White Paper

Cloud-Native Software

# Next-Generation Network Protection for Container Workloads

## Deployed on Red Hat® OpenShift®, the Palo Alto Networks CN-Series Container Firewall protects containerized workloads located across multi-cloud environments based on Intel® architecture.

Table of Contents



Network transformation continues to accelerate, driven by forces that include 5G telecom, cloud-native computing and the expanding importance of edge processing. For more than a decade, this evolution toward cloud-native computing has moved from proprietary, hardware-centric solutions to open, software-defined ones, as shown in Figure 1. Hardware has moved from single-purpose, single-vendor custom systems located exclusively in data centers to general-purpose, open-standards servers located anywhere.

Software models have shifted from monolithic applications hosted in data centers to chains of microservices that are encapsulated with their dependencies in containers. High-cost, inflexible software has largely given way to thriving open-source community innovation. Workloads migrate freely among on-premises resources, public clouds, and edge servers across highly distributed cloud-native infrastructures.

By taking advantage of elastic capacity from public cloud instances, organizations avoid legacy requirements to maintain idle servers for workload peaks and redundant ones for failover, development and testing. Infrastructure modernization is also instrumental for IT adoption of DevOps, DevSecOps and continuous integration and delivery (CI/CD). Together, these changes afford network operators unprecedented flexibility, agility and cost efficiency as they bring new services into production

In 5G networks, network functions that have traditionally been provided by dedicated, single-purpose hardware have been replaced by software-based, containerized virtualized services. This shift enables the use of general-purpose commercial off-the-shelf servers in place of more expensive specialized equipment. Processing data close to where it is produced and/or consumed supports low-latency services required for communications and safety applications. It also enables the use of massive data sets such as those generated in IoT while avoiding the cost and security exposure of backhauling massive data sets back to the data center or cloud for centralized processing.

This paper presents a solution stack based on the Palo Alto CN-Series Container Firewall protecting workloads operating on Red Hat OpenShift Container Platform, hosted on general-purpose servers based on Intel architecture. The CN-Series next-generation firewall (NGFW) is CNF-certified on Red Hat OpenShift, providing assurances that it has been tested to verify deployment best practices and that the unified combination of products is jointly supported by both companies.

