



Technical white paper

HP ProLiant SL2500 with Maxta Storage Platform

Hyper-Converged Deployment Guide

Updated: 10/31/2014

Marcus D'Andrea and Hongjun Ma, HP DIA

Table of contents

Introduction.....	4
Technology Overview	4
SL2500 Hardware Overview.....	5
Compute Node and Chassis Drive Cage.....	6
Storage Controller	6
Boot Device.....	7
Networking	7
SL2500 with Maxta Storage Platform Deployment Network Diagram	8
Maxta Storage Platform Overview.....	9
Understanding The Maxta Storage Platform Architecture	9
Maxta Storage Platform Prerequisites and Requirements	12
Vmware vSphere Requirements	12
Maxta Hardware Requirements	12
Storage Requirements.....	13
Hard Disk Drive Requirements	13
Network Requirements.....	13
Network Interface Cards.....	13
Supported Virtual Switch Types.....	14
HP ProLiant SL2500 Maxta Installation and Configuration.....	14
Configuring the SL2500 HP H220 HBA.....	14
Setting up the Vmware VSphere Environment for Maxta Storage Platform	17
Installing the Maxta Storage Platform.....	19
Maxta Administration.....	26
Applying the Maxta Storage Platform License.....	26
Add a Node to Maxta.....	26
Working with Maxta Clones and Snapshots	27
Creating Custom Pagesizes for VMDKs	31
Maxta Monitoring, and Performance Optimization	35
Monitoring and Troubleshooting	35
Maxta Statistics.....	35
Looking at Replicas and Sync State	37
Performance Optimization	38
SL2500 BIOS Power Optimization Settings.....	38
ESXI Performance Optimization Settings	38
Performance Testing	40
Conclusion	41
Appendix.....	42

Maxta CLI Installation.....42

Introduction

The purpose of this Deployment Guide is twofold: 1). Introduce the technical reader to Maxta Storage Platform using HP's SL2500 Servers and 2). Provide a proven and tested installation tutorial. The intended audience of this paper is the IT technical staff including the network, server and storage designers, architects, and operations teams. The reader of this document is expected to be fluent in virtualization technologies as well as moderate levels of server, storage and networking expertise. In addition, the reader is expected to have a working knowledge of Vmware ESXi release 5.5 administration.

Technology Overview

This section outlines the 2 components used in the solution overview:


- Maxta is a software-defined storage solution that provides organizations the choice to deploy hyper-convergence on any x86 server, using any hypervisor, and any combination of storage devices.
- The HP ProLiant SL2500 server line provides all the features expected in an enterprise server, with the added benefit of being optimized for efficiency, density and flexibility. The ProLiant SL2500 doubles the density, as compared to standard rack servers, to increase available data center floor space, improve performance while lowering energy consumption.

The paper will help HP customers deploy Maxta on HP ProLiant SL2500, along with highlighting the Maxta design and general concept and implementation options unique for HP ProLiant SL2500 platform.


SL2500 Hardware Overview

This section details the SL2500 Infrastructure offering.¹

SL2500 Front View (24 SFF Drive)



SL2500 Back View (4x1U Node)



The HP ProLiant SL2500 scalable system offers a very dense solution with up to four independent HP ProLiant SL210t G8 hot-pluggable server nodes in a standard 2U chassis. Each SL210t Gen8 Server node can be serviced individually without impacting the operation of other nodes sharing the same chassis to provide increased server uptime.

HP ProLiant SL2500 supports:

- Four SL210t-Gen8 server nodes in SL2500 2U Chassis
- 24 Small Form Factor (SFF) or 12 Large Form Factor (LFF) SAS, SATA or SSD drives with 6 SFF (or 3 LFF) drive bays dedicated to one SL210t-Gen8 server

HP ProLiant SL210t-Gen8 servers support:

- Up to two Intel Xeon process E5-2600 v2 CPUs
- Up to 512GB of RAM with 16 DIMM slots
- One FlexibleLOM slot and one low profile PCIe 3.0 expansion slot
- Two gigabit Ethernet and one management LAN interfaces on the motherboard

¹ This paper uses a 1U server node for maximum server density and the SFF drive cage with six dedicated SFF hard disk drives (HDD) or solid state drives (SSD) per server node

For detailed hardware specifications, please refer to [SL2500 QuickSpecs](#).

Each of the four SL210t-Gen8 nodes tested included the following components:

Component	Specification
Processor	Dual-Socket Intel Xeon CPU E5-2620 v2 @ 2.10GHz. 6 cores per socket with Hyper-Threading enabled
Memory	96 GB DDR3, 6x16GB RDIMM @ 1600MHz
Network	HP FlexFabric 10G 2-port 554FLR-SFP+ Adapter
Storage Controller	HP H220 SAS HBA
SSDs	2x100GB 6G SATA HP Enterprise Mainstream Endurance SSD (691862-B21)
HDDs	4x1.2TB 6G SAS 10K RPM HP SAS drives (697574-B21)

Note: The SSD drives should be replaced with 480GB SSD Drives for best practice Maxta sizing. (717971-B21). The 1.2TB 6G SAS 10K drives that the test was conducted on have been replaced with (718162-B21).

Major HP ProLiant SL2500 component details:

Compute Node and Chassis Drive Cage

The HP ProLiant SL210t Gen8 hot-pluggable server node is available in 1U and 2U options (1U used for this paper). The drive cage is available in small form factor (SFF) and large form factor (LFF) options. The HP ProLiant SL2500 chassis can hold up to four HP ProLiant SL210t Gen8 hot-pluggable 1U server nodes and up to 24 SFF disk drives. Each 1U server node can support up to six SFF disk drives and the drives may be SAS, SATA, SSD or a combination of these drive types.

This paper will focus on the 1U server node for maximum server density and the SFF drive cage with six dedicated SFF hard disk drives (HDD) or solid state drives (SSD) per server node.

Storage Controller

The base model storage controller for the HP ProLiant SL210t Gen8 server is configured with the HP B120i software RAID controller². However for our testing we used a very cost-effective and preferred storage controller – the HP H220 Host Bus Adapter.



Figure 1 HP H220 Host Bus Adapter

Boot Device

This Deployment Guide boots VMware ESXi from one 1.2TB SAS drive located in drive bay one. This drive is dedicated to booting VMware ESXi and the remainder of the drive is used for Maxta. All 1.2TB SAS drives will be used for the Maxta shared datastore and the SSD drives will be used for caching and metadata.

Networking

The HP ProLiant SL210 Gen8 server includes two integrated 1Gb Ethernet NICs for user data networking. VMware recommends 10 Gb networking for performance reasons for data storage networking. Therefore, we have added one Dual Port 10 Gb FCOE adapter in FlexibleLOM form factor to preserve our PCI slot for the HP H220 Host Bus Adapter (described above in the Storage Controller section)

The HP FlexFabric 554FLR-SFP+ was used in our testing but users have the options to choose other HP 10 Gb/FCOE adapters.

The following 10 Gb adapters are available and listed in the [SL2500 QuickSpecs](#).

FlexibleLOM Adapters

HP Ethernet 1Gb 4-port 331FLR Adapter	629135-B21
HP Ethernet 1Gb 4-port 366FLR Adapter	665240-B21
HP FlexFabric 10Gb 2-port 526FLR-SFP+ Adapter	629138-B21
HP FlexFabric 10Gb 2-port 534FLR-SFP+ Adapter	700751-B21
HP FlexFabric 10Gb 2-port 554FLR-SFP+ Adapter	629142-B21
HP Ethernet 10Gb 2-port 560FLR-SFP+ Adapter	665243-B21
HP Ethernet 10Gb 2-port 570FLR-SFP+ Adapter	717491-B21



Figure 2: HP FlexFabric 10G 2-port 554FLR-SFP+ Adapter

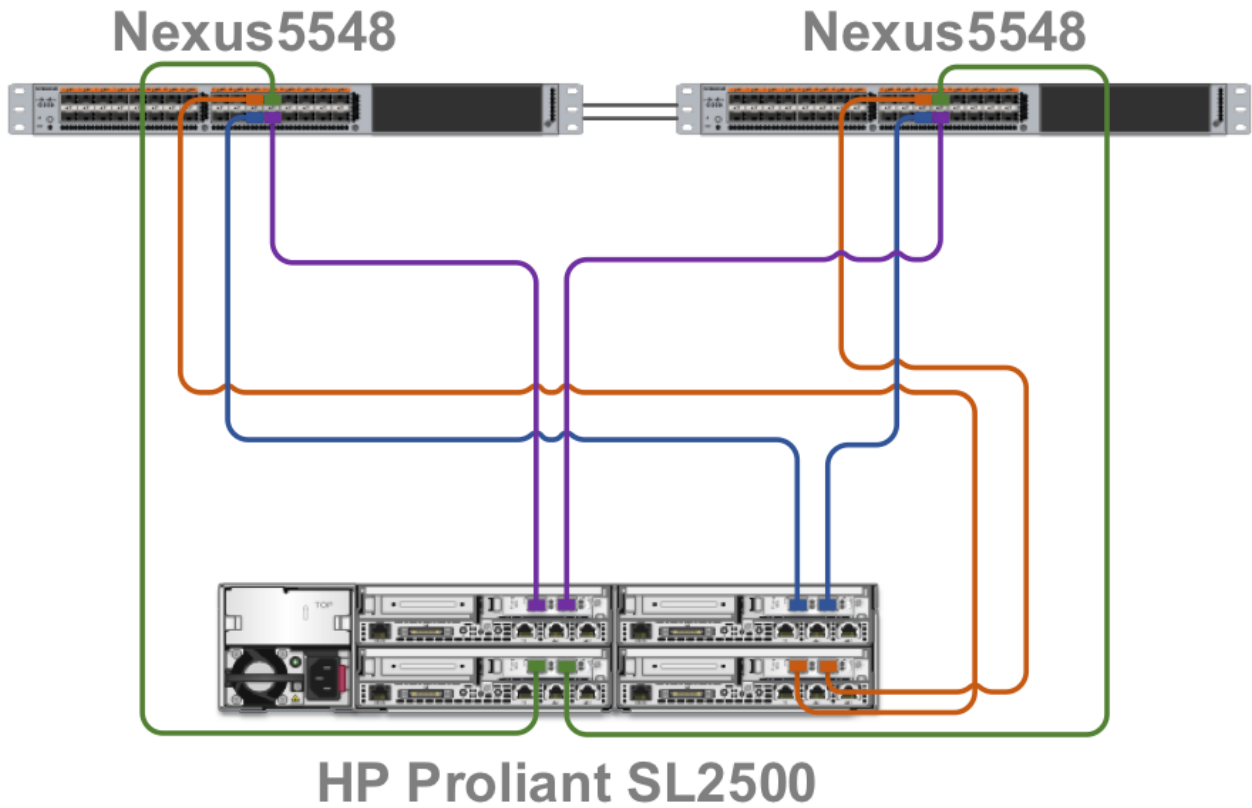
Note HP network adapter naming: The second digit in the model name indicates the manufacturer. Therefore, for our testing we used the Emulex based adapter.

- 52x: Qlogic
- 53x: Broadcom
- 55x: Emulex
- 56x: Intel
- 57x: SolarFare

SL2500 with Maxta Storage Platform Deployment Network Diagram

For this Deployment Guide, the HP ProLiant SL2500 is a strong hyper-converged platform for Maxta Storage Platform. This architecture provides a scale-out solution for compute, memory, and storage and eliminates complex SAN networking and centralized storage.

Notice the simplicity in network connectivity; each server node utilizes two 10 Gb links to the redundant upstream switches (Nexus 5548 in this example).



Maxta Storage Platform Overview

Maxta Storage Platform is a hypervisor agnostic, software-only implementation of enterprise storage solution for virtualized environments providing VM-level storage abstraction and full integration into server virtualization management scheme and user interface. Maxta Storage eliminates the need for storage arrays (either SAN or NAS) delivering shared storage by aggregating dispersed storage resources from multiple servers. It dramatically simplifies IT by eliminating storage provisioning, storage management and storage networking, enabling administrators to focus on high value add initiatives.

Maxta Storage delivers VM-centric enterprise-class data services to support advanced capabilities such as live migration of virtual machines, dynamic load balancing, high availability, data protection, and disaster recovery.

Maxta Storage enables significant capital savings by converging compute and storage resources on standard commodity servers, without compromising performance or scalability for virtual workloads. Maxta Storage can coexist alongside traditional storage arrays and servers thereby providing investment protection. The versatile architecture provides the ability to scale capacity and performance independently, on-demand, without having to over-provision resources.

The key benefits of Maxta Storage Platform are:

Dramatically Simplify IT

Maxta Storage Platform dramatically simplifies IT by eliminating the need for storage provisioning and storage management. The Maxta datastore and all data services such as local replication, snapshots and zero-copy clones are configured and managed from the VMware vSphere Client at VM granularity. This simplification eliminates the day-to-day tasks of storage management and enables administrators to focus on managing applications and VMs.

Deliver Enterprise-class Data Services

Maxta Storage Platform eliminates single point of failure by delivering best-in-class data resiliency, integrity and high availability leveraging strong checksums and data replication for all stored data. Time, performance and capacity efficient snapshots and clones provide the first layer of data protection and the ability to rapidly provision VMs. Maxta Storage Platform seamlessly integrates with all the advanced capabilities of the server virtualization software.

Maximize Cost Savings

Maxta Storage Platform enables significant capital savings by converging compute and storage resources on standard commodity servers, without compromising performance or scalability. This provides considerable up-front capital savings and even greater savings on upgrades compared to the capital expenses associated with purchasing and expanding storage arrays or storage appliances. In addition, Maxta Storage Platform leverages commodity disk drives, enterprise class SSD's with super capacitors, snapshots, zero-copy clones, thin provisioning, in-line compression to increase storage efficiency and reduce storage expenses. Maxta Storage Platform can coexist alongside traditional storage array and servers thereby providing investment protection. By significantly simplifying IT, increasing IT efficiency, and enabling administrators to focus on managing applications and VMs, Maxta Storage Platform enables dramatic reduction in operating expenses.

Understanding The Maxta Storage Platform Architecture

Below is a high level illustration of how the Maxta Storage Platform fits into virtualized environments.

