

# Windows 2012 NPIV on UCS Configuration Example



Document ID: 116471

Contributed by Felipe de Mello and Marcin Latosiewicz, Cisco TAC Engineers.

Jul 21, 2014

## Contents

### Introduction

#### Prerequisites

- Requirements

- Components Used

#### Configure

#### Live Migration

#### Quick Migration

#### Verify Live Migration

#### Troubleshoot

- Common Problems

- MPIO

#### Related Information

## Introduction

This document describes how to configure Windows Server 2012 N\_Port ID Virtualization (NPIV) on Unified Computing System (UCS) Version 2.1(2a). With this feature, a Virtual Machine (VM) that runs on a server can share a single adapter, and still have independent access to its own protected storage.

## Prerequisites

### Requirements

Cisco recommends that you have knowledge of these topics:

- Windows Fabric Network Interface Controller (fNIC) Driver Compatible with UCS Manager (UCSM) Version 2.1(2)
- UCSM Version 2.1(2) Virtual Interface Card (VIC) Firmware Image
- UCSM Version 2.1(2) on Fabric Interconnect/ I/O Modules (IOMs)
- Hyper-V 2012 and Windows 2012 Guests

### Components Used

The information in this document is based on these software and hardware versions:

- Hyper-V Version 3.0
- Windows Server 2012
- NetApp Storage
- UCS Chassis, Fabric Interconnects, and B-Series Servers

- Cisco Nexus 5000 Series Switches

The information in this document was created from the devices in a specific lab environment. All of the devices used in this document started with a cleared (default) configuration. If your network is live, make sure that you understand the potential impact of any command.

## Configure

Complete these steps in order to configure NPIV:

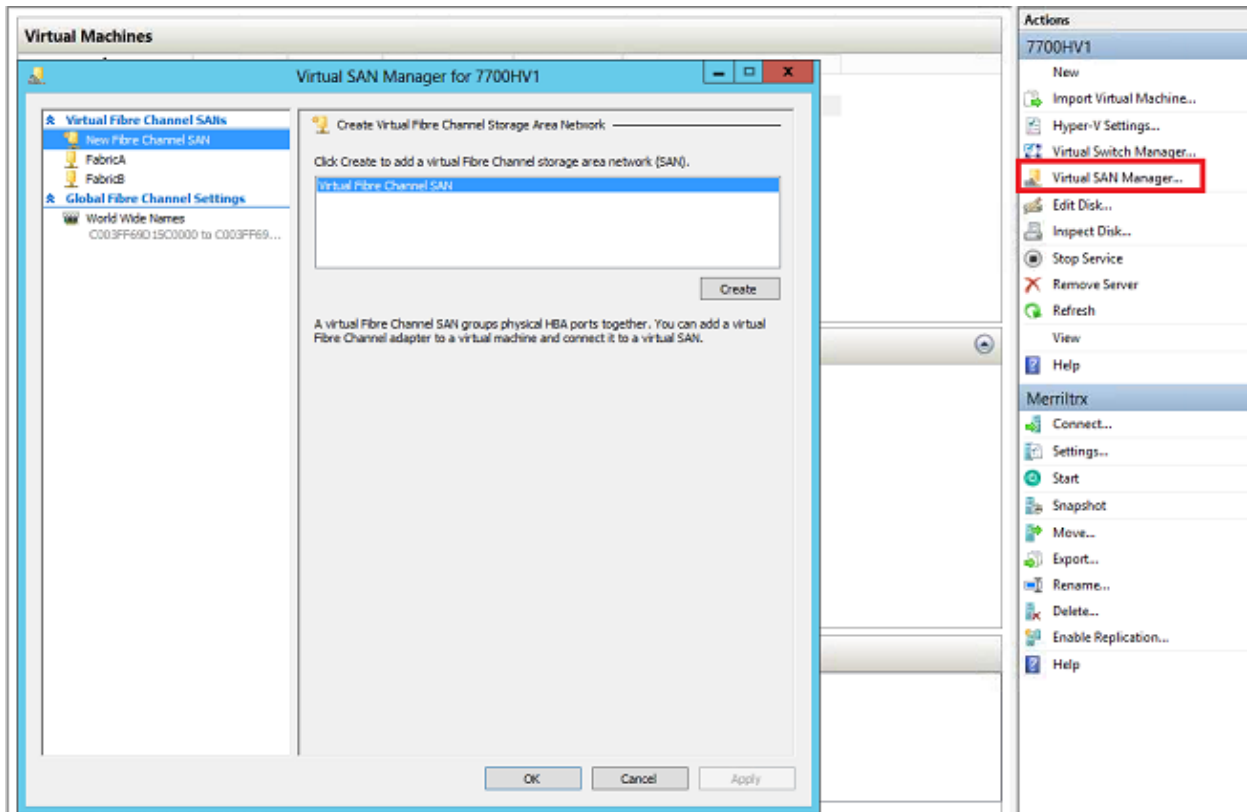
1. From a UCS perspective, you must configure your service profile with two Virtual Host Bus Adapters (vHBAs), one for each fabric. This image shows the vHBAs for one service profile. You can correlate the World Wide Port Names (WWPNs) with the service profile in the output that is shown later in the document.

**Note:** When you use Storage Area Network (SAN) storage in order to boot Hyper-V hosts, it is **highly recommended** that a separate set of vHBAs be used for VM traffic and for device booting. This example outlines a basic configuration with two vHBAs.

