

KubeSphere and VMware Cloud Native Solution

Table of Contents

Executive Summary	3
Business Case	3
Business Values.....	3
Key Results	3
Audience	4
Technology Overview	4
VMware vSphere	4
VMware vSAN	4
VMware CNS.....	4
Antrea	4
KubeSphere	4
Solution Configuration	5
Architecture Diagram	5
Hardware Resources.....	5
Software Resources	6
Network Configuration.....	6
Solution Validation	7
Test Overview	7
Integrate KubeSphere with vSphere CSI	7
Integrate KubeSphere with Antrea	19
Best Practices	34
Conclusion	35
Reference	35
About the Author	35

Executive Summary

Business Case

Running container workloads on VMware vSphere® requires a solution that is cost-effective, highly scalable, and easy to manage. Although Kubernetes delivers a decent solution for container orchestration, its steep learning curve and CLI-based working mode put a huge burden on operation and maintenance (O&M) personnel. On top of Kubernetes, KubeSphere is a container platform that can work together with VMware Cloud Native Storage (CNS) and VMware Antrea to build a one-stop modern application solution.

KubeSphere is an application-oriented multi-tenant container platform, which provides full-stack automated O&M capabilities and simplifies the DevOps workflows for enterprises. KubeSphere can run on Linux VMs powered by VMware vSphere or VMware Tanzu Kubernetes Grid and visualize the management of Kubernetes resources. This solution enables features such as multi-cluster management, logging, monitoring, alerting, notifications, audits, events, and metering and billing. Moreover, KubeSphere provides platform-level capabilities, such as App Store, DevOps, edge computing, and microservices governance.

KubeSphere can work together with VMware CNS and Antrea to achieve infrastructure automation and platform extensibility, which simplifies and automates Day 0 to Day 2 operations with Kubernetes. This solution provides a complete, automated, and one-stop platform for developers and O&M personnel to modernize and manage applications during the whole lifecycle. KubeSphere also helps IT personnel to reduce the learning curve and O&M costs, improves operation efficiency, and fuels digital transformation.

In this solution, we provide infrastructure administrators, developers, and maintenance personnel with the design, deployment, and best practices of KubeSphere on VMware vSphere with VMware vSAN™, to help enterprises run and manage modern applications.

Business Values

Deploying KubeSphere on VMware vSphere with vSAN features the following benefits:

- Quick deployment. You can deploy KubeSphere on native Kubernetes powered by VMware vSphere and vSAN at ease. vSphere Container Storage Interface (CSI) and Antrea are preinstalled in the KubeSphere App Store, so you can configure storage and networking in a few clicks.
- Independent O&M for IaaS and PaaS. For infrastructure administrators, you only need to operate and maintain Fibre Channel Storage Area Network (FC SAN) or Network Attached Storage (NAS) on vSphere or vSAN. For platform administrators, KubeSphere allows you to focus on how to manage container workloads and DevOps workflows.
- Cluster lifecycle management. KubeSphere simplifies the lifecycle management of Kubernetes clusters on vSphere with vSAN, and provides features such as cluster upgrade, certificate renewal, automatic backup, and parameter management.
- Storage and network security. You can use CNS for VMware vSphere to integrate storage of different types, including DAS, SAN, and NAS, and use a CSI for persistent storage on KubeSphere. Antrea powers KubeSphere to enable multi-layer network policies and traffic observability.
- Out-of-the-box PaaS. KubeSphere on VMware vSphere with vSAN is native to Kubernetes and provides out-of-the-box features, such as DevOps powered by Jenkins or Argo CD, microservices governance powered by Istio or Spring Cloud, containerized database management, and edge computing.
- Observability. KubeSphere provides various features for observability, such as unified logging, monitoring, and alerting, diverse notification channels, events, audits, and billing and metering.
- Multi-cluster or multi-cloud delivery. KubeSphere tailors delivery to meet business needs in active-active or geo-redundancy scenarios. To manage multiple clusters across clouds, KubeSphere supports cluster federation and enables observability for unified O&M.

Key Results

Powered by VMware vSphere and vSAN, KubeSphere achieves the following key results:

- Simplify resource management from the following aspects: Kubernetes-native O&M on VMware vSphere with vSAN, collaboration in open source ecosystems, integration with DevOps tools, and service governance.

- Increase sales for VMware and QingCloud and work together with VMware vSphere, vSAN, and NSX to provide a one-stop cloud native solution for existing enterprise users.

Audience

This solution is suitable for IT administrators, developers, O&M personnel, and engineers and experts engaged in planning, design, and deployment of VMware vSphere, vSAN, Kubernetes, and DevOps workloads. Before you read the following content, make sure you have some knowledge of VMware vSphere, vSAN, Kubernetes, and Docker.

Technology Overview

The following components work collaboratively to implement this solution:

- vSphere
- vSAN
- VMware CNS
- Antrea
- KubeSphere

VMware vSphere

VMware vSphere is a virtualization platform powered by VMware, which can transform a data center into a converged computing infrastructure that provides CPU, storage, and networking resources. VMware vSphere manages the infrastructure as a unified environment and provides you with the tools to manage the data center.

VMware vSAN

VMware vSAN uses a software-defined approach to creating shared storage for virtual machines (VMs). This service can virtualize local physical storage resources provided by ESXi servers and turn them into storage pools in which resources can be provisioned to VMs and applications on demand. vSAN is implemented in the ESXi hypervisor, and vSAN can be configured to work as a hybrid cluster or an all-flash cluster.

VMware CNS

VMware CNS for vSphere integrates with Kubernetes and allows Kubernetes to configure storage on demand on vSphere in a fully automated and scalable manner. This service also enables volume observability for administrators through vCenter.

VMware CNS allows administrators to run, monitor, and manage container and VM storage on the same platform. This simplifies container infrastructure storage, lifecycle management, and operations for administrators. VM and container storage provisioning on a unified platform reduces administrator learning costs and enables consistent operations across workloads and clouds. CNS can help you reduce time in managing infrastructure and focus more on building applications that deliver business value.

Antrea

Antrea is a Kubernetes-native network solution that uses Open vSwitch as the data plane at Layer 3 or 4 to provide network and security services for Kubernetes. Antrea can work as a Container Network Interface (CNI) plug-in to enable various network and security features, such as the configuration of static IP addresses, IP Address Management (IPAM), multi-cluster networking, overlay and underlay network construction, traffic encryption, layered and multi-dimensional network policies, network flow analysis and aggregation, egress gateways, LoadBalancer, and Traceflow.

KubeSphere

Built on top of Kubernetes, KubeSphere is a distributed operating system for cloud native applications. KubeSphere provides multiple features to build an enterprise-grade Kubernetes environment, such as multi-cloud and multi-cluster management, resource management, DevOps, application lifecycle management, microservices governance or service meshes, log query and collection, services and networks, multi-tenant management, monitoring and alerting, event and audit query, storage management, access control, GPU support, network policies, image registry management, and security management. KubeSphere abstracts away complex technical details of the underlying infrastructure, and helps enterprises seamlessly deploy, update, migrate, and manage existing containerized applications on various types of infrastructure. This way, KubeSphere allows developers to focus on application

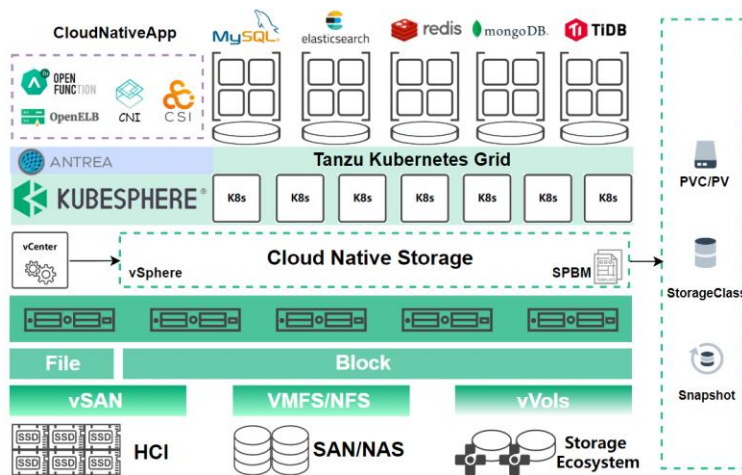
development, and O&M personnel can leverage enterprise-grade features to speed up DevOps workflows and delivery processes. The features include observability, troubleshooting, unified monitoring and log query, storage and network management, and easy-to-use CI/CD pipelines.

Solution Configuration

This section describes how to implement the solution from the following aspects:

- Architecture
- Hardware resources
- Software resources
- Network configuration

Architecture Diagram



Hardware Resources

Table 1. Hardware resources for KubeSphere on Linux

PROPERTY	SPECIFICATION
OS	Ubuntu 16.04/18.04/20.04 Debian Buster, Stretch CentOS 7.x Red Hat Enterprise Linux 7 SUSE Linux Enterprise Server 15 /openSUSE Leap 15.2
CPU	> 2 vCPUs
RAM	> 4 GB
CRI	Docker, Containerd, CRI-O

Dependency	Socat, Conntrack
Disks	> 40 GB

Software Resources

Table 2. Software resources for KubeSphere on Kubernetes

SOFTWARE	VERSION
Kubernetes	v1.19.x, v1.20.x, v1.21.x, v1.22.x, v1.23.x, and v1.24.x
StorageClass	Exist

Network Configuration

Table 3. Ports used by KubeSphere

SERVICE	PROTOCOL	ACTION	SOURCEPORT	DESTPORT
SSH	TCP	ALLOW	22	
ETCD	TCP	ALLOW	2379	2380
APISERVER	TCP	ALLOW	6443	
NODEPORT	TCP	ALLOW	30000	32767
MASTER	TCP	ALLOW	10250	10258
DNS	TCP	ALLOW	53	
DNS	UDP	ALLOW	53	
LOCAL-REGISTRY	TCP	ALLOW	5000	
LOCAL-APT	TCP	ALLOW	5080	
RPCBIND	TCP	ALLOW	111	
METRICS-SERVER	TCP	ALLOW	8443	

Solution Validation

Test Overview

This section validates the feasibility of integrating KubeSphere with vSphere CSI and Antrea.

Integrate KubeSphere with vSphere CSI

This section guides you through how to install native Kubernetes and Kubesphere on VMs powered by VMware vSphere. In this example, VMware CNS is used for persistent storage, and KubeSphere helps you create workloads backed by the vSphere CSI driver. For example, you can create a stateful MySQL application with a persistent volume claim (PVC), test volume attaching, create volume snapshots, and delete or recover volumes.