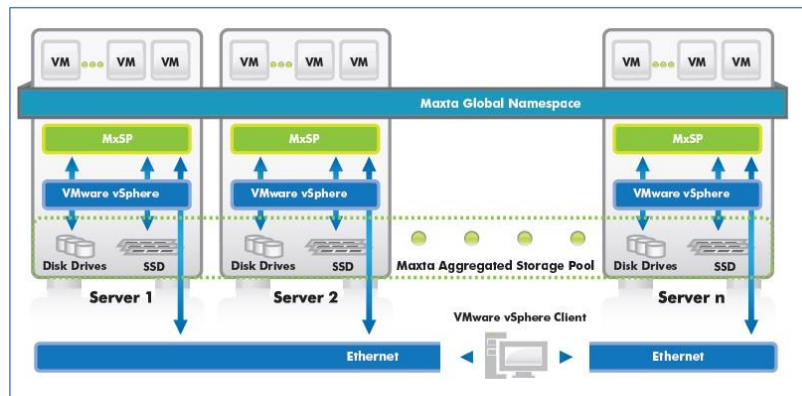


Figure 1: Maxta Storage Platform Architecture

As shown in **Figure 1**, an instance of the Maxta Storage Platform software is installed on each of the servers that are part of a server virtualization cluster. Maxta Storage Platform creates a Maxta Global Namespace that all the virtual machines (that are part of the cluster) can access. The servers running an instance of Maxta Storage Platform can be of two types:

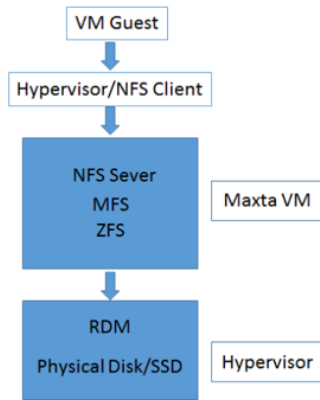
Converged compute/storage servers – these have access to the Maxta Storage Platform storage pool and also contribute storage to it.

Compute-only servers – these have access to Maxta Storage Platform storage pool but do not contribute storage to it. This approach allows for simplistic scale of computer and memory resources when additional storage is not needed.

Maxta Storage Platform intelligently maps VMs to storage resources, optimizing data layout for virtual workloads. Maxta Storage Platform management functionality is seamlessly integrated into the virtualization user interface (vCenter Server), eliminating the need for storage management interface. Maxta Storage Platform leverages any combination of magnetic disk drives and SSDs on standard servers to deliver flash performance and hard disk capacity at an attractive price for all storage workloads in a virtualized environment.

Below is and I/O Path general overview around how data would be accessed in Maxta.

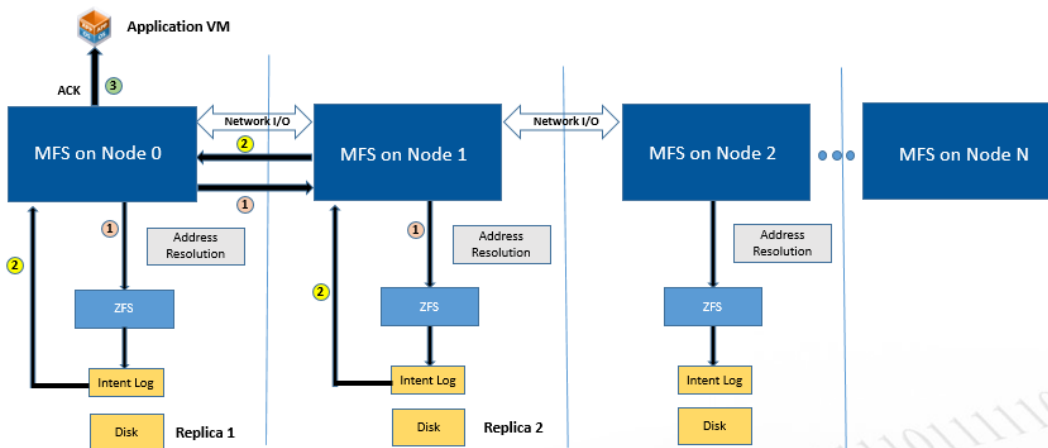
Maxta Storage Platform I/O Path



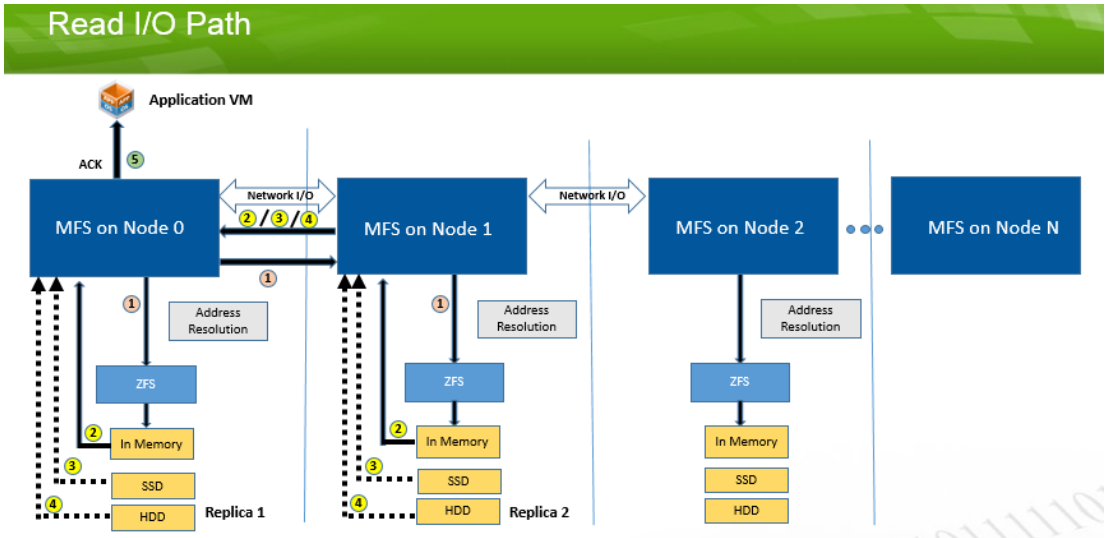
✓ Maxta Storage Platform creates a Datastore and presents it over NFS to the VMware vSphere Hypervisor

A Maxta Write I/O will be processed through local Maxta Storage Platform VM resides on the ESXi host closest to the VM Guest.

Write I/O Path



A Maxta Read I/O will be processed through local Maxta Storage Platform VM resides on the ESXi host closest to the VM Guest.



Maxta Metadev and SSD Caching

Maxta Metadev is a class of device in the Maxta file system. Metadevs store only the file system metadata to accelerate the read/write performance of metadata IOs. Maxta deployments typically use small page sizes (4K or 8K) for space efficiency and performance. Smaller page sizes result in larger amounts of metadata and that create random read/write IO performance issues since metadata can't be cached entirely in memory.

Metadev is a way to separate metadata from data allowing the metadata to be stored on high performance devices (like SSDs). This approach significantly accelerates the metadata read and write operations. Metadevs stores metadata such as file indirect blocks and file system's space allocation records providing key information to the actual data itself.

The HDD to Metadev (SSD) ratio depends specifically on the page size defined during the Maxta Storage Platform deployment. If the deployment is using a 4K page size, the recommended Metadev capacity is 5% of the node's HDD capacity. If the deployment is using 8K page size, then it's 3% of the node's capacity.

There are two types of Metadevs: Shared and Exclusive. Exclusive Metadevs are allocated on SSDs that do not have read/write-back cache partitions. Shared Metadevs can be created on SSDs that also host read/write-back cache partitions. Exclusive Metadevs are created to reduce the wear on the Metadev SSDs. eMCL/PCIe SSDs that support at least 10 drive writes per day for 5 years (this is typical) can be configured to use shared Metadevs. Metadevs can be either shared or exclusive. Users cannot configure both at the same time. During the Maxta Storage Platform installation users will be able to select either an Exclusive or a Shared Metadev configuration. If 2 SSDs are available and a user selects "Enable Metadev" an Exclusive Metadev will be configured with 1 SSD assigned to Exclusive Metadev and the other to read/write-back cache.

If users cannot meet the minimum Metadev SSD ratio requirements Maxta will still utilize the available SSD's automatically for read/write-back cache.

Maxta Storage Platform Prerequisites and Requirements

This section describes the Maxta Storage Platform requirements for this deployment guide. This section includes the VMware Software elements, the storage requirements and the network elements. Wherever appropriate we differentiate between the virtualized devices such as Virtual Switches and the physical devices e.g. SATA disks.

Note: For the purposes of this deployment guide we tested the Maxta Storage Platform on VSphere but other hypervisors are supported.

VMware vSphere Requirements

Maxta Storage Platform requires as a basis the following software components:

1. VMware ESXi version 5.0 U1 or later. We tested with VMware ESXi 5.5U2.
2. VMware vCenter Server 5.0 U1 or later. We test with VMware vCenter Server 5.5U2 (vCenter Server Appliance or Windows based vCenter).
3. VMware vSphere Distributed Switches (VDS) or VMware standard switches (VSS).
4. Maxta Storage Platform version 2.4.1-35-gc721657.
5. Maxta Storage Platform License.

Maxta Hardware Requirements

Component	Quantity
Minimum Number of Servers	3
Converged Compute/Storage Server Specifications	Quantity
Minimum Number of Converged Compute/Storage Servers	2

Converged Compute/Storage Server Specifications

Component	Specification
Processor	Dual-Socket Intel Xeon CPU 2.00GHz. 4 cores per socket with Hyper-Threading enabled or higher performance
Memory	32 GB or higher
Network	GigE port or higher
Storage Controller	Any controller that will allow disabling of RAID
SSDs	Enterprise Mainstream SSD 60GB or Higher
HDDs	7.2K RPM SATA/SAS Drives 200GB or Higher

Compute Only Server Specification

Component	Specification
Processor	Dual-Socket Intel Xeon CPU 2.00GHz. 4 cores per socket with Hyper-Threading enabled or higher performance
Memory	32 GB or higher
Network	GigE port or higher
Storage Controller	Any controller that will allow disabling of RAID
HDDs	7.2K RPM SATA/SAS Drives 200GB or Higher

Note: The Hardware for Maxta on VMware has to be listed in the VMware Hardware Compatibility List (HCL)

<http://www.Vmware.com/resources/compatibility/search.php>

Storage Requirements

Disk Controllers

Each vSphere host that contributes storage to the Virtual SAN cluster requires a disk controller. This can be a SAS or SATA host bus adapter (HBA) or a RAID controller. However, the RAID controller must function in one of two modes:

- Pass-through mode
- RAID0 mode

Pass-through mode, commonly referred to as JBOD or HBA mode, is the preferred configuration for Maxta Storage Platform because it enables Maxta to manage the RAID configuration settings.

The HP H220 Host Bus Adapter was selected so that Maxta can directly manage all HDDs and SSDs.

Flash-Based Devices

Each vSphere host must have at least one flash-based device—SAS, SATA, or PCI Express SSD—to participate in the Maxta Storage Platform cluster. Flash-based devices provide for caching and metadata.

It's important to note that Maxta read/write performance is mostly dependent on SSD class and flash capacity ratio based on anticipated storage consumption.

SSDs are categorized into five classes by VMware:

- Class A: 2,500-5,000 writes per second
- Class B: 5,000-10,000 writes per second
- Class C: 10,000-20,000 writes per second
- Class D: 20,000-30,000 writes per second
- Class E: 30,000+ writes per second

In the lab testing, each HP ProLiant SL210t Gen8 server had six drives. Two out of six drives were 100GB SSD (691862-B21). This SSD is rated as Class-C category (10-20K writes). For optimal performance it is recommended to use Class-E SSDs (30K+ writes). The two SSDs were installed in drive bays five and six. One SSD was used for caching and the other SSD was used for Metadata. vSphere ESXi was installed and set to boot on drive bay one.

Note: Flash-based devices do not contribute to the overall size of the distributed Maxta shared datastore. They count only toward the capacity of the Maxta SAN caching/metadev tier.

Maxta recommendation is a 10 percent ratio of flash/Solid-State versus magnetic disks for read/write-back caching.

Hard Disk Drive Requirements

Each VMware vSphere host must have at least one SAS, near-line SAS (NL-SAS), or SATA magnetic hard-disk drive (HDD) to participate in the Maxta cluster. HDDs account for the storage capacity of the Maxta shared datastore. Additional magnetic disks increase the overall capacity and can also improve virtual machine performance because the virtual machine storage objects may be striped across multiple physical volumes/spindles.

In our lab testing, each HP ProLiant SL210t Gen8 server had six drives. Four out of the six drives were 1.2TB HDD (697574-B21). HDDs were installed in drive bays 1-4. All 4 HDDs were used for Maxta capacity.

Network Requirements

This section outlines the network interface cards (NIC) and (virtual) switch interfaces.

Network Interface Cards

Each VMware vSphere host must have at least one network adapter. It must be 1Gb Ethernet or 10Gb Ethernet capable, but VMware recommends 10Gb.

The HP Flexfabric 10Gb 2-port 554FLR-SFP+ adapter was used in our lab testing.

Supported Virtual Switch Types

Maxta is supported on both the VMware vSphere Distributed Switch (VDS) and the vSphere standard switch (VSS).

HP ProLiant SL2500 Maxta Installation and Configuration

This section details the Maxta Storage Platform installation, setup and configuration on the HP ProLiant SL2500 servers.

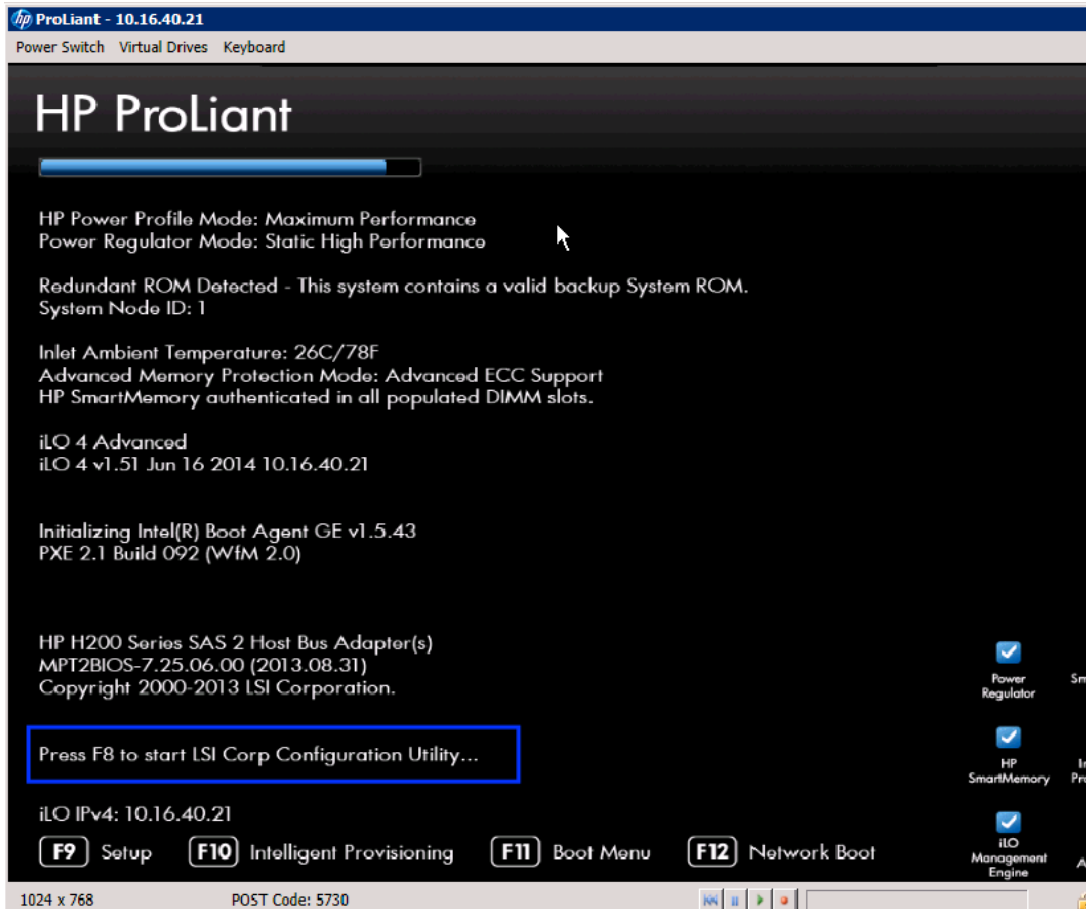
The Maxta Storage Platform installation consists of the following steps:

- Installing VMware ESXi 5.5U2 on a HP ProLiant SL210t Gen8 server local disk.
- Configure VMware ESXi hosts with VMkernel interfaces for management purposes.
- Installing VMware vCenter 5.5U2 and creating a datacenter.
- Create a 4-Node vSphere Cluster.
- Assign IP Address for Maxta .

