

# Definitive guide to Red Hat OpenShift security

## OpenShift Container Platform

Red Hat® OpenShift® Container Platform is a Kubernetes platform for operationalizing container workloads remotely or as a hosted service. OpenShift Container Platform provides more consistent security, built-in monitoring, centralized policy management, and compatibility with Kubernetes workloads. The rapid adoption of open source projects can introduce vulnerabilities in standard Kubernetes environments. OpenShift Container Platform supports these projects, allowing users to gain open source advantages with a managed product's stability and security. Red Hat OpenShift offerings include five managed and two hosted options.

This guide will focus on the Red Hat OpenShift Container Platform 4.5, which is designed to be self-managed within your infrastructure environment with various deployment options.

## OpenShift architecture

OpenShift Container Platform is based on Kubernetes, and while Kubernetes provides container orchestration capabilities, pod resiliency, services definitions, and deployment constructs, there are many other components required to make it work. For example, Kubernetes does not provide default container network interface (CNI) or default monitoring implementations. It is up to the cluster administrator to bring additional tools to operate and manage the Kubernetes cluster and any applications running. For security teams, this presents new challenges—for example, these teams need to create new policies and vet images, configurations, and account access for any new applications that will be deployed into the cluster.

These additional, necessary operational capabilities are provided with OpenShift Container Platform and are pluggable so that administrators can customize components and services to meet their infrastructure needs.

OpenShift Container Platform's architecture requires three different types of nodes within each cluster to ensure highly available deployments.

### Control plane nodes

These nodes run the core Kubernetes control plane functions and provide additional services such as a self-service web console and developer- and operations-focused dashboards.

In most cloud environments, the control plane nodes are hidden from users and managed by providers for high availability, regular upgrades, and added security updates. With OpenShift Container Platform, administrators manage, view, and interact with the control plane nodes directly, which



[facebook.com/redhatinc](https://facebook.com/redhatinc)

@RedHat

[linkedin.com/company/red-hat](https://linkedin.com/company/red-hat)

means that they need to set up their clusters for high availability and adequate security. To be compliant with industry-standard best practices, a minimum of three control-plane nodes should be configured to allow for accessibility to the control plane in a node outage event.

### Infrastructure nodes

These are nodes dedicated to hosting additional functionality such as [OpenShift routes](#) and the [OpenShift internal registry](#). Infrastructure nodes host administrator and network-focused services that are managed separately from your containerized applications.

### App nodes or nodes

These are the OpenShift Container Platform nodes used to run your containerized applications. These are similar to Kubernetes worker nodes and run various monitoring and networking services required across a cluster.

## Cluster design

### Cloud IAM, accounts, and limits

When using a cloud provider, you will want to enforce tight identity and access management (IAM) control of individual clusters and other cloud resources sharing a project. Limit access to resources by applying the principle of least privilege. Understand each provider's account roles and limitations before setting up access to any OpenShift Container Platform clusters.

OpenShift Container Platform helps this process by providing in-depth documentation on the installation process, including installation in [AWS](#), [Azure](#), [GCP](#), and [IBM Z](#).

### Private clusters

Strict network isolation, which prevents unauthorized external ingress to OpenShift cluster API endpoints, nodes, or pod containers, comprises a critical piece of cluster security. By default, OpenShift clusters have Kubernetes cluster API endpoints and nodes with public IP addresses. By default, OpenShift Container Platform is provisioned using publicly-accessible DNS and endpoints. The DNS, Ingress Controller, and API server can be set to private after installing the cluster. Additionally, OpenShift Container Platform may expose operations-focused dashboards for the admins and developers. Ideally, these dashboards will be running on infrastructure nodes away from your high-priority workloads.

The private cluster options vary based on the infrastructure environment. However, there are in-depth guides for setting up a private cluster through various providers. OpenShift Container Platform provides [the installation methods and network setup options that are currently supported](#).

After creating your private cluster, you may need to perform extra configuration steps to ensure your cluster's components are correctly set. Also, upgrades to the cluster may [require internet access and extra considerations](#).

