



OpenShift Container Platform 4.12

Release notes

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

OpenShift Container Platform 4.12 Release notes

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

Legal Notice

Copyright © 2025 Red Hat, Inc.

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

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

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

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

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

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

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

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

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

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

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

All other trademarks are the property of their respective owners.

Abstract

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

Table of Contents

CHAPTER 1. OPENSIFT CONTAINER PLATFORM 4.12 RELEASE NOTES	12
1.1. ABOUT THIS RELEASE	12
1.2. OPENSIFT CONTAINER PLATFORM LAYERED AND DEPENDENT COMPONENT SUPPORT AND COMPATIBILITY	12
1.3. NEW FEATURES AND ENHANCEMENTS	12
1.3.1. Red Hat Enterprise Linux CoreOS (RHCOS)	12
1.3.1.1. Default consoles for new clusters are now determined by the installation platform	13
1.3.1.2. IBM Secure Execution on IBM Z and LinuxONE (Technology Preview)	13
1.3.1.3. RHCOS now uses RHEL 8.6	13
1.3.2. Installation and upgrade	13
1.3.2.1. Assisted Installer SaaS provides platform integration support for Nutanix	13
1.3.2.2. Specify the load balancer type in AWS during installation	13
1.3.2.3. Extend worker nodes to the edge of AWS when installing into an existing Virtual Private Cloud (VPC) with Local Zone subnets.	14
1.3.2.4. Google Cloud Platform Marketplace offering	14
1.3.2.5. Troubleshooting bootstrap failures during installation on GCP and Azure	14
1.3.2.6. IBM Cloud VPC general availability	14
1.3.2.7. Required administrator acknowledgment when upgrading from OpenShift Container Platform 4.11 to 4.12	14
1.3.2.8. Enabling a feature set when installing a cluster	15
1.3.2.9. OpenShift Container Platform on ARM	15
1.3.2.10. Mirroring file-based catalog Operator images in OCI format with the oc-mirror CLI plugin (Technology Preview)	15
1.3.2.11. Installing an OpenShift Container Platform cluster on GCP into a shared VPC (Technology Preview)	15
1.3.2.12. Consistent IP address for Ironic API in bare-metal installations without a provisioning network	15
1.3.2.13. Installing OpenShift Container Platform on GCP using service account authentication	15
1.3.2.14. propagateUserTags parameter for AWS resources provisioned by the OpenShift Container Platform cluster	16
1.3.2.15. Ironic container images use RHEL 9 base image	16
1.3.2.16. Cloud provider configuration updates for clusters that run on RHOSP	16
1.3.2.17. Support for workloads on RHOSP distributed compute nodes	16
1.3.2.18. OpenShift Container Platform on AWS Outposts (Technology Preview)	16
1.3.2.19. Agent-based installation supports two input modes	16
1.3.2.20. Agent-based installation supports installing OpenShift Container Platform clusters in FIPS compliant mode	17
1.3.2.21. Deploy an Agent-based OpenShift Container Platform cluster in a disconnected environment	17
1.3.2.22. Explanation of field to build and push graph-data	17
1.3.2.23. Agent-based installation supports single and dual stack networking	17
1.3.2.24. Agent deployed OpenShift Container Platform cluster can be used as a hub cluster	17
1.3.2.25. Agent-based installation performs installation validations	18
1.3.2.26. Configure static networking in a Agent-based installation	18
1.3.2.27. CLI based automated deployment in an Agent-based installation	18
1.3.2.28. Agent-based installation supports host specific configuration at the instalation time	18
1.3.2.29. Agent-based installation supports DHCP	18
1.3.2.30. Installing a cluster on Nutanix with limited internet access	18
1.3.3. Post-installation configuration	19
1.3.3.1. CSI driver installation on vSphere clusters	19
1.3.3.2. Cluster Capabilities	19
1.3.3.3. OpenShift Container Platform with multi-architecture compute machines (Technology Preview)	19
1.3.4. Web console	20
1.3.4.1. Administrator Perspective	20

1.3.4.1.1. Multi-architecture compute machines on the OpenShift Container Platform web console	20
1.3.4.1.2. Dynamic plugin generally available	20
1.3.4.2. Developer Perspective	20
1.3.4.2.1. Helm page improvements	21
1.3.4.2.2. Negative matchers in Alertmanager	21
1.3.5. OpenShift CLI (oc)	22
1.3.5.1. Managing plugins for the OpenShift CLI with Krew (Technology Preview)	22
1.3.6. IBM Z and LinuxONE	22
Notable enhancements	22
IBM Secure Execution (Technology Preview)	22
Supported features	22
Restrictions	24
1.3.7. IBM Power	24
Notable enhancements	24
Supported features	25
Restrictions	26
1.3.8. Images	27
1.3.9. Security and compliance	27
1.3.9.1. Security Profiles Operator	27
1.3.10. Networking	27
1.3.10.1. Support for dual-stack addressing for the API VIP and Ingress VIP	27
1.3.10.2. Red Hat OpenShift Networking	27
1.3.10.3. OVN-Kubernetes is now the default networking plugin	28
1.3.10.4. Ingress Node Firewall Operator	28
1.3.10.5. Enhancements to networking metrics	29
1.3.10.6. Multi-zone Installer Provisioned Infrastructure VMware vSphere installation (Technology Preview)	29
1.3.10.7. Kubernetes NMState in VMware vSphere now supported	29
1.3.10.8. Kubernetes NMState in OpenStack now supported	29
1.3.10.9. External DNS Operator	30
1.3.10.10. The ingressClassName field is required for each ingress object	30
1.3.10.11. Capturing metrics and telemetry associated with the use of routes and shards	30
1.3.10.12. AWS Load Balancer Operator	30
1.3.10.13. Ingress Controller Autoscaling (Technology Preview)	31
1.3.10.14. HAProxy maxConnections default is now 50,000	31
1.3.10.15. Configuration of an Ingress Controller for manual DNS management	31
1.3.10.16. Supported hardware for SR-IOV (Single Root I/O Virtualization)	31
1.3.10.17. Supported hardware for OvS (Open vSwitch) Hardware Offload	31
1.3.10.18. Multi-network-policy supported for SR-IOV (Technology Preview)	32
1.3.10.19. Switch between AWS load balancer types without deleting the Ingress Controller	32
1.3.10.20. IPv6 unsolicited neighbor advertisements and IPv4 gratuitous address resolution protocol now default on the SR-IOV CNI plugin	32
1.3.10.21. Support for CoreDNS cache tuning	32
1.3.10.22. OVN-Kubernetes supports configuration of internal subnet	32
1.3.10.23. Egress IP support on Red Hat OpenStack Platform (RHOSP)	32
1.3.10.24. OpenShift SDN to OVN-Kubernetes feature migration support	33
1.3.10.25. Egress firewall audit logging	33
1.3.10.26. Advertise MetalLB from a given address pool from a subset of nodes	33
1.3.10.27. Additional deployment specifications for MetalLB	33
1.3.10.28. Node IP selection improvements	33
1.3.10.29. CoreDNS update to version 1.10.0	34
1.3.10.30. Support for a configurable reload interval in HAProxy	34
1.3.10.31. Network Observability Operator updates	34

1.3.10.32. IPv6 for secondary network interfaces on RHOSP	34
1.3.10.33. UDP support for load balancers on RHOSP	34
1.3.10.34. Deploy the SR-IOV Operator for hosted control planes (Technology Preview)	34
1.3.10.35. Support for IPv6 virtual IP (VIP) addresses for the Ingress VIP and API VIP services on bare metal	34
1.3.10.36. Support for switching the Bluefield-2 network device from data processing unit (DPU) mode to network interface controller (NIC) mode (Technology Preview)	35
1.3.10.37. Support for enabling hybrid networking after cluster installation	35
1.3.10.38. Support for allocateLoadBalancerNodePorts in the Network API service object	35
1.3.11. Storage	35
1.3.11.1. Persistent storage using the GCP Filestore Driver Operator (Technology Preview)	35
1.3.11.2. Automatic CSI migration for AWS Elastic Block Storage auto migration is generally available	35
1.3.11.3. Automatic CSI migration for GCP PD auto migration is generally available	36
1.3.11.4. Updating from OpenShift Container Platform 4.12 to 4.13 and later with vSphere in-tree PVs	36
1.3.11.5. Storage capacity tracking for pod scheduling is generally available	36
1.3.11.6. VMware vSphere CSI topology is generally available	36
1.3.11.7. Local ephemeral storage resource management is generally available	36
1.3.11.8. Volume populators (Technology Preview)	37
1.3.11.9. VMware vSphere CSI Driver Operator requirements	37
1.3.11.10. Azure File supporting NFS is generally available	37
1.3.12. Operator lifecycle	37
1.3.12.1. Platform Operators (Technology Preview)	37
1.3.12.2. Controlling where an Operator is installed	38
1.3.12.3. Pod security admission synchronization for user-created openshift-* namespaces	38
1.3.13. Operator development	38
1.3.13.1. Configuring the security context of a catalog pod	38
1.3.13.2. Validating bundle manifests for APIs removed from Kubernetes 1.25	38
1.3.14. Machine API	39
1.3.14.1. Control plane machine sets	39
1.3.14.2. Specifying cluster autoscaler log level verbosity	39
1.3.14.3. Enabling Azure boot diagnostics	39
1.3.15. Machine Config Operator	39
1.3.15.1. RHCOS image layering	39
1.3.16. Nodes	39
1.3.16.1. Updating the interface-specific safe list (Technology Preview)	39
1.3.16.2. Cron job time zones (Technology Preview)	39
1.3.16.3. Linux Control Group version 2 promoted to Technology Preview	40
1.3.16.4. crun container runtime (Technology Preview)	40
1.3.16.5. Self Node Remediation Operator enhancements	40
1.3.16.6. Node Health Check Operator enhancements	40
1.3.17. Monitoring	40
1.3.17.1. Updates to monitoring stack components and dependencies	40
1.3.17.2. Changes to alerting rules	41
1.3.17.3. New option to specify pod topology spread constraints for monitoring components	41
1.3.17.4. New option to improve data consistency for Prometheus Adapter	41
1.3.17.5. Update to Alertmanager configuration for additional secret keys	41
1.3.18. New Network Observability Operator	42
1.3.19. Scalability and performance	42
1.3.19.1. Disabling realtime using workload hints removes Receive Packet Steering from the cluster	42
1.3.19.2. Tuned profile	42
1.3.19.3. Support for new kernel features and options	42
1.3.19.4. Power-saving configurations	42
1.3.19.5. Expanding Single-node OpenShift clusters with worker nodes using GitOps ZTP (Technology	42

