



OpenShift Container Platform 4.19

Release notes

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

OpenShift Container Platform 4.19 Release notes

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

Legal Notice

Copyright © Red Hat.

Except as otherwise noted below, the text of and illustrations in this documentation are licensed by Red Hat under the Creative Commons Attribution–Share Alike 3.0 Unported license . 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, the Red Hat logo, JBoss, Hibernate, and RHCE are trademarks or registered trademarks of Red Hat, LLC. or its subsidiaries in the United States and other countries.

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

XFS is a trademark or registered trademark of Hewlett Packard Enterprise Development LP or its subsidiaries in the United States and other countries.

The OpenStack[®] Word Mark and OpenStack logo are trademarks or registered trademarks of the Linux Foundation, used under license.

All other trademarks are the property of their respective owners.

Abstract

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

Table of Contents

CHAPTER 1. OPENSIFT CONTAINER PLATFORM 4.19 RELEASE NOTES	8
1.1. ABOUT THIS RELEASE	8
1.2. OPENSIFT CONTAINER PLATFORM LAYERED AND DEPENDENT COMPONENT SUPPORT AND COMPATIBILITY	8
1.3. NEW FEATURES AND ENHANCEMENTS	9
1.3.1. Authentication and authorization	9
1.3.1.1. Enabling direct authentication with an external OIDC identity provider (Technology Preview)	9
1.3.1.2. Enable ServiceAccountTokenNodeBinding Kubernetes feature by default	9
1.3.2. Documentation	9
1.3.2.1. Consolidated etcd documentation	9
1.3.2.2. Tutorials guide	9
1.3.3. Edge computing	9
1.3.3.1. Using RHACM PolicyGenerator resources to manage GitOps ZTP cluster policies (General Availability)	10
1.3.3.2. Configuring a local arbiter node (Technology Preview)	10
1.3.3.3. Coordinating reboots for configuration changes	10
1.3.4. Extensions (OLM v1)	10
1.3.4.1. Preflight permissions check for cluster extensions (Technology Preview)	10
1.3.4.2. Deploying a cluster extension in a specific namespace (Technology Preview)	10
1.3.5. Hardware accelerators	11
1.3.5.1. Dynamic Accelerator Slicer Operator (Technology Preview)	11
1.3.6. Hosted control planes	11
1.3.6.1. Hosted control planes on Red Hat OpenStack Platform (RHOSP) 17.1 (Technology Preview)	11
1.3.7. IBM Power	11
1.3.8. IBM Z and IBM LinuxONE	11
1.3.9. IBM Power, IBM Z, and IBM LinuxONE support matrix	11
1.3.10. Insights Operator	16
1.3.10.1. Insights Runtime Extractor is generally available	16
1.3.11. Installation and update	16
1.3.11.1. Cluster API replaces Terraform on IBM Cloud installations	16
1.3.11.2. Installing a cluster on Microsoft Azure with virtual network encryption	16
1.3.11.3. Installing a cluster on AWS in the Malaysia and Thailand regions	16
1.3.11.4. Cluster API replaces Terraform on Microsoft Azure Stack Hub installations	16
1.3.11.5. Support added for additional Microsoft Azure instance types	16
1.3.11.6. Outbound access for VMs in Microsoft Azure will be retired	17
1.3.11.7. Additional Confidential Computing platforms for Google Cloud	17
1.3.11.8. Installing a cluster on Google Cloud with a user-provisioned DNS (Technology Preview)	17
1.3.11.9. Installing a cluster on VMware vSphere with multiple disks (Technology Preview)	18
1.3.11.10. Enabling boot diagnostics collection during installation on Microsoft Azure	18
1.3.11.11. Required administrator acknowledgment when updating from OpenShift Container Platform 4.18 to 4.19	18
1.3.11.12. OpenShift zones support for vSphere host groups (Technology Preview)	18
1.3.11.13. Nutanix support for the Agent-based Installer	18
1.3.11.14. Support for VMware vSphere Foundation 9 and VMware Cloud Foundation 9	19
1.3.12. Machine Config Operator	19
1.3.12.1. New naming for features	19
1.3.12.2. Image mode for OpenShift is now generally available	19
1.3.12.3. Boot image management is now default for Google Cloud and Amazon Web Services (AWS)	20
1.3.12.4. Changes to the Machine Config Operator certificates	20
1.3.13. Machine management	21
1.3.13.1. Migrating resources between the Cluster API and the Machine API (Technology Preview)	21

1.3.13.2. Custom prefixes for control plane machine names	21
1.3.13.3. Configuring Capacity Reservations on Amazon Web Services clusters	21
1.3.13.4. Support for multiple VMware vSphere data disks (Technology Preview)	21
1.3.14. Monitoring	21
1.3.14.1. Updates to monitoring stack components and dependencies	21
1.3.14.2. Changes to alerting rules	22
1.3.14.3. Prometheus v3 upgrade	22
1.3.14.4. Metrics collection profiles is generally available	23
1.3.14.5. Added cluster proxy support for external Alertmanager instances	23
1.3.14.6. Strict validation for the Cluster Monitoring Operator is improved	23
1.3.15. Networking	23
1.3.15.1. Support for route advertisements for cluster user-defined networks (CUDNs) with Border Gateway Protocol (BGP)	23
1.3.15.2. Creating a route with externally managed certificate (General Availability)	23
1.3.15.3. Support for the BGP routing protocol	23
1.3.15.4. Support for using Gateway API to configure cluster ingress traffic (General Availability)	24
1.3.15.5. Support for managing Gateway API custom resource definition (CRD) lifecycle	25
1.3.15.6. Updates to Gateway API custom resource definitions (CRDs)	25
1.3.15.7. Allocate API and ingress load balancers to specific subnets	25
1.3.15.8. Dual-port NICs for improved redundancy in PTP ordinary clocks (Technology Preview)	25
1.3.15.9. Support for conditional webhook matching in the SR-IOV Network Operator	25
1.3.15.10. Enabling DPU device management with the DPU Operator	25
1.3.15.11. Localnet topology for user-defined networks (Generally Available)	26
1.3.15.12. Enable port isolation for a Linux bridge NAD (Generally Available)	26
1.3.15.13. Fast IPAM configuration for the Whereabouts IPAM CNI plugin (Technology Preview)	26
1.3.15.14. Unnumbered BGP peering (Technology Preview)	26
1.3.15.15. Create a custom DNS host name to resolve DNS connectivity issues	26
1.3.15.16. Removal of PTP events REST API v1 and events consumer application sidecar	26
1.3.15.17. Re-add a previously deleted secret with RouteExternalCertificate feature gate enabled	26
1.3.16. OpenShift CLI (oc)	27
1.3.16.1. Mirroring and verifying image signatures in oc-mirror plugin v2	27
1.3.17. Operator development	27
1.3.17.1. Supported Operator base images	27
1.3.18. Postinstallation configuration	27
1.3.18.1. Using bare metal as a service (Technology Preview)	27
1.3.18.2. Amazon Web Services Security Token Service (STS) can be enabled on existing clusters	27
1.3.19. Red Hat Enterprise Linux CoreOS (RHCOS)	27
1.3.19.1. RHCOS uses RHEL 9.6	27
1.3.20. Scalability and performance	28
1.3.20.1. Performance profile kernel page size configuration	28
1.3.20.2. Updates to the cluster-compare plugin	28
1.3.20.3. Tuning hosted control planes using a performance profile	28
1.3.21. Security	28
1.3.21.1. Control plane now supports TLS 1.3 and the Modern TLS security profile	28
1.3.21.2. The External Secrets Operator for Red Hat OpenShift (Technology Preview)	28
1.3.22. Storage	28
1.3.22.1. Support for the Secrets Store CSI driver in disconnected environments	28
1.3.22.2. Azure File cross-subscription support is generally available	29
1.3.22.3. Volume Attributes Classes (Technology Preview)	29
1.3.22.4. New CLI command to show PVC usage (Technology Preview)	29
1.3.22.5. CSI volume resizing recovery is generally available	29
1.3.22.6. Support for resizing vSphere in-tree migrated volumes is generally available	30
1.3.22.7. Disabling and enabling storage on vSphere is generally available	30

1.3.22.8. Increasing the maximum number of volumes per node for vSphere (Technology Preview)	30
1.3.22.9. Migrating CNS volumes between datastores for vSphere is fully supported	30
1.3.22.10. NFS export options for Filestore storage class is generally available.	30
1.3.23. Web console	31
1.3.23.1. Patternfly 6 upgrade	31
1.4. NOTABLE TECHNICAL CHANGES	31
1.4.1. Pods deploy with readOnlyRootFilesystem set to true	31
1.4.2. Extended loopback certificate validity to three years for kube-apiserver	31
1.4.3. Readiness probes exclude etcd checks	31
1.4.4. Installer automatically removes leftover Cloud Native Storage (CNS) volumes	32
1.4.5. Red Hat Enterprise Linux CoreOS (RHCOS) versioning uses Red Hat Enterprise Linux (RHEL) instead of OpenShift Container Platform	32
1.4.6. VMware vSphere 7 and VMware Cloud Foundation 4 end of general support	32
1.5. DEPRECATED AND REMOVED FEATURES	32
1.5.1. Bare metal monitoring deprecated and removed features	33
1.5.2. Images deprecated and removed features	33
1.5.3. Installation deprecated and removed features	33
1.5.4. Networking deprecated and removed features	33
1.5.5. Node deprecated and removed features	34
1.5.6. OpenShift CLI (oc) deprecated and removed features	34
1.5.7. Operator lifecycle and development deprecated and removed features	34
1.5.8. Storage deprecated and removed featuresmco-rn-tp-fix	35
1.5.9. Updating clusters deprecated and removed features	35
1.5.10. Web console deprecated and removed features	35
1.5.11. Workloads deprecated and removed features	35
1.5.12. Deprecated features	35
1.5.12.1. oc adm pod-network command deprecated	35
1.5.12.2. useModal hook for dynamic plugin SDK	36
1.5.12.3. Kubernetes API deprecation	36
1.5.13. Removed features	36
1.5.13.1. cgroup v1 has been removed	36
1.5.13.2. Package-based RHEL compute machines	36
1.5.13.3. APIs removed from Kubernetes 1.32	36
1.5.13.4. Operator SDK CLI and related scaffolding and testing tools	37
1.6. BUG FIXES	37
1.6.1. API Server and Authentication	37
1.6.2. Bare Metal Hardware Provisioning	38
1.6.3. Cloud Compute	39
1.6.4. Cloud Credential Operator	42
1.6.5. Cluster Autoscaler	42
1.6.6. Cluster Resource Override Admission Operator	42
1.6.7. Cluster Version Operator	43
1.6.8. etcd	43
1.6.9. ImageStreams	43
1.6.10. Installer	43
1.6.11. Machine Config Operator	47
1.6.12. Management Console	48
1.6.13. Monitoring	50
1.6.14. Networking	50
1.6.15. Node	53
1.6.16. Node Tuning Operator (NTO)	54
1.6.17. Observability	54
1.6.18. oc-mirror	55

1.6.19. OpenShift CLI (oc)	56
1.6.20. Operator Lifecycle Manager (OLM)	56
1.6.21. Operator Controller Manager	57
1.6.22. Performance Addon Operator	57
1.6.23. Samples Operator	58
1.6.24. Storage	58
1.6.25. Red Hat Enterprise Linux CoreOS (RHCOS)	59
1.7. TECHNOLOGY PREVIEW FEATURES STATUS	59
1.7.1. Authentication and authorization Technology Preview features	60
1.7.2. Edge computing Technology Preview features	60
1.7.3. Extensions Technology Preview features	60
1.7.4. Installation Technology Preview features	61
1.7.5. Machine Config Operator Technology Preview features	62
1.7.6. Machine management Technology Preview features	62
1.7.7. Monitoring Technology Preview features	63
1.7.8. Multi-Architecture Technology Preview features	63
1.7.9. Networking Technology Preview features	64
1.7.10. Node Technology Preview features	65
1.7.11. OpenShift CLI (oc) Technology Preview features	65
1.7.12. Operator lifecycle and development Technology Preview features	66
1.7.13. Red Hat OpenStack Platform (RHOSP) Technology Preview features	66
1.7.14. Scalability and performance Technology Preview features	66
1.7.15. Storage Technology Preview features	67
1.7.16. Web console Technology Preview features	68
1.8. KNOWN ISSUES	68
1.9. ASYNCHRONOUS ERRATA UPDATES	71
1.9.1. RHSA-2026:4434 - OpenShift Container Platform 4.19.26 fixed issues and security update	72
1.9.1.1. Fixed issues	72
1.9.1.2. Updating	74
1.9.2. RHBA-2026:3394 - OpenShift Container Platform 4.19.25 fixed issues and security update	74
1.9.2.1. Fixed issues	74
1.9.2.2. Updating	74
1.9.3. RHSA-2026:2651 - OpenShift Container Platform 4.19.24 fixed issues and security update	75
1.9.3.1. Fixed issues	75
1.9.3.2. Updating	76
1.9.4. RHSA-2026:1552 - OpenShift Container Platform 4.19.23 bug fix advisory	76
1.9.5. Enhancements	76
1.9.6. Fixed issues	77
1.9.7. Updating	77
1.9.8. RHBA-2026:0682 - OpenShift Container Platform 4.19.22 bug fix advisory	77
1.9.9. Fixed issues	78
1.9.10. Updating	80
1.9.11. RHBA-2025:22786 - OpenShift Container Platform 4.19.21 bug fix advisory	80
1.9.12. Bug fixes	80
1.9.13. Updating	80
1.9.14. RHBA-2025:22278 - OpenShift Container Platform 4.19.20 bug fix advisory	80
1.9.15. Bug fixes	81
1.9.16. Updating	81
1.9.17. RHBA-2025:21363 - OpenShift Container Platform 4.19.19 bug fix advisory	82
1.9.18. Bug fixes	82
1.9.19. Updating	83
1.9.20. RHBA-2025:19301 - OpenShift Container Platform 4.19.18 image release and bug fix advisory	83
1.9.20.1. Bug fixes	83

1.9.20.2. Updating	85
1.9.21. RHSA-2025:18233 - OpenShift Container Platform 4.19.17 image release and bug fix advisory	85
1.9.21.1. Bug fixes	85
1.9.21.2. Updating	86
1.9.22. RHBA-2025:17663 - OpenShift Container Platform 4.19.16 image release and bug fix advisory	86
1.9.22.1. Bug fixes	86
1.9.22.2. Updating	87
1.9.23. RHBA-2025:17237 - OpenShift Container Platform 4.19.15 image release and bug fix advisory	87
1.9.23.1. Bug fixes	87
1.9.23.2. Updating	87
1.9.24. RHBA-2025:16693 - OpenShift Container Platform 4.19.14 image release and bug fix advisory	87
1.9.24.1. Bug fixes	88
1.9.24.2. Updating	88
1.9.25. RHBA-2025:16148 - OpenShift Container Platform 4.19.13 image release, bug fix, and security update advisory	88
1.9.25.1. Bug fixes	89
1.9.25.2. Updating	89
1.9.26. RHBA-2025:15694 - OpenShift Container Platform 4.19.12 image release, bug fix, and security update advisory	89
1.9.26.1. Enhancements	89
1.9.26.2. Bug fixes	90
1.9.26.3. Updating	90
1.9.27. RHBA-2025:15293 - OpenShift Container Platform 4.19.11 image release, bug fix, and security update advisory	91
1.9.27.1. Enhancements	91
1.9.27.2. Bug fixes	91
1.9.27.3. Updating	92
1.9.28. RHSA-2025:14823 - OpenShift Container Platform 4.19.10 image release, bug fix, and security update advisory	92
1.9.28.1. Enhancements	92
1.9.28.2. Bug fixes	93
1.9.28.3. Updating	94
1.9.29. RHSA-2025:13848 - OpenShift Container Platform 4.19.9 image release, bug fix, and security update advisory	94
1.9.29.1. Enhancements	95
1.9.29.2. Bug fixes	95
1.9.29.3. Updating	96
1.9.30. RHSA-2025:12341 - OpenShift Container Platform 4.19.7 image release, bug fix, and security update advisory	96
1.9.30.1. Enhancements	96
1.9.30.2. Bug fixes	97
1.9.30.3. Updating	98
1.9.31. RHSA-2025:11673 - OpenShift Container Platform 4.19.6 image release, bug fix, and security update advisory	98
1.9.31.1. Enhancements	98
1.9.31.2. Bug fixes	98
1.9.31.3. Updating	99
1.9.32. RHSA-2025:11363 - OpenShift Container Platform 4.19.5 image release, bug fix, and security update advisory	99
1.9.32.1. Bug fixes	99
1.9.32.2. Updating	100
1.9.33. RHSA-2025:10771 - OpenShift Container Platform 4.19.4 image release, bug fix, and security update advisory	100
1.9.33.1. Bug fixes	100

1.9.33.2. Updating	101
1.9.34. RHBA-2025:10290 - OpenShift Container Platform 4.19.3 image release, bug fix, and security update advisory	101
1.9.34.1. Bug fixes	102
1.9.34.2. Updating	102
1.9.35. RHSA-2025:9750 - OpenShift Container Platform 4.19.2 image release, bug fix, and security update advisory	102
1.9.35.1. Bug fixes	103
1.9.35.2. Updating	104
1.9.36. RHSA-2025:9278 - OpenShift Container Platform 4.19.1 image release, bug fix, and security update advisory	104
1.9.36.1. Bug fixes	104
1.9.36.2. Updating	106
1.9.37. RHSA-2024:11038 - OpenShift Container Platform 4.19.0 image release, bug fix, and security update advisory	106
1.9.37.1. Updating	106
CHAPTER 2. ADDITIONAL RELEASE NOTES	107

CHAPTER 1. OPENSIFT CONTAINER PLATFORM 4.19 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. 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-2024:11038](#)) is now available. This release uses [Kubernetes 1.32](#) with CRI-O runtime. New features, changes, and known issues that pertain to OpenShift Container Platform 4.19 are included in this topic.

OpenShift Container Platform 4.19 clusters are available at <https://console.redhat.com/openshift>. From the Red Hat Hybrid Cloud Console, you can deploy OpenShift Container Platform clusters to either on-premises or cloud environments.

You must use RHCOS machines for the control plane and for the compute machines.

The support lifecycle for odd-numbered releases, such as OpenShift Container Platform 4.19, on all supported architectures, including **x86_64**, 64-bit ARM (**aarch64**), IBM Power® (**ppc64le**), and IBM Z® (**s390x**) architectures is 18 months. For more information about support for all versions, see the [Red Hat OpenShift Container Platform Life Cycle Policy](#).

Commencing with the OpenShift Container Platform 4.14 release, Red Hat is simplifying the administration and management of Red Hat shipped cluster Operators with the introduction of three new life cycle classifications; Platform Aligned, Platform Agnostic, and Rolling Stream. These life cycle classifications provide additional ease and transparency for cluster administrators to understand the life cycle policies of each Operator and form cluster maintenance and upgrade plans with predictable support boundaries. For more information, see [OpenShift Operator Life Cycles](#).

OpenShift Container Platform is designed for FIPS. When running Red Hat Enterprise Linux (RHEL) or Red Hat Enterprise Linux CoreOS (RHCOS) booted in FIPS mode, OpenShift Container Platform core components use the RHEL cryptographic libraries that have been submitted to NIST for FIPS 140-2/140-3 Validation on only the **x86_64**, **ppc64le**, and **s390x** architectures.

For more information about the NIST validation program, see [Cryptographic Module Validation Program](#). For the latest NIST status for the individual versions of RHEL cryptographic libraries that have been submitted for validation, see [Compliance Activities and Government Standards](#).

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

The following new features are supported on IBM Power with OpenShift Container Platform 4.19:

- Support for IBM Power®11

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

1.3.1. Authentication and authorization

1.3.1.1. Enabling direct authentication with an external OIDC identity provider (Technology Preview)

With this release, you can enable direct integration with an external OpenID Connect (OIDC) identity provider to issue tokens for authentication. This bypasses the built-in OAuth server and uses the external identity provider directly.

By integrating directly with an external OIDC provider, you can leverage the advanced capabilities of your preferred OIDC provider instead of being limited by the capabilities of the built-in OAuth server. Your organization can manage users and groups from a single interface, while also streamlining authentication across multiple clusters and in hybrid environments. You can also integrate with existing tools and solutions.

Direct authentication is available as a Technology Preview feature.

For more information, see [Enabling direct authentication with an external OIDC identity provider](#) .

1.3.1.2. Enable `ServiceAccountTokenNodeBinding` Kubernetes feature by default

In OpenShift Container Platform 4.19, the **`ServiceAccountTokenNodeBinding`** feature is now enabled by default, aligning with upstream Kubernetes behavior. This feature allows service account tokens to be bound directly to node objects in addition to the existing binding options. Benefits of this change include enhanced security through automatic token invalidation when bound nodes are deleted and better protection against token replay attacks across different nodes.

1.3.2. Documentation

1.3.2.1. Consolidated `etcd` documentation

This release includes an *etcd* section, which consolidates all of the existing documentation about `etcd` for OpenShift Container Platform. For more information, see [Overview of `etcd`](#).

1.3.2.2. Tutorials guide

OpenShift Container Platform 4.19 now includes a *Tutorials* guide, which takes the place of the *Getting started* guide in previous releases. The existing tutorials were refreshed and the guide now focuses solely on hands-on tutorial content. It also provides a jumping off point to other recommended hands-on learning resources for OpenShift Container Platform across Red Hat.

For more information, see [Tutorials](#).

1.3.3. Edge computing

1.3.3.1. Using RHACM PolicyGenerator resources to manage GitOps ZTP cluster policies (General Availability)

You can now use **PolicyGenerator** resources and Red Hat Advanced Cluster Management (RHACM) to deploy policies for managed clusters with GitOps ZTP. The **PolicyGenerator** API is part of the [Open Cluster Management](#) standard and provides a generic way of patching resources, which is not possible with the **PolicyGenTemplate** API. Using **PolicyGenTemplate** resources to manage and deploy policies will be deprecated in an upcoming OpenShift Container Platform release.

For more information, see [Configuring managed cluster policies by using PolicyGenerator resources](#) .

1.3.3.2. Configuring a local arbiter node (Technology Preview)

You can configure an OpenShift Container Platform cluster with two control plane nodes and one local arbiter node so to retain high availability (HA) while reducing infrastructure costs for your cluster. This configuration is only supported for bare-metal installations.

A local arbiter node is a lower-cost, co-located machine that participates in control plane quorum decisions. Unlike a standard control plane node, the arbiter node does not run the full set of control plane services. You can use this configuration to maintain HA in your cluster with only two fully provisioned control plane nodes instead of three.

To enable this feature, you must define the arbiter machine pool in the **install-config.yaml** file and enable the **TechPreviewNoUpgrade** feature set.

Configuring a local arbiter node is available as a Technology Preview feature. For more information, see [Configuring a local arbiter node](#) .

1.3.3.3. Coordinating reboots for configuration changes

This release adds reboot policies to ZTP reference that can be applied by Topology Aware Lifecycle Manager (TALM) to coordinate reboots across a fleet of spoke clusters when configuration changes require a reboot, such as deferred tuning changes. TALM reboots all nodes in the targeted **MachineConfigPool** object on the selected clusters when the reboot policy is applied.

Instead of rebooting nodes after each individual change, you can apply all configuration updates through policies and then trigger a single, coordinated reboot.

For more information, see [Coordinating reboots for configuration changes](#).

1.3.4. Extensions (OLM v1)

1.3.4.1. Preflight permissions check for cluster extensions (Technology Preview)

With this release, the Operator Controller performs a dry run of the installation process when you try to install an extension. This dry run verifies that the specified service account has the required role-based access control (RBAC) rules for the roles and bindings defined by the bundle.

If the service account is missing any required RBAC rules, the preflight check fails before the actual installation proceeds and generates a report.

For more information, see [Preflight permissions check for cluster extensions \(Technology Preview\)](#)

1.3.4.2. Deploying a cluster extension in a specific namespace (Technology Preview)

With this release, you can deploy an extension in a specific namespace by using the **OwnNamespace** or **SingleNamespace** install modes as a Technology Preview feature for **registry+v1** Operator bundles.

For more information, see [Deploying a cluster extension in a specific namespace \(Technology Preview\)](#)

1.3.5. Hardware accelerators

1.3.5.1. Dynamic Accelerator Slicer Operator (Technology Preview)

With this release, you can use the Dynamic Accelerator Slicer (DAS) Operator to dynamically slice GPU accelerators in OpenShift Container Platform, instead of relying on statically sliced GPUs defined when the node is booted. This allows you to dynamically slice GPUs based on specific workload demands, ensuring efficient resource utilization.

For more information, see [Dynamic Accelerator Slicer \(DAS\) Operator](#).

1.3.6. Hosted control planes

Because hosted control planes releases asynchronously from OpenShift Container Platform, it has its own release notes. For more information, see [Hosted control planes release notes](#).

1.3.6.1. Hosted control planes on Red Hat OpenStack Platform (RHOSP) 17.1 (Technology Preview)

Hosted control planes on RHOSP 17.1 are now supported as a Technology Preview.

For more information, see [Deploying hosted control planes on OpenStack](#).

1.3.7. IBM Power

The IBM Power® release on OpenShift Container Platform 4.19 adds improvements and new capabilities to OpenShift Container Platform components.

This release introduces support for the following features on IBM Power:

- Expand Compliance Operator support with profiles for Defense Information Systems Agency Security Technical Implementation Guide (DISA STIG)

1.3.8. IBM Z and IBM LinuxONE

The IBM Z® and IBM® LinuxONE release on OpenShift Container Platform 4.19 adds improvements and new capabilities to OpenShift Container Platform components.

This release introduces support for the following features on IBM Z® and IBM® LinuxONE:

- Support for IBM® z17 and IBM® LinuxONE 5
- Boot volume Linux Unified Key Setup (LUKS) encryption via IBM® Crypto Express (CEX)

1.3.9. IBM Power, IBM Z, and IBM LinuxONE support matrix

Starting in OpenShift Container Platform 4.14, Extended Update Support (EUS) is extended to the IBM Power® and the IBM Z® platform. For more information, see the [OpenShift EUS Overview](#).

Table 1.1. CSI Volumes

Feature	IBM Power®	IBM Z® and IBM® LinuxONE
Cloning	Supported	Supported
Expansion	Supported	Supported
Snapshot	Supported	Supported

Table 1.2. Multus CNI plugins

Feature	IBM Power®	IBM Z® and IBM® LinuxONE
Bridge	Supported	Supported
Host-device	Supported	Supported
IPAM	Supported	Supported
IPVLAN	Supported	Supported

Table 1.3. OpenShift Container Platform features

Feature	IBM Power®	IBM Z® and IBM® LinuxONE
Adding compute nodes to on-premise clusters using OpenShift CLI (oc)	Supported	Supported
Alternate authentication providers	Supported	Supported
Agent-based Installer	Supported	Supported
Assisted Installer	Supported	Supported
Automatic Device Discovery with Local Storage Operator	Unsupported	Supported
Automatic repair of damaged machines with machine health checking	Unsupported	Unsupported
Cloud controller manager for IBM Cloud®	Supported	Unsupported
Controlling overcommit and managing container density on nodes	Unsupported	Unsupported

Feature	IBM Power®	IBM Z® and IBM® LinuxONE
CPU manager	Supported	Supported
Cron jobs	Supported	Supported
Descheduler	Supported	Supported
Egress IP	Supported	Supported
Encrypting data stored in etcd	Supported	Supported
FIPS cryptography	Supported	Supported
Helm	Supported	Supported
Horizontal pod autoscaling	Supported	Supported
Hosted control planes	Supported	Supported
IBM Secure Execution	Unsupported	Supported
Installer-provisioned Infrastructure Enablement for IBM Power® Virtual Server	Supported	Unsupported
Installing on a single node	Supported	Supported
IPv6	Supported	Supported
Monitoring for user-defined projects	Supported	Supported
Multi-architecture compute nodes	Supported	Supported
Multi-architecture control plane	Supported	Supported
Multipathing	Supported	Supported
Network-Bound Disk Encryption - External Tang Server	Supported	Supported
Non-volatile memory express drives (NVMe)	Supported	Unsupported
nx-gzip for Power10 (Hardware Acceleration)	Supported	Unsupported
oc-mirror plugin	Supported	Supported
OpenShift CLI (oc) plugins	Supported	Supported

Feature	IBM Power®	IBM Z® and IBM® LinuxONE
Operator API	Supported	Supported
OpenShift Virtualization	Unsupported	Supported
OVN-Kubernetes, including IPsec encryption	Supported	Supported
PodDisruptionBudget	Supported	Supported
Precision Time Protocol (PTP) hardware	Unsupported	Unsupported
Red Hat OpenShift Local	Unsupported	Unsupported
Scheduler profiles	Supported	Supported
Secure Boot	Unsupported	Supported
Stream Control Transmission Protocol (SCTP)	Supported	Supported
Support for multiple network interfaces	Supported	Supported
The openshift-install utility to support various SMT levels on IBM Power® (Hardware Acceleration)	Supported	Unsupported
Three-node cluster support	Supported	Supported
Topology Manager	Supported	Unsupported
z/VM Emulated FBA devices on SCSI disks	Unsupported	Supported
4K FCP block device	Supported	Supported

Table 1.4. Operators

Feature	IBM Power®	IBM Z® and IBM® LinuxONE
cert-manager Operator for Red Hat OpenShift	Supported	Supported
Cluster Logging Operator	Supported	Supported
Cluster Resource Override Operator	Supported	Supported
Compliance Operator	Supported	Supported

Feature	IBM Power®	IBM Z® and IBM® LinuxONE
Cost Management Metrics Operator	Supported	Supported
File Integrity Operator	Supported	Supported
HyperShift Operator	Supported	Supported
IBM Power® Virtual Server Block CSI Driver Operator	Supported	Unsupported
Ingress Node Firewall Operator	Supported	Supported
Local Storage Operator	Supported	Supported
MetalLB Operator	Supported	Supported
Network Observability Operator	Supported	Supported
NFD Operator	Supported	Supported
NMState Operator	Supported	Supported
OpenShift Elasticsearch Operator	Supported	Supported
Vertical Pod Autoscaler Operator	Supported	Supported

Table 1.5. Persistent storage options

Feature	IBM Power®	IBM Z® and IBM® LinuxONE
Persistent storage using iSCSI	Supported ^[1]	Supported ^{[1],[2]}
Persistent storage using local volumes (LSO)	Supported ^[1]	Supported ^{[1],[2]}
Persistent storage using hostPath	Supported ^[1]	Supported ^{[1],[2]}
Persistent storage using Fibre Channel	Supported ^[1]	Supported ^{[1],[2]}
Persistent storage using Raw Block	Supported ^[1]	Supported ^{[1],[2]}
Persistent storage using EDEV/FBA	Supported ^[1]	Supported ^{[1],[2]}

1. Persistent shared storage must be provisioned by using either Red Hat OpenShift Data Foundation or other supported storage protocols.
2. Persistent non-shared storage must be provisioned by using local storage, such as iSCSI, FC, or by using LSO with DASD, FCP, or EDEV/FBA.

1.3.10. Insights Operator

1.3.10.1. Insights Runtime Extractor is generally available

In OpenShift Container Platform 4.18, the Insights Operator introduced the *Insights Runtime Extractor* workload data collection feature as a Technology Preview feature to help Red Hat better understand the workload of your containers. Now, in version 4.19, the feature is generally available. The Insights Runtime Extractor feature gathers runtime workload data and sends it to Red Hat.

1.3.11. Installation and update

1.3.11.1. Cluster API replaces Terraform on IBM Cloud installations

In OpenShift Container Platform 4.19, the installation program uses the Cluster API instead of Terraform to provision cluster infrastructure during installations on IBM Cloud.

