Tintri BY DDN

TECHNICAL WHITE PAPER

Tintri VMstore with VMware Best Practices Guide

Best Practices for Deploying the Tintri VMstore™ in VMware vSphere™ Environments

Revision History

Version	Date	Description	Author
3.1	4/2/2019	Amended sections: Jumbo Frames, vSphere Advanced Settings	Tomer Hagay
3.0	10/5/2017	Updated	Tintri Technical Marketing
2.1	10/25/2016	Updated	Rob Girard
1.9	02/03/2015	Document previously titled: "**Tintri NFS and vSphere Best Practices**"	Tintri Technical Marketing

Table 1 - Revision history

Contents

Introduction

The Tintri VMstore is storage specifically built for virtualized infrastructures and offers the ability to directly integrate with VMware vCenter, as well as multiple other hypervisor management platforms. This guide highlights the key considerations and configuration settings that promote a high-performance and reliable virtualized environment for connecting vSphere assets and clusters with the Tintri VMstore platform.

The primary focus of this paper will revolve around the integration and best practices for VMware vSphere virtualized enterprise infrastructure when used with the Tintri VMstore platform. There will however, be a need throughout the document to reference a number of the complimentary technology offerings which comprise the greater Tintri product suite. These include (but are not limited to) the Tintri Operating System (TxOS) which manages the VMstore appliance and integration across connected hypervisors, Tintri Global Center (TGC) which provides centralized management and control for Tintri Scale-out deployments, and Tintri Analytics which provides historical and predictive data for monitoring and forecasting of storage resources.

Intended Audience

This document is intended to assist virtualization administrators, storage engineers and data center architects with responsibilities for the design, deployment and management of VMware-based virtual infrastructures that are hosted on the Tintri VMstore.

Consolidated List of Practices

This section provides a consolidated list of best practices to assist in optimal management of the Tintri VMstore in your VMware virtualized environment. Click the text on any of the recommendations to jump to the section that corresponds to each recommendation for additional information.

✧ Full details concerning compatibility and support across the various Tintri operating system (OS) and VMware vSphere product suite options can be found in the Tintri OS and External Compatibility guide available on the Tintri support portal, or by contacting your local Tintri representative.

✧ Isolate the storage traffic across your infrastructure by assigning a dedicated VMkernel interface on each vSphere host that will access the Tintri VMstore

✧ Ideally, each VMkernel interface should have their own dedicated NIC pair for redundancy and isolation, but where physical resources preclude this, it is advised to assign adapters based upon best practice for security and availability

✧ Each Tintri VMstore participating in a replication schema should run the same version of Tintri OS

✧ A complete list of the required TCP/UDP ports for Tintri network services is available in the Appendix of the Tintri VMstore Admin Guide. Enterprise architects and network administrators should review the list of services and required ports before deploying Tintri Vmstore platform solutions

✧ Ensure that when using Ethernet Jumbo Frames that the vSwitch settings, the physical network switch, and Tintri VMstore are all configured to use jumbo frames. Explicitly configure your physical Ethernet switches and vSwitches with an MTU size of at least 9,000 bytes (typically 9,216) when using Jumbo Frames.

✧ In lieu of dedicated physical networks, VLANs can offer a flexible means to implement logically isolated networks that share common physical resources

✧ Adjust the ESXi host Advanced Settings for Nfs.MaxVolumes, Net.TcpipHeapMax and Net.TcpipHeapSize to their maximum values on EVERY host

✧ Confirm that the ESXi host Advanced Settings for NFS.Heartbeat conforms to recommended values on EVERY host

✧ Tintri recommends using Thin virtual disks as the provisioning format for virtual disks in all scenarios unless specific clustering functionlaity is required

✧ The Tintri VAAI plug-in must be installed on each vSphere host that is intended to leverage the supported functionality. If not installed, it is recommended that the Tintri VMstore UI or Tintri vSphere Plugin be used to create space-efficient clones instead of the vSphere Client.

✧ Installing the Tintri vCenter Plug-In will allow vCenter administrators and users monitor the storage-specific metrics of their Tintri VMs as well as enact per-VM storage related actions from within the vCenter web client interface

✧ Thin provisioning is required to fully realize the capacity efficiencies associated with modern storage platforms that support data reduction technologies

✧ Thick provisioned vDisks in their default state will not benefit from any space savings on the Tintri VMstore platform

✧ Utilizing the ‘Maximize Space Savings’ option on your VMstore in conjunction with NFS VAAI support at the hypervisor layer will allow Thick provisioned virtual disks to enjoy all the data reduction benefits of Tintri’s data-aware deduplication and compression abilities.

Overview

The Tintri VMstore platform is a fully scalable enterprise-storage building block which offers the ability to integrate directly with industry leading hypervisor-management platforms. The Tintri VMstore provides the benefits of flash performance with an underlying architecture that allows each and every storage action to be performed and monitored at the granularity of the Virtual Machine (VM).

The Tintri VMstore platform extends and simplifies the management of VMs through intrinsic VM-object awareness that reaches from the top of the virtualization and computing stack, all the way down to the storage media on the Tintri VMstore. The Tintri VMstore utilizes a custom NFS (Network File Systems) access protocol, specifically optimized for hypervisor-based integration.

Tintri's VM-level granularity of monitoring, management, and orchestration is available to storage and virtualization administrators when using either the VMstore user interface (UI), vSphere web client, or through a rich and complete set of web services for automation and integration. This list of available actions is comprehensive and includes; I/O scheduling, cloning, snapshots, replication, copy data management, disaster recovery, analytics, and quality of service (QoS) policies, amongst others. The key is that each of these actions (and more) can be defined and enacted per VM, as opposed to legacy storage methods which are constrained to LUN-level actions.

Figure 1: Front profile of a Tintri VMstore storage appliance

Simplicity is a key design mantra across the available Tintri suite of products. This extends from the initial installation, configuration, and management, all the way through to the ability to grow your deployment and scale-out as business needs dictate. The Tintri VMstore is designed so that technology professionals and administrators with a working knowledge of VMware should be able to successfully deploy Tintri's purpose-built virtualization-aware storage as easily as they would a vSphere Server.

✧ *Full details concerning compatibility and support across the various Tintri operating system (OS) and VMware vSphere product suite options can be found in the Tintri OS and External Compatibility guide available on the Tintri support portal, or by contacting your local Tintri representative.*

VMstore Networking

As the Tintri VMstore platform utilizes IP based networking for management, storage, and replication traffic – a key recommendation is to ensure proper configuration of both the physical and logical networks. Each VMstore requires a minimum of at least two active physical network connections and another two standby connections for supported operation. These connections are to be used for the VMstore appliance's Management and Storage I/O (Data) interfaces. The Tintri VMstore also offers additional dedicated network interfaces which can be assigned to carry Replication traffic. The configuration and function of each will be discussed throughout the following section.

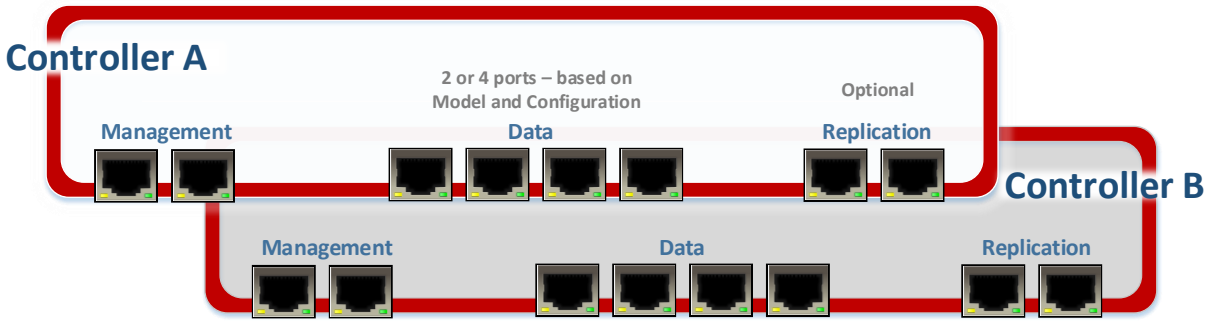

Figure 2: An abstracted representation of the available ports and their functions on the dual-controller Tintri VMstore storage appliance.

The Tintri VMstore Management network is used by administrators for system configuration, management, and monitoring purposes. This connection is required to be in place to allow initial setup of a VMstore appliance. During the initial configuration sequence of a VMstore appliance, a pre-defined IP address to be used for the Management interface will be requested – associating this IP address with a fully qualified domain hostname (FQDN) that is resolvable with the appropriate DNS (Domain Name Services) servers is optional, but advised. The management interface for the Tintri VMstore does not have to be in the same management network being used for the ESX/ESXi servers which will comprise the virtualized environment, but it does have to be routable to the vCenter servers with which it will integrate.

The Storage network refers to the NFS (Network File System) traffic which will handle all requests to storage or I/O (input/output) between the hosted virtual machines (via the hypervisor layer) and the Tintri VMstore storage layer. Please note, this network can also be referred to as the Data network throughout this and other Tintri documents. In addition, the terms Network File Systems (NFS) and NAS (Network Attached Storage) may be used interchangeably in certain dialogs.

✧ *Isolate the storage traffic across your infrastructure by assigning a dedicated VMkernel interface on each vSphere host that will access the Tintri VMstore*

As per any IP (Internet Protocol) based storage configuration, it is recommended to isolate the Data traffic and assign a dedicated VMkernel interface on each vSphere host that will access the Tintri VMstore. The VMkernel port can be thought of as a logical interface for a vSphere host that is then assigned to a vSphere virtual switch. The virtual switch is then assigned one or more physical network interfaces which provide the uplink to the physical network(s).

This storage-specific port group can be placed on a new or existing virtual switch which can use either vSphere virtual switching format; vSphere Standard Switch (vSS, configured independently on each host) or vSphere Distributed Switch (vDS, configured centrally in vCenter across all hosts). A maximum of 64 data network IP addresses can be assigned per VMstore, but for simplicity of management, Tintri recommends that just one is used.

Tintri best practices also recommend that some form of segregation should be applied to the storage-specific host-side VMkernel port – this can be taken care of using either physical or logical means. Mixing other traffic such as user VM access, fault tolerance, vMotion, and additional storage protocols such as iSCSI on the same VMkernel port is possible, but not advised.

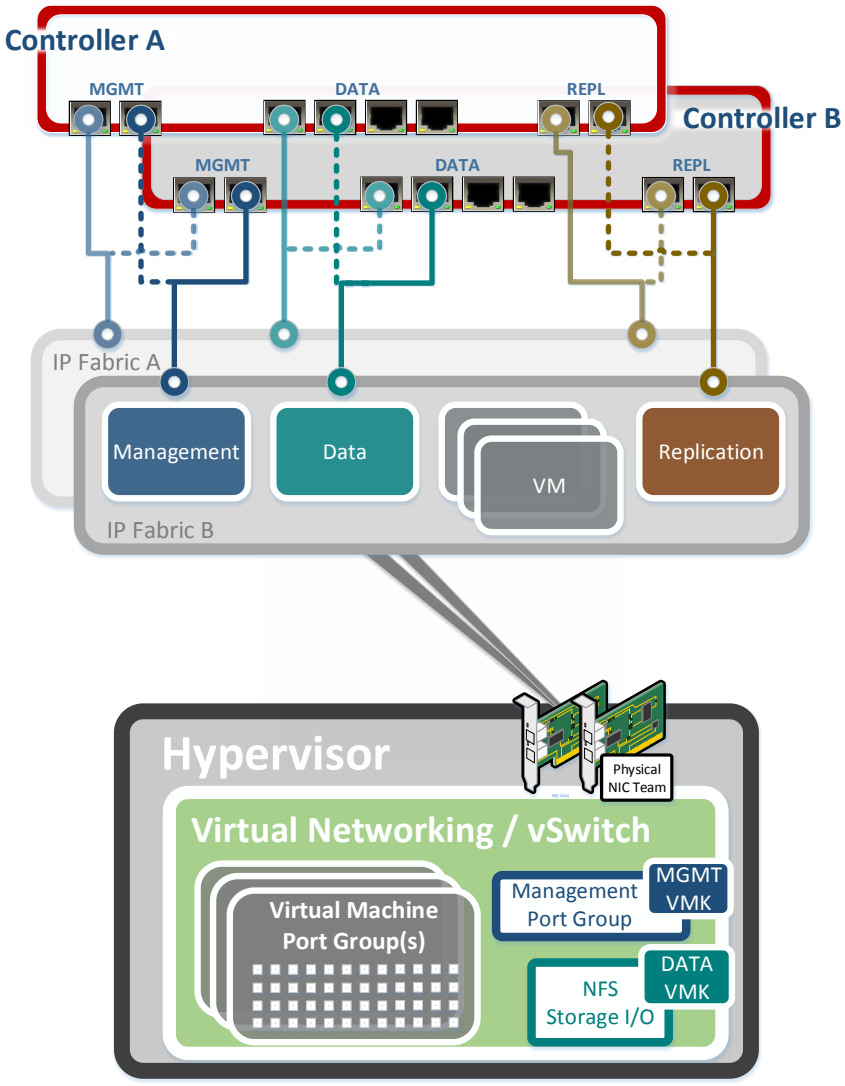

Figure 3: A logical representation of a basic networking topology for a typical data center stack comprising top to bottom; Tintri VMstore platform and associated network functions; IP switching layer with example VLAN definitions; host layer comprising hypervisor managed compute, memory & virtual networking.

A Replication network which can be used for data-aware optimized per-VM replication traffic between Tintri VMstore appliances is an optional configuration. On VMstore systems that are equipped with dedicated network interfaces for replication, this traffic can be segregated from the data and management network paths.

Dedicating interfaces for the purposes of VM replication ensures that the operational VM I/O and (VMstore to VMstore) replication traffic have dedicated active/standby network pathways. Without dedicated network interfaces, replication is still available, however VM replication traffic will be required to travel over the VMstore's Data or Management interfaces.

The VMstore-specific networks will not be the only traffic that needs to be considered when taking a holistic view of a VMware-based virtualized datacenter. Of primary significance will be the VM (Virtual Machine) traffic, those communication packets that are sourced from, or delivered to the hosted virtual machines and hosted business critical applications. VM networks should be isolated at the hypervisor and switch layers based on desired service levels. The VM networks do not need a direct connection to the Tintri VMstore, as storage access is performed via the Data Network which should already have a dedicated VMkernel and port group at the hypervisor level.

At its most basic, each vSphere host should define at least three VMkernel network interfaces. These VMkernel ports should be dedicated to management, storage, and vMotion traffic respectively. Ideally, each of these three interfaces should have their own dedicated network adapter, but where physical resources preclude this, it is advised to assign adapters based upon best practice for security and availability. At a minimum, one physical network interface card (NIC) should be reserved for failover purposes.

