



TECHNICAL
WHITE
PAPER

Cisco UCS Deployment and Best Practice Guide

For Tintri VMstore™ in Cisco UCS B-Series Blade
Server Environments



www.tintri.com

Revision History

Version	Date	Description	Author
1.1	07/08/2016	Updated version	Rob Girard
1.0	05/20/2016	Initial Release	Rob Girard

Table 1 - Revision history

Contents

Executive Summary.....	5
Consolidated List of Practices.....	6
Intended Audience.....	7
Assumptions.....	7
Vendor Guidelines	7
Tintri VMstore Guides	7
Cisco UCS Documentation	7
Solution Overview	8
High Level Objectives.....	8
Tintri VMstore	9
Cisco UCS Series-B Blade Servers, Fabric Interconnects, and UCS Manager Software	10
VMware vSphere.....	10
Design Overview	10
Fabrics.....	10
VMstore NIC Teaming Behavior.....	12
Configuration	12
Cisco UCS Manager – Back Up Your Configuration	12
Cisco UCS – Default Maintenance Policies	13
Cisco UCS – Configure Appliance Ports.....	14
Pin Group Configuration on Fabric Interconnects.....	17
Network Configuration – Upstream Switches.....	17
Storage VLAN Configuration – LAN Cloud	17
Connect the Tintri VMstore to Cisco UCS	18
Rack the VMstore	18
Cabling	18
Configure the VMstore.....	20
Configure the Admin Network.....	20
Complete the “Out-of-the-Box” Configuration Process.....	20
Configure the Hosts.....	22
Configure UCS Servers.....	22
vSphere Host Configuration.....	26
Configure a New Distributed Port Group for a dvSwitch.....	26
Configure a VMkernel port.....	28

Mount the Datastore	31
Test Redundancy.....	33
Migrate VMs to the VMstore.....	33
Install the VAAI Plugin	33
Tintri Upgrades.....	34
Support.....	34
Conclusion.....	34
Appendix A: Design Considerations	36
LACP	36
Rationale against configuring LACP within an FI	36
Configuring LACP using dedicated IP storage switches	36
Native VLANs	37
Jumbo Frames.....	37
Appendix B – vStandard Switch Configuration.....	39
Appendix C - Nexus 5K Switch Configuration	43
Appendix D – Failure Scenarios & Testing Redundancy.....	45
Network Path Fault Tolerance	45
Host Failures.....	46
Human Error.....	46

Executive Summary

Cisco UCS B-Series Blade Servers are a popular server choice. A typical UCS configuration includes a Fabric Interconnect (FI) with two separate fabrics, and there are some important considerations for configuring Tintri VMstore storage systems in the UCS environment.

This guide describes the Tintri best practices for a UCS environment with VMware. Tintri recommends cabling the VMstore such that one port on each controller is configured on each UCS fabric. Fabric A is configured to preferentially carry storage traffic under normal operating conditions.

Distributed virtual switching (dvSwitch) is described as the best practice for switching in the VMware environment. However, a separate appendix describes the configuration of vstandard switching. Additional appendices describe design considerations for LACP, native VLAN use, and jumbo frames as well as configuration for the Cisco Nexus 5K switch.

Consolidated List of Practices

This section summarizes the recommended best practices for deploying a Tintri VMstore with Cisco UCS B-Series Blade Servers. Click on a recommendation to jump to the section of the document that corresponds to that specific best practice.

DO: Configure an upstream switch to allow broadcast traffic for the storage VLAN to flow on BOTH Fabrics in order to tolerate all failover scenarios without interruption.

DO: Unless you are using dedicated storage switches, leave LACP disabled. See LACP in the appendix for more detail.

DO: Backup up your UCS configuration before making changes. Better safe than sorry!

DO: Within UCS Manager, change the Reboot Policy within the default maintenance policy from “Immediate” to “User Ack” to avoid unplanned downtime.

DO: For each VMstore being installed, configure four (4) available UCS FI (fabric interconnect) ports as appliance ports to connect to the VMstore data ports.

DO: Ensure that all VLANs used by vSphere hosts, especially Storage VLANs, are configured end-to-end through the upstream switches.

DO: Configure the Storage VLAN on the UCS LAN Cloud to provide access to upstream ports and servers

DO: Connect the VMstore so that the “A” 10GbE data ports from each controller are connected to Fabric A, and the “B” 10 GbE data ports are connected to Fabric B (See Figure 8.)

DO: Choose the Optical (SFP+) 10 GbE adapter option when ordering a VMstore for use with UCS.

DO: Use TwinAx cables OR fiber cables with compatible SFP+ transceivers and LC to LC connectors to connect the VMstore to the FIs.

DO: Configure your VMstore according to the instructions in the Tintri Reference Guide.

DO: Modify your UCS Service Profile Template(s) or UCS Service Profile(s) to allow access to the newly created storage VLAN.

DO: On every vSphere Host, configure a new vmkernel port on the same subnet & VLAN as the VMstore.

DO: On every vSphere Host, mount a new datastore. Tintri recommends configuring one datastore per VMstore.

DO: Configure your vSphere hosts to use either a distributed virtual switch (dvSwitch) or vStandard switch (vSwitch).

DO: Install the Tintri VAAI plugin on each vSphere server.

DO: If you decide to use Jumbo Frames, ensure that MTU is configured to 9,216 throughout the **ENTIRE** path for the storage VLAN, including the interswitch links of upstream switches.

Intended Audience

This Best Practice Guide assists individuals who want to deploy production Cisco UCS B-series blade servers with VMware and Tintri VMstore™ storage systems. Prior knowledge of Cisco UCS, VMware virtualization, and networking will help in understanding the concepts covered in this best practice paper.

Assumptions

This paper provides best practices for deploying Tintri VMstore storage systems with pre-existing installations of Cisco UCS B-series blade servers and VMware vSphere. This document is not intended to replace the vendor-specific best practices provided by Cisco and VMware for their respective platforms. We recommend that you download and follow specific best practices provided by those companies where they are not otherwise covered in this document.

The reference system for these best practices employed VMware vSphere 6, Tintri OS 4.x, and UCS Manager. UCS B-series blades were connected through UCS 2208XP Fabric Extenders (FEX) to Cisco UCS 6248UP Fabric Interconnects (FI) to a Tintri VMstore storage system via 10GbE networking as described later. The best practices should work with any supported version of Tintri OS, VMware vSphere, and UCS Manager.

This document assumes you are working with a fully configured, highly-available VMware Infrastructure. The content in this document will provide specific guidance around configuring storage and related components to ensure it is also highly-available and suitable for mission-critical Tier 1 applications. Additional recommendations and tips on the design, deployment, and management of virtual infrastructures powered by Tintri VMstore storage appliances are provided in the [Tintri VMstore with VMware Best Practice Guide](#).

Vendor Guidelines

The following vendor guidelines were used in this deployment of Cisco UCS with VMware and Tintri. Every attempt should be made to follow the suggestions in these guides, while keeping in mind that these vendor-specific best practices were made without knowledge of your unique deployment architecture.

Tintri VMstore Guides

Tintri has published a best practice guide for deploying Tintri with VMware. The best practices in this guide apply to all versions of VMware vSphere supported by Tintri OS versions, up to and including vSphere 6. The Reference Guides shown below are useful references for initial installation and configuration of VMstore hardware.

- [Tintri VMstore with VMware® Best Practice Guide](#)
- **Tintri VMstore Administration Manual.** Accessible from the VMstore management UI via online help
- **Tintri VMstore T800 and T5000 Series Reference Guides.** Included with your VMstore. Digital versions are available for download in the [Tintri Support portal](#) under **Downloads – Hardware Platforms Documentation**.

Cisco UCS Documentation

Cisco updates its product documentation frequently. To find the latest documentation for your UCS B-Series systems and UCS Manager visit the following landing pages:

- [Cisco UCS B-Series Blade Servers](#)
- [Cisco UCS Manager](#)
- [Cisco UCS Hardware Compatibility List \(HCL\)](#) - validates that Tintri has and continues to undergo rigorous testing to ensure Tintri VMstores are a highly available storage solution that work across all UCS software releases since 2.x.
-

Solution Overview

This section introduces the infrastructure solutions used in this best practices guide: Tintri VMstore, Cisco UCS Series-B Blade Servers, and VMware vSphere 6. Review this section if you are unfamiliar with these products.

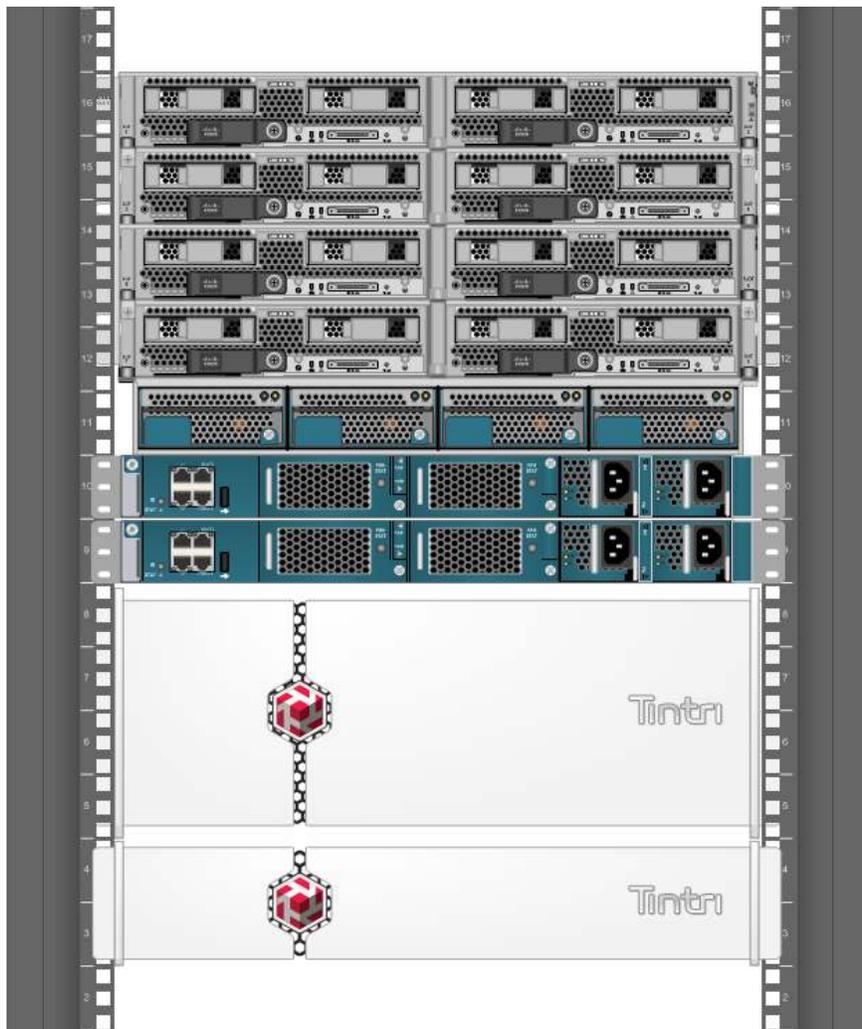


Figure 1 – Front view of a rack containing Cisco UCS with Tintri T800 & T5000 series VMstores

High Level Objectives

This purpose of this best practices guide is to help integrate Tintri storage with UCS to create a solution that is:

- Simple
- Scalable

- Performant
- Fault tolerant
- Resource Efficient

Tintri VMstore

The Tintri VMstore family provides a predictable and efficient storage environment designed for virtualized applications and desktops. The VMstore learns and adapts to your VM environment:

- **Set-up in minutes:** No complex storage configuration or tuning required. You only manage auto-aligned VMs and vDisks; there are no LUNs or volumes to manage.
- **Predictable performance:** The Tintri VMstore delivers consistent performance for all VMs. The Tintri VMstore T800 Hybrid-Flash Series combines HDD capacity and flash performance. The Tintri VMstore T5000 All-Flash Series provides an all-flash architecture designed to maximize performance for every IO operation.
- **High VM density:** Serve thousands of different types of VMs from a single VMstore with VM-level QoS and performance isolation.
- **Instant performance bottleneck visualization:** Real-time VM and vDisk-level insight on IO, throughput, end-to-end latency, and other key metrics enables rapid VDI performance diagnosis.

Tintri storage gives you a global view of all VMs stored, allowing you to identify performance and capacity trends without requiring a deep understanding of underlying storage. You can quickly identify performance hot spots at the hypervisor, network, and storage level with comprehensive performance visualization.

In a blade environment such as Cisco UCS, servers within a chassis share network uplinks, which can become saturated depending on the load. The **per-VM network latency** (see Figure 2) metric provided by the VMstore is an excellent indicator of whether such issues exist or not, providing you with the peace of mind that your infrastructure plumbing is clear and problem free.

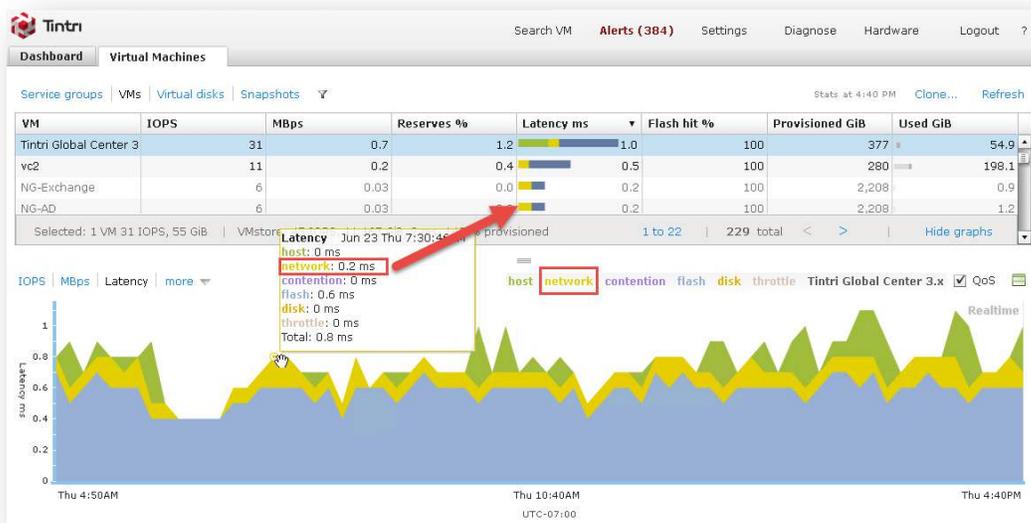


Figure 2 - Per-VM latency breakdown, including non-storage factors such as Network & Host Latency

All VMstore appliances use an active/standby dual-controller architecture. When combined with a fully redundant infrastructure with no single points of failure, such as the Cisco UCS configuration described in this document, VMstore upgrades can be performed online and uninterrupted. Follow the advice within this document to ensure all components within the data path are configured correctly to prevent interruptions during upgrades, as well as all other failure scenarios. More information on [upgrades](#) is provided later in this document.

Cisco UCS Series-B Blade Servers, Fabric Interconnects, and UCS Manager Software

Cisco UCS B-Series Blade Servers incorporate industry-standard server technologies. The UCS design reduces complexity, both at the hardware and management levels, across a distributed compute environment. Tintri storage further reduces the complexity of the overall environment, allowing you to focus on what matters most: your business and the applications that power it.

All Cisco UCS Blade Servers come with [Cisco UCS Manager](#). Cisco UCS with UCS Manager provides:

- Embedded integration of LAN, SAN, and management
- Server profiles and templates for policy-driven server provisioning

The Cisco UCS Fabric Interconnect (FI) is a core part of the Cisco Unified Computing System. Typically deployed in redundant pairs, Cisco Fabric Interconnects provide uniform access to both networks and storage with:

- Superior bandwidth
- High port density and low port-to-port latency
- Unified ports capable of line-rate, low-latency, lossless 1/10 Gigabit Ethernet
- Centralized management with Cisco UCS Manager

VMware vSphere

vSphere is the industry-leading virtualization platform. It enables users to run business-critical applications with confidence and respond quickly to business needs. VMware vSphere accelerates the shift to cloud computing for existing data centers and underpins compatible public cloud offerings, forming the foundation for the industry's best hybrid cloud model.

Design Overview

Fabrics

An important concept to understand in UCS is that it provides two Fabrics (A & B) that are independent and isolated from one another. While there is an internal network bridge between them, it is for management heartbeats ONLY and does NOT allow data to pass between fabrics. This is largely due to the constraints of Fibre Channel (FC), which requires separate fabrics. For Ethernet networking used in this guide, it is preferable to allow both fabrics to communicate with one another. This requires allowing the VLANs to flow through the upstream switches. Traffic flowing out of the fabric into the upstream switch is commonly referred to as "Northbound" traffic.

DO: Configure an upstream switch to allow broadcast traffic for the storage VLAN to flow on BOTH Fabrics in order to tolerate all failover scenarios without interruption.