Create a Kubernetes Cluster

To create a Kubernetes cluster at ease, we recommend that you use the open source tool KubeKey. KubeKey is an open source tool powered by QingCloud. This tool can help you create KubeSphere clusters and install KubeSphere at ease. You can use this tool to add cluster nodes, install Harbor registries, deploy cluster networks, storage plug-ins, and etcd clusters, renew cluster certificates, and quickly upgrade clusters.

To obtain KubeKey, run the following command:

```
curl -sL https://get-kk.kubesphere.io | VERSION=v3.0.2 sh -
```

Docker allows you to run a MySQL application with one command, and KubeKey allows you to run a Kubernetes cluster and KubeSphere with one command. Run the following command to install Kubernetes and KubeSphere:

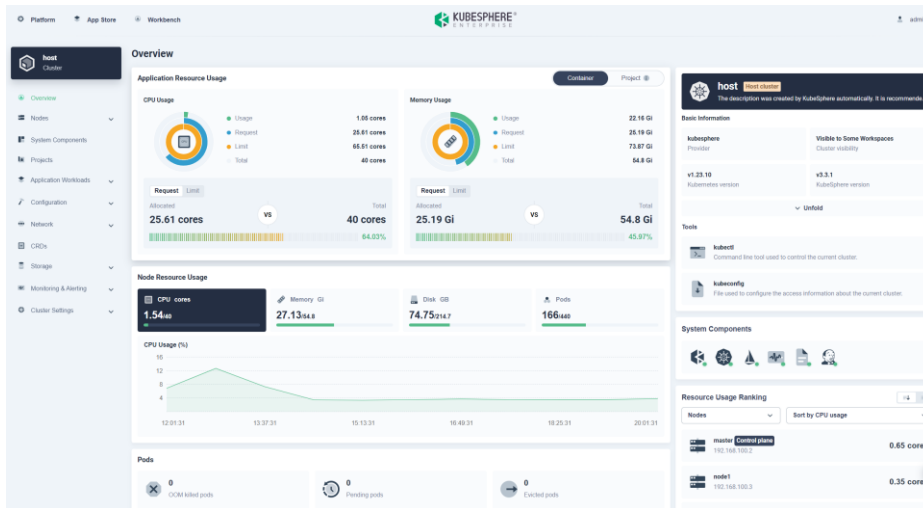
```
kk create cluster --with-kubernetes v1.23.8 --with-kubesphere v3.3.1
```

If you create a cluster with multiple nodes, run **kk create config** to create a configuration file. After you configure node information, run **kk create cluster -f config.yaml** to apply the file.

Sample code:

```
hosts:
- {name: master, address: 192.168.0.2, internalAddress: 192.168.0.2, user: ubuntu, password: Testing123}
- {name: node1, address: 192.168.0.3, internalAddress: 192.168.0.3, user: ubuntu, password: Testing123}
roleGroups:
  etcd:
  - master
  control-plane:
  - master
  worker:
  - node1
```

Wait until Kubernetes and KubeSphere are installed. This process might take about 20 minutes.



Install the storage plug-in for vSphere

vSphere CNS consists of the vCenter control plane and the storage plug-in for Kubernetes. The control plane is available for vCenter 6.7U3 and later, so you need only to install the CSI plug-in in the Kubernetes cluster backed by VMware vSphere.

Before you create a VM for Kubernetes, make sure the following conditions are met:

- ✓ VMware Tools is installed.
- ✓ `disk.EnableUUID` is set to true.
- ✓ The hardware version is 15 or later.
- ✓ A VMware Paravirtual SCSI controller is used for the VM.

Before you install the vSphere CSI driver, check its compatibility with vCenter and Kubernetes on the official website.

Install vsphere-cloud-controller-manager

Add labels to all nodes:

```
kubectl taint node k8s node.cloudprovider.kubernetes.io/uninitialized=true:NoSchedule
```

Download `vsphere-cloud-controller-manager.yaml`:

```
wget https://raw.githubusercontent.com/kubernetes/cloud-provider-vsphere/release-1.23/releases/v1.23/vsphere-cloud-controller-manager.yaml
```

Modify the vCenter-related configuration in Secret and ConfigMap:

```
apiVersion: v1
kind: Secret
metadata:
  name: vsphere-cloud-secret
  namespace: kube-system
stringData:
  10.0.0.1.username: "<ENTER_YOUR_VCENTER_USERNAME>"
  10.0.0.1.password: "<ENTER_YOUR_VCENTER_PASSWORD>"
```

```
apiVersion: v1
kind: ConfigMap
metadata:
  name: vsphere-cloud-config
  namespace: kube-system
```



```
data:
  vcenter:
    <your-vcenter-name-here>:
    server: 10.0.0.1
    user: <use-your-vcenter-user-here>
    password: <use-your-vcenter-password-here>
  datacenters:
    - Datacenter01
```

Run the following command to apply the file:

```
❏ → kubectl apply -f vsphere-con-ma.yaml
serviceaccount/cloud-controller-manager created
secret/vsphere-cloud-secret created
configmap/vsphere-cloud-config created
rolebinding.rbac.authorization.k8s.io/servicecatalog.k8s.io/apiserver-authentication-reader created
clusterrolebinding.rbac.authorization.k8s.io/system:cloud-controller-manager created
clusterrole.rbac.authorization.k8s.io/system:cloud-controller-manager created
daemonset.apps/vsphere-cloud-controller-manager created
```

```
❏ → kubectl get pods -A | grep vsphere
kube-system          vsphere-cloud-controller-manager-km68c    1/1   Running   0          3m4s
```

Create a namespace named **vmware-system-csi**:

```
❏ → kubectl create ns vmware-system-csi
namespace/vmware-system-csi created
```

Create a configuration file for the CSI driver:

```
cat /etc/kubernetes/csi-vsphere.conf
[Global]
cluster-id = "<cluster-id>"#Enter the cluster name
```

```
[VirtualCenter "<IP or FQDN>"] #Enter the IP address or FQDN of vCenter
insecure-flag = "<true or false>"#Specify the CA file and fingerprint if you set the value to false.
user = "<username>"
password = "<password>"
port = "<port>"
datacenters = "<datacenter1-path>"#Enter the path of the data center where the cluster nodes reside.
```

Generate a secret:

```
❏ → kubectl create secret generic vsphere-config-secret --from-file=csi-vsphere.conf --namespace=vmware-system-csi
```

Install **vsphere-csi-driver**:

```
❏ → kubectl apply -f https://raw.githubusercontent.com/kubernetes-sigs/vsphere-csi-driver/v2.6.0/manifests/vanilla/vsphere-csi-driver.yaml
```

Wait until all pods are running properly:

```
[root@k8s kubernetes]# kubectl get pods -A
NAMESPACE          NAME                                     READY   STATUS    RESTARTS   AGE
kube-system        calico-kube-controllers-676c86494f-sfg48 1/1     Running   0           15h
kube-system        calico-node-jdgdq                         1/1     Running   0           15h
kube-system        coredns-757cd945b-kbhpx                  1/1     Running   0           15h
kube-system        coredns-757cd945b-lq5sm                  1/1     Running   0           15h
kube-system        kube-apiserver-k8s                        1/1     Running   0           15h
kube-system        kube-controller-manager-k8s              1/1     Running   0           15h
kube-system        kube-proxy-45pf9                         1/1     Running   0           15h
kube-system        kube-scheduler-k8s                       1/1     Running   0           15h
kube-system        metrics-server-5468b66d4c-zbfbw          1/1     Running   0           5h44m
kube-system        nodelocaldns-fpd8h                       1/1     Running   0           15h
kube-system        openebs-localpv-provisioner-7ff868566c-hknkk 1/1     Running   0           5h7m
kube-system        snapshot-controller-0                    1/1     Running   0           15h
kube-system        vsphere-cloud-controller-manager-km68c    1/1     Running   0           4h57m
kubesphere-controls-system default-http-backend-659cc67b6b-t7898    1/1     Running   0           15h
kubesphere-controls-system kubectrl-admin-7966644f4b-2x4dp          1/1     Running   0           15h
kubesphere-monitoring-system alertmanager-main-0                      2/2     Running   0           15h
kubesphere-monitoring-system kube-state-metrics-75f7c75f86-6cn42       3/3     Running   0           15h
kubesphere-monitoring-system node-exporter-srmfr                       2/2     Running   0           15h
kubesphere-monitoring-system notification-manager-deployment-cdd656fd-gktb8 2/2     Running   0           15h
kubesphere-monitoring-system notification-manager-operator-7f7c564948-4xp67 2/2     Running   0           15h
kubesphere-monitoring-system prometheus-k8s-0                          2/2     Running   0           15h
kubesphere-monitoring-system prometheus-operator-684988fc5c-8lx44     2/2     Running   0           15h
kubesphere-system    ks-apiserver-5bc97d4496-fjlbb            1/1     Running   0           15h
kubesphere-system    ks-console-5ff9d8f9d-kv29v               1/1     Running   0           15h
kubesphere-system    ks-controller-manager-7647fb7bf-9jdx4    1/1     Running   0           15h
kubesphere-system    ks-installer-65bc964898-x8hv8            1/1     Running   0           15h
kubesphere-system    minio-69b778655d-shtsc                   1/1     Running   0           5h43m
vmware-system-csi    vsphere-csi-controller-6d5688c8d5-hj456  7/7     Running   0           25m
vmware-system-csi    vsphere-csi-node-mpfqp                   3/3     Running   0           3m32s
```

Query the status of csidriver and csinode:

```
🔗 → kubectl get csidriver
```

```
NAME          ATTACHREQUIRED  PODINFOONMOUNT  STORAGECAPACITY  TOKENREQUESTS  REQUIRESREUBLISH  MODES  AGE
csi.vsphere.vmware.com  true           false           false           <unset>        false             Persistent  85m
```

```
🔗 → kubectl get CSINODE
```

```
NAME  DRIVERS  AGE
k8s   1        16h
```

The vSphere CSI driver is installed. Then, configure the underlying storage and StorageClass.