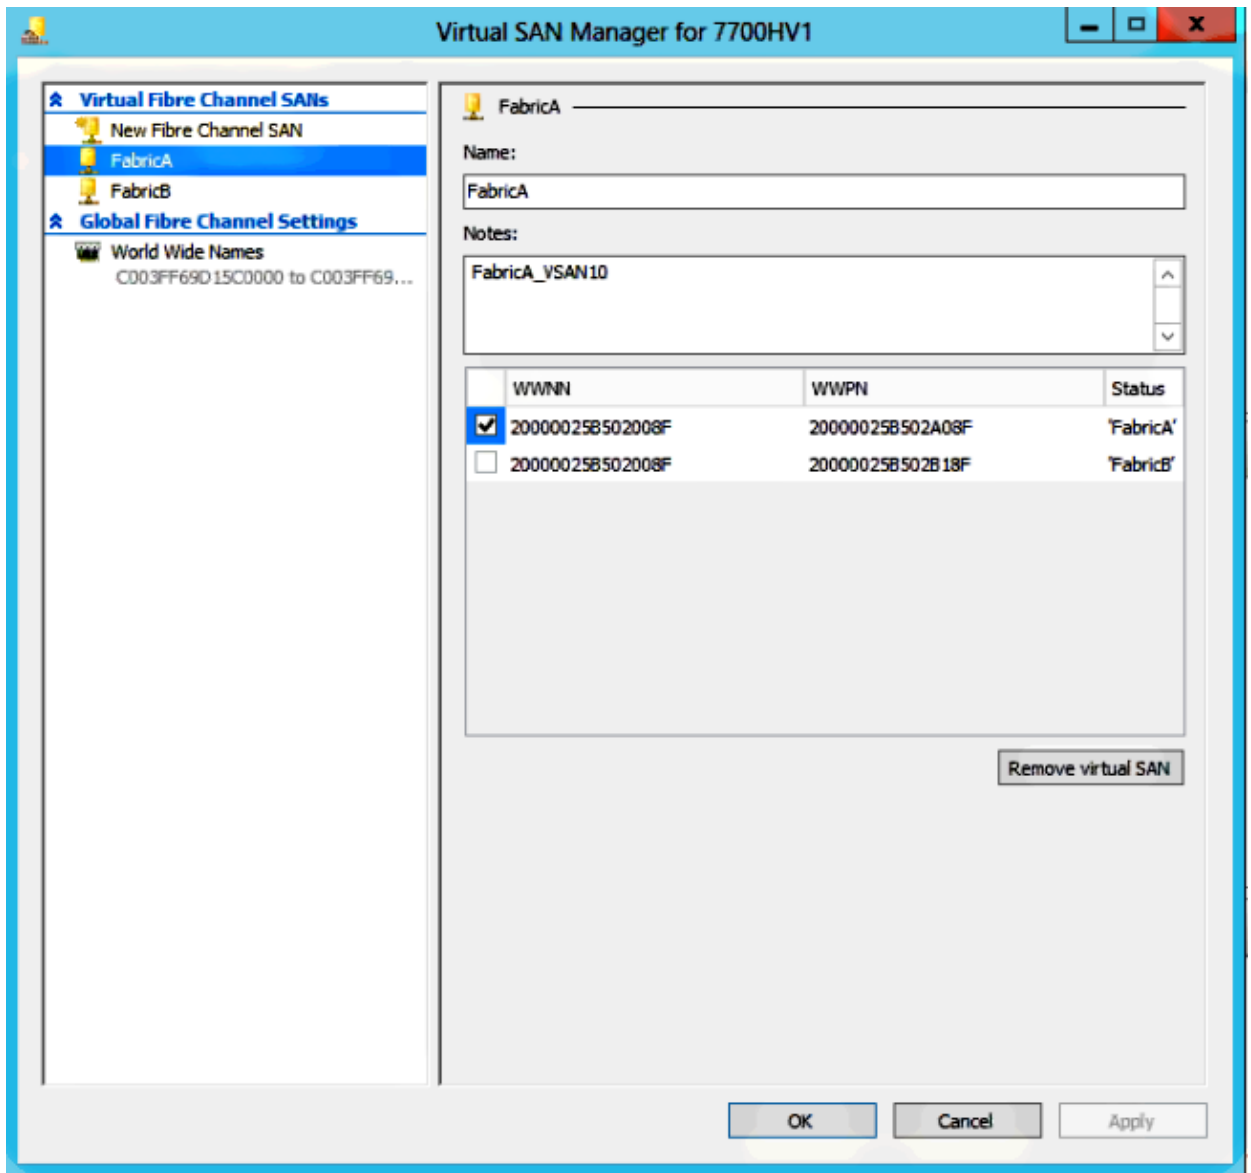
The screenshot displays the UCS Service Profile configuration interface. On the left, a tree view shows the hierarchy: Equipment > Servers > LAN > SAN > VM > Admin. The 'Service Profiles' section is expanded to show a 'root' node containing two hardware profiles: 7700HW1 and 7700HW2. Under 7700HW1, the 'vHBAs' section is expanded, showing two vHBAs: vHBA vHBA0-FabricA and vHBA vHBA1-FabricB. The right pane shows the configuration for these vHBAs, including the World Wide Node Name, Local Disk Configuration Policy, and SAN Connectivity Policy. Below the configuration, a table lists the vHBAs with their WWPNs, Desired Order, Actual Order, and Fabric ID.

Name	WWPN	Desired Order	Actual Order	Fabric ID
vHBA vHBA0-FabricA	20:00:00:25:85:02:A0:8F	3	5	A
vHBA vHBA1-FabricB	20:00:00:25:85:02:B1:8F	4	6	B

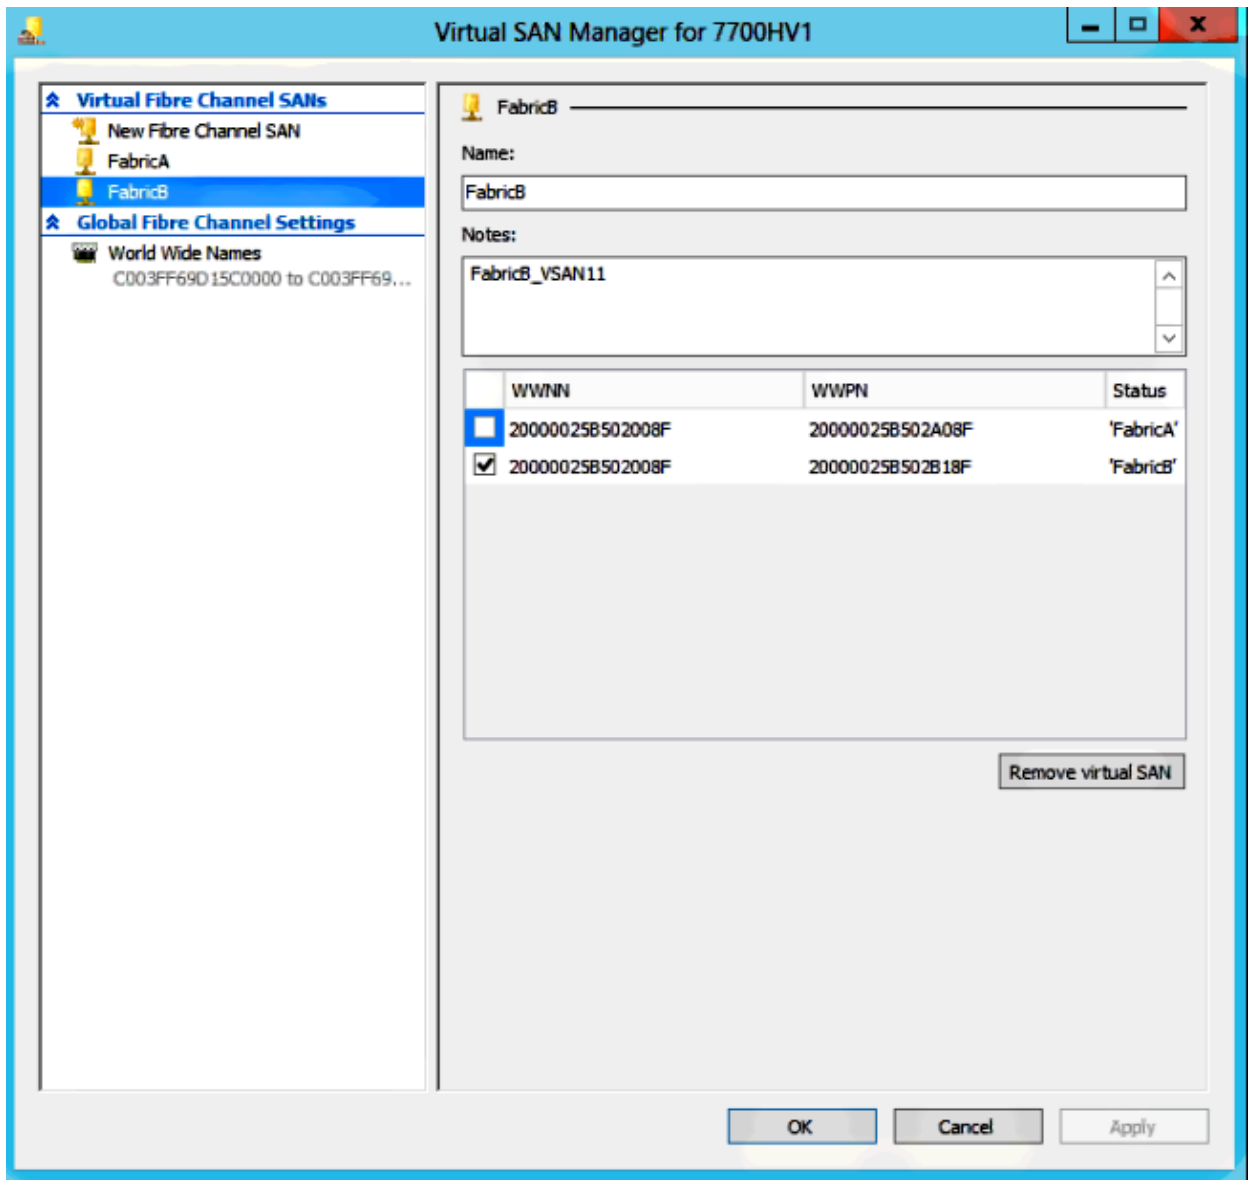
2. Configure your Virtual Storage Area Network (VSAN) from Hyper-V Manager. You must create two VSANs, one for each fabric. When you create a VSAN in other Hyper-V hosts, ensure that you use the same names; otherwise, Live Migration does not work. Click the **host**, and then click **Virtual SAN Manager**.



