Deploying and Managing the VMware NSX Application Platform

Modified on 16 NOV 2023

VMware NSX-T Data Center 3.2

VMware NSX Intelligence 3.2

vmware®
by Broadcom

You can find the most up-to-date technical documentation on the VMware by Broadcom website at:

https://docs.vmware.com/

VMware by Broadcom
3401 Hillview Ave.
Palo Alto, CA 94304
www.vmware.com

Contents

Deploying and Managing the VMware NSX Application Platform 1

The *Deploying and Managing the VMware NSX Application Platform* document describes how to deploy, configure, upgrade, and manage the VMware NSX® Application Platform.

Intended Audience

This information is intended for enterprise system administrators who must deploy or manage the NSX Application Platform and activate the NSX applications that are hosted on the platform. A familiarity with the administration of VMware NSX-T Data Center™ and familiarity with Kubernetes are assumed.

Related Documentation

If necessary, refer to the VMware NSX Documentation set for versions 3.2 or later when you are deploying or upgrading the NSX Application Platform. You use the VMware NSX® Manager™ user interface when you deploy and upgrade the NSX Application Platform.

You can also refer to the following documentation for the NSX features that are hosted on the NSX Application Platform.

- *Activating and Upgrading VMware NSX Intelligence* document for version 3.2 or later for information on activating and upgrading the NSX Intelligence feature.

  This document is delivered with the NSX Intelligence documentation set at https://docs.vmware.com/en/VMware-NSX-Intelligence/index.html.

- *VMware NSX Network Detection and Response Activation and Administration Guide* for information about activating and administering the VMware NSX® Network Detection and Response™ feature.

  This document is delivered with the NSX Intelligence documentation set at https://docs.vmware.com/en/VMware-NSX-Intelligence/index.html.

- For information about the VMware NSX® Malware Prevention feature, see the Security section of the *NSX-T Data Center Administration Guide* for version 3.2 or later. That document is delivered with the VMware NSX Documentation set.

Getting Started with the NSX Application Platform

2

Before you start to deploy the VMware NSX® Application Platform, get familiar with an overview of its purpose and the prerequisites you must meet to successfully deploy the platform.

Overview

The NSX Application Platform is a modern microservices platform that hosts the following NSX features that collect, ingest, and correlate network traffic data in your NSX-T environment.

- VMware NSX® Intelligence™
- VMware NSX® Network Detection and Response™
- VMware NSX® Malware Prevention
- VMware NSX® Metrics

As network traffic data is produced, captured, and analyzed, the NSX Application Platform provides the platform that can be scaled out to meet the needs of these data-intensive features and the core services that support them.

Following is a list of some of the core services utilized by these NSX features. These services can be scaled out as the need arises.

- Messaging
- Analytics
- Data Storage
- Metrics

The NSX Application Platform is available beginning with NSX-T Data Center 3.2. After you meet the minimum system prerequisites and prepare for any existing analytics data that you want migrated from previous NSX Intelligence installation, you can deploy the platform using the NSX Manager user interface.

Once the platform is deployed, you can activate any of the NSX feature features previously mentioned. The system prerequisites assigned for each feature must also be met in order to activate it. See Chapter 4 NSX Features Available on the NSX Application Platform for brief descriptions about these features.

Read the following topics next:

- NSX Application Platform Deployment Prerequisites
- License Requirement for NSX Application Platform Deployment
- NSX Application Platform System Requirements
- Upload the NSX Application Platform Docker Images and Helm Charts to a Private Container Registry
- Generate a TKG Cluster on Supervisor Configuration File with a Non-Expiring Token

NSX Application Platform Deployment Prerequisites

To install the NSX Application Platform successfully and to activate the NSX features that it hosts, you must prepare the deployment environment so that it meets the minimum required resources.

You must satisfy the prerequisites listed in the following sections before you start deploying the NSX Application Platform.

NSX-T Data Center version requirement

Confirm that the NSX-T Data Center product version you are using is compatible with the NSX Application Platform version that you plan to deploy, along with its related NSX-T Data Center features (NSX Intelligence, NSX Network Detection and Response, NSX Malware Prevention, and NSX Metrics).

The versioning of the NSX-T Data Center features that are hosted on the NSX Application Platform matches the NSX Application Platform version number, and not the NSX-T Data Center product version number.

Important In an NSX Federation environment, you can deploy the NSX Application Platform on Local Managers only. You cannot deploy the NSX Application Platform using Global Managers. You can access the NSX Application Platform using a Local Manager only.

To determine which NSX Application Platform version you can deploy with which NSX-T Data Center version, use the following compatibility matrix.

NSX-T Data Center version	Compatible NSX Application Platform version
3.2.x	3.2.0, 3.2.1

Use the following information in determining which documentation to use for your specific NSX Application Platform activation workflow.

- If you need to install a brand new NSX-T Data Center installation, see the *NSX-T Data Center Installation Guide* for version 3.2.x or later in the VMware NSX Documentation set for installation instructions.
- For information about upgrading from an earlier NSX Application Platform version 3.2.x installation, see Upgrade the NSX Application Platform.

- If you are upgrading from NSX-T Data Center 3.1.x or earlier without NSX Intelligence installed, see the *NSX-T Data Center Upgrade Guide* in the VMware NSX Documentation set.
- If you are upgrading from NSX-T Data Center 3.1.x or earlier with an installation of NSX Intelligence 1.2.x or earlier, you must prepare your current NSX Intelligence installation before you upgrade to NSX Intelligence 3.2.x, and NSX-T Data Center 3.2.x. See the *Activating and Upgrading VMware NSX Intelligence* documentation for version 3.2 in the VMware NSX Intelligence Documentation set.

Valid NSX-T or NSX Data Center license requirement

To deploy the NSX Application Platform, the current NSX Manager session in use must have a valid license in effect during the NSX Application Platform deployment.

See License Requirement for NSX Application Platform Deployment for the list of valid licenses.

Valid NSX-T Data Center user role

To deploy the NSX Application Platform, you must have Enterprise Admin role privileges.

Valid CA-signed certificates

- If your NSX Manager appliance uses CA-signed certificates with partial chain on the NSX Manager Unified Appliance cluster, you must replace the certificate with a full certificate chain. See VMware Knowledge Base article 78317 for more information.
- When using multiple NSX Manager appliances, your environment must meet one of the following certificate prerequisites.
  - All the appliances must share the same SSL certificate.
  - A dedicated SSL certificate must be issued for each appliance, where the certificate Common Name (CN) must be unique across all nodes.
  - When using a Virtual IP (VIP), the cluster certificate must either be the same as shared by all individual appliances or must be unique from all the nodes.

Required resources for the Kubernetes cluster

- VMware supports an NSX Application Platform deployment on a Tanzu Kubernetes Grid (TKG) Cluster on Supervisor or an Upstream Kubernetes cluster.

  Important Upstream Kubernetes refers to the vanilla, open-source Kubernetes maintained by the Cloud Native Computing Foundation and does not cover any distributions or releases of Kubernetes which are not explicitly listed in the following table.

  For TKG Cluster on Supervisor installation and configuration information, see Installing and Configuring vSphere with Tanzu (version 8.0) or VMware vSphere Documentation website for other versions.

  VMware tested and supports the following versions.

NSX Application Platform version	TKG Cluster on Supervisor version	Upstream Kubernetes cluster version
3.2.0, 3.2.1	■ 1.17.17, 1.18.19 (See the following Note) ■ 1.19.11, 1.20.7, 1.21.2, 1.21.6	1.17, 1.18, 1.19, 1.20, 1.21

Note The NSX Application Platform **Scale Up** operation is not supported on TKG Cluster on Supervisor versions 1.17.x and 1.18.x. See Failed to Increase the Volume Size of the Data Storage Disk for details.

- Your infrastructure administrator must configure the TKG Cluster on Supervisor or upstream Kubernetes cluster on which you can deploy the NSX Application Platform and the NSX features that the platform hosts. Enough resources must be allocated to the TKG Cluster on Supervisor or upstream Kubernetes cluster to deploy the NSX Application Platform pods. Because each supported NSX feature has specific resource requirements, determine which of the hosted NSX feature you plan to use.

  See the NSX Application Platform System Requirements topic for details about the supported form factors and their resource requirements.

- **Important** The Kubernetes guest cluster used for the NSX Application Platform must use a default Service Domain of `cluster.local`. This is the default value and is defined in the cluster configuration:

  ```
  settings: network: serviceDomain: cluster.local
  ```

  For NSX Application Platform, do not change this value or set a non-default service domain.

- Your infrastructure administrator must also install and configure the following infrastructures in advance.
  - Container Network Interface (CNI), such as Antrea, Calico, and Flannel.
  - Container Storage Interface (CSI). You must have an available storage class in the TKG Cluster on Supervisor or upstream Kubernetes cluster to provision dynamic volumes. To scale up the data storage volume, the storage class must support volume resize.

Internet access requirement

Ensure that your NSX-T Data Center system can access the public VMware-hosted registry and repository where you can obtain the packaged NSX Application Platform Helm chart and Docker images. The direct Internet access is only required during the installation and upgrade operations. This access is limited to the outbound access on TCP Port 443 (HTTPS) to `https://projects.registry.vmware.com` for the purpose of accessing the NSX Application Platform installation Helm charts and Docker images. No inbound access or permanent outbound access is required.

The outbound Internet access is required for both the NSX-T Data Center Unified Appliance VMs and NSX Application Platform guest cluster worker nodes.

If you configured your NSX-T Data Center environment to use an Internet proxy server using the **System > General Settings > Internet Proxy Server** tab, note that the NSX Application Platform can not be deployed using an Internet proxy server. If your Kubernetes cluster does not have access to the Internet or you have security restrictions, see the next optional requirement for an optional Private container registry with chart repository service.

(Optional Requirement) Private container registry with chart repository service

To simplify the NSX Application Platform deployment process, use the VMware-hosted registry and repository. This deployment process uses an outbound connection only and does not retain customer data.

(Optional) If your Kubernetes cluster does not have access to the Internet or you have security restrictions, your infrastructure administrator must set up a private container registry with a chart repository service. Use this private container registry to upload the NSX Application Platform Helm charts and Docker images required to deploy the NSX Application Platform. VMware used Harbor to validate the deployment process that uses a private container registry, however, the NSX Application Platform deployment is standards-based. See Upload the NSX Application Platform Docker Images and Helm Charts to a Private Container Registry for details.

(Optional Requirement) URL for a private container registry

If you are using a private container registry, obtain from your infrastructure administrator the URL for that registry. You use this URL during the deployment process.

Required Kubernetes configuration file

You must also obtain the Kubernetes configuration file from your infrastructure administrator. You need the `kubeconfig` file during the NSX Application Platform deployment for the NSX Manager to securely access your TKG Cluster on Supervisor or upstream Kubernetes cluster. The `kubeconfig` file must have all the privileges to access all the resources of the TKG Cluster on Supervisor or upstream Kubernetes cluster.

Important The default `kubeconfig` file in a VMware vSphere® with Tanzu Kubernetes Guest Cluster contains a token which expires after ten hours by default. While this expired token does not impact functionality, it results in a warning message regarding out-of-date credentials. To avoid the warning, before you deploy the NSX Application Platform on a TKG Cluster on Supervisor, work with your infrastructure administrator to create a long-lived token you can use during the platform deployment. See Generate a TKG Cluster on Supervisor Configuration File with a Non-Expiring Token for details on how to extract the token.

Required Service Name or Interface Service Name (FQDN)

During the NSX Application Platform deployment, you provide a fully qualified domain name (FQDN) for the **Service Name** text box in an NSX-T Data Center 3.2.0 deployment or for the **Interface Service Name** text box in an NSX-T Data Center 3.2.1 or later deployment.

The **Service Name** or **Interface Service Name** value is used as the HTTPS endpoint to connect to the NSX Application Platform.

To obtain the FQDN value, use one of the following workflows.

- You must configure FQDN with a static IP address in the DNS server before the NSX Application Platform deployment. The TKG Cluster on Supervisor or upstream Kubernetes cluster infrastructure must be able to assign a static IP address. The following are the supported Kubernetes environments.
    - MetalLB - an external load balancer for upstream Kubernetes cluster.
    - VMware Tanzu® Kubernetes for VMware vSphere® 7.0 U2 and VMware vSphere 7.0 U3 with NSX-T Data Center.
    - VMware vSphere with Tanzu using vSphere networking. See the VMware vSphere document, Enable Workload Management with vSphere Networking, for more information.
- If you have the External DNS installed (see the Kubernetes SIGs - External DNS webpage), you only have to provide the FQDN when prompted for a Service Name. Your Kubernetes infrastructure automatically configures the FQDN with a dynamic IP address in the DNS server.

  The Workload Control Plane (WCP) with External DNS installed is the supported Kubernetes environment.

Required Messaging Service Name (for NSX-T Data Center 3.2.1 or later deployments)

The Messaging Service Name value is an FQDN for the HTTPS endpoint that is used to receive the streamlined data from the NSX-T Data Center data sources.

Required ports and protocols

Verify that the required ports on your Kubernetes cluster host are open for the NSX Application Platform to access. See the VMware Ports and Protocols webpage.

Required communication from your Kubernetes cluster nodes

Confirm that the Kubernetes cluster nodes you are using can reach the NSX Manager appliance.

System times synchronization requirement

Synchronize the system times on the TKG Cluster on Supervisor or upstream Kubernetes cluster nodes and the NSX Manager appliance.

License Requirement for NSX Application Platform Deployment

To deploy the NSX Application Platform, your NSX Manager session must be using a valid license during the NSX Application Platform deployment.

You can use one of the following licenses.

- NSX Data Center Evaluation
- NSX-T Evaluation
- NSX-T Enterprise Plus
- NSX Data Center Enterprise Plus
- NSX Advanced Threat Prevention Add-On with one of the following required base licenses:
  - NSX Data Center Advanced
  - NSX Data Center Enterprise Plus
  - NSX-T Advanced
  - NSX-T Enterprise Plus
- NSX Distributed Firewall licenses that are available in NSX-T Data Center v3.2 and later:
  - NSX Distributed Firewall
  - NSX Distributed Firewall with Threat Prevention
  - NSX Distributed Firewall with Advanced Threat Prevention
  - NSX Threat Prevention Add-On for Distributed Firewall
  - NSX Advanced Threat Prevention Add-On for Distributed Firewall, NSX-T Data Center Advanced or NSX-T Data Center Enterprise Plus
- NSX Gateway Firewall licenses that are available in NSX-T Data Center v3.2 and later:
  - NSX Gateway Firewall – VM
  - NSX Gateway Firewall with Threat Prevention– VM
  - NSX Gateway Firewall with Advanced Threat Prevention– VM
  - NSX Threat Prevention Add-On for Gateway Firewall– VM
  - NSX Advanced Threat Prevention Add-On for Gateway Firewall– VM

NSX Application Platform System Requirements

This section lists the form factors that the NSX Application Platform supports, along with the minimum resources required for each. Additional system requirements prior to deployment are also provided.

Minimum System Requirements

In addition to the information listed in NSX Application Platform Deployment Prerequisites, use the following table as a guide when working with your infrastructure administrator to prepare for deploying the NSX Application Platform and the NSX-T Data Center features hosted on the platform. The form factor you select determines which NSX-T Data Center features you can activate or install on the platform.

Before you deploy the NSX Application Platform, determine the size of the VMware TKG Cluster on Supervisor or upstream Kubernetes cluster and the minimum number of nodes that your infrastructure administrator must allocate.

Form Factor	Minimum and Recommended # of Nodes in a TKG Cluster on Supervisor or Upstream Kubernetes Cluster	vCPU	Memory	Storage	Ephemeral Storage	Supported NSX Features
Standard	■ A minimum of 1 control plane node is required but 3 control plane nodes are recommended. ■ 3 or more worker nodes (See the Important note in the About the support for a minimum of 3 worker nodes section later in this topic.)	■ 2 vCPUs for the control plane node (See more information later in this topic about using `guaranteed-small` VM) ■ 4 vCPU per worker node	■ 4 GB RAM for the control plane node (See more information later in this topic about using `guaranteed-small` VM) ■ 16 GB RAM per worker node	200 GB per NSX Application Platform instance (See more information later in this topic.)	64 GB	■ NSX Network Detection and Response ■ NSX Malware Prevention ■ NSX Metrics (See the About the support for a minimum of 3 worker nodes later in this topic.)
Advanced	■ A minimum of 1 control plane node is required but 3 control plane nodes are recommended. ■ 3 or more worker nodes (See the Important note in the About the support for a minimum of 3 worker nodes section later in this topic.)	■ 2 vCPUs for the control plane node (See more information later in this topic about using `guaranteed-small` VM) ■ 16 vCPU per worker node	■ 4 GB RAM for the control plane node (See more information later in this topic about using `guaranteed-small` VM) ■ 64 GB RAM per worker node	1 TB per NSX Application Platform instance (See more information later in this topic.)	64 GB	■ NSX Intelligence ■ NSX Network Detection and Response ■ NSX Malware Prevention ■ NSX Metrics (See the About the support for a minimum of 3 worker nodes later in this topic.)
Evaluation (See the following *Note for Evaluation Form Factor.)	1 control plane node and 1 worker node	■ 2 vCPUs for the control plane node (See more information later in this	■ 4 GB RAM for the control plane node (See more information later in this	1 TB per NSX Application Platform instance	64 GB	■ NSX Intelligence ■ NSX Network Detection and Response

Form Factor	Minimum and Recommended # of Nodes in a TKG Cluster on Supervisor or Upstream Kubernetes Cluster	vCPU	Memory	Storage	Ephemeral Storage	Supported NSX Features
		topic about using `guaranteed-small` VM) - 16 vCPU per worker node	topic about using `guaranteed-small` VM) - 64 GB RAM per worker node			- NSX Malware Prevention - NSX Metrics

About the support for a minimum of 3 worker nodes

Important When using a minimum of 3 Advanced Form Factor worker nodes in a TKG Cluster on Supervisor or Upstream Kubernetes Cluster, only the following NSX-T Data Center features that are hosted on the NSX Application Platform can be activated.

