



OpenShift Container Platform 4.9

Release notes

Highlights of what is new and what has changed with this OpenShift Container Platform release

OpenShift Container Platform 4.9 Release notes

Highlights of what is new and what has changed with this OpenShift Container Platform release

Legal Notice

Copyright © 2021 Red Hat, Inc.

The text of and illustrations in this document are licensed by Red Hat under a Creative Commons Attribution–Share Alike 3.0 Unported license ("CC-BY-SA"). An explanation of CC-BY-SA is available at

<http://creativecommons.org/licenses/by-sa/3.0/>

. In accordance with CC-BY-SA, if you distribute this document or an adaptation of it, you must provide the URL for the original version.

Red Hat, as the licensor of this document, waives the right to enforce, and agrees not to assert, Section 4d of CC-BY-SA to the fullest extent permitted by applicable law.

Red Hat, Red Hat Enterprise Linux, the Shadowman logo, the Red Hat logo, JBoss, OpenShift, Fedora, the Infinity logo, and RHCE are trademarks of Red Hat, Inc., registered in the United States and other countries.

Linux[®] is the registered trademark of Linus Torvalds in the United States and other countries.

Java[®] is a registered trademark of Oracle and/or its affiliates.

XFS[®] is a trademark of Silicon Graphics International Corp. or its subsidiaries in the United States and/or other countries.

MySQL[®] is a registered trademark of MySQL AB in the United States, the European Union and other countries.

Node.js[®] is an official trademark of Joyent. Red Hat is not formally related to or endorsed by the official Joyent Node.js open source or commercial project.

The OpenStack[®] Word Mark and OpenStack logo are either registered trademarks/service marks or trademarks/service marks of the OpenStack Foundation, in the United States and other countries and are used with the OpenStack Foundation's permission. We are not affiliated with, endorsed or sponsored by the OpenStack Foundation, or the OpenStack community.

All other trademarks are the property of their respective owners.

Abstract

The release notes for OpenShift Container Platform summarize all new features and enhancements, notable technical changes, major corrections from the previous version, and any known bugs upon general availability.

Table of Contents

CHAPTER 1. OPENSIFT CONTAINER PLATFORM 4.9 RELEASE NOTES	6
1.1. ABOUT THIS RELEASE	6
1.2. OPENSIFT CONTAINER PLATFORM LAYERED AND DEPENDENT COMPONENT SUPPORT AND COMPATIBILITY	6
1.3. NEW FEATURES AND ENHANCEMENTS	6
1.3.1. Red Hat Enterprise Linux CoreOS (RHCOS)	6
1.3.1.1. Installation Ignition config is removed upon boot	6
1.3.2. Installation and upgrade	6
1.3.2.1. Installing a cluster on Microsoft Azure Stack Hub using user-provisioned infrastructure	7
1.3.2.2. Pausing machine health checks before updating the cluster	7
1.3.2.3. Increased size of Azure subnets within the machine CIDR	7
1.3.2.4. Support for AWS regions in China	7
1.3.2.5. Expanding the cluster with Virtual Media on the baremetal network	7
1.3.2.6. Required administrator acknowledgment when upgrading from OpenShift Container Platform 4.8 to 4.9	7
1.3.2.7. Support for installation on RHOSP deployments that use PCI passthrough	8
1.3.2.8. Upgrading etcd version 3.4 to 3.5	8
1.3.2.9. Installing a cluster on IBM Cloud using installer-provisioned infrastructure	8
1.3.2.10. Improved support for Fujitsu hardware on installer-provisioned clusters	8
1.3.3. Web console	8
1.3.3.1. Accessing node logs from the Node page	8
1.3.3.2. Break down cluster utilization by node type	9
1.3.3.3. User preferences	9
1.3.3.4. Hide default projects from project list	9
1.3.3.5. Adding user preferences in the web console	9
1.3.3.6. Developer perspective	9
1.3.4. IBM Z and LinuxONE	9
Notable enhancements	10
Supported features	10
Restrictions	11
1.3.5. IBM Power Systems	11
Notable enhancements	11
Supported features	12
Restrictions	12
1.3.6. Security and compliance	13
1.3.6.1. Configuring the audit log policy with custom rules	13
1.3.6.2. Disabling audit logging	13
1.3.6.3. Customizing the OAuth server URL	13
1.3.6.4. Network-Bound Disk Encryption (NBDE)	13
1.3.7. etcd	14
1.3.7.1. Automatic rotation of etcd certificates	14
1.3.7.2. Additional TLS security profile setting on the API server	14
1.3.8. Networking	14
1.3.8.1. Enhancements to linuxptp services	14
1.3.8.2. Monitoring PTP fast events with the PTP fast event notification framework	14
1.3.8.3. OVN-Kubernetes cluster network provider egress IP feature balances across nodes	14
1.3.8.4. SR-IOV containerized Data Plane Development Kit (DPDK) is GA	14
1.3.8.5. SR-IOV support for using vhost-net with Fast Datapath DPDK applications	14
1.3.8.6. SR-IOV support for single node clusters	14
1.3.8.7. Supported hardware for SR-IOV	15
1.3.8.8. MetalLB load balancer	15

1.3.8.9. CNI VRF plug-in is generally available	15
1.3.8.10. Ingress controller timeout configuration parameters	15
1.3.8.11. Mutual TLS Authentication	16
1.3.8.12. Customizing HAProxy error code response pages	16
1.3.8.13. The provisioningNetworkInterface configuration setting is optional	16
1.3.8.14. DNS Operator managementState	16
1.3.8.15. Load balancer configuration as a cloud provider option for clusters on RHOSP	16
1.3.8.16. Support added for TLS 1.3 and the Modern profile	16
1.3.8.17. Global admission plug-in for HTTP Strict Transport Security requirements	17
1.3.8.18. Ingress empty requests policy	17
1.3.8.19. Create network policies in the web console	17
1.3.9. Storage	17
1.3.9.1. Persistent storage using AWS EBS CSI driver operator is generally available	17
1.3.9.2. Persistent storage using the Azure Stack Hub CSI Driver Operator (general availability)	17
1.3.9.3. Persistent storage using the AWS EFS CSI Driver Operator (Technology Preview)	17
1.3.9.4. Automatic CSI migration supports GCE (Technology Preview)	18
1.3.9.5. Automatic CSI migration supports Azure Disk (Technology Preview)	18
1.3.9.6. VMWare vSphere CSI Driver Operator creates storage policy automatically (Technology Preview)	18
1.3.9.7. New metrics provided for Local Storage Operator	18
1.3.9.8. oVirt CSI driver resizing feature is now available	18
1.3.10. Registry	18
1.3.10.1. Image Registry uses Azure Blob Storage on Azure Stack Hub installations	19
1.3.11. Operator lifecycle	19
1.3.11.1. Operator Lifecycle Manager upgraded to Kubernetes 1.22	19
1.3.11.2. File-based catalogs	19
1.3.11.3. Operator Lifecycle Manager support for Single Node OpenShift	19
1.3.11.4. Enhanced error reporting for cluster administrators	19
1.3.11.4.1. Updating Operator group status conditions	19
1.3.11.4.2. Indicating the reason for install plan failures	20
1.3.11.4.3. Indicating resolution conflicts on subscription statuses	20
1.3.11.5. Image template for custom catalog sources	20
1.3.12. Operator development	20
1.3.12.1. High-availability or single node cluster detection and support	20
1.3.12.2. Operator support for network proxies	20
1.3.12.3. Validating bundle manifests for APIs removed from Kubernetes 1.22	20
1.3.13. Builds	21
1.3.14. Images	21
1.3.14.1. Wildcard domains as registry sources	21
1.3.15. Machine API	21
1.3.15.1. Red Hat Enterprise Linux (RHEL) 8 now supported for compute machines	21
1.3.16. Nodes	21
1.3.16.1. Scheduler profiles GA	21
1.3.16.2. New descheduler profiles and customization	22
1.3.16.3. Multiple logins to the same registry	22
1.3.16.4. Enhanced monitoring of node resources	22
1.3.16.5. Deploy node health checks with the Node Health Check Operator (Technology Preview)	22
1.3.17. Red Hat OpenShift Logging	22
1.3.18. Monitoring	23
1.3.18.1. Monitoring stack components and dependencies	23
1.3.18.2. Alerting rules	23
1.3.18.3. Alertmanager	24
1.3.18.4. Prometheus	24
1.3.18.5. Grafana	25

1.3.19. Metering	25
1.3.20. Scalability and performance	25
1.3.20.1. Special Resource Operator (Technology Preview)	25
1.3.20.2. Memory Manager feature (Technology Preview)	25
1.3.20.3. Additional tools for latency testing	25
1.3.20.4. Cluster maximums	26
1.3.20.5. Zero touch provisioning (Technology Preview)	26
1.3.21. Insights Operator	26
1.3.21.1. Importing RHEL Simple Content Access certificates (Technology Preview)	26
1.3.21.2. Insights Operator data collection enhancements	26
1.3.22. Authentication and authorization	26
1.3.22.1. Support for Microsoft Azure Stack Hub with Cloud Credential Operator in manual mode	26
1.3.23. OpenShift sandboxed containers support on OpenShift Container Platform (Technology Preview)	27
1.4. NOTABLE TECHNICAL CHANGES	27
Automatic defragmentation for etcd data	27
Octavia OVN NodePort changes	27
OpenStack Platform LoadBalancer configuration changes	27
Ingress Controller upgraded to HAProxy 2.2.15	27
CoreDNS update to version 1.8.4	27
Implementation of cloud controller managers for cloud providers	27
Performing a canary rollout update	27
Support for large Operator bundles	27
Reduced resource usage for Operator Lifecycle Manager	28
1.4.1. Default update channel for Operators from "Extras" advisories	28
Operator SDK v1.10.1	28
1.5. DEPRECATED AND REMOVED FEATURES	28
1.5.1. Deprecated features	29
1.5.1.1. SQLite database format for Operator catalogs	30
1.5.1.2. vSphere 6.7 Update 2 and earlier cluster installation and virtual hardware version 13 are now deprecated	30
1.5.1.3. The instance_type_id installation configuration parameter for Red Hat Virtualization (RHV)	30
1.5.2. Removed features	30
1.5.2.1. Metering	30
1.5.2.2. Beta APIs removed from Kubernetes 1.22	30
1.5.2.3. Descheduler v1beta1 API removed	31
1.5.2.4. Use of dhclient in RHCOS removed	31
1.5.2.5. Cease updating the lastTriggeredImageID field and ignore it	32
1.5.2.6. Use of v1 without a group for apiVersion for OpenShift Container Platform resources	32
1.6. BUG FIXES	32
API server and authentication	32
Bare Metal Hardware Provisioning	32
Builds	33
Cloud Compute	33
Cluster Version Operator	34
Console Storage Plug-in	35
Image Registry	35
Installer	35
Kubernetes API server	36
Networking	36
Node	37
OpenShift CLI (oc)	37
Operator Lifecycle Manager (OLM)	38
OpenShift API server	40

OpenShift Update Service	40
Red Hat Enterprise Linux CoreOS (RHCOS)	40
Routing	40
Samples	41
Storage	41
Web console (Administrator perspective)	41
Web console (Developer perspective)	44
1.7. TECHNOLOGY PREVIEW FEATURES	44
1.8. KNOWN ISSUES	46
1.9. ASYNCHRONOUS ERRATA UPDATES	51
1.9.1. RHSA-2021:3759 - OpenShift Container Platform 4.9.0 image release, bug fix, and security update advisory	51
1.9.2. RHBA-2021:3935 - OpenShift Container Platform 4.9.4 bug fix and security update	52
1.9.2.1. Enhancements	52
1.9.2.2. Bug fixes	52
1.9.2.3. Upgrading	52
1.9.3. RHBA-2021:4005 - OpenShift Container Platform 4.9.5 bug fix update	52
1.9.3.1. Known Issues	53
1.9.3.2. Bug fixes	53
1.9.3.3. Upgrading	53
1.9.4. RHBA-2021:4119 - OpenShift Container Platform 4.9.6 bug fix and security update	53
1.9.4.1. Known Issues	53
1.9.4.2. Bug fixes	53
1.9.4.3. Upgrading	54
1.9.5. RHBA-2021:4579 - OpenShift Container Platform 4.9.7 bug fix update	54
1.9.5.1. Features	54
1.9.5.1.1. Updates from Kubernetes 1.22.2	54
1.9.5.2. Upgrading	54
1.9.6. RHBA-2021:4712 - OpenShift Container Platform 4.9.8 bug fix update	54
1.9.6.1. Bug fixes	54
1.9.6.2. Upgrading	54
CHAPTER 2. OPENSIFT CONTAINER PLATFORM VERSIONING POLICY	56

CHAPTER 1. OPENSIFT CONTAINER PLATFORM 4.9 RELEASE NOTES

Red Hat OpenShift Container Platform provides developers and IT organizations with a hybrid cloud application platform for deploying both new and existing applications on secure, scalable resources with minimal configuration and management overhead. OpenShift Container Platform supports a wide selection of programming languages and frameworks, such as Java, JavaScript, Python, Ruby, and PHP.

Built on Red Hat Enterprise Linux (RHEL) and Kubernetes, OpenShift Container Platform provides a more secure and scalable multi-tenant operating system for today's enterprise-class applications, while delivering integrated application runtimes and libraries. OpenShift Container Platform enables organizations to meet security, privacy, compliance, and governance requirements.

1.1. ABOUT THIS RELEASE

OpenShift Container Platform ([RHSA-2021:3759](#)) is now available. This release uses [Kubernetes 1.22](#) with CRI-O runtime. New features, changes, and known issues that pertain to OpenShift Container Platform 4.9 are included in this topic.

OpenShift Container Platform 4.9 clusters are available at <https://cloud.redhat.com/openshift>. The Red Hat OpenShift Cluster Manager application for OpenShift Container Platform allows you to deploy OpenShift clusters to either on-premise or cloud environments.

OpenShift Container Platform 4.9 is supported on Red Hat Enterprise Linux (RHEL) 7.9 and 8.4, as well as on Red Hat Enterprise Linux CoreOS (RHCOS) 4.9.

You must use RHCOS machines for the control plane, and you can use either RHCOS or Red Hat Enterprise Linux (RHEL) 7.9 or 8.4 for compute machines.

1.2. OPENSIFT CONTAINER PLATFORM LAYERED AND DEPENDENT COMPONENT SUPPORT AND COMPATIBILITY

The scope of support for layered and dependent components of OpenShift Container Platform changes independently of the OpenShift Container Platform version. To determine the current support status and compatibility for an add-on, refer to its release notes. For more information, see the [Red Hat OpenShift Container Platform Life Cycle Policy](#).

1.3. NEW FEATURES AND ENHANCEMENTS

This release adds improvements related to the following components and concepts.

1.3.1. Red Hat Enterprise Linux CoreOS (RHCOS)

1.3.1.1. Installation Ignition config is removed upon boot

Nodes installed with the **coreos-installer** program previously retained the installation Ignition config in the **/boot/ignition/config.ign** file. Starting with the OpenShift Container Platform 4.9 installation image, that file is removed when the node is provisioned. This change does not affect clusters that were installed on previous OpenShift Container Platform versions because they still use an older bootimage.

1.3.2. Installation and upgrade

1.3.2.1. Installing a cluster on Microsoft Azure Stack Hub using user-provisioned infrastructure

OpenShift Container Platform 4.9 introduces support for installing a cluster on Azure Stack Hub using user-provisioned infrastructure.

You can incorporate example Azure Resource Manager (ARM) templates provided by Red Hat to assist in the deployment process, or create your own. You are also free to create the required resources through other methods; the ARM templates are just an example.

See [Installing a cluster on Azure Stack Hub using ARM templates](#) for details.

