OpenShift Container Platform 4.17

Release notes

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

Last Updated: 2026-03-24

OpenShift Container Platform 4.17 Release notes

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

Legal Notice

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. OPENSHIFT CONTAINER PLATFORM 4.17 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-2024:3718) is now available. This release uses Kubernetes 1.30 with CRI-O runtime. New features, changes, and known issues that pertain to OpenShift Container Platform 4.17 are included in this topic.

OpenShift Container Platform 4.17 clusters are available at https://console.redhat.com/openshift. With the Red Hat OpenShift Cluster Manager application for OpenShift Container Platform, you can deploy OpenShift Container Platform clusters to either on-premises or cloud environments.

OpenShift Container Platform 4.17 is supported on Red Hat Enterprise Linux (RHEL) 8.8 and a later version of Red Hat Enterprise Linux (RHEL) 8 that is released before End of Life of OpenShift Container Platform 4.17. OpenShift Container Platform 4.17 is also supported on Red Hat Enterprise Linux CoreOS (RHCOS) 4.17. To understand RHEL versions used by RHCOS, see RHEL Versions Utilized by Red Hat Enterprise Linux CoreOS (RHCOS) and OpenShift Container Platform (Knowledgebase article).

You must use RHCOS machines for the control plane, and you can use either RHCOS or RHEL for compute machines. RHEL machines are deprecated in OpenShift Container Platform 4.16 and will be removed in a future release.

The support lifecycle for odd-numbered releases, such as OpenShift Container Platform 4.17, 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. OPENSHIFT CONTAINER PLATFORM LAYERED AND DEPENDENT COMPONENT SUPPORT AND COMPATIBILITY

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

1.3. NEW FEATURES AND ENHANCEMENTS

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

1.3.1. Cluster Resource Override Admission Operator

1.3.1.1. Moving the Cluster Resource Override Operator

By default, the installation process creates a Cluster Resource Override Operator pod on a worker node and two Cluster Resource Override pods on control plane nodes. You can move these pods to other nodes, such as an infrastructure node, as needed. For more information, see Moving the Cluster Resource Override Operator pods.

1.3.1.2. Cluster Resource Override Operator pod is owned by a deployment object

The Cluster Resource Override Operator pod is now owned by a deployment object. Previously, the Operator was owned by a daemon set object. Using a deployment for the Operator addresses a number of issues, including additional security and add the ability to run the pods on worker nodes.

1.3.2. Extensions (OLM v1)

1.3.2.1. Operator Lifecycle Manager (OLM) v1 documentation moved to new Extensions guide (Technology Preview)

The documentation for OLM v1, which has been in Technology Preview starting in OpenShift Container Platform 4.14, is now moved and reworked as a separate guide called Extensions. Previously, OLM v1 documentation was a subsection of the existing Operators guide, which otherwise documents the existing OLM feature set.

The updated location and guide name reflect a more focused documentation experience and aims to differentiate between OLM v1 and existing OLM.

1.3.2.2. OLM v1 Technology Preview features

This Technology Preview phase of OLM v1 introduces the following features:

Custom resource definition (CRD) upgrade safety

When you update a CRD that is provided by a cluster extension, OLM v1 now runs a CRD upgrade safety preflight check to ensure backwards compatibility with previous versions of that CRD. The CRD update must pass the validation checks before the change is allowed to progress on a cluster. For more information, see Custom resource definition (CRD) upgrade safety .

Single object ownership for cluster extensions

In OLM v1, a Kubernetes object can only be owned by a single **ClusterExtension** object at a time. This ensures that objects within an OpenShift Container Platform cluster are managed consistently and prevents conflicts between multiple cluster extensions attempting to control the same object. For more information, see Object ownership for cluster extensions .

Enhanced security

OLM v1 now requires a dedicated service account for installing, updating, and managing cluster extensions. Additionally, catalogd uses HTTPS encryption to secure catalog server responses. For more information, see Creating a service account to manage cluster extensions .

Improved status conditions

In this release, OLM v1 includes improved status conditions and error messaging via the **ClusterExtension** API.

1.3.2.3. OLM v1 supported extensions and known issue

Currently, Operator Lifecycle Manager (OLM) v1 supports installing cluster extensions that meet all of the following criteria:

- The extension must use the **registry+v1** bundle format introduced in existing OLM.
- The extension must support installation via the **AllNamespaces** install mode.
- The extension must not use webhooks.
- The extension must not declare dependencies by using any of the following file-based catalog properties:
  - **olm.gvk.required**
  - **olm.package.required**
  - **olm.constraint**

OLM v1 checks that the extension you want to install meets these constraints. If the extension that you want to install does not meet these constraints, an error message is printed in the cluster extension's conditions.

Operator Lifecycle Manager (OLM) v1 does not support the **OperatorConditions** API introduced in existing OLM.

If an extension relies on only the **OperatorConditions** API to manage updates, the extension might not install correctly. Most extensions that rely on this API fail at start time, but some might fail during reconciliation.

As a workaround, you can pin your extension to a specific version. When you want to update your extension, consult the extension's documentation to find out when it is safe to pin the extension to a new version.

IMPORTANT

Currently, Operator Lifecycle Manager (OLM) v1 cannot authenticate private registries, such as the Red Hat-provided Operator catalogs. This is a known issue. As a result, the OLM v1 procedures that rely on having the Red Hat Operators catalog installed do not work. (OCPBUGS-36364)

1.3.3. Edge computing

1.3.3.1. Managing host firmware settings with GitOps ZTP

You can now configure host firmware settings for managed clusters that you deploy with GitOps ZTP. You save host profile YAML files alongside **SiteConfig** custom resources (CRs) that you use to deploy the managed clusters. GitOps ZTP uses the host profiles to configure firmware settings in the managed cluster hosts during deployment. On the hub cluster, you can use **FirmwareSchema** CRs to discover managed cluster host firmware schema, and **HostFirmwareSettings** CRs and retrieve managed clusters firmware settings.

For more information, see Managing host firmware settings with GitOps ZTP .

1.3.3.2. Image-based upgrade enhancements

With this release, the image-based upgrade introduces the following enhancements:

- Simplifies the upgrade process for a large group of managed clusters by adding the **ImageBasedGroupUpgrade** API on the hub
- Labels the managed clusters for action completion when using the **ImageBasedGroupUpgrade** API
- Improves seed cluster validation before the seed image generation
- Automatically cleans up the container storage disk if usage reaches a certain threshold on the managed clusters
- Adds comprehensive event history in the new **status.history** field of the **ImageBasedUpgrade** CR

For more information about the **ImageBasedGroupUpgrade** API, see Managing the image-based upgrade at scale using the ImageBasedGroupUpgrade CR on the hub.

1.3.3.3. Disk encryption with TPM and PCR protection (Technology Preview)

With this release, you can enable disk encryption with Trusted Platform Module (TPM) and Platform Configuration Registers (PCRs) protection. You can use the **diskEncryption** field in the **SiteConfig** custom resource (CR) to configure the disk encryption. Configuring the **SiteConfig** CR enables disk encryption at the time of cluster installation.

For more information, see Enabling disk encryption with TPM and PCR protection .

1.3.3.4. IPsec encryption for multi-node clusters using GitOps ZTP and SiteConfig resources

You can now enable IPsec encryption in managed multi-node clusters that you deploy with GitOps ZTP and Red Hat Advanced Cluster Management (RHACM). You can encrypt traffic between the managed cluster and IPsec endpoints external to the managed cluster. All network traffic between nodes on the OVN-Kubernetes cluster network is encrypted with IPsec in Transport mode.

For more information, see Configuring IPsec encryption for multi-node clusters using GitOps ZTP and SiteConfig resources.

1.3.3.5. Image-based installation for single-node OpenShift clusters

Image-based installations streamline the installation and deployment process for single-node OpenShift clusters by significantly reducing installation and deployment times.

Using an image-based workflow, you can preinstall instances of single-node OpenShift on target hosts. These preinstalled hosts can be rapidly reconfigured and deployed at the far edge of the network, including in disconnected environments, with minimal intervention.

For more information, see Understanding an image-based installation and deployment for single-node OpenShift clusters.

1.3.4. IBM Z and IBM LinuxONE

With this release, IBM Z® and IBM® LinuxONE are now compatible with OpenShift Container Platform 4.17. You can perform the installation with z/VM, LPAR, or Red Hat Enterprise Linux (RHEL) Kernel-based Virtual Machine (KVM). For installation instructions, see Preparing to install on IBM Z and IBM LinuxONE.

IMPORTANT

Compute nodes must run Red Hat Enterprise Linux CoreOS (RHCOS).

1.3.4.1. IBM Z and IBM LinuxONE notable enhancements

The IBM Z® and IBM® LinuxONE release on OpenShift Container Platform 4.17 adds improvements and new capabilities to OpenShift Container Platform components and concepts.

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

- CPU manager
- Multiarch Tuning Operator
- Non-volatile memory express (NVMe) support for LPAR
- Secondary Scheduler Operator
- Tuning etcd latency tolerances

1.3.5. IBM Power

IBM Power® is now compatible with OpenShift Container Platform 4.17. For installation instructions, see the following documentation:

- Installing a cluster on IBM Power®
- Installing a cluster on IBM Power® in a restricted network

IMPORTANT

Compute nodes must run Red Hat Enterprise Linux CoreOS (RHCOS).

1.3.5.1. IBM Power notable enhancements

The IBM Power® release on OpenShift Container Platform 4.17 adds improvements and new capabilities to OpenShift Container Platform components.

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

- Multiarch Tuning Operator
- Secondary Scheduler Operator
- Tuning etcd latency tolerances
- Installer Provisioned Infrastructure for IBM PowerVS - move to Cluster API

1.3.6. 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. OpenShift Container Platform features

Feature	IBM Power®	IBM Z® and IBM® LinuxONE
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
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

Feature	IBM Power®	IBM Z® and IBM® LinuxONE
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
Operator API	Supported	Supported
OpenShift Virtualization	Unsupported	Unsupported
OVN-Kubernetes, including IPsec encryption	Supported	Supported
PodDisruptionBudget	Supported	Supported
Precision Time Protocol (PTP) hardware	Unsupported	Unsupported

Feature	IBM Power®	IBM Z® and IBM® LinuxONE
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	Supported
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 1.2. 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.

Table 1.3. 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
Cost Management Metrics Operator	Supported	Supported
File Integrity Operator	Supported	Supported
HyperShift Operator	Technology Preview	Technology Preview
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
Multiarch Tuning Operator	Supported	Supported
Network Observability Operator	Supported	Supported
NFD Operator	Supported	Supported
NMState Operator	Supported	Supported
OpenShift Elasticsearch Operator	Supported	Supported
Secondary Scheduler Operator	Supported	Supported
Vertical Pod Autoscaler Operator	Supported	Supported

Table 1.4. 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 1.5. CSI Volumes

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

1.3.7. Insights Operator

The Insights Operator now collects more OpenShift Container Platform container log data from namespaces prefixed with either the **openshift-** or **kube-** prefix and generates recommendations much faster. Enhancements have also been made to give you more flexibility in how the data to be collected gets defined for your service.

1.3.7.1. Rapid Recommendations

This release introduces a new feature called Rapid Recommendations, which provides a more dynamic and version-independent mechanism for remotely configuring the rules that determine which data the Insights Operator collects.

Rapid Recommendations builds on the existing conditional data gathering mechanism. The Insights Operator connects to a secure remote endpoint service running on **/console.redhat.com** to retrieve definitions that contain the rules for determining which container log messages are filtered and collected by Red Hat.

The conditional data-gathering definitions, also referred to as rules, get configured through an attribute named **conditionalGathererEndpoint** in the **pod.yml** configuration file.

```
conditionalGathererEndpoint: https://console.redhat.com/api/gathering/v2/%s/gathering_rules
```

NOTE

Previously, the rules for determining the data that the Insights Operator collects were hard-coded and tied to the corresponding OpenShift Container Platform version.

The preconfigured endpoint URL now provides a placeholder (**%s**) for defining a target version of OpenShift Container Platform.

1.3.7.2. More data collected and recommendations added

The Insights Operator now gathers more data to detect the following scenarios, which other applications can use to generate remedial recommendations to proactively manage your OpenShift Container Platform deployments:

- Detects pods and namespaces that use the deprecated OpenShift SDN CNI plugin and generates a recommendation for the possible actions you should take depending on the data collected from your deployment.
- Collects custom resource definitions (CRD) from RHOSP.
- Collects the **haproxy_exporter_server_threshold** metric to detect the problem and remediation reported in OCPBUGS-36687.
- Collects data to detect custom Prometheus Alertmanager instances that are not in the **openshift-monitoring** namespace because they could potentially impact the management of corresponding resources.
- Detects the upcoming expiry of the default Ingress Controller expiration certificate, which other applications and services can use to generate recommendations to renew the certificate before the expiry date.
    - Before this update, the Insights Operator gathered information about all Ingress Controller certificates, including their **NotBefore** and **NotAfter** dates. This data is now compiled into a **JSON** file located at **aggregated/ingress_controllers_certs.json** for easier monitoring of certificate validity across the cluster. (OCPBUGS-35727)

1.3.8. Installation and update

1.3.8.1. User-defined labels and tags for GCP

With this update, the user-defined labels and tags for Google Cloud is Generally Available.

For more information, see Managing user-defined labels and tags for GCP.

1.3.8.2. Installing a cluster on Nutanix with compute machines using GPUs

With this update, you can install a cluster on Nutanix with compute machines that use GPUs for processing. You attach a GPU to compute nodes with the **compute.platform.nutanix.gpus** parameters in the **install-config.yaml** file.

For more information, see Installation configuration parameters for Nutanix.

1.3.8.3. Installing a cluster on Nutanix with compute nodes using multiple disks

With this update, you can install a cluster on Nutanix with compute machines that have multiple disks attached to them. You attach multiple disks to compute nodes with the **compute.platform.nutanix.dataDisks** parameters in the **install-config.yaml** file.

For more information, see Installation configuration parameters for Nutanix.

1.3.8.4. Installing a cluster on Azure in the Central Spain region

You can now install an OpenShift Container Platform cluster on Azure in the Central Spain region, **spaincentral**.

For more information, see Supported Azure regions.

1.3.8.5. Installing a cluster with the support for configuring multi-architecture compute machines

With this release, you can install an Amazon Web Services (AWS) cluster and Google Cloud cluster with the support for configuring multi-architecture compute machines. While installing the cluster, you can specify different CPU architectures for the control plane and compute machines in the following ways:

- 64-bit x86 compute machines and 64-bit ARM control plane machines
- 64-bit ARM compute machines and 64-bit x86 control plane machines

An OpenShift Container Platform cluster with multi-architecture compute machines supports compute machines with different architectures. For more information, see the following documentation:

- Installing a cluster with multi-architecture support (AWS: Installer-provisioned infrastructure)
- Installing a cluster with multi-architecture support (AWS: User-provisioned infrastructure)
- Installing a cluster with multi-architecture support (Google Cloud)

1.3.8.6. Installing a cluster on Nutanix with Flow Virtual Networking

In OpenShift Container Platform 4.17, you can install a cluster on Nutanix with Flow Virtual Networking enabled. Flow Virtual Networking is a software-defined networking solution for Nutanix AHV clusters that provides multi-tenant isolation, self-service provisioning, and IP address preservation using VPCs, subnets, and other virtual components that are separate from the physical network. To perform this installation, enable Flow Virtual Networking in your Nutanix AHV environment before installing.

For more information, see Flow Virtual Networking overview .

1.3.8.7. Cluster API replaces Terraform for Microsoft Azure installations

In OpenShift Container Platform 4.17, the installation program uses Cluster API instead of Terraform to provision cluster infrastructure during installations on Azure.

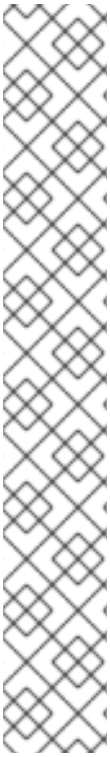

NOTE

With the replacement of Terraform, the following permissions are required if you use a service principal with limited privileges:

- **Microsoft.Network/loadBalancers/inboundNatRules/read**
- **Microsoft.Network/loadBalancers/inboundNatRules/write**
- **Microsoft.Network/loadBalancers/inboundNatRules/join/action**
- **Microsoft.Network/loadBalancers/inboundNatRules/delete**
- **Microsoft.Network/routeTables/read**
- **Microsoft.Network/routeTables/write**
- **Microsoft.Network/routeTables/join/action**

For more information on required permissions, see Required Azure permissions for installer-provisioned infrastructure.

1.3.8.8. Installing a cluster on Google Cloud by using an existing service account

With this update, you can install a cluster on Google Cloud by using an existing service account, allowing you to minimize the permissions that you grant to the service account the installation program uses. You can specify this service account in the **compute.platform.gcp.serviceAccount** and **controlPlane.platform.gcp.serviceAccount** parameters in the **install-config.yaml** file. For more information, see Available installation configuration parameters for Google Cloud .

1.3.8.9. Installing a cluster on AWS by using an existing IAM profile

With this release, you can install OpenShift Container Platform on Amazon Web Services (AWS) by using an existing identity and access management (IAM) instance profile. For more information, see Optional AWS configuration parameters.

1.3.8.10. Installing a cluster on Google Cloud using the N4 machine series

With this release, you can deploy a cluster on Google Cloud using the N4 machine series for compute or control plane machines. The supported disk type of N4 machine series is **hyperdisk-balanced**. For more information, see Installation configuration parameters for GCP.

1.3.8.11. Cluster API replaces Terraform for Google Cloud installations

With this release, the installation program uses Cluster API instead of Terraform to provision cluster infrastructure during installations on Google Cloud.

1.3.8.12. Three-node cluster support for RHOSP

Deploying a three-node cluster on installer-provisioned infrastructure is now supported on Red Hat OpenStack Platform (RHOSP).

For more information, see Installing a three-node cluster on OpenStack .

1.3.8.13. Deploying Red Hat OpenStack Platform (RHOSP) with root volume and etcd on local disk (Generally Available)

You can now move etcd from a root volume (Cinder) to a dedicated ephemeral local disk as a Day 2 deployment with this generally available feature.

For more information, see Deploying on OpenStack with rootVolume and etcd on local disk .

1.3.9. Operator lifecycle

1.3.9.1. New guide location and release notes section for Operator Lifecycle Manager (OLM) v1 (Technology Preview)

For release notes about OLM v1 in OpenShift Container Platform 4.17 and later, including its new guide location starting this release, see the new features and enhancements section for Extensions (OLM v1).

This "Operator lifecycle" section will continue to document new features and enhancements for existing OLM in future releases.

1.3.9.2. Web console warnings for deprecated Operators

When deprecated packages, channels, or versions are defined for Operators in a catalog, the OpenShift Container Platform web console now displays warning badges for the affected elements of the Operator, including any custom deprecation messages, on both the pre- and post-installation pages of the OperatorHub.

For more information on the deprecation schema for Operator catalogs, see see Operator Framework packaging format → Schemas → olm.deprecations schema.

1.3.10. Operator development

1.3.10.1. Token authentication for Operators on cloud providers: GCP Workload Identity

With this release, Operators managed by Operator Lifecycle Manager (OLM) can support token authentication when running on Google Cloud clusters configured for GCP Workload Identity. Updates to the Cloud Credential Operator (CCO) enable semi-automated provisioning of certain short-term credentials, provided that the Operator author has enabled their Operator to support GCP Workload Identity.

For more information, see CCO-based workflow for OLM-managed Operators with GCP Workload Identity.

1.3.11. OpenShift CLI (oc)

1.3.11.1. oc-mirror to include the HyperShift KubeVirt CoreOS container

With this release, oc-mirror now includes the Red Hat Enterprise Linux CoreOS (RHCOS) image for the HyperShift KubeVirt provider when mirroring the OpenShift Container Platform release payload.

The **kubeVirtContainer** flag, which is set to false by default, must be set to **true** in the **imageSetConfig.yaml** file to extract the KubeVirt Container RHCOS. This ensures support for disconnected environments by including the required image for KubeVirt virtual machines acting as nodes for hosted clusters.

1.3.12. Machine Config Operator

1.3.12.1. Control plane TLS security profiles supported by the MCO

The Machine Config Operator (MCO) and Machine Config Server now use the TLS security profile that is configured for the control plane components. For more information, see Configuring the TLS security profile for the control plane.

1.3.12.2. Updated boot images for AWS now supported (Technology Preview)

Updated boot images are now supported as a Technology Preview feature for Amazon Web Services (AWS) clusters. This feature allows you configure your cluster to update the node boot image whenever you update your cluster. By default, the boot image in your cluster is not updated along with your cluster. For more information, see Updated boot images.

1.3.12.3. Updated boot images for GCP clusters promoted to GA

Updated boot images has been promoted to GA for Google Cloud Platform (GCP) clusters. For more information, see Updated boot images.

1.3.12.4. Node disruption policies promoted to GA

The node disruption policies feature has been promoted to GA. A node disruption policy allows you to define a set of Ignition config objects changes that would require little or no disruption to your workloads. For more information, see Using node disruption policies to minimize disruption from machine config changes.

1.3.13. Machine management

1.3.13.1. Supporting AWS Placement Group Partition Number

This release introduces the **placementGroupPartition** field for OpenShift Container Platform **MachineSet** on Amazon Web Services (AWS). With this feature, you can specify a partition number within an existing placement group, enabling precise instance allocation and improved fault tolerance. For example, see Assigning machines to placement groups for Elastic Fabric Adapter instances by using machine sets.

1.3.13.2. Configuring Capacity Reservation by using machine sets

OpenShift Container Platform release 4.17 introduces support for on-demand Capacity Reservation with Capacity Reservation groups on Microsoft Azure clusters. For more information, see *Configuring Capacity Reservation by using machine sets* for compute or 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 0.27.0

- Prometheus Operator to 0.75.2
- Prometheus to 2.53.1
- prom-label-proxy to 0.11.0
- kube-state-metrics to 2.13.0
- node-exporter to 1.8.2
- Thanos to 0.35.1

1.3.14.2. Changes to alerting rules

NOTE

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

- Added the **PrometheusKubernetesListWatchFailures** alert to warn users about Prometheus and Kubernetes API failures, such as unreachable API and permissions issues, which can lead into silent service discovery failures.

1.3.14.3. Updated Prometheus to tolerate jitters at scrape time for user-defined projects

With this update, the Prometheus configuration for monitoring for user-defined projects now tolerates jitters at scrape time. This update optimizes data compression for monitoring deployments that show sub-optimal chunk compression for data storage, which reduces the disk space used by the time series database in these deployments.

1.3.14.4. Network Observability Operator

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

1.3.15. Nodes

1.3.15.1. New CRIO command behavior

Beginning in OpenShift Container Platform 4.17, when a node is rebooted, the **crio wipe** command checks that the CRI-O binary exited cleanly. Those images that did not exit cleanly are targeted as corrupted and removed. This behavior prevents CRI-O from failing to start due to half-pulled images or other unsynced files. In OpenShift Container Platform 4.15 and 4.16, the **crio wipe** command removed all images when a node was rebooted. The **crio wipe** command's new behavior increases efficiency while still reducing the risk of image corruption when a node is rebooted.

1.3.15.2. New flags added for must-gather command

OpenShift Container Platform release 4.17 adds two new flags for use with the **oc adm must-gather** command to limit the timespan of the information gathered. Only one of the following flags can be used at a time. Plugins are encouraged but not required to support these flags.

- **--since**: Only return logs newer than a relative duration, such as 5s, 2m, or 3h. Defaults to all logs.
- **--since-time**: Only return logs after a specific date, expressed in the RFC3339 format. Defaults to all logs.

For a full list of flags to use with the **oc adm must-gather command**, see Must-gather flags.

1.3.15.3. Linux user namespaces now supported for pods (Technology Preview)

OpenShift Container Platform release 4.17 adds support for deploying pods and containers into Linux user namespaces. Running pods and containers in individual user namespaces can mitigate several vulnerabilities that a compromised container can pose to other pods and the node itself. For more information, see Running pods in Linux user namespaces .

1.3.15.4. CRI-O metrics port now uses TLS

OpenShift Container Platform monitoring now uses a TLS-backed endpoint to fetch CRI-O container runtime metrics. These certificates are managed by the system and not the user. OpenShift Container Platform monitoring queries have been updated to the new port. For information on the certificates used by monitoring, see Monitoring and OpenShift Logging Operator component certificates .

1.3.15.5. Adding compute nodes to on-premise clusters

With this release, you can add compute nodes by using the OpenShift CLI (**oc**) to generate an ISO image, which can then be used to boot one or more nodes in your target cluster. This process can be used regardless of how you installed your cluster.

For more information, see Adding worker nodes to an on-premise cluster .

1.3.16. Postinstallation configuration

1.3.16.1. 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 .

1.3.17. Networking

1.3.17.1. Dual-NIC Intel E810 Logan Beach as PTP grandmaster clock

You can now configure **linuxptp** services as a grandmaster clock (T-GM) for dual Intel E810 Logan Beach network interface controllers (NICs). You can configure the **linuxptp** services as a T-GM for the following dual E810 NICs:

- Intel E810-XXVDA4T Westport Channel NICs
- Intel E810-CQDA2T Logan Beach NICs

The host system clock is synchronized from the NIC that is connected to the Global Navigation Satellite Systems (GNSS) time source. The second NIC is synced to the 1PPS timing output provided by the NIC that is connected to GNSS. For more information, see Configuring linuxptp services as a grandmaster

clock for dual E810 NICs.

1.3.17.2. Masquerade subnet change for new clusters

For OpenShift Container Platform 4.17 and later versions, clusters use 169.254.0.0/17 for IPv4 and fd69::/112 for IPv6 as the default masquerade subnet. These ranges should be avoided by users. For upgraded clusters, there is no change to the default masquerade subnet. For information on how to change the masquerade subnet as a Day 2 operation, see Configuring the OVN-Kubernetes masquerade subnet as a Day 2 operation

1.3.17.3. Enabling the SR-IOV network metrics exporter

With this release, you can query the Single Root I/O Virtualization (SR-IOV) virtual function (VF) metrics by using the OpenShift Container Platform web console to monitor the networking activity of the SR-IOV pods. When you query the SR-IOV VF metrics by using the web console, the SR-IOV network metrics exporter fetches and returns the VF network statistics along with the name and namespace of the pod that the VF is attached to.

For more information, see Enabling the SR-IOV network metrics exporter.

1.3.17.4. MetalLB updates for Border Gateway Protocol

With this release, MetalLB includes a new field for the Border Gateway Protocol (BGP) peer custom resource. You can use the **dynamicASN** field to detect the Autonomous System Number (ASN) to use for the remote end of a BGP session. This is an alternative to explicitly setting an ASN in the **spec.peerASN** field.

1.3.17.5. Microsoft Azure for the Kubernetes NMState Operator

Red Hat support exists for using the Kubernetes NMState Operator on Microsoft Azure but in a limited capacity. Support is limited to configuring DNS servers on your system as a postinstallation task.

For more information, see About the Kubernetes NMState Operator.

1.3.17.6. View metrics collected by the Kubernetes NMState Operator

The Kubernetes NMState Operator, **kubernetes-nmstate-operator**, can collect metrics from the **kubernetes_nmstate_features_applied** component and expose them as ready-to-use metrics. You can view these metrics by using the **Administrator** and **Developer** perspectives.

For more information, see About the Kubernetes NMState Operator.

1.3.17.7. New PTP fast events REST API version 2 available

A new PTP fast events O-RAN Release 3 compliant REST API version 2 is available. Now, you can develop PTP event consumer applications that receive host hardware PTP events directly from the PTP Operator-managed pod.

The PTP events REST API v1 and PTP events consumer application sidecar is deprecated.

NOTE

In O-RAN O-Cloud Notification API Specification for Event Consumers 3.0 , the resource is defined as a hierarchical path for the subsystem that produces the notifications. The PTP events REST API v2 does not have a global subscription for all lower hierarchy resources contained in the resource path. You subscribe consumer applications to the various available event types separately.

For more information, see Developing PTP event consumer applications with the REST API v2 .

1.3.17.8. Automatic leap seconds handling for PTP grandmaster clocks

The PTP Operator now automatically updates the leap second file by using Global Positioning System (GPS) announcements.

Leap second information is stored in an automatically generated **ConfigMap** resource named **leap-configmap** in the **openshift-ptp** namespace.

For more information, see Configuring dynamic leap seconds handling for PTP grandmaster clocks .

1.3.17.9. NIC partitioning for SR-IOV devices (Generally Available)

With this update, the ability to enable NIC partitioning for Single Root I/O Virtualization (SR-IOV) devices at install time is Generally Available.

For more information, see NIC partitioning for SR-IOV devices .

1.3.17.10. Host network settings for SR-IOV VFs (Generally Available)

With this update, the ability to update host network settings for Single Root I/O Virtualization (SR-IOV) network virtual functions in an existing cluster is Generally Available.

For more information, see Node network configuration policy for virtual functions .

1.3.17.11. User-defined network segmentation (Technology Preview)

With OpenShift Container Platform 4.17, users can create multiple networks and declare them as primary or secondary networks for their workloads through the technology preview of the **UserDefinedNetwork** (UDN) custom resource definition (CRD). With UDN, users can isolate namespaces without configuring and managing complex network policies.

For more information, see Understanding user-defined networks

1.3.17.12. CoreDNS update to version 1.11.3

OpenShift Container Platform 4.17 now includes CoreDNS version 1.11.3.