When a network device in Fabric A needs to communicate with a network device on the same subnet located in Fabric B, that communication must flow through the upstream switch, otherwise the two network devices can't communicate. A common example of this is the NFS traffic exchanged between a vmkernel port in a vSphere host and an NFS server (ie. the active data port on a Tintri VMstore). If the active NIC used by vSphere is in Fabric A, and the active data port on the VMstore is in Fabric B, the communication can only take place by traversing through Fabric A, northbound to the upstream switches, and then back down into Fabric B (and traverse the same path on the return).

Keeping storage traffic within the same fabric between the host and storage (which is desirable because it has the fewest hops and lowest latency), requires coordination so that vSphere hosts AND the VMstore prefer the SAME fabric as their active link. Hosts will be configured to prefer the other fabric for all other traffic (VM networking, management, vMotion, etc.). In this way, ALL traffic has the benefit of staying "in-fabric". But we'll do it in a way that keeps our solution fault-tolerant to any single failure, including an entire fabric!

See Figure 3 for a visual representation of both fabrics.

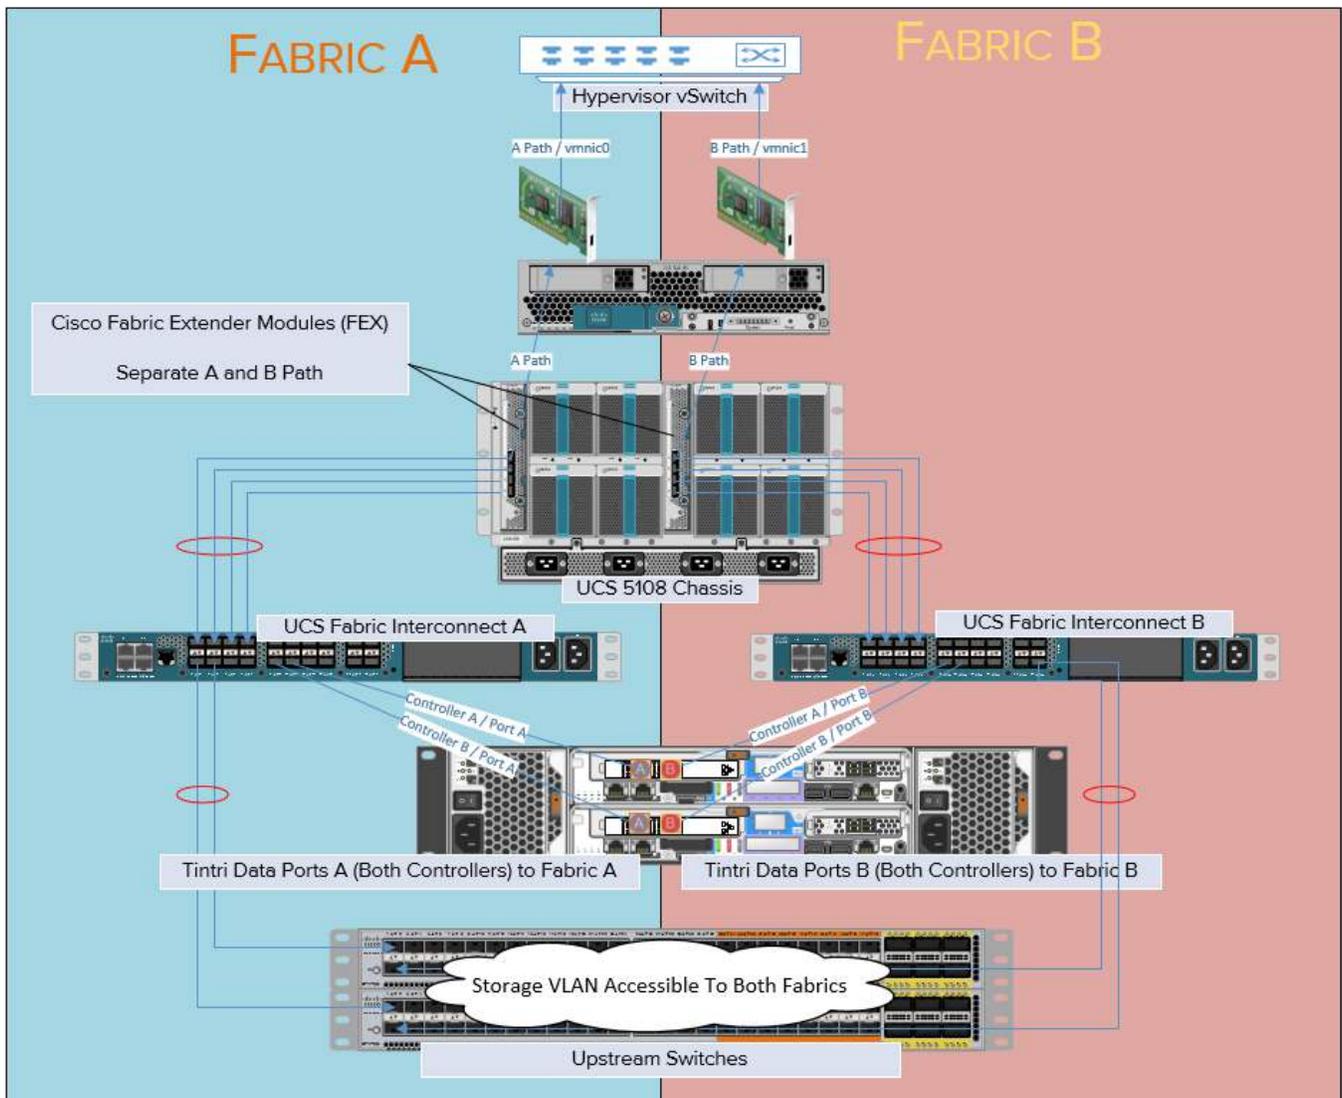


Figure 3 – Overview of connectivity for Fabrics A and B

VMstore NIC Teaming Behavior

When you configure a VMstore for use with UCS, connect it directly to the FIs, as shown in Figure 3 (which is covered in much greater detail later in this document). Leave **LACP disabled** (see Figure 22), which is the default state when data and admin ports operate as Active and Standby. Plug data port A into Fabric A, and data port B into Fabric B. **In a normal functioning state (link up), data port A is ALWAYS preferred as the active port.** If the link goes down, data port B will become active. **However, when data port A's link state returns, it becomes the active port again.**

DO: Unless you are using dedicated storage switches, leave LACP disabled. See [LACP](#) in the appendix for more detail.

With this VMstore behavior in mind, the design covered in this document sets a preference in the hypervisor hosts to favor Fabric A, which is where the VMstore will be active, regardless of whether it's running on controller A or controller B. The only time we'll end up with northbound traffic is in a failure scenario where the cable on data port A (active) becomes disconnected and data port B takes over. In this scenario, storage traffic will continue to flow uninterrupted through the upstream switch. **In all other scenarios (failure and good operational states), traffic remains in-fabric.**

*NOTE: Unfortunately, **LACP cannot be configured BETWEEN fabrics.** As such, do **NOT** use LACP unless you connect your VMstore to external IP storage switches instead of connecting them directly into the Fabric Interconnects. This is a valid and supported option, and is covered in more detail in a section on [LACP found within the Appendix.](#)*

Configuration

This section will cover configuration of the various technologies that make up the solution, including configuration changes which are performed within the respective management UIs: Cisco UCS Manager, VMware vCenter client and the Tintri VMstore Management UI. Physical cabling and upstream switch configuration are also covered in this section.

Cisco UCS Manager – Back Up Your Configuration

Before proceeding, it is strongly recommended that you make a backup of your existing configuration. This can be done within the **Cisco UCS Manager**, on the **Admin** tab. Ensure your filter is set to **ALL** and you should have a backup option, as shown in Figure 4. Refer to Cisco documentation for more information on backing up and restoring the UCS configuration.

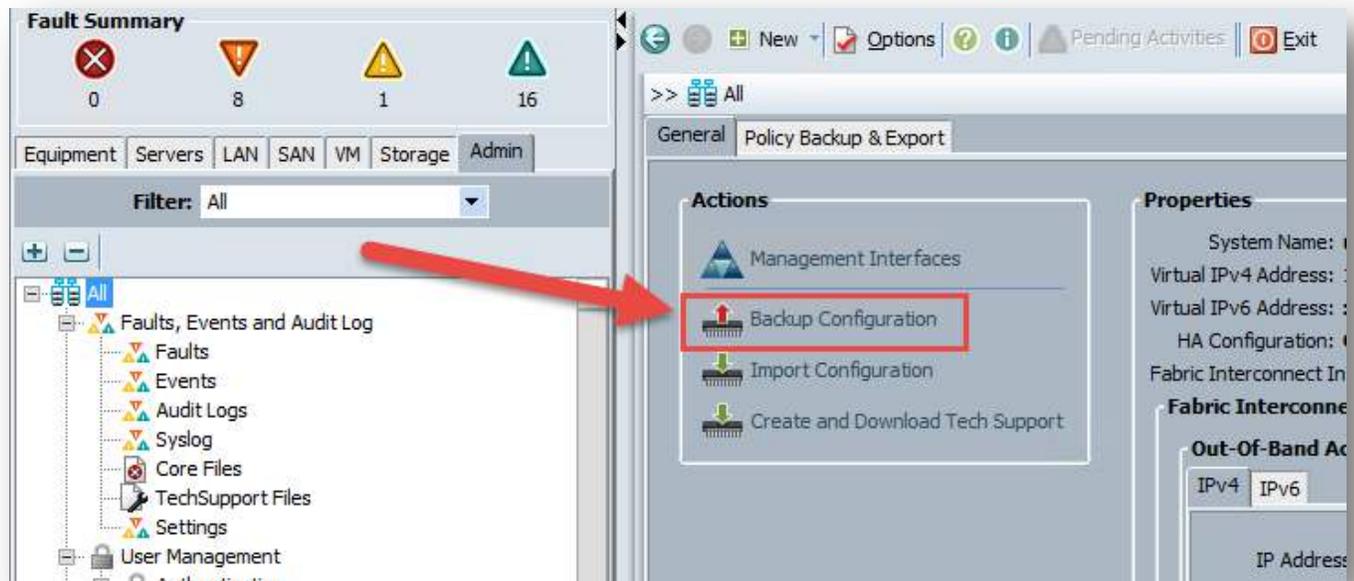


Figure 4 –UCS Manager – Configuration Backup

***DO:** Backup up your UCS configuration before making changes. Better safe than sorry!*

Cisco UCS – Default Maintenance Policies

The default maintenance policy in UCS Manager is to reboot servers **immediately**. Sooner or later, you'll end up making a change that seems trivial (such as adjusting a server profile template) and one or more servers will reboot without prompting. Tintri recommends you change the default maintenance policy to **User Ack**, shown in Figure 5. This option will prompt you to confirm each server is safe to reboot in advance.

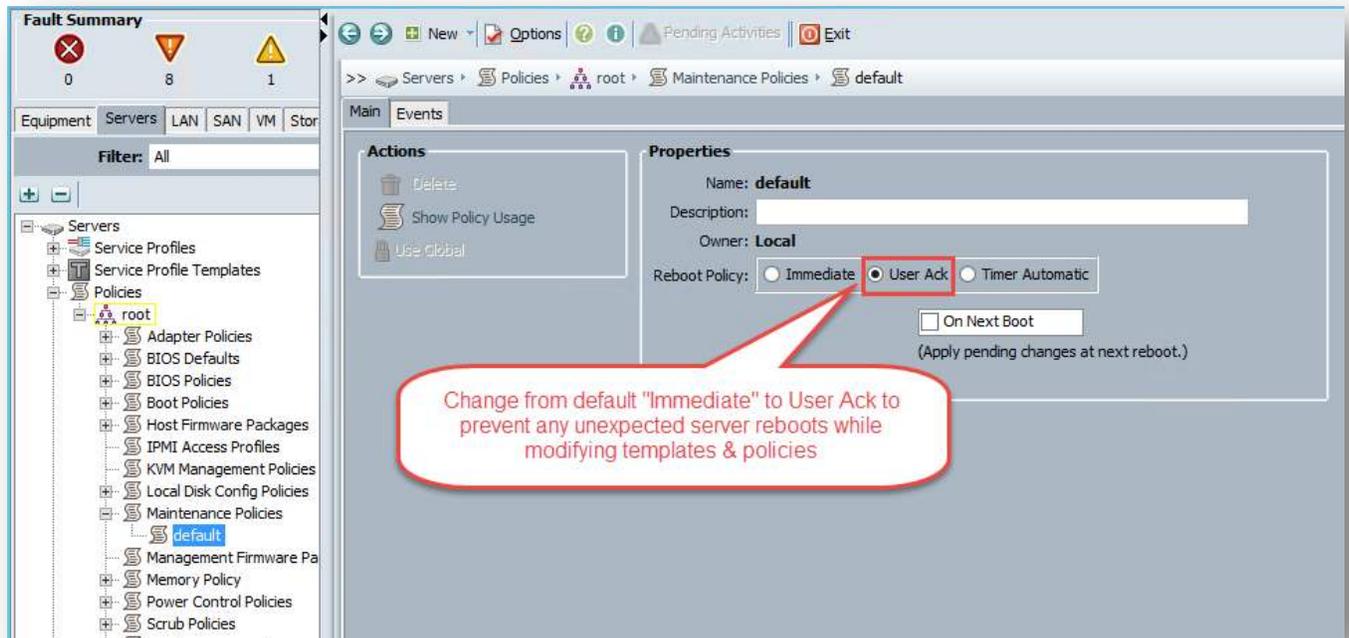


Figure 5 – Cisco UCS Manager - Default Maintenance Policy – Reboot Policy

DO: Within UCS Manager, change the Reboot Policy within the default maintenance policy from “Immediate” to “User Ack” to avoid unplanned downtime.

Cisco UCS – Configure Appliance Ports

The first step in adding a Tintri VMstore to your existing UCS environment is to configure the necessary UCS Fabric Interconnect (FI) ports on the UCS FI. Each VMstore will be connected directly to the FI.

Note: If you are connecting a T5080 with the quad-port 10GbE card option, you’ll need 8 FI ports. Consider connecting to dedicated IP Storage switches using LACP. Refer to Appendix A: Design Considerations for LACP connectivity options.

You will need four ports per VMstore being configured: two ports to connect to the A fabric and two ports to connect to the B fabric for each controller in a VMstore. All ports must be configured as “Appliance Ports”.

DO: For each VMstore being installed, configure four (4) available UCS FI (fabric interconnect) ports as appliance ports to connect to the VMstore data ports.

Before you are able to configure Appliance Ports, the Fabric must first be configured in “**End-Host Mode**” (versus *Switch Mode*). We’ve tested both modes successfully and recommend End-Host Mode, which is the default configuration for the FIs and provides more flexibility. Instructions provided within this document all imply that End-Host Mode. For more information on switching modes, refer to [Cisco’s Documentation](#).

To accomplish this using Cisco UCS Manager software:

1. Select the Equipment tab in the left-hand pane then expand the view as follows:

Fabric Interconnects → Fabric Interconnects A → Module X (fixed or expansion) → Ethernet Ports

2. In the right hand pane you will see the Ethernet Ports view as shown in Figure 6. Choose two unused Ethernet Ports from the list.