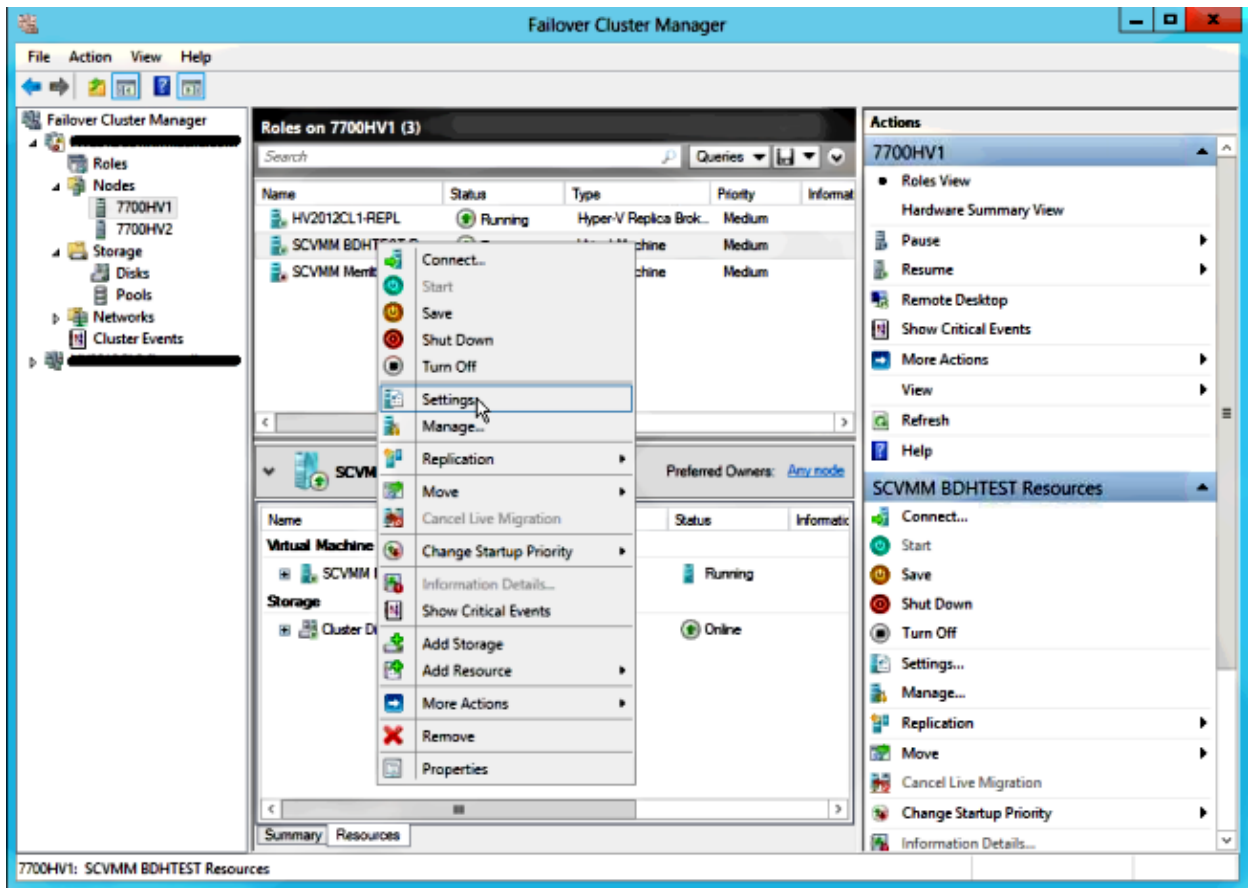
3. Create a new **Fibre Channel SAN** called FabricA, and select the **World Wide Node Name (WWNN)/WWPN** that corresponds to vHBA0–FabricA.



