

IBM Storage SAN Volume Controller  
IBM Storwize V7000



# Guidelines for the Attachment of VMware iSCSI Hosts to SAN Volume Controller and Storwize V7000 - Errata

**Version 6.1.0.0 and Higher**

February 8<sup>th</sup> 2012

## Table of Contents

Table of Contents .....	2
Introduction .....	3
Who should use this guide .....	3
Last update .....	3
Change History .....	3
System prerequisite .....	3
Single session VMware iSCSI Support with SVC or Storwize V7000 version 6.1.x 6.2.x and 6.3.0.0 .....	4
iMware SCSI Version Support .....	5
Configuration Guidelines .....	6
Software Mis-configuration .....	6
Configuring a Single Software initiator Session .....	8
Detailed Description of Supported and Non-Supported Environments .....	12
Illustration 1 : Supported Configuration .....	12
Illustration 2: Unsupported Configuration .....	14
Illustration 3 : Supported Configuration .....	16
Illustration 4: Unsupported Configuration .....	18
Multi Session VMware iSCSI Support with SVC and Storwize V7000 version 6.3.0.1 and higher .....	19
Maximum Sessions .....	19
Connectivity .....	19
Fibre Channel to iSCSI comparison .....	20
VMware iSCSI Specific behavior Details .....	22

## Introduction

This document outlines the current guidelines and restrictions for the attachment of VMware iSCSI hosts to SVC and Storwize V7000.

## Who should use this guide

This errata should be used by anyone using the IBM System Storage SAN Volume Controller Host Attachment Guide or Information Center and IBM Storwize V7000 Host Attachment Guide or Information Center who wishes to attach VMware iSCSI Hosts

## Last update

This document was last updated: February 8<sup>th</sup> 2012

## Change History

The following revisions have been made to this document:

<b>Date</b>	<b>Sections Modified</b>
October 18 <sup>th</sup> 2011	New Document
November 18 <sup>th</sup> 2011	Update for 6.3 ( 6.2.x and Higher )
February 8 <sup>th</sup> 2012	Update for 6.3.0.1 and higher VMware multi-session support

## System prerequisite

The minimum supported software version for SAN Volume Controller and Storwize V7000 VMware iSCSI single session support is version 6.1.x. The minimum supported software version for SAN Volume Controller and Storwize V7000 VMware iSCSI multi-session support is version 6.3.0.1

## **Single session VMware iSCSI Support with SVC or Storwize V7000 version 6.1.x 6.2.x and 6.3.0.0**

- **Only single session is supported between VMware ESXi Server and a SVC Node or a Storwize V7000 Node**
- **Multi-session environments are not supported with 6.1.x, 6.2.x or 6.3.0. Please refer to the 6.3.0.1 section for multi-session support.**
- **Support is with the VMware iSCSI Software initiator only**
- **Each node in the SVC system or Storwize V7000 system corresponds to a single unique IQN.**
- **Each VMware host can have multiple Ethernet ports and each port can be configured with one or more IPs.**
- **Each SVC node can have multiple 1G/10G ports. We can configure one target IPv4 and one target IPv6 on each these ports**
- **iSCSI initiators use Static/Dynamic Discovery method to determine the target IPs configured on the Target IQN (Each SVC/Storwize V7000 Node ).**
- **Initiators discover all target IPs configured on the SVC/Storwize V7000 Node target IQN.**
- **Once target IPs are discovered VMware iSCSI tries to log-in to the target IQN via all of the Target IPs discovered, only if IPs are accessible (ping-able).**
- **An iSCSI session is established when an iSCSI initiator logs onto a target IP over TCP.**

## **VMware SCSI Version Support**

Please visit the SAN Volume Controller or Storwize V7000 Support pages for the latest information on the supported version of VMware.

SAN Volume Controller

<https://www-304.ibm.com/support/docview.wss?uid=ssg1S1003907>

Storwize V7000

<https://www-304.ibm.com/support/docview.wss?uid=ssg1S1003908>

## Configuration Guidelines

### Software Mis-configuration

Note: VMware initiators will discover all target IPs that are present on the Storwize V7000 or SAN Volume Controller system. Once target IPs are discovered, by default, VMware tries to create iSCSI session on all the accessible Targets.