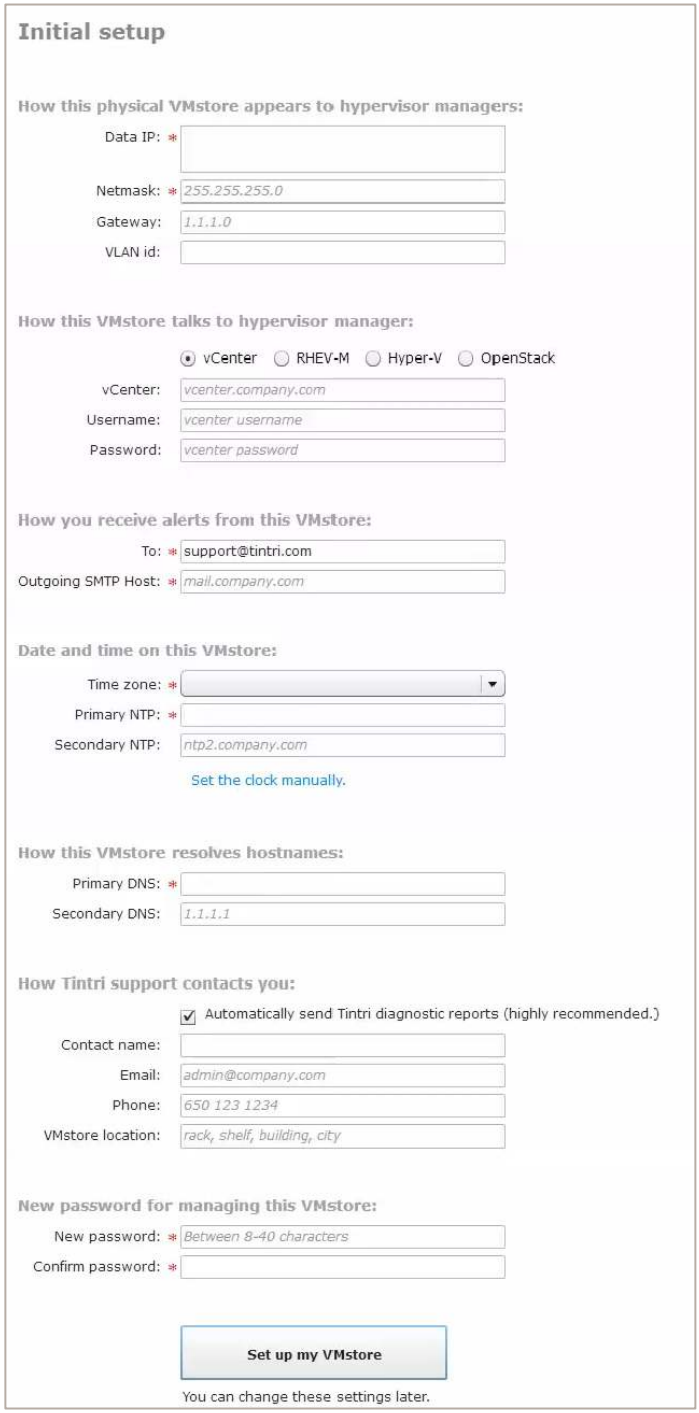

Figure 4: Tintri VMstore initial setup interface requesting data IP and additional initialization details (following the assignment of Management IP)

✧ *Ideally, each VMkernel interface should have their own dedicated NIC pair for redundancy and isolation, but where physical resources preclude this, it is advised to assign adapters based upon best practice for security and availability*

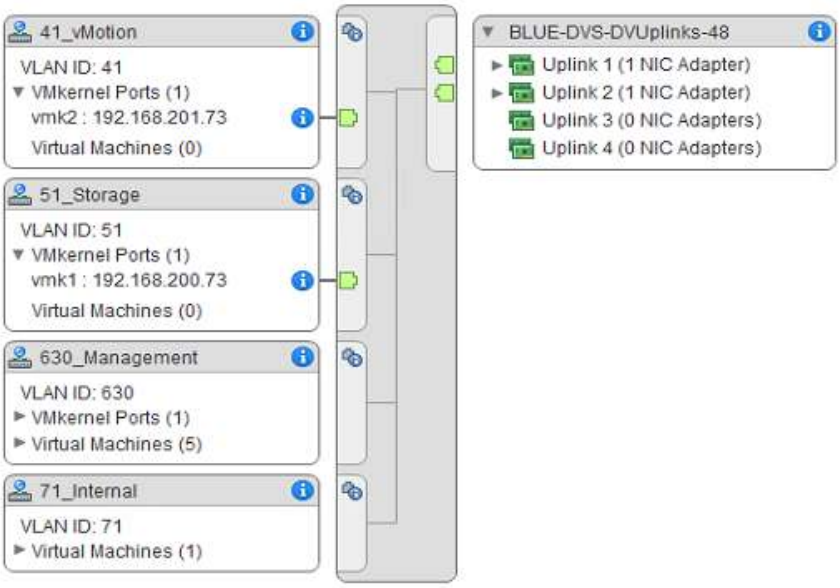

Figure 5: VMware vSphere virtual networking configuration displaying a single Distributed Switch with redundant physical NICs hosting logically segmented logical networks via the use of VLANs.

Redundancy

Connecting the Tintri VMstore to redundant networks is not required for basic operation, but it is a common best practice and highly recommended. Highly-available deployments help protect against storage controller events, network path outages, and complete network switch failures which can disrupt an entire enterprise virtualized infrastructure. To achieve high availability, the LAN on which the NFS traffic will run needs to be designed with availability, downtime-avoidance, isolation, and no single-point of failure (SPOF) in mind.

For simplicity and high-availability, the VMstore's Management and Data interface ports are internally linked as active/standby failover pairs out-of-the-box. Port pairs cannot be split into separately usable interfaces. Either a single port of a given pair of ports is connected leaving the other port unused, or both ports are connected and automatically linked as an internal failover pair by the VMstore.

A controller failover may still occur without fully redundant networking as long as each controller in the HA (High Availability) pair has one management and one data interface port connected. Redundant networking ensures maximum availability in the case of a network disruption. In the case of a network path failure, the Tintri VMstore will automatically activate standby network ports when the link status for an active port transitions to a down state.

The networks to which the VMstore and vSphere assets are connected can consist of multiple Ethernet switches and switch ports to maximize network path availability. Where possible, enterprise architects should look to implement a redundant switching network which can be employed to provide full high availability for each of the core networking functions.

Tintri recommends that a dedicated physical network or VLAN (Virtual Local Area Network) should be employed to carry the traffic between ESX/ESXi Server and Tintri VMstore. In lieu of a dedicated physical network, VLANs may also offer a flexible means to implement a logically isolated storage network. This can be achieved providing that the switch or switches that carry the VLAN traffic have sufficient bandwidth to provide an unrestricted flow of data between vSphere hosts and the Tintri VMstore – refer to the subsequent section: 'VLAN configuration' for more information. An illustrative

example of a minimal-scale, but highly-available networking configuration which combines both redundant switching and VLAN-based logical segmentation is shown in Figure 6.

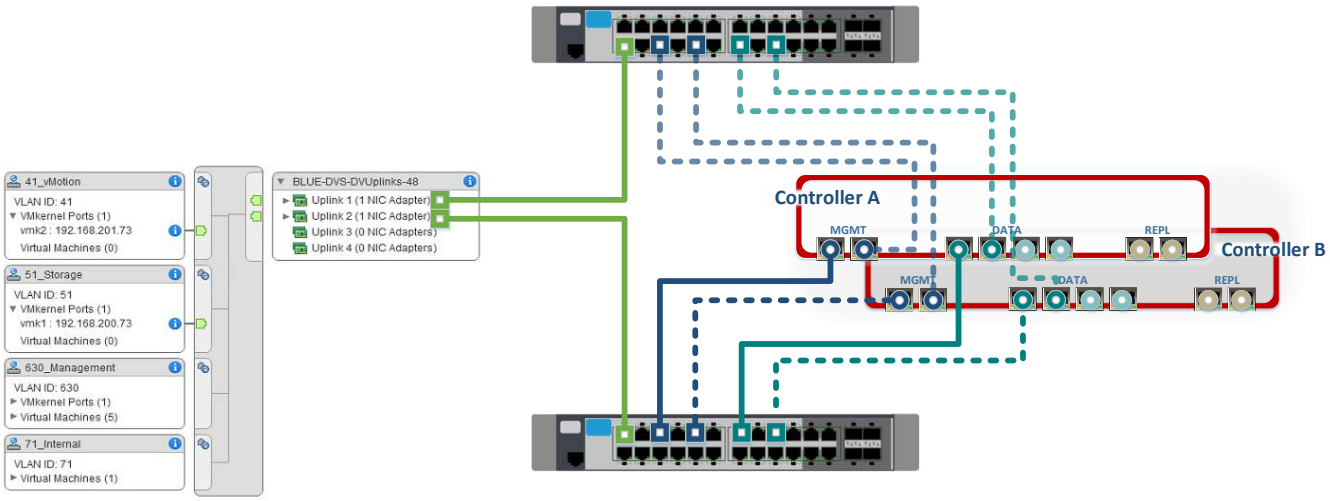

Figure 6: An example of a minimal-scale highly available networking configuration using redundant switching and VLAN trunking via a single NIC pair at the hypervisor layer.

The following example depicts a generic example of a Tintri VMstore appliance connected within a fully redundant network infrastructures, with multiple switches for the management and data networks. The use of redundant switches shown in these examples is optional, but preferred for the Data (Storage I/O) network.

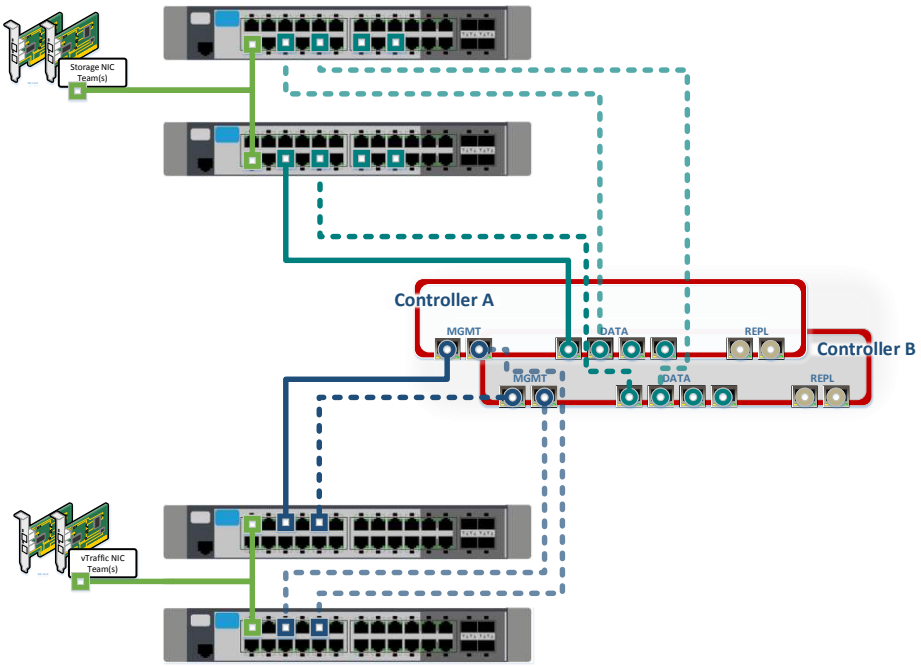

Figure 7: An example of a minimal-scale highly available networking configuration using dedicated physical switching for isolation of Data network traffic.

In the example shown in Figure 7, only two of the optional four Data ports available for Storage I/O traffic are shown as connected. This is for clarity of illustration only, as end-users who will look to implement a physically isolated network for Data traffic may also wish to increase available bandwidth or implement additional contingencies for redundancy. This can be done by increasing the Data port count to four and also through aggregating these ports into a single port-channel group. These constituent paths can be distributed amongst multiple switches using the same logic displayed in both Figure 6 and Figure 7. More information can be found in the subsequent 'LACP section' in this paper.

Replication

Each of the available Tintri VMstore platform versions provide the capability to deploy dedicated ports for use with VMstore to VMstore replication of VMs and their constituent virtual disks. This allows both inter and intra-site replication traffic to be segregated from the data and management network paths and offers data center architects and network administrators maximum flexibility in how to configure and best utilize available networking paths. Without dedicated network interfaces, both synchronous and asynchronous replication options are still available and fully-functional, however VM replication traffic will be required to travel over either the VMstore's Data or Management interfaces.

✧ *Each Tintri VMstore participating in a replication schema should run the same version of Tintri OS*

Replication between unrelated VMstore models (including between both hybrid and all-flash models) is explicitly supported, although Tintri does recommend that each site participating in a replication schema should run the same version of Tintri OS. The points discussed and illustrated in the following section will focus on one-to-one replication, but these same concepts can be applied when replicating across multiple sites to a remote DR (Disaster Recovery) location (many-to-one) or vice versa, from one source to up to four locations (one-to-many).

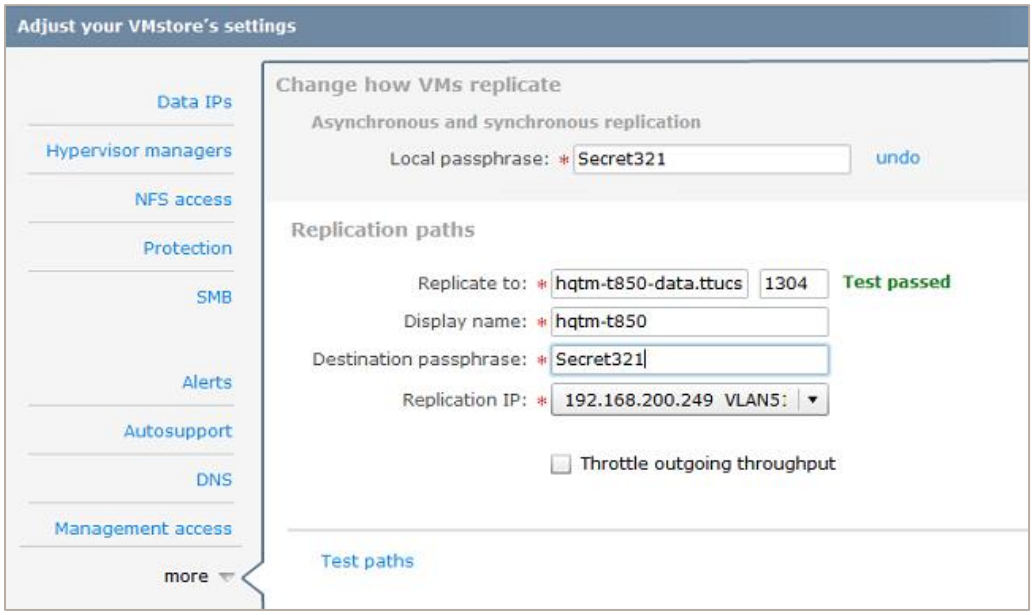

Figure 8: Configuring, amending, or testing VMstore Replication paths can be done by choosing 'Settings – Replication' in the VMstore user interface

✧ *A complete list of the required TCP/UDP ports for Tintri network services is available in the Appendix of the Tintri VMstore Admin Guide. Enterprise architects and network administrators should review the list of services and required ports before deploying Tintri Vmstore platform solutions*

All replication traffic between VMstores is done using TCP/IP. By default destination port 1338 is used for synchronous replication, but this can be changed as required. If changing the default port number used, this must be made to be different from the default TCP port used by Tintri for data-aware snapshot-based (asynchronous) replication – port 1304. Data center firewalls should be configured appropriately to allow replication traffic on the specified TCP ports. These TCP ports can be configured directly from the Tintri VMstore user interface by choosing 'Settings – Replication'. This interface also offers the ability to test replication path connectivity – refer to Figure 8.

It should also be noted that Tintri offers array-based replication on a per-VM basis. Per-VM replication is a simple, yet powerful concept. Instead of replicating the entire LUN (Logical Unit Number, a legacy Storage container) which would normally contain tens to hundreds of whole and partial VMs, Tintri users need only replicate the particular VM(s) and component vDisks of interest. Additional layers of orchestration and control are available via VM (or Service Group) snapshot replication scheduling and Tintri's data-aware mode of operation for replication.

Administrators can define a VM's (or group of) snapshot and replication schedule with an RPO (Recovery Point Objective) which can be as low as one minute. These actions can be performed via Tintri Global Center Service Groups, through VM-level actions in the VMstore user interface, or alternatively from within the VMware vCenter Server web client through integration of the Tintri vCenter Plug-In – refer to Figure 9.

