Deployment Guide: VMware Cloud Foundation on Lenovo ThinkAgile VX

Last update: 22 March 2024
Version 1.7

Describes deployment steps for VMware Cloud Foundation on Lenovo ThinkAgile VX appliances

Includes VMware Tanzu Kubernetes platform for modern applications development

Includes details about hybrid cloud connectivity to Amazon Web Services and Microsoft Azure

Contains Lenovo XClarity integrators for VMware SDDC products

Chandrakandh Mouleeswaran

Cristian Ghetau

LENOVO PRESS

Table of Contents

1 Introduction 1

2 Business problem and business value 2

2.1 Business problem 2

2.2 Business value 2

3 Requirements 4

3.1 Functional requirements 4

3.2 Non-functional requirements 5

4 Architectural overview 7

5 Component model 8

5.1 VMware SDDC Components 8

5.2 VMware vSAN 10

5.3 VMware NSX-T Data Center 13

5.4 Hybrid Clouds 14

5.5 VMware Tanzu Kubernetes Platform 15

5.6 VMware Licensing 19

5.7 HyTrust Security 20

6 Operational model 24

6.2 Edge cluster servers 28

6.3 Management cluster servers 29

6.4 Systems management for Lenovo servers 32

7 Deploying SDDC 38

7.1 VMware Validated Design 38

7.2 VMware Cloud Foundation 38

7.3 Lenovo VX Appliance 39

8 Deployment example 185

8.2 IP/VLAN mapping 187

8.3 Cluster Deployment 188

9 Conclusion 191

Resources 192

Document history 193

1 Introduction

This document describes the reference design of VMware Cloud Foundation (VCF) on Lenovo® ThinkAgile VX servers. VCF on ThinkAgile VX is a way to implement a hybrid cloud solution as a rack based integrated system. This solution is built using ThinkAgile VX hardware from Lenovo, VMware Software Defined Data Center (SDDC) software capabilities and Lenovo XClarity integrators. These three major components come together to give the customers a turnkey hybrid cloud solution with tight integration for ease of management. It provides customers a hyperconverged infrastructure (HCI) solution with automated life cycle management (LCM) capabilities. This document also covers the different components required for implementing an on-premises VMware Cloud Foundation appliance along with a description of various ThinkAgile VX servers available from Lenovo for the customer to pick the right sized solution for their business needs.

The intended audience of this document are IT professionals, technical architects, sales engineers, and consultants to assist in planning, designing, and implementing SDDC products. General understanding of server virtualization, cloud services and VMware software is expected to get the most out of the paper.

This reference design covers the following VMware products:

- vSphere 8.0.2 which provides compute virtualization
- vSAN 8.0, which provides software defined storage (SDS)
- VMware Cloud Foundation 5.1 which automates the entire SDDC system lifecycle and simplifies software operations.
- NSX-T Data Center 4.1.2.1 which provides network virtualization and security by using software defined networking (SDN) and supports private, public, and multi-clouds.
- Aria Suite 8.14, which provides cloud management capabilities for private, public and hybrid clouds with support for multiple hypervisors
- Tanzu Kubernetes Grid 1.5 which provides a container platform to run Kubernetes 1.22 in vSphere to build and deploy modern applications leveraging support from the opensource ecosystem.
- VMware HCX 4.2 which provides infrastructure abstraction and management allowing multi-cloud connectivity and hybrid workflows for Enterprise & Provider Clouds

This document provides an overview of the business problem that is addressed by VCF and embedded SDDC products and the business value that is provided by the SDDC products and Lenovo ThinkAgile VX certified nodes for hybrid cloud and modern applications deployment. A description of customer requirements is followed by an architectural overview of the solution and a description of the logical components. The operational model describes the architecture for deploying into small to medium enterprises. Performance and sizing information is provided with the best practices and networking considerations for implementing SDDC products.

See also the Reference Architecture for VMware vCloud Suite (lenovopress.com/lp0660) which uses network shared storage instead of VMware vSAN.

2 Business problem and business value

This chapter provides a summary of the business problems that this reference design is intended to help address, and the value that this solution can provide.

2.1 Business problem

With rising costs and complexity, it is becoming increasingly harder to manage IT infrastructure in a data center to address private cloud, hybrid cloud and container workloads. As it changes over time, the infrastructure becomes more fragile and more difficult to know the impacts of making changes. Overlaid on the infrastructure issues are the business demands to both reduce costs and at the same time provide a platform to develop more flexible polyglot applications that can meet the business and end-user demands for agility, stability, performance, availability, and easier upgradability.

2.2 Business value

VMware Cloud Foundation (VCF) is a hybrid cloud platform to deploy VMware SDDC for private cloud based on the VMware Validated Design and to integrate with public clouds running VMware SDDC clouds. It provides software defined services for compute, storage, networking, and cloud management to run different workloads. It simplifies installation, upgrade and patch management of SDDC components through lifecycle management either through online or offline.

VCF built on ThinkAgile VX hardware and embedded with VMware SDDC provides all the hardware and software needed for building an enterprise infrastructure platform to support virtualized and containerized workloads that is flexible, easy to manage and easy to change for future needs. By virtualizing compute, storage and networking, SDDC is less dependent on physical hardware. Together with the addition of policy driven configuration, lifecycle management and on demand provisioning, SDDC makes it easier to manage, extend and upgrade the underlying infrastructure to address monolith and microservices architectures. The Lenovo ThinkAgile VX certified nodes and appliances solution for VMware SDDC provides businesses with an affordable, interoperable, and reliable industry-leading cloud solution to manage all of their virtualized and containerized workloads.

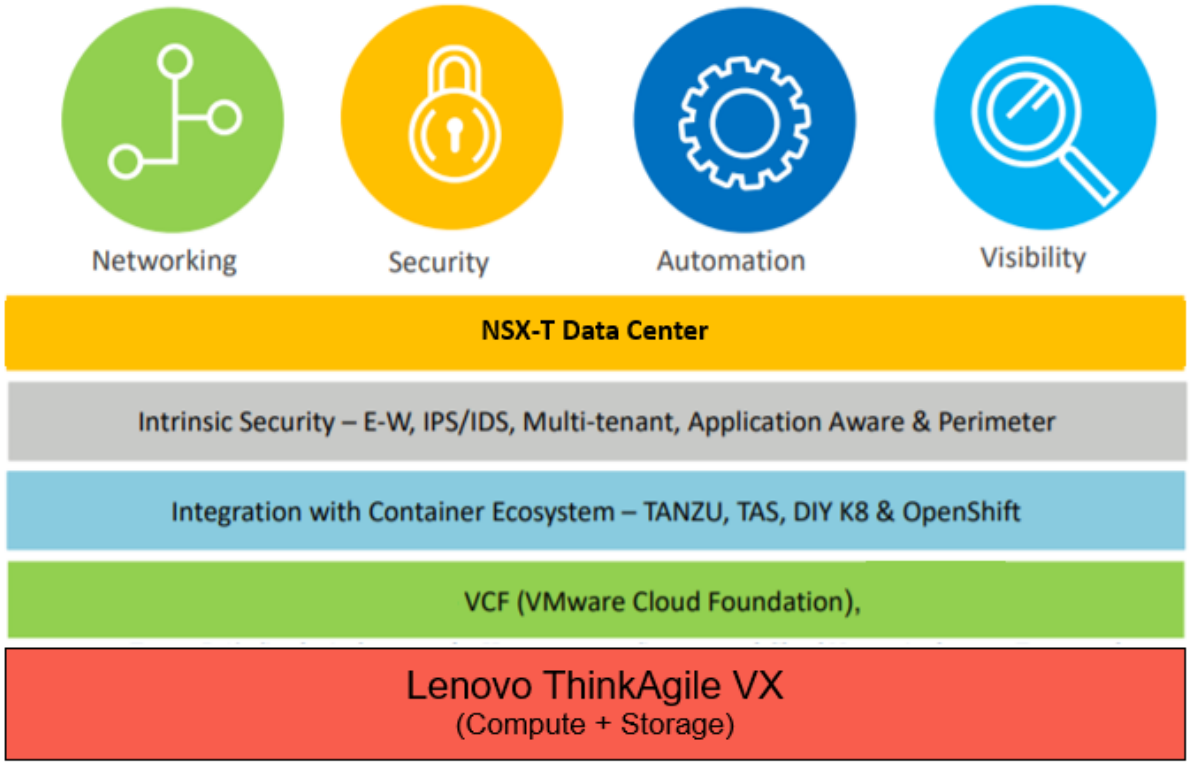

Figure 1: Lenovo ThinkAgile VX for VMware SDDC

3 Requirements

This chapter describes the functional and non-functional requirements for this reference design.

3.1 Functional requirements

The following section describes the functional requirements that are needed for typical multi cloud deployments. Figure 2 shows a simplified use-case model for hybrid cloud deployments.

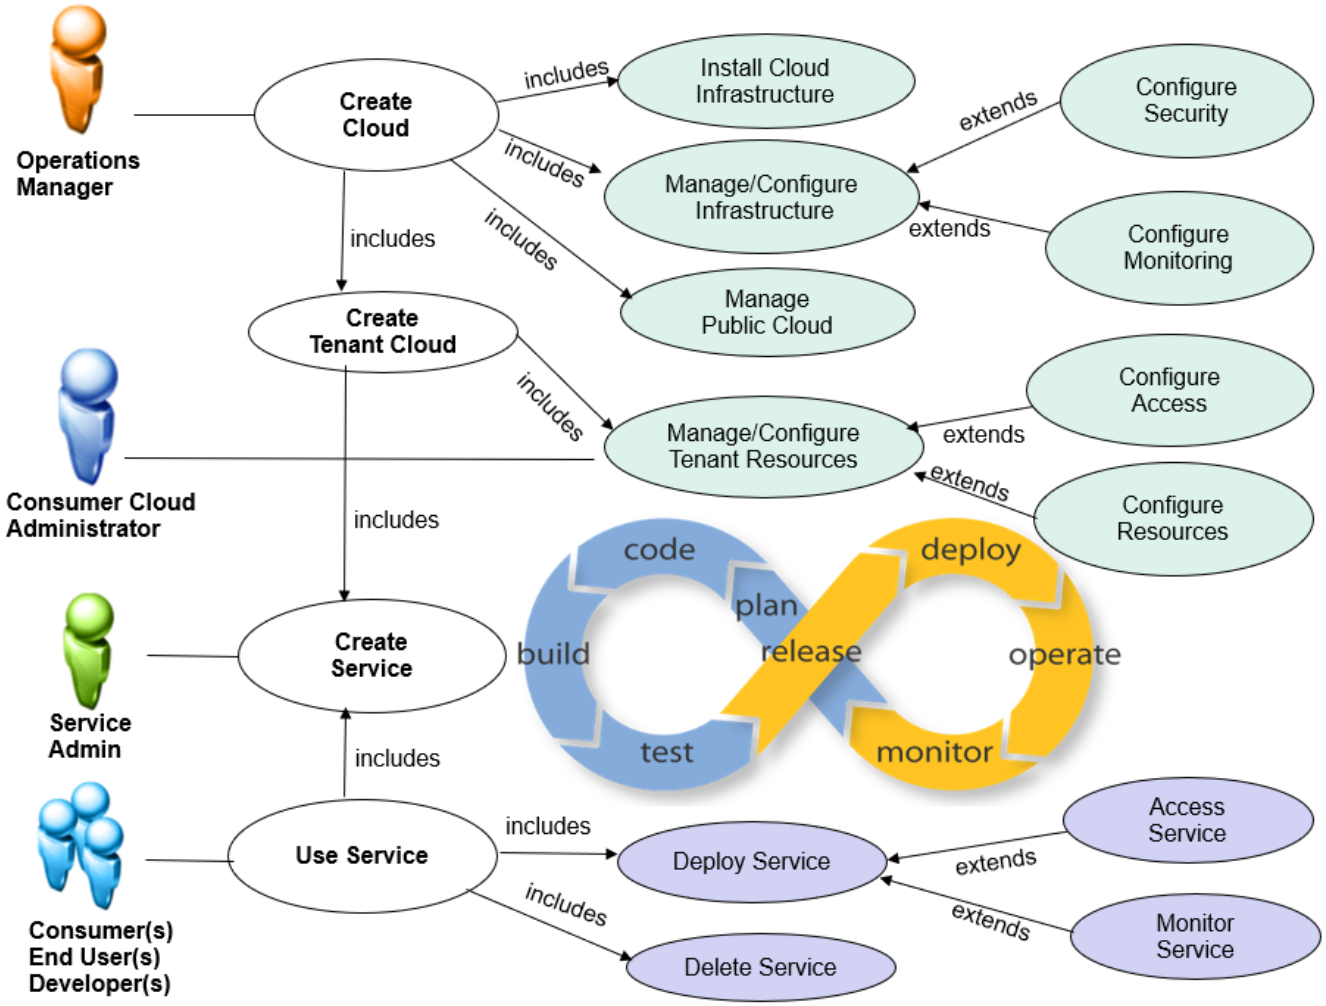

Figure 2: Use case model

Table 1 lists the functional requirements.

Table 1: Functional requirements

Requirement name	Description
Virtualization	Solution supports compute, storage, and network virtualization
Containerization	Mange and deploy containers on the virtualized infrastructure
Monitoring, event and capacity management	Monitors the health of the cloud infrastructure, collection and management of exception events, and capacity planning
Self-service automation	Solution provides on boarding, provisioning, and management of services and VMs from a service catalog
Approval and workflow	Provides the capability to approve, modify, deny, and delegate service requests
Cloud administration	Provides capabilities to administer a cloud environment, such as adding storage or computational resources in the cloud pool or defining new segregated networks
Image management	Provides capabilities to create VMs and containers, establish version control, search for and compare images, and delete images from the virtual images templates repositories
Service management	Provides capabilities to create services, establish version control, search for services, and delete services from the service templates catalog repositories
Access and authorization Controls	Provides the capabilities to create users and groups and to establish authorization to certain features in the cloud, such as tenant cloud administration, service developer, and user service requester
Virtual Machine Migration	Migrate applications, virtual machine and templates between private and public clouds.
Migrate Security Policies	Migrate network and security policies such as firewall rules to public cloud and vice versa,
Network Extension	Retain virtual machines network properties (L2 and L3) across clouds.
Catalog Management	Maintain common catalog for templates across clouds.
Hybrid Cloud Integration	Supports connectivity to seamlessly migrate and manage workloads in multi cloud environments.
Opensource ecosystem	Supports integration and flexibility to leverage open-source software in the platform.
DevSecOps	An advanced approach to security that simplifies and automates container operations across multi-clouds.

3.2 Non-functional requirements

Table 2 lists the non-functional requirements that are needed for typical cloud deployments.

Table 2: Non-functional requirements

Requirement name	Description
Backup/Recovery	Solution support for integrated backup
Ease of installation	Reduced complexity for solution deployment
Ease of management/operations	Simple management of infrastructure and cloud software
Supportability	Available vendor support

Requirement name	Description
Scalability	Solution components scale with increase in number of concurrent users, VMs/services provisioned per minute or per hour
Flexibility	Solution supports variable deployment methodologies
Security	Solution provides ways to secure customer data
Reliability, availability, and serviceability (RAS)	High availability and resiliency of cloud management and managed infrastructure

4 Architectural overview

This chapter gives an architectural overview of SDDC products. Figure 3 gives an overview of how those products are deployed into management, edge and compute and additional compute clusters and seamlessly integrated with different public clouds. This separation of function into these clusters allows for scaling in larger environments.

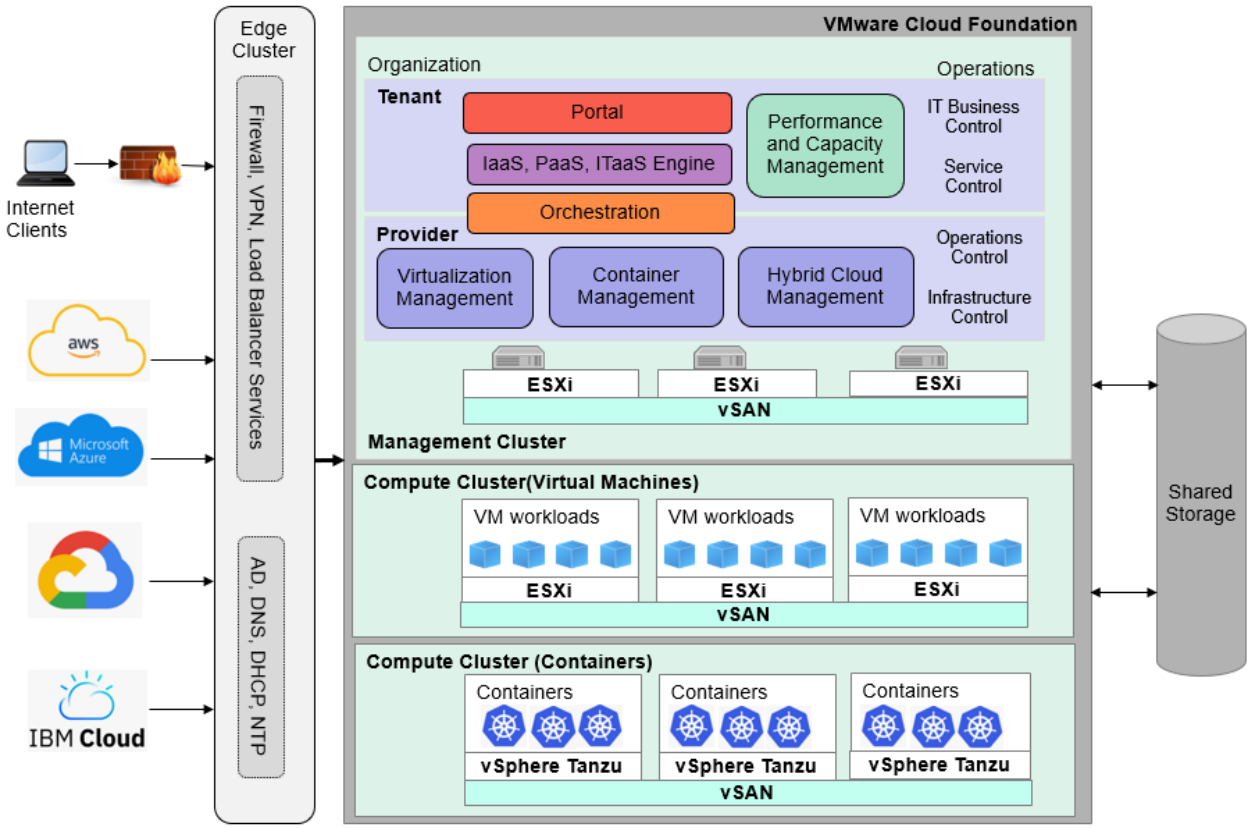

Figure 3: Conceptual design of a SDDC environment

The management cluster runs the components required to support SDDC and is used for management of virtualization and container platforms, public cloud management, monitoring, and infrastructure services. A management cluster provides resource isolation which helps these services to operate at their best possible performance level.

Dedicated edge cluster required for large environments and for small medium deployments, the edge services can coexist in either management or compute clusters. Edge provides protected capacity by which internal data center networks connect via gateways to external networks. Networking edge services and network traffic management occur in this cluster and all external facing network connectivity ends in this cluster. The shared edge and compute cluster also supports the delivery of all other (non-edge) customer workloads and there can be one or more compute clusters, depending on the customer environment. Multiple compute clusters can be for different organizations or tenants, different workload types, or to spread the load in a large enterprise.

5 Component model

This chapter describes the component model for VMware SDDC and optionally extending it into public clouds with hybrid cloud connections. Lastly the HyTrust suite of software is described which provides additional security protection features.

5.1 VMware SDDC Components

Figure 4 shows an overview of the major components of the VMware SDDC.

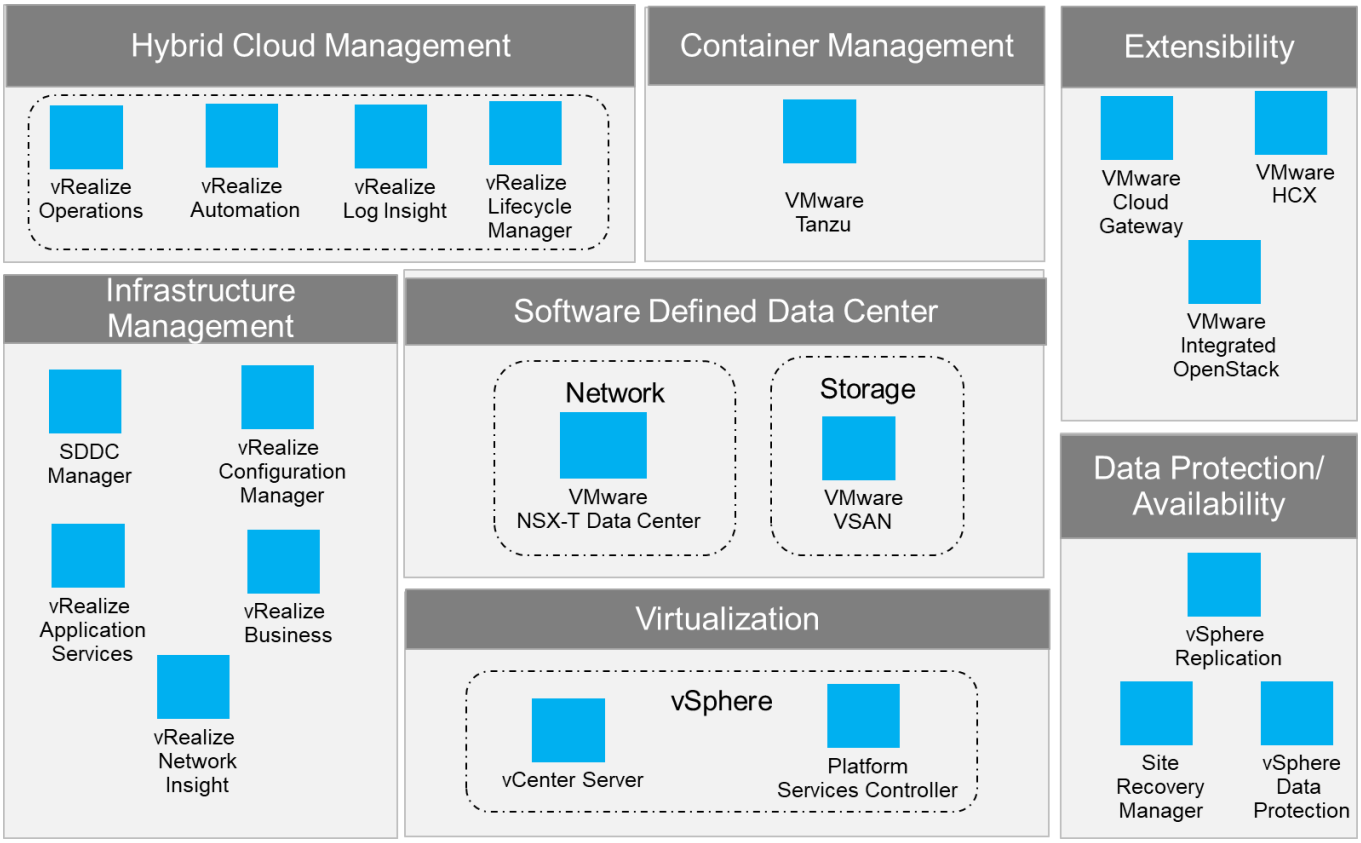

Figure 4: SDDC components

The VMware SDDC features the following components:

ESXi hypervisor	Provides bare-metal virtualization of servers so you can consolidate your applications on less hardware.
vCenter Server	Provides a centralized platform for managing vSphere environments and includes vSphere replication and vSphere data protection.
Platform Services Controller (PSC)	Provides a set of common infrastructure services that encompasses single sign-on (SSO), licensing, and a certificate authority (CA).

VMware Cloud Foundation (VCF)	Suite of components to deploy and manage your software-defined data center (SDDC)
SDDC Manager	Provides management interface to VCF. It performs deployment of ESXi, Aria Suite, NSX-T and lifecycle management operations.
Cloud Builder	Used to deploy and configure the first cluster of the management domain and transfer inventory and control to SDDC Manager
Aria Suite Lifecycle Manager	Provides deployment options such as install, configure, import, and upgrade Aria Suite environments and perform drift analysis and view the health of those environments
Aria Automation	Provides a self-service, policy-enabled IT and application services catalog for deploying and provisioning of business-relevant cloud services across private and public clouds, physical infrastructure, hypervisors, and public cloud providers.
Aria Operations	Provides a set of components for automation of operations including infrastructure health, configurations and compliance, application discovery, and monitoring of hardware and software.
• Aria Operations Manager	Provides comprehensive visibility and insights into the performance, capacity and health of your infrastructure.
• Aria Automation Config	Provides automation of configuration and compliance management across your virtual, physical, and cloud environments, which assesses them for operational and security compliance.
Aria Business for Cloud	Provides transparency and control over the costs and quality of IT services that are critical for private (vCloud Suite) or hybrid cloud (Aria Suite) success.
Aria Operations for Logs	Provides analytics capabilities to unstructured data and log management, which gives operational intelligence and deep, enterprise-wide visibility across all tiers of the IT infrastructure and applications. Standard for Aria Suite.
Aria Operations for Networks	Provides end-to-end management and helps you gain visibility for NSX, VMware SD-WAN, VMware Cloud on AWS, Tanzu Kubernetes Grid
vCenter Site Recovery Manager (SRM)	Provides disaster recovery capability with which you can perform automated orchestration and non-disruptive testing for virtualized applications by using ESXi hypervisor only. SRM is standard for vCloud Suite and optional for Aria Suite.
NSX-T Datacenter	NSX provides virtualization of networking in software and is part of VMware's vision of the SDDC. For more information, see "VMware NSX" on page 13.

VMware Hybrid Cloud Extension (HCX)	Provides Hybrid cloud connectivity between on premise VMware cloud and public cloud SDDC
Tanzu	Kubernetes based container platform for vSphere and supports development of modern applications.

The SDDC products also have dependencies on the following external components:

Identity source	Identity sources (Active Directory, OpenLDAP, or Local OS) or similar is required to implement and operate the vCloud Suite or Aria Suite infrastructure.
DNS	DNS must be configured for connectivity between vCenter Server, Active Directory, ESXi hosts, and the VMs
DHCP/TFTP	PXE boot is required for vSphere Auto Deploy functionality.
Time synchronization	Accurate time keeping and time synchronization is critical for a healthy infrastructure. All components (including ESXi hosts, vCenter Server, the SAN, physical network infrastructure, and VM guest operating systems) must have accurate time keeping.
Microsoft SQL Server database	Many of the SDDC components come with embedded PostgreSQL database or they can use Microsoft SQL Server as external database depending on the component and the intended environment.

Other software components such as Lenovo XClarity Administrator are not shown. As well as providing management of Lenovo hardware, XClarity Administrator also has plugins for VMware vCenter, VMware Aria Orchestrator, and VMware Aria Operations for Logs which are further described in “Systems management for Lenovo servers” on page 32.

5.2 VMware vSAN

VMware vSAN is a Software Defined Storage (SDS) solution embedded in the ESXi hypervisor and provides flexible configurations with mix of SSD, NVMe and HDDs. VMware vSAN All Flash pools flash devices for caching and capacity tiers and vSAN Hybrid uses flash for cache and magnetic disks for capacity across three or more 10 GbE connected servers into a single shared datastore that is resilient and simple to manage.

VMware vSAN can be scaled to 64 servers, with each server supporting up to five disk groups, with each disk group consisting of a one solid-state drives (SSDs) or NVMe drives for cache and up to seven SSDs or hard disk drives (HDDs) for capacity. Performance and capacity can be easily increased by adding components, such as disks, disk groups, flash devices, or servers.

The flash cache is used to accelerate reads and writes. Frequently read data is kept in read cache; writes are coalesced in cache and destaged to disk efficiently, which greatly improves application performance. vSAN All Flash uses cache for write back cache only and reads happens through capacity drives.

VMware vSAN manages data in the form of flexible data containers that are called *objects*. The following types of objects for VMs are available:

- VM Home
- VM swap (`.vswp`)
- VMDK (`.vmdk`)
- Snapshots (`.vmsn`)

Internally, VM objects are split into multiple components that are based on performance and availability requirements that are defined in the VM storage profile. These components are distributed across multiple hosts in a cluster to tolerate simultaneous failures and meet performance requirements. VMware vSAN uses a distributed RAID architecture to distribute data across the cluster. Components are distributed with the use of the following two storage policies:

- Number of stripes per object. It uses RAID 0 method.
- Number of failures to tolerate. It uses either RAID-1 or RAID-5/6 method. RAID-5/6 is currently supported for an all flash configuration only.

VMware vSAN uses the Storage Policy-based Management (SPBM) function in vSphere to enable policy driven VM provisioning, and uses vSphere APIs for Storage Awareness (VASA) to make available vSAN storage capabilities to vCenter. This approach means that storage resources are dynamically provisioned based on requested policy, and not pre-allocated as with many traditional storage solutions. Storage services are precisely aligned to VM boundaries; change the policy, and vSAN implements the changes for the selected VMs. Table 3 lists the vSAN storage policies.

Table 3: vSAN storage policies

Storage Policy	Description	Default	Maximum
Failure Tolerance Method	Defines a method used to tolerate failures. RAID-1 uses mirroring and RAID 5/6 uses parity blocks (erasure encoding) to provide space efficiency. RAID-5/6 is supported only for All Flash configurations. RAID 5 requires minimum 4 hosts and RAID 6 requires minimum 6 hosts. When RAID 5/6 is chosen, RAID 5 is used when FTT=1 and RAID 6 is used when FTT=2.	RAID-1	N/A
Primary level of failures to tolerate	Defines the number of host, disk, or network failures a VM object can tolerate. For *n* failures tolerated, n+1 copies of the VM object are created and 2n+1 hosts with storage are required. For example with a FTT=1, RAID-1 uses 2x the storage and RAID-5/6 uses 1.33x the storage. When FTT=2, RAID-1 uses 3x the storage and RAID-5/6 uses 1.5x the storage.	1	3

Secondary level of failures to tolerate	Works only for stretched clusters and defines the number of disk or host failures a storage object can tolerate for each of the sites. A storage object with the primary level of failures “m” and secondary level of failures “n” can tolerate “n” host or disk failures in addition to “m” site failures. Supported values are 0 to 3 depending on the fault tolerance method (erasure coding can tolerate up to 2 failures). For each of the sites the number of required hosts in order to tolerate “n” failures is “2n+1” for mirroring and 4 or 6 for erasure coding(failures would be 1 or 2 respectively)	0	3
Number of disk stripes per object	The number of HDDs across which each replica of a VM object is striped. A value higher than 1 might result in better performance, but can result in higher use of resources.	1	12
Object space reservation	Percentage of the logical size of the object that should be reserved (or thick provisioned) during VM creation. The rest of the storage object is thin provisioned. If your disk is thick provisioned, 100% is reserved automatically. When deduplication and compression is enabled, this should be set to either 0% (do not apply) or 100%.	0%	100%
Flash read cache reservation	SSD capacity reserved as read cache for the VM object. Specified as a percentage of the logical size of the object. Should be used only to address read performance issues. Reserved flash capacity cannot be used by other objects. Unreserved flash is shared fairly among all objects.	0%	100%
Force provisioning	If the option is set to Yes, the object is provisioned, even if the storage policy cannot be satisfied by the data store. Use this parameter in bootstrapping scenarios and during an outage when standard provisioning is no longer possible. The default of No is acceptable for most production environments.	No	N/A
IOPS limit for object	Defines IOPS limit for a disk and assumes a default block size of 32 KB. Read, write and cache operations are all considered equivalent. When the IOPS exceeds the limit, then IO is throttled.	0	User Defined
Disable object checksum	Detects corruption caused by hardware/software components including memory, drives, etc. during the read or write operations. Object checksums carry a small disk IO, memory and compute overhead and can be disabled on a per object basis.	No	Yes
Data locality	Specify the data location. Either the preferred fault domain or Non-preferred fault domain in a stretched cluster, or set to Host local to pin the VMs folder and VMDKs to the host it was created on. This policy is only valid for objects with the primary level of failures to tolerate = 0. Default value: None	None	N/A

5.3 VMware NSX-T Data Center

VMware NSX-T™ Data Center is an SDN solution that allows the creation of overlay networks with the same capabilities that are available in the physical network. Clients can build multi-tier application networks and implement micro-segmentation to mitigate against threats that penetrate through the perimeter firewall. VMware NSX can be used with VMware vSphere hypervisor and also with several other hypervisors.

When deployed, VMware NSX-T is a collection of virtual machines that work collectively to support the overlay network. These components are distributed across multiple hosts or clusters and can tolerate simultaneous failures while providing optimal performance. Table 4 lists the NSX-T components.

Table 4: NSX-T Components

Component	Description
NSX Manager	management plane for the NSX-T Data Center and provides configuration and orchestration of logical switching and routing, edge services, security services and distributed firewall.
NSX Policy Manager	Provides policy-based access to NSX-T Data center services.
Cloud Service Manager	Manages all public cloud NSX-T environment communications.
NSX Controller	Distributed state management system that controls virtual networks and overlay transport tunnels.
Transport Node	The ESXi hosts are transport nodes and the communication happens through one or more VTEP endpoints on the hosts.
Virtual Tunnel Endpoint (VTEP)	VMkernel interface that is created by the NSX-T manager during the initial preparation of the ESXi Host to participate in the overlay network.
Edge Services Gateway	The Edge Services Gateway gives you access to all NSX Edge services, such as firewall, NAT, DHCP, VPN, load balancing, and high availability. Each Edge Services Gateway can be configured for single or multiple services and have a total of 10 uplink and internal network interfaces. The internal interfaces connect to secured port groups and act as the gateway for all protected virtual machines in the port group.
Logical Switch	Provides a representation of Layer 2 switched connectivity across many hosts with Layer 3 IP reachability between them. It used to isolate tenants from each other.
Distributed Router	East-West routing and it is handled by transport nodes.
Service Router	Edge nodes serve stateful centralized services NAT, DHCP server, VPN, Gateway Firewall, Bridging, Service Interface, Metadata Proxy for OpenStack. Provides north-south routing.
Physical Router	A physical router that is logically connected to each ESXi host in the data center.
Two-Tier routing	Multi-tier routing can be design using DR, SR and physical routers across gateways

Component	Description
Virtual Routing Forwarding (VRF)	virtualization method that consists of creating multiple logical routing instances within a physical routing appliance. It provides a complete control plane isolation between routing instances.
Distributed Firewall (DFW)	provides stateful protection of the workload at the vNIC level and enforcement occurs in the hypervisor kernel, helping deliver micro-segmentation
Load Balancer	Provides Layer 4 and Layer 7 load balancing features

Table 5 lists the various logical networks in which these components are deployed.

Table 5: NSX Component Logical Networks

Logical Network	NSX Component/Service
Management Plane	NSX Manager, Policy Manager, Cloud Service Manager
Control Plane	NSX Controllers
Data Plane	NSX VIBs, NSX Edge, NSX Firewall, NSX Logical (Distributed) Router, Transport Zones

Figure 5 shows the standard set of icons that are defined by VMware to represent the various NSX-T components.

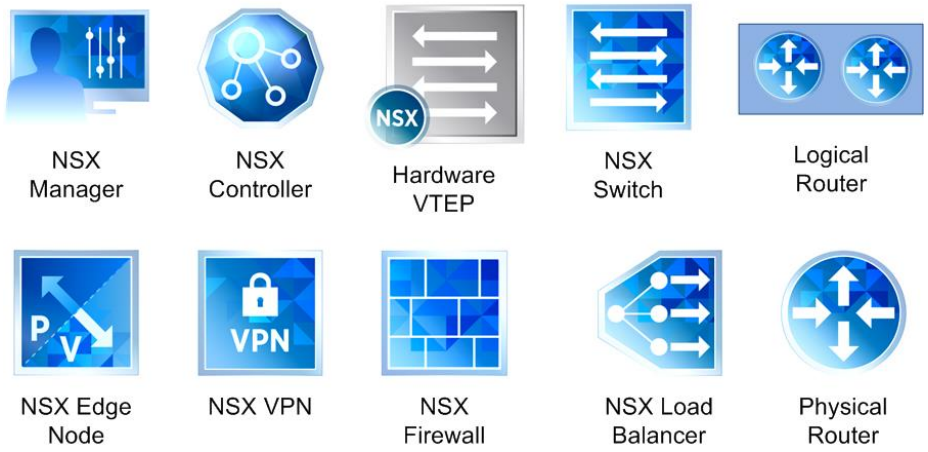

Figure 5: NSX-T Standardized Icons

5.4 Hybrid Clouds

VMware SDDC can run either in on-premises or on any other public clouds such as Amazon Web Services (AWS), Microsoft Azure, IBM Cloud and Google Cloud Platform. VMware Aria can manage workloads across clouds and workloads can be seamlessly provisioned and migrated across different SDDC environments.

5.4.1 VMware Hybrid Cloud Extension (HCX)

Enables on-premises SDDC workloads to migrate and rebalance to different public clouds running VMware Cloud. The migration can be done live or batch or scheduled and Aria Operations for Networks helps to monitor the migration. NSX Hybrid Connect can be used to migrate virtual machines between two on-premises VMware

SDDC cloud. HCX supports various features for proxy and WAN optimization to improve throughput and do migration at scale.

Table 6: VMware Hybrid Cloud Extension support

Source Cloud	Components	Target Cloud
On Premise VCF	HCX	On Premise VCF
On Premise VCF	HCX, NSX Hybrid Connect	VMware Cloud on Amazon Web Services (AWS)
On Premise VCF	VMware NSX® Advanced or Enterprise through IBM Cloud	IBM Cloud for VMware Solutions
On Premise VCF	HCX Connector, HCX Cloud Manager Appliance	Google Cloud VMware Engine,
On Premise VCF	VMware HCX Connector, Azure VMware Solution HCX Cloud Manager	Azure VMware Solution

5.5 VMware Tanzu Kubernetes Platform

A Tanzu Kubernetes Cluster is a full distribution of the open-source Kubernetes container orchestration platform that is built, signed, and supported by VMware. vSphere with Tanzu offers a VM Service functionality that enables DevOps engineers to deploy and run VMs, in addition to containers, in a common, shared Kubernetes environment. By using vSphere with Tanzu the vSphere Administrator can turn a vSphere cluster to a platform for running Kubernetes workloads in dedicated resource pools. The vSphere administrator can manage and monitor vSphere Pods, VMs, and Tanzu Kubernetes clusters by using the vSphere Client.

5.5.1 vSphere with Tanzu

Both, containers and VMs, share the same vSphere Namespace resources and can be managed through a single vSphere with Tanzu interface. The VM Service addresses the needs of DevOps teams that use Kubernetes but have existing VM-based workloads that cannot be easily containerized. It also helps users reduce the overhead of managing a non-Kubernetes platform alongside a container platform. When running containers and VMs on a Kubernetes platform, DevOps teams can consolidate their workload footprint to just one platform. **Error! Reference source not found.** below shows the virtualization and container components in vSphere with Tanzu architecture.

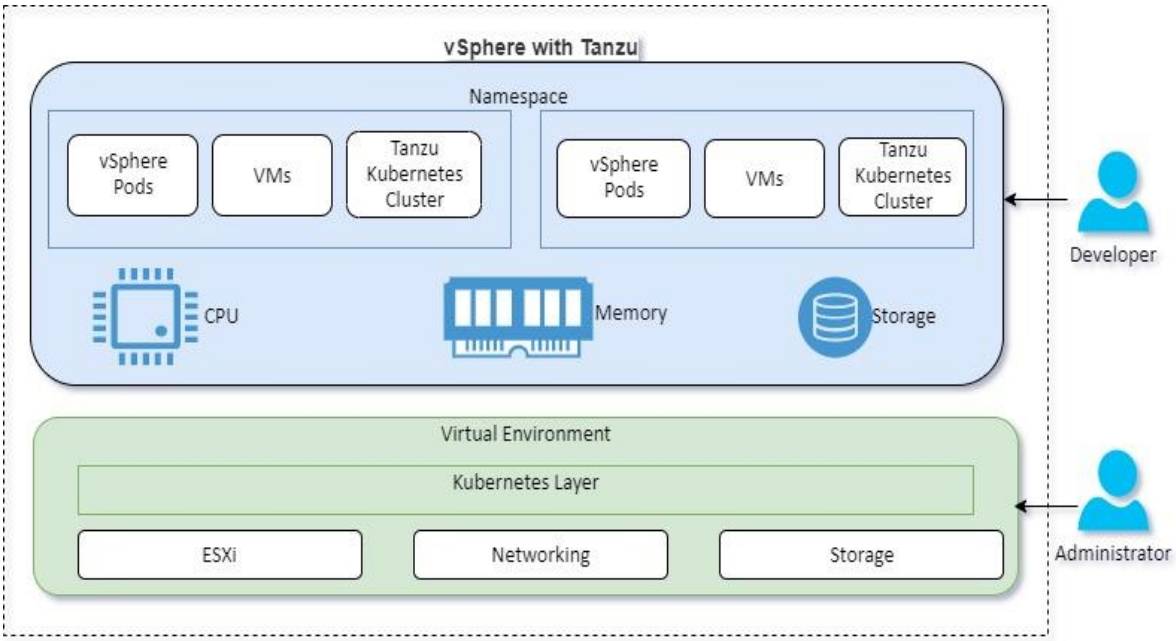

Figure 6: vSphere with Tanzu

The Table 7 shows the list of core components in Tanzu

Table 7: vSphere Tanzu Components

Component	Description
Supervisor Cluster	A cluster that is enabled for vSphere with Tanzu is called a Supervisor Cluster. It runs on top of an SDDC layer that consists of ESXi for compute, NSX-T Data Center or vSphere networking, and vSAN or another shared storage solution
Kubernetes control plane VM	Three Kubernetes control plane VMs in total are created on the hosts that are part of the Supervisor Cluster and three control plane VMs are load balanced
Tanzu Kubernetes Grid Service	Kubernetes control plane runs directly on the hypervisor layer
vSphere Pod	A vSphere Pod is a VM with a small footprint that runs one or more Linux containers. It is equivalent to Kubernetes pod. vSphere Pods are Open Container Initiative (OCI) compatible and can run containers from any operating system as long as these containers are also OCI compatible
vSphere Namespaces	Provides shared resource pools to run containers to isolate applications and tenants. The vSphere administrator can set limits for CPU, memory, storage, as well as the number of Kubernetes objects that can run within the namespace. The vSphere administrator can set limits for CPU, memory, storage, as well as the number of Kubernetes objects that can run within the namespace
Tanzu Kubernetes Cluster	Kubernetes clusters created by Tanzu Kubernetes Grid Service to run workloads

Component	Description
Spherelet	An additional process called Spherelet is created on each host. It is a kubelet that is ported natively to ESXi and allows the ESXi host to become part of the Kubernetes cluster
Container Runtime Executive (CRX)	CRX includes a para-virtualized Linux kernel that works together with the hypervisor. CRX uses the same hardware virtualization techniques as VMs and it has a VM boundary around it. A direct boot technique is used, which allows the Linux guest of CRX to initiate the main init process without passing through kernel initialization. This allows vSphere Pods to boot nearly as fast as containers
Tanzu CLI	A command line interface to access and run commands to manage Kubernetes cluster and containers.

5.5.2 Tanzu Kubernetes Shared Services

The Tanzu Kubernetes platform leverages SDDC components and many opensource components to manage and provide operational services for containers running on Kubernetes cluster. The plugins can be chosen based on the compatibility preferences to use across multi cloud environments. The VCF platform hides all complexities in configuring these components and provides GUI based one click deployment for Tanzu.

Table 8 shows the list of shared services components supported on Tanzu Kubernetes platform.

Table 8: Tanzu Kubernetes Core Services

Services	Description
Infrastructure platform	vSphere 6.7U3, vSphere 7.x, VMware Cloud on AWS, Azure VMware Solution
Cluster Lifecycle Management	Core Cluster API (v0.3.14), Cluster API Provider vSphere (v0.7.6)
Kubernetes node OS distributed with TKG	Photon OS 3, Ubuntu 20.04
Bring your own image	Photon OS 3, Red Hat Enterprise Linux 7, Ubuntu 18.04, Ubuntu 20.04
Container runtime	Containerd (v1.4.3)
Container networking	Antrea (v0.11.3), Calico (v3.11.3)
Container registry	Harbor (v2.1.3)
Ingress	NSX Advanced Load Balancer Enterprise (v20.1.3), Contour (v1.12.0)
Load Balancing	NSX Advanced Load Balancer Essentials, HA Proxy
Storage	vSphere Container Storage Interface (v2.1.0) and vSphere Cloud Native Storage

Services	Description
Infrastructure platform	vSphere 6.7U3, vSphere 7.x, VMware Cloud on AWS, Azure VMware Solution
Authentication	LDAP or OIDC via Pinniped (v0.4.1) and Dex
Observability and Monitoring	Fluent Bit (v1.6.9), Prometheus (v2.18.1), Grafana (v7.3.5), Tanzu Mission Control*
Backup and migration	Velero (v1.5.3)
Service Mesh	VMware Tanzu Service Mesh*
Policy and Management	Tanzu Mission Control*
Image Build	Tanzu Build Service*
Data Flow	Spring Cloud Data Flow*
Database	Tanzu Data Service*
Image Catalog	Tanzu Application Catalog*
API Gateway	Spring Cloud Gateway*

*These services are part of VMware Tanzu Advanced Edition which is not included in VCF editions.

5.5.3 Tanzu Mission Control

Tanzu Mission Control provides centralized management and operations for multi cloud Kubernetes deployments which enables developers to work seamlessly across different environment without compromising security and governance. Tanzu Mission Control also well integrated with Tanzu observability and Service Mesh and provides cluster lifecycle management, data protection, policy management and centralized authentication and authorization capabilities which enables operators and infrastructure teams to manage efficiently. Tanzu Mission Control is offered through VMware Cloud Services which provides also Tanzu Application Services and Data services as subscriptions.

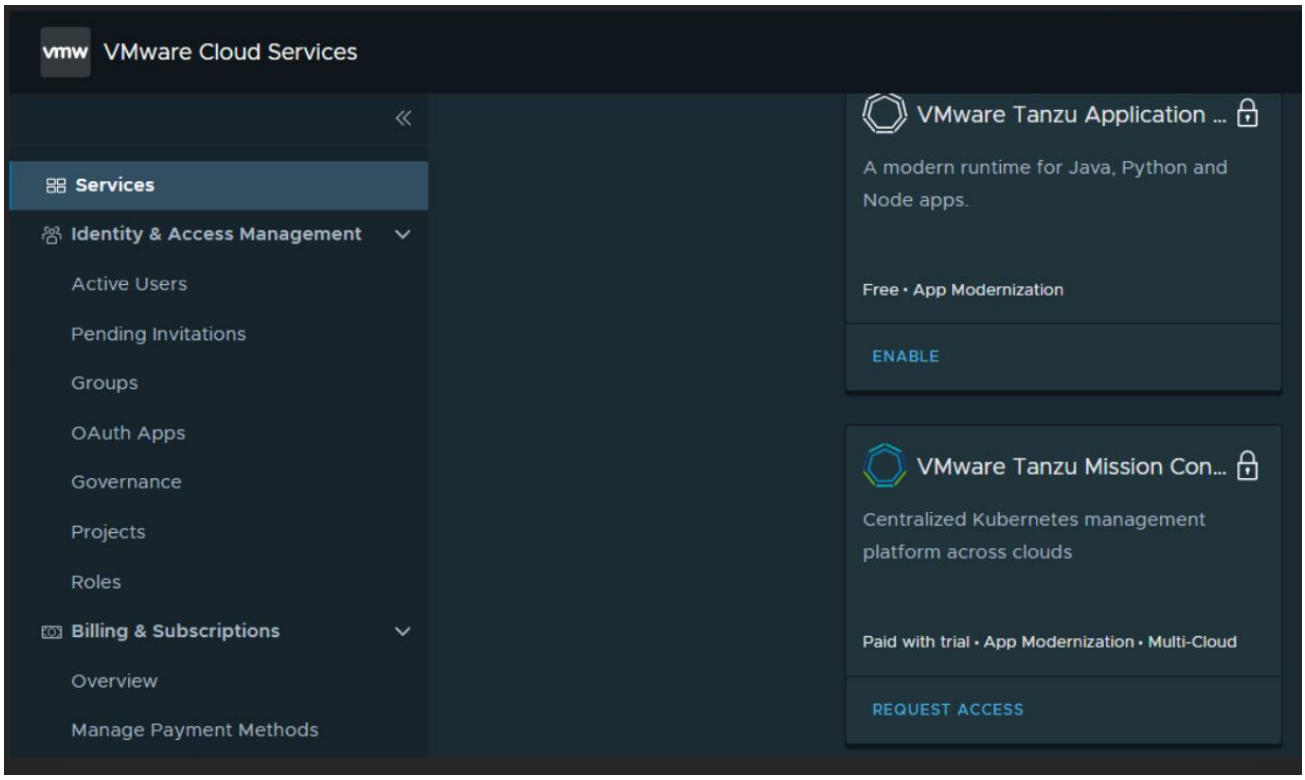

Figure 7: Tanzu Mission Control and VMware Cloud Services

5.6 VMware Licensing

The licensing for vSphere is based on a CPU metric and licensing for other products is based on the number of OS instances. Other components, such as NSX, have their own separate licenses and are optional add-ons. Table 9 lists the standard and optional components that are provided with VCF editions. However, add on licenses can be added as long as they meet compatibility.

Table 9: VMware VCF Editions

Component	**Starter**	**Standard with Tanzu**	**Advanced with Tanzu**	**Enterprise with Tanzu**
SDDC Manager				
vSphere	Enterprise Plus	Enterprise Plus	Enterprise Plus	Enterprise Plus
vSAN	Advanced	Advanced	Advanced	Enterprise
NSX-T	Advanced	Advanced	Advanced	Enterprise Plus
Aria Operations for Networks	Advanced		Advanced	Enterprise
Aria Suite	Standard		Enterprise	Enterprise
Tanzu		Standard	Standard	Standard
Tanzu standard can be upgraded to Tanzu Advanced separately				

5.7 HyTrust Security

HyTrust provides a suite of security-oriented products for vSphere environment. These products are HyTrust KeyControl, DataControl, and CloudControl. Note that the HyTrust products are currently supported on ESXi 7.0U2 and NSX-T Data Center 3.1..2.

5.7.1 HyTrust KeyControl

HyTrust KeyControl (HTKC) enables enterprises to easily manage all their encryption keys at scale, how often they rotate them, and how they are shared securely. HyTrust KeyControl capabilities include:

- VMWare Certified Key Manager Server (KMS) for:
  - vSphere 7.0u2 and vSAN 7.0
- Universal key management for KMIP-compatible encryption agents
- Enterprise scalability and performance
- KeyControl can run in an active-active, high availability cluster
- FIPS 140-2 Level 1 validation and FIPS 140-2 Level 3 hardware security module (HSM)

5.7.2 HyTrust DataControl

HyTrust DataControl (HTDC) secures multi-cloud workloads throughout their lifecycle. DataControl helps manage workloads and encryption keys from a central location to reduce complexity, comply with regulations such as the GDPR.

DataControl provides granular encryption for better multi-cloud security. The protection boundary does not stop at the hypervisor or at the data store; VMs are individually encrypted. Inside the VM, unique keys can be assigned to encrypt individual partitions, including the boot (OS) disk. Encryption and rekeying can be done on the fly and there is no need to take workloads off-line.

Table 10 compares the data encryption features of vSphere, vSAN, and HyTrust DataControl/KeyControl.

Table 10: Comparison of Encryption Features

Encryption	vSphere VM Encryption	vSAN Encryption	HyTrust DataControl
Protection level	Data at rest and in motion	Data at rest	Data at rest
Encryption Approach	Hypervisor does the encryption	Disk based encryption	In Guest encryption
Components	KMS, vCenter, ESXi Host	KMS, vCenter, ESXi Hosts in vSAN Cluster, Disks	KMS, HyTrust DataControl Agent
Encryption Cipher	AES-XTS-256	AES-XTS-256	AES-XTS-512,AES-XTS-256, AES 128
Encrypted objects	Virtual machine files, virtual disk files, and ESXi core dump files	All files in the vSAN datastore	All data in the drives

Encryption	vSphere VM Encryption	vSAN Encryption	HyTrust DataControl
Interface	vSphere Web Client, vSphere Web Services SDK	vSphere Web Client	HyTrust DataControl UI in the Guest OS. HyTrust KeyControl UI to manage VM Set, VMs and users
Enabling Option	Per VM level through vSphere Encryption Storage Policy	Enabled at cluster or VSAN datastore level	Enabled within Guest OS
Access Control	Users with vSphere Cryptographic Operations Privileges	Users with vSphere Cryptographic Operations Privileges	Guest OS User uses KeyControl admin user. Authorization can also be done by HyTrust CloudControl
Interoperability Limitations	vSphere Fault Tolerance, vSphere Replication, Content Library	N/A	N/A
Platform Support	All Guest OS running on the Hypervisor	All Guest OS running on the Hypervisor	Most Windows and Linux flavors and version running on vSphere, KVM, Hyper-V, or XenSever

5.7.3 HyTrust CloudControl

HyTrust CloudControl (HTCC) provides a variety of security and policy enhancements without impacting the existing GUI of vSphere, NSX and ESXi. CloudControl is deployed as a transparent proxy and mediates the actions taken by administrators using familiar interfaces. CloudControl provides the following security features:

- **Role Based Access Control** (RBAC) to control which functions have access to what resources and allows a much closer alignment of access rights to governance and compliance requirements.
- **Policy Control including Two Man Rule** to define and more importantly enforce policy including requiring secondary approval for potentially disruptive actions, reducing potential impact of human error or intentional malevolent behaviour.
- **Access Control including Two Factor Authentication** to significantly enhance the overall security posture of an organization without the traditional weaknesses of using even strong passwords.
- **Forensic grade logs** to provide an in-depth perspective on what has happened as well as what has not happened in your virtual environment.

Table 11 compares the access control features of vCenter and HyTrust CloudControl.

Table 11: Comparison of access control features

Access Control Feature	vCenter	HyTrust CloudControl
vSphere Web Client Access	vCenter URL	Published IP (PIP) associated with vCenter

Access Control Feature	vCenter	HyTrust CloudControl
Authentication	vCenter SSO, IWA	vCenter SSO, IWA, HTCC Service Account, Two factor authentication with RSA Secure ID, RADIUS, or TACACS+
Authorization	Predefined permissions to access various vCenter components	Uses permissions defined in vCenter
vCenter Users	SSO users from multiple AD Domain and vSphere local domain. Predefined solution users for vSphere services.	Users from Single AD Domain which includes configured HTCC Service Account
vCenter User Access Setup	Directory users/group need to be added in vCenter SSO users/group	Directory users need to be added to respective HTCC directory group which is associated with HTCC role
User Groups	14 predefined SSO groups. Directory users/group is mapped to SSO groups.	16 predefined rules for vSphere. HTCC directory group is mapped to HTCC rule.
Role Based Access Control	14 predefined roles with respective privileges	16 predefined roles for vSphere with appropriate privileges
Custom Roles Creation	Supported	Supported
Secondary Approval	Not Available	Available for set of compute and network operations
Auditing	Integrated with Aria Operations for Logs. Auditing dashboard is available based on the event type. User's session details can be monitored in vSphere web client.	Has its own Log Viewer and dashboard. Logs can be redirected to use Aria Operations for Logs as syslog server.

5.7.4 Compliance Management

An important part of security is compliance management. VMware Aria Configuration Manager has twenty built-in compliance templates and others can be added. HyTrust CloudControl (HTCC) supports customizing built-in compliance templates but does not provide any out of the box.

Table 12 compares the compliance management features of Aria Configuration Manager and HTCC.

Table 12: Comparison of compliance management features

Compliance Management Feature	Aria Configuration Manager	HTCC
ESXi Host Compliance	Yes	Yes
Guest Virtual Machine Compliance	Yes	Limited
NSX Manager Compliance	No	Yes

Patching assessment and Deployment	Yes	No
Active Directory Compliance	Yes	No
Software Asset Management	Yes	No
Integration with Aria Operation Manager	Yes	No
Manage Virtual Machines	Yes	No

6 Operational model

This chapter describes the options for mapping the logical components of SDDC onto Lenovo ThinkAgile VX servers. The following section describes the hardware components in a SDDC deployment.

6.1.1 Servers

You can use various rack-based Lenovo ThinkAgile VX server platforms to implement edge, management, or compute clusters with VMware vSAN and supports All Flash and Hybrid configurations.

Server Model	Processor	Drives	Memory	GPU
VX2330 1U Appliance	2x Intel Xeon SP Gen 3	4x 3.5” SAS/SATA 4x 3.5” NVMe	Up to 4TB	No
VX3330 1U Appliance	2x Intel Xeon SP Gen 3	12x 2.5" SAS/SATA 12x 2.5" NVMe	Up to 4TB	No
VX3530-G 2U Appliance	2x Intel Xeon SP Gen 3	24 x 2.5" (HS)	Up to 4TB	Yes
VX3575-G 2U Appliance	2 x 64 Core AMD EPYC™ 7003 Series	24x 2.5" SAS/SATA	Up to 4TB	Yes
VX5530 2U Appliance	2x Intel Xeon SP Gen 3	16 x 3.5” (HS)	Up to 4TB	No
VX7330-N 1U Appliance	2x Intel Xeon SP Gen 3	12x 2.5" NVMe	Up to 4TB	No
VX7530 2U Appliance	2x Intel Xeon SP Gen 3	40x 2.5" SAS/SATA 32x 2.5" NVMe	Up to 4TB	No
VX7820 4U Appliance	4x Intel Xeon SP Gen 3	Up to 24x SFF (HS)	Up to 4TB	No
VX5575 2U Appliance	2x 64 Core AMD EPYC™ 7003 Series	16x 3.5" SAS/SATA	Up to 4TB	
VX7576 2U Appliance	2x 64 Core AMD EPYC™ 7003 Series	35x 2.5" SAS/SATA 16x 3.5" SAS/SATA 32x 2.5" NVMe	Up to 4TB	Yes
VX7575 2U Appliance	2x 64 Core AMD EPYC™ 7003 Series	35x 2.5" SAS/SATA 32x 2.5" NVMe	Up to 4TB	No
Lenovo ThinkAgile 1U Certified Node VX3331, ThinkSystem SR630v2	1x or 2x Intel Xeon SP Gen 3	12x 2.5" SAS/SATA 12x 2.5" NVMe 4x 3.5" SAS/SATA	Up to 4TB	

Lenovo ThinkAgile VX 2U Certified Nodes VX7531, ThinkSystem SR650v2	2x Intel Xeon SP Gen 3	40x 2.5" SAS/SATA 32x 2.5" NVMe 16x 3.5" SAS/SATA	Up to 4TB	
ThinkAgile VX650 V3 2U Integrated System and VX650 V3 2U Certified Node	**1x or 2x 60 Core Intel Xeon SP Gen 4**	**Front Bays: 12x3.5" or 24x2.5"** **Mid Bays: 4x3.5" or 8x2.5"** **Rear bays: 4x3.5" or 8x2.5" (supports 2x7mm hot-swap drives bays)**	Up to 8TB	Yes
ThinkAgile VX630 V3 1U Integrated System and Certified Node	**1x or 2x 60 Core Intel Xeon SP Gen 4**	**Front Bays: 12x2.5" SAS or 12x2.5" NVMe** **Rear bays: 4x3.5" SAS or 4x3.5" NVMe (supports 2x7mm hot-swap drives bays)**	Up to 8TB	Yes
ThinkAgile VX650 V3 2U Integrated System and VX650 V3 2U Certified Node	**1x or 2x 64 Core Intel Xeon SP Gen 5**	**Front Bays: 12x3.5" or 24x2.5"** **Mid Bays: 4x3.5" or 8x2.5"** **Rear bays: 4x3.5" or 8x2.5" (supports 2x7mm hot-swap drives bays)**	Up to 8TB	Yes
ThinkAgile VX630 V3 1U Integrated System and Certified Node	**1x or 2x 64 Core Intel Xeon SP Gen 5**	**Front Bays: 12x2.5" SAS or 12x2.5" NVMe** **Rear bays: 4x3.5" SAS or 4x3.5" NVMe (supports 2x7mm**	Up to 8TB	Yes

		hot-swap drives bays)		
Lenovo ThinkAgile VX 2U Certified Nodes VX7576, ThinkSystem SR665	2x AMD EPYC 7003 Milan	35x 2.5-inch SAS/SATA 16x 3.5-inch SAS/SATA 32x 2.5-inch NVMe	Up to 4TB	
Lenovo ThinkAgile VX 1U Integrated System and Certified Nodes VX635, ThinkSystem SR635	1x AMD EPYC 9004 Genoa	Front bays: 10 x 2.5-inch SAS/SATA or NVMe Rear bays: 2x 2.5-inch SAS/SATA 2x (7mm) SATA/NVME	Up to 1.5TB	Yes
Lenovo ThinkAgile VX 1U Integrated System and Certified Nodes VX645 V3, ThinkSystem SR645 V3	2x AMD EPYC 9004 Genoa	Front bays: 10 x 2.5-inch SAS/SATA or NVM Rear bays: 8x 2.5-inch SAS/SATA 2x (7mm) SATA/NVME EDSFF (New): 16x E3.S thick Hot-Swap	Up to 6TB	Yes
Lenovo ThinkAgile VX 2U Integrated System and Certified Nodes VX655 V3, ThinkSystem SR655 V3	1x AMD EPYC 9004 Genoa	Front bays: 24x 2.5-inch SAS/SATA or NVMe 12x 3.5-inch SAS/SATA or Anybay Mid bays: 8x 2.5-inch SAS/SATA or NVMe 4x 3.5-inch SAS/SATA Rear bays: 8x 2.5-inch SAS/SATA 4x 2.5-inch Anybay 4x 3.5-inch SAS/SATA 2x (7mm) SATA/NVME	Up to 1.5TB	Yes
Lenovo ThinkAgile VX 2U Integrated System and	2x AMD EPYC 9004 Genoa	Front bays: 24x 2.5-inch SAS/SATA	Up to 6TB	Yes

Certified Nodes VX665 V3, ThinkSystem SR665 V3		12x 3.5-inch SAS/SATA or Anybay Mid bays: 8x 2.5-inch SAS/SATA or NVMe 4x 3.5-inch SAS/SATA Rear bays: 8x 2.5-inch SAS/SATA 4x 2.5-inch Anybay 4x 3.5-inch SAS/SATA 2x (7mm) SATA/NVMe Front I/O Chassis: 8x 2.5-inch Anybay and 8x 2.5-inch SAS/SATA		

6.2 Edge cluster servers

The edge cluster runs NSX services for all tenants in the SDDC infrastructure, provides internal and external routing, and also runs tenant workloads.

The shared edge and compute cluster uses its own dedicated vCenter server and NSX-T manager which are deployed in the management cluster. The NSX controllers and edge gateway services VMs are deployed on the shared cluster. The tenant VMs can be deployed in the shared edge and compute cluster or in a separate compute cluster leveraging the vCenter server and NSX services in the shared edge and compute cluster.

6.2.1 Edge and Infrastructure Services VMs

The VMs used for infrastructure services such as Active Directory, DNS/DHCP, firewalls, proxy and anti-virus are deployed in the shared edge and compute cluster. Table 13 lists each infrastructure service VM with the recommended sizes in terms of virtual CPUs, RAM, storage, and networking.

Table 13: Infrastructure services VMs

VM description	CPU (vCPUs)	Memory (GB)	Storage (GB)	Network bandwidth	High availability
AD, DHCP, DNS server	2	4	70	1 GbE	clustered
http proxy server	2	4	30	1 GbE	clustered
NSX Controller (odd # deployment; min 3)	4	4	20	1 GbE	Built-in/vSphere HA

Table 14 lists the NSX service VMs with the recommended sizes in terms of virtual CPUs, RAM, storage, and networking.

Table 14: Edge services VMs for NSX

VM description	CPU (vCPUs)	Memory (GB)	Storage (GB)	Network bandwidth	High availability
Compact (also used for logical router)	1	0.5	0.5	1 GbE	Yes, Optional
Large	2	1	0.5	1 GbE	Yes, Optional
Quad Large	4	1	0.5	1 GbE	Yes, Optional
X-Large	6	8	4.5	1 GbE	Yes, Optional

The actual VM size (compact, large, quad-large, and X-large) depends on the number of type of services that are deployed in the VM. A logical router is always deployed by using a compact VM. A quad large is required for a firewall and an X-large is used for more than one service (for example, firewall, load balancer, and router).

6.2.2 Hybrid cloud VMs

Table 15 lists the cloud connectivity VMs with the recommended sizes in terms of virtual CPUs, RAM, storage, networking, and location. Note that these VMs do not have options for high availability.

Table 15: Cloud connectivity VMs

VM description	CPU (vCPUs)	Memory (GB)	Storage (GB)	Network bandwidth	Location
VMware HCX	2	4	300	1 GbE	On-Premise

6.3 Management cluster servers

The number of VMware SDDC components in the management cluster increases as capabilities are added. This section addresses the SDDC management components that could be used. Third party add-ons must be sized separately.

6.3.1 Management cluster VMs

There are several considerations that contribute to an end-to-end sizing of an entire VMware vCloud environment including Lenovo software for systems management. This section is intended to provide some high-level guidance for management cluster configuration sizing. The recommended number of virtual CPUs, memory size, storage size, and network bandwidth are given for each VM and the VMs are grouped by each major component or appliance.

An essential part of the infrastructure is load balancing of the server VMs and recognizing when a server is down and failing over to another server. The following cases are available for VMs in the management cluster:

- vSphere HA: vCenter automatically restarts the VM on another server, but there is some downtime while the VM starts up.
- Microsoft SQL server clustering: The SQL server cluster automatically handles failover.
- Clustering within component to provide built-in high availability.

Load balancing: An external load balancer such as a Big-IP switch from F5 and/or VMware NSX load balancers can be used.

Table 16 lists each management cluster VM for vSphere with its recommended size in terms of virtual CPUs, RAM, storage, and networking.

Table 16: Management cluster VMs for vSphere

VM description	CPU (vCPUs)	Memory (GB)	Storage (GB)	Network bandwidth	High availability
SDDC Manager	4	16	1000	1 GbE	vSphere HA
vCenter Server Appliance(1) Management Cluster	8	24	50	1 GbE	load balancer
vCenter Server Appliance(2) Edge and Compute Cluster	8	24	50	1 GbE	load balancer

VM description	CPU (vCPUs)	Memory (GB)	Storage (GB)	Network bandwidth	High availability
SDDC Manager	4	16	1000	1 GbE	vSphere HA
vCenter Server Database (MS SQL)	4	8	200	1 GbE	SQL AlwaysOn Availability Group
vSphere Replication	2	4	20	1 GbE	not required
vSphere Data Protection	4	4	1600	1 GbE	not required
Aria Orchestrator Appliance	2	3	12	1 GbE	Clustered

Table 17 lists each management cluster VM for Aria Automation with its size in terms of virtual CPUs, RAM, storage, and networking.

Table 17: Management cluster VMs for Aria Automation

VM description	CPU (vCPUs)	Memory (GB)	Storage (GB)	Network bandwidth	High availability
Aria Suite Lifecycle Manager	4	16	135	1 GbE	N/A
Aria Automation Appliance	4	16	30	1 GbE	load balancer
IaaS Database (MS SQL)	8	16	100	1 GbE	SQL AlwaysOn Availability Group
Infrastructure Web Server	2	4	40	1 GbE	load balancer
Infrastructure Manager Server	2	4	40	1 GbE	load balancer
Distributed Execution Manager (DEM)	2	6	40	1 GbE	load balancer
vSphere Proxy Agent	2	4	40	1 GbE	load balancer
Aria Application Services	8	16	50	1 GbE	vSphere HA

Table 18 lists each management cluster VM for Aria Operations Manager with its size in terms of virtual CPUs, RAM, storage, and networking.

Table 18: Management cluster VMs for Aria Operations Manager

VM description	CPU (vCPUs)	Memory (GB)	Storage (GB)	Network bandwidth	High availability
Aria Operations Manager – Master	4	16	500	1 GbE	clustered
Aria Operations Manager – Data	4	16	500	1 GbE	not required
Aria Configuration Manager – Collector	4	16	150	1 GbE	load balancer

VM description	CPU (vCPUs)	Memory (GB)	Storage (GB)	Network bandwidth	High availability
Aria Configuration Manager Database (MS SQL)	4	16	1000	1 GbE	SQL AlwaysOn Availability Group
Aria Hyperic Server	8	12	16	1 GbE	load balancer
Aria Hyperic Server - Postgres DB	8	12	75	1 GbE	load balancer
Aria Infrastructure Navigator	2	4	24	1 GbE	not required

Table 19 lists the management VMs that are needed for NSX.

Table 19: NSX-T Management cluster VMs

VM description	CPU (vCPUs)	Memory (GB)	Storage (GB)	Network bandwidth	High availability
NSX-T Manager Management Cluster	4	12	300	1 GbE	vSphere HA
NSX-T Controller Management Cluster (odd # deployment; min 3)	4	4	20	1 GbE	Built-in/vSphere HA
NSX-T Manager Edge and Compute Cluster	4	12	60	1 GbE	vSphere HA

Table 20 lists each management cluster VM for HyTrust with its size in terms of virtual CPUs, RAM, storage, and networking.

Table 20: Management cluster VMs for HyTrust

VM description	CPU (vCPUs)	Memory (GB)	Storage (GB)	Network bandwidth	High availability
HyTrust CloudControl	4	16	70	1 GbE	Clustered
HyTrust KeyControl	2	8	20	1 GbE	Clustered

Table 21 lists the VMs that are needed for Lenovo software for systems management.

Table 21: Lenovo System Management VMs

VM description	CPU (vCPUs)	Memory (GB)	Storage (GB)	Network bandwidth	High availability

Lenovo XClarity Administrator	2	4	64	1 GbE	not required
Lenovo XClarity Orchestrator	4	16	500	1 GbE	not required
Lenovo XClarity Integrator (Windows OS)	1	2	30	1 GbE	not required

6.4 Systems management for Lenovo servers

Lenovo XClarity™ family of system management software and tools provide centralized resource management, monitoring and analytics solution that reduces complexity, speeds up response, and enhances the availability of Lenovo® server systems and solutions. The XClarity integrator plugins are designed to work with VMware VCF components as an extension to simplify the operations. For more information, see this website: https://www.lenovo.com/us/en/data-center/software/management/.

6.4.1 Lenovo XClarity Administrator(LXCA)

The Lenovo XClarity Administrator provides agent-free hardware management for Lenovo's ThinkAgile, ThinkSystem® rack servers, System x® rack servers, and Flex System™ compute nodes and components, including the Chassis Management Module (CMM) and Flex System I/O modules. Figure 8 shows the Lenovo XClarity administrator interface, in which Flex System components and rack servers are managed and are seen on the dashboard. Lenovo XClarity Administrator is a virtual appliance that is quickly imported into a virtualized environment server configuration. Lenovo XClarity Administrator supports auto discovery of endpoints, inventory, monitoring, firmware compliance, firmware updates, Windows device driver updates, configuration management and compliance, user management, deployment of operating systems and hypervisors to bare metal servers.

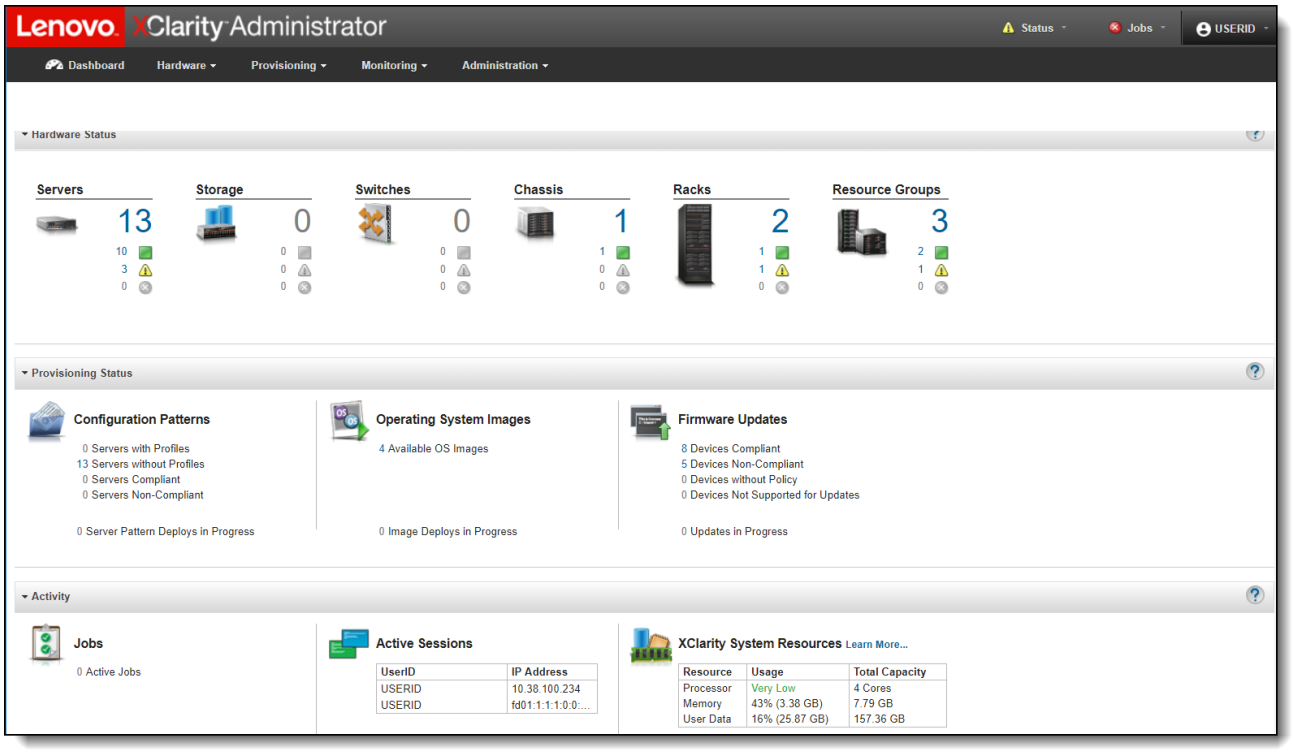

Figure 8: XClarity Administrator dashboard

6.4.2 Lenovo XClarity Orchestrator(LXCO)

XClarity Orchestrator provides a single interface to monitor and manage multiple Lenovo XClarity Administrators and the devices managed by them. LXCO supports deploying updates to Lenovo Lenovo XClarity Administrator and firmware updates to devices that are managed. LXCO can connect to third-party services (such as Splunk) for business intelligence machine learning and predictive analytics to collect resource utilization data and uses metric data to predict failures, create reports and custom alert rules that, when enabled, raise alerts when specific conditions exist in your environment.

6.4.3 Lenovo XClarity Integrators (LXCI) for VMware

Lenovo provides XClarity integration modules for VMware vCenter, VMware Aria Automation, VMware Aria Orchestrator and VMware Aria Operations for Logs.

By using the Lenovo XClarity Integrator for VMware vCenter, administrators can consolidate physical resource management in VMware vCenter, which reduces the time that is required for routine system administration.By using the Lenovo XClarity Integrator for VMware vCenter, administrators can consolidate physical resource management in VMware vCenter, which reduces the time that is required for routine system administration.

The Lenovo XClarity Integrator for VMware vCenter provides the following features and benefits:

- Extends Lenovo XClarity Administrator features to the virtualization management console
- Enables management of legacy infrastructure from the virtualization management console
- Reduces workload downtime by dynamically triggering workload migration in clustered environments during rolling server reboots or firmware updates, and predicted hardware failures

Figure 9 shows Lenovo XClarity Integrator deployed in the vCenter and displays ThinkAgile VX nodes.

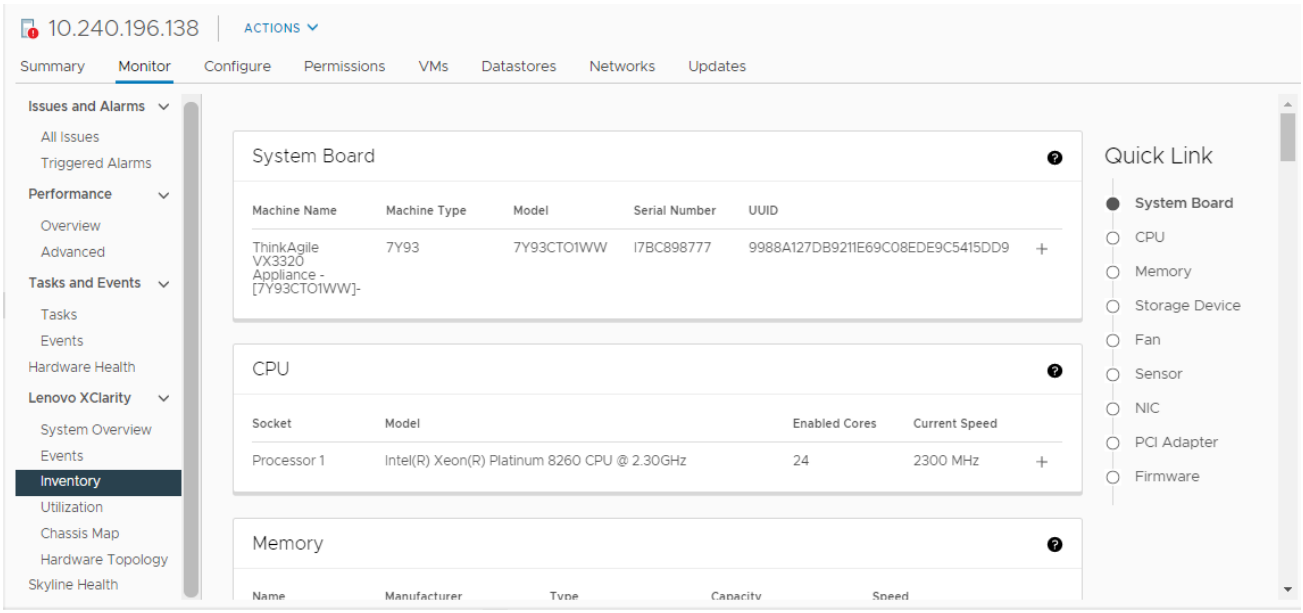

Figure 9: Lenovo XClarity Integrator for VMware vCenter

The Lenovo XClarity Integrator for VMware Aria Orchestrator provides IT administrators with the ability to coordinate physical server provisioning features of Lenovo XClarity Pro with broader Aria Orchestrator workflows. Lenovo XClarity Integrator for VMware Aria Orchestrator provides a library of simple yet robust and customizable workflow routines and actions designed to automate complex, repetitive IT infrastructure tasks such as system discovery and configuration, hypervisor installation, and addition of new hosts to vCenter. Figure 10 shows the Lenovo XClarity Integrator for Aria Orchestrator workflow interface.

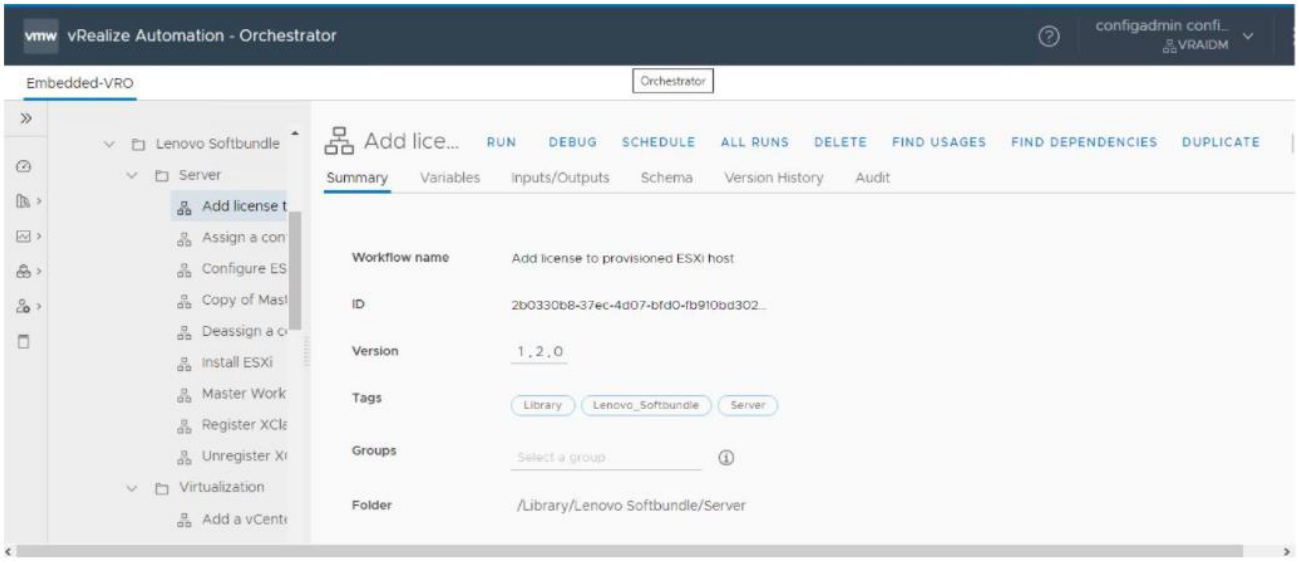

Figure 10: Lenovo XClarity Integrator for VMware Aria Orchestrator interface

The Lenovo XClarity Administrator Content Pack for VMware Aria Operations for Logs simplifies the collection and forwarding of Lenovo XClarity Administrator logs to VMware Aria Operations for Logs for powerful processing and analytics, and displaying insightful information in an intuitive format.

The VMs for VMware vCenter, Aria Orchestrator, Lenovo XClarity Administrator and Lenovo XClarity Administrator Integrator should have access to the management network used for managing servers, storage and networking.

Lenovo XClarity Integrator for Aria Automation provides a set of blueprints to provision infrastructure services based on Lenovo servers, network switches and vSphere. This eases provisioning a new Lenovo server with vSphere installed, network isolation parameters configured on the Lenovo switches, apply vSphere distributed switch configurations and adding the server to the existing or new vSphere Cluster. The Lenovo Aria content pack for Aria Automation needs to imported into Aria Orchestrator and then the Blueprints package is imported using the Aria Cloud Client command line utility by Tenant Administrators and it creates catalog items automatically. The catalog items are created under Lenovo Servers, Lenovo Network, and Lenovo Virtualization services. Figure 11 shows Lenovo XClarity Integrator template items for Aria Automation.

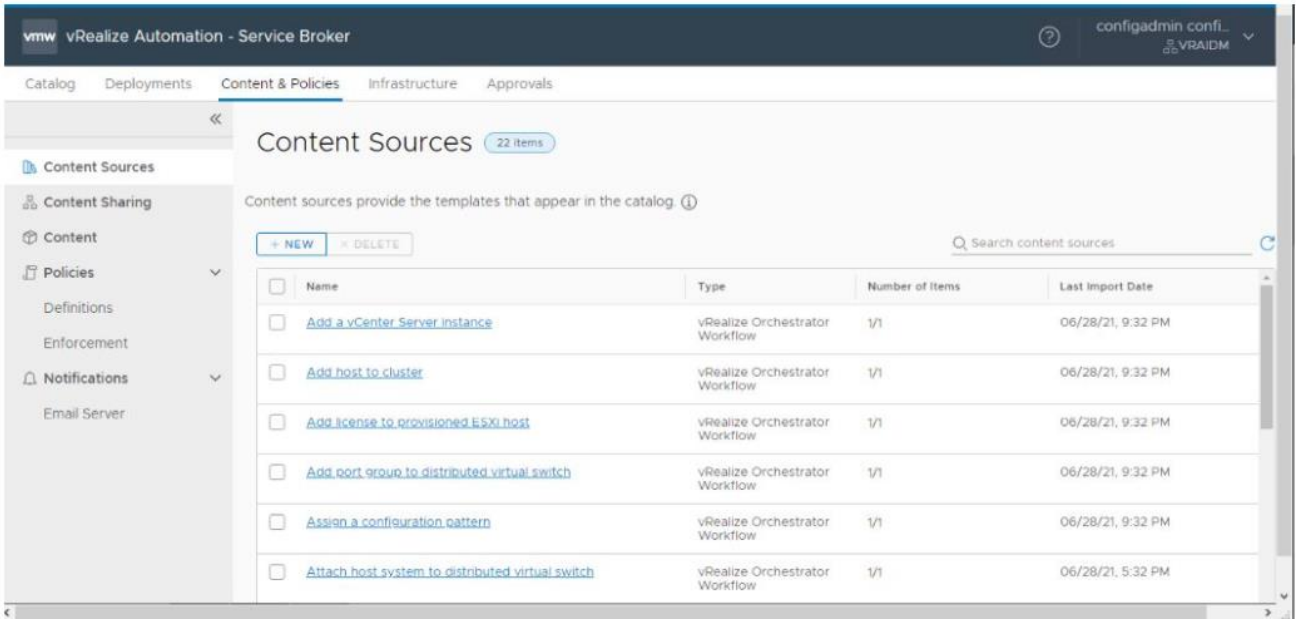

Figure 11: Lenovo XClarity Integrator for Aria Automation template Items

Lenovo XClarity Adapter for Aria Operations Manager provides a global view of the relationship between resources, such as connected chassis, servers, power supplies, and ESXi connectivity. The plugin helps to monitor the hardware events in a Lenovo XClarity Administrator-managed environment. Quickly identify trends based on hardware events received, including hardware failures, power/thermal thresholds that exceeded, and PFAs (predicted failure alerts). These events categorize by source, type of hardware surfacing the events, and whether service is required. This information can help identify issues in your data centers so that you can react before more serious issues occur. Figure 12 shows the XClarity Adapter for Aria Operations Manager interface summary tab contains alerts and recommended actions.

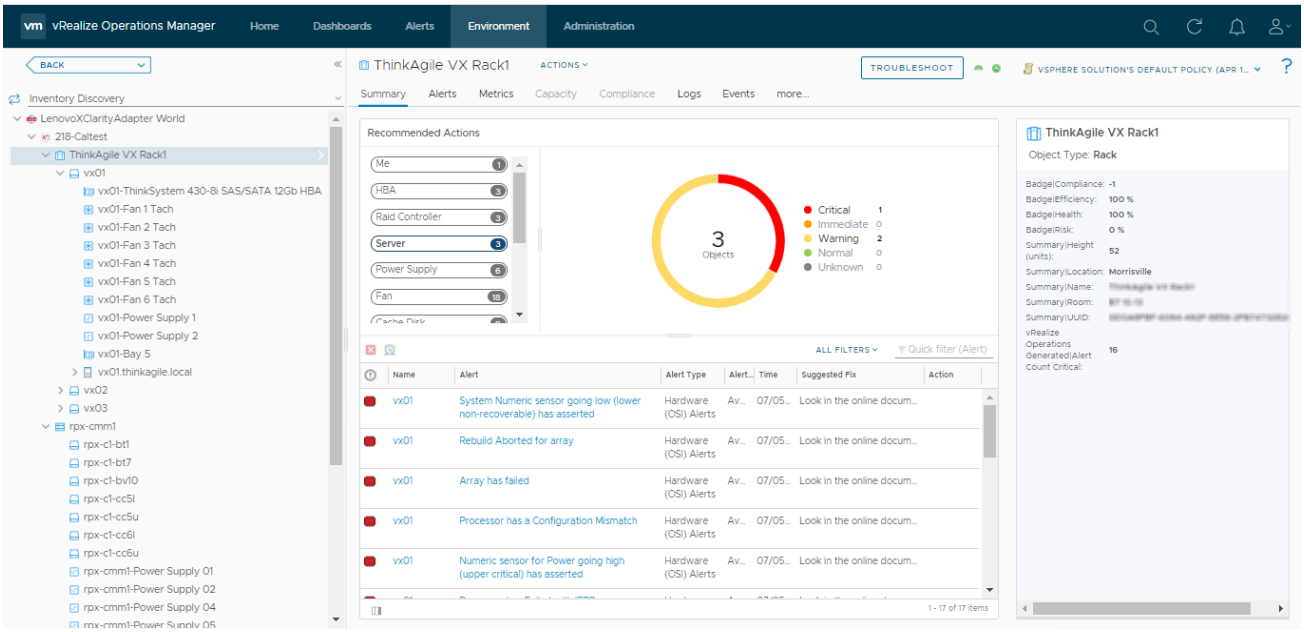

Figure 12: Lenovo XClarity Integrator for Aria Operations Manager

The Aria Operations for Logs content pack for provides analysis of events from the Lenovo XClarity Administrator, Lenovo XClarity Orchestrator, and the resources managed by XClarity. These insights helps to monitor hardware events, resource alerts, auditing security changes, firmware upgrades and configuration management. Figure 13 shows the events insight page for Lenovo XClarity content pack for Aria Operations for Logs.

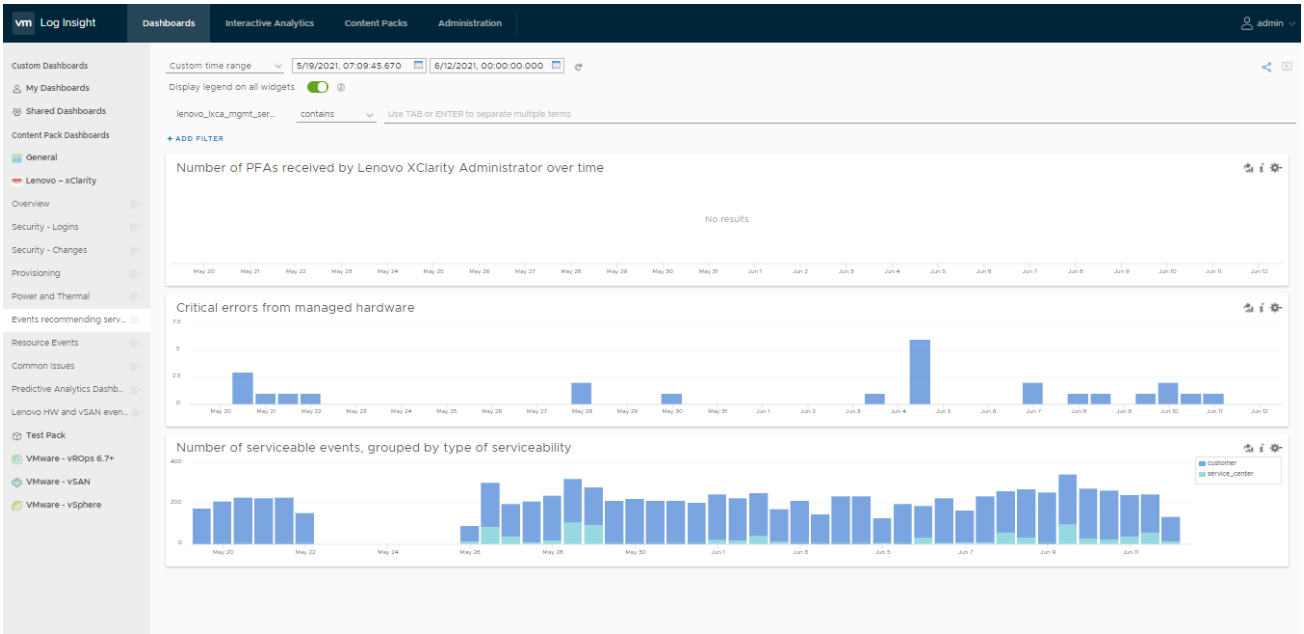

Figure 13: Lenovo XClarity Integrator for Aria Operations for Logs

6.9.3 Lenovo XClarity plugins compatibility

Table 22 below lists current versions of Lenovo integration plugins and the required or supported VMware vCenter and Aria Suite products.

Table 22: Plug-in compatibility

Component Name	Version	Supported Product Versions
Lenovo XClarity Administrator (LXCA)	3.4	VMware vCenter 6.0U2/6.5/6.7, ESXi 6.0U2/6.5 U1/6.7/7.0U2
Lenovo XClarity Integrator (LXCI) for vCenter	7.4	Lenovo XClarity Administrator 1.4.x, 2.x VMware vCenter 5.x U1/U2/U3, 6.0 U1/U2/U3, 6.5 U1/U2,6.7(U1,U2,U3), 7.0(U1, U2,U3)
Lenovo XClarity Administrator content pack for VMWare Aria Operations for Logs	1.0	Lenovo XClarity Administrator 1.1 or higher VMware Aria Operations for Logs 2.5 or higher
Lenovo XClarity Integrator for VMware Aria Automation	1.2	VMware Aria Automation 8.3 or higher
Lenovo XClarity Integrator for VMware Aria Orchestrator	1.2	VMware Aria Automation 7.0 VMware Aria Orchestrator 6.0/7.0
Lenovo Network Plugin for VMware Aria Orchestrator	1.4.0	VMware Aria Orchestrator 7.4.x
Lenovo XClarity Content Pack for Aria Operations Manager	1.2	Aria Operations Manager 8.0, 8.1, 8.2, and 8.3

7 Deploying SDDC

This chapter provides an introduction to deploying SDDC in your data center.

7.1 VMware Validated Design

The VVD documentation (version 6.1) provides a family of solutions for data center designs that span compute, storage, networking, and management, serving as a blueprint for an SDDC implementation.

This reference design is based on VVD. For more details on VVD, please see this website: vmware.com/support/pubs/vmware-validated-design-pubs.html.

7.2 VMware Cloud Foundation

VMware Cloud Foundation (VCF) is a hybrid cloud platform to deploy VMware SDDC for private cloud based on the VMware Validated Design and to integrate with public clouds running VMware SDDC clouds. It provides software defined services for compute, storage, networking, and cloud management to run different workloads. It simplifies installation, upgrade and patch management of SDDC components through lifecycle management either through online or offline.

VCF supports deploying SDDC components on broad range of physical servers (vSAN Ready Nodes) to have flexible customer defined heterogeneous infrastructure to support variety of workloads.

7.2.1 SDDC Manager

The SDDC Manager provides the core management software for VCF. It automates the installation and lifecycle management of the vSphere, vSAN, and NSX from bring-up and configuration to patching and upgrading, making it simple for the cloud admin to build and maintain the SDDC. SDDC Manager also automates the installation and configuration of Aria Operations for Logs, Aria Operations, and Aria Automation by using Aria Suite Lifecycle Manager. SDDC Manger uses same vCenter sso login. The cloud administrator uses vCenter Server as the primary management interface for the virtualized environment.

7.2.2 Workload Domain

A workload domain is a dedicated environment with servers, storage and networking managed by dedicated vCenter and NSX Manager. The management workload domain is created automatically and virtual infrastructure workload domains are created by cloud administrators based on requirements. The resource maximums, limits and scalability for each workload domain is same as the limits applicable for vCenter. The SDDC Manager deploys and configures one vCenter Server and NSX manager per workload domain automatically when the workload domain is created.

Figure 14 shows an example of a management workload domain and two virtual infrastructure workload domains.

