



# OpenShift Container Platform 4.11

## Release notes

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



## OpenShift Container Platform 4.11 Release notes

---

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

## Legal Notice

Copyright © 2023 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<sup>®</sup> is the registered trademark of Linus Torvalds in the United States and other countries.

Java<sup>®</sup> is a registered trademark of Oracle and/or its affiliates.

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

MySQL<sup>®</sup> is a registered trademark of MySQL AB in the United States, the European Union and other countries.

Node.js<sup>®</sup> 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<sup>®</sup> 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

<b>CHAPTER 1. OPENSIFT CONTAINER PLATFORM 4.11 RELEASE NOTES</b>	<b>9</b>
1.1. ABOUT THIS RELEASE	9
1.2. OPENSIFT CONTAINER PLATFORM LAYERED AND DEPENDENT COMPONENT SUPPORT AND COMPATIBILITY	9
1.3. NEW FEATURES AND ENHANCEMENTS	9
1.3.1. Red Hat Enterprise Linux CoreOS (RHCOS)	9
1.3.1.1. Improved support for NVMe over Fabrics	9
1.3.1.2. Investigate kernel crashes on AMD64 machines with kdump	9
1.3.1.3. Investigate kernel crashes on ARM64 machines with kdump (Technology Preview)	10
1.3.1.4. RHCOS now uses RHEL 8.6	10
1.3.1.5. Updated RHCOS registry URL	10
1.3.2. Installation and upgrade	10
1.3.2.1. RHEL 9 support for the OpenShift installer	10
1.3.2.2. New minimum system requirements for installing OpenShift Container Platform on a single node	10
1.3.2.3. OpenShift Container Platform on ARM	10
1.3.2.4. Troubleshooting bootstrap failures during installation on AWS	11
1.3.2.5. Support for Microsoft Hyper-V generation version 2	11
1.3.2.6. Default AWS and VMware vSphere compute node resources	11
1.3.2.7. Support for the AWS SC2S region	11
1.3.2.8. Installing a cluster on Nutanix using installer-provisioned infrastructure	11
1.3.2.9. Installing OpenShift Container Platform using Azure Ultra SSD	11
1.3.2.10. Added support for bootstrapExternalStaticIP and bootstrapExternalStaticGateway configuration settings	11
1.3.2.11. Configuring Fujitsu hardware	12
1.3.2.12. Disconnected mirroring with the oc-mirror CLI plugin is now generally available	12
1.3.2.13. Installing a cluster on Azure using user-managed encryption keys	12
1.3.2.14. Accelerated Networking for Azure enabled by default	12
1.3.2.15. AWS VPC endpoints and restricted installations	13
1.3.2.16. Additional customization when installing OpenShift Container Platform	13
1.3.2.17. Azure Marketplace offering	13
1.3.2.18. AWS Marketplace offering	13
1.3.2.19. CSI driver installation on vSphere clusters	13
1.3.3. Post-installation configuration	13
1.3.3.1. Cluster capabilities	14
1.3.3.2. OpenShift Container Platform clusters with multi-architecture compute machines (Technology Preview)	14
1.3.4. Web console	14
1.3.4.1. Developer Perspective	14
1.3.4.2. Dynamic plugin updates	14
1.3.4.3. Support for dark mode theme	15
1.3.4.4. Display operand instances for all managed namespaces on the Installed Operator page	15
1.3.4.5. Conditional updates	15
1.3.4.6. Pod disruption budgets (PDBs)	15
1.3.5. OpenShift CLI (oc)	15
1.3.5.1. RHEL 9 support for the OpenShift CLI (oc)	15
1.3.6. IBM Z and LinuxONE	15
Notable enhancements	16
Supported features	16
Restrictions	17
1.3.7. IBM Power	18
Notable enhancements	18

Supported features	18
Restrictions	19
1.3.8. Security and compliance	20
1.3.8.1. Audit logs now include OAuth server audit events	20
1.3.9. Networking	20
1.3.9.1. Pod-level bonding for secondary networks	20
1.3.9.2. New option for Ingress Controllers with the hostnetwork endpoint	20
1.3.9.3. Multi-node configuration for control plane and worker nodes	20
1.3.9.4. Support for configuring Classic Load Balancer Timeouts on AWS	20
1.3.9.5. Update to HAProxy 2.2.24	21
1.3.9.6. Support for configuring maximum number of connections for HAProxy processes	21
1.3.9.7. Setting the Ingress Controller health check interval	21
1.3.9.8. Support for configuring interface-level safe network sysctls	21
1.3.9.9. Support for CoreDNS forwarding DNS requests over TLS	21
1.3.9.10. Internal traffic support for OVN-Kubernetes	21
1.3.9.11. Support for AWS Load Balancer Operator (Technology Preview)	22
1.3.9.12. Enhancement to the Route API	22
1.3.9.13. External DNS Operator	22
1.3.9.14. PTP support for dual NIC boundary clocks	22
1.3.9.15. PTP events enhancements	22
1.3.9.16. SR-IOV support for Pensando DSC cards	23
1.3.9.17. SR-IOV support for Mellanox MT2892 cards	23
1.3.9.18. OpenShift Container Platform CIDR Ranges for Networks	23
1.3.9.19. OVN-Kubernetes network provider: Enable IPsec at runtime	23
1.3.9.20. Support for additional MetalLB CRDs and control over logging verbosity	23
1.3.9.21. Ability to create a route using the destination CA certificate in the Ingress annotation	24
1.3.9.22. Hosted control planes (Technology Preview)	24
1.3.9.23. IPv6 single and dual-stack support on user-provisioned bare metal infrastructure with the OVN-Kubernetes cluster network provider	24
1.3.9.24. OVS hardware offloading on RHOSP	24
1.3.9.25. NFV user experience improvements on RHOSP	24
1.3.9.26. Installations on Red Hat OpenStack Platform, VMware vSphere, or oVirt now configure keepalived with unicast as the default	25
1.3.9.27. Network Observability Operator to observe network traffic flow	25
1.3.9.27.1. Network Observability Operator updates	25
1.3.10. Storage	25
1.3.10.1. Persistent storage using Microsoft Azure File CSI Driver Operator is generally available	25
1.3.10.2. Automatic CSI migration for OpenStack Cinder is generally available	25
1.3.10.3. Automatic CSI migration for Microsoft Azure Disk is generally available	26
1.3.10.4. Expanding CSI volumes is generally available	26
1.3.10.5. Support for CSI generic ephemeral volumes is generally available	26
1.3.10.6. VMware vSphere supports resize and snapshots	26
1.3.11. Registry	27
1.3.11.1. Image Registry Operator distribution across availability zones	27
1.3.11.2. Red Hat OpenShift Data Foundation registry storage	27
1.3.12. Operator lifecycle	27
1.3.12.1. File-based catalog format	27
1.3.13. Operator development	27
1.3.13.1. Java-based Operators (Technology Preview)	27
1.3.13.2. Operator SDK support for file-based catalogs	27
1.3.13.3. Validating Operator bundles	28
1.3.14. Jenkins	28
1.3.15. Machine API	28

1.3.15.1. Configuration options for the Amazon EC2 Instance Metadata Service	28
1.3.15.2. Machine API support for Azure ultra disks	28
1.3.15.3. Configuration options for Google Cloud Platform persistent disk types	28
1.3.15.4. Machine API support for Nutanix clusters	28
1.3.15.5. Managing machines with the Cluster API (Technology Preview)	29
1.3.16. Machine Config Operator	29
1.3.16.1. MCO now updates nodes by zone and age	29
1.3.16.2. Enhanced notification for paused Machine Config Pools upon certificate renewal	29
1.3.17. Nodes	29
1.3.17.1. Self Node Remediation Operator replaces the Poison Pill Operator	29
1.3.17.2. Worker nodes for single-node OpenShift clusters	29
1.3.17.3. Descheduler now defaults to simulating pod evictions	30
1.3.17.4. New descheduler customizations	30
1.3.17.5. Node Maintenance Operator enhancements	30
1.3.18. Logging	30
1.3.18.1. Red Hat OpenShift on RHV Logging (Technology Preview)	30
1.3.19. Monitoring	30
1.3.19.1. Updates to monitoring stack components and dependencies	30
1.3.19.2. Changes to alerting rules	31
1.3.19.3. Enable alert routing for user workload monitoring	31
1.3.19.4. Enable a dedicated Alertmanager instance for user-defined alerts	31
1.3.19.5. Use additional authentication settings for remote write configuration	31
1.3.19.6. Create, browse, and manage PromQL queries more easily in the web console	32
1.3.19.7. Scrape interval doubled for ServiceMonitors in single node deployments	32
1.3.19.8. Create alerting rules based on platform monitoring metrics (Technology Preview)	32
1.3.19.9. Add cluster ID labels to remote write storage	32
1.3.19.10. Query metrics using the federation endpoint for user workload monitoring	32
1.3.19.11. Enable body size limit for metrics scraping for default platform monitoring	32
1.3.19.12. Configure retention size settings for metrics storage	32
1.3.19.13. Configure the retention time period for Thanos Ruler in user-defined projects	32
1.3.20. Scalability and performance	33
1.3.20.1. Workload hints for the Node Tuning Operator	33
1.3.20.2. Enhancement to scale-up operations for etcd clusters	33
1.3.20.3. Configuring Ingress Controller (router) Liveness, Readiness, and Startup probes	33
1.3.20.4. New power reduction CPU capability	33
1.3.20.5. Node Observability Operator (Technology Preview)	33
1.3.20.6. Performance Addon Operator functions moved to the Node Tuning Operator	34
1.3.20.7. Low latency tuning documentation updates	34
1.3.20.8. Hub and spoke cluster support	34
1.3.20.9. Enhanced SRO cluster upgrade support	34
1.3.20.10. Enhanced debugging and logging for SRO	34
1.3.20.11. Support for external registries	34
1.3.21. Insights Operator	35
1.3.21.1. Insights Operator data collection enhancements	35
1.3.22. Authentication and authorization	35
1.3.22.1. Additional supported OIDC providers	35
1.3.22.2. Pod security admission	35
1.4. NOTABLE TECHNICAL CHANGES	36
Network Observability operator for observing network flows	36
Update in default value for setting the router load-balancing algorithm	36
LegacyServiceAccountTokenNoAutoGeneration is on by default	36
Operator SDK 1.22.2	37
Cluster Operators no longer referred to as platform Operators	37

1.5. DEPRECATED AND REMOVED FEATURES	38
1.5.1. Deprecated features	39
1.5.1.1. OpenShift CLI (oc) commands and flags for requesting tokens are deprecated	39
1.5.1.2. Red Hat Virtualization (RHV) as a host platform for OpenShift Container Platform will be deprecated	39
1.5.1.3. Support for vSphere 7.0 Update 1 or earlier is deprecated	39
1.5.1.4. Support for ESXi 7.0 Update 1 or earlier is deprecated	39
1.5.1.5. Support for pidsLimit and logSizeMax CRI-O parameters will be deprecated	39
1.5.2. Removed features	40
1.5.2.1. RHEL 7 support for the OpenShift CLI (oc) has been removed	40
1.5.2.2. OpenShift CLI (oc) commands have been removed	40
1.5.2.3. Grafana component removed from monitoring stack	40
1.5.2.4. Prometheus and Grafana user interface access removed from monitoring stack	40
1.5.2.5. Support for virtual hardware version 13 is removed	40
1.5.2.6. Support for vSphere 6.7 Update 2 or earlier is removed	40
1.5.2.7. Support for ESXi 6.7 Update 2 or earlier is removed	40
1.5.2.8. Support for snapshot v1beta1 API endpoint is removed	41
1.5.2.9. Support for manually deploying a custom scheduler has been removed	41
1.5.2.10. Support for deploying single-node OpenShift with OpenShiftSDN has been removed	41
1.5.2.11. Removal of Jenkins images from install payload	41
1.5.3. Future Kubernetes API removals	41
1.6. BUG FIXES	41
Bare Metal Hardware Provisioning	42
Builds	42
Cloud Compute	43
Cluster Version Operator	45
Console Metal3 Plugin	45
Domain Name System (DNS)	45
Image Registry	46
Installer	47
Kubernetes API server	49
Kubernetes Scheduler	50
Machine Config Operator	50
Compliance Operator	50
Management Console	50
Monitoring	50
Networking	52
Networking performance improvements	53
Node	54
OpenShift CLI (oc)	54
Kubernetes Controller Manager	55
Operator Lifecycle Manager (OLM)	55
Operator SDK	56
OpenShift API server	57
Red Hat Enterprise Linux CoreOS (RHCOS)	57
Performance Addon Operator	57
Routing	57
Scalability and performance	59
Storage	59
Web console (Developer perspective)	60
1.7. TECHNOLOGY PREVIEW FEATURES	62
1.8. KNOWN ISSUES	66
1.9. ASYNCHRONOUS ERRATA UPDATES	70



1.9.1. RHSA-2022:5069 - OpenShift Container Platform 4.11.0 image release, bug fix, and security update advisory	71
1.9.2. RHSA-2022:6103 - OpenShift Container Platform 4.11.1 bug fix and security update	71
1.9.2.1. Features	71
1.9.2.1.1. General availability of pod-level bonding for secondary networks	71
1.9.2.2. Bug fixes	71
1.9.2.3. Updating	72
1.9.3. RHBA-2022:6143 - OpenShift Container Platform 4.11.2 bug fix update	72
1.9.3.1. Updating	72
1.9.4. RHSA-2022:6287 - OpenShift Container Platform 4.11.3 bug fix update and security update	72
1.9.4.1. Features	72
1.9.5. Scalability and performance	72
1.9.5.1. Updating	73
1.9.6. RHBA-2022:6376 - OpenShift Container Platform 4.11.4 bug fix update	73
1.9.6.1. Updating	73
1.9.7. RHSA-2022:6536 - OpenShift Container Platform 4.11.5 bug fix and security update	73
1.9.7.1. Known issues	73
1.9.7.2. Bug fixes	73
1.9.7.3. Updating	73
1.9.8. RHBA-2022:6659 - OpenShift Container Platform 4.11.6 bug fix update	73
1.9.8.1. OpenShift Container Platform 4.11 RAN new features	74
1.9.8.1.1. Recovering clusters after a failed update	74
1.9.8.1.2. Disable chronyd in the PolicyGenTemplate custom resource (CR)	74
1.9.8.2. OpenShift Container Platform 4.11 RAN known issues	74
1.9.8.3. Updating	78
1.9.9. RHSA-2022:6732 - OpenShift Container Platform 4.11.7 bug fix update	79
1.9.9.1. Updating	79
1.9.10. RHBA-2022:6809 - OpenShift Container Platform 4.11.8 bug fix update	79
1.9.10.1. Bug fixes	79
1.9.10.2. Updating	79
1.9.11. RHBA-2022:6897 - OpenShift Container Platform 4.11.9 bug fix update	79
1.9.11.1. Updating	79
1.9.12. RHSA-2022:7201 - OpenShift Container Platform 4.11.12 bug fix and security update	80
1.9.12.1. Known issues	80
1.9.12.2. Notable technical changes	80
1.9.12.3. Updating	80
1.9.13. RHBA-2022:7201 - OpenShift Container Platform 4.11.13 bug fix update	80
1.9.13.1. Notable technical changes	80
1.9.13.2. Updating	81
1.9.14. RHSA-2022:8535 - OpenShift Container Platform 4.11.16 bug fix and security update	81
1.9.14.1. Notable technical changes	81
1.9.14.2. Bug fixes	81
1.9.14.3. Updating	81
1.9.15. RHBA-2022:8627 - OpenShift Container Platform 4.11.17 bug fix and security update	81
1.9.15.1. Updating	81
1.9.16. RHBA-2022:8698 - OpenShift Container Platform 4.11.18 bug fix update	82
1.9.16.1. Enhancements	82
1.9.16.2. Notable technical changes	82
1.9.16.3. Bug fixes	82
1.9.16.4. Updating	82
1.9.17. RHSA-2022:8893 - OpenShift Container Platform 4.11.20 bug fix and security update	82
1.9.17.1. Bug fixes	83
1.9.17.2. Known issues	83

1.9.17.3. Updating	83
1.9.18. RHSA-2022:9107 - OpenShift Container Platform 4.11.21 bug fix and security update	83
1.9.18.1. Bug fixes	83
1.9.18.2. Updating	84
1.9.19. RHBA-2023:0027 - OpenShift Container Platform 4.11.22 bug fix update	84
1.9.19.1. Bug fixes	84
1.9.19.2. Known issues	84
1.9.19.3. Updating	84
1.9.20. RHSA-2023:0069 - OpenShift Container Platform 4.11.24 bug fix and security update	85
1.9.20.1. Updating	85
1.9.21. RHSA-2023:0245 - OpenShift Container Platform 4.11.25 bug fix and security update	85
1.9.21.1. Updating	85
1.9.22. RHSA-2023:0565 - OpenShift Container Platform 4.11.26 bug fix and security update	85
1.9.22.1. Known issues	86
1.9.22.2. Bug fixes	86
1.9.22.3. Updating	86
1.9.23. RHSA-2023:0651 - OpenShift Container Platform 4.11.27 bug fix and security update	86
1.9.23.1. Bug fixes	86
1.9.23.2. Updating	86
1.9.24. RHSA-2023:0774 - OpenShift Container Platform 4.11.28 bug fix and security update	86
1.9.24.1. Updating	87
1.9.25. RHSA-2023:0895 - OpenShift Container Platform 4.11.29 bug fix and security update	87
1.9.25.1. Updating	87
1.9.26. RHSA-2023:1030 - OpenShift Container Platform 4.11.30 bug fix and security update	87
1.9.26.1. Bug fixes	87
1.9.26.2. Updating	87
1.9.27. RHSA-2023:1158 - OpenShift Container Platform 4.11.31 bug fix and security update	88
1.9.27.1. Updating	88
1.9.28. RHBA-2023:1296 - OpenShift Container Platform 4.11.32 bug fix and security update	88
1.9.28.1. Updating	88
1.9.29. RHBA-2023:1396 - OpenShift Container Platform 4.11.33 bug fix	88
1.9.29.1. Updating	88
1.9.30. RHSA-2023:1504 - OpenShift Container Platform 4.11.34 bug fix and security update	89
1.9.30.1. Updating	89
1.9.31. RHBA-2023:1650 - OpenShift Container Platform 4.11.35 bug fix	89
1.9.31.1. Bug fixes	89
1.9.31.2. Updating	89
1.9.32. RHBA-2023:1733 - OpenShift Container Platform 4.11.36 bug fix	89
1.9.32.1. Updating	90
1.9.33. RHBA-2023:1760 - OpenShift Container Platform 4.11.37 bug fix	90
1.9.33.1. Updating	90
1.9.34. RHBA-2023:1863 - OpenShift Container Platform 4.11.38 bug fix update	90
1.9.34.1. Updating	90
1.9.35. RHSA-2023:2014 - OpenShift Container Platform 4.11.39 bug fix and security update	90
1.9.35.1. Bug fixes	91
1.9.35.2. Updating	91
1.9.36. RHBA-2023:2694 - OpenShift Container Platform 4.11.40 bug fix update	91
1.9.36.1. Bug fixes	91
1.9.36.2. Updating	91
1.9.37. RHBA-2023:3213 - OpenShift Container Platform 4.11.41 bug fix update	91
1.9.37.1. Updating	92
1.9.38. RHSA-2023:3309 - OpenShift Container Platform 4.11.42 bug fix and security update	92
1.9.38.1. Updating	92

---

1.9.39. RHSA-2023:3542 OpenShift Container Platform 4.11.43 bug fix and security update	92
1.9.39.1. Updating	92
1.9.40. RHSA-2023:3915 - OpenShift Container Platform 4.11.44 bug fixes and security update	92
1.9.40.1. Bug fixes	93
1.9.40.2. Updating	93
1.9.41. RHSA-2023:4053 OpenShift Container Platform 4.11.45 bug fix and security update	93
1.9.41.1. Updating	93
1.9.42. RHSA-2023:4310 OpenShift Container Platform 4.11.46 bug fix and security update	94
1.9.42.1. Updating	94
1.9.43. RHBA-2023:4614 OpenShift Container Platform 4.11.47 bug fix update	94
1.9.43.1. Updating	94
1.9.44. RHBA-2023:4752 OpenShift Container Platform 4.11.48 bug fix update	94
1.9.44.1. Updating	94
1.9.45. RHSA-2023:5001 OpenShift Container Platform 4.11.49 bug fix update	94
1.9.45.1. Updating	95



# CHAPTER 1. OPENSIFT CONTAINER PLATFORM 4.11 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 multitenant 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-2022:5069](#)) is now available. This release uses [Kubernetes 1.24](#) with CRI-O runtime. New features, changes, and known issues that pertain to OpenShift Container Platform 4.11 are included in this topic.

OpenShift Container Platform 4.11 clusters are available at <https://console.redhat.com/openshift>. With the Red Hat OpenShift Cluster Manager application for OpenShift Container Platform, you can deploy OpenShift clusters to either on-premises or cloud environments.

OpenShift Container Platform 4.11 is supported on Red Hat Enterprise Linux (RHEL) 8.5 through 8.7, as well as on Red Hat Enterprise Linux CoreOS (RHCOS) 4.11.

You must use RHCOS machines for the control plane, and you can use either RHCOS or RHEL 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. Improved support for NVMe over Fabrics

OpenShift Container Platform 4.11 introduces the **nvme-cli** package that provides an interface for managing NVMe devices.

#### 1.3.1.2. Investigate kernel crashes on AMD64 machines with kdump

RHCOS now supports **kdump** for the **x86\_64** architecture in OpenShift Container Platform 4.11. Support for **kdump** on other architectures remains in Technology Preview.

### 1.3.1.3. Investigate kernel crashes on ARM64 machines with kdump (Technology Preview)

RHCOS now supports **kdump** for the **arm64** architecture in OpenShift Container Platform 4.11 as a Technology Preview.

### 1.3.1.4. RHCOS now uses RHEL 8.6

RHCOS now uses Red Hat Enterprise Linux (RHEL) 8.6 packages in OpenShift Container Platform 4.11 and above. This enables you to have the latest fixes, features, and enhancements, as well as the latest hardware support and driver updates.

### 1.3.1.5. Updated RHCOS registry URL

The redirector hostname for downloading RHCOS boot images is now **rhcos.mirror.openshift.com**. You must configure your firewall to grant access to the boot images. For more information, see [Configuring your firewall for OpenShift Container Platform](#).

## 1.3.2. Installation and upgrade

### 1.3.2.1. RHEL 9 support for the OpenShift installer

Using Red Hat Enterprise Linux (RHEL) 9 with the OpenShift installer (openshift-install) is now supported.

For more information, see the "Obtaining the installation program" section of the installation documentation for your platform.

### 1.3.2.2. New minimum system requirements for installing OpenShift Container Platform on a single node

This release updates the minimum system requirements for installing OpenShift Container Platform on a single node. When installing OpenShift Container Platform on a single node, you should configure a minimum of 16 GB of RAM. Specific workload requirements can require additional RAM. The complete list of supported platforms has been updated to include bare metal, vSphere, Red Hat OpenStack Platform (RHOSP), and Red Hat Virtualization platforms. In all cases, you must specify the **platform.none: {}** parameter in the **install-config.yaml** configuration file when the **openshift-installer** binary is being used to install single-node OpenShift.

### 1.3.2.3. OpenShift Container Platform on ARM

OpenShift Container Platform 4.11 is now supported on ARM architecture based AWS user-provisioned infrastructure and bare-metal installer-provisioned infrastructure. For more information about instance availability and installation documentation, see [Supported installation methods for different platforms](#).

The following features are supported for OpenShift Container Platform on ARM:

- Disconnected installation support
- Elastic file system (EFS) for AWS
- Local storage operator on bare metal
- Internet Small Computer Systems Interface (iSCSI) for bare metal

The following Operators are supported for OpenShift Container Platform on ARM:

- Special resource operator (SRO)

#### 1.3.2.4. Troubleshooting bootstrap failures during installation on AWS

The installation program now gathers serial console logs from the bootstrap and control plane hosts on AWS. This log data is added to the standard bootstrap log bundle.

For more information, see [Troubleshooting installation issues](#).

#### 1.3.2.5. Support for Microsoft Hyper-V generation version 2

By default, the installation program now deploys a Microsoft Azure cluster using Hyper-V generation version 2 virtual machines (VMs). If the installation program detects that the instance type selected for the VMs does not support version 2, it uses version 1 for the deployment.

#### 1.3.2.6. Default AWS and VMware vSphere compute node resources

Beginning with OpenShift Container Platform 4.11, by default, the installation program now deploys AWS and VMware vSphere compute nodes with 4 vCPUs and 16 GB of virtual RAM.

#### 1.3.2.7. Support for the AWS SC2S region

OpenShift Container Platform 4.11 introduces support for the AWS Secret Commercial Cloud Services (SC2S) region. You can now install and update OpenShift Container Platform clusters in the **us-isob-east-1** SC2S region.

For more information, see [Installing a cluster on AWS into a Secret or Top Secret Region](#)

#### 1.3.2.8. Installing a cluster on Nutanix using installer-provisioned infrastructure

OpenShift Container Platform 4.11 introduces support for installing a cluster on Nutanix using installer-provisioned infrastructure. This type of installation lets you use the installation program to deploy a cluster on infrastructure that the installation program provisions and the cluster maintains.

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

#### 1.3.2.9. Installing OpenShift Container Platform using Azure Ultra SSD

You can now enable Ultra SSD storage when installing OpenShift Container Platform on Azure. This feature requires that both the Azure region and zone where you install OpenShift Container Platform offer Ultra storage.

For more information, see [Additional Azure configuration parameters](#).

#### 1.3.2.10. Added support for `bootstrapExternalStaticIP` and `bootstrapExternalStaticGateway` configuration settings

When deploying an installer-provisioned OpenShift Container Platform cluster on bare metal with static IP addresses and no DHCP server on the **baremetal** network, you must specify a static IP address for the bootstrap VM and the static IP address of the gateway for the bootstrap VM. OpenShift Container Platform 4.11 provides the **`bootstrapExternalStaticIP`** and the **`bootstrapExternalStaticGateway`** configuration settings, which you can set in the **`install-config.yaml`** file before deployment. The

introduction of these settings replaces the workaround procedure [Assigning a bootstrap VM an IP address on the baremetal network without a DHCP server](#) from the OpenShift Container Platform 4.10 release.

For more information, see [Configuring the install-config.yaml file](#) and [Additional install-config parameters](#).

### 1.3.2.11. Configuring Fujitsu hardware

OpenShift Container Platform 4.11 introduces support for configuring the BIOS and RAID arrays of control plane nodes when installing OpenShift Container Platform on bare metal with Fujitsu hardware. In OpenShift Container Platform 4.10, configuring the BIOS and RAID arrays on Fujitsu hardware was limited to worker nodes.