1.3.2.2. Pausing machine health checks before updating the cluster

During the upgrade process, nodes in the cluster might become temporarily unavailable. In the case of worker nodes, the machine health check might identify such nodes as unhealthy and reboot them. To avoid rebooting such nodes, OpenShift Container Platform 4.9 introduces the **cluster.x-k8s.io/paused=""** annotation to let you pause the **MachineHealthCheck** resources before updating the cluster.

For more information, see [Pausing a MachineHealthCheck resource](#).

1.3.2.3. Increased size of Azure subnets within the machine CIDR

The OpenShift Container Platform installation program for Microsoft Azure now creates subnets as large as possible within the machine CIDR. This lets the cluster use a machine CIDR that is appropriately sized to accommodate the number of nodes in the cluster.

1.3.2.4. Support for AWS regions in China

OpenShift Container Platform 4.9 introduces support for AWS regions in China. You can now install and update OpenShift Container Platform clusters in the **cn-north-1** (Beijing) and **cn-northwest-1** (Ningxia) regions.

For more information, see [Installing a cluster on AWS China](#).

1.3.2.5. Expanding the cluster with Virtual Media on the baremetal network

In OpenShift Container Platform 4.9, you can expand an installer provisioned cluster deployed using the **provisioning** network by using Virtual Media on the **baremetal** network. You can use this feature when the **ProvisioningNetwork** configuration setting is set to **Managed**. To use this feature, you must set the **virtualMediaViaExternalNetwork** configuration setting to **true** in the **provisioning** custom resource (CR). You must also edit the machine set to use the API VIP address. See [Preparing to deploy with Virtual Media on the baremetal network](#) for details.

1.3.2.6. Required administrator acknowledgment when upgrading from OpenShift Container Platform 4.8 to 4.9

OpenShift Container Platform 4.9 uses Kubernetes 1.22, which removed a [significant number of deprecated v1beta1 APIs](#).

OpenShift Container Platform 4.8.14 introduced a requirement that an administrator must provide a manual acknowledgment before the cluster can be upgraded from OpenShift Container Platform 4.8 to 4.9. This is to help prevent issues after upgrading to OpenShift Container Platform 4.9, where APIs that have been removed are still in use by workloads, tools, or other components running on or interacting

with the cluster. Administrators must evaluate their cluster for any APIs in use that will be removed and migrate the affected components to use the appropriate new API version. After this is done, the administrator can provide the administrator acknowledgment.

All OpenShift Container Platform 4.8 clusters require this administrator acknowledgment before they can be upgraded to OpenShift Container Platform 4.9.

For more information, see [Preparing to update to OpenShift Container Platform 4.9](#) .

1.3.2.7. Support for installation on RHOSP deployments that use PCI passthrough

OpenShift Container Platform 4.9 introduces support for installation on Red Hat OpenStack Platform (RHOSP) deployments that rely on [PCI passthrough](#).

1.3.2.8. Upgrading etcd version 3.4 to 3.5

OpenShift Container Platform 4.9 supports etcd 3.5. Before you upgrade the cluster, verify that a valid etcd backup exists. An etcd backup ensures that the cluster can be restored if an upgrade failure occurs. In OpenShift Container Platform 4.9, etcd upgrades are automatic. Depending on the cluster's transition state to version 4.9, an etcd backup might be available. However, verifying that a backup exists before the cluster upgrade starts is recommended.

1.3.2.9. Installing a cluster on IBM Cloud using installer-provisioned infrastructure

OpenShift Container Platform 4.9 introduces support for installing a cluster on IBM Cloud® using installer-provisioned infrastructure. The procedure is nearly identical to installer-provisioned infrastructure on bare metal with these differences:

- Installer-provisioned installation of OpenShift Container Platform 4.9 on IBM Cloud requires the **provisioning** network, IPMI, and PXE boot. Red Hat does not support deployment with Redfish and virtual media on IBM Cloud.
- You must create and configure public and private VLANs on the IBM Cloud.
- IBM Cloud nodes must be available before starting the installation process. So you must create the IBM Cloud nodes first.
- You must prepare the provisioner node.
- You must install and configure a DHCP server on the public **baremetal** network.
- You must configure the **install-config.yaml** file so that each node points to the BMC using IPMI, and sets the IPMI privilege level to **OPERATOR**.

See [Deploying installer-provisioned clusters on IBM Cloud](#) for details.

1.3.2.10. Improved support for Fujitsu hardware on installer-provisioned clusters

OpenShift Container Platform 4.9 adds BIOS configuration support for worker nodes when deploying installer-provisioned clusters on Fujitsu hardware and using the Fujitsu integrated Remote Management Controller (iRMC). See [Configuring BIOS for worker node](#) for details.

1.3.3. Web console

1.3.3.1. Accessing node logs from the Node page

With this update, administrators now have the ability to access node logs from the **Node** page. To review the node logs, you can switch between individual log files and journal log units by clicking the **Logs** tab.

1.3.3.2. Break down cluster utilization by node type

You now have the ability to filter by node type in the **Cluster utilization** card on the cluster dashboard. Additional node types will appear in the list when created.

1.3.3.3. User preferences

This update adds a **User Preferences** page for customizing settings, such as default project, perspective, and topology view.

1.3.3.4. Hide default projects from project list

With this update, you can hide **default projects** from the **Projects** dropdown in the web console masthead. You can still toggle to show **default projects** before you search and filter.

1.3.3.5. Adding user preferences in the web console

With this update, you can now add user preferences in the web console. Users can select their default perspective, project, topology, and other preferences.

1.3.3.6. Developer perspective

- You can now import a devfile, a Dockerfile, or a builder image through your Git repository to further customize your deployment. You can also edit the file import type and select a different strategy for importing the file.
- You can now add tasks in a pipeline using **Add task** and **Quick Search** using the updated user interface of the **Pipeline builder** in the developer console. This enhanced experience allows users to add tasks from the **Tekton Hub**.
- To edit your build configurations, you use the **Edit BuildConfig** option in the **Builds** view of the **Developer** perspective. Users can use a **Form view** and a **YAML view** to edit the build configurations.
- You can use the context menu in the topology **Graph view** to add services or create a connection with operator-backed services to the projects.
- You can use the **+Add** actions in the context menu of the topology **Graph view** to add services or remove a service in the application group.
- Initial support for **pipeline as code** is now available in the **Pipelines Repository list** view, enabled by the OpenShift Pipelines Operator.
- Usability enhancement have been made to the **Application Monitoring** section in the **Observe** page of the topology.

1.3.4. IBM Z and LinuxONE

With this release, IBM Z and LinuxONE are now compatible with OpenShift Container Platform 4.9. The installation can be performed with z/VM or RHEL KVM. For installation instructions, see the following documentation:

- [Installing a cluster with z/VM on IBM Z and LinuxONE](#)
- [Installing a cluster with z/VM on IBM Z and LinuxONE in a restricted network](#)
- [Installing a cluster with RHEL KVM on IBM Z and LinuxONE](#)
- [Installing a cluster with RHEL KVM on IBM Z and LinuxONE in a restricted network](#)

Notable enhancements

The following new features are supported on IBM Z and LinuxONE with OpenShift Container Platform 4.9:

- Helm
- Support for multiple network interfaces
- Service Binding Operator

Supported features

The following features are also supported on IBM Z and LinuxONE:

- Currently, the following Operators are supported:
 - Cluster Logging Operator
 - NFD Operator
 - OpenShift Elasticsearch Operator
 - Local Storage Operator
 - Service Binding Operator
- Encrypting data stored in etcd
- Multipathing
- Persistent storage using iSCSI
- Persistent storage using local volumes (Local Storage Operator)
- Persistent storage using hostPath
- Persistent storage using Fibre Channel
- Persistent storage using Raw Block
- OVN-Kubernetes
- Three-node cluster support
- z/VM Emulated FBA devices on SCSI disks
- 4K FCP block device

These features are available only for OpenShift Container Platform on IBM Z and LinuxONE for 4.9:

- HyperPAV enabled on IBM Z and LinuxONE for the virtual machines for FICON attached ECKD storage

Restrictions

Note the following restrictions for OpenShift Container Platform on IBM Z and LinuxONE:

- The following OpenShift Container Platform Technology Preview features are unsupported:
 - Precision Time Protocol (PTP) hardware
- The following OpenShift Container Platform features are unsupported:
 - Automatic repair of damaged machines with machine health checking
 - CodeReady Containers (CRC)
 - Controlling overcommit and managing container density on nodes
 - CSI volume cloning
 - CSI volume snapshots
 - FIPS cryptography
 - Multus CNI plug-in
 - NVMe
 - OpenShift Metering
 - OpenShift Virtualization
 - Tang mode disk encryption during OpenShift Container Platform deployment
- Worker nodes must run Red Hat Enterprise Linux CoreOS (RHCOS)
- Persistent shared storage must be provisioned by using either NFS or other supported storage protocols
- Persistent non-shared storage must be provisioned using local storage, like iSCSI, FC, or using LSO with DASD, FCP, or EDEV/FBA

1.3.5. IBM Power Systems

With this release, IBM Power Systems are now compatible with OpenShift Container Platform 4.9. For installation instructions, see the following documentation:

- [Installing a cluster on IBM Power Systems](#)
- [Installing a cluster on IBM Power Systems in a restricted network](#)

Notable enhancements

The following new features are supported on IBM Power Systems with OpenShift Container Platform 4.9:

- Helm
- Support for Power10

- Support for multiple network interfaces
- Service Binding Operator

Supported features

The following features are also supported on IBM Power Systems:

- Currently, the following Operators are supported:
 - Cluster Logging Operator
 - NFD Operator
 - OpenShift Elasticsearch Operator
 - Local Storage Operator
 - SR-IOV Network Operator
 - Service Binding Operator
- Multipathing
- Persistent storage using iSCSI
- Persistent storage using local volumes (Local Storage Operator)
- Persistent storage using hostPath
- Persistent storage using Fibre Channel
- Persistent storage using Raw Block
- OVN-Kubernetes
- 4K Disk Support
- NVMe
- Encrypting data stored in etcd
- Three-node cluster support
- Multus SR-IOV

Restrictions

Note the following restrictions for OpenShift Container Platform on IBM Power Systems:

- The following OpenShift Container Platform Technology Preview features are unsupported:
 - Precision Time Protocol (PTP) hardware
- The following OpenShift Container Platform features are unsupported:
 - Automatic repair of damaged machines with machine health checking
 - CodeReady Containers (CRC)
 - Controlling overcommit and managing container density on nodes

- FIPS cryptography
- OpenShift Metering
- OpenShift Virtualization
- Tang mode disk encryption during OpenShift Container Platform deployment
- Worker nodes must run Red Hat Enterprise Linux CoreOS (RHCOS)
- Persistent storage must be of the Filesystem type that uses local volumes, Network File System (NFS), or Container Storage Interface (CSI)

1.3.6. Security and compliance

1.3.6.1. Configuring the audit log policy with custom rules

You now have more fine-grained control over the audit logging level for OpenShift Container Platform. You can use custom rules to specify a different audit policy profile (**Default**, **WriteRequestBodies**, **AllRequestBodies**, or **None**) for different groups.

For more information, see [Configuring the audit log policy with custom rules](#).

1.3.6.2. Disabling audit logging

You can now disable audit logging for OpenShift Container Platform by using the **None** audit policy profile.



WARNING

It is not recommended to disable audit logging unless you are fully aware of the risks of not logging data that can be beneficial when troubleshooting issues. If you disable audit logging and a support situation arises, you might need to enable audit logging and reproduce the issue in order to troubleshoot properly.

For more information, see [Disabling audit logging](#).

1.3.6.3. Customizing the OAuth server URL

You can now customize the URL for the internal OAuth server. For more information, see [Customizing the internal OAuth server URL](#).

1.3.6.4. Network-Bound Disk Encryption (NBDE)

OpenShift Container Platform 4.9 provides new procedures for ongoing maintenance of NBDE-configured systems. NBDE allows you to encrypt root volumes of hard drives on physical and virtual machines without having to manually enter a password when restarting machines. For more information, see [About disk encryption technology](#).

1.3.7. etcd

1.3.7.1. Automatic rotation of etcd certificates

In OpenShift Container Platform 4.9, etcd certificates are automatically rotated and are managed by the system.

1.3.7.2. Additional TLS security profile setting on the API server

The Kubernetes API server TLS security profile setting is now also honored by etcd.

1.3.8. Networking

1.3.8.1. Enhancements to linuxptp services

OpenShift Container Platform 4.9 introduces the following updates to PTP:

- New **ptp4lConf** field
- New option to configure **linuxptp** services as a boundary clock

For more information, see [Configuring linuxptp services as boundary clock](#).

1.3.8.2. Monitoring PTP fast events with the PTP fast event notification framework

Fast event notifications for PTP events are now available for bare-metal clusters. The PTP Operator generates event notifications for every configured PTP-capable network interface. Events are made available through a REST API for applications running on the same node. Fast event notifications are transported by an Advanced Message Queuing Protocol (AMQP) message bus provided by the AMQ Interconnect Operator.

For more information, see [About PTP and clock synchronization error events](#).

1.3.8.3. OVN-Kubernetes cluster network provider egress IP feature balances across nodes

The egress IP feature of OVN-Kubernetes now balances network traffic approximately equally across nodes for a given namespace, if that namespace is assigned multiple egress IP addresses. Each IP address must reside on a different node. For more information, refer to [Configuring egress IPs for a project](#) for OVN-Kubernetes.

1.3.8.4. SR-IOV containerized Data Plane Development Kit (DPDK) is GA

The containerized Data Plane Development Kit (DPDK) is now GA in OpenShift Container Platform 4.9. For more information, see [Using virtual functions \(VFs\) with DPDK and RDMA modes](#).

1.3.8.5. SR-IOV support for using vhost-net with Fast Datapath DPDK applications

SR-IOV now supports vhost-net for use with Fast Datapath DPDK applications on Intel and Mellanox NICs. You can enable this feature by configuring the **SriovNetworkNodePolicy** resource. For more information, see [SR-IOV network node configuration object](#).

1.3.8.6. SR-IOV support for single node clusters

Single node clusters support SR-IOV hardware and the SR-IOV Network Operator. Be aware that configuring an SR-IOV network device causes the single node to reboot and that you must configure the **disableDrain** field for the Operator. For more information, see [Configuring the SR-IOV Network Operator](#).

1.3.8.7. Supported hardware for SR-IOV

OpenShift Container Platform 4.9 adds support for additional Broadcom and Intel hardware.

- Broadcom BCM57414 and BCM57508
- Intel E810-CQDA2, E810-XXVDA2, and E810-XXVDA4

For more information, see the [supported devices](#).

1.3.8.8. MetalLB load balancer

This release introduces the MetalLB Operator. After installing and configuring the MetalLB Operator, you can deploy MetalLB to provide a native load balancer implementation for services on bare-metal clusters. Other on-premise infrastructures that are like bare metal can also benefit.

The Operator introduces a custom resource, **AddressPool**. You configure address pools with ranges of IP addresses that MetalLB can assign to services. When you add a service of type **LoadBalancer**, MetalLB assigns an IP address from a pool.

For this release, Red Hat only supports using MetalLB in layer 2 mode.

For more information, see [About MetalLB and the MetalLB Operator](#).

1.3.8.9. CNI VRF plug-in is generally available

The CNI VRF plug-in was previously introduced as a Technology Preview feature in OpenShift Container Platform 4.7 and is now generally available in OpenShift Container Platform 4.9.

For more information, see [Assigning a secondary network to a VRF](#).

