OpenShift Container Platform 4.18

Release notes

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

Last Updated: 2026-01-19

OpenShift Container Platform 4.18 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 4.18 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-2024:6122) is now available. This release uses Kubernetes 1.31 with CRI-O runtime. New features, changes, and known issues that pertain to OpenShift Container Platform 4.18 are included in this topic.

OpenShift Container Platform 4.18 clusters are available at https://console.redhat.com/openshift. From the Red Hat Hybrid Cloud Console, you can deploy OpenShift Container Platform clusters to either on-premises or cloud environments.

OpenShift Container Platform 4.18 is supported on Red Hat Enterprise Linux (RHEL) 8.8 and a later version of RHEL 8 that is released before End of Life of OpenShift Container Platform 4.18. OpenShift Container Platform 4.18 is also supported on Red Hat Enterprise Linux CoreOS (RHCOS) 4.18. 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.

Starting from OpenShift Container Platform 4.14, the Extended Update Support (EUS) phase for even-numbered releases increases the total available lifecycle to 24 months on all supported architectures, including **x86_64**, 64-bit ARM (**aarch64**), IBM Power® (**ppc64le**), and IBM Z® (**s390x**) architectures. Beyond this, Red Hat also offers a 12-month additional EUS add-on, denoted as *Additional EUS Term 2*, that extends the total available lifecycle from 24 months to 36 months. The Additional EUS Term 2 is available on all architecture variants of OpenShift Container Platform. 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. Authentication and authorization

1.3.1.1. Rotating OIDC bound service account signer keys

With this release, you can use the Cloud Credential Operator (CCO) utility (**ccoctl**) to rotate the OpenID Connect (OIDC) bound service account signer key for clusters installed on the following cloud providers:

- Amazon Web Services (AWS) with Security Token Service (STS)
- Google Cloud with GCP Workload Identity
- Microsoft Azure with Workload ID

1.3.2. Backup and restore

1.3.2.1. Hibernating a cluster for up to 90 days

With this release, you can now hibernate your OpenShift Container Platform cluster for up to 90 days and expect the cluster to recover successfully. Before this release, you could only hibernate for up to 30 days.

For more information, see Hibernating an OpenShift Container Platform cluster .

1.3.2.2. Enhanced etcd backup and restore documentation

The etcd disaster recovery documentation was updated and simplified for quicker recovery of the cluster, both in a normal disaster recovery situation and in cases where a full cluster restoration from a previous backup is not necessary.

Two scripts, **quorum-restore.sh** and **cluster-restore.sh**, are introduced to complete many of the steps in the recovery procedure.

In addition, a procedure was added to more quickly recover the cluster when at least one good node exists. If any of the surviving nodes meets specific criteria, you can use it to run the recovery.

For more information, see About disaster recovery.

1.3.3. Edge computing

1.3.3.1. Shutting down and restarting single-node OpenShift clusters up to 1 year after cluster installation

With this release, you can shut down and restart single-node OpenShift clusters up to 1 year after cluster installation. If certificates expired while the cluster was shut down, you must approve certificate signing requests (CSRs) upon restarting the cluster.

Before this update, you could shut down and restart single-node OpenShift clusters for only 120 days after cluster installation.

> **IMPORTANT**
>
> Evacuate all workload pods from the single-node OpenShift cluster before you shut it down.

For more information, see Shutting down the cluster gracefully.

1.3.4. Extensions (OLM v1)

1.3.4.1. Operator Lifecycle Manager (OLM) v1 (General Availability)

Operator Lifecycle Manager (OLM) has been included with OpenShift Container Platform 4 since its initial release and has helped enable and grow a substantial ecosystem of solutions and advanced workloads running as Operators.

OpenShift Container Platform 4.18 introduces *OLM v1*, the next-generation Operator Lifecycle Manager, as a General Availability (GA) feature, designed to improve how you manage Operators on OpenShift Container Platform.

With OLM v1 now generally available, starting in OpenShift Container Platform 4.18, the existing version of OLM that has been included since the launch of OpenShift Container Platform 4 is now known as *OLM (Classic)*.

Previously available as a Technology Preview feature only, the updated framework in OLM v1 evolves many of the concepts that have been part of OLM (Classic) by simplifying Operator management, enhancing security, and boosting reliability.

> **IMPORTANT**
>
> - Starting in OpenShift Container Platform 4.18, OLM v1 is now enabled by default, alongside OLM (Classic). OLM v1 is a cluster capability that administrators can optionally disable before installation of OpenShift Container Platform.
> - OLM (Classic) remains fully supported throughout the OpenShift Container Platform 4 lifecycle.

Simplified API

OLM v1 simplifies Operator management with a new, user-friendly API: the **ClusterExtension** object. By managing Operators as integral extensions of the cluster, OLM v1 caters to the special lifecycle requirements of custom resource definition (CRDs). This design aligns more closely with Kubernetes

principles, treating Operators, which consist of custom controllers and CRDs, as cluster-wide singletons.
OpenShift Container Platform continues to give you access to the latest Operator packages, patches, and updates through default Red Hat Operator catalogs, which are enabled by default for OLM v1 in OpenShift Container Platform 4.18. With OLM v1, you can install an Operator package by creating and applying a **ClusterExtension** API object in your cluster. By interacting with **ClusterExtension** objects, you can manage the lifecycle of Operator packages, quickly understand their status, and troubleshoot issues.

Streamlined declarative workflows

Leveraging the simplified API, you can define your desired Operator states in a declarative way and, when integrating with tools like Git and Zero Touch Provisioning, let OLM v1 automatically maintain those states. This minimizes human error and unlocks a wider range of use cases.

Uninterrupted operations with continuous reconciliation and optional rollbacks

OLM v1 enhances reliability through continuous reconciliation. Rather than relying on single attempts, OLM v1 proactively addresses Operator installation and update failures, automatically retrying until the issue is resolved. This eliminates the manual steps previously required, such as deleting **InstallPlan** API objects, and greatly simplifies the resolution of off-cluster issues, such as missing container images or catalog problems.
In addition, OLM v1 provides optional rollbacks, allowing you to revert Operator version updates under specific conditions after carefully assessing any potential risks.

Granular update control for deployments

With granular update control, you can select a specific Operator version or define a range of acceptable versions. For example, if you have tested and approved version **1.2.3** of an Operator in a stage environment, instead of hoping the latest version works just as well in production, you can use version pinning. By specifying **1.2.3** as the desired version, you can ensure that is the exact version that will be deployed for a safe and predictable update.
Alternatively, automatic z-stream updates provide a seamless and secure experience by automatically applying security fixes without manual intervention, minimizing operational disruptions.

Enhanced security with user-provided service accounts

OLM v1 prioritizes security by minimizing its permission requirements and providing greater control over access. By using user-provided **ServiceAccount** objects for Operator lifecycle operations, OLM v1 access is restricted to only the necessary permissions, significantly reducing the control plane attack surface and improving overall security. In this way, OLM v1 adopts a least-privilege model to minimize the impact of a compromise.

NOTE

The documentation for OLM v1 exists as a stand-alone guide called Extensions. Previously, OLM v1 documentation was a subsection of the Operators guide, which otherwise documents the OLM (Classic) feature set.

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

1.3.4.2. OLM v1 supported extensions

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 OLM (Classic).
- 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.

1.3.4.3. Disconnected environment support in OLM v1

To support cluster administrators that prioritize high security by running their clusters in internet-disconnected environments, especially for mission-critical production workloads, OLM v1 supports these disconnected environments, starting in OpenShift Container Platform 4.18.

After using the oc-mirror plugin for the OpenShift CLI (**oc**) to mirror the images required for your cluster to a mirror registry in your fully or partially disconnected environments, OLM v1 can function properly in these environments by utilizing the sets of resources generated by either oc-mirror plugin v1 or v2.

For more information, see Disconnected environment support in OLM v1 .

1.3.4.4. Improved catalog selection in OLM v1

With this release, you can perform the following actions to control the selection of catalog content when you install or update a cluster extension:

- Specify labels to select the catalog
- Use match expressions to filter across catalogs
- Set catalog priority

For more information, see Catalog content resolution .

1.3.4.5. Basic support for proxied environments and trusted CA certificates

With this release, Operator Controller and catalogd can now run in proxied environments and include basic support for trusted CA certificates.

1.3.4.6. Compatibility with OpenShift Container Platform versions

Before cluster administrators can update their OpenShift Container Platform cluster to its next minor version, they must ensure that all installed Operators are updated to a bundle version that is compatible with the next minor version (4.y+1) of a cluster.

Starting in OpenShift Container Platform 4.18, OLM v1 supports the **olm.maxOpenShiftVersion** annotation in the cluster service version (CSV) of an Operator, similar to the behavior in OLM (Classic), to prevent administrators from updating the cluster before updating the installed Operator to a compatible version.

For more information, see Compatibility with OpenShift Container Platform versions .

1.3.4.7. User access to extension resources

After a cluster extension has been installed and is being managed by Operator Lifecycle Manager (OLM) v1, the extension can often provide **CustomResourceDefinition** objects (CRDs) that expose new API resources on the cluster. Cluster administrators typically have full management access to these resources by default, whereas non-cluster administrator users, or *regular users*, might lack sufficient permissions.

OLM v1 does not automatically configure or manage role-based access control (RBAC) for regular users to interact with the APIs provided by installed extensions. Cluster administrators must define the required RBAC policy to create, view, or edit these custom resources (CRs) for such users.

For more information, see User access to extension resources .

1.3.4.8. Runtime validation of container images using sigstore signatures in OLM v1 (Technology Preview)

Starting in OpenShift Container Platform 4.18, OLM v1 support for handling runtime validation of sigstore signatures for container images is available as a Technology Preview (TP) feature.

1.3.4.9. OLM v1 known issues

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

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.

1.3.4.10. Deprecation of SiteConfig v1

SiteConfig v1 is deprecated starting with OpenShift Container Platform 4.18. Equivalent and improved functionality is now available through the SiteConfig Operator using the **ClusterInstance** custom resource. For more information, see the Red Hat Knowledge Base solution Procedure to transition from SiteConfig CRs to the ClusterInstance API.

For more information about the SiteConfig Operator, see SiteConfig.

1.3.5. Hosted control planes

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

1.3.6. IBM Power

The IBM Power® release on OpenShift Container Platform 4.18 adds improvements and new capabilities to OpenShift Container Platform components.

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

- Added four new data centers to PowerVS Installer Provisioned Infrastructure deployments
- Adding compute nodes to on-premise clusters using OpenShift CLI (**oc**)

1.3.7. IBM Z and IBM LinuxONE

With this release, IBM Z® and IBM® LinuxONE are now compatible with OpenShift Container Platform 4.18. You can perform the installation with z/VM, LPAR, or Red Hat Enterprise Linux (RHEL) Kernel-based Virtual Machine (KVM). For installation instructions, see Installation methods.

IMPORTANT

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

1.3.7.1. IBM Z and IBM LinuxONE notable enhancements

The IBM Z® and IBM® LinuxONE release on OpenShift Container Platform 4.18 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:

- Adding compute nodes to on-premise clusters using OpenShift CLI (**oc**)

1.3.8. 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
Adding compute nodes to on-premise clusters using OpenShift CLI (**oc**)	Supported	Supported
Alternate authentication providers	Supported	Supported
Agent-based Installer	Supported	Supported
Assisted Installer	Supported	Supported
Automatic Device Discovery with Local Storage Operator	Unsupported	Supported

Feature	IBM Power®	IBM Z® and IBM® LinuxONE
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
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

Feature	IBM Power®	IBM Z® and IBM® LinuxONE
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	Supported
OVN-Kubernetes, including IPsec encryption	Supported	Supported
PodDisruptionBudget	Supported	Supported
Precision Time Protocol (PTP) hardware	Unsupported	Unsupported
Red Hat OpenShift Local	Unsupported	Unsupported
Scheduler profiles	Supported	Supported
Secure Boot	Unsupported	Supported
Stream Control Transmission Protocol (SCTP)	Supported	Supported
Support for multiple network interfaces	Supported	Supported
The **openshift-install** utility to support various SMT levels on IBM Power® (Hardware Acceleration)	Supported	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	Supported	Supported
IBM Power® Virtual Server Block CSI Driver Operator	Supported	Unsupported
Ingress Node Firewall Operator	Supported	Supported
Local Storage Operator	Supported	Supported

Feature	IBM Power®	IBM Z® and IBM® LinuxONE
MetalLB Operator	Supported	Supported
Network Observability Operator	Supported	Supported
NFD Operator	Supported	Supported
NMState Operator	Supported	Supported
OpenShift Elasticsearch Operator	Supported	Supported
Vertical Pod Autoscaler Operator	Supported	Supported

Table 1.4. 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

1.3.9. Insights Operator

1.3.9.1. Insights Runtime Extractor (Technology Preview)

In this release, the Insights Operator introduces the workload data collection *Insights Runtime Extractor* feature to help Red Hat better understand the workload of your containers. Available as a Technology Preview, the Insights Runtime Extractor feature gathers runtime workload data and sends it to Red Hat.

Red Hat uses the collected runtime workload data to gain insights that can help you make investment decisions that will drive and optimize how you use your OpenShift Container Platform containers. For more information, see Enabling features using feature gates.

1.3.9.2. Rapid Recommendations

In this release, enhancements have been made to the Rapid Recommendations mechanism for remotely configuring the rules that determine the data that the Insights Operator collects.

The Rapid Recommendations feature is version-independent, and 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 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
>
> In earlier iterations, 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.

1.3.9.3. 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:

- Collects resources from the **nmstate.io/v1** API group.
- Collects data from **clusterrole.rbac.authorization.k8s.io/v1** instances.

1.3.10. Installation and update

1.3.10.1. New version of the Cluster API Provider IBM Cloud

The installation program now uses a newer version of the Cluster API Provider IBM Cloud provider that includes Transit Gateway fixes. Because of the cost of Transit Gateways in IBM Cloud, you can now use the OpenShift Container Platform to create a Transit Gateway when creating an OpenShift Container Platform cluster. For more information, see (OCPBUGS-37588) and (OCPBUGS-41938).

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

With this release, you can install a cluster on Azure using encrypted virtual networks. You are required to use Azure virtual machines that have the **premiumIO** parameter set to **true**, and that do not support NVMe storage. See Microsoft's documentation about Creating a virtual network with encryption and

Requirements and Limitations for more information.

1.3.10.3. Configuring the ovn-kubernetes join subnet during cluster installation

With this release, you can configure the IPv4 join subnet that is used internally by **ovn-kubernetes** when installing a cluster. You can set the **internalJoinSubnet** parameter in the **install-config.yaml** file and deploy the cluster into an existing Virtual Private Cloud (VPC).

For more information, see Network configuration parameters.

1.3.10.4. Introducing the oc adm upgrade recommend command (Technology Preview)

When updating your cluster, the **oc adm upgrade** command returns a list of the next available versions. As long as you are using 4.18 **oc** client binary, you can use the **oc adm upgrade recommend** command to narrow down your suggestions and recommend a new target release before you launch your update. This feature is available for OpenShift Container Platform version 4.16 and newer clusters that are connected to an update service.

For more information, see Updating a cluster by using the CLI

Feature	4.16	4.17	4.18
oc adm upgrade status	Technology Preview	Technology Preview	Technology Preview
oc adm upgrade recommend	Not Available	Not Available	Technology Preview

1.3.10.5. Support for Nutanix Cloud Clusters (NC2) on Amazon Web Services (AWS) and NC2 on Microsoft Azure

With this release, you can install OpenShift Container Platform on Nutanix Cloud Clusters (NC2) on AWS or NC2 on Azure.

For more information, see Infrastructure requirements.

1.3.10.6. Installing a cluster on Google Cloud using the C4 and C4A machine series

With this release, you can deploy a cluster on Google Cloud using the C4 and C4A machine series for compute or control plane machines. The supported disk type of these machines is **hyperdisk-balanced**. If you use an instance type that requires Hyperdisk storage, all of the nodes in your cluster must support Hyperdisk storage, and you must change the default storage class to use Hyperdisk storage.

For more information about configuring machine types, see Installation configuration parameters for GCP, C4 machine series (Compute Engine docs), and C4A machine series (Compute Engine docs).

1.3.10.7. Provide your own private hosted zone when installing a cluster on Google Cloud

With this release, you can provide your own private hosted zone when installing a cluster on Google Cloud into a shared VPC. If you do, the requirements for the bring your own (BYO) zone are that the zone must use a DNS name such as **<cluster_name>.<base_domain>.** and that you bind the zone to the VPC network of the cluster.

For more information, see Prerequisites for installing a cluster on GCP into a shared VPC and Prerequisites for installing a cluster into a shared VPC on GCP using Deployment Manager templates .

1.3.10.8. Installing a cluster on Nutanix by using a preloaded RHCOS image object

With this release, you can install a cluster on Nutanix by using the named, preloaded RHCOS image object from the private cloud or the public cloud. Rather than creating and uploading a RHCOS image object for each OpenShift Container Platform cluster, you can use the **preloadedOSImageName** parameter in the **install-config.yaml** file.

For more information, see Additional Nutanix configuration parameters.

1.3.10.9. Single-stack IPv6 clusters on RHOSP

You can now deploy single-stack IPv6 clusters on RHOSP.

You must configure RHOSP prior to deploying your OpenShift Container Platform cluster. For more information, see Configuring a cluster with single-stack IPv6 networking .

1.3.10.10. Installing a cluster on Nutanix with multiple subnets

With this release, you can install a Nutanix cluster with more than one subnet for the Prism Element into which you are deploying an OpenShift Container Platform cluster.

For more information, see Configuring failure domains and Additional Nutanix configuration parameters.

For an existing Nutanix cluster, you can add multiple subnets by using compute or control plane machine sets.

1.3.10.11. Installing a cluster on VMware vSphere with multiple network interface controllers (Technology Preview)

With this release, you can install a VMware vSphere cluster with multiple network interface controllers (NICs) for a node.

For more information, see Configuring multiple NICs.

For an existing vSphere cluster, you can add multiple subnets by using compute machine sets.

1.3.10.12. Configuring 4 and 5 node control planes with the Agent-based Installer

With this release, if you are using the Agent-based Installer, you can now configure your cluster to be installed with either 4 or 5 nodes in the control plane. This feature is enabled by setting the **controlPlane.replicas** parameter to either **4** or **5** in the **install-config.yaml** file.

For more information, see Optional configuration parameters for the Agent-based Installer.

1.3.10.13. Minimal ISO image support for the Agent-based Installer

With this release, the Agent-based Installer supports creating a minimal ISO image on all supported platforms. Previously, minimal ISO images were supported only on the **external** platform.

This feature is enabled using the **minimalISO** parameter in the **agent-config.yaml** file.

For more information, see Optional configuration parameters for the Agent-based Installer.

1.3.10.14. Internet Small Computer System Interface (iSCSI) boot support for the Agent-based Installer

With this release, the Agent-based Installer supports creating assets that can be used to boot an OpenShift Container Platform cluster from an iSCSI target.

For more information, see Preparing installation assets for iSCSI booting .

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

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

NOTE

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

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

1.3.11. Machine Config Operator

1.3.11.1. Updated boot images for AWS clusters promoted to GA

Updated boot images has been promoted to GA for Amazon Web Services (AWS) clusters. For more information, see Updated boot images.

1.3.11.2. Expanded machine config nodes information (Technology Preview)

The machine config nodes custom resource, which you can use to monitor the progress of machine configuration updates to nodes, now presents more information on the update. The output of the **oc get machineconfignodes** command now reports on the following and other conditions. You can use these statuses to follow the update, or troubleshoot the node if it experiences an error during the update:

- If each node was cordoned and uncordoned
- If each node was drained
- If each node was rebooted
- If a node had a CRI-O reload
- If a node had the operating system and node files updated

1.3.11.3. On-cluster layering changes (Technology Preview)

There are several important changes to the on-cluster layering feature:

- You can now install extensions onto an on-cluster customer layered image by using a **MachineConfig** object.
- Updating the Containerfile in a **MachineOSConfig** object now triggers a build to be performed.
- You can now revert an on-cluster custom layered image back to the base image by removing a label from the **MachineOSConfig** object.
- The **must-gather** for the Machine Config Operator now includes data on the **MachineOSConfig** and **MachineOSBuild** objects.

For more information about on-cluster layering, see Using on-cluster layering to apply a custom layered image.

1.3.12. Machine management

1.3.12.1. Managing machines with the Cluster API for Microsoft Azure (Technology Preview)

This release introduces the ability to manage machines by using the upstream Cluster API, integrated into OpenShift Container Platform, as a Technology Preview for Microsoft Azure clusters. This capability is in addition or an alternative to managing machines with the Machine API. For more information, see About the Cluster API.

1.3.13. Management console

1.3.13.1. Checkbox for enabling cluster monitoring is marked by default

With this update, the checkbox for enabling cluster monitoring is now checked by default when installing the OpenShift Lightspeed Operator. (OCPBUGS-42381)

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:

- Metrics Server to 0.7.2
- Prometheus to 2.55.1
- Prometheus Operator to 0.78.1
- Thanos to 0.36.1

1.3.14.2. Added scrape and evaluation intervals for user workload monitoring Prometheus

With this update, you can configure the intervals between consecutive scrapes and between rule evaluations for Prometheus for user workload monitoring.

1.3.14.3. Added early validation for the monitoring configurations in monitoring config maps

This update introduces early validation for changes to monitoring configurations in **cluster-monitoring-config** and **user-workload-monitoring-config** config maps to provide shorter feedback loops and enhance user experience.

1.3.14.4. Added the proxy environment variables to Alertmanager containers

With this update, Alertmanager uses the proxy environment variables. Therefore, if you configured an HTTP cluster-wide proxy, you can enable proxying by setting the **proxy_from_environment** parameter to **true** in your alert receivers or at the global config level in Alertmanager.

1.3.14.5. Added cross-project user workload alerting and recording rules

With this update, you can create user workload alerting and recording rules that query multiple projects at the same time.

1.3.14.6. Correlating cluster metrics with RHOSO metrics

You can now correlate observability metrics for clusters that run on Red Hat OpenStack Services on OpenShift (RHOSO). By collecting metrics from both environments, you can monitor and troubleshoot issues across the infrastructure and application layers.

For more information, see Monitoring clusters that run on RHOSO .

1.3.15. 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 .

1.3.16. Networking

1.3.16.1. Deploying the SR-IOV Network Operator on a cluster that runs on ARM architecture

You can now deploy the SR-IOV Network Operator on a cluster that runs on ARM architecture. (OCPBUGS-56496)

1.3.16.2. Holdover in a grandmaster clock with GNSS as the source

With this release, you can configure the holdover behavior in a grandmaster (T-GM) clock with Global Navigation Satellite System (GNSS) as the source. Holdover allows the T-GM clock to maintain synchronization performance when the GNSS source is unavailable. During this period, the T-GM clock relies on its internal oscillator and holdover parameters to reduce timing disruptions.

You can define the holdover behavior by configuring the following holdover parameters in the **PTPConfig** custom resource (CR):

- **MaxInSpecOffset**
- **LocalHoldoverTimeout**
- **LocalMaxHoldoverOffSet**

For more information, see Holdover in a grandmaster clock with GNSS as the source .

1.3.16.3. Support for configuring a multi-network policy for IPVLAN and Bond CNI

With this release, you can configure a multi-network policy for the following network types:

- IP Virtual Local Area Network (IPVLAN)
- Bond Container Network Interface (CNI) over SR-IOV

For more information, see Configuring multi-network policy

1.3.16.4. Updated terminology for whitelist and blacklist annotations

The terminology for the **ip_whitelist** and **ip_blacklist** annotations have been updated to **ip_allowlist** and **ip_denylist**, respectively. Currently, OpenShift Container Platform still supports the **ip_whitelist** and **ip_blacklist** annotations. However, these annotations are planned for removal in a future release.

1.3.16.5. Checking OVN-Kubernetes network traffic with OVS sampling using the CLI

OVN-Kubernetes network traffic can be viewed with OVS sampling via the CLI for the following network APIs:

- **NetworkPolicy**
- **AdminNetworkPolicy**
- **BaselineNetworkPolicy**
- **UserDefinedNetwork** isolation
- **EgressFirewall**
- Multicast ACLs.

Checking OVN-Kubernetes network traffic with OVS sampling using the CLI is intended to help with packet tracing. It can also be used while the Network Observability Operator is installed.

For more information, see Checking OVN-Kubernetes network traffic with OVS sampling using the CLI .

1.3.16.6. User-defined network segmentation (Generally Available)

With OpenShift Container Platform 4.18, user-defined network segmentation is generally available. User-defined networks (UDN) introduce enhanced network segmentation capabilities by allowing administrators to define custom network topologies using namespace-scoped UserDefinedNetwork and cluster-scoped ClusterUserDefinedNetwork custom resources.

With UDNs, administrators can create tailored network topologies with enhanced isolation, IP address management for workloads, and advanced networking features. Supporting both Layer 2 and Layer 3 topology types, user-defined network segmentation enables a wide range of network architectures and topologies, enhancing network flexibility, security, and performance. For more information on supported features, see UDN support matrix.

Use cases of UDN include providing virtual machines (VMs) with a lifetime duration for static IP addresses assignment as well as a Layer 2 primary pod network so that users can live migrate VMs between nodes. These features are all fully equipped in OpenShift Virtualization. Users can use UDNs to create a stronger, native multi-tenant environment, allowing you to secure your overlay Kubernetes network, which is otherwise open by default. For more information, see About user-defined networks.

1.3.16.7. The dynamic configuration manager is enabled by default (Technology Preview)

You can reduce your memory footprint by using the dynamic configuration manager on Ingress Controllers. The dynamic configuration manager propagates endpoint changes through a dynamic API. This process enables the underlying routers to adapt to changes (scale ups and scale downs) without reloads.

To use the dynamic configuration manager, enable the **TechPreviewNoUpgrade** feature set by running the following command:

```
$ oc patch featuregates cluster -p '{"spec": {"featureSet": "TechPreviewNoUpgrade"}}' --type=merge
```

1.3.16.8. Additional environments for the network flow matrix

With this release, you can view network information for ingress flows to OpenShift Container Platform services in the following environments:

- OpenShift Container Platform on bare metal
- Single-node OpenShift on bare metal
- OpenShift Container Platform on Amazon Web Services (AWS)
- Single-node OpenShift on AWS

For more information, see OpenShift Container Platform network flow matrix .

1.3.16.9. 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.

1.3.16.10. Configuring an RDMA subsytem for SR-IOV

With this release, you can configure a Remote Direct Memory Access (RDMA) Container Network Interface (CNI) on Single Root I/O Virtualization (SR-IOV) to enable high-performance, low-latency communication between containers. When you combine RDMA with SR-IOV, you provide a mechanism to expose hardware counters of Mellanox Ethernet devices to be used inside Data Plane Development Kit (DPDK) applications.