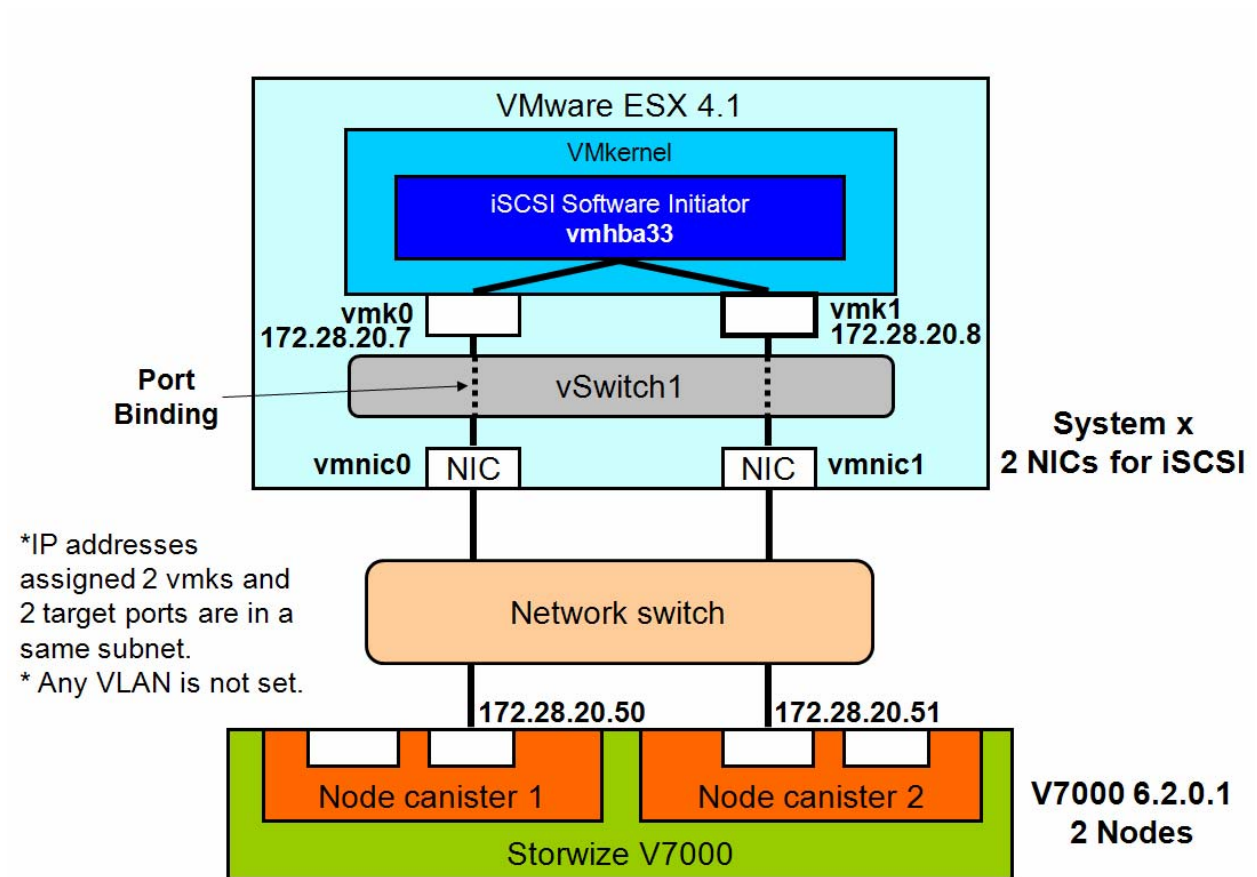


Fig 1 VMware setup with a Storwize V7000

The black lines from NIC to Network Switch in Fig 1 indicate actual physical links (Ethernet Links) between server and two Storwize V7000/SVC Nodes.

In the context of Fig 1 VMware will discover 4 paths to the Storwize V7000 system automatically.

Initiator IPs	Target IPs	Paths
172.28.20.7	172.28.20.50 on IQN1	Path1
172.28.20.7	172.28.20.51 on IQN2	Path2
172.28.20.8	172.28.20.50 on IQN1	Path3
172.28.20.8	172.28.20.51 on IQN2	Path4

This means the ESX 4.1 will see four paths per iSCSI lun or Device. In other words the VMware host has two iSCSI sessions/paths on each SVC/Storwize V7000 Node, one from 172.28.20.7 and the other from 172.28.20.8. **This is not supported and is a current limitation.** Ensure that each Node has single login session from ESX 4.1.

This is outlined in the steps below:

In a VMware environment, we should use 'Static Discovery' in the initiator property.