- NSX Intelligence
- NSX Metrics for NSX Application Platform (activated by default)

To activate the following NSX-T Data Center features, in addition to NSX Intelligence and NSX Metrics for NSX Application Platform, a minimum of 4 worker nodes in the guest cluster is required.

- NSX Network Detection and Response
- NSX Suspicious Traffic for Network Traffic Analysis
- NSX Malware Prevention
- NSX Metrics

All the supported features listed in Table 1: Minimum System Requirements for the Standard Form Factor size can be activated on a 3-worker node guest cluster.

About the Evaluation Form Factor

The Evaluation form factor is applicable only in non-production deployments for use in evaluations or demonstrations. It has limited data retention, no scale out or high availability support, and no support for upgrades. If you want to run multiple services on the Evaluation form factor for an extended period of time, increase the worker node resources to 24 vCPU and 128 GB of RAM. If necessary, contact VMware support for more sizing details.

Availability and Resiliency Best Practices

To avoid data loss when the worker nodes fail, do not use the storage classes that have any persistence volume local to the worker nodes. Consider using a remote, independent, and distributed storage class, such as VMware vSAN volumes.

Important Note for Production Deployments The following best practices apply to any production deployment for either Standard or Advanced form factors.

- For increased availability and resiliency, 3 control plane nodes are required.
- VM classes used for the NSX Application Platform control plane nodes and worker nodes must use a guaranteed reservation of 100% for CPU and Memory resources to avoid any resource overcommitment.

There must be no resource contention for Storage I/O operations per second (IOPS) or Network bandwidth. The vSphere Storage and Network I/O Control are platform features that can help prioritize resources for the NSX Application Platform.

Load Balancing Requirement

When deploying the NSX Application Platform, you must configure your TKG Cluster on Supervisor or upstream Kubernetes cluster to have a load balancer (LB) IP pool with at least five IP addresses. To finish an NSX Application Platform deployment successfully, the platform requires at least five available IP addresses. If you plan to scale out your NSX Application Platform deployment later, your TKG Cluster on Supervisor or Kubernetes cluster LB IP pool must contain one more IP address per Kubernetes node used by the platform. Consider configuring your TKG Cluster on Supervisor or upstream Kubernetes cluster LB IP pool with a total of 15 IP addresses, since VMware only supports a maximum of 10 additional Kubernetes nodes after scaling out the platform.

Additional Volume Requirement

For NSX Application Platform, your guest cluster worker nodes require an additional volume of at least 64 GiB for ephemeral storage.

To specify the disk and storage parameters for each node type, use the information in the following table. For more information, see the v1alpha3 Example: TKC with Default Storage and Node Volumes topic in the *Using Tanzu Kubernetes Grid 2.0 on Supervisor with vSphere with Tanzu 8* documentation.

Node Type	Volume Name	Volume Capacity	Volume mountPath
Worker Node	`containerd`	64 GiB	`/var/lib/containerd`

Control Plane Node Size Requirement

If you are deploying the NSX Application Platform and NSX-T Data Center features, like NSX Intelligence, on a TKG Cluster on Supervisor, the default virtual machine class type of `guaranteed-small` size (2 vCPUs and 4 GB RAM) might not be sufficient for the TKG Cluster on Supervisor control plane node. Consider using the following class type that has a bigger size for the upstream Kubernetes cluster or TKG Cluster on Supervisor control plane node.

- `guaranteed-medium` (2 vCPUs and 8 GB RAM)

Consult the Virtual Machine Classes for Tanzu Kubernetes Clusters documentation for more information.

Upload the NSX Application Platform Docker Images and Helm Charts to a Private Container Registry

If your Kubernetes cluster does not have Internet access or you have specific security restrictions, work with your infrastructure administrator to upload the NSX Application Platform Helm charts and Docker images to a private container registry that you can access and use to deploy the NSX Application Platform.

Note The following steps were validated using a private Harbor container registry. If you are using another container registry, you might need to adjust some steps for that registry.

Your infrastructure administrator (or anyone who has access to the TKG Cluster on Supervisor or upstream Kubernetes cluster in which the private container registry is installed) must perform the following steps.

Prerequisites

- Your infrastructure administrator must install a private container registry, such as Harbor, with chart repository service. See Harbor Installation with Chart Repository Service. They will use this private container registry to host the NSX Application Platform Helm charts and Docker images.

  Important If you are using a VMware TKG Cluster on Supervisor, do not use its embedded Harbor container registry for hosting the NSX Application Platform Helm charts and Docker images. Your infrastructure administrator must set up a separate Harbor container registry.

- Beginning with NSX-T Data Center 3.2.3.1 release, the VMware-hosted NSX Application Platform registry and repository locations no longer support ChartMuseum-compatible private container registry, such as Harbor. If you need to continue using the ChartMuseum provided by Harbor, the Harbor version must be earlier than version 2.8.1.
- For a production environment, your infrastructure administrator must also obtain a CA certificate (signed by a reputable public Certificate Authority) to configure HTTPS access to the private Harbor container registry that they installed. For more information, see the Configure HTTPS Access to Harbor webpage..

- Ensure that the `Docker` tool is installed and configured correctly on the system that will be used for uploading the images and charts. The `Login succeeded` must be displayed after running the following command.

  ```
  docker login <private-registry-URL> --username <private-registry-account-name> --password
  <private-registry-account-password>
  ```

- Ensure that the same CA certificate used by your Harbor container registry is already installed in your `Docker` tool.
- Ensure that the trust is established between your private Harbor CA certificate and your NSX-T Data Center Unified Appliance (UA).
- Verify that the `curl` tool is installed on the system that will be used for uploading the images. Use the following command to verify. The curl version is displayed if the `curl` is installed.

  ```
  curl --version
  ```

Procedure

1. Download the NSX Application Platform deployment bundle from the VMware Product Download portal for NSX-T Data Center 3.2.x. Save the bundle to a system from which you can upload the Helm charts and Docker images to your private container registry.

   Use the NSX-T Data Center download page appropriate for the NSX Application Platform version available or the version that you want to deploy.

2. Extract the contents of the NSX Application Platform deployment bundle using the following command. The *<version-number>* is the specific version number and the build number of the bundle. For example, `VMware-NSX-Application-Platform-3.2.1.1.0.20140674.tgz`.

   ```
   tar xvf VMware-NSX-Application-Platform-<version-number>.tgz
   ```

   This step might take several minutes to finish.

3. Locate and edit the `upload_oci_artifacts_to_private_harbor.sh` file with a text editor. If you need to use a ChartMuseum-compatible private container registry, edit the `upload_artifacts_to_private_harbor.sh` script file.

   You use the `upload_oci_artifacts_to_private_harbor.sh` or the `upload_artifacts_to_private_harbor.sh` script file to upload the extracted NSX Application Platform Helm charts and Docker images.

   a. Set the `DOCKER_REPO` property to the URL for your private container registry.

      For example, `DOCKER_REPO=harbor-repo.mycompany.com/nsx_intelligence`

   b. Set the `DOCKER_USERNAME` property to the user name of the private container registry account.

c Set the DOCKER_PASSWORD property to the password of the private container registry account.

d Save the changes in the script file.

4 Change the executable permission for the modified shell script file using the following command.

For an OCI-compatible private container registry, use the following command.

```
chmod +x upload_oci_artifacts_to_private_harbor.sh
```

For a ChartMuseum-compatible private container registry, use the following command.

```
chmod +x upload_artifacts_to_private_harbor.sh
```

5 Run the script file using the following command.

For an OCI-compatible private container registry, use the following command.

```
./upload_oci_artifacts_to_private_harbor.sh
```

For a ChartMuseum-compatible private container registry, use the following command.

```
./upload_artifacts_to_private_harbor.sh
```

Results

The system uploads the NSX Application Platform Helm charts and Docker images to your private container registry.

What to do next

The NSX enterprise administrator must continue with ensuring that the prerequisites listed in NSX Application Platform Deployment Prerequisites are met before continuing with the NSX Application Platform deployment.

Generate a TKG Cluster on Supervisor Configuration File with a Non-Expiring Token

The default kubeconfig file in a VMware vSphere with Tanzu Kubernetes Guest Cluster contains a token which expires after ten hours by default and results in a warning message. To avoid the warning, work with your Kubernetes infrastructure administrator to generate a valid TKG Cluster on Supervisor configuration file with a non-expiring token that you can use during the NSX Application Platform deployment.

When the token in the default kubeconfig file expires, you see the following warning message in the NSX Manager UI for the NSX Application Platform.

```
Unable to connect, system has encountered a connectivity issue due to the
expiry of Kubernetes Configuration. Update the Kubernetes Configuration to
resolve.
```

The warning does not have an impact on the functionality of the NSX Application Platform nor any of the NSX security features currently activated. However, if you do not replace the default token ten hours after an NSX Application Platform deployment on the Tanzu Kubernetes Guest Cluster, you must generate a valid (not expired) token every time you perform the following operations:

- Deploy the NSX Application Platform
- Upgrade the NSX Application Platform
- Delete the NSX Application Platform

To generate a TKG Cluster on Supervisor configuration file with a non-expiring token that you can use during the NSX Application Platform deployment, work with your Kubernetes infrastructure administrator using the following procedure.

Procedure

1 Log in the vSphere with Tanzu Kubernetes Guest Cluster using the following command.

```
kubectl vsphere login --server <supervisor-cluster_ip> -u <user> --tanzu-kubernetes-
cluster-name <tkg-cluster-name> --tanzu-kubernetes-cluster-namespace <namespace>
```

The parameters are as follows:

- *<supervisor-cluster_ip>* is the Control Plane Node Address which can be found in the vSphere Client by selecting **Workload Management > Supervisor Cluster**.
- *<user>* is the account that has administrator access to the TKG Cluster on Supervisor.
- *<tkg-cluster-name>* is the name of the TKG Cluster on Supervisor.
- *<namespace>* is the vSphere namespace where this cluster resides.

For example,

```
kubectl vsphere login --server 192.111.33.22 -u administrator@vsphere.local --tanzu-
kubernetes-cluster-name napp-tkg-cluster --tanzu-kubernetes-cluster-namespace napp
```

2 Run each of the following commands separately to generate an administrator service account and create a cluster role binding.

```
kubectl create serviceaccount napp-admin -n kube-system

kubectl create clusterrolebinding napp-admin --serviceaccount=kube-system:napp-admin --
clusterrole=cluster-admin
```

3 (Required) (For Kubernetes version 1.24 and later) Manually create the authentication token for the administrator service account. Use the following information.

   a Create a YAML file with a service account. Use the following content for an example YAML file named, `napp-admin.yaml`.

   ```
   apiVersion: v1
   kind: Secret
   type: kubernetes.io/service-account-token
   metadata:
      name: napp-admin
      namespace: kube-system
      annotations:
         kubernetes.io/service-account.name: "napp-admin"
   ```

   b Use the following command to create the authentication token with a service account.

   ```
   kubectl apply -f <filename create above.yaml>
   ```

   Using the example YAML file in the previous step, the command to use is as follows.

   ```
   kubectl apply -f napp-admin.yaml
   ```

   The authentication token or secret is generated.

4 To obtain the authentication token for the administrator service account and the cluster certificate authority, run the following commands separately.

For supported Kubernetes version 1.24 and later, use the following commands.

```
SECRET=$(kubectl get secrets napp-admin -n kube-system -ojsonpath='{.metadata.name}')

TOKEN=$(kubectl get secret $SECRET -n kube-system -ojsonpath='{.data.token}' | base64 -d)

kubectl get secrets $SECRET -n kube-system -o jsonpath='{.data.ca\.crt}' | base64 -d > ./
ca.crt
```

For supported Kubernetes versions prior to version 1.24, use the following commands.

```
SECRET=$(kubectl get serviceaccount napp-admin -n kube-system
-ojsonpath='{.secrets[].name}')

TOKEN=$(kubectl get secret $SECRET -n kube-system -ojsonpath='{.data.token}' | base64 -d)

kubectl get secrets $SECRET -n kube-system -o jsonpath='{.data.ca\.crt}' | base64 -d > ./
ca.crt
```

5 Get the TKG Cluster on Supervisor URL. Run the following commands separately at the command prompt.

```
CONTEXT=$(kubectl config view -o jsonpath='{.current-context}')

CLUSTER=$(kubectl config view -o jsonpath='{.contexts[?(@.name ==
```

```
"'"$CONTEXT"'")].context.cluster}')

URL=$(kubectl config view -o jsonpath='{.clusters[?(@.name ==
"'"$CLUSTER"'")].cluster.server}')
```

6 Generate a configuration file, with a non-expiring token, for the TKG Cluster on Supervisor.

```
TO_BE_CREATED_KUBECONFIG_FILE="<file-name>"
```

The parameter *<file-name>* is the name of `kubeconfig` file you are trying to create.

```
kubectl config --kubeconfig=$TO_BE_CREATED_KUBECONFIG_FILE set-cluster $CLUSTER --
server=$URL --certificate-authority=./ca.crt --embed-certs=true

kubectl config --kubeconfig=$TO_BE_CREATED_KUBECONFIG_FILE set-credentials napp-admin --
token=$TOKEN

kubectl config --kubeconfig=$TO_BE_CREATED_KUBECONFIG_FILE set-context $CONTEXT --
cluster=$CLUSTER --user=napp-admin

kubectl config --kubeconfig=$TO_BE_CREATED_KUBECONFIG_FILE use-context $CONTEXT
```

7 (Optional) Delete the `ca.crt`, which is a temporary file created during the generation of the new `kubeconfig` file.

8 Use the newly generated `kubeconfig` file during the NSX Application Platform deployment.

Deploying the NSX Application Platform

3

Deploying the NSX Application Platform and installing its components successfully require that you meet multiple prerequisites and complete several deployment steps. Carefully review the deployment checklist before you proceed with the platform deployment and complete the provided procedures in the sections that follow.

Read the following topics next:

- NSX Application Platform Deployment Checklist
- Deploy the NSX Application Platform

NSX Application Platform Deployment Checklist

Use the following checklist to track your progress with the NSX Application Platform deployment workflow and the activation of the NSX features that the platform hosts.

1 Install VMware NSX-T Data Center™ 3.2 or later.

Find your specific installation scenario in the following table and use the corresponding installation procedure.

Installation Scenario	Procedure
Fresh installation of NSX-T Data Center version 3.2 or later	Install NSX-T Data Center version 3.2 or later. For details, see *NSX-T Data Center Installation Guide* in the NSX Documentation set.
NSX-T Data Center 3.1 or earlier installation exists and does not have NSX Intelligence 1.x appliance installed.	Upgrade to NSX-T Data Center 3.2 or later. For details, see *NSX-T Data Center Upgrade Guide* in the NSX Documentation set.
NSX-T Data Center 3.1 or earlier exists with NSX Intelligence 1.x appliance installed.	1 Prepare your NSX Intelligence 1.x installation for the upgrade to NSX Intelligence 3.2. See the "Upgrading NSX Intelligence" section in *Activating and Upgrading VMware NSX Intelligence* for version 3.2 or later in the NSX Intelligence Documentation set. 2 Upgrade to NSX-T Data Center 3.2 or later. For details, see the *NSX-T Data Center Upgrade Guide* for version 3.2 or later in the NSX Documentation set.

2 Ensure that you have a valid license in effect in your NSX-T Data Center 3.2 or later installation. See License Requirement for NSX Application Platform Deployment for details.

3 Verify the compatibility of the NSX Application Platform version that you plan to deploy with the NSX Manager version that you plan to use for the deployment.

   See NSX Application Platform Deployment Prerequisites for the compatibility information.

   Important In an NSX Federation environment, you can deploy the NSX Application Platform on Local Managers only. You cannot deploy the NSX Application Platform using Global Managers. You can access the NSX Application Platform using a Local Manager only.

