WHITE PAPER

VMware® iSCSI on TrueNAS®

 | www.truenas.com | marketing@ixsystems.com

CONTENTS

1 Executive Summary

This guide is designed to provide VMware administrators and TrueNAS storage array administrators the knowledge necessary to configure the VMware and TrueNAS software with iSCSI hardware typically involved in an enterprise IT environment.

We will provide background information regarding iSCSI, IQNs, and Portals which can be found in chapter two. Chapter three provides a step-by-step guide for setting up vSphere and TrueNAS.

2 iSCSI Overview

iSCSI is an acronym that stands for "Internet Small Computer Systems Interface", and it represents **a set of standards** for using Internet-based protocols for linking aggregations of binary data storage devices. The draft standards were submitted in March 2000 and iSCSI has since seen a tremendous growth in adoption into enterprise IT environments.

One perspective of iSCSI can be understood as the "encapsulation" of SCSI commands and storage data within the "session" stack, which is then further encapsulated within the "transport" stack, which, yet again, is further encapsulated within the "network" stack, and so on and so forth. Transmitting data this way allows for block-level access to storage devices over LANs, WANs, and even "the Internet" itself (although don't expect performance to be amazing if your data traffic is traversing the Internet).

The table below shows where iSCSI sits in the OSI network stack:

OSI Layer Number	OSI Layer Name	Activity as it relates to iSCSI
7	Application	An application tells the CPU that it needs to write data to non-volatile storage.
6	Presentation	A SCSI Command, and/or SCSI Response, and/or SCSI data payload is created to hold the application data and communicate it to non-volatile storage.
5	Session	Communication between the source and the destination devices begin. This communication establishes when the conversation begins, what we will talk about, and when the conversation ends. This entire dialoge represents the "session". The SCSI Command, or SCSI Response, or SCSI data payload containing the application data is encapsulated within an iSCSI Protocol Data Unit (PDU).
4	Transport	The iSCSI PDU is encapsulated within a TCP segment.
3	Network	The TCP segment is encapsulated within a IP packet.
2	Data	The IP packet is encapsulated within the Ethernet frame.
1	Physical	The Ethernet frame is transmitted as bits (zero's and one's).

2.1 iSCSI Initiators and Targets

iSCSI introduces the concept of “initiators” and “targets” which act as sources and destinations respectively. iSCSI initiators and targets follow a client / server model. Below is a diagram of a typical iSCSI network. The TrueNAS storage array acts as the iSCSI target and can be accessed by many of the different iSCSI initiator types, including software and hardware-accelerated initiators.

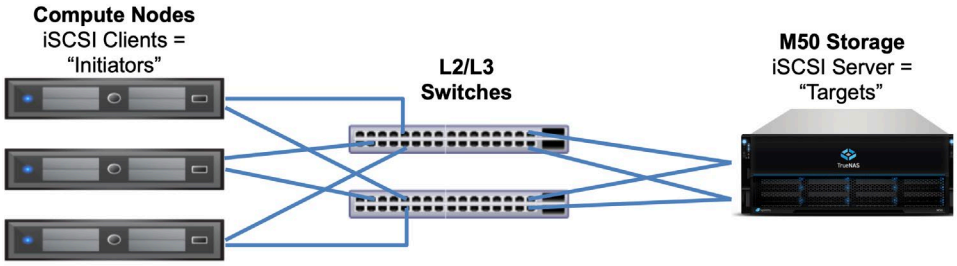

2.2 iSCSI Qualified Names

The iSCSI protocol standards require that iSCSI initiators and targets be represented as iSCSI nodes. It also requires that each node be given a unique iSCSI name. In order to represent these unique nodes via their names, iSCSI requires the use of one of two naming conventions and formats, IQN or EUI. iSCSI also allows the use of iSCSI aliases which are not required to be unique and can be helpful in managing nodes. Since VMware uses IQN in their software iSCSI implementation, that is what we will focus on below.

iSCSI IQN is an acronym that stands for “iSCSI Qualified Name”. It is comprised of the following naming schema with a preamble, node name, and unique identifier:

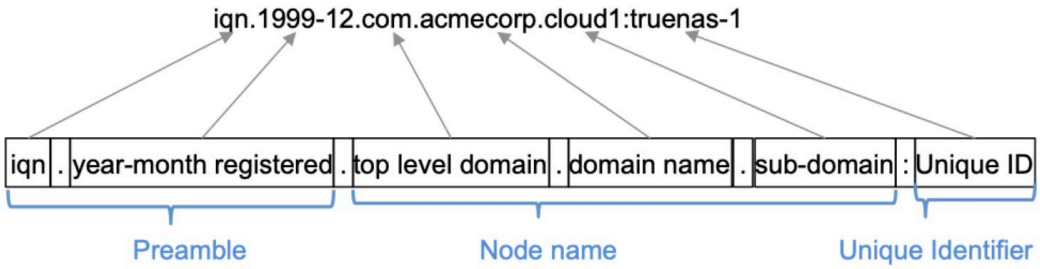

2.3 iSCSI Failover

A VMware datastore backed by iSCSI-based storage will consist of at least three distinct pieces: the storage host, the switching infrastructure, and the VMware host itself. In order to maximize service availability, each of these elements needs to be able to tolerate some level of failure without significantly disrupting iSCSI traffic.

TrueNAS systems support high-availability (HA) through dual-controllers running in active/standby mode. A properly-configured HA TrueNAS system can offer up to 5 × 9’s of system availability. TrueNAS also fully supports asymmetric logical unit access (ALUA) on iSCSI to significantly reduce failover time.

Network switching infrastructure can be made redundant and fault-tolerant through a number of methods, but multipathing is recommended as the best practice for iSCSI networks.

VMware's official documentation details several ways the virtualization host(s) can be made redundant, so we will not cover that here.

The sections below will discuss the proper network switching and storage host configuration required to implement a high-availability iSCSI network in more detail.

2.4 iSCSI Discovery, Authentication, and Access Control

For a VMware ESXi host to communicate with an iSCSI capable storage array, the iSCSI protocol must be configured to provide: **D**iscovery, **A**uthentication, and **A**ccess **C**ontrol **(DAAC)**. Let's take a closer look at these three DAAC categories.

2.4.1 Discovery

iSCSI offers two methods of target discovery: dynamic and static. Dynamic discovery lets the storage array respond automatically to the host initiator's "SendTargets" request. Static discovery requires an administrator to manually add a list of the iSCSI targets to the initiator. Either method of discovery is fine, but dynamic discovery can make the iSCSI setup process easier.