Preview)	43
1.3.19.6. Factory-precaching-cli tool to reduce OpenShift Container Platform and Operator deployment times (Technology Preview)	43
1.3.19.7. Zero touch provisioning (ZTP) integration of the factory-precaching-cli tool (Technology Preview)	43
1.3.19.8. Node tuning in a hosted cluster (Technology Preview)	43
1.3.19.9. Kernel module management Operator	43
1.3.19.10. Hub and spoke cluster support (Technology Preview)	43
1.3.19.11. Topology Aware Lifecycle Manager (TALM)	44
1.3.19.12. Mount namespace encapsulation (Technology Preview)	44
1.3.19.13. Changing the workload partitioning CPU set in single-node OpenShift clusters that are deployed with GitOps ZTP	44
1.3.19.14. RHACM hub template functions now available for use with GitOps ZTP	44
1.3.19.15. ArgoCD managed cluster limits	44
1.3.19.16. GitOps ZTP support for configuring policy compliance evaluation timeouts in PolicyGenTemplate CRs	45
1.3.19.17. Specifying the platform type for managed clusters	45
1.3.19.18. Configuring the hub cluster to use unauthenticated registries	45
1.3.19.19. Ironic agent mirroring in disconnected GitOps ZTP installations	45
1.3.19.20. Configuring kernel arguments for the Discovery ISO by using GitOps ZTP	45
1.3.19.21. Deploy heterogeneous spoke clusters from a hub cluster	46
1.3.19.22. HTTP transport replaces AMQP for PTP and bare-metal events (Technology Preview)	46
1.3.20. Insights Operator	46
1.3.20.1. Insights alerts	46
1.3.20.2. Insights Operator data collection enhancements	46
1.3.21. Authentication and authorization	46
1.3.21.1. Application credentials on RHOSP	46
1.3.22. Hosted control planes (Technology Preview)	47
1.3.22.1. HyperShift API beta release now available	47
1.3.22.2. Versioning for hosted control planes	47
1.3.22.3. Backing up and restoring etcd on a hosted cluster	47
1.3.22.4. Disaster recovery for a hosted cluster within an AWS region	47
1.3.23. Red Hat Virtualization (RHV)	48
1.4. NOTABLE TECHNICAL CHANGES	48
AWS Security Token Service regional endpoints	48
cert-manager Operator general availability	48
Credentials requests directory parameter for deleting GCP resources with the Cloud Credential Operator utility	48
Future restricted enforcement for pod security admission	48
Catalog sources and restricted pod security admission enforcement	49
Operator SDK 1.25.4	49
LVM Operator is now called Logical Volume Manager Storage	49
End of support for RHOSP 16.1	50
1.5. DEPRECATED AND REMOVED FEATURES	50
Operator deprecated and removed features	50
Images deprecated and removed features	50
Monitoring deprecated and removed features	50
Installation deprecated and removed features	51
Updating clusters deprecated and removed features	51
Storage deprecated and removed features	51
Authentication and authorization deprecated and removed features	51
Specialized hardware and driver enablement deprecated and removed features	52
Multi-architecture deprecated and removed features	52

Networking deprecated and removed features	52
Web console deprecated and removed features	53
1.5.1. Deprecated features	53
1.5.1.1. Red Hat Virtualization (RHV) as a host platform for OpenShift Container Platform will be deprecated	53
1.5.1.2. Wildcard DNS queries for the cluster.local domain are deprecated	53
1.5.1.3. Specific hardware models on ppc64le, s390x, and x86_64 v1 CPU architectures are deprecated	53
1.5.1.4. Kuryr support for clusters that run on RHOSP	54
1.5.2. Removed features	54
1.5.2.1. Beta APIs removed from Kubernetes 1.25	54
1.5.2.2. Empty file and stdout support for the oc registry login command	54
1.5.2.3. RHEL 7 support for the OpenShift CLI (oc) has been removed	54
1.5.2.4. OpenShift CLI (oc) commands have been removed	54
1.5.2.5. Grafana component removed from monitoring stack	55
1.5.2.6. Prometheus and Grafana user interface access removed from monitoring stack	55
1.5.2.7. Support for virtual hardware version 13 is removed	55
1.5.2.8. Support for snapshot v1beta1 API endpoint is removed	55
1.5.2.9. Support for manually deploying a custom scheduler has been removed	55
1.5.2.10. Support for deploying single-node OpenShift with OpenShiftSDN has been removed	55
1.5.2.11. Removal of Jenkins images from install payload	55
1.5.3. Future Kubernetes API removals	56
1.6. BUG FIXES	56
API Server and Authentication	56
Bare Metal Hardware Provisioning	56
Builds	57
Cloud Compute	57
Developer Console	60
Image Registry	60
Installer	60
Kubernetes Controller Manager	62
Kubernetes Scheduler	62
Machine Config Operator	62
Management Console	63
Monitoring	64
Networking	65
Node	67
OpenShift CLI (oc)	67
Operator Lifecycle Manager (OLM)	67
Operator SDK	68
File Integrity Operator	68
Compliance Operator	69
OpenShift API server	69
Red Hat Enterprise Linux CoreOS (RHCOS)	69
Scalability and performance	70
Storage	71
Web console (Developer perspective)	71
1.7. TECHNOLOGY PREVIEW FEATURES	72
Networking Technology Preview features	72
Storage Technology Preview features	73
Installation Technology Preview features	74
Node Technology Preview features	75
Multi-Architecture Technology Preview features	75
Serverless Technology Preview features	76

Specialized hardware and driver enablement Technology Preview features	76
Web console Technology Preview features	76
Scalability and performance Technology Preview features	77
Operator Technology Preview features	77
Monitoring Technology Preview features	78
Red Hat OpenStack Platform (RHOSP) Technology Preview features	78
Architecture Technology Preview features	78
Machine management Technology Preview features	79
Authentication and authorization Technology Preview features	79
1.8. KNOWN ISSUES	80
1.9. ASYNCHRONOUS ERRATA UPDATES	89
1.9.1. RHSA-2022:7399 - OpenShift Container Platform 4.12.0 image release, bug fix, and security update advisory	90
1.9.1.1. Features	90
1.9.1.1.1. General availability of pod-level bonding for secondary networks	90
1.9.2. RHSA-2023:0449 - OpenShift Container Platform 4.12.1 bug fix and security update	90
1.9.2.1. Bug fixes	91
1.9.2.2. Updating	91
1.9.3. RHSA-2023:0569 - OpenShift Container Platform 4.12.2 bug fix and security update	91
1.9.3.1. Updating	91
1.9.4. RHSA-2023:0728 - OpenShift Container Platform 4.12.3 bug fix and security update	91
1.9.4.1. Bug fixes	91
1.9.4.2. Updating	92
1.9.5. RHSA-2023:0769 - OpenShift Container Platform 4.12.4 bug fix and security update	92
1.9.5.1. Updating	92
1.9.6. RHSA-2023:0890 - OpenShift Container Platform 4.12.5 bug fix and security update	92
1.9.6.1. Bug fixes	92
1.9.6.2. Updating	93
1.9.7. RHSA-2023:1034 - OpenShift Container Platform 4.12.6 bug fix and security update	93
1.9.7.1. Updating	93
1.9.8. RHBA-2023:1163 - OpenShift Container Platform 4.12.7 bug fix update	93
1.9.8.1. Updating	93
1.9.9. RHBA-2023:1269 - OpenShift Container Platform 4.12.8 bug fix and security update	93
1.9.9.1. Updating	94
1.9.10. RHSA-2023:1409 - OpenShift Container Platform 4.12.9 bug fix and security update	94
1.9.10.1. Bug fixes	94
1.9.10.2. Updating	94
1.9.11. RHBA-2023:1508 - OpenShift Container Platform 4.12.10 bug fix update	94
1.9.11.1. Updating	95
1.9.12. RHSA-2023:1645 - OpenShift Container Platform 4.12.11 bug fix and security update	95
1.9.12.1. Features	95
1.9.12.1.1. New flag for the oc-mirror plugin: --max-nested-paths	95
1.9.12.1.2. New flag for the oc-mirror plugin: --skip-pruning	95
1.9.12.2. Bug fixes	95
1.9.12.3. Updating	95
1.9.13. RHBA-2023:1734 - OpenShift Container Platform 4.12.12 bug fix	96
1.9.13.1. Updating	96
1.9.14. RHBA-2023:1750 - OpenShift Container Platform 4.12.13 bug fix update	96
1.9.14.1. Features	96
1.9.14.1.1. Pod security admission restricted enforcement (Technology Preview)	96
1.9.14.2. Updating	96
1.9.15. RHBA-2023:1858 - OpenShift Container Platform 4.12.14 bug fix update	97
1.9.15.1. Features	97