1.3.16.11. Support configuring the SR-IOV Network Operator on a Secure-Boot-enabled environment for Mellanox cards

With this release, you can configure the Single Root I/O Virtualization (SR-IOV) Network Operator when the system has secure boot enabled. The SR-IOV Operator is configured after you first manually configure the firmware for Mellanox devices. With secure boot enabled, the resilience of your system is enhanced, and a crucial layer of defense for the overall security of your computer is provided.

For more information, see Configuring the SR-IOV Network Operator on Mellanox cards when Secure Boot is enabled.

1.3.16.12. Support for pre-created RHOSP floating IP addresses in the Ingress Controller

With this release, you can now specify pre-created floating IP addresses in the Ingress Controller for your clusters running on RHOSP.

For more information, see Specifying a floating IP address in the Ingress Controller .

1.3.16.13. SR-IOV Network Operator support extension

The SR-IOV Network Operator now supports Intel NetSec Accelerator Cards and Marvell Octeon 10 DPUs. (OCPBUGS-43451)

1.3.16.14. Using a Linux bridge interface as the OVS default port connection

The OVN-Kubernetes plugin can now use a Linux bridge interface as the Open vSwitch (OVS) default port connection. This means that a network interface controller, such as SmartNIC, can now bridge the underlying network with a host. (OCPBUGS-39226)

1.3.16.15. Cluster Network Operator exposing network overlap metrics for an issue

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-39096)

1.3.16.16. Network attachments support dynamic reconfiguration

Previously, the **NetworkAttachmentDefinition** CR was immutable. With this release, you can edit an existing **NetworkAttachmentDefinition** CR. Support for editing makes it easier to accommodate changes in the underlying network infrastructure, such as adjusting the MTU of a network interface.

You must ensure that the configurations of each **NetworkAttachmentDefinition** CR that reference the same network **name** and **type: ovn-k8s-cni-overlay** are in sync. Only when these values are in sync is the network attachment update successful. If the configurations are not in sync, the behavior is undefined because there is no guarantee about which **NetworkAttachmentDefinition** CR OpenShift Container Platform uses for the configuration.

You still must restart any workloads that use the network attachment definition for the network changes to take effect for those pods.

1.3.17. Nodes

1.3.17.1. crun is now the default container runtime

crun is now the default container runtime for new containers created in OpenShift Container Platform. The runC runtime is still supported and you can change the default runtime to runC, if needed. For more information on crun, see About the container engine and container runtime. For information on changing the default to runC, see Creating a ContainerRuntimeConfig CR to edit CRI-O parameters .

Updating from OpenShift Container Platform 4.17.z to OpenShift Container Platform 4.18 does not change your container runtime.

1.3.17.2. sigstore support (Technology Preview)

Available as a Technology Preview, you can use the sigstore project with OpenShift Container Platform to improve supply chain security. You can create signature policies at the cluster-wide level or for a specific namespace. For more information, see Manage secure signatures with sigstore .

1.3.17.3. Enhancements to process for adding nodes

Enhancements have been added to the process for adding worker nodes to an on-premise cluster that was introduced in OpenShift Container Platform 4.17. With this release, you can now generate Preboot Execution Environment (PXE) assets instead of an ISO image file, and you can configure reports to be generated regardless of whether the node creation process fails or not.

1.3.17.4. Node Tuning Operator properly selects kernel arguments

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

1.3.17.5. Default container runtime is not always set properly

The default container runtime that is set by the cluster Node Tuning Operator is always inherited from the cluster, and is not hard-coded by the Operator. Starting with this release, the default value is **crun**. (OCPBUGS-45450)

1.3.18. OpenShift CLI (oc)

1.3.18.1. oc-mirror plugin v2 (Generally Available)

oc-mirror plugin v2 is now generally available. To use it, add the **--v2** flag when running oc-mirror commands. The previous version (oc-mirror plugin v1), which runs when the **--v2** flag is not set, is now deprecated. It is recommended to transition to oc-mirror plugin v2 for continued support and improvements.

For more information, see Mirroring images for a disconnected installation by using the oc-mirror plugin v2.

oc-mirror plugin v2 now supports mirroring helm charts. Also, oc-mirror plugin v2 can now be used in environments where **HTTP/S** proxy is enabled, ensuring broader compatibility with enterprise setups.

oc-mirror plugin v2 introduces v1 retro-compatible filtering of Operator catalogs and generates filtered catalogs. This feature allows cluster administrators to view only the Operators that have been mirrored, rather than the complete list from the origin catalog.

1.3.19. Operator lifecycle

1.3.19.1. Existing version of Operator Lifecycle Manager now known as OLM (Classic)

With the release of Operator Lifecycle Manager (OLM) v1 as a General Availability (GA) feature, starting in OpenShift Container Platform 4.18, the existing version of OLM that has been included since the launch of OpenShift Container Platform 4 is now known as *OLM (Classic)*.

> NOTE
>
> OLM (Classic) remains enabled by default and fully supported throughout the OpenShift Container Platform 4 lifecycle.

For more information on the GA release of OLM v1, see the Extensions (OLM v1) release note sections. For full documentation focused on OLM v1, see the stand-alone Extensions guide.

For full documentation focused on OLM (Classic), continue referring to the Operators guide.

1.3.20. Oracle(R) Cloud Infrastructure (OCI)

1.3.20.1. Bare-metal support on Oracle(R) Cloud Infrastructure (OCI)

OpenShift Container Platform cluster installations on Oracle® Cloud Infrastructure (OCI) are now supported for bare-metal machines. You can install bare-metal clusters on OCI by using either the Assisted Installer or the Agent-based Installer. To install a bare-metal cluster on OCI, choose one of the following installation options:

- Installing a cluster on Oracle Cloud Infrastructure (OCI) by using the Assisted Installer
- Installing a cluster on Oracle Cloud Infrastructure (OCI) by using the Agent-based Installer

1.3.21. Postinstallation configuration

1.3.21.1. Migrating the x86 control plane to arm64 architecture on Amazon Web Services

With this release, you can migrate the control plane in your cluster from **x86** to **arm64** architecture on Amazon Web Services (AWS). For more information, see Migrating the x86 control plane to arm64 architecture on Amazon Web Services.

1.3.21.2. Configuring the image stream import mode behavior (Technology Preview)

This feature introduces a new field, **imageStreamImportMode**, in the **image.config.openshift.io/cluster** resource. The **imageStreamImportMode** field controls the import mode behavior of image streams. You can set the **imageStreamImportMode** field to either of the following values:

- **Legacy**
- **PreserveOriginal**

For more information, see Image controller configuration parameters.

You must enable the **TechPreviewNoUpgrade** feature set in the **FeatureGate** custom resource (CR) to enable the **imageStreamImportMode** feature. For more information, see Understanding feature gates.

1.3.22. Red Hat Enterprise Linux CoreOS (RHCOS)

1.3.22.1. RHCOS uses RHEL 9.4

RHCOS uses Red Hat Enterprise Linux (RHEL) 9.4 packages in OpenShift Container Platform 4.18. These packages ensure that your OpenShift Container Platform instances receive the latest fixes, features, enhancements, hardware support, and driver updates.

1.3.23. Registry

1.3.23.1. Read-only registry enhancements

In previous versions of OpenShift Container Platform, storage mounted as read-only returned no specific metrics or information about storage errors. This could result in silent failures of a registry when the storage backend was read-only. With this release, the following alerts have been added to return storage information when the backend is set to read-only:

Alert Name	Message
ImageRegistryStorageReadOnly	The image registry storage is read-only and no images will be committed to storage.
ImageRegistryStorageFull	The image registry storage disk is full and no images will be committed to storage.

1.3.24. Scalability and performance

1.3.24.1. Cluster validation with the cluster-compare plugin

The **cluster-compare** plugin is an OpenShift CLI (**oc**) plugin that compares a cluster configuration with a target configuration. The plugin reports configuration differences while suppressing expected variations by using configurable validation rules and templates.

For example, the plugin can highlight unexpected differences, such as mismatched field values, missing resources, or version discrepancies, while ignoring expected differences, such as optional components or hardware-specific fields. This focused comparison makes it easier to assess cluster compliance with the target configuration.

You can use the **cluster-compare** plugin in development, production, and support scenarios.

For more information about the **cluster-compare** plugin, see Overview of the cluster-compare plugin.

1.3.24.2. Node Tuning Operator: Deferred Tuning Updates

In this release, the Node Tuning Operator introduces support for deferring tuning updates. Administrators can schedule updates to be applied during a maintenance window with this feature.

For more information, see Deferring application of tuning changes.

1.3.24.3. NUMA Resources Operator now uses default SELinux policy

With this release, the NUMA Resources Operator no longer creates a custom SELinux policy to enable the installation of Operator components on a target node. Instead, the Operator uses a built-in container SELinux policy. This change removes the additional node reboot that was previously required when applying a custom SELinux policy during an installation.

IMPORTANT

In clusters with an existing NUMA-aware scheduler configuration, upgrading to OpenShift Container Platform 4.18 might result in an additional reboot for each configured node. For further information about how to manage an upgrade in this scenario and limit disruption, see the Red Hat Knowledgebase article Managing an upgrade to OpenShift Container Platform 4.18 or later for a cluster with an existing NUMA-aware scheduler configuration

1.3.24.4. Node Tuning Operator platform detection

With this release, when you apply a performance profile, the Node Tuning Operator detects the platform and configures kernel arguments and other platform-specific options accordingly. This release adds support for detecting the following platforms:

- AMD64
- AArch64
- Intel 64

1.3.24.5. Support for worker nodes with AMD EPYC Zen 4 CPUs

With this release, you can use the **PerformanceProfile** custom resource (CR) to configure worker nodes on machines equipped with AMD EPYC Zen 4 CPUs, such as Genoa and Bergamo. These CPUs are fully supported when configured with a single NUMA domain (NPS=1).

IMPORTANT

The per pod power management feature is not functional on AMD EPYC Zen 4 CPUs.

1.3.25. Storage

1.3.25.1. Over-provisioning ratio update after LVMCluster custom resource creation

Previously, the **thinPoolConfig.overprovisionRatio** field in the **LVMCluster** custom resource (CR) could be configured only during the creation of the **LVMCluster** CR. With this release, you can now update the **thinPoolConfig.overprovisionRatio** field even after creating the **LVMCluster** CR.

1.3.25.2. Support for configuring metadata size for the thin pool

This feature provides the following new optional fields in the **LVMCluster** custom resource (CR):

- **thinPoolConfig.metadataSizeCalculationPolicy**: Specifies the policy to calculate the metadata size for the underlying volume group. You can set this field to either **Static** or **Host**. By default, this field is set to **Host**.
- **thinPoolConfig.metadataSize**: Specifies the metadata size for the thin pool. You can configure this field only when the **MetadataSizeCalculationPolicy** field is set to **Static**.

For more information, see About the LVMCluster custom resource.

1.3.25.3. Persistent storage using CIFS/SMB CSI Driver Operator is generally available

OpenShift Container Platform is capable of provisioning persistent volumes (PVs) with a Container Storage Interface (CSI) driver for the Common Internet File System (CIFS) dialect/Server Message Block (SMB) protocol. The CIFS/SMB CSI Driver Operator that manages this driver was introduced in OpenShift Container Platform 4.16 with Technology Preview status. In OpenShift Container Platform 4.18, it is now generally available.

For more information, see CIFS/SMB CSI Driver Operator.

1.3.25.4. Secret Store CSI Driver Operator is generally available

The Secrets Store Container Storage Interface (CSI) Driver Operator, **secrets-store.csi.k8s.io**, allows OpenShift Container Platform to mount multiple secrets, keys, and certificates stored in enterprise-grade external secrets stores into pods as an inline ephemeral volume. The Secrets Store CSI Driver Operator communicates with the provider using gRPC to fetch the mount contents from the specified external secrets store. After the volume is attached, the data in it is mounted into the container's file system. The Secrets Store CSI Driver Operator was available in OpenShift Container Platform 4.14 as a Technology Preview feature. OpenShift Container Platform 4.18 introduces this feature as generally available.

For more information about the Secrets Store CSI driver, see Secrets Store CSI Driver Operator.

For information about using the Secrets Store CSI Driver Operator to mount secrets from an external secrets store to a CSI volume, see Providing sensitive data to pods by using an external secrets store .

1.3.25.5. Persistent volume last phase transition time parameter is generally available

OpenShift Container Platform 4.16 introduced a new parameter, **LastPhaseTransitionTime**, which has a timestamp that is updated every time a persistent volume (PV) transitions to a different phase (**pv.Status.Phase**). For OpenShift Container Platform 4.18, this feature is generally available.

For more information about using the persistent volume last phase transition time parameter, see Last phase transition time.

1.3.25.6. Multiple vCenter support for vSphere CSI is generally available

OpenShift Container Platform 4.17 introduced the ability to deploy OpenShift Container Platform across multiple vSphere clusters (vCenters) as a Technology Preview feature. In OpenShift Container Platform 4.18, Multiple vCenter support is now generally available.

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

1.3.25.7. Always honor persistent volume reclaim policy (Technical Preview)

Prior to OpenShift Container Platform 4.18, the persistent volume (PV) reclaim policy was not always applied.

For a bound PV and persistent volume claim (PVC) pair, the ordering of PV-PVC deletion determined whether the PV delete reclaim policy was applied or not. The PV applied the reclaim policy if the PVC was deleted prior to deleting the PV. However, if the PV was deleted prior to deleting the PVC, then the reclaim policy was not applied. As a result of that behavior, the associated storage asset in the external infrastructure was not removed.

With OpenShift Container Platform 4.18, the PV reclaim policy is consistently always applied. This feature has Technical Preview status.

For more information, see Reclaim policy for persistent volumes .

1.3.25.8. Improved ability to easily remove LVs or LVSs for LSO is generally available

For the Local Storage Operator (LSO), OpenShift Container Platform 4.18 improves the ability to remove Local Volumes (LVs) and Local Volume Sets (LVSs) by automatically removing artifacts, thus reducing the number of steps required.

For more information, see Removing a local volume or local volume set .

1.3.25.9. CSI volume group snapshots (Technology Preview)

OpenShift Container Platform 4.18 introduces Container Storage Interface (CSI) volume group snapshots as a Technology Preview feature. This feature needs to be supported by the CSI driver. CSI volume group snapshots use a label selector to group multiple persistent volume claims (PVCs) for snapshotting. A volume group snapshot represents copies from multiple volumes that are taken at the same point-in-time. This can be useful for applications that contain multiple volumes.

OpenShift Data Foundation supports volume group snapshots.

For more information about CSI volume group snapshots, see CSI volume group snapshots .

1.3.25.10. GCP PD CSI driver supports the C3 instance type for bare metal and N4 machine series is generally available

The Google Cloud Platform Persistent Disk (GCP PD) Container Storage Interface (CSI) driver supports the C3 instance type for bare metal and N4 machine series. The C3 instance type and N4 machine series support the hyperdisk-balanced disks.

Additionally, hyperdisk storage pools are supported for large-scale storage. A hyperdisk storage pool is a purchased collection of capacity, throughput, and IOPS, which you can then provision for your applications as needed.

For OpenShift Container Platform 4.18, this feature is generally available.

For more information, see C3 instance type for bare metal and N4 machine series .

1.3.25.11. OpenStack Manila expanding persistent volumes is generally available

In OpenShift Container Platform 4.18, OpenStack Manila supports expanding Container Storage Interface (CSI) persistent volumes (PVs). This feature is generally available.

For more information, see Expanding persistent volumes and CSI drivers supported by OpenShift Container Platform.

1.3.25.12. GCP Filestore supporting Workload Identity is generally available

In OpenShift Container Platform 4.18, 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. For OpenShift Container Platform 4.18, this feature is generally available.

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

1.3.26. Web console

1.3.26.1. Administrator perspective

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

- A new setting for hiding the **Getting started resources** card on the **Overview** page allowing for maximum use of the dashboard.
- A **Start Job** option was added to the CronJob **List** and **Details** pages, so you can start individual CronJobs manually directly in the web console without having to use the **oc** CLI.
- The **Import YAML** button in the masthead is now a **Quick Create** button that you can use for the rapid deployment of workloads by importing from YAML, Git, or using container images.
- You can build your own generative-AI chat bot with a chat bot sample. The generative-AI chat bot sample is deployed with Helm and includes a full CI/CD pipeline. You can also run this sample on your cluster with no CPUs.
- You can import YAML into the console using OpenShift Lightspeed.

1.3.26.1.1. Content Security Policy (CSP)

With this release, the console Content Security Policy (CSP) is deployed in report-only mode. CSP violations will be logged in the browser console, but the associated CSP directives will not be enforced. Dynamic plugin creators can add their own policies.

Additionally, you can report any plugins that break security policies. Administrators have the ability to disable any plugin breaking those policies. CSP violations will be logged in the browser console, but the associated CSP directives will not be enforced. This feature is behind a **feature-gate**, so you will need to manually enable it.

For more information, see Content Security Policy (CSP) and Enabling feature sets using the web console.

1.3.26.2. Developer Perspective

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

- Added a OpenShift Container Platform toolkit, Quarkus tools and JBoss EAP, and a Language Server Protocol Plugin for Visual Studio Code and IntelliJ.
- Previously, when moving from light mode to dark mode in the Monaco editor, the console remained in dark mode. With this update, the Monaco code editor will match the selected theme.

1.4. NOTABLE TECHNICAL CHANGES

1.4.1. Uninstalling the SR-IOV Network Operator changed

From OpenShift Container Platform 4.18, to successfully uninstall the SR-IOV Network Operator, you need to delete the **sriovoperatorconfigs** custom resource and custom resource definition too.

For more information, see Uninstalling the SR-IOV Network Operator.

1.4.2. Changes to the iSCSI initiator name and service

Previously, the **/etc/iscsi/initiatorname.iscsi** file was present by default on RHCOS images. With this release, the **initiatorname.iscsi** file is no longer present by default. Instead, it is created at run time when the **iscsi.service** and subsequent **iscsi-init.service** services start. This service is not enabled by default and might affect any CSI drivers that rely on reading the contents of the **initiatorname.iscsi** file prior to starting the service.

1.4.3. Operator SDK 1.38.0

OpenShift Container Platform 4.18 supports Operator SDK 1.38.0. See Installing the Operator SDK CLI to install or update to this latest version.

Operator SDK 1.38.0 now supports Kubernetes 1.30 and uses Kubebuilder v4.

Metrics endpoints are now secured using native Kubebuilder metrics configuration instead of **kube-rbac-proxy**, which is now removed.

The following support has also been removed from Operator SDK:

- Scaffolding tools for Hybrid Helm-based Operator projects
- Scaffolding tools for Java-based Operator projects

If you have Operator projects that were previously created or maintained with Operator SDK 1.36.1, update your projects to keep compatibility with Operator SDK 1.38.0:

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

1.4.4. 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.

1.4.5. 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 4.18, 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	4.16	4.17	4.18
Bare Metal Event Relay Operator	Deprecated	Removed	Removed

1.5.2. Images deprecated and removed features

Table 1.7. Images deprecated and removed tracker

Feature	4.16	4.17	4.18
Cluster Samples Operator	Deprecated	Deprecated	Deprecated

1.5.3. Installation deprecated and removed features

Table 1.8. Installation deprecated and removed tracker

Feature	4.16	4.17	4.18
--cloud parameter for **oc adm release extract**	Deprecated	Deprecated	Deprecated
CoreDNS wildcard queries for the **cluster.local** domain	Deprecated	Deprecated	Deprecated
compute.platform.openstack.rootVolume.type for RHOSP	Deprecated	Deprecated	Deprecated
controlPlane.platform.openstack.rootVolume.type for RHOSP	Deprecated	Deprecated	Deprecated
ingressVIP and **apiVIP** settings in the **install-config.yaml** file for installer-provisioned infrastructure clusters	Deprecated	Deprecated	Deprecated
Package-based RHEL compute machines	Deprecated	Deprecated	Deprecated
platform.aws.preserveBootstrapIgnition parameter for Amazon Web Services (AWS)	Deprecated	Deprecated	Deprecated

Feature	4.16	4.17	4.18
Installing a cluster on AWS with compute nodes in AWS Outposts	Deprecated	Deprecated	Deprecated

1.5.4. Machine management deprecated and removed features

Table 1.9. Machine management deprecated and removed tracker

Feature	4.16	4.17	4.18
Managing machine with Machine API for Alibaba Cloud	Removed	Removed	Removed
Cloud controller manager for Alibaba Cloud	Removed	Removed	Removed

1.5.5. Monitoring deprecated and removed features

Table 1.10. Monitoring deprecated and removed tracker

Feature	4.16	4.17	4.18
Alertmanager v1 API	Deprecated	Removed	Removed

1.5.6. Networking deprecated and removed features

Table 1.11. Networking deprecated and removed tracker

Feature	4.16	4.17	4.18
OpenShift SDN network plugin	Deprecated	Removed	Removed
iptables	Deprecated	Deprecated	Deprecated

1.5.7. Node deprecated and removed features

Table 1.12. Node deprecated and removed tracker

Feature	4.16	4.17	4.18
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

Feature	4.16	4.17	4.18
cgroup v1	Deprecated	Deprecated	Deprecated

1.5.8. OpenShift CLI (oc) deprecated and removed features

Table 1.13. OpenShift CLI (oc) deprecated and removed tracker

Feature	4.16	4.17	4.18
oc-mirror plugin v1	General Availability	General Availability	Deprecated

1.5.9. Operator lifecycle and development deprecated and removed features

Table 1.14. Operator lifecycle and development deprecated and removed tracker

Feature	4.16	4.17	4.18
Operator SDK	Deprecated	Deprecated	Deprecated
Scaffolding tools for Ansible-based Operator projects	Deprecated	Deprecated	Deprecated
Scaffolding tools for Helm-based Operator projects	Deprecated	Deprecated	Deprecated
Scaffolding tools for Go-based Operator projects	Deprecated	Deprecated	Deprecated
Scaffolding tools for Hybrid Helm-based Operator projects	Deprecated	Deprecated	Removed
Scaffolding tools for Java-based Operator projects	Deprecated	Deprecated	Removed
SQLite database format for Operator catalogs	Deprecated	Deprecated	Deprecated

1.5.10. Storage deprecated and removed features

Table 1.15. Storage deprecated and removed tracker

Feature	4.16	4.17	4.18
Persistent storage using FlexVolume	Deprecated	Deprecated	Deprecated
Shared Resources CSI Driver Operator	Technical Preview	Deprecated	Removed
AliCloud Disk CSI Driver Operator	General Availability	Removed	Removed

1.5.11. Web console deprecated and removed features

Table 1.16. Web console deprecated and removed tracker

Feature	4.16	4.17	4.18
Patternfly 4	Deprecated	Deprecated	Deprecated
React Router 5	Deprecated	Deprecated	Deprecated

1.5.12. Workloads deprecated and removed features

Table 1.17. Workloads deprecated and removed tracker

Feature	4.16	4.17	4.18
DeploymentConfig objects	Deprecated	Deprecated	Deprecated

1.5.13. Deprecated features

1.5.13.1. Kubernetes API deprecation

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

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

1.5.14. Removed features

1.5.14.1. The Shared Resource CSI Driver is removed

The Shared Resource CSI Driver feature was deprecated in OpenShift Container Platform 4.17, and is now removed from OpenShift Container Platform 4.18. This feature is now generally available in Builds for Red Hat OpenShift 1.1. To use this feature, ensure you are using Builds for Red Hat OpenShift 1.1 or later.

1.5.14.2. The selected bundles feature is removed in oc-mirror v2

The selected bundles feature is removed from the oc-mirror v2 Generally Available release. This change prevents issues where specifying the wrong Operator bundle version could break the Operators in a cluster. (OCPBUGS-49419)

1.5.15. Notice of future deprecation

1.5.15.1. Future Kubernetes API removals

The next minor release of OpenShift Container Platform is expected to use Kubernetes 1.32. Kubernetes 1.32 removed a deprecated API.

See the Deprecated API Migration Guide in the upstream Kubernetes documentation for the list of planned Kubernetes API removals.