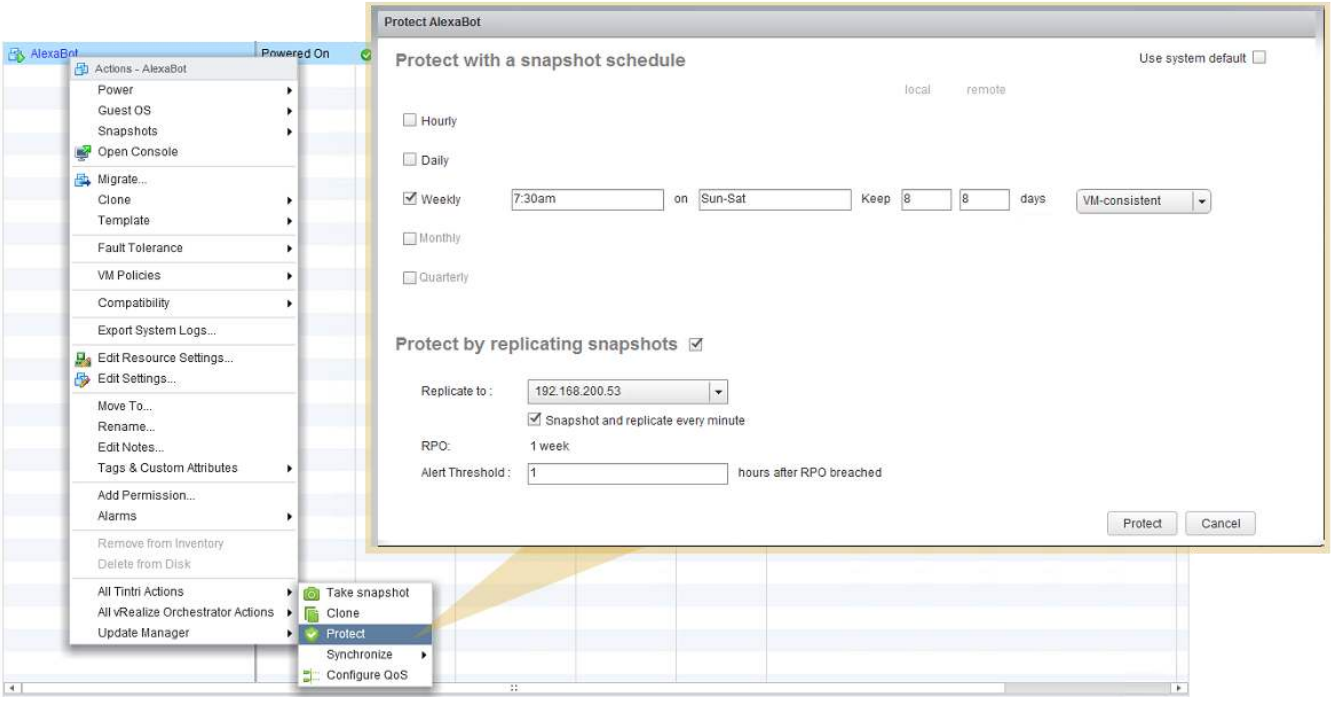

Figure 9: Controlling a VM's protection and replication scheduling using Tintri's vCenter Plug-In functionality.

In the case of setting a replication RPO of one minute, network bandwidth usage for replication would be at its highest due to the maximum frequency of replication being implemented. Disregarding the fact that Tintri VMstore snapshot replication minimizes replication resource usage by only sending unique (post-deduplication) and compressed data over a replication link, some users may want to add an extra layer of protection in relation to network bandwidth over-consumption. To do so, Tintri offers the ability to limit the throughput used by replication during user-defined time windows so that it does not interfere with periods of peak production traffic– refer to Figure 10. This affords Tintri users absolute control over the network resource usage dedicated to protecting the business's virtual machines, data and applications.

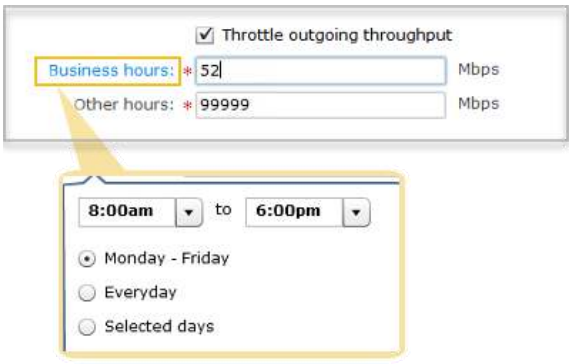

Figure 10: Tintri offers the ability to set upper limits on replication bandwidth usage during specific time periods via the 'Settings - Replication' tab in the VMstore UI

As mentioned already, the optional use of dedicated network interfaces for replication provides a network path between Tintri VMstore devices for replication. Dedicating interfaces for the purposes of VM replication ensures that the operational VM I/O and (VMstore to VMstore) replication traffic have dedicated active/standby network pathways. The VM replication data will flow from VMstore to VMstore through the enterprise network infrastructure, without involving any resource usage at the vSphere hypervisor or management layer. Modifications to vSphere virtual networking configurations to support Tintri VMstore per-VM replication are unnecessary – VM replication occurs without involving and burdening your vSphere servers. An example configuration is shown in Figure 11.

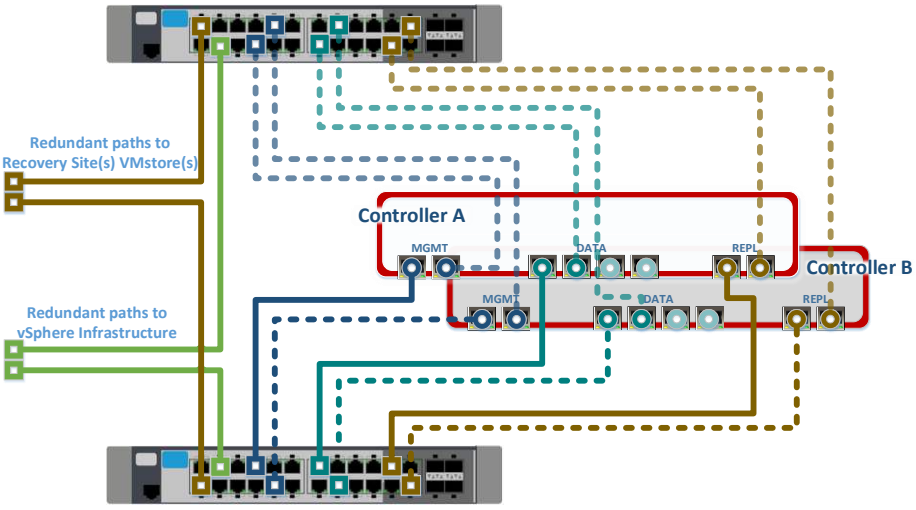

Figure 11: An illustration of a minimum-scale highly available networking configuration for a Tintri VMstore with dedicated Replication ports and paths

Detailed information on the topic of Data Protection and Data Recovery (DP/DR) techniques and considerations for the Tintri VMstore platform can be found in the 'Tintri Data Protection Overview and Best Practices' white paper (linked here, with URL also listed in the appendix).

LACP

Link Aggregation Control Protocol (LACP) can be used combine several physical NICs (Network Interface Cards) into a single logical interface. This can be used to provide a means for network traffic to negotiate a set of active ports for the purposes of bandwidth aggregation and load balancing, while at the same time maintaining support for network redundancy.

LACP is configurable independently on all VMstore network ports. This includes the Management and Storage networks, as well as the optional Replication network (if the applicable network card is present). It should be noted that the Tintri VMstore platform does not support mixing Ethernet link speeds in an LACP port-channel. In addition, all ports in a port-channel should have full-duplex mode enabled and identical configurations, such as using the same types of cable and transceiver.

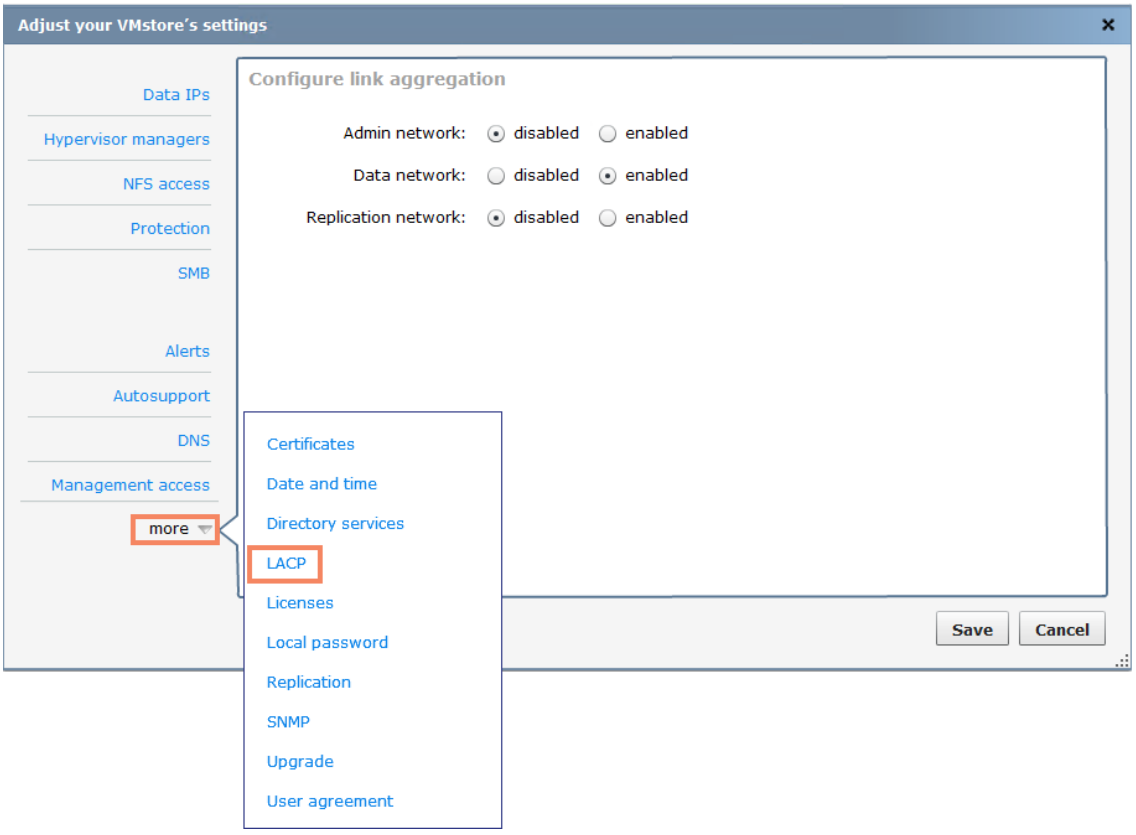

Figure 12: To enable LACP on the VMstore UI, go to 'Settings' and choose 'More', then the 'LACP' option

It should be noted that LACP port channels should be limited to the paths of a single VMstore controller only, i.e. port channel groups that span storage controllers in the same VMstore appliance are not supported. Please refer to Figure 13 which illustrates an unsupported LACP configuration with the Port-Channel spanning VMstore storage controllers.

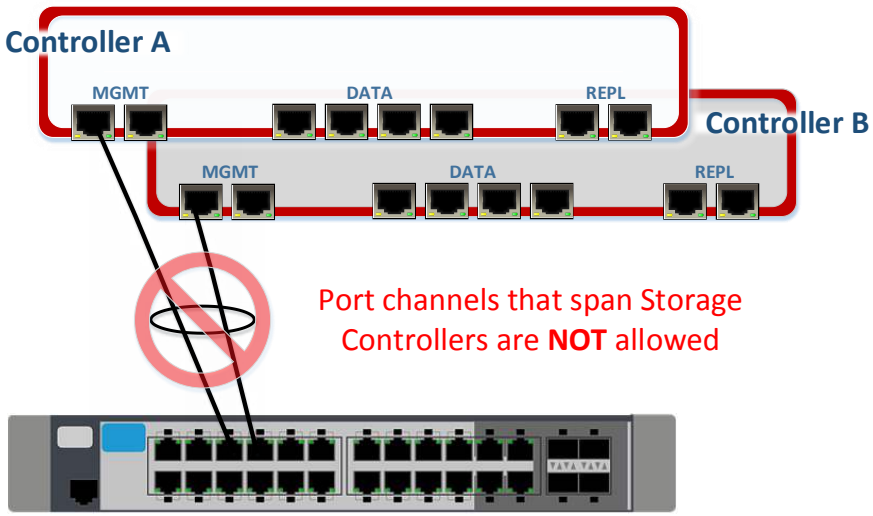

Figure 13: Example of an unsupported Tintri VMstore LACP port-channel configuration

In addition to delivering an increase in available network bandwidth which can be split across multiple dependent hosts, another recommended use of LACP with the Tintri VMstore is for the purpose of network redundancy across two or more network switches. In such a configuration, the ports of a specific VMstore network group (Storage/Management/Replication) will be combined into a single LACP port-channel but split between multiple physical network switches – refer to Figure 14. An LACP configuration such as this leads to half peak bandwidth in the event of a switch failure. Additionally, if a switch were to fail, there is no Tintri VMstore controller failover event.

To support a distributed LACP configuration, the switches used must also be confirmed to support MC-LAG (Multi-Chassis Link Aggregation Group) – please refer to your switch vendors' documentation for more information and details of the vendor specific implementation and naming convention for this protocol. Please note, the Tintri VMstore supports both layer 2 and layer 3 load balancing.

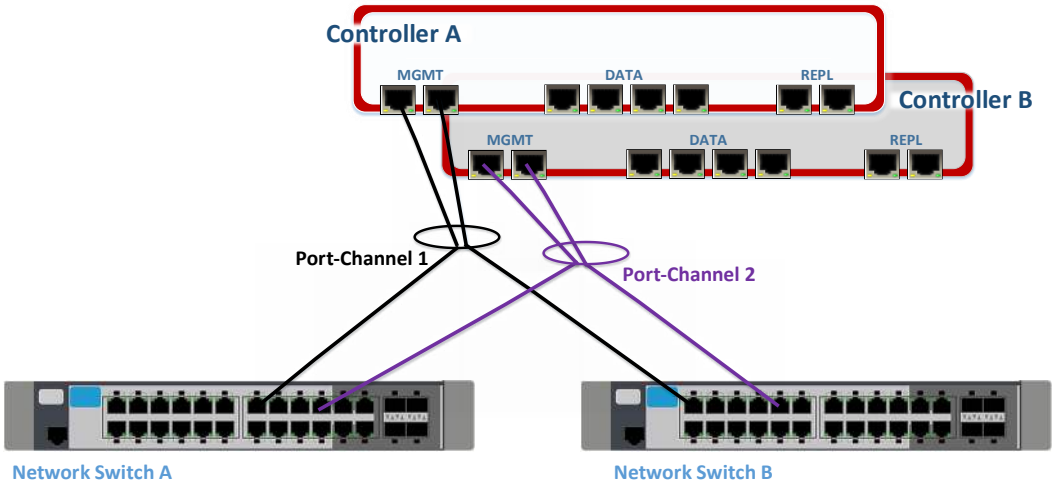

Figure 14: A Tintri VMstore with LACP enabled, network ports combined into port-channel and distributed across multiple network switches

In situations where LACP port-channels cannot be distributed across multiple switches, an additional supported LACP configuration is possible. This can be achieved by configuring LACP port-channels independently and connecting these to separate switch chassis. Figure 15 displays an example of a VMstore network port from Controller A configured with LACP on port-channel 1 on network switch A. In this configuration, if a network switch fails, the attached active Controller will failover. The VMstore becomes available through the surviving switch/controller. The LACP configuration in Figure 15 also keeps full bandwidth available, but at a greater risk of saturating interswitch links during normal operation.

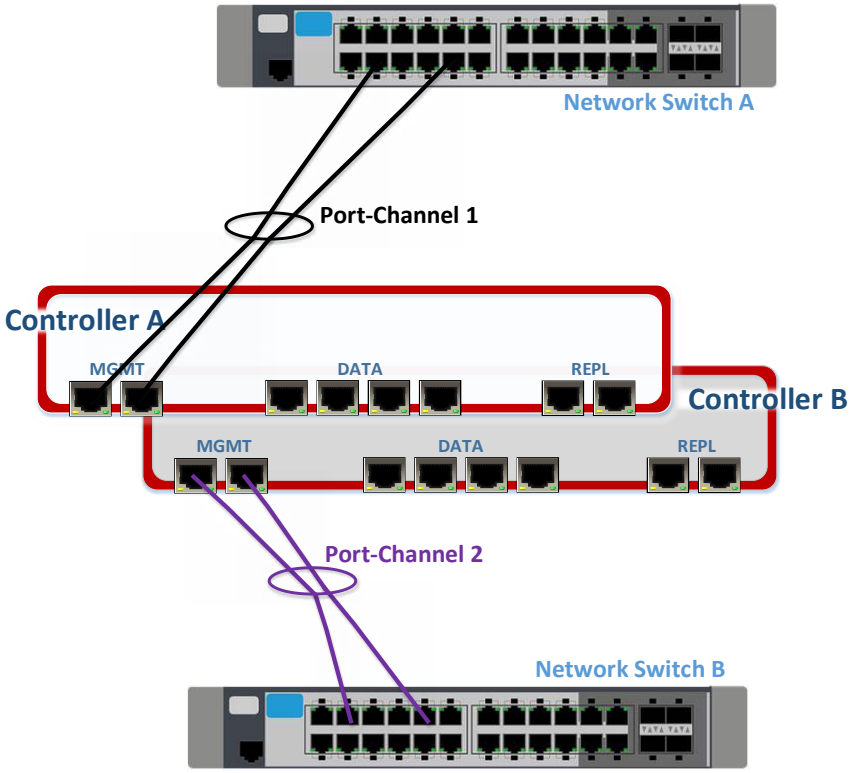

Figure 15: Tintri supported LACP configuration when LACP can't be distributed across multiple physical switches

LACP can be used to ensure the Tintri VMstore NICs are connected correctly and provide an additional point of validation of networking setup. For example, if the VMstore's physical NICs are not connected to the correct ports across the supporting switches, the LACP trunk between the physical switches (in dynamic mode) will not become active. Any issues will result in the VMstore ports (which are part of the port-channel) not becoming active – refer to Figure 16.