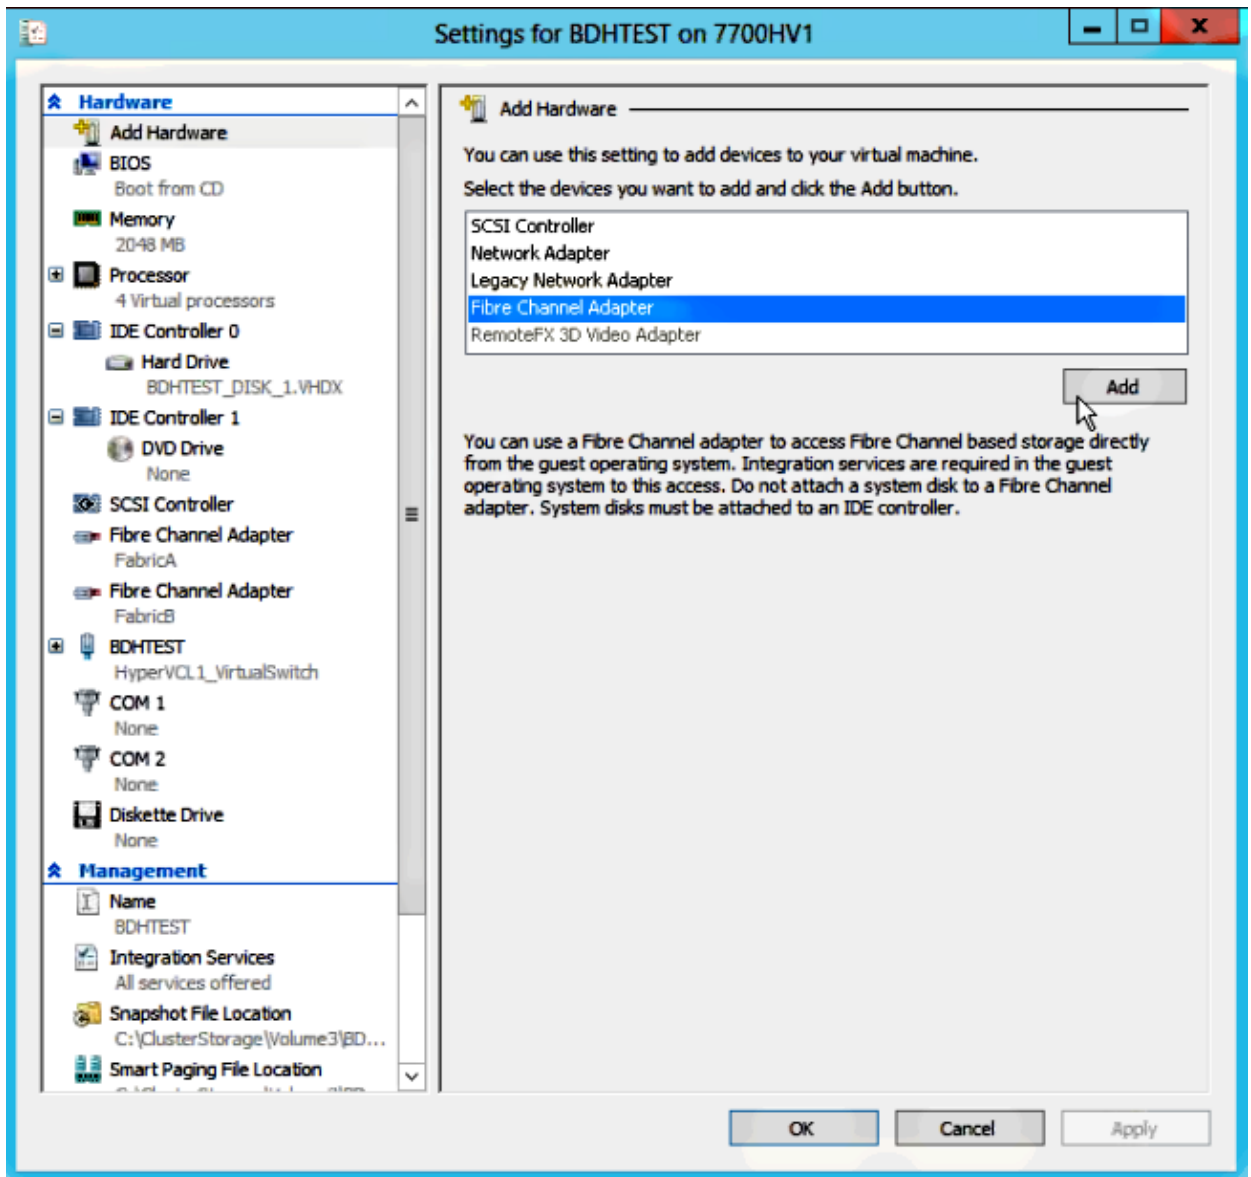
4. Add *FabricB*, and select the *WWNN/WWPN* that corresponds to vHBA1–FabricB.



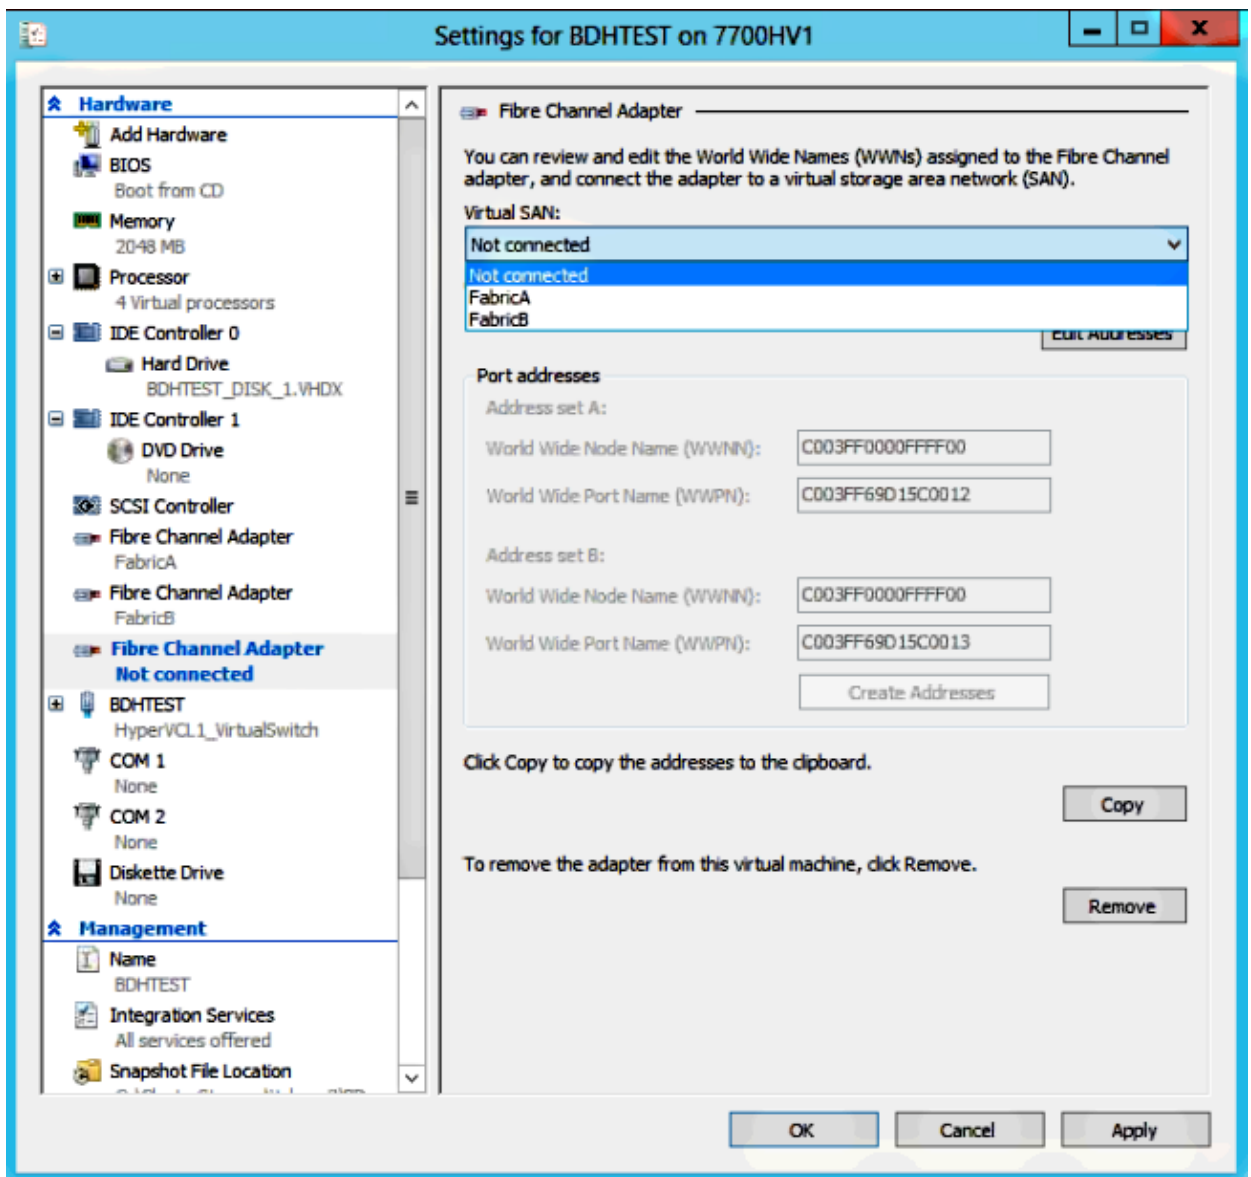
5. Configure the settings for the Windows Server 2012 VM, and add two fibre channel adapters. With the VM powered off, right-click and choose *Settings*. This is done with the Failover Cluster Manager because these hosts are part of a cluster.