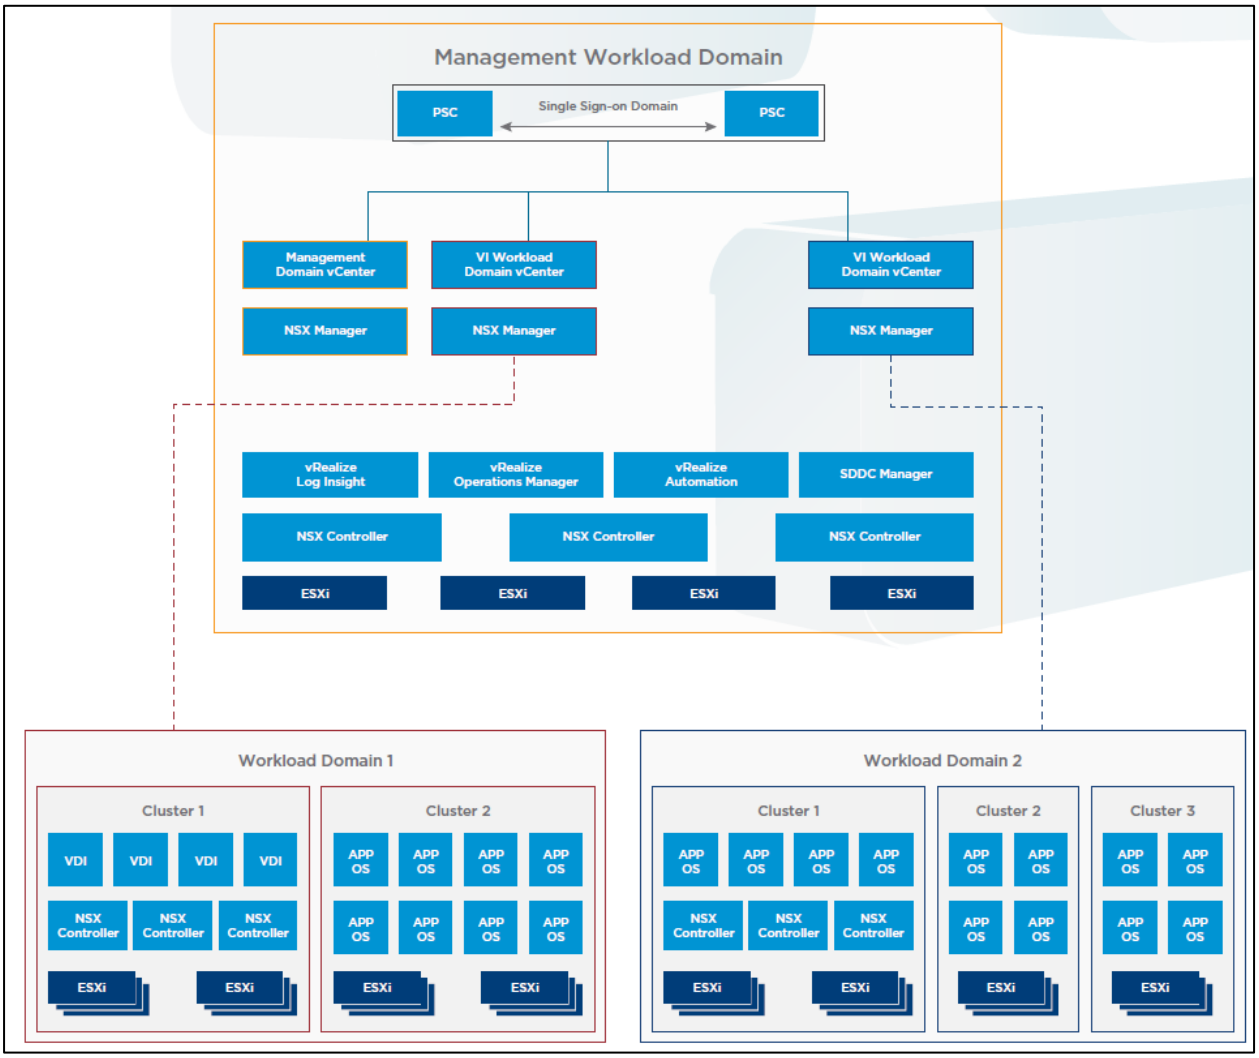

Figure 14: VCF Workload domains

7.3 Lenovo VX Appliance

Lenovo ThinkAgile VX appliances are preloaded with a wizard-based deployment tool to accelerates the greenfield vSAN deployment or new clusters with ESXi 7.0u2 or later. A 4-Node vSAN cluster can be deployed in less than an hour and it works with All Flash and Hybrid vSAN deployments. It discovers the Lenovo ThinkAgile VX nodes over the network, installs ESXi, deploys vSAN and vCenter and install Aria plugins in the vCenter. Figure 15 shows logical network architecture for deploying ThinkAgile VX cluster to setup vSAN using VX Deployer tool. Please refer this page to use VX Deployer to setup vSAN cluster on VX Appliances and verify the deployed components
https://thinkagile.lenovofiles.com/help/index.jsp?topic=%2Fcom.lenovo.thinkagile.vx%2Fcluster_deployment_with_vx_deployer.html

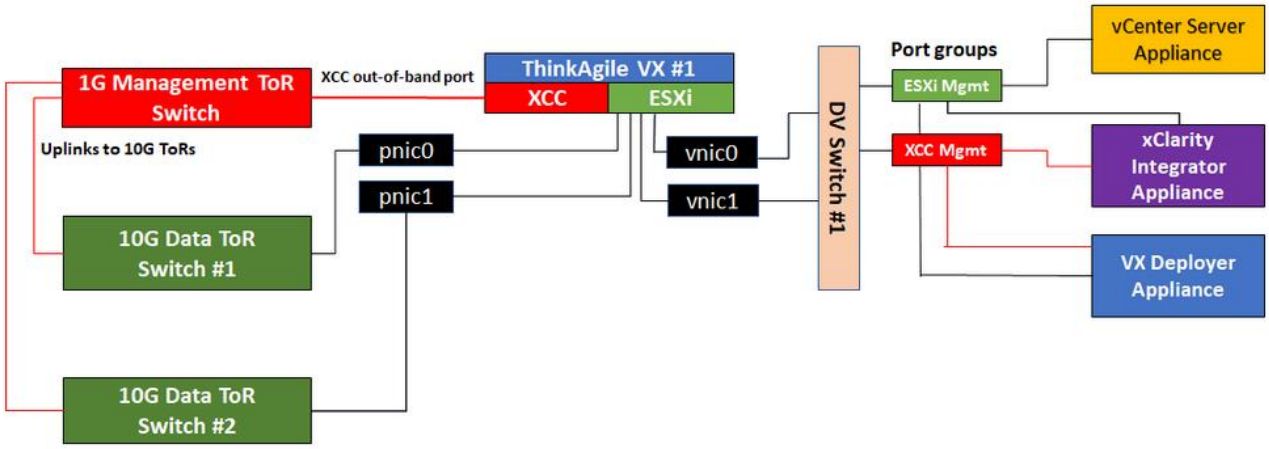

Figure 15: Logical network architecture for Lenovo ThinkAgile VX cluster

7.3.1 Deploying VCF with ThinkAgile VX Appliances

VCF can be installed on Lenovo ThinkAgile VX certified nodes or VX appliances as both have been vSAN certified. Lenovo has validated the install of VCF 3.5.

Table 23 describes the steps to install a complete SDDC environment using VCF 3.5 and ThinkAgile VX appliance.

Table 23: VCF 3.5 installation steps with ThinkAgile VX

#	SDDC Deployment Sequence	Deployed Component/Feature	Comment
1	Run VX Appliance on Lenovo ThinkAgile VX Servers	Deploy ESXi, vSAN, vCenter	Management cluster and compute cluster *(Lenovo XClarity Administrator can also be used to image ESXi servers manually)*
2	Install Cloud Builder Appliance		
3	Run Cloud Builder to deploy the SDDC manager and the Management Workload Domain	SDDC Manager, NSX-T Manager	Input file is used with all configured parameter
4	Deploy VI Workload Domain(s) with SDDC Manager	vCenter, NSX-T	
5	Deploy Aria Suite (Management Workload Domain)	Aria Operations, Aria Suite Lifecycle Manager, Aria Automation, Aria Load balancers (NSX Edges)	

6	Deploy Tanzu using VCF	Create Edge services, Tanzu Supervisor Cluster and Kubernetes Cluster	

7.3.1.1 Run VX Appliance on Lenovo ThinkAgile VX Servers

https://thinkagile.lenovofiles.com/help/index.jsp?topic=%2Fcom.lenovo.thinkagile.vx%2Fcluster_deployment_with_vx_deployer.html

7.3.1.2 Deploy Cloud Builder Appliance

- Download the VMware Cloud Builder .ova file from VMware Product Download website
- Deploy the .ova file in vCenter:

Select OVF Template (VMware-Cloud-Builder-4.3.1.0-18624509_OVF10.ova)

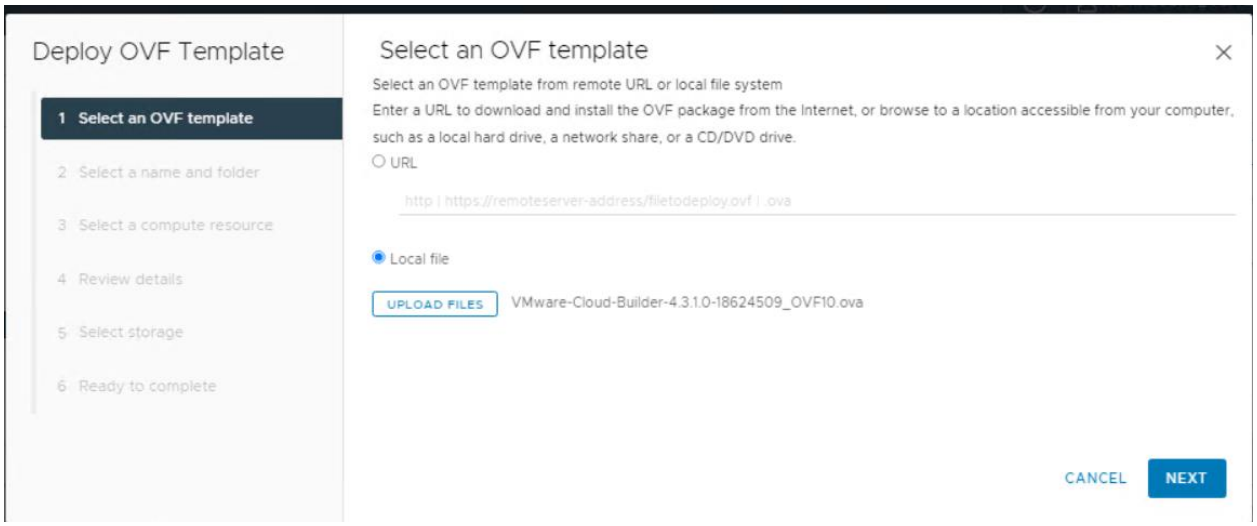

Select a Name and a Folder:

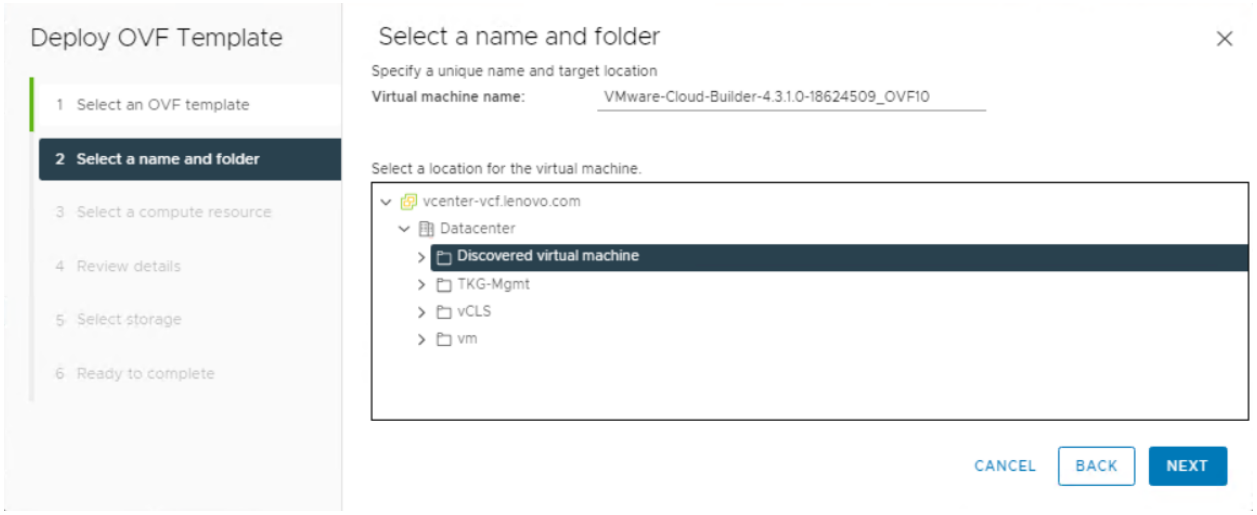

Select a Compute resource:

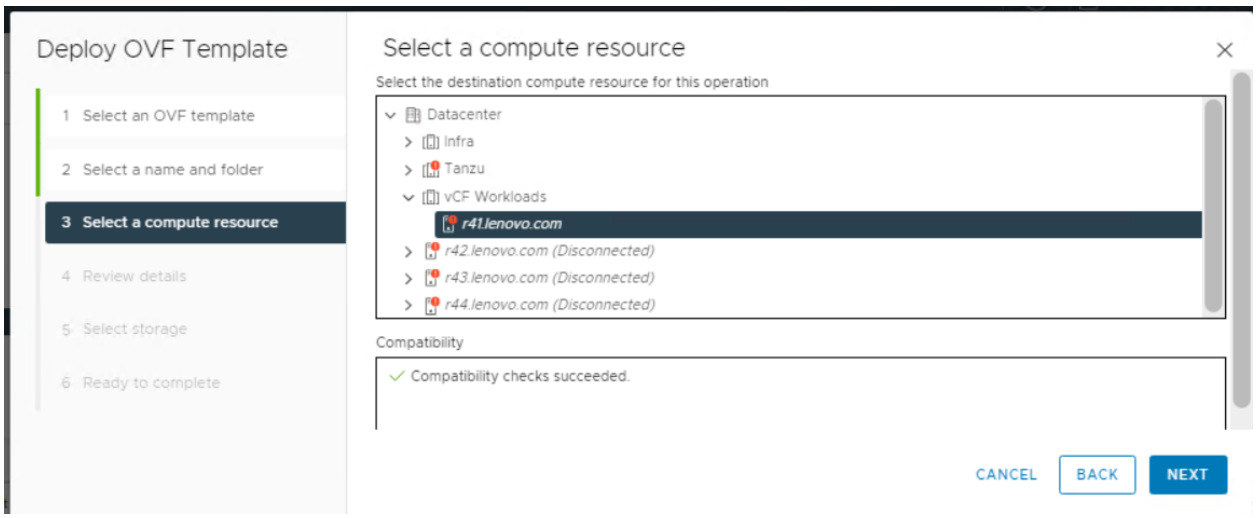

Review details:

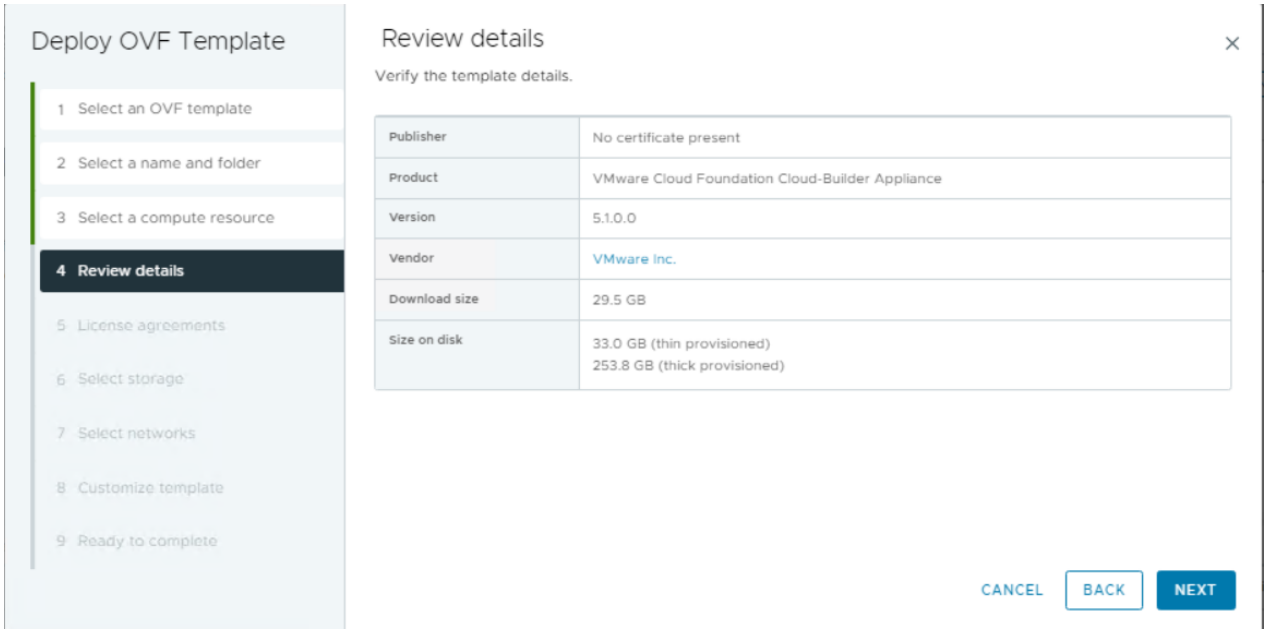

Accept all license agreements:

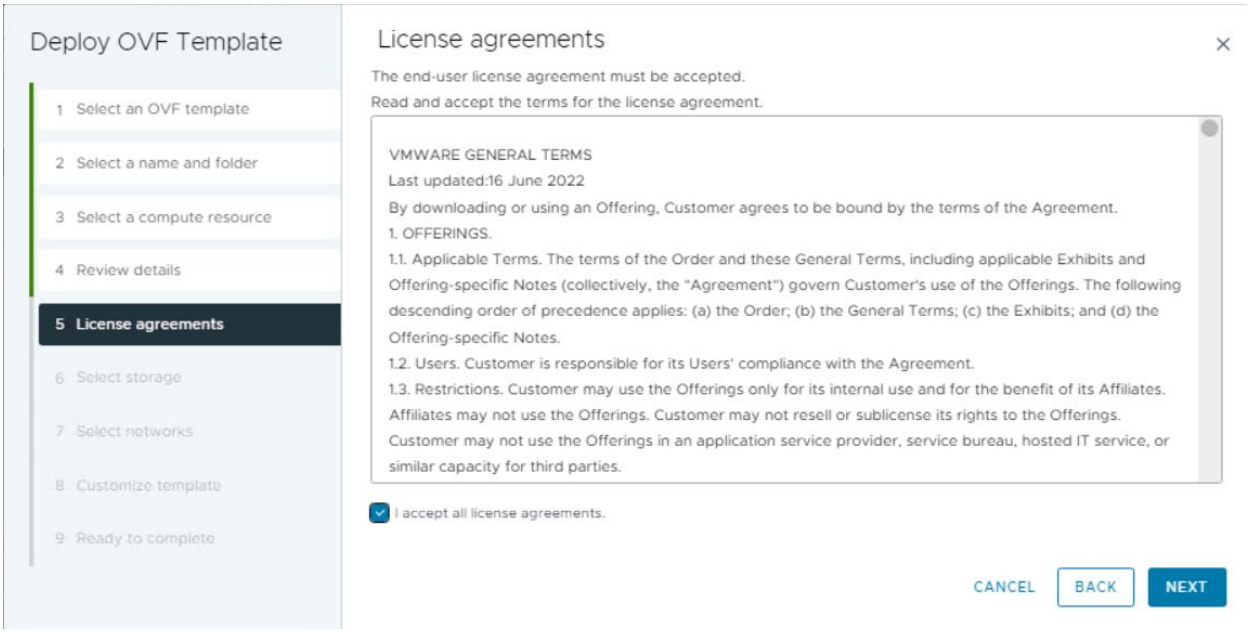

Select a valid storage:

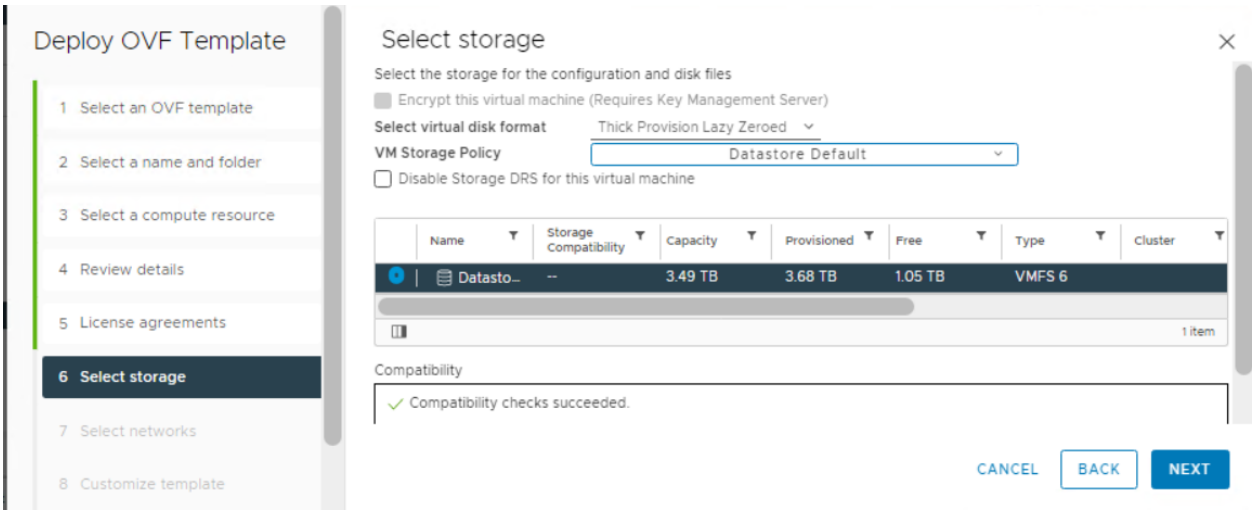

Select a network:

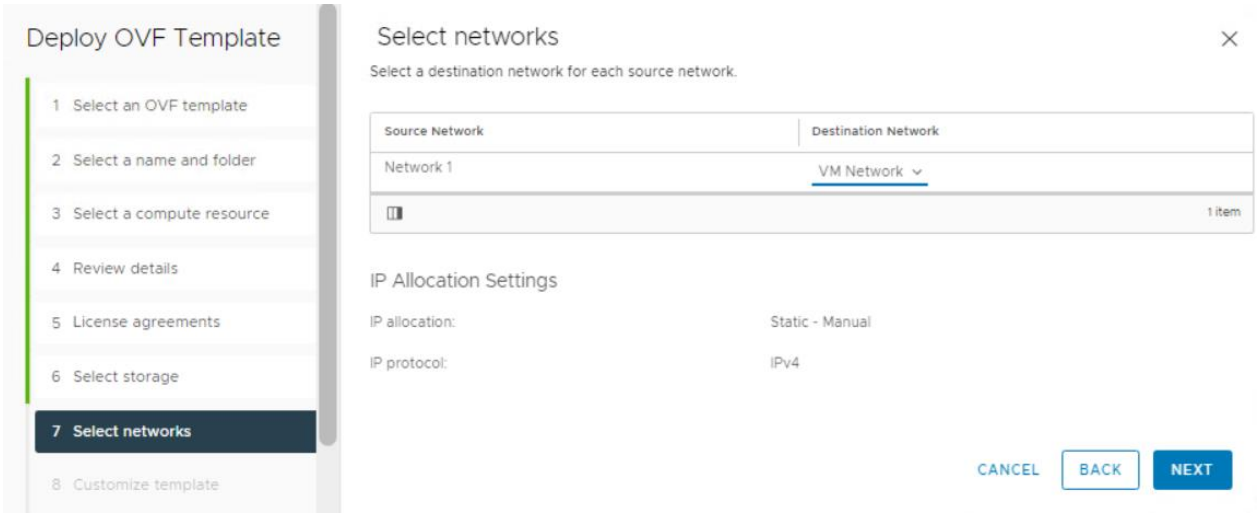

- Customize the Cloud Builder template:
  - Set the admin password
  - Set the root password
  - Set the hostname of the appliance
  - Enter an IP address for the network interface
  - Enter the subnet mask
  - Enter the default gateway
  - Enter the DNS server(s) (maximum 2 entries)
  - Enter the DNS Domain Name
  - Enter the DNS Domain Search Path (optional)
  - Enter the NTP Server

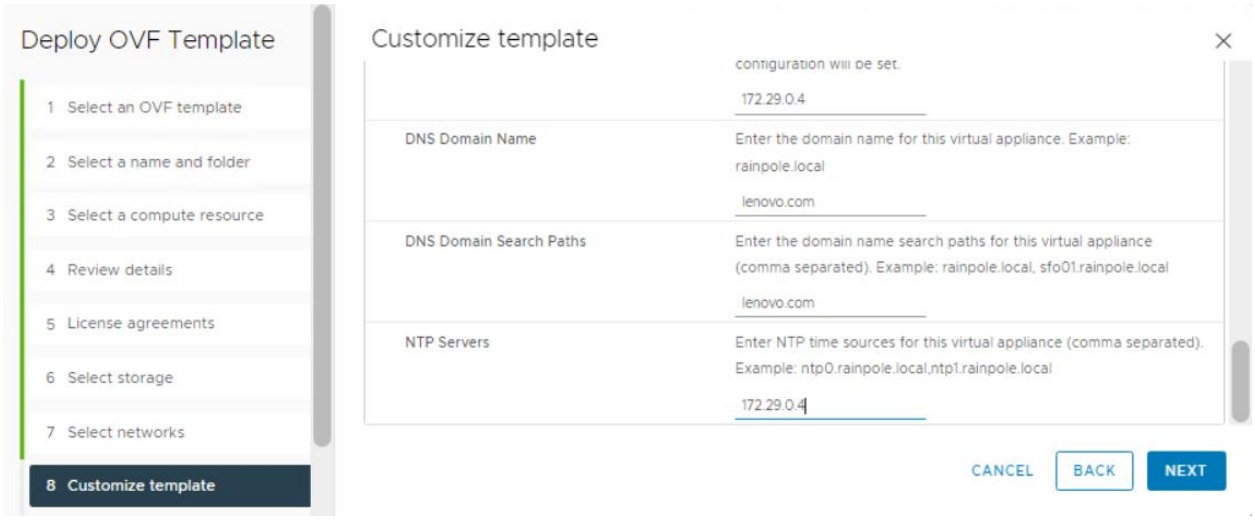

Review the settings and finish the customization:

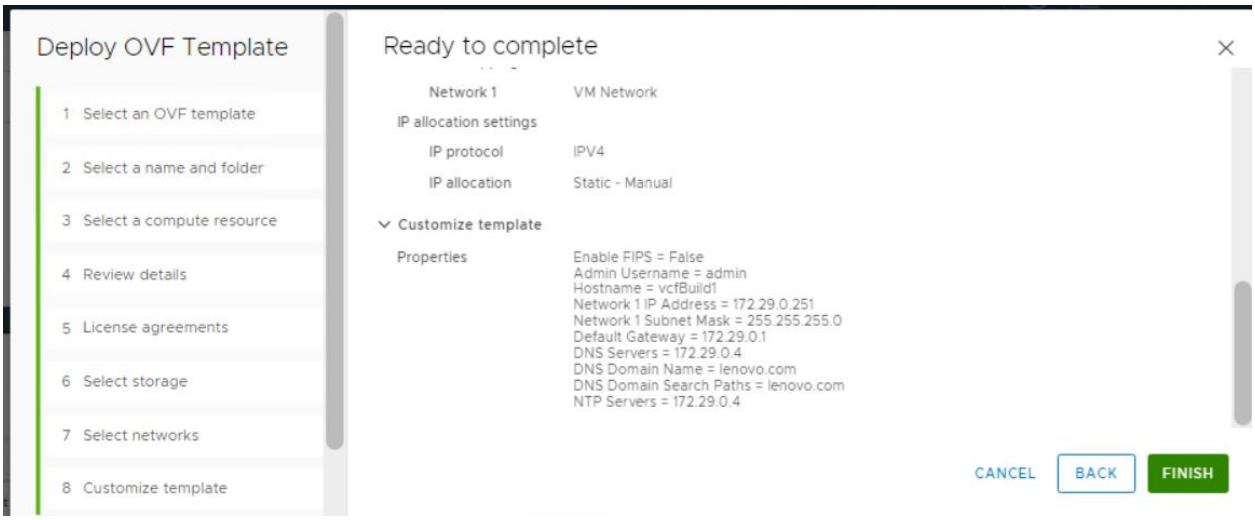

After the Cloud Builder VM has been deployed:

- login to each host and regenerate the self-signed certificates using these commands (mandatory if the FQDNs have been changed):

  #/sbin/generate-certificates

  #reboot

- login to the appliance using the admin user and password and obtain the Security Thumbprints for each ESXi hosts):

  #ssh-keygen -lf <(ssh-keyscan **hostname** 2>/dev/null)

#openssl s_client -connect **hostname**:443 < /dev/null 2> /dev/null | openssl x509 -sha256 -fingerprint -noout -in /dev/stdin

```
root@vcfBuild [ ~ ]# ssh-keygen -lf <(ssh-keyscan r05.lenovo.com 2>/dev/null)
3072 SHA256:Om5Utfs/SRzKIY238QtcPVAQnADWtPMjcD2yz8f/5IE r05.lenovo.com (RSA)
256 SHA256:FhwjgX26LtTwwCBU4UyhpaaprjPRqCCczzeEFt7XuBc r05.lenovo.com (ECDSA)
root@vcfBuild [ ~ ]# openssl s_client -connect r05.lenovo.com:443 < /dev/null 2> /dev/null | openssl x509 -sha256 -fingerprint -noou
t -in /dev/stdin
SHA256 Fingerprint=67:CC:28:AE:D9:49:AA:7A:80:CE:53:58:9C:66:FC:3F:D6:A1:70:58:2F:CE:99:54:25:69:32:98:4E:2D:63:65
root@vcfBuild [ ~ ]#
```

7.3.1.3 Run Cloud Builder to deploy SDDC Manager and the Management Workload Domain

Login to the Cloud Builder web interface (https://<Cloud_Builder_IP>) with the user admin and password configured during the CB deployment:

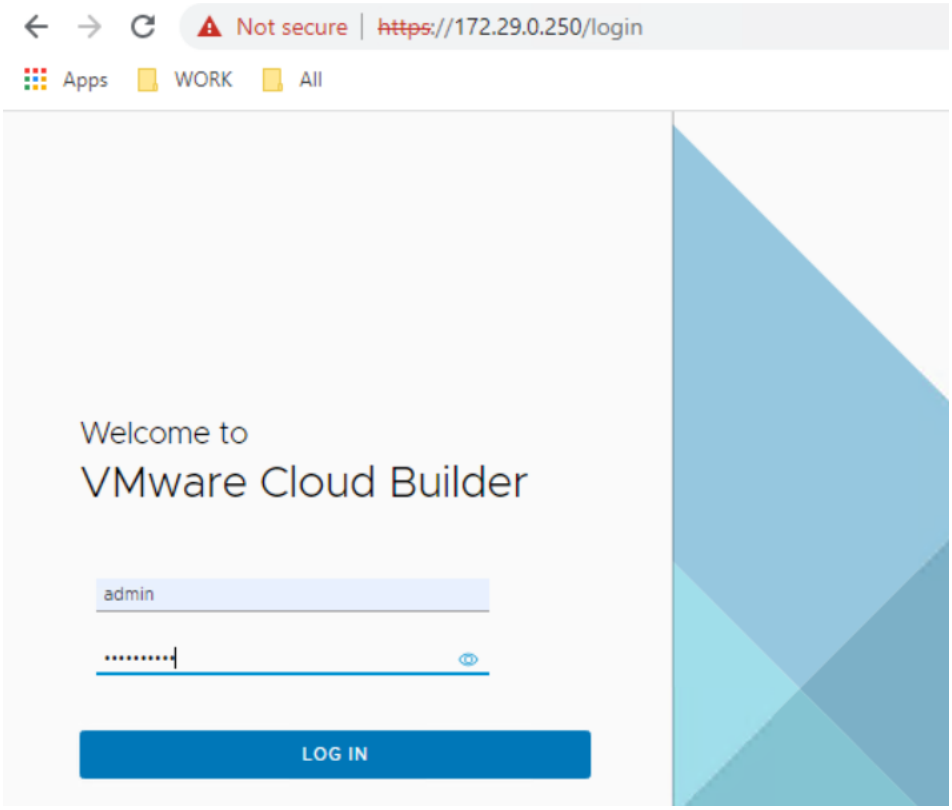

Accept the EULA:

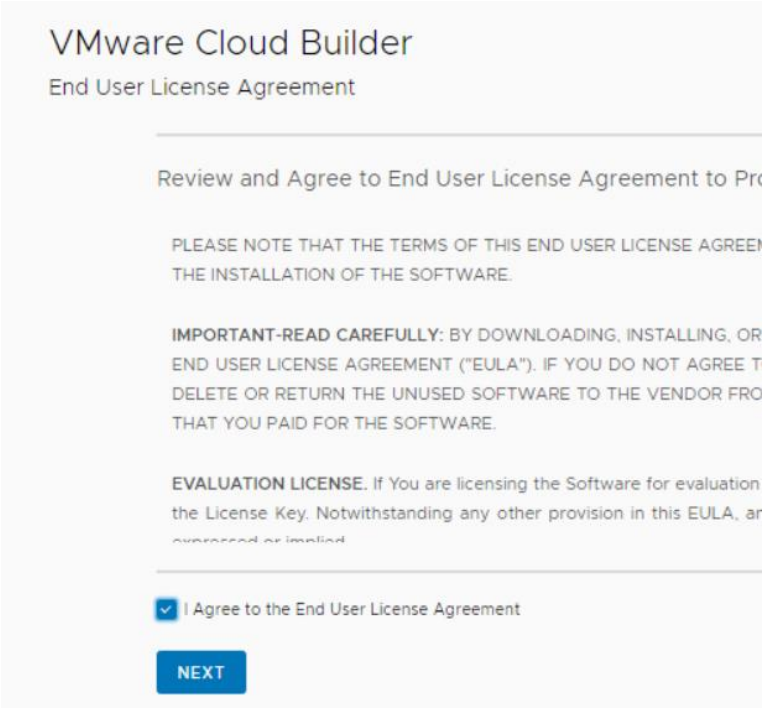

Select VMware Cloud Foundation:

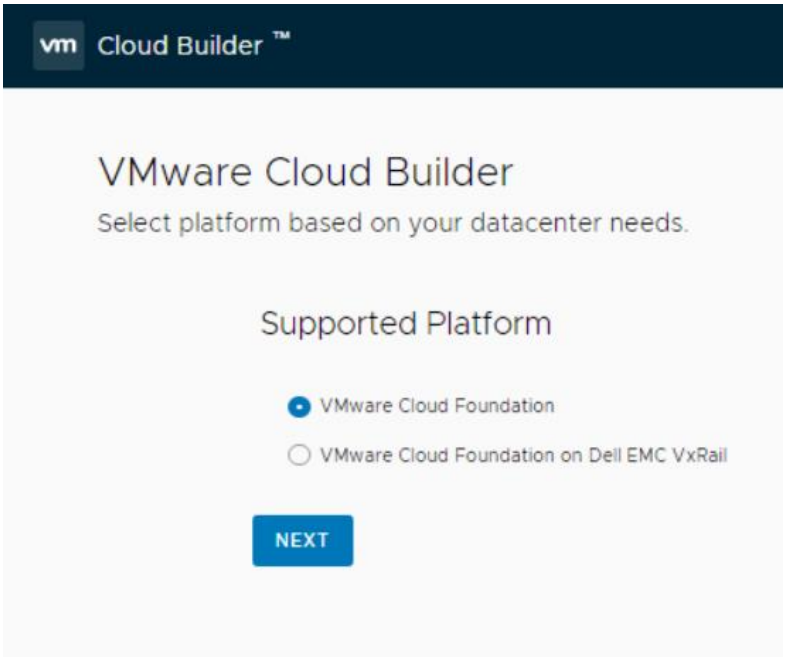

Review and confirm the prerequisites:

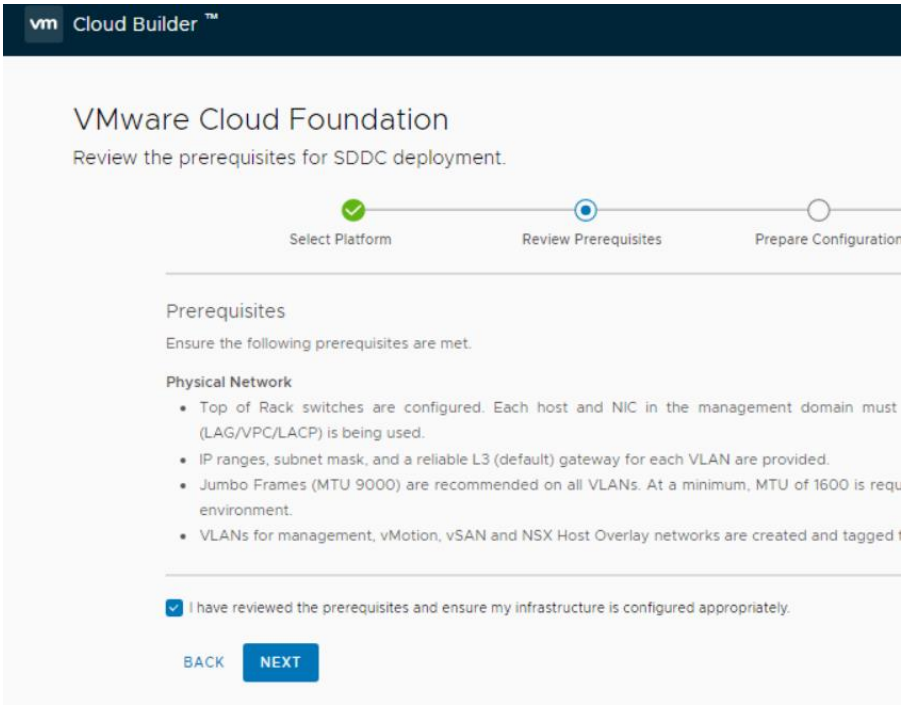

Download the Workbook from VMware Customer Connect (vcf-ems-deployment-parameters.xlsx) then click on Next:

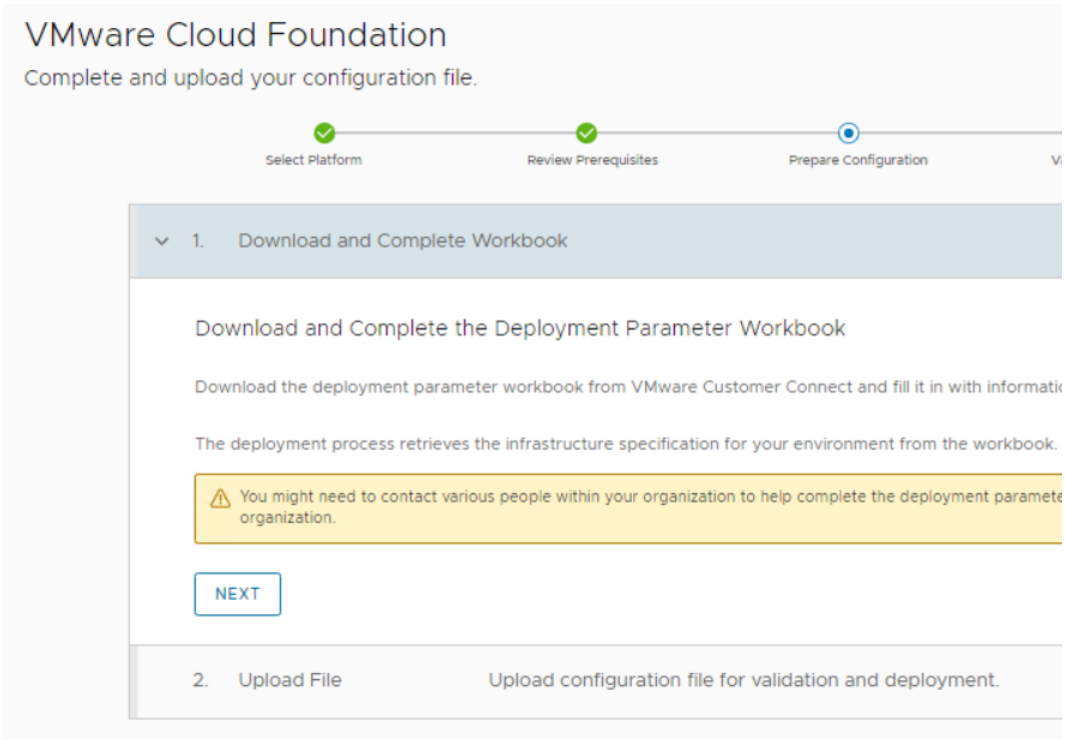

- Configure the vcf-ems-deployment-parameter.xlsx for Cloud Builder

On Credentials tab setup complex passwords for all the components:

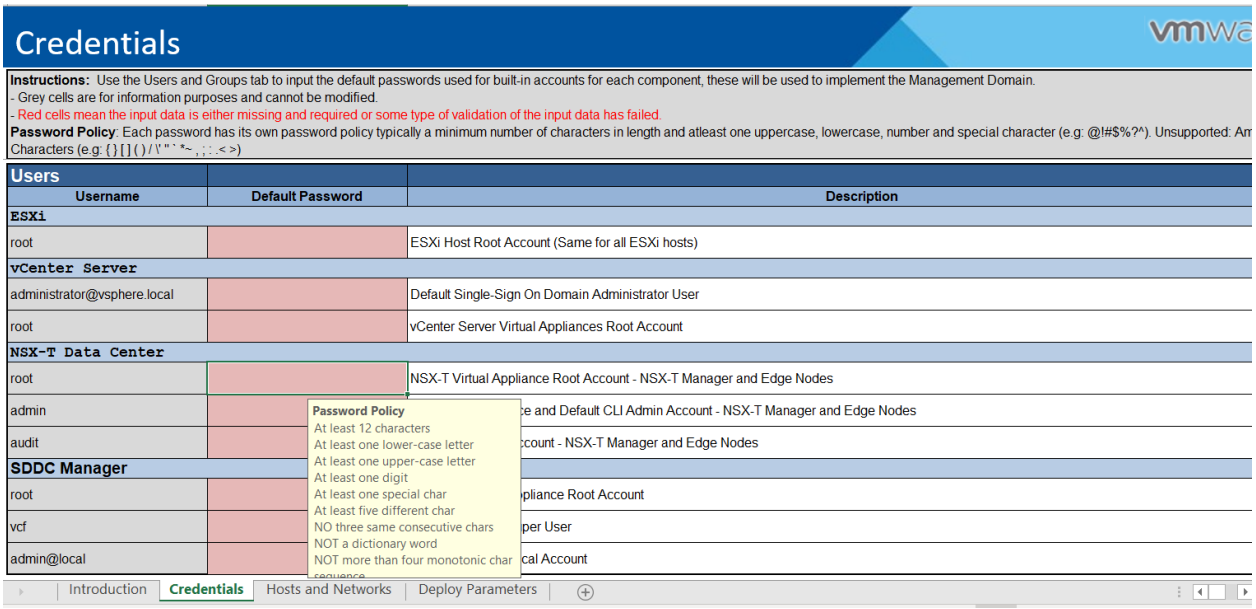

On Hosts and Networks tab fill in the following parameters:

Management Domain Networks:

- Management VLAN ID (ESXi hosts management network)
- VM Management Network VLAN ID (vCenter, SDDC Manager, NSX-T VMs etc. network)
- vMotion Network VLAN ID
- vSAN Network VLAN ID
- Portgroup Name for Management VLAN
- Portgroup Name for vMotion VLAN
- Portgroup Name for vSAN VLAN
- CIDR Notation for Management VLAN (Subnet)
- CIDR Notation for vMotion VLAN (Subnet)
- CIDR Notation for vSAN VLAN (Subnet)
- Gateways for Management, vMotion and vSAN VLANs (subnets) - configured on the Physical Switch
- MTUs for Management, vMotion and vSAN - configured on the Physical Switch
- vSphere Standard Switch Name (already configured in vCenter)
- Primary vSphere Distributed Switch – Name (this will be configured during the vCF deployment)
- Primary vSphere Distributed Switch – pNICs (already configured on the ESXi hosts)
- Primary vSphere Distributed Switch - MTU Size (this will be configured during the vCF deployment)
- vSphere Distributed Switch Profile – Profile-1 (2 physical interfaces for each ESXi host)

Management Domain ESXi Hosts:

- Hostname Length (FQDN of the Management ESXi Hosts)
- IPs for each ESXi Management Host
- vMotion Start IP address (included) for ESXi Management Hosts
- vMotion End IP address (included) for ESXi Management Hosts
- vSAN Start IP address (included) for ESXi Management Hosts
- vSAN End IP address (included) for ESXi Management Hosts
- Security Thumbprints – YES (These can be filled in after the Cloud Builder VM is deployed):

SSH RSA Key Fingerprints (SHA256)

Login to the Cloud Builder VM and issue the following command for each <hostname> (ESXi FQDN):

#ssh-keygen -lf <(ssh-keyscan **hostname** 2>/dev/null)

SSL Thumbprints (SHA256)

Login to the Cloud Builder VM and issue the following command for each <hostname> (ESXi FQDN):

#openssl s_client -connect **hostname**:443 < /dev/null 2> /dev/null | openssl x509 -sha256 -fingerprint -noout -in /dev/stdin

NSX-T Host Overlay Network – Static IP Pool in NSX-T:

- VLAN ID – Configured on the Physical Switch
- Configure NSX-T Host Overlay Using a Static IP Pool – Yes
- Pool Description
- Pool Name
- CIDR Notation (subnet)
- Gateway
- NSX-T Host Overlay Start IP
- NSX-T Host Overlay End IP

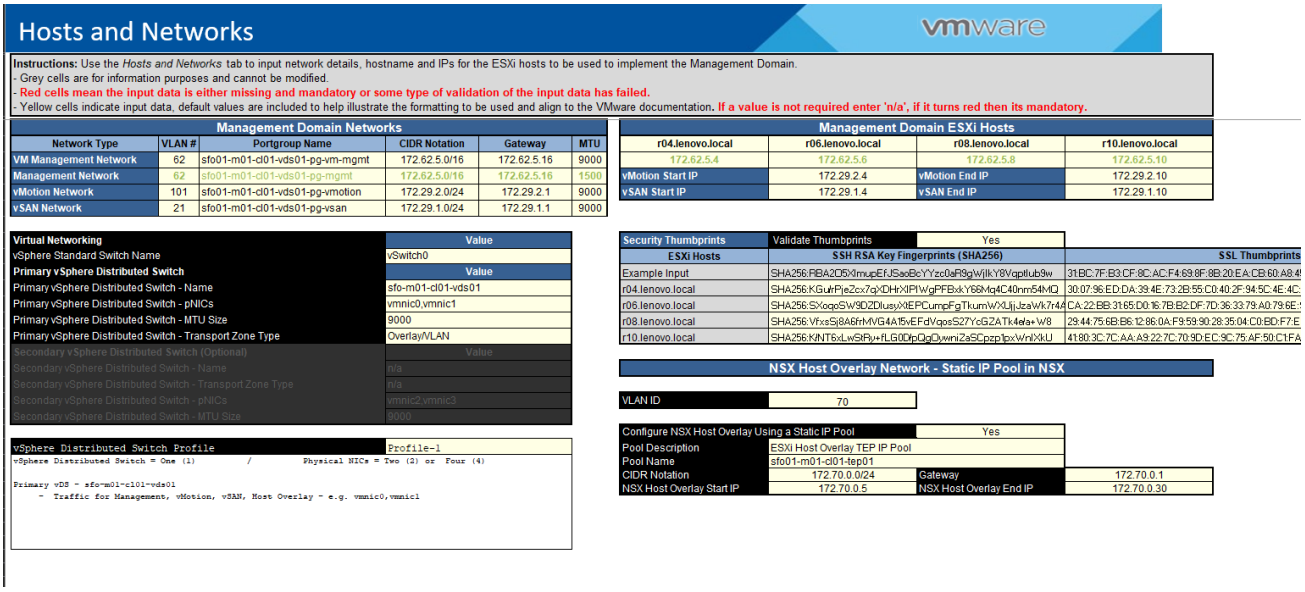

Hosts and Networks

vmware

Instructions: Use the *Hosts and Networks* tab to input network details, hostname and IPs for the ESXi hosts to be used to implement the Management Domain.
- Grey cells are for information purposes and cannot be modified.
- Red cells mean the input data is either missing and mandatory or some type of validation of the input data has failed.
- Yellow cells indicate input data, default values are included to help illustrate the formatting to be used and align to the VMware documentation. If a value is not required enter 'n/a', if it turns red then its mandatory.

Management Domain Networks

Network Type	VLAN #	Portgroup Name	CIDR Notation	Gateway	MTU
VM Management Network	62	sfo01-m01-cl01-vds01-pg-vm-mgmt	172.62.5.0/16	172.62.5.16	9000
Management Network	62	sfo01-m01-cl01-vds01-pg-mgmt	172.62.5.0/16	172.62.5.16	1500
vMotion Network	101	sfo01-m01-cl01-vds01-pg-vmotion	172.29.2.0/24	172.29.2.1	9000
vSAN Network	21	sfo01-m01-cl01-vds01-pg-vsan	172.29.1.0/24	172.29.1.1	9000

Management Domain ESXi Hosts

r04.lenovo.local	r06.lenovo.local	r08.lenovo.local	r10.lenovo.local
172.62.5.4	172.62.5.6	172.62.5.8	172.62.5.10
vMotion Start IP	172.29.2.4	vMotion End IP	172.29.2.10
vSAN Start IP	172.29.1.4	vSAN End IP	172.29.1.10

Virtual Networking	Value
vSphere Standard Switch Name	vSwitch0
Primary vSphere Distributed Switch	**Value**
Primary vSphere Distributed Switch - Name	sfo-m01-cl01-vds01
Primary vSphere Distributed Switch - pNICs	vmnic0,vmnic1
Primary vSphere Distributed Switch - MTU Size	9000
Primary vSphere Distributed Switch - Transport Zone Type	Overlay/VLAN
Secondary vSphere Distributed Switch (Optional)	Value
Secondary vSphere Distributed Switch - Name	n/a
Secondary vSphere Distributed Switch - Transport Zone Type	n/a
Secondary vSphere Distributed Switch - pNICs	vmnic2,vmnic3
Secondary vSphere Distributed Switch - MTU Size	9000

vSphere Distributed Switch Profile	Profile-1

vSphere Distributed Switch = One (1) / Physical NICs = Two (2) or Four (4)

Primary vDS = sfo-m01-cl01-vds01
- Traffic for Management, vMotion, vSAN, Host Overlay - e.g. vmnic0,vmnic1

Security Thumbprints	Validate Thumbprints	Yes	
ESXi Hosts	SSH RSA Key Fingerprints (SHA256)		SSL Thumbprints
Example Input	SHA256:RBA2O5VdmupEfJSaoBcYYzc0aR9gWjIkY8Vqptlub9w		31:BC:7F:B3:CF:8C:AC:F4:69:8F:9B:20:EA:CB:60:A8:4
r04.lenovo.local	SHA256:KGufPje2cx7qXDHrXIP1WgPFBxkY66Mq4C40nm54MQ		30:07:96:ED:DA:39:4E:73:2B:55:C0:40:2F:94:5C:4E:4C:
r06.lenovo.local	SHA256:SXoqoSW9D2DIusyX0EPCumpFgTkumWXLjjJzaWk7r44		CA:22:BB:31:65:D0:16:7B:B2:DF:7D:36:33:79:A0:79:6E:
r08.lenovo.local	SHA256:VfxsSj8A6frMVG4A15vEFdVqosS27YcGZATk4eIa+W8		29:44:75:6B:B6:12:86:0A:F9:59:90:28:35:04:C0:BD:F7:E
r10.lenovo.local	SHA256:KfNT6xLwStRy+fLG0DpQgDywni2aSCpzp1pxWnlXkU		41:80:3C:7C:AA:A9:22:7C:70:9D:EC:9C:75:AF:50:C1:FA

NSX Host Overlay Network - Static IP Pool in NSX

VLAN ID	70

Configure NSX Host Overlay Using a Static IP Pool	Yes		
Pool Description	ESXi Host Overlay TEP IP Pool		
Pool Name	sfo01-m01-cl01-tep01		
CIDR Notation	172.70.0.0/24	Gateway	172.70.0.1
NSX Host Overlay Start IP	172.70.0.5	NSX Host Overlay End IP	172.70.0.30

On the Deployment Parameters tab fill in the following parameters that the Cloud Builder will use during the deployment of vCF Management Domain:

Existing Infrastructure Details:

- DNS Server and DNS Zone Defined
- NTP Servers

License Keys for

- ESXi hosts
- vSAN
- vCenter
- NSX-T Datacenter
- SDDC Manager

vSphere Infrastructure:

- vCenter Server Hostname and IP Address
- vCenter Server Appliance Size (Default Small)
- vCenter Server Appliance Storage Size (default)
- vCenter Datacenter and Cluster
- Datacenter Name
- Cluster Name
- Cluster EVC Setting
- Select the VCF Architecture to be deployed: Standard (without the NSX-T Edge Cluster)
- vSAN Datastore Name

- Enable vSAN Deduplication and Compression: No
- Join Existing Single Sign-On Domain: No

NSX-T Data Center:

- NSX-T Management Cluster VIP (virtual IP)
- NSX-T Virtual Appliance Node #1 Hostname and IP Address
- NSX-T Virtual Appliance Node #2 Hostname and IP Address
- NSX-T Virtual Appliance Node #3 Hostname and IP Address
- NSX-T Virtual Appliance Size (Default Medium)

SDDC Manager:

- SDDC Manager Hostname
- SDDC Manager IP Address
- Network Pool Name
- Cloud Foundation Management Domain Name

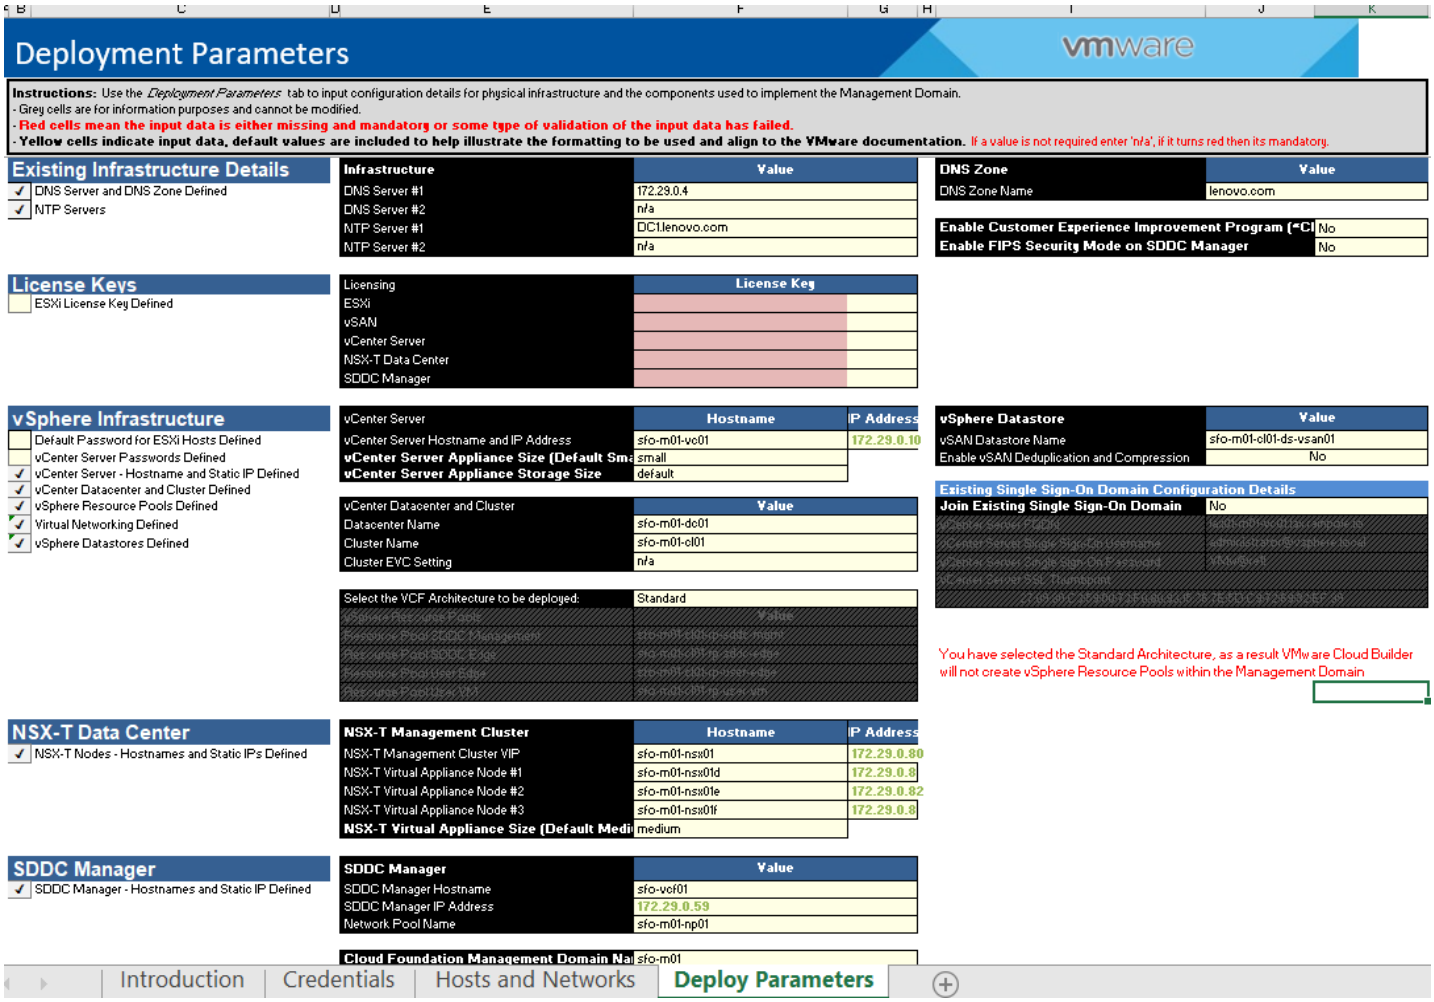

Configure DNS, DHCP

The DNS Server must have interfaces configured in the following VLANs:

- Management VLAN
- Egress VLAN
- Ingress VLAN

Create the following Host(A) entries with PTR(reverse):

- All esxi hosts (defined in vcf-ems-deployment-parameter.xlsx)
- vCenter Server Hostname and IP Address (defined in vcf-ems-deployment-parameter.xlsx)
- NSX-T Management Cluster VIP (virtual IP defined in vcf-ems-deployment-parameter.xlsx)
- NSX-T Management Cluster Nodes (3 nodes defined in vcf-ems-deployment-parameter.xlsx)
- SDDC Manager Hostname (defined in vcf-ems-deployment-parameter.xlsx)
- NSX-T Edge Cluster VIP
- NSX-T Edge Cluster Nodes (3 nodes)
- NSX-T Edge01
- NSX-T Edge02

Create a DHCP Scope for NSX-T Management Cluster and Edge Cluster according to the vcf-ems-deployment-parameter.xlsx configuration file for Cloud builder.

Add the Address Pool according to the vcf-ems-deployment-parameter.xlsx configuration file

Prerequisites (before deploying the SDDC Manager):

- Verify the current vCF version with the versions of the constituent products (Correlating VMware Cloud Foundation version with the versions of its constituent products https://kb.vmware.com/s/article/52520)
- Verify the ESXi host time configuration
  Add the following line to /etc/ntp.conf on each host and the vCF Cloud Builder VM: tos maxdist 30
- Add the Lenovo Customization Addon for Lenovo ThinkSystem (e.g. LVO.702.10.7) on the ESXi host image before deploying the updates
- Verify that VM Network VLANs on the ESXi hosts match the Parameter Workbook VLANs for Management Network
- Verify if only vmnic0 is used on each host for communication on Management VLAN
- vSwitch0 on each host must have 9000 MTU

Complete the Deployment Parameter Workbook then click on Next:

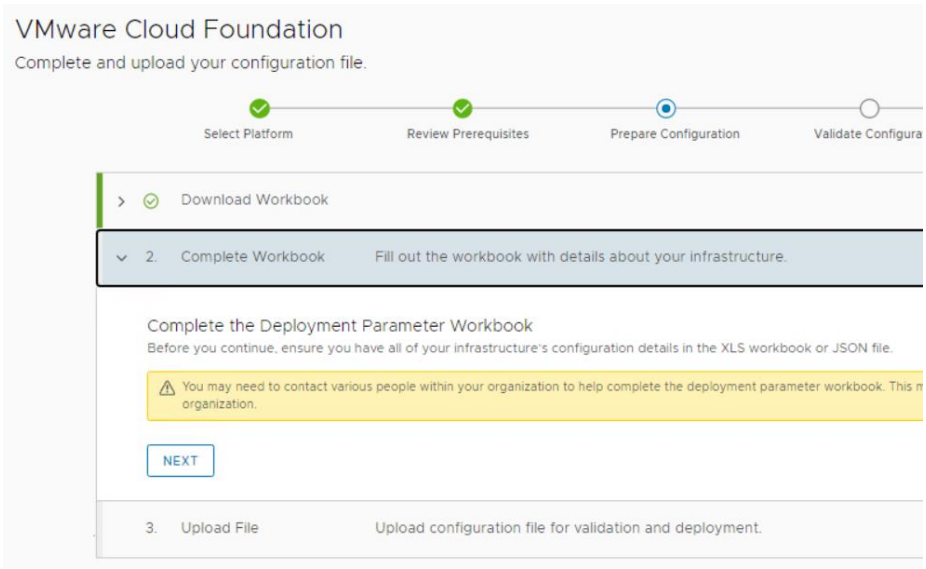

Upload the Workbook file after completing all the parameters then click on Next:

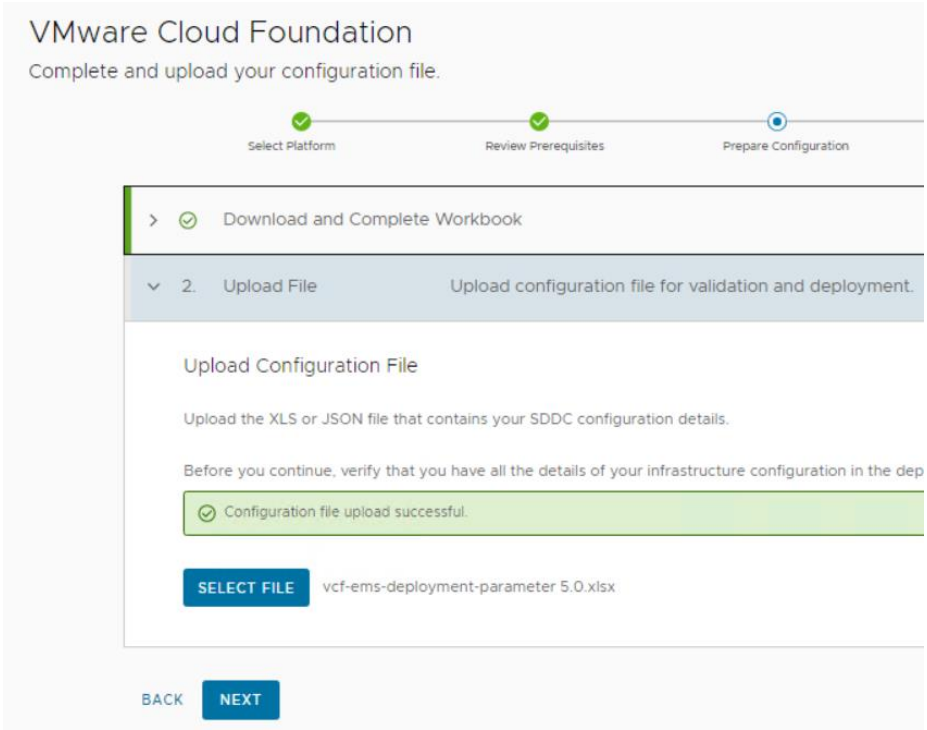

After the configuration file is successfully validated click on Next:

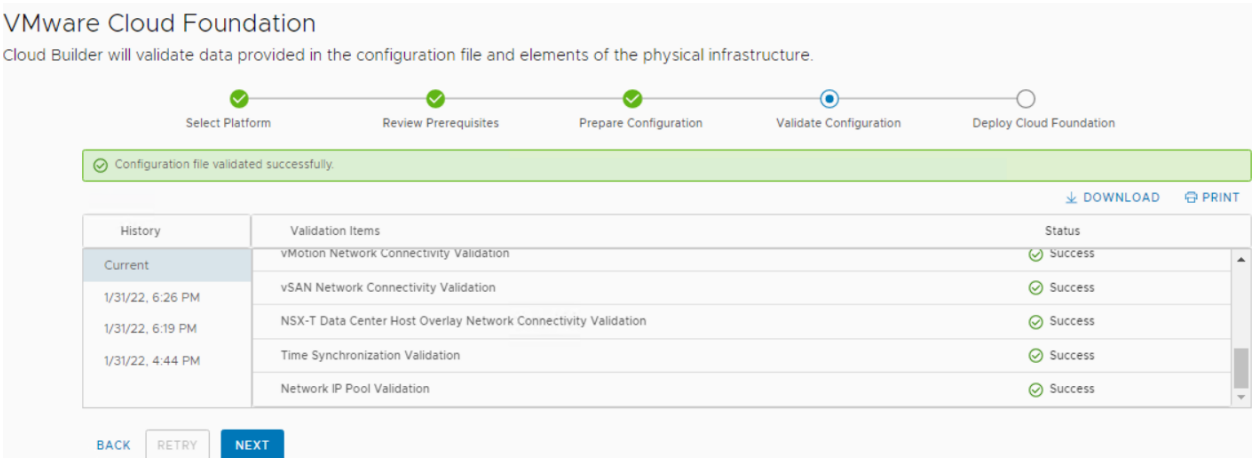

*** If any of the checks fails you can Retry the validation after fixing the issues

Click on Deploy SDDC

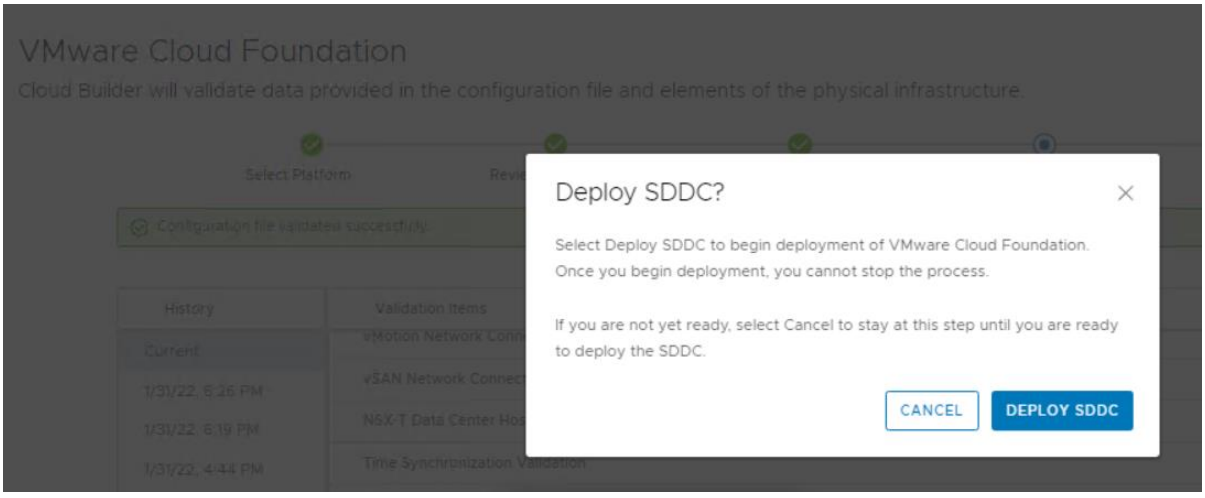

After the deployment finished successfully click Finish:

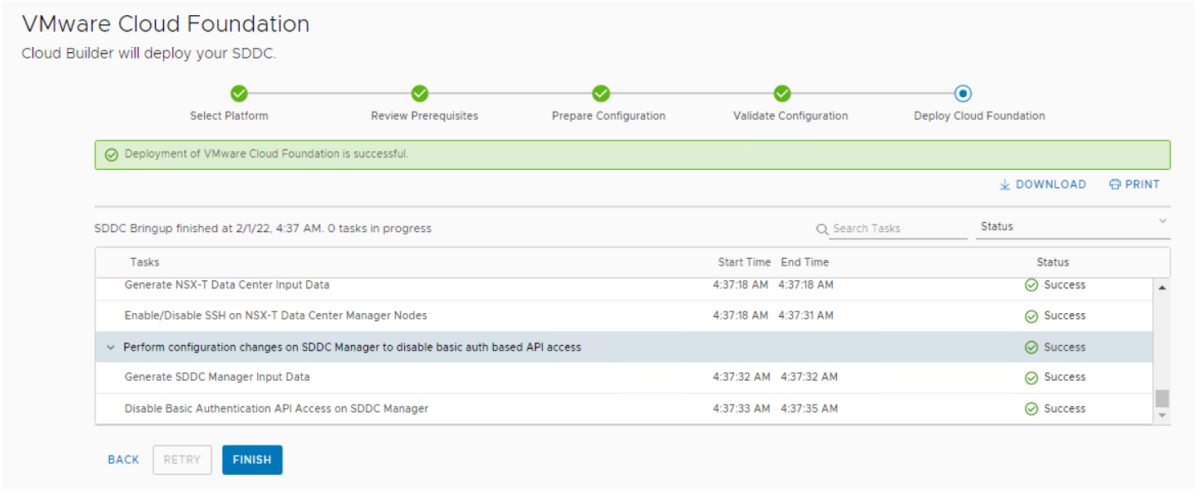

Launch SDDC Manager:

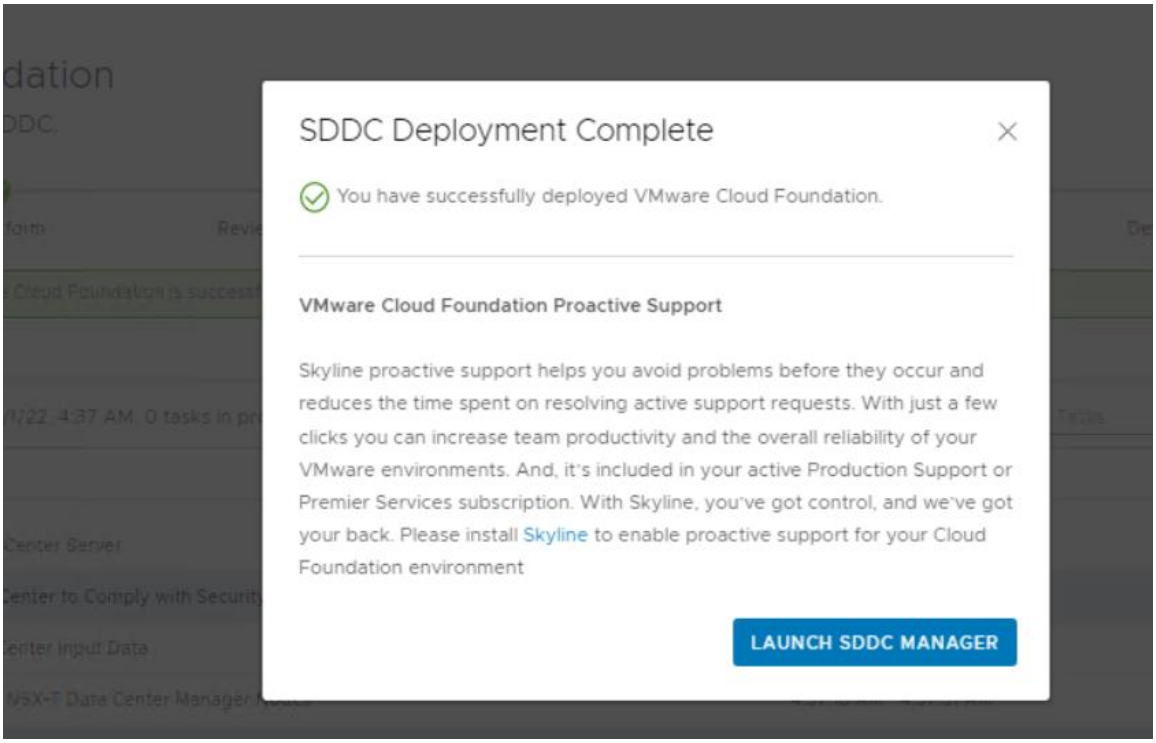

Login to the SDDC Manager using administrator@vsphere.local user and password:

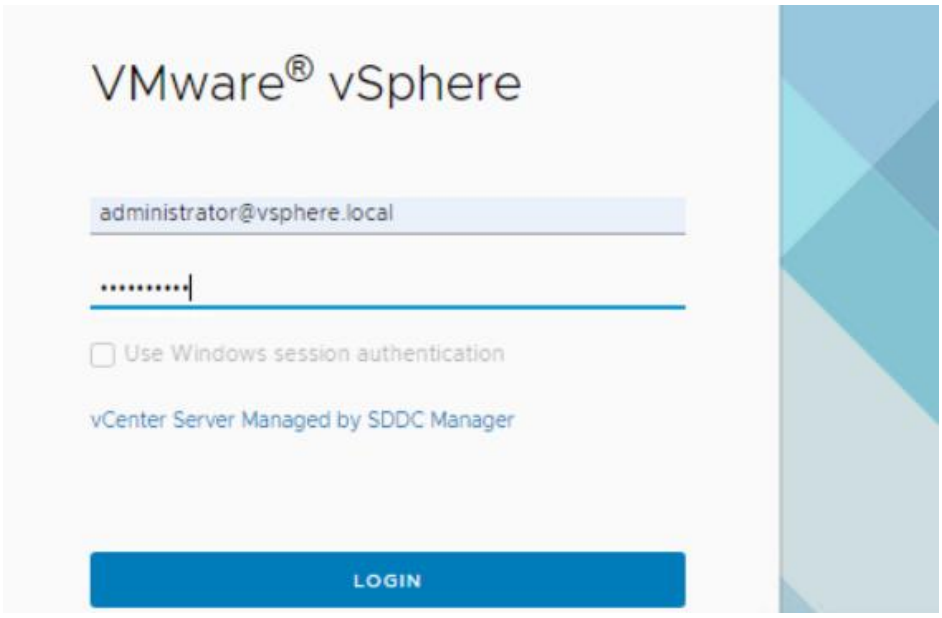

Note: If any of the components fails to install, the deployment can be restarted from the beginning after deleting the 'execution' and 'resource' tables from the PostgreSQL database on the Cloud Builder VM:

- SSH to VMware Cloud Builder appliance and connect to the PostgreSQL database using the admin user and run the commands under the root account:

  admin@vcfBuild [~]$su -

  root@vcfBuild [~]# sudo psql -U postgres -d bringup -h localhost
  bringup=# delete from execution;

  bringup=# delete from "Resource";

  bringup=# \q

- Login to the Cloud Builder web interface (https://<Cloud_Builder_IP>) with the user admin and restart the process

7.3.1.4 Deploy VI Workload Domain(s) with SDDC Manager

In order to deploy vSphere with Tanzu using SDDC Manager in a Standard configuration, a VI Workload domain must be created using a minimum of 3 ESXi hosts that must be commissioned in SDDC Manager.

- Create a new Network Pool:

In SDDC Manager go to Administration > Network Settings and click on Create Network Pool button:

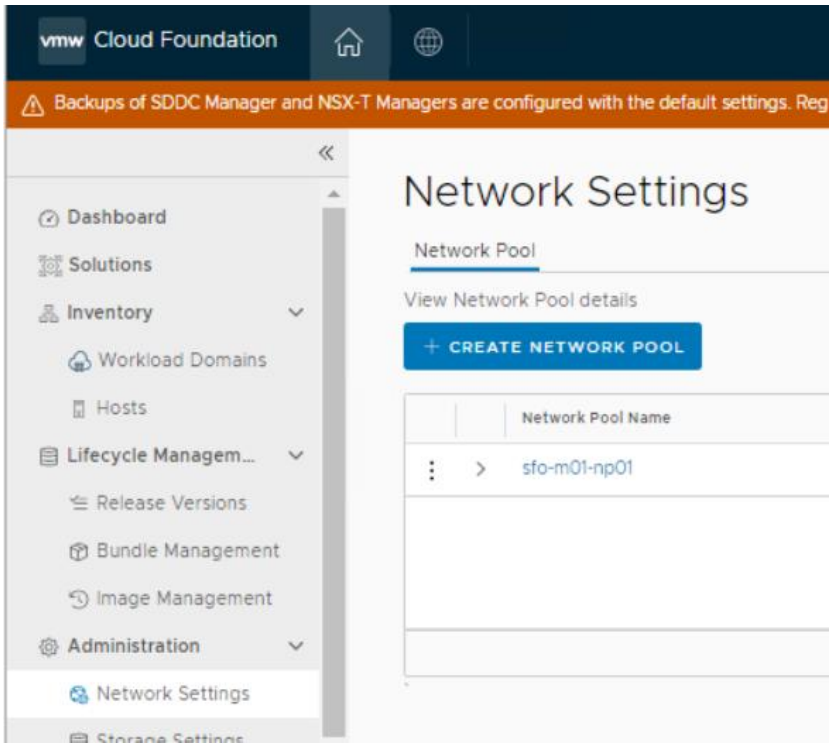

In the Create Network Pool window fill in the following parameters:

- Network Pool name
- Network Type: check VSAN and vMotion
- In the vSAN Network Information:
  - vSAN VLAN ID
  - MTU 9000
  - Network subnet
  - Subnet mask
  - Default Gateway
  - Included IP address range (to match the number of ESXi hosts to be commissioned) – click Add
- In the vMotion Network Information:
  - vMotion VLAN ID
  - MTU 9000
  - Network subnet

- Subnet mask
- Default Gateway
- Included IP address range (to match the number of ESXi hosts to be commissioned) – click Add

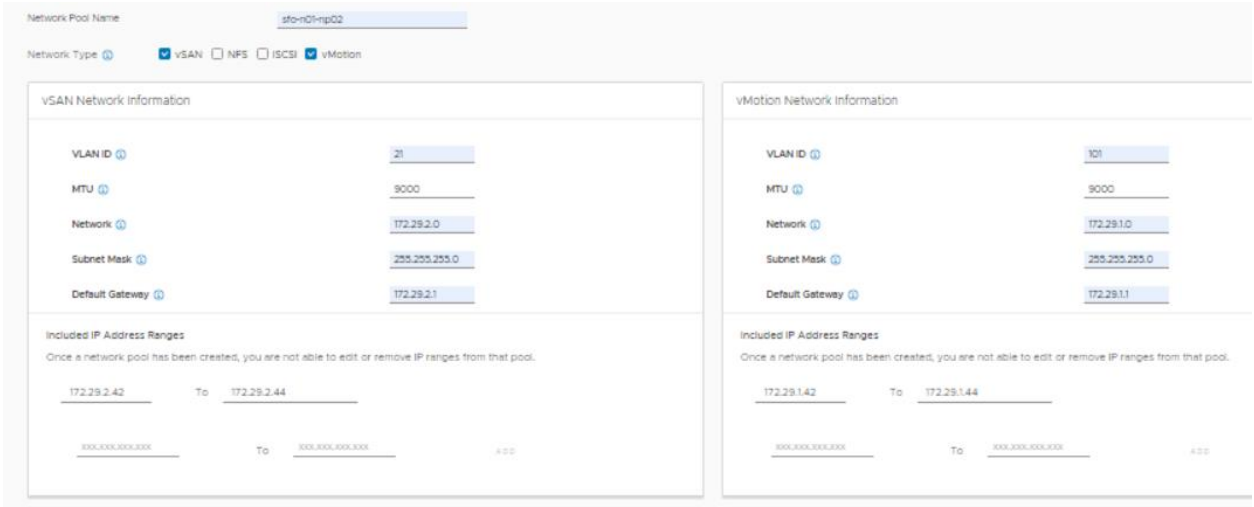

Click on Save to create the Network Pool:

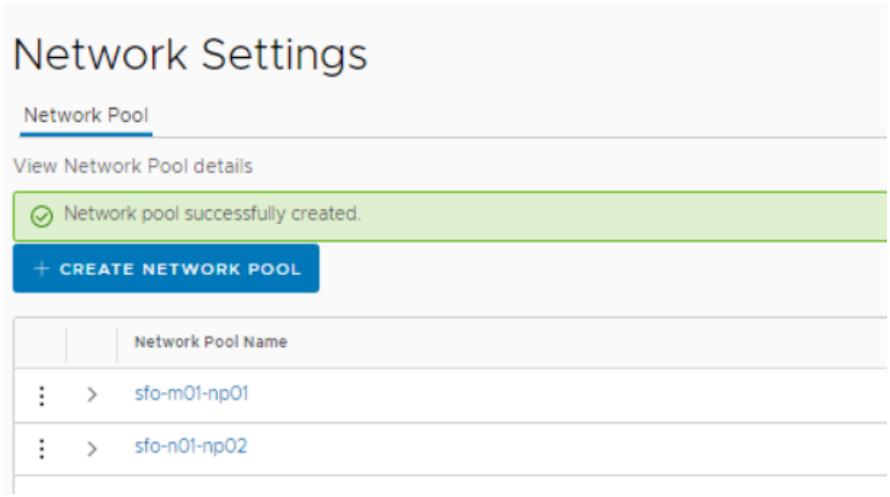

Commission the ESXi hosts

In SDDC Manager go to Inventory > Hosts and click on Commission Host button and Select all in the Checklist then click on Proceed:

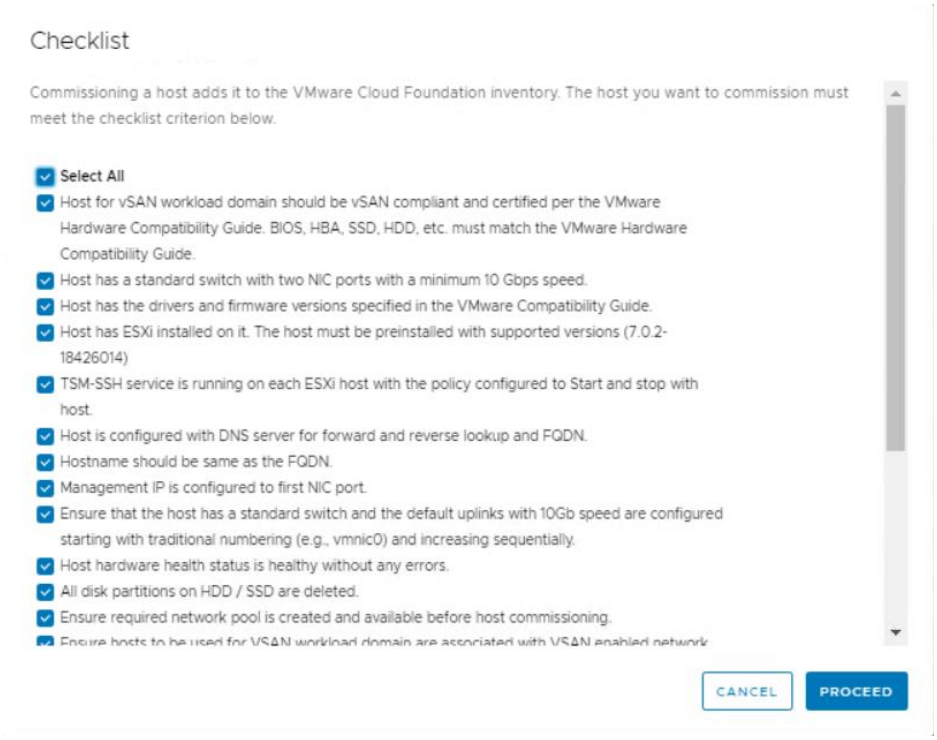

In the Commission Host windows add each ESXi Host:

- Host FQDN
- Storage Type (vSAN)
- Network Pool Name that was created in the previous steps
- Username (root) and password of the hosts

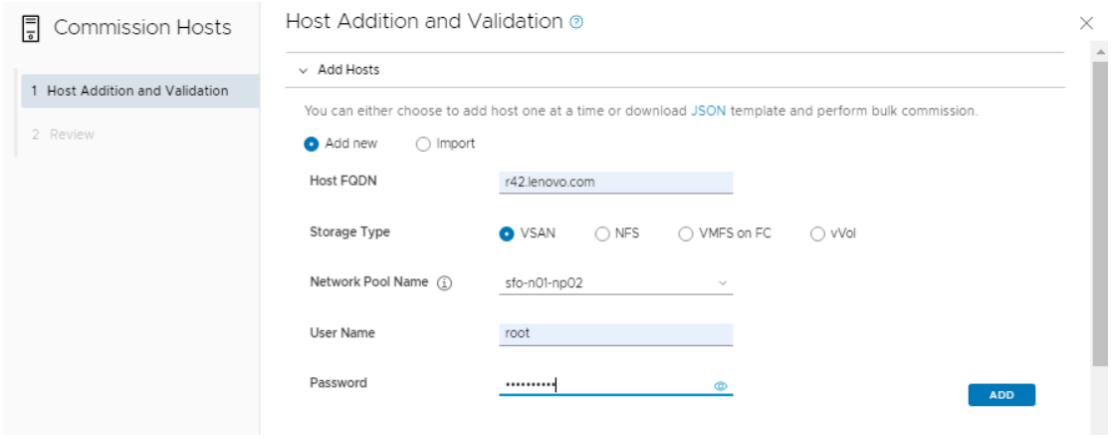

In the same window, after the hosts have been added, select all hosts and click on the Confirm Fingerprint checkbox then click on Validate ALL button:

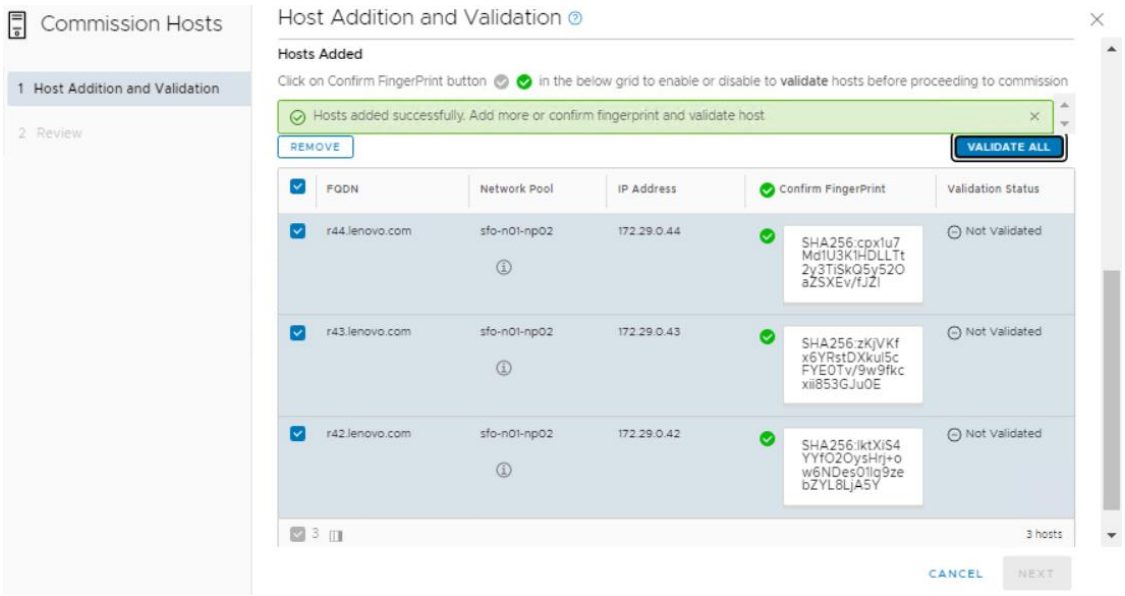

After the Hosts are successfully validated click on Next:

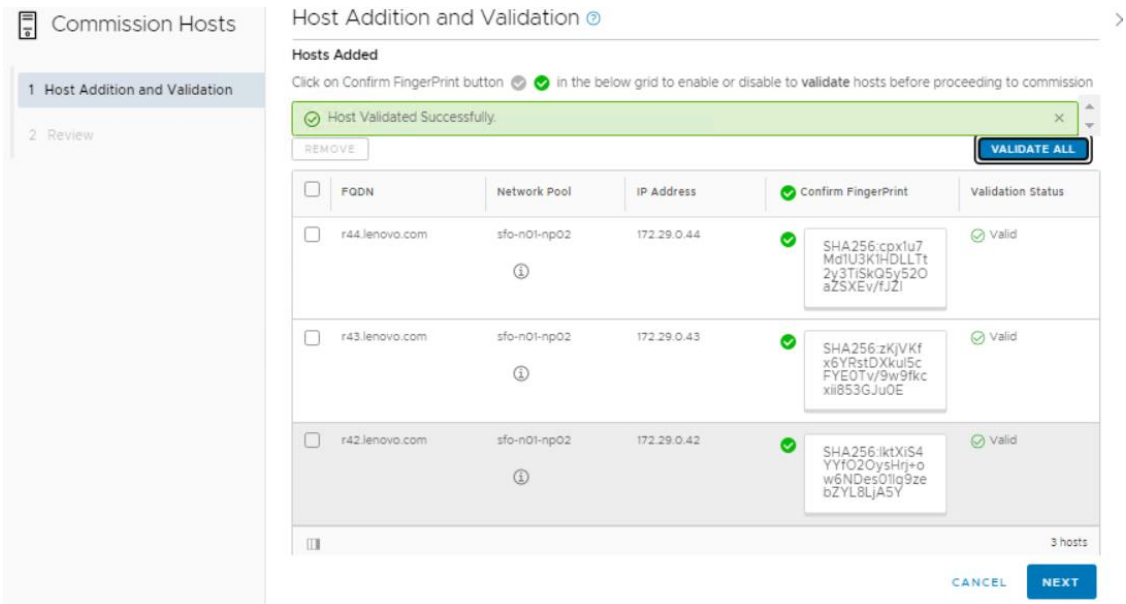

Review the Validated hosts then click on Commission button:

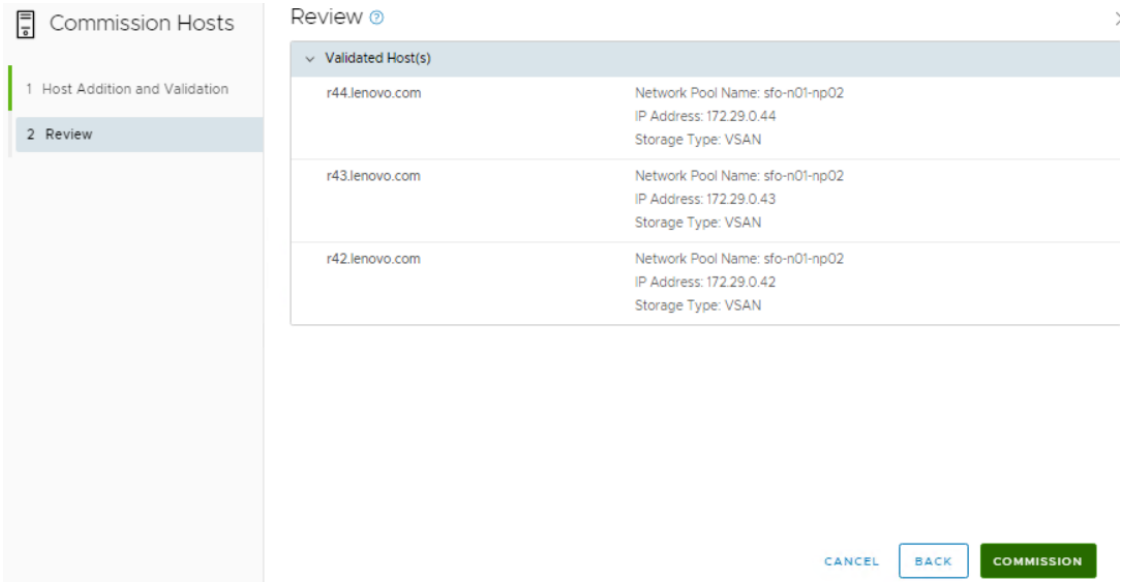

- Create the VI – Workload Domain

Add the NSX-T License in the SDDC Manager:

In the SDDC Manager go to Administration > Licensing and click on the Add License Key button:

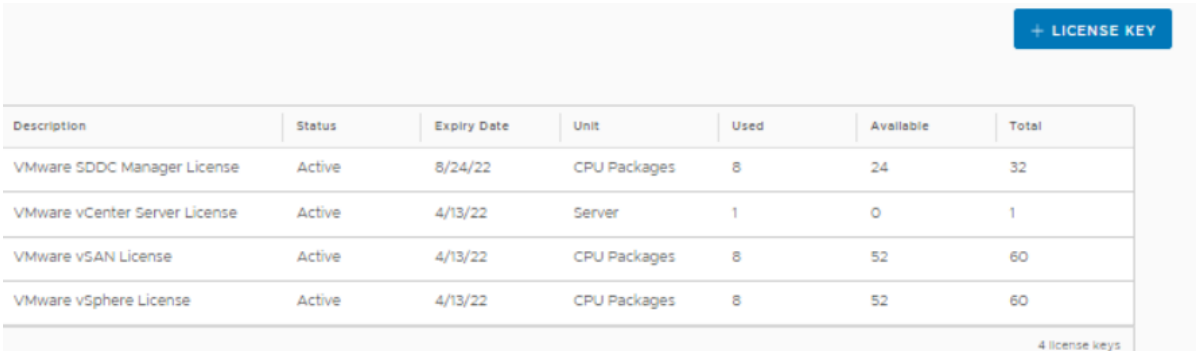

Add the appropriate NSX-T Datacenter license key:

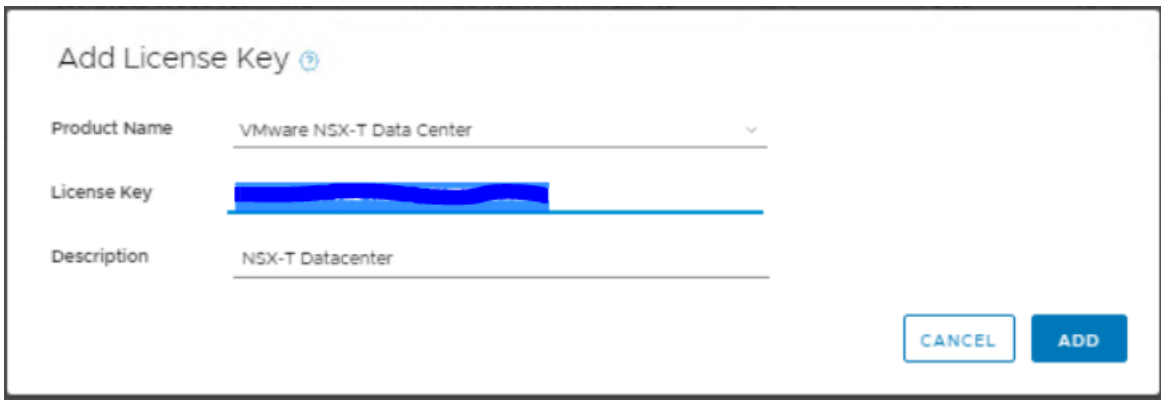

In the SDDC Manager go to Inventory > Workload Domains then click on the Workload Domain button and select VI – Workload Domain:

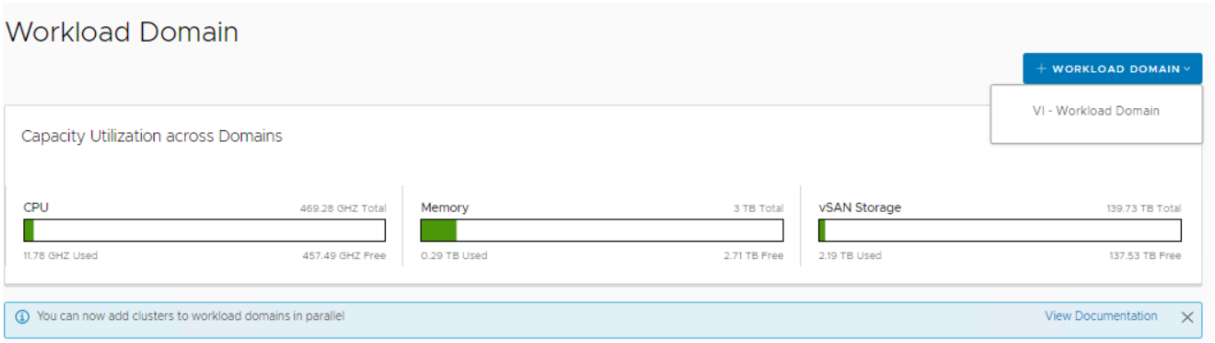

In the Storage Selection window select the appropriate Storage (vSAN) then click on the Begin button:

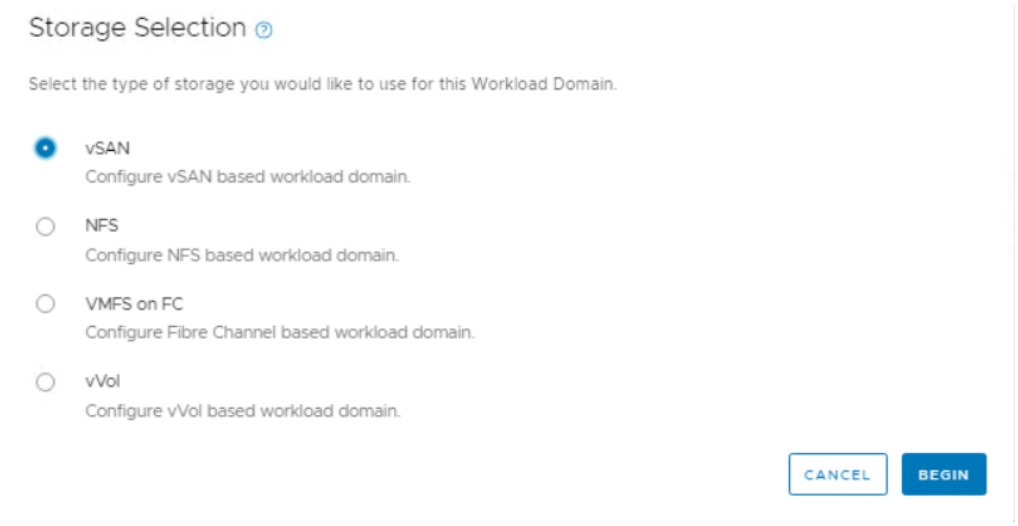

In the Name window provide a Name for the Virtual Infrastructure and the Organization, select Join Management SSO Domain then click on the Next button:

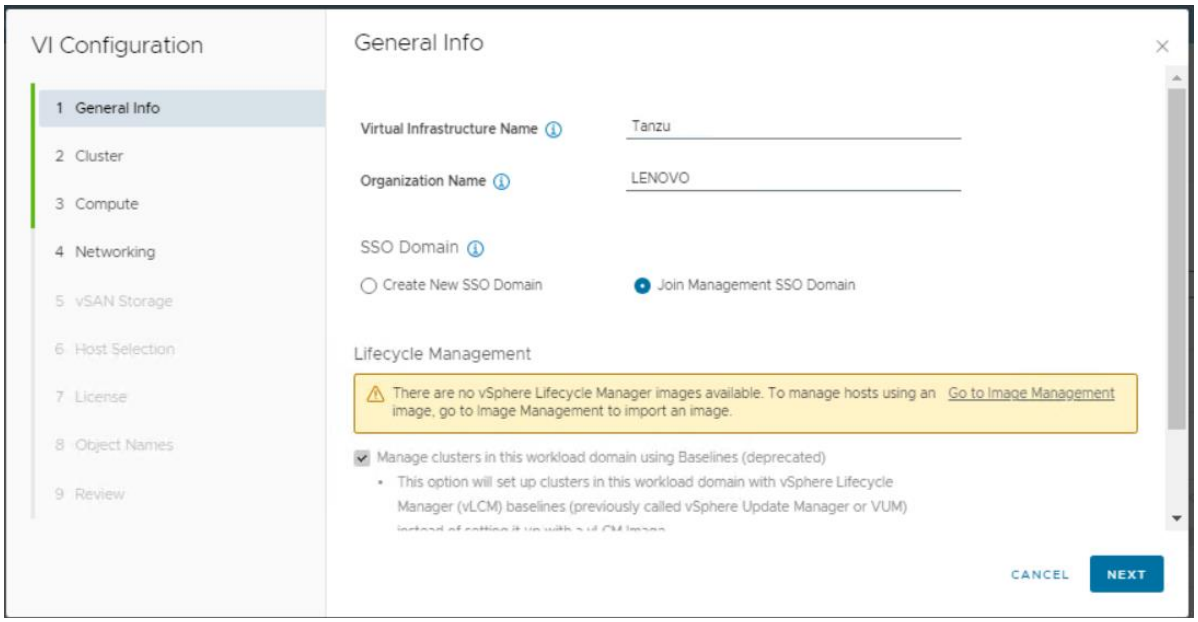

In the Cluster windows provide a Name for the Cluster, then click on the Next button:

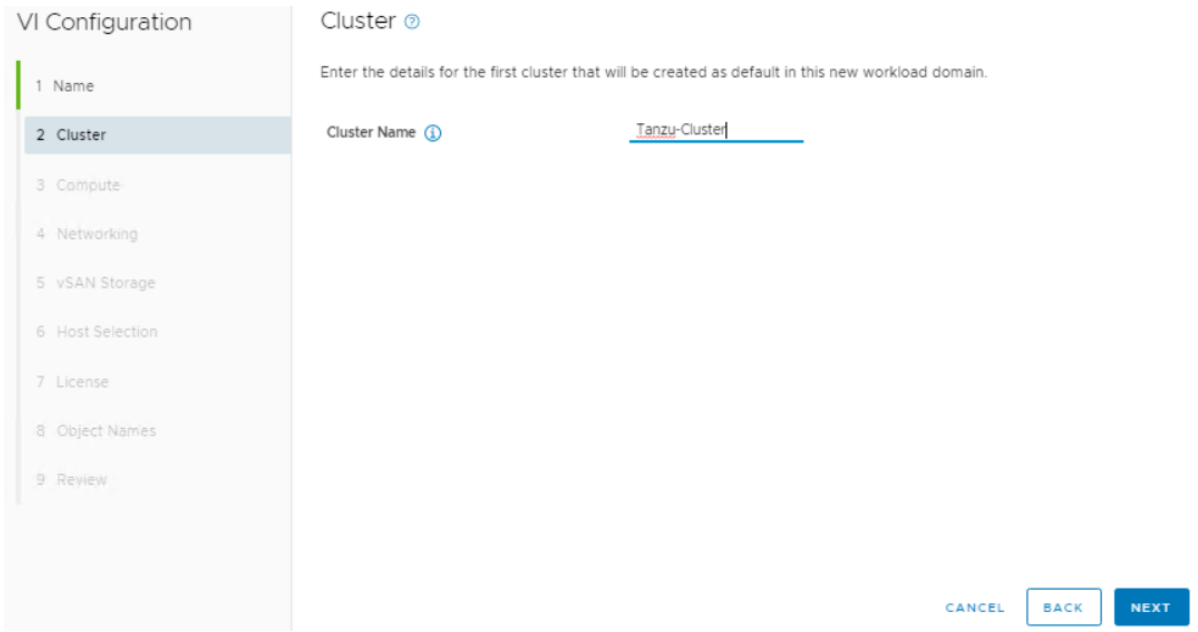

In the Compute window provide the following parameters for Compute, then click on the Next button:

- vCenter FQDN (already created in the DNS)
- vCenter IP address (should autocomplete if the DNS is properly configured)
- vCenter subnet mask
- vCenter Default gateway
- vCenter 'root' password

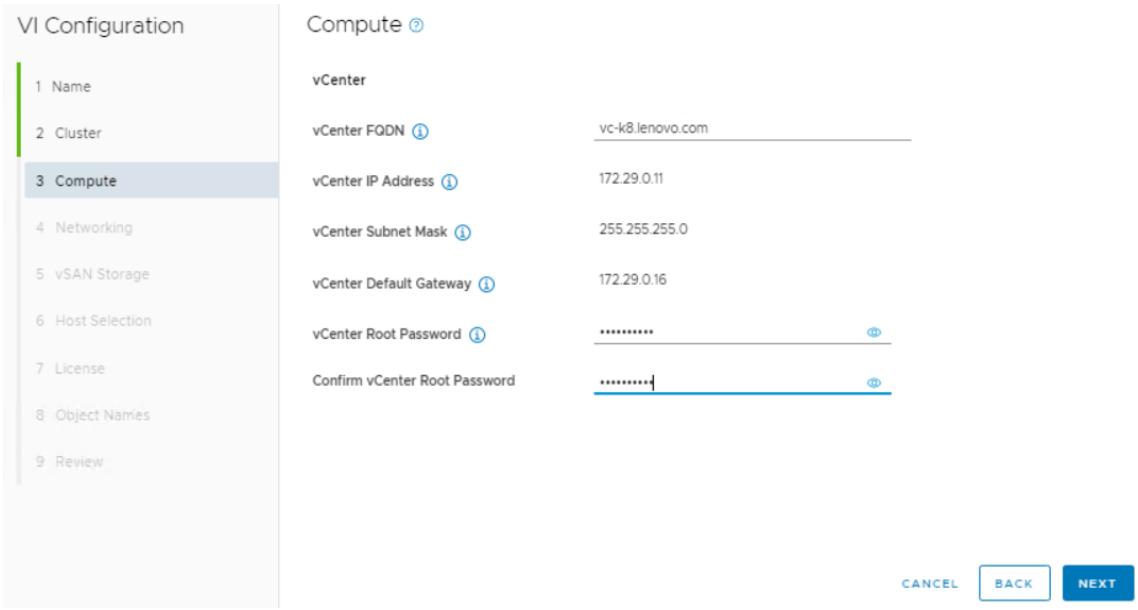

In the Networking window provide the following parameters for the NSX-T cluster manager, and click on the Next button after completing:

- NSX-T Cluster FQDN (already created in DNS)
- NSX-T Cluster IP (should autocomplete if the DNS is correctly configured)
- FQDN1 – the FQDN for the first NSX-T node in the cluster (already configured in the DNS)
- IP Address 1 – the IP address for the first NSX-T node in the cluster (should autocomplete)
- FQDN2 – the FQDN for the second NSX-T node in the cluster (already configured in the DNS)
- IP Address 1 – the IP address for the second NSX-T node in the cluster (should autocomplete)
- FQDN1 – the FQDN for the third NSX-T node in the cluster (already configured in the DNS)
- IP Address 1 – the IP address for the third NSX-T node in the cluster (should autocomplete)
- Admin Password for the NSX-T Manager
- VLAN ID for the NSX-T Host Overlay Network (can be the same VLAN used for the Management Domain NSX-T Cluster)
- IP allocation - Static Pool
- Pool Name
- Description
- CIDR – the subnet in CIDR notation

- IP Range – must not overlap existing IPs used for NSX-T Host Overlay Network in the Management Domain
- Gateway IP

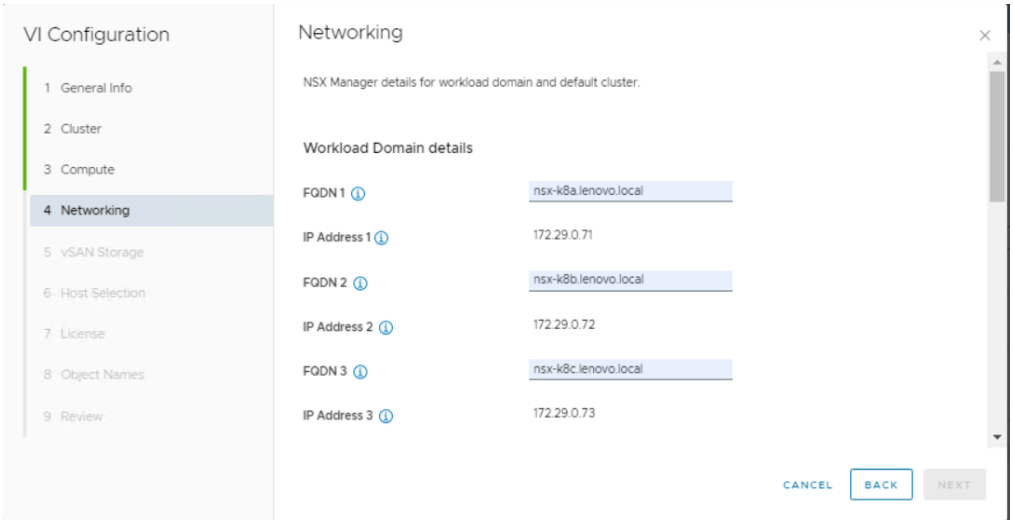

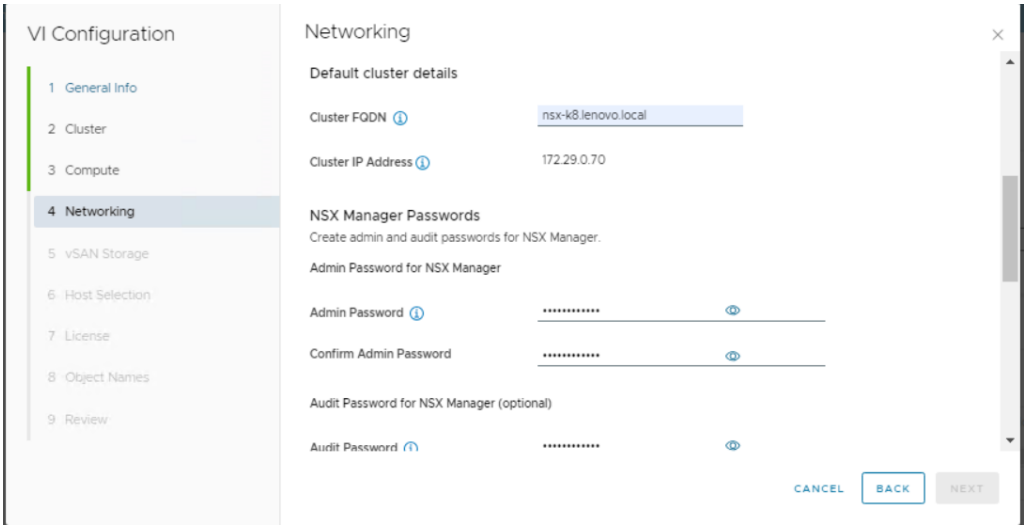

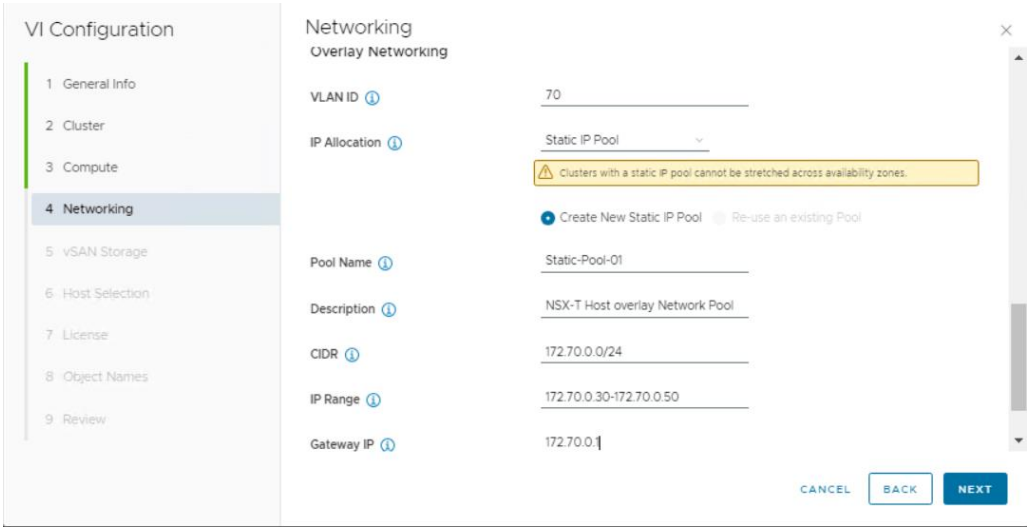

In the next window configure the desired vSAN parameters, then click on the Next button:

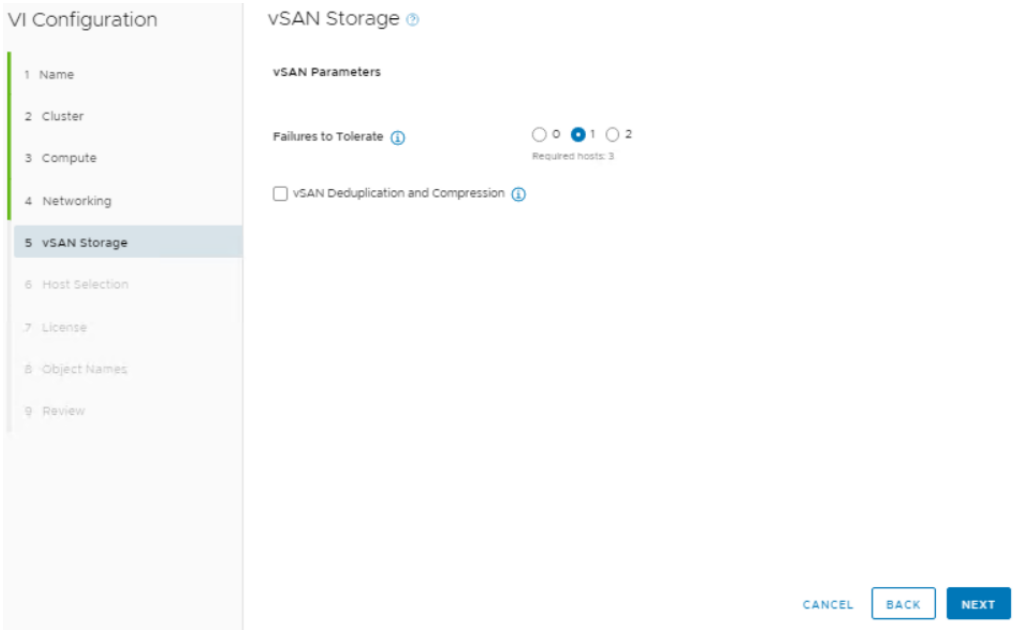

In the Host Selection window select all the hosts and click on the Next button:

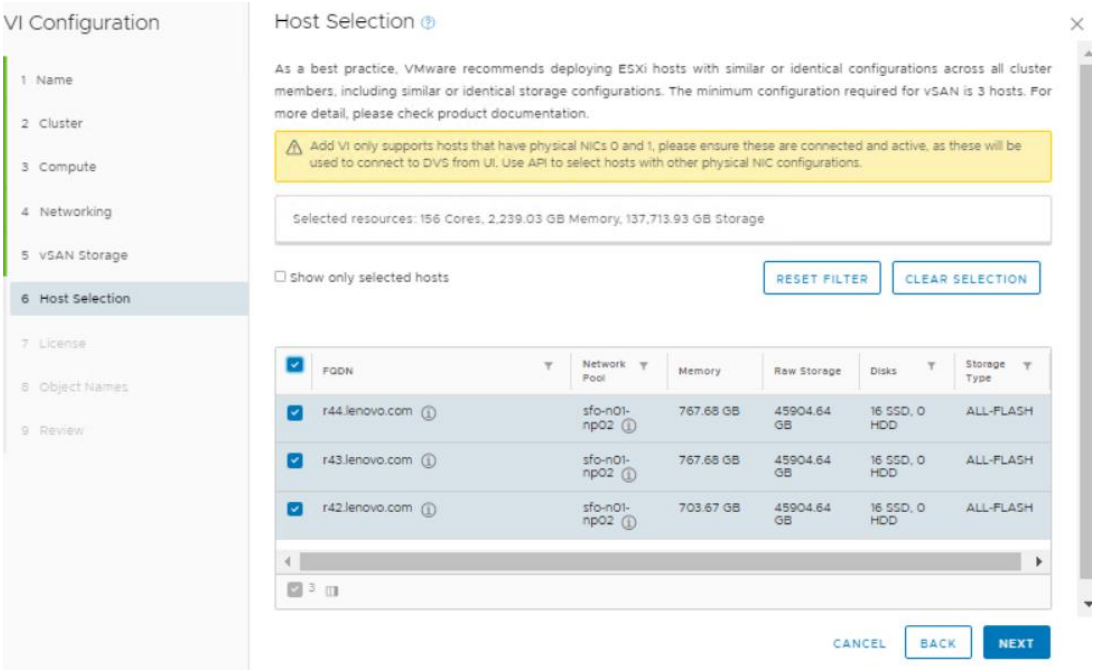

In the License window select the Licenses for each component:

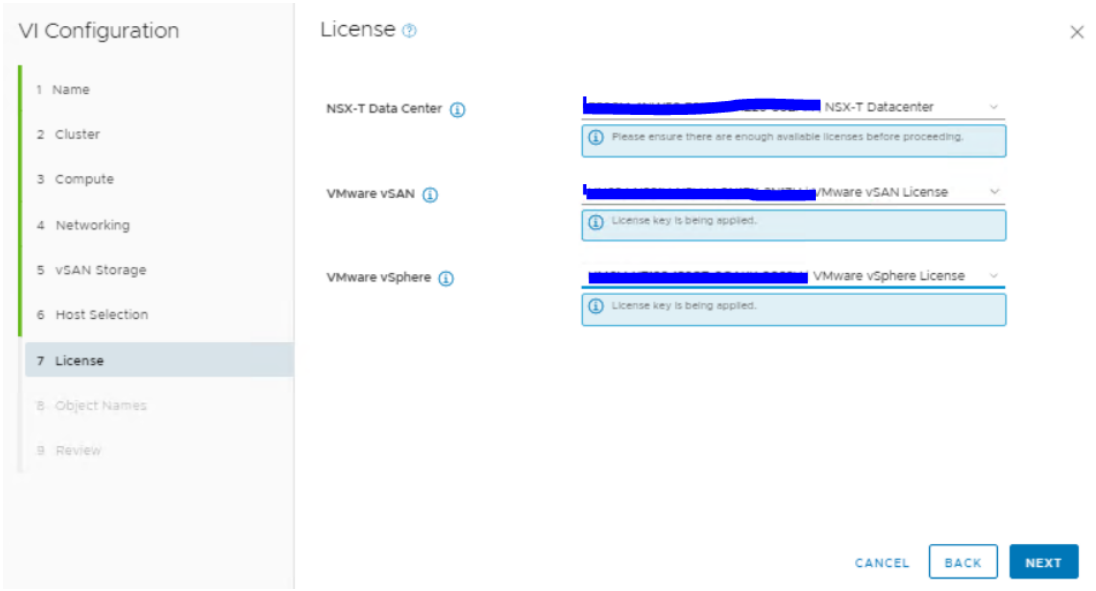

In the Object Names window review the parameters used then click on the Next button:

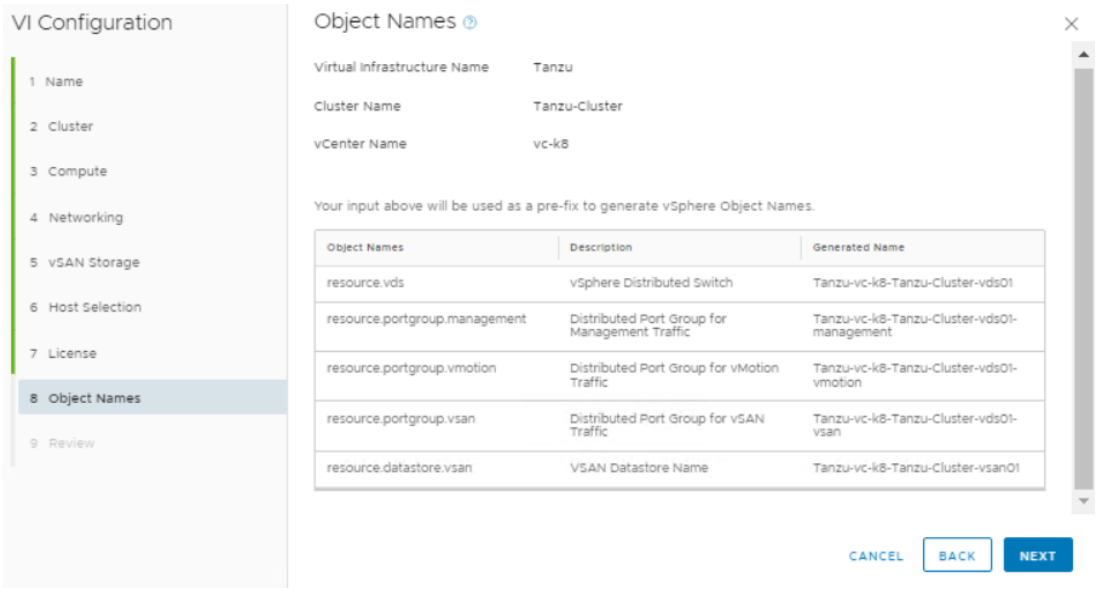

In the Review windows verify that all the parameters are correctly configured then click on the FINISH button:

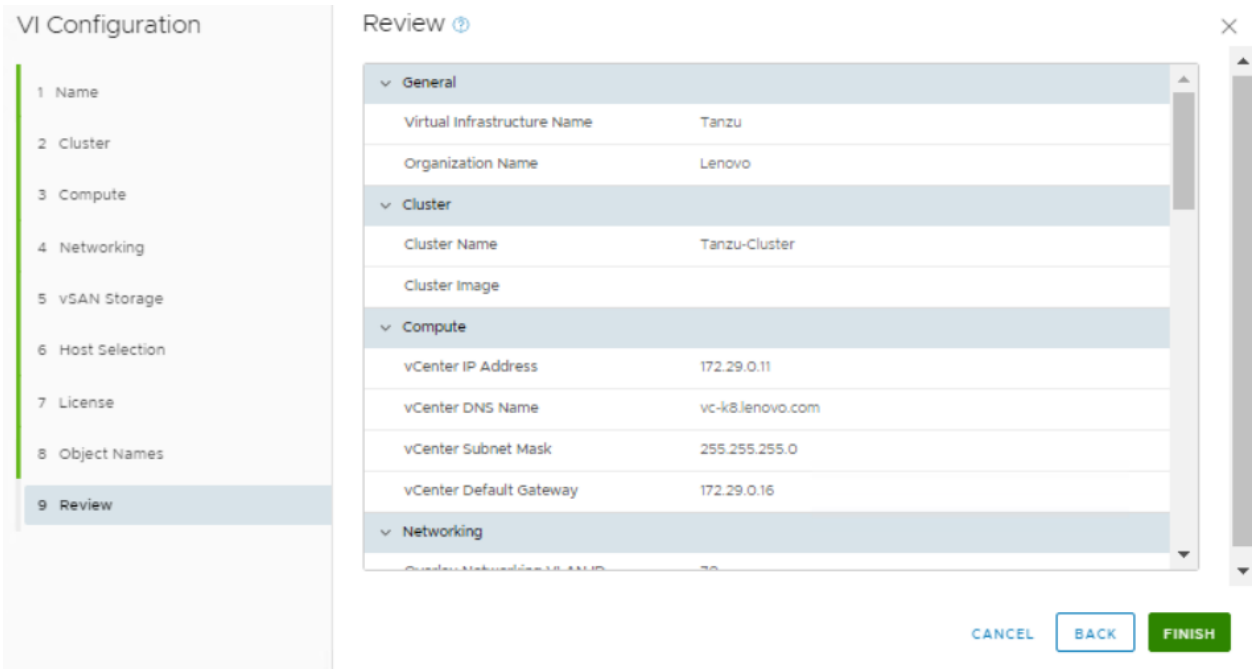

7.3.1.5 Deploy Aria Suite

In order to deploy Aria Suite products an Edge cluster and AVNs (Application Virtual Networking) must first be deployed. AVN is a software-defined networking concept based on NSX-T Data Center that allows the hosting of management applications on NSX segments.

Overlay-Backed NSX Segments

Overlay-backed segments provide flexibility for workload placement by removing the dependence on traditional data center networks. Using overlay-backed segments improves the security and mobility of management applications and reduces the integration effort with existing networks. Overlay-backed segments are created in an overlay transport zone.

In an overlay-backed segment, traffic between two VMs on different hosts but attached to the same overlay segment have their layer-2 traffic carried by a tunnel between the hosts. NSX-T Data Center instantiates and maintains this IP tunnel without the need for any segment-specific configuration in the physical infrastructure. As a result, the virtual network infrastructure is decoupled from the physical network infrastructure. That is, you can create segments dynamically without any configuration of the physical network infrastructure.

VLAN-Backed NSX Segments

VLAN-backed segments leverage the physical data center networks to isolate management applications, while still taking advantage of NSX-T Data Center to manage these networks. VLAN-backed network segments ensure the security of management applications without requiring support for overlay networking. VLAN-backed segments are created in a VLAN transport zone.

A VLAN-backed segment is a layer-2 broadcast domain that is implemented as a traditional VLAN in the physical infrastructure. This means that traffic between two VMs on two different hosts but attached to the same VLAN-backed segment is carried over a VLAN between the two hosts. The resulting constraint is that you must provision an appropriate VLAN in the physical infrastructure for those two VMs to communicate at layer-2 over a VLAN-backed segment.

Aria Suite Component	NSX Segment
Aria Operations for Logs	Region-A
Aria Operations Manager	X-Region
Workspace ONE Access	X-Region
Aria Automation	X-Region
Aria Suite Lifecycle Manager	X-Region

Adding the Edge Cluster on the Workload Management domain:

In SDDC Manager go to Inventory > Workload Domains and click on the 3 dots near the newly created WD and select Add Edge Cluster:

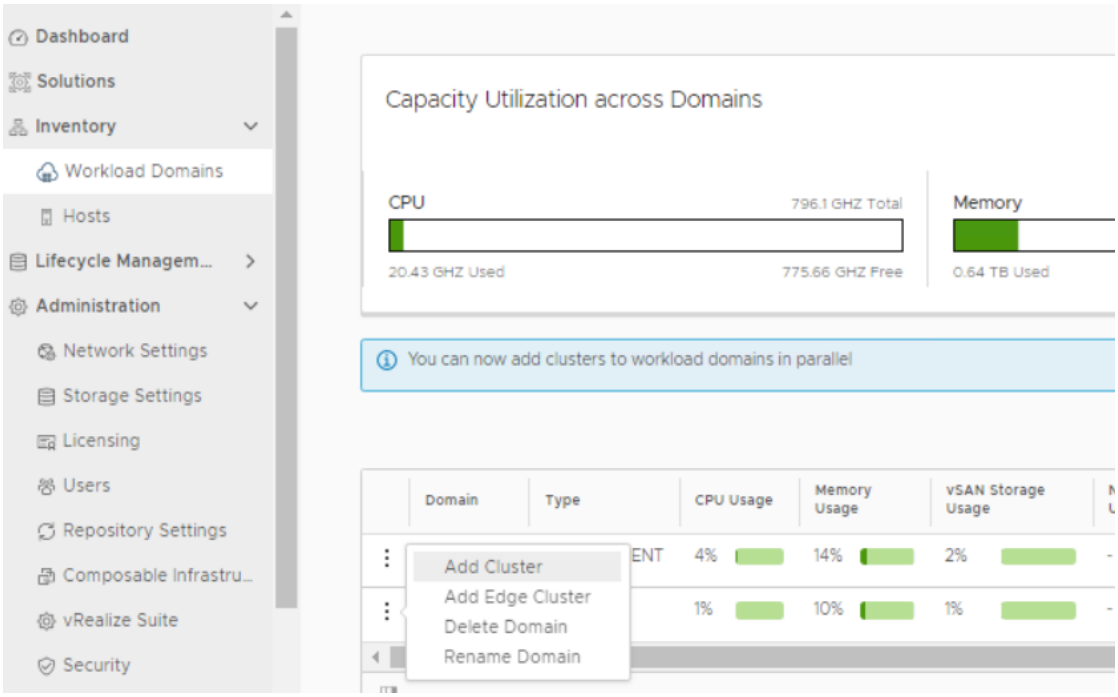

In the Edge Cluster Prerequisites Select All after verifying that all prerequisites have been met and click on the Begin button:

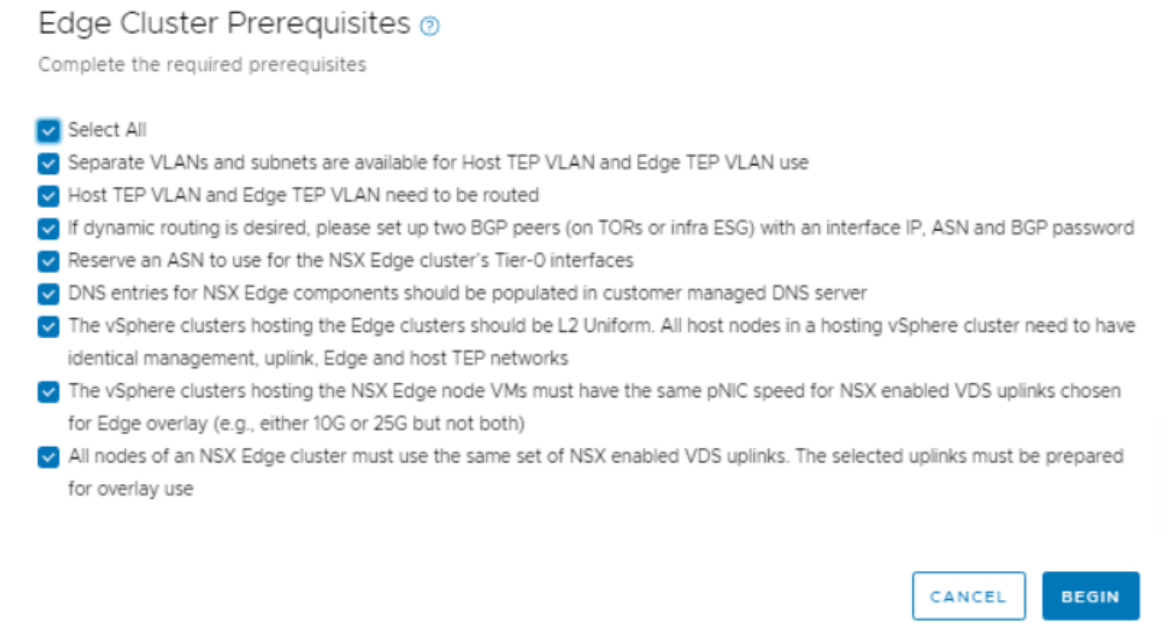

In the General Info window provide the following parameters and click on the Next button:

- Edge Cluster Name
- MTU: 9000
- Tier-0 router name
- Tier-1 router name
- Edge Cluster Profile Type: Default
- Create passwords for Edge root, Edge admin and Edge audit accounts

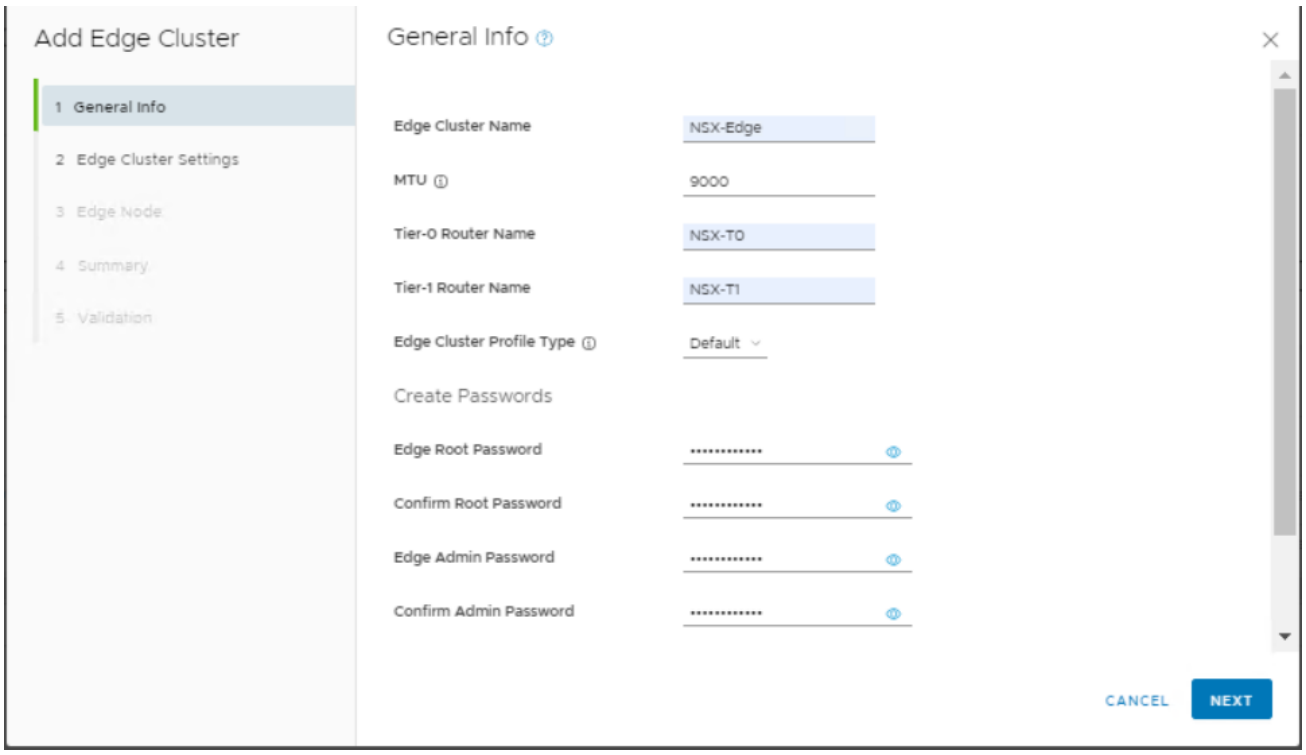

In the Edge Cluster Settings windows select Application Virtual Networks and provide the following parameters:

- ASN (make sure it matches the remote-as ASN configured for BGP on the physical switch)
- Edge Form Factor (medium)
- Tier-0 Service High Availability (Active-Active)
- Tier-0 Routing type (eBGP)

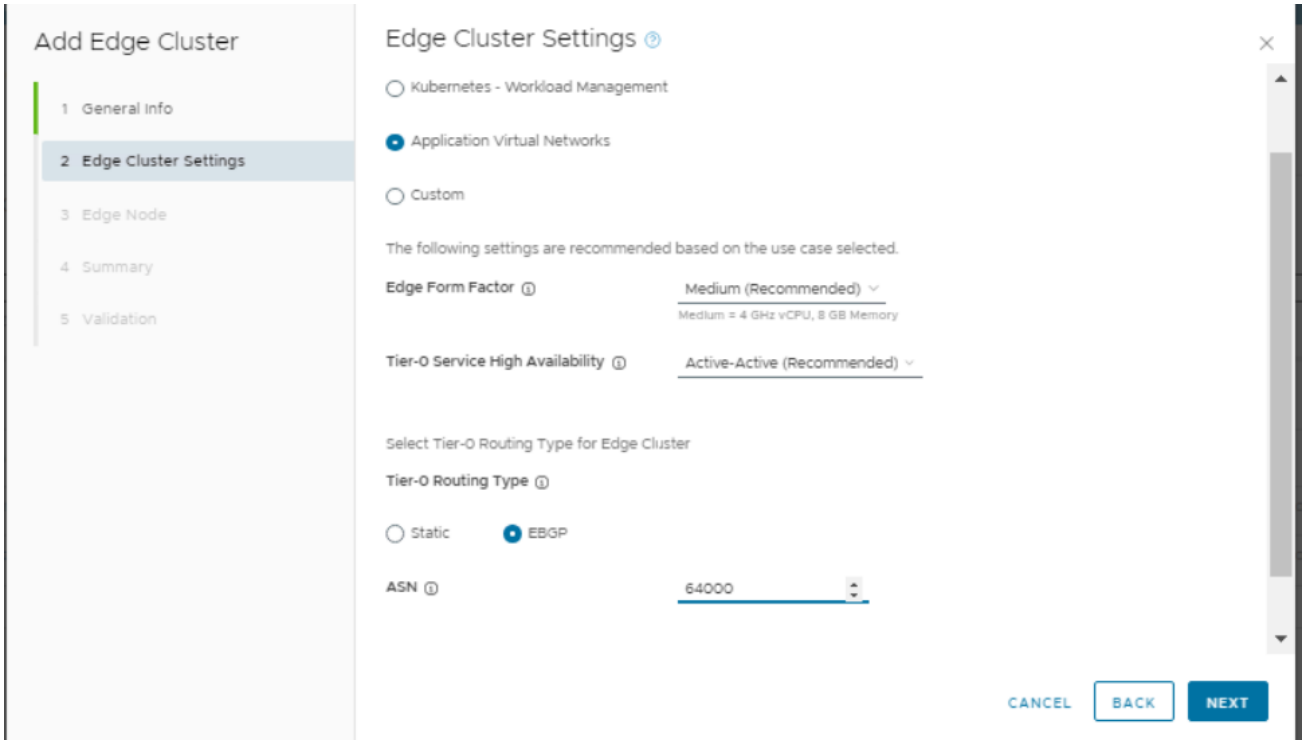

In the Edge Node window provide the following parameters:

- Edge Node Name (FQDN)
- Select the Management Domain cluster
- Cluster Type (L2 uniform – esxi hosts have the same networks for mgmt., TEP etc.)
- Management IP (CIDR)
- Management gateway
- EDGE TEP1 IP CIDR
- EDGE TEP2 IP CIDR
- EDGE TEP Gateway
- EDGE TEP VLAN ID
- Tier-0 Uplink VLAN ID
- Tier-0 Uplink Interface IP CIDR
- BGP Peer IP CIDR

- BGP Peer ASN
- BGP Peer Password

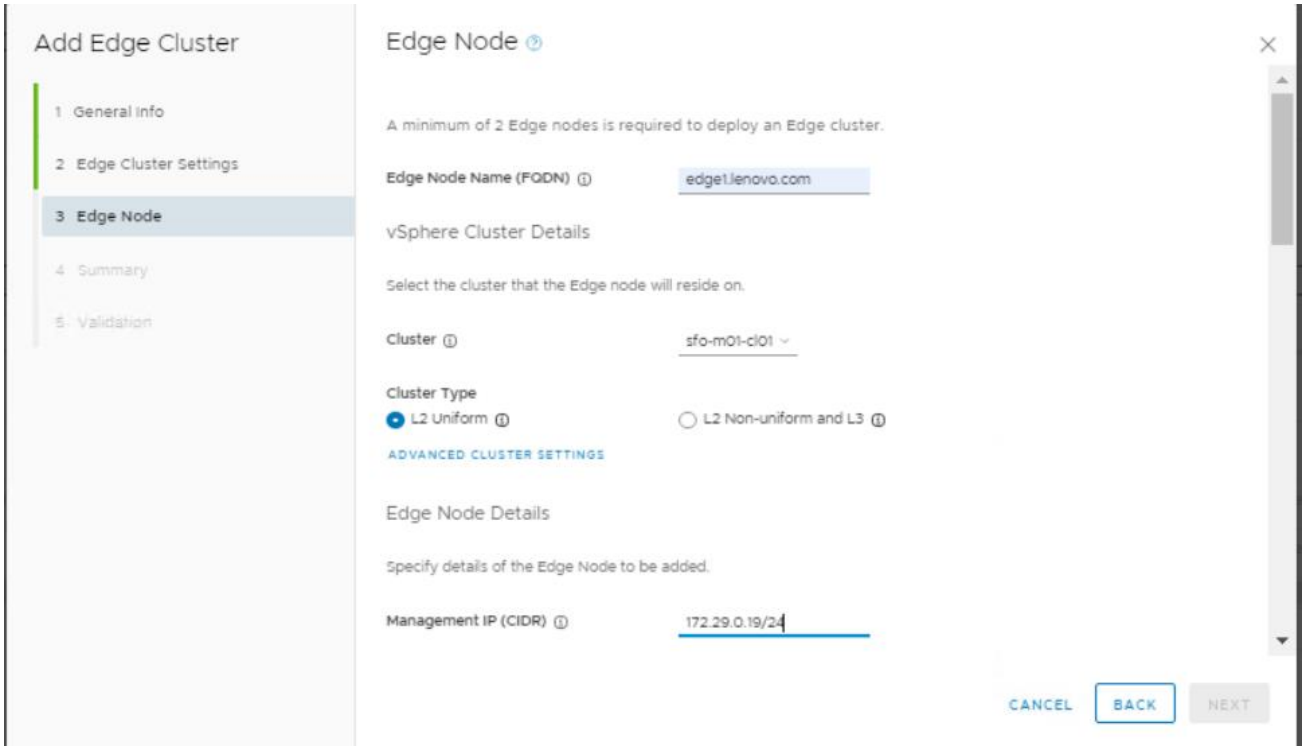

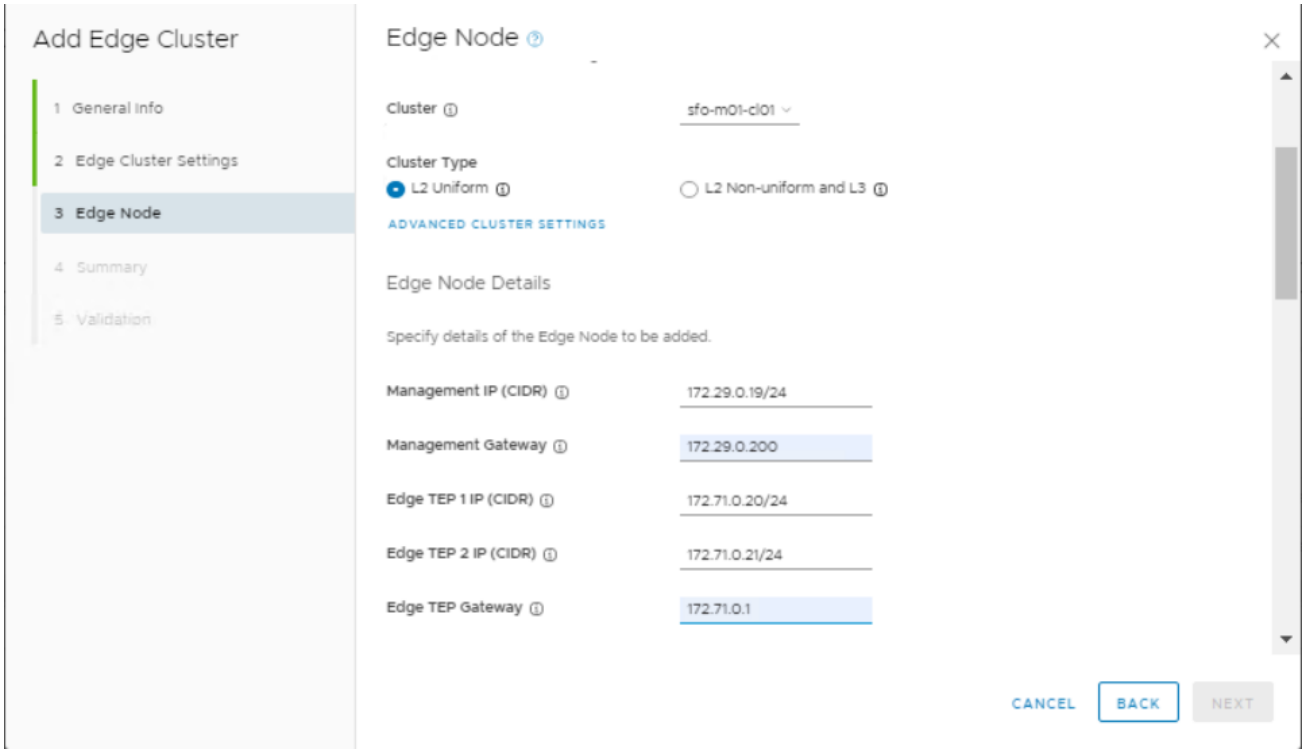

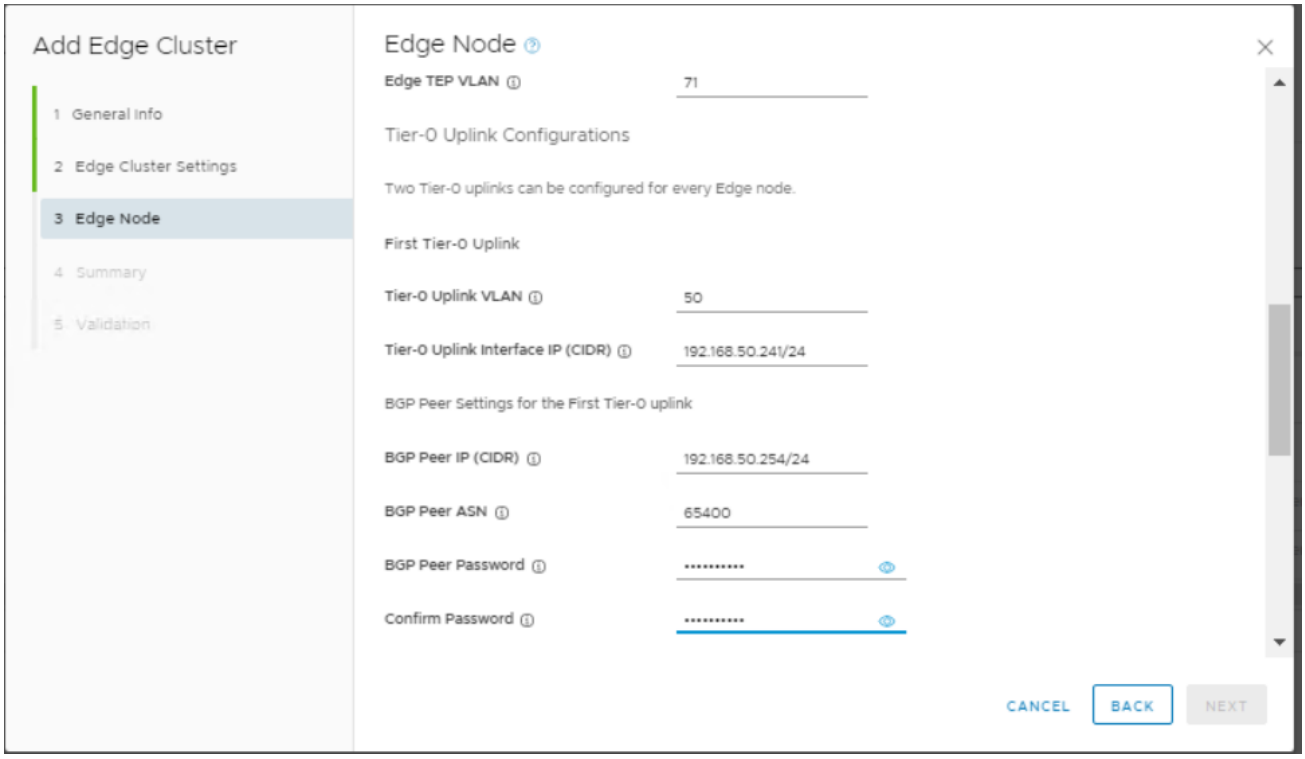

In the Summary windows verify that the Edge Cluster and Nodes have been properly configured then click on the Next button:

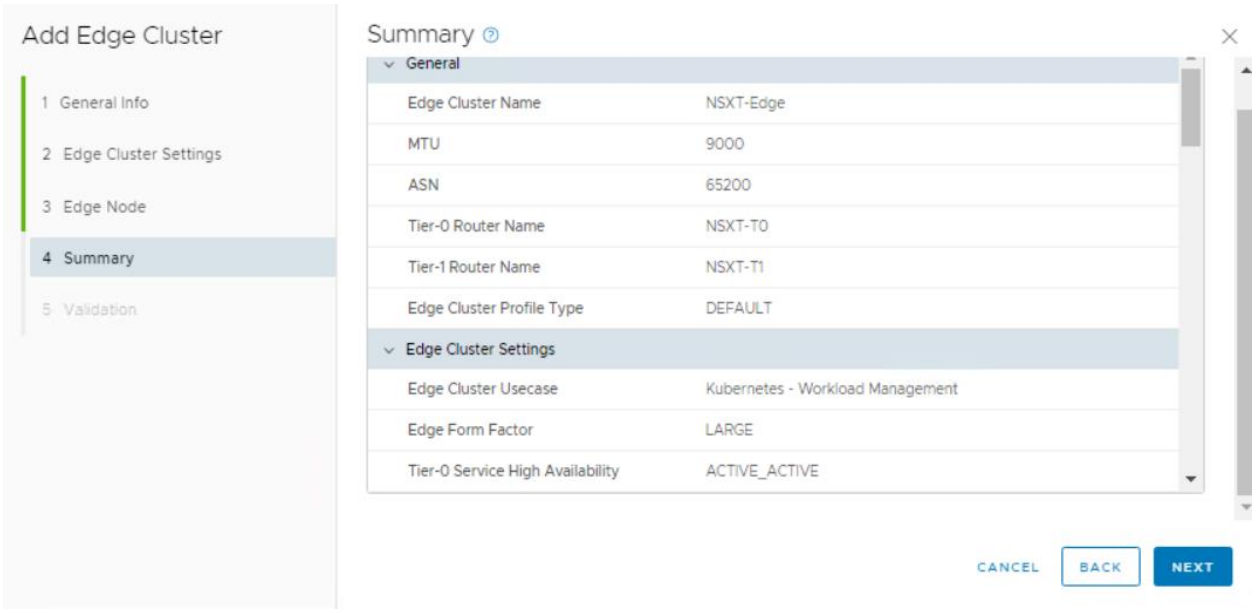

In the Validation window some checks are automatically done, click on the FINISH button if all validations have been successful, otherwise revise the previous settings that failed:

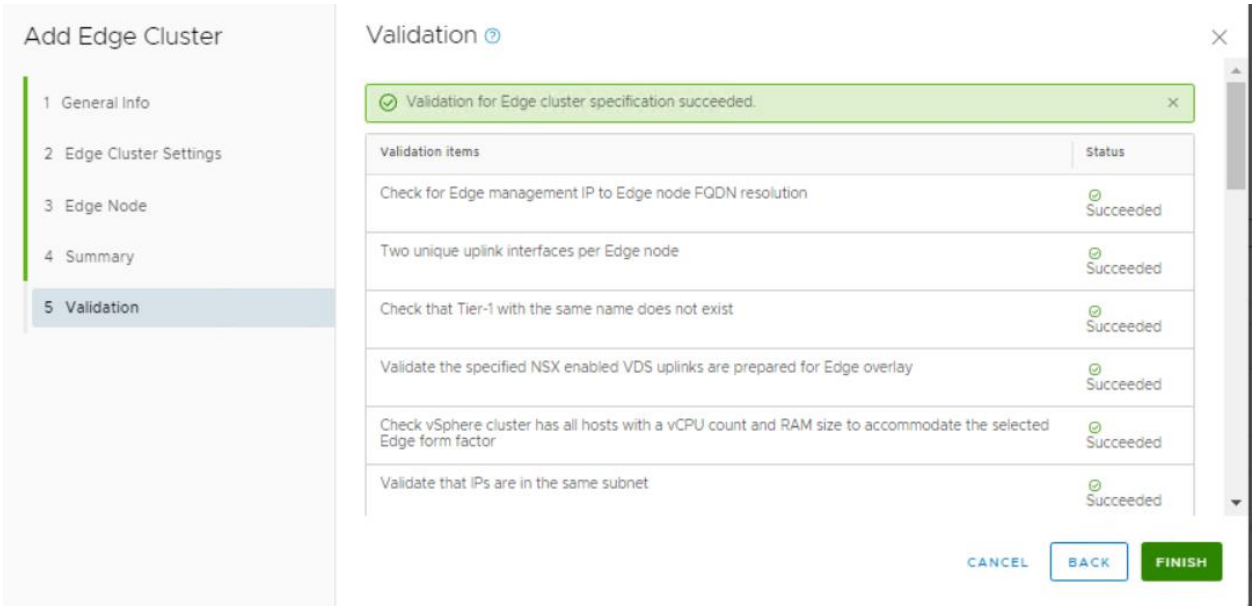

Add AVNs networks:

On the Workload Management Cluster click the 3 dots and select Add AVNs:

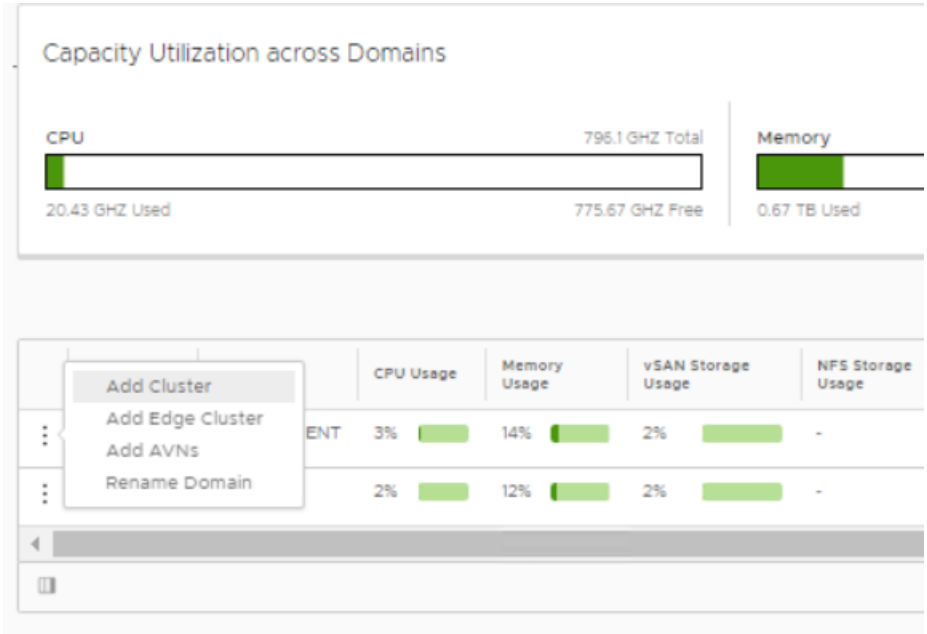

In General windows select VLAN-backed NSX segment:

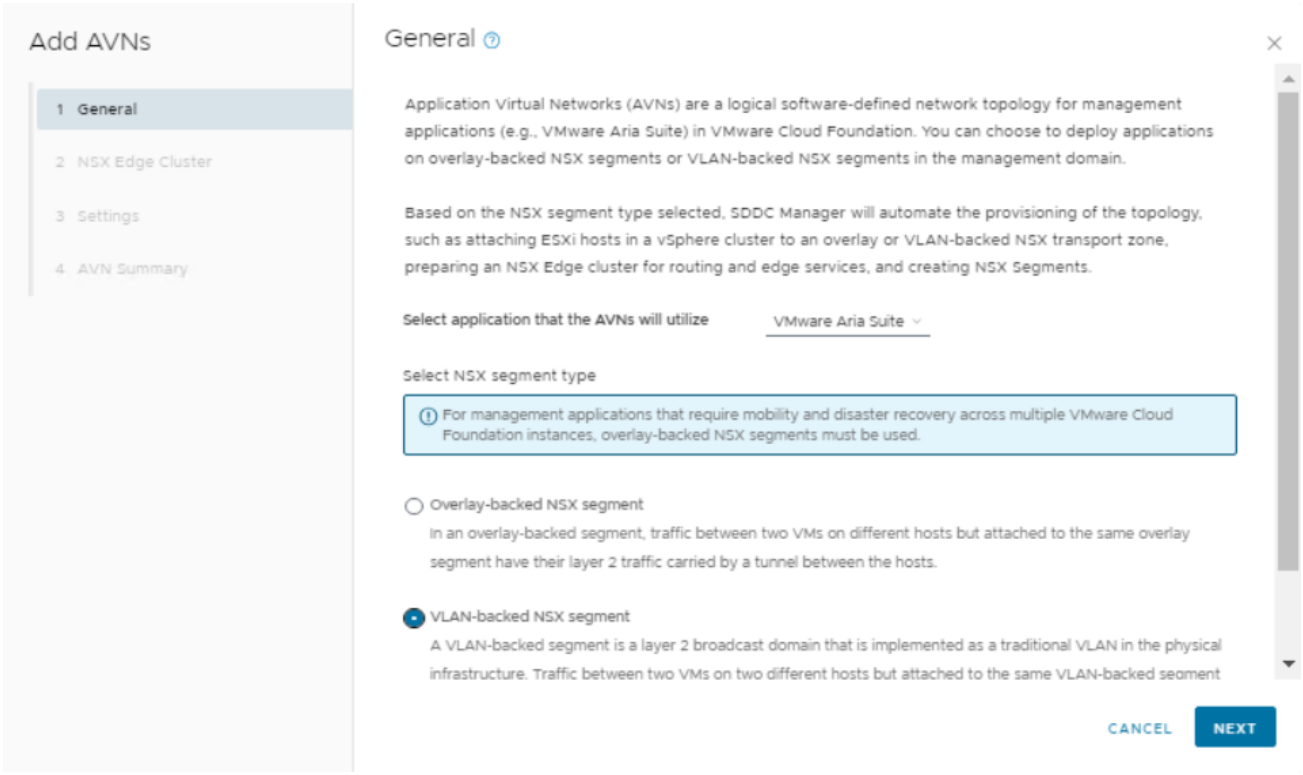

In the NSX Edge Cluster windows select the NSX-Edge cluster deployed on the Management Cluster:

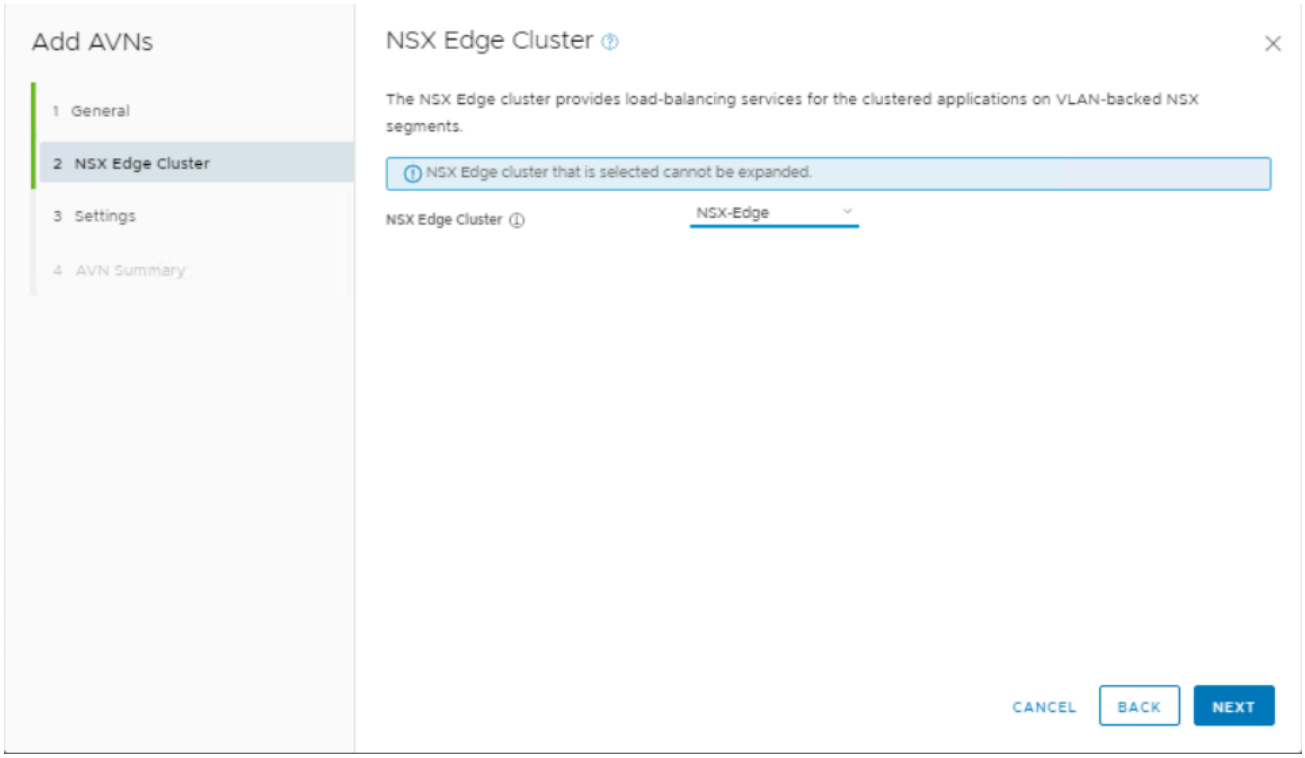

In the Settings window provide the following parameters:

- Region-A Name
- Region-A Network (the first IP must be omitted since it's the subnet gateway IP)
- Region-A network subnet mask
- Region-A network gateway IP address
- Region-A MTU
- Region-A VLAN ID
- X-Region Name
- X-Region Network (the first IP must be omitted since it's the subnet gateway IP)
- X-Region network subnet mask
- X-Region network gateway IP address
- X-Region MTU
- X-Region VLAN ID

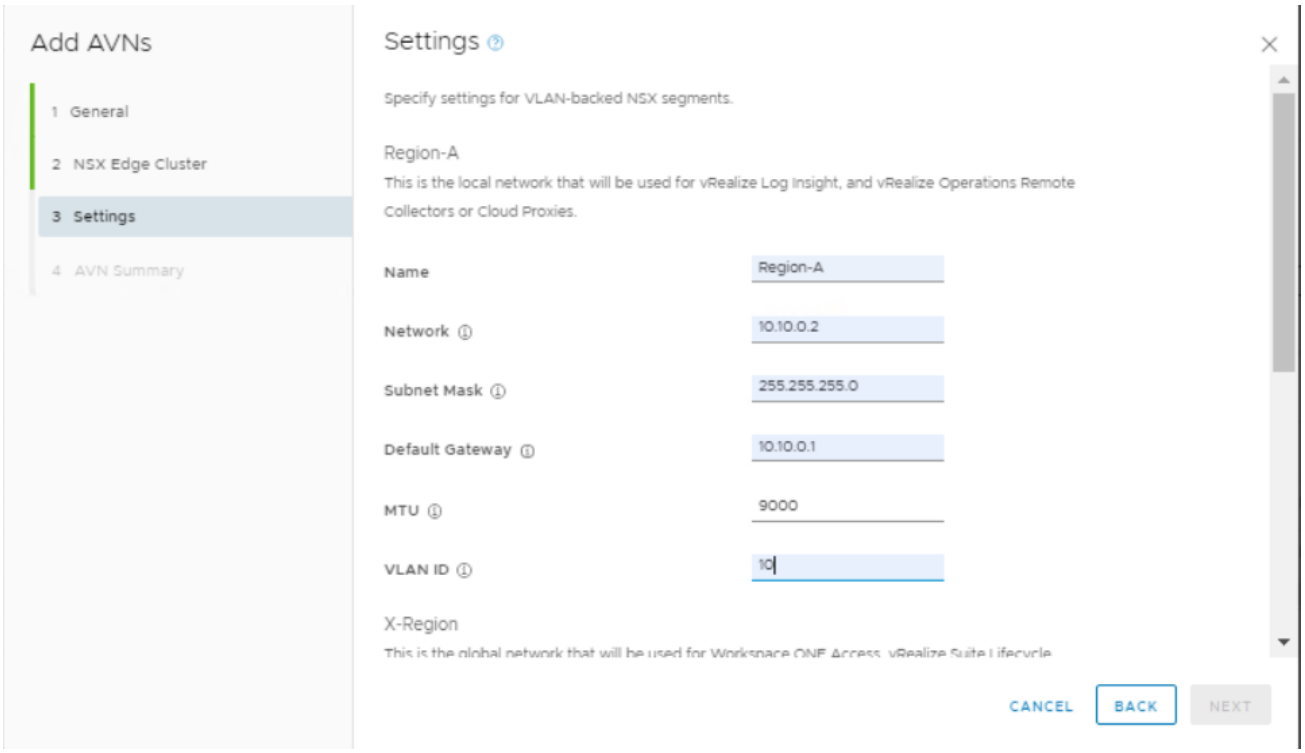
Add AVNs
1 General
2 NSX Edge Cluster
3 Settings
4 AVN Summary
Settings
Specify settings for VLAN-backed NSX segments.
Region-A
This is the local network that will be used for vRealize Log Insight, and vRealize Operations Remote Collectors or Cloud Proxies.
Name
Region-A
Network
10.10.0.2
Subnet Mask
255.255.255.0
Default Gateway
10.10.0.1
MTU
9000
VLAN ID
10
X-Region
CANCEL
BACK
NEXT

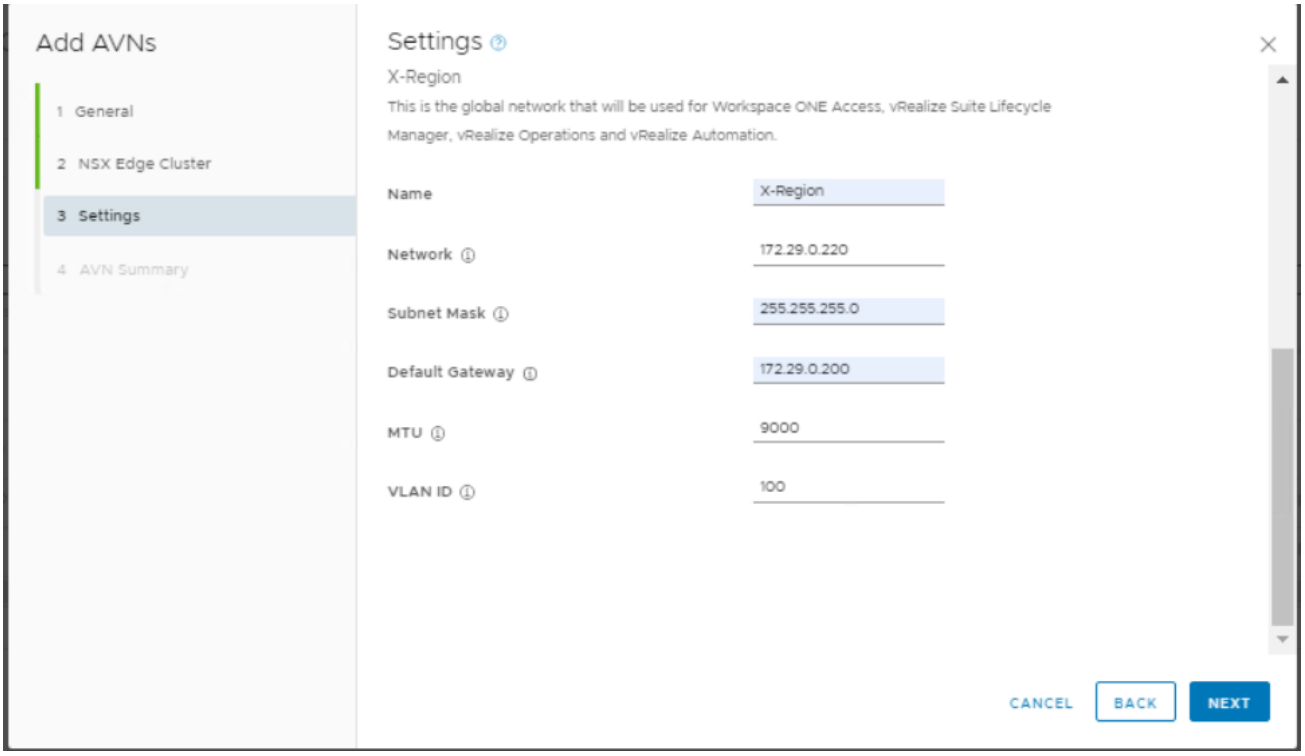
Add AVNs
1 General
2 NSX Edge Cluster
3 Settings
4 AVN Summary
Settings
X-Region
This is the global network that will be used for Workspace ONE Access, vRealize Suite Lifecycle Manager, vRealize Operations and vRealize Automation.
Name
X-Region
Network
172.29.0.220
Subnet Mask
255.255.255.0
Default Gateway
172.29.0.200
MTU
9000
VLAN ID
100
CANCEL
BACK
NEXT

Review the settings:

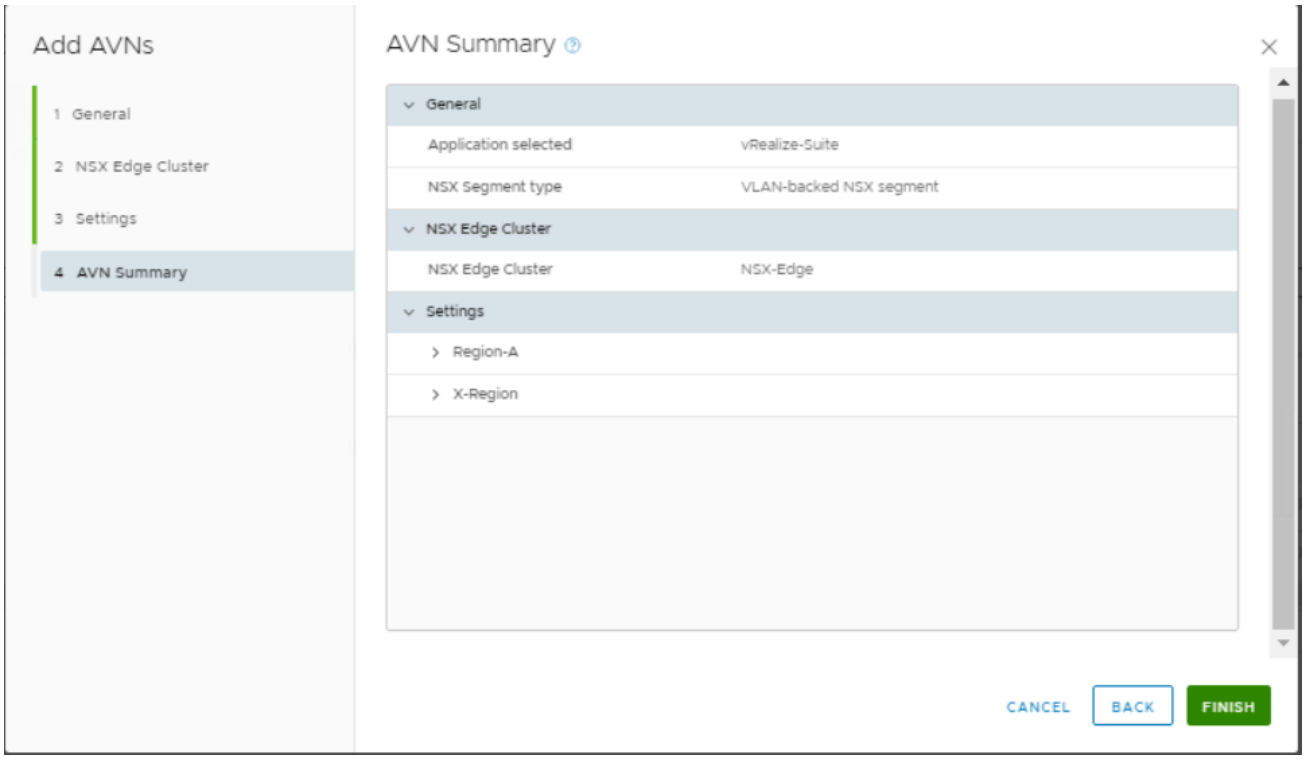

To verify that the AVNs regions have been correctly deployed:

Login to the SDDC Manager appliance using the ‘vcf’ username and password and issue the following commands:

#psql -h localhost -U postgres -d platform

#SELECT * FROM avn;

To update the subnets use the following commands:

#update avn SET gateway='<new_gateway_ip>' where name='<avn_name>';

#update avn SET subnet='<new_subnet>' where name='<avn_name>';

#update avn SET subnet_mask='<new_subnet_mask>' where name='<avn_name>'

Deploy Aria Lifecycle Manager:

In SDDC Manager go to Aria Suite:

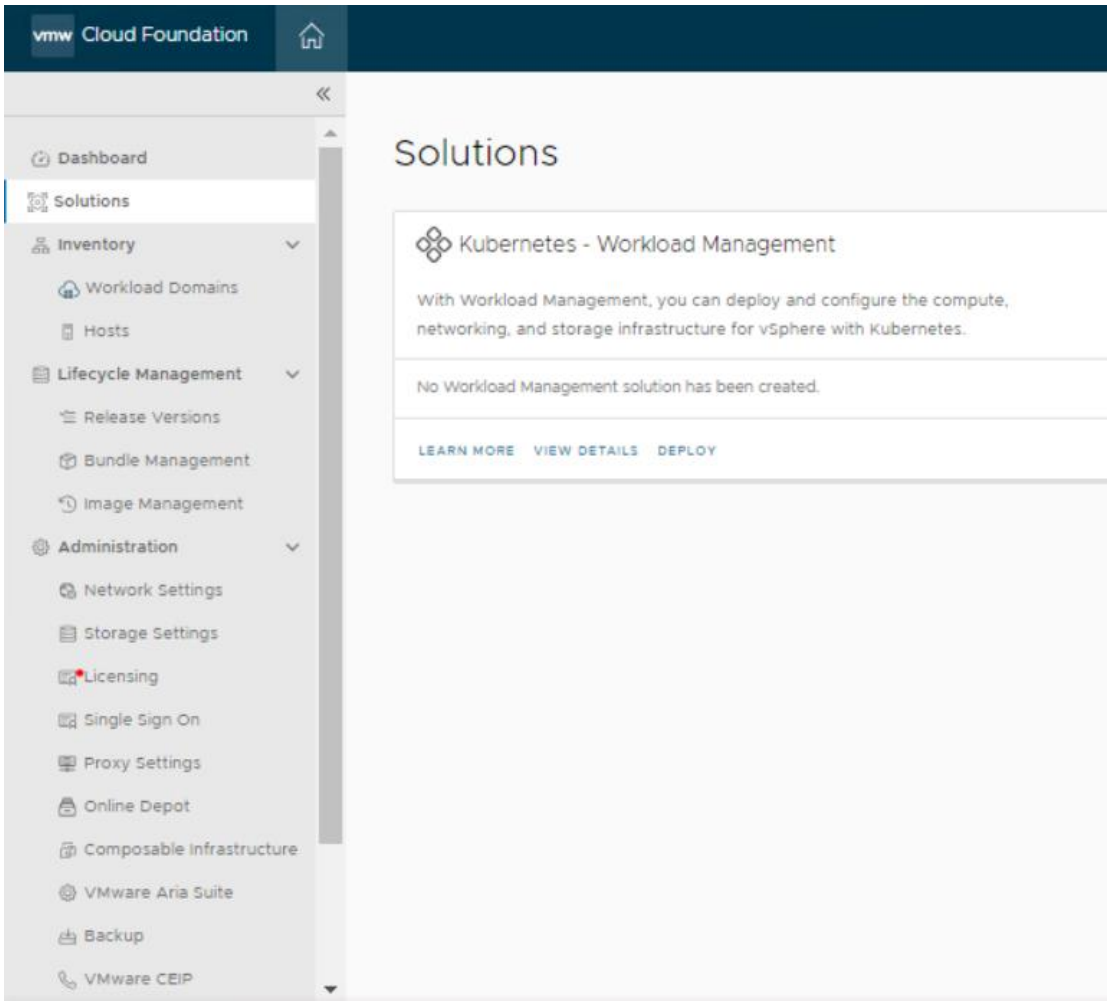

In the Aria Suite, after all the prerequisites have been met, we can click on the Deploy button:

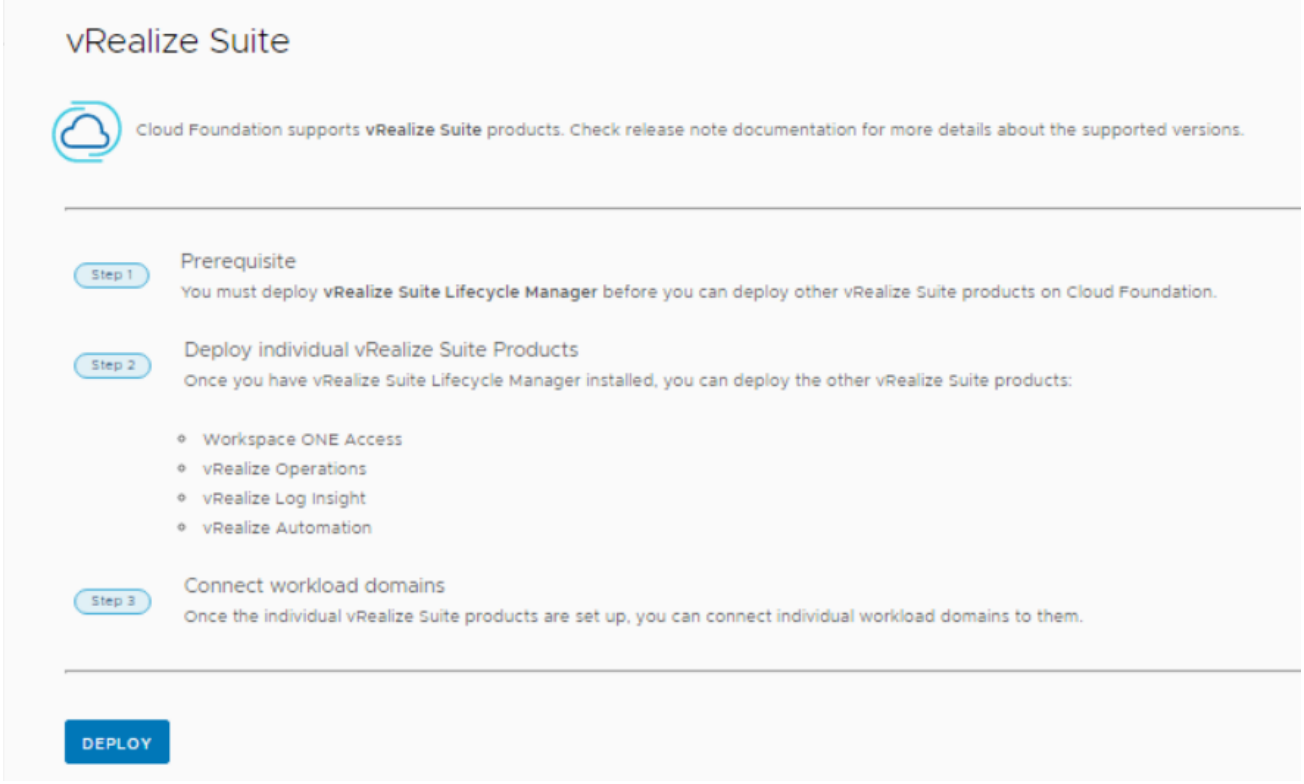

In the Aria LCM Installation Prerequisites window check the Select All button and click Begin:

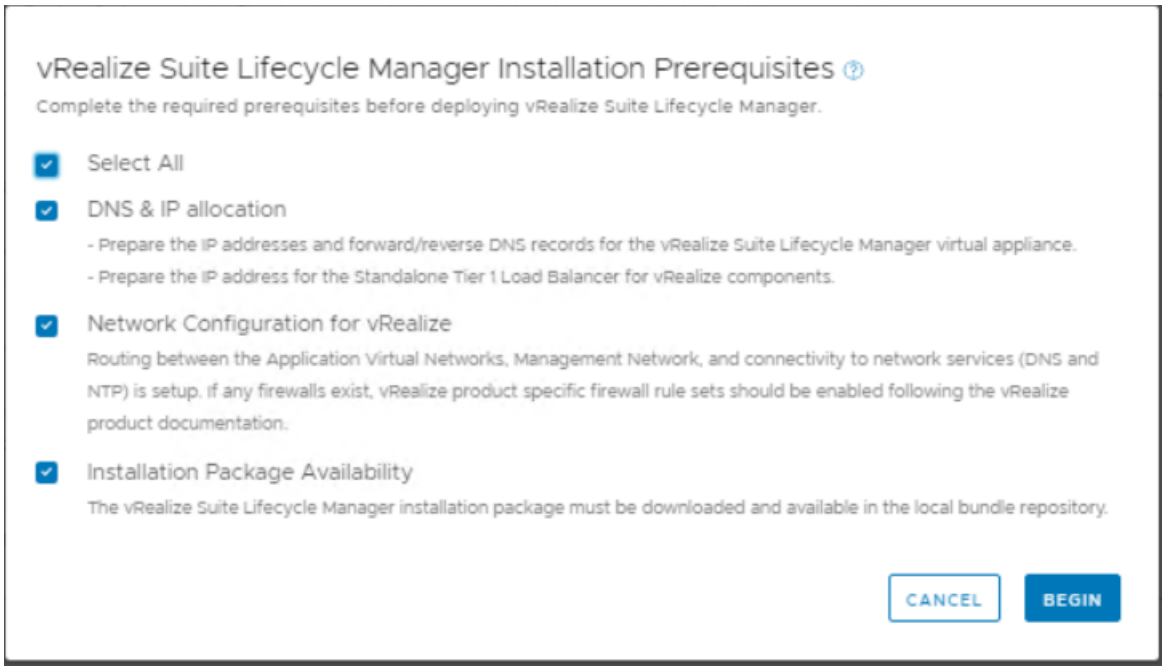

In the Network Settings verify if the X-Region parameters are correct:

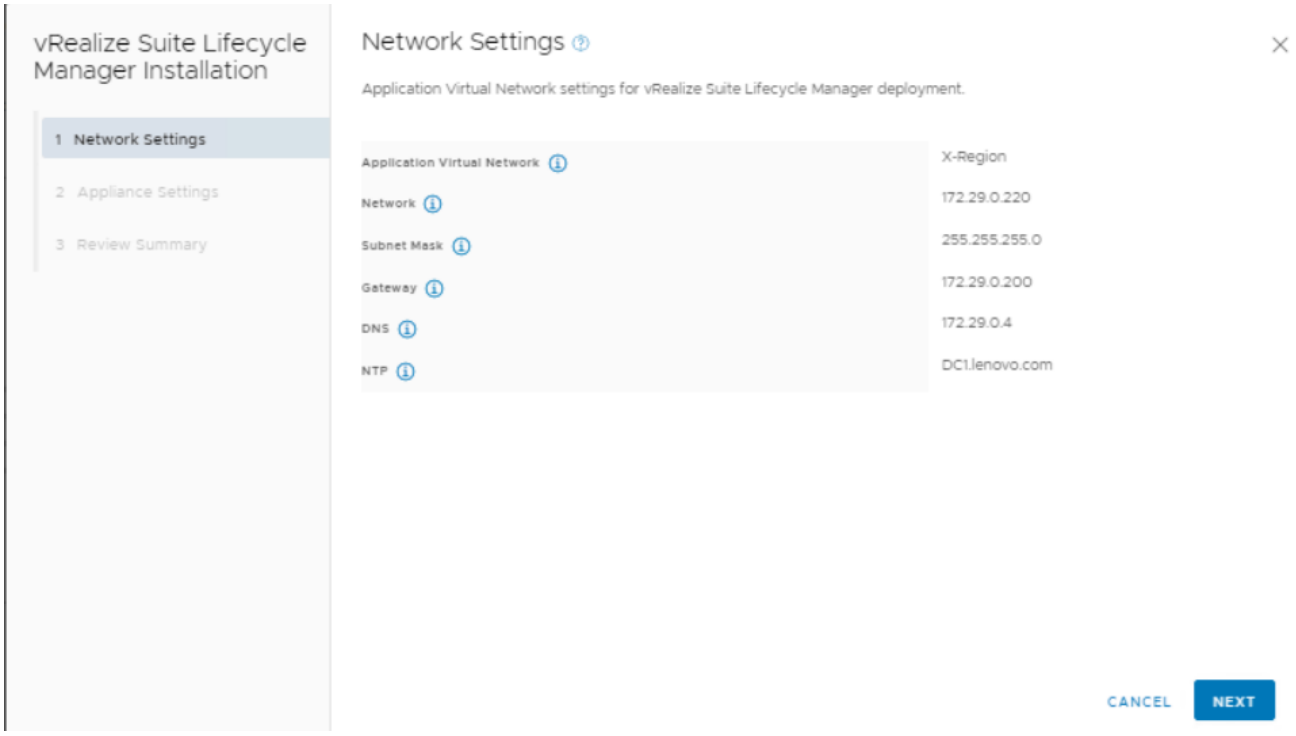

In the Virtual Appliance Settings provide the following parameters and click Next:

- The FQDN of the vRLCM appliance (already configured in DNS)
- NSX-T Tier-1 Gateway
- System Administrator password
- SSH Root Account password

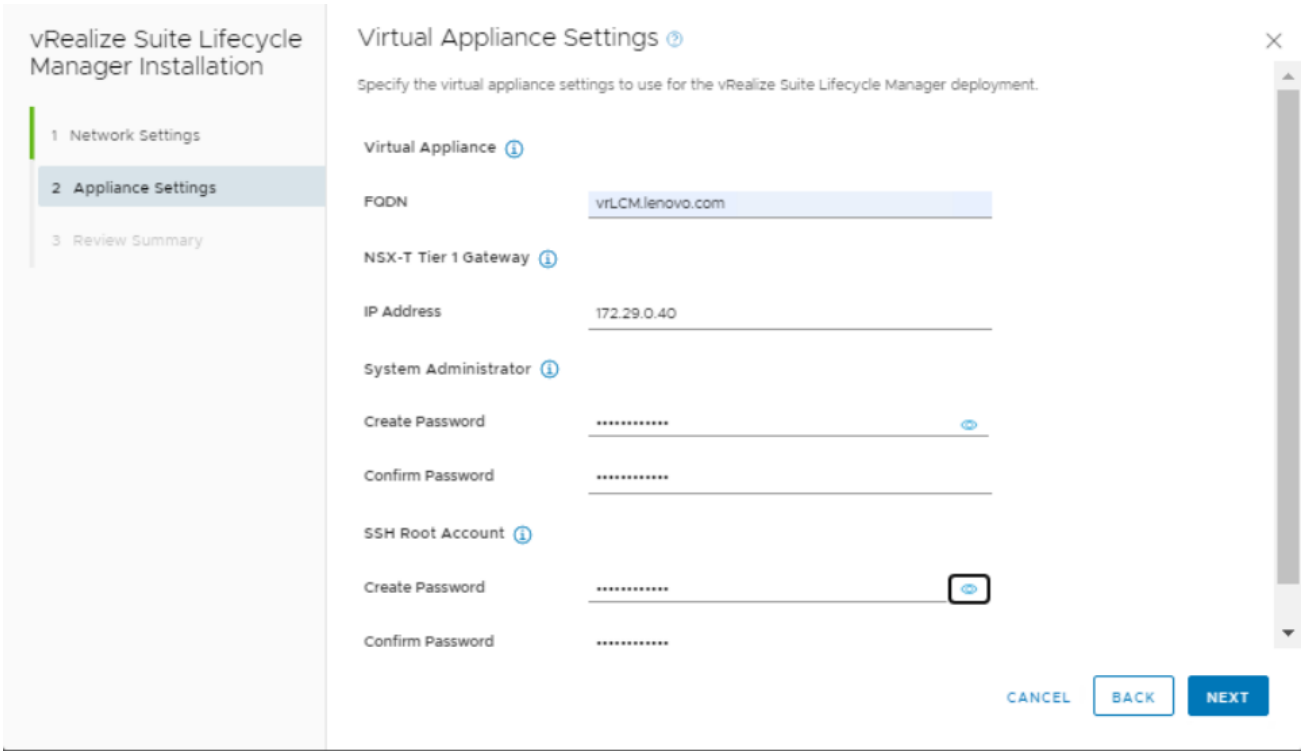

In the Review Summary windows check if the parameters provided are correct and click Finish:

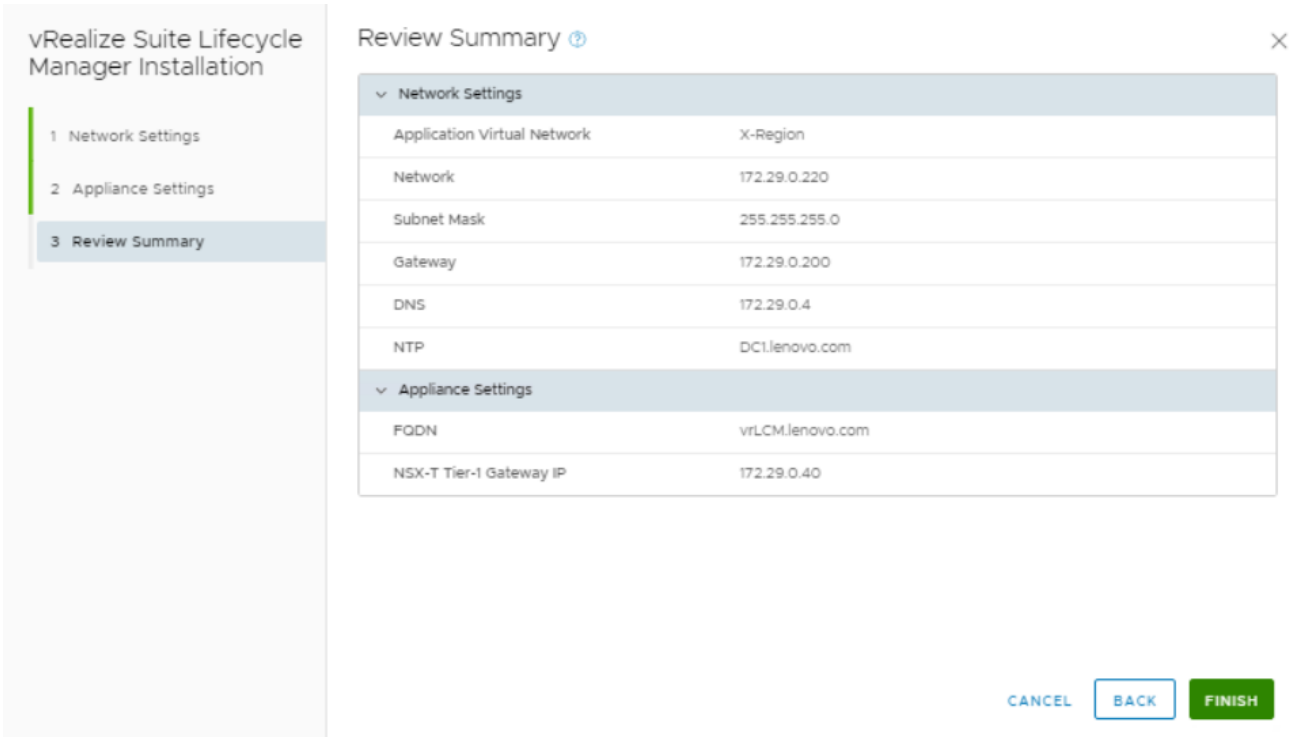

Login to the Aria Lifecycle Manager Dashboard by using the link provided using the vcfadmin@local (System Administrator) user:

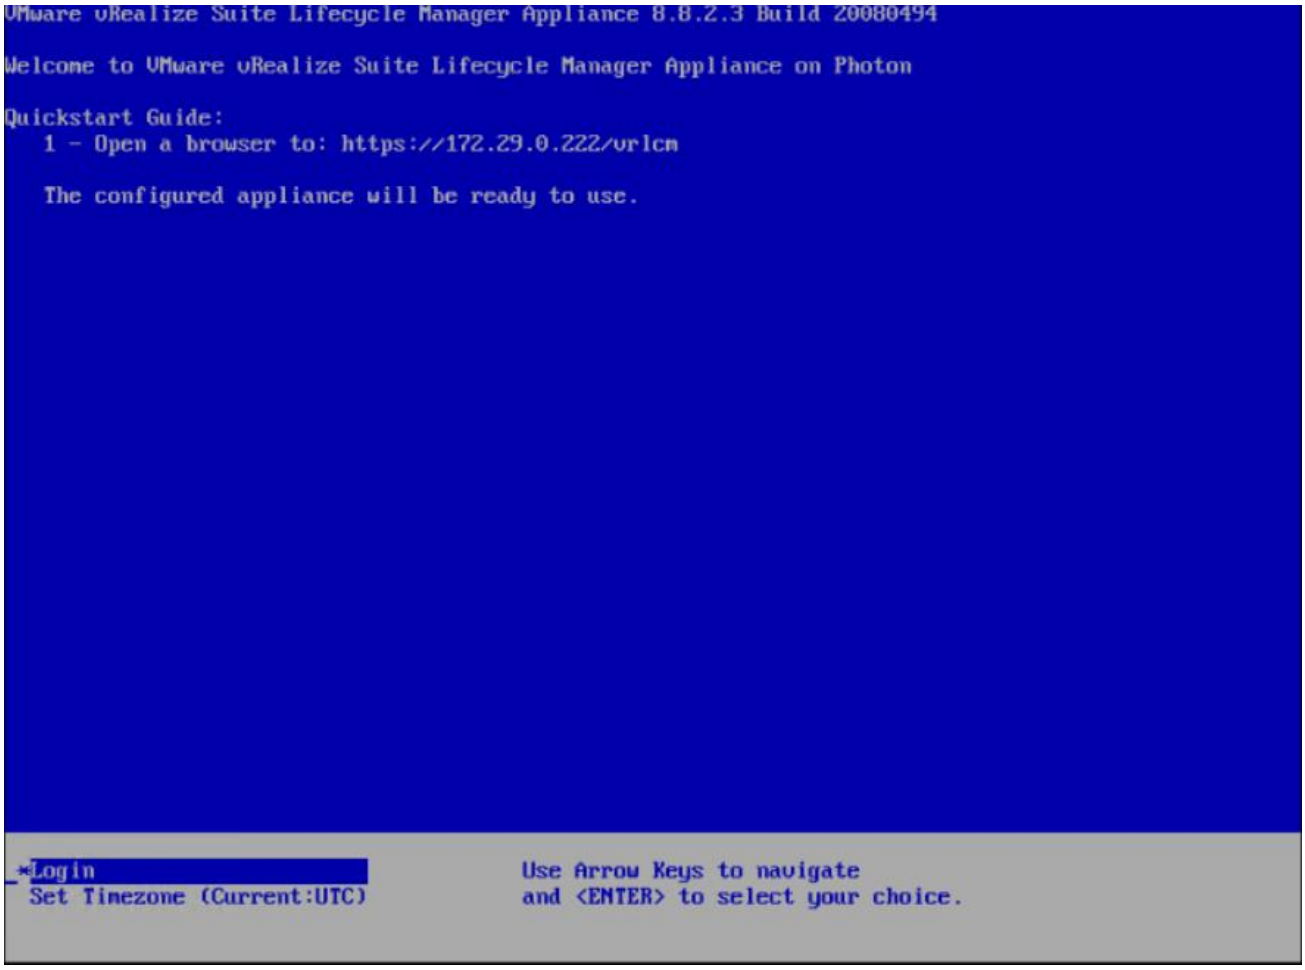

Login to VMware Aria Suite Lifecycle (use the credentials for vcfadmin@local configured during the vRLCM deployment)

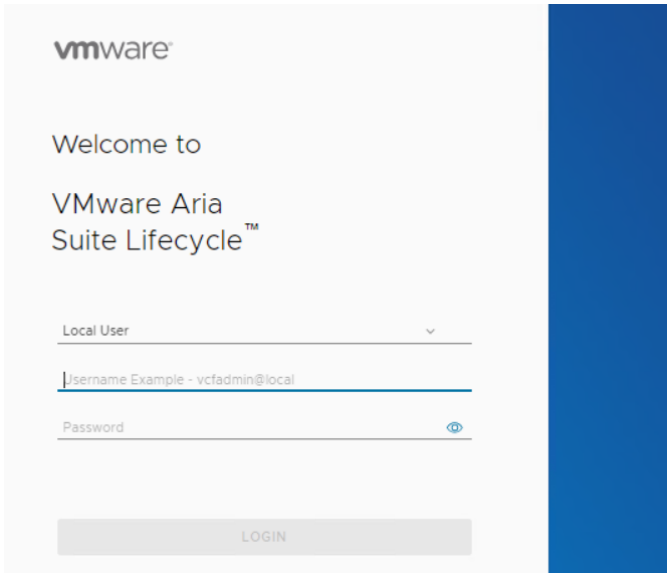

:

From the Dashboard we can access the following Components:

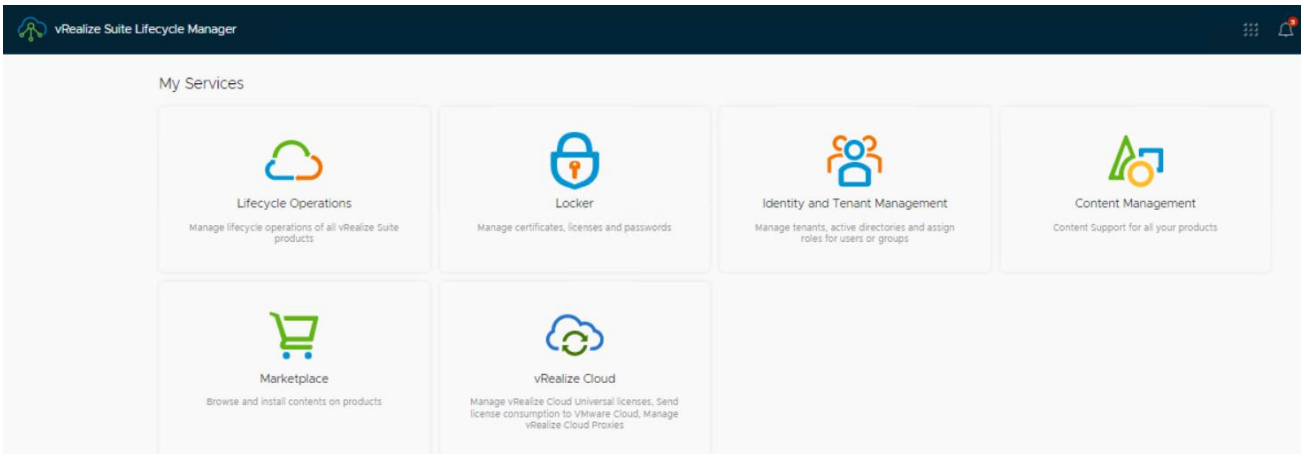

In order to install the other Aria components we must add the binaries to Aria LCM repository. To achieve this, a VMware account must be provided for Aria LCM:

In the Dashboard above, select LifeCycle Operations (login using the System Administrator vcfadmin@local):

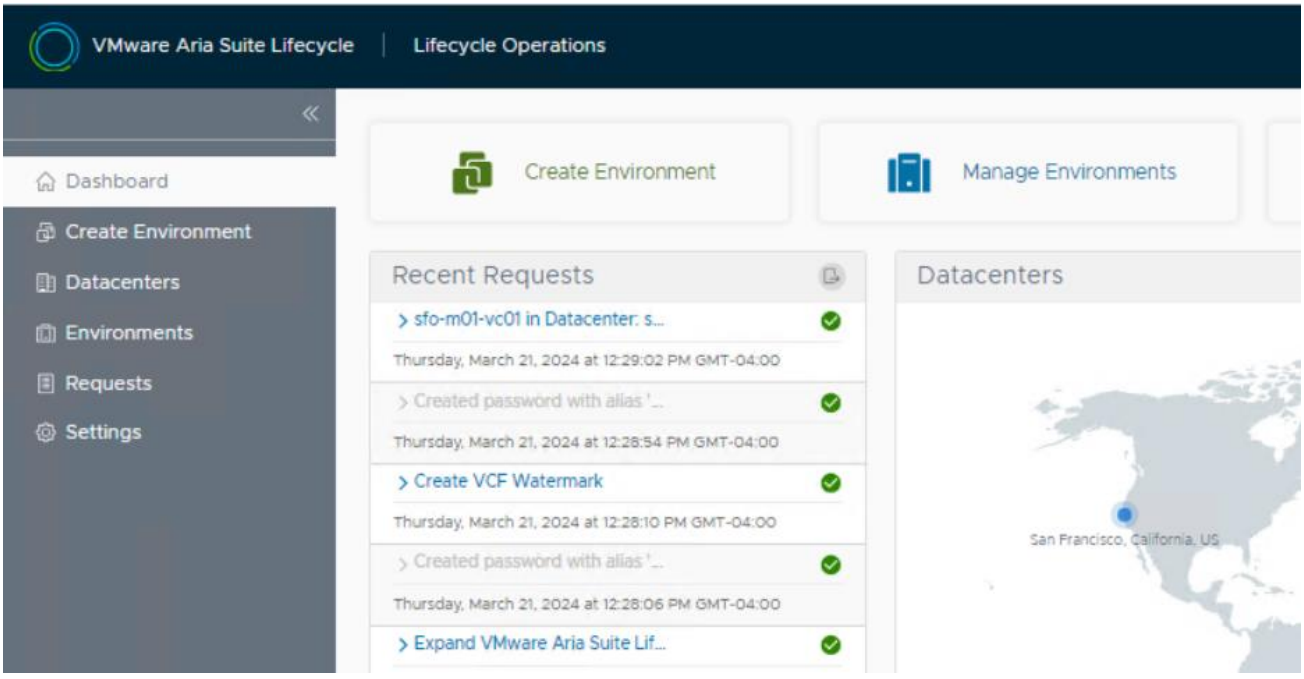

In the Lifecycle Operation dashboard go to Settings and select My VMware:

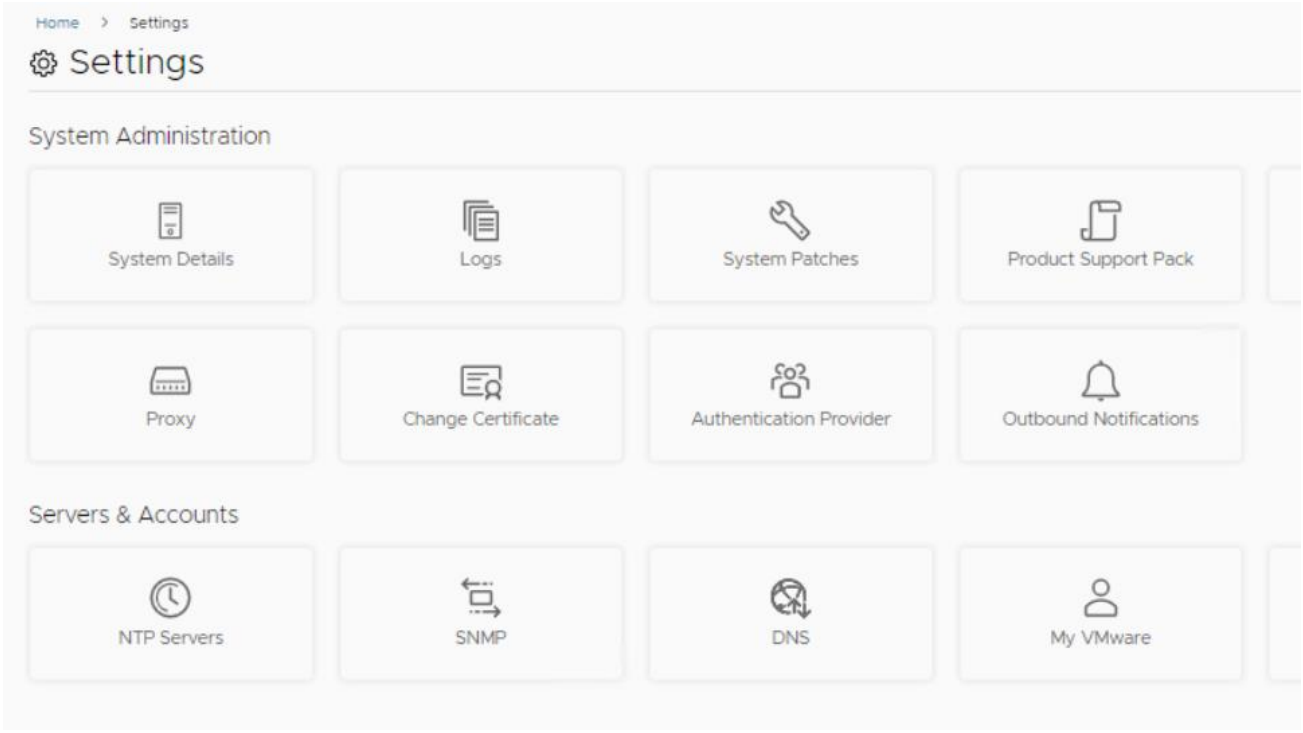

In the My VMware windows click on Add My VMware Account:

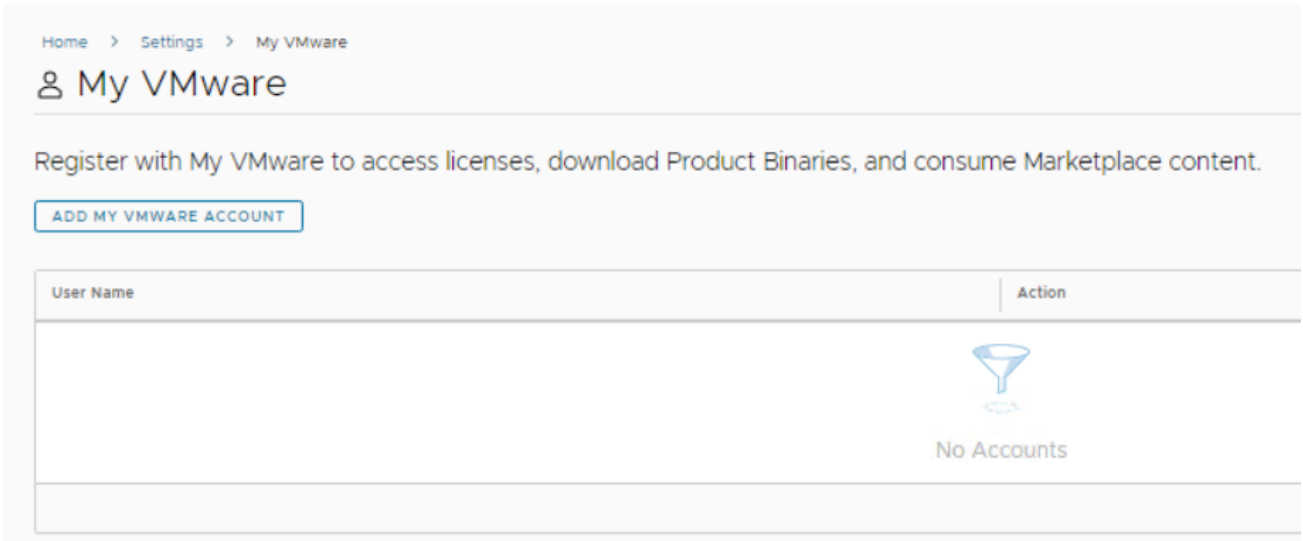

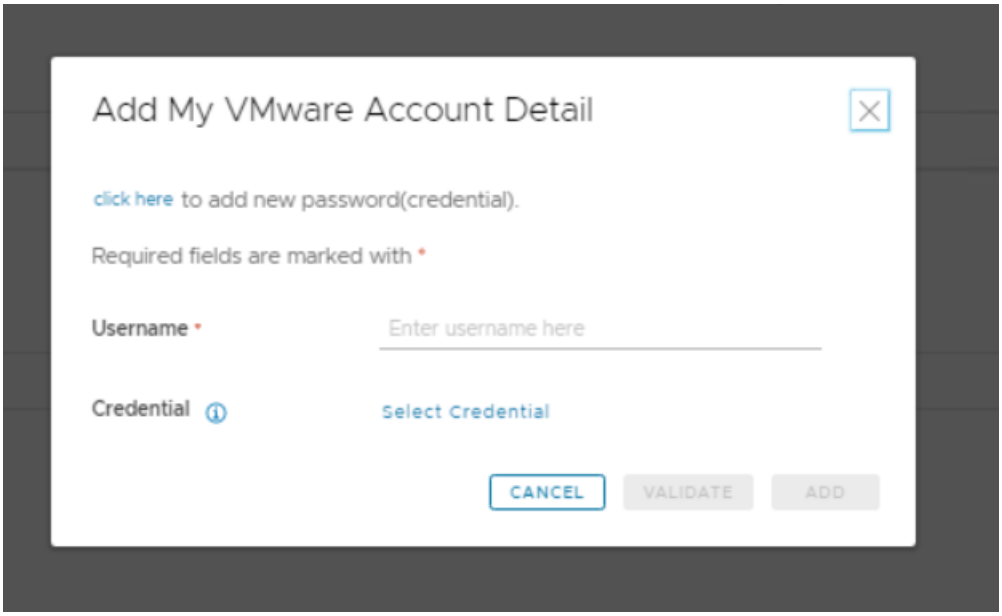

Click on ‘click here’ hyperlink to add new credentials:

Provide the following parameters:

- Password Alias
- My VMWare Account password
- My VMware Account username

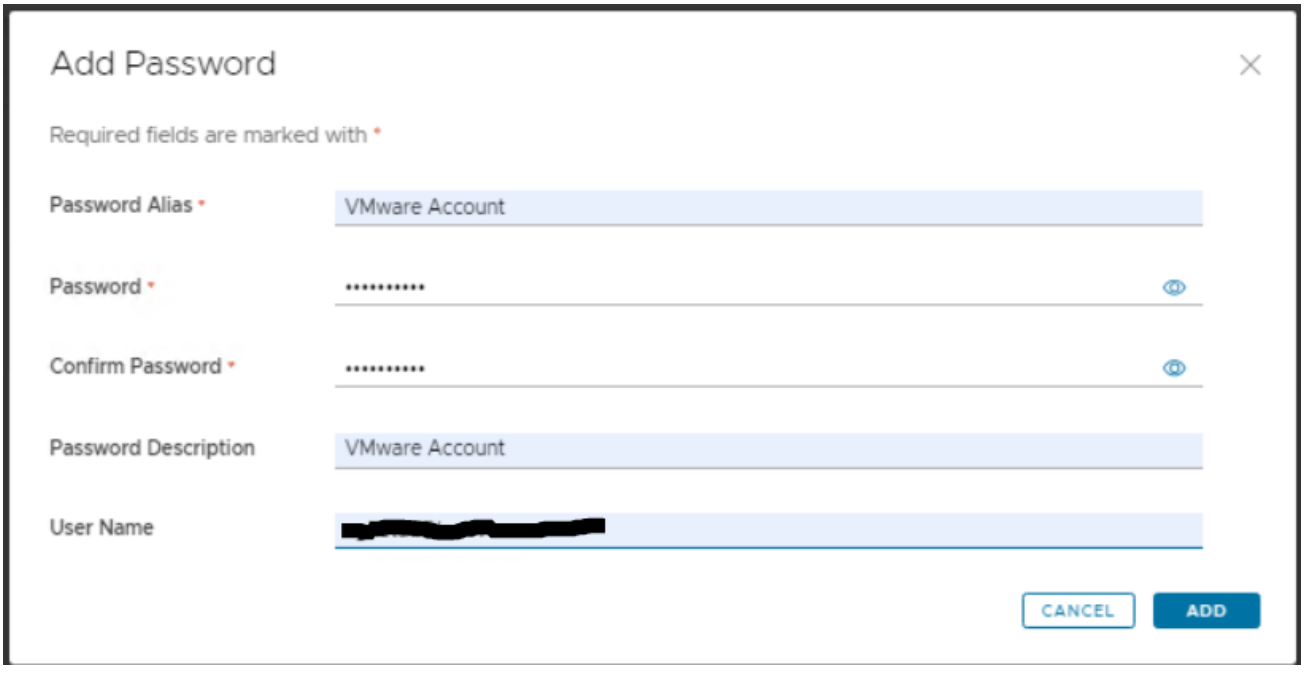

Provide the VMWare user account and click on Select Credential hyperlink:

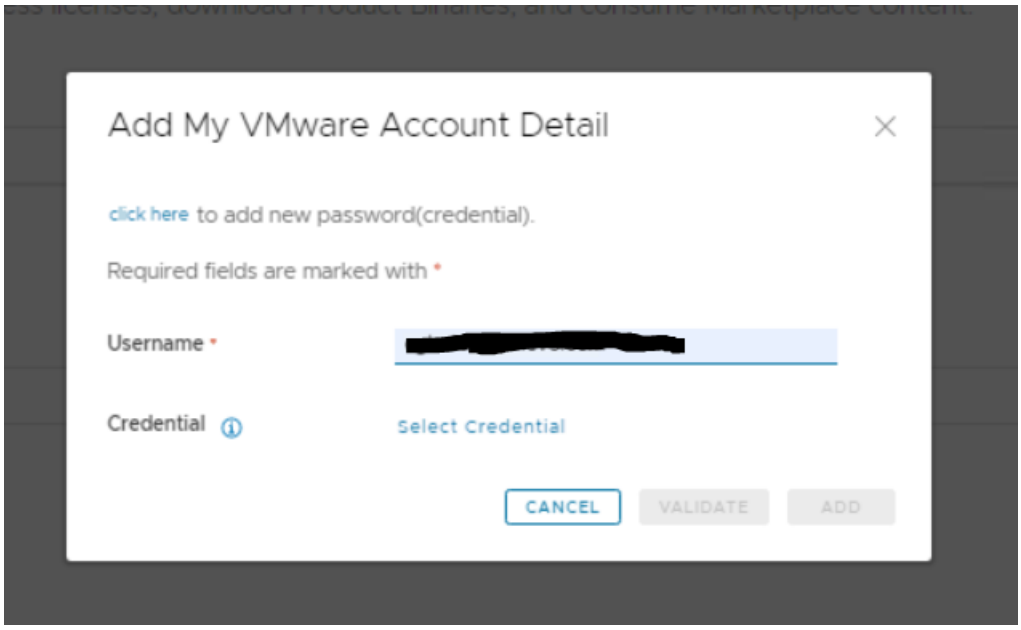

Click on the Password Alias created (VMware Account in this case):

Click on Validate button (vRLCM must have access to the Internet):

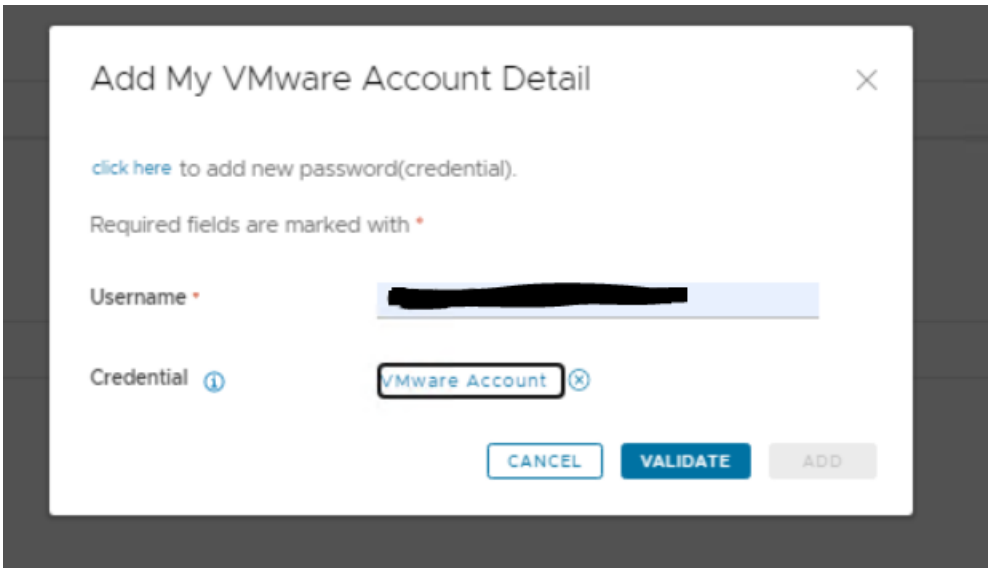

If the credentials are validated successfully click in Add button:

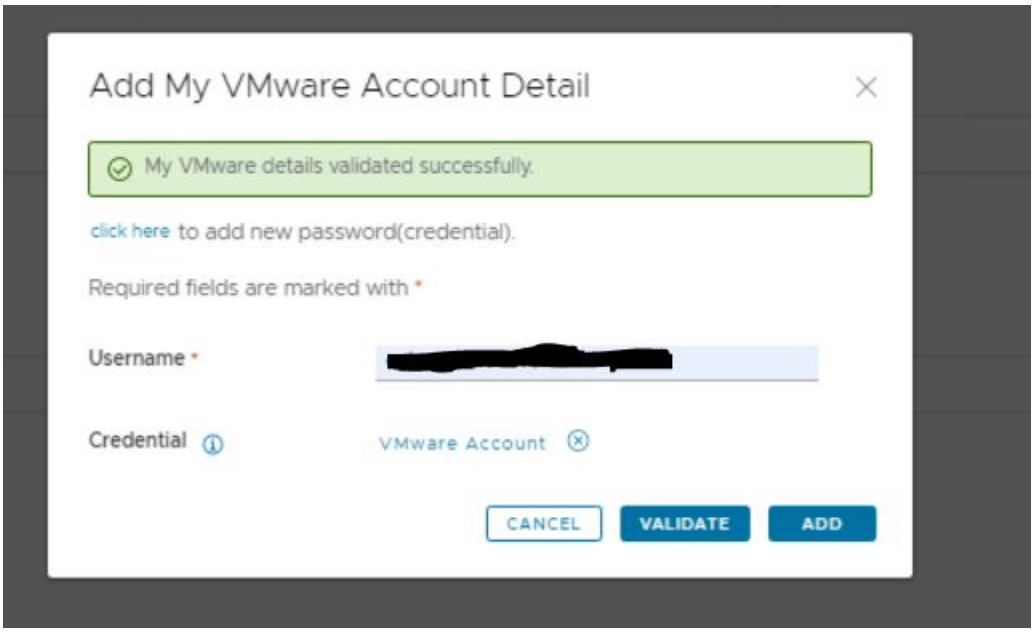

In the vRLCM Lifecycle Operations dashboard verify that the account has been successfully added. You can now add the required binaries for the Aria Suite components.

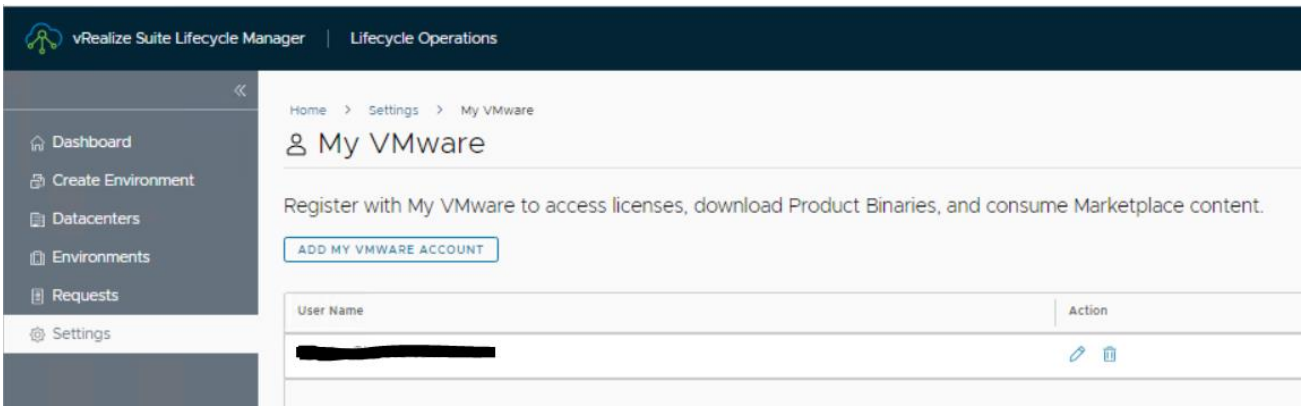

Deploy Aria Workspace One Access (Identity Management)

In the vRLCM Lifecycle Operations go to Settings and select Binary Mapping:

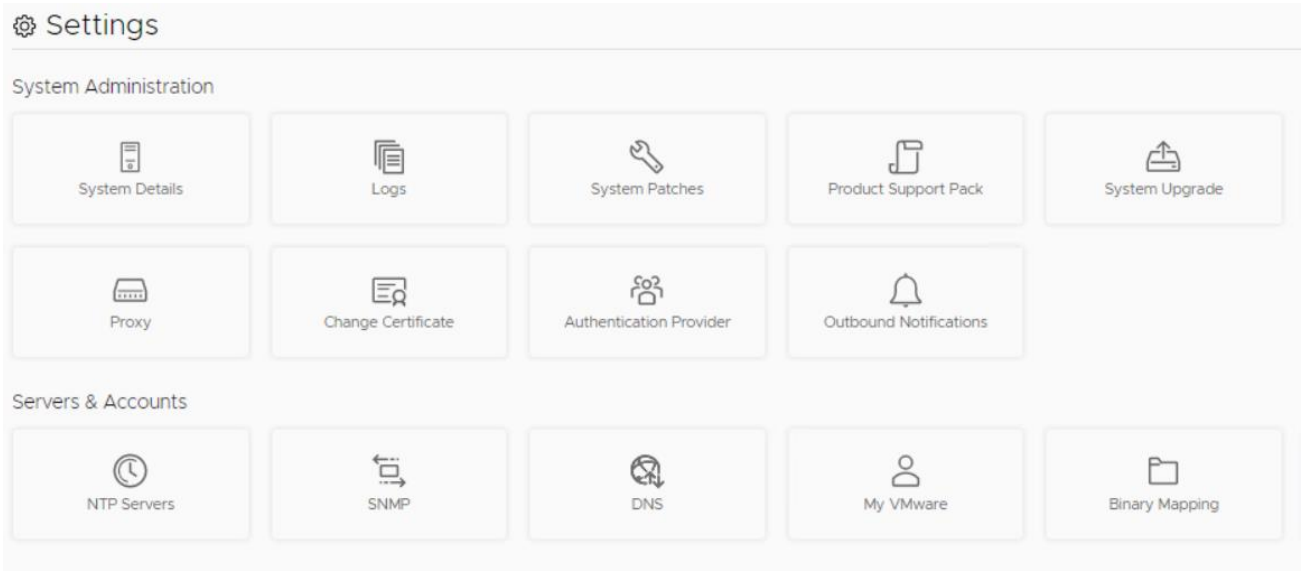

In the Binary Mapping windows click on Add Binaries:

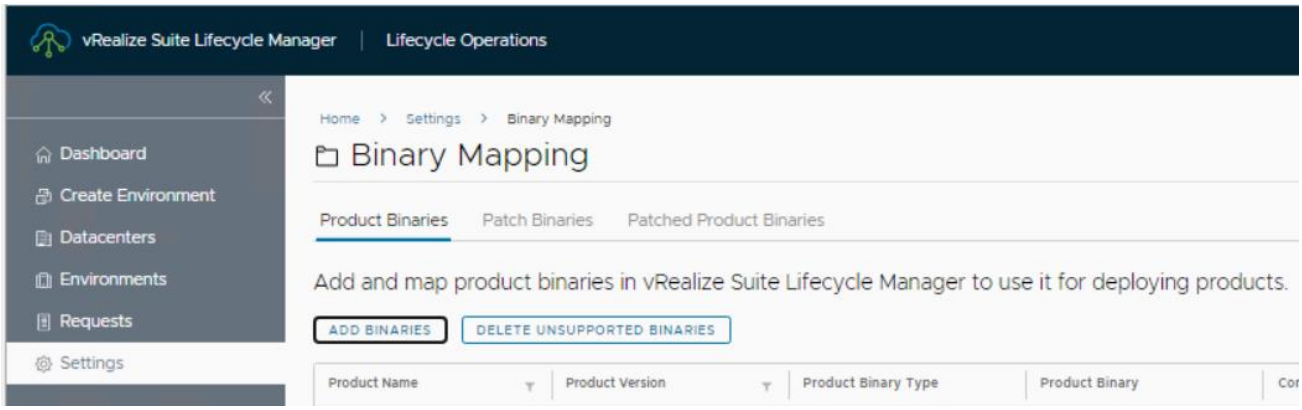

In the Add Product binary window select My VMware and click on Discover button:

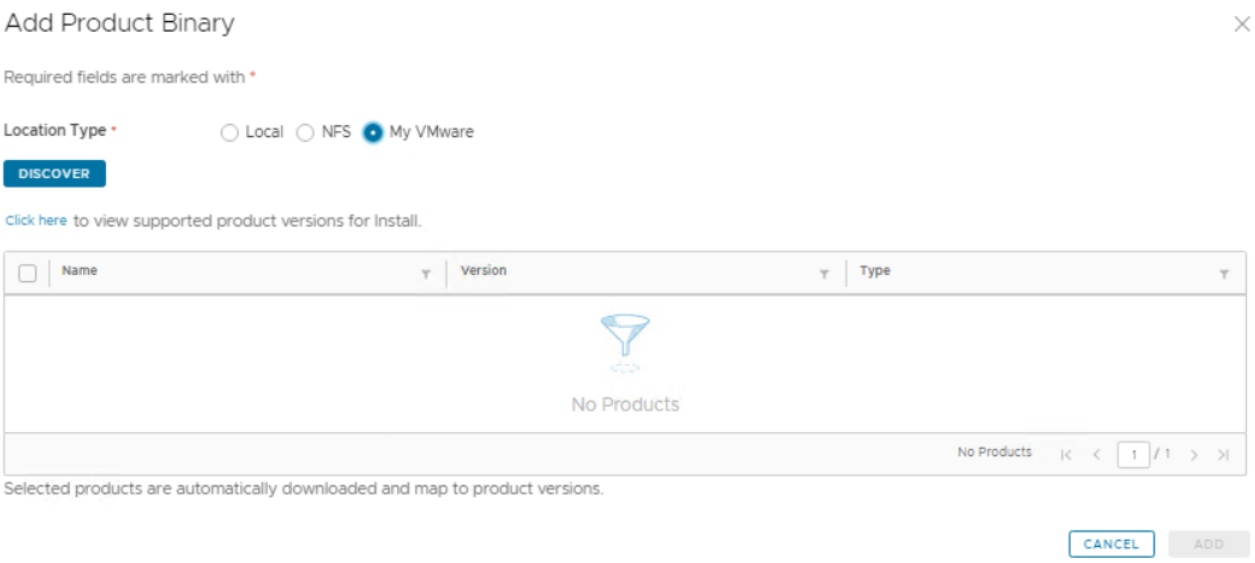

In the Add Product Binary windows select the VMware Identity Manager binaries and click Add:

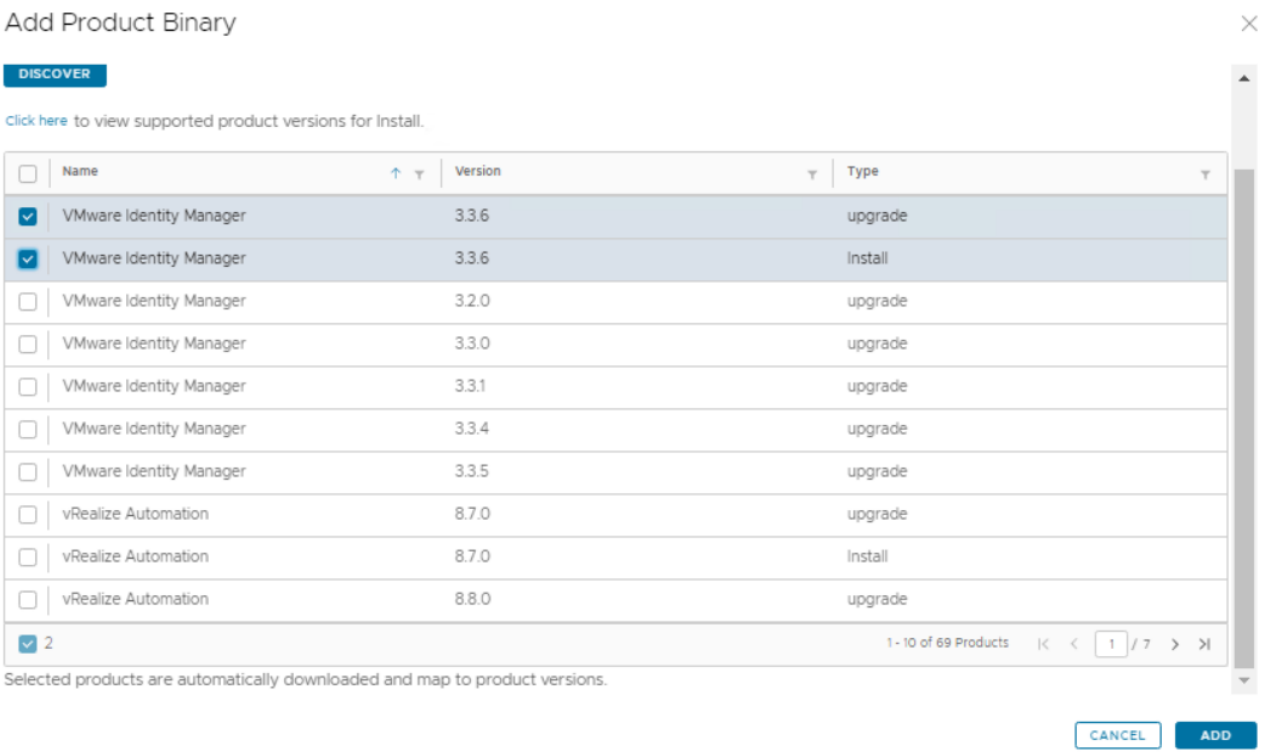

In the vRLCM Lifecycle Operation page click on Requests to check the status of the request:

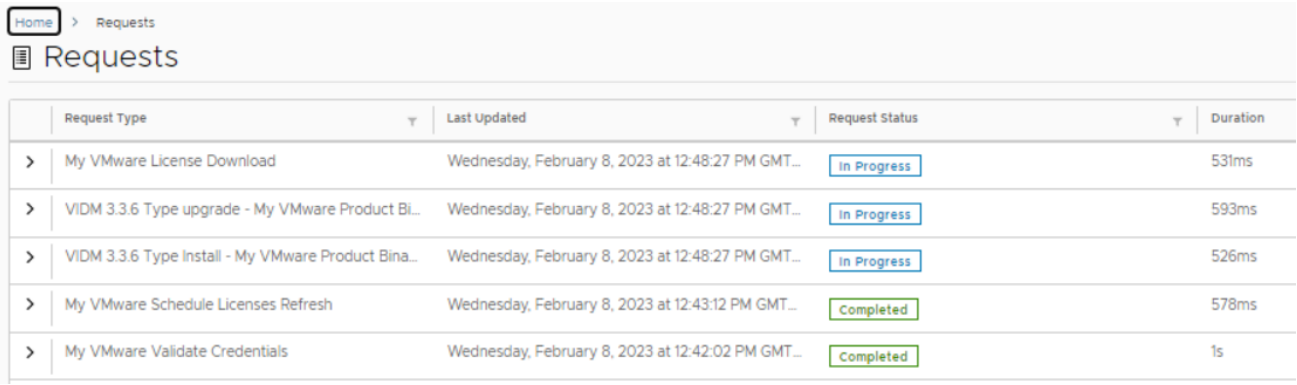

After the requests are in Completed status verify that the binaries are successfully mapped:

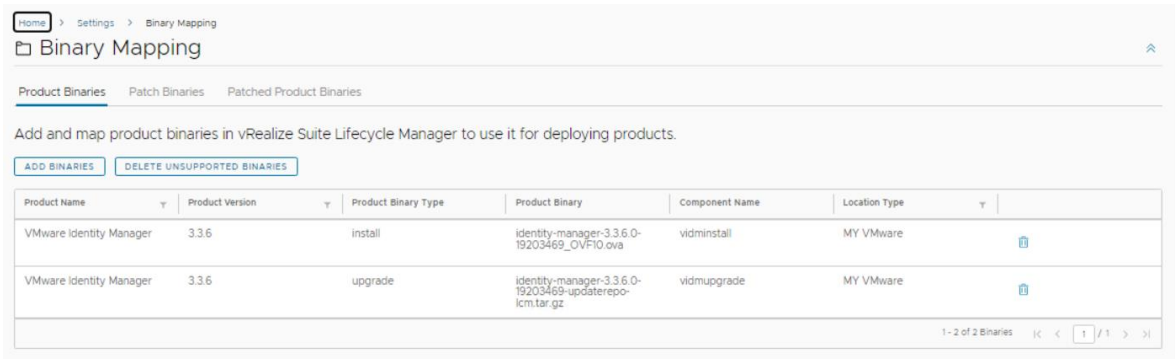

In the SDDC Manager console go to Aria Suite and select Deploy Workspace ONE Access. This will trigger a ‘Create a new globalenvironment’ request in the vRLCM.

Since it’s a pre-validated request we cannot change the Environment name (in vRLCM Lifecycle Manager). Click on the + (plus) sign next to Select Default Password hyperlink to add an administrator account:

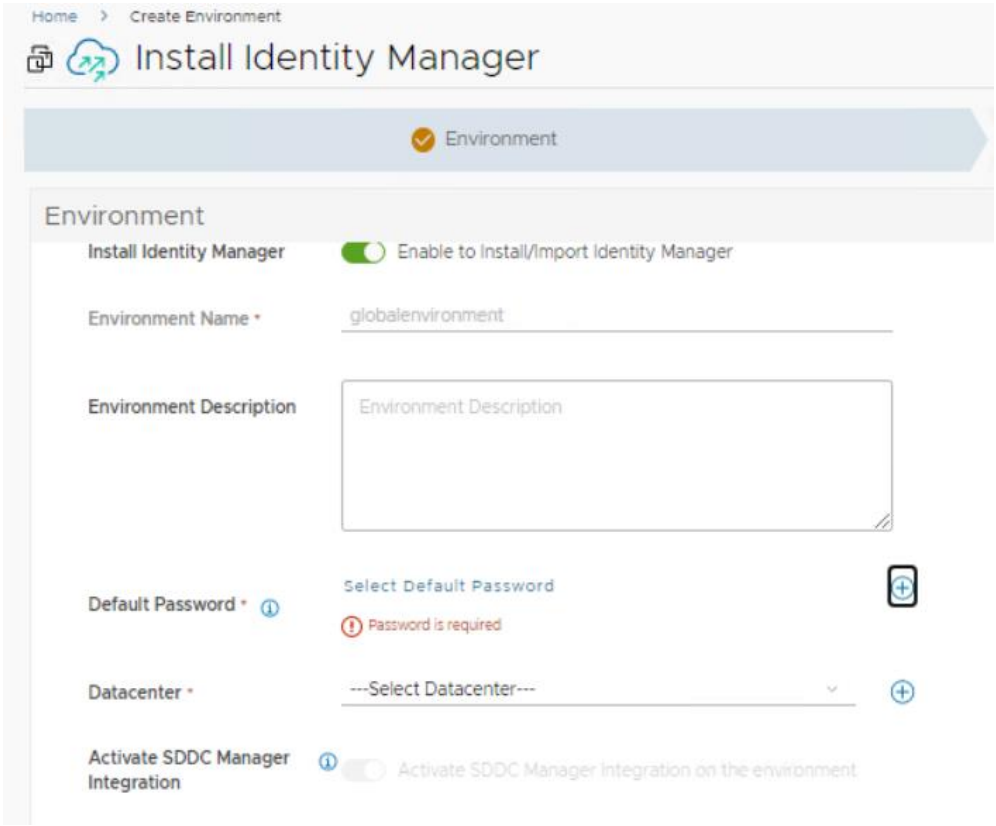

In the Add Password windows provide a Password Alias, a password, a password description and a user name then click Add:

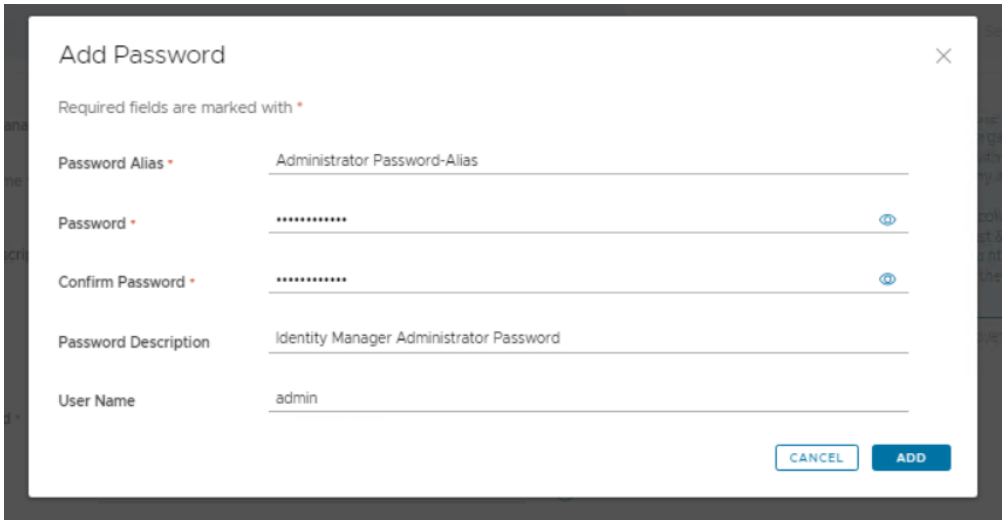

Select the newly created Password:

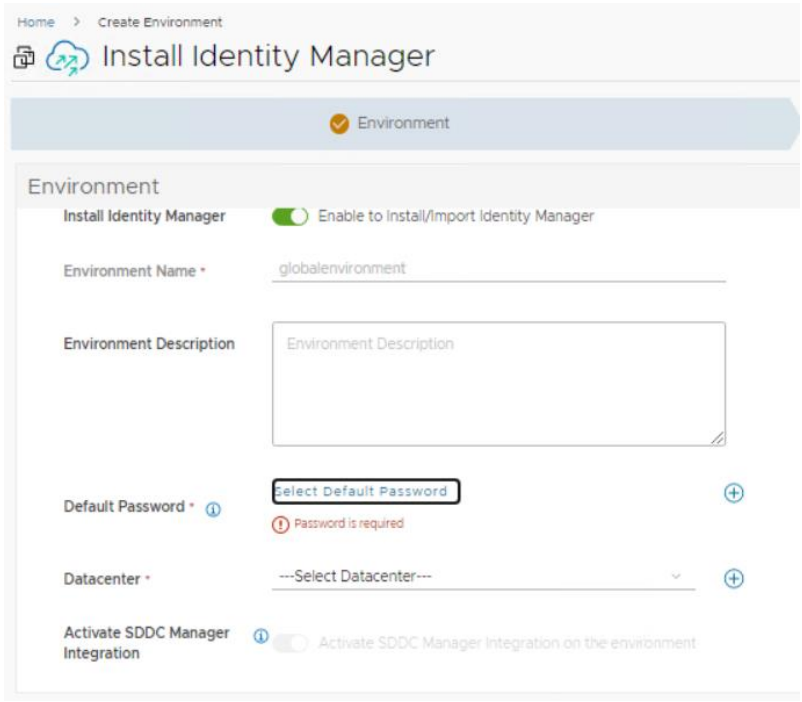

In the Default Password window select the Administrator Password-Alias created previously:

Select the Datacenter (vSphere datacenter associated with the Workload Management cluster) and click Next:

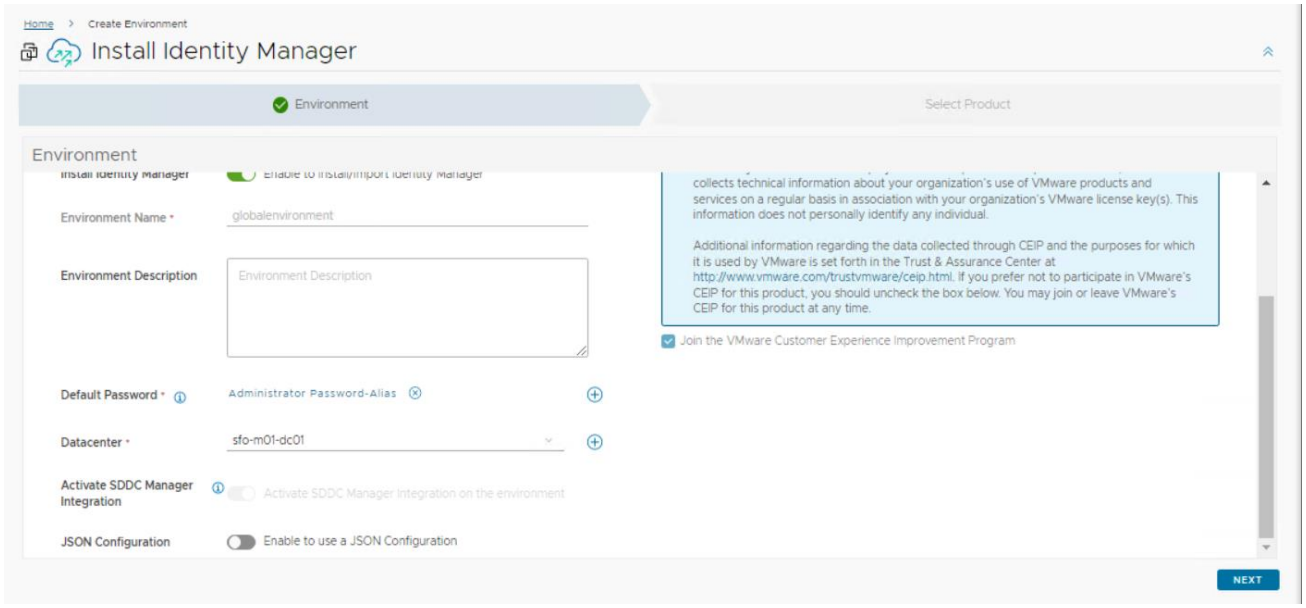

In the Select Product window, click on the checkbox from the VMware Identity Manager, select New Install, select the version and the Deployment Type (Standard – non clustered) then click Next:

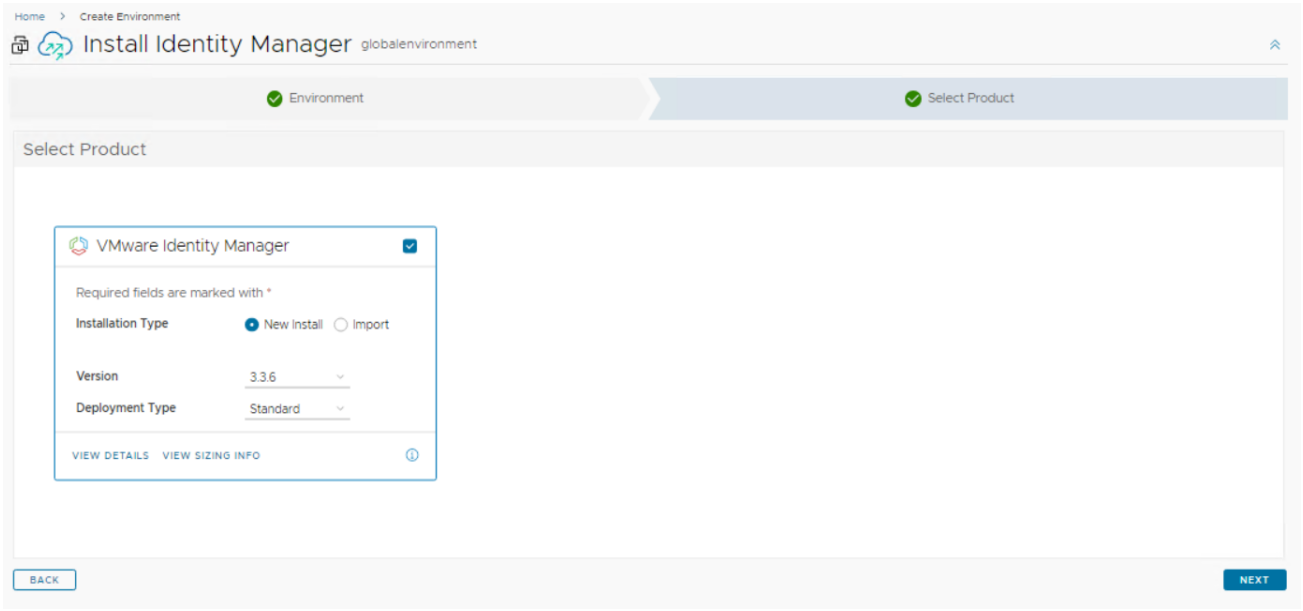

In the EULA windows check the ‘I agree to the terms & conditions’ button and click Next:

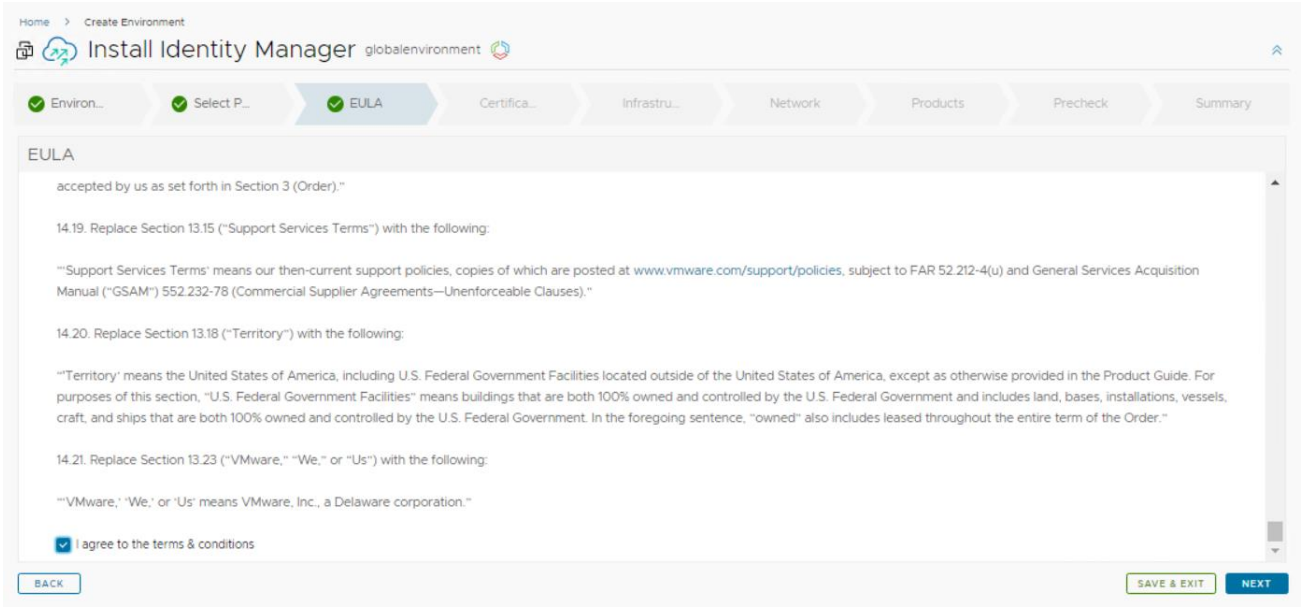

In the Certificate windows click on the + (plus) sign and select Generate Certificate:

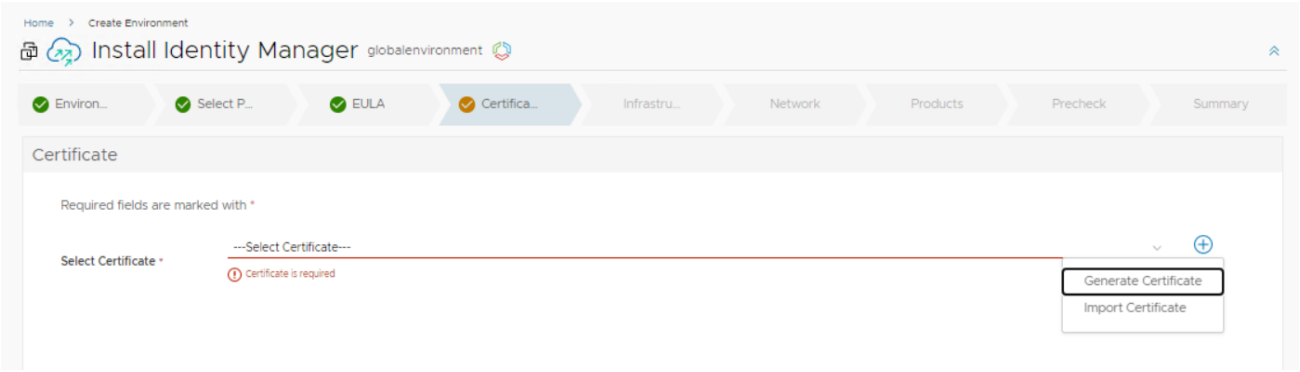

In the Generate Certificate, fill in the required fields:

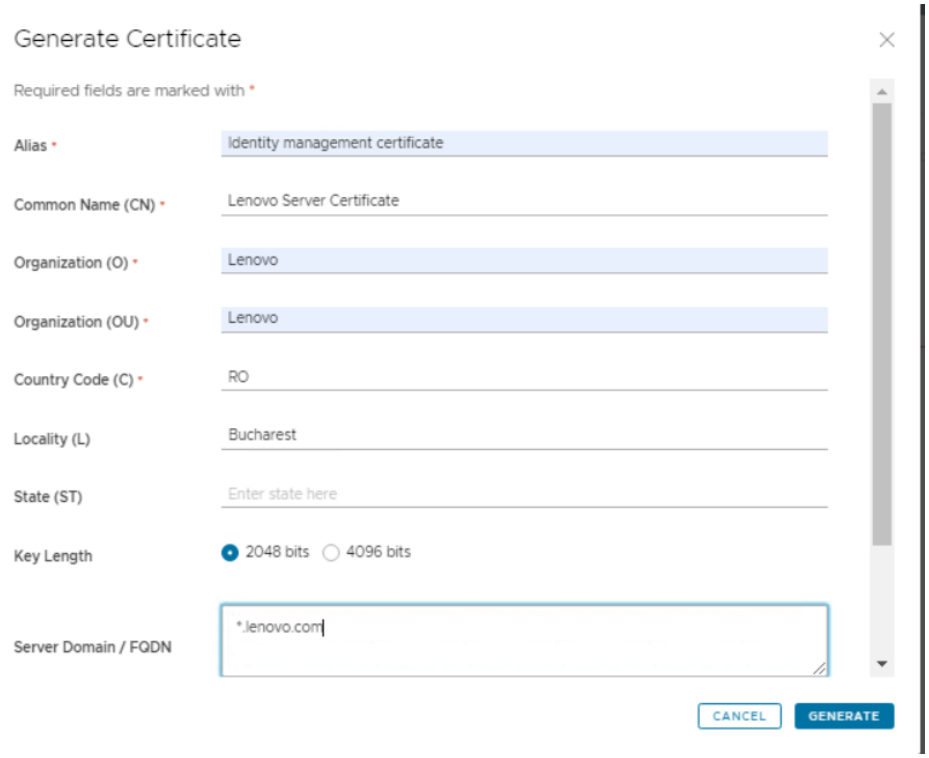

Select the newly created certificate then click Next:

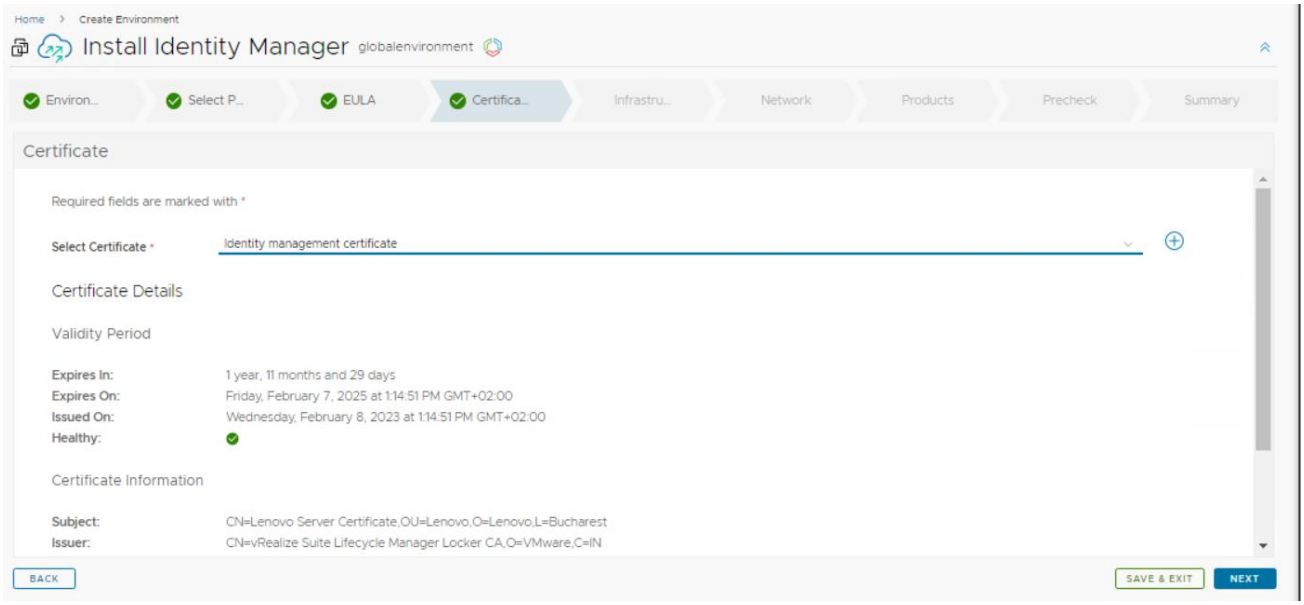

In the Infrastructure window select the Cluster (could be already selected), the Resource Pool (could be already selected) and the Disk Mode (Thin in this case) then click Next:

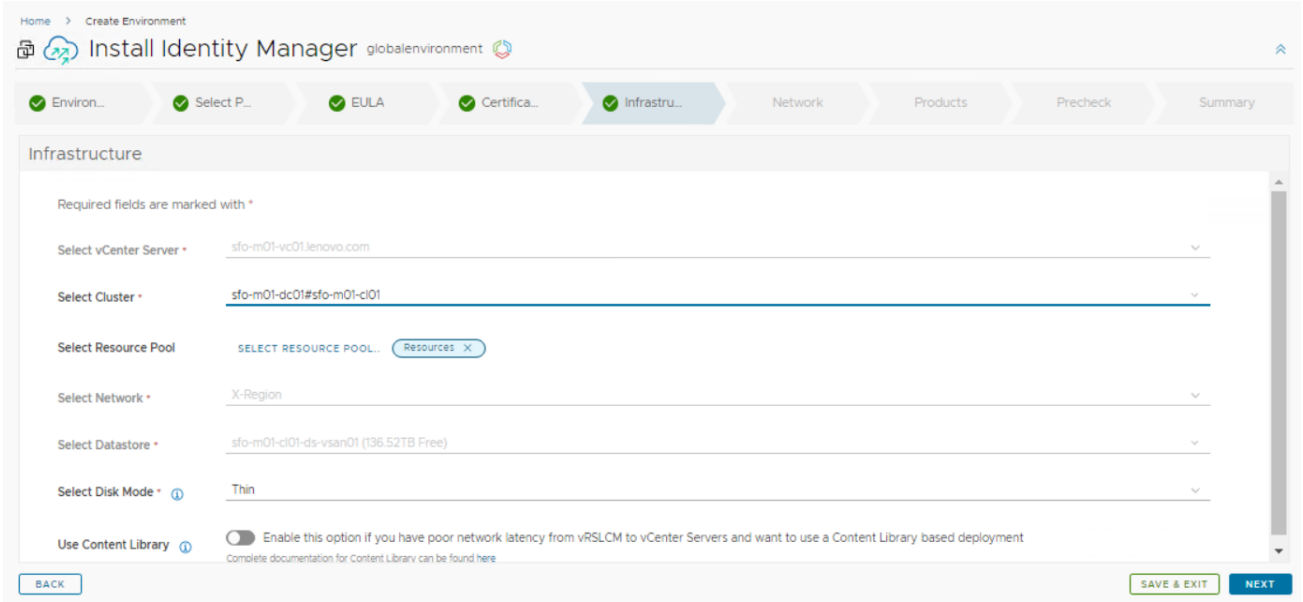

In the Network window verify if the X-Region is selected for the deployment and a DNS server is added then click Next:

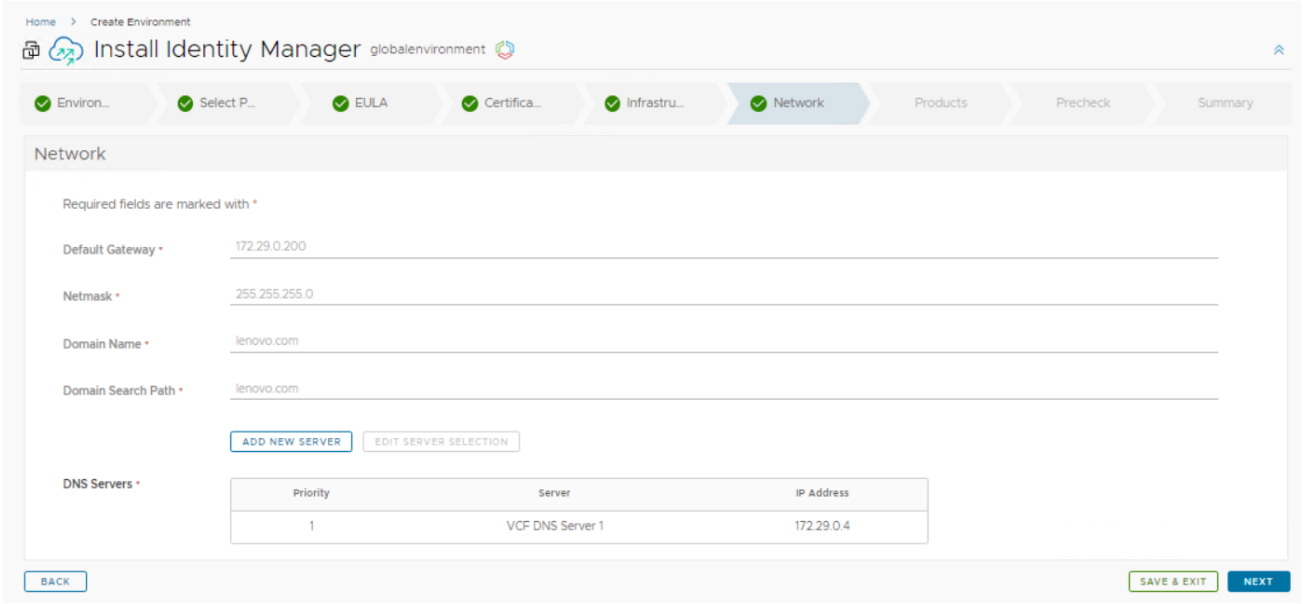

In the Products window select the Certificate, choose a Node Size (Medium)

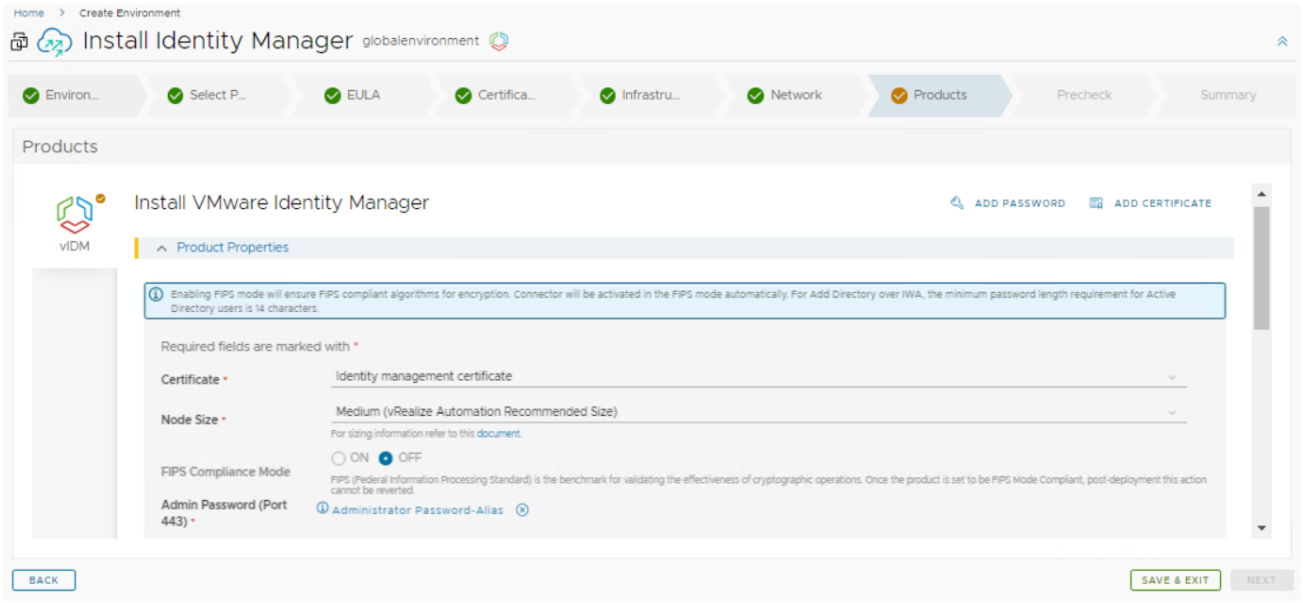

In the same window select the Admin Password (port 443), the default configuration admin e-mail, the default configuration admin username and the default configuration admin password:

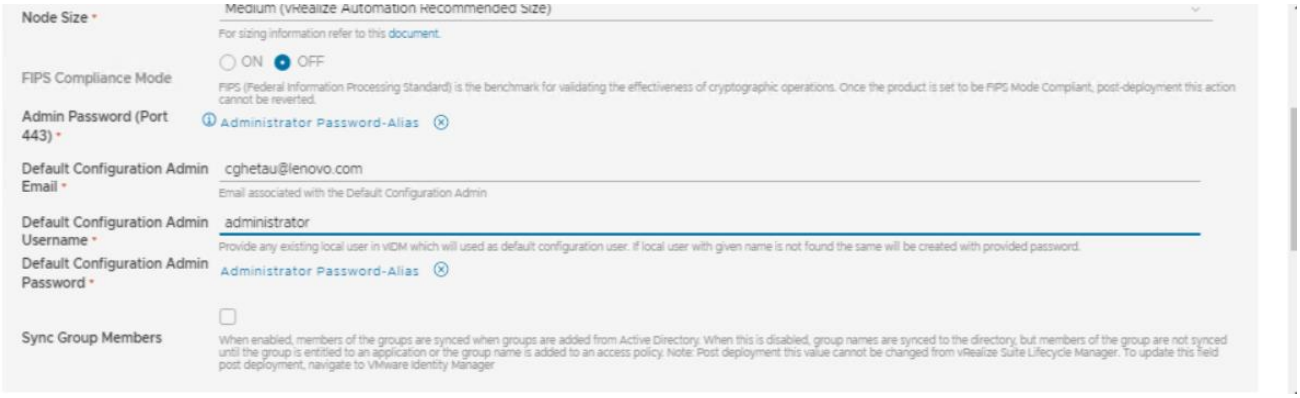

In the same window fill in the VM Name (as it will appear in the vSphere), the FQDN of the appliance and the IP address, then click Next:

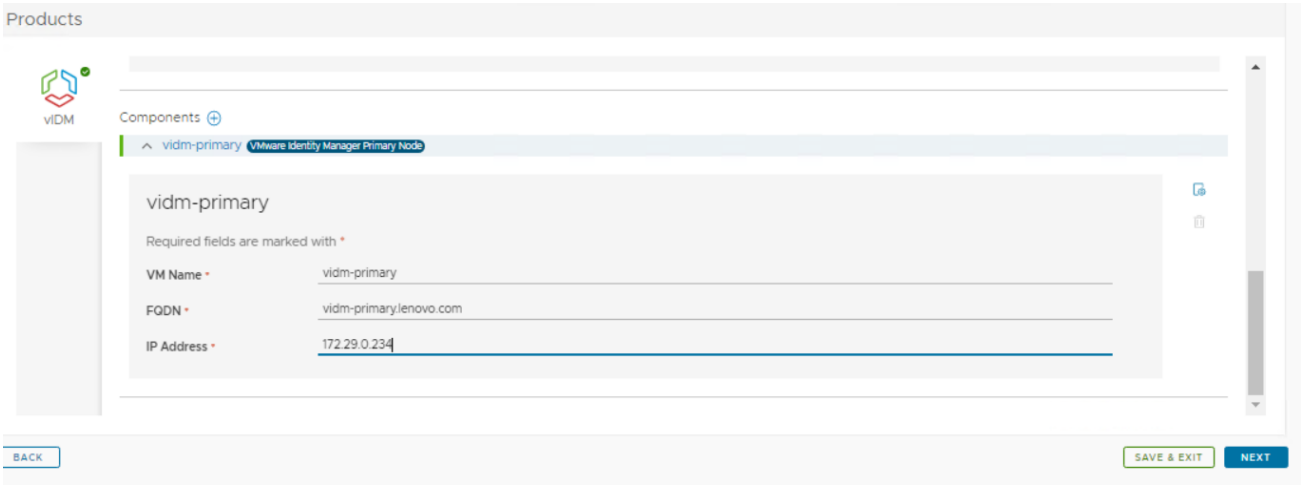

In the Precheck window click on the Run Precheck button:

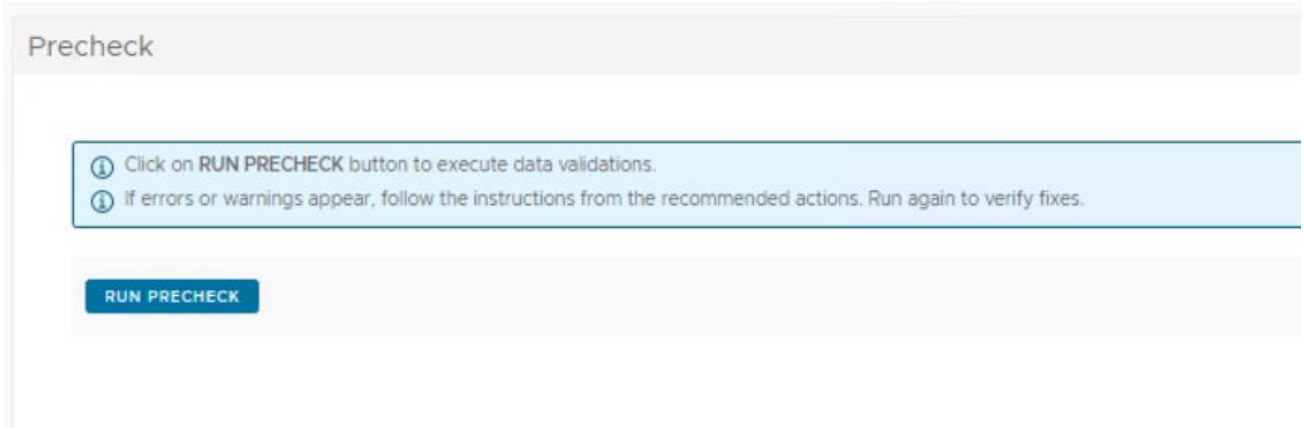

After all validations are passed click on the Next button:

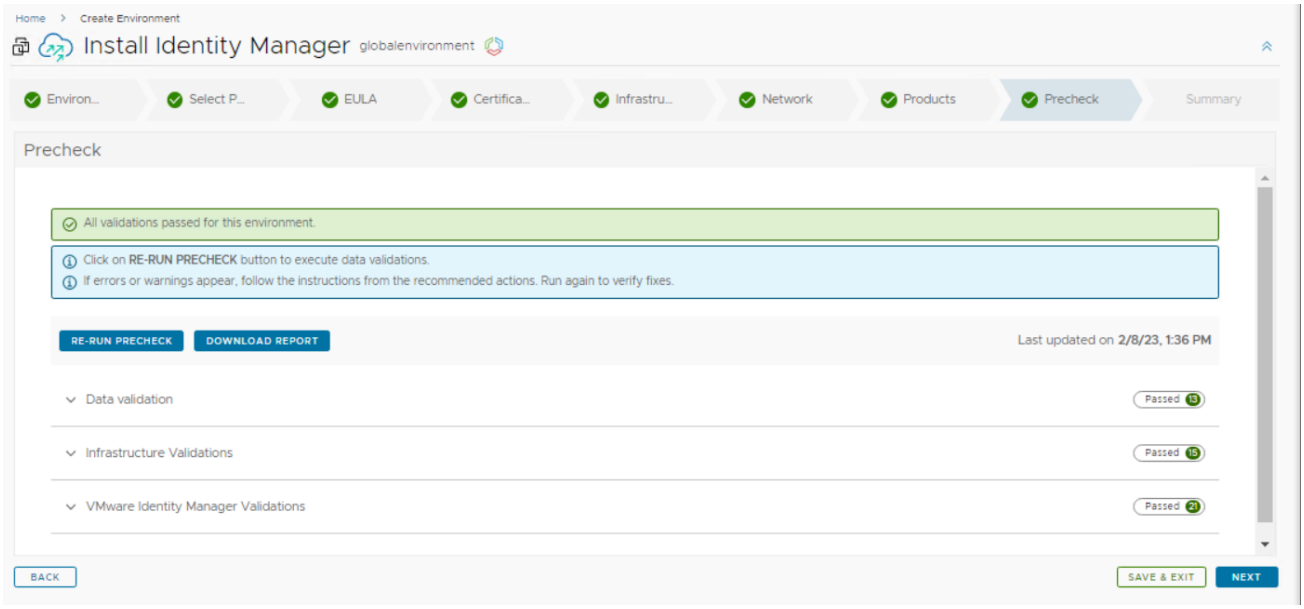

In the Summary window verify if all the data have been correctly provided then click Submit:

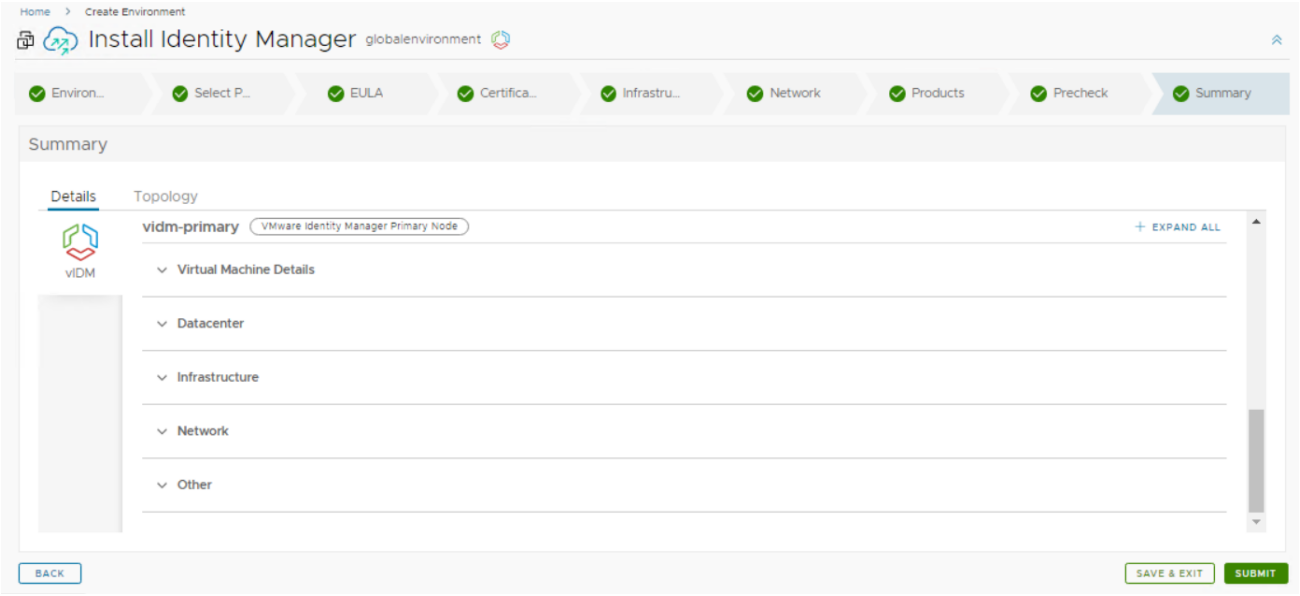

The Request Details window will be automatically switched to in order to monitor each deployment Stage:

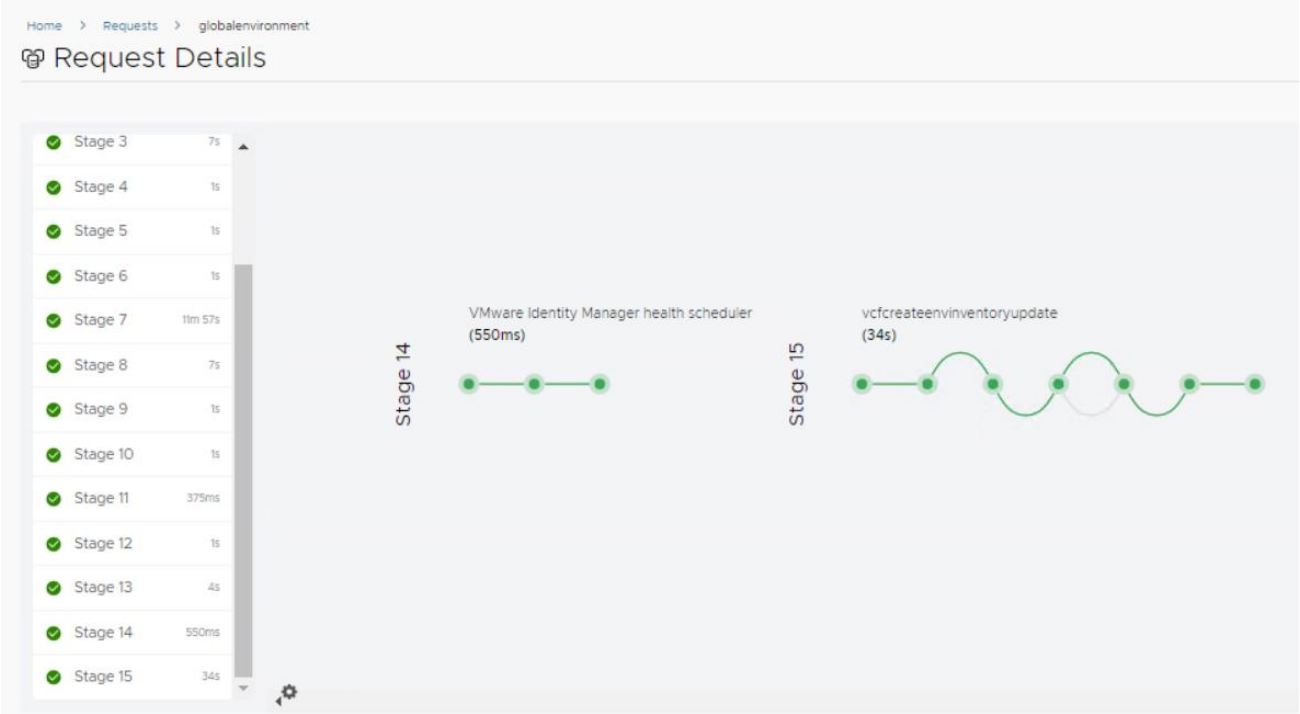

Check the SDDC Manager Console to verify that the Aria Workspace One Access was successfully deployed and shown as Active:

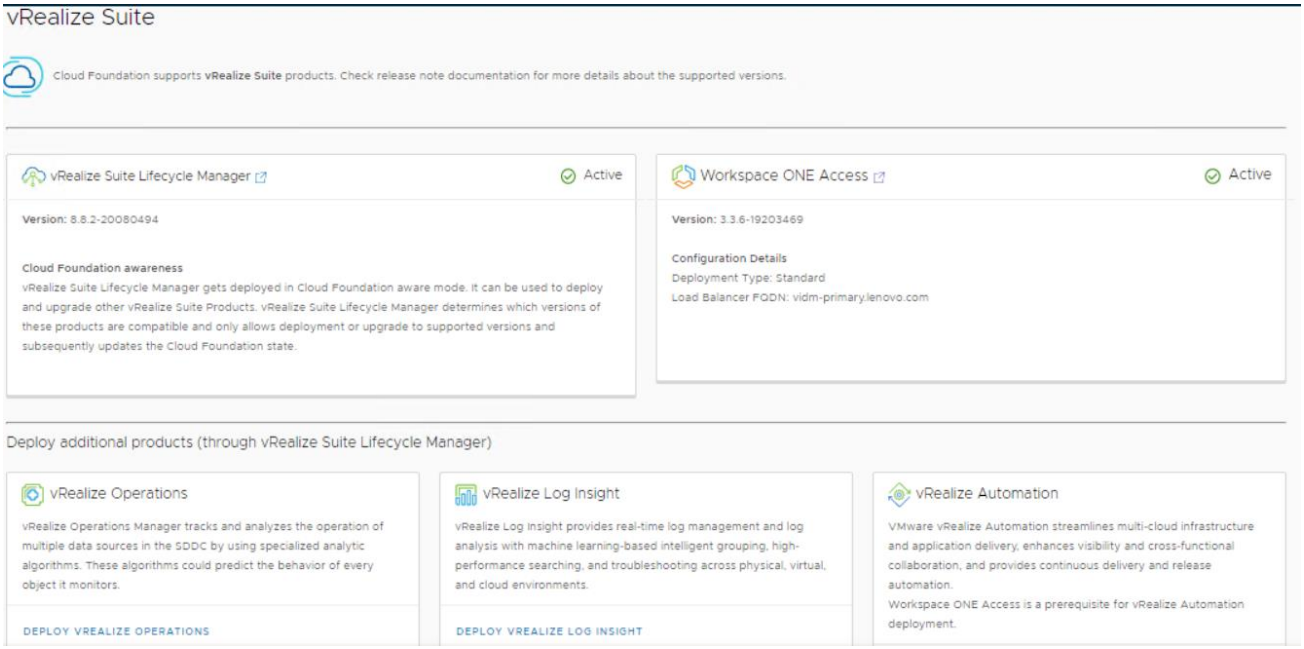

Deploy Aria Operation Manager

Add the Aria Operations binaries from Aria LCM Lifecycle Operations:

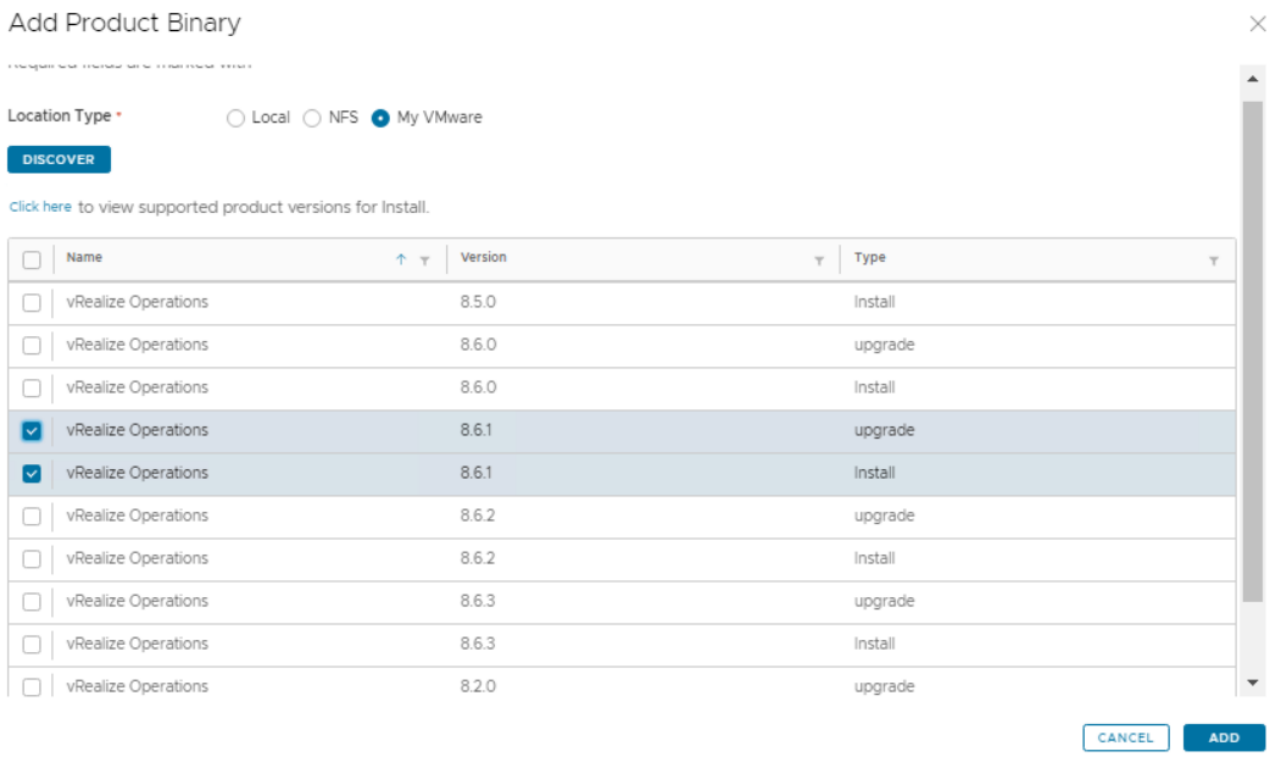

Check the request in the vRLCM Request window:

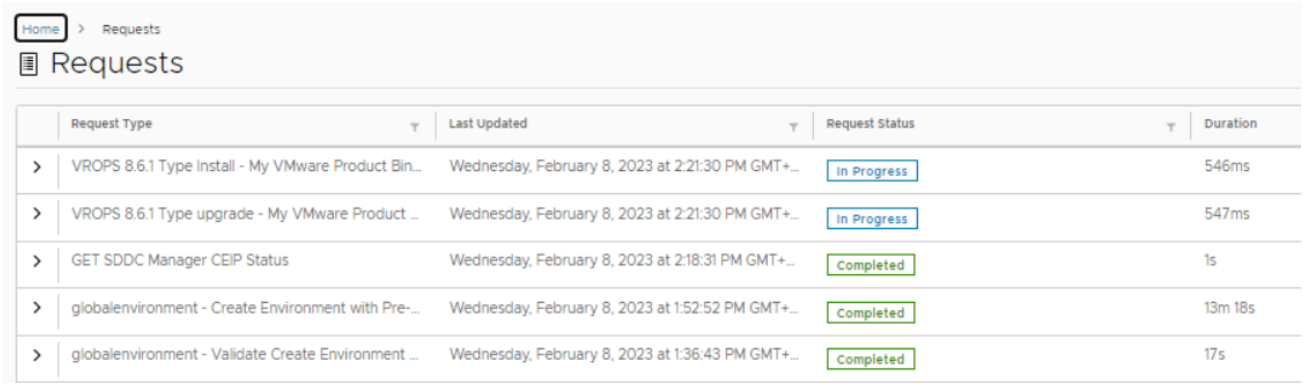

In the SDDC Manager go to Aria Suite and click on the Deploy Aria Operations:

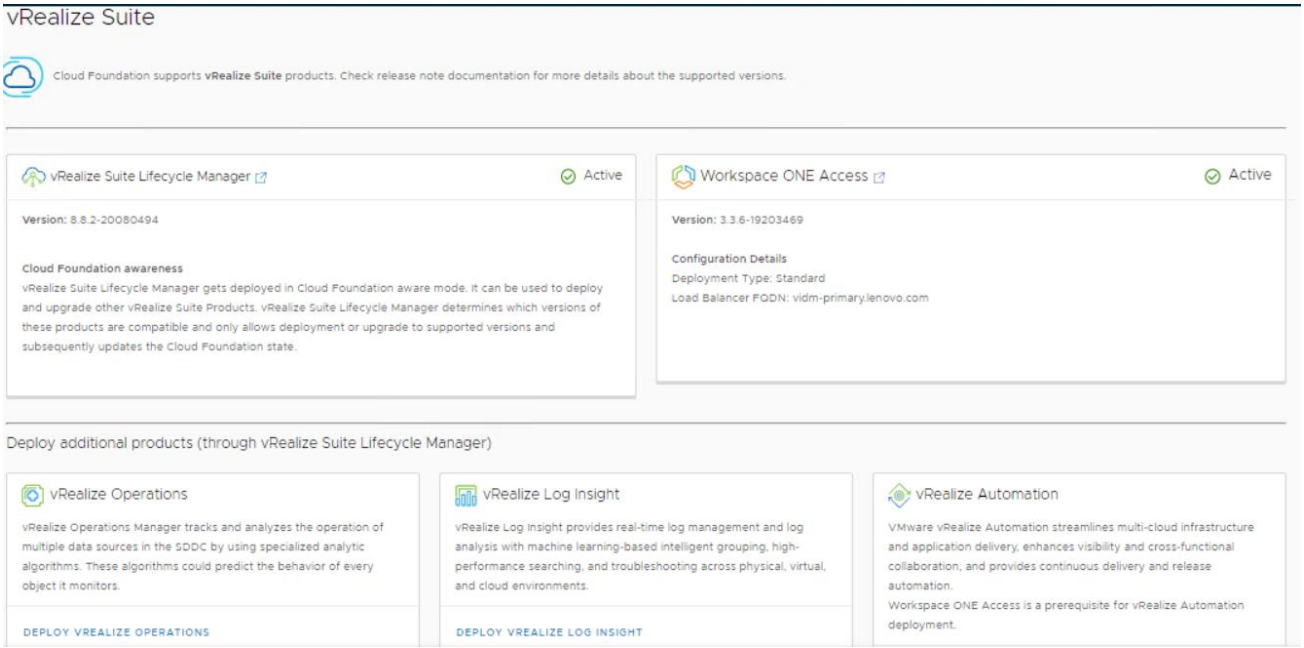

This will trigger a Create a new Environment in the vRLCM Lifecycle Operation. The same can be achieved by choosing Create Environment from vRLCM|LO Dashboard:

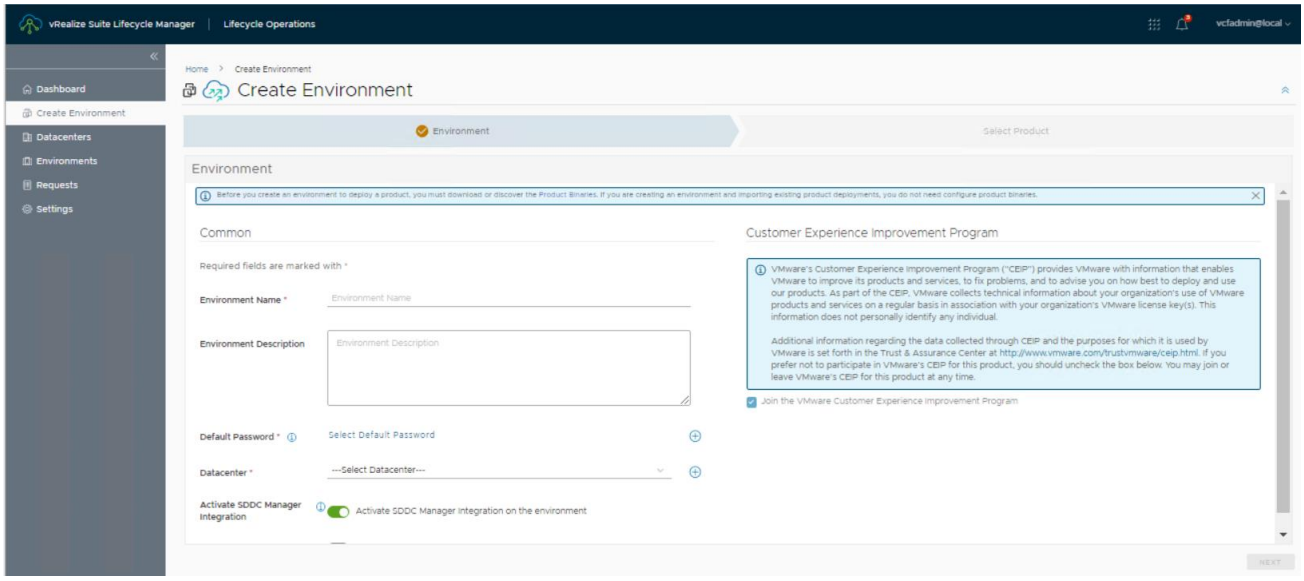

When using the SDDC Manager to deploy vROps, the Environment Name and the Datacenter are already filled in:

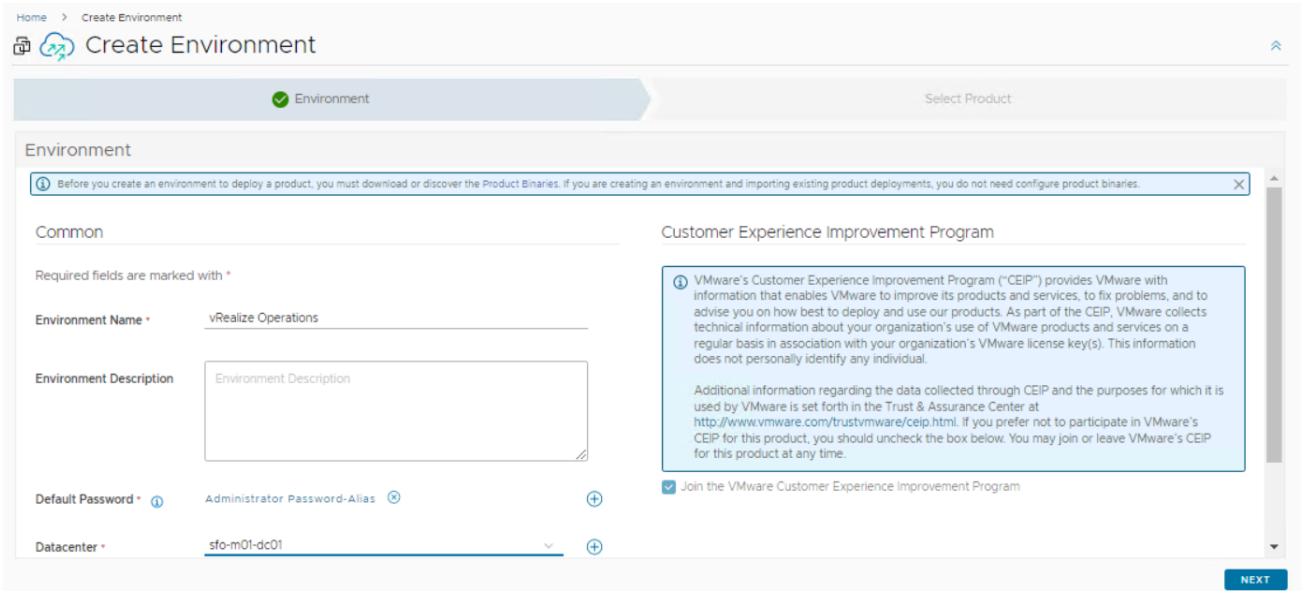

Select Activate SDDC Manager Integration on the environment and click Next:

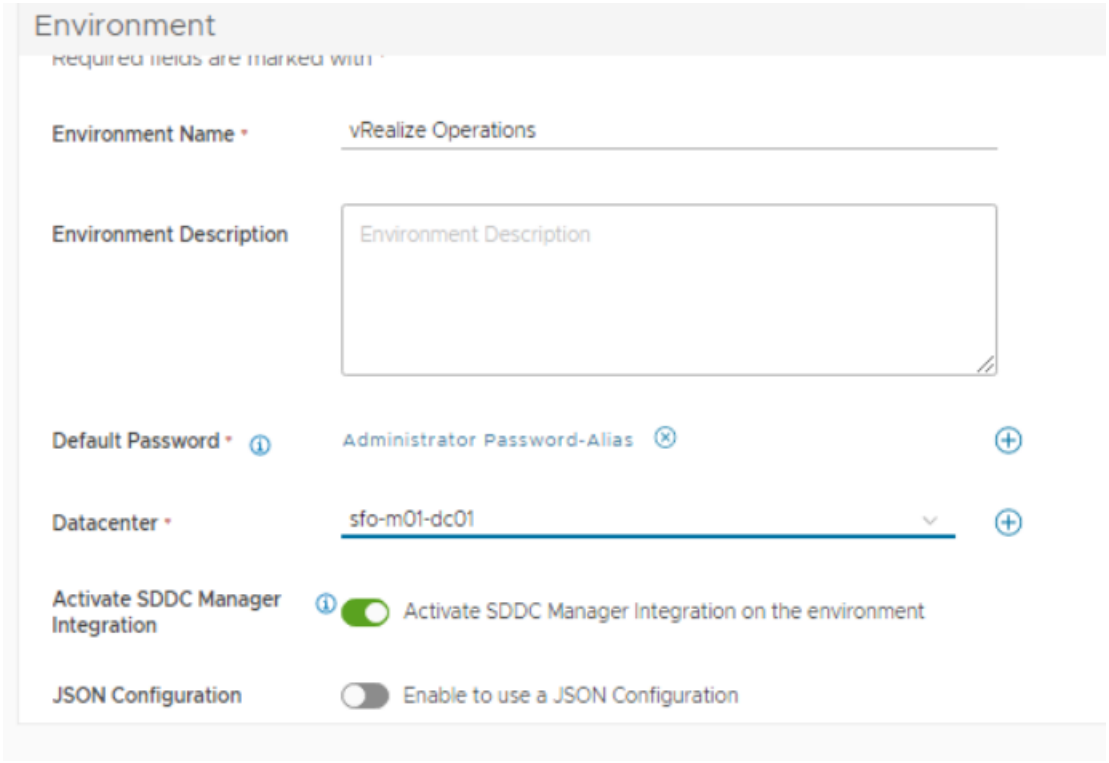

Select the Aria Operations by checking the box associated to it, select New Install, choose the version, the Deployment Type (Medium) and the Node Count (2 in this case) then click Next:

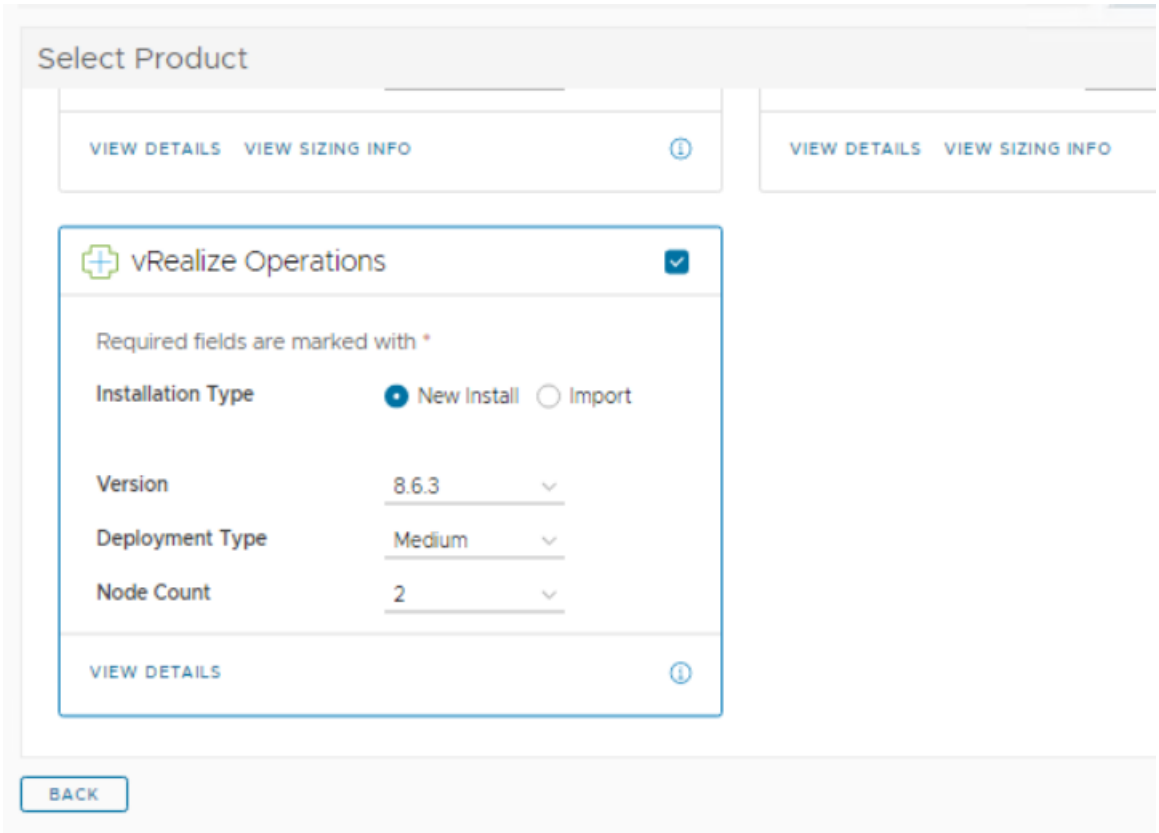

Select the ‘I agree to the terms and conditions’ and click Next:

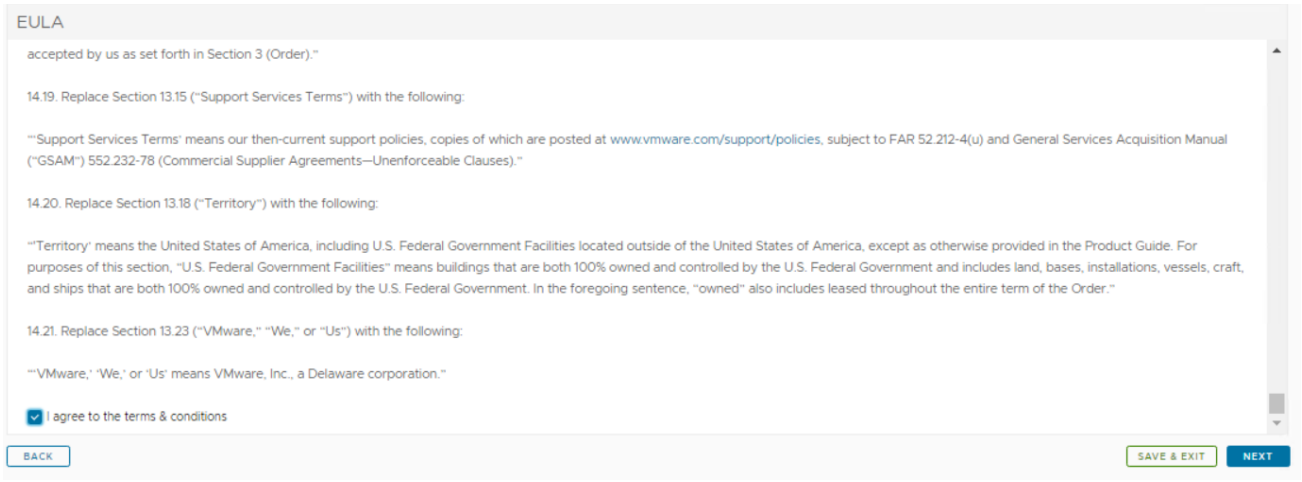

In the License window click on Add to add an existing license:

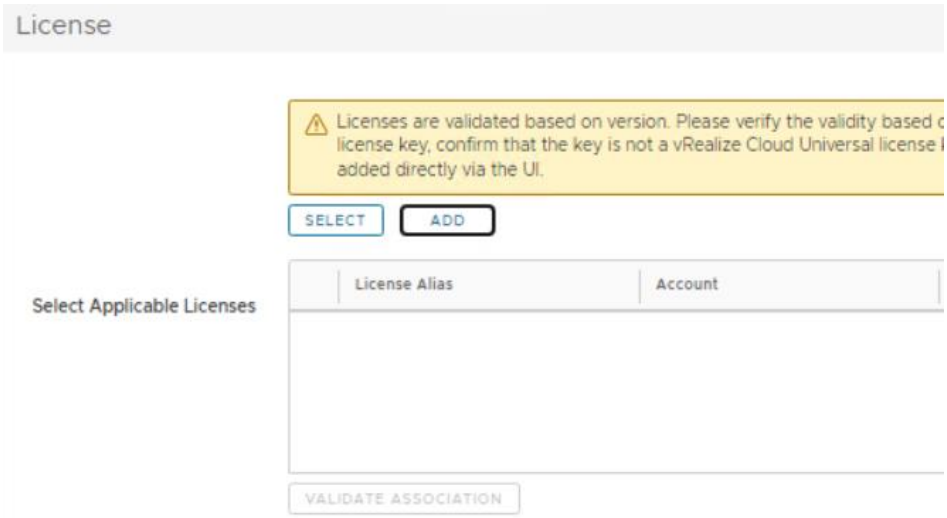

In the Add License window type in a License Alias and provide the License Key for Aria Suite then click on Validate button then click on Add button:

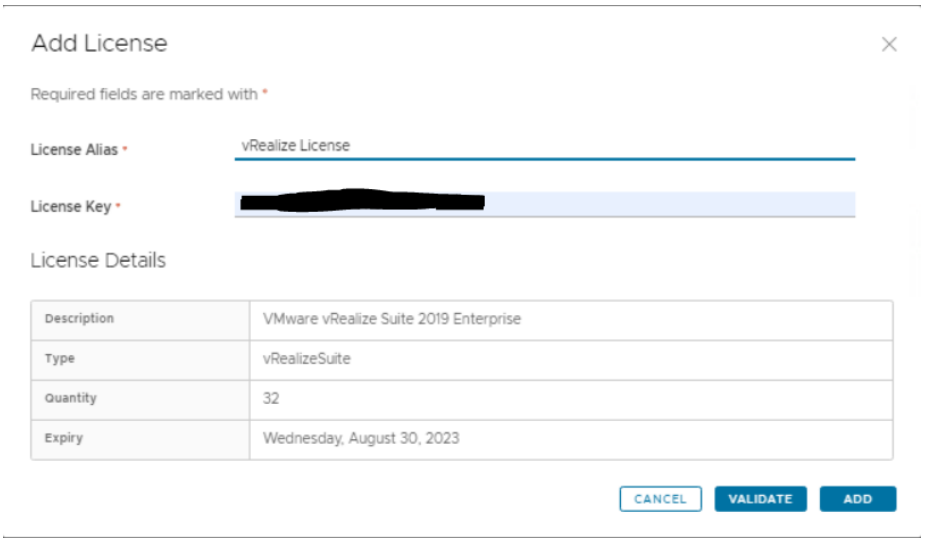

In the same window click in Select button to select the newly add license:

In the Select Applicable Licenses select the Aria License and click Update:

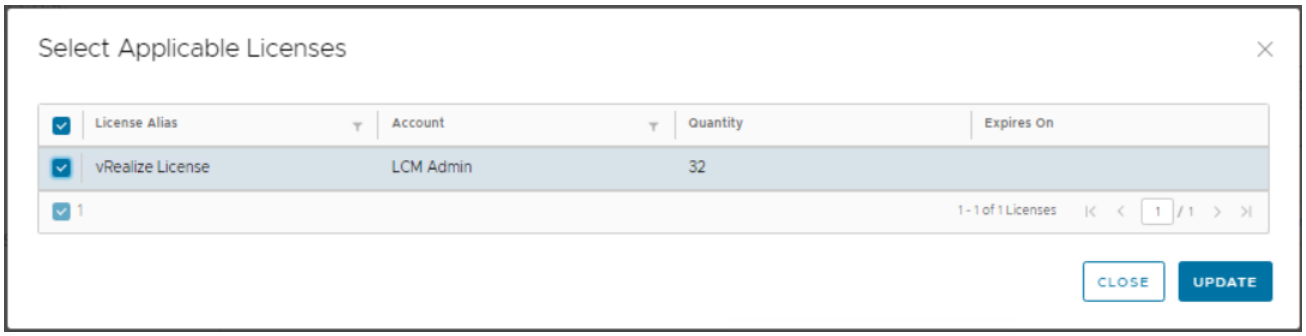

After the license is selected click on Validate Association then click Next:

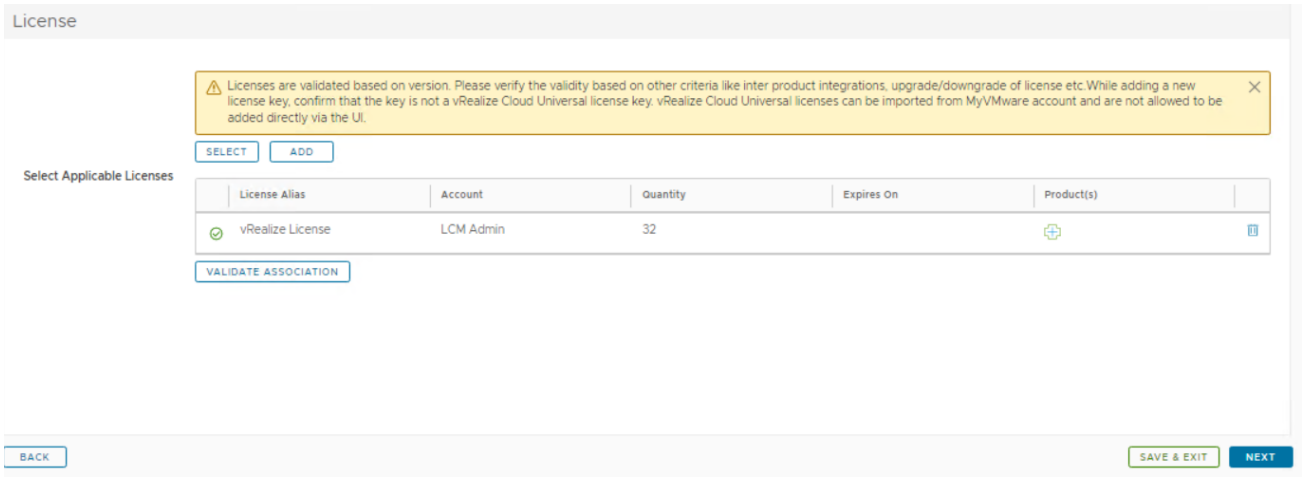

In the Certificate window click the + (plus) sign and select Generate Certificate:

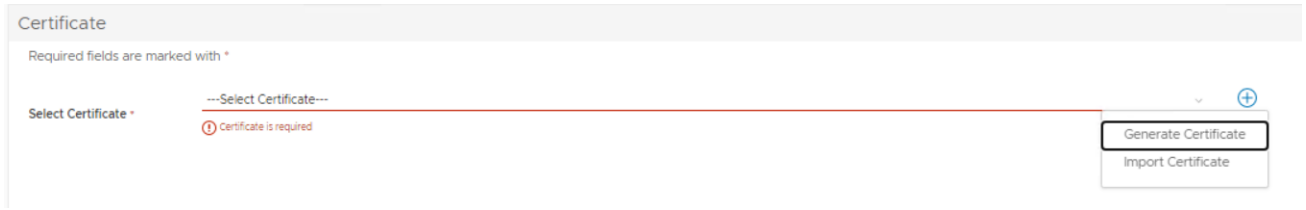

Fill in the requires certificate fields then click Generate:

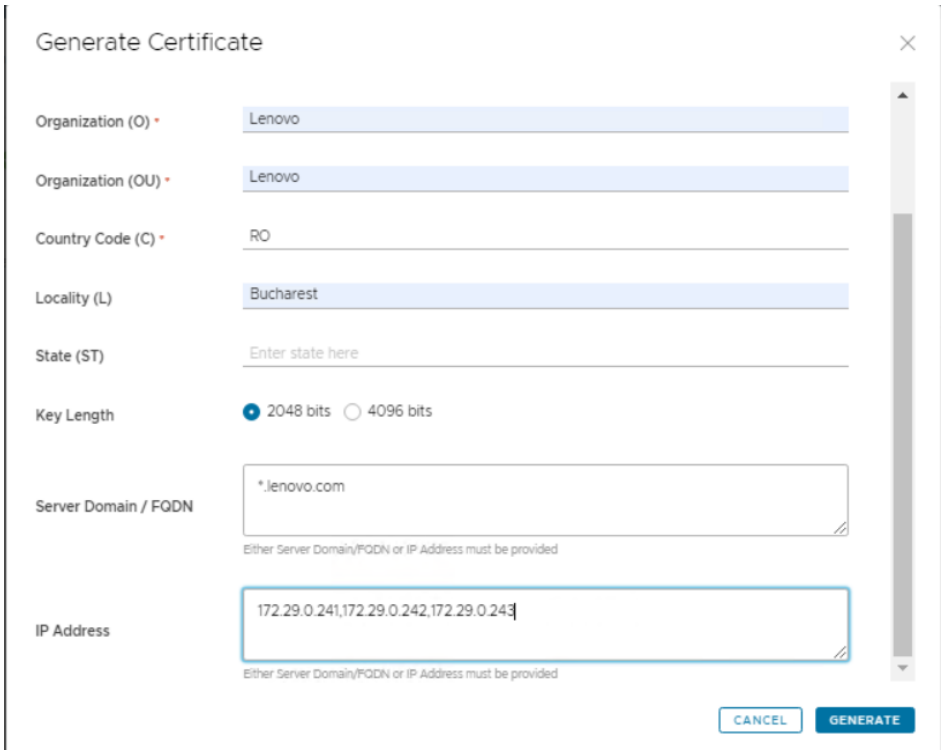

Select the newly generated Certificate then click Next:

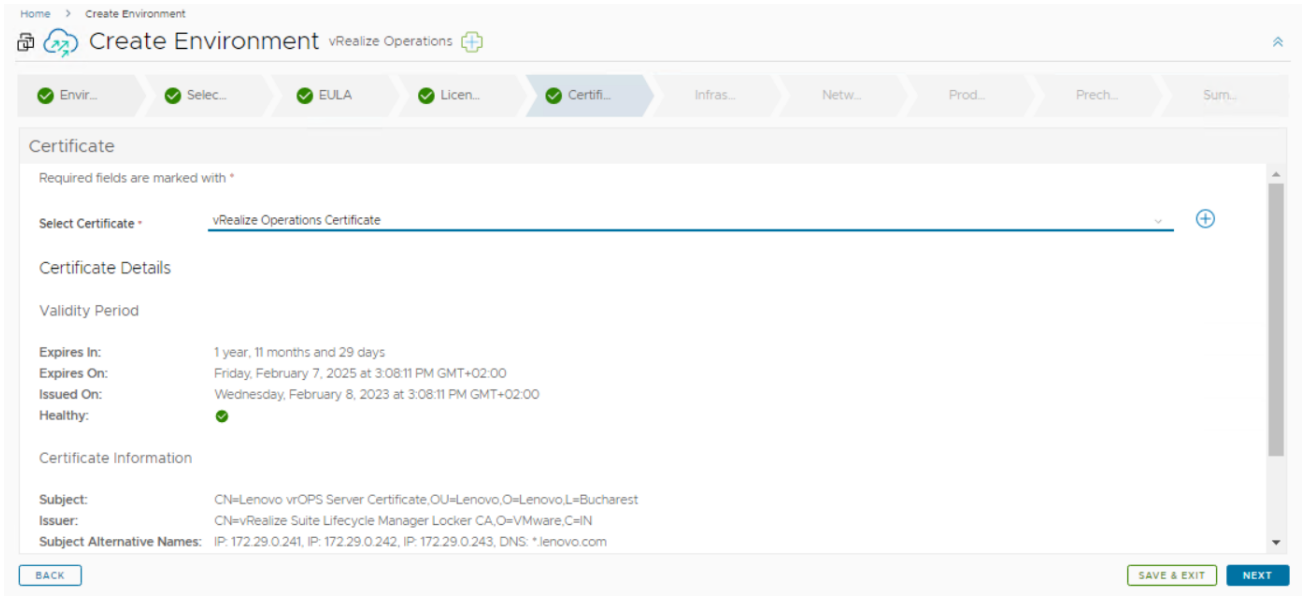

In the Infrastructure window select the Cluster, the Resource Pool and the disk Mode then click Next:

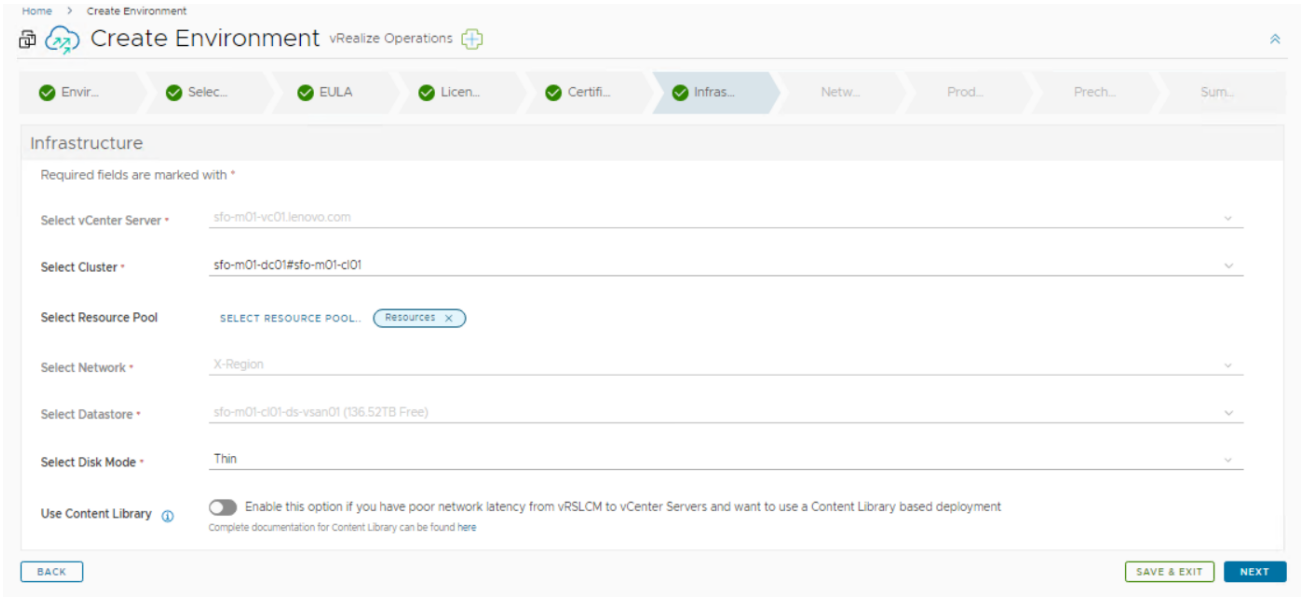

In the Network window verify that the X-Region is selected for the deployment and the corresponding DNS server is used then click Next:

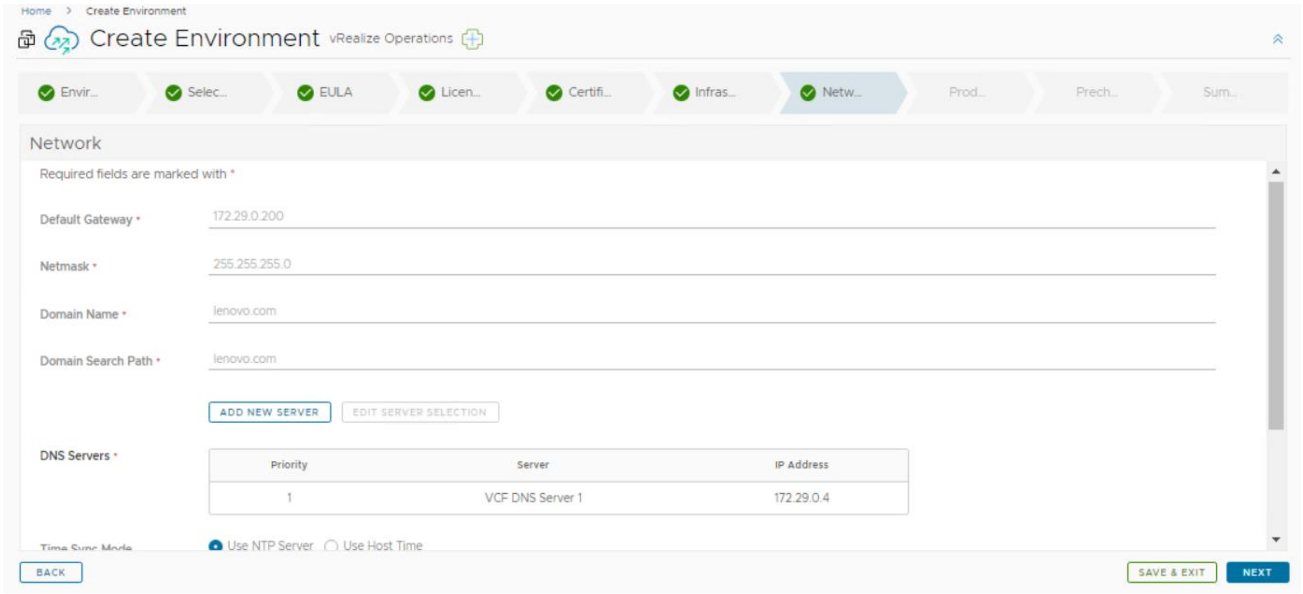

In the Products window Select the TLS version that needs to be disabled, the Certificate for vROPS, the NTP Server used:

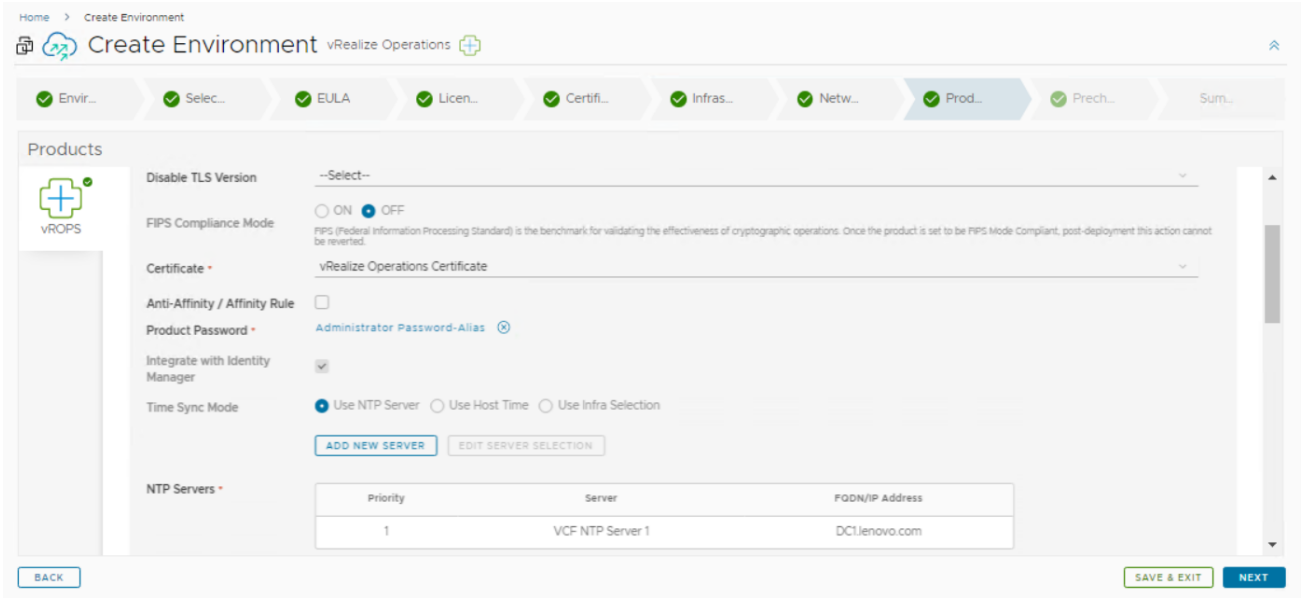

In the same window fill in the VM Name of master node (as it will appear in the vSphere), the FQDN of the master VM (should be already present in the DNS) and the IP address

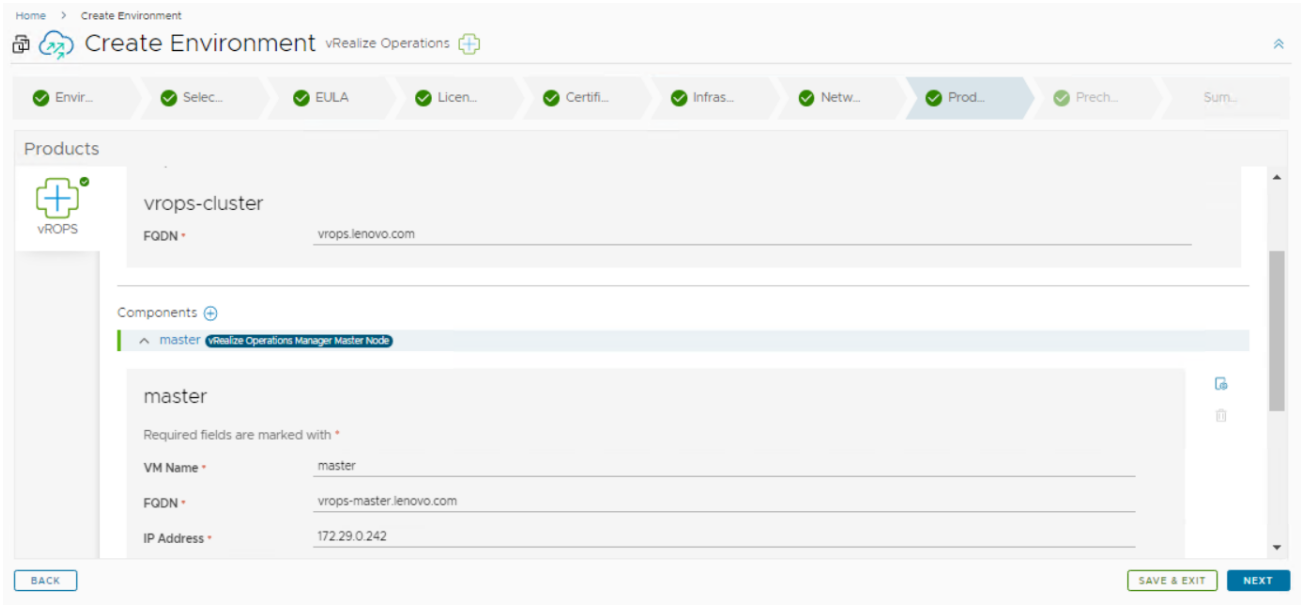

In the same window fill in the VM Name of the replica node, the FQDN (should already be present in the DNS) and the IP address, then click Next:

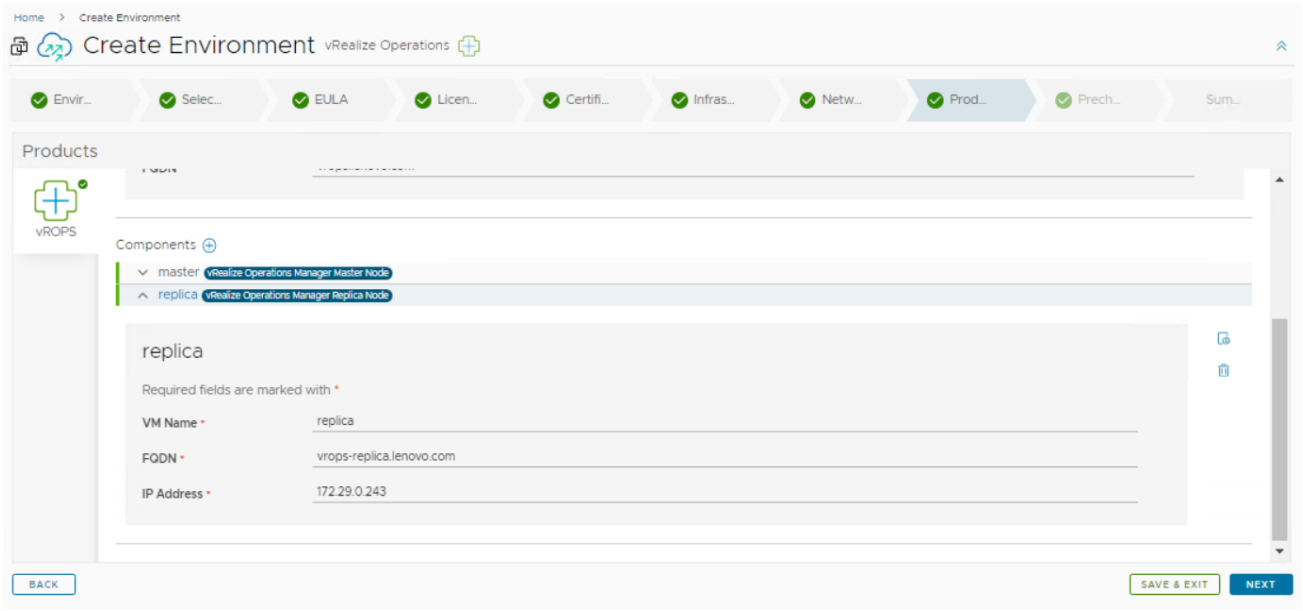

In the Precheck window click on Run Precheck button, then click Next after all validations are Passed:

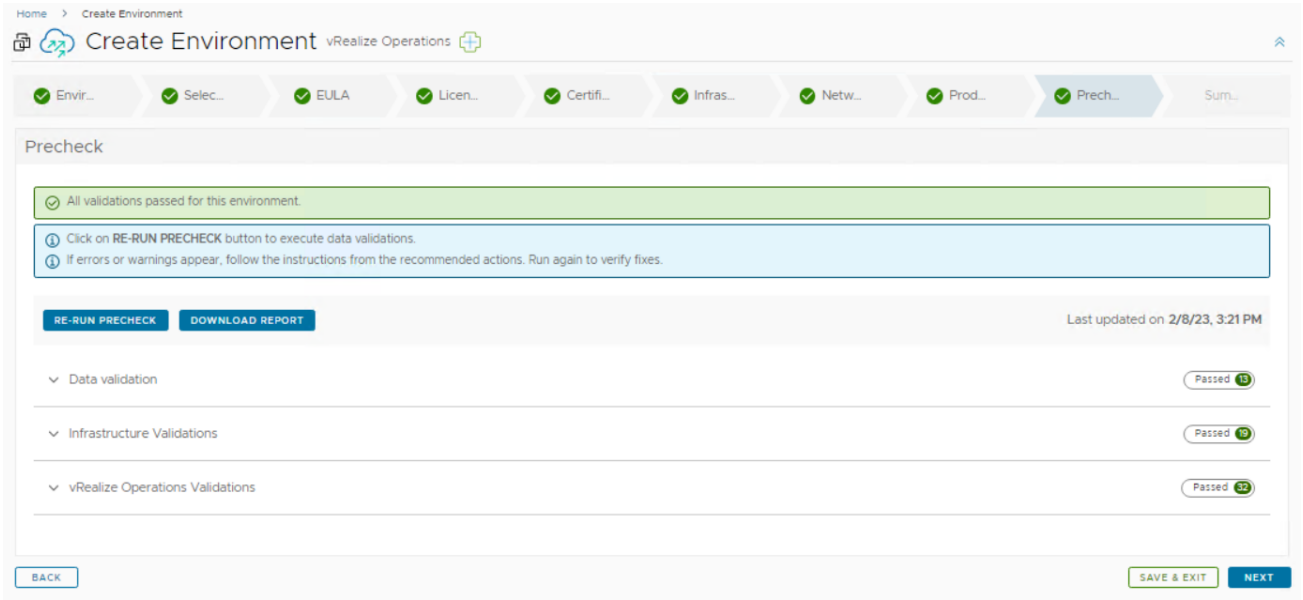

In the Summary windows verify that all parameters are correctly provided and click Submit:

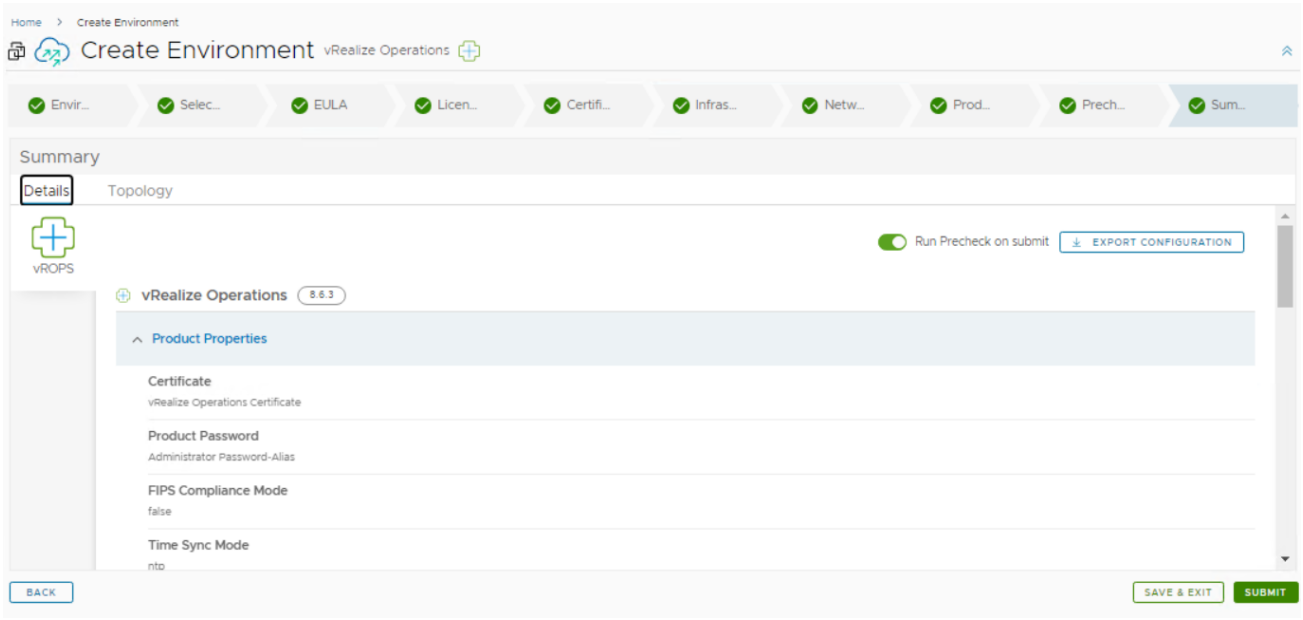

The Request Details window will be automatically switched to in order to monitor each deployment Stage:

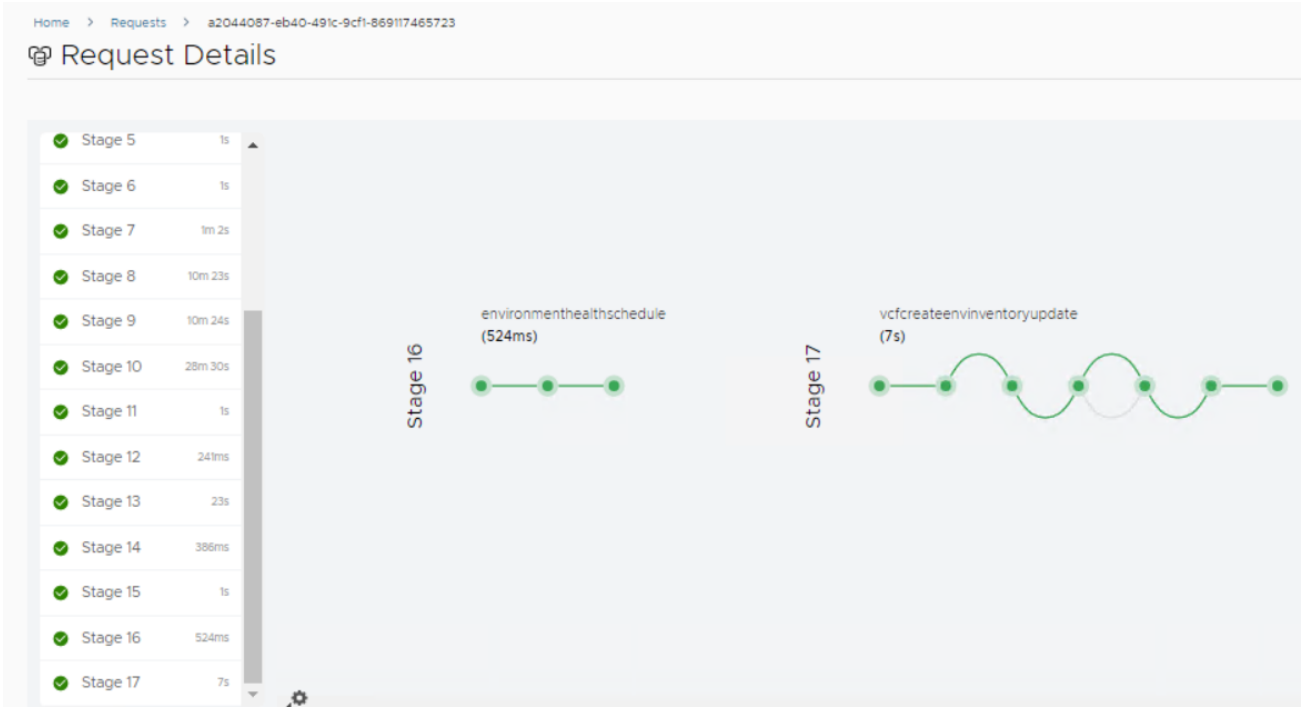

Deploy Aria Automation and Orchestrator:

Add the required product binaries:

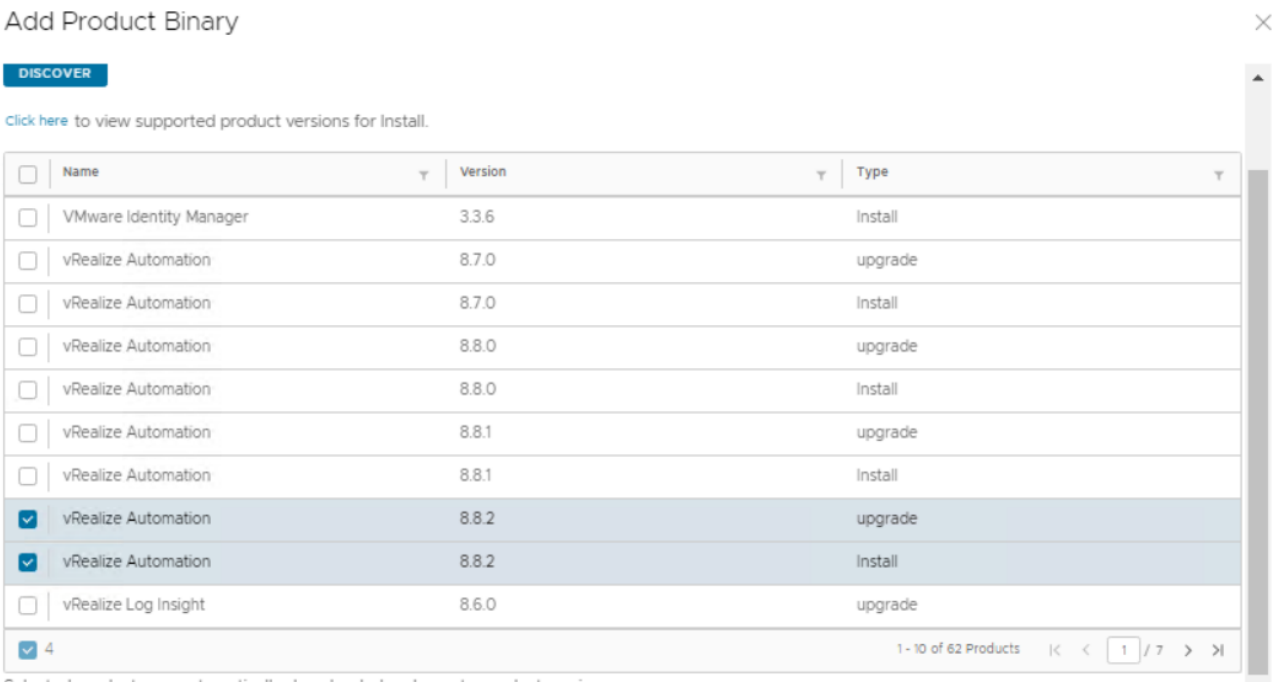

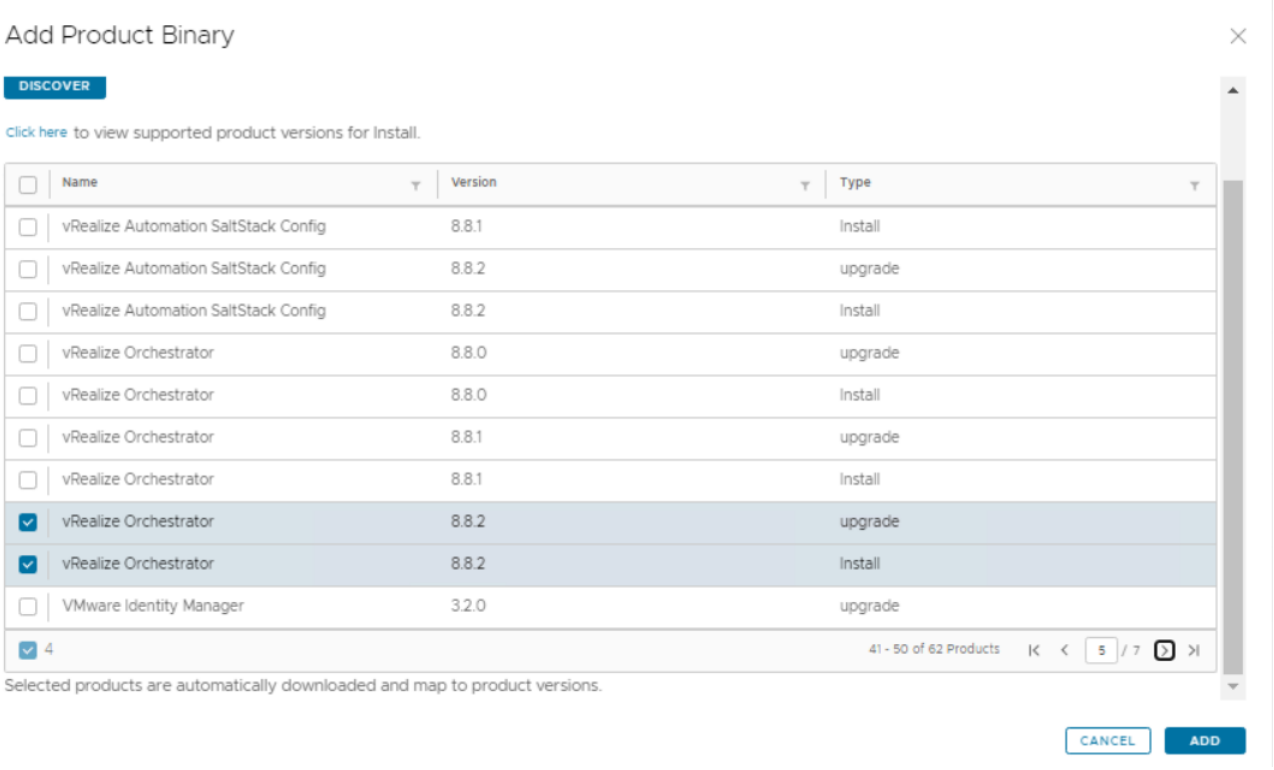

In the vRLCM | LO Environments window click the 3 dots next to Aria Operations and select +Add Product:

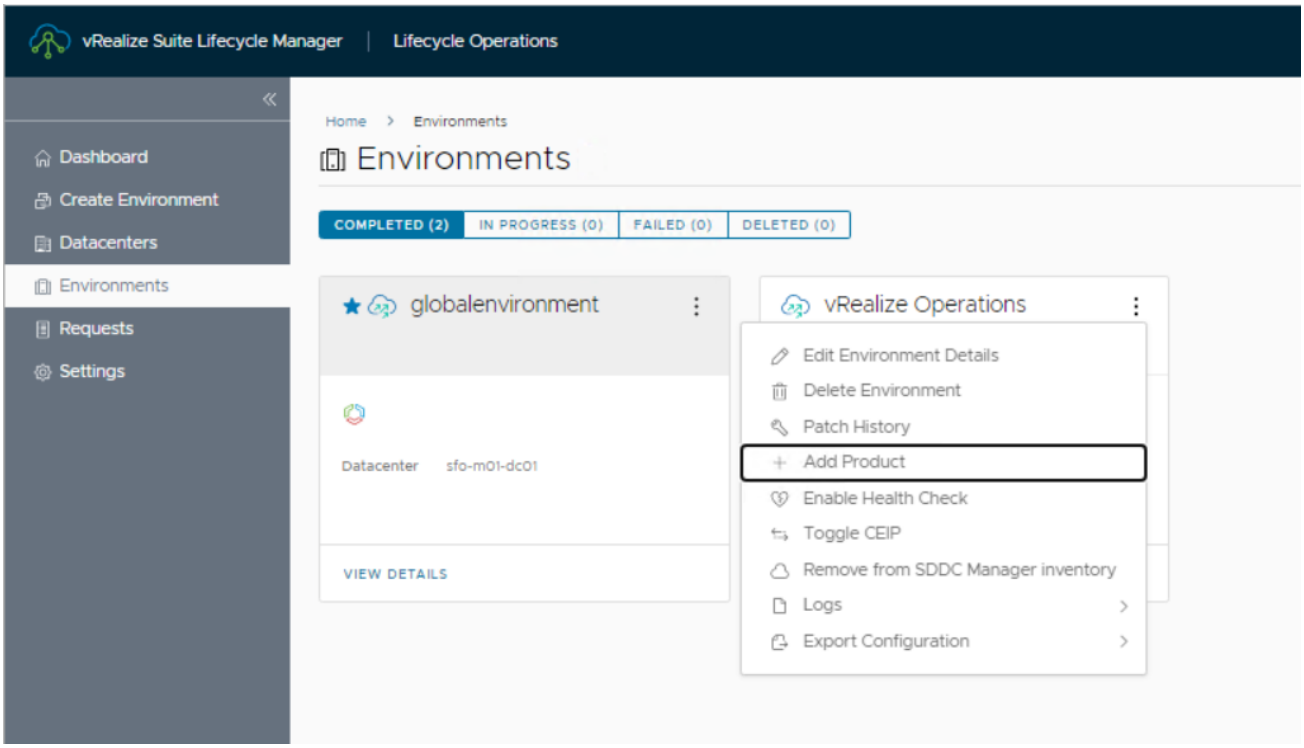

Select Aria Automation and Aria Orchestrator and click Next:

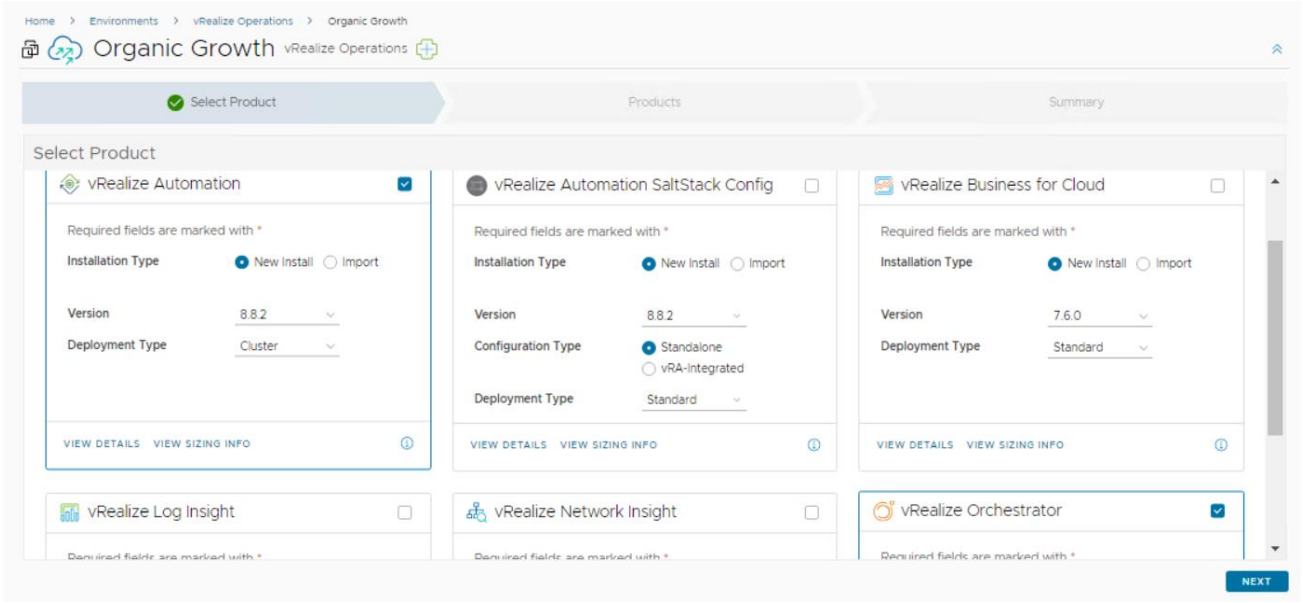

Select the ‘I agree the terms & conditions’ checkbox then click Next:

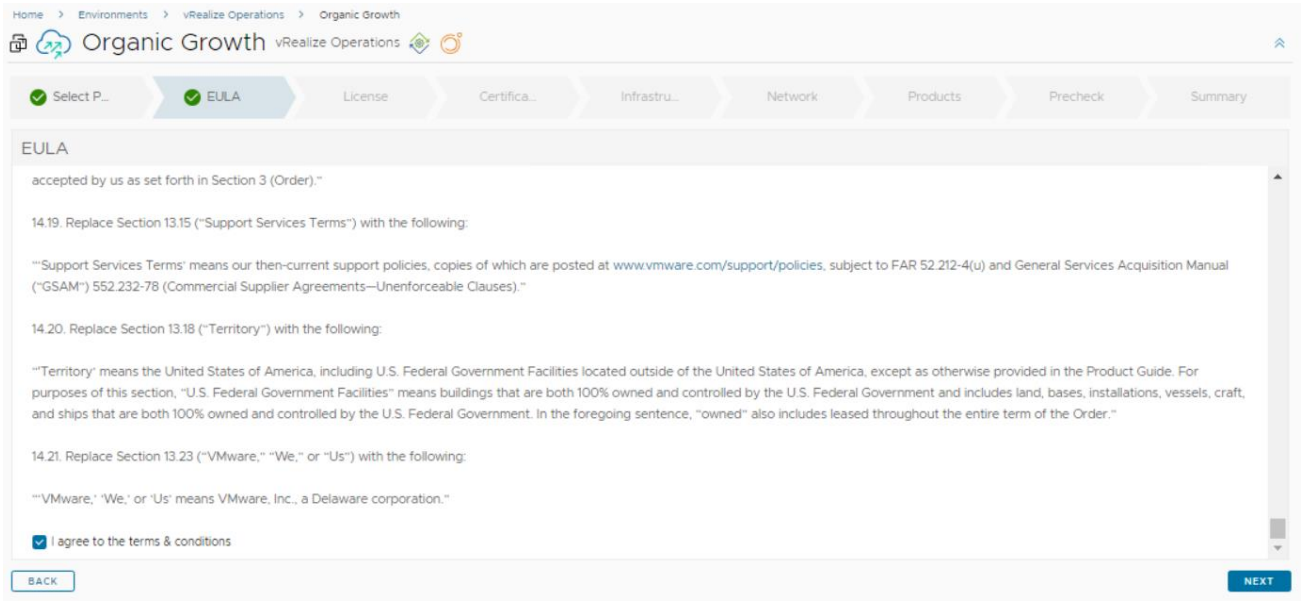

In the License windows click on Select:

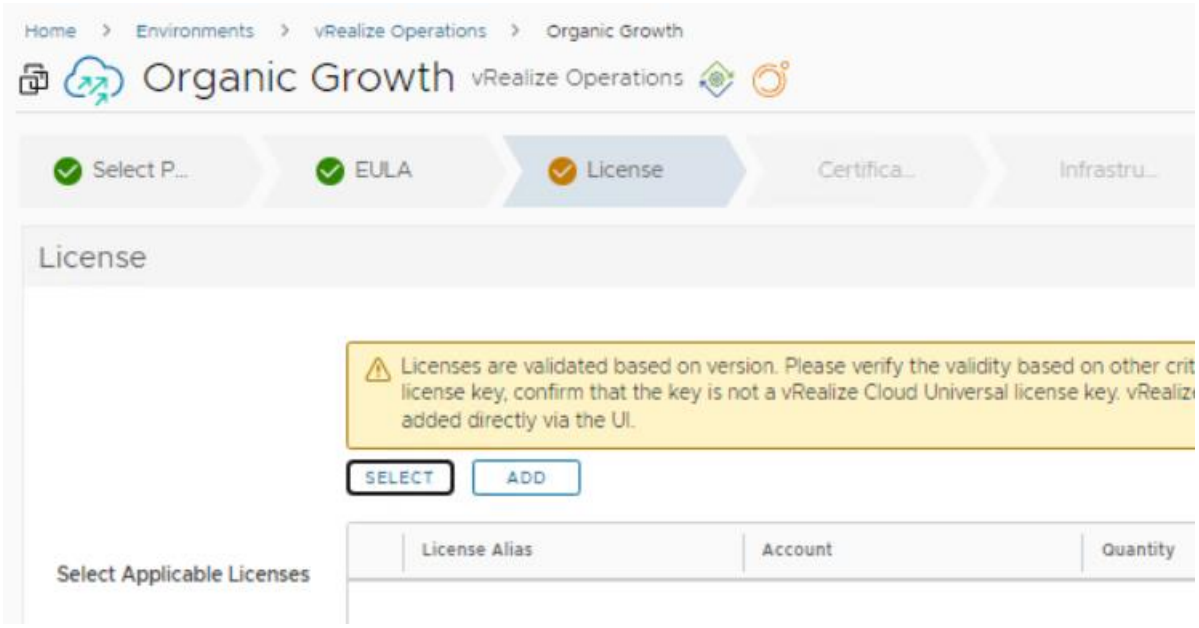

Select the Aria License and click Update:

Click on Validate Association button then click Next:

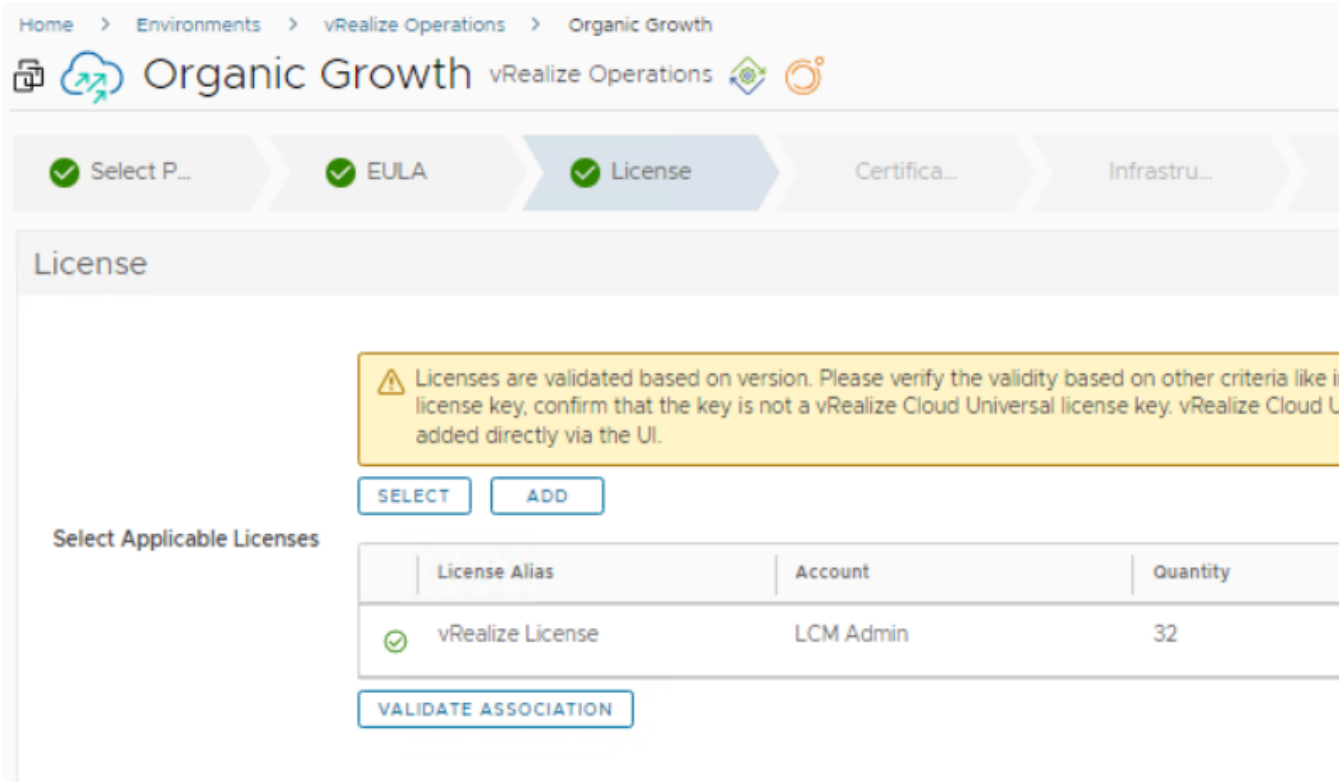

In the Certificate window click the + (plus) sign and select Generate Certificate:

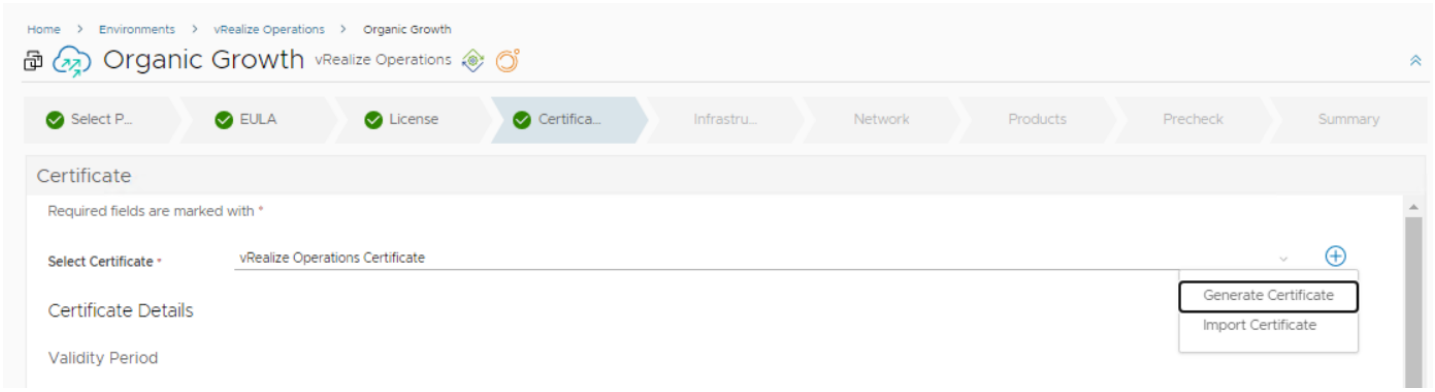

In the Generate Certificate fill in the required fields (IP addresses for vRA Primary node, Secondary01 Node, Secondary02 Node, vRO Primary Node):

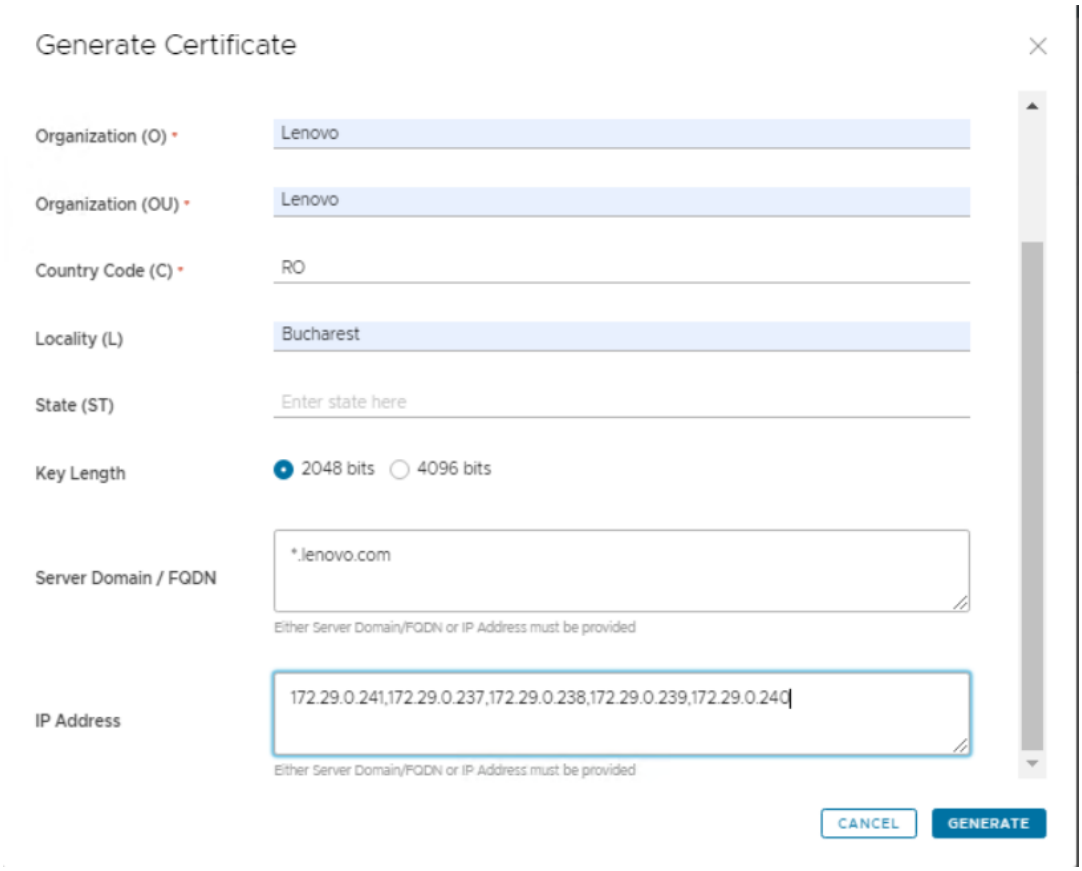

In the Product windows select the Node Size (Medium), the Certificate and Use Default settings for Internal Pods and Services:

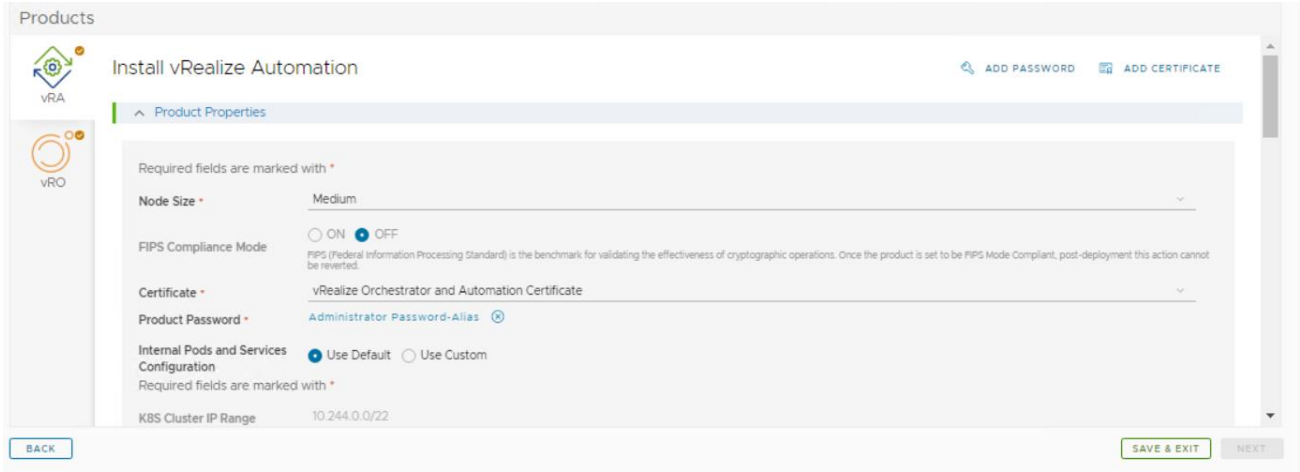

In the same window verify that the Administrator Password-Alias is selected and the NTP Server is correctly configured for X-Region:

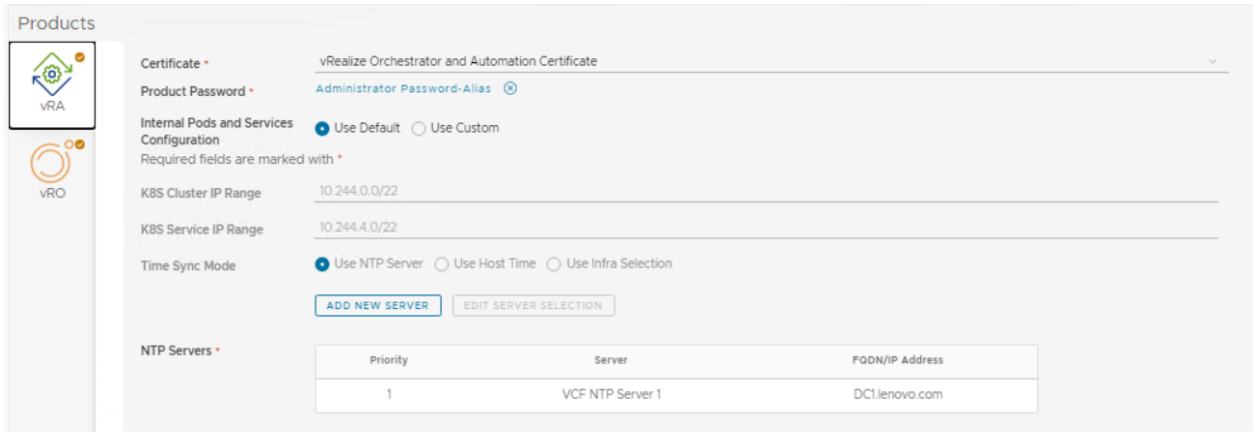

In the same window fill in the VM Name (as it will appear in vSphere), the FQDN and IP address of the primary vRA node:

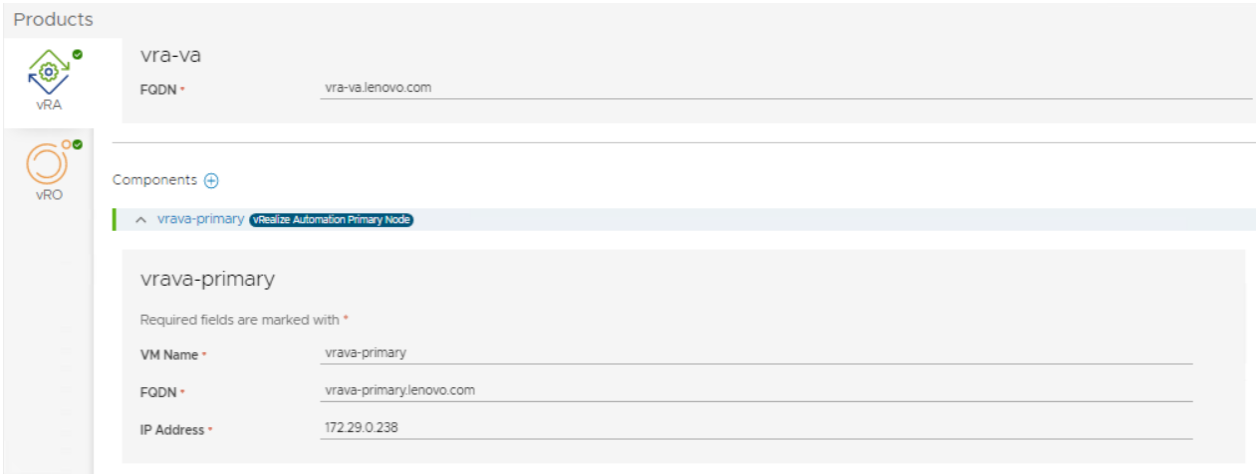

In the same window fill in the VM Name (as it will appear in vSphere), the FQDN and IP address of the secondary01 and secondary02 vRA nodes:

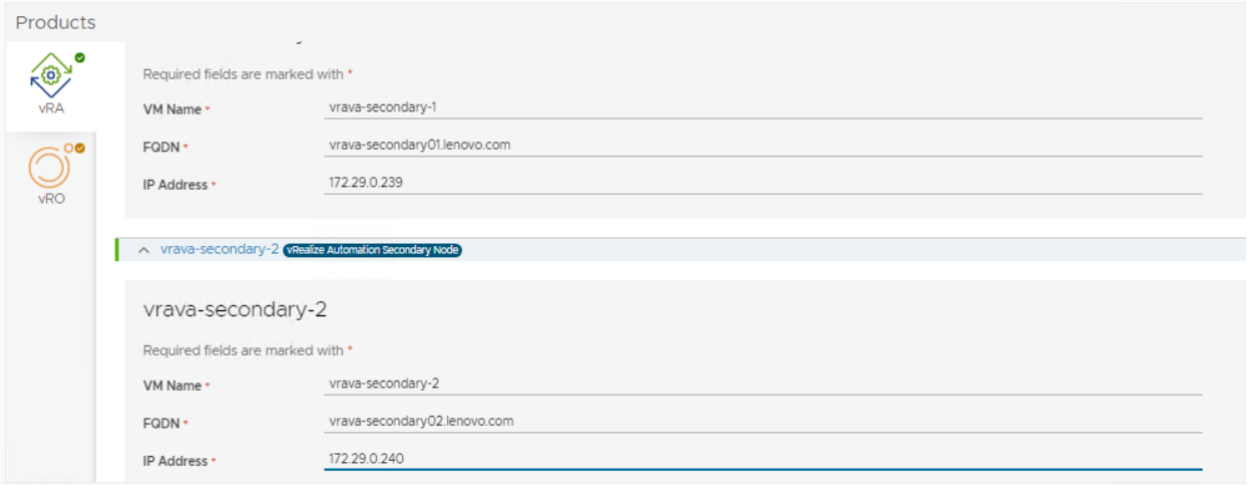

In the same windows select the Certificate for the vRO and vRA:

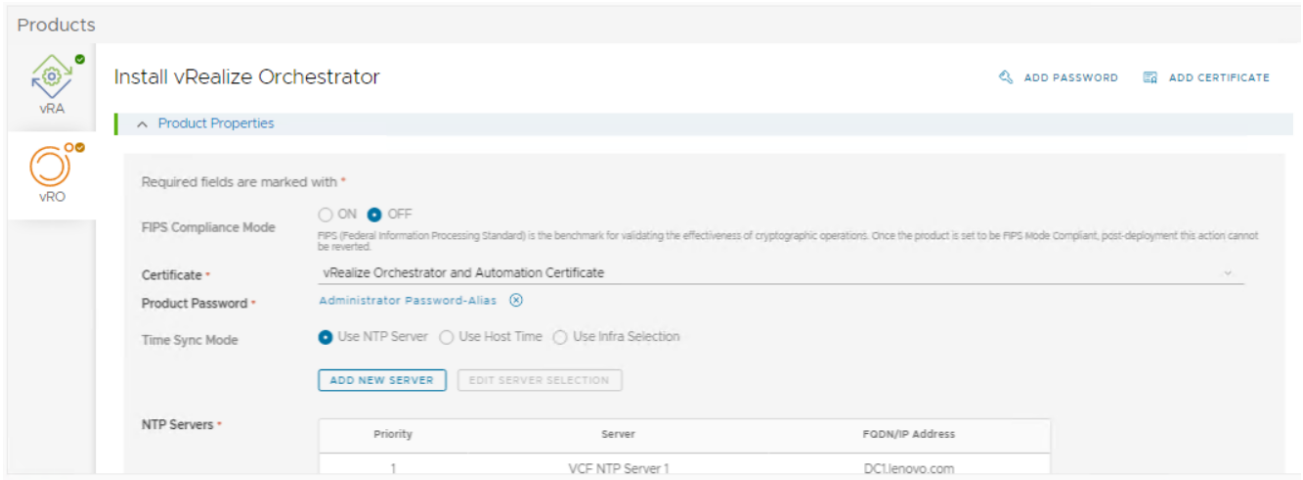

In the same window fill in the VM Name (as it will appear in vSphere), the FQDN and IP Address for the primary vRO node VM then click Next:

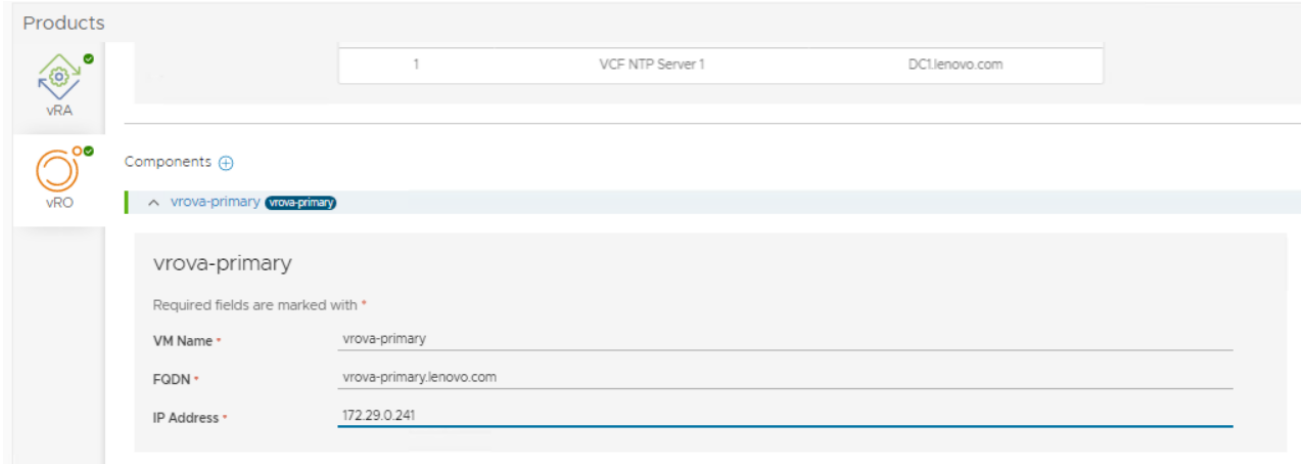

In the Precheck window click in Run Precheck button and click Next after all validations are Passed:

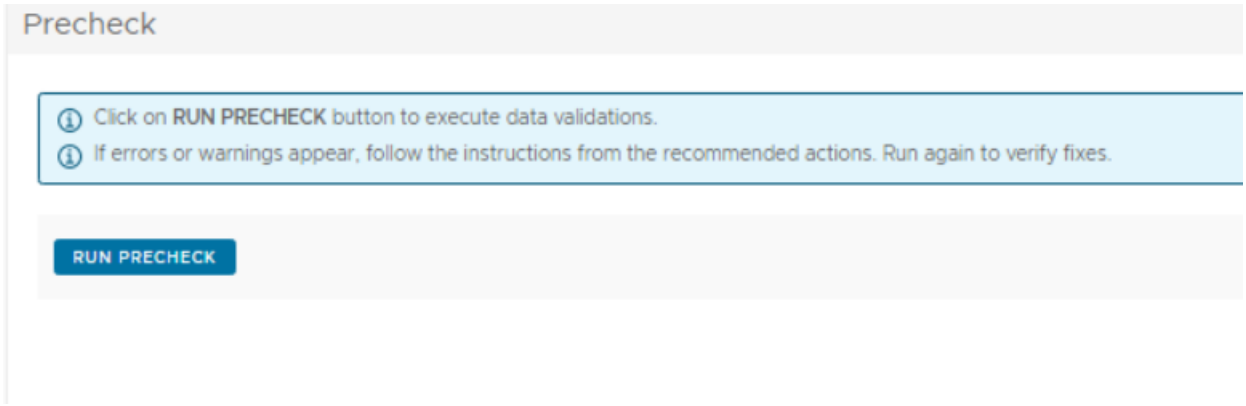

Click Submit after checking that all parameters are correctly provided:

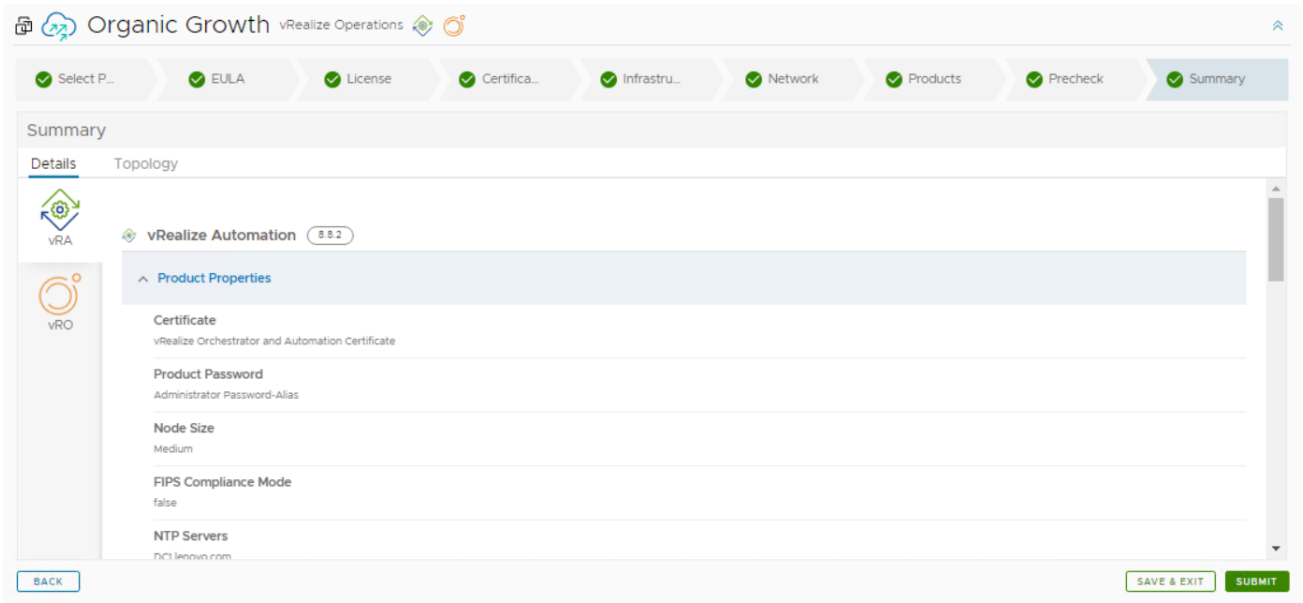

The Request Details window will be automatically switched to in order to monitor each deployment Stage:

Deploy Aria Operations for Logs:

Since the Aria Operations for Logs will be deployed in **Region-A** and the VLAN-backed NSX segment was selected when AVN network was deployed, for the **Region-A** (VAN10 in this case) VMs to communicate to **X-Region** (VLAN100 in this case) VMs, we must add a **Region-A** IP to the vRLCM VM. We might also need to add a **X-Region** IP to vRLI VM if we want to access it from X-Region subnet. This can be done in vSphere by adding a Network Card to the VMs and choosing the appropriate VLANs. We then connect to the VMs console and add the IP addresses for the newly added interfaces.