1.3.17.13. eBPF manager Operator (Tech Preview)

The eBPF manager Operator, available as a Technology Preview, allows you to securely deploy and manage eBPF programs. It facilitates the secure loading, unloading, modifying, and monitoring of eBPF programs in OpenShift Container Platform clusters. For more information about deploying the bpfman Operator, see About the eBPF Manager Operator.

1.3.17.14. eBPF program support for Ingress Node Firewall Operator (Technology Preview)

Secure management of eBPF programs for the Ingress Node Firewall Operator is available as a Technology Preview. Use of this feature requires installation of the eBPF manager Operator, also available as a Technology Preview. For more information, see Ingress Node Firewall Operator integration.

1.3.17.15. Changes to MetalLB

With this update, MetalLB uses **FRR-K8s** as the default backend. Previously, this was an optional feature available in Technology Preview. For more information, see Configuring the integration of MetalLB and FRR-K8s.

MetalLB also includes a new field for the Border Gateway Protocol (BGP) peer custom resource, **connectTime**. You can use this field to specify how long BGP waits between connection attempts to a neighbor. For more information, see About the BGP peer custom resource.

1.3.17.16. Exposing MTU for vfio-pci SR-IOV devices

With this release, maximum transmission unit (MTU) on virtual function using the **vfio-pci** driver is available in the network-status pod annotation, and inside the container.

For more information, see Exposing MTU for vfio-pci SR-IOV devices to pod .

1.3.17.17. MetalLB metrics naming update

With this release, the naming convention for MetalLB BGP and BFD metrics was updated:

- The naming for BGP metrics was updated from **metallb_bgp_<metric_name>** to **frrk8s_bgp_<metric_name>**.
- The naming for BFD metrics was updated from **metallb_bfd_<metric_name>** to **frrk8s_bfd_<metric_name>**.

To view all the metrics in the new format, see MetalLB metrics for BGP and BFD.

1.3.18. Registry

1.3.18.1. New chunkSizeMiB configuration parameter for S3 registry storage

A new, optional configuration parameter, **chunkSizeMiB**, is now available for deployments using S3 API-compatible backend storage. When configured, it determines the size of the multipart upload chunks for the S3 API. The default value is **10** MiB, with a minimum of **5** MiB.

For more information, see Image Registry Operator configuration parameters for AWS S3 .

1.3.19. Red Hat Enterprise Linux CoreOS (RHCOS)

1.3.19.1. RHCOS uses RHEL 9.4

RHCOS uses Red Hat Enterprise Linux (RHEL) 9.4 packages in OpenShift Container Platform 4.17. These packages ensure that your OpenShift Container Platform instance receives the latest fixes, features, enhancements, hardware support, and driver updates.

1.3.19.2. Support for the DNF package manager

With this release, you can now use DNF to install additional packages to your customized Red Hat Enterprise Linux CoreOS (RHCOS) builds. For more information, see Red Hat Enterprise Linux CoreOS (RHCOS) image layering.

1.3.20. Storage

1.3.20.1. AWS EFS CSI storage usage metrics is generally available

Amazon Web Services (AWS) Elastic File Service (EFS) usage metrics allow you to monitor how much space is used by EFS volumes. This feature is generally available.

IMPORTANT

Turning on these metrics can lead to performance degradation because the CSI driver walks through the whole volume. Therefore, this option is disabled by default. Administrators must explicitly enable this feature.

For more information, see AWS EFS storage CSI usage metrics .

1.3.20.2. Preventing unauthorized volume mode conversion is generally available

Previously, there was no validation of whether the mode of an original volume (filesystem or raw block), whose snapshot was taken, matches the mode of a newly created volume. This presented a security gap that could allow malicious users to potentially exploit an as-yet-unknown vulnerability in the host operating system.

Nevertheless, some users have a legitimate need to perform such conversions. This feature allows cluster administrators to provide these rights (ability to perform update or patch operations on **VolumeSnapshotContents objects**) only to trusted users or applications, such as backup vendors.

To convert a volume mode, an authorized user needs to change **snapshot.storage.kubernetes.io/allow-volume-mode-change: "true"** for VolumeSnapshotContent of the snapshot source.

This feature is supported as generally available.

1.3.20.3. Automatic deletion of resources for GCP Filestore is generally available

In earlier versions of OpenShift Container Platform, when destroying a cluster, Google Compute Platform (GCP) Filestore Storage did not delete all of the cloud resources belonging to that cluster. This required manually deleting all of the persistent volume claims (PVCs) that used the Filestore storage class before destroying the cluster.

With OpenShift Container Platform 4.17, when destroying a cluster the OpenShift Container Platform installer should generally delete all of the cloud resources that belong to that cluster, and therefore manual deletion of PVCs should not be required. However, due to the special nature of the Google Compute Platform (GCP) Filestore resources, the automated cleanup process might not remove all of the resources in some rare cases. This feature is supported as generally available.

For more information, see Destroying clusters and GCP Filestore .

1.3.20.4. Azure File CSI supports snapshots (Technology Preview)

OpenShift Container Platform 4.17 introduces volume snapshot support for the Microsoft Azure File Container Storage Interface (CSI) Driver Operator. This capability is supported as a Technology Preview feature.

For more information, see CSI drivers supported by OpenShift Container Platform and CSI volume snapshots.

1.3.20.5. Multiple vCenter support for vSphere CSI (Technology Preview)

OpenShift Container Platform v4.17 introduces the ability to deploy OpenShift Container Platform across multiple vSphere clusters (vCenters). This feature is supported with Technology Preview status.

Multiple vCenters can only be configured during installation. The maximum number of supported vCenter clusters is three.

For more information, see Multiple vCenter support for vSphere CSI and Installation configuration parameters for vSphere.

1.3.20.6. Disabling and enabling storage on vSphere (Technology Preview)

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 is supported at the Technology Preview level.

For more information, see Disabling and enabling storage on vSphere .

1.3.20.7. RWX/RWO SELinux Mount (Developer Preview)

Pods might take a very long time to start when the volume contains a large number of files. To avoid SELinux labeling issues while keeping SELinux confining, you can enable the ReadWriteMany/ReadWriteOnce (RWX/RWO) SELinux Mount feature. Be advised that the RWX/RWO SELinux Mount feature is a Developer Preview feature. It is not supported by Red Hat, and you should not enable this feature set on production or clusters that you plan to maintain over time.

IMPORTANT

RWX/RWO SELinux Mount is a Developer Preview feature only. Developer Preview features are not supported by Red Hat in any way and are not functionally complete or production-ready. Do not use Developer Preview features for production or business-critical workloads. Developer Preview features provide early access to upcoming product features in advance of their possible inclusion in a Red Hat product offering, enabling customers to test functionality and provide feedback during the development process. These features might not have any documentation, are subject to change or removal at any time, and testing is limited. Red Hat might provide ways to submit feedback on Developer Preview features without an associated SLA.

For more information about the RWX/RWO SELinux Mount feature, including how to enable it, see RWX/RWO SELinux Mount feature Knowledge Centered Service article .

1.3.20.8. Migrating CNS volumes between datastores with cns-migration (Developer Preview)

In OpenShift Container Platform 4.17, if you are running out of space in your current datastore, or want to move to a more performant datastore, you can migrate volumes between datastores. Be advised that this feature is a Developer Preview feature. It is not supported by Red Hat.

IMPORTANT

Migrating CNS Volumes Between Datastores is a Developer Preview feature only. Developer Preview features are not supported by Red Hat in any way and are not functionally complete or production-ready. Do not use Developer Preview features for production or business-critical workloads. Developer Preview features provide early access to upcoming product features in advance of their possible inclusion in a Red Hat product offering, enabling customers to test functionality and provide feedback during the development process. These features might not have any documentation, are subject to change or removal at any time, and testing is limited. Red Hat might provide ways to submit feedback on Developer Preview features without an associated SLA.

For more information about cns-migration, see Moving CNS volumes between datastores .

1.3.20.9. Google Secret Manager is now available for the Secrets Store CSI Driver Operator (Technology Preview)

You can now use the Secrets Store CSI Driver Operator to mount secrets from Google Secret Manager to a Container Storage Interface (CSI) volume in OpenShift Container Platform. The Secrets Store CSI Driver Operator is available as a Technology Preview feature.

For the full list of available secrets store providers, see Secrets store providers .

For information about using the Secrets Store CSI Driver Operator to mount secrets from Google Secret Manager, see Mounting secrets from Google Secret Manager .

1.3.21. Scalability and performance

1.3.21.1. Kernel Module Management Operator

In this release, the firmware search path has been updated to copy the contents of the specified path into the path specified in **worker.setFirmwareClassPath** (default: **/var/lib/firmware**). For more information, see Example Module CR.

1.3.21.2. Node scaling for etcd

In this release, if your cluster is installed on a bare metal platform, you can scale a cluster control plane up to 5 nodes as a postinstallation task. The etcd Operator scales accordingly to account for the additional control plane nodes. For more information, see Node scaling for etcd.

1.3.21.3. Support for compute nodes with AMD EPYC Zen 4 CPUs

From release 4.17.10, you can use the **PerformanceProfile** custom resource (CR) to configure compute nodes on machines equipped with AMD EPYC Zen 4 CPUs, such as Genoa and Bergamo. Only single NUMA domain (NPS=1) configurations are supported. Per-pod power management is currently not supported on AMD.

1.3.22. Security

1.3.22.1. Automatic rotation of signer certificates

With this release, all **etcd** certificates originate from a new namespace: **openshift-etcd**. When a new signer certificate is close to its expiration date, the following actions occur:

1. An automatic rotation of the signer certificate activates.
2. The certificate bundle updates.
3. All certificates regenerate with the new signers.

Manual rotation of signer certificates is still supported by deleting the specific secret and waiting for the status pod rollout to complete.

1.3.22.2. Sigstore signature image verification

With this release, Technology Preview clusters use Sigstore signatures to verify images that were retrieved using a pull spec that references the **quay.io/openshift-release-dev/ocp-release** repository.

Currently, if you are mirroring images, you must also mirror **quay.io/openshift-release-dev/ocp-release:<release_image_digest_with_dash>.sig** Sigstore signatures in order for the image verification to succeed.

1.3.23. Web console

1.3.23.1. OpenShift Lightspeed Operator is available in the web console

Starting with OpenShift Container Platform 4.16, OpenShift Lightspeed Operator is available for use in the web console. With this release, a hover button was added to help you discover OpenShift Lightspeed. Once you click the hover button, the chat window appears with instructions on how to enable and install OpenShift Lightspeed on the cluster. You can hide the OpenShift Lightspeed button when changing the default user preferences.

1.3.23.2. Administrator perspective

This release introduces the following updates to the **Administrator** perspective of the web console:

- Deprecated Operators are displayed in the OperatorHub before and after installation along with a warning notification that the Operator is deprecated.
- You can check contnet of the configuration files for **MachineConfig** objects without having to manually retrieve its contents.
- An alert was added to the **Operator details** page and the **Operator installation** page if your cluster is on Google Cloud with Workload Identity Foundation (WIF).
- A page for ShipWright **BuildStrategy** was added to the **Shipwright** page with a **ClusterBuildStrategy** and **BuildStrategy** tab.

1.3.23.2.1. Customize the Create project modal using dynamic plugins

With this release, a new extension point was added, so dynamic plugin creators can pass a component that renders in place of the default **Create Project** modal.

For more information on OpenShift Container Platform Console dynamic plugin SDK extensions, see Dynamic plugin extension types .

1.3.23.2.2. External OpenID Connect (OIDC) token issuer is now functional in the web console

With this update, the web console works as expected when internal **oauth-server** resource and **oauth-apiserver** resource are removed and replaced with an external OpenID Connect (OIDC) issuer.

1.3.23.3. Developer Perspective

This release introduces the following updates to the **Developer** perspective of the web console:

- When you use one of the add flows to create a new deployment, the **Import from Git** or **Container images** automatically opens them in the sidebar.
- You can easily select the desired **Git Type** without having to use the list if OpenShift Container Platform is unable to identify its type.
- **Import from Git** supports GitEA, an open-source alternitve to GitHub.
- A warning notification displays on the **Topology** page if the **PodDisruptionBudget** limit is reached.
- When importing applications through the **Import from Git** flows you can use Shipwright Build strategies such as S2I, buildpack, and buildah strategy for building the image.

1.4. NOTABLE TECHNICAL CHANGES

OpenShift Container Platform 4.17 introduces the following notable technical changes:

1.4.1. Operator SDK 1.36.1

OpenShift Container Platform 4.17 supports Operator SDK 1.36.1. See Installing the Operator SDK CLI to install or update to this latest version.

NOTE

Operator SDK 1.36.1 now supports Kubernetes 1.29 and uses a Red Hat Enterprise Linux (RHEL) 9 base image.

If you have Operator projects that were previously created or maintained with Operator SDK 1.31.0, update your projects to keep compatibility with Operator SDK 1.36.1.

- Updating Go-based Operator projects
- Updating Ansible-based Operator projects
- Updating Helm-based Operator projects
- Updating Hybrid Helm-based Operator projects
- Updating Java-based Operator projects

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. 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 4.17, 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	4.15	4.16	4.17
Bare Metal Event Relay Operator	Removed	Removed	Removed

1.5.2. Images deprecated and removed features

Table 1.7. Cluster Samples Operator deprecated and removed tracker

Feature	4.15	4.16	4.17
Cluster Samples Operator	General Availability	Deprecated	Deprecated

1.5.3. Installation deprecated and removed features

Table 1.8. Installation deprecated and removed tracker

Feature	4.15	4.16	4.17
--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	General Availability	Deprecated	Deprecated
platform.aws.preserveBootstrapIgnition parameter for Amazon Web Services (AWS)	General Availability	Deprecated	Deprecated
Terraform infrastructure provider for Amazon Web Services (AWS), VMware vSphere and Nutanix	General Availability	Removed	Removed
Installing a cluster on Alibaba Cloud with installer-provisioned infrastructure	Technology Preview	Removed	Removed
Installing a cluster on AWS with compute nodes in AWS Outposts	Deprecated	Deprecated	Deprecated

1.5.4. Operator lifecycle and development deprecated and removed features

Table 1.9. Operator lifecycle and development deprecated and removed tracker

Feature	4.15	4.16	4.17
Operator SDK	General Availability	Deprecated	Deprecated
Scaffolding tools for Ansible-based Operator projects	General Availability	Deprecated	Deprecated
Scaffolding tools for Helm-based Operator projects	General Availability	Deprecated	Deprecated
Scaffolding tools for Go-based Operator projects	General Availability	Deprecated	Deprecated

Feature	4.15	4.16	4.17
Scaffolding tools for Hybrid Helm-based Operator projects	Technology Preview	Deprecated	Deprecated
Scaffolding tools for Java-based Operator projects	Technology Preview	Deprecated	Deprecated
Platform Operators	Technology Preview	Removed	Removed
Plain bundles	Technology Preview	Removed	Removed
SQLite database format for Operator catalogs	Deprecated	Deprecated	Deprecated

1.5.5. Machine management deprecated and removed features

Table 1.10. Machine management deprecated and removed tracker

Feature	4.15	4.16	4.17
Managing machine with Machine API for Alibaba Cloud	Technology Preview	Removed	Removed
Cloud controller manager for Alibaba Cloud	Technology Preview	Removed	Removed

1.5.6. Monitoring deprecated and removed features

Table 1.11. Monitoring deprecated and removed tracker

Feature	4.15	4.16	4.17
dedicatedServiceMonitors setting that enables dedicated service monitors for core platform monitoring	Deprecated	Removed	Removed
prometheus-adapter component that queries resource metrics from Prometheus and exposes them in the metrics API	Deprecated	Removed	Removed
Alertmanager v1 API	Deprecated	Deprecated	Removed

1.5.7. Networking deprecated and removed features

Table 1.12. Networking deprecated and removed tracker

Feature	4.15	4.16	4.17
OpenShift SDN network plugin	Deprecated	Deprecated	Removed
iptables	Deprecated	Deprecated	Deprecated
Limited live migration to OVN-Kubernetes from OpenShift SDN	Not Available	General Availability	Removed
PTP events REST API v1 and PTP events consumer application sidecar	General Availability	General Availability	Deprecated

1.5.8. Storage deprecated and removed features

Table 1.13. Storage deprecated and removed tracker

Feature	4.15	4.16	4.17
Persistent storage using FlexVolume	Deprecated	Deprecated	Deprecated
AliCloud Disk CSI Driver Operator	General Availability	Removed	Removed
Shared Resources CSI Driver [1]	Deprecated	Deprecated	Deprecated

1. The Shared Resource CSI Driver feature is now generally available in Builds for Red Hat OpenShift 1.1. This feature is now deprecated in OpenShift Container Platform. To use this feature, ensure you are using Builds for Red Hat OpenShift 1.1 or a more recent version.

1.5.9. Node deprecated and removed features

Table 1.14. Node deprecated and removed tracker

Feature	4.15	4.16	4.17
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	General Availability	Deprecated	Deprecated

1.5.10. Web console deprecated and removed features

Table 1.15. Web console deprecated and removed tracker

Feature	4.15	4.16	4.17
Patternfly 4	Deprecated	Deprecated	Deprecated
React Router 5	Deprecated	Deprecated	Deprecated

1.5.11. Workloads deprecated and removed features

Table 1.16. Workloads deprecated and removed tracker

Feature	4.15	4.16	4.17
DeploymentConfig objects	Deprecated	Deprecated	Deprecated

1.5.12. Deprecated features

1.5.12.1. Announcement of the deprecation of extending compute nodes into AWS Outposts for clusters deployed on AWS Public Cloud

With this release, extending compute nodes into AWS Outposts for clusters deployed on AWS Public Cloud is deprecated. The ability to deploy compute nodes into AWS Outposts after installation, as an extension of an existing OpenShift Container Platform cluster operating in a public AWS region, will be removed with the release of OpenShift Container Platform version 4.20.

For more information, see Extending an AWS VPC cluster into an AWS Outpost .

1.5.12.2. The preserveBootstrapIgnition parameter for AWS

The **preserveBootstrapIgnition** parameter for AWS in the **install-config.yaml** file has been deprecated. You can use the **bestEffortDeleteIgnition** parameter instead. (OCPBUGS-33661)

1.5.12.3. kube-apiserver no longer gets a valid cloud configuration object

In OpenShift Container Platform 4.17, **kube-apiserver** no longer gets a valid cloud configuration object. As a result, the **PersistentVolumeLabel** admission plugin rejects in-tree Google Compute Engine (GCE) persistent disk persistent volumes (PD PVs), that do not have the correct topology. (OCPBUGS-34544)

1.5.12.4. Deprecation of Patternfly 4 and React Router 5

In OpenShift Container Platform 4.16, Patternfly 4 and React Router 5 were deprecated. The deprecated static remains the same for OpenShift Container Platform 4.17. All plugins should migrate to Patternfly 5 and React Router 6 as soon as possible. (OCPBUGS-34538)

1.5.13. Removed features

1.5.13.1. Removed support for TLS 1.2 custom profile ciphers

OpenShift Container Platform no longer supports the following Transport Layer Security (TLS) 1.2 cipher suites. Including these cipher suites as a part of a custom TLS profile for TLS v1.2 does not give you the results that you expect.

- **TLS_RSA_WITH_AES_128_GCM_SHA256**
- **TLS_RSA_WITH_AES_256_GCM_SHA384**
- **TLS_RSA_WITH_AES_128_CBC_SHA**
- **TLS_RSA_WITH_AES_256_CBC_SHA**
- **TLS_RSA_WITH_3DES_EDE_CBC_SHA**

If you must use TLS v1.2, use only v1.2 cipher suites from the Intermediate TLS profile type to ensure that v1.2 is chosen correctly in the TLS handshake. **Intermediate** is the default TLS v1.2 profile and includes the following cipher suites:

- **TLS_AES_128_GCM_SHA256**
- **TLS_AES_256_GCM_SHA384**
- **TLS_CHACHA20_POLY1305_SHA256**
- **ECDHE-ECDSA-AES128-GCM-SHA256**
- **ECDHE-RSA-AES128-GCM-SHA256**
- **ECDHE-ECDSA-AES256-GCM-SHA384**
- **ECDHE-RSA-AES256-GCM-SHA384**
- **ECDHE-ECDSA-CHACHA20-POLY1305**
- **ECDHE-RSA-CHACHA20-POLY1305**
- **DHE-RSA-AES128-GCM-SHA256**
- **DHE-RSA-AES256-GCM-SHA384**

For a wider range of stronger ciphers in TLS v1.2, use the Intermediate TLS profile type rather than a Custom TLS profile type. If you do not require TLS v1.2, use the Modern profile type, which has the strongest ciphers.

1.5.13.2. OpenShift SDN network plugin (Removed)

OpenShift SDN network plugin was deprecated in 4.15 and 4.16. With this release, the SDN network plugin is no longer supported and the content has been removed from the documentation.

1.5.13.3. Removal of RukPak (Technology Preview)

RukPak was introduced as a Technology Preview feature in OpenShift Container Platform 4.12. Starting in OpenShift Container Platform 4.14, it was used as a component in the Technology Preview of Operator Lifecycle Manager (OLM) v1.

Starting in OpenShift Container Platform 4.17, RukPak is now removed and relevant functionality relied upon by OLM v1 has been moved to other components.

1.5.13.4. The Alertmanager v1 API

In Alertmanager **v0.27.0**, the Alertmanager v1 API is removed and is no longer supported. Any requests to **alertmanager-main /api/v1/** endpoints, such as **api/v1/alerts**, will fail. 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.

1.6. BUG FIXES

1.6.1. Bare Metal Hardware Provisioning

- Previously, attempting to configure RAID on specific hardware models by using Redfish might have resulted in the following error: **The attribute StorageControllers/Name is missing from the resource**. With this update, the validation logic no longer requires the **Name** field, because the field is not mandated by the Redfish standard. (OCPBUGS-38465)

- Previously, the management interface for the iDRAC9 Redfish management interface in the Redfish Bare Metal Operator (BMO) module was incorrectly set to iPXE. This caused the error **Could not find the following interface in the ironic.hardware.interfaces.management entrypoint: ipxe** and the deployment failed on Dell Remote Access Controller (iDRAC)-based servers. With this release, the issue is resolved. (OCPBUGS-37261)

1.6.2. Builds

- Previously, builds could not set the **GIT_LFS_SKIP_SMUDGE** environment variable and use its value when cloning the source code. This caused builds to fail for some Git repositories with LFS files. With this release, the build is able to set this environment variable and use it during the **git clone** step of the build, which resolves the issue. (OCPBUGS-33215)

- Previously, if the developer or cluster admin used lowercase environment variable names for proxy information, these environment variables were carried into the build output container image. At runtime, the proxy settings were active and had to be unset. With this release, lowercase versions of the **_PROXY** environment variables are prevented from leaking into built container images. Now, **buildDefaults** are only kept during the build and settings created for the build process only are removed before pushing the image in the registry. (OCPBUGS-12699)

1.6.3. Cloud Compute

- Previously, a machine controller failed to save the VMware vSphere task ID of an instance template clone operation. This caused the machine to go into the **Provisioning** state and to power off. With this release, the VMware vSphere machine controller can detect and recover from this state. (OCPBUGS-1735)

- Previously, the **machine-api** Operator reacted when it deleted a server that was in an **ERROR** state. This happened because the server did not pass a port list. With this release, deleting a machine stuck in an **ERROR** state does not cause an Operator reaction. (OCPBUGS-33806)

- Previously, you could not configure capacity reservation on a Microsoft Azure Workload Identity cluster because of missing permissions. With this release, the **Microsoft.Compute/capacityReservationGroups/deploy/action** permission is added as a

default credential request in the **<infra-name>-openshift-machine-api-azure-cloud-credentials** custom role, so that you can now configure capacity reservation as expected. (OCPBUGS-37154)

- Previously, an optional internal function of the cluster autoscaler caused repeated log entries when it was not implemented. The issue is resolved in this release. (OCPBUGS-33592)

- Previously, a node associated with a restarting machine briefly having a status of **Ready=Unknown** triggered the **UnavailableReplicas** condition in the Control Plane Machine Set Operator. This condition caused the Operator to enter the **Available=False** state and trigger alerts because that state indicates a nonfunctional component that requires immediate administrator intervention. This alert should not have been triggered for the brief and expected unavailabilty during a restart. With this release, a grace period for node unreadiness is added to avoid triggering unnecessary alerts. (OCPBUGS-20061)

- Previously, when an OpenShift Container Platform cluster was installed with no capabilities and later enabled the Build capability, the related Build cluster configuration custom resource definition (CRD) was not created. With this release, the Build cluster configuration CRD and its default instance are created. This allows the Build capability to be fully configured and customized. (OCPBUGS-34395)

- Previously, role bindings related to the Image Registry, Build, and **DeploymentConfig** capabilities were created in every namespace, even if the capabilities were disabled. With this release, role bindings is only created if the capability is enabled on the cluster. (OCPBUGS-34077)

1.6.4. Cloud Credential Operator

- Previously, secrets in the cluster were fetched in a single call. When there was a large number of secrets, the API timed out. With this release, the Cloud Credential Operator fetches secrets in batches limited to 100 secrets. This change prevents timeouts when there is a large number of secrets in the cluster. (OCPBUGS-41233)

- Previously, the Cloud Credential Operator reported an error when the **awsSTSIAMRoleARN** role was not present on a cluster that used manual mode with AWS Security Token Service. With this release, the Cloud Credential Operator no longer reports this as an error. (OCPBUGS-33566)

- Previously, when checking whether passthrough permissions are sufficient, the Cloud Credential Operator sometimes received a response from the Google Cloud API that a permission is invalid for a project. This response caused the Operator to become degraded and installation to fail. With this release, the Operator is updated to handle this error gracefully. (OCPBUGS-36140)

1.6.5. Cluster Version Operator

- Previously, a rarely occurring race condition between Go routines caused the Cluster Version Operator (CVO) to panic after the CVO started. With this release, the Go routines synchronization is improved and the issue is resolved. (OCPBUGS-32678)

1.6.6. Developer Console

- Previously, on some browsers, some icons in the samples catalog were stretched, making it hard to read. With this update, the icons were resized correctly, and now the icons are no longer stretched and easier to read. (OCPBUGS-34516)

- Previously, s2i build strategy was not explicitly mentioned in the **func.yml**. Therefore you could not create OpenShift Serverless functions with the repository. Additionally, error messages were not available if s2i is not mentioned or if **func.yml**. As a result, identifying the reason of failures was not apparent. With this update, if the s2i build strategy is not mentioned, users can still create a function. If it is not s2i, users cannot create a function. The error messages are now different for both the cases. (OCPBUGS-33733)

- Previously, when using a Quick Start guided tour in the OpenShift Container Platform web console, it took multiple clicks of the **Next** button to skip to the next step if the **check your work** dialog was ignored. With this update, it only takes one click, regardless of the state of the **check your work** box. (OCPBUGS-25929)

1.6.7. Driver ToolKit (DTK)

- Previously, DTK incorrectly included the same values for **KERNEL_VERSION** and **RT_KERNEL_VERSION** that exist in the **/etc/driver-toolkit-release.json** configuration file. With this update, the **RT_KERNEL_VERSION** is displayed correctly. (OCPBUGS-33699)

1.6.8. etcd Cluster Operator

- Previous versions of the etcd Operator checked the health of etcd members in serial with an all-member timeout that matched the single-member timeout. As a result, one slow member check could consume the entire timeout and cause later member checks to fail, regardless of the health of that later member. In this release, the etcd Operator checks the health of members in parallel, so the health and speed of one member's check does not affect the other members' checks. (OCPBUGS-36301)

- Previously, the health checks for the etcd Operator were not ordered. As a consequence, the health check sometimes failed even though all etcd members were healthy. The health-check failure triggered a scale-down event that caused the Operator to prematurely remove a healthy member. With this release, the health checks in the Operator are ordered. As a result, the health checks correctly reflect the health of etcd members and an incorrect scale-down event does not occur. (OCPBUGS-36462)

1.6.9. Hosted control planes

To view bug fixes for hosted control planes on OpenShift Container Platform 4.17, see Bug fixes.

1.6.10. Image Registry

- Previously, the internal image registry would not correctly authenticate users on clusters configured with external OpenID Connect (OIDC) users. Consequently, this made it impossible for users to push or pull images to and from the internal image registry. With this update, the internal image registry starts by using the **SelfSubjectReview** API, dropping use of the **openshift specific user** API, which is not available on clusters configured with external OIDC users. As a result, it is now possible to successfully authenticate with the internal image registry again. (OCPBUGS-35335)

- Previously, the image registry was unable to run due to a permissions error in the certificate directory. This issue has been resolved. (OCPBUGS-38885)

- Previously, when enabling **virtualHostedStyle** with **regionEndpoint** set in image registry Operator config, the image registry would ignore the virtual hosted style config and would fail to start. This update fixes the issue by using a new upstream distribution configuration, which is

force path style, in favor of the downstream only version, which is virtual hosted style. (OCPBUGS-32710)

- Previously, when OpenShift Container Platform was deployed on Azure clusters with Workload ID, storage accounts created for the cluster and the image registry had **Storage Account Key Access** enabled by default, which could pose security risks to the deployment.
With this update, shared access keys are disabled by default on new installations that use Workload ID, enhancing security by preventing the use of shared access keys.

IMPORTANT

Shared access keys should only be disabled if the cluster is configured to use Workload ID. Disabling shared access keys on a cluster not configured with Microsoft Entra Workload ID can cause the Image Registry Operator to become degraded.

For existing storage accounts created before this update, shared access keys are not automatically disabled. Administrators must manually disable shared access key support on these storage accounts to prevent the use of shared keys. For more information about disabling shared access keys, see Prevent Shared Key authorization for an Azure Storage account .