2.4.2 Authentication

iSCSI authentication is handled via the Challenge Handshake Authentication Protocol, or CHAP. CHAP uses a shared secret between targets and initiators to let them validate each other's authenticity. By default, no CHAP-based authentication is performed by the VMware iSCSI initiator. If you do decide to use CHAP, authentication can either be unidirectional (where only the target authenticates the initiator) or bidirectional (where both the iSCSI initiator and the iSCSI target are required to authenticate to each other prior to transmitting iSCSI data).

VMware iSCSI initiators operating with unidirectional CHAP can be configured in two behavior modes. In "Required" mode, an iSCSI adapter will give precedence to non-CHAP connections, but if the iSCSI target requires it, the connection will use CHAP instead. Required mode is only supported by Software iSCSI and Dependent Hardware iSCSI adapters. Alternatively, initiators can run in "Prohibited" mode, where an iSCSI adapter will give precedence to CHAP connections, but if the iSCSI target does not support CHAP, the initiator can still connect.

Bidirectional CHAP (called "mutual CHAP" in TrueNAS) offers greater security by ensuring that both sides of the iSCSI connection authenticate against each other. Unidirectional CHAP does not let the iSCSI initiator authenticate the target, and running without CHAP obviously disables all authentication. For this reason, bidirectional CHAP is usually recommended but requires additional configuration and comes with greater administrative overhead when troubleshooting iSCSI connections.

2.4.3 Access Control

Access control policies are set up within a storage array to ensure only certain initiators can connect to the target (even if they possess the correct CHAP password). Access control can be performed using the initiator's name (IQN), its IP address, or its CHAP username.

2.5 iSCSI Jumbo Frames

"Jumbo Frames" are the name given to Ethernet frames that exceed the default 1500 byte size. This parameter is typically referenced by the nomenclature of maximum transmission unit, or MTU. MTU that exceeds the default 1500 bytes necessitates that all devices transmitting Ethernet frames between the source and destination support the specific Jumbo Frame MTU setting, which means that NICs, dependent hardware iSCSI, independent hardware iSCSI cards, ingress and egress Ethernet switch ports, and the NICs of the storage array must all support the same Jumbo Frame MTU value. So, how does one decide if they should use Jumbo Frames?

Administrative time is consumed configuring Jumbo Frames and troubleshooting if/when things go sideways. Some network switches might also have ASICs optimized for processing MTU 1500 frames while others might be optimized for larger frames. Systems administrators should also account for the impact on host CPU utilization. Although Jumbo Frames are designed to increase data throughput, it may measurably increase latency (as is the case with some un-optimized switch ASICs); latency is typically more important than throughput in a VMware environment. Some iSCSI applications might see a net benefit running Jumbo Frames despite possible increased latency. Systems administrators should test Jumbo Frames on their workload with lab infrastructure as much as possible before updating the MTU on their production network.

3 VMware and TrueNAS iSCSI Setup

The setup of VMware iSCSI to TrueNAS requires that ESXi hosts be set up as initiators and TrueNAS storage arrays are set up as targets. This section covers proper configuration of iSCSI on VMware and TrueNAS.

3.1 Configuration of VMware ESXi iSCSI

3.1.1 Enable Software iSCSI

In vCenter, select an ESXi host, click, "Configure", then select "Storage Adapters" under the "Storage" section.

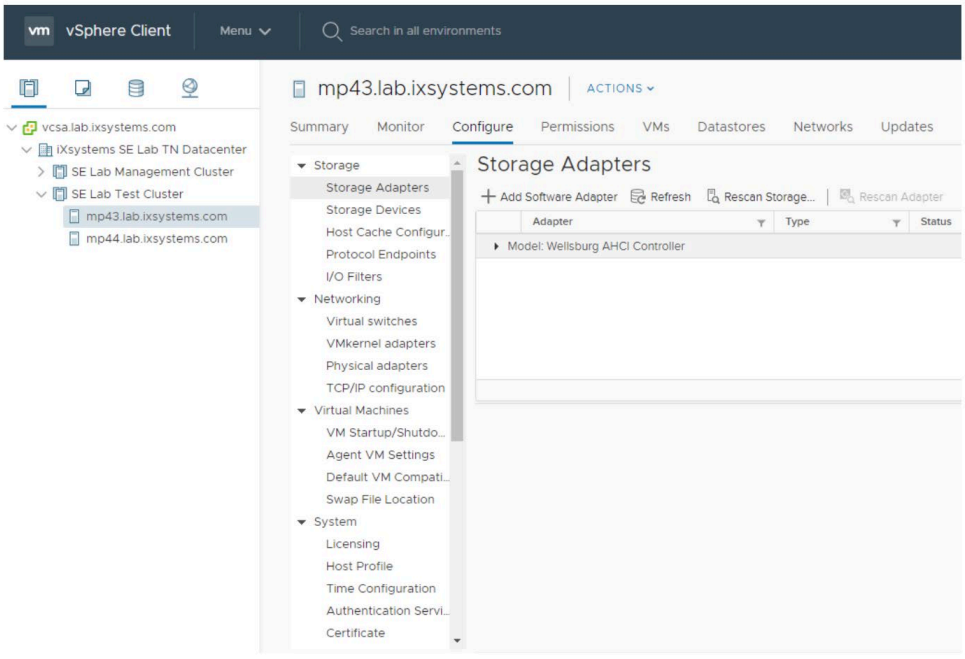

Next, click the “Add Software Adapter” button and select “Add software iSCSI adapter”. Click “OK”. After the software iSCSI adapter has been added you should be able to see your IQN under the “Identifier” column. You will want to record that entire character string as we will be using it when we configure iSCSI on TrueNAS.

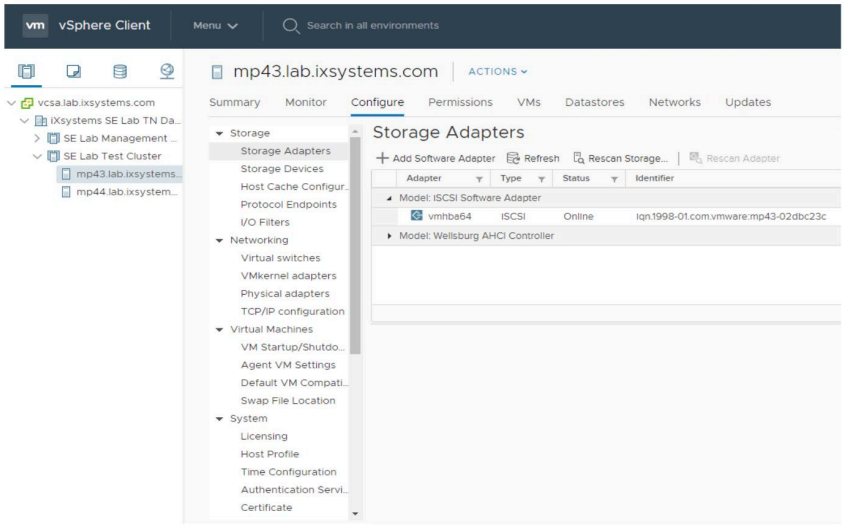

3.1.2 Configure Network Adapters, vSwitches, and vPortGroups

There are many ways to configure virtual networking within a VMware environment, each with their own pros and cons. Instead of covering each method, we will simply suggest a configuration that aligns to VMware and iXsystems best practices.