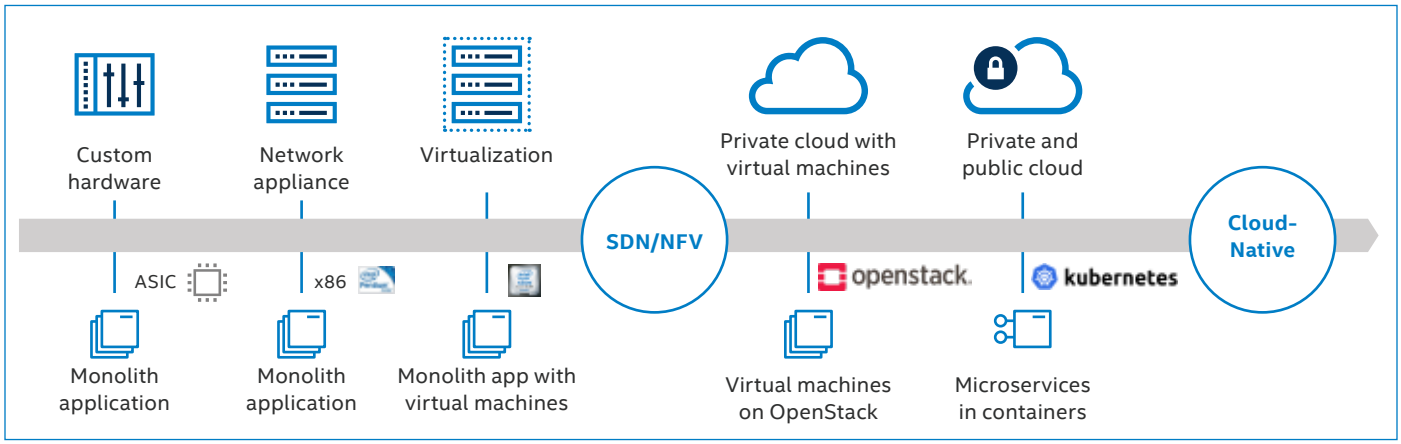


**Figure 1.** Evolution of cloud-native computing.

## 1. Shifting Security Requirements in a Cloud-Native World

Whether applications are running on bare-metal servers, virtual machines or containers, they use the same network stack and protocols. Containerized apps therefore face the same threats that have traditionally plagued legacy apps running on bare metal and virtual machines. Patching security vulnerabilities is often manual and time-consuming, and it can take weeks or months to patch hundreds of vulnerable applications in a deployment. While agent-based deploy-time (shift-left) security products help identify and patch known vulnerabilities at scale, applications remain helpless against unknown and unpatched vulnerabilities. For example, the recent high-profile unknown critical vulnerability Log4j affected some 3 billion devices.

Zero-trust network access (ZTNA) has become the industry standard for protecting data and workloads regardless of where they are hosted. All entities are regarded as untrusted, and every transaction on the network must be individually authenticated and authorized. CN-Series allows customers to deploy the NGFW inside the Kubernetes cluster along with applications, providing the same level of network security for both containerized applications and traditional applications deployed on VMs and bare-metal servers. This approach enables ZTNA security models regardless of the underlying hardware, including third-party clouds and hosted resources.

Evolving firewall requirements are illustrated in Figure 2, where ordering and payments applications run on different pods represented by hexagons, which are hosted on bare-metal or cloud nodes. Traditional approaches to securing the cluster would place a physical or virtual firewall outside the cluster to govern data ingress and egress, as shown in the figure. Under this model, traffic exiting from the cluster would be associated with the node's IP address, obscuring the identity of its pod-level/application-level source. That lack of granularity prevents identifying the application or namespace associated with a given packet, forcing security to enforce the policy at the node level and making Layer 7 protections impossible between multiple applications running on the same Kubernetes cluster. Apart from east-west Layer 7 traffic protection, customers can't enforce application-level policy to prevent container-specific inbound attacks as well as prevent data exfiltration and unwanted outbound connections for container applications.

This example illustrates the necessity of bringing NGFW protection inside the cluster, where it can differentiate between traffic coming from the separate applications and provide the true source IP address of outgoing data. Application-level visibility enables enforcement of true Layer-7 security at the application, namespace and pod level, within the container cluster. A software-based container firewall is central to this model.

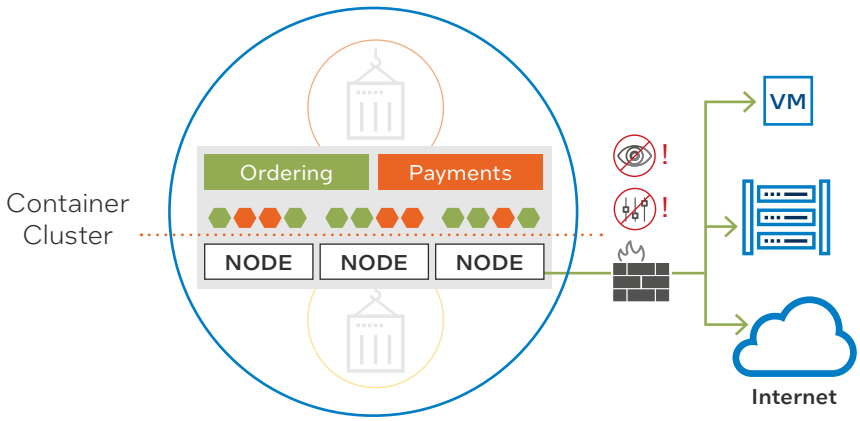


**Figure 2.** Cluster-based firewall protection.

## 2. Palo Alto Networks Cloud Security

The Palo Alto Networks CN-Series Container Firewall is purpose-built to protect Kubernetes environments from network-based attacks. By moving security inside the cluster, CN-series enables Layer-7 protection and granular visibility over traffic sources and destinations at the node, pod and namespace levels. It also provides inline threat protection for containerized applications and scales dynamically based on automated processes.