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

1.6. BUG FIXES

1.6.1. API Server and Authentication

- Previously, API validation did not prevent an authorized client from decreasing the current revision of a static pod operand, such as kube-apiserver, or prevent the operand from progressing concurrently on two nodes. With this release, requests that attempt to do either are now rejected. (OCPBUGS-48502)

- Previously, the oauth-server would crash when configuring an oath identity provider (IDP) with a callback path that contained spaces. With this release, the issue is resolved.(OCPBUGS-44099)

1.6.2. Bare Metal Hardware Provisioning

- 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-49699)

- Previously, in bare-metal configurations where the provisioning network is disabled but the **bootstrapProvisioningIP** field is set, the bare-metal provisioning components might 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 were 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-36869)

- 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-36492)

- Previously, a check for unexpected IP addresses on the provisioning interface during metal3 pod startup was triggered. This issue occurred because of the presence of an IP addresses supplied by DHCP from a previous version of the pod that existed on another node. With this release, a pod startup check now looks only for IP addresses that exist outside the provisioning network subnet, so that a metal3 pod starts immediately, even when if the node has moved to a different node. (OCPBUGS-38507)

- Previously, enabling a provisioning network by editing the cluster-wide **Provisioning** resource was only possible on installer-provisioned infrastructure clusters with platform type **baremetal**. On bare metal, single-node OpenShift, and user-provisioned infrastructure clusters, editing this resource resulted in a validation error. With this release, the excessive validation check has been removed, and enabling a provisioning network is now possible on bare-metal clusters with platform type **none**. As with installer-provisioned infrastructure clusters, users are responsible for making sure that all networking requirements are met for this operation. (OCPBUGS-43371)

1.6.3. Cloud Compute

- Previously, the availability set fault domain count was hardcoded to **2**. This value works in most regions in 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. (OCPBUGS-48659)

- Previously, an updated zone API error message from Google Cloud with increased granularity caused the machine controller to mistakenly mark the machine as valid with a temporary cloud error instead of recognizing it as an invalid machine configuration error. This prevented the invalid machine from transitioning to a **failed** state. With this update, the machine controller handles the new error messages correctly, and machines with an invalid zone or project ID now transition properly to a failed state. (OCPBUGS-47790)

- Previously, the certificate signing request (CSR) approver included certificates from other systems within its calculations for whether it was overwhelmed and should stop approving certificates. In larger clusters, with other subsystems using CSRs, the CSR approver counted unrelated unapproved CSRs towards its total and prevented further approvals. With this release, the CSR approver only includes CSRs that it can approve, by using the **signerName** property as a filter. As a result, the CSR approver only prevents new approvals when there are a large number of unapproved CSRs for the relevant **signerName** values. (OCPBUGS-46425)

- Previously, some cluster autoscaler metrics were not initialized, and therefore were not available. With this release, these metrics are initialized and available. (OCPBUGS-46416)

- Previously, if an informer watch stream missed an event because of a temporary disconnection, the informer might return a special signal type after it reconnected to the network, especially when the informer recognizes that an EndpointSlice object was deleted during the temporary disconnection. The returned signal type indicated that the state of the event has stalled and that the object was deleted. The returned signal type was not accurate and might have caused confusion for a OpenShift Container Platform user. With this release, the Cloud Controller Manager (CCM) handles unexpected signal types so that OpenShift Container Platform users do not receive confusing information from returned types. (OCPBUGS-45972)

- Previously, when the AWS DHCP option set was configured to use a custom domain name that contains a trailing period (**.**), OpenShift Container Platform installation failed. With this release, the logic that extracts the hostname of EC2 instances and turns them into Kubelet node names is updated to trim trailing periods so that the resulting Kubernetes object name is valid. Trailing periods in the DHCP option set no longer cause installation to fail. (OCPBUGS-45889)

- Previously, installation of an AWS cluster failed in certain environments on existing subnets when the **publicIp** parameter for the **MachineSet** object was explicity set to **false**. With this release, a configuration value set for **publicIp** no longer causes issues when the installation program provisions machines for your AWS cluster in certain environment. (OCPBUGS-45130)

- Previously, enabling a provisioning network by editing the cluster-wide **Provisioning** resource was only possible on clusters with platform type **baremetal**, such as ones created by the installer-provisioned infrastructure installation program. On bare metal single-node OpenShift and user-provisioned infrastructure clusters, this would result in a validation error. The excessive validation has been removed, and enabling a provisioning network is now possible on baremetal clusters with platform type **none**. As with installer-provisioned infrastructure, users are responsible for making sure that all networking requirements are met for this operation. (OCPBUGS-43371)

- 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 4.18 from an earlier version, the values defined by the installation program are used. (OCPBUGS-42660)

- Previously, the cluster autoscaler would occasionally leave a node with a **PreferNoSchedule** taint during deletion. With this release, the maximum bulk deletion limit is disabled so that nodes with this taint no longer remain after deletion. (OCPBUGS-42132)

- 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-41936)

- 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-36871)

- Previously, the cloud node manager had permission to update any node object when it needed to update only the node on which it was running. With this release, restrictions have been put in place to prevent the node manager from one node updating the node object of another node. (OCPBUGS-22190)

1.6.4. Cloud Credential Operator

- Previously, the **aws-sdk-go-v2** software development kit (SDK) failed to authenticate an **AssumeRoleWithWebIdentity** API operation on an Amazon Web Services (AWS) Security Token Service (STS) cluster. With this release, **pod-identity-webhook** now includes a default region so that this issue no longer persists. (OCPBUGS-45937)

- Previously, secrets in the cluster were fetched in a single call. When there were a large number of secrets, this caused the API to time out. With this release, the Cloud Credential Operator fetches secrets in batches limited to 100 secrets. This change prevents timeouts when there are large number of secrets in the cluster. (OCPBUGS-39531)

1.6.5. Cluster Resource Override Admission Operator

- Previously, if you specified the **forceSelinuxRelabel** field in a **ClusterResourceOverride** custom resource (CR), and then modified it afterwards, the change would not be reflected in the **clusterresourceoverride-configuration** config map, which is used to apply the SELinux re-labeling workaround feature. With this update, the Cluster Resource Override Operator can track the change to the **forceSelinuxRelabel** feature in order to reconcile the config map object. As a result, the config map object is correctly updated when you change the **ClusterResourceOverride** CR field. (OCPBUGS-48692)

1.6.6. Cluster Version Operator

- Previously, a custom security context constraint (SCC) impacted any pod that was generated by the Cluster Version Operator from receiving a cluster version upgrade. With this release, OpenShift Container Platform now sets a default SCC to each pod, so that any custom SCC created does not impact a pod. (OCPBUGS-46410)

- 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 in the **ClusterVersion Failing** condition message. With this release, the errors that are propogated to the **ClusterVersion Failing** condition message are filtered. (OCPBUGS-15200)

1.6.7. Developer Console

- Previously, if a **PipelineRun** was using a resolver, rerunning that **PipelineRun** resulted in an error. With this fix, a user can rerun **PipelineRun** if it is using resolver. (OCPBUGS-45228)

- Previously, on the if you edited a deployment config in **Form view**, the **ImagePullSecrets** values were duplicated. With this update, editing the form does not add duplicate entries. (OCPBUGS-45227)

- Previously, when you searched on the **OperatorHub** or another catalog, you would experience periods of latency between each key press. With this update, the input on the catalog search bars are debounced. (OCPBUGS-43799)

- Previously, no option existed to close the **Getting started resources** section in the **Administrator** perspective. With this change, user can close the **Getting started resources** section. (OCPBUGS-38860)

- Previously, when cronjobs were created, the creation of pods happens too quickly, causing the component that fetches new pods off the cronjob to fail. With this update, a 3 second delay was added before starting to fetch the pods of the cronjob. (OCPBUGS-37584)

- Previously, resources created when a new user is created were not removed automatically when the user was deleted. This caused clutter on the cluster with configuration maps, roles, and role-bindings. With this update, **ownerRefs** was added to the resources, so they are cleared once the user is deleted and the cluster no longer clutters with users. (OCPBUGS-37560)

- Previously, when importing a Git repository using the serverless import strategy, the environment variables from the **func.yaml** were not automatically loaded into the form. With this update, the environment variables are now loaded upon import. (OCPBUGS-34764)

- Previously, users would erroneously see an option to import a repository using the pipeline build strategy when the devfile import strategy was selected; however, this was not possible. With this update, the pipeline strategy has been removed when the devfile import strategy is selected. (OCPBUGS-32526)

- Previously, when using a custom template, you could not enter multi-line parameters, such as private keys. With this release, you can switch between single-line and multi-line modes so you can fill out template fields with multi-line inputs. (OCPBUGS-23080)

1.6.8. Image Registry

- Previously, you could not install a 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-49693)

- Previously, when the Image Registry Operator was configured with **networkAccess: Internal** in Microsoft Azure, it would not be possible to successfully set **managementState** to **Removed** in the Operator configuration. This occurred because of an authorization error when the Operator tried to delete the storage container. With this update, the Image Registry Operator continues with the deletion of the storage account, which automatically deletes the storage container, resulting in a successful change into the **Removed** state. (OCPBUGS-42732)

- Previously, when configuring the image registry to use an Microsoft Azure storage account located in a resource group other than the cluster's resource group, the Image Registry Operator would become degraded due to a validation error. This update changes the Image Registry Operator to allow for authentication by only storage account key without validating for other authentication requirements. (OCPBUGS-42514)

- Previously, installation with the OpenShift installer used the cluster API. Virtual networks created by the cluster API use a different tag template. Consequently, setting **.spec.storage.azure.networkAccess.type: Internal** in the Image Registry Operator's **config.yaml** file resulted in the Image Registry Operator unable to discover the virtual network. With this update, the Image Registry Operator searches for both new and old tag templates, resolving the issue. (OCPBUGS-42196)

- Previously, 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, fixing the panic. (OCPBUGS-39108)

1.6.9. Installer

- 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-38118)

- 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-33311)

- Previously, when installing a Google Cloud cluster where instances required that IP forwarding was not set, the installation failed. With this release, IP forwarding is disabled for all Google Cloud machines and the issue is resolved. (OCPBUGS-49842)

- Previously, when installing a cluster on AWS in existing subnets, for bring your own virtual private cloud (BYO VPC) in edge zones, the installation program did not tag the subnet edge resource with **kubernetes.io/cluster/<InfraID>:shared**. With this release, all subnets that are used in the **install-config.yaml** file contain the required tags. (OCPBUGS-49792)

- Previously, a cluster that was created on Amazon Web Services (AWS) could fail to deprovision the cluster without the permissions to release the EIP address, **ec2:ReleaseAddress**. This issue occurred when the cluster was created with the minimum permissions in an existing virtual private cloud (VPC), including an unmanaged VPC or bring your own (BYO) VPC, and BYO Public IPv4 Pool address. With this release, the **ec2:ReleaseAddress** permission is exported to the Identity and Access Management (IAM) policy generated during installation. (OCPBUGS-49735)

- Previously, when installing a cluster on Nutanix, the installation program could fail with a timeout while uploading images to Prism Central. This occurred in some slower Prism Central environments when the Prism API attempted to load the Red Hat Enterprise Linux CoreOS (RHCOS) image. The Prism API call timeout value was 5 minutes. With this release, the Prism API call timeout value is a configurable parameter **platform.nutanix.prismAPICallTimeout** in the **install-config.yaml** file and the default timeout value is 10 minutes. (OCPBUGS-49148)

- Previously, the **oc adm node-image monitor** command failed because of a temporary API server disconnection and then displayed an error or End of File message. With this release, the installation program ignores a temporary API server disconnection and the monitor command tries to connect to the API server again. (OCPBUGS-48714)

- Previously, when you deleted backend service resources on Google Cloud, some resources to be deleted were not found. For example, the associated forwarding rules, health checks, and firewall rules were not deleted. With this release, the installation program tries to find the backend service by name first, then searches for forwarding rules, health checks, and firewall rules before it determines if those results match a backend service. The algorithm for associating resources is reversed and the appropriate resources are deleted. There are no leaked backend service resources and the issue is resolved. When you delete a private cluster, the forwarding rules, backend services, health checks, and firewall rules created by the Ingress Operator are not deleted. (OCPBUGS-48611)

- Previously, OpenShift Container Platform was not compliant with PCI-DSS/BAFIN regulations. With this release, the cross-tenant object replication in Microsoft Azure is unavailable. Consequently, the chance of unauthorized data access is reduced and the strict adherence to data governance policies is ensured. (OCPBUGS-48118)

- Previously, when you installed OpenShift Container Platform on Amazon Web Services (AWS) and specified an edge machine pool without an instance type, in some instances it caused the edge node to fail. With this release, if you specify an edge machine pool without an instance type you must use the permission **ec2:DescribeInstanceTypeOfferings**. The permission derives the correct instance type available, based on the AWS Local Zones or Wavelength Zones locations used. (OCPBUGS-47502)

- Previously, when the API server disconnected temporarily, the command **oc adm node-image monitor** reported an end of file (EOF) error. With this release, when the API server disconnects temporarily, the monitor command does not fail. (OCPBUGS-46391)

- Previously, when you specified the **HostedZoneRole** permission in the **install-config.yaml** file while creating a shared Virtual Private Cloud (VPC), you also had to specify the **sts:AssumeRole** permission. Otherwise, it caused an error. With this release, if you specify the **HostedZoneRole** permission the installation program validates that the **sts:AssumeRole** permission is present. (OCPBUGS-46046)

- Previously, when the **publicIpv4Pool** configuration parameter was used during installation the permissions **ec2:AllocateAddress** and **ec2:AssociateAddress** were not validated. As a consequence, permission failures could occur during installation. With this release, the required permissions are validated before the cluster is installed and the issue is resolved. (OCPBUGS-45711)

- Previously, during a disconnected installation, when the **imageContentSources** parameter was configured for more than one mirror for a source, the command to create the agent ISO image could fail, depending on the sequence of the mirror configuration. With this release, multiple mirrors are handled correctly when the agent ISO is created and the issue is resolved. (OCPBUGS-45630)

- Previously, when installing a cluster using the Cluster API on installer-provisioned infrastructure, the user provided a **machineNetwork** parameter. With this release, the installation program uses a random **machineNetwork** parameter. (OCPBUGS-45485)

- Previously, during an installation on Amazon Web Services (AWS), the installation program used the wrong load balancer when searching for the **hostedZone** ID, which caused an error. With the release, the correct load balancer is used and the issue is resolved. (OCPBUGS-45301)

- Previously, endpoint overrides in IBM Power Virtual Server were not conditional. As a consequence, endpoint overrides were created incorrectly and caused failures in Virtual Private Environments (VPE). With this release, endpoint overrides are conditional only for disconnected installations. (OCPBUGS-44922)

- Previously, during a shared Virtual Private Cloud (VPC) installation, the installation program added the records to a private DNS zone created by the installation program instead of adding the records to the cluster's private DNS zone. As a consequence, the installation failed. With this release, the installation program searches for an existing private DNS zone and, if found, pairs that zone with the network that is supplied by the **install-config.yaml** file and the issue is resolved. (OCPBUGS-44641)

- Previously, the **oc adm drain --delete-local-data** command was not supported in the 4.18 **oc** CLI tool. With this release, the command has been updated to **oc adm drain --delete-emptydir-data**. (OCPBUGS-44318)

- Previously, US East (**wdc04**), US South (**dal13**), Sydney (**syd05**), and Toronto (**tor01**) regions were not supported for IBM Power Virtual Server. With this release, these regions, which include **PowerEdgeRouter** (PER) capabilities, are supported for IBM Power Virtual Server. (OCPBUGS-44312)

- Previously, during a Google Cloud installation, when the installation program was creating filters with large numbers of returned data, for example for subnets, it exceeded the quota for the maximum number times that a resource can be filtered in a specific period. With this release, all relevant filtering is moved to the client so that the filter quotas are not exceeded and the issue is resolved. (OCPBUGS-44193)

- Previously, during an Amazon Web Services (AWS) installation, the installation program validated all the tags in the **install-config.yaml** file only when you set **propogateTags** to true. With this release, the installation program validates all the tags in the **install-config.yaml** file. (OCPBUGS-44171)

- Previously, if the **RendezvousIP** value matched a substring in the **next-hop-address** field of a compute node configuration, it reported a validation error. The **RendezvousIP** value must match a control plane host address only. With this release, a substring comparison for

RendezvousIP value is used against a control plane host address only, so that the error no longer exists. (OCPBUGS-44167)

- Previously, when you deleted a cluster in IBM Power Virtual Server, the Transit Gateway connections were cleaned up. With this release, if the **tgName** parameter is set, Red Hat OpenStack Platform (RHOSP) does not clean up the Transit Gateway connection when you delete a cluster. (OCPBUGS-44162)

- 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-44068)

- Previously, the Cluster API Provider IBM Cloud module was very verbose. With this release, the verbosity of the module is reduced, and this will affect the output of the **.openshift_install.log** file. (OCPBUGS-44022)

- Previously, when you deployed a cluster on a IBM Power Virtual Server zone, the load balancers were slow to create. As a consequence, the cluster failed. With this release, the Cluster API Provider IBM Cloud no longer has to wait until all load balancers are ready and the issue is resolved. (OCPBUGS-43923)

- Previously, for the Agent-based Installer, all host validation status logs referred to the name of the first registered host. As a consequence, when a host validation failed, it was not possible to determine the problem host. With this release, the correct host is identified in each log message and now the host validation logs correctly show the host to which they correspond, and the issue is resolved. (OCPBUGS-43768)

- Previously, when you used the **oc adm node-image create** command to generate the image while running the Agent-based Installer and the step fails, the accompanying error message did not show the container log. The **oc adm node-image create** command uses a container to generate the image. When the image generation step fails, the basic error message does not show the underlying issue that caused the image generation failure. With this release, to help troubleshooting, the **oc adm node-image create** command now shows the container log, so the underlying issue is displayed. (OCPBUGS-43757)

- Previously, the Agent-based Installer failed to parse the **cloud_controller_manager** parameter in the **install-config.yaml** configuration file. This resulted in the Assisted Service API failing because it received an empty string, and this in turn caused the installation of the cluster to fail on Oracle® Cloud Infrastructure (OCI). With this release, an update to the parsing logic ensures that the Agent-based Installer correctly interprets the **cloud_controller_manager** parameter so that the Assisted Service API receives the correct string value. As a result, the Agent-based Installer can now installer a cluster on OCI. (OCPBUGS-43674)

- Previously, an update to Azure SDK for Go removed the **SendCertificateChain** option and this changed the behavior of sending certificates. As a consequence, the full certificate chain was not sent. With this release, the option to send a full certification chain is available and the issue is resolved. (OCPBUGS-43567)

- Previously, when installing a cluster on Google Cloud using the Cluster API implementation, the installation program did not distinguish between internal and external load balancers while creating firewall rules. As a consequence, the firewall rule for internal load balancers was open to all IP address sources, that is, **0.0.0.0/0**. With this release, the Cluster API Provider GCP is updated to restrict firewall rules to the machine CIDR when using an internal load balancer. The firewall rule for internal load balancers is correctly limited to machine networks, that is, nodes in the cluster and the issue is resolved. (OCPBUGS-43520)

- Previously, when installing a cluster on IBM Power Virtual Server, the required security group rules were not created. With this release, the missing security group rules for installation are identified and created and the issue is resolved. (OCPBUGS-43518)
- Previously, when you tried to add a compute node with the **oc adm node-image** command by using an instance that was previously created with Red Hat OpenStack Platform (RHOSP), the operation failed. With this release, the issue is resolved by correctly setting the user-managed networking configuration. (OCPBUGS-43513)
- Previously, when destroying a cluster on Google Cloud, a forwarding rule incorrectly blocked the installation program. As a consequence, the destroy process failed to complete. With this release, the issue is resolved by the installation program setting its state correctly and marking all destroyed resources as deleted. (OCPBUGS-42789)
- Previously, when configuring the Agent-Based Installer installation in a disconnected environment with more than one mirror for the same source, the installation might fail. This occurred because one of the mirrors was not checked. With this release, all mirrors are used when multiple mirrors are defined for the same source and the issue is resolved. (OCPBUGS-42705)
- Previously, you could not change the **AdditionalTrustBundlePolicy** parameter in the **install-config.yaml** file for the Agent-based Installer. The parameter was always set to **ProxyOnly**. With this release, you can set **AdditionalTrustBundlePolicy** to other values, for example, **Always**. By default, the parameter is set to **ProxyOnly**. (OCPBUGS-42670)
- Previously, when you installed a cluster and tried to add a compute node with the **oc adm node-image** command, it failed because the date, time, or both might have been inaccurate. With this release, the issue is resolved by applying the same Network Time Protocol (NTP) configuration in the target cluster **MachineConfig** chrony resource to the node ephemeral live environment (OCPBUGS-42544)
- Previously, during installation the name of the artifact that the **oc adm node-image create** command generated did not include **<arch>** in its file name. As a consequence, the file name was inconsistent with other generated ISOs. With this release, a patch fixes the name of the artifact that is generated by the **oc adm node-image create** command by also including the referenced architecture as part of the file name and the issue is resolved. (OCPBUGS-42528)
- Previously, the Agent-based Installer set the **assisted-service** object to a debug logging mode. Unintentionally, the **pprof** module in the **assisted-service** object, which uses port **6060**, was then turned on. As a consequence, there was a port conflict and the Cloud Credential Operator (CCO) did not run. When requested by the VMware vSphere Cloud Controller Manager (CCM), vSphere secrets were not generated, the RHOSP CCM failed to initialize the nodes, and the cluster installation was blocked. With this release, the **pprof** module in the **assisted-service** object does not run when invoked by the Agent-based Installer. As a result, the CCO runs correctly and cluster installations on vSphere that use the Agent-based Installer succeed. (OCPBUGS-42525)
- Previously, when a compute node was trying to join a cluster the rendezvous node rebooted before the process completed. As the compute node could not communicate as expected 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-41811)

- Previously, when using the Assisted Installer, selecting a multi-architecture image for **s390x** CPU architecture on Red Hat Hybrid Cloud Console could cause the installation to fail. The installation program reported an error that the new cluster was not created because the skip MCO reboot was not compatible with **s390x** CPU architecture. With this release, the issue is resolved. (OCPBUGS-41716)

- Previously, a coding issue caused the Ansible script on RHOSP user-provisioned infrastructure installation to fail during the provisioning of compact clusters. This occurred when IPv6 was enabled for a three-node cluster. With this release, the issue is resolved and you can provision compact three-node clusters. (OCPBUGS-41538)

- Previously, a coding issue caused the Ansible script on RHOSP user-provisioned installation infrastructure to fail during the provisioning of compact clusters. This occurred when IPv6 was enabled for a three-node cluster. With this release, the issue is resolved and you can provision compact three-node clusters on RHOSP for user-provisioned installation infrastructure. (OCPBUGS-39402)

- 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 to accommodate older versions of Ansible and the issue is resolved. (OCPBUGS-39285)

- Previously, when you installed a cluster there were issues using a compute node because the date, time, or both might have been inaccurate. With this release, a patch is applied to the live ISO time synchronization. The patch configures the **/etc/chrony.conf** file with the list of the additional Network Time Protocol (NTP) servers that the user provides in the **agent-config.yaml** file, so that you can use a compute node without experiencing a cluster installation issue. (OCPBUGS-39231)

- 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-39081)

- 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 includes the cluster proxy resource settings, if available, to ensure the command is run successfully and the issue is resolved. (OCPBUGS-38990)

- Previously, when installing a cluster on Google Cloud into a shared Virtual Private Cloud (VPC) with a bring your own (BYO) hosted zone, the installation could fail due to an error creating the private managed zone. With this release, a fix ensures that where there is a preexisting private managed zone the installation program skips creating a new one and the issue is resolved. (OCPBUGS-38966)

- Previously, an installer-provisioned installation on VMware vSphere to run OpenShift Container Platform 4.16 in a disconnected environment failed when the template could not be downloaded. With this release, the template is downloaded correctly and the issue is resolved. (OCPBUGS-38918)

- 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-38802)

- Previously, a rare condition on 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 Cluster API Provider vSphere, and the issue is resolved. (OCPBUGS-38657)

- Previously, when a folder was undefined and the data center was located in a data center folder, a wrong folder structure was created starting from the root of the vCenter server. By using the **Govmomi DatacenterFolders.VmFolder**, it used the 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-38599)

- Previously, the installation program on Google Cloud filtered addresses to find and delete internal addresses only. The addition of Cluster API Provider Google Cloud Platform (GCP) provisioned resources included changes to address resources. With this release, Cluster API Provider GCP creates external addresses and these must be included in a cluster cleanup operation. (OCPBUGS-38571)

- 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 that the specified cluster architecture is compatible with OpenShift Container Platform, leading to successful installations. (OCPBUGS-38479)

- 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-38466)

- 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-388118)

- 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-37945)

- 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, when deprecated and supported fields were used with the installation program on VMware vSphere a validation error message was reported. With this release, warning messages are added specifying that using deprecated and supported fields are not recommended with the installation program on VMware vSphere. (OCPBUGS-37628)

- Previously, if you tried to install a second cluster using existing Azure Virtual Networks (VNet) on Microsoft Azure, the installation failed. Where the front end IP address of the API server load balancer was not specified, the Cluster API fixed the address to **10.0.0.100**. As this IP address was already taken by the first cluster, this resulted in the second load balancer failing to install. With this release, a dynamic IP address checks whether the default IP address is available. If it is unavailable, the dynamic IP selects the next available address and you can install the second cluster successfully with a different load balancer IP. (OCPBUGS-37442)

- 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-36494)

- Previously, when installing a cluster on IBM Cloud the installation program checked the first group of subnets, that is 50, only when searching for subnet details by name. With this release, pagination support is provided to search all subnets. (OCPBUGS-36236)

- Previously, when installing Cluster API Provider Google Cloud Platform (GCP) into a shared Virtual Private Cloud (VPC) without the required permission **compute.firewalls.create** the installation failed because no firewall rules were created. With this release, a fix ensures that a rule to create the firewall is skipped during installation and the issue is resolved. (OCPBUGS-35262)

- Previously, for the Agent-Based installer, the networking layout defined through nmstate might result in a configuration error if all hosts do not have an entry in the interfaces section that matches an entry in the **networkConfig** section. However, if the entry in the **networkConfig** section uses a physical interface name then the entry in the interfaces section is not required. This fix ensures that the configuration will not result in an error if an entry in the **networkConfig** section has a physical interface name and does not have a corresponding entry in the interfaces table. (OCPBUGS-34849)

- Previously, the container tools module was enabled by default on the RHEL node. With this release, the container-tools module is disabled to install the correct package between conflicting repositories. (OCPBUGS-34844)

1.6.10. Insights Operator

- Previously, during entitled builds on a Red Hat OpenShift Container Platform cluster running on IBM Z hardware, repositories were not enabled. This issue has been resolved. You can now enable repositories during entitled builds on a Red Hat OpenShift Container Platform cluster running on IBM Z hardware. (OCPBUGS-32233)

1.6.11. Machine Config Operator

- 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-42324)

- 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-34647)

1.6.12. Management Console

- Previously, clicking the "Don’t show again" link in the Lightspeed modal dialog did not correctly navigate to the general **User Preference** tab when one of the other **User Preference** tabs was displayed. After this update, clicking the "Don't show again" link correctly navigates to the general **User Preference** tab. (OCPBUGS-48106)

- Previously, multiple external link icons might show in the primary action button of the OperatorHub modal. With this update, only a single external link icon appears. (OCPBUGS-47742)

- Previously, the web console was disabled when the authorization type was set to **None** in the cluster authentication configuration. With this update, the web console no longer disables when the authorization type was set to **None**. (OCPBUGS-46068)

- Previously, the **MachineConfig Details** tab displayed an error when one or more **spec.config.storage.file** did not include optional data. With this update, the error no longer occurs and the **Details** tab renders as expected. (OCPBUGS-44049)

- Previously, an extra name property was passed into resource list page extensions used to list related Operands on the **CSV details** page. As a result, the Operand list was filtered by the cluster service version (CSV) name and often returned an empty list. With this update, Operands are listed as expected. (OCPBUGS-42796)

- Previously, the **Sample** tab did not show when creating a new ConfigMap with one or more ConfigMap ConsoleYAMLSamples present on the cluster. After this update, the **Sample** tab shows with one or more ConfigMap ConsoleYAMLSamples present. (OCPBUGS-41492)

- Previously, the **Events** page resource type filter incorrectly reported the number of resources when three or more resources were selected. With this update, the filter always reports the correct number of resources. (OCPBUGS-38701)

- Previously, the version number text in the updates graph on the **Cluster Settings** page appeared as black text on a dark background while viewing the page using Firefox in dark mode. With this update, the text appears as white text. (OCPBUGS-37988)

- Previously, **Alerting** pages did not show resource information in their empty state. With this update, resource information is available on the **Alerting** pages. (OCPBUGS-36921)

- Previously, the Operator Lifecycle Manager (OLM) CSV annotation contained unexpected JSON, which was successfully parsed, but then threw a runtime error when attempting to use the resulting value. With this update, JSON values from OLM annotations are validated before use, errors are logged, and the console does not fail when unexpected JSON is received in an annotation. (OCPBUGS-35744)

- Previously, silenced alerts were visible on the **Overview** page of the OpenShift Container Platform web console. This occurred because the alerts did not include any external labels. With this release, silenced alerts include the external labels so they are filtered out and are not viewable. (OCPBUGS-31367)

1.6.13. Monitoring

- Previously, if the SMTP **smarthost** or **from** fields under the **emailConfigs** object were not specified at the global or receiver level in the **AlertmanagerConfig** custom resource (CR), Alertmanager would crash because these fields are required. With this release, the Prometheus Operator fails reconciliation if these fields are not specified. Therefore, the Prometheus Operator no longer pushes invalid configurations to Alertmanager, preventing it from crashing. (OCPBUGS-48050)