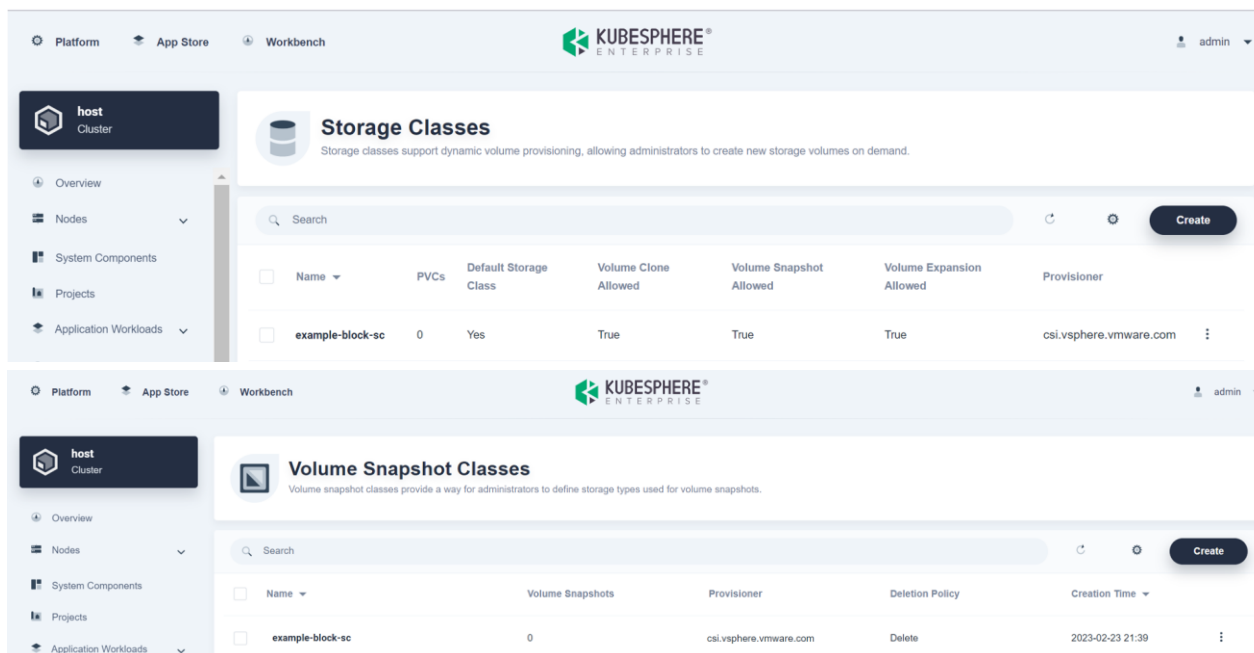
Create a StorageClass

CNS makes Kubernetes aware of how to dynamically provision PVs based on storage policies or StorageClasses.

Create a StorageClass:

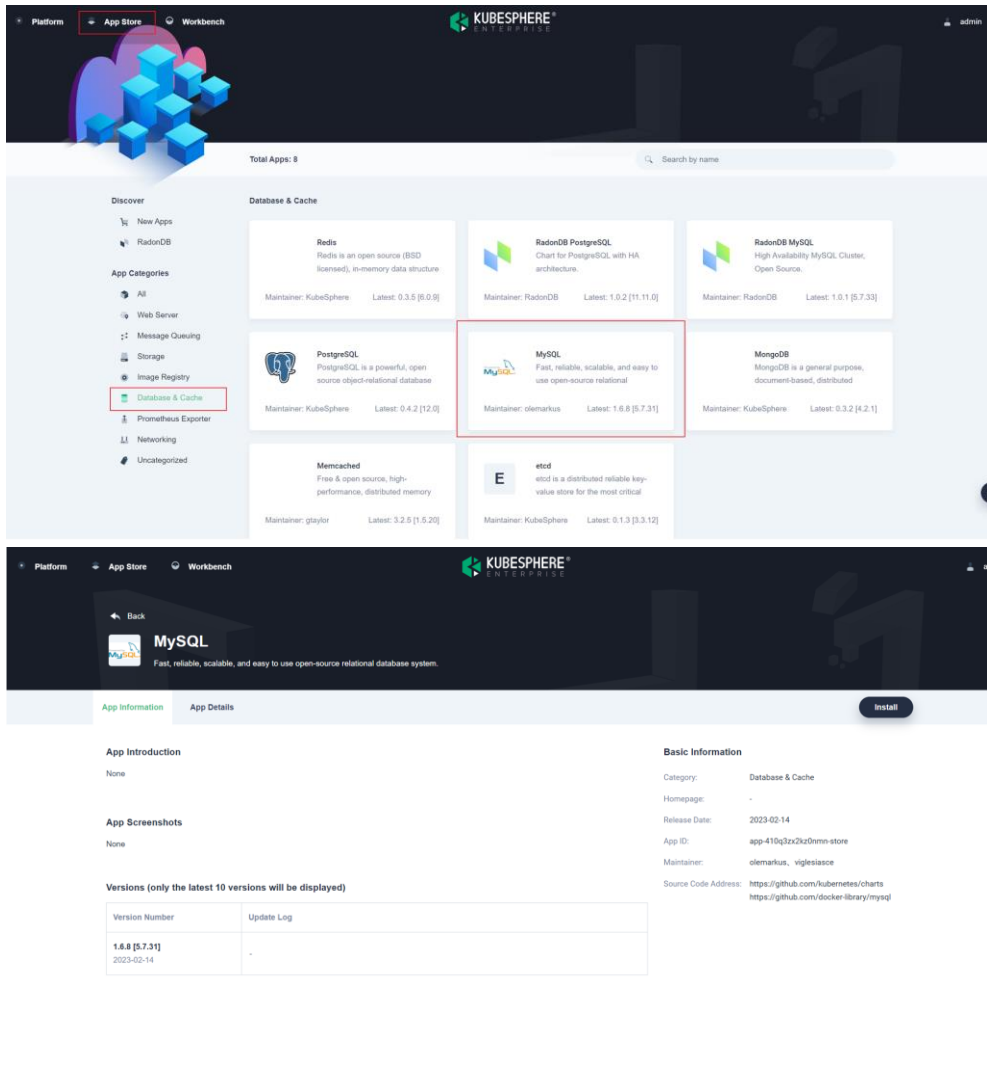
```
kind: StorageClass
apiVersion: storage.k8s.io/v1
metadata:
  name: example-block-sc
  annotations:
    storageclass.kubernetes.io/is-default-class: "true"
provisioner: csi.vsphere.vmware.com
parameters:
  storagepolicyname: "vSAN Default Storage Policy" #The storage policy.
  # datastoreurl: "ds:///vmfs/volumes/vsan:52cdfa80721ff516-ea1e993113acfc77/" #The data storage, for example, NFS mounted to ESXi,
  # FC/IP SAN, or local storage.
  # csi.storage.k8s.io/fstype: "ext4" #The formatted file type.
```

In the KubeSphere web console, check the storage class and snapshot class.



Run a Stateful Application

In the App Store of KubeSphere, deploy a MySQL application to check whether PVs work for vSphere.

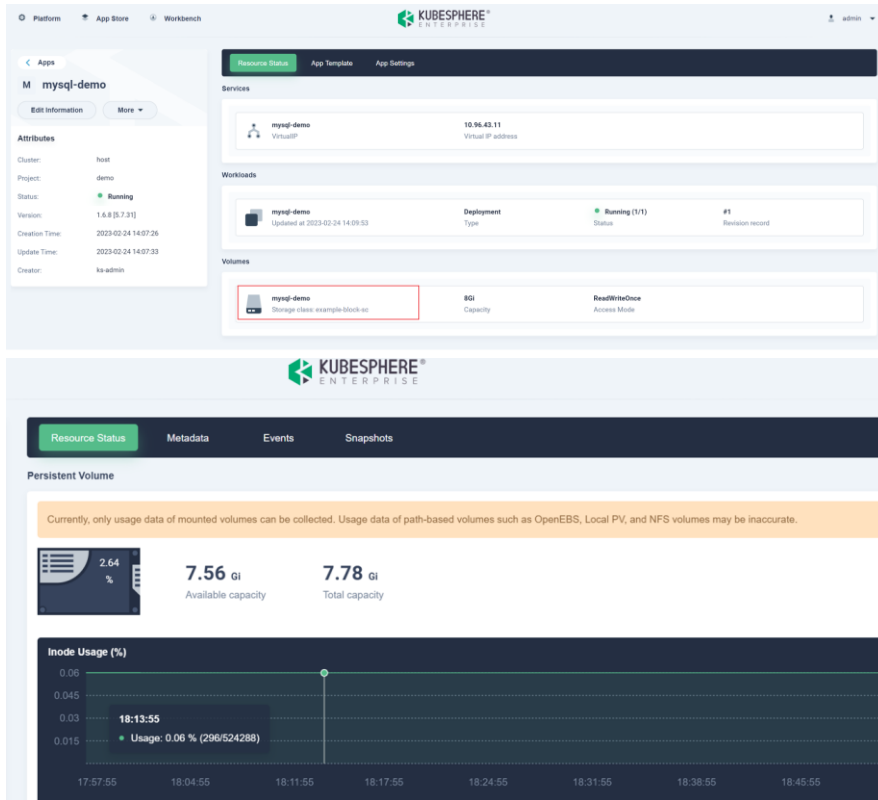


Set `storageClass` to `example-block-sc`, which was created based on the vSphere CSI driver:

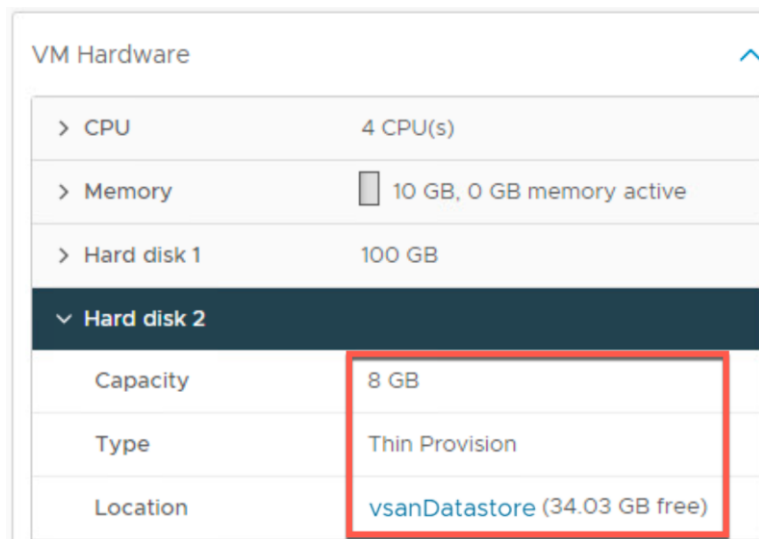
```

97 initialDelaySeconds: 5
98 periodSeconds: 10
99 timeoutSeconds: 1
100 successThreshold: 1
101 failureThreshold: 3
102
103 ## Persist data to a persistent volume
104 persistence:
105   enabled: true
106   ## database data Persistent Volume Storage Class
107   ## If defined, storageClassName: <storageClass>
108   ## If set to "-", storageClassName: "", which disables dynamic provisioning
109   ## If undefined (the default) or set to null, no storageClassName spec is
110   ## set, choosing the default provisioner. (gp2 on AWS, standard on
111   ## GKE, AWS & OpenStack)
112   ##
113   storageClass: "example-block-sc"
114   accessMode: ReadWriteOnce
115   size: 8Gi
116   annotations: {}
117
118 ## Use an alternate scheduler, e.g. "stork".
119 ## ref: https://kubernetes.io/docs/tasks/administer-cluster/configure-multiple-schedulers/
120 ##
121 # schedulerName:
    
```

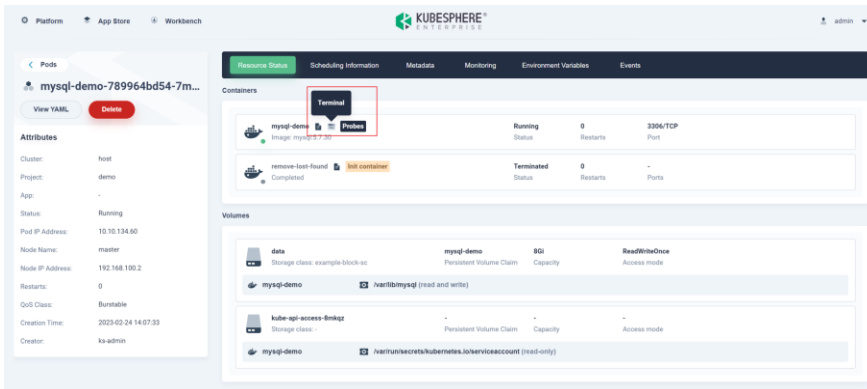
The following figure shows that the application with a PVC is created.



Go to vCenter, and you can view that an 8GB disk has been mounted to the VM. The capacity is the same as that defined in the PVC, and the backend storage is the same as that defined in the StorageClass.



Open Terminal to connect to the MySQL database:



Create a table and insert example values:

```
mysql> create database test;
Query OK, 1 row affected (0.00 sec)

mysql> use test;
Database changed
mysql> create table my_id(id int(4));
Query OK, 0 rows affected (0.01 sec)

mysql> insert my_id values(9527);
Query OK, 1 row affected (0.02 sec)

mysql> select * from my_id;
+-----+
| id   |
+-----+
| 9527 |
+-----+
1 row in set (0.00 sec)
```

Create a snapshot for the PVC:

The screenshot shows the Kubernetes Workbench interface for a Persistent Volume Claim named 'mysql-demo'. The left sidebar contains 'Attributes' such as Project, Status, Capacity, Access Mode, Storage Class, Provisioner, Persistent Volume, Creation Time, and Creator. A 'More' menu is open, highlighting the 'Create Snapshot' option. The main panel shows 'Resource Status' with available and total capacity, and 'Mounted Pods' with a table listing the pod 'mysql-demo-789964b54-7m2j6'.

Pod Name	Node	Pod IP address	CPU	Memory
mysql-demo-789964b54-7m2j6	master (192.168.100.2)	10.10.134.60	4 m	224.55 Mi

The 'Create Snapshot' dialog box is shown with the following details:

- Name:
- Volume Snapshot Class:

Instructions: The name can contain only lowercase letters, numbers, and hyphens (-), and must start and end with a lowercase letter or number. The maximum length is 253 characters. Select a snapshot class to create a snapshot of a specific type.

The screenshot shows the 'Snapshots' tab for the 'mysql-demo' Persistent Volume Claim. A table lists the snapshot 'mysql-snap' with a status of 'Created successfully', a capacity of '8Gi', and a creation time of '2023-02-24 21:17:28'.

Snapshot Name	Status	Capacity	Creation Time
mysql-snap	Created successfully	8Gi	2023-02-24 21:17:28

Create Persistent Volume Claim
Edit YAML

Basic Information
Done

Storage Settings
Current

Advanced Settings
To do

Creation Method

From Volume Snapshot
Select a snapshot to create a volume.

Volume Snapshot *

mysql-snap	8Gi	2023-02-24T13:17:28Z
example-block-sc	Capacity	Creation Time

Access Mode *

ReadWriteOnce(RWO)
ACCESS_MODE_RWO

ReadOnlyMany(ROX)
ACCESS_MODE_ROX

ReadWriteMany(RWX)
ACCESS_MODE_RWX

Cancel Previous Next

Insert data to the MySQL database:

```
mysql> use test;
Reading table information for completion of table and column names
You can turn off this feature to get a quicker startup with -A

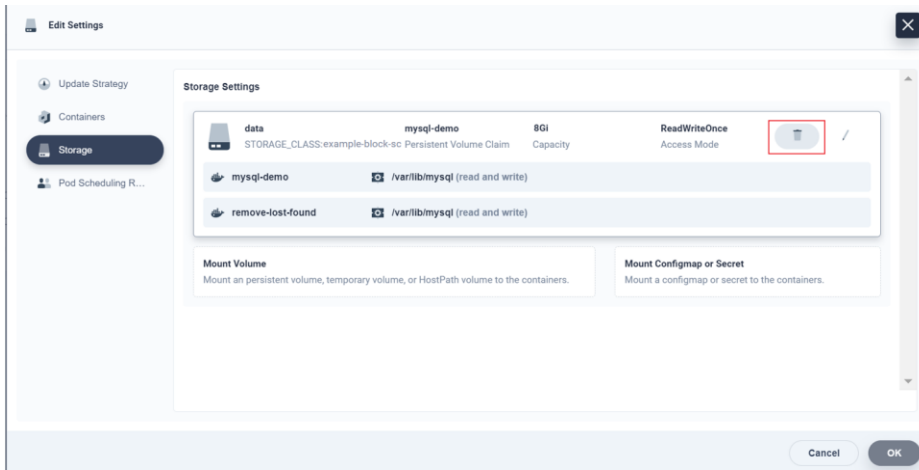
Database changed
mysql> create table your_id(id int(4));
Query OK, 0 rows affected (0.03 sec)

mysql> insert your_id values(9537);
Query OK, 1 row affected (0.01 sec)

mysql> select * from your_id;
+-----+
| id   |
+-----+
| 9537 |
+-----+
1 row in set (0.00 sec)

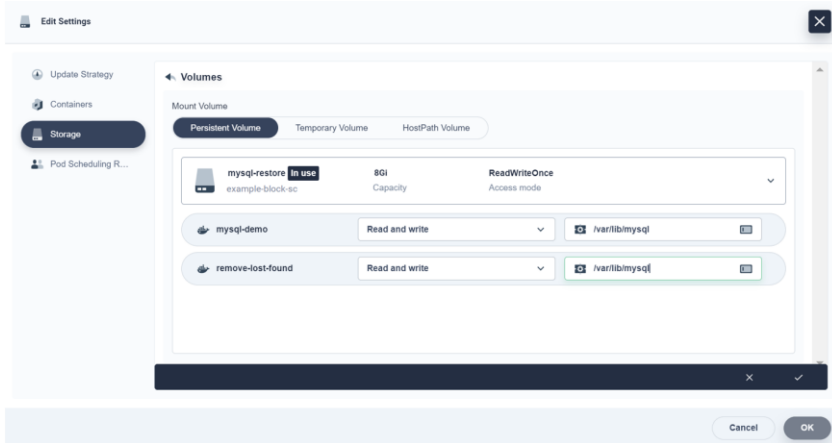
mysql> select * from my_id;
+-----+
| id   |
+-----+
| 9527 |
+-----+
1 row in set (0.00 sec)
```

Delete the PVC that is bound. Then, verify that the PVC cannot be found in the MySQL database.



```
mysql> use test;
ERROR 1049 (42000): Unknown database 'test'
mysql>
```

Use the snapshot created for the PVC, and connect to the MySQL database to query data:



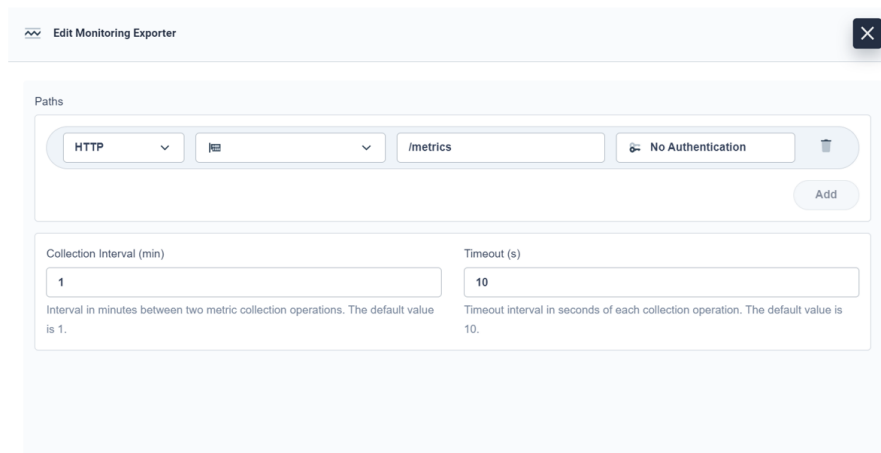
```
mysql> use test;
Reading table information for completion of table and column names
You can turn off this feature to get a quicker startup with -A

Database changed
mysql> select * from my_id;
+-----+
| id    |
+-----+
| 9527  |
+-----+
1 row in set (0.00 sec)

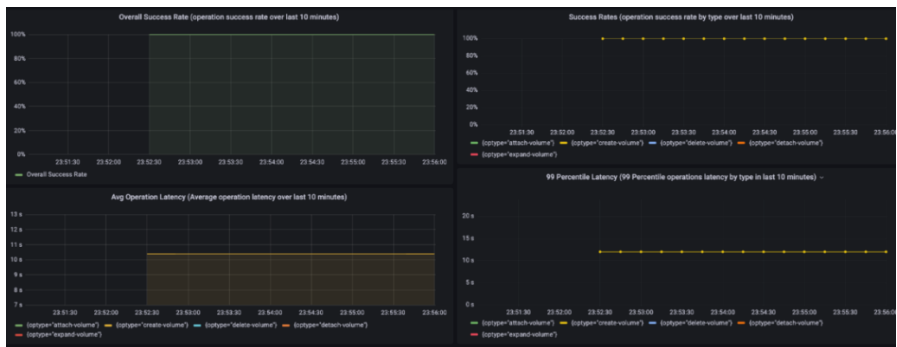
mysql> select * from your_id;
ERROR 1146 (42S02): Table 'test.your_id' doesn't exist
mysql> _
```