After logging in to the appliances console using vSphere issue the following commands to add IPs to the newly added interfaces:

[root@vrlcm ~]# ifconfig <ethX> 10.10.0.15 netmask 255.255.255.0

Add the vRLI product binaries form vRLCM | LO console:

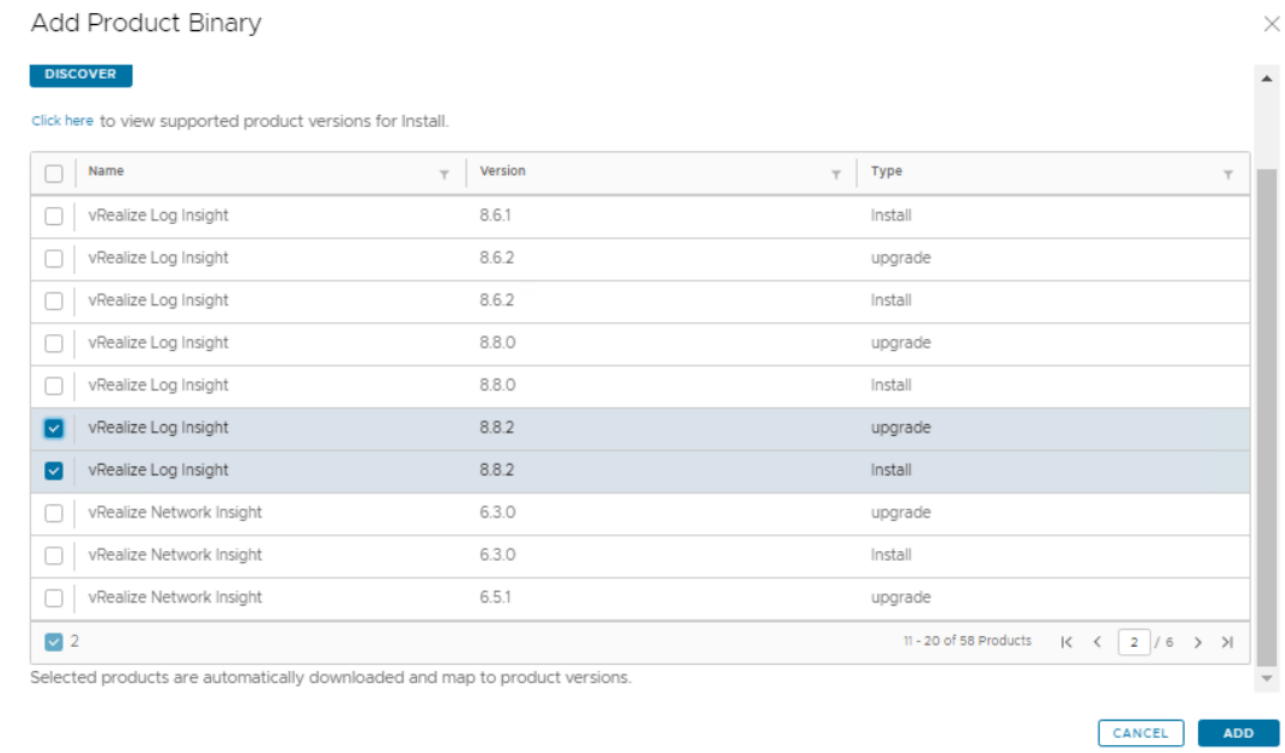

In the vRLCM | LO Environments window click the 3 dots net to Aria Operations:

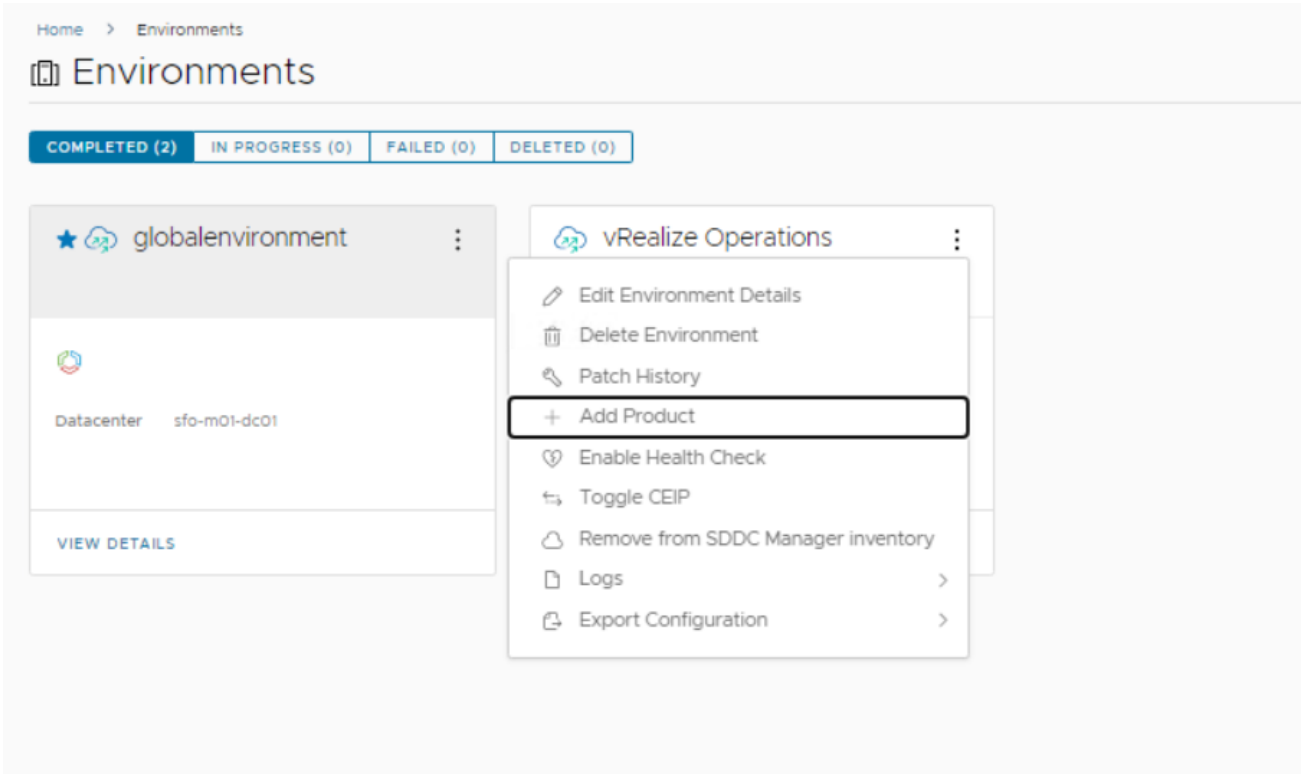

In the Select Product window click the Aria Operations for Logs checkbox, select New Install, Version and Deployment Type (Cluster), then click Next:

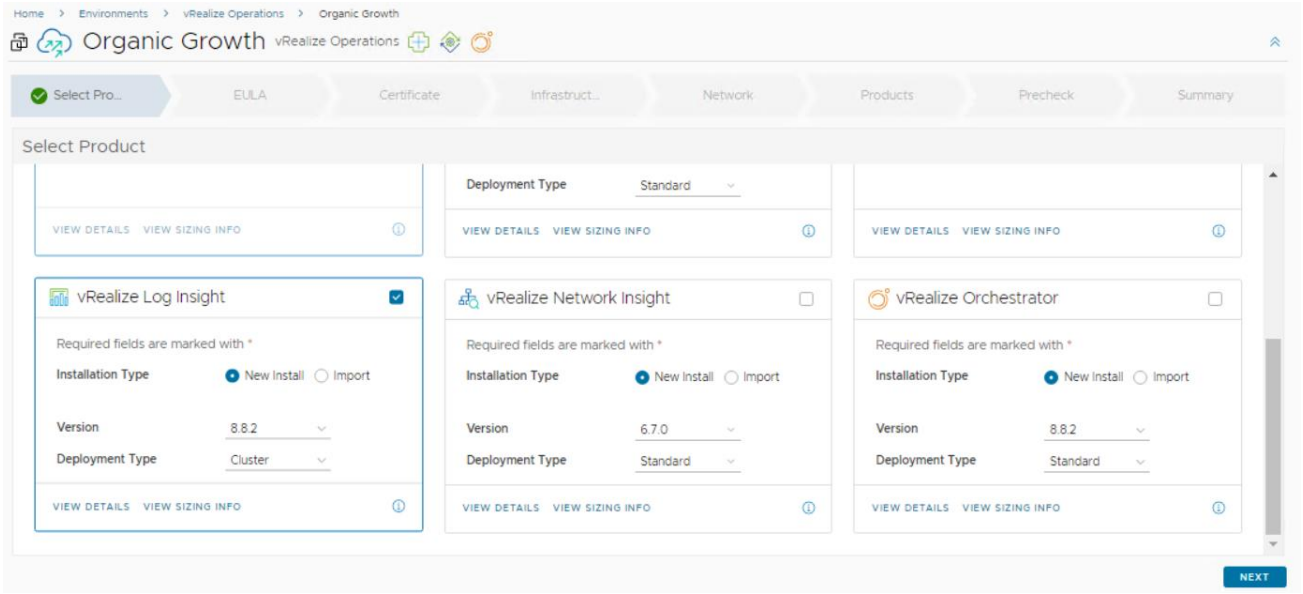

In the EULA windows select the ‘I agree to the terms & conditions’ then click Next:

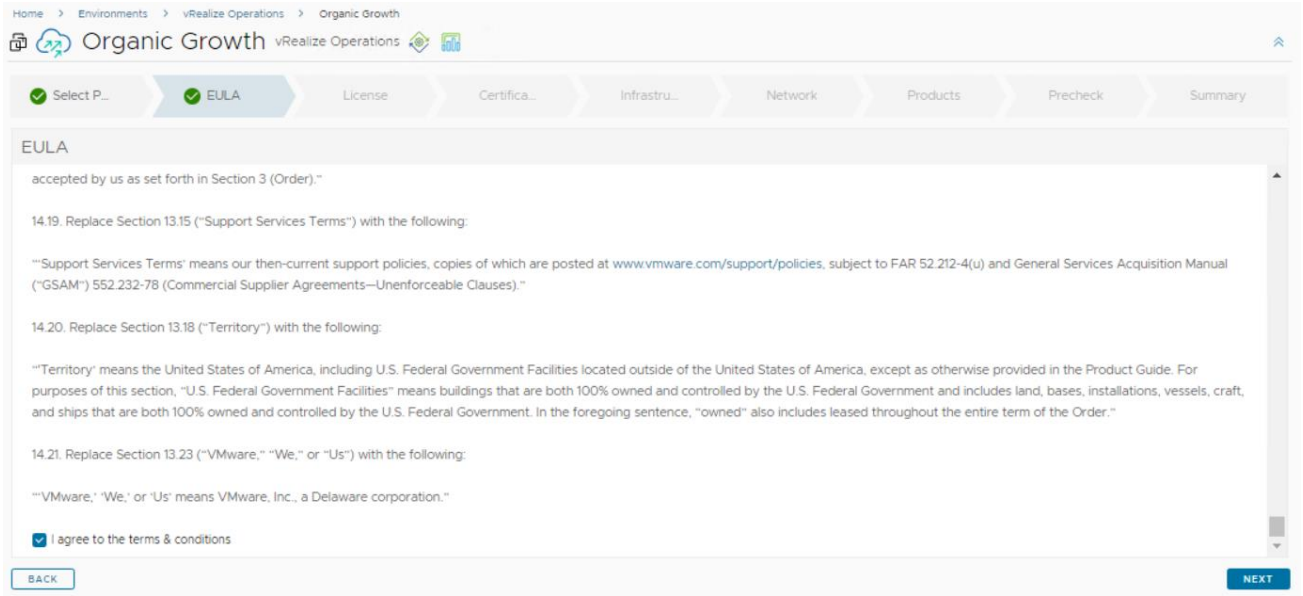

In the License window click on Select:

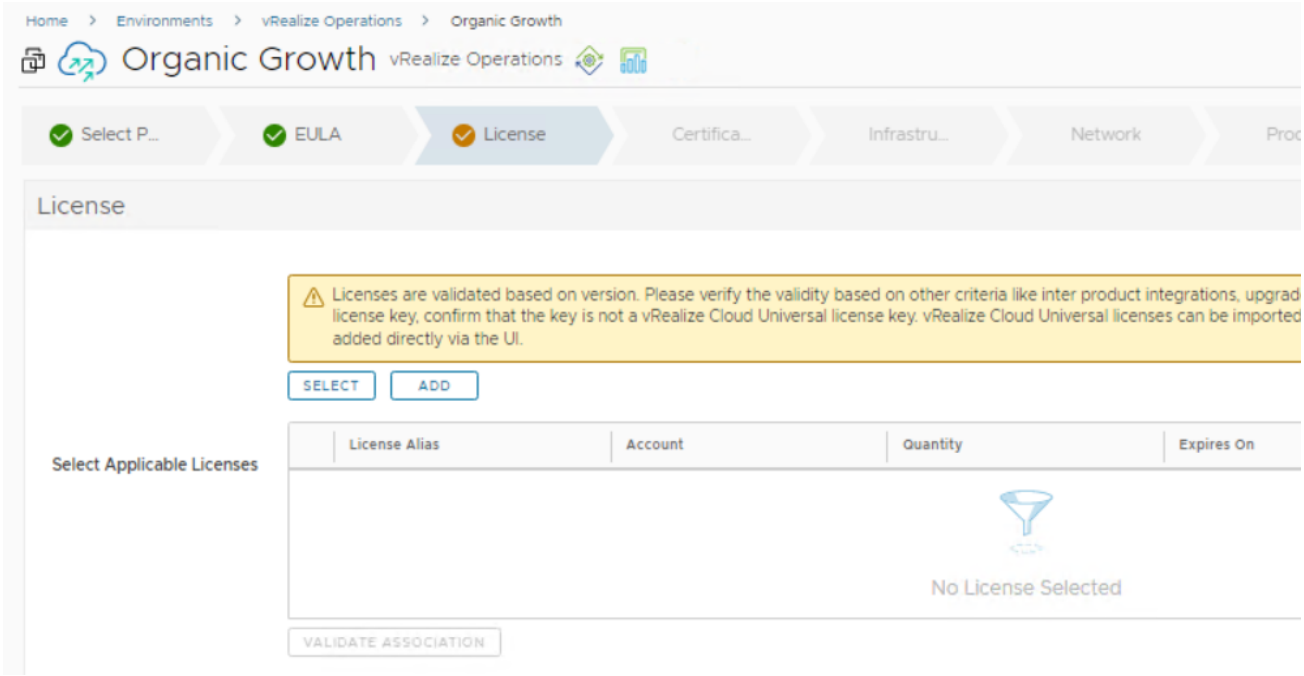

In the Select Applicable Licenses window select the Aria License and click Update:

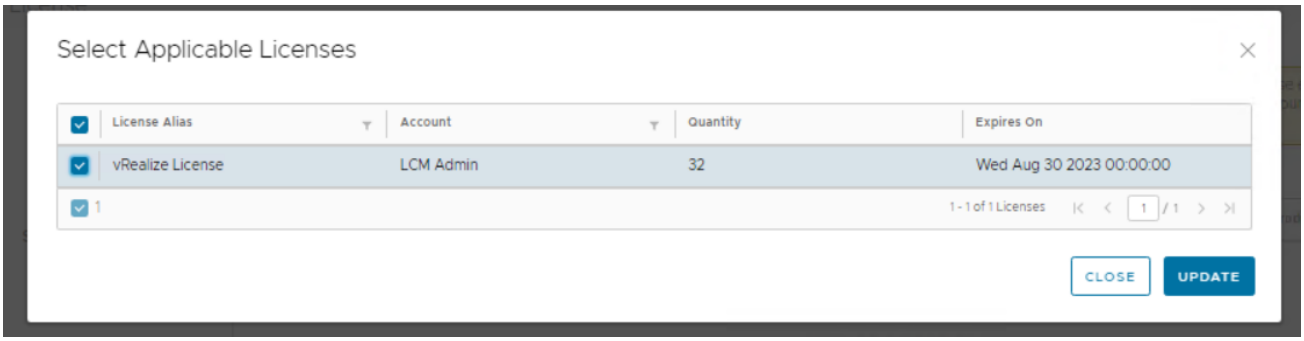

In the License window click Validate Association then click Next after the validation is successful:

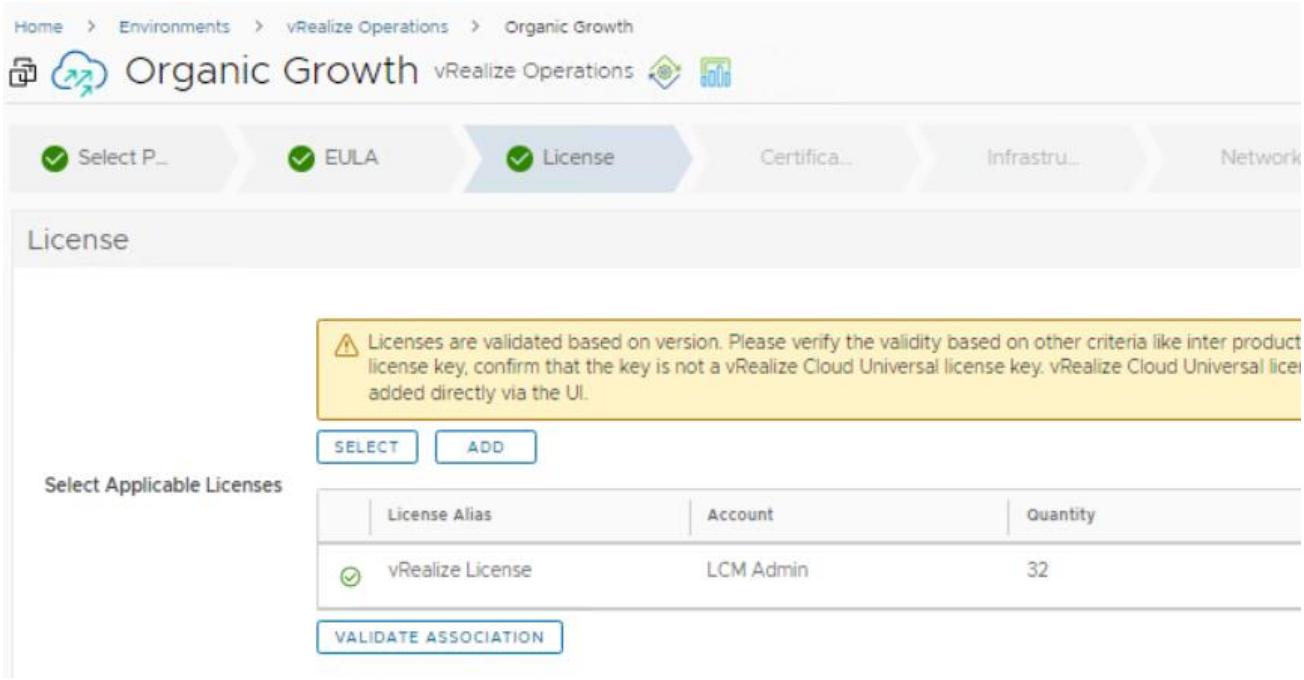

In the Certificate windows click on the + (plus) sign and select Generate Certificate:

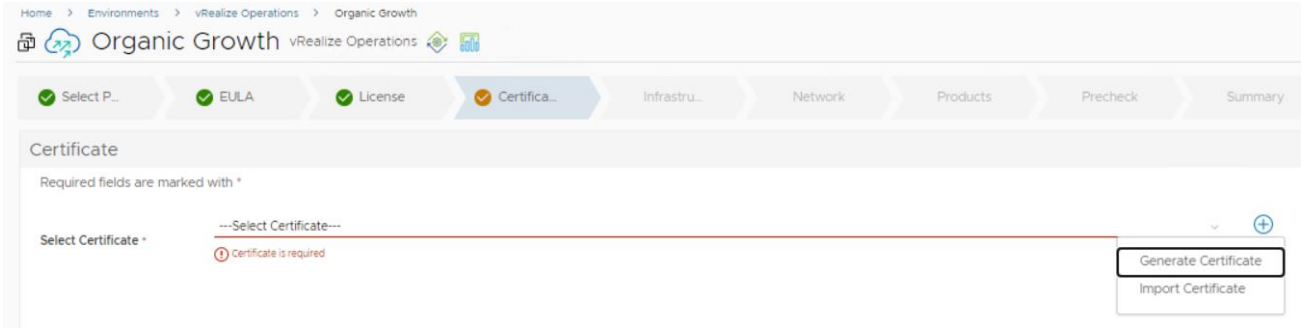

In the Generate Certificate window fill in the required fields (IP Addresses for vrli cluster, master node and the worker nodes) then click Generate:

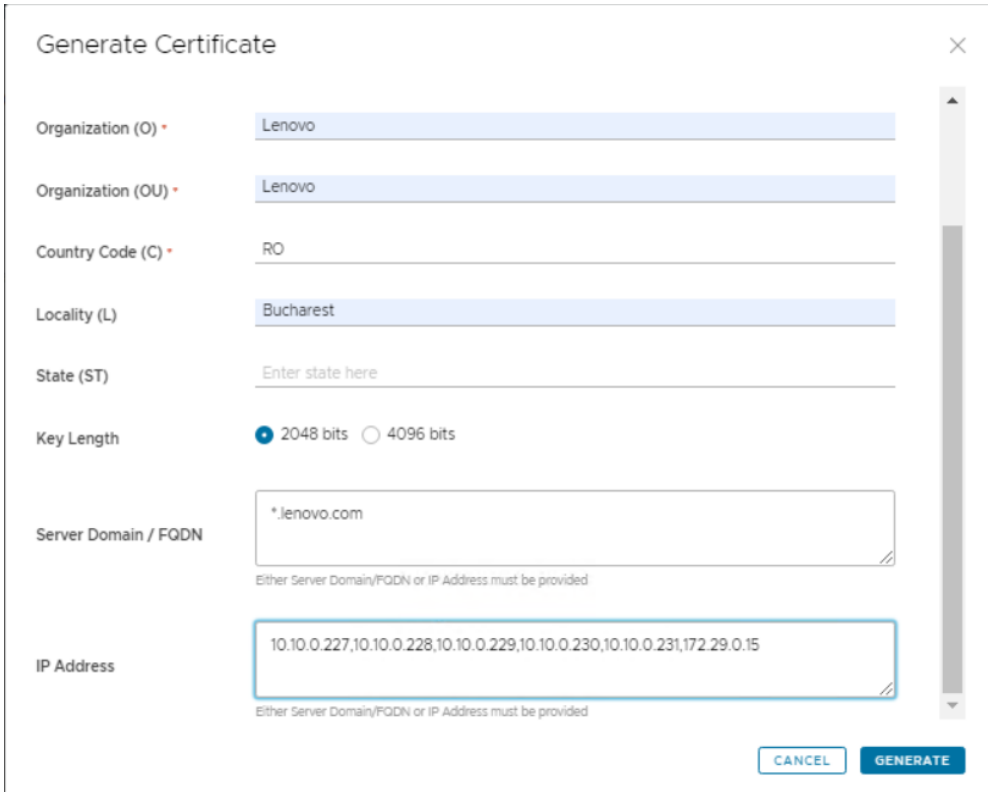

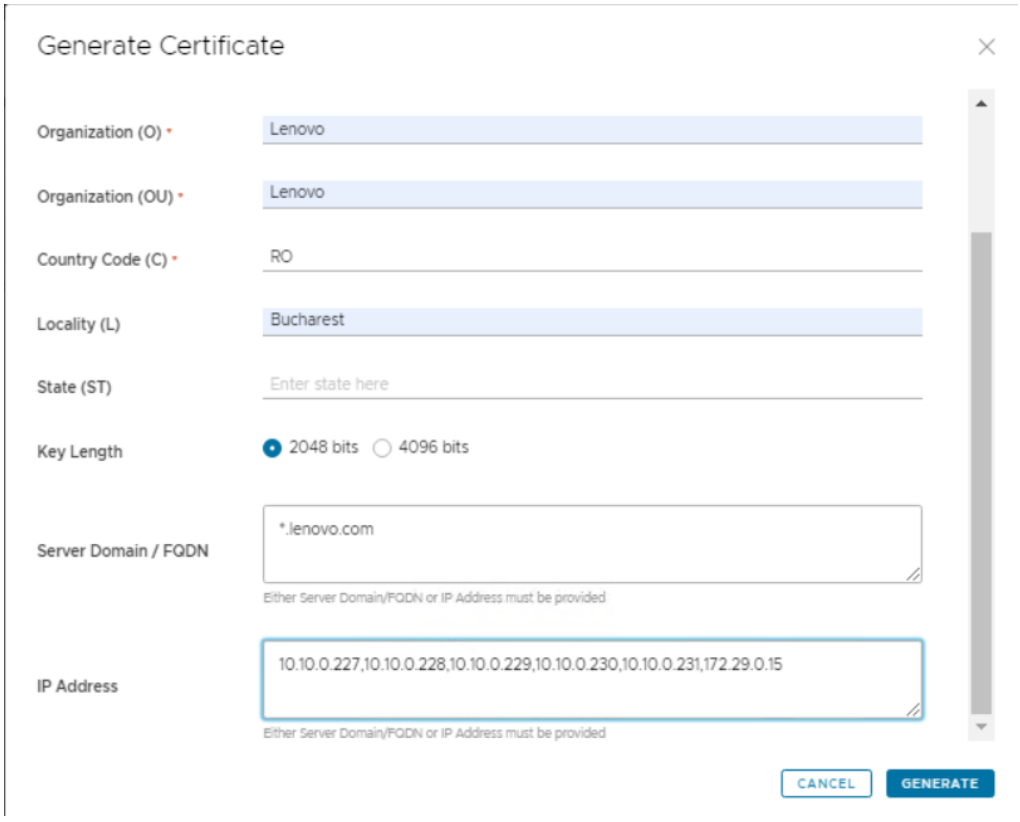

In the Infrastructure window select the vSphere cluster, the Resource Pool and the Disk Mode then click Next:

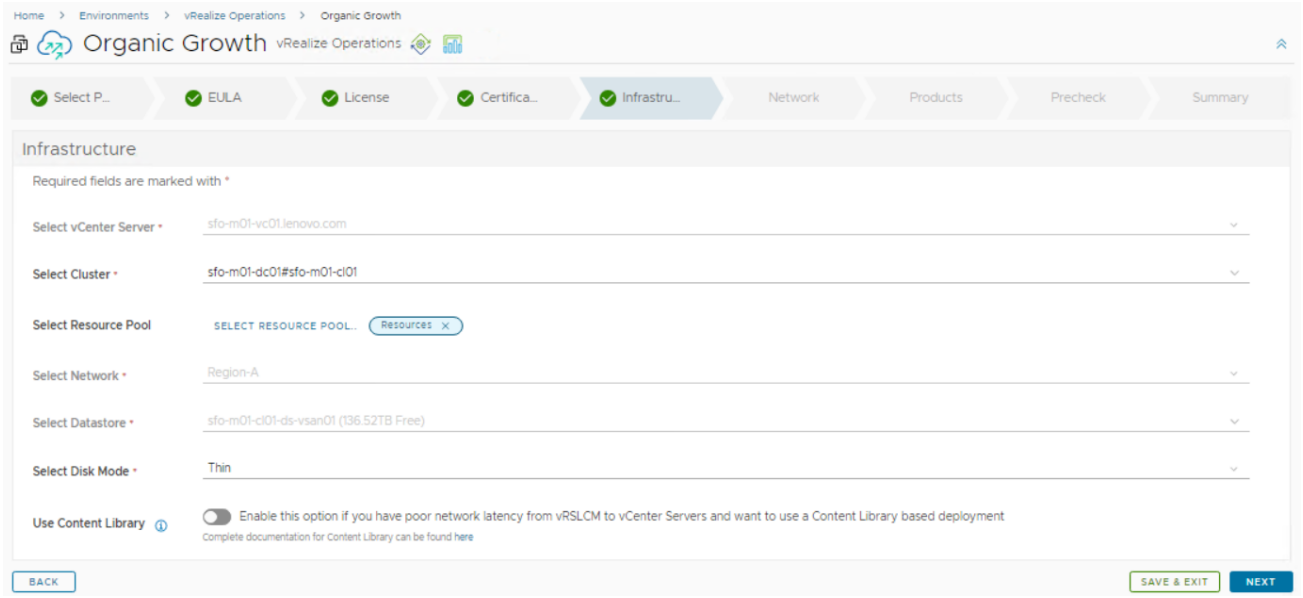

In the Network verify that Region-A network has the correct parameters and add a Region-A DNS Server by clicking on Add New Server:

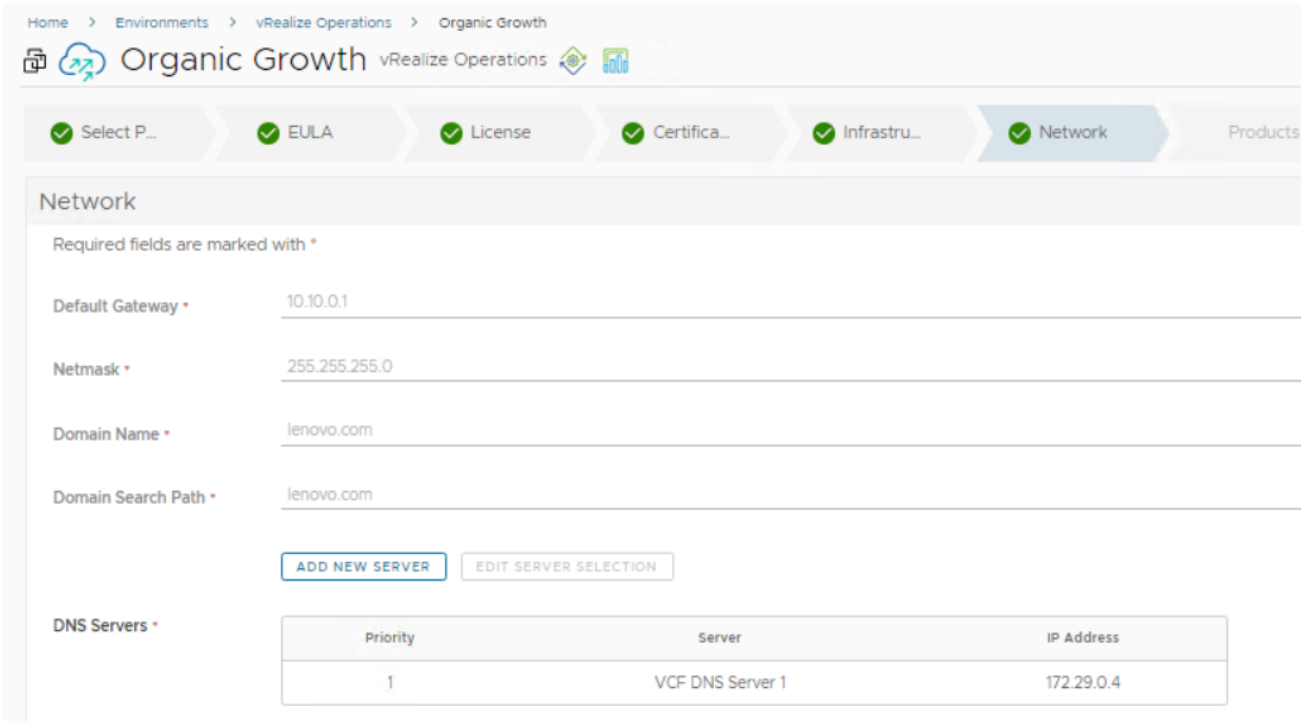

In the Add Server windows fill in a Name for the Server and the IP address of the DNS Server then click Submit. If vRLCM is not able to reach the DNS am error will be thrown. Make sure the vRLCM has an interface in the A-Region (VLAN10 in this case).

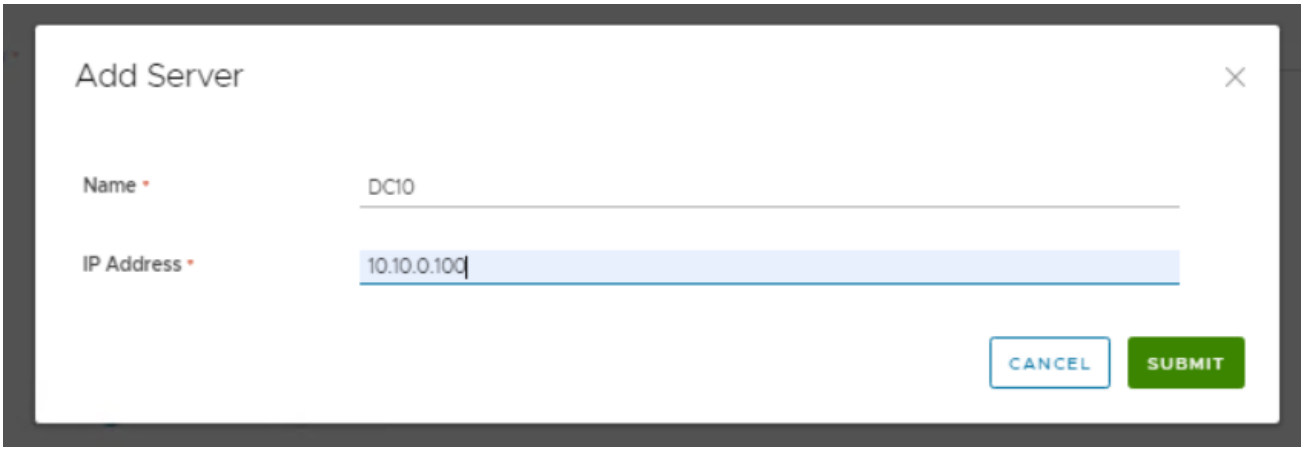

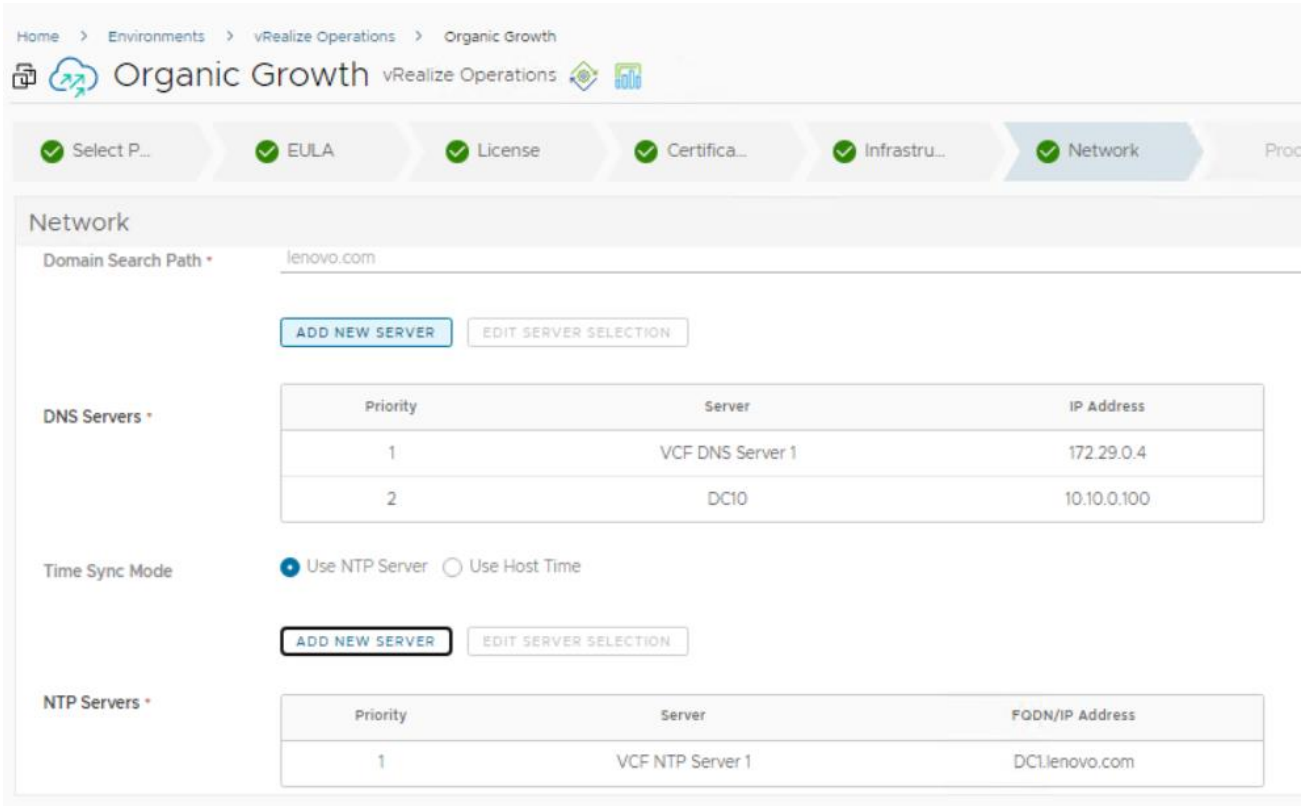

Note that a NTP Server must be added in the same way.

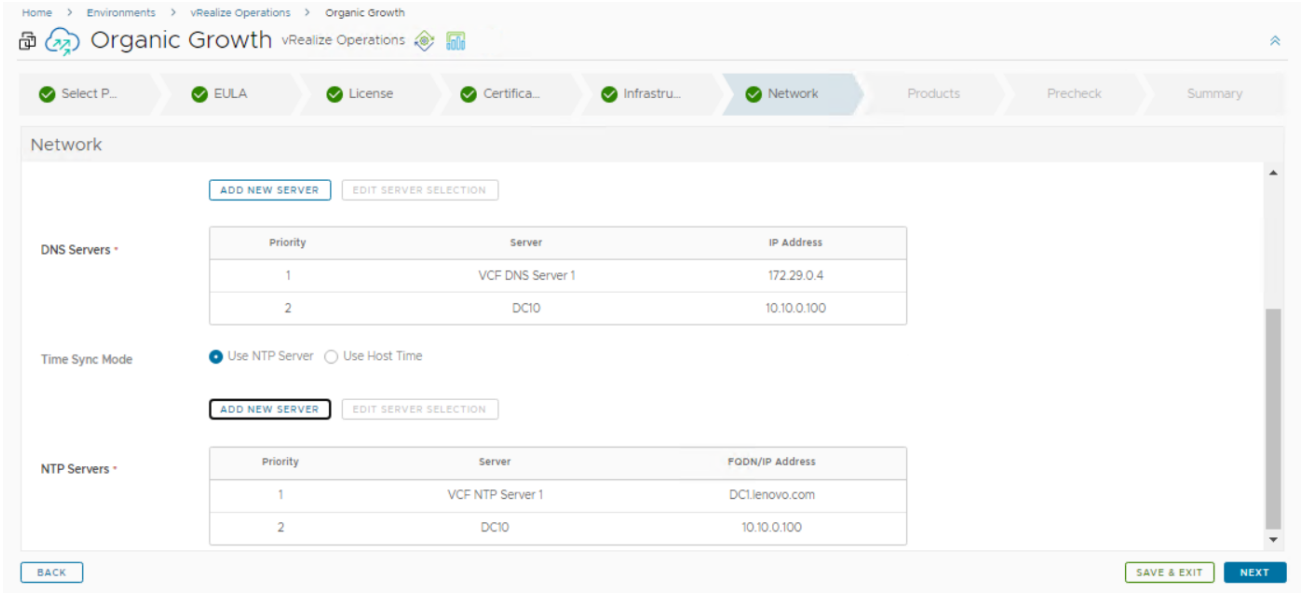

In the Products window select the Node Size (medium):

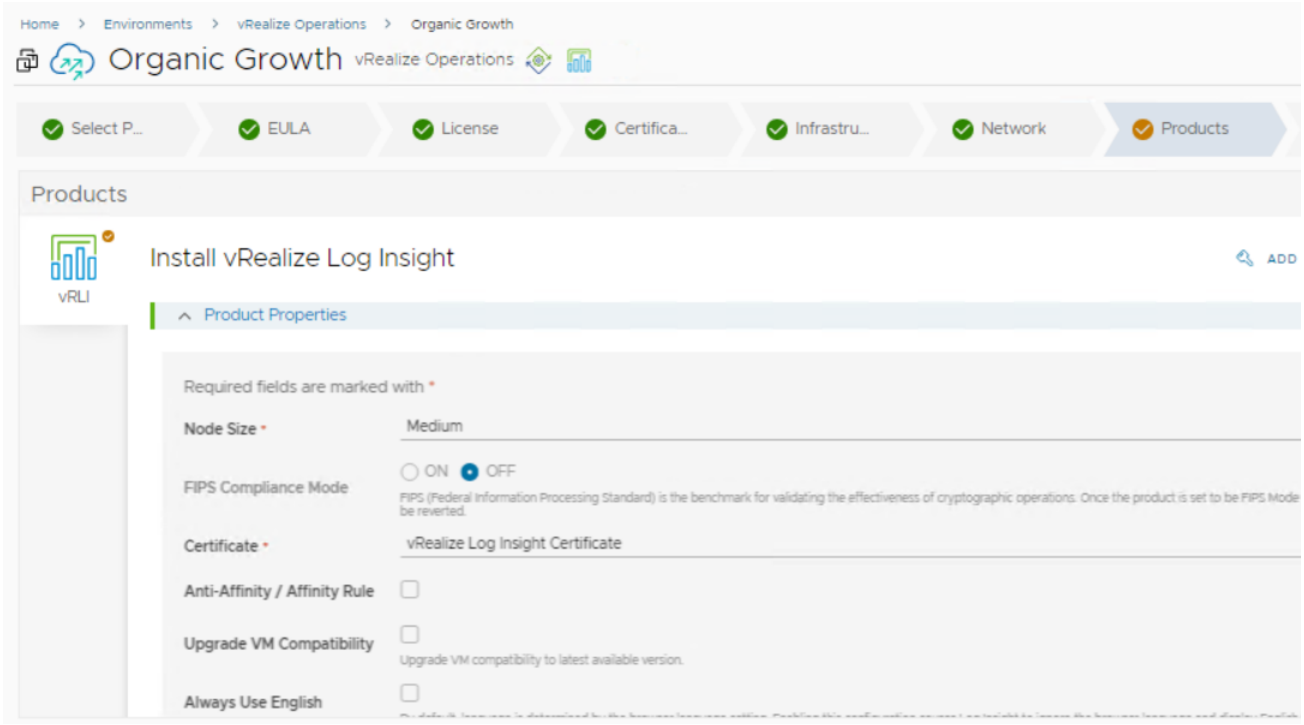

In the same window select an admin e-mail address, check Integrate with Identity Manager box:

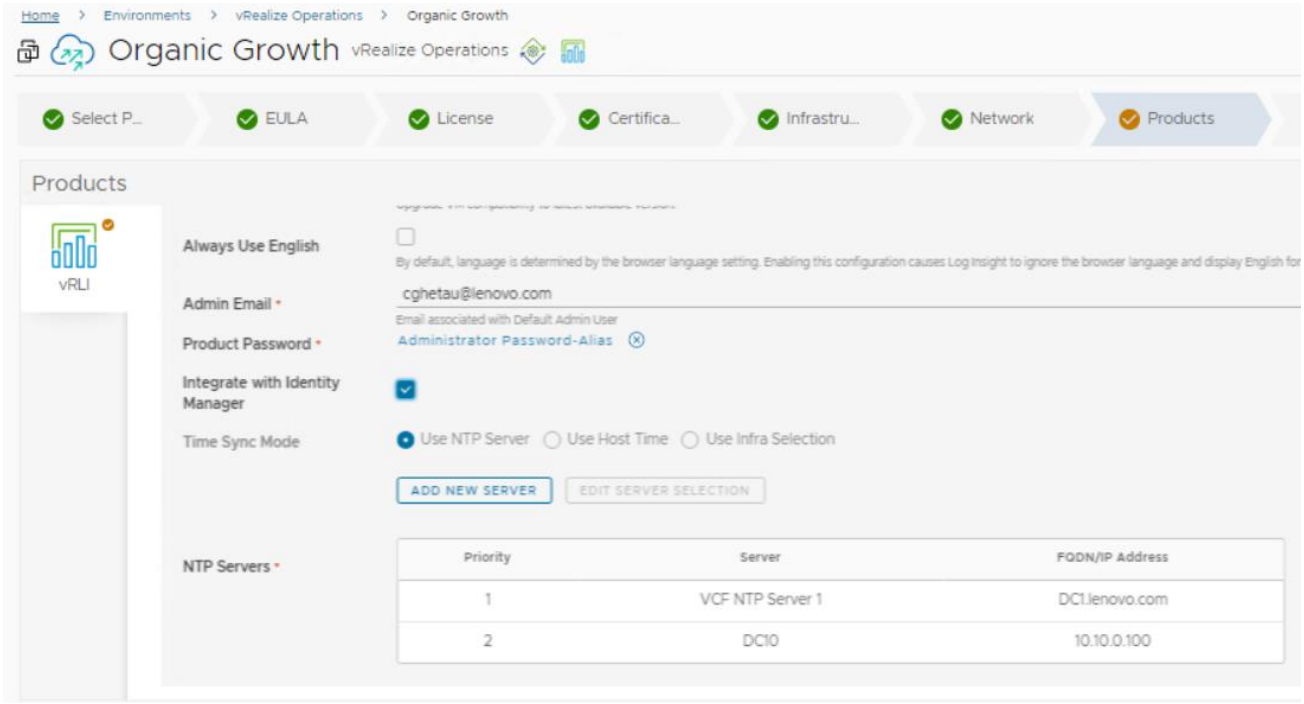

In the same window fill in the FQDN of the vRLI cluster, the VIP of the cluster, the VM Name (as it will appear in the vSphere), the FQDN and IP address of the master node:

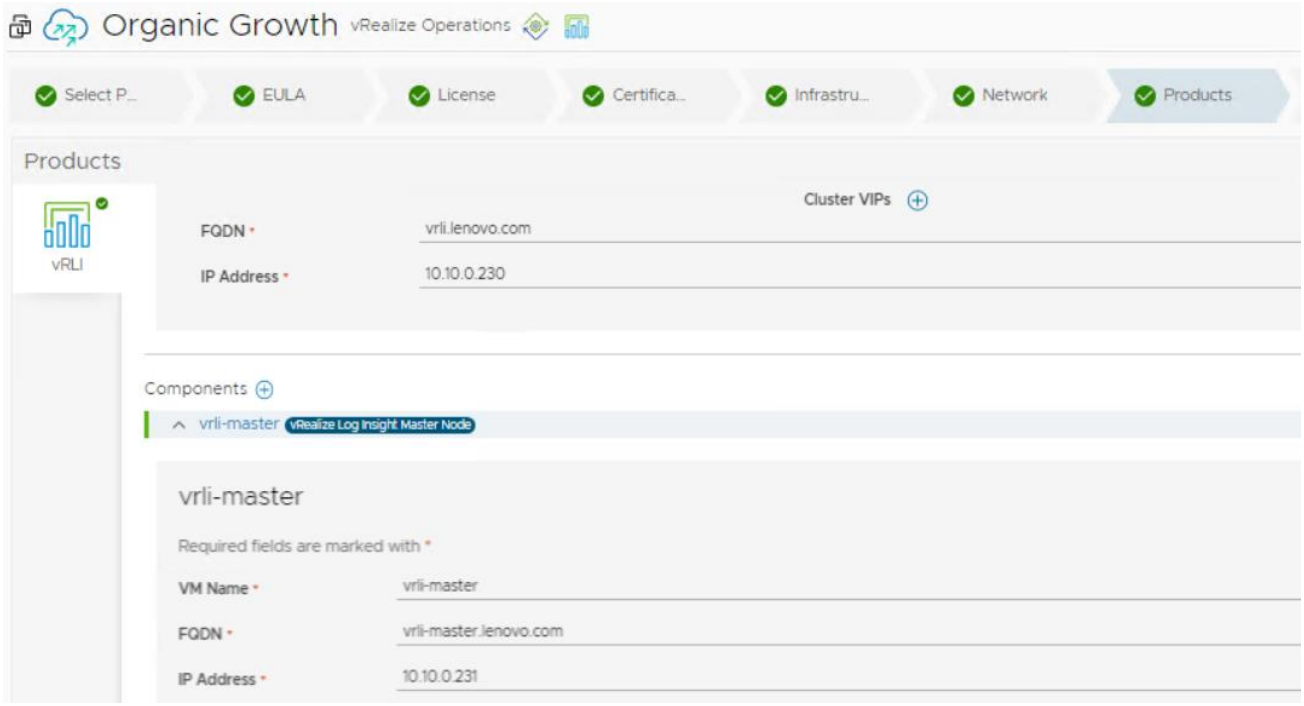

In the same window fill in the VM Names. The FQDNs and IP addresses for the worker nodes and click Next:

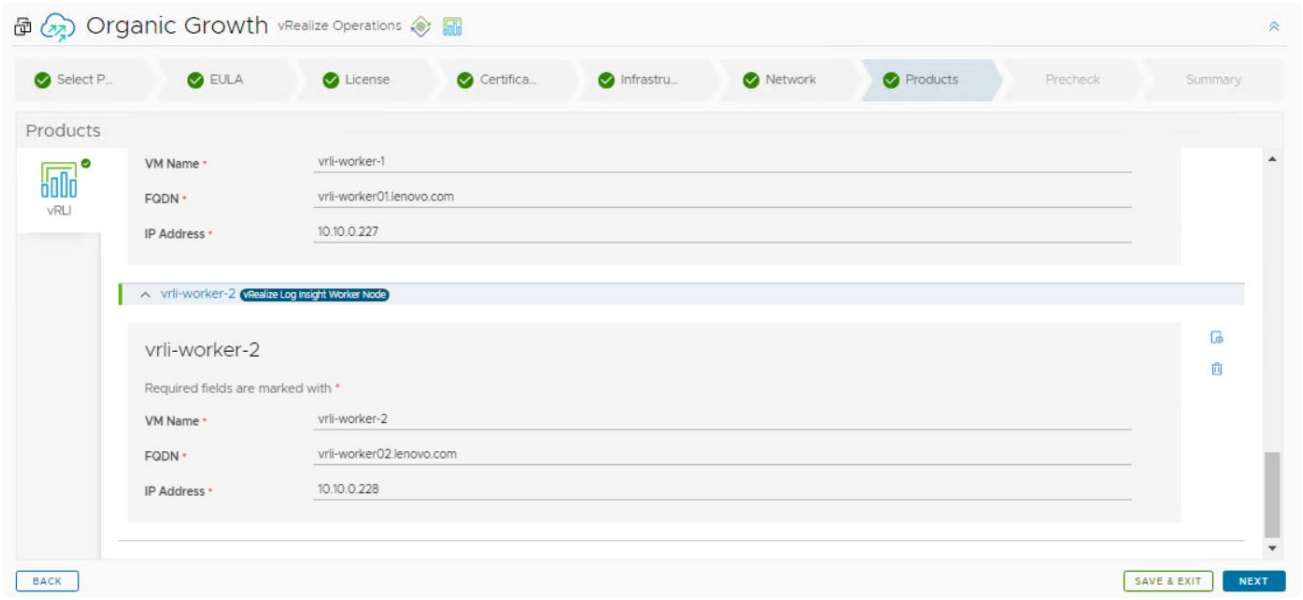

In the Precheck window click on Run Precheck:

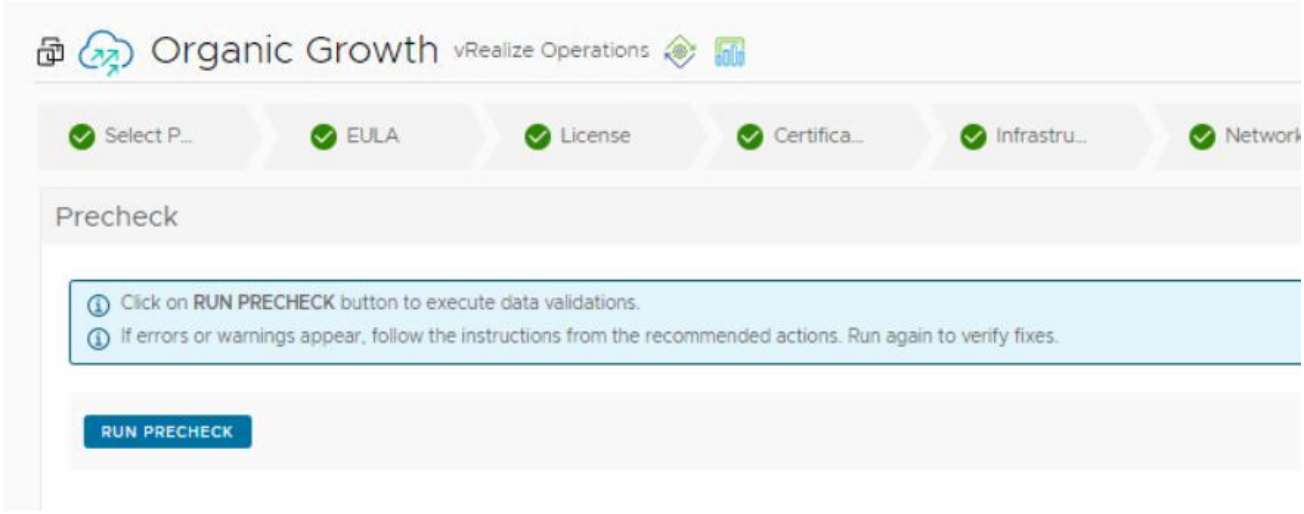

Click Next after all validations are successful:

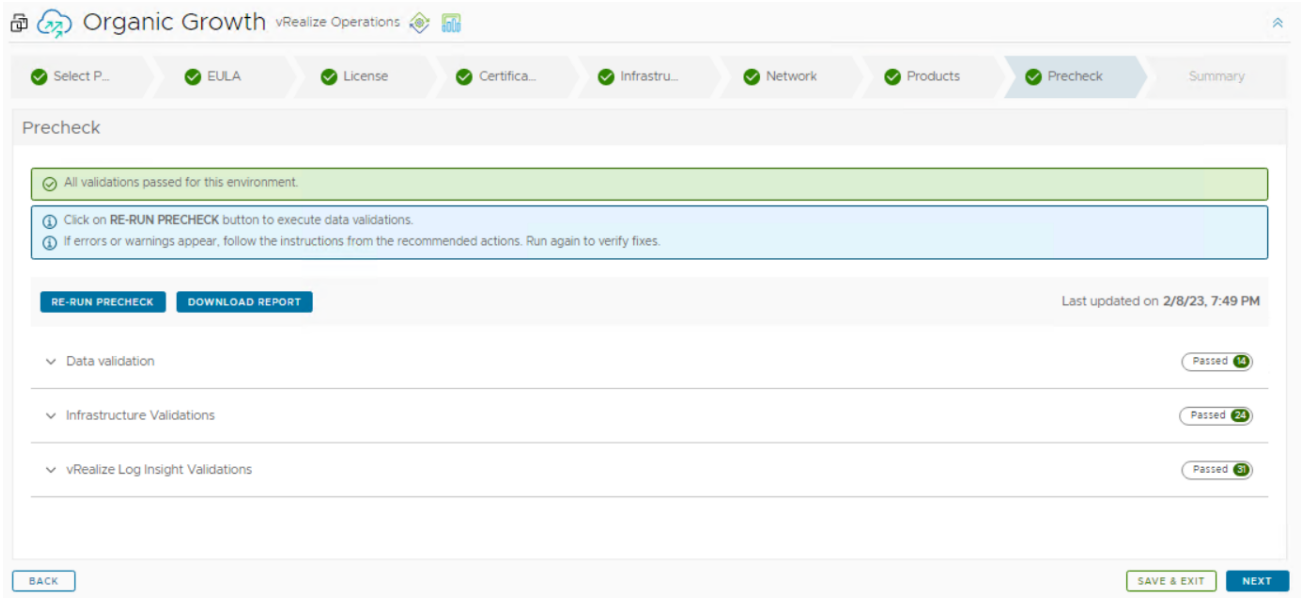

In the Summary window check if all parameters are correct then click Submit:

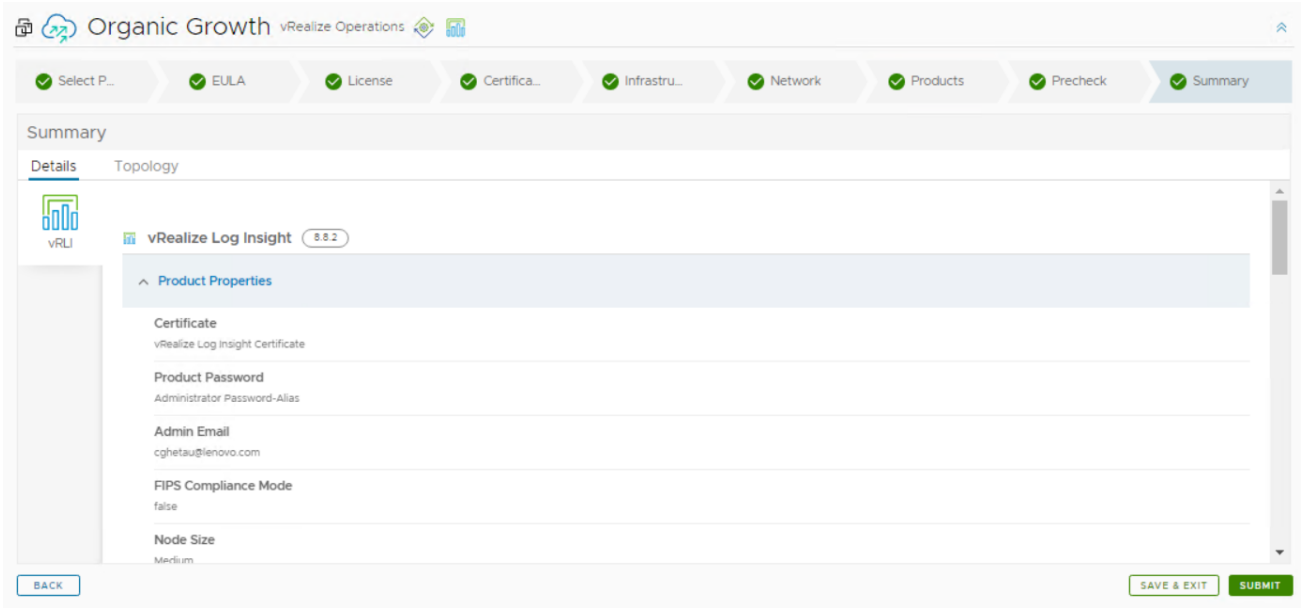

The Request Details window will be automatically switched to in order to monitor each deployment Stage:

Check the SDDC Manager to verify if the components are successfully installed and Active:

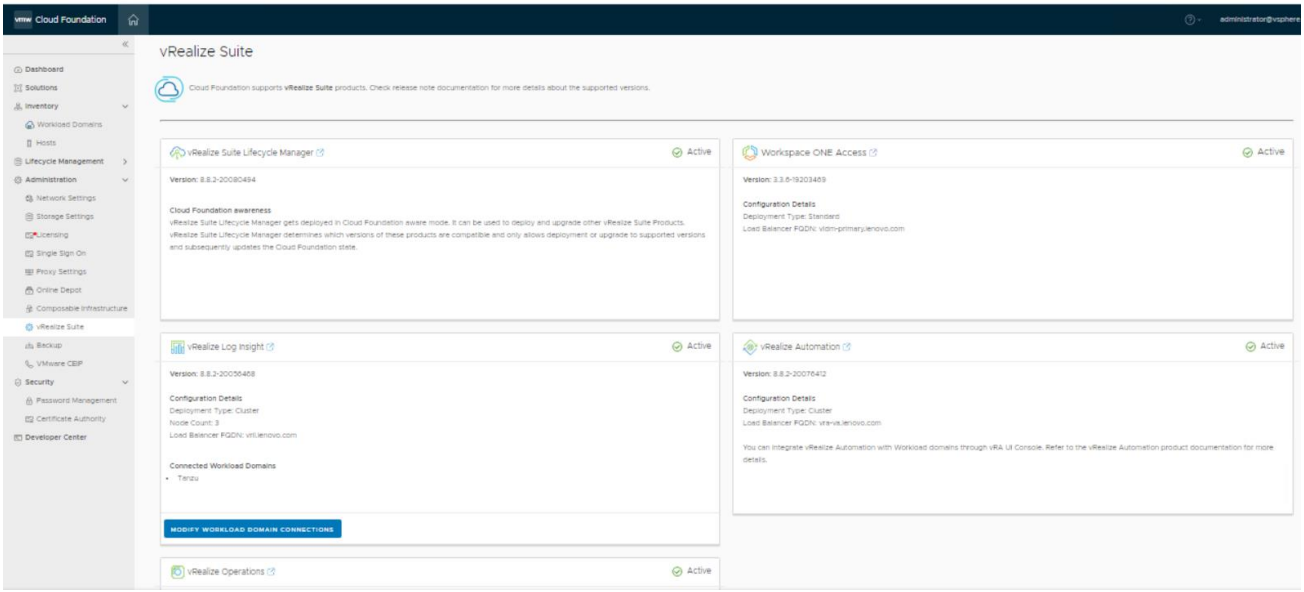

7.3.1.6 Deploy vSphere with Tanzu

Add the NSX-T Edge cluster to the VI – Workload Domain

In SDDC Manager go to Inventory > Workload Domains and click on the 3 dots near the newly created WD and select Add Edge Cluster:

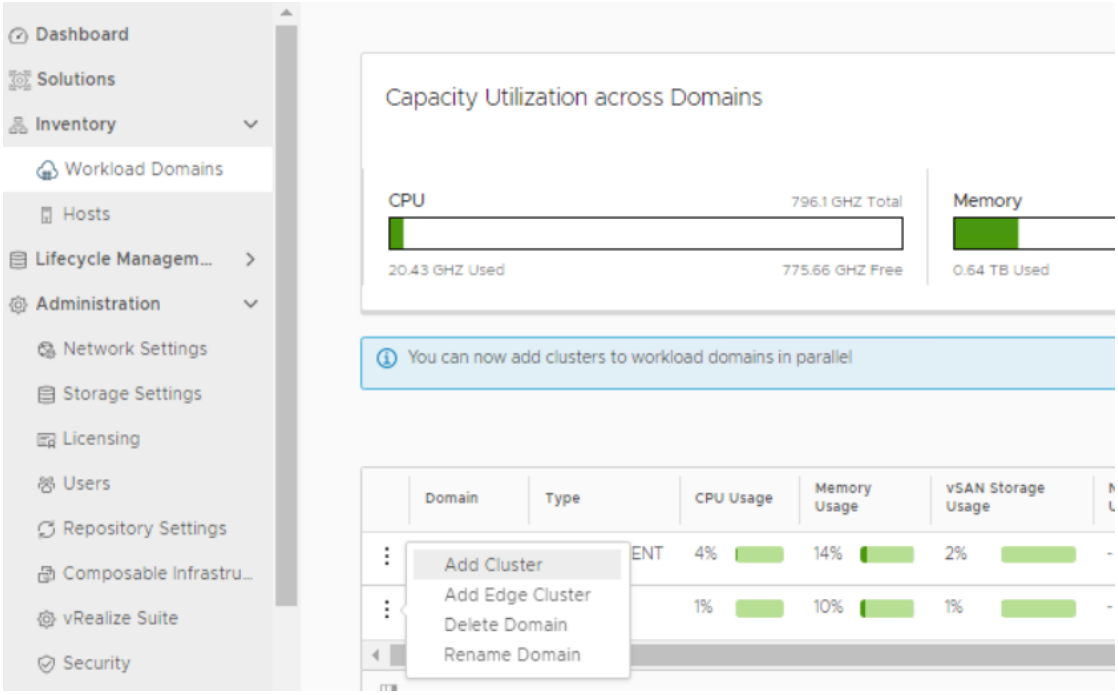

In the Edge Cluster Prerequisites Select All after verifying that all prerequisites have been met and click on the Begin button:

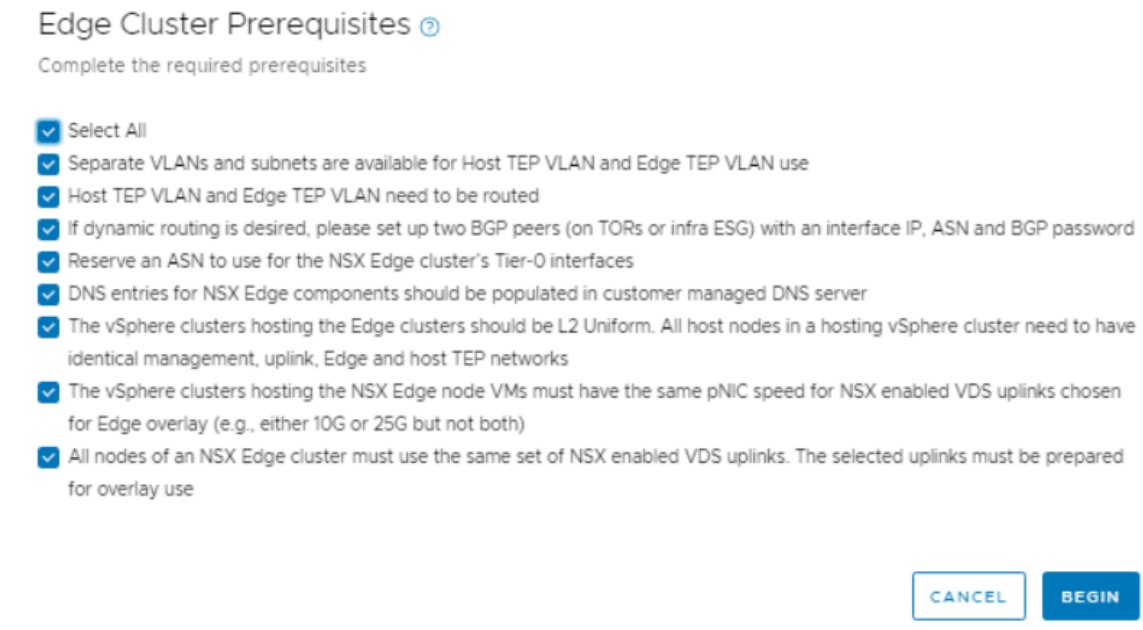

In the General Info window provide the following parameters and click on the Next button:

- Edge Cluster Name
- MTU: 9000
- ASN (make sure it matches the remote-as ASN configured for BGP on the physical switch)
- Tier-0 router name
- Tier-1 router name
- Edge Cluster Profile Type: Default
- Create passwords for Edge root, Edge admin and Edge audit accounts

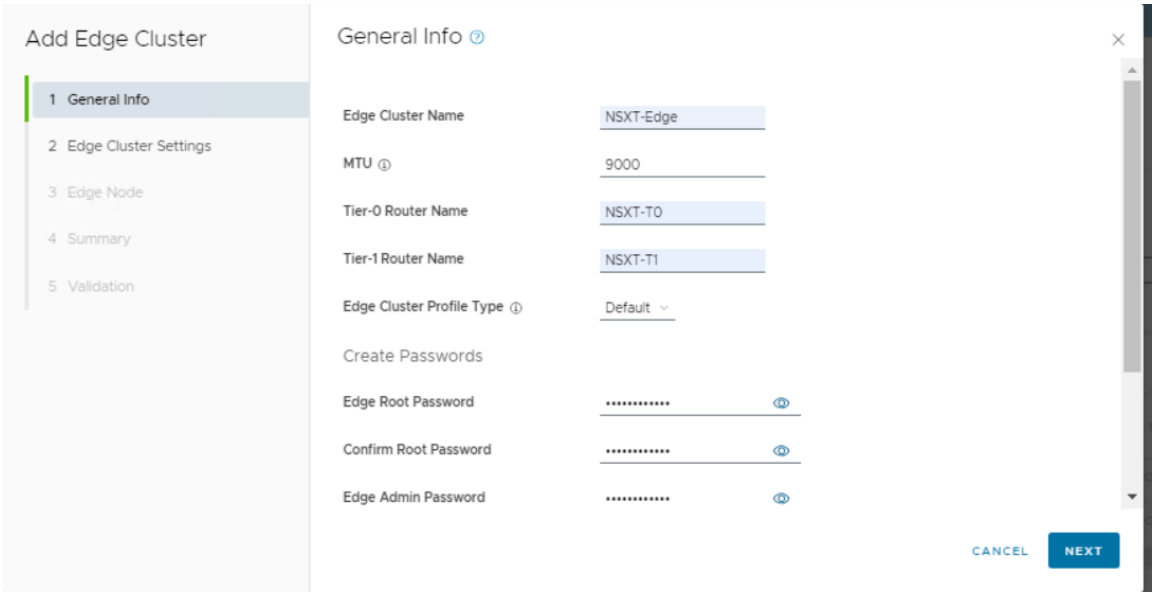

In the Edge Cluster Settings windows provide the following parameters and click on the Next button:

- Select Kubernetes – Workload Management
- Select Tier-0 Routing Type: EBGP
- BGP ASN – as configured prior on the first physical switch BGP configuration

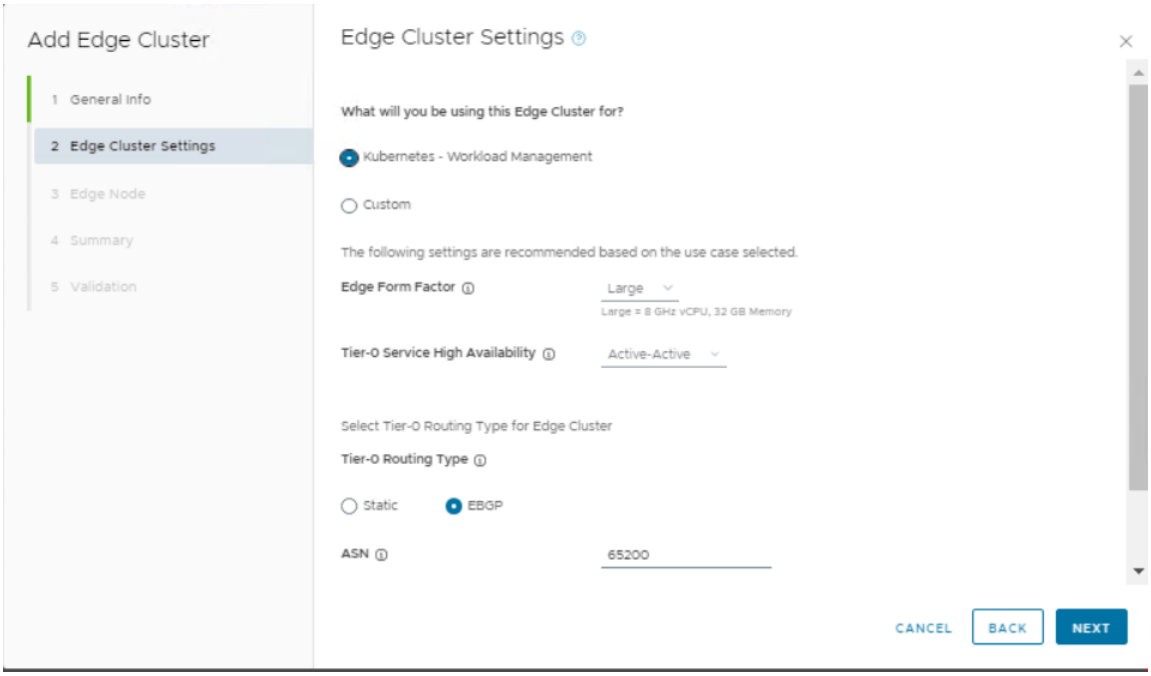

In the Edge Node windows provide the following parameters for minimum 2 nodes, after adding the Edge Nodes click on the Next button:

- Edge Node #1 FQDN (already created in the DNS)
- Management IP (CIDR) – IP should be assigned from the Management subnet/vlan
- Management Gateway
- EDGE TEP 1 IP (CIDR) – IP must be assigned from the EDGE TEP subnet/vlan
- EDGE TEP 2 IP (CIDR) - IP must be assigned from the EDGE TEP subnet/vlan
- EDGE TEP Gateway IP
- EDGE TEP VLAN – must be routable on the physical switch
- Select the Cluster created in the SDDC Manager for Tanzu deployment
- Select Cluster type: L2 uniform
- Primary Tier-0 Uplink VLAN
- Primary Tier-0 Uplink Interface IP (CIDR)
- BGP Peer ASN – as configured prior on the first second switch BGP configuration
- BGP Peer IP (CIDR) – the first physical switch IP interface which must be in the same VLAN as the Tier-0 Uplink interface
- BGP Peer password – as configured prior on the first physical switch BGP configuration
- Secondary Tier-0 Uplink VLAN
- Secondary Tier-0 Uplink Interface IP (CIDR)
- BGP Peer IP (CIDR) – the first physical switch IP interface which must be in the same VLAN as the Tier-0 Uplink interface (same as Primary)
- BGP Peer ASN – as configured prior on the first physical switch BGP configuration (same as the Primary Uplink)
- BGP Peer password - as configured prior on the first physical switch BGP configuration

Edge Node #1:

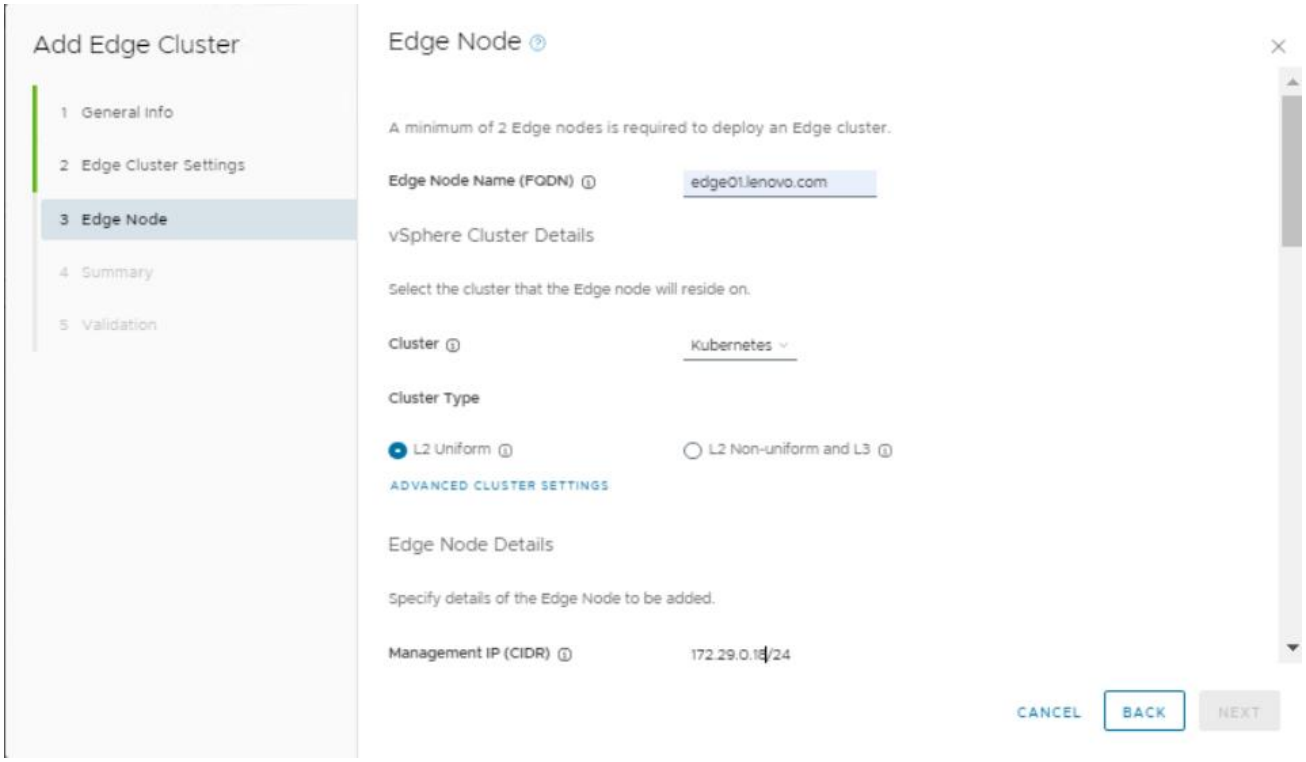
Add Edge Cluster
1 General Info
2 Edge Cluster Settings
3 Edge Node
4 Summary
5 Validation
Edge Node
A minimum of 2 Edge nodes is required to deploy an Edge cluster.
Edge Node Name (FQDN)
edge01.lenovo.com
vSphere Cluster Details
Select the cluster that the Edge node will reside on.
Cluster
Kubernetes
Cluster Type
L2 Uniform
L2 Non-uniform and L3
ADVANCED CLUSTER SETTINGS
Edge Node Details
Specify details of the Edge Node to be added.
Management IP (CIDR)
172.29.0.18/24
CANCEL
BACK
NEXT

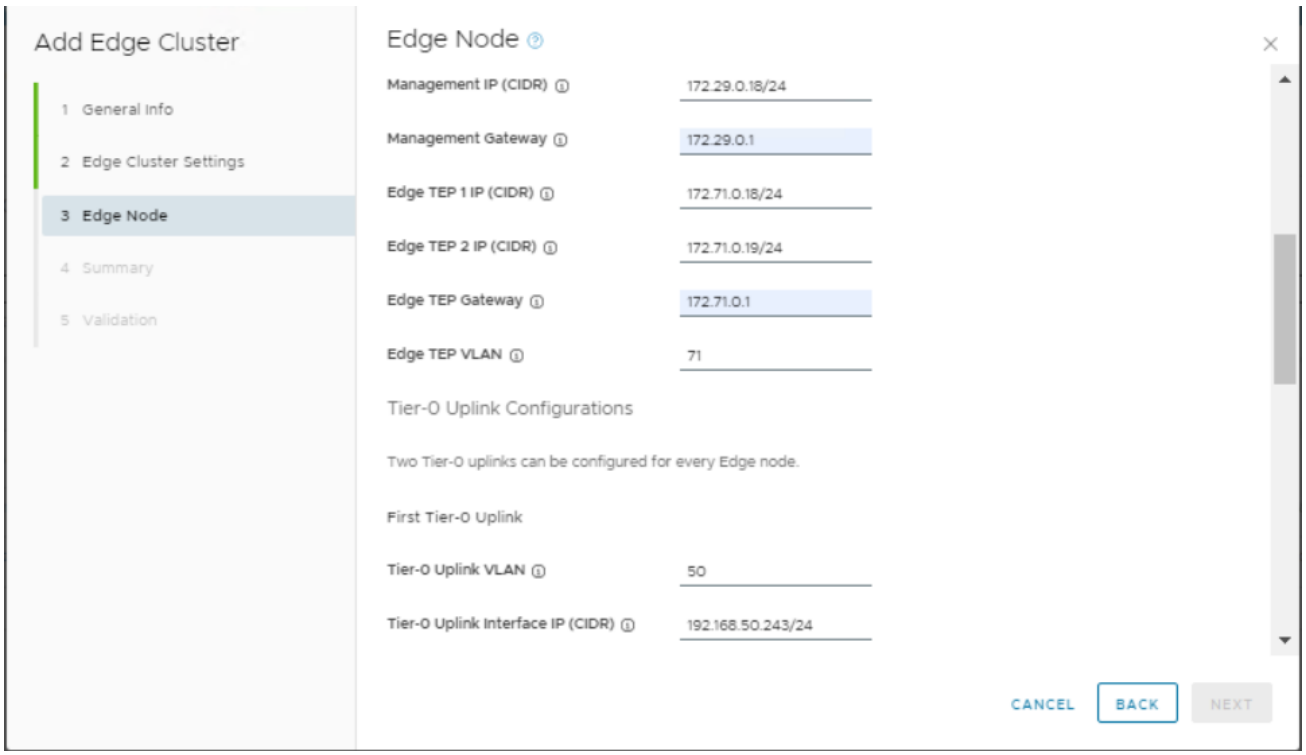
Add Edge Cluster
1 General Info
2 Edge Cluster Settings
3 Edge Node
4 Summary
5 Validation
Edge Node
Management IP (CIDR)
172.29.0.18/24
Management Gateway
172.29.0.1
Edge TEP 1 IP (CIDR)
172.71.0.18/24
Edge TEP 2 IP (CIDR)
172.71.0.19/24
Edge TEP Gateway
172.71.0.1
Edge TEP VLAN
71
Tier-0 Uplink Configurations
Two Tier-0 uplinks can be configured for every Edge node.
First Tier-0 Uplink
Tier-0 Uplink VLAN
50
Tier-0 Uplink Interface IP (CIDR)
192.168.50.243/24
CANCEL
BACK
NEXT

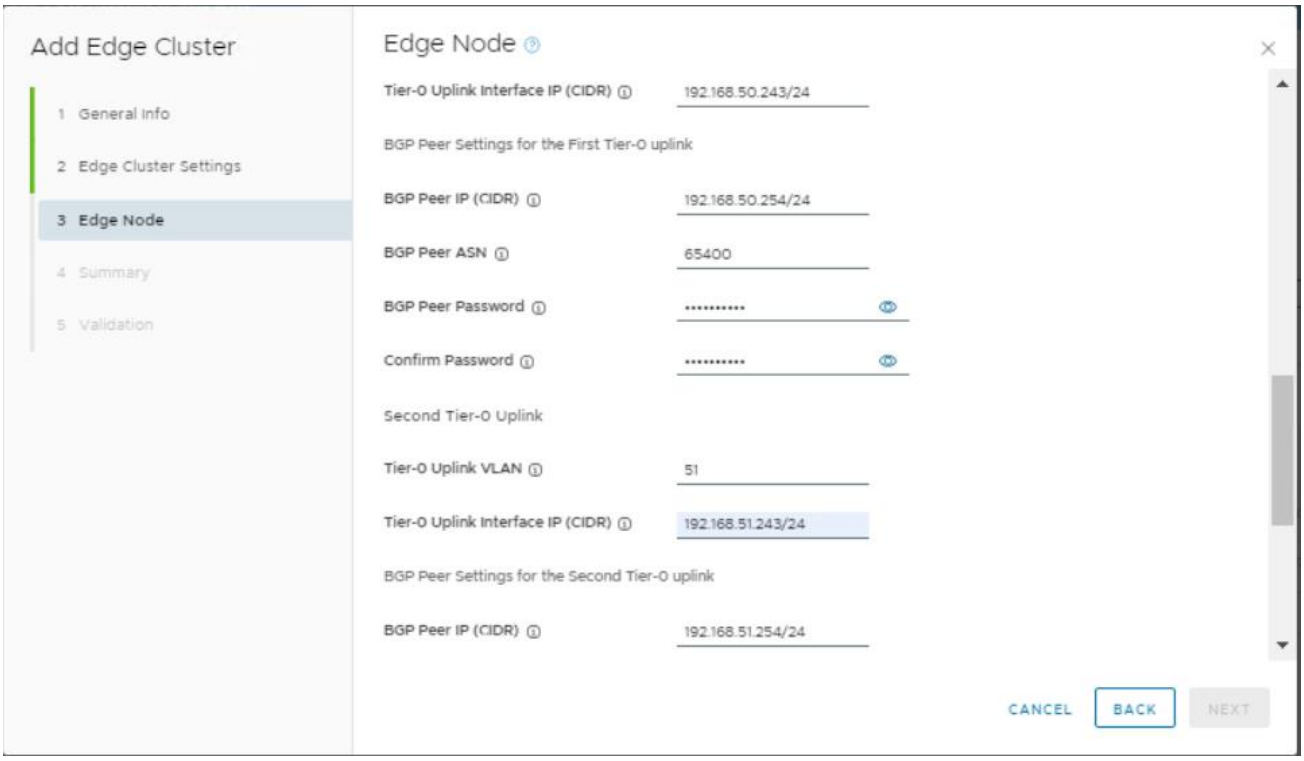

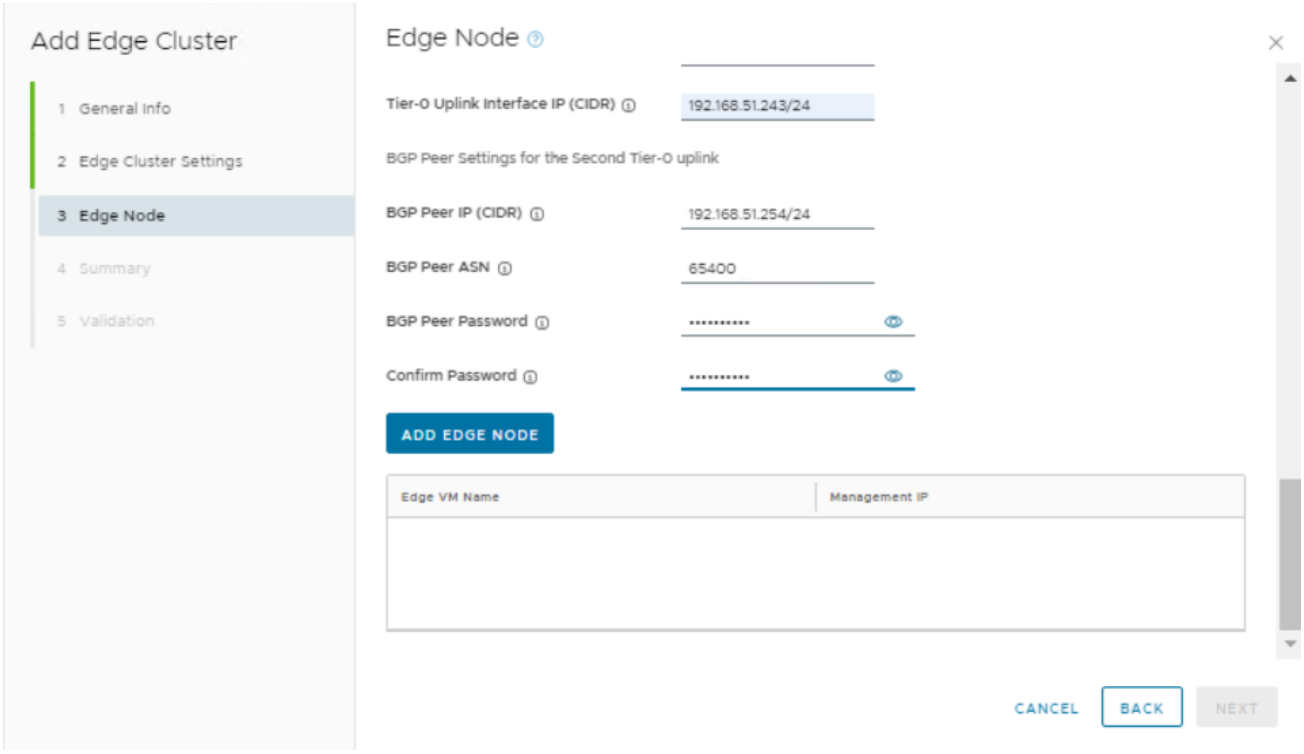

- Edge Node #2 FQDN (already created in the DNS)
- Management IP (CIDR) – IP should be assigned from the Management subnet/vlan
- Management Gateway
- EDGE TEP 1 IP (CIDR) – IP must be assigned from the EDGE TEP subnet/vlan

- EDGE TEP 2 IP (CIDR) - IP must be assigned from the EDGE TEP subnet/vlan
- EDGE TEP Gateway IP
- EDGE TEP VLAN – must be routable on the physical switch
- Select the Cluster created in the SDDC Manager for Tanzu deployment
- Select Cluster type: L2 uniform
- Primary Tier-0 Uplink VLAN
- Primary Tier-0 Uplink Interface IP (CIDR)
- BGP Peer IP (CIDR) – the second physical switch IP interface which must be in the same VLAN as the Tier-0 Uplink interface
- BGP Peer ASN – as configured prior on the first second switch BGP configuration
- BGP Peer password – as configured prior on the second physical switch BGP configuration
- Secondary Tier-0 Uplink VLAN
- Secondary Tier-0 Uplink Interface IP (CIDR)
- BGP Peer IP (CIDR) – the second physical switch IP interface which must be in the same VLAN as the Tier-0 Uplink interface (same as Primary)
- BGP Peer ASN – as configured prior on the second physical switch BGP configuration (same as the Primary Uplink)
- BGP Peer password - as configured prior on the second physical switch BGP configuration

Click on ADD EDGE NODE button and then on ADD MORE EDGE NODES to add the second node:

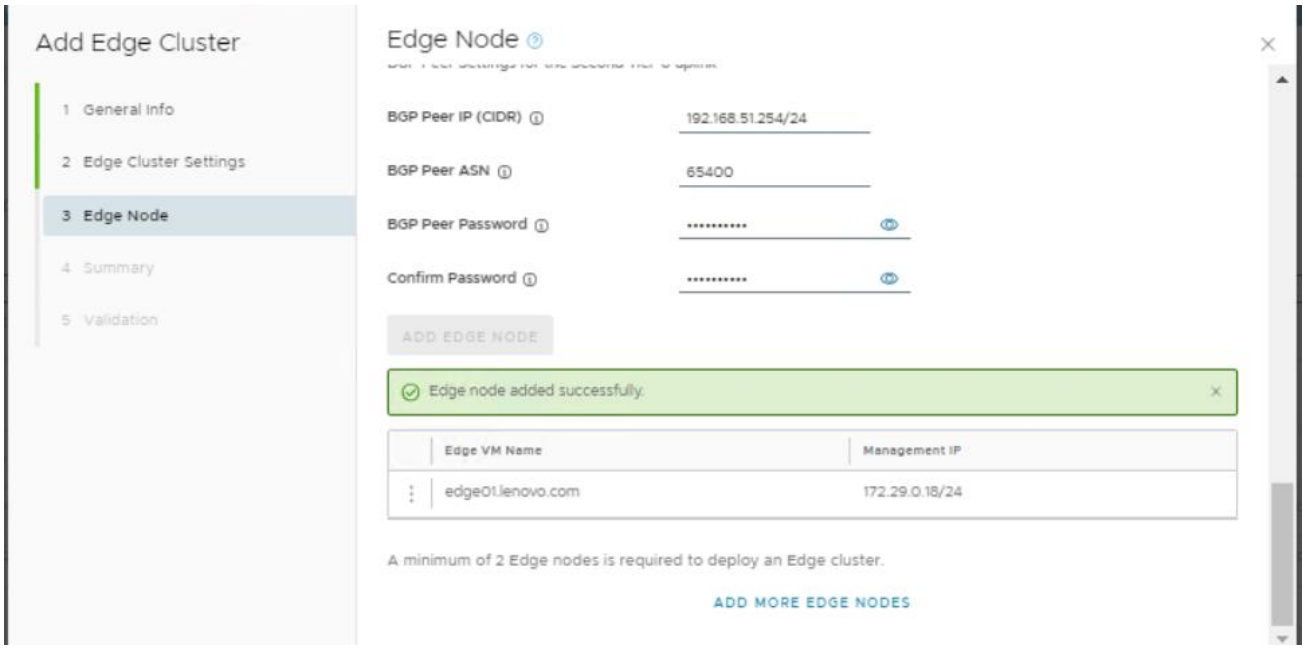

Edge Node #2:

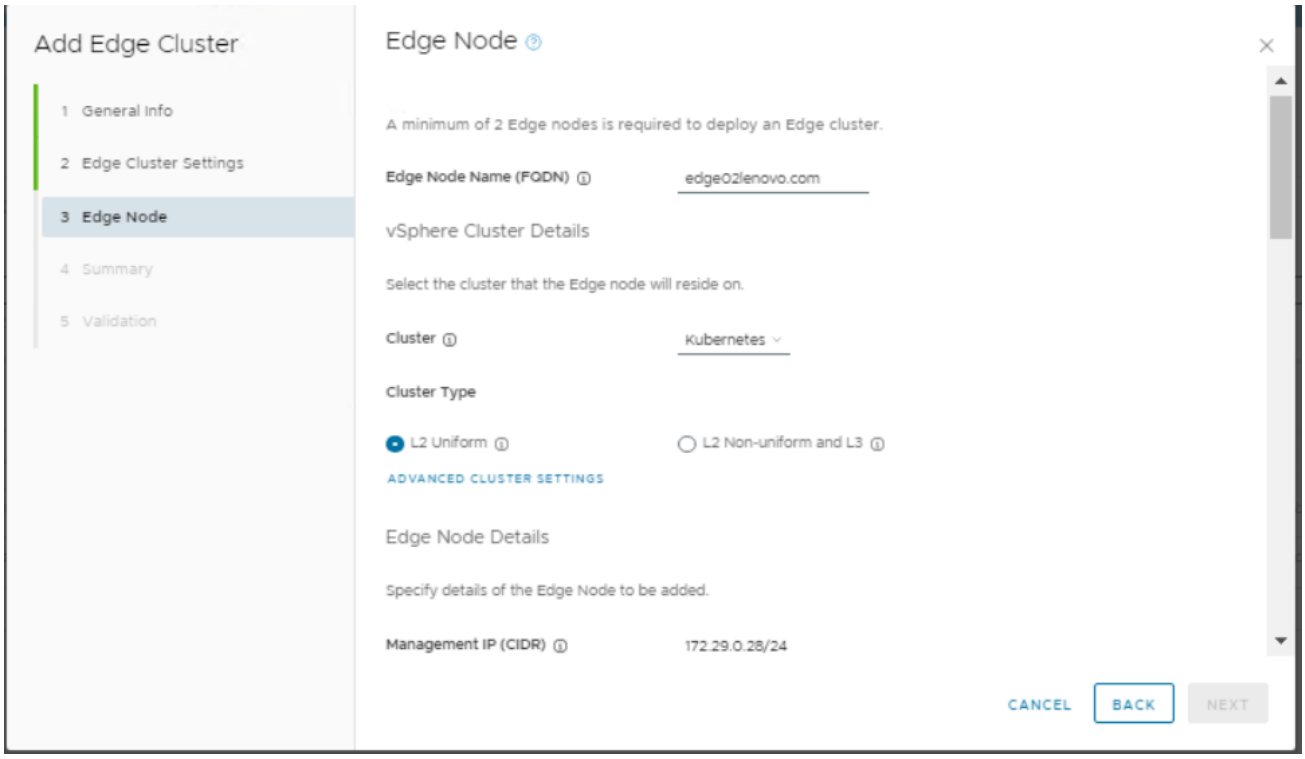
Add Edge Cluster
1 General Info
2 Edge Cluster Settings
3 Edge Node
4 Summary
5 Validation
Edge Node
A minimum of 2 Edge nodes is required to deploy an Edge cluster.
Edge Node Name (FQDN)
edge02lenovo.com
vSphere Cluster Details
Select the cluster that the Edge node will reside on.
Cluster
Kubernetes
Cluster Type
L2 Uniform
L2 Non-uniform and L3
ADVANCED CLUSTER SETTINGS
Edge Node Details
Specify details of the Edge Node to be added.
Management IP (CIDR)
172.29.0.28/24
CANCEL
BACK
NEXT

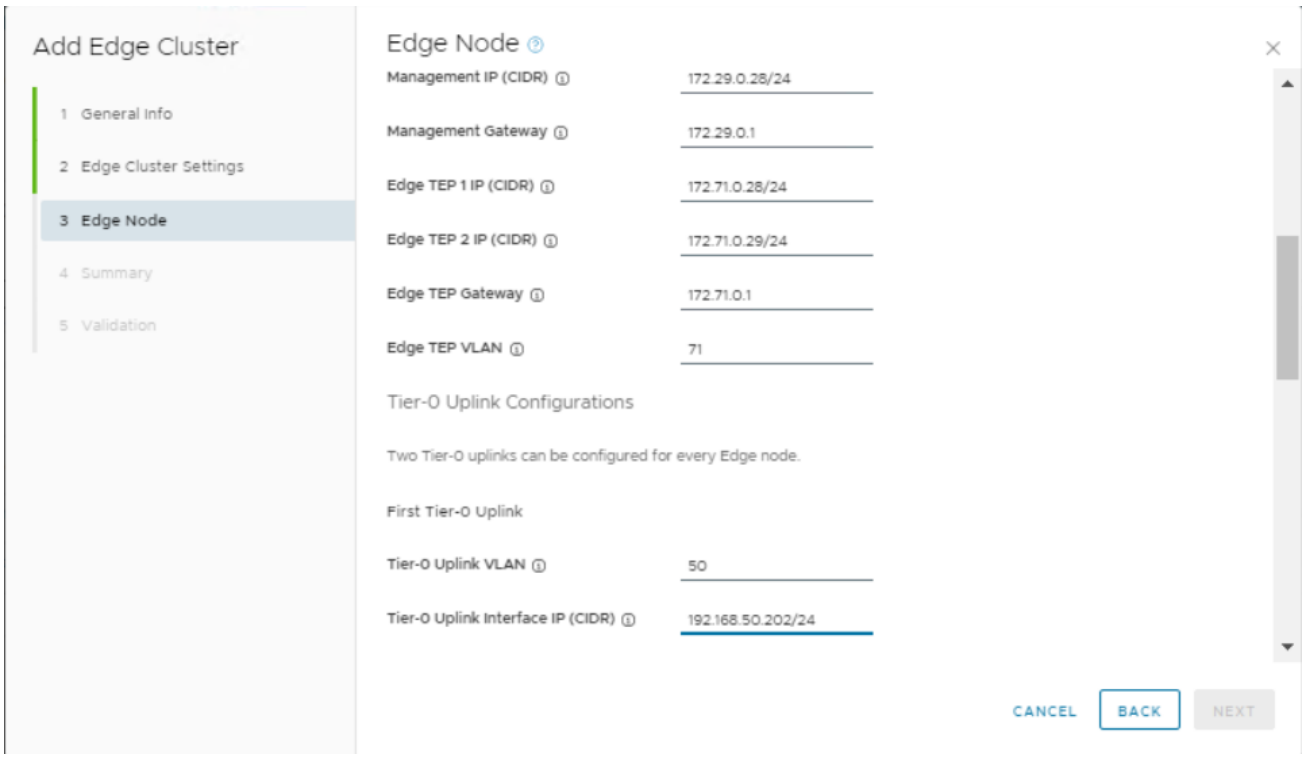
Add Edge Cluster
1 General Info
2 Edge Cluster Settings
3 Edge Node
4 Summary
5 Validation
Edge Node
Management IP (CIDR)
172.29.0.28/24
Management Gateway
172.29.0.1
Edge TEP 1 IP (CIDR)
172.71.0.28/24
Edge TEP 2 IP (CIDR)
172.71.0.29/24
Edge TEP Gateway
172.71.0.1
Edge TEP VLAN
71
Tier-0 Uplink Configurations
Two Tier-0 uplinks can be configured for every Edge node.
First Tier-0 Uplink
Tier-0 Uplink VLAN
50
Tier-0 Uplink Interface IP (CIDR)
192.168.50.202/24
CANCEL
BACK
NEXT