### Setting up a bastion host

A bastion host provides access to a private network from an external network and is a simple way to add an extra layer of security to your OpenShift cluster. A bastion host minimizes the chances of unauthorized access to your OpenShift Container Platform cluster by allowing for more tightly tuned access. Benefits of a bastion host include:

- Separate log-in accounts for everyone accessing the bastion host.
- Auditing of user access and time.
- Specific node access.

A bastion host is a useful way to augment security to your cluster. Restricting access to only specific nodes through the cluster using the bastions .ssh/config allows for private network access and can restrict users from tampering with nodes deemed off-limits.

**Note:** When using a cloud provider for deployment, use software-defined networks that are available. The proper implementation of cloud IAM accounts, firewall rules, and private networking will significantly reduce the attack surface.

### VPC networks

When deploying your OpenShift cluster, you will want to take advantage of the various cloud providers' built-in networking and security protections. This will vary depending on the environment; however, there are defaults and best practices to keep in mind during setup.

1. Create a single, virtual private cloud (VPC) network for each cluster and allow access accordingly.
2. Set up firewall rules to allow for only the ports required. These include:
  - a. [AWS restricted network setup](#).
  - b. [Azure private network setup](#).
  - c. [GCP restricted network setups](#).

### Securing etcd

By default, data stored in etcd is not encrypted at rest in OpenShift Container Platform. Etcd encryption can be enabled in the cluster to effectively provide an additional layer of data security and can be debugged in your cluster to help protect the loss of sensitive data if an etcd backup is exposed to incorrect parties. Because OpenShift Container Platform [recommends an etcd backup during any upgrade](#), encrypting etcd should be a standard practice in your organization.

When you enable etcd encryption, these server resources are encrypted:

- Secrets
- ConfigMaps
- Routes
- OAuth access tokens
- OAuth authorized tokens

When etcd encryption is enabled, encryption keys are created. These keys are rotated every week, and the admin must have these keys to restore from an etcd backup.

### Node images

Compromised nodes create a danger to your entire cluster and its workloads. Using minimal base operating system (OS) images and configuring read-only file systems provides two critical ways to protect your nodes against many attacks and limit their potential blast radius. With minimal images,

attackers have limited tools to use, and if they cannot write or overwrite configuration files and binaries on the node's root file system, they cannot hijack the system as easily nor install their malicious tools.

Providers are increasingly making available minimal, container-optimized OS images such as AWS Bottlerocket and GCP's Container-Optimized OS (COS). However, it is best to leverage Red Hat OpenShift's relationship with the cloud providers and use the most recent Red Hat Enterprise Linux® CoreOS for all of your OpenShift Container Platform cluster's nodes. Red Hat Enterprise Linux CoreOS is the default operating system for all cluster machines; however, you can create worker machines that use Red Hat Enterprise Linux as their operating system.

Red Hat Enterprise Linux CoreOS is designed to be as immutable as possible, allowing for only a few system settings to be changed. These settings are configured remotely, with the help of a specific operator developed by Red Hat OpenShift. This scenario means no user will need to access a node directly, and any changes to the node will need to be directly authorized through the use of the Red Hat machine operator.

## CRI-O

Red Hat Enterprise Linux CoreOS also uses CRI-O as its default container runtime. CRI-O focuses only on features needed by Kubernetes platforms. It also provides a smaller footprint and reduced attack surface than is possible with container engines that include a superset of functionality beyond Kubernetes-centric features. Because OpenShift Container Platform is based on Kubernetes, it benefits from these features as well. By not including extra features for direct command-line use or other orchestration facilities, CRI-O's footprint is smaller, and therefore potential vulnerabilities are reduced.

## Authentication and authorization

### Control access to cluster resources

In addition to using cloud provider IAM roles and authorization, the OpenShift Container Platform control plane includes a builtin OAuth server. This server allows administrators to secure API access control via authentication and authorization regardless of the cluster's deployment location. OAuth 2.0 is the industry-standard protocol for authorization, which works over HTTPS and authorizes devices, servers, etc. with access tokens rather than credentials.

