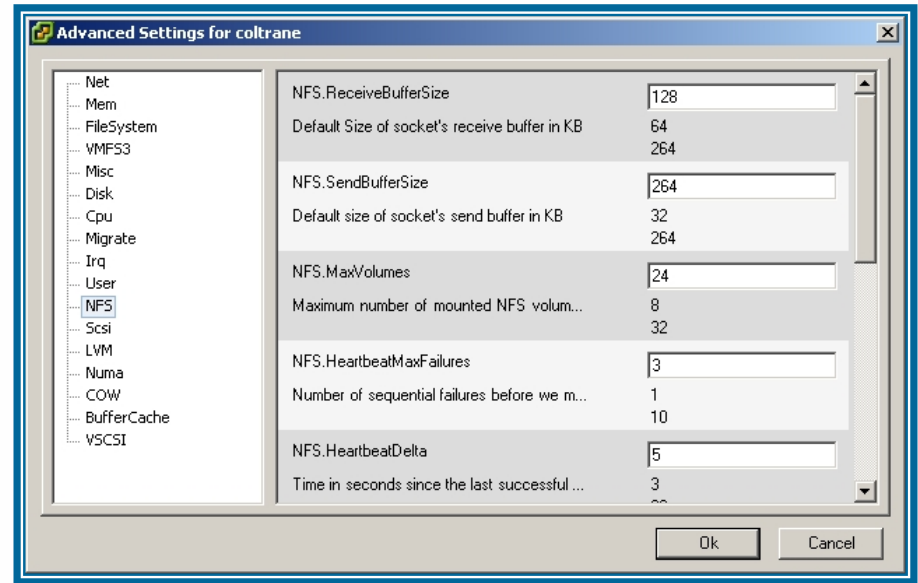
A terminal window titled 'root@giza:~' with standard window controls. The terminal displays the output of the command 'vmkfstools -P /vmfs/volumes/NetApp-Home'. The output shows that the NFS-1.00 file system is spanning 1 partition, with a label 'NetApp-Home', mode 'public', and capacity information. It also lists the UUID and the partition spanned.

```
root@giza:~  
[root@giza root]# vmkfstools -P /vmfs/volumes/NetApp-Home  
NFS-1.00 file system spanning 1 partitions.  
File system label (if any): NetApp-Home  
Mode: public  
Capacity 97430245376 (23786681 file blocks * 4096), 41162186752 (10049362 blocks) avail  
UUID: 325cd719-166b8372-0000-000000000000  
Partitions spanned:  
    nfs:NetApp-Home  
[root@giza root]#
```

- Mounted in the VMkernel NOT on the Service Console
- No need to modify /etc/fstab
- Recognized by vmkfstools as an NFS volume
- Virtual Machines can be stored there

NAS Tips and Tricks

- ESX Server needs full access to NFS datastores to create directories, set permissions
 - turn off **root squash**
- 8 NFS mounts per ESX Server allowed by default. To increase
 - Select host from inventory, right-click “Advanced Settings” and select “NFS”
 - Adjust “NFS.MaxVolumes”



- Avoid VM swapping to NFS volumes
 - Edit VM config file to add:
 - **sched.swap.dir = /vmfs/volumes/<volume_name>/<dir_name>"**
 - Where <volume_name> is a VMFS3 volume

Troubleshooting NAS

- Verify NAS filer/host configuration
- Make sure ESX Server is on the (Read/Write) host list on the filer
- Use ethereal to monitor NFS traffic
- On ESX Server use `tcpdump` to get a trace
- Verify the NFS connection properties on ESX Server
- To list Linux NFS host throughput, use:
 - > `hdparm -tT /dev/sd<x>`

Questions?

Presentation Download

Please remember to complete your
session evaluation form
and return it to the room monitors
as you exit the session

The presentation for this session can be downloaded at
<http://www.vmware.com/vmtn/vmworld/sessions/>

Enter the following to download (case-sensitive):

Username: cbv_rep
Password: cbvfor9v9r

Some or all of the features in this document may be representative of feature areas under development. Feature commitments must not be included in contracts, purchase orders, or sales agreements of any kind. Technical feasibility and market demand will affect final delivery.

VMWORLD 2006