1.9.15.1.1. Cloud provider OpenStack is updated to 1.25	97
1.9.15.2. Updating	97
1.9.16. RHBA-2023:2037 - OpenShift Container Platform 4.12.15 bug fix update	97
1.9.16.1. Bug fixes	97
1.9.16.2. Updating	97
1.9.17. RHSA-2023:2110 - OpenShift Container Platform 4.12.16 bug fix and security update	98
1.9.17.1. Bug fixes	98
1.9.17.2. Updating	98
1.9.18. RHBA-2023:2699 - OpenShift Container Platform 4.12.17 bug fix update	98
1.9.18.1. Bug fixes	98
1.9.18.2. Updating	98
1.9.19. RHBA-2023:3208 - OpenShift Container Platform 4.12.18 bug fix update	98
1.9.19.1. Bug fixes	99
1.9.19.2. Updating	99
1.9.20. RHSA-2023:3287 - OpenShift Container Platform 4.12.19 bug fix and security update	99
1.9.20.1. Updating	99
1.9.21. RHSA-2023:3410 - OpenShift Container Platform 4.12.20 bug fix update	99
1.9.21.1. Bug fixes	100
1.9.21.2. Updating	100
1.9.22. RHBA-2023:3546 - OpenShift Container Platform 4.12.21 bug fix and security update	100
1.9.22.1. Bug fixes	100
1.9.22.2. Updating	100
1.9.23. RHSA-2023:3615 - OpenShift Container Platform 4.12.22 bug fix and security update	100
1.9.23.1. Bug fixes	101
1.9.23.2. Updating	101
1.9.24. RHSA-2023:3925 - OpenShift Container Platform 4.12.23 bug fix and security update	101
1.9.24.1. Updating	101
1.9.25. RHBA-2023:3977 - OpenShift Container Platform 4.12.24 bug fix and security update	102
1.9.25.1. Features	102
1.9.25.2. NUMA-aware scheduling with the NUMA Resources Operator is generally available	102
1.9.25.3. Updating	102
1.9.26. RHBA-2023:4048 - OpenShift Container Platform 4.12.25 bug fix update	102
1.9.26.1. Updating	103
1.9.27. RHBA-2023:4221 - OpenShift Container Platform 4.12.26 bug fix update	103
1.9.27.1. Updating	103
1.9.28. RHBA-2023:4319 - OpenShift Container Platform 4.12.27 bug fix update	103
1.9.28.1. Updating	103
1.9.29. RHBA-2023:4440 - OpenShift Container Platform 4.12.28 bug fix update	103
1.9.29.1. Updating	104
1.9.30. RHBA-2023:4608 - OpenShift Container Platform 4.12.29 bug fix update	104
1.9.30.1. Bug fixes	104
1.9.30.2. Updating	104
1.9.31. RHSA-2023:4671 - OpenShift Container Platform 4.12.30 bug fix update	104
1.9.31.1. Updating	104
1.9.32. RHBA-2023:4756 - OpenShift Container Platform 4.12.31 bug fix update	105
1.9.32.1. Updating	105
1.9.33. RHBA-2023:4900 - OpenShift Container Platform 4.12.32 bug fix update	105
1.9.33.1. Bug fix	105
1.9.33.2. Updating	105
1.9.34. RHBA-2023:5016 - OpenShift Container Platform 4.12.33 bug fix update	105
1.9.34.1. Updating	106
1.9.35. RHBA-2023:5151 - OpenShift Container Platform 4.12.34 bug fix update	106
1.9.35.1. Bug fixes	106

1.9.35.2. Updating	106
1.9.36. RHBA-2023:5321 - OpenShift Container Platform 4.12.35 bug fix update	106
1.9.36.1. Updating	107
1.9.37. RHSA-2023:5390 - OpenShift Container Platform 4.12.36 bug fix and security update	107
1.9.37.1. Updating	107
1.9.38. RHBA-2023:5450 - OpenShift Container Platform 4.12.37 bug fix update	107
1.9.38.1. Updating	107
1.9.39. RHSA-2023:5677 - OpenShift Container Platform 4.12.39 bug fix and security update	107
1.9.39.1. Bug fixes	108
1.9.39.2. Updating	108
1.9.40. RHSA-2023:5896 - OpenShift Container Platform 4.12.40 bug fix and security update	108
1.9.40.1. Updating	108
1.9.41. RHSA-2023:6126 - OpenShift Container Platform 4.12.41 bug fix and security update	108
1.9.41.1. Updating	108
1.9.42. RHSA-2023:6276 - OpenShift Container Platform 4.12.42 bug fix and security update	108
1.9.42.1. Feature	109
1.9.42.1.1. APIServer.config.openshift.io is now tracked by Insights Operator	109
1.9.42.2. Bug fixes	109
1.9.42.3. Updating	109
1.9.43. RHSA-2023:6842 - OpenShift Container Platform 4.12.43 bug fix and security update	109
1.9.43.1. Updating	109
1.9.44. RHSA-2023:6894 - OpenShift Container Platform 4.12.44 bug fix and security update	110
1.9.44.1. Updating	110
1.9.45. RHSA-2023:7608 - OpenShift Container Platform 4.12.45 bug fix and security update	110
1.9.45.1. Bug fixes	110
1.9.45.2. Updating	110
1.9.46. RHSA-2023:7823 - OpenShift Container Platform 4.12.46 bug fix and security update	111
1.9.46.1. Bug fixes	111
1.9.46.2. Updating	111
1.9.47. RHSA-2024:0198 - OpenShift Container Platform 4.12.47 bug fix and security update	111
1.9.47.1. Bug fixes	111
1.9.47.2. Updating	111
1.9.48. RHSA-2024:0485 - OpenShift Container Platform 4.12.48 bug fix and security update	112
1.9.48.1. Updating	112
1.9.49. RHSA-2024:0664 - OpenShift Container Platform 4.12.49 bug fix and security update	112
1.9.49.1. Bug fixes	112
1.9.49.2. Updating	112
1.9.50. RHSA-2024:0833 - OpenShift Container Platform 4.12.50 bug fix and security update	112
1.9.50.1. Bug fixes	113
1.9.50.2. Updating	113
1.9.51. RHSA-2024:1052 - OpenShift Container Platform 4.12.51 bug fix and security update	113
1.9.51.1. Bug fixes	113
1.9.51.2. Updating	113
1.9.52. RHSA-2024:1265 - OpenShift Container Platform 4.12.53 bug fix update	113
1.9.52.1. Updating	114
1.9.53. RHSA-2024:1572 - OpenShift Container Platform 4.12.54 bug fix and security update	114
1.9.53.1. Updating	114
1.9.54. RHSA-2024:1679 - OpenShift Container Platform 4.12.55 bug fix and security update	114
1.9.54.1. Bug fixes	114
1.9.54.2. Updating	114
1.9.55. RHSA-2024:1896 - OpenShift Container Platform 4.12.56 bug fix and security update	115
1.9.55.1. Bug fixes	115
1.9.55.2. Updating	115

1.9.56. RHSA-2024:2782 - OpenShift Container Platform 4.12.57 bug fix and security update	115
1.9.56.1. Bug fixes	115
1.9.56.2. Updating	116
1.9.57. RHSA-2024:3349 - OpenShift Container Platform 4.12.58 bug fix and security update	116
1.9.57.1. Bug fixes	116
1.9.57.2. Known issue	116
1.9.57.3. Updating	117
1.9.58. RHSA-2024:3713 - OpenShift Container Platform 4.12.59 bug fix and security update	117
1.9.58.1. Bug fixes	117
1.9.58.2. Updating	117
1.9.59. RHSA-2024:4006 - OpenShift Container Platform 4.12.60 bug fix and security update	117
1.9.59.1. Bug fixes	118
1.9.59.2. Updating	118
1.9.60. RHSA-2024:4677 - OpenShift Container Platform 4.12.61 bug fix and security update	118
1.9.60.1. Bug fixes	118
1.9.60.2. Enhancements	119
1.9.60.3. Updating	119
1.9.61. RHSA-2024:5200 - OpenShift Container Platform 4.12.63 bug fix and security update	119
1.9.61.1. Updating	119
1.9.62. RHSA-2024:5808 - OpenShift Container Platform 4.12.64 bug fix and security update	119
1.9.62.1. Updating	120
1.9.63. RHSA-2024:6642 - OpenShift Container Platform 4.12.65 bug fix and security update	120
1.9.63.1. Bug fixes	120
1.9.63.2. Updating	120
1.9.64. RHSA-2024:6705 - OpenShift Container Platform 4.12.66 bug fix and security update	120
1.9.64.1. Updating	120
1.9.65. RHSA-2024:7590 - OpenShift Container Platform 4.12.67 bug fix and security update	120
1.9.65.1. Bug fixes	121
1.9.65.2. Updating	121
1.9.66. RHSA-2024:8692 - OpenShift Container Platform 4.12.68 bug fix and security update	121
1.9.66.1. Bug fixes	121
1.9.66.2. Updating	122
1.9.67. RHBA-2024:8996 - OpenShift Container Platform 4.12.69 bug fix	122
1.9.67.1. Updating	122
1.9.68. RHBA-2024:10533 - OpenShift Container Platform 4.12.70 bug fix and security update	122
1.9.68.1. Updating	122
1.9.69. RHSA-2025:0014 - OpenShift Container Platform 4.12.71 bug fix and security update	122
1.9.69.1. Updating	123
1.9.70. RHSA-2025:0832 - OpenShift Container Platform 4.12.72 bug fix and security update	123
1.9.70.1. Bug fixes	123
1.9.70.2. Updating	123
1.9.71. RHSA-2025:1242 - OpenShift Container Platform 4.12.73 bug fix and security update	123
1.9.71.1. Bug fixes	124
1.9.71.2. Updating	124
1.9.72. RHSA-2025:2441 - OpenShift Container Platform 4.12.74 bug fix and security update	124
1.9.72.1. Updating	124
1.9.73. RHSA-2025:3573 - OpenShift Container Platform 4.12.75 bug fix and security update	124
1.9.73.1. Updating	124
1.9.74. RHSA-2025:4409 - OpenShift Container Platform 4.12.76 bug fix and security update	125
1.9.74.1. Bug fixes	125
1.9.74.2. Updating	125
1.9.75. RHSA-2025:8303 - OpenShift Container Platform 4.12.77 bug fix and security update	125
1.9.75.1. Bug fixes	125

