VMware Container Networking with Antrea

At a glance

VMware Container Networking™ with Antrea™ offers users signed images and binaries, along with full enterprise support for Project Antrea, a Cloud Native Computing Foundation sandbox project. VMware Container Networking integrates with managed Kubernetes services to further enhance Kubernetes network policies. It also supports Windows and Linux workloads on Kubernetes across multiple clouds.

Key benefits

- Simplifies Kubernetes networking by using a unified network stack across Kubernetes, regardless of cloud platform
- Improves application service load balancing performance through native Open vSwitch service proxy implementation
- Ensures secure pod connectivity with enforcement of Kubernetes network policies and advanced native policies
- Improves container security by encrypting traffic between pods despite running on untrusted fabrics
- Improves visibility and equips operations with network diagnostic tools such as Traceflow

Business challenges with current Kubernetes networking

Community support lacks predefined SLAs

Enterprises benefit from collaborative engineering and receive the latest innovations from open source projects. However, it is a challenge for any enterprise to rely solely on community support to run their operations, namely because community support is a best effort and cannot provide a predefined service-level agreement (SLA).

Solution fragmentation, incompatibility and incompleteness

Container network interfaces (CNIs) may become incompatible with the applied Kubernetes version. There is a minimum Kubernetes version that each CNI supports, and older versions of Kubernetes are typically not supported.

Project viability, loss of contributors

Open source projects can sometimes languish due to low user adoption or the loss of core contributors. Antrea has active contributors from Intel, Mellanox/Nvidia and VMware. One of the surefire ways to maintain project viability is through widespread user adoption. Another way is to design the open source project into a broader managed Kubernetes solution that already has an installed user base, such as Amazon Elastic Kubernetes Service (Amazon EKS), Azure Kubernetes Service (AKS), or Google Kubernetes Engine (GKE).

VMware Container Networking with Antrea solves these problems

VMware Container Networking with Antrea provides the assurance of signed images and binaries with full enterprise support backed by VMware. Antrea maintains active contributors from the community, including Intel, Mellanox, Nvidia and VMware. Because Antrea is designed into the VMware Tanzu® portfolio and VMware vSphere® 7 with VMware Tanzu, there already exists an installed user base for VMware Container Networking. Customers with valid licenses for VMware NSX® Advanced receive VMware Container Networking at no extra charge.

VMware Container Networking provides support for the latest conformant Kubernetes and stable releases of Antrea. It closely follows with open source and the release cadence of Kubernetes.

Benefits of VMware Container Networking with Antrea

- Simplify Kubernetes networking with a unified networking stack across multiple managed Kubernetes providers. You can use Antrea across your on-premises clouds, public clouds and edge clouds.
- Improve application performance for Windows and Linux workloads with load balancing enhancements through Open vSwitch. Antrea accelerates packet processing performance by offloading the network data plane to SmartNICs for execution. Aided by SmartNICs, Antrea provides secure, high-performance networking to support CPU-intensive use cases, such as big data and machine learning.
- Operate easily as Antrea seamlessly integrates with Prometheus and monitors CRDs to observe control and data plane health. Platform operators can use diagnostic features, such as Traceflow, to aid in troubleshooting and root cause analysis.
- Improve container security by encrypting traffic between pods despite running on untrusted fabrics.
- Get comprehensive, enterprise-class support, backed by VMware Support, for the most stable releases of Antrea that comply with Cloud Native Computing Foundation specifications.

What is Project Antrea?

Antrea is a Kubernetes-native project that implements the CNI and Kubernetes NetworkPolicy to provide network connectivity and security for pod workloads.

Antrea uses Open vSwitch as the networking data plane in every Kubernetes node. Due to the programmable characteristic of Open vSwitch, Antrea extends the benefits of programmable networks and performance from Open vSwitch to Kubernetes.