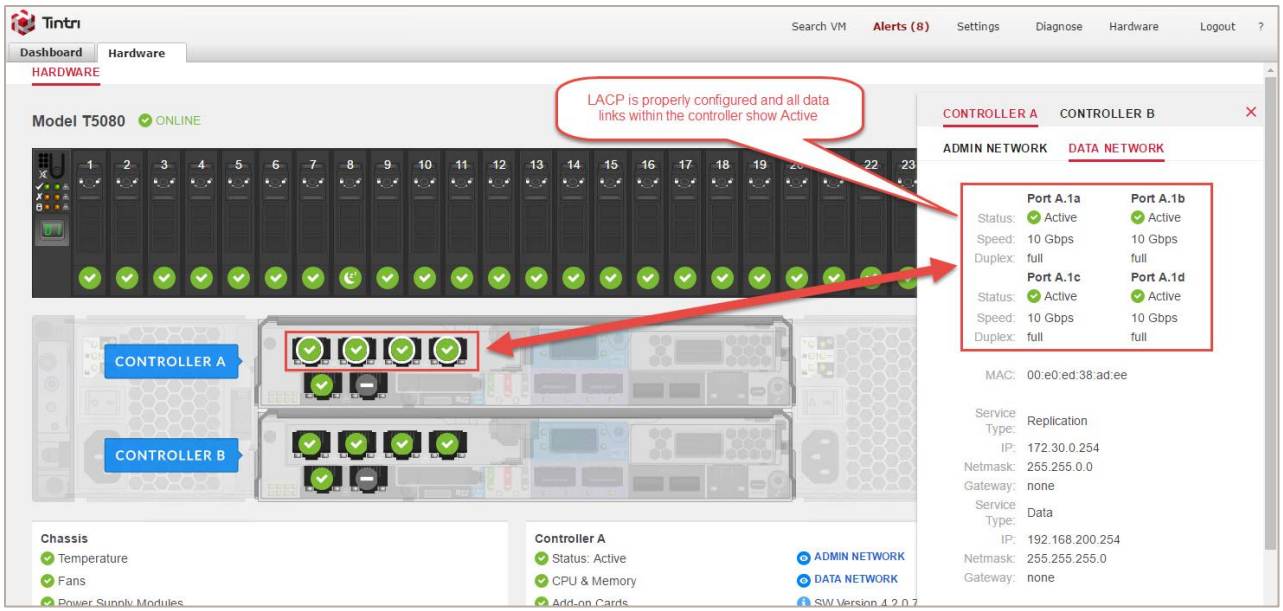

Figure 16: A Tintri VMstore (T5080 model) shown with Quad-port 10Gb card installed and LACP successfully enabled

Jumbo Frames

The Tintri VMstore platform supports the use of Ethernet Jumbo Frames – frames that can carry a payload greater than the standard size of 1,500 bytes. The MTU (Maximum Transmission Unit) supported by the Tintri VMstore is 9,000 bytes. When using Jumbo Frames, it is of the upmost importance to ensure that Jumbo Frames with consistent MTU settings are enabled across the entire networking path. When enabling Jumbo Frames, administrators should confirm that each of the VMstore's data interface, network switches including interswitch links (ISLs), and connected vSphere hosts are configured with corresponding Jumbo Frames settings.

Most networking components will support a MTU of greater than 9,000 – typically 9,216. The MTU utilized throughout the Ethernet fabric should never be less than that used for the endpoints. For more information, please refer to your switch vendor's documentation on this subject.

✧ *Ensure that when using Ethernet Jumbo Frames that the vSwitch settings, the physical network switch, and Tintri VMstore are all configured to use jumbo frames. Explicitly configure your physical Ethernet switches and vSwitches with an MTU size of at least 9,000 bytes (typically 9,216) when using Jumbo Frames.*

Misconfiguration of jumbo frames can result in an increased number of dropped packets across the network, resulting in sub-optimal performance. Due to the high frequency of issues resulting from misconfiguration seen in customer deployments, Tintri recommends against configuring Jumbo Frames unless you have a specific reason to do so. If required, Jumbo Frames can be enabled for a Tintri VMstore network interface by selecting 'Settings' within the VMstore's web management interface – refer to Figure 17.

Additional information on how to confirm a correct Jumbo Frame configuration for your VMstore data paths can be found in Section D of the Appendix in this document.

VLAN configuration

A VLAN is a logical broadcast domain that is used to segregate or bound traffic within a network. The reasons for using VLANs vary but in many environments, VLANs are considered common practice and are used to help simplify, control, and better secure enterprise virtual infrastructure.

✧ *In lieu of dedicated physical networks, VLANs can offer a flexible means to implement logically isolated networks that share common physical resources*

VLAN enabled ports are generally categorized in one of two ways, tagged or untagged. VLAN tagging (802.1Q) is supported on all Tintri VMstore interfaces. This allows one or more IP addresses and VLAN IDs to be specified on the VMstore's network interfaces. The VMstore's physical Management, Storage and Replication network ports can be connected to Ethernet switch trunk ports and pass traffic for multiple VLANs.

When configuring VMstore networking to connect to Ethernet switch ports configured as access ports, the VLAN field must remain empty.

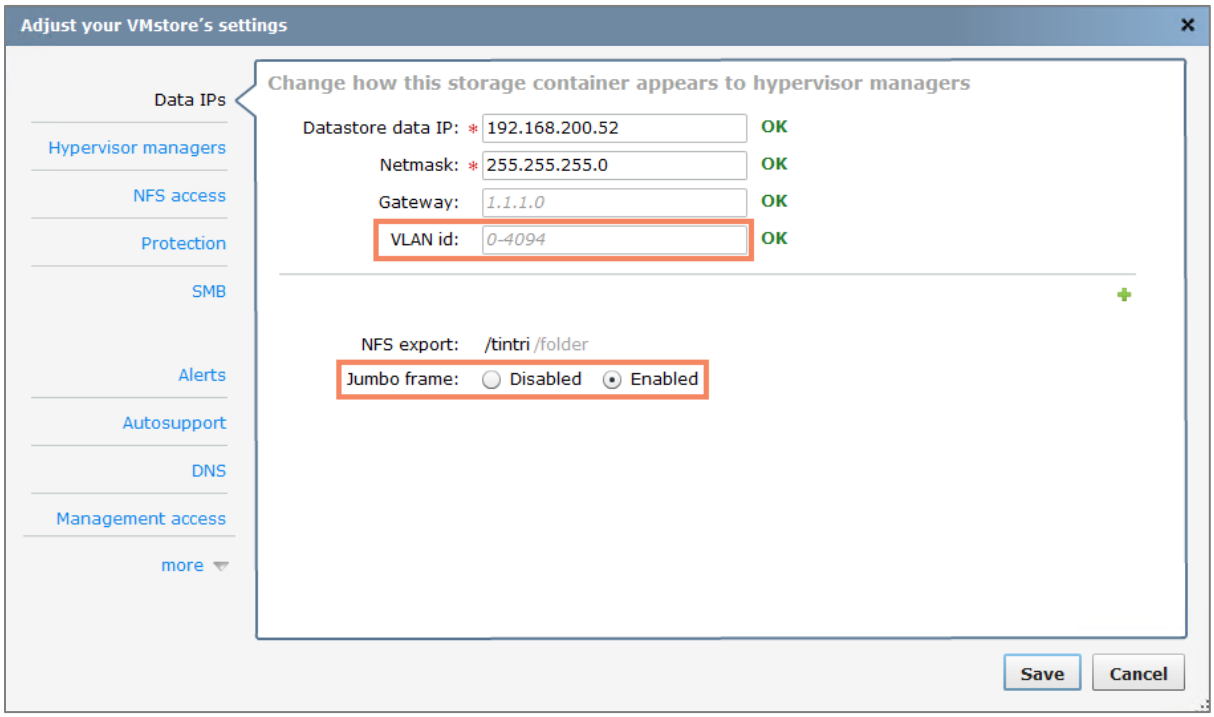

Figure 17: Use the Tintri VMstore 'Settings' tab to set VLAN IDs and Jumbo Frame configuration

vSphere Advanced Settings

This section discusses a number of advanced settings specific to the integration of NFS storage with vSphere hypervisors. Consistent with the VMware's KB article 2239, when you have multiple NFS Datastores per vSphere host, set the following three option values together:

- Nfs.MaxVolumes
- Net.TcpipHeapMax
- Net.TcpipHeapSize

NFS.MaxVolumes limits the number of concurrent NFS datastores your ESXi host can concurrently mount at a given time. *NFS.MaxVolumes* value should be modified to allow an increased number of NFS datastores. Tintri recommends increasing the *NFS.MaxVolumes* parameter to 256.

When *NFS.MaxVolumes* is increased, two additional parameters related to the in-host memory allocation required to support ESXi TCP/IP networking are also recommended be increased in unison. These parameters are *Net.TcpipHeapMax* and *Net.TcpipHeapSize* are relate to the amount of heap memory which will be allocated for managing VMkernel TCP/IP network connectivity. Tintri recommends setting both of these values to the maximum permissible values – since ESXi version 6.0, *Net.TcpipHeapMax* can be set to 1536 and *Net.TcpipHeapSize* can be set to 32 (both of these values representing the memory allocation amounts in MB).

✧ *Adjust the ESXi host Advanced Settings for Nfs.MaxVolumes, Net.TcpipHeapMax and Net.TcpipHeapSize to their maximum values on EVERY host*

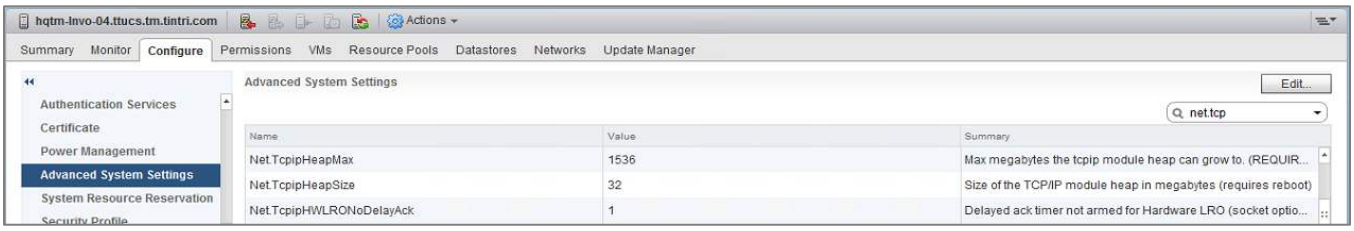

Figure 18: Adjusting the recommended advanced settings via the vCenter web client interface

Each of these values can be configured from within the vCenter web client through the 'Advanced System Settings' in the host configuration tab (see Figure 18) or alternatively, via the functionality provided by the Tintri vCenter Plug-In (see Figure 19).

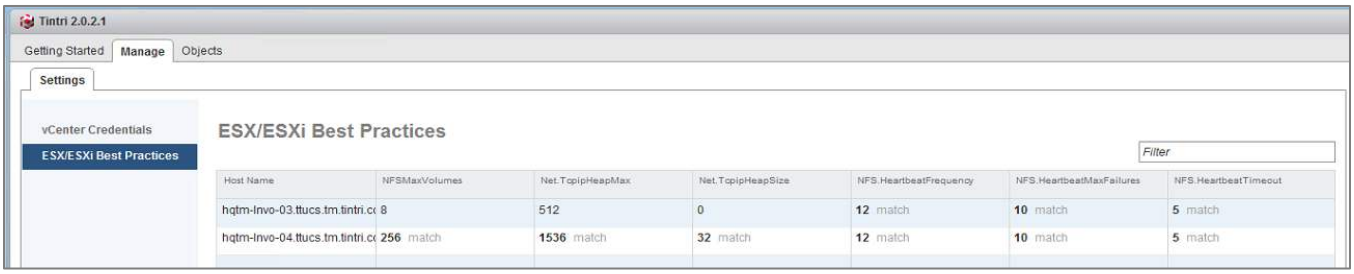

Figure 19: Using the Tintri vCenter plug-in functionality to automatically assign recommended advanced settings for all vSphere hosts

It is advised to always adjust the *NFS.MaxVolumes*, *Net.TcpipHeapMax* and *Net.TcpipHeapSize* settings together. Changing *Net.TcpipHeapMax* and *Net.TcpipHeapSize* requires a reboot of the vSphere host.

The *NFS.Heartbeat* settings affect how ESX/ESXi monitors the availability of its NFS datastore volumes. Heartbeat settings must be identical on all ESX servers or data corruption may result. Note that heartbeat settings should already be set to the correct values on ESXi versions later than 5.x. For information on the recommended *NFS.Heartbeat* settings for older versions of ESX, please refer to the Appendix section 'NFS.Heartbeat'.

✧ *Confirm that the ESXi host Advanced Settings for NFS.Heartbeat conforms to recommended values on EVERY host*

---

vSphere Virtual Disk Formats

Three distinct options are available when creating or cloning NFS-based virtual disks (vDisks) for use with a new or existing virtual machine. These three options fall into one of two categories; Thin, where logical capacity for the disk is allocated as it is consumed, or Thick, where the entire logical capacity of the virtual disk is allocated at the time of creation. Thick disks have two consumption models: the first is Thick Eager Zero, where the full allocated capacity is completely zeroed at the time of creation; the second is Thick Lazy Zero, where that same fully allocated logical capacity is physically consumed or zeroed as it is accessed.

By default, virtual disks that are to reside on NFS datastores should be provisioned using the Thin format and Tintri recommends using Thin virtual disks as the provisioning format for all situations that do not explicitly specify otherwise.

If a virtual disk is to support clustering solutions for the purposes of Fault Tolerance (for example Oracle RAC or Microsoft Clustering), you cannot provision the disk using the Thin format using VMware vSphere. In these scenarios, the virtual disk should be provisioned using the Thick Eager Zero option and the vSphere 'Multi-Writer' setting should be utilized. This can be done at the time of creation via the vCenter user interface (as of ESXi 6.x, refer to Figure 20) or by amending a powered-off virtual machine's VMX file to include the line *"scsiX.Y.sharing = "multi-writer"* (replace X.Y – refer to VMware KB1034165).