CN-Series firewall deployments consist of a management-plane pod and a set of dataplane pods. The management pod runs as a Kubernetes service, while the dataplane pods can be deployed in either distributed or clustered mode:

- **Distributed mode** instantiates the dataplane as a daemon set on every individual node, which places it as close as possible to workloads. A single command can deploy CN-Series firewalls to every node in a cluster.
- **Clustered mode** runs the firewall dataplane as a Kubernetes service on a dedicated security node or nodes in each cluster. Autoscaling enables robust firewall protection even in highly dynamic environments, making this approach more scalable and cost-effective for large environments than distributed mode.

Comprehensive security posture typically requires both hardware and software firewalls. To optimize deployment flexibility and streamline management and maintenance, the CN-Series is built to be centrally provisioned, monitored and managed from a centralized management plane called Panorama, along with other Palo Alto Networks firewalls including Hardware and VM-series Firewall. Panorama enables security teams to easily create consolidated policies and protect workloads across physical, virtual and containerized environments from a single pane of glass, helping drive efficiency at any scale.

## 3. Red Hat OpenShift Container Platform

To enable the hybrid cloud, Red Hat OpenShift increases the enterprise readiness of Kubernetes, as shown in Figure 3. The underlying community-powered Kubernetes runs containers in a unified way from the network core, to the cloud, to the edge. Red Hat OpenShift expands on that foundation to provide a full-lifecycle platform to build, deploy, and run applications, including traditional n-tier apps, cloud-native microservices and ISV packaged apps. The platform provides application security capabilities, end-to-end automation and management, and integrated developer toolsets and processes.