For more information, see [Configuring the BIOS](#) and [Configuring the RAID](#).

### 1.3.2.12. Disconnected mirroring with the oc-mirror CLI plugin is now generally available

You can use the oc-mirror OpenShift CLI (**oc**) plugin to mirror images in a disconnected environment. This feature was previously introduced as a Technology Preview in OpenShift Container Platform 4.10 and is now generally available in OpenShift Container Platform 4.11.

This release of the oc-mirror plugin includes the following new features:

- Pruning images from the target mirror registry
- Specifying version ranges for Operator packages and OpenShift Container Platform releases
- Generating supporting artifacts for OpenShift Update Service (OSUS) usage
- Obtaining a template for the initial image set configuration



#### IMPORTANT

If you used the Technology Preview version of the oc-mirror plugin for OpenShift Container Platform 4.10, it is not possible to migrate your mirror registry to OpenShift Container Platform 4.11. You must download the new oc-mirror plugin, use a new storage back end, and use a new top-level namespace on the target mirror registry.

For more information, see [Mirroring images for a disconnected installation using the oc-mirror plugin](#).

### 1.3.2.13. Installing a cluster on Azure using user-managed encryption keys

OpenShift Container Platform 4.11 introduces support for installing a cluster on Azure with user-managed disk encryption.

For more information, see [Enabling user-managed encryption for Azure](#).

### 1.3.2.14. Accelerated Networking for Azure enabled by default

OpenShift Container Platform 4.11 on Azure provides accelerated networking for control plane and compute nodes. Accelerated networking is enabled by default for supported instance types in an installer-provisioned infrastructure installation.

For more information, see [Openshift 4 on Azure - accelerated networking](#).



### 1.3.2.15. AWS VPC endpoints and restricted installations

You are no longer required to configure AWS VPC endpoints when installing a restricted OpenShift Container Platform cluster on AWS. While configuring VPC endpoints remains an option, you can also choose to configure a proxy without VPC endpoints or configure a proxy with VPC endpoints.

For more information, see [Requirements for using your VPC](#).

### 1.3.2.16. Additional customization when installing OpenShift Container Platform

OpenShift Container Platform 4.11 allows you to disable the installation of the **baremetal** and **marketplace** Operators, and the **openshift-samples** content that is stored in the **openshift** namespace. You can disable these features by adding the **baselineCapabilitySet** and **additionalEnabledCapabilities** parameters to the **install-config.yaml** configuration file prior to installation. If you disable any of these capabilities during the installation, you can enable them after the cluster is installed. After a capability has been enabled, it cannot be disabled again.

For more information, see the "Installation configuration parameters" section of the installation documentation for your platform.

### 1.3.2.17. Azure Marketplace offering

OpenShift Container Platform is now available on the Azure Marketplace. The Azure Marketplace offering is available to customers who procure OpenShift Container Platform in North America and EMEA.

For more information, see [Installing OpenShift using Azure Marketplace](#).

### 1.3.2.18. AWS Marketplace offering

OpenShift Container Platform is now available on the AWS Marketplace. The AWS Marketplace offering is available to customers who procure OpenShift Container Platform in North America.

For more information, see [Installing OpenShift using AWS Marketplace](#).

### 1.3.2.19. CSI driver installation on vSphere clusters

To install a CSI driver on a cluster running on vSphere, you must have the following components installed:

- Virtual hardware version 15 or later
- vSphere version 7.0 Update 2 or later, up to but not including version 8. vSphere 8 is not supported.
- VMware ESXi version 7.0 Update 2 or later

Components with versions earlier than those above are deprecated or removed. Deprecated versions are still fully supported, but Red Hat recommends that you use ESXi 7.0 Update 2 or later and vSphere 7.0 Update 2 up to but not including version 8. vSphere 8 is not supported.

For more information, see [Deprecated and removed features](#).

## 1.3.3. Post-installation configuration

### 1.3.3.1. Cluster capabilities

As a cluster administrator, you can enable cluster capabilities to select or deselect one or more optional components before installation or post installation.

For more information, see [Cluster capabilities](#).

### 1.3.3.2. OpenShift Container Platform clusters with multi-architecture compute machines (Technology Preview)

OpenShift Container Platform 4.11 introduces clusters with multi-architecture compute machines support using Azure installer-provisioned infrastructure in Technology Preview. This feature offers, as a day-two operation, the ability to add **arm64** compute nodes to an existing **x86\_64** Azure cluster that is installer provisioned with a multi-architecture installer binary. You can add **arm64** compute nodes to your cluster by creating a custom Azure machine set that uses a manually generated **arm64** boot image. Control planes on **arm64** architectures are not currently supported. For more information, see [Configuring a multi-architecture cluster](#).



#### NOTE

You can manually upgrade your cluster to the latest multi-architecture release image by using the release **image-pullsec**. For more information, see [Upgrading your multi-architecture compute machines](#).

### 1.3.4. Web console

#### 1.3.4.1. Developer Perspective

- With this update, in the developer perspective, you can add your GitHub repository containing pipelines to the OpenShift Container Platform cluster. You can now run pipelines and tasks from your GitHub repository on the cluster when relevant Git events, such as push or pull requests are triggered.
  - In the administrator perspective, you can configure your GitHub application with the OpenShift cluster to use a pipeline as code. With this configuration, you can execute a set of tasks required for build deployment.
- With this update, you can create a customized pipeline using your own set of curated tasks. You can search, install, and upgrade your tasks directly from the developer console.
- With this update, in the web terminal you can now have multiple tabs, view bash history, and the web terminal remains open until you close the browser window or tab.
- With this update, in the **Add+** page of the developer perspective, a new menu added to share project and Helm Chart repositories that allows to add or remove users to the project.

#### 1.3.4.2. Dynamic plugin updates

With this update, you can use the new **console.openshift.io/use-i18next** annotation to determine if the **ConsolePlugin** contains localization resources. If the annotation is set to **"true"**, the localization resources from the i18n namespace named after the dynamic plugin, are loaded. If the annotation is set to any other value or is missing on the **ConsolePlugin** resource, localization resources are not loaded.

For more information, see [Overview of dynamic plugins](#).

### 1.3.4.3. Support for dark mode theme

The OpenShift Container Platform web console now supports the dark mode theme. On the **User Preferences** page, select your preferred theme to view the web console in.

### 1.3.4.4. Display operand instances for all managed namespaces on the Installed Operator page

With this update, the **Operator → Installed Operator** page will show all Operators across all namespaces. You are still able to view only the instances in the selected namespace within the project selector. When viewing the operand instances, a new switching control allows all operand instances from either all namespaces or only the current namespace to be seen.

### 1.3.4.5. Conditional updates

With this update, if conditional updates are available, you can enable **Include supported but not recommended versions** in the **Select new version** dropdown of the **Update cluster** modal to populate the dropdown list with conditional updates. If a **Supported but not recommended** version is selected, an alert will appear below the dropdown menu displaying potential issues with the version.

### 1.3.4.6. Pod disruption budgets (PDBs)

This update provides support for pod disruption budgets (PDBs) to the OpenShift Container Platform web console. From **Workloads → PodDisruptionBudgets**, you can create PDBs for pod resources. You can select **maxUnavailable** and **minAvailable** from the availability requirement list and set the value of pods running. Alternatively, pod disruption budgets can be created from **pod controller resources** list and **Detail** pages. For example, from **Workloads → Deployments** click **Add PodDisruptionBudget**.

For more information, see [Pod preemption and other scheduler settings](#).

## 1.3.5. OpenShift CLI (oc)

### 1.3.5.1. RHEL 9 support for the OpenShift CLI (oc)

Using Red Hat Enterprise Linux (RHEL) 9 with the OpenShift CLI (**oc**) is now supported.



#### NOTE

It is not supported to install the OpenShift CLI (**oc**) as an RPM for Red Hat Enterprise Linux (RHEL) 9. You must install the OpenShift CLI for RHEL 9 by downloading the binary.

For more information, see [Installing the OpenShift CLI](#).

## 1.3.6. IBM Z and LinuxONE

With this release, IBM Z and LinuxONE are now compatible with OpenShift Container Platform 4.11. 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 4.11:

- Alternate Authentication Provider
- Automatic Device Discovery with Local Storage Operator
- CSI Volumes
  - Cloning
  - Expansion
  - Snapshot
- File Integrity Operator
- Monitoring for user-defined projects
- Operator API
- OC CLI plugin

### Supported features

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

- Currently, the following Operators are supported:
  - Cluster Logging Operator
  - Compliance Operator
  - Local Storage Operator
  - NFD Operator
  - NMState Operator
  - OpenShift Elasticsearch Operator
  - Service Binding Operator
  - Vertical Pod Autoscaler Operator
- The following Multus CNI plugins are supported:
  - Bridge
  - Host-device
  - IPAM
  - IPVLAN

- Encrypting data stored in etcd
- Helm
- Horizontal pod autoscaling
- 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, including IPsec encryption
- Support for multiple network interfaces
- 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 4.11:

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

### Restrictions

The following restrictions impact 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
  - Red Hat OpenShift Local
  - Controlling overcommit and managing container density on nodes
  - FIPS cryptography
  - NVMe
  - OpenShift Metering
  - OpenShift Virtualization
  - Tang mode disk encryption during OpenShift Container Platform deployment

- Compute nodes must run Red Hat Enterprise Linux CoreOS (RHCOS)
- Persistent shared storage must be provisioned by using either Red Hat OpenShift Data Foundation 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

### 1.3.7. IBM Power

With this release, IBM Power is now compatible with OpenShift Container Platform 4.11. For installation instructions, see the following documentation:

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

#### Notable enhancements

The following new features are supported on IBM Power with OpenShift Container Platform 4.11:

- Alternate Authentication Provider
- CSI Volumes
  - Cloning
  - Expansion
  - Snapshot
- File Integrity Operator
- IPv6
- Monitoring for user-defined projects
- Operator API
- OC CLI plugin

#### Supported features

The following features are also supported on IBM Power:

- Currently, the following Operators are supported:
  - Cluster Logging Operator
  - Compliance Operator
  - Local Storage Operator
  - NFD Operator
  - NMState Operator
  - OpenShift Elasticsearch Operator
  - SR-IOV Network Operator

- Service Binding Operator
- Vertical Pod Autoscaler Operator
- The following Multus CNI plugins are supported:
  - Bridge
  - Host-device
  - IPAM
  - IPVLAN
- Encrypting data stored in etcd
- Helm
- Horizontal pod autoscaling
- Multipathing
- Multus SR-IOV
- OVN-Kubernetes, including IPsec encryption
- 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
- Support for multiple network interfaces
- Support for Power10
- Three-node cluster support
- 4K Disk Support

### Restrictions

The following restrictions impact OpenShift Container Platform on IBM Power:

- 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
  - Red Hat OpenShift Local
  - Controlling overcommit and managing container density on nodes

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

### 1.3.8. Security and compliance

#### 1.3.8.1. Audit logs now include OAuth server audit events

OAuth server audit events, annotated with login events, are now logged at the metadata level in the audit logs. The login events include failed login attempts.

For more information, see [About audit log policy profiles](#).

### 1.3.9. Networking

#### 1.3.9.1. Pod-level bonding for secondary networks

Bonding at the pod level is vital to enable workloads inside pods that require high availability and more throughput. With pod-level bonding, you can create a bond interface from multiple single root I/O virtualization (SR-IOV) virtual function interfaces in kernel mode interface. The SR-IOV virtual functions are passed into the pod and attached to a kernel driver.

Scenarios where pod-level bonding is required include creating a bond interface from multiple SR-IOV virtual functions on different physical functions. Creating a bond interface from two different physical functions on the host can be used to achieve high availability at pod level.

For more information, see [Configuring a bond interface from two SR-IOV interfaces](#)

#### 1.3.9.2. New option for Ingress Controllers with the hostnetwork endpoint

This update introduces a new option for Ingress Controllers with the **hostnetwork** endpoint strategy. You can now host multiple Ingress Controllers on the same worker node using **httpPort**, **httpsPort**, and **statsPort** binding ports.

#### 1.3.9.3. Multi-node configuration for control plane and worker nodes

You can simultaneously apply a single configuration to multiple, bare-metal, installer-provisioned infrastructure nodes in a cluster. Applying a single configuration to multiple nodes reduces the risk of misconfiguration due to the single-provisioning process.

This mechanism is only available for initial deployments when the **install-config** file is used.

#### 1.3.9.4. Support for configuring Classic Load Balancer Timeouts on AWS



You can now configure idle connection timeouts for AWS Classic Load Balancers (CLBs) in the Ingress Controller.

For more information, see [Configuring Classic Load Balancer timeouts](#)

#### 1.3.9.5. Update to HAProxy 2.2.24

OpenShift Container Platform updated to HAProxy 2.2.24.

#### 1.3.9.6. Support for configuring maximum number of connections for HAProxy processes

You can now set the maximum number of simultaneous connections that can be established per HAProxy process in the Ingress Controller to any value between 2000 and 2,000,000.

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

#### 1.3.9.7. Setting the Ingress Controller health check interval

With this update, a cluster administrator can set the health check interval to define how long the router waits between two consecutive health checks. This value is applied globally as a default for all routes. The default value is 5 seconds.

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

#### 1.3.9.8. Support for configuring interface-level safe network sysctls

Use the new **tuning-cni** meta plugin to set an interface level safe network sysctls that only applies to a specific interface. For example, you can change the behavior of **accept\_redirects** on a particular network interface by configuring the **tuning-cni** plugin. A complete list of the interface-specific safe sysctls that can be set is available in the documentation.

In addition to this enhancement the set of system-wide safe sysctls that can be set has increased to support **net.ipv4.ping\_group\_range** and **net.ipv4.ip\_unprivileged\_port\_start**.

For more information about configuring the **tuning-cni** plugin, see [Setting interface level network sysctls](#).

For more information about the newly supported interface level network safe sysctls and updates to the list of supported system-wide safe sysctls, see [Using sysctls in containers](#).

#### 1.3.9.9. Support for CoreDNS forwarding DNS requests over TLS

When working in a highly regulated environment, you might need the ability to secure domain name system (DNS) traffic when forwarding requests to upstream resolvers so that you can ensure additional DNS traffic and data privacy. Cluster administrators can now configure transport layer security (TLS) for forwarded DNS queries. This feature applies only to the DNS Operator and not the CoreDNS instance managed by the Machine Config Operator.

For more information, see [Using DNS forwarding](#).

#### 1.3.9.10. Internal traffic support for OVN-Kubernetes

As a cluster administrator, you can configure **internalTrafficPolicy=Local** on a Kubernetes service object when using the OVN-Kubernetes Container Network Interface (CNI) cluster network provider. This feature allows cluster administrators to route traffic to an endpoint on the same node that the

traffic originated from. If there are no local node endpoints, traffic will be dropped.

For more information, see [Service Internal Traffic Policy](#).

### 1.3.9.11. Support for AWS Load Balancer Operator (Technology Preview)

As a cluster administrator, you can install the AWS Load Balancer Operator from the OperatorHub by using the OpenShift Container Platform web console or CLI. The AWS Load Balancer Operator is in Technology Preview.

For more information, see [Installing AWS Load Balancer Operator](#).

### 1.3.9.12. Enhancement to the Route API

Previously, you could not specify the subdomain of a route, and the **spec.host** field was required to set the host name. You can now specify the **spec.subdomain** field and omit the **spec.host** field of a route. The router deployment that exposes the route will use the **spec.subdomain** value to determine the host name.

You can use this enhancement to simplify sharding by enabling a route to have multiple, distinct host names determined by each router deployment that exposes the route.

### 1.3.9.13. External DNS Operator

In OpenShift Container Platform 4.11, the External DNS Operator is available for AWS Route53, Azure DNS, GCP DNS, and Infoblox in General Availability (GA) status. The External DNS Operator is still in Technology Preview (TP) status for BlueCat and AWS Route53 on GovCloud. With this update, the External DNS Operator provides the following enhancements:

- You can create DNS records on a DNS zone on Infoblox.
- By default, the External DNS Operator creates the operands in the namespace **external-dns-operator**. You do not have to manually create a namespace for the operands and Role-based access control (RBAC) prior to installation.
- You can use the status of the route to retrieve the DNS FQDN names.
- The proxy support for BlueCat DNS provider is now available.
- You can enable the automatic DNS configuration deployment while using the BlueCat DNS provider.

Ensure that you migrate from TP to GA. The upstream version of **ExternalDNS** for an OpenShift Container Platform 4.11 is **v0.12.0** and for the TP is **v0.10.2**. For more information, see [About the External DNS Operator](#).

### 1.3.9.14. PTP support for dual NIC boundary clocks

You can now configure dual network interfaces (NICs) as boundary clocks with **PtpConfig** profiles for each NIC channel.

For more information, see [Using PTP with dual NIC hardware](#).

### 1.3.9.15. PTP events enhancements

A new PTP events API endpoint is available, **api/cloudNotifications/v1/publishers**. Using this endpoint, you can get PTP **os-clock-sync-state**, **ptp-clock-class-change**, and **lock-state** details for the cluster node.

For more information, see [Subscribing DU applications to PTP events REST API reference](#) .

#### 1.3.9.16. SR-IOV support for Pensando DSC cards

SR-IOV support is now available for [Pensando DSC cards](#) . OpenShift SR-IOV is supported, but you must set a static, Virtual Function (VF) media access control (MAC) address using the SR-IOV CNI config file when using SR-IOV.

#### 1.3.9.17. SR-IOV support for Mellanox MT2892 cards

SR-IOV support is now available for [Mellanox MT2892 cards](#) .

#### 1.3.9.18. OpenShift Container Platform CIDR Ranges for Networks

Be aware that CIDR ranges for networks are not adjustable after cluster installation. Red Hat does not provide direct guidance on determining the range because it requires careful consideration of the number of pods created.

#### 1.3.9.19. OVN-Kubernetes network provider: Enable IPsec at runtime

If you are using the OVN-Kubernetes cluster network provider, you can now enable IPsec encryption after cluster installation. For more information about how to enable IPsec, see [Configuring IPsec encryption](#).

#### 1.3.9.20. Support for additional MetalLB CRDs and control over logging verbosity

Additional MetalLB custom resource definitions (CRDs) have been added to support more complex configurations.

The following CRDs have been added:

- **IPAddressPools**
- **L2Advertisement**
- **BGPAdvertisement**
- **Community**

With these enhancements, you can use the Operator for more complex configurations. For example, you can use the enhancements to isolate nodes or segment the network. In addition, the enhancements made to the FRRouting (FRR) logging component allow you to control the verbosity of the logs generated.

**NOTE**

The CRDs documented for OpenShift Container Platform 4.10 and the existing ways to configure MetalLB as described in [About MetalLB and the MetalLB Operator](#) are still supported but deprecated. The **AddressPool** configuration is deprecated.

In 4.10, layer 2 and BGP IP addresses using the **AddressPool** were assigned from different address pools. In OpenShift Container Platform 4.11, layer 2 and BGP IP addresses can be assigned from the same address pool.

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

**1.3.9.21. Ability to create a route using the destination CA certificate in the Ingress annotation**

The **route.openshift.io/destination-ca-certificate-secret** annotation can now be used on an Ingress object to define a route with a custom certificate (CA).

See [Creating a route using the destination CA certificate in the Ingress annotation](#) for more information.

**1.3.9.22. Hosted control planes (Technology Preview)**

With hosted control planes for OpenShift Container Platform, you can host clusters at scale to reduce management costs, optimize cluster deployment time, and separate management and workload concerns. You can enable this deployment model as a Technology Preview feature when you install the multicluster engine for Kubernetes Operator version 2.0. For more information, see [Overview of hosted control planes \(Technology Preview\)](#).

Open Virtual Network (OVN) was redesigned to host its control plane and data store alongside the cluster's control plane. With hosted control planes, OVN supports split control planes.

**1.3.9.23. IPv6 single and dual-stack support on user-provisioned bare metal infrastructure with the OVN-Kubernetes cluster network provider**

For clusters on user-provisioned [bare metal infrastructure](#), the OVN-Kubernetes cluster network provider supports both IPv4 and IPv6 address families.

**1.3.9.24. OVS hardware offloading on RHOSP**

For clusters that run on RHOSP, [Open vSwitch \(OVS\)](#) hardware offloading is now generally available.

For more information, see [Enabling OVS hardware offloading](#).

**1.3.9.25. NFV user experience improvements on RHOSP**

For clusters that run on RHOSP, the network functions virtualization deployment experience is improved. Changes for this release include:

- Network data fetching from metadata service URLs rather than a configuration drive
- Automatic VFIO loading with no-IOMMU for all discovered devices
- DPDK vHost user ports

These changes are reflected in simplified post-installation and network configuration documentation.

### 1.3.9.26. Installations on Red Hat OpenStack Platform, VMware vSphere, or oVirt now configure keepalived with unicast as the default

For OpenShift Container Platform installer-provisioned installations on Red Hat OpenStack Platform (RHOSP), VMware vSphere, or oVirt, keepalived is now configured with unicast as the default instead of multicast. You no longer need to allow multicast traffic. The unicast migration occurs a few minutes after the cluster finishes upgrading because all of the nodes must migrate at the same time. Having both multicast and unicast clusters at the same time should not result in any issues because keepalived treats unicast and multicast as completely separate.

### 1.3.9.27. Network Observability Operator to observe network traffic flow

As an administrator, you can now install the Network Observability Operator to observe the network traffic for your OpenShift Container Platform cluster in the console. You can view and monitor the network traffic data in different graphical representations. The Network Observability Operator uses eBPF technology to create the network flows. The network flows are enriched with OpenShift Container Platform information, and stored in Loki. You can use the network traffic information for detailed troubleshooting and analysis.

The Network Observability Operator is General Availability (GA) status in the 4.12 release of OpenShift Container Platform and is also supported in the 4.10 and 4.11 versions.

For more information, see [Network Observability](#).

#### 1.3.9.27.1. Network Observability Operator updates

The Network Observability Operator releases updates independently from the OpenShift Container Platform minor version release stream. Updates are available through a single, rolling stream which is supported on all currently supported versions of OpenShift Container Platform 4. Information regarding new features, enhancements, and bug fixes for the Network Observability Operator can be found in the [Network Observability release notes](#).

## 1.3.10. Storage

### 1.3.10.1. Persistent storage using Microsoft Azure File CSI Driver Operator is generally available

OpenShift Container Platform is capable of provisioning persistent volumes (PVs) using the Container Storage Interface (CSI) driver for Azure File. This feature was previously introduced as a Technology Preview feature in OpenShift Container Platform 4.10 and is now generally available and enabled by default in OpenShift Container Platform 4.11.

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

### 1.3.10.2. Automatic CSI migration for OpenStack Cinder is generally available

Starting with OpenShift Container Platform 4.8, automatic migration for in-tree volume plugins to their equivalent Container Storage Interface (CSI) drivers became available as a Technology Preview feature. Support for Cinder was provided in this feature in OpenShift Container Platform 4.8, and OpenShift Container Platform 4.11 now supports automatic migration for Cinder as generally available. CSI migration for Cinder is now enabled by default and requires no action by an administrator.

This feature automatically translates in-tree objects to their counterpart CSI representations and should be completely transparent to users. Translated objects are not stored on disk, and user data is not migrated.

While storage class referencing to the in-tree storage plugin will continue working, it is recommended that you switch the default storage class to the CSI storage class.

For more information, see [CSI automatic migration](#).

### 1.3.10.3. Automatic CSI migration for Microsoft Azure Disk is generally available

Starting with OpenShift Container Platform 4.8, automatic migration for in-tree volume plugins to their equivalent Container Storage Interface (CSI) drivers became available as a Technology Preview feature. Support for Azure Disk was provided in this feature in OpenShift Container Platform 4.9, and OpenShift Container Platform 4.11 now supports automatic migration for Azure Disk as generally available. CSI migration for Azure Disk is now enabled by default and requires no action by an administrator.

This feature automatically translates in-tree objects to their counterpart CSI representations and should be completely transparent to users. Translated objects are not stored on disk, and user data is not migrated.

While storage class referencing to the in-tree storage plugin will continue working, it is recommended that you switch the default storage class to the CSI storage class.

For more information, see [CSI automatic migration](#).

### 1.3.10.4. Expanding CSI volumes is generally available

Starting with OpenShift Container Platform 4.3, expansion of Container Storage Interface (CSI) storage volumes after they have already been created became available as a Technology Preview feature, and is now generally available in OpenShift Container Platform 4.11.

For more information, see [Expanding CSI volumes](#).

### 1.3.10.5. Support for CSI generic ephemeral volumes is generally available

OpenShift Container Platform 4.11 supports Container Storage Interface (CSI) generic ephemeral volumes as generally available. Generic ephemeral volumes are a type of ephemeral volume that can be provided by all storage drivers that also support persistent volumes and dynamic provisioning.

For more information, see [Generic ephemeral volumes](#).

### 1.3.10.6. VMware vSphere supports resize and snapshots

OpenShift Container Platform 4.11 supports volume resize and snapshots for vSphere Container Storage Interface (CSI) Driver Operator with the following restrictions:

- Snapshots:
  - Requires vSphere version 7.0 Update 3 or later, up to but not including version 8. vSphere 8 is not supported for both vCenter Server and ESXi.
  - Does not support fileshare volumes.
- Resize:
  - Offline volume expansion: minimum required vSphere version is 6.7 Update 3 P06
  - Online volume expansion: minimum required vSphere version is 7.0 Update 2.

For more information, see [CSI drivers supported by OpenShift Container Platform](#) .

### 1.3.11. Registry

#### 1.3.11.1. Image Registry Operator distribution across availability zones

The default configuration of the Image Registry Operator now spreads image registry pods across topology zones to prevent delayed recovery times in case of a complete zone failure where all pods are impacted.

For more information, see [Image Registry Operator distribution across availability zones](#) .

#### 1.3.11.2. Red Hat OpenShift Data Foundation registry storage

Red Hat OpenShift Data Foundation registry storage supported in OpenShift Container Platform 4.11.

OpenShift Data Foundation integrates multiple storage types that you can use with the internal image registry including:

- Ceph, which is a shared and distributed file system with on-premises object storage
- NooBaa, which provides a Multicloud Object Gateway

### 1.3.12. Operator lifecycle

#### 1.3.12.1. File-based catalog format

The default Red Hat-provided Operator catalogs for OpenShift Container Platform 4.11 releases in the file-based catalog format. OpenShift Container Platform 4.6 through 4.10 released in the SQLite database format. File-based catalogs are the latest iteration of the catalog format in Operator Lifecycle Manager (OLM). It is a plain text-based file in JSON or YAML and is a declarative configuration evolution of the earlier SQLite database format. Cluster administrators and users will not see any change to their install workflows and Operator consumption with the new catalog format.