The image below is an example of how to configure iSCSI Network Adapters, vSwitches, and vPortGroups via vCenter:

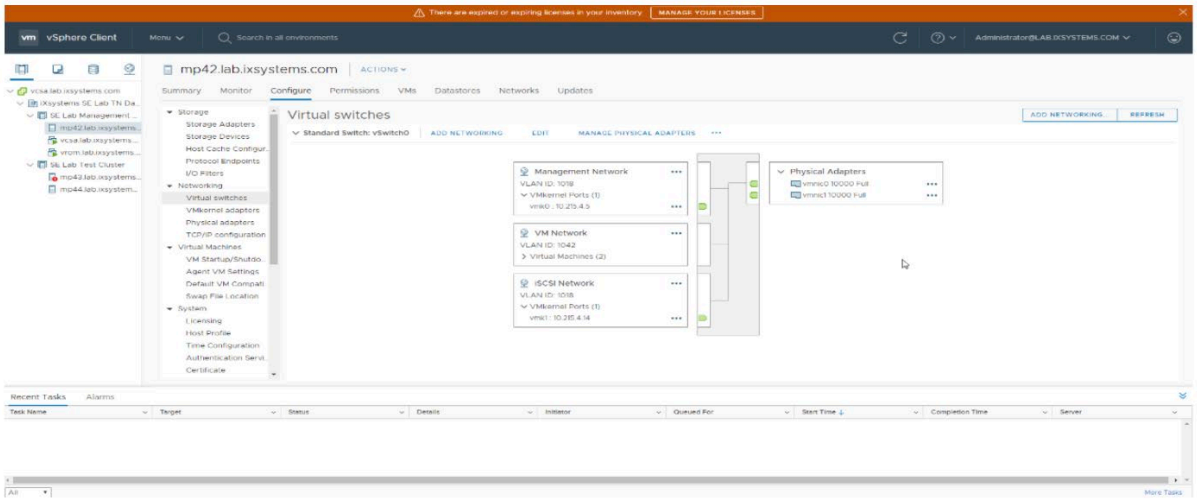

For additional information, please consult the vCenter User Guides.

3.1.3 Configure Port Binding and Multi-Pathing

Port binding configuration is wholly dependent upon what kind of configurations are applied to your network adapters, vSwitches, vPortGroups, and multipathing policies. The image below contains a specific port binding configuration but it only applies if your environment is following similar multipathing policies and virtual network configurations.

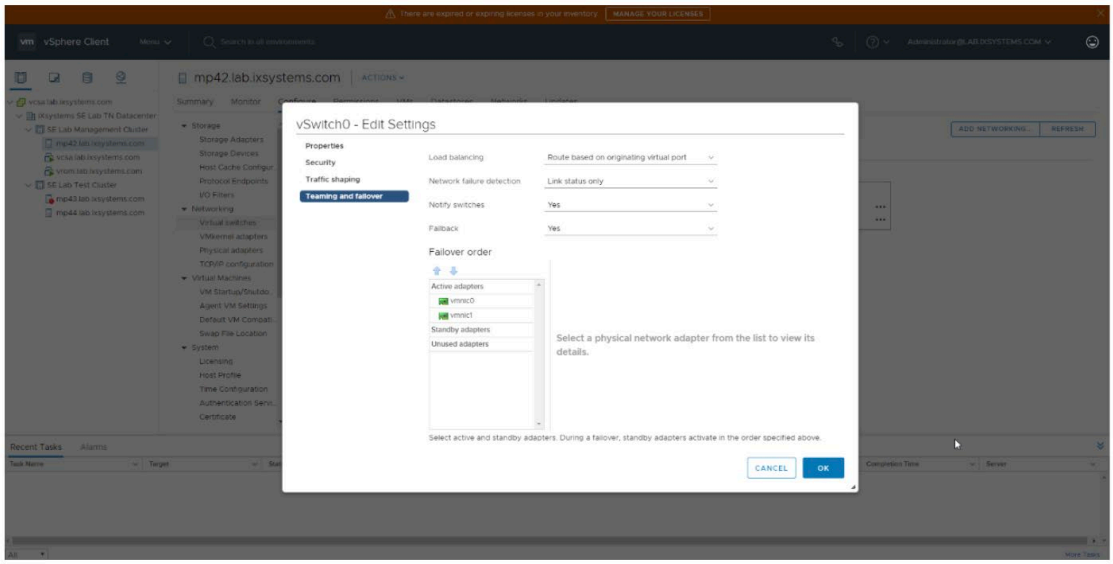

iXsystems recommends using the round robin multipathing policy (shown below).

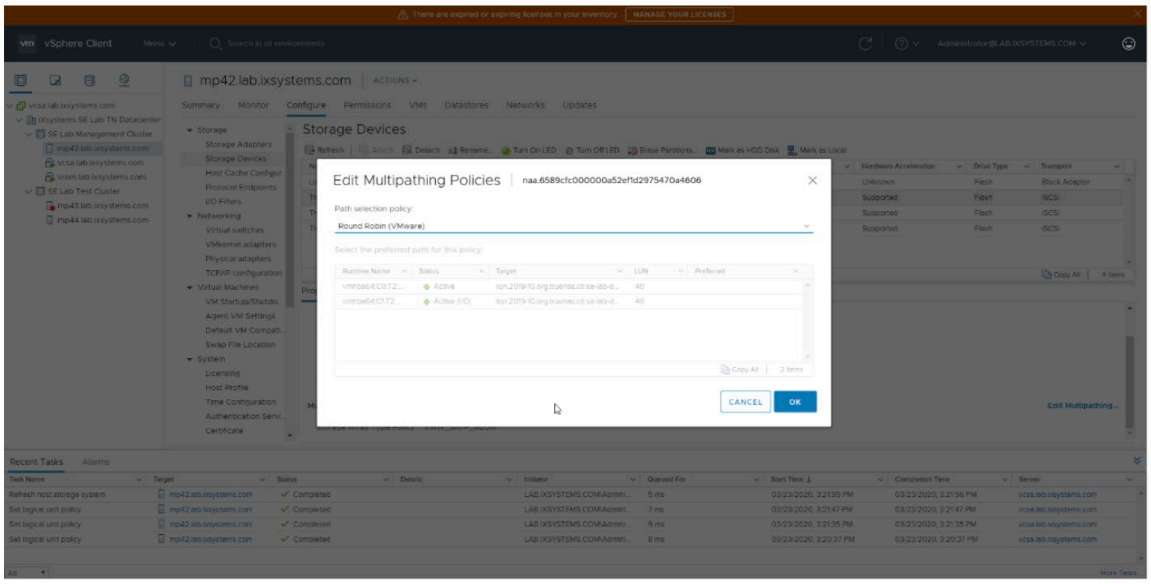

3.2 iSCSI Configuration on TrueNAS

3.2.1 Initial TrueNAS Setup

Starting with TrueNAS v11.3, the TrueNAS management GUI provides a wizard to assist users in iSCSI configuration. There is a 14 step sequence to enable TrueNAS as a vSphere datastore as shown in the table below.