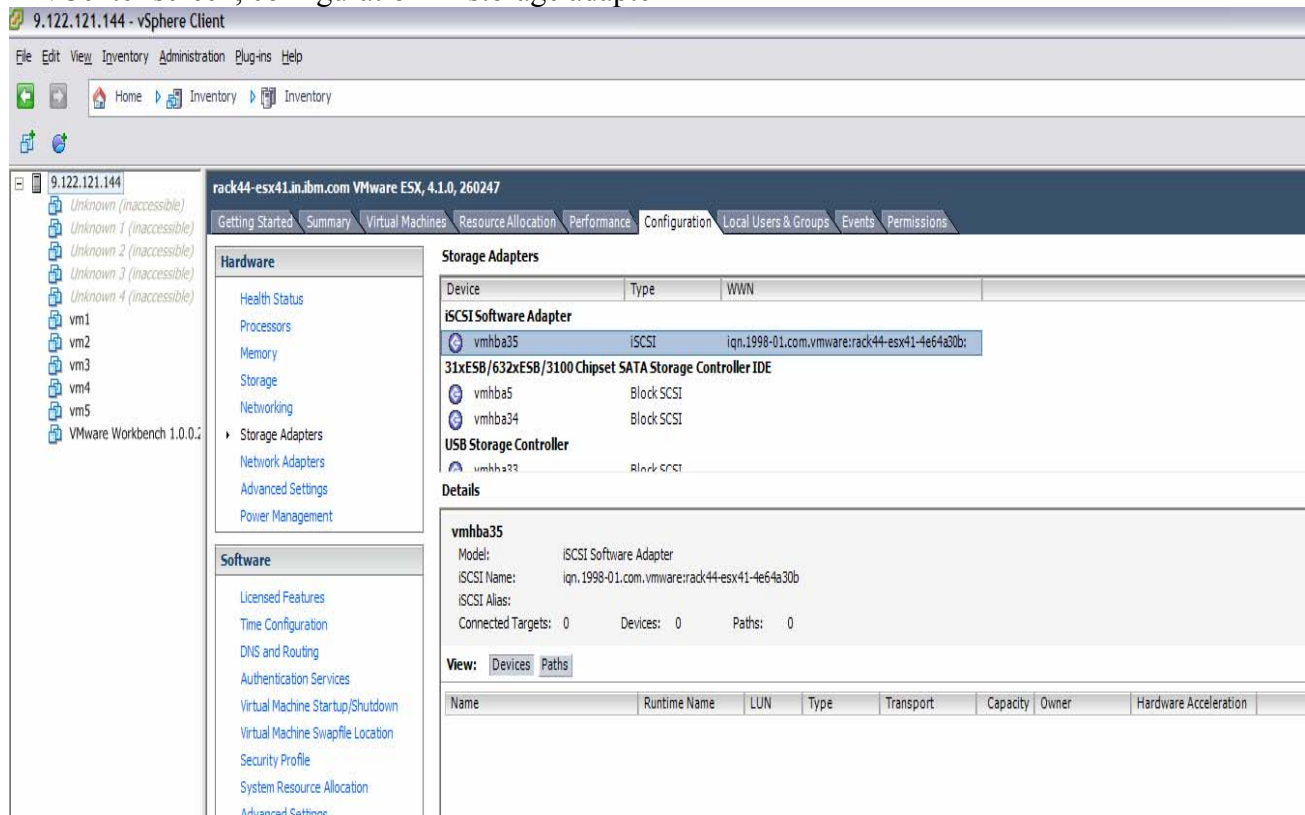
Select 'Dynamic Discovery', your ESX will see all target IPs on the Storwize V7000/SVC. Then two IPs on the same node, a total four IPs from one V7000 can be recognized from ESX. This will cause multi-session to happen on ESX servers by default. This is only supported with SVC or Storwize V7000 Version 6.3.0.1 and higher

## Configuring a Single Software initiator Session

In order to configure single session, use 'Static Discovery' and set only one IP address from each SVC/Storwize V7000 node.

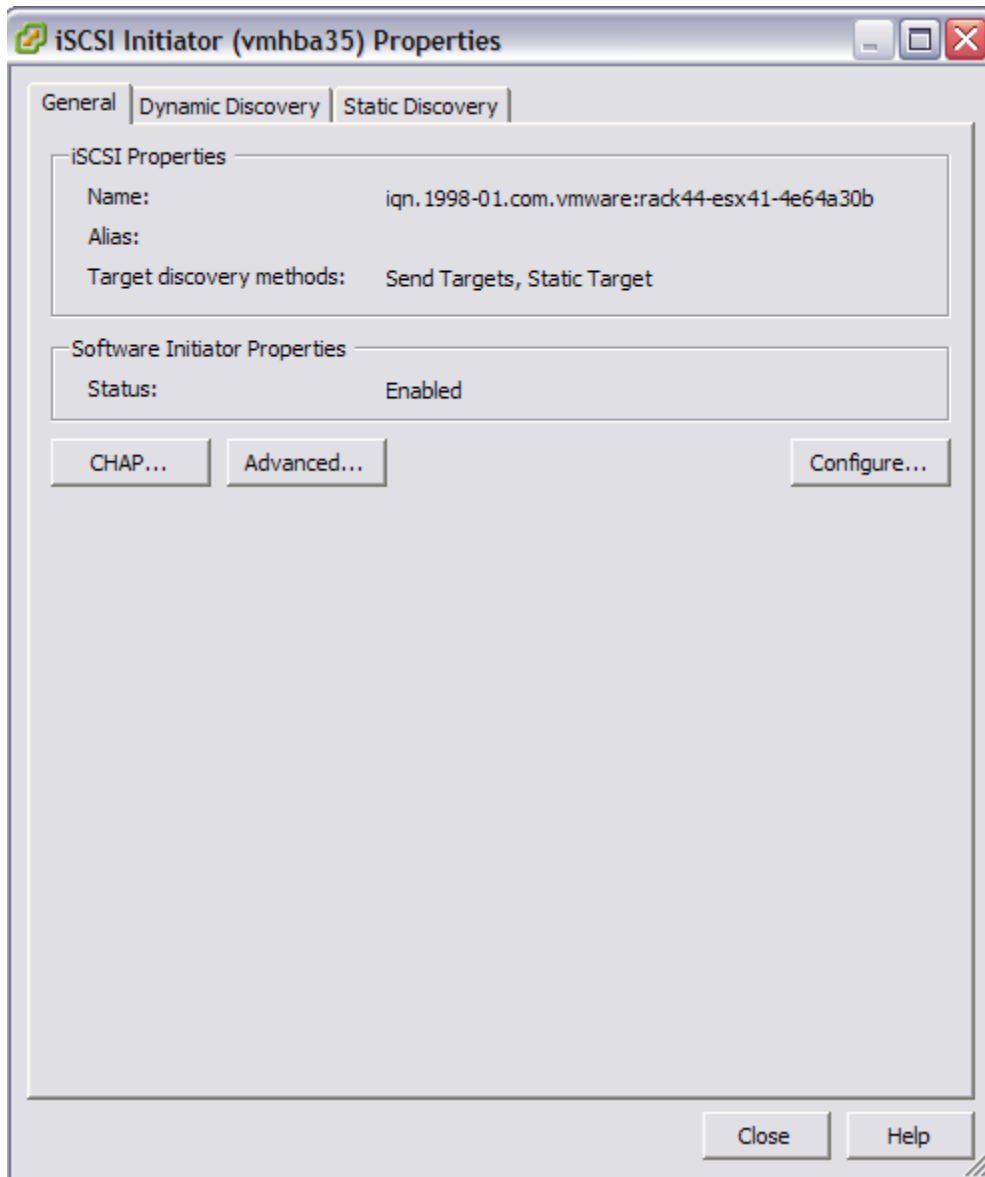
Ex:

- 1) In vCenter screen, configuration -> storage adapter




- 2) Next select iSCSI Software Adapter  
Right click, then select and open 'property'





In properties, select 'Static Discovery'.



Use the Add button to ensure there is only one Target IP per Target IQN.

Once Target IPs are discovered then the VMware iSCSI logs onto all of them, VMware will rescan for the devices against each path. For each Storwize V7000 /SVC iSCSI device/lun exported to VMware there must be at-most two paths, one each for the Storwize V7000/SVC node.

**Additional Information:**

- In a clustered setup, where multiple ESX hosts are involved, each of the ESX servers/hosts need to have an iSCSI session independently, with the restrictions described above.
- In configurations where multiple Ethernet cards are involved VMware Teaming capability can be utilized to get maximum throughput.

**References:**

- 1) [http://www..com/pdf/vsphere4/r40/vsp\\_40\\_iscsi\\_san\\_cfg.pdf](http://www..com/pdf/vsphere4/r40/vsp_40_iscsi_san_cfg.pdf)

## Detailed Description of Supported and Non-Supported Environments

The following section gives examples of supported and non supported configurations

### Illustration 1 : Supported Configuration

Consider the configuration below with the following details:

- VMware ESX 4.1 server with just one NIC Port and a single IPv4 configured on it.
- A Storwize V7000 system having a single Target IPv4 configured on an Ethernet Port.

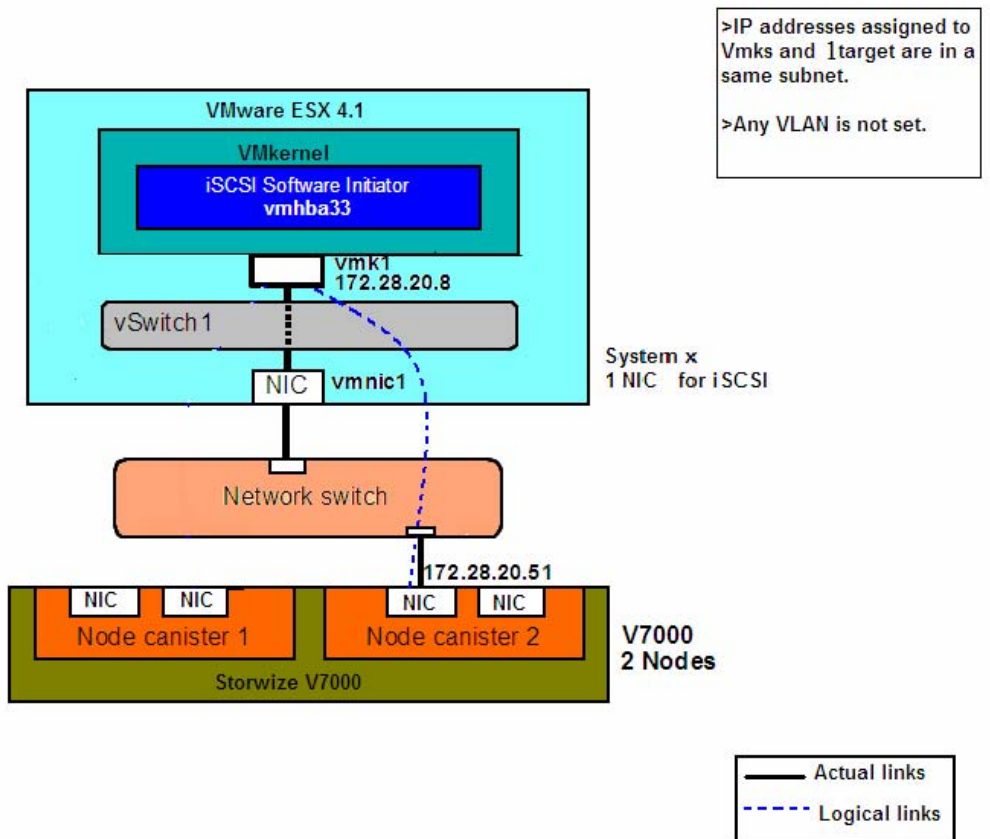


Fig 2 with VMware a single NIC Port

Fig 2 is a Supported Configuration.