1.3.11.2. Installing a cluster on Microsoft Azure with virtual network encryption

With this release, you can install a cluster on Azure using encrypted virtual networks. You are required to use Azure virtual machines that have the **premiumIO** parameter set to **true**. See Microsoft's documentation about [Creating a virtual network with encryption](#) and [Requirements and Limitations](#) for more information.

1.3.11.3. Installing a cluster on AWS in the Malaysia and Thailand regions

You can now install an OpenShift Container Platform cluster on Amazon Web Services (AWS) in the Malaysia (**ap-southeast-5**) and Thailand (**ap-southeast-7**) regions.

For more information, see [Supported Amazon Web Services \(AWS\) regions](#).

1.3.11.4. Cluster API replaces Terraform on Microsoft Azure Stack Hub installations

In OpenShift Container Platform 4.19, the installation program uses the Cluster API instead of Terraform to provision clusters during installer-provisioned infrastructure installations on Microsoft Azure Stack Hub.

1.3.11.5. Support added for additional Microsoft Azure instance types

Additional Microsoft Azure instance types for machine types based on 64-bit x86 architecture have been tested with OpenShift Container Platform 4.19.

For the Dxcv6 machine series, the following instance types have been tested:

- **StandardDdsv6Family**
- **StandardDldsv6Family**
- **StandardDlsv6Family**

- **StandardDsv6Family**

For the Lsv4 and Lasv4 machine series, the following instance types have been tested:

- **standardLasv4Family**
- **standardLsv4Family**

For the ND and NV machine series, the following instance types have been tested:

- **StandardNVadsV710v5Family**
- **Standard NDASv4_A100 Family**

For more information, see [Tested instance types for Azure](#) and [Azure documentation](#) (Microsoft documentation).

1.3.11.6. Outbound access for VMs in Microsoft Azure will be retired

On 30 September 2025, the default outbound access connectivity for all new virtual machines (VMs) in Microsoft Azure will be retired. To enhance security, Azure is moving towards a secure-by-default model where default outbound access to the internet will be turned off. However, configuration changes to OpenShift Container Platform are not required. By default, the installation program creates an outbound rule for the load balancer.

For more information, see [Azure Updates](#) (Microsoft documentation), [Azure's outbound connectivity methods](#) (Microsoft documentation), and [Preparing to install a cluster on Azure](#).

1.3.11.7. Additional Confidential Computing platforms for Google Cloud

With this release, you can use additional Confidential Computing platforms on Google Cloud. The new supported platforms, which can be enabled in the **install-config.yaml** file prior to installation, or configured after installation by using machine sets and control plane machine sets, are as follows:

- **AMDEncryptedVirtualization**, which enables Confidential Computing with AMD Secure Encrypted Virtualization (AMD SEV)
- **AMDEncryptedVirtualizationNestedPaging**, which enables Confidential Computing with AMD Secure Encrypted Virtualization Secure Nested Paging (AMD SEV-SNP)
- **IntelTrustedDomainExtensions**, which enables Confidential Computing with Intel Trusted Domain Extensions (Intel TDX)

For more information, see [Installation configuration parameters for Google Cloud](#), [Configuring Confidential VM by using machine sets \(control plane\)](#), and [Configuring Confidential VM by using machine sets \(compute\)](#).

1.3.11.8. Installing a cluster on Google Cloud with a user-provisioned DNS (Technology Preview)

With this release, you can enable a user-provisioned domain name server (DNS) instead of the default cluster-provisioned DNS solution. For example, your organization's security policies might not allow the use of public DNS services such as Google Cloud DNS. You can manage your DNS only for the IP

addresses of the API and Ingress servers. If you use this feature, you must provide your own DNS solution that includes records for **api.<cluster_name>.<base_domain>**, and ***.apps.<cluster_name>.<base_domain>**. Enabling a user-provisioned DNS is available as a Technology Preview feature.

For more information, see [Enabling user-managed DNS](#).

1.3.11.9. Installing a cluster on VMware vSphere with multiple disks (Technology Preview)

With this release, you can install a cluster on VMware vSphere with multiple storage disks as a Technology Preview feature. You can assign these additional disks to special functions within the cluster, such as etcd storage.

For more information, see [Optional vSphere configuration parameters](#).

1.3.11.10. Enabling boot diagnostics collection during installation on Microsoft Azure

With this release, you can enable boot diagnostics collection when you install a cluster on Microsoft Azure. Boot diagnostics is a debugging feature for Azure virtual machines (VMs) to identify VM boot failures. You can set the **bootDiagnostics** parameter in the **install-config.yaml** file for compute machines, for control plane machines, or for all machines.

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

1.3.11.11. Required administrator acknowledgment when updating from OpenShift Container Platform 4.18 to 4.19

OpenShift Container Platform 4.19 uses Kubernetes 1.32, which removed several [deprecated APIs](#).

A cluster administrator must provide manual acknowledgment before the cluster can be updated from OpenShift Container Platform 4.18 to 4.19. This is to help prevent issues after updating to OpenShift Container Platform 4.19, where APIs that have been removed are still in use by workloads, tools, or other components running on or interacting with the cluster. Administrators must evaluate their cluster for any APIs in use that will be removed and migrate the affected components to use the appropriate new API version. After this is done, the administrator can provide the administrator acknowledgment.

All OpenShift Container Platform 4.18 clusters require this administrator acknowledgment before they can be updated to OpenShift Container Platform 4.19.

For more information, see [Preparing to update to OpenShift Container Platform 4.19](#).

1.3.11.12. OpenShift zones support for vSphere host groups (Technology Preview)

With this release, you can map OpenShift Container Platform failure domains to VMware vSphere host groups. This enables you to make use of the high availability offered by a vSphere stretched cluster configuration. This feature is available as a Technology Preview in OpenShift Container Platform 4.19.

For information on configuring host groups at installation, see [VMware vSphere host group enablement](#).

For information on configuring host groups for existing clusters, see [Specifying multiple host groups for your cluster on vSphere](#).


1.3.11.13. Nutanix support for the Agent-based Installer

With this release, you can now use the Agent-based Installer to install a cluster on Nutanix. Installing a cluster on Nutanix with the Agent-based Installer is enabled by setting the **platform** parameter to **nutanix** in the **install-config.yaml** file.

For more information, see [Required configuration parameters](#) in the Agent-based Installer documentation.

1.3.11.14. Support for VMware vSphere Foundation 9 and VMware Cloud Foundation 9

You can now install OpenShift Container Platform on VMware vSphere Foundation (VVF) 9 and VMware Cloud Foundation (VCF) 9.



NOTE

The following additional VCF and VVF components are outside the scope of Red Hat support:

- Management: VCF Operations, VCF Automation, VCF Fleet Management, and VCF Identity Broker.
- Networking: VMware NSX Container Plugin (NCP).
- Migration: VMware HCX.

1.3.12. Machine Config Operator

1.3.12.1. New naming for features

Red Hat Enterprise Linux CoreOS (RHCOS) image layering is now called *image mode for OpenShift*. As a part of this change, *on-cluster layering* is now called *on-cluster image mode* and *out-of-cluster layering* is now *out-of-cluster image mode*.

The *updated boot images* feature is now called *boot image management*.

1.3.12.2. Image mode for OpenShift is now generally available

Image mode for OpenShift, formerly called on-cluster layering, is now Generally Available (GA). The following changes have been introduced with the promotion to GA:

- The API version is now **machineconfiguration.openshift.io/v1**. The new version includes the following changes:
 - The **baseImagePullSecret** parameter is now optional. If not specified, the default **global-pull-secret-copy** is used.
 - The **buildInputs** parameter is no longer required. All parameters previously under the **buildInputs** parameter are promoted one level.
 - The **containerfileArch** parameter now supports multiple architectures. Previously, only **noarch** was supported.
 - The required **imageBuilderType** is now **Job**. Previously, the required builder was **PodImageBuilder**.
 - The **renderedImagePushspec** parameter is now **renderedImagePushSpec**.

- The **buildOutputs** and **currentImagePullSecret** parameters are no longer required.
- The output of the **oc describe MachineOSConfig** and **oc describe MachineOSBuild** commands have multiple differences.
- The **global-pull-secret-copy** is automatically added to the **openshift-machine-config-operator** namespace.
- You can now revert an on-cluster custom layered image back to the base image by removing a label from the **MachineOSConfig** object
- You can now automatically delete an on-cluster custom layered image by deleting the associated **MachineOSBuild** object.
- The **must-gather** for the Machine Config Operator now includes data on the **MachineOSConfig** and **MachineOSBuild** objects.
- On-cluster layering is now supported in disconnected environments.
- On-cluster layering is now supported in single node OpenShift (SNO) clusters.

1.3.12.3. Boot image management is now default for Google Cloud and Amazon Web Services (AWS)

The boot image management feature, previously called updated boot images, is now the default behavior in Google Cloud and Amazon Web Services (AWS) clusters. As such, after updating to OpenShift Container Platform 4.19, the boot images in your cluster are automatically updated to version 4.19. With subsequent updates, the Machine Config Operator (MCO) again updates the boot images in your cluster. A boot images is associated with a machine set and is used when scaling new nodes. Any new nodes you create after updating are based on the new version. Current nodes are not affected by this feature.

Before upgrading to 4.19, you must opt-out of this default behavior or acknowledge this change before proceeding. For more information, see [Disabling boot image management](#).



NOTE

The managed boot images feature is available for only Google Cloud and AWS clusters. For all other platforms, the MCO does not update the boot image with each cluster update.

1.3.12.4. Changes to the Machine Config Operator certificates

The Machine Config Server (MCS) CA bundle created by the installation program is now stored in the **machine-config-server-ca** config map in the **openshift-machine-config-operator** namespace. The bundle was previously stored in the **root-ca** configmap in the **kube-system namespace**. The **root-ca** configmap is no longer used in a cluster that is updated to OpenShift Container Platform 4.19. This change was made to make it clear that this CA bundle is managed by the Machine Config Operator (MCO).

The MCS signing key is stored in the **machine-config-server-ca** secret in the **openshift-machine-config-operator** namespace.

The MCS CA and MCS cert are valid for 10 years and are automatically rotated by the MCO at approximately 8 years. Upon update to OpenShift Container Platform 4.19, the CA signing key is not

present. As a result, the CA bundle is immediately considered expired when the MCO certificate controller comes up. This expiration causes an immediate certificate rotation, even if the cluster is not 10 years old. After that point, the next rotation takes place at the standard 8 year period.



NOTE

This automatic certificate rotation applies only to clusters that use machine sets. For clusters that do not use machine sets, such as vSphere user-provisioned infrastructure clusters, you are required to manually rotate these certificates. For more information on manual certificate rotation, see the Red Hat Knowledgebase article [Regenerating CA certificates for the Machine Config Server](#).

For more information about the MCO certificates, see [Machine Config Operator certificates](#).

1.3.13. Machine management

1.3.13.1. Migrating resources between the Cluster API and the Machine API (Technology Preview)

With this release, you can migrate some resources between the Cluster API and the Machine API on Amazon Web Services (AWS) as a Technology Preview feature. For more information, see [Migrating Machine API resources to Cluster API resources](#).

To support this capability, the OpenShift Container Platform Cluster API documentation now includes additional configuration details for [AWS](#) clusters.

1.3.13.2. Custom prefixes for control plane machine names

With this release, you can customize the prefix of machine names for machines created by the control plane machine set. This feature is enabled by modifying the `spec.machineNamePrefix` parameter of the `ControlPlaneMachineSet` custom resource.

For more information, see [Adding a custom prefix to control plane machine names](#).

1.3.13.3. Configuring Capacity Reservations on Amazon Web Services clusters

With this release, you can deploy machines that use Capacity Reservations, including On-Demand Capacity Reservations and Capacity Blocks for ML, on Amazon Web Services clusters.

You can configure these features with [compute](#) and [control plane](#) machine sets.

1.3.13.4. Support for multiple VMware vSphere data disks (Technology Preview)

With this release, you can add up to 29 disks to the virtual machine (VM) controller for your vSphere cluster as a Technology Preview feature. This capability is available for [compute](#) and [control plane](#) machine sets.

1.3.14. Monitoring


The in-cluster monitoring stack for this release includes the following new and modified features:

1.3.14.1. Updates to monitoring stack components and dependencies

This release includes the following version updates for in-cluster monitoring stack components and dependencies:

- Alertmanager to 0.28.1
- Prometheus to 3.2.1
- Prometheus Operator to 0.81.0
- Thanos to 0.37.2
- kube-state-metrics to 2.15.0
- node-exporter to 1.9.1

1.3.14.2. Changes to alerting rules



NOTE

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

- Added the **PrometheusPossibleNarrowSelectors** alert to warn users when PromQL queries or metric relabel configurations use selectors that could be too restrictive and might not take into account that values on the **le** label of classic histograms or the **quantile** label of summaries are floats in Prometheus v3. For more information, see the "Prometheus v3 upgrade" section.

1.3.14.3. Prometheus v3 upgrade

This release introduces a major update to the Prometheus component, transitioning from v2 to v3. The monitoring stack and other core components include all of the necessary adjustments to ensure a smooth upgrade. However, some user-managed configurations might require modifications. The key changes include the following items:

- The values of the **le** label for classic histograms and the **quantile** label for summaries are normalized during ingestion. For example, the **example_bucket{le="10"}** metric selector is ingested as **example_bucket{le="10.0"}**. As a result, alerts, recording rules, dashboards, and relabeling configurations that reference label values as whole numbers, for example, **le="10"**, might no longer work as intended.
To mitigate the issue, update your selectors:
 - If your queries need to cover data from both before and after the Prometheus upgrade, ensure both values are considered, for example, use a regular expression, **example_bucket{le=~"10(.0)?"}**.
 - For queries that only cover data after the upgrade, use float values, for example, **le="10.0"**.
- Configurations that send alerts to additional Alertmanager instances through **additionalAlertmanagerConfigs** by using the Alertmanager v1 API are no longer supported. To mitigate the issue, upgrade any affected Alertmanager instances to support the v2 API, which is supported since Alertmanager **v0.16.0**, and update your monitoring configuration to use the v2 scheme.

For more information about the changes between Prometheus v2 and v3, see [Prometheus 3.0 migration guide](#).

1.3.14.4. Metrics collection profiles is generally available

OpenShift Container Platform 4.13 introduced the ability to set a metrics collection profile for default platform monitoring to collect either the default amount of metrics data or a minimal amount of metrics data. In OpenShift Container Platform 4.19, metrics collection profiles are now generally available.

For more information, see [About metrics collection profiles](#) and [Choosing a metrics collection profile](#).

1.3.14.5. Added cluster proxy support for external Alertmanager instances

With this release, external Alertmanager instances now use the cluster-wide HTTP proxy settings for communication. The Cluster Monitoring Operator (CMO) reads the cluster-wide proxy settings and configures the appropriate proxy URL for the Alertmanager endpoints.

1.3.14.6. Strict validation for the Cluster Monitoring Operator is improved

With this release, the strict validation introduced in OpenShift Container Platform 4.18 is improved. Error messages now clearly identify the affected field, and validation is case-sensitive to ensure more accurate and consistent configuration.

For more information, see [\(OCPBUGS-42671\)](#) and [\(OCPBUGS-54516\)](#).

1.3.15. Networking

1.3.15.1. Support for route advertisements for cluster user-defined networks (CUDNs) with Border Gateway Protocol (BGP)

With route advertisements enabled, the OVN-Kubernetes network plugin supports the direct advertisement of routes for pods and services associated with cluster user-defined networks (CUDNs) to the provider network. This feature enables some of the following benefits:

- Learns routes to pods dynamically
- Advertises routes dynamically
- Enables layer 3 notifications of EgressIP failovers in addition to the layer 2 ones based on gratuitous ARPs.
- Supports external route reflectors, which reduces the number of BGP connections required in large networks

For more information, see [About route advertisements](#).

1.3.15.2. Creating a route with externally managed certificate (General Availability)

With this release, OpenShift Container Platform routes can be configured with third-party certificate management solutions, utilizing the `.spec.tls.externalCertificate` field in the route API. This allows you to reference externally managed TLS certificates through secrets, streamlining the process by eliminating manual certificate management. By using externally managed certificates, you reduce errors, ensure a smoother certificate update process, and enable the OpenShift router to promptly serve renewed certificates. For more information, see [Creating a route with externally managed certificate](#).

1.3.15.3. Support for the BGP routing protocol

The Cluster Network Operator (CNO) now supports enabling Border Gateway Protocol (BGP) routing.

With BGP, you can import and export routes to the underlying provider network and use multi-homing, link redundancy, and fast convergence. BGP configuration is managed with the **FRRConfiguration** custom resource (CR).

When upgrading from an earlier version of OpenShift Container Platform in which you installed the MetalLB Operator, you must manually migrate your custom frr-k8s configurations from the **metallb-system** namespace to the **openshift-frr-k8s** namespace. To move these CRs, enter the following commands:

1. To create the **openshift-frr-k8s** namespace, enter the following command:

```
$ oc create namespace openshift-frr-k8s
```

2. To automate the migration, create a **migrate.sh** file with the following content:

```
#!/bin/bash
OLD_NAMESPACE="metallb-system"
NEW_NAMESPACE="openshift-frr-k8s"
FILTER_OUT="metallb-"
oc get frrconfigurations.frrk8s.metallb.io -n "${OLD_NAMESPACE}" -o json | \
jq -r '.items[] | select(.metadata.name | test("'"${FILTER_OUT}"'" | not)' | \
jq -r '.metadata.namespace = "'"${NEW_NAMESPACE}"'" | \
oc create -f -
```

3. To run the migration script, enter the following command:

```
$ bash migrate.sh
```

4. To verify that the migration succeeded, enter the following command:

```
$ oc get frrconfigurations.frrk8s.metallb.io -n openshift-frr-k8s
```

After the migration is complete, you can remove the **FRR-K8s** custom resources from the **metallb-system** namespace.

For more information, see [About BGP routing](#).

1.3.15.4. Support for using Gateway API to configure cluster ingress traffic (General Availability)

With this release, support for managing ingress cluster traffic using Gateway API resources is Generally Available. Gateway API provides a robust networking solution within the transport layer, L4, and the application layer, L7, for OpenShift Container Platform clusters using a standardized open source ecosystem.

For more information, see [Gateway API with OpenShift Container Platform networking](#).



IMPORTANT

Gateway API resources must conform to the supported OpenShift Container Platform API surface. This means you cannot use another vendor-specific resource, such as Istio's VirtualService, with OpenShift Container Platform's implementation of Gateway API. For more information, see [Gateway API implementation for OpenShift Container Platform](#).

1.3.15.5. Support for managing Gateway API custom resource definition (CRD) lifecycle

With this release, OpenShift Container Platform manages the lifecycle of Gateway API CRDs. This means that the Ingress Operator handles the required versioning and management of resources. Any Gateway API resources created in a previous OpenShift Container Platform version must be re-created and redeployed so that it conforms to the specifications required by the Ingress Operator.

For more information, see [Preparing for Gateway API management succession by the Ingress Operator](#).

1.3.15.6. Updates to Gateway API custom resource definitions (CRDs)

OpenShift Container Platform 4.19 updates Red Hat OpenShift Service Mesh to version 3.0.2, and Gateway API to version 1.2.1. See the [Service Mesh 3.0.0 release notes](#) and the [Gateway API 1.2.1 changelog](#) for more information.

1.3.15.7. Allocate API and ingress load balancers to specific subnets

With this release, you can now allocate load balancers to customize deployments when installing an OpenShift Container Platform cluster on AWS. This feature ensures optimal traffic distribution, high application availability, uninterrupted service, and network segmentation.

For more information, see [Installation configuration parameters on AWS](#) and [Allocating load balancers to specific subnets](#).

1.3.15.8. Dual-port NICs for improved redundancy in PTP ordinary clocks (Technology Preview)

With this release, you can use a dual-port network interface controller (NIC) to improve redundancy for Precision Time Protocol (PTP) ordinary clocks. Available as a Technology Preview, in a dual-port NIC configuration for an ordinary clock, if one port fails, the standby port takes over, maintaining PTP timing synchronization.



NOTE

You can configure PTP ordinary clocks with added redundancy on **x86** architecture nodes with dual-port NICs only.

For more information, see [Using dual-port NICs to improve redundancy for PTP ordinary clocks](#).

1.3.15.9. Support for conditional webhook matching in the SR-IOV Network Operator

You can now enable the **featureGates.resourceInjectorMatchCondition** feature in the **SriovOperatorConfig** object to limit the scope of the Network Resources Injector webhook. If this feature is enabled, the webhook applies only to pods with the secondary network annotation **k8s.v1.cni.cncf.io/networks**.

If this feature is disabled, the webhook's **failurePolicy** is set to **Ignore** by default. This configuration can cause pods requesting SR-IOV networks to be deployed without the required resource injection if the webhook is unavailable. If this feature is enabled and the webhook is unavailable, pods without the annotation are still deployed, preventing unnecessary disruptions to other workloads.

For more information, see [About the Network Resources Injector](#).

1.3.15.10. Enabling DPU device management with the DPU Operator

With this release, OpenShift Container Platform introduces the Data Processing Unit (DPU) Operator and using the Operator to manage DPU devices. The DPU Operator manages components on compute nodes that have configured DPUs, such as enabling the offloading of data networking, storage, and security workloads. Enabling DPU device management leads to improved cluster performance, reduced latency, and enhanced security, that overall contribute to a more efficient cluster infrastructure. For more information, see [About DPU and the DPU Operator](#).

1.3.15.11. Localnet topology for user-defined networks (Generally Available)

Administrators can now use the **ClusterUserDefinedNetwork** custom resource to deploy secondary networks on a **Localnet** topology. This feature allows pods and virtual machines connected to the localnet network to egress to the physical network. For more information, see [Creating a ClusterUserDefinedNetwork CR for a Localnet topology](#).

1.3.15.12. Enable port isolation for a Linux bridge NAD (Generally Available)

You can enable port isolation for a Linux bridge network attachment definition (NAD) so that virtual machines (VMs) or pods that run on the same virtual LAN (VLAN) can operate in isolation from one another. For more information, see [Enabling port isolation for a Linux bridge NAD](#).

1.3.15.13. Fast IPAM configuration for the Whereabouts IPAM CNI plugin (Technology Preview)

To improve the performance of Whereabouts, especially if nodes in your cluster run a high amount of pods, you can now enable the Fast IP Address Management (IPAM) feature. The Fast IPAM feature uses **nodeslice pools**, which are managed by the Whereabouts Controller, to optimize IP address allocation for nodes. For more information, see [Fast IPAM configuration for the Whereabouts IPAM CNI plugin](#).

1.3.15.14. Unnumbered BGP peering (Technology Preview)

With this release, OpenShift Container Platform introduces unnumbered BGP peering. Available as a Technology Preview feature, you can use the **spec.interface** field of the BGP peer custom resource to configure unnumbered BGP peering.

1.3.15.15. Create a custom DNS host name to resolve DNS connectivity issues

In a disconnected environment where the external DNS server cannot be reached, you can resolve Kubernetes NMState Operator health probe issues by specifying a custom DNS host name in the **NMState** custom resource definition (CRD). For more information, see [Creating a custom DNS host name to resolve DNS connectivity issues](#).

1.3.15.16. Removal of PTP events REST API v1 and events consumer application sidecar

With this release, the PTP events REST API v1 and events consumer application sidecar support are removed.

You must use the O-RAN compliant PTP events REST API v2 instead.

For more information, see [Developing PTP event consumer applications with the REST API v2](#).

1.3.15.17. Re-add a previously deleted secret with RouteExternalCertificate feature gate enabled

If you enabled the **RouteExternalCertificate** feature gate for your cluster, you can now re-add a previously deleted secret. ([OCPBUGS-33958](#))

1.3.16. OpenShift CLI (oc)

1.3.16.1. Mirroring and verifying image signatures in oc-mirror plugin v2

Starting with OpenShift Container Platform 4.19, the oc-mirror plugin v2 supports mirroring and verifying cosign tag-based signatures for container images.

1.3.17. Operator development

1.3.17.1. Supported Operator base images

The following base images for Operator projects are updated for compatibility with OpenShift Container Platform 4.19. The runtime functionality and configuration APIs for these base images are supported for bug fixes and for addressing CVEs.

- The base image for Ansible-based Operator projects
- The base image for Helm-based Operator projects

For more information, see [Updating the base image for existing Ansible- or Helm-based Operator projects for OpenShift Container Platform 4.19 and later](#) (Red Hat Knowledgebase).

1.3.18. Postinstallation configuration

1.3.18.1. Using bare metal as a service (Technology Preview)

In OpenShift Container Platform 4.19, you can deploy non-OpenShift Container Platform nodes by using bare metal as a service (BMaaS). BMaaS nodes can run workloads that might not be suitable for containerization or virtualization. For example, workloads such as applications that require direct hardware access, conduct high-performance computing tasks or are legacy applications and operate independently of the cluster are suitable for deployment by using BMaaS.

For more information, see [Using bare metal as a service](#) .

1.3.18.2. Amazon Web Services Security Token Service (STS) can be enabled on existing clusters

With this release, you can configure your AWS OpenShift Container Platform cluster to use STS even if you did not do so during installation.

For more information, see [Enabling AWS Security Token Service \(STS\) on an existing cluster](#) .

1.3.19. Red Hat Enterprise Linux CoreOS (RHCOS)

1.3.19.1. RHCOS uses RHEL 9.6

RHCOS uses Red Hat Enterprise Linux (RHEL) 9.6 packages in OpenShift Container Platform 4.19. These packages ensure that your OpenShift Container Platform instances receive the latest fixes, features, enhancements, hardware support, and driver updates.

1.3.20. Scalability and performance

1.3.20.1. Performance profile kernel page size configuration

With this update, you can specify larger kernel page sizes to improve performance for memory-intensive, high-performance workloads on ARM infrastructure nodes with the realtime kernel disabled. For more information, see [Configuring kernel page sizes](#).

1.3.20.2. Updates to the cluster-compare plugin

This release includes the following usability and functional updates to the **cluster-compare** plugin:

- Match capture groups more effectively: You can now match across and between templates more accurately with improved capture group handling.
- Generate JUnit output: You can use the **-o junit** flag to output results in **junit** format, making it easier to integrate with testing or CI/CD systems.
- **sprig** function support: The **cluster-compare** plugin supports all **sprig** library functions, except for the **env** and **expandenv** functions. For the full list of sprig library functions, see [Sprig Function Documentation](#).

For a complete list of available template functions, see [Reference template functions](#)

1.3.20.3. Tuning hosted control planes using a performance profile

With this update, you can now tune nodes in hosted control planes for low latency by applying a performance profile. For more information, see [Creating a performance profile for hosted control planes](#).

1.3.21. Security

1.3.21.1. Control plane now supports TLS 1.3 and the Modern TLS security profile

With this release, the control plane supports TLS 1.3. You can now use the **Modern** TLS security profile for the control plane.

For more information, see [Configuring the TLS security profile for the control plane](#).

1.3.21.2. The External Secrets Operator for Red Hat OpenShift (Technology Preview)

With this release, you can use the External Secrets Operator for Red Hat OpenShift to authenticate with the external secrets store, retrieve secrets, and inject the retrieved secrets into a native Kubernetes secret. The External Secrets Operator for Red Hat OpenShift is available as a Technology Preview.

For more information, see [External Secrets Operator for Red Hat OpenShift overview](#)

1.3.22. Storage

1.3.22.1. Support for the Secrets Store CSI driver in disconnected environments

With this release, the secrets store providers support by using the Secrets Store CSI driver in disconnected clusters.

For more information, see [Support for disconnected environments](#).

1.3.22.2. Azure File cross-subscription support is generally available

Cross-subscription support allows you to have an OpenShift Container Platform cluster in one Azure subscription and mount your Azure file share in another Azure subscription using the Azure File Container Storage Interface (CSI) driver. The subscriptions must be in the same tenant.

This feature is generally available in OpenShift Container Platform 4.19.

For more information, see [AWS EFS CSI cross account support](#).

1.3.22.3. Volume Attributes Classes (Technology Preview)

Volume Attributes Classes provide a way for administrators to describe "classes" of storage they offer. Different classes might correspond to different quality-of-service levels.

Volume Attributes Classes in OpenShift Container Platform 4.19 is available only with AWS Elastic Block Storage (EBS) and Google Cloud Platform (GCP) persistent disk (PD) Container Storage Interface (CSI).

You can apply a Volume Attributes Classes to a persistent volume claim (PVC). If a new Volume Attributes Class becomes available in the cluster, you can update the PVC with the new Volume Attributes Classes if needed.

Volume Attributes Classes have parameters that describe volumes belonging to them. If a parameter is omitted, the default is used at volume provisioning. If a user applies the PVC with a different Volume Attributes Class with omitted parameters, the default value of the parameters might be used depending on the CSI driver implementation. For more information, see the related CSI driver documentation.

Volume Attributes Classes is available in OpenShift Container Platform 4.19 with Technology Preview status.

For more information, see [Volume Attributes Classes](#).

1.3.22.4. New CLI command to show PVC usage (Technology Preview)

OpenShift Container Platform 4.19 introduces a new command to view persistent volume claim usage. This feature has Technology Preview status.

For more information, see [Viewing PVC usage statistics](#).

1.3.22.5. CSI volume resizing recovery is generally available

Previously, you might expand a persistent volume claim (PVC) to a size that is not supported by the underlying storage provider. In this case, the expansion controller typically tries forever to expand the volume and keeps failing.

This new feature allows you to recover and provide another resize value for the PVC. Resizing recovery is supported as generally available in OpenShift Container Platform 4.19.

For more information about resizing volumes, see [Expanding persistent volumes](#).

For more information about recovering when resizing volumes, see [Recovering from failure when expanding volumes](#).

1.3.22.6. Support for resizing vSphere in-tree migrated volumes is generally available

Previously, VMware vSphere persistent volumes that were migrated from in-tree to Container Storage Interface (CSI) could not be resized. With OpenShift Container Platform 4.19, resizing migrated volumes is supported. This feature is generally available.

For more information about resizing volumes, see [Expanding persistent volumes](#).

1.3.22.7. Disabling and enabling storage on vSphere is generally available

Cluster administrators might want to disable the VMware vSphere Container Storage Interface (CSI) Driver as a Day 2 operation, so the vSphere CSI Driver does not interface with your vSphere setup.

This feature was introduced in OpenShift Container Platform 4.17 with Technology Preview status. This feature is now supported as generally available in OpenShift Container Platform 4.19.

For more information, see [Disabling and enabling storage on vSphere](#).

1.3.22.8. Increasing the maximum number of volumes per node for vSphere (Technology Preview)

For VMware vSphere version 7, OpenShift Container Platform restricts the maximum number of volumes per node to 59.

However, with OpenShift Container Platform 4.19 for vSphere version 8 or later, you can increase the allowable number of volumes per node to a maximum of 255. Otherwise, the default value remains at 59.

This feature has Technology Preview status.

For more information, see [Increasing maximum volumes per node for vSphere](#).

1.3.22.9. Migrating CNS volumes between datastores for vSphere is fully supported

If you are running out of space in your current datastore, or want to move to a more performant datastore, you can migrate VMware vSphere Cloud Native Storage (CNS) volumes between datastores. This applies to both attached and detached volumes.

OpenShift Container Platform now fully supports migration of CNS volume using the vCenter UI. Migrated volumes should work as expected and should not result in non-functional persistent volumes. CNS volumes can also be migrated while in use by pods.

This feature was introduced as a Development Preview in OpenShift Container Platform 4.17, but is now fully supported in 4.19.

Migrating CNS volumes between datastores requires VMware vSphere 8.0.2 or later or vSphere 7.0 Update 30 or later.

For more information, see [Migrating CNS volumes between datastores for vSphere](#).

1.3.22.10. NFS export options for Filestore storage class is generally available.

By default, a Filestore instance grants root level read/write access to all clients that share the same Google Cloud project and virtual private cloud (VPC) network. Network File System (NFS) export options can limit this access to certain IP ranges and specific user/group IDs for the Filestore instance. When creating a storage class, you can set these options using the **nfs-export-options-on-create** parameter.

