Tanzu for Kubernetes Operations on VxRail

Building Modern Applications across Clouds

vmware | DELL Technologies

Table of contents

Executive Summary

In this era of modern app development, the complexity of multi-layer design and distributed architectures, drive the need for the operations model to evolve. Kubernetes has emerged as the de facto container orchestration platform and the number of containerized applications running in production continues to grow. As such, customers are increasingly interested in a consistent runtime across their deployments regardless of where they reside. They are also looking for consistent operations and management with fine grain visibility into their Kubernetes and application frameworks. They want their environments to be secure, easy to deploy, manage and upgrade, while managing the application sprawl of microservices

Cloud native application development leverages DevOps and CI/DC (Continuous Integration/Continuous Delivery) practices that streamline the delivery of production ready microservice applications. Cloud native application model suits many workloads, and an increasing number of companies are “born in the cloud” or migrating to the cloud. When companies build and operate applications using a cloud native architecture, they bring new ideas to market and respond to customer demands faster.

This reference architecture presents a design and implementation methodology to integrate on-premises private cloud powered by Tanzu for Kubernetes Operations, running on Dell VxRail with Tanzu for Kubernetes Operations on VMware Cloud on AWS. This integration is further extended to other Kubernetes cloud offerings, such as Amazon EKS, that bring all Kubernetes clusters under a common management domain using Tanzu Mission Control.

On-demand, multi-cloud connectivity and management can be a challenge technically, as well as costly for most customers. First, connections to the cloud providers must be established and maintained separately, which adds complexity to an already complex paradigm. Second, when applications are spread over smaller clusters across clouds to provide separate fault domains, management of these clusters on an ongoing basis is not trivial. In this reference architecture we present a solution using a third-party cloud connectivity provider such as Equinix which can manage connection enumeration to most major cloud providers simultaneously. Such a solution would be beneficial for customers connecting to multiple cloud providers, however, is not a requirement for this reference architecture. With centralized policy and cluster management, and fine grain insight into cluster operations, Tanzu Mission Control and Tanzu Observability provide a common management methodology across all cloud instances.

For end-to-end connectivity, cross-site load balancing and ingress, NSX Advanced Load Balancer provides a robust networking stack that can support global DNS services as Kubernetes deployment instances grow from on-premises private cloud to multi-cloud environment. This reference architecture uses this NSX Advanced Load Balancer GSLB (Global Server Load Balancing) capability to load balance application instances across clouds. AKO (NSX Advanced Load Balancer Kubernetes Operator) and AMKO (Avi Multi-Cluster Operator) provides ingress services across Kubernetes deployments.

For application security, NSX Advanced Load Balancer features an Intelligent Web Application Firewall (iWAF) that covers OWASP CRS protection, support for compliance regulations such as PCI DSS, HIPAA, and GDPR, and signature-based detection. It deploys positive security model and application learning to prevent web application attacks. Additionally, built-in analytics provide actionable insights on performance, end-user interactions and security events in a single dashboard (Avi App Insights) with end-to-end visibility.

Dell VxRail delivers a turnkey experience and is fully integrated, pre-configured, and pre-tested solution. Tanzu for Kubernetes Operations on VxRail is a future proof solution that simplifies transformation journey to modern applications for most customers. Whether its move from legacy to cloud native applications, repatriating cloud native applications to on-premises private cloud, or architecting distributed application on multi-cloud environment, Tanzu Kubernetes on VxRail is the all-encompassing solution.

Audience

This white paper is intended for architects, engineers, consultants, and IT administrators who design, implement, and manage modern application environment on-premises or in the cloud. Readers with strong understanding of technologies such as VMware NSX Advanced Load Balancer, vSphere with Tanzu, VMware vSAN, and cloud native concepts will benefit from the content in this paper.

Architecture Overview

The solution is built on multi-cloud architecture, including On-premises private cloud, Amazon EKS and VMware Cloud on AWS, with the three site instances making up the global DNS namespace. Amazon EKS was included in the architecture to demonstrate public cloud management and integration capabilities of Tanzu portfolio of products. Another public cloud, such as Azure or Google (GCP) can also be integrated with relevant ease. NSX Advanced Load Balancer (formerly known as Avi) manages the global DNS zone. User queries for applications to the corporate DNS are directed to the appropriate site holding the application. Instances of the applications are installed on multiple sites for load balancing and high availability. Load balancing based on geo-location or priority, provides improved performance by directing the user to the nearest or highest priority site holding the desired application. In this architecture the primary site is the on-premises private-cloud infrastructure and software built on Tanzu Kubernetes Grid service. This primary site holds the Active directory domain infrastructure and corporate DNS services. This site also has the leader GSLB (Global Server Load Balancing) service instance. Avi multi-cluster Kubernetes operator (AMKO) is installed here, which manages and coordinates ingress and load balancing from Avi Kubernetes operators (AKO) from all sites. VMware Cloud on AWS (VMC) and Amazon EKS makeup the other two GSLB follower instances. These follower sites only require AKO (Avi Kubernetes operator) installation. Tanzu Service Mesh, a part of Tanzu for Kubernetes Operations provided end-to-end connectivity, security, and insights for microservices running across clouds making up the global application namespace.

Figure 1: Logical Architecture Overview

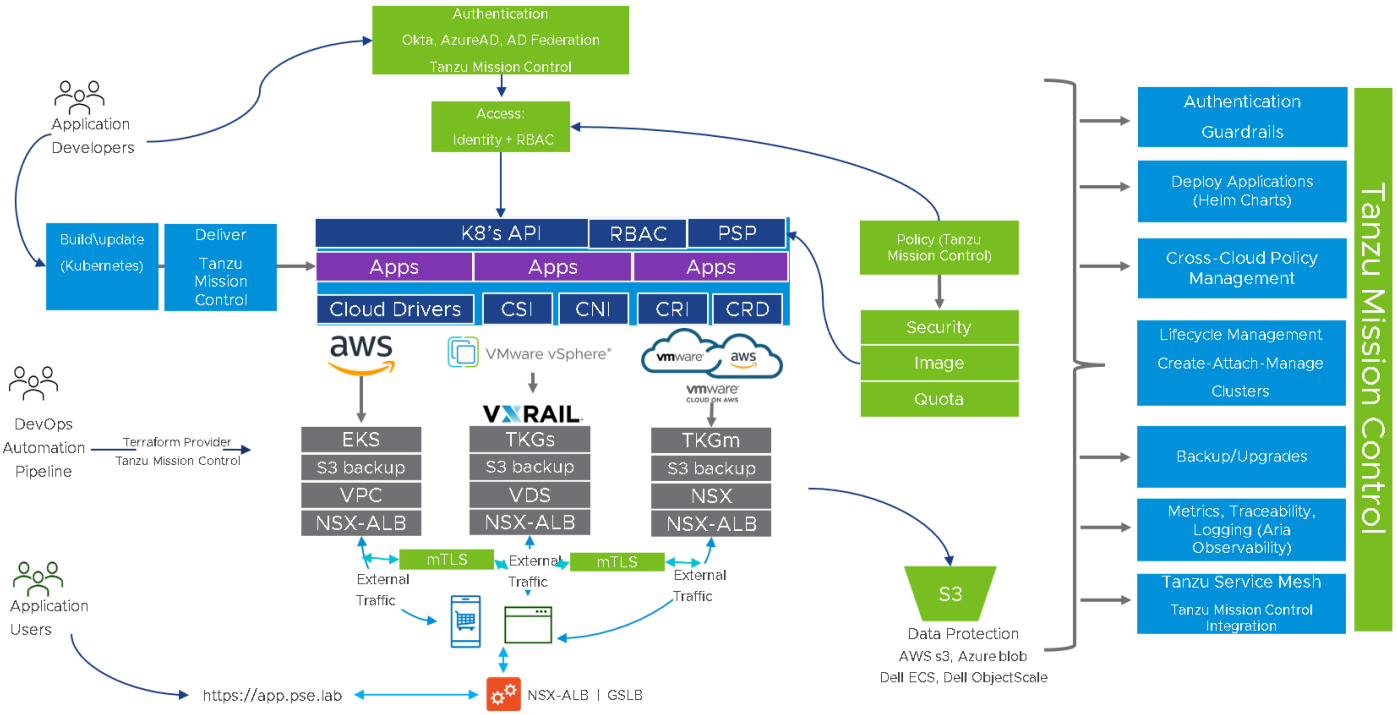

Tanzu Mission Control (Tanzu Mission Control) managed the complete lifecycle of Kubernetes clusters across sites including Amazon EKS clusters. Tanzu Mission Control also provides centralized policy management and developer self-service access to all three sites. Tanzu Observability now part of VMware Aria, provided fine grain insight and observability for VMware Tanzu and Amazon EKS clusters. Tanzu Mission Control also provided data-protection through Velero and Restic open-source software using S3 compatible storage, on-premises or in the cloud.

An intermediate site connecting the on-premises private cloud to multiple cloud providers is hosted at Equinix. On-demand connections to multiple cloud providers are enumerated from Equinix, per customer performance and cost requirements.

Note: An intermediate cloud services provider is not a requirement for this reference architecture. It is however an option for customers who want to take advantage of these services for their multi-cloud connectivity. Alternatively,

customers can connect to their cloud providers individually via VPN or direct connections. For more information on how Equinix is configured to provide multi-cloud connectivity, please see Multi-Cloud Network Connectivity with Equinix Overview.

Key Components

This solution is built upon a solid foundation using Dell VxRail cluster made up of four V570 model HCI nodes. When configured with VMware vSAN and NSX Advanced Load Balancer, Dell VxRail provides an enterprise grade software defined datacenter architecture that is agile, easy to manage and secure. vSphere with Tanzu enhances these underlying qualities and delivers a developer-ready, modern application platform for upstream Kubernetes clusters. From a manageability perspective, Tanzu Mission Control and Tanzu Observability provides a solution that is future-proof and extensible from on-premises to the cloud. A description of the key components follows.

Dell VxRail

Whether accelerating data center modernization, deploying a hybrid cloud, or creating a developer-ready Kubernetes platform, VxRail delivers a turnkey experience that enables customers to continuously innovate. The only hyperconverged system jointly engineered by Dell Technologies and VMware, it is fully integrated, pre-configured, and pre-tested, automating lifecycle management and simplifying operations. Powered by VMware vSAN or VMware Cloud Foundation, VxRail transforms HCI networking and simplifies VMware cloud adoption, while meeting any HCI use case - including support for the most demanding workloads and applications. Learn more.

NSX Advanced Load Balancer

VMware NSX Advanced Load Balancer provides multi-cloud load balancing, web application firewall and application analytics across on-premises data centers and any cloud. The software-defined platform delivers applications consistently across bare metal servers, virtual machines, and containers to ensure a fast, scalable, and secure application experience. Learn more.

Tanzu Mission Control

VMware Tanzu Mission Control is a centralized management hub, with a robust policy engine, which simplifies multi-cloud and multi-cluster Kubernetes management. Whether you are new to Kubernetes, or quite experienced, Tanzu Mission Control helps platform operators reduce complexity, increase consistency, and offer a better developer experience. Learn more.

Tanzu Observability (Aria)

VMware Tanzu Observability by Wavefront now part of VMware Aria is an observability platform specifically designed for enterprises needing monitoring, observability, and analytics for their cloud-native applications and environments. DevOps, SRE and developer teams use Tanzu Observability to proactively alert on, rapidly troubleshoot and optimize performance of their modern applications running on the enterprise multi-cloud. Learn more.

Tanzu Kubernetes Grid

Tanzu Kubernetes Gid Standard has everything an enterprise needs to make best use of Kubernetes as part of its vSphere-based infrastructure. Kubernetes is embedded in the vSphere control plane, addressing the needs of both operators and developers. Operators can support virtual machines and containers side-by-side on a unified platform, group these elements into applications and simplify management. Developers gain self-service access to resources using Kubernetes APIs and speed development processes.

Tanzu Service Mesh

Tanzu Service Mesh provides advanced, end-to-end connectivity, security, and insights for modern applications—across application end-users, microservices, APIs, and data—enabling compliance with Service Level Objectives (SLOs) and data protection and privacy regulations. More information can be found at VMware Tanzu Service Mesh.

VMware Cloud on AWS

VMware Cloud on AWS is an integrated cloud offering jointly developed by Amazon Web Services (AWS) and VMware. You can deliver a highly scalable and secure service by migrating and extending your on-premises VMware vSphere-based environments to the AWS Cloud running on Amazon Elastic Compute Cloud (Amazon EC2). Learn more.

Amazon Elastic Kubernetes Service

Amazon Elastic Kubernetes Service (Amazon EKS) is a managed service that you can use to run Kubernetes on AWS without needing to install, operate, and maintain your own Kubernetes control plane or nodes.

Key Software Components

The following table shows key software component versions used in this reference architecture.

Table 1: Key software components

Software Specifications		
On-Premises Component	**Version**	**Notes**
– VxRail Manager	– 7.0.350-27409467	
– VMware ESXi	– VMware ESXi, 7.0.3, 19193900	4 x VxRail V570 nodes
– VMware vCenter Server	– vSphere Client version 7.0.3.00600	
– NSX Advanced Load Balancer	– 20.1.7 Enterprise	3 x Avi appliances for HA
– vSphere with Tanzu	– Tanzu Kubernetes Grid	K8s version v1.20.12+vmware.1
– Pod Networking (CNI)	– Antrea Advanced	1 Core
VMware Cloud on AWS	**Version**	**Notes**
– VMware ESXi	– VMware ESXi, 7.0.3, 19888012	3 x ESXi nodes
– VMware (SDDC) vCenter Server	– vSphere Client version 7.0.3.20000	
– NSX Advanced Load Balancer	– 21.1.4 Enterprise	No orchestrator mode
– Tanzu Kubernetes Grid	– Multi-Cloud (TKGm) 1.5.3->1.6.0 *	K8s version v1.22.9 ->1.23.8 *
– Service Installer for VMware Tanzu	– 1.3	
– Pod Networking (CNI)	– Antrea Advanced	1 Core
Amazon EKS	**Version**	**Notes**
– Kubernetes version	– v1.22.9	
– NSX Advanced Load Balancer	– 21.1.4	Enterprise license
– Pod Networking (CNI)	– Amazon Native VPC CNI	
SaaS Services	**Version**	**Notes**
– Tanzu Mission Control	– Advanced	1 Core
– Tanzu Observability		20 PPS
– Tanzu Service Mesh	Advanced	

* On VMware Cloud on AWS, Kubernetes version was upgraded from 1.5.3 to 1.6.0 to match latest available version at the time of the release of this second version of the RA. This also validated the upgrade process without issue in the reference environment. Instructions to upgrade the management and workload clusters can be found here.

Solution Configuration

Any solution, especially a multi-cloud solution as complex as presented in this document can be configurated in multiple ways depending on the requirements at hand. This document presents the solution configuration in a modular fashion where each site configuration is independent of each other except for GSLB and Avi DNS service configuration which depends on the number of sites configured. Distributed application functionality depends on these services across sites. The flow chart below shows high-level workflow used to configure the sites in this reference architecture.

Figure 2: Configuration workflow

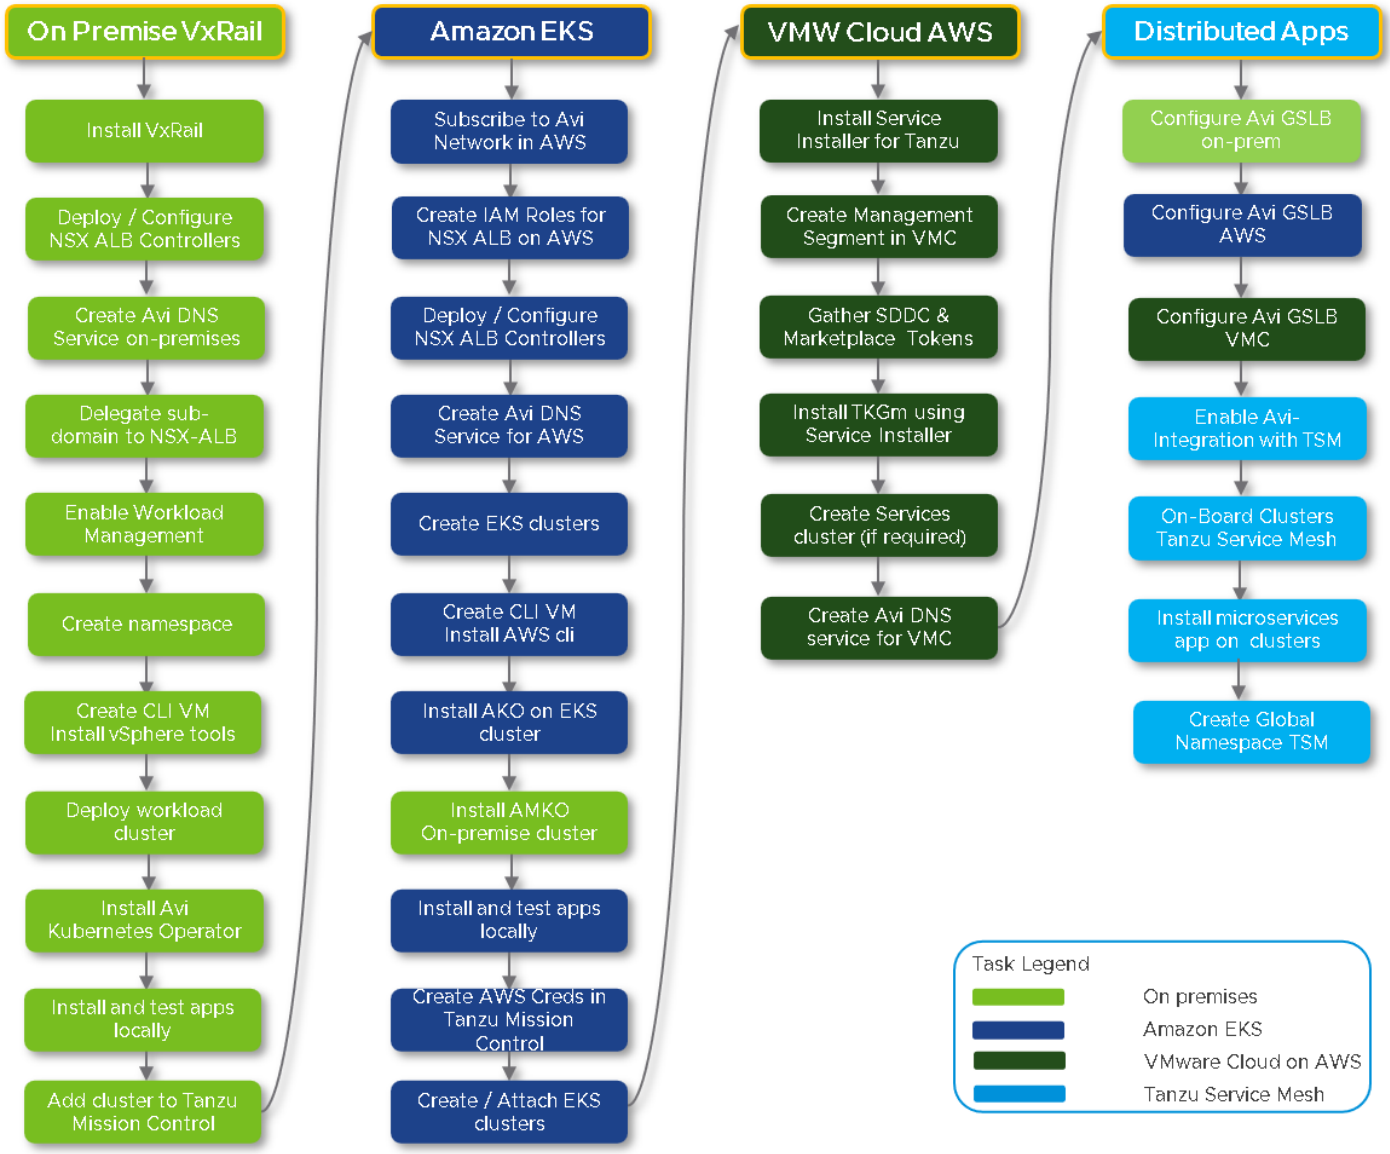

Solution Network Overview

As depicted in figure 2, the on-premises private cloud (datacenter) is connected to Equinix via VMware SD-WAN. As compared to MPLS circuits which are expensive, VMware SD-WAN (formerly known as VeloCloud) provides a cost effective, secure, and zero-touch deployment option for WAN (Wide-Area-Network). Connections to AWS VPC and VMware Cloud on AWS are configured via AWS DirectConnect. DirectConnect bandwidth can be configured per customer requirements from 50Mbps to 10 Gbps. For this reference architecture 500 Mbps was configured for these connections.

Figure 3: Logical Lab Network

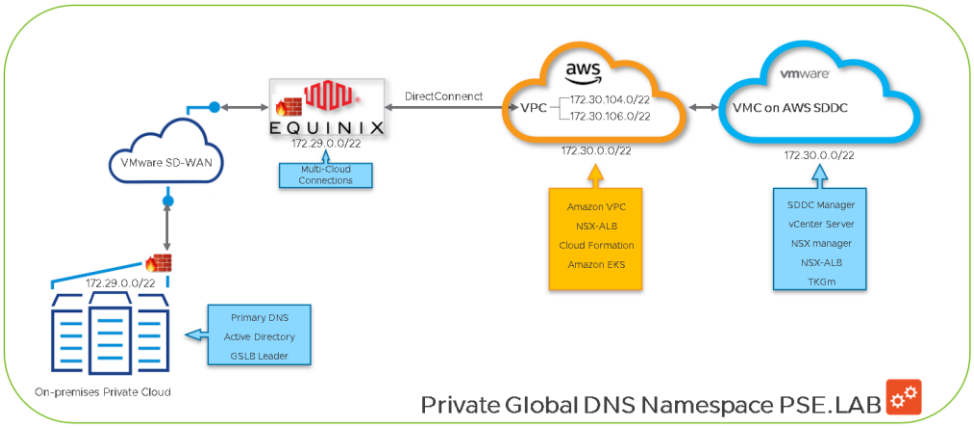

Note: An intermediate cloud services provider is not a requirement for this reference architecture. It is however an option for customers who want to take advantage of these services for their multi-cloud connectivity. Alternatively, customers can connect to their cloud providers individually via VPN or direct connections. For more information on how Equinix is configured to provide multi-cloud connectivity, please see Multi-Cloud Network Connectivity with Equinix Overview.

NSX Advanced Load Balancer provide GSLB functionality in this reference architecture. Global server loading balancing (GSLB) is the process of balancing an application's load across instances of the application that have been deployed to multiple locations. Load balancing can be performed based on user's geo-location or round-robin algorithms. With GSLB, when a Kubernetes application is installed, a virtual service is created with application's URL. Users access the application using its URL. The user is directed to the appropriate site based on algorithm and preference set by the GSLB administrator. Figure 4 below shows the GSLB workflow.

Figure 4: Global server load balancing

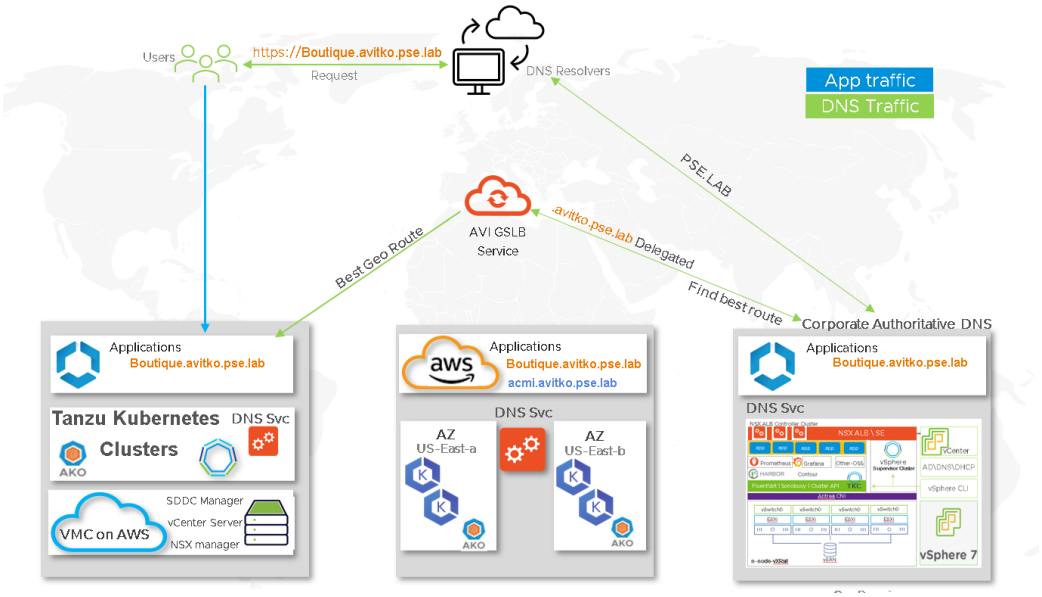

In instances when an application or site is unavailable the requests are serviced by the active sites. For this reference architecture the authoritative DNS sever was in the on-prem datacenter. For redundancy, secondary DNS servers were also installed on the Equinix site. If desired, the corporate DNS server or DNS zones can also be hosted on Amazon Route53. Figure 5 shows the workflow when a site is not available in GSLB environment.

Figure 5: Site failure in GSLB

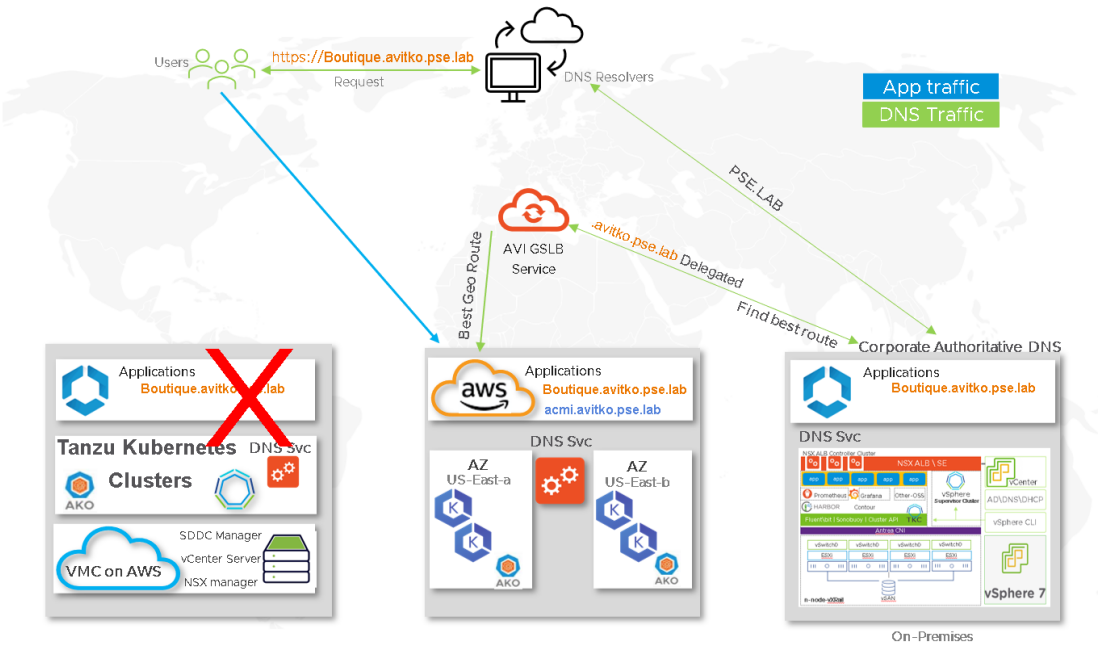

On-Premises Private Cloud

A four node Dell VxRail V570 cluster makes up the infrastructure foundation of this on-premises modern application solution. vSphere with Tanzu provides the capability to run Kubernetes workloads natively on the ESXi hypervisor and create upstream compliant Kubernetes clusters on demand. The NSX Advanced Load Balancer provides dynamically scaling load balancing endpoints for Tanzu Kubernetes clusters provisioned by the Tanzu Kubernetes Grid Service. Along with its Avi Kubernetes Operator and Avi Multi-Cluster Kubernetes Operator, NSX Advanced Load Balancer provides L4 and L7 ingress and load balancing to the deployed workloads. This site also serves as the "leader" GSLB site. VMware vCenter Server along with Dell VxRail Manager makes up the local infrastructure management domain. Harbor is used as the local registry and can be installed manually or via Tanzu Mission Control. In addition to Tanzu Observability, local monitoring and diagnostic, tools such as, Prometheus, Grafana, Fluent are also installed.

VMware vSAN provides enterprise class hyperconverged storage, which is consistent across deployments and integrates fully with VMware Tanzu. From a storage perspective vSAN future-proofs the solution with its integration with object storage types such as Dell ECS, Dell ObjectScale and others.

Figure 6: On-premises private cloud architecture

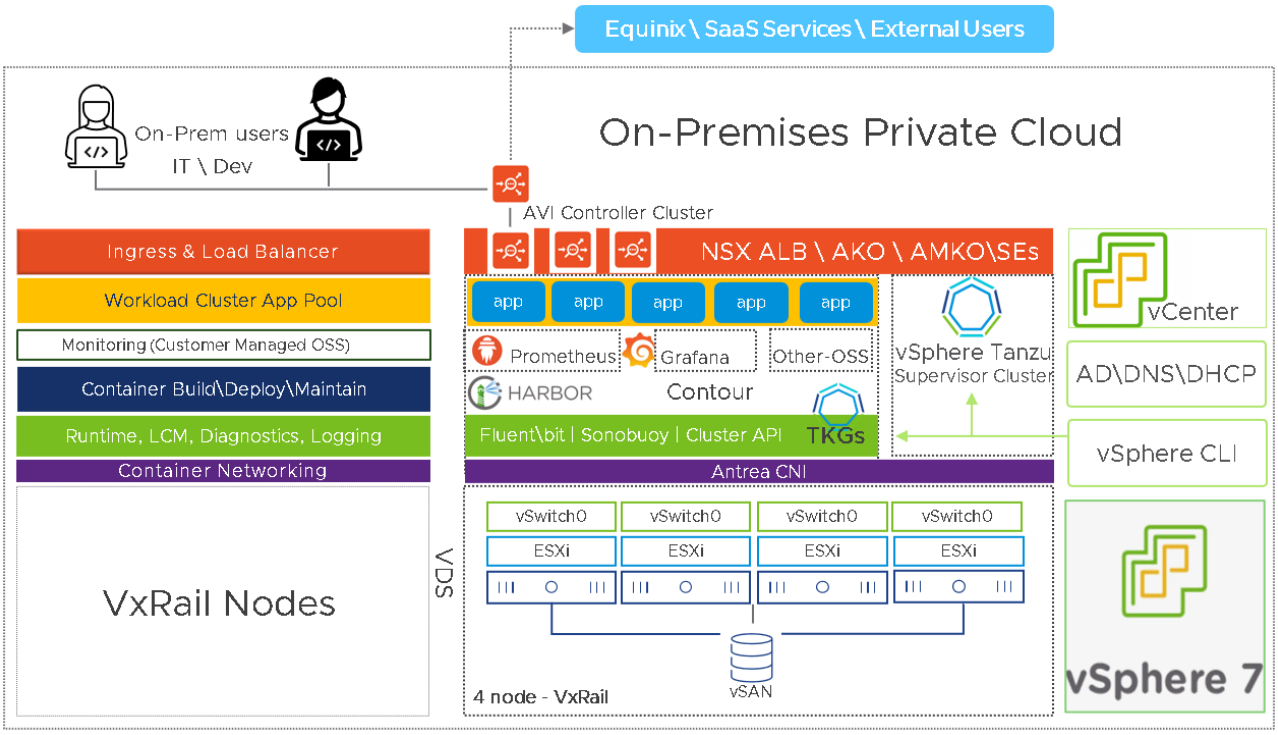

On-premises network overview

Dell VxRail deployment creates the Virtual Distributed Switch with minimum required port groups, such as vCenter and VxRail Management. Additional port groups for vSphere with Tanzu were created for NSX ALB and supervisor node management, front-end, and workload networks. Placing these networks on separate port groups provides isolation and enables application of granular security\firewall policies. Figure 7 depicts the high-level logical diagram of the network stack configured in the lab.

Figure 7: Logical Network Architecture

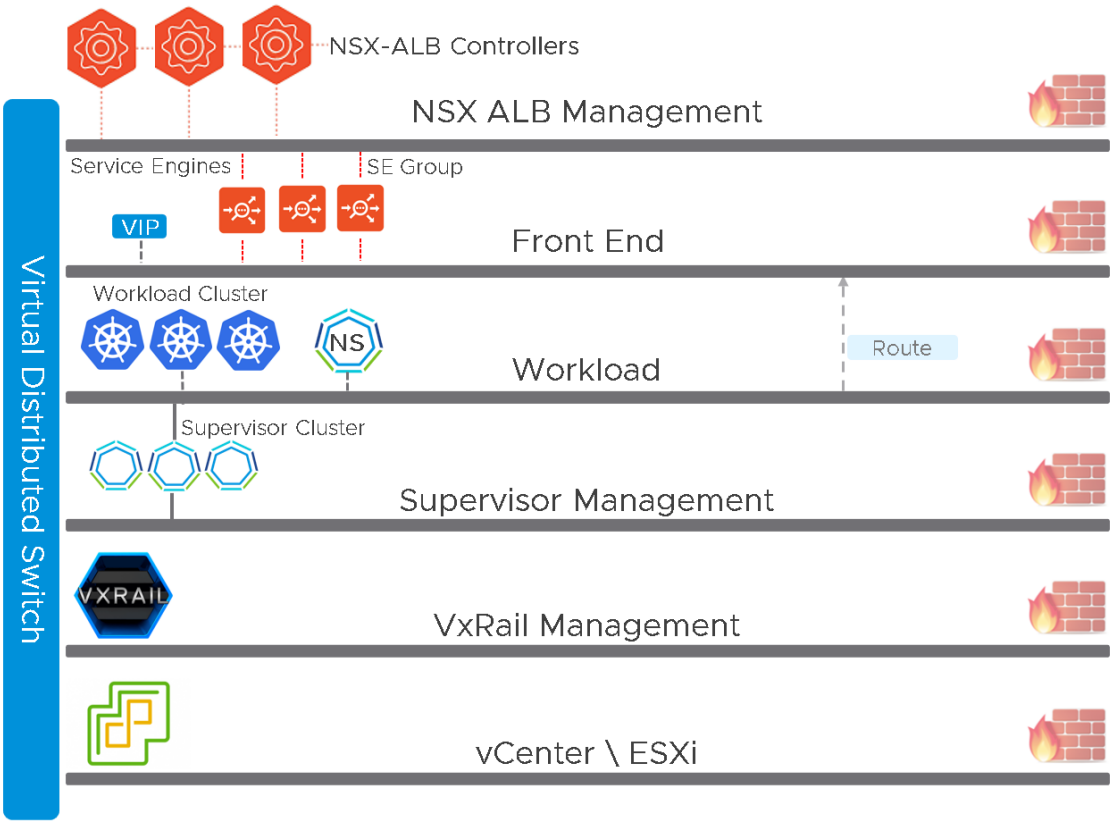

Note: The VIP and workload network should be routable. NSX Advanced Load Balancer management and vCenter Server management networks should also be able to communicate with each other.

In addition to networks provisioned with Dell VxRail deployment additional network were created for traffic segmentation. Table 3 lists these additional required network\port groups with a brief description.

Table 3: Tanzu Kubernetes networks

Network	Description
NSX ALB Management	NSX Advanced Load Balancer controllers and Service Engines connect to this network
Supervisor Management	TKGs Supervisor nodes are placed on this network
Front End	This is where the users connect to and holds the virtual services and VIPs
Workload	TKG workload cluster control plane and worker nodes connect here

Dell VxRail

Joint engineering between Dell and VMware leads to a curated and optimized VxRail hyperconverged experience. This deep integration combined with the simplicity of the VxRail HCI System Software enables seamless adoption of new technology and features, and provides an ideal platform across core, edge, and cloud.

- ✓ Consistent ease of use with automated full stack lifecycle management
- ✓ Simplify with a consistent operational model across your infrastructure landscape.
- ✓ Single point of support with 97% of all cases resolved in house.

Installation

Note: Prior to starting VxRail cluster installation, ensure that hardware is setup properly including Top-Of-Rack switches, and that the nodes are imaged with the desired VxRail version specific image. This image includes ESXi, vSAN, hardware firmware/drivers, and VxRail HCI System Software that will be deployed and configured automatically based on the

desired VxRail cluster JSON file configuration parameters. Please consult Dell VxRail support documentation for more details.

- Once the VxRail nodes are powered on, from your jump host, go to https://192.168.10.200 to connect to the VxRail Manager. The VxRail Deployment Wizard welcome screen displays, as shown in figure below. Click GET STARTED on the welcome screen. (NOTE: The following screenshot instructions may vary from the actual deployment configuration used in this reference architecture. These deployment images are shown for general awareness purposes only.)

Figure 8: VxRail Installation.

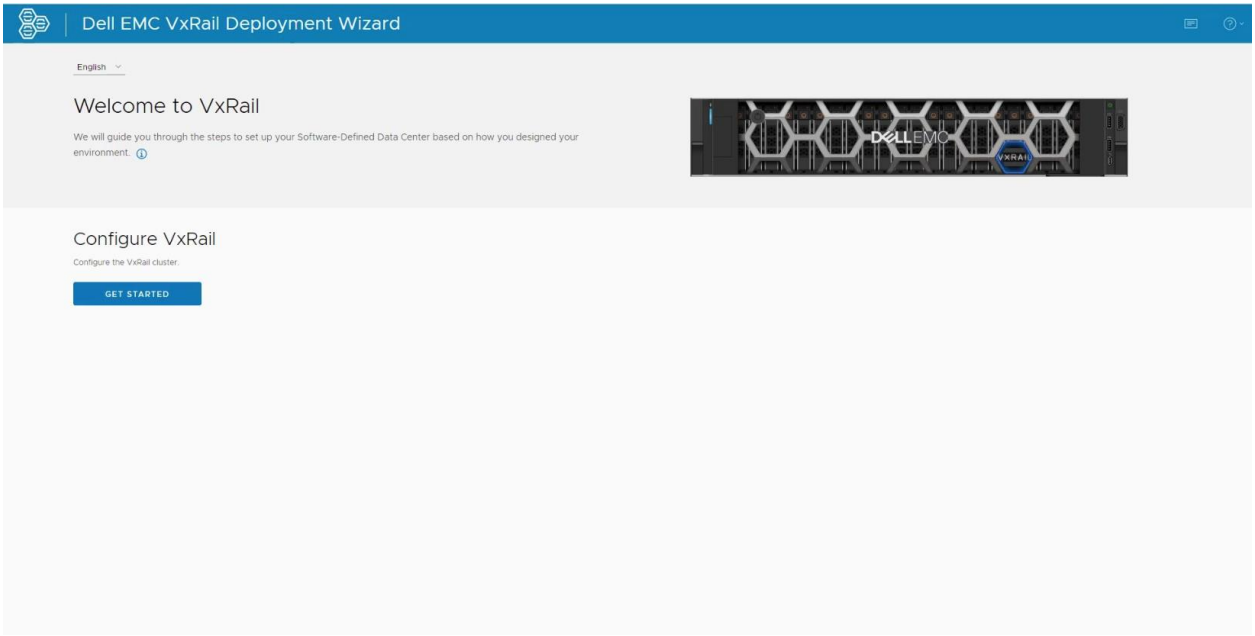

- On the EULA page, review the terms provided, and if you agree, click ACCEPT.

Figure 9: Accept EULA.

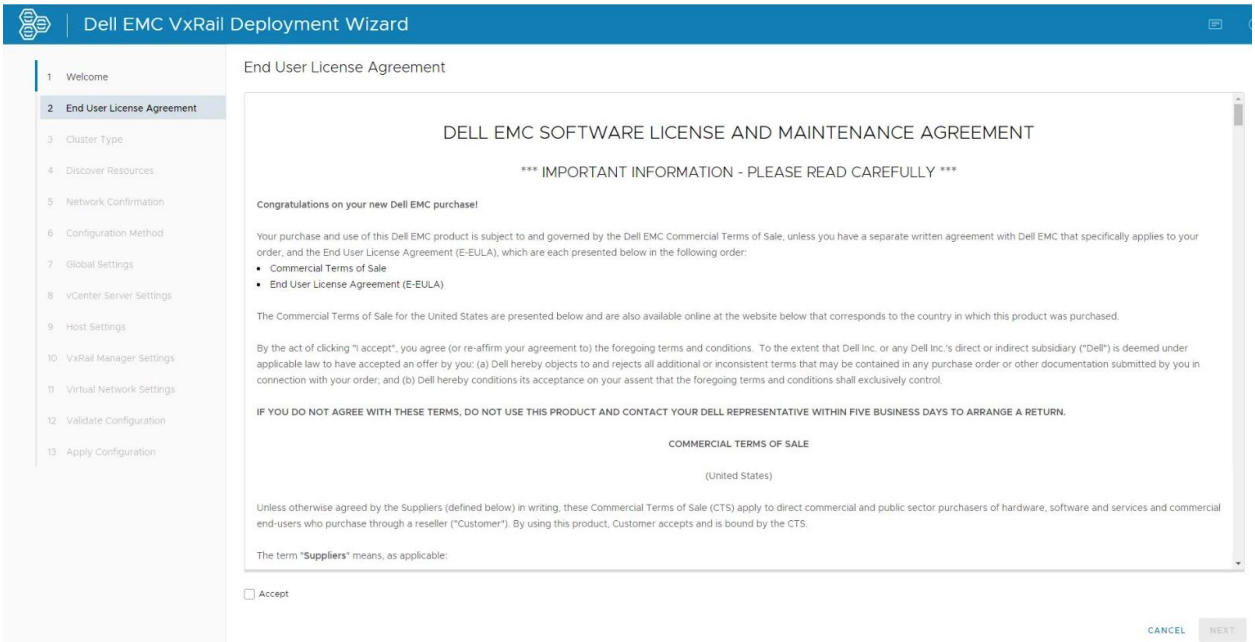

- The VxRail cluster type page displays. Select Standard Cluster. This will deploy a VxRail with vSAN HCI cluster deployment type which is what we are using in this reference architecture.

Figure 10: Select cluster type.

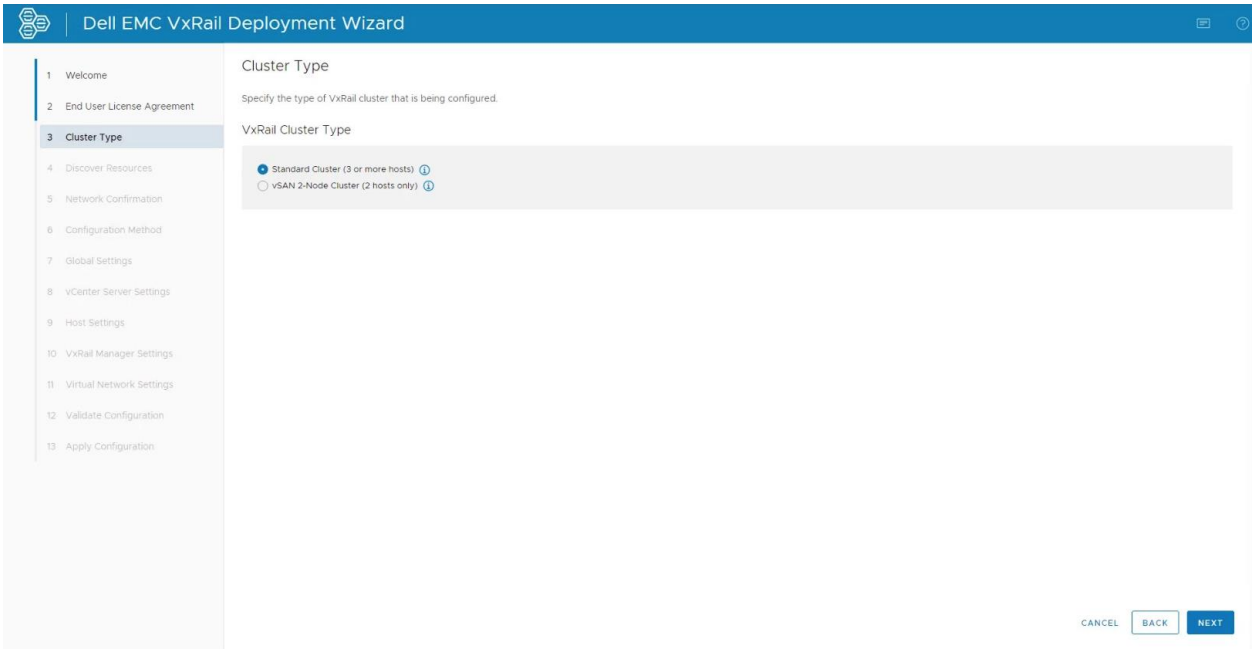

- When all VxRail nodes are discovered click NEXT.

Figure 10: VxRail nodes discovered.

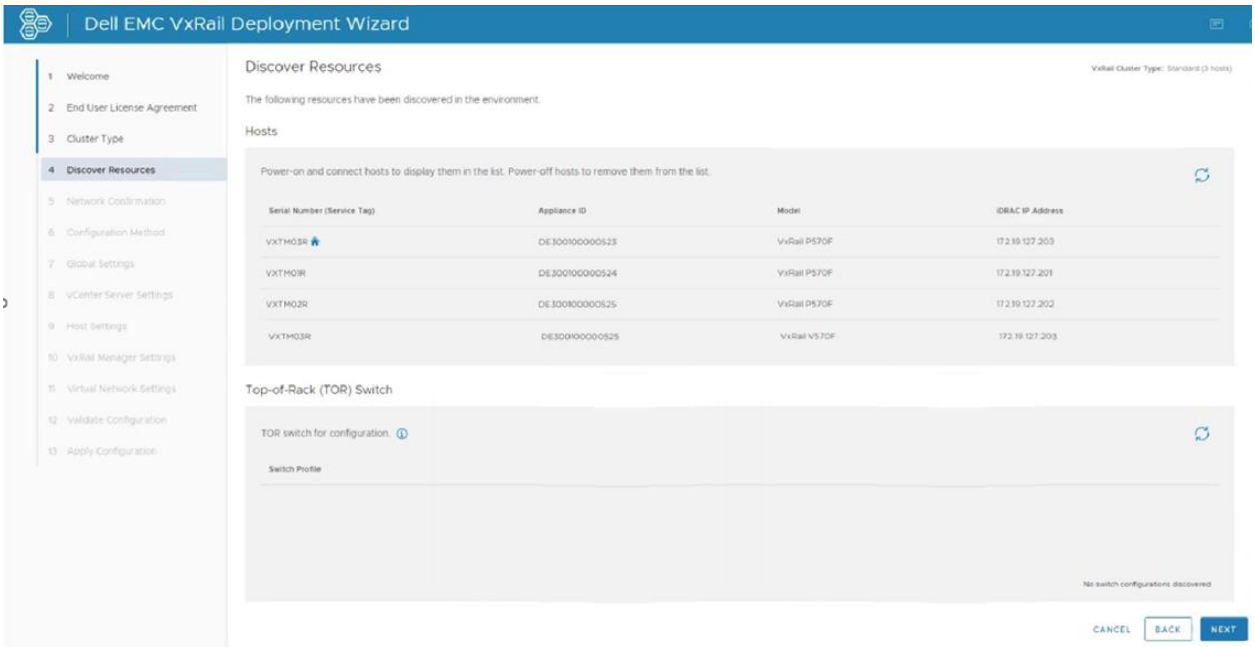

- Acknowledge that network is configured per best practices, by checking the two boxes.

Figure 11: VxRail network confirmation.

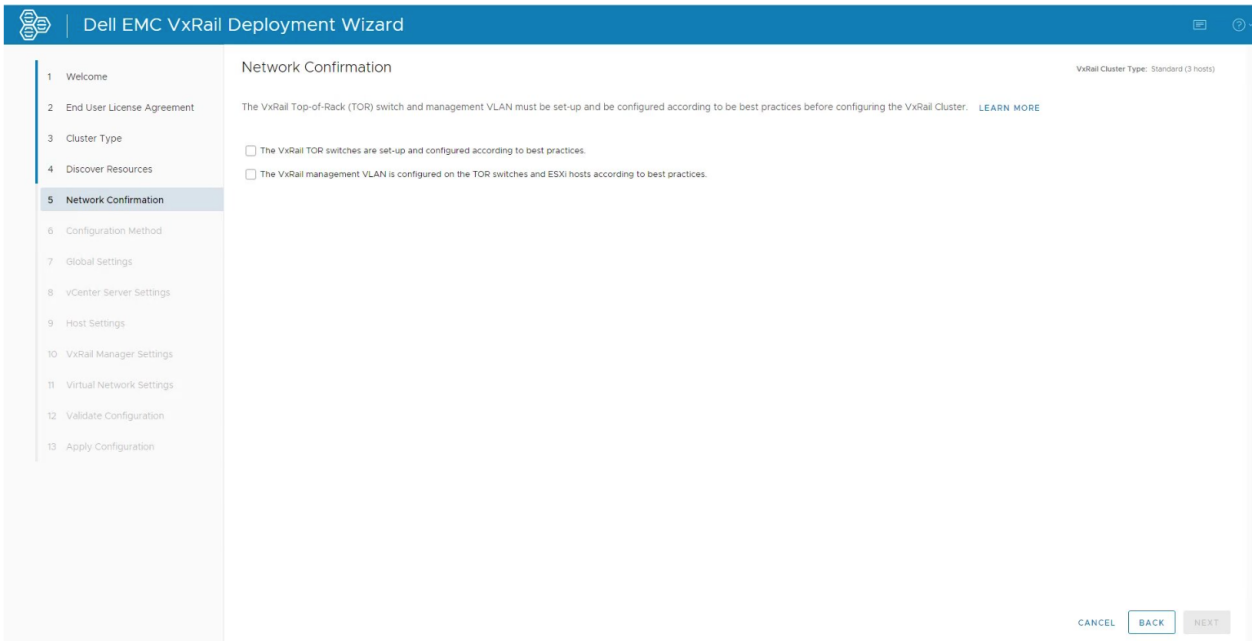

- There are two configuration methods. Users can choose to provide inputs for each step of the process or use a preconfigured JSON file. A JSON file with preconfigured cluster configuration parameters was used for this reference architecture. Click upload.

Figure 12: VxRail configuration methods.

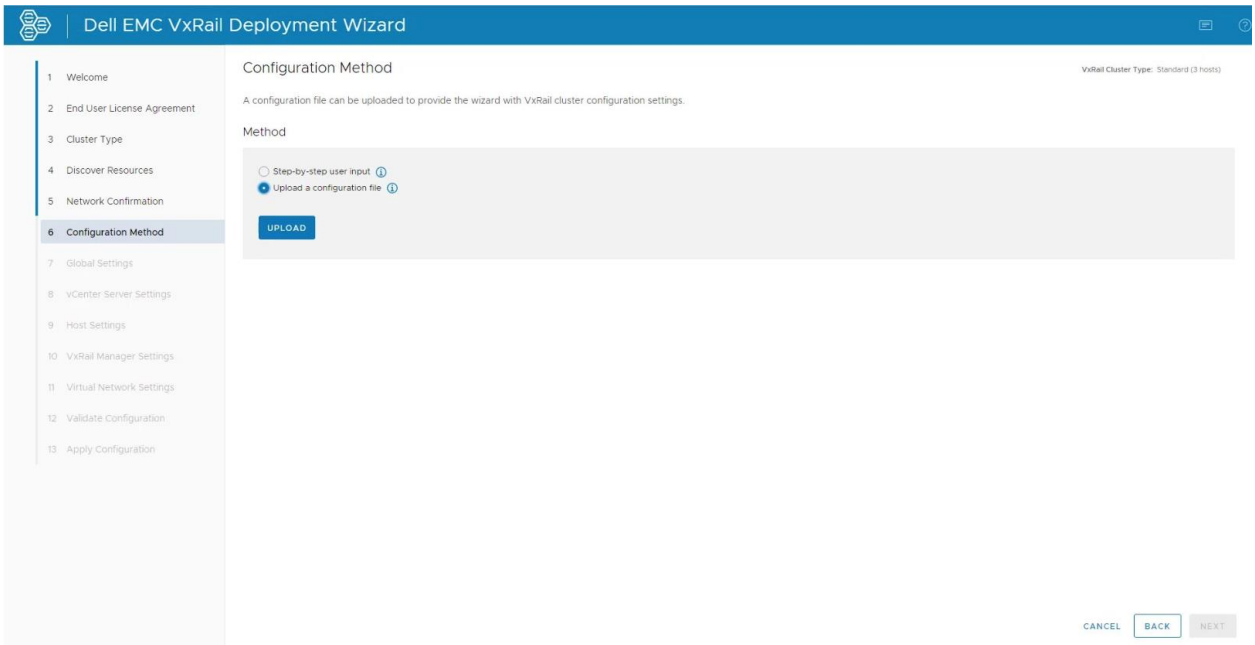

- Browse and select the preconfigured JSON file. Click "Open" to upload VxRail configuration file.

Figure 13: VxRail configuration file.

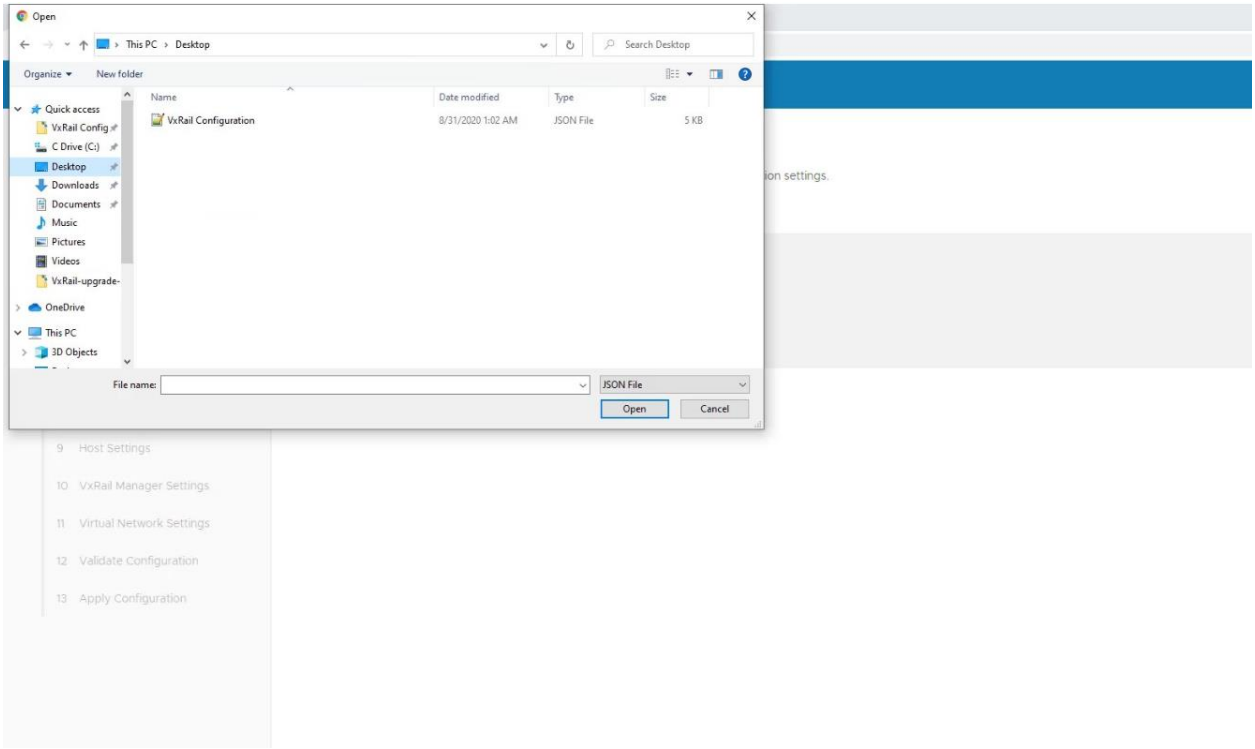

- All required cluster configuration parameters have now automatically populated. Validate and confirm the global settings and click NEXT.

Figure 14: VxRail global settings.

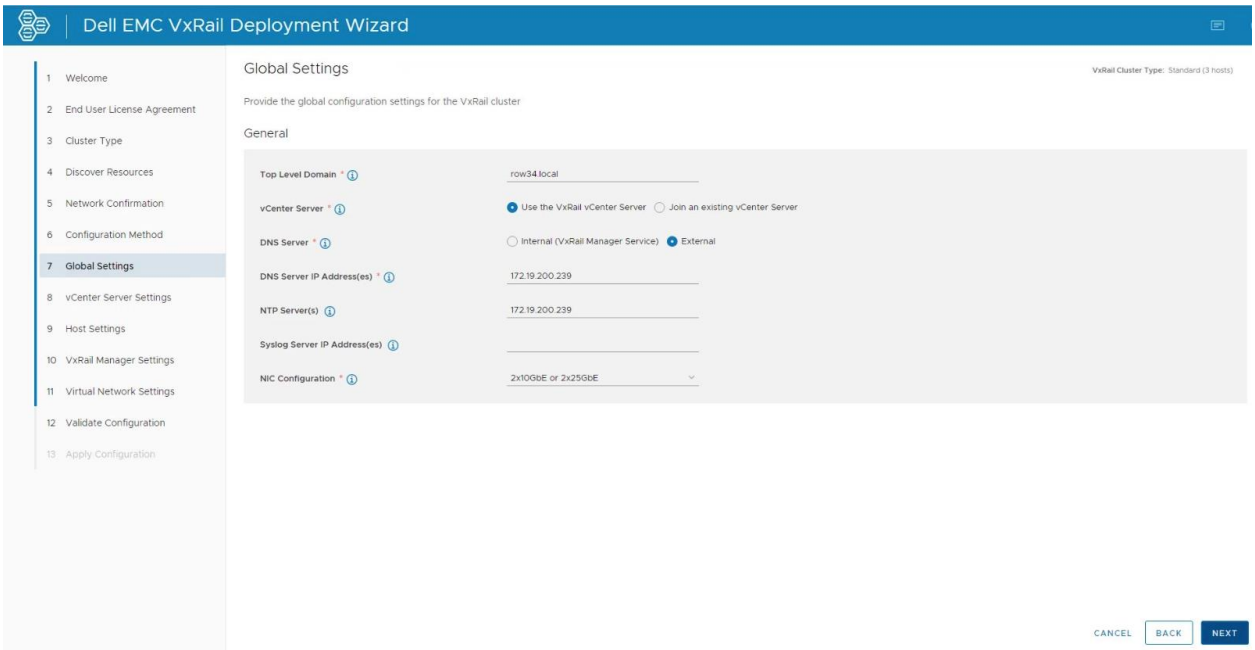

- Validate and confirm the vCenter settings and click NEXT.

Figure 15: VxRail vCenter settings.

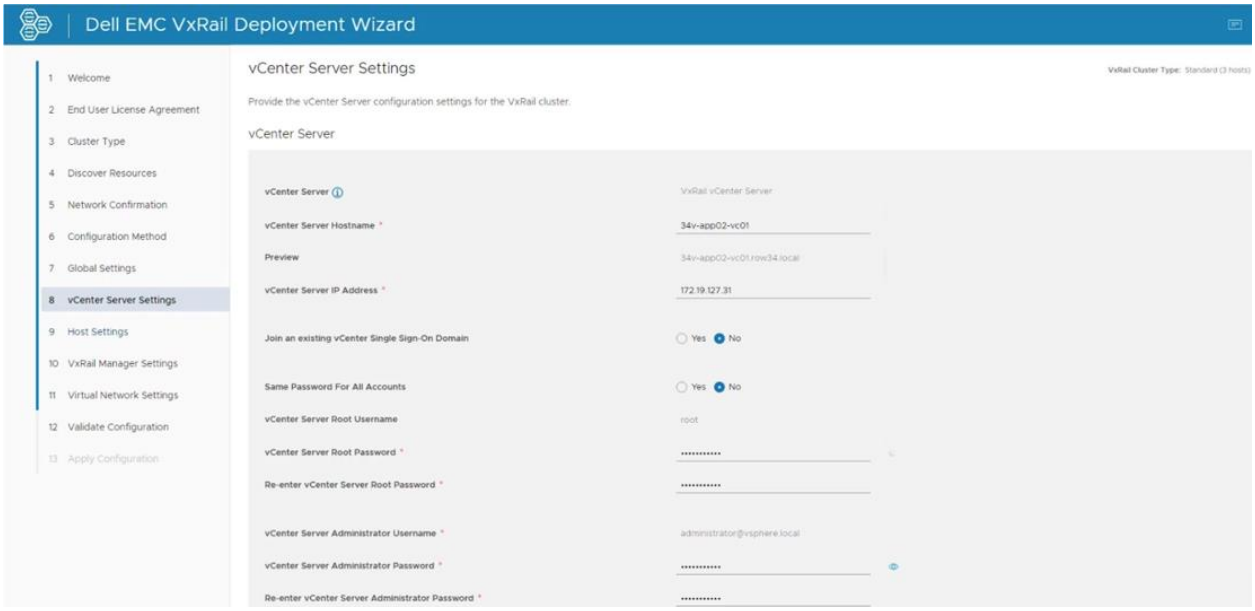

- Validate and confirm the individual ESXi hosts settings and click NEXT.

Figure 16: VxRail host settings.

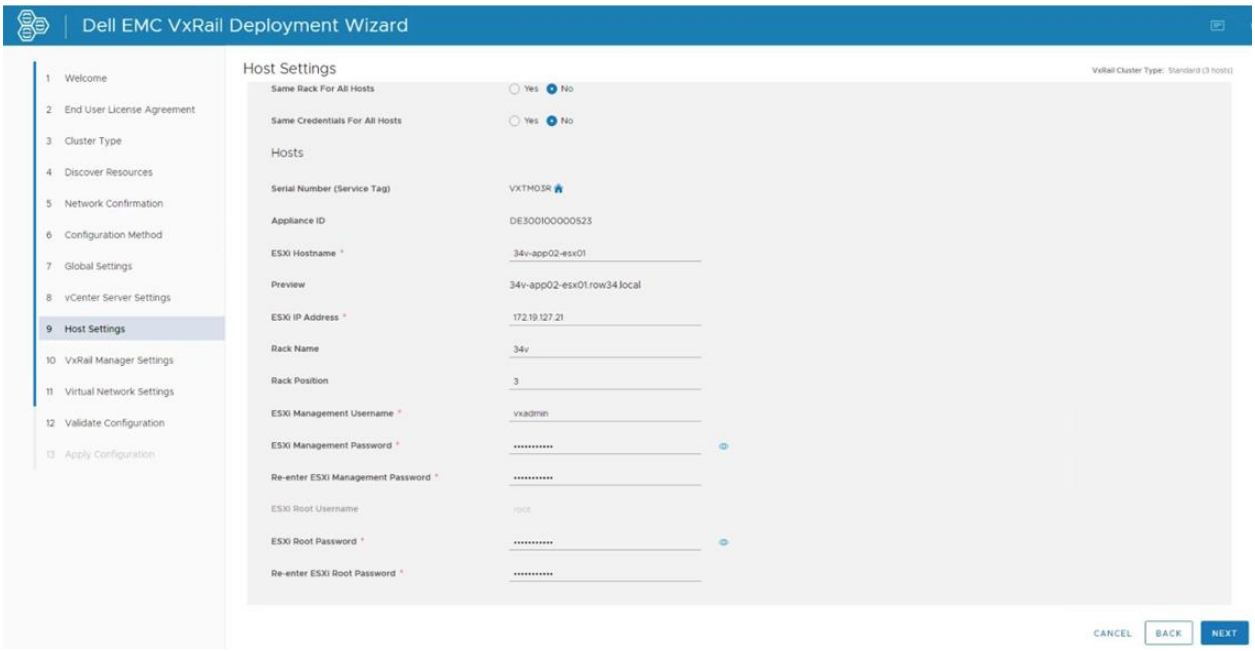

- Validate and confirm the VxRail Manager settings and click NEXT.

Figure 17: VxRail Manager settings.

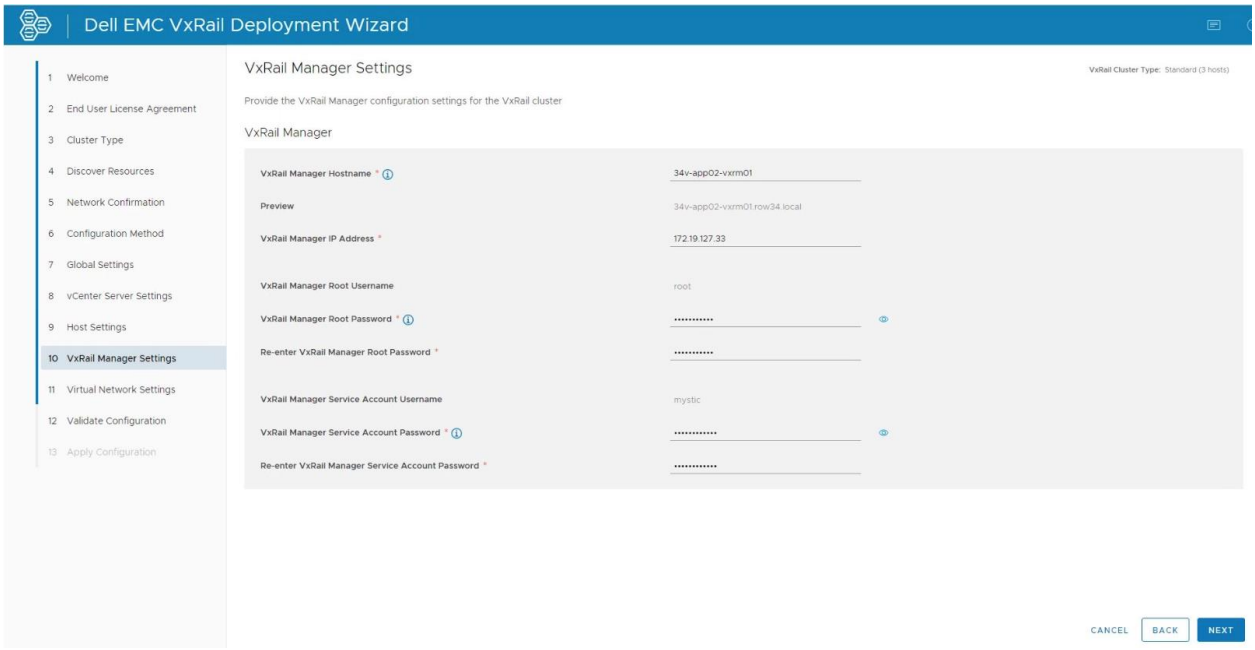

- Validate and confirm the virtual network settings. Click NEXT.

Figure 18: VxRail virtual network settings.

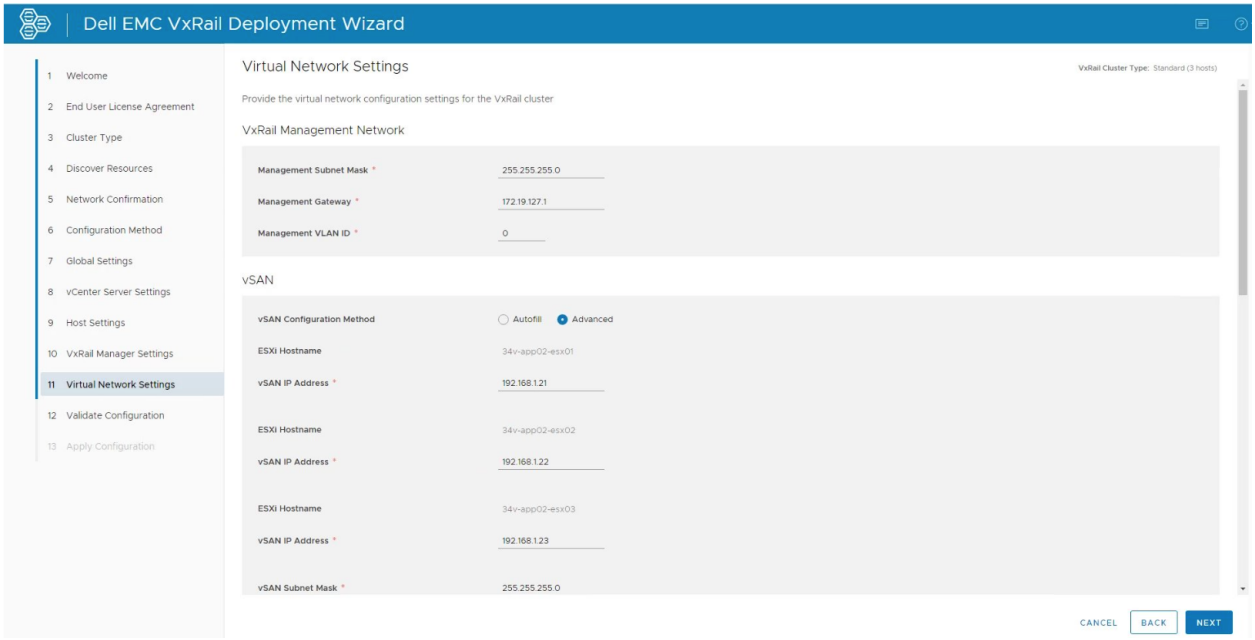

- Next, execute an automated VxRail Manager cluster configuration validation check to ensure all parameters have been entered correctly. Click on the Validate Configuration button.

Figure 19: VxRail validate settings.

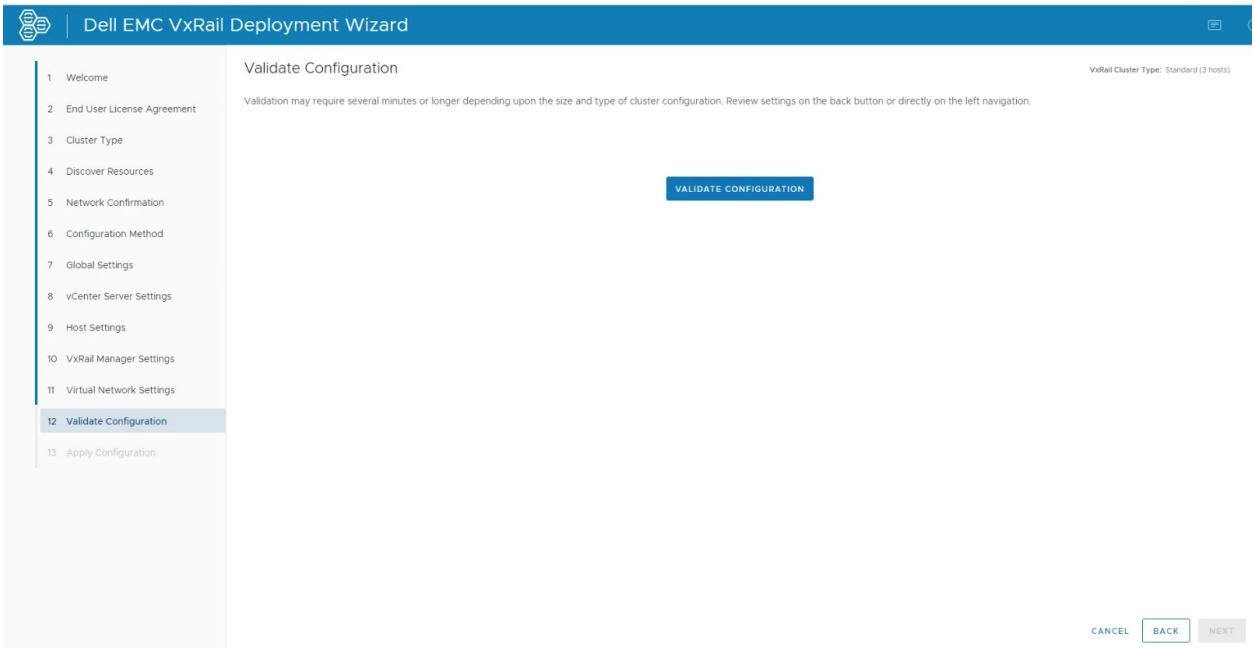

- VxRail Manager will automatically validate the configuration input provided on previous screens.

Figure 20: VxRail validation process.

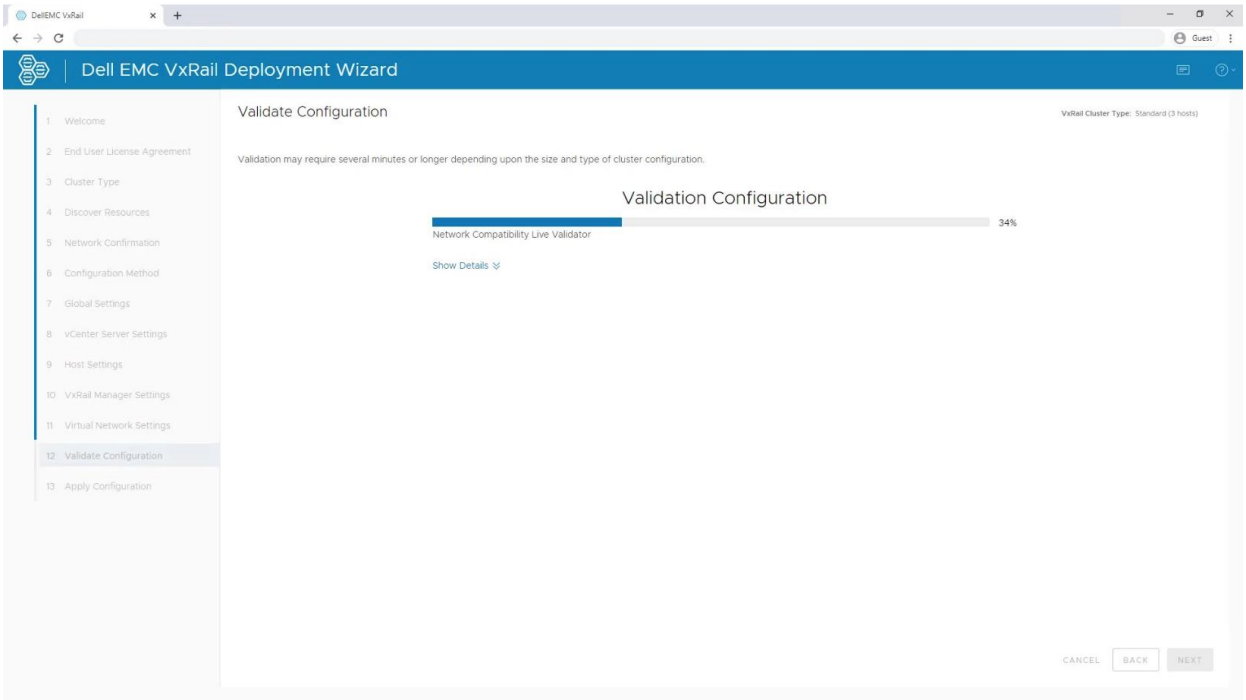

- The configuration JSON file can be downloaded once the validation has completed.

Figure 21: VxRail validation complete.