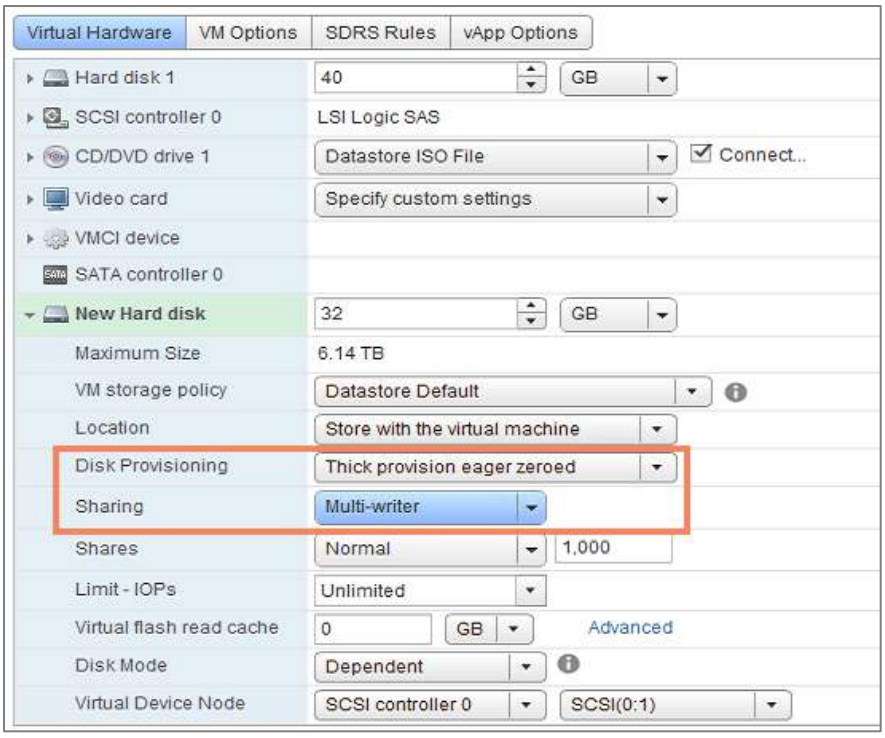

Figure 20: Choosing the 'Multi-Writer' option during the creation of a new virtual disk in the vCenter web client interface

Please note that Thick provisioned virtual disks are only available when the NFS VAAI command set is supported. Please refer to the section 'Tintri VAAI provider' on how to enable VAAI for use with Tintri VMstore hosted datastores.

✧ *Tintri recommends using Thin virtual disks as the provisioning format for virtual disks in all scenarios unless specific clustering functionlaity is required*

For additional information on the various space savings considerations associated with NFS virtual disks hosted on the Tintri VMstore platform, please refer to Appendix section B of this document.

Tintri VAAI Provider

VMware's vSphere Storage APIs for Array Integration (VAAI) are utilized to deliver enhanced integration between storage and vSphere. Using VAAI, vSphere can offload a number of storage-specific operations to a supporting storage platform such as the Tintri VMstore. This feature delivers efficiencies to the virtualized environment by reducing the overall load and number of operations taking place on host-based resources.

As is standard, VMware's VAAI operations for NAS are only available through the use of a specific plug-in which is required to be installed at the hypervisor level. Tintri have developed a supporting VAAI plug-in and this is available free to Tintri customers, with all supporting documentation available from the Tintri support site. Tintri's VAAI plug-in supports each of the three core VAAI NAS primitives – these include: Full File Clone, Extended Statistics, and Reserve Space.

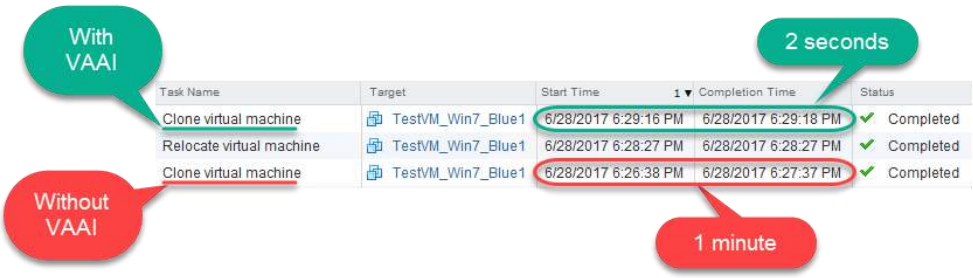

Figure 21: Examining the difference in VM Clone completion time when using Tintri VAAI enabled clone functionality versus without. Default VM with active Windows 7 image and 32GB vDisk.

Full File Clone, or Full Copy, is similar to the Extended Copy (XCOPY) hardware acceleration primitive provided for block storage arrays. This primitive enables virtual disks to be cloned by the NAS device rather than by using the hypervisor Data Mover, which consumes ESXi host CPU and memory resources as well as network bandwidth. The reduction in completion times for clone operations using Tintri VAAI functionality can be significant, as seen in Figure 21.

The Extended Statistics primitive enables visibility into space usage on NAS datastores. This is especially useful for thin-provisioned datastores (such as the Tintri VMstore) because it enables vSphere to display actual space usage statistics in situations where oversubscription was used.

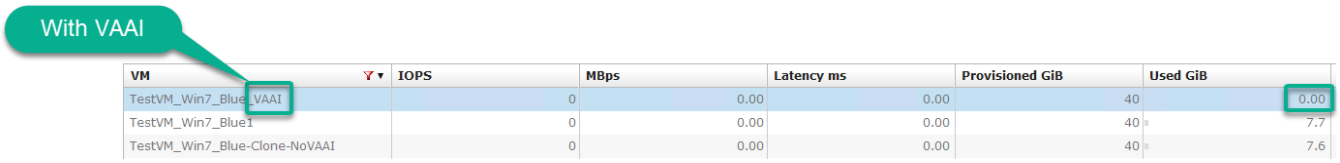

Figure 22: Displaying capacity consumption for cloned VMs which were created with Tintri VAAI enabled.

The Reserve Space reserve space operation enables the creation of thick VMDK files on NAS datastores, so administrators can reserve all of the space required by a VMDK, even when the datastore is NAS. Tintri best practices recommend that all VMDK files are created as Thin, however there are situations where Thick provisioned VMDK files are required – refer to section: Virtual Disk Formats.

✧ *The Tintri VAAI plug-in must be installed on each vSphere host that is intended to leverage the supported functionality. If not installed, it is recommended that the Tintri VMstore UI or Tintri vSphere Plugin be used to create space-efficient clones instead of the vSphere Client.*

Installation can be performed via ESXCLI or through VMware Update Manager (VUM). Full instructions on both methods are included in the version appropriate Tintri VAAI Plug-In Release Notes document – available from the Tintri support portal.

- NOTE:
  - PowerCLI functionality can also be employed to install the Plug-In using either the Install-VMHostPatch command, or alternatively through remote ESXCLI functions as shown in the example below – see Figure 23. After the plug-in is installed, a host reboot is required.

```
PowerCLI C:\Windows\system32> $ESXHOST

Name                 ConnectionState PowerState NumCpu CpuUsageMhz CpuTotalMhz   MemoryUsageGB   MemoryTotalGB Version
----                 --------------- ---------- ------ ----------- -----------   -------------   ------------- -------
hqtm-lnvo-03.ttuc... Connected       PoweredOn      28          62       67172           3.669         255.185   6.5.0


PowerCLI C:\Windows\system32> $ESXCLI = Get-EsxCli -VMHost $ESXHOST.name
PowerCLI C:\Windows\system32> $ESXCLI.software.vib.list() | select Vendor | findstr -i Tintri
PowerCLI C:\Windows\system32> Set-VMHost $ESXHOST -State maintenance

Name                 ConnectionState PowerState NumCpu CpuUsageMhz CpuTotalMhz   MemoryUsageGB   MemoryTotalGB Version
----                 --------------- ---------- ------ ----------- -----------   -------------   ------------- -------
hqtm-lnvo-03.ttuc... Maintenance     PoweredOn      28         353       67172           3.417         255.185   6.5.0


PowerCLI C:\Windows\system32> $ESXCLI.software.vib.install("/tmp/TINTRI-ESX-5.x-6.x-TintriVaaiNasPlugin-1.0.0.8-offline_bundle.zip",$false,$true,$true,$true,$false,$null,$null,$null)


Message        : The update completed successfully, but the system needs to be rebooted for the changes to be effective.
RebootRequired : true
VIBsInstalled  : {Tintri_bootbank_vmware-esx-TintriVaaiNasPlugin_1.0.0.8-2.8}
VIBsRemoved    :
VIBsSkipped    :
```

Figure 23: Installing the Tintri VAAI VIB remotely via VMware PowerCLI commands.

For VMware Horizon VDI (Virtual Desktop Environments) that employ VMware View Composer servers, cloning operations will rely on VCAI (View Composer API) functionality to offload resource intensive provisioning tasks to Tintri VMstore. VCAI requires that the involved ESXi servers support VAAI NAS functionality and as such, also require the Tintri VAAI Plug-In is installed. For more information, please refer to the View Composer API for Array Integration (VCAI) from VMware® on Tintri® VMstore™ Platform document.

Monitoring VAAI with NFS

At times, there may be requirements to examine the on-host VAAI NFS activity in closer detail. To do so, users should tail the VPXA log file, while filtering for entries containing the string “VAAI”. Figure 24 displays the command structure and subsequent sampled output of a Tintri VAAI enabled clone operation completing. This information can be used to determine the appropriate timestamps and lines within the log file that are relevant to the VAAI enabled operation.

```
[root@HQTM-LNVO-04:/var/log] cat /var/log/vpxa.log | grep -i VAAI
2017-07-27T23:48:07.223Z info vpxa[878FB70] [Originator@6876 sub=Libs opID=ProvisioningWizard-addMulti-405117-ngc-4f-01-c1]  VAAI-NAS :: SvaNasPlugin: SUCCESSES: RsrvSpace [0] Cln-Ful
2017-07-27T23:48:07.223Z info vpxa[878FB70] [Originator@6876 sub=Libs opID=ProvisioningWizard-addMulti-405117-ngc-4f-01-c1]  VAAI-NAS :: SvaNasPlugin: FAILURES: RsrvSpace [0] Cln-Full
2017-07-27T23:48:07.223Z info vpxa[878FB70] [Originator@6876 sub=Libs opID=ProvisioningWizard-addMulti-405117-ngc-4f-01-c1]  VAAI-NAS :: TintriVAAINasPlugin: SUCCESSES: RsrvSpace [0]
2017-07-27T23:48:07.223Z info vpxa[878FB70] [Originator@6876 sub=Libs opID=ProvisioningWizard-addMulti-405117-ngc-4f-01-c1]  VAAI-NAS :: TintriVAAINasPlugin: FAILURES: RsrvSpace [0] C
2017-07-27T23:48:07.223Z info vpxa[878FB70] [Originator@6876 sub=Libs opID=ProvisioningWizard-addMulti-405117-ngc-4f-01-c1]  VAAI-NAS :: vmfsNasPlugin: SUCCESSES: RsrvSpace [0] Cln-Fu
2017-07-27T23:48:07.223Z info vpxa[878FB70] [Originator@6876 sub=Libs opID=ProvisioningWizard-addMulti-405117-ngc-4f-01-c1]  VAAI-NAS :: vmfsNasPlugin: FAILURES: RsrvSpace [0] Cln-Ful
2017-07-27T23:48:07.223Z info vpxa[878FB70] [Originator@6876 sub=Libs opID=ProvisioningWizard-addMulti-405117-ngc-4f-01-c1]  VAAI-NAS :: NAS Mapping Used successfully for 2 times
2017-07-27T23:48:07.930Z info vpxa[878FB70] [Originator@6876 sub=Libs opID=ProvisioningWizard-addMulti-405117-ngc-4f-01-c1] TintriVAAI: CloneFile[3CCA460] src /vmfs/volumes/c01b0983-
s/volumes/c01b0983-45fa2a7a/tingle1006_02/tingle1006_02-flat.vmdk flags 10: err 0/0
2017-07-27T23:48:07.930Z info vpxa[878FB70] [Originator@6876 sub=Libs opID=ProvisioningWizard-addMulti-405117-ngc-4f-01-c1] VAAI-NAS [TintriVAAINasPlugin : /vmfs/volumes/c01b0983-45fa
006_02/tingle1006_02-flat.vmdk] succeeded.
2017-07-27T23:48:08.594Z info vpxa[878FB70] [Originator@6876 sub=Libs opID=ProvisioningWizard-addMulti-405117-ngc-4f-01-c1] TintriVAAI: CloneFile[3AFFEA0] src /vmfs/volumes/c01b0983-
mfs/volumes/c01b0983-45fa2a7a/tingle1006_02/tingle1006_02_1-flat.vmdk flags 10: err 0/0
2017-07-27T23:48:08.594Z info vpxa[878FB70] [Originator@6876 sub=Libs opID=ProvisioningWizard-addMulti-405117-ngc-4f-01-c1] VAAI-NAS [TintriVAAINasPlugin : /vmfs/volumes/c01b0983-45fa
006_02/tingle1006_02_1-flat.vmdk] succeeded.
```

Figure 24: Examining Tintri VAAI related actions on a vSphere host can be performed by monitoring the VXPA log

Tintri vCenter Plug-In

The Tintri vCenter Plug-In (VCP) is a vCenter enhancement which is installed within a vCenter Server instance and allows users of Tintri VMstore storage to introduce Tintri-specific monitoring and

management capabilities directly into the vSphere vCenter web client. The Tintri VCP is free to all users of the Tintri VMstore platform and can be downloaded from the Tintri support portal. All supporting documentation with instructions for installation and administration can be found in the same location.

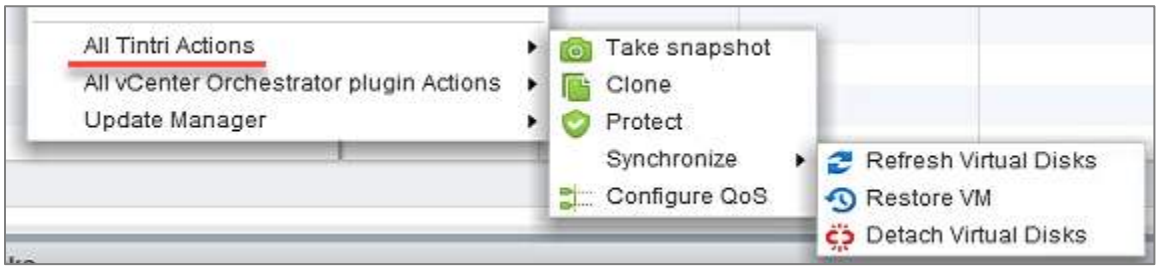

Figure 25: A selection of the active storage-specific actions which can be run against Tintri hosted VMs when using the Tintri vCenter Plug-In

---

✧ *Installing the Tintri vCenter Plug-In will allow vCenter administrators and users monitor the storage-specific metrics of their Tintri VMs as well as enact per-VM storage related actions from within the vCenter web client interface*

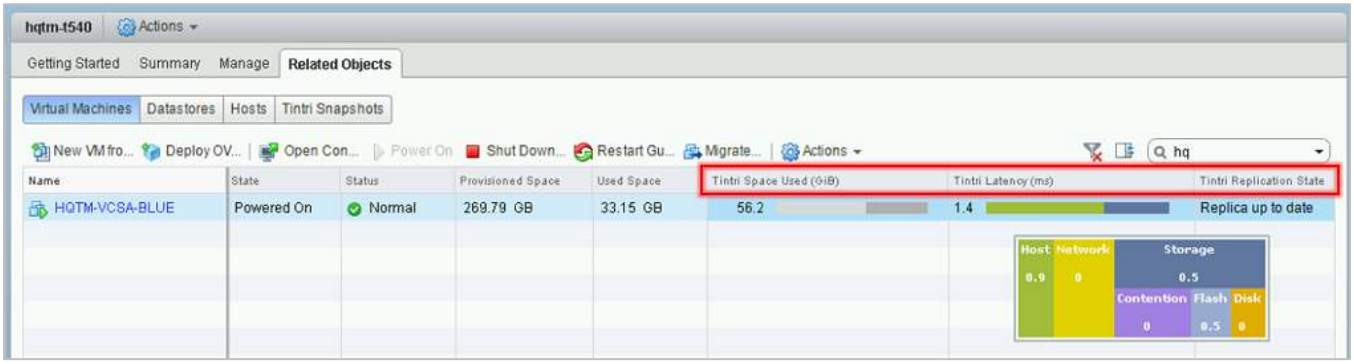

Figure 26: Through use of the Tintri VCP, full stack performance metrics for Tintri VMstore hosted VMs become available for display in the vCenter web client. In addition, additional information with regards capacity and replication is available.

