STRATA™
BY PALO ALTO NETWORKS

CN-Series Container Firewall

Prevent threats in Kubernetes environments with the industry's leading Next-Generation Firewall

The Palo Alto Networks CN-Series is the industry's first next-generation firewall delivered in a container form factor and natively integrated with Kubernetes®. The container firewall prevents network-based threats from spreading across Kubernetes namespace boundaries.

CN-Series container firewalls deliver:

- Threat prevention and advanced network security services to protect Kubernetes namespace boundaries
- Automated deployment and configuration via Kubernetes for frictionless network security
- Visibility into native Kubernetes metadata (e.g., namespace) for context-based security policies
- Centralized management with Panorama

Overview

Conventional next-generation firewalls can only be deployed at the edge of a Kubernetes environment and therefore cannot determine the specific pod where traffic originates. To overcome this challenge, CN-Series Container Next-Generation Firewalls are deployed on each node of a Kubernetes cluster, giving them precise visibility into container traffic. The CN-Series container firewall delivers Layer 7 visibility and control while enabling the enforcement of advanced security services. This protection can be enforced on allowed traffic traversing namespace boundaries—whether outbound, inbound, or east-west—between pods, and even between containerized applications and legacy workloads, such as virtual machines (VMs) and bare metal servers.

CN-Series firewalls are easy to deploy using Kubernetes orchestration to simplify integration of network security into continuous integration/continuous development (CI/CD) processes. Ongoing management of CN-Series firewalls is centralized in the same management console as all Palo Alto Networks firewalls, which gives network security teams a single pane of glass through which to manage the overall network security posture of their organizations.

How the CN-Series Works

CN-Series firewalls deploy as two sets of pods: one for the management plane (CN-MGMT) and another for the firewall dataplane (CN-NGFW). The firewall dataplane runs as a daemon set, allowing a single command from within Kubernetes to deploy firewalls on all nodes in a Kubernetes cluster at once. The management plane simply runs as a Kubernetes service.

Management of CN-Series firewalls is done through the Panorama™ network security management console. A Kubernetes plugin within Panorama provides contextual information about containers in an environment, and this seamlessly enables context-based network security policies. For example, Kubernetes namespaces can be used to define a traffic source in a firewall policy.

Customers can deploy CN-Series firewalls in Kubernetes environments hosted on-premises or in public clouds. CN-Series firewalls can also be deployed into cloud-managed Kubernetes offerings, including Google Kubernetes Engine (GKE), Azure Kubernetes Service (AKS), and Amazon Elastic Kubernetes Service (EKS).

Deployment via Kubernetes package managers such as Helm is also available and community-supported.

CN-Series Use Cases

Prevent Data Exfiltration from Kubernetes Environments

CN-Series firewalls offer a multitude of security capabilities to prevent exfiltration of sensitive data from Kubernetes environments. Traffic content inspection—including inspection of encrypted SSL traffic—ensures that packets containing malicious payloads are identified and remediated. URL Filtering bars outbound connections to potentially nefarious websites, including malicious code repositories.

Prevent Lateral Spread of Threats Across Kubernetes Namespace Boundaries

Trust boundaries between applications are logical locations to enforce segmentation policies that prevent the lateral movement of threats. In many Kubernetes environments, the Kubernetes namespace is the trust boundary. CN-Series firewalls can enforce Threat Prevention policies between Kubernetes namespaces, as well as between a Kubernetes namespace and other workload types (e.g., VMs and bare metal servers), to deter threats from moving between your cloud native applications and your legacy infrastructure.

Prevent Known and Unknown Inbound Threats

Just like many applications, attacks can use any port, which limits the effectiveness of port-based network security controls. With their application-centric security policies, CN-Series firewalls augment basic port-based access controls and inspect network traffic to ensure only allowed applications are permitted across open ports.

By enabling our integrated cloud-delivered subscriptions services, you can enhance your security capabilities without compromising productivity. Turning on our Threat Prevention and WildFire® malware prevention services on the CN-Series firewall protects your Kubernetes environment against any file-based threats, including exploits, malware, spyware, and previously unknown threats, attempting to sneak through open ports. In addition, deploying our URL Filtering and DNS Security services protects your environment from web-based threats, including phishing, command and control, and data theft.

CN-Series Key Capabilities

Whatever the security needs of your container environment, the CN-Series is built to deliver.

Inline Network Security Visibility and Control

- **Threat prevention and sandboxing:** Threat Prevention and WildFire services can be enabled on CN-Series firewalls to block exploits, prevent malware, and stop both known and unknown advanced threats.
- **Exfiltration prevention and URL filtering:** The CN-Series enables content inspection and SSL Decryption, preventing sensitive information from leaving your network. URL Filtering uses machine learning to categorize URLs and block access to malicious sites that deliver malware or steal credentials. Automation ensures protections are always up to date.
- **Flexible tag-based policy model:** CN-Series firewall policies can be defined by application, user, content, native Kubernetes labels, and other metadata to deliver flexible policies aligned with business needs.