Stage	Step	Description
Enable iSCSI	1a	Login to TrueNAS and go to Dashboard
	1b	Select Services from Dashboard View
	1c	Enable iSCSI and setup for automatic restart
	1d	Select Sharing and Block Shares (iSCSI)
	1e	Configure iSCSI Target with IQN base name and ALUA
Configure Storage	2a	Select Storage section
	2b	Add a new pool
	2c	Add a ZVOL dataset for iSCSI
	2d	Name the dataset and select sync=always
Create and Share iSCSI LUNs	3a	Select Sharing→Block and start iSCSI Wizard
	3b	Name Datastore and share to ESX Hosts
	3c	Setup iSCSI Portal and Discovery Authentication Method
	3d	Authorize the ESXi initiators
	3e	Confirm the iSCSI target configuration

Let's walk through each step to best understand what is required.

Step 1a: Log in to the browser management interface of TrueNAS using administrative credentials.

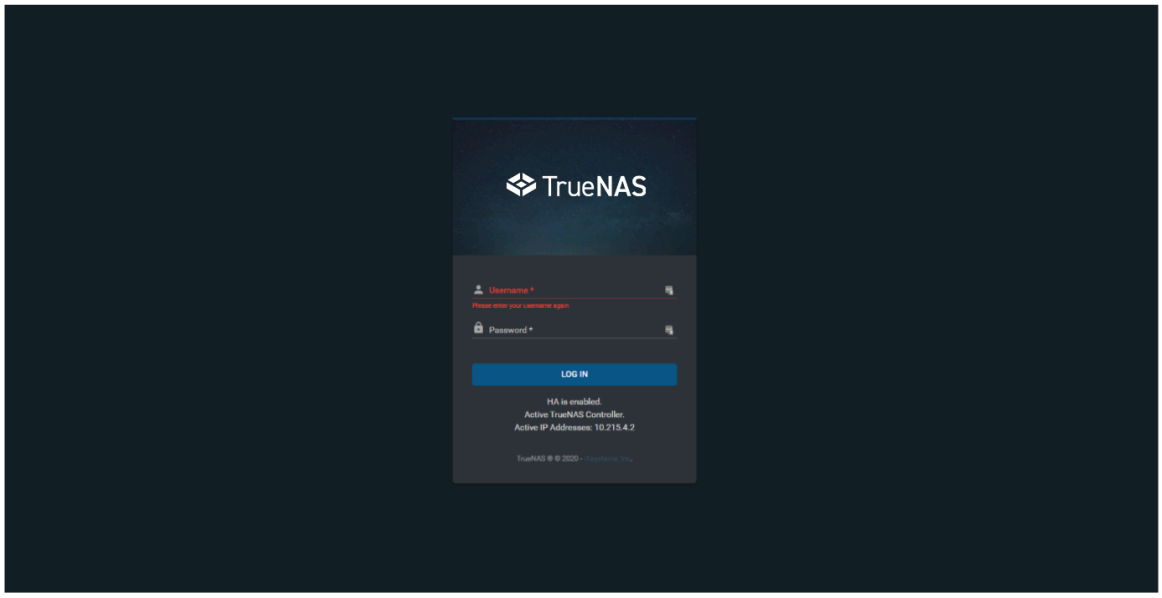

Step 1b: On the left-hand side of the “Dashboard” view, there is a column of options. If this column is hidden, select the hamburger button in the top left. Scroll down and select the option titled, “Services”.

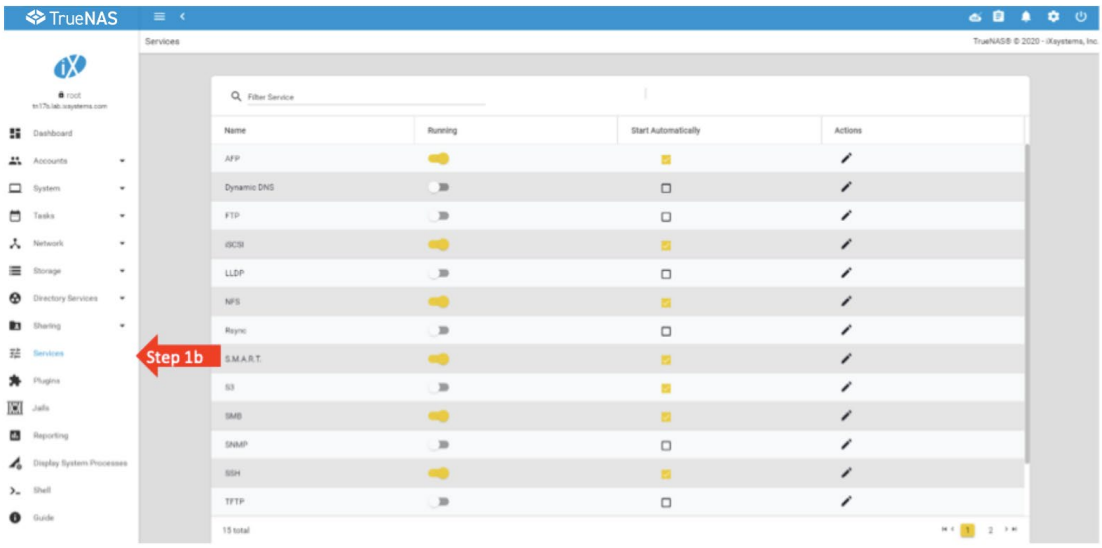

Step 1c: From the list of services, find the service titled, “iSCSI”, and toggle the switch to “Running”. Select the “Start Automatically” check box.