The VMware iSCSI initiator discovers a single target IP on node 2. It then tries to log-in to the target IQN via the discovered IP. By the nature of configuration, only a single session on the Storwize V7000/SVC node exists. VMware just sees a single path for each of the iSCSI luns exported by the Storwize V7000/SVC, hence there is no need for NMP on the VMware ESX side.

Once node 2 goes down, the Target IP fails over to Nnode 1. In such an event sessions and IOs will be serviced by node 1. Once node 2 is back online the IP failback to node 2, Session and IOs are then managed and serviced by node 2.

**This configuration is supported by Storwize V7000 and SAN Volume Controller**

## Illustration 2: Unsupported Configuration

Consider the configuration below with the following details:

- VMware ESX 4.1 server with just one NIC Port and one IPv4 and IPv6 configured on it.
- A Storwize V7000 system having one Target IPv4 and one Target IPv6 configured on an Ethernet Port.

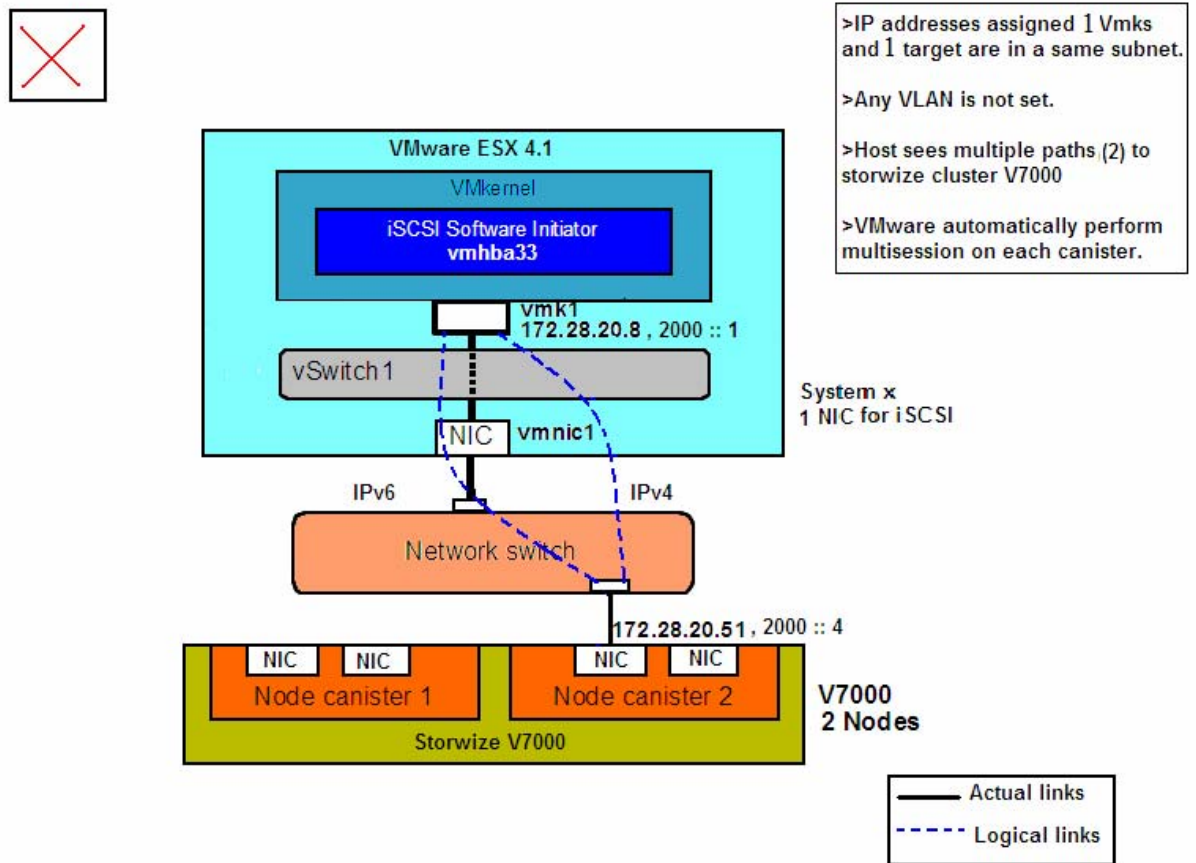


Fig 3 VMware with a single NIC Port and 2 Configured IPs

Fig 3 is an un-supported Configuration.

The VMware iSCSI initiator discovers two target IPs (One IPv4 and one IPv6) on node 2. It then tries to log-in to the target IQN via the discovered IPs. By the nature of configuration, two sessions exist on the Storwize V7000/SVC node, one IPv4 session and one IPv6.