4 Configure the ports and protocols required for deploying and using the NSX Application Platform. See the VMware Ports and Protocols webpage.

5 Ask your infrastructure administrator to install and configure Tanzu Kubernetes Grid (TKG) Cluster on Supervisor or an Upstream Kubernetes cluster. After the cluster installation and configuration are successfully completed, ask them to allocate the TKG Cluster on Supervisor or upstream Kubernetes cluster resources necessary to deploy the NSX Application Platform and to activate any of the hosted NSX features that you want to use. See the information about resources for TKG Cluster on Supervisor or upstream Kubernetes cluster in NSX Application Platform Deployment Prerequisites.

6 Ensure that you have satisfied all the NSX Application Platform deployment prerequisites. See NSX Application Platform Deployment Prerequisites for more information.

7 Confirm that the NSX Manager appliance and the TKG Cluster on Supervisor or upstream Kubernetes cluster are reachable to each other, and that their system times are synchronized on both.

8 Deploy the NSX Application Platform. See Deploy the NSX Application Platform.

9 Activate one or more of the available NSX features that are hosted on the NSX Application Platform and begin using the activated features. See Chapter 4 NSX Features Available on the NSX Application Platform for more information.

Deploy the NSX Application Platform

After you have an NSX-T Data Center version 3.2 or later installed and all the deployment prerequisites are met, you can proceed to deploy the NSX Application Platform.

Prerequisites

You must meet all of the deployment prerequisites, including the form factor system requirements. See the NSX Application Platform Deployment Prerequisites topic for details.

Caution If you migrated your NSX Intelligence 1.2.x traffic flow data, do not deploy the NSX Application Platform using the Evaluation form factor. Using the Evaluation form factor forces the system to deploy the NSX Application Platform without migrating your traffic flow data from the previous NSX Intelligence 1.2.x installation and causes the loss of information about that previous installation.

Step 1: Prepare to Deploy

To deploy the NSX Application Platform, provide the Helm repository and Docker registry information.

Procedure

1. From your browser, log in with Enterprise Admin privileges to an NSX Manager at https://<nsx-manager-ip-address>.
2. Navigate to **System > NSX Application Platform** in the Configuration section.
3. Click **Deploy NSX Application Platform**.
4. Verify the **Helm Repository** URL and **Docker Registry** path.

   Starting with NSX-T Data Center 3.2.3.1, the **Helm Repository** text box, by default, has the value **`oci://projects.registry.vmware.com/nsx_application_platform/helm-charts`**. This is the public VMware-hosted Helm repository from which the system obtains the packaged NSX Application Platform Helm chart.

   The **Docker Registry** path has the `projects.registry.vmware.com/nsx_application_platform/clustering` value. This is the public VMware-hosted registry location from which the system obtains the NSX Application Platform docker images.

   Using these public VMware-hosted NSX Application Platform registry and repository locations simplify the deployment process. This deployment process is an outbound connection only and does not retain any customer data.

   If you are unable to use the recommended deployment process and access the public VMware-hosted locations, work with your infrastructure administrator to upload the NSX Application Platform Helm chart and Docker images to your company's private Helm repository and Docker registry locations. Both privately-hosted locations must be accessible from the Kubernetes cluster and the NSX Manager appliance you are using for the NSX Application Platform deployment. See Upload the NSX Application Platform Docker Images and Helm Charts to a Private Container Registry.

   If you are using a private Helm repository and Docker registry location, use the following steps.

   a. In the **Helm Repository** text box, enter the private registry URL.

      - For OCI-compatible Helm private repository, use the format `oci://`*`<your-private-registry-server-fqdn>`*`/<your-private-registry-name>/helm-charts`.
      - For ChartMuseum-compatible Helm private repository, use the format `https://`*`<your-private-registry-server-fqdn>`*`/chartrepo/`*`<your-private-registry-name>`*.

   b. In the **Docker Registry** text box, enter your private registry location. For either OCI-compatible or ChartMuseum-compatible Helm registry, use the format *`<your-private-registry-server-fqdn>`*`/`*`<your-private-registry-name>`*`/clustering`.

      There is no leading **`https`** or **`oci`** in that path value.

5 Click **Save URL**.

This step might take some time to complete as the system gathers the NSX Application Platform details from the Helm charts and Docker registry locations.

6 In the **Platform Target Version** text box, verify that the correct NSX Application Platform version is selected for the deployment.

The system derives the list of versions from the Helm repository.

7 Click **Next**.

Step 2: Provide Configuration Details

To deploy the NSX Application Platform, you must provide the configuration information about the TKG Cluster on Supervisor or upstream Kubernetes cluster resources that your infrastructure administrator created for you.

You must have the `kubeconfig` file that you obtained from your infrastructure administrator. This file contains configuration information for your TKG Cluster on Supervisor or upstream Kubernetes environment and provides access information.

Procedure

1 In the **Upload file** text box, click **Select** and navigate to the location of the `kubeconfig` file provided to you by your infrastructure administrator.

2 Click **Upload**.

This step can take some time to complete while the system verifies the Kubernetes configuration file contents.

3 (Optional) If you see the error message `Server version and client version are incompatible, upload the latest Kubernetes Tools version to resolve the error`, upload a compatible version of the Kubernetes tools bundle.

You can use the Kubernetes Tools bundle provided in the VMware Product Download portal for the NSX-T Data Center version that you are using. When you download the Kubernetes Tools bundle, the default name is `kubernetes-tools-`*buildversion*`.tar.gz`. For example, `kubernetes-tools-1.20.11-00_3.5.4-1.tar.gz`. Do not rename the file when you download it. The file is signed with a VMware private key.

a Either select **Upload Local File** or **Upload Remote File**.

b If you selected **Upload Local File**, click **Select** and navigate to the location of the Kubernetes Tools file.

c If you selected **Upload Remote File**, enter the URL from which the system can obtain the compatible Kubernetes Tools file. For example, enter the URL of the `kubernetes-tools-`*buildversion*`.tar.gz` file that you downloaded.

d Click **Upload**.

4 Verify the **Cluster Type** information is correct.

This information refers to the type of Kubernetes environment. Currently, `Standard` is the only type supported.

5 Verify the **Storage Class** information is correct.

The system obtains the storage class values from the Kubernetes configuration file and makes them available in the drop-down menu.

6 Enter a valid fully qualified domain name (FQDN) value for the **Service Name** text box in an NSX-T Data Center 3.2.0 deployment or for the **Interface Service Name** text box in an NSX-T Data Center 3.2.1 or later deployment.

The **Service Name** or **Interface Service Name** value is used as the HTTPS endpoint to connect to the NSX Application Platform. See details in the Service Name (FQDN) section of the NSX Application Platform Deployment Prerequisites topic.

7 For an NSX-T Data Center 3.2.1 or later deployment, enter a valid FQDN value for the **Messaging Service Name** text box.

The **Messaging Service Name** value is the HTTPS endpoint used to receive the streamlined data from the NSX data sources.

8 Select the **Form Factor** appropriate for your needs. Review the NSX Application Platform System Requirements for details.

9 (Optional) If you select the **Evaluation** form factor, confirm your intention to use the deployment for non-production use only.

a Read the information displayed in the **Evaluation** dialog box.

b Select the **Confirm** check box to acknowledge that you plan to use the NSX Application Platform deployment for proof of concept and non-production use only.

c Click **Select.**

10 Back in the **Configuration** tab, click **Next**.

Step 3: Precheck the Platform

The system needs to check the configuration information that have been obtained before proceeding with the NSX Application Platform deployment.

Procedure

1 In the **Precheck Platform** tab, click **Run Prechecks**.

The system displays the progress status for each precheck performed.

2 If there are any errors displayed in the **Details** column, click the link provided for the error. Obtain the details and make the necessary corrections to resolve the reported errors. See Chapter 6 Troubleshooting Issues with the NSX Application Platform for more information.

3 Click **Next**.

Step 4: Review & Deploy

The NSX Application Platform deployment wizard gives you the chance to review and edit any of the configuration details that the system has obtained.

Procedure

1. In the **Review & Deploy** tab, review the information displayed in the Platform, Configuration and Prechecks sections. Click the **Edit** link for the corresponding section where changes are needed.

   When you click **Edit**, you are taken back to the tab where you can update the information.

2. If all the information looks correct, click **Deploy**.

   The system proceeds with the final deployment steps and provides progress information in the UI. The steps can take some time to complete.

   Caution After the system migrates the NSX Intelligence 1.2.x traffic flow data, if the NSX Application Platform deployment fails during the NSX Metrics activation, do NOT press **Cancel**. Doing so deletes the persistent storage that contains data migration information and causes the loss of information about the NSX Intelligence 1.2.x installation. Use another browser tab or new browser window to resolve the NSX Metrics issue before returning to the original NSX Application Platform deployment window and clicking **Retry**. For information about troubleshooting the NSX Metrics issue, see Chapter 6 Troubleshooting Issues with the NSX Application Platform.

Results

The system successfully deploys the NSX Application Platform and updates the UI with the details about the platform, such as alarms, cluster information, and so on. The following image shows a sample of what the UI looks like after a successful deployment.

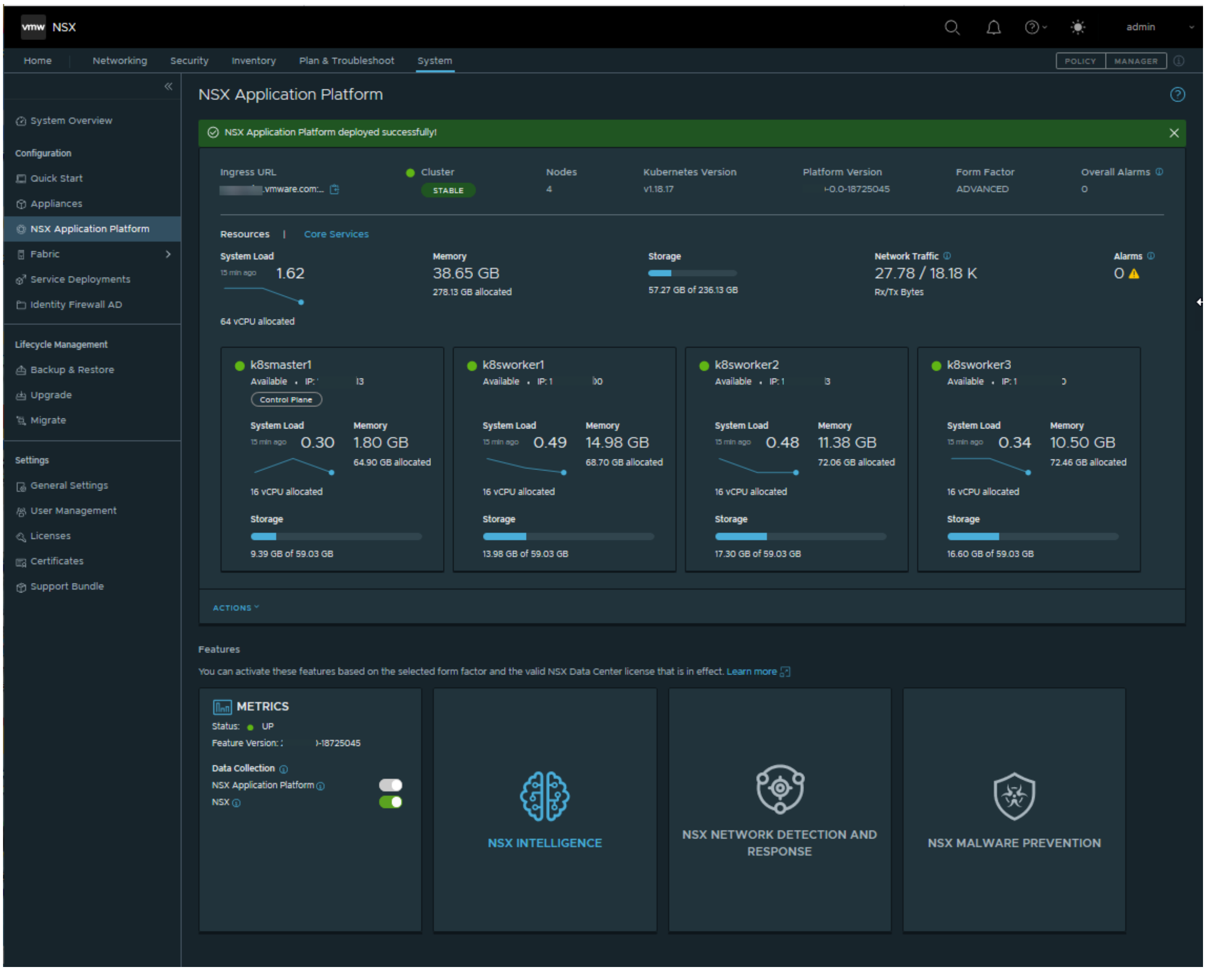

NSX Application Platform

NSX Application Platform deployed successfully!

Ingress URL: [redacted].vmware.com:...
Cluster: STABLE
Nodes: 4
Kubernetes Version: v1.18.17
Platform Version: ı-0.0-18725045
Form Factor: ADVANCED
Overall Alarms: 0

Resources | Core Services

System Load: 15 min ago 1.62 — 64 vCPU allocated
Memory: 38.65 GB — 278.13 GB allocated
Storage: 57.27 GB of 236.13 GB
Network Traffic: 27.78 / 18.18 K Rx/Tx Bytes
Alarms: 0 ⚠

k8smaster1 — Available • IP: [redacted] l3 — Control Plane
System Load: 15 min ago 0.30 — 16 vCPU allocated
Memory: 1.80 GB — 64.90 GB allocated
Storage: 9.39 GB of 59.03 GB

k8sworker1 — Available • IP: 1 [redacted])0
System Load: 15 min ago 0.49 — 16 vCPU allocated
Memory: 14.98 GB — 68.70 GB allocated
Storage: 13.98 GB of 59.03 GB

k8sworker2 — Available • IP: 1 [redacted] 3
System Load: 15 min ago 0.48 — 16 vCPU allocated
Memory: 11.38 GB — 72.06 GB allocated
Storage: 17.30 GB of 59.03 GB

k8sworker3 — Available • IP: 1 [redacted] ɔ
System Load: 15 min ago 0.34 — 16 vCPU allocated
Memory: 10.50 GB — 72.46 GB allocated
Storage: 16.60 GB of 59.03 GB

Features

You can activate these features based on the selected form factor and the valid NSX Data Center license that is in effect. Learn more

METRICS
Status: ● UP
Feature Version: : [redacted])-18725045
Data Collection
NSX Application Platform: ○ off
NSX: ● on

What to do next

You can now activate any of the available NSX features that can be hosted on the NSX Application Platform and is available for the form factor and NSX license that you are currently using. See Chapter 4 NSX Features Available on the NSX Application Platform for more information.

NSX Features Available on the NSX Application Platform

4

After you successfully deploy the NSX Application Platform, you can activate the NSX features that are available to be hosted on the platform.

The following features are either automatically activated or are available for activation after a successful NSX Application Platform deployment. Each feature activation is dependent on meeting the minimum system requirements for that NSX feature. See NSX Application Platform System Requirements for details.

For details about activating an NSX feature, consult the specific documentation mentioned for that feature in the following table.

Hosted NSX Feature	Description
VMware NSX® Metrics	The NSX Metrics feature collects data that allows it to monitor key statistics across the entities in your NSX-T and the NSX Application Platform environments. This feature is automatically activated after a successful NSX Application Platform deployment. By default, the data collection feature is always activated for the NSX Application Platform and cannot be turned off. Data collection for NSX entities is turned on by default, but you can optionally turn it off by toggling **NSX** to off. For information on how to use the NSX Metrics feature, see the APIs to Fetch Time-Series Metrics topic in the *NSX-T Data Center Administration Guide* for version 3.2. For additional API information, see http://developer.vmware.com/apis/nsx-intelligence-&-application-platform/.
VMware NSX® Intelligence™	The NSX Intelligence feature aggregates the network traffic flows in your NSX-T environment to provide deep traffic flow visibility, firewall policy recommendations, and suspicious traffic detection. To allow for scale-out capability, this feature is now hosted on the NSX Application Platform beginning with NSX Intelligence 3.2. For information on how to activate this feature, see the *Activating and Upgrading VMware NSX Intelligence* document for version 3.2 or later at https://docs.vmware.com/en/VMware-NSX-Intelligence/index.html. For a production environment, this feature requires that you select the Advanced form factor during the NSX Application Platform deployment.

Hosted NSX Feature	Description
VMware NSX® Network Detection and Response™	The NSX Network Detection and Response feature sends threat alert data to the VMware NSX® Advanced Threat Prevention cloud services, which then performs correlation and visualization on those data using the NSX Network Detection and Response user interface. For activation information, see the *VMware NSX Network Detection and Response Activation and Administration Guide* at https://docs.vmware.com/en/VMware-NSX-Intelligence/index.html.
VMware NSX® Malware Prevention	This feature detects and prevents malicious files (malware) from entering into your NSX-T Data Center environment and from spreading laterally across the data center. It uses NSX Advanced Threat Prevention cloud services to fetch periodic detection updates and to upload the data for further analysis. For activation information, see the NSX Malware Prevention information in the *NSX-T Data Center Administration Guide* for version 3.2 or later at https://docs.vmware.com/en/VMware-NSX-T-Data-Center/index.html.