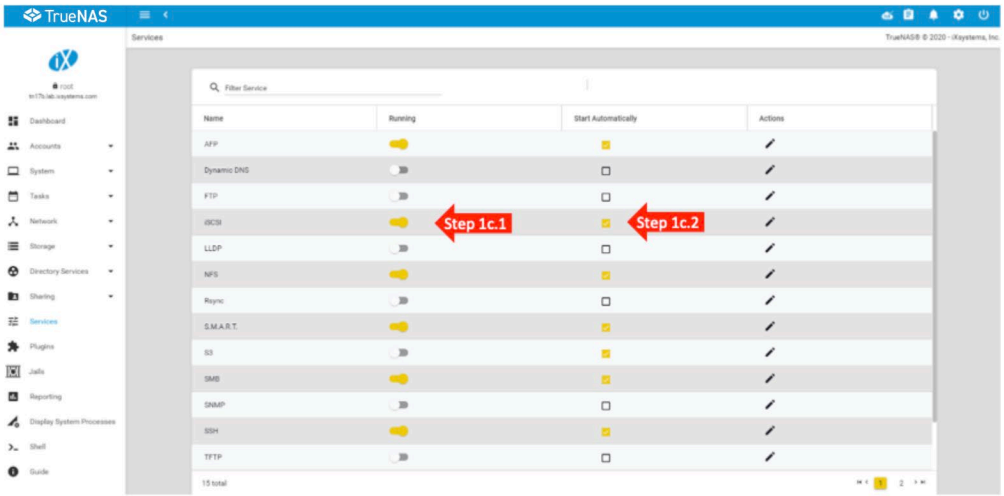

Step 1d: Expand the option titled “Sharing” on the left menu bar then select “Block Shares (iSCSI)”.

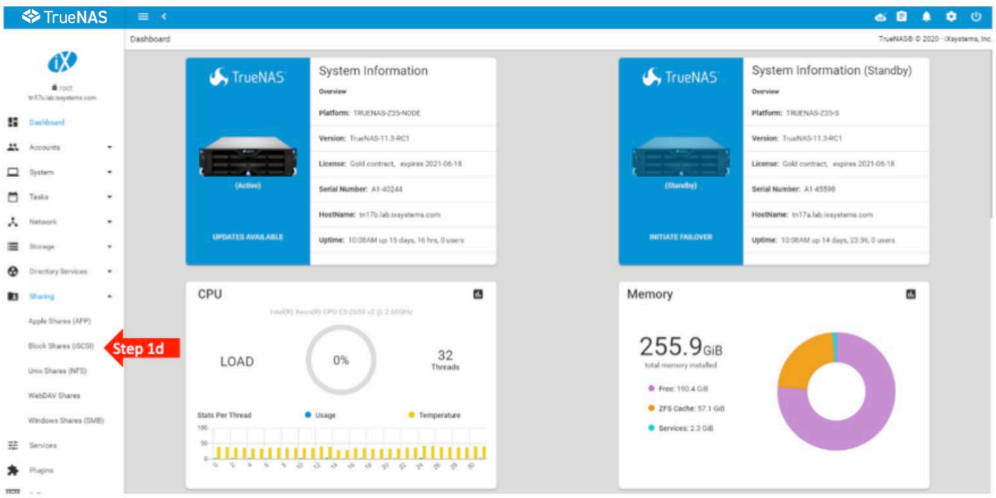

Step 1e: After selecting the “Block Shares (iSCSI)” option, you should be able to populate the “Target Global Configuration” tab fields, “Base Name”, (which must follow IQN naming convention standards). Select the checkbox beside “Enable iSCSI ALUA” only if you have a dual-controller HA TrueNAS system. Finally, click “Save”.

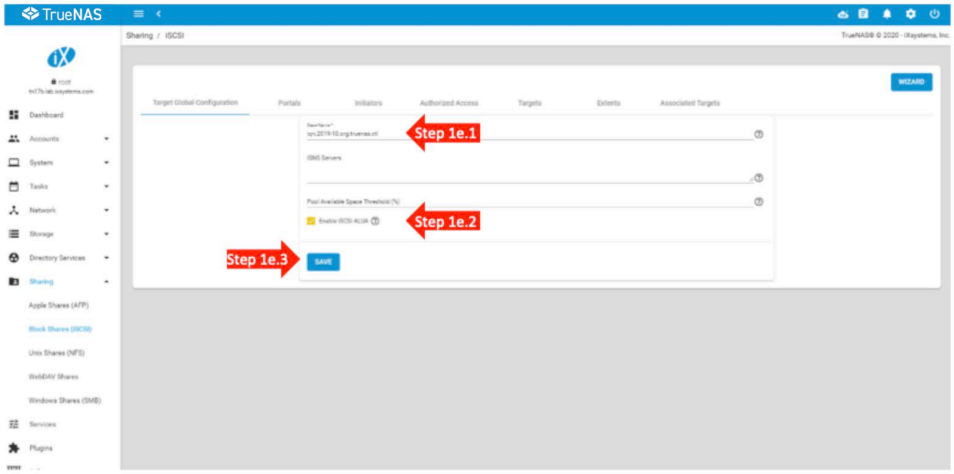

3.2.2 Create ZFS Pools and Datasets

TrueNAS HA systems are typically delivered with preconfigured ZFS pools. In this case, steps 2 a,b,c,d are unnecessary. Proceed to step 3a and createing the iSCSI LUNs.

Step 2a: After clicking “Save” in Step 1e, expand the “Storage” section on the left-hand column. Then select the “Pools” section under Storage. Next, select the “Add” button in the upper right-hand corner. It should be located on your screen in a similar place to the image below:

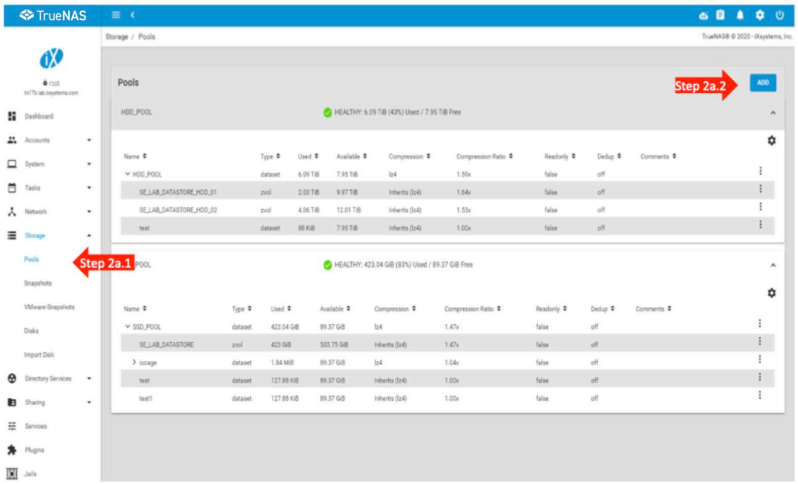

Step 2b: After clicking the “Add” button in Step 2a, click the “Create Pool” button. It should be located on your screen in a similar place to the image below:

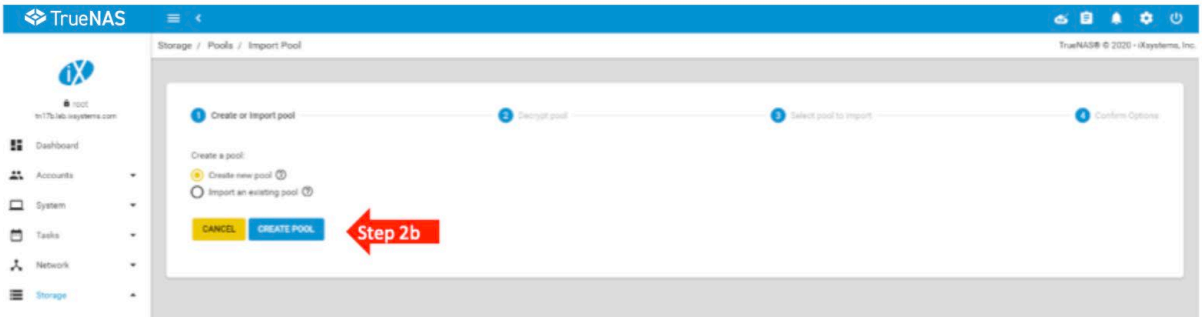

The image below shows the Pool Manager which will look unique to your environment based on the quantity and types of drives (both SSDs and HDDs) within your TrueNAS storage array.

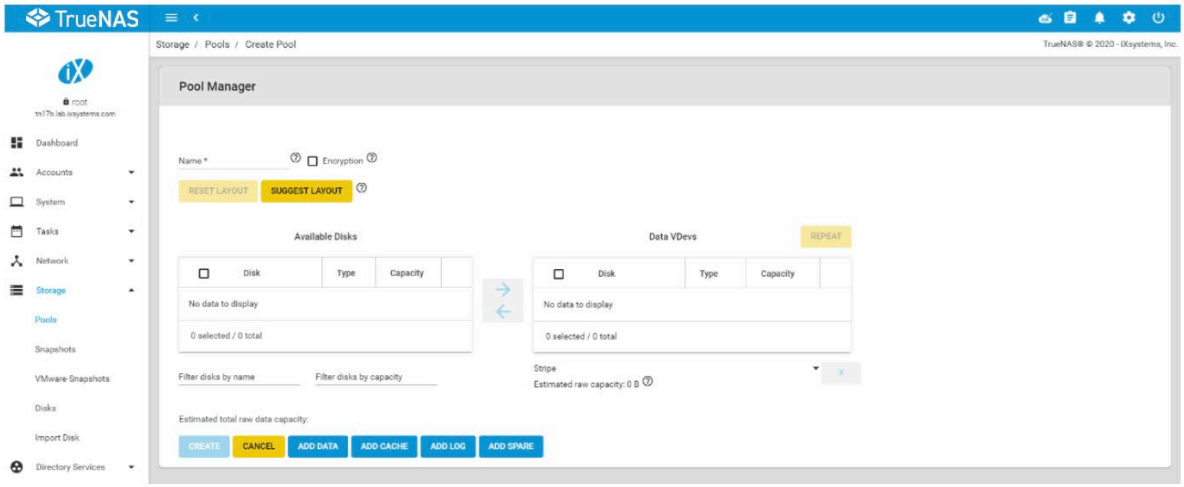

Step 2c: On the same line as the name of the ZFS Pool, click on the ellipsis on the right-hand side and select the "Add Zvol" option.

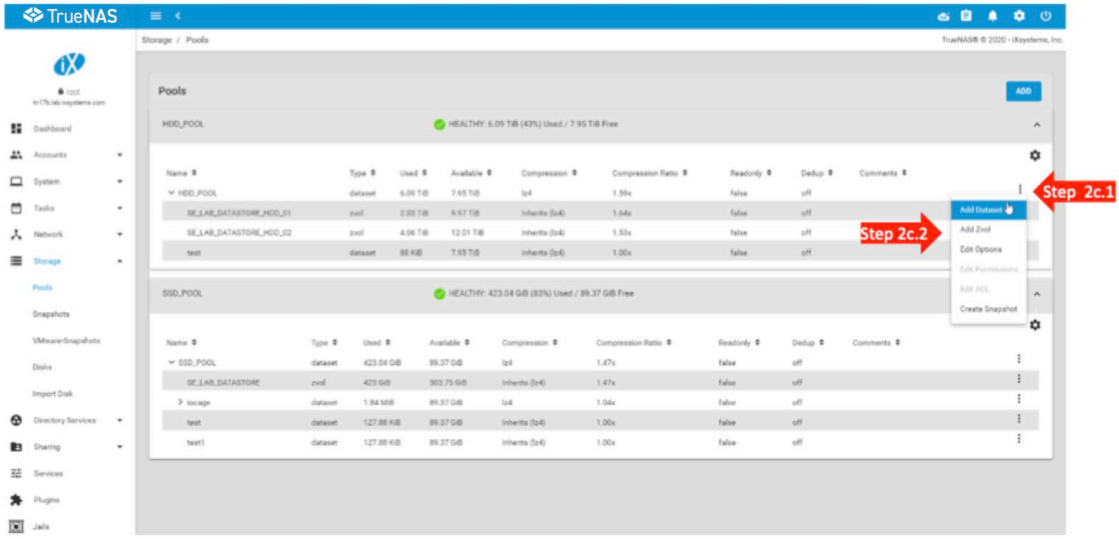

Step 2d: Create the Name of the Dataset and select the Dataset Options. For virtualization workloads, it is important to set "sync=always" to eliminate any chances for data corruption or loss. For typical VMware use-cases, a block size of 16 KB is recommended.