VMware sees two paths for each of the iSCSI luns exported by the Storwize V7000 /SVC node. **This configuration is NOT supported by Storwize V7000 or SAN Volume Controller.**

**Resolution:**

Remove either of the IPs from the target node 2 so that only one logical link is possible to ensure there is single session on node 2.





exported by the Storwize V7000 . We support all the Multipath Policies ( Round Robin, Fixed etc ).

Once either node goes down, the Target IP fails over to the other node. In such an event sessions and IOs will be serviced by partner the node . Once the failed node is back online the IP failback happens from the partner node Session and IOs are then managed and serviced by the original node.

## Illustration 4: Unsupported Configuration

Consider the configuration below with the following details:

- VMware ESX 4.1 server with just one NIC Port and one IPv4 configured on it.
- A Storwize V7000 system having two Target IPv4 IPs configured on the Ethernet Interfaces.

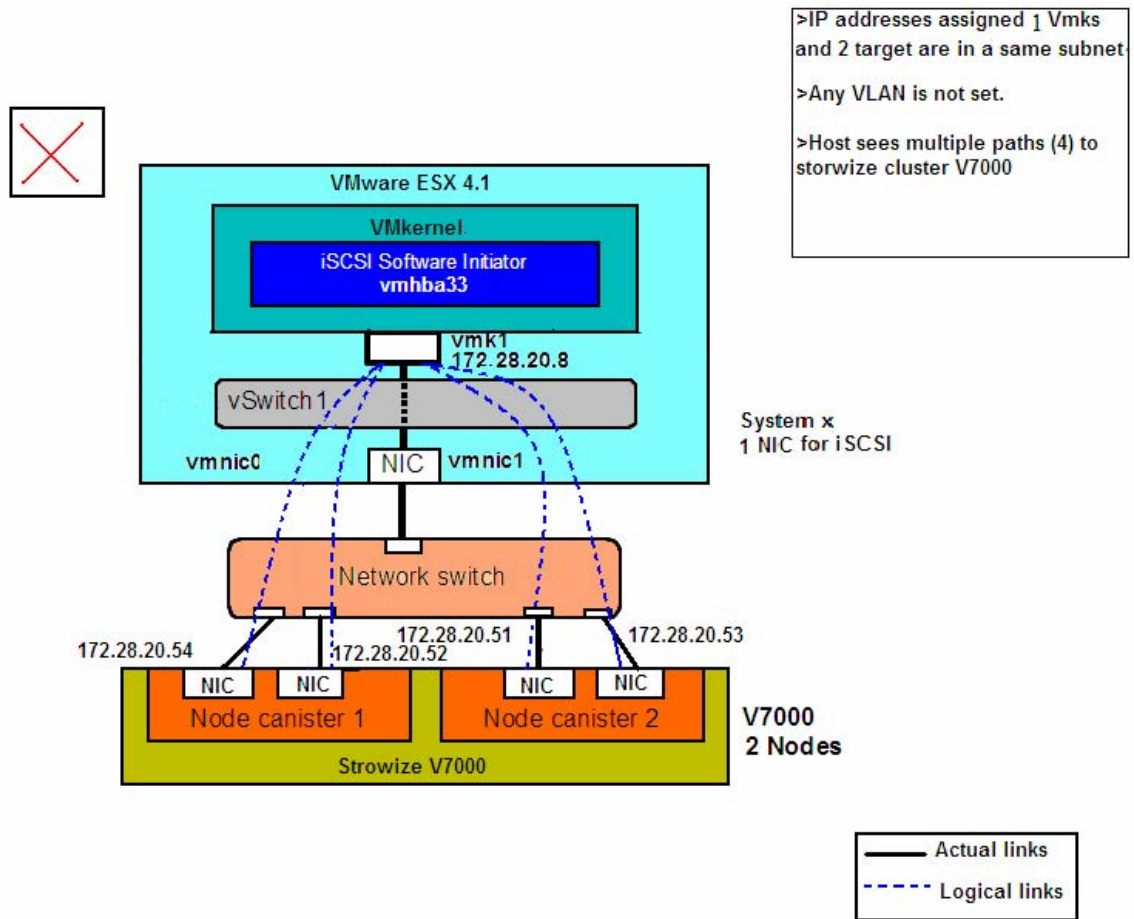


Fig 5 VMware with a single NIC Port and 2 Configured IPs on each Storwize V7000 Node

Fig 5 is an un-supported configuration.

The VMware iSCSI initiator discovers two target IPs on each node. It then tries to log-in to the target IQN via the discovered IPs. By the nature of this configuration, two sessions exist on the Storwize V7000/SVC node.

VMware sees more than one path for each of the iSCSI luns exported by each **Storwize V7000/SVC node. This configuration is **NOT** supported by Storwize V7000 or SAN Volume Controller**

#### **Resolution:**

Ensure that there is a single logical path between each target node and the VMware ESX server. This will ensure that there is single session on each of the nodes.