NFS export options is supported as generally available in OpenShift Container Platform 4.19.

For more information, see [NFS export options](#).

1.3.23. Web console

Starting with OpenShift Container Platform 4.19, the perspectives in the web console have unified to simplify navigation, reduce context switching, streamline tasks, and provide users with a more cohesive OpenShift Container Platform experience.

With this unified design, there is no longer a **Developer** perspective in the default view; however, *all* OpenShift Container Platform web console features are discoverable to all users. If you are not the cluster owner, you might need to request permission for certain features from the cluster owner. The **Developer** perspective can still be manually enabled if you prefer.

The **Getting Started** pane in the web console provides resources such as, a tour of the console, information on setting up your cluster, a quick start for enabling the **Developer** perspective, and links to explore new features and capabilities.

1.3.23.1. Patternfly 6 upgrade

The web console now uses Patternfly 6. Support for Patternfly 4 in the web console is no longer available.

This release also introduces the following updates to the web console. You can now do the following actions:

- Specify distinct console logos for both light and dark themes using the **logos** field in the **.spec.customization.logos** configuration, allowing for more comprehensive branding.
- Easily delete identity providers (IDPs) directly from the web console, streamlining authentication configuration without manual YAML file edits.
- Easily set the default **StorageClass** directly in the web console.
- Quickly find specific jobs in the web console by sorting the **Created** column by creation date and time.

1.4. NOTABLE TECHNICAL CHANGES

1.4.1. Pods deploy with `readOnlyRootFilesystem` set to `true`

With this release, Cloud Credential Operator pods now deploy with the **readOnlyRootFilesystem** security context setting set to **true**. This enhances security by ensuring that the container root file system is mounted as read-only.

1.4.2. Extended loopback certificate validity to three years for kube-apiserver

Previously, the self-signed loopback certificate for the Kubernetes API Server expired after one year. With this release, the expiration date of the certificate is extended to three years.

1.4.3. Readiness probes exclude etcd checks

The readiness probes for the API server have been modified to exclude etcd checks. This prevents client connections from being closed if etcd is temporarily unavailable. This means that client connections persist through brief etcd unavailability and minimizes temporary API server outages.

1.4.4. Installer automatically removes leftover Cloud Native Storage (CNS) volumes

The OpenShift installation program now automatically detects and removes leftover persistent storage volumes on VMware vSphere when you delete a cluster. This prevents orphaned volumes from consuming disk space and creating unnecessary alerts in vCenter.

1.4.5. Red Hat Enterprise Linux CoreOS (RHCOS) versioning uses Red Hat Enterprise Linux (RHEL) instead of OpenShift Container Platform

As part of aligning with Image Mode for RHEL, RHCOS is now built as a layer on top of a shared RHEL base image. The most noticeable change for users is around versioning. For example, **VERSION_ID** in **/etc/os-release** now reflects the version of RHEL, such as RHEL 9.6, rather than the version of OpenShift Container Platform, such as OpenShift Container Platform 4.19. This version change might show up in other places, such as in the output of the command **rpm-ostree status**, or in boot loader entries. **OPENSHIFT_VERSION** in **/etc/os-release** on the node image still uses the version of OpenShift Container Platform and is unaffected by this change.

Because RHCOS now uses the RHEL base image, any new operating system features added to that base image will usually be inherited by all OpenShift Container Platform releases that use this base image.

1.4.6. VMware vSphere 7 and VMware Cloud Foundation 4 end of general support

Broadcom has ended general support for VMware vSphere 7 and VMware Cloud Foundation (VCF) 4. If your existing OpenShift Container Platform cluster is running on either of these platforms, you must plan to migrate or upgrade your VMware infrastructure to a supported version. OpenShift Container Platform supports installation on vSphere 8 Update 1 or later, or VCF 5 or later.

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 4.19, refer to the table below. Additional details for more functionality that has been deprecated and removed are listed after the table.

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

- *Not Available*
- *Technology Preview*
- *General Availability*
- *Deprecated*
- *Removed*

1.5.1. Bare metal monitoring deprecated and removed features

Table 1.6. Bare Metal Event Relay Operator tracker

Feature	4.17	4.18	4.19
Bare Metal Event Relay Operator	Removed	Removed	Removed

1.5.2. Images deprecated and removed features

Table 1.7. Images deprecated and removed tracker

Feature	4.17	4.18	4.19
Cluster Samples Operator	Deprecated	Deprecated	Deprecated

1.5.3. Installation deprecated and removed features

Table 1.8. Installation deprecated and removed tracker

Feature	4.17	4.18	4.19
--cloud parameter for oc adm release extract	Deprecated	Deprecated	Deprecated
CoreDNS wildcard queries for the cluster.local domain	Deprecated	Deprecated	Deprecated
compute.platform.openstack.rootVolume.type for RHOSP	Deprecated	Deprecated	Deprecated
controlPlane.platform.openstack.rootVolume.type for RHOSP	Deprecated	Deprecated	Deprecated
ingressVIP and apiVIP settings in the install-config.yaml file for installer-provisioned infrastructure clusters	Deprecated	Deprecated	Deprecated
Package-based RHEL compute machines	Deprecated	Deprecated	Removed
platform.aws.preserveBootstrapIgnition parameter for Amazon Web Services (AWS)	Deprecated	Deprecated	Deprecated
Installing a cluster on AWS with compute nodes in AWS Outposts	Deprecated	Deprecated	Deprecated

1.5.4. Networking deprecated and removed features

Table 1.9. Networking deprecated and removed tracker

Feature	4.17	4.18	4.19
iptables	Deprecated	Deprecated	Deprecated

1.5.5. Node deprecated and removed features

Table 1.10. Node deprecated and removed tracker

Feature	4.17	4.18	4.19
ImageContentSourcePolicy (ICSP) objects	Deprecated	Deprecated	Deprecated
Kubernetes topology label failure-domain.beta.kubernetes.io/zone	Deprecated	Deprecated	Deprecated
Kubernetes topology label failure-domain.beta.kubernetes.io/region	Deprecated	Deprecated	Deprecated
cgroup v1	Deprecated	Deprecated	Removed

1.5.6. OpenShift CLI (oc) deprecated and removed features

Table 1.11. OpenShift CLI (oc) deprecated and removed tracker

Feature	4.17	4.18	4.19
oc-mirror plugin v1	General Availability	Deprecated	Deprecated

1.5.7. Operator lifecycle and development deprecated and removed features

Table 1.12. Operator lifecycle and development deprecated and removed tracker

Feature	4.17	4.18	4.19
Operator SDK	Deprecated	Deprecated	Removed
Scaffolding tools for Ansible-based Operator projects	Deprecated	Deprecated	Removed
Scaffolding tools for Helm-based Operator projects	Deprecated	Deprecated	Removed
Scaffolding tools for Go-based Operator projects	Deprecated	Deprecated	Removed
Scaffolding tools for Hybrid Helm-based Operator projects	Deprecated	Removed	Removed
Scaffolding tools for Java-based Operator projects	Deprecated	Removed	Removed

Feature	4.17	4.18	4.19
SQLite database format for Operator catalogs	Deprecated	Deprecated	Deprecated

1.5.8. Storage deprecated and removed featuresmco-rn-tp-fix

Table 1.13. Storage deprecated and removed tracker

Feature	4.17	4.18	4.19
Persistent storage using FlexVolume	Deprecated	Deprecated	Deprecated
AliCloud Disk CSI Driver Operator	Removed	Removed	Removed
Shared Resources CSI Driver Operator	Deprecated	Removed	Removed

1.5.9. Updating clusters deprecated and removed features

Table 1.14. Updating clusters deprecated and removed tracker

Feature	4.17	4.18	4.19
---------	------	------	------

1.5.10. Web console deprecated and removed features

Table 1.15. Web console deprecated and removed tracker

Feature	4.17	4.18	4.19
useModal hook for dynamic plugin SDK	General Availability	General Availability	Deprecated
Patternfly 4	Deprecated	Deprecated	Removed

1.5.11. Workloads deprecated and removed features

Table 1.16. Workloads deprecated and removed tracker

Feature	4.17	4.18	4.19
DeploymentConfig objects	Deprecated	Deprecated	Deprecated

1.5.12. Deprecated features

1.5.12.1. oc adm pod-network command deprecated

The **oc adm pod-network** command for working with OpenShift SDN multitenant mode has been removed from the **oc adm --help** output. If the **oc adm pod-network** command is used, an error message is displayed to tell users that it has been deprecated.

1.5.12.2. useModal hook for dynamic plugin SDK

With this release, support for the **useModal** hook in dynamic plugins are deprecated.

Starting with this release, use the **useOverlay** API hook to launch modals

1.5.12.3. Kubernetes API deprecation

OpenShift Container Platform 4.17 inadvertently reintroduced a removed Kubernetes API, **admissionregistration.k8s.io/v1beta1**. This API is deprecated and is planned for removal in a future OpenShift Container Platform release. Migrate any instances of this API to **admissionregistration.k8s.io/v1**.

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

1.5.13. Removed features

1.5.13.1. cgroup v1 has been removed

cgroup v1, which was deprecated in OpenShift Container Platform 4.16, is no longer supported and has been removed from OpenShift Container Platform. If your cluster is using cgroup v1, you must configure cgroup v2 before you can upgrade to OpenShift Container Platform 4.19. All workloads must now be compatible with cgroup v2.

For information on configuring cgroup v2 in your cluster, see [Configuring Linux cgroup](#) in the OpenShift Container Platform version 4.18 documentation.

For more information on cgroup v2, see [About Linux cgroup version 2](#) and [Red Hat Enterprise Linux 9 changes in the context of Red Hat OpenShift workloads](#) (Red Hat blog).

1.5.13.2. Package-based RHEL compute machines

With this release, support for the installation of packaged-based RHEL worker nodes is removed.

RHCOS image layering replaces this feature and supports installing additional packages on the base operating system of your worker nodes.

For information on how to identify and remove RHEL nodes in your cluster, see [Preparing to update from OpenShift Container Platform 4.18 to a newer version](#). For more information on image layering, see [RHCOS image layering](#).

1.5.13.3. APIs removed from Kubernetes 1.32

Kubernetes 1.32 removed the following deprecated APIs, so you must migrate manifests and API clients to use the appropriate API version. For more information about migrating removed APIs, see the [Kubernetes documentation](#).

Table 1.17. APIs removed from Kubernetes 1.32

Resource	Removed API	Migrate to	Notable changes
FlowSchema	flowcontrol.apiserver.k8s.io/v1beta3	flowcontrol.apiserver.k8s.io/v1	No
PriorityLevelConfiguration	flowcontrol.apiserver.k8s.io/v1beta3	flowcontrol.apiserver.k8s.io/v1	Yes

1.5.13.4. Operator SDK CLI and related scaffolding and testing tools

With this release, the Red Hat-supported version of the Operator SDK CLI tool, including the related scaffolding and testing tools for Operator projects, is no longer released with OpenShift Container Platform.

Red Hat will provide bug fixes and support for versions of the Operator SDK that were released with earlier versions of OpenShift Container Platform according to the [Product Life Cycles for OpenShift Container Platform 4](#) (Red Hat Customer Portal).

Operator authors with existing Operator projects can use [the version of the Operator SDK CLI tool released with OpenShift Container Platform 4.18](#) to maintain their projects and create Operator releases that target newer versions of OpenShift Container Platform. For more information, see [Updating the base image for existing Ansible- or Helm-based Operator projects for OpenShift Container Platform 4.19 and later](#) (Red Hat Knowledgebase).

For more information about the unsupported, community-maintained, version of the Operator SDK, see [Operator SDK \(Operator Framework\)](#).

1.6. BUG FIXES

1.6.1. API Server and Authentication

- Previously, contents of the **MachineConfig** and **ControllerConfig** resources from group **machineconfiguration.openshift.io** were not excluded from audit logs. With this release, they are excluded from audit logs because they might contain secrets. ([OCPBUGS-55709](#))
- Previously, the kube-apiserver service level objective (SLO) alert expression incorrectly summed read and write success ratios independently of total request volume. This led to misleading burn rate calculations during disruptions. With this release, the fix adjusts the calculation to properly weight success ratios by total request count. This results in accurate and reliable alerting based on the true proportion of successful requests. ([OCPBUGS-49764](#))
- Previously, cluster bootstrap removal could break kube-apiserver readiness if etcd access was lost which could lead to downtime. With this release, each kube-apiserver has 2 stable etcd endpoints before removing bootstrap which maintains availability during rollout. ([OCPBUGS-48673](#))
- Previously, the Static Pod Operator API allowed invalid node statuses with unset **currentRevision** and multiple nonzero **targetRevision** entries, which led to failures in node and installer controllers. With this release, new validation rules were added to enforce correct revision fields to ensure stable and consistent static pod status handling. ([OCPBUGS-46380](#))
- Previously, the node controller applied stale **NodeStatus** data from its lister, unintentionally

overwriting recent updates from other controllers. With this release, the fix uses managed fields to let controllers update separate entries without conflict which preserves accurate and concurrent node status updates. ([OCPBUGS-46372](#))

- Previously, a fixed five minute timeout for removing the etcd bootstrap member started too early. This led to premature failures in HA clusters despite sufficient overall time. With this release, the narrow timeout is removed to rely on overall bootstrap progress instead, which ensures reliable and quorum-safe etcd bootstrap removal. ([OCPBUGS-46363](#))
- Previously, bootstrapping would unblock after detecting two kube-apiserver endpoints, including the bootstrap instance, causing periods of 0% availability as rollouts occur with only one permanent instance. With this release, the teardown is delayed until multiple permanent instances are ready. This ensures continuous kube-apiserver availability during rollout. ([OCPBUGS-46010](#))
- Previously, when the temporary control plane was down, the **networkConfig.status.ServiceNetwork** was not populated, and when generated certificates did not have the Kubernetes service IP in the SANs, the clients would fail to connect to the kube-apiserver through the default kubernetes service. With this release, a guard has been added to skip certificated generation if **networkConfig.status.ServiceNetwork** is nil. Client connections will be stable and valid. ([OCPBUGS-45943](#))
- Previously, the installer deleted the bootstrap machine before the etcd member was removed. This led to quorum loss in HA clusters. With this release, the check from SNO is extended to all topologies, using the etcd operator's condition as a safe removal sign, which ensures etcd cluster stability during bootstrap teardown. ([OCPBUGS-45482](#))
- Previously, the openshift-apiserver could panic when both image and error fields were unset during CRD request handling, this led to runtime crashes and instability in the API server under certain conditions. With this release, a guard is added to ensure no panic occurs by safely handling the case when both fields are unset, resulting in a more robust and stable CRD request handling process without crashes. ([OCPBUGS-45861](#))

1.6.2. Bare Metal Hardware Provisioning

- Previously, NetworkManager logs from the Ironic Python Agent (IPA) were not included in the ramdisk logs; instead only **dmesg** logs were included in the ramdisk logs. With this release, the ramdisk logs that exist in the **metal3-ramdisk-logs** container of a metal3 pod now contain the entire journal from the host instead of just **dmesg** logs and IPA. ([OCPBUGS-56042](#))
- Previously, ramdisk logs did not include clear file separators, causing the content from one file to merge into random lines of another file. Because of this issue, distinguishing what content belonged to which file was difficult. With this release, file entries now include file separators so that each file is clearly indicated from the contents of the other file being merged into a ramdisk log file. ([OCPBUGS-55743](#))
- Previously, if you forgot to include a Redfish system ID, such as **redfish://host/redfish/v1/** instead of **redfish://host/redfish/v1/Self**, in a Baseboard Management Console (BMC) URL, a JSON parsing issue existed in Ironic. With this release, BMO can now handle URLs without a Redfish system ID as a valid address without causing a JSON parsing issue. ([OCPBUGS-56026](#))
- Previously, a race condition existed during provisioning which, in case of a slow DHCP response, could cause different hostnames to be used for machine and node objects. This could prevent CSRs of worker nodes from being automatically approved. With this release, the race condition was fixed and CSRs of worker nodes are now properly approved. ([OCPBUGS-55315](#))

- Previously, certain models of SuperMicro machines, such as **ars-111gl-nhr**, use a different virtual media device string than other SuperMicro machines, which could cause virtual media boot attempts to fail on these servers. With this release, an extra conditional check was added to check for the specific model affected and adjust behavior accordingly, so that SuperMicro models such as **ars-111gl-nhr** can now boot from virtual media. ([OCPBUGS-56639](#))
- Previously, after deleting a **BaremetalHost** that has a related **Datalmage**, the **Datalmage** was still present. With this release, the **Datalmage** is deleted if it exists after the **BaremetalHost** has been deleted. ([OCPBUGS-51294](#))

1.6.3. Cloud Compute

- When upgrading Google Cloud clusters that use a boot disk that is not compatible with UEFI, you cannot enable Shielded VM support. Previously, this prevented the creation of new compute machines. With this release, disks with known UEFI incompatibility have Shielded VM support disabled. This primarily affects customers upgrading from OpenShift Container Platform version 4.12 to 4.13 using the Google Cloud marketplace images. ([OCPBUGS-17079](#))
- Previously, VMs in a cluster that ran on Azure failed because the attached network interface controller (NIC) was in a **ProvisioningFailed** state. With this release, the Machine API controller checks the provisioning status of a NIC and refreshes the VMs on a regular basis to prevent this issue. ([OCPBUGS-31515](#))
- Previously, in larger clusters that had other subsystems using certificate signing requests (CSRs), the CSR approver counted unrelated, unapproved CSRs towards its total and prevented further approvals. With this release, the CSR approver uses a **signerName** property as a filter and only includes CSRs that it can approve. As a result, the CSR approver only prevents new approvals when there are a large number of unapproved CSRs for the relevant **signerName** values. ([OCPBUGS-36404](#))
- Previously, the Machine API controller read only the zone number to populate machine zone information. For machines in Azure regions that only support availability sets, the set number represents the zone, so the Machine API controller did not populate their zone information. With this release, the Machine API controller references the Azure fault domain property. This property works for availability sets and availability zones, so the controller correctly reads the fault domain in each case and machines always report a zone. ([OCPBUGS-38570](#))
- Previously, increased granularity in Google Cloud zone API error messages caused the machine controller to mistakenly mark some machines with invalid configurations as valid with a temporary cloud error. This behavior prevented invalid machines from transitioning to a failed state. With this release, the machine controller handles the more granular error messages correctly so that machines with an invalid zone or project ID correctly move to a failed state. ([OCPBUGS-43531](#))
- Previously, some permissions required for linked actions were missing. Linked actions create the subresources necessary for other Azure resources that the cloud controller manager and OpenShift Container Platform require. With this release, the cloud controller manager for Azure has the following permissions for linked actions:
 - **Microsoft.Network/applicationGateways/backendAddressPools/join/action**
 - **Microsoft.Network/applicationSecurityGroups/joinIpConfiguration/action**
 - **Microsoft.Network/applicationSecurityGroups/joinNetworkSecurityRule/action**
 - **Microsoft.Network/ddosProtectionPlans/join/action**

- **Microsoft.Network/gatewayLoadBalancerAliases/join/action**
- **Microsoft.Network/loadBalancers/backendAddressPools/join/action**
- **Microsoft.Network/loadBalancers/frontendIPConfigurations/join/action**
- **Microsoft.Network/loadBalancers/inboundNatRules/join/action**
- **Microsoft.Network/networkInterfaces/join/action**
- **Microsoft.Network/networkSecurityGroups/join/action**
- **Microsoft.Network/publicIPAddresses/join/action**
- **Microsoft.Network/publicIPPrefixes/join/action**
- **Microsoft.Network/virtualNetworks/subnets/join/action**

([OCPBUGS-44126](#))

- Previously, some permissions required for linked actions were missing. Linked actions create the subresources necessary for other Azure resources that the Machine API and OpenShift Container Platform require. With this release, the Machine API provider for Azure has the following permissions for linked actions:
 - **Microsoft.Compute/disks/beginGetAccess/action**
 - **Microsoft.KeyVault/vaults/deploy/action**
 - **Microsoft.ManagedIdentity/userAssignedIdentities/assign/action**
 - **Microsoft.Network/applicationGateways/backendAddressPools/join/action**
 - **Microsoft.Network/applicationSecurityGroups/joinIpConfiguration/action**
 - **Microsoft.Network/applicationSecurityGroups/joinNetworkSecurityRule/action**
 - **Microsoft.Network/ddosProtectionPlans/join/action**
 - **Microsoft.Network/gatewayLoadBalancerAliases/join/action**
 - **Microsoft.Network/loadBalancers/backendAddressPools/join/action**
 - **Microsoft.Network/loadBalancers/frontendIPConfigurations/join/action**
 - **Microsoft.Network/loadBalancers/inboundNatPools/join/action**
 - **Microsoft.Network/loadBalancers/inboundNatRules/join/action**
 - **Microsoft.Network/networkInterfaces/join/action**
 - **Microsoft.Network/networkSecurityGroups/join/action**
 - **Microsoft.Network/publicIPAddresses/join/action**
 - **Microsoft.Network/publicIPPrefixes/join/action**

- **Microsoft.Network/virtualNetworks/subnets/join/action**

([OCPBUGS-44130](#))

- Previously, installing an AWS cluster failed in certain environments on existing subnets when the **publicip** parameter in the compute machine set CR was set to **false**. With this release, a fix ensures that a configuration value set for **publicip** no longer causes issues when the installation program provisions machines for your AWS cluster in certain environments. ([OCPBUGS-44373](#))
- Previously, Google Cloud clusters that used non-UEFI disks failed to load. This release adds a check to ensure that disks are UEFI-compatible before enabling features that require UEFI, such as secure boot. This change adds **compute.images.get** and **compute.images.getFromFamily** permissions requirements. As a result, you can use non-UEFI disks if you do not need these features. ([OCPBUGS-44671](#))
- Previously, when the AWS **DHCPOptionSet** parameter was configured to use a custom domain name that contains a trailing period (.), OpenShift Container Platform installation failed. With this release, the logic that extracts the hostname of EC2 instances and turns them into kubelet node names trims trailing periods so that the resulting Kubernetes object name is valid. Trailing periods in this parameter no longer cause installation to fail. ([OCPBUGS-45306](#))
- Previously, the number of Azure availability set fault domains used a fixed value of **2**. This setting works in most Azure regions because fault domain counts are typically at least 2. However, this setting failed in the **centraluseuap** and **eastusstg** regions. With this release, the number of availability set fault domains in a region is set dynamically. ([OCPBUGS-45663](#))
- Previously, the Azure cloud controller manager panicked when there was a temporary API server disconnection. With this release, the Azure cloud controller manager correctly recovers from temporary disconnection. ([OCPBUGS-45859](#))
- Previously, some services became stuck in a pending state due to incorrect or missing annotations. With this release, validation added to the Azure **service.beta.kubernetes.io/azure-load-balancer-tcp-idle-timeout** and Google Cloud **cloud.google.com/network-tier** annotations resolves the issue. ([OCPBUGS-48481](#))
- Previously, the method used to fetch the provider ID from AWS could fail to provide this value to the kubelet when needed. As a result, sometimes machines could get stuck in different states and fail to complete initialization. With this release, the provider ID is consistently set when the kubelet starts up. ([OCPBUGS-50905](#))
- Previously, an incorrect endpoint in the Azure cloud controller manager caused installations on Microsoft Azure Government Cloud to fail. The issue is resolved in this release. ([OCPBUGS-50969](#))
- Previously, the Machine API sometimes detected an unhealthy control plane node during cluster creation on IBM Cloud and attempted to replace the node. This effectively destroyed the cluster. With this release, the Machine API only attempts to replace unhealthy compute nodes during cluster creation and does not attempt to replace unhealthy control plane nodes. ([OCPBUGS-51864](#))
- Previously, Azure spot machines that were evicted before their node became ready could get stuck in the **provisioned** state. With this release, Azure spot instances now use a delete-eviction policy. This policy ensures that the machines correctly move to the **failed** state upon preemption. ([OCPBUGS-54617](#))

- Previously, a bug fix altered the availability set configuration by changing the fault domain count to use the maximum available value instead of a fixed value of **2**. This inadvertently caused scaling issues for compute machine sets created before the bug fix, as the controller attempted to change immutable availability sets. With this release, availability sets are no longer modified after creation, allowing affected compute machine sets to scale properly. ([OCPBUGS-56653](#))
- Previously, the **openshift-cnv** namespace components did not feature the **openshift.io/required-scc** annotation. Workloads were not requesting their required security content constraints (SCCs). With this release, the **openshift.io/required-scc** annotation is added to the **openshift-cnv** namespace components so that workloads can request the required SCCs. ([OCPBUGS-49657](#))

1.6.4. Cloud Credential Operator

- Previously, the **aws-sdk-go-v2** software development kit (SDK) failed to authenticate an **AssumeRoleWithWebIdentity** API operation on an Amazon Web Services (AWS) Security Token Service (STS) cluster. With this release, **pod-identity-webhook** now includes a default region so that this issue no longer persists. ([OCPBUGS-41727](#))

1.6.5. Cluster Autoscaler

- Previously, when a Machine Set was scaled down and had reached its minimum size, the Cluster Autoscaler could leave the last remaining node with a no schedule taint that prevented use of a node. This issue was caused by a counting error in the Cluster Autoscaler. With this release, the counting error has been fixed so that the Cluster Autoscaler works as expected when a Machine Set is scaled down and has reached its minimum size. ([OCPBUGS-54231](#))
- Previously, some cluster autoscaler metrics were not initialized, and therefore were not available. With this release, these metrics are initialized and available. ([OCPBUGS-25852](#))
- Previously the Cluster Autoscaler could stop scaling because of a failed machine in a machine set. This condition occurred because of inaccuracies in the way the Cluster Autoscaler counts machines in various non-running phases. With this release, the inaccuracies have been fixed, so that the Cluster Autoscaler accurately counts machines. ([OCPBUGS-11115](#))

1.6.6. Cluster Resource Override Admission Operator

- Previously, the Cluster Resource Admission Override Operator failed to delete old secrets during upgrading from OpenShift Container Platform 4.16 to OpenShift Container Platform 4.17. This situation caused the Cluster Resource Override Admission Operator webhook to stop working and prevented pods from being created in namespaces that had the Cluster Resource Override Admission Operator enabled. With this release, old secrets are deleted, error handling by the Cluster Resource Override Admission Operator is improved, and the issue with creating pods in namespaces is resolved. ([OCPBUGS-54886](#))
- Previously, if you deleted the **clusterresourceoverride-operator** service or uninstalled the Cluster Resource Admission Override Operator, the **v1.admission.autoscaling.openshift.io** API service was unreachable and prevented needed cluster functions, such as installing other Operators on the cluster. With this release, a fix ensures that if the Cluster Resource Admission Override Operator is uninstalled, the **v1.admission.autoscaling.openshift.io** API service is also deleted so that the cluster functions are not impacted. ([OCPBUGS-48115](#))
- Previously, if you specified a **forceSelinuxRelabel** parameter in a **ClusterResourceOverride** CR and then changed the parameter to another value, the changed value would not be reflected in the **clusterresourceoverride-configuration** Config Map. This Config Map is required for

applying the selinux relabelling workaround feature to your cluster. With this release, this issue is fixed so that when the **forceSelinuxRelabel** parameter is changed, the **clusterresourceoverride-configuration** Config Map received the update. ([OCPBUGS-44649](#))

1.6.7. Cluster Version Operator

- Previously, the status of the **ClusterVersion** condition could be changed from **ImplicitlyEnabled** to **ImplicitlyEnabledCapabilities**. With this release, the **ClusterVersion** condition type has been fixed and changed from **ImplicitlyEnabled** to **ImplicitlyEnabledCapabilities**. ([OCPBUGS-56771](#))
- Previously, a custom Security Context Constraint (SCC) impacted any pod that was generated by the Cluster Version Operator from receiving a cluster version upgrade. With this release, OpenShift Container Platform now sets a default SCC to each pod, so that any custom SCC created does not impact a pod. ([OCPBUGS-31462](#))
- Previously, when a Cluster Operator takes a long time to upgrade, Cluster Version Operator does not report anything as it cannot determine if the upgrade is still progressing or already stuck. With this release, a new unknown status is added for the failing condition in status of the Cluster Version reported by Cluster Version Operator to remind the cluster administrators to check the cluster and avoid waiting on a blocked Cluster Operator upgrade. ([OCPBUGS-23514](#))

1.6.8. etcd

- Before this update, during rolling cluster updates from etcd 3.5.19 to a release of 3.6, the wrong membership data could be propagated to new members. As a consequence, cluster updates failed with an error about too many learner members in the cluster. With this release, etcd is updated to 3.5.24, which includes fixes so that the membership-related errors no longer occur. ([OCPBUGS-63473](#))

1.6.9. ImageStreams

- Previously, image import blocked registries that would fail if those registries were configured with **NeverContactSource**, even when mirror registries were set up. With this update, image importing is no longer blocked when a registry has mirrors configured. This ensures that image imports succeed even if the original source was set to **NeverContactSource** in the **ImageDigestMirrorSet** or **ImageTagMirrorSet** resources. ([OCPBUGS-44432](#))

1.6.10. Installer

- Previously, if you attempted to install an Amazon Web Services (AWS) cluster with minimum privileges and you did not specify an instance type in the **install-config.yaml** file, installation of the cluster failed. This issue happened because the installation program could not find supported instance types that the cluster could use in supported availability zones. For example, the **m6i.xlarge** default instance type was unavailable in **ap-southeast-4** and **eu-south-2** availability zones. With this release, the **openshift-install** program now requires the **ec2:DescribeInstanceTypeOfferings** AWS permission to prevent the installation of the cluster from failing in situations where **m6i.xlarge** or another supported instance type is unavailable in a supported availability zone. ([OCPBUGS-46596](#))
- Previously, the installation program did not prevent users from attempting to install a single-node cluster on bare metal, which resulted in a failed installation. With this update, the installation program prevents single-node cluster installations on unsupported platforms. ([OCPBUGS-56811](#))

- Previously, when you diagnosed issues related to running the **openshift-install destroy cluster** command for VMware vSphere, the logging information provided insufficient detail. As a consequence, it was unclear why clusters were not removed from virtual machines (VMs). With this release, when you destroy a cluster, enhanced debug logging is provided and the issue is resolved. ([OCPBUGS-56372](#))
- Previously, when installing into an existing virtual private cloud (VPC) on Amazon Web Services (AWS), a potential mismatch could occur in the subnet information in the AWS Availability Zone between the machine set custom resources for control plane nodes and their corresponding AWS EC2 instances. As a consequence, where the control plane nodes were spread across three Availability Zones and one was recreated the discrepancy could result in an unbalanced control plane as two nodes occurred within the same Availability Zone. With this release, it is ensured that the subnet Availability Zone information in the machine set custom resources and in the EC2 instances match and the issue is resolved. ([OCPBUGS-55492](#))
- Previously, when installing a cluster with the **OVNKubernetes** network plugin, the installation could fail if the plugin is specified as **OVNkubernetes** with a lowercase "k". With this update, the installation program correctly interprets the plugin name regardless of case. ([OCPBUGS-54606](#))
- When a Proxy is configured, the installation program adds the **machineNetwork** CIDR to the **noProxy** field. Previously, if the **machineNetwork** CIDR had also been configured by the user in the **noProxy** field, this would result in a duplicate entry, which is not allowed by ignition and could prevent the host from booting properly. With this release, the installation program will not add the **machineNetwork** CIDR to the **noProxy** field if it has already been set. ([OCPBUGS-53183](#))
- Previously, API and ingress VIPs were automatically assigned even when a user-managed load balancer was in use. This behavior was unintended. Now, API and ingress VIPs are no longer automatically assigned. If these values are not explicitly set in the **install-config.yaml** file, the installation fails with an error, prompting the user to provide them. ([OCPBUGS-53140](#))
- Previously, when using the Agent-based Installer, the WWN of Fibre Channel (FC) multipath volumes was not detected during hardware discovery. As a result, when the **wwn** root device hint was specified, all multipath FC volumes were excluded by it. With this release, the WWN is now collected for multipath FC volumes, so when more than one multipath volume is present, users can select between them by using the **wwn** root device hint. ([OCPBUGS-52994](#))
- Previously, when installing a cluster on Azure, the installation program did not include support for NVMe or SCSI, which prevented the use of VM instance families that require it. With this update, the installation program can make use of VM instance families that require NVMe or SCSI support. ([OCPBUGS-52658](#))
- Previously, when installing a cluster on Google Cloud with a user-provided encryption key, the installation program could fail to find the key ring. With this update, the installation program finds the user-provided encryption key ring so the installation does not fail. ([OCPBUGS-52203](#))
- Previously, when installing a cluster on Google Cloud, the installation could fail if network instability prevented the fetching of Google Cloud tags during installation. With this update, the installation program has been improved to tolerate network instability during installation. ([OCPBUGS-50919](#))
- Previously, the installer was not checking for ESXi hosts that were powered off within a VMware vSphere cluster, which caused the installation to fail because the OVA could not be uploaded. With this release, the installer now checks the power status of each ESXi host and skips any that

are powered off, which resolves the issue and allows the OVA to be imported successfully. ([OCPBUGS-50649](#))