OCPBUGS-39428

1.6.11. Installer

- Previously, extracting the IP address from the Cluster API Machine object only returned a single address. On VMware vSphere, the returned address would always be an IPv6 address and this caused issues with the **must-gather** implementation if the address was non-routable. With this release, the Cluster API Machine object returns all IP addresses, including IPv4, so that the **must-gather** issue no longer occurs on VMware vSphere. (OCPBUGS-37427)

- Previously, when installing a cluster on IBM Cloud® into an existing VPC, the installation program retrieved an unsupported VPC region. Attempting to install into a supported VPC region that follows the unsupported VPC region alphabetically caused the installation program to crash. With this release, the installation program is updated to ignore any VPC regions that are not fully available during resource lookups. (OCPBUGS-14963)

- Previously, the installation program attempted to download the OVA on VMware vSphere whether the template field was defined or not. With this update, the issue is resolved. The installation program verifies if the template field is defined. If the template field is not defined, the OVA is downloaded. If the template field is defined, the OVA is not downloaded. (OCPBUGS-39240)

- Previously, enabling custom feature gates sometimes caused installation on an AWS cluster to fail if the feature gate **ClusterAPIInstallAWS=true** was not enabled. With this release, the **ClusterAPIInstallAWS=true** feature gate is not required. (OCPBUGS-34708)

- Previously, some processes could be left running if the installation program exited due to infrastructure provisioning failures. With this update, all installation-related processes are terminated when the installation program terminates. (OCPBUGS-36378)

- Previously, the installation program required permission to create and delete IAM roles when installing a cluster on AWS even when an existing IAM role was provided. With this update, the installation program only requires these permissions when it is creating IAM roles. (OCPBUGS-36390)

- Previously, long cluster names were trimmed without warning the user. With this update, the installation program warns the user when trimming long cluster names. (OCPBUGS-33840)

- Previously, the **openshift-install** CLI sometimes failed to connect to the bootstrap node when collecting bootstrap gather logs. The installation program reported an error message such as **The bootstrap machine did not execute the release-image.service systemd unit**. With this release and after the bootstrap gather logs issue occurs, the installation program now reports **Invalid log bundle or the bootstrap machine could not be reached and bootstrap logs were not collected**, which is a more accurate error message. (OCPBUGS-34953)

- Previously, when installing a cluster on AWS, subnets that the installation program created were incorrectly tagged with the **kubernetes.io/cluster/<clusterID>: shared** tag. With this update, these subnets are correctly tagged with the **kubernetes.io/cluster/<clusterID>: owned** tag. (OCPBUGS-36904)

- Previously, the local etcd data store that is saved during installation was not deleted if the installation failed, consuming extra space on the installation host. With this update, the data store is deleted if infrastructure provisioning failures prevent a successful installation. (OCPBUGS-36284)

- Previously, when a folder was undefined and the data center was located in a data center folder, an wrong folder structure was created starting from the root of the vCenter server. By using the Govmomi **DatacenterFolders.VmFolder**, it used the a wrong path. With this release, the folder structure uses the data center inventory path and joins it with the virtual machine (VM) and cluster ID value, and the issue is resolved. (OCPBUGS-38616)

- Previously, when templates are defined for each failure domain, the installation program required an external connection to download the OVA in VMware vSphere. With this release, the issue is resolved. (OCPBUGS-39239)

- Previously, installing a cluster with a Dynamic Host Configuration Protocol (DHCP) network on Nutanix caused a failure. With this release, this issue is resolved. (OCPBUGS-38934)

- Previously, due to an EFI Secure Boot failure in the SCOS, when the FCOS pivoted to the SCOS the virtual machine (VM) failed to boot. With this release, the Secure Boot is disabled only when the Secure Boot is enabled in the **coreos.ovf** configuration file, and the issue is resolved. (OCPBUGS-37736)

- Previously, if you specified an unsupported architecture in the **install-config.yaml** file the installation program would fail with a **connection refused** message. With this update, the installation program correctly validates the cluster architecture parameter, leading to successful installations. (OCPBUGS-38841)

- Previously, a rare condition om VMware vSphere Cluster API machines caused the vCenter session management to time out unexpectedly. With this release, the Keep Alive support is disabled in the current and later versions of CAPV, and the issue is resolved. (OCPBUGS-38677)

- Previously, the installation program on Amazon Web Services (AWS) used multiple IPv4 public IP addresses that Amazon has started charging for. With this release, support is provided for bring your own (BYO) public IPv4 pools in OpenShift Container Platform so that users have control of IP addresses that are used by their services. Where the BYO public IPv4 pools feature is enabled, two new permissions, **ec2:DescribePublicIpv4Pools** and **ec2:DisassociateAddress**, are required, and the issue is resolved. (OCPBUGS-35504)

- Previously, when users provided public subnets while using existing subnets and creating a

private cluster, the installation program occasionally exposed on the public internet the load balancers that were created in public subnets. This invalidated the reason for a private cluster. With this release, the issue is resolved by displaying a warning during a private installation that providing public subnets might break the private clusters and, to prevent this, users must fix their inputs. (OCPBUGS-38963)

- Previously, during installation the **oc adm node-image create** command used the kube-system/cluster-config-v1 resource to determine the platform type. With this release, the installation program uses the infrastructure resource, which provides more accurate information about the platform type. (OCPBUGS-39092)

- Previously, the **oc adm node-image create** command failed when run against a cluster in a restricted environment with a proxy because the command ignored the cluster-wide proxy setting. With this release, when the command is run it checks the cluster proxy resource settings, where available, to ensure the command is run successfully and the issue is resolved. (OCPBUGS-39090)

- Previously, when installing a cluster with the Agent-based installer, the assisted-installer process could timeout when attempting to add control plane nodes to the cluster. With this update, the assisted-installer process loads fresh data from the assisted-service process, preventing the timeout. (OCPBUGS-36779)

- Previously, when the VMware vSphere vCenter cluster contained an ESXi host that did not have a standard port group defined and the installation program tried to select that host to import the OVA, the import failed and the error **Invalid Configuration for device 0** was reported. With this release, the installation program verifies whether a standard port group for an ESXi host is defined and, if not, continues until it locates an ESXi host with a defined standard port group, or reports an error message if it fails to locate one, resolving the issue. (OCPBUGS-38560)

- Previously, extracting the IP address from the Cluster API Machine object only returned a single IP address. On VMware vSphere, the returned address would always be an IPv6 address and this caused issues with the **must-gather** implementation if the address was non-routable. With this release, the Cluster API Machine object returns all IP addresses, including IPv4, so that the **must-gather** issue no longer occurs on VMware vSphere. (OCPBUGS-37607)

- Previously, when installing a cluster on AWS, Elastic Kubernetes Service (EKS) messages could appear in the installation logs even when EKS was meant to be disabled. With this update, EKS log messages have been disabled. (OCPBUGS-35752)

- Previously, unexpected output would appear in the terminal when creating an installer-provisioned infrastructure cluster. With this release, the issue has been resolved and the unexpected output no longer shows. (OCPBUGS-35547)

- Previously, when installing a cluster on AWS after deleting a cluster with the **./openshift-install destroy cluster** command, the installation would fail with an error stating that there might already be a running cluster. With this update, all leftover artifacts are removed when the cluster is destroyed, resulting in successful installations afterwards. (OCPBUGS-35542)

- Previously, when installing a cluster on AWS, load balancer ingress rules were continuously revoked and re-authorized, causing unnecessary API calls and delays in cluster provisioning. With this update, load balancer ingress rules are no longer revoked during installation, reducing API traffic and installation delays. (OCPBUGS-35440)

- Previously, when setting **platform.openstack.controlPlanePort.network** without a **fixedIPs** value, the installation program would output a misleading error message about the network missing subnets. With this release, the installation program validates that the **install-config** field

controlPlanePort has a valid subnet filter set because it is a required value. (OCPBUGS-37104)

- Previously, adding IPv6 support for user-provisioned installation platforms caused an issue with naming Red Hat OpenStack Platform (RHOSP) resources, especially when you run two user-provisioned installation clusters on the same Red Hat OpenStack Platform (RHOSP) platform. This happened because the two clusters share the same names for network, subnets, and router resources. With this release, all the resources names for a cluster remain unique for that cluster so no interfere occurs. (OCPBUGS-33973)

- Previously, when installing a cluster on IBM Power® Virtual Server with installer-provisioned infrastructure, the installation could fail due to load balancer timeouts. With this update, the installation program waits for the load balancer to be available instead of timing out. (OCPBUGS-34869)

- Previously, when using the Assisted Installer, using a password that contained the colon character (**:**) resulted in a failed installation. With this update, pull secrets containing a colon in the password do not cause the Assisted Installer to fail. (OCPBUGS-31727)

- Previously, solid state drives (SSD) that used SATA hardware were identified as removable. The Assisted Installer for OpenShift Container Platform reported that no eligible disks were found and the installation stopped. With this release, removable disks are eligible for installation. (OCPBUGS-33404)

- Previously, when installing a cluster on bare metal using installer provisioned infrastructure, the installation could time out if the network to the bootstrap virtual machine is slow. With this update, the timeout duration has been increased to cover a wider range of network performance scenarios. (OCPBUGS-41500)

- Previously, when installing a cluster on IBM Power® Virtual Server, the installation program did not list the **e980** system type in the **madrid** region. With this update, the installation program correctly lists this region. (OCPBUGS-38439)

- Previously, after installing a single-node OpenShift cluster, the monitoring system could produce an alert that applied to clusters with multiple nodes. With this update, single-node OpenShift clusters only produce monitoring alerts that apply to single-node OpenShift clusters. (OCPBUGS-35833)

- Previously, when installing a cluster on IBM Power® Virtual Server, the installation could fail due to a DHCP server network collision. With this update, the installation program selects a random number to generate the DHCP network to avoid collision. (OCPBUGS-33912)

- Previously, the installation program used the Neutron API endpoint to tag security groups. This API does not support special characters, so some Red Hat OpenStack Platform (RHOSP) clusters failed to install on RHOSP. With this release, the installation program uses an alternative endpoint to tag security groups so that the issue no longer persists. (OCPBUGS-36913)

- Previously, setting an invalid Universally Unique Identifier (UUID) for the **additionalNetworkIDs** parameter of a machine pool in your **install-config** configuration file could result in the installation program exiting from installing the cluster. With this release, the installation program checks the validity of the **additionalNetworkIDs** parameter before the program continuing with installing the cluster so that this issue no longer persists. (OCPBUGS-35420)

- Previously, for IBM Power® Virtual Server installer-provisioned infrastructure clusters, if no network name existed for a Dynamic Host Configuration Protocol (DHCP), the destroy code would skip deleting the DHCP resource. With this release, a test now checks if a DHCP is in an **ERROR** state, so that the DHCP resource is deleted. (OCPBUGS-35039)

1.6.12. Insights Operator

- Previously, in some Hypershift hosted clusters, the IO archive contained the hostname even with network obfuscation enabled. This issue has been resolved, and IO archives no longer contain hostnames when they are obfuscated. (OCPBUGS-33082)

1.6.13. Machine Config Operator

- Previously, in a cluster that runs OpenShift Container Platform 4.16 with the Telco RAN DU reference configuration, long duration **cyclictest** or **timerlat** tests could fail with maximum latencies detected above **20** us. This issue occured because the **psi** kernel command line argument was being set to **1** by default when cgroup v2 is enabled. With this release, the issue is fixed by setting **psi=0** in the kernel arguments when enabling cgroup v2. The **cyclictest** latency issue reported in OCPBUGS-34022 is now also fixed. (OCPBUGS-37271)

- Previously, if a cluster admin creates a new **MachineOSConfig** object that references a legacy pull secret, the canonicalized version of this secret that gets created is not updated whenever the original pull secret changes. With this release, the issue is resolved. (OCPBUGS-34079)

- Previously, the **/etc/mco/internal-registry-pull-secret.json** secret was being managed by the Machine Config Operator (MCO). Due to a recent change, this secret rotates on an hourly basis. Whenever the MCO detected a change to this secret, it rolled the secret out to each node in the cluster, which resulted in disruptions. With this fix, a different internal mechanism processes changes to the internal registry pull secret to avoid rolling out repeated MachineConfig updates. (OCPBUGS-33913)

- Previously, if you created more than one **MachineOSConfig** object that required a canonicalized secret, only the first object would build. With this fix, the build controller handles multiple **MachineOSBuilds** that use the same canonicalized secret. (OCPBUGS-33671)

- Previously, if machine config pools (MCP) had a higher **maxUnavailable** value than the cluster's number of unavailable nodes, cordoned nodes were able to be erroneously selected as an update candidate. This fix adds a node readiness check in the node controller so that cordoned nodes are queued for an update. (OCPBUGS-33397)

- Previously, nodes could be drained twice if the node was queued multiple times in the drain controller. This behaviour might have been due to increased activity on the node object by on-cluster layering functionality. With this fix, a node queued for drain only drains once. (OCPBUGS-33134)

- Previously, a potential panic was seen in Machine Config Controller and Machine Build Controller objects if a de-reference accidentally deleted **MachineOSConfig**/**MachineOSBuild** to read the build status. The panic is controlled with additional error conditions to warn for allowed MachineOSConfig deletions. (OCPBUGS-33129)

- Previously, after upgrading from OpenShift Container Platform 4.1 or 4.2 to version 4.15, some machines could get stuck during provisioning and never became available. This was because the **machine-config-daemon-firstboot** service was failing due to an incompatible **machine-config-daemon** binary on those nodes. With this release, the correct **machine-config-daemon** binary is copied to nodes before booting. (OCPBUGS-28974)

- Previously, if you attempted to configure on-cluster Red Hat Enterprise Linux CoreOS (RHCOS) image layering on a non-RHCOS node, the node became degraded. With this fix, in this situation, an error message is produced in the node logs, but the node is not degraded. (OCPBUGS-19537)

1.6.14. Management Console

- Previously, the **Cluster overview** page included a **View all steps in documentation** link that resulted in a 404 error for Red Hat OpenShift Service on AWS and Red Hat OpenShift Dedicated clusters. With this update, the link does not appear for Red Hat OpenShift Service on AWS and Red Hat OpenShift Dedicated clusters. (OCPBUGS-37054)

- Previously, a warning was not provided when you were on a Google Cloud cluster that supports GCP Workload Identity and that the Operator supports it. With this release, logic was added to support GCP Workload Identity and Federated Identity Operator installs, so now you are alerted when you are on a Google Cloud cluster. (OCPBUGS-38591)

- Previously, the version number text in the **Updates** graph on the **Cluster Settings** page appeared as black text on a dark background when using Firefox in dark mode. With this update, the text appears as white text. (OCPBUGS-38427)

- Previously, dynamic plugins using PatternFly 4 referenced variables that are not available in OpenShift Container Platform 4.15 and later. This was causing contrast issues for Red Hat Advanced Cluster Management (RHACM) in dark mode. With this update, older chart styles are now available to support PatternFly 4 charts used by dynamic plugins. (OCPBUGS-36816)

- Previously, when the **Display Admission Webhook** warning implementation presented issues with some incorrect code. With this update, the unnecessary warning message has been removed. (OCPBUGS-35940)

- Previously, the global sync lock that applied to all HTTP servers spawned goroutines with a sync lock that is specific to each of the refresh tokens. With this release, the global refresh sync lock on a cluster with an external OIDC environment was replaced with a sync that refreshes for each token. As a result, refresh token performance is improved by 30% to 50%. (OCPBUGS-35080)

- Previously, a warning was not displayed for the **minAvailable** warning in **PodDisruptionBudget** create and edit form. With this update, code logic for displaying the **minAvailable** warning was added, and the **minAvailable** warning is displayed if violated. (OCPBUGS-34937)

- Previously, the **OperandDetails** page displayed information for the first CRD that matched by name. After this fix, the **OperandDetails** page displays information for the CRD that matches by name and the version of the operand. (OCPBUGS-34901)

- Previously, one inactive or idle browser tab caused session expiration for all other tabs. With this change, activity in any tab will prevent session expiration even if there is one inactive or idle browser tab. (OCPBUGS-34387)

- Previously, text areas were not resizable. With this update, you are now able to resize text areas. (OCPBUGS-34200)

- Previously, the **Debug container** link was not displayed for pods with a **Completed** status. With this change, the link now appears. (OCPBUGS-33631)

- Previously, the OpenShift Container Platform web console did not show **filesystem** metrics on the **Nodes list** page due to incorrect Prometheus query. With this update, **filesystem** metrics are correctly displayed. (OCPBUGS-33136)

- Previously, pseudolocalization was not working due to a configuration issue. After this fix, pseudolocalization works again. (OCPBUGS-30218)

- Previously, console pods would crash loop if the **--user-auth** flag was set to **disabled**. With this update, the console backend properly handles this value. (OCPBUGS-29510)

- Previously, utilization cards displayed a **limit** that incorrectly implied a relationship between capacity and limits. With this update, the position of **limit** was changed and the wording updated. (OCPBUGS-23332)

- Previously, in some edge cases, the wrong resource could be fetched when using websockets to watch a namespaced resource without providing a namespace. With this update, a validation to the resource watch logic was added to prevent the websocket request and log an error under this condition. (OCPBUGS-19855)

- Previously, perspective switching was not properly handled. With this update, perspectives that are passed with URL search parameters or plugin route page extensions now correctly switch the perspective and retain the correct URL path. (OCPBUGS-19048)

1.6.15. Networking

- Previously, the SR-IOV Network Operator was listing the **SriovNetworkNodePolicies** resources in random order. This caused the **sriov-device-plugin** pod to enter a continuous restart loop. With this release, the SR-IOV Network Operator lists policies in a deterministic order so that the **sriov-device-plugin** pod does not enter a continuous restart loop. (OCPBUGS-36243)

- Previously, an interface created inside a new pod would remain inactive and the Gratuitous Address Resolution Protocol (GARP) notification would be generated. The notification did not reach the cluster and this prevented ARP tables of other pods inside the cluster from updating the MAC address of the new pod. This situation caused cluster traffic to stall until ARP table entries expired. With this release, a GARP notification is now sent after the interface inside a pod is active so that the GARP notification reaches the cluster. As a result, surrounding pods can identify the new pod earlier than they could with the previous behavior. (OCPBUGS-30549)

- Previously, enabling FIPS for a cluster caused SR-IOV device plugin pods to fail. With this release, SR-IOV device plugin pods have FIPS enabled so that when you enable FIPS for the cluster, the pods do not fail. (OCPBUGS-41131)

- Previously, a race condition was generated after rebooting an OpenShift Container Platform node that used a performance profile with a small number of reserved CPUs. This occurred because Single Root I/O Virtualization (SR-IOV) virtual functions (VFs) shared the same MAC address and any pods that used the VFs would experience communication issues. With this release, an update to the SR-IOV Network Operator config daemon ensures that the Operator checks that no duplicate MAC addresses do not exist on VFs. (OCPBUGS-33137)

- Previously, if you deleted the **sriovOperatorConfig** custom resource (CR), you could not create a new **sriovOperatorConfig** CR. With this release, the Single Root I/O Virtualization (SR-IOV) Network Operator removes validating webhooks when you delete the **sriovOperatorConfig** CR, so that you can create a new **sriovOperatorConfig** CR. (OCPBUGS-37567)

- Previously, when you switched your cluster to use a different load balancer, the Ingress Operator did not remove the values from the **classicLoadBalancer** and **networkLoadBalancer** parameters in the **IngressController** custom resource (CR) status. This situation caused the status of the CR to report wrong information from the **classicLoadBalancer** and **networkLoadBalancer** parameters. With this release, after you

switch your cluster to use a different load balancer, the Ingress Operator removes values from these parameters so that the CR reports a more accurate and less confusing message status. (OCPBUGS-38646)

- Previously, no multicast packets reached their intended target nodes when a multicast sender and a multicast receiver existed on the same node. This happened because of an OVN-Kubernetes RPM package update. With this release, this regression is fixed in the OVN-Kubernetes RPM package, so that the issue no longer persists. (OCPBUGS-34778)

- Previously, when you created a **LoadBalancer** service for the Ingress Operator, a log message was generated that stated the change was not effective. This log message should only trigger for a change to an **Infra** custom resource. With this release, this log message is no longer generated when you create a **LoadBalancer** service for the Ingress Operator. (OCPBUGS-34413)

- Previously, the **DNSNameResolver** controller sent DNS requests to CoreDNS pods for DNS names that had IP addresses with expired time-to-live (TTL) values. This caused a continuous generation of DNS requests and memory leak issues for those pods. With this release, the **DNSNameResolver** controller waits until it receives the updated list of IP addresses and TTL values for a DNS name before sending any more requests to the DNS name. As a result, the controller no longer generates erroneous requests and sends them to pods. CoreDNS pods can now respond to DNS requests in a timely manner and update the **DNSNameResolver** objects with the latest IP addresses and TTLs. (OCPBUGS-33750)

- Previously, when you used the **must-gather** tool, a Multus Container Network Interface (CNI) log file, **multus.log**, was stored in a node's file system. This situation caused the tool to generate unnecessary debug pods in a node. With this release, the Multus CNI no longer creates a **multus.log** file, and instead uses a CNI plugin pattern to inspect any logs for Multus DaemonSet pods in the **openshift-multus** namespace. (OCPBUGS-33959)

- Previously, an alert for **OVNKubernetesNorthdInactive** would not fire in circumstances where it should fire. With this release, the issue is fixed so that the alert for **OVNKubernetesNorthdInactive** fires as expected. (OCPBUGS-33758)

- Previously, for all pods where the default route has been customized, a missing route for the Kubernetes-OVN masquerade address caused each pod to be unable to connect to itself through a service for which it acts as a backend. With this release, the missing route for Kubernetes-OVN masquerade address is added to pods so that the issue no longer occurs. (OCPBUGS-36865)

- Previously, the **iptables-alerter** pod did not handle errors from the **crictl** command-line interface, which could cause the pod to incorrectly log events from **host-network** pods or cause pod restarts. With this release, the errors are handled correctly so that these issues no longer persist. (OCPBUGS-37713)

- Previously, if you created a hosted cluster by using a proxy for the purposes of making the cluster reach a control plane from a compute node, the compute node would be unavailable to the cluster. With this release, the proxy settings are updated for the node so that the node can use a proxy to successfully communicate with the control plane. (OCPBUGS-37786)

- Previously, when a cluster failed to install on an on-premise platform with a configured load balancer, the **LoadBalancer** service's **LoadBalancerReady** condition received the **SyncLoadBalancerFailed** status. The status generated the following message:

  > The kube-controller-manager logs might contain more details.

This message is wrong because the logs are stored in the **cloud-controller-manager** namespace of a project. With this release, the **SyncLoadBalancerFailed** status now communicates the correct message:

> The cloud-controller-manager logs may contain more details.

(OCPBUGS-31664)

- Previously, you could not control log levels for the internal component that selects IP addresses for cluster nodes. With this release, you can now enable debug log levels so that you can either increase or decrease log levels on-demand. To adjust log levels, you must create a config map manifest file with a configuration similar to the following:

```
apiVersion: v1
data:
  enable-nodeip-debug: "true"
kind: ConfigMap
metadata:
  name: logging
  namespace: openshift-vsphere-infra
# ...
```

(OCPBUGS-32348)

- Previously, the Ingress Operator could not successfully update the canary route because the Operator did not have permission to update **spec.host** or **spec.subdomain** fields on an existing route. With this release, the required permission is added to the cluster role for the Operator's service account and the Ingress Operator can update the canary route. (OCPBUGS-36465)

- Previously, administrator privileges were required to run some networking containers, such as Keepalived, on supported on-premise platforms. With this release, these containers no longer require administrator privileges to run them on supported on-premise platforms. (OCPBUGS-36175)

- Previously, if your **NodeNetworkConfigurationPolicy** (NNCP) custom resource (CR) is set to use the default spanning tree protocol (STP) implementation, the CR configuration file would show **stp.enabled: true**, but the OpenShift Container Platform web console cleared the STP checkbox. With this release, the web console only clears the STEP checkbox after you define **stp.enabled: false** in the NNCP CR YAML file. (OCPBUGS-36238)

- Previously, the Ingress Controller status was incorrectly displayed as **Degraded=False** because of a timing update issue with the **CanaryRepetitiveFailures** condition. With this release, the Ingress Controller status is correctly marked as **Degraded=True** for the appropriate length of time that the **CanaryRepetitiveFailures** condition exists. (OCPBUGS-39220)

1.6.16. Node

- Previously, the Container Runtime Config controller did not detect whether a mirror configuration was in use before adding the scope from a **ClusterImagePolicy** CR to the **/etc/containers/registries.d/sigstore-registries.yaml** file. As a consequence, image verification failed with a **Not looking for sigstore attachments** message. With this fix, images are pulled from the mirror registry as expected. (OCPBUGS-36344)

- Previously, a group ID was not added to the **/etc/group** directory within a container when the **spec.securityContext.runAsGroup** attribute was set in the pod specification. With this release, this issue is fixed. (OCPBUGS-39478)

- Previously, because of a critical regression on RHEL 9.4 kernels earlier than **5.14.0-427.26.1.el9_4**, the **mglru** feature had memory management disabled. In this release, the regression issue is fixed so that the **mglru** feature is now enabled in OpenShift Container Platform 4.17. (OCPBUGS-35436)

1.6.17. Node Tuning Operator (NTO)

- Previously, due to an internal bug, the Node Tuning Operator incorrectly computed CPU masks for interrupt and network-handling CPU affinity if a machine had more than 256 CPUs. This prevented proper CPU isolation on those machines and resulted in **systemd** unit failures. With this release, the Node Tuning Operator computes the masks correctly. (OCPBUGS-39164)

- Previously, the Open vSwitch (OVS) pinning procedure set the CPU affinity of the main thread, but other CPU threads did not pick up this affinity if they had already been created. As a consequence, some OVS threads did not run on the correct CPU set, which might interfere with the performance of pods with a Quality of Service (QoS) class of **Guaranteed**. With this update, the OVS pinning procedure updates the affinity of all the OVS threads, ensuring that all OVS threads run on the correct CPU set. (OCPBUGS-35347)

1.6.18. Observability

- Previously, when you log on under the **Administrator** perspective on the OpenShift Container Platform web console and use the **Observe → Alerting** function, an **S is not a function** displayed on alert metrics graph. This issue happened because of a missing function validation check. With this release, the function validation check is added so the alert metric chart displays collected metrics. (OCPBUGS-37291)

1.6.19. OpenShift CLI (oc)

- Previously, when using oc-mirror plugin v2 with the **--delete** flag to remove Operator catalogs from mirror registries, the process failed with the following error:

  > 2024/08/02 12:18:03 [ERROR]: [OperatorImageCollector] pinging container registry localhost:55000: Get "https://localhost:55000/v2/": http: server gave HTTP response to HTTPS client.

  This occurred because oc-mirror plugin v2 was querying the local cache using HTTPS instead of HTTP. With this update, the HTTP client is now properly configured before the query, resolving the issue. (OCPBUGS-41503)

- Previously, when using the oc-mirror plugin v2 in mirror-to-disk mode, catalog images and contents were stored in **subfolders** under **working-dir**, based on the image digest. During the disk-to-mirror process in fully disconnected environments, the plugin tried to resolve the catalog image tag through the source registry, which was unavailable, leading to such errors:

  > [ERROR] : [OperatorImageCollector] pinging container registry registry.redhat.io: Get "http://registry.redhat.io/v2/": dial tcp 23.217.255.152:80: i/o timeout

  With this update, the plugin checks the local cache during the disk-to-mirror process to determine the digest, avoiding the need to query the registry. (OCPBUGS-36214)

- Previously, when using oc-mirror plugin v2 in mirror-to-disk mode in disconnected environments, the plugin was unable to access **api.openshift.com** to download **graph.tar.gz**, resulting in mirroring failures. With this update, the plugin now searches the local cache for the graph image in disconnected environments where the **UPDATE_URL_OVERRIDE** environment variable is set. If the graph image is missing, the plugin skips it without failing. (OCPBUGS-38469)

- Previously, oc-mirror plugin v2 failed to mirror Operator catalogs from disk-to-mirror in fully disconnected environments. This issue also affected catalogs that specified a **targetCatalog** in the **ImageSetConfiguration** file. With this update, the plugin can now successfully mirror catalogs in fully disconnected environments, and the **targetCatalog** functionality works as expected. (OCPBUGS-34521)

- Previously, with the oc-mirror plugin v2, there was no validation for the **-v2** vs **--v2** flags for the **oc-mirror** command. As a result, users who mistakenly used **-v2**, which sets the log level to 2, instead of **--v2**, which switches to oc-mirror plugin v2, received unclear error messages. With this update, flag validation is provided. If the **-v2** flag is used while the **ImageSetConfig** is using the **v2alpha1** API and **--v2** is not specified, an error message is displayed. The following message is now enabled that provides a clear guidance to the user:

  > [ERROR]: Detected a v2 ImageSetConfiguration, please use --v2 instead of -v2.