Managing the NSX Application Platform 5

After deploying the NSX Application Platform, you can monitor the resources and services it uses, upgrade it to a newer version, add more resources to scale out the platform, increase the data storage it uses, obtain support bundles, and perform other supported administrative tasks.

To manage the NSX Application Platform, use the procedures described in this section.

- Check NSX Application Platform Status

  You can monitor the status of the deployed NSX Application Platform using the NSX Manager UI.

- Manage the NSX Application Platform Alarms

  The NSX Application Platform sends alarm notifications to alert you about specific events that might require your immediate attention. You can monitor alarms that are active for the NSX Application Platform using the NSX Manager UI.

- Managing Certificates in the NSX Application Platform

  Using the NSX Manager user interface (UI) or command line interface (CLI), you can manage certain certificates used by the NSX Application Platform or the NSX features that the platform hosts.

- Collect the Support Bundles for the NSX Application Platform Using the UI

  If you encounter an issue while using the NSX Application Platform and must provide the VMware support team with a support bundle, collect the bundle using the NSX Manager UI. You can download the bundle to your local system or upload them to a remote file server.

- Manage the Shut Down and Start Up of the vSphere with Tanzu Workload Domain

  If your vSphere with Tanzu Workload Domain is required to be shut down or started up, ensure that your infrastructure administrator performs these operations in the specified order to avoid the risk of any inconsistency or data loss.

- Upgrade the NSX Application Platform

  You can upgrade the NSX Application Platform to a later build version using the NSX Manager UI.

- Scale Out the NSX Application Platform

  If the resources that are currently allocated for the NSX Application Platform reach the threshold values set by the system, the system generates an alarm. To accommodate the needs of the NSX Application Platform core services, you must scale out the platform.

- Manage the NSX Application Platform Persistent Data Storage

  The NSX Application Platform persistent data storage retains the data collected and generated by the NSX Application Platform even when the platform is offline. You can increase the persistent volume of the data storage used by the NSX Application Platform as the storage needs of the platform services increase.

- Update the Form Factor Used in the NSX Application Platform

  You can update the form factor that you selected when you last deployed the NSX Application Platform. You can only update the form factor from a Standard form factor to an Advanced form factor.

- Update the NSX Application Platform Settings

  You can modify some of the settings that you used when you initially deployed and configured the NSX Application Platform.

- Delete the NSX Application Platform

  If for some reason you must redeploy the NSX Application Platform or free up the cluster that it uses, you must first delete any existing platform deployment you already have.

Check NSX Application Platform Status

You can monitor the status of the deployed NSX Application Platform using the NSX Manager UI.

Prerequisites

You must have Enterprise Admin privileges.

Procedure

1 From your browser, log in with Enterprise Admin privileges to an NSX Manager at `https://<nsx-manager-ip-address>`.

2 Navigate to **System > NSX Application Platform**.

3 Locate the **Platform** status section that is highlighted with a red box in the following image.

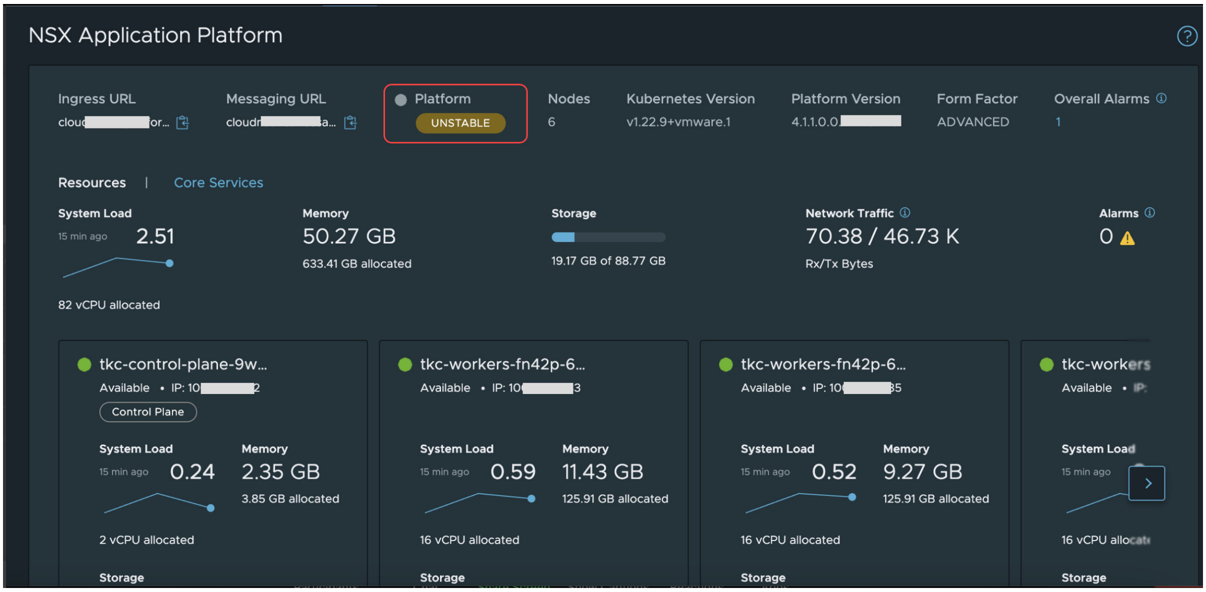

NSX Application Platform

Ingress URL
clou[redacted]or...

Messaging URL
cloudr[redacted]a...

Platform
UNSTABLE

Nodes
6

Kubernetes Version
v1.22.9+vmware.1

Platform Version
4.1.1.0.0[redacted]

Form Factor
ADVANCED

Overall Alarms
1

Resources | Core Services

System Load
15 min ago 2.51
82 vCPU allocated

Memory
50.27 GB
633.41 GB allocated

Storage
19.17 GB of 88.77 GB

Network Traffic
70.38 / 46.73 K
Rx/Tx Bytes

Alarms
0

tkc-control-plane-9w...
Available • IP: 10[redacted]2
Control Plane
System Load
15 min ago 0.24
2 vCPU allocated
Memory
2.35 GB
3.85 GB allocated
Storage

tkc-workers-fn42p-6...
Available • IP: 10[redacted]3
System Load
15 min ago 0.59
16 vCPU allocated
Memory
11.43 GB
125.91 GB allocated
Storage

tkc-workers-fn42p-6...
Available • IP: 10[redacted]5
System Load
15 min ago 0.52
16 vCPU allocated
Memory
9.27 GB
125.91 GB allocated
Storage

tkc-workers
Available • IP:
System Load
15 min ago
16 vCPU allocati
Storage

The following table lists the possible `Status` that might appear on the NSX Application Platform page. Use the information in the Description column to help you understand what the displayed `Status` indicates and, if necessary, what possible actions you can take to try to resolve any issues.

Status	Description
`Stable`	All NSX Application Platform services are up and healthy. Resources like CPU, disk, and memory are not constrained, network connectivity is available, and there are no delays in data processing.
`Unstable`	One or more NSX Application Platform services are not fully operational and the platform is exhibiting partial functionality. Although the services are running and are partially accessible, they are not performing optimally. The possible causes for this status include configuration issues, resource limitations, or problems with the underlying components. To address these issues, investigate the relevant error messages or log files and rectify the root causes. This might involve adjusting configurations, allocating additional resources, or troubleshooting malfunctioning components. Check if any alarm notifications exist. See Manage the NSX Application Platform Alarms for details. For additional information that might assist your investigation, see the topics in Chapter 5 Managing the NSX Application Platform or Chapter 6 Troubleshooting Issues with the NSX Application Platform.

Status	Description
`Down`	None of the NSX Application Platform services are functioning. The entire system is completely unavailable or offline. This status typically occurs due to critical system issues, failures, or planned maintenance-related downtime. To resolve this status, investigate the cause of the service outage and take appropriate actions to restore functionality. Check if any alarm notifications exist. See Manage the NSX Application Platform Alarms for details. For additional information that might assist your investigation, see the topics in Chapter 5 Managing the NSX Application Platform or Chapter 6 Troubleshooting Issues with the NSX Application Platform.
`Degraded`	This indicates that one or more NSX Application Platform services or resources are down, resulting in reduced performance or functionality. Check if any alarm notifications exist. See Manage the NSX Application Platform Alarms for details. For additional information that might assist your investigation, see the topics in Chapter 5 Managing the NSX Application Platform or Chapter 6 Troubleshooting Issues with the NSX Application Platform.

Manage the NSX Application Platform Alarms

The NSX Application Platform sends alarm notifications to alert you about specific events that might require your immediate attention. You can monitor alarms that are active for the NSX Application Platform using the NSX Manager UI.

The system automatically resolves some of the alarms it detects. To manage the alarms that require your attention, use the following steps.

Prerequisites

To manage the alarms, you must have Enterprise Admin privileges.

Procedure

1 From your browser, log in with Enterprise Admin privileges to an NSX Manager at https://<nsx-manager-ip-address>.

2 Navigate to **System > NSX Application Platform**.

3 To see the total number of alarms detected, locate the **Overall Alarms** section in the upper-right side of the NSX Application Platform pane.

   If alarms exist, the displayed number link indicates the total number of detected alarms in the resources the NSX Application Platform uses or the core services it supports.

4 (Optional) To view the open alarms specific to the resources or core services, click the **Resources** or the **Core Services** tab on the left side of the pane.

   Alarm links also appear for each individual resource or core service that has an existing open alarm.

5 Review the active alarms.

a To view the list of the active alarms, click the alarm links provided in the UI.

The system displays the corresponding dialog box depending on which alarm link you click. For example, the **Resources Alarms** dialog box displays only the open alarms detected in the resources that the NSX Application Platform uses.

b To display the details for an alarm, expand the row for the alarm in the dialog box.

You can review the description of the alarm type, severity level, times it was reported, and the recommended action to resolve the alarm.

c To close the dialog box, click **Close**.

6 To take one of the available actions you can apply to an `Open` alarm, navigate to **Home > Alarms** page.

a To narrow the list of alarms displayed, use the filter for the NSX Application Platform Communication and NSX Application Platform Health features.

b Locate the alarm that you want to manage and select the check box in the leftmost column.

c Click ⋮ and select one of the following actions that you want to apply to the selected alarm.

You can move an alarm to one of the following states depending on the actions allowed for the current alarm state.

Action	Description
Open	The alarm is placed in an active, but unacknowledged state.
Acknowledge	The alarm is acknowledged, but it remains open. Its **Last Reported Time** value continues to be updated until you move the alarm to another state.
Suppress	When you suppress an alarm, the system prompts you to specify how many hours you want the alarm suppressed. If after the specified suppression time is reached and the alarm condition remains the same, the alarm state is returned to the `Open` status. If during the suppression period the alarm condition is resolved, the system automatically changes the alarm state to `Resolved`.
Resolve	You can manually resolve an alarm. Once you manually resolved an alarm, you can no longer change its state. If you resolve an alarm manually, but the problem persists, the system opens another similar alarm. The system continuously monitors the NSX Application Platform and can auto-resolve an alarm.

The system updates the **Alarm State** value for the alarm.

Managing Certificates in the NSX Application Platform

Using the NSX Manager user interface (UI) or command line interface (CLI), you can manage certain certificates used by the NSX Application Platform or the NSX features that the platform hosts.

Manage the CA-Signed Certificates Used in the NSX Application Platform

When the NSX Application Platform is deployed, a default self-signed certificate is used. You can replace that default certificate with a CA-signed certificate and assign it to the NSX Application Platform.

You can replace the default self-signed certificate with either a CA-signed certificate with a private key or a CA-signed certificate with a CSR. When the certificate is being imported, all the services used by the NSX Application Platform become unavailable.

Prerequisites

- You must have Enterprise Admin account privileges.
- Ensure that no active alarms exist on the NSX Application Platform.
- Verify that the NSX Data Center license in effect meets the minimum required.
- Verify that you have a valid certificate with a private key or a certificate with a certificate signing request (CSR). You must generate the CSR using the NSX Manager UI, as described in the following steps.

Procedure

1. From your browser, log in with Enterprise Admin privileges to an NSX Manager at https://<nsx-manager-ip-address>.
2. Navigate to **System > NSX Application Platform**.
3. In the bottom-left corner of the NSX Application Platform section, click **Actions** and select **Manage CA Certificate** from the drop-down menu.
4. If you are importing a CA-signed certificate with a private key, use the following steps.
   a. In the **Certificate with Private Key** tab, click **Import**.
   b. Enter a name for the certificate.
   c. Enter the contents of the certificate in the **Certificate Contents** text box. Click **Browse**, navigate to the location of the CA-signed certificate, and copy its contents.
   d. Enter the private key in the **Private key** text box, by clicking **Browse** and navigating to where you stored the private key. Copy its contents in the **Private key** text box.
   e. Click **Import**.

      The CA-signed certificate is imported and assigned to the NSX Application Platform.

5 If you are importing a CA-signed certificate with CSR, use the following steps.

   a In the **Certificate with CSR** tab, click **Generate CSR**.

   b In the Generate CSR dialog box, enter the required information and click **Generate**.

      If for some reason you want to regenerate the CSR, click the delete icon next to the name of the existing CSR and click **Generate CSR** again.

   c After the CSR is generated, click **Download CSR PEM** and submit the generated CSR form to a Certificate Authority (CA).

      The CA must sign and return a full chain certificate.

   d After you receive the CA-signed digital identity full chain certificate, click **Import** in the Certificate with CSR section of the **Manage CA Certificate** dialog box.

   e Click **Import**.

      The CA-signed certificate with CSR is imported and assigned to the NSX Application Platform.

6 If you want to use another CA-signed certificate, click **Replace**.

7 If you want to delete the CA-signed certificate, click the delete icon (trash can). When prompted, confirm that you want the CA-signed certificate deleted.

   A new default self-signed certificate is assigned by the system.

Manage the Kafka Messaging Client Certificate

The Kafka messaging client self-signed certificate is used by the NSX Intelligence common agent that is running in the NSX Manager unified appliance. It only supports RSA encryption. You can replace that default self-signed certificate with a CA-signed certificate.

You can replace the default self-signed certificate with either a CA-signed certificate with a private key or a CA-signed certificate with a CSR. When the certificate is being imported, the Kafka messaging service used by the NSX Manager unified appliance becomes unavailable.

Prerequisites

- You must have Enterprise Admin account privileges.
- Ensure that no active alarms exist on the NSX Application Platform.
- Verify that you have a valid certificate with a private key or a certificate with a certificate signing request (CSR). You must generate the CSR using the NSX Manager UI.

Procedure

1 Log into the NSX Manager appliance as with an Enterprise Admin user account.

2 Import the CA-signed certificate on the NSX Manager by running the following command at the system command prompt.

```
curl -v -H "Content-Type: application/json" -ku 'username:password' 'https://<manager-ip-
aeddress>/api/v1/trust-management/certificates?action=import' -d
{
     "pem_encoded": "xxx",
     "private_key": "yyyy"
}
```

3 Apply the CA-signed certificate as the Kafka messaging client certificate by running the following command.

```
curl -v -H "Content-Type: application/json"
-ku 'username:password' 'https://<manager-ip>/api/v1/trust-management/certificates/
<certificate-id>?action=apply_certificate&service_type=K8S_MSG_CLIENT'
```

Collect the Support Bundles for the NSX Application Platform Using the UI

If you encounter an issue while using the NSX Application Platform and must provide the VMware support team with a support bundle, collect the bundle using the NSX Manager UI. You can download the bundle to your local system or upload them to a remote file server.

Prerequisites

- You must have Enterprise Admin account privileges.
- The NSX Application Platform must be successfully deployed.

Procedure

1 From your browser, log in with Enterprise Admin privileges to an NSX Manager at https://<nsx-manager-ip-address>.

2 Navigate to **System > Support Bundle**.

3 In the **Request Bundle** tab, select **NSX Application Platform** from the **Type** drop-down menu.

4 In the **Available** pane, select the **Available** check box if you want support bundles collected for all the services.

Optionally select specific the NSX Application Platform services for which you want a support bundle created.

5 To move the selected nodes to the **Selected** pane, click the **>** icon.

6 In the **Log age (days)** text box, keep the default value of **`All`** or enter the specific number of days' worth of logs that you want included in the support bundle.

7 (Optional) To specify that you want the core files and audit log files to be included in the support bundle, click the **Include core files and audit logs** toggle to **Yes**.

Ensure that you read the information under the toggle and understand what it means when you include or exclude the core files and audit logs.