For more information, see [File-based catalogs](#) .

### 1.3.13. Operator development

#### 1.3.13.1. Java-based Operators (Technology Preview)

Starting in OpenShift Container Platform 4.11 as a Technology Preview feature, the Operator SDK includes the tools and libraries to develop a Java-based Operator. Operator developers can take advantage of Java programming language support in the Operator SDK to build a Java-based Operator and manage its lifecycle.

For more information, see [Getting started with Operator SDK for Java-based Operators](#) .

#### 1.3.13.2. Operator SDK support for file-based catalogs

As of OpenShift Container Platform 4.11, the **run bundle** command supports the file-based catalog format for Operator catalogs by default. The deprecated SQLite database format for Operator catalogs continues to be supported; however, it will be removed in a future release.

For more information, see [Working with bundle images](#).

### 1.3.13.3. Validating Operator bundles

As an Operator author, you can run the **bundle validate** command in the Operator SDK to validate the content and format of an Operator bundle. In addition to the default test, you can run optional validators to test for issues in your bundle, such as an empty CRD description or unsupported Operator Lifecycle Manager (OLM) resources.

For more information, see [Validating Operator bundles](#). In earlier versions of OpenShift Container Platform, the Performance Addon Operator provided automatic, low latency performance tuning for applications. In OpenShift Container Platform 4.11, these functions are part of the Node Tuning Operator. The Node Tuning Operator is part of the standard installation for OpenShift Container Platform 4.11. If you upgrade to OpenShift Container Platform 4.11, the Node Tuning Operator removes the Performance Addon Operator and all related artifacts on startup.

For more information, see [Node Tuning Operator](#).

### 1.3.14. Jenkins

- This enhancement adds a new Jenkins environment variable, **JAVA\_FIPS\_OPTIONS**, that controls how the JVM operates when running on a FIPS node. For more information, see [OpenJDK support article \(BZ#2066019\)](#)

### 1.3.15. Machine API

#### 1.3.15.1. Configuration options for the Amazon EC2 Instance Metadata Service

You can now use machine sets to create compute machines that use a specific version of the Amazon EC2 Instance Metadata Service (IMDS). Machine sets can create compute machines that allow the use of both IMDSv1 and IMDSv2 or compute machines that require the use of IMDSv2.

For more information, see [Machine set options for the Amazon EC2 Instance Metadata Service](#) .

#### 1.3.15.2. Machine API support for Azure ultra disks

You can now create a machine set running on Azure that deploys machines with ultra disks. You can either deploy machines with ultra disks as data disks, or by using persistent volume claims (PVCs) that use in-tree or Container Storage Interface (CSI) PVCs.

For more information, see the following topics:

- [Machine sets that deploy machines with ultra disks as data disks](#)
- [Machine sets that deploy machines with ultra disks using CSI PVCs](#)
- [Machine sets that deploy machines with ultra disks using in-tree PVCs](#)

#### 1.3.15.3. Configuration options for Google Cloud Platform persistent disk types

The **pd-balanced** persistent disk type for the Google Cloud Platform (GCP) Compute Engine is now supported. For more information, see [Configuring persistent disk types by using machine sets](#) .

#### 1.3.15.4. Machine API support for Nutanix clusters



The new platform support for Nutanix clusters includes the ability to manage machines using Machine API machine sets. For more information, see [Creating a machine set on Nutanix](#).

### 1.3.15.5. Managing machines with the Cluster API (Technology Preview)

OpenShift Container Platform 4.11 introduces the ability to manage machines by using the upstream Cluster API, integrated into OpenShift Container Platform, as a Technology Preview for AWS and GCP clusters. This capability is in addition or an alternative to managing machines with the Machine API. For more information, see [Managing machines with the Cluster API](#).

## 1.3.16. Machine Config Operator

### 1.3.16.1. MCO now updates nodes by zone and age

The Machine Config Operator (MCO) now updates the affected nodes alphabetically by zone, based on the **topology.kubernetes.io/zone** label. If a zone has more than one node, the oldest nodes are updated first. For nodes that do not use zones, such as in bare metal deployments, the nodes are upgraded by age, with the oldest nodes updated first. Previously, the MCO did not consider zones or node age.

For more information, see [Machine config overview](#).

### 1.3.16.2. Enhanced notification for paused Machine Config Pools upon certificate renewal

You will now receive alerts in the Alerting UI of the OpenShift Container Platform web console if the MCO attempts to renew an expired **kube-apiserver-to-kubelet-signer** CA certificate on a machine config pool (MCP) that is paused. If the MCPs are paused, the MCO cannot push the newly rotated certificates to those nodes, which can result in failures.

For more information, see [Pausing the machine config pools](#).

## 1.3.17. Nodes

### 1.3.17.1. Self Node Remediation Operator replaces the Poison Pill Operator

OpenShift Container Platform 4.11 introduces the Self Node Remediation Operator that replaces the Poison Pill Operator.

The Self Node Remediation Operator provides the following enhancements:

- Introduces separate remediation templates based on the remediation strategy.
- Captures last error message if remediation is unsuccessful.
- Improves metrics for Self Node Remediation Operator's configuration parameters by providing minimum values for the parameters.

For more information, see [Remediating nodes with the Self Node Remediation Operator](#).

### 1.3.17.2. Worker nodes for single-node OpenShift clusters

You can now add worker nodes to single-node OpenShift clusters. This is useful for deployments in resource-constrained environments or at the network edge when you need to add additional capacity to your cluster.

For more information, see [Worker nodes for single-node OpenShift clusters](#).

### 1.3.17.3. Descheduler now defaults to simulating pod evictions

By default, the descheduler now runs in predictive mode, which means that it only simulates pod evictions. You can review the descheduler metrics to view details about pods that would be evicted.

To evict pods instead of simulating the evictions, change the descheduler mode to automatic.

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

### 1.3.17.4. New descheduler customizations

This release introduces the following customizations for the descheduler:

- Priority threshold filtering: Set the priority threshold either by class name (**thresholdPriorityClassName**) or by numeric value (**thresholdPriority**) to not evict pods that have a priority that is equal to or greater than that value.
- Namespace filtering: Set a list of user-created namespaces to include or exclude from descheduler operations. Note that protected namespaces (**openshift-\***, **kube-system**, **hypershift**) are always excluded.
- Thresholds for the **LowNodeUtilization** strategy: Set experimental thresholds for underutilization and overutilization for the **LowNodeUtilization** strategy.

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

### 1.3.17.5. Node Maintenance Operator enhancements

The Node Maintenance Operator provides the following enhancements:

- Additional feedback, **drainProgress** and **lastUpdate**, is now provided regarding the status of **NodeMaintenance** CR tasks.
- For clusters with bare-metal nodes, an easier method is now available on the web console where you can place a node into maintenance mode, and resume a node from maintenance mode.

For more information, see [Using the Node Maintenance Operator to place nodes in maintenance mode](#).

## 1.3.18. Logging

### 1.3.18.1. Red Hat OpenShift on RHV Logging (Technology Preview)

OpenShift Container Platform 4.11 introduces a new connector for the RHV API that adds automated log messages for all installations and oVirt components in a cluster.

## 1.3.19. Monitoring


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

### 1.3.19.1. Updates to monitoring stack components and dependencies

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

- Alertmanager to 0.24.0
- kube-state-metrics to 2.5.0
- Prometheus to 2.36.2
- Prometheus operator to 0.57.0
- Thanos to 0.26.0

### 1.3.19.2. Changes to alerting rules



#### NOTE

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

- **New**
  - Added the **KubePersistentVolumeInodesFillingUp** alert, which functions similarly to the existing **KubePersistentVolumeFillingUp** alert but is applied to inodes rather than volume space.
  - Added the **PrometheusScrapeBodySizeLimitHit** alert to detect targets hitting the body size limit.
  - Added the **PrometheusScrapeSampleLimitHit** alert to detect targets hitting the sample limit.
- **Changed**
  - Fixed the **KubeDaemonSetRolloutStuck** alert to use the updated metric **kube\_daemonset\_status\_updated\_number\_scheduled** from **kube-state-metrics**.
  - Replaced the **KubeJobCompletion** alert with **KubeJobNotCompleted**. The new **KubeJobNotCompleted** alert avoids false positives when an earlier job failed but the most recent job succeeded.
  - Updated the **NodeNetworkInterfaceFlapping** alert to exclude **tunbr** interfaces from the alert expression.

### 1.3.19.3. Enable alert routing for user workload monitoring

A cluster administrator can now enable alert routing for user workload monitoring so that developers and other users can configure custom alerts and alert routing for their user-defined projects.

### 1.3.19.4. Enable a dedicated Alertmanager instance for user-defined alerts

You now have the option to enable a separate instance of Alertmanager dedicated to sending alerts only for user-defined projects. This feature can help reduce the load on the default platform Alertmanager instance and can better separate user-defined alerts from default platform alerts.

### 1.3.19.5. Use additional authentication settings for remote write configuration

You can now use the following authentication methods to access a remote write endpoint: AWS Signature Version 4, custom Authorization header, and OAuth 2.0. Before this release, you could only use TLS client and basic authentication.

### 1.3.19.6. Create, browse, and manage PromQL queries more easily in the web console

The Query Browser on the **Observe** → **Metrics** page of the OpenShift Container Platform web console adds various enhancements to improve your ability to create, browse, and manage PromQL queries. For example, administrators can now duplicate existing queries and use autocomplete suggestions when building and editing queries.

### 1.3.19.7. Scrape interval doubled for ServiceMonitors in single node deployments

The scrape interval has been doubled for all Cluster Monitoring Operator (CMO) controlled ServiceMonitors on single-node OpenShift Container Platform deployments. The maximum interval is now two minutes.

### 1.3.19.8. Create alerting rules based on platform monitoring metrics (Technology Preview)

This release introduces a Technology Preview feature in which administrators can create alerting rules based on existing platform monitoring metrics. This feature helps administrators to more quickly and easily create new alerting rules specific to their environments.

### 1.3.19.9. Add cluster ID labels to remote write storage

You can now add cluster ID labels to metrics being sent to remote write storage. You can then query these labels to identify the source cluster for a metric and distinguish that metric data from similar metric data sent by other clusters.

### 1.3.19.10. Query metrics using the federation endpoint for user workload monitoring

You can now use the Prometheus **/federate** endpoint to scrape user-defined metrics from a network location outside the cluster. Before this release, you could only access the federation endpoint to scrape metrics in default platform monitoring.

### 1.3.19.11. Enable body size limit for metrics scraping for default platform monitoring

You can now set the **enforcedBodySizeLimit** config map option for default platform monitoring to enable a body size limit on metrics scraping. This setting triggers the new **PrometheusScrapeBodySizeLimitHit** alert when at least one Prometheus scrape target replies with a response body larger than the configured **enforcedBodySizeLimit**. The setting can limit the impact that a malicious target can have on both the Prometheus component and on the cluster as a whole.

### 1.3.19.12. Configure retention size settings for metrics storage

You can now configure the maximum amount of disk space reserved for retained metrics storage for both default platform monitoring and user-workload monitoring. Before this release, you could not configure this setting.

### 1.3.19.13. Configure the retention time period for Thanos Ruler in user-defined projects

You can now configure the retention time period for Thanos Ruler data in user-defined projects. Before this release, you could not change the default value of **24h**.

## 1.3.20. Scalability and performance

### 1.3.20.1. Workload hints for the Node Tuning Operator

OpenShift Container Platform 4.11 supports a hinting mechanism for the Node Tuning Operator that can tune **PerformanceProfile** to meet the demands of different industry environments. In this release, workload hints are available for **highPowerConsumption** (very low latency at the cost of increased power consumption) and **realtime** (priority given to optimum latency). A combination of true/false settings for these hints can be used to deal with application-specific workload profiles and requirements.

### 1.3.20.2. Enhancement to scale-up operations for etcd clusters

The Raft algorithm allows for scaling of etcd members using a new **learner** state. As a result, maintaining cluster quorum, adding and removing new members, and promoting **learners** occur without disrupting the cluster operation.

OpenShift Container Platform 4.11 introduces the **imageStorage** option as part of the **AgentServiceConfig** custom resource. This option improves application performance by allowing the user to specify persistent storage claim details for use with the image service.

The **releaseImage** parameter of the **ClusterImageSet** custom resource now supports operating system image version identification. The discovery ISO is based on an operating system image version as the releaseImage, or the latest version if the specified version is unavailable.

The **openshiftVersion** parameter of the **AgentServiceConfig** custom resource (CR) now supports either "x.y" (major.minor) or "x.y.z" (major.minor.patch) formats.

### 1.3.20.3. Configuring Ingress Controller (router) Liveness, Readiness, and Startup probes

OpenShift Container Platform 4.11 introduces the ability to configure the timeout values for the kubelet's liveness, readiness, and startup probes for router deployments that are managed by the OpenShift Container Platform ingress operator. The ability to set larger timeout values can reduce the risk of unnecessary and unwanted restarts that are caused by short default timeout of 1 second.

For more information see, [Configuring Ingress Controller \(router\) Liveness, Readiness, and Startup probes](#)

### 1.3.20.4. New power reduction CPU capability

You can decrease power consumption through the Node Tuning Operator by specifying CPUs in the **offlined** field in the performance profile. For more information, see [Reducing power consumption by taking CPUs offline](#).

### 1.3.20.5. Node Observability Operator (Technology Preview)

OpenShift Container Platform 4.11 introduces the Node Observability Operator in Technology Preview.

The Node Observability Operator provides the ability to:

- Deploy Node Observability Agents on the worker nodes.
- Trigger a CRI-O and Kubelet profiling.
- Make the profiling data files available for further analysis.

For more information, see [Requesting CRI-O and Kubelet profiling data using the Node Observability Operator](#).

### 1.3.20.6. Performance Addon Operator functions moved to the Node Tuning Operator

In earlier versions of OpenShift Container Platform, the Performance Addon Operator provided automatic, low latency performance tuning for applications. In OpenShift Container Platform 4.11, these functions are part of the Node Tuning Operator. The Node Tuning Operator is part of the standard installation for OpenShift Container Platform 4.11. If you upgrade to OpenShift Container Platform 4.11, the Node Tuning Operator removes the Performance Addon Operator and all related artifacts on startup.

For more information, see [Node Tuning Operator](#).



#### NOTE

You must still use the **performance-addon-operator-must-gather** image when running the **must-gather** command with the Performance Profile Creator. For more information, see [Gathering data about your cluster using must-gather](#)

### 1.3.20.7. Low latency tuning documentation updates

In earlier version of OpenShift Container Platform, documentation for low latency tuning included references to the Performance Addon Operator. Since the Node Tuning Operator now provides low latency tuning, the documentation title changed from "Performance Addon Operator for low latency nodes" to "Low latency tuning", and multiple cross-references to this document were updated accordingly. For more information, see [Low latency tuning](#).

### 1.3.20.8. Hub and spoke cluster support

For hub and spoke deployments in which spoke clusters require out-of-tree driver support, you can use the Special Resource Operator (SRO) deployed in the hub cluster to manage the deployment of the required kernel module(s) to one or more managed clusters. This uses Red Hat Advanced Cluster Management (RHACM) and no longer requires the use of Node Feature Discovery (NFD). For more information, see [Building and running the simple-kmod SpecialResource for a hub-and-spoke topology](#).

### 1.3.20.9. Enhanced SRO cluster upgrade support

When upgrading a cluster in which special resources are being managed, you can run the pre-upgrade custom resource to verify that a new driver container exists to support a kernel update. This helps avoid possible disruption of managed special resources. The documentation for this feature is currently unavailable and is targeted for release at a later date.

### 1.3.20.10. Enhanced debugging and logging for SRO

The Special Resource Operator (SRO) includes a consistent log output format with greater message detail for troubleshooting.

### 1.3.20.11. Support for external registries

Before this update, SRO did not support connecting to registries in disconnected environments. This release provides support for disconnected environments for which the driver container is hosted in a registry outside of the OpenShift Container Platform cluster.

## 1.3.21. Insights Operator

### 1.3.21.1. Insights Operator data collection enhancements

In OpenShift Container Platform 4.11, the Insights Operator collects the following additional information:

- The **images.config.openshift.io** resource definition
- **kube-controller-manager** container logs when the **"Internal error occurred: error resolving resource"** or **"syncing garbage collector with updated resources from discovery"** error messages are present
- **storageclusters.ocs.openshift.io/v1** resources

With this additional information, Red Hat improves OpenShift Container Platform functionality and enhances Insights Advisor recommendations.

## 1.3.22. Authentication and authorization

### 1.3.22.1. Additional supported OIDC providers

The following OpenID Connect (OIDC) providers are now tested and supported with OpenShift Container Platform:

- Active Directory Federation Services for Windows Server



#### NOTE

Currently, it is not supported to use Active Directory Federation Services for Windows Server with OpenShift Container Platform when custom claims are used.

- Microsoft identity platform (Azure Active Directory v2.0)



#### NOTE

Currently, it is not supported to use Microsoft identity platform when group names are required to be synced.

For the full list of OIDC providers, see [Supported OIDC providers](#).

### 1.3.22.2. Pod security admission

[Pod security admission](#) is now enabled on OpenShift Container Platform.

Pod admission is enforced by both pod security and security context constraints (SCC) admissions. Pod security admission runs globally with **privileged** enforcement and **restricted** audit logging and API warnings.

A controller monitors the [SCC-related permissions](#) of service accounts in user-created namespaces and automatically labels these namespaces with [pod security admission](#) **warn** and **audit** labels.

To improve workload security in accordance with the **restricted** pod security profile, this release introduces SCCs that enforce pod security in accordance with new pod security admission controls. These SCCs are:

- **restricted-v2**
- **hostnetwork-v2**
- **nonroot-v2**

They correspond to older, similarly named SCCs, but with the following enhancements:

- **ALL** capabilities are dropped from containers. Previously, only the **KILL**, **MKNOD**, **SETUID**, and **SETGID** capabilities were dropped.
- The **NET\_BIND\_SERVICE** capability can now be added explicitly.
- **seccompProfile** is defaulted to **runtime/default** if unset. In previous releases, this field had to be empty.
- **allowPrivilegeEscalation** must be unset or set to **false** in security contexts. Previously, a **true** value was allowed.

In OpenShift Container Platform 4.11, the **restricted-v2** SCC is now the default granted SCC available to users for new installations, instead of the **restricted** SCC. In new clusters, the **restricted-v2** SCC is used for any authenticated user in place of the **restricted** SCC. The **restricted** SCC is no longer available to users of new clusters, unless the access is explicitly granted. In clusters originally installed in OpenShift Container Platform 4.10 or earlier, all authenticated users can use the **restricted** SCC when upgrading to OpenShift Container Platform 4.11 and later. This keeps the default security permissions for OpenShift Container Platform as secure-by-default and aligns with pod security admission and pod security standards from the upstream Kubernetes project, while keeping the permissions of upgraded clusters consistent with their previous state.

You can enable synchronization for most namespaces and disable synchronization for all namespaces.

For this release, namespaces that are prefixed with **openshift-** do not have restricted enforcement. Restricted enforcement for such namespaces is planned for inclusion in a future release.

For more information, see [Understanding and managing pod security admission](#) .

## 1.4. NOTABLE TECHNICAL CHANGES

OpenShift Container Platform 4.11 introduces the following notable technical changes.

### Network Observability operator for observing network flows

The Network Observability Operator is General Availability (GA) status in the 4.12 release of OpenShift Container Platform and is also supported in OpenShift Container Platform 4.11.

For more information, see [Network Observability](#).

### Update in default value for setting the router load-balancing algorithm

The **haproxy.router.openshift.io/balance** variable, which sets the router load-balancing algorithm, now defaults to the value **random** instead of **leastconn**. See [Route-specific annotations](#) for more information.

**LegacyServiceAccountTokenNoAutoGeneration** is on by default



In previous releases, two service account token secrets were generated when a service account was created:

- A service account token secret for authenticating to the internal OpenShift Container Platform registry
- A service account token secret for accessing the Kubernetes API

Starting with OpenShift Container Platform 4.11, this second service account token secret for accessing the Kubernetes API is no longer created. This is because the

**LegacyServiceAccountTokenNoAutoGeneration** upstream Kubernetes feature gate was enabled in Kubernetes 1.24, which stops the automatic generation of secret-based service account tokens for accessing the Kubernetes API. After upgrading to OpenShift Container Platform 4.11, any existing service account token secrets are not deleted and continue to function.



## NOTE

Do not rely on these automatically generated secrets for your own use; they might be removed in a future OpenShift Container Platform release.

Workloads are automatically injected with a projected volume to obtain a bound service account token. If your workload needs an additional service account token, add an additional projected volume in your workload manifest. For more information, see [Using bound service account tokens](#).

You can also manually create a service account token secret to obtain a token, if the security exposure of a non-expiring token in a readable API object is acceptable to you. For more information, see [Creating a service account token secret](#).

## Operator SDK 1.22.2

OpenShift Container Platform 4.11 supports Operator SDK 1.22.2. See [Installing the Operator SDK CLI](#) to install or update to this latest version.



## NOTE

Operator SDK 1.22.2 supports Kubernetes 1.24.

If you have Operator projects that were previously created or maintained with Operator SDK 1.16.0, update your projects to keep compatibility with Operator SDK 1.22.2.

- [Updating Go-based Operator projects](#)
- [Updating Ansible-based Operator projects](#)
- [Updating Helm-based Operator projects](#)
- [Updating Hybrid Helm-based Operator projects](#)

## Cluster Operators no longer referred to as platform Operators

OpenShift Container Platform documentation previously referred to cluster Operators interchangeably with the alternative naming "platform Operators". Because this double naming could cause confusion about the type of Operators being discussed, the term "platform Operator" is no longer used when referring to cluster Operators, which are represented by **ClusterOperator** API objects. The documentation sets for OpenShift Container Platform 4.11 and earlier have been updated to now solely use the term "cluster Operator".

For example, see [Cluster Operators reference](#).

## 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 4.11, refer to the table below. Additional details for more 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*
- **DEP:** *Deprecated*
- **REM:** *Removed*

**Table 1.1. Deprecated and removed features tracker**

Feature	OCP 4.9	OCP 4.10	OCP 4.11
SQLite database format for Operator catalogs	DEP	DEP	DEP
<b>ImageChangesInProgress</b> condition for Cluster Samples Operator	DEP	DEP	DEP
<b>MigrationInProgress</b> condition for Cluster Samples Operator	DEP	DEP	DEP
Cluster Loader	DEP	REM	REM
Bring your own RHEL 7 compute machines	DEP	REM	REM
Jenkins Operator	DEP	REM	REM
Grafana component in monitoring stack	-	DEP	REM
Access to Prometheus and Grafana UIs in monitoring stack		DEP	REM
vSphere 6.7 Update 2 or earlier	DEP	DEP	REM
vSphere 7.0 Update 1 or earlier	-	-	DEP
Virtual hardware version 13	DEP	DEP	REM
VMware ESXi 6.7 Update 2 or earlier	DEP	DEP	REM
VMware ESXi 7.0 Update 1 or earlier	-	-	DEP

Feature	OCP 4.9	OCP 4.10	OCP 4.11
Snapshot.storage.k8s.io/v1beta1 API endpoint	DEP	DEP	REM
Minting credentials for Microsoft Azure clusters	GA	REM	REM
Persistent storage using FlexVolume	-	DEP	DEP
Automatic generation of service account token secrets	GA	GA	REM
Removal of Jenkins images from install payload	GA	GA	REM
Multicluster console (Technology Preview)	-	REM	REM

### 1.5.1. Deprecated features

#### 1.5.1.1. OpenShift CLI (oc) commands and flags for requesting tokens are deprecated

The following **oc** commands and flags for requesting tokens are now deprecated:

- The **oc serviceaccounts create-kubeconfig** command
- The **oc serviceaccounts get-token** command
- The **oc serviceaccounts new-token** command
- The **--service-account/-z** flag for the **oc registry login** command

Use the **oc create token** command instead to request tokens.

#### 1.5.1.2. Red Hat Virtualization (RHV) as a host platform for OpenShift Container Platform will be deprecated

Red Hat Virtualization (RHV) will be deprecated in an upcoming release of OpenShift Container Platform. Support for OpenShift Container Platform on RHV will be removed from a future OpenShift Container Platform release, currently planned as OpenShift Container Platform 4.14.

#### 1.5.1.3. Support for vSphere 7.0 Update 1 or earlier is deprecated

In OpenShift Container Platform 4.11, support for VMware vSphere 7.0 Update 1 or earlier is deprecated. While vSphere 7.0 Update 1 or earlier remains fully supported, Red Hat recommends that you use vSphere 7.0 Update 2 or later, up to but not including version 8. vSphere 8 is not supported.

#### 1.5.1.4. Support for ESXi 7.0 Update 1 or earlier is deprecated

In OpenShift Container Platform 4.11, support for VMware ESXi 7.0 Update 1 or earlier is deprecated. While ESXi 7.0 Update 1 or earlier remains fully supported, Red Hat recommends that you use ESXi 7.0 Update 2 or later.

#### 1.5.1.5. Support for `pidsLimit` and `logSizeMax` CRI-O parameters will be deprecated

In OpenShift Container Platform 4.11, the **pidsLimit** and **logSizeMax** fields in the **ContainerRuntimeConfig** CR will be deprecated and removed in a future release. Use the **podPidsLimit** and **containerLogMaxSize** fields in the **KubeletConfig** CR instead. The default value of the **podPidsLimit** field is **4096**.

## 1.5.2. Removed features

### 1.5.2.1. RHEL 7 support for the OpenShift CLI (oc) has been removed

Support for using Red Hat Enterprise Linux (RHEL) 7 with the OpenShift CLI (**oc**) has been removed. If you use the OpenShift CLI (**oc**) with RHEL, you must use RHEL 8 or later.

### 1.5.2.2. OpenShift CLI (oc) commands have been removed

The following OpenShift CLI (**oc**) commands were removed with this release:

- **oc adm migrate etcd-ttl**
- **oc adm migrate image-references**
- **oc adm migrate legacy-hpa**
- **oc adm migrate storage**

### 1.5.2.3. Grafana component removed from monitoring stack

The Grafana component is no longer a part of the OpenShift Container Platform 4.11 monitoring stack. As an alternative, go to **Observe** → **Dashboards** in the OpenShift Container Platform web console to view monitoring dashboards.

### 1.5.2.4. Prometheus and Grafana user interface access removed from monitoring stack

Access to the third-party Prometheus and Grafana user interfaces have been removed from the OpenShift Container Platform 4.11 monitoring stack. As an alternative, click **Observe** in the OpenShift Container Platform web console to view alerting, metrics, dashboards, and metrics targets for monitoring components.