Note: It is assumed that users have a vCenter Server available to manage this installation footprint and that these hosts will be joined to this vCenter along with the complete creation of a vSphere Cluster.

Configuring the SL2500 HP H220 HBA

- Loading ESXi on HP ProLiant SL210t Gen8 servers is accomplished by using iLO virtual media³ or PXE booting. Important note: mark a specific HDD or SSD in the HP H220 Host Bus Adapter (HBA) BIOS as a boot drive before VMware ESXi installation in order for the server to boot from the installed VMware ESXi image.
- During initial setup/installation press “F9” during the booting process to disable the embedded HP B120i controller.
- After the HP B120i has been disabled, press “F8” during the server booting process to configure the HP H220 HBA BIOS.



- Select Option 2 for the SAS configuration Utility.

```
HP H200 Series SAS 2 Host Bus Adapter(s)
MPT2BIOS-7.25.06.00 (2013.08.31)
Copyright 2000-2013 LSI Corporation.

Please wait, invoking SAS Configuration Menu...

Please select a configuration option:
  1. Enable tape based One Button Disaster Recovery (OBDR)
  2. Invoke SAS Configuration Utility
```

- Press "Enter" to get into H220 config mode.

```
LSI Corp Config Utility for HP H200 v7.25.06.00 (2013.08.31)
Adapter List Global Properties
Adapter      PCI   PCI   PCI   PCI   FW Revision   Status   Boot
              Bus  Dev  Fnc  Slot
H220        06   00   00   01   15.10.06.00-IT  Enabled  0
```

- Select "SAS Topology" and press "Enter".

```
LSI Corp Config Utility for HP H200 v7.25.06.00 (2013.08.31)
Adapter Properties -- SAS2308-2

Adapter      H220
PCI Slot     01
PCI Address(Bus/Dev) 06:00
MPT Firmware Revision 15.10.06.00-IT
SAS Address  500605B0:06E0A81A
NUDATA Version 0F.00.00.15
Status       Enabled
Boot Order   0
Boot Support [Enabled BIOS & OS]

SAS Topology

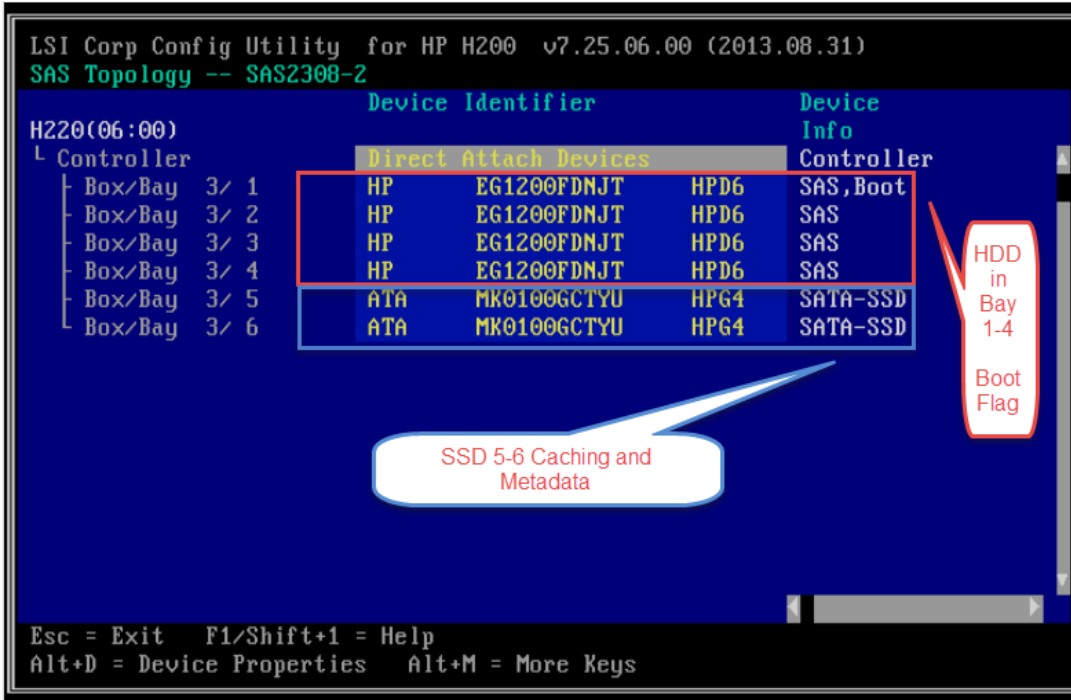
Advanced Adapter Properties
```

- Press "Enter" for a list of "Direct Attach Devices"

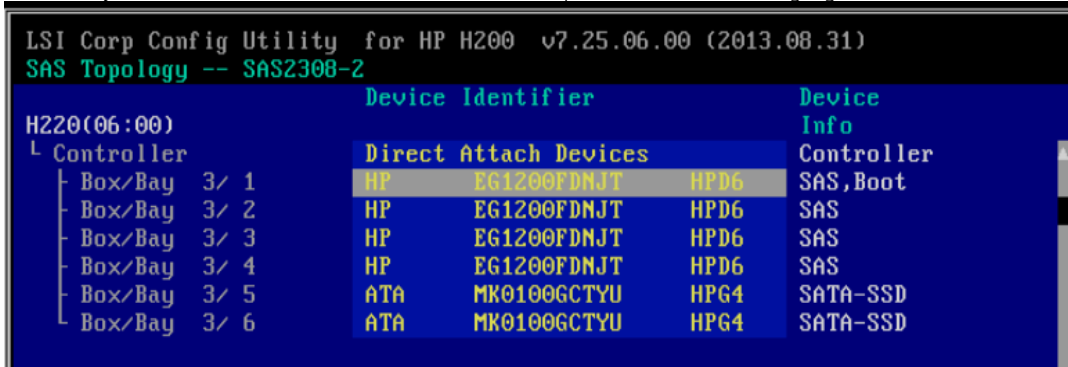
```
LSI Corp Config Utility for HP H200 v7.25.06.00 (2013.08.31)
SAS Topology -- SAS2308-2

Device Identifier      Device
H220(06:00)           Info
└ Controller          Direct Attach Devices Controller
```

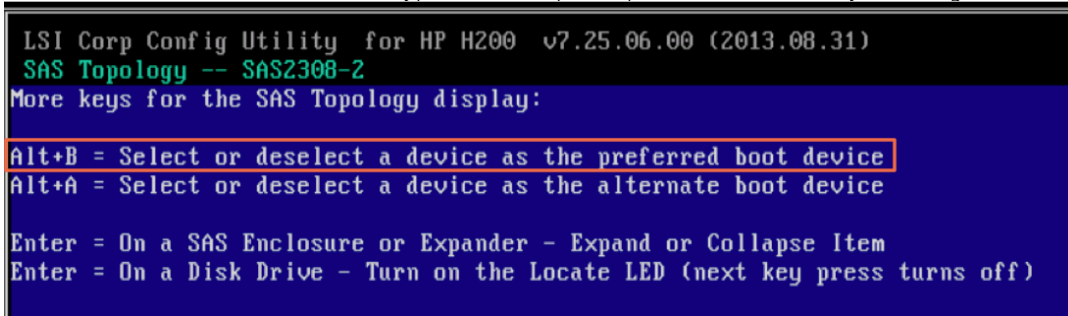
- Notice that HDDs are installed in bays 1-4 and SSDs in bay 5 and 6. The following screen shows that the HD in bay 1 has been marked as the "boot" disk.




- Users may want to choose another drive for a boot drive, please scroll down to highlight the drive of choice.



- “Alt-B” to mark the drive as boot drive. Type “Alt-M” to open help menu for more hotkeys including “Alt-B”.



- During Vmware ESXi image installation make sure the install of the ESXi image to the correct drive being marked as the boot drive in the H220 BIOS takes place. Users should choose a low capacity HDD or SSD for ESXi installation and boot.



Setting up the Vmware VSphere Environment for Maxta Storage Platform

Before installing Maxta Storage Platform, the Vmware infrastructure must be configured. The high-level steps are described below. For detailed information on each of these steps, refer to Vmware documentation at www.Vmware.com/support/pubs.

Install Vmware VSphere

The next step is to install Vmware vSphere 5.5 U2 on each of the servers.

For detailed information about vSphere installation requirements and process, refer to Vmware’s vSphere Installation and Setup guide, available at <http://www.Vmware.com/support/pubs>

Install vCenter Server

vCenter Server enables users to centrally manage vSphere hosts or virtual machines. Users can install vCenter Server on either a physical or a virtual Windows machine. The host that runs the vCenter Server can be part of the Maxta Storage Platform. For the purposes of testing we installed the Vmware VCenter 5.5 U2 Appliance on a local disk part of the soon to be Maxta cluster.

For detailed information about vCenter Server installation requirements and process, refer to Vmware’s vCenter Server and Host Management document available at <http://www.Vmware.com/support/pubs>

Create a datacenter and then create a cluster within the datacenter

With the above set up in place, users can now create a datacenter and create a cluster within vCenter Server. A datacenter is a logical entity that contains a collection of all the different types of objects needed in virtual infrastructure: hosts, virtual machines, networks, and datastores.

A cluster is a collection of vSphere hosts and associated virtual machines. When a user add a host to a cluster, the host’s resources become part of the cluster’s resources. The cluster manages the resources of all hosts. When a user creates a cluster, it can be enabled for Vmware High Availability (HA) and Distributed Resources Scheduler (DRS).

Note: Users can also leverage an existing datacenter and add a cluster to it, or use existing clusters to install Maxta Storage Platform.

To create a datacenter

Open a vSphere Client session to a vCenter Server.

Go to Home > Inventory > Hosts and Clusters.

Select File > New > Datacenter. A new datacenter object appears in the inventory.

Enter a name for the datacenter and press Enter.

To create a cluster

In the vSphere Client, select the datacenter to add a cluster.

Select File > New > Cluster. A new Cluster wizard appears in the inventory.

Enter a name for the cluster, leave other parameters with default values, and complete the wizard.

Add hosts to the cluster

When a user creates a cluster, it does not include any hosts or virtual machines. Users can add hosts to a cluster in vCenter by following the steps listed below. Maxta Storage platform requires a cluster with a minimum of three hosts and can be scaled up to eight hosts. When users add a host to the cluster, vCenter gathers all the host information such as network card, storage and creates its own repository, which is called an inventory. The inventory is used to keep track of all this information in vCenter.

To add hosts to the cluster

Go to Home > Inventory > Hosts and Clusters.

Select the cluster in the inventory.

Select File > New > Add Host.

On the Connection Settings page of the Add Host Wizard, enter the IP address and root credentials of the vSphere host to add and click Next.

On the Host Summary page, review the information about the host and click Next.

On the Assign License page, select Assign a new license key to this host option if the vSphere host does not have an assigned license, and assign a new license.

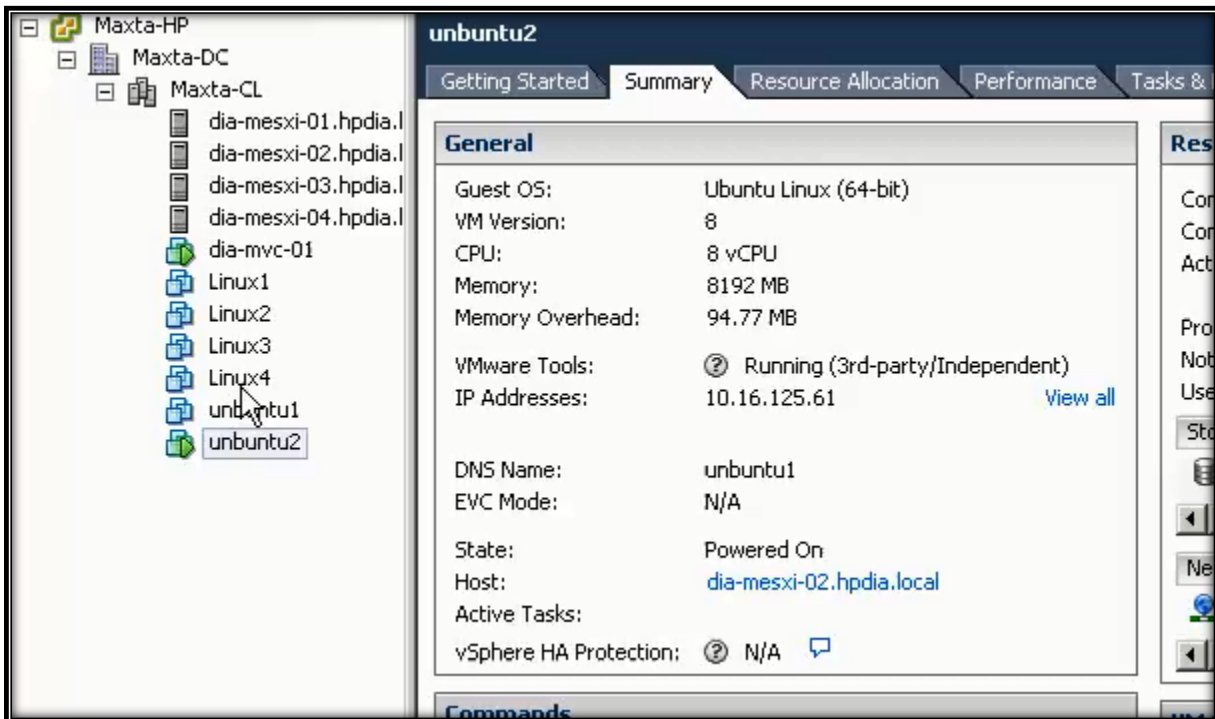
Click Next.