  (OCPBUGS-33121)

- Previously, oc-mirror plugin v2 did not automatically perform retries when it encountered errors on registries, such as timeouts, expired authentication tokens, HTTP 500 errors, and so on. With this update, retries for these errors are implemented, and users can configure retry behavior with the following flags:

  - **--retry-times**: Specifies the number of retry attempts. Default is 2.
  - **--retry-delay**: Sets the delay between retries. Default is 1 second.
  - **--image-timeout**: Defines the timeout period for mirroring an image. Default is 10 minutes.
  - **--max-parallel-downloads**: Controls the maximum number of layers to pull simultaneously during a single copy operation. Default is 6.
    (OCPBUGS-34021)

- Previously, when using the oc-mirror plugin v2 with the **--rebuild-catalogs** flag, the catalog cache was regenerated locally, which caused failures either due to compatibility issues with the **opm** binary and the platform or due to cache integrity problems on the cluster. With this update, the **--rebuild-catalogs** flag defaults to true, so catalogs are rebuilt without regenerating the internal cache. Additionally, the image command has been modified to generate the cache during pod startup, which may delay pod initialization. (OCPBUGS-37667)

- Previously, the oc-mirror plugin v2 did not use the system proxy configuration to recover signatures for releases when running behind a proxy with system proxy settings. With this release, the system proxy settings are now applied during the signature recovery process. (OCPBUGS-37055)

- Previously, oc-mirror plugin v2 would stop the mirroring process when it encountered Operators using bundle versions that were not compliant with semantic versioning, which also prevented the creation of cluster resources like IDMS, ITMS, and **CatalogSource** objects. With this fix, the

plugin now skips these problematic images instead of halting the process. If an image uses incorrect semantic versioning, a warning message is displayed in the console with the relevant image details. (OCPBUGS-33081)

- Previously, oc-mirror plugin v2 did not generate **ImageDigestMirrorSet** (IDMS) or **ImageTagMirrorSet** (ITMS) files when mirroring failed due to network issues or invalid Operator catalogs. With this update, the **oc-mirror** continues mirroring other images when Operator or additional images fail, and stops only when release images fail. Cluster resources are generated based on successfully mirrored images, and all errors are collected in a log file for review. (OCPBUGS-34020)

- Previously, OpenShift Container Platform release images were not visible in certain registries, such as Red Hat Quay. This prevented users from installing OpenShift Container Platform due to the missing release images. With this update, release images are always tagged to ensure they appear in registries like Red Hat Quay, enabling proper installation. (OCPBUGS-36410)

- Previously, the **oc adm must-gather** command took a long time to gather CPU-related performance data in large clusters. With this release, the data is gathered in parallel instead of sequentially, which shortens the data collection time. (OCPBUGS-34360)

- Previously, the **oc set env** command incorrectly changed the API version of **Route** and **DeploymentConfig** objects, for example, **apps.openshift.io/v1** became **v1**. This caused the command to exit with **unable to recognize no matches for kind** errors. With this release, the error is fixed so that the **os set env** command keeps the correct API version in **Route** and **DeploymentConfig** objects. (OCPBUGS-32108)

- Previously, when a **must-gather** operation failed for any reason and the user manually deleted the leftover namespace, a cluster role binding created by the **must-gather** command would remain in the cluster. With this release, when the temporary **must-gather** namespace is deleted, the associated cluster role binding is automatically deleted with it. (OCPBUGS-31848)

- Previously, when using the **--v2** flag with the oc-mirror plugin v2, if no images were mirrored and some were skipped, empty **imds.yaml** and **itms.yaml** files were generated. With this release, the custom resource generation is only triggered when at least one image is successfully mirrored, preventing the creation of empty files. (OCPBUGS-33775)

1.6.20. Operator Lifecycle Manager (OLM)

- Previously, clusters with many custom resources (CRs) experienced timeouts from the API server and stranded updates where the only workaround was to uninstall and then reinstall the stranded Operators. This occurred because OLM evaluated potential updates by using a dynamic client lister. With this fix, OLM uses a paging lister for custom resource definitions (CRDs) to avoid timeouts and stranded updates. (OCPBUGS-41549)

- Previously, catalog source pods could not recover from a cluster node failure when the **registryPoll** parameter was unset. With this fix, OLM updates its logic for checking for dead pods. As a result, catalog source pods now recover from node failures as expected. (OCPBUGS-39574)

- Previously, if you tried to install a previously-deleted Operator after an OpenShift Container Platform update, the installation might fail. This occurred because OLM could not find previously created bundle unpack jobs. With this fix, OLM correctly installs previously installed Operators. (OCPBUGS-32439)

- Previously, when a new version of a custom resource definition (CRD) specified a new

conversion strategy, this conversion strategy was expected to successfully convert resources. However, OLM cannot run the new conversion strategies for CRD validation without actually performing the update operation. With this release, OLM generates a warning message during the update process when CRD validations fail with the existing conversion strategy, and the new conversion strategy is specified in the new version of the CRD. (OCPBUGS-31522)

- Previously, if the **spec.grpcPodConfig.securityContextConfig** field in **CatalogSource** objects was unset within namespaces with a **PodSecurityAdmission** (PSA) level value of **restricted**, the catalog pod would not pass PSA validation. With this release, the OLM Catalog Operator now configures the catalog pod with the **securityContexts** necessary to pass PSA validation. (OCPBUGS-29729)

- Previously, the **catalogd-controller-manager** pod might not have been deployed to a node despite being in the scheduling queue, and the OLM Operator would fail to install. With this fix, CPU requests are reduced for the related resources, and the issue no longer occurs. (OCPBUGS-29705)

- Previously, the Catalog Operator sometimes attempted to connect to deleted catalog sources that were stored in the cache. With this fix, the Catalog Operator queries a client to list the catalog sources on a cluster. (OCPBUGS-8659)

1.6.21. Red Hat Enterprise Linux CoreOS (RHCOS)

- Previously, LUKS encryption on a system using 512 emulation disks caused provisioning to fail at the **ignition-ostree-growfs** step because of an **sfdisk** alignment issue. With this release, the **ignition-ostree-growfs** script detects this situation and fixes the alignment automatically. As a result, the system no longer fails during provisioning. (OCPBUGS-35410)

- Previously, a bug in the **growpart** utility caused a LUKS device to lock. This caused the system to boot into an emergency mode. With this release, the call to the **growpart** utility is removed and the system successfully boots without issue. (OCPBUGS-33124)

- Previously, if a new deployment was done at the OSTree level on the host, which is identical to the current deployment on a different stateroot, OSTree identified them as equal. This behavior prevented the bootloader from updating when the **set-default** command was invoked, because OSTree did not recognize the two stateroots as a differentiating factor for deployments. With this release, OSTree's logic is modified to consider the stateroots. As a result, OSTree properly sets the default deployment to a new deployment that has different stateroots. (OCPBUGS-30276)

1.6.22. Storage

- Previously, the Secrets Store Container Storage Interface (CSI) Driver on hosted control planes clusters failed to mount secrets because of an issue when using the hosted control planes command-line interface, **hcp**, to create OpenID Connect (OIDC) infrastructure on Amazon Web Services. With this release, the issue has been fixed so that the driver can now mount volumes. (OCPBUGS-18711)

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. Networking Technology Preview features

Table 1.17. Networking Technology Preview tracker

Feature	4.15	4.16	4.17
Ingress Node Firewall Operator	General Availability	General Availability	General Availability
eBPF manager Operator	N/A	N/A	Technology Preview
Advertise by using L2 mode the MetalLB service from a subset of nodes, using a specific pool of IP addresses	Technology Preview	Technology Preview	Technology Preview
Multi-network policies for SR-IOV networks	General Availability	General Availability	General Availability
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	Technology Preview
Admin Network Policy (**AdminNetworkPolicy**)	Technology Preview	General Availability	General Availability
IPsec external traffic (north-south)	General Availability	General Availability	General Availability
Host network settings for SR-IOV VFs	Technology Preview	Technology Preview	General Availability

Feature	4.15	4.16	4.17
Integration of MetalLB and FRR-K8s	Not Available	Technology Preview	General Availability
Dual-NIC Intel E810 PTP boundary clock with highly available system clock	Not Available	General Availability	General Availability
Intel E810 Westport Channel NIC as PTP grandmaster clock	Technology Preview	General Availability	General Availability
Dual-NIC Intel E810 Westport Channel as PTP grandmaster clock	Technology Preview	General Availability	General Availability
Automatic leap seconds handling for PTP grandmaster clocks	Not Available	Not Available	General Availability
PTP events REST API v2	Not Available	Not Available	General Availability
Configure the **br-ex** bridge needed for OVN-Kubernetes to use NMState	Not Available	General Availability	General Availability
Overlapping IP configuration for multi-tenant networks with Whereabouts	Not Available	General Availability	General Availability
User defined network segmentation	Not Available	Not Available	Technology Preview

1.7.2. Storage Technology Preview features

Table 1.18. Storage Technology Preview tracker

Feature	4.15	4.16	4.17
AWS EFS storage CSI usage metrics	Not Available	Not Available	General Availability
Automatic device discovery and provisioning with Local Storage Operator	Technology Preview	Technology Preview	Technology Preview
Azure File CSI snapshot support	Not Available	Not Available	Technology Preview
IBM Power® Virtual Server Block CSI Driver Operator	General Availability	General Availability	General Availability

Feature	4.15	4.16	4.17
Read Write Once Pod access mode	Technology Preview	General Availability	General Availability
Secrets Store CSI Driver Operator	Technology Preview	Technology Preview	Technology Preview
CIFS/SMB CSI Driver Operator	Not Available	Technology Preview	Technology Preview
VMWare vSphere multiple vCenter support	Not Available	Not Available	Technology Preview
Disabling/enabling storage on vSphere	Not Available	Not Available	Technology Preview
RWX/RWO SELinux Mount	Not Available	Not Available	Developer Preview
Migrating CNS Volumes Between Datastores	Not Available	Not Available	Developer Preview

1.7.3. Installation Technology Preview features

Table 1.19. Installation Technology Preview tracker

Feature	4.15	4.16	4.17
Installing OpenShift Container Platform on Oracle® Cloud Infrastructure (OCI) with VMs	General Availability	General Availability	General Availability
Installing OpenShift Container Platform on Oracle® Cloud Infrastructure (OCI) on bare metal	Developer Preview	Developer Preview	Developer Preview
Adding kernel modules to nodes with kvc	Technology Preview	Technology Preview	Technology Preview
Enabling NIC partitioning for SR-IOV devices	Technology Preview	Technology Preview	General Availability
User-defined labels and tags for Google Cloud	Technology Preview	Technology Preview	General Availability
Installing a cluster on Alibaba Cloud by using installer-provisioned infrastructure	Technology Preview	Not Available	Not Available

Feature	4.15	4.16	4.17
Installing a cluster on Alibaba Cloud by using Assisted Installer	Not Available	Technology Preview	Technology Preview
Mount shared entitlements in BuildConfigs in RHEL	Technology Preview	Technology Preview	Technology Preview
Selectable Cluster Inventory	Technology Preview	Technology Preview	Technology Preview
Static IP addresses with VMware vSphere (IPI only)	Technology Preview	General Availability	General Availability
Support for iSCSI devices in RHCOS	Technology Preview	General Availability	General Availability
Installing a cluster on Google Cloud using the Cluster API implementation	Not Available	Technology Preview	General Availability
Support for Intel® VROC-enabled RAID devices in RHCOS	Technology Preview	General Availability	General Availability

1.7.4. Node Technology Preview features

Table 1.20. Nodes Technology Preview tracker

Feature	4.15	4.16	4.17
MaxUnavailableStatefulSet featureset	Technology Preview	Technology Preview	Technology Preview
Linux user namespace support	Not Available	Not Available	Technology Preview

1.7.5. Multi-Architecture Technology Preview features

Table 1.21. Multi-Architecture Technology Preview tracker

Feature	4.15	4.16	4.17
IBM Power® Virtual Server using installer-provisioned infrastructure	General Availability	General Availability	General Availability
kdump on **arm64** architecture	Technology Preview	Technology Preview	Technology Preview

Feature	4.15	4.16	4.17
kdump on **s390x** architecture	Technology Preview	Technology Preview	Technology Preview
kdump on **ppc64le** architecture	Technology Preview	Technology Preview	Technology Preview
Multiarch Tuning Operator	Not available	General Availability	General Availability

1.7.6. Scalability and performance Technology Preview features

Table 1.22. Scalability and performance Technology Preview tracker

Feature	4.15	4.16	4.17
factory-precaching-cli tool	Technology Preview	Technology Preview	Technology Preview
Hyperthreading-aware CPU manager policy	Technology Preview	Technology Preview	Technology Preview
HTTP transport replaces AMQP for PTP and bare-metal events	Technology Preview	General Availability	General Availability
Mount namespace encapsulation	Technology Preview	Technology Preview	Technology Preview
Node Observability Operator	Technology Preview	Technology Preview	Technology Preview
Tuning etcd latency tolerances	Technology Preview	General Availability	General Availability
Increasing the etcd database size	Not Available	Technology Preview	Technology Preview
Using RHACM **PolicyGenerator** resources to manage GitOps ZTP cluster policies	Not Available	Technology Preview	Technology Preview
Pinned Image Sets	Not Available	Technology Preview	Technology Preview

1.7.7. Operator lifecycle and development Technology Preview features

Table 1.23. Operator lifecycle and development Technology Preview tracker

Feature	4.15	4.16	4.17
Operator Lifecycle Manager (OLM) v1	Technology Preview	Technology Preview	Technology Preview
RukPak	Technology Preview	Technology Preview	Removed
Platform Operators	Technology Preview	Removed	Removed
Scaffolding tools for Hybrid Helm-based Operator projects	Technology Preview	Deprecated	Deprecated
Scaffolding tools for Java-based Operator projects	Technology Preview	Deprecated	Deprecated

1.7.8. OpenShift CLI (oc) Technology Preview features

Table 1.24. OpenShift CLI (oc) Technology Preview tracker

Feature	4.15	4.16	4.17
oc-mirror plugin v2	Not Available	Technology Preview	Technology Preview
Enclave support	Not Available	Technology Preview	Technology Preview
Delete functionality	Not Available	Technology Preview	Technology Preview

1.7.9. Monitoring Technology Preview features

Table 1.25. Monitoring Technology Preview tracker

Feature	4.15	4.16	4.17
Metrics Collection Profiles	Technology Preview	Technology Preview	Technology Preview
Metrics Server	Technology Preview	General Availability	General Availability

1.7.10. Monitoring Technology Preview features

Table 1.26. Monitoring Technology Preview tracker

Feature	4.15	4.16	4.17
Red Hat OpenShift Lightspeed in the OpenShift Container Platform web console	Not Available	Developer Preview	Developer Preview

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

Table 1.27. RHOSP Technology Preview tracker

Feature	4.15	4.16	4.17
Dual-stack networking with installer-provisioned infrastructure	General Availability	General Availability	General Availability
Dual-stack networking with user-provisioned infrastructure	General Availability	General Availability	General Availability
RHOSP integration into the Cluster CAPI Operator	Technology Preview	Technology Preview	Technology Preview
Control Plane with **rootVolumes** and **etcd** on local disk	Technology Preview	Technology Preview	General Availability

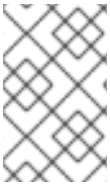

NOTE

To know the status of the hosted control planes features on OpenShift Container Platform 4.17, see Generally Available and Technology Preview features in hosted control planes release notes.

1.7.12. Hosted control planes Technology Preview features

Table 1.28. Hosted control planes Technology Preview tracker

Feature	4.15	4.16
Hosted control planes for OpenShift Container Platform on Amazon Web Services (AWS)	Technology Preview	General Availability
Hosted control planes for OpenShift Container Platform on bare metal	General Availability	General Availability
Hosted control planes for OpenShift Container Platform on OpenShift Virtualization	General Availability	General Availability
Hosted control planes for OpenShift Container Platform using non-bare metal agent machines	Technology Preview	Technology Preview

Feature	4.15	4.16
Hosted control planes for an ARM64 OpenShift Container Platform cluster on Amazon Web Services	Technology Preview	Technology Preview
Hosted control planes for OpenShift Container Platform on IBM Power	Technology Preview	Technology Preview
Hosted control planes for OpenShift Container Platform on IBM Z	Technology Preview	Technology Preview
Hosted control planes for OpenShift Container Platform on RHOSP	Not Available	Not Available

1.7.13. Machine management Technology Preview features

Table 1.29. Machine management Technology Preview tracker

Feature	4.15	4.16	4.17
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
Managing machines with the Cluster API for RHOSP	Technology Preview	Technology Preview	Technology Preview
Managing machines with the Cluster API for VMware vSphere	Not Available	Technology Preview	Technology Preview
Defining a vSphere failure domain for a control plane machine set	Technology Preview	General Availability	General Availability
Cloud controller manager for Alibaba Cloud	Technology Preview	Removed	Removed
Cloud controller manager for Google Cloud	General Availability	General Availability	General Availability
Cloud controller manager for IBM Power® Virtual Server	Technology Preview	Technology Preview	Technology Preview

1.7.14. Authentication and authorization Technology Preview features

Table 1.30. Authentication and authorization Technology Preview tracker

Feature	4.15	4.16	4.17
Pod security admission restricted enforcement	Technology Preview	Technology Preview	Technology Preview

1.7.15. Machine Config Operator Technology Preview features

Table 1.31. Machine Config Operator Technology Preview tracker

Feature	4.15	4.16	4.17
Improved MCO state reporting	Technology Preview	Technology Preview	Technology Preview
On-cluster RHCOS image layering	Not Available	Technology Preview	Technology Preview
Node disruption policies	Not Available	Technology Preview	General Availability
Updating boot images for GCP clusters	Not Available	Technology Preview	General Availability
Updating boot images for AWS clusters	Not Available	Not Available	Technology Preview

1.7.16. Edge computing Technology Preview features

Table 1.32. Edge computing Technology Preview tracker

Feature	4.15	4.16	4.17
Accelerated provisioning of GitOps ZTP	Not Available	Technology Preview	Technology Preview
Enabling disk encryption with TPM and PCR protection	Not Available	Not Available	Technology Preview

1.8. KNOWN ISSUES

- The **oc annotate** command does not work for LDAP group names that contain an equal sign (**=**), because the command uses the equal sign as a delimiter between the annotation name and value. As a workaround, use **oc patch** or **oc edit** to add the annotation. (BZ#1917280)

- A known issue exists when deleting a **NetworkAttachmentDefinition** (NAD) resource created by a **UserDefinedNetwork** resource. You must check to see if a pod is referencing the NAD before deleting the NAD. The pod should be deleted before the NAD. Failure to do so can leave

pods in an unexpected state. (OCPBUGS-39185)

- The DNF package manager included in Red Hat Enterprise Linux CoreOS (RHCOS) images cannot be used at runtime, because DNF relies on additional packages to access entitled nodes in a cluster that are under a Red Hat subscription. As a workaround, use the **rpm-ostree** command instead. (OCPBUGS-35247)

- When installing a cluster on Microsoft Azure, the installation will fail if no **install-config.yaml** file is provided. If an **install-config.yaml** file is provided, and **controlPlane.platform** is present but **controlPlane.platform.azure** is not provided, the installation will fail. (OCPBUGS-42296)
  See Sample customized install-config.yaml file for Azure for a sample configuration file, or set a non-null parameter as in the following example:

```
controlPlane:
  platform:
    azure: {}
```

- When installing multiple clusters on Microsoft Azure, running multiple installations simultaneously from the same installation host will result in only one of the clusters installing successfully. If you run the installations sequentially rather than simultaneously, you can install multiple clusters on Azure from the same installation host. (OCPBUGS-36202)

- When installing a cluster on Microsoft Azure, specifying the **Standard_M8-4ms** instance type for control plane machines results in an error due to that instance type specifying its memory in decimal format instead of integer format. (OCPBUGS-42241)

- At the release of OpenShift Container Platform 4.17, a change in storage account naming caused an issue where the Azure File Container Storage Interface (CSI) driver would fail to mount all volumes when the image registry was configured as private. The mount failures occurred because the CSI driver tried to use the storage account of the Image Registry Operator, which was not configured to allow connections from worker subnets. This issue was resolved in OpenShift Container Platform 4.17.5 and applies to later releases.

- When installing a cluster on Azure, the installation fails if a customer-managed encryption key is specified. (OCPBUGS-42349)

- When an error occurs during mirroring Operators and additional images, the log message "Generating Catalog Source" might still appear, even if no files are generated. (OCPBUGS-42503)

- If you have IPsec enabled on the cluster, on the node hosting the north-south IPsec connection, restarting the **ipsec.service** systemd unit or restarting the **ovn-ipsec-host** pod causes a loss of the IPsec connection. (RHEL-26878)

- Run Once Duration Override Operator (RODOO) cannot be installed on clusters managed by the Hypershift Operator. (OCPBUGS-17533)

- When you run Cloud-native Network Functions (CNF) latency tests on an OpenShift Container Platform cluster, the test can sometimes return results greater than the latency threshold for the test; for example, 20 microseconds for **cyclictest** testing. This results in a test failure. (OCPBUGS-42328)

- Each node group must only match one **MachineConfigPool** object. In some cases, the NUMA Resources Operator can allow a configuration where a node group matches more than one **MachineConfigPool** object. This issue could lead to unexpected behavior in resource management. (OCPBUGS-42523)

- If you plan to deploy the NUMA Resources Operator, avoid using OpenShift Container Platform versions 4.17.7 or 4.17.8. (OCPBUGS-45639)
- When the bond mode in the **NetworkNodeConfigurationPolicy** is changed from **balance-rr** to **active-backup** on kernel bonds that are attached to the **br-ex** interface, the change might fail on arbitrary nodes. As a workaround, create a **NetworkNodeConfigurationPolicy** object without specifying the bond port configuration. (OCPBUGS-42031)
- If the controller pod terminates while cloning, or taking or restoring a volume snapshot, is in progress, the Microsoft Azure File clone or snapshot persistent volume claims (PVCs) remain in the Pending state. To resolve this issue, delete any affected clone or snapshot PVCs, and then recreate those PVCs. (OCPBUGS-35977)
- Deploying a self-managed private hosted cluster on AWS fails because the **bootstrap-kubeconfig** file uses an incorrect KAS port. As a result, the AWS instances are provisioned, but cannot join the hosted cluster as nodes. (OCPBUGS-31840)

1.9. ASYNCHRONOUS ERRATA UPDATES

Security, bug fix, and enhancement updates for OpenShift Container Platform 4.17 are released as asynchronous errata through the Red Hat Network. All OpenShift Container Platform 4.17 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.

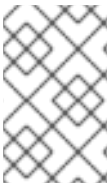

NOTE

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 4.17. Versioned asynchronous releases, for example with the form OpenShift Container Platform 4.17.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-2026:4510 - OpenShift Container Platform 4.17.51 fixed issues and security update

Issued: 18 March 2026

OpenShift Container Platform release 4.17.51 is now available. The list of bug fixes that are included in the update is documented in the RHSA-2026:4510 advisory. The RPM packages that are included in the update are provided by the RHBA-2026:4479 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 4.17.51--pullspecs
```

1.9.1.1. Fixed issues

- 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 blocked or slowed down host removal, creating operational inefficiencies. With this release, if the BMH has the 'detached annotation' and deletion is requested, the BMH transitions to the deleting state, allowing for direct deletion. (OCPBUGS-77185)

- Before this update, a regression introduced in version OpenShift Container Platform 4.15 impacted the **AlertingRule** logic, leading to duplicate **prometheusRules** with a different hash but with the same alerting rule. This regression caused the system to generate duplicate alerts. Because one of these alerts was no longer maintained, this behavior could be confusing and might lead to stale expressions being evaluated. With this release, the underlying regression has been resolved, restoring the intended behavior and logic to the **AlertingRule** component. (OCPBUGS-77275)

- Before this update, the user workload Prometheus Operator did not validate the **webhookURL** secret reference in the Microsoft Teams receiver configuration of the **AlertmanagerConfig** custom resource (CR). As a consequence, an invalid or missing **webhookURL** secret could be accepted, which caused the user workload **Alertmanager** to crash at runtime. With this update, the user workload Prometheus Operator validates the **webhookURL** secret for Microsoft Teams receivers. As a result, invalid configurations are rejected before the configurations can affect the **Alertmanager**. (OCPBUGS-77452)

- Before this update, the system failed to correctly flush stale **conntrack** entries for User Datagram Protocoal (UDP) **NodePort** and **LoadBalancer** services when an **EndpointSlice** was updated, specifically in scenarios where the backend pod resided on a different node. This resulted in traffic being routed to obsolete endpoints, causing intermittent connectivity drops for UDP streams. With this release, the logic has been corrected to ensure **conntrack** entries are properly cleared by filtering for the **nodePort** regardless of the pod's location. (OCPBUGS-77568)

- Before this update, the Microsoft Azure installation process would fail when deploying into pre-existing virtual networks (VNets) that used an IP range different from the installation default. This occurred because the installation program selected an IP address as an internal **loadbalancer** front end IP address from the default subnet rather than the pre-configured subnet. With this release, the installation logic has been updated to dynamically detect and respect the Classless Inter-Domain Routing (CIDR) blocks of pre-existing VNets. (OCPBUGS-77791)

1.9.1.2. Updating

To update an OpenShift Container Platform 4.17 cluster to this latest release, see Updating a cluster using the CLI.

1.9.2. RHSA-2026:3418 - OpenShift Container Platform 4.17.50 fixed issues and security update

Issued: 4 March 2026

OpenShift Container Platform release 4.17.50 is now available. The list of bug fixes that are included in the update is documented in the RHSA-2026:3418 advisory. The RPM packages that are included in the update are provided by the RHSA-2026:3416 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 4.17.50--pullspecs
```

1.9.2.1. Fixed issues

- Before this update, the **collect-profiles** job in {olm0-first} periodically used CPU and disk space. With this release, the **collect-profiles** job is removed from existing OLM, which reduces the periodic CPU and disk use caused by the job process. (OCPBUGS-76953)

1.9.2.2. Updating

To update an OpenShift Container Platform 4.17 cluster to this latest release, see Updating a cluster using the CLI.

1.9.3. RHSA-2026:2672 - OpenShift Container Platform 4.17.49 fixed issues

Issued: 18 February 2026

OpenShift Container Platform release 4.17.49 is now available. The list of fixed issues that are included in the update is documented in the RHBA-2026:2672 advisory. The RPM packages that are included in the update are provided by the RHSA-2026:2670 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 4.17.49--pullspecs
```

1.9.3.1. Fixed issues

- 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 (conntrack) entries could remain on worker nodes. This occurred because the **conntrack** 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, the cleanup process uses the service port when deleting or updating service endpoints. This change ensures that stale conntrack entries are correctly matched and removed. Network connectivity now remains reliable during service endpoint lifecycle changes. (OCPBUGS-71985)

- 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-73799)

- Before this update, the **osDiskImage.source.id** property in the user-provisioned infrastructure (UPI) ARM template was incorrectly placed. As a consequence, it led to installation failures. With this release, the **storageProfile.osDiskImage.source.id** property location in the UPI ARM template is corrected. As a result, the UPI ARM template installation on Azure no longer encounters invalid parameter errors. (OCPBUGS-74555)

- Before this update, Google Cloud installation failed due to unspecified zones in **install-config** for regions with AI zones. As a consequence, Google Cloud installation in specific regions failed, causing installation issues for users. With this release, the Google Cloud installer now supports AI zones in us-south1 and us-central1 regions. As a result, Google Cloud installation does not fail in specific regions due to AI zones. (OCPBUGS-74677)

1.9.3.2. Updating

To update an OpenShift Container Platform 4.17 cluster to this latest release, see Updating a cluster using the CLI.

1.9.4. RHSA-2026:1577 - OpenShift Container Platform 4.17.48 fixed issues

Issued: 04 February 2026

OpenShift Container Platform release 4.17.48 is now available. The list of fixed issues that are included in the update is documented in the RHSA-2026:1577 advisory. The RPM packages that are included in the update are provided by the RHBA-2026:1544 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 4.17.48--pullspecs
```

1.9.4.1. Fixed issues

- Before this update, aggregated API servers on OpenShift Container Platform used in-memory loopback certificates that were valid for only one year. With this release, these servers now use in-memory loopback certificates that are valid for three years. (OCPBUGS-66227)

- Before this update, 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. Consequently, 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. If an instance does not appear within 20 seconds, the machine status changes to **Failed** to prevent instance leaks. (OCPBUGS-73822)

- Before this update, 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, failing authentication. With this release, the proper authentication header is sent when a user and password is specified in the proxy URL. (OCPBUGS-74230)

1.9.4.2. Updating

To update an OpenShift Container Platform 4.17 cluster to this latest release, see Updating a cluster using the CLI.