## **Multi Session VMware iSCSI Support with SVC and Storwize V7000 version 6.3.0.1 and higher**

**Note : Current Support is for Software initiator only unless otherwise stated on the SAN Volume Controller or Storwize V7000 Support Site**

SAN Volume Controller

<https://www-304.ibm.com/support/docview.wss?uid=ssg1S1003907>

Storwize V7000

<https://www-304.ibm.com/support/docview.wss?uid=ssg1S1003908>

### ***Maximum Sessions***

**A maximum of 4 sessions are supported between a the VMware ESXi Server and a SAN Volume Controller or Storwize V7000 node (canister). This means a maximum of 8 paths are available to the Disk/Lun i.e. 4 paths or sessions to each SAN Volume Controller or Storwize V7000 node (canister) from each iSCSI initiator IQN.**

### ***Connectivity***

- Each node in the SVC system or Storwize V7000 System corresponds to a single unique IQN.
- Each VMware host can have multiple Ethernet ports and each port can be configured with one or more IPs.
- Each SVC/Storwize V7000 node can have multiple 1G/10G ports. One target IPv4 and one target IPv6 can be configured on each of these ports.
- VMware iSCSI initiators use Static/Dynamic Discovery to find and communicate with the target IPs configured on the Target IQN (Each SVC/Storwize V7000 node/canister).

- Initiators discover all target IPs configured on the SVC node/Storwize V7000 canister target IQN.
- Once target IPs are discovered the VMware iSCSI initiator attempts to log-in to the target IQN via all of the Target IPs discovered only if IPs are accessible (ping-able).
- An iSCSI session is established when an iSCSI initiator logs into a target IP over TCP.

Note: VMware initiators will discover all target IPs that are present on the Storwize V7000 system. Once target IPs are discovered by default, VMware will attempt to create an iSCSI session on all accessible target IQNs.

### ***Fibre Channel to iSCSI comparison***

Criterion	FC	iSCSI (FC-equivalent in iSCSI)
Identification of the SAN elements.	FC supports: <ul style="list-style-type: none"> <li>➤ target WWPN</li> <li>➤ initiator WWPN</li> <li>➤ Both target and initiator WWPN</li> </ul>	iSCSI IQN iSCSI supports: <ul style="list-style-type: none"> <li>➤ Target IQN</li> <li>➤ Initiator IQN</li> <li>➤ Both Target and initiator functions</li> </ul>
Lun Discovery	FC initiator cards need to be zoned to be able to do FLOGI→PLOGI→PRLI. Post this zoning process luns are visible.	For iSCSI the initiator needs to be able to ping the target IPs. TCP login → iSCSI login Post this process luns are visible.
Pre-condition for setup	Initiator and target WWPN must be zoned together	Target and initiator IPs must be ping-able. Note that the identification element is IQN but the IPs are the mandatory transport over which iSCSI sessions are established. This implies two aspects: <ul style="list-style-type: none"> <li>➤ Even if Ethernet connectivity is established it is not guaranteed that iSCSI sessions will be established.</li> <li>➤ For iSCSI session to be established it is important that IPs are ping-able, i.e. between Initiator→ Target IPs</li> </ul>

Unique Identity	WWPN is unique in FC	IQN is the unique name in iSCSI.
IT nexus	A path is a Nport-id given to the initiator-target.	A path is an iSCSI login between initiator-Target IQN. Each path is a TCP/IP connection. So in scenarios where either initiator /target have multiple IPs the initiator can establish multiple sessions with the target. Each session is identified by IQN and IP tuple. IP tuple consists of one initiator IP and Target IP.
SVC related differences	If an IT nexus fails it will not be utilized and another path is sought.	If an IT nexus suffers a failure due to target constraints such as a node canister reboot the target IQN and hence the target IP will failover to the partner node
Target Discovery	Configures automatically via RCSN calls from the switch once the FC cards are plugged into the fabric.	Discovery of the target is manual and the administrator requests a target IP via one of the accessible target IPs or via iSNS server. <ul style="list-style-type: none"> <li>➤ The initiator IQN requests the target IPs with the Target IQN.</li> <li>➤ \the target IQN responds with the target IPs to the IQN</li> <li>➤ Initiator IQN will then check the accessibility of the target IQN via ping &lt;IP&gt; and then performs the login procedure</li> </ul>

## **VMware iSCSI Specific behavior Details**

VMware has a notion of auto discovery. Which means VMware continuously tries to login on the target IPs advertised by the target IQNs. This has the following implications:

- If the target IP is advertised and the iSCSI login fails. The VMware server will waste CPU cycles attempting to login again via that initiator to target IP.
- If a target IP is down for a period and then comes back on line the VMware host will automatically find the IP and try to establish iSCSI login to the Storage.
- In cases where the fifth or subsequent connection from the same initiator IQN is attempted then the target will fail the login with a resource failure. VMware will be unaware and loop again for an iSCSI login utilizing both the host and storage resource.
- In cases where an initiator has more than one IP (via single or multiple Ethernet connections) VMware will attempt to login via all of the available initiator IPs to the advertised target IPs of the IQN thus causing multiple sessions. So for example if target IQN advertises two target IPs then:
  - If VMware has two initiator IPs Dynamic discovery will attempt to create four iSCSI sessions.
  - If VMware has four initiator IPs then Dynamic discovery will attempt to create eight iSCSI sessions.