On the Virtual Machine Location page, select the datacenter created and click Next.

On the Ready to Complete page, review the selections and click Finish.

Repeat the above steps to add another host.

This is what should have in place from a VSphere perspective prior to the installation of Maxta Storage Platform.



Installing the Maxta Storage Platform


At this point, users have the Vmware infrastructure ready to install the Maxta Storage Platform software. Users have two installation options for installing the Maxta Storage Platform which is either going to be CLI or GUI based. For the purpose of this testing we used the GUI installation method. CLI installation will also be touched on in this Deployment Guide.

During the Maxta installation the user will be prompted for some information but it is important to understand some of the key items to anticipate.

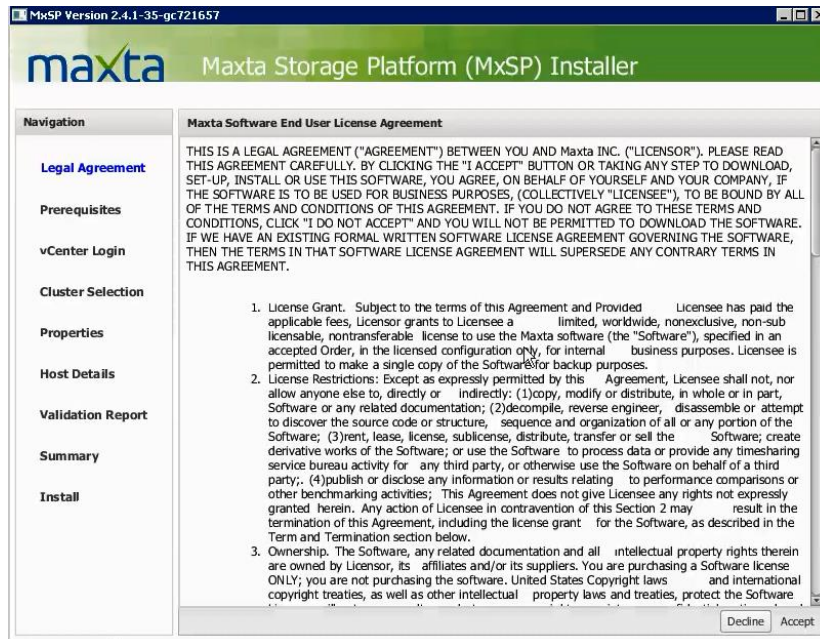
Field Name	Description
Management Server IP	The Management Server IP is the IP address that is used to manage the cluster. Each cluster should have its own unique management IP address.
Storage Subnet Prefix	Maxta needs a group of IP addresses for internal communication. The installer generates a default group of IP addresses but users can also customize if. For the installation we tested we just used the default values provided. Note: During the Installation Maxta will create Vmware port groups attaching the Maxta VM to both an internal network and Management Network. The user will have an option to specify a VLAN ID that will be used for the Management Network. Refer to Figure 1: for a high level understanding.
Enable Local Copies	Maxta automatically maintains two copies of data (<i>on 2 different converged Compute storage servers</i>) to provide data redundancy for protection in case of disk or server failures. If the Enable Local Copies checkbox is selected, data is also mirrored within the vSphere host. This will avoid rebuild over network when drives fail, improving rebuild duration, minimizing network traffic, and protecting against any three drive failures.
Rack Awareness	Hosts in a cluster are distributed across multiple racks or sites to protect them from power loss or failure of the entire rack or site. If the Rack Awareness checkbox is selected, Maxta distributes the replicas of objects (VMs/Snapshots/Clones) across multiple racks or sites, thereby delivering continuous data availability in the event of a rack or site failure. Note that a minimum three rack IDs are required for this option to be enabled. Note: This option is also important to select if the user intended to add additional nodes to this Maxta Cluster after the day one build.
Disk Usage Rule	Select the type of drives that are used. Select "Hybrid" if a combination of Solid State Drives (SSD) and HDD are installed in the servers. Select "All Flash" if only SSDs are installed in the servers.
Page Size	Select the page size to be configured. The default page size is 4K. Aligning the page size with the application page size will deliver better performance Note: This will be the default page size for all VMDKs created in the Maxta Datastore but Maxta give users the ability to customize the page size on a per VMDK basis that will be covered in this Deployment Guide.
Metadev Enabled	Exclusive Metadevs are allocated on SSDs that do not have read/write-back cache partitions. This option will dedicate an SSD for Metadev. Note: A user cannot have both Metadev options at the same time.

Field Name	Description
Shared Metadev	<p>Shared Metadevs can be created on SSDs that also host read/write-back cache partitions. This option allows the user to share the Metadev across all SSD's used in the read/write-back cache. During testing we leveraged a shared Metadev.</p> <p>Note: Exclusive Metadevs are created to reduce the wear on SSDs. eMCL/PCIe SSDs that support at least 10 drive writes per day for 5 years (this is typical) can be configured to use shared Metadevs. By separating out the Metadev from the read/write-back cache additional performance can be achieved as well.</p>

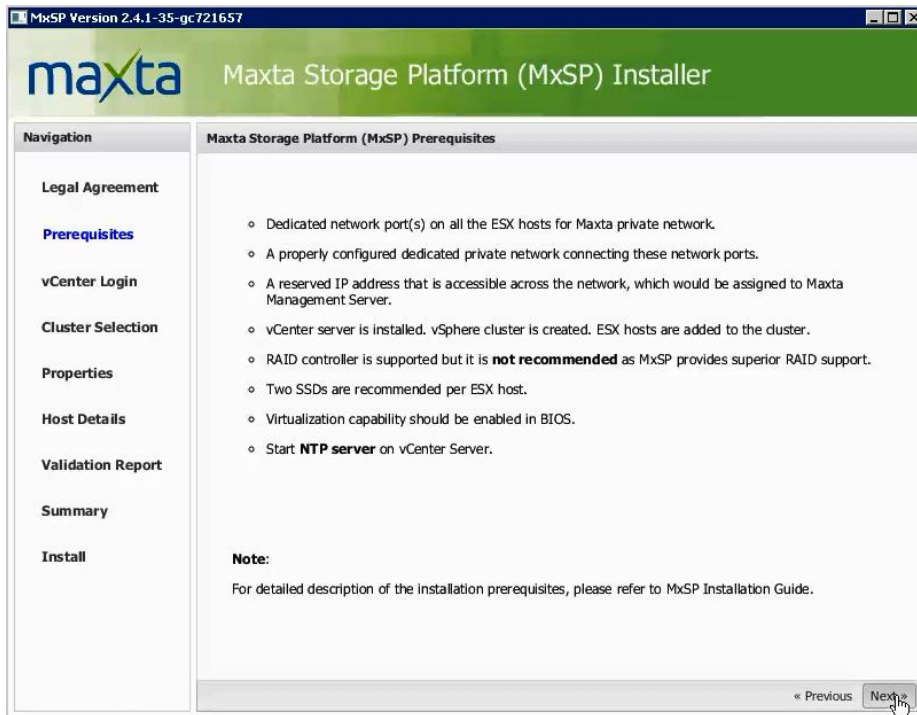
- The Maxta Software Platform package will need to be extracted into a folder prior to clicking on the Maxta executable file. IP connectivity is needed for the workstation or host that the user will run the Maxta executable from and this connectivity needs to extend into the hosts and Vmware VCenter that Maxta will be deployed on.



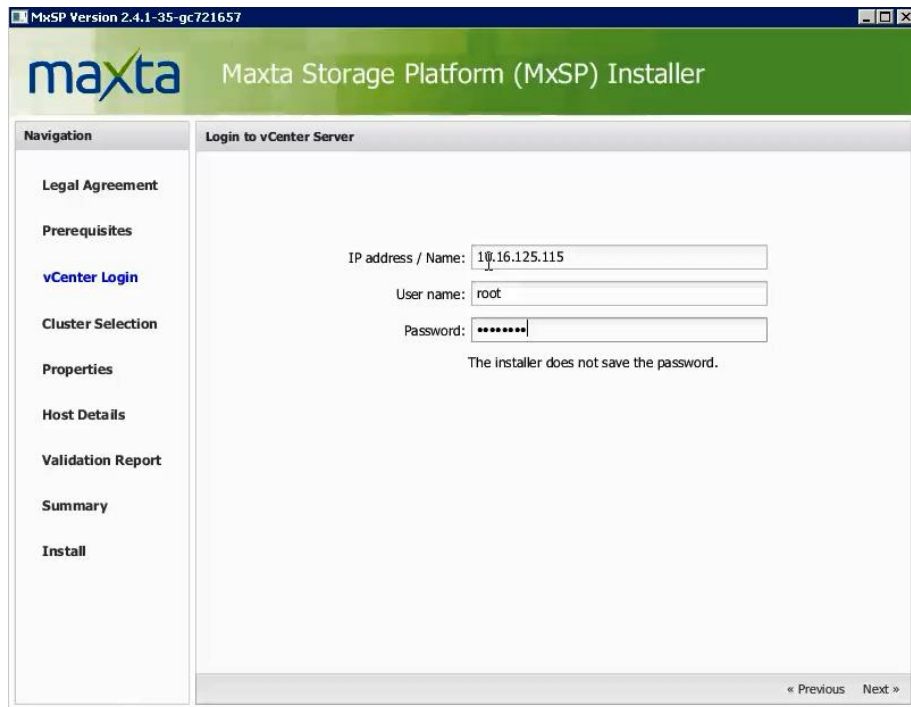
- During the installation users will see the license agreement screen simply click "Accept".



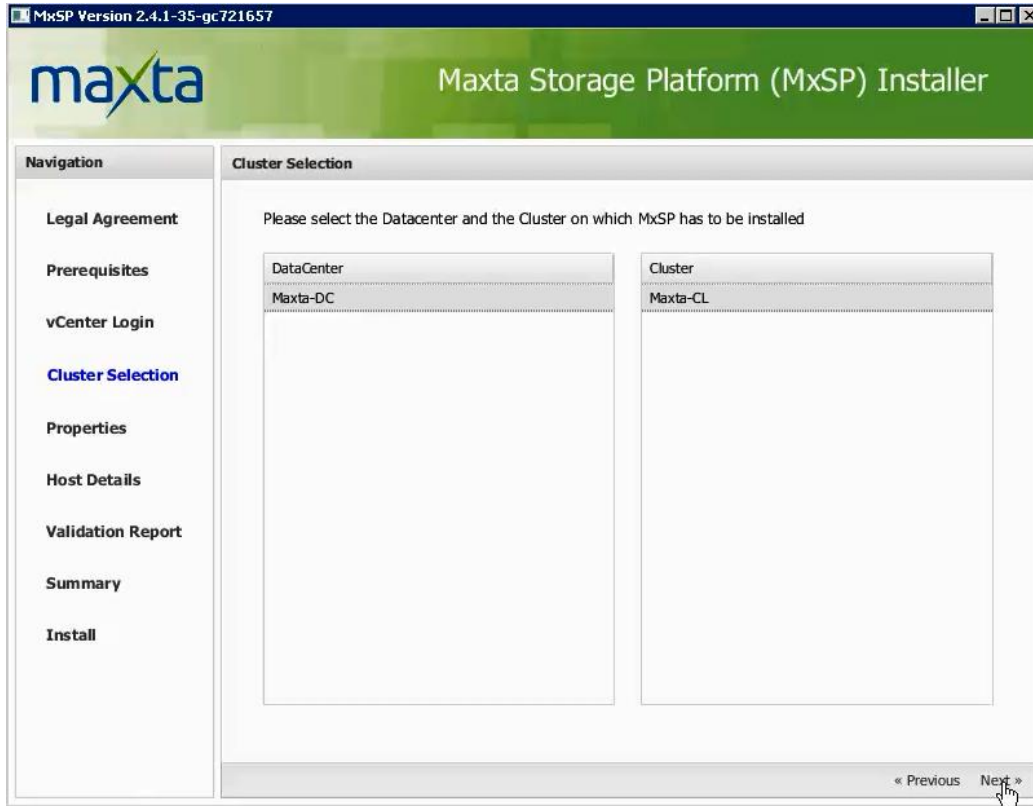
- Maxta will provide users with a quick reminder of the installation prerequisites. The user may want to review this just to make sure something was not missed. Simply click “Next” if everything is in order.



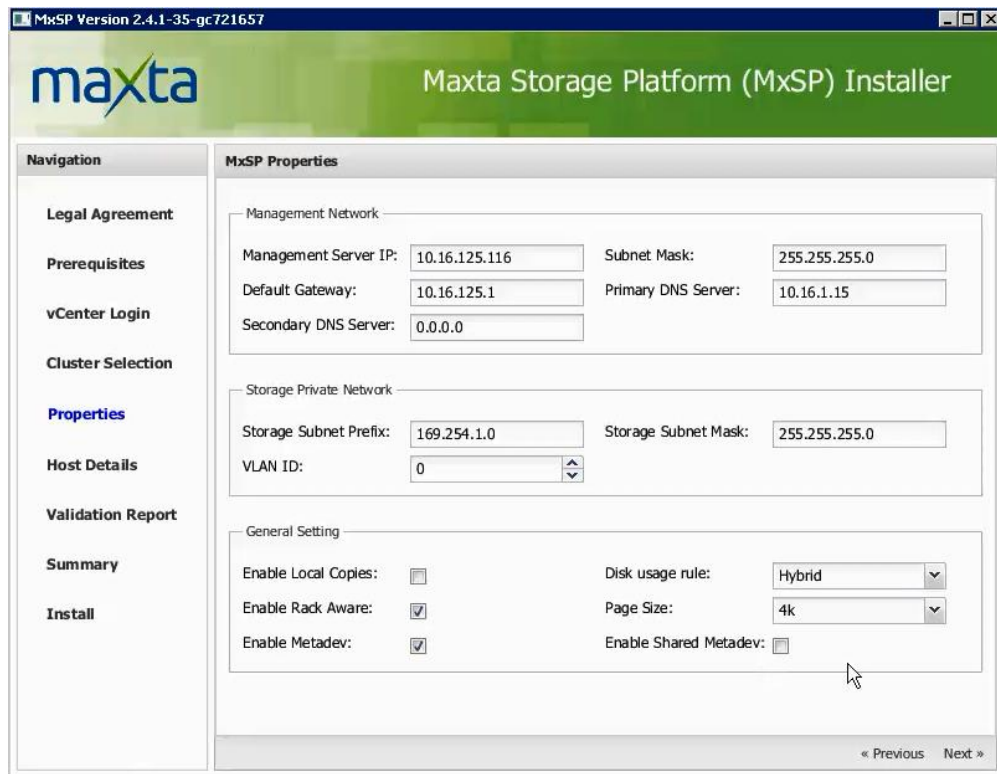
- In the vCenter login screen, provide the login credentials and IP of the vCenter Server that manages the vSphere hosts where Maxta software will be installed.