6. Click *Add Hardware*, select *Fibre Channel Adapter*, and click *Add*.

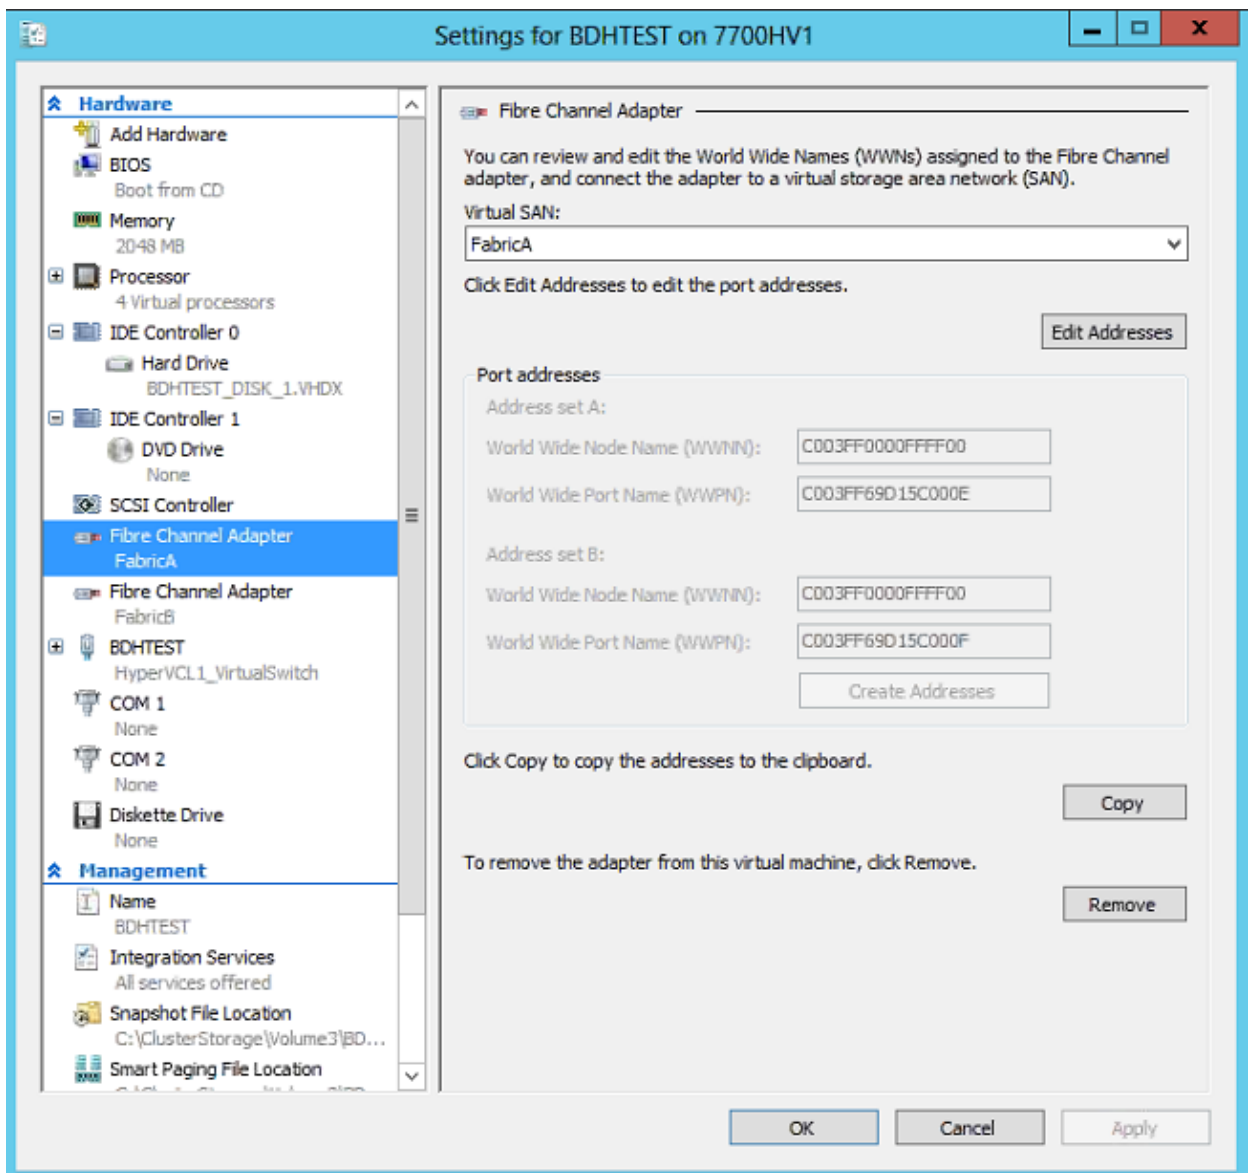


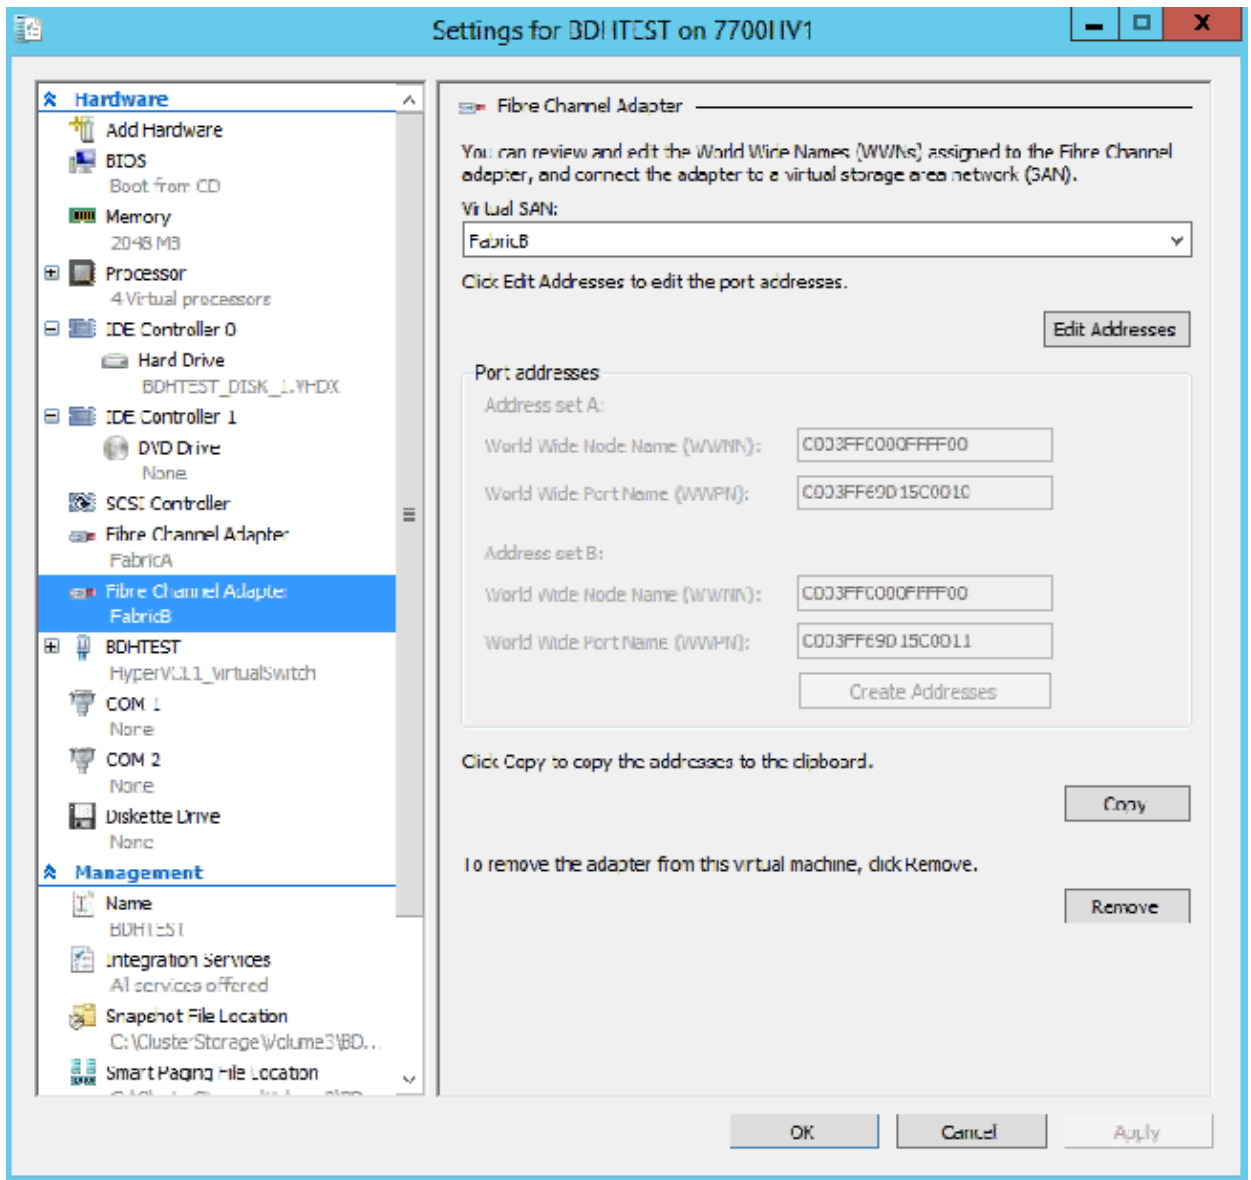
7. Select *Virtual SAN FabricA* for the first adapter, and *Virtual SAN FabricB* for the second adapter.



These images show both adapters and their respective WWNNs/WWPNs.







8. Add zoning on Nexus 5k switches.

Here is an example of how the Nexus 5000 Series switch might look (zoning for WWPNs of the UCS vHBAs are not shown):

```
! Zoning for HYPERVTEST Fabric A
```

```
fcalias name HYPERVTEST vsan 10
    member pwn c0:03:ff:69:d1:5c:00:0e
    member pwn c0:03:ff:69:d1:5c:00:0f

    zone name HYPERVTEST_to_NetAppl vsan 10
    Member fcalias HYPERVTEST
    Member fcalias NetAppl

zoneset name HyperVZoneset1 vsan 10
    member HYPERVTEST_to_NetAppl

zoneset activate name HyperVZoneset1 vsan 10
```

```

!Zoning for HYPERVTEST Fabric B

fcalias name HYPERVTEST vsan 11
  member pwwn c0:03:ff:69:d1:5c:00:10
  member pwwn c0:03:ff:69:d1:5c:00:11

  zone name HYPERVTEST_to_NetApp2 vsan 11