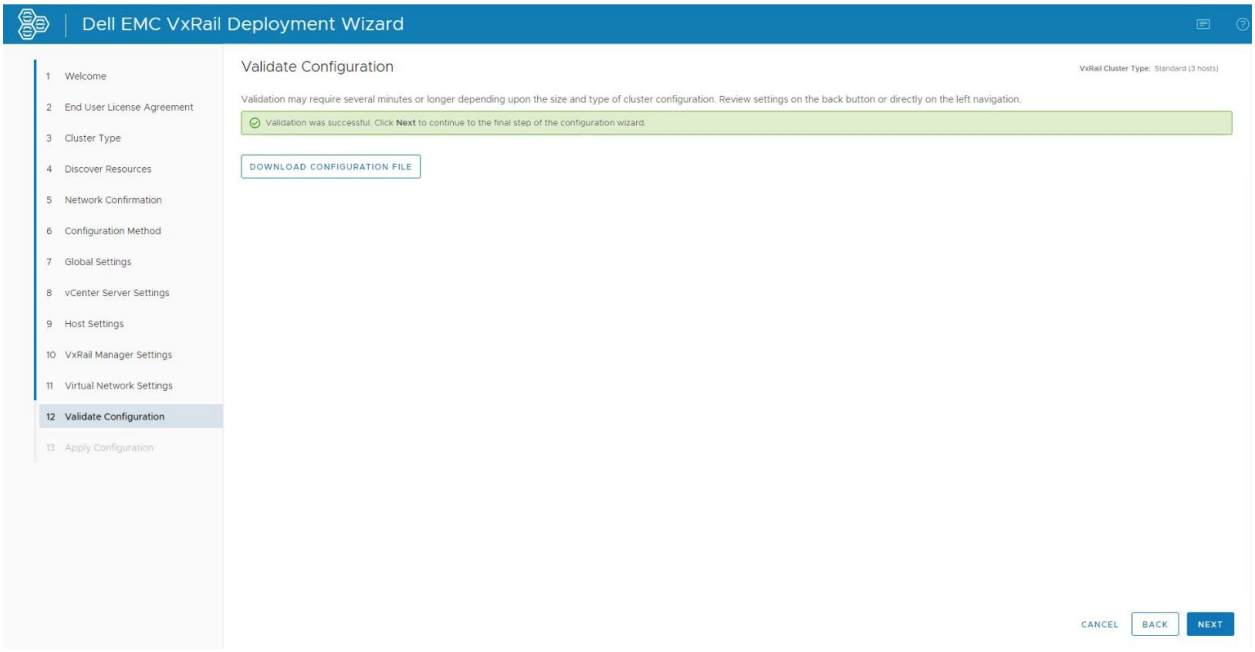

- Once the configuration validation has completed successfully, the cluster configuration can be used to automatically create the cluster.

Figure 22: VxRail apply configuration.

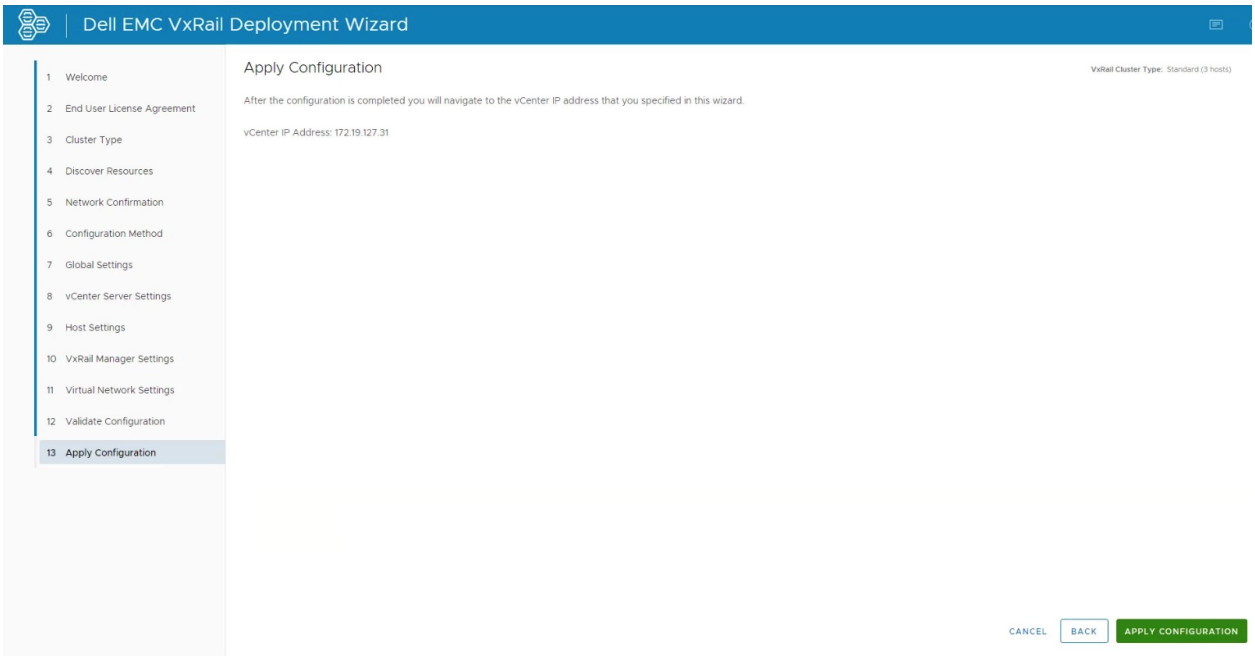

- The process will take a few minutes to complete.

Figure 23: VxRail applying configuration.

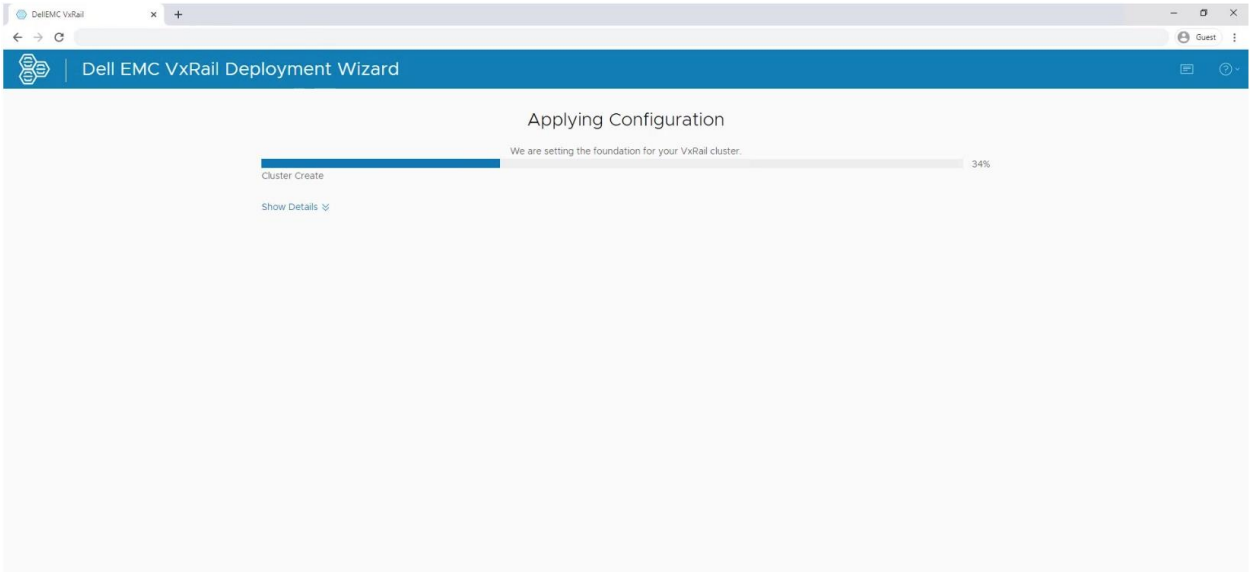

- After a few minutes, the VxRail Installation completes and vCenter (configured with the VxRail Manager vCenter Plugin) can be accessed by clicking "LAUNCH VCENTER" button.

Figure 24: VxRail Installation complete.

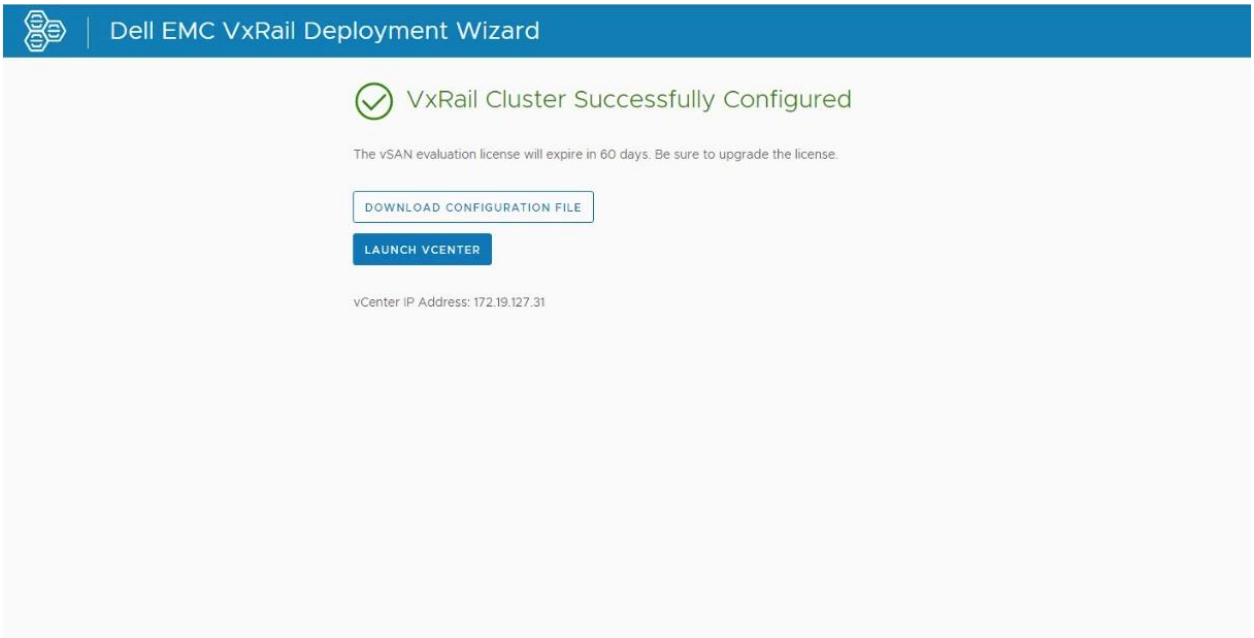

VMware NSX Advanced Load Balancer

NSX Advanced Load Balancer (formerly known as Avi) comes in two editions, Essentials, and Enterprise. To use L7 load balancing with NSX Advanced Load Balancer, the Enterprise edition is required and was used for this reference architecture. The NSX Advanced Load Balancer provides dynamically scaling load balancing endpoints for Tanzu Kubernetes clusters provisioned by the Tanzu Kubernetes Grid Service. Once you have configured the Controller, it automatically provisions load balancing endpoints for you. The Controller creates a virtual service and deploys Service Engine VMs to host that service. This virtual service provides load balancing for the Kubernetes control plane. NSX Advanced Load Balancer has some key components that are explained below.

NSX Advanced Load Balancer Controller: As the name suggests, NSX Advanced Load Balancer Controller controls and manages the provisioning of service engines, coordinating resources across service engines, and aggregating service

engine metrics and logging. It interacts with vCenter Server to automate the load balancing for Kubernetes clusters. It is deployed as an OVA and provides a Web interface and CLI.

NSX Advanced Load Balancer Service Engine: Service Engine runs one or more virtual services and is a data plane component and runs as a virtual machine. Service engines are provisioned and controlled by the controller. The service engines have two interfaces. One connects to the NSX Advanced Load Balancer Controller management network and the second connects to the front-end network from where virtual services are accessed. For service engine sizing guidance, see Sizing Service Engines.

Avi Kubernetes Operator (AKO): Avi Kubernetes Operator runs as a Kubernetes POD in the Supervisor, management, and workload clusters to provide ingress and load balancing.

Avi Multi-Cluster Kubernetes Operator (AMKO): Avi Multi-Cluster Kubernetes Operator runs in a pod in the Tanzu GSLB leader cluster. In conjunction with Avi Kubernetes Operator, Avi Multi-Cluster Kubernetes Operator facilitates multi-cluster application deployment. It maps the same application deployed on multiple clusters to a single GSLB service, extending application ingresses across multi-region and multi-availability zone deployments.

Steps to configure NSX Advanced Load Balancer

NSX Advanced Load Balancer Controller is deployed as VM using the OVA that can be downloaded from https://customerconnect.vmware.com/ using an account that has access to downloading software packages. Once the OVA is downloaded import it to vCenter. For this reference architecture version 20.1.7 was used with Enterprise license. The process of deploying and configuring NSX Advanced Load Balancer follows. Screenshots are used where necessary to emphasize specific configurations.

Prerequisites:

- VxRail cluster is already installed and configured.
- vSphere distributed switch port groups for required networks are created as described in network overview section previously.
- A resource pool is created in vCenter that will hold the NSX Advanced Load Balancer virtual machines.

Controller Deployment: On-Premises

1. Import controller OVA and provide a name for the controller.

Figure 25: Import OVA

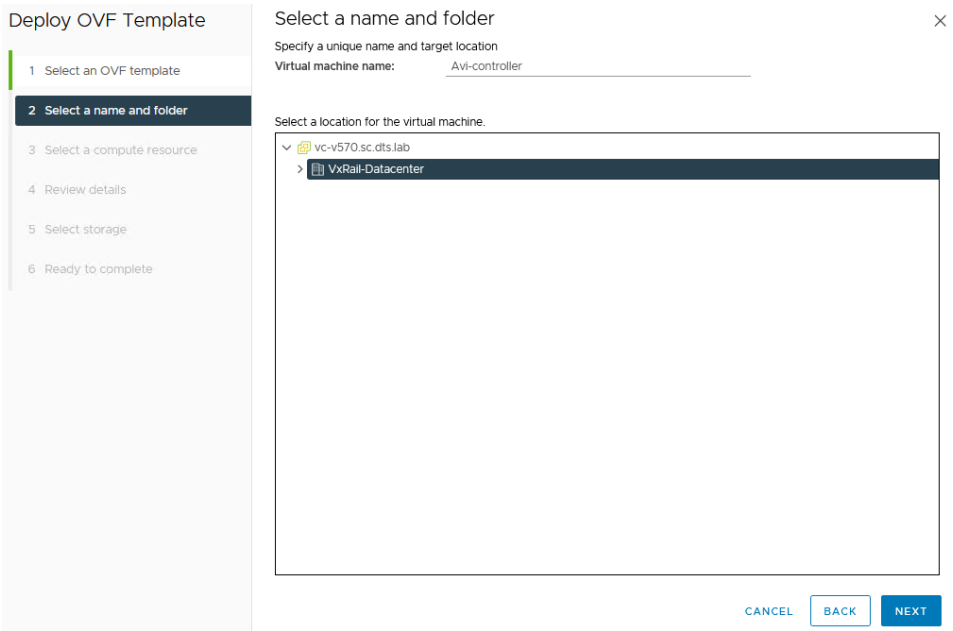

2. Select the resource pool for the controllers.

Figure 26: Select resource pool

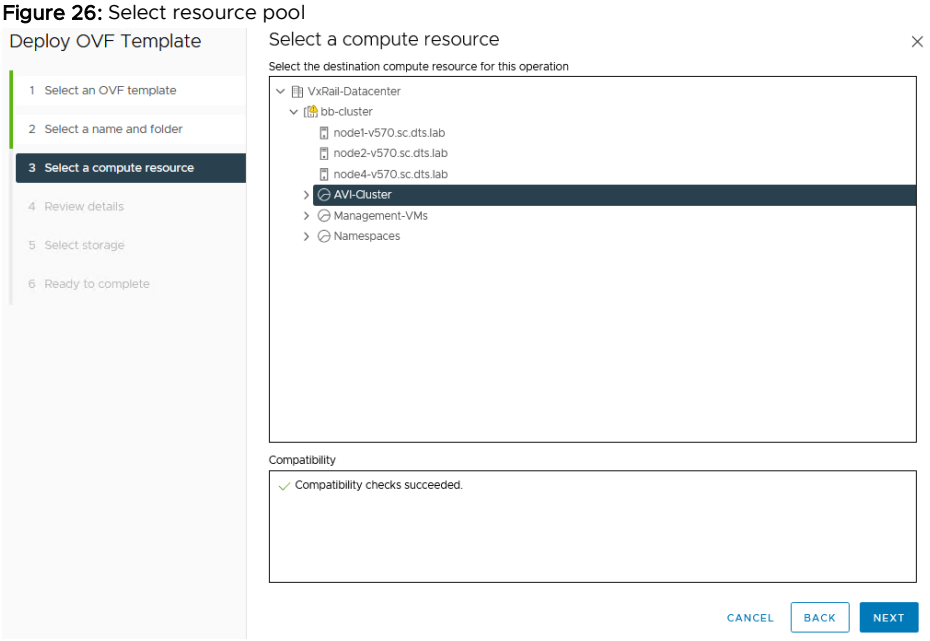

3. For storage, select vSAN datastore that was created during VxRail deployment.

Figure 27: Select storage

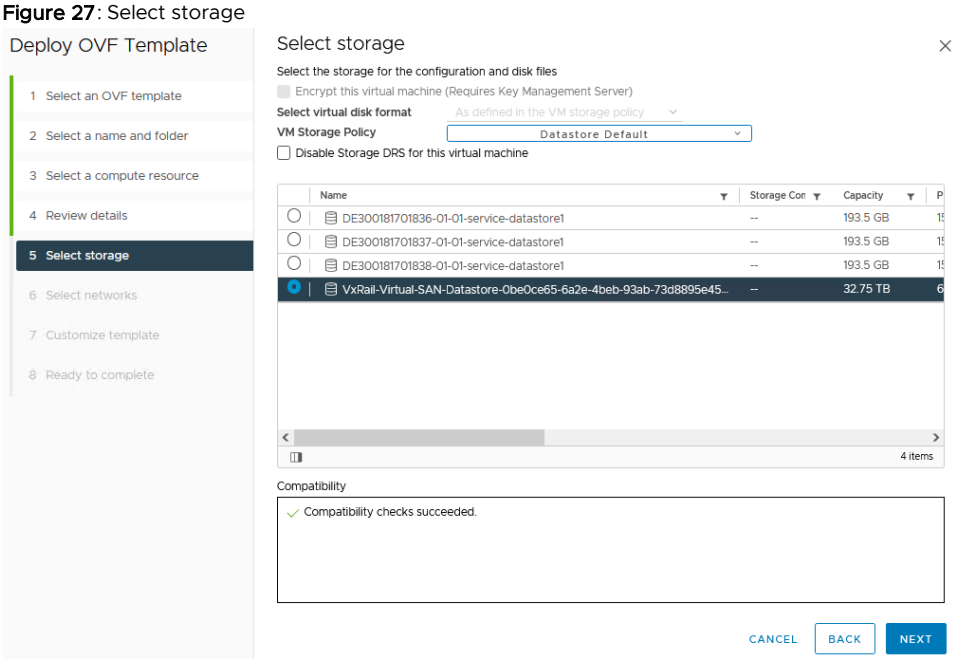

4. Select VDS port group that is designated for NSX Advanced Load Balancer management interface. Ensure that the port group to which NSX Advanced Load Balancer is attached, can communicate with port group to which vCenter Server management network resides.

Figure 28: Select management interface

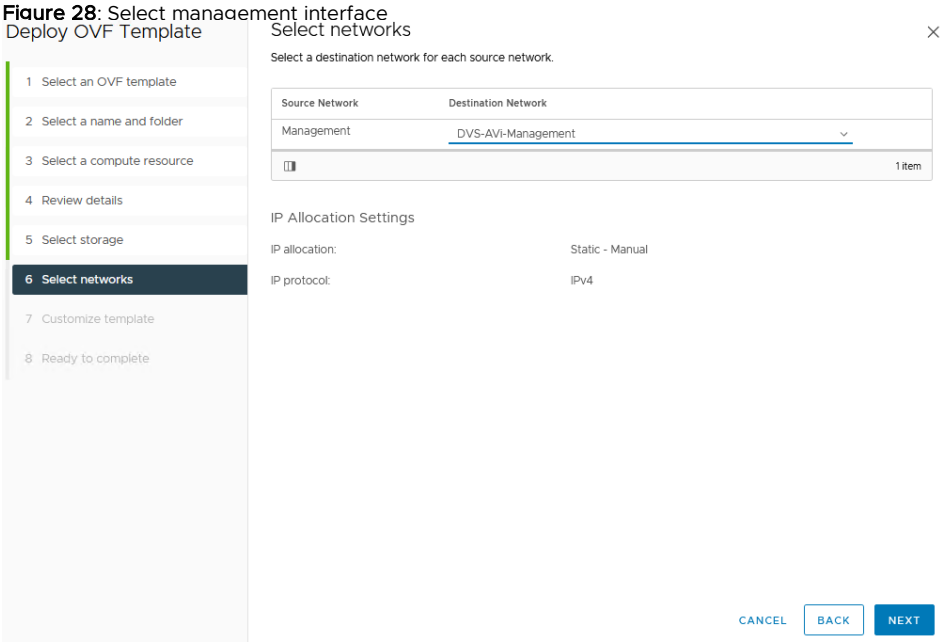

5. On the next screen enter the required information and proceed to finish on the next screen.

Figure 29: Required information

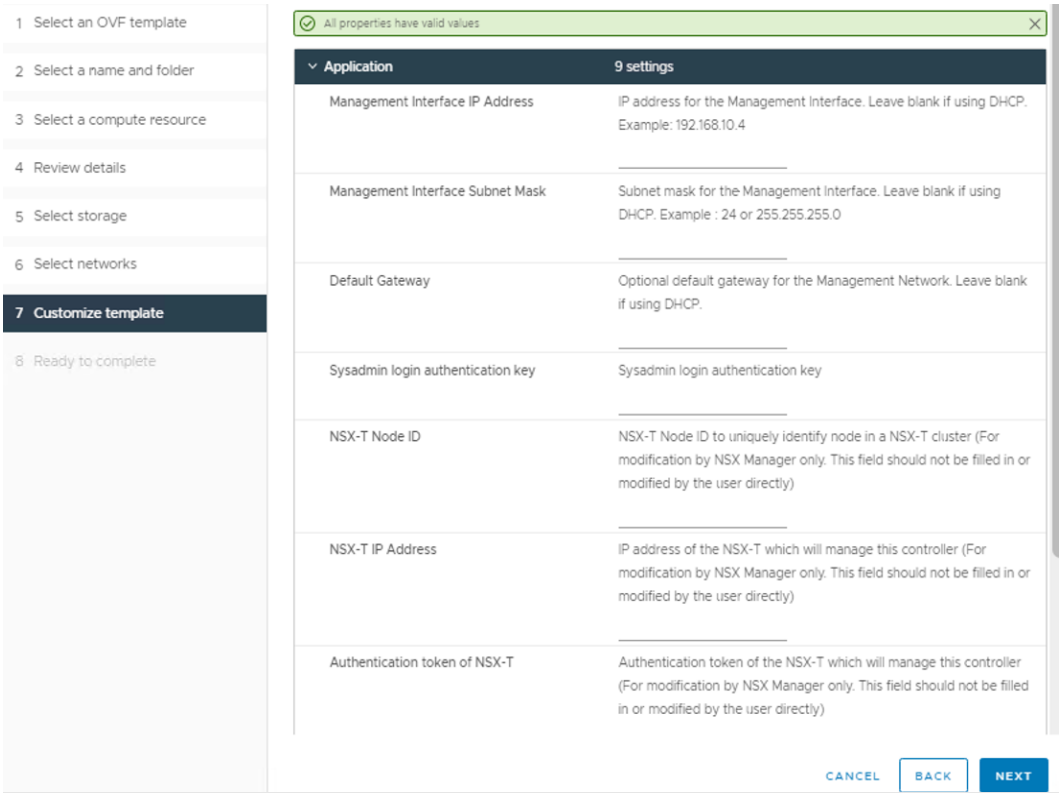

Figure 30: Complete OVF deployment

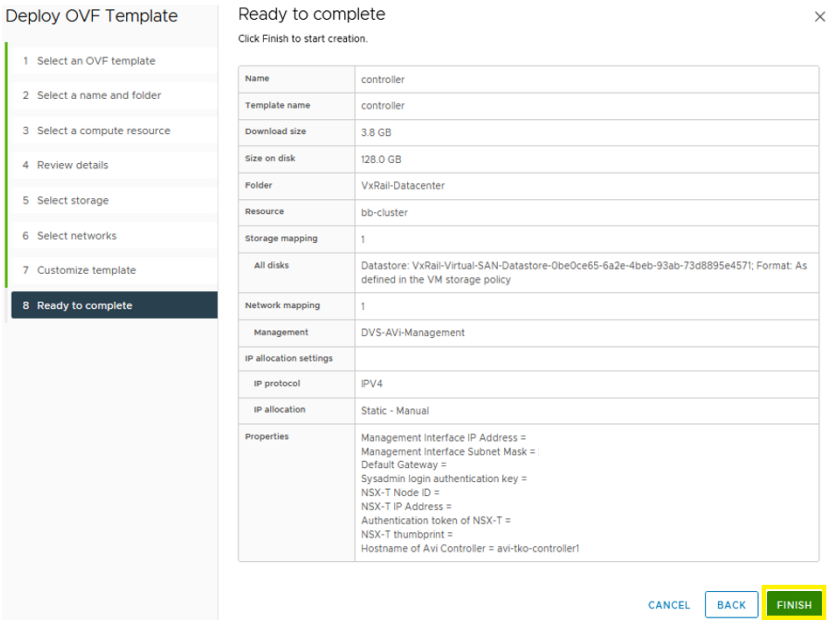

Login and Initial Configuration

Once the controller is deployed and ready, access the admin portal from a browser using the previously configured hostname or IP address. Please note that it takes a few minutes for the controller to be available for login.

1. On the login screen create a new password and create the admin account
2. On the next screen fill in the system settings including choosing a passphrase, DNS, domain, and SMTP information. Choose your multi-tenant settings and click save. Figure 9 depicts this screen.

Figure 31: Initial login setup

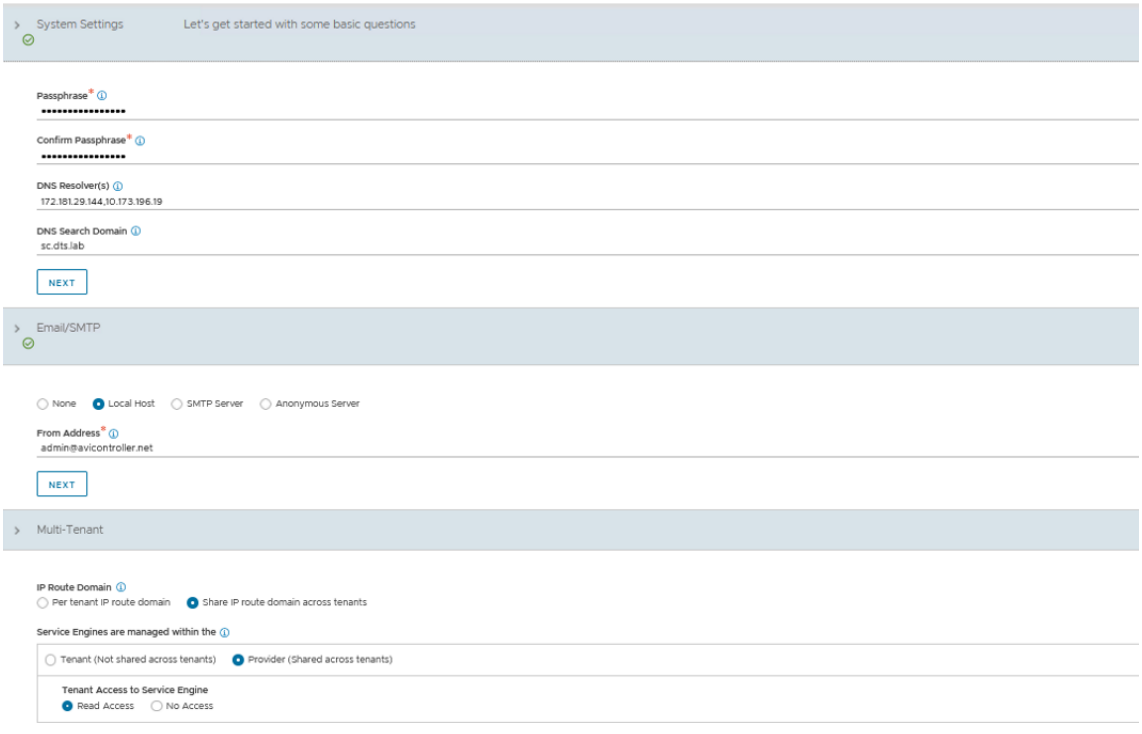

Controller Configuration

Once the controller is deployed, several tasks need to be performed for NSX Advanced Load Balancer to work with Tanzu Kubernetes Grid Service. These tasks are summarized below.

1. Configure the default cloud instance.
2. Configure settings for system access.
3. Configure Service Engine group.
4. Configure the VIP network.
5. Create IPAM Profile
6. Add IPAM to Default-Cloud instance.
7. Create DNS service.
8. Add IPAM and DNS Profile to the default cloud instance.
9. Export the SSL/TLS certificate.
10. Create route between workload and front-end networks.

Configure default cloud instance.

Currently only Default-Cloud instance is supported on vSphere with Tanzu. Access the Default-Cloud settings via Infrastructure ->Clouds and click the pencil icon to edit the cloud configuration.

1. On the Infrastructure tab, fill in the IP address, username, and password for vCenter Server. Ensure that under access permissions “Write” is selected. NSX Advanced Load Balancer requires write permissions to vCenter to creating, modifying, and removing Service Engines or other resources automatically as requirements change.

Figure 32: Add vCenter.

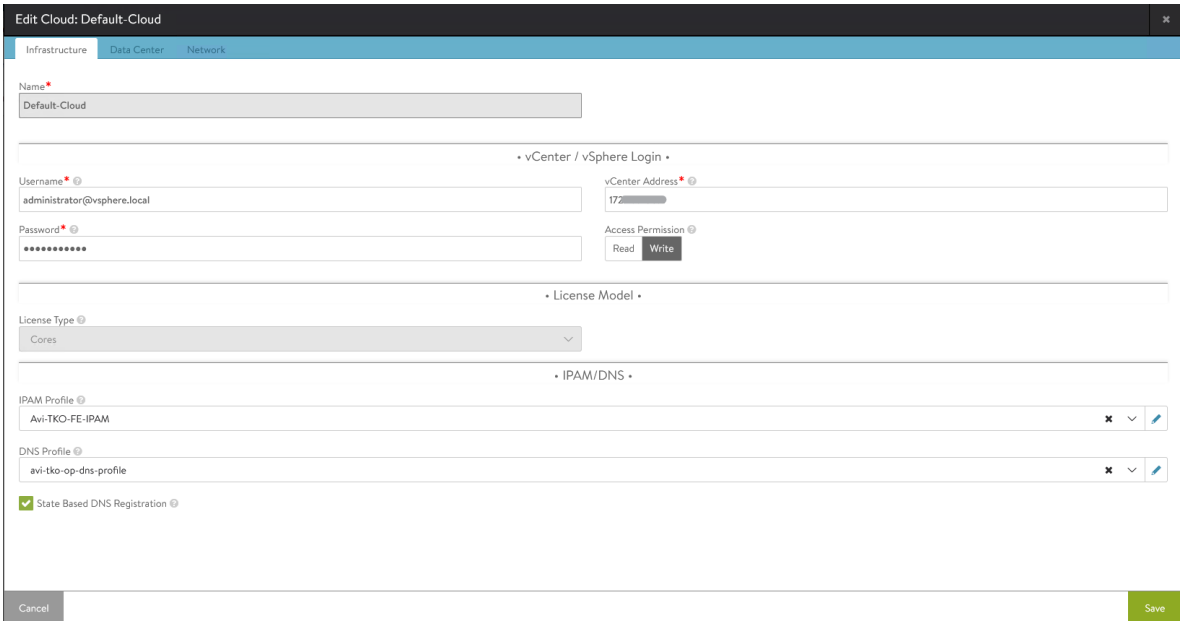

2. On the Data Center tab select your data center.

Figure 33: Select Data Center

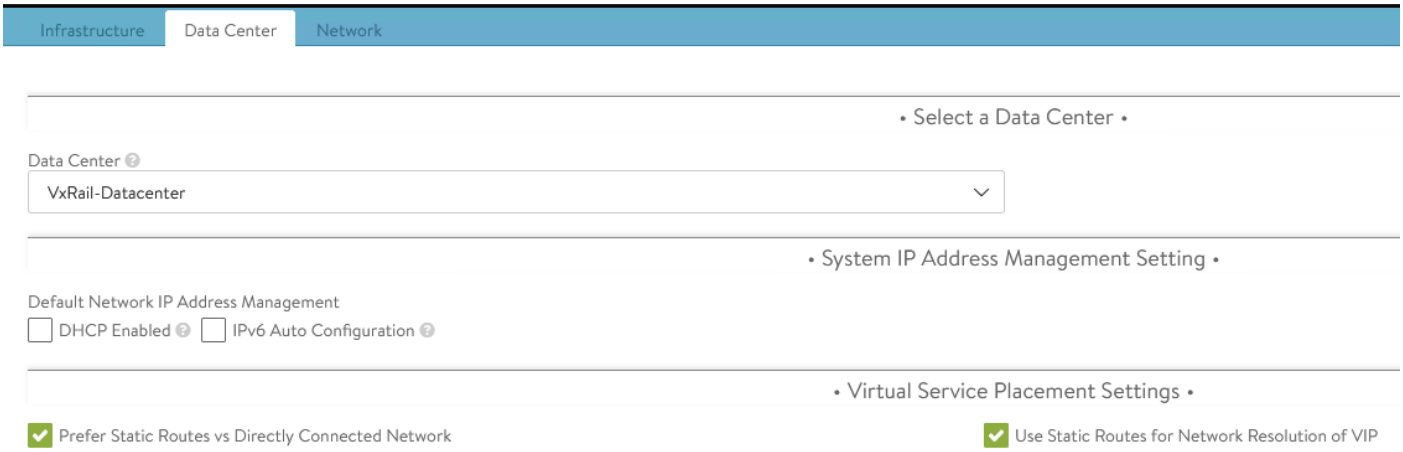

3. On the Network tab, select network designated for NSX Advanced Load Balancer management network, IP subnet, default gateway and an IP range for the static pool.

Figure 34: Select network.

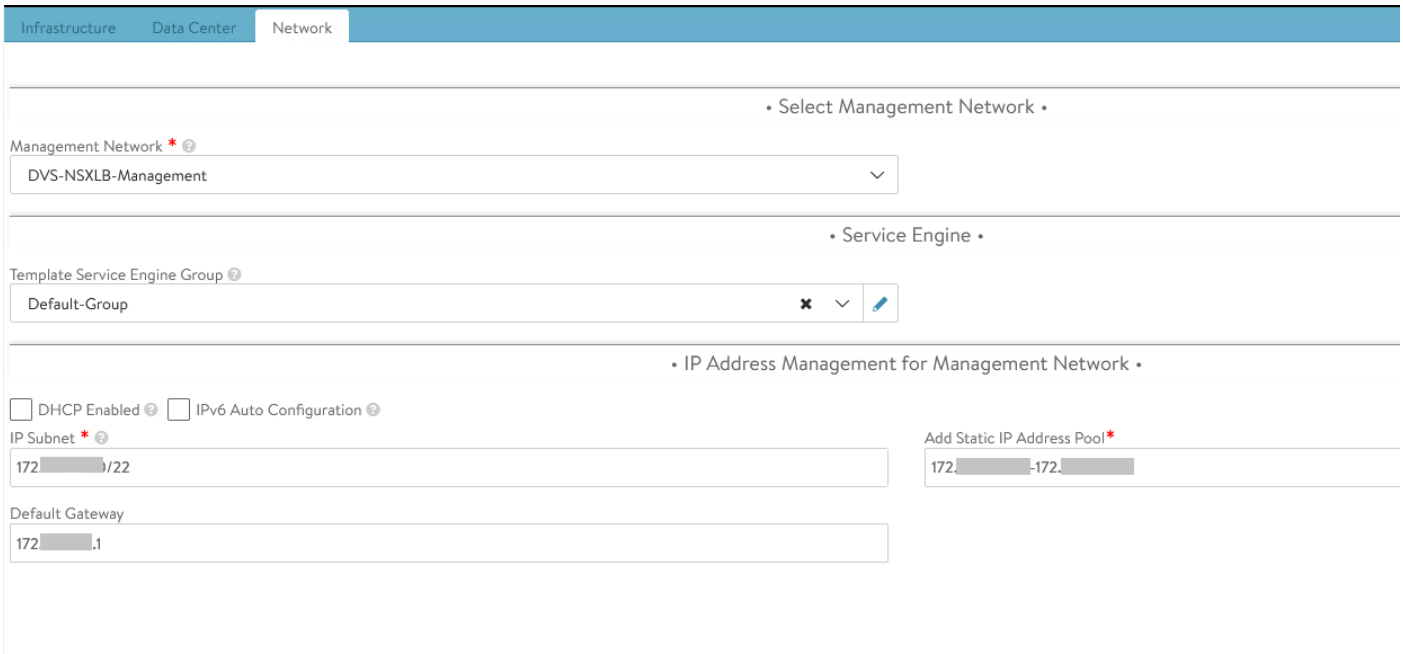

Configure settings for system access.

1. Basic authentication can be set using the following process. Navigate to Administrator > Settings > Access Settings and check “Allow Basic Authentication.”

Figure 35: Set basic authentication.

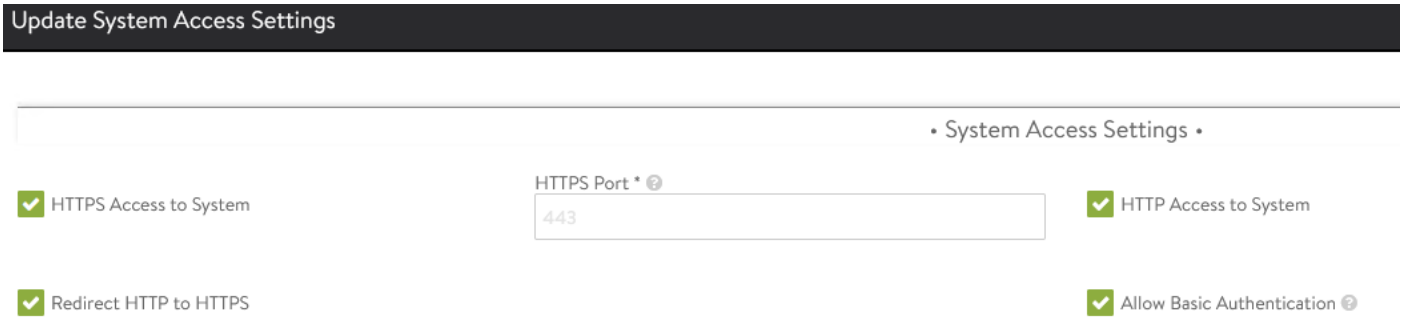

Staying on the same screen delete the existing SSL\TLS certificate and create a new one with your specific organization information. The Controller has a default self-signed certificate. But this certificate does not have the correct SAN (Subject Alternate Name). Certificate must be replaced with a valid external or self-signed certificate that has the correct SAN. Step-by-step instructions visit [NSX Advanced Load Balancer documentation] page.

Figure 36: Create certificate.

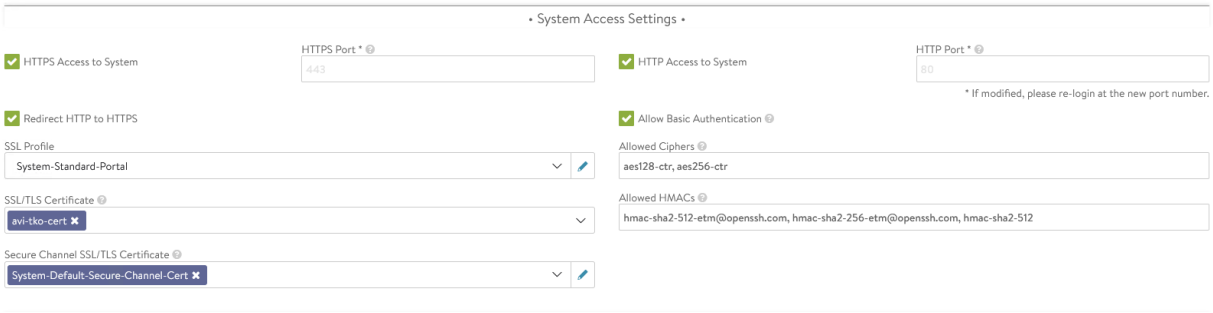

Configure Service Engine group.

1. From Infrastructure > Service Engine Group > Basic Settings, ensure that N+M (buffer) is selected under elastic HA. This is the default mode, where “N” is the minimum number of Service Engines required to place virtual services in a SE group and “M” is the additional Service Engines that the controller spins up to manage Service Engine failures without reducing the capacity of the group.

Figure 37: Create service engine group.

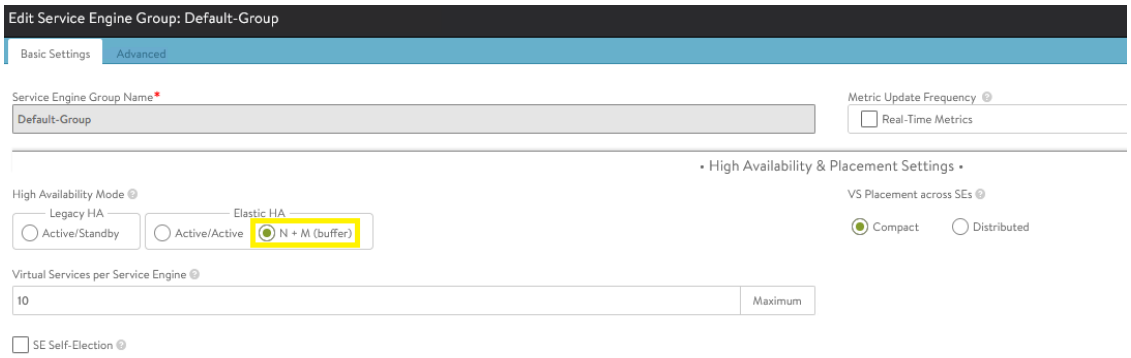

2. In the “Advanced” tab select your cluster and vSAN datastore

Figure 38: Select cluster and vSAN datastore.

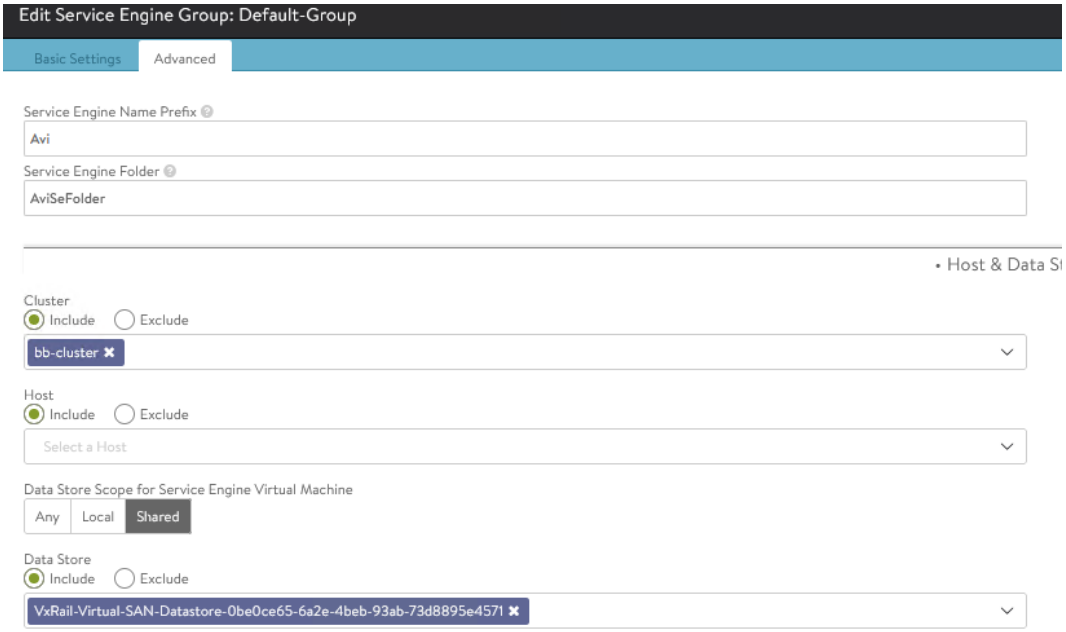

Configure the VIP network.

This network is where various Kubernetes control plane and Kubernetes applications require load balancing services. In this case the VIPs reside on the front-end network.

Figure 39: Configure VIP network.

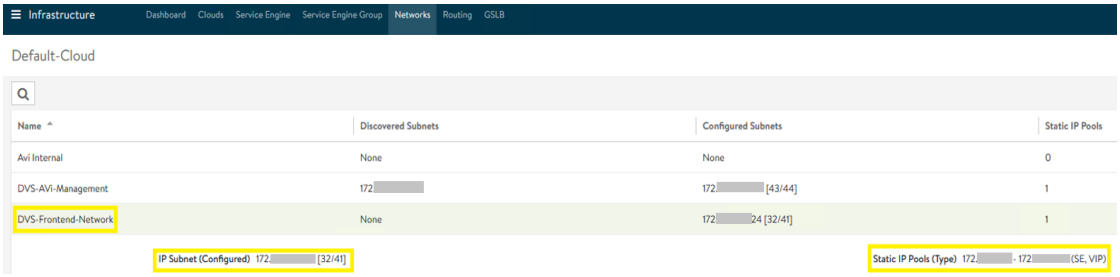

Create IPAM Profile

Create IPAM for the VIP network created earlier to assign IPs to the virtual services. This can be accessed via Templates > Profiles > IPAM/DNS Profiles.

1. Create a new IPAM Profile by using the Create button on the top right-hand side of the screen.

Figure 40: Create IPAM

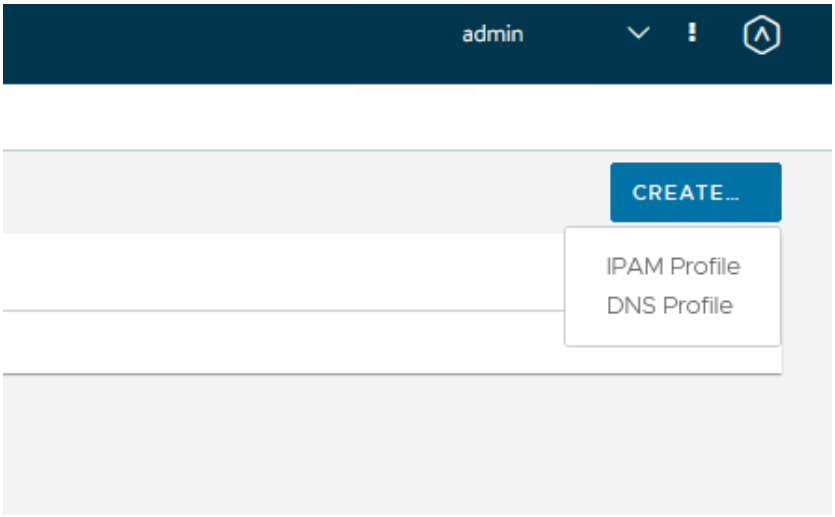

2. Enter a name. In the Type field, select Avi Vantage IPAM, and add a usable network which will be your VIP network.

Figure 41: Create IPAM

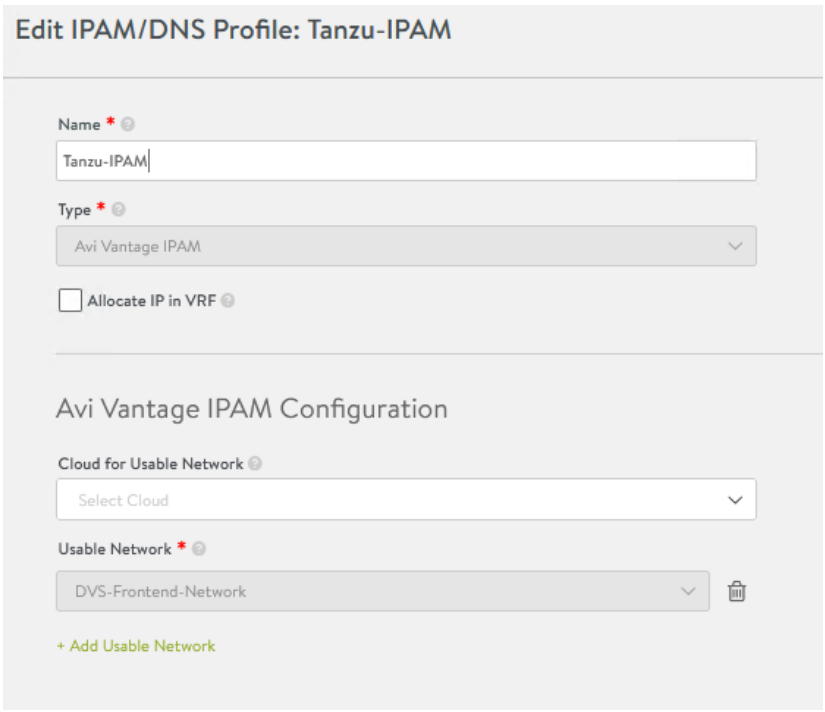

Add IPAM to Default-Cloud instance.

The new IPAM needs to be added to the default-cloud instance.

1. Navigate to Infrastructure > Default-Cloud and edit. Select the newly created IPAM profile from the dropdown list.

Figure 42: Add IPAM to Default-Cloud

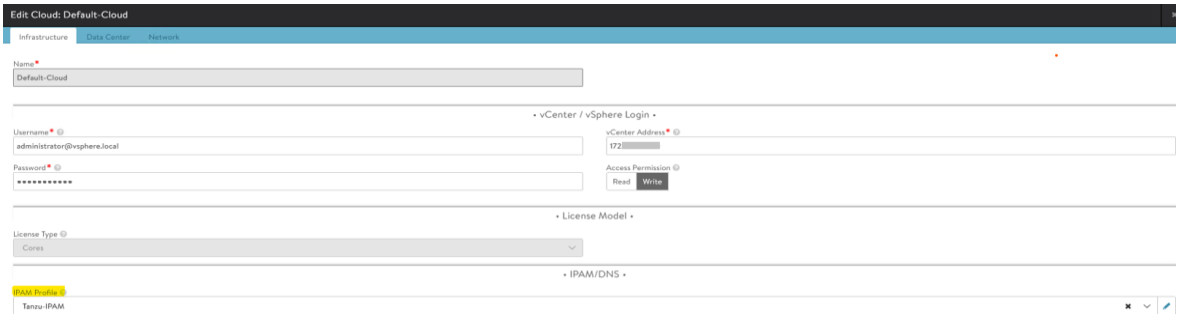

Create DNS Profile

Create DNS profile via Templates > Profiles > IPAM/DNS Profiles.

3. Create a new DNS Profile by using the Create button on the top right-hand side of the screen.

Figure 43: Create DNS profile.

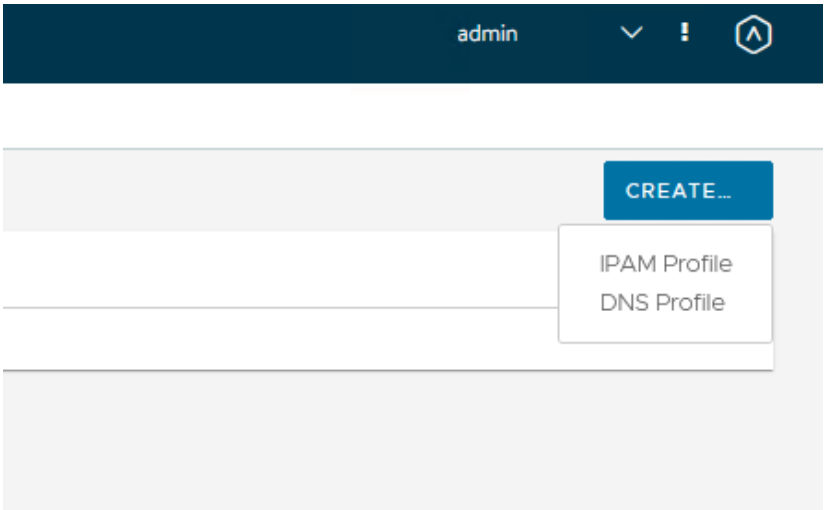

4. Enter a name. In the “Type” field, select Avi Vantage DNS, and add the desired sub-domain. This subdomain will be delegated to NSX Advanced Load Balancer.

Figure 44: Create DNS profile.

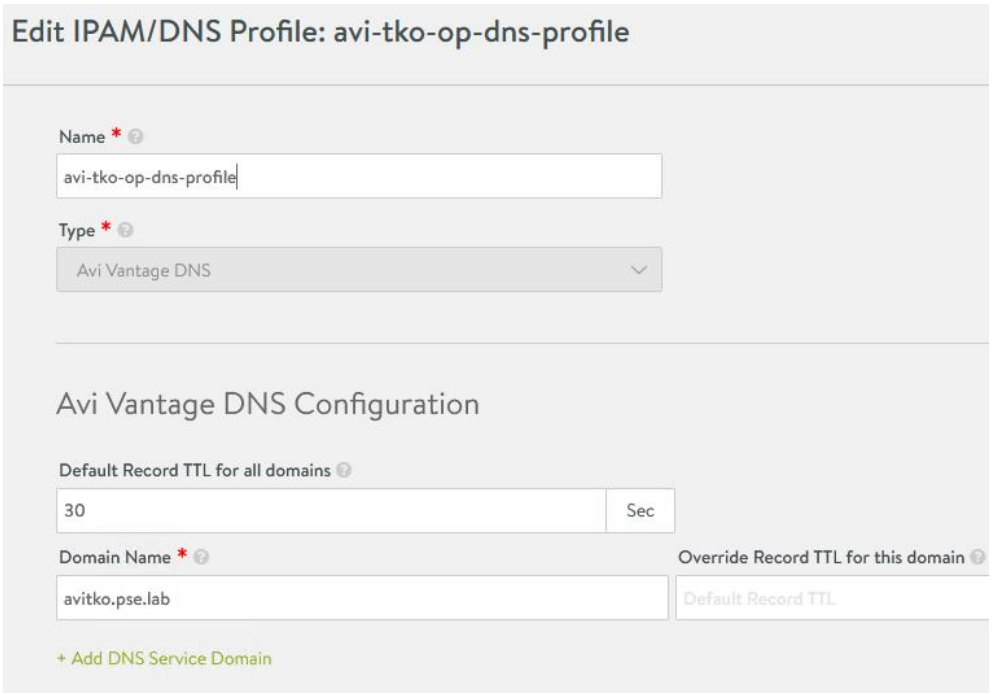

Add IPAM and DNS profile to Default-Cloud instance.

The new IPAM and DNS profile needs to be added to the default-cloud instance.

2. Navigate to Infrastructure > Default-Cloud and edit. Select the newly created IPAM and DNS profile from the dropdown list.

Figure 45: Add IPAM and DNS profiles to Default-Cloud

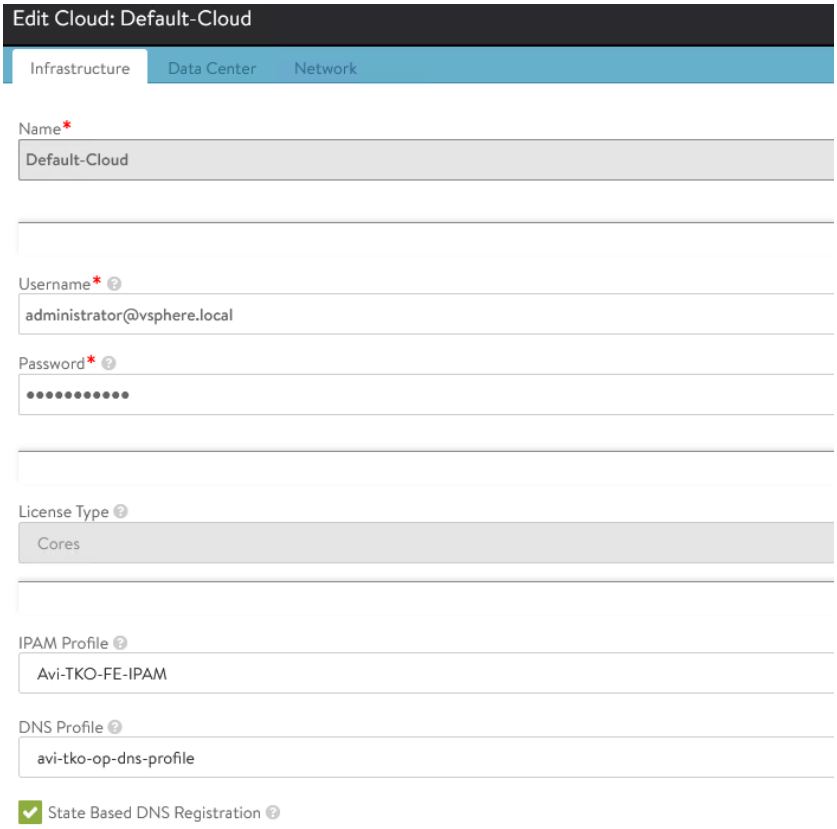

Create DNS service.

NSX Advanced Load Balancer provides generic DNS virtual services that can be implemented with various functionalities to meet different requirements. The DNS virtual service can be used to load balance DNS servers, hosting static DNS entries, virtual service IP address DNS hosting or hosting GSLB service DNS entries. For more information on NSX Advanced Load Balancer features please visit NSX Advanced Load Balancer features.

For this reference architecture a DNS virtual service was created that served as DNS sever for a subdomain of the primary Active Directory domain via DNS delegation. This generic process is outlined below. With this DNS configuration along with IPAM, a DNS entry will be created for services. Delegation of DNS domain will depend on your Active Directory architecture. The domain delegation process for this reference architecture is described later in this document under DNS Domain Delegation.

1. Create DNS virtual service.

   To create DNS virtual service, navigate to Applications > Virtual Services > Create Virtual Service. Give the service a name and select TCP/UDP and application profile. Application profile is of type “System DNS.” Click save.

Figure 46: Create DNS service.

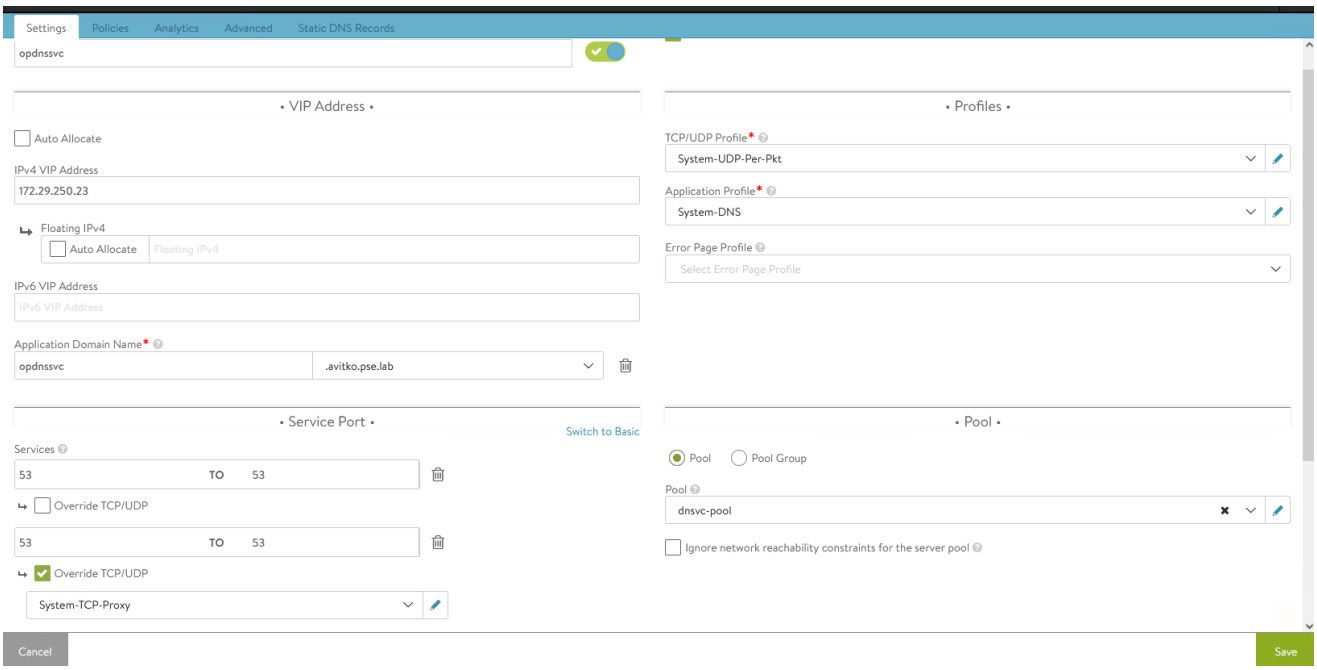

2. Add DNS virtual service to Default-Cloud instance.

   Navigate to Administrator > Settings > DNS Service and select virtual DNS service.

Figure 47: Add virtual service

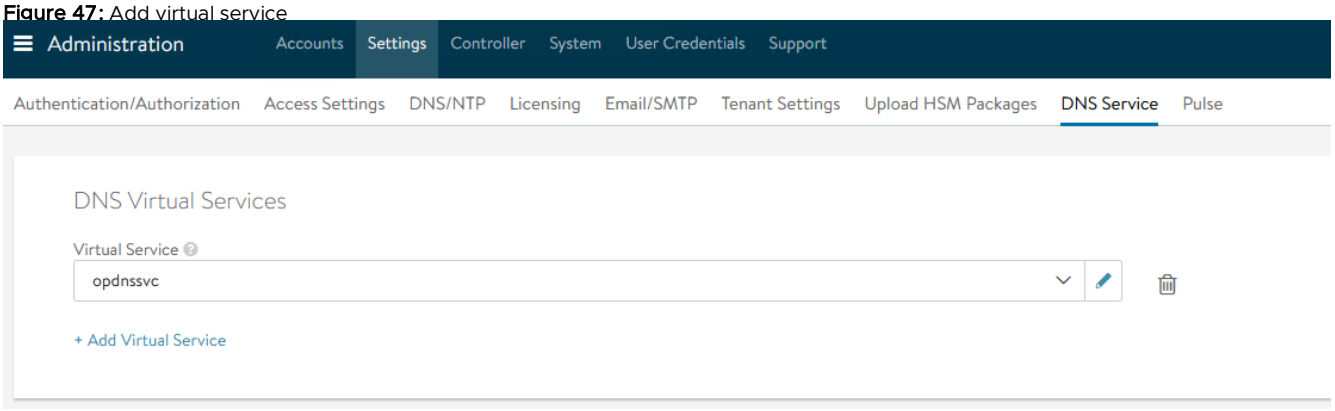

Export the SSL/TLS certificate.

The certificate created in the earlier steps will be needed during Tanzu Workload Management deployment. Following these steps, export the certificate to be used later.

1. Go to Templates > Security and select the certificate created earlier.
2. Click the down arrow on the right-hand side to export the certificate.

Figure 48: Export certificate

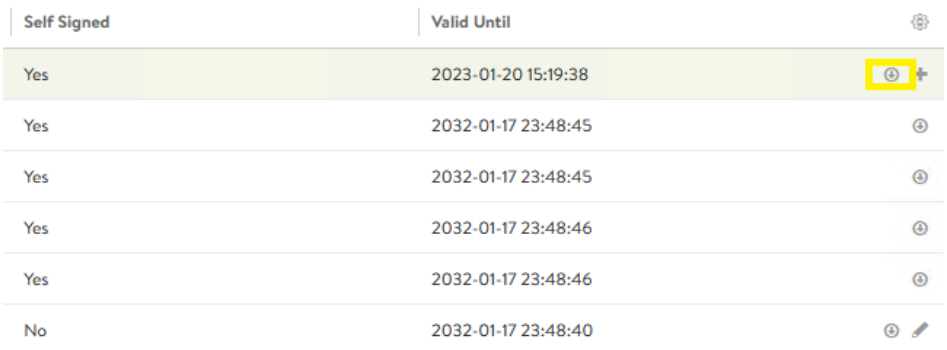

3. Copy the certificate to clipboard.

Figure 49: Copy certificate

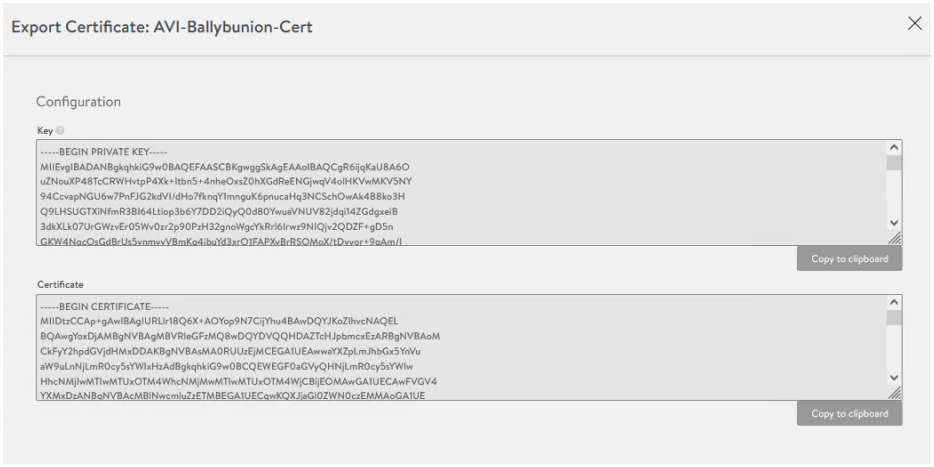

Create route between workload and front-end networks.

If the VIP and workloads are on separate networks, as in the case here, a route needs to be created between the front-end and workload networks.

1. Navigate to Infrastructure > Routing > Static Route tab.
2. Create a new static route.

Figure 50: Create route.

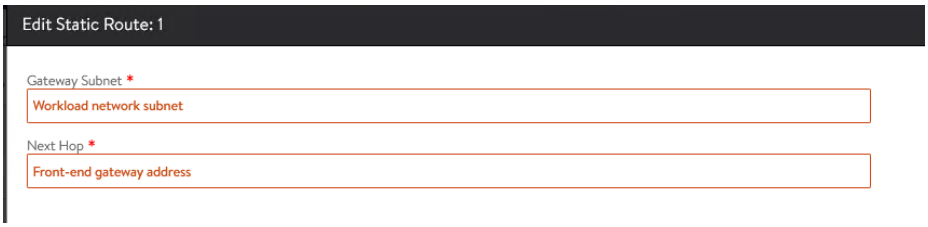

Enable vSphere with Tanzu Workload Management

Once NSX Advanced Load Balancer has been successfully deployed, vSphere with Tanzu Workload Management can be enabled. As a best practice, a Tanzu specific storage policy needs to be defined and storage tagged prior to enabling Workload Management. This storage policy should be different than the default vSAN storage policy that is initially configured by the VxRail HCI System software during initial VxRail cluster creation. For vSphere with Tanzu use cases, storage policies and tags are used to assign storage to Kubernetes cluster nodes and persistent volumes. Using a separate and dedicated storage policy ensures that Tanzu workloads are placed on the desired storage pool, separate from other vSphere workloads. The process for doing this is described below.

Note: Dell VxRail provides options for cluster vCenter Server deployment configurations. VxRail can deploy a vCenter Server that is hosted on the deployed cluster and managed by VxRail HCI System Software as part of its cluster lifecycle management capabilities. This Is referred to as a VxRail managed vCenter on cluster deployment configuration. The other option that is available is to use an existing, customer managed vCenter Server that can be used to manage the Dell VxRail cluster deployment. In this configuration, vCenter server lifecycle management would be the responsibility of the customer to manage. Please consult VxRail documentation for more info. For this reference architecture document, a VxRail managed on cluster vCenter Server deployment configuration was used.

Create storage tag and policy.

1. Select your cluster in vCenter Server and go to Datastores. Select the Datastore to be used for Workload Management.
2. Under “Tags” click “Assign.”

Figure 51: Create tag

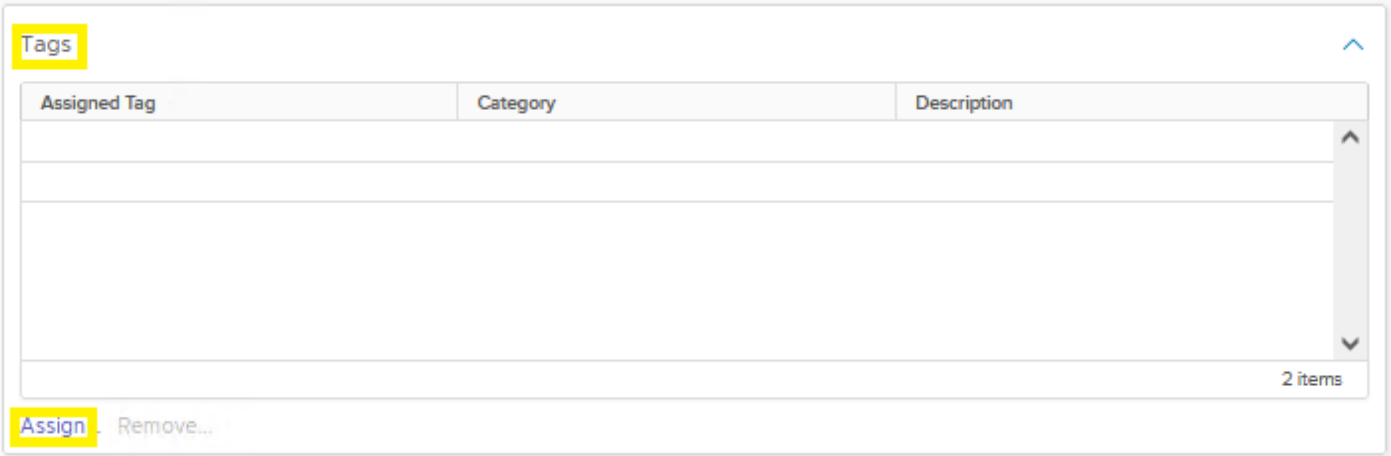

3. On the next screen click “ADD TAG” and Create Tag dialog opens. Enter a name for the tag and select a category. If desired a new category can be created. Click CREATE to create the tag.

Figure 52: Create tag.

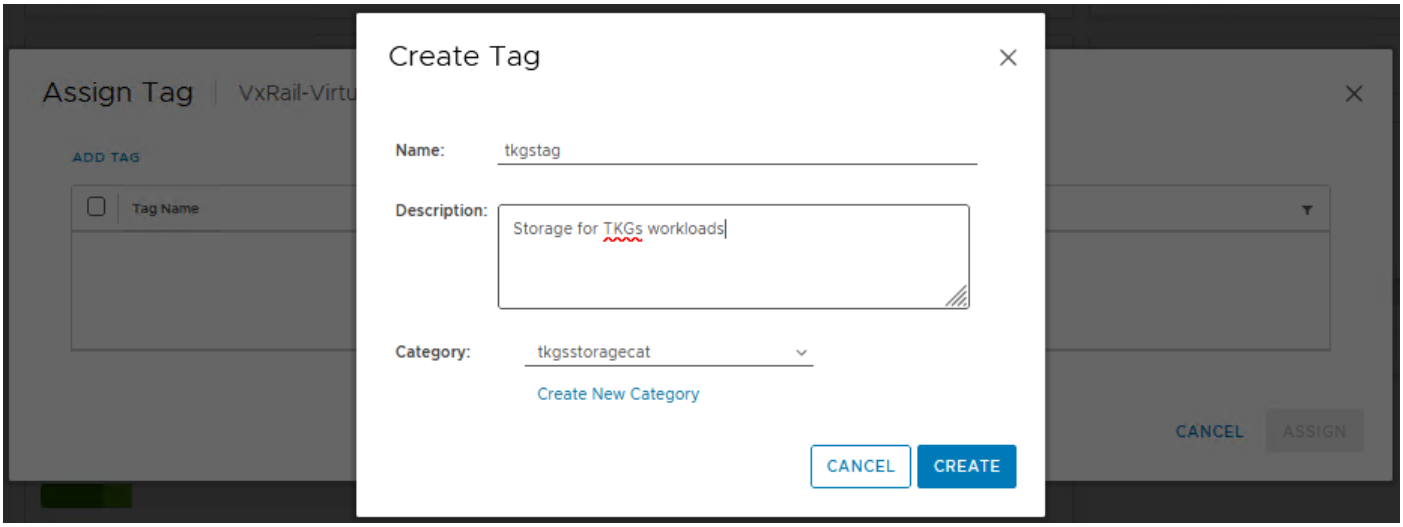

4. In vCenter Server navigate to Menu > Policies and Profiles > VM Storage Policies and select CREATE. On the next screen give the policy a name and click NEXT.
5. For Policy Structure, check the "Enable tag-based placement rules" and click NEXT.

Figure 53: Placement Rules

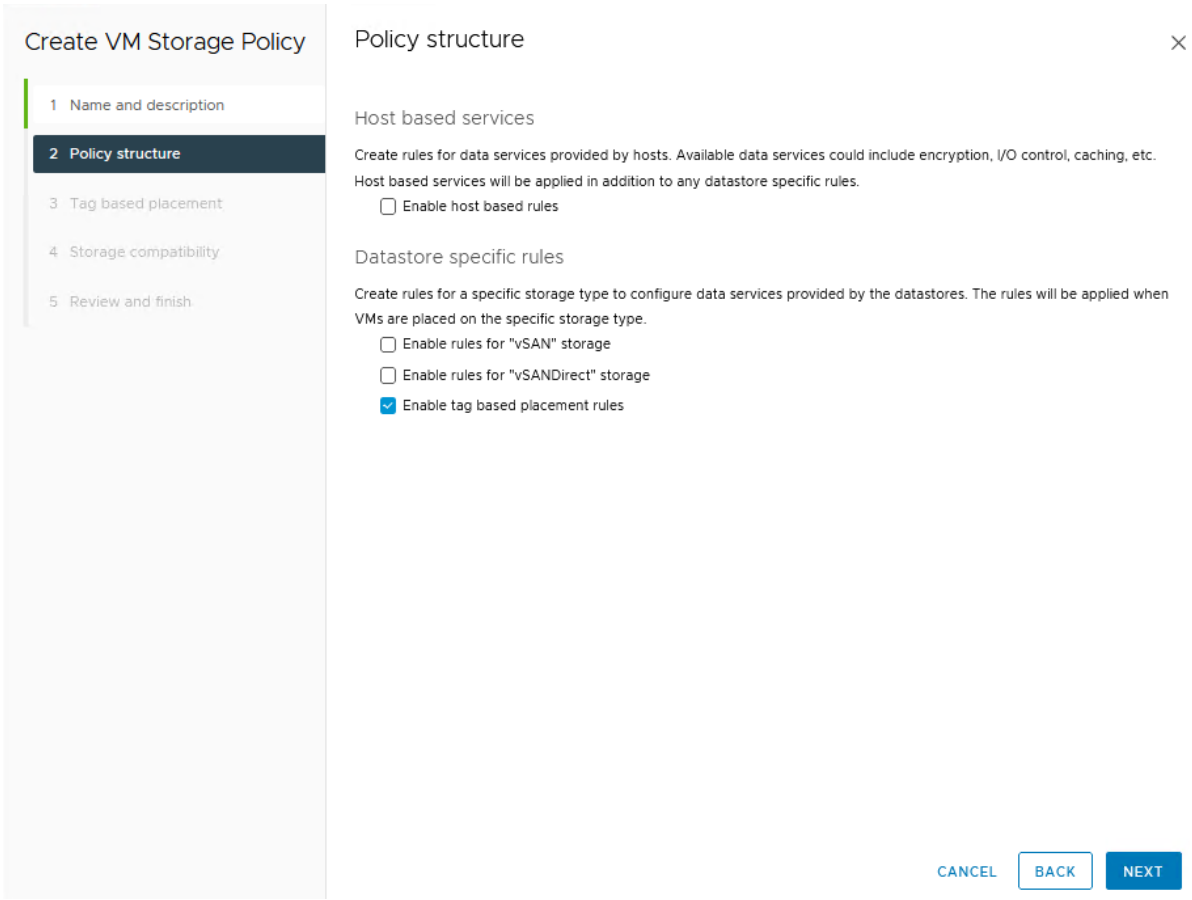

6. Create a rule by selecting the tag category, and "Use storage tagged with" as usage option. Browse and select the tag created earlier. Click NEXT.