- Previously, when using the Agent-based Installer, erroneous error messages regarding **unable to read image** were output when building the Agent ISO image in a disconnected environment. With this release, these erroneous messages have been removed and no longer appear. ([OCPBUGS-50637](#))
- Previously, when installing a cluster on Azure, the installation program would crash with a segmentation fault error if it did not have the correct permissions to check IP address availability. With this update, the installation program correctly identifies the missing permission and fails gracefully. ([OCPBUGS-50534](#))
- Previously, when the **ClusterNetwork** classless inter-domain routing (CIDR) mask value is greater than the **hostPrefix** value and the **networking.ovnKubernetesConfig.ipv4.internalJoinSubnet** section is provided in the **install-config.yaml** file, the installation program failed a validation check and returned a Golang runtime error. With this release, the installation program still fails the validation check and now outputs a descriptive error message that indicates the invalid **hostPrefix** value. ([OCPBUGS-49784](#))
- Previously, when installing a cluster on IBM Cloud®, the installation program failed to install on the **ca-mon** region even though it is available. With this update, the installation program is up to date with the latest available IBM Cloud® regions. ([OCPBUGS-49623](#))
- Previously, after installing a cluster on AWS with minimum permissions in an existing VPC with a user-provided public IPv4 pool, the cluster could not be destroyed due to a missing permission. With this update, the installation program propagates the **ec2:ReleaseAddress** permission so that the cluster can be destroyed. ([OCPBUGS-49594](#))
- Previously, the installer for VMware vSphere did not validate the number of networks provided in the **install-config.yaml** for failure domains. This caused the installation to proceed with an unsupported configuration if more than the maximum of 10 networks were specified, without providing an error. With this release, the installer now validates the number of configured networks, which resolves the issue by preventing the use of a configuration that exceeds the maximum limit. ([OCPBUGS-49351](#))
- Previously, installing a cluster on AWS with existing subnets (BYO VPC) in Local or Wavelength zones resulted in the edge subnets resource missing the **kubernetes.io/cluster/<InfracID>:shared** tag. With this release, a fix ensures that all subnets used in the **install-config.yaml** file have the required tags. ([OCPBUGS-48827](#))
- Previously, an issue prevented configuring multiple subnets in the failure domain of a Nutanix cluster during installation. The issue is resolved in this release. ([OCPBUGS-49885](#))
- Previously, when installing a cluster on AWS, the **ap-southeast-5** region was not available in the installation program survey, even though this region was supported by OpenShift Container Platform. With this update, the **ap-southeast-5** region is available. ([OCPBUGS-47681](#))
- Previously, when destroying a cluster installed on Google Cloud, some resources could be left behind because the installation program did not wait to ensure that all destroy operations completed successfully. With this update, the destroy API waits to ensure that all resources are appropriately deleted. ([OCPBUGS-47489](#))
- Previously, when installing a cluster on AWS in the **us-east-1** regions, the installation could fail if no zone is specified in the **install-config.yaml** file because the **use1-az3** zone does not support

any instance types supported by OpenShift Container Platform. With this update, the installation program prevents the use of the **use1-az3** zone when no zones are specified in the installation configuration file. ([OCPBUGS-47477](#))

- Previously, when installing a cluster on Google Cloud, the installation would fail if you enabled the **constraints/compute.vmCanIpForward** constraint on your project. With this update, the installation program disables this constraint if it is enabled, allowing the installation to succeed. ([OCPBUGS-46571](#))
- Previously, when installing a cluster on Google Cloud, the installation program would fail to detect if the user provided an encryption key ring that did not exist, causing the installation to fail. With this update, the installation program correctly validates the existence of user provided encryption key rings, preventing failure. ([OCPBUGS-46488](#))
- Previously, when destroying a cluster that was installed on Microsoft Azure, the inbound NAT rules and security groups for the bootstrap node were not deleted. With this update, the correct resource group ensures that all resources are deleted when the cluster is destroyed. ([OCPBUGS-45429](#))
- Previously, when installing a cluster on AWS in the **ap-southeast-5** region, the installation could fail due to a malformed load balancer hostname. With this update, the installation program has been improved to form the correct hostname so that installation succeeds. ([OCPBUGS-45289](#))
- Previously, when installing a cluster on Google Cloud, the installation program could fail to locate the service account it created due to a delay in activating the service account on Google's servers. With this update, the installation program waits an appropriate amount of time before attempting to use the created service account. ([OCPBUGS-45280](#))
- Previously, when installing a cluster on AWS, the installation could fail if you specified an edge machine pool but did not specify an instance type. With this update, the installation program requires that an instance type be provided for edge machine pools. ([OCPBUGS-45218](#))
- Previously, when destroying a cluster installed on Google Cloud, PVC disks with the label **kubernetes-io-cluster-<cluster-id>: owned** were not deleted. With this update, the installation program correctly locates and deletes these resources when the cluster is destroyed. ([OCPBUGS-45162](#))
- Previously, during a disconnected installation, when the **imageContentSources** parameter was configured for more than one mirror for a source, the command to create the agent ISO image could fail, depending on the sequence of the mirror configuration. With this release, multiple mirrors are handled correctly when the agent ISO is created and the issue is resolved. ([OCPBUGS-44938](#))
- Previously, when installing a cluster on AWS, if the **publicIPv4Pool** parameter was set but the **ec2:AllocateAddress** permission was not present, the installation would fail. With this update, the installation program requires that this permission is present. ([OCPBUGS-44925](#))
- Previously, during a shared Virtual Private Cloud (VPC) installation, the installer added the records to a private DNS zone created by the installer instead of adding the records to the cluster's private DNS zone. As a consequence, the installation failed. With this release, the installer searches for an existing private DNS zone and, if found, pairs that zone with the network that is supplied by the **install-config.yaml** file and the issue is resolved. ([OCPBUGS-44641](#))
- Previously, you could add white space to Amazon Web Services (AWS) tag names but the installation program did not support them. This situation resulted in the installation program outputting an **ERROR failed to fetch Metadata** message. With this release, the regular

expression for AWS tags now validates any tag name that has white space so that the installation program accepts these tags and no longer outputs an error because of white space. ([OCBBUGS-44199](#))

- Previously, when destroying a cluster that was installing on Google Cloud, forwarding rules, health checks and firewall rules were not deleted, leading to errors. With this update, all resources are deleted when the cluster is destroyed. ([OCBBUGS-43779](#))
- Previously, when installing a cluster on Microsoft Azure, specifying the **Standard_M8-4ms** instance type resulted in an error due to that instance type specifying its memory in decimal format instead of integer format. With this update, the installation program correctly parses the memory value. ([OCBBUGS-42241](#))
- Previously, when installing a cluster on VMware vSphere, installations could fail if the API and Ingress server virtual IPs were outside of the machine network. With this update, the installation program includes the API and Ingress server virtual IPs in the machine network by default. If you specify the API and Ingress server virtual IPs, ensure that they are in the machine network. ([OCBBUGS-36553](#))
- Previously, when installing a cluster on IBM Power Virtual Server, the installation failed if you selected the Madrid zone due to an image import error. With this update, the installation program has been modified to use the correct storage bucket name and continue the installation successfully. ([OCBBUGS-50899](#))
- Previously, when destroying a cluster installed on IBM Power Virtual Server, some resources including network subnets were not deleted. With this update, all network resources are deleted when the cluster is destroyed. ([OCBBUGS-50657](#))
- Previously, when installing a cluster using the Assisted Installer, the installation could fail due to a timeout when pulling images. With this update, the timeout has been increased so that the installation program completes pulling images. ([OCBBUGS-50655](#))
- Previously, in some slower PrismCentral environments, the installation program would fail with a timeout when make the prism-api call to load the RHCOS image. The timeout value was previously 5 minutes. With this release, the prism-api call timeout value is a configurable parameter in the **install-config.yaml** file as **platform.nutanix.prismAPICallTimeout**, with a default value of 10 minutes. ([OCBBUGS-48570](#))
- Previously, an issue prevented configuring multiple subnets in the failure domain of a Nutanix cluster during installation. The issue is resolved in this release. ([OCBBUGS-48044](#))
- Previously, when installing a cluster on IBM Power Virtual Server using installer-provisioned infrastructure, the installation program selected a random machine network instead of using the one provided by the user. With this update, the installation program uses the machine network that the user provides. ([OCBBUGS-45286](#))
- Previously, the temporary directory that was created when the **openshift-install agent create pxe-files** command was run was not removed after the command completed. With this release, the temporary directory is now removed properly after the command completes. ([OCBBUGS-39583](#))

1.6.11. Machine Config Operator

- Previously, the **ContainerRuntimeConfig** incorrectly set the **--root** path for the **runc** runtime. This caused containers to run with an incorrect root path and led to issues with the container operations. With this release, the **--root** path for container runtime is correct and matches the

specified runtime, providing consistent operation. ([OCPBUGS-47629](#))

- Previously, users were not warned if their cluster contained Red Hat Enterprise Linux (RHEL) worker nodes, which are no longer supported in OpenShift Container Platform 4.19 and later. With this release, the Machine Config Operator detects RHEL nodes and notifies users that are not compatible with OpenShift Container Platform 4.19. ([OCPBUGS-54611](#))
- Previously, when the Machine Config Operator (MCO) rebooted a node too quickly after staging an update, the update failed. With this release, the MCO waits for the staging operation to finish before rebooting the system, allowing the update to complete. ([OCPBUGS-51150](#))
- Previously, after deleting a **MachineOSConfig** object, the associated **MachineOSBuild** object was not deleted as expected. This was because the ownership of the **MachineOSBuild** object was not set. With this release, all of the objects are created for the build and all associated objects are deleted when the **MachineOSConfig** object is deleted. ([OCPBUGS-44602](#))

1.6.12. Management Console

- Previously, the **Projects details** in the **Developer Perspective** erroneously did not include breadcrumbs. With this release, the breadcrumbs have been added. ([OCPBUGS-52298](#))
- Previously, when opening the **Project** drop-down list while the web terminal is open, there were visual artifacts. After this update, the artifact was fixed, so you can use the **Project** drop-down list when the web terminal is open. ([OCPBUGS-45325](#))
- Previously, a **PipelineRuns** CR that used a resolver could not be rerun on the OpenShift Container Platform web console. If you attempted to rerun the CR, an "Invalid **PipelineRun** configuration, unable to start Pipeline" was generated. With this release, you can now rerun a **PipelineRuns** CR that uses resolver without experiencing this issue. ([OCPBUGS-44265](#))
- Previously, when you used the **Form View** to edit **Deployment** or **DeploymentConfig** API objects on the OpenShift Container Platform web console, duplicate **ImagePullSecrets** parameters existed in the YAML configuration for either object. With this release, a fix ensures that duplicate **ImagePullSecrets** parameters do not get automatically added for either object. ([OCPBUGS-41974](#))
- Previously, a **TaskRun** for a particular **PipelineRun** was fetched based on the **PipelineRun** name. If any two **PipelineRuns** had the same name, the **TaskRun** for both **PipelineRuns** was fetched and displayed. With this release, the **TaskRun** for the particular **PipelineRun** will be fetched based on **PipelineRun** UID instead of **PipelineRun** name. ([OCPBUGS-36658](#))
- Previously, the **Test Serverless function** button did not respond if there were no pods running. With this update, the button is disabled when there are no pods running. ([OCPBUGS-32406](#))
- Previously, the results of a failed **TaskRun** would not show up in the UI. With this update, the results of a **TaskRun** are always available, regardless of failure. ([OCPBUGS-23924](#))
- Previously, the console alerted users that compute nodes must be updated within 60 days when performing a control plane only update. With this release, the console no longer displays this invalid alert. ([OCPBUGS-56077](#))
- Previously, the **Critical Alerts** section of the **Notification Drawer** could not be collapsed. With this release, the section can be collapsed. ([OCPBUGS-55702](#))
- Previously, when viewing the list of installed Operators, an Operator displayed twice in the list if the currently selected project matched an Operator's default namespace while copied CSVs

were disabled in Operator Lifecycle Manager (OLM). With this release, the Operator displays only once in such cases. ([OCPBUGS-54601](#))

- Previously, the link to **OperatorHub** on the **Installed Operators** page would trigger a hard reload. With this release, this link no longer triggers a hard reload. ([OCPBUGS-54536](#))
- Previously, selecting **All Projects** from the project picker while on the **Create VolumeSnapshot** page resulted in a page not found error. With this release, the VolumeSnapshot list page is correctly displayed. ([OCPBUGS-53227](#))
- Previously, there was incorrect logic for calculating the pod container count causing it to be inaccurate. With this release, the **Ready** and **Started** status in count logic was added so the correct pod container count is displayed, which is consistent with **oc** CLI. ([OCPBUGS-53118](#))
- Previously, the **Select** menu above the **Node Logs** section did not close when opened, unless the **Select** menu's toggle was clicked again or one of the Select's menu items was clicked. With this release, the **Select** menu closes after clicking outside of the menu or by pressing the appropriate key on the keyboard. ([OCPBUGS-52316](#))
- Previously, the shared timestamp component referenced an undefined property when calculating relative times. As a result, most times displayed in the console were not correctly displaying relative strings such as **Just now** or **Less than a minute ago**. With this release, the issue is fixed and relative time strings are correctly rendered in the console. ([OCPBUGS-51202](#))
- Previously, the **Observe** menu only displayed based on the current user and console configuration for monitoring. This caused other items added by observability plugins to be hidden. With this release, the **Observe** menu also displays items from different observability plugins. ([OCPBUGS-50693](#))
- Previously, when logging in to the console for the first time, automatic perspective detection caused the console to ignore the specific URL path that the user clicked to get onto the console, and instead load a different page. With this release, the current path is followed. ([OCPBUGS-50650](#))
- Previously, there was an issue when creating new tabs in the horizontal navigation present in the web console from plugins. With this release, you can use plugins to create tabs on the web console horizontal navigation. ([OCPBUGS-49996](#))
- Previously, the **Cluster Settings** page would not properly render during a cluster update if the **ClusterVersion** did not receive a **Completed** update. With this release, the **Cluster Setting** page properly renders even if the **ClusterVersion** has not received a **Completed** update. ([OCPBUGS-49839](#))
- Previously, the links on the CLI downloads page were not sorted by operating system. With this release, the links are sorted by their operating system in alphabetical order. ([OCPBUGS-48413](#))
- Previously, multiple external link icons could appear in the primary **Action** button of the **OperatorHub** modal. With this release, only a single external link icon is displayed. ([OCPBUGS-46555](#))
- Previously, clicking the **Don't show again** link in the Red Hat OpenShift Lightspeed modal did not correctly navigate to the general **User Preference** tab when one of the other **User Preference** tabs was displayed. After this update, clicking the **Don't show again** link correctly navigates to the general **User Preference** tab. ([OCPBUGS-46511](#))
- Previously, a console plugin could be enabled multiple times in the **Console plugin enablement** modal, resulting in multiple entries for the plugin appearing in the Console Operator

Configuration. With this release, it is no longer possible to enable a plugin if it is already enabled. ([OCPBUGS-44595](#))

- Previously, the OpenShift Container Platform web console login page always allowed you to click the **Login** button. You could still click when no username or password was entered, or if the **Login** button was already clicked. With this release, the **Login** button is disabled so you are unable to click the **Login** button without a username or password. ([OCPBUGS-43610](#))
- Previously, **PackageManifest** by name only was selected on the **Operator installation** status page. In some cases, this caused the incorrect **PackageManifest** to be used for displaying the logo and provider since there can be name collisions. With this release, **PackageManifests** are selected by name and label selector to make sure the correct one is selected for the current installation. As a result, the correct logo and provider are always shown on the operator install status page. ([OCPBUGS-21755](#))

1.6.13. Monitoring

- Previously, if a scrape failed, Prometheus erroneously considered the samples from the very next scrape as duplicates and dropped them. This issue affected only the scrape immediately following a failure, while subsequent scrapes were processed correctly. With this release, the scrape following a failure is now correctly handled, ensuring that no valid samples are mistakenly dropped. ([OCPBUGS-53025](#))

1.6.14. Networking

- Previously, when a pod used the CNI plugin for DHCP address assignment, in conjunction with other CNI plugins, the network interface of the pod might have been unexpectedly deleted. As a consequence, when the DHCP lease expired for the pod, the DHCP proxy entered a loop when trying to re-create a new lease, leading to the node becoming unresponsive. With this release, the DHCP lease maintenance terminates if the network interface does not exist. As a result, interface deletions are handled gracefully, ensuring node stability. ([OCPBUGS-45272](#))
- Previously, the Kubernetes NMState Operator Operator did not create the **nmstate-console-plugin** pod because of an issue with the **pluginPort** template. With this release, a fix to the template ensures that the Operator can now successfully create the **nmstate-console-plugin** pod. ([OCPBUGS-54295](#))
- Previously, the pod controller in the Whereabouts reconciler was not passing the namespace to the leader election function, so the pod controller was not deleting orphaned allocations. This led to repeated log error messages. With this release, the namespace is passed in and the orphaned allocations are deleted properly. ([OCPBUGS-53397](#))
- Previously, the **SriovOperatorConfig** Operator removed any parameters that had default values in the **SriovOperatorConfig** resources. This situation caused certain information to be missing from the output of a resource. With this release, the Operator uses the PATCH method for API servers to preserve parameters with default values so that no information is missing from the output for a resource. ([OCPBUGS-53346](#))
- Previously, the **SriovNetworkNodePolicy** object reconciler executed with every node resource update. This resulted in excessive resource consumption by the SR-IOV Operator pod and an overabundance of log entries. This release changes the behavior so that the reconciler only runs when a node label changes, thereby reducing resource consumption and log entry generation. ([OCPBUGS-52955](#))
- Previously, a cluster with the **clusterNetwork** parameter listing multiple networks of the same IP

address family entered a **crashloopbackoff** state when upgrading to the latest version of OpenShift Container Platform. With this release, a fix ensures that the cluster with this configuration no longer enters the **crashloopbackoff** state during a cluster upgrade. ([OCPBUGS-49994](#))

- Previously, the **resolv-prepender** service triggered earlier than expected. This situation caused the service to fail and resulted in an incorrectly configured setting for the host DNS. With this release, the configuration for the **resolv-prepender** service is updated so that when the service triggers earlier than expected, it no longer causes incorrect configuration of the host DNS settings. ([OCPBUGS-49436](#))
- Previously, the **nmstate-configuration** service was only enabled for deployments that had the **platform** parameter set to **baremetal**. However, you could also use the Assisted Installer to configure a bare-metal deployment by setting the **platform** parameter set to **None**, but the NMState **br-ex** network bridge creation feature did not work with this installation method. With this release, the **nmstate-configuration** service is moved to the base directory in the cluster installation path so that any deployments configured with the **platform** parameter set to **None**, do not impact the NMState **br-ex** network bridge creation feature. ([OCPBUGS-48566](#))
- Previously, for a layer 2 or layer 3 topology network that had the gateway mode set to **local**, OVN-Kubernetes experienced an issue when it was restarted. The issue caused the Egress IP to be selected as the primary IP address for the network. With this release, a fix ensures that this behavior no longer occurs. ([OCPBUGS-46585](#))
- Previously, the DNS-based egress firewall incorrectly prevented creation of any firewall rule that contained a DNS name in uppercase characters. With this release, a fix to the egress firewall means that creation of a firewall rule that contains a DNS name in uppercase occurs. ([OCPBUGS-46564](#))
- Previously, when a pod was running on a node on which an egress on the IPv6 protocol was assigned, the pod was not able to communicate with the OVN-Kubernetes service in a dual-stack cluster. This resulted in the traffic with the IP address family, that the **egressIP** was not applicable to, being dropped. With this release, only the source network address translation (SNAT) for the IP address family that the egress IPs applied to is deleted, eliminating the risk of traffic being dropped. ([OCPBUGS-46543](#))
- Previously, when you used static IP addresses in the customized **br-ex** network bridge configuration of a manifest object, a race condition was added and caused a node reboot operation that further impacted deployment of a cluster. With this release, the **nodeip-configuration** service now starts after the **br-ex** network bridge is up, preventing the race condition and node reboot. ([OCPBUGS-46072](#))
- Previously, the HAProxy router incorrectly assumed that only SHA1 leaf certificates were rejected by HAProxy, causing the router to fail by not rejecting SHA1 intermediate certificates. With this update, the router now inspects and rejects all non-self-signed SHA1 certificates, thereby preventing crashes and improving stability for your cluster stability. ([OCPBUGS-45290](#))
- Previously, when a node restarted the **openvswitch** daemon, the **nmstate-handler** container could not access the OpenVSwitch (OVS) database and this caused all OVS-related NNCP configurations to fail. With this release, the issue is fixed. The **nmstate-handler** container can access the OVS database even after restarting the OVS process on the node. The **nmstate-handler** no longer requires a manual restart. ([OCPBUGS-44596](#))
- Previously, a **MultiNetworkPolicy** API was not enforced when the **protocol** parameter was

specified, but the **port** parameter was not, in the cluster configuration. This situation caused all network traffic to reach the cluster. With this release, the **MultiNetworkPolicy** API policy only allows connections from and to the ports specified with the **protocol** parameter so that only specific traffic reaches the cluster. ([OCPBUGS-44354](#))

- Previously, HAProxy left idle connections open when it reloaded its configuration until either the next time a client sent a request by using the idle connection or the **hard-stop-after** period elapsed. This release adds a new **IdleConnectionTerminationPolicy** API field to control HAProxy behavior for idle connections during reloads. The new default setting is **Immediate**, which means that HAProxy immediately terminates any idle connections when it reloads its configuration. The previous behavior can be specified by using the **Deferred** setting for **IdleConnectionTerminationPolicy**. ([OCPBUGS-43745](#))
- Previously, if an application did not use the Path MTU discovery (PMTUD) mechanism while sending UDP packets larger than the network MTU, an issue with the **OVN** package caused the a packet to be dropped while fragmenting the packet. With this release, the **OVN** package has been fixed so that large UDP packets are properly fragmented and sent over the network. ([OCPBUGS-43649](#))
- Previously, a pod with a secondary interface in an OVN-Kubernetes **Localnet** network that was plugged into a **br-ex** interface bridge was out of reach by other pods on the same node, but used the default network for communication. The communication between pods on different nodes was not impacted. With this release, the communication between a **Localnet** pod and a default network pod running on the same node is possible, however the IP addresses that are used in the **Localnet** network must be within the same subnet as the host network. ([OCPBUGS-43004](#))
- Previously, when specific network changes were made to a running cluster, a **NetworkManager** connection profile was permanently created by the **ovs-configuration** service and the profile was incorrectly saved to storage. This profile file would persist through a reboot operation and caused the **ovs-configuration** service to fail. With this release, the **ovs-configuration** cleanup process is updated to remove any unnecessary files, preventing such files from causing issues after a reboot operation. ([OCPBUGS-41489](#))
- Previously, the **parseIPList** function failed to handle IP address lists that contained both valid and invalid IP addresses or CIDR ranges. This situation caused the function to return an empty string when it encountered an invalid entry and skip processing valid entries. With this release, the **haproxy.router.openshift.io/ip_allowlist** route annotation skips any invalid IP addresses or CIDR ranges so that the **parseIPList** function can process all listed entries. ([OCPBUGS-39403](#))
- Previously, the HAProxy router lacked out-of-bounds validation for the **router.openshift.io/haproxy.health.check.interval** annotation. If you set a value that exceeded the maximum value that the HAProxy router could handle, the **router-default** pod could not reach the **Ready** state. With this release, the router now validates the value for the annotation and excludes values that are out of bounds. The router now functions as expected. ([OCPBUGS-38078](#))
- Previously, in certain situations the gateway IP address for a node changed and caused the **OVN** cluster router, which manages the static route to the cluster subnet, to add a new static route with the new gateway IP address, without deleting the original one. As a result, a stale route still pointed to the switch subnet and this caused intermittent drops during egress traffic transfer. With this release, a patch applied to the **OVN** cluster router ensures that if the gateway

IP address changes, the **OVN** cluster router updates the existing static route with the new gateway IP address. A stale route no longer points to the **OVN** cluster router so that egress traffic flow does not drop. ([OCPBUGS-32754](#))

- Previously, there was no event logged when an error occurred from a failed conversion from ingress to route. With this update, an error for a failed conversion is logged. ([OCPBUGS-29354](#))
- Previously, the PowerVS installer used a hard-coded list of supported machine types. However, this list was not always updated as new types were added. With this release, datacenter is queried to get the current list of supported types. ([OCPBUGS-49940](#))
- Previously, when a `RootDiskHint` was defined and the installation failed with the error **Requested installation disk is not part of the host's valid disks**, it was difficult to determine the valid disk names that could be used as the hint. With this release, logging was added for the list of acceptable disks so that a user can quickly determine what the root disk hint should be. ([OCPBUGS-43578](#))
- Previously, the `oc adm node-image monitor` command returned an EOF error when there was an API server interruption or temporary connection issue. This caused the command to terminate. With this release, the command now detects API server interruptions and temporary connection issues and reconnects to the API server without terminating the command. ([OCPBUGS-38975](#))
- Previously, when you created a virtual machine (VM) and no IP addresses existed in the IP pool, the VM did not start. An error message was generated in the `virt-launcher-<vm_name>` pod, but the message was unclear as to the source of the issue. With this release, for this situation where no IP addresses exist in the IP pool, the `virt-launcher-<vm-name>` pod includes a clear error message that is similar to the following example:

```
Warning ErrorAllocatingPod 4s (x7 over 79s) ovnk-controlplane failed to update pod
localnet-ipam/virt-launcher-vmb-localnet-ipam-hlnmf: failed to assign pod addresses for
localnet-ipam/ipam-localnet-nad/localnet-ipam/virt-launcher-vmb-localnet-ipam-hlnmf: failed
to allocate new IPs for tenantblue-network: subnet address pool exhausted
```

([OCPBUGS-54245](#))

1.6.15. Node

- Previously, if your cluster used Zscaler and scanned all transfers, it could experience a timeout when pulling images. This problem was due to a hard-coded timeout value for image pulls. The pull progress timeout of CRI-O is now increased to 30 seconds. As a result, previously affected clusters should not experience timeouts. ([OCPBUGS-54662](#))
- Previously, containers that used the `container_logreader_t` SELinux domain to watch container logs on the host at the `/var/log` location could not access the logs. This behavior occurred because the logs in the `/var/log/containers` location were symbolic links. With this fix, containers can watch logs as expected. ([OCPBUGS-48555](#))
- Previously, an end-of-file error occurred in the `json.NewDecoder` file when the file was in a loop operation. This error caused inconsistent application updates to namespaced policies that existed in multiple namespaces. This issue could potential cause security vulnerabilities for the cluster. With this release, a new policy buffer is added to the `json.NewDecoder` file when it enters each loop operation, and a test case is added for multiple namespaces. As a result, the policy buffer provides a robust decoding process for JSON policy files so that namespace policies receive updates without any issues. ([OCPBUGS-48195](#))

- Previously, there was an issue in the image reference digest calculation that led to failed container creation based on the **schemaVersion 1** image. This issue prevented new deployment creations. With this release, the image digest calculation has been fixed and new operators can be installed. ([OCPBUGS-42844](#))
- Previously, for a Technology Preview-enabled cluster with Sigstore verification for payload images in the **policy.json** file, the Podman version in the base image did not support Sigstore configuration. This lack of support caused the new node to be unavailable. With this release, the issue is fixed and the node is available. ([OCPBUGS-38809](#))
- Previously, CPUs for the last guaranteed pod admitted to a node remained allocated after the pod was deleted. This behavior caused scheduling domain inconsistencies. With this release, CPUs that are allocated to guaranteed pods return to the pool of available CPU resources as expected, ensuring correct CPU scheduling for subsequent pods. ([OCPBUGS-17792](#))

1.6.16. Node Tuning Operator (NTO)

- Previously, when applying a performance profile to a node, OpenShift Container Platform selected the appropriate profile based on the vendor identifier of the CPU unit on the node. Because of this, if a CPU uses a different vendor identifier that is not recognized, OpenShift Container Platform failed to include the proper profile. For example the identifier might include APM, rather than ARM. With this fix, for CPUs that use the ARM architecture, the Operator now selects a profile based only on the architecture, and not the vendor identifier. As a result, the correct profile is applied. ([OCPBUGS-52352](#))

1.6.17. Observability

- Previously, the **Silence details** page had an incorrect link URL that was missing the **namespace** parameter, which caused users to be unable to silence specific alerts in the dev console for certain versions. This resulted in poor alert management. With this release, the undefined link in **SilencedAlertsList** has been fixed using the active namespace. As a result, the 'No Alert found' error is now resolved, and the **Alert details** page in OpenShift Container Platform Monitoring is navigated to correctly. ([OCPBUGS-48142](#))
- Previously, console updates deprecated PatternFly 4, rendering an incorrect layout of the monitoring plug-in table rendered. With this release, the tables and styles are upgraded to PatternFly 5 and are rendered correctly. ([OCPBUGS-47535](#))
- Previously, a namespace was passed to a full cluster query on the alerts graph, and this caused the tenancy API path to be used. The API lacked permissions to retrieve data so no data was shown on the alerts graph. With this release, the namespace is no longer passed to a full cluster query for an alert graph. A non-tenancy API path is now used because this API has the correct permissions to retrieve data. Data is not available on an alert graph. ([OCPBUGS-45896](#))
- Previously, a Red Hat Advanced Cluster Management (RHACM) Alerting UI refactor update caused an **isEmpty** check to go missing on the **Observe > Metrics** menu. The missing check inverted the behavior of the **Show all Series** and **Hide all Series** states. This release readds **isEmpty** check so that **Show all Series** is now visible when series are hidden and **Hide all Series** is now visible when the series are shown. ([OCPBUGS-45816](#))
- Previously, on the **Observe → Alerting → Silences** tab, the **DateTime** component changed the ordering of an event and its value. Because of this issue, you could not edit the **until** parameter for a silent alert in the web console. With this release, a fix means to the **DateTime** component means that you can now edit the **until** parameter for a silent alert. ([OCPBUGS-45801](#))

- Previously, bounds were based on the first bar in a bar chart. If a bar was larger in size than the first bar, the bar would extend beyond the bar chart boundary. With this release, the bound for a bar chart is based on the largest bar, so no bars extend outside the boundary of a bar chart. ([OCPBUGS-45174](#))

1.6.18. oc-mirror

- Previously, oc-mirror plugin v2 did not display any progress output during the local cache population phase. For mirror configurations involving a large number of images, this could make the process appear unresponsive or stuck. With this update, a progress bar has been added to provide the cache population status, enabling the user to see the current progress of the cache population. ([OCPBUGS-56563](#))
- Previously, when mirroring Operators using oc-mirror plugin v2, some community Operators with long lists of **skips** and **replaces** entries in their channel graphs caused the mirroring process to run out of memory and fail. With this update, oc-mirror plugin v2 improves the filtering logic by avoiding repeated evaluation of entries referenced in multiple **skips** and **replaces** stanzas, resulting in better memory handling during Operator mirroring. ([OCPBUGS-52471](#))
- Previously, when re-running oc-mirror plugin v2 with the same working directory, existing **tar** archive files from previous runs were not removed. This resulted in a mix of outdated and new archives, which could cause mirroring failures when pushing to the target registry. With this update, oc-mirror plugin v2 automatically deletes old **tar** archive files at the beginning of each run, ensuring that the working directory only contains archives from the current execution ([OCPBUGS-56433](#))
- Previously, oc-mirror plugin v2 would terminate with an error if the source registry responded with any of the following HTTP status codes during image copying: 502, 503, 504. With this update, oc-mirror plugin v2 automatically retries the copy operation when encountering these temporary server errors. ([OCPBUGS-56185](#))
- Previously, when mirroring a Helm chart that included container images with both a tag and digest in the reference, oc-mirror plugin v2 failed with the following error:

Docker references with both a tag and digest are currently not supported.