1.3.8.10. Ingress controller timeout configuration parameters

This release introduces six timeout configurations for the Ingress Controller **tuningOptions** parameter:

- **clientTimeout** specifies how long a connection is held open while waiting for a client response.
- **serverFinTimeout** specifies how long a connection is held open while waiting for the server response to the client that is closing the connection.
- **serverTimeout** specifies how long a connection is held open while waiting for a server response.
- **clientFinTimeout** specifies how long a connection is held open while waiting for the client response to the server closing the connection.
- **tlsInspectDelay** specifies how long the router can hold data to find a matching route.
- **tunnelTimeout** specifies how long a tunnel connection, including WebSocket connections, remains open while the tunnel is idle.

For more information, see [Ingress controller configuration parameters](#).

1.3.8.11. Mutual TLS Authentication

You can now configure the Ingress Controller to enable mutual TLS (mTLS) authentication by setting **spec.clientTLS**. The **clientTLS** field specifies configuration for the Ingress Controller to verify client certificates.

For more information, see [Configuring Mutual TLS Authentication](#).

1.3.8.12. Customizing HAProxy error code response pages

Cluster administrators can specify a custom HTTP error code response page for either 503, 404, or both error pages.

For more information, see [Customizing HAProxy error code response pages](#).

1.3.8.13. The provisioningNetworkInterface configuration setting is optional

In OpenShift Container Platform 4.9, the **provisioningNetworkInterface** configuration setting for installer-provisioned clusters is optional. The **provisioningNetworkInterface** configuration setting identifies the NIC name used for the **provisioning** network. In OpenShift Container Platform 4.9, you can alternatively specify the **bootMACAddress** configuration setting in the **install-config.yml** file, which enables Ironic to identify the IP address for the NIC connected to the **provisioning** network and bind to it. You can also omit the **provisioningInterface** configuration setting in the provisioning custom resource so that the provisioning custom resource uses the **bootMACAddress** configuration setting instead.

1.3.8.14. DNS Operator managementState

In OpenShift Container Platform 4.9, you can now change the DNS Operator **managementState**. The **managementState** of the DNS Operator is set to **Managed** by default, which means that the DNS Operator is actively managing its resources. You can change it to **Unmanaged**, which means the DNS Operator is not managing its resources.

The following are use cases for changing the DNS Operator **managementState**:

- You are a developer and want to test a configuration change to see if it fixes an issue in CoreDNS. You can stop the DNS Operator from overwriting the change by setting the **managementState** to **Unmanaged**.
- You are a cluster administrator and have reported an issue with CoreDNS, but need to apply a workaround until the issue is fixed. You can set the **managementState** field of the DNS Operator to **Unmanaged** to apply the workaround.

For more information, see [Changing the DNS Operator managementState](#).

1.3.8.15. Load balancer configuration as a cloud provider option for clusters on RHOSP

For clusters that run on RHOSP, you can now configure Octavia for load balancing as a cloud provider option.

For more information, see [Setting cloud provider options](#).

1.3.8.16. Support added for TLS 1.3 and the Modern profile

This release adds Ingress Controller support for TLS 1.3 and the **Modern** profile in HAProxy.

For more information, see [Ingress Controller TLS security profiles](#) .

1.3.8.17. Global admission plug-in for HTTP Strict Transport Security requirements

Cluster administrators can configure HTTP Strict Transport Security (HSTS) verification on a per-domain basis with the addition of an admission plug-in for the router, called **route.openshift.io/RequiredRouteAnnotations**. If a cluster administrator configures this plug-in to enforce HSTS, then any newly created route must be configured with a compliant HSTS Policy, which is verified against the global setting on the cluster Ingress configuration, called **ingresses.config.openshift.io/cluster**.

For more information, see [HTTP Strict Transport Security](#) .

1.3.8.18. Ingress empty requests policy

In OpenShift Container Platform 4.9 you can now configure the Ingress Controller to log or ignore empty requests by setting the **logEmptyRequests** and **HTTPEmptyRequestsPolicy** fields.

For more information, see [Ingress controller configuration parameters](#).

1.3.8.19. Create network policies in the web console

Logging in to the web console with the **cluster-admin** role now enables you to create new network policies in any namespace in the cluster from a form in the console. Previously, this could only be done directly in YAML.

1.3.9. Storage

1.3.9.1. Persistent storage using AWS EBS CSI driver operator is generally available

OpenShift Container Platform is capable of provisioning persistent volumes (PVs) using the Container Storage Interface (CSI) driver for AWS Elastic Block Store (EBS). This feature was previously introduced as a Technology Preview feature in OpenShift Container Platform 4.5 and is now generally available and enabled by default in OpenShift Container Platform 4.9.

For more information, see [AWS EBS CSI Driver Operator](#) .

1.3.9.2. Persistent storage using the Azure Stack Hub CSI Driver Operator (general availability)

OpenShift Container Platform is capable of provisioning PVs using the CSI driver for Azure Stack Hub Storage. Azure Stack Hub, which is part of the Azure Stack portfolio, allows you to run apps in an on-premises environment and deliver Azure services in your datacenter. The Azure Stack Hub CSI Driver Operator that manages this driver is new for 4.9 and generally available.

For more information, see [Azure Stack Hub CSI Driver Operator](#) .

1.3.9.3. Persistent storage using the AWS EFS CSI Driver Operator (Technology Preview)

OpenShift Container Platform is capable of provisioning PVs using the CSI driver for AWS Elastic File Service (EFS). The AWS EFS CSI Driver Operator that manages this driver is in Technology Preview.

For more information, see [AWS EFS CSI Driver Operator](#) .

1.3.9.4. Automatic CSI migration supports GCE (Technology Preview)

Starting with OpenShift Container Platform 4.8, automatic migration for in-tree volume plug-ins to their equivalent CSI drivers became available as a Technology Preview feature. This feature now supports automatic migration from Google Compute Engine Persistent Disk (GCE PD) in-tree plug-in to the Google Cloud Platform (GCP) Persistent Disk CSI driver.

For more information, see [CSI Automatic Migration](#).

1.3.9.5. Automatic CSI migration supports Azure Disk (Technology Preview)

Starting with OpenShift Container Platform 4.8, automatic migration for in-tree volume plug-ins to their equivalent CSI drivers became available as a Technology Preview feature. This feature now supports automatic migration from the Azure Disk in-tree plug-in to the Azure Disk CSI driver.

For more information, see [CSI Automatic Migration](#).

1.3.9.6. VMWare vSphere CSI Driver Operator creates storage policy automatically (Technology Preview)

The vSphere CSI Operator Driver storage class now uses vSphere's storage policy. OpenShift Container Platform automatically creates a storage policy that targets datastore configured in cloud configuration.

For more information, see [VMWare vSphere CSI Driver Operator](#).

1.3.9.7. New metrics provided for Local Storage Operator

OpenShift Container Platform 4.9 provides the following new metrics for the Local Storage Operator:

- **iso_discovery_disk_count**: total number of discovered devices on each node
- **iso_lvset_provisioned_PV_count**: total number of PVs created by **LocalVolumeSet** objects
- **iso_lvset_unmatched_disk_count**: total number of disks that Local Storage Operator did not select for provisioning because of mismatching criteria
- **iso_lvset_orphaned_symlink_count**: number of devices with PVs that no longer match **LocalVolumeSet** object criteria
- **iso_lv_orphaned_symlink_count**: number of devices with PVs that no longer match **LocalVolume** object criteria
- **iso_lv_provisioned_PV_count**: total number of provisioned PVs for **LocalVolume**

For more information, see [Persistent storage using local volumes](#).

1.3.9.8. oVirt CSI driver resizing feature is now available

OpenShift Container Platform 4.9 adds resizing capability to the oVirt CSI Driver, which allows users to increase the size of their existing persistent volume claims (PVCs). Prior to this feature, users had to create new PVCs with the increased size, and move all of the content from the old persistent volume (PV) to the new PV, which could result in data loss. Now, users can edit the existing PVC and the oVirt CSI Driver will resize the underlying oVirt disk.

1.3.10. Registry

1.3.10.1. Image Registry uses Azure Blob Storage on Azure Stack Hub installations

In OpenShift Container Platform 4.9, the integrated Image Registry uses Azure Blob Storage for clusters installed on Microsoft Azure Stack Hub using user-provisioned infrastructure.

See [Installing a cluster on Azure Stack Hub using ARM templates](#) for details.

1.3.11. Operator lifecycle

The following new features and enhancements relate to running Operators with Operator Lifecycle Manager (OLM).

1.3.11.1. Operator Lifecycle Manager upgraded to Kubernetes 1.22

Starting in OpenShift Container Platform 4.9, Operator Lifecycle Manager (OLM) supports Kubernetes 1.22. As a result, [a significant number of v1beta1 APIs have been removed and updated to v1](#). Operators that depend on the removed **v1beta1** APIs will not run on OpenShift Container Platform 4.9. Cluster administrators should [upgrade their installed Operators](#) to the **latest** channel before upgrading a cluster to OpenShift Container Platform 4.9.



IMPORTANT

Kubernetes 1.22 introduces [several notable changes](#) to **v1** of the **CustomResourceDefinition** API.

1.3.11.2. File-based catalogs

File-based catalogs are the latest iteration of the catalog format in Operator Lifecycle Manager (OLM). The format is a plain text-based (JSON or YAML) and declarative config evolution of the earlier, and now deprecated, [SQLite database format](#), and it is fully backwards compatible. The goal of this format is to enable Operator catalog editing, composability, and extensibility.

For more information about the file-based catalog specification, see [Operator Framework packaging format](#).

For instructions about creating file-based catalogs by using the **opm** CLI, see [Managing custom catalogs](#).

1.3.11.3. Operator Lifecycle Manager support for Single Node OpenShift

Operator Lifecycle Manager (OLM) is now available on Single Node OpenShift (SNO) clusters, enabling self-service Operator installations.

1.3.11.4. Enhanced error reporting for cluster administrators

Because administrators should not require an understanding of the interaction process between the various low-level APIs or access to the Operator Lifecycle Manager (OLM) pod logs to successfully debug such issues, OpenShift Container Platform 4.9 introduces the following enhancements in OLM to provide administrators with more comprehensible error reporting and messages:

1.3.11.4.1. Updating Operator group status conditions

Previously, if a namespace contained multiple Operator groups or could not find a service account, the status of the Operator group would not report an error. With this enhancement, these scenarios now update the status condition of the Operator group to report an error.

1.3.11.4.2. Indicating the reason for install plan failures

Before this release, if an install plan failed, the subscription condition would not state why the failure occurred. Now, if an install plan fails, the subscription status condition indicates the reason for the failure.

1.3.11.4.3. Indicating resolution conflicts on subscription statuses

Because dependency resolution treats all components in a namespace as a single unit, if a resolution failure occurs, all subscriptions on the namespace now indicate the error.

1.3.11.5. Image template for custom catalog sources

To avoid cluster upgrades potentially leaving Operator installations in an unsupported state or without a continued update path, you can enable automatically changing your Operator catalog's index image version as part of cluster upgrades.

Set the **olm.catalogImageTemplate** annotation to your catalog image name and use one or more of the Kubernetes cluster version variables when constructing the template for the image tag.

For more information, see [Image template for custom catalog sources](#).

1.3.12. Operator development

The following new features and enhancements relate to developing Operators with the Operator SDK.

1.3.12.1. High-availability or single node cluster detection and support

An OpenShift Container Platform cluster can be configured in high-availability (HA) mode, which uses multiple nodes, or in non-HA mode, which uses a single node. A single node cluster, also known as Single Node OpenShift (SNO), is likely to have more conservative resource constraints. Therefore, it is important that Operators installed on a single node cluster can adjust accordingly and still run well.

By accessing the cluster high-availability mode API provided in OpenShift Container Platform, Operator authors can use the Operator SDK to enable their Operator to detect a cluster's infrastructure topology, either HA or non-HA mode. Custom Operator logic can be developed that uses the detected cluster topology to automatically switch the resource requirements, both for the Operator and for any Operands or workloads it manages, to a profile that best fits the topology.

For more information, see [High-availability or single node cluster detection and support](#).

1.3.12.2. Operator support for network proxies

Operator authors can now develop Operators that support network proxies. Operators with proxy support inspect the Operator deployment for environment variables and pass the variables on to the required Operands. Cluster administrators configure proxy support for the environment variables that are handled by Operator Lifecycle Manager (OLM). For more information, see the Operator SDK tutorials for developing Operators using [Go](#), [Ansible](#), and [Helm](#).

1.3.12.3. Validating bundle manifests for APIs removed from Kubernetes 1.22

You can now check bundle manifests for APIs removed from Kubernetes 1.22 by using the Operator Framework suite of tests with the **bundle validate** subcommand.

For example:

```
$ operator-sdk bundle validate .<bundle_dir_or_image> \  
--select-optional suite=operatorframework \  
--optional-values=k8s-version=1.22
```

If your bundle manifest includes APIs removed from Kubernetes 1.22, the command displays a warning message. The warning message indicates which APIs you need to migrate and links to the Kubernetes API migration guide.

See the [table of beta APIs removed from Kubernetes 1.22](#) and the [Operator SDK CLI reference](#) for more information.

1.3.13. Builds

As a developer using OpenShift Container Platform for builds, with this update, you can use the following new capabilities:

- You can mount build volumes to give running builds access to information that you do not want to persist in the output container image. Build volumes can provide sensitive information, such as repository credentials, which the build environment or configuration only needs at build-time. Build volumes are different from build inputs, whose data can persist in the output container image.
- You can configure image changes to trigger builds based on information recorded in the BuildConfig status. This way, you can use **ImageChange** triggers with builds in a GitOps workflow.

1.3.14. Images

1.3.14.1. Wildcard domains as registry sources

This release introduces support for using wildcard domains as registry sources in your image registry settings. With a wildcard domain, such as ***.example.com**, you can set your cluster to push and pull images from multiple subdomains without having to manually enter each one. For more information, see [Image controller configuration parameters](#).

1.3.15. Machine API

1.3.15.1. Red Hat Enterprise Linux (RHEL) 8 now supported for compute machines

Starting in OpenShift Container Platform 4.9, you can now use Red Hat Enterprise Linux (RHEL) 8.4 for compute machines. Previously, RHEL 8 was not supported for compute machines.

You cannot upgrade RHEL 7 compute machines to RHEL 8. You must deploy new RHEL 8 hosts, and the old RHEL 7 hosts should be removed.

1.3.16. Nodes

1.3.16.1. Scheduler profiles GA

Scheduling pods using a scheduler profile is now generally available. This is a replacement for configuring a scheduler policy. The following scheduler profiles are available:

- **LowNodeUtilization:** This profile attempts to spread pods evenly across nodes to get low resource usage per node.
- **HighNodeUtilization:** This profile attempts to place as many pods as possible onto as few nodes as possible, to minimize node count with high usage per node.
- **NoScoring:** This is a low-latency profile that strives for the quickest scheduling cycle by disabling all score plug-ins. This might sacrifice better scheduling decisions for faster ones.

For more information, see [Scheduling pods using a scheduler profile](#) .

1.3.16.2. New descheduler profiles and customization

The following descheduler profiles are now available:

- **SoftTopologyAndDuplicates:** This profile is the same as **TopologyAndDuplicates**, except that pods with soft topology constraints, such as **whenUnsatisfiable: ScheduleAnyway**, are also considered for eviction.
- **EvictPodsWithLocalStorage:** This profile allows pods with local storage to be eligible for eviction.
- **EvictPodsWithPVC:** This profile allows pods with persistent volume claims to be eligible for eviction.

You can also customize the pod lifetime value for the **LifecycleAndUtilization** profile.

For more information, see [Evicting pods using the descheduler](#) .

1.3.16.3. Multiple logins to the same registry

When configuring the **docker/config.json** file to allow pods to pull images from private registries, you can now list specific repositories in the same registry, each with credentials specific to that registry path. Previously, you could list only one repository from a given registry. You can also now define a registry with a specific namespace.