Click on ADD EDGE NODE button to add the second node

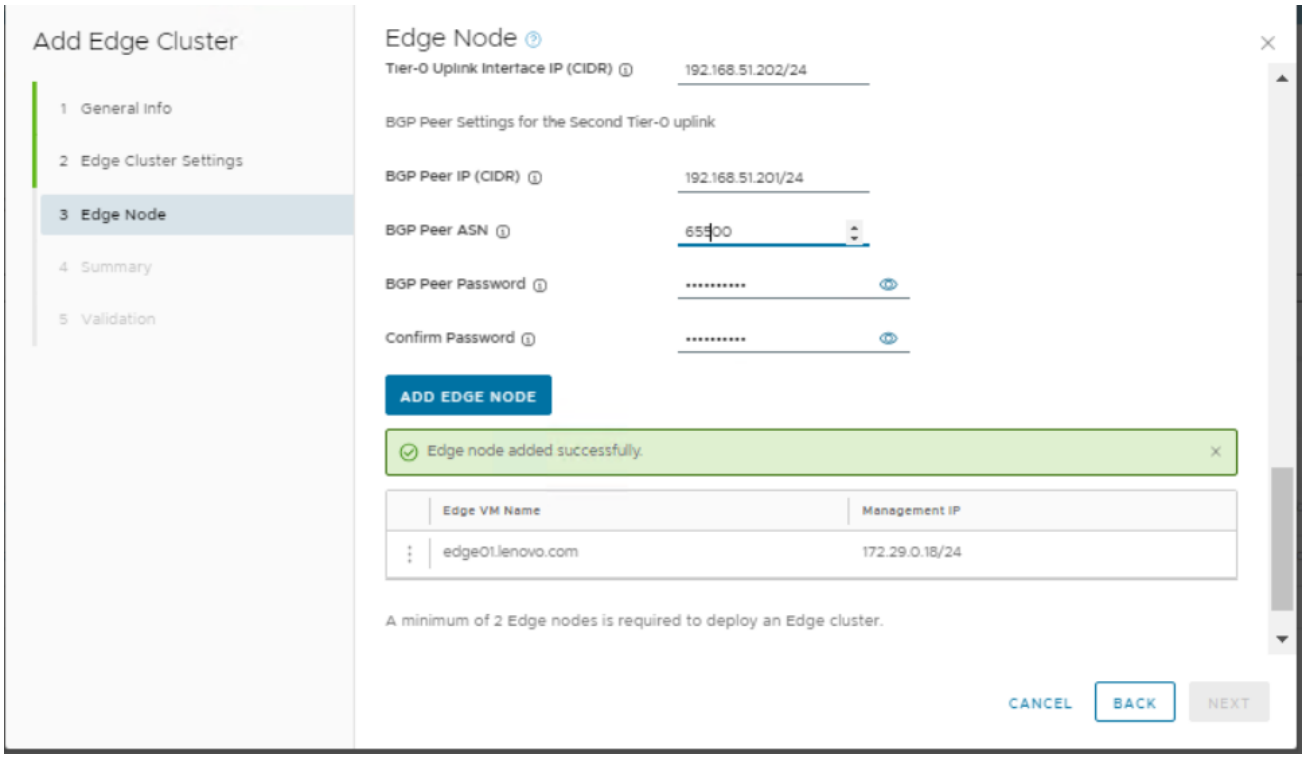

Click the Next button after both Edge Nodes are added:

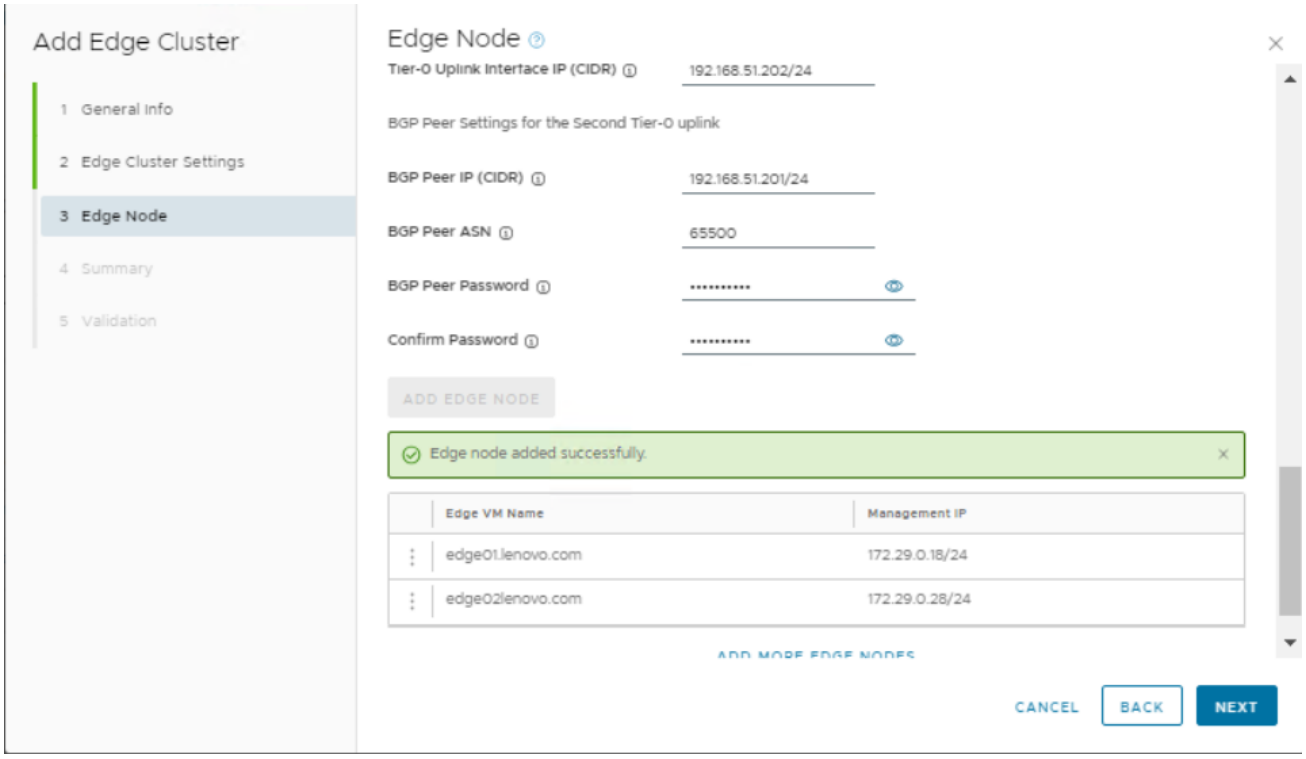

In the Summary windows verify that the Edge Cluster and Nodes have been properly configured then click on the Next button:

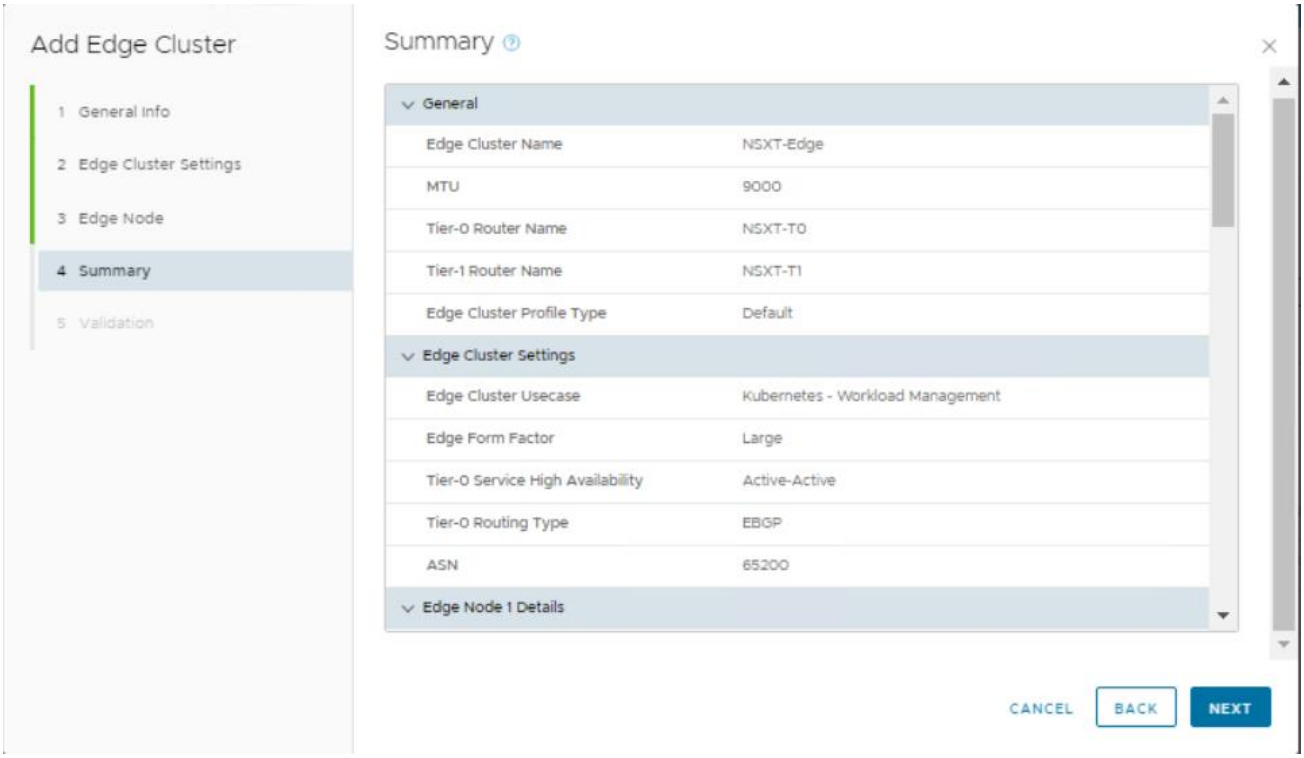

In the Validation window some checks are automatically done, click on the FINISH button if all validations have been successful, otherwise revise the previous settings that failed:

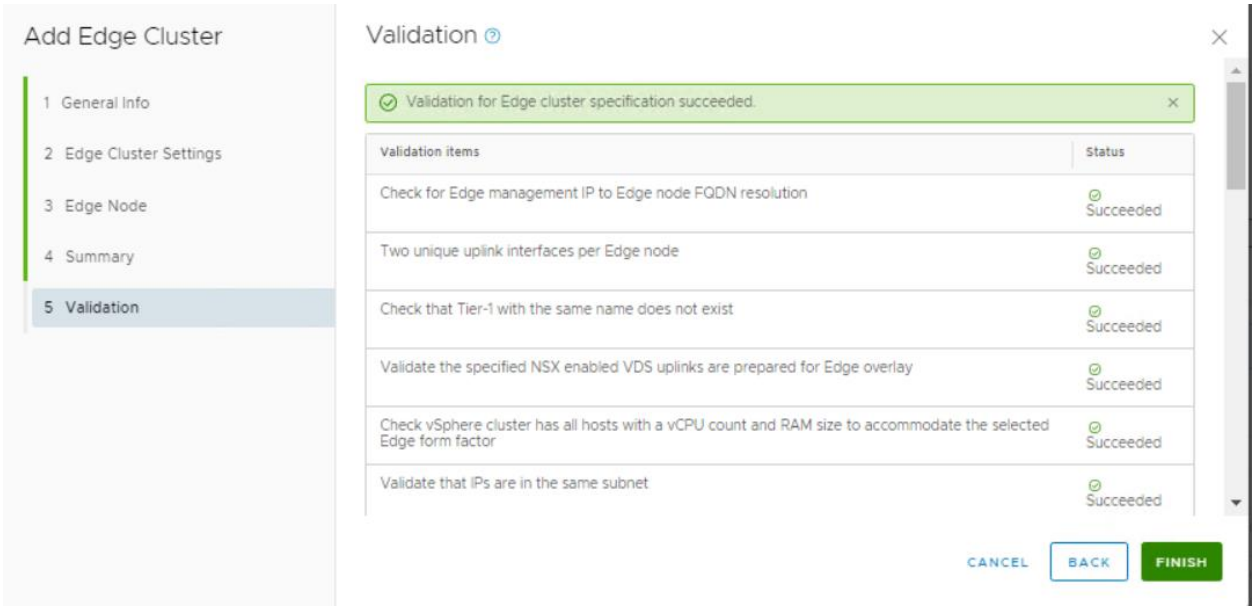

- **Deploy Kubernetes – Workload Management Solution**

In SDDC Manager go to Solutions > Deploy

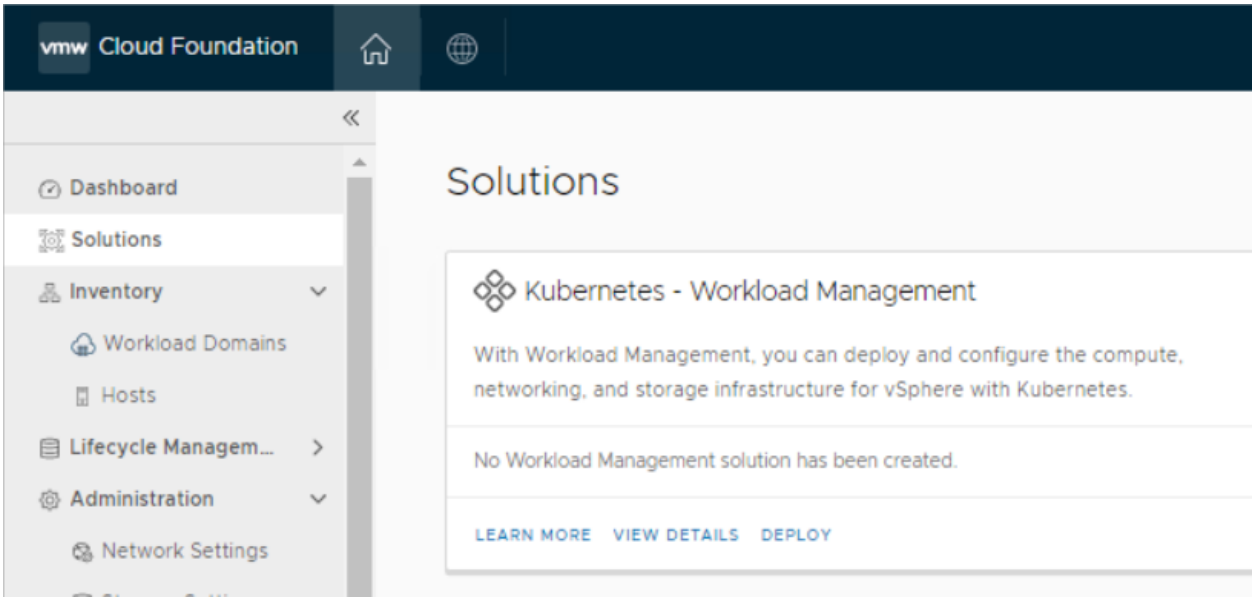

In the Workload Management Deployment Prerequisites windows verify that all the prerequisites have been met and click on Select All checkbox then click on the BEGIN button after Adding the Content Library:

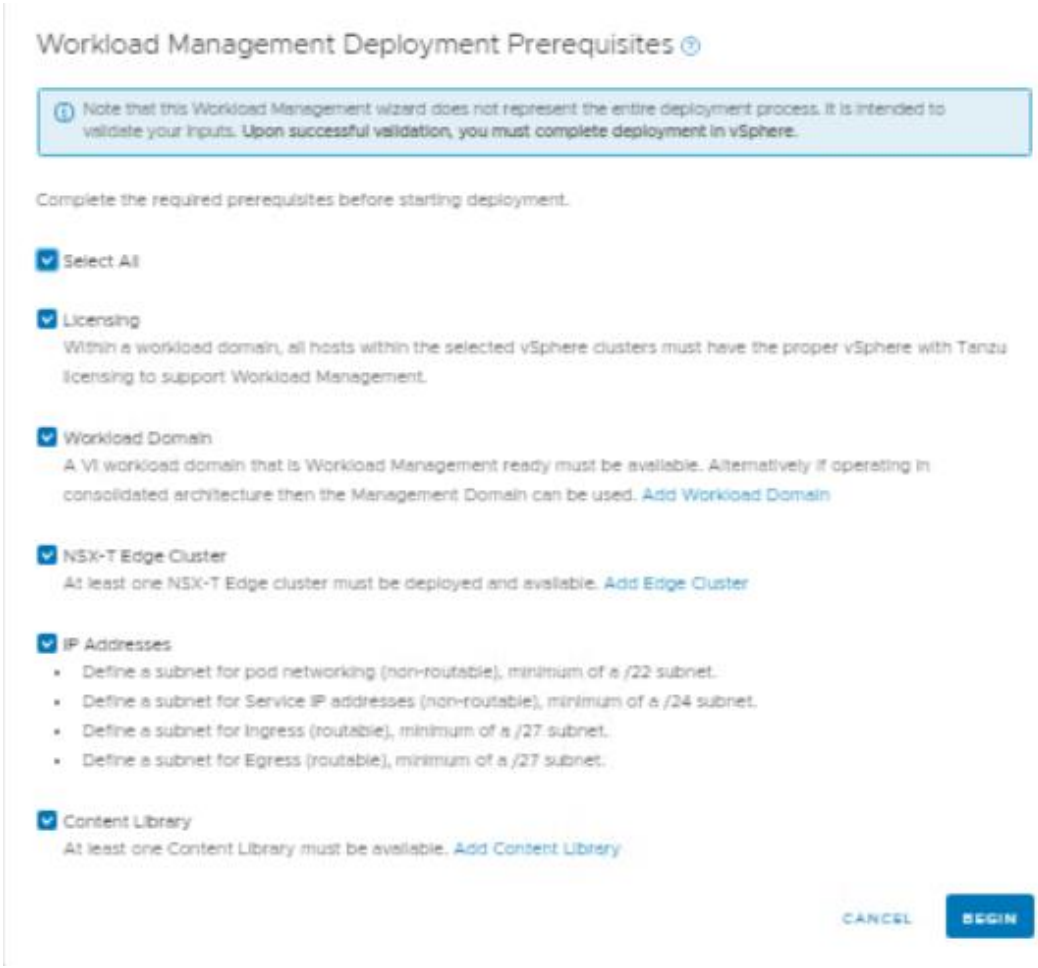

Create a Content Library in vCenter:

Click on Add Content Library to open the vSphere Client then go to Menu > Content Libraries and click on Create button:

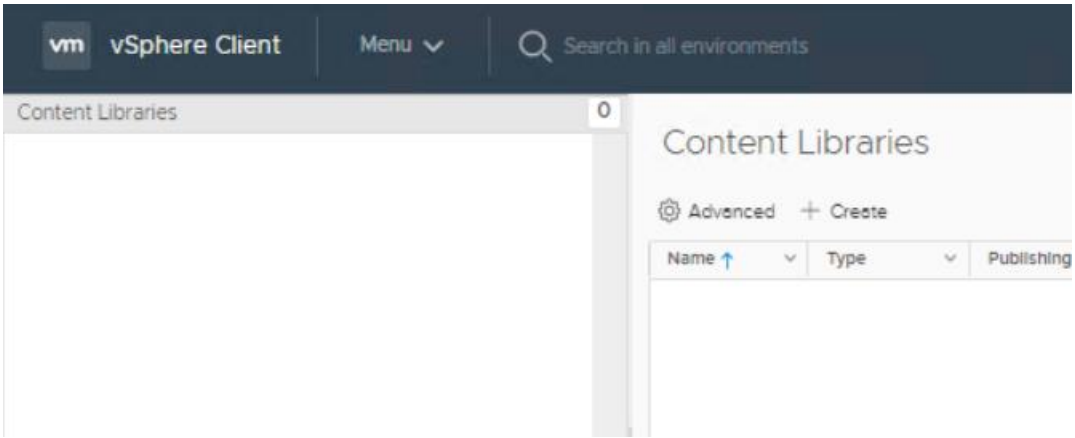

In the Name and Location window provide the following parameters then click Next button:

- Name – Tanzu Kubernetes Library name
- vCenter Server – select the vCenter deployed for Tanzu Kubernetes VI - Workload Domain

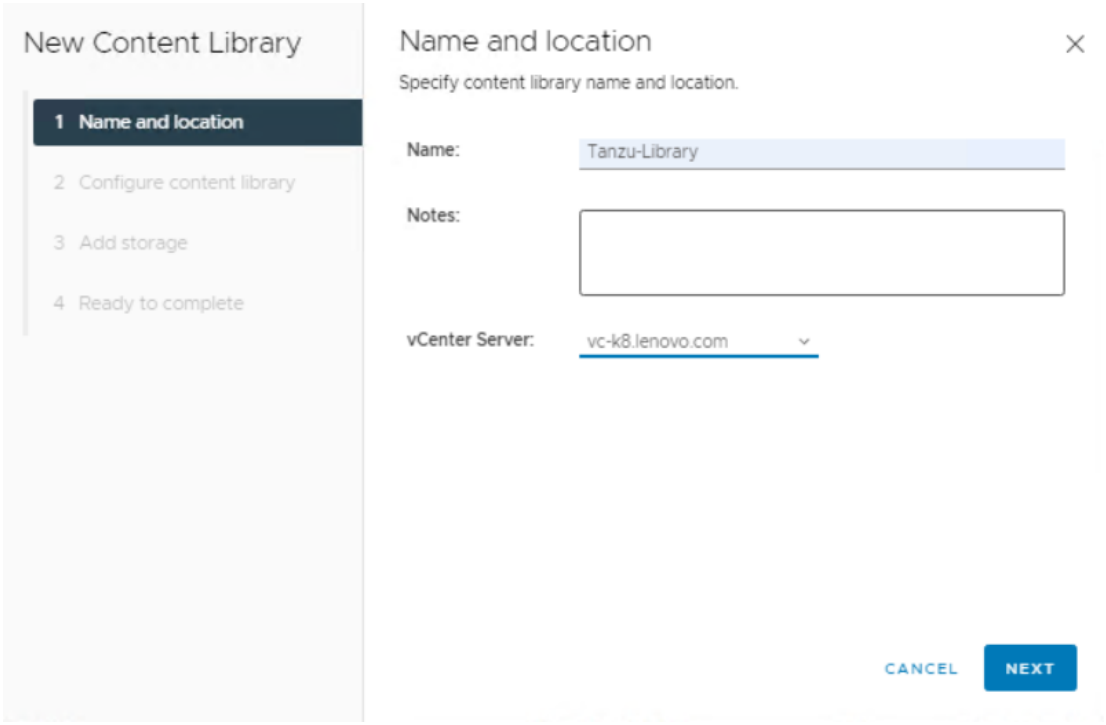

In the Configure Content Library window:

- Select Subscribed content library: https://wp-content.vmware.com/v2/latest/lib.json
- Select Download content: immediately

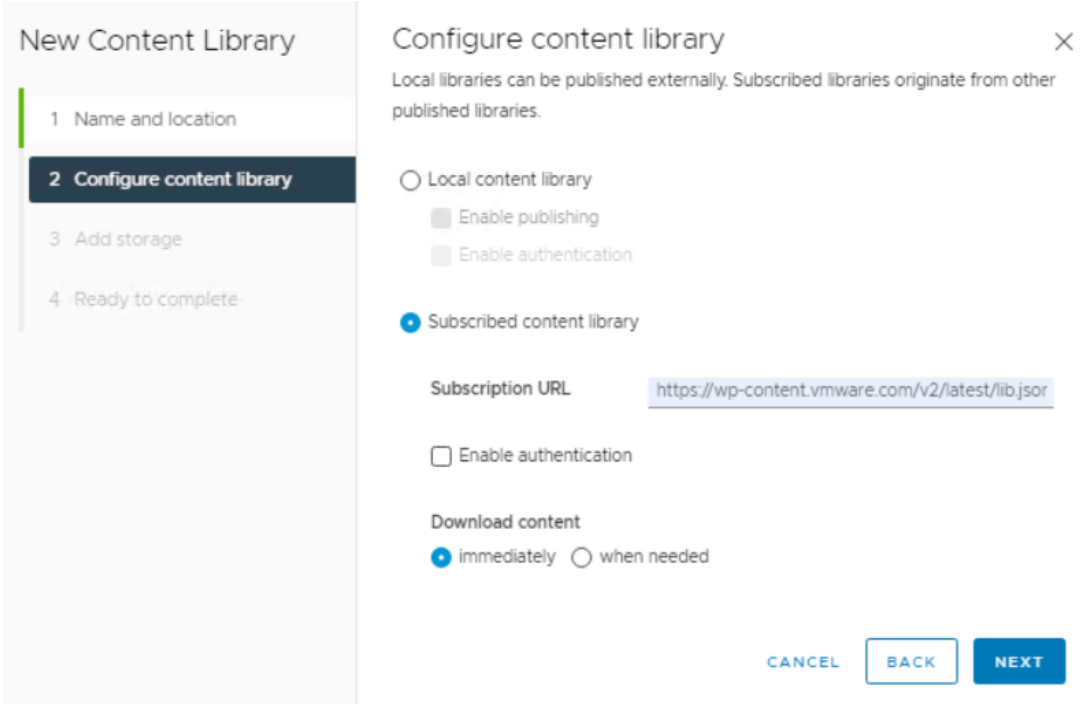

In the Tanzu Library – Authentication window click on Yes button:

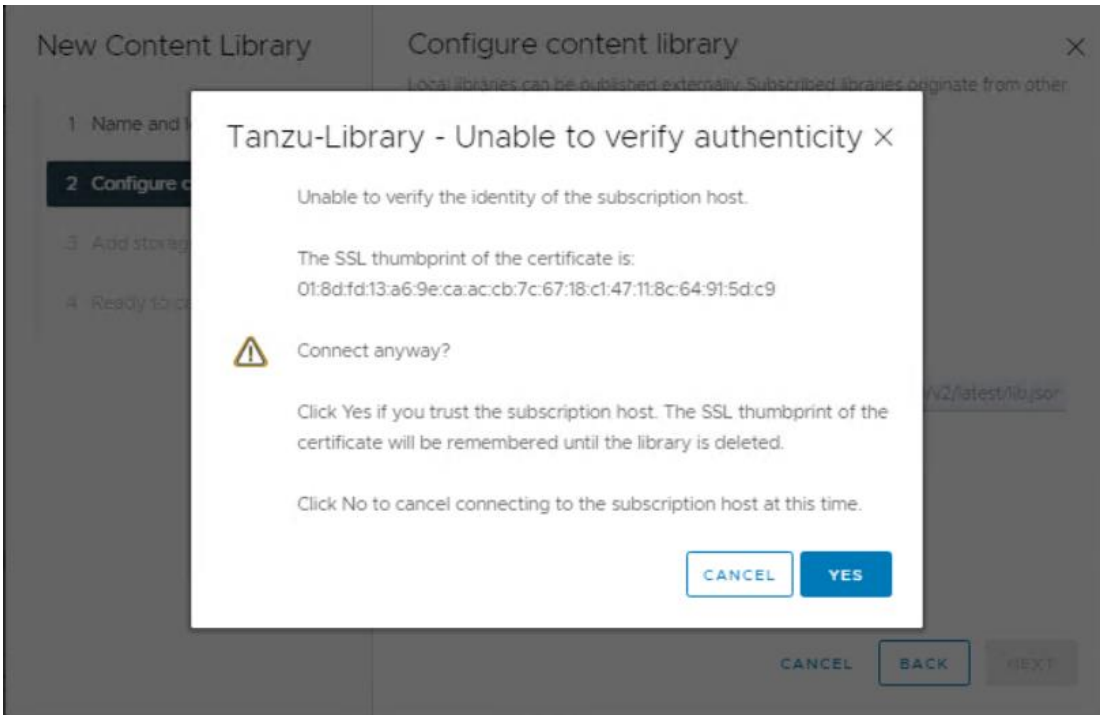

In the Add storage window select the Tanzu vSAN Cluster

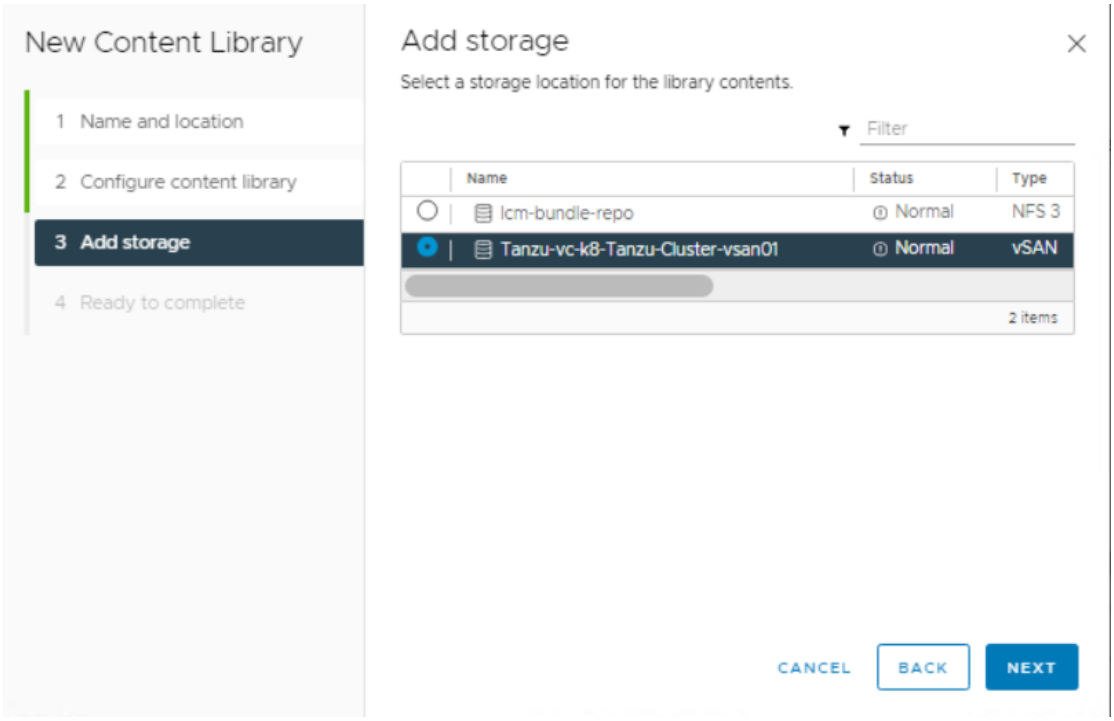

In the Ready to complete window review the settings and click on the FINISH button:

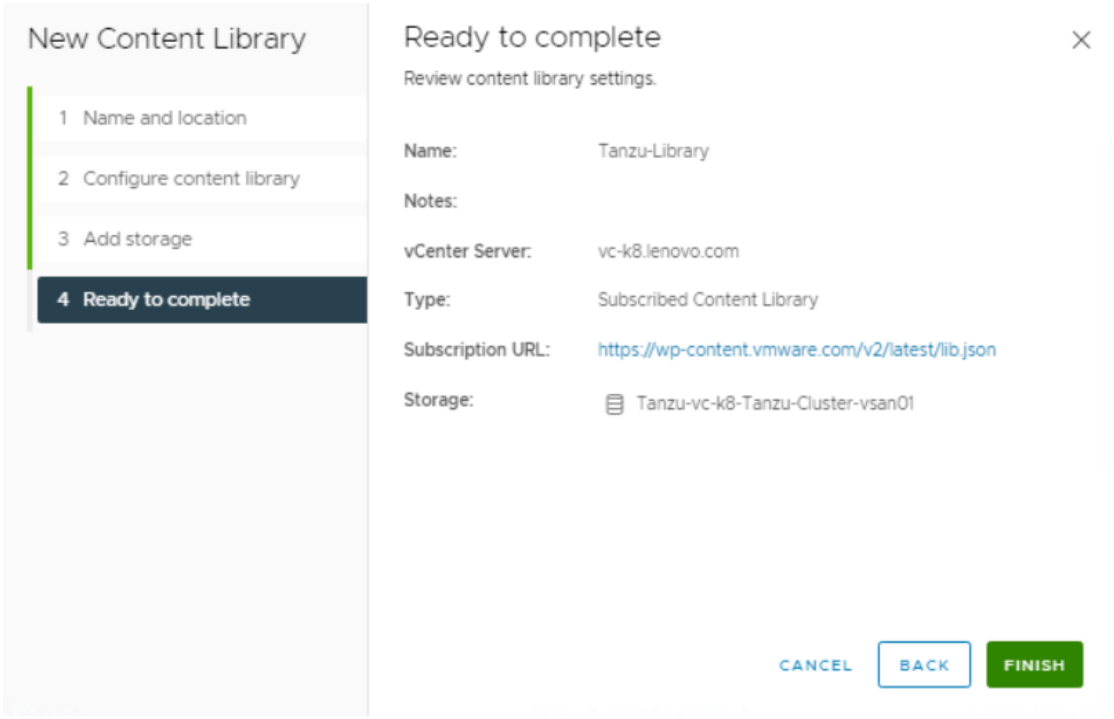

After Adding the Content Library in vSphere Client go back to SDDC Manager and click on the BEGIN button.

In the Select a Cluster window:

- Select the Workload Domain: Tanzu VI - Workload Domain
- Select the Tanzu Cluster

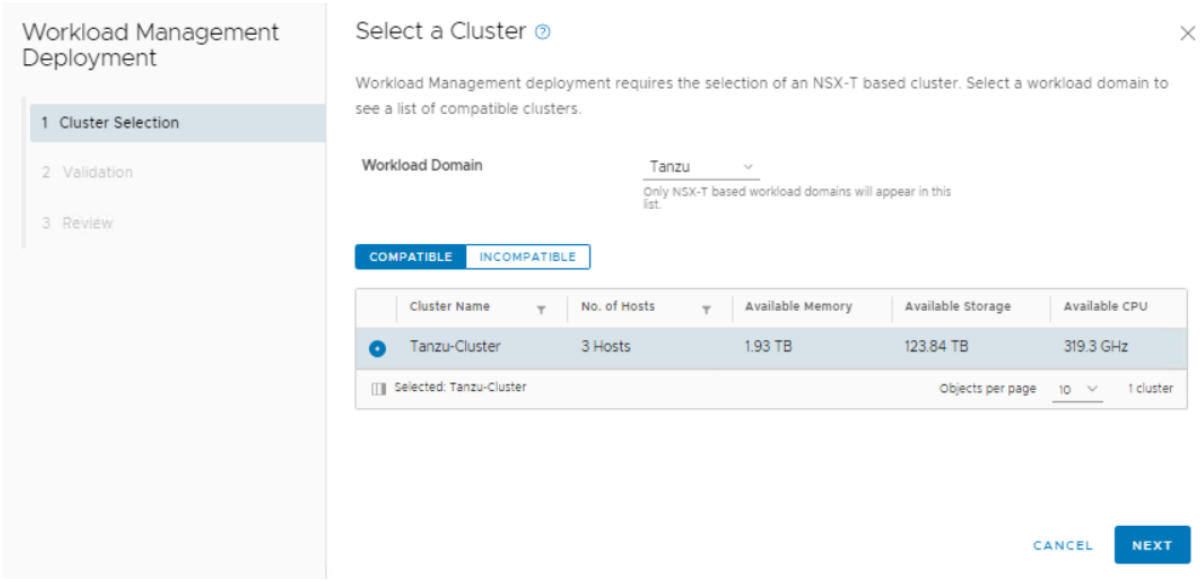

In the Validation windows click on the Next button after al validations are successful:

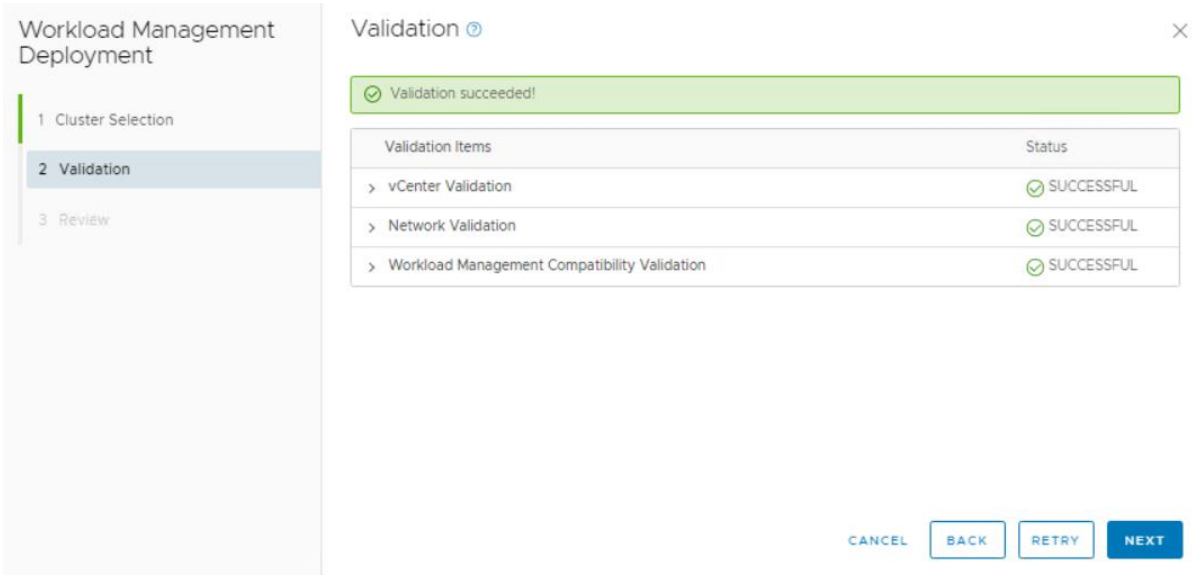

In the Review window click on the COMPLETE ON VSPHERE button:

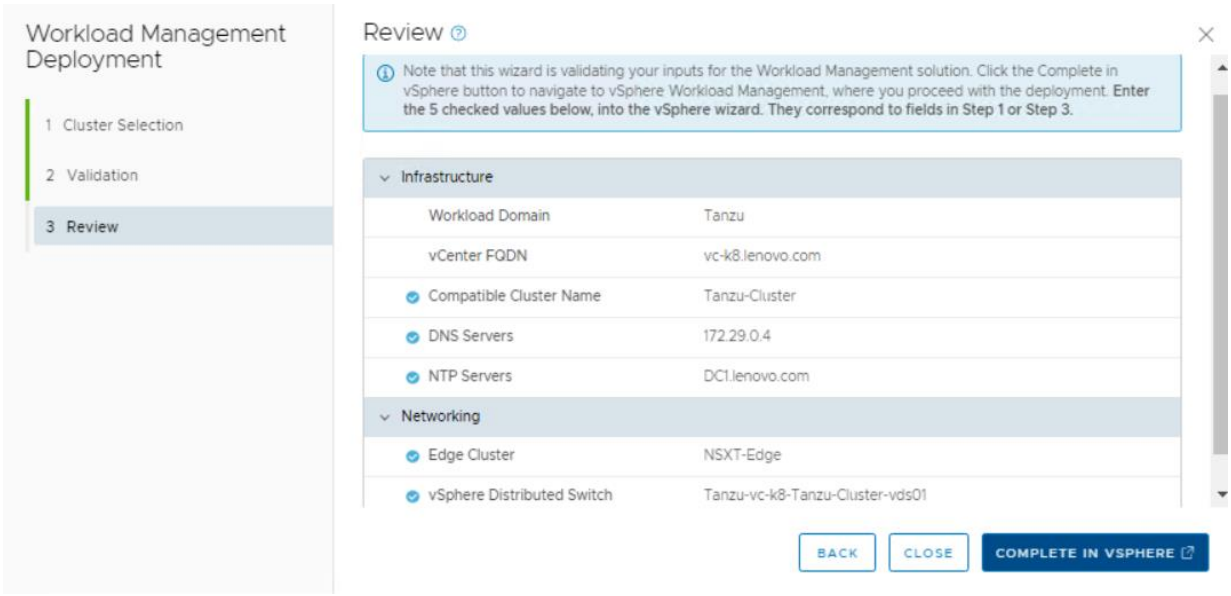

In the vCenter Server and Network window:

- Select a vCenter – select the vCenter deployed for Tanzu Kubernetes VI - Workload Domain
- Select a networking stack option: NSX-T
- Click on the Next button

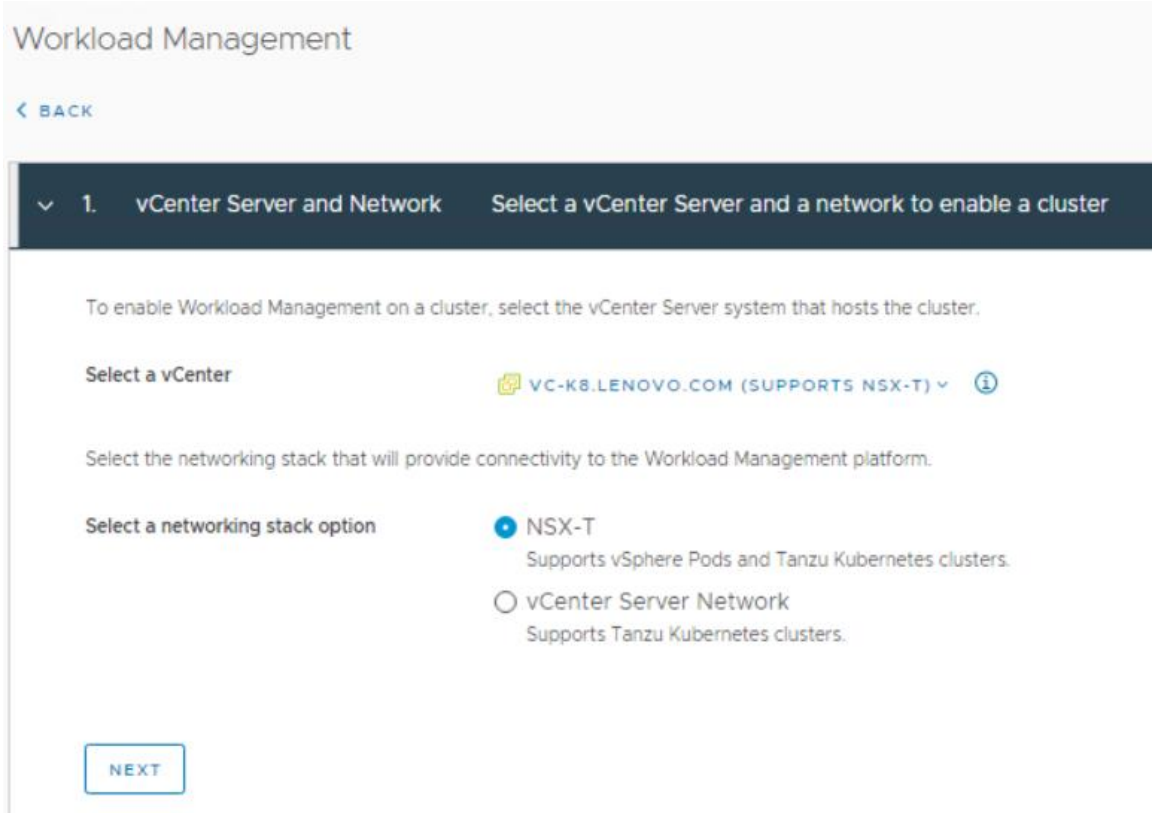

In the Supervisor location window provide a Supervisor name, select the Tanzu Cluster then click on the Next button:

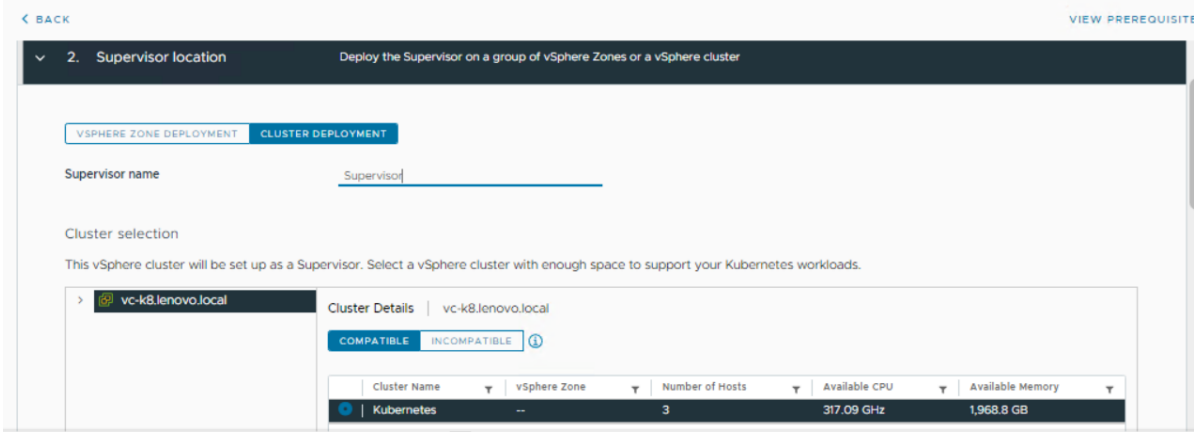

In the same window specify a Zone name to be created:

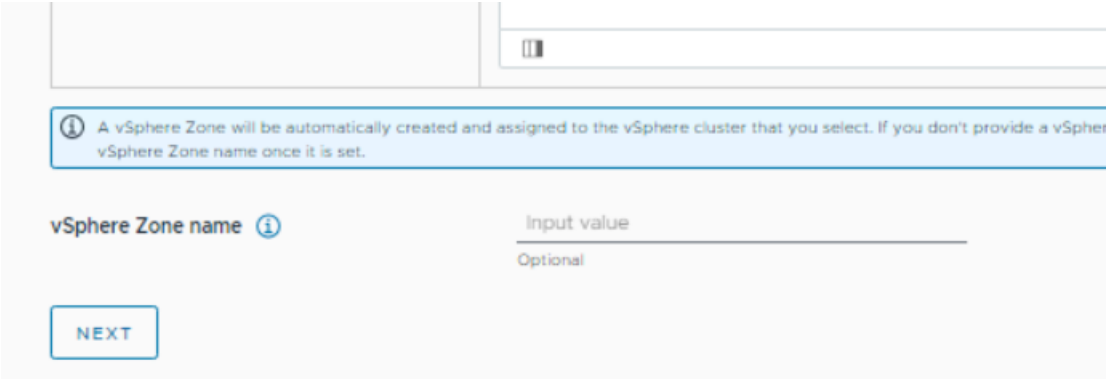

In the Storage window select the appropriate storage policies for each of the component and click on the Next button:

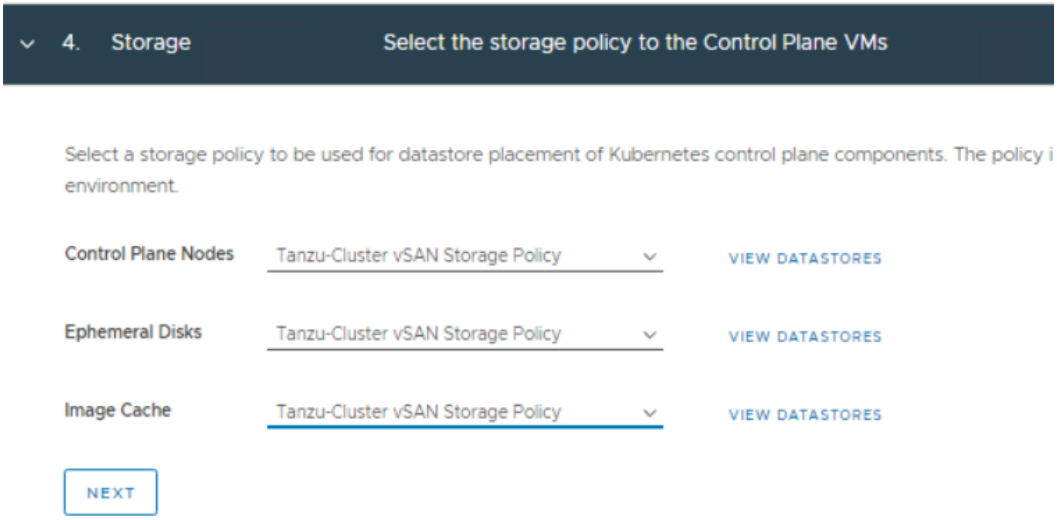

In the Management Network window provide the following parameters then click on the Next button:

- Network Mode – Static
- Network – Select the Management Port Group
- Starting IP Address – 5 IP Address are needed for the Control Plane VMs
- Subnet Mask
- Gateway
- DNS Server
- DNS Search Domain
- NTP Server

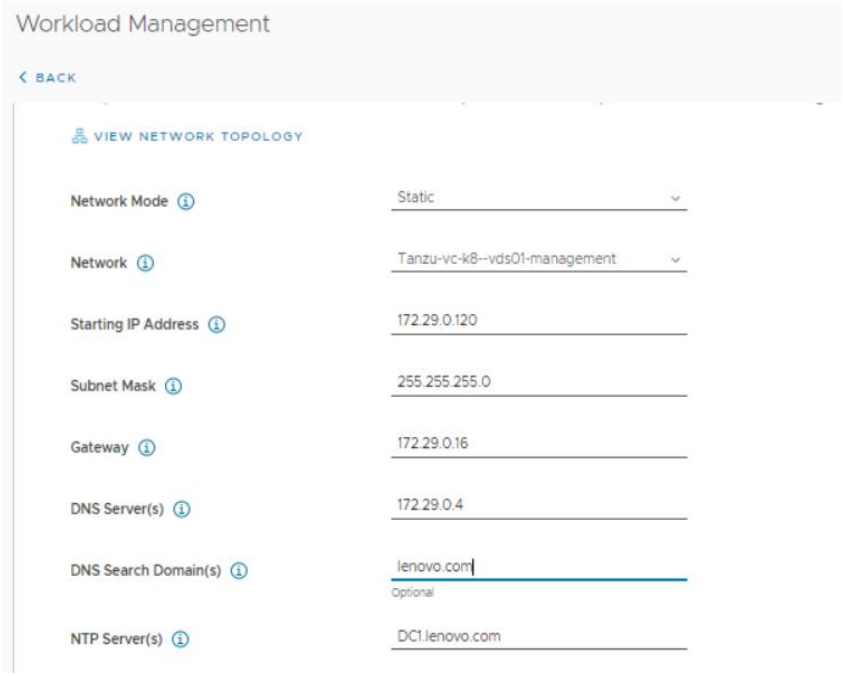

In the Workload Network window:

- Select the vSphere Distributed Switch – Tanzu Cluster VDS
- Select the EDGE Cluster
- DNS Servers – The DNS server must be reachable from the Egress subnet
- Tier-0 Gateway from NSX-T
- NAT Mode – Enabled (choose Disabled for Tanzu Mission Control – see the Note bellow)
- Namespace Network (former POD CIDR) (non-routable) – should only be modified if it already exists in the environment (routable or NATed subnet for Tanzu Mission Control – see the Note bellow)
- Subnet Prefix – the Workload subnets will be carved using this mask
- Service CIDRs (non-routable) – should only be modified if it already exists in the environment
- Ingress CIDRs (routable on the physical switch) – Uplink VLAN configured in the Edge cluster deployment
- Egress CIDRs (routable on the physical switch) – Uplink VLAN configured in the Edge cluster deployment (will not be available if NAT mode is Disabled)

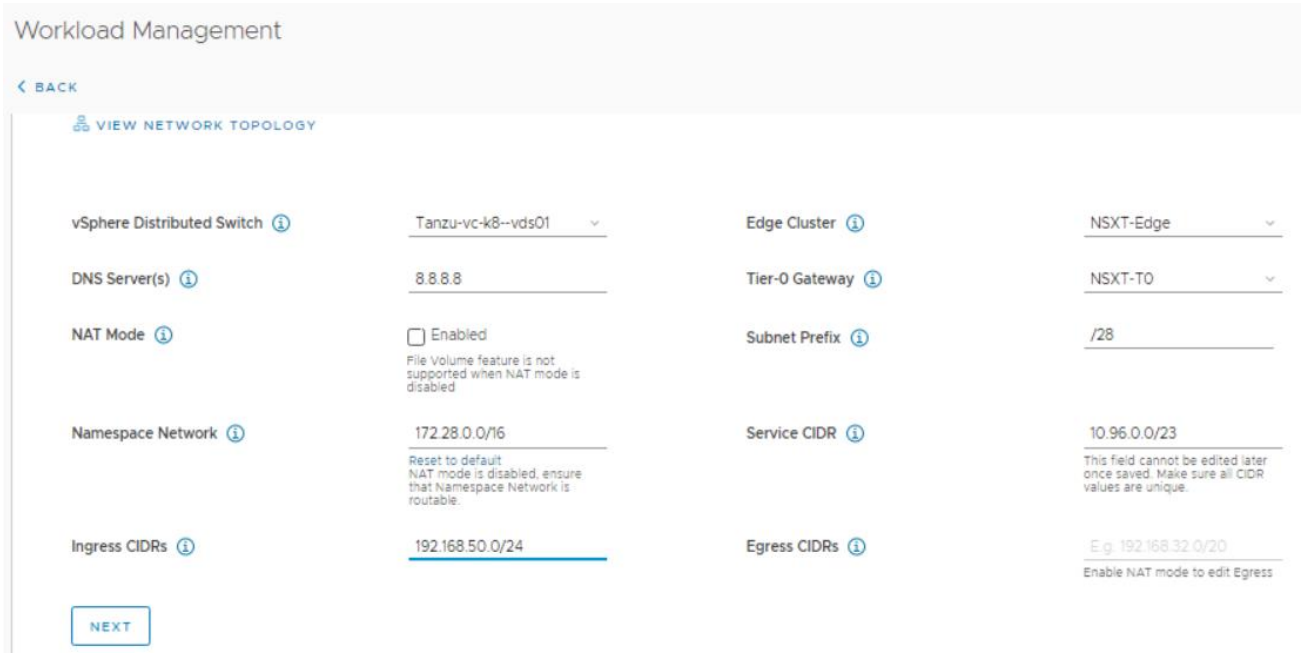

*Note: If Tanzu Mission Control is intended to be deployed, then the following things must be taking into consideration:

- DNS Server for Workload cluster(s) must be set to 8.8.8.8 because any internal DNS server will not be reachable from the VMs that will be deployed by TMC
- NAT Mode must be disabled, for the SupervisorControlPlane VMs deployed via Tanzu Mission Control to be able to pull images from VMware registry, since a private registry cannot be configured
- The Namespace Network must be a routable or a NAT-ed subnet for agent-updater, extension-updater and other Kubernetes services to be able to pull images from VMware registry upon creating clusters via TMC console (the NAT is done on the physical switches that have access to the Internet)

In the Review and Confirm window select the Control Plane Size, specify the API Server DNS Name and click on the FINISH button (check the Export configuration box and save the .zip file):

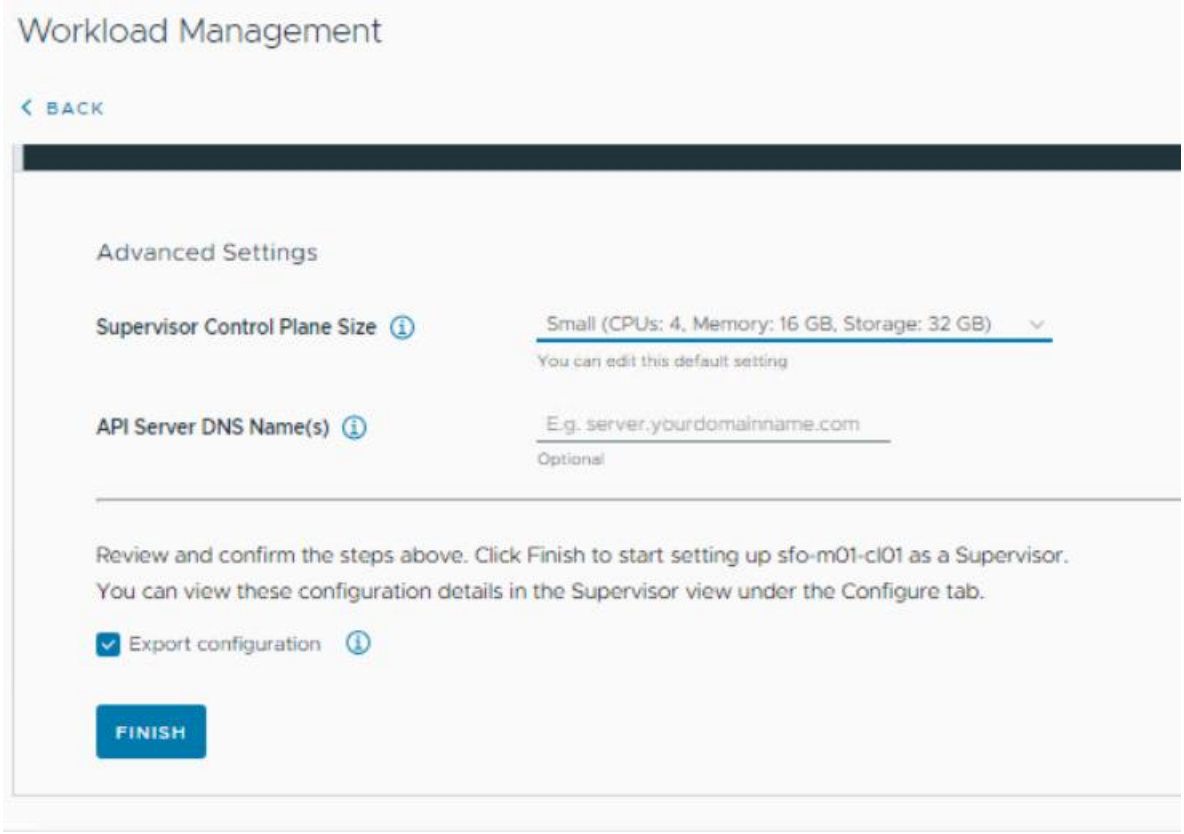

After the Tanzu Kubernetes installation is done successfully:

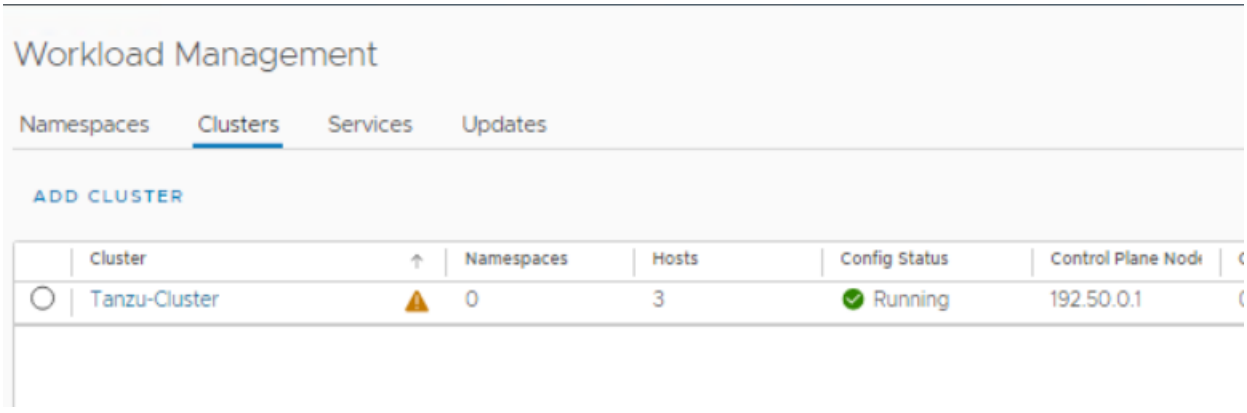

Go to vSphere Client > Hosts and Clusters > Select Tanzu-Cluster > Monitor > Namespaces Overview to verify that Tanzu components have been installed successfully:

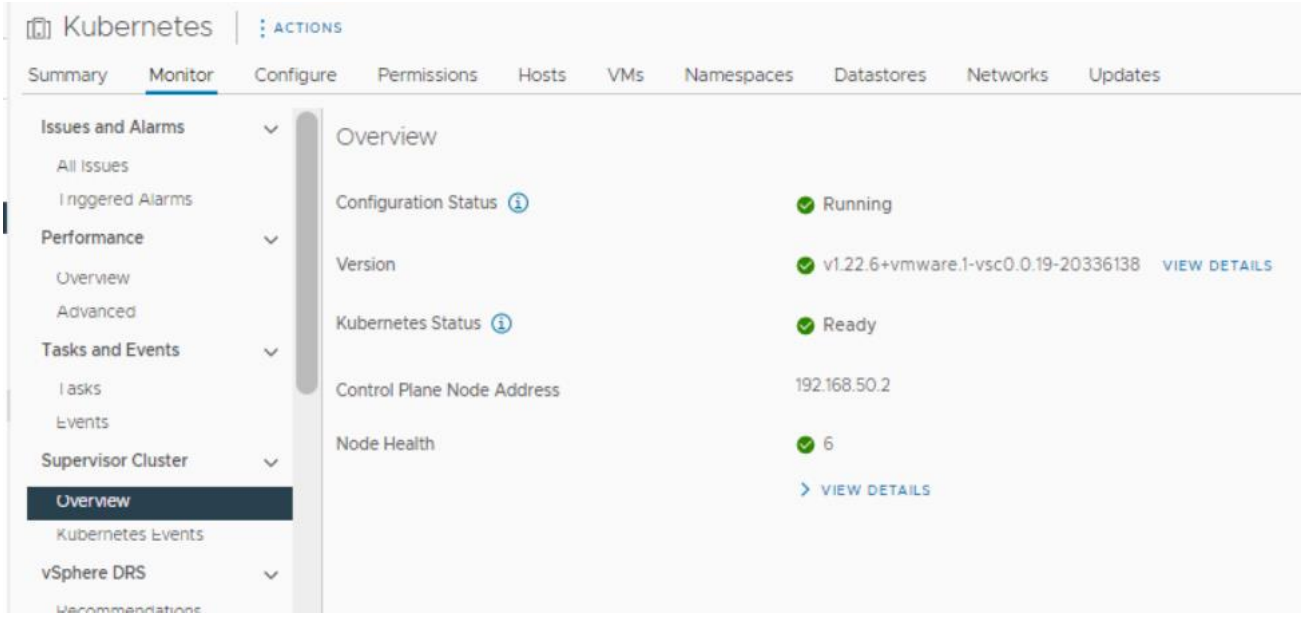

To enable the Embedded Harbor registry, go to vSphere client > Hosts and Clusters > Select Tanzu-Cluster > Configure > Image Registry and click on the ENABLE HARBOR button:

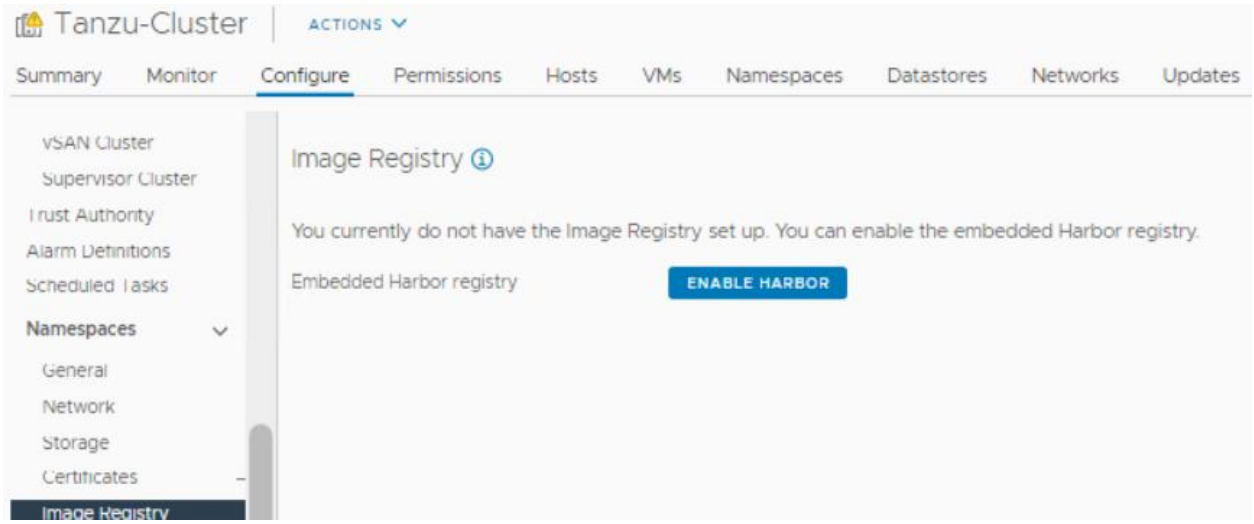

In the Select Storage Policies window select the appropriate storage policy:

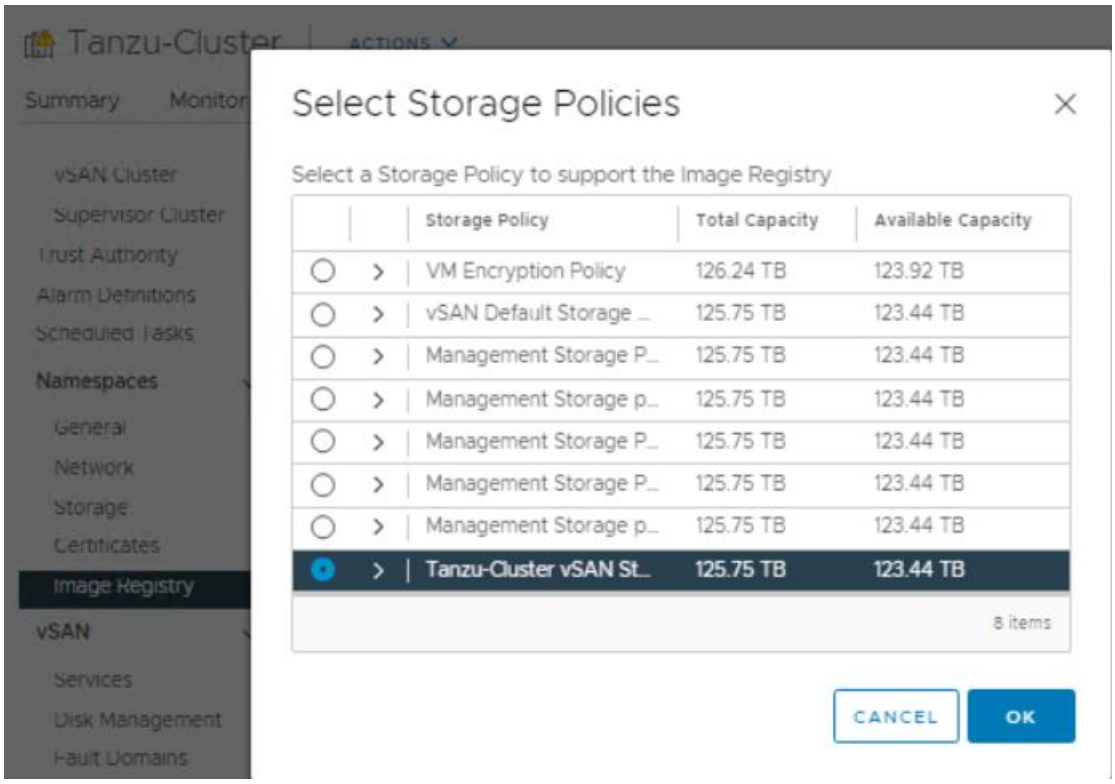

After the Embedded Harbor registry has been successfully enabled:

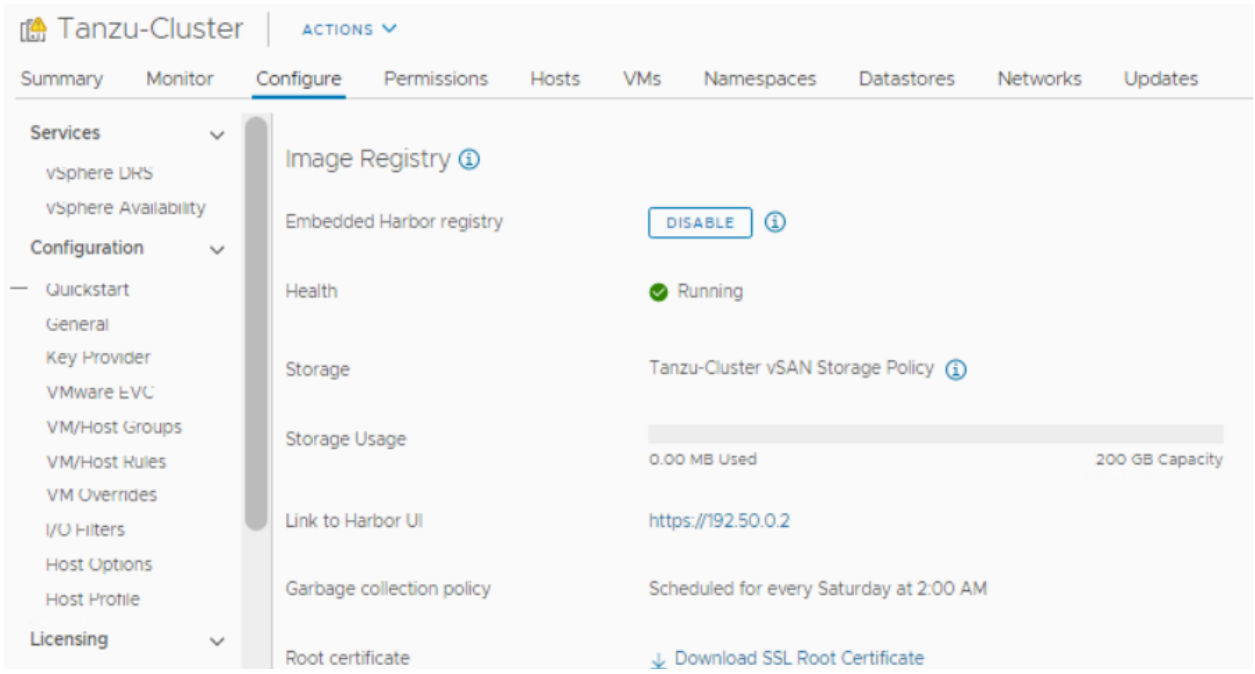

You can login to the Harbor by accessing the Link to Harbor UI:

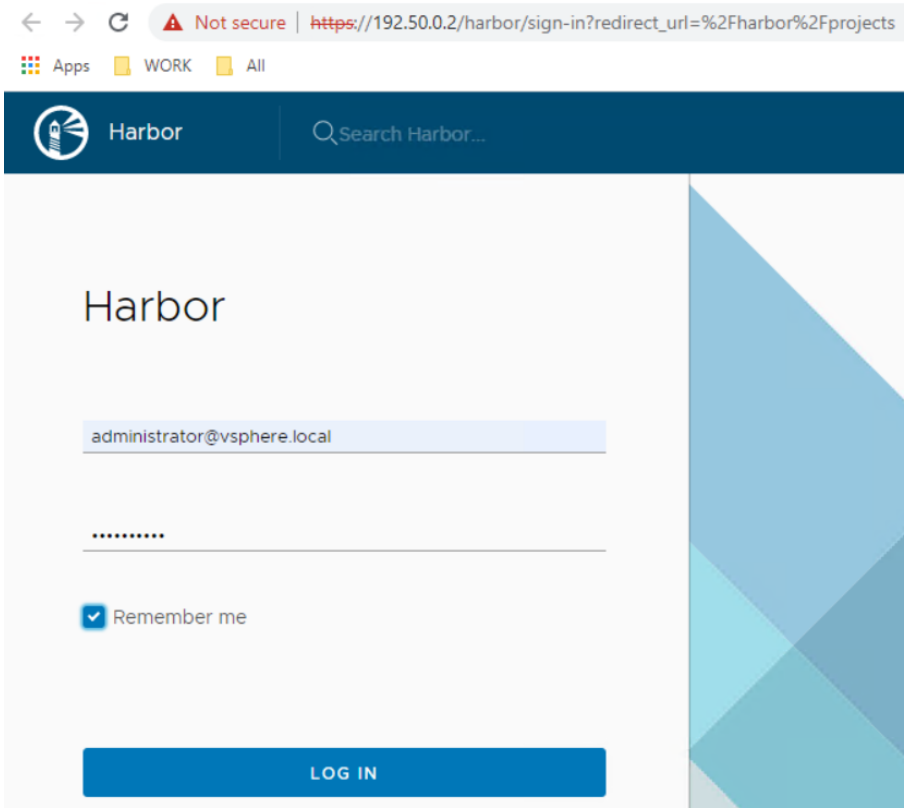

7.3.1.7 VMware Tanzu Kubernetes use case (creating a nginx deployment)

In order to be able to use Tanzu Kubernetes, a Namespace must be created and configured, using vSphere.

In vSphere, go to Menu>Workload Management and click on NEW NAMESPACE button:

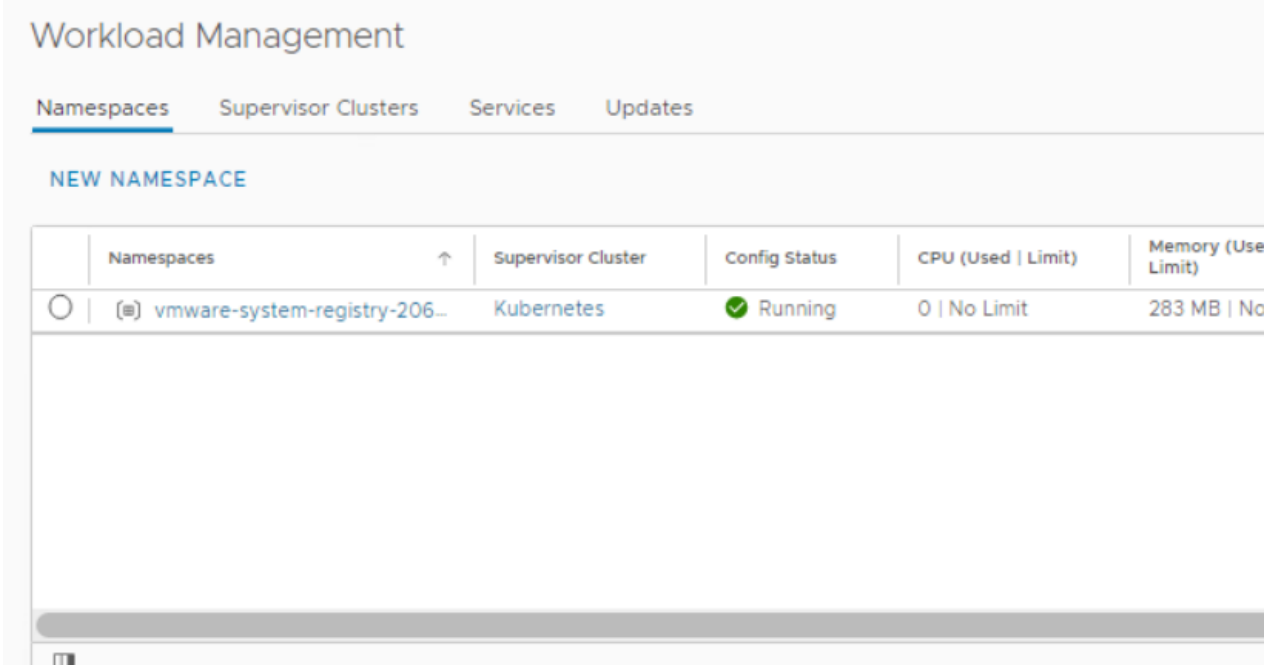

In the Create New Namespace window, select the Kubernetes cluster (in this case Kubernetes), provide a name for the namespace and click on Create button:

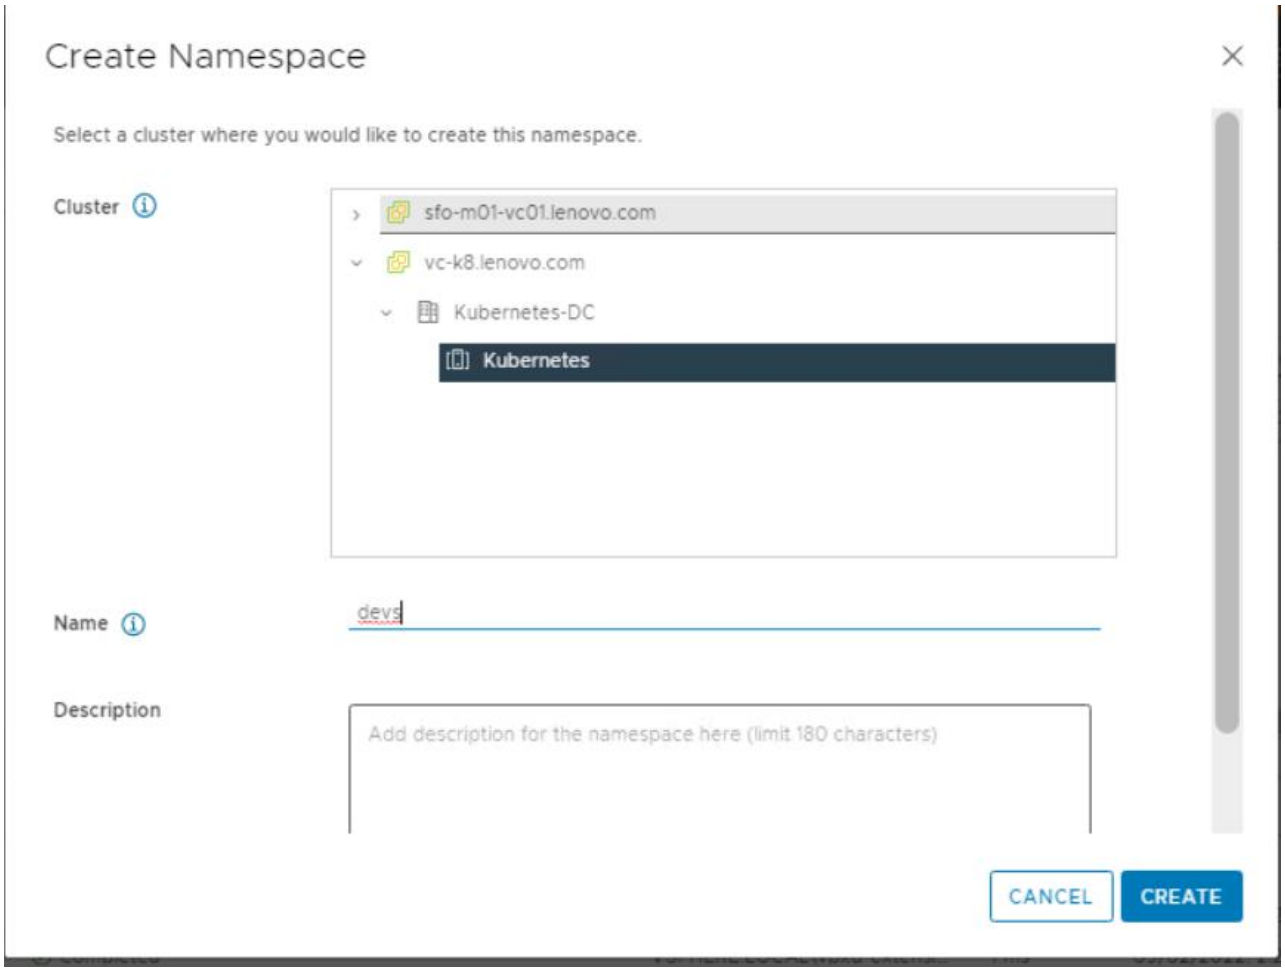

After the Namespace has been created, click on Add Permissions button:

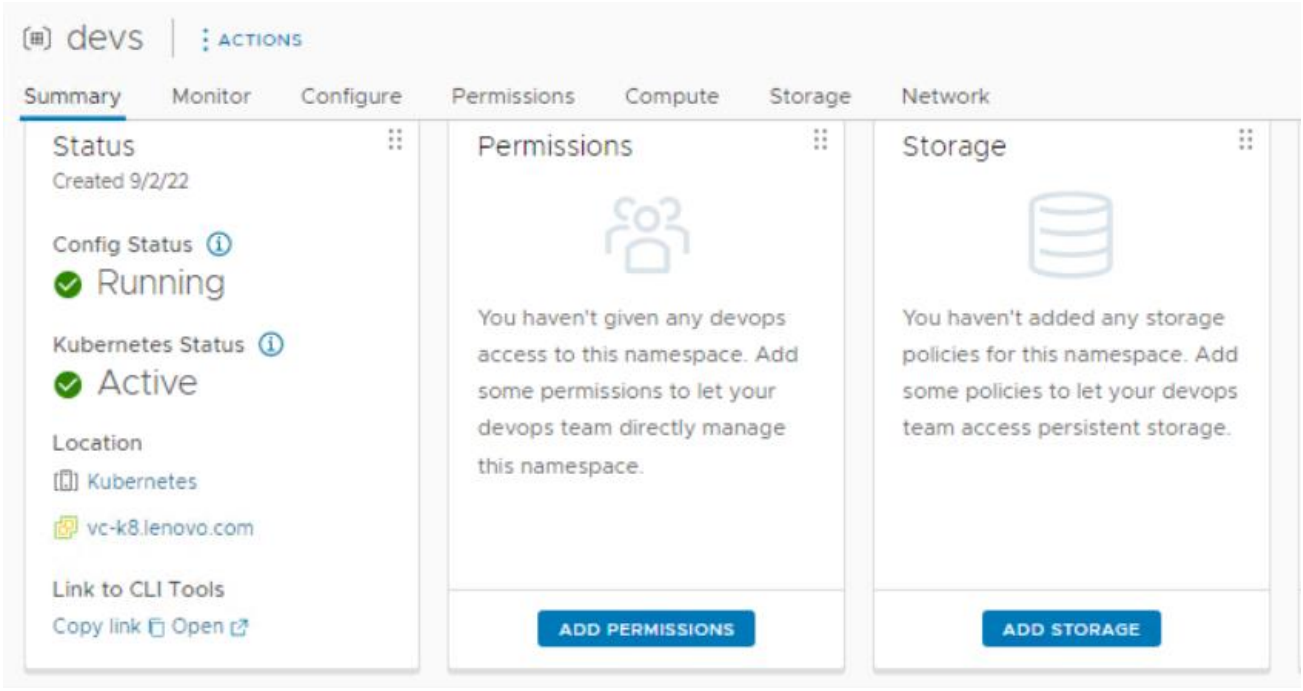

In the Add Permissions window assign the Owner role to Administrator@vsphere.local and click OK:

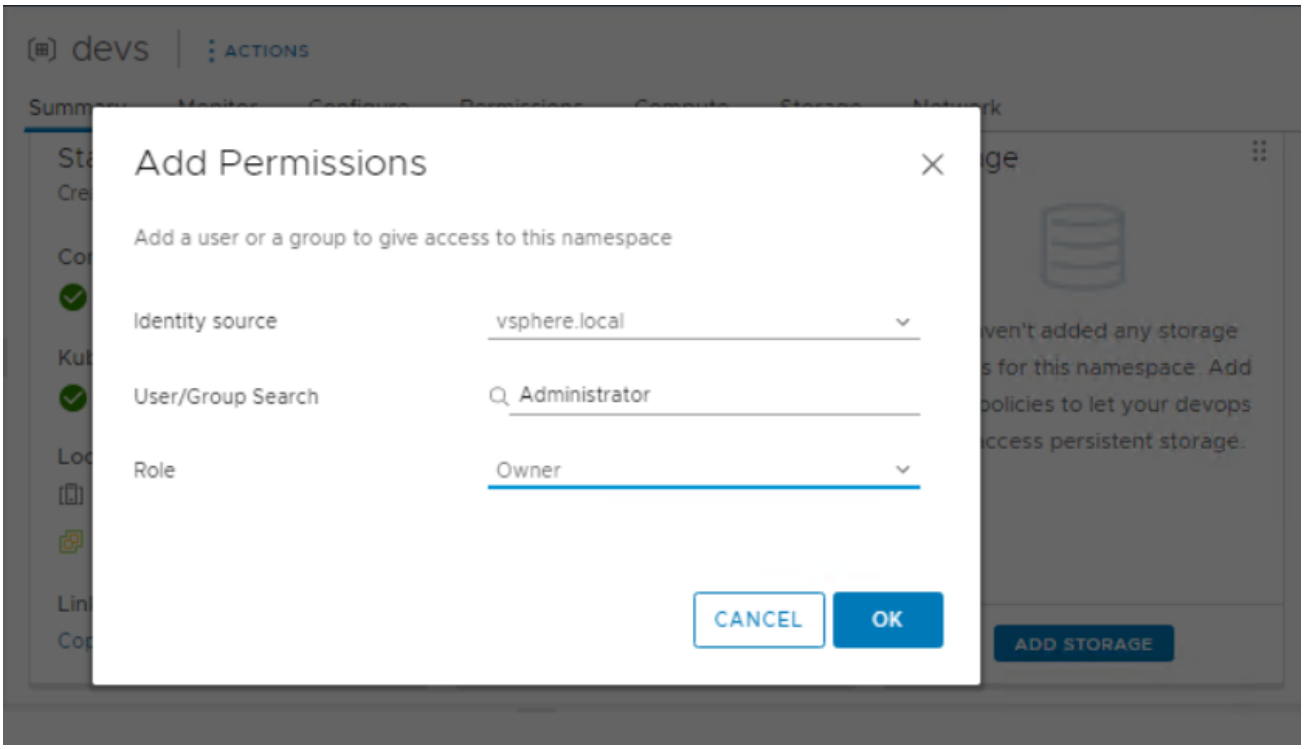

The next step is to add storage for the Namespace, click on Add Storage button:

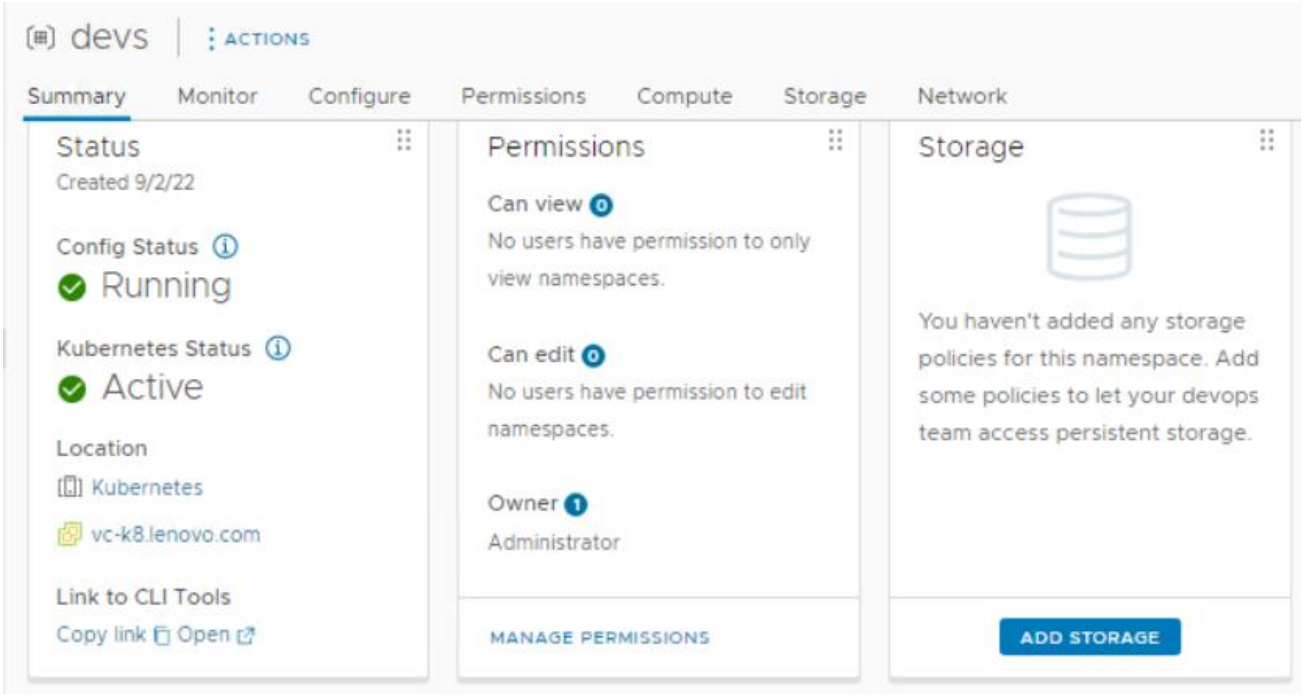

In the Select Storage Policy windows select the Kubernetes vSAN Storage Policy and click OK:

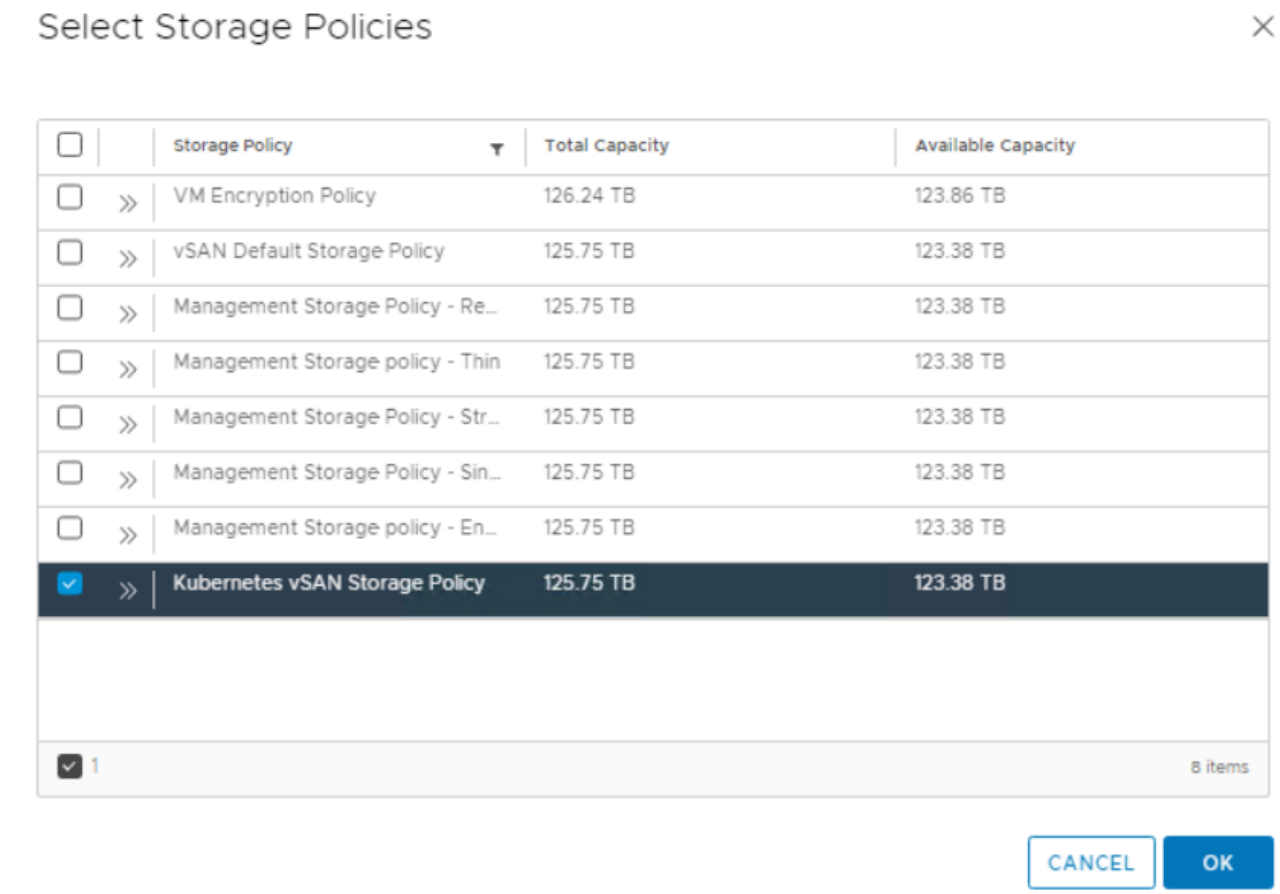

The next step is to add VM Classes for the ‘dev’ namespace, click on VM Service>Associated VM Classes>Add VM Class button:

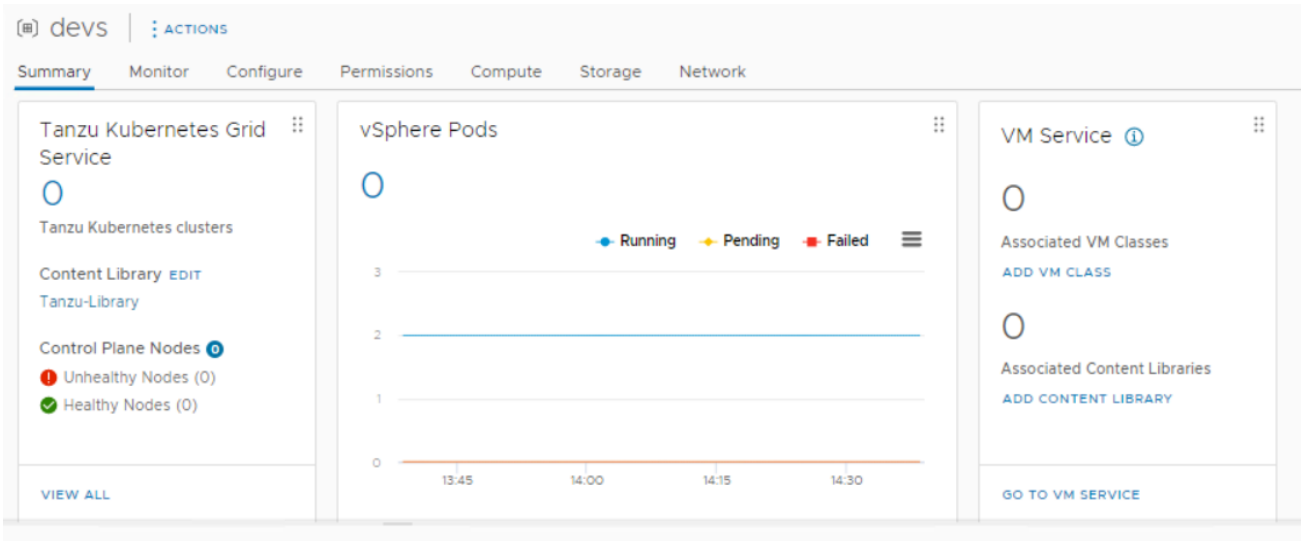

In the Add VM Classes window, select the appropriate VM Classes that best suit your environment and click OK:

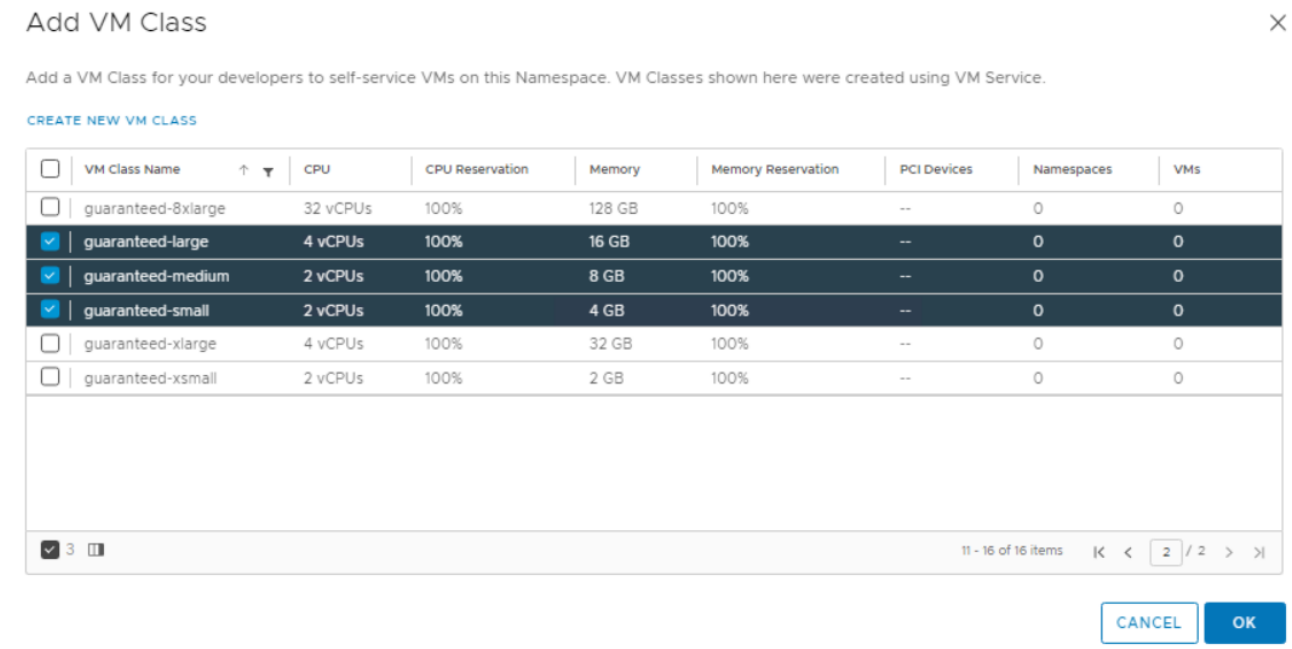

The next step is to add a content library (you can create a new one or use the existing content library), click on VM Service>Associated VM Classes>Add Content Library button:

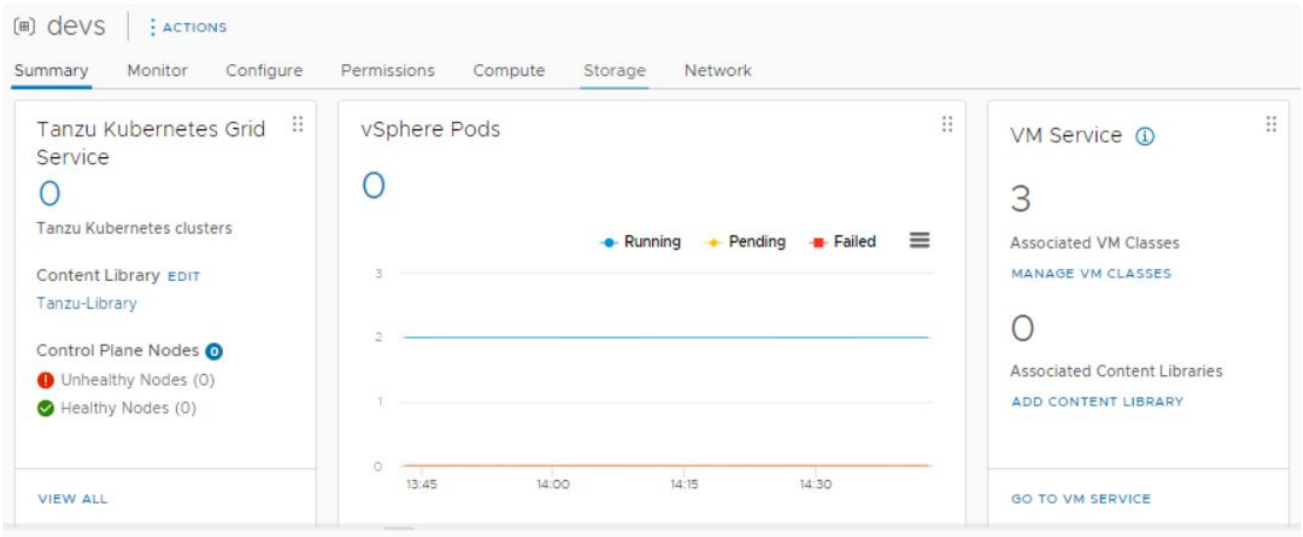

In the Add Content Library window create a new Library or select the existing Library and click OK:

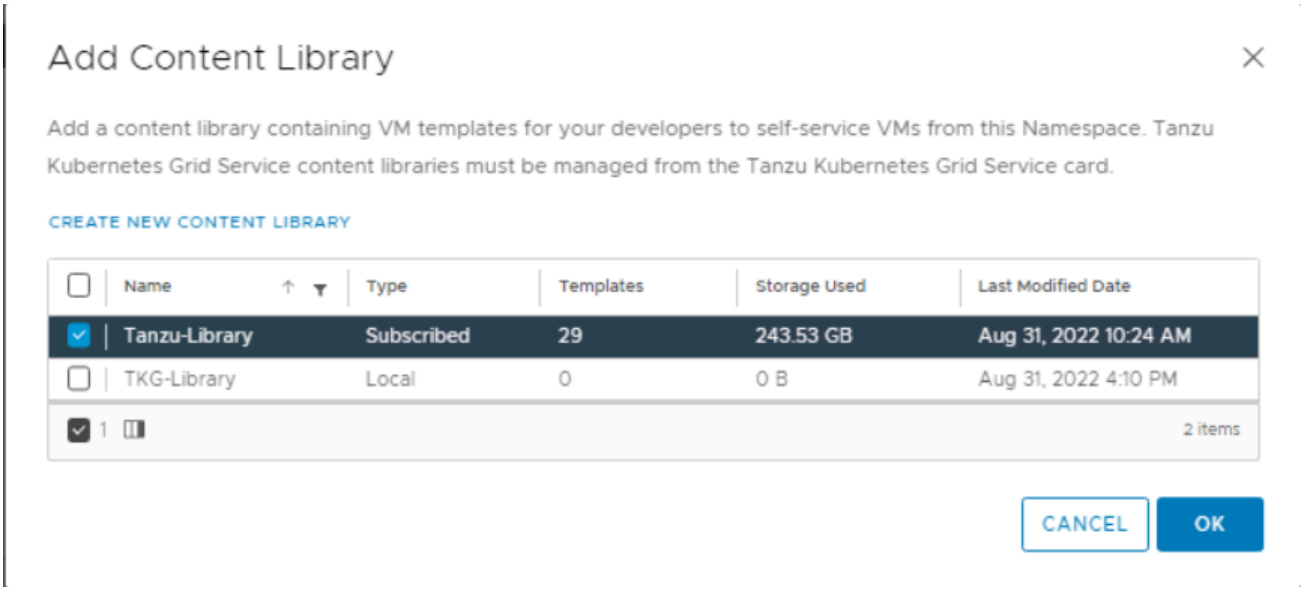

The next step is to install Kubernetes CLI tools, go to Status>Link to CLI tools in the namespace window and click on Open button:

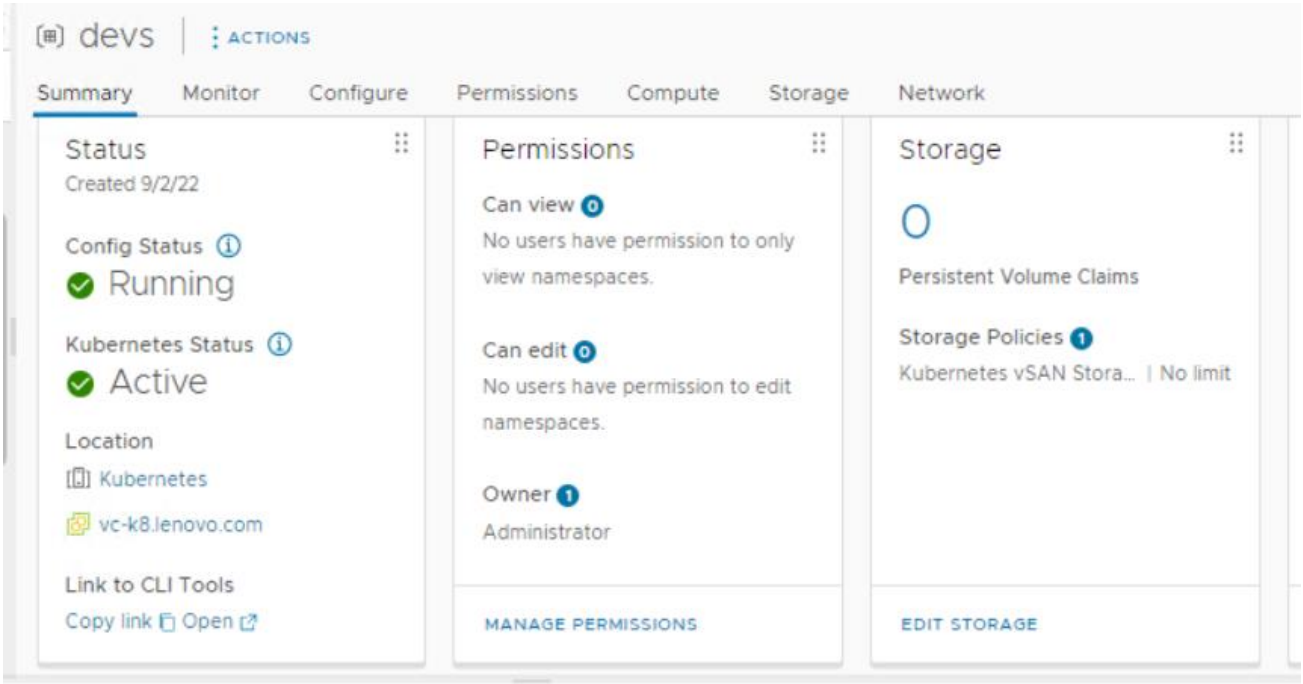

In the newly opened browser window, Select the Operating system (Linux in this case) and click on Download button:

In the same browser window select the Operating System for vSphere Docker Credential Helper (in this case Linux) and click on the Download For Linux button:

Follow the instructions provided in the same browser window to install the CLI tools to the Linux OS.