  Member fcalias HYPERVTEST
  Member fcalias NetApp2

zoneset name HypervZoneset2 vsan 11
  member HYPERVTEST_to_NetApp2

zoneset activate name HyperVZoneset2 vsan 11

```

9. Add the WWPNs to *NetApp* in order to make sure they can access the Logical Unit Numbers (LUNs).

The screenshot displays the 'Initiator Groups' section of the LUN Management interface. It features a table with the following data:


Name	Type	Operating System	ALUA	Initiator Count
BDHTEST	FC/FCoE	Hyper-V	Enabled	4
HV1_7700	FC/FCoE	Windows	Enabled	2
HV2_7700	FC/FCoE	Windows	Enabled	2
HV3_7700	FC/FCoE	Windows	Enabled	2
HV4_7700	FC/FCoE	Windows	Enabled	2
MSPEX1	FC/FCoE	Windows	Enabled	2
RELASQL	FC/FCoE	Windows	Enabled	2

Below the table, the 'Initiators' tab is active, showing a list of WWPNs:

- c0:03:ff:69:d1:5c:00:0e
- c0:03:ff:69:d1:5c:00:11
- c0:03:ff:69:d1:5c:00:10
- c0:03:ff:69:d1:5c:00:0f

## Live Migration

Each VM adapter has two sets of WWNN/WWPN. These are used by Hyper-V during a live migration. This image shows how each of the WWPNs is used during live migration.