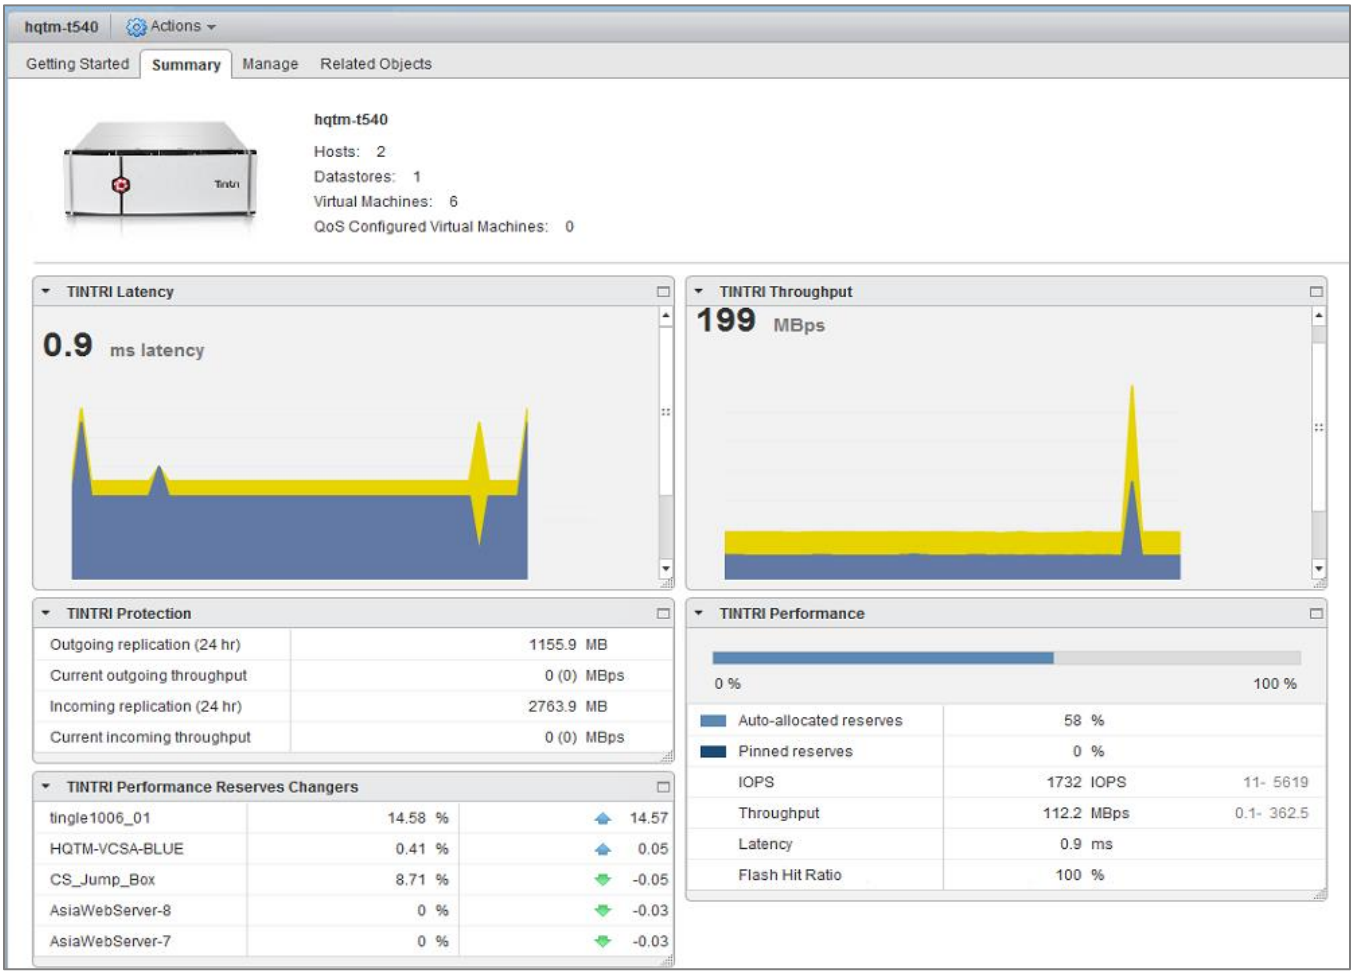

Figure 27: A selection of the Tintri VMstore specific metrics which can be monitored from the vCenter web client when using the Tintri vCenter Plug-In

Once the Tintri VCP is installed with the necessary vCenter and VMstore connections established, vSphere administrators will be provided with access to Tintri-specific per-VM information, metrics, and operations from within the vSphere web client. These capabilities include the display of both real-time and historical per-VM performance, confirmation of the current protection status for each VM, and a number of additional VM-level actions specific to Tintri's vCenter integration capabilities.

Please refer to the Figures in this sub-section for examples of the VMstore and hosted VM related metrics and actions available to vCenter administrators when using the Tintri vCenter Plug-In. The information and actions available to specific vCenter user groups can be limited via vCenter Role Based Access Control – please refer to sub-section Access Control for more information.

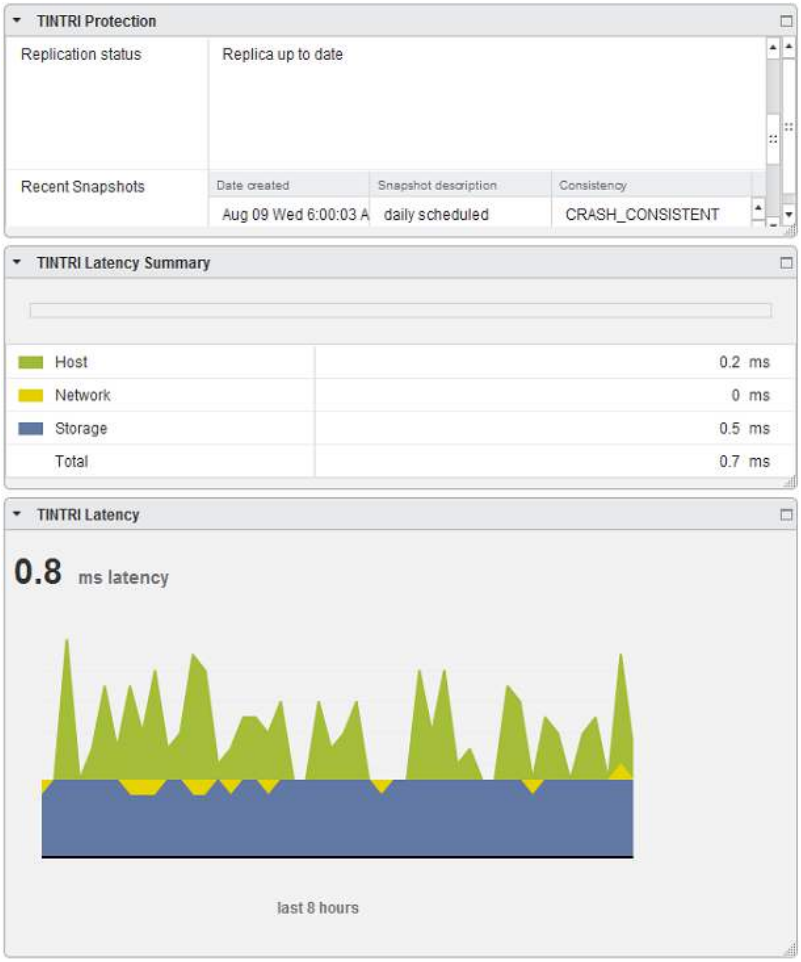

Figure 28: Examining a Tintri hosted VM's array-based protection status and performance statistics in the vCenter we client interface

Access Control

The Tintri vCenter Web Client Plug-In offers role-based access control to restrict what Tintri-related actions authorized users of an enterprise's vCenter can perform. Custom Tintri roles and privileges integrate with VMware's vSphere RBAC (Role Based Access Control) mechanism to enable granular control over Tintri storage polling and actions. Refer to Figure 29 which shows the menu location within vCenter as well as a listing of the various privileges available to vCenter users and administrators.

LDAP (Lightweight Directory Access Protocol) RBAC for the Tintri VMstore interface is also available via integration with an organization's Active Directory structure. Enterprise administrators who wish to implement LDAP integration with Tintri should refer the relevant Administrator Guide, available from the Tintri support portal.

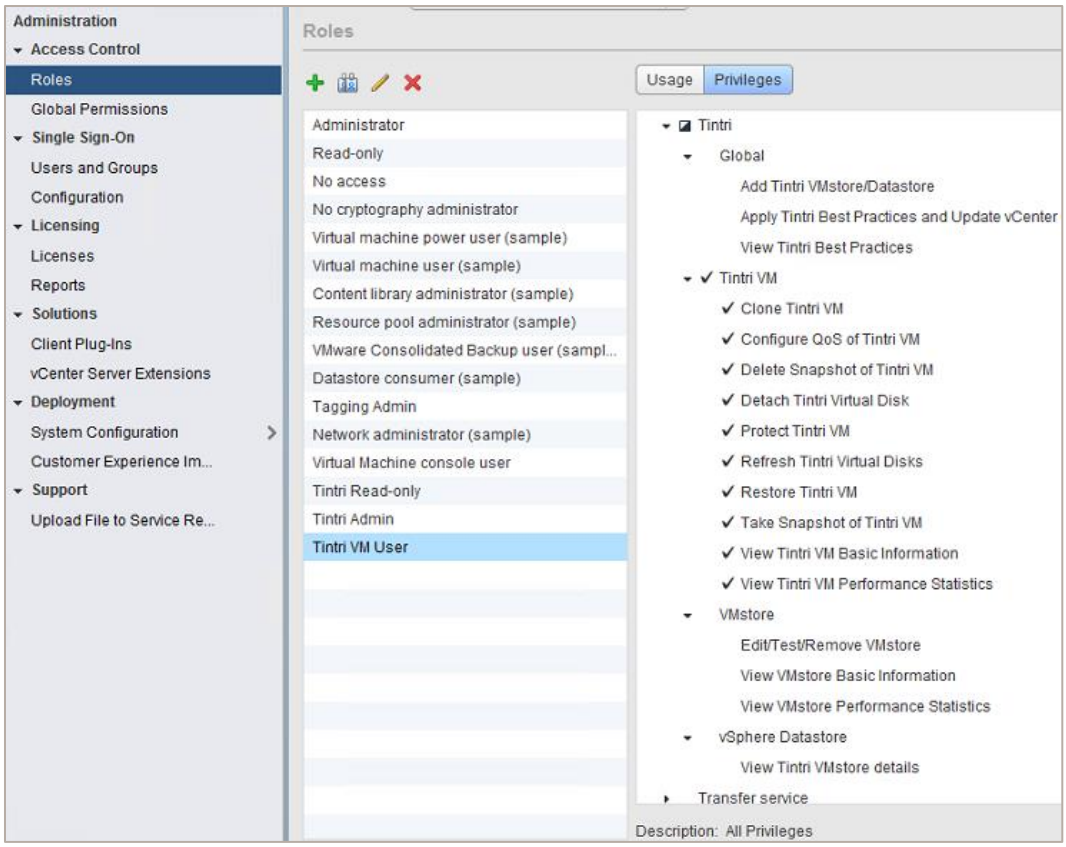

Figure 29: Tintri-specific privileges available within vSphere's Role-Based Access-Control which can be assigned/removed from designated vCenter user and administrator groups

VMware Tools

VMware Tools™, which is installed within each VM's guest operating system, ensures that the guest OS is optimized for running on vSphere virtualized infrastructure. VMware Tools not only enhances the performance and management of virtual machines, but also maintains appropriate disk access timeout values for the guest OS. This ensures that any possible failover events that can occur in the underlying storage network are handled gracefully by each VM's guest OS.

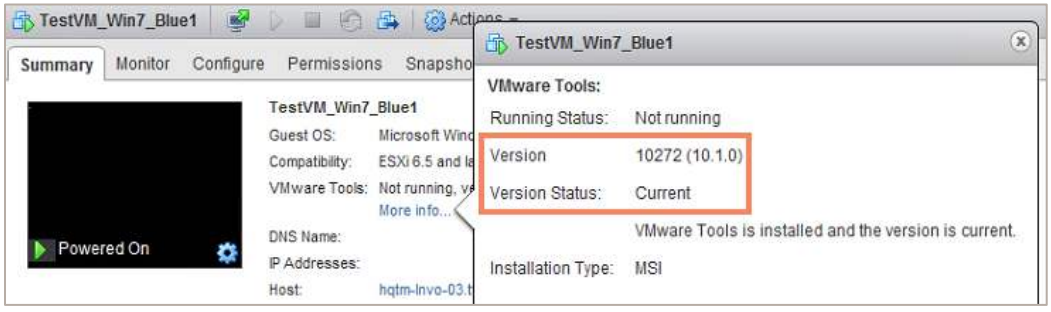

Figure 30: Confirming a VM's VMware Tools version can be done from the VM Summary page in the vCenter web client interface

It is generally considered a best practice to install an up to date version of VMware Tools in each VM guest OS. Verifying the version of VMware tools installed in a VM is a straightforward task and can be

done by looking at the VM summary tab in the vCenter web client, or scripted via PowerCLI as shown in Figure 31.

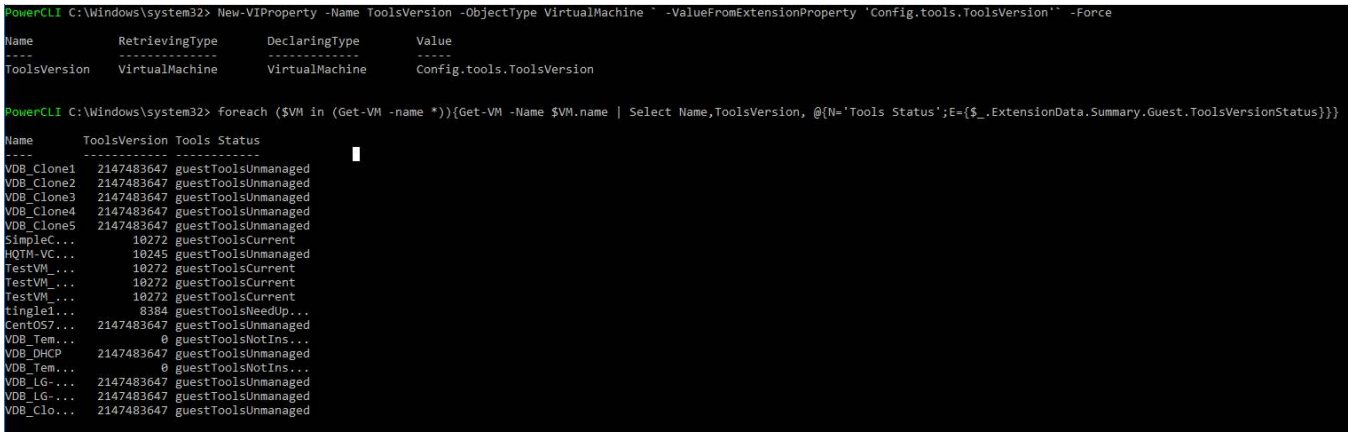

Figure 31: A PowerCLI command that can be used to remotely poll a group of VM's VMware Tools status and version

Conclusion

This document was produced to assist enterprise architects, data-center administrators, and Tintri end-users better understand, integrate, and realize the operational efficiencies available to VMware based virtual infrastructures hosted on Tintri VMstore platform.

Within the content presented here, there has been guidance on how to best engineer a Tintri hosted vSphere environment. This guidance can never be considered final and absolute due to the continuing evolution of virtualized enterprise solutions and infrastructures and each environments' own unique criteria and business challenges. However, it will provide a solid grounding to those looking to address in-house hosted hypervisor-integrated storage projects and solutions that offer the capabilities and efficiencies of cloud based methodologies and architectures. As such, it is always best practice to contact your local Tintri representative to assist in your journey to taking advantage of the opportunities provided by the Tintri product suite.

Appendix

A: Reference

Tintri Homepage - https://www.tintri.com/

Tintri Solutions & White Papers Portal – https://www.tintri.com/resources

Tintri Technical Documentation & Knowledge Base - https://knowledge.tintri.com/

Tintri Video Presentations – https://www.youtube.com/user/TintriInc

Tintri Technical Demonstrations – https://www.youtube.com/channel/UCP1IFC6OYgluF2di11mEIDg

Tintri Sales Enquires - https://www.tintri.com/contact-a-specialist

Tintri Support Portal – https://support.tintri.com

Tintri API, Automation & Scripting – https://github/Tintri

Tintri Twitter Page – https://twitter.com/Tintri

Tintri Data Protection and Best Practices White paper:
https://www.tintri.com/sites/default/files/field/pdf/whitepapers/tintri-data-protection-overview-and-best-practices-white-paper.pdf

VMware Documentation Portal – https://docs.vmware.com/

B: Tintri VMstore Related Information

Virtual Machines Space Saving Considerations

Depending on the virtual disk provisioning format used, differing logical and physical space consumption models will exist. The following section will identify a number of the key capacity considerations relating to the available virtual disk formats when using Tintri VMstore hosted datastores.