1.3.16.4. Enhanced monitoring of node resources

Node-related metrics and alerts have been enhanced to give you an earlier indication of when the stability of a node is compromised.

1.3.16.5. Deploy node health checks with the Node Health Check Operator (Technology Preview)

You can use the Node Health Check Operator to deploy the **NodeHealthCheck** controller. The controller identifies unhealthy nodes and uses the Poison Pill Operator to remediate the unhealthy nodes.

1.3.17. Red Hat OpenShift Logging

In OpenShift Container Platform 4.7, *Cluster Logging* became *Red Hat OpenShift Logging* . For more information, see [Release notes for Red Hat OpenShift Logging](#) .

1.3.18. Monitoring

The monitoring stack for this release includes the following new and modified features.

1.3.18.1. Monitoring stack components and dependencies

Updates to versions of monitoring stack components and dependencies include the following:

- Prometheus to 2.29.2
- The Prometheus Operator to 0.49.0
- The Prometheus Adapter to 0.9.0
- Alertmanager to 0.22.2
- Thanos to 0.22.0

1.3.18.2. Alerting rules

- **New**
 - **HighlyAvailableWorkloadIncorrectlySpread** informs you about a potential problem when two instances of a highly available monitoring component are running on the same node and have persistent volumes attached.
 - **NodeFileDescriptorLimit** triggers an alert when a node kernel is running out of available file descriptors. A warning level alert fires at greater than 70% usage, and a critical level alert fires at greater than 90% usage.
 - **PrometheusLabelLimitHit** detects when a target exceeds the defined label limits.
 - **PrometheusTargetSyncFailure** detects when Prometheus fails to synchronize targets.
 - All critical alerting rules contain links to runbooks.
- **Enhanced**
 - **AlertmanagerReceiversNotConfigured** and **KubePodCrashLooping** now contain fewer false positives.
 - **KubeCPUOvercommit** and **KubeMemoryOvercommit** are now more robust in non-homogeneous environments.
 - The **for** duration setting of the **NodeFilesystemAlmostOutOfSpace** alerting rule has changed from one hour to 30 minutes so that the system more quickly detects when disk space runs low.
 - **KubeDeploymentReplicasMismatch** now fires as expected. In previous versions, this alert did not fire.
 - The following alerts now contain a **namespace** label:
 - **AlertmanagerReceiversNotConfigured**
 - **KubeClientErrors**

- **KubeCPUOvercommit**
- **KubeletDown**
- **KubeMemoryOvercommit**
- **MultipleContainersOOMKilled**
- **ThanosQueryGrpcClientErrorRate**
- **ThanosQueryGrpcServerErrorRate**
- **ThanosQueryHighDNSFailures**
- **ThanosQueryHttpRequestQueryErrorRateHigh**
- **ThanosQueryHttpRequestQueryRangeErrorRateHigh**
- **ThanosSidecarPrometheusDown**
- **Watchdog**



NOTE

Red Hat does not guarantee backward compatibility for metrics, recording rules, or alerting rules.

1.3.18.3. Alertmanager

- You can add and configure additional external Alertmanagers for both platform and user-defined project monitoring stacks.
- You can disable the local Alertmanager instance.

1.3.18.4. Prometheus

- You can enable and configure remote write storage for both platform monitoring and user-defined projects in Prometheus. This feature enables you to send ingested metrics to long-term storage.
- To reduce the overall memory consumption of Prometheus, the following cAdvisor metrics with both an empty **pod** and **namespace** label have been dropped:
 - **container_fs_.***
 - **container_spec_.***
 - **container_blkio_device_usage_total**
 - **container_file_descriptors**
 - **container_sockets**
 - **container_threads_max**
 - **container_threads**

- **container_start_time_seconds**
- **container_last_seen**
- When persistent storage is not configured for platform monitoring, upgrades and cluster disruptions can lead to data loss. A warning message has been added to the **Degraded** condition when the system detects that persistent storage is not configured for platform monitoring.
- You can exclude individual user-defined projects from the **openshift-user-workload-monitoring** project by adding the **openshift.io/user-monitoring: "false"** label to them.
- You can configure an **enforcedTargetLimit** parameter for the **openshift-user-workload-monitoring** project to set an overall limit on the number of targets scraped.

1.3.18.5. Grafana

Because running the default Grafana dashboard can take resources from user workloads, you can disable the Grafana dashboard deployment.

1.3.19. Metering

This release removes the OpenShift Container Platform Metering Operator.

1.3.20. Scalability and performance

1.3.20.1. Special Resource Operator (Technology Preview)

You can now use the Special Resource Operator (SRO) to help manage the deployment of kernel modules and drivers on an existing OpenShift Container Platform cluster. This is currently a Technology Preview feature.

For more information, see [About the Special Resource Operator](#).

1.3.20.2. Memory Manager feature (Technology Preview)

The Memory Manager feature is now enabled by default for all pods running on the node that is configured with one of the following Topology Manager policies:

- **single-numa-node**
- **restricted**

For more information, see [Topology Manager policies](#).

1.3.20.3. Additional tools for latency testing

OpenShift Container Platform 4.9 introduces two additional tools to measure system latency:

- **hwlatdetect** measures the baseline that the bare hardware can achieve
- **cyclictest** schedules a repeated timer after **hwlatdetect** passes validation and measures the difference between the desired and the actual trigger times

For more information, see [Running the latency tests](#).

1.3.20.4. Cluster maximums

Updated guidance around [cluster maximums](#) for OpenShift Container Platform 4.9 is now available.



IMPORTANT

No large scale testing for performance against OVN-Kubernetes testing was executed for this release.

Use the [OpenShift Container Platform Limit Calculator](#) to estimate cluster limits for your environment.

1.3.20.5. Zero touch provisioning (Technology Preview)

OpenShift Container Platform 4.9 supports zero touch provisioning (ZTP), which allows you to provision new edge sites with declarative configurations of bare metal equipment at remote sites. ZTP uses the GitOps deployment set of practices for infrastructure deployment. GitOps achieves these tasks using declarative specifications stored in Git repositories, such as YAML files and other defined patterns, to provide a framework for deploying the infrastructure. The declarative output is leveraged by the Open Cluster Manager (OCM) for multisite deployment. For more information, see [Provisioning edge sites at scale](#).

1.3.21. Insights Operator

1.3.21.1. Importing RHEL Simple Content Access certificates (Technology Preview)

In OpenShift Container Platform 4.9, Insights Operator can import RHEL Simple Content Access (SCA) certificates from Red Hat OpenShift Cluster Manager.

For more information, see [Importing RHEL Simple Content Access certificates with Insights Operator](#) .

1.3.21.2. Insights Operator data collection enhancements

In OpenShift Container Platform 4.9, the Insights Operator collects the following additional information:

- All of the **MachineConfig** resource definitions from a cluster.
- The names of the **PodSecurityPolicies** installed in a cluster.
- If installed, the **ClusterLogging** resource definition.
- If the **SamplesImagestreamImportFailing** alert is firing, then the **ImageStream** definitions and the last 100 lines of container logs from the **openshift-cluster-samples-operator** namespace.

With this additional information, Red Hat can provide improved remediation steps in Insights Advisor.

1.3.22. Authentication and authorization

1.3.22.1. Support for Microsoft Azure Stack Hub with Cloud Credential Operator in manual mode

With this release, installations on Microsoft Azure Stack Hub can be performed by configuring the Cloud Credential Operator (CCO) in manual mode.

For more information, see [Using manual mode](#).

1.3.23. OpenShift sandboxed containers support on OpenShift Container Platform (Technology Preview)

To review OpenShift sandboxed containers new features, bug fixes, known issues, and asynchronous errata updates, see [OpenShift sandboxed containers 1.1 release notes](#).

1.4. NOTABLE TECHNICAL CHANGES

OpenShift Container Platform 4.9 introduces the following notable technical changes.

Automatic defragmentation for etcd data

In OpenShift Container Platform 4.9, etcd data is automatically defragmented by the etcd Operator.

Octavia OVN NodePort changes

Previously, on Red Hat OpenStack Platform (RHOSP) deployments, opening traffic on NodePorts was constrained to the CIDR of the node's subnet. In order to support LoadBalancer services using the Octavia Open Virtual Network (OVN) provider, the security group rules that allow NodePort traffic to master and worker nodes are now changed to open **0.0.0.0/0**.

OpenStack Platform LoadBalancer configuration changes

The Red Hat OpenStack Platform (RHOSP) cloud provider LoadBalancer configuration now defaults to **use-octavia=True**. An exception to this rule is a deployment with Kuryr, in which case **use-octavia** is set to **false**, because Kuryr handles LoadBalancer services on its own.

Ingress Controller upgraded to HAProxy 2.2.15

The OpenShift Container Platform Ingress Controller is upgraded to HAProxy version 2.2.15.

CoreDNS update to version 1.8.4

In OpenShift Container Platform 4.9, CoreDNS uses version 1.8.4, which includes bug fixes.

Implementation of cloud controller managers for cloud providers

The Kubernetes controller manager that manages cloud provider deployments does not include support for Azure Stack Hub as a provider. Because using cloud controller managers is the preferred method for interacting with underlying cloud platforms, there is no plan to add this support. As a result, the Azure Stack Hub implementation in OpenShift Container Platform uses cloud controller managers.

In addition, this release supports using cloud controller managers for Amazon Web Services (AWS), Microsoft Azure, and Red Hat OpenStack Platform (RHOSP) as a [Technology Preview](#). Any new cloud platform support that is added to OpenShift Container Platform will also use cloud controller managers.

To learn more about the cloud controller manager, see the [Kubernetes documentation on this component](#).

To manage the cloud controller manager and cloud node manager deployments and lifecycles, this release introduces the Cluster Cloud Controller Manager Operator.

For more information, see the [Cluster Cloud Controller Manager Operator](#) entry in the *Red Hat Operators reference*.

Performing a canary rollout update

With OpenShift Container Platform 4.9, a new process to perform a canary rollout update has been introduced. For a detailed overview of this process, see [Performing a canary rollout update](#).

Support for large Operator bundles

Operator Lifecycle Manager (OLM) now compresses Operator bundles with large amounts of metadata, such as large custom resource definition (CRD) manifests, to stay below the 1 MB limit set by etcd.

Reduced resource usage for Operator Lifecycle Manager

Operator Lifecycle Management (OLM) catalog pods are now more efficient and use less RAM.

1.4.1. Default update channel for Operators from "Extras" advisories

Operators that ship with OpenShift Container Platform "Extras" advisories, such as [RHBA-2021:3760](#), are published in Red Hat-provided catalogs and run on Operator Lifecycle Manager (OLM). Starting with OpenShift Container Platform 4.9, these Operators are now included in a **stable** update channel in addition to the version-specific **4.9** channel.

For OpenShift Container Platform 4.9 and future releases, **stable** will be the default channel for these Operators. Cluster administrators should use the **stable** channel so that changing update channels for these Operators in OLM is no longer necessary with future cluster upgrades.

For more information about OLM-based Operators, see [Red Hat-provided Operator catalogs](#) and [Understanding OperatorHub](#). For more information about update channels in OLM, see [Upgrading installed Operators](#).

Operator SDK v1.10.1

OpenShift Container Platform 4.9 supports Operator SDK v1.10.1. See [Installing the Operator SDK CLI](#) to install or update to this latest version.



NOTE

Operator SDK v1.10.1 supports Kubernetes 1.21.

If you have any Operator projects that were previously created or maintained with Operator SDK v1.8.0, see [Upgrading projects for newer Operator SDK versions](#) to ensure your projects are upgraded to maintain compatibility with Operator SDK v1.10.1.

1.5. DEPRECATED AND REMOVED FEATURES

Some features available in previous releases have been deprecated or removed.

Deprecated functionality is still included in OpenShift Container Platform and continues to be supported; however, it will be removed in a future release of this product and is not recommended for new deployments. For the most recent list of major functionality deprecated and removed within OpenShift Container Platform 4.9, refer to the table below. Additional details for more fine-grained functionality that has been deprecated and removed are listed after the table.

In the table, features are marked with the following statuses:

- **GA:** *General Availability*
- **TP:** *Technology Preview*
- **DEP:** *Deprecated*
- **REM:** *Removed*

Table 1.1. Deprecated and removed features tracker

Feature	OCP 4.7	OCP 4.8	OCP 4.9
Package manifest format (Operator Framework)	DEP	REM	REM
SQLite database format for Operator catalogs	GA	GA	DEP
oc adm catalog build	DEP	REM	REM
--filter-by-os flag for oc adm catalog mirror	DEP	REM	REM
v1beta1 CRDs	DEP	DEP	REM
Docker Registry v1 API	DEP	DEP	REM
Metering Operator	DEP	DEP	REM
Scheduler policy	DEP	DEP	DEP
ImageChangesInProgress condition for Cluster Samples Operator	DEP	DEP	DEP
MigrationInProgress condition for Cluster Samples Operator	DEP	DEP	DEP
Use of v1 without a group in apiVersion for OpenShift Container Platform resources	DEP	DEP	REM
Use of dhclient in RHCOS	DEP	DEP	REM
Cluster Loader	GA	DEP	DEP
Bring your own RHEL 7 compute machines	DEP	DEP	DEP
lastTriggeredImageID field in the BuildConfig spec for Builds	GA	DEP	REM
Jenkins Operator	TP	DEP	DEP
HPA custom metrics adapter based on Prometheus	TP	REM	REM
vSphere 6.7 Update 2 or earlier and virtual hardware version 13	GA	GA	DEP
The instance_type_id installation configuration parameter for Red Hat Virtualization (RHV)	DEP	DEP	DEP

1.5.1. Deprecated features

1.5.1.1. SQLite database format for Operator catalogs

The SQLite database format used by Operator Lifecycle Manager (OLM) for catalogs and index images has been deprecated, including the related **opm** CLI commands. Cluster administrators and catalog maintainers are encouraged to familiarize themselves with the new [file-based catalog format](#) introduced in OpenShift Container Platform 4.9 and begin migrating catalog workflows.



NOTE

The default [Red Hat-provided Operator catalogs](#) for OpenShift Container Platform 4.6 and later are currently still shipped in the SQLite database format.

1.5.1.2. vSphere 6.7 Update 2 and earlier cluster installation and virtual hardware version 13 are now deprecated

Installing a cluster on VMware vSphere version 6.7 Update 2 or earlier and virtual hardware version 13 is now deprecated. Support for these versions will end in a future version of OpenShift Container Platform.

Hardware version 15 is now the default for vSphere virtual machines in OpenShift Container Platform. Hardware version 15 will be the only supported version in a future version of OpenShift Container Platform.

1.5.1.3. The `instance_type_id` installation configuration parameter for Red Hat Virtualization (RHV)

The `instance_type_id` installation configuration parameter is deprecated and will be removed in a future release.

1.5.2. Removed features

1.5.2.1. Metering

This release removes the OpenShift Container Platform Metering Operator feature.

1.5.2.2. Beta APIs removed from Kubernetes 1.22

Kubernetes 1.22 removed the following deprecated **v1beta1** APIs. Migrate manifests and API clients to use the **v1** API version. For more information about migrating removed APIs, see the [Kubernetes documentation](#).