- Previously, the Cluster Monitoring Operator (CMO) did not mark configurations in **cluster-monitoring-config** and **user-workload-monitoring-config** config maps as invalid for unknown (for example, no longer supported) or duplicated fields. With this release, stricter validation is added that helps identify such errors. (OCPBUGS-42671)

- Previously, it was not possible for a user to query the user workload monitoring Thanos API endpoint with **POST** requests. With this update, a cluster admin can bind a new **pod-metrics-reader** cluster role with a role binding or cluster role binding to allow **POST** queries for a user or service account. (OCPBUGS-41158)

- Previously, an invalid config map configuration for core platform monitoring, user workload monitoring, or both caused Cluster Monitoring Operator (CMO) to report an **InvalidConfiguration** error. With this release, if only the user workload monitoring configuration is invalid, CMO reports **UserWorkloadInvalidConfiguration**, making it clear where the issue is located. (OCPBUGS-33863)

- Previously, **telemeter-client containers** showed a **TelemeterClientFailures Warnings** message in multiple clusters. With this release, a runbook is added for the **TelemeterClientFailures** alert to explain the cause of the alert triggering and the alert provides resolution steps. (OCPBUGS-33285)

- Previously, **AlertmanagerConfig** objects with invalid child routes generated invalid Alertmanager configuration leading to Alertmanager disruption. With this release, Prometheus Operator rejects such **AlertmanagerConfig** objects, and users receive a warning about the invalid child routes in logs. (OCPBUGS-30122)

- Previously, the **config-reloader** for Prometheus for user-defined projects would fail if unset environment variables were used in the **ServiceMonitor** configuration, which resulted in Prometheus pods failing. With this release, the reloader no longer fails when an unset environment variable is encountered. Instead, unset environment variables are left as they are, while set environment variables are expanded as usual. Any expansion errors, suppressed or otherwise, can be tracked through the **reloader_config_environment_variable_expansion_errors** variable. (OCPBUGS-23252)

1.6.14. Networking

- Previously, enabling encapsulated security payload (ESP) offload hardware when using IPSec on Open vSwitch attached interfaces would break connectivity in your cluster. To resolve this issue, OpenShift Container Platform by default disables ESP offload hardware on Open vSwitch attached interfaces. This fixes the issue. (OCPBUGS-42987)

- Previously, if you deleted the default **sriovOperatorConfig** custom resource (CR), you could not recreate the default **sriovOperatorConfig** CR, because the **ValidatingWebhookConfiguration** was not initially deleted. 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-41897)

- 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-41352)

- Previously, bonds that were configured in **active-backup** mode would have IPsec Encapsulating Security Payload (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-39438)

- 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 4.18 4.13 and above, through manual intervention, or by upgrading to a release that includes this fix. (OCPBUGS-38012)

- Previously, the Ingress Controller status incorrectly displayed as **Degraded=False** because of a migration time 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-37491)

- 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 the traffic with the IP family, that the egressIP is not applicable to, being dropped. With this release, only the source network address translation (SNAT) for the IP family that the egress IPs applied to is deleted, eliminating the risk of traffic being dropped. (OCPBUGS-37193)

- 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-23795)

- 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-13181)

- Previously, when either the **clusterNetwork** or **serviceNetwork** IP address pools overlapped with the default **transit_switch_subnet 100.88.0.0/16** IP address and the custom value of **transit_switch_subnet** did not take effect, **ovnkube-node** pods crashed after the live migration operation. With this release, the custom value of **transit_switch_subnet** can be passed to **ovnkube node** pods, so that this issue no longer persists. (OCPBUGS-43740)

- Previously, a change in OVN-Kubernetes that standardized the **appProtocol** value **h2c** to **kubernetes.io/h2c** was not recognized by OpenShift router. Consequently, specifying **appProtocol: kubernetes.io/h2c** on a service did not cause OpenShift router to use clear-text HTTP/2 to connect to the service endpoints. With this release, OpenShift router was changed to handle **appProtocol: kubernetes.io/h2c** the same way as it handles **appProtocol: h2c** resolving the issue. (OCPBUGS-42972)

- Previously, instructions that guided the user after changing the **LoadBalancer** parameter from **External** to **Internal** were missing for IBM Power Virtual Server, Alibaba Cloud, and Red Hat OpenStack Platform (RHOSP). This caused the Ingress Controller to be put in a permanent **Progressing** state. With this release the message **The IngressController scope was changed from Internal to External** is followed by **To effectuate this change, you must delete the service** resolving the permanent **Progressing** state. (OCPBUGS-39151)

- Previously, there was no event logged when an error occurred from failed conversion from ingress to route conversion. With this update, this error appear in the event logs. (OCPBUGS-29354)

- Previously, an **ovnkube-node** pod on a node that uses cgroup v1 was failing because it could not find the kubelet cgroup path. With this release, an **ovnkube-node** pod no longer fails if the node uses cgroup v1. However, the OVN-Kubernetes network plugin outputs an **UDNKubeletProbesNotSupported** event notification. If you enable cgroup v2 for each node, OVN-Kubernetes no longer outputs the event notification.(OCPBUGS-50513)

- Previously, when you finished the live migration for a kubevirt virtual machine (VM) that uses the Layer 2 topology, an old node still transmits IPv4 egress traffic to the virtual machine. With this release, the OVN-Kubernetes plugin updates the gateway MAC address for a kubevirt virtual machine (VM) during the live migration process so that this issue no longer occurs. (OCPBUGS-49857)

- Previously, the DNS-based egress firewall incorrectly prevented creation of a firewall rule that contained a DNS name in uppercase characters. With this release, an fix to the egress firewall no longer prevents creation of a firewall rule that contains a DNS name in uppercase characters. (OCPBUGS-49589)

- Previously, when you attempted to use the Cluster Network Operator (CNO) to upgrade a cluster with existing **localnet** networks, **ovnkube-control-plane** pods failed 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-44195)

- Previously, the SR-IOV Network Operator could not retrieve metadata when cloud-native network (CNF) workers were deployed with a configuration drive on Red Hat OpenStack Platform (RHOSP). A configuration drive is often unmounted after a boot operation on immutable systems, so now the Operator dynamically mounts a configuration drive when required. The Operator can now retrieve the metadata and then unmount the configuration drive. This means that you no longer need to manually mount and unmount the configuration drive. (OCPBUGS-41829)

- 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-38217)

- Previously, a duplicate feature gate, **ExternalRouteCertificate**, was added to the **FeatureGate** CR. With this release, **ExternalRouteCertificate** is removed because a OpenShift Container Platform cluster does not use this feature gate. (OCPBUGS-36479)

- Previously, after a user created a route, the user needed both **create** and **update** permissions on the **routes/custom-host** sub-resource to edit the **.spec.tls.externalCertificate** field of a route. With this release, this permission requirement has been fixed, so that a user only needs the **create** permission to edit the **.spec.tls.externalCertificate** field of a route. The **update** permission is now marked as an optional permission. (OCPBUGS-34373)

1.6.15. Node

- Previously, the **cadvisor** code that collected and reported container network metrics contained a bug that caused inaccurate results. With this release, the container network metrics are correctly reported. (OCPBUGS-38515)

1.6.16. Node Tuning Operator (NTO)

- Previously, CPU masks for interrupt and network handling CPU affinity were computed incorrectly on machines with more than 256 CPUs. This issue prevented proper CPU isolation and caused **systemd** unit failures during internal node configuration. This fix ensures accurate CPU affinity calculations, enabling correct CPU isolation on machines with more than 256 CPUs. (OCPBUGS-36431)

- Previously, entering an invalid value in any **cpuset** field under **spec.cpu** in the **PerformanceProfile** resource caused the webhook validation to crash. With this release, improved error handling for the **PerformanceProfile** validation webhook ensures that invalid values for these fields return an informative error. (OCPBUGS-45616)

- 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-47678)

- Previously, configuring the Node Tuning Operator (NTO) using **PerformanceProfiles** created the **ocp-tuned-one-shot systemd** service, which ran before kubelet and blocked its execution. The **systemd** service invoked Podman, which used the NTO image. When the NTO image was not present, Podman tried to fetch the image. With this release, support for cluster-wide proxy environment variables defined in **/etc/mco/proxy.env** is added. This support allows Podman to pull the NTO image in environments that need to use **http(s)** proxy for out-of-cluster connections. (OCPBUGS-39005)

1.6.17. Observability

- Previously, a namespace was passed to a full cluster query on the alerts graph, and this caused the tenancy API path to be used. The API lacked permissions to retrieve data so no data was shown on the alerts graph. With this release, the namespace is no longer passed to a full cluster query for an alert graph. A non-tenancy API path is now used because this API has the correct permissions to retrieve data. Data is not available on an alert graph. (OCPBUGS-46371)

- Previously, bounds were based on the first bar in a bar chart. If a bar was larger in size than the first bar, the bar would extend beyond the bar chart boundary. With this release, the bound for a bar chart is based on the largest bar, so no bars extend outside the boundary of a bar chart. (OCPBUGS-46059)

- Previously, a Red Hat Advanced Cluster Management (RHACM) Alerting UI refactor update caused an **isEmpty** check to go missing on the **Observe → Metrics** menu. The missing check inverted the behavior of the **Show all Series** and **Hide all Series** states. This release readds **isEmpty** check so that **Show all Series** is now visible when series are hidden and **Hide all Series** is now visible when the series are shown. (OCPBUGS-46047)

- Previously, on the **Observe → Alerting → Silences** tab, the **DateTime** component changed the ordering of an event and its value. Because of this issue, you could not edit the **until** parameter for a silent alert in either the **Developer** or the **Administrator** perspective. With this release, a fix means to the **DateTime** component means that you can now edit the **until** parameter for a silent alert. (OCPBUGS-46021)

- Previously, when using the **Developer** perspective with custom editors, clicking the **n** key caused the **Namespace** menu unexpectedly opened. The issue happened because the keyboard shortcut did not account for custom editors. With this release, the **Namespace** menu accounts for custom editors and does not open if you type the **n** key. (OCPBUGS-38775)

- Previously, on the **Observe → Alerting → Silences** tab, the **creator** field was not autopopulated and was not designated as mandatory. This issue happened when the API made the field empty from OpenShift Container Platform 4.15 and onwards. With this update, the field is marked as mandatory and populated with the current user for correct validation. (OCPBUGS-35048)

1.6.18. oc-mirror

- Previously, when using **oc-mirror --v2 delete --generate** command, the contents of the **working-dir/cluster-resources** directory were cleared. With this fix, the **working-dir/cluster-resources** directory is not cleaned when the delete feature is used. (OCPBUGS-48430)

- Previously, release images were signed using a **SHA-1** key. On RHEL 9 FIPS STIG-compliant machines, verification of release signatures using the old **SHA-1** key failed due to security restrictions on weak keys. With this release, release images are signed using a new **SHA-256** trusted key so that the release signatures no longer fail. (OCPBUGS-48314)

- Previously, when using the **--force-cache-delete** flag to delete images from a remote registry, the deletion process did not work as expected. With this update, the issue has been resolved, ensuring that images are deleted properly when the flag is used. (OCPBUGS-47690)

- Previously, oc-mirror plugin v2 could not delete the graph image when the mirroring uses a partially disconnected mirroring workflow (mirror-to-mirror). With this update, graph images can now be deleted regardless of the mirroring workflow used. (OCPBUGS-46145)

- Previously, if the same image was used by multiple OpenShift Container Platform release components, oc-mirror plugin v2 attempted to delete the image multiple times, but failed after the first attempt. This issue has been resolved by ensuring oc-mirror plugin v2 generates a list of unique images during the delete **--generate** phase. (OCPBUGS-45299)

- Previously, **oci** catalogs on disk were not mirrored correctly in the oc-mirror plugin v2. With this update, **oci** catalogs are now successfully mirrored. (OCPBUGS-44225)

- Previously, if you reran the **oc-mirror** command, the rebuild of the **oci** catalog failed and an error was generated. With this release, if you rerun the **oc-mirror** command, the wrokspace file is deleted so that the failed catalog issue does not happen. (OCPBUGS-45171)

- Previously, if you ran the **oc adm node-image create** command on the first attempt, sometimes an **image can't be pulled** error message was generated. With this release, a retry mechanism addresses temporary failures when pulling the image from the release payload. (OCPBUGS-44388)

- Previously, duplicate entries could appear in the signature **ConfigMap YAML** and **JSON** files created in the **clusterresource** object, leading to issues when applying them to the cluster. This update ensures that the generated files do not contain duplicates. (OCPBUGS-42428)

- Previously, the release signature **ConfigMap** for oc-mirror plugin v2 was incorrectly stored in an archived TAR file instead of in the **cluster-resources** folder. This caused **mirror2disk** to fail. With this release. the release signature **ConfigMap** for oc-mirror plugin v2 that is in JSON format or YAML format, compatible with oc-mirror plugin v1, now get stored in the **cluster-resources** folder. (OCPBUGS-38343) and (OCPBUGS-38233)

- Previously, using an invalid log-level flag caused oc-mirror plugin v2 to panic. This update ensures that the oc-mirror plugin v2 handles invalid log levels gracefully. Additionally, the **loglevel** flag has been renamed to **log-level** to align with tools like Podman for the convenience of the user. (OCPBUGS-37740)

1.6.19. OpenShift CLI (oc)

- Previously, the **oc adm node-image create --pxe generated** command did not create only the Preboot Execution Environment (PXE) artifacts. Instead, the command created the PXE artifacts with other artifacts from a **node-joiner** pod and stored them all in the wrong subdirectory. Additionally, the PXE artifacts were incorrectly prefixed with **agent** instead of **node**. With this release, generated PXE artifacts are stored in the correct directory and receive the correct prefix. (OCPBUGS-46449)

- Previously, requests to the **deploymentconfig/scale** subresource would fail when there was an admission webhook matching the request. With this release, the issue is resolved and requests to the **deploymentconfig/scale** subresource will succeed. (OCPBUGS-41136)

1.6.20. Operator Lifecycle Manager (OLM)

- Previously, concurrent reconciliation of the same namespace in Operator Lifecycle Manager (OLM) Classic led to **ConstraintsNotSatisfiable** errors on subscriptions. This update resolves the issue. (OCPBUGS-48660)

- Previously, excessive catalog source snapshots caused severe performance regressions. This update fixes the issue. (OCPBUGS-48644)

- Previously, when the kubelet terminated catalog registry pods with the **TerminationByKubelet** message, the registry pods were not recreated by the catalog Operator. This update fixes the issue. (OCPBUGS-46474)

- Previously, OLM (Classic) failed to upgrade Operator cluster service versions (CSVs) due to a TLS validation error. This update fixes the issue. (OCPBUGS-43581)

- Previously, service account tokens for Operator groups failed to generate automatically in Operator Lifecycle Manager (OLM) Classic. This update fixes the issue. (OCPBUGS-42360)

- Previously when Operator Lifecycle Manager (OLM) v1 validated custom resource definition (CRD) upgrades, the message output when detecting changed default values was rendered in bytes instead of human-readable language. With this update, related messages are now updated to show human-readable values. (OCPBUGS-41726)

- Previously, the status update function did not return an error when a connection error occurred in the Catalog Operator. As a result, the Operator might crash because the IP address returned a **nil** status. This update resolves the issue so that an error message is returned and the Operator no longer crashes. (OCPBUGS-37637)

- Previously, catalog source registry pods did not recover from cluster node failures. This update fixes the issue. (OCPBUGS-36661)

- Previously, Operators with many custom resources (CRs) exceeded API server timeouts. As a result, the install plan for the Operator got stuck in a pending state. This update fixes the issue by adding a page view for list CRs deployed on the cluster. (OCPBUGS-35358)

1.6.21. Performance Addon Operator

- 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-45903)

- Previously, if you specified a long string of isolated CPUs in a performance profile, such as **0,1,2,…,512**, the **tuned**, Machine Config Operator and **rpm-ostree** components failed to process the string as expected. As a consequence, after you applied the performance profile, the expected kernel arguments were missing. The system failed silently with no reported errors. With this release, the string for isolated CPUs in a performance profile is converted to sequential ranges, such as **0-512**. As a result, the kernel arguments are applied as expected in most scenarios. (OCPBUGS-45472)

  NOTE

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

1.6.22. Red Hat Enterprise Linux CoreOS (RHCOS)

- Previously, the **kdump** initramfs would stop responding when trying to open a local encrypted disk. This occurred even when the **kdump** destination was a remote machine that did not need access to the local disk. With this release, the issue is fixed and the **kdump** initramfs successfully opens a local encrypted disk. (OCPBUGS-43040)

- Previously, explicitly disabling FIPS mode with **fips=0** caused some systemd services, that assume FIPS mode was requested, to run and consequently fail. This issue resulted in RHCOS failing to boot. With this release, the relevant systemd services now only run if FIPS mode is enabled by specifying **fips=1**. As a result, RHCOS now correctly boots without FIPS mode enabled when **fips=0** is specified. (OCPBUGS-39536)

1.6.23. Scalability and performance

- Previously, you could configure the NUMA Resources Operator to map a **nodeGroup** to more than one **MachineConfigPool**. This implementation is contrary to the intended design of the Operator, which assumed a one-to-one mapping between a **nodeGroup** and a **MachineConfigPool**. With this release, if a **nodeGroup** maps to more than one **MachineConfigPool**, the Operator accepts the configuration, but the Operator state moves to **Degraded**. To retain the previous behavior, you can apply the **config.node.openshift-kni.io/multiple-pools-per-tree: enabled** annotation to the NUMA Resources Operator. However, the ability to assign a **nodeGroup** to more than one **MachineConfigPool** will be removed in a future release. (OCPBUGS-42523)

1.6.24. Storage

- Previously, Portworx plugin Container Storage Interface (CSI) migration failed without the inclusion of an upstream patch. With this release, the Portworx plugin CSI translation now copies the secret name and namespace to Kubernetes version to 1.31 so that an upstream patch is not required. (OCPBUGS-49437)

- Previously, the VSphere Problem Detector Operator waited up to 24 hours to reflect a change in the **clustercsidrivers.managementState** parameter from **Managed** to **Removed** for a VMware vSphere cluster. With this release, the VSphere Problem Detector Operator now reflects this state change in about 1 hour. (OCPBUGS-39358)

- Previously, the Azure File Driver attempted to reuse existing storage accounts. With this release, the Azure File Driver creates storage accounts during dynamic provisioning. This means that updated clusters using newly-created Persistent Volumes (PVs) also use a new storage account. PVs that were previously provisioned continue using the same storage account used before the cluster update. (OCPBUGS-38922)

- Previously, the configuration loader logged YAML **unmarshall** errors when the **INI** succeeded. With this release, the **unmarshall** errors are no longer logged when the **INI** succeeds. (OCPBUGS-38368)

- Previously, the Storage Operator counted an incorrect number of control plane nodes that existed in a cluster. This count is needed for the Operator to determine the number of replicas for controllers. With this release, the Storage Operator now counts the correct number of control plane nodes, leading to a more accurate count of replica controllers. (OCPBUGS-36233)

- Previously, the **manila-csi-driver** and node registrar pods had missing health checks because of a configuration issue. With this release, the health checks are now added to both of these resources. (OCPBUGS-29240)

1.7. TECHNOLOGY PREVIEW FEATURES STATUS

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

Technology Preview Features Support Scope

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

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

1.7.1. Authentication and authorization Technology Preview features

Table 1.18. Authentication and authorization Technology Preview tracker

Feature	4.16	4.17	4.18
Pod security admission restricted enforcement	Technology Preview	Technology Preview	Technology Preview

1.7.2. Edge computing Technology Preview features

Table 1.19. Edge computing Technology Preview tracker

Feature	4.16	4.17	4.18
Accelerated provisioning of GitOps ZTP	Technology Preview	Technology Preview	Technology Preview
Enabling disk encryption with TPM and PCR protection	Not Available	Technology Preview	Technology Preview

1.7.3. Installation Technology Preview features

Table 1.20. Installation Technology Preview tracker

Feature	4.16	4.17	4.18
Adding kernel modules to nodes with kvc	Technology Preview	Technology Preview	Technology Preview
Enabling NIC partitioning for SR-IOV devices	Technology Preview	General Availability	General Availability
User-defined labels and tags for Google Cloud	Technology Preview	General Availability	General Availability
Installing a cluster on Alibaba Cloud by using Assisted Installer	Technology Preview	Technology Preview	Technology Preview
Mount shared entitlements in BuildConfigs in RHEL	Technology Preview	Technology Preview	Technology Preview
OpenShift Container Platform on Oracle® Cloud Infrastructure (OCI)	General Availability	General Availability	General Availability
Selectable Cluster Inventory	Technology Preview	Technology Preview	Technology Preview
Installing a cluster on Google Cloud using the Cluster API implementation	Technology Preview	General Availability	General Availability

Feature	4.16	4.17	4.18
OpenShift Container Platform on Oracle Compute Cloud@Customer (C3)	Not Available	Not Available	General Availability
OpenShift Container Platform on Oracle Private Cloud Appliance (PCA)	Not Available	Not Available	General Availability
Installing a cluster on VMware vSphere with multiple network interface controllers	Not Available	Not Available	Technology Preview

1.7.4. Machine Config Operator Technology Preview features

Table 1.21. Machine Config Operator Technology Preview tracker

Feature	4.16	4.17	4.18
Improved MCO state reporting (**oc get machineconfignode**)	Technology Preview	Technology Preview	Technology Preview
On-cluster RHCOS image layering	Technology Preview	Technology Preview	General Availability
Node disruption policies	Technology Preview	General Availability	General Availability
Updating boot images for GCP clusters	Technology Preview	General Availability	General Availability
Updating boot images for AWS clusters	Technology Preview	Technology Preview	General Availability
Pinned Image Sets	Technology Preview	Technology Preview	Technology Preview

1.7.5. Machine management Technology Preview features

Table 1.22. Machine management Technology Preview tracker

Feature	4.16	4.17	4.18
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

Feature	4.16	4.17	4.18
Managing machines with the Cluster API for IBM Power® Virtual Server	Technology Preview	Technology Preview	Technology Preview
Managing machines with the Cluster API for Microsoft Azure	Not Available	Not Available	Technology Preview
Managing machines with the Cluster API for RHOSP	Technology Preview	Technology Preview	Technology Preview
Managing machines with the Cluster API for VMware vSphere	Technology Preview	Technology Preview	Technology Preview
Cloud controller manager for IBM Power® Virtual Server	Technology Preview	Technology Preview	Technology Preview
Defining a vSphere failure domain for a control plane machine set	General Availability	General Availability	General Availability
Cloud controller manager for Alibaba Cloud	Removed	Removed	Removed
Adding multiple subnets to an existing VMware vSphere cluster by using compute machine sets	Not Available	Not Available	Technology Preview

1.7.6. Monitoring Technology Preview features

Table 1.23. Monitoring Technology Preview tracker

Feature	4.16	4.17	4.18
Metrics Collection Profiles	Technology Preview	Technology Preview	Technology Preview

1.7.7. Web console Technology Preview features

Table 1.24. Web console Technology Preview tracker

Feature	4.16	4.17	4.18
Red Hat OpenShift Lightspeed in the OpenShift Container Platform web console	Technology Preview	Technology Preview	Technology Peview

1.7.8. Multi-Architecture Technology Preview features

Table 1.25. Multi-Architecture Technology Preview tracker

Feature	4.16	4.17	4.18
kdump on **arm64** architecture	Technology Preview	Technology Preview	Technology Preview
kdump on **s390x** architecture	Technology Preview	Technology Preview	Technology Preview
kdump on **ppc64le** architecture	Technology Preview	Technology Preview	Technology Preview
Multiarch Tuning Operator	General Availability	General Availability	General Availability
Support for configuring the image stream import mode behavior	Not Available	Not Available	Technology Preview

1.7.9. Networking Technology Preview features

Table 1.26. Networking Technology Preview tracker

Feature	4.16	4.17	4.18
eBPF manager Operator	N/A	Technology Preview	Technology Preview
Advertise using L2 mode the MetalLB service from a subset of nodes, using a specific pool of IP addresses	Technology Preview	Technology Preview	Technology Preview
Updating the interface-specific safe sysctls list	Technology Preview	Technology Preview	Technology Preview
Egress service custom resource	Technology Preview	Technology Preview	Technology Preview
VRF specification in **BGPPeer** custom resource	Technology Preview	Technology Preview	Technology Preview
VRF specification in **NodeNetworkConfigurationPolicy** custom resource	Technology Preview	Technology Preview	Technology Preview
Host network settings for SR-IOV VFs	Technology Preview	General Availability	General Availability
Integration of MetalLB and FRR-K8s	Technology Preview	General Availability	General Availability

Feature	4.16	4.17	4.18
Automatic leap seconds handling for PTP grandmaster clocks	Not Available	General Availability	General Availability
PTP events REST API v2	Not Available	General Availability	General Availability
Customized **br-ex** bridge needed by OVN-Kubernetes to use NMState	General Availability	General Availability	General Availability
Live migration to OVN-Kubernetes from OpenShift SDN	Not Available	General Availability	Not Available
User defined network segmentation	Not Available	Technology Preview	General Availablity
Dynamic configuration manager	Not Available	Not Available	Technology Preview
SR-IOV Network Operator support for Intel C741 Emmitsburg Chipset	Not Available	Not Available	Technology Preview
SR-IOV Network Operator support on ARM architecture	Not Available	Not Available	General Availability

1.7.10. Node Technology Preview features

Table 1.27. Nodes Technology Preview tracker

Feature	4.16	4.17	4.18
MaxUnavailableStatefulSet featureset	Technology Preview	Technology Preview	Technology Preview
sigstore support	Not Available	Technology Preview	Technology Preview

1.7.11. OpenShift CLI (oc) Technology Preview features

Table 1.28. OpenShift CLI (oc) Technology Preview tracker

Feature	4.16	4.17	4.18
oc-mirror plugin v2	Technology Preview	Technology Preview	General Availability

Feature	4.16	4.17	4.18
oc-mirror plugin v2 enclave support	Technology Preview	Technology Preview	General Availability
oc-mirror plugin v2 delete functionality	Technology Preview	Technology Preview	General Availability

1.7.12. Extensions Technology Preview features

Table 1.29. Extensions Technology Preview tracker

Feature	4.16	4.17	4.18
Operator Lifecycle Manager (OLM) v1	Technology Preview	Technology Preview	General Availability
OLM v1 runtime validation of container images using sigstore signatures	Not Available	Not Available	Technology Preview

1.7.13. Operator lifecycle and development Technology Preview features

Table 1.30. Operator lifecycle and development Technology Preview tracker

Feature	4.16	4.17	4.18
Operator Lifecycle Manager (OLM) v1	Technology Preview	Technology Preview	General Availability
Scaffolding tools for Hybrid Helm-based Operator projects	Deprecated	Deprecated	Removed
Scaffolding tools for Java-based Operator projects	Deprecated	Deprecated	Removed