8 (Optional) Select the **Upload bundle to a remote file server** check box if you want to upload the support bundle to a remote file server.

a Provide the remote file server's IP address or host name; port, and protocol to use.

b Enter the user name and password for the remote file server.

c Enter the absolute destination path to where the bundle is to be uploaded in the remote server.

d If you want the NSX Manager to upload the bundle, toggle **Manager upload** to **Enabled**.

9 Click **Start Bundle Collection**.

10 Monitor the status of the bundle collection process.

The **Status** page shows the progress of the support bundle collection. When the collection has finished successfully, the size of the bundle is displayed next to **Support Bundle**. The **Details** table displays the info about all the support bundles that completed successfully or did not finish.

11 To store the support bundle to a local folder, click **Download**.

If you selected the **Upload the bundle to a remote file server** check box earlier, the support bundle is uploaded to the file server you had specified.

What to do next

For information about the support bundle files, see Understanding NSX Application Platform Support Bundle File Paths.

Understanding NSX Application Platform Support Bundle File Paths

The NSX Application Platform support bundles include log information for the features that NSX Application Platform hosts. These supported features are NSX Intelligence, NSX Network Detection and Response, NSX Malware Prevention, and NSX Metrics.

The following table lists each of the support log file paths included in the bundle and the NSX Application Platform-hosted feature that uses it.

File Path	Features Using the File Path
/var/log/napp/supportbundle/platform_service_dbg.log	NSX Application Platform
/var/log/napp/supportbundle/<K8s_worker_node_name>/kafka*.log	NSX Application Platform
/var/log/napp/supportbundle/<K8s_worker_node_name>/zookeeper*.log	NSX Application Platform

File Path	Features Using the File Path
/var/log/napp/supportbundle/ <K8s_worker_node_name>/ spark-app-rawflow-driver*.log	NSX Application Platform
/var/log/napp/supportbundle/ <K8s_worker_node_name>/authserver*.log	NSX Application Platform
/var/log/napp/supportbundle/ <K8s_worker_node_name>/cluster-api*.log	NSX Application Platform
/var/log/napp/supportbundle/ <K8s_worker_node_name>/common-agent*.log	NSX Application Platform
/var/log/napp/supportbundle/ <K8s_worker_node_name>/projectcontour*.log	NSX Application Platform
/var/log/napp/supportbundle/ <K8s_worker_node_name>/routing-controller*.log	NSX Application Platform
/var/log/napp/supportbundle/ <K8s_worker_node_name>/trust-manager*.log	NSX Application Platform
/var/log/napp/supportbundle/ <K8s_worker_node_name>/upgrade-coordinator*.log	NSX Application Platform
/var/log/napp/supportbundle/ <K8s_worker_node_name>/postgresql*.log	NSX Application Platform
/var/log/napp/supportbundle/ <K8s_worker_node_name>/visualization*.log	NSX Intelligence Visualization
/var/log/napp/supportbundle/ <K8s_worker_node_name>/recommendation*.log	NSX Intelligence Recommendation
/var/log/napp/supportbundle/ <K8s_worker_node_name>/rawflowcorrelator*.log	NSX Intelligence
/var/log/napp/supportbundle/ <K8s_worker_node_name>/ overflowcorrelator*.log	NSX Intelligence
/var/log/napp/supportbundle/ <K8s_worker_node_name>/*druid*.log	NSX Intelligence
/var/log/napp/supportbundle/ <K8s_worker_node_name>/nsx-config*.log	NSX Intelligence
/var/log/napp/supportbundle/ <K8s_worker_node_name>/redis*.log	NSX Intelligence
/var/log/napp/supportbundle/ <K8s_worker_node_name>/minio*.log	NSX Intelligence
/var/log/napp/supportbundle/ <K8s_worker_node_name>/nsx-ndr*.log	NSX Network Detection and Response
/var/log/napp/supportbundle/ <K8s_worker_node_name>/sa-*.log	NSX Malware Prevention

File Path	Features Using the File Path
/var/log/napp/supportbundle/<K8s_worker_node_name>/malware*.log	NSX Malware Prevention
/var/log/napp/supportbundle/<K8s_worker_node_name>/metrics*.log	NSX Metrics

Manage the Shut Down and Start Up of the vSphere with Tanzu Workload Domain

If your vSphere with Tanzu Workload Domain is required to be shut down or started up, ensure that your infrastructure administrator performs these operations in the specified order to avoid the risk of any inconsistency or data loss.

The shutdown and startup operations in your vSphere with Tanzu environment must be performed in a specified order. Follow the information provided in the "Shut Down and Start Up the vSphere with Tanzu Workload Domain" topic in the *vSphere with Tanzu Configuration and Management* documentation.

Note NSX Application Platform environments running on an Upstream Kubernetes cluster follow a similar sequence of operations (without the Tanzu components).

Upgrade the NSX Application Platform

You can upgrade the NSX Application Platform to a later build version using the NSX Manager UI.

The upgrade process retains the form factor being used for the current platform deployment. If you deployed the platform using an Evaluation form factor and want to continue using an Evaluation form factor, but use a later platform version, you must delete your current NSX Application Platform deployment first. You then redeploy the platform using a later platform version.

Upgrading to a newer version of the NSX Application Platform involves multiple steps. You must first configure and deploy the Upgrade Coordinator before you can proceed with upgrading the platform and each of the currently activated NSX-T Data Center features. The Upgrade Coordinator orchestrates all of the upgrade steps, and the system provides status on the UI as it upgrades each component.

Important Beginning with NSX-T Data Center 3.2.3.1 release, the VMware-hosted NSX Application Platform registry and repository locations no longer support ChartMuseum-compatible private container registry, such as Harbor.

If you are currently using an NSX-T Data Center version prior to 3.2.3.1, consider upgrading to NSX-T Data Center version 3.2.3.1 before attempting to upgrade your current NSX Application Platform deployment.

If you are unable to access the public VMware-hosted NSX Application Platform registry and repository locations, work with your infrastructure administrator to upload the NSX Application Platform Helm chart and Docker images to your company's private Helm repository and Docker registry locations. Both privately-hosted locations must be accessible from the Kubernetes cluster and the NSX Manager appliance you are using for the NSX Application Platform deployment. If you need to continue using the ChartMuseum provided by Harbor, the Harbor version must be earlier than version 2.8.1. See Upload the NSX Application Platform Docker Images and Helm Charts to a Private Container Registry for details.

Prerequisites

- Review the NSX-T Data Center Release Notes for any known upgrade issue and workaround documented for the NSX Application Platform.
- Ensure that there are no open alarms detected on the NSX Application Platform.
- Verify that you have met all of the prerequisites and system requirements listed in NSX Application Platform Deployment Prerequisites.
- You must have Enterprise Admin privileges.

Procedure

1 From your browser, log in with Enterprise Admin privileges to an NSX Manager at `https://<nsx-manager-ip-address>`.

2 Navigate to **System > Upgrade**.

3 Deploy the Upgrade Coordinator.

a In the NSX Application Platform card, click **Upgrade**.

This step can take some time as the system obtains the information from the VMware-hosted Helm repository. When the information is obtained successfully, the **Prepare for Upgrade** tab is displayed with the text boxes prepopulated with the information obtained for **Helm Repository**, **Docker Registry**, and **Platform Target Version**, as illustrated in the following image. Note that the values shown in the image are the default values for the VMware-hosted Helm repository and Docker registry locations, and the latest available NSX Application Platform version.

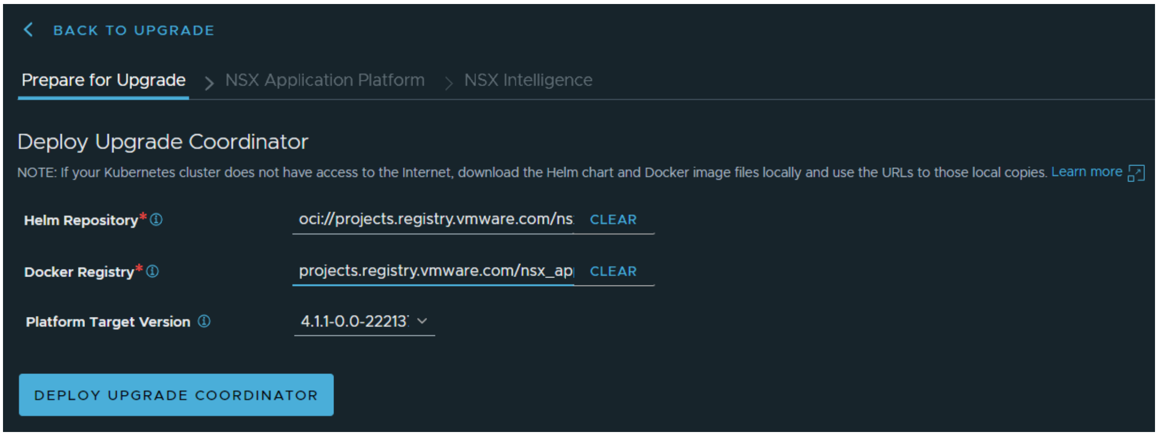
< BACK TO UPGRADE
Prepare for Upgrade > NSX Application Platform > NSX Intelligence

Deploy Upgrade Coordinator

NOTE: If your Kubernetes cluster does not have access to the Internet, download the Helm chart and Docker image files locally and use the URLs to those local copies. Learn more

Helm Repository* ⓘ oci://projects.registry.vmware.com/ns CLEAR

Docker Registry* ⓘ projects.registry.vmware.com/nsx_ap CLEAR

Platform Target Version ⓘ 4.1.1-0.0-22213

DEPLOY UPGRADE COORDINATOR

b (Optional) If you are using a private Helm repository and Docker registry location, provide the private locations of the required Helm charts and Docker images.

Use the following information for the **Helm Repository** text box.

- For OCI-compatible Helm private repository, use the format `oci://`*`<your-private-registry-server-fqdn>`*`/<your-private-registry-name>/helm-charts`.
- For ChartMuseum-compatible Helm private repository, use the format `https://`*`<your-private-registry-server-fqdn>`*`/chartrepo/`*`<your-private-registry-name>`*.

For the **Docker Registry** text box, use the format *`<your-private-registry-server-fqdn>`*`/`*`<your-private-registry-name>`*`/clustering`.

Click **Save URL**.

c In the **Platform Target Version** text box, verify that the build version that you want to use for the upgrade is selected.

d Click **Deploy Upgrade Coordinator**.

This step can also take some time as the system deploys the Upgrade Coordinator to your TKG Cluster on Supervisor pod or upstream Kubernetes pod.

After the Upgrade Coordinator deployment completes, the **Prepare** tab is displayed. The **Status** section displays the `Success` status.

4 In the **Deploy Upgrade Coordinator** section located in the upper half of the **Prepare** tab, verify that the values shown in **Helm Repository**, **Docker Registry**, and **Platform Target Version** text boxes are correct.

   If you must modify any of the values, click **Delete** next to the **Note** located after the **Status** section and redeploy a new Upgrade Coordinator.

5 Review the **Summary** section located in the lower half of the **Prepare** tab.

   The **NSX Application Platform** card displays the status information for the platform. The `Upgrade Completed` indicates that the Upgrade Coordinator has been upgraded with the target NSX Application Platform version successfully. The card shows the current version and the target version to which the platform will be upgraded. The card also shows the precheck status.

   If other NSX-T Data Center features that are hosted on the **NSX Application Platform** are activated, those features are also checked and scheduled for the upgrade. A separate feature card for each activated feature is also displayed. For example, the NSX Intelligence feature card appears in the following image because it is currently an activated NSX-T Data Centerfeature on the NSX Application Platform. The system upgrades the activated features after the platform upgrade finishes successfully.

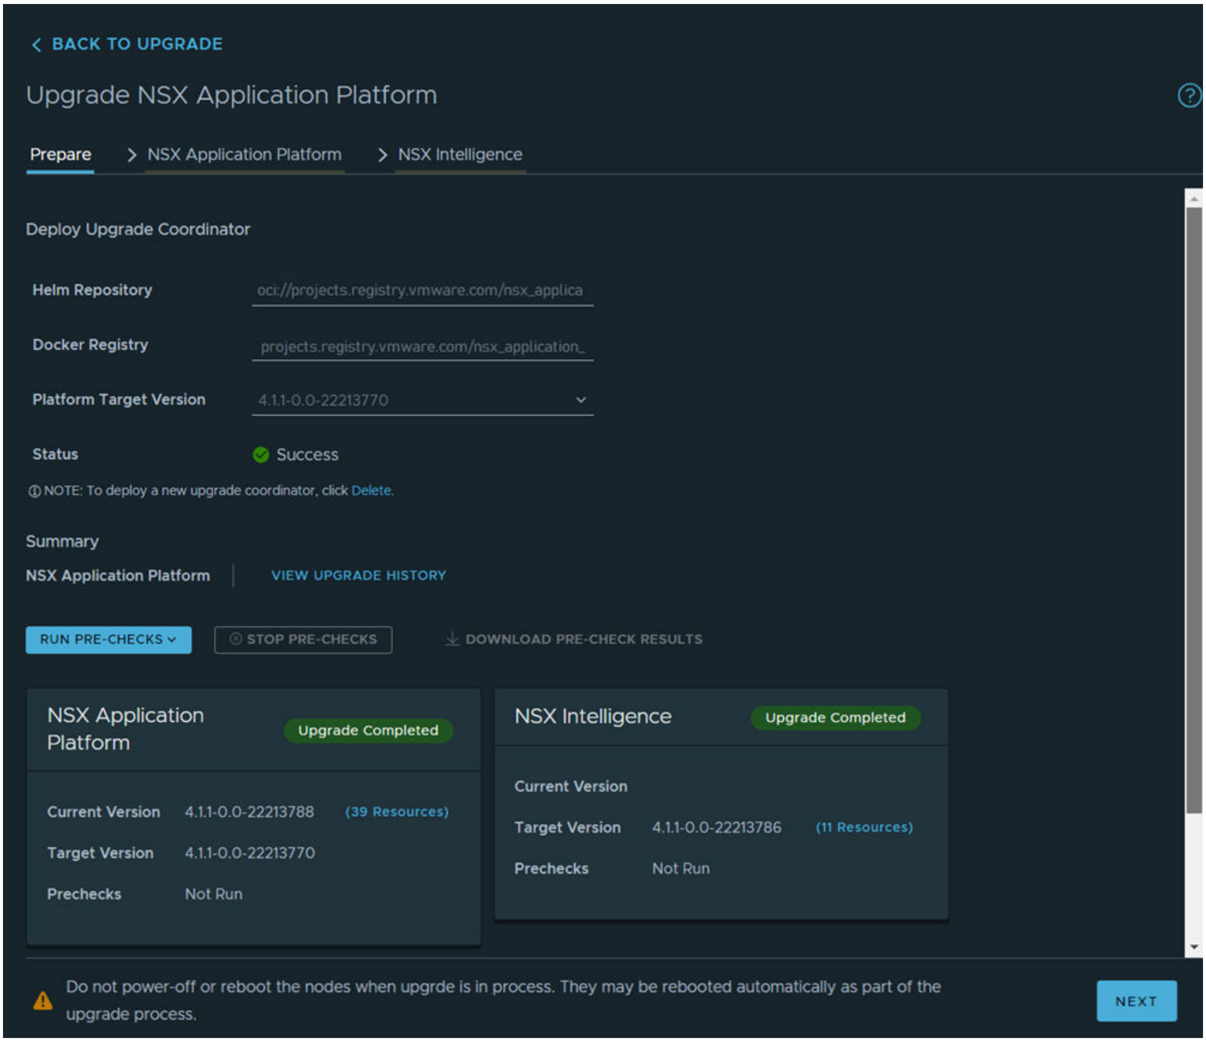

Upgrade NSX Application Platform

Deploy Upgrade Coordinator

Helm Repository	oci://projects.registry.vmware.com/nsx_applica
Docker Registry	projects.registry.vmware.com/nsx_application_
Platform Target Version	4.1.1-0.0-22213770
Status	✓ Success

ⓘ NOTE: To deploy a new upgrade coordinator, click Delete.

Summary

NSX Application Platform

NSX Application Platform — Upgrade Completed

Current Version	4.1.1-0.0-22213788	(39 Resources)
Target Version	4.1.1-0.0-22213770	
Prechecks	Not Run	

NSX Intelligence — Upgrade Completed

Current Version		
Target Version	4.1.1-0.0-22213786	(11 Resources)
Prechecks	Not Run	

⚠ Do not power-off or reboot the nodes when upgrde is in process. They may be rebooted automatically as part of the upgrade process.

6 If all the Upgrade Coordinator values are correct, click **Run Pre-Checks** and select **All Prechecks** from the drop-down menu.

To optionally precheck specific components only, click **Run Prechecks** and from the drop-down menu, select the name of the component that you want to precheck.

The system performs all the prechecks for all the components that are scheduled for the upgrade. The prechecks help detect and resolve potential problems early in the upgrade process, which can make the upgrade process run more smoothly. The system updates the component cards with their prechecks status.

If the system identifies any issues during the precheck, you can click **Download Pre-check Results** and use the information in the downloaded file to help investigate the reported issues.