### 1.5.2.5. Support for virtual hardware version 13 is removed

In OpenShift Container Platform 4.11, support for virtual hardware version 13 is removed. Support for virtual hardware version 13 was deprecated in OpenShift Container Platform 4.9. Red Hat recommends that you use virtual hardware version 15 or later.

### 1.5.2.6. Support for vSphere 6.7 Update 2 or earlier is removed

In OpenShift Container Platform 4.11, support for VMware vSphere 6.7 Update 2 or earlier is removed. Support for vSphere 6.7 Update 2 or earlier was deprecated in OpenShift Container Platform 4.9. Red Hat recommends that you use vSphere 7.0 Update 2 or later, up to but not including version 8. vSphere 8 is not supported.

### 1.5.2.7. Support for ESXi 6.7 Update 2 or earlier is removed

In OpenShift Container Platform 4.11, support for VMware ESXi 6.7 Update 2 or earlier is removed. Support for ESXi 6.7 Update 2 or earlier was deprecated in OpenShift Container Platform 4.10. Red Hat recommends that you use ESXi 7.0 Update 2 or later.

#### 1.5.2.8. Support for snapshot v1beta1 API endpoint is removed

In OpenShift Container Platform 4.11, support for **snapshot.storage.k8s.io/v1beta1** API endpoint is removed. Support for **snapshot.storage.k8s.io/v1beta1** API endpoint was deprecated in OpenShift Container Platform 4.7. Red Hat recommends that you use **snapshot.storage.k8s.io/v1**. All objects created as **v1beta1** are available through the v1 endpoint.

#### 1.5.2.9. Support for manually deploying a custom scheduler has been removed

Support for deploying custom schedulers manually has been removed with this release. Use the [Secondary Scheduler Operator for Red Hat OpenShift](#) instead to deploy a custom secondary scheduler in OpenShift Container Platform.

#### 1.5.2.10. Support for deploying single-node OpenShift with OpenShiftSDN has been removed

Support for deploying single-node OpenShift clusters with OpenShiftSDN has been removed with this release. OVN-Kubernetes is the default networking solution for single-node OpenShift deployments.

#### 1.5.2.11. Removal of Jenkins images from install payload

- OpenShift Container Platform 4.11 moves the "OpenShift Jenkins" and "OpenShift Agent Base" images to the **ocp-tools-4** repository at **registry.redhat.io** so that Red Hat can produce and update the images outside the OpenShift Container Platform lifecycle. Previously, these images were in the OpenShift Container Platform install payload and the **openshift4** repository at **registry.redhat.io**. For more information, see [OpenShift Jenkins](#).
- OpenShift Container Platform 4.11 removes "OpenShift Jenkins Maven" and "NodeJS Agent" images from its payload. Previously, OpenShift Container Platform 4.10 deprecated these images. Red Hat no longer produces these images, and they are not available from the **ocp-tools-4** repository at **registry.redhat.io**.  
However, upgrading to OpenShift Container Platform 4.11 does not remove "OpenShift Jenkins Maven" and "NodeJS Agent" images from 4.10 and earlier releases. And Red Hat provides bug fixes and support for these images through the end of the 4.10 release lifecycle, in accordance with the [OpenShift Container Platform lifecycle policy](#).

For more information, see [OpenShift Jenkins](#).

### 1.5.3. Future Kubernetes API removals

The next minor release of OpenShift Container Platform is expected to use Kubernetes 1.25. Currently, Kubernetes 1.25 is scheduled to remove several deprecated **v1beta1** and **v2beta1** APIs.

See the [Deprecated API Migration Guide](#) in the upstream Kubernetes documentation for the list of planned Kubernetes API removals.

See [Navigating Kubernetes API deprecations and removals](#) for information about how to check your cluster for Kubernetes APIs that are planned for removal.

## 1.6. BUG FIXES

## Bare Metal Hardware Provisioning

- Previously, when writing the RHCOS image onto some disks, the **qemu-img** was allocating space onto the entire disk, including sparse areas. This extended the time of the writing process on some hardware. This update disables the **qemu-img** sparse in image creation. As a result, the image writing would no longer take a long time on affected hardware. ([BZ#2002009](#))
- Previously, if the **rotational** field was set for **RootDeviceHints**, the host could fail the provision. With this update, the **rotational** field in **RootDeviceHints** is properly copied and checked. As a result, the provisioning will succeed when using the **rotational** field. ([BZ#2053721](#))
- Previously, Ironic was unable to use virtual media to provision Nokia OE 20 servers because the BMC required the **TransferProtocolType** attribute to be explicitly set in the request despite this being an optional attribute. Additionally, the BMC also required the use of a dedicated **RedFish** settings resource to override boot orders, whereas most BMCs just use the **system** resource. This error occurred because Nokia OE 20 strictly requires an optional **TransferProtocolType** attribute for vMedia attachments and requires the use of the **RedFish** settings resource for overriding boot sequences. Consequently, virtual media based provisioning would fail on Nokia OE 20. There are two workarounds for this issue:
  1. When the vMedia attachment request fails with an error indicating that the **TransferProtocolType** attribute is missing, retry the request and explicitly specify this attribute.
  2. Check for the presence of the RedFish settings resource in the system. If it is present, use it for the boot sequence override.

As a result to these workarounds, virtual media based provisioning will succeed on Nokia OE 20 machines. ([BZ#2059567](#))

- Previously, the Ironic API inspector image failed to clean disks that were part of passive multipath setups when using OpenShift Container Platform bare-metal IPI deployments. This update fixes the failures when active or passive storage arrays are in use. As a result, it is now possible to use OpenShift Container Platform bare metal IPI when customers want to use multipath setups that are active or passive. ([BZ#2089309](#))
- Previously, Ironic failed to match **wwn** serial numbers to multi-path devices. Consequently, **wwn** serial numbers for device mapper devices could not be used in the **rootDeviceHint** parameter in the **install-config.yaml** configuration file. With this update, Ironic now recognizes **wwn** serial numbers as unique identifiers for multi-path devices. As a result, it is now possible to use **wwn** serial numbers for device mapper devices for the **install-config.yaml** file. ([BZ#2098392](#))
- Before this update, when a Redfish system features a Settings URI, the Ironic provisioning service always attempts to use this URI to make changes to boot-related BIOS settings. However, bare-metal provisioning fails if the Baseboard Management Controller (BMC) features a Settings URI but does not support changing a particular BIOS setting by using this Settings URI. In OpenShift Container Platform 4.11 and later, if a system features a Settings URI, Ironic verifies that it can change a particular BIOS setting by using the Settings URI before proceeding. Otherwise, Ironic implements the change by using the System URI. This additional logic ensures that Ironic can apply boot-related BIOS setting changes and bare-metal provisioning can succeed. ([OCPBUGS-2052](#))

## Builds

- Previously, using a forward slash (/) in the **ImageLabel** name of a **BuildConfig** instance resulted in an error. This fix resolves the issue by changing the utility used for validation. As a result, you can use a forward slash in the **ImageLabel** name of a **BuildConfig** instance. ([BZ#2105167](#))

- Previously, when using the **\$ oc new-app --search <image\_stream\_name>** command, you could receive an incorrect message related to **docker.io** images. This resulted in confusion for the user because OpenShift Container Platform does not use image streams that point to **docker.io**. This fix adds code checks to prevent the reference to **docker.io**. As a result, the output from that command does not include the message. ([BZ#2049889](#))
- Previously, Shared Resource CSI Driver metrics were not exported to the Telemetry service. As a result, usage metrics for the Shared Resource CSI Driver could not be analyzed. With this fix, Shared Resource CSI Driver metrics are exposed to the Telemetry service. As a result, usage metrics for the Shared Resource CSI Driver can be collected and analyzed. ([BZ#2058225](#))
- By default, Buildah prints steps to the log file, including the contents of environment variables, which might include [build input secrets](#). Although you can use the **--quiet** build argument to suppress printing of those environment variables, this argument isn't available if you use the source-to-image (S2I) build strategy. The current release fixes this issue. To suppress printing of environment variables, set the **BUILDAH\_QUIET** environment variable in your build configuration:

```
sourceStrategy:
...
env:
- name: "BUILDAH_QUIET"
  value: "true"
```

- Before this update, using the **\$ oc new-app --search <image\_stream\_name>** command provided a warning that you may not have access to the container image **"docker.io/library/<image\_name>:<tag>"**. This caused confusion about OpenShift Container Platform having image streams that pointed to **docker.io**. This update fixes the issue by adding code checks to prevent that confusing reference to 'docker.io'. Now, the output from that command does not include messages about **docker.io** ([BZ#2049889](#))

## Cloud Compute

- **CertificateSigningRequest** (CSR) resource renewal is handled by the Kubernetes controller manager and correctly left pending by the Cluster Machine Approver Operator, which increases the value of the **mapi\_current\_pending\_csr** metric to **1**. Previously, when the Kubernetes controller manager approved the CSR, the Operator ignored it and left the metric unchanged. As a result, the **mapi\_current\_pending\_csr** metric was stuck at **1** until the Operator next reconciled. With this release, CSR approvals from other controllers are always reconciled to update metrics and the value of the **mapi\_current\_pending\_csr** metric is updated after every reconcile. ([BZ#2047702](#))
- Previously, only regions contained in a list of known regions within the AWS SDK were validated, and specifying any other region caused an error. This meant that, as new regions were added, they could not be used until the SDK was updated to contain the new region information. With this release, regions are validated with a less strict setting that warns the user when a region is not recognized. As a result, new regions might cause warning messages but can be used immediately. ([BZ#2065510](#))
- Previously, the Cluster Machine Approver Operator appended the **"Approved"** status condition to its condition list. As a result, the Kubernetes API server logged errors containing the message **[SHOULD NOT HAPPEN] failed to update managedFields**. With this release, the Operator is updated to check its conditions before appending to the list, and only update the condition when necessary. As a result, the conditions are no longer duplicated in the **CertificateSigningRequest** resource and the Kubernetes API server no longer logs errors about the duplicate. ([BZ#1978303](#))

- Previously, a defect in the Cisco ACI neutron implementation that was present in Red Hat OpenStack Platform (RHOSP) version 16, caused the query for subnets belonging to a given network to return unexpected results. As a result, the RHOSP Cluster API provider might try to provision instances with duplicated ports on the same subnet, leading to a failed provisioning. With this release, additional filtering in the RHOSP Cluster API provider ensures there is no more than one port per subnet, and it is now possible to deploy OpenShift Container Platform on RHOSP version 16 with Cisco ACI. ([BZ#2033862](#))
- Previously, the Red Hat OpenStack Platform (RHOSP) Machine API provider did not use the proxy environment variable directives, causing installation behind an HTTP or HTTPS proxy to fail. With this release, the provider obeys proxy directives and functions correctly in a restricted environment in which egress traffic is only allowed through a proxy. ([BZ#2046133](#))
- Previously, when upgrading from OpenShift Container Platform 4.9 to 4.10, inconsistencies between multiple controllers resulted in incorrect version numbers. As a result, the version number was not consistent. With this release, consistent reading of version numbers has been enacted and the release version is now stable in the cluster operator status. ([BZ#2059716](#))
- Previously, leaks of load balancer targets on AWS Machine API providers may occur. This is because IP-based load balancer attachments may remain within the load balancer registration when replacing control plane machines. With this release, IP-based load balancer attachments are removed from the load balancer before the Amazon EC2 instance is removed from AWS. As a result, the leaks are avoided. ([BZ#2065160](#))
- Previously, during an upgrade, a new machine created through the Machine API defaulted to HW-13 which caused the cluster to degrade. With this release, the machine controller checks the virtual machine's hardware version during machine creation from the template clone. The machine goes into a **failed** state if the template's hardware version is less than 15, which is the minimal supported hardware version for OpenShift Container Platform 4.11 and above versions. ([BZ#2059338](#))
- Previously, the procedural name generator for Azure availability sets exceeded the 80 character maximum limit. This might cause the Machine API to reuse the same sets during the name truncation, rather than creating multiple availability sets. With this release, the procedural name generator is updated to ensure that the name is no longer than 80 characters and the cluster name is not duplicated in the set name. As a result, Azure availability sets are no longer truncated in unexpected ways by the procedural name generator. ([BZ#2093044](#))
- Because the Cluster Autoscaler Operator was deploying the cluster autoscaler without setting any leader election parameters, the cluster autoscaler could unexpectedly fail and restart after a cluster restart. With this fix, the Cluster Autoscaler Operator now deploys the cluster autoscaler with well-defined leader election flags. As a result, the cluster autoscaler operates as expected after restarts. ([BZ#2063194](#))
- Previously, certificate signing request (CSR) renewal was handled by **kube-controller-manager** and correctly left pending by the machine approver, which increased **mapi\_current\_pending\_csr** to **1**. Afterwards, the **kube-controller-manager** approved the CSR, but the machine approver would ignore it, which left the metric unchanged. Consequently, the **mapi\_current\_pending\_csr** was stuck at **1** until another machine approver reconciled it. With this update, CSR approvals are reconciled from other controllers to update metrics properly. As a result, **mapi\_current\_pending\_csr** is always up-to-date after every reconcile. ([BZ#2072195](#))
- Previously, the Machine API Operator did not report as degraded if an insufficient number of worker nodes started upon cluster installation, even though other Operators were reported as degraded. With this release, the Machine API Operator is now reported as degraded and errors



are posted in the installation log in this scenario. As a result, with the Machine API Operator appears in the list of failed Operators, and users now know to examine the state of the machines as the reason why there are insufficient worker nodes. ([BZ#1994820](#))

- Because the Machine API Operator did not honor the proxy environment variable directives, installation behind an HTTP or HTTPS proxy failed. With this fix, the HTTP transport logic of Machine API Operator now obeys proxy directives. As a result, the Machine API Operator now works in a restricted environment where egress traffic is only allowed through a proxy. ([BZ#2082667](#))
- Previously, installations on vSphere for clusters with thousands of tags and heavy API loads were failing. Now, the machine controllers only query tags that are related to a particular OpenShift Container Platform installation. As a result, OpenShift Container Platform installs correctly on vSphere for these clusters. ([BZ#2097153](#))
- Because the kubelet needs to contact the vCenter to obtain node zone labels, if the vCenter credentials are stored in a secret, the kubelet could not start because the kube client was not created on time. As a result, if you rebooted the nodes, as happens when you edit the **cloud-provider-config** config map, the nodes did not come up. With this fix, the kubelet now obtains zone labels after node registration if the credentials are stored in the secret. As a result, the nodes reboot as expected. ([BZ#1902307](#))
- Previously, the lack of a **timeout** option might have led to controller blockage, which could cause vCenter to never, or very slowly, respond. This release adds a **timeout** option for vCenter clients within the vSphere machine controller. ([BZ#2083237](#))

### Cluster Version Operator

- Previously, the Cluster Version Operator (CVO) when it encountered an error during an upgrade would stop reconciling current release manifests. With this update, release loading is separated from reconciling so the latter does not block the former and a new condition **ReleaseAccepted** has been added to clarify the status of the release load. ( [BZ#1822752](#))

### Console Metal3 Plugin

- Previously, there was a missing option in the UI to set the **bootMode** strategy. Consequently, the UI would always use the default (UEFI) boot strategy, which caused problems booting up some types of bare metals deployments. This update exposes a new field in the **Add Bare Metal** host form to choose the appropriate boot mode strategy. As a result, the bare metal machine properly starts. ([BZ#2091030](#))
- Previously, the node maintenance feature was moved to a new project which resulted in API changes. Consequently, node maintenance stopped working. This update fixes the code to work properly with new APIs. As a result, node maintenance works again. ([BZ#2090621](#))
- Previously, in some cases, day 2 workers on clusters created by the assisted installer are missing underlying bare metal hosts and machine resources. Consequently, the UI failed when trying to show details of the worker node. With this update, the bare metal hosts and machine resources are handled more gracefully, and the UI shows all available details. ([BZ#2090993](#))

### Domain Name System (DNS)

- Topology aware hints is a new feature in OpenShift Container Platform 4.11 that allows the **EndpointSlice** controller to specify hints to the container network interface (CNI) for how it should route traffic to a service's endpoints. The DNS operator was not enabling topology

aware hints for the cluster DNS service. Consequently, the CNI network providers did not keep DNS traffic local to the zone or node. With this fix, the DNS operator was updated to specify topology aware hints on the cluster DNS service. ([BZ#2095941](#))

- Previously, the kubelet created a default **/etc/resolv.conf** for pods based on the **/etc/resolv.conf** from the node host. Consequently, the malformed **resolv.conf** file could cause resolvers to fail to parse **resolv.conf**, therefore breaking DNS resolution in pods. With this update, the kubelet accepts a **resolv.conf** file and pods get a valid **resolv.conf** file. ([BZ#2063414](#))
- Previously, the DNS Operator did not set the **cluster-autoscaler.kubernetes.io/enable-ds-eviction** annotation on DNS pods. Consequently, the cluster autoscaler did not remove the DNS pod from a node before removing the node. With this update, the DNS operator was changed to add the **cluster-autoscaler.kubernetes.io/enable-ds-eviction** annotation to DNS pods. The DNS pod can now be removed from a node before removing the node by the cluster autoscaler. ([BZ#2061244](#))

## Image Registry

- Previously, the image registry used an **ImageContentSourcePolicy** (ICSP) as a source only when it was an exact match. The same source file was expected to be pulled through for any subrepositories. The ICSP name and path did not match for subrepositories. As a result, the image was not used. Now, the ICSP is applied successfully to the subrepositories, which can use the mirrored images. ([BZ#2014240](#))
- Previously, if the pruner failed, the image registry Operator is reported as degraded until the pruner successfully runs. With this update, the Operator is more resilient to the pruner failures. ([BZ#1990125](#))
- Previously, the registry used exact match when applying **ImageContentSourcePolicy**(ICSP). With this update, ICSP is applied for subrepositories and mirrors now work as expected. ([BZ#2014240](#))
- Previously, the Image Registry Operator did not work with AWS S3 on RHOSP. With this update, the Image Registry Operator trusts AWS S3 on all platforms. ([BZ#2007611](#))
- Previously, OpenShift Container Platform image registry did not work with Ceph Radosgw. With this update, the image registry works with Ceph Radosgw. ([BZ#1976782](#))
- Previously, the Image Registry Operator interpreted **429** error messages from upstream registries as if the data was not available. The Operator returned a **404 Not Found** message instead of **429 Too Many Requests**. With this update, the proper **429 Too Many Requests** message is returned so that administrators know to retry the request. ([BZ#1923536](#))
- Previously, the Image Registry Operator could not be used to configure CloudFront. With this update, the Image Registry Operator can configure CloudFront. ([BZ#2065224](#))
- Previously, the Image Registry Operator pushed images when KMS encryption was enabled, but not pulled. With this update, images can be pushed and pulled with KMS encryption enabled. ([BZ#2048214](#))
- Previously, the Image Registry Operator could not pull public images anonymously due to credentials not being provided. With this update, clients can pull public images anonymously. ([BZ#2060605](#))
- Previously, the Image Registry Operator could not use the **ap-southeast-3** AWS region. With this update, the registry can be configured to **ap-southeast-3**. ([BZ#2065552](#))

## Installer

- Previously, if users specified an OpenShift Container Platform cluster name with a period, the installation failed. This update adds a validation check to the installation program, which returns an error if a period is present in the cluster name. ([BZ#2084580](#))
- Previously, users could select the AWS **us-gov-east-1** region when using the installation program to create the **install-config.yaml** file. This caused the deployment to fail because the installation program could only be used to create the **install-config.yaml** file for a public AWS region. This update removes all of the AWS regions from the installation program that are not supported by the public AWS cloud. ([BZ#2048222](#))
- Previously, users could not select the **ap-north-east-3** region when using the installation program to create the **install-config.yaml** file. The AWS SDK that caused this problem has been updated, and users can now select the **ap-north-east-3** region. ([BZ#1996544](#))
- Previously, installing a private (internal) OpenShift Container Platform cluster on Azure Stack Hub failed because the installation program did not create the DNS record for the API virtual IP address. This update removes the invalid check that caused this problem. The installation program now correctly creates DNS records for private cluster. ([BZ#2061549](#))
- Previously, uninstalling an IBM Cloud VPC cluster might have caused unexpected results. When a user uninstalled a cluster (cluster 1), the DNS records of another cluster (cluster 2) were removed when either the name of cluster 1 (example) was a subset of the name of cluster 2 (myexample) or both clusters shared a base domain. This update corrects this behavior. Only those resources specific to the cluster that is being uninstalled are now removed. ([BZ#2060617](#))
- Previously, Azure Stack Hub did not support any disk types other than Standard\_LRS. This update adds the ability to customize disk types, which allows clusters to have a default disk type with no manual customizations. This resulted in making the switch from hard coding the disk type to accepting the inputs from the user and validating it against the Stack Hub APIs. ([BZ#2061544](#))
- Previously, when destroying a cluster, the ID of the private route5 hosted zone for the cluster was incorrectly reported when the DNS records were deleted from the hosted zone. This caused incorrect host zone IDs to be reported in the log of the destroyer. This update uses the correct host zone ID in the log. As a result, the log shows the correct hosted zone ID when destroying the DNS record in the base domain's hosted zone. ([BZ#1965969](#))
- Previously, system proxy settings were not considered when requesting an AWS custom service endpoint. This update configures AWS custom service endpoint validation to consider the system proxy settings along with a **HEAD** request to the AWS custom service endpoint. As a result, the AWS custom service endpoint can be accessed from the user's machine. ([BZ#2048451](#))
- Previously, the installation program used any Terraform provider in the **\$PATH** on the installer host. Therefore, the installation would fail if there were Terraform providers in the **\$PATH** that used an incorrect version or provider rather than the providers embedded in the installation program. With this update, the installation program embeds providers to a known directory and sets Terraform to use the known directory. As a result, installation will be successful because the installation program will always use the providers in the known directory. ([BZ#1932812](#))
- Previously, there was an eventual consistency issue in the AWS Terraform provider when updating new load balancers. Therefore, the installation would fail when trying to access the new load balancers. With this fix, the installation program is now updated to an upstream Terraform

provider, which ensures eventual consistency. As a result, the installation does not fail.

([BZ#1898265](#))

- Previously, the installation program had a list of required APIs that check for quota and permissions, and the list contained some unnecessary APIs that failed if the user did not provide the permissions. With this update, The list of APIs is split into **required** and **optional** where the **optional** APIs are not accessible. A warning message is displayed for the **optional** APIs. ([BZ#2084280](#))
- Previously, the **.apps** entry did not have the tag **kubernetes.io\_cluster<infraID>** that was used by the installation program to delete code from the database that would isolate all the resources created for a given cluster and delete them. With this update, a tag was added to the cluster Ingress Operator during creation time which makes the entry visible to delete. ([BZ#2055601](#))
- Previously, the **openshift-install** command would fail when using the internal publishing strategy. With this update, the **openshift-install** command no longer fails. ([BZ#2047670](#))
- Previously, there was an eventual consistency issue in the AWS Terraform provider when updating to newly created Virtual Private Clouds (VPCs). Consequently, the installation program would fail when attempting to access VPCs. With this update, the installation program is updated to the upstream Terraform provider and the install does not fail. ([BZ#2043080](#))
- Previously, there was an eventual consistency issue in the AWS Terraform provider when updating to newly created network interfaces. Consequently, installs would fail to access the network interfaces. With this update, the Terraform provider is updated to accept eventual consistency and installation does not fail. ([BZ#2047741](#))
- Previously, the **corespersocket** value could be higher than the **numCores** value when installing a VMware cluster using installer-provisioned infrastructure. This could cause unexpected results during installation. With this update, users receive a warning to correct these values before the cluster is created. ([BZ#2034147](#))
- Previously, a rare bug caused inconsistency during AWS cluster installations. The installation program has been updated to avoid this scenario. ([BZ#2046277](#))
- Previously, the number of supported user-defined tags was 8, and reserved OpenShift Container Platform tags were 2 for AWS resources. With this release, the number of supported user-defined tags is now 25 and reserved OpenShift Container Platform tags are 25 for AWS resources. You can now add up to 25 user tags during installation. ([CFE#592](#))
- Previously, the bootstrap machine used a default size and instance type when installing on Azure using installer-provisioned infrastructure. With this update, the bootstrap machine uses the size and instance type of the control plane machines. You can now control the size and instance type of the bootstrap machine by modifying the control plane settings in your installation configuration. ([BZ#2026356](#))
- Previously, the installation program provided Terraform with an ambiguous network name. This could lead to Terraform being unable to determine the correct network to use. With this update, the installation program provides Terraform with a unique network ID so that the installation succeeds. ([BZ#1918005](#))
- Previously, control plane machines could be created before the dependent NAT gateway was created when installing a cluster on AWS, causing the installation to fail. With this update, Terraform ensures that the NAT gateway has been created before the control plane machines are created. This ensures that the installation will succeed. ([BZ#2049108](#))

- Previously, if you used an OVN network rather than the default OSN network, the scale-up task failed because it took longer than the maximum amount of time required. This update doubles the amount of retries during the scale-up task so that the task can be completed. ([BZ#2090151](#))
- Previously, when you tried to delete multiple clusters from the database in parallel, the deletion process failed because of a bug in the **vmware** and **govmomi** libraries. The bug caused one of the cluster's tags to be deleted, which resulted in a 404 error when the deletion process could not find the tag. This update ignores tags that aren't found and continues the deletion process so that it finishes without error. ([BZ#2021041](#))
- With this update, OpenShift Container Platform on Red Hat Virtualization (RHV) supports preallocated disks for control plane and worker nodes for installer-provisioned infrastructure. In high-load environments, preallocated disks benefit performance for etcd and other components. ([BZ#2035334](#))
- Previously, when you tried to delete multiple clusters in parallel, the process failed because of a bug in the **vmware/govmomi** library. The bug caused a cluster tag to be deleted, which resulted in a **404** error when the deletion process could not find the tag. This update ignores tags that are not found and continues to delete so that it finishes without error. ([BZ#2021041](#))
- Previously, installation methods for VMware vSphere included validation that checked for network existence during the creation of configuration files. This caused an error for user-provisioned infrastructure and other installation methods where the network can be created as part of provisioning the infrastructure, in which case the network might not exist when the config files are generated. This fix updates the installation program to only perform network validation on installer-provisioned infrastructure installs. As a result, user-provisioned infrastructure and other installation methods can generate configuration files regardless of network existence. ([BZ#2050767](#))
- Previously, **runc** had an inversioned dependency on **libseccomp** 2.5 or later, which causes issues when the operating system was installed using version 8.3 or later and not fully updated to 8.4 or later. With this update, the RHEL host installs successfully, avoiding issues with early versions of the package. ([BZ#2060147](#))
- Previously, the installation program specified network resources to terraform by the relative path. When a network resource was nested in a folder, the terraform provider could not find the resource. With this update, network resources are now specified by ID, and the installation succeeds. ([BZ#2063829](#))
- Previously, vSphere RHCOS images had no **/etc/resolv.conf** file. This caused the default **networkmanager** settings to display an error for **/etc/resolv.conf**. With this update, the **rc-manager=unmanaged** value is set, and the **networkmanager** settings do not direct to **/etc/resolv.conf**. ([BZ#2029438](#))
- Previously, the installation program did not create a cloud provider configuration because it is not required by Amazon Web Services (AWS). This caused the Kubernetes API server to provide an error without a cloud provider configuration. With this update, an empty cloud provider configuration for AWS is created, and the Kubernetes API server can run successfully. ([BZ#1926975](#))