Monitoring

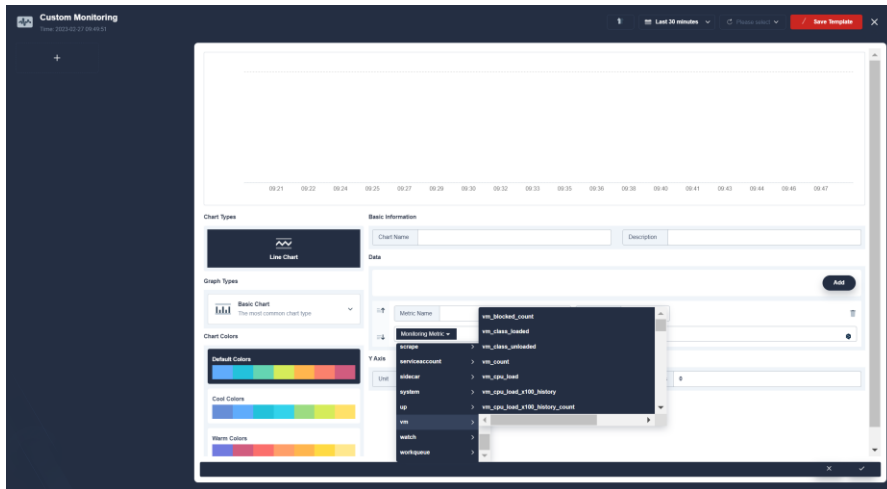
The vSphere CSI driver provides Prometheus-based monitoring metrics, including metrics for CNS components and storage. In the Kubernetes web console, you can use the monitoring exporter to create ServiceMonitor:



vSphere provides Grafana monitoring templates available for use:



In the Kubernetes web console, you can configure custom dashboards based on monitoring metrics, such as the number of times that block storage is successfully created.



Administrators can also view CNS monitoring data in vCenter.

Note: This section introduces the Cloud Native Storage (CNS) features in vSAN for standard or 'vanilla' Kubernetes clusters. See the link below for more information on CNS for vSphere:

<https://blogs.vmware.com/virtualblocks/2019/08/14/introducing-cloud-native-storage-for-vsphere/>

Integrate KubeSphere with Antrea

Deploy Antrea and antctl

```
kubectl apply -f https://raw.githubusercontent.com/antrea-io/antrea/main/build/yamls/antrea.yml
```

```
kubectl get pod -n kube-system | grep antrea
```

```
antrea-agent-4gbn8           2/2   Running   36d
antrea-agent-n8v2k           2/2   Running   36d
antrea-agent-wznw6           2/2   Running   36d
antrea-controller-879cc648-jgp5x 1/1   Running   40d
```

Use KubeKey to deploy a Kubernetes cluster and Antrea at ease

Create a cluster configuration file:

```
./kk create config -f config.yaml
```

```
vi config.yaml
```

```
apiVersion: kubekey.kubesphere.io/v1alpha2
```

```
kind: Cluster
```

```
metadata:
```

```
  name: example
```

```
spec:
```

```
  hosts:
```

```
  - {name: node1, address: 172.16.0.2, internalAddress: 172.16.0.2, privateKeyPath: "~/ssh/id_rsa"}
  - {name: node2, address: 172.16.0.3, internalAddress: 172.16.0.3, privateKeyPath: "~/ssh/id_rsa"}
  - {name: node3, address: 172.16.0.4, internalAddress: 172.16.0.4, privateKeyPath: "~/ssh/id_rsa"}
```

```

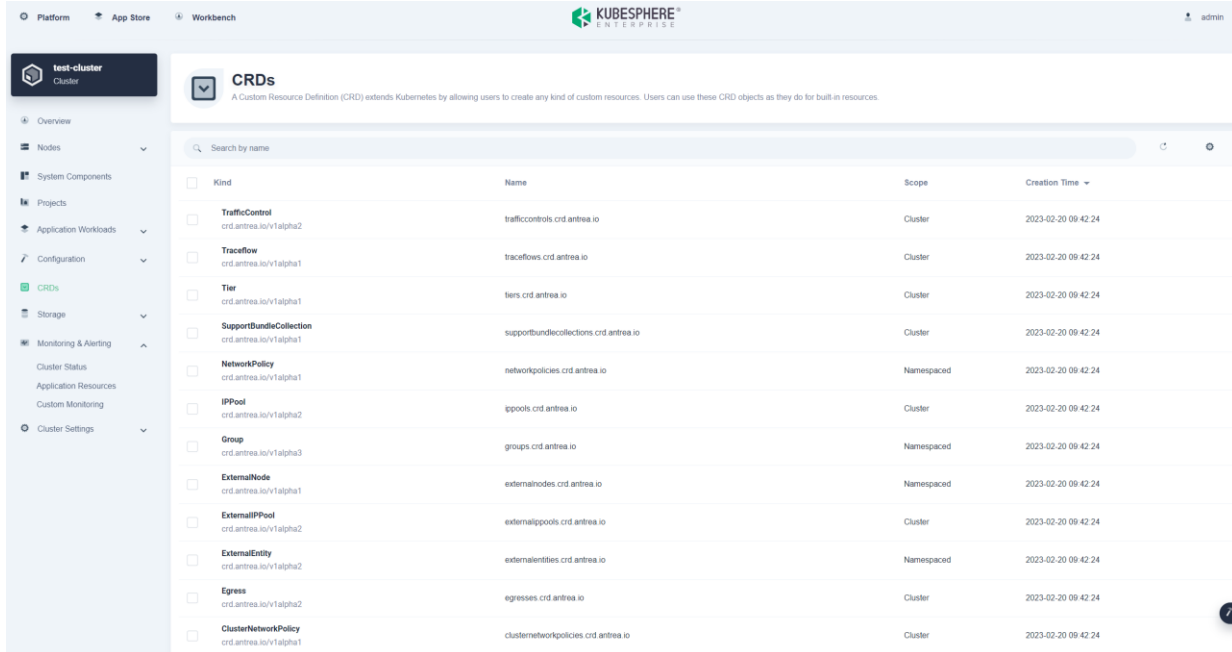
...
network:
  plugin: none
  kubePodsCIDR: 10.16.0.0/16
  kubeServiceCIDR: 10.96.0.0/16
...
addons:
  - name: antrea
    namespace: kube-system
sources:
  chart:
    name: antrea
    repo: https://charts.antrea.io

```

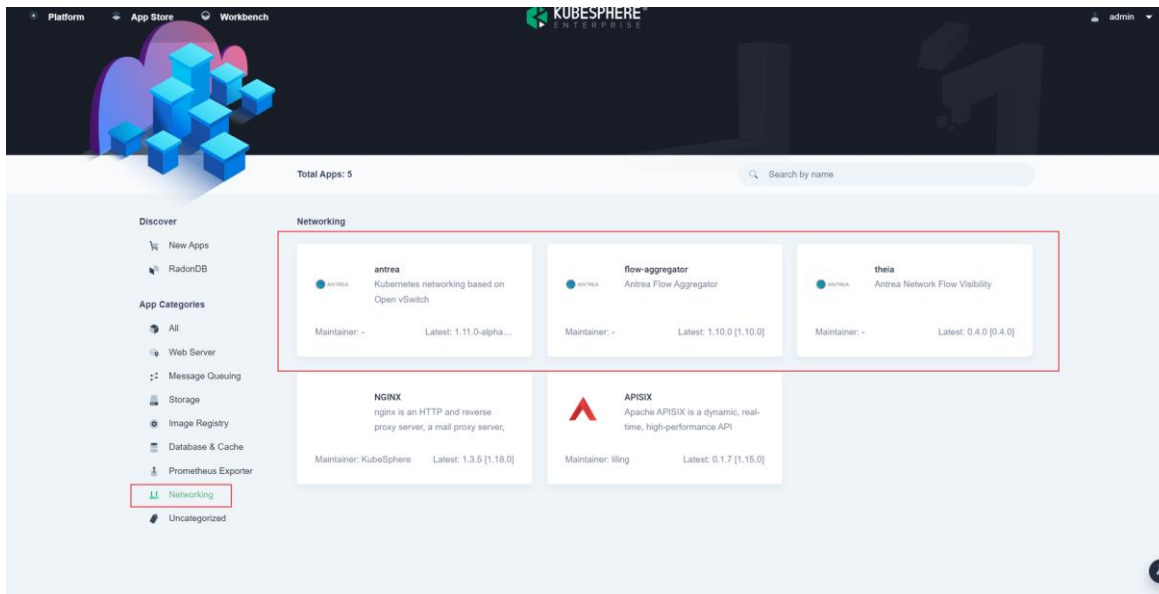
Run the `kk create cluster -f config.yaml` command to deploy the Kubernetes cluster and Antrea.

KubeKey supports using yaml files or Helm charts to install Antrea as an add-on. Support for network plugin is in assessment and might be implemented in future versions.