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

Table 1.31. RHOSP Technology Preview tracker

Feature	4.16	4.17	4.18
RHOSP integration into the Cluster CAPI Operator	Technology Preview	Technology Preview	Technology Preview
Control Plane with **rootVolumes** and **etcd** on local disk	Technology Preview	General Availability	General Availability

1.7.15. Scalability and performance Technology Preview features

Table 1.32. Scalability and performance Technology Preview tracker

Feature	4.16	4.17	4.18
factory-precaching-cli tool	Technology Preview	Technology Preview	Technology Preview
Hyperthreading-aware CPU manager policy	Technology Preview	Technology Preview	Technology Preview
Mount namespace encapsulation	Technology Preview	Technology Preview	Technology Preview
Node Observability Operator	Technology Preview	Technology Preview	Technology Preview
Increasing the etcd database size	Technology Preview	Technology Preview	Technology Preview
Using RHACM **PolicyGenerator** resources to manage GitOps ZTP cluster policies	Technology Preview	Technology Preview	Technology Preview

1.7.16. Storage Technology Preview features

Table 1.33. Storage Technology Preview tracker

Feature	4.16	4.17	4.18
AWS EFS storage CSI usage metrics	Not Available	General Availability	General Availability
Automatic device discovery and provisioning with Local Storage Operator	Technology Preview	Technology Preview	Technology Preview
Azure File CSI snapshot support	Not Available	Technology Preview	Technology Preview
Read Write Once Pod access mode	General Availability	General Availability	General Availability
Shared Resources CSI Driver in OpenShift Builds	Technology Preview	Technology Preview	Technology Preview
Secrets Store CSI Driver Operator	Technology Preview	Technology Preview	General Availability
CIFS/SMB CSI Driver Operator	Technology Preview	Technology Preview	General Availability

Feature	4.16	4.17	4.18
VMware vSphere multiple vCenter support	Not Available	Technology Preview	General Availability
Disabling/enabling storage on vSphere	Not Available	Technology Preview	Technology Preview
RWX/RWO SELinux Mount	Not Available	Developer Preview	Developer Preview
Migrating CNS Volumes Between Datastores	Not Available	Developer Preview	Developer Preview
CSI volume group snapshots	Not Available	Not Available	Technology Preview
GCP PD supports C3/N4 instance types and hyperdisk-balanced disks	Not Available	Not Available	General Availability
GCP Filestore supports Workload Identity	Not Available	General Availability	General Availability
OpenStack Manila support for CSI resize	Not Available	Not Available	General Availability

1.8. KNOWN ISSUES

- Previously, when you attempted to set the policy for a Google Cloud service account, the API reported a **400: Bad Request** validation error. When you create a service account, it might take up to 60 seconds for the account to become active, and this causes the validation error. If this error occurs, create a service account with a true exponential backoff that lasts at least 60 seconds. (OCPBUGS-48187)

- An installation can succeed when installing a cluster on a Google Cloud shared virtual private network (VPC) using the minimum permissions and without specifying the`controlPlane.platform.gcp.serviceAccount` in the **install-config.yaml** file. Firewall rules in Kubernetes (K8s) are created in the shared VPC, but destroying the cluster will not delete these firewall rules in K8s because the host project lacks the permissions. (OCPBUGS-38689)

- oc-mirror plugin v2 currently returns an exit status of **0**, meaning "success", even when mirroring errors occur. As a result, do not rely on the exit status in automated workflows. Until this issue is resolved, manually check the **mirroring_errors_XXX_XXX.txt** file generated by **oc-mirror** for errors. (OCPBUGS-49880)

- 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)

- There is a known issue in OpenShift Container Platform version 4.18 that prevents configuring multiple subnets in the failure domain of a Nutanix cluster during installation. There is no workaround for this issue. (OCPBUGS-49885)

- The following known issues exist for configuring multiple subnets for an existing Nutanix cluster by using a control plane machine set:

    - Adding subnets above the existing subnet in the **subnets** stanza causes a control plane node to become stuck in the **Deleting** state. As a workaround, only add subnets below the existing subnet in the **subnets** stanza.

    - Sometimes, after adding a subnet, the updated control plane machines appear in the Nutanix console but the OpenShift Container Platform cluster is unreachable. There is no workaround for this issue.

  These issues occur on clusters that use a control plane machine set to configure subnets regardless of whether subnets are specified in a failure domain or the provider specification. (OCPBUGS-50904)

- There is a known issue with RHEL 8 worker nodes that use **cgroupv1** Linux Control Groups (cgroup). The following is an example of the error message displayed for impacted nodes: **UDN are not supported on the node ip-10-0-51-120.us-east-2.compute.internal as it uses cgroup v1.** As a workaround, users should migrate worker nodes from **cgroupv1** to **cgroupv2**. (OCPBUGS-49933)

- The current PTP grandmaster clock (T-GM) implementation has a single National Marine Electronics Association (NMEA) sentence generator sourced from the GNSS without a backup NMEA sentence generator. If NMEA sentences are lost before reaching the e810 NIC, the T-GM cannot synchronize the devices in the network synchronization chain and the PTP Operator reports an error. A proposed fix is to report a **FREERUN** event when the NMEA string is lost. Until this limitation is addressed, T-GM does not support PTP clock holdover state. (OCPBUGS-19838)

- There is a known issue with a Layer 2 network topology on clusters running on Google Cloud Platform (GCP). At this time, the egress IP addresses being used in the Layer 2 network that is created by a **UserDefinedNetwork** (UDN) resource are using the wrong source IP address. Consequentially, UDN is not supported on Layer 2 on GCP. Currently, there is no fix for this issue. (OCPBUGS-48301)

- There is a known issue with user-defined networks (UDN) that causes OVN-Kubernetes to delete any routing table ID equal or higher to 1000 that it does not manage. Consequently, any Virtual Routing and Forwarding (VRF) instance created outside OVN-Kubernetes is deleted. This issue impacts users who have created user-defined VRFs with a table ID greater than or equal to 1000. As a workaround, users must change their VRFs to a table ID lower than 1000 as these are reserved for OpenShift Container Platform. (OCPBUGS-50855)

- If you attempted to log in to a OpenShift Container Platform 4.17 server by using the OpenShift CLI (**oc**) that you installed as part of the OpenShift Container Platform 4.18, you would see the following warning message in your terminal:

```
Warning: unknown field "metadata"
You don't have any projects. You can try to create a new project, by running

    oc new-project <projectname>
```

This warning message is a known issue but does not indicate any functionality issues with OpenShift Container Platform. You can safely ignore the warning message and continue to use OpenShift Container Platform as intended. (OCPBUGS-44833)

- There is a known issue in OpenShift Container Platform 4.18 which causes the cluster's masquerade subnet to be set to **169.254.169.0/29** if the **ovnkube-node** daemon set is deleted. When the masquerade subnet is set to **169.254.169.0/29**, **UserDefinedNetwork** custom resources (CRs) cannot be created.

  NOTE

  - If your masquerade subnet has been configured at Day 2 by making changes to the **network.operator** CR, it will not be reverted to **169.254.169.0/29**.
  - If a cluster has been upgraded from OpenShift Container Platform 4.16, the masquerade subnet remains **169.254.169.0/29** for backward compatibility. The masquerade subnet should be changed to a subnet with more IPs, for example, **169.254.0.0/17**, to use the user-defined networks feature.

  This known issue occurs after performing one of the following actions:

Action	Consequence
You have restarted the **ovnkube-node DaemonSet** object.	The masquerade subnet is set to **169.254.169.0/29**, which does not support **UserDefinedNetwork** CRs.
You have deleted the **ovnkube-node DaemonSet** object.	The masquerade subnet is set to **169.254.169.0/29**, which does not support **UserDefinedNetwork** CRs. Additionally, **ovnkube-node** pods crash and remain in a **CrashLoopBackOff** state.

  As a temporary workaround, you can delete the **UserDefinedNetwork** CR and then restart all **ovnkube-node** pods by running the following command:

  ```
  $ oc delete pod -l app=ovnkube-node -n openshift-ovn-kubernetes
  ```

  The **ovnkube-node** pods automatically restart, which re-stabilizes the cluster. Then, you can set the masquerade subnet to a larger IP address, for example, **169.254.0.0/17** for IPv4. As a result, **NetworkAttachmentDefinition** or **UserDefinedNetwork** CRs can be created.

  IMPORTANT

  Do not delete the **ovnkube-node DaemonSet** object when deleting **ovnkube-node** pods. Doing so sets the masquerade subnet to **169.254.169.0/29**.

  For more information, see Configuring the OVN-Kubernetes masquerade subnet as a Day 2 operation.

  (OCPBUGS-49662)

- Adding or removing nodes from the cluster can cause ownership contention over the node status. This can cause new nodes to take an extended period of time to appear. As a workaround, you can restart the **kube-apiserver-operator** pod in the **openshift-kube-apiserver-operator** namespace to expedite the process. (OCPBUGS-50587)

- For dual-stack networking clusters that run on RHOSP, when a Virtual IP (VIP) that is attached to a Floating IP (FIP) moves between master nodes, the association between VIP and FIP might stop working if the new master is on a different compute node. This issue occurs because OVN assumes that both IPv4 and IPv6 addresses on a shared Neutron port belong to the same node. (OCPBUGS-50599)

- Disk encryption with PCR 1 and PCR 7 protection fails on systems that automatically create additional Extensible Firmware Interface (EFI) entries during the boot operation. These extra entries modify EFI variables, preventing server attestation with PCR 1. (**OCPBUGS-54593**)

- 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)

- 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-49826)

- Due to an issue with Kubernetes, the CPU Manager is unable to return CPU resources from the last pod admitted to a node to the pool of available CPU resources. These resources are allocatable if a subsequent pod is admitted to the node. However, this pod then becomes the last pod, and again, the CPU manager cannot return this pod’s resources to the available pool. This issue affects CPU load-balancing features, which depend on the CPU Manager releasing CPUs to the available pool. Consequently, non-guaranteed pods might run with a reduced number of CPUs. As a workaround, schedule a pod with a **best-effort** Quality of Service (QOS) on the affected node. This pod will be the last admitted pod and this ensures the resources will be correctly released to the available pool. (OCPBUGS-46428)

- When a pod uses the CNI plugin for DHCP address assignment in conjunction with other CNI plugins, the network interface for the pod might be unexpectedly deleted. As a result, when the DHCP lease for the pod expires, the DHCP proxy enters a loop when trying to re-create a new lease, leading to the node becoming unresponsive. There is currently no workaround. (OCPBUGS-45272)

- When using PXE boot to add a worker node to an on-premise cluster , sometimes the host fails to reboot from the disk properly, preventing the installation from completing. As a workaround, you must manually reboot the failed host from the disk. (OCPBUGS-45116)

- The GCP PD CSI driver does not support hyperdisk-balanced volumes with RWX mode. Attempting to provision hyperdisk-balanced volumes with RWX mode using the GCP PD CSI driver produces errors and does not mount the volumes with the desired access mode. (OCPBUGS-44769)

- Currently, a GCP PD cluster with c3-standard-2, c3-standard-4, n4-standard-2, and n4-standard-4 nodes can erroneously exceed the maximum attachable disk number, which should be 16. This issue may prevent you from successfully creating or attaching volumes to your pods. (OCPBUGS-39258)

1.9. ASYNCHRONOUS ERRATA UPDATES

Security, bug fix, and enhancement updates for OpenShift Container Platform 4.18 are released as asynchronous errata through the Red Hat Network. All OpenShift Container Platform 4.18 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 4.18. Versioned asynchronous releases, for example with the form OpenShift Container Platform 4.18.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-2026:0338 - OpenShift Container Platform 4.18.31 fixed issues and security update

Issued: 14 January 2026

OpenShift Container Platform release 4.18.31 is now available. The list of bug fixes that are included in the update is documented in the RHSA-2026:0338 advisory. The RPM packages that are included in the update are provided by the RHSA-2026:0331 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 4.18.31 --pullspecs
```

1.9.2. Fixed issues

You can ensure your environment remains stable by reviewing these fixed issues to verify resolved bugs.

The following issues are fixed for this release:

- Before this update, the installation program created images that did not support instance types requiring NVMe or SCSI. This caused installations to fail when using those specific instance types. With this release, the installation program now supports NVMe and SCSI instance types. You can now use these instances for your OpenShift deployments. (OCPBUGS-52659)

- Before this update, a null pointer dereference occurred in the **waitFor** method due to failed Google Cloud (GCP) API calls or uninitialized clients. As a consequence, users experienced a panic error during cluster destruction on GCP. With this release, the null pointer dereference in the Cluster Uninstaller utility is fixed, preventing panic errors during OpenShift Container Platform cluster destruction on GCP. As a result, cluster destruction on GCP no longer triggers a panic error with the updated OpenShift Container Platform 4.18 release. (OCPBUGS-63494)

- Before this update, the **EgressFirewall** resource retained managed fields after a machine was deleted. As a consequence, large clusters experienced etcd storage buildup, which increased the risk of API server instability or failure. With this release, the installation program does not retain information for deleted machines, which prevents unnecessary etcd growth. As a result, you can manage large clusters in OpenShift Container Platform 4.18.0 without risking API server breakdowns due to stale **EgressFirewall** data. (OCPBUGS-66140)

- Before this update, a validation error occurred during the creation of user accounts with special characters in the email address. As a consequence, user account creation failed. With this release, the validation error is resolved. As a result, user account creation with special characters in the email address is successful. (OCPBUGS-66381)

- Before this update, the Ironic API advertised an unreachable IP address even when a routable path existed. As a consequence, users could not access the Ironic API because the advertised IP address was non-routable. With this release, the installation program verifies the routability of the advertised IP address for the Ironic API. As a result, the Ironic API is consistently reachable, ensuring a reliable user experience in OpenShift Container Platform 4.18.0. (OCPBUGS-68360)

- Before this update, **HAProxy** pods frequently exceeded the **maxConn** limit because idle connections remained open when using the **idle-close-on-response** setting. As a consequence, **HAProxy** pods entered a **CrashLoop** state, resulting in performance degradation and persistent reconciliation failures. With this release, the **IdleConnectionTerminationPolicy** is activated to ensure idle connections are closed correctly. As a result, **HAProxy** pods in OpenShift Container Platform 4.18.0 no longer exceed the **maxConn** limit or experience crashes related to connection buildup. (OCPBUGS-70315)

1.9.3. Updating

To update an OpenShift Container Platform 4.18 cluster to this latest release, see Updating a cluster using the CLI.

1.9.4. RHBA-2025:22696 - OpenShift Container Platform 4.18.30 bug fix and security update

Issued: 10 December 2025

OpenShift Container Platform release 4.18.30, which includes security updates, is now available. The list of bug fixes that are included in the update is documented in the RHBA-2025:22696 advisory. The RPM packages that are included in the update are provided by the RHBA-2025:22694 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 4.18.30 --pullspecs
```

1.9.5. Enhancements

This release contains the following enhancements:

- The KubeVirt Container Storage Interface (CSI) driver now supports volume expansion. This means that users can now dynamically increase the size of their persistent volumes in their tenant cluster. This capability greatly simplifies storage management, allowing for a more flexible and scalable infrastructure. (OCPBUGS-61700)

1.9.6. Bug fixes

The following bugs are fixed with this release:

- Before this update, the **RevisionController** process in core Operators prevented upgrades to new revisions during a cluster upgrade if the **operator.openshift.io/revision-ready** annotation was missing on the latest available revision ConfigMap. With this release, the **RevisionController** process was updated to upgrade to a new revision even if the **operator.openshift.io/revision-ready** annotation is missing. As a result, core Operators can successfully transition to new revisions regardless of the annotation’s presence. (OCPBUGS-58412)

- Before this update, if NetworkManager was restarted or crashed on a node with a **br-ex** interface managed by NMState, the node lost network connectivity. With this release, a fallback check in the dispatcher script was added to detect NMState-managed **br-ex** interfaces by checking for the **br-ex-br** bridge ID when the standard **br-ex** bridge ID is not found. As a result, nodes with this interface type do not lose network connectivity when NetworkManager restarts or crashes. (OCPBUGS-62169)

- Before this update, creating a network bond could hide the original MAC addresses of some interfaces, which caused provisioning to fail. With this release, the **ironic-python-agent** now discovers all MAC addresses within a bond, ensuring nodes can be correctly identified and provisioned no matter which MAC address is used for the **BareMetalHost bootMACAddress** value. (OCPBUGS-62456)

- Before this update, the Cloud Credential Operator utility (**ccoctl**) 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-65479)

- Before this update, during the mirror operation, **oc-mirror** inadvertently set the executable flag on some synchronized files that did not contain executable code or scripts, potentially causing unexpected execution. With this release, unintended executable flags have been removed from the synchronized files. As a result, correct file permissions are set, preventing unintended execution of synchronized files. (OCPBUGS-65510)

- Before this update, a race condition in the Redfish Power interface caused power operations to fail during simultaneous access. As a consequence, users were unable to manage power states reliably. With this release, the race condition in the Redfish Power interface has been resolved, ensuring successful power operations. As a result, users can now manage power states reliably. (OCPBUGS-65573)

- 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 1 update domain when the fault domain count is 1. With this release, the logic was updated to explicitly set the **platformUpdateDomainCount** to 1 whenever the **platformFaultDomainCount** is determined to be 1. As a result, Machine Availability Sets are created with valid parameter combinations, allowing machines to successfully provision in Azure regions with a single fault domain. (OCPBUGS-65882)

- Before this update, a rounding error caused metrics charts to display gaps even when the underlying data was continuous. With this release, the rounding error has been corrected, ensuring that charts display without gaps when the data is complete. (OCPBUGS-65958)

- Before this update, a regression in **runc** prevented pods which had the **shareProcessNamespace** object set to **true** from running properly. With this release, that 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-65976)

1.9.7. Updating

To update an OpenShift Container Platform 4.18 cluster to this latest release, see Updating a cluster using the CLI.

1.9.8. RHBA-2025:21797 - OpenShift Container Platform 4.18.29 bug fix advisory

Issued: 26 November 2025

OpenShift Container Platform release 4.18.29 is now available. The list of bug fixes that are included in the update is documented in the RHBA-2025:21797 advisory. The RPM packages that are included in the update are provided by the RHBA-2025:21794 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 4.18.29--pullspecs
```

1.9.9. Bug fixes

The following bugs are fixed with this release:

- Before this update, when a new secret was created or updated in any namespace, Alertmanager was reconciling even if that secret was not referenced in the **AlertmanagerConfig** resource. As a consequence, the Prometheus Operator generated excessive API calls, causing increased CPU usage on control plane nodes. With this release, Alertmanager only reconciles secrets that the **AlertmanagerConfig** resource explicitly references. (OCPBUGS-61965)

- Before this release, the **ccoctl** utility would automatically generate a new keypair if the private key was not found, even when users intentionally provided only the public key according to documented security procedures. This behavior caused a problem, as the newly generated keys would not match the cluster’s keys, resulting in service outages for users following the correct process. With this update, the utility was changed to ensure a new keypair is never generated when the **--public-key-file** parameter is specified, and this parameter was added to all create-all functions for consistency. As a result, specifying the public key file now guarantees the provided key is used, ensuring the cluster continues to function as expected without interruption. (OCPBUGS-63548)

- 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-64579)

- 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-64608)

- Before this update, the node log length was unlimited. As a consequence, an extremely large log could prevent the display of the log or cause the browser to crash. With this release, the node log length is limited to 1,000 lines. As a result, the log displays correctly. (OCPBUGS-64682)

- 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-65533)

1.9.10. Updating

To update an OpenShift Container Platform 4.18 cluster to this latest release, see Updating a cluster using the CLI.

1.9.11. RHBA-2025:19865 - OpenShift Container Platform 4.18.28 bug fix and security update

Issued: 12 November 2025

OpenShift Container Platform release 4.18.28 is now available. The list of bug fixes that are included in the update is documented in the RHBA-2025:19865 advisory. The RPM packages that are included in the update are provided by the RHBA-2025:19863 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 4.18.28 --pullspecs
```

1.9.11.1. Enhancements

- With this update, the **remoteWrite[].oauth2.proxyFromEnvironment** setting can now be used to configure a cluster-wide proxy in 4.18.z. This improvement backports a feature previously available only in 4.19 and later builds, allowing for more flexible and consistent proxy configurations. (OCPBUGS-63410)

1.9.11.2. Bug fixes

- 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** configuration file, the installation fails with an error, prompting you to provide them. (OCPBUGS-53235)

- 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-61773)

- Before this update, Redfish transactions in some hardware models would fail due to the Baseboard Management Controller (BMC) sending an empty ETag. As a consequence, users could not use the **HostFirmwareSettings** custom resource (CR). With this release, the Redfish transaction with empty ETag issue has been resolved and returns the correct **ETag instances without warnings. As a result, the Redfish transaction no longer fails, allowing users to use the `HostFirmwareSettings** CR. (OCPBUGS-62647)

- Before this update, inconsistent updates to the driver-config ConfigMap in the hosted cluster namespace caused the driver-config ConfigMap content to flap, resulting in inconsistent storage class enforcement and affecting user experience. With this release, the driver-config ConfigMap stability has been restored, preventing the flapping of storage classes in the hosted cluster namespace. (OCPBUGS-62808)

- Before this update, the controller created and deleted a file with a random name when setting up a session to Amazon Web Services (AWS), which caused the controller to continuously allocate more memory to cache the session. With this release, the controller now uses the same file name instead of a random one, allowing the kernel to re-use the **dentry** instead of requesting a new one for each session. As a result, excessive memory allocation is resolved. (OCPBUGS-63138)

- 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. As a result, the logs no longer overflow, as these specific messages are now less verbose by default. (OCPBUGS-63324)

- Before this update, when a user 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-63450)

- 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 and ensuring greater network stability. (OCPBUGS-63459)

- Before this update, the Azure machine provider was not passing the **dataDisks** configuration from the compute machine set 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-63669)

1.9.11.3. Updating

To update an OpenShift Container Platform 4.18 cluster to this latest release, see Updating a cluster using the CLI.

1.9.12. RHSA-2025:19047 - OpenShift Container Platform 4.18.27 bug fix and security update

Issued: 29 October 2025

OpenShift Container Platform release 4.18.27 is now available. The list of bug fixes that are included in the update is documented in the RHSA-2025:19047 advisory. The RPM packages that are included in the update are provided by the RHBA-2025:19045 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 4.18.27 --pullspecs
```

1.9.12.1. Bug fixes

- Before this update, if the OVN-Kubernetes controller was not processing updates from the Kubernetes API server and configuring the open virtual network (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 even though that IP might have moved to another node. With this release, these GARPs are blocked when the OVN-Kubernetes controller is not processing updates. (OCPBUGS-62671)

- Before this update, the Cluster Version Operator (CVO) in 4.19.9 and 4.18.23 started to require bearer token authentication in metrics requests. As a consequence, hosted control planes clusters were broken because the metrics scraper did not provide client authentication. With this release, the CVO does not require client authentication for metrics requests. As a result, access to CVO metrics scraping is recovered on hosted control planes clusters. (OCPBUGS-62869)

- 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-63041)

- Before this update, users without a project saw only part of the **Roles** list because of insufficient role-based access control (RBAC) permissions. With this release, the access logic is fixed. As a result, these users cannot open the **Roles** page, which keeps sensitive data secure. (OCPBUGS-63247)

- Before this update, during an update from 4.18.21 to 4.19.6, 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-63346)

1.9.12.2. Updating

To update an OpenShift Container Platform 4.18 cluster to this latest release, see Updating a cluster using the CLI.

1.9.13. RHSA-2025:17657 - OpenShift Container Platform 4.18.26 bug fix and security update

Issued: 15 October 2025

OpenShift Container Platform release 4.18.26 is now available. The list of bug fixes that are included in the update is documented in the RHSA-2025:17657 advisory. The RPM packages that are included in the update are provided by the RHBA-2025:17655 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 4.18.26 --pullspecs
```

1.9.13.1. Bug fixes

- Before this update, scaling large numbers of nodes was slow because each machine was reconciled sequentially when the queue was large. With this release, up to 10 machines are reconciled concurrently, improving the speed of scale events. (OCPBUGS-59387)

- Before this update, pods with secondary interfaces on an OVN-Kubernetes Localnet network could communicate with other pods on the same node only if their Localnet IP addresses were within the same subnet as the host network. With this release, Localnet IP addresses can be drawn from any subnet. In this case, an external router connects the Localnet subnet to the host network. (OCPBUGS-61455)

- Before this update, stale IP addresses in the **Address_Set** for DNS Egress Firewall rules were not removed, causing memory leaks. With this release, IP addresses are removed from the **Address_Set** five seconds after their TTL expires, preventing memory growth. (OCPBUGS-61748)

- Before this update, the **MIRRORED_RELEASE_IMAGE** environment variable for the ignition server could fluctuate due to race conditions, causing unnecessary pod restarts. With this release, the **MIRRORED_RELEASE_IMAGE** values are consistent, stabilizing ignition server deployments. (OCPBUGS-61904)

- Before this update, control plane nodes created before OpenShift Container Platform version 4.12 did not include 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. (OCPBUGS-62321)

- Before this update, visiting the **/auth/error** page could display an improper error. With this release, the page renders a proper error message. (OCPBUGS-62326)

- Before this update, omitting the **--v1** or **--v2** flags when using oc-mirror could lead to inconsistent behavior. With this release, specifying **--v1** or **--v2** is mandatory. (OCPBUGS-62432)

- Before this update, resizing a PVC immediately after creation could fail because the PV was temporarily not found. With this release, you can resize PVCs immediately after creation without errors. (OCPBUGS-62467)

- Before this update, the Machine Config Operator (MCO) did not search the **/etc/docker/certs.d** directory for required certificates, causing the Operator Controller and catalogd to fail. With this release, the MCO can access certificates in this directory, and the Operator Controller and catalogd now start successfully. (OCPBUGS-54175)

- Before this update, an external actor could uncordon a node while the Machine Config Operator (MCO) was draining it, causing the MCO and scheduler to compete in scheduling and removing pods. With this release, the MCO detects external uncordons and continues the drain process without interference. (OCPBUGS-62637)

- Before this update, you had to provide the oc-mirror version for every bug report, which could delay issue resolution. With this release, oc-mirror v2 outputs its version in both the standard output and log, making troubleshooting faster. (OCPBUGS-62696)

1.9.13.2. Updating

To update an OpenShift Container Platform 4.18 cluster to this latest release, see Updating a cluster using the CLI.