1.9.5. RHSA-2026:0715 - OpenShift Container Platform 4.17.47 fixed issues

Issued: 21 January 2026

OpenShift Container Platform release 4.17.47 is now available. The list of fixed issues that are included in the update is documented in the RHSA-2026:0715 advisory. The RPM packages that are included in the update are provided by the RHBA-2025:0701 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 4.17.47--pullspecs
```

1.9.5.1. Enhancements

- You can now collect logs that contain command-line information from **virt-launcher** pods. These logs serve as a centralized record of the JSON-encoded command-line options used to start virtual machines. You can use this data to troubleshoot and analyze virtual machine startup issues. This feature is backported to specific OpenShift Container Platform versions. (OCPBUGS-61775)

1.9.5.2. Fixed issues

- Before this release, pods on the same node could not reach a pod's secondary **Localnet** interface on a **br-ex** bridge, so the pod defaulted to the primary network. With this release, communication between a **Localnet** pod and a pod running on the same node is possible, if the **Localnet** network's IP addresses are on the same subnet as the host network. (OCPBUGS-59381)

- Before this update, bonded network configurations with **mode=active-backup** and **fail_over_mac=follow** failed due to a race condition, causing interfaces to be activated multiple times and leading to unpredictable states and identical MAC addresses. As a consequence, these configurations flapped, causing high availability issues. With this release, the bonded network configuration issue is fixed by modifying **configure-ovs.sh** to check for both **activating** and **active** states. As a result, bonded network configurations with **active-backup** mode and **fail_over_mac=follow** does not flap, ensuring high availability. (OCPBUGS-60890)

- Before this update, Prometheus **Remote-Write** caused alerts due to dropped samples in relabeling configuration. As a consequence, the Prometheus **Remote-Write Behind** alert was active. With this release, the **Remote-Write** issue is resolved, and dropped samples do not trigger the alert. As a result, the Prometheus **Remote-Write Behind** alert does not become active when samples are dropped due to relabeling configuration. (OCPBUGS-61766)

- Before this update, **NMState** service failed due to a dependency issue with **NetworkManager-wait-online** in bare metal deployments. As a consequence, the **NMState** service failed in OpenShift Container Platform deployments, causing **br-ex** configuration issues and deployment

failures. With this release, the **NetworkManager-wait-online** dependency issue is resolved, and ensures that the **NMState** service does not fail in OpenShift Container Platform deployments, and network configuration is applied correctly. (OCPBUGS-61859)

1.9.5.3. Updating

To update an OpenShift Container Platform 4.17 cluster to this latest release, see Updating a cluster using the CLI.

1.9.6. RHBA-2026:23120 - OpenShift Container Platform 4.17.46 fixed issues

Issued: 06 January 2026

OpenShift Container Platform release 4.17.46 is now available. The list of fixed issues that are included in the update is documented in the RHBA-2026:23120 advisory. The RPM packages that are included in the update are provided by the RHBA-2025:23118 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 4.17.46--pullspecs
```

1.9.6.1. Fixed issues

- Before this update, when an IDMS or ICSP in the management cluster defined a source that pointed to **registry.redhat.io** or **registry.redhat.io/redhat**, and the mirror registry did not contain the required OLM catalog images, the provisioning of the **HostedCluster** object stalled due to unauthorized image pulls. As a consequence, the **HostedCluster** object was not deployed, and was blocked from pulling essential catalog images from the mirrored registry. With this release, the provisioning explicitly fails and blocks if a required image cannot be pulled due to authorization errors. In addition, the logic for registry overrides is improved to allow matches on the root of the registry, such as **registry.redhat.io**, for OLM CatalogSource image resolution. Also, a fallback mechanism is introduced to use the original image reference if the registry override does not yield a working image. As a result, the **HostedCluster** object is deployed, even in scenarios where the mirror registry lacks the required OLM catalog images, because the system correctly falls back to pull from the original source when appropriate. (OCPBUGS-57123).

- Before this update, 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 gathers the information correctly. (OCPBUGS-60607)

- Before this update, when network interfaces were bonded, the bond interface presented itself using the MAC address of one of the interfaces that it used. The other interface in the bond also presented itself using that MAC address, but also retained its permanent MAC address. With this release, this permanent MAC address is reported and discovered by ironic-python-agent, so that if it is used as a boot MAC address, the ironic agent can look up the node that contains the MAC address. (OCPBUGS-62490)

- Before this update, when you directly navigated to a page created by a web console dynamic plugin, the web console might have redirected you to a different URL. With this release, the URL redirect is removed. (OCPBUGS-65933)

- Before this update, the Azure Machine API provider attempted to use the default **platformUpdateDomainCount** of **5** even in regions that are restricted to a single fault domain. This caused machine creation to fail for all node types in affected regions because Azure only supports one update domain when the fault domain count is **1**. With this release, the logic was updated to explicitly set the **platformUpdateDomainCount** parameter to **1** whenever the **platformFaultDomainCount** parameter is determined to be **1**. As a result, Machine Availability Sets are created with valid parameter combinations, which allows machines to successfully provision in Azure regions with a single fault domain. (OCPBUGS-65954)

1.9.6.2. Updating

To update an OpenShift Container Platform 4.17 cluster to this latest release, see Updating a cluster using the CLI.

1.9.7. RHBA-2025:22266 - OpenShift Container Platform 4.17.45 fixed issues

Issued: 3 December 2025

OpenShift Container Platform release 4.17.45 is now available. The list of fixed issues that are included in the update is documented in the RHBA-2025:22266 advisory. The RPM packages that are included in the update are provided by the RHBA-2025:22264 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 4.17.45--pullspecs
```

1.9.7.1. Fixed issues

- Before this update, URLs passed to the **AlertmanagerConfig** receiver, specifically for Slack and Discord configurations, were not validated. With this release, URL validation is implemented, ensuring that only valid URLs are accepted and invalid ones are rejected. (OCPBUGS-58408)

- Before this update, a Kube State Metrics (KSM) deny-list had typographical errors as demonstrated in the following example:

  ```
  --metric-denylist=
          ^kube_secret_labels$,
          ^kube_.+_annotations$
          ^kube_customresource_.+_annotations_info$,
          ^kube_customresource_.+_labels_info$,
  ```

  With this release, a missing comma is now included as shown in the following example:

  ```
  --metric-denylist=
          ^kube_secret_labels$,
          ^kube_.+_annotations$,
          ^kube_customresource_.+_annotations_info$,
          ^kube_customresource_.+_labels_info$,
  ```

  As a result, the entries are correctly separated. (OCPBUGS-64580)

- Before this update, the **must-gather** pod could be scheduled on a node marked with a

NotReady taint, resulting in deployment to an unavailable node and subsequent log collection failures. With this release, the scheduler now accounts for node taints and automatically applies a node selector to the pod specification. This change ensures that **must-gather** pods are not scheduled on tainted nodes, thereby preventing log collection failures. (OCPBUGS-64615)

- Before this update, during failover, the system's duplicate address detection (DAD) could incorrectly disable the Egress IPv6 address if it was briefly present on both nodes, breaking the connection. With this release, the Egress IPv6 is configured to skip the DAD check during failover, guaranteeing uninterrupted egress IPv6 traffic after an Egress IP address successfully moves to a different node. As a result, greater network stability is ensured. (OCPBUGS-64674)

- Before this update, the **Observe → Metric** page used the cluster-wide metrics API even when you did not have cluster-wide metrics API permissions. As a consequence, the query input showed an error and the autocomplete for the query input did not work without cluster-wide metrics API access. With this release, the **namespace-tenancy** metrics API is used if you do not have cluster-wide metrics API permissions, As a result, an error does not occur and autocomplete is available for the metrics within the selected namespace. (OCPBUGS-64942)

- 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-65480)

- Before this update, a regression in **runc** prevented pods which had the **shareProcessNamespace** object set to **true** from running properly. With this release, the regression has been corrected. As a result, the underlying issue is resolved, and pods using the **shareProcessNamespace** object can now start and function as expected. (OCPBUGS-65977)

1.9.7.2. Updating

To update an OpenShift Container Platform 4.17 cluster to this latest release, see Updating a cluster using the CLI.

1.9.8. RHBA-2025:21225 - OpenShift Container Platform 4.17.44 fixed issues

Issued: 19 November 2025

OpenShift Container Platform release 4.17.44 is now available. The list of fixed issues that are included in the update is documented in the RHBA-2025:21225 advisory. The RPM packages that are included in the update are provided by the RHBA-2025:21221 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 4.17.44--pullspecs
```

1.9.8.1. Fixed issues

- Before this update, API and Ingress Virtual IP (VIP) addresses were automatically assigned even when a user-managed load balancer was in use. With this release, the API and Ingress VIPs are

no longer automatically assigned. If these values are not explicitly set in the **install-config.yaml**, the installation fails with an error, prompting you to provide them. (OCPBUGS-53236)

- Before this update, the **Authentication Error** message page was not properly rendering, causing an empty page and preventing error messages from being displayed to users. With this release, the **Authentication Error** message page now displays content, enhancing the user experience. (OCPBUGS-62631)

- Before this update, if the open virtual network (OVN)-Kubernetes controller was not processing updates from the Kubernetes API server and configuring the OVN databases on each node, then the OVN-Controller, which consumed this database, might have connected to the database before the OVN-Kubernetes controller had configured them. As a consequence, the OVN-Controller synced with a stale OVN database, consumed source network address translations (SNATs) that were configured to support the egress IP, and proceeded to the gratuitous address resolution protocol (GARP) for the associated IP address even though that IP address might have moved to another node. With this release, these GARPs are blocked when the OVN-Kubernetes controller is not processing updates. (OCPBUGS-63154)

- Before this update, when you ran the **ocp-tuned-one-shot.service** systemd unit that was owned by the Node Tuning Operator (NTO), a dependency failure might have occurred for the kubelet. As a consequence, the kubelet did not start. With this release, running the` ocp-tuned-one-shot.service` unit does not cause a dependency failure. As a result, the kubelet starts when you run the unit. (OCPBUGS-63504)

- Before this update, gRPC connection logs were set at a highly verbose log level. This generated an excessive number of messages, which caused the logs to overflow. With this release, the gRPC connection logs have been moved to the V(4) log level. Consequently, the logs no longer overflow, as these specific messages are now less verbose by default. (OCPBUGS-63682)

- Before this update, the Azure machine provider was not passing the **dataDisks** configuration from the **MachineSet** into the virtual machine creation API request for the Azure Stack Hub. As a consequence, new machines were created without the specified data disks because the configuration was silently ignored during the VM creation process. With this release, the VM creation for the Azure Stack Hub is updated to include the **dataDisks** configuration. An additional update manually implements the behavior of the **deletionPolicy: Delete** parameter in the controller because the Azure Stack Hub does not natively support this option. As a result, data disks are correctly provisioned on the Azure Stack Hub VMs. The **Delete** policy is also functionally supported, which ensures that disks are properly removed when their machines are removed. (OCPBUGS-63700)

1.9.8.2. Updating

To update an OpenShift Container Platform 4.17 cluster to this latest release, see Updating a cluster using the CLI.

1.9.9. RHBA-2025:19314 - OpenShift Container Platform 4.17.43 fixed issues and security update

Issued: 5 November 2025

OpenShift Container Platform release 4.17.43 is now available. The list of fixed issues that are included in the update is documented in the RHBA-2025:19314 advisory. The RPM packages that are included in the update are provided by the RHBA-2025:19312 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 4.17.43 --pullspecs
```

1.9.9.1. Fixed issues

- Before this update, the stale IP addresses in the **address_set** corresponding to the Domain Name System (DNS) Egress Firewall rule were not removed. This resulted in an ever-growing **address_set** leading to memory leak issues. With this release, the stale IP address are removed from the **address_set** after a 5 second grace period. (OCPBUGS-61749)

- Before this update, the dashboard query template was not updated with the newly selected namespace. As a consequence, there was a difference between the namespace parameter and the query, which caused a rejection of the request. With this release, the dashboard namespace variable is synced with the console's selected namespace in the console. As a result, the synced dashboard variable is filled in the query template, and the backend returns the correct results. (OCPBUGS-62282)

- Before this update, the linked URL was in the developer perspective, but the perspective was not switched when you clicked the link. As a consequence, a blank page was shown. With this release, the perspective changes when you click the link and the page is correctly shown. (OCPBUGS-63212)

- Before this update, the Machine Config Operator (MCO) failed due to multiple labels in the **capacity.cluster-autoscaler.kubernetes.io/labels** annotation in one or more machine sets. With this release, the MCO accepts multiple labels in the **capacity.cluster-autoscaler.kubernetes.io/labels** annotation. As result, the MCO does not fail during the update to 4.19.6. (OCPBUGS-63364)

- Before this update, 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. The **istag** object remains intact when using the **--dry-run=server** option. (OCPBUGS-63392)

- Before this update, the **Roles** list page could display invalid data if a regular user had yet to create a project. With this release, the **Roles** list page correctly displays. (OCPBUGS-63397)

- 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-63460)

- Before this update, control plane nodes created before OpenShift Container Platform 4.12 would not have had the **node-role.kubernetes.io/control-plane** label. With this release, the Machine Config Operator (MCO) adds the label whenever it uncordons a control plane node, if the label was missing. (OCPBUGS-63540)

1.9.9.2. Updating

To update an OpenShift Container Platform 4.17 cluster to this latest release, see Updating a cluster using the CLI.

1.9.10. RHBA-2025:18235 - OpenShift Container Platform 4.17.42 fixed issues and security update

Issued: 22 October 2025

OpenShift Container Platform release 4.17.42 is now available. The list of fixed issues that are included in the update is documented in the RHBA-2025:18235 advisory. There are no RPM packages for this release.

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 4.17.42 --pullspecs
```

1.9.10.1. Fixed issues

- Before this update, multiple mirrors in the hosted control planes payload caused image lookup failures, and led to hosted control planes creation failures when multiple mirror images were handled. With this release, the hosted control planes payload supports multiple mirrors, correctly handles unavailable mirrors, and results in a successful cluster creation. (OCPBUGS-57143)
- Before this update, a fixed upstream issue in OpenShift Container Platform version 4.17 was not reproducible and was triggered by consuming the issue for downstream. This led to data inconsistency for users. With this release, the upstream issue in the **external-resizer** mechanism is fixed, and reduces potential consumption for downstream. As a result, users do not experience the issue in OpenShift Container Platform version 4.17. (OCPBUGS-62466)
- Before this update, the incorrect **hostPath** configuration for the **/etc/docker** directory caused Operator Lifecycle Manager (OLM) v1 pod issues. As a consequence, OLM v1 pods failed to work due to a **hostPath** type check error. With this release, the fix ensures the correct directory for **/etc/docker** in **hostPath**, and resolves OLM v1 pod issues. As a result, the OpenShift Container Platform cluster operates smoothly. (OCPBUGS-62741)

1.9.10.2. Updating

To update an OpenShift Container Platform 4.17 cluster to this latest release, see Updating a cluster using the CLI.

1.9.11. RHSA-2025:17232 - OpenShift Container Platform 4.17.41 fixed issues and security update

Issued: 08 October 2025

OpenShift Container Platform release 4.17.41 is now available. The list of fixed issues that are included in the update is documented in the RHSA-2025:17232 advisory. The RPM packages that are included in the update are provided by the RHBA-2025:17230 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 4.17.41 --pullspecs
```

1.9.11.1. Fixed issues

- Before this release, pods on the same node could not reach a pod's secondary **Localnet** interface on a **br-ex** bridge, so the pod defaulted to the primary network. With this release, communication between a **Localnet** pod and a pod running on the same node is possible, when the **Localnet** network's IP addresses are on the same subnet as the host network. (OCPBUGS-59381)

- Before this release, when the **configure-ovs.sh** script activated connection profiles, it triggered an error state on the switch side that caused ports to be disabled. With this release, the **configure-ovs.sh** script waits for normal activation to occur and does not change the profile states, unless it is necessary. This change prevents the link flapping that previously caused problems. (OCPBUGS-60890)

- Before this release, the **PrometheusRemoteWriteBehind** alerts could activate, even when the remote endpoint did not receive any data. With this release, these alerts do not activate when the remote endpoint does not received data. (OCPBUGS-61766)

- Before this release, command line logs from **virt-launcher** pods were not collected across the Kubernetes cluster, making troubleshooting virtual machines difficult. With this release, command line logs from **virt-launcher** pods are collected and saved in JSON format at **namespaces/{namespace-name}/pods/{pod-name}/virt-launcher.json**, so debugging is possible. (OCPBUGS-61775)

- Before this release, a **systemd** service dependency issue in **nmstate-2.2.48** prevented the service from starting on certain nodes. With this release, the fixed **systemd** unit from upstream is deployed, and allows the service to start correctly while new **nmstate** packages containing the fix are pending. (OCPBUGS-61859)

1.9.11.2. Updating

To update an OpenShift Container Platform 4.17 cluster to this latest release, see Updating a cluster using the CLI.

1.9.12. RHBA-2025:16133 - OpenShift Container Platform 4.17.40 fixed issues and security update

Issued: 24 September 2025

OpenShift Container Platform release 4.17.40 is now available. The list of fixed issues that are included in the update is documented in the RHBA-2025:16133 advisory. The RPM packages that are included in the update are provided by the RHBA-2025:16131 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 4.17.40 --pullspecs
```

1.9.12.1. Enhancement

- With this update, **cluster-etcd-operator** now implements a multi-stage notification system for the **etcdDatabaseQuotaLowSpace** alert to proactively manage etcd storage quotas. This enhancement is designed to prevent 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. (OCPBUGS-61337)

1.9.12.2. Fixed issues

- Before this update, image import blocked registries that would fail if those registries were configured with the **NeverContactSource** value, even when mirror registries were set up. With this update, image importing is not blocked when a registry has mirrors configured. This fix ensures that image imports succeed even if the original source was set to **NeverContactSource** in the **ImageDigestMirrorSet** or **ImageTagMirrorSet** resources. (OCPBUGS-53382)

- Before this update, an outdated version of the Azure API prevented the specification of a Capacity Reservation Group for a **MachineSet**, if that group resided in a different subscription than the one originating the server creation. With this release, the most recent version of the Azure API is used, which allows a Capacity Reservation Group for a **MachineSet** to be specified, even when that group is located in a separate subscription from the creation point for the server. (OCPBUGS-56168)

- Before this update, when a **MachineSet** was scaled down and had reached its minimum size, the Cluster Autoscaler could leave the last remaining node with a **NoSchedule** 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 **MachineSet** is scaled down and has reached its minimum size. (OCPBUGS-59266)

- Before this update, when you disabled the OpenShift Container Platform image registry, existing pull secrets caused hung secret deletion during the deletion of a namespace. As a consequence, you could not delete **Dockercfg** secrets after you disabled the OpenShift Container Platform image registry. With this release, existing pull secrets do not block namespace deletion after you remove the registry. (OCPBUGS-61199)

- Before this update, an incorrect **semver** parsing issue caused build failures. With this release, the **semver** parsing issue in the **machine-api-provider-powervs** package is fixed. As a result, the semantic version parsing error in the **PowerVS machine-api-provider** parameter is resolved, which ensures correct version handling and improved stability. (OCPBUGS-61204)

- Before this update, the default **node-monitor-grace-period** value was 50 seconds. As a consequence, nodes did not stay ready for the duration of time that Kubernetes components needed to reconnect, coordinate, and complete their requests. With this release, the default **node-monitor-grace-period** value is 55 seconds. As a result, deployments have enough time to be completed. (OCPBUGS-61290)

1.9.12.3. Updating

To update an OpenShift Container Platform 4.17 cluster to this latest release, see Updating a cluster using the CLI.

1.9.13. RHBA-2025:15344 - OpenShift Container Platform 4.17.39 fixed issues and security update

Issued: 10 September 2025

OpenShift Container Platform release 4.17.39 is now available. The list of fixed issues that are included in the update is documented in the RHBA-2025:15344 advisory. The RPM packages that are included in the update are provided by the RHBA-2025:15323 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 4.17.39 --pullspecs
```

1.9.13.1. Enhancements

- OLM support for tailored network policies in Operators

In OpenShift Container Platform 4.17.39, Operator Lifecycle Manage (OLM) allows 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 4.17.39 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)

- This enhancement reconfigures the **cluster-policy-controller** to only accept connections from the **localhost**, avoiding exposing the network port (10357) outside of the node network. (OCPBUGS-60249)

1.9.13.2. Fixed issues

- Before this update, the '/metrics/usage' endpoint was recently updated to require authentication, which included Cross-Site Request Forgery (CSRF) protections. Requests made to the endpoint began failing with a **forbidden** response because the requests did not include the newly required CSRF token in the cookie. With this release, a CSRF token was added to the request cookie for all calls to the '/metrics/usage' endpoint eliminating the **forbidden** errors. (OCPBUGS-58365)
- Before this update, when a cluster Operator takes a long time to upgrade, the Cluster Version Operator does not report anything because it cannot determine if the upgrade is still progressing or is already stuck. With this release, a new unknown status is added for the failing condition in the status of the cluster vVersion reported by Cluster Version Operator. This alerts the cluster administrators to check the cluster and avoid waiting on a blocked cluster Operator upgrade. (OCPBUGS-58451)
- 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 suboptimal performance. With this release, downloaded workloads consistently schedule on control plane nodes, which improves resource allocation. (OCPBUGS-60298)

- Before this update, hostnames were not obfuscated sufficiently in the Insights Archive on hosted clusters, which was caused by an anonymization method that relied on client host addresses. With this release, all hostnames are now correctly obfuscated in the Insights Archive for hosted clusters. (OCPBUGS-60395)

- Before this update, filtering nodes by role in OpenShift Container Platform web console in versions 4.16 and 4.17 showed all nodes instead of filtered roles, due to an issue. As a consequence, users saw all nodes in the cluster, causing confusion and potential management issues. With this release, filtering nodes now shows only affected roles, improving cluster visibility. As a result, the user experience is improved. (OCPBUGS-60441)

- Before this update, the Azure Disk CSI Driver Operator entered a degraded state after its pod experienced a panic, specifically an **assignment to entry in nil map** error and a Remote Procedure Call (RPC) keep-alive ping timeout. This failure prevented the Operator from reconciling its static resources, creating a significant risk of failures during future cluster upgrades. With this release, the **clustercsidriver** custom resource is deleted, forcing the Operator to re-create and reconcile the object, resolving the panics and ensuring the cluster's stability. (OCPBUGS-60597)

- Before this update, when using multiple recommenders for the Vertical Pod Autoscaler (VPA), a bug caused the default recommender to erroneously delete **VPACheckpoint** objects that belonged to a VPA associated with a non-default recommender. With this release, the default recommender no longer deletes **VPACheckpoint** objects that belong to a non-default recommender. (OCPBUGS-60608)

- Before this update, the HTTP client for an identity provider (IdP) validation did not include the system trust bundles in the bundle it used to validate IdP endpoint requests. This caused IdP validation to fail if the IdP used a publicly trusted endpoint. With this release, the Control Plane Operator trust bundle includes both the system trust bundle and any user-provided additional trust bundles when validating IdP endpoints. (OCPBUGS-61101)

1.9.13.3. Updating

To update an OpenShift Container Platform 4.17 cluster to this latest release, see Updating a cluster using the CLI.

1.9.14. RHSA-2025:14060 - OpenShift Container Platform 4.17.38 bug fix update and security

Issued: 27 August 2025

OpenShift Container Platform release 4.17.38 is now available. The list of fixed issues that are included in the update is documented in the RHSA-2025:14060 advisory. The RPM packages that are included in the update are provided by the RHBA-2025:13976 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 4.17.38 --pullspecs
```

1.9.14.1. Enhancement

- This enhancement adds the **sosreport** command to the tools **imagestream** within OpenShift Container Platform, streamlining debugging for Telco Operators in disconnected environments. (OCPBUGS-56734)

1.9.14.2. Fixed issues

- Before this release, the resources containing the configuration for the vSphere connection would break because of a mismatch between the user interface and API. With this release, the resources do not break because the user interface uses the updated API definition. (OCPBUGS-58337)

- Before this release, when a Machine Set was scaled down and had reached its minimum size, the Cluster Autoscaler could leave the last remaining node with a **NoSchedulle** taint that prevented use of the node. This was caused by a counting error in the Cluster Autoscaler. With this release, the counting error is fixed so that the Cluster Autoscaler works as expected when a Machine Set scales down and reaches its minimum size. (OCPBUGS-59266)

- Before this release, pods on the same node could not reach a pod's secondary **Localnet** interface on a **br-ex** bridge, so the pod defaulted to the primary network. With this release, communication between a **Localnet** pod and a pod running on the same node is possible, provided the **Localnet** network's IP addresses are on the same subnet as the host network. (OCPBUGS-59381)

- Before this release, in multi-zone clusters with a single worker per zone, the Monitoring Operator could degrade if the two nodes running its Prometheus pods rebooted sequentially and each took longer than 15 minutes to recover. With this release, the timeout is extended to 20 minutes reducing the likelihood of the Monitoring Operator entering a degraded state on common cluster topologies. (OCPBUGS-60017)

- Before this release, the image registry would, in some cases, panic when attempting to purge failed uploads from s3-compatible storage providers. This was caused by the image registry's s3 driver mishandling empty directory paths. With this update, the image registry properly handles empty directory paths. (OCPBUGS-60090)

- Before this update, when the Machine Config Operator (MCO) updated the **CoreDNS** template during an upgrade, the next **rpm-ostree** image pull operation would fail with DNS lookup errors while the **CoreDNS** pod was temporarily unavailable. With this release, the operating system update operation now has a retry mechanism allowing image pull retries, ensuring that node upgrades can now proceed and complete successfully. (OCPBUGS-60239)

- Before this release, there was a period when the event data was not yet available when the **cloud-event-proxy** container or pod rebooted. This caused the **getCurrenState** function to incorrectly return a **clockclass** of 0. With this release, the **getCurrentState** function no longer returns a wrong **clockclass** and instead returns an HTTP **400 Bad Request** or **404 Not Found Error**. (OCPBUGS-60267)

1.9.14.3. Updating

To update an OpenShift Container Platform 4.17 cluster to this latest release, see Updating a cluster using the CLI.

1.9.15. RHSA-2025:12437 - OpenShift Container Platform 4.17.37 fixed issues and security update

Issued: 06 August 2025

OpenShift Container Platform release 4.17.37 is now available. The list of fixed issues that are included in the update is documented in the RHSA-2025:12437 advisory. The RPM packages that are included in the update are provided by the RHBA-2025:12438 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 4.17.37 --pullspecs
```

1.9.15.1. Fixed issues

- Before this update, the **catalog-operator** captured snapshots every five minutes, which caused CPU spikes when dealing with many namespaces, subscriptions, and large catalog sources. This increased load on the catalog source pods and prevented users from installing or upgrading Operators. With this release, the catalog snapshot cache lifetime has been increased to 30 minutes allowing enough time for the catalog source to resolve attempts without causing an undue load and stabilizing the Operator installation and upgrade process. (OCPBUGS-57428)

- Before this update, forward slashes were permitted in **console.tab/horizontalNav href** values. Starting in 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 work as expected. (OCPBUGS-59265)

- Before this update, the **Observe → Metrics → query → QueryKebab → Export as csv** drop-down item did not handle an undefined title element. As a consequence, users were unable to export the CSV file for certain queries on the **Metrics** tab of OpenShift Lister version 4.17. With this release, the metrics download for all queries correctly handles the object properties in the drop-down menu items allowing for successful CSV exports. (OCPBUGS-52592)

1.9.15.2. Updating

To update an OpenShift Container Platform 4.17 cluster to this latest release, see Updating a cluster using the CLI.

1.9.16. RHSA-2025:11359 - OpenShift Container Platform 4.17.36 fixed issues and security update

Issued: 23 July 2025

OpenShift Container Platform release 4.17.36 is now available. The list of fixed issues that are included in the update is documented in the RHSA-2025:11359 advisory. The RPM packages that are included in the update are provided by the RHBA-2025:11360 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 4.17.36 --pullspecs
```

1.9.16.1. Fixed issues

- Before this update, the initial loading time for the network plugin pages increased by 10 seconds. With this release, the initial loading time is reduced by correcting the network plugin page delay. (OCPBUGS-58318)
- Before this update, multiple plugins using the same **useModal** hook caused their modals to overwrite each other and led to a loss of functionality for multiple plugins in the user interface. With this release, unique identifiers are used and the modals do not overwrite each other. (OCPBUGS-58224)
- Before this update, a non-default boot image was not updated by the Machine Config Operator (MCO) without degradation, resulting in system instability. With this release, the MCO degradation for non-default boot image updates is fixed, and boot image update issues do not occur. (OCPBUGS-58219)
- Before this update, build failures occurred because of missing service account secrets retrieval. With this release, a fix prevents errors when retrieving service account secrets, and eliminates failures because of a **CannotRetrieveServiceAccount** error. (OCPBUGS-57950)
- Before this update, an OperatorGroup resource reconciliation caused unnecessary **ClusterRole** updates because of the changing order of aggregation rule selectors. As a result, unnecessary etcd writes and authentication cache invalidation occurred. With this release, a specific order in **ClusterRole** aggregation rule selectors reduces unnecessary API server writes. (OCPBUGS-57438)
- Before this update, the Keepalived script for an ingress Virtual IP (VIP) check failed because of a missing **SYS_CHROOT** permission in **chroot**. As a consequence, core ingress services were inaccessible due to incorrect VIP placement. With this release, **chroot** permissions are added to the Keepalived script for an ingress VIP check. As a result, incorrect **chk_default_ingress** permissions are fixed, resulting in correct ingress VIP placement. (OCPBUGS-56625)

1.9.16.2. Updating

To update an OpenShift Container Platform 4.17 cluster to this latest release, see Updating a cluster using the CLI.