Login to the Kubernetes server using kubectl, by issuing the following command (on Linux system) and provide the administrator@vsphere.local password:

server=IP of the SupervisorControlPlaneVM Cluster Management IP

```
# kubectl vsphere login --server=172.29.0.120 --vsphere-username
administrator@vsphere.local --insecure-skip-tls-verify

Logged in successfully.

You have access to the following contexts:

   172.29.0.120

   devs

If the context you wish to use is not in this list, you may need to try

logging in again later, or contact your cluster administrator.


To change context, use `kubectl config use-context <workload name>`
```

Choose a context by issuing the following command:

```
# kubectl config use-context 172.29.0.120

Switched to context "172.29.0.120".
```

The next step is to add the Nginx image to the Embedded Harbor. In order to do that we have to login to the Harbor using ‘docker’. The Harbor certificate must be uploaded to the /etc/docker/certs.d/<IP of the Harbor> directory.

Login to the Harbor in order to download the harbor certificate:

Go to Projects>devs>Repositories and click on Registry Certificate button:

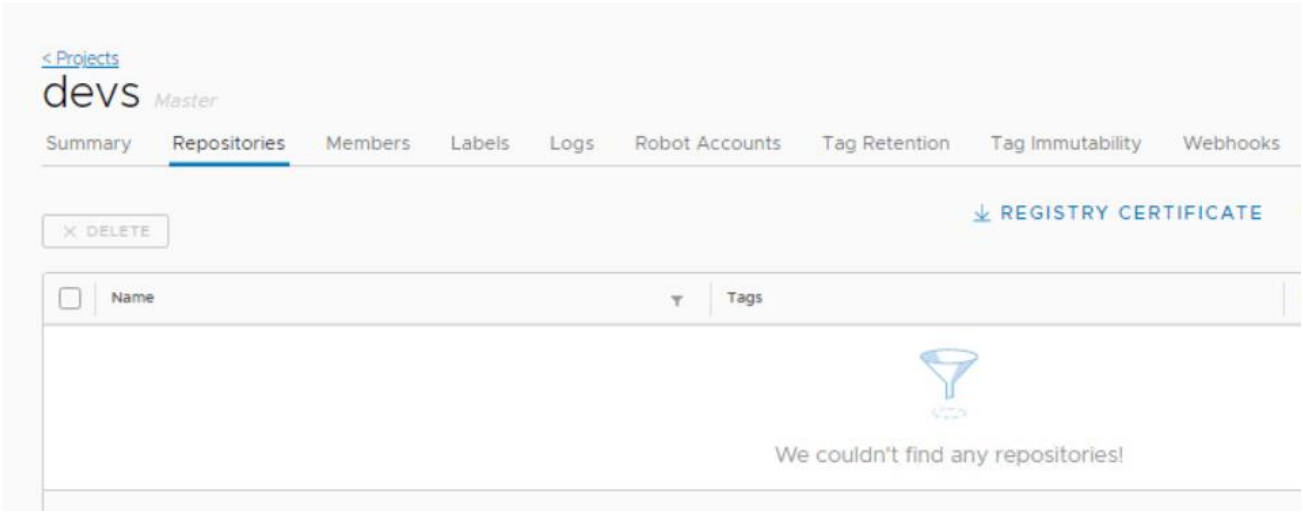

Create a directory under the /etc/docker/certs.d/ with the same name as the Harbor IP:

```
# mkdir /etc/docker/certs.d/192.50.0.4
```

Create a new file called ca.crt:

```
# touch /etc/docker/certs.d/192.50.0.4/ca.crt
```

Copy the harbor certificate details to the newly created ca.crt file using a text editor.

Start the docker service:

```
# systemctl start docker
```

Login to the harbor using the ‘docker-credential-vsphere’ command:

```
# docker login https://192.168.50.4/devs
Username: administrator@vsphere.local
Password: INFO[0012] Fetched username and password
INFO[0012] Fetched auth token
INFO[0012] Saved auth token
```

Pull the ‘nginx’ image using ‘docker’ command:

```
# docker pull nginx
Using default tag: latest
latest: Pulling from library/nginx
7a6db449b51b: Pull complete
ca1981974b58: Pull complete
d4019c921e20: Pull complete
7cb804d746d4: Pull complete
e7a561826262: Pull complete
7247f6e5c182: Pull complete
Digest: sha256:b95a99feebf7797479e0c5eb5ec0bdfa5d9f504bc94da550c2f58e839ea6914f
Status: Downloaded newer image for nginx:latest
docker.io/library/nginx:latest
# docker images
REPOSITORY                  TAG             IMAGE ID      CREATED        SIZE
nginx                       latest          2b7d6430f78d  10 days ago    142MB
```

The image must be tagged in order to be pushed to the Harbor:

```
# docker tag 2b7d6430f78d 192.50.0.4/devs/nginx:latest
# docker images


REPOSITORY                  TAG             IMAGE ID      CREATED        SIZE
```

```
192.50.0.4/devs/nginx       latest           2b7d6430f78d   10 days ago     142MB
nginx                       latest           2b7d6430f78d   10 days ago     142MB
```

Push the image to the Harbor:

```
# docker push 192.50.0.4/devs/nginx:latest
```

Verify that the image has been pushed to the Harbor:

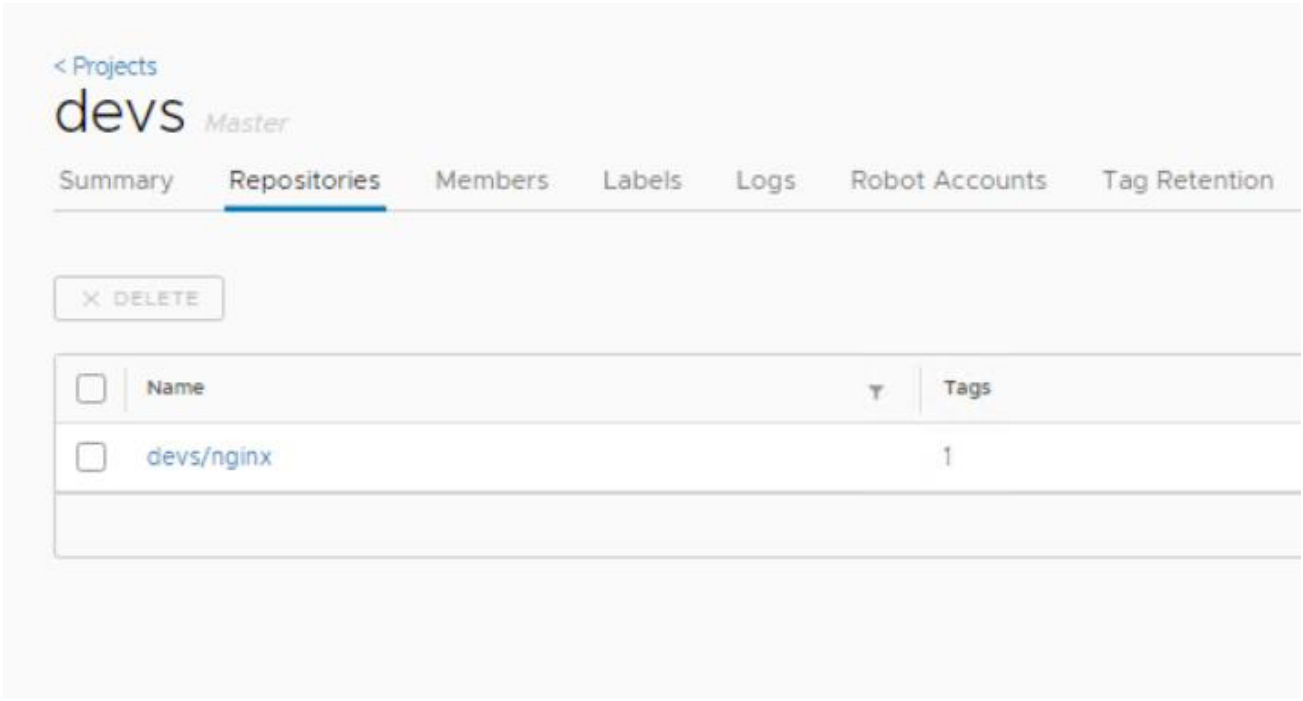

Create a deployment using the image:

```
# kubectl create deployment webserver --image 192.50.0.4/devs/nginx:latest -n
devs
deployment.apps/webserver created
```

Verify that the deployment has been created and is Ready:

```
# kubectl get deployment -n devs
NAME       READY   UP-TO-DATE   AVAILABLE   AGE
webserver   1/1     1            1           86s
```

Scale up the webserver deployment by creating a replica:

```
# kubectl scale deployment webserver --replicas 2 -n devs
deployment.apps/webserver scaled
```

```
# kubectl get deployment -n devs
NAME       READY   UP-TO-DATE   AVAILABLE   AGE
webserver   2/2     2            2           4m23s
```

We now have 2 webservers pods running

```
# kubectl get pods -n devs
NAME                         READY   STATUS    RESTARTS   AGE
webserver-85fb564965-jx7nc   1/1     Running   0          6m47s
webserver-85fb564965-zzbgv   1/1     Running   0          3m39s
```

The webserver must be exposed so it can be accessible:

```
# kubectl expose deployment webserver --port=80 --type=LoadBalancer -n devs
```

Check if the webserver is running and take note of the ‘external’ IP

```
# kubectl get services -n devs
NAME       TYPE          CLUSTER-IP    EXTERNAL-IP   PORT(S)       AGE
webserver   LoadBalancer   10.96.1.252   192.50.0.5    80:31145/TCP   55s
```

Verify that the webserver is running by opening a browser and use the External-IP:

7.3.2 Deploy Tanzu Mission Control

Login to VMware Cloud Services and launch VMware Tanzu Mission Control console by clicking on Launch Service:

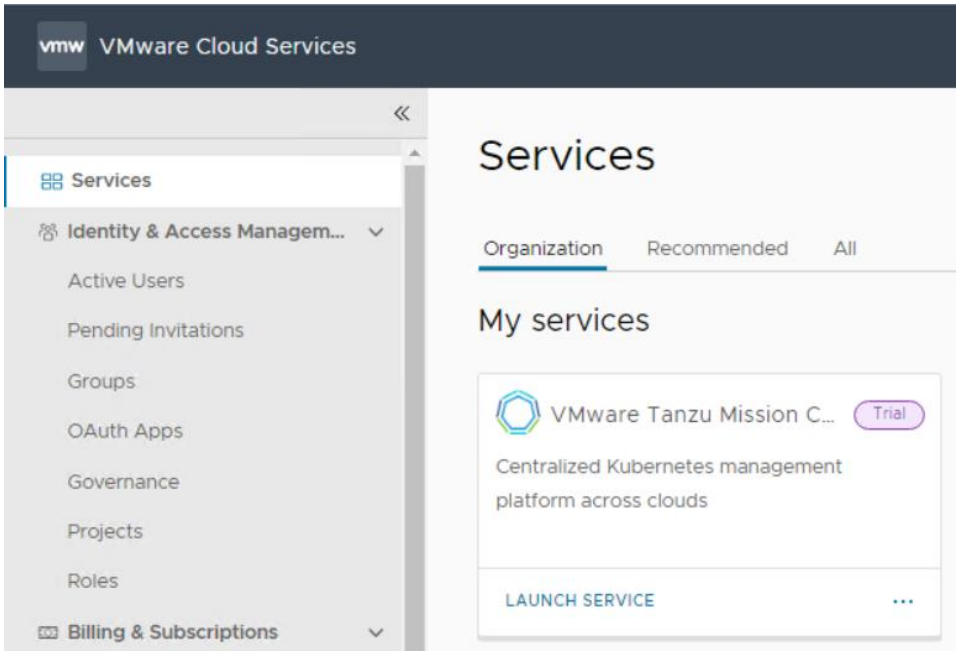

In TMC Console go to Administration>Management clusters>Register Management Cluster>vSphere with Tanzu:

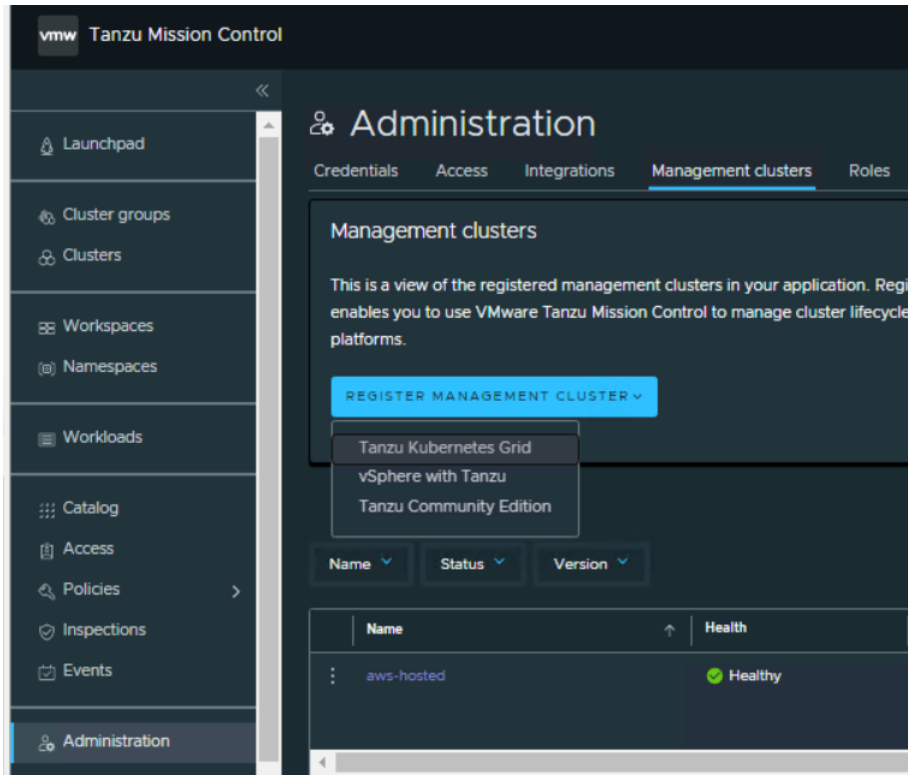

In the Register management cluster window provide a name for the management cluster and select a group then click Next:

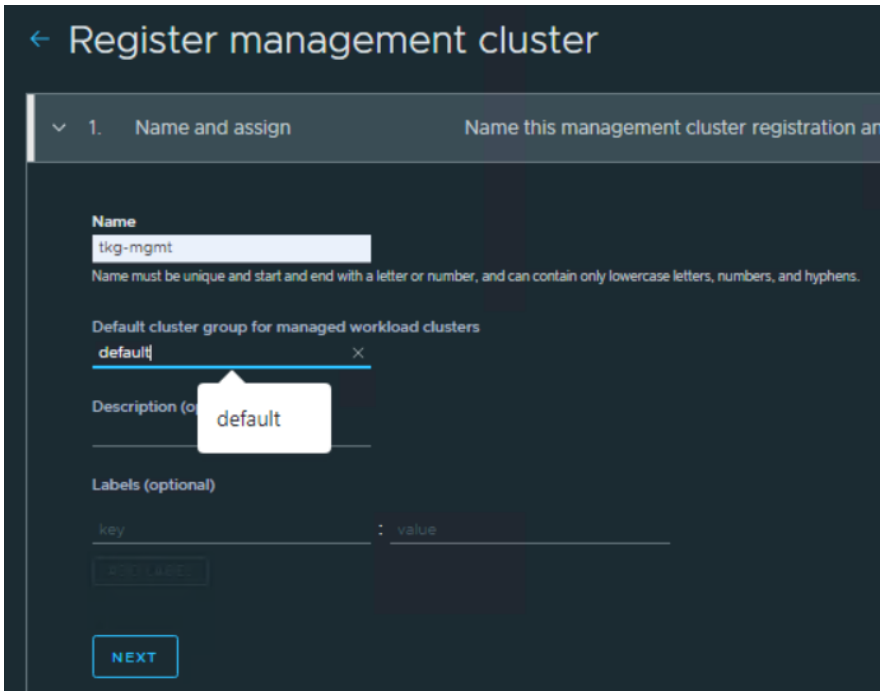

In the Proxy configuration window click Next:

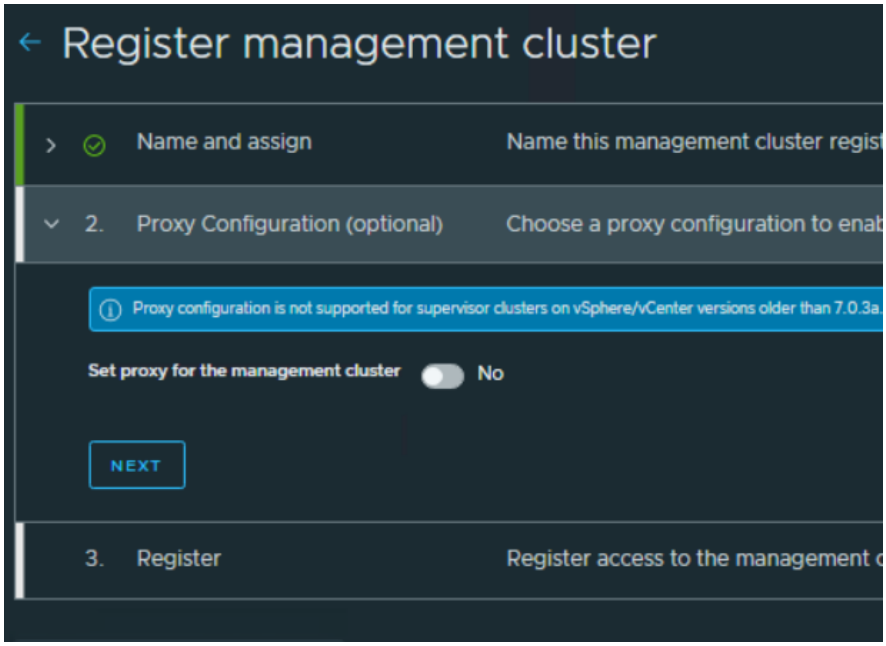

In the Register window copy the registration URL provided:

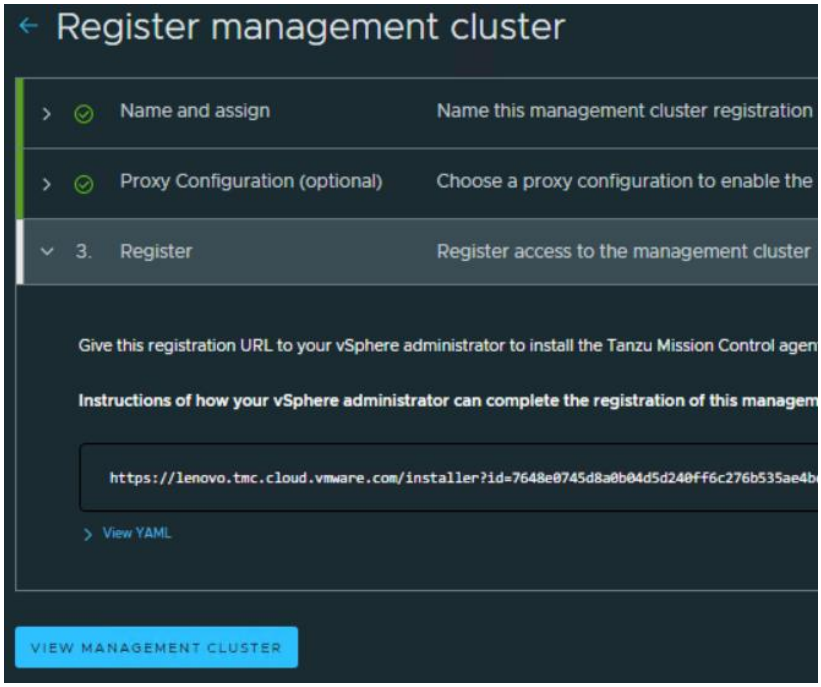

Use the registration URL in the vSphere – click on the Workload Management>Supervisor>Tanzu Mission Control, copy the link in the Registration URL window and click on the Register button:

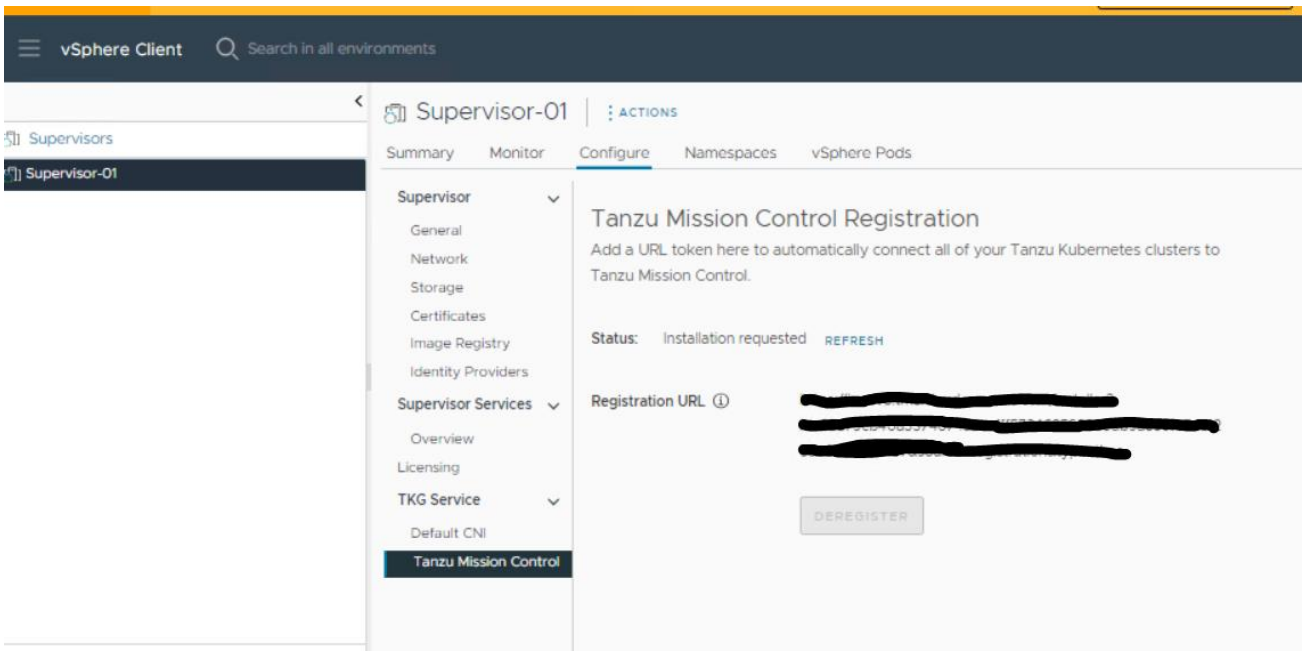

After the registration is successfully done:

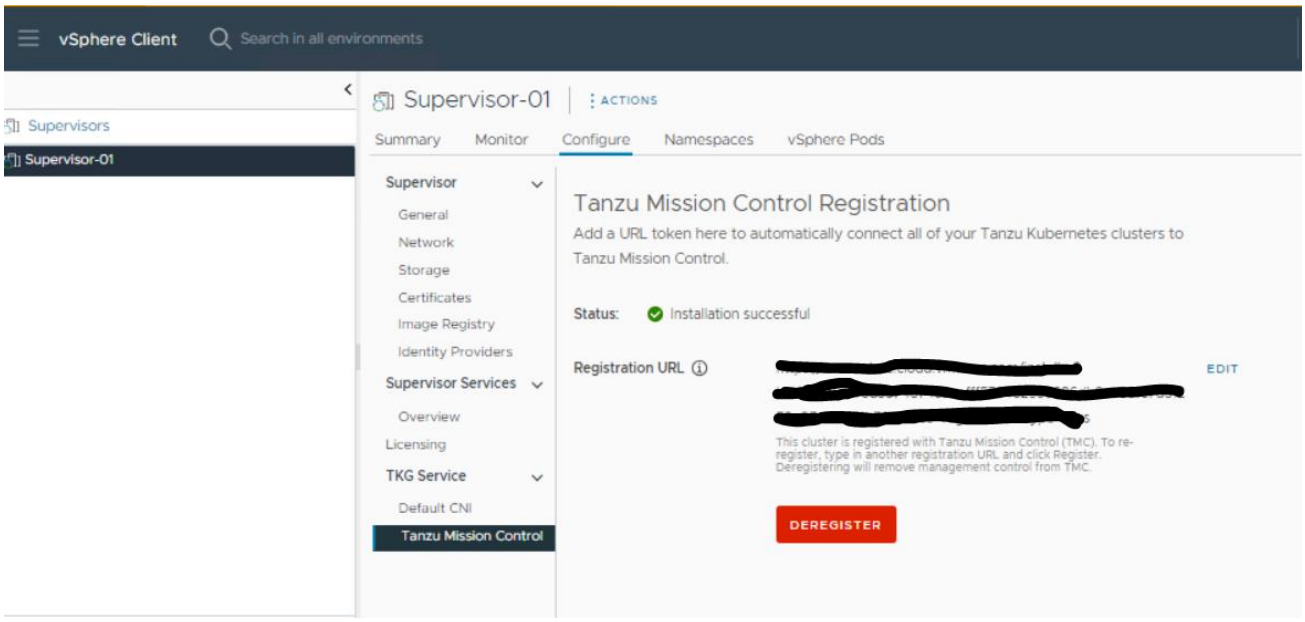

Check in the TMC console that the management cluster status is Ready and Healthy:

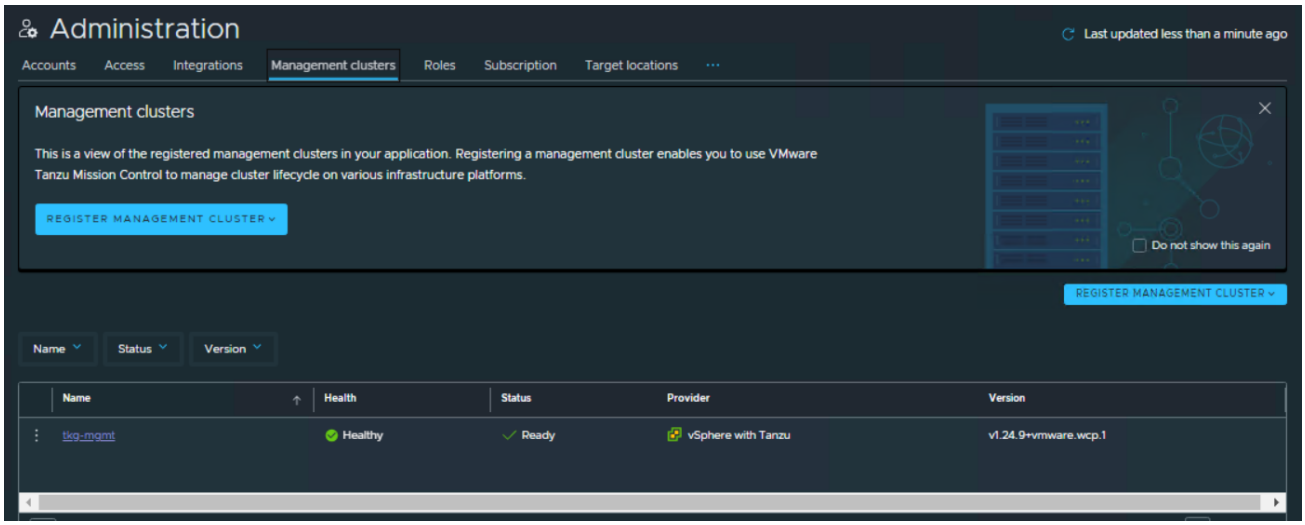

To deploy a vSphere with Tanzu Cluster go to Clusters:

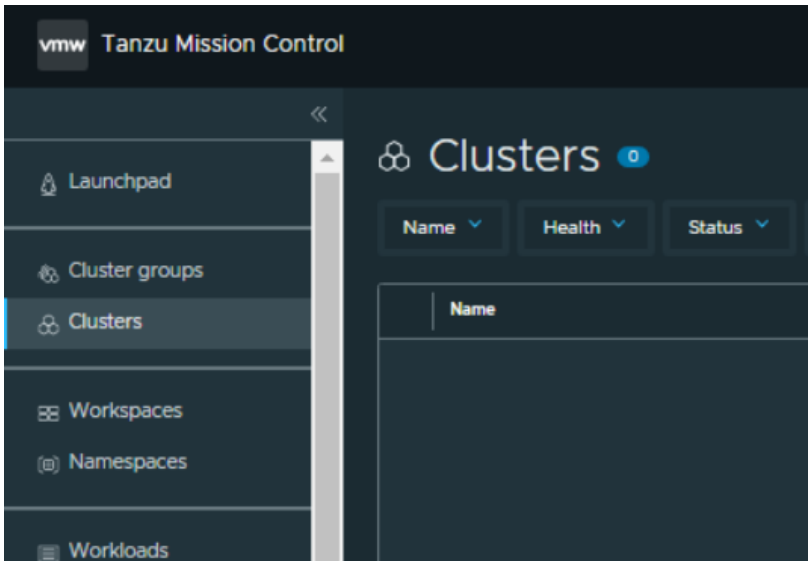

Click on Create Cluster and select Create Tanzu Kubernetes Grid cluster:

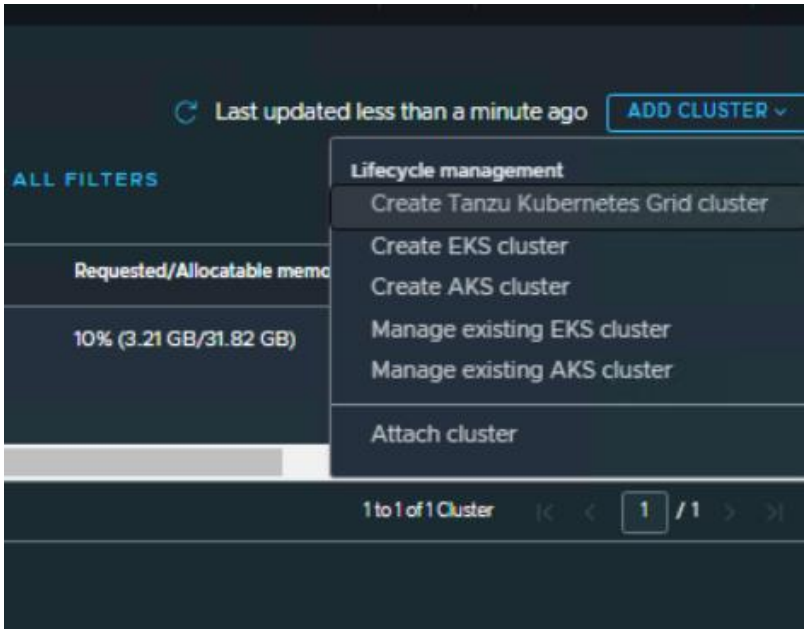

In the Create cluster window select the management cluster and click on ‘Continue to create cluster’:

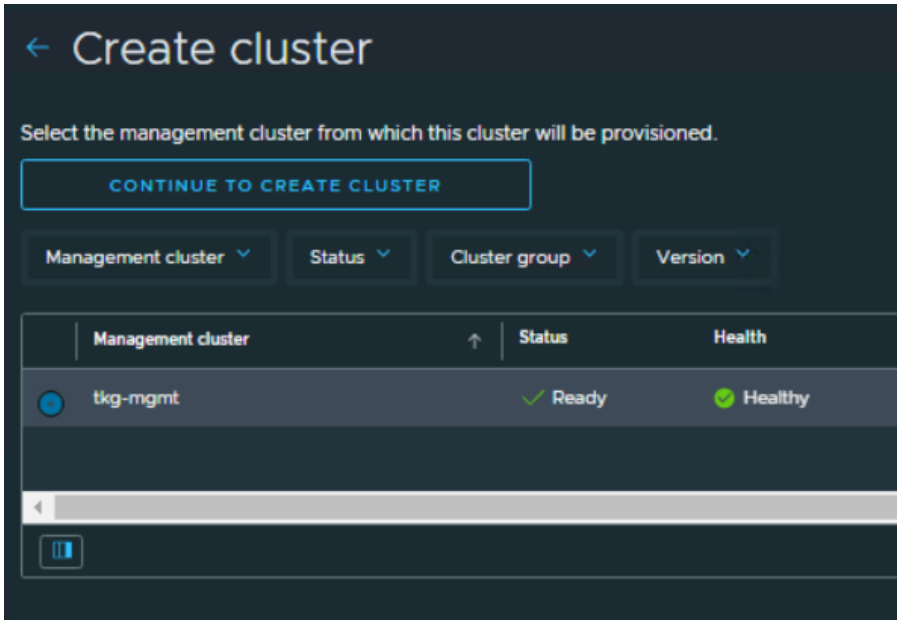

Select a Provisioner (this should be a vSphere with Tanzu Namespace created in vSphere) and click Next:

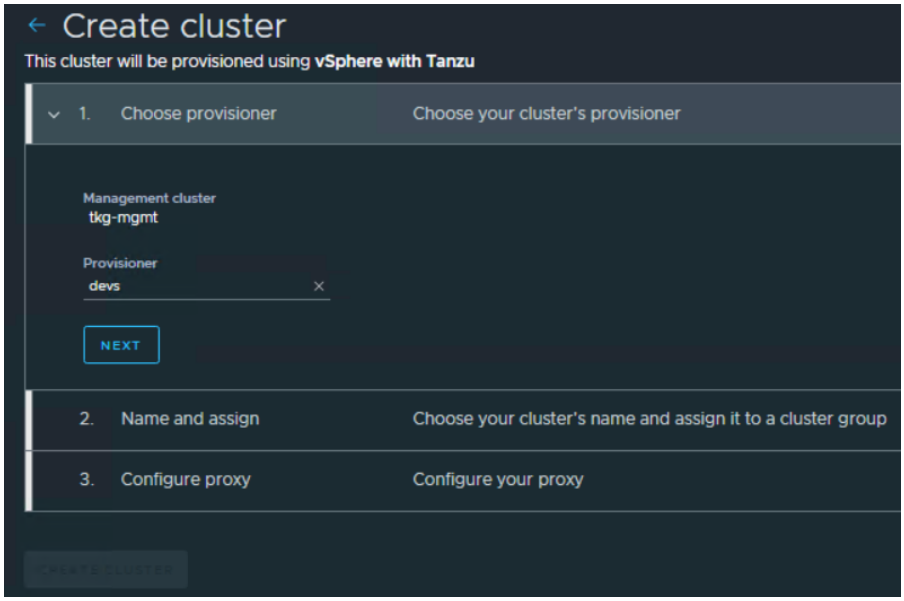

Specify a name for the cluster and select a TMC group and click Next:

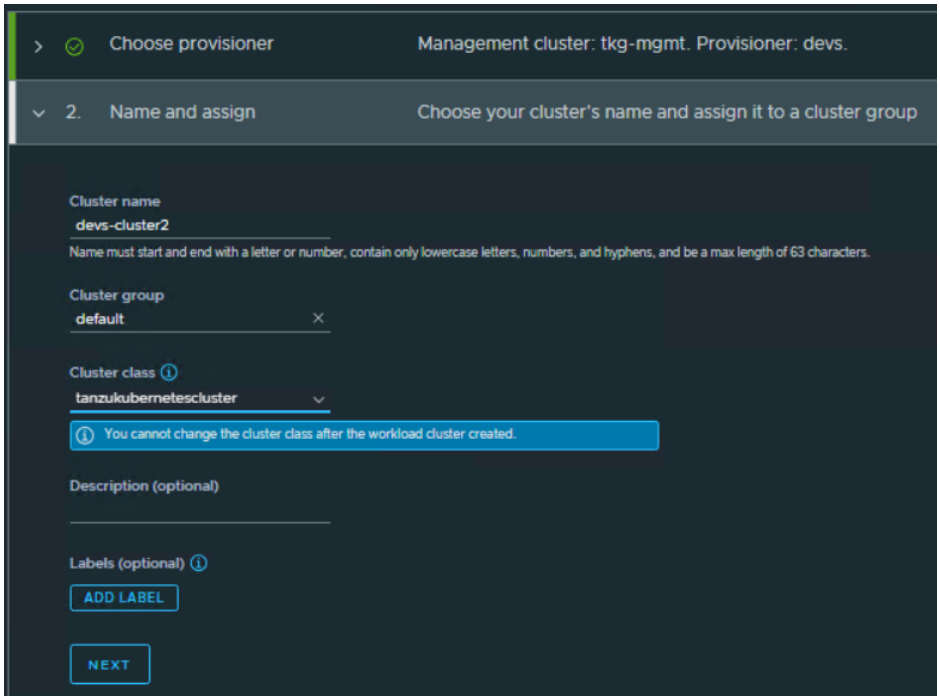

Configure a proxy (optional)

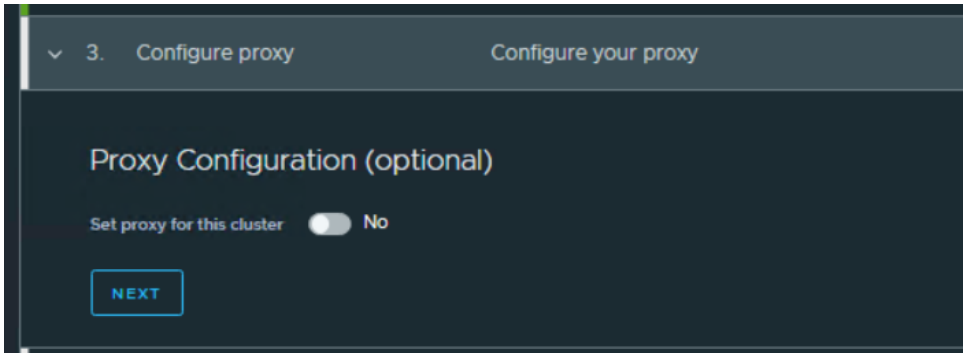

Configure network and storage policy

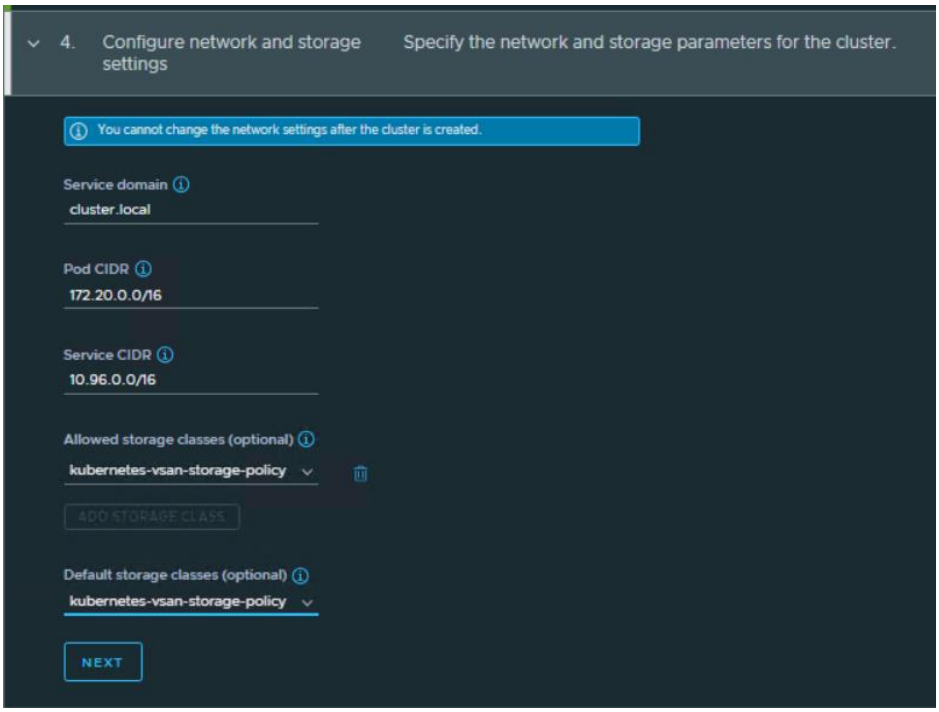

Select between single node or highly available control plane VMs then click Next:

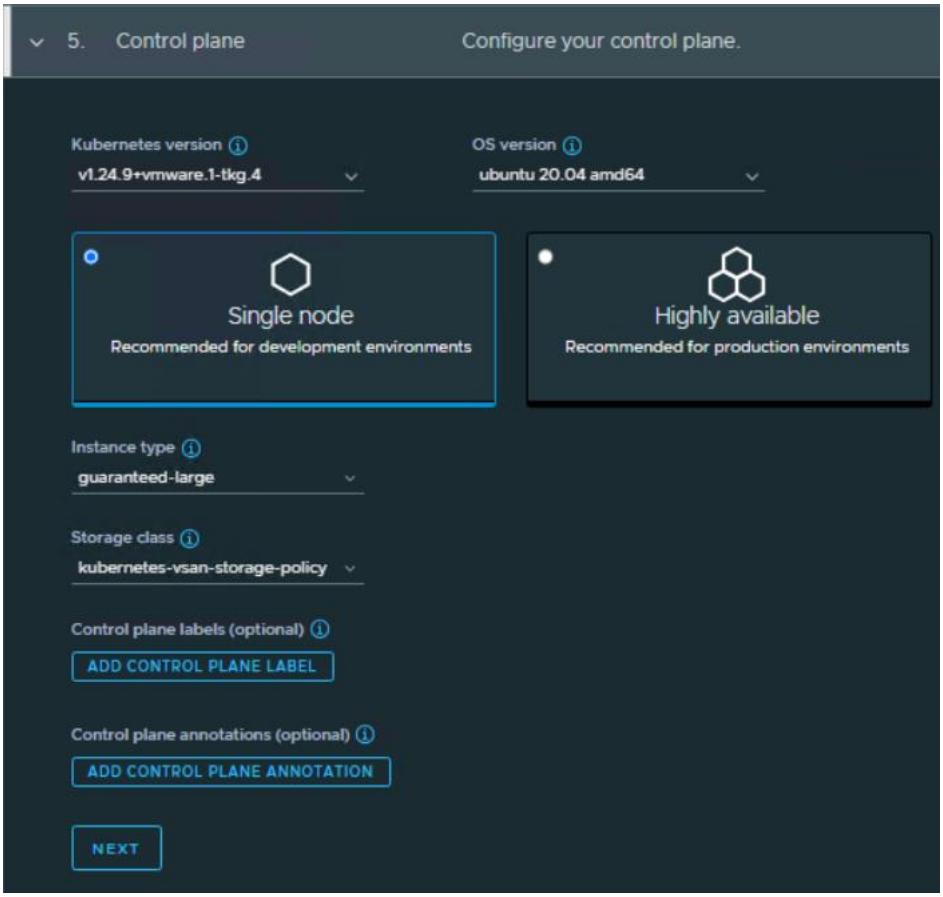

Configure default volumes:

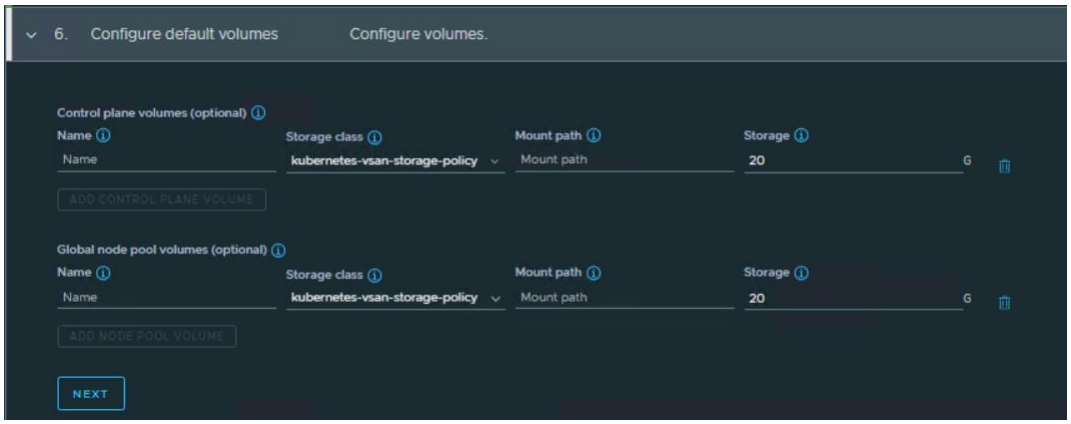

Configure node pool:

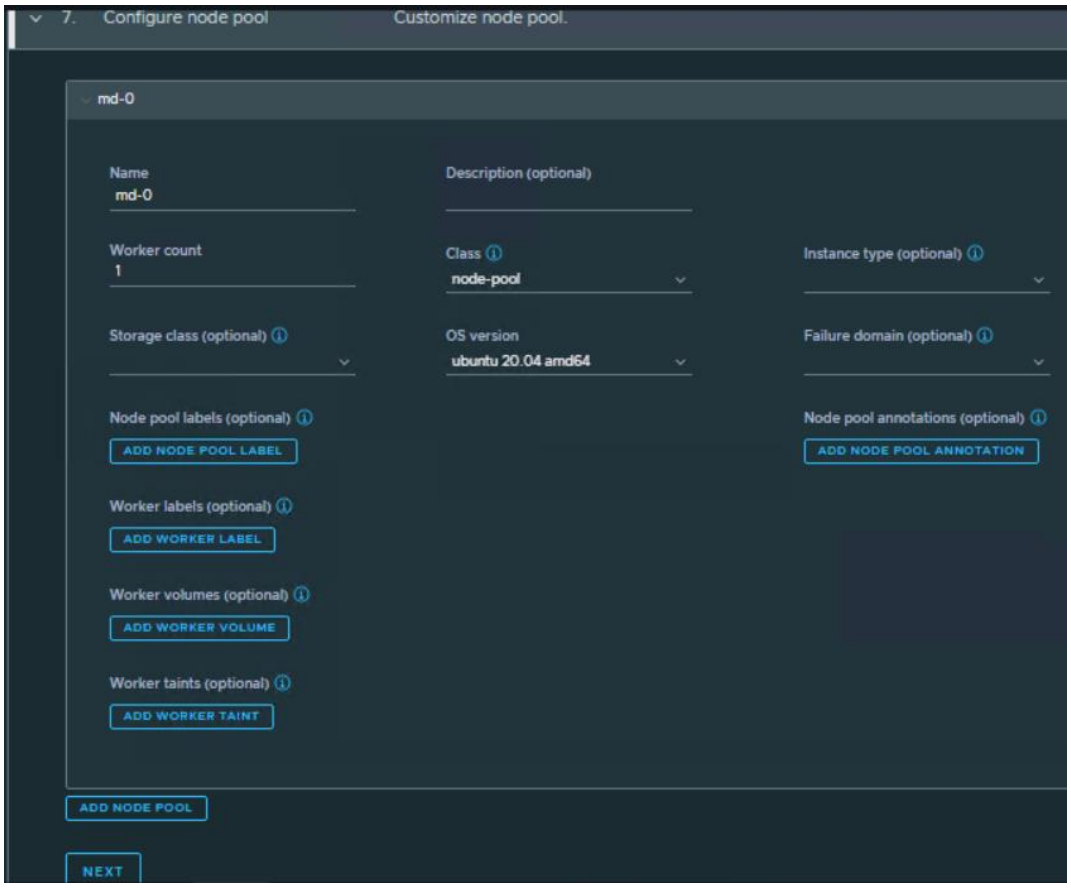

Configure additional cluster configurations then click on CREATE CLUSTER button:

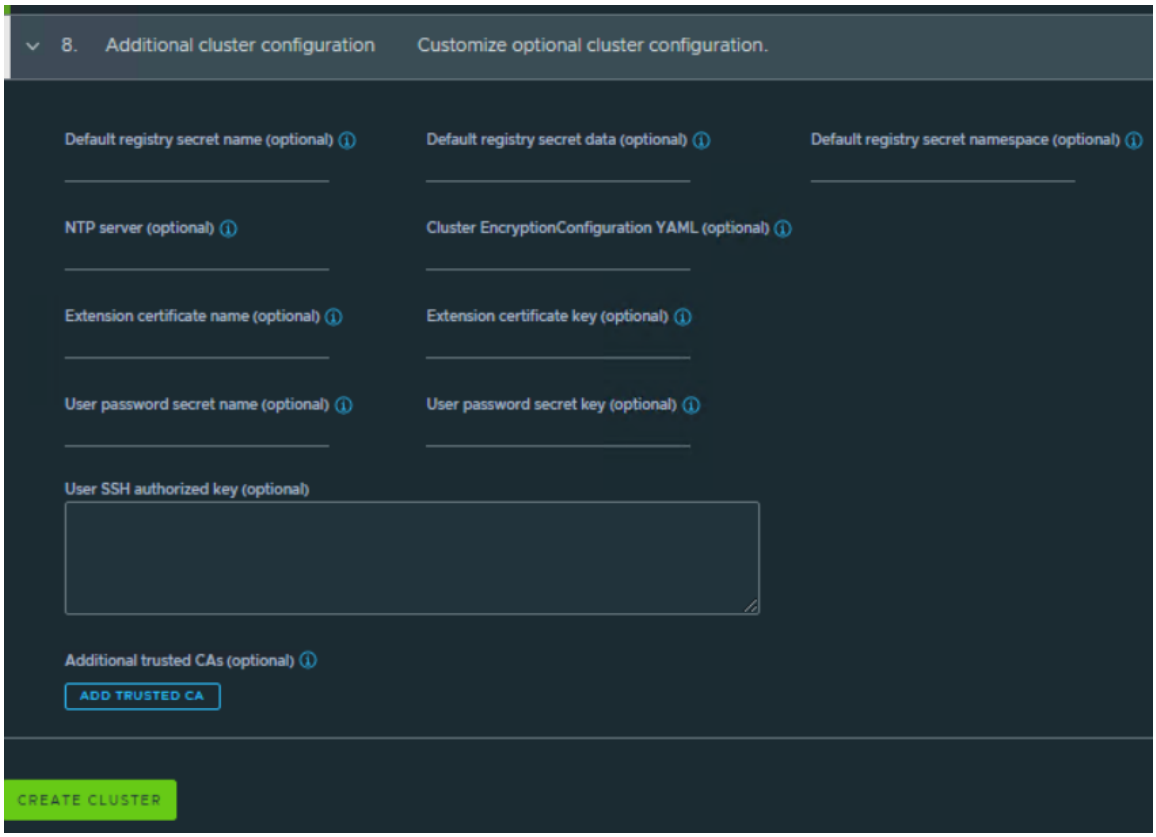

After the Cluster has been successfully deployed you can check the status:

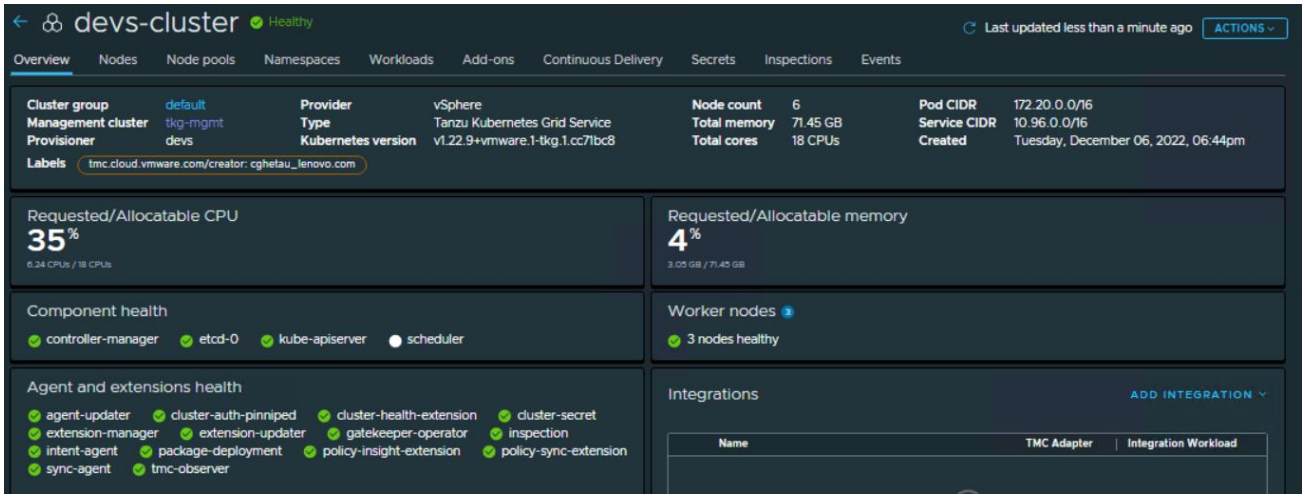

If all the components are healthy, we can deploy a Workload on the cluster.

In the TMC console go to Catalog, select the newly created cluster and choose cert-manager:

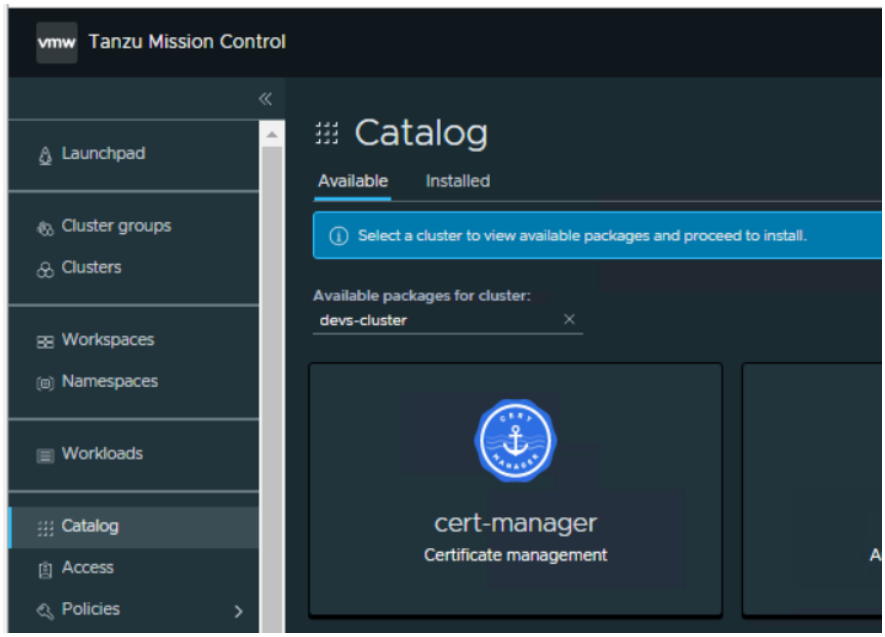

In the cert-manager workload window click on Install Package in the right corner:

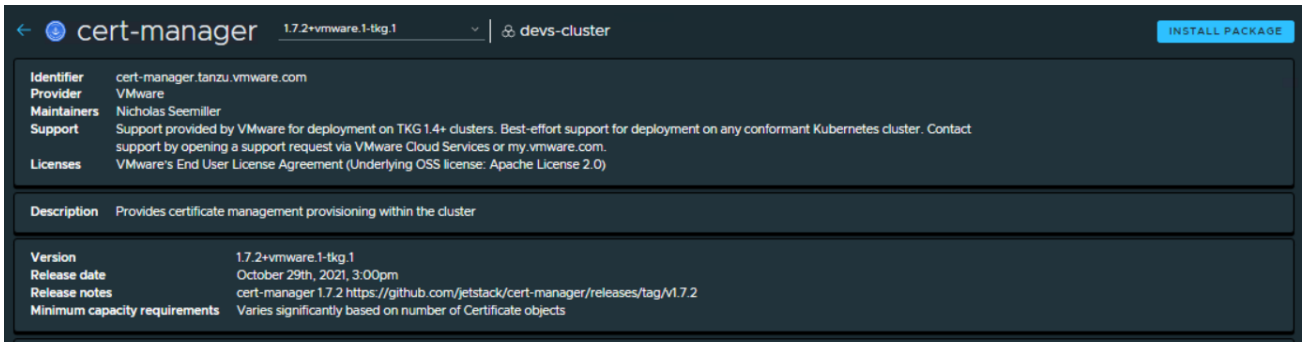

In the Install cert-manager window click on the Install Package:

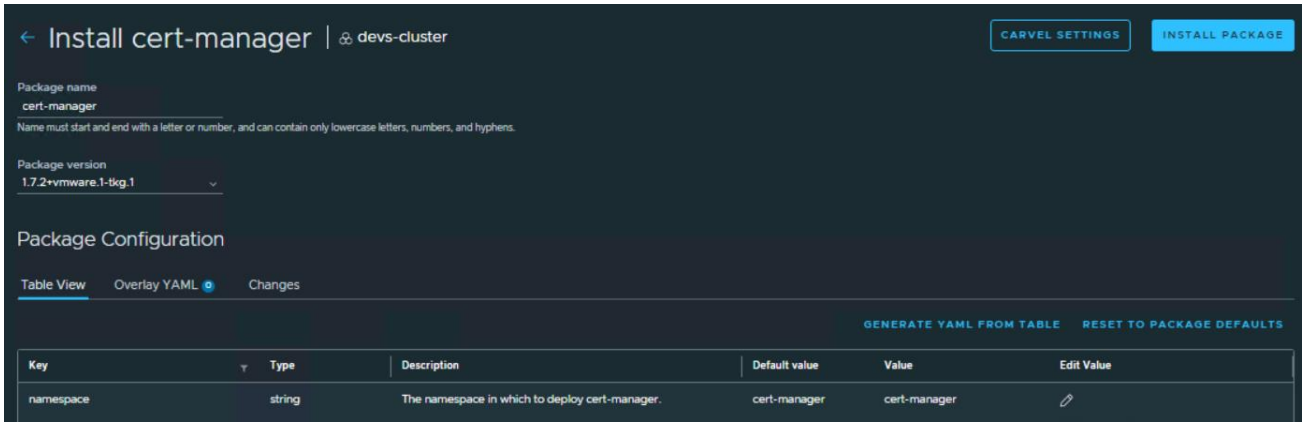

After the cert-manager package has been successfully installed it will be shown in the Add-ons tab of the cluster:

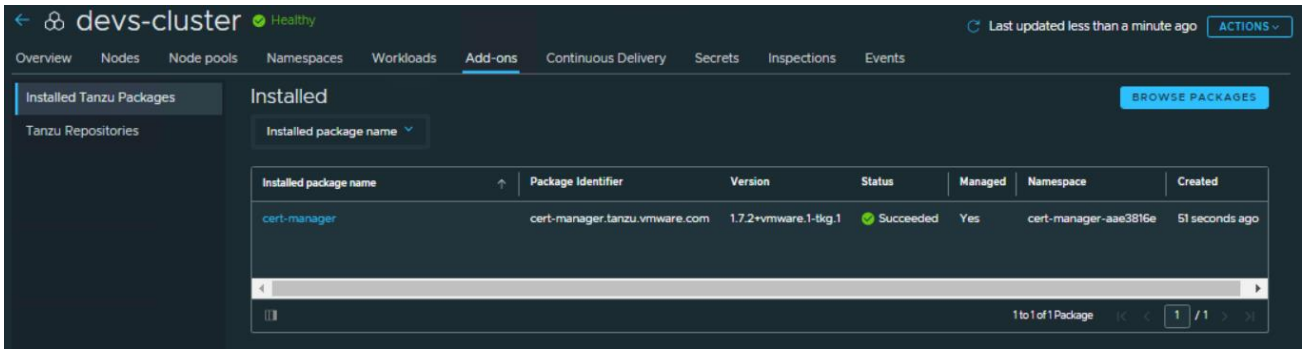

In the Workloads tab you can see the respective Deployments and Replica sets:

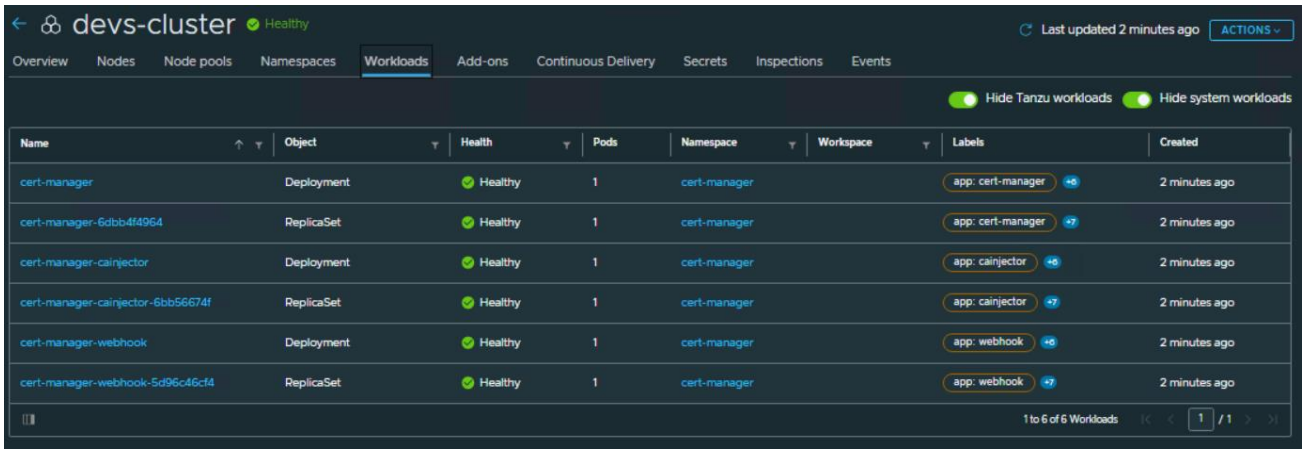

8 Deployment example

This chapter describes an example deployment of Aria Suite 8.14, VMware Tanzu 1.x, vSAN 8.0, and NSX-T 4.1.2.1 following the guidance in the VMware Validated Design (VVD) documentation. Four physical servers are used for each of the shared edge and compute, management, and additional compute clusters. Lenovo ThinkAgile VX servers are used for the shared edge and compute cluster and management cluster. Hardware views

The various hardware views are described in this section.

8.1.1 Rack view

Figure 16 shows a view of the rack with the twelve servers and switches.

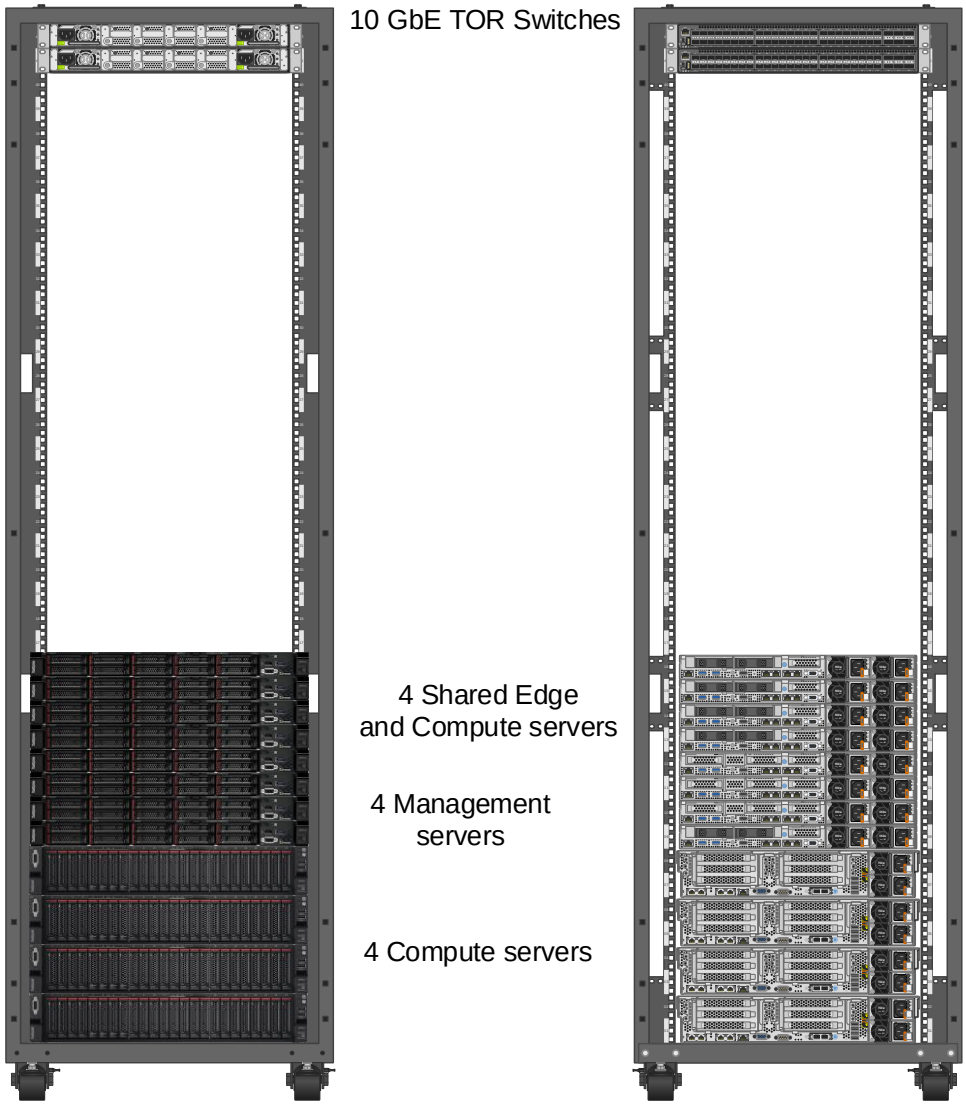

Figure 16: Rack Layout

8.1.2 Network view

Figure 17 shows a view of the physical 10 GbE network and connections to the external internet.

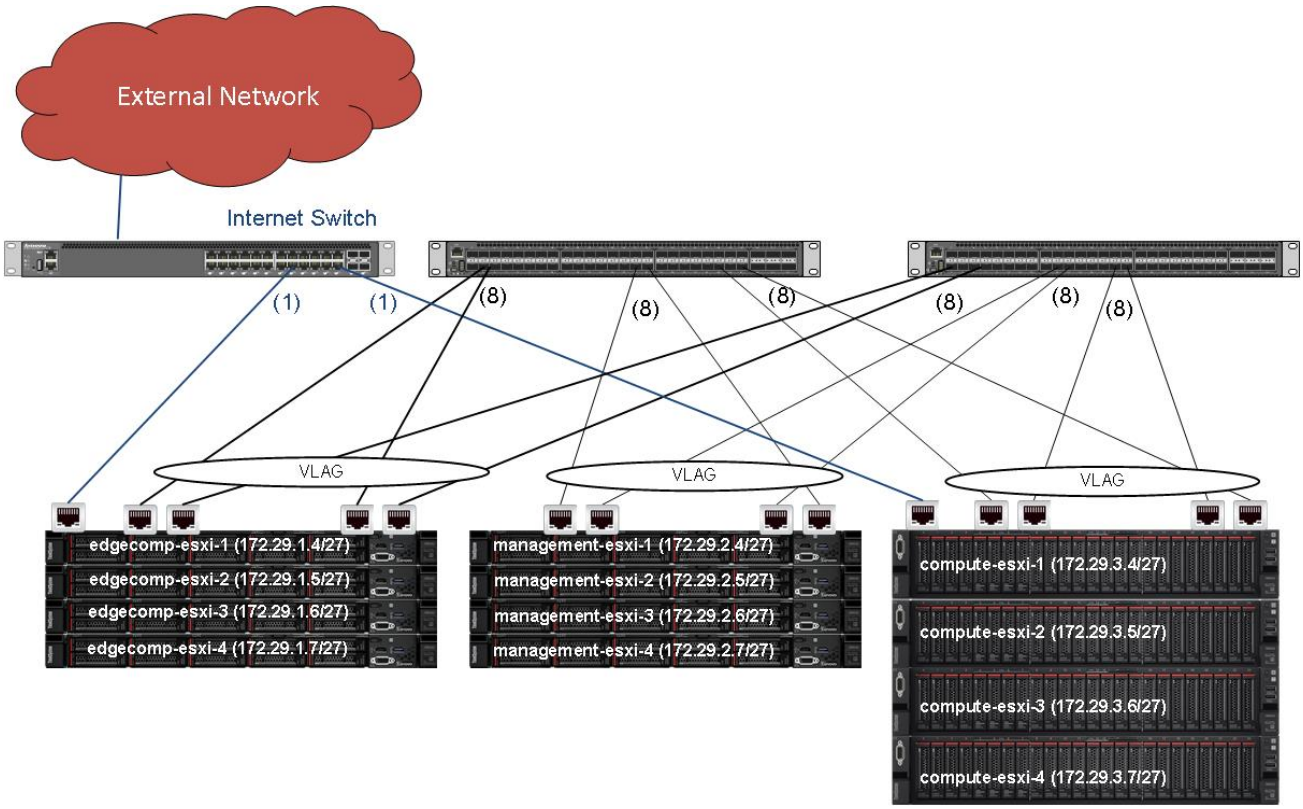

Figure 17: Networking Overview

For the shared edge and compute, management and additional compute clusters, the nodes use VLAG technology and as such are using a LAG configuration within the vSphere Distributed Switches. It is recommended to use VLAG for all the clusters connected to the same set of switches.

The servers in the shared edge and compute cluster and the additional compute cluster are connected to a 1G switch. This switch in turn is connected to the internet via a gateway and firewall (not shown).

8.2 IP/VLAN mapping

This example deployment uses the following nine VLANs:

- Management
- vMotion
- FT
- Storage
- VTEP
- vSAN
- vRA1
- vRA2 (for second region)
- Compute VMs
- vSphere Pod

Table 24 lists example IP address ranges for the VLANs in each cluster where RID means Rack ID.

Table 24: Network Segments

Traffic	**Shared Edge and Compute (RID 1)**		**Management (RID 2)**		**Compute (RID 3)**	
	Subnet	**VLAN**	**Subnet**	**VLAN**	**Subnet**	**VLAN**
Manage	172.29.1.0/27	101	172.29.2.0/27	201	172.29.3.0/27	301
vMotion	172.29.1.32/27	102	172.29.2.32/27	202	172.29.3.32/27	302
FT	172.29.1.64/27	103	172.29.2.64/27	203	172.29.3.64/27	303
Storage	172.29.1.96/27	104	172.29.2.96/27	204	172.29.3.96/27	304
TEP	172.29.1.128/27	105	172.29.2.128/27	205	172.29.3.128/27	305
vSAN	172.29.1.160/27	106	172.29.2.160/27	206	172.29.3.160/27	306
vRA1	N/A	107	172.29.2.192/27	207	N/A	307
vRA2	N/A	108	172.29.2.224/27	208	N/A	308
Comp VMs	172.29.2.192/27	109	N/A	209	172.29.2.192/27	309
vSphere Pod	172.29.2.224/27	110	N/A		172.29.3.224/27	310

In this example, each cluster needs a minimum of five network segments within the 172.29.RID.x address range. Each segment does not require more than 30 IP addresses; therefore, a 255.255.255.224 (/27) netmask provides enough addresses. The vSphere pods and virtual machines uses dedicated VLAN to address appropriate workloads running on them. The same VLAN IDs can be used across racks with different IP segments. In this example, that option is not available because the switches and routers are shared across the three clusters. For more information about how customers can plan their network segments, see the VMware NSX-T Design Guide.

8.3 Cluster Deployment

This section describes list of underlay and overlay virtualized networking used for the clusters. Multiple transport zones are used to segregate the clusters and logical switches that participate in each cluster. With NSX-T, there are flexible options chosen to use either underlay or overlay for different tenants and workloads. NSX-T and VMware Validated Design provides flexibility to deploy edge, management and compute VMs on the same cluster or shared cluster or dedicated cluster.

8.3.1 Deploying vSphere with Tanzu with dedicated cluster

vSphere with Tanzu deployment is done from VCF console. vSphere with Tanzu can be deployed in two clusters, one cluster for the Management and Edge functions, and another one dedicated to Workload Management. Figure 18 shows shared management and edge cluster and dedicated Kubernetes cluster.

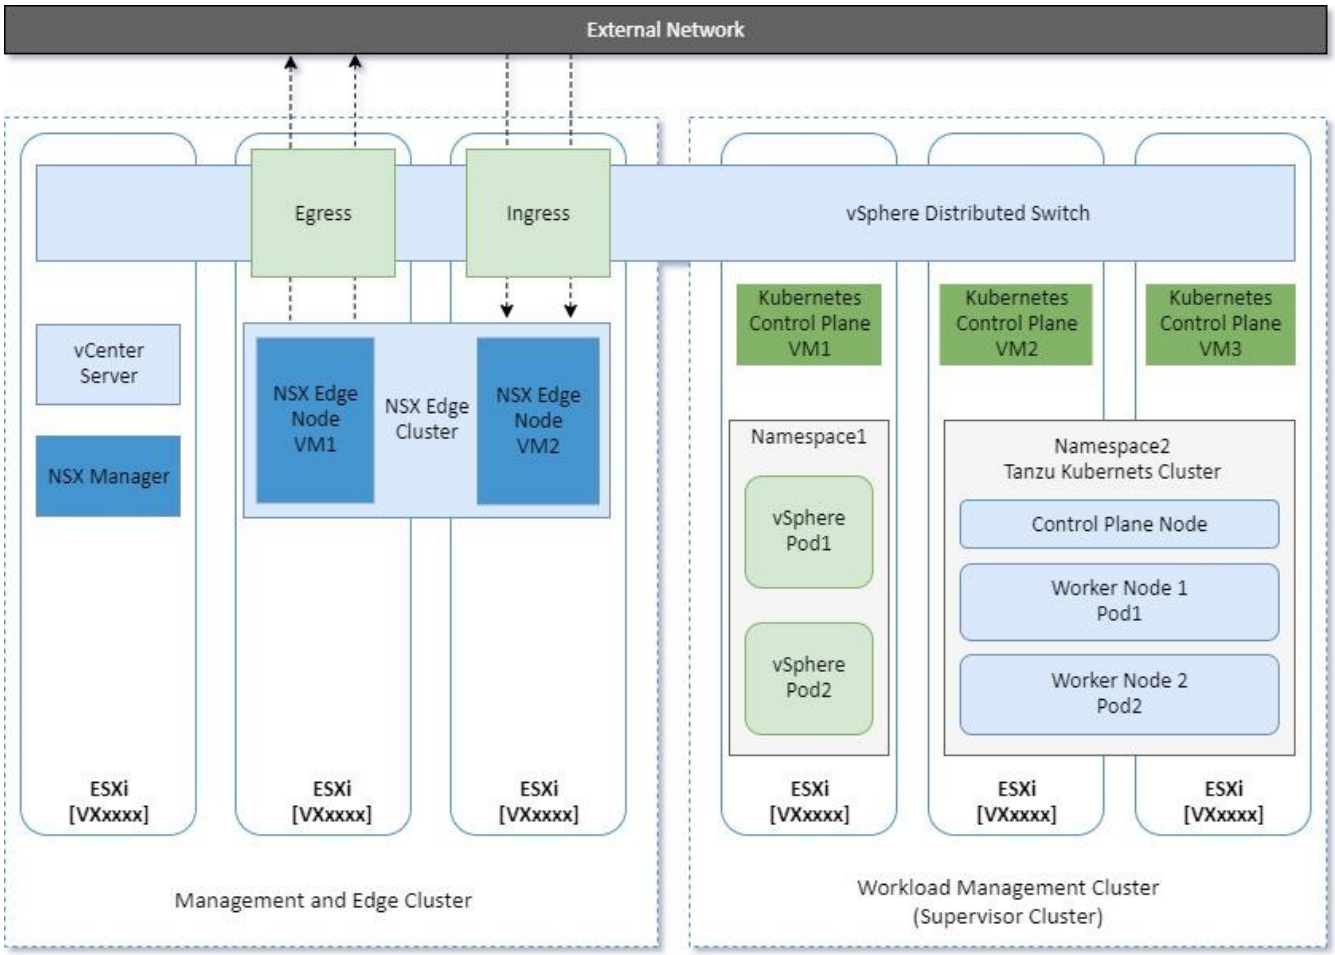

Figure 18: vSphere Tanzu with Dedicated Clusters

Deploying Tanzu Kubernetes Grid(TKG) is network sensitive operation and proper configurations need to be set in the SDDC Manager deployment input file. Since it is an automated deployment, any small configuration

issue will cause failure. The prerequisites are a running VI Workload Domain for Tanzu Kebernetes Cluster, NSX-T Edge on the VI Workload domain, a subnet for pod networking (non-routable, minimum /22), a subnet for Service IP (non-routable, minimum /24), a subnet for ingress(routable, minimum /27), a subnet for egress(routable, minimum/27) and MTU set to 9000 for all VLANs in physical switches. Also the content library needs to be available in vSphere to download and install components.

8.3.2 Deploying vSphere with Tanzu Consolidated Architecture

vSphere with Tanzu can be deployed in a single vSAN cluster where VCF workload management domain is running. This consolidated architecture hosts management, workload and edge components on a single cluster which ideally suit for development and SMB environment, but it can also be used for large environments where isolation needed based on teams or organization group. Figure 19 shows consolidated architecture to run all domains on the same cluster.

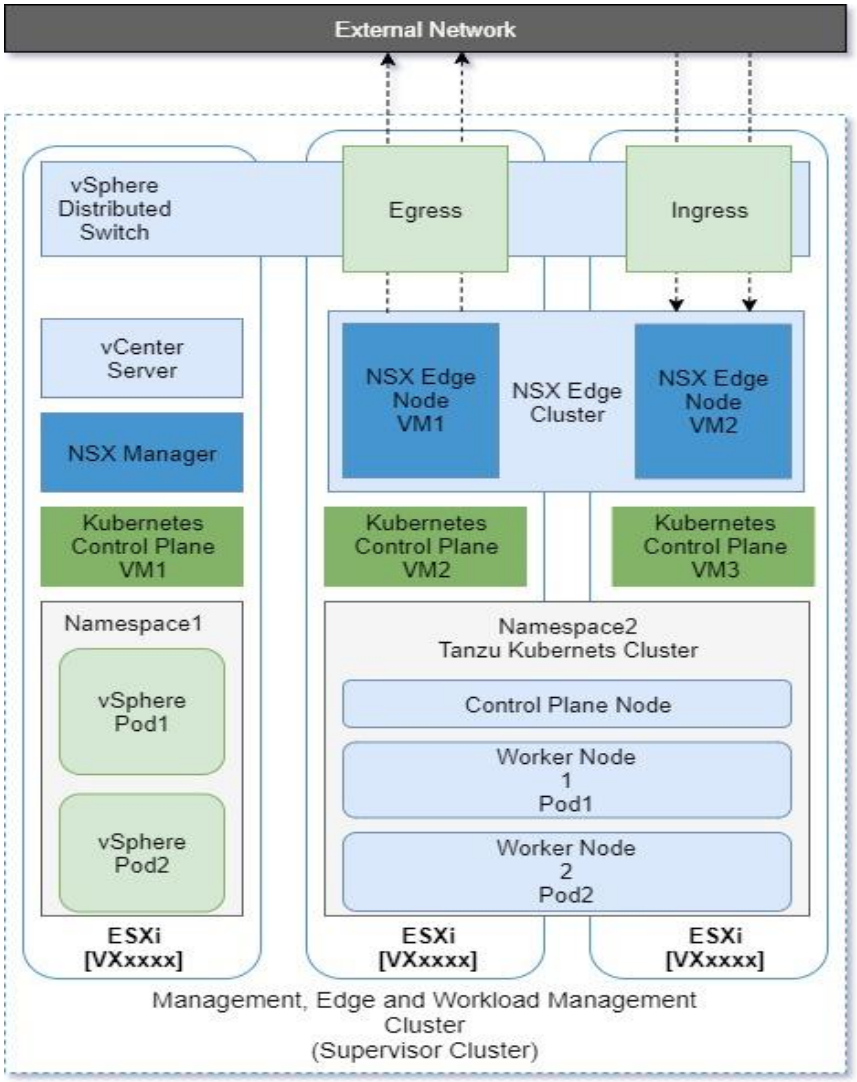

Figure 19: vSphere Tanzu with consolidated architecture

8.3.3 vSphere with Tanzu Deployment Best Practices

The SDDC Manager automates the deployment of Tanzu and the following networking considerations taken care before starting the deployment.

- The MTU should be set to 9000 for all VLANs except the management VLAN
- The NSX-T Edge cluster should use EBGP
- The routable VLANs should be configured on physical switches for ingress and egress traffic for Workload Domain
- The first IP of the subnet should not be used as default gateway IP on the physical switches for ingress and egress VLANs. Because this is being used by NSX-T for SNAT rule for subnets.
- The DNS server should be accessible from egress VLAN
- At least 5 IP addresses need to be reserved for the control plane VMs on the management VLAN
- The VLANs used for NSX-T Edge VTEP and Host TEP should be inter-routable.

9 Conclusion

The combination of Lenovo ThinkAgile VX nodes and VMware Cloud Foundation provides an ideal hybrid cloud platform for a customer to start their application modernization journey.

Resources

For more information about the topics that are described in this document, see the following resources:

- Software Defined Data Center:
  vmware.com/software-defined-datacenter
- VMware Validated Designs Documentation (VVD):
  vmware.com/support/pubs/vmware-validated-design-pubs.html
- vSphere Hypervisor (ESXi):
  vmware.com/products/vsphere-hypervisor
- vCenter Server:
  vmware.com/products/vcenter-server
- vSAN:
  vmware.com/products/virtual-san
- VMware Compatibility Guide (VCG)
  vmware.com/resources/compatibility
- NSX:
  vmware.com/products/nsx
- VMware NSX-T Reference Design Guide
- https://nsx.techzone.vmware.com/resource/nsx-t-reference-design-guide-3-0Aria Suite:
  vmware.com/products/Aria-suite.html
- Aria Automation:
  vmware.com/products/Aria-automation
- Aria Automation Reference Architecture:
  vmware.com/files/pdf/products/vCloud/Aria-Automation-6x-Reference-Architecture.pdf
- Aria Operations:
  vmware.com/products/Aria-operations
- Aria Business:
  vmware.com/products/Aria-business
- Aria Operations for Logs:
  vmware.com/products/aria-operations-for-logs

Document history

Version 1.0	29 September 2021	• First version for Lenovo ThinkAgile VX with VCF
Version 1.1	11 January 2022	• Added more ThinkAgile VX appliances • Added VMware Tanzu Advance edition features • Revised SDDC deployment components • Added VMware Tanzu deployment best practices
Version 1.2	31 August 2022	• Added VMware Tanzu use case – nginx deployment
Version 1.3	13 December 2022	• Added VMware Tanzu Mission Control deployment • Updated 7.3.3 Chapter with the newly SupervisorControlPlane VM installation wizard
Version 1.4	21 March 2023	• Added 7.3.1.5 Deploy Aria Suite
Version 1.5	7 August 2023	• Updated Chapters 7.3.1.3, 7.3.1.4, 7.3.1.5 and 7.3.1.6 with the new VMware Cloud Foundation 5.0 (vSphere 8.0.1)
Version 1.6	20 September 2023	• Updated Chapters 7.3.1.3, 7.3.1.4, 7.3.1.5 and 7.3.1.6 with the new VMware Cloud Foundation 5.0 (vSphere 8.0.1)
Version 1.7	22 March 2024	• Updated Chapters 7.3.1.2, 7.3.1.3, 7.3.1.5 and 7.3.1.6 with the new VMware Cloud Foundation 5.1 (vSphere 8.0.2)

Trademarks and special notices

© Copyright Lenovo 2024.

References in this document to Lenovo products or services do not imply that Lenovo intends to make them available in every country.

The following terms are trademarks of Lenovo in the United States, other countries, or both:

Lenovo®
Flex System
System x®
ThinkAgile
ThinkSystem
XClarity®

The following terms are trademarks of other companies:

Intel® and Xeon® are trademarks of Intel Corporation or its subsidiaries.

Linux® is the trademark of Linus Torvalds in the U.S. and other countries.

Active Directory®, Azure®, Hyper-V®, Microsoft®, SQL Server®, and Windows® are trademarks of Microsoft Corporation in the United States, other countries, or both.

Other company, product, or service names may be trademarks or service marks of others.

Information is provided "AS IS" without warranty of any kind.

All customer examples described are presented as illustrations of how those customers have used Lenovo products and the results they may have achieved. Actual environmental costs and performance characteristics may vary by customer.

Information concerning non-Lenovo products was obtained from a supplier of these products, published announcement material, or other publicly available sources and does not constitute an endorsement of such products by Lenovo. Sources for non-Lenovo list prices and performance numbers are taken from publicly available information, including vendor announcements and vendor worldwide homepages. Lenovo has not tested these products and cannot confirm the accuracy of performance, capability, or any other claims related to non-Lenovo products. Questions on the capability of non-Lenovo products should be addressed to the supplier of those products.

All statements regarding Lenovo future direction and intent are subject to change or withdrawal without notice, and represent goals and objectives only. Contact your local Lenovo office or Lenovo authorized reseller for the full text of the specific Statement of Direction.

Some information addresses anticipated future capabilities. Such information is not intended as a definitive statement of a commitment to specific levels of performance, function or delivery schedules with respect to any future products. Such commitments are only made in Lenovo product announcements. The information is presented here to communicate Lenovo’s current investment and development activities as a good faith effort to help with our customers' future planning.

Performance is based on measurements and projections using standard Lenovo benchmarks in a controlled environment. The actual throughput or performance that any user will experience will vary depending upon considerations such as the amount of multiprogramming in the user's job stream, the I/O configuration, the storage configuration, and the workload processed. Therefore, no assurance can be given that an individual user will achieve throughput or performance improvements equivalent to the ratios stated here.

Photographs shown are of engineering prototypes. Changes may be incorporated in production models.

Any references in this information to non-Lenovo websites are provided for convenience only and do not in any manner serve as an endorsement of those websites. The materials at those websites are not part of the materials for this Lenovo product and use of those websites is at your own risk.