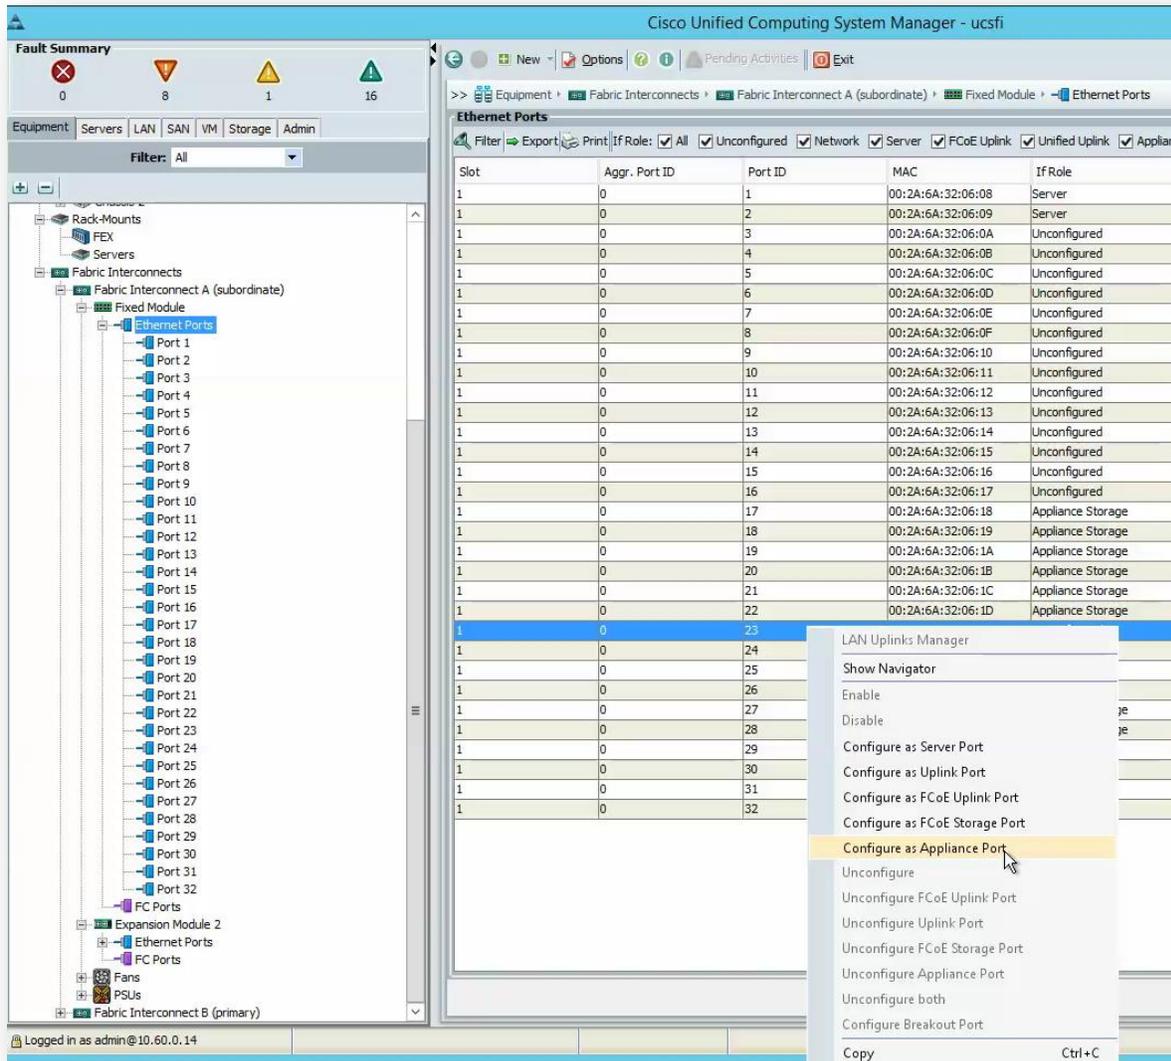


Figure 6 - Configuring an Ethernet Port as an Appliance Port in Cisco UCS Manager

3. For each port, access the context-sensitive menu, select “Configure as Appliance Port”, and complete the dialog using the following settings (See Figure 7):
 - a. Priority: Platinum (unless you know you have other traffic that must have priority over VMstore traffic)
 - b. If there is an existing VLAN you wish to use for storage, do so. Otherwise, for the first port ONLY, create a VLAN for use by Tintri network traffic:

- c. Select the “Create VLAN” button
- d. Give the VLAN a unique name (such as "Tintri_Storage") and unique VLAN ID (be sure and write them down for later reference)
- e. Select the check boxes for the VLANs you want VMstore to access including the newly created VLAN. (Selecting more than one VLAN provides greater flexibility for future configuration needs.)

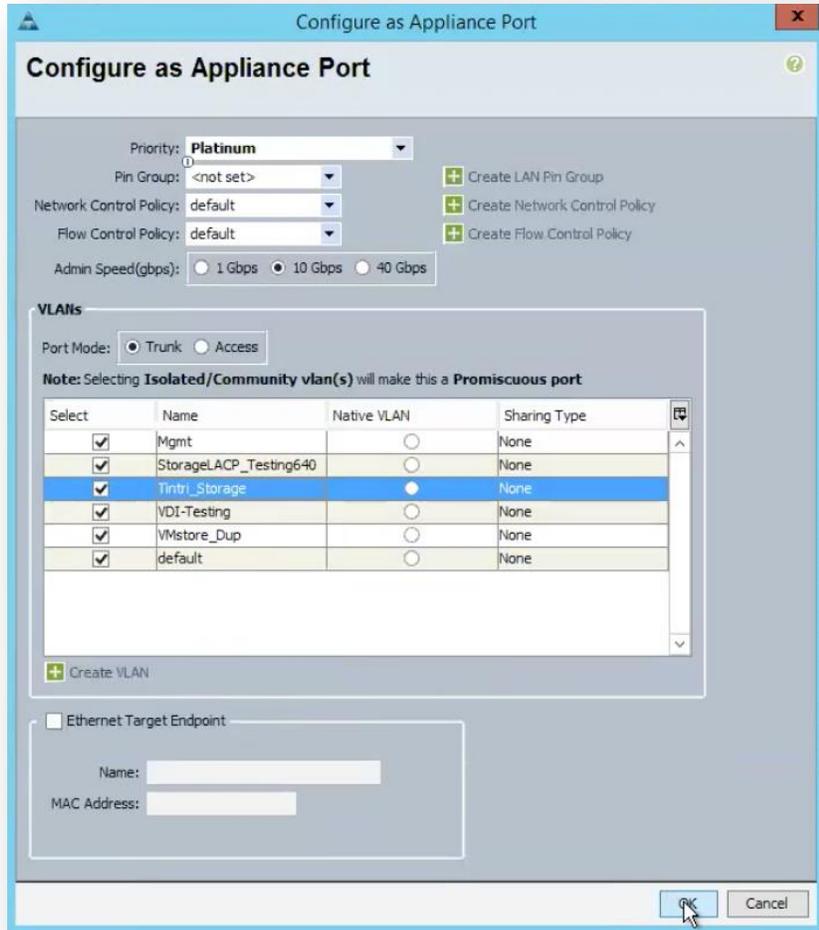


Figure 7 – Completed dialog for configuring an Appliance Port in Cisco UCS Manager.

- f. It is **NOT** necessary to select a native VLAN. Tintri recommends **not** doing so. Clear the radio buttons from the dialog and tag at the hypervisor and/or VMstore level. See Native VLANs for the explanation for this and additional details.
4. Complete the same steps shown above for the second port on Fabric A. It is not necessary to create the VLAN. It will already exist in the VLANs list as a result of creating it in the prior step.
5. Switch to Fabric B and repeat the process for the two ports on Fabric B. If possible, use the same port numbers as used on Fabric A. On the Equipment tab in the left-hand pane expand the view as follows:

Interconnects→Fabric Interconnects B→Module X (fixed or expansion)→Ethernet Ports

***Note:** Where possible, use the same port numbers on both Fabric A and Fabric B to simplify administration. For example, if Controller A – Data Port A plugs into Ethernet port 1/13 on FI A, use port 1/13 on FI B to connect Controller A – Data Port B. Repeat using ports 1/14 on both FI A and B to connect Controller B.*

Appliance Port Limitations – At the time of writing this, Cisco’s documentation states that there is a maximum of only 4 Appliance ports per fabric configurable on 6200 series fabrics in version 2.2.4 and onwards: http://www.cisco.com/c/en/us/td/docs/unified_computing/ucs/sw/configuration_limits/2-2/b_UCS_Configuration_Limits_2_2.html.

This would impose a limit of only 2 VMstores per UCS Domain (to service up to 160 hosts!). After several discussions with Cisco support, and validation in our own lab, we learned **there is no maximum limit of appliance ports**. The only limitation is the number of physical FI ports and FI port licenses you have available. We had 4 VMstores configured as appliance ports in our lab (10 total ports per fabric). While you can add as many appliance ports as you like, beyond several VMstores, you may want to consider dedicated IP Storage switches to connect your VMstores to.

Pin Group Configuration on Fabric Interconnects

There should be 2 pin Groups – One per fabric. Each pin Group should be configured to designate the port-channel that is used to uplink to the upstream switch. What this will do is mark the vNIC exposed to the host as down in the event the FI is isolated from the rest of the network. When vSphere detects the failure, the teaming rules we put in place for Active/Standby links will work as intended, and spare you from any unplanned downtime. For more information on Pin Groups, refer to Cisco’s documentation.

Network Configuration – Upstream Switches

Configure the upstream switches to allow VLANs to communicate between fabric A and fabric B. This can be done by configuring the ports the FIs are uplinked to as Trunk ports. As network switches vary widely from environment to environment, steps and terminology will be different and won’t be covered in this document.

Since Cisco Nexus 5K switches were in the lab test bed and are commonly used by many customers to connect their UCS FIs, a brief summary of what is required for Cisco Nexus is provided in **Appendix C - Nexus 5K Switch Configuration** for convenience.

Refer to Cisco documentation for full configuration details and best practices.

***DO:** Ensure that all VLANs used by vSphere hosts, especially Storage VLANs, are configured end-to-end through the upstream switches.*

While Tintri doesn’t endorse using Jumbo Frames, we do support their use. If Jumbo frames (MTU=9000) are to be configured, ensure that upstream switches have this configured for the storage VLAN throughout, **including inter-switch links!** See the **Jumbo Frames** section in the Appendix for more information and the design rationale.

Storage VLAN Configuration – LAN Cloud

The next step is to extend the Storage VLAN you created in the Appliance Cloud to make it available to the LAN Cloud. This will make the VLAN accessible to upstream switches and external servers, as well as

give you the option to allow your blade servers to access the VLAN. To accomplish this using Cisco UCS Manager:

1. Select the LAN tab in the left-hand pane then expand the view as follows:

LAN→LAN Cloud →VLANs

2. If the storage VLAN doesn't already exist, select "New VLAN".
3. In the dialog box that appears, use the same VLAN name and unique VLAN ID as before and click okay.

DO: Configure the Storage VLAN on the UCS LAN Cloud to provide access to upstream ports and servers

Connect the Tintri VMstore to Cisco UCS

Rack the VMstore

If you're installing a new VMstore, refer to the Txxxx Reference Guide for racking and initial setup guidance. A hard copy of the guide should be included in your VMstore packaging, but keep in mind that newer versions may be available. For the latest version of VMstore Txxxx Reference Guide, a digital copy can be downloaded from the support portal, as described in the **Tintri VMstore Guides** section.

The guide contains a Pre-Installation / Site Preparation section that will help you gather necessary IP addresses and other information ahead of time.

Cabling

With ports and VLANs configured, you are now ready to physically connect each new VMstore to the ports you just configured on the FI. The critical thing to note is that one port from each VMstore controller is connected to Fabric A and the other port is connected to Fabric B as illustrated in Figure 8.

DO: Connect the VMstore so that the "A" 10GbE data ports from each controller are connected to Fabric A, and the "B" 10 GbE data ports are connected to Fabric B (See Figure 8.)

Both the T800 and T5000 series VMstore systems ship with either Optical (SFP+) or Copper (RJ45) 10 GbE options. Be sure to specify the Optical (SFP+) option in order to connect your VMstore directly to the FIs.

DO: Choose the Optical (SFP+) 10 GbE adapter option when ordering a VMstore for use with UCS.

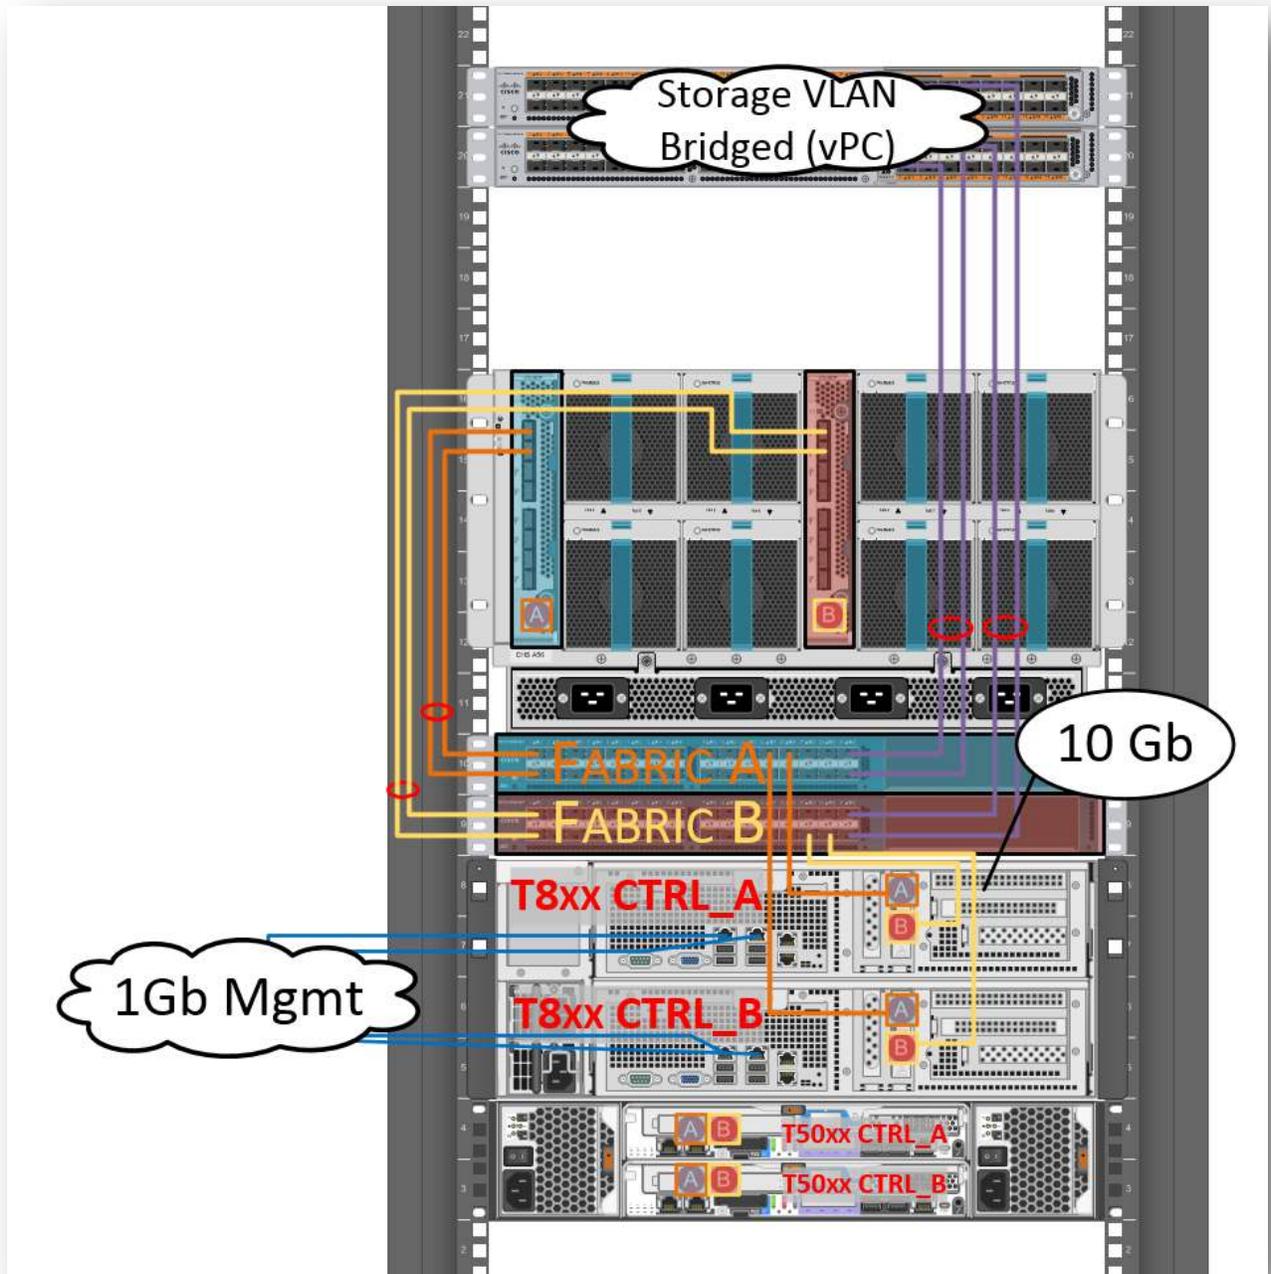


Figure 8– The recommended way to cable a VMstore to the UCS FI. T800 (cabled) and T5000 (uncabled) shown for reference

NOTE: The T850 and T880 VMstore models ship with 10 GbE data ports. The T820 VMstore has 1 GbE data ports with an optional upgrade to 10 GbE; the 10 GbE option is strongly recommended when connecting a T820 directly to Cisco UCS FIs.

When connecting the VMstore directly to the FIs, you can use Fiber cables with LC to LC connectors to connect between the supplied SFP+ transceiver in the VMstore and Cisco-branded SFP+ adapters in the FIs, or you can use TwinAx cables (passive – up to 5m). Tintri only supports Intel-based TwinAx cables that comply with SFF-8431 v4.1 and SFF-8472 v10.4 specifications. TwinAx cables don't have the same range as fiber cables, but are a much more economical option if you don't already have SFP+ transceivers for your FI ports.

To use TwinAx cables, simply remove the SFP+ transceivers that shipped with the VMstore and slide the built-in transceivers into the SFP+ ports until the cable clicks and locks into place. If the cable does not lock into place, turn it over and try again.