With this update, oc-mirror plugin v2 supports Helm charts that reference images using both tag and digest. The tool mirrors the image using the digest as the source and applies the tag at the destination. ([OCPBUGS-54891](#))

- Previously, during image cleanup, oc-mirror plugin v2 would stop the deletion process if any error occurred while removing an image. With this release, oc-mirror plugin v2 continues attempting to delete remaining images even after encountering errors. After the process completes, it displays a list of any failed deletions. ([OCPBUGS-54653](#))
- Previously, it was possible to mirror an empty catalog during the mirror-to-disk (m2d) phase if an invalid Operator was specified in the **ImageSetConfiguration** file. This led to a failure during the subsequent disk-to-mirror (d2m) phase. With this release, oc-mirror plugin v2 prevents mirroring empty catalogs by validating operator references in the configuration, ensuring a more reliable mirroring process. ([OCPBUGS-52588](#))
- Previously, when using oc-mirror plugin v2 with the **--dry-run** flag, the **cluster-resources** folder inside the working directory was cleared. As a result, previously generated files such as **idms-oc-mirror.yaml** and **itms-oc-mirror.yaml** were deleted. With this release, the **cluster-resources**

folder is no longer cleared during dry-run operations, preserving any previously generated configuration files. ([OCPBUGS-50963](#))

- Previously, the `oc-mirror` plugin v2 returned an exit status of `0` (success) even when mirroring errors occurred. As a result, `oc-mirror` plugin v2 run failures in automated workflows could go undetected. With this release, `oc-mirror` plugin v2 has been updated to return a non-`0` exit status when mirror failures occur. Despite this fix, users should not rely solely on the exit status in automated workflows. Users are advised to manually check the **mirroring_errors_XXX_XXX.txt** file generated by `oc-mirror` plugin v2 to identify any potential issues. ([OCPBUGS-49880](#))
- Previously, when mirroring using an internal `oc-mirror` reserved keyword—such as **release-images**—in the destination or **--from** path flags, the operation could fail or behave unexpectedly. With this release, `oc-mirror` plugin v2 correctly handles reserved keywords used in destination or source paths. ([OCPBUGS-42862](#))

1.6.19. OpenShift CLI (oc)

- Previously, if you tried to add a node to a disconnected environment using the **oc adm node-image** command, private registry images were inaccessible to the command, causing node addition failure. This error only occurred if the cluster was initially installed with an installer binary downloaded from ([mirror.openshift.com](#)). With this release, a fix has been implemented that enables successful image pull and node creation in disconnected environments. ([OCPBUGS-53106](#))
- For clusters that were installed with the Agent-based Installer for versions 4.15.0 to 4.15.26, root certificates that were built in from CoreOS were added to the user-ca-bundle, even though they were not explicitly specified by the user. In previous releases, when adding a node to one of these clusters using the **oc adm node-image create** command, the **additionalTrustBundle** obtained from the cluster's user-ca-bundle was too large to process, resulting in a failure to add the node. With this release, the built-in certificates are filtered out when generating the **additionalTrustBundle**, so that only explicitly user-configured certificates are included, and nodes can be added successfully. ([OCPBUGS-43990](#))
- Previously, a bug on **oc adm inspect --all-namespaces** command construction meant that `must-gather` was not correctly gathering information about leases, **csistoragecapacities**, and the assisted-installer namespace. With this release, the issue is fixed and `must-gather` will gather the information correctly. ([OCPBUGS-44857](#))
- Previously, the **oc adm node-image create --pxe generated** command did not create only the Preboot Execution Environment (PXE) artifacts. Instead, the command created the PXE artifacts with other artifacts from a **node-joiner** pod and stored them all in the wrong subdirectory. Additionally, the PXE artifacts were incorrectly prefixed with **agent** instead of **node**. With this release, generated PXE artifacts are stored in the correct directory and receive the correct prefix. ([OCPBUGS-45311](#))

1.6.20. Operator Lifecycle Manager (OLM)

- Previously, if an Operator did not have the required **olm.managed=true** label, the Operator might fail and enter a **CrashLoopBackOff** state. When this happened, the logs did not report the status as an error. As a result, the failure was difficult to diagnose. With this update, this type of failure is reported as an error. ([OCPBUGS-56034](#))
- Previously, the Machine Config Operator (MCO) did not search the **/etc/docker/certs.d** directory for certificates required to mount images. As a result, Operator Controller and

catalogd failed to start because they did not have access to the certificates hosted in this directory. With this update, the issue is resolved. ([OCPBUGS-54175](#))

- Before this release, cluster extension updates sometimes failed with the following error from the **CRDUpgradeCheck** resource: **unknown change, refusing to determine that change is safe**. This error occurred due to the way that OLM v1 calculated the difference between version schemas. This update fixes the issue. ([OCPBUGS-53019](#))
- Previously, Operator Controller sometimes failed to mount CA certificates properly. As a result, the Operator Controller failed connect to catalogd due to a TLS certificate validation error. This update fixes the issue. ([OCPBUGS-49860](#))
- Previously, OLM v1 did not wait for certificates to reach a ready state before mounting Operator Controller and catalogd pods. These updates fix the issue. [OCPBUGS-48830](#) and ([OCPBUGS-49418](#))
- Previously, OLM v1 did not apply all of the metadata provided by cluster extension authors in Operator bundles. As a result, OLM v1 did not apply properties, such as update constraints, that were specified in the **metadata/properties.yaml** file. This update fixes the issue. ([OCPBUGS-44808](#))


1.6.21. Operator Controller Manager

- Previously, the **HTTP_PROXY**, **http_proxy**, **HTTPS_PROXY**, **https_proxy**, **NO_PROXY**, and **no_proxy** variables were set on build containers regardless of default proxy settings. With this release, the variables are only added if they are defined in defaults and are not null. ([OCPBUGS-55642](#))
- Previously, an image pull secret generated for the internal Image Registry would not be regenerated until after the embedded credentials had expired. This resulted in a small period of time in which the image pull secrets were invalid. With this release, the image pull secrets are refreshed before the embedded credentials have expired. ([OCPBUGS-50507](#))
- Previously, OLM v1 did not search the **/etc/docker/** directory for the certificates required to mount images. As a result, OLM v1 failed to mount custom certificates. This update fixes the issue. ([OCPBUGS-48795](#))
- Previously, OLM v1 would send an error message during temporary outages that occur during routine cluster maintenance, such as leader election. This update fixes the issue. ([OCPBUGS-48765](#))
- Previously, Operator Lifecycle Manager (OLM) Classic incorrectly reported failures to **Subscription** resources during concurrent attempts to reconcile Operators in the same namespace. When this occurred, Operators failed to install. This update fixes the issue. ([OCPBUGS-48486](#))
- Previously, OLM (Classic) took a snapshot of the catalog source for every installed Operator when it reconciled a subscription. This behavior resulted in high CPU usage. With this update, OLM (Classic) caches catalog sources and limits calls to the gRPC Remote Procedure Calls (gRPC) server to resolve the issue. ([OCPBUGS-48468](#))

1.6.22. Performance Addon Operator

- Previously, if you specified a long string of isolated CPUs in a performance profile, such as **0,1,2, ...,512**, the **tuned**, Machine Config Operator and **rpm-ostree** components failed to process the

string as expected. As a consequence, after you applied the performance profile, the expected kernel arguments were missing. The system failed silently with no reported errors. With this release, the string for isolated CPUs in a performance profile is converted to sequential ranges, such as **0-512**. As a result, the kernel arguments are applied as expected in most scenarios. ([OCPBUGS-45264](#))



NOTE

The issue might still occur with some combinations of input for isolated CPUs in a performance profile, such as a long list of odd numbers **1,3,5,...,511**.

- Previously, the Performance Profile Creator (PPC) failed to build a performance profile for compute nodes that had different core ID numbering (core per socket) for their logical processors and the nodes existed under the same node pool. For example, the PPC failed in a situation for two compute nodes that have logical processors **2** and **18**, where one node groups them as core ID **2** and the other node groups them as core ID **9**.
With this release, PPC no longer fails to create the performance profile because PPC can now build a performance profile for a cluster that has compute nodes that each have different core ID numbering for their logical processors. The PPC now outputs a warning message that indicates to use the generated performance profile with caution, because different core ID numbering might impact system optimization and isolated management of tasks. ([OCPBUGS-44372](#))

1.6.23. Samples Operator

- Previously, the Samples Operator updated the **lastTransitionTime** spec in the **Progressing** condition even when the condition did not change. This made the Operator show as less stable than it was. With this release, the **lastTransitionTime** spec updates only when the **Progressing** condition changes. ([OCPBUGS-54591](#))
- Previously, unsorted image stream names in the **Progressing** condition caused unnecessary updates. This caused excessive user updates and reduced system performance. With this release, the **activeImageStreams** function sorts falling image imports. This action improves Cluster Samples Operator efficiency, reduces unnecessary updates, and enhances overall performance. ([OCPBUGS-54590](#))
- Previously, the Samples Operator established a watch for all cluster Operators, which caused the sync loop of the Samples Operator to run when any Operator changed. With this release, the Samples Operator watches only the Operators that it needs to monitor. ([OCPBUGS-54589](#))

1.6.24. Storage

- Previously, using the **oc adm top pvc** command would not show usage statistics for persistent volume claims (PVCs) for clusters with restricted network configurations, such as clusters with a proxy or clusters in a disconnected environment. With this release, usage statistics can be acquired for clusters in these environments. ([OCPBUGS-54168](#))
- Previously, the VMware vSphere CSI driver Operator entered panic mode if the vCenter address was incorrect. With this release, the issue has been resolved. ([OCPBUGS-43273](#))
- Previously, a Google Cloud Persistent Disk cluster with C3-standard-2, C3-standard-4, N4-standard-2, and N4-standard-4 nodes could erroneously exceed the maximum attachable disk number, which should be 16, which could prevent you from successfully creating or attaching

volumes to your pods. With this release, the maximum is not exceeded, and thus does not interfere with successfully creating or attaching volumes to your pods. ([OCPBUGS-39258](#))

- Previously, when persistent volumes (PVs) are deleted, the Local Storage Operator (LSO) did not reliably recreate the symlinks. With this release, when creating PVs, previously specified symlinks are picked before finding new symlinks. ([OCPBUGS-31059](#))
- Previously, when the Cloud Credential Operator (CCO) did not provide credentials for Container Storage Interface (CSI) driver Operators, the CSI driver Operators remain in **Progressing=true** indefinitely, with a message indicating the **operator is waiting for deployment/unavailable**. With this release, when progressing state lasts 15 minutes or longer, the Operator changes to **Degraded=True**. ([OCPBUGS-24588](#))
- Previously, compute nodes with names that are 53 characters long and when using hostpath Container Storage Interface (CSI) driver cause volume provisioning to fail when using **--enable-node-deployment flag** on external-provisioner. With this release, this issue is resolved and there is no restriction on compute node name lengths. ([OCPBUGS-49805](#))
- Previously, on Azure Red Hat OpenShift when using hosted control planes to create hosted clusters, the Azure Disk Container Storage Interface (CSI) driver would not provision volumes successfully. With this release, this issue has been resolved, and the Azure Disk CSI driver can provision volumes successfully. ([OCPBUGS-46575](#))
- Previously, Internet Small Computer System Interface (iSCSI) and Fibre Channel devices that were attached to a multipath device did not resolve correctly when these devices were partitioned. With this release, a fix ensures that partitioned multipath storage devices can now correctly resolve. ([OCPBUGS-46038](#))
- Previously, when creating a hosted cluster with specified labels, the AWS EBS driver, Driver Operator, snapshot controller, and snapshot webhook pods do not get these specified labels propagated to them. With this release, the specified labels are propagated. ([OCPBUGS-45073](#))
- Previously, the Manila Container Storage Interface (CSI) driver had services running on unintended hosts. This occurred because the Manila CSI driver uses a single binary for both the controller and node (worker) services. With this release, the CSI driver controller pods run only controller services, and CSI driver node pods run only node services. ([OCPBUGS-54447](#))
- Previously, the Container Storage Interface (CSI) Operator was issuing warnings in the log about missing items that would become fatal in the future. With this release, the warnings are no longer issued. ([OCPBUGS-44374](#))
- Previously, the VMWare vSphere CSI driver Operator would panic if the vCenter address was incorrect. With this release, the issue has been resolved. ([OCPBUGS-43273](#))

1.6.25. Red Hat Enterprise Linux CoreOS (RHCOS)

- Previously, the **GRUB** bootloader was not automatically updated on RHCOS nodes. As a result, when nodes were created on RHEL 8 and were subsequently updated to RHEL, **GRUB** could not load the kernel as it uses a format that is not supported by older **GRUB** versions. With this release, a **GRUB** bootloader update is forced on nodes during updates to OpenShift Container Platform 4.18 so that the issue does not occur on OpenShift Container Platform 4.19. ([OCPBUGS-55144](#))

1.7. TECHNOLOGY PREVIEW FEATURES STATUS

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 following tables, features are marked with the following statuses:

- *Not Available*
- *Technology Preview*
- *General Availability*
- *Deprecated*
- *Removed*

1.7.1. Authentication and authorization Technology Preview features

Table 1.18. Authentication and authorization Technology Preview tracker

Feature	4.17	4.18	4.19
Pod security admission restricted enforcement	Technology Preview	Technology Preview	Technology Preview
Direct authentication with an external OIDC identity provider	Not Available	Not Available	Technology Preview

1.7.2. Edge computing Technology Preview features

Table 1.19. Edge computing Technology Preview tracker

Feature	4.17	4.18	4.19
Accelerated provisioning of GitOps ZTP	Technology Preview	Technology Preview	Technology Preview
Enabling disk encryption with TPM and PCR protection	Technology Preview	Technology Preview	Technology Preview
Configuring a local arbiter node	Not Available	Not Available	Technology Preview

1.7.3. Extensions Technology Preview features

Table 1.20. Extensions Technology Preview tracker

Feature	4.17	4.18	4.19
Operator Lifecycle Manager (OLM) v1	Technology Preview	General Availability	General Availability
OLM v1 runtime validation of container images using sigstore signatures	Not Available	Technology Preview	Technology Preview
OLM v1 permissions preflight check for cluster extensions	Not Available	Not Available	Technology Preview
OLM v1 deploying a cluster extension in a specified namespace	Not Available	Not Available	Technology Preview

1.7.4. Installation Technology Preview features

Table 1.21. Installation Technology Preview tracker

Feature	4.17	4.18	4.19
Adding kernel modules to nodes with kvc	Technology Preview	Technology Preview	Technology Preview
Enabling NIC partitioning for SR-IOV devices	General Availability	General Availability	General Availability
User-defined labels and tags for Google Cloud	General Availability	General Availability	General Availability
Installing a cluster on Alibaba Cloud by using Assisted Installer	Technology Preview	Technology Preview	Technology Preview
Installing a cluster on Microsoft Azure with confidential VMs	Not Available	Technology Preview	General Availability
Mount shared entitlements in BuildConfigs in RHEL	Technology Preview	Technology Preview	Technology Preview
OpenShift zones support for vSphere host groups	Not Available	Not Available	Technology Preview
Selectable Cluster Inventory	Technology Preview	Technology Preview	Technology Preview
Installing a cluster on Google Cloud using the Cluster API implementation	General Availability	General Availability	General Availability
Enabling a user-provisioned DNS on Google Cloud	Not Available	Not Available	Technology Preview

Feature	4.17	4.18	4.19
Installing a cluster on VMware vSphere with multiple network interface controllers	Not Available	Technology Preview	Technology Preview
Using bare metal as a service	Not Available	Not Available	Technology Preview

1.7.5. Machine Config Operator Technology Preview features

Table 1.22. Machine Config Operator Technology Preview tracker

Feature	4.17	4.18	4.19
Improved MCO state reporting (oc get machineconfignode)	Technology Preview	Technology Preview	General Availability
Image mode for OpenShift/On-cluster RHCOS image layering	Technology Preview	General Availability ¹	General Availability
Pinned Image Sets	Technology Preview	Technology Preview	General Availability ²

1. This feature is GA starting in OpenShift Container Platform 4.18.20. Earlier 4.18.x versions remain in Technology Preview.
2. This feature is GA starting in OpenShift Container Platform 4.19.12. Earlier 4.19.x versions remain in Technology Preview.

1.7.6. Machine management Technology Preview features

Table 1.23. Machine management Technology Preview tracker

Feature	4.17	4.18	4.19
Managing machines with the Cluster API for Amazon Web Services	Technology Preview	Technology Preview	Technology Preview
Managing machines with the Cluster API for Google Cloud	Technology Preview	Technology Preview	Technology Preview
Managing machines with the Cluster API for IBM Power® Virtual Server	Technology Preview	Technology Preview	Technology Preview

Feature	4.17	4.18	4.19
Managing machines with the Cluster API for Microsoft Azure	Not Available	Technology Preview	Technology Preview
Managing machines with the Cluster API for RHOSP	Technology Preview	Technology Preview	Technology Preview
Managing machines with the Cluster API for VMware vSphere	Technology Preview	Technology Preview	Technology Preview
Managing machines with the Cluster API for bare metal	Not Available	Not Available	Technology Preview
Cloud controller manager for IBM Power® Virtual Server	Technology Preview	Technology Preview	Technology Preview
Adding multiple subnets to an existing VMware vSphere cluster by using compute machine sets	Not Available	Technology Preview	Technology Preview
Configuring Trusted Launch for Microsoft Azure virtual machines by using machine sets	Technology Preview	Technology Preview	General Availability
Configuring Azure confidential virtual machines by using machine sets	Technology Preview	Technology Preview	General Availability

1.7.7. Monitoring Technology Preview features

Table 1.24. Monitoring Technology Preview tracker

Feature	4.17	4.18	4.19
Metrics Collection Profiles	Technology Preview	Technology Preview	General Availability

1.7.8. Multi-Architecture Technology Preview features

Table 1.25. Multi-Architecture Technology Preview tracker

Feature	4.17	4.18	4.19
kdump on arm64 architecture	Technology Preview	Technology Preview	Technology Preview
kdump on s390x architecture	Technology Preview	Technology Preview	Technology Preview

Feature	4.17	4.18	4.19
kdump on ppc64le architecture	Technology Preview	Technology Preview	Technology Preview
Support for configuring the image stream import mode behavior	Not Available	Technology Preview	Technology Preview

1.7.9. Networking Technology Preview features

Table 1.26. Networking Technology Preview tracker

Feature	4.17	4.18	4.19
eBPF manager Operator	Technology Preview	Technology Preview	Technology Preview
Advertise using L2 mode the MetalLB service from a subset of nodes, using a specific pool of IP addresses	Technology Preview	Technology Preview	Technology Preview
Updating the interface-specific safe sysctls list	Technology Preview	Technology Preview	Technology Preview
Egress service custom resource	Technology Preview	Technology Preview	Technology Preview
VRF specification in BGPPeer custom resource	Technology Preview	Technology Preview	Technology Preview
VRF specification in NodeNetworkConfigurationPolicy custom resource	Technology Preview	Technology Preview	General Availability
Host network settings for SR-IOV VFs	General Availability	General Availability	General Availability
Integration of MetalLB and FRR-K8s	General Availability	General Availability	General Availability
Automatic leap seconds handling for PTP grandmaster clocks	General Availability	General Availability	General Availability
PTP events REST API v2	General Availability	General Availability	General Availability
OVN-Kubernetes customized br-ex bridge on bare metal	General Availability	General Availability	General Availability

Feature	4.17	4.18	4.19
OVN-Kubernetes customized br-ex bridge on vSphere and RHOSP	Technology Preview	Technology Preview	Technology Preview
Live migration to OVN-Kubernetes from OpenShift SDN	General Availability	Not Available	Not Available
User-defined network segmentation	Technology Preview	General Availability	General Availability
Dynamic configuration manager	Not Available	Technology Preview	Technology Preview
SR-IOV Network Operator support for Intel C741 Emmitsburg Chipset	Not Available	Technology Preview	Technology Preview
SR-IOV Network Operator support on ARM architecture	Not Available	General Availability	General Availability
Gateway API and Istio for Ingress management	Not Available	Technology Preview	General Availability
Dual-port NIC for PTP ordinary clock	Not Available	Not Available	Technology Preview
DPU Operator	Not Available	Not Available	Technology Preview
Fast IPAM for the Whereabouts IPAM CNI plugin	Not Available	Not Available	Technology Preview
Unnumbered BGP peering	Not Available	Not Available	Technology Preview

1.7.10. Node Technology Preview features

Table 1.27. Nodes Technology Preview tracker

Feature	4.17	4.18	4.19
MaxUnavailableStatefulSet featureset	Technology Preview	Technology Preview	Technology Preview
sigstore support	Technology Preview	Technology Preview	Technology Preview

1.7.11. OpenShift CLI (oc) Technology Preview features

Table 1.28. OpenShift CLI (oc) Technology Preview tracker

Feature	4.17	4.18	4.19
oc-mirror plugin v2	Technology Preview	General Availability	General Availability
oc-mirror plugin v2 enclave support	Technology Preview	General Availability	General Availability
oc-mirror plugin v2 delete functionality	Technology Preview	General Availability	General Availability

1.7.12. Operator lifecycle and development Technology Preview features

Table 1.29. Operator lifecycle and development Technology Preview tracker

Feature	4.17	4.18	4.19
Operator Lifecycle Manager (OLM) v1	Technology Preview	General Availability	General Availability
Scaffolding tools for Hybrid Helm-based Operator projects	Deprecated	Removed	Removed
Scaffolding tools for Java-based Operator projects	Deprecated	Removed	Removed

1.7.13. Red Hat OpenStack Platform (RHOSP) Technology Preview features

Table 1.30. RHOSP Technology Preview tracker

Feature	4.17	4.18	4.19
RHOSP integration into the Cluster CAPI Operator	Technology Preview	Technology Preview	Technology Preview
Control plane with rootVolumes and etcd on local disk	General Availability	General Availability	General Availability
Hosted control planes on RHOSP 17.1	Not Available	Not Available	Technology Preview

1.7.14. Scalability and performance Technology Preview features

Table 1.31. Scalability and performance Technology Preview tracker

Feature	4.17	4.18	4.19
---------	------	------	------

Feature	4.17	4.18	4.19
factory-precaching-cli tool	Technology Preview	Technology Preview	Technology Preview
Hyperthreading-aware CPU manager policy	Technology Preview	Technology Preview	Technology Preview
Mount namespace encapsulation	Technology Preview	Technology Preview	Technology Preview
Node Observability Operator	Technology Preview	Technology Preview	Technology Preview
Increasing the etcd database size	Technology Preview	Technology Preview	Technology Preview
Using RHACM PolicyGenerator resources to manage GitOps ZTP cluster policies	Technology Preview	Technology Preview	General Availability
NUMA-aware scheduling supported on hosted control planes	Not Available	Not Available	Technology Preview

1.7.15. Storage Technology Preview features

Table 1.32. Storage Technology Preview tracker

Feature	4.17	4.18	4.19
AWS EFS storage CSI usage metrics	General Availability	General Availability	General Availability
Automatic device discovery and provisioning with Local Storage Operator	Technology Preview	Technology Preview	Technology Preview
Azure File CSI snapshot support	Technology Preview	Technology Preview	Technology Preview
Azure File cross-subscription support	Not Available	Not Available	General Availability
Shared Resources CSI Driver in OpenShift Builds	Technology Preview	Technology Preview	Technology Preview
Secrets Store CSI Driver Operator	Technology Preview	General Availability	General Availability

Feature	4.17	4.18	4.19
CIFS/SMB CSI Driver Operator	Technology Preview	General Availability	General Availability
VMware vSphere multiple vCenter support	Technology Preview	Technology Preview	General Availability
Disabling/enabling storage on vSphere	Technology Preview	Technology Preview	General Availability
Increasing max number of volumes per node for vSphere	Not Available	Not Available	Technology Preview
RWX/RWO SELinux Mount	Developer Preview	Developer Preview	Developer Preview
Migrating CNS Volumes Between Datastores	Developer Preview	Developer Preview	General Availability
CSI volume group snapshots	Not Available	Technology Preview	Technology Preview
GCP PD supports C3/N4 instance types and hyperdisk-balanced disks	Not Available	General Availability	General Availability
GCP Filestore supports Workload Identity	General Availability	General Availability	General Availability
OpenStack Manila support for CSI resize	Not Available	General Availability	General Availability
Volume Attribute Classes	Not Available	Not Available	Technology Preview

1.7.16. Web console Technology Preview features

Table 1.33. Web console Technology Preview tracker

Feature	4.17	4.18	4.19
Red Hat OpenShift Lightspeed in the OpenShift Container Platform web console	Technology Preview	Technology Preview	Technology Preview

1.8. KNOWN ISSUES

- In OpenShift Container Platform 4.19, clusters using IPsec for network encryption might experience intermittent loss of pod-to-pod connectivity. This prevents some pods on certain nodes from reaching services on other nodes, resulting in connection timeouts. Internal testing

could not reproduce this issue on clusters with 120 nodes or less. There is no workaround for this issue. ([OCPBUGS-55453](#))

- OpenShift Container Platform clusters that are installed on AWS in the Mexico Central region, **mx-central-1**, cannot be destroyed. ([OCPBUGS-56020](#))
- When installing a cluster on Azure, if you set any of the **compute.platform.azure.identity.type**, **controlplane.platform.azure.identity.type**, or **platform.azure.defaultMachinePlatform.identity.type** field values to **None**, your cluster is unable to pull images from the Azure Container Registry. You can avoid this issue by either providing a user-assigned identity, or by leaving the identity field blank. In both cases, the installation program generates a user-assigned identity. ([OCPBUGS-56008](#))
- Previously, the kubelet would not account for probes that ran in the **syncPod** method, which periodically checks the state of a pod and does a readiness probe outside of the normal probe period. With this release, a bug is fixed for when the kubelet incorrectly calculates **readinessProbe** periods. However, pod authors might see that the readiness latency of pods configured with readiness probes might increase. This behavior is more accurate to the configured probe. For more information, see ([OCPBUGS-50522](#))
- When installing a cluster on AWS, if you do not configure AWS credentials before running any **openshift-install create** command, the installation program fails. ([OCPBUGS-56658](#))
- The **must-gather** tool does not collect IPsec information for a cluster that was upgraded from OpenShift Container Platform 4.14. This issue occurs because the **ipsecConfig** configuration in the **networks.operator.openshift.io cluster** CR has an empty construct, **{}**. The empty construct is passed to the upgraded version of OpenShift Container Platform. As a workaround for this issue, run the following command with the following **ipsecConfig** configuration in the Cluster Network Operator (CNO) CR:

```
$ oc patch networks.operator.openshift.io cluster --type=merge -p \
  {
  "spec":{
    "defaultNetwork":{
      "ovnKubernetesConfig":{
        "ipsecConfig":{
          "mode":"Full"
        }
      }
    }
  }
}
```

After you run the command, the CNO collects **must-gather** logs that you can inspect.

([OCPBUGS-52367](#))

- There is a known issue with Gateway API and Amazon Web Services (AWS), Google Cloud, and Microsoft Azure private clusters. The load balancer that is provisioned for a gateway is always configured to be external, which can cause errors or unexpected behavior:
 - In an AWS private cluster, the load balancer becomes stuck in the **pending** state and reports the error: **Error syncing load balancer: failed to ensure load balancer: could not find any suitable subnets for creating the ELB.**
 - In Google Cloud and Azure private clusters, the load balancer is provisioned with an external IP address, when it should not have an external IP address.

There is no supported workaround for this issue. ([OCPBUGS-57440](#))

- The **istiod** component, which supports Gateway API, previously created an **istio-ca-root-cert** ConfigMap in all namespaces watched by OpenShift Service Mesh (OSSM). A fix is included in OSSM 3.0.1 (part of OpenShift Container Platform 4.19) that limits this behavior, so ConfigMaps are only created in namespaces that contain a Gateway API resource. However, this fix does not automatically clean up existing ConfigMaps from namespaces that do not contain a Gateway. As a result, clusters might still contain numerous unnecessary **istio-ca-root-cert** ConfigMaps after an upgrade. These extra ConfigMaps are benign and can be safely ignored, or they can be manually deleted from namespaces that do not host a Gateway. ([OCPBUGS-43093](#))
- When the same hostname is specified in two or more different **Gateway** resources, the system does not report a conflict. The upstream Gateway API specification treats each **Gateway** as a separate independent proxy. This is the intended behavior of the upstream Gateway API specification. The cluster will not prevent this configuration. The routing of traffic to a specific Gateway is determined by external configuration, such as which Gateway's service IP address is used in the DNS record for the hostname. ([NE-2180](#))
- Gateway API is not supported on on-premise platforms, including Bare Metal, vSphere, Nutanix, KubeVirt, and clusters using the **None** platform type. The Gateway API implementation automatically creates a service of **type: LoadBalancer**. On-premise platforms do not have a default cloud controller manager or other provisioner to assign an external IP address to this service. As a result, the **Gateway** resource will never be marked as "accepted" and will remain inaccessible from outside the cluster. ([OCPBUGS-37385](#))
- If the **appProtocol** field is not explicitly defined in a Kubernetes **Service**, the Gateway API implementation infers the network protocol from the **Service** port name. For example, a port named **https-web** is interpreted as using the HTTPS protocol. If the port name does not match the actual protocol used by the application, this inference can cause unexpected traffic behavior or connection failures. To avoid this issue, explicitly define the **appProtocol** field in your **Service** resources or ensure that port names strictly match the protocol in use. ([OSSM-7231](#))
- The reverse proxy implementation for Gateway API (based on Istio and Envoy) currently demonstrates lower performance compared to the default HAProxy-based Ingress Controller. HAProxy has been extensively tuned for OpenShift Container Platform over many releases. Consequently, users might observe higher latency or lower throughput when using Gateway API compared to standard OpenShift Routes. Performance optimizations for the Gateway API implementation are ongoing. ([OSSM-6842](#))
- OpenShift Service Mesh (OSSM) 3.x, which provides the Gateway API implementation, cannot coexist with OSSM 2.x on the same cluster, except during a supported migration process. Installing both versions simultaneously for standard operations is not supported and might lead to resource conflicts; for example, the Ingress Operator might report a **degraded** status. Additionally, the Network Observability Operator (NIO) no longer supports OSSM 2.x. It is recommended that users use OSSM 3.x for ongoing support and compatibility. Users must choose a single Service Mesh version for their cluster or follow the official migration procedure when moving from 2.x to 3.x. ([OSSM-5407](#))
- The Gateway API implementation currently does not support the automatic idling of services. Unlike standard OpenShift Routes, which can idle services to conserve resources when they are not in use, workloads managed by Gateway API do not automatically scale to zero or enter an idle state based on traffic inactivity.
- The Gateway API implementation currently lacks dedicated integration with the OpenShift Container Platform web console and the OpenShift CLI (**oc**). Consequently, Gateway API resources, such as **Gateway**, **GatewayClass**, and **HTTPRoute**, do not appear in specific console

dashboards or views. To manage these resources, users must utilize standard **oc** commands (for example, **oc get gateway** or **oc edit httproute**) and apply configurations using YAML manifests.