As an administrator, OAuth can be configured to authenticate using an identity provider, such as LDAP, GitHub, or Google. Administrators can obtain OAuth access tokens to authenticate themselves to the API as well. This feature can be enabled by cluster creation or after creation.

### Manage API access

Applications can have multiple, independent API services with different endpoints. Always aim to restrict access to endpoints and services and only grant the minimal access required. In addition to the OAuth server, OpenShift Container Platform includes a containerized version of its [3scale API gateway](#).

Red Hat 3scale API Management gives you a variety of standard options for API authentication and security, which can be used alone or in combination to issue credentials and control access: standard API keys, application ID and key pair, and OAuth 2.0. Using the 3scale API gateway allows for fine-grained control over authorization and exposure for the cluster's API services. For example, quotas and throttling can minimize Denial-of-Service (DoS) attacks.

## Rotate cluster certificates

Kubernetes and OpenShift clusters rely on several secure certificate chains and credentials for security. If sensitive keys or certificates are compromised, the integrity and safety of the entire cluster and its workloads may be placed at risk. Additionally, many security policies and compliance certifications require regular rotation of encryption keys and credentials.

OpenShift Container Platform uses REST-based HTTPS communication with encryption via TLS certificates. These certificates are configured during installation for the components that require HTTPS traffic:

- API server and controllers
- Etcd
- Nodes
- Registry
- Router

OpenShift Container Platform manages these certificates for the administrators and provides more control, such as allowing administrators to [rotate certificates manually](#).

## Remove the `kube-admin` account

By default, a service account is mounted to every pod in an OpenShift Container Platform cluster, allowing containers to send requests to the Kubernetes API server. An attacker who gains access to a pod can obtain the corresponding service account token. With role-based access control (RBAC) enabled by default in an OpenShift Container Platform cluster service account, privileges are determined by role bindings. If these grant elevated privileges, an attacker could send a request to the Kubernetes API server to compromise cluster resources.

Organizations can mitigate this threat vector by configuring Kubernetes RBAC and adopting the least privilege model for service accounts and their role bindings. A core example of this model is [removing the `kube-admin` user](#), with OpenShift Container Platform itself recommending its removal to improve cluster security.

## Networking

The concept of zero-trust security has emerged to address the new security challenges of cloud-native architecture. These challenges include:

- The sharing of cloud infrastructure among workloads with different levels of trust.
- Smaller microservices increasing complexity and enlarging the attack surface of applications.

Microservice architecture creates a more extensive network attack surface. To address this issue, administrators and developers will have to ensure both external networks and internal software-defined networks are securely configured.

## Securing service load balancers

OpenShift Container Platform, at a minimum, requires two load balancers: one to load balance the control plane (the control plane API endpoints) and one for the data plane (the application routers). If a load balancer is created using a cloud provider, the load balancer will be internet-facing and may

have no firewall restrictions. In most on-premise deployments, appliance-based load balancers (such as F5 or Netscaler) are used. Both types of load balancers will need to be configured by the administrator. If the load balancer needs to be internet-facing but should not be open to all IP addresses, you can add the field `loadBalancerSourceRanges` to the service specification to limit the IP address blocks allowed to connect. Verify that your load balancer supports this functionality.

### Enable network policy

By default, network traffic in an OpenShift cluster is allowed between pods and can leave the cluster network altogether. Creating restrictions to allow only necessary service-to-service and cluster ingress and egress connections decrease the number of potential targets for malicious or misconfigured pods and limit their ability to exploit the cluster resources.

The OpenShift Software Defined Network (OpenShift SDN) can control network traffic to and from the cluster's pods by implementing the standard [Kubernetes Network Policy](#) API. Network Policies can control both ingress traffic and block or allow individual IP blocks. NetworkPolicy objects are additive, which means you can combine multiple NetworkPolicy objects to satisfy complex network requirements. Other container network interface (CNI) implementations allow for egress rules also to be set. OpenShift SDN does not currently support that functionality.