✧ *Thin provisioning is required to fully realize the capacity efficiencies associated with modern storage platforms that support data reduction technologies*

The default and recommended provisioning format for virtual disks on NFS hosted datastores is Thin. This means that the vDisk will be provisioned in a manner that presents the in-guest Operating System and hosting hypervisor with the full capacity, but only allocates and consumes this logical and physical capacity at the management and underlying storage layers as the space is actually accessed and used.

The Thin format is ideal for virtualized environments where multiple cohabiting virtual machines all require excess capacity to be available, but they are unlikely to fully consume that capacity simultaneously. Thin provisioning is required to fully realize the capacity efficiencies associated with modern storage platforms that support data reduction technologies through data-aware storage techniques. Thin provisioned virtual disks fully benefit from all space saving features on the VMstore,

including on-demand logical provisioning, compression, and deduplication (all where applicable, based upon VMstore model).

The alternative provisioning model of Thick requires the up-front reservation of physical capacity at the storage layer to be equal to the provisioned size of the virtual disk. This up-front reservation is by design and confirms with the fundamental reasoning for using Thick provisioned virtual disks. The core function of thick provisioned vDisks is to avoid any possible over-consumption of available capacity. To do this, storage management systems should enact a means to confirm reservation all physical capacity at the time of creation, regardless of the format of Thick provisioning used or the storage access profile following the creation. For this reason, Thick provisioned vDisks in their default state will not benefit from any space savings on the Tintri VMstore platform.

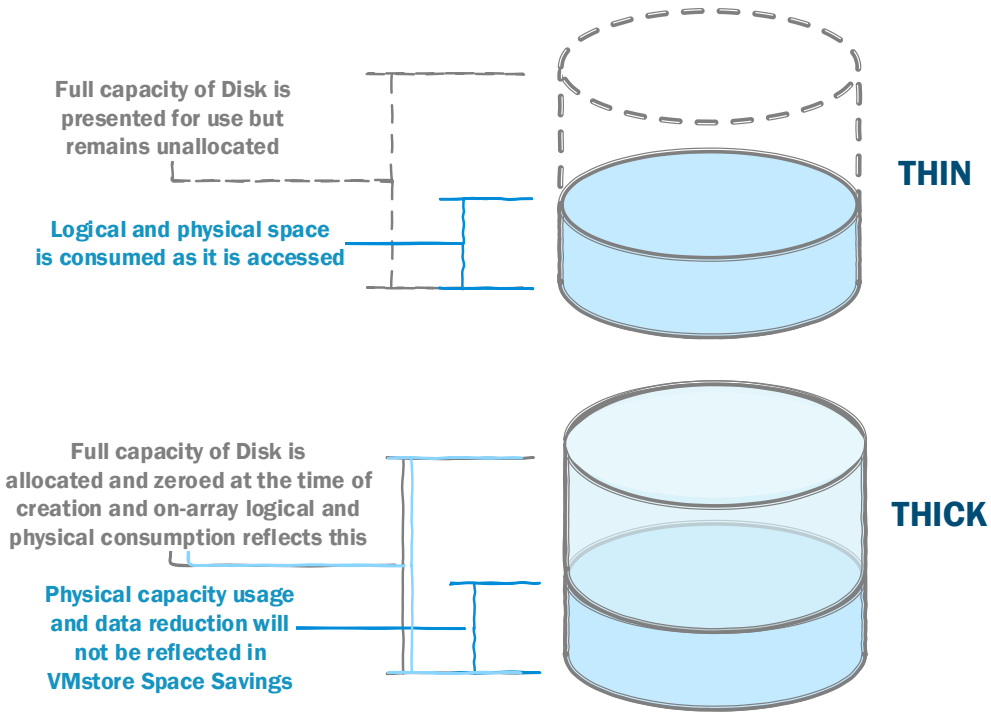

Figure 32: Illustrating the differences in capacity consumption models for Thick and Thin virtual disk formats on the Tintri VMstore platform

---

✧ *Thick provisioned vDisks in their default state will not benefit from any space savings on the Tintri VMstore platform*

Tintri end-users have the ability to override the default space consumption behaviors associated with thick vDisks and enjoy the same space savings available to Thin vDisks. This can be done while still preserving the presentation of a genuine THICK provisioned vDisk to the hypervisor layer. Tintri recommends that users should set their VMstore to 'maximize space savings' by ignoring the Thick provisioning methodology (while still supporting Thick provisioning requests). This can be done via the VMstore settings page in the VMstore GUI – refer to Figure 33.

- *Utilizing the 'Maximize Space Savings' option on your VMstore in conjunction with NFS VAAI support at the hypervisor layer will allow Thick provisioned virtual disks to enjoy all the data reduction benefits of Tintri's data-aware deduplication and compression abilities.*

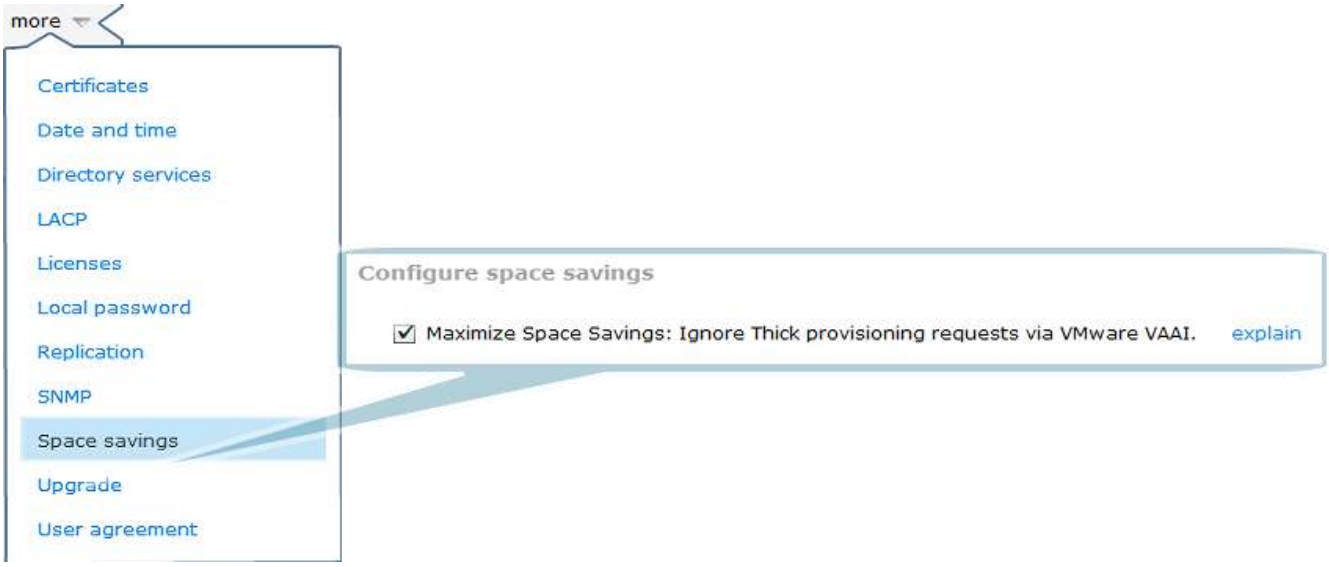

Figure 33: Configuring the Tintri VMstore to accept and process requests for Thick provisioned vDisks while maintaining Thin provision mechanics and efficiencies

To better understand the space savings associated with any particular virtual disk (or group of), it is first required to define the various capacity consumption essentials as they relate to Tintri VMstore managed storage. Because the Tintri VMstore platform provides thinly provisioned storage to the host layer, there will typically be a difference between the logically allocated and logically consumed capacity on the array. In addition, the Tintri VMstore is a data-aware storage appliance, meaning that deduplication (used to avoid repeatedly storing the same data) and compression (used to minimize the amount of data written to the storage media) will occur. These operations happen inline to the data path and will result in differences being seen in the logically consumed capacity and the physically consumed capacity. These differences will be referred to in terms of 'Space Savings' throughout this document and a graphical representation of these Space Savings factors can be seen in Figure 34.

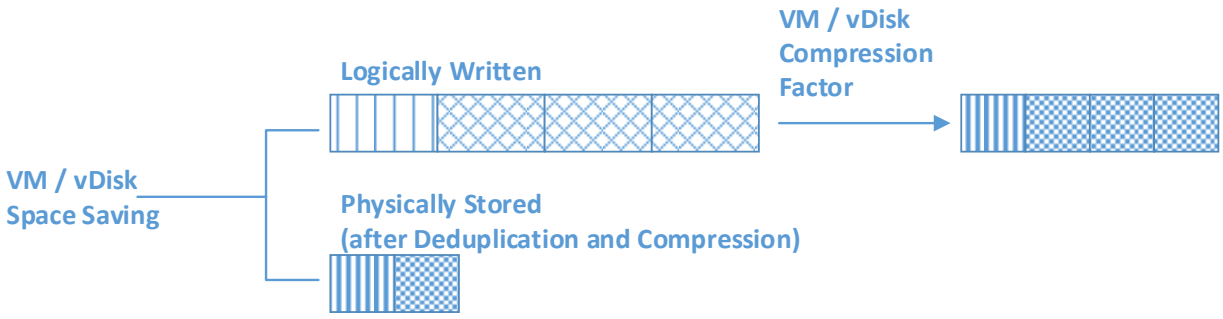

Figure 34: Illustrating the data-aware I/O processing activities, their sequence and impact to VM / vDisk Space Savings on the Tintri VMstore platform

For users of the Tintri VMstore platform, the capacity levels of interest are as follows:

- **Logical Footprint**
  - The logical size of data accessible to the client, space reserved due to Thick Provisioning, and space used by snapshots.
- **Logical Stored**
  - The logical size of the data that is stored and space reserved due to thick provisioning. The deduplication process filters data in the logical footprint to remove duplicate data. After duplicates are removed, the remaining data is what will be logically stored.
- **Physically Stored**
  - After compression, the actual size of the logically stored data and space reserved due to Thick Provisioning.

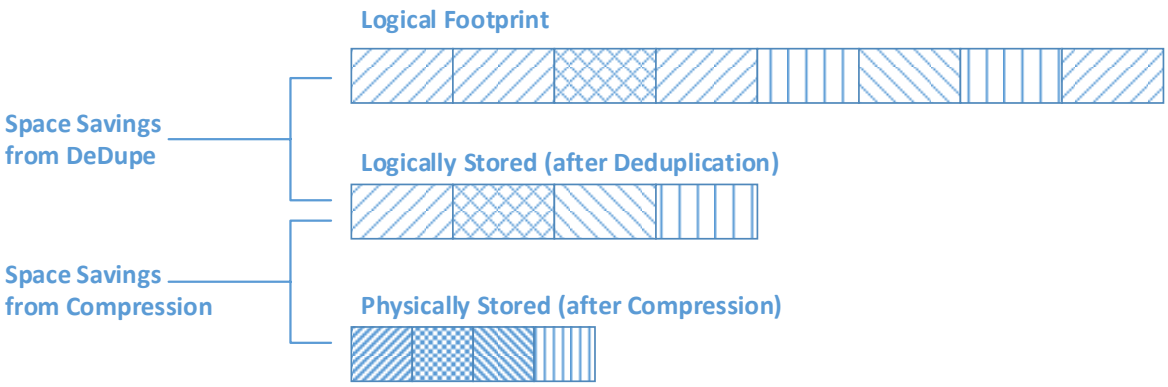

Figure 35: Tintri is a data-aware storage platform which only stores unique and compressed data to media

:

Capacity and space-saving metrics for each of the presented datastores, hosted VMs and their constituent vDisks can be examined via the VMstore user interface. Datastore level space savings can be found on the main dashboard under the physical space gauge, as a link labeled space savings. To view the space savings at the Tintri VMstore system level, click on space savings link and the right-hand column will expand to show the space savings on the entire system. The Space savings factor measures the total space savings on the VMstore due to compression and deduplication, and is computed as a ratio of the logical footprint and physical space used.

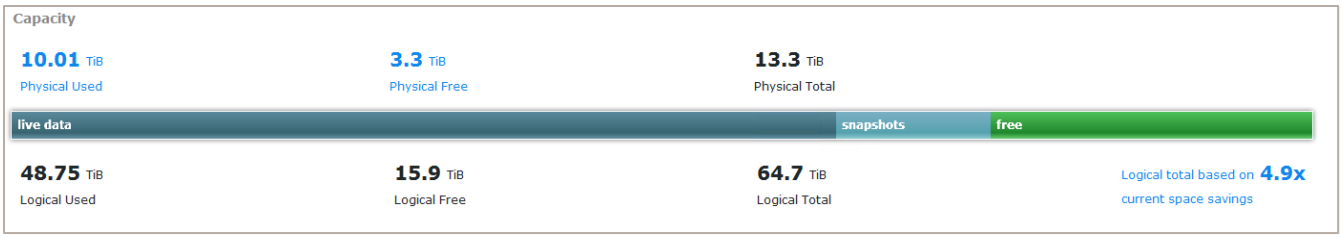

Figure 36: Examining the capacity usage dashboard on the Tintri VMstore user interface

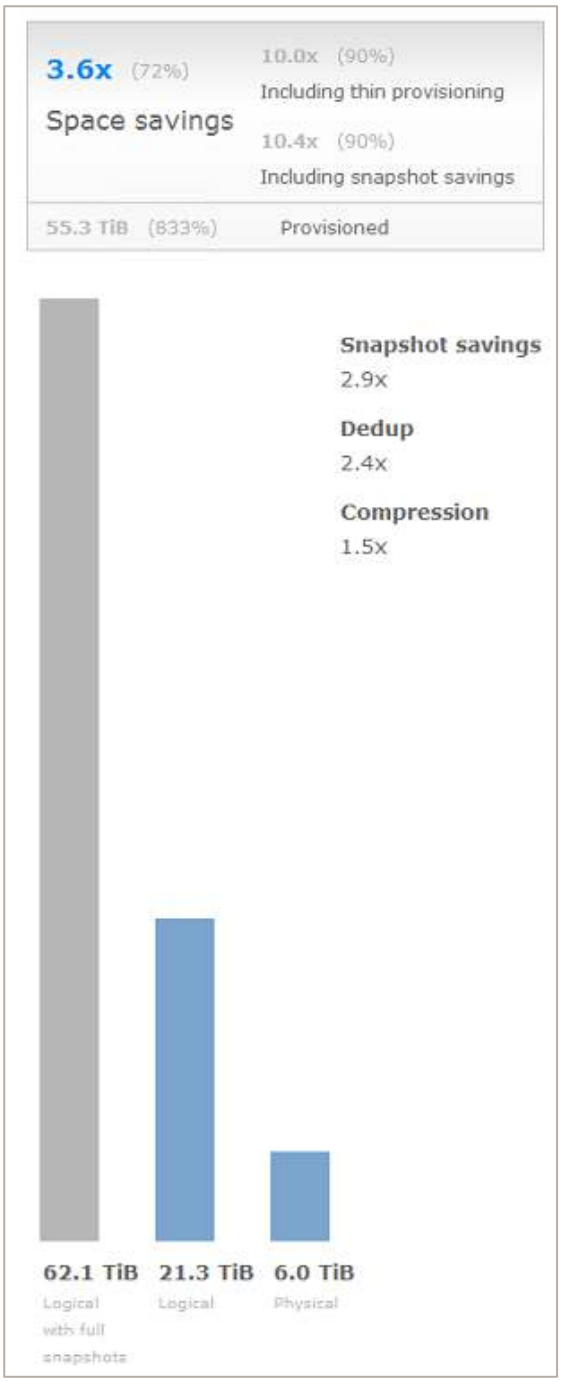

Figure 38: (Left)

A detailed presentation of the Space Savings information for the hosted datastores. This breakout window can be used to examine Thin Provisioning details as well as Snapshot savings

The logical footprint of the datastore is shown in the logical bar in the space breakout window. Physical stored data is shown in the Physical space gauge in the main dashboard, as well as in the physical bar in the space breakout window. There are two space savings number displayed within the breakout window – one prominent which details the combined savings realized through inline deduplication and compression of stored data, and another less prominent, which combined the standard space savings factor with the thin provisioning over-allocation factor.