Source: Hyper-V Virtual Fibre Channel Overview

It is important to note that there is a moment of overlap in which both WWPNs are logged in to the fabric.

This guarantees continuous work on the storage without interruption, even in case of migration failure.

The *Verify Live Migration* section shows the flogi database during the live migration process so you can see both WWPNs for the adapter flogi during the process.

## Quick Migration

Unlike live migration, quick migration temporarily suspends the VM that is moved.

Due to this, there is no reason to log in both WWPNs from a set. Instead, the VM can log out from one node and log in from a new node.

## Verify Live Migration

If everything is configured correctly, you should see a flogi entry in the flogi database for the UCS vHBAs and the VM fibre channel adapters.

```
NEXUS1# show flogi database
```

Interface	VSAN	FCID	Port Name	Node Name
fc1/31	10	0x930001	50:0a:09:83:8d:80:b7:ae	50:0a:09:80:8d:80:b7:ae
fc1/32	10	0x930000	50:0a:09:84:9d:80:b7:ae	50:0a:09:80:8d:80:b7:ae
San-po31	10	0x930002	24:1f:54:7f:ee:57:1d:c0	20:0a:54:7f:ee:57:1d:c1
San-po31	10	0x930003	20:00:00:25:b5:02:a0:8f	20:00:00:25:b5:02:00:8f <vHBA0-FabricA
San-po31	10	0x930004	20:00:00:25:b5:02:a0:9f	20:00:00:25:b5:02:00:9f
San-po31	10	0x930005	20:00:00:25:b5:02:a0:6f	20:00:00:25:b5:02:00:6f
San-po31	10	0x930006	20:00:00:25:b5:02:a0:7f	20:00:00:25:b5:02:00:7f
San-po31	10	0x930007	20:00:00:25:b5:02:a0:4f	20:00:00:25:b5:02:00:4f
San-po31	10	0x930008	20:00:00:25:b5:02:a0:5f	20:00:00:25:b5:02:00:5f
San-po31	10	0x930009	c0:03:ff:69:d1:5c:00:0e	c0:03:ff:00:00:ff:ff:00 <Set A for Adapter FabricA

```
NEXUS2# show flogi database
```

```
-----
```

Interface	VSAN	FCID	Port Name	Node Name
fc1/31	11	0x9f0001	50:0a:09:84:8d:80:b7:ae	50:0a:09:80:8d:80:b7:ae
fc1/32	11	0x9f0000	50:0a:09:83:9d:80:b7:ae	50:0a:09:80:8d:80:b7:ae
San-po32	11	0x9f0002	24:20:54:7f:ee:57:1a:80	20:0b:54:7f:ee:57:1a:81
San-po32	11	0x9f0003	20:00:00:25:b5:02:b1:8f	20:00:00:25:b5:02:00:8f <vHBA1-FabricB
San-po32	11	0x9f0004	20:00:00:25:b5:02:b1:9f	20:00:00:25:b5:02:00:9f
San-po32	11	0x9f0005	20:00:00:25:b5:02:b1:6f	20:00:00:25:b5:02:00:6f
San-po32	11	0x9f0006	20:00:00:25:b5:02:b1:7f	20:00:00:25:b5:02:00:7f
San-po32	11	0x9f0007	20:00:00:25:b5:02:b1:4f	20:00:00:25:b5:02:00:4f
San-po32	11	0x9f0008	20:00:00:25:b5:02:b1:5f	20:00:00:25:b5:02:00:5f
San-po32	11	0x9f000b	c0:03:ff:69:d1:5c:00:10	c0:03:ff:00:00:ff:ff:00 <Set A for Adapter FabricB

In order to show the LUN, open the Disk Management in the guest VM, and enter the *rescan disks* command. If the LUN appears twice, Multipath I/O (MPIO) is not enabled.

During a live migration, you should see the WWPN for both Address Set A and Address Set B in each of the switches.

NEXUS1# *show flogi database*

```
-----
```

Interface	VSAN	FCID	Port Name	Node Name
fc1/31	10	0x930001	50:0a:09:83:8d:80:b7:ae	50:0a:09:80:8d:80:b7:ae
fc1/32	10	0x930000	50:0a:09:84:9d:80:b7:ae	50:0a:09:80:8d:80:b7:ae
San-po31	10	0x930002	24:1f:54:7f:ee:57:1d:c0	20:0a:54:7f:ee:57:1d:c1
San-po31	10	0x930003	20:00:00:25:b5:02:a0:8f	20:00:00:25:b5:02:00:8f
San-po31	10	0x930004	20:00:00:25:b5:02:a0:9f	20:00:00:25:b5:02:00:9f
San-po31	10	0x930005	20:00:00:25:b5:02:a0:6f	20:00:00:25:b5:02:00:6f
San-po31	10	0x930006	20:00:00:25:b5:02:a0:7f	20:00:00:25:b5:02:00:7f
San-po31	10	0x930007	20:00:00:25:b5:02:a0:4f	20:00:00:25:b5:02:00:4f
San-po31	10	0x930008	20:00:00:25:b5:02:a0:5f	20:00:00:25:b5:02:00:5f
San-po31	10	0x930009	c0:03:ff:69:d1:5c:00:0e	c0:03:ff:00:00:ff:ff:00 <Address Set A
San-po31	10	0x93000a	c0:03:ff:69:d1:5c:00:0f	c0:03:ff:00:00:ff:ff:00 <Address Set B

NEXUS2# *show flogi database*

```
-----
```

Interface	VSAN	FCID	Port Name	Node Name
fc1/31	11	0x9f0001	50:0a:09:84:8d:80:b7:ae	50:0a:09:80:8d:80:b7:ae
fc1/32	11	0x9f0000	50:0a:09:83:9d:80:b7:ae	50:0a:09:80:8d:80:b7:ae
San-po32	11	0x9f0002	24:20:54:7f:ee:57:1a:80	20:0b:54:7f:ee:57:1a:81
San-po32	11	0x9f0003	20:00:00:25:b5:02:b1:8f	20:00:00:25:b5:02:00:8f
San-po32	11	0x9f0004	20:00:00:25:b5:02:b1:9f	20:00:00:25:b5:02:00:9f
San-po32	11	0x9f0005	20:00:00:25:b5:02:b1:6f	20:00:00:25:b5:02:00:6f
San-po32	11	0x9f0006	20:00:00:25:b5:02:b1:7f	20:00:00:25:b5:02:00:7f
San-po32	11	0x9f0007	20:00:00:25:b5:02:b1:4f	20:00:00:25:b5:02:00:4f
San-po32	11	0x9f0008	20:00:00:25:b5:02:b1:5f	20:00:00:25:b5:02:00:5f
San-po32	11	0x9f000b	c0:03:ff:69:d1:5c:00:10	c0:03:ff:00:00:ff:ff:00 <Address Set A
San-po32	11	0x9f000c	c0:03:ff:69:d1:5c:00:11	c0:03:ff:00:00:ff:ff:00 <Address Set B

## Troubleshoot

This section provides information you can use in order to troubleshoot your configuration.