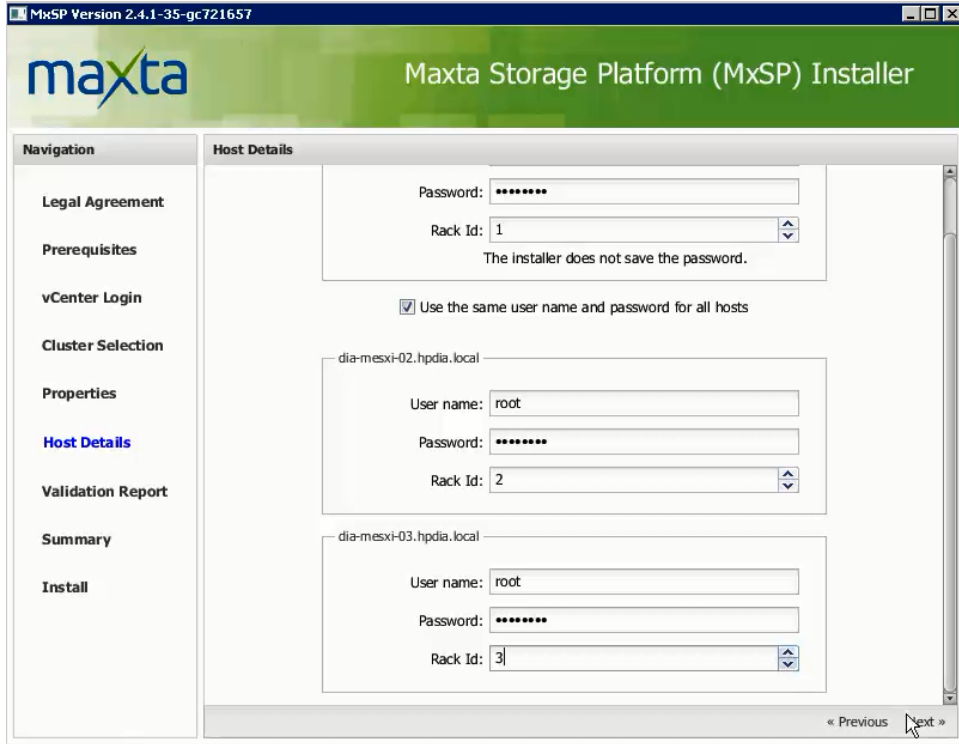
- The select Cluster screen appears so select one Datacenter and Cluster a user would like to install Maxta on and click “Next”.



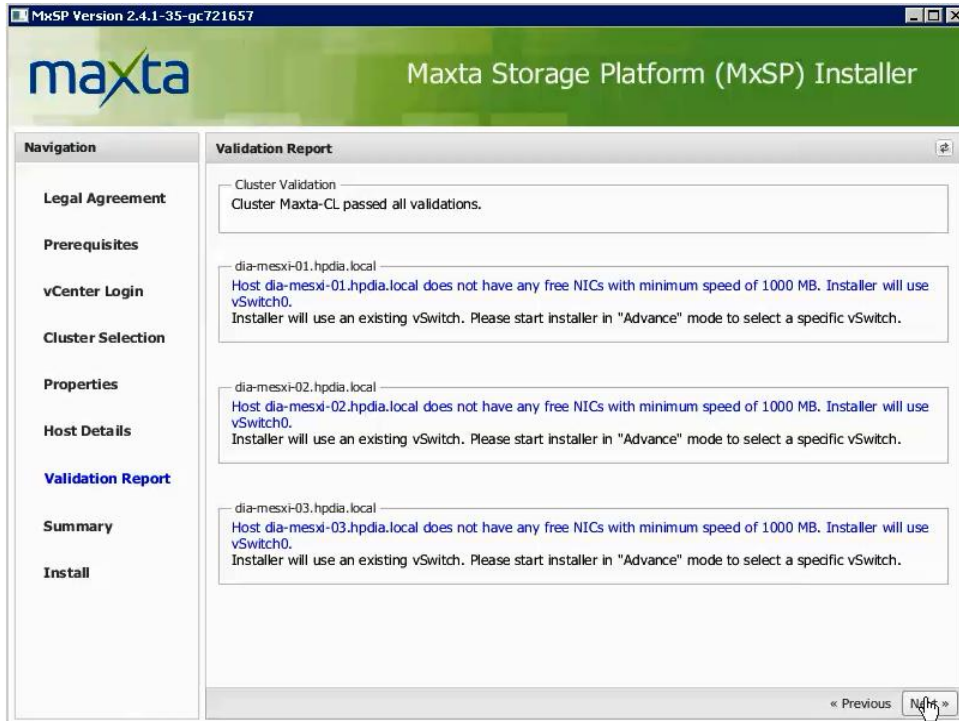
- In the Properties screen, enter configuration information of the Maxta Storage Platform and click “Next”.



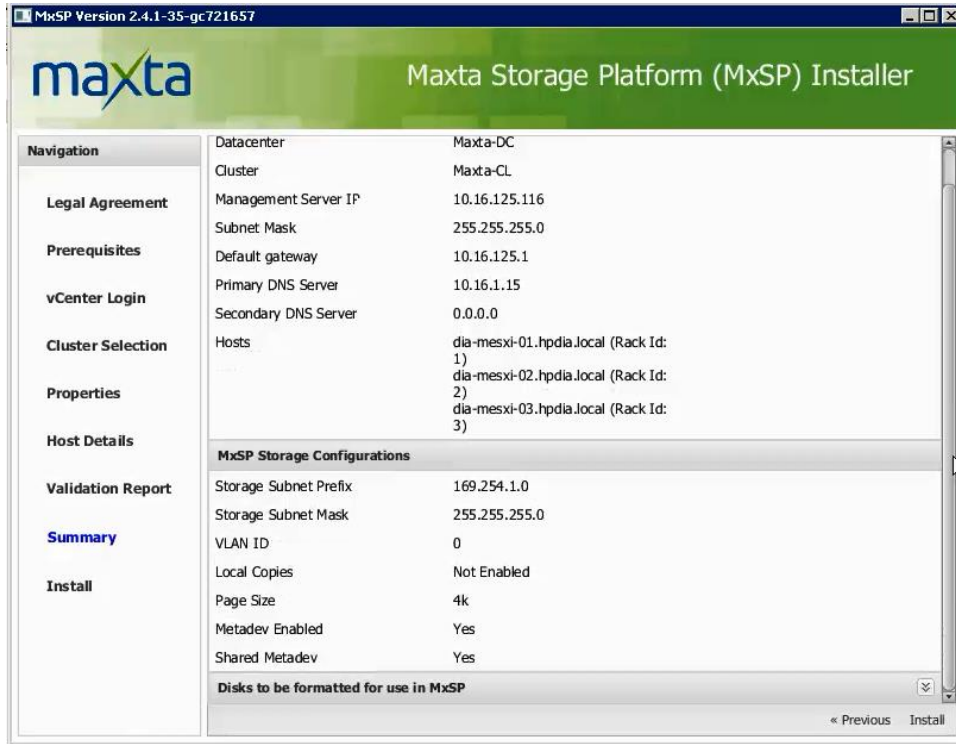
- In the Host Details screen, select the host credentials that are required to validate the configuration of hosts and create the Maxta virtual machines. If Rack awareness is enabled, the installer prompts to select the Rack Id for each host. For the purposes of testing rack awareness is enabled as the desire to test adding an additional compute node as well as expanding capacity by adding a converged compute node was a use case. After the users has made their selections in regards to Rack Id click “Next”.



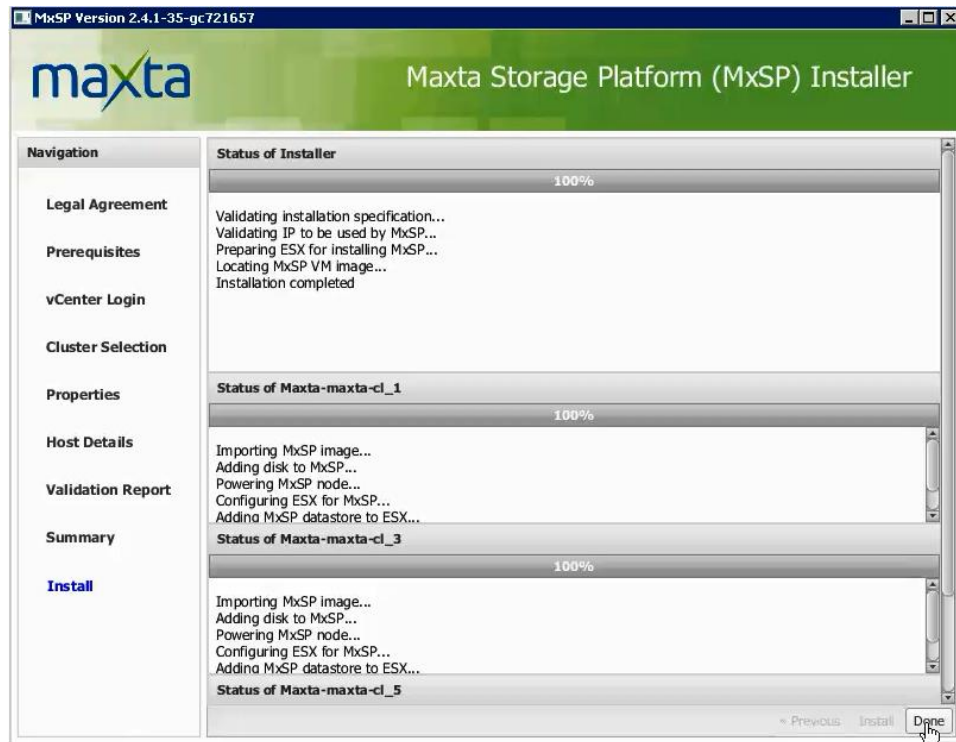
- The Validation Report screen appears as show below simply click “Next”.




- At this point everything has been validated and the Maxta Installer is ready to begin the installation just click "Install" and it will start configuring the VSphere Cluster for Maxta.



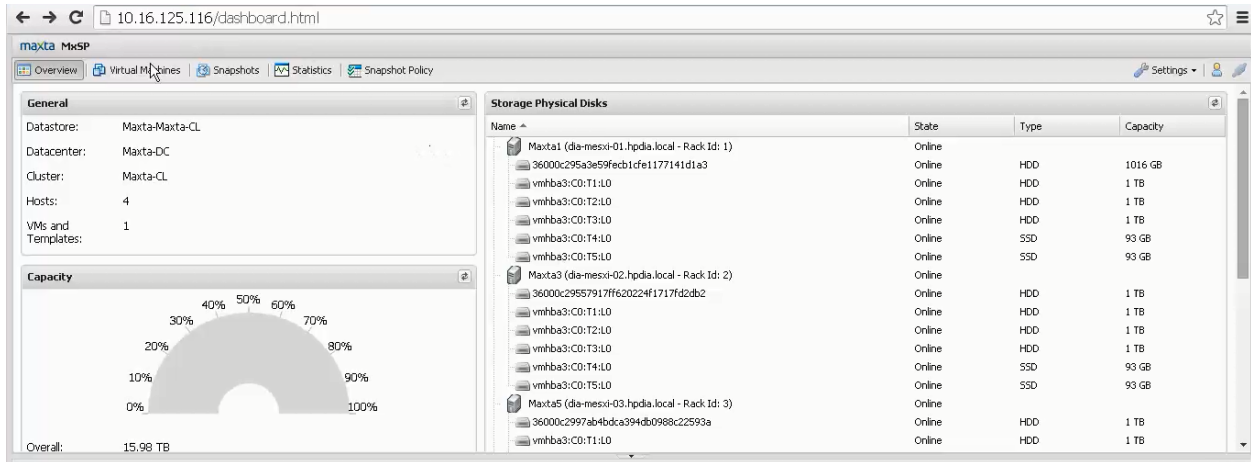
- The status of the installer page will update for the progress bars for each node in the Maxta configuration until they installer reaches 100% at this point the Maxta installation will be finished and the user can click "Done".



- After the installation is complete the user can verify the operation of Maxta by opening up a web browser and navigating to <http://MANGEMENTIP> and logging with the VCenter credentials.



- When users login to the Maxta web interface, they can see the Datastore name that was created along with verifying the number of nodes and status of Maxta nodes. Users can also launch the VSphere Client as well and verify that the Maxta datastore is presented to the hosts.

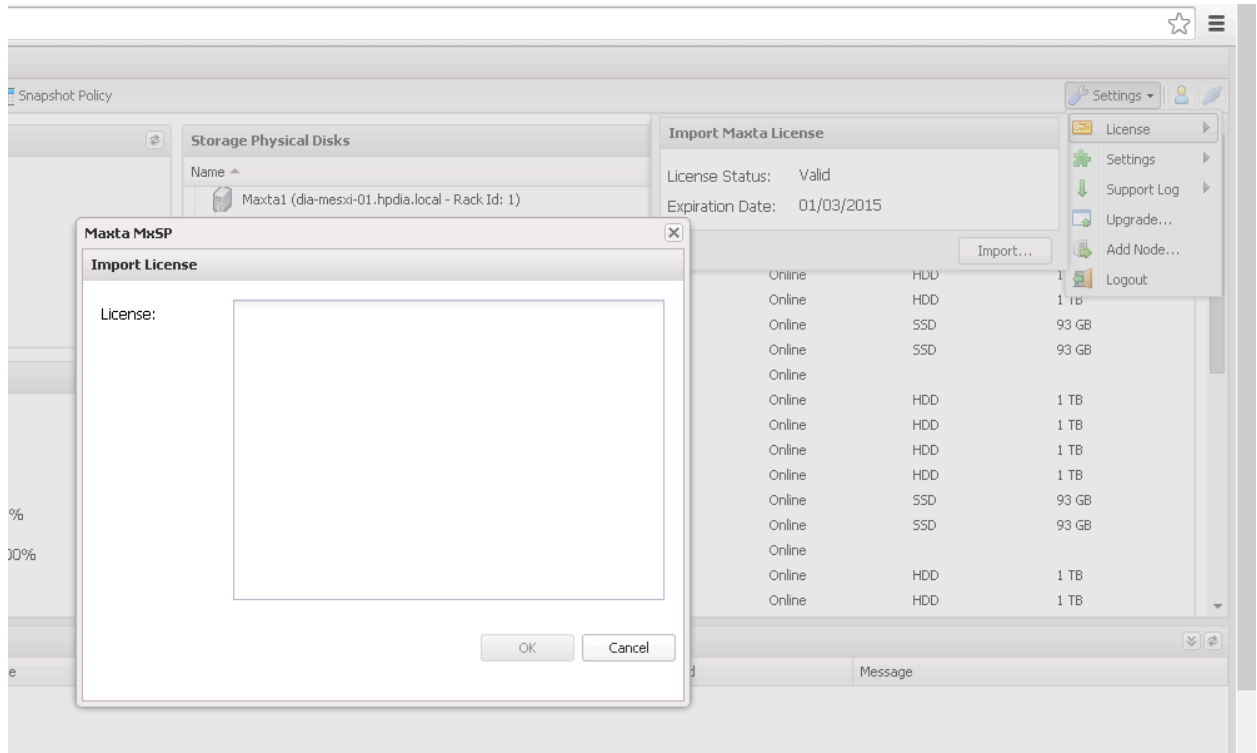


Maxta Administration

Applying the Maxta Storage Platform License

After installing the Maxta Storage Platform, import/apply license to use Maxta features, such as managing snapshots and clones, or managing other operations of Maxta. To import the Maxta Storage Platform license:

1. In the vSphere client, click the Maxta tab.
2. On the Maxta tab, select Settings on the upper-right corner of the screen.