Of key importance on the space savings breakout window is the ability to see what virtual disk provisioning methods for the VMs stored on the array. When a user hovers over the physical capacity usage bar, an informational pop-up will appear which displays what volume and percentage of overall capacity on the array has been allocated to Thick provisioned virtual machines. As mentioned already, these thick vDisks will not enjoy the data-aware space savings possible with Tintri VMstore intelligent storage.

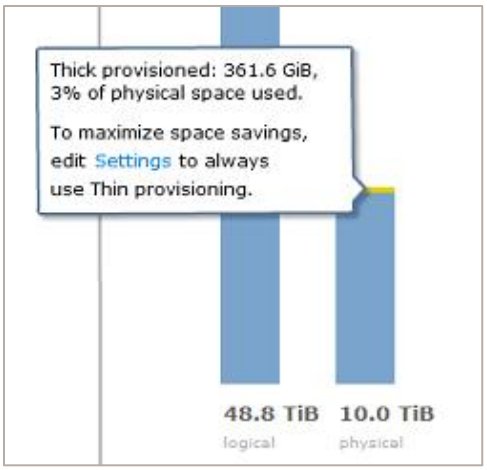

Figure 37: (Top) For VMstore systems with Thick Provisioned vDisks (and no Maximize Space Savings setting enabled), the Thick capacity consumption can be seen in the Physical capacity column

For Thick vDisks, Tintri recommends enabling the 'Maximize Space Savings' option, or preferably converting these to Thin virtual Disks where possible. The conversion operation can be instigated via the VMstore interface by right clicking on a VM identified as Thick and choosing the 'Convert to thin provisioned" option – see Figure 39.

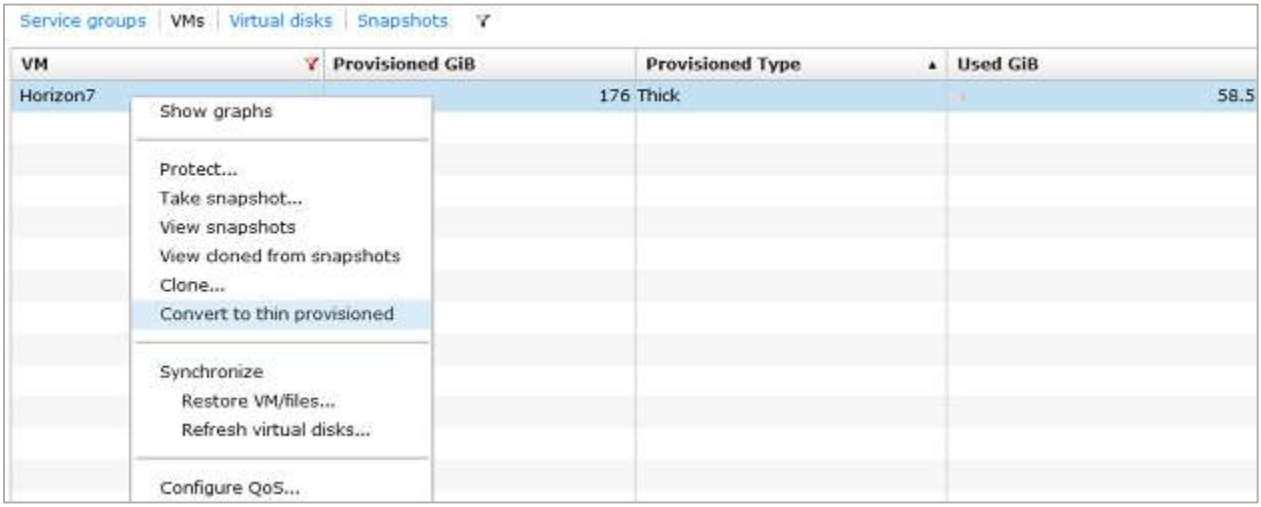

Figure 39: Converting a Thick provisioned VM to Thin can be done from the Tintri VMstore UI

C: vSphere Related Information

NFS.Heartbeat

The following settings should be applied to Tintri-connected hosts running VMware hypervisors which are older than ESXi version 5.0. For VMware ESXi version 5 and higher, these settings are the default and no alteration is required. If any of these settings are required to be altered, please ensure that all settings are changed in unison. If not, issues may occur.

Option Name	Value
NFS.HeartbeatFrequency	12
NFS.HeartbeatMaxFailures	10
NFS.HeartbeatTimeout	5

VM MAC Conflict alerts following Clone operation

When using various recent versions of vCenter Server, administrators may experience a vCenter generated alert warning of MAC address conflicts for recently cloned virtual machines. These alerts indicate that a duplicate MAC address has been detected amongst the group of managed VMs and should be resolved to avoid packet loss and other networking problems.

Status	Issue	Type
Alert	VM MAC Conflict	Triggered Alarm

Figure 40: VM MAC Conflict alert as reported by vCenter

These alerts can be triggered during a VM cloning operation which is performed using an external actor (i.e. not directly via vCenter) and occur early in the clone operation sequence when for a brief moment, a duplicate MAC address will exist between source and target VM – before ultimately being resolved prior to the clone becoming ready. Unfortunately, the VM MAC address conflict alarm remains triggered (or the alarm status will remain as red) even after the MAC address conflict issue has been resolved. This behavior is per VMware vCenter design.

As advised by VMware, the only recommended course of action is to either; manually acknowledge or clear the alarm after the issue is reported, or disable the alarm entirely. Refer to the 'Networking Issues' section in the VMware vCenter Server version 6.0 U1b release notes for further information.

Due to the nature of the error, it is always advisable to confirm if the alert condition is genuine or not, but for those that wish to fully disable the alarm in their environment, this action can be done via the 'Monitor' tab after choosing an appropriate vCenter object in the vCenter web client interface. To disable the alarm, filter the 'Alarm Definitions' for "MAC" and then edit the appropriate alarm definition (refer to Figure 41).

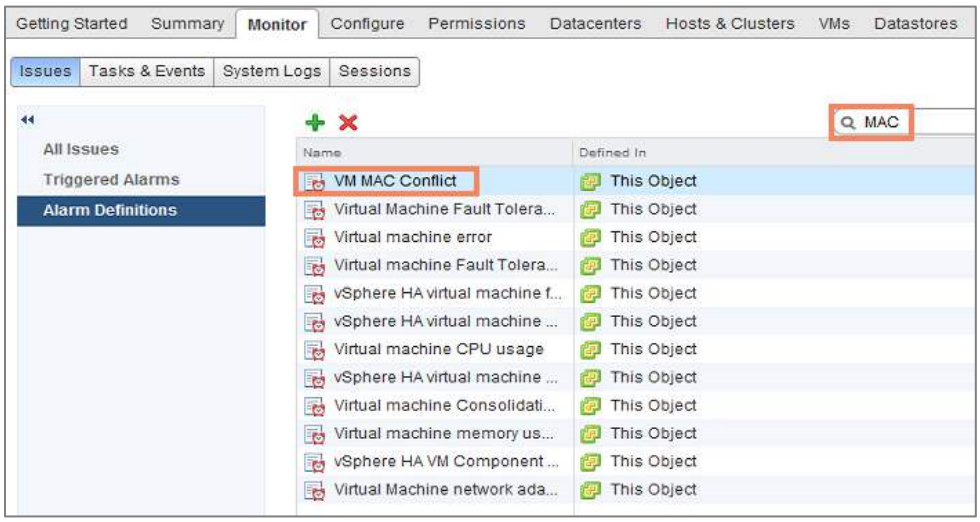

Figure 41: Filtering vCenter Alert Definitions to alter or disable the alarm behavior

D: Jumbo Frame Validation

To assist in validating that Jumbo Frames are configured correctly end-to-end, Tintri have added a diagnosis utility to the Tintri OS (since version 4.2). This utility is shown in Figure 42. While this is useful for analysis, 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 thoroughly as possible to ensure all paths allow jumbo frames end-to-end.

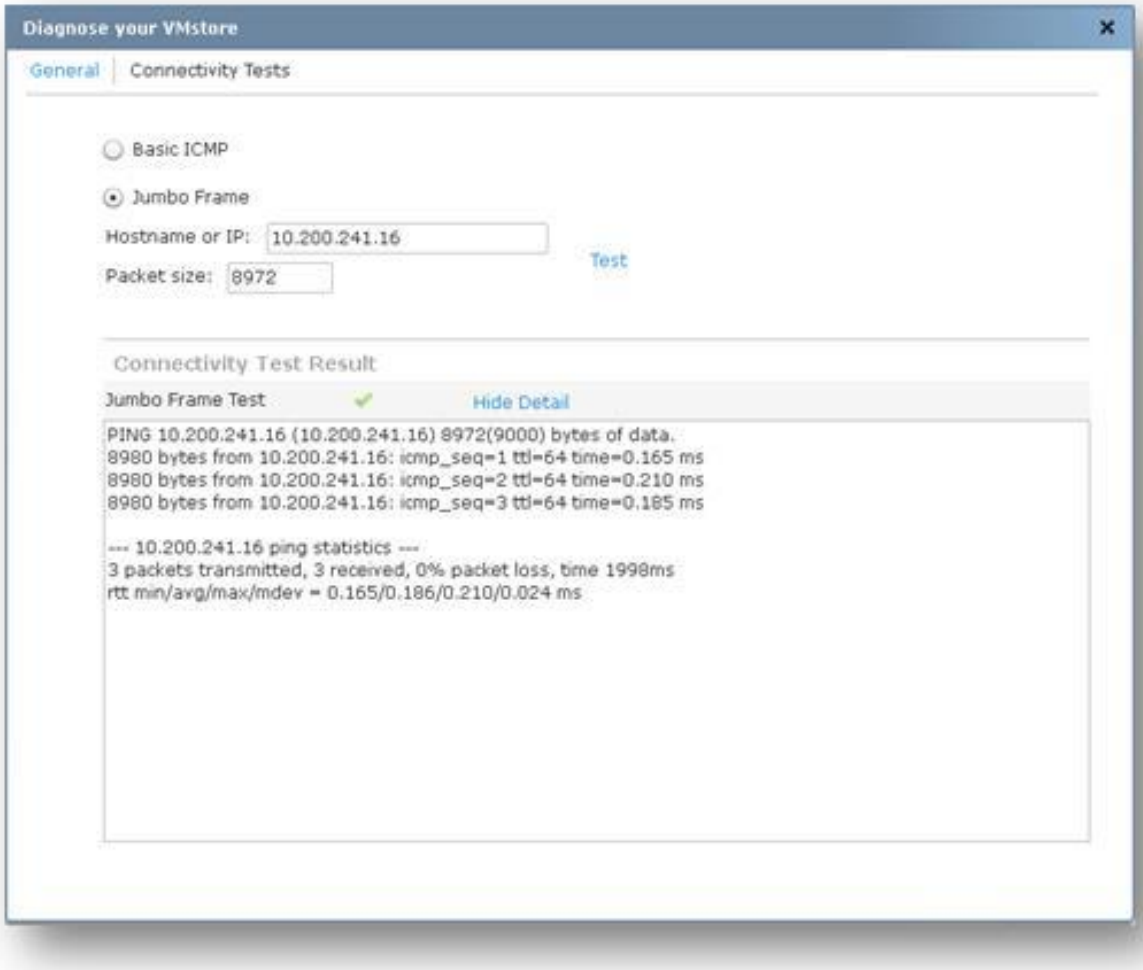

Figure 42: Network Diagnostics tool that is available in VMstore version 4.2 and later

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 43.

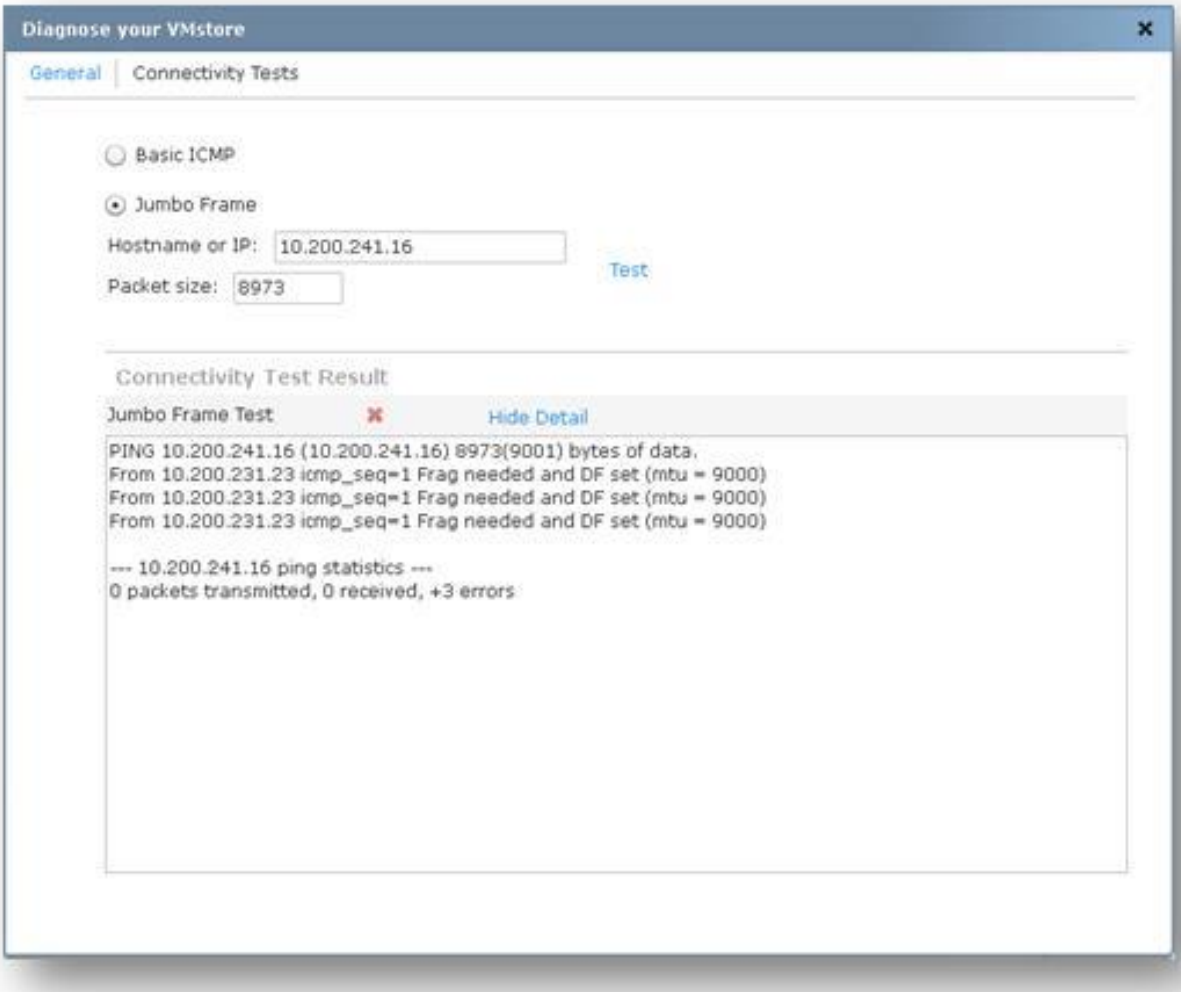

Figure 43: Testing Jumbo Frames - "Do Not Fragment" test failed

Tintri, Tintri VMstore, the Tintri Logo and FlashFirst are trademarks or registered trademarks of Tintri, Inc. All other trademarks or service marks are the property of their respective holders and are hereby acknowledged.

© 2019 Tintri by DDN. All rights reserved.

2929 Patrick Henry Drive
Santa Clara, CA 95054
info@tintri.com