- The Gateway API implementation currently does not provide a mechanism for users to configure advanced Istio or Envoy options. Specifically, capabilities such as configuring access logs or tuning low-level proxy settings are not exposed. Users must rely on the default configuration provided by the Gateway API controller.
- The Gateway API implementation currently does not support Cluster User-Defined Networks (CUDNs). Gateway API resources can only be used with the default cluster network and cannot be attached to user-defined secondary networks.
- The **HTTPRoute** resource in the Gateway API implementation currently does not support **Passthrough** or **Re-encrypt** TLS termination strategies. Unlike standard OpenShift Routes, which offer flexible termination options, **HTTPRoute** currently supports only Edge termination (terminating TLS at the gateway). Users requiring **Passthrough** or **Re-encrypt** functionality cannot currently achieve this using **HTTPRoute** resources.
- In the event of a crash, the **mlx5_core** NIC driver causes an out-of-memory issue and **kdump** does not save the **vmcore** file in **/var/crash**. To save the **vmcore** file, use the **crashkernel** setting to reserve 1024 MB of memory for the **kdump** kernel. ([OCPBUGS-54520](#), [RHEL-90663](#))
- There is a known latency issue on 4th Gen Intel Xeon processors. ([OCPBUGS-42495](#))
- Currently, pods that use a **guaranteed** QoS class and request whole CPUs might not restart automatically after a node reboot or kubelet restart. The issue might occur in nodes configured with a static CPU Manager policy and using the **full-ppcus-only** specification, and when most or all CPUs on the node are already allocated by such workloads. As a workaround, manually delete and re-create the affected pods. ([OCPBUGS-43280](#))
- Currently, when a **irqbalance** service runs on a specific AArch64 machine, a buffer overflow issue might cause the service to crash. As a consequence, latency sensitive workloads might be affected by unmanaged interrupts that are not properly distributed across CPUs, leading to performance degradation. There is currently no workaround for this issue. ([RHEL-89986](#))
- Currently, on clusters with SR-IOV network virtual functions configured, a race condition might occur between system services responsible for network device renaming and the TuneD service managed by the Node Tuning Operator. As a consequence, the TuneD profile might become degraded after the node restarts, leading to performance degradation. As a workaround, restart the TuneD pod to restore the profile state. ([OCPBUGS-41934](#))
- NFS volumes exported from VMware vSAN Files cannot be mounted by clusters running OpenShift Container Platform 4.19 due to RHEL-83435. To avoid this issue, ensure that you are running VMware ESXi and vSAN at the latest patch versions of 8.0 P05, or later. ([OCPBUGS-55978](#))

1.9. ASYNCHRONOUS ERRATA UPDATES

Security, bug fix, and enhancement updates for OpenShift Container Platform 4.19 are released as asynchronous errata through the Red Hat Network. All OpenShift Container Platform 4.19 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 4.19. Versioned asynchronous releases, for example with the form OpenShift Container Platform 4.19.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-2026:4434 - OpenShift Container Platform 4.19.26 fixed issues and security update

Issued: 18 March 2026

OpenShift Container Platform release 4.19.26 is now available. The list of fixed issues that are included in the update is documented in the [RHSA-2026:4434](#) advisory. The RPM packages that are included in the update are provided by the [RHBA-2026:4427](#) advisory.

Space precluded documenting all of the container images for this release in the advisory.

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

```
$ oc adm release info 4.19.26 --pullspecs
```

1.9.1.1. Fixed issues

- Before this update, the daemon set failed to update the **kubeconfig** token due to an incorrect file path reference. As a consequence, user tokens expired and caused service disruptions. With this release, the daemon set is updated to correctly renew the **kubeconfig** token, preventing future service disruptions. ([OCPBUGS-67208](#))
- Before this update, a loadbalancer service creation without an IP address during a hosted control planes cluster upgrade from OpenShift Container Platform 4.18 to OpenShift Container Platform 4.19 blocked the upgrade process. As a consequence, a loadbalancer service creation blocked the cluster upgrade, causing an upgrade failure. With this release, a loadbalancer service creation that was blocked during upgrade is resolved, enabling a smooth upgrade process. ([OCPBUGS-75932](#))
- Before this update, unauthorized access to the **8445** port on VMware vSphere for the **vmware-vmisphere-csi-driver-operator-metrics** metric occurred. As a consequence, users encountered a **500** error when they accessed metric endpoints on vSphere Infrastructure as a Service (IAAS).

With this release, the fix for the **8445** port access issue in VMware **vsphere-csi-driver-operator-metrics** is included. As a result, users do not get a **500** error when accessing the **vmware-vsphere-csi-driver-operator-metrics** endpoint. ([OCPBUGS-77007](#))

- Before this update, duplicate default routes in the Open Virtual Network (OVN) Northbound database caused compute nodes to get stuck in the **NotReady** state. As a consequence, 259 out of 359 nodes consistently failed to join the network. With this release, the duplicate default routes in the OVN Northbound database are fixed, allowing compute nodes to join the network. As a result, the number of stuck nodes decreases, and cluster functionality improves. ([OCPBUGS-77082](#))
- Before this update, a **V4SubnetAllocationThresholdExceeded** alert was triggered due to insufficient subnet allocation with a host prefix of **23** in a cluster with 11 nodes. As a consequence, the **V4SubnetAllocationThresholdExceeded** alert triggered with only 11 nodes, causing potential network instability. With this release, cluster network configuration is adjusted to increase the **V4SubnetAllocationThreshold** alert. As a result, **V4SubnetAllocationThresholdExceeded** alerts do not trigger for clusters with 11 nodes and the specified **clusterNetwork** configuration. ([OCPBUGS-77096](#))
- Before OpenShift Container Platform 4.19, the **ingress-to-route** metrics code misidentified managed ingresses as **unmanaged** in the same namespace, causing **UnmanagedRoutes** alerts. With this release, a metrics code fix resolves the unmanaged routes alert issue. ([OCPBUGS-77172](#))
- Before this update, a missing Kubernetes secret referenced in a custom **AlertManagerConfig** caused a **Prometheus Operator configuration generation failure**. As a consequence, user-defined **AlertManagerConfig** referencing non-existent Kubernetes Secret caused monitoring Operator degradation. With this release, the **AlertManagerConfig** issue is resolved, and non-existent secrets do not impact **AlertManagerConfig** reconciliation. ([OCPBUGS-77269](#))
- Before this update, duplicate **PrometheusRule** objects were created due to the controller updating rules without replacing existing ones. This caused inconsistency in rule application, affecting monitoring performance. With this release, duplicate **PrometheusRule** objects are removed, ensuring that updated rules replace existing ones. As a result, duplicate **PrometheusRules** are resolved, ensuring consistent rule versions across clusters. ([OCPBUGS-77273](#))
- Before this update, frequent updates by **hypershift-controlplane-manager** caused **ignition-server** pods to frequently restart, causing service disruptions. With this release, the frequency of **ignition-server** pod restarts is resolved by reducing **hypershift-controlplane-manager** updates. As a result, service stability for users is improved. ([OCPBUGS-77367](#))
- Before this update, Open Virtual Network (OVN) database updates were not propagating to all nodes, and caused stale pod IP addresses to receive User Datagram Protocol (UDP) traffic. As a consequence, UDP traffic was routed to stale pod IP addresses, and caused connections to fail. With this release, UDP traffic is correctly routed to the updated pod endpoints, improving service availability and reducing errors. ([OCPBUGS-77442](#))
- Before this update, image mirroring failed due to the **use-sigstore-attachments** option set to **false** in the **default.yaml** file, causing missing signatures. As a consequence, users experienced a mirror failure and **image not found** errors. With this release, sigstore attachments for the Docker registry are disabled to correct the mirroring failure. As a result, images mirror successfully without errors due to missing signatures. ([OCPBUGS-77663](#))

- Before this update, frequent updates in deployment and the container environment caused **ignition-server** pods to frequently restart. As a consequence, service interruptions occurred. With this release, the **ignition-server** pod restarts are fixed, and no more updates are applied. As a result, service stability is improved. ([OCPBUGS-77884](#))

1.9.1.2. Updating

To update an OpenShift Container Platform 4.19 cluster to this latest release, see [Updating a cluster using the CLI](#).

1.9.2. RHBA-2026:3394 - OpenShift Container Platform 4.19.25 fixed issues and security update

Issued: 04 March 2026

OpenShift Container Platform release 4.19.25 is now available. The list of fixed issues that are included in the update is documented in the [RHBA-2026:3394](#) advisory. The RPM packages that are included in the update are provided by the [RHSA-2026:3391](#) advisory.

Space precluded documenting all of the container images for this release in the advisory.

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

```
$ oc adm release info 4.19.25 --pullspecs
```

1.9.2.1. Fixed issues

The following issues are fixed for this release:

- Before this update, incorrect subnet was association with automatic v-net subnet creation. As a consequence, previously configured subnets were affected, leading to Azure installation failure. In this release, correct subnet assignment is done in pre-configured v-nets. As a result, the Azure installation issue is resolved. ([OCPBUGS-73781](#))
- Before this update, CoreOS networking timeout occurred during automation of agent-based OpenShift Container Platform installations on VLAN interfaces in bare metal environments. As a consequence, CoreOS nodes failed to automatically configure networking, preventing installation on bare metal systems. With this release, the **pre-network-manager-config** service timeout is increased, resolving the networking issue in OpenShift Container Platform installations. As a result, during agent-based OpenShift Container Platform installation, CoreOS nodes now automatically configure and start networking, eliminating the need for manual intervention. ([OCPBUGS-74940](#))
- Before this update, the **collect-profiles** job was affected by maintenance difficulties, causing confusion for customers and support engineers during troubleshooting. With this release, the **collect-profiles** job is removed, improving user experience and reducing maintenance effort. ([OCPBUGS-76587](#))

1.9.2.2. Updating

To update an OpenShift Container Platform 4.19 cluster to this latest release, see [Updating a cluster using the CLI](#).

1.9.3. RHSA-2026:2651 - OpenShift Container Platform 4.19.24 fixed issues and security update

Issued: 18 February 2026

OpenShift Container Platform release 4.19.24 is now available. The list of fixed issues that are included in the update is documented in the [RHSA-2026:2651](#) advisory. The RPM packages that are included in the update are provided by the [RHBA-2026:2632](#) advisory.

Space precluded documenting all of the container images for this release in the advisory.

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

```
$ oc adm release info 4.19.24 --pullspecs
```

1.9.3.1. Fixed issues

The following issues are fixed for this release:

- Before this update, the **oauth** pods in the **openshift-authentication** namespaces could get stuck while rolling out changes when one of the nodes these pods were running on was not ready or not available. This was causing authentication to completely stop until the blocked rollout concludes. With this release, the pods can proceed with the rolling update even with an unhealthy node that is down or unavailable ([OCPBUGS-65602](#))
- Before this update, the Systemd use in the container entry point prevented the **ConfigMap** mount from working correctly, which caused broken file permissions. As a consequence, users could not access config files due to the broken file permissions in the containers. With this release, the Systemd entry point issue is resolved, which allows the config map mount with the correct file permissions. As a result, the file permissions are correct for the containers. ([OCPBUGS-69671](#))
- Before this update, clusters using multi-architecture release payload images that had the Cluster Version object's **spec.desiredUpdate.architecture** field set to the **Multi** value experienced a significant issue where the Cluster Version resource would not be populated with update recommendations for later releases. With this release, a comparison of the incorrect values in the code has been fixed, ensuring that the bug no longer presents itself and update recommendations now appear as expected. ([OCPBUGS-70181](#))
- Before this update, the catalog sync triggered high I/O on masters where etcd ran. As a consequence, an etcd leader election was triggered. The triggered election reset TTL counters on keys and prevented etcd events from being cleared. With this release, the default catalog polling interval is increased from ten minutes to for hours. As a result, the load on catalog sources is reduced. ([OCPBUGS-73882](#))
- Before this update, clusters using **localnet** IPAM-less topologies had the Label Switched Paths (LSPs) removed during the update process. This caused connectivity disruptions and potential data loss for traffic relying on those specific network paths. With this release, the update logic has been adjusted to ensure that LSP configurations are preserved and remain intact throughout the update. This results in consistent network stability and path persistence without the risk of accidental removal. ([OCPBUGS-74269](#))
- Before this update, when navigating to the Software Catalog page with the **Devfiles** category disabled, the page would show an Devfile related error in a disconnected cluster. With this release, this error is no longer shown. ([OCPBUGS-74341](#))

- Before this update, the deletion sequence attempted to remove machine instances while their primary network interface cards (NICs) were still attached, which conflicted with environments enforcing strict policies against deleting a primary NIC on an active server. This created a circular dependency that caused machine deletions to fail or hang indefinitely. With this release, the order of deletion for machines and ports has been corrected to ensure that resources are decommissioned in the proper sequence. As a result, machines can now be successfully deleted even in restricted environments, ensuring reliable resource cleanup and adherence to local security policies. ([OCPBUGS-74586](#))
- Before this update, static pod manifests for control plane components lacked explicit priority values, causing the **kubelet** to treat them as low-priority during graceful node shutdowns. On single-node OpenShift, this resulted in control plane components terminating immediately, leading to extended shutdown times, forced terminations, and the potential for storage corruption. With this release, explicit priority values have been added to the control plane static pod manifests, ensuring these components remain available during a graceful shutdown until all workload pods have terminated, preventing storage corruption and ensuring a cleaner shutdown process. ([OCPBUGS-74628](#))
- Before this update, when installing a cluster with the Agent-based Installer and configuring **installConfig** overrides such as enabling FIPS mode, a failure to apply these overrides would not cause the application of other overrides to be retried. As a result, any error in the overrides caused all other overrides to be ignored. With this release this application of overrides is checked for completion, so that either all of the overrides are applied successfully or the installation does not proceed. ([OCPBUGS-74630](#))
- Before this update, when installing a cluster on Google Cloud in the **us-south1** and **us-central1** regions, if you did not specify standard zones in the **install-config.yaml** file, the installation program would select an AI zone as the default zone. As a result, the AI zone would likely have insufficient resources for a cluster and the installation would fail. With this release, AI zones are excluded from selection by the installation program, and installations no longer fail if you do not specify a zone when installing a cluster in the **us-south1** and **us-central1** regions. ([OCPBUGS-74675](#))

1.9.3.2. Updating

To update an OpenShift Container Platform 4.19 cluster to this latest release, see [Updating a cluster using the CLI](#).

1.9.4. RHSA-2026:1552 - OpenShift Container Platform 4.19.23 bug fix advisory

Issued: 04 February 2026

OpenShift Container Platform release 4.19.23 is now available. The list of bug fixes that are included in the update is documented in the [RHSA-2026:1552](#) advisory. The RPM packages that are included in the update are provided by the [RHBA-2026:1538](#) advisory.

Space precluded documenting all of the container images for this release in the advisory.

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

```
$ oc adm release info 4.19.23 --pullspecs
```

1.9.5. Enhancements

The following enhancement is included with this release:

- With this update, the **kubernetes-nmstate** Operator now automatically retries failed node network configuration policy (NNCP) applications with exponential backoff (up to five retries, with one- to thirty-second delays). This enhancement reduces false failures from transient network issues and improves Operator reliability. ([OCPBUGS-55353](#))

1.9.6. Fixed issues

The following issues are fixed with this release:

- Before this update, the {ocp} Web console relied on the full name for user identification, which created potential ambiguity when users shared similar names. With this release, the display logic is updated to prioritize unique identifiers. As a result, the console now shows the user name instead of the full name as the display name. ([OCPBUGS-70335](#))
- Before this update, the CA certificate reference to **HelmChartRepository** caused chart installation failure. As a consequence, Helm chart installation failed because it could not find the chart. With this release, the issue has been fixed. As a result, the CA certificate configuration no longer breaks the Helm chart installation. ([OCPBUGS-72399](#))
- Before this update, Amazon Web Services (AWS) APIs returned inconsistent results regarding the existence of a **Machine**. The safeguards designed to handle this inconsistency checked the stored instance ID in the incorrect location. As a consequence, during AWS API instability, virtual machines (VMs) leaked and attempted to join the cluster indefinitely. With this release, the system uses the correct provider ID for consistency checks. As a result, if an instance does not appear within 20 seconds, the machine status changes to **Failed** to prevent instance leaks. ([OCPBUGS-73729](#))
- Before this update, cluster namespaces with a "nodes" suffix in their name caused the Performance Profile Creator (PPC) to fail when mistaking incorrect **namespace** directories with the must-gather **nodes** directories in the must-gather processed by the PPC. With this release, the PPC now correctly excludes the **namespaces** directory when processing the must-gather data to create a suggested **PerformanceProfile**. ([OCPBUGS-73797](#))
- Before this update, the **Cluster Autoscaler** API object did not have role-based access control (RBAC) permissions for the **VolumeAttachment** API object which caused a problem on upgrade because of changes in the **Cluster Autoscaler** object that require those permissions. With this release, the **Cluster Autoscaler** is deployed with proper permissions for **VolumeAttachment** objects. ([OCPBUGS-73954](#))

1.9.7. Updating

To update an OpenShift Container Platform 4.19 cluster to this latest release, see [Updating a cluster using the CLI](#).

1.9.8. RHBA-2026:0682 - OpenShift Container Platform 4.19.22 bug fix advisory

Issued: 21 January 2026

OpenShift Container Platform release 4.19.22 is now available. The list of bug fixes that are included in the update is documented in the [RHBA-2026:0682](#) advisory. The RPM packages that are included in the update are provided by the [RHBA-2026:0673](#) advisory.

Space precluded documenting all of the container images for this release in the advisory.

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

```
$ oc adm release info 4.19.22 --pullspecs
```

1.9.9. Fixed issues

The following issues are fixed with this release:

- Before this update, the Route API validation allowed routes to **spec.paths** containing spaces or "#" characters, which the OpenShift Container Platform router failed to filter when generating the HAProxy configuration. This led to a significant consequence where users could create routes with invalid paths that prevented HAProxy from starting. With this release, the Route API validation now prohibits the use of space and "#" characters in **spec.path** values, ensuring that users can no longer create routes with these characters. ([OCPBUGS-61662](#))
- Before this release, resizing or cloning small Google Cloud Hyperdisk volumes would fail due to an Input/Output Operations Per Second (IOPS) validation error from the Google Cloud API. This occurred because the container storage interface (CSI) driver did not automatically adjust the provisioned IOPS to meet the minimum requirements of the new volume size. With this release, the driver has been updated to correctly calculate and provide the required IOPS during volume expansion operations. Users can now successfully resize and clone these smaller Google Cloud Hyperdisk volumes. ([OCPBUGS-62116](#))
- Before this update, rounding errors occurred when the Vertical Pod Autoscaler (VPA) Recommender loaded histograms from checkpoints following a restart. This caused the VPA Recommender to produce invalid memory recommendations with a value of 0, which could lead to pods receiving insufficient resources and failing to perform correctly. With this release, the logic for loading histogram data has been corrected to eliminate these rounding errors. As a result, the VPA Recommender now consistently provides accurate memory recommendations even after a restart, ensuring the stability and proper scaling of managed workloads. ([OCPBUGS-63455](#))
- Before this update, when a **MachineDeployment** was in the process of upgrading its machines and the Cluster Autoscaler was also scaling the **MachineDeployment**, it was possible for the Cluster Autoscaler to remove new machines by scaling down the **MachineDeployment** for idle nodes. With this release, new functionality has been implemented so that scale down does not occur when a **MachineDeployment** is in the process of upgrading its machines. ([OCPBUGS-63603](#))
- Before this update, a Cluster Ingress Operator pod could restart with existing **IngressController** resources in **Available** or **Degraded** status causing the **ingress_controller_conditions** metric to disappear from the Operator's **/metrics** endpoint. As a result, users were unable to monitor the **IngressController** status following a pod restart. With this release, the **IngressControllerConditions** metric is now set during every reconciliation cycle, regardless of whether an Ingress controller status update occurred, ensuring reliable and continuous monitoring of the **IngressController** health. ([OCPBUGS-65723](#))
- Before this release, any unrelated changes to a **netpol** resource triggered a full reconcile of the object, including deleting and re-adding rules. With this release, a **netpol** object fully reconciles when required. Otherwise, it is skipped. ([OCPBUGS-65956](#))
- Before this update, **AdminNetworkPolicy**, **AdminPolicyBasedRouteListers**, **EgressFirewall**, **EgressQoS**, and **NetworkQoS** objects were retaining **managedFields** status entries for nodes that had been deleted, leading to buildup of stale data in etcd for large clusters with frequent node churn. With this release, the cleanup logic is fixed for all the aforementioned resource types. ([OCPBUGS-66139](#))

- Before this update, when opening a terminal to a running pod, the session was disconnected whenever the annotations of the pod changed. With this release, the terminal session no longer disconnects when this metadata is changed. ([OCPBUGS-66179](#))
- Before this update, the filtering mechanism was not handling the scenario where a range was actually a specific version, which resulted in the system mirroring more Operator versions than requested by the customer. With this release, we have fixed the logic to ignore a range when only one version is required, ensuring that only the specific requested version is mirrored. ([OCPBUGS-66410](#))
- Before this update, a bug prevented the **ValidatingAdmissionPolicy** resource from applying to certain OpenShift Container Platform API resources, such as **BuildConfig** and **DeploymentConfig**. This meant that custom admission policies were not enforced on these specific resources, potentially allowing configurations that did not meet organizational standards to be created or updated. With this release, the validation logic has been corrected to ensure that the `ValidatingAdmissionPolicy` resource now correctly identifies and applies to all intended OpenShift Container Platform resources. As a result, users can consistently enforce policies across their entire cluster, including the BuildConfig and DeploymentConfig resources. (OCPBUGS-66920)`
- Before this update, changes in image mode for OpenShift increased space usage on the rendezvous node's ephemeral temporary file system to approximately 9.4GB during cluster bootstrapping. Because the ephemeral temporary file system is capped at 50% of available RAM, installation would fail on hosts with less than 19GiB of memory due to insufficient space for container images. With this release, the additional data has been moved to a separate temporary file system. As a result, any rendezvous host meeting the minimum RAM requirement for a control plane node (16GB) now has sufficient capacity to successfully bootstrap a cluster. ([OCPBUGS-66999](#))
- Before this update, users could set **prometheus**, **prometheus_replica**, or **cluster** as Prometheus external labels to the **cluster-monitoring-config** and **user-workload-monitoring-config** config maps. This was not recommended and could cause issues with the cluster. With this release, when you run **oc adm upgrade**, **Upgradeable** is **False** if reserved Prometheus **externalLabels** are used, and the upgrade is blocked. ([OCPBUGS-67003](#))
- Before this update, the deployment agent was unable to find a valid IP address when the Ironic API was running on a non-standard port because the reachability test only used the hostname. This limitation resulted in a **LookupAgentIPError**, preventing successful connections and halting the deployment process in environments with custom port configurations. With this release, the reachability test has been updated to use the full API URL, including the specific port number, instead of relying solely on the hostname. As a result, the agent can now consistently identify the correct IP address and connect to Ironic APIs on any port, ensuring smoother deployments and eliminating the previous lookup errors. ([OCPBUGS-67304](#))
- Before this update, OpenShift Container Platform 4.16 and later versions failed to respect the timeout **http-keep-alive** setting due to a known upstream High HAProxy bug, preventing users from effectively managing connection persistence. This lack of control resulted in inconsistent connection behaviors, where long-lived sessions might be terminated unexpectedly or held open longer than normal. With this release, the **HTTPKeepAliveTimeout** tuning option has been integrated into the **IngressController** API, providing a formal way for customers to configure and enforce this specific timeout. As a result, cluster administrators now possess the granular control necessary to align connection persistence with specific application needs. ([OCPBUGS-68378](#))

- Before this update, the **Operand** details page would incorrectly show information using only half of the screen's display area in the OpenShift web console. With this release, the **Operand** details take up the full page width as expected. ([OCPBUGS-69677](#))
- Before this update, **iptables-alerter** pods experienced high CPU usage in some clusters due to an issue in the 4.18.20 upgrade. As a consequence, high CPU usage impacted **iptables-alerter pods**, causing performance degradation. With this release, **iptables-alerter** CPU usage has been reduced by optimizing code in version 4.18.21. As a result, high CPU usage for **iptables-alerter** pods has been resolved, improving cluster performance. ([OCPBUGS-70328](#))
- Before this update, when service endpoints were deleted and re-created in OpenShift Container Platform clusters using OVN-Kubernetes networking and the service port differed from the endpoint port, stale User Datagram Protocol (UDP) connection tracking (**contrack**) entries could remain on worker nodes. This occurred because the **contrack** cleanup logic incorrectly used the endpoint port, which is the target port on the pod, instead of the externally-facing service port that clients connect to when attempting to delete stale entries. With this release, when service endpoints are deleted or updated, the cleanup process correctly uses the service port to match and remove stale **contrack** entries. This change ensures that network connectivity continues to work reliably across endpoint lifecycle events. ([OCPBUGS-70345](#))

1.9.10. Updating

To update an OpenShift Container Platform 4.19 cluster to this latest release, see [Updating a cluster using the CLI](#).

1.9.11. RHBA-2025:22786 - OpenShift Container Platform 4.19.21 bug fix advisory

Issued: 10 December 2025

OpenShift Container Platform release 4.19.21 is now available. The list of bug fixes that are included in the update is documented in the [RHBA-2025:22786](#) advisory. The RPM packages that are included in the update are provided by the [RHBA-2025:22766](#) advisory.

Space precluded documenting all of the container images for this release in the advisory.

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

```
$ oc adm release info 4.19.21 --pullspecs
```

1.9.12. Bug fixes

There are no notable bug fixes in this release.

1.9.13. Updating

To update an OpenShift Container Platform 4.19 cluster to this latest release, see [Updating a cluster using the CLI](#).

1.9.14. RHBA-2025:22278 - OpenShift Container Platform 4.19.20 bug fix advisory

Issued: 2 December 2025

OpenShift Container Platform release 4.19.20 is now available. The list of bug fixes that are included in the update is documented in the [RHBA-2025:22278](#) advisory. The RPM packages that are included in the update are provided by the [RHBA-2025:22276](#) advisory.

Space precluded documenting all of the container images for this release in the advisory.

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

```
$ oc adm release info 4.19.20 --pullspecs
```

1.9.15. Bug fixes

The following bugs are fixed with this release:

- Before this update, if NetworkManager was restarted or crashed on a node with a **br-ex** interface managed by NMState, the node lost network connectivity. With this release, a fallback check in the dispatcher script was added to detect NMState-managed **br-ex** interfaces by checking for the **br-ex-br** bridge ID when the standard **br-ex** bridge ID is not found. As a result, nodes with this interface type do not lose network connectivity when NetworkManager restarts or crashes. ([OCPBUGS-62168](#))
- Before this update, during the mirror operation, **oc-mirror** inadvertently set the executable program flag on some synchronized files that did not contain executable program code or scripts, potentially causing unexpected execution. With this release, unintended executable program flags have been removed from the synchronized files. As a result, correct file permissions are set, preventing unintended execution of synchronized files. ([OCPBUGS-64683](#))
- Before this update, when directly navigating to a page created by a web console dynamic plugin, the web console might redirect to a different URL. With this release, the URL redirect has been removed. ([OCPBUGS-64834](#))
- Before this update, the **ccoctl** utility did not support pagination when retrieving **CloudFront** distributions. As a result, if the distribution to be deleted was not included in the first batch of results, the **CloudFront** distribution and its associated origin access identity could not be deleted successfully during the **ccoctl** Amazon Web Services (AWS) delete operation. With this release, pagination support is added to the **ccoctl** utility when fetching **CloudFront** distributions, ensuring that the distribution can be located and deleted properly. ([OCPBUGS-65478](#))
- Before this update, a race condition in the Redfish Power interface caused power operations to fail during simultaneous access. As a consequence, users were unable to manage power states reliably. With this release, the race condition in the Redfish Power interface has been resolved, ensuring successful power operations. As a result, users can now manage power states reliably. ([OCPBUGS-65572](#))
- Before this update, it was impossible to schedule a **must-gather** pod to a specific worker node when the **--node-name** argument was used as the pod's node affinity accepted only control plane nodes. With this release, the **must-gather** logic is updated to avoid setting node affinity when the **--node-name** argument is set. ([OCPBUGS-65594](#))

1.9.16. Updating

To update an OpenShift Container Platform 4.19 cluster to this latest release, see [Updating a cluster using the CLI](#).

1.9.17. RHBA-2025:21363 - OpenShift Container Platform 4.19.19 bug fix advisory

Issued: 19 November 2025

OpenShift Container Platform release 4.19.19 is now available. The list of bug fixes that are included in the update is documented in the [RHBA-2025:21363](#) advisory. The RPM packages that are included in the update are provided by the [RHBA-2025:21361](#) advisory.

Space precluded documenting all of the container images for this release in the advisory.

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

```
$ oc adm release info 4.19.19 --pullspecs
```

1.9.18. Bug fixes

The following bugs are fixed with this release:

- Before this update, French users experienced a login screen in Spanish due to a locale mismatch. As a consequence, the user interface displayed an incorrect language: Spanish instead of French. With this release, the login screen language is corrected for French users. As a result, the login screen displays in French for French users, and improves localization. ([OCPBUGS-58892](#))
- Before this update, a bonded network configuration caused missing **bootMACAddress** configuration setting errors, and prevented host registrations with the Ironic API service. As a consequence, users could not register hosts with the Ironic service due to missing setting. With this release, bonded network MAC consistency is restored for the Ironic service agent registration. As a result, bonded network MAC consistency is correct, and the host registration in the Assisted Installer is successful. ([OCPBUGS-62441](#))
- Before this update, the Control Plane Operator (CPO) incorrectly used the overridden image for the Cluster Version Operator (CVO). As a consequence, the user experienced a CVO deployment with an incorrect image. With this release, the CVO correctly uses image overrides for the CVO. As a result, the correct image is used for CVO deployment. ([OCPBUGS-62959](#))
- Before this update, the User Workload Monitoring Prometheus Operator excessively accessed the Kubernetes API because of an improper Secret reference object in the **AlertmanagerConfig** custom resource definition. As a result, User Workload Monitoring overloaded the Kubernetes API, and caused increased primary node CPU use. With this release, excessive Secret object **GET** requests in the User Workload Monitoring Prometheus Operator are reduced. As a result, the API load on primary nodes is optimized. ([OCPBUGS-63197](#))
- Before this update, the kubelet server certificate was not updated after a certificate rotation due to unauthorized access. As a consequence, the cluster failed to start in a healthy state because of an unauthorized access after the rotation. With this release, the kubelet server certificate is updated after a rotation. As a result, the certificate rotation in OpenShift Container Platform clusters is successful, which ensures secure communication and a healthy cluster state. ([OCPBUGS-63342](#))
- Before this update, the data disk configuration from a MachineSet Custom Resource Definition (CRD) was not passed to the **StorageProfile** object when a virtual machine (VM) was created on Azure Stack. This action caused the configuration to be ignored. As a consequence, user custom disk configurations were applied to VMs. With this release, data disk configuration in a

MachineSet CRD is passed to the **StorageProfile** object, which ensures consistent handling on Azure Stack. As a result, data disk configurations are passed to the **StorageProfile** object, proper setup on Azure Stack is ensured, and VM creation is improved. ([OCPBUGS-63578](#))

- Before this update, the communication matrix project failed to create endpoint slices for open ports 9193 and 9194 on the primary node due to a missing service connection. As a consequence, missing endpoint slices for open ports caused inaccurate communication matrixes. With this release, the service is connected to open ports 9193 and 9194, resolving the missing endpoint slices, which ensures accurate communication matrixes for OpenShift Container Platform users. ([OCPBUGS-63586](#))
- Before this update, large journal downloads caused browser crashes in the node log viewer. As a consequence, journal logs overloaded and crashed the log viewer, affecting user experience. With this release, the log download size is limited, which prevents browser crashes and displays journal lines in the log viewer. ([OCPBUGS-63607](#))
- Before this update, the denylist metric incorrectly formatted the regular expression for the Kubernetes custom resource by omitting the **annotations** field. As a consequence, users experienced missing metrics due to an incorrect denylist configuration. With this release, unnecessary entries are removed from the metric denylist. As a result, registry metrics include missing annotations, and data accuracy is improved. ([OCPBUGS-64578](#))
- Before this update, the node scheduled a pod on a tainted node with no specific tolerance. As a consequence, the must-gather pod was scheduled on an unavailable node, and caused the log collection to fail. With this release, the pod is scheduled on a tainted node only if it includes a specific tolerance for the taint on the node where it should be scheduled. ([OCPBUGS-64585](#))

1.9.19. Updating

To update an OpenShift Container Platform 4.19 cluster to this latest release, see [Updating a cluster using the CLI](#).