## Kubernetes API server

- Previously, long running requests used for streaming were taken into account for the **KubeAPIErrorBudgetBurn** calculation. Consequently, The alert from **KubeAPIErrorBudgetBurn** would be triggered and cause false positives. This update excludes

long running requests from the **KubeAPIErrorBudgetBurn** calculation. As a result, false positives are reduced on the **KubeAPIErrorBudgetBurn** metric. ([BZ#1982704](#))

## Kubernetes Scheduler

- With OpenShift Container Platform 4.11, the hosted control plane namespace is excluded from eviction when the descheduler is installed on a cluster that has hosted control planes enabled. As a result, pods are no longer evicted from the hosted control plane namespace when the descheduler is installed. ([BZ#2000653](#))
- Previously, resources incorrectly specified the API version in the owner reference of the **kubedescheduler** custom resource (CR). Consequently, the owner reference was invalid, and the affected resources would not be deleted when the **kubedescheduler** CR ran. This update specifies the correct API version in all owner references. As a result, all resources with an owner reference to the **kubedescheduler** CR are deleted after the CR is deleted. ([BZ#1957012](#))

## Machine Config Operator

- Because **keyFile** is not configured as the default plugin for NetworkManager on RHEL nodes, RHEL nodes might not reach the ready state after a reboot. With this fix, **keyFile** is set as the default NetworkManager plugin on all cluster nodes. As a result, nodes correctly reach the **ready** state after a reboot. ([BZ#2060133](#))
- Because vSphere UPI clusters do not set the **PlatformStatus.VSphere** parameter at installation, the parameter was set to **nil**. This caused the MCO logs to be filled with unnecessary and repetitive messages that this parameter cannot have the value **nil**. This fix removes the warning, which was added to resolve a separate issue. As a result, the logs no longer list this message for vSphere UPI installations. ([BZ#2080267](#))
- Previously, when you attempted to create a cluster with both FIPS and **realTimeKernel** the Machine Config Operator (MCO) degraded due to a merge logic issue within the code. With this update, the MCO no longer degrades when creating a cluster with both FIPS and **realTimeKernel**. ([BZ#2096496](#))

## Compliance Operator

- Previously, the Compliance Operator held references to machine configuration data, which significantly increased memory usage. Consequently, the Compliance Operator would fail with **CrashLoopBackoffs** because of out-of-memory exceptions. As a workaround, you should use an updated version of the Compliance Operator, for example, 0.1.53, which includes better handling of large machine configuration data sets in memory. As a result, the Compliance Operator continues to run when dealing with large machine configuration data sets. ([BZ#2094854](#))

## Management Console

- Previously, the web console was not properly authenticating permissions when approving **InstallPlans**. This resulted in a possible unhandled error. With this update, permissions were altered for consistency and error messages are now displayed correctly in the web console. ([BZ#2006067](#))

## Monitoring

- Before this update, dashboards in the OpenShift Container Platform web console that contained queries using a container label for **container\_fs\*** metrics returned no data points because the container labels had been dropped due to high cardinality. This update resolves the issue, and these dashboards now display data as expected. ([BZ#2037513](#))



- Before this update, the **prometheus-operator** component allowed any time value for **ScrapeTimeout** in the config map. If you set **ScrapeTimeout** to a value greater than the **ScrapeInterval** value, Prometheus would stop loading the config map settings and fail to apply all subsequent configuration changes. With this update, if the **ScrapeTimeout** value specified is greater than the **ScrapeInterval** value, the system logs the settings as invalid, but continues loading the other config map settings. ([BZ#2037762](#))
- Before this update, in the **CPU Utilisation** panel on the **Kubernetes / Compute Resources / Cluster** dashboard in the OpenShift Container Platform web console, the formula used to calculate the CPU utilization of a node could incorrectly display invalid negative values. With this update, the formula has been updated, and the **CPU Utilisation** panel now shows correct values. ([BZ#2040635](#))
- Before this update, data from the **prometheus-adapter** component could not be accessed during the automatic update that occurs every 15 days because the update process removed old pods before the new pods became available. With this release, the automatic update process now only removes old pods after the new pods are able to serve requests so that data from the old pods continues to be available during the update process. ([BZ#2048333](#))
- Before this update, the following metrics were incorrectly missing from **kube-state-metrics**: **kube\_pod\_container\_status\_terminated\_reason**, **kube\_pod\_init\_container\_status\_terminated\_reason**, and **kube\_pod\_status\_scheduled\_time**. With this release, **kube-state-metrics** correctly displays these metrics so that they are available. ([BZ#2050120](#))
- Before this update, if invalid write relabel config map settings existed for the **prometheus-operator** component, the configuration would still load all subsequent settings. With this release, the component checks for valid write relabel settings when loading the configuration. If invalid settings exist, an error is logged, and the configuration loading process stops. ([BZ#2051470](#))
- Before this update, the **init-config-reloader** container for the Prometheus pods requested **100m** of CPU and **50Mi** of memory, even though in practice the container needed fewer resources. With this update, the container requests **1m** of CPU and **10Mi** of memory. These settings are consistent with the settings of the **config-reloader** container. ([BZ#2057025](#))
- Before this update, when an administrator enabled user workload monitoring, the **user-workload-monitoring-config** config map was not automatically created. Because non-administrator users with the **user-workload-monitoring-config-edit** role did not have permission to create the config map manually, they required an administrator to create it. With this update, the **user-workload-monitoring-config** config map is now automatically created when an administrator enables user workload monitoring and is available to edit by users with the appropriate role. ([BZ#2065577](#))
- Before this update, after you deleted a deployment, the Cluster Monitoring Operator (CMO) did not wait for the deletion to be completed, which caused reconciliation errors. With this update, the CMO now waits until deployments are deleted before recreating them, which resolves this issue. ([BZ#2069068](#))
- Before this update, for user workload monitoring, if you configured external labels for metrics in Prometheus, the CMO did not correctly propagate these labels to Thanos Ruler. If you queried external metrics for user-defined projects, not provided by the user workload monitoring instance of Prometheus, you would sometimes not see external labels for these metrics even though you had configured Prometheus to add them. With this update, the CMO now properly propagates the external labels that you configured in Prometheus to Thanos Ruler, and you can see the labels when you query external metrics. Therefore, for user-defined projects, if you queried external metrics not provided by the user workload monitoring instance of Prometheus,

you would sometimes not see external labels for these metrics even though you had configured Prometheus to add them. With this update, the CMO now properly propagates the external labels that you configured in Prometheus to Thanos Ruler, and you can see the labels when you query external metrics. ([BZ#2073112](#))

- Before this update, the **tunbr** interface incorrectly triggered the **NodeNetworkInterfaceFlapping** alert. With this update, the **tunbr** interface is now included in the list of interfaces that the alert ignores and no longer causes the alert to trigger incorrectly. ([BZ#2090838](#))
- Previously, the Prometheus Operator allowed invalid re-label configurations. With this update, the Prometheus Operator validates re-labeled configurations. ([BZ#2051407](#))

## Networking

- Previously, when using the bond CNI plugin for an additional network attachment, it was not compatible with Multus. When the bond CNI plugin was used in conjunction with the Whereabouts IPAM plugin for a network attachment definition, assigned IP addresses were incorrectly reconciled. Now a network attachment definition that uses the bond CNI plugin works correctly with the Whereabouts IPAM plugin for IP address assignment. ([BZ#2082360](#))
- Previously, when using the OVN-Kubernetes cluster network provider with multiple default gateways, the wrong gateway was selected, causing the OVN-Kubernetes pods to stop unexpectedly. Now the correct default gateway is selected such that these pods no longer fails. ([BZ#2040933](#))
- For clusters using the OVN-Kubernetes cluster network provider, previously if the NetworkManager service restarted on a node, that node lost network connectivity. Now network connectivity survives a restart of the NetworkManager service. ([BZ#2048352](#))
- Previously a **goroutine** handling cache updates could stall writing to an unbuffered channel while holding a **mutex**. With this update, these race conditions were solved. ( [BZ#2052398](#))
- Previously, for ovn-kubernetes, setting up **br-ex** on boot with a bond or team interface caused a mismatch on media access control (MAC) addresses between the **br-ex** and the bond interface. Consequently, on bare metal or some virtual platforms all of the traffic was dropped due to network interface controller (NIC) driver dropping traffic due to an unexpected **br-ex** MAC address. With this update, **br-ex** and the bond interface use the same MAC address resulting in no dropped traffic. ([BZ#2103080](#))
- Previously, users with **cluster-reader** role could not read custom resources from kubernetes-nmstate, such as **NodeNetworkConfigurationPolicy**. With this update, users with **cluster-reader** role can read kubernetes-nmstate custom resources. ( [BZ#2022745](#))
- Previously, contrack entries for **LoadBalancer** IPs were not removed when the service endpoints were removed causing connections to fail. With this update, contrack entries do not cause connections to fail. ([BZ#2061002](#))
- Previously, a missing **jq package** caused the scale up of a cluster with RHEL nodes to fail on node deployment. With this update, `jq package` is installed on deployment and the scale up of a cluster with RHEL nodes succeeds. ([BZ#2052393](#))
- Previously, OVN-Kubernetes spent excessive time when a service configuration change was made. This caused a noticeable latency in service configuration changes. With this update, OVN-Kubernetes is optimized to reduce the latency of service configuration changes. ([BZ#2070674](#))



- Previously, the IP reconciliation CronJob for Whereabouts IPAM CNI would fail due to API connectivity issues, which caused CronJob to intermittently fail. With this update, CronJob launched on Whereabouts IPAM CNI use the api-internal server address and an extended api timeout to prevent these connectivity issues. ([BZ#2048575](#))
- With this update, OpenShift Container Platform clusters with Kubernetes-NMstate installed now include in the `must-gathers` Kubernetes-NMstate resources. This improves issue handling by including resources from Kubernetes-NMstate in **must-gathers**. ([BZ#2062355](#))
- There is currently a known issue with host routes being ignored when load-balancer services are configured with cluster traffic policy. Consequently, egress traffic of load balancer services is steered to the default gateway and not towards the best matching route that is present on the host routing table. As a workaround, set the load balancer type Services to **Local** traffic policy. ([BZ#2060159](#))
- Previously, **PodDisruptionBudget** specifications did not fit for single-node OpenShift clusters, which limited upgrade capabilities due to not all pods being able to be evicted. With this update, **PodDisruptionBudget** specification is adjusted based on cluster topology making Kubernetes-NMState operator capable of upgrading on single-node OpenShift clusters. ([BZ#2075491](#))
- Previously, when setting up **br-ex** bridge on boot, the DHCP client id and IPv6 address generation mode configuration did not work properly causing an unexpected IP address on **br-ex**. With this update, DHCP client id and IPv6 generation mode configuration are now properly set on **br-ex**. ([BZ#2058030](#))
- Previously, **egress-router-cni** pods lacked some cluster internal routes due to the reliance on the **gateway** field of the CNI definition to delete the default route and inject its own. With this update, **egress-router-cni** pods inject the correct routing information so that pods reach external and internal cluster destinations. ([BZ#2075475](#))
- Previously, a pod disruption budget was used to create an OVN raft quorum on single-node OpenShift. This raised an unhelpful **PodDisruptionBudgetAtLimit** alert on single-node OpenShift clusters. With this update, the **PodDisruptionBudgetAtLimit** alert is no longer raised on these clusters. ([BZ#2037721](#))
- Previously, a race condition in the **NetworkManager** restart sometimes prevented DHCP resolution to successfully complete setting up the **br-ex** bridge on node boot when using OVN-Kubernetes. With this update, **NetworkManager** is no longer restarted when setting up **br-ex**, eliminating the race condition. ([BZ#2055433](#))
- Previously, the **PtpConfigSlave** source custom resource (CR) was set to the unsupported network transport UDPv4, causing errors on Distributed Unit (DU) nodes. This fix updates the **PtpConfigSlave** source CR to use network transport L2 rather than UDPv4. As a result, errors are no longer found on the DU nodes. ([BZ#2060492](#))
- Previously, all OpenShift Container Platform policies logging configurations were updated when a network policy was updated. This caused a noticeable latency on some concurrent or later network policies. With this update, all network policies are no longer updated when adding a new policy, eliminating latency. ([BZ#2089392](#))

## Networking performance improvements

- Previously, a **systemd** service set a default Recieve Packet Steering (RPS) mask according to the reserved CPUs list in the performance profile for all network devices visible from udev excluding virtual devices. A **crio** hook script set the RPS mask of all network devices visible from **/sys/devices** for guaranteed pods. This resulted in multiple impacts to network performance. In this update, the **systemd** service only sets the default RPS mask for virtual interfaces under

**/sys/devices/virtual**. The **crio** hook script now also excludes physical devices. This configuration mitigates issues such as overloading of processes, lengthy polling intervals, and spikes in latency. ([BZ#2081852](#))

## Node

- Previously, the pod manager handled the registration and de-registration of pod secrets and config maps. Because of this, pod secrets would sometimes fail to be mounted within a pod. With this fix, the pod ID is included in the key that the kubelet uses to manage registered pods. As a result, secrets are properly mounted as expected. ([BZ#1999325](#))
- Because of a memory leak in the garbage collection process, pods might not be able to start on a node due to lack of memory. With this fix memory is no longer leaking in the garbage collection process and nodes should start as expected. ([BZ#2065749](#))
- Because of a change in the upstream Kubernetes, the kubelet was not running readiness probes on terminated pods. As a result, a load balancer or controller could react more slowly to a terminating pod, which might have resulted in errors. With this fix, readiness probes are again being performed on pod termination. ([BZ#2089933](#))
- Because of a bug, the kubelet could incorrectly reject pods that have **OutOfCpu** errors, if the pods were rapidly scheduled after other pods were reported as complete in the API. With this fix, the kubelet now waits to report the phase of a pod as terminal in the API until all running containers have stopped and no new containers have been started. Short-lived pods may take slightly longer, approximately 1s, to report either success or failure after this change. ([BZ#2022507](#))
- Because recent versions of the **prometheus-adapter** are sending additional pod metrics, the Vertical Pod Autoscaler (VPA) recommender is producing a large number of unnecessary and repetitive messages. With this fix, the VPA recognizes and ignores the extra metrics. As a result, these messages are no longer being produced. ([BZ#2102947](#))

## OpenShift CLI (oc)

- Previously, **oc** catalog mirroring failed if an older, deprecated image version was used as the source. Now, the image manifest version is detected automatically and mirroring works successfully. ([BZ#2049133](#))
- Previously, it was hard to understand from the logs when fallback inspect occurred. The logs are now improved to make this more explicit. As a result, the output of **must-gather run** is much more clear. ([BZ#2035717](#))
- Previously, if you ran **must-gather** with invalid arguments, it did not consistently report the error and instead might attempt to collect data even when that is not possible. Now if **must-gather** is called with invalid options, it provides useful error output. ([BZ#1999891](#))
- Previously, if the **oc adm catalog mirror** command resulted in errors, it still continued and returned a **0** exit code. A **--continue-on-error** flag is now available that allows users to determine whether the command should continue if there are errors, or exit with a non-zero exit code. ([BZ#2088483](#))
- With this update, a **--subresource** flag was added to the **oc adm policy who-can** command to check who can perform a specified action on a subresource. ([BZ#1905850](#))
- Previously, users were unable to use tab completion for the **oc project** command. Now, hitting tab after **oc project** properly lists projects. ([BZ#2080416](#))

- Previously, startup probes were not removed from debug pods, which could cause issues with the debug pods if the startup probe failed. The **--keep-startup** flag has been added, which is **false** by default, meaning that startup probes are removed by default from debug pods. ([BZ#2056122](#))
- Previously, there was no timeout specified after invoking **oc debug node**, so users were never logged out of the cluster. A **TMOUT** environment variable has been added, so that after the specified time of inactivity, the session is automatically terminated. ([BZ#2043314](#))
- With this update, **oc login** now shows the URL for the web console even when users are logged out. ([BZ#1957668](#))
- Previously, the **oc rsync** command displayed a wrong error output when the container was not found. With this release, the **oc rsync** command displays the correct error message when the specific container is not running. ([BZ#2057633](#))
- Previously, large images could not be pruned if they were new to the cluster. This caused recent images to be omitted when filtering the over sized images. With this release, you can now prune images that exceed the given size. ([BZ#2083999](#))
- Previously, there was a typo in the **gather** script. As a result, insights data was not collected properly. With this release, the typo is corrected and Insights data is now properly collected through the **must-gather**. ([BZ#2106543](#))
- Previously, you could not apply the **EgressNetworkPolicy** resource type in your cluster through the **oc** CLI. With this release, you can now create, update, and delete an **EgressNetworkPolicy** resource. ([BZ#2071614](#))

### Kubernetes Controller Manager

- Previously, the beta feature for tracking jobs using pod finalizers was enabled by default. In some cases, pods were not always removed because the finalizers on them were not removed. With this update, the feature gate **JobTrackingWithFinalizers** is disabled by default. As a result, no pods should be left behind during removal. ([BZ#2075621](#))
- Previously, a **PodDisruptionBudgetAtLimit** alert would occur whenever the CR replica count was zero. With this update, the alert will no longer fire if there is no application to disrupt or when the replica count is zero. ([BZ#2053622](#))

### Operator Lifecycle Manager (OLM)

- Before this update, invalid subscription labels were created when a resource name exceeded 63 characters. Truncating labels that exceed the 63-character limit resolves the issue, and the subscription resource no longer rejects the Kubernetes API. ([BZ#2016425](#))
- Before this update, catalog source pods for the Marketplace Operator prevented nodes from draining. As a result, the Cluster Autoscaler could not scale down effectively. With this update, adding the **cluster-autoscaler.kubernetes.io/safe-to-evict** annotation to the catalog source pods fixes the issue, and the Cluster Autoscaler can scale down effectively. ([BZ#2053343](#))
- Before this update, the **collect-profiles** job could take a long time to complete in certain circumstances, such as when a pod could not be scheduled. As a result, if enough jobs were scheduled but unable to run, the number of scheduled jobs exceeded pod quota limits. With this update, only one **collect-profiles** pod exists at a time, and the **collect-profiles** job does not exceed pod quota limits. ([BZ#2055861](#))
- Before this update, the package server was not aware of pod topology when defining its leader

election duration, renewal deadline, and retry periods. As a result, the package server strained topologies with limited resources, such as single-node environments. This update introduces a **leaderElection** package that sets reasonable lease duration, renewal deadlines, and retry periods. This fix reduces strain on clusters with limited resources. ([BZ#2048563](#))

- Previously, there was a bad catalog source in the **openshift-marketplace** namespace. Because of this, all subscriptions were blocked. With this update, if there is a bad catalog source in the **openshift-marketplace** namespace, users can subscribe to an operator from a quality catalog source of their own namespace with the original annotation. As a result, if there is a bad catalog source in the local namespace, the user cannot subscribe to any operator in the namespace. ([BZ#2076323](#))
- Previously, info-level logs were generated during **operator-marketplace** project polling, which caused log spam. This update uses the command line flag to reduce the log line to the debug level, and adds more control of the log levels for the user. As a result, this reduces log spam. ([BZ#2057558](#))
- Previously, each component managed by the Cluster Version Operator (CVO) consisted of YAML files defined in the **/manifest** directory in the root of a repository for a project. When removing a YAML file from the **/manifest** directory, you needed to add the **release.openshift.io/delete: "true"** annotation, otherwise the CVO would not delete the resources from the cluster. This update reintroduces any resources that were removed from the **/manifest** directory and adds the **release.openshift.io/delete: "true"** annotation so that the CVO cleans up the resources. As a result, resources that are no longer required for the OLM component are removed from the cluster. ([BZ#1975543](#))
- Previously, the **CheckRegistryServer** function used by gRPC catalog sources did not confirm the existence of the service account associated with the catalog source. This caused the existence of an unhealthy catalog source with no service account. With this update, the gRPC **CheckRegistryServer** function checks if the service account exists and recreates the service if it is not found. As a result, the OLM recreates service accounts owned by gRPC catalog sources if they do not exist. ([BZ#2074612](#))
- Previously, in an error message that occurred when users ran **opm index prune** against a file-based catalog image, imprecise language made it unclear that this command does not support that catalog format. This update clarifies the error message so users understand that the command **opm index prune** only supports SQLite-based images. ([BZ#2039135](#))
- Previously, there was a broken thread safety around the Operator API. Consequently, Operator resources were not properly deleted. With this update, Operator resources are correctly deleted. ([BZ#2015023](#))
- Previously, pod failures were artificially extending the validity period of certificates causing them to incorrectly rotate. With this update, the certificate validity period is correctly determined and the certificates are correctly rotated. ([BZ#2020484](#))
- In OpenShift Container Platform 4.11 the default cluster-wide pod security admission policy is set to **baseline** for all namespaces and the default warning level is set to **restricted**. Before this update, Operator Lifecycle Manager displayed pod security admission warnings in the **operator-marketplace** namespace. With this fix, reducing the warning level to **baseline** resolves the issue. ([BZ#2088541](#))

## Operator SDK

- Before this update, the Operator SDK used upstream images rather than downstream supported images to scaffold Hybrid Helm-based Operators. With this update, the Operator SDK uses supported downstream images to scaffold Hybrid Helm-based Operators.

([BZ#2039135](#))

- With OpenShift Container Platform 4.11, Operator SDK allows for **arm64** Operator images to be built. As a result, Operator SDK now supports building Operator images that target **arm64**. ([BZ#2035899](#))
- Previously, {product-tile} running Hybrid Helm Operators scaffolded with Operator SDK would use upstream images rather than supported downstream images. With this update, scaffolding a Hybrid Helm Operator uses downstream images. ([BZ#2066615](#))

### OpenShift API server

- Because multiple Authentication Operator controllers were synchronizing at the same time, the Authentication Operator was taking too long to react to changes to its configuration. This feature adds jitter to the regular synchronization periods so that the Authentication Operator controllers do not compete for resources. As a result, it now takes less time for the Authentication Operator to react to configuration changes. ([BZ#1958198](#))
- With OpenShift Container Platform 4.11, authentication attempts from external identity providers are now logged to the audit logs. As a result, you can view successful, failed, and errored login attempts from external identity providers in the audit logs. ([BZ#2086465](#))

### Red Hat Enterprise Linux CoreOS (RHCOS)

- Before this update, if a machine was booted through PXE and the **BOOTIF** argument was on the kernel command line, the machine would boot with DHCP enabled on only a single interface. With this update, the machine boots with DHCP enabled on all interfaces even if the **BOOTIF** argument is provided. ([BZ#2032717](#))
- Previously, nodes that were provisioned from VMware OVA images did not delete the Ignition config after initial provisioning. Consequently, this created security issues when secrets are stored within the Ignition config. With this update, the Ignition config is now deleted from the VMware hypervisor after initial provisioning on new nodes and when upgrading from a previous OpenShift Container Platform release on existing nodes. ([BZ#2082274](#))
- Previously, any arguments provided to the **toolbox** command were ignored when the command was first invoked. This fix updates the toolbox script to initiate the **podman container create** command followed by the **podman start** and **podman exec** commands. It also modifies the script to handle multiple arguments and whitespaces as an array. As a result, the arguments passed to the **toolbox** command are executed every time as expected. ([BZ#2039589](#))

### Performance Addon Operator

- Previously, the CNF cyclicttest runner should have provided the **--mainaffinity** argument, which told the binary which thread it should run on, however the cyclicttest runner was missing the **--mainaffinity** argument. This update adds the **--mainaffinity** argument to the cyclicttest runner, so that it is properly passed to the **cyclitest** command. ([BZ#2051540](#))
- Previously, the **oslat** container specs lacked the **cpu-quota.crio.io: "disable"** annotation, which resulted in high latency. Consequently, the **cpu-quay.crio.io: "disable"** annotation was missing from the pod definition during creation. With this update, the **cpu-quota.crio.io: "disable"** annotation is appended during pod creation and, as a result, appears in the **oslat** pod's **specification** field. ([BZ#2061676](#))

### Routing

- Previously, the Ingress Operator did not validate whether a Kubernetes service object in the OpenShift Ingress namespace was created or owned by the Ingress Controller it was trying to

reconcile with. Therefore, the Ingress Operator would modify or remove Kubernetes services that had the same name and namespace, regardless of ownership, causing unexpected behavior. With this update, the Ingress Operator can now check the ownership of existing Kubernetes services before attempting to modify or remove services. If ownership does not match, the Ingress Operator shows an error and does not take any action. As a result, the Ingress Operator cannot modify or delete a custom Kubernetes service with the same name as the OpenShift Ingress namespace that it wants to modify or remove. ([BZ#2054200](#))

- Previously, OpenShift Container Platform 4.8 added an API for customizing platform routes. This API includes status and spec fields in the cluster ingress configuration for reporting the current host names of customizable routes and the user's desired host names for these routes, respectively. The API also defined constraints for these values. These constraints were restrictive and excluded some valid potential host names. Consequently, the restrictive validation for the API prevented users from specifying custom host names that should have been permitted and prevented users from being able to install clusters with domains that should have been permitted. With this update, the constraints on host names were relaxed to allow all host names that are valid for routes and OpenShift Container Platform allows users to use cluster domains with TLDs that contain decimal digits. ([BZ#2039256](#))
- Previously, the Ingress Operator did not check whether Ingress Controllers configured with cluster **spec.domain** parameter matched the **spec.baseDomain** parameter. This caused the Operator to create DNS records and set **DNSManaged** conditions to **false**. With this fix, the Ingress Operator now checks whether the **spec.domain** parameter matches with the cluster **spec.baseDomain**. As a result, for custom Ingress Controllers, the Ingress Operator does not create DNS records and sets **DNSManaged** conditions to false. ([BZ#2041616](#))
- Previously, in OpenShift Container Platform 4.10, the HAProxy must-gather function could take up to an hour to run. This can happen when routers in the terminating state delay the **oc cp** command. The delay lasts until the pod is terminated. With the new release, a 10 minute limit on the **oc op** command prevents longer delays. ([BZ#2104701](#))
- Previously, the Ingress Operator did not clear the route status when Ingress Controllers were deleted, showing that the route was still in the operator after its deletion. This fix clears the route status when an Ingress Controller is deleted, resulting in the route being cleared in the operator after its deletion. ([BZ#1944851](#))
- Previously, the output for the **oc explain router.status.ingress.conditions** command explain route status showed **Currently only Ready** rather than **Admitted** due to incorrect wording in the Application Programming Interface (API). This fix corrects the wording in the API. As a result, the command output is correct. ([BZ#2041133](#))
- Previously, the Ingress Operator detected that the user modified an annotation that the Operator manages on the **LoadBalancer-type** service. Consequently, the Operator set the Ingress Cluster Operator's **Upgradeable** status condition to **False** to block upgrades, and the Ingress Operator erroneously set the **Upgradeable** status condition to **False**, blocking upgrades, if the service had no annotations. Now, the logic that checks the service's annotations correctly handles empty annotations, and the Ingress Operator no longer erroneously blocks upgrades. ([BZ#2097555](#))
- Previously, the Ingress Operator removed a finalizer that the Operator added to **LoadBalancer-type** services from previous versions of OpenShift Container Platform. With this update, the Ingress Operator no longer removes finalizers. ([BZ#2069457](#))
- The Ingress Operator performs health checks against the ingress canary route. Before this update, **keepalive** daemons are enabled on the connection, which caused the Ingress Operator not to close the TCP Connection to the load balancer (LB) after the health check was completed. A new connection was created for the next health check instead of using the existing



connection. Consequently, the connection built on the load balancer and creating too many connections on the load balancer. With this update, **keepalive** daemons are disabled when connecting to the canary route, and a new connection is made and closed each time the canary probe is run. ([BZ#2037447](#))

- Previously, the Ingress Controller did not set the **allowPrivilegeEscalation** value to false in the router deployment, which caused the router pods to be selected into the incorrect Security Context Constraint (SCC) and created conflicts with custom SCCs. This fix sets the **allowPrivilegeEscalation** value to **true**, ensuring that router pods are selected into the correct SCC and avoid conflicts with custom SCCs. ([BZ#2007246](#))
- Previously, when the canary route was not admitted to an Ingress Controller, the Ingress Operator status condition did not show as **degraded**. Consequently, the canary route could have shown as **valid** when its status condition should show as **not admitted**. With this update, the Ingress Operator status more accurately reflects the status of the canary controller. ([BZ#2021446](#))
- Previously, the **openshift-router** process briefly ignored the **SIGTERM** shutdown signal. This caused containers to ignore a Kubernetes shutdown request, resulting in one hour shutdown times. With this update, the router now responds to **SIGTERM** signals. ([BZ#2076297](#))
- Previously, when an Ingress Controller for an admitted route was deleted or a sharding configuration was added, a false status of **admitted** was given. With this update, the Ingress Controller clears the status of an **unadmitted** route, avoiding the false status scenario. ([BZ#1944851](#))
- Previously, OpenShift Container Platform clusters installed using version 4.7 or earlier maintained a **service.beta.kubernetes.io/aws-load-balancer-internal** annotation value of **0.0.0.0/0**. Clusters that were installed using 4.8 or later have the annotation value **true**. AWS cloud-provider implementations that check for the annotation value **true** would return the wrong result if the value was **0.0.0.0/0**. As a result, cluster upgrades to 4.10 would not complete. With this update, the annotation value is normalized to **true** so that the cluster upgrades can complete. ([BZ#2055470](#))

## Scalability and performance

- Before this update, SRO installed Node Feature Discovery (NFD) by default, regardless of whether NFD had already been installed. If NFD had been installed, this would cause the SRO deployment to fail. SRO no longer deploys NFD by default.