1.9.75.2. Updating	126
1.9.76. RHSA-2025:10270 - OpenShift Container Platform 4.12.78 bug fix and security update	126
1.9.76.1. Bug fixes	126
1.9.76.2. Updating	126
1.9.77. RHSA-2025:12323 - OpenShift Container Platform 4.12.79 bug fix and security update	126
1.9.77.1. Updating	126
1.9.78. RHSA-2025:15332 - OpenShift Container Platform 4.12.80 bug fix and security update	127
1.9.78.1. Updating	127

CHAPTER 1. OPENSIFT CONTAINER PLATFORM 4.12 RELEASE NOTES

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

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

1.1. ABOUT THIS RELEASE

OpenShift Container Platform ([RHSA-2022:7399](#)) is now available. This release uses [Kubernetes 1.25](#) with CRI-O runtime. New features, changes, and known issues that pertain to OpenShift Container Platform 4.12 are included in this topic.

OpenShift Container Platform 4.12 clusters are available at <https://console.redhat.com/openshift>. With the Red Hat OpenShift Cluster Manager application for OpenShift Container Platform, you can deploy OpenShift clusters to either on-premises or cloud environments.

OpenShift Container Platform 4.12 is supported on Red Hat Enterprise Linux (RHEL) 8.6 and a later version of RHEL 8 that is released before End of Life of OpenShift Container Platform 4.12.

You must use RHCOS machines for the control plane, and you can use either RHCOS or RHEL for compute machines.

Starting with OpenShift Container Platform 4.12 an additional six months of Extended Update Support (EUS) phase on even numbered releases from 18 months to two years. For more information, see the [Red Hat OpenShift Container Platform Life Cycle Policy](#).

OpenShift Container Platform 4.8 is an Extended Update Support (EUS) release. More information on Red Hat OpenShift EUS is available in [OpenShift Life Cycle](#) and [OpenShift EUS Overview](#).

Maintenance support ends for version 4.8 in January 2023 and goes to extended life phase. For more information, see the [Red Hat OpenShift Container Platform Life Cycle Policy](#).

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

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

1.3. NEW FEATURES AND ENHANCEMENTS

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

1.3.1. Red Hat Enterprise Linux CoreOS (RHCOS)

1.3.1.1. Default consoles for new clusters are now determined by the installation platform

Red Hat Enterprise Linux CoreOS (RHCOS) nodes installed from an OpenShift Container Platform 4.12 boot image now use a platform-specific default console. The default consoles on cloud platforms correspond to the specific system consoles expected by that cloud provider. VMware and OpenStack images now use a primary graphical console and a secondary serial console. Other bare metal installations now use only the graphical console by default, and do not enable a serial console. Installations performed with **coreos-installer** can override existing defaults and enable the serial console.

Existing nodes are not affected. New nodes on existing clusters are not likely to be affected because they are typically installed from the boot image that was originally used to install the cluster.

For information about how to enable the serial console, see the following documentation:

- [Default console configuration](#) .
- [Modifying a live install ISO image to enable the serial console](#) .
- [Modifying a live install PXE environment to enable the serial console](#) .

1.3.1.2. IBM Secure Execution on IBM Z and LinuxONE (Technology Preview)

OpenShift Container Platform now supports configuring Red Hat Enterprise Linux CoreOS (RHCOS) nodes for IBM Secure Execution on IBM Z and LinuxONE (s390x architecture) as a Technology Preview feature. IBM Secure Execution is a hardware enhancement that protects memory boundaries for KVM guests. IBM Secure Execution provides the highest level of isolation and security for cluster workloads, and you can enable it by using an IBM Secure Execution-ready QCOW2 boot image.

To use IBM Secure Execution, you must have host keys for your host machine(s) and they must be specified in your Ignition configuration file. IBM Secure Execution automatically encrypts your boot volumes using LUKS encryption.

For more information, see [Installing RHCOS using IBM Secure Execution](#) .

1.3.1.3. RHCOS now uses RHEL 8.6

RHCOS now uses Red Hat Enterprise Linux (RHEL) 8.6 packages in OpenShift Container Platform 4.12. This enables you to have the latest fixes, features, and enhancements, as well as the latest hardware support and driver updates. OpenShift Container Platform 4.10 is an Extended Update Support (EUS) release that will continue to use RHEL 8.4 EUS packages for the entirety of its lifecycle.

1.3.2. Installation and upgrade

1.3.2.1. Assisted Installer SaaS provides platform integration support for Nutanix

Assisted Installer SaaS on console.redhat.com supports installation of OpenShift Container Platform on the Nutanix platform with Machine API integration using either the Assisted Installer user interface or the REST API. Integration enables Nutanix Prism users to manage their infrastructure from a single interface, and enables auto-scaling. There are a few additional installation steps to enable Nutanix integration with Assisted Installer SaaS. See the Assisted Installer documentation for details.

1.3.2.2. Specify the load balancer type in AWS during installation

Beginning with OpenShift Container Platform 4.12, you can specify either Network Load Balancer (NLB) or Classic as a persistent load balancer type in AWS during installation. Afterwards, if an Ingress Controller is deleted, the load balancer type persists with the lbType configured during installation.

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

1.3.2.3. Extend worker nodes to the edge of AWS when installing into an existing Virtual Private Cloud (VPC) with Local Zone subnets.

With this update you can install OpenShift Container Platform to an existing VPC with installer-provisioned infrastructure, extending the worker nodes to Local Zones subnets. The installation program will provision worker nodes on the edge of the AWS network that are specifically designated for user applications by using NoSchedule taints. Applications deployed on the Local Zones locations deliver low latency for end users.

For more information, see [Installing a cluster using AWS Local Zones](#).

1.3.2.4. Google Cloud Platform Marketplace offering

OpenShift Container Platform is now available on the GCP Marketplace. Installing an OpenShift Container Platform with a GCP Marketplace image lets you create self-managed cluster deployments that are billed on pay-per-use basis (hourly, per core) through GCP, while still being supported directly by Red Hat.

For more information about installing using installer-provisioned infrastructure, see [Using a GCP Marketplace image](#). For more information about installing a using user-provisioned infrastructure, see [Creating additional worker machines in GCP](#).

1.3.2.5. Troubleshooting bootstrap failures during installation on GCP and Azure

The installer now gathers serial console logs from the bootstrap and control plane hosts on GCP and Azure. This log data is added to the standard bootstrap log bundle.

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

1.3.2.6. IBM Cloud VPC general availability

IBM Cloud VPC is now generally available in OpenShift Container Platform 4.12.

For more information about installing a cluster, see [Preparing to install on IBM Cloud VPC](#).

1.3.2.7. Required administrator acknowledgment when upgrading from OpenShift Container Platform 4.11 to 4.12

OpenShift Container Platform 4.12 uses Kubernetes 1.25, which removed [several deprecated APIs](#).

A cluster administrator must provide a manual acknowledgment before the cluster can be upgraded from OpenShift Container Platform 4.11 to 4.12. This is to help prevent issues after upgrading to OpenShift Container Platform 4.12, where APIs that have been removed are still in use by workloads, tools, or other components running on or interacting with the cluster. Administrators must evaluate their cluster for any APIs in use that will be removed and migrate the affected components to use the appropriate new API version. After this is done, the administrator can provide the administrator acknowledgment.

All OpenShift Container Platform 4.11 clusters require this administrator acknowledgment before they can be upgraded to OpenShift Container Platform 4.12.

For more information, see [Preparing to update to OpenShift Container Platform 4.12](#) .

1.3.2.8. Enabling a feature set when installing a cluster

Beginning with OpenShift Container Platform 4.12, you can enable a feature set as part of the installation process. A feature set is a collection of OpenShift Container Platform features that are not enabled by default.