DO: Use TwinAx cables OR fiber cables with compatible SFP+ transceivers and LC to LC connectors to connect the VMstore to the FIs.

Configure the VMstore

DO: Configure your VMstore according to the instructions in the Tintri Reference Guide.

With the physical cabling complete, the next step is to configure the VMstore.

Configure the Admin Network

The first step is to connect to the VMstore console and configure the admin network. This is explained in the Reference Guide.

Complete the “Out-of-the-Box” Configuration Process

Next, you will complete the web-based “out-of-the-box” configuration. When you configure the Data IP, be sure and configure it on the storage VLAN you are using for the VMstore.

Note: As a general rule, there's no need to configure multiple IP addresses and multiple VLANs for a VMstore. Multiple IP addresses may be used if you're connecting a VMstore to multiple, separate environments, such as in a Multi-tenant environment for a cloud provider or large enterprise.

You will also be asked to configure the following items, as shown in Figure 9.

- vCenter(s) this VMstore will connect to.
- Contact(s) for alerts and SMTP information
- NTP server (or set date and time)
- Primary and secondary DNS servers
- Support contact
- New password

Initial setup

How this physical VMstore appears to hypervisor managers:

Data IP: *

Netmask: *

Gateway:

VLAN id:

How this VMstore talks to hypervisor manager:

vCenter RHEV-M Hyper-V OpenStack

vCenter:

Username:

Password:

How you receive alerts from this VMstore:

To: *

Outgoing SMTP Host: *

Date and time on this VMstore:

Time zone: *

Primary NTP: *

Secondary NTP:

[Set the clock manually.](#)

How this VMstore resolves hostnames:

Primary DNS: *

Secondary DNS:

How Tintri support contacts you:

Automatically send Tintri diagnostic reports (highly recommended.)

Contact name:

Email:

Phone:

VMstore location:

New password for managing this VMstore:

New password: *

Confirm password: *

Set up my VMstore

You can change these settings later.

Figure 9— Configuration wizard that is shown at first login for a new VMstore installation only

Configure the Hosts

DO: Modify your UCS Service Profile Template(s) or UCS Service Profile(s) to allow access to the newly created storage VLAN.

DO: On every vSphere Host, configure a new vmkernel port on the same subnet & VLAN as the VMstore.

DO: On every vSphere Host, mount a new datastore. Tintri recommends configuring one datastore per VMstore.

Configure UCS Servers

The next step is to make sure that the desired set of UCS hosts can reach the VMstore via the VLAN created or chosen earlier. This is accomplished through UCS Manager by updating either appropriate Service Profile Template(s) or Service Profile(s). This procedure is for updating a Service Profile Template.

At a minimum you need connectivity on both fabrics to the storage vLAN. If you only have Tintri storage, the Tintri recommends a simple configuration using a single vNIC on each fabric for a total of 2 vNICs per host. The following figures show how these vNICs should be configured.

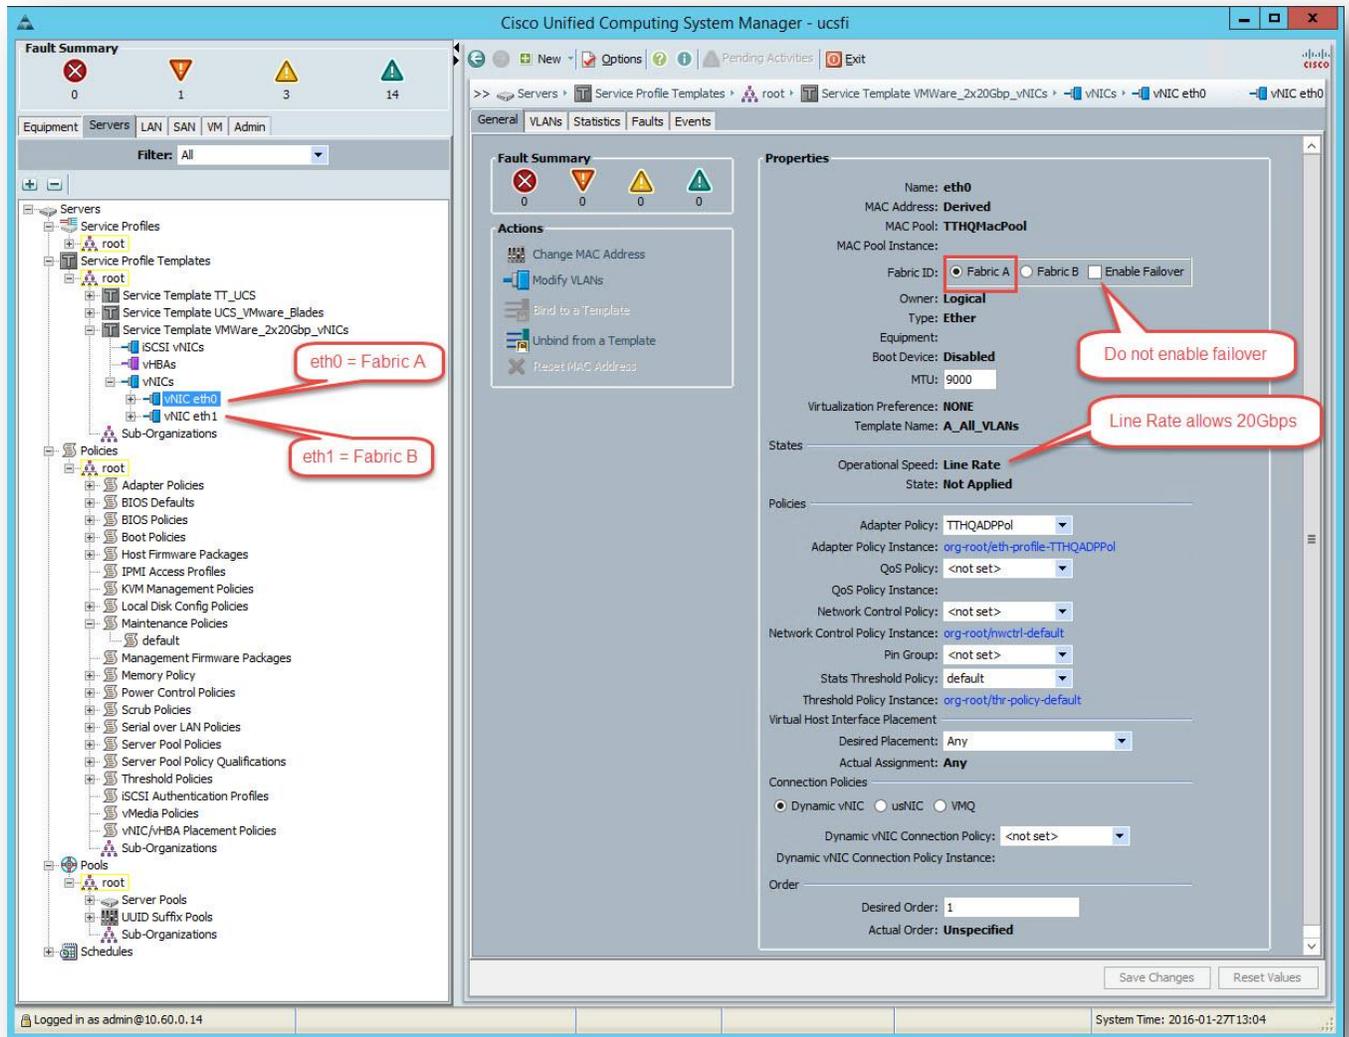


Figure 10. Service Profile Template showing vNIC configuration for Fabric A.

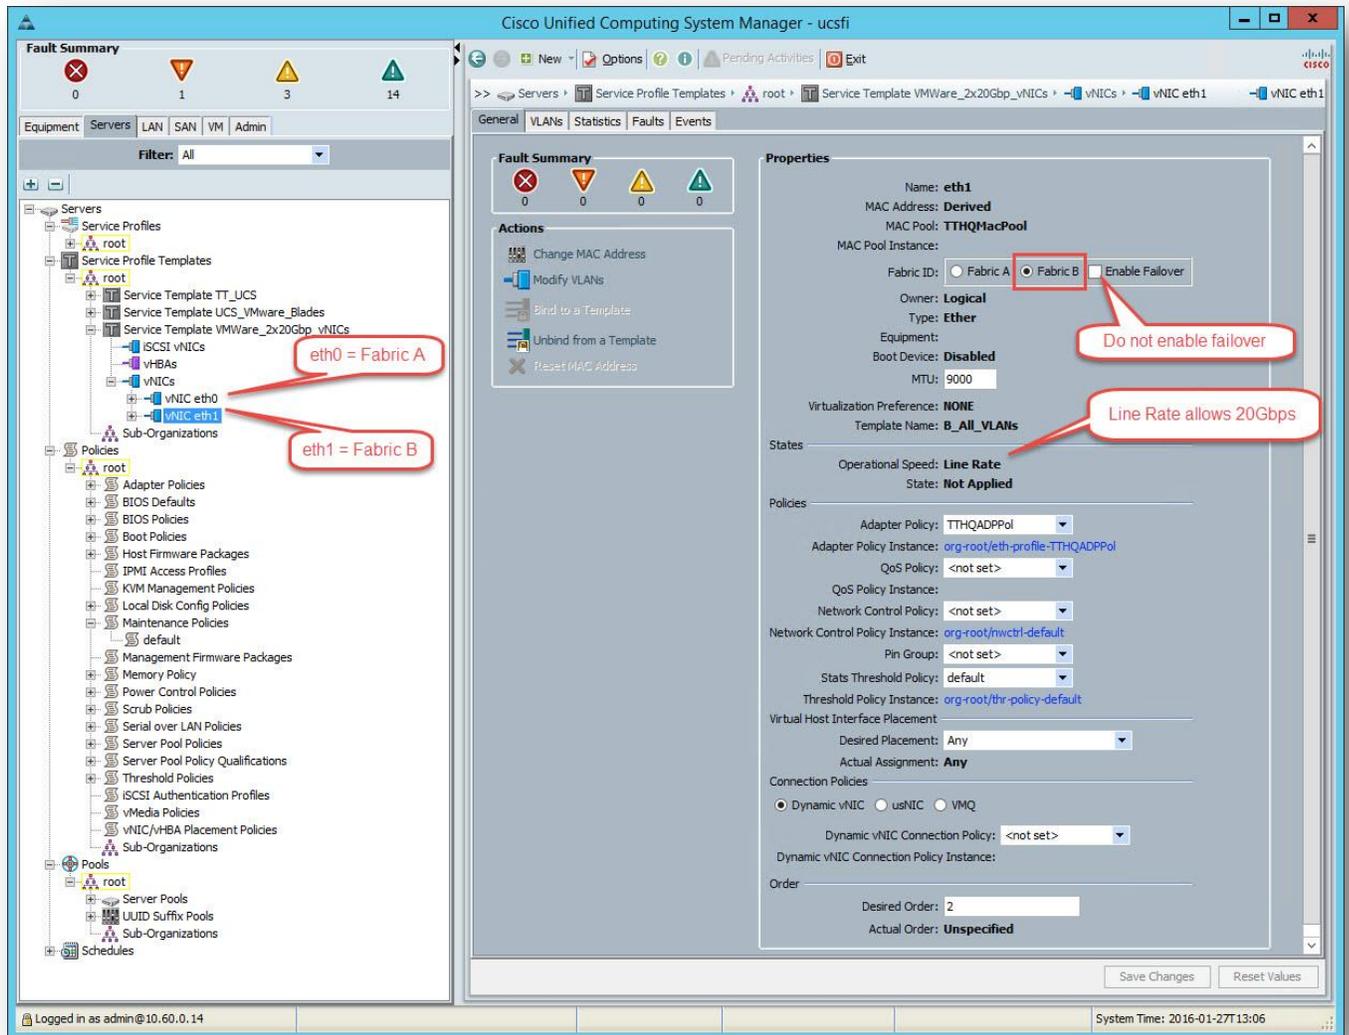


Figure 11 - Service Profile Template showing vNIC configuration for Fabric B.

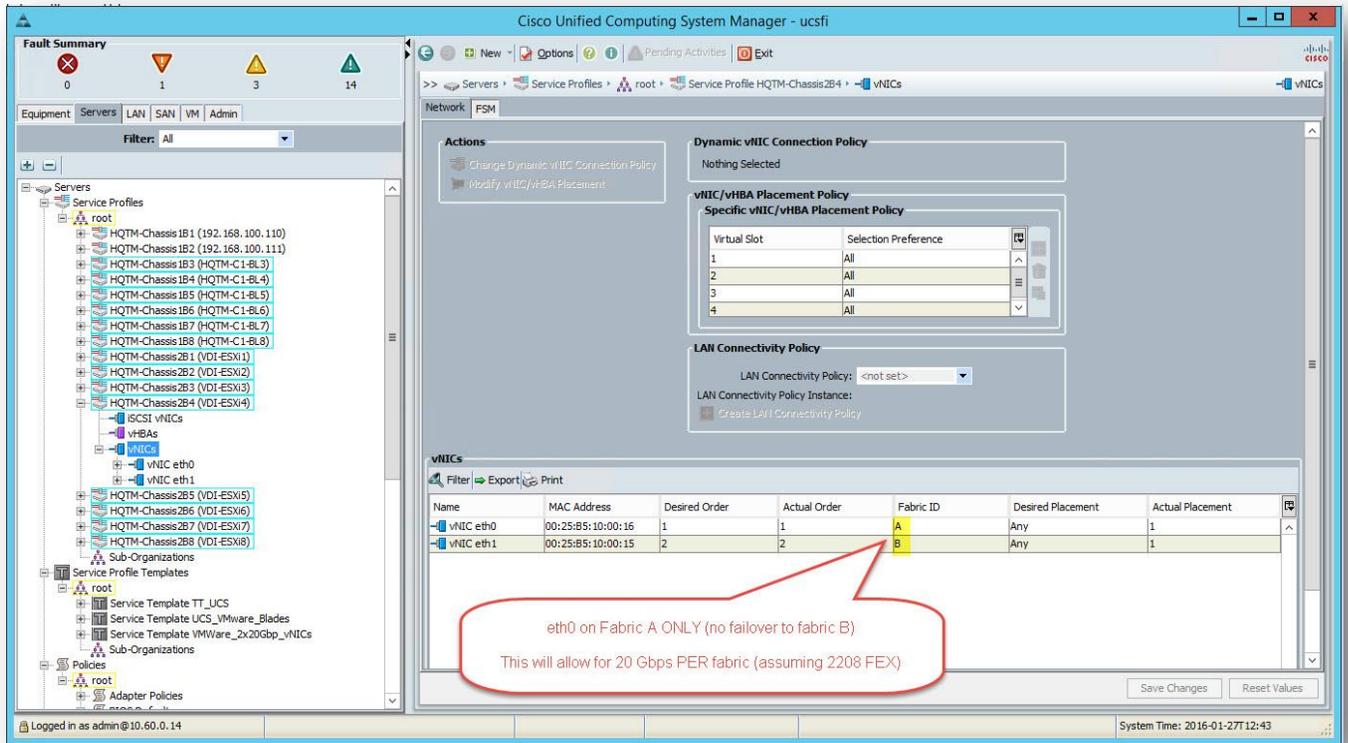


Figure 12 - Service Profile showing vNICs.

Once vNICs have been configured, use the following procedure to add the vNIC to the VLAN created previously.

1. In UCS Manager, select the Servers tab in the left-hand pane and then expand the view as follows for vNIC eth0:

Servers → Service Profile Templates → root → "Your Target Template" → vNICs → vNIC eth0

Where "Your Target Template" is the name of the template you want to modify.