3. Select License and then click Import on the Import Maxta License screen.
4. On the License page, provide the correct license sent by Maxta and click OK.

Note: Users can also view the validity of the license on the Import Maxta License screen.

Add a Node to Maxta

The user can add vSphere hosts to the Maxta storage platform. If Maxta installation was performed with Rack awareness enabled, then in the Add Node screen, a new option "Rack" is displayed. Rack drop-down list displays all the Rack IDs, created as part of installation. The new node being added will be part of one of the Rack IDs that was already created so users need to consider the Rack layout in regards to sites and how this cluster was originally designed. Example: A user may have an existing 5 node cluster split between two datacenters with Rack ID 1-2 sitting in DC 1 and Rack ID 3-5 sitting in DC 2 if the user is adding a new node in DC 2 they would want to select Rack 3 in regards to redundancy and replica consideration.

Note: The vSphere hosts should have already been added to the cluster in the vSphere client. With add node operation, the new vSphere host is added to the Maxta storage platform. The minimum hardware requirement mentioned earlier in the document applies when adding new nodes to the cluster.


Users can add two different types of nodes to Maxta which are a Converged node or a Compute only node.

Converged compute/storage servers – these have access to the Maxta Storage Platform storage pool and also contribute storage to it.

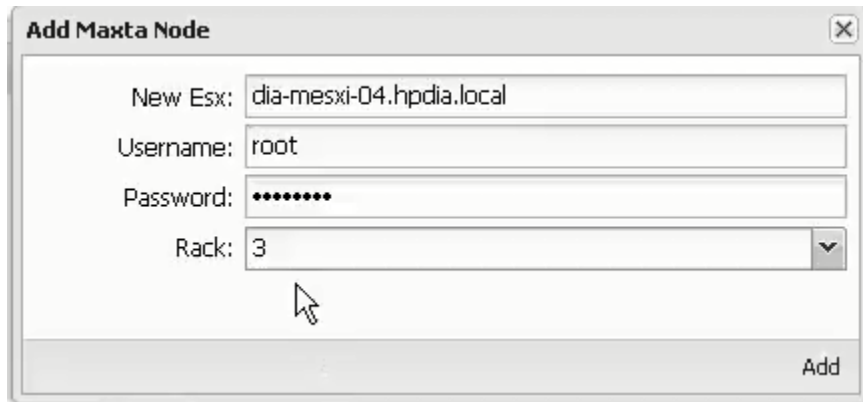
Compute-only servers – these have access to Maxta Storage Platform storage pool but do not contribute storage to it. Maxta will reconfigure the host vswitch or DVS while adding in access to the Maxta datastore. From an installation perspective this host must be part of the same subnet from an IP perspective or the Maxta node add will fail.

To add a new node:

1. In the vSphere client, click the Maxta tab.
2. On the Maxta tab, select Settings, and then click Add Node.



3. In the Add Node screen, enter the ESXi host IP or host name being added to the cluster.
4. Provide the user name and password of the ESXi host being added.
5. From the Rack drop-down list, select the rack to which this host should be added.



6. Click Add. It takes a few minutes for the ESXi host to be added to the Maxta storage platform.

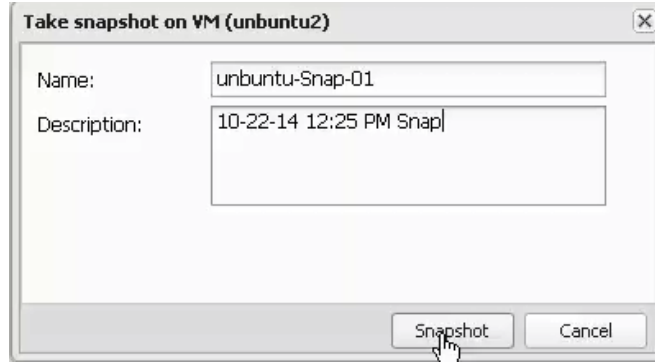
Working with Maxta Clones and Snapshots

Maxta Snapshots

Users can take a snapshot of a virtual machine at any time while a virtual machine is powered on, powered off or suspended. One method of taking snapshots involves using the VMware vSphere Client menus. Another method involves taking snapshots from the Maxta tab. Taking snapshots native in VMware and comparing this to how quickly Maxta can snapshot is something that really stood out during testing. Maxta provides a number of value added features around snapshots and even allows for snapshot policies and retention.

To take a snapshot from Maxta tab:

1. On the Maxta tab, click the Virtual Machines tab.
2. From the list of virtual machines, select a virtual machine for which the user needs to create a snapshot.
3. Right-click the virtual machine and click Take Snapshot.



4. Type a name and description for the snapshot.
5. Click Snapshot to complete the process of taking the snapshot.

Creating a Maxta Snapshot Policy

Users can create a policy to create snapshots as per the defined schedule. Snapshots can be created on hourly, daily, weekly, or monthly intervals. More than one policy can be created for VM.

With each snapshot schedule, users can also specify the number of scheduled snapshots to be retained. Older snapshots that exceed the retention limit are automatically deleted. Additional manual snapshots can be created and retained.

To create a snapshot policy:

1. Go to Maxta tab in the vSphere client.
2. Select the Snapshot Policy tab. The Snapshot Policy List is displayed, which lists all the policies that have been created.
3. Click Add Policy to create a new policy. The Create Snapshot Policy screen is displayed, in which the user must enter the following information.


4. Click Save to save the details and Close to close the screen.

Creating Maxta Clones

A clone is a copy of an existing virtual machine. Users can create a clone using the VMware vSphere Client menu or from the Maxta tab. In the current release of Maxta users take a clone that clone will originate on the original Maxta converge node. Its replica will exist on another converged node that was part of the VM that was cloned. In short when deploying large numbers of clones one may want to take a native VMware clone first and migrate that clone over to another host prior to taking a Maxta clone to optimally distribute the disk load. During the testing of Maxta it was noted that the performance in regards to cloning was significantly faster than that of a native VMware clone operation.

To clone a virtual machine from Maxta tab:

1. In the vSphere Client, select View > Inventory > Hosts and Clusters.
2. From the list of clusters on the left navigation pane, select the cluster on which Maxta has been installed.
3. In the right-side panel, click the Maxta tab and then the Snapshots tab.
4. From the list of snapshots, select a snapshot of the virtual machine for which needs to be cloned.



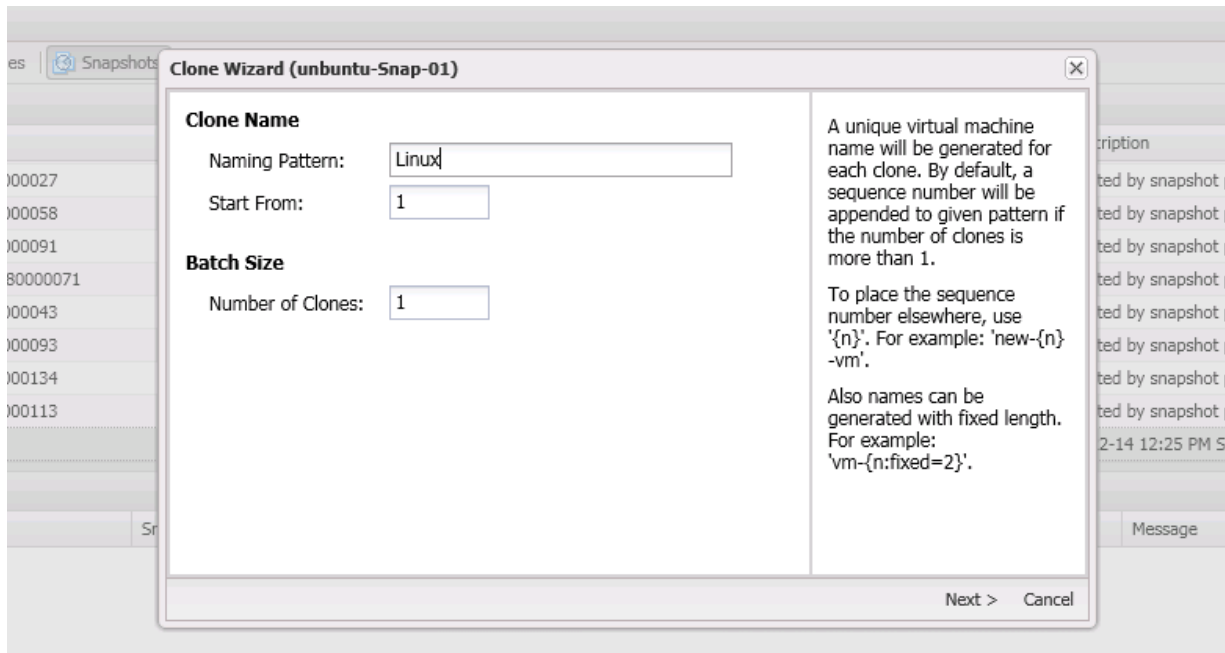
5. Right-click the snapshot and select Clone <snapshot name>.
6. Type a name for the clone.
7. Select the host location where the user wants to store the virtual machine files.
8. Click Clone to complete the process of creating the clone.

Creating Multiple Maxta Clones

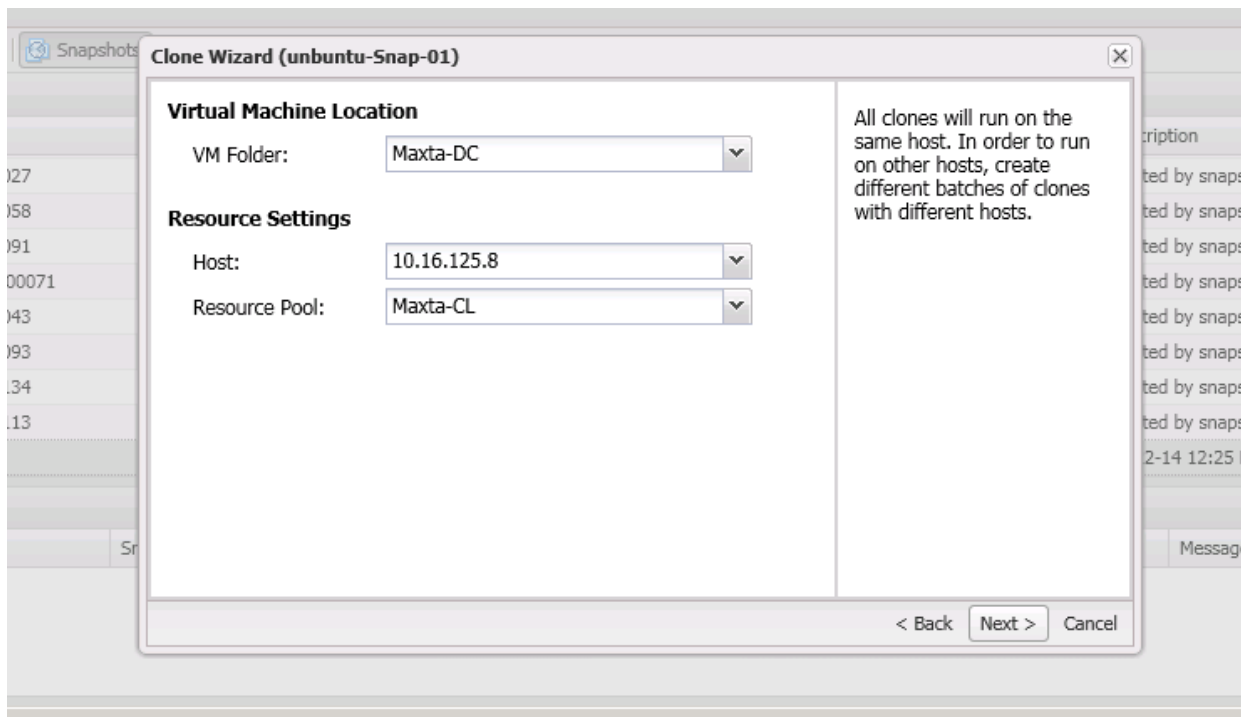
We tested creating multiple Maxta clones using a wizard that allow for high volume clone creating in a minimal timeframe with minimal performance impact. However, using this method of cloning requires and understanding of how Maxta manages replicas and how the cloning process works with those replicas.

To create multiple clones from a snapshot:

1. Go to Maxta tab in the vSphere client.
2. Select the Snapshots tab. Right-click a snapshot and select Clone Wizard. The Clone Wizard opens.
3. In the first screen, provide the following information:
 - a. Naming Pattern: - Enter a prefix for the clones to be created.
 - b. Start From: - Enter the sequence number to append to the name of the clone.
 - c. Number of clones: - Specify the number of clones to be created. For example if the Naming Patter is "My-Name-", the starting number is 5, the number of clones is 2, then after the end of the operation, two clones are created with names "My-Name-5" and "My-Name-6".