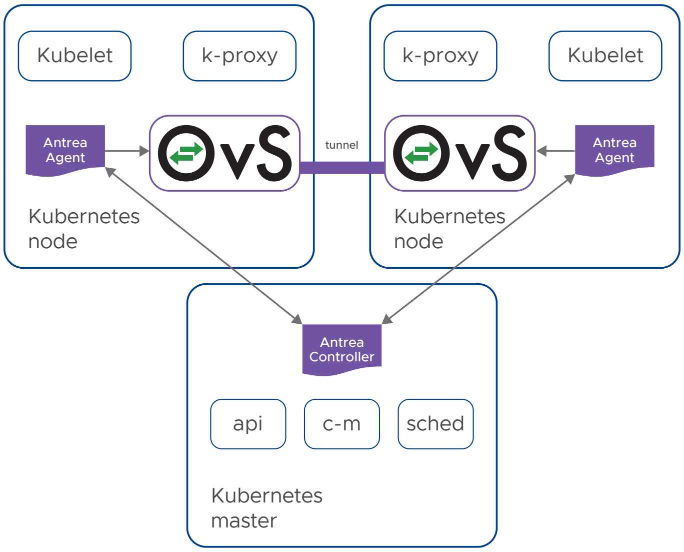

Figure 1: Antrea and Open vSwitch.

What is Open vSwitch?

Open vSwitch is a high-performance, programmable virtual switch that supports Windows and Linux workloads.

What is Kubernetes NetworkPolicy?

By default, all Kubernetes pods can communicate with each other. Applying a NetworkPolicy to a given pod isolates it, meaning it can only send traffic to, or receive traffic from, a pod that has been explicitly selected.

Antrea implements more than just the CNI. Other CNIs for Kubernetes merely provide network connectivity, but Antrea provides NetworkPolicy enforcement that empowers admins to implement fine-grain controls over pod traffic.

Antrea can augment pod networking solutions in managed Kubernetes services, such as AKS, Amazon EKS and GKE.

Use cases for VMware Container Networking for Antrea

Complete container networking and security solutions for VMware Tanzu and vSphere with Tanzu

Designed into the VMware Tanzu portfolio, VMware Container Networking is the default networking solution for VMware Tanzu Kubernetes Grid™, Tanzu Kubernetes Grid Integrated Edition, and vSphere 7 with Tanzu.

NetworkPolicy enforcement for managed Kubernetes services

The NetworkPolicy feature in Kubernetes allows you to define rules for ingress and egress traffic between pods in the cluster.

VMware Container Networking can enforce network policies for managed Kubernetes services, such as Amazon EKS, AKS and GKE. This enables users to augment their existing network policy implementation with advanced policy-tiering options.

Hardware offload for CPU-intensive workloads

Accelerate CPU-intensive workloads, such as big data and machine learning, by offloading to hardware. NICs have been preconfigured for specific functions or SmartNICs, ensuring high performance and flexible data processing.

Container security for Kubernetes

Provide policy enforcement on managed Kubernetes clusters from public cloud providers, and encrypt traffic between pods.

Feature	Feature status	Project Antrea	VMware Container Networking Standard with Antrea	VMware Container Networking Advanced with Antrea	VMware Container Networking Enterprise with Antrea
Connectivity					
High-performance, pod-to-pod networking IPv4 and IPv6 (dual stack)	GA	•	•	•	•
Choice of routed, encapsulated (overlay), hybrid and cloud provider routing mode of operation	GA	•	•	•	•
Choice of encapsulation (overlay) modes (Geneve, VXLAN, STT, GRE)	GA	•	•	•	•
High-performance, low-latency Kubernetes service implementation in Open vSwitch (kube-proxy/iptables replacement)	Beta/GA	•	•	•	•
Advanced load balancer integration through pod-specific NodePort (node port local)	Beta/GA	•	•	•	•
Platform support					
Windows containers network dataplane	GA	•	•	•	•
ARM Kubernetes node support	Alpha	•	N/A	N/A	N/A
OpenShift Operator	GA				•
Amazon EKS	GA	•	•	•	•
Azure Kubernetes Service, AKS engine	GA	•	•	•	•
Google Kubernetes Engine	GA	•	•	•	•