1.9.14. RHBA-2025:16732 - OpenShift Container Platform 4.18.25 bug fix update

Issued: 01 October 2025

OpenShift Container Platform release 4.18.25 is now available. The list of bug fixes that are included in the update is documented in the RHBA-2025:16732 advisory. The RPM packages that are included in the update are provided by the RHSA-2025:16729 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 4.18.25 --pullspecs
```

1.9.14.1. Enhancements

- With this update, OpenShift Container Platform 4.18.25 is upgraded to Kubernetes 1.31.12, ensuring a secure, stable, and up-to-date platform. This enhancement improves stability, enhances security, and boosts performance by incorporating the latest upstream changes and fixes. (OCPBUGS-60511)

- With this update, the collection of command line logs from **virt-launcher** pods across an OpenShift Container Platform cluster are enabled. JSON-encoded logs are saved in the **namespaces/<namespace_name>/pods/<pod_name>/virt-launcher.json** path, facilitating troubleshooting and debugging of virtual machines. (OCPBUGS-61656)

1.9.14.2. Bug fixes

- Before this update, the **agent-based-installer** set the permissions for the **/var/lib/etcd/member** etcd directory as 0755 when using single-node OpenShift deployment instead of 0700, which is correctly set on a multi-node deployment. With this release, the **/var/lib/etcd/member** etcd directory permissions are set to 0700 for single-node OpenShift deployments. (OCPBUGS-61529)

- Before this update, the **PrometheusRemoteWriteBehind** alert fired if the remote endpoint has not received any data. With this release, the **PrometheusRemoteWriteBehind** alert no longer fires if the remote endpoint has not yet received any data. (OCPBUGS-61706)

- Before this update, an NMState service failure occurred in OpenShift Container Platform deployments due to a **NetworkManager-wait-online** dependency in bare metal deployments. As a consequence, end users experienced deployment failures due to a network misconfiguration. With this release, the NMState service dependency issue has been resolved, and **NetworkManager-wait-online** is no longer required for **br-ex** configuration. (OCPBUGS-61840)

1.9.14.3. Updating

To update an OpenShift Container Platform 4.18 cluster to this latest release, see Updating a cluster using the CLI.

1.9.15. RHBA-2025:15714 - OpenShift Container Platform 4.18.24 bug fix update

Issued: 17 September 2025

OpenShift Container Platform release 4.18.24 is now available. The list of bug fixes that are included in the update is documented in the RHBA-2025:15714 advisory. The RPM packages that are included in the update are provided by the RHBA-2025:15712 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 4.18.24 --pullspecs
```

1.9.15.1. Enhancements

- With this update, the **cluster-etcd-operator** in OpenShift Container Platform platforms is proactive, modifying the **etcdDatabaseQuotaLowSpace alert**. It triggers at many levels, for example, information, warning, critical, when the etcd quota usage approaches 95%. This enhancement provides cluster administrators with time to address potential issues before the API server is impacted, resulting in a more stable and manageable environment. (OCPBUGS-61235)

- The name of the **MachineOSConfig** object that was used with Red Hat Enterprise Linux CoreOS (RHCOS) image layering must now be the same as the machine configuration pool where the custom layered image is deployed. Previously, you could use any name. This change was made to prevent attempts to use many **MachineOSConfig** objects with each machine configuration pool. For more information, see Image mode for OpenShift.

- With this update, the system ensures balanced distribution of live API connections across many Kubernetes API servers when primary nodes are online. This prevents any single primary node from reaching 100% CPU use, improving system performance and stability during primary node restarts or quorum establishment. This change maintains optimal performance and stability of the system by evenly distributing live API connections and preventing a single Kubernetes API server from handling the majority of connections during primary node or Kubernetes API server outages. (OCPBUGS-61039)

1.9.15.2. Known issues

- A Berkley Packet Filter (BPF) program creation failure due to a denied permission in the container runtime version 1.23 on OpenShift Container Platform 4.18.22 caused a failed Graphics Processing Unit (GPU) stack deployment. To correct this failure, the **no-cgroups = true** variable was added to the GPU worker node **/var/usrlocal/nvidia/toolkit/.config/nvidia-container-runtime/config.toml** file. As a result, the GPU stack deploys successfully on OpenShift Container Platform 4.18.22. (OCPBUGS-60663)

1.9.15.3. Bug fixes

- Before this update, LDAP integration created excessive health check alerts due to many config maps for each user. As a consequence, the user interface was cluttered with excessive alerts. With this release, excessive health check alerts from the user config maps are removed. As a result, users experience reduced console alerts and improved usability. (OCPBUGS-50983)

- Before this update, an incorrect cloud configuration path in Microsoft Azure Stack Hub caused the Container Storage Interface (CSI) driver to use an incorrect environment configuration. As a consequence, users experienced unhealthy CSI driver pods. With this release, the incorrect environment configuration for the CSI operator on Azure Stack Hub is fixed and the CSI driver correctly reads the cloud configuration. (OCPBUGS-55053)

- Before this update, multiple mirrors in the **imagecontentsourcepolicy** policy caused a Hosted control planes payload error during an image lookup. As a consequence, the Hosted control planes payload failed to handle multiple mirrors, and caused cluster creation failures. With this release, the Hosted control planes payload supports multiple mirrors and avoids creation errors. (OCPBUGS-57142)

- Before this update, bonded network configurations with **mode=active-backup** and **fail_over_mac=follow** variables failed because of a race condition in the **configure-ovs.sh** file. This condition caused interfaces to change to an unpredictable state. As a consequence, bonded network flapping occurred and affected high availability. With this release, the race condition that handles the **configure-ovs.sh** file is improved, and ensures that bonded network configurations with the **fail_over_mac=follow** variable functions without flapping. (OCPBUGS-57357)

- Before this update, an incorrect URL in the **Alertmanager** configuration created logs in **user-workload** pods on OpenShift Container Platform 4.16 clusters. As a consequence, user workload monitoring failed due to an invalid Slack URL in the configuration. With this release, invalid URLs are allowed in the **Alertmanager** configuration. As a result, the **Alertmanager** configuration errors are resolved, which improves monitoring stability on clusters earlier than OpenShift Container Platform 4.19. (OCPBUGS-58194)

- Before this update, multiple IP addresses on the primary interface caused the API server to connect to an unmatched etcd IP address during a single node OpenShift Container Platform 4.18 deployment. This action resulted in an API server pod failure during deployment and cluster initialization issues. With this release, a deployment with multiple IP addresses on the primary interface is fixed, ensuring the API server connects to etcd using the correct IP address in the certificate, and prevents a failure during a single node deployment. (OCPBUGS-59992)

- Before this update, insufficient obfuscation of new network data types in Hosted control planes exposed user data. As a consequence, sensitive information was not protected. With this release, obfuscation for new network data types is implemented. As a result, obfuscated network data in Hosted control planes improves data privacy. (OCPBUGS-60301)

- Before this update, a hosted cluster rejected certificates with multiple Storage Area Network (SAN) entries due to conflicting Domain Name System (DNS) names. As a consequence, users encountered invalid configuration errors after deploying hosted clusters with the certificates. With this release, the cluster accepts certificates with multiple SAN entries. As a result, deployment errors are eliminated. (OCPBUGS-60485)

- Before this update, the Cluster Network Operator (CNO) missed setting the **release.openshift.io/version** annotation on a **DaemonSet** API object and in the **openshift-multus** namespace during an upgrade from OpenShift Container Platform 4.18.6 to OpenShift Container Platform 4.18.22. As a consequence, the upgrade failed due to missing annotations. With this release, the CNO sets the annotation for the **DaemonSet** object and for the deployment during an upgrade. As a result, upgrading a cluster completes without stopping due to missing annotations in the **openshift-multus** namespace. (OCPBUGS-60795)

- Before this update, improper machine deletion handling during cluster autoscaling caused the last node to retain the **ToBeDeletedByClusterAutoscaler** taint. As a consequence, resource allocation was affected during cluster scaling. With this release, the **ToBeDeletedByClusterAutoscaler** taint removal after scaling down a machine set is implemented. As a result, the last node does not retain the unwanted taint after scaling down a machine set. (OCPBUGS-60908)

- Before this update, disabling the OpenShift Container Platform image registry left legacy pull secrets with persistent finalizers, which prevented their deletion. As a consequence, users were unable to delete **Dockercfg** secrets after disabling the image registry. With this release, legacy pull secrets do not block namespace deletion after the registry is removed. As a result, pull secret finalizers do not block namespace deletion when the registry is disabled. (OCPBUGS-61002)

- Before this update, IBM Cloud was omitted from the list of platforms that supported single-node installations in the validation code. As a consequence, users could not install a single-node configuration on IBM Cloud because of a validation error. With this release, IBM Cloud support for single-node installations is enabled. As a result, users can complete single-node installations on IBM Cloud. (OCPBUGS-61178)

1.9.15.4. Updating

To update an OpenShift Container Platform 4.18 cluster to this latest release, see Updating a cluster using the CLI.

1.9.16. RHSA-2025:14820 - OpenShift Container Platform 4.18.23 bug fix and security update

Issued: 03 September 2025

OpenShift Container Platform release 4.18.23 is now available. The list of bug fixes that are included in the update is documented in the RHSA-2025:14820 advisory. The RPM packages that are included in the update are provided by the RHBA-2025:14816 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 4.18.23 --pullspecs
```

1.9.16.1. Enhancements

- In OpenShift Container Platform 4.18.23, Operator Lifecycle Manager (OLM) Classic and OLM v1 allow Operators to include network policy manifests in their resource bundles. These tailored network policies protect against data leaks and harden against many attack vectors on OpenShift Container Platform clusters.

  #### TIP

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

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

  Update to OpenShift Container Platform 4.18.23 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)

- 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 wired to the 'oc delete istag' command. As a result, the accidental deletion of image objects is prevented and the **istag** object remains intact when using the **--dry-run=server** option. (OCPBUGS-58461)

- Before this update, the **cluster-policy-controller** container was exposing the **10357** port for all networks (the bind address was set to **0.0.0.0**). The port was exposed outside the host network for the node because the KCM pod manifest set **hostNetwork** to **true**. This port is used solely for the probe of the container. With this enhancement, the bind address was updated to listen on the **localhost** only. As result, the node security is improved because the port is not exposed outside the node network. (OCPBUGS-60131)

1.9.16.2. Bug fixes

- 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 will gather the information correctly. (OCPBUGS-46437)

- Before this update, certain traffic patterns with large packets running between OpenShift Container Platform nodes and pods triggered an OpenShift Container Platform host to send Internet Control Message Protocol (ICMP) needs frag to another OpenShift Container Platform host. This situation lowered the viable maximum transmission unit (MTU) in the cluster. As a consequence, executing the **ip route show cache** command displayed a cached route with a lower MTU than the physical link. Packets were dropped and OpenShift Container Platform components were degrading because the host did not send pod-to-pod traffic with the large packets. With this release, NF Tables rules prevent the OpenShift Container Platform nodes from lowering their MTU in response to these traffic patterns. (OCPBUGS-58288)

- Before this update, when a Cluster Operator takes a long time to upgrade, Cluster Version Operator does not report anything as it cannot determine if the upgrade is still progressing or already stuck. With this release, a new unknown status is added for the failing condition in status of the Cluster Version reported by Cluster Version Operator to remind the cluster administrators to check the cluster and avoid waiting on a blocked Cluster Operator upgrade. (OCPBUGS-58450)

- Before this update, intermittent egress IP handling due to inconsistent state updates caused packet drops in user traffic. With this release, egress IP handling consistency has been improved. (OCPBUGS-59371)

- Before this update, SELinux status check were not performed before running SELinux related commands. As a consequence, if the target RHEL machine did not have SELinux enabled, the related steps failed. With this release, the SELinux status check is added to ensure that the steps are successfully executed. (OCPBUGS-59844)

- Before this update, downloads on control plane nodes were inconsistently scheduled because of a mismatch between the node selector for downloads and the console pods. As a consequence, downloads were scheduled on random nodes, which caused potential resource contention and sub-optimal performance. With this release, downloaded workloads consistently schedule on control plane nodes, which improves resource allocation. (OCPBUGS-59897)

- Before this update, the delete workflow erroneously displayed **workflow mode: diskToMirror / delete**, leading to user confusion regarding the correct workflow mode. With this release, **workflow mode: delete** displays during delete operations. (OCPBUGS-59966)

- Before this update, the **netavark** package did not install because it did not contain the 4.18 RHEL8 repositories. With this release, the **netavark** package is succesfully installed from the **container-tools** module. (OCPBUGS-59973)

- Before this update, 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 an incorrect **clockclass** and instead returns an HTTP **400 Bad Request** or **404 Not Found Error**. (OCPBUGS-59984)

- Before this update, OCP updates that shipped a change to **coredns** templates would restart the static pod and pre-reboot the node. This issue occurred before the image pull for the base operation system image update. As a consequence, a race occurred where the **rpm-ostree**, the operating system (OS) update manager, failed the image pull because of network errors and stall. With this release, a retry in the Machine Config Operator (MCO) OS update operation is added to work around the race condition due to the restarts of the **coredns** pod. (OCPBUGS-60034)

- Before this update, when you used multiple recommenders for the Vertical Pod Autoscaler (VPA), the default VPA recommender would erroneously garbage collect **VPACheckpoint** objects that belonged to a VPA object which was associated with a non-default recommender. With this release, the default recommender does not garbage collect the checkpoints owned by the other recommenders. (OCPBUGS-60235)

- Before this update, the **Events** page incorrectly displayed **{error}** instead of an error message. With this release, the error message is displayed. (OCPBUGS-60278)

- Before this update , on dual-stack clusters with IPv6 as the primary networking stack, the bare-metal Installer-Provisioned Infrastructure (IP) would incorrectly provide an IPv4 URL for the virtual media ISO image. This caused installation failures on Baseboard Management Controllers (BMCs) that were configured exclusively for IPv6 networking, as they could not reach the IPv4 address. With this release, the installer logic was updated to always provide an IPv6 URL when a BMC is using IPv6 addressing and the installation process now completes successfully. (OCPBUGS-60592)

1.9.16.3. Updating

To update an OpenShift Container Platform 4.18 cluster to this latest release, see Updating a cluster using the CLI.

1.9.17. RHSA-2025:13325 - OpenShift Container Platform 4.18.22 bug fix and security update

Issued: 13 August 2025

OpenShift Container Platform release 4.18.22 is now available. The list of bug fixes that are included in the update is documented in the RHSA-2025:13325 advisory. The RPM packages that are included in the update are provided by the RHBA-2025:13326 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 4.18.22 --pullspecs
```

1.9.17.1. Enhancements

- The readiness probes for the API server (/**readyz** endpoint) have been modified to exclude etcd checks. This modification prevents client connections from being closed if etcd is temporarily unavailable. As a result, etcd will be ready again before a client connection times out, enabling the client connections to persist through a brief etcd unavailability. This persistence minimizes temporary API server outages. (OCPBUGS-49749)

1.9.17.2. Known issues

- Stale Source Network Address Translations (SNATs) or routing policies might occur in the following circumstances:
  - You are upgrading from 4.17 to 4.18 during an update of the OVN-Kubernetes image.
  - During the upgrade, if a pod on another system that was selected by an egress IP was deleted when the **ovnkube-node** pod was not running.

(OCPBUGS-59531)

1.9.17.3. Bug fixes

- Before this update, destroying a cluster in the unsupported region **mx-central-1** caused the destroyer to fail to find a partition and not exit. As a consequence, you could not destroy an OpenShift Container Platform cluster in the **mx-central-1** region due to constant error reporting. With this release, the destroyer does not report errors for the unsupported region **mx-central-1** which enables the successful destruction of a cluster. (OCPBUGS-56177)

- Before this update, combined specification and status updates lists triggered unnecessary firmware upgrades, which caused system downtime. With this release, a firmware upgrade optimization skips unnecessary firmware upgrades. (OCPBUGS-56766)

- Before this update, the **console-telemetry** plugin received a **Forbidden** error due to using the wrong API endpoint for tracking usage. As a consequence, the **Forbidden** console-telemetry-plugin usage tracking error occurred. With this release, the **console-telemetry** plugin posts usage data to **/api/metrics/usage** instead of **/metrics/usage**. As a result, the **console-telemetry** plugin does not receive a **Forbidden** error, which ensures accurate usage tracking. (OCPBUGS-58364)

- Before this update, the installation program failed when Amazon Web Services (AWS) credentials were not found and the survey was attempting to list all AWS regions preventing users from creating the **install-config** YAML file. With this release, the installation program no longer fails when AWS credentials are not set, allow users to input them during the survey. (OCPBUGS-59155)

- 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-59503)

- Before this update, the **oc-mirror** did not detect Helm Chart images that used an aliased sub-chart. As a consequence, the Helm Chart images were missing after mirroring. With this release, the **oc-mirror** detects and mirrors Helm Chart images with an aliased sub-chart. (OCPBUGS-59798)

- Before this update, **netavark** could not be downloaded from the **container-tools** module. With this release, the **container-tools** module is enabled for **netvark**. As a result,**netavark** can be downloaded from the module. (OCPBUGS-59843)

- Before this update, when you cloned a TAR file with zero length, the **oc-mirror** ran indefinitely due to an empty archive file. As a consequence, no progress occurred when you mirrored a 0-byte TAR file. With this release, 0-byte TAR files are detected and reported as errors, which prevents the **oc-mirror** from hanging. (OCPBUGS-59864)

- Before this update, in multi-zone clusters with only a single compute node per zone, if the Monitoring Operator’s Prometheus pods were scheduled to nodes that reboot back-to-back and both reboots took longer than 15 minutes to return to service, the Monitoring Operator might have degraded. With this release, the time-out has been extended to 20 minutes to prevent the Monitoring Operator from entering a degraded state on common cluster topologies. Clusters where the two nodes with Prometheus pods reboot back-to-back and take more than 20 minutes might still report a degraded state until the second node and Prometheus pod return to a normal state.(OCPBUGS-59962)

1.9.17.4. Updating

To update an OpenShift Container Platform 4.18 cluster to this latest release, see Updating a cluster using the CLI.

1.9.18. RHSA-2025:11677 - OpenShift Container Platform 4.18.21 bug fix and security update

Issued: 30 July 2025

OpenShift Container Platform release 4.18.21 is now available. The list of bug fixes that are included in the update is documented in the RHSA-2025:11677 advisory. The RPM packages that are included in the update are provided by the RHSA-2025:11678 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 4.18.21 --pullspecs
```

1.9.18.1. Bug fixes

- Before this update, the **Observe** > **Metrics** > **query** > **QueryKebab** > **Export as csv**drop-down item did not handle a undefined title element. As a consequence, you were unable to export the CSV file for certain queries on the **Metrics** tab of OpenShift Lister versions 4.16, 4.17, and 4.18. With this release, the metrics download for all queries correctly handles object properties in the drop-down menu items. As a result, the CSV export for all queries works on the **Metrics** page. (OCPBUGS-54314)

- Before this update, on outdated version of the Microsoft Azure API prevented specifying 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 server's creation point. (OCPBUGS-56167)

- Before this update, oc-mirror v2 did not use the custom credentials supplied by the **--authfile** parameter for the creation of a graph image. As a consequnce, authentication failures prevented the graph image creation. With this release, the **--authfile** parameter correctly uses the custom credentials for both mirroring and graph image creation. As a result, no authorization failures occur during the graph image creation. (OCPBUGS-57068)

- Before this update, when on-prem installer-provisioned infrastructure (IPI) deployments used the Cilium container network interface (CNI), the firewall rule that redirected traffic to the load balancer was ineffective. With this release, the rule works with the Cilium CNI and **OVNKubernetes**. (OCPBUGS-57782)

- Before this update, the build controller looked for secrets that were linked for general use, not specifically for the image pull. With this release, when searching for default image pull secrets, the builds use **ImagePullSecrets** that are linked to the service account. (OCPBUGS-57949)

- Before this update, **oc-mirror v2** sent a large number of requests to container registries. As a consequence, container registries rejected some requests with a **too many requests** error message. With this release, the default for the **maxParallelLayerDownloads** and

maxParallelImageDownloads flags is reduced and the retry-times are increased. The retry-delay is also set to **0** to enable exp backoff. As a result, fewer requests are sent to container registries which results in fewer rejections of the requests. (OCPBUGS-58280)

- Before this update, a regression in OpenShift Container Platform 4.15 caused forward slashes to not work in the 'href' values for the **console.tab/horizontalNav`parameter. With this release, forward slashes in the `href** values for the **console.tab/horizontalNav** parameter work as expected. (OCPBUGS-58457)

- Before this update, when you ran the oc-mirror v2 disk-to-mirror workflow without valid mirror tar files, the returned error messages did not correctly identify the problem. With this release, the oc-mirror v2 workflow returns an error message that states **no tar archives matching "mirror_[0-9]{6}\.tar" found in "<directory>"**. (OCPBUGS-59235)

- Before this update, when a machine set was scaled down and had reached its minimum size, the Cluster Autoscaler could leave the last remaining node with a no schedule taint that prevented use of a node. This issue was caused by a counting error in the Cluster Autoscaler. With this release, the counting error has been fixed so that the Cluster Autoscaler works as expected when a machine set is scaled down and has reached its minimum size. (OCPBUGS-59260)

- Before this update, bundle unpack jobs did not inherit control-plane tolerances from the catalog-operator that created them. As a consequence, the bundle unpack jobs ran on only worker nodes. If no worker nodes were available due to taints, then admins are unable to install or upgrade Operators on the cluster. With this release, control-plane tolerations are adopted for bundle unpack jobs so that the jobs are executed on primary nodes as part of the control plane. (OCPBUGS-59421)

1.9.18.2. Updating

To update an OpenShift Container Platform 4.18 cluster to this latest release, see Updating a cluster using the CLI.

1.9.19. RHSA-2025:10767 - OpenShift Container Platform 4.18.20 bug fix and security update

Issued: 17 July 2025

OpenShift Container Platform release 4.18.20 is now available. The list of bug fixes that are included in the update is documented in the RHSA-2025:10767 advisory. The RPM packages that are included in the update are provided by the RHSA-2025:10768 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 4.18.20 --pullspecs
```

1.9.19.1. New features and enhancements

1.9.19.1.1. On-cluster layering changes is now generally available

There are several important changes to the on-cluster layering feature:

- The API version is now **machineconfiguration.openshift.io/v1**. The new version includes the following changes:
  - The **baseImagePullSecret** parameter is now optional. If not specified, the default **global-pull-secret-copy** is used.
  - The **buildInputs** parameter is no longer required. All parameters previously under the **buildInputs** parameter are promoted one level.
  - The **containerfileArch** parameter now supports multiple architectures. Previously, only **noarch** was supported.
  - The required **imageBuilderType** is now **Job**. Previously, the required builder was **PodImageBuilder**.
  - The **renderedImagePushspec** parameter is now **renderedImagePushSpec**.
  - The **buildOutputs** and **currentImagePullSecret** parameters are no longer required.
- You can manually rebuild your custom layered image by applying an annotation to the **MachineOSConfig** object.
- You can now automatically delete an on-cluster custom layered image by deleting the associated **MachineOSBuild** object.
- On-cluster layering is now supported in disconnected environments.

For more information, see Using on-cluster layering to apply a custom layered image .

1.9.19.2. Bug fixes

- Before this update, a subset of endpoints on the console backend were gated by **TokenReview** requests to the API server. In some cases, the API server would throttle these requests, causing slower load times in the UI. With this release, the **TokenReview** gating was removed from all but one of our endpoints resulting in improved performance. (OCPBUGS-58317)
- Before this update, useful error messages were not generated when the`oc adm node-image create` command failed. With this release, the **oc adm node-image create** command generates error messages when the command fails. (OCPBUGS-58233)
- Before this update, the HAProxy configuration used the /**version** endpoint for health checks causing unreliable health checks to be generated. With this release, the liveness probe is customized to use /**livez?exclude=etcd&exclude=log** on IBM Cloud for more accurate health checks. This change avoids disruptions due to inappropriate probe configurations on hosted control planes, while retaining /**version** for other platforms. (OCPBUGS-58126)
- Before this update, the Machine Config Operator (MCO) updated boot images without verifying that the current boot image was from the marketplace. As a consequence, the MCO overrode a marketplace boot image with a standard OpenShift Container Platform installer image. With this release, the MCO references a lookup table in Amazon Web Services (AWS) that contains all standard installer Amazon Machine Images (AMIs) in OpenShift Container Platform before it updates the boot image. In Google Cloud, the MCO checks the URL header before updating the boot image. As a result, the MCO does not update machine sets that have a marketplace boot image. (OCPBUGS-58044)
- Before this update, a validation issue in the **oc-mirror** plugin caused the command to reject the

file://. reference. Using **file://.** for a content path generated an error message stating **content filepath is tainted**. With this release, **oc-mirror** properly validates the **.** directory reference. (OCPBUGS-57970)

- Before this update, the **oc-mirror v2** plugin was not using the correct filtered catalog during its operations. As a consequence, unspecified operators were included in the configuration and the plugin tried to connect to the catalog registry during disk-to-mirror workflows, even in air-gapped environments. With this release, the correct filtered catalog is used. (OCPBUGS-57964)

- Before this update, modals that used the same **useModal** hook instance overwrote each other. As a consequence, the OpenShift Container Platform Lightspeed user interface disappeared. With this release, modals have unique IDs. As a result, modal conflicts are resolved, which enables the simultaneous user interface display for the **Lightspeed**, **Troubleshooting Panel**, and **Networking** pages. (OCPBUGS-57931)

- Before this update, images were ignored and recreated on every reconcile call, which caused new caches and prevented the use of cached images. As a consequence, memory usage of hosted control planes quickly grew, which created performance issues. With this release, images in hosted control planes are cached effectively by creating and using a global registry provider instead of recreating registry and release providers on every reconcile call. As a result, memory usage is optimized in hosted control planes. (OCPBUGS-57818)

- Before this update, **OLMv1** was used to install Operators with the **olm.maxOpenShiftVersion** set to **4.19**. Because of an issue with the **OLMv1** parsing logic for floating-point formatted **olm.maxOpenShiftVersion`values, the system failed to prevent upgrades to OpenShift Container Platform 4.20. With this release, the parsing logic for `olm.maxOpenShiftVersion** is corrected. As a result, this correction prevents upgrades to OpenShift Container Platform 4.20 when Operators that include **olm.maxOpenShiftVersion:4.19** are installed. (OCPBUGS-57767)

- Before this update, the catalog-operator captured catalog snapshots with a frequency of five minutes. When using many namespaces and subscriptions, and with larger catalog sources available in 4.15, 4.16, the snapshots were failing, but cascaded across the catalog sources which caused CPU loads to spike. This additional load caused an inability to upgrade and install operators. With this release, the cache lifetime is 30 minutes, which provides sufficient time for attempts to be resolved without undue load on the catalog source pods. (OCPBUGS-57427)

- Before this update, deleting a common user data network (CUDN) resource in a namespace with an existing, endpoint-less Service caused the **ovnkube-node** pod restart to fail. With this release, the **ovnkube-node** pod restarts successfully after you delete CUDN resources with existing, endpoint-less services in the targeted namespace. (OCPBUGS-57318)

- Before this update, the **Create PodDisruptionBudget** page contained a typo. With this release, the typo is corrected. (OCPBUGS-57213)

- Before this update, when re-running oc-mirror plugin v2 with the same working directory, existing **tar** archive files from previous runs were not removed. This resulted in a mix of outdated and new archives, which could cause mirroring failures when pushing to the target registry. With this release, oc-mirror plugin v2 automatically deletes old **tar** archive files at the beginning of each run, ensuring that the working directory contains only archives from the current execution. (OCPBUGS-57197)

- Before this update, oc-mirror v2 rejected valid **ImageSetConfiguration** parameter values that contained the words **mirror** or **delete**. With this release, oc-mirror v2 now correctly validates the words **delete** and **mirror** in the **ImageSetConfiguration** parameter and rejects only invalid configurations. (OCPBUGS-57124)

- Before this update, the **cadvisor** endpoint on the kubelet reported invalid metrics values, which combined the counter data with the data from different devices into one metric. With this release, the **cadvisor** endpoint reports valid metrics. (OCPBUGS-57070)

- Before this update, one of the **keepalived** health check scripts failed due to missing permissions. This failure sometimes caused the ingress VIP to be misplaced when shared ingress services were in use. With this release, the necessary permission is added back to the container. As a result, the health check works correctly. (OCPBUGS-56624)

- Before this update, **hostPath** volumes in the **kube-rbac-proxy-crio** pod were configured with read-write access, which violated security best practices for Kubernetes security. As a consequence, unauthorized modification of system files occurred because of the read/write **hostPath** mounts. With this release, **hostPath** mounts in the `kube-rbac-proxy-crio`pod are read-only to improve security. (OCPBUGS-55246)

- Before this update, for clusters that were installed with the Agent-based Installer for versions 4.15.0 to 4.15.26, root certificates that were built in from CoreOS were added to the user-ca-bundle, even though they were not explicitly specified by the user. In previous releases, when you added a node to one of these clusters using the **oc adm node-image create** command, the **additionalTrustBundle** value obtained from the cluster's **user-ca-bundle** was too large to process, resulting in a failure to add the node. With this release, the built-in certificates are filtered out when generating the **additionalTrustBundle** value, so that only explicitly user-configured certificates are included, and nodes can be added successfully. (OCPBUGS-54744)

- Before this update, the **HostedCluster** command failed to create network policies because of an invalid IP address format in the **nodePort** configuration. As a consequence, the creation of the hosted cluster failed due to an incorrect IP address determination. This release adds log messages for IP address type determination in the **virt** launcher network policy reconciliation. The additional log messages resolve the incorrect IP address determination in the **kubevirt** hosted cluster. As a result, the **HostedCluster** command successfully creates the network policies and the hosted cluster. (OCPBUGS-46629)

- Before this update, Kubernetes VM with dynamic IPv6 configuration used primary user-defined network (UDN) layer2, which caused the default IPv6 gateway to be multipath. As a consequence, user traffic flowed incorrectly between nodes because of multiple default IPv6 gateways. With this release, the default IPv6 gateway is correct for dynamic VM configurations due to using the new Open Virtual Network (OVN) transit router topology. As a result, the default IPv6 gateway points to the correct node, which reduces inter-node traffic and improves network stability during node maintenance. (OCPBUGS-46401)

1.9.19.3. Updating

To update an OpenShift Container Platform 4.18 cluster to this latest release, see Updating a cluster using the CLI.