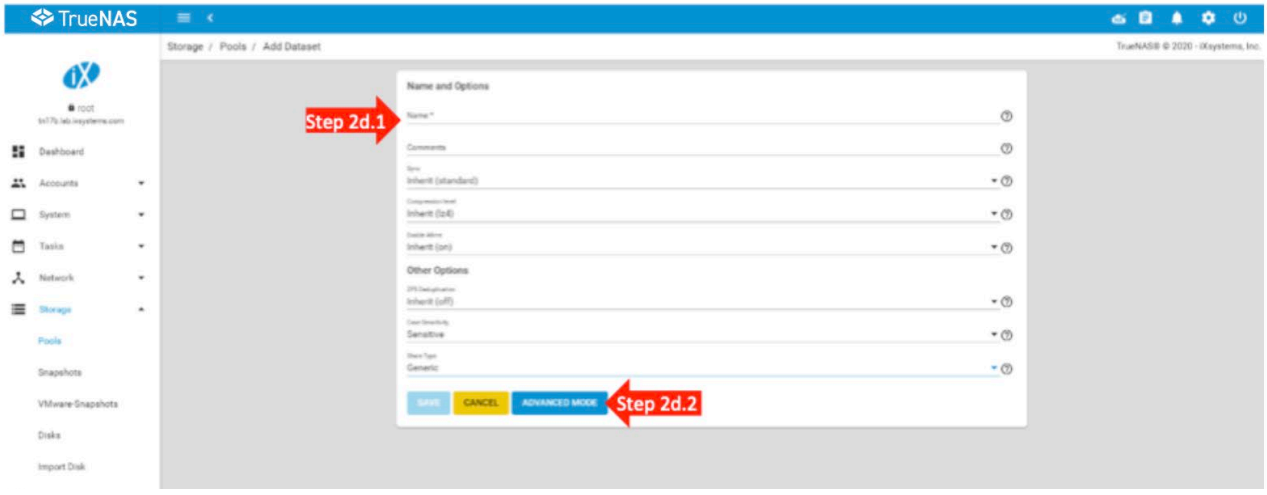

3.2.3 Create and Share iSCSI LUNs

After creating the ZFS Pool and Zvols according to your VMware needs, we can proceed with initiating the iSCSI Wizard.

Step 3a: In the left-hand column, expand “Sharing”, then select “Block Shares (iSCSI)”. Click on the “Wizard” button on the upper right-hand side of the screen.

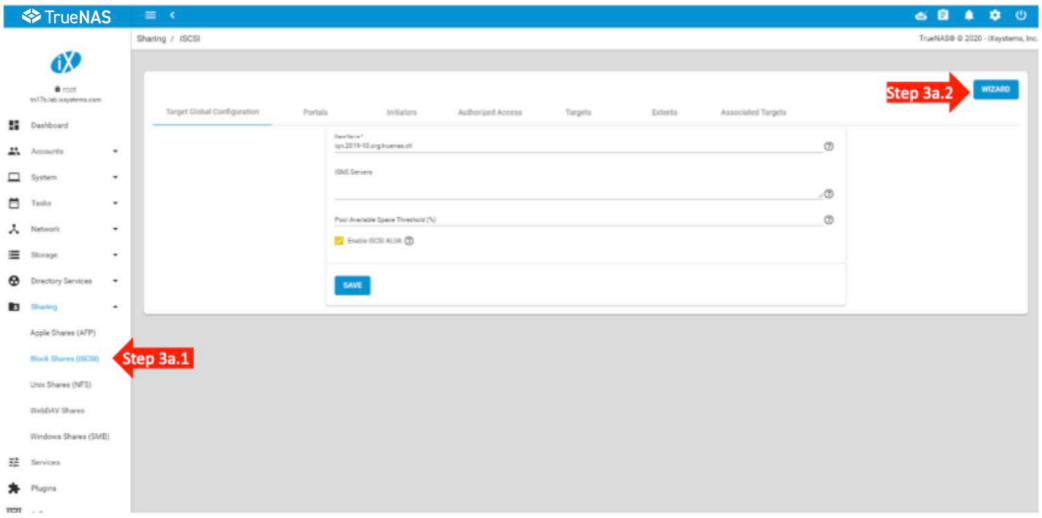

Step 3b: Give the Datastore a unique Name. Select “Device” under the Type section. Select the Pool/Zvol that you already created. You can also create a new Zvol from this screen. Be sure to select the “VMware…” option under the “What are you using this for” section. Click “Next” when you are ready to proceed.

Step 3c: From the Drop-down arrow on the far-right of the Portal line item, select “Create New” (or select an existing Portal). Start by defining the Discovery Authentication Method (CHAP, Mutual CHAP, or None)., If you select CHAP or Mutual CHAP, you will need to provide the CHAP Group ID, the user name, and the shared secret key (between 12 and 16 characters). If you have an existing Discovery Authentication Group defined, you can re-use it. Finish by defining the target’s IP Address, and IP Port (this should default to 3260). If you want a single Portal to accept connections to multiple interfaces, click “Add”. When you are done adding interfaces to the portal, click “Next”.

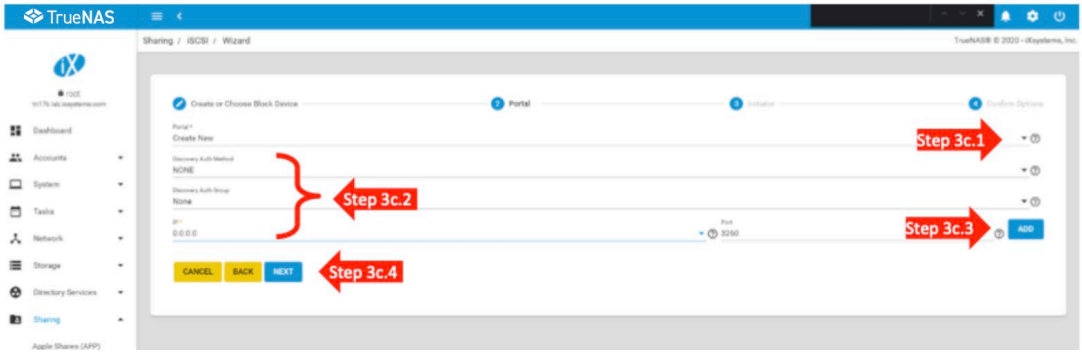

Step 3d: Populate the VMware Initiator names (IQNs) for each of the Initiators that you want to connect to this Portal (separated by a space). Note that the IQN must follow the iSCSI protocol IQN format; the image below contains an example. If these IQNs are left blank, all initiators on the authorized network will be accepted. This greatly simplifies cluster expansion but could potentially introduce a security risk. Add any “Authorized Networks” if needed. Click “Next” when you are finished.