Figure 54: Create rule

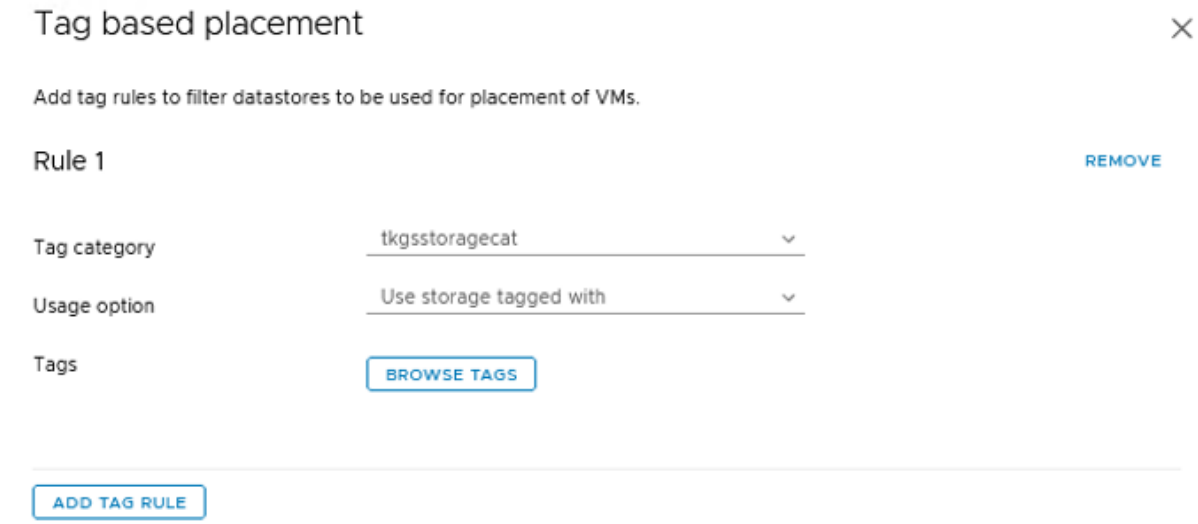

7. Select storage and click next and finish the policy creation.

Figure 55: Select storage

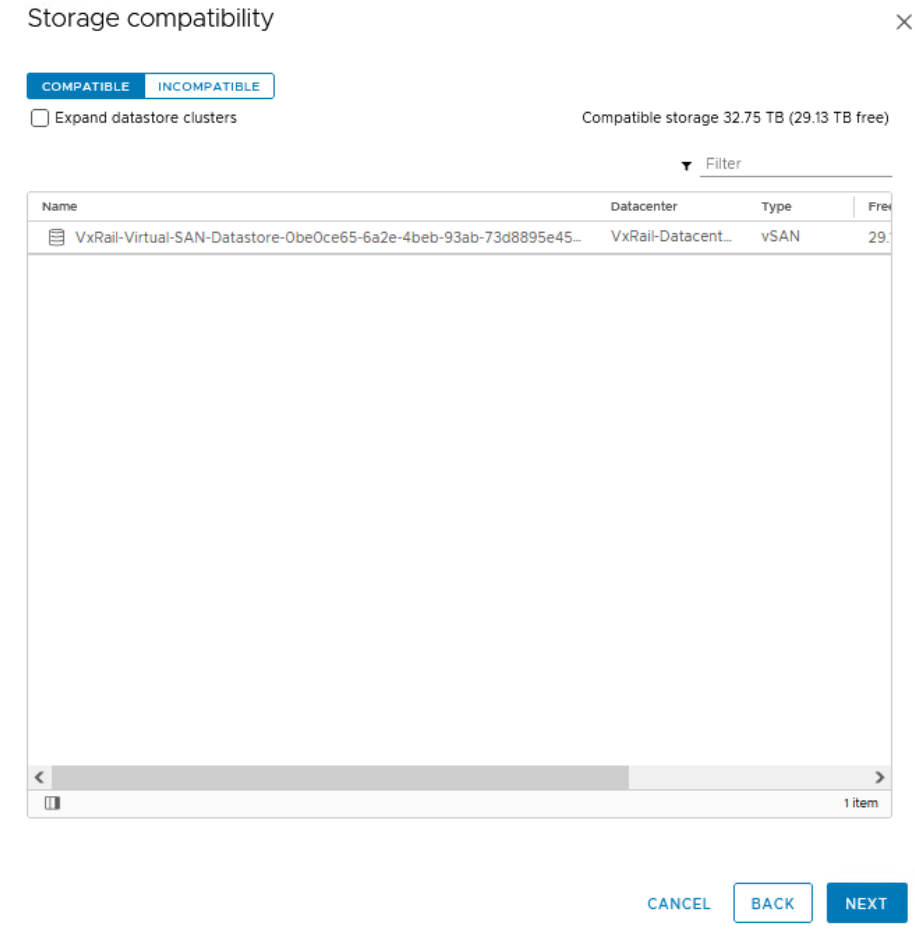

Create content library.

A content library is required by Tanzu Kubernetes Grid Service that will hold the images required by vSphere to deploy Supervisor and workload clusters. The subscription URL used for this content library is https://wp-content.vmware.com/v2/latest/lib.json. Create a content library with the given subscription URL. Please note that it will take some time before content is downloaded and available in the library.

Enable Workload Management

1. Navigate to Menu > Workload Management. Review the prerequisites for setting up Supervisor cluster and ensure that they are met before proceeding. Click "Get Started."

Figure 56: Enable Workload Management

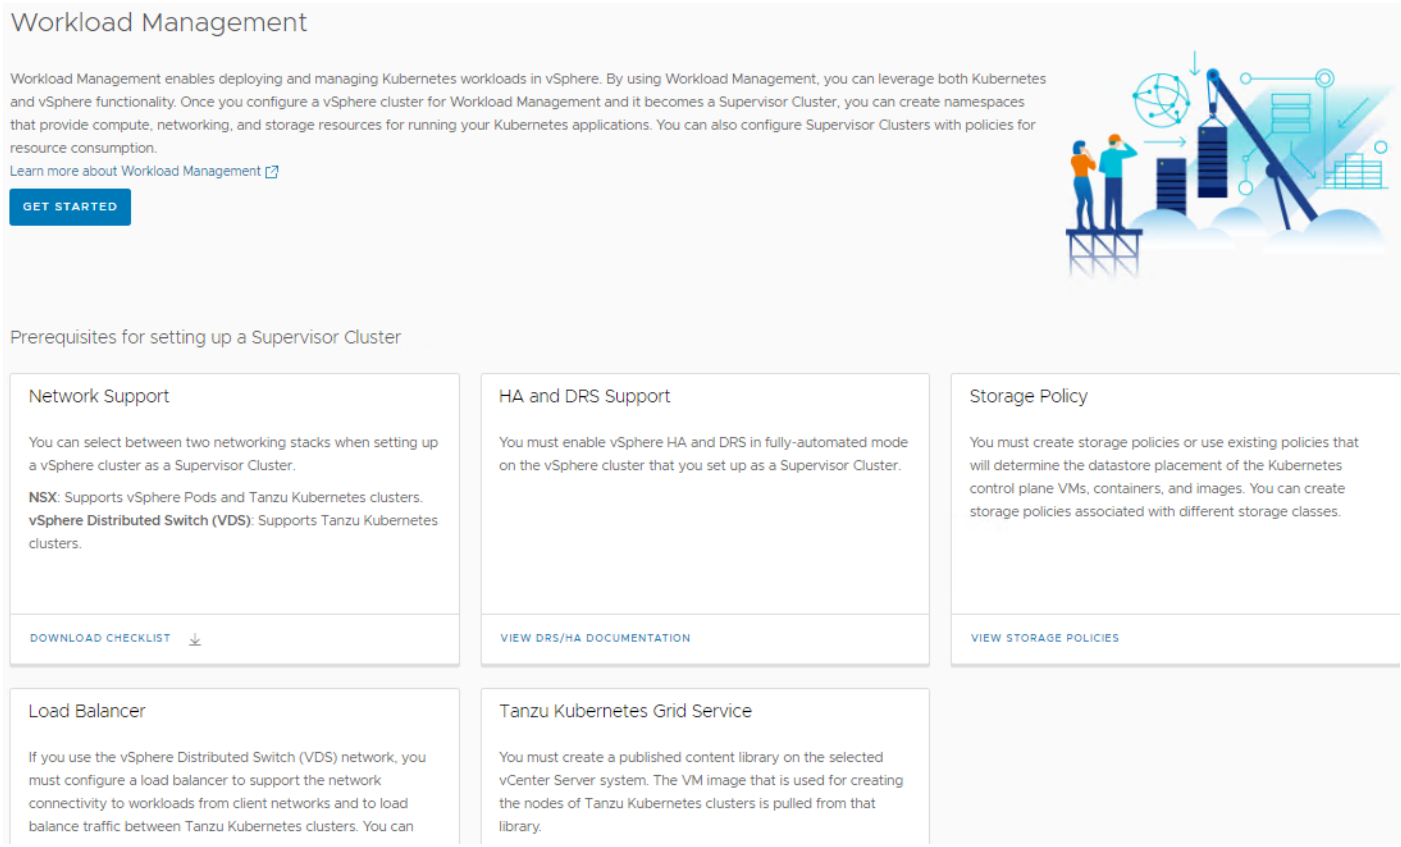

2. Select vSphere Distributed Switch and click next.

Figure 57: vCenter Server and Network

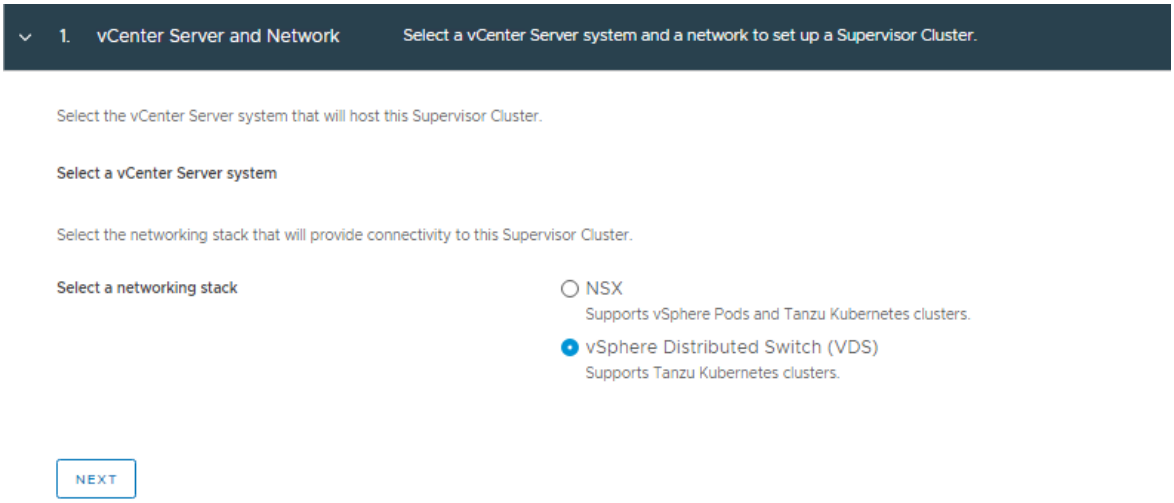

3. Select the cluster.

Figure 58: Select cluster.

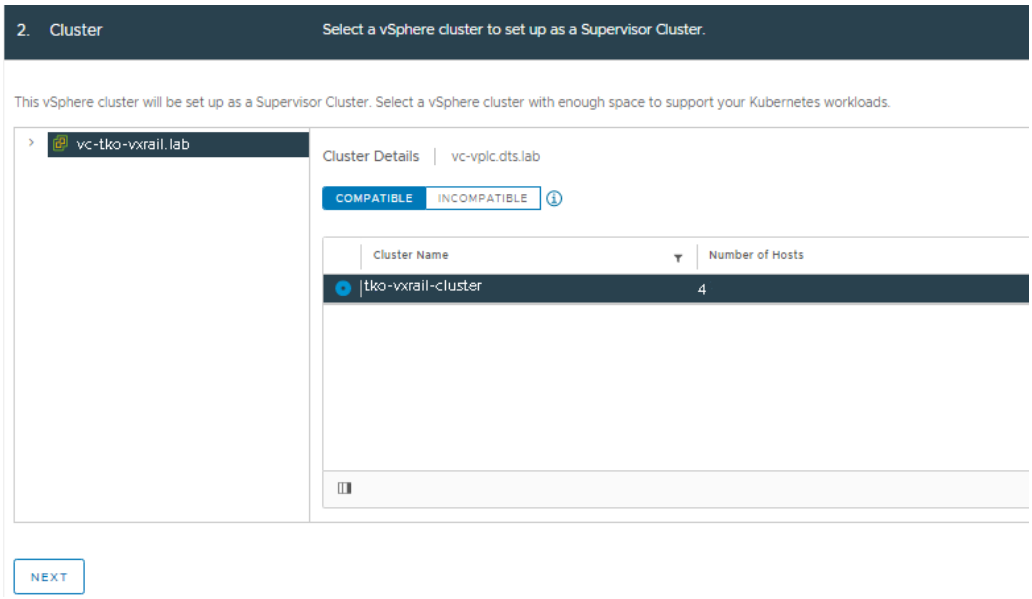

4. Select the storage policy created previously.

Figure 59: Storage Policy

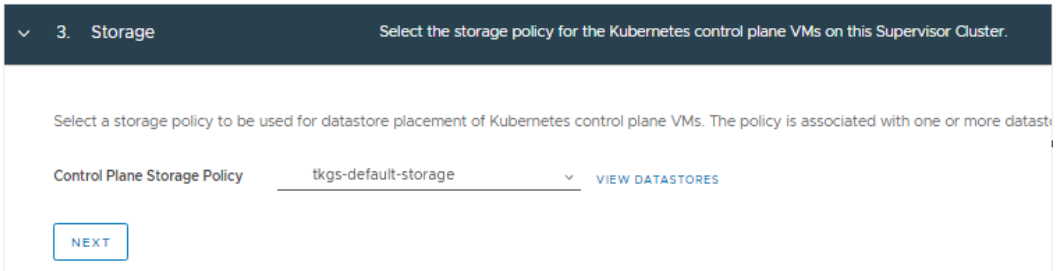

5. On the next screen fill out the NSX Advanced Load Balancer details and copy and paste the certificate exported earlier. Ensure to use "< IP address>:443 " format when entering the IP address for the controller. Click NEXT.

Figure 60: NSX Advanced Load Balancer details

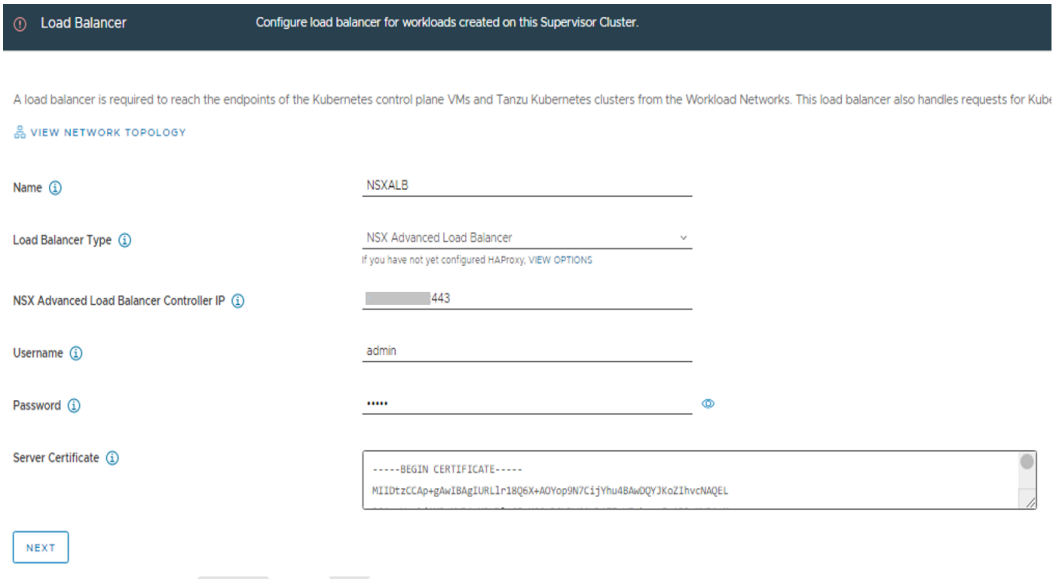

6. Fill in the management network details. Either DHCP or Static assignment can be used. When using static IP address assignment, ensure to reserve a block of five IP addresses for control plane VMs in the Supervisor cluster. When using DHCP, ensure that the DHCP server in your environment supports client identifiers to provide IP addresses for Supervisor Cluster control plane VMs and floating IP. The DHCP server must also be configured with compatible DNS server(s), NTP server(s), and DNS search domain(s). Click NEXT.

Figure 61: Configure management network.

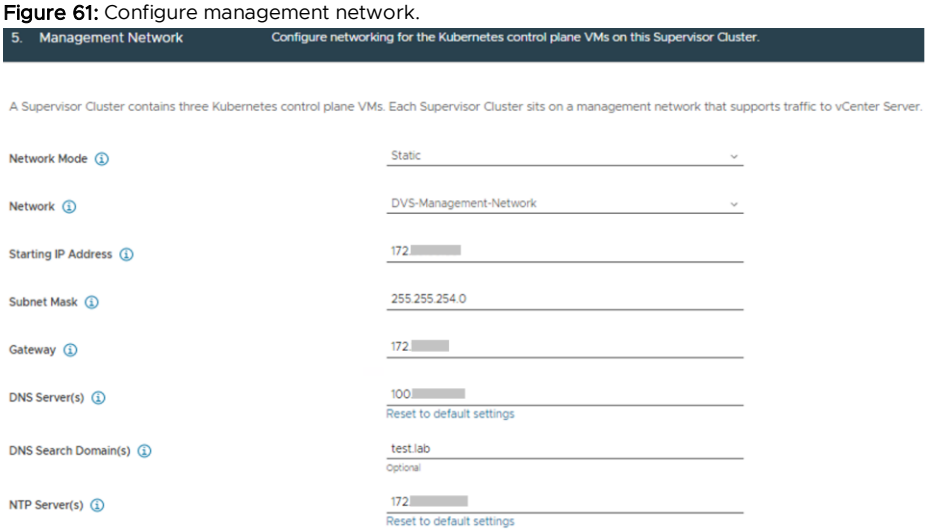

7. vSphere namespaces on this Supervisor Cluster require Workload Networks to provide connectivity to the nodes of Tanzu Kubernetes clusters and the workloads that run inside them. Internal IP addresses are used to allocate Kubernetes services of type ClusterIP. These IP addresses are internal to the cluster but should not conflict with any other IP range. Configure the workload network information page for your specific network. Click NEXT.

Figure 62: Configure workload network

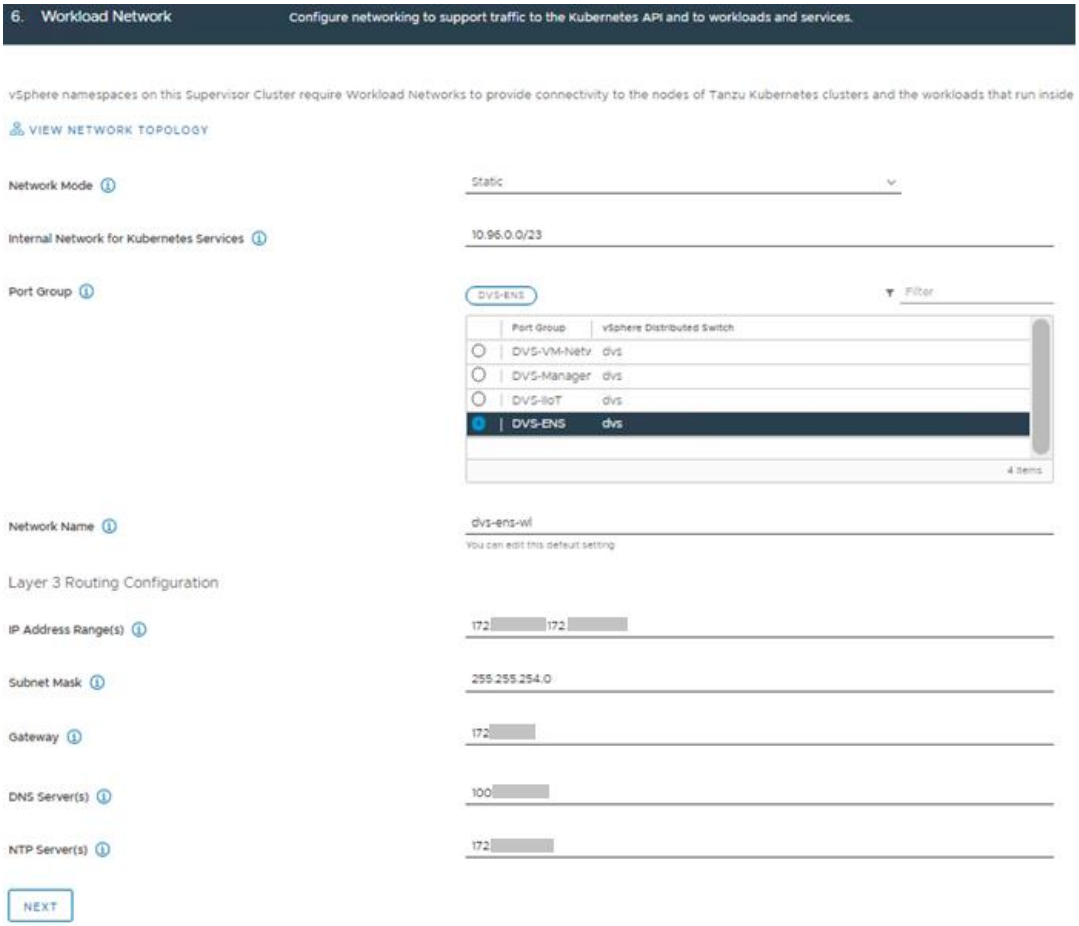

8. Add content library.

Figure 63: Add content library

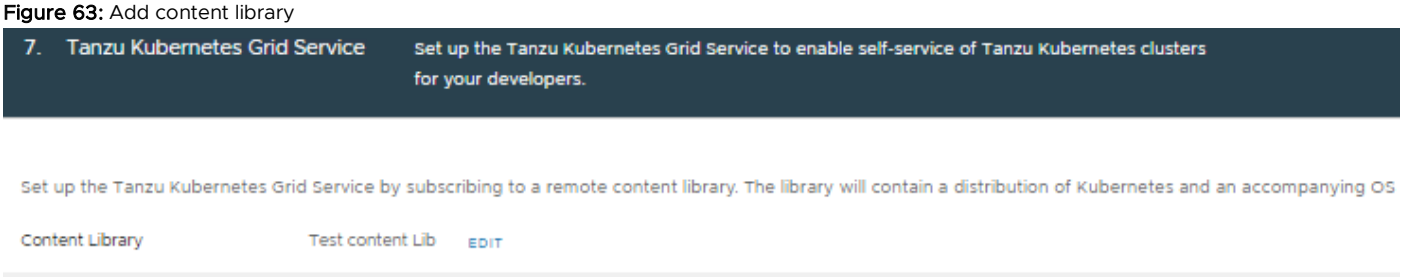

9. Select the size of the control plane VM per your requirements and optionally enter DNS name designated for Kubernetes API server. For production deployments with Tanzu Mission Control integration, a large form factor is recommended for Supervisory control plane nodes. Click Finish to start the configuration process.

Figure 64: Control plane size and API server

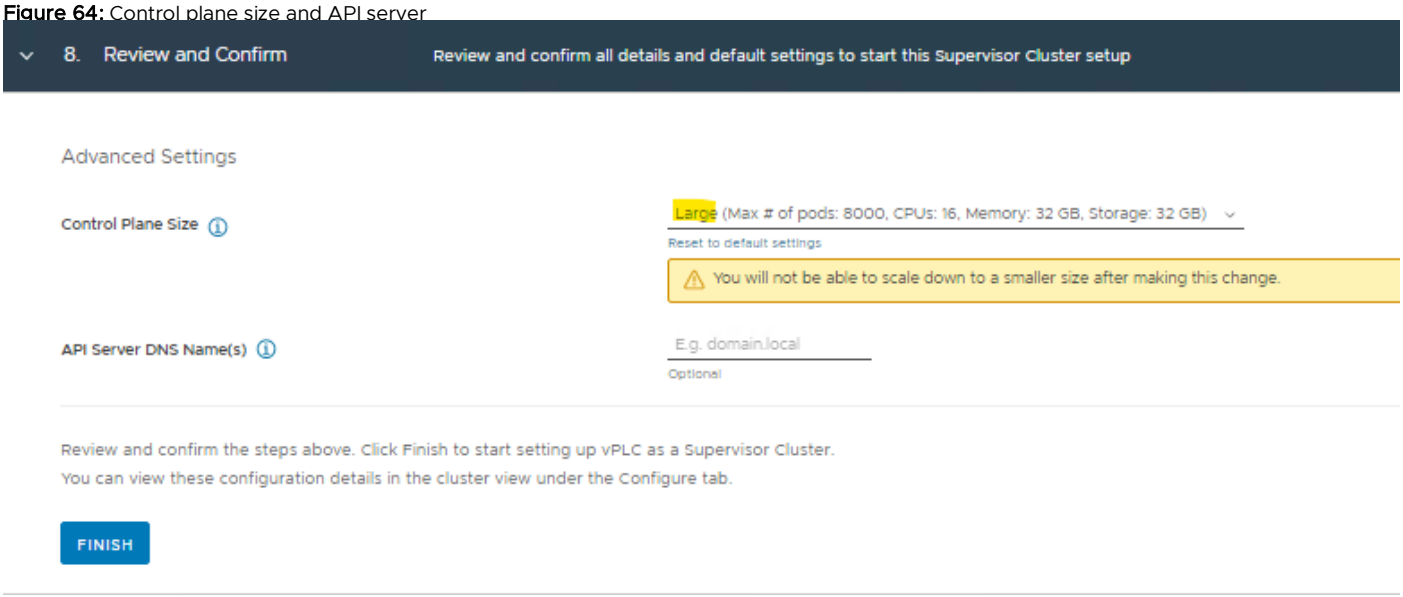

Authentication and access

Authentication to Tanzu Kubernetes clusters can be accomplished in different ways depending on your architecture, user authentication and access requirements. In an on-premises environment the simple and reliable method is to use vCenter SSO to authenticate users or to add a local identity source such as Active Directory over LDAPS. For this reference architecture Active Directory authentication with LDAPS was used to authenticate users. The domain controller was configured with Certificate Authority and the controller certificate was exported to be used in the identity source configuration process. In a multi-cloud environment, an external identity source such as Azure Active Directory or Okta can be incorporated for user authentication. VMware cloud services also provide a way to authenticate via federated domains. For more information visit VMware Cloud Services Enterprise Federation page.

What follows are some steps that the vSphere administrator will perform to give access to the domain users in the namespace created for initial workload cluster deployment.

vSphere Administrator Tasks.

1. Add Active Directory as identity source to vCenter Server.
2. Create namespace for DevOps admins and developers to deploy clusters to.
3. Assign permissions to DevOps engineers in the namespace.
4. Assign storage policies, virtual machine classes and quotas to namespace.
5. Provide namespace access information to DevOps and/or Developers.

Add Active Directory as identity source to vCenter.

1. vSphere menu > Administration > Single Sign On > Configuration > Identity Provide > Identity Sources and click ADD.
2. Fill in the required information for the domain, upload the domain controller certificate and connect to the domain controller using port LDAPS port (636)

Note: Use of domain names with ". Local" is not supported. Please see KB article. For information on configuring LDAPS in your environment please see Microsoft Documentation.

Figure 65: Configure identity source.

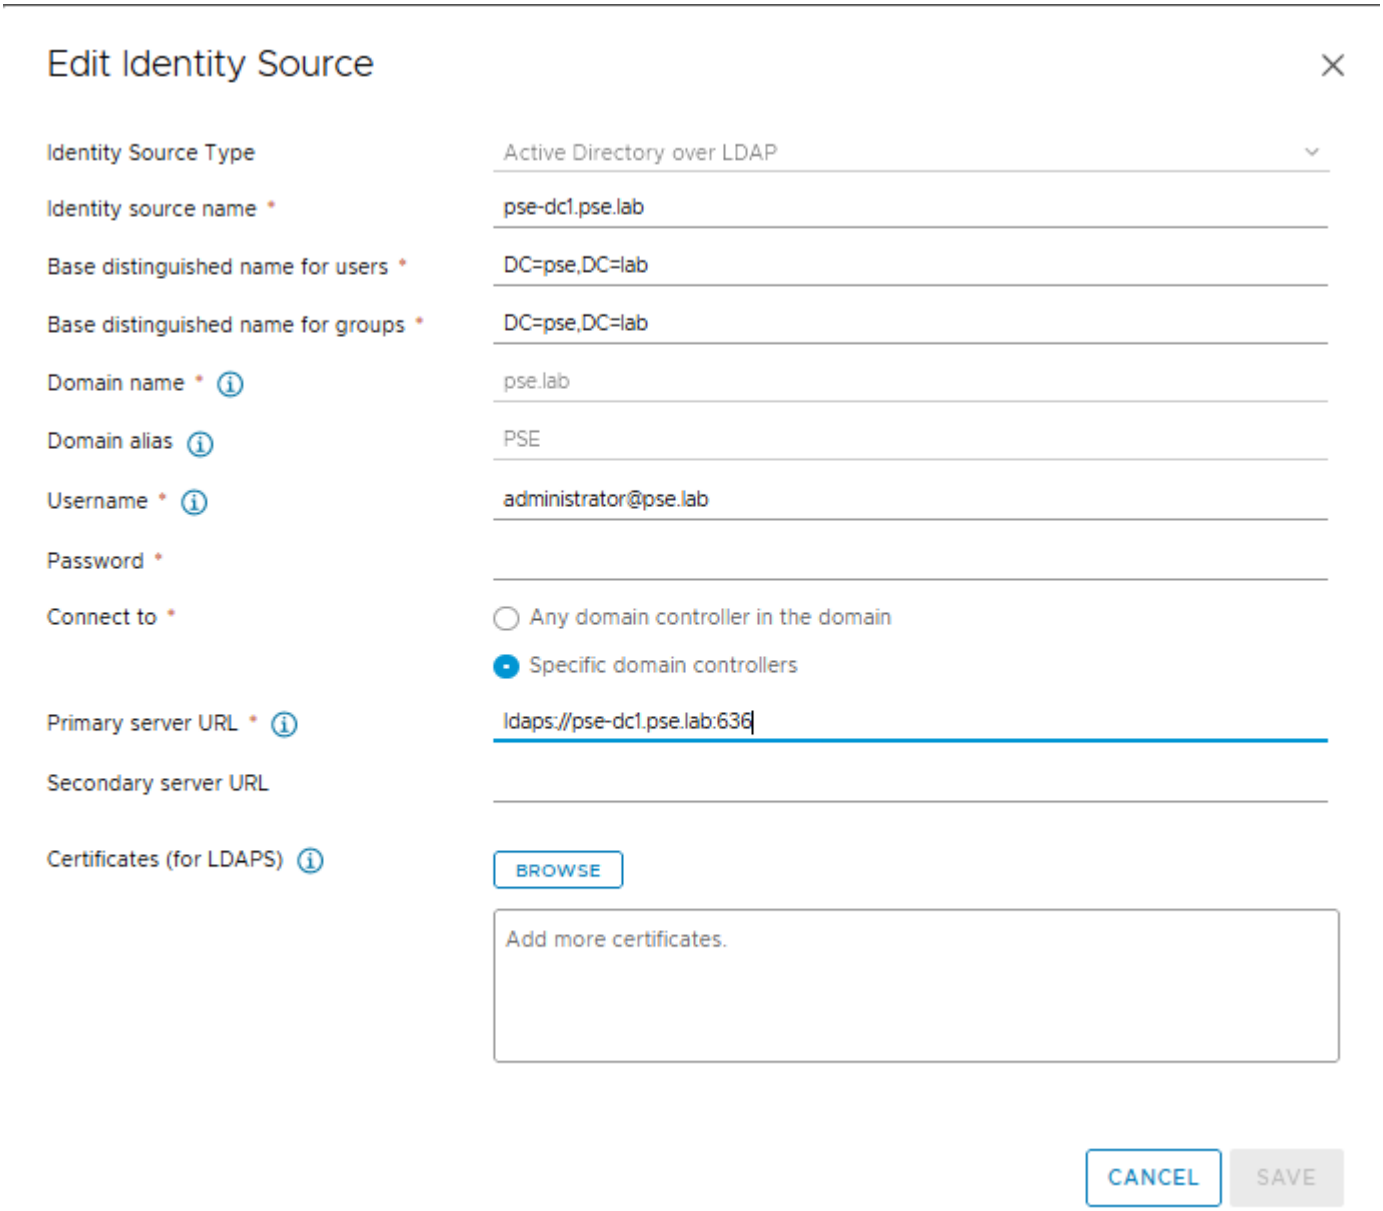

3. To verify the Active Directory integration was successful, navigate to Users and Groups. The domain now should be visible in the drop-down list and query to find a user should be successful.

Figure 66: Validate AD integration.

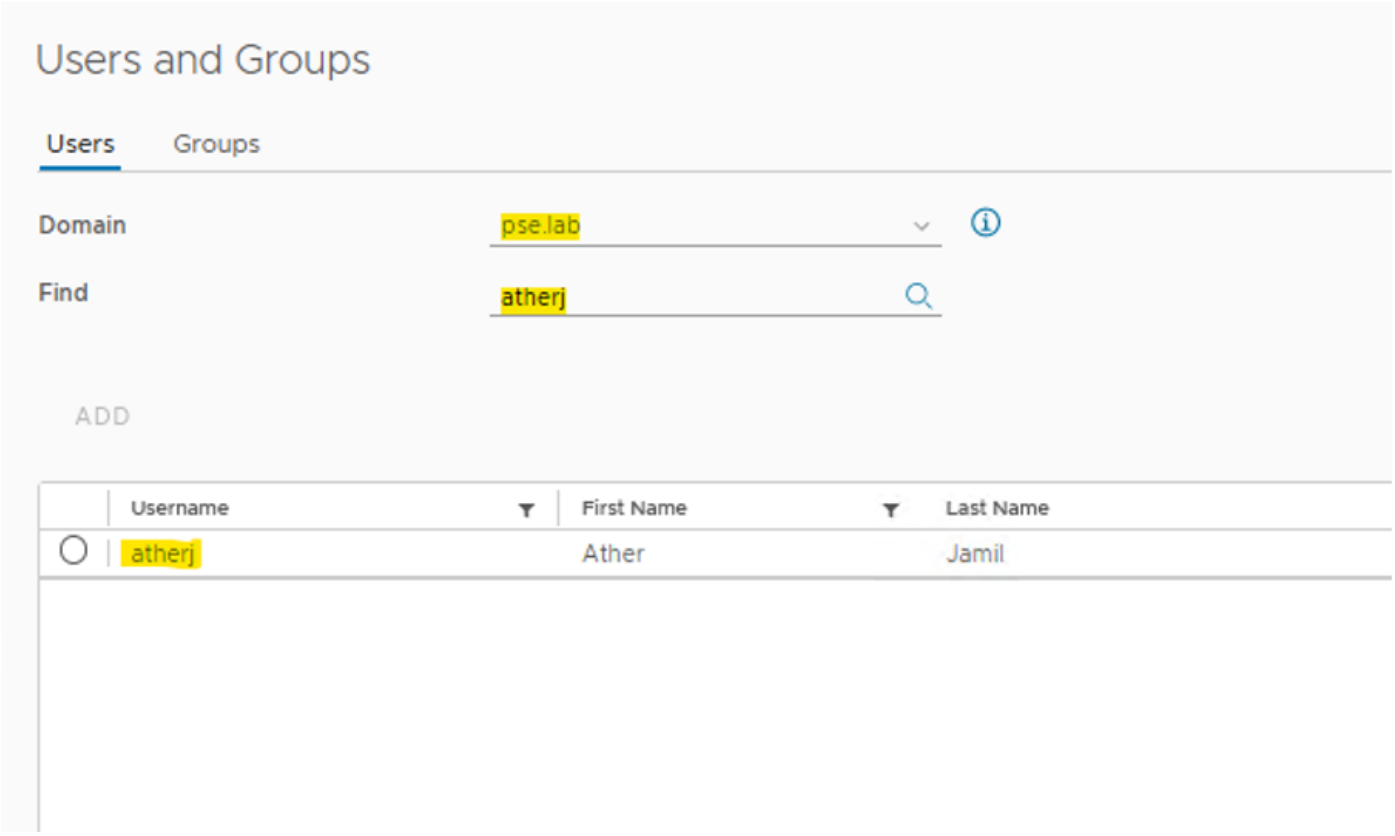

Create namespace.

Tanzu Kubernetes Grid Service uses namespaces to provide tenant separation and isolation. Namespaces are defined on the Supervisor cluster and can be configured with user permissions, resource quotas and storage policies. Depending on requirements you assign VM classes and content libraries to the namespaces to download latest Tanzu Kubernetes releases and VM images. The number of namespaces created depends on organizational requirements. For this reference architecture a single namespace was created.

Figure 67: Namespace configuration

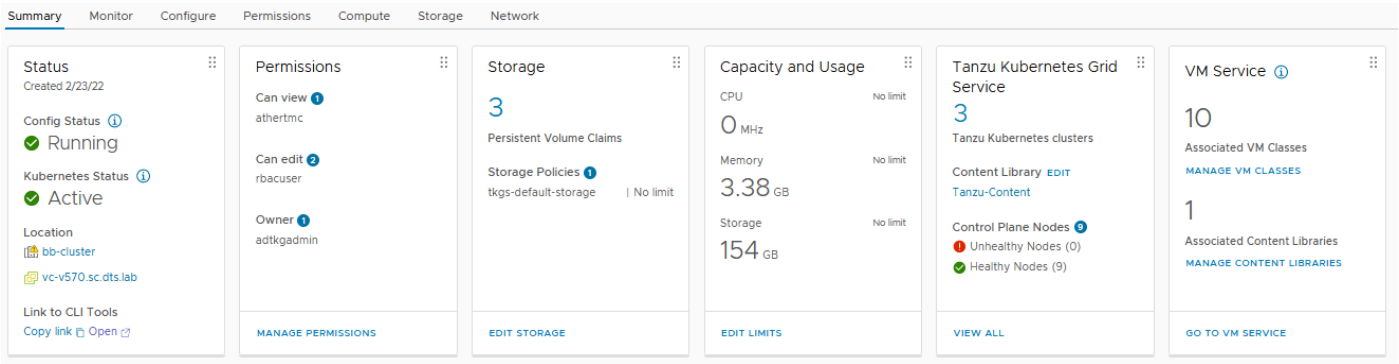

Provide namespace and login information to users.

Once the namespace is configured, the administrator needs to provide the DevOps team with relevant information such as username and password, vCenter Server certificate, as well as the namespace URL, so they can begin to create clusters and deploy workloads on them. The user will install the certificate on the access machine where he or she intend to run Kubernetes commands. The namespace URL can be obtained from the namespace configuration status page as show in figure 43.

Figure 68: Access URL

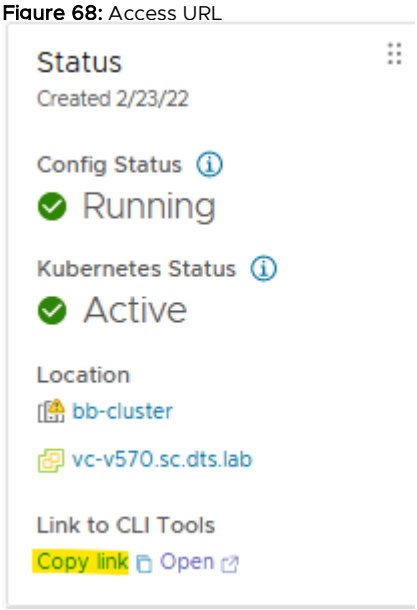

The URL provides instructions to the user to download vSphere CLI tools to access the namespace.

Figure 69: Kubernetes CLI tools

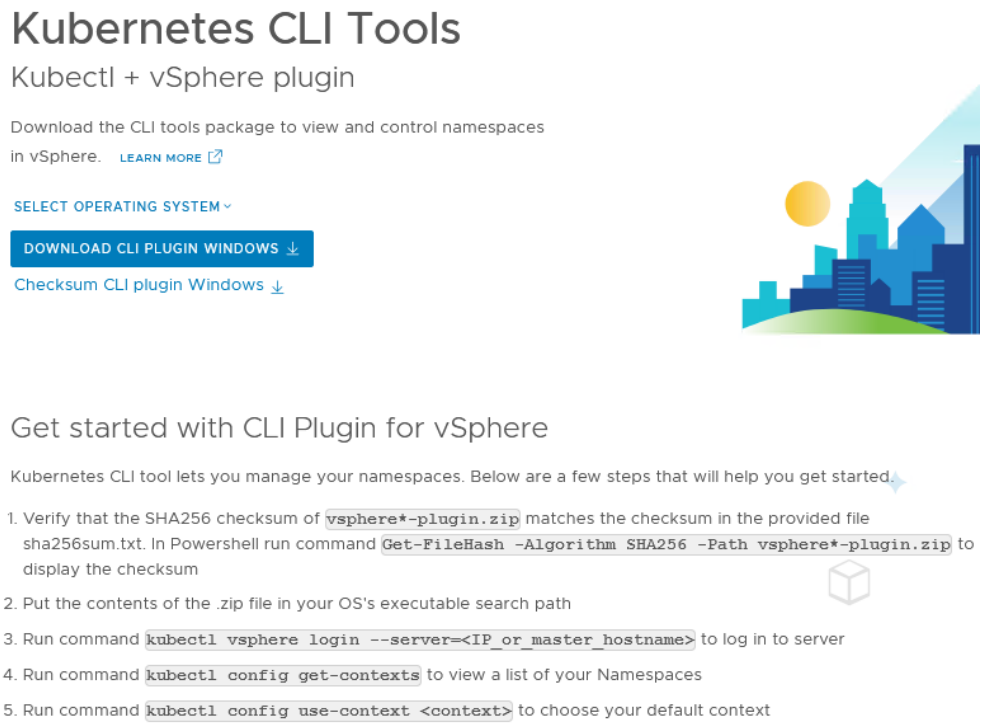

VMware Cloud (VMC) on AWS

VMware Cloud on AWS is an integrated cloud offering jointly developed by Amazon Web Services (AWS) and VMware. You can deliver a highly scalable and secure service by migrating and extending your on-premises VMware vSphere-based environments to the AWS Cloud running on Amazon Elastic Compute Cloud (Amazon EC2).

Figure 70: VMware cloud on AWS

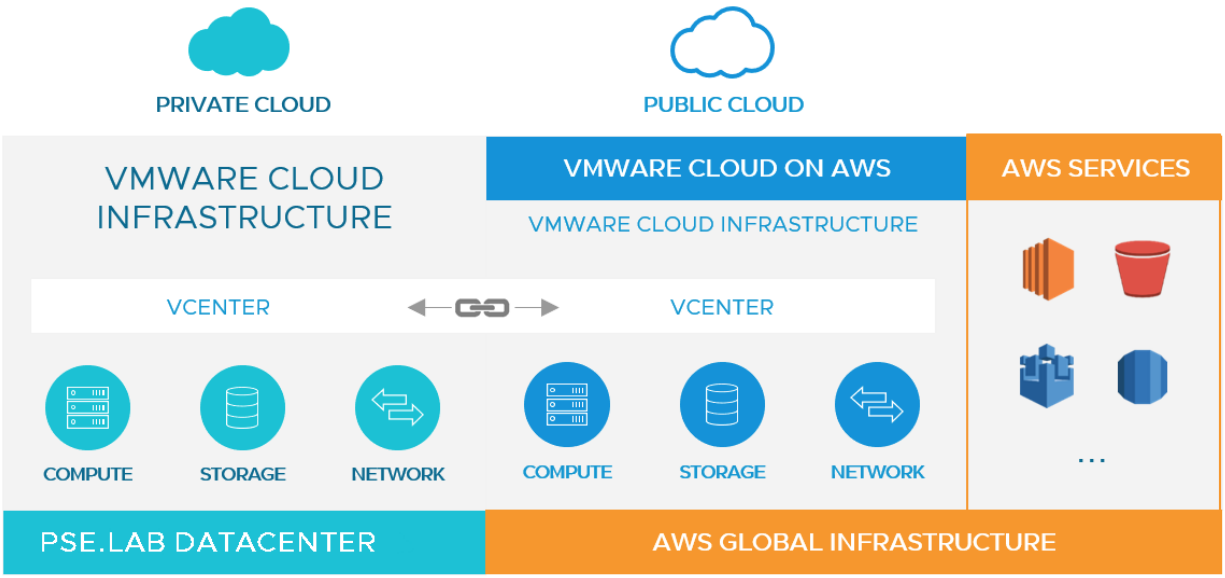

Note: This document does not discuss provisioning VMware Cloud on AWS. For on-boarding SDDC clusters on AWS please refer to the VMware Cloud on AWS On-Boarding Guide

Note: At the time of writing of this paper, VMC only supports deployment of vSphere with Tanzu with NSX-T. Deployment with NSX Advanced load balancer only is not supported. It can however still be used as a load balancer along with NSX-T.

Installing Tanzu for Kubernetes Operations on VMC

Installation of Tanzu Kubernetes Grid clusters is now made easy by Service Installer for Tanzu. For this reference architecture, Service installer for Tanzu (SIVT) version 1.3 was used to deploy Tanzu Kubernetes Grid (multi-cloud) clusters on VMC. Service installer for Tanzu documentation has step by step deployments guides that can be used to deploy Tanzu Kubernetes Grid on vSphere as well as on VMC. For detailed step by step process, please refer to Service Installer for Tanzu deployment guide for VMC.

For this reference architecture Service Installer for Tanzu (SVIT) version 1.3 was installed on the VMC cluster. SVIT only requires a single segment to be created to be used for Tanzu Kubernetes Grid management. SVIT creates the remaining segments based on SVIT reference design for deploying Tanzu Kubernetes Grid on VMC.

For example, in the below snapshot, segment “tko-tkgm-mgmt-seg” was manually created for Tanzu for Kubernetes Operations (TKO) management traffic. The remaining segment are created by SVIT based on the input provided during the deployment process.

Figure 71: VMC network segment creation

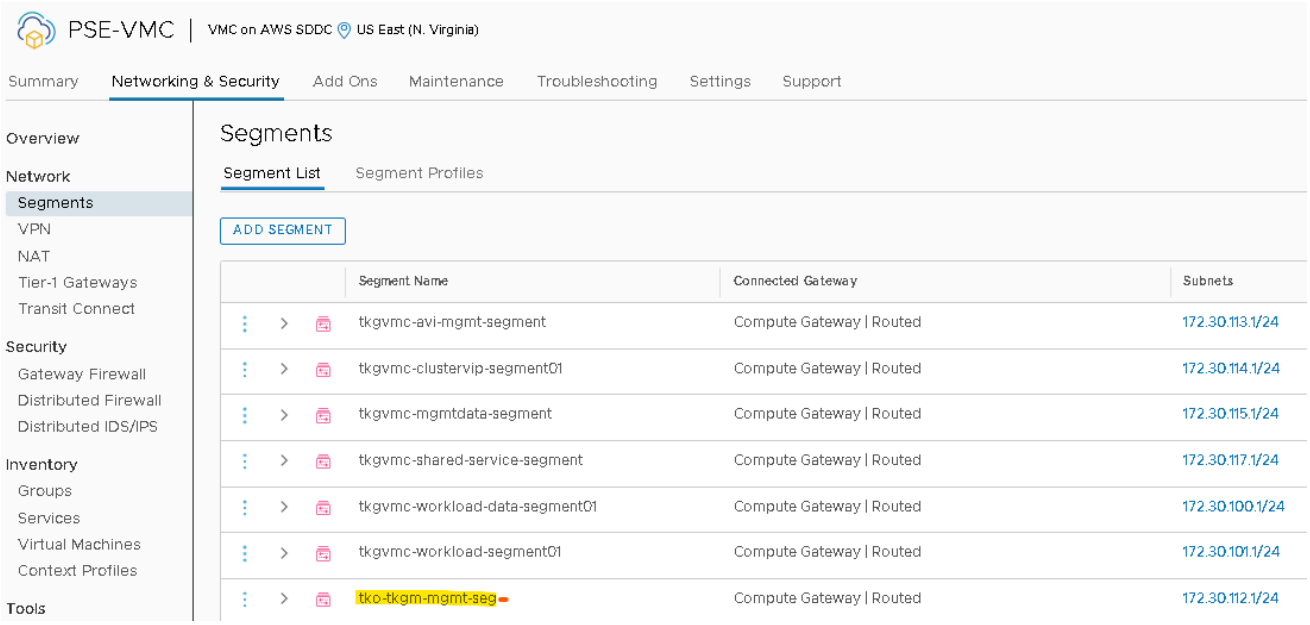

SVIT will also create groups, Gateway Firewall rules and assign them to the groups.

Figure 72: Groups

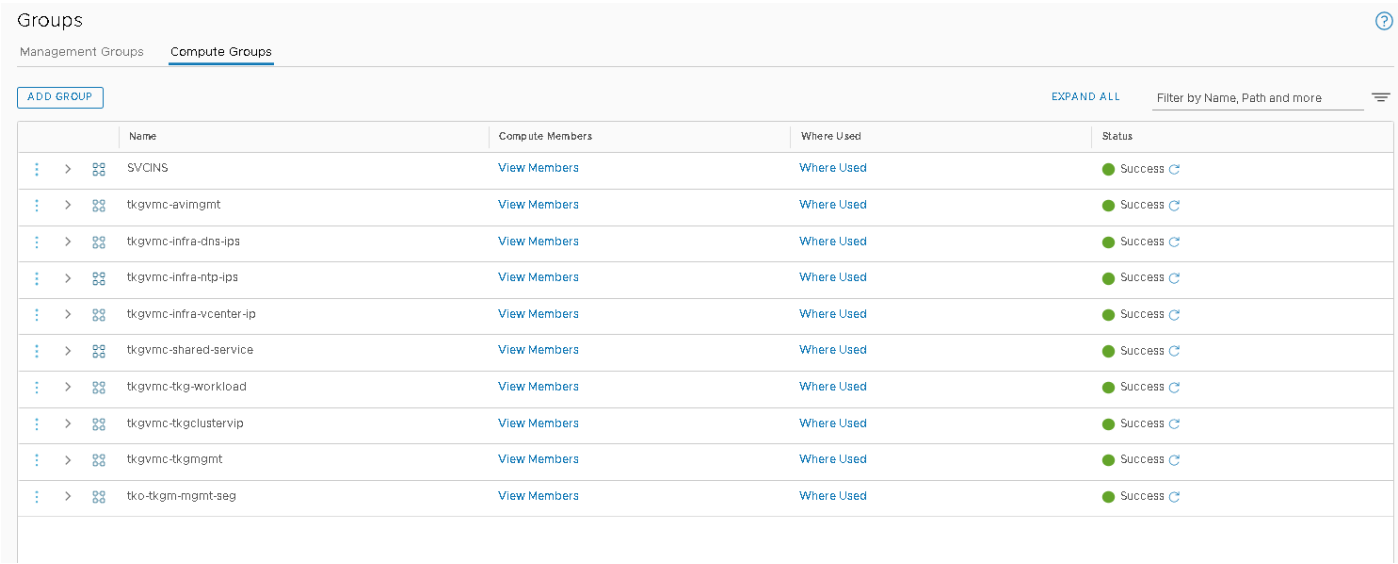

Figure 73: Firewall rules

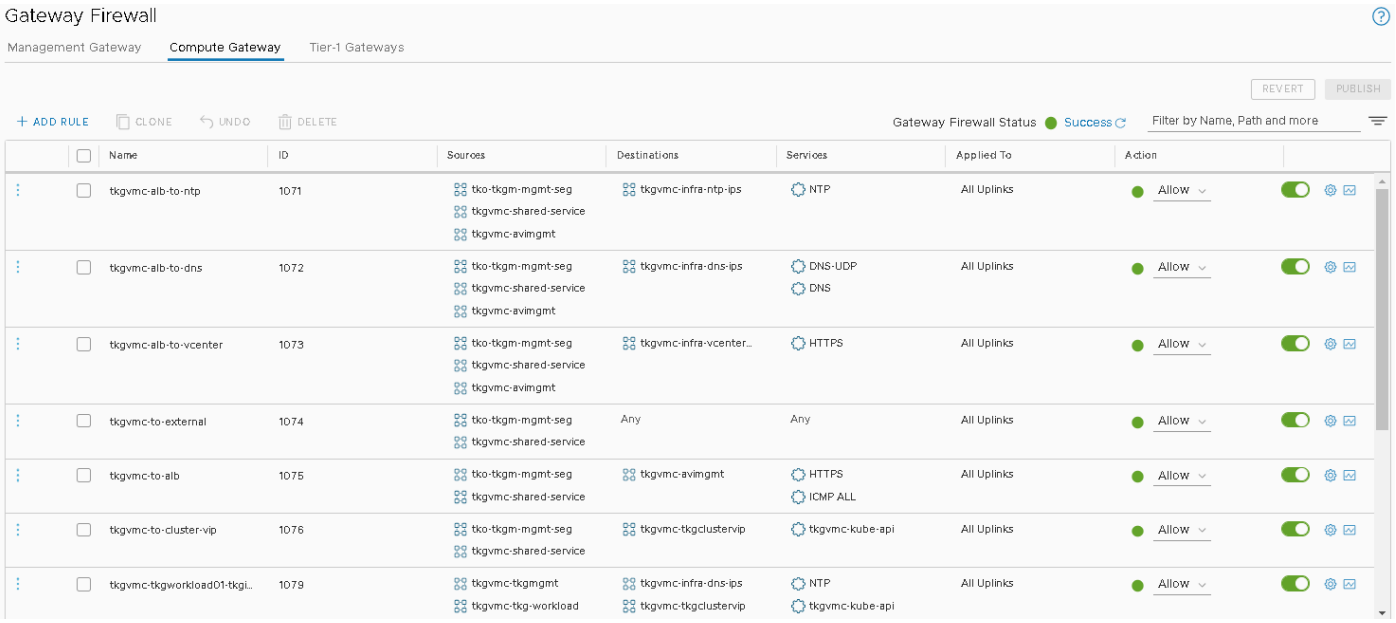

High-level steps to deploy Tanzu for Kubernetes Operations with NSX Advanced Load Balancer on VMC using SVIT as follows.

1. Download SVIT OVA (Open Virtual Appliance) from VMware Marketplace. A marketplace account is required.
2. Create a SVIT virtual machine using the OVA. Ensure that the virtual machine resides on the management network that you manually created.
3. Obtain your SDDC token as well as a marketplace token. These tokens will be needed for SVIT to configure SDDC and to download NSX Advanced Load Balancer OVA and other Kubernetes images from VMware Marketplace. Below is a snapshot of where these tokens are entered in SVIT user interface.

Figure 74: Tokens

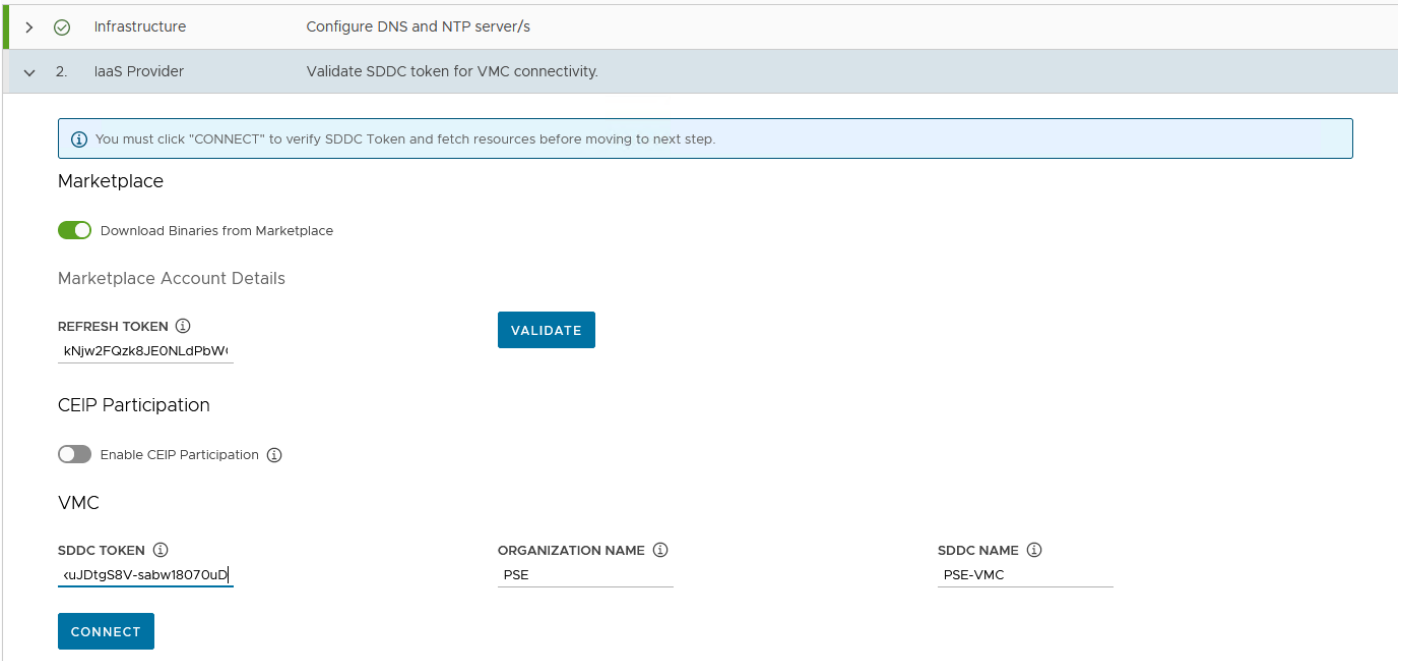

Note: In instances where SVIT cannot download NSX Advanced Load Balancer controller and Kubernetes images from the VMware Marketplace, a local content library needs to be created and images and controller. OVA uploaded. SVIT will install the controller from the local content library as shown in figure below.

Figure 75: local content library option

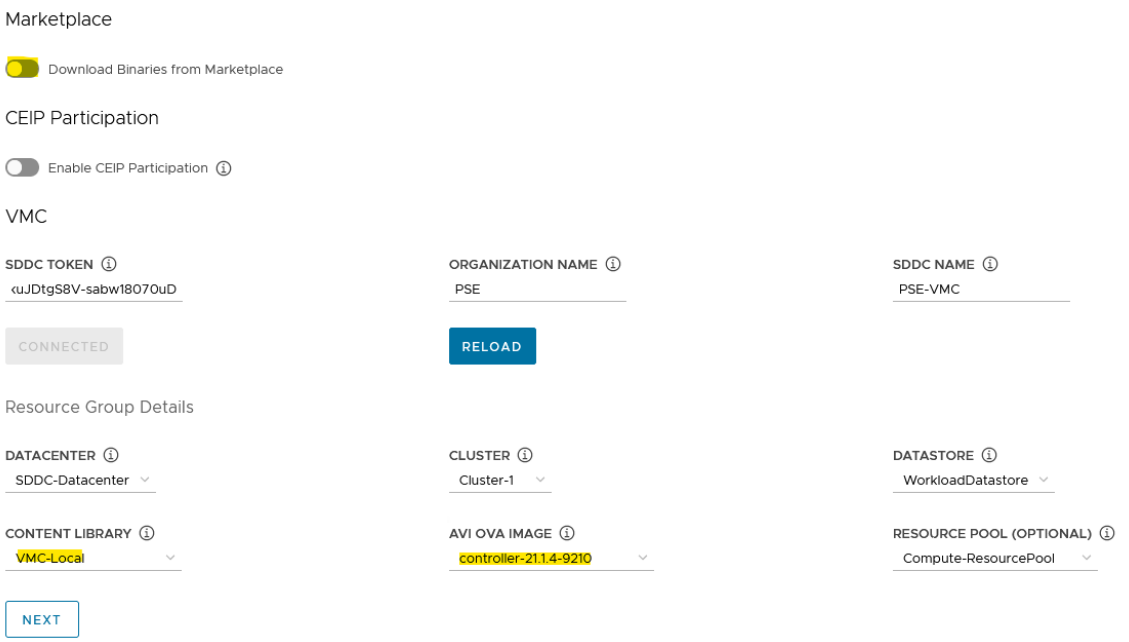

4. Enter your environment specific parameters on the remaining SVIT screen and review the configuration. On the review page you will have an option to view and safe the configuration .yaml file to your local disk or to save it to SVIT virtual machine. The default location for SVIT deployment yaml files is "/opt/vmware/arcas/src/"
5. Finally, run the command with desired parameters to deploy Tanzu for Kubernetes Operations. The command and parameters are in the SVIT deployment guide mentioned previously.

If deployment is successfully completed, appropriate resource groups and clusters will be created and viewable in vCenter.

Figure 76: VMC Tanzu Clusters

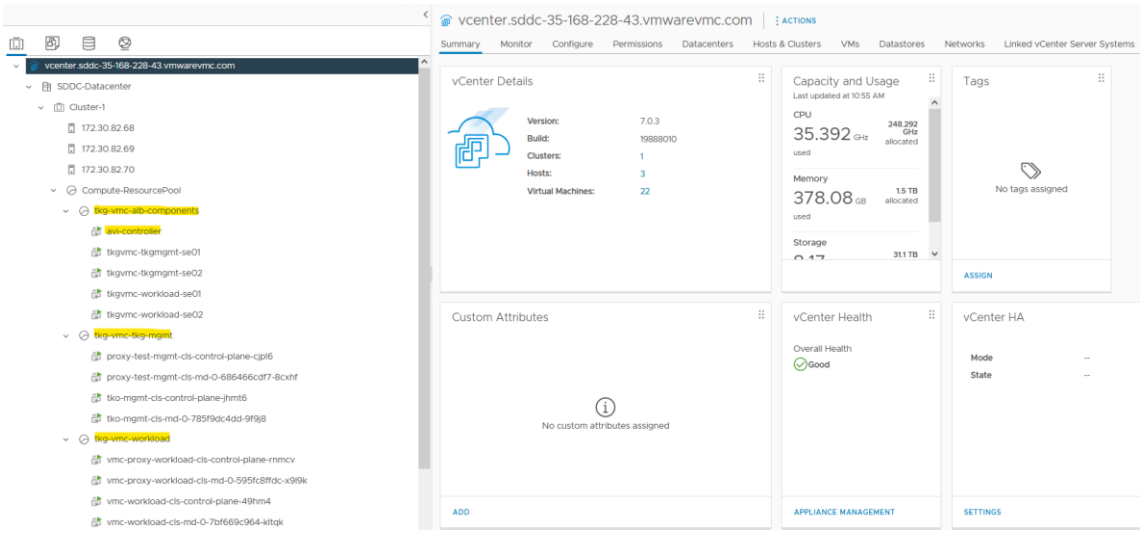

Note: Tanzu CLI and accompanying tools will be required to manage clusters. Visit Install Tanzu CLI for installation instructions.

Amazon Elastic Kubernetes Service (EKS)

Amazon Elastic Kubernetes Service is a managed service that can be used to run Kubernetes on AWS without needing to install, operate, and maintain your own Kubernetes control plane or nodes. Amazon EKS runs a single tenant control plane for each cluster. This means that the same control plane infrastructure cannot be shard across clusters. A customer

has the option to create worker nodes as self-managed Amazon EC2 nodes or to deploy their applications workloads to AWS Fargate, which is a serverless compute engine. For this reference architecture Amazon EC2 nodes were used. For more information on Fargate visit AWS Fargate page.

Amazon EKS clusters can be deployed either via "eksctl" utility or via AWS EKS user interface via AWS console. When eksctl utility is used with default settings, a new Amazon Cloud Formation stack is created which also includes an associated VPC (Virtual Private Cloud) and subnets. This is a viable option when there is a need for a separate VPC for EKS clusters. For this reference architecture AWS console was used to create EKS clusters and nodes and integrated into an existing VPC. This process is outlined below.

Note: This process requires that an AWS VPC exists, and subnets configured per requirements. For this reference architecture a VPC was created with subnet connections and routing in place to the internal private network, which included connections to VMC on AWS as well as the on-premises environment through Equinix. This configuration is depicted in the previously mentioned guide Multi-Cloud Network Connectivity with Equinix Overview and figure 3. For more information on how to create AWS VPC please consult AWS documentation.

Creating EKS clusters

Amazon EKS make calls to other AWS services on your behalf to manage the resources that you use with the service. Prior to creating EKS clusters, an IAM (Identity Access Management) role needs to exist or be created for EKS service. Determine if your account already has a role named "eksClusterRole." If it does not exist, then create one as follows. The example below shows the role created for this reference architecture.

Create IAM Role

1. From IAM > Roles menu create a role. Select "AWS Service" and search for "EKS" in "Use cases for other AWS services" search field. Select EKS Clusters and click next.

Figure 77: Create IAM Role

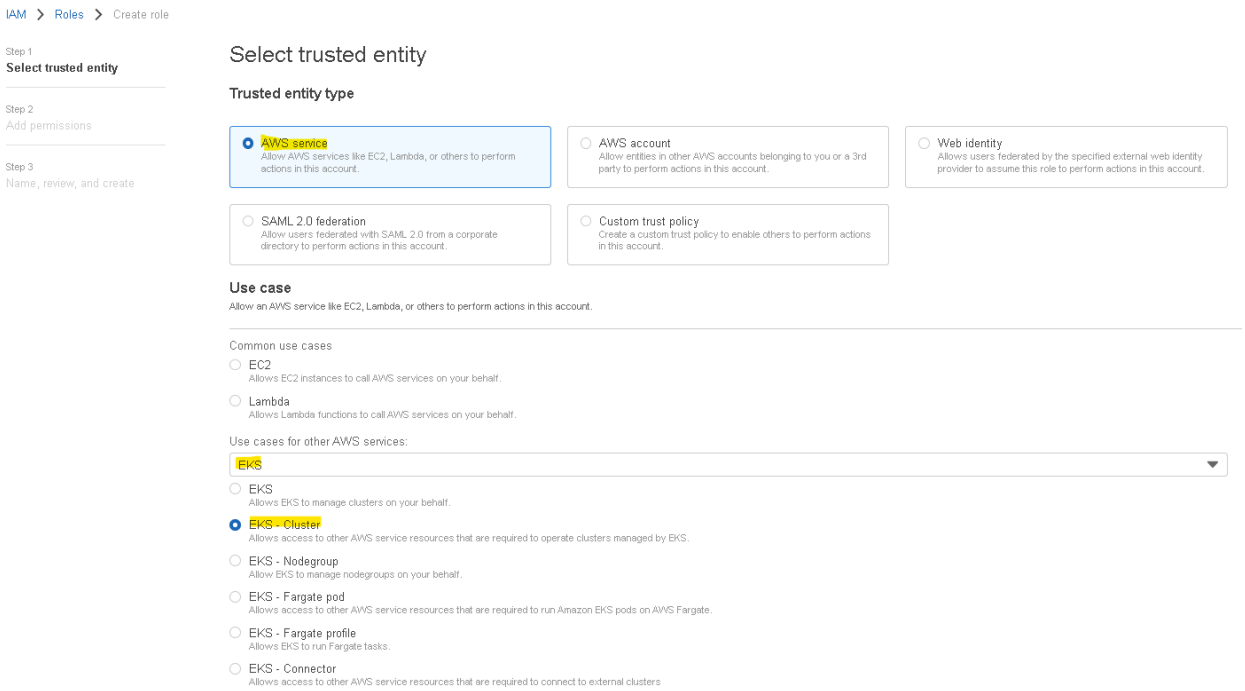

2. On "Add Permissions screen, Amazon EKSClusterPolicy will be automatically added. Click Next.

Figure 78: Add permissions.

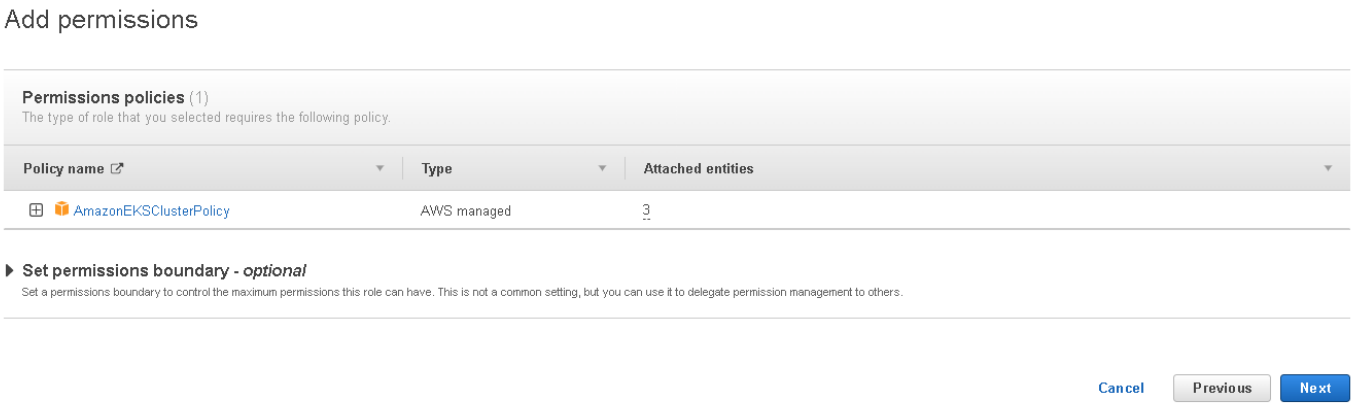

3. Give the role a meaningful name and click Create Role.

Figure 79: Create role.

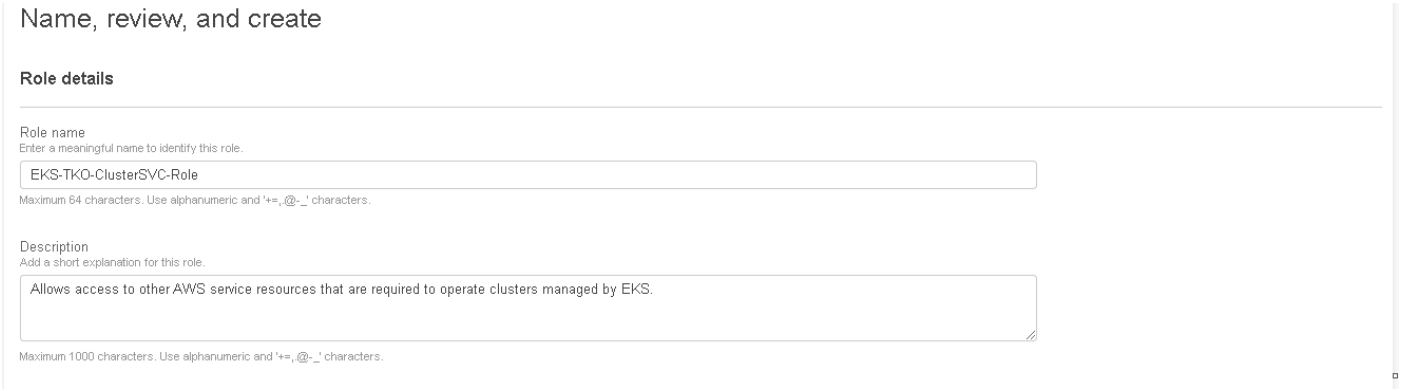

4. Edit the role just created and select “Attach Policies.”

Figure 80: Attach policies.

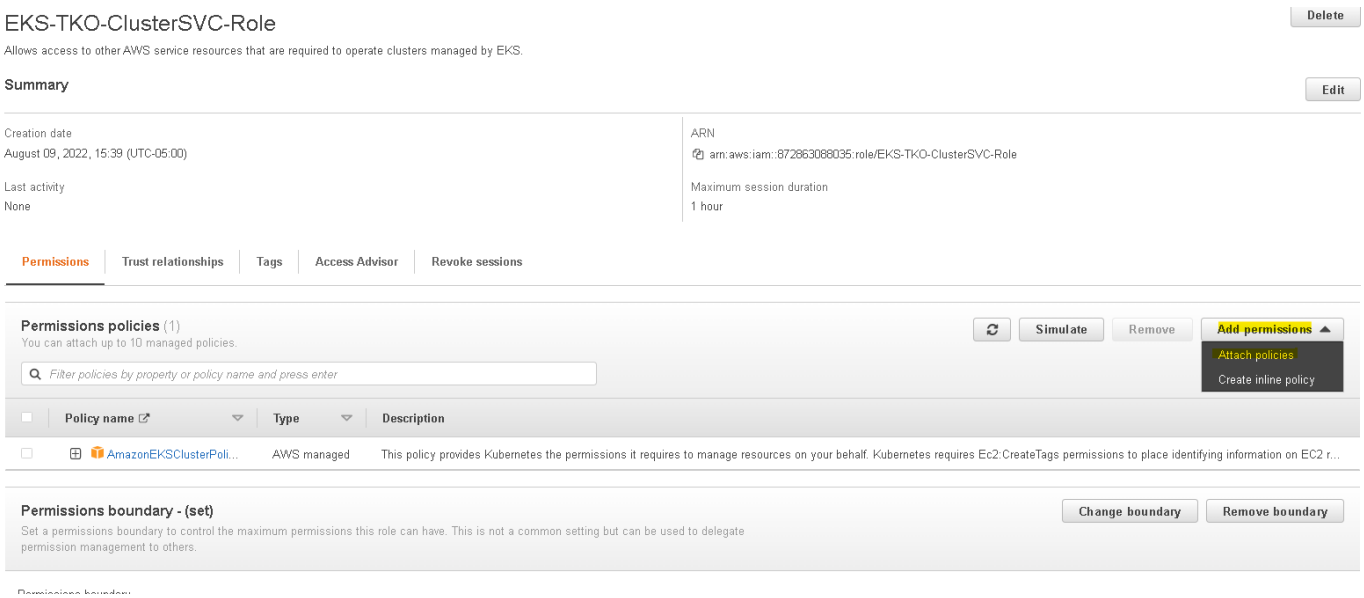

5. On the next screen search for “AmazonEKSVPCResourceController” and select “Attach Policy”

Figure 81: Attach polices.

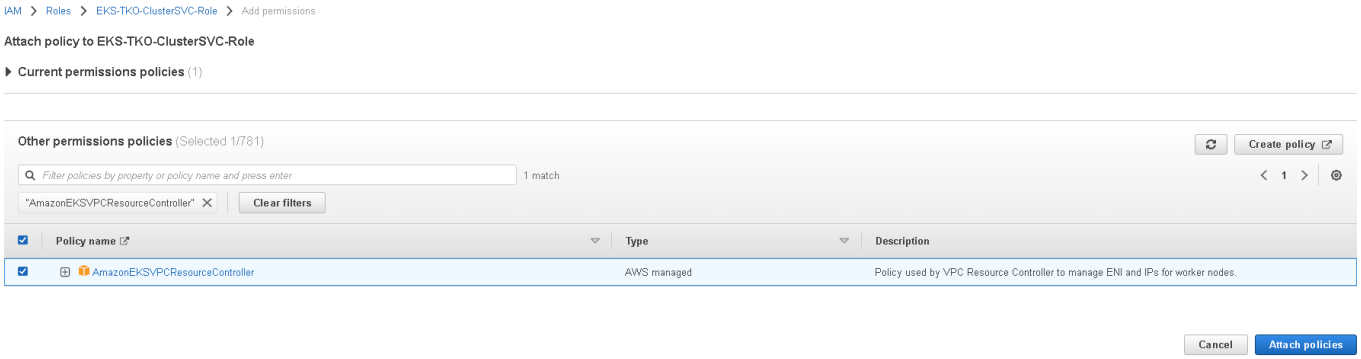

6. If Amazon CloudWatch monitoring is desired an inline policy needs to be created for CloudWatch to receive the clusters metrics. Create an incline policy by selecting “Create Inline Policy.”