All supported versions of OpenShift Container Platform come with Network Policies enabled by default. However, the cluster will still allow for all pod traffic to be accepted. Make sure to deny all traffic by default and create additive rules to limit pod traffic only to what is required. Test the policies to make sure they block only unwanted traffic while allowing required traffic.

### Master authorized networks

To protect against future vulnerabilities in the OpenShift API server and Kubernetes API server, limit network access to API endpoints to trusted IP addresses. Regardless of the OpenShift Container Platform clusters, administrators need to create rules for access to the cluster's API endpoints.

### Container images

Following a few best practices for building secure container images will minimize running containers' exploitability and simplify security updates. Images that contain only the files required for the application's runtime make it much more difficult for attackers to compromise or exploit the containers in your cluster. Avoid using container images with frequently exploited or vulnerable tools like Linux package managers, web or other network clients, or UNIX shells.

#### What to do:

1. Use a minimal, secure base image.
2. Only install tools needed for the container's application. Debugging tools should be omitted from production containers.
3. Remove exploitable tools as part of the initialization process with init-containers.
4. Remove the package manager from the image as a Dockerfile build step.
5. Keep images up to date. This practice includes watching for new versions of both the base image and any third-party tools you install.
6. Maintain compatibility with proper tags. Avoid the use of the "latest" tag to help downstream users avoid compatibility issues.

## Manage images with a registry

OpenShift Container Platform offers an integrated image registry located on the infrastructure nodes of the cluster. This setup allows an organization to avoid third-party hosting and public image storing services, such as Docker Hub. By keeping all necessary images within the cluster, organizations can avoid dependency on a third-party service and any outages associated with them.

Organizations can enact more tuned policies for image use and vetting by taking advantage of a local image repository. Only after an image has passed the proper controls should it be allowed to be used to deploy containers into your cluster.

## Use verified images from the container catalog

To help select secure base images, Red Hat offers a container image catalog. Images in the catalog are maintained by Red Hat, meaning they are frequently scanned and updated. The container catalog also provides an image's history; every time it is updated, a new tag is created. According to a simple scale A to F, the image is graded based on creation date and known security flaws.

## Use an image scanner

Using containers free of known software security vulnerabilities requires ongoing vigilance. All the images deployed to a cluster should be scanned regularly by an image scanner that keeps an up-to-date database of CVEs (common vulnerabilities and exposures). Several open source and proprietary options exist, but be sure to choose a solution that scans for vulnerabilities in operating system packages and in third-party runtime libraries for the programming languages your applications use.

To address CVEs when they are found in your internally maintained images, your organization should have policies for updating and replacing images known to have serious, fixable vulnerabilities for images that are already deployed. Image scanning should be part of your CI/CD pipeline process, and images with high severity, fixable CVEs should generate an alert and fail a build.

## OpenShift Container Security Operator

OpenShift Container Platform allows for a [Container Security Operator](#) (CSO) to be installed in the cluster. This Operator will view results from the OpenShift Container Platform web console for container images used to launch running pods in the cluster. The CSO can be directed to look at all pods or specific namespaces. It will then query the container registry where the image came from for any vulnerabilities. The Operator will also access vulnerabilities based on how long the container has been used without an upgrade or patch.

The CSO's findings can be viewed through the OpenShift Container Platform dashboard that outlines the cluster's overall health and specific images that require attention.

## Runtime security for workloads

### Projects and namespaces

Kubernetes namespaces provide scoping for cluster objects, allowing fine-grained cluster object management. Kubernetes role-based access control (RBAC) rules for most resource types apply at the namespace level. Controls, like network policies and many add-on tools, and frameworks, like service meshes, are also often scoped to the namespace level.

OpenShift expands on the Kubernetes namespace functionality with OpenShift projects. A project is simply a Kubernetes namespace with additional annotations, and just like namespaces, users must be granted access to create and use project resources.

Projects allow the user to add more metadata to a Kubernetes namespace and provide more context in the OpenShift Container Platform web console. Projects can have a separate name, displayName, and description, with displayName and description being option inputs to give a more specific name and description in the web console.