## Storage

- Previously, the Alibaba Container Storage Interface (CSI) driver that shipped with OpenShift Container Platform returned errors when a user created a persistent volume claim (PVC) smaller than 20 GiB. This was due to Alibaba Cloud only supporting volumes larger than 20 GiB. With this update, Alibaba CSI driver automatically increases all volume sizes to at least 20 GiB and smaller PVCs are now dynamically provisioned. This can result in increased costs. Administrators can use quota on PVC count for each namespace in restricted environments to limit costs. ([BZ#2057495](#))
- In earlier versions, the Local Storage Operator (LSO) added an owner reference to the persistent volumes (PVs) it created, such that deletion of the node would also issue a delete request for the PV. This could result in the PV remaining in the **Terminating** state while still attached to a pod. LSO no longer creates that OwnerReference, which means cluster administrators must delete any unused PVs after a node is removed from the cluster. ([BZ#2061447](#)) For more information, see [Persistent storage using local volumes](#).

- This fix ensures that read-only-many volumes are provisioned appropriately by GCP CSI Driver as read-only. ([BZ#1968253](#))
- OpenShift Container Platform allows installation of vSphere CSI drivers shipped by either the upstream community or VMware. Although Red Hat does not support this driver, cluster administrators can still install and use it because it has more features than the vSphere CSI driver shipped by Red Hat. OpenShift Container Platform can be upgraded to 4.11 with the upstream and VMware vSphere CSI driver, however, you will be warned about the presence of a third party CSI driver. For more information, see ([BZ#2089419](#)) and ([BZ#2052071](#)).
- With this update, the default credentials request for Amazon Web Services (AWS) is modified to allow mounting of encrypted volumes using customer managed keys from Key Management Service (KMS). Administrators who created credentials requests in manual mode with the Cloud Credential Operator (CCO) will need to apply those changes manually if they intend to mount encrypted volumes using customer managed keys on AWS. Other administrators should not be impacted by this change. ([BZ#2049872](#))
- Previously, after deleting an OpenShift Container Platform cluster deployed on IBM Cloud, the back-end storage volumes were not deleted. This prevented the cluster resources from being completely removed. This fix adds support in the installation program and the Container Storage Interface (CSI) driver, resulting in back-end volume deletion after a cluster is removed. ([BZ#2047732](#))

### Web console (Developer perspective)

- Before this update, the **Git Import** form displayed an error message when an invalid devfile repository (older than devfile v2.2) is entered. With this update, the error message states that devfiles older than v2.2 are not supported. ([BZ#2046435](#))
- Before this update, if the **ConsoleLink CR** (openshift-blog) was not available in the cluster, the blog link was undefined. Clicking on the blog link did not redirect to the OpenShift blog. With this update, a fall back link to <https://developers.redhat.com/products/openshift/whats-new> is added even if the **ConsoleLink CR** (openshift-blog) is not present in the cluster. ([BZ#2050637](#))
- Before this update, the API version for the kafka CR was updated. This version did not support the old version, so an empty **Bootstrap server** was displayed on **Create Event Source - KafkaSource** even if it was created. With this update, the updated APIs for the Kafka CR support the old versions and renders the **Bootstrap server** list in **Create Event Source - KafkaSource** form. ([BZ#2058623](#))
- Before this update, when you used the **Import from Git** form to import a private Git repository, the correct import type and a builder image were not identified because the secret to fetch the private repository details was not decoded. With this update, the **Import from Git** form decodes the secret to fetch the private repository details. ([BZ#2053501](#))
- Before this update, from the developer perspective, the **Observe** dashboard opened for the most recently viewed workload rather than the one you selected in the **Topology** view. This issue happens because the session prefers the redux store instead of the query parameters in the URL. With this update, the **Observe** dashboard renders components based on the query parameters in the URL. ([BZ#2052953](#))
- Before this update, the **Pipeline** used to start with the hardcoded value **gp2** as the default storage class even if it did not exist on the cluster. With this update, you can use the default specified storage class name instead of a hardcoded value. ([BZ#2084635](#))
- Before this update, while running high-volume pipeline logs, the auto-scroll functionality does



not work and logs display older messages. Running high-volume pipeline logs generates a large number of calls to the **scrollIntoView** method. With this update, high-volume pipeline logs do not generate any calls to the **scrollIntoView** method and gives a smooth auto-scroll functionality. ([BZ#2014161](#))

- Before this update, when creating a **RoleBinding** using the **Create RoleBinding** form, the subject name was mandatory. A missing subject name fails to load the **Project Access** tab. With this update, the **RoleBinding** without the **Subject Name** property is not listed in the **Project Access** tab. ([BZ#2051558](#))
- Before this update, the sink and trigger for event sources showed all the resources, even though those are standalone or part of backing the **k-native service, Broker, or KameletBinding**. The addressed resources used to show in the sink dropdown list. With this update, a filter has been added to show only standalone resources as sink. ([BZ#2054285](#))
- Before this update, the empty tabs in the sidebar of the topology view were not filtered out before rendering. This displayed invalid tabs for **Workloads** in the topology view. With this update, the empty tabs are filtered properly before rendering. ([BZ#2049483](#))
- Before this update, when a pipeline is started using the **Start Last Run** button, the **started-by** annotation of the created **PipelineRun** was not updated to the correct username so the triggered by section did not show the correct username. With this update, the **started-by** annotation value is updated to the correct username and the triggered by section shows the username of the correct user that started the pipeline. ([BZ#2046618](#))
- Before this update, the **ProjectHelmChartRepository** CR does not show up in the cluster. Consequently, the API schema for this CR has not been initialized in the cluster yet. With this update the **ProjectHelmChartRepository** shows up in the cluster. ([BZ#2054197](#))
- Before this update, when you navigate using a keyboard in the topology, the selected items were not highlighted. With this update, navigation using a keyboard highlights and updates styles to the selected items. ([BZ#2039277](#))
- Before this update, the layout of the web terminal opened outside of the default view and could not be resized. With this update, the web terminal opens inside the default view and resizes properly. ([BZ#2022253](#))
- Before this update, some of the sidebar items did not contain namespace context. Consequently, when links were opened from another browser, or opening links from a different active namespace, the web console does not switch to the correct namespace. With this update, the correct namespace is selected when opening the URL. ([BZ#2039647](#))
- Previously, when using the console to instantiate a template, its parameters were stored as a secret resource. When the template was removed, the secret remained. This consequently created an unnecessary build up of secrets in a cluster. With this update, an ownership reference is added to the secret that maps to the template instance. Now, when the template instance is removed, the secret is also removed. ([BZ#2015042](#))
- With this update, the **jsonData** property has been deprecated and replaced with **data** in the **ping** source. ([BZ#2084438](#))
- Previously, the topology view in the OpenShift Container Platform web console would fail or lag for clusters with more than 100 nodes. With this update, the topology view shows a **LimitExceeded** state for clusters with more than 100 nodes. An option is provided to view the resources by using the **Search** page instead. Alternatively, you can click **Show topology anyway** to continue to load the topology view. ([BZ#2060329](#))

- Previously, if a service exposed multiple service ports and the route target port was **8080**, attempts to change the target port would result in another service port being updated instead of the port **8080** service port. With this update, the service port corresponding to the active target port is replaced when a new target port is set. ([BZ#2077943](#))
- Previously, the **git** detection used to manage instance APIs to get repository information did not work for repositories from self hosted GitHub and Bitbucket. With this update, detection for self hosted GitHub and Bitbucket instance repositories works. ([BZ#2038244](#))
- Previously, **apiVersion** was not passed in the correct format to the **Resource** drop-down menu in the **EventSource** creation form. This caused **InContext** to not be selected under **EventSource** creation, which excluded it from the **Resource** drop-down menu. With this update, the **Resource** drop-down menu includes **Resource** from **InContext**. ([BZ#2070020](#))
- Previously, the **Pipeline metrics** page displayed all API calls for the metrics query and failed with a **404** error. With this update, the **prometheus-tenancy** API is used to get the metrics data for the pipeline. Now, the pipeline metrics page displays all data and graphs to the non-admin user with at least view access to the namespace. ([BZ#2041769](#))
- Previously, you could access a quick search and add modal with the Ctrl+space keyboard shortcut, but you could not close them by using the same keyboard shortcut. With this update, you can close the quick search and add modal using the Ctrl+space keyboard shortcut. ([BZ#2093586](#))
- Previously, the resources created for user settings were not removed if the user was deleted. Consequently, the created resources in the **openshift-console-user-settings** namespace was never removed. With this update, an **ownerReference** is added to your metadata at creation. This allows automatic removal of the resources when the user no longer exists. ([BZ#2019564](#))

## 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 1.2. Technology Preview tracker

Feature	OCP 4.9	OCP 4.10	OCP 4.11
PTP single NIC hardware configured as boundary clock	-	TP	GA
PTP dual NIC hardware configured as boundary clock	-	-	TP

Feature	OCP 4.9	OCP 4.10	OCP 4.11
PTP events with ordinary clock	TP	GA	GA
PTP events with boundary clock	-	TP	GA
Shared Resources CSI Driver and Build CSI Volumes in OpenShift Builds	-	TP	TP
Service Binding	TP	GA	GA
CSI volume expansion	TP	TP	GA
CSI AliCloud Disk Driver Operator	-	GA	GA
CSI Azure Disk Driver Operator	TP	GA	GA
CSI Azure File Driver Operator	-	TP	GA
CSI Azure Stack Hub Driver Operator	GA	GA	GA
CSI GCP PD Driver Operator	GA	GA	GA
CSI IBM VPC Block Driver Operator	-	GA	GA
CSI AWS EFS Driver Operator	TP	GA	GA
CSI vSphere Driver Operator	TP	GA	GA
CSI automatic migration (AWS EBS, Azure file, GCP disk, VMware vSphere)	TP	TP	TP
CSI automatic migration (Azure Disk, OpenStack Cinder)	TP	TP	GA
CSI inline ephemeral volumes	TP	TP	TP
CSI generic ephemeral volumes	-	-	GA
Shared Resource CSI Driver	-	TP	TP
Automatic device discovery and provisioning with Local Storage Operator	TP	TP	TP
OpenShift Pipelines	GA	GA	GA
OpenShift GitOps	GA	GA	GA

Feature	OCP 4.9	OCP 4.10	OCP 4.11
OpenShift sandboxed containers	TP	GA	GA
Adding kernel modules to nodes with kvc	TP	TP	TP
Non-preempting priority classes	TP	TP	GA
Kubernetes NMState Operator	TP	GA	GA
Assisted Installer	TP	GA	GA
<b>kdump</b> on <b>x86_64</b> architecture	TP	TP	GA
<b>kdump</b> on <b>arm64</b> architecture	-	-	TP
<b>kdump</b> on <b>s390x</b> architecture	TP	TP	TP
<b>kdump</b> on <b>ppc64le</b> architecture	TP	TP	TP
OpenShift on ARM platforms	-	GA	GA
Serverless functions	TP	TP	TP
Memory Manager	GA	GA	GA
Cloud controller manager for Alibaba Cloud	-	TP	TP
Cloud controller manager for Amazon Web Services	TP	TP	TP
Cloud controller manager for Google Cloud Platform	-	TP	TP
Cloud controller manager for IBM Cloud	-	GA	GA
Cloud controller manager for Microsoft Azure	TP	TP	TP
Cloud controller manager for Red Hat OpenStack Platform (RHOSP)	TP	TP	TP
Cloud controller manager for VMware vSphere	-	TP	TP
Driver Toolkit	TP	TP	TP
Special Resource Operator (SRO)	TP	TP	TP
Simple Content Access	TP	GA	GA

Feature	OCP 4.9	OCP 4.10	OCP 4.11
Node Health Check Operator	TP	TP	GA
Pod-level bonding for secondary networks	-	GA	<a href="#">GA</a>
IPv6 dual stack	GA	GA	GA
Selectable Cluster Inventory	-	TP	TP
Heterogeneous clusters	-	-	TP
Hyperthreading-aware CPU manager policy	-	TP	
Heterogeneous Clusters	-	-	TP
Dynamic Plugins	-	TP	TP
Hybrid Helm Operator	-	TP	TP
Alert routing for user-defined projects monitoring	-	TP	GA
Disconnected mirroring with the oc-mirror CLI plugin	-	TP	GA
Mount shared entitlements in BuildConfigs in RHEL	-	TP	TP
Mount shared secrets in RHEL	-	GA	GA
Support for RHOSP DCN	-	TP	TP
Support for external cloud providers for clusters on RHOSP	-	TP	TP
OVS hardware offloading for clusters on RHOSP	-	TP	GA
External DNS Operator	-	GA	GA
Web Terminal Operator	TP	GA	GA
Alerting rules based on platform monitoring metrics	-	-	TP
AWS Load Balancer Operator	-	-	TP
Node Observability Operator	-	-	TP
Java-based Operator	-	-	TP
Hosted control planes for OpenShift Container Platform	-	-	TP

Feature	OCP 4.9	OCP 4.10	OCP 4.11
Managing machines with the Cluster API	-	-	TP
Topology Aware Lifecycle Manager	-	TP	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 4.11, you can either revoke or continue to allow unauthenticated access. Unless there is a specific need for unauthenticated access, you should revoke 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](#) )
- In the monitoring stack, if you have enabled and deployed a dedicated Alertmanager instance for user-defined alerts, you cannot silence alerts in the **Developer** perspective in the OpenShift Container Platform web console. This issue has been fixed in 4.11.8. ([BZ#2100860](#))
- On a newly installed OpenShift Container Platform 4.11 cluster, platform monitoring alerts do not have the **openshift\_io\_alert\_source="platform"** label. This issue does not affect clusters upgraded from a previous minor version. There is currently no workaround exists for this issue. ([BZ#2103127](#))
- In the Red Hat OpenStack Platform (RHOSP), a potential issue can affect Kuryr when the ports pools are populated with a variety of bulk port creation requests made at the same time. During these bulk requests, if the IP allocation for one of the IP addresses fails, the Neutron retries the operation for all ports. This issue was resolved in a previous errata; however, you can set a small value for the ports pools batch to avoid large bulk port creation requests. ([BZ#2024690](#))
- The oc-mirror CLI plugin cannot mirror OpenShift Container Platform catalogs earlier than version 4.9. ([BZ#2097210](#))
- The oc-mirror CLI plugin can fail to upload an image set to the target mirror registry if the **archiveSize** value in the image set configuration is smaller than the container image size, causing the catalog directory to span multiple archives. ([BZ#2106461](#))
- In OpenShift Container Platform 4.11, the MetalLB Operator scope has changed from namespace to cluster and this results in upgrading from a prior version failing. As a workaround, remove the previous version of the MetalLB Operator. Do not delete the namespace or the instance of the MetalLB custom resource, and deploy the new Operator version. This keeps MetalLB running and configured.

For more information, see [Upgrading the MetalLB Operator](#). ([BZ#2100180](#))

- Deleting the bidirectional forwarding detection (BFD) profile and removing the **bfdProfile** added to the border gateway protocol (BGP) peer resource does not disable the BFD. Instead, the BGP peer starts using the default BFD profile. To disable BFD from a BGP peer resource, delete the BGP peer configuration and recreate it without a BFD profile. ([BZ#2050824](#))
- The OpenShift CLI (**oc**) for OpenShift Container Platform 4.11 does not work properly on macOS due to a change in error handling of untrusted certificates in Go 1.18 libraries. Due to this change, **oc login** and other **oc** commands can fail with a **certificate is not trusted** error without proceeding further when running on macOS. Until the error handling is properly fixed in Go 1.18 (tracked by [Go issue #52010](#)), the workaround is to use the OpenShift Container Platform 4.10 **oc** CLI instead. Note that when using the OpenShift Container Platform 4.10 **oc** CLI with an OpenShift Container Platform 4.11 cluster, it is no longer possible to get a token by using the **oc serviceaccounts get-token <service\_account>** command. ([BZ#2097830](#)) ([BZ#2109799](#))
- There is currently a known issue in the **Add Helm Chart Repositories** form to extend the Developer Catalog of a project. The **Quick Start** guides shows that you can add the **ProjectHelmChartRepository** CR in the desired namespace whereas it does not mention that to perform this you need permission from the kubeadmin. ([BZ#2054197](#))

- There is currently a known issue. If you create an instance of **ProjectHelmChartRepository** custom resource (CR) that uses TLS verification, you cannot list the repositories and perform Helm-related operations. There is currently no workaround for this issue. ([HELM-343](#))
- When running OpenShift Container Platform on bare-metal IBM Power, there is a known issue that the Petitboot bootloader is unable to populate boot configurations for some RHCOS live images. In such cases, after booting nodes with PXE to install RHCOS, the expected live image disk configuration might not be visible.

As a workaround, you can boot manually with **kexec** from the Petitboot shell.

Identify the disk that holds the live image, in this example **nvme0n1p3**, and run the following commands:

```
# cd /var/petitboot/mnt/dev/nvme0n1p3/ostree/rhcos-*/
# kexec -l vmlinuz-*.ppc64le -i initramfs-*.img -c "ignition.firstboot rd.neednet=1 ip=dhcp
$(grep options /var/petitboot/mnt/dev/nvme0n1p3/loader/entries/ostree-1-rhcos.conf | sed
's,^options ,,' )" && kexec -e
```

([BZ#2107674](#))

- In a disconnected environment, SRO does not try to fetch DTK from the main registry. It fetches from the mirror registry instead. ([BZ#2102724](#))
- The process counter displays incorrect information for the **phc2sys** process on the interface where **phc2sys** is not running. There is currently no workaround for this issue. ([OCPBUGSM-46005](#))
- When a network interface controller (NIC) in a node with a dual NIC PTP configuration is shut down, a faulty event is generated for both PTP interfaces. There is currently no workaround for this issue. ([OCPBUGSM-46004](#))
- In low-band systems, the system clock does not synchronize with the PTP ordinary clock after the grandmaster clock is disconnected for a few hours and recovered. There is currently no workaround for this issue. ([OCPBUGSM-45173](#))
- Previously, if you were using the OVN-Kubernetes cluster network provider, if a service with **type=LoadBalancer** was configured with **internalTrafficPolicy=cluster** set, then all traffic was routed to the default gateway, even if the host routing table included better routes to use. Now the best route is used rather than always using the default gateway. ([BZ#2060159](#))
- When an OVN cluster has more than 75 worker nodes, simultaneously creating 2000 or more services and route objects can cause pods created at the same time to hang in the **ContainerCreating** status. If this problem occurs, entering the **oc describe pod <podname>** command will show events with the following warning: `'FailedCreatePodSandBox...failed to configure pod interface: timed out waiting for OVS port binding (ovn-installed)'`. There is currently no workaround for this issue. ([BZ#2084062](#))
- There is currently a known issue with OVN-Kubernetes that whenever the **NetworkManager** service restarts the node will lose network connectivity and must be recovered. ([BZ#2074009](#))
- The default security context constraint (SCC) might cause a pod using generic ephemeral volumes to remain in the **Pending** state. To work around this issue, you can create a custom SCC. For more information, see [Pods with Generic Ephemeral Volumes fail with SCC errors](#). For the workaround, see [Allowing the use of generic ephemeral volumes](#) ([BZ#2100429](#))
- If you have OpenShift sandboxed containers, when upgrading a cluster, you might encounter an issue where the Machine Config Operator (MCO) pod changes to a **CrashLoopBackOff** state



and the **openshift.io/scc** annotation of the pod displays **sandboxed-containers-operator-scc** instead of the default **hostmount-anyuid** value.

If this happens, temporarily change the **seLinuxOptions** strategy in the **sandboxed-containers-operator-scc** SCC to the less restrictive **RunAsAny**, so that the admission process does not prefer it over the **hostmount-anyuid** SCC.

1. Change the **seLinuxOptions** strategy by running the following command:

```
$ oc patch scc sandboxed-containers-operator-scc --type=merge --patch
'{"seLinuxContext":{"type": "RunAsAny"}}'
```

2. Restart the MCO pod by running the following commands:

```
$ oc scale deployments/machine-config-operator -n openshift-machine-config-operator --
replicas=0
```

```
$ oc scale deployments/machine-config-operator -n openshift-machine-config-operator --
replicas=1
```

3. Revert the **seLinuxOptions** strategy of the **sandboxed-containers-operator-scc** to its original value of **MustRunAs** by running the following command:

```
$ oc patch scc sandboxed-containers-operator-scc --type=merge --patch
'{"seLinuxContext":{"type": "MustRunAs"}}'
```

4. Verify that the **hostmount-anyuid** SCC is applied to the MCO pod by running the following command:

```
$ oc get pods -n openshift-machine-config-operator -l k8s-app=machine-config-operator -
o yaml | grep scc
openshift.io/scc: hostmount-anyuid
```

([KATA-1373](#))

- The pipeline metrics API does not support the required **pipelinerun/taskrun** histogram values from RHOSP 1.6 and beyond. Consequently, the **Metrics** tab in the **Pipeline → Details** page is removed instead of displaying incorrect values. There is currently no workaround for this issue. (link: [BZ#2074767](#))
- Some Alibabacloud services are not placing all resources of a cluster into a specified resource group. Consequently, some resources created by the OpenShift Container Platform installation program are placed in the default resource group. There is currently no workaround for this issue. ([BZ#2096692](#))
- After rebooting each cluster node, cluster Operators **network** and **kube-apiserver** turns to **degraded** after rebooting each node of the cluster and the cluster turns unhealthy. There is currently no workaround for this issue. ([BZ#2102011](#))
- When **resourceGroupID** is specified in **install-config.yaml**, an error is displayed when deleting bootstrap resources and OpenShift Container Platform installation fails. As a workaround for this issue, do not specify **resourceGroupID** in **install-config.yaml**. ([BZ#2100746](#))

- There is a known issue with scale-up on RHEL compute nodes. The new nodes could turn into **Ready**, but Ingress pods cannot turn into **Running** on these nodes, and scale-up does not succeed. As a workaround, scale-up with RHCOS nodes does work. ([BZ#2056387](#))
- After creating a **machineset** in a Google Cloud Platform (GCP), the **capg-controller-manager** machine is stuck in **Provisioning**. There is currently no workaround for this issue. ([BZ#2107999](#))
- There is a known issue on Nutanix where persistent volumes (PVs) created by a cluster are not cleaned up by the **destroy cluster** command. As a workaround for this issue, you need to clean up the PVs manually. ([BZ#2108700](#))
- There is a known issue with Nutanix installation where the installation fails if you use 4096-bit certificates with Prism Central 2022.x. Instead, use 2048-bit certificates. ([KCS](#))
- There is currently a known issue when creating and editing **egressqos** with incorrect syntax or values success. The incorrect values in **egressqos** should not create successfully. There is currently no workaround for this issue. ([BZ#2097579](#))
- Due to the inclusion of old images in some image indexes, running **oc adm catalog mirror** and **oc image mirror** might result in the following error: **error: unable to retrieve source image**. As a temporary workaround, you can use the **--skip-missing** option to bypass the error and continue downloading the image index. For more information, see [Service Mesh Operator mirroring failed](#).
- It is not possible to create a macvlan on the physical function (PF) when a virtual function (VF) already exists. This issue affects the Intel E810 NIC. ([BZ#2120585](#))
- If a cluster that was deployed through ZTP has policies that do not become compliant, and no **ClusterGroupUpdates** object is present, you must restart the TALM pods. Restarting TALM creates the proper **ClusterGroupUpdates** object, which enforces the policy compliance. ([OCPBUGS-4065](#))
- Currently, a certificate compliance issue, specifically outputted as **x509: certificate is not standards compliant**, exists when you run the installation program on macOS for the purposes of installing an OpenShift Container Platform cluster on VMware vSphere. This issue relates to a known issue with the **golang** compiler in that the compiler does not recognize newly supported macOS certificate standards. No workaround exists for this issue. ([OSDOCS-5694](#))
- Currently, when using a persistent volume (PV) that contains a very large number of files, the pod might not start or can take an excessive amount of time to start. For more information, see this [knowledge base article](#). ([BZ1987112](#))

## 1.9. ASYNCHRONOUS ERRATA UPDATES


Security, bug fix, and enhancement updates for OpenShift Container Platform 4.11 are released as asynchronous errata through the Red Hat Network. All OpenShift Container Platform 4.11 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 through 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 4.11. Versioned asynchronous releases, for example with the form OpenShift Container Platform 4.11.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-2022:5069 - OpenShift Container Platform 4.11.0 image release, bug fix, and security update advisory

Issued: 2022-08-10

OpenShift Container Platform release 4.11.0, which includes security updates, is now available. The list of bug fixes that are included in the update is documented in the [RHSA-2022:5069](#) advisory. The RPM packages that are included in the update are provided by the [RHSA-2022:5068](#) advisory.