1.9.17. RHSA-2025:10294 - OpenShift Container Platform 4.17.35 fixed issues and security update

Issued: 09 July 2025

OpenShift Container Platform release 4.17.35 is now available. The list of fixed issues that are included in the update is documented in the RHSA-2025:10294 advisory. The RPM packages that are included in the update are provided by the RHSA-2025:10295 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 4.17.35 --pullspecs
```

1.9.17.1. Enhancements

- This enhancement extends the expiration date of the self-signed **loopback** certificate for the Kubernetes API Server from one year to three years. (OCPBUGS-57196)

1.9.17.2. Fixed issues

- Previously, the **oc adm node-image create** command incorrectly modified the existing permissions of the target assets folder when the command saved the artifacts on the disk. With this release, a bug fix ensures that the copying operation for the command preserves the destination folder permissions. (OCPBUGS-58091)

- 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, the subnet Availability Zone information in the machine set custom resources and in the EC2 instances now match and the issue is resolved. (OCPBUGS-57293)

- Previously, the kubelet stopped reporting metrics if a **stat** call stalled from the kernel. For example, in instances where a **stat** call on the disk was run on the Network File System (NFS). With this release, the kubelet reports metrics even if a disk is stuck. (OCPBUGS-57289)

- Previously, the /**metrics** endpoint failed to correctly parse a bearer token from the authorization header on internal Prometheus scrape requests. This caused the **TokenReviews** to fail and a **TargetDown** alert was triggered for the console metrics endpoint. With this release, the /**metrics** endpoint correctly parses the bearer token from the authorization header, the **TokenReview** step works as intended, and the **TargetDown** alert no longer displays. (OCPBUGS-57182)

- Previously, an **iptables-alerter** pod had to make several calls to the **crictl** command-line interface (CLI) for each pod that existed in a node to fetch information for the cluster. These calls required high CPU usage that impacted cluster performance. With this release, an **iptables-alerter** pod only needs to make a single call to **crictl** to fetch information for all pods that exist in a node. (OCPBUGS-55518)

- Previously, clusters that did not have the **IdleConnectionTerminationPolicy** API setting in the Ingress Controller API had the **idle-close-on-response** HAProxy setting enabled by default. This resulted in idle connections being closed immediately upon a response. With this release, the **IdleConnectionTerminationPolicy** API setting was added to the Ingress Controller API with **Deferred** as the default, enabling the HAProxy setting and keeping idle connections open until the last response is handled after a soft stop. (OCPBUGS-49702)

1.9.17.3. Updating

To update an OpenShift Container Platform 4.17 cluster to this latest release, see Updating a cluster using the CLI.

1.9.18. RHBA-2025:9289 - OpenShift Container Platform 4.17.34 fixed issues

Issued: 25 June 2025

OpenShift Container Platform release 4.17.34 is now available. The list of fixed issues that are included in the update is documented in the RHBA-2025:9289 advisory. The RPM packages that are included in the update are provided by the RHBA-2025:9290 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 4.17.34 --pullspecs
```

1.9.18.1. Known issues

- A known issue exists where a Technology Preview-enabled cluster has Sigstore verification for payload images in the **policy.json** file, but the Podman version in the base image does not support Sigstore configuration, so the new node is not available. As a workaround, the node starts running when the Podman version in the base image does not support Sigstore, so use the default **policy.json** file that does not have Sigstore verification if the base image is 4.11 or earlier. (OCPBUGS-52313)

1.9.18.2. Fixed issues

- Previously, if you tried to update a hosted cluster that used in-place updates, the proxy variables were not honored and the update failed. With this release, the pod that performs in-place upgrades honors the cluster proxy settings. As a result, updates now work for hosted clusters that use in-place updates. (OCPBUGS-57432)

- Previously, when you defined multiple bring-your-own (BYO) subnet CIDRs for the **machineNetwork** parameter in the **install-config.yaml** configuration file, the installation failed at the bootstrap stage. This situation occurred because the control plane nodes were blocked from reaching the machine config server (MCS) to get their necessary setup configurations. The root cause was an overly strict AWS security group rule that limited MCS access to only the first specified machine network CIDR. With this release, a fix to the AWS security group means that the installation succeeds when multiple CIDRs are specified in the **machineNetwork** parameter of the **install-config.yaml**. (OCPBUGS-57292)

- Previously, a Machine Config Operator (MCO) incorrectly set an **Upgradeable=False** condition to all new nodes that were added to a cluster. A **PoolUpdating** reason was provided for the **Upgradeable=False** condition. With this release, the MCO now correctly sets an **Upgradeable=True** condition to all new nodes that get added to a cluster, which resolves the issue. (OCPBUGS-57135)

- Previously, the installation program 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-56448)

- Previously, in certain situations the gateway IP address for a node changed and caused the **OVN** cluster router to add a new static route with the new gateway IP address, without deleting the original one. The **OVN** cluster router manages the static route to the cluster subnet. 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-56443)

- 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-56244)

1.9.18.3. Updating

To update an OpenShift Container Platform 4.17 cluster to this latest release, see Updating a cluster using the CLI.

1.9.19. RHSA-2025:8552 - OpenShift Container Platform 4.17.33 fixed issues and security update

Issued: 11 June 2025

OpenShift Container Platform release 4.17.33 is now available. The list of fixed issues that are included in the update is documented in the RHSA-2025:8552 advisory. The RPM packages that are included in the update are provided by the RHBA-2025:8553 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 4.17.33 --pullspecs
```

1.9.19.1. Known issues

- If an egress IP is moved to a different node while a pod has a connection established to an external system, the pod will not receive any traffic from the external system on the same connection until it sends some traffic through it. Currently, there is no workaround for this issue. To keep pods connected, pods must regularly send some traffic through any connection they open via any egress IP in order to keep both directions of the connection working during egress IP failovers. (OCPBUGS-58355)

1.9.19.2. Fixed issues

- Previously, the OpenShift Container Platform web console sent you an alert that compute nodes must be updated within 60 days after you updated control plane nodes. This action request was invalid. With this update, the OpenShift Container Platform web console no longer sends you this invalid alert. (OCPBUGS-56375)

1.9.19.3. Updating

To update an OpenShift Container Platform 4.17 cluster to this latest release, see Updating a cluster using the CLI.

1.9.20. RHSA-2025:8280 - OpenShift Container Platform 4.17.32 fixed issues and security update

Issued: 04 June 2025

OpenShift Container Platform release 4.17.31 is now available. The list of fixed issues that are included in the update is documented in the RHSA-2025:8280 advisory. The RPM packages that are included in the update are provided by the RHBA-2025:8281 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 4.17.32 --pullspecs
```

1.9.20.1. Fixed issues

- Previously, a bug fix altered the availability set configuration by changing the fault domain count to use the maximum available value instead of being fixed at 2. This inadvertently caused scaling issues for **MachineSet** objects created before the bug fix, because the controller attempted to modify immutable availability sets. With this release, availability sets are no longer modified after creation, allowing affected **MachineSet** objects to scale properly. (OCPBUGS-56655)

- Previously, the Samples Operator updated the Progressing condition's **lastTransitionTime** value even when the condition's status did not actually change. This led to potential installation errors and perceived instability for end users. With this release, the Operator is prevented from updating the **lastTransitionTime** value unless there is a status change. This enhances Operator stability, minimizes installer errors, and ensures a smoother user experience. (OCPBUGS-55800)

- Previously, the Samples Operator watched all cluster Operators in the cluster, which triggered the Cluster Samples Operator sync loop to run unnecessarily. This behavior negatively impacted overall performance. With this release, the Cluster Samples Operator only watches specific cluster Operators. (OCPBUGS-55795)

- Previously, the Grandmaster Timekeeper (T-GM) operation unexpectedly set the Precision Time Protocol (PTP) announce message internal signal flags incorrectly. This caused the loss of time synchronization across the network. With this release, the PTP announce message flags are correctly initialized, ensuring accurate and standardized time synchronization information is disseminated across the network. (OCPBUGS-55740)

- Previously, image pull timeouts occurred due to the Zscaler platform scanning all data transfers. This resulted in timed out image pulls. With this release, the image pull timeout is increased to 30 seconds, allowing successful updates. (OCPBUGS-54664)

1.9.20.2. Updating

To update an OpenShift Container Platform 4.17 cluster to this latest release, see Updating a cluster using the CLI.

1.9.21. RHBA-2025:8108 - OpenShift Container Platform 4.17.31 fixed issues

Issued: 28 May 2025

OpenShift Container Platform release 4.17.31 is now available. The list of fixed issues that are included in the update is documented in the RHBA-2025:8108 advisory. There are no RPM packages for this release.

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 4.17.31 --pullspecs
```

1.9.21.1. Fixed issues

- Previously, OpenShift Container Platform 4.15 and later versions managed by OpenShift Lifecycle Manager (OLM) were required to have the **olm.managed: "true"** label. In some cases, the solution failed to start and entered a **CrashLoopBackOff** state if the label was missing. The logs for this scenario were displayed as informative, which made it more challenging to identify the root cause. For this release, the log level is changed to error to make the issue clearer and easier to diagnose when the label is missing. (OCPBUGS-56250)

- Previously, if the default proxy environment variables were set to null on build containers, some applications in the container would not run. With this release, the proxy environment variables are added to the build container only if they are defined and the default values are not null. (OCPBUGS-55826)

- Previously, an Operator updated the Progressing condition's **lastTransitionTime** value even when the condition's status did not actually change. This led to potential installation errors and perceived instability for end users. With this release, the Operator is prevented from updating the **lastTransitionTime** value unless there is a status change. This enhances Operator stability, minimizes installer errors, and ensures a smoother user experience. (OCPBUGS-55800)

- Previously, the Cluster Samples Operator watched all cluster Operators in the cluster, which triggered the Cluster Samples Operator sync loop to run unnecessarily. This behavior negatively impacted overall performance. With this release, the Cluster Samples Operator only watches specific cluster Operators. (OCPBUGS-55795)

- Previously, the Grandmaster Timekeeper (T-GM) operation unexpectedly set the Precision Time Protocol (PTP) announce message internal signal flags incorrectly. This caused the loss of time synchronization across the network. With this release, the PTP announce message internal signal flags are correctly initialized, ensuring accurate and standardized time synchronization information is disseminated across the network. (OCPBUGS-55740)

- Previously, when creating a multi-network policy without specifying a protocol, the Open Virtual Network (OVN) would crash. With this release, the protocol is assumed to be the Transmission Control Protocol (TCP) if none is specified, preventing OVN crashes. (OCPBUGS-52480)

1.9.21.2. Updating

To update an OpenShift Container Platform 4.17 cluster to this latest release, see Updating a cluster using the CLI.

1.9.22. RHSA-2025:7669 - OpenShift Container Platform 4.17.30 bug fix and security update advisory

Issued: 21 May 2025

OpenShift Container Platform release 4.17.30 is now available. The list of bug fixes that are included in the update is documented in the RHSA-2025:7669 advisory. The RPM packages that are included in the update are provided by the RHBA-2025:7671 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 4.17.30 --pullspecs
```

1.9.22.1. Bug fixes

- Previously, no event was logged when an error occurred from a failed ingress-to-route conversion. With this update, this error appears in the **event** logs. (OCPBUGS-55943)
- Previously, a bug caused excessive updates to the **progressing** condition because of unsorted failing image imports. This led to unnecessary resource consumption for users. With this release, the excessive updates are fixed in the Samples Operator by sorting failing image imports. As a result, the Operator performance is improved by reducing unnecessary updates for the image imports. (OCPBUGS-55894)
- Previously, when Microsoft Azure Spot virtual machine (VM) de-allocation occurred during provisioning, the machine controller entered a loop, leading to Spot VM provisioning failures and unavailable nodes. With this release, the **deallocate eviction** policy is replaced with the **delete eviction policy** in Azure Spot VM provisioning. As a result, the machine controller is more resilient and it no longer enters a loop during provisioning. (OCPBUGS-55729)

1.9.22.2. Updating

To update an OpenShift Container Platform 4.17 cluster to this latest release, see Updating a cluster using the CLI.

1.9.23. RHSA-2025:4723 - OpenShift Container Platform 4.17.29 bug fix and security update advisory

Issued: 15 May 2025

OpenShift Container Platform release 4.17.29 is now available. The list of bug fixes that are included in the update is documented in the RHSA-2025:4723 advisory. The RPM packages that are included in the update are provided by the RHBA-2025:4725 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 4.17.29 --pullspecs
```

1.9.23.1. Bug fixes

- Previously, a race condition in the hostname caused inconsistencies between the node and machine hostnames. With this release, the race condition is resolved, which ensures consistent hostnames in the Ignition configuration file during operating system installation. (OCPBUGS-55680).
- Previously, the Subject Alternative Name (SAN) of the custom certificate that the user added to the **hc.spec.configuration.apiServer.servingCerts.namedCertificates** field conflicted with the hostname set in the **hc.spec.services.servicePublishingStrategy** field for the Kubernetes agent server (KAS). As a result, the KAS certificate was not added to the set of certificates to

generate a new payload, causing certificate validation issues for nodes that joined the hosted cluster. With this release, the validation fails earlier so that the user is warned about the issue with the conflicting SANs. (OCPBUGS-55500).

- Previously, the boot image updates failed when the Amazon Machine Image (AMI) for a specified region was not found. This issue occurred because AMIs for all regions are not published in the **scos.json** file for the installation program. With this release, update failures are prevented by using the **us-east-1** region by default for any unavailable regions during Amazon Web Services (AWS) boot image updates. (OCPBUGS-55490).

- Previously, when viewing the list of installed Operators, an Operator appeared twice in the list. This duplication occurred when the currently selected project matched an Operator's default namespace while copied cluster service versions (CSVs) were disabled in the Operator Lifecycle Manager (OLM). With this release, the Operator appears once. (OCPBUGS-55415).

1.9.23.2. Updating

To update an OpenShift Container Platform 4.17 cluster to this latest release, see Updating a cluster using the CLI.

1.9.24. RHSA-2025:4431 - OpenShift Container Platform 4.17.28 bug fix and security update advisory

Issued: 9 May 2025

OpenShift Container Platform release 4.17.28 is now available. The list of bug fixes that are included in the update is documented in the RHSA-2025:4431 advisory. The RPM packages that are included in the update are provided by the RHBA-2025:4433 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 4.17.28 --pullspecs
```

1.9.24.1. Bug fixes

- Previously, if you had permission to view nodes but not Certificate Signing Requests (CSR), you could not access the **Nodes list** page. With this release, permissions to view CSRs are no longer required to access the **Nodes list** page. (OCPBUGS-55202)

- Previously, after you deleted the **ClusterResourceOverride** custom resource (CR) or you uninstalled the Cluster Resource Override Operator, which also removes the **ClusterResourceOverride** CR, the **v1.admission.autoscaling.openshift.io** API service becomes unreachable. This situation impacted other cluster functions, such as other Operator installations from succeeding. With this release, when you delete the Cluster Resource Override Operator, the **v1.admission.autoscaling.openshift.io** API service is also removed. As a result, you can now install other Operators without experiencing installation failures. (OCPBUGS-55355)

- Previously, when you attempted to upgrade the Cluster Resource Override Operator from OpenShift Container Platform 4.16 to 4.17, the Cluster Resource Override webhook stopped functioning. This situation prevented pods from getting created in namespaces that had the Cluster Resource Override enabled. With this release, a stale secret is deleted so that OpenShift

Container Platform regenerates the secret with the correct parameters and values during an upgrade operation. As a result, the Operator upgrade succeeds and you can now create pods in any namespaces that have the Cluster Resource Override enabled. (OCPBUGS-55239)

- Previously, the Assisted Installer failed to detect World Wide Name (WWN) details during Fibre Channel multipath volumes hardware discovery. As a result, a Fibre Channel multipath disk could not be matched with a WWN root device. This meant that when you specified a **wwn** root device hint, the hint excluded all Fibre Channel multipath disks. With this release, the Assisted Installer now detects WWN details during Fibre Channel multipath disk discovery. If multiple Fibre Channel multipath disks exist, you can now use the **wwn** root device hint to choose a primary disk for your cluster. (OCPBUGS-55184)

- Previously, the **mtu-migration** service did not work correctly when you used **nmstate** to manage a **br-ex** bridge because of a missing service dependency. With this release, the service dependency is now added so that a network configuration that uses **nmstate** to manage a **br-ex** is correct before the migration process begins. (OCPBUGS-54830)

1.9.24.2. Updating

To update an OpenShift Container Platform 4.17 cluster to this latest release, see Updating a cluster using the CLI.

1.9.25. RHSA-2025:4204 - OpenShift Container Platform 4.17.27 bug fix and security update advisory

Issued: 6 May 2025

OpenShift Container Platform release 4.17.27 is now available. The list of bug fixes that are included in the update is documented in the RHSA-2025:4204 advisory. The RPM packages that are included in the update are provided by the RHBA-2025:4206 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 4.17.27 --pullspecs
```

1.9.25.1. Known issues

- There is a known issue when the grandmaster clock (T-GM) transitions to the Locked state too soon. This happens before the Digital Phase-Locked Loop (DPLL) completes its transition to the **Locked-HO-Acquired** state, and after the Global Navigation Satellite Systems (GNSS) time source is restored. (**OCPBUGS-54534**)

1.9.25.2. Bug fixes

- Previously, when you selected the load balancer, the installation program picked a fixed internet protocol (IP) address, (**10.0.0.100**), and attached the address to the load balancer even if the IP was outside of the range of the machine network or virtual network. With this release, the installation program checks for an available IP in the provided control plane subnet or machine network and elects an IP that is not reserved if the default IP is not within the range. (OCPBUGS-55224)

- Previously, if a scrape failed, Prometheus erroneously considered the samples from the very

next scrape as duplicates and dropped them. This issue impacted 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-54941)

- Previously, for an Ingress resource with an **IngressWithoutClassName** alert, the Ingress Controller did not delete the alert along with deletion of the resource. The alert continued to show on the OpenShift Container Platform web console. With this release, the Ingress Controller resets the **openshift_ingress_to_route_controller_ingress_without_class_name** metric to **0** before the controller deletes the Ingress resource, so that the alert is deleted and no longer shows on the web console. (OCPBUGS-53077)

- Previously, during cluster creation a control plane node was replaced when it was detected as unhealthy. This replacement irrevocably disabled the cluster and prevented the creation of the cluster. With this fix, the node is not inadvertently replaced, ensuring the stabilization of the control plane and the successful creation of the cluster. (OCPBUGS-52957)

- Previously, the Single Root I/O Virtualization (SR-IOV) virtual function (VF) did not revert any unexpected value changes to the maximum transmission unit (MTU) value when a pod was deleted. This issue occurred if the application inside the pod had its MTU value changed; in turn, the pod would also have its MTU value changed. With this release, the SR-IOV Container Network Interface (CNI) now reverts any unexpected MTU value changes to the original value so that this issue no longer exists. (OCPBUGS-54392)

1.9.25.3. Updating

To update an OpenShift Container Platform 4.17 cluster to this latest release, see Updating a cluster using the CLI.

1.9.26. RHSA-2025:4012 - OpenShift Container Platform 4.17.26 bug fix and security update advisory

Issued: 24 April 2025

OpenShift Container Platform release 4.17.26 is now available. The list of bug fixes that are included in the update is documented in the RHSA-2025:4012 advisory. The RPM packages that are included in the update are provided by the RHBA-2025:4014 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 4.17.26 --pullspecs
```

1.9.26.1. Bug fixes

- Previously, when you attempted to create a validating webhook for a resource that was managed by the **oauth** API server, the validating webhook was not created. This issue occurred because of a communication issue with the **oauth** API server and the data plane. With this release, a Konnectivity proxy sidecar has been added to bridge communications between the **oauth** API server and the data plane so that you can now create a validating webhook for any resource that the **oauth** API server manages. (OCPBUGS-54841)

- Previously, virtual machines (VMs) in a cluster that ran on Microsoft Azure failed because the

attached network interface controller (NIC) was in a **ProvisioningFailed** state. With this release, the Machine API controller now checks the provisioning status of a NIC and refreshes the VMs on a regular basis to prevent this issue. (OCPBUGS-54393)

- Previously, the installation program malfunctioned if it attempted to retrieve Google Cloud tags over an unstable network, or when it could not reach the GCP server. With this release, the issue is resolved. (OCPBUGS-51210)

- Previously, a User Datagram Protocol (UDP) packet that was larger than the maximum transmission unit (MTU) value set for the cluster, could not be sent to the endpoint of the packet by using a service. With this release, the pod IP address is used instead of the service IP address regardless of the packet size, so that the UDP packet can be sent to the endpoint. (OCPBUGS-50579)

1.9.26.2. Updating

To update an OpenShift Container Platform 4.17 cluster to this latest release, see Updating a cluster using the CLI.

1.9.27. RHSA-2025:3798 - OpenShift Container Platform 4.17.25 bug fix and security update advisory

Issued: 16 April 2025

OpenShift Container Platform release 4.17.25 is now available. The list of bug fixes that are included in the update is documented in the RHSA-2025:3798 advisory. The RPM packages that are included in the update are provided by the RHBA-2025:3800 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 4.17.25 --pullspecs
```

1.9.27.1. Bug fixes

- Previously, containers that use the SELinux domain of **container_logreader_t** for the purposes of viewing container logs on a host at **/var/log** could not access logs in the **/var/log/containers** subdirectory. This issue happened because of a missing symbolic link. With this release, a symbolic link is created for **/var/log/containers** so that containers can access the logs in **/var/log/containers**. (OCPBUGS-54343)

- Previously, the cluster autoscaler stopped scaling when a machine failed in a machine set. This situation happened 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 has a more accurate count. (OCPBUGS-54325)

- Previously, the **Alerts** page on the Developer perspective of the web console stopped querying the Prometheus tenancy path. This issue caused an **Error loading silences from alert manager** banner to show on the page. With this release, the page now queries the Prometheus tenancy path and the page retrieves silent alert data from the Developer perspective data store so that the banner no longer shows on the page. (OCPBUGS-54211)

- Previously, a missing machine config for the container runtime configuration prevented a cluster update operation from succeeding because of a container runtime controller failure. With this

release, the missing machine config is now ignored so that a cluster operation can succeed. (OCPBUGS-52188)

1.9.27.2. Updating

To update an OpenShift Container Platform 4.17 cluster to this latest release, see Updating a cluster using the CLI.

1.9.28. RHSA-2025:3565 - OpenShift Container Platform 4.17.24 bug fix and security update advisory

Issued: 9 April 2025

OpenShift Container Platform release 4.17.24 is now available. The list of bug fixes that are included in the update is documented in the RHSA-2025:3565 advisory. The RPM packages that are included in the update are provided by the RHBA-2025:3567 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 4.17.24 --pullspecs
```

1.9.28.1. Bug fixes

- Previously, an update to the IBM Cloud® Cloud Internet Services (CIS) implementation impacted the upstream Terraform plugin. If you attempted to create an external-facing cluster on IBM Cloud®, an error occurred. With this release, you can create an external cluster on OpenShift Container Platform without the plugin issue. (OCPBUGS-54357)
- Previously, when users tried building the agent ISO in a disconnected setup, an error occurred. With this release, the setup completes without an error. (OCPBUGS-53378)
- Previously, the **ovn-ipsec-host** pod failed with a crash loop on RHEL worker nodes because of a missing shared library during the container execution. With this release, the **ovn-ipsec-host** pod successfully starts on the worker node without an error. (OCPBUGS-52951)
- Previously, the Operator Lifecycle Manager (OLM) CSV annotation contained unexpected JSON data, which was successfully parsed, but then resulted in a runtime error when attempting to use the resulting value. With this release, JSON values from OLM annotations are validated before use, errors are logged, and the console does not fail when an unexpected JSON is received in an annotation. (OCPBUGS-51277)

1.9.28.2. Updating

To update an OpenShift Container Platform 4.17 cluster to this latest release, see Updating a cluster using the CLI.

1.9.29. RHSA-2025:3297 - OpenShift Container Platform 4.17.23 bug fix and security update advisory

Issued: 3 April 2025

OpenShift Container Platform release 4.17.23 is now available. The list of bug fixes that are included in the update is documented in the RHSA-2025:3297 advisory. The RPM packages that are included in the update are provided by the RHBA-2025:3299 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 4.17.23 --pullspecs
```

1.9.29.1. Bug fixes

- Previously, the Operator Marketplace and the Operator Lifecycle Manager (OLM) used an older version, v1.24, of the **pod-security.kubernetes.io/** label. With this release, the namespace where Operator Marketplace is deployed now uses the Pod Security Admission (PSA) label marked as **latest**. (OCPBUGS-53283)

- Previously, the **openshift-install agent create pxe-files** command created temporary directories in **/tmp/agent** and the command did not remove these directories upon command completion. With this release, the command now removes the directories upon completion, so no you do not need to manual deleted the directories. (OCPBUGS-52961)

- Previously, a code migration operation failed to process external labels correctly on the **Alert detail** page on the **Administrator perspective** of the web console. These external labels are required to prevent silenced alert notifications from getting added to notification bell icon. Because the **Alert detail** page did not handle external labels correctly, the notification bell provided links to these **Alert detail** pages that generated a **no matching alerts found** message when you clicked on a link. With this release, the **Alert detail** page accepts external labels so clicking on an alert in the notification bell links to the correct **Alert detail** page. (OCPBUGS-51117)

- Previously, when you created a cluster that includes the following **kubevirt** CR configuration, you received a **failed to reconcile virt launcher policy: could not determine if <address_name>** is an IPv4 or IPv6 address` error message:

  ```
  # ...
  - service: APIServer
    servicePublishingStrategy:
      type: NodePort
      nodePort:
        address: <address_name>
        port: 305030
  # ...
  ```

  This error message was generated because network policies were not properly deployed on virtual machine (VM) namespaces. With this release, a fix means that you can add a host name address to the **nodePort.address** configuration of the CR so that network policies can be properly deployed on VMs. (OCPBUGS-48439)

- Previously, the Single Root I/O Virtualization (SR-IOV) network config daemon unbinded network drivers from the physical function (PF) interface instead of unbinding the drivers from the virtual function (VF) interface when SR-IOV was configured with an InfiniBand (IB) type. This unbinding workflow removed the IB interface from the node, and this sitaution made the IB interface non-functional. With this release, a fix to the SR-IOV network config daemon ensures

that the IB interface remains functional when it correctly unbinds the VF network interface. Additionally, the SR-IOV Network Operator targets the network drivers of a VF interface instead of the PF interface when configuring SR-IOV with an IB type. (OCPBUGS-53254)

1.9.29.2. Updating

To update an OpenShift Container Platform 4.17 cluster to this latest release, see Updating a cluster using the CLI.

1.9.30. RHSA-2025:3059 - OpenShift Container Platform 4.17.22 bug fix and security update advisory

Issued: 26 March 2025

OpenShift Container Platform release 4.17.22 is now available. The list of bug fixes that are included in the update is documented in the RHSA-2025:3059 advisory. The RPM packages that are included in the update are provided by the RHSA-2025:3061 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 4.17.22 --pullspecs
```

1.9.30.1. Bug fixes

- Previously, during cluster shutdown, a race condition prevented a stage **ostree** deployment from finalizing if the deployment was moved to a staging location during a reboot operation. With this release, a fix removes the race condition from the **ostree** deployment so that the staged deployment can finalize even during a reboot operation. (OCPBUGS-53225)

1.9.30.2. Updating

To update an OpenShift Container Platform 4.17 cluster to this latest release, see Updating a cluster using the CLI.

1.9.31. RHSA-2025:2696 - OpenShift Container Platform 4.17.21 bug fix and security update advisory

Issued: 19 March 2025

OpenShift Container Platform release 4.17.21 is now available. The list of bug fixes that are included in the update is documented in the RHSA-2025:2696 advisory. The RPM packages that are included in the update are provided by the RHBA-2025:2698 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 4.17.21 --pullspecs
```

1.9.31.1. Bug fixes

- Previously, the **trusted-ca-bundle-managed** config map was a mandatory component. If you attempted to use a custom Public Key Infrastructure (PKI), the deployment would fail because the OpenShift API server expected the presence of the **trusted-ca-bundle-managed** config map. With this release, you can deploy clusters without the **trusted-ca-bundle-managed** config map when you use a custom PKI. (OCPBUGS-52657)

- Previously, the **Observe** section on the web console did not show items contributed from plugins unless certain flags related to the monitoring were set. However, these flags prevented other plugins, such as logging, distributed tracing, network observability, and so on, from adding items to the **Observe** section. With this release, the monitoring flags are removed so that other plugins can add items to the **Observe** section. (OCPBUGS-52205)

- Previously, a custom Security Context Constraint (SCC) impacted pods that were 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-50589)

- Previously, you could not create **NodePool** resources with ARM64 architecture on non-AWS or Azure platforms. This bug resulted in validation errors that prevented the addition of bare-metal compute nodes and caused Common Expression Language (CEL) validation blocks when creating a **NodePool** resource. The fix modifies the **NodePool** spec validation rules to allow ARM64 architecture on non-AWS or Azure by setting **None** for the **self.platform.type** section. You can now create **NodePool** with ARM64 architecture specifications on non-AWS or Azure bare-metal platforms. (OCPBUGS-46440)

- Previously, if you deleted a bare-metal host with a related data image, the data image stayed present. With this release, the issue is resolved, and the data image is deleted with the bare-metal host as expected. (OCPBUGS-42387)

1.9.31.2. Updating

To update an OpenShift Container Platform 4.17 cluster to this latest release, see Updating a cluster using the CLI.

1.9.32. RHSA-2025:2445 - OpenShift Container Platform 4.17.20 bug fix and security update advisory

Issued: 12 March 2025

OpenShift Container Platform release 4.17.20 is now available. The list of bug fixes that are included in the update is documented in the RHSA-2025:2445 advisory. The RPM packages that are included in the update are provided by the RHBA-2025:2447 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 4.17.20 --pullspecs
```

1.9.32.1. Bug fixes

- Previously, Local Storage Operator (LSO) ignored existing Small Computer System Interface (SCSI) symlinks during persistent volumes (PV) creation. With this release, the LSO no longer ignores these symlinks because it gathers these symlinks before finding new symlinks when

creating a PV. (OCPBUGS-51291)

- Previously, when the OVN-Kubernetes network plugin and the Kubernetes-NMState Operator interacted with each other, unexpected connection profiles persisted on disk storage. These connection profiles sometimes caused the **ovs-configuration** service to fail when restarted. With this release, the unnecessary connection profiles are now cleaned before the **ovs-configuration** starts so that this issue no longer occurs. (OCPBUGS-52257)

- Previously, the **vmware-vsphere-csi-driver-operator** Container Storage Interface (CSI) driver entered panic mode when the VMware vSphere vCenter address was incorrect or missing. With this release, the CSI driver does not go into panic mode if the VMware vSphere vCenter address is incorrect or missing. (OCPBUGS-52207)

- 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-51292)

- Previously, the alerts links on the **Alert rules** page of the **Developer** perspective included external labels to invalid links. This happened because the URL for the **Alerts** page did not expect external labels. With this release, the external labels are no longer added to the alerts URL on the **Alert rules** page so the alerts links are accurate. (OCPBUGS-51126)

- Previously, for a **kubevirt-csi** pod that ran on a node in a hosted cluster, the persistent volume claim (PVC) from the hosted cluster was removed from the virtual machine (VM) after the VM was rebooted. However, the **VolumeAttachment** resource was not removed and this caused issues for the cluster as it expected the PVC to be attached to the VM. With this release, after a VM is rebooted, the **VolumeAttachment** resource is removed so that the cluster issues no longer occur. (OCPBUGS-44623)

- Previously, in bare-metal configurations where the provisioning network was disabled but the **bootstrapProvisioningIP** field was set, the bare-metal provisioning components would fail to start. These failures occur when the provisioning process reconfigures the external network interface on the bootstrap VM during the process of pulling container images. With this release, dependencies are added to ensure that interface reconfiguration only occurs when the network is idle, preventing conflicts with other processes. As a result, the bare-metal provisioning components now start reliably, even when the **bootstrapProvisioningIP** field is set and the provisioning network is disabled. (OCPBUGS-43528)

1.9.32.2. Updating

To update an OpenShift Container Platform 4.17 cluster to this latest release, see Updating a cluster using the CLI.