Table 1.2. **v1beta1** APIs removed from Kubernetes 1.22

Resource	API	Notable changes
APIService	apiregistration.k8s.io/v1beta1	No
CertificateSigningRequest	certificates.k8s.io/v1beta1	Yes
ClusterRole	rbac.authorization.k8s.io/v1beta1	No
ClusterRoleBinding	rbac.authorization.k8s.io/v1beta1	No

Resource	API	Notable changes
CSIDriver	storage.k8s.io/v1beta1	No
CSINode	storage.k8s.io/v1beta1	No
CustomResourceDefinition	apiextensions.k8s.io/v1beta1	Yes
Ingress	extensions/v1beta1	Yes
Ingress	networking.k8s.io/v1beta1	Yes
IngressClass	networking.k8s.io/v1beta1	No
Lease	coordination.k8s.io/v1beta1	No
LocalSubjectAccessReview	authorization.k8s.io/v1beta1	Yes
MutatingWebhookConfiguration	admissionregistration.k8s.io/v1beta1	Yes
PriorityClass	scheduling.k8s.io/v1beta1	No
Role	rbac.authorization.k8s.io/v1beta1	No
RoleBinding	rbac.authorization.k8s.io/v1beta1	No
SelfSubjectAccessReview	authorization.k8s.io/v1beta1	Yes
StorageClass	storage.k8s.io/v1beta1	No
SubjectAccessReview	authorization.k8s.io/v1beta1	Yes
TokenReview	authentication.k8s.io/v1beta1	No
ValidatingWebhookConfiguration	admissionregistration.k8s.io/v1beta1	Yes
VolumeAttachment	storage.k8s.io/v1beta1	No

1.5.2.3. Descheduler v1beta1 API removed

The deprecated **v1beta1** API for the descheduler has been removed in OpenShift Container Platform 4.9. Migrate any resources using the descheduler **v1beta1** API version to **v1**.

1.5.2.4. Use of dhclient in RHCOS removed

The deprecated **dhclient** binary has been removed from RHCOS. Starting with OpenShift Container Platform 4.6, RHCOS switched to using **NetworkManager** in the **initramfs** to configure networking during early boot. Use the **NetworkManager** internal DHCP client for networking configuration instead. See [BZ#1908462](#) for more information.

1.5.2.5. Cease updating the `lastTriggeredImageID` field and ignore it

The current release stops updating the `buildConfig.spec.triggers[i].imageChange.lastTriggeredImageID` field when the `ImageStreamTag` referenced by `buildConfig.spec.triggers[i].imageChange` points to a new image. Instead, this release updates the `buildConfig.status.imageChangeTriggers[i].lastTriggeredImageID` field.

Additionally, the Build Image Change Trigger controller ignores the `buildConfig.spec.triggers[i].imageChange.lastTriggeredImageID` field.

Now, the Build Image Change Trigger controller starts a build based on the `buildConfig.status.imageChangeTriggers[i].lastTriggeredImageID` field and how it compares to the image ID now referenced by the `ImageStreamTag` referenced in the `buildConfig.spec.triggers[i].imageChange`.

Therefore, update scripts and jobs that inspect `buildConfig.spec.triggers[i].imageChange.lastTriggeredImageID` accordingly. ([BUILD-190](#))

1.5.2.6. Use of `v1` without a group for `apiVersion` for OpenShift Container Platform resources

Support for using `v1` without a group for `apiVersion` for OpenShift Container Platform resources has been removed. Every resource that includes `*.openshift.io` must match the `apiVersion` value found in the [API index](#).

1.6. BUG FIXES

API server and authentication

- Previously, encryption conditions could remain indefinitely and be reported as a degraded condition for some Operators. Stale encryption conditions are now cleared properly and no longer improperly reported. ([BZ#1974520](#))
- Previously, the CA for API server client certificates was rotated early in the lifetime of a cluster, which prevented the Authentication Operator from creating a certificate signing request (CSR) because a previous CSR with the same name still existed. The Kubernetes API server was unable to authenticate itself to the OAuth API server when sending **TokenReview** requests, which caused authentication to fail. Generated names are now used when creating CSRs by the Authentication Operator, so an early rotation of the CA for API server client certificates no longer causes authentication failures. ([BZ#1978193](#))

Bare Metal Hardware Provisioning

- Previously, `metal3` pods could not download an Red Hat Enterprise Linux CoreOS (RHCOS) image due to the sequencing of creating `initContainers`. This issue is fixed by reordering the creation of the `initContainers`, so that the `metal-static-ip-set` `initContainer` is created before the `metal3-machine-os-downloader` `initContainer`. The RHCOS image now downloads as expected. ([BZ#1973724](#))
- Previously, when using installer-provisioned installation on bare metal with a host configured to use `idrac-virtualmedia`, the `bios_interface` for that host got set to `idrac-wsman` by default.

This resulted in the BIOS settings being unavailable and an exception occurring. This issue is fixed by using **idrac-redfish** for the default **bios_interface** when using **idrac-virtualmedia**. ([BZ#1928816](#))

- Previously, in UEFI mode, the **ironic-python-agent** created a UEFI bootloader entry after downloading the RHCOS image. When using an RHCOS image based on RHEL 8.4, the image could fail to boot using this entry and output a BIOS error screen. This is fixed by the **ironic-python-agent** configuring the boot entry based on a CSV file located in the image, instead of using a fixed boot entry. The image boots properly without error. ([BZ#1966129](#))
- Previously, if **provisioningHostIP** had been set in **install-config**, it was assigned to the metal3 pod, even in cases where the provisioning network had been disabled. This has been fixed. ([BZ#1972753](#))
- Previously, the assisted installer could not provision Supermicro X11/X12-based systems because of a mismatch in the sushy resource library. The mismatch resulted in an installation issue by being unable to attach virtual media to the **Inserted** and **WriteProtected** attributes, and not being allowed in the **VirtualMedia.InsertMedia** request body. This issue is fixed by modifying the sushy resource library, and adding a condition to stop sending these optional attributes when not strictly required, thus allowing the installation to progress past this point. ([BZ#1986238](#))
- Previously, some error types in the provisioned state caused the host to be deprovisioned. This occurred after a restart of the metal3 pod if the image provisioned to a bare metal host became unavailable. In this case, the host would enter the deprovisioning state. This issue is fixed by modifying the action of the error in the provisioned state so that if the image becomes unavailable, the error will be reported but deprovisioning will not be initiated. ([BZ#1972374](#))

Builds

- In OpenShift Container Platform and later, the fix for bug [BZ#1884270](#) incorrectly pruned SSH protocol URLs in an attempt to provide SCP-styled URL capabilities. This error caused the **oc new-build** command not to pick an automatic source clone secret: the build could not use the **build.openshift.io/sbuild.openshift.io/source-secret-match-uri-1ource-secret-match-uri-1** annotation to map SSH keys with the associated secrets, and therefore could not perform git cloning. This update reverts the changes from [BZ#1884270](#) so that builds can use the annotation and perform git cloning.
- Before this update, various allowed and block registry configuration options of the cluster image configuration could block the Cluster Samples Operator from creating image streams. When that happened, the samples operator marked itself as **degraded**, which impacted the general OpenShift Container Platform install and upgrade status. The Cluster Samples Operator can bootstrap itself as **removed** in a variety of circumstances. With this update, these circumstances include when the [image controller configuration parameters](#) prevent the creation of image streams by using the default image registry or by using the image registry specified by the [samplesRegistry setting](#). The Operator status also indicates when the cluster image configuration prevents the creation of the sample image streams.

Cloud Compute

- Previously, when a root volume was created for a new server, and that server was not successfully created, the automatic deletion for the volume was not triggered because there was no deletion of a server associated with the volume. In some situations, this led to the creation of many extra volumes, and caused errors if the volume quota was reached. With this release, newly created root volumes are deleted when server creation call fails. ([BZ#1943378](#))

- Previously, when using the default value for **instanceType**, the Machine API created **m4.large** instances on AWS. This is different than the **m5.large** instance type for machines that are created by the OpenShift Container Platform installer. With this release, the Machine API creates **m5.large** instances for new machines on AWS when the default value is specified. ([BZ#1953063](#))
- Previously, the machine set definitions of compute nodes could not specify whether a port should be trunked. This was a problem for technologies that require the user to configure trunked and non-trunked ports for the same machines. This release adds a new field, **spec.Port.Trunk = bool**, which gives the user more flexibility to determine which ports result in trunks. If no value is specified, **spec.Port.Trunk** inherits the value of **spec.Trunk** and the name of the trunk created matches the name of the port used. ([BZ#1964540](#))
- Previously, the Machine API Operator constantly attached new targets even if they were already attached. The excessive calls to the AWS API resulted in a high number of errors. With this release, the Operator checks whether a load balancer attachment is required before attempting the attachment process. This change reduces the frequency of failed API requests. ([BZ#1965080](#))
- Previously, when using automatic pinning for a VM, the name of the property was **disabled**, **existing**, or **adjust**. With this release, the name better describes each policy, and **existing** was removed because it is blocked on oVirt. The new property names are **none** and **resize_and_pin**, which align with the oVirt user interface. ([BZ#1972747](#))
- Previously, the cluster autoscaler was unable to access the **csidrivers.storage.k8s.io** or **csistoragecapacities.storage.k8s.io** resources, which resulted in permissions errors. This fix updates the role assigned to the cluster autoscaler so that it includes permissions for these resources. ([BZ#1973567](#))
- Previously, it was possible to delete a machine with a node that has been deleted. This caused the machine to remain in a deleting phase indefinitely. This fix allows you to delete machines in this state properly. ([BZ#1977369](#))
- When using **boot-from-volume** image, creating a new instance leaks volumes if the machine controller is rebooted. This caused the previously created volume to never be cleaned up. This fix ensures that the volume created previously is either pruned or reused. ([BZ#1983612](#))
- Previously, the Red Hat Virtualization (RHV) provider ignored NICs with **br-ex** names for machines. Since a network type of **OVNKubernetes** creates a NIC with a **br-ex** name, this resulted in the machine never getting an IP address on OVN-Kubernetes. With this fix, it is now possible to install OpenShift Container Platform on RHV with network set to **OVNKubernetes**. ([BZ#1984481](#))
- Previously, when deployed on Red Hat OpenStack Platform (RHOSP) with a combination of proxy and custom CA certificate, a cluster would not become fully operational. This fix passes the proxy settings to the HTTP transport used when connecting with a custom CA certificate, ensuring that all cluster components work as expected. ([BZ#1986540](#))

Cluster Version Operator

- Previously, the Cluster Version Operator (CVO) did not respect the **noProxy** property in the proxy configuration resource. As a result, the CVO was denied access to update recommendations or release signatures when only unproxied connections completed. Now, when the proxy resource requests direct, unproxied access, the CVO reaches the upstream update service and signature stores directly. ([BZ#1978749](#))
- Previously, the Cluster Version Operator (CVO) loaded its proxy configuration from proxy

resource specification properties instead of from status properties that were verified by the Network Operator. As a result, any incorrectly configured values would prevent the CVO from reaching the upstream update service or signature stores. Now, the CVO loads its proxy configuration only from the verified status properties. ([BZ#1978774](#))

- Previously, the Cluster Version Operator (CVO) did not remove volume mounts that were added outside of the manifest. As a result, pod creation could fail during a volume failure. Now, CVO removes all volume mounts that do not appear in the manifest. ([BZ#2004568](#))

Console Storage Plug-in

- Previously, when working with Ceph storage, the Console Storage Plug-in unnecessarily included a redundant use of a namespace parameter. This bug had no customer-visible impact; however, the plug-in has been updated to avoid the redundant use of the namespace. ([BZ#1982682](#))

Image Registry

- The Operator to check if the registry should use custom tolerations was checking **spec.nodeSelector** instead of **spec.tolerations**. The custom tolerations from **spec.tolerations** are applied only when **spec.nodeSelector** is set. This fix uses the field **spec.tolerations** to check for the presence of custom tolerations. Now, the Operator uses custom tolerations if **spec.tolerations** are set. ([BZ#1973318](#))
- The **spec.managementState** in **configs.imageregistry** is set to **Removed**, which caused the image pruner pod to generate warnings about deprecated CronJob in **v1.21** and later, and that **batch/v1** should be used. This fix updates **batch/v1beta1** with **batch/v1** in OpenShift Container Platform **oc**. Now, warnings about the deprecated CronJob in image pruner pods no longer appear. ([BZ#1976112](#))

Installer

- Previously, the network interface on Azure control plane nodes was missing a hyphen in the interface name. This was inconsistent compared to other platforms, which caused issues. The missing hyphen has been added. Now all control plane nodes are named the same, regardless of the platform. ([BZ#1882490](#))
- You can now configure the **autoPinningPolicy** and **hugepages** fields in the **install-config.yaml** file for oVirt. The **autoPinningPolicy** field allows you to automatically set the non-uniform memory access (NUMA) pinning settings and CPU topology changes for the cluster. The **hugepages** field allows you to set the Hugepages of the hypervisor. ([BZ#1925203](#))
- Previously, the installation program did not output any errors when the Ed25519 SSH key type was used with FIPS enabled, even though it could not be used. Now the installation program validates SSH key types, outputting an error when an SSH key type is not supported with FIPS enabled; only RSA and ECDSA SSH key types are allowed when FIPS is enabled. ([BZ#1962414](#))
- In certain conditions, Red Hat OpenStack Platform (RHOSP) network trunks do not contain a tag to indicate that the trunk belongs to the cluster. Consequently, cluster deletion missed the trunk ports and got stuck in a loop until they timed out. Deleting the cluster now deletes trunks for which the tagged port is a parent. ([BZ#1971518](#))
- Previously, when uninstalling a cluster on Red Hat OpenStack Platform (RHOSP), the installer used an inefficient algorithm to delete resources. The inefficient algorithm caused the uninstall process to require more time than was necessary. The installer is updated with a more efficient algorithm that should uninstall the cluster more quickly. ([BZ#1974598](#))

- Previously, if the **AWS_SHARED_CREDENTIALS_FILE** environment variable was set to an empty file, the installer prompted for credentials and then created a **aws/credentials** file, ignoring the value of the environment variable and possibly overwriting existing credentials. With this fix, the installer is updated to store credentials in the specified file. If the specified file has invalid credentials, the installer produces an error instead of overwriting the file and risking information loss. ([BZ#1974640](#))
- Previously, users encountered a vague error message when they deleted a cluster on Azure that shared resources with another cluster, making it difficult to understand why the deletion failed. This update adds an error message that explains why the failure occurs. ([BZ#1976016](#))
- Previously, because of a typo, Kuryr deployments were being checked against the wrong requirements, meaning that installations with Kuryr could succeed even if they did not meet the minimum requirements for Kuryr. This fix eliminates the error, allowing the installer to check the right requirements. ([BZ#1978213](#))
- Before this update, the ingress checks for **keepalived** did not include fall and raise directives, which meant that a single failed check could cause an ingress virtual IP failover. This bug fix introduces fall and raise directives to prevent such failovers. ([BZ#1982766](#))

Kubernetes API server

- Previously, when a deployment and image stream were created at the same time, a race condition could occur which caused the deployment controller to create replica sets in an infinite loop. The responsibilities of the API server's image policy plug-in were lowered and concurrent creation of a deployment and image stream no longer causes infinite replica sets. ([BZ#1925180](#)), ([BZ#1976775](#))
- Previously, there was a race between the installer pod and the cert-syncer container, which were writing to the same path. This could leave some certificates empty and prevent the server from running. Kubernetes API server certificates are now written in an atomic way to prevent races between multiple processes. ([BZ#1971624](#))