7 Click **Next**.

The **NSX Application Platform** tab displays a grid of all the groups of components that comprise the platform. You can expand each row to see all of the units for each component group that will be upgraded.

8 Click **Upgrade**.

The system upgrades each group that comprise the NSX Application Platform. This step can take some time to finish. You can leave the **Upgrade** UI screen and return to it by navigating back to **System > Upgrade** page and clicking **Continue With Upgrade**.

There are multiple ways to track the progress of the upgrade.

a To view the logs generated as the upgraded progresses, click **Recent Logs**.

b To monitor the upgrade status for each component group, use the **Group Status** column.

c To see the upgrade status for each item in a particular group, expand the grid row for that group and verify the status shown for each group item.

If an error occurs for a group upgrade, expand the row for the group and click the **Failed** link to see the reason for the failure. Use that information to resolve the reported problem and to work with your infrastructure administrator or VMware support. When you have resolved the cause of the failure, click **Retry** to try to complete the upgrade.

9 When the NSX Application Platform is successfully upgraded, click **Next** and in the tab for the NSX-T Data Center feature (for example, NSX Intelligence), click **Update**.

After this feature is successfully upgraded, repeat this step for each of the remaining NSX-T Data Center features that must be upgraded.

10 After you have upgraded all of the NSX-T Data Center features activated on the NSX Application Platform, navigate to **System > NSX Application Platform**. Verify the **Platform Version** and **Feature Version** details for each of the activated features are correct.

Scale Out the NSX Application Platform

If the resources that are currently allocated for the NSX Application Platform reach the threshold values set by the system, the system generates an alarm. To accommodate the needs of the NSX Application Platform core services, you must scale out the platform.

Use the following table to determine the minimum number of nodes required to scale out the core services and to learn what improvements are obtained after successfully scaling out the core services.

Core Service Name	Minimum # of nodes required for a Scale Out operation	What improves after scaling out the service
Messaging	At least five nodes	Scaling out the Messaging service increases messaging broker instance in your TKG Cluster on Supervisor or upstream Kubernetes cluster.
Analytics	At least four nodes	Scaling out the Analytics service increases the overall data pipeline processing throughput of the NSX Intelligence features.

Core Service Name	Minimum # of nodes required for a Scale Out operation	What improves after scaling out the service
Data Storage	At least eight nodes	Scaling out Data Storage service adds more object storage broker instances in your TKG Cluster on Supervisor or upstream Kubernetes cluster.
Metrics	At least three nodes	Scaling out the Metrics service increases the metrics data processing throughput of the NSX Metrics feature.

Caution Select the **All** check box in the **Scale Out** dialog box and let the system decide which of the cores services must be scaled out. Arbitrarily scaling out a core service category might lead to more resources being used without achieving any performance improvement. Consult the VMware support team before scaling out a single core service category.

Prerequisites

- All existing nodes in your TKG Cluster on Supervisor or upstream Kubernetes cluster must be in a healthy and ready state before you can scale out the NSX Application Platform.
- Before proceeding with the scale-out procedures, ensure that your infrastructure administrator has already allocated the minimum number of nodes required for scaling out the NSX Application Platform services.

Procedure

1. From your browser, log in with Enterprise Admin privileges to an NSX Manager at https://<nsx-manager-ip-address>.
2. Navigate to **System > NSX Application Platform**.
3. In the bottom-left corner of the NSX Application Platform section of the UI page, click **Actions** and select **Scale Out** from the drop-down menu.

   Note The **Scale Out** action is only supported if you deployed the NSX Application Platform using the Advanced form factor. The action is not supported for Standard or Evaluation form factor deployments.

   If all of the services are scaled out already, the **Scale Out** button is disabled on the drop-down menu. In this case, it indicates that your cluster nodes have reached the maximum number of nodes allocated. You must first request for your infrastructure administrator to add five more nodes to your current cluster before you can continue with the next steps. To scale out all of the services, you must have a total of eight worker nodes in your cluster.
4. Select the **All** checkbox.

5 In the **Advanced Options** section, ensure that all of the services available for the scale-out action are selected.

   Unless specifically advised by the VMware support team, ensure that all of the core services are selected so that the system can decide which of the core services must be scaled out. Scaling out one core service arbitrarily can lead to more resources being used without any improvement to the system performance. Before proceeding with single-category service scale out procedure, consult the VMware support team or confirm that you know clearly what can happen if you scale out a single-category service.

6 Click **Scale Out**.

   The UI displays the progress of the scale out operation.

Results

If there are no errors encountered, the message banner with `Scale out completed successfully!` is displayed.

Manage the NSX Application Platform Persistent Data Storage

The NSX Application Platform persistent data storage retains the data collected and generated by the NSX Application Platform even when the platform is offline. You can increase the persistent volume of the data storage used by the NSX Application Platform as the storage needs of the platform services increase.

The system will generate an alarm for the Data Storage core service if it has crossed the threshold value.

Important You cannot request to increase the data storage using a volume size that is smaller than what was specified in the last failed request to increase the data storage. You can only specify a volume size that is equal or larger than the last requested volume size. If your new request uses a smaller volume size, the system will use the previously requested volume size again to process your request.

Prerequisites

- You must have Enterprise Admin privileges.
- Ensure that the cluster nodes that the NSX Application Platform uses are in a healthy state and that no open alarms exist.
- Obtain from your infrastructure administrator the number of available disk pools and verify the current limit of the data storage volume. If necessary, request for an increase in the number of disks used by your TKG Cluster on Supervisor or upstream Kubernetes cluster nodes.

- Ensure that the TKG Cluster on Supervisor or upstream Kubernetes cluster that the NSX Application Platform uses support volume resizing. If the cluster does not support volume expansion, work with your infrastructure administrator to enable volume expansion. Without volume resizing support, the scale-up operation can fail, but it will not affect the existing cluster that the NSX Application Platform is already using.

Procedure

1. From your browser, log in with Enterprise Admin privileges to an NSX Manager at https://<nsx-manager-ip-address>.
2. Navigate to **System > NSX Application Platform** and click **Core Services**.
3. In the bottom-left corner of the **Data Storage** section, click **Manage Data Storage**.
4. In the **Manage Data Storage** dialog box, enter the new data storage value.

   This value is the number of available disk pools that you obtained from your infrastructure administrator.
5. Click **Update**.

   The system proceeds to update the volume size of each persistent storage used by the NSX Application Platform core services.

Update the Form Factor Used in the NSX Application Platform

You can update the form factor that you selected when you last deployed the NSX Application Platform. You can only update the form factor from a Standard form factor to an Advanced form factor.

Note To update from using an Evaluation form factor, you must first delete your evaluation NSX Application Platform deployment and redeploy the NSX Application Platform to use either the Standard or Advanced form factor.

Prerequisites

- You must have Enterprise Admin privileges.
- Ensure that no active alarms exist on the NSX Application Platform.
- Verify that the minimum number of nodes required for the Advanced form factor has been configured by your infrastructure administrator. See NSX Application Platform System Requirements for more information.

Procedure

1. From your browser, log in with Enterprise Admin privileges to an NSX Manager at https://<nsx-manager-ip-address>.
2. Navigate to **System > NSX Application Platform**.

3 In the bottom-left corner of the NSX Application Platform section, click **Actions** and select **Update Form Factor** from the drop-down menu.

4 Review the information provided in the **Update Form Factor** dialog box and perform any additional tasks before proceeding to the next steps.

5 Click **Go to Settings**.

6 In the NSX Application Platform Settings page, navigate to the **Form Factor** section and select the new form factor. For example, select **Advanced**.

7 Click **Next**.

8 In the **Precheck Platform** tab, click **Run Prechecks**.

The system updates precheck table with the status of each precheck performed. If the system encounters any issues, it provides error or warning links in the **Details** column. Review the information provided for the error and address them, if required, before continuing with the form factor update.

9 After the system completes all the prechecks successfully, click **Next**.

10 In the **Review & Update** tab, review the details and if everything looks correct, click **Update**.

The system displays the progress bar while updating the form factor. If an error occurs, the system displays an error message with more information.

If there are any errors, take the necessary actions to correct the issue and retry updating the form factor again.

11 Click **Close** when the system displays a message when the form factor update completed successfully.

Results

In the **Advanced in NSX Application** pane, the **Form Factor** value now displays Advanced.

Update the NSX Application Platform Settings

You can modify some of the settings that you used when you initially deployed and configured the NSX Application Platform.

The following lists examples of platform settings that you might have to update.

- The Kubernetes configuration file, if you encountered issues connecting to the TKG Cluster on Supervisor or upstream Kubernetes cluster used by the platform.
- The Kubernetes Tools, if you receive an error message that there are compatibility issues with the TKG Cluster on Supervisor or upstream Kubernetes cluster.
- The Helm chart and Docker registry URLs, if you are using your own private Helm repository and Docker registry repository.

Prerequisites

- You must have Enterprise Admin privileges.
- Ensure that no active alarms exist on the NSX Application Platform.
- Ensure that the NSX Data Center license in effect meets the minimum required.

Procedure

1. From your browser, log in with Enterprise Admin privileges to an NSX Manager at https://<nsx-manager-ip-address> that is associated with your NSX Application Platform deployment.
2. Navigate to **System > NSX Application Platform**.
3. In the bottom-left corner of the NSX Application Platform section, click **Actions** and select **Settings** from the drop-down menu.
4. In the **Settings** page, click **Edit** for the Settings section (Platform, Configuration, or Prechecks) that you want to modify.
5. If you are editing the **Platform** section or **Configuration** section, make the changes you want in the values that are available for modification.

   Not all settings are available for you to modify in this workflow. When you point to values that are not editable, the tooltip shows `Edit not allowed` or you are unable to enter anything in the text box. For example, you cannot change the **Platform Target Version** value in the **Prepare to Deploy** tab.
6. Click **Next** and make any other settings that are available for modification on the subsequent tabs.
7. In the **Precheck Platform** tab, click **Run Prechecks** and then **Next** after all the prechecks completed successfully.
8. In the **Summary** tab, review the settings information and click **Save** if you made changes or **Close** if you did not.

Delete the NSX Application Platform

If for some reason you must redeploy the NSX Application Platform or free up the cluster that it uses, you must first delete any existing platform deployment you already have.

When you delete the NSX Application Platform, the platform gets deleted permanently.

Prerequisites

- You must have Enterprise Admin privileges.
- Delete all the NSX features that you activated on the NSX Application Platform.

Procedure

1. From your browser, log in with Enterprise Admin privileges to an NSX Manager at https://<nsx-manager-ip-address>.
2. Navigate to **System > NSX Application Platform > .**
3. Click **Actions** in the bottom-left corner of the NSX Application Platform section and select **Delete** from the pop-up menu.
4. Click **Yes, Delete** to confirm.

   The deletion process can take a while as the system deletes all the Kubernetes resources it created in your TKG Cluster on Supervisor or upstream Kubernetes cluster when you deployed the NSX Application Platform.

Results

If the deletion is successful, a message banner with `NSX Application Platform deleted successfully!` appears on the UI.

Troubleshooting Issues with the NSX Application Platform

6

This section provides information that might assist you in resolving problems you might encounter when deploying or managing the NSX Application Platform.

Read the following topics next:

NSX Application Platform Deployment Failed

The NSX Application Platform deployment operation failed due to a timeout.

Problem

An attempt to deploy the NSX Application Platform failed to complete and the operation timed out.

Cause

The NSX Application Platform deployment process includes the installation of several services used for Messaging, Data Storage, Analytics and more. In the event these services do not start within 20 minutes of deployment initiation, the deployment process is halted with a timeout error. Possible causes for the timeout error are infrastructure issues, such as CPU contention, Storage contention, or network connectivity issues.

Solution

1. Log in to the NSX Manager appliance with the root account.
2. Use the following command to verify that the `nsxi-platform` pods are deployed successfully.

   ```
   napp-k get pods --selector='app.kubernetes.io/instance=nsxi-platform'
   ```

   If the pod status is either `Running` or `Completed` that means the pod started successfully.

   If one of the following pod status continues to be returned, it can indicate a pod did not start successfully.

   - `Init/Podinitializing/CrashLoopBackOff`
   - `Pending`
   - `ContainerCreating`
   - `ImagePullBackOff`
   - `ErrImagePull`
   - `ImageInspectError`
   - `CreateContainerConfigError`
3. If any of the pods do not start within 20 minutes of deployment initiation and the deployment halts with a timeout error, investigate if there are any issues with the infrastructure, such as CPU contention, storage contention, or network connectivity issue. Resolve the issues and retry the NSX Application Platform deployment from the NSX Manager UI.

4. If the deployment continues to time out, use the following command to inspect the log files for each pod that fails to start up.

   Obtain the *<POD_NAME>* from the output of the command in step 2.

   ```
   napp-k logs <POD_NAME>
   ```

5. Resolve the errors reported in the log files and retry the NSX Application Platform deployment from the NSX Manager UI.
6. If the deployment continues to fail, contact your VMware support for assistance.

NSX Application Platform Upgrade Failed

The NSX Application Platform upgrade failed due to an error.

Problem

An attempt to upgrade the NSX Application Platform failed with the error `Deploying Upgrade Coordinator: Upgrade Coordinator post deployment plugin call failed`.

Cause

If the connectivity between the Kubernetes control plane node and an existing Pod is lost, the Kubernetes certificate manager is unable to successfully send the certificate to the existing Pod and the Egress API calls sent to the NSX Manager fail. This problem can also occur if there is a connectivity issue from the Update Coordinator node to the Kubernetes control plane node.

Solution

To resolve the issue, perform the following steps.

1. Log in to the NSX Manager appliance with the `root` account.
2. At the system prompt, run the following command to restart the pods.

   ```
   napp-k rollout restart deployment cluster-api trust-manager
   ```

3. Wait for the `cluster-api` and `trust-manager` pods to restart successfully.
4. Once the pods are back online, restart the Upgrade Coordinator deployment using the following command.

   ```
   napp-k rollout restart deployment upgrade-coordinator
   ```

5. Wait for the Upgrade Coordinator to restart successfully.
6. Retry the NSX Application Platform upgrade again using the **System > Upgrade** tab in the NSX Manager UI. See Upgrade the NSX Application Platform for details.

Unstable NSX Application Platform Upgrade UI

The Upgrade user interface can become unstable during an upgrade of the NSX Application Platform.

Problem

The NSX Application Platform upgrade UI becomes unstable for 15–16 minutes if you change the NSX Application Platform upgrade plan in between the NSX Application Platform upgrade and the upgrade of one of the hosted NSX-T Data Center applications. For example, if you start the upgrade of the NSX Application Platform and the NSX Malware Prevention feature application, but you uninstall the NSX Malware Prevention before it is upgraded, the NSX Application Platform upgrade UI becomes unstable.

Cause

During the upgrade process, Kubernetes restarts the NSX Application Platform Pod after eight minutes and the Upgrade Coordinator can take around 7–8 minutes to start up. During that time, the Upgrade user interface is unstable for about 16 minutes.

Solution

Wait for 15–16 minutes to restart the Kubernetes Pod and the Upgrade Coordinator.

Collect the NSX Application Platform Support Bundles Using the CLI

You are unable to collect the NSX Application Platform support bundles using the NSX Manager UI.

Problem

You attempted to deploy or delete the NSX Application Platform but it failed before finishing the deploy or delete operation. Since the NSX Application Platform is no longer accessible, there is no way to access the support bundles using the NSX Manager UI to find out what caused the operation to fail.

Cause

There can be multiple reasons as to why the deploy or delete operation failed.

Solution

To find out the reason for the deploy or delete operation failure, use the following information to collect the support bundles using the NSX Manager CLI.

1 Log into the NSX Manager appliance CLI as an admin user.

2. Before generating the support bundle, enter the following command at the CLI prompt.

   ```
   set napp kubeconfig
   ```

3. Enter the `kubeconfig` file content and press Ctrl+D.

   This action overrides any existing `kubeconfig` file used by the NSX Application Platform.

4. Get the support bundle file using the following command, where *support-bundle-filename* is the filename that you want the system to use for the generated support bundle file.

   ```
   get support-bundle file support-bundle-filename.tgz
   ```

   If you entered a valid `kubeconfig` file content, you can find the NSX Application Platform support log files in the `/napp` folder after you unpack the *`support-bundle-filename`*`.tgz` file.

   If the `kubeconfig` file content that you entered is invalid, the `/napp` folder is empty.

5. To transfer the support bundle file securely from the NSX Manager host to a remote host location, use the following command.

   ```
   copy file support-bundle-filename.tgz url remote-host-url
   ```

   The *`remote-host-url`* has the format `scp://`*`<username>@<ip-address>:<dir>`*, where *`<username>`* is a valid user account that has an access to the remote host, *`<ip-address>`* is the IP address of the remote host, and *`<dir>`* is the file folder to where the support bundle file is to be transferred.

Attempt to Add or Modify the Helm Repository Failed

An attempt to add the Helm repository information failed during the NSX Application Platform deployment preparation.