Figure 82: Inline Policies

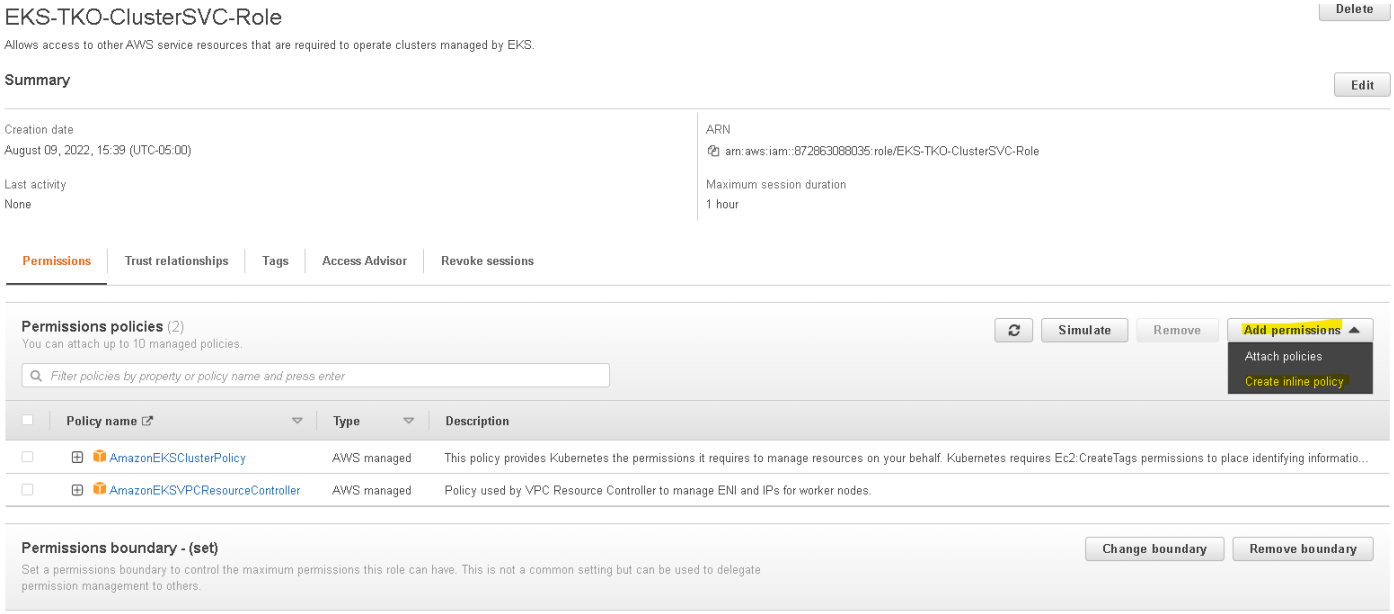

7. On the next screen select “JSON” and enter or paste the policy definition below and attach policy.

```
Inline cloud watch
{
   "Version": "2012-10-17",
   "Statement": [
      {
         "Action": [
            "cloudwatch:PutMetricData"
         ],
         "Resource": "*",
         "Effect": "Allow"
      }
   ]
}
```

Figure 83: Inline polices.

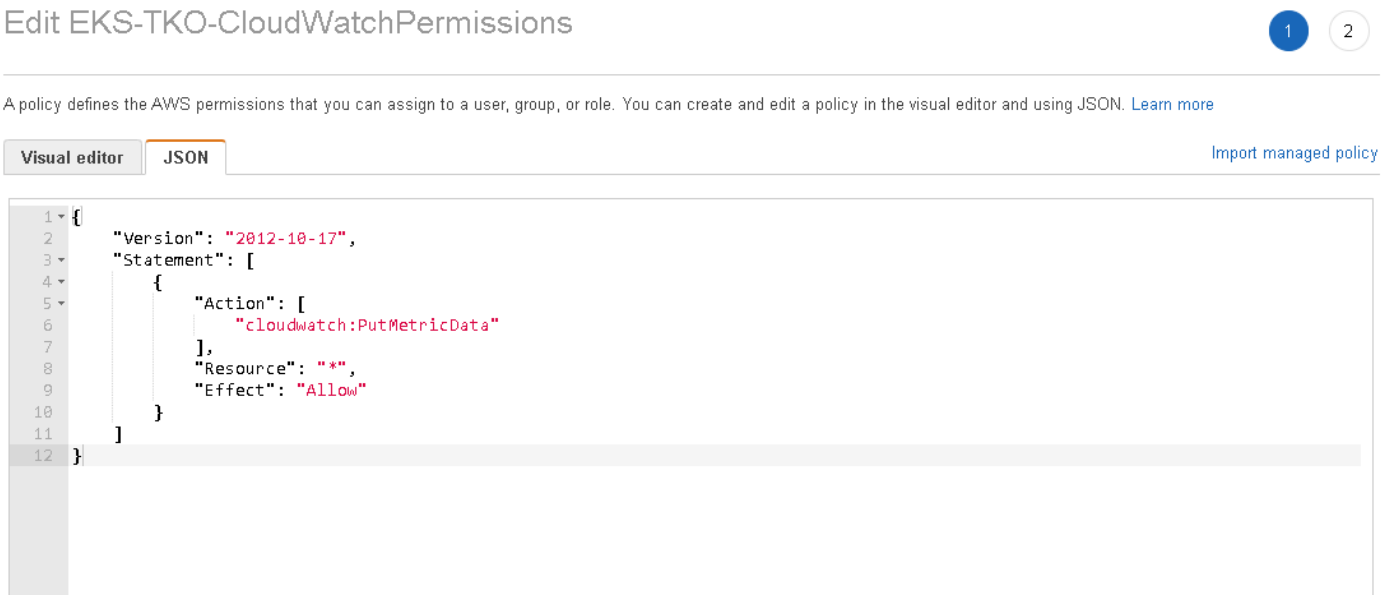

8. Once created, the IAM role should have the following policies. EKS clusters can now be created.

Figure 84: IAM role created.

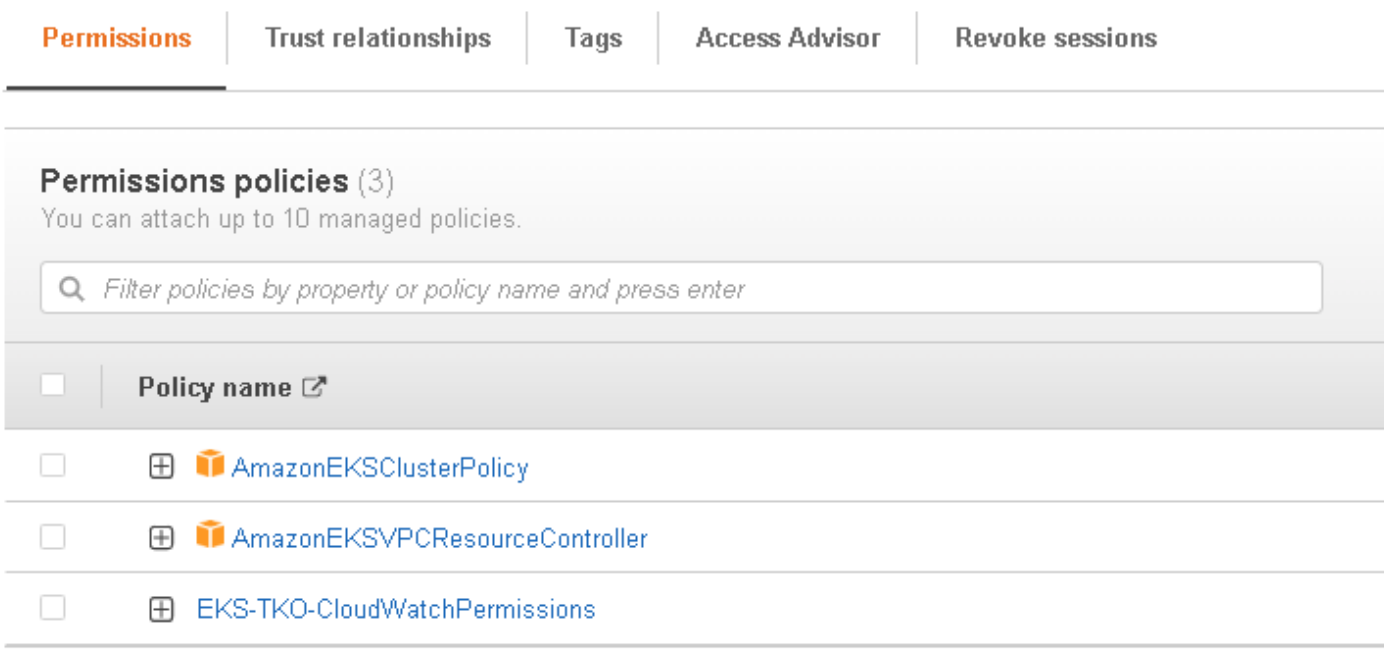

Create EKS Cluster

1. In Amazon console, navigate to EKS > Clusters > Create EKS cluster. Assign cluster a name, choose Kubernetes version and select cluster service role created earlier. Click Next.

Figure 85: Create EKS cluster.

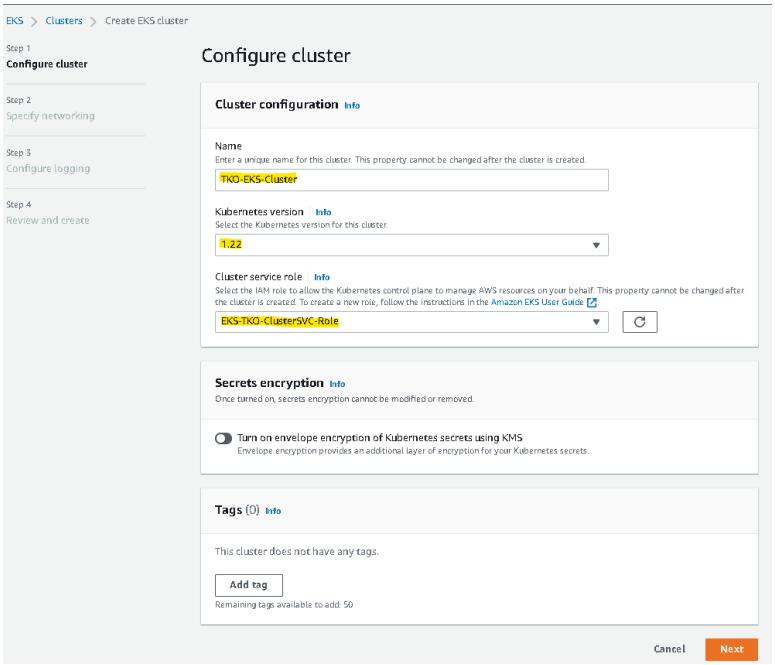

2. Specify VPC and subnets to be used. EKS cluster creation wizard adds all subnets available in the VPC by default. Remove any unwanted subnets. For this reference architecture only two subnets in two availability zones are required.

Figure 86: Set networking.

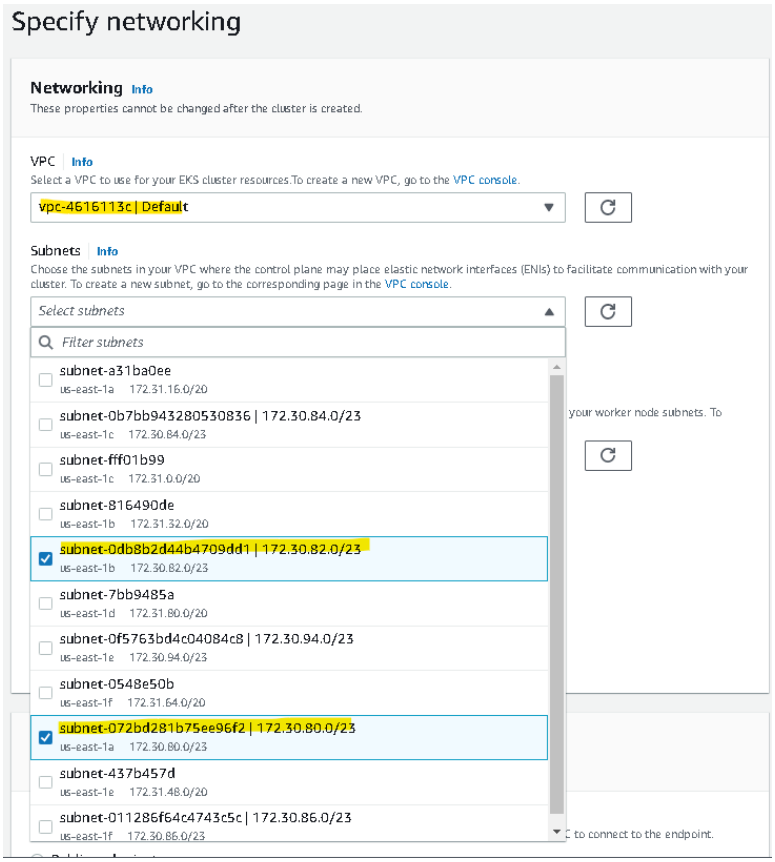

3. Select the desired security group. The default security group for the VPC was selected here.

Figure 87: Security groups

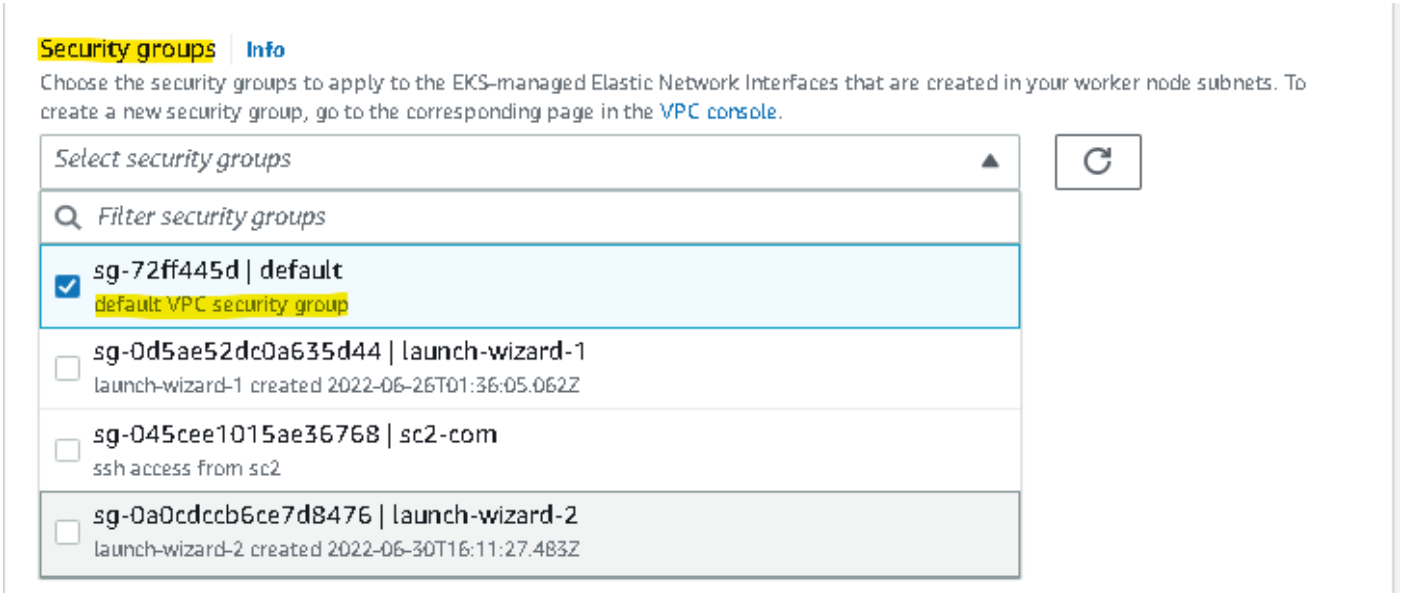

4. Select the Cluster endpoint access type. This depends on your environment and cluster requirements. Amazon EKS creates an endpoint for the managed Kubernetes API server that you use to communicate with your cluster (using Kubernetes management tools such as kubectl). By default, this API server endpoint is public to the internet, and access to the API server is secured using a combination of AWS Identity and Access Management (IAM) and native Kubernetes Role Based Access Control (RBAC).

Figure 88: Endpoint access

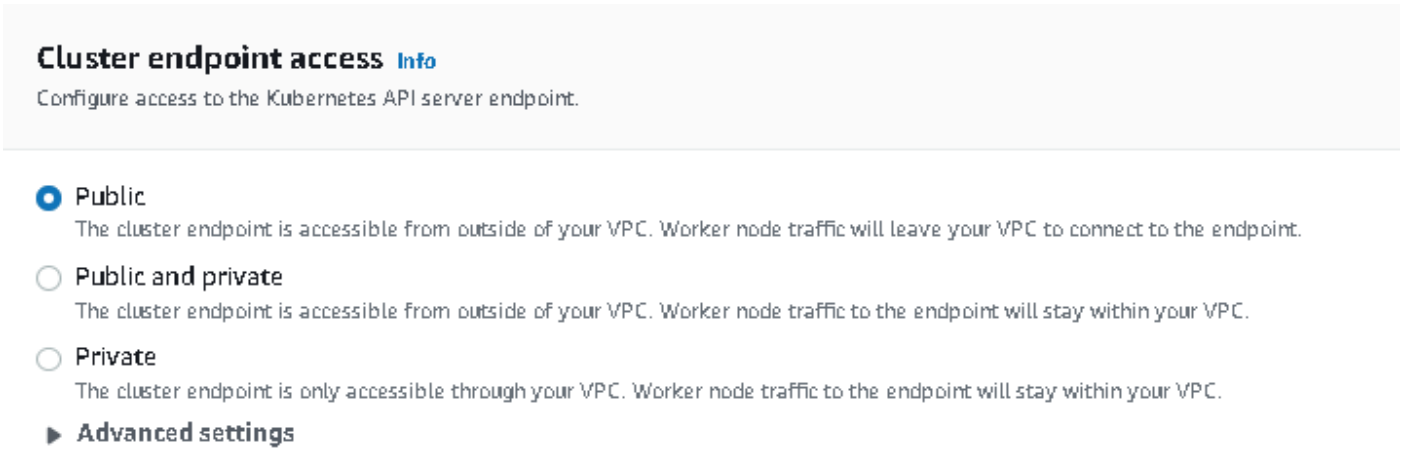

5. Select any specific versions of CNI, CoreDNS and Kube-proxy you require. Defaults were used as shown below.

Figure 89: Add-ons

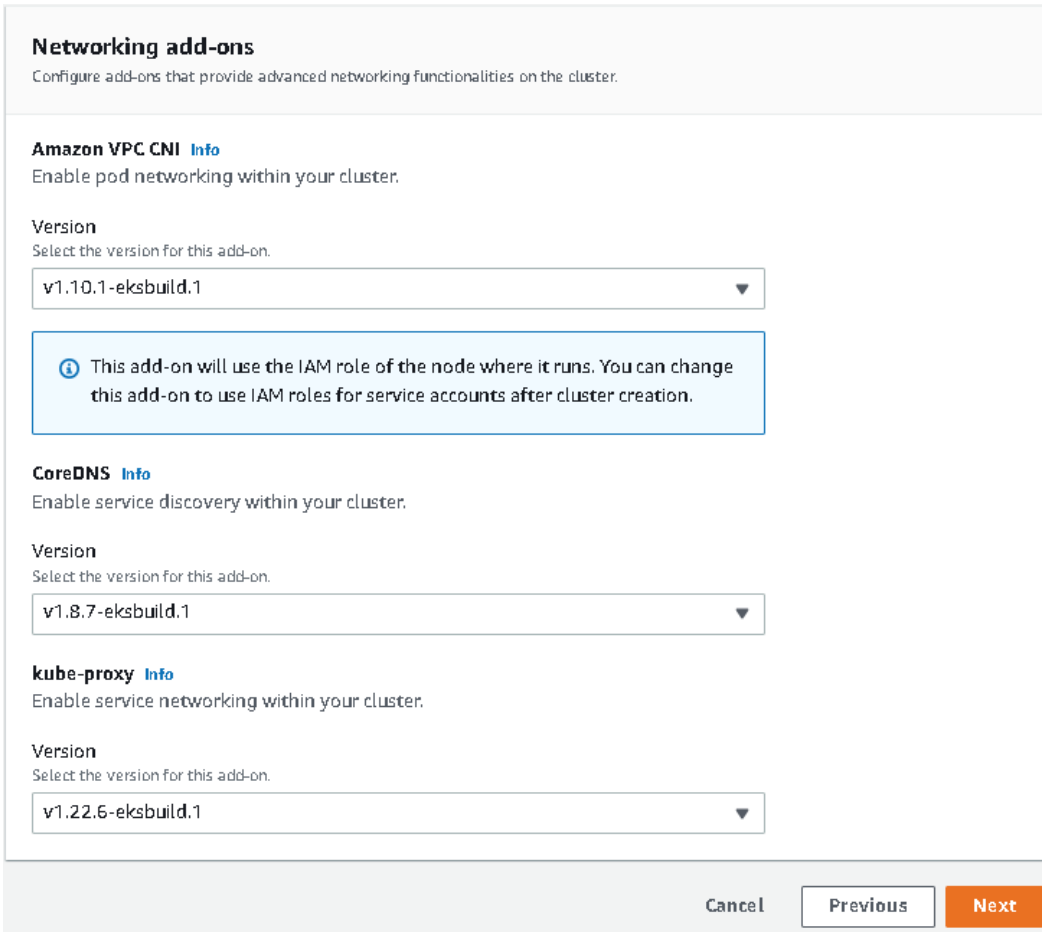

6. Configure Logging options.

Figure 90: Logging

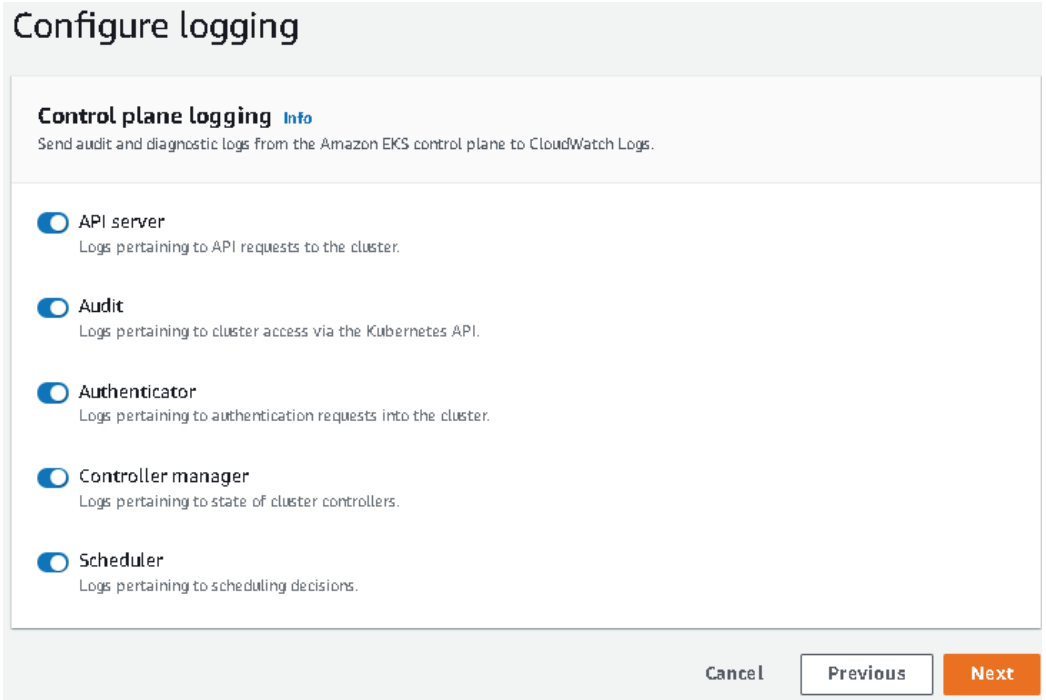

7. Finally review configuration and create the cluster.

Figure 91: Finish cluster creation

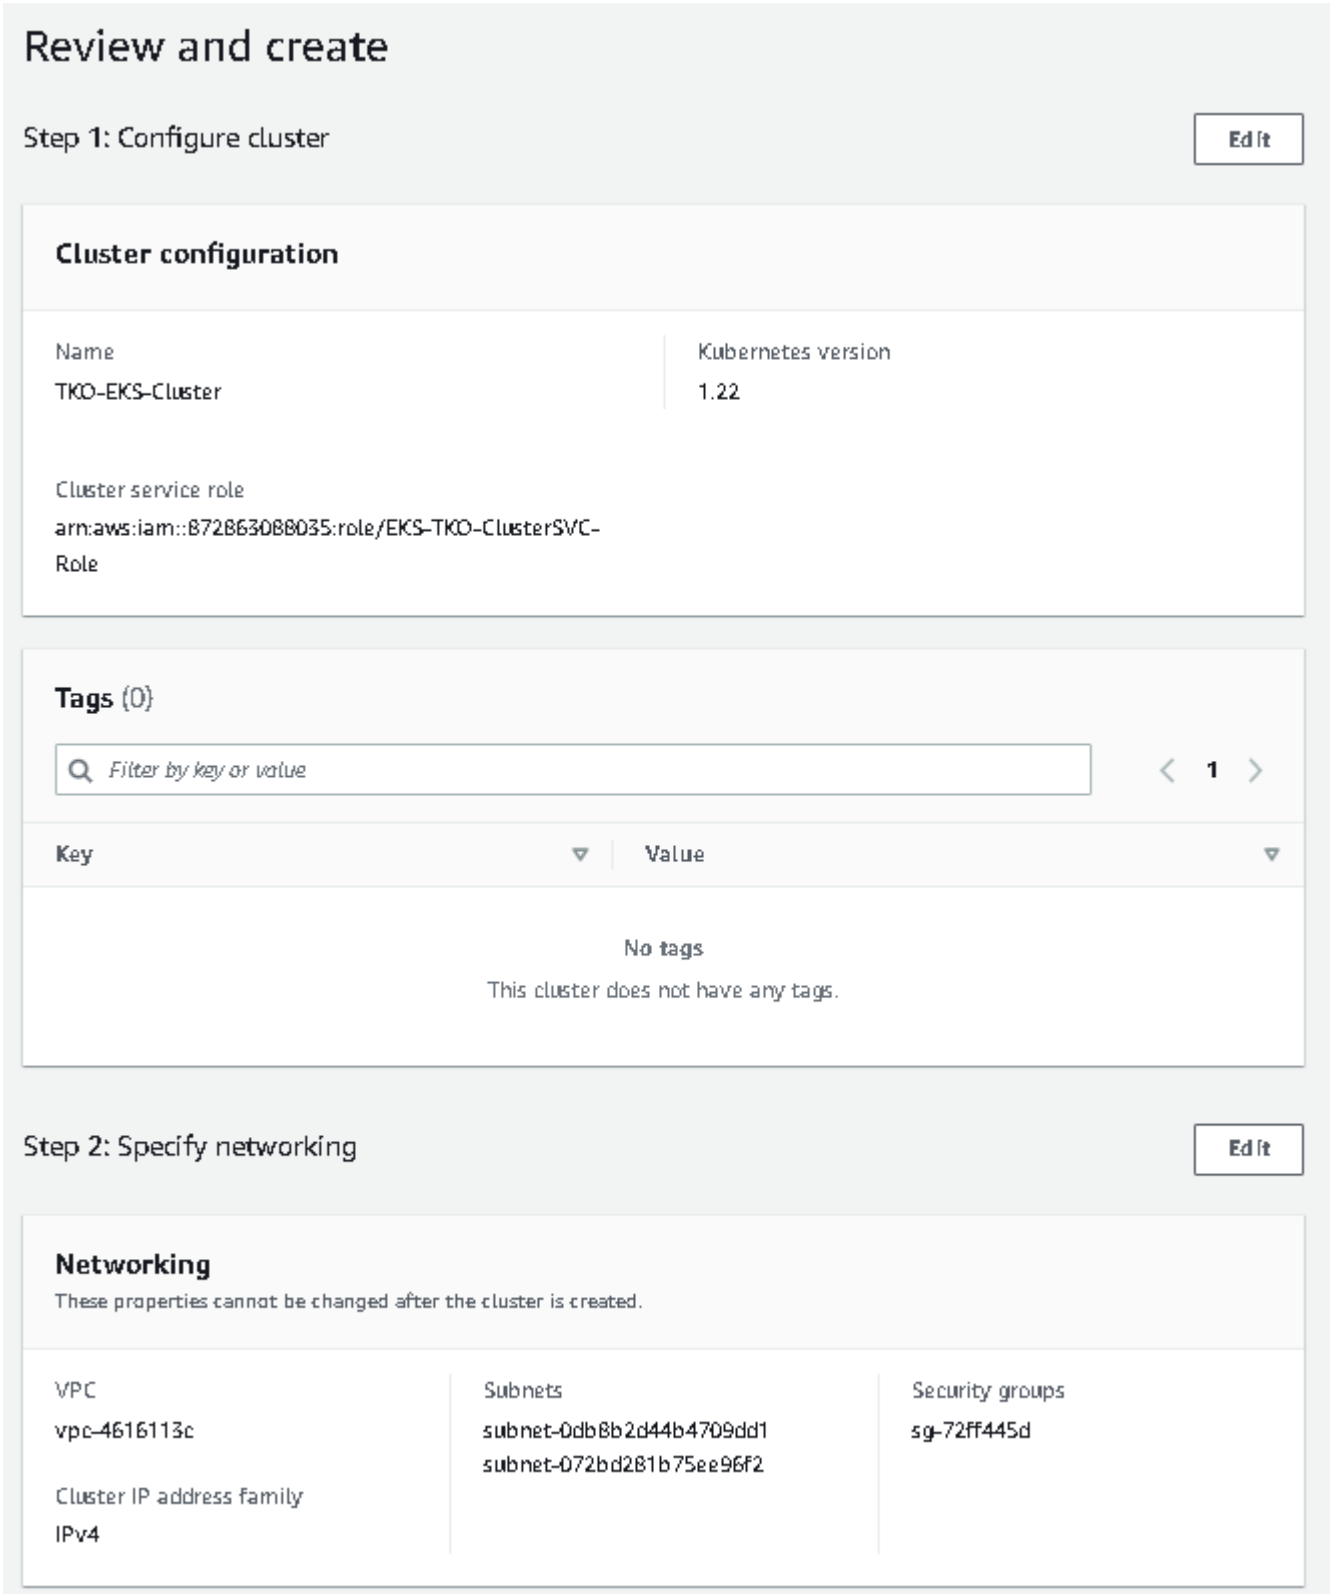

Create Node Group

A node group now needs to be created where your workloads will run. These will be Amazon EC2 AMI (Amazon Machine Image) instances. Prior to creating a node group, a "Node Instance Role" needs to be created via IAM > Roles > console.

1. Create node instance role and add the required policies for the EKS nodes. These policies will be used in the next step when creating node group. These are AWS managed policies and can be added by searching via the "Permissions policies" search bar.

Figure 92: Node instance policies

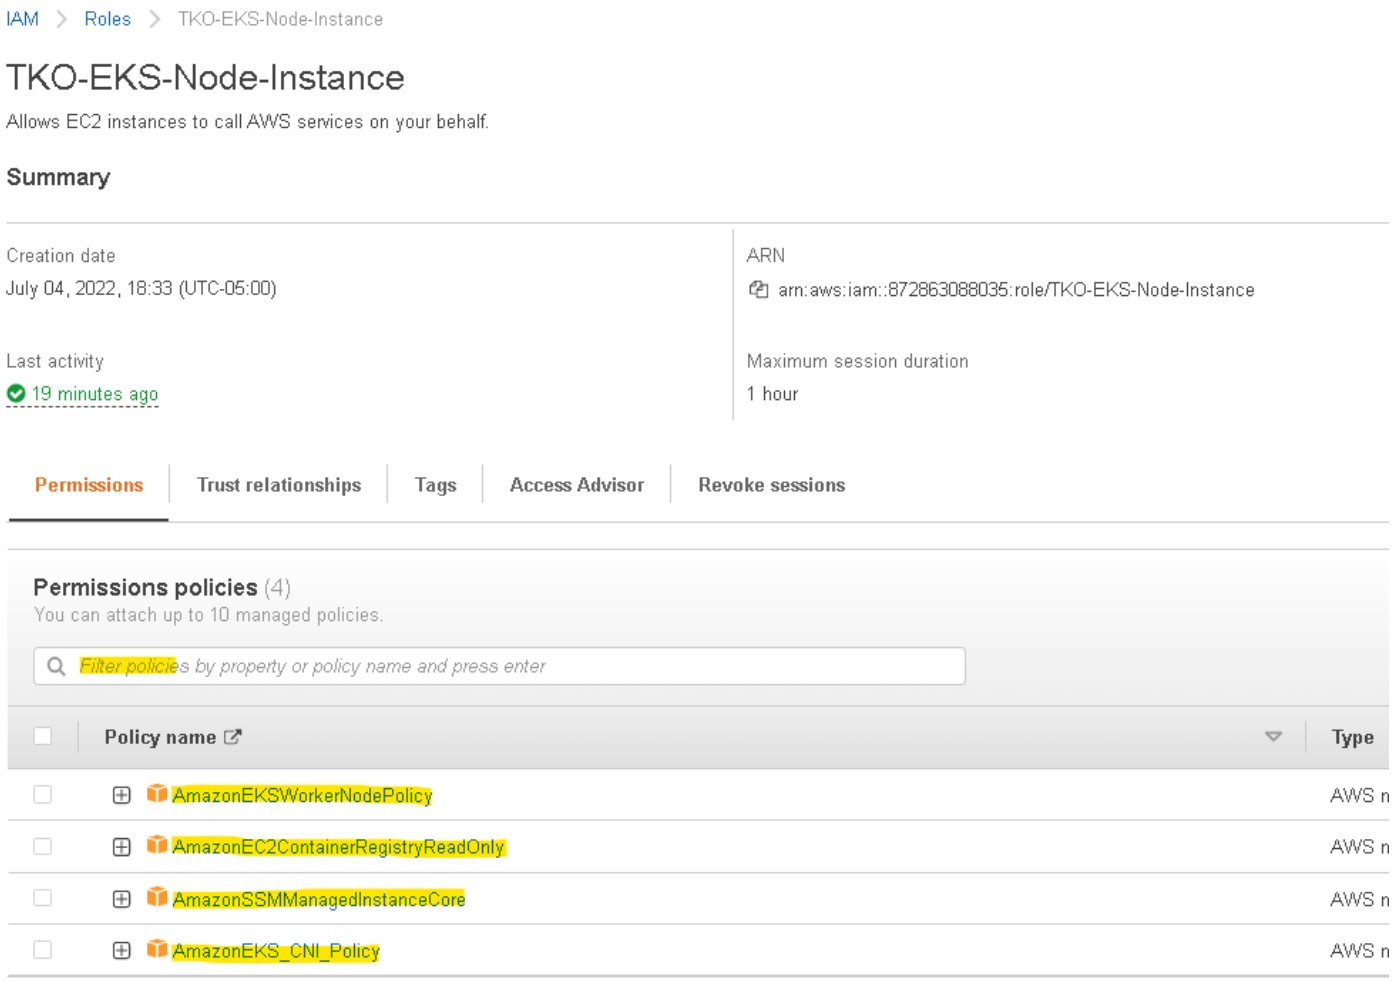

2. Navigate to your EKS cluster and click Add node group. Give the node group a name and select the IAM role created earlier. Leave everything else at default values and click next.

Figure 93: Create node group.

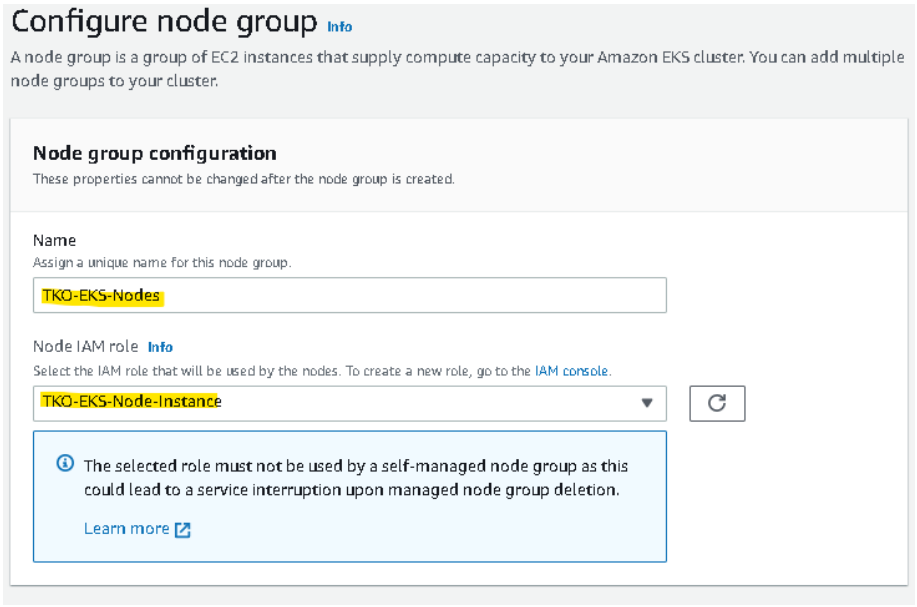

3. Select your desired node configurations and click next.

Figure 94: Set configurations.

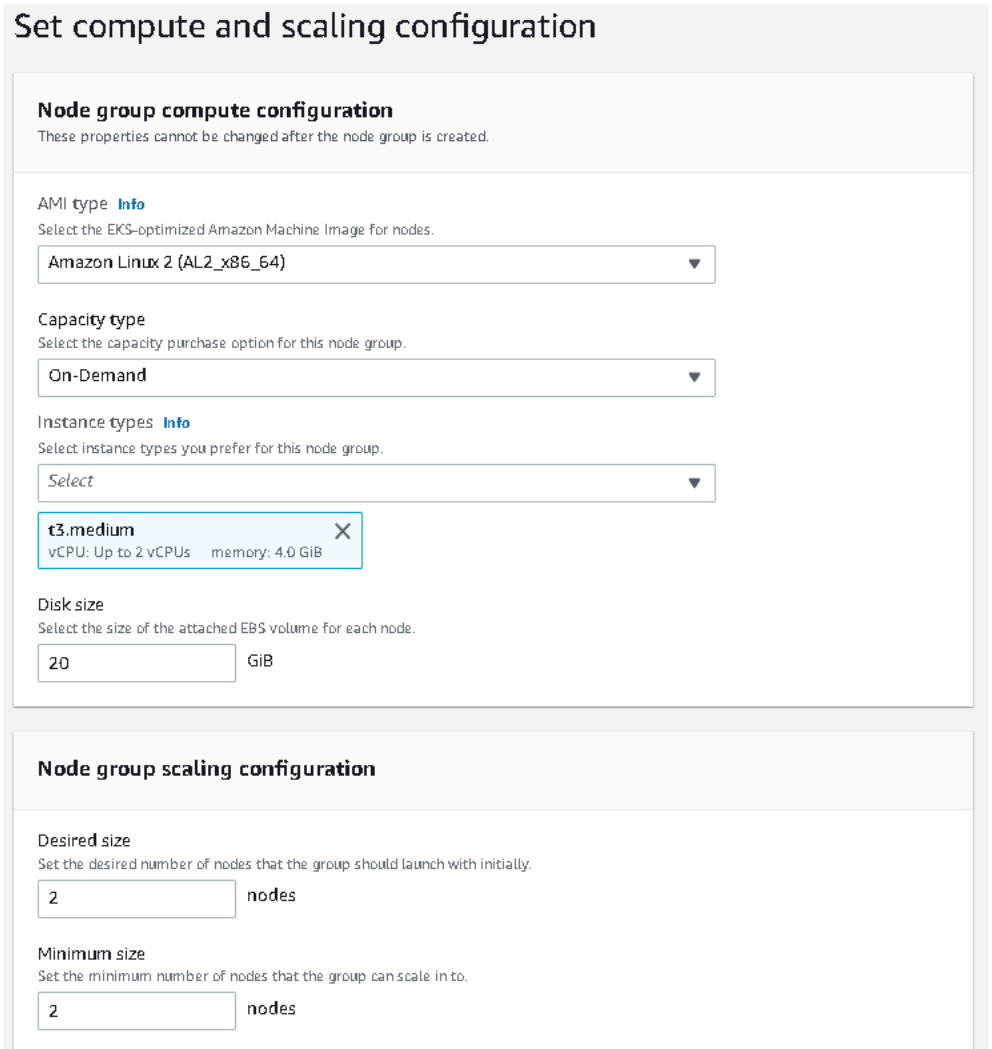

4. On the next page the subnets selected during cluster creation will automatically be selected. If node access via SSH, enable the option. On the next screen review and create the node group.

Note: SSH option requires that SSH key pair is already defined. In addition, when enabling this option, managed node groups will create a security group with port 22 inbound access. If launching your worker in a public subnet, it is strongly recommended to restrict the source IP address ranges.

Figure 95: Cluster networking

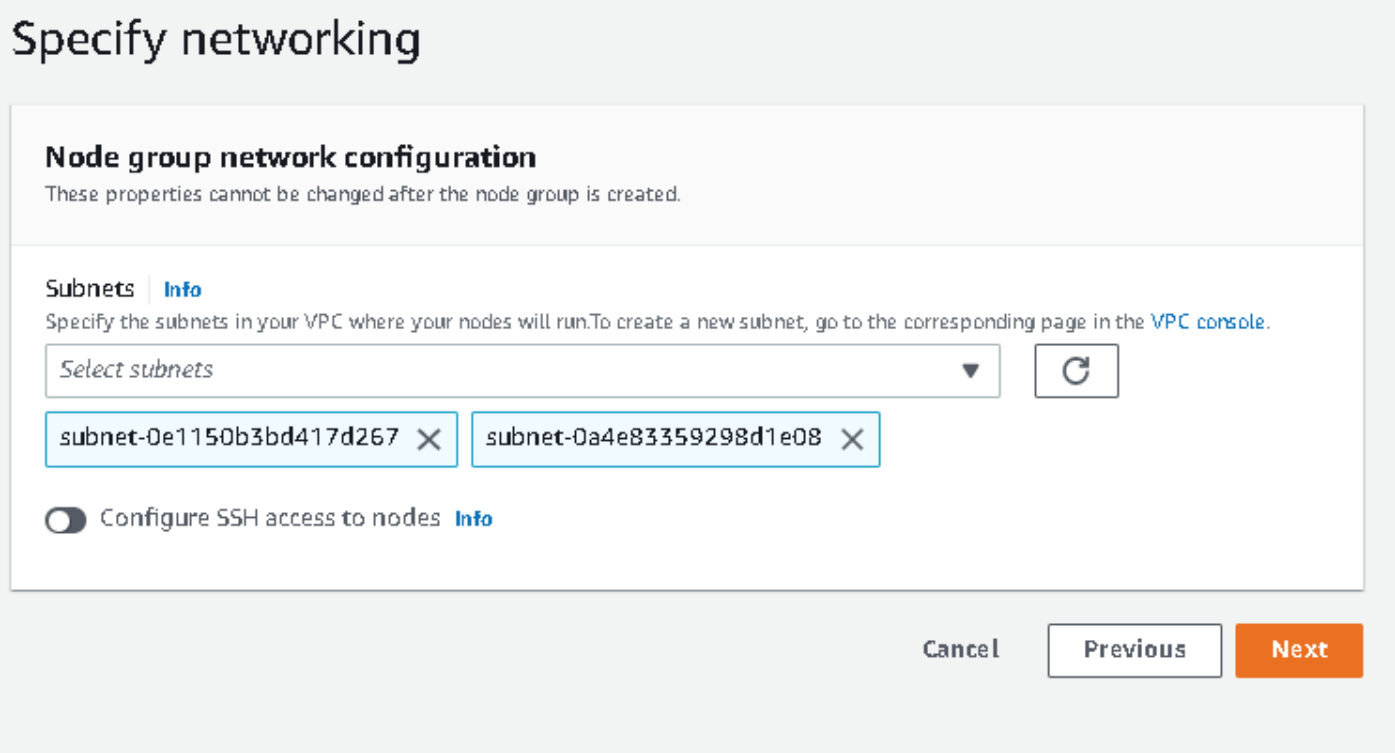

5. Once the node group has been created, node status can be verified on the cluster's "Compute" tab.

Figure 96: Node status

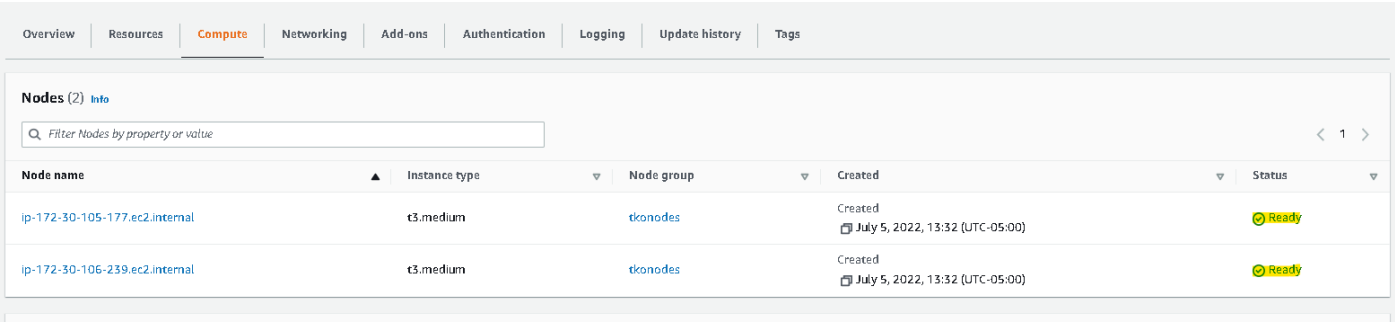

Installing NSX Advanced Load Balancer on AWS

NSX Advanced Load Balancer (Avi) is available as a subscription from Amazon Marketplace.

Figure 97: NSX Advanced Load Balancer (Avi) subscription

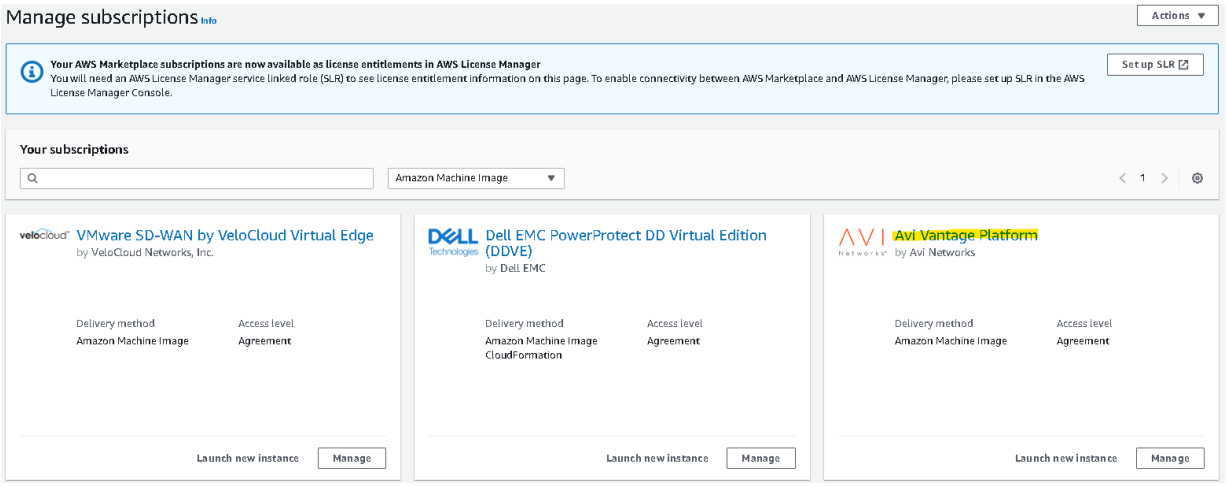

NSX Advanced load balancer runs as an Amazon AMI on the VPC and provides services comparable to a traditional datacenter installation. Deployment process is as follows. Subscribe to the service prior to launching the AMI installation.

1. With a valid subscription, launch a new instance.

Figure 98: Launch Avi

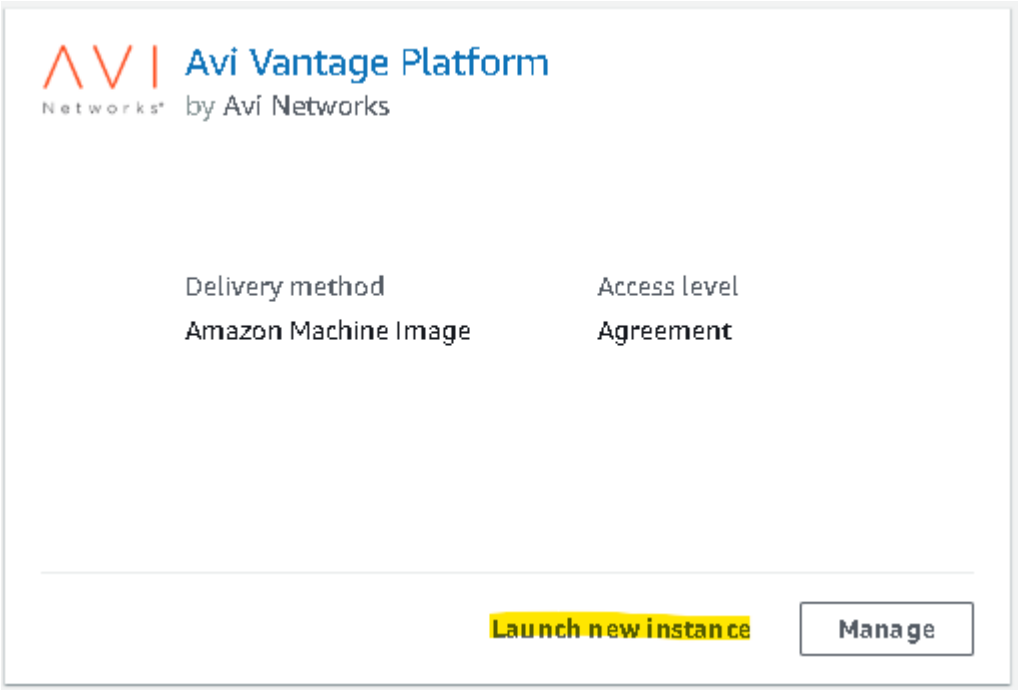

2. Select the region and version. By default, only the latest version will be displayed. To change the version, click on the “full AWS Marketplace website” link.

Figure 99: Launch Avi

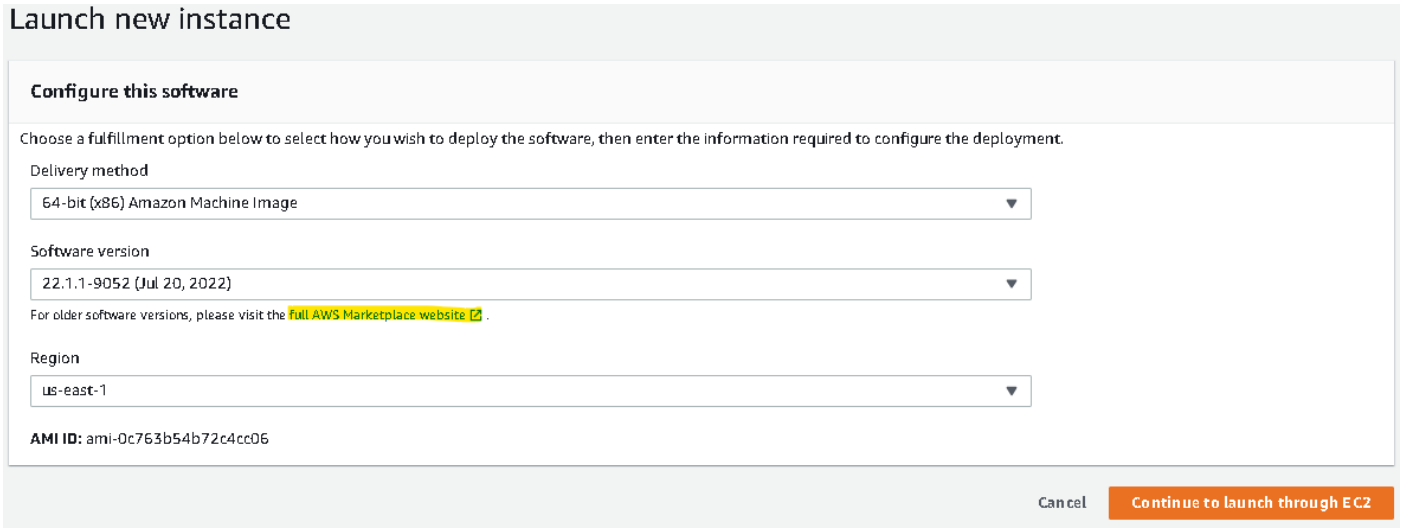

3. For this reference architecture 21.1.4 was used. Click "Continue to launch."

Figure 100: Launch Avi

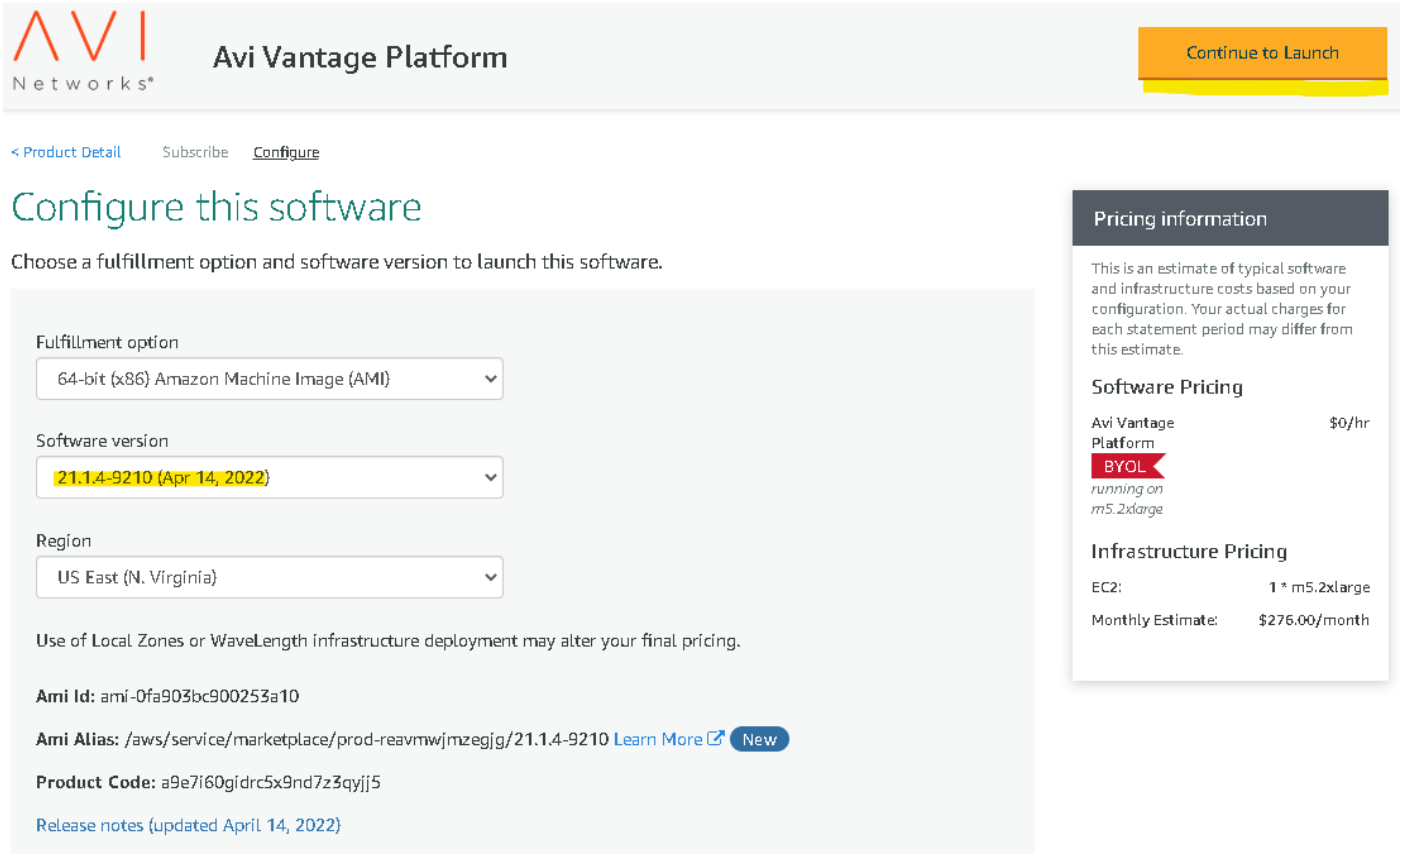

4. Select "Launch through EC2" and click "Launch."

Figure 101: Launch Avi

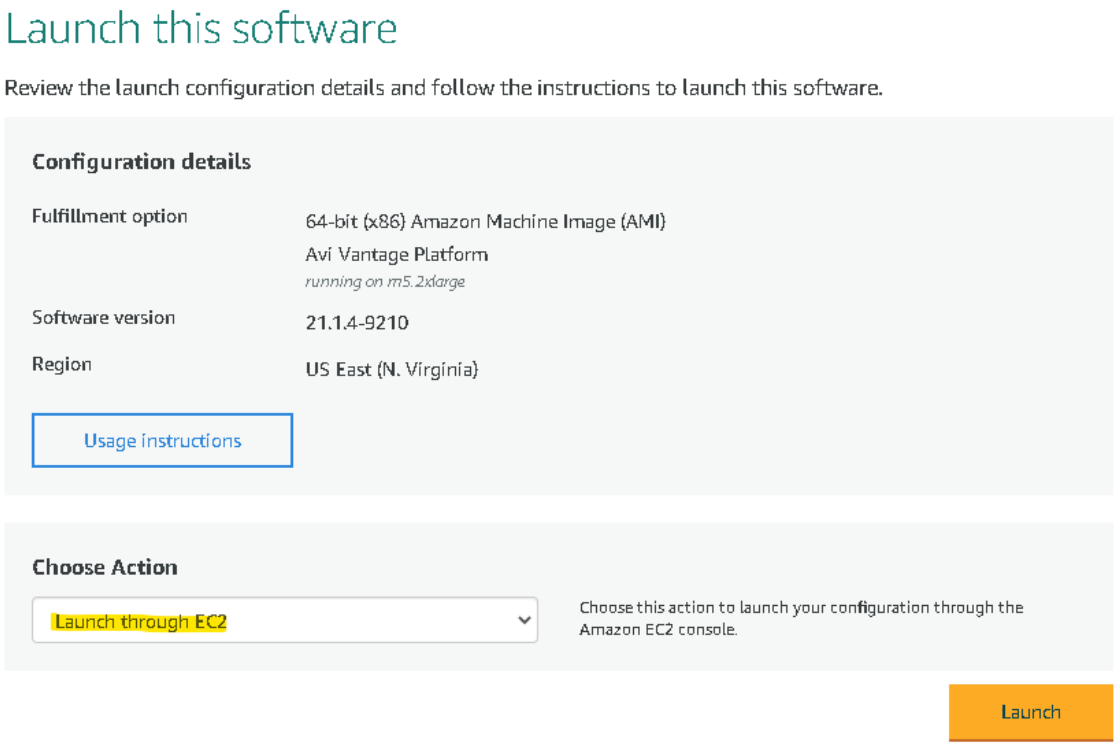

5. On the next page, give the instance a name.

Figure 102: Launch Avi

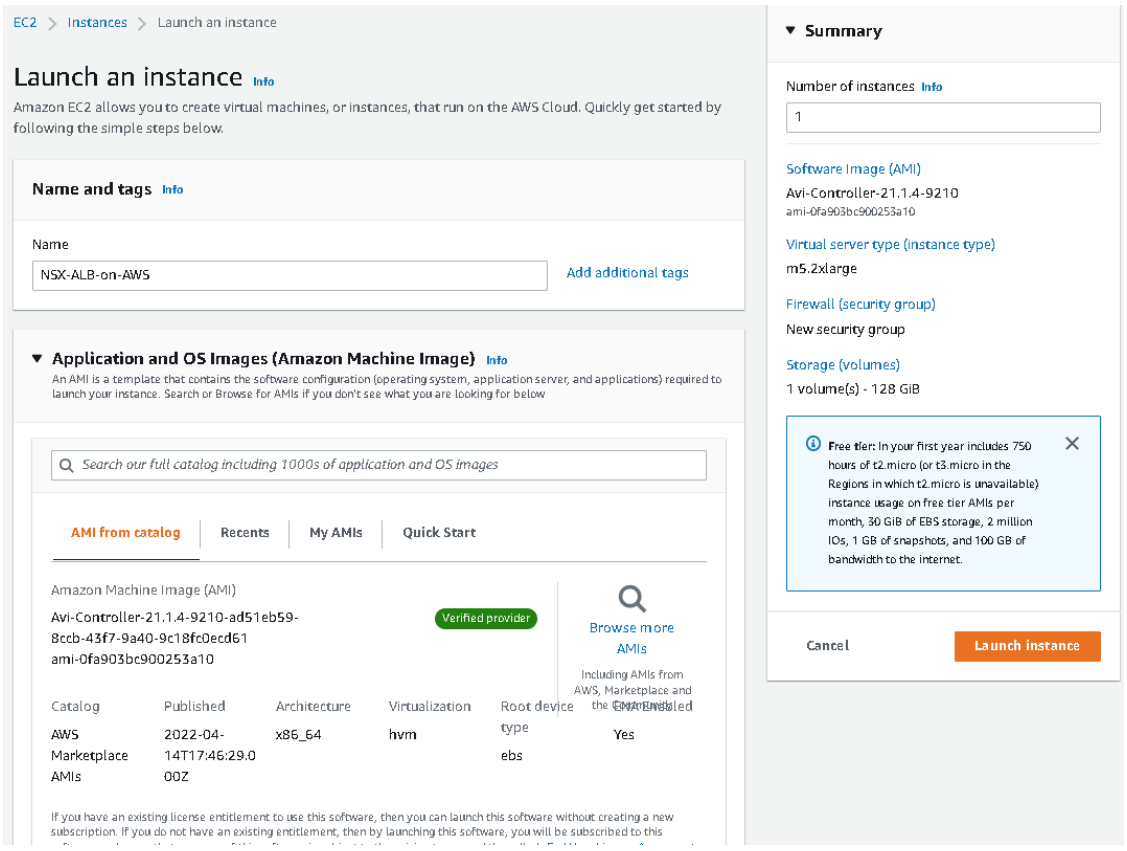

6. On the same page select a SSH key pair. Create one if one does not exist. This will be needed when accessing the controller via SSH. In “Network settings,” verify the VPC, subnet options and traffic rules. Inbound SSH traffic

can be filtered based on security groups or CIDR or IP address bases. Default security group for the VPC is selected in this case. Select “Create security group” option. This option will create a security group for the controller based on recommended settings.

Figure 103: Launch Avi

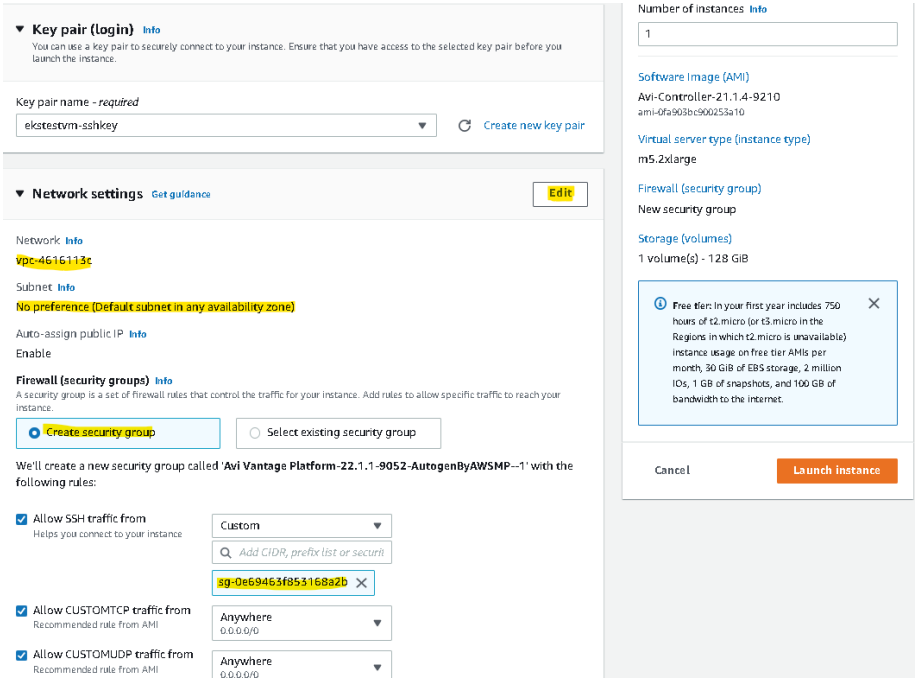

7. Edit the network configuration to modify the default settings for VPC and subnets. This will also show you the name of the security group that will be created as well as VPC and subnet selection options will be available.

Figure 104: Launch Avi

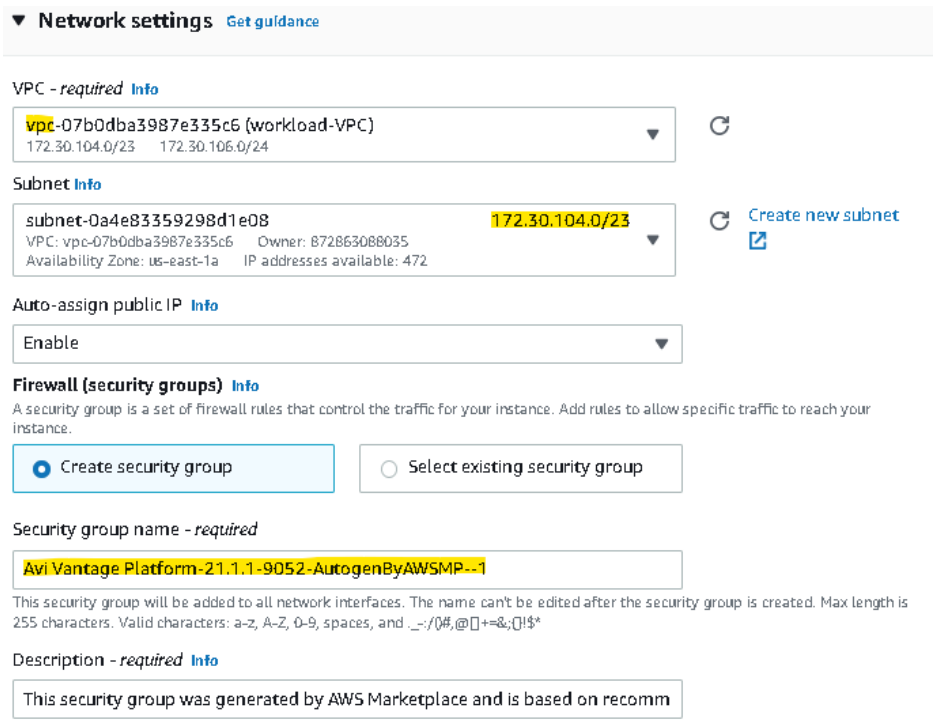

8. Click Launch Instance and the deployment process will start. This will take a few minutes. Once the process is complete; the controller will be available in EC2 console and can be accessed from the public or private IPs assigned during provisioning.

Figure 105: Avi created.

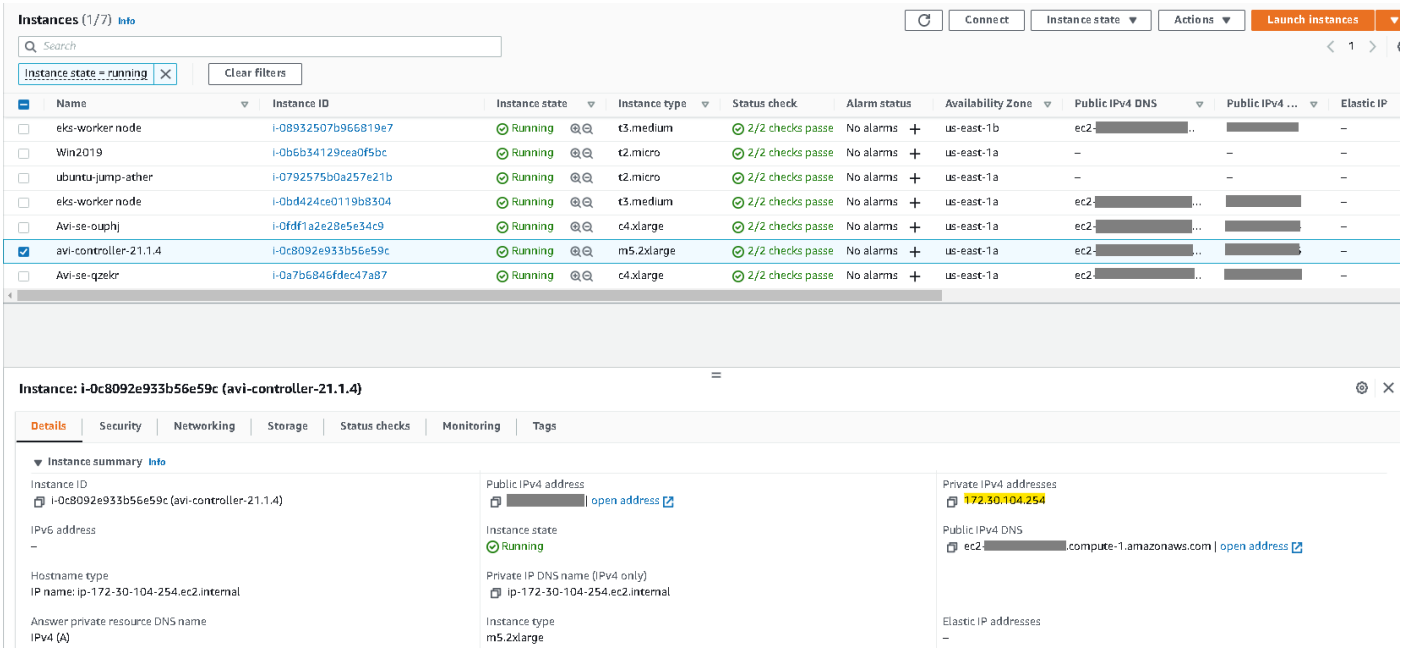

Configuring NSX Advanced Load Balancer on AWS

Prior to configuring NSX Advanced Load Balancer on AWS a credential method must be chosen. There are several methods AWS credential can be assigned to NSX Advanced Load Balancer components.

- **AWS customer account key:** A unique authentication key associated with the AWS account. Access credentials are needed by the NSX Advanced Load Balancer Controller to communicate with AWS APIs.

  Note: AWS cloud configuration with NSX Advanced Load Balancer SaaS only supports the Use Access/Secret Key credentials method.
- **Identity and Access Management (IAM) roles:** IAM roles are the set of policies that define access to resources within AWS. The roles and the policies that define their access are defined in JSON files. This method does not require an AWS account key. Instead, the role and policy files must be downloaded from NSX Advanced Load Balancer and installed using the AWS CLI. Use this method if you do not want to enter AWS credentials.
- **Use Cross-Account AssumeRole:** NSX Advanced Load Balancer can be deployed for Amazon Web Services (AWS) with multiple AWS accounts utilizing the IAM AssumeRole functionality that provides access across AWS accounts to the AWS resources/API from the respective accounts, instead of sharing user Access Key ID and Secret Access Key from different accounts.
- For the detailed information on Cross-Account AssumeRole, refer to AWS Cross-Account AssumeRole Support.

For this deployment **IAM roles** method was used. This method does not require a customer account key. Instructions on creating AWS required roles can be found here These roles are required for NSX Advanced Load Balancer to gain access to AWS objects including AMI images that will be used to deploy service engine. Once the roles are configured, they will appear in AWS IAM > Roles.

Figure 106: Avi IAM refined role.

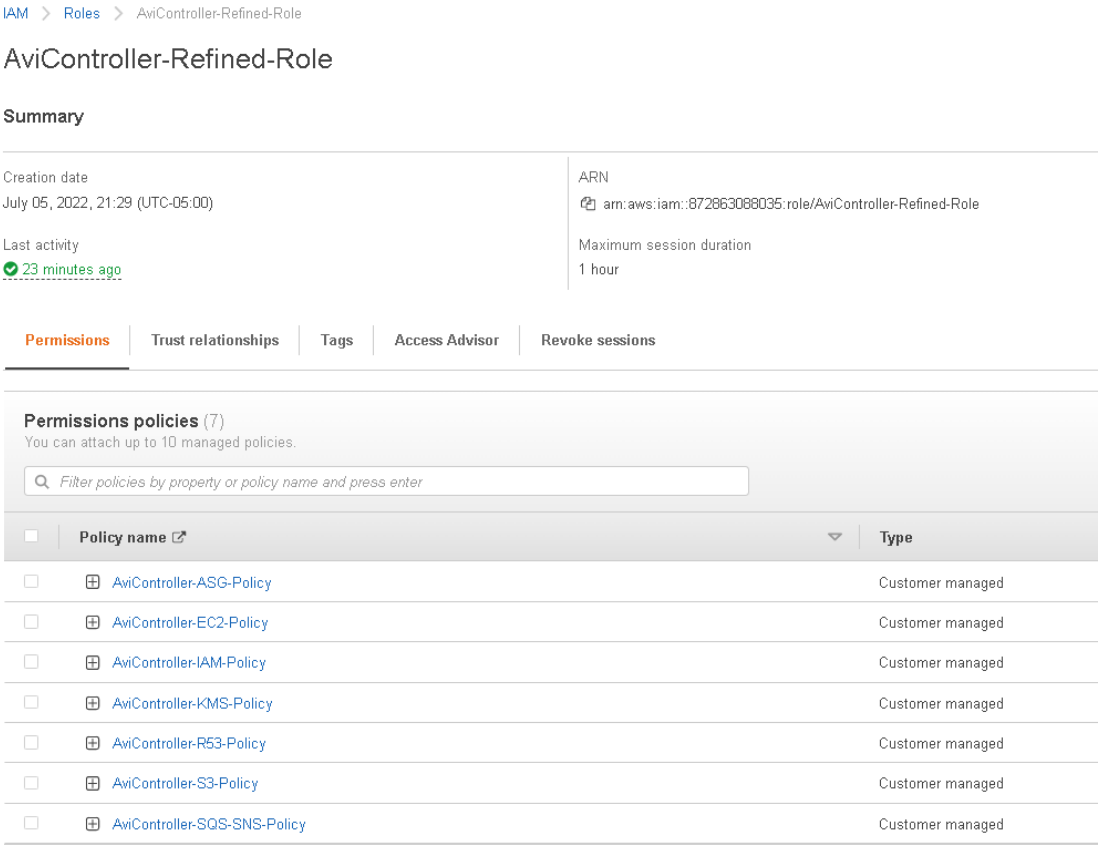

Figure 107: Avi vmimport role

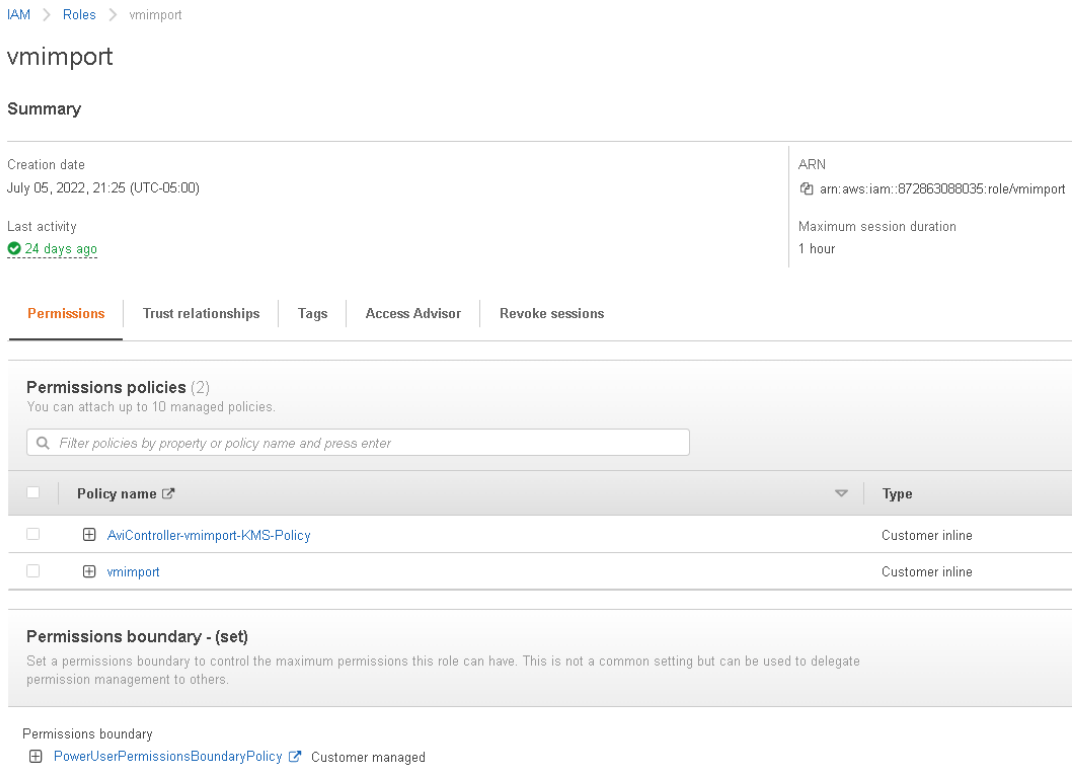

Once the roles are configured. NSX Advanced Load Balancer can be configured using the public or private IP address. Complete the NSX Advanced Load Balancer configuration per environment requirements. Instructions to configure NSX Advanced Load Balancer after provisioning can be accessed here.