1.9.33. RHSA-2025:1912 - OpenShift Container Platform 4.17.19 bug fix and security update advisory

Issued: 5 March 2025

OpenShift Container Platform release 4.17.19 is now available. The list of bug fixes that are included in the update is documented in the RHSA-2025:1912 advisory. The RPM packages that are included in the update are provided by the RHSA-2025:1914 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 4.17.19 --pullspecs
```

1.9.33.1. Bug fixes

- Previously, when the OpenShift cluster was created with a secure proxy enabled, and the certificate was set in **configuration.proxy.trustCA**, the cluster failed to complete provisioning. With this release, you can create a cluster with secure proxy enabled with the certificate set in **configuration.proxy.trustCA**. Also, the fix prevents an issue that prevented **oauth** from connecting to Cloud APIs using the management cluster proxy. (OCPBUGS-51098)

- Previously, when you deleted a Dynamic Host Configuration Protocol (DHCP) network on an IBM Power Virtual Server cluster, subresources still existed. With this release, when you delete a DHCP network, the subresources deletion occurs before continuing with the delete operation. (OCPBUGS-50967)

- Previously, when a worker node tried to join a cluster, the rendezvous node rebooted before the process completed. Because the worker node could not communicate with the rendezvous node, the installation was not successful. With this release, a patch is applied that fixes the racing condition that caused the rendezvous node to reboot prematurely and the issue is resolved. (OCPBUGS-50011)

- Previously, the DNS-based egress firewall incorrectly disallowed creation of rules containing DNS names in uppercase. With this release, the issue is fixed and the egress firewall is created with uppercase DNS names. (OCPBUGS-49961)

- Previously, all host validation status logs referred to the name of the first host registered. When a host validation failed, it was not possible to decide the host with an issue. With this release, the correct host is identified in each log message and the host validation logs correctly how the host that they are linked to. (OCPBUGS-44058)

- Previously, when the VMware vSphere vCenter cluster contained an ESXi host that did not have a standard port group defined and the installation program tried to select that host to import the Open Virtual Appliance (OVA), the import failed and the error **Invalid Configuration for device 0** was reported. With this release, the installation program verifies whether a standard port group for an ESXi host is defined and, if not, continues verifying until it locates an ESXi host with a defined standard port group, or reports an error message if it fails to locate one. (OCPBUGS-37945)

1.9.33.2. Updating

To update an OpenShift Container Platform 4.17 cluster to this latest release, see Updating a cluster using the CLI.

1.9.34. RHSA-2025:1703 - OpenShift Container Platform 4.17.18 bug fix and security update advisory

Issued: 26 February 2025

OpenShift Container Platform release 4.17.18 is now available. The list of bug fixes that are included in the update is documented in the RHSA-2025:1703 advisory. The RPM packages that are included in the update are provided by the RHBA-2025:1706 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 4.17.18 --pullspecs
```

1.9.34.1. Bug fixes

- Previously, the control plane Operator did not honor the set PROXY environment variables when it checked the API endpoint availability. With this release, the issue is resolved. (OCPBUGS-50596)
- Previously, when installing a cluster on Amazon Web Services (AWS) in existing subnets that were located in edge zones, such as a AWS Local Zone or a Wavelength Zone, the **kubernetes.io/cluster/<InfraID>:shared** tag was missing in the subnet resources of the edge zone. With this release, a fix ensures that all subnets that are used in the **install-config.yaml** configuration file have the required tag. (OCPBUGS-49975)
- Previously, incorrect addresses were being passed to the Kubernetes **EndpointSlice** on a cluster, and this issue prevented the installation of the MetalLB Operator on an Agent-based cluster in an IPv6 disconnected environment. With this release, a fix modifies the address evaluation method. Red Hat Marketplace pods can now successfully connect to the cluster API server, so that the installation of MetalLB Operator and handling of ingress traffic in IPv6 disconnected environments can occur. (OCPBUGS-46665)
- Previously, the method to validate the container image architecture did not go through the image metadata provider. As a consequence, the image overrides did not take effect. With this release, the methods on the image metadata provider were modified to allow multi-architecture validations, and those methods were propagated through all components for image validation steps. As a result, the issue is resolved. (OCPBUGS-46664)

1.9.34.2. Updating

To update an OpenShift Container Platform 4.17 cluster to this latest release, see Updating a cluster using the CLI.

1.9.35. RHSA-2025:1403 - OpenShift Container Platform 4.17.17 bug fix and security update advisory

Issued: 18 February 2025

OpenShift Container Platform release 4.17.17 is now available. The list of bug fixes that are included in the update is documented in the RHSA-2025:1403 advisory. The RPM packages that are included in the update are provided by the RHBA-2025:1405 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 4.17.17 --pullspecs
```

1.9.35.1. Bug fixes

- Previously, the availability set fault domain count was hardcoded to **2**. This value works in most regions on Microsoft Azure because the fault domain counts are typically at least **2**, but failed in the **centraluseuap** and **eastusstg** regions. With this release, the availability set fault domain

count in a region is set dynamically so that this issue no longer occurs. (OCPBUGS-50017)

- Previously, when installing a cluster on Google Cloud, a cluster installation failed when an instance required IP Forwarding to be disabled. With this release, IP Forwarding is disabled for all Google Cloud machines before cluster installation so that the cluster installation issue no longer occurs. (OCPBUGS-49993)
- Previously, you could not install an cluster on AWS in the **ap-southeast-5** region or other regions because the OpenShift Container Platform internal registry did not support these regions. With this release, the internal registry is updated to include the following regions so that this issue no longer occurs:
  - **ap-southeast-5**
  - **ap-southeast-7**
  - **ca-west-1**
  - **il-central-1**
  - **mx-central-1**
    (OCPBUGS-49695)
- Previously, when a pod was running on a node on which egress IPv6 is assigned, the pod was not able to communicate with the Kubernetes service in a dual-stack cluster. This resulted in traffic with the IP family, that the **egressIP** object is not applicable to, being dropped. With this release, only the source network address translation (SNAT) for the IP family that the egress IP applied to is deleted, eliminating the risk of traffic being dropped. (OCPBUGS-48828)
- Previously, when you attempted to use the hosted control planes CLI to create a cluster in a disconnected environment, the installation command failed. An issue existed with the registry that hosts the command. With this release, a fix to the command registry means that you can use the hosted control planes CLI to create a cluster in a disconnected environment. (OCPBUGS-48170)

1.9.35.2. Updating

To update an OpenShift Container Platform 4.17 cluster to this latest release, see Updating a cluster using the CLI.

1.9.36. RHSA-2025:1120 - OpenShift Container Platform 4.17.16 bug fix and security update advisory

Issued: 11 February 2025

OpenShift Container Platform release 4.17.16 is now available. The list of bug fixes that are included in the update is documented in the RHSA-2025:1120 advisory. The RPM packages that are included in the update are provided by the RHSA-2025:1122 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 4.17.16 --pullspecs
```

1.9.36.1. Bug fixes

- Previously, the Bare Metal Operator (BMO) created the **HostFirmwareComponents** custom resource for all Bare-metal hosts (BMH), including ones based on the intelligent platform management interface (IPMI), which did not support it. With this release, **HostFirmwareComponents** custom resources are only created for BMH that support it. (OCPBUGS-49701)

- Previously, importing manifest lists could cause an API crash if the source registry returned an invalid sub-manifest result. With this update, the API flags an error on the imported tag instead of crashing. (OCPBUGS-49399)

- Previously, the Konnectivity proxy used by the **openshift-apiserver** in the control plane resolved registry names with cloud API suffixes on the control plane and then attempted to access them through the data plane.
  A hosted cluster that used the no-egress feature in ROSA, and a container registry that was accessible through an Amazon Virtual Private Cloud (VPC) endpoint was created but failed to install because **imagestreams** that use the container registry did not resolve. With this release, the Konnectivity proxy resolves and routes hostnames consistently. (OCPBUGS-46465)

- Previously, if you ran a build that required a trust bundle to access registries, it did not pick up the bundle configured in the cluster proxy. The builds failed if they referenced a registry required for a custom trust bundle. With this release, builds that require the trust bundle specified in the proxy configuration succeed, and the issue is resolved. (OCPBUGS-45268)

- Previously, when you attempted to use the hosted control planes CLI to create a hosted control planes cluster, the installation failed because of a release image check on multi-arch images. With this release, an update to the hosted control planes CLI codebase fixes the issue so that the release image check does not fail when checking for multi-arch images. (OCPBUGS-44927)

- Previously, installing an AWS cluster in either the Commercial Cloud Services (C2S) region or the Secret Commercial Cloud Services (SC2S) region failed because the installation program added unsupported security groups to the load balancer. With this release, the installation program no longer adds unsupported security groups to the load balancer for a cluster that needs to be installed in either the C2S region or SC2S region. (OCPBUGS-42763)

1.9.36.2. Updating

To update an OpenShift Container Platform 4.17 cluster to this latest release, see Updating a cluster using the CLI.

1.9.37. RHSA-2025:0876 - OpenShift Container Platform 4.17.15 bug fix and security update advisory

Issued: 5 February 2025

OpenShift Container Platform release 4.17.15 is now available. The list of bug fixes that are included in the update is documented in the RHSA-2025:0876 advisory. The RPM packages that are included in the update are provided by the RHSA-2025:0878 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 4.17.15 --pullspecs
```

1.9.37.1. Bug fixes

- Previously, when you used the installation program to install a cluster in a Prism Central environment, the installation failed because a **prism-api** call that loads an RHCOS image timed out. This issue happened because the **prismAPICallTimeout** parameter was set to **5** minutes. With this release, the **prismAPICallTimeout** parameter in the **install-config.yaml** configuration file now defaults to **10** minutes. You can also configure the parameter if you need a longer timeout for a **prism-api** call. (OCPBUGS-49362)

- Previously, every time a subcription was reconciled, the OLM catalog Operator requested a full view of the catalog metadata from the catalog source pod of the subscription. These requests caused performance issues for the catalog pods. With this release, the OLM catalog Operator now uses a local cache that is refreshed periodically and reused by all subscription reconciliations, so that the performance issue for the catalog pods no longer persists. (OCPBUGS-48695)

- Previously, if you specified a **forceSelinuxRelabel** field in the **ClusterResourceOverride** CR and then modified the CR at a later stage, the Cluster Resource Override Operator did not apply the update to the associated **ConfigMap** resource. This **ConfigMap** resource is important for an SELinux relabeling feature, **forceSelinuxRelabel**. With this release, the Cluster Resource Override Operator now applies and tracks any **ClusterResourceOverride** CR changes to the **ConfigMap** resource. (OCPBUGS-48691)

1.9.37.2. Updating

To update an OpenShift Container Platform 4.17 cluster to this latest release, see Updating a cluster using the CLI.

1.9.38. RHSA-2025:0654 - OpenShift Container Platform 4.17.14 bug fix and security update advisory

Issued: 28 January 2025

OpenShift Container Platform release 4.17.14 is now available. The list of bug fixes that are included in the update is documented in the RHSA-2025:0654 advisory. The RPM packages that are included in the update are provided by the RHSA-2025:0656 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 4.17.14 --pullspecs
```

1.9.38.1. Bug fixes

- Previously, some cluster autoscaler metrics were not initialized and were unavailable. With this release, the cluster autoscaler metrics are initialized and available. (OCPBUGS-48606)

- Previously, you could not add a new worker by using the **oc adm node-image create** command if the node date or time was inaccurate. With this release, the issue is resolved by applying the same NTP configuration that is in the target cluster **machineconfig chrony** resource to the node ephemeral live environment. (OCPBUGS-45344)

- Previously, you could not use all available machine types in a zone because all zones in a region were assumed to have the same set of machine types. With this release, all machine types are available in all enabled zones. (OCPBUGS-46432)

- Previously, the installation program was not compliant with PCI-DSS/BAFIN regulations. With this release, the cross-tenant replication in Microsoft Azure is disabled, which reduces the chance of unauthorized data access and ensures strict adherence to data governance policies. (OCPBUGS-48119)

- Previously, when you clicked the **Don't show again** link in the Red Hat Ansible Lightspeed modal, it did not display the correct **General User Preference** tab when one of the other **User Preference** tabs was open. With this release, clicking the **Don't show again** link goes to the correct **General User Preference** tab. (OCPBUGS-48227)

- Previously, when a Google Cloud Platform (GCP) service account was created, the account would not always be immediately available. When the account was not available for updates, the installation program received failures when adding permissions to the account. According to **Retry failed requests**, a service account might be created, but is not active for up to 60 seconds. With this release, the service account is updated on an exponential backoff to give the account enough time to update correctly. (OCPBUGS-48359)

- Previously, on RHEL 9 FIPS STIG compliant machines, The SHA-1 key caused the release signature verification to fail because of the restriction to use weak keys. With this release, the key used by the oc-mirror plugin for release signature verification is changed and release images are signed by a new SHA256 trusted-key that is different from the old SHA-1 key. (OCPBUGS-48363)

- Previously, the Operator Lifecycle Manager (OLM) would sometimes concurrently resolve the same namespace in a cluster. This led to subscriptions reaching a terminal state of **ConstraintsNotSatisfiable**, because two concurrent processes interacted with a subscription and this caused a CSV file to become unassociated. With this release, OLM can resolve concurrent namespaces for a subscription so no CSV remains unassociated. (OCPBUGS-45845)

1.9.38.2. Updating

To update an OpenShift Container Platform 4.17 cluster to this latest release, see Updating a cluster using the CLI.

1.9.39. RHSA-2025:0115 - OpenShift Container Platform 4.17.12 bug fix and security update advisory

Issued: 14 January 2025

OpenShift Container Platform release 4.17.12 is now available. The list of bug fixes that are included in the update is documented in the RHSA-2025:0115 advisory. The RPM packages that are included in the update are provided by the RHBA-2025:0118 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 4.17.12 --pullspecs
```

1.9.39.1. Updating

To update an OpenShift Container Platform 4.17 cluster to this latest release, see Updating a cluster using the CLI.

1.9.40. RHBA-2025:0023 - OpenShift Container Platform 4.17.11 bug fix and security update advisory

Issued: 8 January 2025

OpenShift Container Platform release 4.17.11 is now available. The list of bug fixes that are included in the update is documented in the RHBA-2025:0023 advisory. The RPM packages that are included in the update are provided by the RHBA-2025:0026 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 4.17.11 --pullspecs
```

1.9.40.1. Enhancements

1.9.40.1.1. GCP Filestore supporting Workload Identity is generally available

Google Compute Platform (GCP) Filestore Container Storage Interface (CSI) storage supports Workload Identity. This allows users to access Google Cloud resources using federated identities instead of a service account key. This feature is generally available.

For more information, see Google Compute Platform Filestore CSI Driver Operator .

1.9.40.2. Bug fixes

- Previously, the certificate signing request (CSR) approver included certificates from other systems when it calculated if it should stop approving certificates when the system was overloaded. In larger clusters, where other subsystems used CSRs, the CSR approver determined that there were many unapproved CSRs and prevented additional approvals. With this release, the CSR approver prevents new approvals when there are many CSRs for the **signerName** values that it observes, but has not been able to approve. The CSR approver now only includes CSRs that it can approve, using the **signerName** property as a filter. (OCPBUGS-46429)

- Previously, a hard eviction of a pod in a node caused a pod to enter a termination grace period instead of instantly shutting down and deleted by the kubelet. Each pod that enters a termination grace period exhausts the node resources. With this release, a bug fix ensures that a pod enters a one-second termination grace period so the kubelet can shut down and then delete the pod. (OCPBUGS-46364)

- Previously, the permissions **ec2:AllocateAddress** and **ec2:AssociateAddress** were not verified when the **PublicIpv4Pool** feature was used, which resulted in permission failures during the installation. With this release, the required permissions are validated before the cluster is installed. (OCPBUGS-46360)

- Previously, users could enter an invalid string for any CPU set in the performance profile, resulting in a broken cluster. With this release, the fix ensures that only valid strings can be entered, eliminating the risk of cluster breakage. (OCPBUGS-45964)

- Previously, in certain scenarios, an event was missed by the informer watch stream. If an object was deleted while this disconnection occurred, the informer returned an unexpected type, which caused a stale state. As a result, the incorrect returned type caused a problem. With this release, the unexpected types are correctly handled, and the temporary disconnection possibilities are successful. (OCPBUGS-46039)

1.9.40.3. Updating

To update an OpenShift Container Platform 4.17 cluster to this latest release, see Updating a cluster using the CLI.

1.9.41. RHBA-2024:11522 - OpenShift Container Platform 4.17.10 bug fix and security update advisory

Issued: 2 January 2025

OpenShift Container Platform release 4.17.10 is now available. The list of bug fixes that are included in the update is documented in the RHBA-2024:11522 advisory. The RPM packages that are included in the update are provided by the RHBA-2024:11525 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 4.17.10 --pullspecs
```

1.9.41.1. Enhancements

1.9.41.1.1. Node Tuning Operator architecture detection

The Node Tuning Operator can now properly select kernel arguments and management options for Intel and AMD CPUs. (OCPBUGS-43664)

1.9.41.2. Bug fixes

- Previously, when the webhook token authenticator was enabled and had the authorization type set to **None**, the OpenShift Container Platform web console would consistently crash. With this release, a bug fix ensures that this configuration does not cause the OpenShift Container Platform web console to crash. (OCPBUGS-46390)

- Previously, the **SiteConfig** custom resource (CR) configuration, which configures ingress rules and sevices for a cluster, caused the **BareMetalHost** CR to remain in a deleted state instead of being deleted as part of the cluster cleanup operation. With this release, this issue no longer occurs provided that you update to the GitOps Operator to version 1.13 or a later version. (OCPBUGS-46071)

- Previously, when you attempted to use Operator Lifecycle Manager (OLM) to upgrade an Operator, the upgrade was blocked and an **error validating existing CRs against new CRD's schema** message was generated. An issue existed with OLM, whereby OLM erroneously identified incompatibility issues validating existing custom resources (CRs) against the new Operator version's custom resource definitions (CRDs). With this release, the validation is corrected so that Operator upgrades are no longer blocked. (OCPBUGS-46054)

- Previously, a **PipelineRuns** CR that used a resolver could not be rerun on the OpenShift

Container Platform web console. If you attempted ro rerun the CR, an **Invalid PipelineRun configuration, unable to start Pipeline** message was generated. With this release, you can now rerun a **PipelineRuns** CR that uses resolver without experience this issue. (OCPBUGS-45949)

- 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-45948)

- Previously, the **aws-sdk-go-v2** software development kit (SDK) failed to authenticate an **AssumeRoleWithWebIdentity** API operation on an AWS Security Token Service (STS) cluster. With this release, the pod identity webhook now includes a default region so that this issue no longer persists. (OCPBUGS-45938)

- Previously, installation of an AWS cluster failed in certain environments on existing subnets when the **publicIp** parameter of the **MachineSet** object was explicity 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-45186)

- Previously, an additional filtering property was passed into the component that is used to list operands on the **Operator details** page. The additional property caused the list to always be empty if it was extended by a dynamic plugin. With this release, the extra property has been removed so that available operands are listed as expected. (OCPBUGS-45667)

- Previously, when you ran the **oc adm node-image create** command, the command would sometimes fail and output an **image can't be pulled** error message. With this release, a fix adds a retry mechanism to the command so that if the command fails to pull images from a release workload, a retry operation ensures the command runs as expected. (OCPBUGS-45517)

- Previously, an IBM Power® Virtual Server cluster installation failed on installer-provisioned infrastructure because the installation program used a random network type instead of using of the specified network type that was specified in the **install-config.yaml** configuration file. With this release, the installation program now uses the network type that is specified in the **install-config.yaml** so that this issue no longer persists. (OCPBUGS-45484)

- 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-44644)

1.9.41.3. Updating

To update an OpenShift Container Platform 4.17 cluster to this latest release, see Updating a cluster using the CLI.

1.9.42. RHBA-2024:11010 - OpenShift Container Platform 4.17.9 bug fix and security update advisory

Issued: 19 December 2024

OpenShift Container Platform release 4.17.9 is now available. The list of bug fixes that are included in the update is documented in the RHBA-2024:11010 advisory. The RPM packages that are included in the update are provided by the RHBA-2024:11013 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 4.17.9 --pullspecs
```

1.9.42.1. Known issues

- If you plan to deploy the NUMA Resources Operator, avoid using OpenShift Container Platform versions 4.17.7 or 4.17.8. (OCPBUGS-45639)

1.9.42.2. Bug fixes

- Previously, the Google Cloud **Project Number input** field was incorrectly labeled as **GCP Pool ID**. With this release, the Google Cloud **Project Number input** field is correctly labeled. (OCPBUGS-46000)
- Previously, there was a maximum bulk delete limit of 10. This limit caused an issue with the **PreferNoSchedule** taint. With this release, the maximum bulk delete rate is disabled. (OCPBUGS-45929)
- Previously, users wanted to configure their Amazon Web Services DHCP option set with a custom domain name that contained a trailing period. When the hostname of EC2 instances were converted to Kubelet node names, the trailing period was not removed. Trailing periods are not allowed in a Kubernetes object name. With this release, trailing periods are allowed in a domain name in a DHCP option set. (OCPBUGS-45918)
- Previously, when a long string of individual CPUs were in the Performance Profile, the machine configurations were not processed. With this release, the user input process is updated to use a sequence of numbers or a range of numbers on the kernel command line. (OCPBUGS-45627)
- Previously, when accessing the image from the release payload to run the **oc adm node-image** command, the command failed. With this release, a retry mechanism is added to correct the temporary failures when the image is accessed. (OCPBUGS-45517)
- Previously, the first reboot failed while running Agent-based Installer using FCP or NVME storage devices for multiple images on s390x hardware. With this release, this issue is resolved and the reboot completes. (OCPBUGS-44904)
- Previously, a missing permission caused cluster deprovision to fail when using a custom identity and access management (IAM) profile. With this release, the list of required permissions includes **tag:UntagResource** and the cluster deprovision completes. (OCPBUGS-44848)
- Previously, when you created a hosted cluster by using a shared VPC where the private DNS hosted zones existed in the cluster creator account, the private link controller failed to create the **route53** DNS records in the local zone. With this release, the ingress shared role adds

records to the private link controller. The VPC endpoint is used to share the role to create the VPC endpoint in the VPC owner account. A hosted cluster is created in a shared VPC configuration, where the private hosted zones exist in the cluster creator account. (OCPBUGS-44630)

- Previously, the **kdump initramfs** stopped responding when opening a local encrypted disk, even when the kdump destination was a remote machine that did not need to access the local machine. With this release, this issue is fixed and the **kdump initranfs** successfully opens a local encrypted disk. (OCPBUGS-43079)

- Previously, the Cluster Version Operator (CVO) did not filter internal errors that were propogated to the **ClusterVersion Failing condition** message. As a result, errors that did not negatively impact the update were shown for the **ClusterVersion Failing condition** message. With this release, the errors that are propogated to the **ClusterVersion Failing condition** message are filtered. (OCPBUGS-39558)

1.9.42.3. Updating

To update an OpenShift Container Platform 4.17 cluster to this latest release, see Updating a cluster using the CLI.

1.9.43. RHSA-2024:10818 - OpenShift Container Platform 4.17.8 bug fix and security update advisory

Issued: 11 December 2024

OpenShift Container Platform release 4.17.8 is now available. The list of bug fixes that are included in the update is documented in the RHSA-2024:10818 advisory. The RPM packages that are included in the update are provided by the RHBA-2024:10821 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 4.17.8 --pullspecs
```

1.9.43.1. Bug fixes

- Previously, the provider ID for the **IBMPowerVSCluster** object did not populate properly due to an improper retrieval of the IBM Cloud Workspace ID. As a result, certificate signing requests (CSRs) were pending in the hosted cluster. With this release, the provider ID successfully populates and the issue is resolved. (OCPBUGS-44880)

- Previously, you could not remove a finally pipeline task from the edit Pipeline form if you created a pipeline with only one finally task. With this change, you can remove the finally task from the edit Pipeline form and the issue is resolved. (OCPBUGS-44873)

- Previously, if you used the **oc adm node-image create** command and the image generation step failed, it reported a simple error and did not show the log for the container. As a result, the error message did not show the underlying issue that caused the image generation step to fail. With this release, the **oc adm node-image create** command shows the log for the container. (OCPBUGS-44508)

- Previously, if you created load balancers for a cluster that you wanted to install on IBM Power®

and the creation of the load balancers timed out, the installation of the cluster failed and did not report the error. The cluster failed because both internal and external DNS load balancer names were not created. With this release, if internal and external DNS load balancer names do not exist during cluster installation, the installation program produces an error notification, and you can add the names so that the cluster installation process continues. (OCPBUGS-44247)

- Previously, the IDs that were used to determine the number of rows in a Dashboard table were not unique, and some rows were combined if their IDs were the same. With this release, the ID uses more information to prevent duplicate IDs and the table can display each expected row. (OCPBUGS-43441)

1.9.43.2. Updating

To update an OpenShift Container Platform 4.17 cluster to this latest release, see Updating a cluster using the CLI.

1.9.44. RHSA-2024:10518 - OpenShift Container Platform 4.17.7 bug fix and security update advisory

Issued: 3 December 2024

OpenShift Container Platform release 4.17.7 is now available. The list of bug fixes that are included in the update is documented in the RHSA-2024:10518 advisory. The RPM packages that are included in the update are provided by the RHBA-2024:10521 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 4.17.7 --pullspecs
```

1.9.44.1. Enhancements

1.9.44.1.1. Deprecated clusterTasks OpenShift Pipelines version 1.17

- The OpenShift Container Platform 4.17 release deprecates the **clusterTasks** resource from Red Hat OpenShift Pipelines version 1.17. The release also removes **clusterTasks** resource dependencies from the static plugin on the OpenShift Pipelines page of the OpenShift Container Platform web console. (OCPBUGS-44183)

1.9.44.2. Bug fixes

- Previously, you could not enter multi-line parameters in a custom template, such as private keys. With this release, you can switch between single-line and multi-line modes in a custom template so that you can enter multi-line inputs in a template field. (OCPBUGS-44699)

- Previously, when you attempted to use the Cluster Network Operator (CNO) to upgrade a cluster with existing **localnet** networks, **ovnkube-control-plane** pods would fail to run. This happened because the **ovnkube-cluster-manager** container could not process an OVN-Kubernetes **localnet** topology network that did not have subnets defined. With this release, a fix ensures that the **ovnkube-cluster-manager** container can process an OVN-Kubernetes **localnet** topology network that does not have subnets defined. (OCPBUGS-43454)

- Previously, when you attempted to scale a **DeploymentConfig** object with an admission

webhook that the object's **deploymentconfigs/scale** subresource, the **apiserver** failed to handle the request. This impacted the **DeploymentConfig** object as it could not be scaled. With this release, a fix ensures that this issue no longer occurs. (OCPBUGS-42752)

1.9.44.3. Updating

To update an OpenShift Container Platform 4.17 cluster to this latest release, see Updating a cluster using the CLI.