You can view the container images in this release by running the following command:

```
$ oc adm release info 4.11.0 --pullspecs
```

### 1.9.2. RHSA-2022:6103 - OpenShift Container Platform 4.11.1 bug fix and security update

Issued: 2022-08-23

OpenShift Container Platform release 4.11.1, which includes security updates, is now available. The list of bug fixes that are included in the update is documented in the [RHSA-2022:6103](#) advisory. The RPM packages that are included in the update are provided by the [RHSA-2022:6102](#) advisory.

You can view the container images in this release by running the following command:

```
$ oc adm release info 4.11.1 --pullspecs
```

#### 1.9.2.1. Features

##### 1.9.2.1.1. General availability of pod-level bonding for secondary networks

With this update, [Using pod-level bonding](#) is now generally available.

##### 1.9.2.2. Bug fixes

- Previously, the functionality of Bond-CNI was limited to only active-backup mode. In this update, the supported bonding modes are as follows:
  - **balance-rr** - 0
  - **active-backup** - 1
  - **balance-xor** - 2

([BZ#2102047](#))

### 1.9.2.3. Updating

To update an existing OpenShift Container Platform 4.11 cluster to this latest release, see [Updating a cluster using the CLI](#).

## 1.9.3. RHBA-2022:6143 - OpenShift Container Platform 4.11.2 bug fix update

Issued: 2022-08-29

OpenShift Container Platform release 4.11.2 is now available. The list of bug fixes that are included in the update is documented in the [RHBA-2022:6143](#) advisory. The RPM packages that are included in the update are provided by the [RHBA-2022:6142](#) advisory.

You can view the container images in this release by running the following command:

```
$ oc adm release info 4.11.2 --pullspecs
```

### 1.9.3.1. Updating

To update an existing OpenShift Container Platform 4.11 cluster to this latest release, see [Updating a cluster using the CLI](#).

## 1.9.4. RHSA-2022:6287 - OpenShift Container Platform 4.11.3 bug fix update and security update

Issued: 2022-09-06

OpenShift Container Platform release 4.11.3, which includes security updates, is now available. The list of bug fixes that are included in the update is documented in the [RHSA-2022:6287](#) advisory. The RPM packages that are included in the update are provided by the [RHBA-2022:6286](#) advisory.

You can view the container images in this release by running the following command:

```
$ oc adm release info 4.11.3 --pullspecs
```

### 1.9.4.1. Features

## 1.9.5. Scalability and performance

Beginning with OpenShift Container Platform 4.11.3, users no longer need to set the Root FS image URL (**rootFSUrl**) in the **agent\_service\_config.yaml** file. The **rootFSUrl** is now handled automatically.

### 1.9.5.1. Updating

To update an existing OpenShift Container Platform 4.11 cluster to this latest release, see [Updating a cluster using the CLI](#).

## 1.9.6. RHBA-2022:6376 - OpenShift Container Platform 4.11.4 bug fix update

Issued: 2022-09-12

OpenShift Container Platform release 4.11.4 is now available. The list of bug fixes that are included in the update is documented in the [RHBA-2022:6376](#) advisory. The RPM packages that are included in the update are provided by the [RHBA-2022:6375](#) advisory.

You can view the container images in this release by running the following command:

```
$ oc adm release info 4.11.4 --pullspecs
```

### 1.9.6.1. Updating

To update an existing OpenShift Container Platform 4.11 cluster to this latest release, see [Updating a cluster using the CLI](#).

## 1.9.7. RHSA-2022:6536 - OpenShift Container Platform 4.11.5 bug fix and security update

Issued: 2022-09-20

OpenShift Container Platform release 4.11.5, which includes security updates, is now available. The list of bug fixes that are included in the update is documented in the [RHSA-2022:6536](#) advisory. The RPM packages that are included in the update are provided by the [RHSA-2022:6535](#) advisory.

You can view the container images in this release by running the following command:

```
$ oc adm release info 4.11.5 --pullspecs
```

### 1.9.7.1. Known issues

- Sharding the default Ingress Controller breaks the OpenShift Container Platform factory routes such as **canary**, **oauth**, and **console**. As a workaround, you can manually add matching labels and expressions to the routes. ([BZ#2024946](#))

### 1.9.7.2. Bug fixes

- Previously, an update of the **routeSelector** cleared the route status of the Ingress Controller prior to the router deployment. As a result, the route status repopulated incorrectly. With this update, the route status is cleared with a **routeSelector** update. ([BZ#2110528](#))

### 1.9.7.3. Updating

To update an existing OpenShift Container Platform 4.11 cluster to this latest release, see [Updating a cluster using the CLI](#).

## 1.9.8. RHBA-2022:6659 - OpenShift Container Platform 4.11.6 bug fix update

Issued: 2022-09-28

OpenShift Container Platform release 4.11.6 is now available. The list of bug fixes that are included in the update is documented in the [RHBA-2022:6659](#) advisory. The RPM packages that are included in the update are provided by the [RHBA-2022:6658](#) advisory.

You can view the container images in this release by running the following command:

```
$ oc adm release info 4.11.6 --pullspecs
```

### 1.9.8.1. OpenShift Container Platform 4.11 RAN new features

#### 1.9.8.1.1. Recovering clusters after a failed update

For single-node OpenShift, you can use the Topology Aware Lifecycle Manager (TALM) Operator to create a backup of a current deployment before an OpenShift Container Platform version update. If the update fails, use the backup to return the cluster to its pre-update state. For more information, see [Creating a backup of cluster resources before upgrade](#).

#### 1.9.8.1.2. Disable chronyd in the `PolicyGenTemplate` custom resource (CR)

You must disable **chronyd** if you update to OpenShift Container Platform 4.11 from earlier versions. To disable **chronyd**, add the following line in the `[service]` section under `.spec.profile.data` of the **TunedPerformancePatch.yaml** file. The **TunedPerformancePatch.yaml** file is referenced in the group **PolicyGenTemplate** CR:

```
[service]
service.chronyd=stop,disable
```

For more information, see [Recommended cluster kernel configuration](#).

### 1.9.8.2. OpenShift Container Platform 4.11 RAN known issues

- After deploying a single-node OpenShift cluster with the RAN DU profile, an error **openvswitch: cpu\_id mismatch with handler threads** is continuously generated in the Open vSwitch kernel logs. ([OCPBUGSM-46165](#))
- If secure boot is currently disabled and you try to enable it using ZTP, the cluster installation does not start. When secure boot is enabled through ZTP, the boot options are configured before the virtual CD is attached. Therefore, the first boot from the existing hard disk has the secure boot turned on. The cluster installation gets stuck because the system never boots from the CD. ([OCPBUGSM-45085](#))
- If an invalid subscription channel is specified in the subscription policy that is used to perform a cluster upgrade, the Topology Aware Lifecycle Manager indicates a successful upgrade right after the policy is enforced because the **Subscription** state remains **AtLatestKnown**. ([OCPBUGSM-43618](#))
- If the SRIOV-FEC Operator is installed through separate policies and the **CatalogSource** reuses the name of a default catalog source, the Operator installation might fail due to conflicts with the management of the default catalog sources. To avoid this issue, the SRIOV-FEC **CatalogSource** CR should be added to the **common-config-policy** and the Operator

subscription should be added to the **common-subscriptions-policy**. As a workaround when using separate policies to install the SRIOV-FEC Operator, you must ensure that the **CatalogSource** for this Operator is uniquely named. ([OCBUGSM-39859](#))

- The **SiteConfig** disk partition definition fails when applied to multiple nodes in a cluster. When a **SiteConfig** CR is used to provision a compact cluster, creating a valid **diskPartition** config on multiple nodes fails with a Kustomize plugin error. ([OCBUGSM-44403](#))
- When you patch the **ArgoCD** resource with the ZTP container, the patch points to a tag which follows the latest container version for that release version. If you want to pin the ZTP container to a specific version within the release, the patch file **argocd-openshift-gitops-patch.json** should be updated to point to the specific version. ([OCBUGSM-44261](#))
- Applying a **BMCEventSubscription** CR fails to create a Redfish events subscription. After the subscription YAML file is created and applied, no active Redfish subscription is visible. As a workaround, call the API directly and create the subscription. For example,
  1. Get the authentication token by running the following command:

```
$ curl -i --insecure --request POST --header "OData-Version: 4.0" \
--header "Content-Type: application/json" -d '{"UserName": <BMC_USERNAME>, \
"Password": <BMC_PASSWORD>}'
https://<BMC_IP>/redfish/v1/SessionService/Sessions/ |grep 'X-Auth-Token'
```

#### Example output

```
X-Auth-Token: 1234abcd5678efgh9012ijkl3456mnop
```

2. Using the authentication token, create the Redfish events subscription:

```
$ curl -X POST -i --insecure --header "X-Auth-Token:
1234abcd5678efgh9012ijkl3456mnop" \
-H 'Content-Type: application/json' --data-raw '{"Protocol": "Redfish", "Context": \
"Public", "Destination": "https://hw-event-proxy-openshift-hw-
events.apps.example.com/webhook", \
"EventTypes": ["Alert"]}' https://<BMC_IP>/redfish/v1/EventService/Subscriptions
```

You should receive a **201 Created** response and a header with **Location:** [https://<BMC\\_IP>/redfish/v1/EventService/Subscriptions/35](#) that indicates that the Redfish events subscription is successfully created. ([OCBUGSM-43707](#))

- When using the GitOps ZTP pipeline to install a single-node OpenShift cluster in a disconnected environment, there should be two **CatalogSource** CRs applied in the cluster. One of the **CatalogSource** CRs gets deleted following multiple node reboots. As a workaround, you can change the default names, such as **certified-operators** and **redhat-operators**, of the catalog sources. ([OCBUGSM-46245](#))
- Several clusters fail to update during scale testing. After starting an OpenShift Container Platform 4.9.26 update with the Telecom vDU configuration applied, a cluster update to 4.10.13 initiated with the **ClusterVersion** CR fails at 15% and waits for cluster Operators updating to the target version. ([OCBUGSM-44655](#))
- Sometimes, during ZTP GitOps pipeline single-node OpenShift installations, the Operator Lifecycle Manager registry-server container fails to reach the **READY** state. When creating subscriptions with a new **CatalogSource**, the **CatalogSource** remains in a

**TRANSIENT\_FAILURE** state. ([OCPBUGSM-44041](#))

- When you apply a tuned override to a pod and delete the tuned pod to force a restart, the pod should restart and the system should run normally. Instead, a **systemd** hang can occur, causing the system to stop responding. ([RHELPLAN-131021](#))
- When booting ZT Systems machines from a live ISO with static IPv6 address configuration, **NetworkManager** exits successfully before the interface link becomes ready. This leaves the network interface without a configuration. As a workaround, edit the RHCOS ISO referenced in the **AgentServiceConfig** CR to append **rd.net.timeout.carrier** to the kernel parameters in the **grub.cfg** file:

1. Pull the **rhcos-live** ISO image for the release being installed. The URL can be retrieved from the **AgentServiceConfig** CR on the hub cluster by running the following command:

```
$ oc get AgentServiceConfig agent -o yaml
```

2. Mount the image in the **/mnt/iso/** directory:

```
$ mount rhcos-live.x86_64.iso /mnt/iso/
```

3. Create the **iso-grub-cfg/** directory and change to the directory:

```
$ mkdir iso-grub-cfg/; pushd iso-grub-cfg/
```

4. Copy the contents of the **/mnt/iso/** directory to your working directory:

```
$ rsync -avH /mnt/iso/* .
```

5. Open the GRUB configuration file:

```
$ vim EFI/redhat/grub.cfg
```

- a. Append the **rd.net.timeout.carrier=20** string to the **linux** boot line.

6. Return to the initial working directory by running the following command:

```
$ popd
```

7. Generate the ISO file from the **iso-grub-cfg** directory:

```
$ mkisofs -JR -graft-points -o rhcos-carrier-timeout.iso iso-grub-cfg
```

8. Push the updated ISO image to the server that is accessible by the hub cluster.

9. On the hub cluster, update the **osImages** entry in the **AgentServiceConfig** CR for the installing release to point to the updated ISO image:

```
$ oc edit AgentServiceConfig agent
```

10. Update the **url** field to point to the URL of the updated ISO image.

([OCPBUGSM-46336](#))



- A kernel failure occurs after rebooting a bare-metal SNO with DU profile and workload test application. As a workaround, you can add an additional kernel parameter to the performance profile:

```
apiVersion: performance.openshift.io/v2
kind: PerformanceProfile
spec:
  additionalKernelArgs:
    - rcutree.kthread_prio=11
```

([RHELPLAN-123262](#))

- During ZTP cluster deployments, bare-metal host image provisioning might fail with an error referencing the HTTP 412 error code:

```
Deploy step deploy.deploy failed with HTTPError: HTTP PATCH
https://10.16.230.34/redfish/v1/Managers/1/VirtualMedia/EXT1 returned code 412.
Base.1.8.GeneralError: A general error has occurred. See ExtendedInfo for more
information. Extended information: [
```

```
{'MessageSeverity': 'Critical', 'MessageArgs': [], 'MessageId': 'Base.1.8.PreconditionFailed',
'Resolution': 'Try the operation again using the appropriate ETag.', '@odata.type':
'#Message.v1_1_0.Message', 'Message': 'The ETag supplied did not match the ETag
required to change this resource.'}]
```

This issue might affect various server models, including Gen8 HP machines and Gen9 HP machines running older firmware. For Gen9 HP machines, upgrading to the latest iLO firmware might resolve the issue. For Gen8 HP and others machines, there is currently no workaround for this issue. ([OCPBUGS-1246](#))

- On certain Lenovo models, for example the SE450, bare-metal host image provisioning during ZTP cluster deployments might fail with the HTTP 400 status code and a **PropertyNotWriteable** error:

```
HTTP response for PATCH https://192.168.26.178/redfish/v1/Systems/1/Pending: status
code: 400, error: Base.1.8.GeneralError: A general error has occurred. See ExtendedInfo for
more information., extended: [{'MessageArgs': ['BootSourceOverrideEnabled'], 'Resolution':
'Remove the property from the request body and resubmit the request if the operation failed.',
'MessageId': 'Base.1.8.PropertyNotWritable', 'Message': 'The property
BootSourceOverrideEnabled is a read only property and cannot be assigned a value.',
 '@odata.type': '#Message.v1_1_0.Message', 'MessageSeverity': 'Warning'}]
```

Currently, there is no workaround for this issue. ([OCPBUGSM-46305](#))

- During a primary node replacement in a bare-metal cluster, the new primary host gets stuck in a **Provisioning** state, but the node reports into the cluster as **Ready**. ([OCPBUGSM-45772](#))
- Using Red Hat Advanced Cluster Management (RHACM), spoke cluster deployments on Dell PowerEdge R640 servers are blocked when the virtual media does not disconnect the ISO in the iDRAC console after writing the image to the disk. As a workaround, disconnect the ISO manually through the Virtual Media tab in the iDRAC console. ([OCPBUGSM-45884](#))
- On a dual-stack networking environment, devices and connections get stuck in an **ip-check** state with a default **dhcp6** profile generated by the **nm-initrd-generator** utility. Due to this issue, the **/etc/resolve.conf** file is not generated. As a workaround, restart the

**NetworkManager** service. This generates the missing `/etc/resolve.conf` file and the installation can continue. ([RHELPLAN-127788](#), [OCBUGS-70](#))

- When using NVIDIA-branded Mellanox NICs on Dell hardware, incoming packets that are larger than the preset F5 application receive buffer (currently set to 8K) arrive with an incorrect VLAN tag. This results in irregular truncated packets being delivered by the NIC. ([RHELPLAN-123058](#))
- Sometimes, a single-node OpenShift node configured with static IP and deployed using the GitOps ZTP pipeline becomes unreachable during Day 2 operator configuration. The OpenShift Container Platform cluster installation completes successfully, and the cluster is healthy. After reboot, the node becomes unreachable with the network interface down. ([OCBUGSM-46688](#))
- After removing the **SriovNetworkNodePolicy** policy from the Git repository, the **SriovNetworkNodePolicy** resource managed by the removed policy remains on the spoke cluster. ([OCBUGSM-34614](#))
- When upgrading the PTP Operator from 4.10 to 4.11, the OpenShift subscription reports the error no channel heads (entries not replaced by another entry) found in channel "**stable**" of package "**ptp-operator**". However, the Operator upgrades successfully. ([OCBUGSM-46114](#))
- Currently, all **PtpConfig** CRs in the ZTP source CRs includes the **phc2sysOpts** option. Therefore, even if the user does not include **phc2sysOpts** in the user **PolicyGenTemplate** CR, the **phc2sysOpts** option is added to the spoke PTP configuration. If PTP with dual NIC is configured through ZTP, the user needs to update one **PtpConfig** CR to remove the **phc2sysOpts** option after ZTP is completed. ([OCBUGSM-47798](#))
- When secure boot is enabled **stallid**, service fails to start because it is not able to open the `/sys/kernel/debug/sched_features` file. ([OCBUGSM-1466](#))
- When secure boot is enabled, the **kdump** service might fail to start with **kexec: failed to load kdump kernel** error. To workaround this issue, add ``efi=runtime`` to the kernel arguments. ([OCBUGSM-97](#))
- If an SNO cluster is upgraded from OCP 4.10 to OCP 4.11, the SNO cluster might get rebooted three times during the upgrade process. ([OCBUGSM-46704](#))
- If a Supermicro server is deployed through ZTP, the incorrect boot device might be chosen preventing the install from starting. ([OCBUGSM-369](#))
- The default **dns-default** pod is missing the **"target.workload.openshift.io/management:"** annotation. As a result, when the workload partitioning feature is enabled on SNO, the pod resources do not get mutated and pinned to the reserved CPU set. As a workaround, the cluster administrator can add the annotation manually using the following command:

```
$ oc annotate pod dns-default
target.workload.openshift.io/management='{\"effect\":\"PreferredDuringScheduling\"}' -n
openshift-dns
```

([OCBUGSM-753](#))

### 1.9.8.3. Updating

To update an existing OpenShift Container Platform 4.11 cluster to this latest release, see [Updating a cluster using the CLI](#).

### 1.9.9. RHSA-2022:6732 - OpenShift Container Platform 4.11.7 bug fix update

Issued: 2022-10-03

OpenShift Container Platform release 4.11.7 is now available. The list of bug fixes that are included in the update is documented in the [RHBA-2022:6732](#) advisory. There are no RPM packages for this update.

You can view the container images in this release by running the following command:

```
$ oc adm release info 4.11.7 --pullspecs
```

#### 1.9.9.1. Updating

To update an existing OpenShift Container Platform 4.11 cluster to this latest release, see [Updating a cluster using the CLI](#).

### 1.9.10. RHBA-2022:6809 - OpenShift Container Platform 4.11.8 bug fix update

Issued: 2022-10-12

OpenShift Container Platform release 4.11.8 is now available. The list of bug fixes that are included in the update is documented in the [RHBA-2022:6809](#) advisory. The RPM packages that are included in the update are provided by the [RHBA-2022:6808](#) advisory.

You can view the container images in this release by running the following command:

```
$ oc adm release info 4.11.8 --pullspecs
```

#### 1.9.10.1. Bug fixes

- Previously, in the monitoring stack, if you have enabled and deployed a dedicated Alertmanager instance for user-defined alerts, you cannot silence alerts in the Developer perspective in the OpenShift Container Platform web console. With this update, user-defined alerts can be silenced from the Developer perspective. ([OCPBUGS-1790](#))

#### 1.9.10.2. Updating

To update an existing OpenShift Container Platform 4.11 cluster to this latest release, see [Updating a cluster using the CLI](#).

### 1.9.11. RHBA-2022:6897 - OpenShift Container Platform 4.11.9 bug fix update

Issued: 2022-10-17

OpenShift Container Platform release 4.11.9 is now available. The list of bug fixes that are included in the update is documented in the [RHBA-2022:6897](#) advisory. The RPM packages that are included in the update are provided by the [RHBA-2022:6896](#) advisory.

You can view the container images in this release by running the following command:

```
$ oc adm release info 4.11.9 --pullspecs
```

#### 1.9.11.1. Updating

To update an existing OpenShift Container Platform 4.11 cluster to this latest release, see [Updating a cluster using the CLI](#).

### 1.9.12. RHSA-2022:7201 - OpenShift Container Platform 4.11.12 bug fix and security update

Issued: 2022-11-02

OpenShift Container Platform release 4.11.12, which includes security updates, is now available. The list of bug fixes that are included in the update is documented in the [RHSA-2022:7201](#) advisory. The RPM packages that are included in the update are provided by the [RHBA-2022:7200](#) advisory.

You can view the container images in this release by running the following command:

```
$ oc adm release info 4.11.12 --pullspecs
```

#### 1.9.12.1. Known issues

- If a cluster is incrementally updated from a version less than or equal to 4.9, the **openshift-dns namespace** may not contain the required pod-security labels required for future version updates. ([OCPBUGS-1549](#))

#### 1.9.12.2. Notable technical changes

- With this release, when the service account issuer is changed to a custom one, existing bound service tokens are no longer invalidated immediately. Instead, when the service account issuer is changed, the previous service account issuer continues to be trusted for 24 hours.

For more information, see [Configuring bound service account tokens using volume projection](#) .

#### 1.9.12.3. Updating

To update an existing OpenShift Container Platform 4.11 cluster to this latest release, see [Updating a cluster using the CLI](#).

### 1.9.13. RHBA-2022:7201 - OpenShift Container Platform 4.11.13 bug fix update

Issued: 2022-11-09

OpenShift Container Platform release 4.11.13 is now available. The list of bug fixes that are included in the update is documented in the [RHBA-2022:7290](#) advisory. The RPM packages that are included in the update are provided by the [RHBA-2022:7289](#) advisory.

You can view the container images in this release by running the following command:

```
$ oc adm release info 4.11.13 --pullspecs
```

#### 1.9.13.1. Notable technical changes

- The Cloud Credential Operator utility (**ccctl**) now creates secrets that use regional endpoints for the [AWS Security Token Service \(AWS STS\)](#) . This approach aligns with AWS recommended best practices.

### 1.9.13.2. Updating

To update an existing OpenShift Container Platform 4.11 cluster to this latest release, see [Updating a cluster using the CLI](#).

## 1.9.14. RHSA-2022:8535 - OpenShift Container Platform 4.11.16 bug fix and security update

Issued: 2022-11-24

OpenShift Container Platform release 4.11.16, which includes security updates, is now available. There is no IBM powerbuild for this release. The list of bug fixes that are included in the update is documented in the [RHSA-2022:8535](#) advisory. The RPM packages that are included in the update are provided by the [RHSA-2022:8534](#) advisory.

You can view the container images in this release by running the following command:

```
$ oc adm release info 4.11.16 --pullspecs
```

### 1.9.14.1. Notable technical changes

- With this release, when you [delete GCP resources with the Cloud Credential Operator utility](#), you must specify the directory containing the files for the component **CredentialsRequest** objects.

### 1.9.14.2. Bug fixes

- Previously, if you used upper-case letters when specifying an Azure Disk Encryption Set (DES) or Resource Group (RG) name, the validation would fail. With this release, you can now use upper and lower case letters in DES and RG names. ([OCPBUGS#4826](#))

### 1.9.14.3. Updating

To update an existing OpenShift Container Platform 4.11 cluster to this latest release, see [Updating a cluster using the CLI](#).

## 1.9.15. RHBA-2022:8627 - OpenShift Container Platform 4.11.17 bug fix and security update

Issued: 2022-11-28

OpenShift Container Platform release 4.11.17, which includes security updates, is now available. There is no IBM powerbuild for this release. The list of bug fixes that are included in the update is documented in the [RHBA-2022:8627](#) advisory. The RPM packages that are included in the update are provided by the [RHSA-2022:8626](#) advisory.

You can view the container images in this release by running the following command:

```
$ oc adm release info 4.11.17 --pullspecs
```

### 1.9.15.1. Updating

To update an existing OpenShift Container Platform 4.11 cluster to this latest release, see [Updating a cluster using the CLI](#).

## 1.9.16. RHBA-2022:8698 - OpenShift Container Platform 4.11.18 bug fix update

Issued: 2022-12-05

OpenShift Container Platform release 4.11.18 is now available. The list of bug fixes that are included in the update is documented in the [RHBA-2022:8698](#) advisory. The RPM packages that are included in the update are provided by the [RHBA-2022:8697](#) advisory.

You can view the container images in this release by running the following command:

```
$ oc adm release info 4.11.18 --pullspecs
```

### 1.9.16.1. Enhancements

- IPv6 unsolicited neighbor advertisements and IPv4 gratuitous address resolution protocol now default on the SR-IOV CNI plugin. Pods created with the Single Root I/O Virtualization (SR-IOV) CNI plugin, where the IP address management CNI plugin has assigned IPs, now send IPv6 unsolicited neighbor advertisements and/or IPv4 gratuitous address resolution protocol by default onto the network. This enhancement notifies hosts of the new pod's MAC address for a particular IP to refresh ARP/NDP caches with the correct information. For more information, see [Supported devices](#).