1.9.20. RHSA-2025:9725 - OpenShift Container Platform 4.18.19 bug fix and security update

Issued: 2 July 2025

OpenShift Container Platform release 4.18.19 is now available. The list of bug fixes that are included in the update is documented in the RHSA-2025:9725 advisory. The RPM packages that are included in the update are provided by the RHSA-2025:9726 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 4.18.19 --pullspecs
```

1.9.20.1. Bug fixes

- Previously, 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-57580)
- 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 keeps the destination folder permissions. (OCPBUGS-57507)
- Previously, using Alertmanager apiVersion v1 in the **openshift-monitoring/cluster-monitoring-config** or **openshift-user-workload-monitoring/user-workload-monitoring-config** parameters caused OpenShift Container Platform 4.19 to fail early with an **InvalidConfigXXX** error. This issue occurred because OpenShift Container Platform 4.19 is using Prometheus v3 which does not support Alertmanager apiVersion v1. With this release, Cluster Monitoring Operator (CMO) sets a value of **upgradable=false** if apiVersion v1 is detected in the config maps to prevent the **InvalidConfigXXX** error from occurring after upgrading to OpenShift Container Platform 4.19. As a result, clusters upgrade through OpenShift Container Platform 4.18.x before moving to OpenShift Container Platform 4.19. (OCPBUGS-56251)

1.9.20.2. Updating

To update an OpenShift Container Platform 4.18 cluster to this latest release, see Updating a cluster using the CLI.

1.9.21. RHSA-2025:9269 - OpenShift Container Platform 4.18.18 bug fix and security update

Issued: 25 June 2025

OpenShift Container Platform release 4.18.18 is now available. The list of bug fixes that are included in the update is documented in the RHSA-2025:9269 advisory. The RPM packages that are included in the update are provided by the RHBA-2025:9270 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 4.18.18 --pullspecs
```

1.9.21.1. Bug fixes

- Previously, the Started column was hidden on reduced screen sizes, causing the **VirtualizedTable** component to break due to a missing sort function. As a consequence, this broken table component prevented users from consistently viewing **pipelinerun** list pages. With this release, the **VirtualizedTable** component now handles missing sort functions for default columns for reduced screen sizes. As a result, you can consistently view **pipelinerun** list pages, regardless of screen size. (OCPBUGS-57353)

- Previously, when you changed the order of selectors in the **ClusterRole** parameter for the **OperatorGroup** in Operator Lifecycle Management (OLM), unnecessary etcd writes and auth cache invalidation degraded performance. With this release, an update to OLM prevents unnecessary etcd writes and auth cache invalidation when you change the selector order in the **ClusterRole** parameter. (OCPBUGS-57314)

- Previously, the agent-based installer ignored the custom **additionalTrustBundlePolicy** parameter because of a missing field in the **install-config.yaml** file. Consequently, cluster installations sometimes did not comply with specified settings due to ignored overrides. With this release, the **additionalTrustBundlePolicy** config overrides are now properly applied in the **install-config.yaml** file for the assisted-service. As a result, you can correctly set the **additionalTrustBundlePolicy** parameter, and other installation configuration overrides are correctly applied. (OCPBUGS-57306)

- 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-57273)

- Previously, when installing into an existing virtual private cloud (VPC) on Amazon Web Services (AWS), a potential mismatch could occur in the subnet information in the AWS Availability Zone between the machine set custom resources for control plane nodes and their corresponding AWS EC2 instances. As a consequence, where the control plane nodes were spread across three Availability Zones and one was recreated the discrepancy could result in an unbalanced control plane as two nodes occurred within the same Availability Zone. With this release, it is ensured that the subnet availability zone (AZ) information in the machine set custom resources and in the EC2 instances match and the issue is resolved. (OCPBUGS-57220)

- 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 which was run on the Network File System (NFS)). With this release, the kubelet reports metrics even if a disk is stuck. (OCPBUGS-57219)

- Previously, the **/metrics** endpoint was not correctly parsing a bearer token from the **Authorization** header on internal Prometheus scrape requests. Consequently, the **TokenReview** failed and all the scrape requests returned a 401 response. With this release, the metrics endpoint handler is updated to correctly parse bearer tokens in the **Authorization** header for the **TokenReview**. This update resolves the **TargetDown** alert in the OpenShift Container Platform web console. (OCPBUGS-57181)

- 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-57139)

- Previously, 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-56955)

- 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. (OCPBUGS-56835)

- Previously, a Machine Config Operator (MCO) incorrectly set an **Upgradeable=False** condition to all new nodes that were added to a cluster. The condition stated a **PoolUpdating** reason for set condition. With this release, the MCO now correctly sets **Upgradeable=True** condition to all new nodes that get added to a cluster so that the issue no longer exists. (OCPBUGS-56517)

- Previously, the IDMS or ICSP resources from the management cluster were processed without considering that a user might specify only the root registry name as a mirror or source for image replacement. As a consequence, any IDMS or ICSP entries that used only the root registry name did not work as expected. With this release, the mirror replacement logic now correctly handles cases where only the root registry name is provided. As a result, the issue no longer occurs, and the root registry mirror replacements are now supported. (OCPBUGS-56166)

- Previously, during image cleanup, oc-mirror plugin v2 would stop the deletion process if any error occurred while removing an image. With this release, oc-mirror plugin v2 continues attempting to delete remaining images even after encountering errors. After the process completes, it displays a list of any failed deletions. (OCPBUGS-56125)

1.9.21.2. Updating

To update an OpenShift Container Platform 4.18 cluster to this latest release, see Updating a cluster using the CLI.

1.9.22. RHSA-2025:8560 - OpenShift Container Platform 4.18.17 bug fix and security update

Issued: 10 June 2025

OpenShift Container Platform release 4.18.17 is now available. The list of bug fixes that are included in the update is documented in the RHSA-2025:8560 advisory. The RPM packages that are included in the update are provided by the RHBA-2025:8561 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 4.18.17 --pullspecs
```

1.9.22.1. Bug fixes

- Previously, a single-node OpenShift deployment on OpenShift Container Platform 4.11 failed on a Red Hat OpenStack Platform (RHOSP) provider because of an unsupported installation on the platform. With this release, single-node OpenShift deployments support installations on RHOSP, which enhances installation flexibility. (OCPBUGS-56864)

- Previously, the **disk2mirror** process did not display logs during the cache registry population, which caused an incomplete process. With this release, the working and cache directories are verified before adding the extracted mirror archives. This update improves visibility during the **disk2mirror** process and reduces user uncertainty about an incomplete process. (OCPBUGS-56659)

1.9.22.2. Updating

To update an OpenShift Container Platform 4.18 cluster to this latest release, see Updating a cluster using the CLI.

1.9.23. RHSA-2025:8284 - OpenShift Container Platform 4.18.16 bug fix update

Issued: 03 June 2025

OpenShift Container Platform release 4.18.16 is now available. The list of bug fixes that are included in the update is documented in the RHSA-2025:8284 advisory. The RPM packages that are included in the update are provided by the RHBA-2025:8285 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 4.18.16 --pullspecs
```

1.9.23.1. Bug fixes

- Previously, in non-Zonal Azure regions, a bug in the dynamically computed fault domain count update caused scaling failures. This issue prevented upgrades to OpenShift Container Platform version 4.15.48 and later from scaling existing machine sets. With this release, an update is implemented that rectifies the dynamic fault domain calculation and prevents scaling failures in non-Zonal regions. This fix ensures smooth operation of machine set scaling after upgrading to OpenShift Container Platform version 4.15.48 and later. (OCPBUGS-56654)

- Previously, API server subject alternative name (SAN) validation was performed regardless of public key infrastructure (PKI) reconciliation status, causing potential connectivity issues with hosted control plane clusters due to invalid SANs. With this release, a fix removes validation for OpenShift Container Platform API server SANs when PKI reconciliation is disabled. The hosted control plane performance is improved by eliminating unnecessary validation, especially when the PKI reconciliation is not managed. (OCPBUGS-56627)

- Previously, a bug in the Bare Metal Operator (BMO) caused JSON parsing errors because of a missing Redfish system ID in Baseboard Management Controller (BMC) URLs. This issue caused users to receive errors when the system ID was left out of the URLs. With this release, the BMO handles URLs without a Redfish system ID as addresses without a system ID. This fix improves software handling of missing a Redfish system ID in BMC URLs. (OCPBUGS-56431)

- Previously, during Ironic Python Agent (IPA) deployments, the absence of NetworkManager logs in RAM disk logs hindered effective debugging, impacting network issue resolution. With this release, NetworkManager logs are included in RAM disk logs for IPA debugging. This results in enhanced IPA logs that provide comprehensive **NetworkManager** data for improved debugging. (OCPBUGS-56097)

- Previously, Helm did not support Docker image mirroring with a tag and digest. This resulted in failed Helm repository mirroring that caused image duplication and inconsistencies in deployment. With this release, a fix addresses Docker references in Helm repository mirroring that allows tag and digest, and improves successful image mirroring. (OCPBUGS-56043)

- Previously, an issue started from an incorrect bucket name that was used for Red Hat Enterprise Linux CoreOS (RHCOS). Users could not create OpenShift Container Platform clusters because of a failed RHCOS image import. This problem was resolved by correcting the cluster creation in the Madrid zone for the **PowerVS** installer-provisioned infrastructure Cluster CAPI Operator by using the correct bucket name. With this release, the user can create OpenShift Container Platform clusters in the Madrid zone by using the installation program. (OCPBUGS-53142)

- Previously, intermittent resource leaks in the **MachineOSConfig** (MOSC) to **MachineOSBuild** (MOSB) connection due to missing owner references during job creation resulted in potential resource exhaustion, affecting pod updates. With this release, the owner references are added in the job creation to ensure the consistent removal of MOSB resources when MOSC is deleted, preventing resource leaks. (OCPBUGS-52189)

1.9.23.2. Updating

To update an OpenShift Container Platform 4.18 cluster to this latest release, see Updating a cluster using the CLI.

1.9.24. RHBA-2025:8104 - OpenShift Container Platform 4.18.15 bug fix update

Issued: 27 May 2025

OpenShift Container Platform release 4.18.15 is now available. The list of bug fixes that are included in the update is documented in the RHBA-2025:8104 advisory. The RPM packages that are included in the update are provided by the RHBA-2025:8106 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 4.18.15 --pullspecs
```

1.9.24.1. Bug fixes

- Previously, the web console incorrectly displayed a 60-day update limitation message in version 4.16. This outdated message misled users in OpenShift Container Platform 4.16 and later versions. With this release, the 60-day update warning is removed from the web console, which ensures an accurate and an up-to-date user experience. (OCPBUGS-56255)

- Previously, secrets and credentials in **MachineConfigs** secrets were added to audit logs, which exposed sensitive data. With this release, an update ignores **MachineConfig** secrets in audit logs, which results in the removal of sensitive data from the logs, and ensures improved data security. (OCPBUGS-56030)

- Previously, the token rotation mechanism incorrectly created a time frame without a valid token for image authentication that resulted in temporary authentication failures in the image registry.

This failure affected pod scheduling and build processes. With this release, an update enhances the token refresh mechanism in OpenShift Container Platform to prevent invalid tokens. This improvement reduces authentication failures and ensures a smoother operation of the image registry and related processes. (OCPBUGS-54304)

- Previously, a host shutdown caused a failure during an Open Virtual Appliance (OVA) import in a VMware vSphere cluster. An update was made to ensure that the vSphere ESXi host was not powered off or in maintenance mode during the import process. With this release, the update allows for a successful OVA import without disruption. (OCPBUGS-50690)

1.9.24.2. Updating

To update an OpenShift Container Platform 4.18 cluster to this latest release, see Updating a cluster using the CLI.

1.9.25. RHSA-2025:7863 - OpenShift Container Platform 4.18.14 bug fix update and security update

Issued: 20 May 2025

OpenShift Container Platform release 4.18.14 is now available. The list of bug fixes that are included in the update is documented in the RHSA-2025:7863 advisory. The RPM packages that are included in the update are provided by the RHBA-2025:7865 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 4.18.14 --pullspecs
```

1.9.25.1. Bug fixes

- Previously, the omission of the **OLMManagedLabelKey** label on objects resulted in cluster operation failures. With this release, an update improves pod stability and ensures that the Operator Lifecycle Manager operates properly. (OCPBUGS-56098)

- Previously, an invalid **.tar** extraction format resulted in an improper file separation in the **ramdisk** logs, and caused file separators to appear randomly. With this release, an updated **ramdisk** log file processes **.tar** entries individually. This fix improves log readability, making them easier to interpret. (OCPBUGS-55938)

- Previously, incorrectly formatted proxy variables in an external binary resulted in build failures. With this release, an update removes proxy environment variables from the build pod and prevents any build failures. (OCPBUGS-55699)

- Previously, no event was logged when an error occurred from failed conversion from ingress to route. With this update, this error appear in the event logs. (OCPBUGS-55338)

- Previously, a missing **afterburn** package resulted in the failure of the **gcp-hostname.service**, which caused the **scale-up** job to fail, impacting end-user deployments. With this release, the **afterburn** package is installed in the RHEL **scale-up** job. This fix enables a successful **scale-up** action, resolving the **gcp-hostname** service failure. (OCPBUGS-55158)

- 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-55016)

- 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-54663)

- Previously, you could add white space to Amazon Web Services (AWS) tag names, but the installation program did not support them. This situation resulted in the installation program returning an **ERROR failed to fetch Metadata** message. With this release, the regular expression for AWS tags now validates any tag name that has white space. The installation program accepts these tags and no longer returns an error because of white space. (OCPBUGS-53221)

- Previously, cluster nodes repeatedly lost communication due to improper remote port binding by Open Virtual Network (OVN)-Kubernetes. This affected pod communication across nodes. With this release, the remote port binding functionality is updated to be handled by OVN directly, improving the reliability of cluster node communication. (OCPBUGS-51144)

1.9.25.2. Updating

To update an OpenShift Container Platform 4.18 cluster to this latest release, see Updating a cluster using the CLI.

1.9.26. RHSA-2025:4712 - OpenShift Container Platform 4.18.13 bug fix update and security update

Issued: 14 May 2025

OpenShift Container Platform release 4.18.13 is now available. The list of bug fixes that are included in the update is documented in the RHSA-2025:4712 advisory. The RPM packages that are included in the update are provided by the RHBA-2025:4714 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 4.18.13 --pullspecs
```

1.9.26.1. Known issues

- When a pod uses the CNI plugin for DHCP address assignment, in conjunction with other Container Network Interface (CNI) plugins, the pod’s network interface might be unexpectedly deleted. As a result, when the pod’s DHCP lease expires, the DHCP proxy enters a loop when trying to re-create a new lease, leading to the node becoming unresponsive. There is currently no known workaround. (OCPBUGS-55354)

1.9.26.2. Bug fixes

- Previously, unsorted image stream names in the **Progressing** condition caused unnecessary updates. This led to excessive user updates and potential performance degradation. With this release, the failing image imports are sorted within the **activeImageStreams** function. This change improves Cluster Samples Operator efficiency, reduces unnecessary updates, and enhances overall performance. (OCPBUGS-55783)

- 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-55782)
- 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-55781)
- Previously, when adding a node in a disconnected environment, private registry images were inaccessible for the **oc adm node-image** command. As a result, issues with pulling images prevented adding a node. This error only occurs if the cluster is initially installed with an installation program binary downloaded from **mirror.openshift.com**. The issue is resolved in this release. (OCPBUGS-55449)
- Previously, there was an issue in the image reference digest calculation that led to a failed container creation based on the SchemavVersion 1 image. This prevented new deployment creations. With this release, the image digest calculation is fixed and the new Operators can be installed. (OCPBUGS-55435)
- Previously, Microsoft Azure Spot Virtual Machines (VMs) that were evicted before their node became ready, could get stuck in the **provisioned** state. With this release, Azure Spot VMs now use a delete eviction policy, ensuring that the VMs correctly transition to the **failed** state if they are preempted. (OCPBUGS-55373)
- Previously, the **oc-mirror** plugin returned an exit status of **0**, indicating success, even when mirroring errors occurred in automated workflows. As a result, customers could not rely on the exit status in the automated workflows. With this release, the **oc mirror** plugin returns a null exit status, indicating failure, when there are errors. (OCPBUGS-54626)
- Previously, an image pull secret generated for the internal image registry was not regenerated until after the embedded credentials expired. This resulted in a small period of time during which the image pull secrets were invalid. With this release, the image pull secrets are refreshed before the embedded credentials expire. (OCPBUGS-54304)
- Previously, the Machine Config Operator (MCO) in OpenShift Container Platform 4.18 was not updated to reflect that package-based Red Hat Enterprise Linux (RHEL) support was removed in 4.19. With this release, this Operator ensures compatibility by blocking updates to OpenShift Container Platform 4.19 on clusters with packaged-based RHEL nodes. (OCPBUGS-53427)

1.9.26.3. Updating

To update an OpenShift Container Platform 4.18 cluster to this latest release, see Updating a cluster using the CLI.

1.9.27. RHSA-2025:4427 - OpenShift Container Platform 4.18.12 bug fix update and security update

Issued: 08 May 2025

OpenShift Container Platform release 4.18.12 is now available. The list of bug fixes that are included in the update is documented in the RHSA-2025:4427 advisory. The RPM packages that are included in the update are provided by the RHBA-2025:4429 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 4.18.12 --pullspecs
```

1.9.27.1. Bug fixes

- Previously, a race condition in the hostname that handled the code caused inconsistencies between the boot and machine hostnames. With this release, the race condition is resolved, which ensures consistent hostnames in the Ignition configuration file during the operating system installation. (OCPBUGS-55364)

- Previously, a version of the installation program produced boot image update failures and region-specific user issues because of missing Amazon Machine Image (AMI) IDs in certain regions. With this release, when a region AMI is not found, the region defaults to the **us-east-1** AMI, and the installation program has reliable, default AMIs for all regions. (OCPBUGS-55290)

- Previously, when viewing the list of installed Operators and the currently selected project matches an Operator's default namespace while copied cluster service versions (CSV) are disabled in the Operator Lifecycle Manager (OLM), an Operator appeared twice in the list. With this release, the Operator appears one time. (OCPBUGS-55195)

- Previously, certain IP addresses in the **namedCertificates** server configuration conflicted with the internal API URLs. This condition caused users to experience configuration issues with the **HostedCluster** custom resource because of a subject alternative name (SAN) mismatch in the certificates. With this release, the conflicting SANs in the Kasm Workspaces agent (KAS) server certificates are resolved, ensuring proper configuration and improved service functionality. (OCPBUGS-54946)

- Previously, insufficient memory requests for proxy containers, such as **socks5-proxy**, **konnectivity-proxy**, **http-proxy**, and **client-token-minter**, often caused performance issues. With this release, memory requests for these containers are increased to 30 Megabytes and the steady-state performance is improved by providing more memory to the proxy containers. (OCPBUGS-54737)

1.9.27.2. Updating

To update an OpenShift Container Platform 4.18 cluster to this latest release, see Updating a cluster using the CLI.

1.9.28. RHSA-2025:4211 - OpenShift Container Platform 4.18.11 bug fix update and security update

Issued: 01 May 2025

OpenShift Container Platform release 4.18.11 is now available. The list of bug fixes that are included in the update is documented in the RHSA-2025:4211 advisory. The RPM packages that are included in the update are provided by the RHBA-2025:4213 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 4.18.11 --pullspecs
```

1.9.28.1. Bug fixes

- Previously, service deletion improperly handled API service association, resulting in API service unavailability. When the **ClusterResourceOverride** resource was deleted, the **admission.autoscaling.openshift.io/v1** API service was unreachable, affecting Operator installations. With this release, deleting the **ClusterResourceOverride** resource removes associated API services, allowing the operators to retrieve the list of server APIs for successful installation. (OCPBUGS-55242)

- Previously, during the Cluster Resource Override Operator (CROO) upgrade from OpenShift Container Platform version 4.16 to 4.17, old secrets were not deleted, causing the CROO to fail after the OpenShift Container Platform upgrade. With this release, the pod creation and namespace deletion during the OpenShift Container Platform 4.17 upgrade completes successfully, resolving CROO errors. (OCPBUGS-55240)

- Previously, missing catalogd certificate authorities (CA) resulted in failed Operator Lifecycle Manager (OLM) v1 installations. With this release, an updated Operator Controller uses a new directory for CA certificates. This change improves the stability of the system, ensuring the correct installation of cluster extensions, and enhancing the user experience. (OCPBUGS-55172)

- Previously, a private IP address mismatch for the load balancer led to the failure of fetching Azure IP availability, causing a private IP address conflict in an OpenShift 4.17 deployment. This issue was resolved by checking for IP address availability within the control plane subnet. The fix resulted in resolving the error in Azure IP address availability for OpenShift 4.17 deployment, ensuring that private IP addresses are now validated within subnet ranges. (OCPBUGS-54947)

- Previously, in 4.18 4.16, a migration failure occurred after a reboot for users moving from OpenShift SDN to **OVNKubernetes** due to an interface configuration issue. The failure occurred because the **mtu-migration** service that was active before the **wait-for-primary-ip** service in the NMState-managed **br-ex**. With this release, the order of these services are reversed to ensure a successful migration and to prevent the **mtu-migration** service failure that occurs after the first reboot. (OCPBUGS-54817)

- Previously, missing cluster role permissions for Cluster Network Operator (CNO) in the **monitoring.coreos.com** and **monitoring.rhobs** APIs caused monitoring issues because of insufficient permissions. With this release, permissions for CNO to manage the **servicemonitors** and **prometheusrules** objects exist. The CNO patches **servicemonitor** and **prometheusrules** objects in the **monitoring.coreos.com** API group, which corrects the monitoring issues. (OCPBUGS-54698)

1.9.28.2. Updating

To update an OpenShift Container Platform 4.18 cluster to this latest release, see Updating a cluster using the CLI.

1.9.29. RHSA-2025:4019 - OpenShift Container Platform 4.18.10 bug fix update and security update

Issued: 22 April 2025

OpenShift Container Platform release 4.18.10 is now available. The list of bug fixes that are included in the update is documented in the RHSA-2025:4019 advisory. The RPM packages that are included in the update are provided by the RHBA-2025:4021 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 4.18.10 --pullspecs
```

1.9.29.1. Enhancements

- In the bootstrap phase of the installation process, the Transport Layer Security (TLS) between the **metal3 httpd** server and the node’s Baseboard Management Controller (BMC) is enabled by default in OpenShift Container Platform 4.18 and later. The **httpd** server is on port 6183 instead of port 6180 when TLS is enabled. Disable the TLS setting by adding 'disableVirtualMediaTLS: true' to the provisioning custom resource (CR) file that is created on the disk. (OCPBUGS-39404)

1.9.29.2. Bug fixes

- Previously, the Prometheus remote-write proxy configuration was not correctly applied to the Prometheus user workload custom resource (CR), which caused communication and data collection problems in the cluster. With this release, the user workload monitoring (UWM) Prometheus configurations, including user workload Prometheus, correctly inherit the proxy settings from the cluster proxy resource. (OCPBUGS-38655)

- Previously, when running Red Hat Enterprise Linux CoreOS (RHCOS) in an active environment, the **rpm-ostree-fix-shadow-mode.service** systemd service that used to run caused that service to fail. With this release, the **rpm-ostree-fix-shadow-mode.service** systemd service does not activate when RHCOS does not run from an installed environment. (OCPBUGS-41625)

- Previously, an incorrect component import in the **SimpleSelect.tsx** file caused an undefined function **r** function in the **react-dom.production.min.js** file. This component caused error messages on the **Dashboards** and **Metrics** pages related to dropdown lists. With this release, the dropdown lists on the affected pages function correctly, eliminating the error message. (OCPBUGS-42845)

- Previously, an error in the rotation logic of the image pull secret controller’s secret token caused a temporary, invalid token for authentication. As a consequence, the image pull process was disrupted. With this release, the updated image pull secret controller eliminates the period when the token is not valid while the token rotates. As a result, the image pull process is smooth and continuous. (OCPBUGS-54304)

- Previously, an error occurred in hosted control planes-managed clusters because of the omission of the **shutdown-watch-termination-grace-period** setting in the **kube-apiserver** configuration. This error led to the unstable shutdown of applications in hosted control planes-managed clusters. With this release, an update improves the shutdown process of applications in hosted control planes-managed clusters, providing a grace period for the **kube-apiserver** configuration. During a shutdown, the application stability is improved and potential errors are decreased. (OCPBUGS-53404)

- Previously, an issue with the version of the **github.com/sherine-k/catalog-filter** element stopped, causing instability in the mirroring process. With this release, the **github.com/sherine-k/catalog-filter** element in the **go.mod** file is updated, which solves the problem and ensures a stable and reliable mirroring process. (OCPBUGS-54727)

- Previously, an iteration counter increment omission in the **scrapeCache** setting led to an incorrect series count for subsequent scrapes. As a result, monitoring was interrupted and data could potentially be lost during the Prometheus scrape process. With this release, an update ensures uninterrupted monitoring, because Prometheus continues scraping and processing data while parsing errors. (OCPBUGS-54940)

1.9.29.3. Updating

To update an OpenShift Container Platform 4.18 cluster to this latest release, see Updating a cluster using the CLI.