Global DNS Namespace and Ingress

NSX Advanced Load Balancer is the key component that facilities automated application and resource discovery by building a unified and global DNS namespace. When an application is created, a service URL is automatically created by NSX Advanced Load Balancer components through which users can access the applications. AKO and AMKO are needed for multi-cluster ingress for services.

AKO and AMKO installation

Both AKO and AMKO are required for multi-cluster ingress and load balancing. AKO is installed on all participating clusters, where AMKO is installed on the leader cluster. An optional second instance of AMKO can be installed on a second cluster for redundancy. For this reference architecture a single instance of AMKO was installed that coordinated with AKO instances on all clusters.

AKO and AMKO can be installed Helm Charts. Helm is a pre-requisite and installation instructions can be found here. A high-level process is as follows.

1. Install Helm.
2. Create avi-system namespace "*kubectl create ns avi-system*."
3. Add AKO Helm chart "*helm repo add ako https://projects.registry.vmware.com/chartrepo/ako*"
4. Export AKO parameters to values.yaml file "helm show values ako/ako --version 1.7.1 > values.yaml"
5. Edit the value.yaml file per your environmental requirements. A sample yaml file used to deploy AKO on EKS clusters is in Appendix A for reference.
6. Install AKO: "*helm install ako/ako --generate-name --version <AKO version> -f values.yaml --set ControllerSettings.controllerHost=<IP of AVi controller> --set avicredentials.username=admin --set avicredentials.password=<controller password> --namespace=avi-system*"

AMKO is installed on the GSLB leader cluster in similar manner. For more information on value file parameters visit Install and Configure AMKO page. This page also provides information on AMKO federation, if required by the environment. For this reference architecture non-federated AMKO architecture was used.

Note: Yaml files exported during the above process sometimes do not output the yaml in a correct format. Double check the correct yaml format of these yaml files before installation, by using a yaml verification tool such as yamllint or another similar application.

Note: AWS creates external load balancer for service type "loadbalancer" by default. To avoid the creation of AWS load balancer, annotate the service with *" service.beta.kubernetes.io/aws-load-balancer-internal: "true" "*

Configuring GSLB with NSX Advanced Load Balancer

Configuring NSX Advanced Load balancer to serve as a global DNS namespace provider is a multi-step process and depends on the use case at hand. Below are some of the use cases for use of GSLB functionality.

GSLB Use Cases:

- Optimal application experience for geographically distributed users
  - Multiple applications are deployed in multiple data centers.
  - Avi GSLB can steer user traffic to the most optimal location.
- Application high availability across data center failures
  - Applications are deployed in multiple data centers.
  - In case of a data center failure, application instances running in the remaining data center(s) can take over the user traffic.
- Disaster recovery
  - Applications are deployed in two data centers.
  - While both are healthy, all traffic is directed to the primary DC.

    - If the primary DC fails, the global DNS directs all user traffic to the other.
- Hybrid cloud with “cloud bursting”
    - Applications are deployed across private and public clouds.
    - When/if an application experiences an unusually high request load, Avi GSLB “bursts” to the public cloud site to absorb the load.

The GSLB configuration discussed in this reference architecture addresses all the above use cases. “Cloud bursting” however was not tested but can be configured with the same architecture. What follows are the steps to configure GSLB on on-premises private cloud and Amazon EKS clusters. VMC or another public cloud instance can be added using the same steps.

Note: In the configuration example, pse.lab is the primary domain and avitko.pse.lab is the subdomain. In this architecture DNS authoritative server for **pse.lab** domain lives on the on-premises datacenter which is also the GSLB leader site. This is not a requirement. Authoritative DNS server can be on a follower site or integrated with external DNS service such as Route53. Subdomain **avitko.pse.lab** will be delegated to NSX Advanced Load Balancers which will serve requests to applications hosted on the sub-domain. For example, request to “app.avitko.pse.lab” will be resolved by the NSX Advanced Load Balancers.

Configuring Sites for GSLB

Create DNS virtual service:

Following are the steps for GSLB configuration for avi controllers on on-premises and on AWS. The leader controller resides on-premises.

1. On the on-premises Avi controller, create a DNS profile specifying the sub-domain avitko.pse.lab.

Figure 108: DNS profile

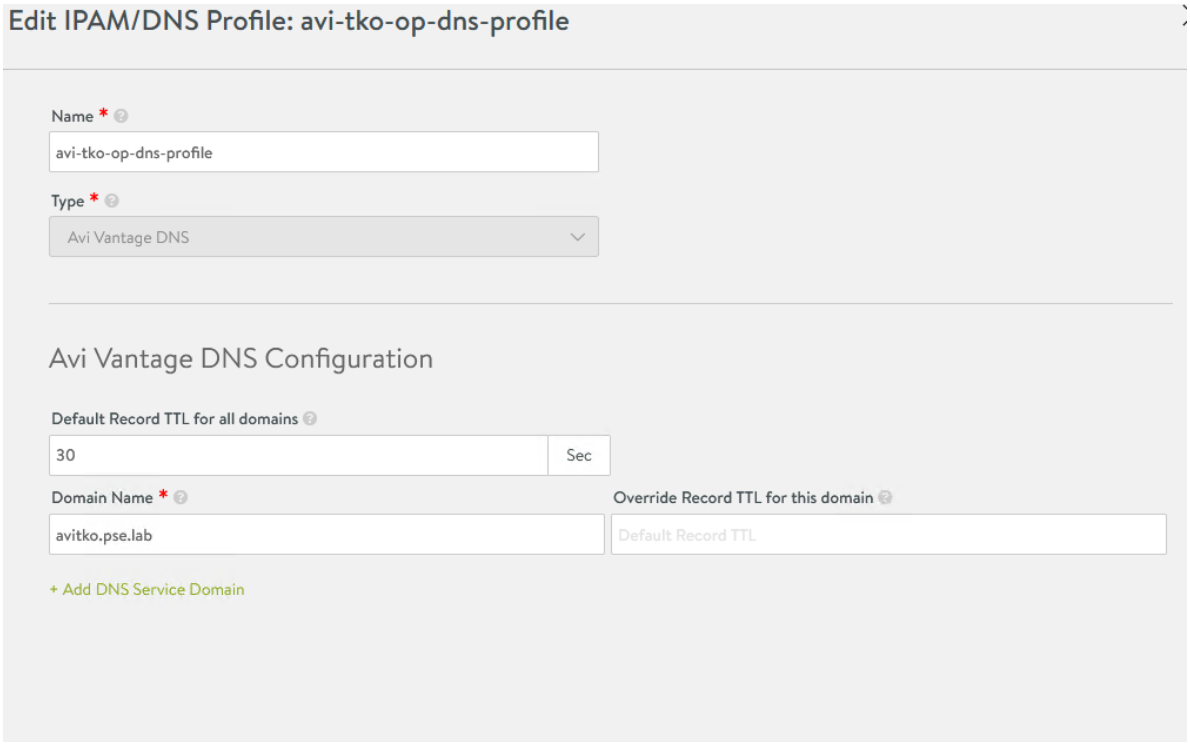

2. Create a service engine group for the DNS service. It is a recommended practice that separate service engine group be created for the DNS service. It is also good practice to pre-fix the service engines names with a distinguishable string. Example “Avidns- “

Figure 109: Service engine group

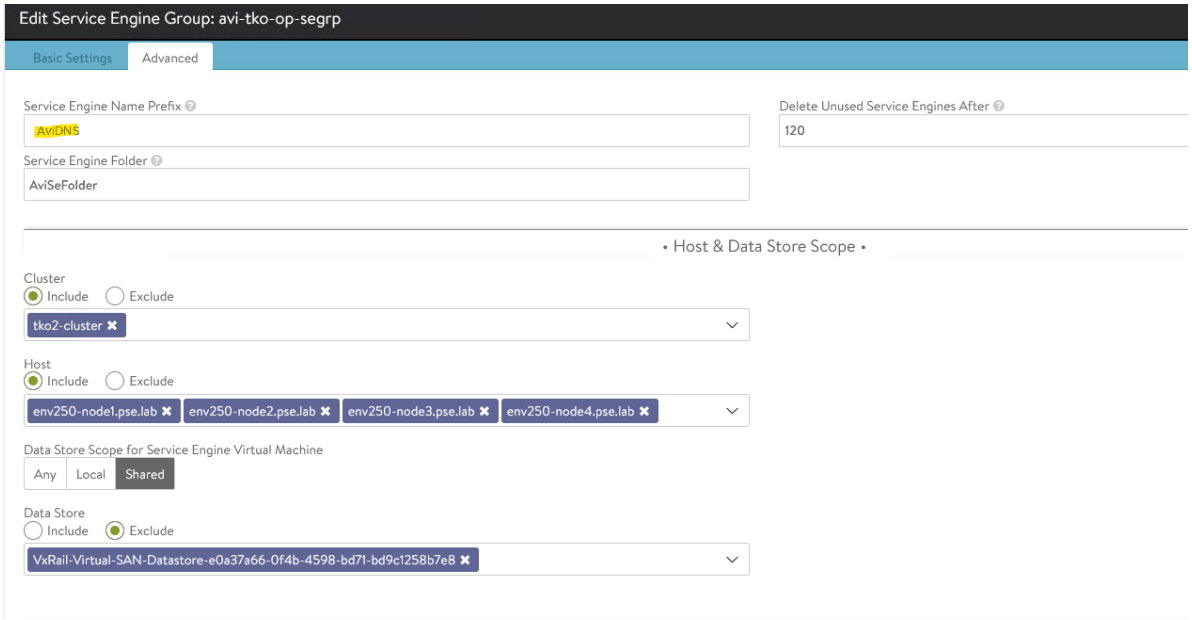

3. Create a pool of DNS servers

Figure 110: DNS pool

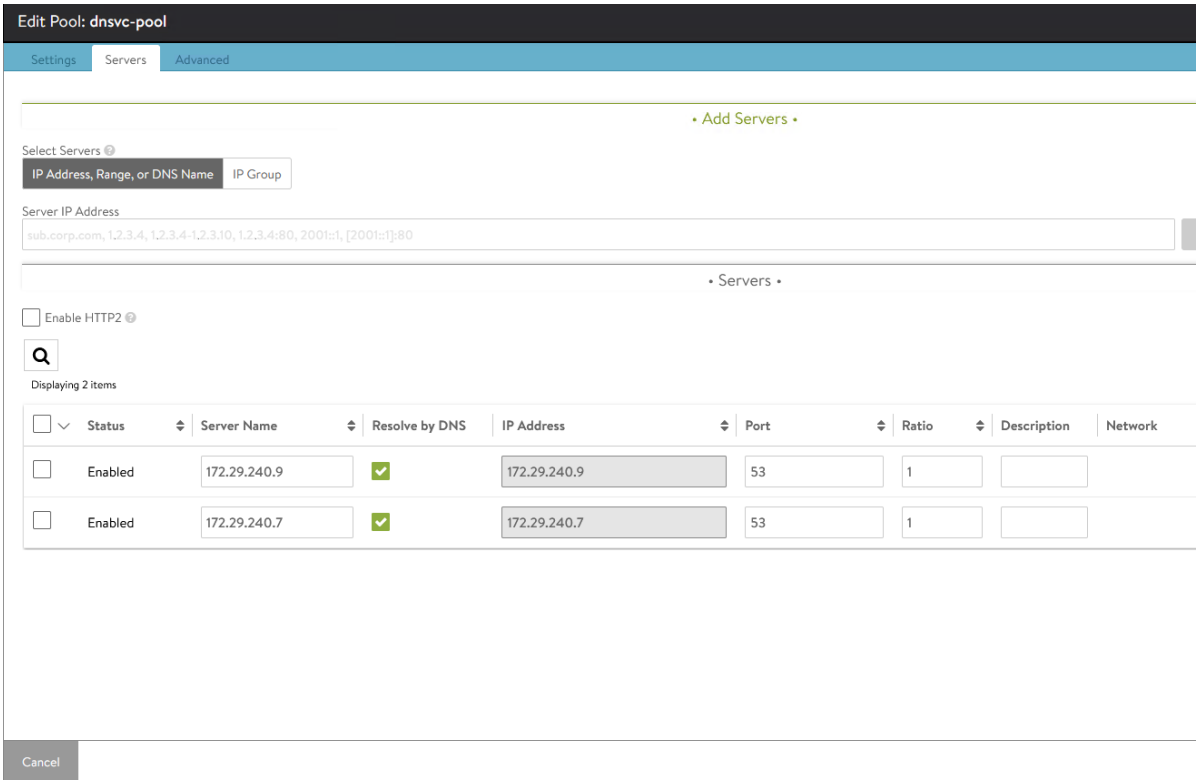

4. Create a virtual service for DNS

Figure 111: Create DNS service.

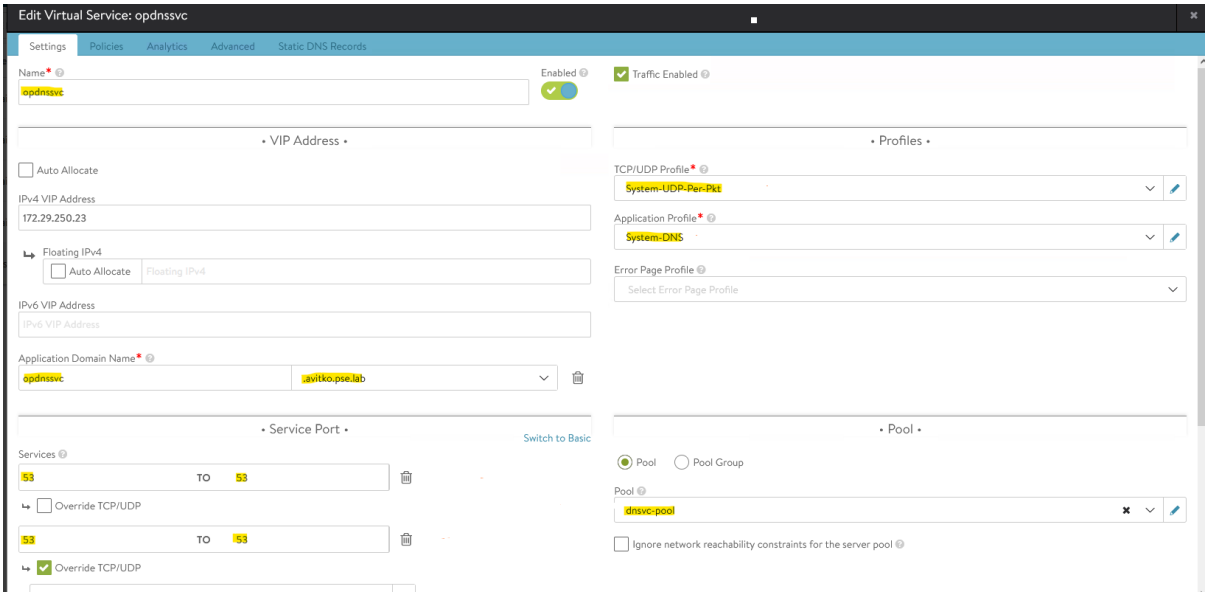

5. Add the pool created earlier to the virtual service

Figure 112: Create DNS service.

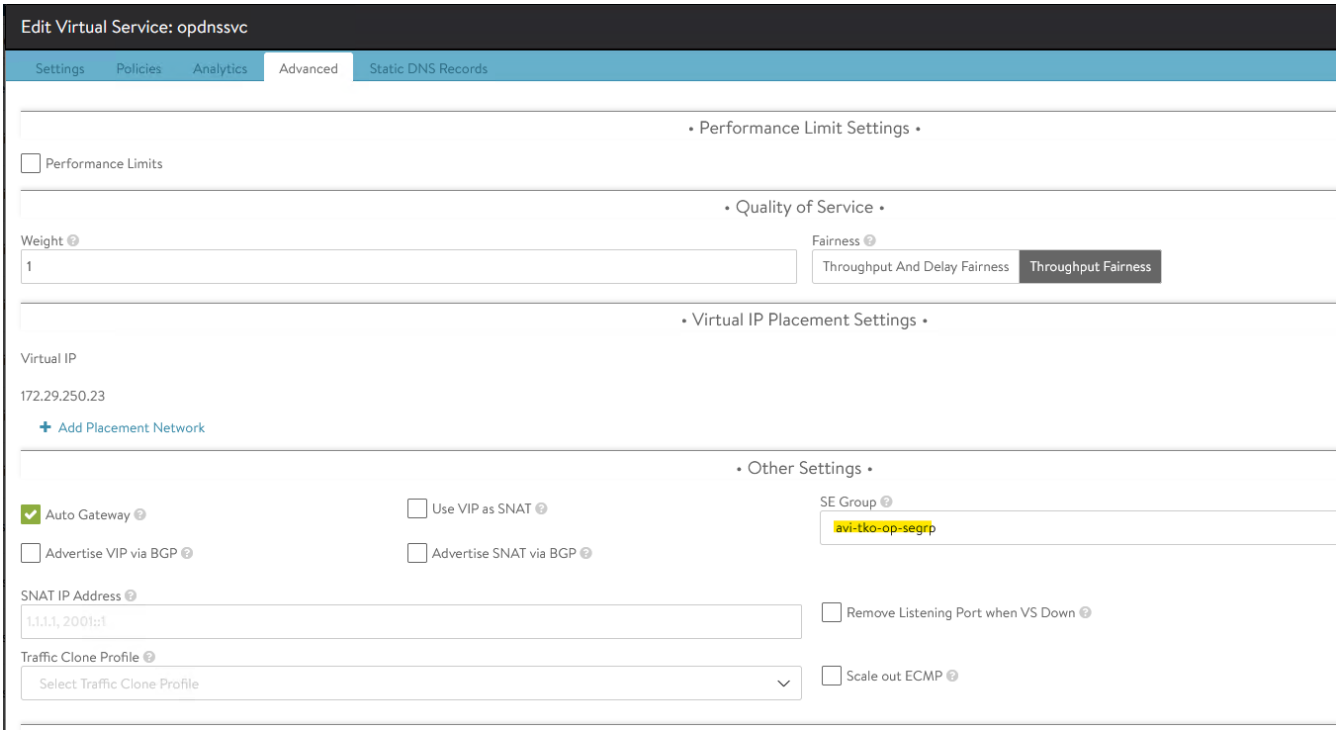

6. Ensure that virtual service is up and running.

Figure 113: Verify DNS service.

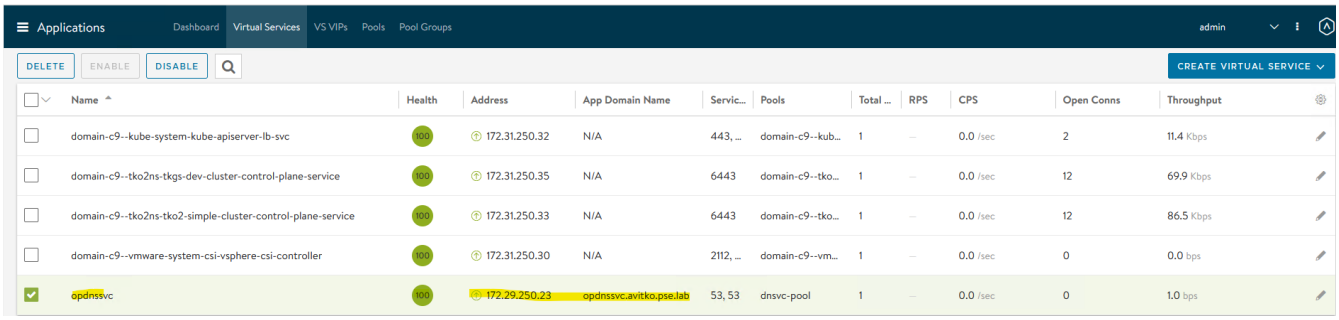

7. Follow the same steps on NSX Advanced Load Balancer on AWS and create a virtual service for DNS.

Figure 114: Create DNS service on AWS.

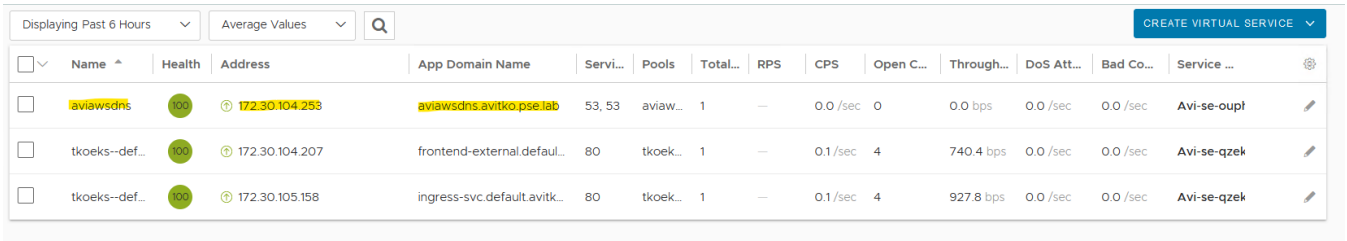

8. Note the IP addresses of the services. These will be used when domain delegation is created.

Configure GSLB Sites

1. On the on-premises controller, which will be the leader, configure GSLB via Infrastructure > GSLB. Click the pencil icon.

Figure 115: Enable GSLB

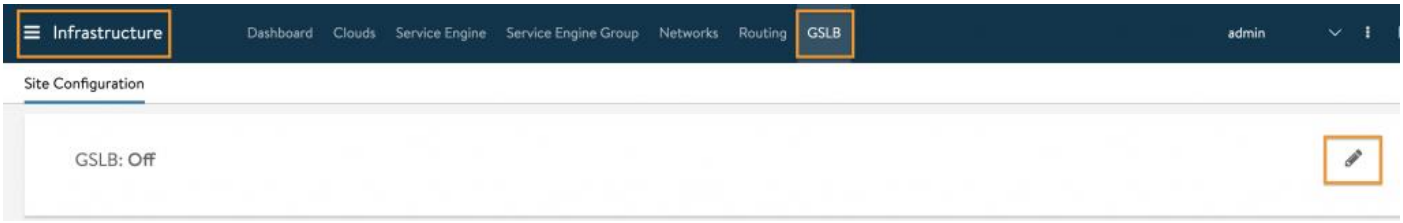

2. Give it a name. Enter the credentials for the controller and IP address. Enter the subdomain to be delegated to the load balancer.

Figure 116: Configure GSLB sites.

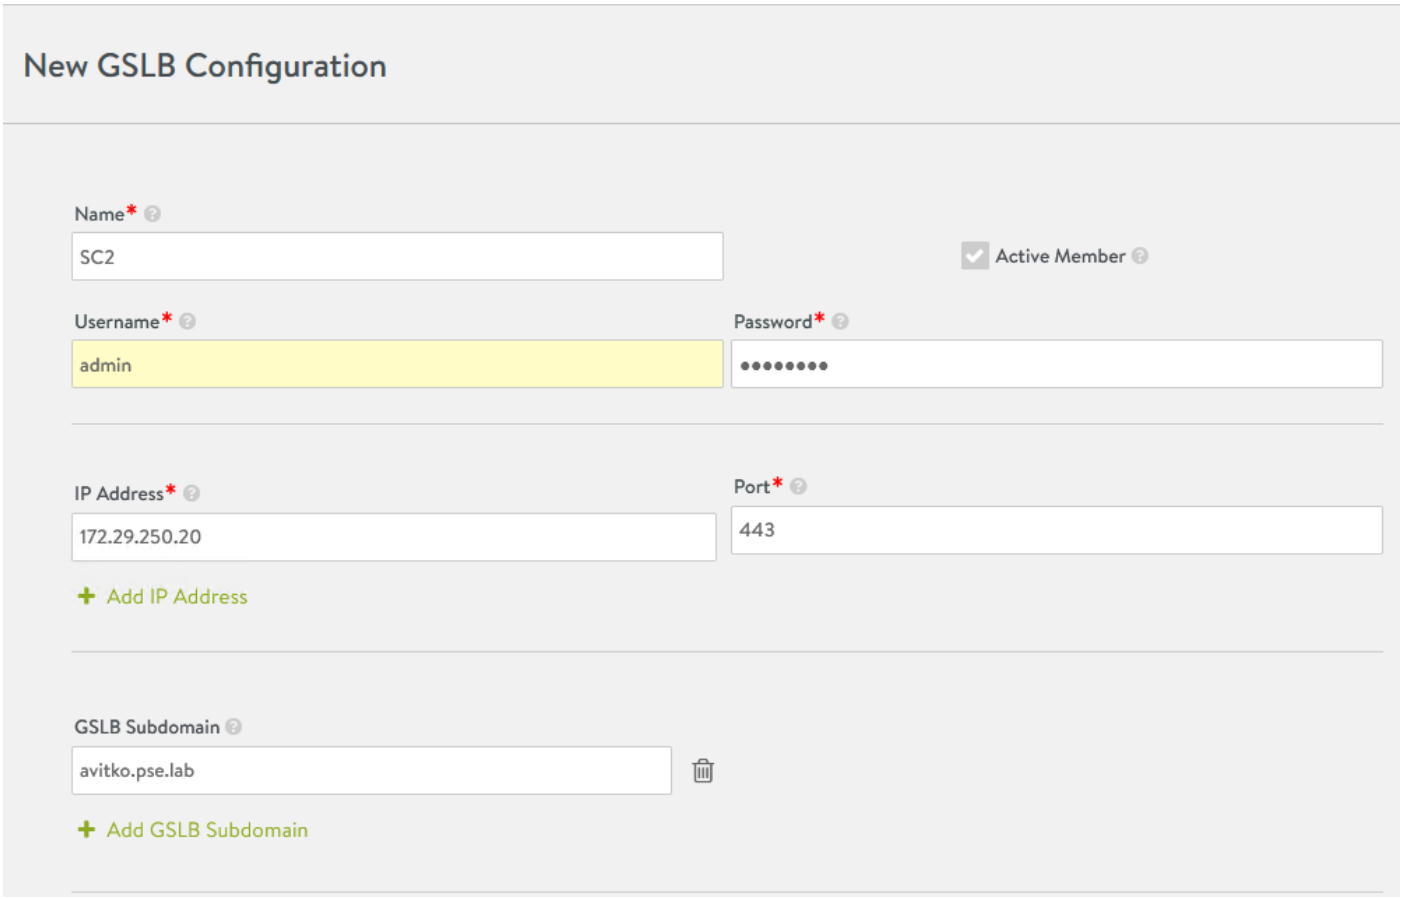

3. Under advanced settings configure the geo location parameters if load balancing based on Geo location is required. Click "Save and Set DNS Virtual Services."

Figure 117: Configure GSLB sites.

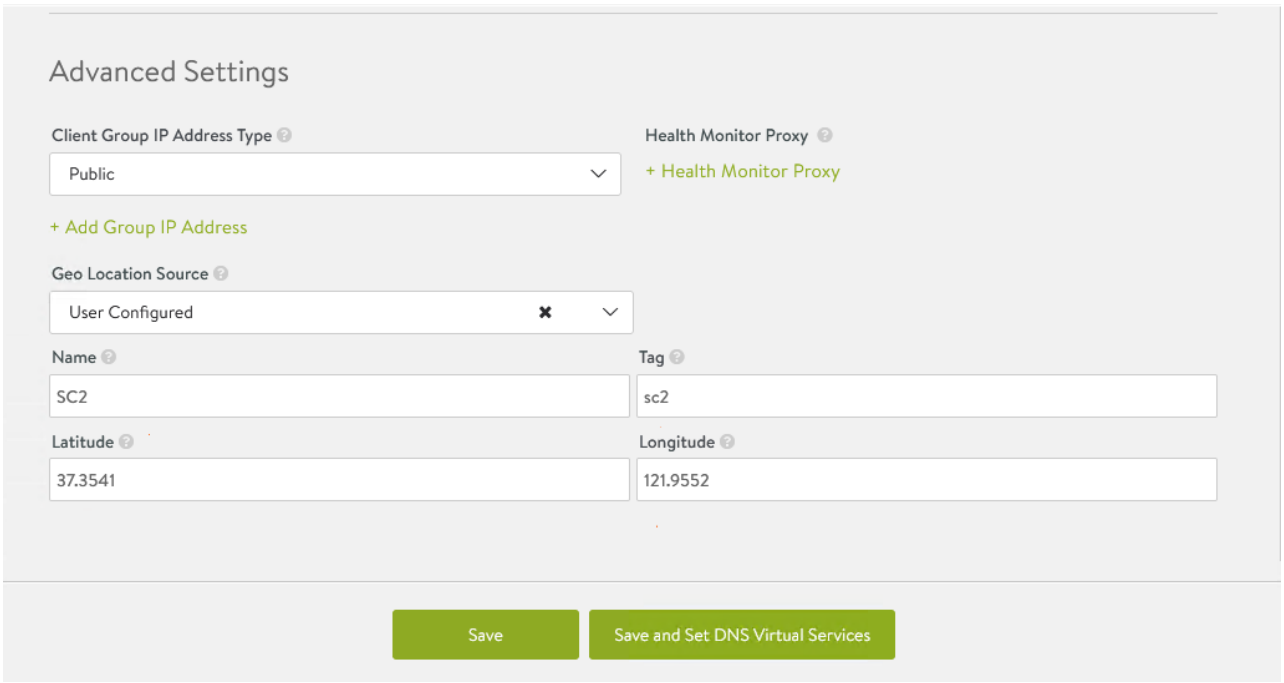

4. On the next screen, select the subdomain and virtual service created earlier for on-premises DNS. Click "Save."

Figure 118: Configure GSLB sites.

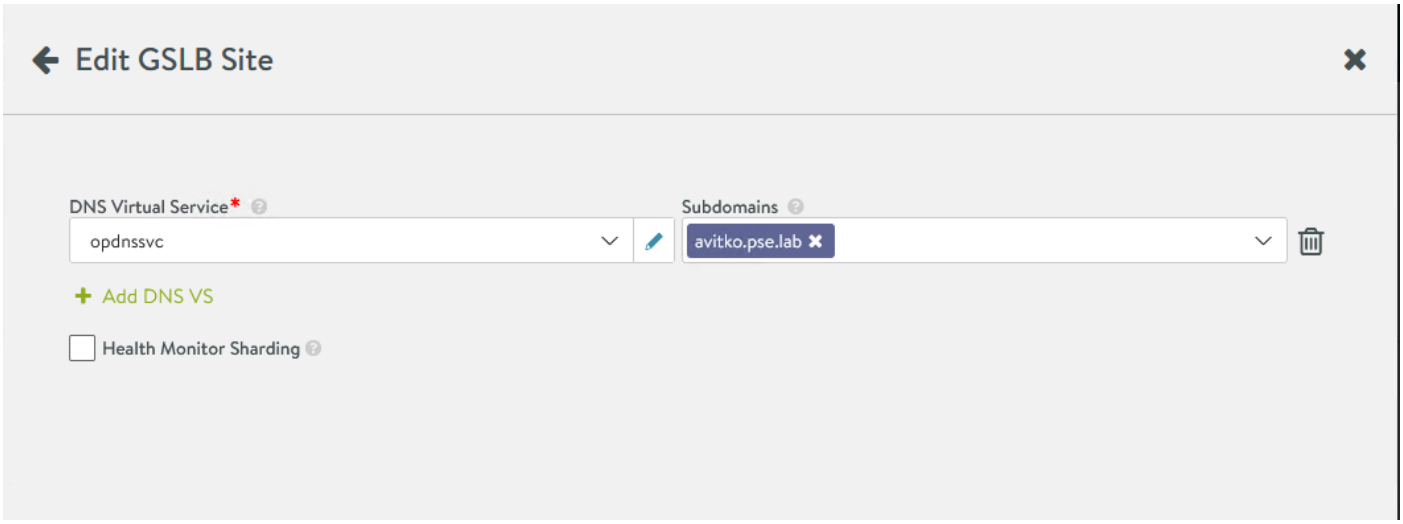

5. GSLB for on-premises is created. While on the same screen click “Add new site.”

Figure 119: Configure GSLB sites.

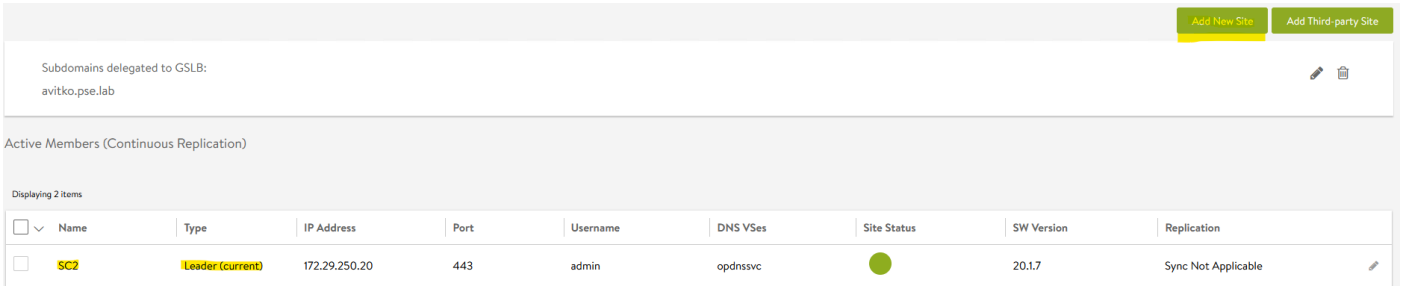

6. Enter the information for the NSX Advanced Load Balancer on AWS including geo location parameters if required. Ensure that “Active Member” is checked.

Figure 120: Configure GSLB sites.

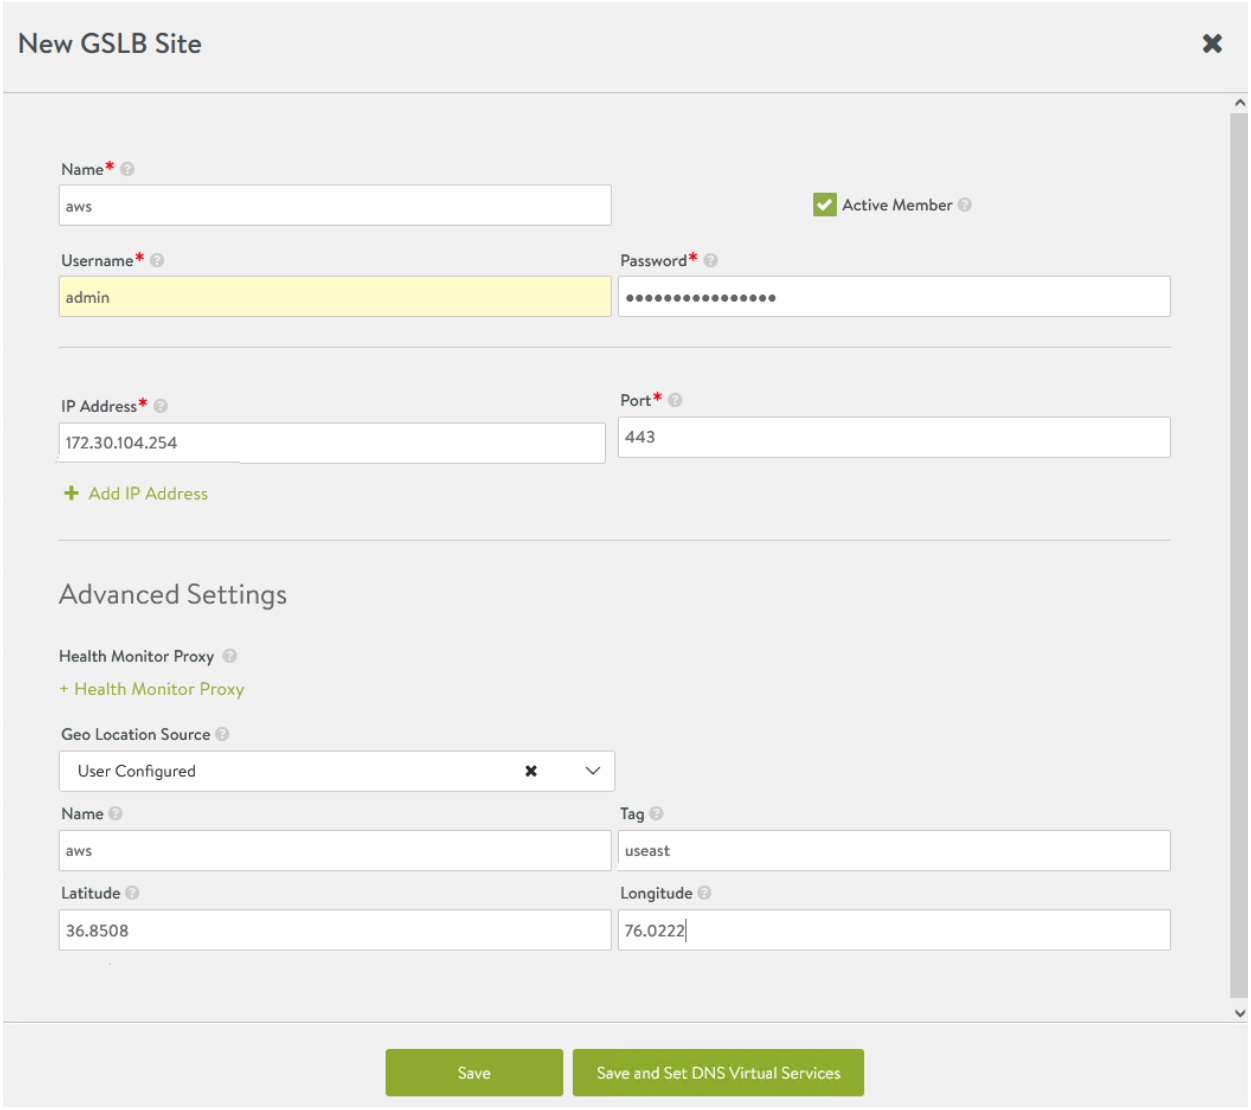

7. Click “Save and Set DNS Virtual Services” and select the subdomain and virtual created on controller on AWS. Click Save.

Figure 121: Configure GSLB sites.

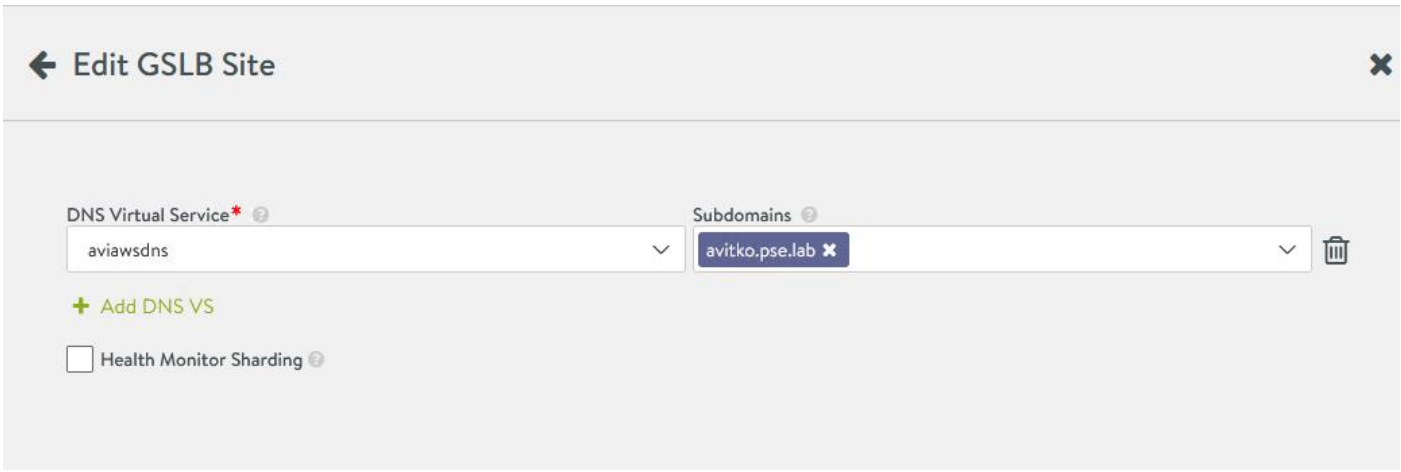

8. GSLB for both on-premises and AWS are created and should be running. AWS service will show “In Synch” when synch is successful.

Figure 122: Configure GSLB sites.

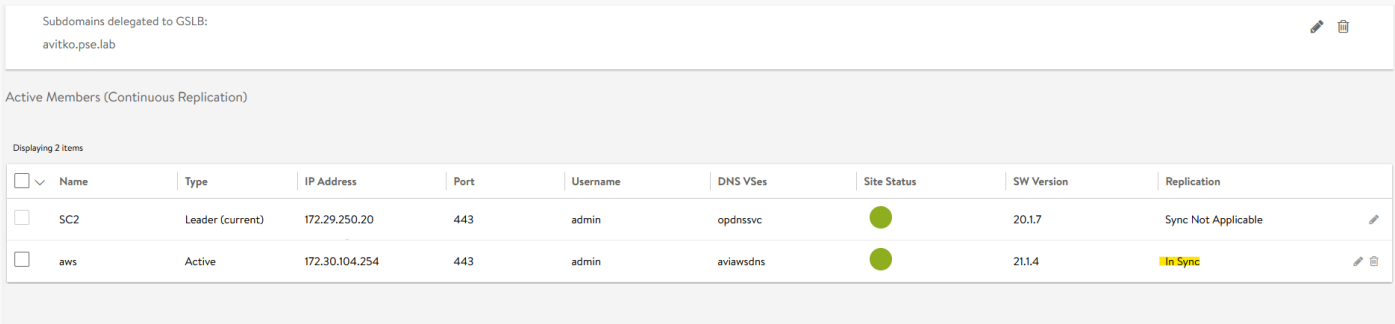

Configuring GSLB Services

Once the GSLB sites are configured and synchronized, GSLB service for desired applications or services can be created. Create a GSLB service to test and verify functionality as follows.

Note: When a Kubernetes application is installed on multiple sites, and the app selector label matches the label assigned during AMKO installation, AMKO will automatically create the GSLB service for the application. Following process creates the creation of GSLB service to illustrate the functionality.

1. On the on-premises controller navigate to Applications > GSLB Service. Give the service a name and select the sub-domain. Select a health monitoring option and a load balancing algorithm. Load balancing can be based on service priority or geo location. Click “Add Pool.”

Figure 123: Configure GSLB services.

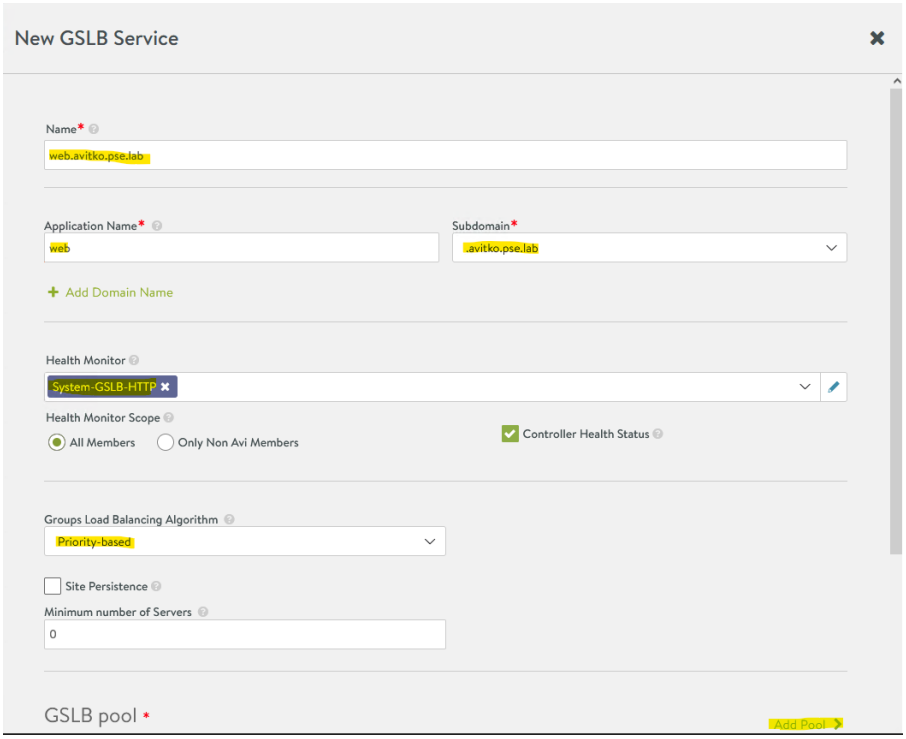

2. Give the pool a name and select load balancing algorithm. Under “Pool Member” site and virtual service created earlier. Click done.

Figure 124: Configure GSLB services.

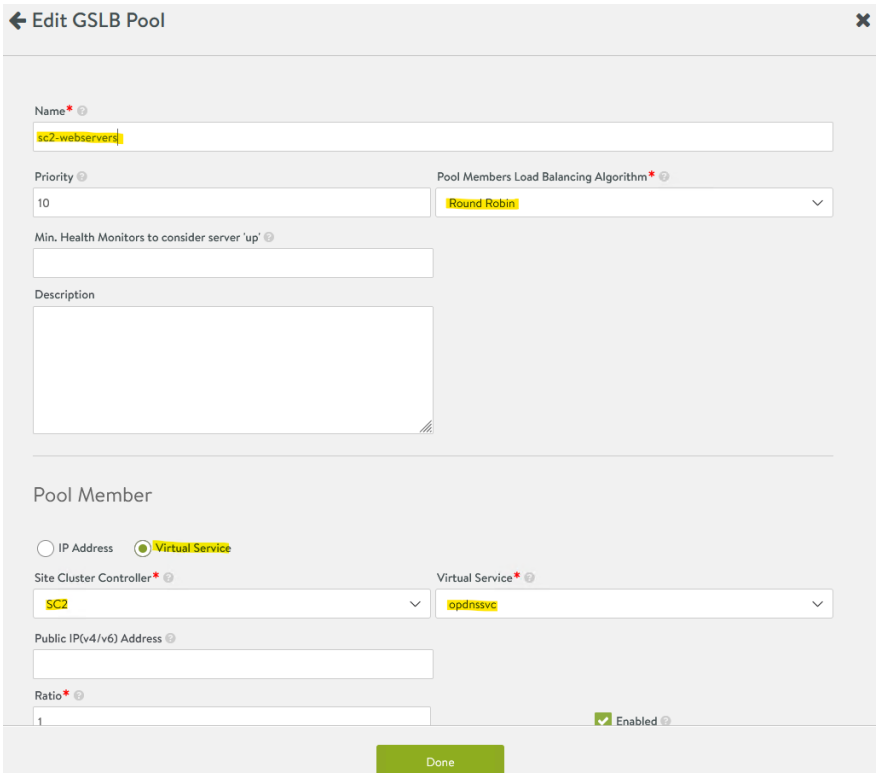

3. Create another pool for AWS in similar fashion.

Figure 125: Configure GSLB services.

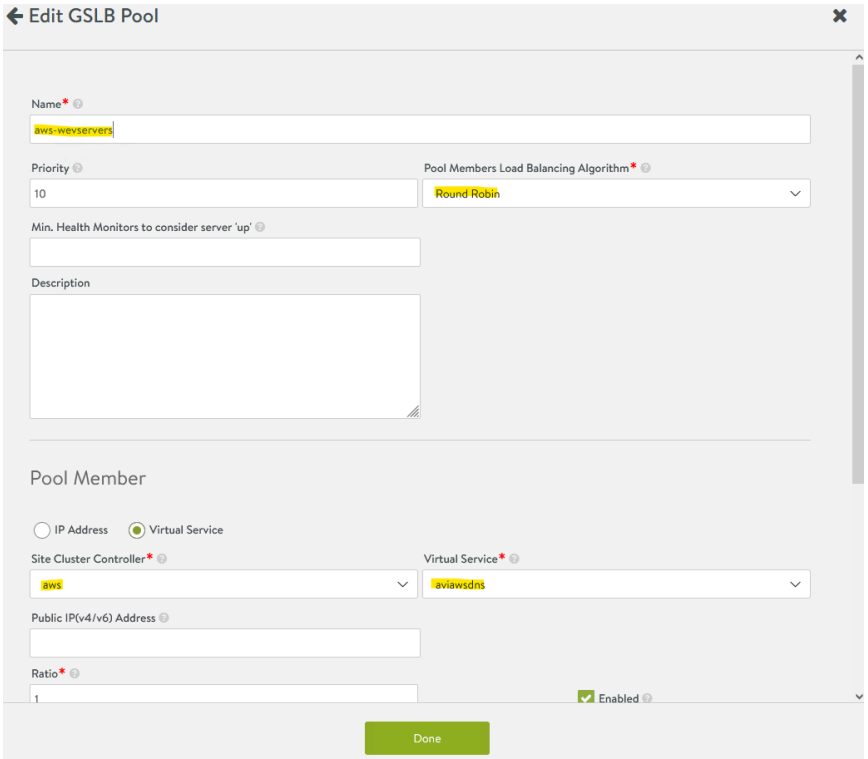

4. You should now have two pools showing in the virtual GSLB service. Click save.

Figure 126: Configure GSLB services.

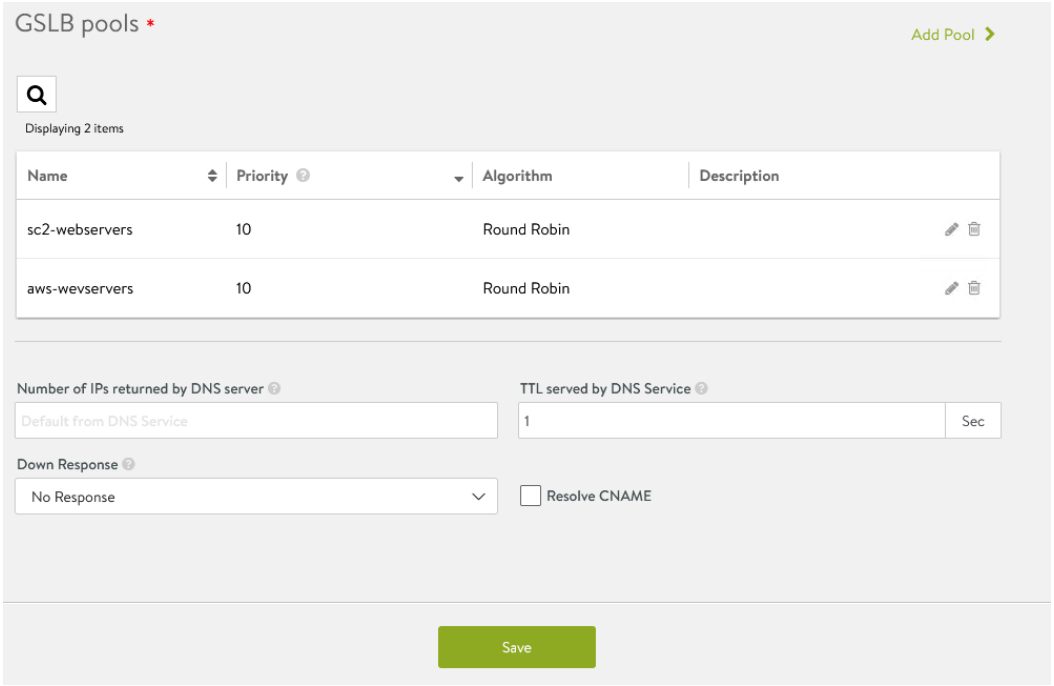

5. The GSLB service is created. This service is on the "avitko.pse.lab" domain which will be delegated to NSX Advanced Load Balancer. Hence next step is to create a delegation for this sub-domain.

DNS Domain Delegation

The Avi DNS virtual service is a generic DNS infrastructure that can implement the following functionality.

- DNS Load Balancing
- Hosting Manual or Static DNS Entries
- Virtual Service IP Address DNS Hosting
- Hosting GSLB Service DNS Entries

NSX Advanced Load Balancer DNS service can be deployed in a couple of ways. It can be deployed as an authoritative name server for a sub-domain delegated to it or as a primary DNS server for the domain. In the latter case, any requests that do not match DNS records in NSX Advanced Load Balancer are "proxied" to the corporate DNS server. For this reference architecture sub-domain delegation was used as described below. For more information on NSX Advanced Load Balancer DNS architecture please visit Avi DNS Architecture and Features page.

To create DNS domain delegation, the IP addresses of the DNS services created earlier for on-premises and AWS sites will be needed. This reference architecture uses Microsoft Active Directory Domain DNS.

1. In DNS manager, create an A Record of the two DNS services that were created.
2. On the Domain controller open the DNS manager and create a "New Delegation."

Figure 127: DNS domain delegation

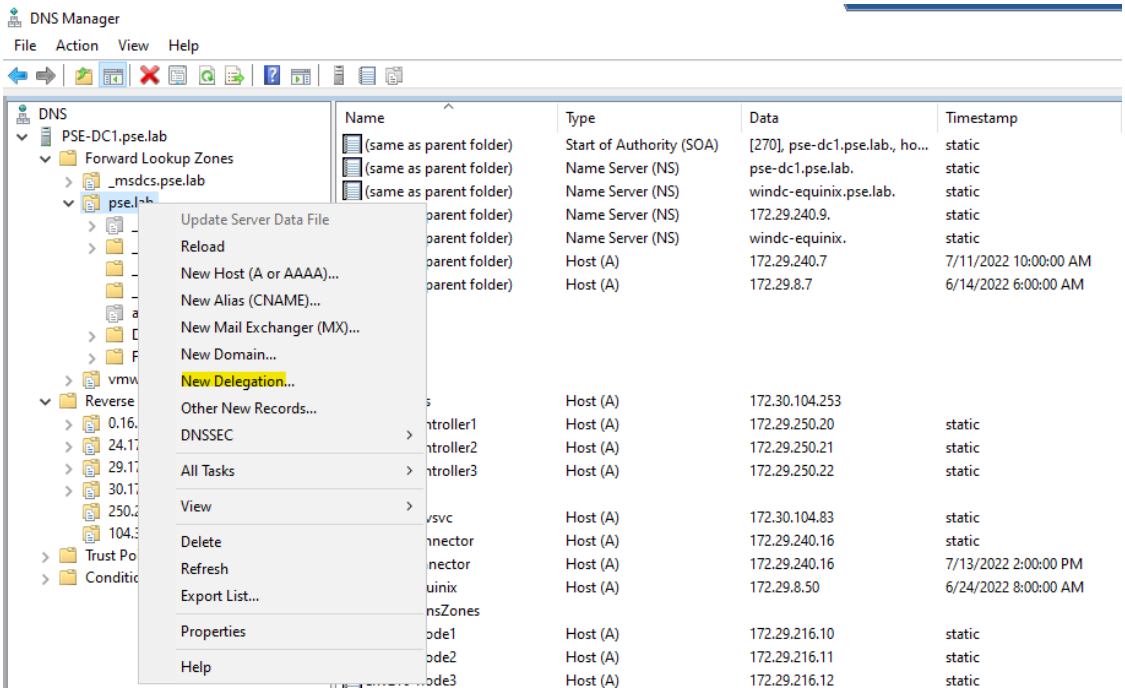

3. Enter the name of the sub-domain. In this case its “avitko.” Click Next

Figure 128: DNS domain delegation

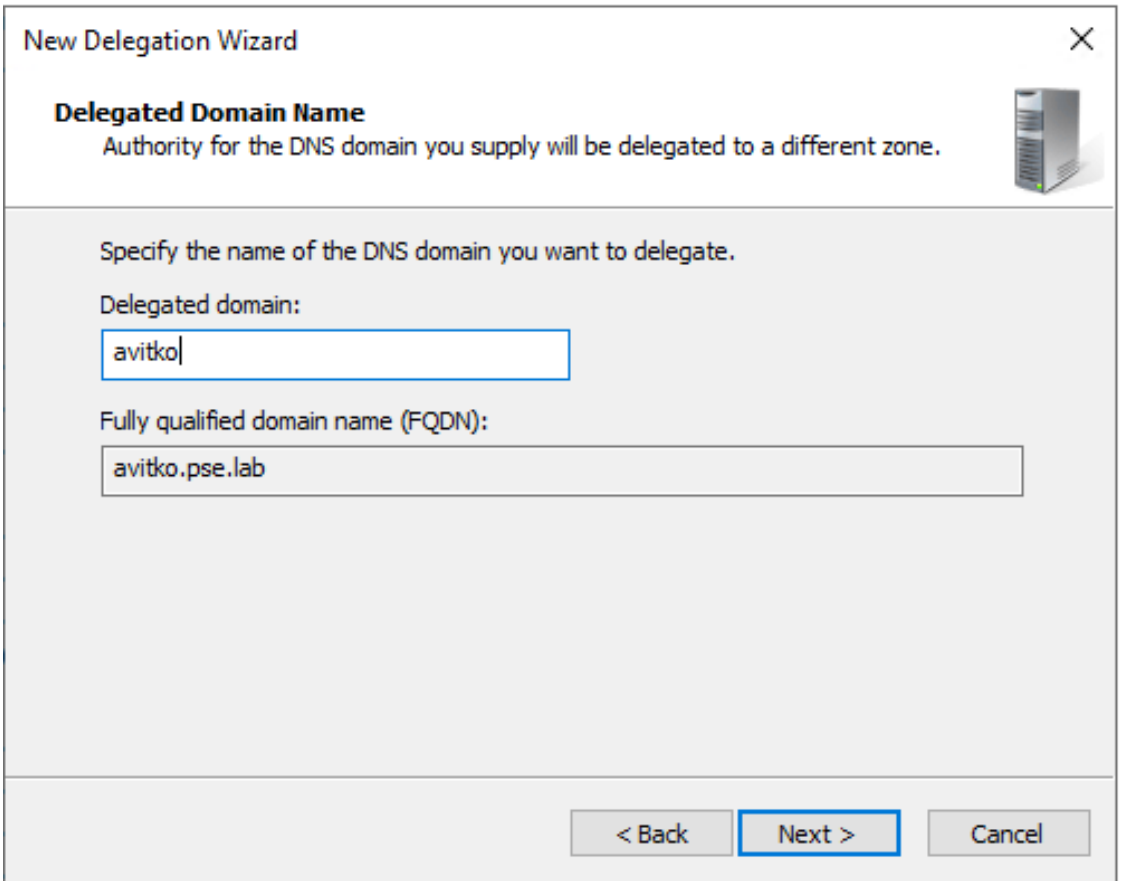

4. Add the two DNS services to the delegation and click Next and finish the process.

Figure 129: DNS domain delegation

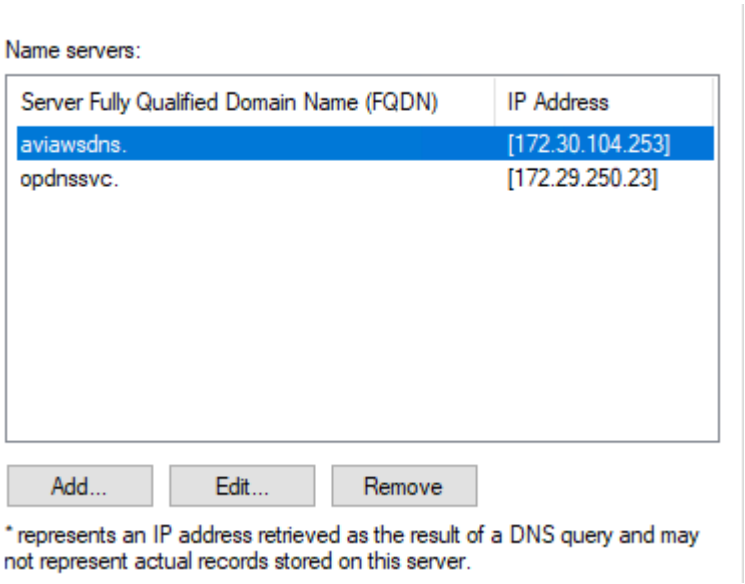

5. The delegation is complete.
6. Test the delegation by pinging the service URL. You should get a response from one of the two DNS services. This indicates that domain delegation is functioning properly. Since we selected round-robin algorithm, subsequent ping should be answered by the service on the other site. This indicates that load balancing configuration is functional. Geo location option can be set when desired as primary response type with round-robin as second option. In this case, users will be directed to the nearest application instance based on their location.

Figure 130: Test DNS domain delegation

```
aj@vmc-ubuntu-jump:~$ ping web.avitko.pse.lab
PING web.avitko.pse.lab (172.30.104.253) 56(84) bytes of data.
64 bytes from aviawsdns.pse.lab (172.30.104.253): icmp_seq=1 ttl=58 time=1.95 ms
64 bytes from aviawsdns.pse.lab (172.30.104.253): icmp_seq=2 ttl=58 time=1.73 ms
64 bytes from aviawsdns.pse.lab (172.30.104.253): icmp_seq=3 ttl=58 time=1.75 ms
64 bytes from aviawsdns.pse.lab (172.30.104.253): icmp_seq=4 ttl=58 time=1.92 ms
^C
--- web.avitko.pse.lab ping statistics ---
4 packets transmitted, 4 received, 0% packet loss, time 3004ms
rtt min/avg/max/mdev = 1.729/1.838/1.950/0.098 ms
aj@vmc-ubuntu-jump:~$ ping web.avitko.pse.lab
PING web.avitko.pse.lab (172.29.250.23) 56(84) bytes of data.
64 bytes from opdnssvc.pse.lab (172.29.250.23): icmp_seq=1 ttl=56 time=75.8 ms
64 bytes from opdnssvc.pse.lab (172.29.250.23): icmp_seq=2 ttl=56 time=75.1 ms
64 bytes from opdnssvc.pse.lab (172.29.250.23): icmp_seq=3 ttl=56 time=75.1 ms
64 bytes from opdnssvc.pse.lab (172.29.250.23): icmp_seq=4 ttl=56 time=75.1 ms
^C
```

Lifecyle Management via Tanzu Mission Control

As companies grow their cloud native environments to multiple cloud providers, platform consistency and manageability becomes a challenge. Each cloud provider has its own management portal and lifecycle management of such environment can become a nightmare. Enterprises need a solution to help platform operators efficiently expand control and provide Kubernetes environments with guardrails so DevOps teams can have consistency and developers can operate autonomously, in a self-service fashion. For user authentication, an identity source such as Microsoft Active Directory or another a 3rd party identity source needs to be federated with Tanzu Mission Control. Please see "Self-Service Federation Setup" in Tanzu Mission Control Documentation.

VMware Tanzu Mission Control is a centralized management hub with cluster lifecycle management and a unified policy engine that simplifies multi-cloud and multi-cluster Kubernetes management across teams in the enterprise.

Administrator can perform several tasks to manage their on-premises or multi-cloud environments. Some of the tasks that an administrator needs to perform to administer their environment is listed and explained below.

1. Create a cluster group.
2. Add management cluster to Tanzu Mission Control
3. Create Kubernetes workload clusters.
4. Attach existing Kubernetes clusters.
5. Install Tanzu toolkit packages and applications via helm charts.
6. Configure policies and policies templates.
7. DevOps access and automation via Tanzu Mission Control CLI
8. Enable continuous delivery (CD) via Git repository integration.
9. Create, attach, or delete Amazon EKS clusters.
10. Run conformance and security inspections on clusters.
11. Enable data protection for clusters.

Note: It is not the intent of this document to cover all aspects and features of Tanzu Mission Control. For more details please see Tanzu Mission Control documentation.

Create a cluster group

Creating a cluster group for different deployments or site is an optional step. The advantage is that it organizes different cluster types and policies can be applied to all cluster at the group level. Create cluster groups from the left menu pane in Tanzu Mission Control portal.

Add management cluster to Tanzu Mission Control

For Tanzu Mission Control to manage the Tanzu Kubernetes Grid environment, the management cluster need to be registered to it. The following steps depict the management cluster registration process.

1. In the Tanzu Mission Control portal, navigate to Administration > Management clusters and click on Register Management Cluster and select type of management cluster you are registering.

Figure 131: Register management cluster

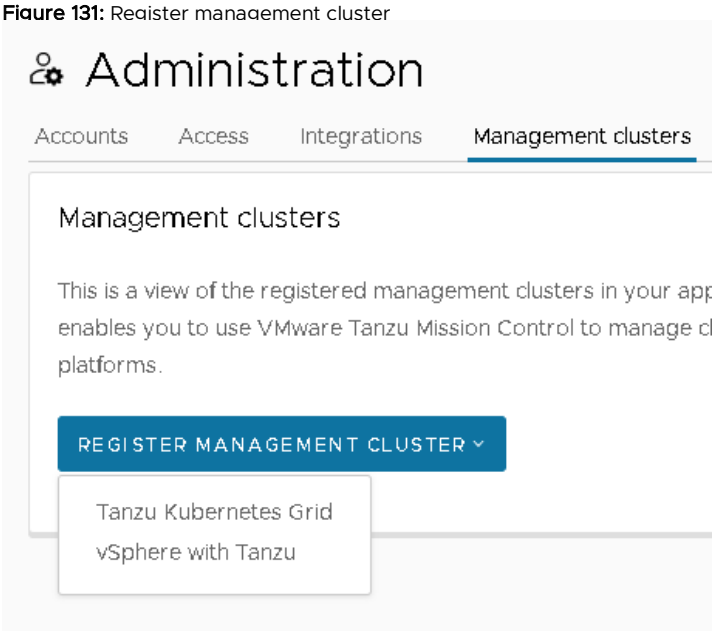

2. Enter name, cluster group, description, and label information if desired. Labels help organize various Tanzu Mission Control objects and that can be sorted and displayed easily.

Figure 132: Register management cluster

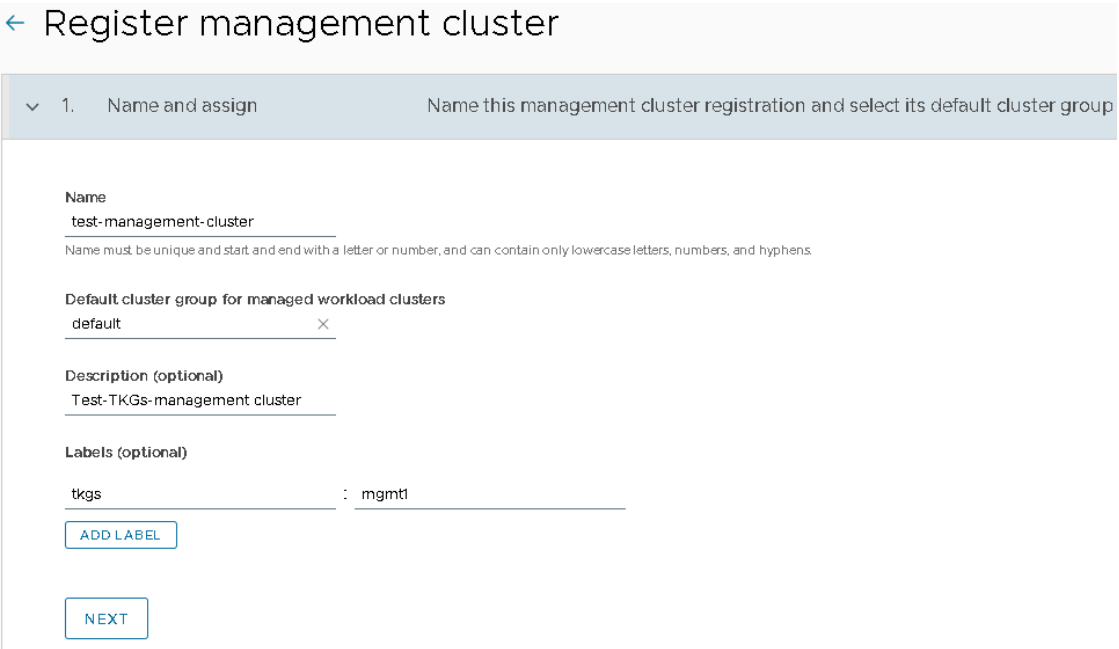

3. Enter proxy information if your management cluster is behind a proxy.

Figure 133: Enter proxy

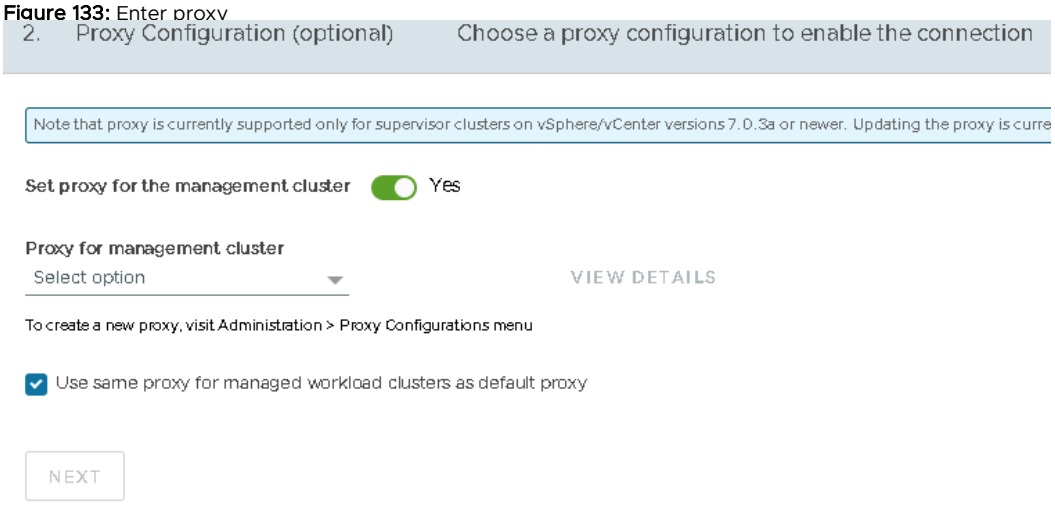

4. Copy and provide the registration URL that has the registration key to the vSphere administrator. The vSphere administrator will perform the next step in registering the management cluster to Tanzu Mission Control.

Figure 134: Copy registration URL

5. As a vSphere administrator, login to the Supervisor cluster and list namespaces. Take note of the Tanzu Mission Control service namespace.

Figure 135: Tanzu Mission Control namespace

```
aj@aj-vspherecli:~/$ kubectl get ns
NAME                                   STATUS   AGE
bbns                                   Active   56m
default                                Active   7d20h
kube-node-lease                        Active   7d20h
kube-public                            Active   7d20h
kube-system                            Active   7d20h
svc-tmc-c9                             Active   7d20h
vmware-system-ako                      Active   7d20h
```

6. Create and apply .yaml file using the registration URL and svc-tmc-xx namespace as shown below.
   Sample yaml:
   apiVersion: installers.tmc.cloud.vmware.com/v1alpha1
   kind: AgentInstall
   metadata:
   name: tmc-agent-installer-config
   namespace: svc-tmc-c9
   spec:
   operation: INSTALL
   registrationLink: https://org.tmc.cloud.vmware.com/installer?id= 17e139c2ba3551axxxxxxxxx
7. Apply the yaml via kubectl create -f <filename.yaml> to complete the registration process.
8. In Tanzu Mission Control console, verify that connection to the Supervisor cluster is successful and cluster is added and functional.

Figure 136: Verify connection

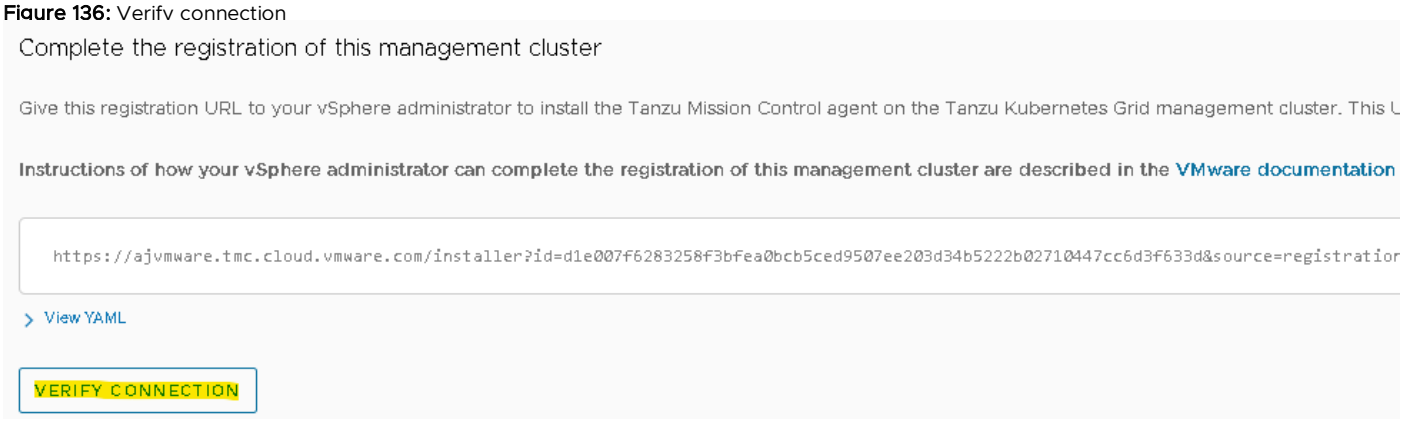

Create Kubernetes workload clusters.