For more information about enabling a feature set during installation, see [Enabling OpenShift Container Platform features using feature gates](#).

1.3.2.9. OpenShift Container Platform on ARM

OpenShift Container Platform 4.12 is now supported on ARM architecture-based Azure installer-provisioned infrastructure. AWS Graviton 3 processors are now available for cluster deployments and are also supported on OpenShift Container Platform 4.11. For more information about instance availability and installation documentation, see [Supported installation methods for different platforms](#)

1.3.2.10. Mirroring file-based catalog Operator images in OCI format with the oc-mirror CLI plugin (Technology Preview)

Using the oc-mirror CLI plugin to mirror file-based catalog Operator images in OCI format instead of Docker v2 format is now available as a [Technology Preview](#) .

For more information, see [Mirroring file-based catalog Operator images in OCI format](#) .

1.3.2.11. Installing an OpenShift Container Platform cluster on GCP into a shared VPC (Technology Preview)

In OpenShift Container Platform 4.12, you can install a cluster on GCP into a shared VPC as a [Technology Preview](#). In this installation method, the cluster is configured to use a VPC from a different GCP project. A shared VPC enables an organization to connect resources from multiple projects to a common VPC network. You can communicate within the organization securely and efficiently by using internal IP addresses from that network.

For more information, see [Installing a cluster on GCP into a shared VPC](#) .

1.3.2.12. Consistent IP address for Ironic API in bare-metal installations without a provisioning network

With this update, in bare-metal installations without a provisioning network, the Ironic API service is accessible through a proxy server. This proxy server provides a consistent IP address for the Ironic API service. If the Metal3 pod that contains **metal3-ironic** relocates to another pod, the consistent proxy address ensures constant communication with the Ironic API service.

1.3.2.13. Installing OpenShift Container Platform on GCP using service account authentication

In OpenShift Container Platform 4.12, you can install a cluster on GCP using a virtual machine with a service account attached to it. This allows you to perform an installation without needing to use a service account JSON file.

For more information, see [Creating a GCP service account](#).

1.3.2.14. **propagateUserTags** parameter for AWS resources provisioned by the OpenShift Container Platform cluster

In OpenShift Container Platform 4.12, the **propagateUserTags** parameter is a flag that directs in-cluster Operators to include the specified user tags in the tags of the AWS resources that the Operators create.

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

1.3.2.15. **Ironic container images use RHEL 9 base image**

In earlier versions of OpenShift Container Platform, Ironic container images used Red Hat Enterprise Linux (RHEL) 8 as the base image. From OpenShift Container Platform 4.12, Ironic container images use RHEL 9 as the base image. The RHEL 9 base image adds support for CentOS Stream 9, Python 3.8, and Python 3.9 in Ironic components.

For more information about the Ironic provisioning service, see [Deploying installer-provisioned clusters on bare metal](#).

1.3.2.16. **Cloud provider configuration updates for clusters that run on RHOSP**

In OpenShift Container Platform 4.12, clusters that run on Red Hat OpenStack Platform (RHOSP) are switched from the legacy OpenStack cloud provider to the external Cloud Controller Manager (CCM). This change follows the move in Kubernetes from in-tree, legacy cloud providers to external cloud providers that are implemented by using the [Cloud Controller Manager](#).

For more information, see [The OpenStack Cloud Controller Manager](#).

1.3.2.17. **Support for workloads on RHOSP distributed compute nodes**

In OpenShift Container Platform 4.12, cluster deployments to Red Hat OpenStack Platform (RHOSP) clouds that have distributed compute node (DCN) architecture were validated. A reference architecture for these deployments is forthcoming.

For a brief overview of this type of deployment, see the blog post [Deploying Your Cluster at the Edge With OpenStack](#).

1.3.2.18. **OpenShift Container Platform on AWS Outposts (Technology Preview)**

OpenShift Container Platform 4.12 is now supported on the AWS Outposts platform as a [Technology Preview](#). With AWS Outposts you can deploy edge-based worker nodes, while using AWS Regions for the control plane nodes. For more information, see [Installing a cluster on AWS with remote workers on AWS Outposts](#).

1.3.2.19. **Agent-based installation supports two input modes**

The Agent-based installation supports two input modes:

- **install-config.yaml** file
- **agent-config.yaml** file

Optional

Optional

- Zero Touch Provisioning (ZTP) manifests

With the preferred mode, you can configure the **install-config.yaml** file and specify Agent-based specific settings in the **agent-config.yaml** file. For more information, see [About the Agent-based OpenShift Container Platform Installer](#).

1.3.2.20. Agent-based installation supports installing OpenShift Container Platform clusters in FIPS compliant mode

Agent-based OpenShift Container Platform Installer supports OpenShift Container Platform clusters in Federal Information Processing Standards (FIPS) compliant mode. You must set the value of the **fips** field to **True** in the **install-config.yaml** file. For more information, see [About FIPS compliance](#).

1.3.2.21. Deploy an Agent-based OpenShift Container Platform cluster in a disconnected environment

You can perform an Agent-based installation in a disconnected environment. To create an image that is used in a disconnected environment, the **imageContentSources** section in the **install-config.yaml** file must contain the mirror information or **registries.conf** file if you are using ZTP manifests. The actual configuration settings to use in these files are supplied by either the **oc adm release mirror** or **oc mirror** command. For more information, see [Understanding disconnected installation mirroring](#).

1.3.2.22. Explanation of field to build and push graph-data

When creating the image set configuration, you can add the **graph: true** field to build and push the graph-data image to the mirror registry. The graph-data image is required to create OpenShift Update Service (OSUS). The **graph: true** field also generates the **UpdateService** custom resource manifest. The **oc** command-line interface (CLI) can use the **UpdateService** custom resource manifest to create OSUS.

1.3.2.23. Agent-based installation supports single and dual stack networking

You can create the agent ISO image with the following IP address configurations:

- IPv4
- IPv6
- IPv4 and IPv6 in parallel (dual-stack)



NOTE

IPv6 is supported only on bare metal platforms.

For more information, see [Dual and single IP stack clusters](#).

1.3.2.24. Agent deployed OpenShift Container Platform cluster can be used as a hub cluster

You can install the multicluster engine for Kubernetes Operator and deploy a hub cluster with the Agent-based OpenShift Container Platform Installer. For more information, see [Preparing an Agent-based installed cluster for the multicluster engine for Kubernetes Operator](#).

1.3.2.25. Agent-based installation performs installation validations

The Agent-based OpenShift Container Platform Installer performs validations on:

- Installation image generation: The user-provided manifests are checked for validity and compatibility.
- Installation: The installation service checks the hardware available for installation and emits validation events that can be retrieved with the **openshift-install agent wait-for** subcommands.

For more information, see [Installation validations](#).

1.3.2.26. Configure static networking in a Agent-based installation

With the Agent-based OpenShift Container Platform Installer, you can configure static IP addresses for IPv4, IPv6, or dual-stack (both IPv4 and IPv6) for all the hosts prior to creating the agent ISO image. You can add the static addresses to the **hosts** section of the **agent-config.yaml** file or in the **NMStateConfig.yaml** file if you are using the ZTP manifests. Note that the configuration of the addresses must follow the syntax rules for NMState as described in [NMState state examples](#).



NOTE

IPv6 is supported only on bare metal platforms.

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

1.3.2.27. CLI based automated deployment in an Agent-based installation

With the Agent-based OpenShift Container Platform Installer, you can define your installation configurations, generate an ISO for all the nodes, and then have an unattended installation by booting the target systems with the generated ISO. For more information, see [Installing a OpenShift Container Platform cluster with the Agent-based OpenShift Container Platform Installer](#).

1.3.2.28. Agent-based installation supports host specific configuration at the installation time

You can configure the hostname, network configuration in NMState format, root device hints, and role in an Agent-based installation.

For more information, see [About root device hints](#).

1.3.2.29. Agent-based installation supports DHCP

With the Agent-based OpenShift Container Platform Installer, you can deploy to environments where you rely on DHCP to configure networking for all the nodes, as long as you know the IP that at least one of the systems will receive. This IP is required so that all nodes use it as a meeting point. For more information, see [DHCP](#).

1.3.2.30. Installing a cluster on Nutanix with limited internet access

You can now install a cluster on Nutanix when the environment has limited access to to the internet, as in the case of a disconnected or restricted network cluster. With this type of installation, you create a registry that mirrors the contents of the OpenShift Container Platform image registry and contains the

installation media. You can create this registry on a mirror host, which can access both the internet and your closed network.

For more information, see [About disconnected installation mirroring](#) and [Installing a cluster on Nutanix in a restricted network](#).

1.3.3. Post-installation configuration

1.3.3.1. CSI driver installation on vSphere clusters

To install a CSI driver on a cluster running on vSphere, the following requirements must be met:

- Virtual machines of hardware version 15 or later
- VMware vSphere version 7.0 Update 2 or later, which includes version 8.0.
- vCenter 7.0 Update 2 or later, which includes version 8.0.
- No third-party CSI driver already installed in the cluster
If a third-party CSI driver is present in the cluster, OpenShift Container Platform does not overwrite it.

Components with versions earlier than those above are still supported, but are deprecated. These versions are still fully supported, but version 4.12 of OpenShift Container Platform requires vSphere virtual hardware version 15 or later. For more information, see [Deprecated and removed features](#).

Failing to meet the above requirements prevents OpenShift Container Platform from upgrading to OpenShift Container Platform 4.13 or later.