2. Select "Modify VLANS" from the Actions menu and select the checkbox for the VLAN you created above, as shown in Figure 13

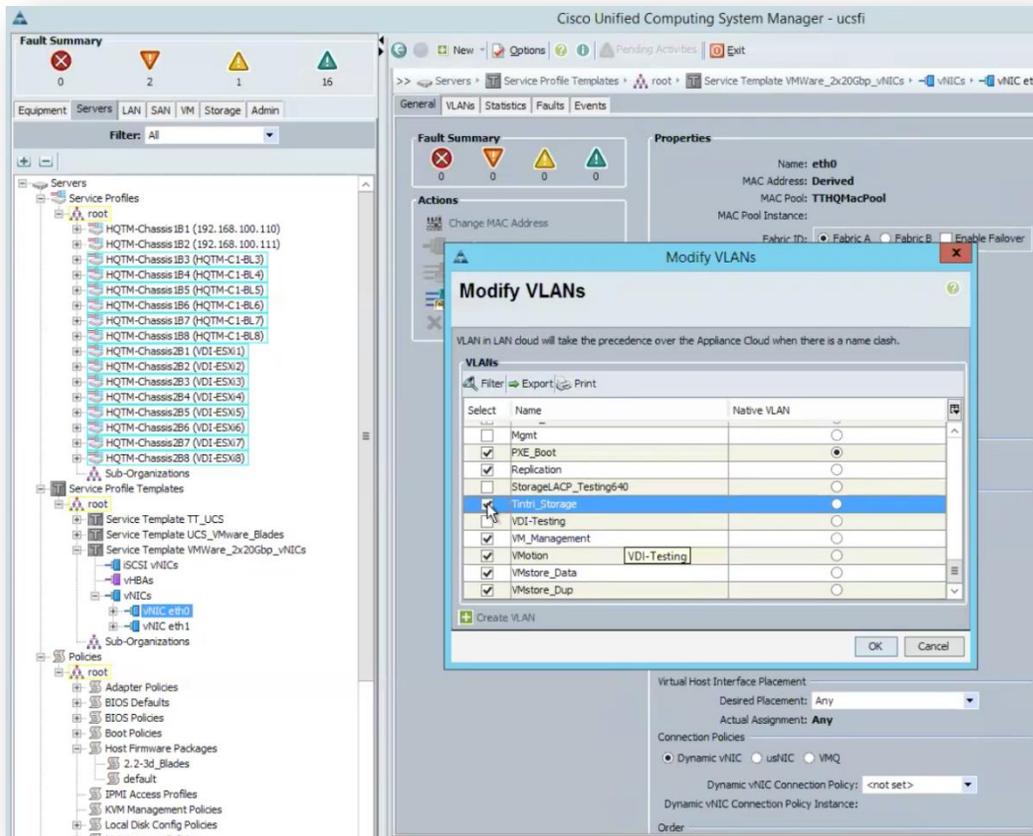


Figure 13 – Adding the Tintri VLAN to a UCS Service Profile

3. Repeat the process for vNIC eth1 by expanding the view in the left-hand pane as follows:

Servers→Service Profile Templates→root→"Your Target Template"→vNICs→vNIC eth1

4. Repeat the above process for other templates if needed.

vSphere Host Configuration

DO: Configure your vSphere hosts to use either a distributed virtual switch (dvSwitch) or vStandard switch (vSwitch).

The next step is to configure vSphere hosts using the vSphere web client or the legacy vSphere windows client. If you already have an existing storage VLAN, you can use it without configuring a new vmkernel. Otherwise, you must configure vSphere hosts to use either dvSwitch or vStandard switching. The procedure for using dvSwitch is shown below. If you are using vStandard switches, refer to Appendix B – vStandard Switch Configuration.

Configure a New Distributed Port Group for a dvSwitch

1. Login to the web client.
2. From the Navigator pane on the left, select "Networks."
3. Expand the view as follows for dvSwitch configuration:

“top-level”→”data center”→dvSwitch

4. From the context-sensitive menu for dvSwitch select:

Distributed Port Group→New Distributed Port Group

5. Give the new Distributed Port Group a unique name (Such as “Storage”).
6. Configure the settings as shown in Figure 14. Be sure and use the VLAN ID for the VLAN you created previously.

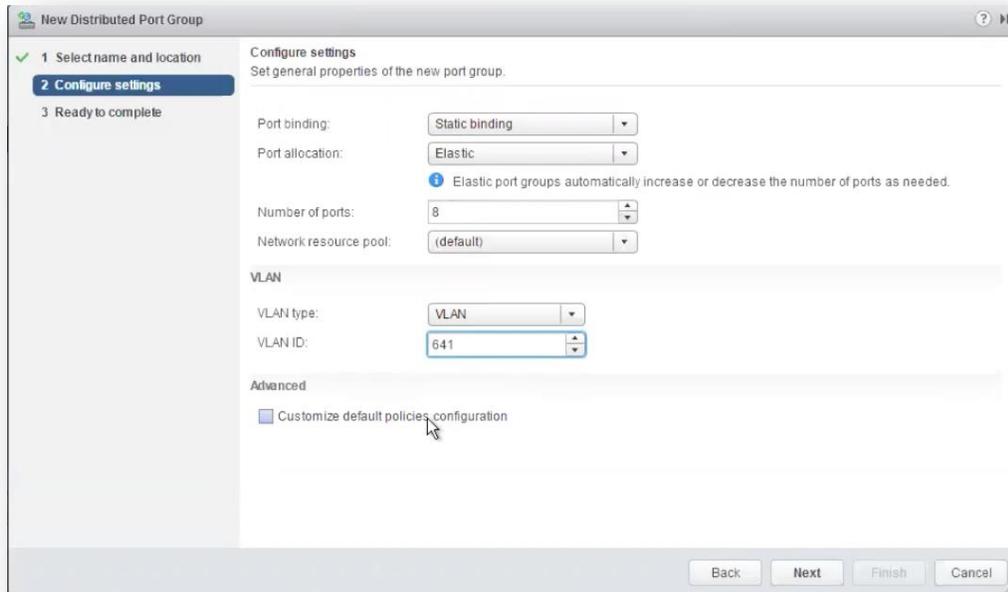


Figure 14 – Completed dialog for New Distributed Port Group when using a dvSwitch.

7. Return to dvSwitch and select “Distributed Port Groups”
8. On the pane that appears, locate the New Port Group you just created.
9. From the context-sensitive menu for that Port Group, select “Edit Settings...”
10. Select “Teaming and failover” from the left hand menu.
11. Configure the uplinks so that dvUplink10G_1 is active and dvUplink10G_2 is standby.
12. Configure any other uplinks so that they are unused as shown in the left side of Figure 15 (full client) and Figure 16 (web client).

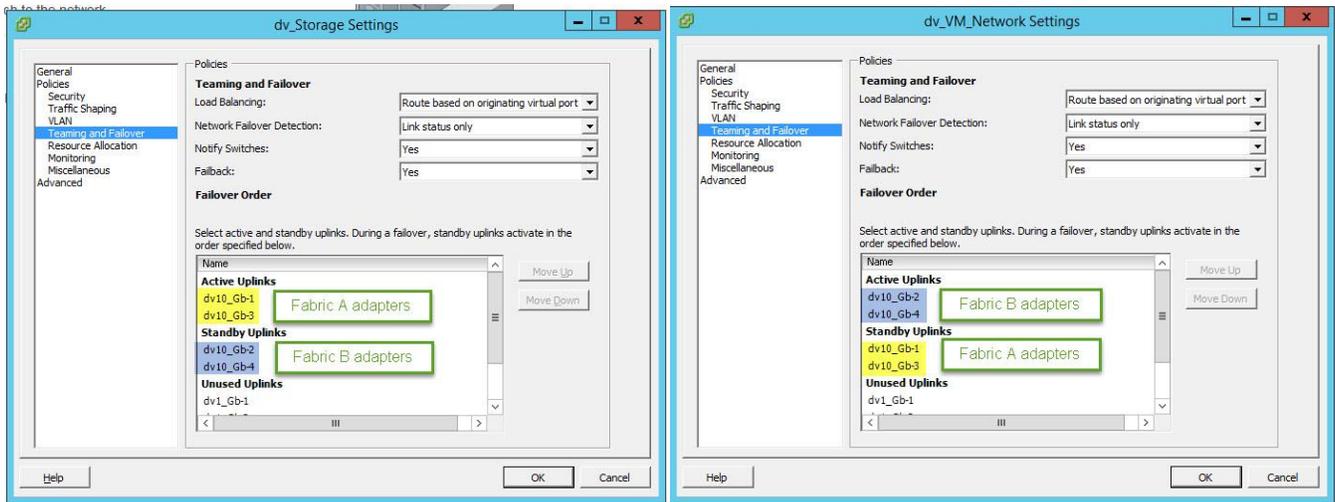


Figure 15 – Teaming configuration of the network used for VM networking (commonly referred to as “VM Network” or “Production Network”) versus teaming setup of the network used for Storage. Note that VM network uses Fabric B while storage uses Fabric A.

This ensures that the Tintri NICs on Fabric A are active and the NICs on Fabric B are on standby to be used if a failure occurs.

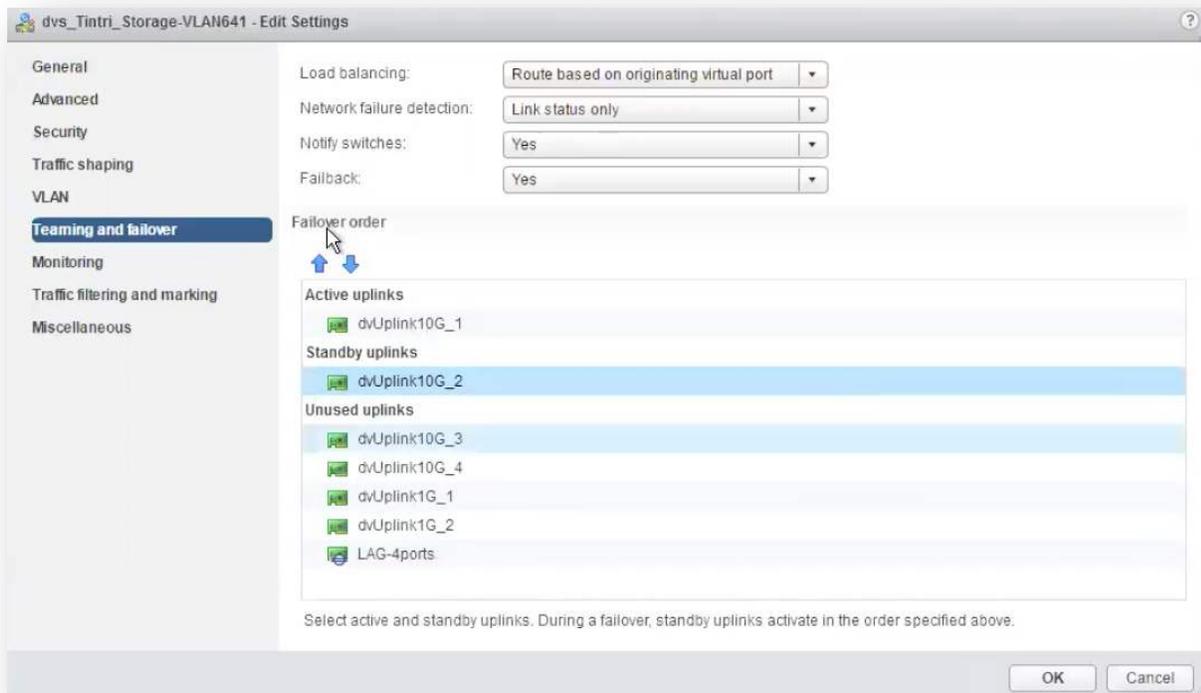


Figure 16 – Storage Teaming configuration using a dvSwitch, as seen in vSphere Web Client

Configure a VMkernel port

The next step when using a dvSwitch is to configure a new VMkernel adapter and Datastore. An overview of the effective Active/Standby paths is shown in Figure 17. The goal is to configure storage traffic on Fabric A and all other traffic on Fabric B by default.

View: vSphere Standard Switch vSphere Distributed Switch

Networking

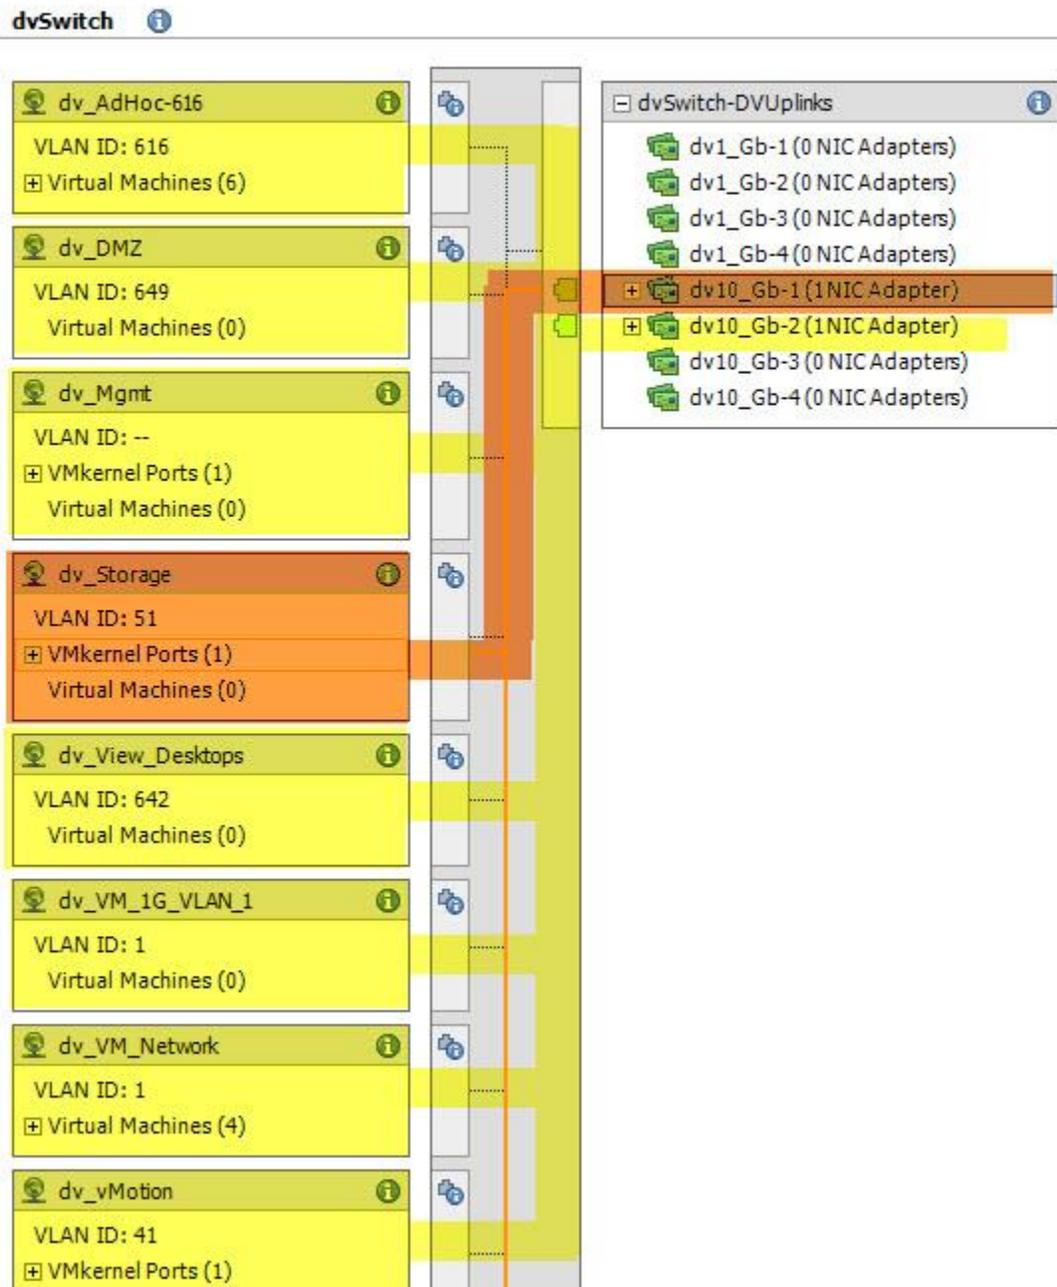


Figure 17 – Overview of dvSwitch. Orange shading represents Fabric A data path as active for Storage, Yellow represents Fabric B (all other networks)

This can be accomplished via the vSphere Web Client or Windows Client. Details for both are as follows:

From the vSphere Web Client:

1. Select the first host in the cluster of interest.
2. Select the “Manage” then “Networking” tabs in the right hand pane and then select “VMkernel adapters”.

3. Select the “+” icon above the list of VMkernel adapters to create a new adapter and complete the wizard.
 - a. Set the connection type to “VMkernel Network Adapter”.
 - b. For target device, choose the radio button for “Select an existing network”.(See Figure 18)

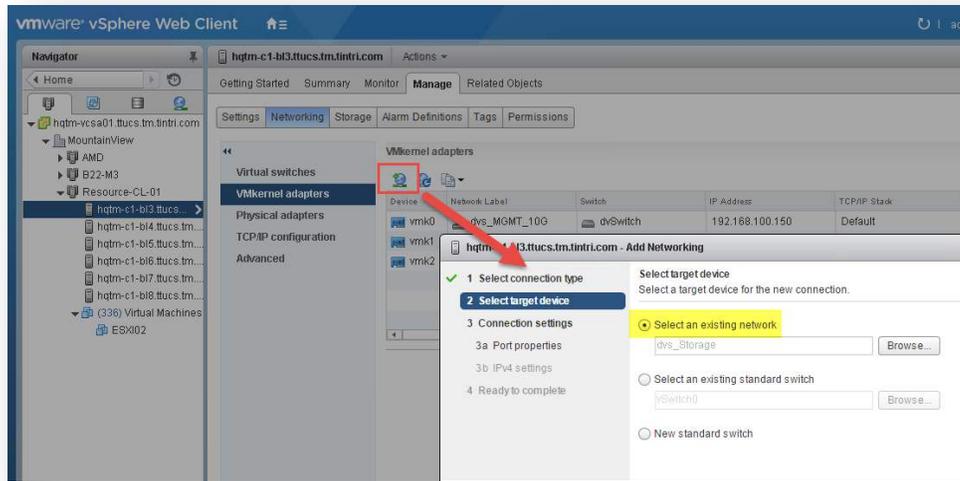


Figure 18 - Creating a new VMkernel port in the vSphere Web Client

- c. Browse the list and select the port group you created in “Configure a New Distributed Port Group for a dvSwitch” above (ie Storage).
 - d. Give the host a unique IP address on the same subnet as the data network of the VMstore.
 - e. Select Finish.
4. Repeat for each host in the cluster, as well as additional hosts you want to connect to the VMstore.