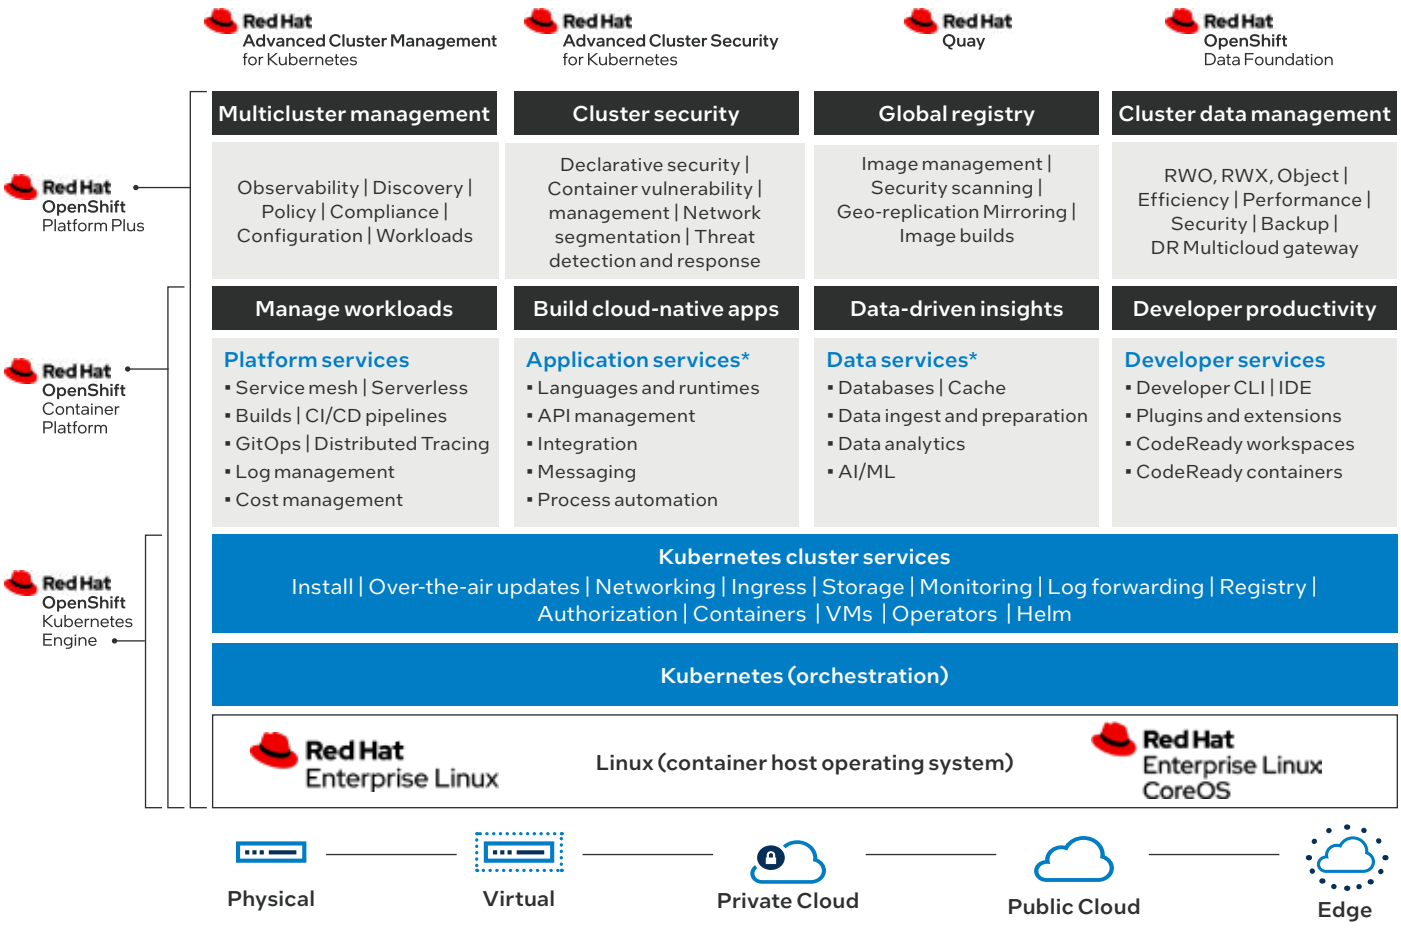


* Red Hat OpenShift® includes supported runtimes for popular languages/frameworks/databases. Additional capabilities listed are from the Red Hat Application Services and Red Hat Data Services portfolios.

** Disaster recovery, volume and multicloud encryption, key management service, and support for multiple clusters and off-cluster workloads requires OpenShift Data Foundation Advanced

**Figure 3.** Providing an enterprise-ready container platform with Red Hat OpenShift.

Red Hat OpenShift enables development organizations to deploy applications anywhere, with a consistent environment, so that architecture can be scaled using existing teams and practices. It offers flexibility of choice across a choice of topology options to meet varying organizational and site requirements. And it participates in a broad Red Hat and partner ecosystem that provides building blocks that can be combined to accommodate individual use cases.

Human operational knowledge such as processes for packaging, deploying and managing containerized applications can be instantiated in code as Red Hat OpenShift Operators. They enable automation of these processes without human intervention, which fosters repeatability, increases efficiency and reduces the incidence of error. The solution implements the following Operators, which cluster administrators can install using the OpenShift command-line interface or the web console:

- **Cluster Network Operator** deploys and manages cluster network components, including a choice of container network interface (CNI) plug-ins compatible with cluster and functional requirements.
- **Performance Addon Operator** enables advanced performance tuning on sets of nodes, including control of features such as CPU allocation and reservation, memory hugepages and IRQ management.
- **SR-IOV Network Operator** creates and manages the SR-IOV stack for OpenShift, including discovery of SR-IOV network devices, configuring the OpenShift device plug-in, and creating needed custom objects.

## 4. Open-Standards Servers Based on Intel Architecture

Implementing the CN-Series firewall on Red Hat OpenShift and Intel architecture enables customers to achieve outstanding performance and efficiency at any scale. Intel® Xeon® D processors are optimized for high-density compute at low thermal design power (TDP). The platform enables enterprises and service providers to achieve high throughput and low latency within the typical space and power constraints at the distributed edge. Intel Xeon Scalable processors provide data center-level flexible performance from a wide range of core counts and feature sets, as well as the ability to scale from two to eight sockets per system.