Networking

- When using the OVN-Kubernetes cluster network provider, the logical flow cache was configured without any memory limit. As a result, in some situations high memory pressure could cause a node to become unusable. With this update, the logical flow cache is configured with a 1 GB memory limit by default. ([BZ#1961757](#))
- When using the OVN-Kubernetes cluster network provider, any network policies created in a OpenShift Container Platform 4.5 cluster that was subsequently upgraded might allow or drop unexpected traffic. In later versions of OpenShift Container Platform, OVN-Kubernetes uses a different convention for managing IP address sets, and any network policies created in OpenShift Container Platform 4.5 did not use this convention. Now, during upgrades all network policies are migrated to the new convention. ([BZ#1962387](#))
- For the OVN-Kubernetes cluster network provider, when using **must-gather** to retrieve Open vSwitch (OVS) logs, the **INFO** log level was missing from the gathered logging data. Now all log levels are included in OVS logging data. ([BZ#1970129](#))
- Previously, performance testing revealed that the service controller metrics had a significant increase in cardinality due to a label requirement. As a result, memory usage was elevated on Open Virtual Network (OVN) Prometheus pods. With this update, the label requirement is removed. Service controller cardinality metrics and memory usage are now reduced. ([BZ#1974967](#))

- Previously, **ovnkube-trace** required `iproute` to be installed in the source and/or destination pod because it needed to detect the interfaces `link` index. This causes **ovnkube-trace** to fail on pods if there is no `iproute` installed. Now, you can get the `link` index from `/sys/class/net/<interface>/iflink` instead of `iproute`. As a result, **ovnkube-trace** no longer requires `iproute` to be installed in source and destination pods. ([BZ#1978137](#))
- Previously, the Cluster Network Operator (CNO) deployed a service monitor for the **network-check-source** service to get discovered by Prometheus without correct annotations and role-based access control (RBAC). As a result, the service and its metrics never populated in Prometheus. Now, the correct annotations and RBAC are added to the namespace of **network-check-source** service. Now, metrics of service **network-check-source** get scraped by Prometheus. ([BZ#1986061](#))
- Previously, when using IPv6 DHCP, node interface addresses might be leased with a `/128` prefix. Consequently, OVN-Kubernetes uses the same prefix to infer the node's network and routes any other address traffic, including traffic to other cluster nodes, through the gateway. With this update, OVN-Kubernetes inspects the node's routing table and checks for the wider routing entry for the node's interface address and uses that prefix to infer the node's network. As a result, traffic to other cluster nodes is no longer routed through the gateway. ([BZ#1980135](#))
- Previously, when a cluster used the OVN-Kubernetes Container Network Interface provider, attempting to add an egress router with IPv6 address failed. With the fix, support for IPv6 is added to the egress router CNI plug-in and adding adding egress routers succeeds. ([BZ#1989688](#))

Node

- Previously, in containers, CRI-O did not create a symbolic link from `/proc/mounts` file to the `/etc/mtab` file. As a consequence, the user could not view the list of the mounted devices in the container's `/etc/mtab` file. CRI-O now adds the symbolic link. As a result, users can view the container's mounted devices. ([BZ#1868221](#))
- Previously, if pods were deleted quickly after creation, the kubelet might not clean up the pods properly. This resulted in pods being stuck in a terminating state, and could impact the availability of upgrades. This fix improves the pod lifecycle logic to avoid this problem. ([BZ#1952224](#))
- Previously, the **SystemMemoryExceedsReserved** alert would fire when the system memory usage exceeded 90% of the reserved memory. As a result, clusters could fire an excessive number of alerts. The threshold for this alarm was changed to fire at 95% of reserved memory. ([BZ#1980844](#))
- Previously, a bug in CRI-O caused CRI-O to leak a child PID of a process it created. As a result, if under load, `systemd` could create a significant number of zombie processes. This could lead to node failure if the node ran out of PIDs. CRI-O was fixed to prevent the leakage. As a result, these zombie processes are no longer being created. ([BZ#2003197](#))

OpenShift CLI (oc)

- Previously, the `oc` command-line tool was crashing while mirroring the registry, causing a **slice bounds out of range** panic runtime error because of an unchecked index operation on a slice when using the `--max-components` argument. With this fix, a check has been added to ensure that the components check does not request an out-of-range index value so that the `oc` tool no longer panics when using the `--max-components` argument. ([BZ#1786835](#))
- Previously, the `oc describe quota` command showed inconsistent units in **Used** memory for the **ClusterResourceQuota** value, which was unpredictable and difficult to read. With this fix, the

Used memory now always uses the same unit as the **Hard** memory so that the **oc describe quota** command shows predictable values. ([BZ#1955292](#))

- Previously, the **oc logs** command did not work with pipeline builds because of a missing client setup. The client setup has been fixed in the **oc logs** command so that it now works with pipeline builds. ([BZ#1973643](#))

Operator Lifecycle Manager (OLM)

- Previously, the Operator Lifecycle Manager (OLM) upgradeable condition message was unclear when installed Operators set **olm.maxOpenShiftVersion** to a minor OpenShift Container Platform version less than or equal to the current version. This resulted in an incorrect error message that has been fixed to specify that only minor and major version upgrades are blocked when **olm.maxOpenShiftVersion** is set to version different than the current OpenShift Container Platform version. ([BZ#1992677](#))
- Previously, the **opm** command failed to deprecate bundles when they were present in the index. Consequently, bundles truncated as part of another deprecation in the same call were reported as missing. This update adds a check for bundles before any deprecation takes place to differentiate between a bundle that is not present and one that has been truncated. As a result, deprecated bundles along the same upgrade path are no longer reported as missing. ([BZ#1950534](#))
- A transient error can occur when Operator Lifecycle Manager (OLM) attempts to update a custom resource definition (CRD) object in the cluster. This caused OLM to permanently fail the install plan containing the CRD. This bug fix updates OLM to retry CRD updates on resource-modified conflict errors. As a result, OLM is now more resilient to this class of transient errors. Install plans no longer permanently fail on conflict errors that OLM is able to retry and resolve. ([BZ#1923111](#))
- The **opm index|registry add** commands attempted to verify the existence of Operator bundles in an index that are replaced, regardless of whether they were already truncated from the index. The commands would consistently fail after a bundle was deprecated for a given package. This bug fix updates the **opm** CLI to handle this edge case and no longer verify the existence of truncated bundles. As a result, the commands no longer fail after a bundle is deprecated for a given package. ([BZ#1952101](#))
- Operator Lifecycle Manager (OLM) can now allow priority classes to be projected into registry pods using labels in catalog source resources. The default catalog sources are important components in namespaces managed by the cluster, which mandate priority classes. With this enhancement, all default catalog sources in the **openshift-marketplace** namespace have a **system-cluster-critical** priority class. ([BZ#1954869](#))
- The Marketplace Operator was using the leader-for-life implementation where a config map holding the leasing owner's identity has owner references placed by the controller's pod. This is problematic in the case where the node that the pod was scheduled on became unavailable, and the pod was unable to be terminated. This made the config map unable to be properly garbage collected so a new leader could be elected. Minor version cluster upgrades were blocked as the newer Marketplace Operator version could not gain leader election. Manual cleanup of the config map holding the leader election lease was required in order to release the lock and complete the upgrade of the Marketplace component. This bug fix switches to using the leader-for-lease leader election implementation. As a result, leader election no longer gets stuck in this scenario. ([BZ#1958888](#))
- Previously, a new **Failed** phase was introduced for install plans. Failure to detect a valid Operator group (OG) or service account (SA) resource for the namespace the install plan was being created in would transition the install plan to the failed state. That is, failure to detect

these resources when the install plan was reconciled the first time was considered a permanent failure. This was a regression from the following previous behavior of install plans:

- Failure to detect OG or SA resources would requeue the install plan for reconciliation.
- Creation of the required resources before the retry limit of the informer queue was reached would transition the install plan from the **Installing** phase to the **Complete** phase, unless the bundle unpacking step failed.

This regression introduced strange behavior for users who had infrastructure built that applied a set of manifests simultaneously to install an Operator that included a subscription, which creates install plans, along with the required OG and SA resources. In those cases, whenever there was a delay in the reconciliation of the OG and SA, the install plan would be transitioned to a state of permanent failure.

This bug fix removes the logic that transitioned the install plan to the **Failed** phase. Instead, the install plan is now requeued for any reconciliation error. As a result, when no OG is detected, the following condition is set:

```
conditions:
- lastTransitionTime: ""2021-06-23T18:16:00Z""
lastUpdateTime: ""2021-06-23T18:16:16Z""
message: attenuated service account query failed - no operator group found that
is managing this namespace
reason: InstallCheckFailed
status: ""False""
type: Installed
```

When a valid OG is created, the following condition is set:

```
conditions:
- lastTransitionTime: ""2021-06-23T18:33:37Z""
lastUpdateTime: ""2021-06-23T18:33:37Z""
status: ""True""
```

([BZ#1960455](#))

- When updating a catalog source, a **Get** call is immediately followed by a **Delete** call on a number of resources related to the catalog source. In some instances, the resource had already been deleted but the resource still existed in the cache. This allowed the **Get** call to succeed, but the following **Delete** call failed as the resource did not exist on the cluster. This bug fix updates Operator Lifecycle Manager (OLM) to ignore the error returned by the **Delete** call if the resource is not found. As a result, OLM no longer reports an error when updating a catalog source due to a caching issue that results in a "Resource Not Found" error from the **Delete** call. ([BZ#1967621](#))
- A cluster service version (CSV) with a name over the limit of 63 characters causes an invalid **ownerref** label. Previously, when Operator Lifecycle Manager (OLM) used the **ownerref** reference to retrieve owned resources, including cluster role bindings, the lister returned all cluster role bindings in the namespaces due to the invalid label. This bug fix updates OLM to use a different method to let the server reject invalid **ownerref** labels instead. As a result, when CSVs have an invalid name, OLM no longer removes cluster role bindings. ([BZ#1970910](#))
- Previously, Operator dependencies were not always persisted after installation time. After installing an Operator that declares dependencies, later updates and installations within the same namespace could fail to honor the previously installed Operator's dependencies. With this

bug fix, dependencies are now persisted, along with all declared properties for the Operator, in an annotation on the Operator's **ClusterServiceVersion** (CSV) object. As a result, the declared dependencies of installed Operators continue to be respected for future installations.

([BZ#1978310](#))

- Previously, when you removed an Operator with a deprecated bundle, the deprecation history was not included in the garbage collection. As a result, if you reinstalled the Operator, the bundle version would show up the deprecated table. This update fixes the issue with better garbage collection for deprecated bundles. ([BZ#1982781](#))
- Previously, the z-stream version of a cluster was used in Operator compatibility calculations. As a result, micro releases of OpenShift Container Platform were blocked. This update fixes the issue by ignoring cluster z-stream versions in Operator compatibility comparisons. ([BZ#1993286](#))

OpenShift API server

- Previously, a single failed request to the discovery endpoint of a service could make an Operator report **Available=False**. To increase resilience, a set of improvements have been introduced to prevent some Operators from reporting **Available=False** during an update due to various transient errors. ([BZ#1948089](#))

OpenShift Update Service

- Previously, when creating an update service application through the web console, an invalid host error occurred. This occurred because the default OpenShift Update Service (OSUS) application name was too long. A shorter default name is now in place and the error no longer occurs. ([BZ#1939788](#))

Red Hat Enterprise Linux CoreOS (RHCOS)

- Previously, systemd was unable to read environment files in **/etc/kubernetes**. The SELinux policy caused this and as a result, the kubelet did not start. The policy has been modified. The kubelet starts and the environment files are read. ([BZ#1969998](#))
- In an s390x Kernel Virtual Machine (KVM) with an ECKD DASD attached, the DASD would appear to be a regular virtio storage device but would become inaccessible if the VTOC was removed. As a result, you could not use DASD as a virtio block device when installing Red Hat Enterprise Linux CoreOS (RHCOS) on the KVM. The **coreos-installer** program has been updated so that it now installs Red Hat Enterprise Linux CoreOS (RHCOS) with a VTOC-format partition table when the installation destination is a virtio storage device such as an ECKD DASD attached to a KVM. ([BZ#1960485](#))
- Previously, the **NetworkManager-wait-online-service** timed out too early, which prevented establishing a connection before the **coreos-installer** program could start. Consequently, if the network took too long to start, the **coreos-installer** program failed to fetch the Ignition config. With this update, the **NetworkManager-wait-online-service** timeout has been increased to its default upstream value. As a result, the **coreos-installer** program no longer fails to fetch the Ignition config. ([BZ#1967483](#))

Routing

- Previously, there was config drift when the Cluster Network Operator (CNO) attempted to sanitize the proxy configuration, specifically the **no_proxy** config. This resulted in a specific IPv6 CIDR missing from the **no_proxy**. This fix implements logic that updates the dual stack (IPV4 and IPV6) for all scenarios. ([BZ#1981975](#))
- Previously, if the **.spec.privateZone** field of the **dns.config.openshift.io** Operator was filled

out incorrectly so that the Ingress Operator was not able to find the private hosted zone, then the Ingress Operator became degraded. However, even after fixing the **.spec.privateZone** field, the Ingress Operator stayed degraded. The Ingress Operator finds the hosted zone and adds the **.apps** resource record, but the Ingress Operator does not reset the degraded status. This fix watches the DNS config object and monitors changes regarding the **spec.privateZone** field. It applies the appropriate logic and updates the Operator status accordingly. The Operator status returns to degraded, or **False**, once the correct **.spec.privateZone** field is set. ([BZ#1942657](#))

Samples

- Previously, the lack of a connection timeout led to lengthy delays. This occurred when the Cluster Samples Operator, with **managementState** set to **Removed**, tested the connection to **registry.redhat.io**. With the addition of a connection timeout, the delay is eliminated. ([BZ#1990140](#))

Storage

- Previously, you could delete **LocalVolumeSets** with in-use PVs, which required manual clean up. This fix ensures that all released PVs are cleaned automatically. ([BZ#1862429](#))
- Previously, the **oc get volumesnapshotcontent** command did not display the namespace for a volume snapshot, which meant that the volume snapshot was not uniquely identified. This command now displays the namespace for the volume snapshot. ([BZ#1965263](#))
- Previously, the Manila CSI Operator used a custom transport when communicating with a Red Hat OpenStack Platform (RHOSP) endpoint that used self-signed certificates. Because this custom transport did not consume the proxy environment variables, the Manila CSI Operator would fail to communicate with Manila. This update ensures that the custom transport consumes the proxy environment variables. As a result, the Manila CSI Operator now works with proxy and custom CA certificates. ([BZ#1960152](#))
- Previously, the Cinder CSI Driver Operator was not using the configured proxy to connect to the Red Hat OpenStack Platform (RHOSP) API, which could cause the installation to fail. With this update, an annotation is included in the Cinder CSI Driver Operator deployment that ensures proxy environment variables are set on the container. As a result, the installation no longer fails. ([BZ#1985391](#))
- The frequency at which the Local Storage Operator inspects newly added block devices has been changed from 5 seconds to 60 seconds to decrease its CPU consumption. ([BZ#1994035](#))
- Previously, communication failure with the Manila CSI Operator would degrade the cluster. With this update, failed communication with the Manila CSI Operator endpoint results in a non-fatal error. As a result, the Manila CSI Operator gets disabled instead of degrading the cluster. ([BZ#2001958](#))
- Previously, the Local Storage Operator deleted orphaned persistent volumes (PVs) with a 10 second delay, and the delay was cumulative. When several persistent volume claims (PVCs) were deleted at the same time, it could take several minutes or hours to get their PVs deleted. Consequently, corresponding local disks were not available for new PVCs for several hours. With this fix, the 10 second delay is removed. As a result, PVs are detected and corresponding local disks are made available for new PVCs sooner. ([BZ#2007684](#))

Web console (Administrator perspective)

- Previously, all rows on the **PF4** table were re-rendering. With this update, the content in **React.memo** was wrapped so the content does not re-render on every scroll event. ([BZ#1856355](#))