When creating OpenShift projects, it is best to plan out how you will separate the projects in the cluster. Since RBAC and network policies integrate with the Kubernetes namespaces, it is essential to evaluate how their implementation will affect your applications. To start, configuring one namespace per application provides the best opportunity to implement excellent security protocols. The fine-grained control using network policies and RBAC will add slightly more management and more secure scalability. Also, avoid using the default namespace in any cluster outside of a development cluster. All applications require their namespace and should not be deployed into default due to ease of deployment.

### **Kubernetes RBAC**

Kubernetes role-based access control provides the standard method for managing authorization for the Kubernetes API endpoints. The practice of creating and managing comprehensive RBAC roles that follow the principle of least privilege and provide some of the most critical protections possible for your OpenShift Container Platform clusters. By using a principle of least privilege and regular audits, we can limit bad actors, internal misconfigurations, and accidents.

RBAC is enabled by default in OpenShift Container Platform clusters and includes the same set of default cluster roles that can be found in Kubernetes clusters. Configuring RBAC effectively and securely requires some understanding of the Kubernetes API. You can start with the official documentation, read about some best practices, and you can also find an in-depth explanation from the [Red Hat OpenShift documentation](#).

When working with OpenShift Container Platform, we want to create all necessary RBAC resource objects for the cluster workloads and test them in a non-production environment. Once your team has a solid working knowledge of RBAC, create some internal policies and guidelines. Make sure you also regularly audit your role permissions and RoleBindings. Pay special attention to minimizing the use of ClusterRoles and ClusterRoleBindings, as these apply globally across all namespaces and to resources that do not support namespaces. (You can use the output of `kubectl api-resources` in your cluster to see which resources are not namespace-scoped.)

### **Limit container runtime privileges**

Most containerized applications will not need any special host privileges on the node to function correctly. By following the principle of least privilege and minimizing your cluster's running containers' capabilities, you can significantly reduce the level of exploitation for malicious containers and of accidental damage by misbehaving applications. With Kubernetes, the PodSpec Security Context is used to define the exact runtime requirements for each workload. With Red Hat OpenShift Container Platform, security context constraints (SCCs) are used to restrict privileges for pods. Similar to how RBAC resources control user access, administrators can use SCCs to control permissions for pods.

SCCs are OpenShift resources; they define a set of conditions (or rules) that a pod must satisfy to be created (or admitted in the cluster). These controls can limit the resources, system calls, and file system access of the pods running in the cluster. Using SCCs, the level of privileges are controlled for the application, and if needed, give them more permissive or more restrictive privileges.

Red Hat [outlines the capabilities of SCCs](#) in their documentation:

- Whether a pod can run privileged containers.
- The capabilities that a container can request.
- The use of host directories as volumes.
- The SELinux context of the container.
- The container user ID.
- The use of host namespaces and networking.
- The allocation of an FSGroup that owns the pod's volumes.
- The configuration of allowable supplemental groups.
- Whether a container requires the use of a read-only root file system.
- The usage of volume types.
- The configuration of allowable seccomp profiles.

By default, for authenticated users, resources deployed in a project inherit a default security context associated with the authenticated users role. An OpenShift cluster contains eight default SCCs that can be applied to authenticated users:

- anyuid
- hostaccess
- hostmount-anyuid
- Hostnetwork
- node-exporter
- non-root
- privileged
- restricted

Make sure not to tamper with these default SCCs because they are used for essential cluster functions. Instead, create new SCCs for specific users and limit their capabilities.

Some guidelines when creating new SCCs include:

1. Do not allow containers to run as root. Running as root creates the most significant risk since root access in a container is equal to root access on the underlying node.
2. Do not use the host network or process space. Again, these settings create the potential for compromising the node and every container running on it.

3. Do not allow privilege escalation.
4. Use a read-only root filesystem in the container.
5. Use the default (masked) /proc filesystem mount.
6. Drop unused Linux capabilities and do not add optional capabilities that your application does not require. (Available capabilities depend on the container runtime in use on the nodes.)
7. Use SELinux options for more fine-grained process controls.
8. Give each application its own Kubernetes service account rather than sharing or using the namespace's default service account.
9. Do not mount the service account token in a container if it does not need to access the Kubernetes API.