1. Navigate to Clusters and click create cluster.

Figure 137: Create cluster

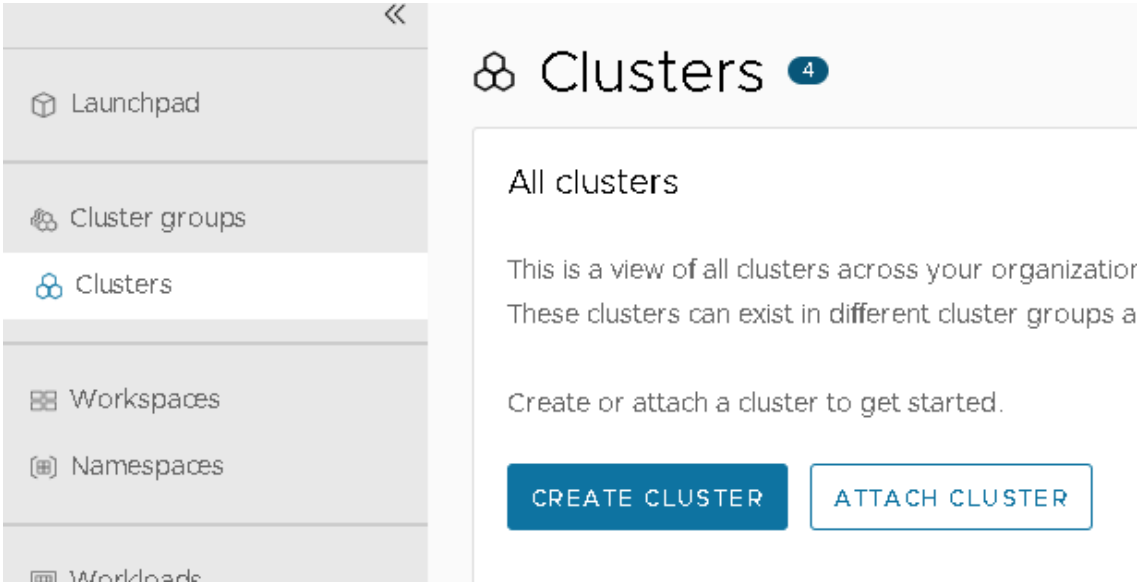

2. Select the management cluster and click continue to create cluster.

Figure 138: Select management cluster

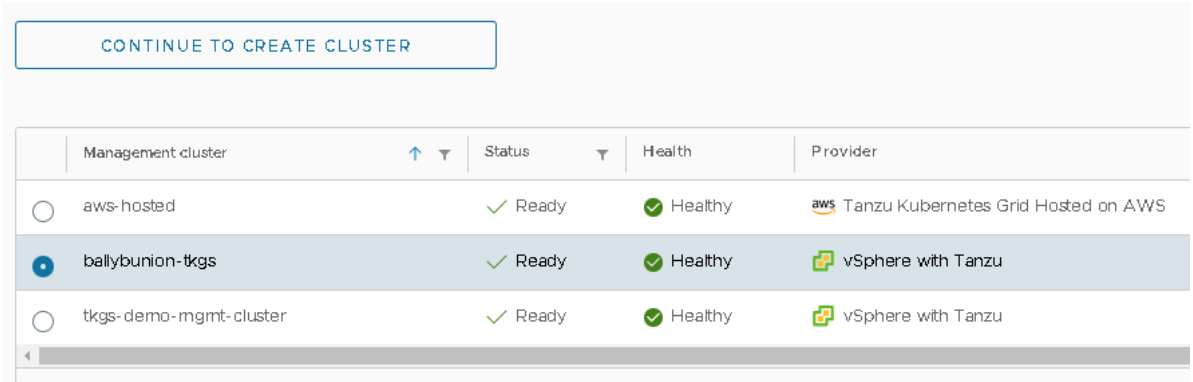

3. Select provisioner which in this case is the namespace you created in workload management.

Figure 139: Select provisioner

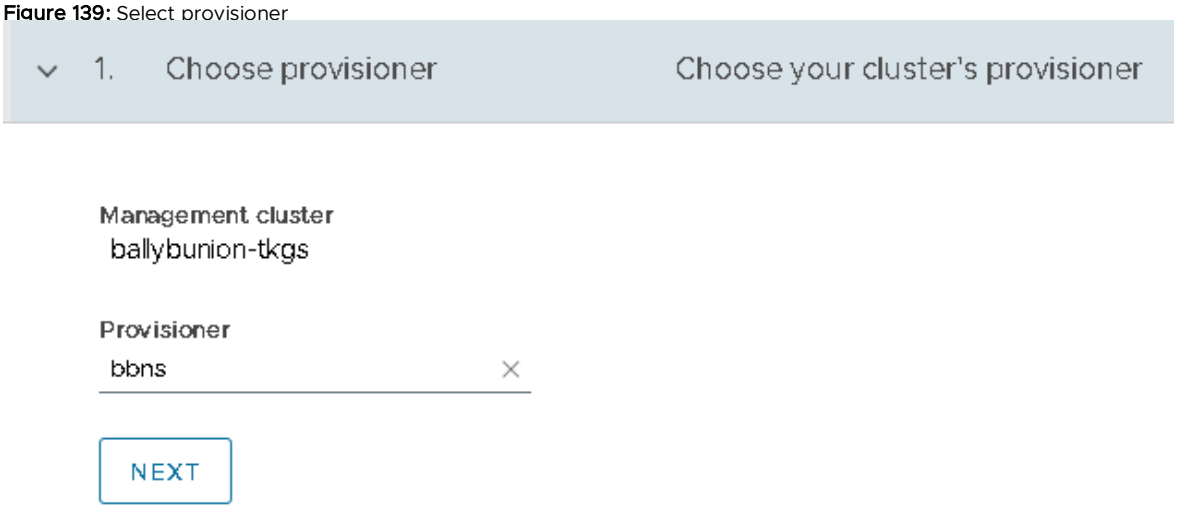

4. On the next screen give cluster a name and select a group.

Figure 140: Cluster name and group

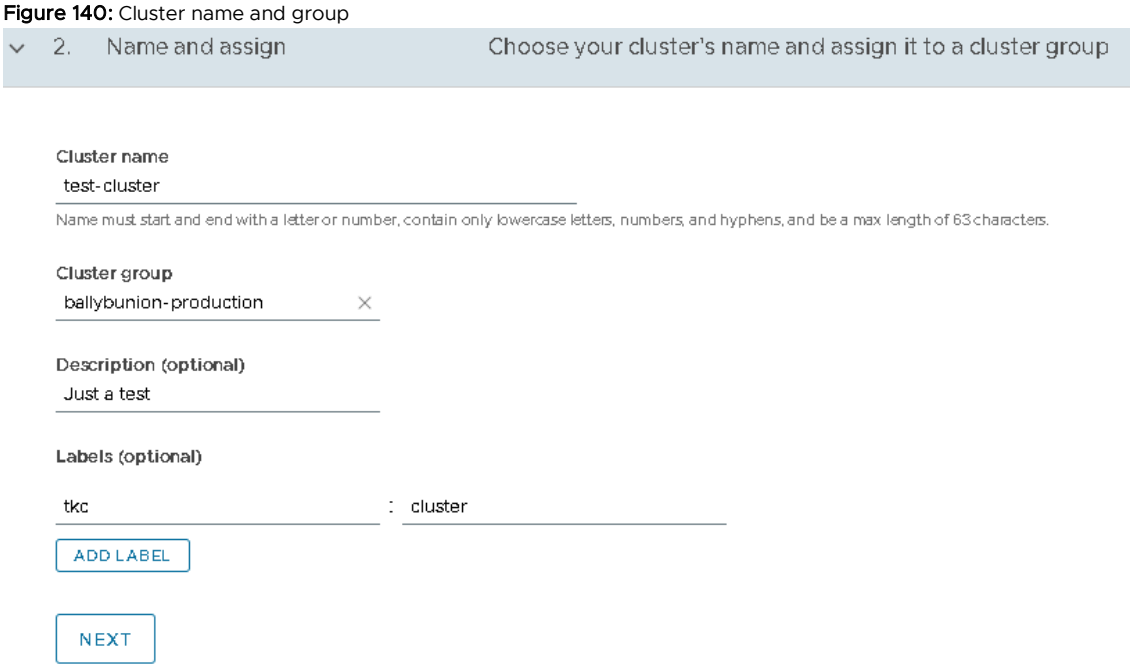

5. Select a Kubernetes version and assign network CIDR and storage class.

Figure 141: Configure parameters.

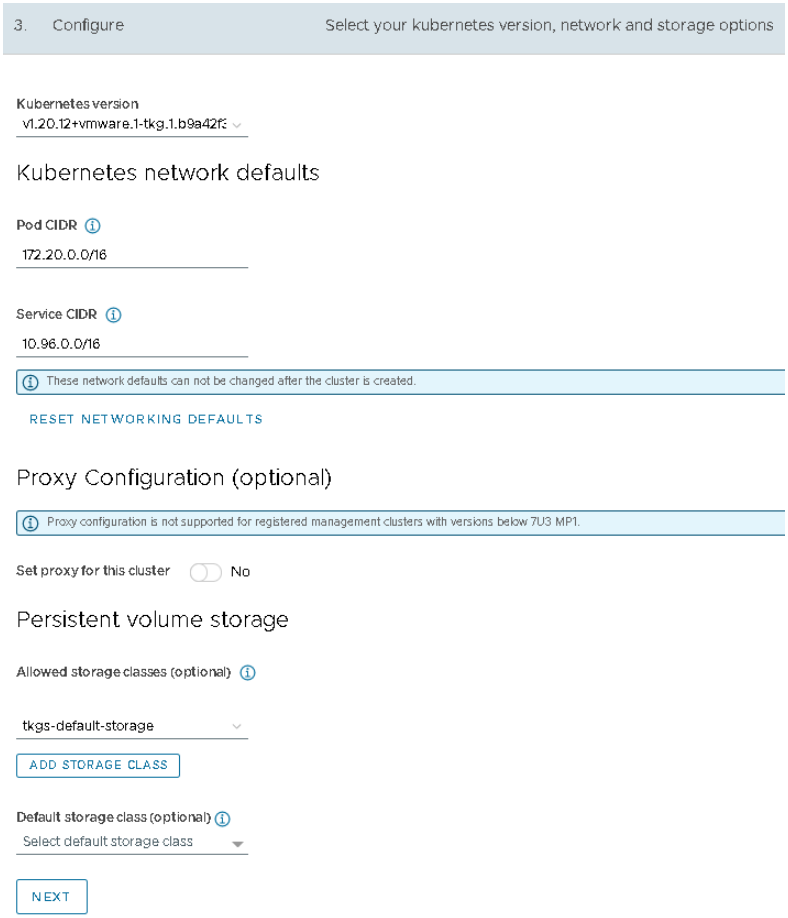

6. On the next screen select a deployment model for your control plane nodes, select a VM class and storage policy. You can also create a volume at this point.

Figure 142: Configure control plane specifications.

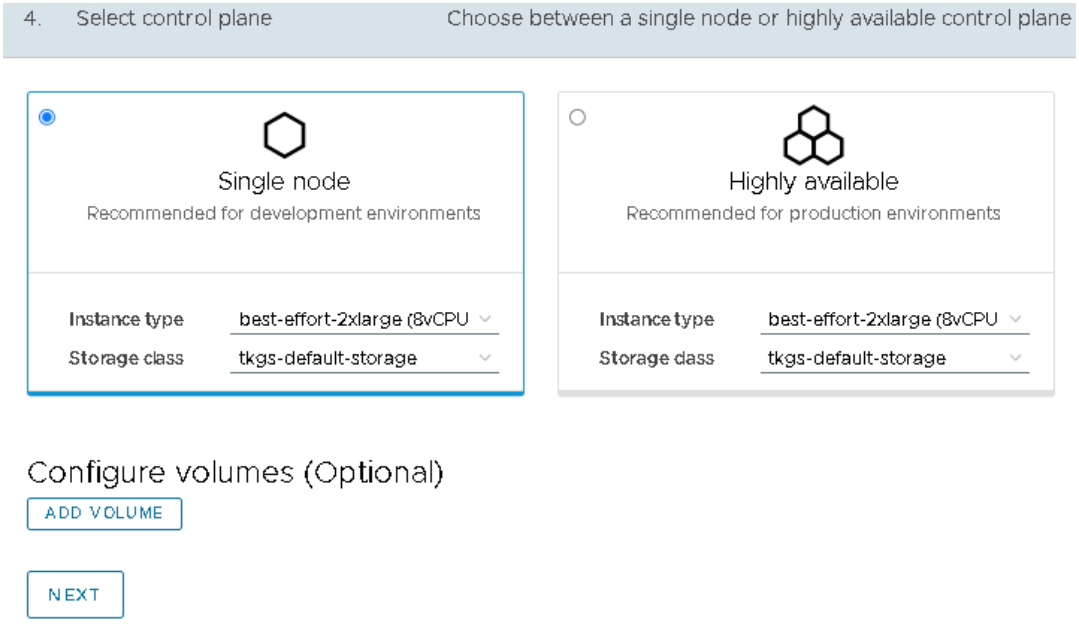

7. Modify the default pool configuration which has one node, to the desired number of worker nodes. Set the VM class and storage policy. Click Create Cluster to start the cluster creation process.

Figure 143: Modify default node pool

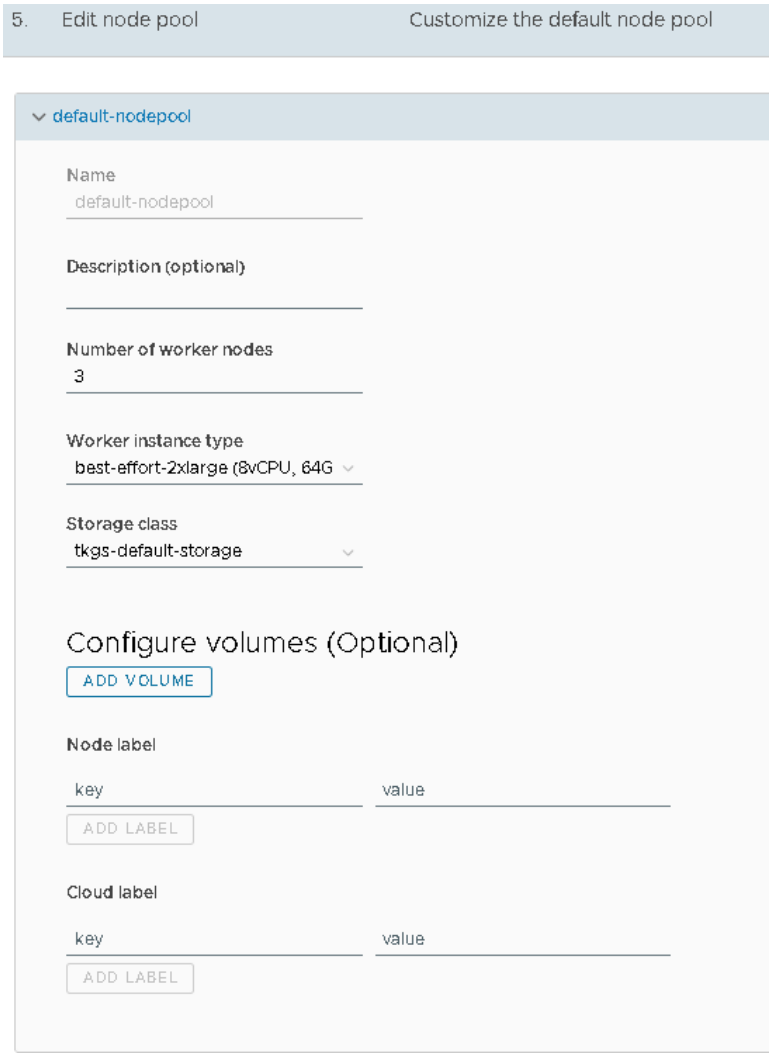

Attach existing Kubernetes clusters.

1. In Tanzu Mission Control navigate to Clusters > Attach Cluster and enter the desired information.

Figure 144: Attached Cluster

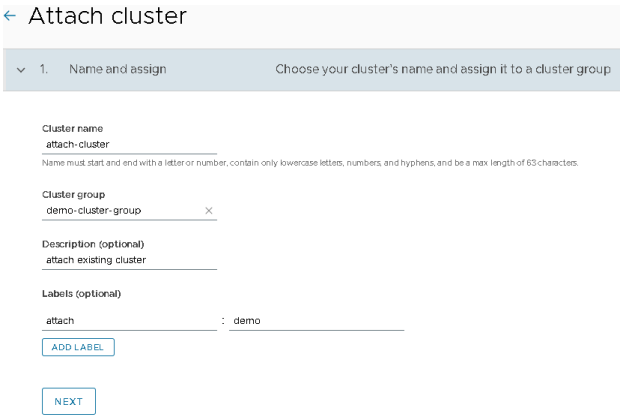

2. On the next screen enter proxy information if you cluster is behind a proxy.
3. On the “Install Agent” step copy the kubectl command.

Figure 145: Install agent.

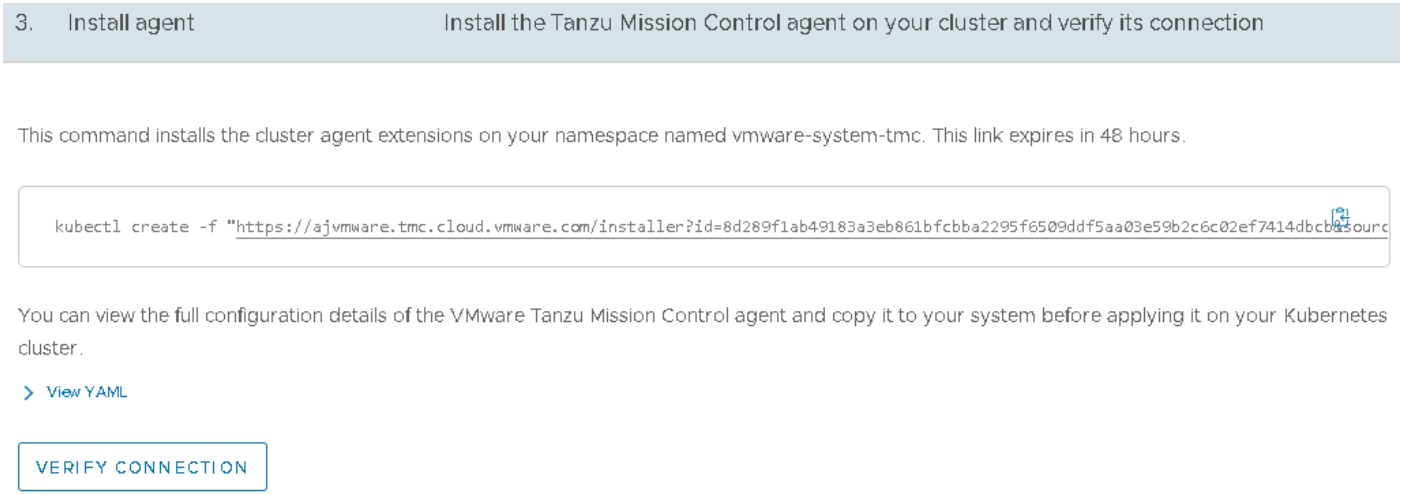

4. Login to the cluster and run the command. Cluster should be added, and policies created.

Figure 146: Run cli command

```
aj@demo-vspherecli:~/cli$ kubectl create -f "https://ajvmware.tmc.cloud.vmware.com/installer?id=8
d289f1ab49183a3eb861bfcbba2295f6509ddf5aa03e59b2c6c02ef7414dbcb&source=attach"
namespace/vmware-system-tmc created
configmap/stack-config created
secret/tmc-access-secret created
customresourcedefinition.apiextensions.k8s.io/agents.clusters.tmc.cloud.vmware.com created
customresourcedefinition.apiextensions.k8s.io/extensionconfigs.intents.tmc.cloud.vmware.com creat
ed
customresourcedefinition.apiextensions.k8s.io/extensionintegrations.clusters.tmc.cloud.vmware.com
 created
customresourcedefinition.apiextensions.k8s.io/extensionresourceowners.clusters.tmc.cloud.vmware.c
om created
customresourcedefinition.apiextensions.k8s.io/extensions.clusters.tmc.cloud.vmware.com created
serviceaccount/extension-manager created
clusterrole.rbac.authorization.k8s.io/extension-manager-role created
clusterrolebinding.rbac.authorization.k8s.io/extension-manager-rolebinding created
service/extension-manager-service created
deployment.apps/extension-manager created
serviceaccount/extension-updater-serviceaccount created
podsecuritypolicy.policy/vmware-system-tmc-agent-restricted created
clusterrole.rbac.authorization.k8s.io/extension-updater-clusterrole created
clusterrole.rbac.authorization.k8s.io/vmware-system-tmc-psp-agent-restricted created
clusterrolebinding.rbac.authorization.k8s.io/extension-updater-clusterrolebinding created
clusterrolebinding.rbac.authorization.k8s.io/vmware-system-tmc-psp-agent-restricted created
service/extension-updater created
deployment.apps/extension-updater created
serviceaccount/agent-updater created
clusterrole.rbac.authorization.k8s.io/agent-updater-role created
clusterrolebinding.rbac.authorization.k8s.io/agent-updater-rolebinding created
deployment.apps/agent-updater created
cronjob.batch/agentupdater-workload created
```

5. In the Tanzu Mission Control console verify that cluster has been attached

Figure 147: Verify

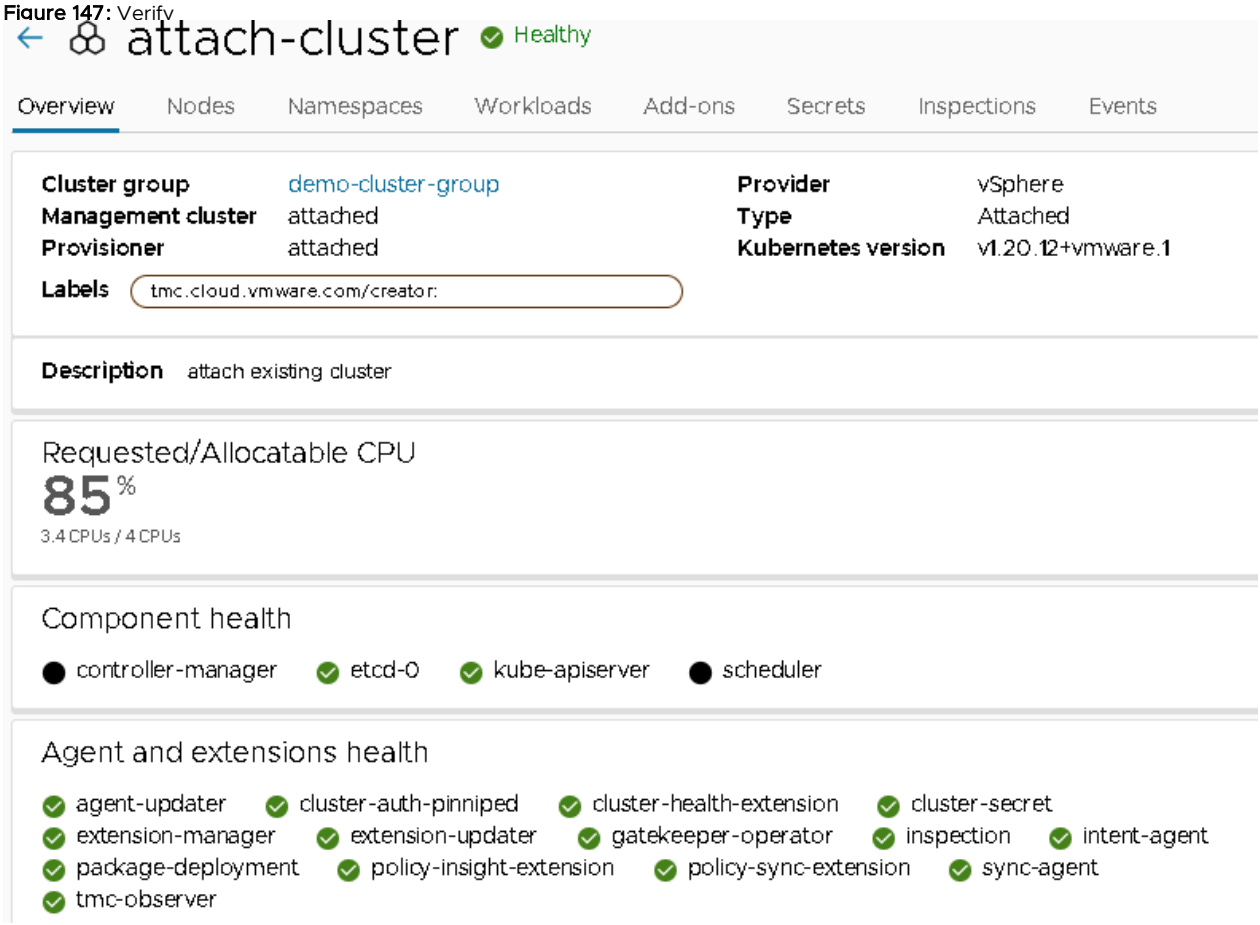

Tanzu toolkit packages and helm charts

Tanzu Mission Control operators can install, delete, and manage packages on Kubernetes clusters. Tanzu Mission Control uses Carvel for package management. The “Catalog” page shows the packages available to be installed on Kubernetes clusters.

Figure 148: Packages

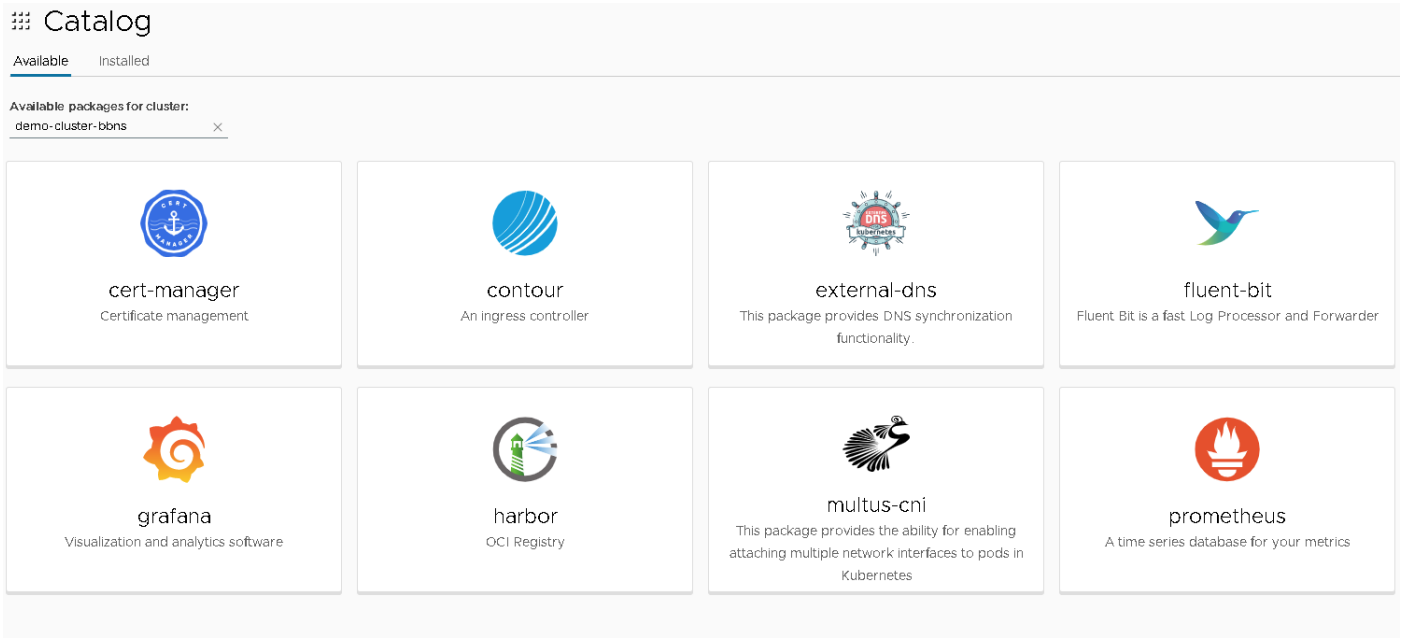

Package repositories available for each cluster can be viewed, enabled, or disabled via Cluster > Add-on tab. Custom package repositories can be added via the “Add Package Repository” button.

Figure 149: Repositories

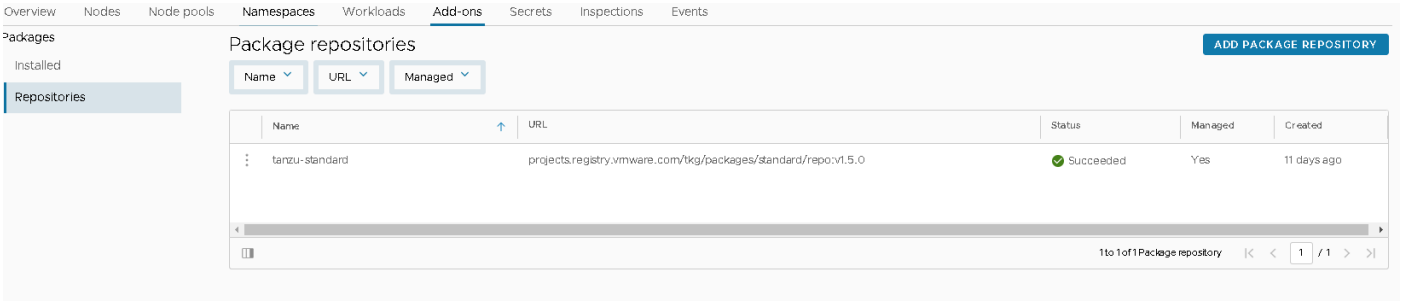

Install Packages

Figure 63 shows the packages available with Tanzu standard repository. Method of deployment is the same for all packages. Some packages however have more customizable fields in Tanzu Mission Control during installation. Below is an example of how to install Prometheus and Grafana using Tanzu Mission Control.

Install Prometheus and Grafana

1. Navigate to Catalog select a cluster and click on Prometheus and select install package.
2. Give the package a name and select a version to be installed from the drop-down list. Under package configuration, fields that have a pencil icon can be modified and configured per your configuration requirement.

Figure 150: Install Prometheus

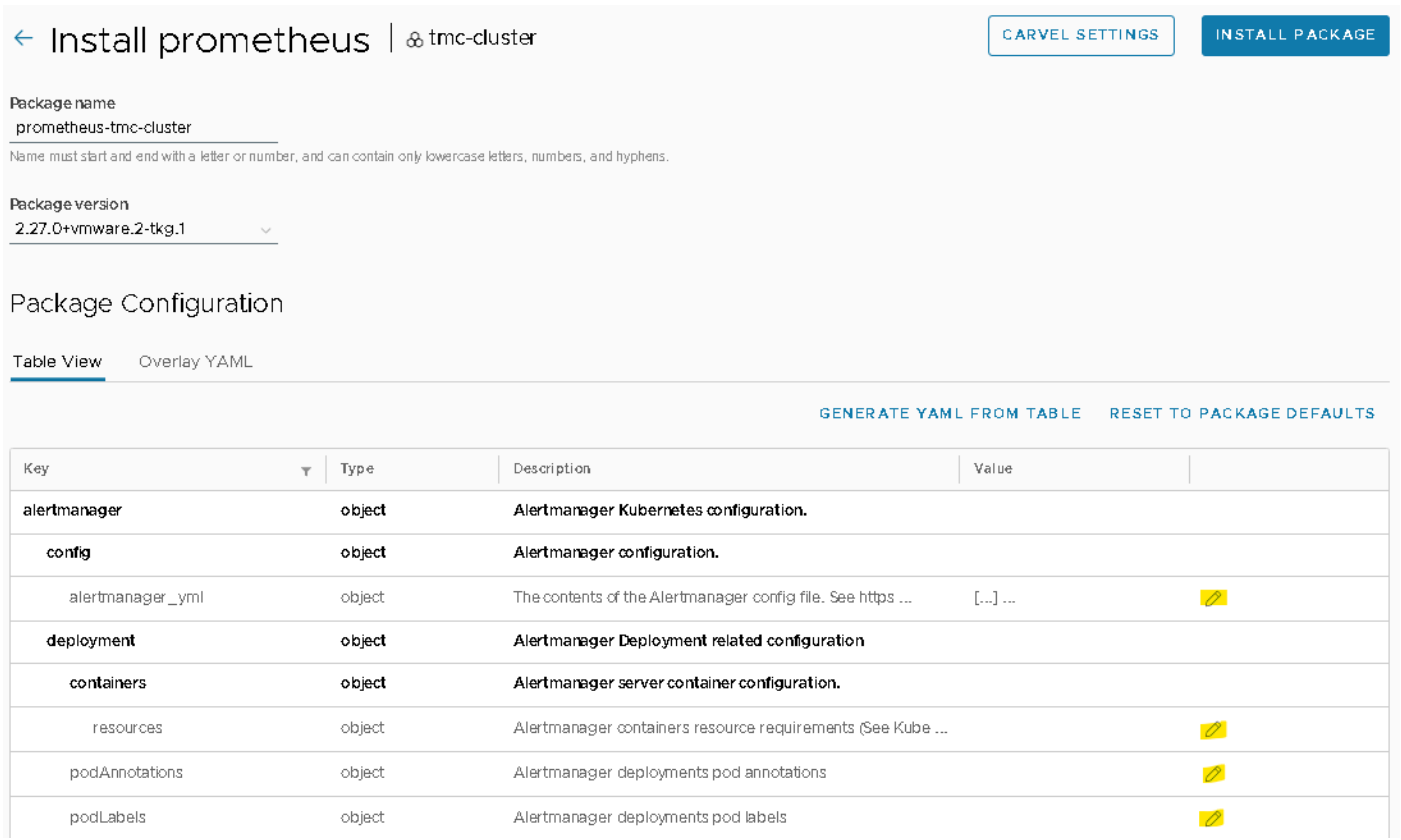

3. Some Carvel Package settings can be modified such as Carvel Resources namespace via the “Carvel Settings” button.

Figure 151: Carvel settings

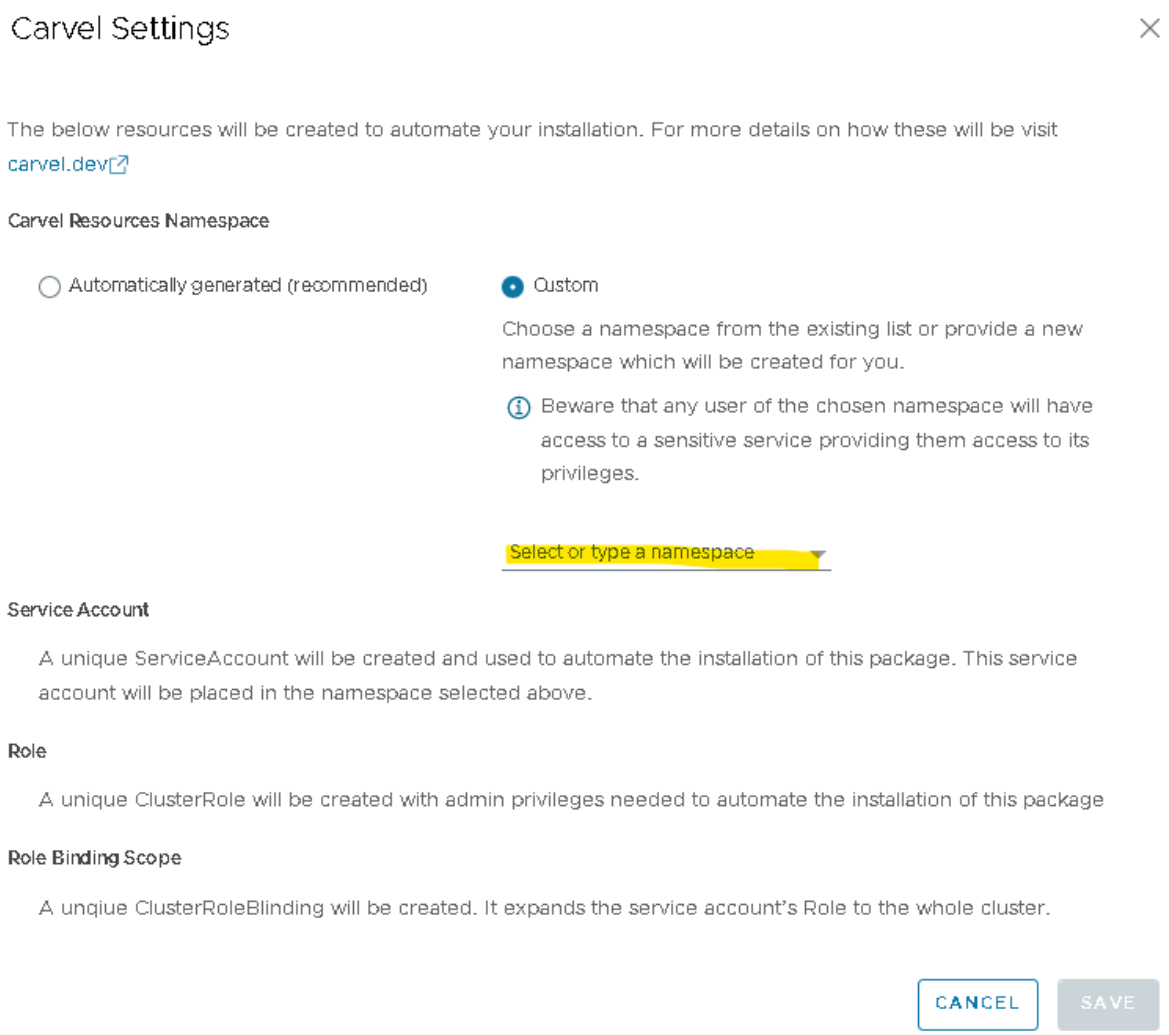

4. Click install package either leaving the settings at default or modify as needed.
5. Once Prometheus is installed successfully, install Grafana similarly.

Figure 152: Prometheus and Grafana

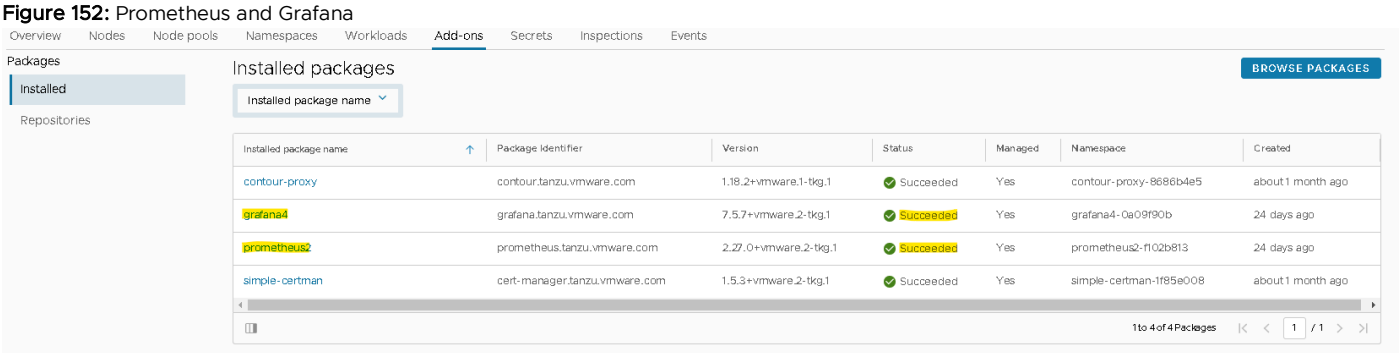

6. Verify that you can access Grafana via its external IP address. Grafana is installed in the "tanzu-system-dashboards" namespace. Use "*kubectl get svc -n tanzu-system-dashboards*" command to get the external IP address Grafana is running on.

Figure 153: Grafana

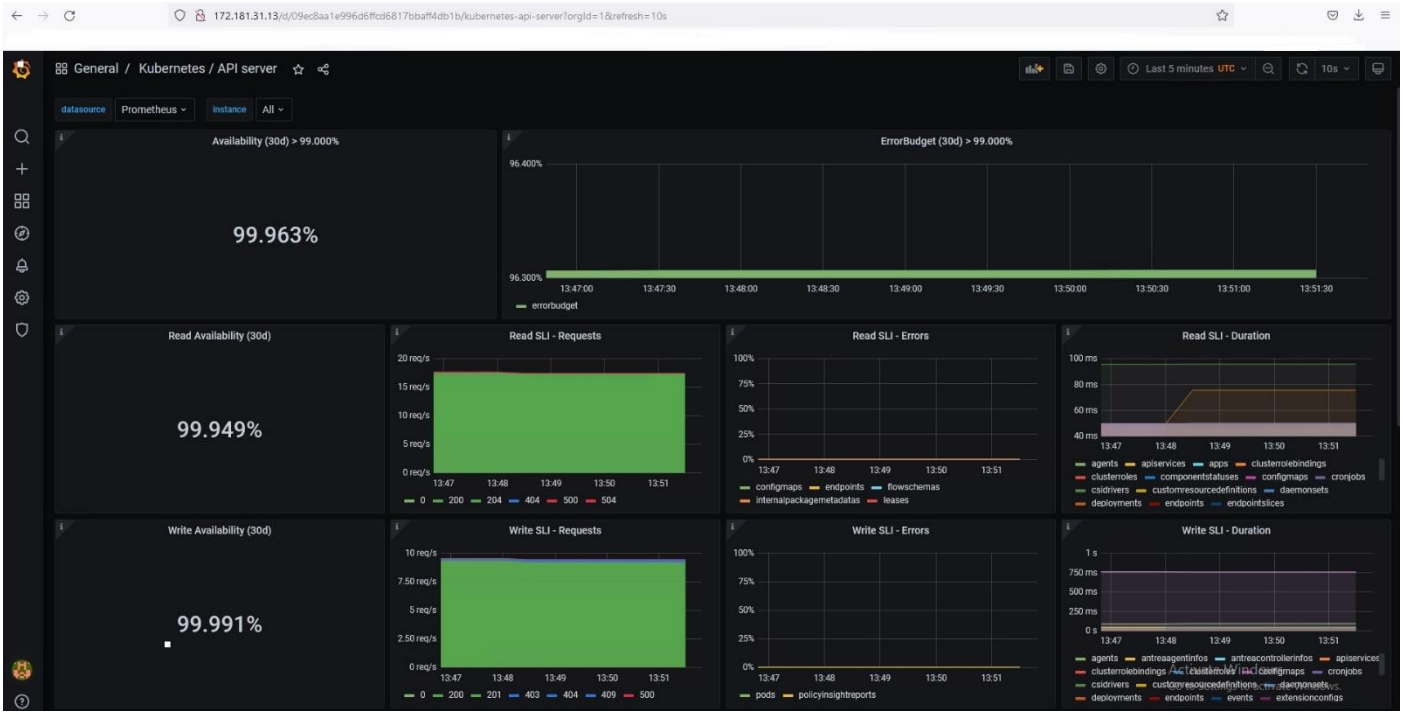

In addition to Tanzu packages and extensions, Tanzu Mission Control can deploy variety of applications from the "Catalog" using helm charts.

Figure 154: Helm Charts

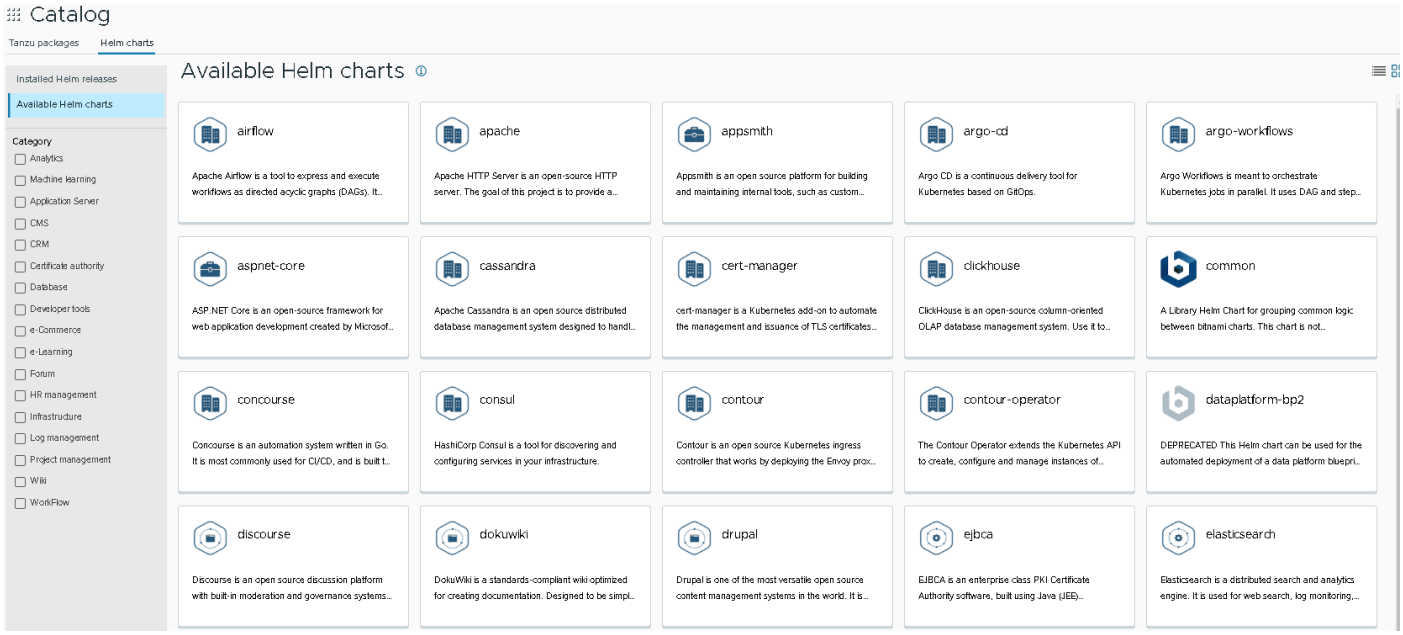

Configure Policies

Various types of policies can be created by the platform administrator to manage operations of Kubernetes environments or other organizational objects. The two policies most relevant to Kubernetes operations are Role

Based Access Control (RBAC) and Security Policies. Please note that security policies are supported on Kubernetes version 1.16 or higher. The application of these policies is discussed in the following section. For more information on policies, roles and role-bindings, please see Policy-Driven Cluster Management.

RBAC and Role binding

Access policies control how users and groups access and manage resources, such as clusters via Tanzu Mission Control. Organizations have predefined roles that govern access to an object based on granted permissions, whereas role binding defines the scope of the access policy to which the role applies. Roles are bound to a given user or group effectively granting permissions to the user or group of users to the desired object. The following example binds a user identity to a cluster via Tanzu Mission Control policy management engine.

1. From left pane in Tanzu Mission Control navigate to Policies > Assignments > Access tab > Clusters and select the cluster or a group of clusters you want to apply the policy to. Expand the cluster name under "Direct access policies."

Figure 155: Apply role binding.

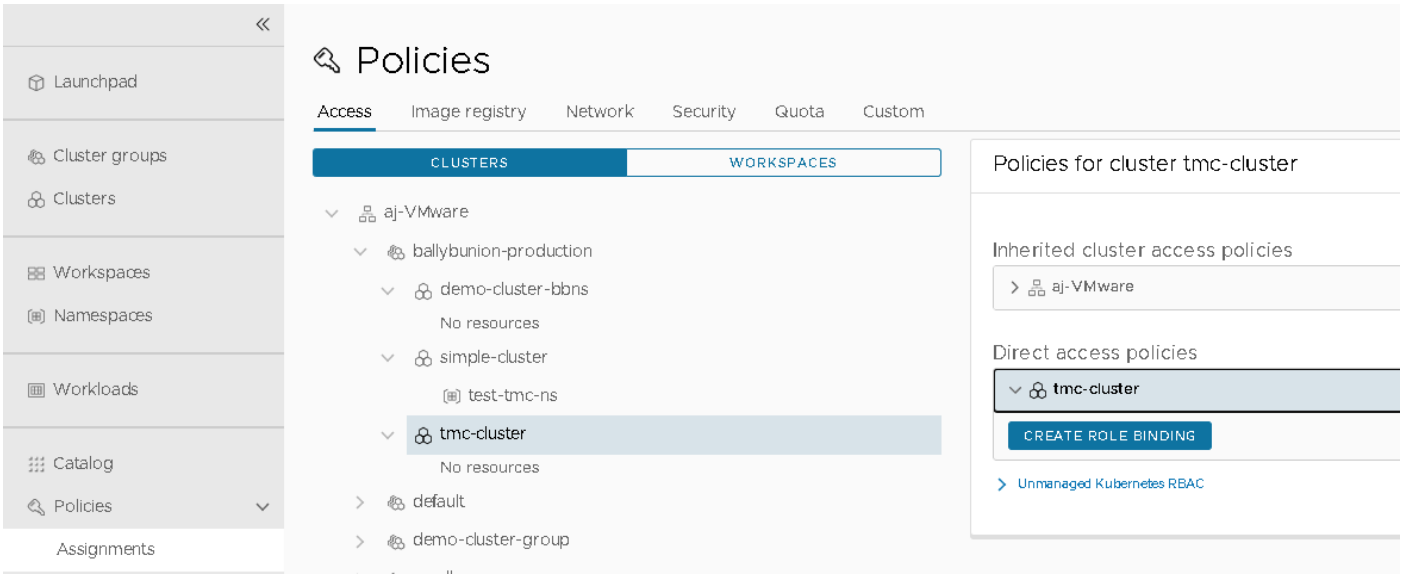

2. Create role binding for a user and assign a cluster level role.

Figure 156: Create role binding.

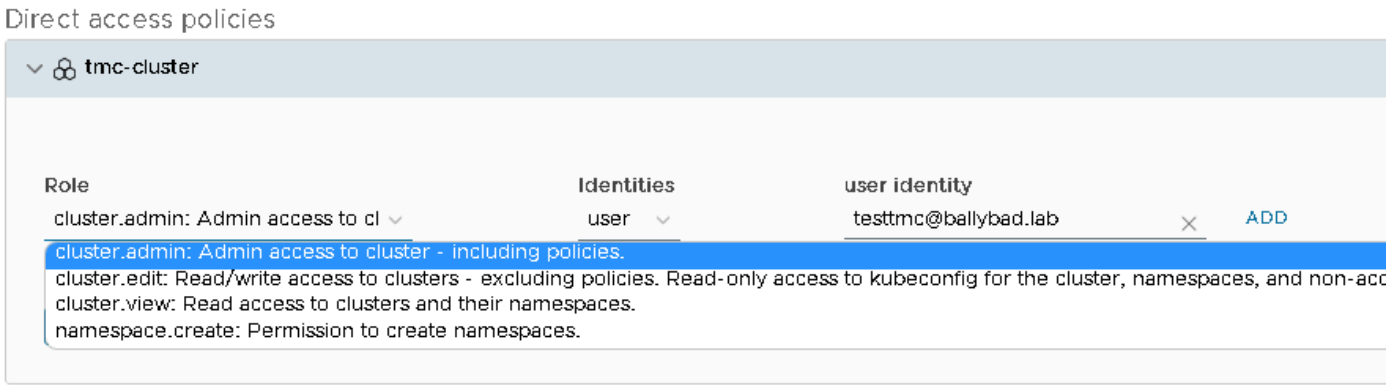

3. Click ADD and SAVE. Role binding will be created.

Figure 157: Role binding created

4. Verify that role binding is created on the cluster correctly. Use "*Kubectl describe*" command to view role binding configured.

Figure 158: Verify

```
aj@tmc-clusters-cli:~/cli/tmcli$ ktmc describe clusterrolebinding rid-c-f5389320-
k8s.io-ClusterRole-cluster-admin
Name:         rid-c-f5389320-bd9e-4c63-a6e9-82422c09bb27-ballybunion-tkgs-bbns-t
Labels:       tmc.cloud.vmware.com/managed=true
Annotations:  tmc.cloud.vmware.com/iam-policy-rid: rid:c:f5389320-bd9e-4c63-a6e9-
              tmc.cloud.vmware.com/iam-policy-version: 75361
Role:
  Kind:  ClusterRole
  Name:  cluster-admin
Subjects:
  Kind  Name                 Namespace
  ----  ----                 ---------
  User  testtmc@ballybad.lab
aj@tmc-clusters-cli:~/cli/tmcli$
```

Security Policies

Security policies allow you to manage the security context in which deployed pods operate in your clusters by imposing constraints on your clusters that define what pods can do and which resources they have access to. Tanzu Mission Control security policies are not implemented using the Kubernetes native "PodSecurityPolicy" object. Tanzu Mission Control uses Gatekeeper project from Open Policy Agent (OPA Gatekeeper). The security-sensitive aspects of the pod specification that they control are, however, the same. For more information, see the OPA Gatekeeper documentation. Tanzu Mission Control with Tanzu Standard only supports pre-defined, "Basic" and "Strict" policies. For custom policy implementation Tanzu Advanced is required. Security Policies can be assigned via Policies > Assignments > Security Tab. Below is an example of how to configure and verify security policies.

1. Select the cluster or group of clusters the policy will be applied to.

Figure 159: Select cluster.

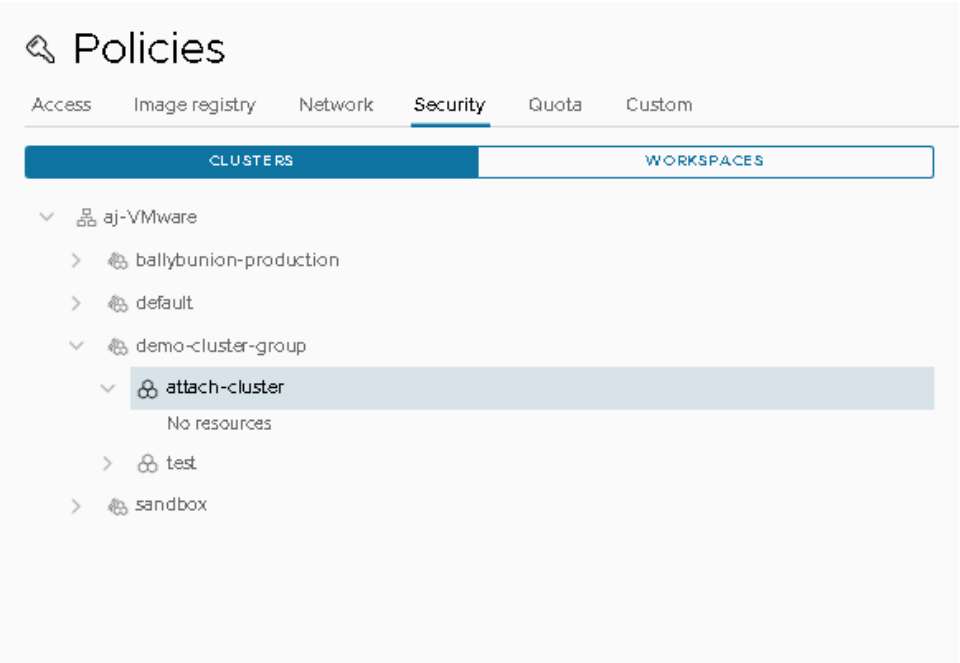

2. Under "Direct Security Policies" click "create Security Policy." Select either Basic or Strict security template per your requirements. Give policy a name and enter label selector information if required.

Figure 160: Create policy.

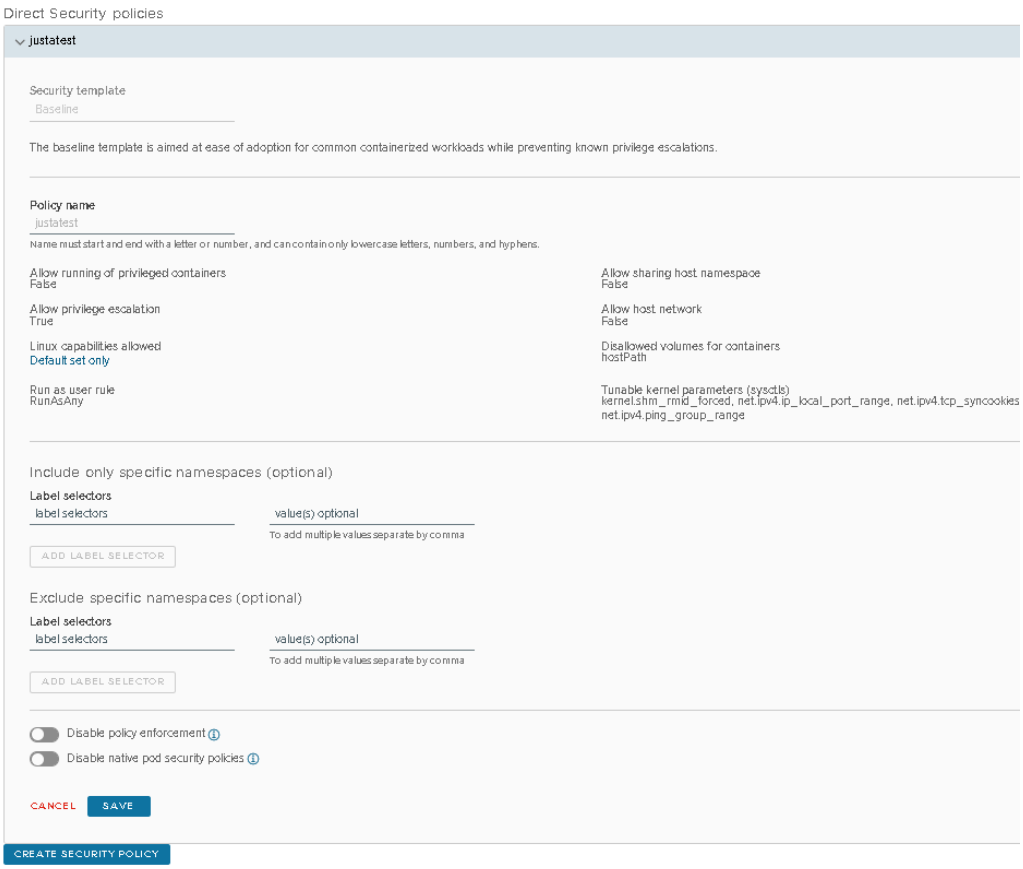

3. Verify that policy is applied to the cluster. Since policies are applied via Gatekeeper constraints and not Kubernetes native POD security policy, you will run command "*kubectl get constraints*" to display applied policies to the cluster. Each constraint that has been applied will be appended by the policy name.

Figure 161: Verify constraints

```
aj@demo-vspherecli:~/cli$ kubectl get constraints
NAME                                                                                                  AGE
vmware-system-tmc-block-privileged-container-v1.constraints.gatekeeper.sh/tmc.cp.justatest           17h

NAME                                                                                                  AGE
vmware-system-tmc-linux-capabilities-v1.constraints.gatekeeper.sh/tmc.cp.justatest                    17h

NAME                                                                                                  AGE
vmware-system-tmc-allowed-users-v1.constraints.gatekeeper.sh/tmc.cp.justatest                         17h

NAME                                                                                                  AGE
vmware-system-tmc-block-host-namespace-v1.constraints.gatekeeper.sh/tmc.cp.justatest                  17h

NAME                                                                                                  AGE
vmware-system-tmc-enforce-host-networking-v1.constraints.gatekeeper.sh/tmc.cp.justatest               17h

NAME                                                                                                  AGE
vmware-system-tmc-forbidden-sysctls-v1.constraints.gatekeeper.sh/tmc.cp.justatest                     17h
```

Quota Policies

Quota policies restricts or set boundaries on usage of cluster resource usage. In Tanzu Mission Control there are three preconfigured templates (small, medium, large) that define common limits on CPU and memory requests. There is also a custom template that allows you specify CPU, memory, and storage limits, as well as limits on a variety of object types, including those listed under Object Count Quota in the Kubernetes documentation.

Below example illustrates the use of quotas for a particular namespace (yelb).

1. A sample namespace is created in a cluster with no quotas attached.

Figure 162: Create namespace.

```
aj@azure-cli:~/cli$ aks "kubectl create ns yelb"
command started at 2023-01-30 20:44:33+00:00, finished at 2023-01-30 20:44:33+00:00 with exitcod
namespace/yelb created

aj@azure-cli:~/cli$ aks "kubectl describe ns yelb"
command started at 2023-01-30 20:44:54+00:00, finished at 2023-01-30 20:44:54+00:00 with exitcod
Name:         yelb
Labels:       kubernetes.io/metadata.name=yelb
Annotations:  <none>
Status:       Active

No resource quota.

No LimitRange resource.

aj@azure-cli:~/cli$
```

2. On Tanzu Mission Control the Policies page, click the Quota tab and use the tree control to navigate to and select the cluster or group object for the quota policy needs to be created.

Figure 163: Quota Policy

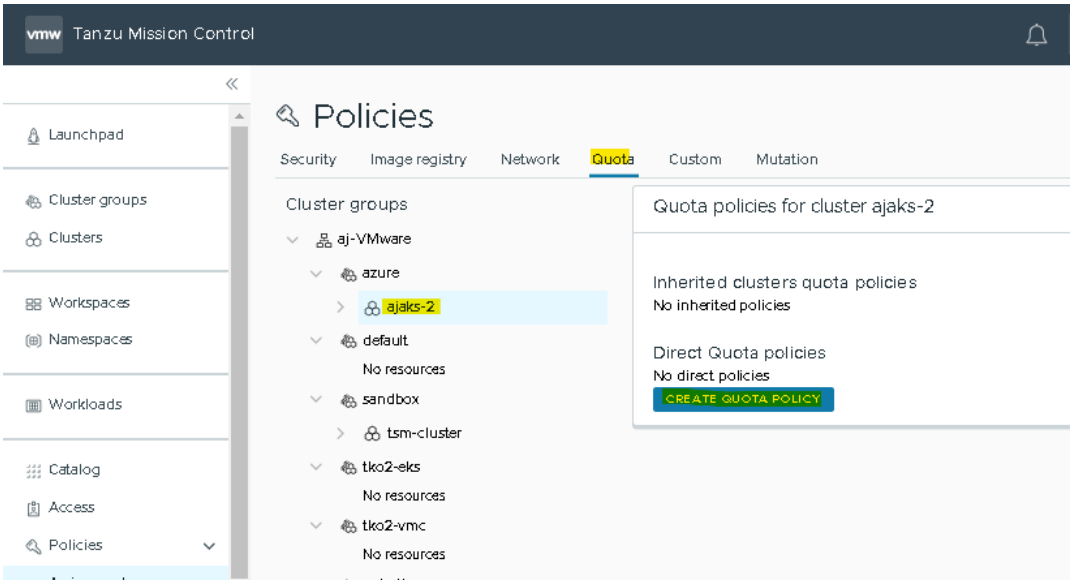

3. Select the policy template to use either from the predefined list or create a custom policy. In this example "small" predefined policy is used.

Figure 164: Quota Policy

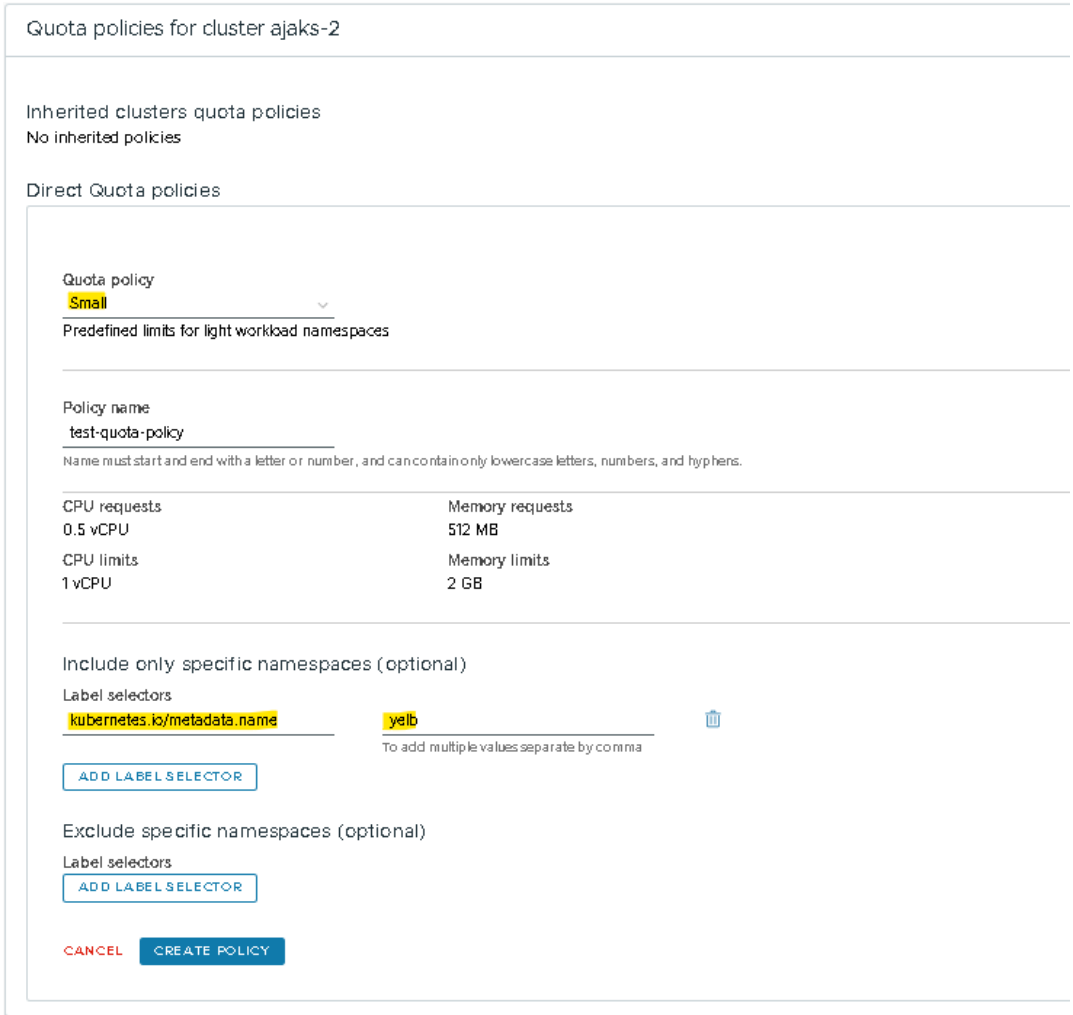

4. Optionally add label selectors to include or exclude in the calculation of aggregate resource usage. In this example the namespace is used to define the aggregate limit is assigned.
5. Optionally repeat this step to add more label selectors for this policy. Click Create Policy.

Screen shot below shows that the quota was applied to the namespace.

Figure 165: Quota Applied

```
aj@azure-cli:~/cli$ aks "kubectl describe ns yelb"
command started at 2023-01-30 20:44:54+00:00, finished at 2023-01-30 20:44:54+00:00 wit
Name:         yelb
Labels:       kubernetes.io/metadata.name=yelb
Annotations:  <none>
Status:       Active

No resource quota.

No LimitRange resource.

aj@azure-cli:~/cli$ aks "kubectl describe ns yelb"
command started at 2023-01-30 21:00:20+00:00, finished at 2023-01-30 21:00:20+00:00 wit
Name:         yelb
Labels:       kubernetes.io/metadata.name=yelb
Annotations:  <none>
Status:       Active

Resource Quotas
  Name:            tmc.cp.test-quota-policy
  Resource         Used  Hard
  --------         ---   ---
  limits.cpu       0     1
  limits.memory    0     2Gi
  requests.cpu     0     500m
  requests.memory  0     512Mi

No LimitRange resource.

aj@azure-cli:~/cli$
```

DevOps access and automation via Tanzu Mission Control CLI

Tanzu Mission Control provides resource management, including clusters via Tanzu Mission Control CLI that can be downloaded via the Tanzu Mission Control portal. In addition, Tanzu Mission Control provides Tanzu Mission Control API and Terraform to manage Tanzu Kubernetes Grid clusters.

Figure 166: Tanzu Mission Control CLI

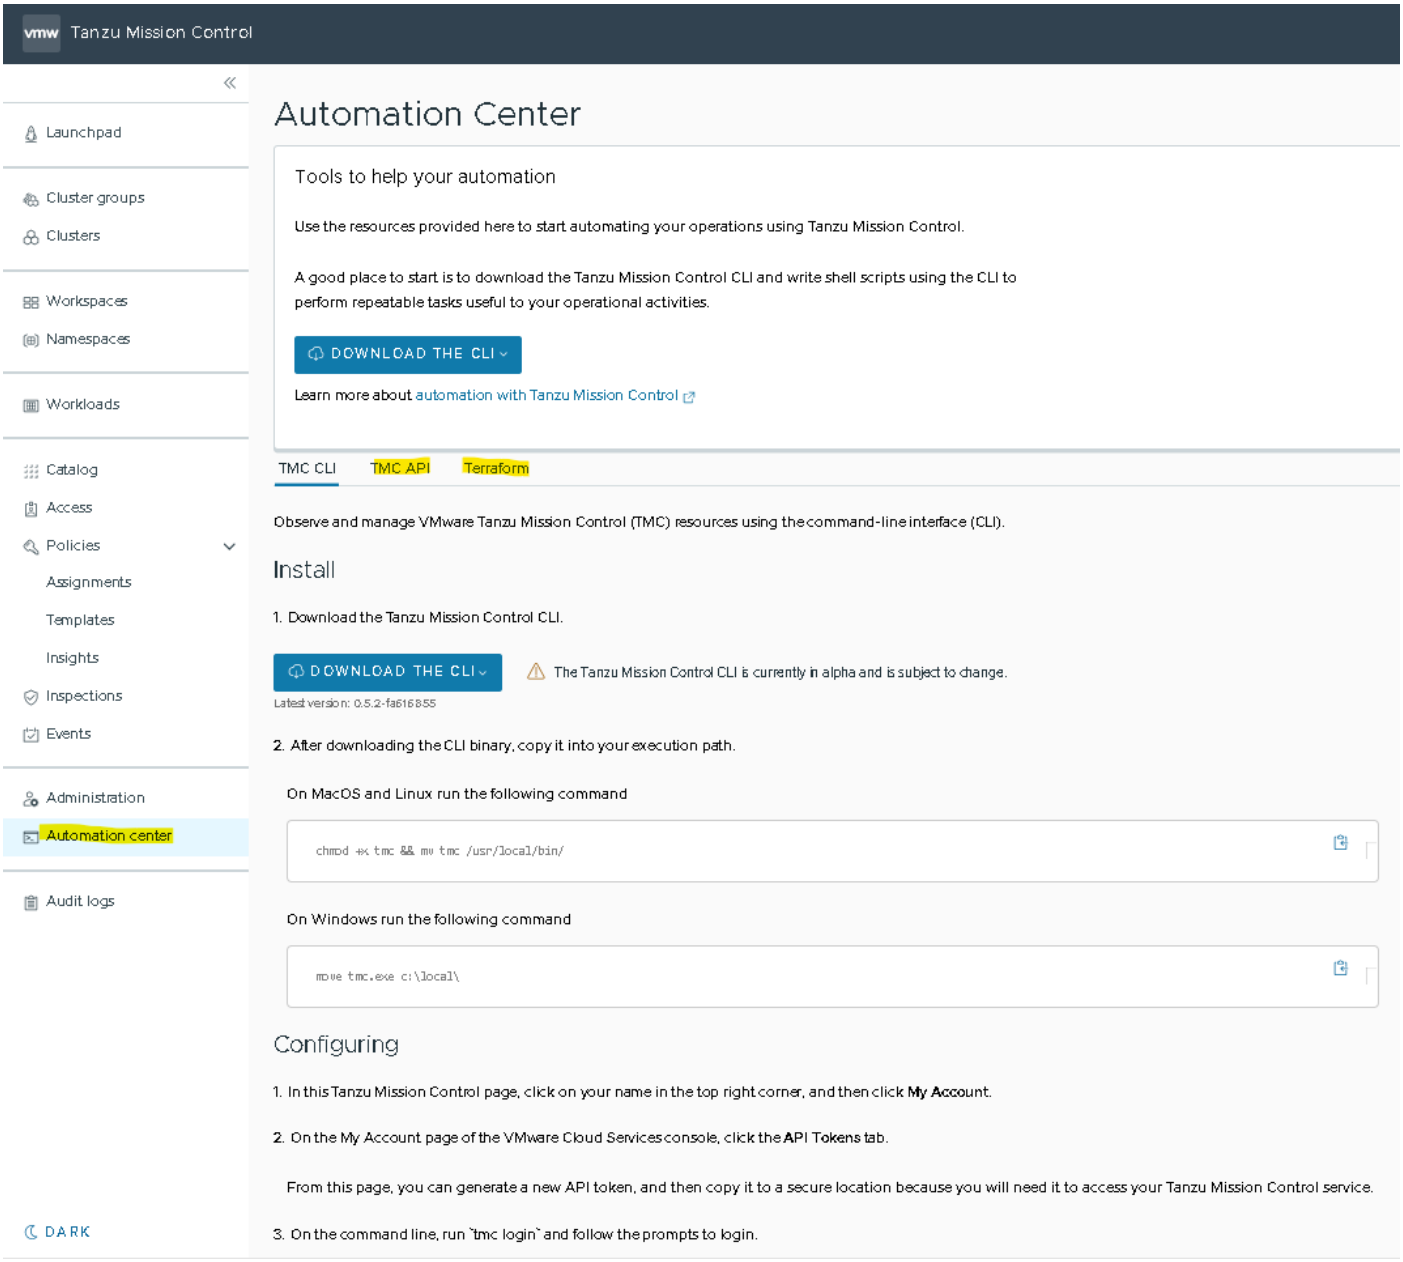

In addition to Tanzu Kubernetes clusters, Tanzu Mission Control can manage complete lifecycle of Amazon EKS clusters as well as existing Azure AKS and Google GKE or any other supported Kubernetes cloud deployment.

Enable Continuous Deliver (CD)

Tanzu Mission Control can now be used to connect Kubernetes clusters to a Git repository, and then manage the cluster's resources declaratively from the repository. Cluster administrators can use Tanzu Mission Control to set up continuous delivery for your clusters. Administrators define the configuration of a cluster (as well as other resources like Helm packages) declaratively using YAML in a Git repository, connect the cluster to the repository, and then synchronize the repository to the cluster. After continuous delivery configuration of a cluster, Tanzu Mission Control drives the continuous delivery of repository objects to the cluster. Continuous delivery can be enabled with or without authentication, depending on requirements.

Tanzu Mission Control uses Flux (an open-source community standard) for continuous delivery. Flux uses "Kustomize" to synchronize YAML to your cluster. Kustomize is a standalone tool used to customize Kubernetes objects. Although it is commonly used to apply overlay YAML to existing resources, Kustomize can also be used to create and manage new

resources. Flux CD runs in your cluster, connects to your repositories, and periodically synchronizes your defined Kustomization files to your cluster.

Figure 167: Configuration Workflow

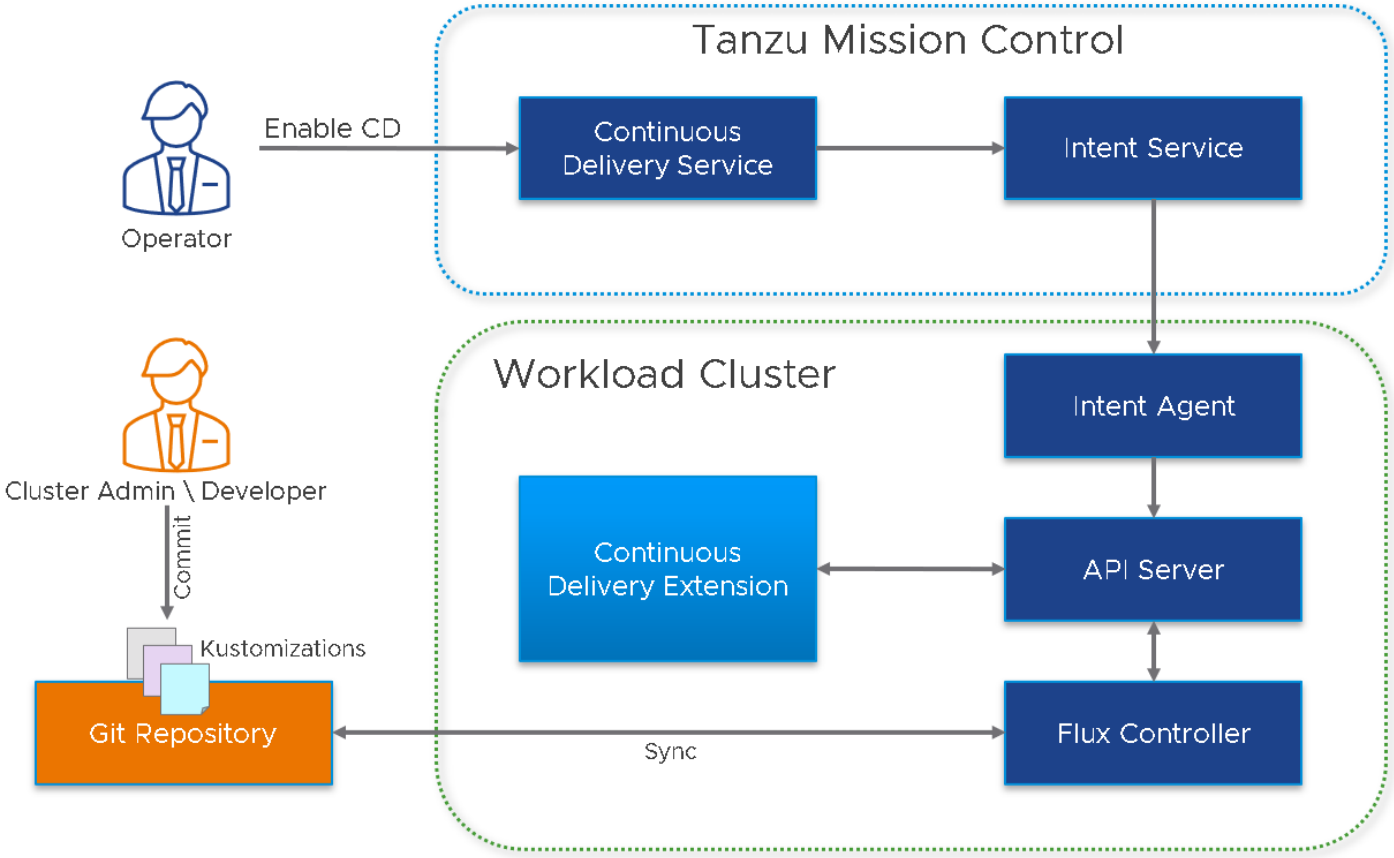

Continuous delivery can be enabled on a cluster from Cluster > “Continuous Delivery” tab.