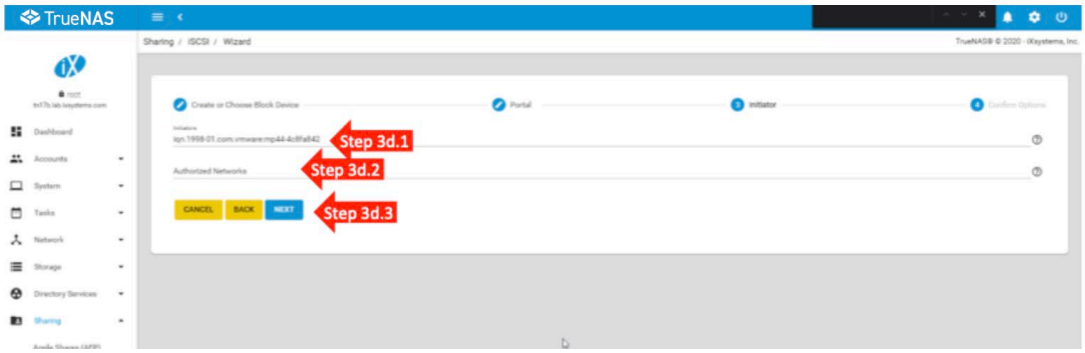

Step 3e: Review the information for accuracy and select “Submit”. If any configuration errors are identified, you will be notified and have a chance to fix them.

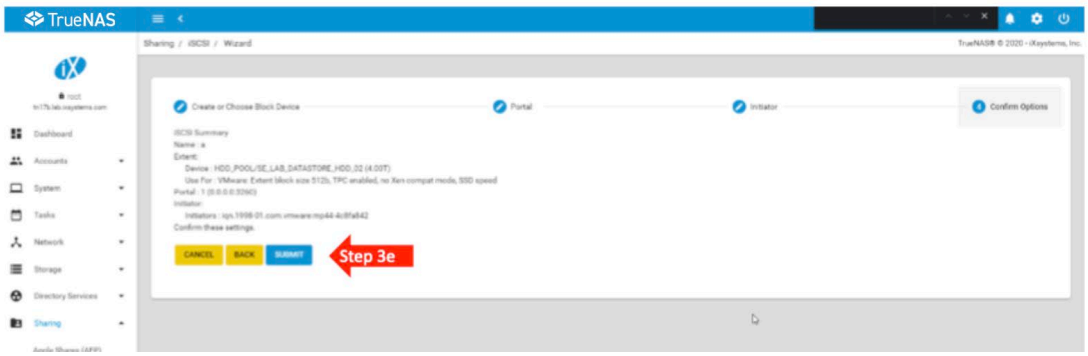

You should now have iSCSI setup on VMware and TrueNAS. Storage IO should be traversing your network via iSCSI.

4 Appendix

4.1 Networking Diagram Overview

The following picture is the recommended wiring configuration on a per-server basis for full network switch, compute host, storage array BMC/IPMI, IP network management, production network traffic, vMotion, svMotion, provisioning, and iSCSI data redundancy.

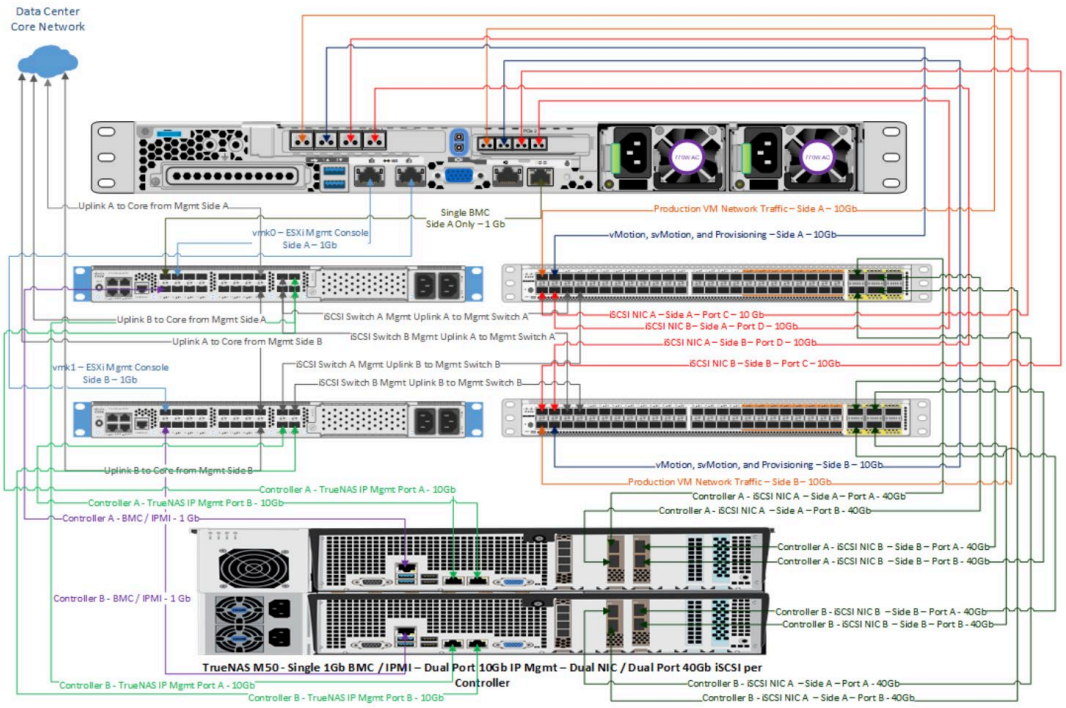

4.2 iXsystems Networking VMware VCenter Plugin for TrueNAS

To improve the automation of VMware and TrueNAS, iXsystems has developed a vCenter plugin for TrueNAS that automates the major tasks for operating a vSphere environment including:

- **Creation and assignment of LUNs**
- **Snapshotting LUNs**
- **Replicating LUNs**

The TrueNAS plugin for vCenter is introduced in **this video.**

Setup is further simplified by having the vCenter Plugin automate many of the configuration tasks. The plugin achieves this automation by using the TrueNAS REST APIs. The workflow for configuring TrueNAS with the vCenter plugin is simplified down to the following dozen steps.

Stage	Step	Description	Tool
Install TrueNAS	1a	Install TrueNAS and Configure Networking	TrueNAS
	1b	Select Storage section and add a new Pool	TrueNAS
Enable iSCSI	2a	Configure iSCSI Target with IQN base name	TrueNAS
	2b	Configure ALUA for HA Operation	TrueNAS
	2c	Add a ZVOL dataset for iSCSI	TrueNAS
	2d	Name the dataset and select sync=always	TrueNAS
Setup vCenter	3a	Install vCenter	vCenter
	3b	Install TrueNAS Plugin for vCenter	TrueNAS vCenter Plugin Deployment Tool
	3c	Configure vCenter settings in the TrueNAS Plugin	Plugin/vCenter
	3d	Add TrueNAS and password to vCenter	Plugin/vCenter
Operate vCenter	4a	Install ESXi hosts and Set-up ISCSI adapter	vCenter
	4b	Operate vCenter - Setup LUNs	vCenter

Once these steps are completed, most additional steps are automated by the plugin.

Contact iXsystems for assistance in downloading and installing the vCenter Plugin for TrueNAS.