1.3.3.2. Cluster Capabilities

The following new cluster capabilities have been added:

- Console
- Insights
- Storage
- CSISnapshot

A new predefined set of cluster capabilities, **v4.12**, has been added. This includes all capabilities from **v4.11**, and the new capabilities added with the current release.

For more information, see link: [Enabling cluster capabilities](#).

1.3.3.3. OpenShift Container Platform with multi-architecture compute machines (Technology Preview)

OpenShift Container Platform 4.12 with multi-architecture compute machines now supports manifest listed images on image streams. For more information about manifest list images, see [Configuring multi-architecture compute machines on an OpenShift Container Platform cluster](#).

On a cluster with multi-architecture compute machines, you can now override the node affinity in the Operator's **Subscription** object to schedule pods on nodes with architectures that the Operator supports. For more information, see [Using node affinity to control where an Operator is installed](#) .

1.3.4. Web console

1.3.4.1. Administrator Perspective

With this release, there are several updates to the **Administrator** perspective of the web console.

- The OpenShift Container Platform web console displays a **ConsoleNotification** if the cluster is upgrading. Once the upgrade is done, the notification is removed.
- A *restart rollout* option for the **Deployment** resource and a *retry rollouts* option for the **DeploymentConfig** resource are available on the **Action** and **Kebab** menus.

1.3.4.1.1. Multi-architecture compute machines on the OpenShift Container Platform web console

The **console-operator** now scans all nodes and builds a set of all architecture types that cluster nodes run on and pass it to the **console-config.yaml**. The **console-operator** can be installed on nodes with architectures of the values **amd64**, **arm64**, **ppc64le**, or **s390x**.

For more information about multi-architecture compute machines, see [Configuring a multi-architecture compute machine on an OpenShift cluster](#).

1.3.4.1.2. Dynamic plugin generally available

This feature was previously introduced as a Technology Preview in OpenShift Container Platform 4.10 and is now generally available in OpenShift Container Platform 4.12. With the dynamic plugin, you can build high quality and unique user experiences natively in the web console. You can:

- Add custom pages.
- Add perspectives beyond administrator and developer.
- Add navigation items.
- Add tabs and actions to resource pages.
- Extend existing pages.

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

1.3.4.2. Developer Perspective

With this release, there are several updates to the **Developer** perspective of the web console. You can perform the following actions:

- Export your application in the ZIP file format to another project or cluster by using the **Export application** option on the **+Add** page.
- Create a Kafka event sink to receive events from a particular source and send them to a Kafka topic.

- Set the default resource preference in the **User Preferences** → **Applications** page. In addition, you can select another resource type to be the default.
 - Optionally, set another resource type from the **Add** page by clicking **Import from Git** → **Advanced options** → **Resource type** and selecting the resource from the drop-down list.
- Make the **status.HostIP** node IP address for pods visible in the **Details** tab of the **Pods** page.
- See the resource quota alert label on the **Topology** and **Add** pages whenever any resource reaches the quota. The alert label link takes you to the **ResourceQuotas** list page. If the alert label link is for a single resource quota, it takes you to the **ResourceQuota details** page.
 - For deployments, an alert is displayed in the topology node side panel if any errors are associated with resource quotas. Also, a yellow border is displayed around the deployment nodes when the resource quota is exceeded.
- Customize the following UI items using the form or YAML view:
 - Perspectives visible to users
 - Quick starts visible to users
 - Cluster roles accessible to a project
 - Actions visible on the **+Add** page
 - Item types in the **Developer Catalog**
- See the common updates to the **Pipeline details** and **PipelineRun details** page visualization by performing the following actions:
 - Use the mouse wheel to change the zoom factor.
 - Hover over the tasks to see the task details.
 - Use the standard icons to zoom in, zoom out, fit to screen, and reset the view.
 - **PipelineRun details** page only: At specific zoom factors, the background color of the tasks changes to indicate the error or warning status. You can hover over the tasks badge to see the total number of tasks and the completed tasks.

1.3.4.2.1. Helm page improvements

In OpenShift Container Platform 4.12, you can do the following from the **Helm** page:

- Create Helm releases and repositories using the **Create** button.
- Create, update, or delete a cluster-scoped or a namespace-scoped Helm chart repository.
- View the list of the existing Helm chart repositories with their scope in the **Repositories** page.
- View the newly created Helm release in the **Helm Releases** page.

1.3.4.2.2. Negative matchers in Alertmanager

With this update, Alertmanager now supports a **Negative matcher** option. Using **Negative matcher**, you can update the **Label value** to a Not Equals matcher. The negative matcher checkbox changes **=** (value equals) into **!=** (value does not equal) and changes **=~** (value matches regular expression) into **!~** (value

does not match regular expression). Also, the **Use RegEx** checkbox label is renamed to **RegEx**.

1.3.5. OpenShift CLI (oc)

1.3.5.1. Managing plugins for the OpenShift CLI with Krew (Technology Preview)

Using Krew to install and manage plugins for the OpenShift CLI (**oc**) is now available as a [Technology Preview](#).

For more information, see [Managing CLI plugins with Krew](#).

1.3.6. IBM Z and LinuxONE

With this release, IBM Z and LinuxONE are now compatible with OpenShift Container Platform 4.12. The installation can be performed with z/VM or RHEL KVM. For installation instructions, see the following documentation:

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

Notable enhancements

The following new features are supported on IBM Z and LinuxONE with OpenShift Container Platform 4.12:

- Cron jobs
- Descheduler
- FIPS cryptography
- IPv6
- PodDisruptionBudget
- Scheduler profiles
- Stream Control Transmission Protocol (SCTP)

IBM Secure Execution (Technology Preview)

OpenShift Container Platform now supports configuring Red Hat Enterprise Linux CoreOS (RHCOS) nodes for IBM Secure Execution on IBM Z and LinuxONE (s390x architecture) as a Technology Preview feature.

For installation instructions, see the following documentation:

- [Installing RHCOS using IBM Secure Execution](#)

Supported features

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

- Currently, the following Operators are supported:

- Cluster Logging Operator
- Compliance Operator
- File Integrity Operator
- Local Storage Operator
- NFD Operator
- NMState Operator
- OpenShift Elasticsearch Operator
- Service Binding Operator
- Vertical Pod Autoscaler Operator
- The following Multus CNI plugins are supported:
 - Bridge
 - Host-device
 - IPAM
 - IPVLAN
- Alternate authentication providers
- Automatic Device Discovery with Local Storage Operator
- CSI Volumes
 - Cloning
 - Expansion
 - Snapshot
- Encrypting data stored in etcd
- Helm
- Horizontal pod autoscaling
- Monitoring for user-defined projects
- Multipathing
- Operator API
- OC CLI plugins
- Persistent storage using iSCSI
- Persistent storage using local volumes (Local Storage Operator)
- Persistent storage using hostPath

- Persistent storage using Fibre Channel
- Persistent storage using Raw Block
- OVN-Kubernetes, including IPsec encryption
- Support for multiple network interfaces
- Three-node cluster support
- z/VM Emulated FBA devices on SCSI disks
- 4K FCP block device

These features are available only for OpenShift Container Platform on IBM Z and LinuxONE for 4.12:

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

Restrictions

The following restrictions impact OpenShift Container Platform on IBM Z and LinuxONE:

- Automatic repair of damaged machines with machine health checking
- Red Hat OpenShift Local
- Controlling overcommit and managing container density on nodes
- NVMe
- OpenShift Metering
- OpenShift Virtualization
- Precision Time Protocol (PTP) hardware
- Tang mode disk encryption during OpenShift Container Platform deployment
- Compute nodes must run Red Hat Enterprise Linux CoreOS (RHCOS)
- Persistent shared storage must be provisioned by using either Red Hat OpenShift Data Foundation or other supported storage protocols
- Persistent non-shared storage must be provisioned using local storage, like iSCSI, FC, or using LSO with DASD, FCP, or EDEV/FBA

1.3.7. IBM Power

With this release, IBM Power is now compatible with OpenShift Container Platform 4.12. For installation instructions, see the following documentation:

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

Notable enhancements

The following new features are supported on IBM Power with OpenShift Container Platform 4.12:

- Cloud controller manager for IBM Cloud
- Cron jobs
- Descheduler
- FIPS cryptography
- PodDisruptionBudget
- Scheduler profiles
- Stream Control Transmission Protocol (SCTP)
- Topology Manager

Supported features

The following features are also supported on IBM Power:

- Currently, the following Operators are supported:
 - Cluster Logging Operator
 - Compliance Operator
 - File Integrity Operator
 - Local Storage Operator
 - NFD Operator
 - NMState Operator
 - OpenShift Elasticsearch Operator
 - SR-IOV Network Operator
 - Service Binding Operator
 - Vertical Pod Autoscaler Operator
- The following Multus CNI plugins are supported:
 - Bridge
 - Host-device
 - IPAM
 - IPVLAN
- Alternate authentication providers
- CSI Volumes
 - Cloning
 - Expansion

- Snapshot
- Encrypting data stored in etcd
- Helm
- Horizontal pod autoscaling
- IPv6
- Monitoring for user-defined projects
- Multipathing
- Multus SR-IOV
- Operator API
- OC CLI plugins
- OVN-Kubernetes, including IPsec encryption
- Persistent storage using iSCSI
- Persistent storage using local volumes (Local Storage Operator)
- Persistent storage using hostPath
- Persistent storage using Fibre Channel
- Persistent storage using Raw Block
- Support for multiple network interfaces
- Support for Power10
- Three-node cluster support
- 4K Disk Support

Restrictions

The following restrictions impact OpenShift Container Platform on IBM Power:

- Automatic repair of damaged machines with machine health checking
- Red Hat OpenShift Local
- Controlling overcommit and managing container density on nodes
- OpenShift Metering
- OpenShift Virtualization
- Precision Time Protocol (PTP) hardware
- Tang mode disk encryption during OpenShift Container Platform deployment
- Compute nodes must run Red Hat Enterprise Linux CoreOS (RHCOS)

- Persistent storage must be of the Filesystem type that uses local volumes, Red Hat OpenShift Data Foundation, Network File System (NFS), or Container Storage Interface (CSI)

1.3.8. Images

A new import value, **importMode**, has been added to the **importPolicy** parameter of image streams. The following fields are available for this value:

- **Legacy: Legacy** is the default value for **importMode**. When active, the manifest list is discarded, and a single sub-manifest is imported. The platform is chosen in the following order of priority:
 1. Tag annotations
 2. Control plane architecture
 3. Linux/AMD64
 4. The first manifest in the list
- **PreserveOriginal**: When active, the original manifest is preserved. For manifest lists, the manifest list and all of its sub-manifests are imported.

