



CN-Series container firewalls deliver:

- **Inline network security** to improve visibility and protect Kubernetes namespace boundaries with security services such as threat protection and URL filtering.
- **Flexible configuration and deployment options** to allow administrators to automate security deployment and leverage the autoscaling capabilities of Kubernetes.
- **Cloud and on-premises support** to give network architects the tools to design public, private, and hybrid cloud architectures.
- **Centralized security management** via Panorama to unify control of hardware and virtual Palo Alto Network firewalls as well as native security provided by Amazon Web Services, Google Cloud, Microsoft Azure Kubernetes Services and Openshift.

CN-Series Container Firewall

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

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

Overview

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

Conventional NGFWs 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 firewalls move security inside the Kubernetes environment, giving them precise visibility into and control over container traffic.

The CN-Series delivers Layer 7 visibility and control while enabling the enforcement of advanced security services, such as Intrusion Prevention. This protection can be enforced on allowed traffic traversing namespace boundaries within or between Kubernetes clusters, including 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, allowing operators to deploy network security using the same processes and technology they use to manage the rest of their environments. Ongoing management of CN-Series firewalls is centralized in the Panorama™ network security management solution—the same management console as all Palo Alto Networks firewalls—giving network security teams a single pane of glass 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 management pod always runs as a Kubernetes service. The dataplane pods can be deployed in two modes: distributed or clustered. In distributed mode, the firewall dataplane runs as a daemon set on each node. Administrators can deploy firewalls on all cluster nodes with a single command, placing security controls as close to workloads as possible. In clustered deployment mode, the firewall dataplane runs as a Kubernetes service in a dedicated security node. When deployed in clustered mode, CN-Series takes advantage of the native autoscaling capabilities of Kubernetes to ensure threat protection in even the most dynamic Kubernetes environments. Clustered deployments are best suited for large Kubernetes environments where a distributed deployment would be resource intensive and cost prohibitive.

CN-Series firewalls are managed through the Panorama 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.

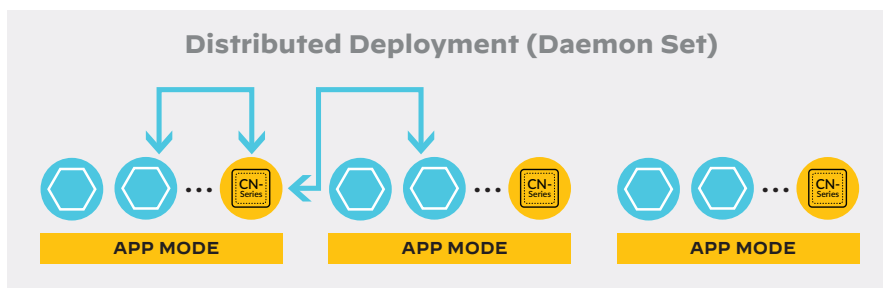


Figure 1: Distributed deployment mode. The firewall dataplane runs as a daemon set on each node.

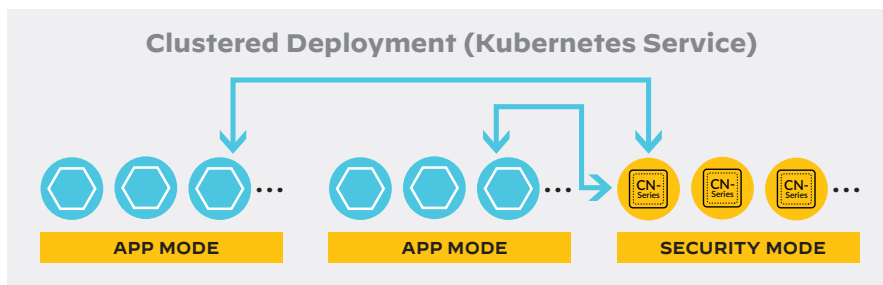


Figure 2: Cluster deployment mode. The firewall dataplane runs as a Kubernetes service in a dedicated security node.

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 TLS-/SSL-encrypted 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 Both Known and Unknown Inbound Threats

Like many applications, attacks can use any port, which limits the effectiveness of port-based network security controls. With 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.

Application Visibility and Inline Threat Prevention

- **Application visibility and control:** Get immediate visibility into application traffic within your Kubernetes environment. Define application-based policies to control application traffic and enforce Zero Trust best practices.
- **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.

Automated Scalability and Configuration

- **Autoscale with Kubernetes:** CN-Series firewalls can leverage the autoscaling capabilities of Kubernetes to ensure protection in even the most dynamic environments.
- **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.
- **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 Deployment and Consistent CNI Integration

- **Flexible deployment options:** Customers can choose to deploy CN-Series firewalls in distributed or clustered modes, depending on their use case, budget and environmental configuration.
- **Simple insertion:** The CN-Series supports multiple container network interface (CNI) plugins for use in different types of Kubernetes deployments.

On-Premises and Cloud Support for Kubernetes

- **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 Security Management

- **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 auditing and compliance.

Size and Scale Security Based on Immediate Needs—In Minutes

Match software firewalls and security services with the speed and flexibility needed for rapidly changing requirements. Maximize your ROI on security investments with the industry’s most flexible way to adopt software NGFWs and security services. Discover unmatched flexibility with easy scaling and sizing of VM-Series virtual and CN-Series container NGFWs, cloud-delivered Security Services, and VM Panorama for management and log collection.

Three simple steps let you choose and deploy the right firewalls and security services you need at any given time:

1. Procure Software NGFW Credits.
2. Allocate or reallocate credits across different deployments to activate your choice of security products and your choice of security services in just minutes.
3. Manage and monitor credits via the Palo Alto Networks customer support portal.

As needs change, Software NGFW Credits can be reallocated to new and other firewall-as-a-platform solutions without having to go through additional procurement cycles.

Table 1: CN-Series Support Matrix	
Product	Version(s)
Containerized PAN-OS	10.0 and above
Panorama Kubernetes Plugin	1.0.0
Container Runtime	Docker, CRI-O
Native Kubernetes	1.14 – 1.18
Cloud Provider-Managed Kubernetes*	OpenShift 4.2, 4.4, 4.5 AWS EKS (1.14 – 1.17) Azure AKS (1.14 – 1.18) GCP GKE (1.14 – 1.17)
Customer-Managed Kubernetes†	
Kubernetes Host VM OS	Ubuntu 16.04, 18.04, RHEL/ CentOS 7.3+
CNI	CNI Spec 0.3.0 and higher, which support CNI chaining (e.g., 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 so far as Kubernetes, CNI and host OS versions are from 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
CPU (max.)	2	4
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)

Find more information about CN-Series container firewalls [here](#).