1.9.20. RHBA-2025:19301 - OpenShift Container Platform 4.19.18 image release and bug fix advisory

Issued: 05 November 2025

OpenShift Container Platform release 4.19.18 is now available. The list of bug fixes that are included in the update is documented in the [RHBA-2025:19301](#) advisory. The RPM packages that are included in the update are provided by the [RHBA-2025:19299](#) advisory.

Space precluded documenting all of the container images for this release in the advisory.

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

```
$ oc adm release info 4.19.18 --pullspecs
```

1.9.20.1. Bug fixes

- Before this update, the **OAuth** route was not accepted by the private router when the DNS record was not registered in the **external-dns** Operator. This action led to improper URL resolution, and the console failed to access the **OAuth** route. As a consequence, the **Console**

ClusterOperator was stuck. With this release, the **OAuth** route admission and URL resolution issues in the hosted cluster is fixed. As a result, the **OAuth** route is accessible and the console access is approved. ([OCBUGS-61407](#))

- Before this release, deleting a **VolumeSnapshot** object for a persistent volume provisioned by the Azure file CSI driver also deleted the underlying fileshare and resulted in data loss. With this release, this issue is corrected by updating the driver to ensure that only the snapshot is removed. The source fileshare is preserved. ([OCBUGS-62911](#))
- Before this update, frequent driver configuration updates that were caused by an inconsistent storage class order in a hosted cluster namespace caused **ConfigMap** content flaps. This issue resulted in inconsistent storage class enforcement and affected user experience. With this release, the **ConfigMap** driver configuration is stabilized, which prevents storage classes from flapping, improves **ConfigMap** driver configuration stability, and prevents frequent flapping of a storage class order in a hosted cluster namespace. ([OCBUGS-62807](#))
- Before this update, Redfish transactions in single-node OpenShift nodes failed due to an empty **eTag** field in **metal3-ironic** container logs. As a consequence, users experienced failed Redfish transactions on single-node OpenShift nodes. With this release, the Redfish transaction **eTag** field issue is resolved, and results in correct **eTags**. As a result, the Redfish transaction does not fail, allowing the telecommunications company to use the **HostFirmwareSettings** parameter in single-node OpenShift nodes. ([OCBUGS-62961](#))
- Before this update, excessive CPU overcommitment that was greater than 200% caused the **KubeCPUOvercommit** alert to stop triggering after 10 minutes. As a consequence, users were unaware of the CPU overcommitment due to the missing alert. With this release, the **KubeCPUOvercommit** alert triggers correctly when CPU limits are overcommitted, and ensures timely resource management and improved cluster stability. ([OCBUGS-62965](#))
- Before this update, the **KubeMemoryOvercommit** alert falsely triggered on small multi-node clusters after memory-consuming spikes occurred within the permitted limits. With this release, the alert expression is adjusted to correctly account for small multi-node clusters. As a result, the **KubeMemoryOvercommit** alert does not falsely trigger after these instances. ([OCBUGS-62966](#))
- Before this update, a Kubernetes **StatefulSet** status replica alert did not activate for invalid pod specifications when a controller failed to create a pod. As a consequence, users received false assurance when the **StatefulSet** failed to create the required number of replicas. With this release, a Kubernetes **StatefulSet** replica count alert triggers for unsuccessful pod creation. As a result, the alert displays correctly when **StatefulSet** replicas do not match the configured amount. ([OCBUGS-62967](#))
- Before this update, the **KubeAggregatedAPIErrors** alert triggered based on the total number of errors across instances, and caused sensitive user alerts for multiple instances of an API. With this release, the alerting function for the **KubeAggregatedAPIErrors** alert is changed to operate at the instance level, reducing false alarms for APIs with multiple instances. ([OCBUGS-62968](#))
- Before this update, alerts were not filtered for cordoned nodes, leading to false positives for nodes under maintenance. As a consequence, users experienced false positives due to the filtering of cordoned nodes in alerts. With this release, the cordoned nodes are filtered from alerts to reduce false positives during maintenance. As a result, maintenance alerts are correct and false positives are reduced for cordoned nodes. ([OCBUGS-62969](#))
- Before this update, destroying an OpenShift Container Platform cluster on Google Cloud

caused a panic error due to a null pointer reference cancellation in the **waitFor** method. This was due to a failed Google Cloud API call or an uninitialized client. As a consequence, users experienced a panic error during cluster destruction on Google Cloud. With this release, the null pointer issue in the cluster uninstaller is fixed, and panic errors during a Google Cloud destruction are prevented. As a result, the panic error during cluster destruction on Google Cloud is resolved, ensuring smooth deletion of resources. ([OCPBUGS-62981](#))

- Before this update, an administrative user with a single namespace role created a pod and encountered a blank page while viewing metrics due to an incorrect perspective in the URL. As a consequence, the user could not view CPU usage metrics. With this release, the administrative user can view pod metrics in the developer perspective, and the user can view CPU usage metrics. ([OCPBUGS-62999](#))
- Before this update, the Oracle Container Storage Interface (CSI) node registrar container logged **gRPC** connection checks every 10 seconds, excessively filling Elasticsearch space. This rapid log filling increased user costs. With this release, the logging frequency of **gRPC** connection checks in the **csi-node-registrar** container is reduced. As a result, the Elasticsearch log capacity is increased, which lowers costs and improves performance. ([OCPBUGS-63193](#))
- Before this update, if a user did not define an EgressIP failover between gateway nodes with a proper IP address, the EgressIP address was not reassigned to the second gateway node after a reboot. This resulted in a communication failure between the pod and the external system. With this release, the EgressIP failover logic is improved for dual-stack environments, ensuring that the EgressIP address is properly reassigned to the available gateway node after a reboot. The communication between pods and external systems after a gateway node reboot is not interrupted. ([OCPBUGS-63234](#))
- Before this update, the upgrade to OpenShift Container Platform 4.18.24 triggered a kubelet failure on primary nodes due to **ocp-tuned-one-shot** service issues. As a consequence, the kubelet failed to start on primary nodes during upgrades. With this release, the kubelet issue is resolved by fixing the **ocp-tuned-one-shot** service. As a result, the kubelet starts on primary nodes after the upgrade. ([OCPBUGS-63418](#))

1.9.20.2. Updating

To update an OpenShift Container Platform 4.19 cluster to this latest release, see [Updating a cluster using the CLI](#).

1.9.21. RHSA-2025:18233 - OpenShift Container Platform 4.19.17 image release and bug fix advisory

Issued: 22 October 2025

OpenShift Container Platform release 4.19.17 is now available. The list of bug fixes that are included in the update is documented in the [RHSA-2025:18233](#) advisory. The RPM packages that are included in the update are provided by the [RHBA-2025:18201](#) advisory.

Space precluded documenting all of the container images for this release in the advisory.

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

```
$ oc adm release info 4.19.17 --pullspecs
```

1.9.21.1. Bug fixes

- Before this update, the Cluster Version Operator (CVO) in 4.19.9 and 4.18.23 started to require bearer token authentication in metrics requests. As a consequence, HyperShift and Hosted clusters were broken because the metrics scraper did not provide client authentication. With this release, the CVO does not require client authentication for metrics requests. As a result, access to CVO metrics scraping is recovered on HyperShift and Hosted clusters. ([OCBUGS-62868](#))
- Before this update, the installation program did not allow IPv6 primary dual-stack installations on platforms that were designated as **None** or **External**. As a consequence, an error or configuration block occurred when you proceeded with a dual-stack installation on these platform types. With this release, you can successfully install IPv6 primary dual-stack configurations on **None** and **External** platforms. ([OCBUGS-62911](#))

1.9.21.2. Updating

To update an OpenShift Container Platform 4.19 cluster to this latest release, see [Updating a cluster using the CLI](#).

1.9.22. RHBA-2025:17663 - OpenShift Container Platform 4.19.16 image release and bug fix advisory

Issued: 14 October 2025

OpenShift Container Platform release 4.19.16 is now available. The list of bug fixes that are included in the update is documented in the [RHBA-2025:17663](#) advisory. The RPM packages that are included in the update are provided by the [RHBA-2025:17660](#) advisory.

Space precluded documenting all of the container images for this release in the advisory.

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

```
$ oc adm release info 4.19.16 --pullspecs
```

1.9.22.1. Bug fixes

- Before this update, the timeout on one etcd member caused context deadlines to exceed. As a consequence, all members were declared unhealthy, even though some were reachable. With this release, if one member times out, other members are no longer incorrectly marked as unhealthy. ([OCBUGS-60941](#))
- Before this update, a pod with a secondary interface in an OVN-Kubernetes local network (mapped to the br-ex bridge) could communicate with pods on the same node that used the default network for connectivity only if the local network IP addresses were within the same subnet as the host network. With this release, you can extract the local network IP addresses from any subnet. In this generalized case, an external router outside the cluster is expected to connect the local network subnet to the host network. ([OCBUGS-61454](#))
- Before this update, an external actor could unordon a node that the Machine Config Operator (MCO) is draining. As a consequence, the MCO and the scheduler would schedule and unschedule pods at the same time, prolonging the drain process. With this fix, the MCO attempts to reordon the node if an external actor unordons it during the drain process. As a result, the MCO and scheduler no longer schedule and remove pods at the same time. ([OCBUGS-62003](#))
- Before this update, the omission of binary version display in **oc-mirror** output hindered debugging, causing delays in identifying required fixes and slowing user experience. With this

release, **oc-mirror** now displays its version in output for easier debugging. As a result, end users can easily identify the **oc-mirror** version for faster debugging. ([OCPBUGS-62311](#))

- Before this update, Prometheus for both the platform and user workload monitoring would negotiate and accept UTF-8 metrics, even though the monitoring stack does not yet fully support UTF-8. With this release, Prometheus no longer accepts UTF-8 metrics. ([OCPBUGS-62429](#))
- Before this update, a race condition would sometimes cause an intermittent failure, or "flake", when a persistent volume claim (PVC) was resized too quickly after being created. As a consequence, this resulted in an error where the system would incorrectly report that the bound persistent volume (PV) could not be found. With this release, the timing issue was fixed, so resizing a PVC right after its creation works. ([OCPBUGS-62468](#))

1.9.22.2. Updating

To update an OpenShift Container Platform 4.19 cluster to this latest release, see [Updating a cluster using the CLI](#).

1.9.23. RHBA-2025:17237 - OpenShift Container Platform 4.19.15 image release and bug fix advisory

Issued: 07 October 2025

OpenShift Container Platform release 4.19.15 is now available. The list of bug fixes that are included in the update is documented in the [RHBA-2025:17237](#) advisory. The RPM packages that are included in the update are provided by the [RHBA-2025:17235](#) advisory.

Space precluded documenting all of the container images for this release in the advisory.

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

```
$ oc adm release info 4.19.15 --pullspecs
```

1.9.23.1. Bug fixes

- With this release, as part of the **oc-mirror** v1 deprecation process, a warning message indicates that the **--v1** or **--v2** flag is mandatory. As a result, **oc-mirror** fails if you do not specify these flags. ([OCPBUGS-62062](#))
- Before this update, the **/auth/error** page did not render correctly. As a consequence, the page was empty and error details did not appear. With this release, the front end error page content appears on the **/auth/error** page. ([OCPBUGS-62083](#))

1.9.23.2. Updating

To update an OpenShift Container Platform 4.19 cluster to this latest release, see [Updating a cluster using the CLI](#).

1.9.24. RHBA-2025:16693 - OpenShift Container Platform 4.19.14 image release and bug fix advisory

Issued: 30 September 2025

OpenShift Container Platform release 4.19.14 is now available. The list of bug fixes that are included in the update is documented in the [RHBA-2025:16693](#) advisory. The RPM packages that are included in the update are provided by the [RHBA-2025:16691](#) advisory.

Space precluded documenting all of the container images for this release in the advisory.

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

```
$ oc adm release info 4.19.14 --pullspecs
```

1.9.24.1. Bug fixes

- Before this update, the ignition server deployment used a global **mirroredReleaseImage** state that could be modified by concurrent image lookup operations, which caused race conditions. As a consequence, the **MIRRORED_RELEASE_IMAGE** environment variable flipped between the original image and its mirror registry, which triggered constant deployment regenerations. With this release, the global mirror state is replaced with an image-specific lookup logic, which ensures deterministic mirror resolution and eliminates defensive filtering for empty registry entries. As a result, ignition server deployments remain stable with consistent **MIRRORED_RELEASE_IMAGE** values, which eliminates unnecessary pod restarts and deployment churn. ([OCPBUGS-61677](#))
- Before this update, the **Expand** button in the **Pod** and **Node logs** page in the web console did not work correctly. As a consequence, you could not provide input at the prompt in the terminal. With this release, the browser is set to full-screen when you click the **Expand** button. As a result, you can successfully provide input in the terminal. ([OCPBUGS-61821](#))
- Before this update, an **NMState** service failure occurred in OpenShift Container Platform deployments because of a **NetworkManager-wait-online** dependency issue in baremetal and multiple network interface controller (NIC) environments. As a consequence, an incorrect network configuration caused deployment failures. With this release, the **NetworkManager-wait-online** dependency for baremetal deployments is updated, which reduces deployment failures and ensures **NMState** service stability. ([OCPBUGS-61835](#))

1.9.24.2. Updating

To update an OpenShift Container Platform 4.19 cluster to this latest release, see [Updating a cluster using the CLI](#).

1.9.25. RHBA-2025:16148 - OpenShift Container Platform 4.19.13 image release, bug fix, and security update advisory

Issued: 23 September 2025

OpenShift Container Platform release 4.19.13, which includes security updates, is now available. The list of bug fixes that are included in the update is documented in the [RHBA-2025:16148](#) advisory. The RPM packages that are included in the update are provided by the [RHBA-2025:16146](#) advisory.

Space precluded documenting all of the container images for this release in the advisory.

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

```
$ oc adm release info 4.19.13 --pullspecs
```

1.9.25.1. Bug fixes

- Before this update, the **config-sync-controller** did not log results due to missing logging statements in the code. As a consequence, users experienced silent failures in the **config-sync-controller**. With this release, the **config-sync-controller** logs results enhancing error diagnosis for users. ([OCPBUGS-56788](#))
- Before this update, calls to retrieve an image manifest and metadata using a tagged image name did not cache the result of the lookup. As a consequence, hosted control plane memory usage quickly grew, which created performance issues. With this release, images in hosted control plane using a named tag or canonical name are cached for 12 hours. As a result, hosted control plane memory usage is optimized. ([OCPBUGS-59933](#))
- Before this update, the **agent-based-installer** set the permissions for the etcd directory **/var/lib/etcd/member** as 0755 when using single-node OpenShift deployment instead of 0700, which is correctly set on a multi-node deployment. With this release, the etcd directory **/var/lib/etcd/member** permissions are set to 0700 for single-node OpenShift deployments. ([OCPBUGS-61313](#))
- Before this update, the **PrometheusRemoteWriteBehind** alert fired if the remote endpoint never received any data. With this release, the **PrometheusRemoteWriteBehind** alert no longer fires if the remote endpoint has not yet received any data. ([OCPBUGS-61486](#))
- Before this update, it was possible for webhook failures to trigger a **kube-apiserver** crash while generating an audit log entry for a request. As a consequence, API server disruptions were possible. With this release, the audit system has been updated so that the **kube-apiserver** no longer crashes and the API disruptions are resolved. ([OCPBUGS-61488](#))
- Before this update, the **Operand details page** in the web console would show additional status items in a third column, which resulted in the content appearing squashed. With this update, the defect has been corrected, so that only two columns display in the details page. ([OCPBUGS-61781](#))

1.9.25.2. Updating

To update an OpenShift Container Platform 4.19 cluster to this latest release, see [Updating a cluster using the CLI](#).

1.9.26. RHBA-2025:15694 - OpenShift Container Platform 4.19.12 image release, bug fix, and security update advisory

Issued: 16 September 2025

OpenShift Container Platform release 4.19.12, which includes security updates, is now available. The list of bug fixes that are included in the update is documented in the [RHBA-2025:1694](#) advisory. The RPM packages that are included in the update are provided by the [RHBA-2025:15692](#) advisory.

Space precluded documenting all of the container images for this release in the advisory.

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

```
$ oc adm release info 4.19.12 --pullspecs
```

1.9.26.1. Enhancements

- With this update, the **cluster-etcd-operator** Operator now implements a multi-stage notification system for the **etcdDatabaseQuotaLowSpace** alert to proactively manage etcd storage quotas. This enhancement prevents API server instability by providing earlier warnings of low database space. As etcd disk space usage reaches 65%, 75% and 85%, administrators now receive alerts with a severity level of info, warning, or critical. ([OCPCBUGS-60443](#))
- With this update, the collection of command line logs from **virt-launcher** pods across a Kubernetes cluster are enabled. JSON-encoded logs are saved at the path **namespaces/<namespace_name>/pods/<pod_name>/virt-launcher.json**, facilitating troubleshooting and debugging of virtual machines. ([OCPCBUGS-61485](#))
- With this update, the machine config nodes custom resource, which you can use to monitor the progress of machine configuration updates to nodes, is now generally available. With the promotion to General Availability, you can view the status of updates to custom machine config pools, in addition to the control plane and worker pools. The functionality for the feature has not changed. However, some of the information in the command output and in the status fields in the **MachineConfigNode** object have been updated. The **must-gather** for the Machine Config Operator includes all **MachineConfigNodes** objects in the cluster. For more information, see [About checking machine config node status](#).
- With this update, the **PinnedImageSet** object, which you can use to get the container images in advance, before they are actually needed, is now generally available. You can associate these images with a machine config pool. In clusters with slow, unreliable connections to an image registry, pinning images ensures that the images are available when needed. The **must-gather** for the Machine Config Operator now includes all **PinnedImageSet** objects in the cluster. For more information, see [Pinning images to nodes](#).

1.9.26.2. Bug fixes

- Before this update, if the cluster was created without an SSH key, creating a node image with the **oc adm node-image create** command failed due to the absence of the **99-worker-ssh** machine configuration. This prevented worker node image creation. With this release, the **machineConfig** for **worker-ssh** is created, enabling node image creation. As a result, node image creation for worker nodes now succeeds. ([OCPCBUGS-60832](#))
- Before this update, executing **ccoctl** multiple times while using the Amazon Web Services (AWS) platform and the **--create-private-s3-bucket** parameter caused the wrong URL to be configured for the OpenID Connect (OIDC) issuer. As a consequence, some cluster Operators could not authenticate to the AWS API. With this release, **ccoctl** properly configures the correct URL for the OIDC issuer. As a result, the cluster Operators continue to authenticate as expected. ([OCPCBUGS-60970](#))
- Before this update, the **MachineHealthCheck** custom resource (CR) did not show the default value for the **maxUnhealthy** field. With this release, the CR also documents the value it defaults to when not set. ([OCPCBUGS-61096](#))
- Before this update, the time it took to apply updates to the **multus-networkpolicy** DaemonSet scaled linearly with the node count. With this release, the DaemonSet has been updated to allow for 10% **maxUnavailable** so that the DaemonSet updates immediately in clusters larger than 10 nodes. ([OCPCBUGS-61460](#))

1.9.26.3. Updating

To update an OpenShift Container Platform 4.19 cluster to this latest release, see [Updating a cluster using the CLI](#).

1.9.27. RHBA-2025:15293 - OpenShift Container Platform 4.19.11 image release, bug fix, and security update advisory

Issued: 09 September 2025

OpenShift Container Platform release 4.19.11, which includes security updates, is now available. The list of bug fixes that are included in the update is documented in the [RHBA-2025:15293](#) advisory. The RPM packages that are included in the update are provided by the [RHSA-2025:15291](#) advisory.

Space precluded documenting all of the container images for this release in the advisory.

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

```
$ oc adm release info 4.19.11 --pullspecs
```

1.9.27.1. Enhancements

- With this update, the Kubernetes API server distribution is optimized, ensuring a balanced load among all primary nodes after a quorum is reestablished. This addresses the issue of a single API server receiving the majority of live connections, and causing high CPU usage. Resource utilization is improved and CPU spikes are reduced during primary node or API server restarts. ([OCPBUGS-60121](#))

1.9.27.2. Bug fixes

- Before this update, disabling the OpenShift image registry retained legacy pull secret finalizers and caused hung secret deletion during the registry removal. This issue blocked cluster deletion. With this release, secret finalizers do not block namespace deletion when the registry is disabled, and ensures cluster deletion. ([OCPBUGS-56614](#))
- Before this update, the **oc mirror** command failed on RHEL 8 systems with a **noexec-mounted /tmp** directory because it could not start the temporary files or scripts. As a consequence, image mirroring was prevented. With this release, the **oc mirror** command includes an exception for the **noexec-mounted /tmp** directory, and mirroring on RHEL 8 systems is successful. As a result, the **oc mirror** command lists output and mirror container images on RHEL 8 systems with a **noexec-mounted /tmp** directory. ([OCPBUGS-59760](#))
- Before this update, temporary **apiserver** downtime caused the **cluster-etcd-operator** to incorrectly report that the **openshift-etcd** namespace did not exist. As a consequence, users saw incorrect messages about a missing namespace during the downtime. With this release, a fix is implemented to improve the error message for a missing etcd namespace. As a result, the error message is corrected, and ensures that the **cluster-etcd-operator** status accurately reflects the issue during temporary **apiserver** downtime. ([OCPBUGS-59802](#))
- Before this update, duplicate link buttons on the **Quickstarts** page appeared only in the **/quickstart** path, and confused users. With this release, the **Quickstart** link buttons display correctly and duplicates are eliminated. ([OCPBUGS-60420](#))
- Before this update, a Hosted control planes cluster rejected certificates with multiple Storage Area Network (SAN) entries due to conflicting DNS names. As a consequence, users encountered certificate deployment errors with multi-SAN hostnames in Hosted control planes clusters. With this release, certificate validation for multiple SAN entries is supported in Hosted control planes clusters. As a result, certificates with multiple SAN entries are accepted, improving Hosted control planes cluster deployments. ([OCPBUGS-60483](#))

- Before this update, the last node retained the **ToBeDeletedByClusterAutoscaler** taint during the scale-down process because of incorrect machine deletion handling. As a consequence, the last node affected the cluster autoscaling efficiency. With this release, the **ToBeDeletedByClusterAutoscaler** taint is removed from the last node after scaling down. The last node does not retain the unwanted taint, and the cluster stability is improved. ([OCPBUGS-60900](#))
- Before this release, the stale IP address entries in the **address_set** configuration element that corresponded to the DNS egress firewall rule were not removed. This resulted in an increasing **address_set**, and led to memory leak issues. With this release, the issue is fixed by removing the IP addresses from the **address_set** after a 5-second grace period that follows the Time to Live (TTL) expiration of the IP addresses. ([OCPBUGS-60979](#))

1.9.27.3. Updating

To update an OpenShift Container Platform 4.19 cluster to this latest release, see [Updating a cluster using the CLI](#).

1.9.28. RHSA-2025:14823 - OpenShift Container Platform 4.19.10 image release, bug fix, and security update advisory

Issued: 02 September 2025

OpenShift Container Platform release 4.19.10, which includes security updates, is now available. The list of bug fixes that are included in the update is documented in the [RHBA-2025:14823](#) advisory. The RPM packages that are included in the update are provided by the [RHBA-2025:14817](#) advisory.

Space precluded documenting all of the container images for this release in the advisory.

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

```
$ oc adm release info 4.19.10 --pullspecs
```

1.9.28.1. Enhancements

- The name of the **MachineOSConfig** object used with on-cluster image mode must now be the same as the machine config pool where you want to deploy the custom layered image. This change prevents using multiple **MachineOSConfig** objects with each machine config pool. ([OCPBUGS-60414](#))
- In OpenShift Container Platform 4.19.10, Operator Lifecycle Manager (OLM) Classic and OLM v1 allow Operators to include network policy manifests in their resource bundles. These tailored network policies protect against data leaks and harden against many attack vectors on OpenShift Container Platform clusters.

TIP

If your current version of OLM does not support tailored network policies, a notification is displayed in the following locations:

- The Red Hat Hybrid Cloud Console
- The web console of the affected cluster

Update to OpenShift Container Platform 4.19.10 or later to enable OLM support for tailored network policies.

For more information, including the planned timeline for releasing Red Hat-provided Operators with tailored network policies, see [Operators shipping with network policies may require OCP cluster upgrade before they can be upgraded \(Red Hat Knowledgebase\)](#).

([OCPBUGS-60525](#) and [OCPBUGS-60521](#))

1.9.28.2. Bug fixes

- Before this update, the Hosted Control Plane (HCP) did not query payload repositories sequentially when the management cluster was configured with many image repositories, causing the hosted cluster deployments to fail in disconnected environments if the first mirror was unavailable. The system errored out instead of searching for the next available image. With this release, the HCP payload iterates through the entire list of mirrors until an available image is found, allowing deployments to succeed as expected. ([OCPBUGS-57141](#))
- Before this update, when deploying single-node OpenShift by using zero-touch provisioning (ZTP) in release 4.19 with many IP addresses configured on the primary interface, the **apiserver** pod would fail to connect to etcd. As a result, the etcd certificate did not include all the configured IP addresses, leading to Transport Layer Security (TLS) authentication errors. With this release, the **apiserver** pod can now successfully connect to etcd in these configurations, allowing single-node OpenShift deployments with many primary interface IP addresses to initialize correctly. ([OCPBUGS-59285](#))
- Before this update, IBM Cloud was not included in the validation check for single-node OpenShift installs, causing a validation error when attempting to install single-node OpenShift on IBM Cloud. With this release, IBM Cloud now supports single-node OpenShift installs improving the installation experience for end users on IBM Cloud. ([OCPBUGS-59607](#))
- Before this update, the **Delete** workflow erroneously displayed **workflow mode: diskToMirror / delete**, leading to user confusion regarding the correct workflow mode. With this release, **workflow mode: delete** displays during delete operations. ([OCPBUGS-59761](#))
- Before this update, sharing duplicated images between different container images resulted in a wrong count calculation in **oc-mirror** for total mirrored images of Helm charts. As a consequence, some Helm images were not mirrored. With this release, the wrong count of mirrored Helm images in **oc-mirror** has been fixed, improving the accuracy of mirrored image counts. ([OCPBUGS-60086](#))
- Before this update, the **HorizontalPodAutoscaler** temporarily scaled **istiod-openshift-gateway** to two replicas, causing Continuous Integration (CI) failure due to the tests expecting only one replica. With this release, **HorizontalPodAutoscaler** scaling is adjusted to support a single replica for **istiod-openshift-gateway**. ([OCPBUGS-60204](#))

- Before this update, an upgrade to a version before 4.15 or a new install of 4.15 deployed **MachineConfigNode** custom resource definitions (CRDs) despite being in Technology Preview. As a result, clusters failed to upgrade due to unneeded CRDs. With this release, Technology Preview **MachineConfigNode** CRDs were removed from default clusters ensuring seamless upgrades. ([OCPBUGS-60265](#))
- Before this update, on dual-stack clusters with IPv6 as the primary networking stack, the bare-metal Installer-Provisioned Infrastructure (IP) would incorrectly supply an IPv4 URL for the virtual media ISO image. This caused installation failures on Baseboard Management Controllers (BMCs) that were configured only for IPv6 networking, as the BMCs could not reach the IPv4 address. With this release, the installation program logic was updated to always supply an IPv6 URL when a BMC is using IPv6 addressing and the installation process now completes successfully. ([OCPBUGS-60402](#))
- Before this update, Amazon Web Services (AWS) **machinesets** could have a null **userDataSecret** name, leading to machines remaining in a provisioning state. With this release, a non-empty **userDataSecret** name is required, preventing unexpected machine behavior. ([OCPBUGS-60427](#))
- Before this update, a limitation prevented a certificate's validity from exceeding that of the signer. This impacted the **localhost-recovery.kubeconfig**, as the node-system-admin-client certificate was incorrectly generated with a one-year lifespan instead of the intended two years, causing the premature expiration of the **localhost-recovery.kubeconfig**. With this release, the signer certificate's validity is extended to three years, ensuring the node-system-admin-client certificate now has a two-year lifespan. ([OCPBUGS-60495](#))
- Before this update, OpenShift Container Platform clusters on AWS that were created with version 4.13 or earlier could not update to version 4.19. Clusters that were created with version 4.14 and later have an AWS **cloud-conf** ConfigMap by default, and this ConfigMap is required starting in OpenShift Container Platform 4.19. With this release, the Cloud Controller Manager Operator is updated to create a default **cloud-conf** ConfigMap when none is present on the cluster. This change enables clusters that were created with version 4.13 or earlier to update to version 4.19. ([OCPBUGS-60950](#))

1.9.28.3. Updating

To update an OpenShift Container Platform 4.19 cluster to this latest release, see [Updating a cluster using the CLI](#).

1.9.29. RHSA-2025:13848 - OpenShift Container Platform 4.19.9 image release, bug fix, and security update advisory

Issued: 19 August 2025

OpenShift Container Platform release 4.19.9, which includes security updates, is now available. The list of bug fixes that are included in the update is documented in the [RHSA-2025:13848](#) advisory. The RPM packages that are included in the update are provided by the [RHBA-2025:13827](#) advisory.

Space precluded documenting all of the container images for this release in the advisory.

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

```
$ oc adm release info 4.19.9 --pullspecs
```

1.9.29.1. Enhancements

- With this update, you can install the NUMA Resources Operator on hosted control planes, enabling NUMA-aware scheduling support. For more information, see [Creating the NUMAResourcesOperator custom resource for hosted control planes](#). This enhancement is available as a Technology Preview feature.