Problem

The system attempted to add the Helm repository URL information, but the operation failed. Before you can deploy the NSX Application Platform, the system must add or update the Helm repository information in the NSX Manager. The system fetches the Helm charts from the Helm repository and installs them on the specified TKG Cluster on Supervisor or upstream Kubernetes cluster.

Cause

The privately hosted Helm repository is not accessible.

Solution

Ensure that the Helm repository is present in your private network and is accessible from the NSX Manager appliance.

Attempt to Upload the `kubeconfig` File Failed

An attempt to upload the kubeconfig file failed during the NSX Application Platform deployment preparation.

Problem

Before you can deploy the NSX Application Platform on the TKG Cluster on Supervisor or upstream Kubernetes cluster, you must upload the `kubeconfig` file provided by your infrastructure administrator. However, your attempt to upload the `kubeconfig` file failed.

Cause

The `kubeconfig` file is not in the correct YAML format.

Solution

Validate and correct the YAML format used in the `kubeconfig` file. Try to re-upload the modified file.

Kubernetes Cluster is Inaccessible

The NSX Manager is unable to connect to the TKG Cluster on Supervisor or upstream Kubernetes cluster using the kubeconfig file that you uploaded.

Problem

While preparing to deploy the NSX Application Platform, the NSX Manager had problems connecting to the TKG Cluster on Supervisor or upstream Kubernetes cluster using the provided `kubeconfig` file.

Cause

The `kubeconfig` file is either expired or the `kubeconfig` file that you uploaded is not for the correct TKG Cluster on Supervisor or upstream Kubernetes cluster.

Solution

1. If the `kubeconfig` file has expired, re-upload a valid and up-to-date `kubeconfig` file.
2. Verify that you have uploaded the correct `kubeconfig` file for the TKG Cluster on Supervisor or upstream Kubernetes cluster and that it is accessible from the NSX Manager appliance.

Service Name/FQDN Information is Invalid

The service name/FQDN information you provided is invalid.

Problem

During the NSX Application Platform deployment preparation, you provided a service name or FQDN information that is invalid. The system needs this information to access the TKG Cluster on Supervisor or upstream Kubernetes cluster to deploy the NSX Application Platform components on it.

Cause

The service name or FQDN is in an invalid format.

Solution

Correct the format of the service name or FQDN, and ensure that it is in a local service domain. Use a format similar to the following.

```
<prefix>-<namespace>.<service-name>.<local-service-domain>.com
```

The correct format depends on the TKG Cluster on Supervisor or upstream Kubernetes environment you are using.

Deployment Precheck Failed

You must execute certain prechecks before you can deploy the NSX Application Platform. Use the information in this section to try to resolve the precheck problems you might encounter.

Kubernetes Tools Synchronization Precheck Failed

The synchronization precheck of the Kubernetes Tools failed.

Problem

The system validates the versions of the Kubernetes Tools used on the NSX Manager and on the TKG Cluster on Supervisor or upstream Kubernetes cluster you specified in an earlier step of the deployment preparation.

Cause

The cause of the precheck failure can be either of the following reasons.

- The difference between the minor version number of the Kubernetes Tools on the NSX Manager appliance and on the TKG Cluster on Supervisor or upstream Kubernetes cluster is more than one. For example, if the version installed on the NSX Manager appliance is version 1.18 and the TKG Cluster on Supervisor or upstream Kubernetes cluster uses version 1.19, the precheck passes. But if one of them uses version 1.18 and the other uses 1.20, the precheck fails.
- If the NSX Manager appliances do not have the same Kubernetes Tools version installed, the precheck fails.

Solution

1. Update the version of the Kubernetes Tools installed on your NSX Manager appliance by downloading the Kubernetes Tools bundle version provided in the VMware NSX Product Downloads webpage for your NSX Manager version.

   The Kubernetes Tools binaries are provided for NSX 3.2 release and later.

2. Upload the Kubernetes Tools binaries to your NSX Manager appliance.

   Uploading the updated version to one NSX Manager updates the rest of the NSX Manager VMs.

Kubernetes Cluster Connection Precheck Failed

The attempt to precheck the TKG Cluster on Supervisor or upstream Kubernetes cluster connection failed.

Problem

The NSX Manager tried to access the TKG Cluster on Supervisor or upstream Kubernetes cluster that you specified, but the cluster was inaccessible.

Cause

The cause of the connection precheck failure can be any of the following reasons.

- The `kubeconfig` file is not in the correct YAML format.
- The `kubeconfig` file is expired.
- The `kubeconfig` file that you uploaded is not for the correct TKG Cluster on Supervisor or upstream Kubernetes cluster.

Solution

To try to resolve the connectivity precheck failure, use the following information.

- Validate and correct the YAML format used in the `kubeconfig` file.
- If the `kubeconfig` file has expired, re-upload a valid and up-to-date `kubeconfig` file.
- Verify that you have uploaded the correct `kubeconfig` file for the TKG Cluster on Supervisor or upstream Kubernetes cluster that is accessible from the NSX Manager.

Service Name/FQDN Validation Precheck Failed

An attempt to validate the service name/FQDN failed.

Problem

The system tried to validate the service name/FQDN you provided and the precheck failed.

Cause

The service name or FQDN you provided is using an invalid format.

Solution

Correct the format of the service name or FQDN, and ensure that it is in a local service domain. Use a format similar to the following.

```
<prefix>-<namespace>.<service-name>.<local-service-domain>.com
```

The correct format depends on the TKG Cluster on Supervisor or upstream Kubernetes environment you are using.

Version Compatibility Precheck Failed

The precheck failed while validating the version compatibility of the NSX Application Platform and NSX Manager.

Problem

The precheck failed while validating the compatibility of the NSX Application Platform version with the version of the NSX Manager you are using for the platform deployment.

Cause

You selected an NSX Application Platform version that is not compatible with the version of your current NSX Manager installation.

Solution

Select another build version that is compatible with your current NSX Manager installation.

Precheck of the Available Kubernetes Resources Failed

The system found a problem while checking the available TKG Cluster on Supervisor or upstream Kubernetes resources.

Problem

The precheck failed while validating that the minimum required resources are available on the TKG Cluster on Supervisor or upstream Kubernetes cluster for the selected form factor. The system verifies that the TKG Cluster on Supervisor or upstream Kubernetes cluster has the minimum number of control node and worker nodes. It also verifies that the minimum memory, CPU, and ephemeral-storage are available on each node.

Cause

The cause of the resource availability precheck failure can be any of the following reasons.

- The TKG Cluster on Supervisor or upstream Kubernetes cluster does not have the minimum required control node and worker nodes.

- The control node and worker nodes do not have the minimum required CPU or memory for the form factor you selected.
- The control node and worker nodes do not have the minimum required ephemeral-storage for the form factor you selected.

Solution

Work with your infrastructure administrator to configure the TKG Cluster on Supervisor or upstream Kubernetes cluster with all of the required resources, as specified in the NSX Application Platform System Requirements.

Time Synchronization Precheck Warning

The system displays a warning indicator after the time synchronization precheck failed.

Problem

The system displays a warning symbol after performing the time synchronization precheck between the NSX Manager appliance and the TKG Cluster on Supervisor or upstream Kubernetes cluster you are using for the NSX Application Platform deployment.

Cause

The system time on the NSX Manager appliance is not synchronized with the system time on the TKG Cluster on Supervisor or upstream Kubernetes cluster.

Solution

Keep the system times synchronized between the NSX Manager appliance and the TKG Cluster on Supervisor or upstream Kubernetes cluster host.

The warning does not block the NSX Application Platform deployment from continuing. However, it is possible for the NSX Application Platform deployment to fail if the system times between the NSX Manager appliance and the TKG Cluster on Supervisor or upstream Kubernetes cluster host are not synchronized.

Precheck of Existing Namespaces Failed

The precheck failed while verifying the existence of the namespaces on the TKG Cluster on Supervisor or upstream Kubernetes cluster.

Problem

The namespaces that the system creates during the NSX Application Platform deployment exist on the TKG Cluster on Supervisor or upstream Kubernetes cluster and can cause a deployment failure.

Cause

The TKG Cluster on Supervisor or upstream Kubernetes cluster has existing namespaces that the system created during a previous NSX Application Platform deployment. It is possibile that the namespaces might have beeen left due to a `FORCE_DELETE` action to undeploy the NSX Application Platform from a previous deployment attempt.

Solution

Delete the existing namespaces from your TKG Cluster on Supervisor or upstream Kubernetes cluster using the information in Clean up the Kubernetes Cluster After an Incomplete Delete Action. The system must create the namespaces during the current NSX Application Platform deployment only.

Upload of the Kubernetes Tools Failed

An attempt to upload a version of the Kubernetes Tools failed.

Problem

The system was unable to complete the upload of the Kubernetes Tools bundle.

Cause

There can be multiple reasons for the upload failure. The following are some possible causes.

- The Kubernetes Tools bundle filename is not in the correct format.
- If you uploaded the Kubernetes Tools bundle file using a remote URL, the NSX Manager is unable to access that remote server.

Solution

- Ensure that you have download the Kubernetes Tools bundle from the VMware Product Download web site and that after the download, there is no tampering performed on the bundle. Do not rename the downloaded Kubernetes Tools bundle file.
- If you uploaded the Kubernetes Tools bundle file to a remote server location using the **Upload Remote File** option to upload bundle, ensure the NSX Manager is able to connect to that remote server.

NSX Application Platform Deployment Failed When Using Harbor Registry with HTTP

An attempt to deploy the NSX Application Platform failed when the system tried to install the Helm chart.

Problem

After completing 10% of the NSX Application Platform deployment process, the system displayed the following error message when it tried to install the Helm chart.

```
NSX Application Platform deployment failed!
See Troubleshooting Documentation.
Helm install chart operation failed. Error: failed post-install: timed out waiting for the
condition.
```

Cause

The deployment process is trying to access a Harbor registry that is configured to use HTTP instead of HTTPS.

Solution

1 For a production environment, your infrastructure administrator must obtain a CA-signed certificate to configure HTTPS access to the private Harbor registry that they installed. See the Configure HTTPS Access to Harbor webpage for more information.

2 (Optional) (Use this step with caution for upstream Kubernetes cluster.) If you want to continue to use a Harbor registry that uses HTTP instead of HTTPS, your infrastructure administrator must use the following information as a workaround. They must apply the workaround on all the control and worker nodes of the upstream Kubernetes cluster that you are using for the NSX Application Platform deployment.

   a Add the following to your `daemon.json` file located at: `/etc/docker/daemon.json` directory.

   ```
   {
   "insecure-registries" : ["Harbor FQDN or Harbor IP"]
   }
   ```

   b Restart the Docker Engine using the following command.

   ```
   systemctl restart docker
   ```

3 (Optional) (Use this step with caution for TKG Cluster on Supervisor.) To continue to use a Harbor registry that uses HTTP instead of HTTPS, your infrastructure administrator must apply the following workaround information on all the control and worker nodes of the TKG Cluster on Supervisor that you are using for the NSX Application Platform deployment.

   a Log in to each Tanzu worker node using the steps described in SSH to Tanzu Kubernetes Cluster Nodes as System User Using a Password.

   b Edit the `config.toml` file using the following command.

   ```
   sudo vim /etc/containerd/config.toml
   ```

c Add the following entries in the `config.toml` for the Harbor registry, where "10.222.44.111" is the URL for an example Harbor registry.

```
 [plugins.cri.registry]
      [plugins.cri.registry.mirrors]
        [plugins.cri.registry.mirrors."docker.io"]
          endpoint = ["https://registry-1.docker.io"]
        [plugins.cri.registry.mirrors."localhost:5000"]
           endpoint = ["http://localhost:5000"]
        [plugins.cri.registry.mirrors."10.222.44.111"]
           endpoint = ["http://10.222.44.111"]
```

d Restart the Docker client using the following command.

```
sudo systemctl restart docker
```

NSX Application Platform Failed During the Helm Cert-Manager Installation

The NSX Application Platform deployment failed while attempting to deploy the Helm `cert-manager` tool.

Problem

The system is unable to continue with the NSX Application Platform deployment due to an error encountered while deploying the `cert-manager` tool. The system displayed the error message `Helm install chart operation failed`.

Cause

A possible cause is the process that attempted to download the NSX Application Platform Helm charts and Docker images stalled, possibly due to slow network performance, and timed out.

Solution

1 Determine if the deployment failure is due to the process timing out while attempting to download the NSX Application Platform Helm charts and Docker images from the public VMware-hosted registry.

a Log in to the NSX Manager as `root`.

b Perform the following debugging steps at the system prompt, where *cert-manager-webhook-pod-ID* is the value returned from the `napp-k get jobs -n cert-manager` command.

```
NSXManager-prompt% napp-k get ns
NSXManager-prompt% napp-k get jobs -n cert-manager
NSXManager-prompt% napp-k describe pod cert-manager-webhook-pod-ID -n cert-manager
```

   When the `napp-k describe pod` command returns the output `Pulling image from ...`, it indicates that the image download process is taking a longer time than normal. The system will retry the download.

   After waiting for several minutes and the `napp-k describe pod` returns the `Running` status, it indicates that the pod is up.

2. After the pod is ready, retry deploying the NSX Application Platform again from the NSX Manager UI.

Kubernetes Pods Are Stuck in the Terminating Status

After the system experiences a very high rate of events in a short period of time, some of the pods in the TKG Cluster on Supervisor or upstream Kubernetes cluster are stuck in the `Terminating` status.

Problem

After the system recovered from a very high rate of events occurring, the **System > NSX Application Platform** UI displays that the NSX Application Platform is in a `Degraded` status. In addition, some of the pods in the TKG Cluster on Supervisor or upstream Kubernetes cluster are stuck in the `Terminating` status for a few minutes or longer.

Cause

Due to some Kubernetes infrastructure issues, some of the pods cannot be deleted correctly because of one of the following reasons.

- A finalizer associated with the stuck pod is not able to complete.
- The stuck pod is not responding to the termination signals.

Solution

Ask your infrastructure administrator to use the following information to manually delete the pods that are stuck in the `Terminating` status.

1. Log in to the control node for your TKG Cluster on Supervisor or upstream Kubernetes cluster.
2. Use the following command to find all of the pods that are in the `Terminating` status.

   ```
   get pod -A | grep Terminating
   ```

3. Force delete the pods with the `Terminating` status, using the following command.

   ```
   kubectl delete pod <pod-name> -n <namespace> --force --grace-period=0
   ```

4 Repeat the following command and verify that the stuck pods have been deleted successfully. If necessary, repeat step 3 again for the pods that continue to be in the `Terminating` status.

```
get pod -A | grep Terminating
```

Clean up the Kubernetes Cluster After an Incomplete Delete Action

The system did not complete deleting the NSX Application Platform components from the TKG Cluster on Supervisor or upstream Kubernetes cluster because the `kubeconfig` file expired before the `delete` operation completed.

Problem

While trying to delete the NSX Application Platform deployment from the TKG Cluster on Supervisor or upstream Kubernetes cluster, the operation failed to finish because the `kubeconfig` file that you are using to connect to the TKG Cluster on Supervisor or upstream Kubernetes cluster expired before the `delete` operation finished. You performed the `force delete` operation when prompted. However, that operation only removed the NSX Application Platform entry from the NSX Manager unified appliance. The previously deployed NSX Application Platform components still exist on the TKG Cluster on Supervisor or upstream Kubernetes cluster. You must delete those components before you can use the same TKG Cluster on Supervisor or upstream Kubernetes cluster for any future NSX Application Platform deployment.

Cause

The TKG Cluster on Supervisor or upstream Kubernetes cluster became inaccessible because the `kubeconfig` file that you used to access the cluster expired or became invalid.

Solution

1 Obtain an updated `kubeconfig` file from your infrastructure administrator so you can access the same TKG Cluster on Supervisor or upstream Kubernetes cluster.

Alternatively, ask your infrastructure administrator to help you remove the NSX Application Platform components that are still remaining on the TKG Cluster on Supervisor or upstream Kubernetes cluster.

2 Log into the NSX Manager appliance node as a **`root`** user.

Note that the `napp-h` command used in later solution steps is an alias in the NSX Manager session for the following command.

```
helm --kubeconfig <path-to-updated-kubeconfig-file> --namespace nsxi-platform
```

3 Copy the updated `kubeconfig` file to the file's current location in the NSX Manager appliance at `/config/vmware/napps/.kube`.

4 If you activated the NSX Network Detection and Response feature on the NSX Application Platform, delete the feature first using one of the following methods.

- From the NSX Manager command prompt, use the following API calls.

  ```
  nsx-manager-prompt> napp-h uninstall cloud-connector
  nsx-manager-prompt> napp-h uninstall nsx-ndr
  ```

- From any other system that has access to the TKG Cluster on Supervisor or upstream Kubernetes cluster and installed Helm chart, use the following API calls.

  ```
  K8s-prompt> helm --kubeconfig <path-to-updated-kubeconfig-file> --namespace nsxi-
  platform uninstall cloud-connector
  K8s-prompt> helm --kubeconfig <path-to-updated-kubeconfig-file> --namespace nsxi-
  platform uninstall nsx-ndr
  ```

If any or all of the previous commands provided in this step fail or time out, retry the delete process using the following API calls.