Feature	Feature status	Project Antrea	VMware Container Networking Standard with Antrea	VMware Container Networking Advanced with Antrea	VMware Container Networking Enterprise with Antrea
Security					
Kubernetes NetworkPolicy support	GA	•	•	•	•
Cluster-wide network policy support (Antrea network policy)	GA	•	•	•	•
Cluster-wide network policy tiering and nesting	GA	•	•	•	•
Role-based access control for policy tiers	GA			•	•
FQDN/DNS-based egress policy with wildcard matching	Alpha	•	•	•	•
Deny network policy	GA	•	•	•	•
Network policy statistics export	GA	•	•	•	•
Advanced security and manageability					
Multi-cluster Antrea network policy management through NSX-T management plane	Requires NSX-T 3.2				•
Central Kubernetes object inventory through NSX-T management plane	Requires NSX-T 3.2				•
Central multi-cluster connectivity troubleshooting via Traceflow through NSX-T management plane	Requires NSX-T 3.2				•

Feature	Feature status	Project Antrea	VMware Container Networking Standard with Antrea	VMware Container Networking Advanced with Antrea	VMware Container Networking Enterprise with Antrea
Manageability					
Traceflow pod-to-pod connectivity visualization	GA	•	•	•	•
ID-aware network metrics and flow export	GA	•	•	•	•
Prometheus load and flow metrics	GA	•	•	•	•
Monitoring CRDs for control plane status and health	GA	•	•	•	•
Octant UI plug-in for traceflows, cluster inventory, and monitoring control plane health	GA	•	•	•	•
Antrea network policy audit logs	GA	•	•	•	•
Advanced troubleshooting CLI	GA	•	•	•	•
Hardening, services and support					
FIPS-compliant product release	GA		•	•	•
VMware Tanzu Observability™ by Wavefront integration	GA			•	•
24x7 enterprise support	GA		•	•	•

Figure 2: VMware Container Networking with Antrea licensing information.

Primer on Kubernetes networking

While Kubernetes does not provide a default networking implementation, it does provide a model for implementing third-party tools, known as a CNI.

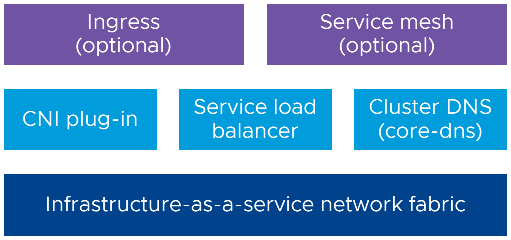

Figure 3: Kubernetes networking in layers.

CNI plug-in

The main responsibility of a CNI plug-in is to build the pod network and provide connectivity.

The Kubernetes service load balancer enables you to define a logical set of pods and an access policy as a service. You can then use services and service types to precisely control how your applications receive traffic. The load balancer service type creates an external load balancer in the public cloud infrastructure and assigns a fixed, external IP to the service. Then, authorized users can access the service via the exposed IP address.

Cluster DNS for services and pods

Kubernetes provides its own DNS service to resolve domain names inside the cluster, which allows pods to communicate with each other. This is implemented by deploying a regular Kubernetes service that handles name resolution inside the cluster and configures individual containers to contact the DNS service and resolve domain names.

Pod to pod: How pods communicate with each other

Each pod has its own unique IP in a flat address space inside the Kubernetes cluster. Direct pod-to-pod communication is possible without any type of proxy or address translation.

Pod to service: How pods communicate with services

Kubernetes services allow users to group pods under a common access policy. For example, a group of pods can be load balanced. In that case, the load balancing services are assigned a virtual IP. Outside pods can communicate by using this virtual IP.

External to service: Incoming traffic from the outside world

Nodes inside a Kubernetes cluster are firewalled from the internet by default, and service IPs are only reachable within the cluster network. There are two possible approaches: route requests using the external IP, or offer an ingress API.

Learn more

Download Antrea from GitHub

Try VMware NSX-T free for 60 days

To allow for incoming traffic, a service can be mapped to one or more external IPs. Incoming requests at the external IP are routed to the node. The node knows which services are mapped to that external IP and which pods are part of the service. The request is then routed to the appropriate pod.

To support more complex policies, Kubernetes provides the ingress API, which offers externally reachable URLs, traffic load balancing, SSL termination and name-based virtual hosting. An ingress is a collection of rules that allow an inbound connection to the service. An ingress controller, typically a load balancer, is responsible for fulfilling the ingress.

An ingress controller will allocate an external IP to satisfy the rules defined by the ingress and forward all requests arriving to that external IP to the service mapped in the ingress specification.

Item No: 1158495aq-ds-vmw-cntnr-ntwkng-antrea-uslet 12/21