1.9.29.2. Bug fixes

- Before this update, boot images for 4.1 and 4.2 failed to work with OpenShift Container Platform 4.19, and caused degraded cluster operation. With this release, a static Grand Unified Bootloader (GRUB) configuration is installed for Extensible Firmware Interface (EFI) and firmware components, and the cluster operates normally during node scaling. ([OCPBUGS-52485](#))
- Before this update, the Google Cloud machine API was blocked by sequential reconciliation processing. As a consequence, users experienced slow scaling of nodes during GCP integration. With this release, the GCP machine API performance is improved by enabling parallel execution of many reconciliation processes. As a result, GCP node scaling performance improves. ([OCPBUGS-59386](#))
- Before this update, users configured OpenShift Container Platform Vertical Pod Autoscaler (VPA) custom recommenders using a version with an upstream issue in VPA. As a consequence, the issue caused instability in VPA updates. With this release, the custom VPA checkpoint garbage collector does not remove untracked checkpoints, and prevents instability in OpenShift Container Platform. As a result, OpenShift Container Platform VPA updates are stable and constant pod rescheduling does not occur. ([OCPBUGS-59638](#))
- Before this update, the machine config daemon failed Domain Name System (DNS) lookups during the OpenShift Container Platform 4.16 manifest application on VMware vSphere infrastructure. As a consequence, user DNS lookups failed during the OpenShift Container Platform 4.16 upgrade, halting upgrades indefinitely. With this release, retries of remote operating system updates with backoff are implemented to avoid failing because of CoreDNS pod restarts during the upgrades. ([OCPBUGS-59899](#))
- Before this update, cluster upgrade failures occurred because of an increased reconcile attempts limit. This failure caused Prometheus pod unavailability and resulted in service degradation. With this release, the Operator allows an additional reconcile attempt before reporting failures. As a result, the cluster upgrade test stability is improved, reducing failure rate, and enhancing upgrade reliability. ([OCPBUGS-59932](#))
- Before this update, the sidecar of a OpenShift Container Platform Precision Time Protocol (PTP) pod unexpectedly restarted after termination, and caused the clock class termination to fail with an **exit code 7** error. As a consequence, the metrics were unavailable. With this release, the sidecar restart does not cause the clock class termination error in a OpenShift Container Platform PTP pod, and does not stop during the restart. ([OCPBUGS-59970](#))
- Before this update, when a user upgraded to OpenShift Container Platform 4.19, the Machine Config Operator (MCO) rotated a Transport Layer Security (TLS) certificate. This caused an issue where nodes could not join the cluster during the scale-up process. With this release, the MCO provides a custom ARO resource that determines the necessary Subject Alternative Name (SAN) IP address, and adds it in the rotated TLS certificate. As a result, nodes can join the cluster during the scale-up process. ([OCPBUGS-59978](#))
- Before this update, an interpolation error in the **ResourceEventStream** code format caused incorrect error messages when users connected to the event stream. With this release, the

interpolation format for error messages in the event stream is correct. As a result, users see accurate error messages when they connect to the event stream. ([OCBUGS-60039](#))

- Before this update, the primary node port in the communication matrix project was unbound, and caused a missing communication flow and service unavailability on the primary node. With this release, the port on the controller manager is closed and is available only from the **localhost**. As a result, the service is bound to the correct port. ([OCBUGS-60132](#))
- Before this update, a **MachineSet** custom resource update failure occurred because of multiple arch annotation labels. As a consequence, the machine update failed. With this release, the update issue is corrected by allowing multiple labels in the **{{capacity.cluster-autoscaler.kubernetes.io/labels}}** annotation, and properly parses the architecture value. As a result, the Machine Config Operator does not fail during an update. ([OCBUGS-60224](#))
- Before this update, the **LeaderWorkerSet** Operator description was outdated. As a consequence, users experienced incorrect descriptions. With this release, the **LeaderWorkerSet** Operator description is updated and the description accurately describes the concept. ([OCBUGS-60225](#))
- Before this update, the **cloud-event-proxy** sidecar process terminated and caused the notification API to remain in a **clockClass=0** state even when the pod recovered. As a consequence, the notification API remained inactive after the sidecar process terminated. With this release, the **cloud-event-proxy** process recovery does not result in a **clockClass=0** state for the notification API. Now, the notification API correctly updates the **clockClass** variable when the **cloud-event-proxy** recovers. ([OCBUGS-60261](#))
- Before this update, insufficient obfuscation of new network data types in an OVN-K hosted cluster exposed sensitive data. As a consequence, user data was exposed. With this release, the anonymizer is updated to discover and obfuscate new network data types, and ensure secure communication. ([OCBUGS-60295](#))

1.9.29.3. Updating

To update an OpenShift Container Platform 4.19 cluster to this latest release, see [Updating a cluster using the CLI](#).

1.9.30. RHSA-2025:12341 - OpenShift Container Platform 4.19.7 image release, bug fix, and security update advisory

Issued: 05 August 2025

OpenShift Container Platform release 4.19.7, which includes security updates, is now available. The list of bug fixes that are included in the update is documented in the [RHSA-2025:12341](#) advisory. The RPM packages that are included in the update are provided by the [RHBA-2025:12342](#) advisory.

Space precluded documenting all of the container images for this release in the advisory.

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

```
$ oc adm release info 4.19.7 --pullspecs
```

1.9.30.1. Enhancements

- The KubeVirt Container Storage Interface (CSI) driver now supports volume expansion. Users

can dynamically increase the size of their persistent volumes in their tenant cluster. This capability simplifies storage management, allowing for more flexible and scalable infrastructure. ([OCPBUGS-58239](#))

1.9.30.2. Bug fixes

- Before this update, a plugin conflict in the console modal occurred due to multiple plugins using the same **CreateProjectModal** extension point. As a consequence, only one plugin extension was used and the list order could not be changed. With this release, an update to the plugin store resolves extensions in the same order that are defined in the console operator configuration. As a result, anyone with permission to update the operator configuration can set the priority of the plugin. ([OCPBUGS-56280](#))
- Before this update, when you clicked **Configure** in an **AlertmanagerReceiversNotConfigured** alert on the **Overview** page, a runtime error occurred. With this release, improved navigation handling ensures that no runtime errors occur when you click **Configure**. ([OCPBUGS-57105](#))
- Before this update, the **/metrics/usage** endpoint was updated to include authentication and Cross-Site Request Forgery (CSRF) protections. Because of this, requests to this endpoint started failing with a "forbidden" error message because the requests lacked the necessary CSRF token in the request cookie. With this release, a CSRF token was added to the **/metrics/usage** request cookie, which resolved the "forbidden" error message. ([OCPBUGS-58331](#))
- Before this update, when you configured an OpenID Connect (OIDC) provider for a **HostedCluster** resource with an Open ID cluster that did not specify a client secret, a default secret name was automatically generated. As a consequence, you could not configure OIDC public clients because these clients cannot use client secrets. With this release, a default secret name is not generated when no client secret is provided. As a result, you can configure OIDC public clients. ([OCPBUGS-58683](#))
- Before this update, when a Bare Metal Host (BMH) was marked as **Provisioned** or **ExternallyProvisioned**, the system would try to deprovision it or power it off first and the **DataImage** attached to the BMH would also prevent deletion. This issue blocked or slowed down host removal, creating operational inefficiencies. With this release, if the BMH has the **detached annotation** status and deletion is requested, the BMH transitions to the deleting state, allowing for direct deletion. ([OCPBUGS-59133](#))
- Before this update, downloads on control plane nodes were inconsistently scheduled because of a mismatch between the node selector for downloads and the console pods. As a consequence, downloads were scheduled on random nodes, which caused potential resource contention and sub-optimal performance. With this release, downloaded workloads consistently schedule on control plane nodes, which improves resource allocation. ([OCPBUGS-59488](#))
- Before this update, a cluster upgrade to OpenShift Container Platform 4.18 caused inconsistent egress IP allocation due to stale Network Address Translation (NAT) handling. This issue occurred only when you deleted an egress IP pod while the OVN-Kubernetes controller for an egress node was down. As a consequence, duplicate Logical Router Policies and egress IP usage occurred, which caused inconsistent traffic flow and outage. With this release, egress IP allocation cleanup ensures consistent and reliable egress IP allocation in OpenShift Container Platform 4.18 clusters. ([OCPBUGS-59530](#))
- Before this update, if you did not have sufficient privileges when you logged into the console, the **get started** message occupied excessive space on pages. This issue prevented the complete display of important status messages such as **no resources found**. As a consequence, truncated versions of the messages were displayed. With this release, the **get**

started message is resized and the page's disable property is removed to use less screen space and to allow scrolling. This fix allows users to view complete statuses and information on all pages. You can now view complete statuses and information on all pages. As a result, the **get started** content remains fully accessible through scrolling, which ensures the visibility of new user guidance and important system messages. ([OCPBUGS-59639](#))

- Before this update, when you cloned a **.tar** file with zero length, the **oc-mirror** ran indefinitely due to an empty archive file. As a consequence, no progress occurred when you mirrored a 0-byte **.tar** file. With this release, 0-byte **.tar** files are detected and reported as errors, which prevents the **oc-mirror** from hanging. ([OCPBUGS-59779](#))
- Before this update, the **oc-mirror** did not detect Helm Chart images that used an aliased sub-chart. As a consequence, the Helm Chart images were missing after mirroring. With this release, the **oc-mirror** detects and mirrors Helm Chart images with an aliased sub-chart. ([OCPBUGS-59799](#))

1.9.30.3. Updating

To update an OpenShift Container Platform 4.19 cluster to this latest release, see [Updating a cluster using the CLI](#).

1.9.31. RHSA-2025:11673 - OpenShift Container Platform 4.19.6 image release, bug fix, and security update advisory

Issued: 29 July 2025

OpenShift Container Platform release 4.19.6, which includes security updates, is now available. The list of bug fixes that are included in the update is documented in the [RHSA-2025:11673](#) advisory. The RPM packages that are included in the update are provided by the [RHBA-2025:11674](#) advisory.

Space precluded documenting all of the container images for this release in the advisory.

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

```
$ oc adm release info 4.19.6 --pullspecs
```

1.9.31.1. Enhancements

- The KubeVirt Container Storage Interface (CSI) driver now supports volume expansion. Users can dynamically increase the size of their persistent volumes in their tenant cluster. This capability simplifies storage management, allowing for more flexible and scalable infrastructure. ([OCPBUGS-58239](#))

1.9.31.2. Bug fixes

- Before this update, the **/metrics/usage** endpoint was updated to include authentication and Cross-Site Request Forgery (CSRF) protections. Because of this, requests to this endpoint started failing with a "forbidden" error message because the requests lacked the necessary CSRF token in the request cookie. With this release, a CSRF token was added to the **/metrics/usage** request cookie, resolving the "forbidden" error message. ([OCPBUGS-58331](#))
- Before this update, the **console.flag/model** extension point did not work, preventing flags from being properly set when their associated model was provided. With this release, the **console.flag/model** works as expected and properly sets a flag when the associated model is

provided. ([OCPBUGS-59513](#))

1.9.31.3. Updating

To update an OpenShift Container Platform 4.19 cluster to this latest release, see [Updating a cluster using the CLI](#).

1.9.32. RHSA-2025:11363 - OpenShift Container Platform 4.19.5 image release, bug fix, and security update advisory

Issued: 22 July 2025

OpenShift Container Platform release 4.19.5, which includes security updates, is now available. The list of bug fixes that are included in the update is documented in the [RHSA-2025:11363](#) advisory. The RPM packages that are included in the update are provided by the [RHBA-2025:11364](#) advisory.

Space precluded documenting all of the container images for this release in the advisory.

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

```
$ oc adm release info 4.19.5 --pullspecs
```

1.9.32.1. Bug fixes

- Before this update, bundle unpack jobs did not inherit control-plane tolerances from the catalog-operator that created them. As a consequence, the bundle unpack jobs ran on only worker nodes. If no worker nodes were available due to taints, then admins were unable to install or upgrade Operators on the cluster. With this release, control-plane tolerations are adopted for bundle unpack jobs so that the jobs are executed on primary nodes as part of the control plane. ([OCPBUGS-59258](#))
- Before this update, intermittent egress internet protocol (IP) handling due to inconsistent state updates in `OVNkubernetes` caused packet drops. These packet drops affected network traffic flow. With this release, `OVNkubernetes` pods consistently use their assigned egress IPs. As a result, dropped packages are reduced and network traffic flow is improved. ([OCPBUGS-59234](#))
- Before this update, the Amazon Web Services (AWS) Cloud Provider did not set the default ping target of **HTTP:10256/healthz** for the AWS Load Balancer. For the LoadBalancer Services that ran on AWS, the Load Balancer object created in AWS had a ping target of **TCP:32518**. As a consequence, the health probes for cluster-wide services did not work and the services were down during upgrades. With this release, the cloud config **ClusterServiceLoadBalancerHealthProbeMode** property is set to **Shared** to ensure that the config is passed to the AWS Cloud Provider. As a result, the AWS Load Balancers have the correct health check ping target of **HTTP:10256/healthz**. ([OCPBUGS-59101](#))
- Before this update, the **MachineConfigOperator** (MCO) installed the **podman-etcd** agent to enable testing while waiting for the RPM Package Manager (RPM) version to reach the repositories. With this release, the agent that was installed by MCO is removed because the RPM version is available. ([OCPBUGS-58894](#))
- Before this update, when you ran the **oc-mirror v2** disk-to-mirror workflow without valid mirror tar files, the returned error messages did not correctly identify the problem. With this release, the **oc-mirror v2** workflow returns an error message that states **no tar archives matching "mirror_[0-9]{6}.tar" found in "<directory>"**. ([OCPBUGS-58341](#))

- Before this update, the build controller searched for secrets that were linked for general use rather than specifically for the image pull. With this release, when the controller searches for the default image pull secrets, the builds use **ImagePullSecrets** that are linked to the service account. ([OCPBUGS-57951](#))
- Before this update, combined specification and status updates lists triggered unnecessary firmware upgrades, which caused system downtime. With this release, a firmware upgrade optimization skips unnecessary firmware upgrades when a Baseboard Management Controller (BMC) URL is added. ([OCPBUGS-56765](#))
- Before this update, when you defined the **blockedImages** value in the **imageSetConfiguration** parameter for **oc-mirror v2**, you were required to provide an extensive list of image references for excluding images from mirroring. This requirement sometimes prevented the exclusion of images from mirroring because the image digests changed between executions. With this release, you can use regular expressions for the **blockedImages** value to facilitate the exclusion of images from mirroring. ([OCPBUGS-56728](#))
- Before this update the **Observe > Metrics > query > QueryKebab > Export as csv** drop-down item did not handle a undefined title element. As a consequence, you could not export the CSV file for certain queries on the **Metrics** tab of OpenShift Container Platform Lister versions 4.16, 4.17, and 4.18. With this release, the metrics download for all queries correctly handles object properties in the drop-down menu items. As a result, the CSV export for all queries works on the **Metrics** page. ([OCPBUGS-52592](#))

1.9.32.2. Updating

To update an OpenShift Container Platform 4.19 cluster to this latest release, see [Updating a cluster using the CLI](#).

1.9.33. RHSA-2025:10771 - OpenShift Container Platform 4.19.4 image release, bug fix, and security update advisory

Issued: 15 July 2025

OpenShift Container Platform release 4.19.4, which includes security updates, is now available. The list of bug fixes that are included in the update is documented in the [RHSA-2025:10771](#) advisory. The RPM packages that are included in the update are provided by the [RHBA-2025:10772](#) advisory.

Space precluded documenting all of the container images for this release in the advisory.

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

```
$ oc adm release info 4.19.4 --pullspecs
```

1.9.33.1. Bug fixes

- Previously, when the Gateway API feature was enabled, it installed an Istio control plane configured with one pod replica and an associated **PodDisruptionBudget** setting. The **PodDisruptionBudget** setting prevented the only pod replica from being evicted, blocking cluster upgrades. With this release, the Ingress Operator prevents the Istio control plane from being configured with the **PodDisruptionBudget** setting allowing cluster upgrades. ([OCPBUGS-58394](#))

- Previously, a runtime error occurred when clicking **Edit HorizontalPodAutoscaler** using the form view. With this release, the **Edit HorizontalPodAutoscaler** form view renders as expected. ([OCPBUGS-58377](#))
- Previously, forward slashes were permitted in **console.tab/horizontalNav href** values. Starting in version 4.15, a regression resulted in forward slashes no longer working correctly when used in **href** values. With this release, forward slashes in **console.tab/horizontalNav href** values continue to work as expected. ([OCPBUGS-58375](#))
- Previously, when a hosted cluster was configured with a proxy URL such as **http://user:pass@host**, the authentication header was not getting forwarded by the Konnectivity proxy to the user proxy, which caused authentication to fail. With this release, the proper authentication header is sent when a user and password is specified in the proxy URL. ([OCPBUGS-58335](#))
- Previously, a subset of endpoints on the console backend were gated by **TokenReview** requests to the API server. In some cases, the API server would throttle these requests, causing slower load times in the UI. With this release, the **TokenReview** gating was removed from all but one of our endpoints resulting in improved performance. ([OCPBUGS-58316](#))
- Previously, the amount of requests that the oc-mirror plugin v2 sent many requests to container registries caused container registries to reject some requests with a **too many requests** error. With this release, the default values for several related parameters were adjusted to result in fewer requests being sent to the container registries. ([OCPBUGS-58279](#))
- Previously, the kubelet server certificate was not updated after certificate rotation due to unauthorized access to the API server causing the cluster to start in an unhealthy state. With this release, the kubelet server certificate is updated after certificate rotation, ensuring a healthy cluster state. ([OCPBUGS-58116](#))
- Previously, when on-premise installer-provisioned infrastructure deployments used the Cilium container network interface (CNI), the firewall rule that redirected traffic to the load balancer was ineffective. With this release, the rule works with the Cilium CNI and **OVNKubernetes**. ([OCPBUGS-57781](#))
- Previously, deleting an **istag** resource with the **--dry-run=server** option unintentionally caused actual deletion of the image from the server. This unexpected deletion occurred due to the dry-run option being implemented incorrectly in the **oc delete istag** command. With this release, the dry-run option is now wired to the **oc delete istag** command, preventing accidental deletion of image objects and the **istag** object remains intact when using the **--dry-run=server** option. ([OCPBUGS-57206](#))
- Previously, an outdated version of the Azure API prevented specifying a Capacity Reservation Group for a machine set, if that group resided in a different subscription than the one originating the server creation. With this release, OpenShift Container Platform uses a newer version of the Azure API that is compatible with this configuration. ([OCPBUGS-56163](#))

1.9.33.2. Updating

To update an OpenShift Container Platform 4.19 cluster to this latest release, see [Updating a cluster using the CLI](#).

1.9.34. RHBA-2025:10290 - OpenShift Container Platform 4.19.3 image release, bug fix, and security update advisory

Issued: 08 July 2025

OpenShift Container Platform release 4.19.3, which includes security updates, is now available. The list of bug fixes that are included in the update is documented in the [RHBA-2025:10290](#) advisory. The RPM packages that are included in the update are provided by the [RHSA-2025:10291](#) advisory.

Space precluded documenting all of the container images for this release in the advisory.

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

```
$ oc adm release info 4.19.3 --pullspecs
```

1.9.34.1. Bug fixes

- Previously, useful error messages were not generated when the **oc adm node-image create** command failed. With this release, the **oc adm node-image create** command provides error messages when the command fails. ([OCPBUGS-58077](#))
- Previously, when on-prem installer-provisioned infrastructure (IPI) deployments used the Cilium container network interface (CNI), the firewall rule that redirected traffic to the load balancer was ineffective. With this release, the rule works with the Cilium CNI and **OVNKubernetes**. ([OCPBUGS-57781](#))
- Previously, when you defined the **blockedImages** value in the **imageSetConfiguration** parameter for **oc-mirror v2**, you were required to provide an extensive list of image references for excluding images from mirroring. This requirement sometimes prevented the exclusion of images from mirroring because the image digests changed between executions. With this release, you can use regular expressions for the **blockedImages** value to facilitate the exclusion of images from mirroring. ([OCPBUGS-56728](#))
- Previously, certain traffic patterns with large packets running between OpenShift Container Platform nodes and pods triggered an OpenShift Container Platform host to send Internet Control Message Protocol (ICMP) needs frag to another OpenShift Container Platform host. This situation lowered the viable maximum transmission unit (MTU) in the cluster. As a consequence, executing the **ip route show cache** command generated a cached route with a lower MTU than the physical link. Packets were dropped and OpenShift Container Platform components were degrading because the host did not send pod-to-pod traffic with the large packets. With this release, NF Tables rules prevent the OpenShift Container Platform nodes from lowering their MTU in response to traffic patterns with large packets. ([OCPBUGS-55997](#))
- Previously, you needed to update vSAN files to 8.0 P05 or later to enable clusters running OpenShift Container Platform 4.19 to mount the network file system (NFS) volumes that were exported from the VMWare vSAN Files. With this release, you do not need to upgrade existing vSAN File Services versions to mount VMWare vSAN File volumes. ([OCPBUGS-55978](#))

1.9.34.2. Updating

To update an OpenShift Container Platform 4.19 cluster to this latest release, see [Updating a cluster using the CLI](#).

1.9.35. RHSA-2025:9750 - OpenShift Container Platform 4.19.2 image release, bug fix, and security update advisory

Issued: 01 July 2025

OpenShift Container Platform release 4.19.2, which includes security updates, is now available. The list of bug fixes that are included in the update is documented in the [RHSA-2025:9750](#) advisory. The RPM packages that are included in the update are provided by the [RHSA-2025:9751](#) advisory.

Space precluded documenting all of the container images for this release in the advisory.

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

```
$ oc adm release info 4.19.2 --pullspecs
```

1.9.35.1. Bug fixes

- Previously, the installation program checked only the first compute machine pool entry in the install configuration to determine whether to disable the Machine Config Operator (MCO) boot image management feature. If multiple compute pools were specified (Amazon Web Services (AWS) edge nodes are the only supported scenario), but another compute machine pool had a custom Amazon Machine Image (AMI), the installation program would not disable the MCO boot image management and the custom AMI would be overwritten by the MCO. With this release, the installation program checks all compute machine pool entries and if a custom image is found, the MCO boot image management is disabled. ([OCPBUGS-58060](#))
- Previously, if a user specified a custom boot image for Amazon Web Services (AWS) or the Google Cloud, the Machine Config Operator (MCO) would overwrite it with the default boot image during installation. With this release, a manifest generation was added for MCO configuration which disables the default boot image during installation if a custom image is specified. ([OCPBUGS-57796](#))
- Previously, a validation issue within **oc-mirror** plugin caused the command to reject the **file://.** reference. Users attempting to use **file://.** for a content path received an error message stating **content filepath is tainted**. With this release, the **oc-mirror** plugin properly validates the '.' directory reference. ([OCPBUGS-57786](#))
- Previously, the **oc-mirror v2** command was not using the correct filtered catalog during its operations, which led to errors such as including more operators than specified in the configuration, and trying to connect to the catalog registry during disk-to-mirror workflows even in air-gapped environments. With this release, the correct filtered catalog is used. ([OCPBUGS-57784](#))
- Previously, the Red Hat OpenShift Lightspeed UI would disappear when the **Create Project** modal was opened or when modals on the Networking pages were triggered. This was due to the modals using the **useModal** hook causing the modals to overwrite each other. With this release, the modals no longer overwrite each other allowing multiple UI elements to be displayed simultaneously. ([OCPBUGS-57755](#))
- Previously, the HAProxy configuration used the **/version** endpoint for health checks causing unreliable health checks to be generated. With this release, the liveness probe is customized to use **/livez?exclude=etcd&exclude=log** on IBM Cloud for more accurate health checks avoiding disruptions due to inappropriate probe configurations on Hypershift, while retaining **/version** for other platforms. ([OCPBUGS-57485](#))
- Previously, the installer failed when AWS credentials were not found and the survey was attempting AWS regions preventing users from creating the **install-config** file. With this release, the installer no longer fails when AWS credentials are not set, allowing users to input them during the survey. ([OCPBUGS-57394](#))

- Previously, cloning a persistent volume claim (PVC) in the web console resulted in an error due to an unsupported unit **B** for storage size. Because of this, users encountered errors when cloning the Red Hat OpenShift console PVC due to incorrect parsing of the storage size unit. With this release, support for **B** as unit of storage size has been removed from the Red Hat OpenShift console PVC. ([OCPBUGS-57391](#))
- Previously, Operator Lifecycle Manager (OLM) v1 was used to install Operators with the **olm.maxOpenShiftVersion** set to **4.19**. Due to an issue with the OLM v1 parsing logic for floating-point formatted **olm.maxOpenShiftVersion`values, the system failed to prevent upgrades to OpenShift Container Platform. With this release, the parsing logic for `olm.maxOpenShiftVersion** has been corrected preventing upgrades to OpenShift Container Platform when Operators with **olm.maxOpenShiftVersion:4.19** are installed. ([OCPBUGS-56852](#))
- Previously, one of the **keepalived** health check scripts was failing due to missing permissions. This could result in the incorrect assignment of the ingress Virtual IP Address (VIP) in environments using shared ingress services. With this release, the necessary permission was added back to the container so the health check now works correctly. ([OCPBUGS-56623](#))

1.9.35.2. Updating

To update an OpenShift Container Platform 4.19 cluster to this latest release, see [Updating a cluster using the CLI](#).

1.9.36. RHSA-2025:9278 - OpenShift Container Platform 4.19.1 image release, bug fix, and security update advisory

Issued: 24 June 2025

OpenShift Container Platform release 4.19.1, which includes security updates, is now available. The list of bug fixes that are included in the update is documented in the [RHSA-2025:9278](#) advisory. The RPM packages that are included in the update are provided by the [RHSA-2025:9279](#) advisory.

Space precluded documenting all of the container images for this release in the advisory.

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

```
$ oc adm release info 4.19.1 --pullspecs
```

1.9.36.1. Bug fixes

- Previously, when you added vCenter cloud credentials for the post installation of the Assisted Installer, a bug was triggered because of an invalid **ConfigMap** object for the cloud provider configuration. As a result, a **missing vcenterplaceholder** error was displayed. With this release, the **ConfigMap** data is correct, and the error is not displayed. ([OCPBUGS-57384](#))
- Previously, a network issue during an API call in a cluster caused a timeout in Operator Lifecycle Manager (OLM) Classic. As a consequence, Operator installations often failed because of timeout issues. With this release, the catalog cache refresh interval is updated to resolve timeout issues. As a result, the likelihood of Operator installation timeouts is reduced. ([OCPBUGS-57352](#))
- Previously, Operator group reconciliation in Operator Lifecycle Manager (OLM) Classic triggered unnecessary **ClusterRole** updates because of the changing order of aggregation rule

selectors. As a result, unnecessary API server writes occurred. With this release, a bug fix ensures the deterministic order of a **ClusterRoleSelectors** array in the aggregation rule, reducing unnecessary API server writes and improving cluster stability. ([OCPBUGS-57279](#))

- Previously, ignoring the **AdditionalTrustBundlePolicy** setting in the assisted-service's installation configuration led to Federal Information Processing Standard (FIPS) and other installation configuration overrides. With this release, the installation configuration includes an **AdditionalTrustBundlePolicy** field, which you can set to ensure that FIPS and other installation configuration overrides function as intended. ([OCPBUGS-57208](#))
- Previously, the authentication process for the **/metrics** endpoint was missing a token review check and caused unauthorized requests. As a result, the OpenShift Container Platform console was prone to **TargetDown** alerts. With this release, the token review for unauthorized requests occurs with the user token provided in the request context. As a result, unauthorized requests to the OpenShift Container Platform console do not cause **TargetDown** alerts. ([OCPBUGS-57180](#))
- Previously, the **Started** column was hidden when screen size was reduced. As a consequence, the **VirtualizedTable** component malfunctioned because of a missing sort function, and the table sorting functionality was affected on the **PipelineRun** list pages. With this release, the table component handles missing sort functions correctly for reduced screen sizes. ([OCPBUGS-57110](#))
- Previously, if you configured a masthead logo for a theme but used the default settings for the rest of the theme, the logo shown on the user interface was inconsistent. With this release, the masthead logo displays a default option for both light and dark themes, improving the interface consistency. ([OCPBUGS-57054](#))
- Previously, the cluster installation failed because of an invalid security group configuration for a Network Load Balancer (NLB). This failure prevented the traffic from both primary subnets for bootstrapping. With this release, the security group allows traffic from both primary subnets for bootstrapping, and the cluster installation does not fail because of security group restrictions on additional primary subnets. ([OCPBUGS-57039](#))
- Previously, users without project access saw an incomplete roles list on the **Roles** page because of improper API group access. With this release, users without project access cannot see an incomplete roles list on the **Roles** page. ([OCPBUGS-56987](#))
- Previously, the **node-image create** command modified directory permissions and caused user directories to lose original permissions during the operation. With this release, the **node-image create** command preserves file permissions during the file-copying process by using the **rsync** tool, and ensures that user directories maintain original permissions during the operation. ([OCPBUGS-56905](#))
- Previously, the **delete** keyword in image names was allowed in the **ImageSetConfiguration** file, which is not supported. As a consequence, users encountered errors while mirroring images. With this release, the error for image names ending with **delete** in the **ImageSetConfiguration** file has been removed. As a result, users can now successfully mirror images with names ending in **delete**. ([OCPBUGS-56798](#))
- Previously, the user interface displayed incorrect alert severity icons for information alerts in the **Observe Alerting** field. With this release, the alert severity icons match in the **Observe Alerting** field. As a result, alert icons match consistently, reducing potential confusion for users. ([OCPBUGS-56470](#))
- Previously, if you used an unauthorized access configuration file in the **oc-mirror** command, an

Unauthorized error was displayed when you synchronized your image sets. With this release, the Docker configuration is updated to use a custom authorization file for authentication. You can successfully synchronize your image sets without encountering the **Unauthorized** error. ([OCBUGS-55701](#))

1.9.36.2. Updating

To update an OpenShift Container Platform 4.19 cluster to this latest release, see [Updating a cluster using the CLI](#).

1.9.37. RHSA-2024:11038 - OpenShift Container Platform 4.19.0 image release, bug fix, and security update advisory

Issued: 17 June 2025

OpenShift Container Platform release 4.19.0, which includes security updates, is now available. The list of bug fixes that are included in the update is documented in the [RHSA-2024:11038](#) advisory. The RPM packages that are included in the update are provided by the [RHEA-2025:2851](#) advisory.

Space precluded documenting all of the container images for this release in the advisory.

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

```
$ oc adm release info 4.19.0 --pullspecs
```

1.9.37.1. Updating

To update an OpenShift Container Platform 4.19 cluster to this latest release, see [Updating a cluster using the CLI](#).

CHAPTER 2. ADDITIONAL RELEASE NOTES

Release notes for additional related components and products not included in the core [OpenShift Container Platform 4.19 release notes](#) are available in the following documentation.



IMPORTANT

The following release notes are for downstream Red Hat products only; upstream or community release notes for related products are not included.

A

[AWS Load Balancer Operator](#)

B

[Builds for Red Hat OpenShift](#)

C

[cert-manager Operator for Red Hat OpenShift](#)
[Cluster Observability Operator \(COO\)](#)

[Compliance Operator](#)

[Custom Metrics Autoscaler Operator](#)

D

[Red Hat Developer Hub Operator](#)

E

[External DNS Operator](#)
[External Secrets Operator for Red Hat OpenShift](#)

F

[File Integrity Operator](#)

H

[Hosted control planes](#)

K

[Kube Descheduler Operator](#)
[Red Hat build of Kueue](#)

L

[Leader Worker Set Operator](#)
[Logging](#)

M

[Migration Toolkit for Containers \(MTC\)](#)

N

[Network Observability Operator](#)
[Network-bound Disk Encryption \(NBDE\) Tang Server Operator](#)

O

[OpenShift API for Data Protection \(OADP\)](#)
[Red Hat OpenShift Dev Spaces](#)

[Red Hat OpenShift Distributed Tracing Platform](#)

[Red Hat OpenShift GitOps](#)

[Red Hat OpenShift Local \(Upstream CRC documentation\)](#)

[Red Hat OpenShift Pipelines](#)

[OpenShift sandboxed containers](#)

[Red Hat OpenShift Serverless](#)

[Red Hat OpenShift Service Mesh 2.x](#)

[Red Hat OpenShift Service Mesh 3.x](#)

[Red Hat OpenShift support for Windows Containers](#)

[Red Hat OpenShift Virtualization](#)

[Red Hat build of OpenTelemetry](#)

P

[Power monitoring for Red Hat OpenShift](#)

R

[Run Once Duration Override Operator](#)

S

[Secondary Scheduler Operator for Red Hat OpenShift](#)
[Security Profiles Operator](#)

Z

[Zero Trust Workload Identity Manager](#)