1.9.30. RHSA-2025:3775 - OpenShift Container Platform 4.18.9 bug fix update and security update

Issued: 15 April 2025

OpenShift Container Platform release 4.18.9 is now available. The list of bug fixes that are included in the update is documented in the RHSA-2025:3775 advisory. The RPM packages that are included in the update are provided by the RHBA-2025:3777 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 4.18.9 --pullspecs
```

1.9.30.1. Bug fixes

- Previously, the **manifest-topology.yaml** file was not added for the topology-related feature in VMware vSphere. With this release, the **manifest-topology.yaml** file is added for the topology-related feature and tested resulting in improved performance and enhanced end-user experience when using the topology feature. (**OCPBUGS-54701**)

- Previously, the OVN-Kubernetes container failed to start because of the incorrect handling of User-Defined Networks (UDN) in **EgressIP** logical router policies. Users experienced intermittent deployment failures on AWS, which led to prolonged downtime and service disruptions. With this release, the OVN-Kubernetes container starts successfully with UDNs configured. (**OCPBUGS-54671**)

- Previously, the identity provider (IdP) reconciler did not consider the additional trust bundle for customer proxies, leading to failed TLS certificate verification and IdP integration failures in hosted clusters. This resulted in service interruptions for end users. With this release, the TLS certificate verification problem is resolved, allowing the IdP to function correctly in hosted clusters with a proxy configuration that specifies an additional trust bundle, resulting in an improved end-user experience with seamless IdP integration. (**OCPBUGS-54627**)

- Previously, the control plane controller did not select the correct Cluster Version Operator (CVO) manifests for the required feature set. With this release, the control plane controller selects the correct CVO manifests, which are deployed for the hosted cluster. (**OCPBUGS-54625**)

- Previously, an issue arose when the expiration timestamp annotation on ignition tokens was reset, which should not occur. This led to the accumulation of outdated tokens, and caused resource mismanagement or security vulnerabilities within the cluster. With this release, this results in an improved end-user experience as the hosted control planes Operator effectively cleans up expired ignition tokens, ensuring efficient resource management and enhancing system security. (OCPBUGS-54624)

- Previously, a bug occurred due to insufficient enforcement of a minimum number of services for IBM Cloud in the **HostedControlPlane** and **HostedCluster** specifications within the code. This issue led to potential data loss or incorrect processing of user-entered data, causing unexpected application behavior. With this release, the problem causing inaccurate data to display in the user interface is corrected, ensuring more reliable and precise information for end users. (OCPBUGS-54609)

- Previously, improper scoping of the Secret Janitor in the Hypershift Operator, which caused improper secret cleanup. This resulted in token secrets accumulating over time, disrupting the secret management process, while using annotation scoping with two instances of the Hypershift Operator. With this release, a fix ensures that the secret cleanup continues as expected in a Red Hat OpenShift Kubernetes Service (ROKS) cluster that is managed by the Hypershift Operator. The large amount of token secrets is eliminated and the proper secret management is maintained. (OCPBUGS-54498)

- Previously, a bug occurred due to the incorrect handling of the etcd URL, preventing access to the **Kyverno** service. This resulted in DNS errors during the **kyverno** validation, preventing users on a OpenShift Container Platform cluster with hosted control planes from creating additional test groups. With this release, users can create additional test groups without encountering DNS errors during the **kyverno** validation. (OCPBUGS-54411)

- Previously, insufficient permissions led to the persistence of the **disk.csi.azure.com/agent-not-ready=value:NoExecute** taint after creating Microsoft Azure disk Container Storage Interface (CSI) driver nodes. With this release, a fix disables the removal of the **not-ready** taint for Azure disk CSI driver nodes, making the scheduler adhere to the **volume-attach-limit** value. (OCPBUGS-54383)

- Previously, containers that used the SELinux domain of **container_logreader_t** to view container logs on a host in the **/var/log** directorycould not access logs in the **/var/log/containers** subdirectory. This was because of a missing symbolic link. With this release, a symbolic link is created so the containers can access the logs in the **/var/log/containers** subdirectory. (OCPBUGS-54342)

- Previously, an image pull secret that was generated for the internal Image Registry was not regenerated until after the embedded credentials expired. This resulted in the image pull secrets being temporarily invalid. With this release, the image pull secrets are refreshed before the embedded credentials expire. (OCPBUGS-54304)

- Previously, the cluster autoscaler stopped scaling because of a failed machine in a machine set. This condition occurred due to inaccuracies in the way the cluster autoscaler counted machines in various non-running phases. With this release, the inaccuracies are fixed, allowing the cluster autoscaler to have an exact count. (OCPBUGS-53241)

- Previously, the telecom grandmaster (T-GM) status was incorrectly announced as **S2** before the digital phase-locked loop (DPLL) was locked during the global navigation satellite system (GNSS) holdover. This caused inaccurate synchronization. With this release, the **DPLL** state decision logic is modified to ensure that the T-GM status moves to **S2** only after both phase offsets are valid and DPLL is in "Locked Holdover Acquired" state. This guarantees that the T-GM status accurately reflects the DPLL state when the GNSS source starts. (OCPBUGS-52956)

- Previously, a User-Defined Network (UDN) pod that is not qualified to use egress IP address had to use its own UDN pod IP address instead of its node IP address as the source IP address in egressing packets. With this release, the UDN pod network is advertised correctly. (OCPBUGS-50965)
- Previously, when the cluster certificate authority (CA) bundle was updated with a custom CA bundle, a delay occurred for the change to reflect in the **insights-runtime-extractor** container. This issue occurred if the Insights Operator gathered data after the CA bundle was updated. With this release, a fix removes the delay and this issue no longer occurs. (OCPBUGS-48790)
- Previously, in OpenShift Container Platform 4.17, a bug occurred because the code did not verify the load balancer IP address in the control plane subnet’s Classless Inter-Domain Routing (CIDR) range. This resulted in the IP address existing outside of the valid range and caused a 400 error during the installation. With this release, this fix prevents the 400 error caused by private IP address conflicts, ensuring a successful deployment of private OpenShift Container Platform clusters on Azure. (OCPBUGS-43724)

1.9.30.2. Updating

To update an OpenShift Container Platform 4.18 cluster to this latest release, see Updating a cluster using the CLI.

1.9.31. RHSA-2025:3577 - OpenShift Container Platform 4.18.8 bug fix update and security update

Issued: 10 April 2025

OpenShift Container Platform release 4.18.8 is now available. The list of bug fixes that are included in the update is documented in the RHSA-2025:3577 advisory. The RPM packages that are included in the update are provided by the RHBA-2025:3579 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 4.18.8 --pullspecs
```

1.9.31.1. Known issues

- IPsec is not supported on Red Hat Enterprise Linux (RHEL) compute nodes because of a **libreswan** incompatiblility issue between a host and an **ovn-ipsec** container that exist in each compute node. (OCPBUGS-52949).

1.9.31.2. Bug fixes

- Previously, sometimes network interface controllers (NICs) attached to virtual machines (VMs) that ran on Microsoft Azure failed because the NICs were 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. If a NIC is in **ProvisioningFailed** state, the VM now fails so that you have a better indication of the issue for troubleshooting purposes. (OCPBUGS-54355)

- Previously, selecting the **All projects** option on the **VolumeSnapshot** page of the web console in the **Administrator** perspective resulted in a **404: Page Not Found** error. With this release, a fix ensures that when you select the **All projects** option on the **VolumeSnapshot** page, the page shows results as expected without an error. (OCPBUGS-54269)

- Previously, the oc-mirror plugin v2 **delete** plugin had a typographical error in its **--help** argument output; the **--generate** listing stated **cahce** instead of **cache**. With this release, the typographical error is fixed to state **cache** in the **--generate** listing description. (OCPBUGS-54205)

- Previously, the output of the **oc-mirror --v2 version** command was missing the version information. With this release, the output from the command now correctly shows the version number. (OCPBUGS-53388)

- 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®, the following error occurred:

  ```
  ERROR Error: Plugin did not respond
  ERROR
  ERROR   with module.cis.ibm_cis_dns_record.kubernetes_api_internal[0],
  ERROR   on cis/main.tf line 27, in resource "ibm_cis_dns_record" "kubernetes_api_internal":
  ERROR   27: resource "ibm_cis_dns_record" "kubernetes_api_internal"
  ```

  With this release, you can use the installation program to create an external cluster on OpenShift Container Platform without the plugin issue. (OCPBUGS-53453)

- Previously, the Container Storage Interface (CSI) driver for the Google Cloud persistent disk (PD) did not support the **hyperdisk-balanced** volume type when the **ReadWriteMany** (RWX) access mode was set. If you attempted to provision a **hyperdisk-balanced** volume with this configuration, an error occurred that suggested mounting the volume with RWX access mode enabled was not possible. With this release, you can now mount a **hyperdisk-balanced** volume when the RWX access mode is enabled so that this issue no longer persists. See the Google Cloud documentation for further limitations when using Hyperdisk volumes in multi-writer mode. (OCPBUGS-44769)

1.9.31.3. Updating

To update an OpenShift Container Platform 4.18 cluster to this latest release, see Updating a cluster using the CLI.

1.9.32. RHBA-2025:3293 - OpenShift Container Platform 4.18.7 bug fix update

Issued: 3 April 2025

OpenShift Container Platform release 4.18.7 is now available. The list of bug fixes that are included in the update is documented in the RHBA-2025:3293 advisory. The RPM packages that are included in the update are provided by the RHBA-2025:3295 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 4.18.7 --pullspecs
```

1.9.32.1. Bug fixes

- Previously, when a proxy was configured, the installation program added the **machineNetwork** Classless Inter-Domain Routing (CIDR) to the **noProxy** field. If the **machineNetwork** CIDR was also configured by the user in **noProxy**, this resulted in a duplicate entry. A duplicate entry was not allowed by ignition and possibly prevented the host from booting properly. With this release, the fix ensures that the installation program does not add the **machineNetwork** CIDR to **noProxy** if it is already set. (OCPBUGS-53183)

- Previously, an **unable to read image** error message occurred when building the agent ISO in a disconnected setup. With this release, the error message does not appear. (OCPBUGS-52515)

- Previously, the code blocked the image import from a blocked registry. With this release, the image import from the registry is not blocked when the registry has mirrors configured. (OCPBUGS-52312)

1.9.32.2. Updating

To update an OpenShift Container Platform 4.18 cluster to this latest release, see Updating a cluster using the CLI.

1.9.33. RHSA-2025:3066 - OpenShift Container Platform 4.18.6 bug fix update and security update

Issued: 25 March 2025

OpenShift Container Platform release 4.18.6 is now available. The list of bug fixes that are included in the update is documented in the RHSA-2025:3066 advisory. The RPM packages that are included in the update are provided by the RHSA-2025:3068 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 4.18.6 --pullspecs
```

1.9.33.1. 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-53149) (OCPBUGS-53108)

- 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-53111)

- Previously, the **audit-logs** container that handles the **SIGTERM** signal timed out. Kubelet needed to send a hard termination signal (**SIGKILL**) to the **audit-logs** container to terminate the **SIGTERM** signal. With this release, a fix to a process ID (PID) alias means that audit log can properly handle a **SIGTERM** signal without the signal timing out. (OCPBUGS-52982)

- Previously, the **apply-bootstrap** container did not handle a **SIGTERM** signal correctly. The container waited for a sleep operation to complete before handling the signal, which then exceeded the **termination-grace-period** of the pod. This situation required a **SIGKILL** signal to force the shutdown operation and allow the pods to finish deletion. With this release, the **apply-bootstrap** container now handles the signal **SIGTERM** correctly so a correct graceful shutdown period happens without the need for a **SIGKILLED** signal. (OCPBUGS-52878)

- Previously, if you mirrored an empty catalog during a mirror-to-disk operation, and this caused the disk-to-mirror operation failed. This empty catalog was generated from an invalid Operator entry in the **ImageSetConfiguration** CR. With this release, you can no longer mirror an empty catalog so a disk-to-mirror operation can succeed. (OCPBUGS-52943)

- Previously, if you upgraded Google Cloud clusters that used a boot disk that was not compatible with UEFI, shielded VM support could not be enabled. This behavior prevented the creation of new machines. With this release, shielded VM support is disabled for disks that are known to be incompatible with UEFI. This change primarily affects customers who are upgrading from OpenShift Container Platform version 4.12 to 4.13 by using the Google Cloud marketplace images. (OCPBUGS-52495)

- Previously, node logs on the OpenShift Container Platform web console did not close when you clicked outside the node logs menu. With this release, the node logs menu now closes when you click outside the node logs menu. (OCPBUGS-52490)

- Previously, when you logged on to the Developer Sandbox from the OpenShift Container Platform web console, the web console ignored the path in the URL and displayed the **all projects** view on the Developer Sandbox instead of the namespace detailed in the URL. With this release, a fix corrects this behavior so the error no longer exists. (OCPBUGS-52406)

- Previously, the **capturegroup** inline diff algorithm in the **cluster-compare** tool failed to match the source text in an object with the **capturegroup** regular expression from a reference template. This issue existed if the source text had a similar structure to a regular expression. With this release, a fix to the **capturegroup** inline diff algorithm means that this matching issue no longer occurs. (OCPBUGS-51306)

- Previously, when you ran **oc-mirror** v2 on a continuous integration (CI) automation cycle and you viewed **oc-mirror** v2 logs on a non-TTY console, the output was missing progress information because of an issue with the progress bar implementation. With this release, **oc-mirror** v2 now disables the progress bar implementation and uses plain text logging instead for redirecting output so that the missing information no longer persists. (OCPBUGS-50996)

1.9.33.2. Updating

To update an OpenShift Container Platform 4.18 cluster to this latest release, see Updating a cluster using the CLI.

1.9.34. RHSA-2025:2705 - OpenShift Container Platform 4.18.5 bug fix update and security update

Issued: 18 March 2025

OpenShift Container Platform release 4.18.5 is now available. The list of bug fixes that are included in the update is documented in the RHSA-2025:2705 advisory. The RPM packages that are included in the update are provided by the RHBA-2025:2707 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 4.18.5 --pullspecs
```

1.9.34.1. Bug fixes

- Previously, during cluster creation, the Machine API started and managed machines in an installer-provisioned infrastructure deployed cluster on IBM Cloud. The Machine API detected an unhealthy control plane node and flagged it for deletion, which destroyed the cluster. With this release, during cluster creation, all control plane nodes are restored. (OCPBUGS-52872)

- Previously, the **managed-trust-bundle** volume mount and 'trusted-ca-bundle` config map were introduced as mandatory components. This requirement caused deployment failures for users who used their own public key infrastructure (PKI). The OpenShift Container Platform API server expected the **trust-ca-bundle-managed** config map. With this release, these components are optional, allowing clusters to deploy successfully without the **trusted-ca-bundle-managed** config map when the custom PKI is in use. (OCPBUGS-52516)

- Previously, etcd compaction blocked the process when it took more than 10 ms to process a batch. With this release, the issue is fixed and the etcd compaction proceeds as expected. (OCPBUGS-51971)

- Previously, Ampere ARM-based CPUs used a different CPU vendor identification identifier than other ARMs. The platform tuning matched the vendor identification and did not identify machines with ARM-based CPUs. With this release, the ARM detection is changed to use the architecture field, and machines with Ampere CPUs are properly tuned. (OCPBUGS-52484)

- Previously, when you ran the **openshift-install agent create pxe-files** command, it created a temporary directory. This directory was not removed when the command completed. With this release, the temporary directory is removed when the command is entered. (OCPBUGS-52429)

- Previously, the performance slowed down when **oc-mirror** started to use the Operator Lifecycle Manager (OLM) login to filter the catalog. With this release, this condition is resolved. (OCPBUGS-52350)

1.9.34.2. Updating

To update an OpenShift Container Platform 4.18 cluster to this latest release, see Updating a cluster using the CLI.

1.9.35. RHSA-2025:2449 - OpenShift Container Platform 4.18.4 bug fix update and security update

Issued: 11 March 2025

OpenShift Container Platform release 4.18.4 is now available. The list of bug fixes that are included in the update is documented in the RHSA-2025:2449 advisory. The RPM packages that are included in the update are provided by the RHBA-2025:2451 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 4.18.4 --pullspecs
```

1.9.35.1. Known issues

- There is currently a known issue where a Technology Preview-enabled cluster has Sigstore verification for payload images in **policy.json**, but the Podman version in the base image does not support Sigstore configuration. As a result, the new node is not available.
  Workaround: The node will start running when the Podman version in the base image does not support Sigstore. Use the default **policy.json** file that does not have Sigstore verification if the base image is 4.11 or earlier. (OCPBUGS-48296)

- There is currently a known issue where the data image is still present after deleting a related bare-metal host.
  Workaround: Delete the data image if it exists after the bare-metal host has been deleted. (OCPBUGS-45250)

1.9.35.2. Bug fixes

- Previously, you could not use catalog or bundle images that contained files with restricted extended attributes. With this release, the issue is resolved. (OCPBUGS-52173)

- Previously, you could not use the **registryOverride** option to override catalog Operator images. With this release, the logic for the control plane Operator is updated, and the issue is resolved. (OCPBUGS-51375)

- Previously, the Installer failed to retrieve Google Cloud Platform (GCP) tags over an unstable network or when it could not reach the GCP server. With this release, the issue is resolved. (OCPBUGS-51211)

- 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-51149)

- Previously, the **Observe** section on the web console did not show items contributed from plugins unless certain flags related to 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-51086)

- Previously, if **kubevirt** and **graphImage** images were not retrieved during the collection phase of **oc-mirror**, the run would succeed even if the images were missing. With this release, the **oc-mirror** run fails as expected if the images are not found. (OCPBUGS-50981)

- Previously, an issue prevented the configuration of multiple subnets in the failure domain of a Nutanix cluster during installation. With this release, the issue is resolved. (OCPBUGS-49885)

- Previously, a new Ingress Controller API was added to manage the **idle-close-on-response** HAProxy setting: **IdleConnectionTerminationPolicy**. If a cluster did not have the **IdleConnectionTerminationPolicy** API field, the **idle-close-on-response** setting was enabled unconditionally. With this release, the default value is **Deferred**, and the issue is resolved. (OCPBUGS-48377)

1.9.35.3. Updating

To update an OpenShift Container Platform 4.18 cluster to this latest release, see Updating a cluster using the CLI.

1.9.36. RHBA-2025:2229 - OpenShift Container Platform 4.18.3 bug fix update

Issued: 6 March 2025

OpenShift Container Platform release 4.18.3 is now available. The list of bug fixes that are included in the update is documented in the RHBA-2025:2229 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 4.18.3 --pullspecs
```

1.9.36.1. Updating

To update an OpenShift Container Platform 4.18 cluster to this latest release, see Updating a cluster using the CLI.

1.9.37. RHBA-2025:1904 - OpenShift Container Platform 4.18.2 image release, bug fix, and security update advisory

Issued: 4 March 2025

OpenShift Container Platform release 4.18.2, which includes security updates, is now available. The list of bug fixes that are included in the update is documented in the RHBA-2025:1904 advisory. The RPM packages that are included in the update are provided by the RHSA-2025:1908 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 4.18.2 --pullspecs
```

1.9.37.1. Bug fixes

- Previously, when using the **dry-run** argument with the oc-mirror v2 command, the **cluster-resources** directory was cleared in error. As a result, files that were generated from a previous mirroring operation, such as **idms-oc-mirror.yaml** and **itms-oc-mirror.yaml**, were deleted. With this release, the **cluster-resources** directory is no longer cleared when you add the **dry-run** argument to the oc-mirror v2 command. (OCPBUGS-51185)

- Previously, the Operator Controller would not accept live updates to registries with a proxy configuration. Unless the controller pods were restarted, this issue caused OLM v1 to read the wrong image URL. With this release, a fix to the Operator Controller means that it accepts live updates to registries with a proxy configuration and the controller pod no longer needs to be restarted. (OCPBUGS-51140)

- Previously, Internet Small Computer System Interface (iSCSI) and Fibre Channel devices that were attached to a multipath device did not resolve correctly when these devices were partitioned. With this relase, a fix ensures that partitioned multipath storage devices can now correctly resolve. (OCPBUGS-51100)

- Previously, mirroring of some Operators from catalog resources caused oc-mirror v1 to fail with an error message that indicated an issue with the **ocischema.DeserializedImageIndex** manifest file. With this release, oc-mirror v1 can handle the **ocischema.DeserializedImageIndex** manfiest file so that this issue no longer occurs. (OCPBUGS-51099)

- Previously, when you created a cluster with secure proxy enabled and the certificate configuration is set to **configuration.proxy.trustCA**, the cluster installation failed. Additionally, the OpenShift OAuth API server could not use the management cluster proxy to reach cloud APIs. With this release, fixes are in place to prevent these issues. (OCPBUGS-51050)

- Previously, when you deleted Dynamic Host Configuration Protocol (DHCP) network on a IBM Power Virtual Server cluster, subresources would still exist. With this release, when you delete a DHCP network, the subresources are also deleted. (OCPBUGS-50870)

- Previously, when you deleted Dynamic Host Configuration Protocol (DHCP) network on an IBM Power Virtual Server cluster, subresources could still exist. With this release, when you delete a DHCP network, the subresources deletion now occur before continuing the destroy operation. (OCPBUGS-50870)

- Previously, the **vmware-vsphere-csi-driver-operator** Container Storage Interface (CSI) driver entered panic mode when the VMware vCenter address was incorrect or missing. With this release, the CSI driver does not go into panic mode if the VMware vCenter address is incorrect or missing. (OCPBUGS-50638)

- Previously, when you used the Agent-based Installer to install a cluster on a host, sometimes the **/dev/sda** device, an Extensible Firmware Interface (EFI) device, failed to mount. With this release, a retry operation is added to the EFI device so it mounts correctly. (OCPBUGS-50621)

- 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-50550)

- 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-49921)

- Previously, when the **ClusterNetwork** classless inter-domain routing (CIDR) mask value is greater than the **hostPrefix** value and the **networking.ovnKubernetesConfig.ipv4.internalJoinSubnet** section is provided in the **install-config.yaml** file, the installation program failed a validation check and returned a Golang runtime error. With this release, the installation program still fails the validation check and now outputs a descriptive error message that indicates the invalid **hostPrefix** value. (OCPBUGS-49864)

- Previously, the router incorrectly assumed that only **SHA1** leaf certificates were rejected by HAProxy. This caused the router to fail as it rejected **SHA1** intermediate certificates. With this release, the router now inspects all non-self-signed certificates and rejects any that use **SHA1**. The router no longer crashes because of the existence of **SHA1** intermediate certificates. Self-signed **SHA1** certificates are no longer rejected. Root CAs can continue to use **SHA1**. (OCPBUGS-49389)

- Previously, Google Cloud did not include a **wait** operation for API calls that destroyed clutser resources. In certain situations, this missing operation caused the installation program to not delete the backend services. With this release, Google Cloud adds a **wait** operation to an API call so that the installation program can delete backend services. (OCPBUGS-49320)

- Previously, on the **Operator Details** page on the web console, **ClusterServiceVersion** (CSV) details did not render. With this release, the CSV details now render on the **Operator Details** page. (OCPBUGS-48736)

- Previously, bundle properties occasionally did not get propagated to the annotations on Helm charts that were created during an Operator installation. With this release, properties are now taken from both the CSV of the bundle and the **metadata.yaml** file or the **properties.yaml** file so that this issue no longer exists. (OCPBUGS-45114)

- 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-51056)

- 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-50512)

1.9.37.2. Updating

To update an OpenShift Container Platform 4.18 cluster to this latest release, see Updating a cluster using the CLI.

1.9.38. RHSA-2024:6122 - OpenShift Container Platform 4.18.1 image release, bug fix, and security update advisory

Issued: 25 February 2025

OpenShift Container Platform release 4.18.1, which includes security updates, is now available. The list of bug fixes that are included in the update is documented in the RHSA-2024:6122 advisory. The RPM packages that are included in the update are provided by the RHEA-2024:6126 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 4.18.1 --pullspecs
```

1.9.38.1. Updating

To update an OpenShift Container Platform 4.17 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 4.18 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

K

Kube Descheduler Operator
Red Hat build of Kueue

L

Leader Worker Set Operator
Logging

M

Migration Toolkit for Containers (MTC)

N

Network Observability Operator

Network-bound Disk Encryption (NBDE) Tang Server Operator

O

OpenShift API for Data Protection (OADP)

Red Hat OpenShift Dev Spaces

Red Hat OpenShift Distributed Tracing Platform

Red Hat OpenShift GitOps

Red Hat OpenShift Local (Upstream CRC documentation)

Red Hat OpenShift Pipelines

OpenShift sandboxed containers

Red Hat OpenShift Serverless

Red Hat OpenShift Service Mesh 2.x

Red Hat OpenShift Service Mesh 3.x

Red Hat OpenShift support for Windows Containers

Red Hat OpenShift Virtualization

Red Hat build of OpenTelemetry

P

Power monitoring for Red Hat OpenShift

R

Run Once Duration Override Operator

S

Secondary Scheduler Operator for Red Hat OpenShift

Security Profiles Operator