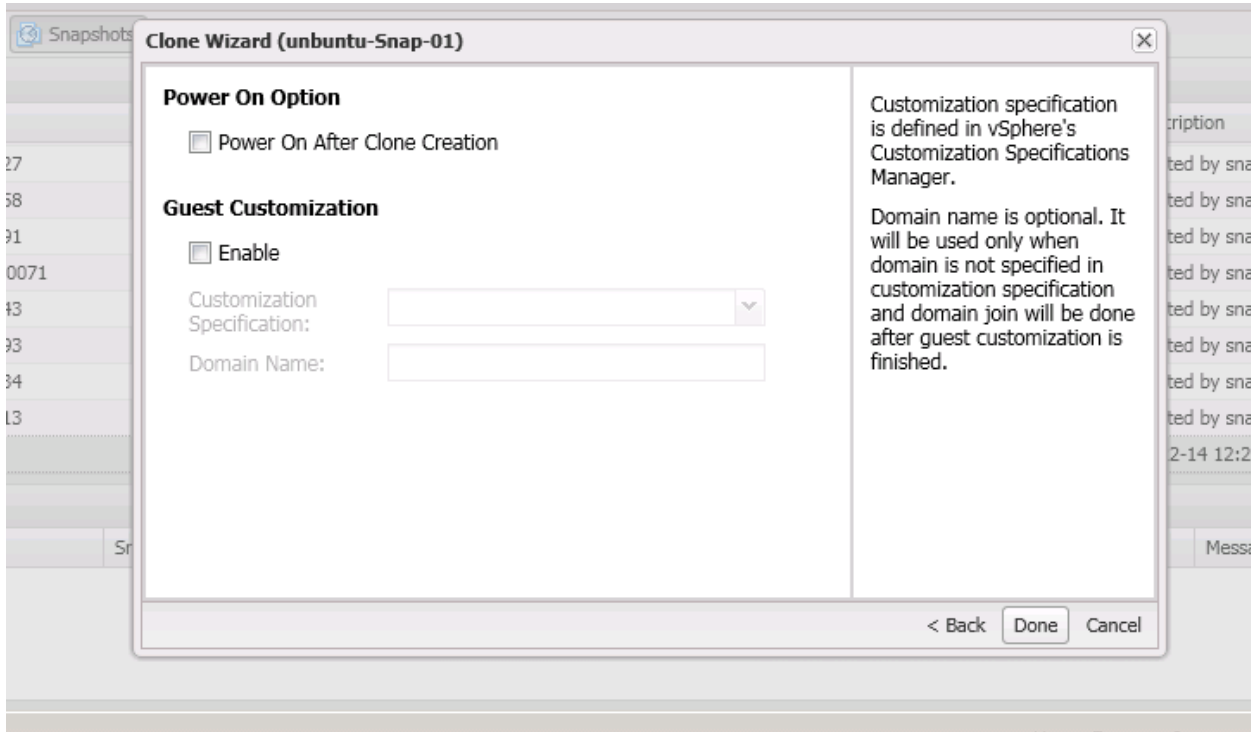


4. In the second screen, provide the following information:
 - a. VM Folder: - Select the datacenter under which the new clones need to be created.
 - b. Host: - Select the host from the drop-down list in which the clones needs to be created. Please note that all the clones will be created under one host. If clones need to be created under different hosts, create different batches.
 - c. Resource Pool: - From the drop-down list, select the resource pool. Clones will use the resources from this pool.



5. In the third screen, provide the following information:
 - a. Power On After Clone creation: - Select this check box if the clones need to be powered on after creation.

- b. Guest Customization: - Select the Guest Customization checkbox if the user wants to use the templates already created using vSphere Customization Specification Manager.



- 6. Click Done to complete the wizard.

Creating Custom Pagesizes for VMDKs

One of the challenges we see is page size misalignment issues and the performance problems they bring into the world of storage. A typical scenario is where a user has a storage array and raid set that is configured with one page size and an Operating System writing to that disk with a completely different page size. This scenario causes misalignment which leads to fragmentation and performance degradation. Maxta solves this problem by giving users the ability to align the VMDK page size with the way the Operating Systems and applications write data to disk by allowing the creation of custom page sizes per VMDK.

To create a VMDK with a custom page size involves a two-step process:

Step I - Create a virtual disk from MxShell using mxvmkfstools.

This step uses the mxvmkfstools command to create a new virtual disk for a virtual machine with a specified pagesize.

Step II – Manually add the newly created virtual disk to the virtual machine through Vmware vCenter.

Step 1: Create a VDisk Using mxvmkfstools

1. Login to MxShell, which is a command-line tool for Maxta. Use the following credentials. Username- console, password maxta
2. Run the following command

```
mxvmkfstools disksize:<virtual disk size> diskname:<vmdk name> pagesize:<page size> vmfolder:<virtual machine folder> where:
```

disksize:<virtual disk size> is the size of virtual disk

diskname:<vmdk name> is the name of virtual disk to be created. For example, Hard.vmdk

pagesize:<pagesize> is the VMDK page size. Must be power of 2 between 4K and 128K

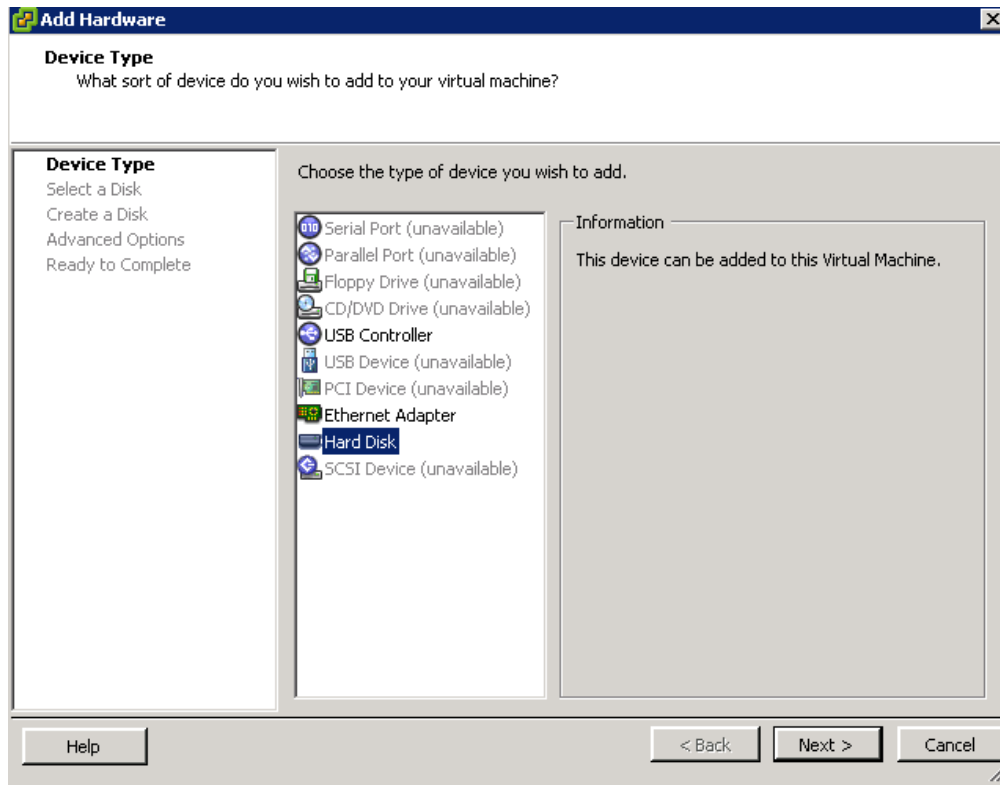
vmfolder:<vm folder> is the virtual machine name. VMDK will be located here

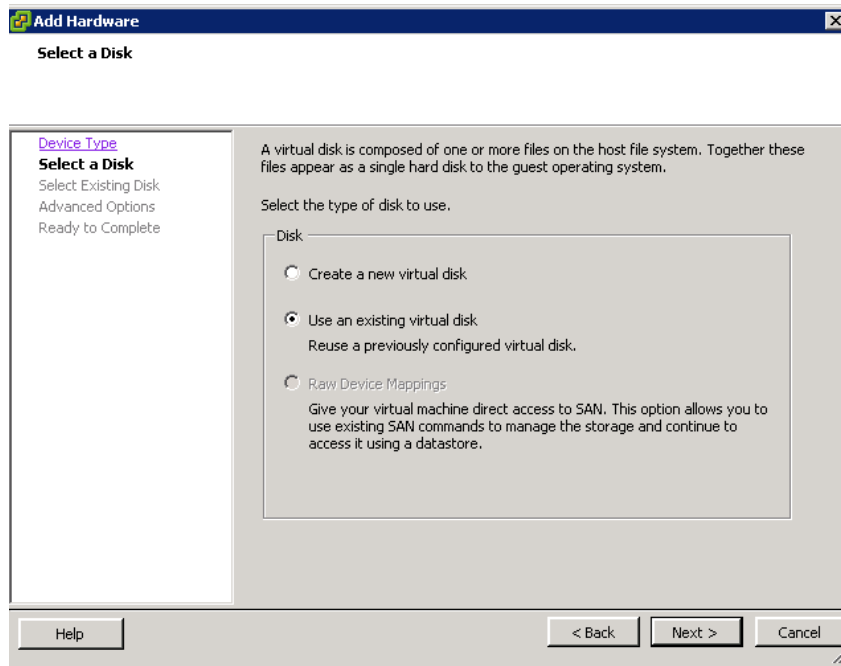
```
usage: mxvmkfstools
-f,--disklocation <disklocation> Folder to save disk to
-h,--help Prints this message
-n,--diskname <diskname> Disk name (without .vmdk or
-flat.vmdk)
-p,--pagesize <pagesize> Page size must be power of 2, larger
than 4 K but no larger than 128 K
-s,--disksize <disksize> Size of disk in bytes
console:~$ mxVmKfsTools.sh -n HPDIA -f unbuntu2 -p 8192 -s 2097152000
```

Example of 8K pagesize: mxVmKfsTools.sh -n HPDIA -f unbuntu2 -p 8192 -s 2097152000

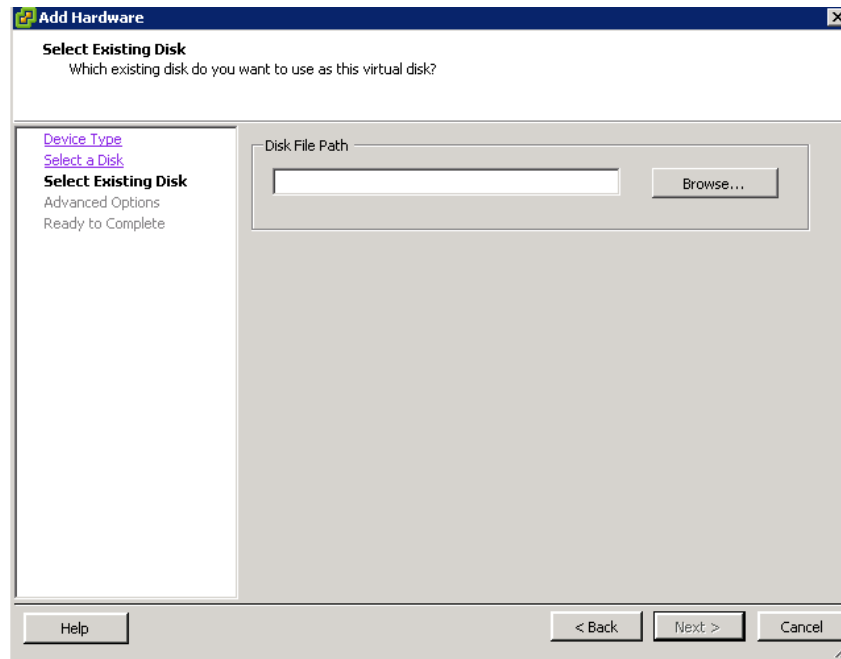
Step 2: Add the vDisk to the virtual machine

1. Connect to the vSphere host which contains the VM to which the user wants to add a new virtual disk.
2. Then, go to the datastore where the user wants to place the new vdisk.
3. Verify if the newly added VMDK is displayed.
4. Add the new vdisk to the virtual machine using the Add Hardware wizard. Follow the wizard as shown.
5. Click Next.

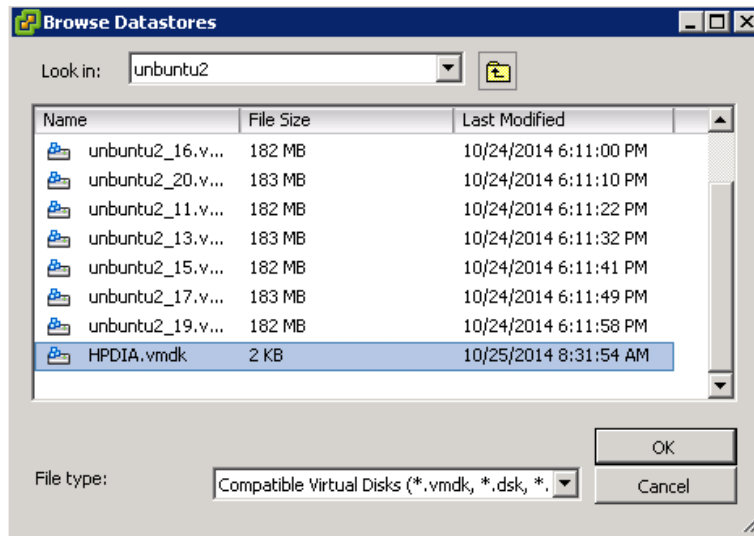




6. Select the option Use an existing virtual disk and click Next.



7. Browse for the disk. Select the Maxta datastore and navigate to the respective virtual machine.



8. Select the newly created VMDK and complete the wizard. The virtual disk is added to the virtual machine.

Maxta Monitoring, and Performance Optimization

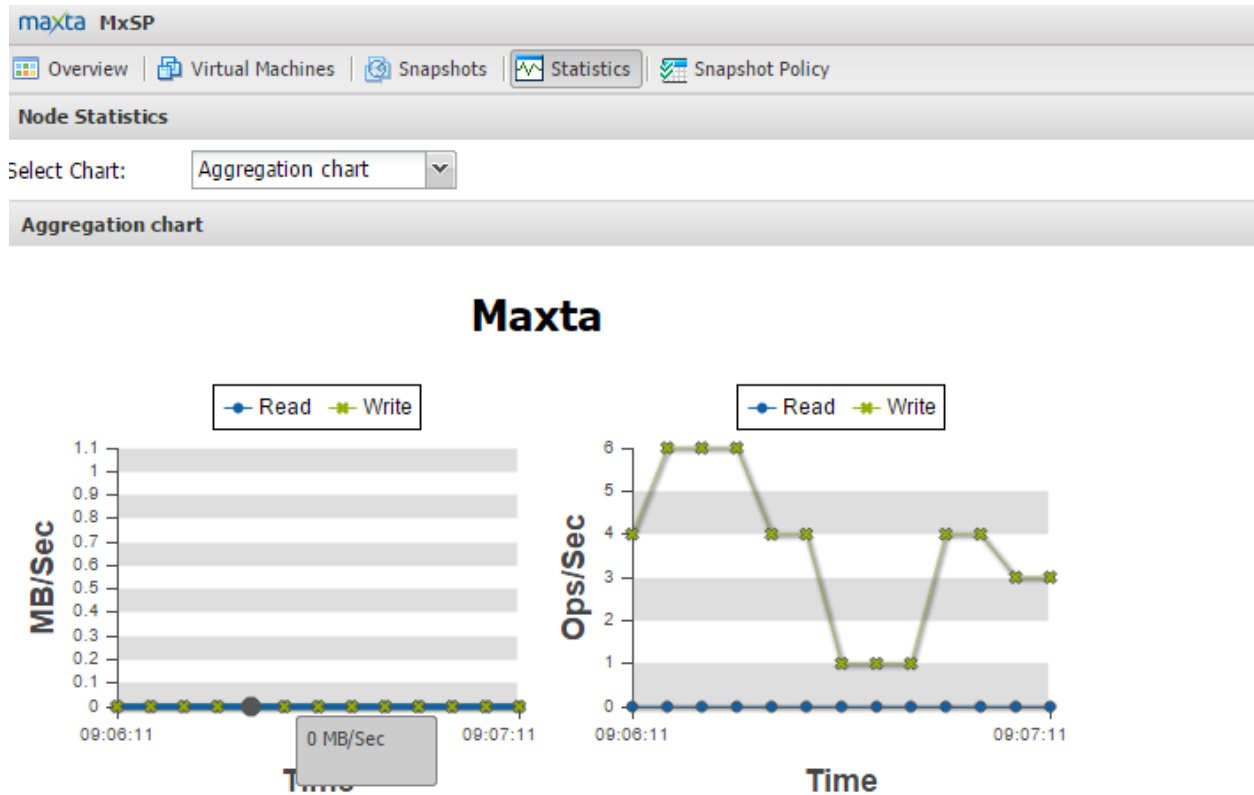
This section provides examples and details for redundancy, monitoring and performance. In this section the reader will see configuration examples with default settings as-well-as explanations with ideas on changing for specific implementations.