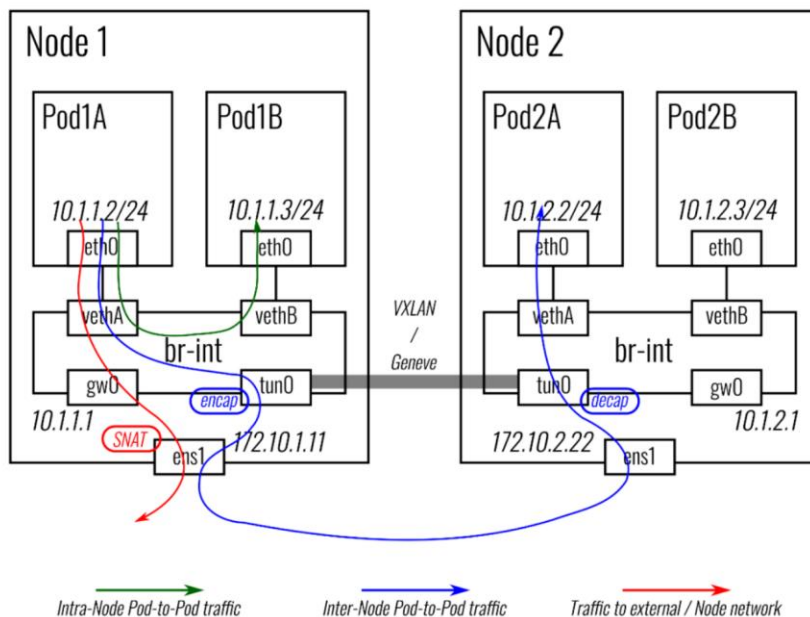
You can view the CRD resources related to Antrea in the KubeSphere web console.



Three projects for Antrea, Flow Aggregator, and Theia will be rolled out on KubeSphere App Store. You can directly install these applications from the App Store in this way.



Traffic walk



The image is cited from the Antrea official website, which can help you familiarize yourself with Antrea's communication mode.

- Intra-node traffic

Packets between two local pods will be forwarded by the OVS bridge directly.

- Inter-node traffic

Packets to a pod on another node will be first forwarded to the **antrea-tun0** port, encapsulated, and sent to the destination node through the tunnel; then they will be decapsulated, injected through the antrea-tun0 port to the OVS bridge, and finally forwarded to the destination pod.

- Pod to external traffic

Packets sent to an external IP or the Nodes' network will be forwarded to the **antrea-gw0** port (as it is the gateway of the local Pod subnet), and will be routed (based on routes configured on the Node) to the appropriate network interface of the Node (for example, a physical network interface for a bare metal node) and sent out to the Node network from there. Antrea Agent creates an iptables (MASQUERADE) rule to perform SNAT on the packets from Pods, so their source IP will be rewritten to the Node's IP before going out.

Security

Antrea ensures security from the following aspects:

- Control plane security

All API communication between Antrea control plane components is encrypted with TLS. You can generate the required certificate manually, or manage certificates by using cert-manager.

- Traffic encryption

Antrea supports encrypting tunnel traffic across nodes with IPsec ESP or WireGuard, to meet audit and security compliance for enterprises.

- Network policies

Network policies are native Kubernetes resources that include a whitelist of allowed egress rules. However, Kubernetes-native network policies do not support node-level control, events and logging, display of traffic rejection, cluster-level control, fine-grained action rules, and policy priorities.

Compared with Kubernetes-native network policies, Antrea abstracts a lot of network details and offers CRDs such as Tier, ClusterGroup, Group, ClusterNetworkPolicy, and NetworkPolicy. It defines policy priorities, groups resources by cluster or namespace, and supports cluster network policies. In contrast, Kubernetes-native network policies are similar to firewall rules for O&M personnel.

Antrea provides Tiers with specific priorities. You can attach a network policy to any Tier.

Emergency > SecurityOps > NetworkOps > Platform > Application > K8s NetworkPolicy > Baseline

In a ClusterGroup resource, you can specify labels, CIDR blocks, services, and member groups as needed.

Sample code:

```
apiVersion: crd.antrea.io/v1alpha3
kind: ClusterGroup
metadata:
  name: demo-cluster-group
spec:
  podSelector:
    matchLabels:
      role: db
  namespaceSelector:
    matchLabels:
      env: prod
  ipBlocks:
    - cidr: 10.0.10.0/24
  serviceReference:
    name: test-service
    namespace: default
  childGroups: [test-cg-sel, test-cg-ip-blocks, test-cg-svc-ref]
```

After the resource is defined, you can configure a network policy. In the following example, the policy priority is set to 5, and the policy applies to the resource group named **test-db**. Each ClusterNetworkPolicy might consist of zero or more ordered set of ingress or egress rules. Each rule, depending on the **action** field of the value, allows, drops, rejects, or passes traffic. Each policy can work

together with multiple selectors and resources, such as ports, ipBlock, services, FQDN, ICMP, IGMP, and HTTP. You can enable logging for resources. Logs are stored in `/var/log/antrea/networkpolicy/np.log`, and you can use tools such as Filebeat to collect and analyze logs. The design of NetworkPolicy is similar to that of ClusterNetworkPolicy. This way, you can implement secure network policies at Layer 3 to Layer 7.

Sample code:

```

apiVersion: crd.antrea.io/v1alpha1
kind: ClusterNetworkPolicy
metadata:
  name: acnp-demo
spec:
  priority: 5
  tier: securityops
  appliedTo:
    - group: "test-db"
  ingress:
    - action: Allow #[["Allow", "Drop", "Reject", "Pass"]]
      from:
        - podSelector:#[["podSelector", "namespaceSelector", "nodeSelector"]]
          matchLabels:
            role: frontend
        ports:
          - protocol: TCP
            port: 8080
            endPort: 9000
          name: AllowFromFrontend
          enableLogging: false
      egress:
        - action: Drop
          to:
            - ipBlock:
                cidr: 10.0.10.0/24
            ports:
              - protocol: TCP
                port: 5978
            # toServices:
            #   - name: svcName
            #     namespace: svcNamespace
            # protocols:
            #   - icmp:
            #       icmpType: 8
            #       icmpCode: 0
            #   l7Protocols:
            #     - http:
            #         path: "/api/v2/*"
            #         host: "foo.bar.com"
            #         method: "GET"

          name: DropToThirdParty
          enableLogging: true

```

IPAM

Antrea provides two types of IPAM capabilities, that is, NodeIPAM and AntreaIPAM.

NodeIPAM is a Kubernetes component, which manages IP address pool allocation per each node, when the node initializes. Antrea NodeIPAM controller can be executed in scenarios where the NodeIPAMController is disabled in kube-controller-manager.

When a Pod's IP is allocated from an IP Pool, the traffic from the Pod to Pods on another Node or from the Pod to external networks will be sent to the underlay network through the Node's transport network interface and will be forwarded or routed by the underlay network.

In real life practice, if you use underlay networks, use AntreaIPAM to allocate IP addresses and connect to physical networks. If you use overlay networks, use Multus for the second NIC and AntreaIPAM for VLAN connections.

Sample code:

```
{
  "cniVersion": "0.3.0",
  "name": "ipv4-net-1",
  "type": "macvlan",
  "master": "eth0",
  "mode": "bridge",
  "ipam": {
    "type": "antrea",
    "ippools": [ "ipv4-pool-1" ]
  }
}
```

```
kind: Pod
metadata:
  annotations:
    ipam.antrea.io/ippools: 'pod-ip-pool1'
```

Egress

When pods access critical applications such as databases outside the cluster, security audits are required. In this scenario, you can adopt egress gateways. Egress is a CRD API that manages external access from the Pods in a cluster. It supports specifying which egress (SNAT) IP the traffic from the selected Pods to the external network should use.

Sample code:

```
apiVersion: crd.antrea.io/v1alpha2
kind: Egress
metadata:
  name: egress-prod-web
spec:
  appliedTo:
    namespaceSelector:
      matchLabels:
        env: prod
    podSelector:
      matchLabels:
        role: web
  egressIP: 10.10.0.8
  externalIPPool: prod-external-ip-pool
```

You can use egressIP and externalIPPool to specify Egress resources. The specified IP address must be accessible over physical networks. High availability is provided automatically when the egressIP is allocated from an externalIPPool; for example, when the externalIPPool is specified. If the Node hosting the egressIP fails, another Node will be selected (from among the remaining Nodes selected by the nodeSelector of the externalIPPool) as the new egress Node of this Egress. It will take over the IP and send layer 2 advertisement (for example, Gratuitous ARP for IPv4) to notify the other hosts and routers on the network that the MAC address associated with the IP has changed.

NodePortLocal

In Kubernetes, you can use NodePort or LoadBalancer to expose services. LoadBalancers such as F5, Server Load Balancer (SLB), and OpenELB can help you create a service like ClusterIP and open a port on each node. NodePorts expose services on static ports that range from 30000 to 32767. However, the design of kube-proxy can cause issues in various scenarios, such as traffic distribution, session persistence, health check on external loads, and numerous ports occupied.

NodePortLocal creates mappings between *NodeIP:Port* and *PodIP:Port*, so each pod can access each other by using NodeIP and NodePort. This feature is suitable for the traditional business that transforms to cloud native, for example, clusters where microservices reside and interconnection between registries and clusters.

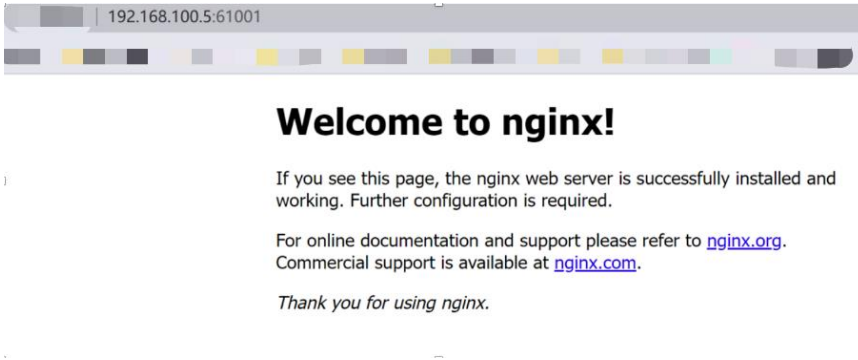
Add service annotations:

```
apiVersion: v1
kind: Service
metadata:
  name: nginx
  annotations:
    nodeportlocal.antrea.io/enabled: "true"
spec:
  ports:
  - name: web
    port: 80
    protocol: TCP
    targetPort: 80
  selector:
    app: nginx
```

View pod annotations:

```
[root@ks1-master ~]# kubectl get pods nginx-85b98978db-fwbmw -oyaml
apiVersion: v1
kind: Pod
metadata:
  annotations:
    nodeportlocal.antrea.io: '[{"podPort":80,"nodeIP":"192.168.100.5","nodePort":61001,"protocol":"tcp","protocols":["tcp"]}]'
  creationTimestamp: "2022-11-28T08:53:35Z"
  generateName: nginx-85b98978db-
  labels:
    app: nginx
    pod-template-hash: 85b98978db
    name: nginx-85b98978db-fwbmw
    namespace: default
```

Access the service in the *NodeIP:NodePort* format:



Query NAT rules:

```
conntrack -L -j |grep 10.10.1.24
tcp 6 431970 ESTABLISHED src=192.168.255.254 dst=192.168.100.5 sport=7149 dport=61001 packets=2 bytes=92 src=10.10.1.24
dst=192.168.255.254 sport=80 dport=7149 packets=1 bytes=52 [ASSURED] mark=0 delta-time=29 use=1
```

TraceFlow

Antrea supports using Traceflow for network diagnosis. You must be familiar with this feature if NSX is not new to you. Creating a new Traceflow CRD triggers the Traceflow module to inject packet into OVS, provide various observation points along the packet's path, and populate these observations into the status field of the Traceflow CRD.