1.9.45. RHBA-2024:10137 - OpenShift Container Platform 4.17.6 bug fix and security update advisory

Issued: 26 November 2024

OpenShift Container Platform release 4.17.6 is now available. The list of bug fixes that are included in the update is documented in the RHBA-2024:10137 advisory. The RPM packages that are included in the update are provided by the RHBA-2024:10140 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 4.17.6 --pullspecs
```

1.9.45.1. Enhancements

1.9.45.1.1. Updating to Kubernetes version 1.30.6

OpenShift Container Platform release 4.17.6 contains the changes that come from the update to Kubernetes version 1.30.6. (OCPBUGS-44512)

1.9.45.2. Bug fixes

- Previously, if you set custom annotations in a custom resource (CR), the SR-IOV Operator would override all the default annotations in the **SriovNetwork** CR. With this release, when you define custom annotations in a CR, the SR-IOV Operator does not override the default annotations. (OCPBUGS-42252)
- Previously, in the Red Hat OpenShift Container Platform web console **Notifications** section, silenced alerts were visible in the notification drawer because the alerts did not include external labels. With this release, the alerts include external labels so that silenced alerts are not visible on the notification drawer. (OCPBUGS-44722)
- Previously, when you clicked the **start lastrun** option in the Red Hat OpenShift Container Platform web console **Edit BuildConfig** page, an error prevented the **lastrun** operation from running. With this release, a fix ensures that **start lastrun** option runs as expected. (OCPBUGS-44587)
- Previously, OpenShift Container Platform 4.17 introduced collapsing and expanding the **Getting started** section on the **Administrator perspective** of the Red Hat OpenShift Container Platform web console. When you either collapsed or expanded the section, you could not close the section. With this release, a fix ensures that you can now close the **Getting started** section. (OCPBUGS-44586)

- Previously, the **MachineConfig** tab on the **Details** page of the Red Hat OpenShift Container Platform web console displayed an error when one or more **spec.config.storage.files** did not include data fields that were marked as optional. With this update, a fix ensures that this error no longer displays if you populated no values in the optional fields. (OCPBUGS-44479)

- Previously, a Hosted control planes cluster that used the IBM® platform was unable to accept authenticatation from an **oc login** command. This behavior caused an error for the web broswer where the browser could not fetch the token from the cluster. With this release, a fix ensures that cloud-based endpoints do not get proxied so that authentication by using the **oc login** command works as expected. (OCPBUGS-44276)

- Previously, if the **RendezvousIP** matched a substring in the **next-hop-address** field of a compute node configuration, a validation error. The **RendezvousIP** must match only a control plane host address. With this release, a substring comparison for **RendezvousIP** is used only against a control plane host address, so that the error no longer exists. (OCPBUGS-44261)

- Previously, when you created load balancers for a cluster that you wanted to install on IBM Power® and the creation of the load balancers timed out, the installation of the cluster failed and did not report the error. The cluster failed because both internal and external DNS load balancer names were not created. With this release, if internal and external DNS load balancer names do not exist during cluster installation, the installation program outputs an error so that you have an opportunity to add the names so that the cluster installation process can continue. (OCPBUGS-44247)

- Previously, when you attempted to run a disk cleanup operation on a physical storage device that used thin provisioning, the cleanup operation failed. With this release, a bug fix now ensures that you can run a disk cleanup operation on a physical storage device without the cleanup operation failing. (OCPBUGS-31570)

- Previously, if the Operator Lifecycle Manager (OLM) was unable to access the secret associated with a service account, OLM would rely on the Kubernetes API server to automatically create a bearer token. With this release, Kubernetes versions 1.22 and later do not automatically create the bearer token. Instead, OLM now uses the **TokenRequest** API to request a new Kubernetes API token. (OCPBUGS-44760)

1.9.45.3. Updating

To update an OpenShift Container Platform 4.17 cluster to this latest release, see Updating a cluster using the CLI.

1.9.46. RHSA-2024:9610 - OpenShift Container Platform 4.17.5 bug fix and security update advisory

Issued: 19 November 2024

OpenShift Container Platform release 4.17.5, which includes security updates, is now available. The list of bug fixes that are included in the update is documented in the RHSA-2024:9610 advisory. The RPM packages that are included in the update are provided by the RHSA-2024:9613 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 4.17.5 --pullspecs
```

1.9.46.1. Enhancements

1.9.46.1.1. Improving the validation criteria for Cluster Monitoring Operator

- With this release, the Cluster Monitoring Operator (CMO) has improved validation criteria. The CMO blocks cluster updates with configurations in **openshift-monitoring/cluster-monitoring-config** or **openshift-user-workload-monitoring/user-workload-monitoring-config** that include unsupported fields or misconfigurations. (OCPBUGS-43690)

1.9.46.2. Bug fixes

- Previously, the Azure File Driver attempted to reuse existing storage accounts. With this release, the Azure File Driver creates storage accounts during dynamic provisioning. For updated clusters, newly-created Persistent Volumes use a new storage account. Persistent Volumes that were previously provisioned continue using the same storage account used before the cluster update. (OCPBUGS-42949)

- Previously, when you used the **must-gather** tool, a Multus Container Network Interface (CNI) log file, **multus.log**, was stored in a node's file system. This situation caused the tool to generate unnecessary debug pods in a node. With this release, the Multus CNI no longer creates a **multus.log** file, and instead uses a CNI plugin pattern to inspect any logs for Multus DaemonSet pods in the **openshift-multus** namespace. (OCPBUGS-42835)

- Previously, the task data did not fully load on ArtifactHub when you attempted to create a pipeline. With this release, the console fully loads the data from ArtifactHub and the issue is resolved. (OCPBUGS-16141)

1.9.46.3. Updating

To update an OpenShift Container Platform 4.17 cluster to this latest release, see Updating a cluster using the CLI.

1.9.47. RHSA-2024:8981 - OpenShift Container Platform 4.17.4 bug fix and security update advisory

Issued: 13 November 2024

OpenShift Container Platform release 4.17.4, which includes security updates, is now available. The list of bug fixes that are included in the update is documented in the RHSA-2024:8981 advisory. The RPM packages that are included in the update are provided by the RHSA-2024:8984 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 4.17.4 --pullspecs
```

1.9.47.1. Enhancements

1.9.47.1.1. Authenticating customer workloads with GCP Workload Identity

With this release, applications in customer workloads on OpenShift Container Platform clusters that use Google Cloud Platform Workload Identity can authenticate by using GCP Workload Identity.

To use this authentication method with your applications, you must complete configuration steps on the cloud provider console and your OpenShift Container Platform cluster.

For more information, see Configuring GCP Workload Identity authentication for applications on Google Cloud.

1.9.47.1.2. ansible-operator upstream version information

- The **ansible-operator** version now shows the corresponding upstream version information. (OCPBUGS-43836)

1.9.47.2. Bug fixes

- Previously, when installing a cluster on an IBM platform and adding an existing VPC to the cluster, the Cluster API Provider IBM Cloud would not add ports 443, 5000, and 6443 to the security group of the VPC. This situation prevented the VPC from being added to the cluster. With this release, a fix ensures that the Cluster API Provider IBM Cloud adds the ports to the security group of the VPC so that the VPC gets added to your cluster. (OCPBUGS-44226)

- Previously, the installation program populated the **network.devices**, **template** and **workspace** fields in the **spec.template.spec.providerSpec.value** section of the VMware vSphere control plane machine set custom resource (CR). These fields should be set in the vSphere failure domain, and the installation program populating them caused unintended behaviors. Updating these fields did not trigger an update to the control plane machines, and these fields were cleared when the control plane machine set was deleted. With this release, the installation program is updated to no longer populate values that are included in the failure domain configuration. If these values are not defined in a failure domain configuration, for instance on a cluster that is updated to OpenShift Container Platform 4.17 from an earlier version, the values defined by the installation program are used. (OCPBUGS-44047)

- Previously, enabling ESP hardware offload using IPSec on attached interfaces in Open vSwitch broke connectivity due to a bug in Open vSwitch. With this release, OpenShift automatically disables ESP hardware offload on the Open vSwitch attached interfaces, and the issue is resolved. (OCPBUGS-43917)

- Previously, if you configured the identity provider (IDP) name for OAuth custom resource (CR) to contain whitespaces, **oauth-server** crashed. With this release, an identity provider (IDP) name that contains whitespaces does not cause **oauth-server** to crash. (OCPBUGS-44118)

- Previously, due to a behavior regression in Go 1.22, **oauth-server** pods crashed if the IDP configuration contained multiple password-based IDPs, such as **htpasswd**, with at least one of them having spaces in its name. Note that if the bootstrap user ,**kubeadmin**, still exists in a cluster, the user also counts as a password-based IDP. With this release, a fix to the **oauth-server** resolves this issue and prevents the server from crashing. (OCPBUGS-43587)

- Previously, when you used the Agent-based Installer to install a cluster on a node that had an incorrect date, the cluster installation failed. With this release, a patch is applied to the Agent-based Installer live ISO time synchronization. The patch configures the **/etc/chrony.conf** file with the list of additional Network Time Protocol (NTP) servers, so that you can set any of these additional NTP servers in the **agent-config.yaml** without experiencing a cluster installation issue. (OCPBUGS-43846)

- Previously, when you installed a private cluster on Google Cloud, the API firewall rule used the source range of **0.0.0.0/0**. This address allowed non-cluster resources unintended access to the private cluster. With this release, the API firewall rule now only allows resources that have source

ranges in the Machine Network to access the private cluster. (OCPBUGS-43786)

- Previously, an invalid or unreachable identity provider (IDP) blocked updates to hosted control planes. With this release, the **ValidIDPConfiguration** condition in the **HostedCluster** object now reports any IDP errors so that these errors do not block updates to hosted control planes. (OCPBUGS-43746)

- Previously, the **attach**, **oc exec**, and **port-forward** commands received an error if you ran the commands through a proxy. With this release, a patch applied to kubectl ensures kubectl can handle any proxy errors with these commands, so that the commands run as expected. (OCPBUGS-43696)

- Previously, Red Hat Enterprise Linux (RHEL) CoreOS templates that were shipped by the Machine Config Operator (MCO) caused node scaling to fail on Red Hat OpenStack Platform (RHOSP). This issue happened because of an issue with **systemd** and the presence of a legacy boot image from older versions of OpenShift Container Platform. With this release, a patch fixes the issue with **systemd** and removes the legacy boot image, so that node scaling can continue as expected. (OCPBUGS-42577)

- Previously, if a Cluster Version Operator (CVO) pod restarted while the pod was initializing a syncing operation, the guard for a blocked upgrade request failed. As a consequence, the blocked upgrade request was unexpectedly accepted. With this release, a CVO pod postpones the reconciliation of requests during initialization, so that the guard of a blocked upgrade requests persists after the CVO pod restarts. (OCPBUGS-42386)

1.9.47.3. Updating

To update an OpenShift Container Platform 4.17 cluster to this latest release, see Updating a cluster using the CLI.

1.9.48. RHSA-2024:8434 - OpenShift Container Platform 4.17.3 bug fix and security update advisory

Issued: 29 October 2024

OpenShift Container Platform release 4.17.3, which includes security updates, is now available. The list of bug fixes that are included in the update is documented in the RHSA-2024:8434 advisory. The RPM packages that are included in the update are provided by the RHSA-2024:8437 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 4.17.3 --pullspecs
```

1.9.48.1. Enhancements

1.9.48.1.1. Exposing network overlap metrics with the Cluster Network Operator

- When you start the limited live migration method and an issue exists with network overlap, the Cluster Network Operator (CNO) can expose network overlap metrics for the issue. This is possible because the **openshift_network_operator_live_migration_blocked** metric now includes the new **NetworkOverlap** label. (OCPBUGS-39121)

1.9.48.1.2. Loading git repository environment variables automatically from your repository

- Previously, when you imported a git repository by using the serverless import strategy, the environment variables from the **func.yaml** file were not automatically loaded into the form. With this update, the environment variables are now loaded upon import. (OCPBUGS-42474)

1.9.48.2. Bug fixes

- Previously, a regression was introduced to the OpenShift installer during a Power VS deployment. As a result, the security group rules required for OpenShift Install were not created. With this release, the issue is resolved. (OCPBUGS-43547)
- Previously, if an image registry Operator was configured with the **networkAccess** field set to **Internal** in Azure, an authorization error prevented the image registry Operator from deleting the storage container, and the **managementState** field from being set to **Removed**. With this release, the Operator can delete the storage account and storage container, and the **managementState** field can successfully set to **Removed**. (OCPBUGS-43350)
- Previously, both active and passive high-availability (HA) deployments ran three replicas instead of the required two. As a result, control planes contained more pods than required, which led to scaling issues. With this release, the number of replicas in active and passive HA deployments is reduced from three to two. (OCPBUGS-42704)
- Previously, the configuration loader logged **yaml** unmarshal errors when the INI succeeded. With this release, the unmarshal errors are no longer logged when the INI succeeds. (OCPBUGS-42327)
- Previously, the Machine Config Operator (MCO)'s vSphere **resolve-prepender** script used **systemd** directives that were incompatible with old bootimage versions used in OpenShift Container Platform 4. With this release, nodes can scale using newer bootimage versions 4.17 4.13 and above, through manual intervention, or by upgrading to a release that includes this fix. (OCPBUGS-42108)
- Previously, the installation program did not validate the maximum transmission unit (MTU) for a custom IPv6 network on Red Hat Enterprise Linux CoreOS (RHCOS). If you specified a low value for the MTU, the installation of the cluster would fail. With this release, the minimum MTU value for IPv6 networks is set to **1380**, where **1280** is the minimum MTU for IPv6 and **100** is the OVN-Kubernetes encapsulation overhead. With this release, the installation program now validates the MTU for a custom IPv6 network on Red Hat Enterprise Linux CoreOS (RHCOS). (OCPBUGS-41812)
- Previously, the **rpm-ostree-fix-shadow-mode.service** service would run when you ran RHCOS in a live environment. As a result, the **rpm-ostree-fix-shadow-mode.service** service logged a failure that did not impact the deployment or live system. With this release, the **rpm-ostree-fix-shadow-mode.service** service will not run if the RHCOS is not running from an installed environment and the issue is resolved. (OCPBUGS-41621)

1.9.48.3. Updating

To update an OpenShift Container Platform 4.17 cluster to this latest release, see Updating a cluster using the CLI.

1.9.49. RHSA-2024:8229 - OpenShift Container Platform 4.17.2 bug fix and security update advisory

Issued: 23 October 2024

OpenShift Container Platform release 4.17.2, which includes security updates, is now available. The list of bug fixes that are included in the update is documented in the RHSA-2024:8229 advisory. The RPM packages that are included in the update are provided by the RHBA-2024:8232 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 4.17.2 --pullspecs
```

1.9.49.1. Enhancements

- The Operator SDK now correctly scaffolds the Dockerfile for **kube-rbac-proxy**. Additionally, the Operator SDK can now use **-rhel9** container images. (OCPBUGS-42953)

- When you start the limited live migration method and an issue exists with network overlap, the Cluster Network Operator (CNO) can now expose network overlap metrics for the issue. This is possible because the **openshift_network_operator_live_migration_blocked** metric now includes the new **NetworkOverlap** label. (OCPBUGS-39121)

1.9.49.2. Bug fixes

- Previously, the approval mechanism for certificate signing requests (CSRs) failed because the node name and internal DNS entry for a CSR did not match in terms of character case differences. With this release, an update to the approval mechanism for CSRs skips case-sensitive checks so that a CSR with a matching node name and internal DNS entry does not fail the check because of character case differences. (OCPBUGS-43312)

- Previously, a port conflict on the Cloud Credential Operator (CCO) and the **assisted-service** object caused cluster installations on VMware vSphere to fail. With this release, the installation program ignores the **pprof** module in the **assisted-service** object so that the port conflict no longer exists. (OCPBUGS-43069)

- Previously, when you attempted to use the **oc import-image** command to import an image in a hosted control planes cluster, the command failed because of access issues with a private image registry. With this release, an update to **openshift-apiserver** pods in a hosted control planes cluster resolves names that use the data plane so that the **oc import-image** command now works as expected with private image registries. (OCPBUGS-43051)

- Previously, for managed services on hosted control planes, audit logs were sent to a local webhook service, **audit-webhook**. This caused issues for hosted control planes pods that sent audit logs through the **konnectivity** service. With this release, **audit-webhook** is added to the list of **no_proxy** hosts so that hosted control planes pods can send auti logs to the **audit-webhook** service. (OCPBUGS-42974)

- Previously, when you used the Agent-based Installer to install a cluster, **assisted-installer-controller** timed out or exited the installation process depending on whether **assisted-service** was unavailable on the rendezvous host. This situation caused the cluster installation to fail during CSR approval checks. With this release, an update to **assisted-installer-controller** ensures that the controller does not timeout or exit if **assisted-service** is unavailable. The CSR approval check now works as expected. (OCPBUGS-42839)

- Previously, running the **openshift-install gather bootstrap --dir <workdir>** command might cause the installation program to skip the analysis of the collected logs. The command would output the following message:

  ```
  Invalid log bundle or the bootstrap machine could not be reached and bootstrap logs were not collected
  ```

  With this release, the installation program can now analyze log handles that the **gather bootstrap --dir <workdir>** argument generates. (OCPBUGS-42806)

- Previously, when you used the **Developer** perspective with custom editors on the OpenShift Container Platform web console, entering the **n** keyboard shortcut caused the namespace menu to unexpectedly open. This issue happened because the keyboard shortcut key did not account for custom editors. With this release, the namespace menu now accounts for custom editors and does not unexpectedly open when you enter the **n** keyboard shortcut. (OCPBUGS-42607)

- Previously, the installation program did not validate the maximum transmission unit (MTU) for a custom IPv6 network on Red Hat Enterprise Linux CoreOS (RHCOS). If you specified a low value for the MTU, installation of the cluster would fail. With this release, the minimum MTU value for IPv6 networks is set to **1380**, where **1280** is the minimum MTU for IPv6 and **100** is the OVN-Kubernetes encapsulation overhead. With this release, the installation program now validates the MTU for a custom IPv6 network on Red Hat Enterprise Linux CoreOS (RHCOS) (OCPBUGS-41812)

- Previously, a hosted control planes cluster that used mirroring release images might result in existing node pools to use the hosted cluster's operating system version instead of the **NodePool** version. With this release, a fix ensures that node pools use their own versions. (OCPBUGS-41552)

- Previously, after you creating a private OpenShift Container Platform cluster on Microsoft Azure, the installation program did not mark the storage account that it created as private. As a result, the storage account was publically available. With this release, the installation program now correctly always marks the storage account as private regardless if the cluster is publically or privately available. (OCPBUGS-42349)

1.9.49.3. Updating

To update an OpenShift Container Platform 4.17 cluster to this latest release, see Updating a cluster using the CLI.

1.9.50. RHSA-2024:7922 - OpenShift Container Platform 4.17.1 bug fix and security update advisory

Issued: 16 October 2024

OpenShift Container Platform release 4.17.1, which includes security updates, is now available. The list of bug fixes that are included in the update is documented in the RHSA-2024:7922 advisory. The RPM packages that are included in the update are provided by the RHSA-2024:7925 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 4.17.1 --pullspecs
```

1.9.50.1. Enhancements

- The Operator SDK now correctly scaffolds the Dockerfile for **ansible-operator**. Additionally, the Operator SDK can now use **-rhel9** container images. (OCPBUGS-42853)
- The Operator SDK now correctly scaffolds the Dockerfile for **helm-operator**. Additionally, the Operator SDK can now use **-rhel9** container images. (OCPBUGS-42786)
- When you install the Red Hat OpenShift Lightspeed, which is a Developer Preview feature only, the checkbox for enabling in-cluster monitoring is now enabled by default. (OCPBUGS-42380)
- The installation program now uses a newer version of the **cluster-api-provider-ibmcloud** provider that includes Transit Gateway fixes. (OCPBUGS-42483)
- The migrating CNS volumes feature, which is a Developer Preview feature only, includes the following enhancements:
  - The feature now checks for the vCenter version before moving CNS volumes to VMware vSphere. (OCPBUGS-42006)
  - The feature keeps migrating CNS volumes even if some volumes do not exist, so that the migration operation does not exit when no volume exists. (OCPBUGS-42008)
- The vSphere CSI Driver Operator can now delete all resources for a vSphere CSI driver that was removed from a cluster. (OCPBUGS-42007)

1.9.50.2. Bug fixes

- Previously, the Single-Root I/O Virtualization (SR-IOV) Operator did not expire the acquired lease during the Operator's shutdown operation. This impacted a new instance of the Operator, because the new instance had to wait for the lease to expire before the new instance was operational. With this release, an update to the Operator shutdown logic ensures that the Operator expires the lease when the Operator is shutting down. (OCPBUGS-37668)
- Previously, when you configured the image registry to use an Microsoft Azure storage account that was located in a resource group other than the cluster's resource group, the Image Registry Operator would become degraded. This occurred because of a validation error. With this release, an update to the Operator allows for authentication only by using a storage account key. Validation of other authentication requirements is not required. (OCPBUGS-42812)
- Previously, if availability zones were not in a specific order in the **install-config.yaml** configuration file, the installation program would wrongly sort the zones before saving the control plane machine set manifests. When the program created the machines, additional control plane virtual machines were created to reconcile the machines into each zone. This caused a resource-constraint issue. With this release, the installation program no longer sorts availability zones so that this issue no longer occurs. (OCPBUGS-42699)
- Previously, the Red Hat OpenShift Container Platform web console did not require the **Creator** field as a mandatory field. API changes specified an empty value for this field, but a user profile could still create silent alerts. With this release, the API marks the **Creator** field as a mandatory field for a user profile. (OCPBUGS-42606)
- Previously, during root certification rotation, the Ingress Operator and DNS Operator failed to start. With this release, an update to the kubeconfigs for the Ingress Operator and DNS Operator ensure that annotations set the conditions for managing the public key infrastructure

(PKI). This update ensures that both Operators can start as expected during root certification rotation. (OCPBUGS-42261)

- Previously, during root certification rotation, the **metrics-server** pod in the data plane failed to start correctly. This happened because of a certificate issue. With this release, the **hostedClusterConfigOperator** resource sends the correct certificate to the data plane so that the **metrics-server** pod starts as expected. (OCPBUGS-42098)

- Previously, the installation program attempted to create a private zone for a cluster that needed to be installed on a Google Cloud shared virtual private network (VPC). This caused the installation of the cluster to fail. With this release, a fix skips the creation of the private zone so that this cluster installation issue no longer exists. (OCPBUGS-42142)

- Previously, when the installation program installed a cluster on a Google Cloud VPC, role bindings for the control plane service account were not removed by the program. With this release, a fix removes the role bindings so that your cluster no longer includes these artifacts. (OCPBUGS-42116)

- Previously, if you enabled on-cluster layering for your cluster and you attempted to configure kernel arguments in the machine configuration, machine config pools (MCPs) and nodes entered a degraded state. This happened because of a configuration mismatch. With this release, a check for kernel arguments for a cluster with OCL-enabled ensures that the arguments are configured and applied to nodes in the cluster. This update prevents any mismatch that previously occurred between the machine configuration and the node configuration. (OCPBUGS-42081)

- Previously, creating cron jobs to create pods for your cluster caused the component that fetches the pods to fail. Because of this issue, the **Topology** page on the OpenShift Container Platform web console failed. With this release, a **3** second delay is configured for the component that fetches pods that are generated from the cron job so that this issue no longer exists. (OCPBUGS-41685)

- Previously, when you deployed an OpenShift Container Platform cluster that listed a **TechPreviewNoUpgrade** feature gate in its configured on RHOSP, the **cluster-capi-operator** pod crashed. This occurred because the Cluster CAPI Operator expected a different API version than the one that was served. With this release, an update to the Cluster CAPI Operator ensures that the Operator uses the correct version of the API so that this issue no longer occurs. (OCPBUGS-41576)

- Previously, using DNF to install additional packages on your customized Red Hat Enterprise Linux CoreOS (RHCOS) builds caused builds to fail because packages could not be located. With this release, the Subscription Manager adds the correct packages to RHCOS so that this issue no longer occurs. (OCPBUGS-41376)

- Previosuly, bonds that were configured in **active-backup** mode would have EFI System Partition (ESP) offload active even if underlying links did not support ESP offload. This caused IPsec associations to fail. With this release, ESP offload is disabled for bonds so that IPsec associations pass. (OCPBUGS-41255)

- Previosuly, resources for a new user account were not removed when the account was deleted. This caused unnecessary information in config maps, roles, and role-bindings. With this release, an **ownerRef** tag is added to these resources, so that when you delete a user account the resources are also deleted from all cluster resources. (OCPBUGS-39601)

- Previosuly, a coding issue caused the Ansible script on RHCOS user-provisioned installation infrastructure to fail. This occurred when IPv6 was enabled for a three-node cluster. With this

release, support exists for installing a three-node cluster with IPv6 enabled on RHCOS. (OCPBUGS-39409)

- Previously, the order of an Ansible Playbook was modified to run before the **metadata.json** file was created, which caused issues with older versions of Ansible. With this release, the playbook is more tolerant of missing files and the issue is resolved. (OCPBUGS-39286)

- Previously, when the Node Tuning Operator (NTO) was configured to use **PerformanceProfiles**, NTO would create an **ocp-tuned-one-shot systemd** service. The **systemd** service would run before kubelet and blocked execution. The **systemd** service invokes Podman which uses an NTO image, but when the NTO image was not present Podman still tried to fetch the image and it would fail. With this release, support is added for cluster-wide proxy environment variables defined in **/etc/mco/proxy.env**. Now, Podman pulls NTO images in environments which need to use proxies for out-of-cluster connections. (OCPBUGS-39124)

- Previously, the resource type filter for the **Events** page on the OpenShift Container Platform web console wrongly reported the number of resources when you selected more than three resources. With this release, the filter now reports the correct number of resources based on your resource selection. (OCPBUGS-39091)

- Previously, Ironic inspection failed if special or invalid characters existed in the serial number of a block device. This occurred because the **lsblk** command failed to escape the characters. With this release, the command now escapes the characters so this issue no longer persists. (OCPBUGS-39013)

- Previously, when you used the Redfish Virtual Media to add an xFusion bare-metal node to your cluster, the node did not get added because of a node registration issue. The issue occurred because the hardware was not 100% compliant with Redfish. With this release, you can now add xFusion bare-metal nodes to your cluster. (OCPBUGS-38784)

- Previously, on the **Developer** perspective on the OpenShift Container Platform web console, when you navigated to **Observe** > **Metrics**, two **Metrics** tabs existed. With this release, the duplicate tab is removed and now exists in the **openshift-monitoring/monitoring-plugin** application that serves the **Metrics** tabs on the web console. (OCPBUGS-38462)

- Previously, the **manila-csi-driver** and node registrar pods had missing healthchecks because of a configuration issue. With this release, the health checks are now added to both of these resources. (OCPBUGS-38457)

- Previously, updating the **additionalTrustBundle** parameter in a hosted control planes cluster configuration did not get applied to compute nodes. With this release, a fix ensures that updates to the **additionalTrustBundle** parameter automatically apply to compute nodes that exist in your hosted control planes cluster. (OCPBUGS-36680)

- Previously, with oc-mirror plugin v2 (Technology Preview), images that referenced both **tag** and **digest** references were not supported. During the disk-to-mirror process for fully disconnected environments, the images were skipped and this caused build issues for the archive file. With this release, oc-mirror plugin v2 now supports an image that includes both of these references. An image is now pulled from the **digest** reference and keeps the **tag** reference for information purposes, while an appropriate warning message is displayed in the console output. (OCPBUGS-42421)

- Previously, the Cluster API used an unsupported tag template for a virtual network installation when compared with the default non-cluster-API-provisioned installation. This caused the Image Registry Operator to enter a degraded state when configured with **networkAccess:**

Internal. With this release, the Image Registry Operator now supports both tag templates so that this issue no longer exists. (OCPBUGS-42394)

- Previously, the Cloud Controller Manager (CCM) liveness probe used on IBM Cloud cluster installations could not use loopback and this caused the probe to continuously restart. With this release, the probe can use loopback so that this issue not longer occurs. (OCPBUGS-41941)

- Previously, when the **globallyDisableIrqLoadBalancing** field was set to **true** in the **PerformanceProfile** object, the isolated CPUs were listed in the **IRQBALANCE_BANNED_CPULIST** variable instead of the **IRQBALANCE_BANNED_CPUS** variable. These variables are stored in **/etc/sysconfig/irqbalance**. Changing the value of the **globallyDisableIrqLoadBalancing** field from **true** to **false** did not update the **IRQBALANCE_BANNED_CPULIST** variable correctly. As a result, the number of CPUs available for load rebalancing did not increase because the isolated CPUs remained in the **IRQBALANCE_BANNED_CPULIST** variable. With this release, a fix ensures that isolated CPUs are now listed in the **IRQBALANCE_BANNED_CPUS** variable, so that the number of CPUs available for load rebalancing increase as expected. (OCPBUGS-42323)

1.9.50.3. Updating

To update an OpenShift Container Platform 4.16 cluster to this latest release, see Updating a cluster using the CLI.

1.9.51. RHSA-2024:3718 - OpenShift Container Platform 4.17.0 image release, bug fix and security update advisory

Issued: 1 October 2024

OpenShift Container Platform release 4.17.0, which includes security updates, is now available. The list of bug fixes that are included in the update is documented in the RHSA-2024:3718 advisory. The RPM packages that are included in the update are provided by the RHSA-2024:3722 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 4.17.0 --pullspecs
```

1.9.51.1. Known issues

- When the **globallyDisableIrqLoadBalancing** field is set to **true** in the **PerformanceProfile** object, the isolated CPUs are listed in the **IRQBALANCE_BANNED_CPULIST** variable instead of the **IRQBALANCE_BANNED_CPUS** variable. However, changing the value of the **globallyDisableIrqLoadBalancing** field from **true** to **false** does not update the **IRQBALANCE_BANNED_CPULIST** variable correctly. As a result, the number of CPUs available for load rebalancing does not increase, as the isolated CPUs remain in the **IRQBALANCE_BANNED_CPULIST** variable.

NOTE

The **IRQBALANCE_BANNED_CPULIST** variable and the **IRQBALANCE_BANNED_CPUS** variable are stored in the **/etc/sysconfig/irqbalance** file.

(OCPBUGS-42323)

1.9.51.2. Updating

To update an OpenShift Container Platform 4.16 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 4.17 release notes are available in the following documentation.

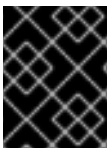

IMPORTANT

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

F

File Integrity Operator

H

Hosted control planes

K

Kube Descheduler Operator

L

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