From the NSX Manager command prompt, use the following API calls.

```
nsx-manager-prompt> napp-h uninstall cloud-connector --no-hooks
nsx-manager-prompt> napp-h uninstall nsx-ndr --no-hooks
```

From any other system that has access to the TKG Cluster on Supervisor or upstream Kubernetes cluster, use the following API calls.

```
K8s-prompt> helm --kubeconfig <path-to-updated-kubeconfig-file> --namespace nsxi-platform
uninstall cloud-connector --no-hooks
K8s-prompt> helm --kubeconfig <path-to-updated-kubeconfig-file> --namespace nsxi-platform
uninstall nsx-ndr --no-hooks
```

5 If you activated the the NSX Malware Prevention feature on the NSX Application Platform, delete it first using one of the following methods.

- From the NSX Manager command prompt, use the following API calls.

  ```
  nsx-manager-prompt> napp-h uninstall cloud-connector
  nsx-manager-prompt> napp-h uninstall reputation-service
  nsx-manager-prompt> napp-h uninstall malware-prevention
  ```

- From any other system that has access to the TKG Cluster on Supervisor or upstream Kubernetes cluster, use the following API calls.

  ```
  K8s-prompt> helm --kubeconfig <path-to-updated-kubeconfig-file> --namespace nsxi-
  platform uninstall cloud-connector
  ```

```
K8s-prompt> helm --kubeconfig <path-to-updated-kubeconfig-file> --namespace nsxi-
platform uninstall reputation-service
K8s-prompt> helm --kubeconfig <path-to-updated-kubeconfig-file> --namespace nsxi-
platform uninstall malware-prevention
```

If any or all of the previous commands provided in this step fail or time out, retry using the following API calls.

From the NSX Manager command prompt, use the following API calls.

```
nsx-manager-prompt> napp-h uninstall cloud-connector --no-hooks
nsx-manager-prompt> napp-h uninstall reputation-service --no-hooks
nsx-manager-prompt> napp-h uninstall malware-prevention --no-hooks
```

From any other system that has access to the TKG Cluster on Supervisor or upstream Kubernetes cluster, use the following API calls.

```
K8s-prompt> helm --kubeconfig <path-to-updated-kubeconfig-file> --namespace nsxi-platform
uninstall cloud-connector --no-hooks
K8s-prompt> helm --kubeconfig <path-to-updated-kubeconfig-file> --namespace nsxi-platform
uninstall reputation-service --no-hooks
K8s-prompt> helm --kubeconfig <path-to-updated-kubeconfig-file> --namespace nsxi-platform
uninstall malware-prevention --no-hooks
```

6 If you activated the NSX Intelligence feature on the NSX Application Platform, delete it first using one of the following methods.

- From the NSX Manager command prompt, use the following API call.

  ```
  nsx-manager-prompt> napp-h uninstall intelligence
  ```

- From any other system that has access to the TKG Cluster on Supervisor or upstream Kubernetes cluster and installed Helm chart, use the following API call.

  ```
  K8s-prompt> helm --kubeconfig <path-to-updated-kubeconfig-file> --namespace nsxi-
  platform uninstall intelligence
  ```

If any or all of the previous commands provided in this step fail or time out, retry the delete process using the following API calls.

From the NSX Manager command prompt, use the following API call.

```
nsx-manager-prompt> napp-h uninstall intelligence --no-hooks
```

From any other system that has access to the TKG Cluster on Supervisor or upstream Kubernetes cluster, use the following API call.

```
K8s-prompt> helm --kubeconfig <path-to-updated-kubeconfig-file> --namespace nsxi-platform
uninstall intelligence --no-hooks
```

7 Delete the NSX Application Platform components. Type the following commands individually at the NSX Manager command prompt and in the order listed.

   a If the Upgrade Coordinator was previously deployed, use the following command to uninstall it.

   ```
   nsx-manager-prompt> napp-h uninstall nsxi-upgrade
   ```

   b Remove the NSX Metrics feature from the TKG Cluster on Supervisor or upstream Kubernetes cluster by typing the following command.

   ```
   nsx-manager-prompt> napp-h uninstall metrics
   ```

   c Uninstall the NSX Application Platform from the TKG Cluster on Supervisor or upstream Kubernetes cluster, and delete its namespace by typing the following commands separately.

   ```
   nsx-manager-prompt> napp-h uninstall nsxi-platform
   nsx-manager-prompt> kubectl --kubeconfig <path-to-updated-kubeconfig-file> delete
   namespace nsxi-platform
   ```

   d Uninstall the project contour component and delete its namespace using the following commands.

   ```
   nsx-manager-prompt> helm --kubeconfig <path-to-updated-kubeconfig-file> --namespace
   projectcontour uninstall projectcontour
   nsx-manager-prompt> kubectl --kubeconfig <path-to-updated-kubeconfig-file> delete
   namespace projectcontour
   ```

   e Uninstall the certificate manager and delete its namespace using the following commands.

   ```
   nsx-manager-prompt> helm --kubeconfig <path-to-updated-kubeconfig-file> --namespace
   cert-manager uninstall cert-manager
   nsx-manager-prompt> kubectl --kubeconfig <path-to-updated-kubeconfig-file> delete
   namespace cert-manager
   ```

Metrics Postgres Database Deployment Failed

An error occurred during the NSX Metrics feature deployment while the system tried to deploy the Postgres database.

Problem

The system encountered a problem while trying to deploy the Postgres database used by the NSX Metrics feature.

Cause

There are multiple reasons for the Metrics Postgres database deployment to fail.

Solution

1. Ensure that the system deployed the NSX Application Platform successfully.
2. If you have access to the NSX Manager command line interface (CLI), use the following steps to investigate any errors recorded in the NSX Metrics logs. If you do not have access to the NSX Manager CLI, use the information provided in step 3.
   a. Log into the NSX Manager appliance as a root user.
   b. Mark the TKG Cluster on Supervisor or upstream Kubernetes configuration for any subsequent `helm` and `kubectl` command invocation.

      ```
      export KUBECONFIG=/config/vmware/napps/.kube/config
      ```

   c. Use the following command to verify that the system has deployed the Metrics helm chart successfully.

      ```
      helm --namespace nsxi-platform list --all --filter 'metrics'
      ```

      The STATUS property must display `deployed`.
   d. Use the following command to verify that the setup pods were deployed and completed successfully.

      ```
      kubectl --namespace nsxi-platform get pods --selector='app.kubernetes.io/
      instance=metrics,app.kubernetes.io/name=postgresql-ha'
      ```

   e. Using the following command, inspect the logs for the setup pods that did not deploy or are not in the `Ready` state.

      ```
      kubectl --namespace nsxi-platform logs <POD_NAME>
      ```

3. If you do not have access to the NSX Manager CLI, collect the NSX Application Platform support bundle using information in Collect the Support Bundles for the NSX Application Platform Using the UI. Inspect the support bundle logs for the TKG Cluster on Supervisor or upstream Kubernetes pods whose name start with `metrics-postgresql-ha`.
4. Resolve the errors reported in the logs.
5. After resolving the errors reported in the logs, retry deploying the NSX Application Platform using the NSX Manager UI.

NSX Metrics Feature Deployment Failed

The NSX Metrics feature deployment failed due to errors.

Problem

The system encountered some errors while attempting to deploy the NSX Metrics feature.

Cause

The NSX Metrics feature deployment requires several Kubernetes components to be available on top of the NSX Application Platform deployment. If the TKG Cluster on Supervisor or upstream Kubernetes cluster is in an unstable state, the NSX Metrics deployment can fail.

Solution

1 Ensure that the system deployed the NSX Application Platform successfully.

2 If you have access to the NSX Manager command line interface (CLI), use the following steps to investigate any errors recorded in the NSX Metrics logs. If you do not have access to the NSX Manager CLI, use the information provided in step 3.

   a Log into the NSX Manager appliance as a root user.

   b Mark the Kubernetes configuration for any subsequent `helm` and `kubectl` command invocation.

   ```
   export KUBECONFIG=/config/vmware/napps/.kube/config
   ```

   c Use the following command to verify that the system has deployed the Metrics helm chart successfully.

   ```
   helm --namespace nsxi-platform list --all --filter 'metrics'
   ```

   The STATUS property must display `deployed`.

   d Use the following command to verify that the setup pods were deployed and completed successfully.

   ```
   kubectl --namespace nsxi-platform get pods --selector='app.kubernetes.io/
   instance=metrics'
   ```

   e Using the following command, inspect the logs for the setup pods that did not deploy or are not in the `Ready` state.

   ```
   kubectl --namespace nsxi-platform logs <POD_NAME>
   ```

3 If you do not have access to the NSX Manager CLI, Collect the NSX Application Platform support bundle using information in Collect the Support Bundles for the NSX Application Platform Using the UI. Inspect the support bundle logs for the TKG Cluster on Supervisor or upstream Kubernetes pods whose names start with `metrics-`.

4 Resolve the errors reported in the logs.

5 After resolving the errors reported in the logs, retry deploying the NSX Application Platform using the NSX Manager UI.

Metrics Feature Post-Deployment Failure

The user interface displays the `Meterics-PostDeployment failed` error message.

Problem

The system encountered an error while invoking post-deployment API calls after deploying the NSX Metrics feature.

Cause

After deploying the NSX Metrics feature, the system sends a post-deployment API request that invokes an API call to the NSX and ClusterApi pod on the NSX Application Platform. The system encountered an error

Solution

1 Ensure that the system deployed the NSX Application Platform successfully.

2 If you have access to the NSX Manager command line interface (CLI), use the following steps to investigate any errors recorded in the NSX Metrics logs. If you do not have access to the NSX Manager CLI, use the information provided in step 3.

  a Log into the NSX Manager appliance as a root user.

  b Mark the Kubernetes configuration for any subsequent `helm` and `kubectl` command invocation.

  ```
  export KUBECONFIG=/config/vmware/napps/.kube/config
  ```

  c Use the following command to verify that the system has deployed the Metrics helm chart successfully.

  ```
  helm --namespace nsxi-platform list --all --filter 'metrics'
  ```

  The STATUS property must display `deployed`.

  d Use the following command to verify that the setup pods were deployed and completed successfully.

  ```
  kubectl --namespace nsxi-platform get pods --selector='app.kubernetes.io/
  instance=metrics'
  ```

  e Using the following command, inspect the logs of the `metrics-app-server` pod. The `POD_NAME` starts with `metrics-app-server-*` for any exceptions or errors.

  ```
  kubectl --namespace nsxi-platform logs <POD_NAME>
  ```

3 If you do not have access to the NSX Manager CLI, collect the NSX Application Platform support bundle using information in Collect the Support Bundles for the NSX Application Platform Using the UI. Inspect the support bundle logs for the TKG Cluster on Supervisor or upstream Kubernetes pods whose names start with `metrics-app-server`.

4 Post-deployment calls `napp/api/v1/metrics/data-collection` API `GET` and `POST` allow you to read and set the source for the data collection. You can toggle the data collection on or off. If there was problem encountered and the post deployment API did finish successfully, you can manually see if the data collection attribute for some of the hosts are not turned on, and then turn them on as desired.

5 Additionally, the `/infra/sites/intelligence/registration GET` and `/infra/sites/intelligence/registration/{cluster-id} POST` APIs are called as part of the post-deployment step. The output of `api/v1/infra/sites/intelligence/registration` has an attribute `is_metrics_enabled`, which should be set during deployment and is an indication that post-deployment completed successfully. If it is not set, then you can manually set the NSX Metrics feature on using the following information.

   a Invoke the `GET napp/api/v1/metrics/data-collection` command, which returns a JSON content, similar to the following.

```
{
    "metrics_toggle_nsx": true,
    "metrics_toggle_nsx_cloud_native": true,
    "metrics_toggle_nsx_config": false
}
```

   b Locate all of the flags that are set to **`false`** and toggle them to **`true`**. More specifically, ensure that the `metrics_toggle_nsx_config` and `metrics_toggle_nsx` attributes are set to **`true`**.

```
{
    "metrics_toggle_nsx": true,
    "metrics_toggle_nsx_cloud_native": true,
    "metrics_toggle_nsx_config": true
```

   c Invoke the `POST napp/api/v1/metrics/data-collection` command using the modified body with flags set to **`true`**.

Failed to Scale Out the NSX Application Platform Services

An attempt to scale out the NSX Application Platform core services failed.

Problem

When you tried to scale out the NSX Application Platform core services, the operation failed.

Cause

There are multiple reasons as to why the scale-out operation failed.

Solution

- Ensure that the prerequisites listed in Scale Out the NSX Application Platform are met.
- Ensure that all the core services are selected in the **Scale Out** dialog box and retry the scale-out operation again.

Failed to Increase the Volume Size of the Data Storage Disk

An attempt to increase the volume size of the data storage disk failed.

Problem

After invoking the **Scale Up** operation on the Data Storage service, the system displays an error message similar to the following.

```
failed to scale up DATA_STORAGE in stage resize persistent volume: failed to wait pvc
resizing: timeout waiting for minio pvc resizing
```

.

Cause

- Your TKG Cluster on Supervisor or upstream Kubernetes cluster does not support volume resizing.
- Your TKG Cluster on Supervisor or upstream Kubernetes cluster might not have enough free disk space to allow for the expansion of the existing volumes.

Solution

- Work with your infrastructure administrator to enable volume expansion on the TKG Cluster on Supervisor or upstream Kubernetes cluster that the NSX Application Platform uses.

  The **Scale Up** operation is currently not supported if your cluster does not support the expansion of volumes that are attached to the running pod. Specifically, the **Scale Up** operation is not supported for TKG Cluster on Supervisor versions 1.17.x and 1.18.x.
- Ask your infrastructure administrator to properly expand the Persistent Volume size of the MinIO StatefulSet container that is currently in use.

User-Imported CA-Signed Certificate is Expired

The CA-signed certificate that you imported has expired.

Problem

The CA-signed certificate that you imported to the NSX Manager appliance host has expired and you are unable to continue working with the NSX Application Platform. You must delegate a self-signed CA-signed certificate to continue working with the NSX Application Platform and the NSX-T Data Center features that it hosts.

Cause

The default expiry for a CA-signed certificate is 825 days. The system self-signed certificates are automatically renewed. If you import a custom CA-signed certificate, you need to maintain the lifecycle of that certificate. If you forget to renew it, then the connection between the NSX Application Platform andNSX Manager unified appliance will break and you will not be able to continue working with the NSX Application Platform and the NSX features that it hosts.

Solution

1 Log into the NSX Manager appliance as the root user.

2 Delegate to a self-signed CA-signed certificate using the following `kubectl` command at the NSX Manager command prompt.

```
system prompt> kubectl patch certificate ca-cert -n cert-manager --type='json' -p='[{"op":
"replace", "path": "/spec/secretName", "value":"ca-key-pair"}]'
```

3 Wait for about 30 seconds, then export the egress certificate by entering the following `kubectl` and `cat` commands at the NSX Manager system prompt, one at a time.

```
system prompt> kubectl get secret -n nsxi-platform egress-tls-cert
-o=jsonpath='{.data.tls\.crt}' | base64 -d - > tls.crt
system prompt> kubectl get secret -n nsxi-platform egress-tls-cert
-o=jsonpath='{.data.ca\.crt}' | base64 -d - > ca.crt
system prompt> cat tls.crt ca.crt > egress.crt
```

4 Import the `egress.crt` into the NSX Manager apppliance using the user interface and otain the certificate UUID.

   a From your browser, log in with Enterprise Admin privileges to an NSX Manager at https://<nsx-manager-ip-address>.

   b Navigate to **System > Certificates**, click **Import**, and select **Certificate** from the drop-down menu.

   c Enter a name for the certificate.

   d Set **Service Certificate** to **No**.

   e In the **Certificate Contents** textbox, paste the contents of the the `egress.crt` file that you created in the previous step.

   f Click **Save**.

   g In the Certificates table, expand the row for the newly added certificate and copy the certificate ID value.

5 Back in the NSX Manager appliance root user session, obtain the PricinpleIdentity `cloudnative_platform_egress` UUID using the following `curl` command at the system prompt.

```
system prompt> curl -ku 'admin:yourAdminPassword' https://127.0.0.1/api/v1/trust-
management/principal-identities
```

6 Bind the PrincipleIdentity with the imported certificate ID, using the following `curl` command at the system prompt.

```
curl -ku 'admin:yourAdminPassword' https://127.0.0.1/api/v1/trust-management/principal-
identities?action=update_certificate -X POST -H "Content-Type: application/json" -H
"X-Allow-Overwrite: true" -d '{"principal_identity_id": "<PI-UUID>", "certificate_id":
"<EGRESS-CERT-ID>"}'
```