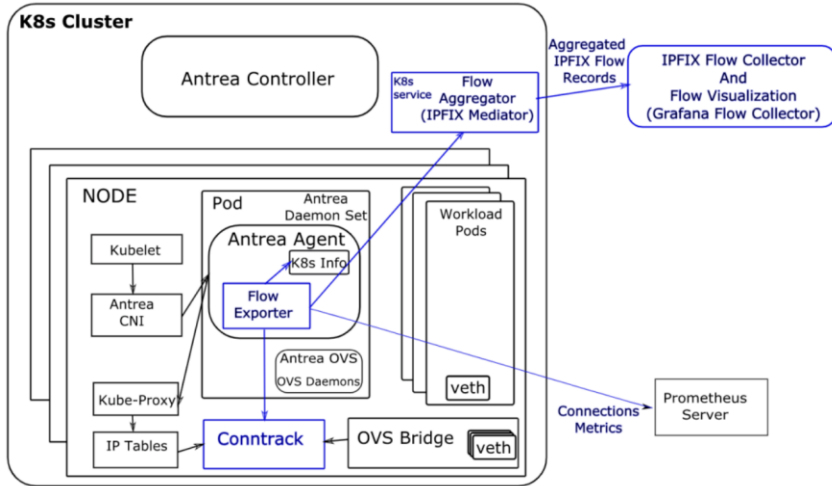
You can use `antctl` or `octant` to visualize Traceflow results:

```
[root@ks1-master ~]# antctl tf -S nginx-85b98978db-ckgtz -D nginx-85b98978db-fwbmw
name: nginx-85b98978db-ckgtz-to-nginx-85b98978db-fwbmw-rsc7pnr6
phase: Succeeded
source: default/nginx-85b98978db-ckgtz
destination: default/nginx-85b98978db-fwbmw
results:
- node: ks1-worker2
  timestamp: 1669629401
  observations:
  - component: Forwarding
    action: Received
  - component: Forwarding
    componentInfo: Output
    action: Delivered
- node: ks1-worker1
  timestamp: 1669629401
  observations:
  - component: SpoofGuard
    action: Forwarded
  - component: Forwarding
    componentInfo: Output
    action: Forwarded
  tunnelDstIP: 192.168.100.5
```


Network observability

To visualize network traffic, Antrea offers Flow Aggregator and Theia. Antrea monitors the flows in Linux conntrack module. These flows are converted to flow records, and then flow records are post-processed before they are sent to the configured external flow

collector. In Antrea, the basic building block for the Network Flow Visibility is the Flow Exporter. Connections from the connection store are exported to the Flow Aggregator Service using the IPFIX protocol and displayed on Grafana through Theia.



Use KubeSphere App Store to deploy Flow Aggregator and Theia:



App Introduction

None

App Screenshots

None

Versions (only the latest 10 versions will be displayed)

Version Number	Update Log
1.10.0 [1.10.0] 2023-02-17	Antrea Flow Aggregator

Basic Information

Category: Networking
 Homepage: -
 Release Date: 2023-02-17
 App ID: app:m1mwn5kovlm1m-store
 Source Code Address: <https://github.com/antrea-io/antrea>



Basic Information


Name * Version *

The name can contain only lowercase letters, numbers, and hyphens (-), must start with a lowercase letter, and must end with a lowercase letter or number. The maximum length is 32 characters.

Description

The description can contain any characters and the maximum length is 256 characters.

Location



App Settings

```

11 address: ""
12 # -- Provide the 32-bit observation domain ID which will uniquely identify this instance of the flow
13 # aggregator to an external flow collector. If omitted, an observation domain ID will be generated
14 # from the persistent cluster UUID generated by antrea.
15 observationDomainID: ""
16 # -- Provide format for records sent to the configured flow collector.
17 # Supported formats are INFLUX and JSON.
18 format: default "INFLUX"
19 # Clickhouse contains clickhouse related configuration options.
20 clickhouse:
21 # -- Determine whether to enable exporting flow records to Clickhouse.
22 enable: true
23 # -- Database is the url to the database, TCP protocol is required.
24 databaseURL: "tcp://clickhouse-clickhouse.flow-visibility.svc:9000"
25 # -- Debug enables debug logs from Clickhouse sql driver.
26 debug: false
27 # -- Compress enables LZ4 compression when committing flow records.
28 compress: true
29 # -- commitInterval is the periodical interval between batch commit of flow records to DB.
30 # Valid time units are "ms", "s" (or "sec"), "m", "h", "d", "w".
31 commitInterval: "5s"
32 # -- Credentials to connect to Clickhouse, they will be stored in a Secret.
33 connectionSecret:
    
```



App Introduction

None

App Screenshots

None

Versions (only the latest 10 versions will be displayed)

Version Number	Update Log
0.4.0 [0.4.0]	Antrea Network Flow Visibility

Basic information

Category: Networking
Homepage: -
Release Date: 2023-02-17
App ID: app-nfmxyq250q5352-store
Source Code Address: <https://github.com/antrea-io/theia>



Basic Information

Name: Version:

The name can contain only lowercase letters, numbers, and hyphens (-), must start with a lowercase letter, and must end with a lowercase letter or number. The maximum length is 32 characters.

Description:

The description can contain any characters and the maximum length is 256 characters.

Location

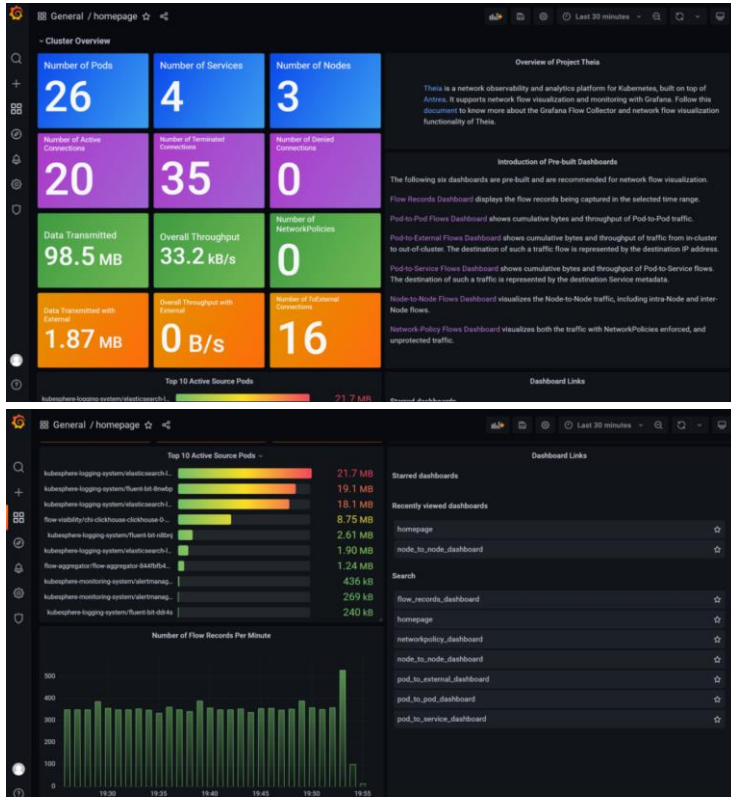
demo Workspace | host Cluster | demo Project

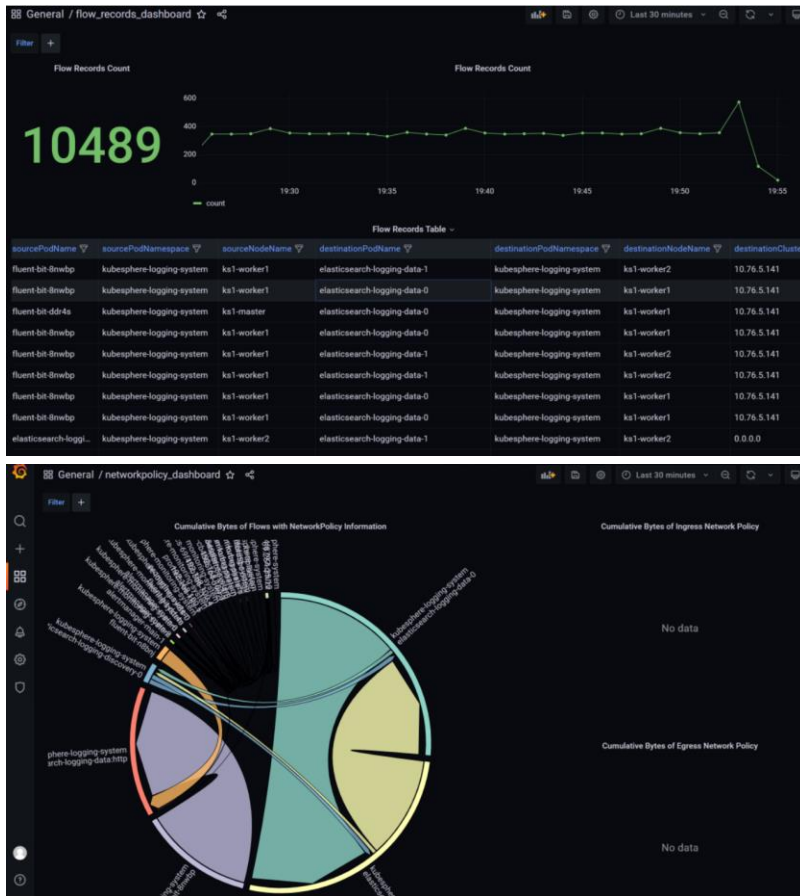


App Settings

```
222 # volumeName: "my-pvc".
223 # To dynamically provision a PersistentVolume, please set
224 # storageClassName: "my-storage-class".
225 # HostPath storage is used if both createPersistentVolume.type and
226 # persistentVolumeClaimSpec are empty.
227 persistentVolumeClaimSpec: {}
228 # StorageClassName: ""
229 # volumeName: ""
230
231 sparkOperator:
232 # ... Determine whether to install Spark Operator, it is required to run Network
233 # Policy Recommendation jobs.
234 enable: true
235 # ... Name of Spark Operator.
236 name: "policy-recommendation"
237 # ... Container image used by Spark Operator.
238 image:
239   repository: "projects.registry.vmware.com/antrea/theia-spark-operator"
240   pullPolicy: "IfNotPresent"
241   tag: "theia2-1.3.3-3.1.1"
242
243 theiaManager:
244 # ... Determine whether to install Theia Manager.
245 enable: false
246 # ... Container image used by Theia Manager.
247 image:
248   repository: "projects.registry.vmware.com/antrea/theia-manager"
```

View statistical data on Grafana:





It also provides suggestions on network policies and security hardening for platform components and system architectures.