### 1.9.16.2. Notable technical changes

- Previously named heterogeneous clusters are now referred to as multi-architecture in OpenShift Container Platform documentation. For more information, see [Configuring a multi-architecture cluster](#)

### 1.9.16.3. Bug fixes

- Previously, some object storage instances responded with **204 No Content** when no content displayed. The Red Hat OpenStack Platform (RHOSP) SDK used in OpenShift Container Platform did not handle 204s correctly. With this update, the installation program works around the issue when there are zero items to list. ([OCPBUGS-4081](#))
- Previously, the rollout times for **kube-apiserver** on a loaded cluster were slow and sometimes exceeded the 5 min rollout timeout. With this update, the rollout times are shorter and within the 5 min threshold. ([OCPBUGS-3182](#))

### 1.9.16.4. Updating

To update an existing OpenShift Container Platform 4.11 cluster to this latest release, see [Updating a cluster using the CLI](#).

## 1.9.17. RHSA-2022:8893 - OpenShift Container Platform 4.11.20 bug fix and security update

Issued: 2022-12-15

OpenShift Container Platform release 4.11.20, which includes security updates, is now available. The list of bug fixes that are included in the update is documented in the [RHSA-2022:8893](#) advisory. The RPM packages that are included in the update are provided by the [RHBA-2022:8892](#) advisory.

You can view the container images in this release by running the following command:

```
$ oc adm release info 4.11.20 --pullspecs
```

### 1.9.17.1. Bug fixes

- Previously, the OpenShift Container Platform installation program presented the user with an incomplete list of regions when installing a cluster on Google Cloud Platform (GCP). With this update, the installation program includes all supported regions. ([OCPBUGS-3023](#))

### 1.9.17.2. Known issues

- Switching the router scope of the default Ingress Controller by setting the **spec.endpointPublishingStrategy.loadBalancer.scope** field results in a degraded Ingress Operator. Consequently, routes that use that endpoint such as the web console URL become inaccessible. As a workaround, restarting one of the router pods brings back several instances under the **loadbalancer** back to the **inService** status. ([OCPBUGS-2554](#))

### 1.9.17.3. Updating

To update an existing OpenShift Container Platform 4.11 cluster to this latest release, see [Updating a cluster using the CLI](#).

## 1.9.18. RHSA-2022:9107 - OpenShift Container Platform 4.11.21 bug fix and security update

Issued: 2023-01-04

OpenShift Container Platform release 4.11.21, which includes security updates, is now available. The list of bug fixes that are included in the update is documented in the [RHSA-2022:9107](#) advisory. There are no RPM packages for this release.

You can view the container images in this release by running the following command:

```
$ oc adm release info 4.11.21 --pullspecs
```

### 1.9.18.1. Bug fixes

- Previously, after rotating Red Hat OpenStack Platform (RHOSP) credentials, the Cinder CSI driver continued to use old credentials until it was restarted out of band. If the old credentials were no longer valid, all volume operations failed. With this update, the Cinder CSI driver is updated automatically when the RHOSP credentials are rotated. ([OCPBUGS-4103](#))
- Previously, in CoreDNS v1.7.1, all upstream cache refreshes used DNSSEC. Bufsize was hardcoded to 2048 bytes for the upstream query, causing some DNS upstream queries to break when there were UDP Payload limits within the networking infrastructure. With this update, OpenShift Container Platform always uses bufsize 512 for upstream cache requests as that is the bufsize specified in the Corefile. Customers might be impacted if they rely on the incorrect functionality of bufsize 2048 for upstream DNS requests. ([OCPBUGS-2901](#))

- Previously, availability sets were created when **vmSize** was invalid in a region that has zones. However, availability sets should only be created in regions that do not have zones. With this update, the correct **vmSize** is provided and an availability set is no longer provided for a machine set. ([OCPBUGS-2123](#))

### 1.9.18.2. Updating

To update an existing OpenShift Container Platform 4.11 cluster to this latest release, see [Updating a cluster using the CLI](#).

## 1.9.19. RHBA-2023:0027 - OpenShift Container Platform 4.11.22 bug fix update

Issued: 2023-01-09

OpenShift Container Platform release 4.11.22 is now available. The list of bug fixes that are included in the update is documented in the [RHBA-2023:0027](#) advisory. There are no RPM packages for this release.

You can view the container images in this release by running the following command:

```
$ oc adm release info 4.11.22 --pullspecs
```

### 1.9.19.1. Bug fixes

- The OpenShift Container Platform 4.11 release updates the **install-config.yaml** file to now list the **me-west1**, (Tel Aviv, Israel) region. After you install the OpenShift Container Platform by running the **openshift-install** binary, you can select the **me-west1** region for your chosen cluster. ([OCPBUGS-4720](#))
- Previously, some object storage instances responded with a **204 No Content** error message when content was supposed to display. The Red Hat OpenStack Platform (RHOSP) SDK used in OpenShift Container Platform does not handle 204s correctly. With this update, the installation program works around the issue when there are zero objects to list in a Swift container. ([OCPBUGS-5078](#))
- Previously, the deployment of OpenShift Container Platform 4.11.ec4 build failed with the latest RHCOS image **412.86.202210072120-0** and **rhel-86** image. As a result, the Red Hat Enterprise Linux CoreOS (RHCOS) node is stuck at booting. With this update, the deployment completes successfully. ([OCPBUGS-2321](#))

### 1.9.19.2. Known issues

- Any 4.11 and later **arm64** cluster that has more than around 470 containers to a node, can cause additional pod creation to fail with the following error:

```
runc create failed: unable to start container process: unable to init seccomp: error loading seccomp filter into kernel: error loading seccomp filter: errno 524"
```

This is due to a CoreOS limitation in the number of seccomp profiles that can be created on the worker node. This most likely occurs on clusters with multiple containers per pod during an upgrade, worker node failure, or pod scaleup. This will be fixed in a later version of OpenShift Container Platform. ([OCPBUGS-2637](#))

### 1.9.19.3. Updating



To update an existing OpenShift Container Platform 4.11 cluster to this latest release, see [Updating a cluster using the CLI](#).

### 1.9.20. RHSA-2023:0069 - OpenShift Container Platform 4.11.24 bug fix and security update

Issued: 2023-01-19

OpenShift Container Platform release 4.11.24, which includes security updates, is now available. The list of bug fixes that are included in the update is documented in the [RHSA-2023:0069](#) advisory. The RPM packages that are included in the update are provided by the [RHBA-2023:0068](#) advisory.

You can view the container images in this release by running the following command:

```
$ oc adm release info 4.11.24 --pullspecs
```

#### 1.9.20.1. Updating

To update an existing OpenShift Container Platform 4.11 cluster to this latest release, see [Updating a cluster using the CLI](#).

### 1.9.21. RHSA-2023:0245 - OpenShift Container Platform 4.11.25 bug fix and security update

Issued: 2023-01-23

OpenShift Container Platform release 4.11.25, which includes security updates, is now available. The list of bug fixes that are included in the update is documented in the [RHSA-2023:0245](#) advisory. The RPM packages that are included in the update are provided by the [RHBA-2023:0244](#) advisory.

You can view the container images in this release by running the following command:

```
$ oc adm release info 4.11.25 --pullspecs
```

#### 1.9.21.1. Updating

To update an existing OpenShift Container Platform 4.11 cluster to this latest release, see [Updating a cluster using the CLI](#).

### 1.9.22. RHSA-2023:0565 - OpenShift Container Platform 4.11.26 bug fix and security update

Issued: 2023-02-07

OpenShift Container Platform release 4.11.26, which includes security updates, is now available. The list of bug fixes that are included in the update is documented in the [RHSA-2023:0565](#) advisory. The RPM packages that are included in the update are provided by the [RHBA-2023:0564](#) advisory.

You can view the container images in this release by running the following command:

```
$ oc adm release info 4.11.26 --pullspecs
```

### 1.9.22.1. Known issues

- This release regresses functionality where SR-IOV secondary network interfaces configured with static MAC address attached to pods might not provide full network connectivity. This issue only impacts SR-IOV virtual functions based on Intel Ethernet Network Adapter X710 products (i40e/iavf Linux kernel drivers). For more information, see [OCBUGS-5139](#).

### 1.9.22.2. Bug fixes

- Previously, the **cluster-image-registry-operator** would default to using a persistent volume claim (PVC) when it failed to reach Swift. With this update, failure to connect to Red Hat OpenStack Platform (RHOSP) API or other incidental failures cause the **cluster-image-registry-operator** to retry the probe. During the retry, the default to PVC only occurs if the RHOSP catalog is correctly found, and it does not contain object storage; or alternatively, if RHOSP catalog is there and the current user does not have permission to list containers. ([OCBUGS-5578](#))
- Previously due to a missing definition for **spec.provider**, the **Operator details** page failed when trying to show **ClusterServiceVersion**. With this update, the user interface works without **spec.provider** and the **Operator detail page** does not fail. ([OCBUGS-6689](#))

### 1.9.22.3. Updating

To update an existing OpenShift Container Platform 4.11 cluster to this latest release, see [Updating a cluster using the CLI](#).

## 1.9.23. RHSA-2023:0651 - OpenShift Container Platform 4.11.27 bug fix and security update

Issued: 2023-02-15

OpenShift Container Platform release 4.11.27, which includes security updates, is now available. The list of bug fixes that are included in the update is documented in the [RHSA-2023:0651](#) advisory. The RPM packages that are included in the update are provided by the [RHBA-2023:0650](#) advisory.

You can view the container images in this release by running the following command:

```
$ oc adm release info 4.11.27 --pullspecs
```

### 1.9.23.1. Bug fixes

- Previously, the topology sidebar did not display updated information. When you updated the resources directly from the topology sidebar, you had to reopen the sidebar to see the changes. With this fix, the updated resources are displayed correctly. ([OCBUGS-5459](#))

### 1.9.23.2. Updating

To update an existing OpenShift Container Platform 4.11 cluster to this latest release, see [Updating a cluster using the CLI](#).

## 1.9.24. RHSA-2023:0774 - OpenShift Container Platform 4.11.28 bug fix and security update

Issued: 2023-02-21

OpenShift Container Platform release 4.11.28, which includes security updates, is now available. The list of bug fixes that are included in the update is documented in the [RHSA-2023:0774](#) advisory. The RPM packages that are included in the update are provided by the [RHBA-2023:0773](#) advisory.

You can view the container images in this release by running the following command:

```
$ oc adm release info 4.11.28 --pullspecs
```

#### 1.9.24.1. Updating

To update an existing OpenShift Container Platform 4.11 cluster to this latest release, see [Updating a cluster using the CLI](#).

### 1.9.25. RHSA-2023:0895 - OpenShift Container Platform 4.11.29 bug fix and security update

Issued: 2023-02-28

OpenShift Container Platform release 4.11.29, which includes security updates, is now available. The list of bug fixes that are included in the update is documented in the [RHSA-2023:0895](#) advisory. There are no RPM packages for this update.

You can view the container images in this release by running the following command:

```
$ oc adm release info 4.11.29 --pullspecs
```

#### 1.9.25.1. Updating

To update an existing OpenShift Container Platform 4.11 cluster to this latest release, see [Updating a cluster using the CLI](#).

### 1.9.26. RHSA-2023:1030 - OpenShift Container Platform 4.11.30 bug fix and security update

Issued: 2023-03-07

OpenShift Container Platform release 4.11.30, which includes security updates, is now available. The list of bug fixes that are included in the update is documented in the [RHSA-2023:1030](#) advisory. The RPM packages that are included in the update are provided by the [RHBA-2023:1029](#) advisory.

You can view the container images in this release by running the following command:

```
$ oc adm release info 4.11.30 --pullspecs
```

#### 1.9.26.1. Bug fixes

- Previously, when creating a **Secret**, the **Start Pipeline** model created an invalid JSON value. As a result, the **Secret** was unusable and the **PipelineRun** could fail. With this fix, the **Start Pipeline** model creates a valid JSON value for the Secret. Now, you can create valid secrets while starting a pipeline. ([OCPBUGS-7494](#))

#### 1.9.26.2. Updating

To update an existing OpenShift Container Platform 4.11 cluster to this latest release, see [Updating a cluster using the CLI](#).

### 1.9.27. RHSA-2023:1158 - OpenShift Container Platform 4.11.31 bug fix and security update

Issued: 2023-03-14

OpenShift Container Platform release 4.11.31, which includes security updates, is now available. The list of bug fixes that are included in the update is documented in the [RHSA-2023:1158](#) advisory. The RPM packages that are included in the update are provided by the [RHBA-2023:1157](#) advisory.

You can view the container images in this release by running the following command:

```
$ oc adm release info 4.11.31 --pullspecs
```

#### 1.9.27.1. Updating

To update an existing OpenShift Container Platform 4.11 cluster to this latest release, see [Updating a cluster using the CLI](#).

### 1.9.28. RHBA-2023:1296 - OpenShift Container Platform 4.11.32 bug fix and security update

Issued: 2023-03-22

OpenShift Container Platform release 4.11.32, which includes security updates, is now available. The list of bug fixes that are included in the update is documented in the [RHBA-2023:1296](#) advisory. The RPM packages that are included in the update are provided by the [RHBA-2023:1295](#) advisory.

You can view the container images in this release by running the following command:

```
$ oc adm release info 4.11.32 --pullspecs
```

#### 1.9.28.1. Updating

To update an existing OpenShift Container Platform 4.11 cluster to this latest release, see [Updating a cluster using the CLI](#).

### 1.9.29. RHBA-2023:1396 - OpenShift Container Platform 4.11.33 bug fix

Issued: 2023-03-28

OpenShift Container Platform release 4.11.33 is now available. The list of bug fixes that are included in the update is documented in the [RHBA-2023:1396](#) advisory. The RPM packages that are included in the update are provided by the [RHBA-2023:1395](#) advisory.

You can view the container images in this release by running the following command:

```
$ oc adm release info 4.11.33 --pullspecs
```

#### 1.9.29.1. Updating

To update an existing OpenShift Container Platform 4.11 cluster to this latest release, see [Updating a cluster using the CLI](#).

### 1.9.30. RHSA-2023:1504 – OpenShift Container Platform 4.11.34 bug fix and security update

Issued: 2023-04-04

OpenShift Container Platform release 4.11.34 is now available. The list of bug fixes that are included in the update is documented in the [RHSA-2023:1504](#) advisory. The RPM packages that are included in the update are provided by the [RHSA-2023:1503](#) advisory.

You can view the container images in this release by running the following command:

```
$ oc adm release info 4.11.34 --pullspecs
```

#### 1.9.30.1. Updating

To update an existing OpenShift Container Platform 4.11 cluster to this latest release, see [Updating a cluster using the CLI](#).

### 1.9.31. RHBA-2023:1650 – OpenShift Container Platform 4.11.35 bug fix

Issued: 2023-04-12

OpenShift Container Platform release 4.11.35 is now available. The list of bug fixes that are included in the update is documented in the [RHBA-2023:1650](#) advisory. The RPM packages that are included in the update are provided by the [RHBA-2023:1649](#) advisory.

You can view the container images in this release by running the following command:

```
$ oc adm release info 4.11.35 --pullspecs
```

#### 1.9.31.1. Bug fixes

- Previously, if the OpenStack **clouds.yaml** file was rotated, you needed to restart **machine-api-provider-openstack** to pick up new cloud credentials. As a result, the ability of a **MachineSet** to scale to zero could be affected. With this change, cloud credentials are no longer cached and **machine-api-provider-openstack** reads the corresponding secret when it is needed. ([OCBUGS-10954](#))

#### 1.9.31.2. Updating

To update an existing OpenShift Container Platform 4.11 cluster to this latest release, see [Updating a cluster using the CLI](#).

### 1.9.32. RHBA-2023:1733 – OpenShift Container Platform 4.11.36 bug fix

Issued: 2023-04-13

OpenShift Container Platform release 4.11.36 is now available. The bug fix that is included in the update is documented in the [RHBA-2023:1733](#) advisory. There are no RPM packages for this update.

You can view the container images in this release by running the following command:

```
$ oc adm release info 4.11.36 --pullspecs
```

### 1.9.32.1. Updating

All OpenShift Container Platform 4.11 users are advised that the only defect fixed in this release is limited to install time; therefore, there is no need to update previously installed clusters to this version.

## 1.9.33. RHBA-2023:1760 - OpenShift Container Platform 4.11.37 bug fix

Issued: 2023-04-19

OpenShift Container Platform release 4.11.37 is now available. The list of bug fixes that are included in the update is documented in the [RHBA-2023:1760](#) advisory. The RPM packages that are included in the update are provided by the [RHBA-2023:1759](#) advisory.

You can view the container images in this release by running the following command:

```
$ oc adm release info 4.11.37 --pullspecs
```

### 1.9.33.1. Updating

To update an existing OpenShift Container Platform 4.11 cluster to this latest release, see [Updating a cluster using the CLI](#).

## 1.9.34. RHBA-2023:1863 - OpenShift Container Platform 4.11.38 bug fix update

Issued: 2023-04-26

OpenShift Container Platform release 4.11.38 is now available. Bug fixes included in the update are listed in the [RHBA-2023:1863](#) advisory. RPM packages included in the update are provided by the [RHBA-2023:1862](#) advisory.

You can view the container images in this release by running the following command:

```
$ oc adm release info 4.11.38 --pullspecs
```

### 1.9.34.1. Updating

To update an existing OpenShift Container Platform 4.11 cluster to this latest release, see [Updating a cluster using the CLI](#).

## 1.9.35. RHSA-2023:2014 - OpenShift Container Platform 4.11.39 bug fix and security update

Issued: 2023-05-02

OpenShift Container Platform release 4.11.39, which includes security updates, is now available. Bug fixes included in the update are listed in the [RHSA-2023:2014](#) advisory. RPM packages included in the update are provided by the [RHBA-2023:2056](#) advisory.

You can view the container images in this release by running the following command:

```
$ oc adm release info 4.11.39 --pullspecs
```

### 1.9.35.1. Bug fixes

- Previously, the Ingress Operator failed to publish the **router-certs** secret on a cluster with a high number of Ingress Controllers because of the secret size limit of 1 MB. Consequently, the Authentication Operator, which uses the **router-certs** secret to access the cluster Ingress domain, might not have had the latest certificate to use for OAuth purposes. With this update, the Ingress Operator publishes the certificate and key only for the Ingress Controller that owns the cluster Ingress domain, so that the secret does not exceed the size limit. This update ensures that the Authentication Operator can read and use an up-to-date certificate for OAuth authentication. ([OCPBUGS-8000](#))

### 1.9.35.2. Updating

To update an existing OpenShift Container Platform 4.11 cluster to this latest release, see [Updating a cluster using the CLI](#).

## 1.9.36. RHBA-2023:2694 - OpenShift Container Platform 4.11.40 bug fix update

Issued: 2023-05-18

OpenShift Container Platform release 4.11.40 is now available. Bug fixes included in the update are listed in the [RHBA-2023:2694](#) advisory. RPM packages included in the update are provided by the [RHBA-2023:2693](#) advisory.

You can view the container images in this release by running the following command:

```
$ oc adm release info 4.11.40 --pullspecs
```

### 1.9.36.1. Bug fixes

- Previously, when you deleted a Knative (**kn**) service on the OpenShift web console, the associated **<kn-service-name>-github-webhook-secret** webhook was not deleted. If you attempted to recreate the Knative service, while retaining the same name as the original service, the operation would fail. With this update, when you delete a Knative (**kn**) service on the OpenShift web console, the associated webhook is deleted at the same time as the service. You can now recreate a Knative service with the same name as the deleted service without the operation failing. ([OCPBUGS-7949](#))

### 1.9.36.2. Updating

To update an existing OpenShift Container Platform 4.11 cluster to this latest release, see [Updating a cluster using the CLI](#).

## 1.9.37. RHBA-2023:3213 - OpenShift Container Platform 4.11.41 bug fix update

Issued: 2023-05-24

OpenShift Container Platform release 4.11.41 is now available. Bug fixes included in the update are listed in the [RHBA-2023:3213](#) advisory. RPM packages included in the update are provided by the [RHBA-2023:3212](#) advisory.

You can view the container images in this release by running the following command:

```
$ oc adm release info 4.11.41 --pullspecs
```

### 1.9.37.1. Updating

To update an existing OpenShift Container Platform 4.11 cluster to this latest release, see [Updating a cluster using the CLI](#).

## 1.9.38. RHSA-2023:3309 - OpenShift Container Platform 4.11.42 bug fix and security update

Issued: 2023-05-31

OpenShift Container Platform release 4.11.42 is now available. Bug fixes included in the update are listed in the [RHSA-2023:3309](#) advisory. RPM packages included in the update are provided by the [RHBA-2023:3308](#) advisory.

You can view the container images in this release by running the following command:

```
$ oc adm release info 4.11.42 --pullspecs
```

### 1.9.38.1. Updating

To update an existing OpenShift Container Platform 4.11 cluster to this latest release, see [Updating a cluster using the CLI](#).

## 1.9.39. RHSA-2023:3542 OpenShift Container Platform 4.11.43 bug fix and security update

Issued: 2023-06-14

OpenShift Container Platform release 4.11.43, which includes security updates, is now available. Bug fixes included in the update are listed in the [RHSA-2023:3542](#) advisory. RPM packages included in the update are provided by the [RHSA-2023:3541](#) advisory.

You can view the container images in this release by running the following command:

```
$ oc adm release info 4.11.43 --pullspecs
```

### 1.9.39.1. Updating

To update an existing OpenShift Container Platform 4.11 cluster to this latest release, see [Updating a cluster using the CLI](#).

## 1.9.40. RHSA-2023:3915 - OpenShift Container Platform 4.11.44 bug fixes and security update

Issued: 2023-07-06



OpenShift Container Platform release 4.11.44, which includes security updates, is now available. This update includes a Red Hat security bulletin for customers who run OpenShift Container Platform in FIPS mode. For more information, see [RHSB-2023:001](#).

Bug fixes included in the update are listed in the [RHSA-2023:3915](#) advisory. RPM packages included in the update are provided by the [RHSA-2023:3914](#) advisory.

You can view the container images in this release by running the following command:

```
$ oc adm release info 4.11.44 --pullspecs
```

#### 1.9.40.1. Bug fixes

- Previously, because client TLS (mTLS) was configured on an ingress controller, and the certificate authority (CA) in the client CA bundle required more than 1MB of certificate revocation lists (CRLs) to be downloaded, the CRL **ConfigMap** object could not be updated due to the size limitation on **ConfigMap** objects. As a result of the missing CRLs, connections with valid client certificates could have been rejected with the error **unknown ca**. With this update, the CRL **ConfigMap** object for each ingress controller no longer exists; instead, CRL **ConfigMap** objects are downloaded directly by each router pod, ensuring that connections with valid client certificates are no longer rejected. ([OCPBUGS-14456](#))
- Previously, because client TLS (mTLS) was configured on an ingress controller, mismatches between the distributing certificate authority (CA) and the issuing CA caused the incorrect certificate revocation list (CRL) to be downloaded. As a result, the incorrect CRL would be downloaded in place of the correct CRL causing connections with valid client certificates to be rejected with the error message **unknown ca**. With this update, downloaded CRLs are now tracked by the CA that distributes them. This ensures that valid client certificates are no longer rejected. ([OCPBUGS-14457](#))

#### 1.9.40.2. Updating

To update an existing OpenShift Container Platform 4.11 cluster to this latest release, see [Updating a cluster using the CLI](#).

### 1.9.41. RHSA-2023:4053 OpenShift Container Platform 4.11.45 bug fix and security update

Issued: 2023-07-19

OpenShift Container Platform release 4.11.45, which includes security updates, is now available. Bug fixes included in the update are listed in the [RHSA-2023:4053](#) advisory. RPM packages included in the update are provided by the [RHBA-2023:4052](#) advisory.

You can view the container images in this release by running the following command:

```
$ oc adm release info 4.11.45 --pullspecs
```

#### 1.9.41.1. Updating

To update an existing OpenShift Container Platform 4.11 cluster to this latest release, see [Updating a cluster using the CLI](#).

## 1.9.42. RHSA-2023:4310 OpenShift Container Platform 4.11.46 bug fix and security update

Issued: 2023-08-02

OpenShift Container Platform release 4.11.46, which includes security updates, is now available. Bug fixes included in the update are listed in the [RHSA-2023:4310](#) advisory. RPM packages included in the update are provided by the [RHSA-2023:4312](#) advisory.

You can view the container images in this release by running the following command:

```
$ oc adm release info 4.11.46 --pullspecs
```

### 1.9.42.1. Updating

To update an existing OpenShift Container Platform 4.11 cluster to this latest release, see [Updating a cluster using the CLI](#).

## 1.9.43. RHBA-2023:4614 OpenShift Container Platform 4.11.47 bug fix update

Issued: 2023-08-16

OpenShift Container Platform release 4.11.47, which includes security updates, is now available. Bug fixes included in the update are listed in the [RHBA-2023:4614](#) advisory. RPM packages included in the update are provided by the [RHBA-2023:4616](#) advisory.

You can view the container images in this release by running the following command:

```
$ oc adm release info 4.11.47 --pullspecs
```

### 1.9.43.1. Updating

To update an existing OpenShift Container Platform 4.11 cluster to this latest release, see [Updating a cluster using the CLI](#).

## 1.9.44. RHBA-2023:4752 OpenShift Container Platform 4.11.48 bug fix update

Issued: 2023-08-31

OpenShift Container Platform release 4.11.48 is now available. Bug fixes included in the update are listed in the [RHBA-2023:4752](#) advisory. There is no RPM package in this update.

You can view the container images in this release by running the following command:

```
$ oc adm release info 4.11.48 --pullspecs
```

### 1.9.44.1. Updating

To update an existing OpenShift Container Platform 4.11 cluster to this latest release, see [Updating a cluster using the CLI](#).

## 1.9.45. RHSA-2023:5001 OpenShift Container Platform 4.11.49 bug fix update

Issued: 2023-09-13

OpenShift Container Platform release 4.11.49, which includes security updates, is now available. Bug fixes included in the update are listed in the [RHSA-2023:5001](#) advisory. RPM packages included in the update are provided by the [RHBA-2023:5003](#) advisory.

You can view the container images in this release by running the following command:

```
$ oc adm release info 4.11.49 --pullspecs
```

### 1.9.45.1. Updating

To update an existing OpenShift Container Platform 4.11 cluster to this latest release, see [Updating a cluster using the CLI](#).