From the vSphere Windows client:

1. Select the first host in the cluster of interest.
2. Select the “Configuration” tab in the right hand pane and then select “Networking”.
3. Change the view to “vSphere Distributed Switch” and then select “Manage Virtual Adapters...”.
 - a. Click “Add” in the Manage Virtual Adapters dialog, as seen in Figure 19:

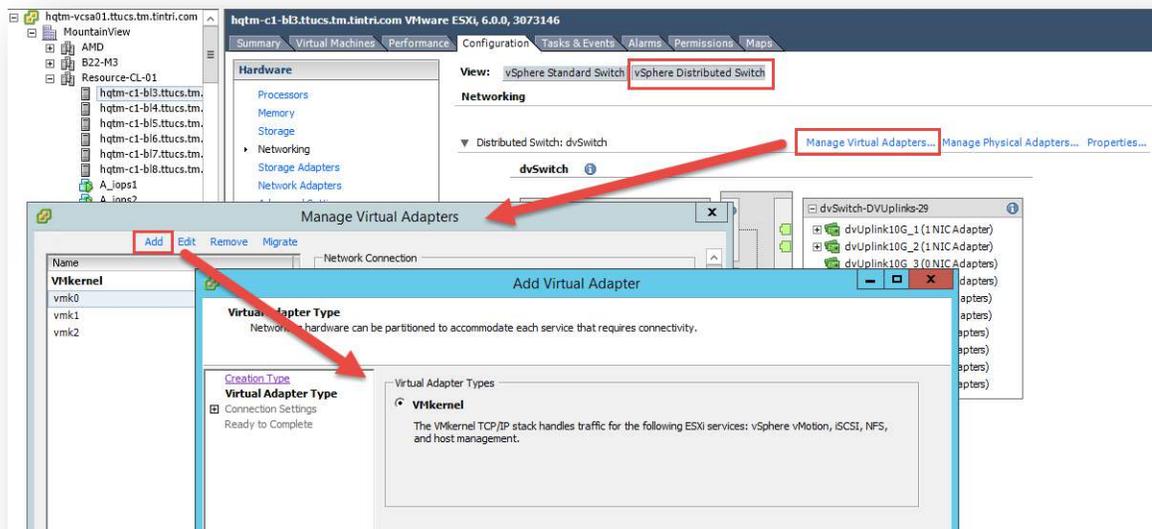


Figure 19 - Creating a new VMkernel port in the vSphere Windows Client

- b. Choose “New virtual adapter”, then click “Next”.
 - c. Select “VMkernel” as the adapter type and click “Next”.
 - d. Select the port group you created earlier (ie. Storage) and leave the checkboxes for vMotion, FT logging and Management clear, then click “Next”.
 - e. Give the host a unique IP address on the same subnet as the data network of the VMstore, and provide the Subnet Mask, then click “Next” and “Finish”.
4. Repeat for each host in the cluster, as well as additional hosts you want to connect to the VMstore.

Mount the Datastore

The vSphere Web Client is the quickest and most efficient ways to add a datastore. It can be used to mount the datastore to all your hosts in the same vCenter in one shot, assuming each host has been configured with a VMkernel port and IP (previous step) on the same network as the VMstore data IP.

1. From the Navigator pane on the left, select “Hosts and Clusters.”
2. Right-click on the Datacenter, Cluster, or Host on the right. From the context-sensitive menu select:

Storage → New Datastore...

3. Complete the pop-up wizard:
 - a. Type: NFS datastore
 - b. Select NFS version 3 when prompted.
 - c. Name and configuration. (See Figure 20 for an example of the dialog.)
 - i. Fill in the a name for your data store
 - ii. Specify the folder: **/tintri** or **/tintri/submount**
 - iii. Specify the data IP of the VMstore
 - iv. Click Next

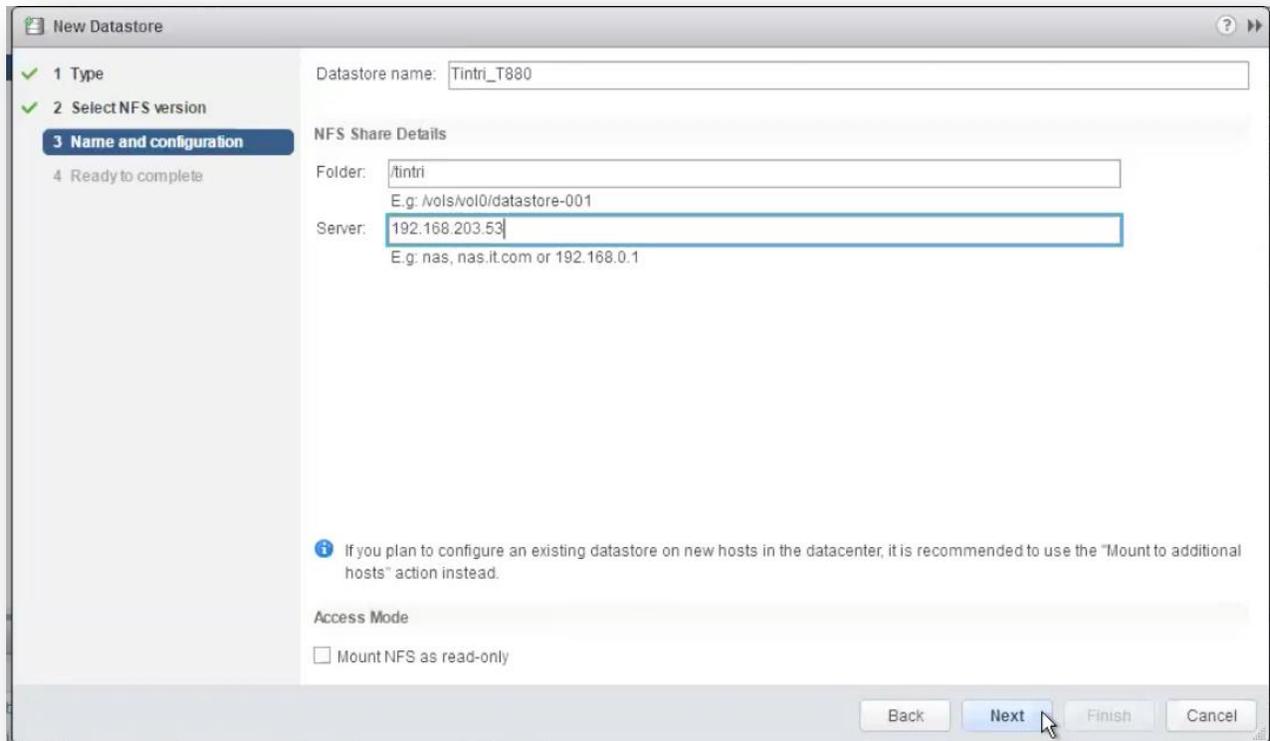


Figure 20 – The New Datastore dialog.

4. If you chose “New Datastore” from the A cluster or the whole data center, you will be prompted for “Host accessibility”. During the host accessibility step, select every host you wish to be able to access the new VMstore and then complete the wizard. (See Figure 21)

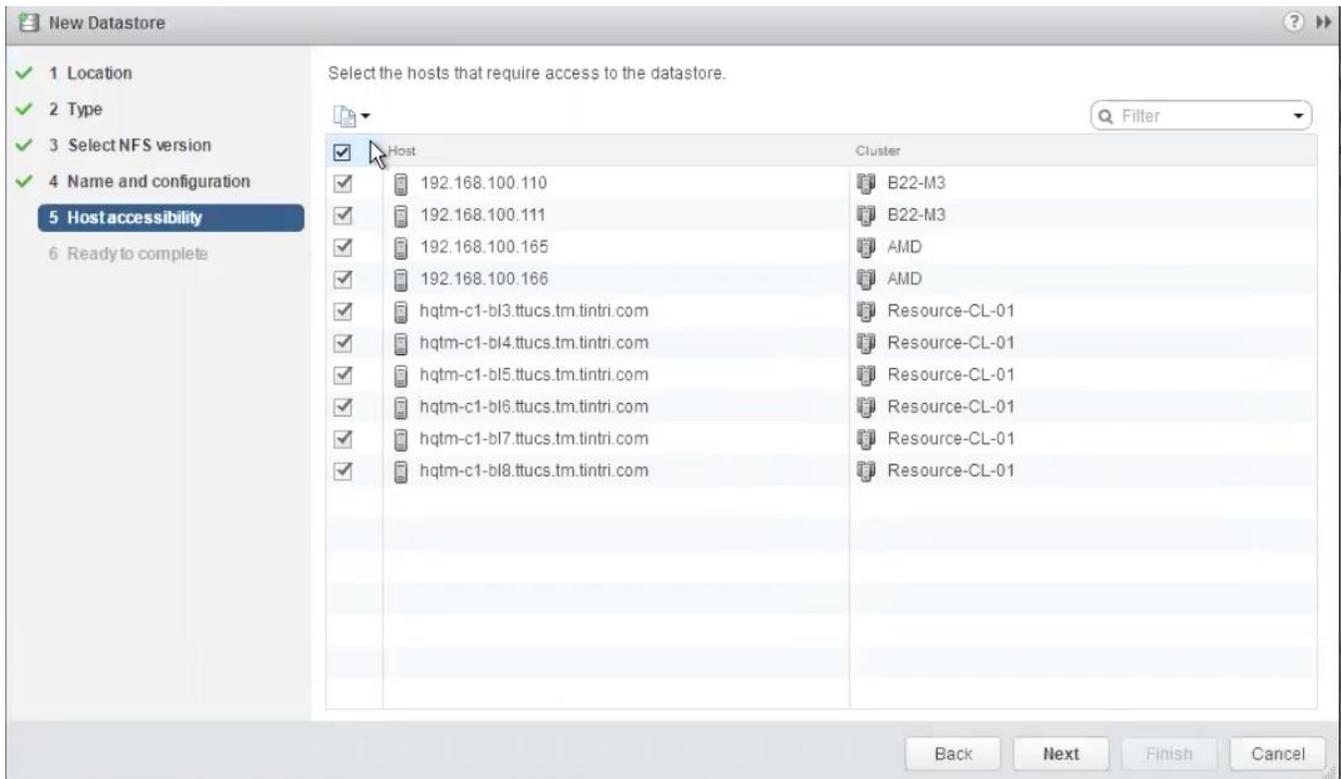


Figure 21 – Selecting multiple hosts in the New Datastore wizard.

5. Click Finish.

Congratulations, your VMstore is now ready to host VMs!

Test Redundancy

IMPORTANT. Before proceeding to deploying production workloads onto your newly provisioned VMstore, we **strongly** recommended that you take this time to test the full redundancy of the system as a whole. Details on components to test can be found below in **Appendix D – Failure Scenarios & Testing Redundancy.**

Migrate VMs to the VMstore

The VMstore should now be fully configured and ready for use. Create new VMs in the Tintri datastore, and/or use Storage vMotion to migrate existing VMs to the new storage system. An optional final step is installation of the Tintri VAAI provider.

Install the VAAI Plugin

DO: Install the Tintri VAAI plugin on each vSphere server.

Tintri’s VAAI provider for vSphere is available free to customers, and is downloadable from the Tintri support site. VAAI allows vSphere (ESX/ESXi) servers to offload VM cloning operations to Tintri VMstore, creating space-efficient “Tintri Clones”. The speed with which Tintri VMstore can deploy new VMs from a selected VM or template (a process referred to as cloning) is significantly faster than the software-only

cloning built into vSphere, and does not consume additional host resources during cloning. Additionally, by using Tintri's VAAI Plug-in, overall VMstore capacity consumption is reduced, allowing you to fit more VMs on the same VMstore with higher space efficiency.

Refer to the [Tintri VMstore with VMware® Best Practice Guide](#) for installation instructions.

Tintri Upgrades

Once fully configured as described in this document, Tintri upgrades can be installed simply and seamlessly without any downtime to VMs. It is strongly recommended that you run the latest GA code available from the [Tintri Support Portal](#).

IMPORTANT. Prior to installing an upgrade for the first time on your newly deployed VMstore in a new UCS deployment, or after any significant network or topology modifications, **we strongly recommend that you perform a manual failover BEFORE upgrading your VMstore** to ensure there is no loss of connectivity, preferably during a maintenance window. This will ensure a smooth failover during the upgrade without the risk of having two controllers operating on different versions of Tintri OS. Once this has been done once, you can be confident this upgrade, and all future upgrades will install smoothly.

After you've tested the failover, you can proceed with upgrading your VMstore. This can be done by navigating to the Hardware tab within the VMstore UI, confirm the VMstore is in a Redundant (good) state, and then click "Failover". Upgrades can also be performed in an automated fashion via a PowerShell script. For additional information on the upgrade procedure, refer to the Tintri VMstore Administration manual, found within Help on every VMstore.

Support

Solution support is available direct through Cisco and Tintri.

To find options for contacting Tintri support go to www.tintri.com/support/contact-support. There you will find a link to the support portal (valid login required) and the latest phone numbers and other options for contacting Tintri from anywhere in the world.

For Cisco, there are multiple ways to open a TAC Service Request. Note: Severity 1 (S1) and Severity 2 (S2) cases MUST be opened by telephone

- **Via the Web:** Request cases quickly using Cisco Support Case Manager. In order to open a case online, the contract must be properly associated to the CCO ID being used to make the request <https://tools.cisco.com/ServiceRequestTool/scm/mgmt/case>
- **Via the Phone:** For US and Canada dial - 1 800 553 2447
- **Via Email:** Please include all relevant information (CCO ID, Contract, Serial Number and problem description) to ensure timely and proper routing of your case. Send email requests to: tac@cisco.com

Conclusion

The Tintri VMstore offers performance and manageability features that make it an excellent choice for use in conjunction with Cisco UCS and VMware vSphere. Following the best practices described in this guide will help you get the most from your deployment. The key to success is proper configuration of UCS Fabric A for preferential use by storage traffic and UCS Fabric B for preferential use by other VM traffic as illustrated in Figure 3 – Overview of connectivity for Fabrics A and B at the beginning of this document.

We hope the information in this best practice guide helps you get the most from your Tintri VMstore environment. Your feedback is valuable to us! If you have any improvement suggestions for this document, please reach out and let us know - Thank you.

Appendix A: Design Considerations

LACP

Link Aggregation Control Protocol (LACP) is a useful technology that allows us to utilize more than a single 10 Gbps data port to meet the performance needs of demanding virtualized workloads. While any single host will not exceed 10 Gbps via LACP, multiple hosts can. Unfortunately Cisco UCS fabrics don't allow us to bond multiple ports with LACP between Fabric Interconnects; as such we don't recommend configuring LACP when directly connecting a VMstore into Cisco FIs.

While it is technologically possible to create an LACP channel when all data ports from the same controller are connected to one FI, and the data ports of the other controller are connected to the other FI, it is not recommended. This section covers the rationale behind our recommendation, as well as a recommendation on how to go about setting up LACP in a more robust way using dedicated IP storage switches.

Rationale against configuring LACP within an FI

The ESX hosts will originally use the same FI, keeping all traffic "in-Fabric", which is fine when the active controller of the VMstore is connected to that same fabric. As soon as there is a controller failover to the other FI, this will no longer be the case. The hosts will not be aware of that change and will continue to use their same preferred fabric, and all traffic will then travel inefficiently northbound through the uplinked 10Gb switch and back down into the other Fabric, which may saturate the uplinks to the core switch and contend with VM networking traffic.

Controllers fail over as part of the upgrade process, which makes the above scenario inevitable. The upgrade process can be performed online and the failover is seamless to the hosts, however the path for traffic will change. When connected as described in this doc, both controllers will favor data ports A, which connect to the same fabric. Keeping in mind that the scenario where LACP is desirable is when we want to accommodate high bandwidth workloads, this is the worst case scenario when it comes to possibly saturating the uplink between the FIs and the upstream switches.

Another scenario where configuring LACP within the same Fabric can have an adverse effect is if a FEX fails in a UCS blade chassis, in which case the VMstore would not know to switch to the other controller in order to keep traffic in-fabric and avoid sending traffic northbound. In a multi-chassis and multi-VMstore scenario, you can imagine how it becomes increasingly more difficult to coordinate all traffic to remain in-fabric when controller failovers on any VMstore or FEX failure on any chassis increase in probability.