Sample code:

```

theia policy-recommendation run
theia policy-recommendation list
CreationTime      CompletionTime  ID                               Status
2022-10-08 06:35:04 2022-10-08 06:38:00 60a79979-19ed-40a7-bc11-2b389a8a43bb COMPLETED
theia policy-recommendation retrieve --id 60a79979-19ed-40a7-bc11-2b389a8a43bb
apiVersion: crd.antrea.io/v1alpha1
kind: ClusterNetworkPolicy
metadata:
  name: recommend-allow-acnp-kube-system-f2g5l
spec:
  appliedTo:
  - namespaceSelector:
    matchLabels:
      kubernetes.io/metadata.name: kube-system
  egress:
  - action: Allow
  to:
  - podSelector: {}
  ingress:
  - action: Allow
    
```

```

    from:
    - podSelector: {}
    priority: 5
    tier: Platform
---
apiVersion: crd.antrea.io/v1alpha1
kind: ClusterNetworkPolicy
metadata:
  name: recommend-allow-acnp-flow-aggregator-uny5c
spec:
  appliedTo:
  - namespaceSelector:
    matchLabels:
      kubernetes.io/metadata.name: flow-aggregator
  egress:
  - action: Allow
    to:
    - podSelector: {}
  ingress:
  - action: Allow
    from:
    - podSelector: {}
  priority: 5
  tier: Platform
---
apiVersion: crd.antrea.io/v1alpha1
kind: ClusterNetworkPolicy
metadata:
  name: recommend-allow-acnp-flow-visibility-ecj61
spec:
  appliedTo:
  - namespaceSelector:
    matchLabels:
      kubernetes.io/metadata.name: flow-visibility
  egress:
  - action: Allow
    to:
    - podSelector: {}
  ingress:
  - action: Allow
    from:
    - podSelector: {}
  priority: 5
  tier: Platform

```

Multi-cluster networking

Antrea Multi-cluster implements connectivity between pods or services. A Multi-cluster ClusterSet is comprised of a single leader cluster and at least two member clusters. Then, pods can access each other via pod IP addresses, and services can access each other via import or export.

Assume that the following two clusters exist:

Cluster A: Apiserver 192.168.100.3:6443, podCIDR 10.10.0.0/16

Cluster B: Apiserver 192.168.100.6, podCIDR 10.20.0.0/16

Enable the multi-cluster feature:

```
kubectl edit cm -n kube-system
# Enable Antrea Multi-cluster Gateway to support cross-cluster traffic.
# This feature is supported only with encap mode.
Multicluster: true
```

Create a namespace named **antrea-multicluster** and use **antctl** to deploy ClusterSet:

```
kubectl create ns antrea-multicluster
antctl mc deploy leadercluster -n antrea-multicluster --antrea-version v1.10.0
antctl mc deploy membercluster -n kube-system --antrea-version v1.10.0
```

Create a ClusterSet in cluster A, which works as the leader cluster and has two member clusters.

```
antctl mc init --clusterset test-clusterset --clusterid test-cluster-leader -n antrea-multicluster --create-token -j join-config.yml
antctl mc join --clusterid test-cluster-leader -n kube-system --config-file join-config.yml
```

Query ClusterClaim and ClusterSet.

```
kubectl get clustersets.multicluster.crd.antrea.io -A
```

NAMESPACE	NAME	LEADER CLUSTER	NAMESPACE	TOTAL CLUSTERS	READY CLUSTERS	AGE
antrea-multicluster	test-clusterset	antrea-multicluster	1	1	4m12s	
kube-system	test-clusterset	antrea-multicluster	1	1	4m4s	

```
kubectl get clusterclaims.multicluster.crd.antrea.io -A
```

NAMESPACE	NAME	VALUE	AGE
antrea-multicluster	clusterset.k8s.io	test-clusterset	4m32s
antrea-multicluster	id.k8s.io	test-cluster-leader	4m32s
kube-system	clusterset.k8s.io	test-clusterset	4m24s
kube-system	id.k8s.io	test-cluster-leader	4m24s

Specify the Gateway Node for cluster A, which is responsible for routing all cross-clusters traffic from the local cluster to other member clusters through tunnels.

```
kubectl annotate node ks1-master multicluster.antrea.io/gateway=true
```

Deploy member cluster B.

```
antctl mc deploy membercluster -k ~/.kube/config1 -n kube-system --antrea-version v1.10.0
```

Add cluster B to ClusterSet.

```
antctl mc deploy membercluster -k ~/.kube/config1 -n kube-system --antrea-version v1.9.0
```

Configure the Gateway Node for cluster B.

```
kubectl annotate node ks2 multicluster.antrea.io/gateway=true
```

Check the cluster status:

```
antctl mc get clusterset -A
```

CLUSTER-ID	NAMESPACE	CLUSTERSET-ID	TYPE	STATUS	REASON
test-cluster-leader	antrea-multicluster	test-clusterset	Ready	True	Connected
test-cluster-leader	kube-system	test-clusterset	IsLeader	True	<NONE>
test-cluster-leader	kube-system	test-clusterset	Ready	True	<NONE>
test-cluster-member	antrea-multicluster	test-clusterset	Ready	True	Connected

Check cluster import:

```
kubectl get clusterinfoimport -n kube-system
```

NAME	CLUSTER ID	SERVICE CIDR	AGE
test-cluster-member-clusterinfo	test-cluster-member	10.86.0.0/16	28m

```
[root@ks2 ~]# kubectl get clusterinfoimport -n kube-system
NAME                CLUSTER ID      SERVICE CIDR  AGE
test-cluster-leader-clusterinfo  test-cluster-leader  10.76.0.0/16  32m
```

In **antrea-mc-controller**, specify **podCIDR** for which you want to enable connectivity:

```
kubectl edit configmap -n kube-system antrea-mc-controller-config
data:
  controller_manager_config.yaml: |
    apiVersion: multicluster.crd.antrea.io/v1alpha1
    kind: MultiClusterConfig
    health:
      healthProbeBindAddress: :8080
    metrics:
      bindAddress: "0"
    webhook:
      port: 9443
    leaderElection:
      leaderElect: false
    serviceCIDR: ""
    podCIDRs:
      - "10.10.0.0/16" #另一边为 10.20.0.0/16
    gatewayIPPrecedence: "private"
    endpointIType: "ClusterIP"
```

Test pod connectivity:

```
kubectl get pods -o wide
NAME    READY   STATUS    RESTARTS   AGE   IP           NODE           NOMINATED NODE   READINESS GATES
box2    1/1     Running   0          5m29s  10.10.2.2    ks1-worker1    <none>           <none>
kubectl get pods -o wide
NAME    READY   STATUS    RESTARTS   AGE   IP           NODE           NOMINATED NODE   READINESS GATES
box1    1/1     Running   0          5m14s  10.20.0.4    ks2            <none>           <none>
kubectl exec box1 ping 10.10.2.2
PING 10.10.2.2 (10.10.2.2): 56 data bytes
64 bytes from 10.10.2.2: seq=0 ttl=61 time=2.832 ms
```

Best Practices

This guide leverages the advantages of VMware cloud-native infrastructure and QingCloud KubeSphere container platform to provide users with a full-lifecycle cloud-native solution. For information about implementation and design, see the following best practices:

- Best practices for cloud native storage
 - [Best practices for deploying VMware CNS](#)
- Best practices for cloud native networking:
 - [Best practices for deploying Antrea](#)
- Best practices for deploying KubeSphere on VMware vSphere
 - [Deploy KubeSphere on VMware vSphere](#)
- Best practices for container management
 - [Manage containers on KubeSphere](#)
- DevOps Best Practices

- o [Best practices for KubeSphere DevOps](#)

Conclusion

Powered by QingCloud, KubeSphere can work together with VMware in a seamless manner to provide cloud native solutions, for example, integration with vSAN and CNS to provide storage solutions, and with Antrea to provide network solutions. VMware features high performance, high availability, stability, and scalability, and KubeSphere is scalable, open source, and native to clouds. This indicates that KubeSphere can help you run containers on VMware and minimize the learning costs and O&M pressure, so you can care about only the upper-layer applications and accelerate digital transformation of your business.

Reference

- [VMware CNS](#)
- [Antrea](#)
- [KubeSphere](#)

About the Author

Wei MA, Solution Architect at KubeSphere, QingCloud

Shuang YU, Product Director at QingCloud

The following reviewers also contributed to the paper contents:

- Fei LIU, Senior Solution Architect at VMware
- Jiali XU, Senior Manager at Cloud Infrastructure Solutions Department, VMware
- Qin XU, Senior Solution Architect at Cloud Infrastructure Solutions Department, VMware
- Ting YIN, Senior Solution Architect at Cloud Infrastructure Solutions Department, VMware
- Min ZHANG, Senior Network Security Product Expert at VMware



vmware®

 QingCloud
Technologies

VMware, Inc. 3401 Hillview Avenue Palo Alto CA 94304 USA Tel 877-486-9273 Fax 650-427-5001 www.vmware.com.
Copyright © 2023 VMware, Inc. All rights reserved. This product is protected by U.S. and international copyright and intellectual property laws. VMware products are covered by one or more patents listed at vmware.com/go/patents. VMware is a registered trademark or trademark of VMware, Inc. and its subsidiaries in the United States and other jurisdictions. All other marks and names mentioned herein may be trademarks of their respective companies. Item No: vmw-wp-tech-temp-word-102-proof 5/19