1.3.9. Security and compliance

1.3.9.1. Security Profiles Operator

The Security Profiles Operator (SPO) is now available for OpenShift Container Platform 4.12 and later.

The SPO provides a way to define secure computing ([seccomp](#)) profiles and SELinux profiles as custom resources, synchronizing profiles to every node in a given namespace.

For more information, see [Security Profiles Operator Overview](#).

1.3.10. Networking

1.3.10.1. Support for dual-stack addressing for the API VIP and Ingress VIP

Assisted Installer supports installation of OpenShift Container Platform 4.12 and later versions with dual stack networking for the API VIP and Ingress VIP on bare metal only. This support introduces two new configuration settings: **api_vips** and **ingress_vips**, which can take a list of IP addresses. The legacy settings, **api_vip** and **ingress_vip** must also be set in OpenShift Container Platform 4.12; however, since they only take one IP address, you must set the IPv4 address when configuring dual stack networking for the API VIP and Ingress VIP with the legacy **api_vip** and **ingress_vip** configuration settings.

The API VIP address and the Ingress VIP address must be of the primary IP address family when using dual-stack networking. Currently, Red Hat does not support dual-stack VIPs or dual-stack networking with IPv6 as the primary IP address family. However, Red Hat does support dual-stack networking with IPv4 as the primary IP address family. Therefore, you must place the IPv4 entries before the IPv6 entries. See the Assisted Installer documentation for details.

1.3.10.2. Red Hat OpenShift Networking

Red Hat OpenShift Networking is an ecosystem of features, plugins, and advanced networking capabilities that extend Kubernetes networking beyond the Kubernetes CNI plugin with the advanced networking-related features that your cluster needs to manage its network traffic for one or multiple hybrid clusters. This ecosystem of networking capabilities integrates ingress, egress, load balancing, high-performance throughput, security, and inter-, and intra-cluster traffic management and provides role-based observability tooling to reduce its natural complexities.

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

1.3.10.3. OVN-Kubernetes is now the default networking plugin

When installing a new cluster the OVN-Kubernetes network plugin is the default networking plugin. For all prior versions of OpenShift Container Platform, OpenShift SDN remains the default networking plugin.

The OVN-Kubernetes network plugin includes a wider array of features than OpenShift SDN, including:

- Support for all existing OpenShift SDN features
- Support for [IPv6 networks](#)
- Support for [Configuring IPsec encryption](#)
- Complete support for the [NetworkPolicy API](#)
- Support for [audit logging of network policy events](#)
- Support for [network flow tracking](#) in NetFlow, sFlow, and IPFIX formats
- Support for [hybrid networks](#) for Windows containers
- Support for [hardware offloading](#) to compatible NICs

There are also enormous scale, performance, and stability improvements in OpenShift Container Platform 4.12 compared to prior versions.

If you are using the OpenShift SDN network plugin, note that:

- Existing and future deployments using OpenShift SDN continues to be supported.
- OpenShift SDN remains the default on OpenShift Container Platform versions earlier than 4.12.
- As of OpenShift Container Platform 4.12, OpenShift SDN is a supported installation-time option.
- OpenShift SDN remains feature frozen.

For more information about OVN-Kubernetes, including a feature comparison matrix with OpenShift SDN, see [About the OVN-Kubernetes network plugin](#).

For information on migrating to OVN-Kubernetes from OpenShift SDN, see [Migrating from the OpenShift SDN network plugin](#).

1.3.10.4. Ingress Node Firewall Operator

This update introduces a new stateless Ingress Node Firewall Operator. You can now configure firewall rules at the node level. For more information, see [Ingress Node Firewall Operator](#).

1.3.10.5. Enhancements to networking metrics

The following metrics are now available for the OVN-Kubernetes network plugin:

- **ovn_controller_southbound_database_connected**
- **ovnkube_master_libovsdb_monitors**
- **ovnkube_master_network_programming_duration_seconds**
- **ovnkube_master_network_programming_ovn_duration_seconds**
- **ovnkube_master_egress_routing_via_host**
- **ovs_vswitchd_interface_resets_total**
- **ovs_vswitchd_interface_rx_dropped_total**
- **ovs_vswitchd_interface_tx_dropped_total**
- **ovs_vswitchd_interface_rx_errors_total**
- **ovs_vswitchd_interface_tx_errors_total**
- **ovs_vswitchd_interface_collisions_total**

The following metric has been removed:

- **ovnkube_master_skipped_nbctl_daemon_total**

1.3.10.6. Multi-zone Installer Provisioned Infrastructure VMware vSphere installation (Technology Preview)

Beginning with OpenShift Container Platform 4.12, the ability to configure multiple vCenter datacenters and multiple vCenter clusters in a single vCenter installation using installer-provisioned infrastructure is now available as a Technology Preview feature. Using vCenter tags, you can use this feature to associate vCenter datacenters and compute clusters with openshift-regions and openshift-zones. These associations define failure domains to enable application workloads to be associated with specific locations and failure domains.

1.3.10.7. Kubernetes NMState in VMware vSphere now supported

Beginning with OpenShift Container Platform 4.12, you can configure the networking settings such as DNS servers or search domains, VLANs, bridges, and interface bonding using the Kubernetes NMState Operator on your VMware vSphere instance.

For more information, see [About the Kubernetes NMState Operator](#).

1.3.10.8. Kubernetes NMState in OpenStack now supported

Beginning with OpenShift Container Platform 4.12, you can configure the networking settings such as DNS servers or search domains, VLANs, bridges, and interface bonding using the Kubernetes NMState Operator on your OpenStack instance.

For more information, see [About the Kubernetes NMState Operator](#).

1.3.10.9. External DNS Operator

In OpenShift Container Platform 4.12, the External DNS Operator modifies the format of the ExternalDNS wildcard TXT records on AzureDNS. The External DNS Operator replaces the asterisk with **any** in ExternalDNS wildcard TXT records. You must avoid the ExternalDNS wildcard A and CNAME records having **any** leftmost subdomain because this might cause a conflict.

The upstream version of **ExternalDNS** for OpenShift Container Platform 4.12 is v0.13.1.

1.3.10.10. The ingressClassName field is required for each ingress object

Beginning with OpenShift Container Platform 4.12, you must specify a class name in the **ingressClassName** field for each ingress object to ensure proper routing and functionality.

1.3.10.11. Capturing metrics and telemetry associated with the use of routes and shards

In OpenShift Container Platform 4.12, the Cluster Ingress Operator exports a new metric named **route_metrics_controller_routes_per_shard**. The **shard_name** label of the metric specifies the name of the shards. This metric gives the total number of routes that are admitted by each shard.

The following metrics are sent through telemetry.

Table 1.1. Metrics sent through telemetry

Name	Recording rule expression	Description
cluster:route_metrics_controller_routes_per_shard:min	min(route_metrics_controller_routes_per_shard)	Tracks the minimum number of routes admitted by any of the shards
cluster:route_metrics_controller_routes_per_shard:max	max(route_metrics_controller_routes_per_shard)	Tracks the maximum number of routes admitted by any of the shards
cluster:route_metrics_controller_routes_per_shard:avg	avg(route_metrics_controller_routes_per_shard)	Tracks the average value of the route_metrics_controller_routes_per_shard metric
cluster:route_metrics_controller_routes_per_shard:median	quantile(0.5, route_metrics_controller_routes_per_shard)	Tracks the median value of the route_metrics_controller_routes_per_shard metric
cluster:openshift_route_info:tls_termination:sum	sum (openshift_route_info by (tls_termination))	Tracks the number of routes for each tls_termination value. The possible values for tls_termination are edge , passthrough and reencrypt

1.3.10.12. AWS Load Balancer Operator

In OpenShift Container Platform 4.12, the AWS Load Balancer controller now implements the Kubernetes Ingress specification for multiple matches. If multiple paths within an Ingress match a

request, the longest matching path takes the precedence. If two paths still match, paths with an exact path type take precedence over a prefix path type.

The AWS Load Balancer Operator sets the **EnableIPTargetType** feature gate to **false**. The AWS Load Balancer controller disables the support for services and ingress resources for **target-type ip**.