### Common Problems

- The *device or driver does not support virtual fibre channel* message displays under the Status column of a Virtual Fibre Channel SAN in *Hyper-V Manager > Virtual SAN Manager* when the

Microsoft Windows 2012 FNIC driver is not at the correct version. Verify the current FNIC driver version by going to *Device Manager > Storage Controllers > Cisco VIC FCoE Storport Miniport > Properties > Driver*. Use the UCS Interoperability Matrix in order to determine which driver is supported based on the blade model, UCS firmware version, and adapter. If necessary, update the driver.

- Under certain conditions, Live Migration fails with the *Synthetic FibreChannel Port: Failed to finish reserving resources* message. A couple of things should be verified:
  - ◆ Whether the WWPNs are added at the storage target – initiator groups in NetApp.
  - ◆ Whether zoning information accounts for access of both sets of WWPNs assigned to VMs.
  - ◆ Whether the latest patches have been applied from Microsoft, which includes KB 2894032.
- Live migration might fail when the device uses the same pair of HBAs for booting and VM traffic. This is described in *Unified Computing System Virtual Machine Live Migration Fails with Virtual Fibre Channel Adapters*.

## MPIO

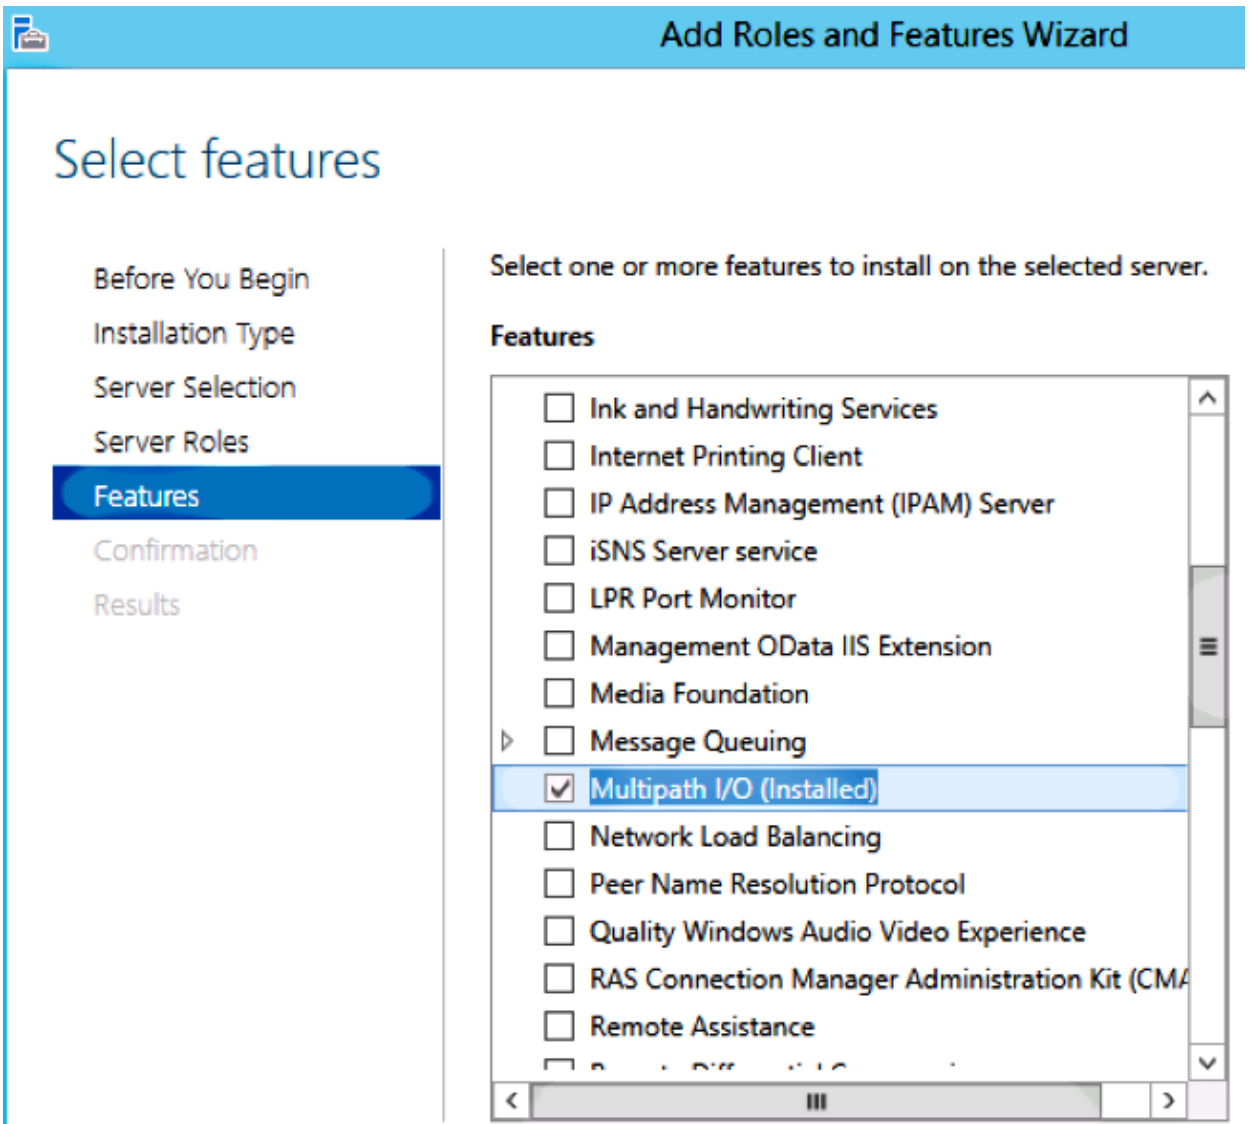
For resiliancy and fault tolerance, MultiPath I/O should be enabled on the operating system.

1. Enable Asymmetric Logical Unit Access (ALUA) on the NetApp for a particular initiator group.

The screenshot shows the 'Edit Initiator Group' window for 'VM1'. The window has two tabs: 'General' and 'Initiators'. The 'General' tab is selected. The fields are as follows:

Name:	VM1
Operating System:	Windows
Type:	FC/FCoE
<input checked="" type="checkbox"/> Enable ALUA (Asymmetric Logical Unit Access)	

2. Enable the MPIO feature on the Microsoft side. From the *Add Roles and Features*, make sure that MPIO is enabled.



## Related Information

- [Hyper-V Virtual Fibre Channel Overview](#)
- [Virtual Machine Live Migration Overview](#)
- [Technical Support & Documentation – Cisco Systems](#)