Figure 168: Continuous delivery

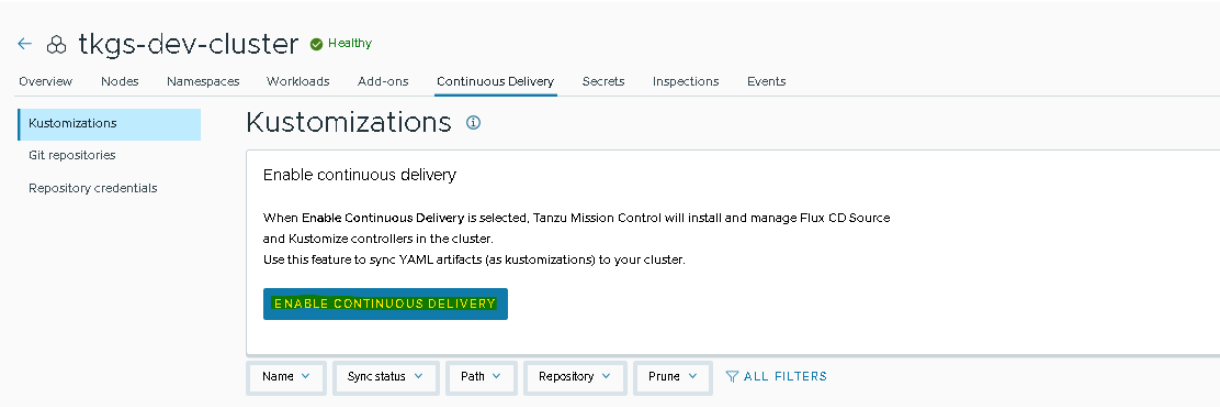

Customizations can now be added if have the repository credentials and Git repositories configured beforehand.

Figure 169: Continuous delivery

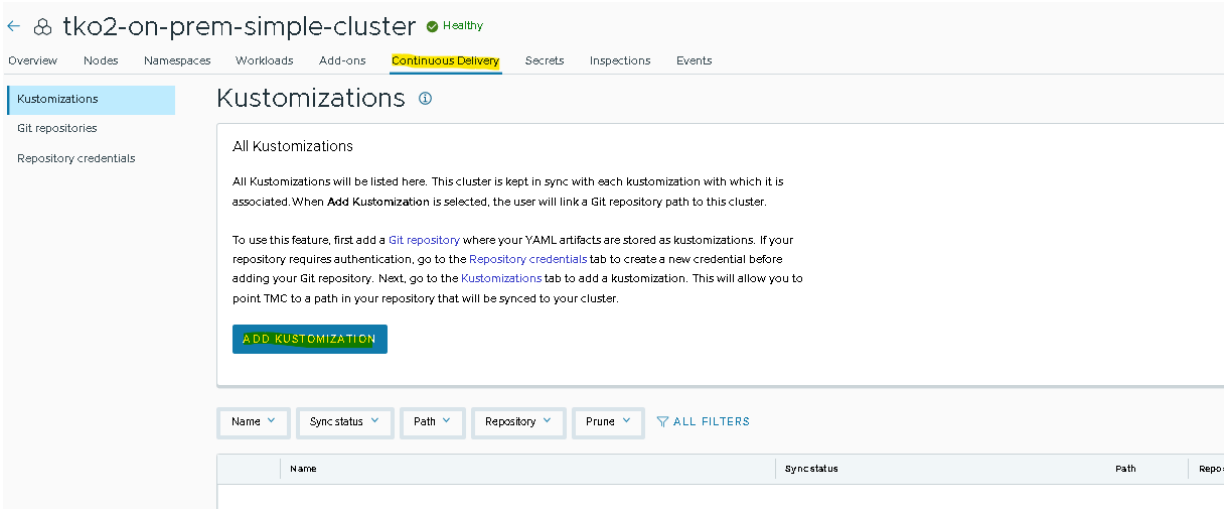

To configure credentials, click on repository credentials.

Figure 170: Continuous delivery

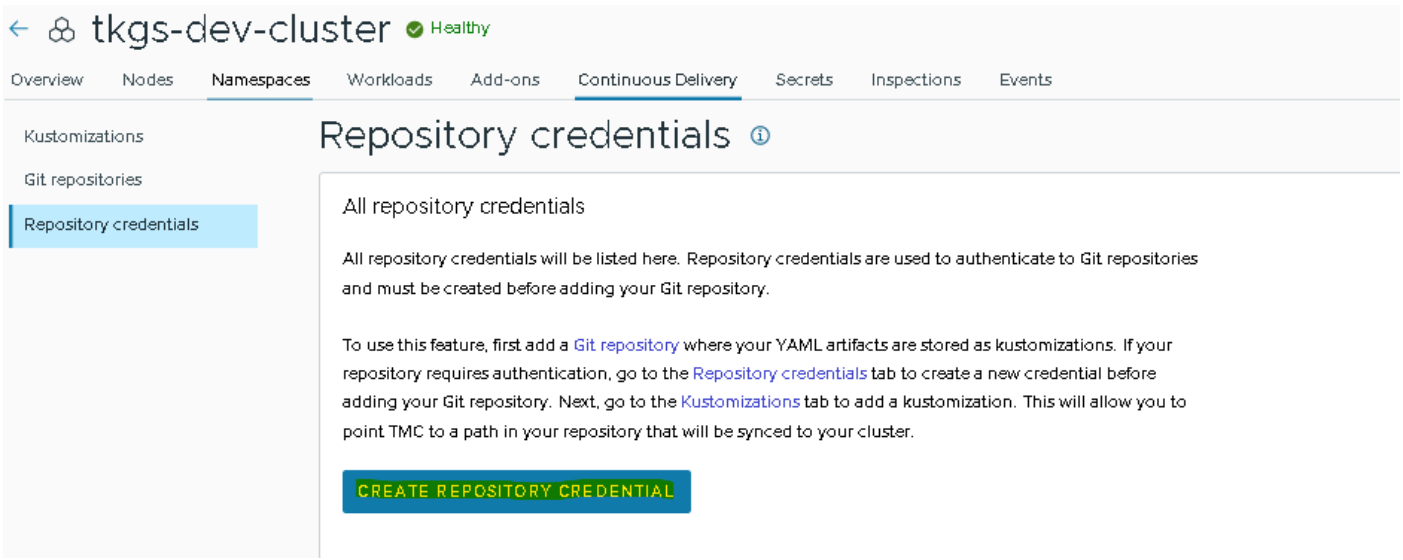

Create credentials either via Gitlab username/password or ssh key.

Figure 171: Continuous delivery

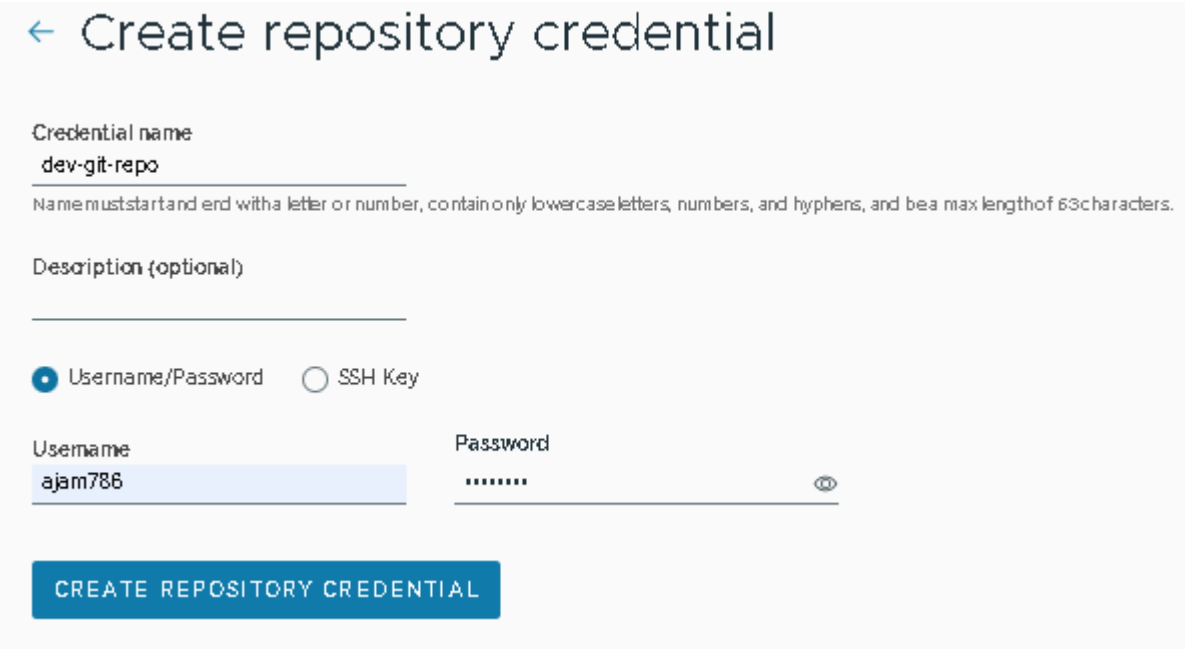

Git repository can now be added by clicking the "Add Git Repository" button and entering the Git repo information.

Figure 172: Continuous delivery

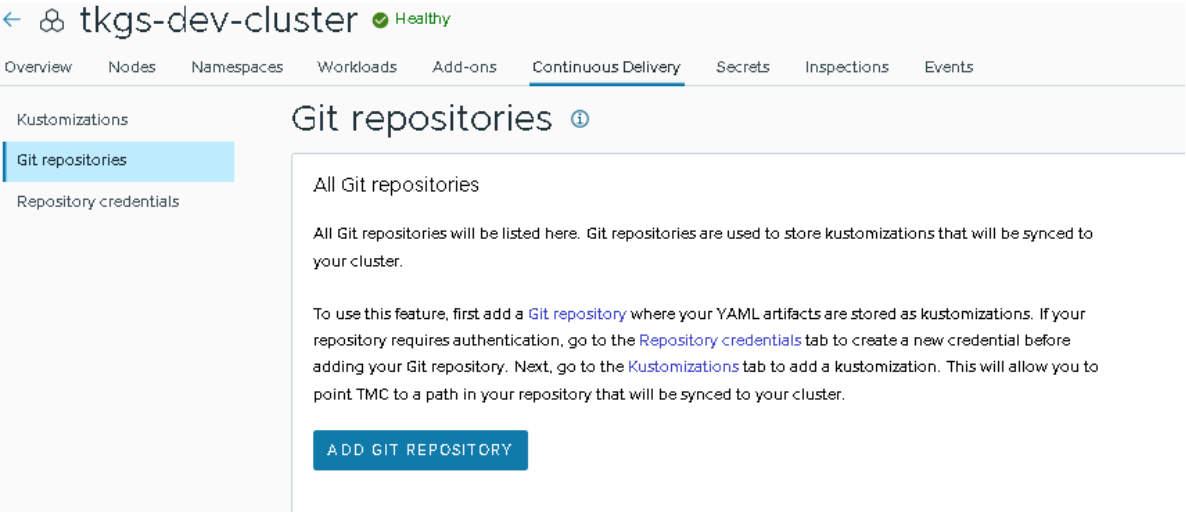

Figure 173: Continuous delivery

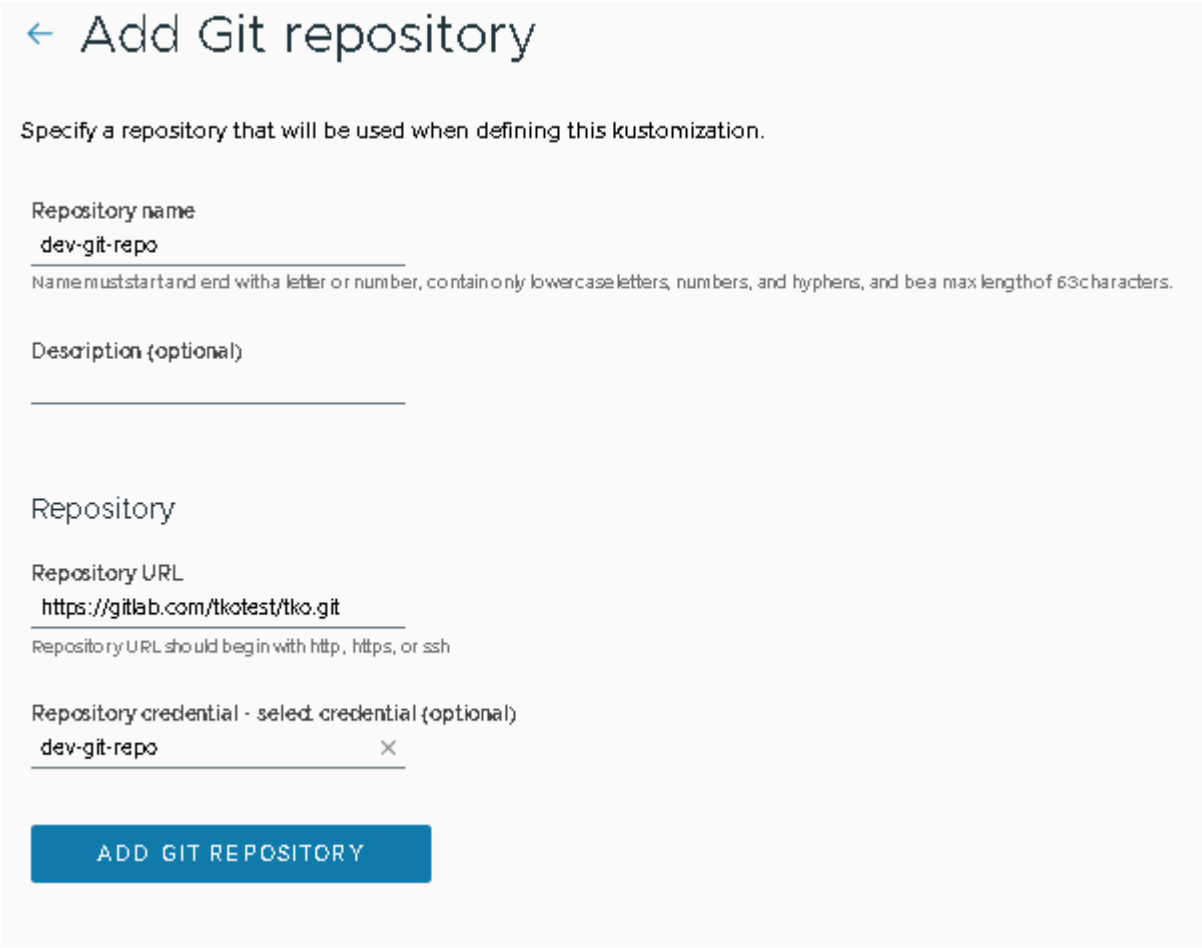

Once the credential verification process is complete, the repository will be ready to be used.

Figure 174: Continuous delivery

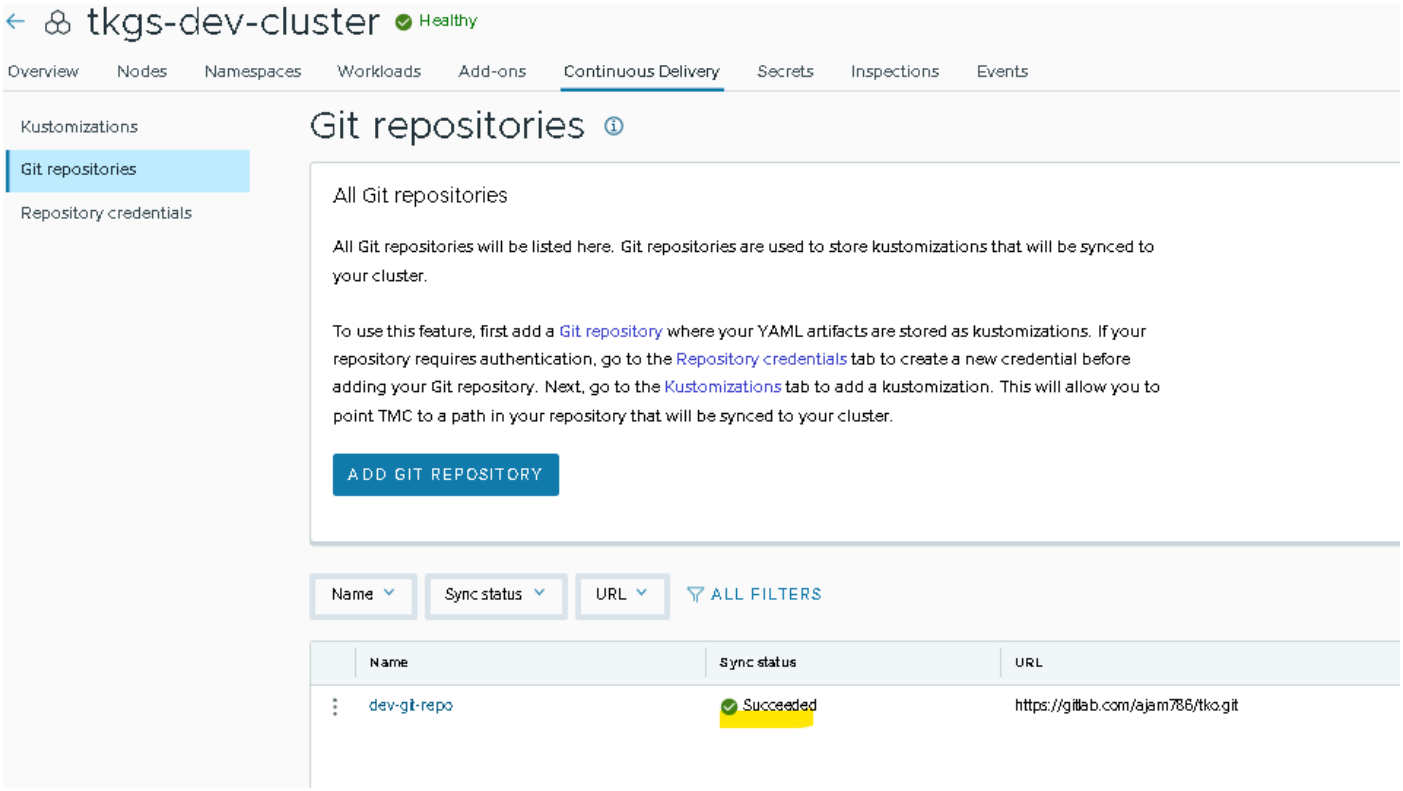

Note: For more information about continuous delivery using Flux, visit https://fluxcd.io/docs/.
For more information about Kustomize, see Declarative Management of Kubernetes Objects Using Kustomize in the Kubernetes documentation.
Please consult Tanzu Mission Control documentation for further details on Continuous Delivery feature and use cases.

Lifecycle management of Amazon EKS clusters

Tanzu Mission Control provides capability to manage complete lifecycle of Amazon EKS clusters. With lifecycle management for Amazon EKS clusters, operations teams will be able to offer more choice to their developers. By centralizing management of multiple Kubernetes cluster types with Tanzu Mission Control, operations teams will be able to efficiently manage their Kubernetes estate through consistent deployment patterns and granular access control and other policies. This capability in preview and intended for general availability soon.

Pre-requisites

- A VPC created with public and private networks.
- User has access to Tanzu Mission Control role cluster.admin role, needed to create credentials.

Create credentials for Amazon EKS lifecycle management.

1. In the Tanzu Mission Control console, click Administration in the left navigation pane.
2. On the Credentials tab of the Administration page, click Create Credential and choose AWS EKS.

Figure 175: Credentials

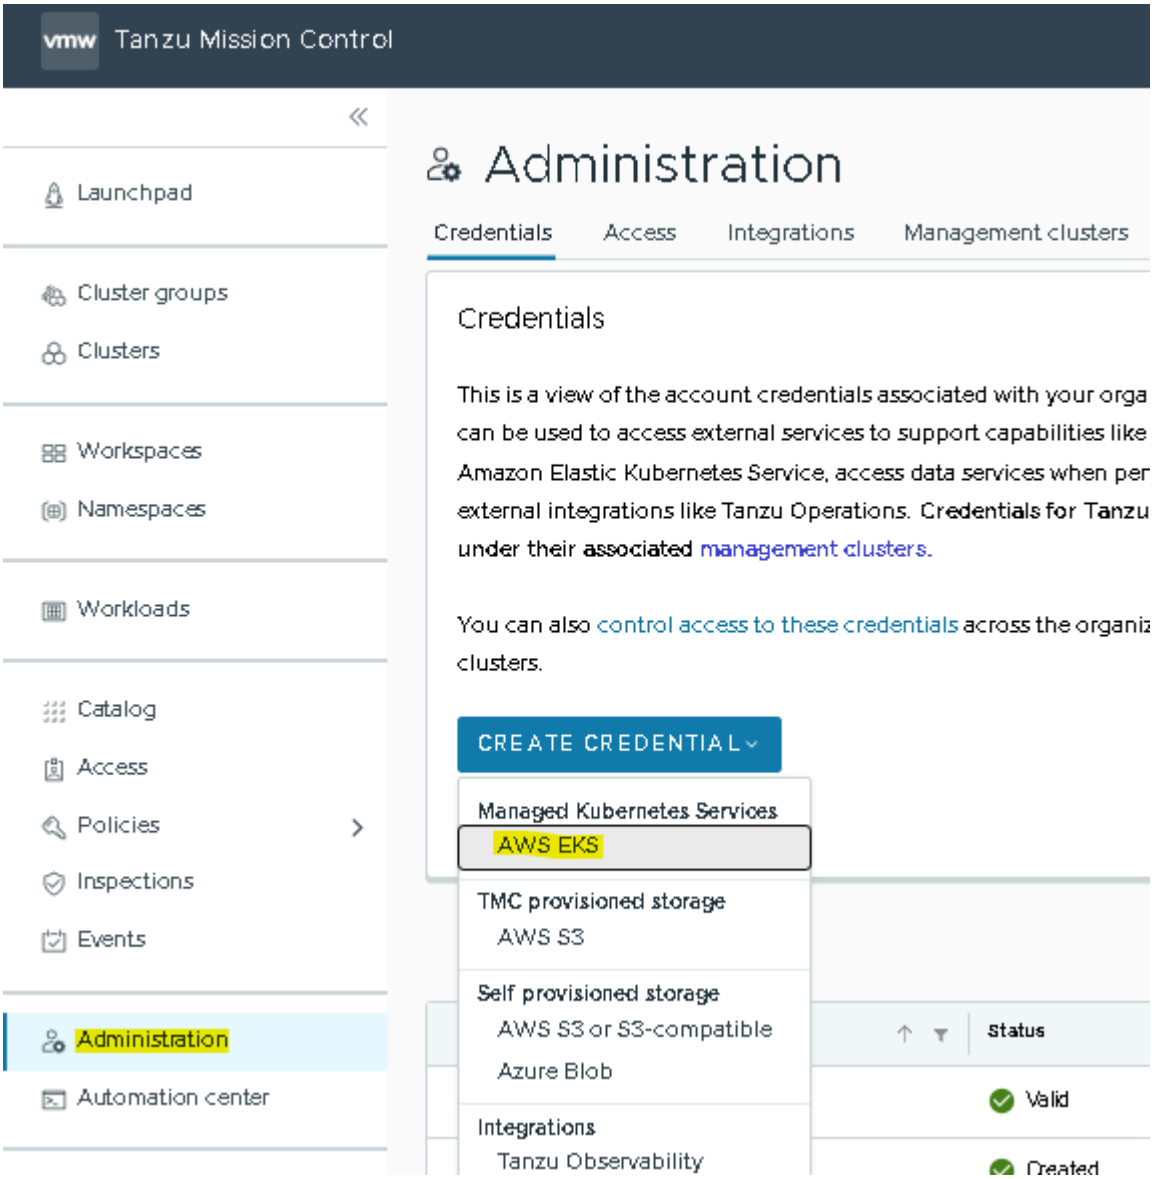

3. On the Create credential page, provide a name for the credential.
4. You can optionally provide a description and labels.
5. Click Next.

Figure 176: EKS Credentials

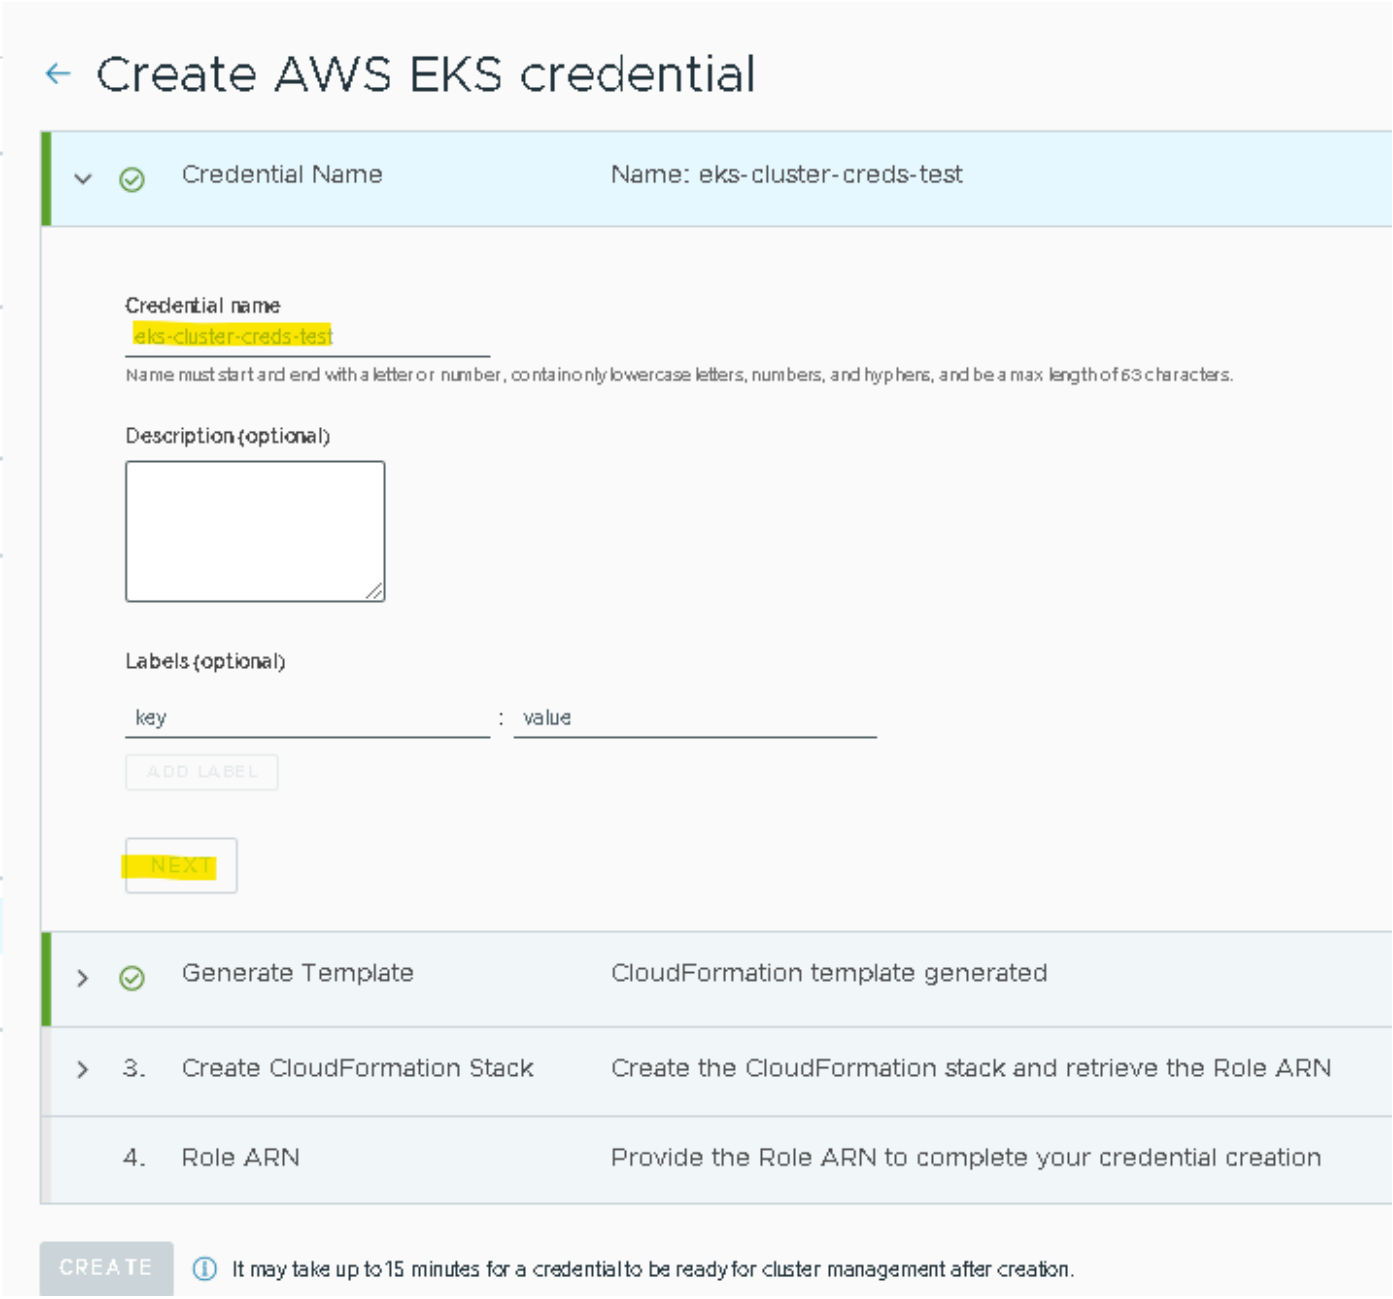

6. Click Generate Template, and then after the template is generated, click Next.

Figure 177: Generate template.

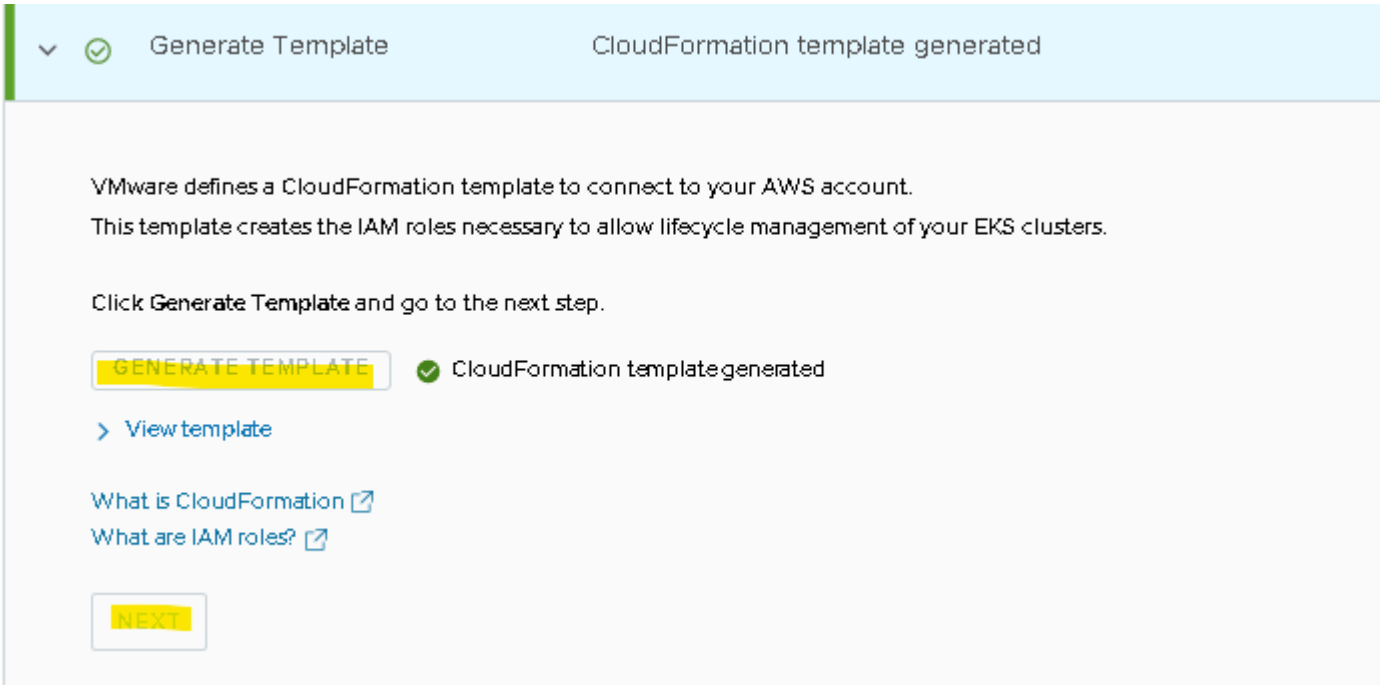

7. Use the generated template in one of two ways to create the AWS CloudFormation stack, either via the AWS CLI or the AWS console UI.

Figure 178: Create Cloud Formation stack.

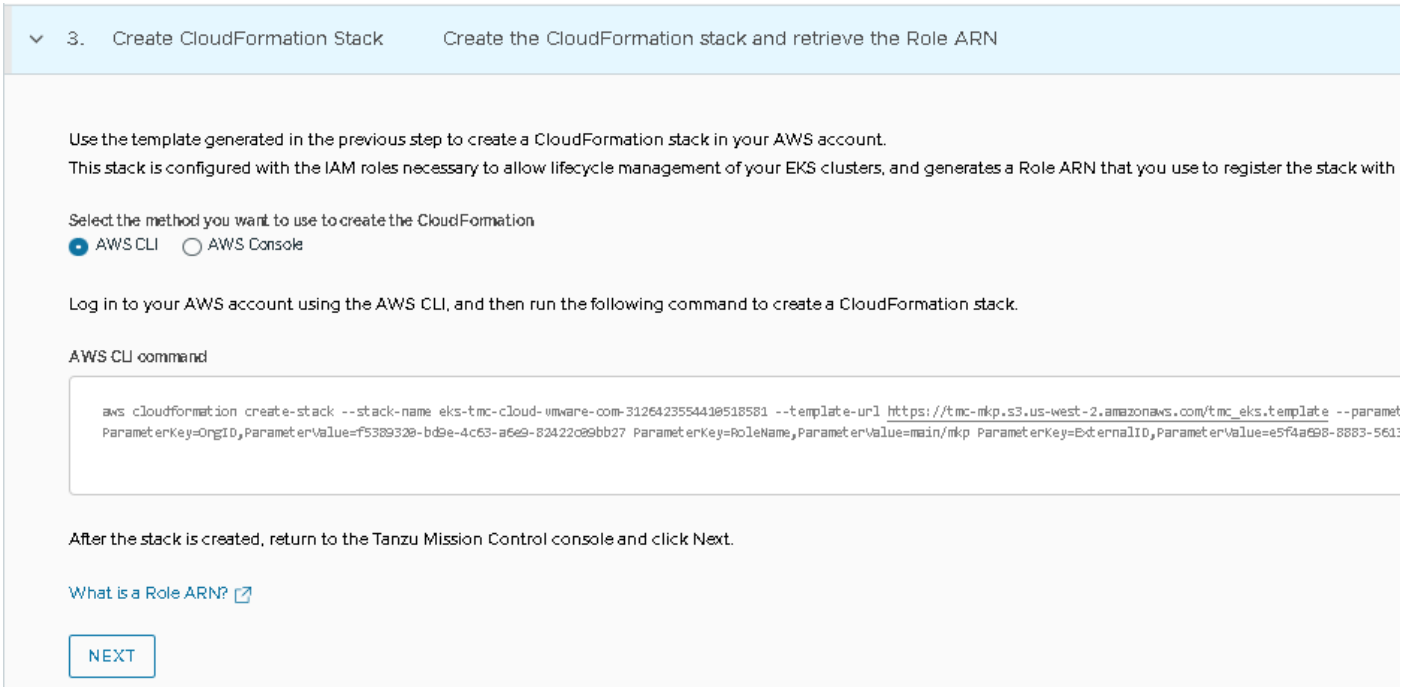

8. Retrieve the Role ARN using the CLI or AWS console UI by navigating to CloudFormation > Stacks > <your stack> > Outputs.

Figure 179: Retrieve credentials via UI.

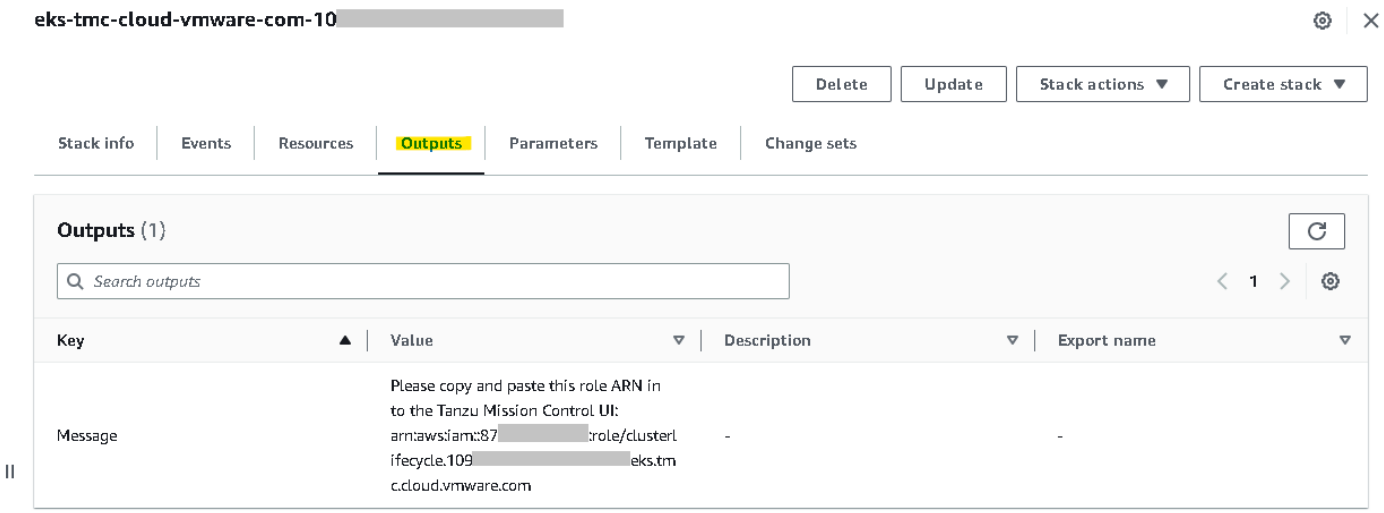

9. Provide Tanzu Mission Control with the ARN role in the last step. Credentials will be created in a few minutes.

Figure 180: Credentials created.

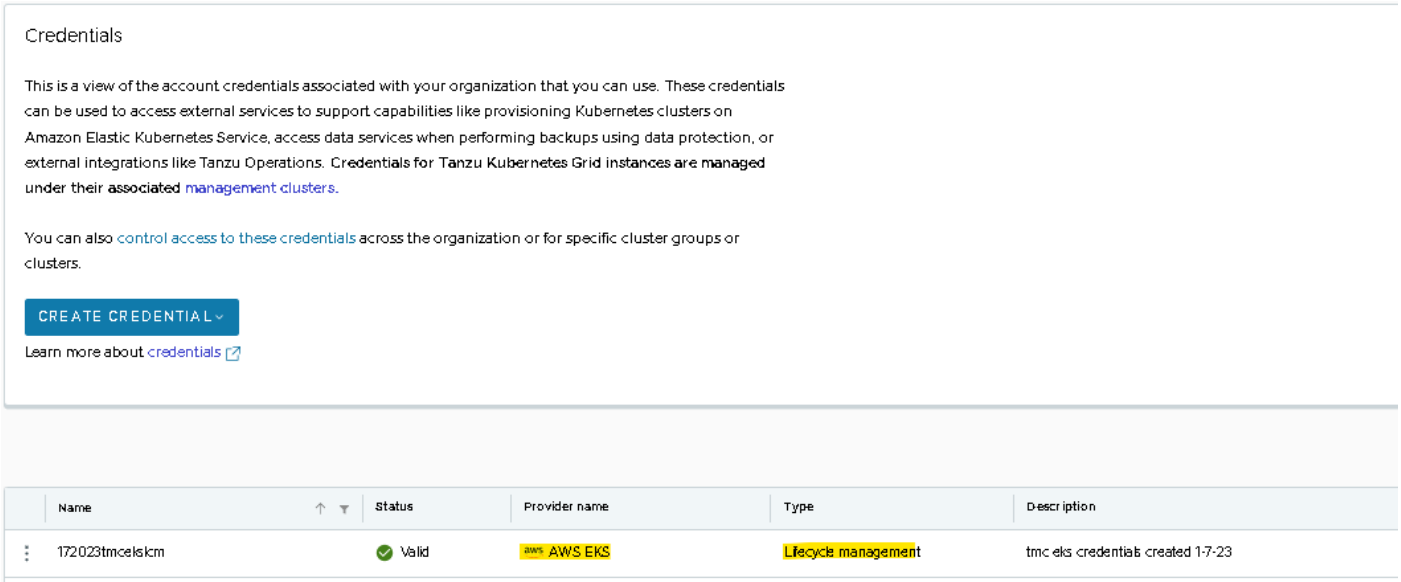

Enable data protection for clusters.

The data protection features of Tanzu Mission Control allow you to create the following types of backups for managed clusters (both attached and provisioned):

- all resources in a cluster
- selected or excluded namespaces in a cluster.
- specific or excluded resources in a cluster identified by a given label.

You can selectively restore the backups you have created, by specifying the following:

- the entire backup
- selected or excluded namespaces from the backup.
- specific or excluded resources from the backup identified by a given label.

Additionally, you can schedule regular backups and manage the storage of backups and volume snapshots you create by specifying a retention period for each backup and deleting backups that are no longer needed.

When you perform a backup for a cluster, Tanzu Mission Control uses Velero to create a backup of the specified Kubernetes resources with snapshots of persistent volume data, and then stores the backup in the location that you specify.

Note: The namespaces kube-system, velero, tkg-system, and vmware-system-tmc are not included in backups.

For the storage of your backups, you can specify a target location that allows Tanzu Mission Control to manage the storage of backups, provisioning resources as necessary according to your specifications. However, if you prefer to manage your own storage for backups, you can also specify a target location that points to a storage location that you create and maintain in your cloud provider account, such as an AWS S3 or S3-compatible storage location or an Azure Blob storage location. With self-provisioned storage, you can leverage existing storage investments for backups, reducing network and cloud storage costs, and apply existing storage policies, quotas, and encryption. For a list of supported S3-compatible providers, see S3-Compatible object store providers in the Velero documentation.

Before you define a backup for a cluster, you must create a target location and credential that you will use to perform the backup.

- The data protection credential specifies the access credentials for the account where your backup is stored. This account can be either your AWS account where Tanzu Mission Control manages backup storage, or an account where you manage backups (the account that contains your AWS S3 or S3-compatible storage or the subscription that contains your Azure Blob storage).
- The data protection target location identifies the place where you want the backup stored and references the associated data protection credential. You can share the target location across multiple cluster groups and clusters.

High-level process of creating backup on Amazon S3 follows. Similar process can be followed to backup resources on another S3 compatible storage or Azure Blob storage. The following example uses Amazon S3 storage and assumes that a S3 bucket already exists.

1. Create Amazon AWS credentials.

Figure 181: Create Credentials.

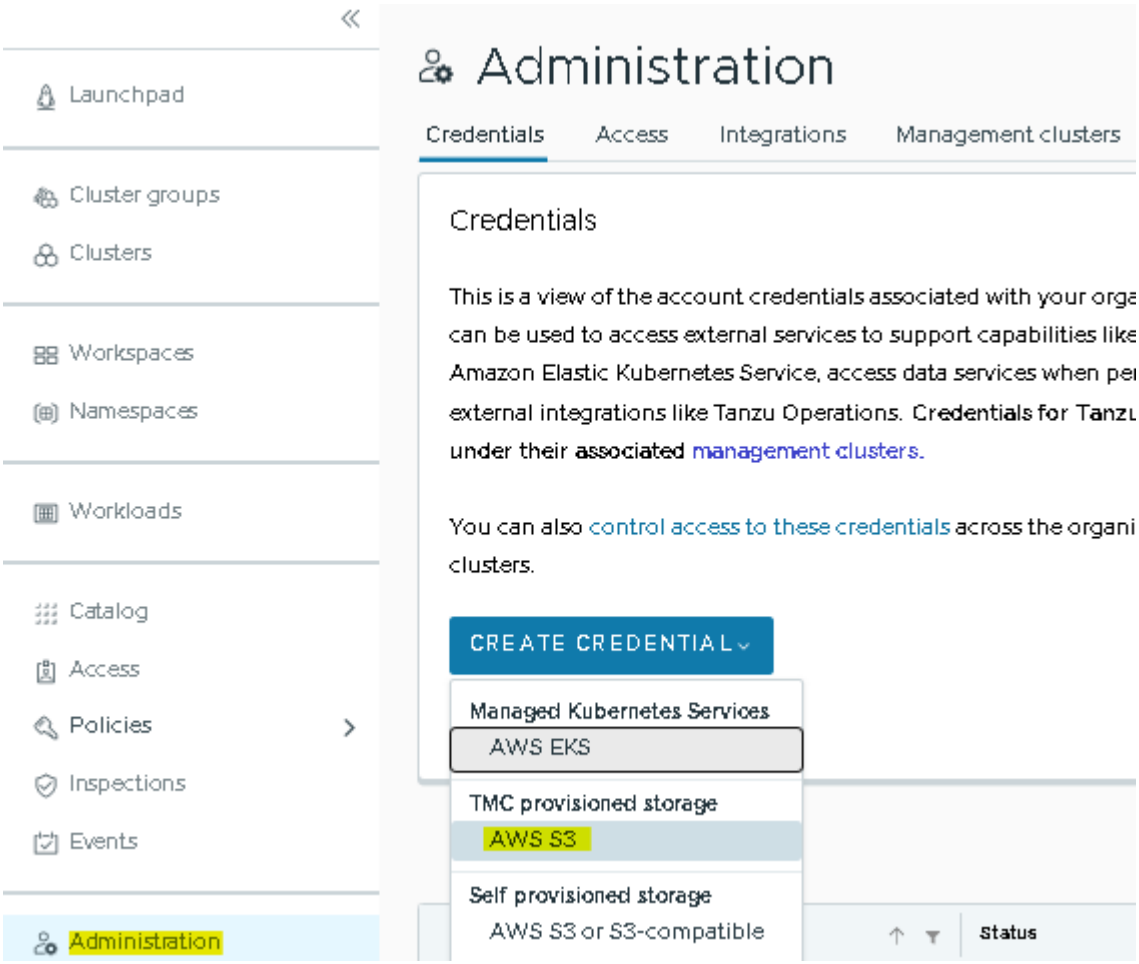

2. In Tanzu Mission Control create a backup target location
    - In the Tanzu Mission Control console, click Administration in the left navigation pane. On the Administration a page, click the Target Locations tab.
    - Click Create Target Location, and then choose the type of storage for the new target location.
    - Select Tanzu Mission Control provisioned storage: AWS S3

Figure 182: Target Location

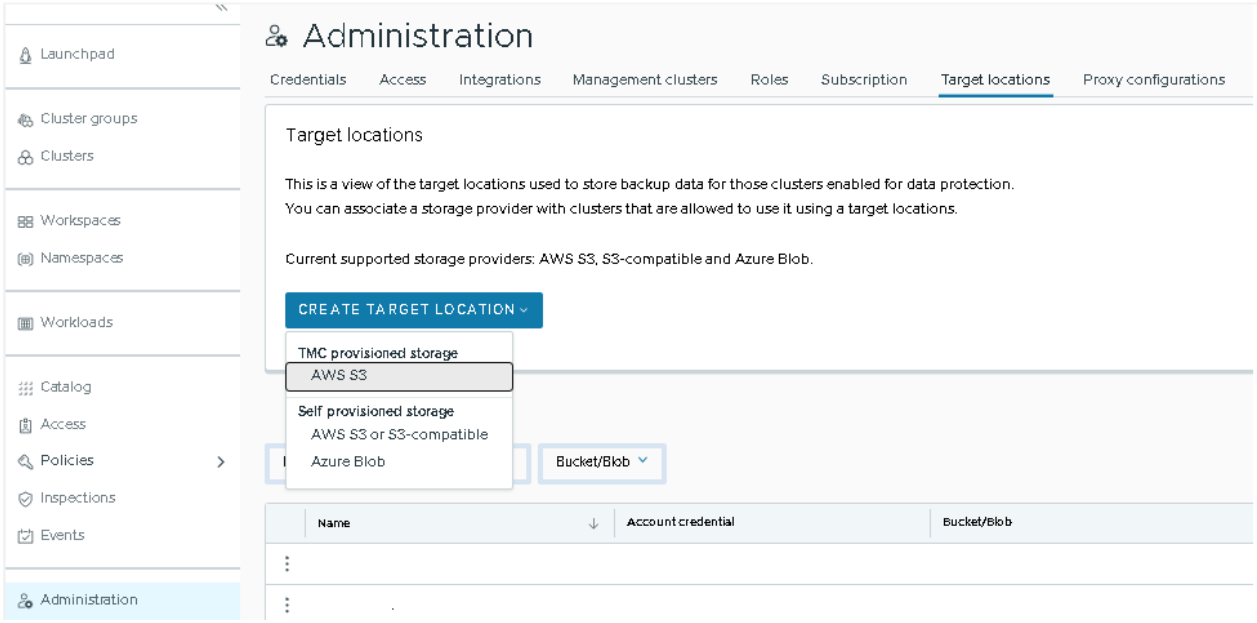

3. Fill-in the required information and create backup location.

Figure 183: Create location.

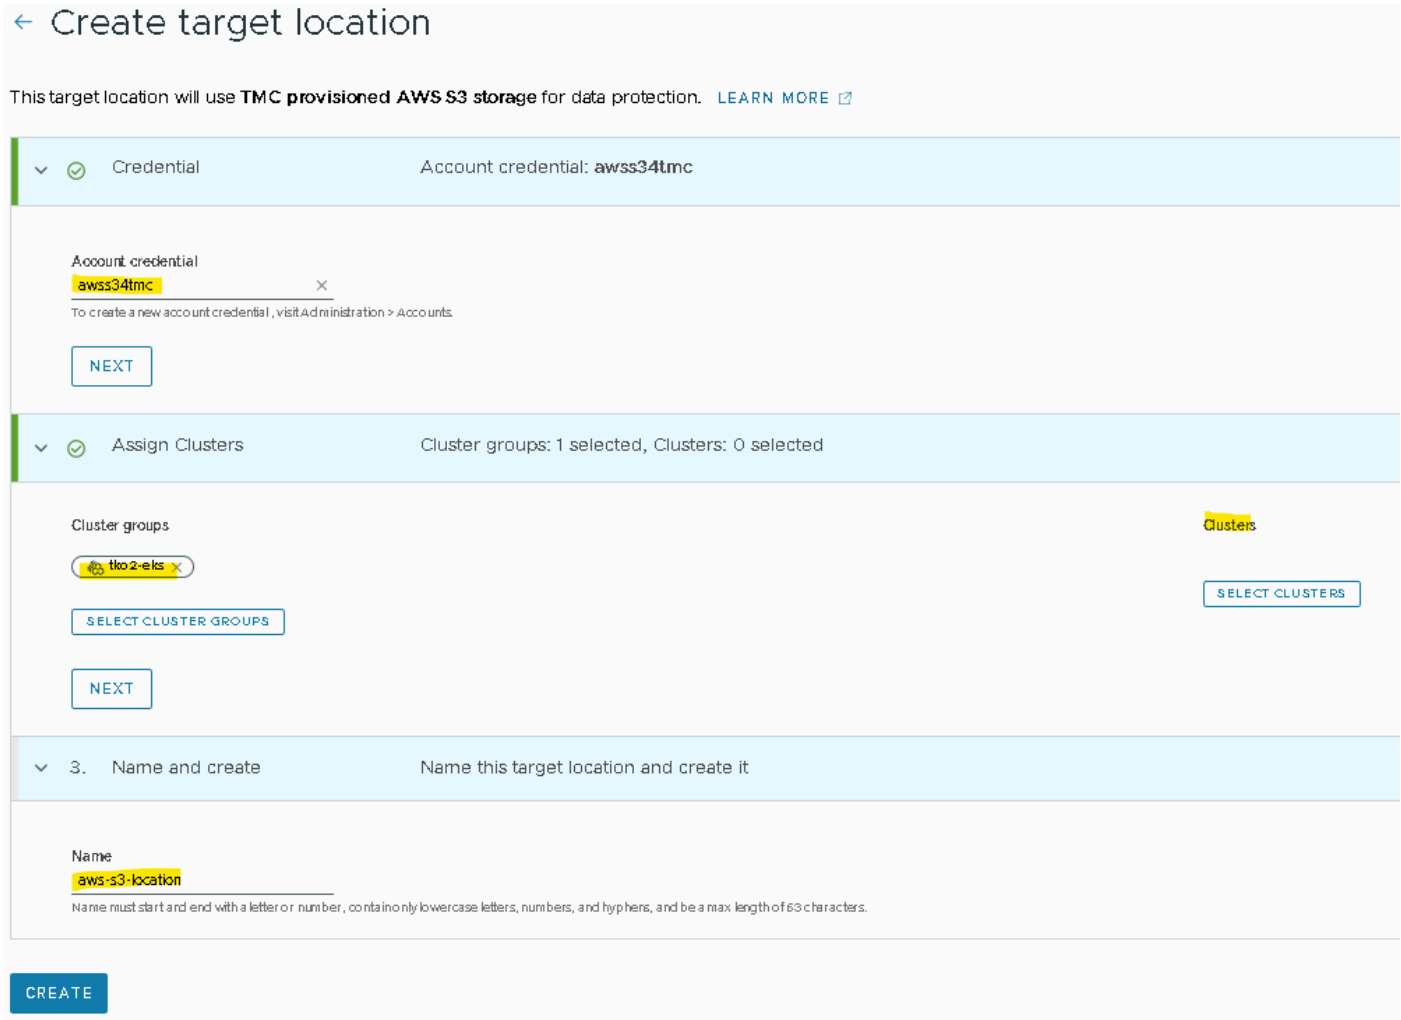

4. Enable data protection on the cluster. Creating required pods takes a few minutes to complete.

Figure 184: Enable Data protection on cluster.

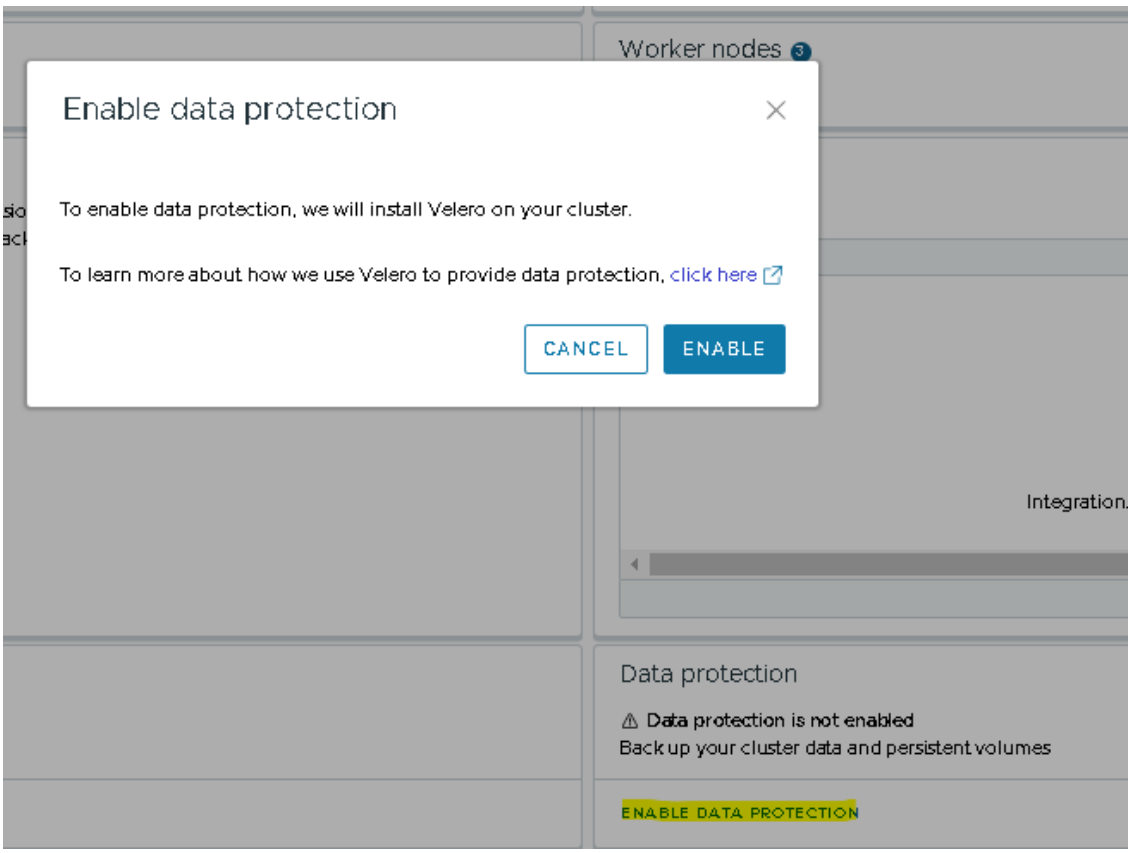

5. After data protection is enabled, create or schedule backup using the backup location created earlier.

Figure 185: Create backup.

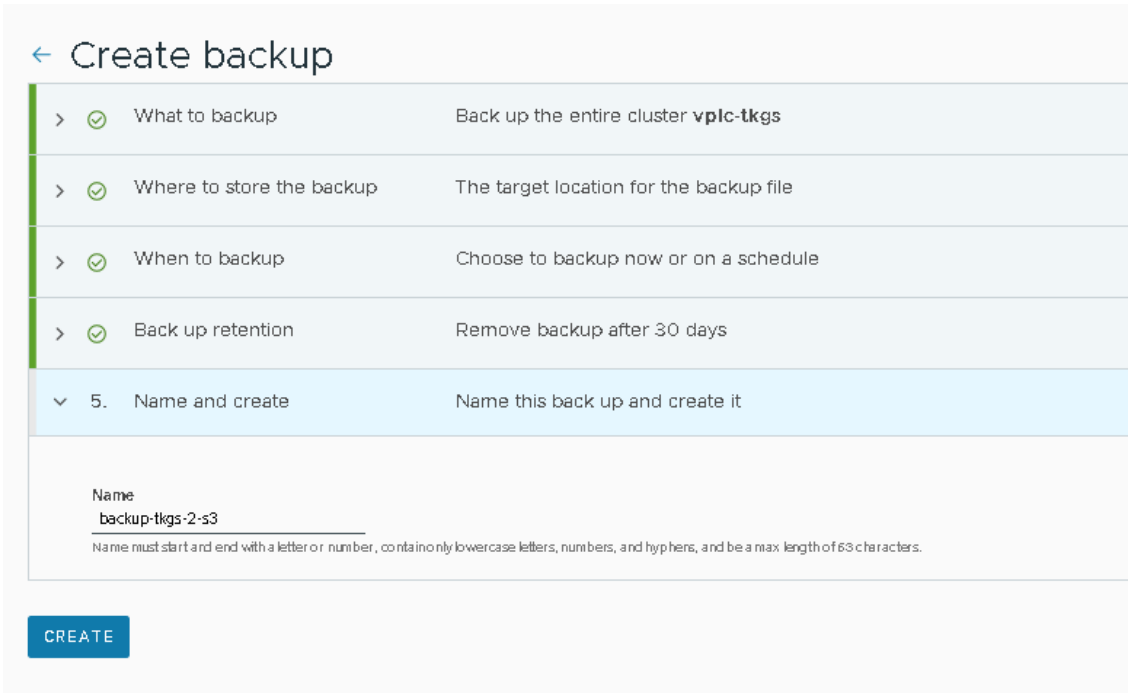

Once the backup is processed the status of the backup can be verified in Tanzu Mission Control user interface or viewing the S3 storage bucket.

Figure 186: Backup complete

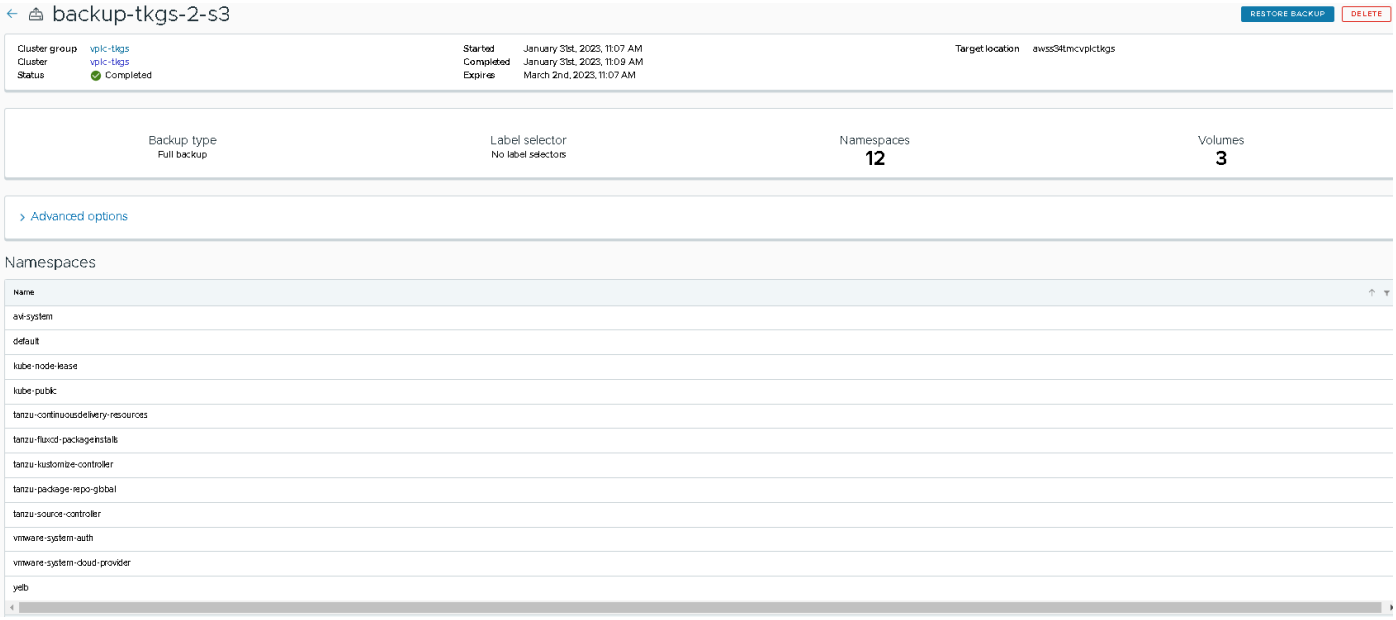

Figure 187: Amazon AWS bucket

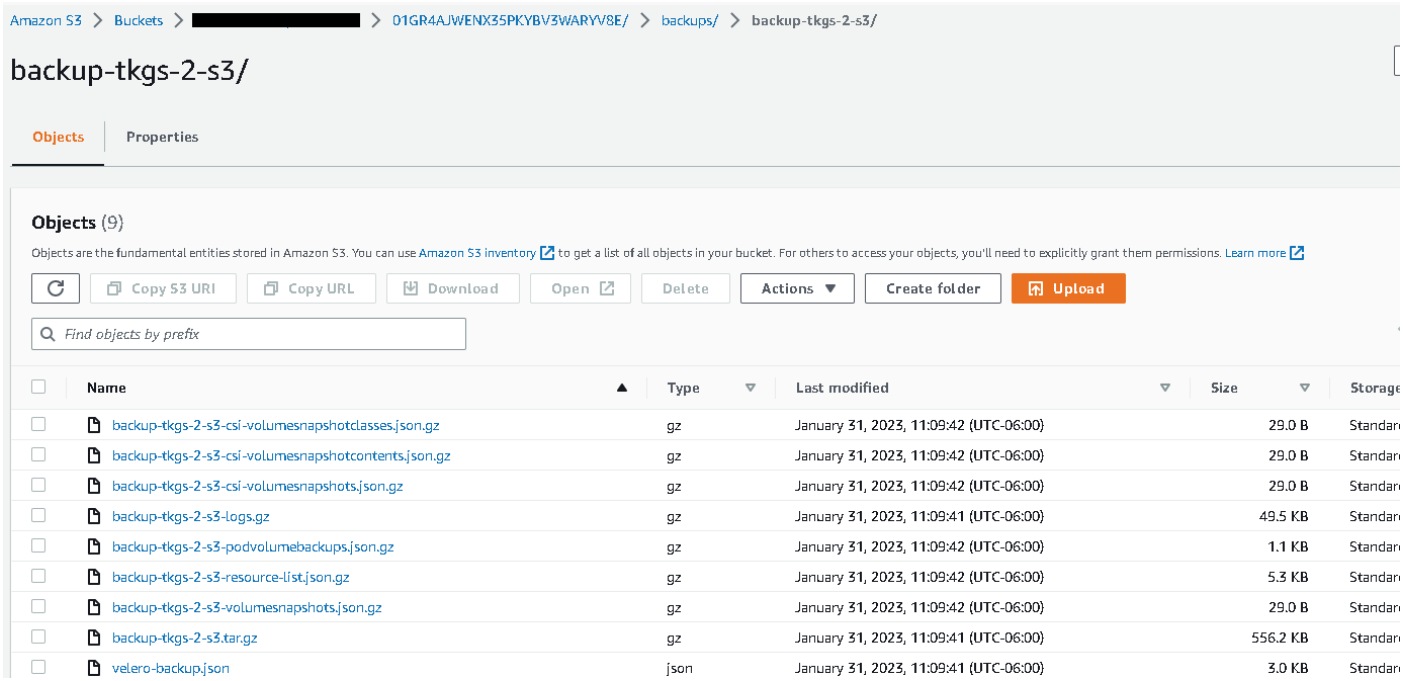

Cluster Inspections

Tanzu Mission Control Advanced edition provides preconfigured cluster inspections using Sonobuoy, an open-source community standard.

The following cluster inspections are available from the Overview and Inspection tabs of the cluster detail page in the Tanzu Mission Control console.

The Conformance inspection validates the binaries running on your cluster and ensures that your cluster is properly installed, configured, and working. Reports can be generated from within Tanzu Mission Control to assess and address any issues that arise. For more information, see the Kubernetes Conformance documentation at Kubernetes conformance.

The CIS benchmark inspection evaluates your cluster against the CIS Benchmark for Kubernetes published by the Center for Internet Security. This inspection type is available in the advanced version of Tanzu Mission Control.

The Lite inspection is a node conformance test that validates whether nodes meet requirements for Kubernetes. For more information, see Validate node setup in the Kubernetes documentation.

Because the cluster inspections provide a point-in-time report of the condition of the cluster, run them periodically (to avoid drifting out of conformance) and any time significant alterations are made, such as after patching or upgrading a cluster.

From the Inspections page in the Tanzu Mission Control console, a list of the most recent inspections that have been run against all the clusters in the organization, along with the results of those inspections. This page also allows you to start a new inspection.

Figure 188: Inspections

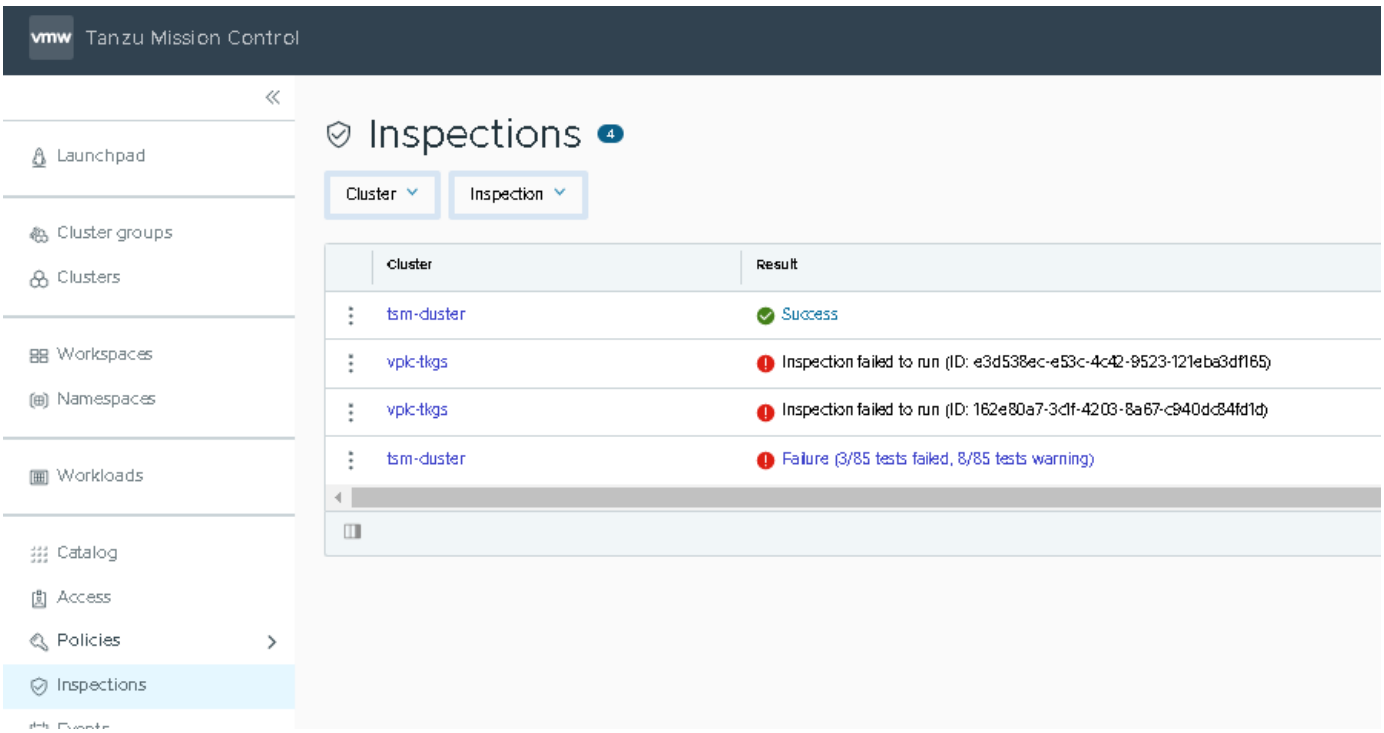

A list of inspections tests and results can then be viewed at the cluster level for details of the test runs.

Figure 189: Cluster inspection details.

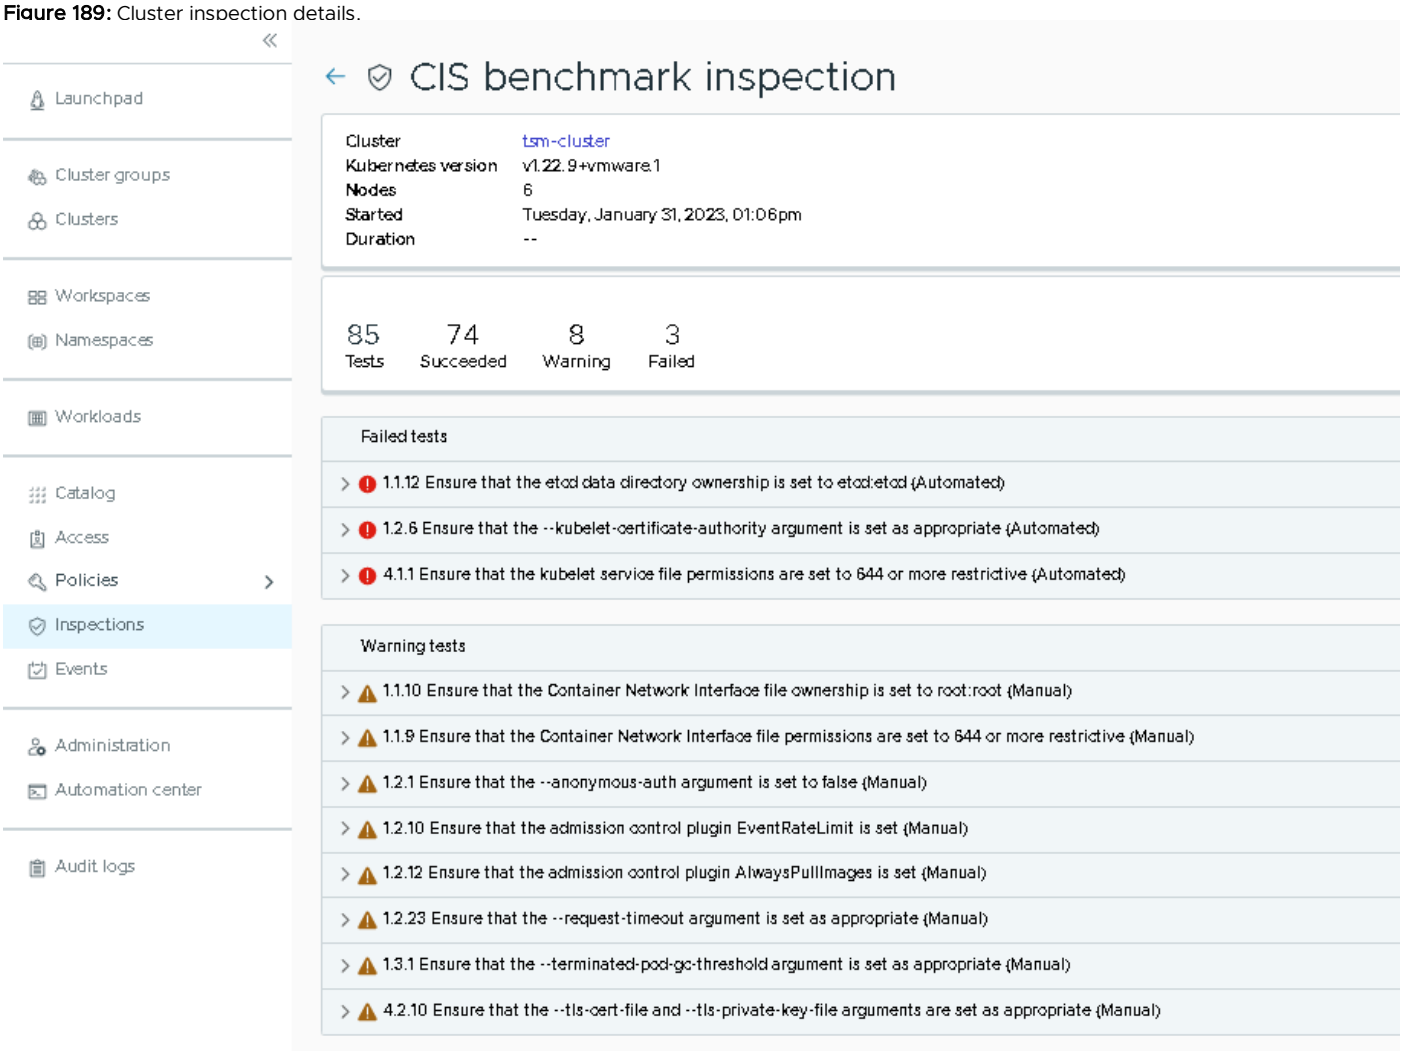

Cloud Native Applications

Cloud native is an approach to building and running applications that exploits the advantages of the cloud computing delivery model. When companies build and operate applications using a cloud native architecture, they bring new ideas to market faster and respond sooner to customer demands. While public cloud has affected the thinking about infrastructure investment in virtually every industry, cloud-like delivery is not exclusive to public environments. Cloud native development is appropriate for both public and private clouds; it is about how applications are created and deployed, not where.

To assess the workings of this multi-cloud reference architecture, a cloud-native application (Online Boutique) developed by Google was installed and scenarios depicted in figure 4 and figure 5 were tested. Online Boutique is a cloud-native microservices demo application. Online Boutique consists of an 11-tier microservices application. The application is a web-based e-commerce app where users can browse items, add them to the cart, and purchase them. Below is an example of a scenario that was evaluated based on Avi GSLB functionality discussed in this document.