Automated Deployment and Configuration with Kubernetes

- **Kubernetes orchestrated deployment:** CN-Series firewalls run as a daemon set, allowing a single command from within Kubernetes to deploy firewalls on all nodes in a cluster at once.
- **DevOps-friendly configuration:** All configuration of CN-Series firewalls is specified in a YAML file and can be easily integrated into infrastructure deployment files for fast, repeatable deployments. Configuration templates can be found in our official CN-Series GitHub repository.
- **Community-supported Kubernetes Helm chart:** For development teams using Helm to manage their Kubernetes applications, a CN-Series Helm chart has been created to simplify firewall deployment and management.

Flexible and Consistent CNI Integration

- **Simple insertion:** The CN-Series supports multiple container network interface (CNI) plugins for use in different types of Kubernetes deployments.

Kubernetes On-Premises and Cloud Support

- **Public cloud:** CN-Series firewalls can be deployed in hosted container environments such as GKE, AKS, Amazon EKS, and Red Hat OpenShift®. For detailed platform support information, see table 1.
- **On-premises:** CN-Series firewalls can also be deployed into Kubernetes environments hosted on-premises.

Centralized Management in Panorama

- **Consistent management:** Manage the CN-Series from Panorama—the same management console you use for your hardware and virtual form factor Palo Alto Networks firewalls.
- **Plugin architecture:** Panorama plugins for GKE, AKS, Amazon EKS, and OpenShift allow you to manage network security for each environment from Panorama.
- **Centralized logging:** Panorama centralizes logging to simplify audit and compliance.

Table 1: CN-Series Support Matrix

Product	Version(s)*
Containerized PAN-OS	10.0
Panorama Kubernetes Plugin	1.0.0
Container Runtime	Docker, CRI-O
Native Kubernetes	1.13, 1.14
Cloud Provider Managed Kubernetes	OpenShift 4.2 Amazon EKS (1.13, 1.14) AKS (1.13, 1.14) GKE (1.13, 1.14)

Table 1: CN-Series Support Matrix (continued)

Customer-Managed Kubernetes†	
Kubernetes Host VM OS	Ubuntu 16.04, 18.04 RHEL/CentOS 7.3+ CoreOS 21XX, 22XX
CNI	CNI Spec 0.3.0 and higher: • Calico • Flannel • Weave

* Recommended versions for Kubernetes, Calico, etc.

† In customer-managed deployments, Kubernetes can be deployed using any orchestrator (e.g., Rancher, Kubespray) and deployed in a public or private cloud as long as Kubernetes, CNI, and Host OS versions are those in table 1.

Key Performance Metrics

Testing was conducted on GKE, with traffic directed between nodes and between pods on the same node in the same cluster.

Table 2: CN-Series Key Performance Metrics

Feature/Attribute	CN-NGFW (1 Core)
Firewall Throughput (App-ID Enabled)	500 Mbps
Threat Prevention Throughput	250 Mbps
Max Sessions	20,000

CPU and Memory Requirements

Table 3 shows system requirements for the cluster in which the CN-Series is deployed. While the CPU, memory, and disk storage will depend on your needs, these are general guidelines.

Table 3: CN-Series CPU and Memory Requirements

Resource	CN-MGMT (StatefulSet Pod for Fault Tolerance)	CN-NGFW (DaemonSet Pod)
Memory (min.)	2 GB	2 GB
Memory (max.)	4 GB	2.5 GB
CPU (min.)	2	1
Disk	50 GB	N/A

Scalability of Components

Table 4 lists scalability numbers of the different CN-Series components, and table 5 lists those of the Panorama Kubernetes plugin.

Table 4: Scalability of CN-Series Components

Capacity/Attribute	CN-Series Scale
Max. CN-MGMT Pairs per K8s Cluster	4
Max. CN-NGFW Pods per CN-MGMT Pair	30
Kubernetes Pods Secured by CN-NGFW (per Kubernetes Node)	30
Max. TCP/IP Sessions on CN-NGFW	20,000
Max. DAG IP Addresses*	2,500
Tags per IP Address	32

Table 5: Scalability of Panorama Kubernetes Plugin Components

Capacity/Attribute	Panorama Kubernetes Plugin Scale
Max. Clusters per Panorama Kubernetes Plugin	16 (across Native Kubernetes, GKE, EKS, AKS)
Max. Pods per Cluster in Panorama Kubernetes Plugin	900
Max. Services per Kubernetes Cluster (Internal + External)	40
Max IPs (Pods + Services) Across Clusters per Device Group in Panorama Kubernetes Plugin	1,560 (MP supports 2,500)

paloalto NETWORKS

3000 Tannery Way
Santa Clara, CA 95054

Main: +1.408.753.4000
Sales: +1.866.320.4788
Support: +1.866.898.9087

www.paloaltonetworks.com