## Monitoring and maintenance

### Cluster upgrades

Support policy:

Keeping up to date with patches for your cluster's OpenShift Container Platform version is an important requirement for your cluster and workload security. OpenShift Container Platform provides full support of the newest released versions until a month after the newest version is released. OpenShift Container Platform will support the maintenance of the three most-recent, minor releases.

Release channels:

OpenShift Container Platform 4.5 offers the following upgrade channels:

- Candidate-4.5
- Fast-4.5
- Stable-4.5

The candidate-4.5 channel contains candidate builds for a z-stream (4.5.z) release. Release candidates contain all the features of the product but are not supported. Use release candidate versions to test feature acceptance and assist in qualifying the next OpenShift Container Platform version. The fast-4.5 channel is updated with new 4.5 versions as soon as Red Hat declares the given version as a general availability release. These releases are fully supported, production quality, and have performed well while available as a release candidate in the candidate-4.5 channel from where they were promoted.

While the fast-4.5 channel contains releases as soon as their errata are published, releases are added to the stable-4.5 channel after a delay. Data is collected from Red Hat site reliability engineering (SRE) teams, Red Hat-support services, and pre-production and production environments that participate in a connected customer program about the release's stability during this delay. You can use the stable-4.5 channel to upgrade from a previous minor version of OpenShift Container Platform.

## Upgrade recommendations

OpenShift Container Platform allows for multiple ways of upgrading a cluster. A minor upgrade can be made through the web console or through the command-line interface. Updating a restricted cluster will need to be done via the command line and will require administrative control and internet access to pull the most recent OpenShift images.

Regardless of the method chosen, make sure to have a recent etcd backup in case of a failed upgrade. Also, ensure all Operators are upgraded to the latest version on their channel. This will ensure that the Operators will be valid for the upgraded cluster.

## Audit logging

Logging events and changes at both the OpenShift level and the node level create a critical audit trail to use for evaluating your cluster's security, especially in case of a breach or attack. Audit logs are available by default and can be easily accessed through the command line. This is vital for restricted clusters that are not using the web interface to view cluster logs.

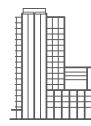
OpenShift Container Platform allows administrators to view audit logs from the node and API [directly from the command line](#). OpenShift Container Platform uses Elasticsearch and Fluentd for logging and storing your cluster events. Make sure to take advantage of these managed open source tools to debug in your cluster effectively.

## Further reading: [Implementing Kubernetes-native security with Red Hat](#)

Security platforms purpose-built to protect Kubernetes offer powerful security and operational advantages. Kubernetes-native security applies controls at the Kubernetes layer, ensuring consistency, automation, and scale. Organizations successfully deploy security as code, enabling security that's built in, not bolted on.

Read this [Kubernetes-native security: What is it and why it matters](#) whitepaper to learn about the key features and benefits of Kubernetes-native security and how it is different from existing container security approaches to deliver protections that are purpose-built for Kubernetes environments.

## About Red Hat



Red Hat is the world's leading provider of open source software solutions, using a community-powered approach to provide reliable and high-performing cloud, Linux, middleware, storage, and virtualization technologies. Red Hat also offers award-winning support, training, and consulting services. As a connective hub in a global network of enterprises, partners, and open source communities, Red Hat helps create relevant, innovative technologies that liberate resources for growth and prepare customers for the future of IT.



[@RedHat](https://facebook.com/redhatinc)  
[linkedin.com/company/red-hat](https://linkedin.com/company/red-hat)

redhat.com  
#F28110\_0421

**NORTH AMERICA**  
1888 REDHAT1

**EUROPE, MIDDLE EAST,  
AND AFRICA**  
00800 7334 2835  
europe@redhat.com

**ASIA PACIFIC**  
+65 6490 4200  
apac@redhat.com

**LATIN AMERICA**  
+54 11 4329 7300  
info-latam@redhat.com