High software and API compatibility across all past and present Intel Xeon processors and the unmatched Intel ecosystem of hardware and software providers offer flexibility and agility in deployment. The spectrum of ecosystem enablement includes partner programs such as Intel Network Builders as well as contributing to and maintaining open source projects. Co-engineering work with Red Hat and Palo Alto Networks helps optimize the benefit of hardware features and capabilities to the joint implementation of components from the three companies.

Both Intel Xeon D and Xeon Scalable processors provide microarchitecture and feature improvements that dramatically increase per-core performance, security, and energy efficiency compared to their predecessors. Shared features across both platforms provide a unified set of security and performance capabilities for deployments at varying scale:

- **Intel AES New Instructions (Intel AES-NI)** enable developers to access hardware acceleration for key processing-intensive portions of the AES algorithm, reducing the overhead associated with the pervasive encryption that is typical in 5G and other ZTNA environments.
- **Intel Total Memory Encryption (Intel TME) and Intel Multi-Key Total Memory Encryption (Intel MKTME)** cryptographically protect memory against attacks using the NIST AES XTS standard without requiring modification to applications. Keys are generated using a hardened random-number generator that is implemented in silicon, beyond the reach of software-based attacks.
- **Intel Deep Learning Boost (Intel DL Boost)** eliminates unneeded precision within the massive calculations required to process machine learning workloads. By simplifying this computation, Intel DL Boost reduces resource requirements to accelerate throughput.

## 5. Use Cases for CN-Series Firewalls on Red Hat OpenShift

With cloud-native security becoming a core requirement for most or all enterprises, an expanding set of use cases has emerged. Because of their rapid growth, 5G networks — which by their nature are highly distributed — are an implementation area of specific note for container firewalls. In addition to carrier networks that are gaining momentum, private 5G networks may be expected to become more prevalent in the next few years.

Cloud-native infrastructure is being adopted more broadly for general-purpose enterprise networks as well. The use cases described in this section focus on the ability of the Palo Alto Networks CN-Series Container Firewall to provide Layer-7 protection in OpenShift environments, recognizing the specific applications associated with a given traffic flow. These capabilities enable networks to accommodate different trust levels, security policies and isolation requirements for individual applications.

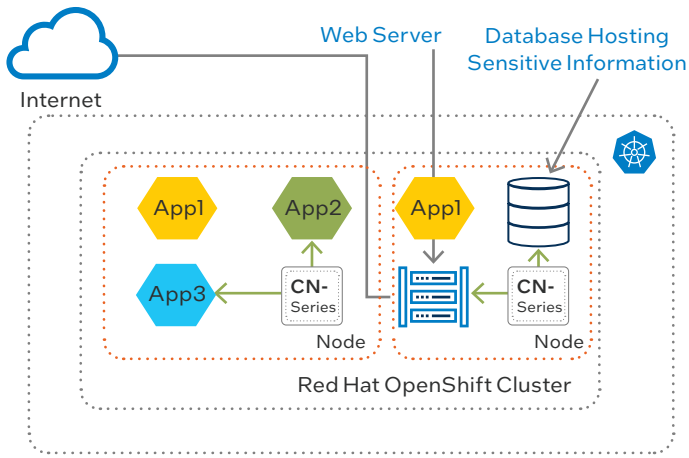


**Figure 4.** Stopping lateral movement of threats.

### 5.1 East-West Layer-7 Traffic Inspection Use Case

The customer use case illustrated in Figure 4 deploys an internet-facing web application and a database application on the same Red Hat OpenShift cluster. The database stores sensitive information such as customer credit card information and patient records. A CN-Series firewall enforces Layer-7 traffic inspection between the web server and the database to help prevent threats from moving laterally between these two applications, helping maintain regulatory compliance standards such as PCI and HIPAA.

