VMware Container Networking with Antrea

AT A GLANCE

VMware Container Networking™ with Antrea™ offers full enterprise support for VMware Kubernetes Service (VKS) in VMware Cloud Foundation. It integrates with managed Kubernetes services to enhance Kubernetes network policies and is the default Container Network Plugin for VMware Kubernetes Service in VCF. It also supports Windows and Linux workloads on Kubernetes across multiple clouds.

KEY BENEFITS

- Delivers pod networking for VKS and integrates deeply with VPC underlay networking and external load balancers
- 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
- Improves network visibility and diagnostics with tools such as Traceflow, PacketCapture, IPFIX, and Support Bundle
- Extends upstream Kubernetes NetworkPolicy with features such as priorities, deny semantics, and richer selectors using Antrea-native policies

Challenges in 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 networking solutions are not only dependent on compatibility with a specific Kubernetes version, but also on alignment with the broader platform stack on which Kubernetes runs. In VKS with Antrea CNI, compatibility across multiple layers, including Kubernetes, CNI, hypervisor, networking platform (NSX), and cluster lifecycle management components is always maintained.

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.

Addressing the Challenges with VMware Container Networking with Antrea

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 VCF & Tanzu by Broadcom, there already exists an installed user base for VMware Container Networking. Customers with valid licenses for VMware Cloud Foundation 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 open-source and the release cadence of Kubernetes. When VKS clusters are deployed with Antrea on NSX-backed infrastructure, Antrea handles pod / service networking semantics in the cluster while NSX provides VPC networking, subnet attachment, egress / ingress connectivity, NAT, external IP realization, and load balancing integration.

DETAILED BENEFITS

- 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 & PacketCapture to aid in troubleshooting and root cause analysis.
- Improve container security by encrypting traffic between pods despite running on untrusted fabrics using common node-to-node encryption frameworks such as Wireguard.
- 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. It 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.

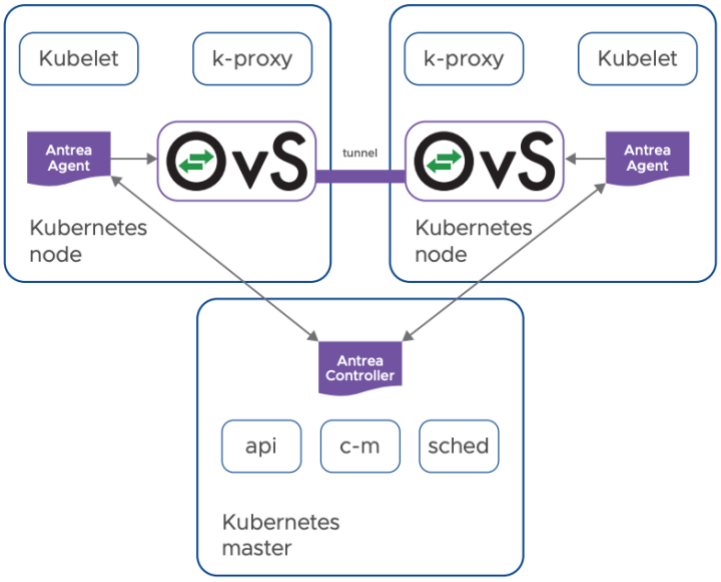

Figure 1: Antrea and Open vSwitch

What is Open vSwitch?

Open vSwitch is a high-performance, programmable virtual switch that supports Windows and Linux workloads. In Antrea deployments, OVS is the per-node software switching data plane responsible for pod attachment, forwarding, encapsulation / decapsulation where used, service proxy handling, and enforcement integration. Antrea also programs OVS flows to implement service handling, policy enforcement, and traffic steering.

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. Antrea's proprietary policy model is also broader than upstream NetworkPolicy alone because AntreaNetworkPolicy and ClusterNetworkPolicy add richer matching and ordering semantics that are useful for enterprise tenants and platform operators.

Use Cases for VMware Container Networking for Antrea

Complete container networking and security solution for VMware Kubernetes Service (VKS)

Designed into the VMware Cloud Foundation platform, Antrea is the default container networking solution for VKS deployments in VMware Cloud Foundation, supporting pod-to-pod networking, and granular distributed firewall with integration with vDefend Distributed Firewall.

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 VKS pods. In VKS, in-cluster policies can be extended by combining with NSX infrastructure security, inventory, and VPCs.

Key Features	
Connectivity	• Choice of routed, encapsulated (overlay), hybrid and cloud provider routing mode of operation • Choice of encapsulation (overlay) modes (Geneve, VXLAN, STT, GRE) • High-performance, low-latency Kubernetes service implementation in Open vSwitch (kube-proxy/iptables replacement) • Advanced load balancer integration through pod-specific NodePort (node port local) • Add multiple networks and support for multicast traffic
Platform Support	• VMware Kubernetes Service (VKS) • Windows containers network data plane • ARM Kubernetes node support

Security	• Kubernetes NetworkPolicy support for pod-level security within VKS clusters running on NSX-backed infrastructure • Cluster-wide network policy support (Antrea Cluster Network Policy) • Cluster-wide network policy tiering and nesting • Role-based access control for policy tiers • FQDN/DNS-based egress policy with wildcard matching • Deny network policy • Network policy statistics export
Advanced security and manageability	• Integration with VMware vDefend Firewall with consistent workflow for creating groups and Distributed Firewall rules/policies for Kubernetes pods and VMs. • Antrea network policy management through NSX management plane • Central Kubernetes object inventory through NSX management plane • Central multi-cluster connectivity troubleshooting via Traceflow through NSX management plane
Manageability	• Traceflow pod-to-pod connectivity visualization • ID-aware network metrics and flow export • Prometheus load and flow metrics • Container IPFIX • Monitoring CRDs for control plane status and health • Octant UI plug-in for Traceflow, cluster inventory, and monitoring control plane health • Antrea network policy audit logs • Advanced troubleshooting through CLI and support bundle
Hardening, Services, and Support	• FIPS-compliant product release • 24x7 enterprise support

For more information or to purchase VMware products

Call 877-4-VMWARE (outside North America, +1-650-427-5000), visit vmware.com/products, or search online for an authorized reseller.

A 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.

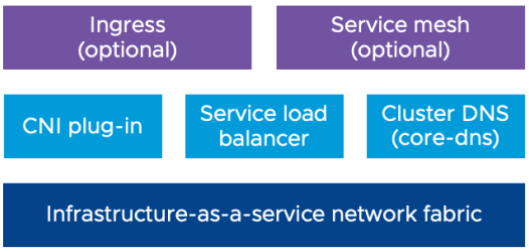

Figure 2: Kubernetes Networking Layers

CNI Plugin

A CNI plugin builds the pod network and provide connectivity. The Kubernetes service load balancer helps define pods and access policy. Services and service types are used to precisely control how applications receive traffic. The load balancer service type creates an external load balancer in the public cloud 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 has its own DNS service for domain names inside the cluster, allowing 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 a 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

Kubernetes nodes are firewalled from the internet by default and service IPs are only reachable within the cluster. To allow 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 allocates an external IP to satisfy ingress rules and forward requests to the service mapped in the ingress specification.

vmware by Broadcom

Item No: vmw-bc-ds-a4-word-2025 Mar-26