Configuring LACP using dedicated IP storage switches

In the case where you desire the use LACP due to demanding workloads, we recommend you use a dedicated IP storage switch, such as Cisco 7K or 9K series switches. This applies to using two x 10Gb data ports per controller for most VMstore models, as well as when using the quad-port 10Gb adapter option in our T5080 All-Flash VMstores. Using dedicated IP storage switches is also desirable when you have a number of hosts you want to connect to the VMstore that are not directly attached to the FIs, such as rack-mount servers with dedicated NICs connected into the storage network.

When configuring the VMstore to connect via dedicated storage switches, it is preferable to use a distributed LACP configuration (ie. vpc for Cisco switches) that allows each controller to connect to multiple switches. For additional information on configuring LACP in this manner, refer to the LACP section within our [Tintri VMstore with VMware Best Practices Guide](#).

After the switch ports are configured as an LACP port channel, enable LACP on the VMstore from within the VMstore UI. The option can be found within Settings – More – LACP, as shown in Figure 22.

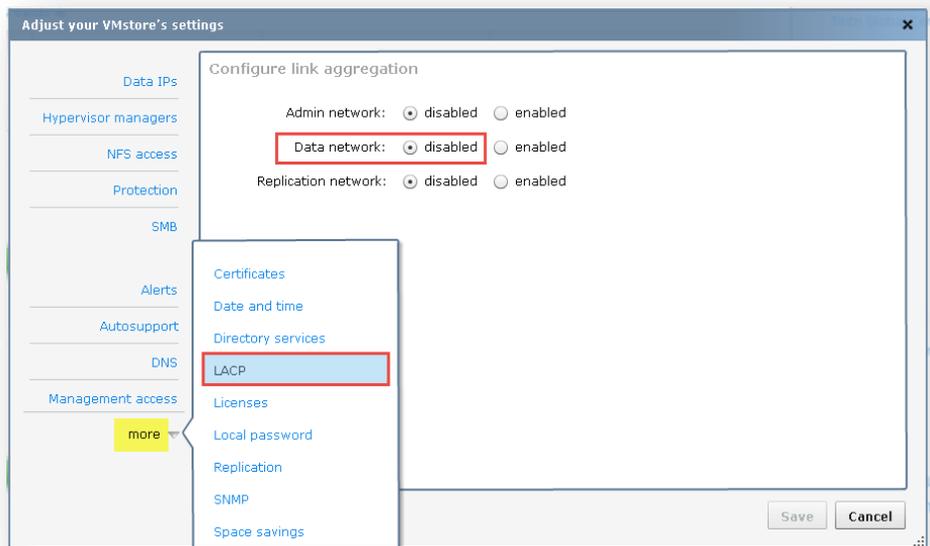


Figure 22 - LACP can be configured under "Settings - more - LACP" from within the VMstore UI

Native VLANs

Any time you create a Native VLAN, it can only be accessed by referring to it as VLAN 0 and not tagging it. This creates the potential for confusion when looking at it from the hosts section of UCS Manager, the Tintri interface, or vSphere because the rest of the world knows it by a different name. Tintri recommends against designating a VLAN as a native VLAN for this reason.

Jumbo Frames

Ethernet jumbo frames, which support a larger Ethernet maximum transmission unit (MTU) size, are supported on the VMstore's data interface for those who use jumbo frames. Tintri VMstore supports an MTU of 9,000 bytes. Use an MTU of 9,216 in all devices on all data paths from Tintri to hosts to accommodate form trucking and Ethernet overhead.

The critical issue with jumbo frames is that it must be enabled everywhere, including interswitch links (ISLs), to ensure consistent behavior. Because of the large number of support calls resulting from misconfiguration of jumbo frames, Tintri recommends **against** configuring jumbo frames unless you have a specific reason to do so.

For more information on Jumbo Frames, refer to the [Tintri VMstore with VMware Best Practice Guide](#).

DO: If you decide to use Jumbo Frames, ensure that MTU is configured to 9,216 throughout the **ENTIRE** path for the storage VLAN, including the interswitch links of upstream switches.

To assist validating that jumbo frames are configured end-to-end, we've added a new utility in Tintri OS version 4.2, as shown in Figure 23 . While it's useful to check, keep in mind that **only the active data path is being tested**, and other scenarios can vary test results, such as hashing algorithms, failovers due

to cable link problems, or operating on a different VMstore controller after failover or software update. Be sure to test as thorough as possible to ensure all paths allow jumbo frames end-to-end.

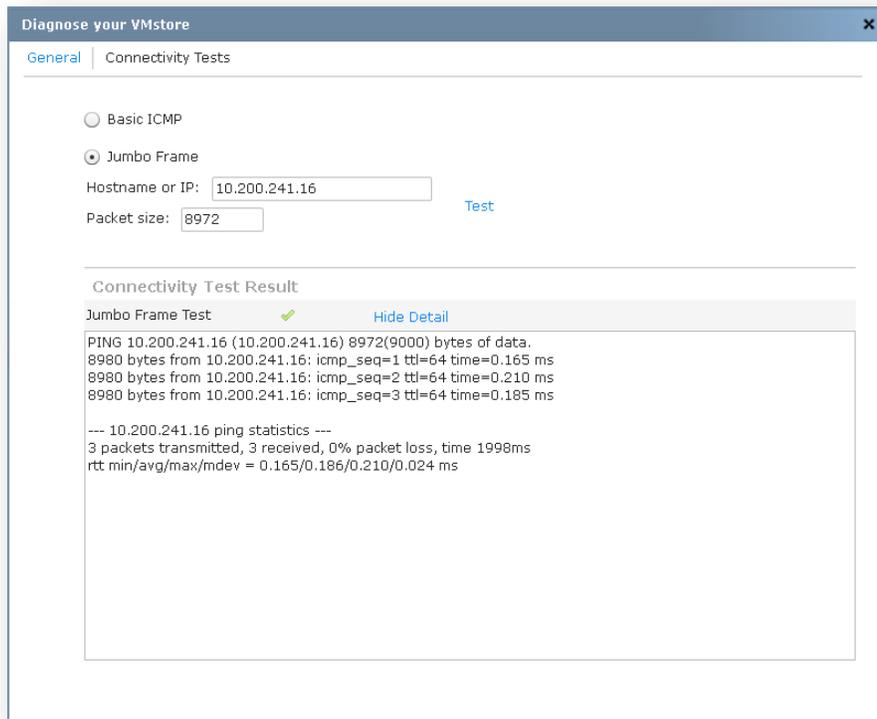


Figure 23 - Network Diagnostics tool that is available in VMstore version 4.2 and later

To test if jumbo frames are configured correctly, **make sure that the Jumbo Frames diagnostics test can pass frames with an MTU of 8972 without fragmenting**. With an end-point MTU set to 9000 on vSphere Hosts and the VMstore, a test using MTU 8973 (slightly larger) will fail the “do not fragment” ping test, which is normal and expected. This can be seen in Figure 24.

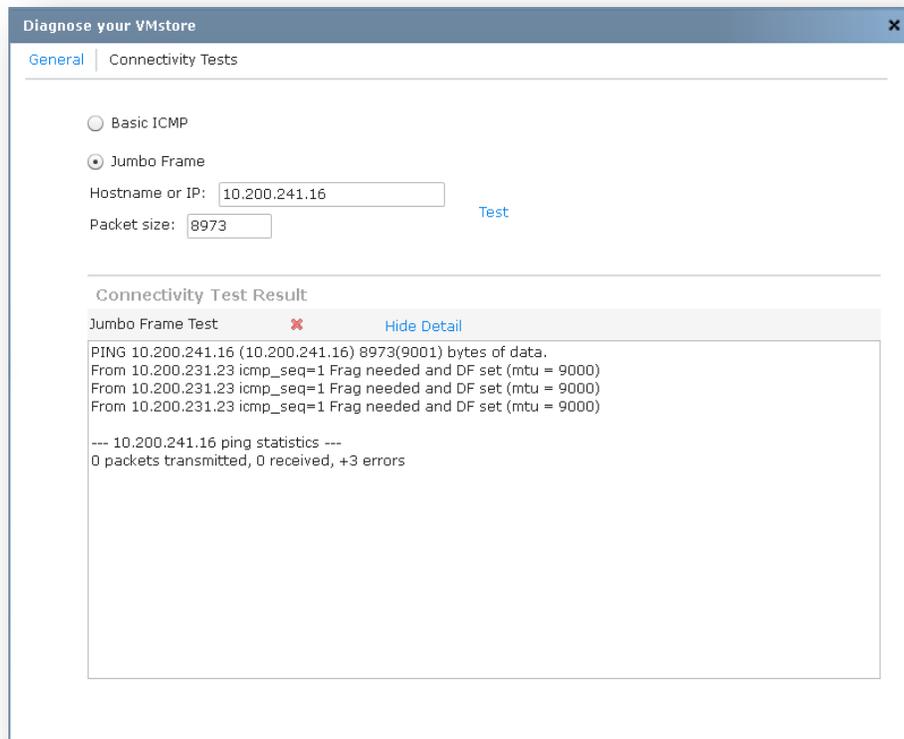


Figure 24 - Testing Jumbo Frames - "Do Not Fragment" test failed

Appendix B – vStandard Switch Configuration

Earlier in the document, we covered configuration of vSphere Distributed Switches (vds), which have many advantages over vStandard switches, including the ability to configure port groups once and have all hosts automatically assume those configuration change. However, vds aren't available in all vSphere editions lower than Enterprise Plus and you may be using vStandard switches in your environment. A major drawback of vStandard switches is that you will need to repeat the configuration for **every host**.

If you are using vStandard switches, follow the steps below. Figure 25 shows an overview of the goals for this configuration process.

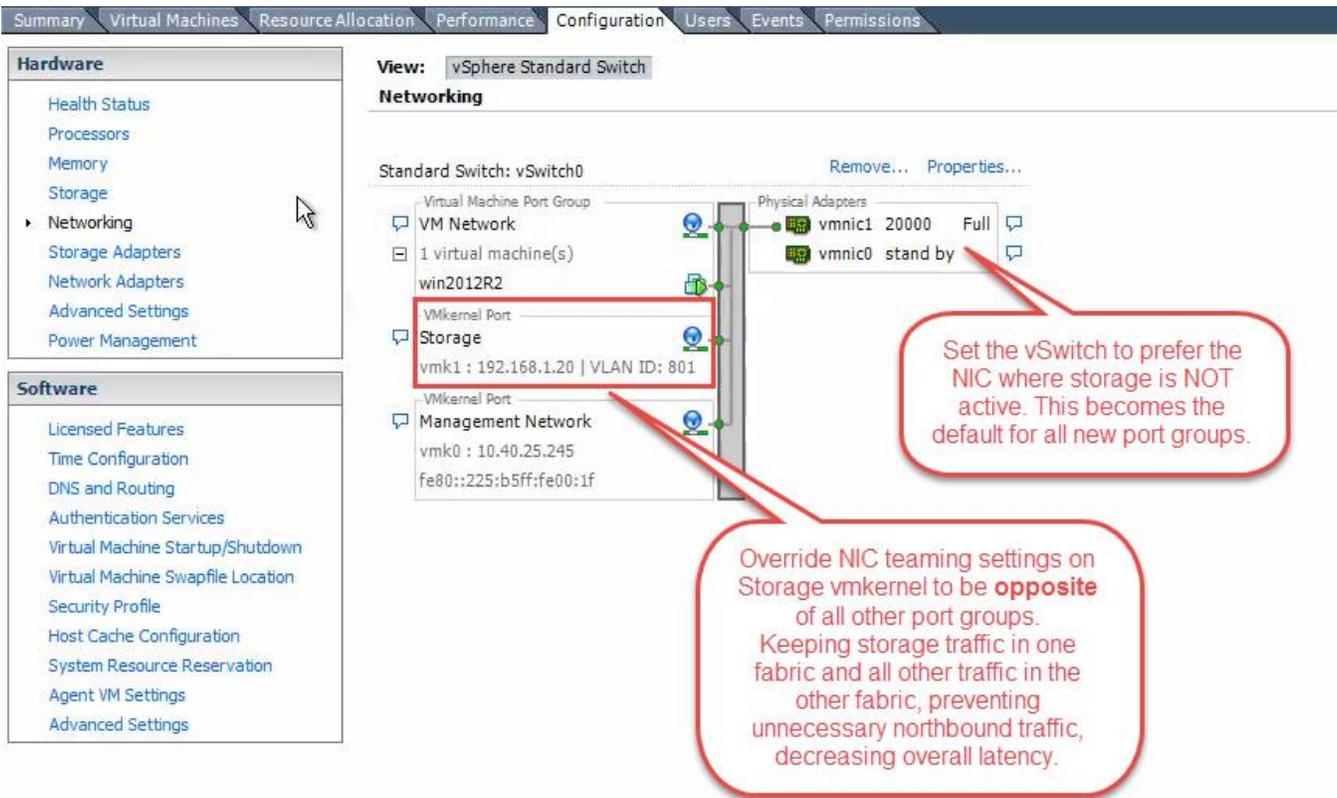


Figure 25 – Overview of vStandard switch configuration. Configure the vSwitch to prefer vmnic1 (Fabric B). Configure the Storage VMkernel Port to prefer vmnic0 (Fabric A).

1. Using the vSphere Client, select “Networking” in the left-hand pane and go to the Configuration tab. (See
2. Select “Add Networking...” from the upper right of the window.
3. On the first screen of the Add Networking Wizard select the “VMkernel” radio button and hit “Next”
4. If you have an existing switch (vSwitch0) select it and hit “Next” again.
5. On the VMkernel – Connection Settings screen, give the port group a unique name and choose a unique VLAN ID.
6. Make sure the Network Type is “IP”. Don’t select any of the checkboxes on this screen. Hit “Next”
7. On the VMkernel – IP Connection Settings screen, give the connection a unique IP address, set the subnet mask, and select “Next” and then “Finish”
8. Next, select “Properties” for vSwitch0 from the Networking Configuration pane to access vSwitch0 properties.
9. Select vSwitch from the left-hand pane and hit the “Edit” button, then select the “NIC Teaming” tab.
10. Use the “Move Up” and “Move Down” buttons to configure vmnic0 so it’s under Standby Adapters and configure vmnic1 so it’s under Active Adapters. (See Figure 26)

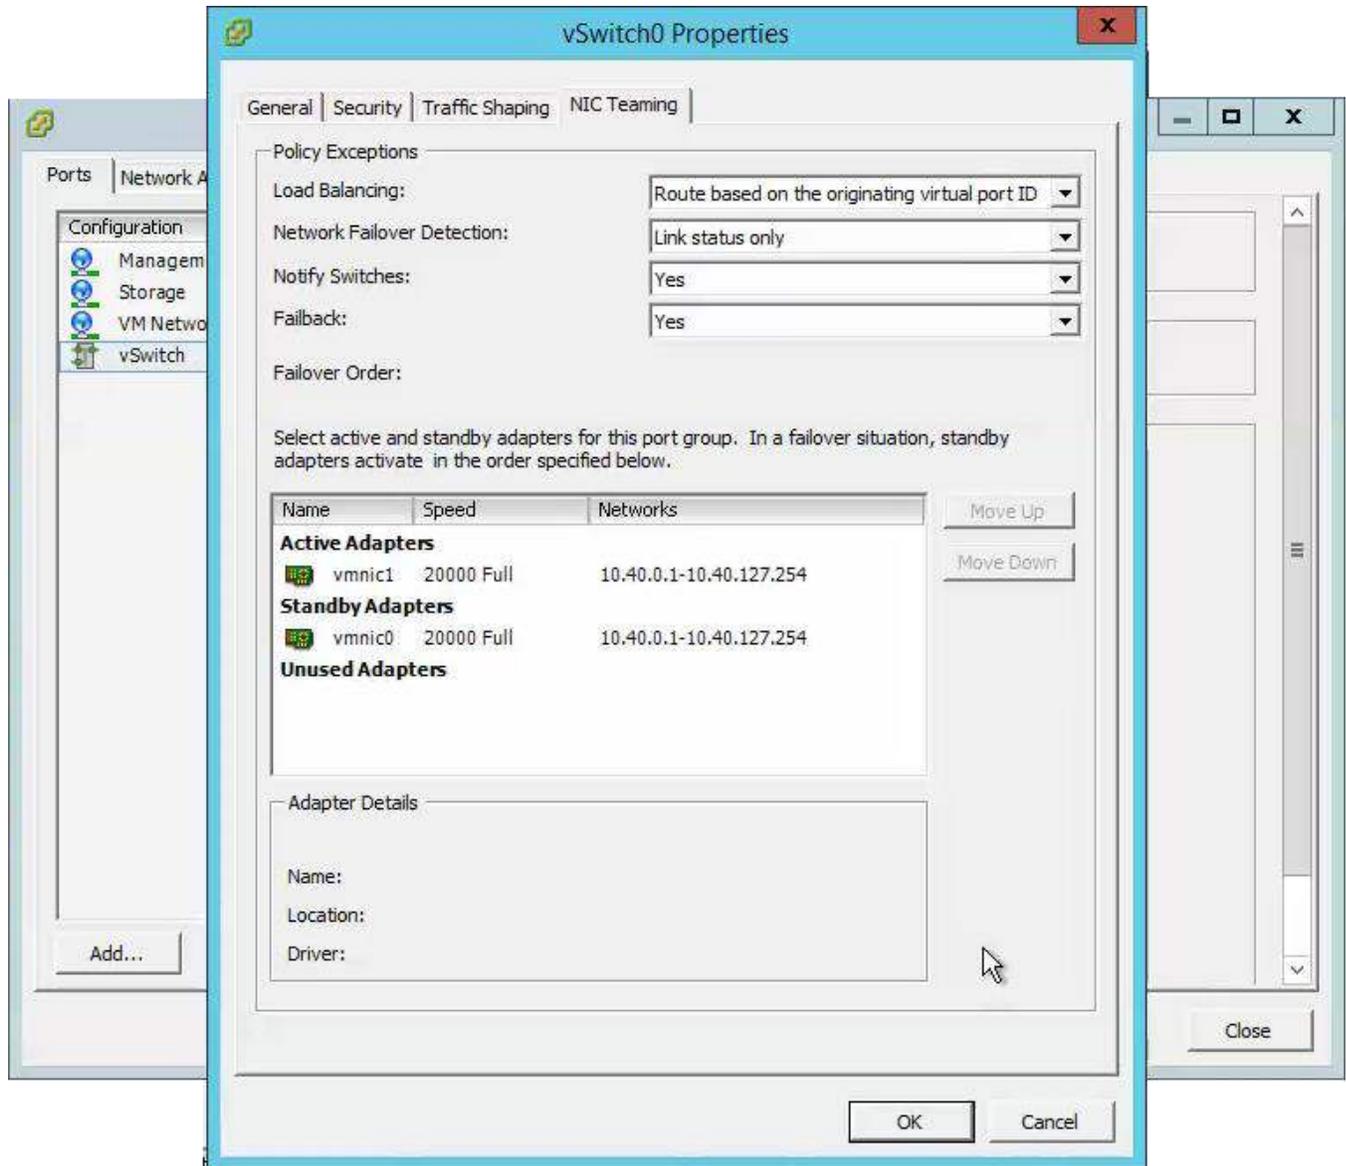


Figure 26 – Proper configuration of vmnic0 and vmnic1 for a vStandard switch.