Monitoring and Troubleshooting

Maxta Statistics

The Statistics tab provides performance (MB/Sec and Ops/Sec) information for the Maxta datastore, Maxta network, and Disk IO for the selected Maxta virtual machine or aggregated Maxta Cluster. To access the Maxta statistics users can navigate to them under the Maxta VSphere tab or they can be accessed them from the Maxta management IP <http://ManagementIP>.

The Aggregated chart view shows how the entire Maxta Cluster is performing.





The Maxta Storage Platform Virtual Machine view shows us as host by host look at statistics.

Node Statistics


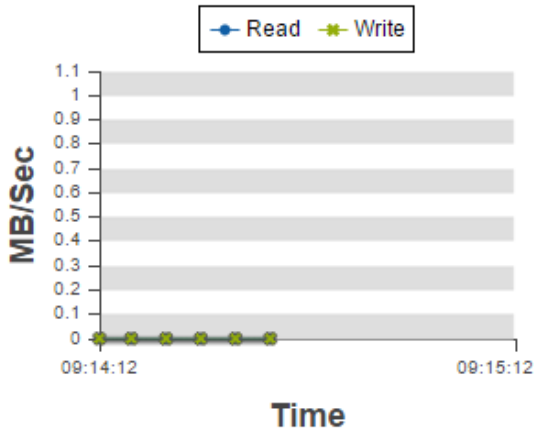
Select Chart: Maxta-Maxta-CL_1

Maxta-Maxta-CL_1

Maxta



Network



Disk IO



Looking at Replicas and Sync State

The Replica Resync section displays the resync status of virtual machines.

Maxta replicates the virtual machines data to a second host in the cluster creating a replica of the virtual machine. If a host fails or the Maxta private network on a host goes down, the data on the failed host is marked as STALE and the data is still available and accessible from the second host. When the host comes back online, it requires a sync with the surviving host. The Replica Resync section displays information of all the virtual machines that are being synced.

This is what a normal sync state looks like and also gives the user a picture of what nodes have a replica of a particular VM. This could be useful in cloning scenarios or adding nodes to the cluster with Maxta Rack Awareness enabled.

Name	Provisioned Capacity	Used Capacity	Replica State (nodes)
Linux1	32.47 GB	2.65 GB	In Sync (Maxta1, M
ubuntu2	32.00 GB	2.65 GB	In Sync (Maxt
Linux3	32.47 GB	2.65 GB	In Sync (Maxta1, Maxta3)
Linux4	32.47 GB	2.65 GB	In Sync (Maxta1, Maxta3)
Linux2	32.47 GB	2.65 GB	In Sync (Maxta1, Maxta3)
ubuntu5	36.00 GB	1.11 GB	In Sync (Maxta1, Maxta3)

The goal is to see In Sync on Maxta Replicas

For testing purposes we powered off Node 3 of the test Maxta cluster which flipped the replica status to stale. When we restored Node 3 after a few minutes the status returned to Sync.

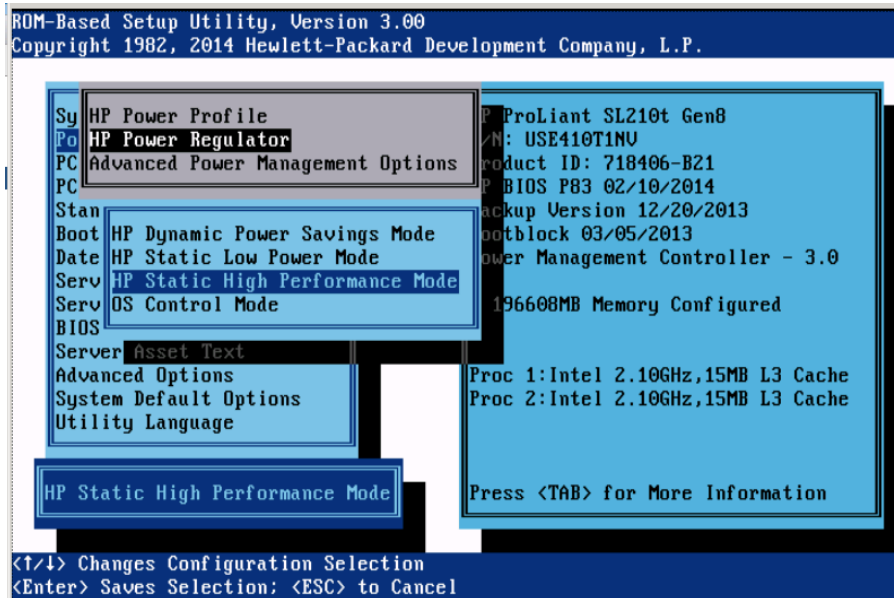
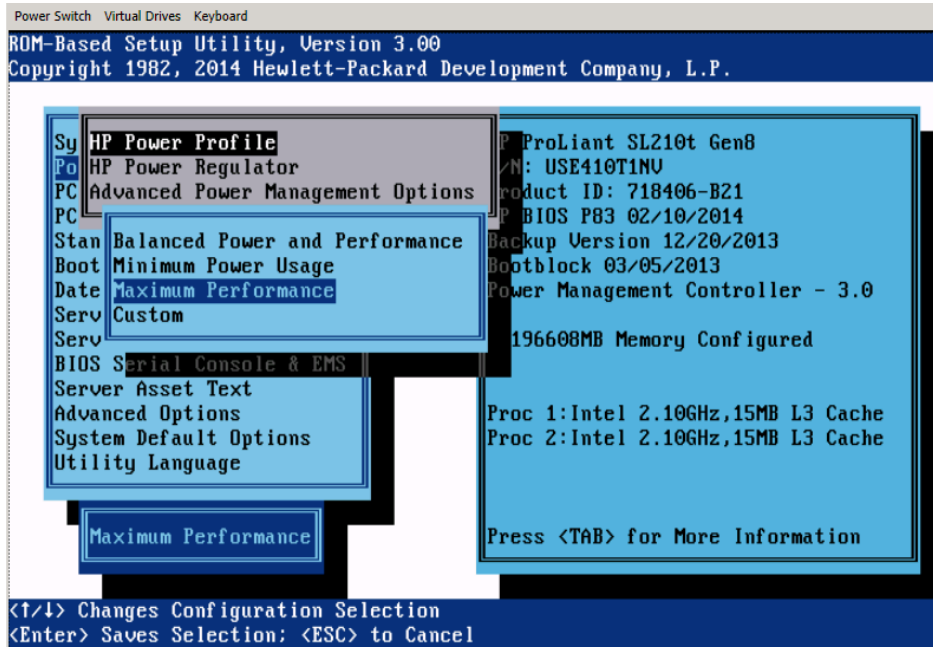
Name	Provisioned Capacity	Used Capacity	Replica State (nodes)
Linux1	32.47 GB	2.65 GB	⚠ Stale (Maxta1, Maxta3)
ubuntu2	32.00 GB	2.65 GB	⚠ Stale (Maxta1, Maxta3)
Linux3	32.47 GB	2.65 GB	In Sync (Maxta1, Maxta3)
Linux4	32.47 GB	2.65 GB	In Sync (Maxta1, Maxta3)
Linux2	32.47 GB	2.65 GB	⚠ Stale (Maxta1, Maxta3)

When a node failure occurs the replica state moves to Stale until the Node is restored

Performance Optimization

SL2500 BIOS Power Optimization Settings

The following capture shows the HP ProLiant SL210t Gen8 server BIOS power setting that would be used to set achieve maximum performance. These setting configured BIOS Power Management for 'Performance', i.e., all power saving features were disabled.



ESXI Performance Optimization Settings

TCP/IP Heap Size

Net.TcplpHeapSize is the size of the memory (in MB) which is allocated up front by the VMkernel to TCP/IP heap. Net.TcplpHeapMax is the maximum amount of memory which can be consumed by TCP/IP as heap. The maximum value for Net.TcplpHeapMax is 512MB in VSphere 5.5. This is a per ESXi host setting and must be done on all hosts.

Vmware ESXi host network heap size in are lab testing was increased from default 256MB to 512MB to maximize performance.

Users will need to SSH into each host and execute the following command “esxcli system settings advanced set -o /Net/TcpipHeapMax -i 512”.

Note: Users will need to reboot for this setting to take effect.

Users can validate this setting using “esxcli system settings advanced list -o /Net/TcpipHeapmax”

```
~ # esxcli system settings advanced list -o /Net/TcpipHeapmax
  Path: /Net/TcpipHeapMax
  Type: integer
  Int Value: 512
  Default Int Value: 256
  Min Value: 32
  Max Value: 512
  String Value:
  Default String Value:
  Valid Characters:
  Description: Max megabytes the tcpip module heap can grow to. (REQUIRES REBOOT!)
~ #
```

Performance Testing

We conducted two simple tests on a Maxta setup we configured.

The Two tests for Maxta performance benchmarking using IOMeter included:


- 100% Read (80% random)
- 70% Read, 30% Write (80% random)

We set all servers to use maximum performance power setting in server BIOS, and increased all hosts network heap size to 512MB.

Each HP ProLiant SL210t Gen8 server node ran a single eight vCPU 32-bit Ubuntu 14.04.01 VM with 16GB VMDK for Ubuntu. We configured 8 * 2GB VMDKs for each guest which were distributed across a Vmware Paravirtual SCSI (PVSCSI) controller.

For each VM, we ran IOMeter with 8 workers. Each worker was configured to work on a single 2GB VMDK. The max disk size per worker was reduced to 2GB because our SSD had much lower capacity than then desired for an optimal design. Outstanding I/Os per worker started at 1 and exponentially increased by 2 up to 16 and all I/O block size and alignment was set to 4K bytes.

The following chart is the summary of the performance. In general, using four HP ProLiant SL210t Gen8 server nodes with Class-C SSD, 100% read test can achieve around 76K IOPs while 70% read can achieve around 24K IOPs. The 70% read test was conducted with a sector size for the disk workers set to 1800 of the total 2GB's of capacity based on the current design of the testing platform and Maxta software release. For maximum Maxta performance, HP customers should choose Class-E SSD along with a larger SSD size. It is also possible to increase the VCPU of each Maxta node to 8 along with increasing the default memory from 8GB to 16GB which should drastically improve the write performance in the current release of Maxta.



Note: We conducted additional testing with performance enhancements outside the scope of the Maxta default parameters which yielded about 130K IOPs for 100% read and 37K IOPs for 70% read.

Conclusion

This deployment guide serves as a solution for deploying a Maxta Storage Platform on the HP SL2500 thus providing a Hyper-Converged solution yielding energy and space savings and requirements in Data Centers. The Maxta Storage Platform is the next generation VM-based storage architectures. After reading this guide and deploying Maxta on the HP ProLiant SL2500, customers benefit from an ideal hyper-converged architecture in terms of cost, space, energy, and compute with systems management.

When compared with the traditional rack-mount form factor, HP ProLiant SL2500 enables administrators to easily scale out a Maxta cluster while providing exceptional performance without compromising compute, storage, and memory requirements.

Appendix

Maxta CLI Installation

The Maxta Software Platform package will need to be extracted into a folder prior to clicking on the Maxta executable file. IP connectivity is needed for the workstation or host that users will run the Maxta executable from and this connectivity needs to extend into the hosts and VMware vCenter that Maxta will be deployed on.

CLI Installation steps:

1. In the installation folder users will find a file called "mxsp_installspec.properties" if the user tried a previous GUI install. If the user has not installed this before they can use the below mxsp_installspec.properties file definition just make sure to name it this.
2. Edit the file and update the following
Replace "****" with the correct password for both vCenter and ESXi
Remove all occurrence of "dataEsx0.rack_id=null", dataEsx1.rack_id=null etc
Replace "ifs_metadev_enabled=False" with "ifs_metadev_enabled=True" for exclusive metadev. Replace "ifs_metadev_shared=False" with "ifs_metadev_shared=True" for shared metadev.
3. Start the installer in CLI mode
4. Open a cmd window
Input the following
<Installer Dir>\jre7\bin\java -cp <Installer Dir>\jre7\lib\jfxrt.jar;<Installer Dir>\mxsp\mxinstall.jar com.maxta.installer.Main -p <Installer Dir>\mxsp_installspec.properties

Note: For settings like compression and deduplication users can modify that behavior by setting the attribute to either True or False in the properties file.

mxsp_installspec.properties file contents

```
#Datacenter(datacenter) in the vCenter
datacenter=Maxta-DC
```

```
#Cluster(cluster) in the Datacenter
cluster=Maxta-CL
```

```
#Virtual Center(vc) IP or name
vc=dia-mvc-01.hpdia.local
```

```
#vCenter username/password
user=root
password=*****
```

```
#Esx to be used creating MxSP nodes
dataEsx0.user=root
dataEsx0.passwd=*****
dataEsx0.rack_id=1
dataEsx1.user=root
dataEsx1.passwd=*****
dataEsx1.rack_id=2
dataEsx2.user=root
dataEsx2.passwd=*****
```

```
dataEsx2.rack_id=3
dataEsx=dia-mesxi-01.hpdia.local,dia-mesxi-02.hpdia.local,dia-mesxi-
03.hpdia.local,

#Maxta network IP
maxta_mgmt_ip=10.16.125.116

#Maxta network mask
maxta_mgmt_mask=255.255.255.0

#Prefix for the private storage IPs.
maxta_net_prefix=169.254.1.0

#Network mask the private storage network.
maxta_net_mask=255.255.255.0

#default gateway
maxta_default_gateway=10.16.125.1
maxta_dns1=10.16.1.15
maxta_dns2=0.0.0.0

#Local Protection
maxta_local_protection=False

#Data Compression
maxta_data_compression=True

#Deduplication
maxta_data_deduplication=False

#Format disks
maxta_format_disks=True
maxta_net_vlan=0
maxta_flash_disk_option=Hybrid

#Maxta LFS Properties
lfs_metadev_enabled=true
lfs_metadev_shared=true
lfs_page_size=4k
```