## Example Supported Configuration

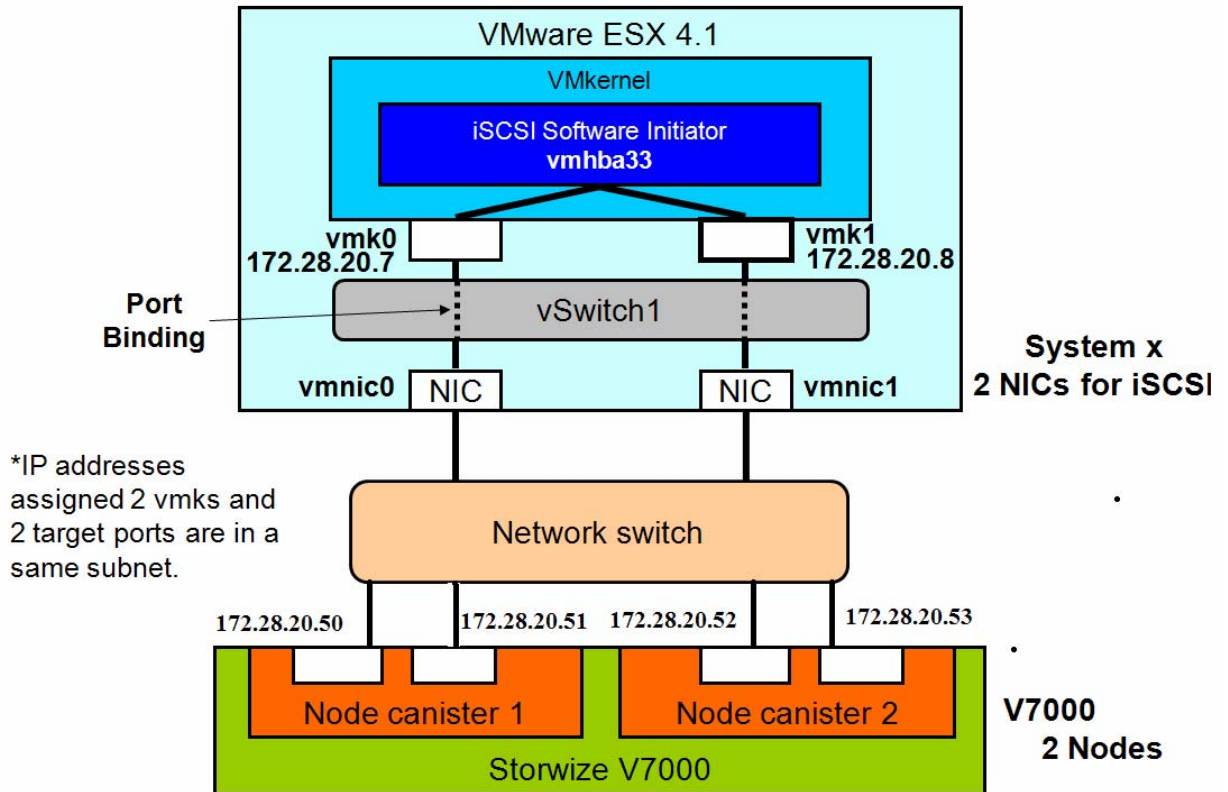


Fig 6 setup with a Storwize V7000 system

The black lines from NIC to network Switch in Fig 6 indicate actual physical links (Ethernet Links) between the server and the two Storwize V7000 canisters.

In the context of the Fig 6 VMware will discover eight paths to Storwize V7000 system automatically; Four paths per Storwize canister.

Initiators IPs	Target IPs	Paths
172.28.20.7	172.28.20.50 on IQN1	Path1
172.28.20.7	172.28.20.51 on IQN1	Path2
172.28.20.8	172.28.20.50 on IQN1	Path3
172.28.20.8	172.28.20.51 on IQN1	Path4
172.28.20.7	172.28.20.52 on IQN2	Path5
172.28.20.7	172.28.20.53 on IQN2	Path6
172.28.20.8	172.28.20.52 on IQN2	Path7

172.28.20.8	172.28.20.53 on IQN2	Path8
-------------	----------------------	-------

This means ESX 4.1 will see eight paths per iSCSI lun or Device. In other words the VMware iSCSI host has four iSCSI sessions/paths on each node/canister, Two from 172.28.20.7 and two from 172.28.20.8. **This current limitation is the only restriction imposed on the number of sessions for each canister for any given initiator IQN.** There is a requirement to ensure that each node/canister has a maximum of four login sessions from VMware ESX 4.1.

Note:

- If the initiators choose to attempt more than four sessions per node/canister the fifth session will receive a resource error from the attempted log in.
- When using servers in a clustered configuration, each of the VMware servers in the cluster will have an individual maximum limit of eight paths from the Storwize V7000 system, 4 paths per Canister.
- In cases where more than four ethernet adapters are used, the Host adapter will need to be Teamed/Bonded to a single IP to ensure that no more than four paths between Host and the Target IQN are created. Please consult the documentation for information on how to configure Teaming/Bonding.