- Previously, the charts in **Cluster Utilization** in the OpenShift Container Platform web console displayed the data time span in a confusing manner. For example, if the six-hour time span option is selected, but data exists only for the final three hours, those three data points are stretched to fill the entire chart. The first three hours are not displayed. This could result in the assumption that the chart is showing the full six-hour time span. To avoid confusion, the charts now show blank space for missing information. In this example, the chart displays the entire six-hour time span with data starting at the fourth hour. The first three hours are blank. ([BZ#1904155](#))
- Previously, **NetworkPolicy** was not translated to Korean or Chinese on the web console. With this update, **NetworkPolicy** is now translated correctly when viewing the web console in Korean or Chinese. ([BZ#1965930](#))
- Previously, an issue with the **Needs Attention** state of the **Console Overview** section showed Operators as **upgrading**, even if they were not upgrading. This update fixes the **Needs Attention** state so that the correct status of an Operator is displayed. ([BZ#1967047](#))
- Previously, the alert for a failed Cluster Service Version (CSV) displayed a generic **status.message** that did not help troubleshoot the failed CSV. With this update, copied CSVs show a helpful message and a link to the original CSV to troubleshoot. ([BZ#1967658](#))
- Previously, a user was unable to use the drop-down options in the masthead with a keyboard. With this update, users are now able to access the drop-down options with a keyboard. ([BZ#1967979](#))
- Previously, a utility function used for matching Operator-owned resources with their owners would return false matches. Consequently, **Managed by** links on Operator-owned resource pages would sometimes link to the incorrect URL. This fix updates the function logic to correctly match owned Operators. As a result, **Managed by** links now link to the correct URL. ([BZ#1970011](#))
- Previously, the **OperatorHub** web console interface would lead users to unrelated install plans. With this update, **OperatorHub** links users to the Operator **Subscription** details tab to view installation progress. ([BZ#1970466](#))
- Previously, items in the **Add** drop-down list on the **OAuth details** page were not internationalized. With this update, these items are internationalized and the user experience for non-English speakers is improved. ([BZ#1970604](#))
- Previously, an invalid localization property prevented some messages from being internationalized. This update removes the invalid property. As a result, these messages are internationalized and the user experience for non-English speakers is improved. ([BZ#1970980](#))
- This update removes a tooltip that appeared when mousing over a resource link on a list page, because the information did not improve the user experience. ([BZ#1971532](#))
- Previously, console pods were deployed with the **preferredDuringSchedulingIgnoredDuringExecution** anti-affinity rule, which sometimes resulted in both console pods being scheduled on the same control plane node. This fix changes the rule to **requiredDuringSchedulingIgnoredDuringExecution** so that the pods must be scheduled on different nodes if the condition matches. ([BZ#1975379](#))
- Previously, uninstalling an Operator failed to remove all of the enabled plugins. With this release, uninstalling an Operator now removes all enabled plugins. ([BZ#1975820](#))
- Previously, front-end Operator Lifecycle Manager (OLM) descriptor handling only used the first

x-descriptor to render a property on an operand details page. Consequently, if multiple x-descriptors were defined for a property and the first one in the list was invalid or unsupported, it would not render as expected. This fix updates the descriptor validation logic to prioritize supported x-descriptors over unsupported x-descriptors. As a result, descriptor-decorated properties are rendered on the **Operand details** page using the first valid and supported x-descriptor found in the list. ([BZ#1976072](#))

- Previously, string data was used for encoded secrets. As a result, binary secret data was not properly uploaded by the web console. This update encodes secrets and uses data instead of string data in the API. As a result, binary secrets now upload correctly. ([BZ#1978724](#))
- Previously, when processes running on the cluster were manually terminated, the terminal **ps - aux** command showed that some processes were not cleared. This caused stray processes to remain, leaving the cluster in an invalid state. This fix ensures that all processes terminate properly on the cluster and do not appear in the list of active processes that are listed on the terminal. ([BZ#1979571](#))
- Previously, when a default pull secret was added to a new project and credentials for multiple registries were uploaded, only the first credential was listed on the **Project Details** page. There was also no indication that the list was truncated. As a result, when a user clicked the project details from **Default pull secret**, only the first credential was listed. This fix ensures that all of the credentials are listed and informs the users that additional credentials exist if they are not listed on the current page. ([BZ#1980704](#))
- Previously, when users changed the default browser language to Simplified Chinese, the cluster utilization resource metrics on the **Overview** page of the web console displayed in a combination of English and Simplified Chinese characters. This fix allows the user to view the cluster utilization resources entirely in the selected language. ([BZ#1982079](#))
- Previously, when the language was changed to Simplified Chinese, the cluster utilization usage statistic did not match the translation in the left menu for **project**, **pod**, and **node**. This fix corrects the Simplified Chinese translation so that the cluster utilization metrics are consistent with the **top consumers** filter. ([BZ#1982090](#))
- Previously, users saw an error instead of the default pull secret from the service account. This resulted in incomplete information on the project details screen. The user had to go to the default ServiceAccount to view the entire list of default pull secrets. This fix allows the user to view the entire list of pull secrets from the default ServiceAccount on the project details page. ([BZ#1983091](#))
- Previously, if you resized the web page for a node or pod while viewing the **Terminal** tab, sometimes the browser displayed two vertical scrollbars. The console is now updated to display one scrollbar only when the window is resized. ([BZ#1983220](#))
- Previously, the web console did not deploy when installing OpenShift Container Platform 4.8.2 using a single node developer profile. If a valid Operator group or service account was not detected for the namespace in which the install plan was being created, the install plan was placed in a failed state. No further attempts were made. With this revision, the failed install plan is set to run again until an Operator group or service account is detected. ([BZ#1986129](#))
- Previously, in the **Events Dashboard**, **More** and **Show Less** were not internationalized, which resulted in poor user experience. With this update, the text is now internationalized. ([BZ#1986754](#))
- Previously, the logic that constructed the fully qualified domain name (FQDN) of a service in the **Console** page was missing. As a result, FQDN information was missing on the service's detail page. This update adds logic that constructs the FQDN so that the service's FQDN information

is now available on the page. ([BZ#1996816](#))

Web console (Developer perspective)

- Previously, kamelets of type **sink** were shown in the catalog for event sources along with source kamelets. In the current release, the catalog for event sources displays only kamelets of type **source**. ([BZ#1971544](#))
- Previously, the log file contained information in a single line without any line breaks. In the current release, the log file contains the expected line breaks with additional line breaks around log headers. ([BZ#1985080](#))

1.7. TECHNOLOGY PREVIEW FEATURES

Some features in this release are currently in Technology Preview. These experimental features are not intended for production use. Note the following scope of support on the Red Hat Customer Portal for these features:

[Technology Preview Features Support Scope](#)

In the table below, features are marked with the following statuses:

- **TP:** *Technology Preview*
- **GA:** *General Availability*
- **-:** *Not Available*
- **DEP:** *Deprecated*

Table 1.3. Technology Preview tracker

Feature	OCP 4.7	OCP 4.8	OCP 4.9
Precision Time Protocol (PTP)	TP	TP	TP
oc CLI plug-ins	TP	GA	GA
Descheduler	GA	GA	GA
HPA for memory utilization	GA	GA	GA
Service Binding	TP	TP	TP
Raw Block with Cinder	TP	GA	GA
CSI volume snapshots	GA	GA	GA
CSI volume expansion	TP	TP	TP
vSphere Problem Detector Operator	GA	GA	GA

Feature	OCP 4.7	OCP 4.8	OCP 4.9
CSI Azure Disk Driver Operator	-	TP	TP
CSI Azure Stack Hub Driver Operator	-	-	GA
CSI GCP PD Driver Operator	TP	GA	GA
CSI OpenStack Cinder Driver Operator	TP	TP	TP
CSI AWS EBS Driver Operator	TP	TP	GA
CSI AWS EFS Driver Operator	-	-	TP
CSI automatic migration	-	TP	TP
CSI inline ephemeral volumes	TP	TP	TP
CSI vSphere Driver Operator	-	TP	TP
Automatic device discovery and provisioning with Local Storage Operator	TP	TP	TP
OpenShift Pipelines	TP	GA	GA
OpenShift GitOps	TP	GA	GA
OpenShift sandboxed containers	-	TP	TP
Vertical Pod Autoscaler	TP	GA	GA
Cron jobs	TP	GA	GA
PodDisruptionBudget	TP	GA	GA
Adding kernel modules to nodes with kvc	TP	TP	TP
Egress router CNI plug-in	TP	GA	GA
Scheduler profiles	TP	TP	GA
Non-preempting priority classes	TP	TP	TP
Kubernetes NMState Operator	TP	TP	TP
Assisted Installer	TP	TP	TP

Feature	OCP 4.7	OCP 4.8	OCP 4.9
AWS Security Token Service (STS)	TP	GA	GA
Kdump	TP	TP	TP
OpenShift Serverless	-	GA	GA
Serverless functions	-	TP	TP
Data Plane Development Kit (DPDK) support	TP	TP	GA
Memory Manager feature	-	-	TP
CNI VRF plug-in	TP	TP	GA
Cluster Cloud Controller Manager Operator	-	-	GA
Cloud controller manager for AWS	-	-	TP
Cloud controller manager for Azure	-	-	TP
Cloud controller manager for OpenStack	-	-	TP
Driver Toolkit	-	TP	TP
Special Resource Operator (SRO)	-	-	TP
Node Health Check Operator	-	-	TP

1.8. KNOWN ISSUES

- In OpenShift Container Platform 4.1, anonymous users could access discovery endpoints. Later releases revoked this access to reduce the possible attack surface for security exploits because some discovery endpoints are forwarded to aggregated API servers. However, unauthenticated access is preserved in upgraded clusters so that existing use cases are not broken. If you are a cluster administrator for a cluster that has been upgraded from OpenShift Container Platform 4.1 to 4.8, you can either revoke or continue to allow unauthenticated access. It is recommended to revoke unauthenticated access unless there is a specific need for it. If you do continue to allow unauthenticated access, be aware of the increased risks.

**WARNING**

If you have applications that rely on unauthenticated access, they might receive HTTP **403** errors if you revoke unauthenticated access.

Use the following script to revoke unauthenticated access to discovery endpoints:

```
## Snippet to remove unauthenticated group from all the cluster role bindings
$ for clusterrolebinding in cluster-status-binding discovery system:basic-user
system:discovery system:openshift:discovery ;
do
### Find the index of unauthenticated group in list of subjects
index=$(oc get clusterrolebinding ${clusterrolebinding} -o json | jq 'select(.subjects!=null) |
.subjects | map(.name=="system:unauthenticated") | index(true)');
### Remove the element at index from subjects array
oc patch clusterrolebinding ${clusterrolebinding} --type=json --patch "[{'op': 'remove','path':
'/subjects/${index}'}]";
done
```

This script removes unauthenticated subjects from the following cluster role bindings:

- **cluster-status-binding**
- **discovery**
- **system:basic-user**
- **system:discovery**
- **system:openshift:discovery**

([BZ#1821771](#))

- The **oc annotate** command does not work for LDAP group names that contain an equal sign (=), because the command uses the equal sign as a delimiter between the annotation name and value. As a workaround, use **oc patch** or **oc edit** to add the annotation. ([BZ#1917280](#))
- Cluster administrators can specify a custom HTTP error code response page for either 503, 404, or both error pages. If you do not specify the correct format for the custom error code response page, a router pod outage occurs and does not resolve. The router does not reload to reflect custom error code page updates. As a workaround, you can use the **oc rsh** command to locally access the router pods. Then run **reload-haproxy** in all the router pods that serve the custom http error code pages:

```
$ oc -n openshift-ingress rsh router-default-6647d984d8-q7b58
sh-4.4$ bash -x /var/lib/haproxy/reload-haproxy
```

Alternatively, you can annotate the route to force a reload. ([BZ#1990020](#)), ([BZ#2003961](#))

- An Open Virtual Network (OVN) bug causes persistent connectivity issues with Octavia load balancers. When Octavia load balancers are created, OVN might not plug them into some

Neutron subnets. These load balancers might be unreachable for some of the Neutron subnets. This problem affects Neutron subnets, which are created for each OpenShift namespace at random when Kuryr is configured. As a result, when this problem occurs the load balancer that implements OpenShift **Service** objects will be unreachable from OpenShift namespaces affected by the issue. Because of this bug, OpenShift Container Platform 4.8 deployments that use Kuryr SDN are not recommended on Red Hat OpenStack Platform (RHOSP) 16.1 with OVN and OVN Octavia configured. This will be fixed in a future release of RHOSP. ([BZ#1937392](#))

- Installations on Red Hat OpenStack Platform (RHOSP) with Kuryr will not work if configured with a cluster-wide proxy when the cluster-wide proxy is required to access RHOSP APIs. ([BZ#1985486](#))
- Due to a race condition, the Red Hat OpenStack Platform (RHOSP) cloud provider might not start properly. Consequently, LoadBalancer services might never get an **EXTERNAL-IP** set. As a temporary workaround, you can restart the kube-controller-manager pods using the procedure described in ([BZ#2004542](#)).
- The **ap-northeast-3** AWS region is not provided as an option by the installation program when installing OpenShift Container Platform, even though it is a supported AWS region. As a temporary workaround, you can select a different AWS region from the installation prompt and then update the region information in the generated **install-config.yaml** file before installing your cluster. ([BZ#1996544](#))
- When installing a cluster on AWS in the **us-east-1** region, you cannot use local AWS zones. As a temporary workaround, you must specify non-local availability zones in the **install-config.yaml** file when installing a cluster. ([BZ#1997059](#))
- You can only install OpenShift Container Platform on Azure Stack Hub with public endpoints, such as the ARM endpoint, that are secured with certificates signed by a publicly trusted certificate authority (CA). Support for internal CAs will be added in a future z-stream release of OpenShift Container Platform. ([BZ#2012173](#))
- Cluster administrators can specify a custom HTTP error code response page for either 503, 404, or both error pages. The router does not reload to reflect custom error code pages updates. As a workaround, rsh in the router pods and run **reload-haproxy** in all the router pods that serve the custom http error code pages:

```
$ oc -n openshift-ingress rsh router-default-6647d984d8-q7b58  
sh-4.4$ bash -x /var/lib/haproxy/reload-haproxy
```

Alternatively, you can annotate the route to force a reload. ([BZ#1990020](#))

- This release contains a known issue. If you customize the hostname and certificate of the OpenShift OAuth route, Jenkins no longer trusts the OAuth server endpoint. As a result, users cannot log in to the Jenkins console if they rely on the OpenShift OAuth integration to manage identity and access. A workaround is not available at this time. ([BZ#1991448](#))
- Because certain high cardinality monitoring metrics were inadvertently dropped, the following container performance input and output metrics are not available in this release: **pod**, **qos**, and **System**.
No workaround exists for this issue. To track these metrics for production workloads, do not upgrade to the initial 4.9 release. ([BZ#2008120](#))
- The Special Resource Operator (SRO) might fail to install on Google Cloud Platform due to a software-defined network policy. As a result, the simple-kmod pod is not created. This is fixed in OpenShift Container Platform [4.9.4](#) release. ([BZ#1996916](#))