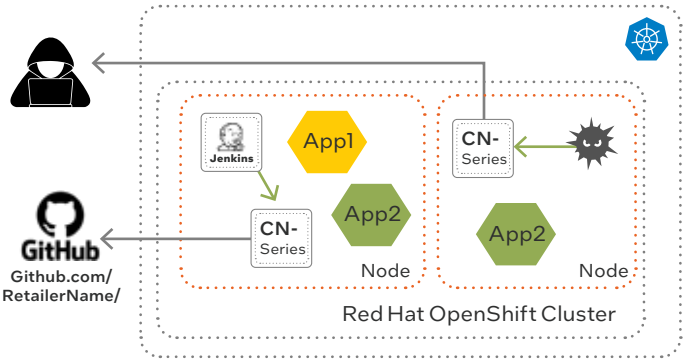


**Figure 5.** Filtering traffic to prevent malicious data egress.

### 5.2 Outbound Traffic Protection Use Case

Figure 5 depicts egress filtering — one of the most prevalent use cases for deployment of the CN-Series firewall — at a large retailer. The left side of the figure shows Jenkins deployed inside the cluster, with Layer-7 policy applied using the firewall to allow downloads from the Jenkins application only from a specific GitHub repository. All other applications running on the same Kubernetes cluster are free to download data from any URL.

On the right side of the figure, the CN-Series firewall is being used to disallow connections to command-and-control servers that would attempt to hijack resources for illicit crypto mining. Here, the CN-Series is being used in conjunction with Palo Alto Networks' cloud-delivered URL filtering and DNS security services, which help provide real-time protection from both known and unknown threats.

### 5.3 Inbound Threat-Prevention Use Case

The topology shown in Figure 6 shows an eCommerce company's implementation of the CN-Series firewall. To maintain regulatory compliance for payment data, the

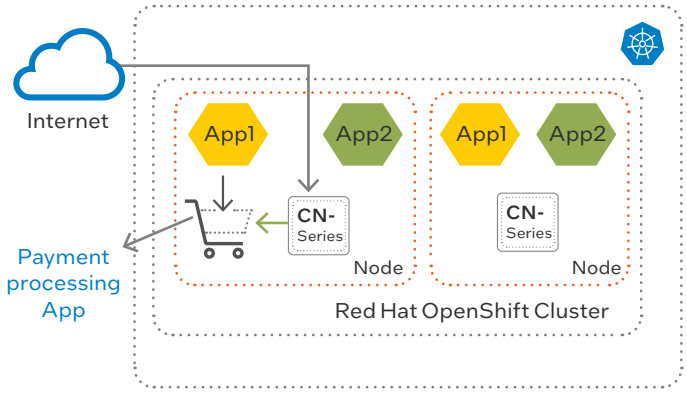


**Figure 6.** Protecting against malware with deep packet inspection.

company needed to have different security policies for different applications running on the same Kubernetes cluster. With CN-series, the customer is able to enforce the strictest policy for payment processing application, while enforcing slightly less-strict policies for other applications running on the same Kubernetes cluster.

In conjunction with Palo Alto Networks cloud-delivered Threat Prevention and Wildfire security services, the CN-Series firewall enforces different policies for different applications, performing deep packet inspection and threat prevention to stop malware and other incoming threats.

## Conclusion

The shift to cloud-native infrastructure offers dramatic improvements in agility and efficiency that include organizational opportunities such as supporting adoption of DevOps, DevSecOps and other modern practices. At the same time, this transformation introduces a significant set of new security challenges, including a lack of visibility inside Kubernetes clusters.

The CN-Series firewall helps address both those sets of challenges by deploying in software inside the cluster, with full CNF certification for Red Hat OpenShift. That certification gives customers confidence as they take advantage of OpenShift automation and robust develop-test-build-run capabilities as well as the security capabilities and performance advantages of underlying Intel architecture. The combination of these building blocks from Palo Alto Networks, Red Hat and Intel help provide future-ready protection for service providers and enterprises as they advance their transformation journeys.

## More Information

Intel Xeon Scalable processors: intel.com/xeonscalable

Palo Alto Networks CN-Series Container Firewalls for Kubernetes: paloaltonetworks.com/network-security/cn-series

Red Hat OpenShift: redhat.com/en/technologies/cloud-computing/openshift


0822/RKM/MESH/349365-001US