Scenario: Application unavailability

With the architecture represented in the reference architecture, the application is installed on multiple sites and can be accessed via a single URL “boutique.avitko.pse.lab.” Which site the user connects to, depends on the load balancing algorithm selected. With Geo location option, the users will be directed to the nearest instance of the application, improving performance and user experience. In the example below, the application instances reside on on-premises private cloud and Amazon AWS EKS clusters. When a user initially connects to the application, he/she are directed to the

on-premises instance of the application since this is the closest Geo location to the user. This is indicated by a label “On-Prem” on the application frontend user interface.

Figure 190: User connection

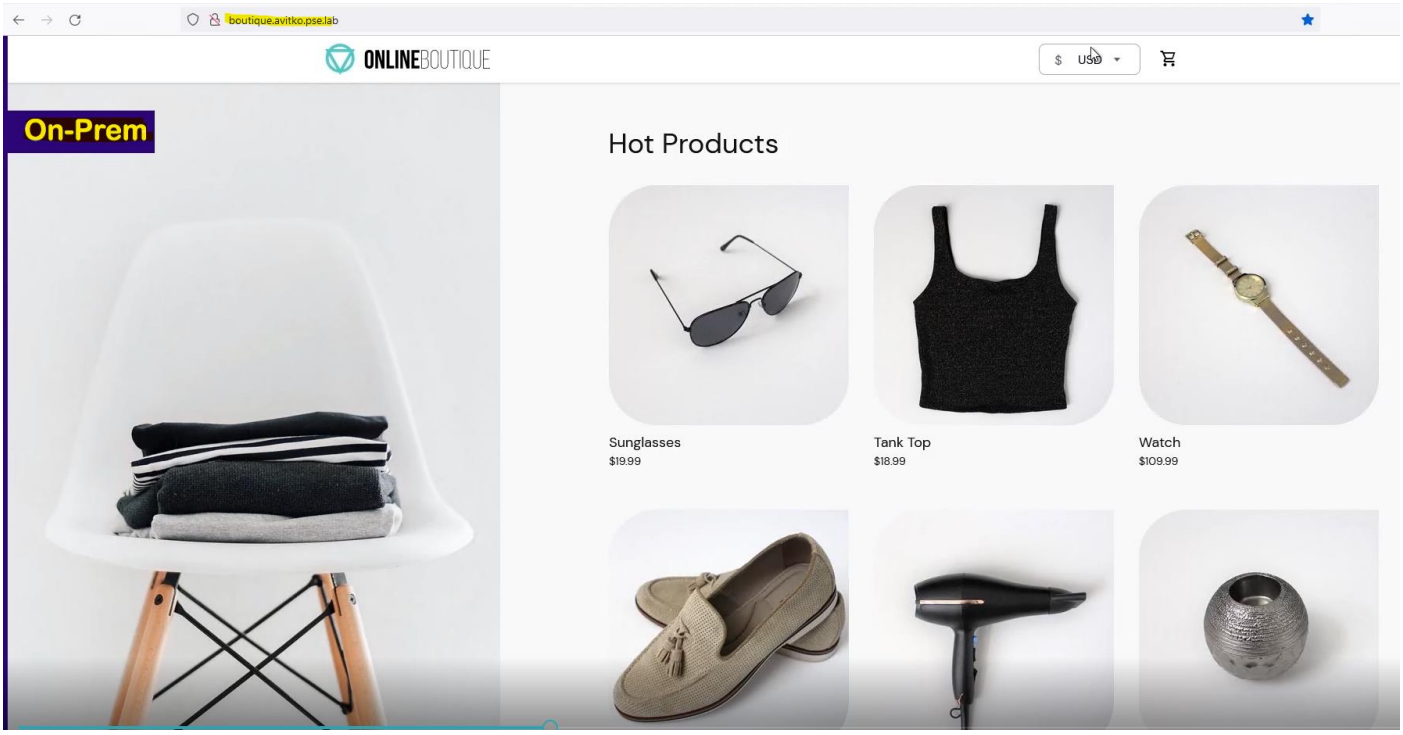

To simulate application unavailability or site failure, the application is uninstalled.

Figure 191: Application unavailable.

```
aj@tmc-clusters-cli:~/cli/boutique/microservices-demo/release$ kubectl delete -f kubernetes-manifests.yaml
deployment.apps "emailservice" deleted
service "emailservice" deleted
deployment.apps "checkoutservice" deleted
service "checkoutservice" deleted
deployment.apps "recommendationservice" deleted
service "recommendationservice" deleted
deployment.apps "frontend" deleted
service "frontend" deleted
service "frontend-external" deleted
deployment.apps "paymentservice" deleted
service "paymentservice" deleted
deployment.apps "productcatalogservice" deleted
service "productcatalogservice" deleted
deployment.apps "cartservice" deleted
service "cartservice" deleted
deployment.apps "loadgenerator" deleted
deployment.apps "currencyservice" deleted
service "currencyservice" deleted
deployment.apps "shippingservice" deleted
service "shippingservice" deleted
deployment.apps "redis-cart" deleted
service "redis-cart" deleted
deployment.apps "adservice" deleted
service "adservice" deleted
aj@tmc-clusters-cli:~/cli/boutique/microservices-demo/release$
```

At this point connectivity is lost to the application. When the user’s browser or client application retirees the connection, the user is redirected to the “AWS” site since the “On-Prem” instance is unavailable indicated by the “AWS” label on the frontend user interface.

Figure 192: Application available

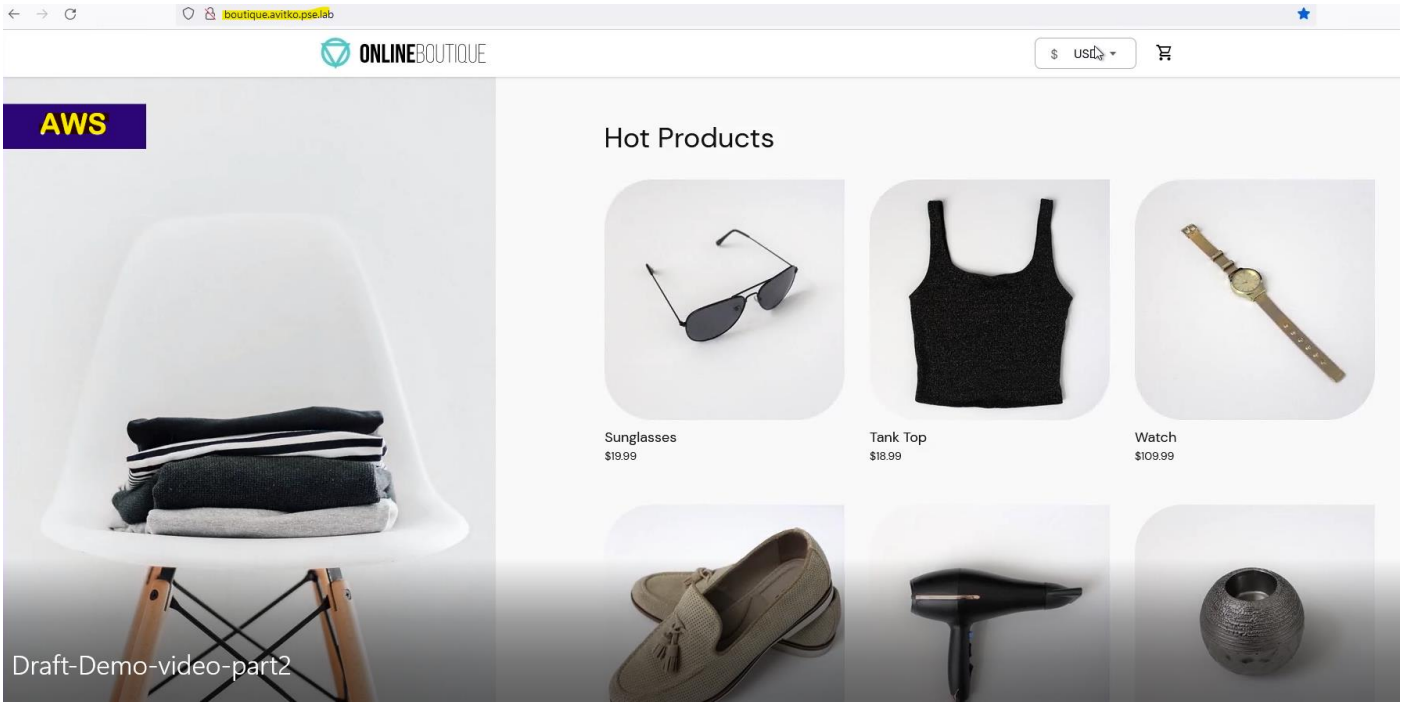

Distributed Microservices via Tanzu Service Mesh

VMware Tanzu Service Mesh is an enterprise-class service mesh solution that provides reliable control and security for microservices, end users, and data across all your clusters and clouds in the most demanding multi-cluster and multi-cloud environments.

To control application traffic, Tanzu Service Mesh provides fine-grained, traffic management policies that give you complete control and visibility into how traffic and API calls flow between your services and across clusters and clouds. To secure communication between services and protect sensitive data, you can use Tanzu Service Mesh to implement a zero-trust security model for cloud-based applications.

Note: Zero-Trust security, PII Data/user security and API security is only available with Enterprise Edition of Tanzu Service Mesh.

You can measure application performance with a configurable service level objective (SLO) definition. For more information, see the Service Level Objectives with Tanzu Service Mesh documentation. As application demands change, you can auto-scale services to maintain SLOs using Tanzu Service Mesh Service Auto-scaler. For more information, see the Service Autoscaling in Tanzu Service Mesh documentation .

Tanzu Service Mesh supports cross-cluster and cross-cloud use cases with global namespaces. With global namespaces, you can securely deploy applications across clusters and clouds and have consistent traffic management policies, application continuity, and security policies across cloud silos and boundaries, regardless of where the applications are running. Global namespaces can be each considered to mark an application boundary and as such provide strongly isolated environments for application teams and business units managing different applications and data.

What follows is a discussion of the major configuration areas of Tanzu Service Mesh as they relate to this reference architecture. The flow of configuration is as follows.

- Onboard Clusters
- Integrate NSX Advanced Load Balancer with Tanzu Service Mesh
- Add DNS and Domains
- Install Applications
- Create Global Namespace for applications.
- Create public service.
- Service Autoscaling
- Service Level Policies

Onboard Clusters

The first step is to onboard clusters that need to be part of Tanzu Service Mesh. This onboarding can be done either through Tanzu Service Mesh or if the cluster is already part of Tanzu Mission Control, the cluster can directly be onboarded from Tanzu Mission Control user interface.

If a proxy server is configured in your corporate environment, when onboarding your cluster, specify that it will connect to Tanzu Service Mesh through a proxy server. All traffic between the cluster and Tanzu Service Mesh will be routed through the proxy server and will be encrypted using Transport Layer Security (TLS). All requests that are sent from the cluster to Tanzu Service Mesh will be authorized using access tokens.

If the Avi controller is behind a private network and cannot be reached directly by the Tanzu Service Mesh global controller, a proxy connection is needed through one of the Kubernetes clusters available on Tanzu Service Mesh. A WebSockets proxy is implemented on all the client Kubernetes clusters onboarded into Tanzu Service Mesh for this purpose. In this way, Tanzu Service Mesh can connect to the Avi controller through the client cluster, which should have connectivity to the Avi controller as well. A **cluster label** must be assigned to the cluster to use the proxy.

1. To onboard a cluster from Tanzu Service Mesh user interface Click “New Workflow” and select “On-board Clusters

Figure 193: Onboard Cluster

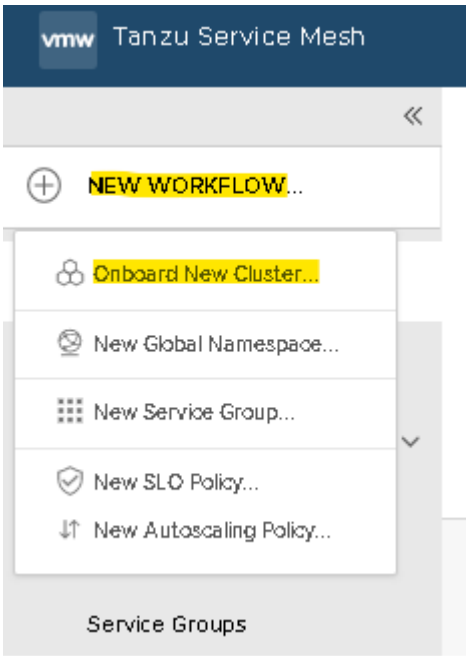

2. Give the cluster a name and create a “proxy location” cluster label and save.

Figure 194: Onboard cluster

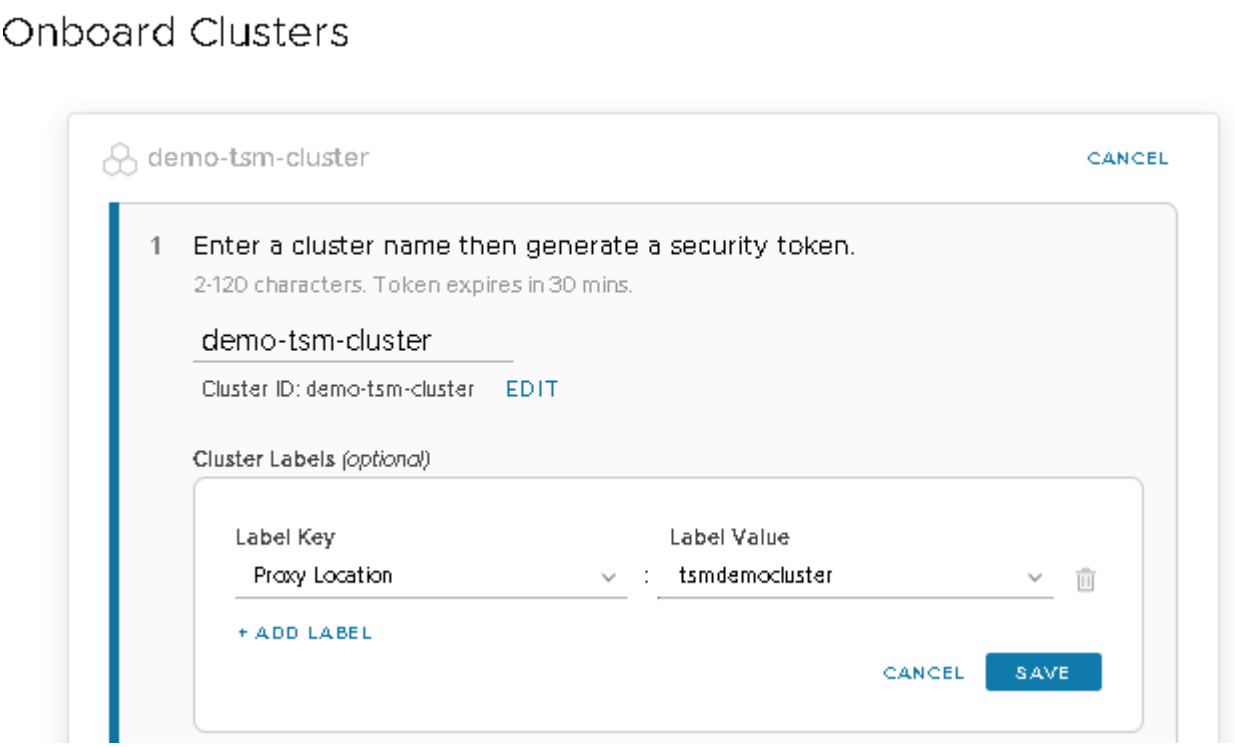

3. Click “Generate Security Token and copy the two commands to be run on the cluster.

Figure 195: Generate Tokens

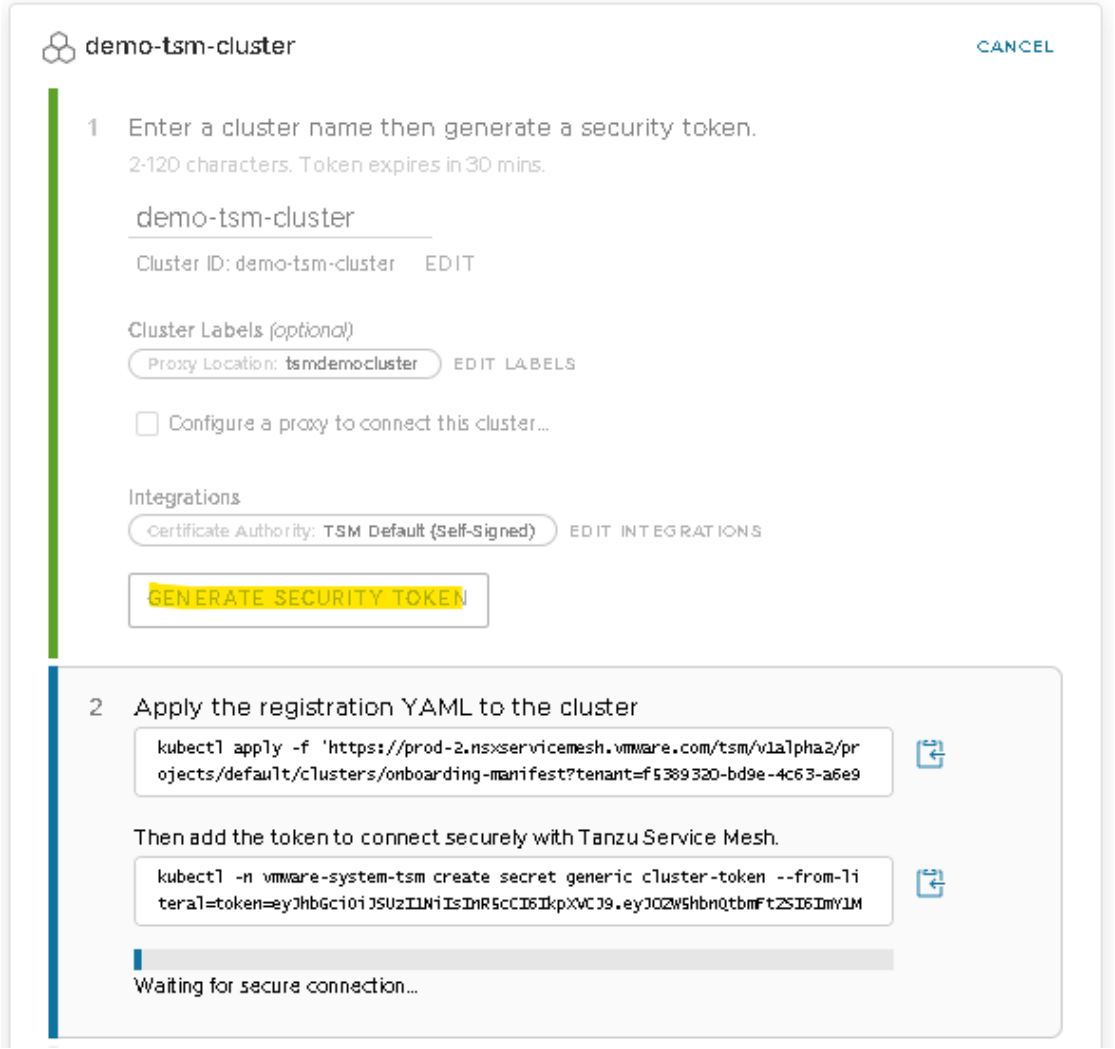

4. Login to the cluster and run the two copied commands.

Figure 196: Run onboarding commands.

```
aj@tkgs-cli:~$ kubectl apply -f  'https://prod-2.nsxservicemesh.vmware.com/tsm/v1alpha2/projects/d
efault/clusters/onboarding-manifest?tenant=f5389320-bd9e-4c63-a6e9-82422c09bb27'

namespace/vmware-system-tsm created
customresourcedefinition.apiextensions.k8s.io/aspclusters.allspark.vmware.com created
customresourcedefinition.apiextensions.k8s.io/clusters.client.cluster.tsm.tanzu.vmware.com created
customresourcedefinition.apiextensions.k8s.io/tsmclusters.tsm.vmware.com created
customresourcedefinition.apiextensions.k8s.io/clusterhealths.client.cluster.tsm.tanzu.vmware.com c
reated
configmap/tsm-agent-operator created
serviceaccount/tsm-agent-operator-deployer created
clusterrole.rbac.authorization.k8s.io/tsm-agent-operator-cluster-role created
role.rbac.authorization.k8s.io/vmware-system-tsm-namespace-admin-role created
clusterrolebinding.rbac.authorization.k8s.io/tsm-agent-operator-crb created
rolebinding.rbac.authorization.k8s.io/tsm-agent-operator-rb created
deployment.apps/tsm-agent-operator created
job.batch/update-scc-job created
aj@tkgs-cli:~$
aj@tkgs-cli:~$ kubectl -n vmware-system-tsm create secret generic cluster-token --from-literal=token=
yNDIyYzA5YmIyNyIsImNsdXN0ZXItbmFtZSI6ImRlbW8tdHNtLWNsdXN0ZXIiLCJlcmkiOiJwcm9kLTItaW50ZXJuYWwubnN4c2Vy
ODkzMjAtYmQ5ZS00YzYzLWE2ZTktODI0MjJjMDliYjI3IiwiYWxsc3BhcmstdmVyc2lvbiI6InYxLjE2LjkiLCJ0c20tdmVyc2lvb
mVmcmVzaClpbnRlcnZhbCI6MCwibTctcGFja2FnZSllcmkiOiIiLCJpYXQiOjE2NzYzODc0MTgsImlzcyI6InRzbSIsInNlYiI6In
XaoAKEvLQon--TO-P02QDjG_wmyt0UxTKhqzAX5cwjNYWTid5Ol9AclAjV5snENY3UjmvUD9IGCnjox7qS2NC7O0uSU3PJlf3S9cC
zcqrNfLV4B5FMM1-joL9TiwjtdNxh-P7jIKJ-h5MQIhvzCnK2AJe979O41D3Cv7kr-w4o-kpm0fDA-IwGObwCdQ59MWsWtjL0d2_s
secret/cluster-token created
aj@tkgs-cli:~$
```

5. Choose to install Tanzu Service Mesh on all namespaces or select namespaces that should be excluded. Click "Install Tanzu Service Mesh."

Figure 197: Install Tanzu Service Mesh

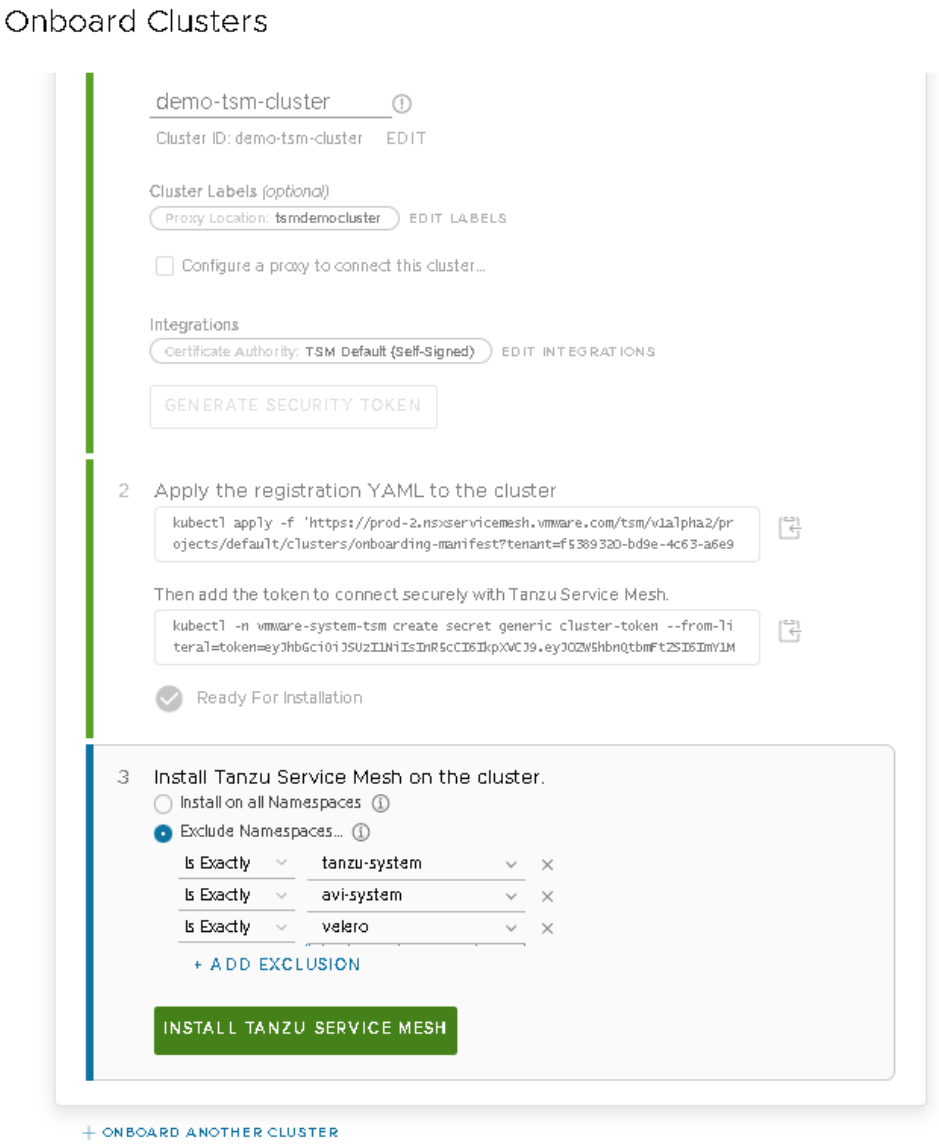

Note: The system namespaces on the cluster, such as kube-system, kube-public, and istio-system, are excluded from Tanzu Service Mesh by default.

6. Verify that all pods under namespaces *"vmware-system-tsm"* and *"istio-system"* are running.

Figure 198: Verify installation.

```
aj@tkgs-cli:~$ k get pods -n vmware-system-tsm
NAME                                    READY   STATUS      RESTARTS   AGE
allspark-ws-proxy-6cf5859b4c-bwcvd      1/1     Running     0          15m
config-service-7bd95cc764-6f8qd         1/1     Running     0          15m
installer-job--1-v8gqc                  0/1     Completed   0          4m3s
k8s-cluster-manager-78f9fb676d-crk6h    1/1     Running     0          15m
telegraf-istio-code-7c584f4797-hjwp2    1/1     Running     0          3m4s
telegraf-istio-count-f989cd6bc-2gb4r    1/1     Running     0          3m4s
telegraf-istio-duration-6d6795c485-f65nl   1/1     Running     0          3m4s
telegraf-nodes-66cb96f9b6-5ps2k         1/1     Running     0          3m3s
tsm-agent-operator-d866f497b-hxl4n      1/1     Running     0          15m
update-scc-job--1-csnx9                 0/1     Completed   0          15m
aj@tkgs-cli:~$ k get pods -n istio-system
NAME                                    READY   STATUS    RESTARTS   AGE
allspark-telegraf-node-7s8kq            1/1     Running   0          4m10s
allspark-telegraf-node-b48q2            1/1     Running   0          4m10s
allspark-telegraf-node-j5kd9            1/1     Running   0          4m10s
allspark-telegraf-node-mckst            1/1     Running   0          4m10s
istio-egressgateway-55bd7974c4-4d6q4    1/1     Running   0          4m36s
istio-egressgateway-55bd7974c4-fqdlk    1/1     Running   0          4m36s
istio-ingressgateway-945bb5ff-pf8wf     1/1     Running   0          4m36s
istio-ingressgateway-945bb5ff-rznlx     1/1     Running   0          4m36s
istio-telemetry-6df84d554d-4z4zt        2/2     Running   0          4m14s
istio-telemetry-6df84d554d-vwvbs        2/2     Running   0          4m14s
istiod-795596f849-tr6x6                 1/1     Running   0          4m44s
istiod-795596f849-twg2g                 1/1     Running   0          4m44s
aj@tkgs-cli:~$
```

Integrate NSX Advanced Load Balancer with Tanzu Service Mesh

In this reference architecture, NSX Advanced Load Balancer plays a central role in connecting sites, clusters and services and providing load balancing and ingress. For Tanzu Service Mesh to provide central management and GSLB, NSX Advanced Load Balancer leader controller needs to be integrated with Tanzu Service Mesh. Applications are exposed to users through a public service configured in a global namespace. NSX Advanced Load Balancer routes user requests to optimal application instances by using the global load balancing configuration specified for the public service. Following procedure outlines the integration process.

1. From Tanzu Service Mesh left navigation pane click Tanzu Admin > Integration
2. On the **Integrations** page, under **All Integrations**, find the Avi card with the **DNS** and **GSLB** labels.

Figure 199: Integration

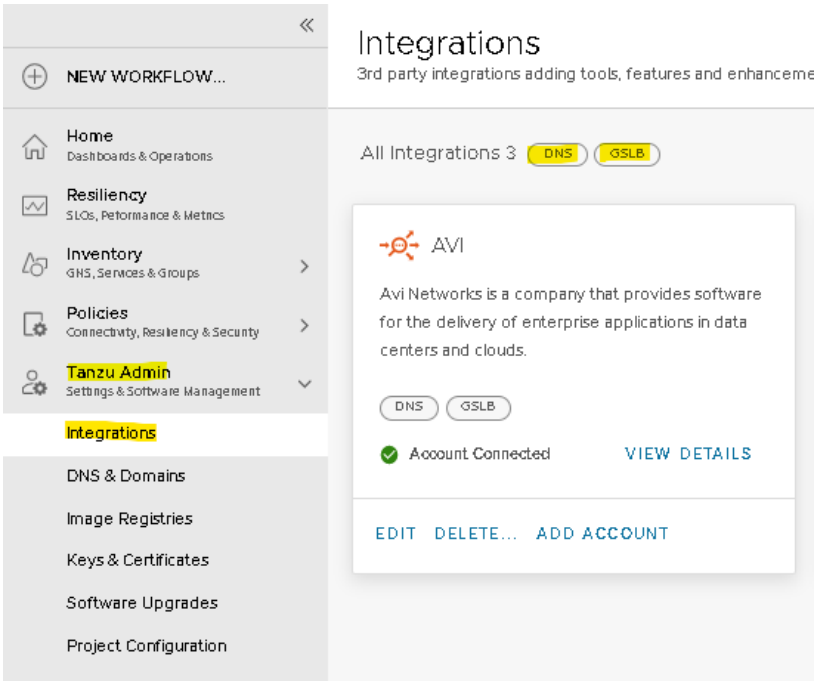

3. Select one of the following options.
    a. If you are creating the first Avi integration account, at the bottom of the card, click **Configure**.

b. If one or more Avi integration accounts exist and you are creating another account, at the bottom of the card, click **Add Account**.

4. Enter required information for the environment including the proxy location created while onboarding the cluster.

Figure 200: NSX Advanced Load Balancer Integration

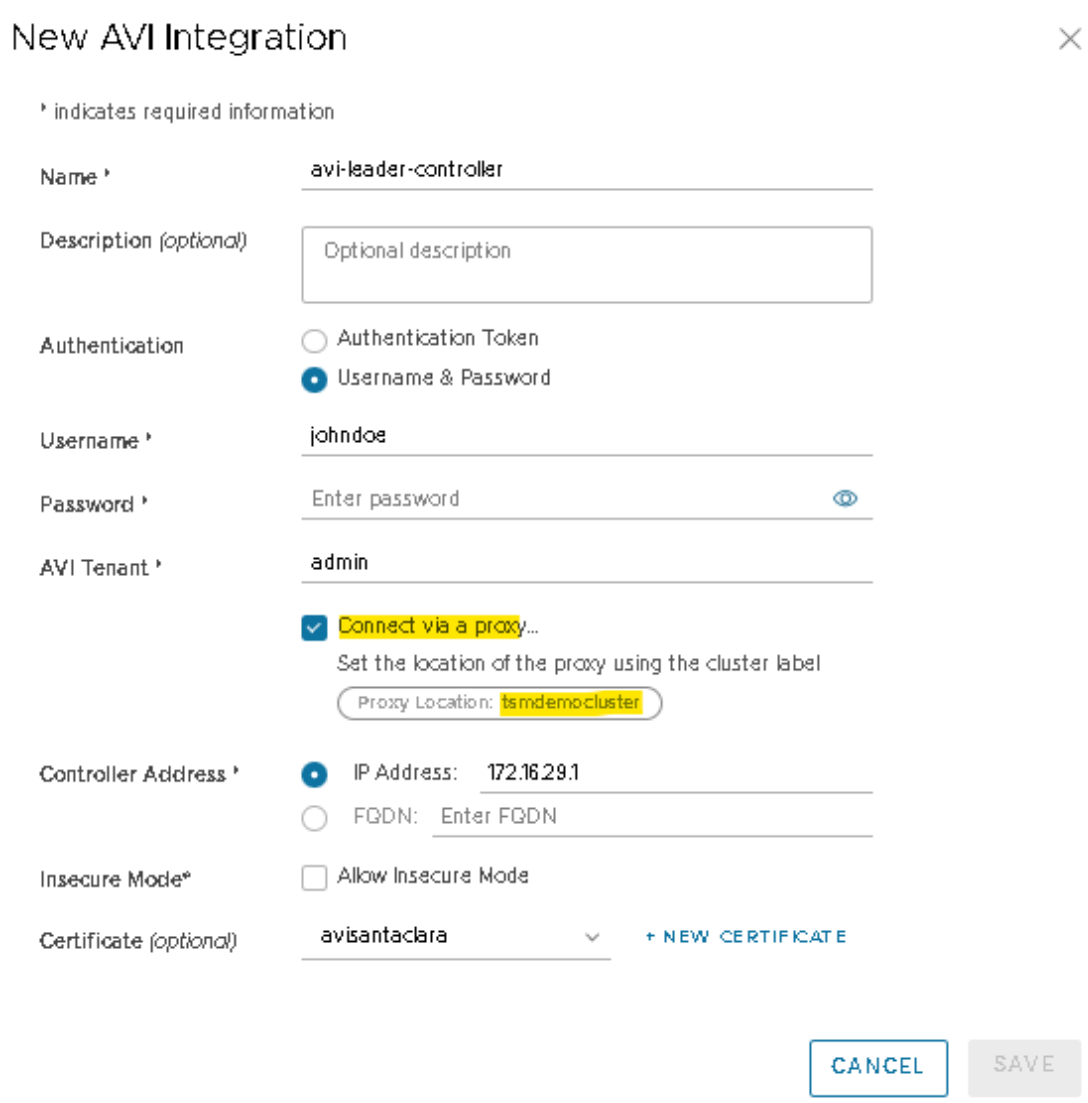

Note: It is recommended to use certificate from NSX Advanced Load Balancer leader controller instead on "*insecure mode*." Insecure mode allows to still use TLS, but do not require globally trusted certificates.

Add DNS and Domains

Once the integration has been created, add a DNS domain the activate it. NSX Advanced Load Balancer account status will still show disconnected (Red) until the DNS and Domain is linked to the integration.

1. Click Tanzu Admin and select DNS and Domains. In the New DNS Account dialog box, select the name of the Avi integration account created in previous step as Domain Provider.

Figure 201: Add domain.

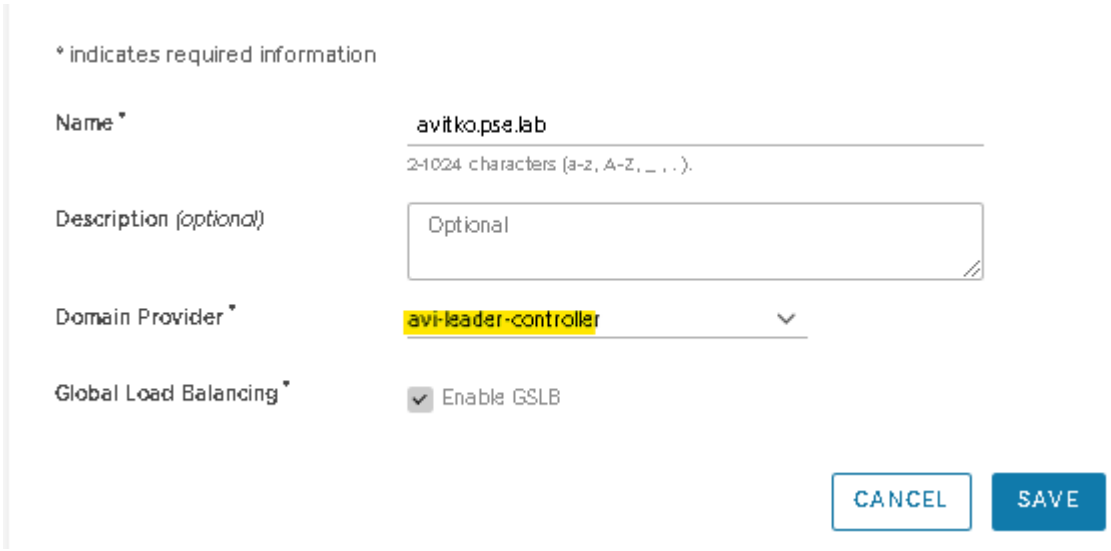

2. Once done, go back to Integrations page and verify that the Avi integration is green.

Figure 202: Integration complete

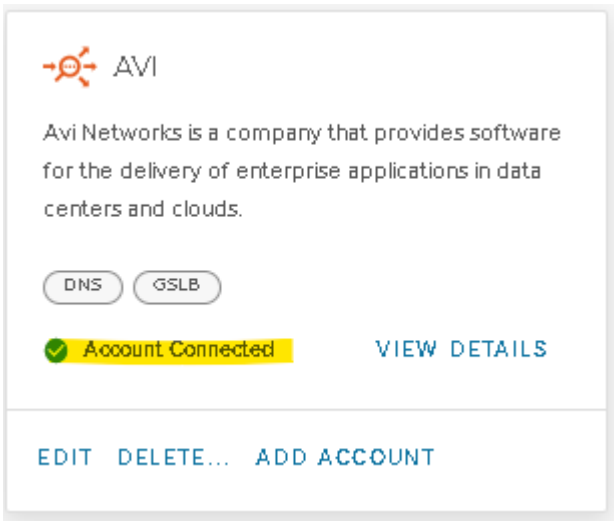

Installing Applications

To enable end-user access to frontend services outside the global namespace, Tanzu Service Mesh exposes the service via "public service" construct. A similar construct named "External service" is also implemented. External services are services that exist outside the VMware Tanzu Service Mesh (for example, third-party database services) but are made accessible by services within a global namespace of the VMware Tanzu Service Mesh. Services can run on virtual machines, external Kubernetes clusters, Tanzu Application Service environments, lambda functions or even on bare metal, and can be accessed over TCP, TLS, HTTP, or HTTPS.

Note: When mapping services to the global namespace where you want to configure a public service for GSLB, verify that all the services reside in namespaces with the same name. This is required by the current version of Tanzu Service Mesh and will change in the future.

Note: If Avi Kubernetes Operator (AKO) is installed on the onboarded clusters where instances of the public service will be deployed, deactivate the L4Settings.autoFQDN configuration setting during installation. If this setting is not deactivated, Tanzu Service Mesh will try to resolve the ingress gateway using the local FQDN rather than the external IP

address, which will only work if the resolvers on the nodes point to Avi DNS. For information about the L4Settings.autoFQDN setting, see the AKO documentation on GitHub.

To enable the cross-cluster communication between the services, application manifest for Kubernetes deployment may have to be edited for the appropriate service on one cluster, to specify the domain name of the global namespace and prefixing the domain name with the name of the service on the other cluster.

For example, if a service called 'frontend' on one cluster needs to communicate to another service called 'cart' on another cluster, the deployment manifests of frontend service on the first cluster needs to be edited to set appropriate variable to "cart.avitko.pse.lab," for example. The "cart" prefix is required for "frontend" service to communicate with "cart" service. In addition, service protocol type and port need to be added, if it does not exist in the application manifests. Example below.

```
apiVersion: v1
kind: Service
metadata:
  name: order-svc
spec:
  ports:
    - appProtocol: http
      number: 3001
```

Create Global Namespace for applications.

With global namespaces in Tanzu Service Mesh, can connect and secure the services across clusters. A global namespace map, discover and connect services automatically across clusters. A global namespace can be shared across a single cluster, multiple clusters, or even clusters in different clouds. Below is a high-level process of creating a global namespace.

1. In the navigation panel on the left, click New Workflow and then click New Global Namespace. On the General Details page of the New Global Namespace wizard, enter a unique name and a domain name for the global namespace.

Figure 203: Global namespace

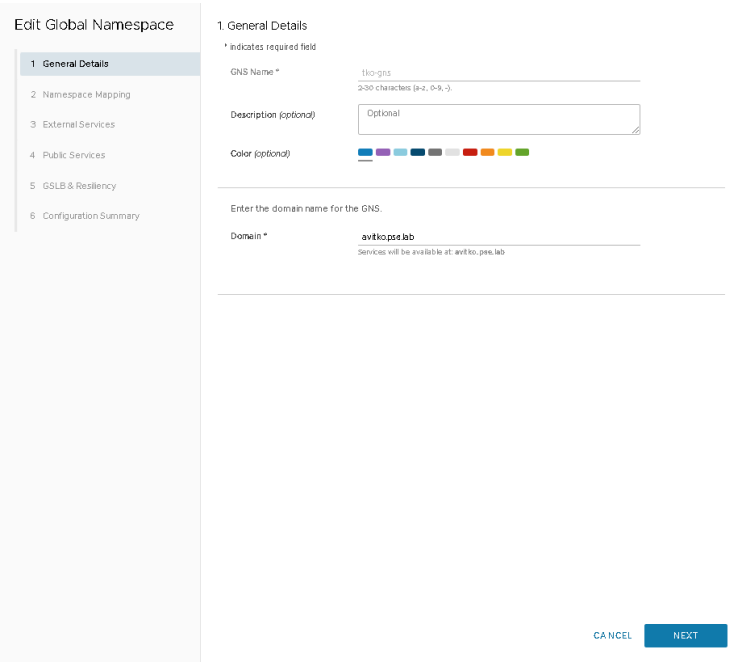

2. On the Namespace Mapping page, to add the services in your application to the global namespace, specify their Kubernetes namespace-cluster pairs. Under Namespace Mapping Rule, in the left drop-down menu, select the namespace on one of your clusters that holds some of the services and in the right drop-down menu, select the

name of the cluster. Click Add Mapping Rule to create multiple namespace mapping rules for the same or different clusters.

Figure 204: Namespace mappings

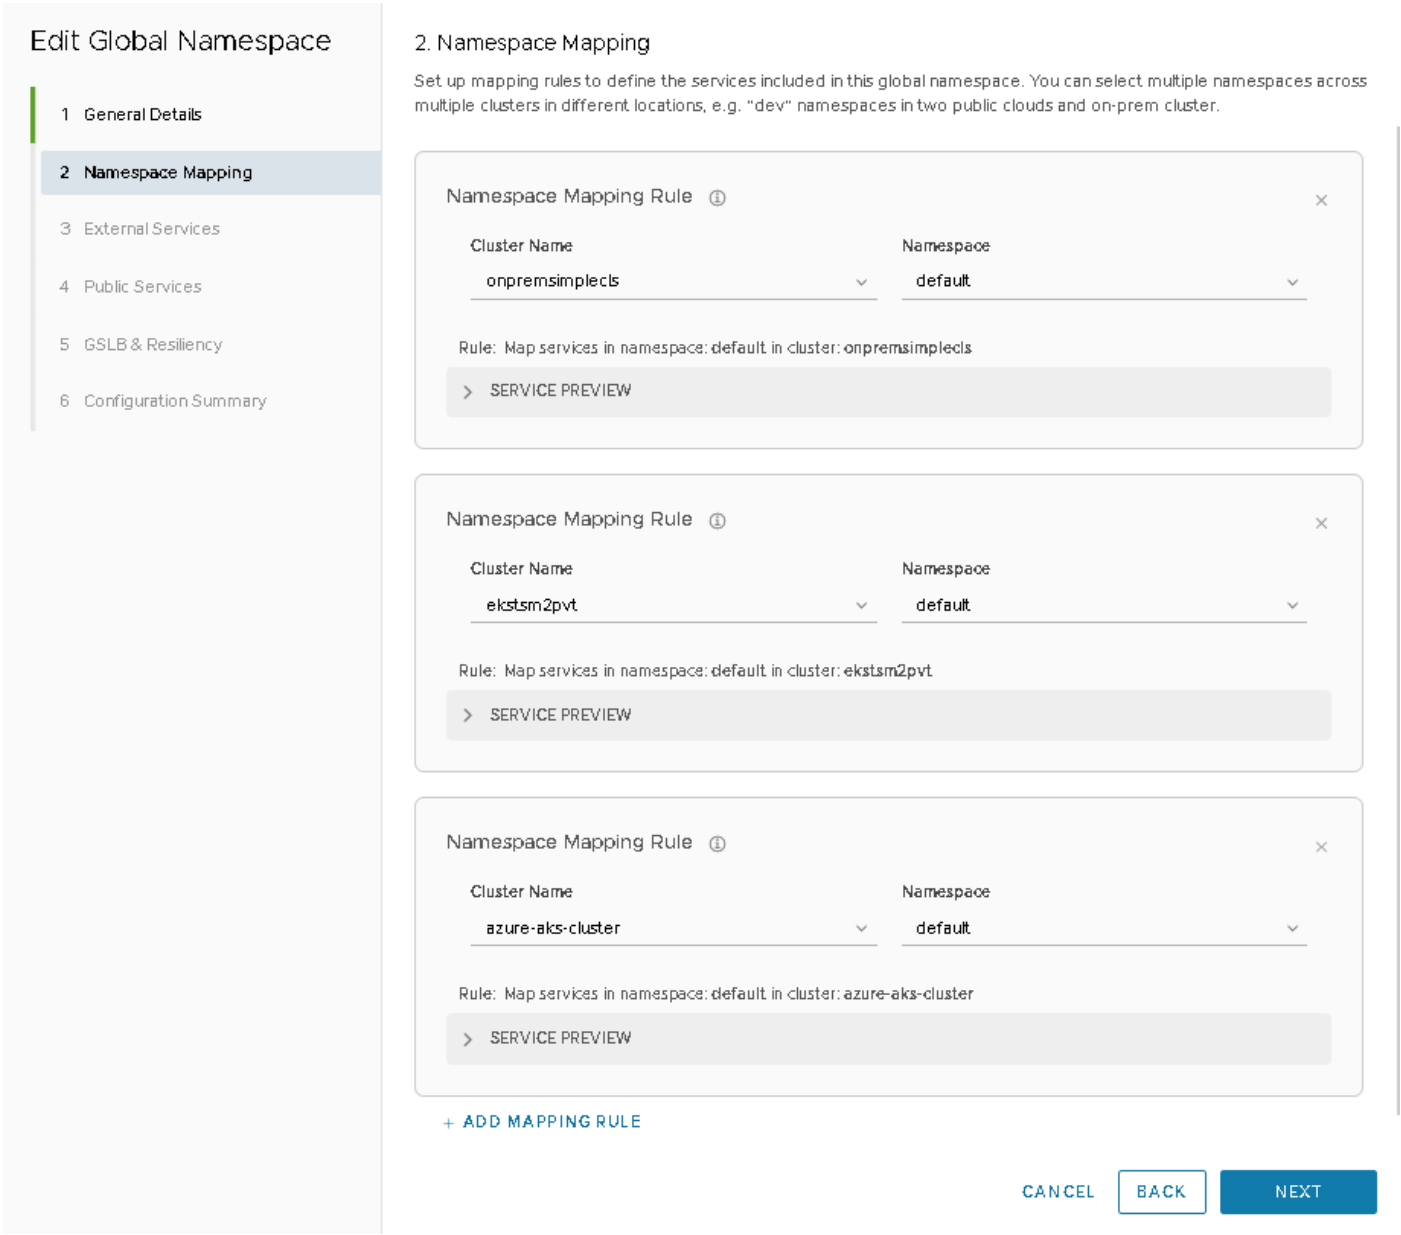

3. To configure external services in the global namespace, click Add Public Service(s), provide the configuration for each public service, and click Next.

Note: No external service was configured for this reference architecture.

Figure 205: External Service

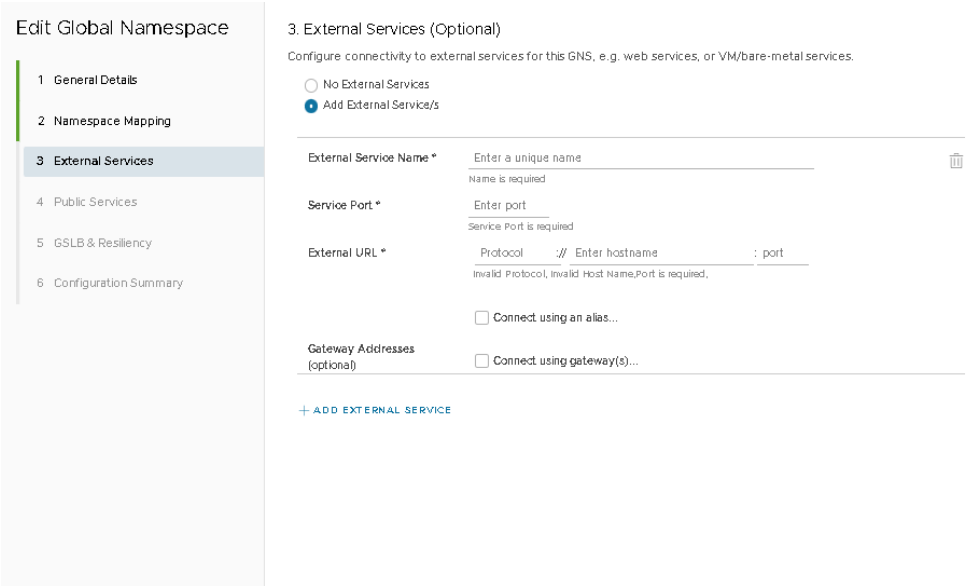

4. To configure public services in the global namespace, click Configure Public Service(s), provide the configuration for each public service, and click Next.

Figure 206: Public service

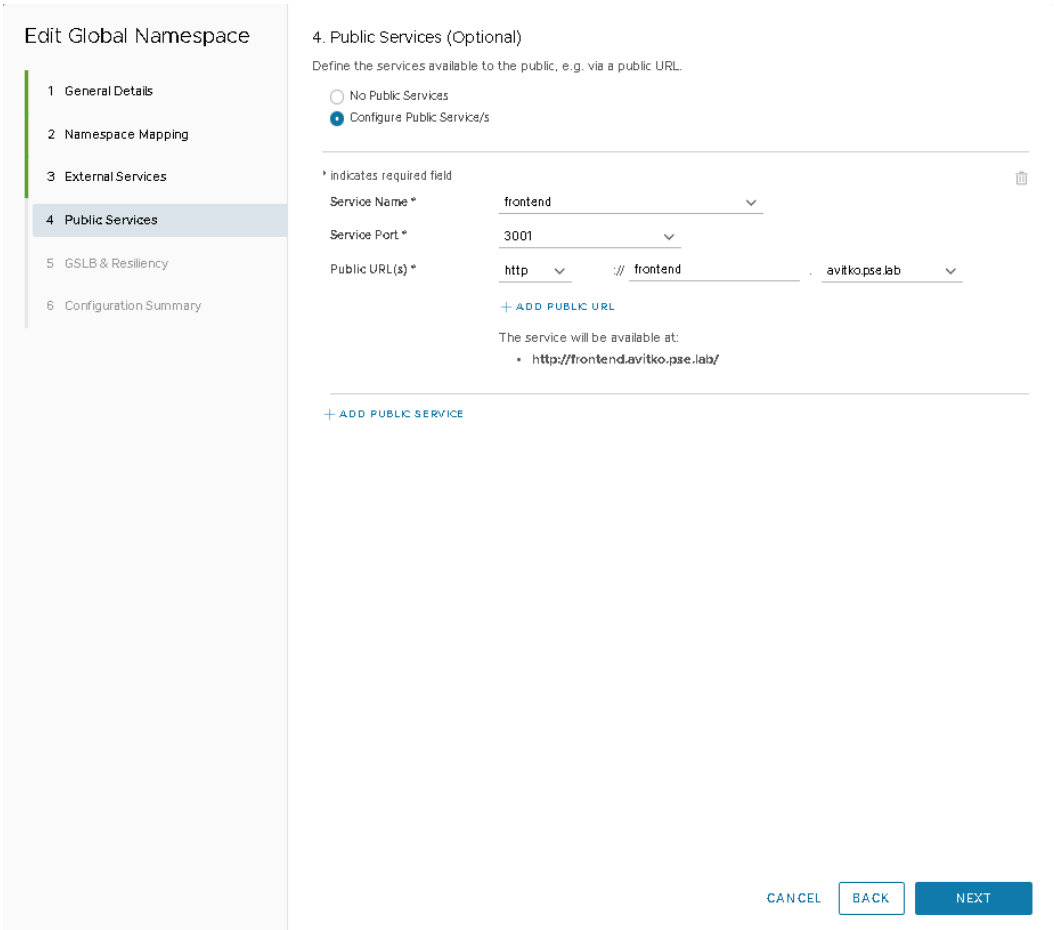

5. On the GLSB & Resiliency page, configure global load balancing scheme for the public services and click Next.

Figure 207: load balancing scheme

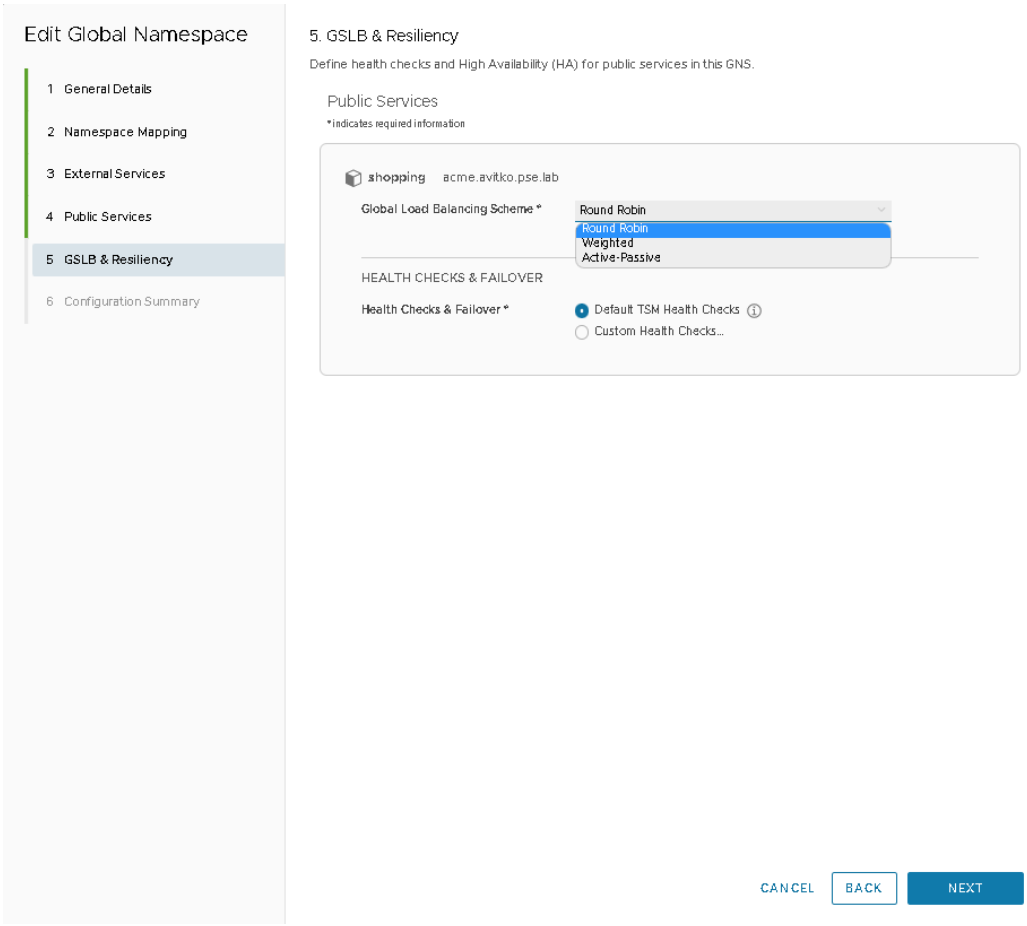

6. On the Configuration Summary page, review the configuration of the global namespace and click Finish.

Service autoscaling policies

Autoscaling represents the ability of a service to automatically scale up or down to efficiently handle changes of the service demand. With Tanzu Service Mesh Service Autoscaler, developers and operators can have automatic scaling of microservices that meet changing levels of demand based on metrics, such as CPU or memory usage. These metrics are available to Tanzu Service Mesh without needing additional code changes or metrics plugins.

Tanzu Service Mesh Autoscaler supports configuring an autoscaling policy for services inside a global namespace through the UI as well as API. For more information, see Configure GNS-Scoped Autoscaling Policy Using Tanzu Service Mesh UI. Tanzu Service Mesh Autoscaler also provides a Kubernetes Custom Resource Definition to configure autoscaling for services directly in cluster namespaces. For more information, see Deploying the Tanzu Service Mesh Service Autoscaler Through CRD. This approach for configuring autoscaling with CRD is available only for org-scoped autoscaling policies. Once an autoscaling policy is configured, Tanzu Service Mesh starts to monitor the Autoscaler metric for the service and scales the service accordingly.

Below is an example of configuring autoscaling a front-end service called “shopping” based on service CPU usage. Figure 208 shows three clusters as part of the global namespace with application services distributed across the three cloud instances. The catalog service is on Amazon EKS cluster. There are two instances of the frontend “shopping” service, one on Azure AKS and one on on-premises private cloud. On-premises private cloud also holds the rest of the application services.

Figure 208: Distributed application

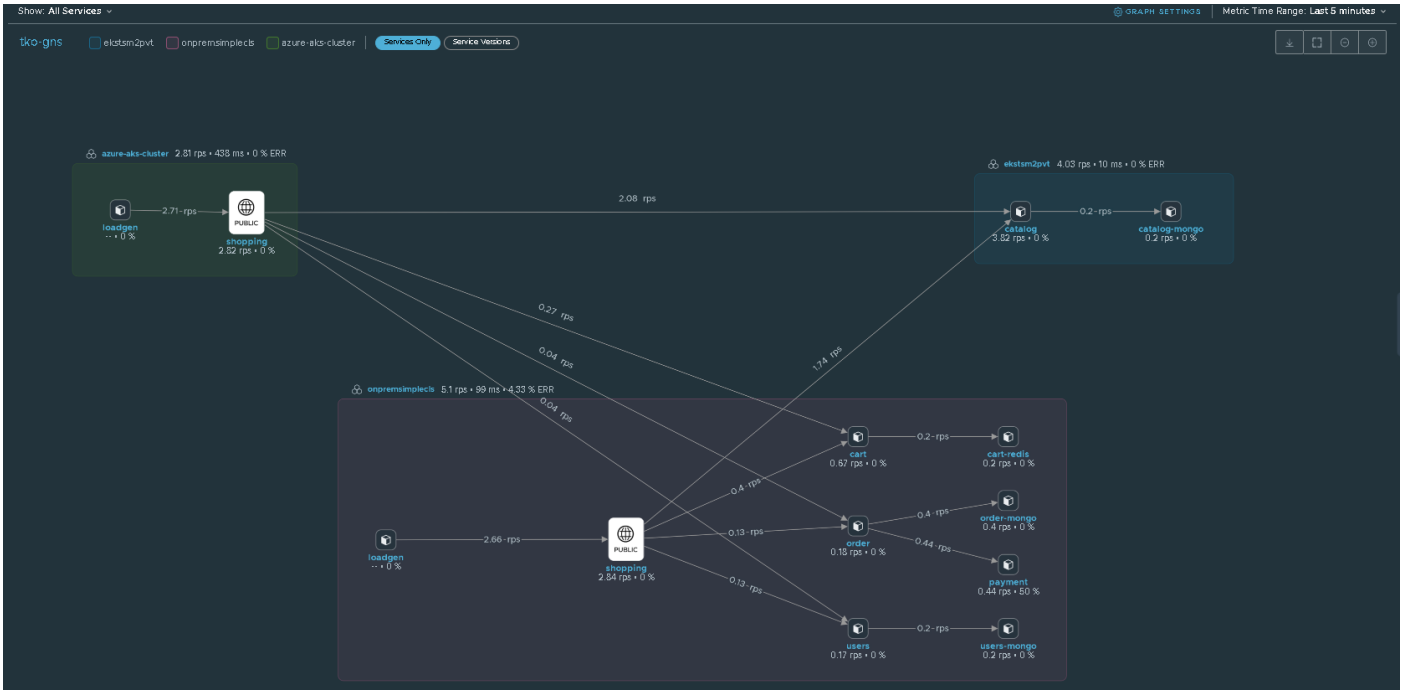

To ensure that the frontend service where the users connect to the applications, performs efficiently, an autoscaling policy will be created to scale shopping service up or down depending on CPU usage of the microservice. A load generator will be used to simulated user traffic by generating load on the frontend service. Figure 209 shows the current CPU usage for both services prior to load generation.

Figure 209: CPU usage

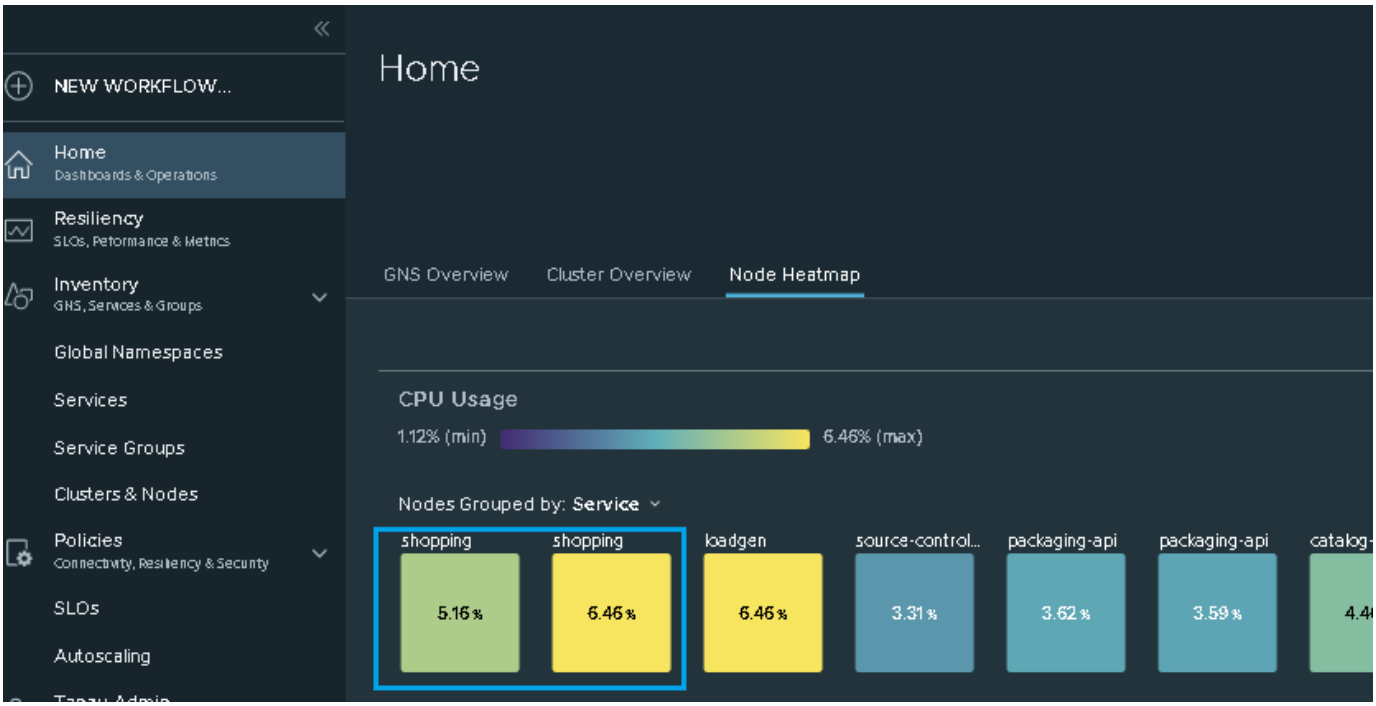

There are three instances of the application currently running as shown below.

Figure 210: Service instances

To simulate service scale-up, an Autoscaling policy is created show below in figure 211.

Figure 211: Simulate service scale-up.

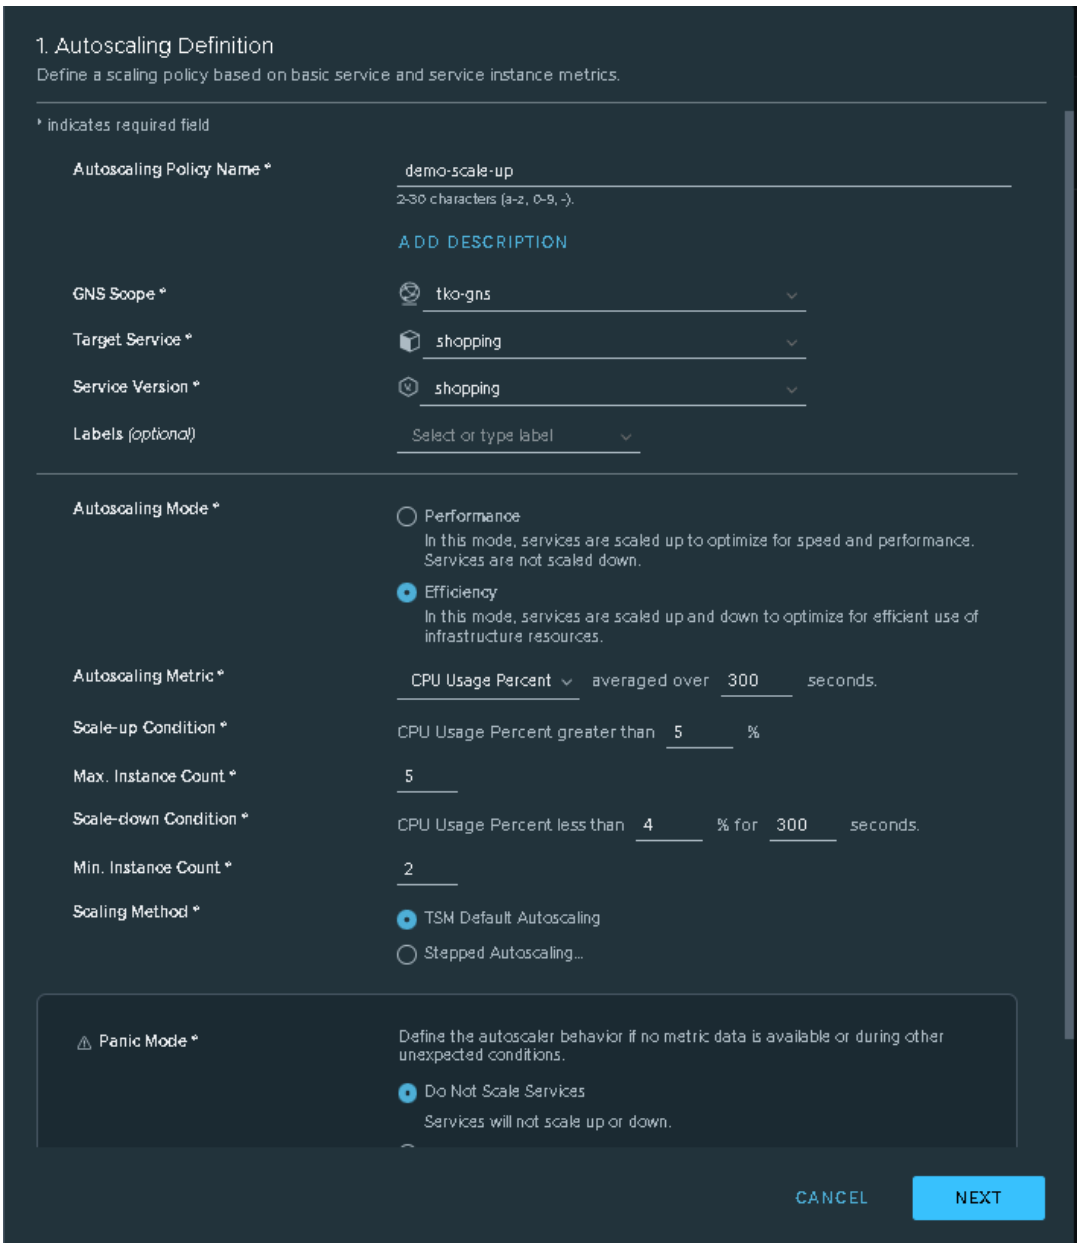

After 5 minutes interval set in the policy, where the CPU usage was above the 5% threshold the on-premises and Azure AKS shopping services are scaled up as shown in figure 212 below.

Figure 212: Services scaled up.

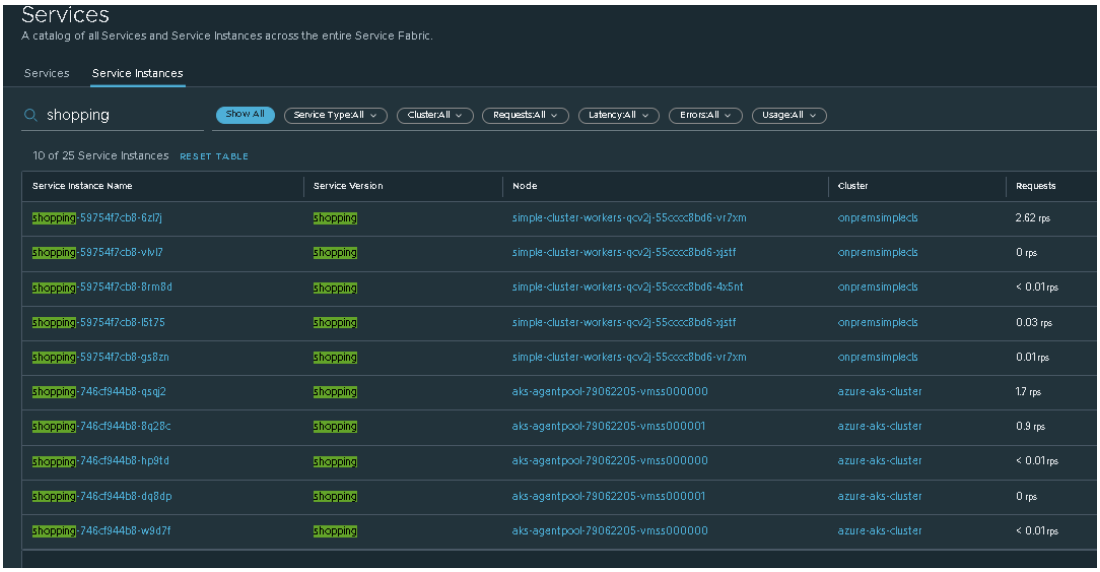

Note: The above is only an example of CPU usage and scaling policy setting accordingly. CPU usage and corresponding policy settings will vary as requirements changes.

Service Level Object Policy

Service level objectives (SLOs) provide a formalized way to describe, measure, and monitor the performance, quality, and reliability of microservice applications. SLOs provide a shared quality benchmark for application and platform teams to reference for gauging service level agreement (SLA) compliance and continuous improvement.

An SLO describes the high-level objective for acceptable operation and health of one or more services over a length of time (for example, a week or a month). Operators can specify, for example, that a service or application should be healthy 99 percent of the time. An SLO of 99 percent permits a service to have an Error Budget of 1 percent of the time which means to be “unhealthy” 1 percent of the time, which allows for realistic downtime, error cases, planned maintenance windows, and service upgrades. Teams can specify which performance characteristics and thresholds are key to the health of their applications. Multiple SLOs can be defined for a single service, reflecting the reality of Quality of Service (QoS) contracts between different classes of end users.

An SLO consists of one or more service level indicators (SLIs). SLOs defined using a combination of SLIs allow teams to describe service health in a more precise and relevant way. SLIs capture important low-level performance characteristics for a particular service. Tanzu Service Mesh collects SLI metrics on ten second intervals for every service instance that is part of the mesh. An example of an SLI would be 99 percent of successful requests respond with latencies faster than 350 ms (99th percentile latency < 350 ms). Another example is an SLI set for a service that responds with error codes for fewer than 0.1 percent of requests (error rate < 0.1%). Tanzu Service Mesh incorporates SLO and SLI measurements by displaying them in real time through its user interface.

Actionable SLOs can also be used in combination to, or drive autoscaling for services. An example of this configuration is found here.

Appendix A: Sample .yaml files

For reference, the yaml files used for AKO and AMKO are provided below. These files are for reference only. Configuration is provided as a sample only and will vary based on the environment.

AKO yaml

```
---
AKOSettings:
  apiServerPort: 8080
  clusterName: tko2-simple-cluster
  cniPlugin: antrea
  deleteConfig: "false"
  disableStaticRouteSync: "false"
  enableEVH: false
  enableEvents: "true"
  fullSyncFrequency: "1800"
  layer7Only: false
  logLevel: WARN
  namespaceSelector:
    labelKey: ""
    labelValue: ""
  primaryInstance: true
  servicesAPI: true
  vipPerNamespace: "false"
ControllerSettings:
  cloudName: Default-Cloud
  controllerHost: "172.29.250.20"
  controllerVersion: "20.1.7"
  serviceEngineGroupName: Default-Group
  tenantName: admin
L4Settings:
  autoFQDN: default # Set to “disabled” if using Tanzu Service Mesh
  defaultDomain: pse.lab
L7Settings:
  defaultIngController: "true"
  enableMCI: "false"
  noPGForSNI: false
  passthroughShardSize: SMALL
  serviceType: ClusterIP
  shardVSSize: LARGE
NetworkSettings:
  bgpPeerLabels: []
  enableRHI: false
  nodeNetworkList: []
  nsxtT1LR: ""
  vipNetworkList:
    - cidr: 172.31.248.0/22
      networkName: DVS-Frontend-Network
avicredentials:
#  authtoken: ~
#  certificateAuthorityData: ~
  password: password
  username: admin
image:
  pullPolicy: IfNotPresent
  repository: projects.registry.vmware.com/ako/ako
logFile: avi.log
mountPath: /log
nodePortSelector:
  key: ""
  value: ""
persistentVolumeClaim: ""
podSecurityContext: {}
rbac:
  pspEnable: false
replicaCount: 1
resources:
  limits:
    cpu: 350m
    memory: 400Mi
  requests:
    cpu: 200m
    memory: 300Mi
```

AMKO Yaml

```
---
configs:
  controllerVersion: "20.1.7"
  gslbLeaderController: "172.29.250.20"
  logLevel: INFO
  memberClusters:
    -       clusterContext: "arn:aws:eks:us-east-1:8728XXXXXXX:cluster/tkoeks"
    -       clusterContext: tko2-simple-cluster
  refreshInterval: 1800
  useCustomGlobalFqdn: false
globalDeploymentPolicy:
  appSelector:
    label:
      amko: "gslb   <example label key-value for an ingress/service type LB>"
  matchClusters:
    -       cluster: "arn:aws:eks:us-east-1:8728XXXXXXX:cluster/tkoeks"
    -       cluster: tko2-simple-cluster
gslbLeaderCredentials:
  password: password
  username: admin
image:
  pullPolicy: IfNotPresent
  repository: projects.registry.vmware.com/ako/amko
logFile: amko.log
mountPath: /
multiClusterIngress:
  enable: false
persistentVolumeClaim: ""
rbac:
  pspEnable: false
replicaCount: 1
resources:
  limits:
    cpu: 250m
    memory: 300Mi
  requests:
    cpu: 100m
    memory: 200Mi
service:
  port: 80
  type: ClusterIP
serviceAccount:
  annotations: {}
  create: true
  name: ~
serviceDiscovery: ~
```

About the Author

Ather Jamil is a member of the technical staff in the Office of the CTO at VMware. Ather brings 25 plus years of experience in the Information Technology industry. Starting his career with Compaq Computers in the early 90' while still in college, Ather was soon leading efforts in building several "industry first" technologies including PC Blades, virtual storage, and architectures and composable infrastructure. Ather also spent part of his career building video analytics and intelligence solution for several organizations in the public sector and has authored several papers in the information technology field. Building innovative and cutting-edge solutions that solve customer problems is Ather's passion and part of his responsibilities at VMware.