- In OpenShift Container Platform 4.9, a user with cluster role cannot scale a deployment or deployment config using the console if they do not have edit rights to the deployment or deployment config. ([BZ#1886888](#))
- In OpenShift Container Platform 4.9, when minimal or no data exists in the **Developer Console**, most of the monitoring charts or graphs (for example, CPU consumption, memory usage, and bandwidth) show a range of -1 to 1. However, none of these values can ever go below zero, so the negative values are incorrect. ([BZ#1904106](#))
- The **ip vrf exec** command does not work due to a **cgroups** mismatch. As a result, this command cannot be used inside OpenShift pods. To use virtual routing and forwarding (VRF), applications must be VRF-aware and bind directly to the VRF interface. ([BZ#1995631](#))
- A nonuniform memory access (NUMA) bug can cause undesired NUMA pinning for the container, which can lead to latency or performance degradation. The Topology Manager can pin the container, with resources that the **single-numa-node** topology management policy can satisfy, to more than one NUMA node. The container is pinned under the guaranteed Quality of Service (QoS) pod. As a workaround, do not start guaranteed QoS pods when the container memory-resource requests are bigger than the **single-numa-node** policy can provide. ([BZ#1999603](#))
- Occasionally, a pod selected for deletion is not deleted. This occurs when a cluster is running out of resources. To reclaim resources, the system selects one or more pods for deletion. With low resources causing slow processing, the deletion operation may exceed the established grace period for deletion, resulting in failure. If this occurs, manually delete the pod. The cluster then reclaims the freed up resources. ([BZ#1997476](#))
- Intermittently, pods can hang in the **ContainerCreating** state and time out while waiting for Open vSwitch (OVS) port binding. The reported event is **failed to configure pod interface: timed out waiting for OVS port binding**. This issue can occur when many pods are created for the OVN-Kubernetes plug-in. ([BZ#2005598](#))
- After rebooting the egress node, the **lr-policy-list** contains errors, such as duplicate records or missed internal IP addresses. The expected result is that the **lr-policy-list** has the same records as before rebooting the egress node. As a workaround, you can restart the **ovn-kubemaster** pods. ([BZ#1995887](#))
- When IP multicast relay is enabled on a logical router that contains distributed gateway ports, multicast traffic is not forwarded correctly on the distributed gateway port. The result is broken IP multicast functionality in OVN-Kubernetes. ([BZ#2010374](#))
- In the **Administrator** perspective of the web console, a page that is supposed to display a list of nodes is rendered before the list of nodes is available, which causes the page to become unresponsive. There is no workaround, but this issue will be addressed in a future release. ([BZ#2013088](#))
- Operator Lifecycle Manager (OLM) uses a combination of timestamp checks and obsolete API calls, which do not work for **skipRange** upgrades, to determine if it is necessary to perform an upgrade for a particular subscription. For Operators that use the **skipRange** upgrade, there is a delay in the upgrade process that can take up to 15 minutes to resolve and can potentially be blocked for even longer.
As a workaround, cluster administrators can delete the **catalog-operator** pod in the **openshift-operator-lifecycle-manager** namespace. This causes the pod to be automatically recreated, which causes the **skipRange** upgrade to trigger. ([BZ#2002276](#))
- Currently, when launching Red Hat Enterprise Linux (RHEL) 8 on Google Cloud Platform with FIPS mode enabled, RHEL 8 fails to download metadata when trying to install packages from

the Red Hat Update Infrastructure (RHUI). As a temporary workaround, you can disable RHUI repositories and use Red Hat Subscription Management to get content. ([BZ#2001464](#)), ([BZ#1997516](#)).

- Following an OpenShift Container Platform single node reboot, all pods are restarted which causes significant load and longer than normal pod creation times. This happens because the Container Network Interface (CNI) is not able to process the **pod add** events quickly enough. The following error message is displayed: **timed out waiting for OVS port binding**. The OpenShift Container Platform single node instance eventually recovers, just slower than expected. ([BZ#1986216](#))
- When MetalLB is run in layer 2 mode with the OVN-Kubernetes Container Network Interface network provider, instead of a single node with a speaker pod responding to an ARP or NDP request, every node in the cluster responds to the request. The unexpected number of ARP responses might resemble an ARP-spoofing attack. Although the experience is different than designed, traffic is routed to the service as long as no software on the hosts or subnet is configured to block ARP. This bug is fixed in a future OpenShift Container Platform release. ([BZ#1987445](#))
- When Tang disk encryption and a static IP address configuration are applied on a VMWare vSphere user-provisioned infrastructure cluster, the nodes fail to boot properly after they are initially provisioned. ([BZ#1975701](#))
- Operators must list any related images for Operator Lifecycle Manager (OLM) to run from a local source. Presently, if the **relatedImages** parameter of the **ClusterServiceVersion** (CSV) object is not defined, **opm render** does not populate the related images. This is planned to be fixed in a later release. ([BZ#2000379](#))
- Previously, Open vSwitch (OVS) ran in a container on each OpenShift Container Platform cluster node and the node exporter agent collected OVS CPU and memory metrics from the nodes. Now, OVS runs on the cluster nodes as systemd units and the metrics are not collected. This is planned to be fixed in a later release. OVS packet metrics are still collected. ([BZ#2002868](#))
- The flag that is used to hide or show the **Storage → Overview** page in the OpenShift Container Platform web console is misconfigured. As a result, the overview page is not visible after deploying a cluster that includes OpenShift Cluster Storage. This is planned to be fixed in a later release. ([BZ#2013132](#))
- In OpenShift Container Platform 4.6 and later, image references for a pull must specify the following Red Hat registries:
 - **registry.redhat.io**
 - **registry.access.redhat.com**
 - **quay.io**

Otherwise, if those registries are not specified, the build pods cannot pull the images.

As a workaround, use fully qualified names, such as **registry.redhat.io/ubi8/ubi:latest** and **registry.access.redhat.com/rhel7.7:latest**, in image pull specifications.

Optionally, you can update the image registry settings by [adding registries that allow image short names](#). ([BZ#2011293](#))

- Prior to OpenShift Container Platform 4.8 the default load balancing algorithm was **leastconn**.

The default was changed to **random** in OpenShift Container Platform 4.8.0 for non-passthrough routes. Switching to **random** is incompatible with environments that need to use long-running websocket connections because it significantly increases memory consumption in those environments. To mitigate this significant memory consumption, the default load balancing algorithm was reverted to **leastconn** in OpenShift Container Platform 4.9. Once there is a solution that does not incur significant memory usage, the default will be changed to **random** in a future OpenShift Container Platform release.

You can check the default setting by entering the following command:

```
$ oc get deployment -n openshift-ingress router-default -o yaml | grep -A 2
ROUTER_LOAD_BALANCE_ALGORITHM
  - name: ROUTER_LOAD_BALANCE_ALGORITHM
    value: leastconn
```

The **random** option is still available. However routes that would benefit from this algorithmic choice must explicitly set that option in an annotation on a per-route basis by entering the following command:

```
$ oc annotate -n <NAMESPACE> route/<ROUTE-NAME>
"haproxy.router.openshift.io/balance=random"
```

([BZ#2015829](#))

1.9. ASYNCHRONOUS ERRATA UPDATES

Security, bug fix, and enhancement updates for OpenShift Container Platform 4.9 are released as asynchronous errata through the Red Hat Network. All OpenShift Container Platform 4.9 errata is [available on the Red Hat Customer Portal](#). See the [OpenShift Container Platform Life Cycle](#) for more information about asynchronous errata.

Red Hat Customer Portal users can enable errata notifications in the account settings for Red Hat Subscription Management (RHSM). When errata notifications are enabled, users are notified via email whenever new errata relevant to their registered systems are released.



NOTE

Red Hat Customer Portal user accounts must have systems registered and consuming OpenShift Container Platform entitlements for OpenShift Container Platform errata notification emails to generate.

This section will continue to be updated over time to provide notes on enhancements and bug fixes for future asynchronous errata releases of OpenShift Container Platform 4.9. Versioned asynchronous releases, for example with the form OpenShift Container Platform 4.9.z, will be detailed in subsections. In addition, releases in which the errata text cannot fit in the space provided by the advisory will be detailed in subsections that follow.



IMPORTANT

For any OpenShift Container Platform release, always review the instructions on [updating your cluster](#) properly.

1.9.1. RHSA-2021:3759 - OpenShift Container Platform 4.9.0 image release, bug fix, and security update advisory

Issued: 2021-10-18

OpenShift Container Platform release 4.9.0, which includes security updates, is now available. The list of bug fixes that are included in the update is documented in the [RHSA-2021:3759](#) advisory. The RPM packages that are included in the update are provided by the [RHSA-2021:3758](#) advisory.

Space precluded documenting all of the container images for this release in the advisory. See the following article for notes on the container images in this release:

[OpenShift Container Platform 4.9.0 container image list](#)

1.9.2. RHBA-2021:3935 - OpenShift Container Platform 4.9.4 bug fix and security update

Issued: 2021-10-26

OpenShift Container Platform release 4.9.4 is now available. The bug fixes that are included in the update are listed in the [RHBA-2021:3935](#) advisory. The RPM packages that are included in the update are provided by the [RHSA-2021:3934](#) advisory.

Space precluded documenting all of the container images for this release in the advisory. See the following article for notes on the container images in this release:

[OpenShift Container Platform 4.9.4 container image list](#)

1.9.2.1. Enhancements

A new conditional gatherer has been implemented for the **SamplesImageStreamImportFailing** alert, which collects logs and image streams of the **openshift-cluster-samples-operator** namespace when fired. The additional data gathering allows for more insight into problems when pulling image streams from an external registry. ([BZ#1966153](#))

1.9.2.2. Bug fixes

- Previously, the **Nodes** page rendered before the list of nodes became available. With this update, the **Nodes** page renders correctly when the node list is available. ([BZ#2013088](#))

1.9.2.3. Upgrading

To upgrade an existing OpenShift Container Platform 4.9 cluster to this latest release, see [Updating a cluster by using the CLI](#) for instructions.

1.9.3. RHBA-2021:4005 - OpenShift Container Platform 4.9.5 bug fix update

Issued: 2021-11-01

OpenShift Container Platform release 4.9.5 is now available. The bug fixes that are included in the update are listed in the [RHBA-2021:4005](#) advisory. The RPM packages that are included in the update are provided by the [RHSA-2021:4004](#) advisory.

Space precluded documenting all of the container images for this release in the advisory. See the following article for notes on the container images in this release:

[OpenShift Container Platform 4.9.5 container image list](#)

1.9.3.1. Known Issues

- The flag that is used to hide or show the **Storage → Overview** page in the OpenShift Container Platform web console is misconfigured. Consequently, the **Overview** page is invisible after deploying a cluster that included OpenShift Cluster Storage. A fix for this bug is planned for a future release. ([BZ#2013132](#))

1.9.3.2. Bug fixes

- With the deprecation of the **lastTriggeredImageID** field for build configs, the image change trigger controller stopped checking the ID field prior to initiating builds. Consequently, if a build config was created and had an image change trigger start while the cluster was running OpenShift Container Platform 4.7 or earlier, it continuously tried to trigger builds. With this update, these unnecessary attempts to trigger builds no longer occur. ([BZ#2006793](#))

1.9.3.3. Upgrading

To upgrade an existing OpenShift Container Platform 4.9 cluster to this latest release, see [Updating a cluster by using the CLI](#) for instructions.

1.9.4. RHBA-2021:4119 - OpenShift Container Platform 4.9.6 bug fix and security update

Issued: 2021-11-10

OpenShift Container Platform release 4.9.6, which includes security updates, is now available. The bug fixes that are included in the update are listed in the [RHBA-2021:4119](#) advisory. The RPM packages that are included in the update are provided by the [RHSA-2021:4118](#) advisory.

Space precluded documenting all of the container images for this release in the advisory. See the following article for notes on the container images in this release:

[OpenShift Container Platform 4.9.6 container image list](#)

1.9.4.1. Known Issues

- The current opt-in obfuscation will not work on clusters with OVN because the **hostsubnets.network.openshift.io** is not currently on OVN clusters. ([BZ#2014633](#))

1.9.4.2. Bug fixes

- Previously, a bug in the lock implementation for the **nmstate-handler** pod caused multiple nodes to gain control. This update fixes the lock implementation so that only one node is in control of the lock. ([BZ#1954309](#))
- Previously, OpenStack flavor validation accepted flavors not meeting the RAM requirements using the wrong unit. With this update, the correct unit is used for comparing minimum RAM against value returned by OpenStack. ([BZ#2009787](#))
- Previously, OpenShift Container Platform deployments on OpenStack failed for compact clusters with undedicated workers due to control plane nodes missing Ingress security group rules. With this update, an Ingress security group was added to OpenStack when control planes are schedulable. ([BZ#2016267](#))

- Previously, some **cAdvisor** metrics were dropped in order to reduce overall memory consumption but the **Utilization** dashboard in the console did not display any results. With this update, the dashboards display correctly again. ([BZ#2018455](#))

1.9.4.3. Upgrading

To upgrade an existing OpenShift Container Platform 4.9 cluster to this latest release, see [Updating a cluster by using the CLI](#) for instructions.

1.9.5. RHBA-2021:4579 - OpenShift Container Platform 4.9.7 bug fix update

Issued: 2021-11-15

OpenShift Container Platform release 4.9.7 is now available. The bug fixes that are included in the update are listed in the [RHBA-2021:4579](#) advisory. There are no RPM packages for this release.

Space precluded documenting all of the container images for this release in the advisory. See the following article for notes on the container images in this release:

[OpenShift Container Platform 4.9.7 container image list](#)

1.9.5.1. Features

1.9.5.1.1. Updates from Kubernetes 1.22.2

This update contains changes from Kubernetes 1.22.2. More information can be found in the following changelog: [1.22.2](#).

1.9.5.2. Upgrading

To upgrade an existing OpenShift Container Platform 4.9 cluster to this latest release, see [Updating a cluster within a minor version by using the CLI](#) for instructions.

1.9.6. RHBA-2021:4712 - OpenShift Container Platform 4.9.8 bug fix update

Issued: 2021-11-22

OpenShift Container Platform release 4.9.8 is now available. The bug fixes that are included in the update are listed in the [RHBA-2021:4712](#) advisory. The RPM packages that are included in the update are provided by the [RHBA-2021:4711](#) advisory.

Space precluded documenting all of the container images for this release in the advisory. See the following article for notes on the container images in this release:

[OpenShift Container Platform 4.9.8 container image list](#)

1.9.6.1. Bug fixes

- Previously, if you added or deleted the **SriovNetworkNodePolicy** custom resource (CR) while any of the **SriovNetworkNodeState** CRs had a **syncStatus** object with a value other than **Succeeded**, the SR-IOV network configuration daemon pod would cordon the node and mark it as unschedulable. This update fixes the problem. ([BZ#2002508](#))

1.9.6.2. Upgrading

To upgrade an existing OpenShift Container Platform 4.9 cluster to this latest release, see [Updating a cluster within a minor version by using the CLI](#) for instructions.

CHAPTER 2. OPENSIFT CONTAINER PLATFORM VERSIONING POLICY

OpenShift Container Platform provides strict backwards compatibility guarantees for all supported APIs, excluding alpha APIs (which may be changed without notice) and beta APIs (which may occasionally be changed in a non-backwards compatible manner).

Red Hat did not publicly release OpenShift Container Platform 4.0 and, instead, released OpenShift Container Platform 4.1 directly after version 3.11.

The OpenShift Container Platform version must match between master and node hosts, excluding temporary mismatches during cluster upgrades. For example, in a 4.9 cluster, all masters must be 4.9 and all nodes must be 4.9. If you installed an earlier version of **oc**, you cannot use it to complete all of the commands in OpenShift Container Platform 4.9. You must download and install the new version of **oc**.

Changes of APIs for non-security related reasons will involve, at minimum, two minor releases (4.1 to 4.2 to 4.3, for example) to allow older **oc** to update. Using new capabilities may require newer **oc**. A 4.3 server may have additional capabilities that a 4.2 **oc** cannot use and a 4.3 **oc** may have additional capabilities that are not supported by a 4.2 server.

Table 2.1. Compatibility Matrix

	X.Y (oc Client)	X.Y+N footnote:versionpolicyn[Where N is a number greater than 1.] (oc Client)
X.Y (Server)	1	3
X.Y+N footnote:versionpolicyn[] (Server)	2	1

- 1** Fully compatible.
- 2** **oc** client may not be able to access server features.
- 3** **oc** client may provide options and features that may not be compatible with the accessed server.