This configuration causes the switch to send all non-storage traffic to vmnic1 on Fabric B, in effect reserving Fabric A for storage traffic. By configuring vSwitch0 this way, every new port group will inherit this configuration by default. (The configuration can be overridden as needed.)

11. Next select “Storage” from the left-hand pane and hit the “Edit” button, then select the “NIC Teaming” tab.
12. Select the checkbox for “Override switch failover order:” as shown in Figure 27.
13. Configure vmnic0 and vmnic1 as shown in Figure 27.

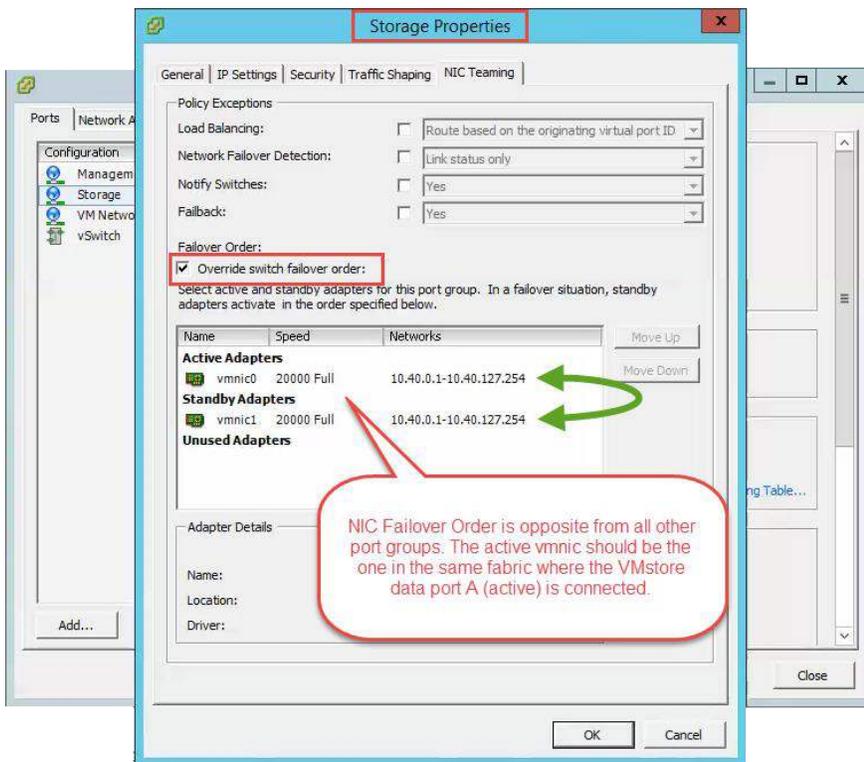


Figure 27 – Proper configuration of vmnic0 and vmnic1 for storage traffic.

14. Repeat these steps on ALL HOSTS.

At any time you can check that the configuration is correct by going to the Configuration tab and clicking the blue icon to the left of the Storage VMkernel port as shown in Figure 28. You can use the same process to verify that the VM Network and Management Network are properly configured. (Not shown.)

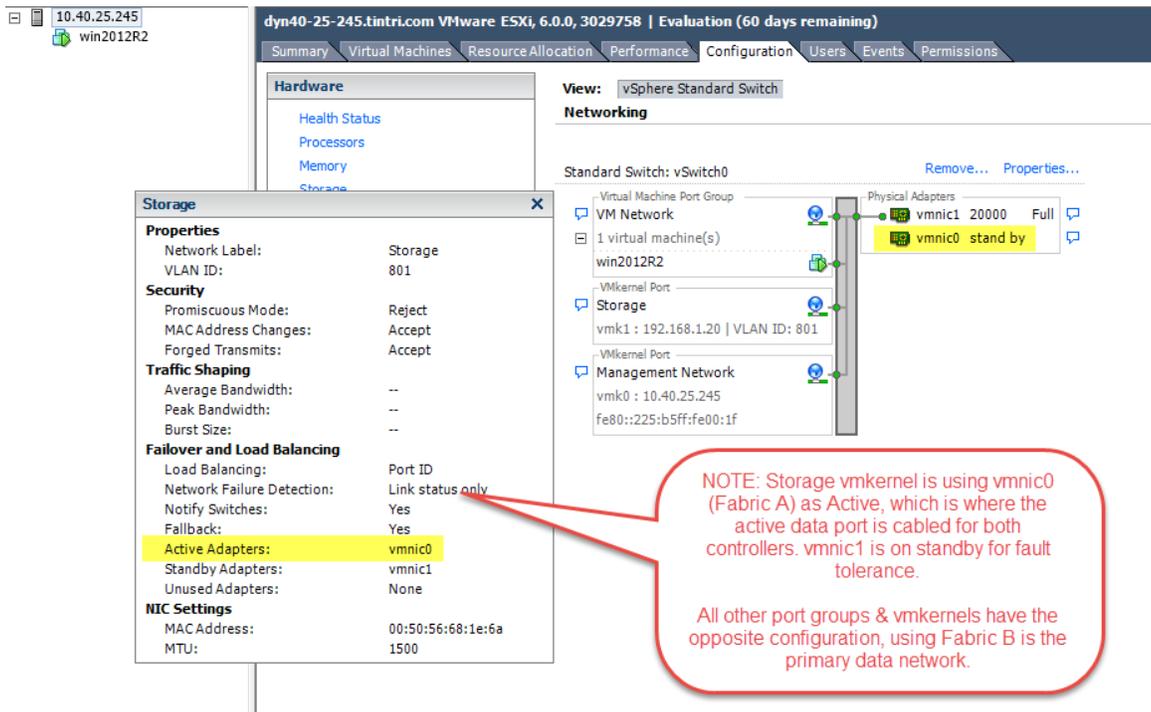


Figure 28 – Verifying correct configuration of the Storage VMkernel port

Appendix C - Nexus 5K Switch Configuration

Configure the Uplinks from the FIs to the Nexus switches as Port-Channels in both UCS Manager and the Nexus Switches.

1. To turn on the port channel feature, connect to each switch and enter the following command:

```
switch(config)#feature vpc
```

2. Next you need to create a VPC domain. Use the following command to check the configuration:

```
#Show vpc brief
```

Legend:

(*) - local vPC is down, forwarding via vPC peer-link

```
vPC domain id          : 123
Peer status            : peer adjacency formed ok
vPC keep-alive status  : peer is alive
Configuration consistency status : success
Per-vlan consistency status : success
Type-2 consistency status : success
vPC role               : primary, operational secondary
Number of vPCs configured : 2
Peer Gateway          : Disabled
Dual-active excluded VLANs : -
Graceful Consistency Check : Enabled
Auto-recovery status  : Disabled
```

vPC Peer-link status

```
-----
id  Port  Status Active vlans
--  ----  -----
1   Po1   up    1,21,31,41,51,61,71,600-649
```

vPC status

```
-----
id  Port      Status Consistency Reason          Active vlans
-----
101 Po101     up    success    success          1,21,31,41,
                    51,61,71,60
                    0-649
102 Po102     up    success    success          1,21,31,41,
                    51,61,71,60
                    0-649
```

3. To create the domain:

```
switch(config)#vpc domain 5 (any number you want, but you only do this once per switch.)
```

WARNING! – If a VPC domain already exists (as the previous step would show), do NOT create a new one! Doing so may loop your network and lead to a major outage.

4. Keep-alives are used out of band to ensure the vPC is working properly. This is usually the management IP address of the switch, but SVI (switch virtual interfaces) can be used as well. The management interface has a VRF (virtual routing and forwarding) instance inherent to it.

```
switch(config-vpc-domain)# peer-keepalive destination 192.168.254.2 (peer switch destination management port)
```

Using an SVI for peer keepalive link:

```
switch(config-vpc-domain)#peer-keepalive destination 172.30.1.2 source 172.30.1.3 vrf default
```

5. Create the peer-link. This is the port bundle between the switches (ie. Inter-switch link). This should be at least two 10Gb ports. This link synchronizes the state of the switches.

```
switch(config)# interface port-channel 20
switch(config-if)# vpc peer-link
```

6. Create the port group to connect the Nexus switches to the Fabric Interconnects:

```
switch(config)#interface e1/1
switch(config-if)channel-group 20
switch(config-if)# interface port-channel 20
switch(config-if)# vpc 100 (must match on both switches)
```

Below is a simple sample as seen from the running config:

Switch 1:

```
feature vpc
vpc domain 123
  peer-keepalive destination 192.168.255.254
port-profile default max-ports 512

interface port-channel1
  switchport mode trunk
  switchport trunk allowed vlan 1,21,31,41,51,61,71,600-649
  spanning-tree port type network
  speed 10000
  vpc peer-link

interface port-channel101
  description Port Channel for UCS
  switchport mode trunk
  spanning-tree port type edge trunk

  speed 10000
  vpc 101

interface port-channel102
  description Port Channel for UCS
  switchport mode trunk
  spanning-tree port type edge trunk
  speed 10000
  vpc 102
```

Switch 2:

```
feature vpc
vpc domain 123
  peer-keepalive destination 192.168.255.253
port-profile default max-ports 512
```

```

interface port-channel1
  switchport mode trunk

  switchport trunk allowed vlan 1,21,31,41,51,61,71,600-649
  spanning-tree port type network
  speed 10000
  vpc peer-link

interface port-channel101
  description Port Channel for UCS
  switchport mode trunk
  spanning-tree port type edge trunk
  speed 10000
  vpc 101

interface port-channel102
  description Port Channel for UCS
  switchport mode trunk
  spanning-tree port type edge trunk
  speed 10000
  vpc 102

```

Appendix D – Failure Scenarios & Testing Redundancy

If you have followed the advice in this guide, your configured UCS Infrastructure with one or more Tintri VMstores should be fully redundant and capable of withstanding interruption or failure of any one component. In other words, there are no single points of failure in the design. Not only is this desirable for withstanding unlucky equipment failure, but a resilient design is also critical to being able to upgrade & perform routine maintenance on all components with the system (UCS & host software updates, firmware updates to hosts, networking switches and FIs, VMstore upgrades, etc.).

Below are some areas to consider and thoroughly test to ensure a non-disruptive experience for your end-users and customers. Ideally, all of these tests should be performed before production workloads occupy the infrastructure. Even once you've fully tested all aspects of your infrastructure and you could take down any component anytime with minimal or no disruption, we strongly recommend that maintenance tasks and other operations that carry a degree of risk should be performed in pre-planned maintenance windows to minimize impact in the event something doesn't go smoothly as planned.

Network Path Fault Tolerance

If you follow this guide, storage traffic will remain in fabric and not flow through the uplinks (ie. Northbound traffic), EXCEPT in one of the following failure scenarios:

1. Cable issue between data port A on the active controller of the VMstore and the FI it is connected. This includes unplugging the cable, cutting it, having an SFP failure, or any other scenario that would render the link state as inoperable. There is no way to inform the hypervisors to use their standby NIC connected to the other fabric, so they will continue to use the primary NIC, which will send traffic into Fabric A, up to the uplinked switch(es), back down to Fabric B, and then into the VMstore. The return traffic follows the same path.
2. FEX A is pulled from the back of the chassis. In this case, the Hypervisor will see vNIC0 as down, and user vNIC1, which is connected to Fabric B. The VMstore will still operate on Fabric A, because data port A is unaffected (link state is good), and communicate on Fabric A, which has to go through the upstream switch to get to Fabric B and the host.

All other failure scenarios are covered without requiring upstream traffic. If all of Fabric A goes down (as is the case when you upgrade firmware in UCS), both the blades AND VMstore will sense the failure, and both will use Fabric B, which won't require traffic to traverse the upstream switch.

If a VMstore controller fails (or you upgrade Tintri software, which results in a controller failover), traffic will prefer Fabric A on the other controller. The hosts will be unaware of the controller failover, and will continue to favor Fabric A for storage. This is why it's important that port A on each controller be connected to Fabric A.

To fully test your network redundancy, you'll want to test all major components:

- **Cables** - Disconnect one at a time and ensure you give ample time for connections to come back up before moving onto the next cable. Administratively shutting down ports can have the effect as physically disconnecting cables and has the advantage of being performed remotely. If you are remote and this is the first time you've performed these tests, ensure that you have an alternate option to re-connect in the event you disconnect yourself and the expected resiliency does not perform as expected. In other words, it's best to have someone available onsite, just in case!
- **Fabric Interconnects** – Ensure you can lose an FI (graceful shutdown or via hard power disconnects). Then bring the FI back up, wait for it to enter stable state and then repeat on the other FI to ensure no single points of failure were overlooked on it.
- **Upstream Switches** – Similar to testing the FIs, ensure your design can tolerate failures of your upstream switches
- **Management Access** – Throughout your tests, ensure you still have management access to all critical components
- **Power** – Redundant power sources are strongly recommended. Pull power cables to ensure multiple power supplies provide the required power to keep components online. Also keep in mind the scenario of losing an electrical circuit by cabling your redundant power supplies to independent power sources wherever possible.

Host Failures

VMware provides us with the ability to create clusters with High Availability and Fault Tolerance. Your design should include enough spare compute (ie. N+1 or greater redundancy) to build your clusters in a way that can tolerate an unexpected host failure as well as planned host outages by leveraging Maintenance Mode. For planned host outages, VMs non-disruptively vMotion to other hosts in the cluster, and then rebalance themselves (using DRS) once the host comes back online. For unexpected outages, VMware HA brings VMs up on surviving hosts quickly and automatically. For critical VMs and VMs that other VMs have dependencies on, we recommend you place a higher restart priority on these to minimize the impact and shorten the time it takes to return to fully operational.

Human Error

More often than not, human error is a greater threat than unlucky component failures. While it's impossible to fully protect yourself from the risks of human error, ensure you have frequent backups your data and configurations. Well-documented Data Protection and Disaster Recovery planning is your last line of defense.... Test recoveries and failovers periodically.

© 2016 Tintri, Inc. All rights reserved. Tintri, Tintri VMstore, Tintri Global Center, ReplicateVM, SecureVM, and SyncVM are trademarks of Tintri, Inc., and may be registered in the U.S. Patent and Trademark Office and in other jurisdictions. All other marks appearing in this publication are the property of their respective owners.

Tintri believes the information in this document is accurate as of its publication date. The information in this publication is provided as is and is subject to change without notice. Tintri makes no representations or warranties of any kind with respect to the information in this publication, and specifically disclaims any implied warranties of merchantability or fitness for a particular purpose.



303 Ravendale Drive
Mountain View CA 94043
+1 650.810.8200
info@tintri.com