The upstream version of **aws-load-balancer-controller** for an OpenShift Container Platform 4.12 is v2.4.4.

1.3.10.13. Ingress Controller Autoscaling (Technology Preview)

You can now use the OpenShift Container Platform Custom Metrics Autoscaler Operator to dynamically scale the default Ingress Controller based on metrics in your deployed cluster, such as the number of worker nodes available. The Custom Metrics Autoscaler is available as a Technology Preview feature.

For more information, see [Autoscaling an Ingress Controller](#).

1.3.10.14. HAProxy maxConnections default is now 50,000

In OpenShift Container Platform 4.12, the default value for the **maxConnections** setting is now 50000. Previously starting with OpenShift Container Platform 4.11, the default value for the **maxConnections** setting was 20000.

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

1.3.10.15. Configuration of an Ingress Controller for manual DNS management

You can now configure an Ingress Controller to stop automatic DNS management and start manual DNS management. Set the **dnsManagementPolicy** parameter to specify automatic or manual DNS management.

For more information, see [Configuring an Ingress Controller to manually manage DNS](#).

1.3.10.16. Supported hardware for SR-IOV (Single Root I/O Virtualization)

OpenShift Container Platform 4.12 adds support for the following SR-IOV devices:

- Intel X710 Base T
- MT2892 Family [ConnectX-6 Dx]
- MT2894 Family [ConnectX-6 Lx]
- MT42822 BlueField-2 in ConnectX-6 NIC mode
- Silicom STS Family

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

1.3.10.17. Supported hardware for OvS (Open vSwitch) Hardware Offload

OpenShift Container Platform 4.12 adds OvS Hardware Offload support for the following devices:

- MT2892 Family [ConnectX-6 Dx]

- MT2894 Family [ConnectX-6 Lx]
- MT42822 BlueField-2 in ConnectX-6 NIC mode

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

1.3.10.18. Multi-network-policy supported for SR-IOV (Technology Preview)

OpenShift Container Platform 4.12 adds support for configuring multi-network policy for SR-IOV devices.

You can now configure multi-network for SR-IOV additional networks. Configuring SR-IOV additional networks is a Technology Preview feature and is only supported with kernel network interface cards (NICs).

For more information, see [Configuring multi-network policy](#).

1.3.10.19. Switch between AWS load balancer types without deleting the Ingress Controller

You can update the Ingress Controller to switch between an AWS Classic Load Balancer (CLB) and an AWS Network Load Balancer (NLB) without deleting the Ingress Controller.

For more information, see [Configuring ingress cluster traffic on AWS](#).

1.3.10.20. IPv6 unsolicited neighbor advertisements and IPv4 gratuitous address resolution protocol now default on the SR-IOV CNI plugin

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

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

1.3.10.21. Support for CoreDNS cache tuning

You can now configure the time-to-live (TTL) duration of both successful and unsuccessful DNS queries cached by CoreDNS.

For more information, see [Tuning the CoreDNS cache](#).

1.3.10.22. OVN-Kubernetes supports configuration of internal subnet

Previously, the subnet that OVN-Kubernetes uses internally was **100.64.0.0/16** for IPv4 and **fd98::/48** for IPv6 and could not be modified. To support instances when these subnets overlap with existing subnets in your infrastructure, you can now change these internal subnets to avoid any overlap.

For more information, see [Cluster Network Operator configuration object](#)

1.3.10.23. Egress IP support on Red Hat OpenStack Platform (RHOSP)

RHOSP, paired with OpenShift Container Platform, now supports automatic attachment and detachment of Egress IP addresses. The traffic from one or more pods in any number of namespaces has a consistent source IP address for services outside of the cluster. This support applies to OpenShift

SDN and OVN-Kubernetes as default network providers.

1.3.10.24. OpenShift SDN to OVN-Kubernetes feature migration support

If you plan to migrate from the OpenShift SDN network plugin to the OVN-Kubernetes network plugin, your configurations for the following capabilities are automatically converted to work with OVN-Kubernetes:

- Egress IP addresses
- Egress firewalls
- Multicast

For more information about how the migration to OVN-Kubernetes works, see [Migrating from the OpenShift SDN cluster network provider](#).

1.3.10.25. Egress firewall audit logging

For the OVN-Kubernetes network plugin, egress firewalls support audit logging using the same mechanism that network policy audit logging uses. For more information, see [Logging for egress firewall and network policy rules](#).

1.3.10.26. Advertise MetalLB from a given address pool from a subset of nodes

With this update, in BGP mode, you can use the node selector to advertise the MetalLB service from a subset of nodes, using a specific pool of IP addresses. This feature was introduced as a Technology Preview feature in OpenShift Container Platform 4.11 and is now generally available in OpenShift Container Platform 4.12 for BGP mode only. L2 mode remains a Technology Preview feature.

For more information, see [Advertising an IP address pool from a subset of nodes](#).

1.3.10.27. Additional deployment specifications for MetalLB

This update provides additional deployment specifications for MetalLB. When you use a custom resource to deploy MetalLB, you can use these additional deployment specifications to manage how MetalLB **speaker** and **controller** pods deploy and run in your cluster. For example, you can use MetalLB deployment specifications to manage where MetalLB pods are deployed, define CPU limits for MetalLB pods, and assign runtime classes to MetalLB pods.

For more information about deployment specifications for MetalLB, see [Deployment specifications for MetalLB](#).

1.3.10.28. Node IP selection improvements

Previously, the **nodeip-configuration** service on a cluster host selected the IP address from the interface that the default route used. If multiple routes were present, the service would select the route with the lowest metric value. As a result, network traffic could be distributed from the incorrect interface.

With OpenShift Container Platform 4.12, a new interface has been added to the **nodeip-configuration** service, which allows users to create a hint file. The hint file contains a variable, **NODEIP_HINT**, that overrides the default IP selection logic and selects a specific node IP address from the subnet **NODEIP_HINT** variable. Using the **NODEIP_HINT** variable allows users to specify which IP address is used, ensuring that network traffic is distributed from the correct interface.

For more information, see [Optional: Overriding the default node IP selection logic](#).

1.3.10.29. CoreDNS update to version 1.10.0

In OpenShift Container Platform 4.12, CoreDNS uses version 1.10.0, which includes the following changes:

- CoreDNS does not expand the query UDP buffer size if it was previously set to a smaller value.
- CoreDNS now always prefixes each log line in Kubernetes client logs with the associated log level.
- CoreDNS now reloads more quickly at an approximate speed of 20ms.

1.3.10.30. Support for a configurable reload interval in HAProxy

With this update, a cluster administrator can configure the reload interval to force HAProxy to reload its configuration less frequently in response to route and endpoint updates. The default minimum HAProxy reload interval is 5 seconds.

For more information, see [Configuring HAProxy reload interval](#).

1.3.10.31. Network Observability Operator updates

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

1.3.10.32. IPv6 for secondary network interfaces on RHOSP

IPv6 for secondary network interfaces is now supported in clusters that run on RHOSP.

For more information, see [Enabling IPv6 connectivity to pods on RHOSP](#).

1.3.10.33. UDP support for load balancers on RHOSP

Resulting from the switch to an external OpenStack cloud provider, UDP is now supported for **LoadBalancer** services for clusters that run on that platform.

1.3.10.34. Deploy the SR-IOV Operator for hosted control planes (Technology Preview)

If you configured and deployed your hosting service cluster, you can now deploy the SR-IOV Operator for a hosted cluster. For more information, see [Deploying the SR-IOV Operator for hosted control planes](#).

1.3.10.35. Support for IPv6 virtual IP (VIP) addresses for the Ingress VIP and API VIP services on bare metal

With this update, in installer-provisioned infrastructure clusters, the **ingressVIP** and **apiVIP** configuration settings in the **install-config.yaml** file are deprecated. Instead, use the **ingressVIPs** and **apiVIPs** configuration settings. These settings support dual-stack networking for applications on bare metal that require IPv4 and IPv6 access to the cluster by using the Ingress VIP and API VIP services.

The **ingressVIPs** and **apiVIPs** configuration settings use a list format to specify an IPv4 address, an IPv6 address, or both IP address formats. The order of the list indicates the primary and secondary VIP address for each service. The primary IP address must be from the IPv4 network when using dual stack networking.

1.3.10.36. Support for switching the Bluefield-2 network device from data processing unit (DPU) mode to network interface controller (NIC) mode (Technology Preview)

With this update, you can switch the BlueField-2 network device from data processing unit (DPU) mode to network interface controller (NIC) mode.

For more information, see [Switching Bluefield-2 from DPU to NIC](#).

1.3.10.37. Support for enabling hybrid networking after cluster installation

Previously, for clusters that use the [OVN-Kubernetes network plugin](#), during cluster installation you could enable hybrid networking so that your cluster supported Windows nodes. Now you can enable hybrid networking after installation. For more information, see [Configuring hybrid networking](#).

1.3.10.38. Support for allocateLoadBalancerNodePorts in the Network API service object

The **ServiceSpec** component in the Network API under the **Service** object describes the attributes that a user creates on a service. The **allocateLoadBalancerNodePorts** attribute within the **ServiceSpec** component is now supported as of OpenShift Container Platform 4.12.28 release. The **allocateLoadBalancerNodePorts** attribute defines whether the **NodePorts** will be automatically allocated for services of the **LoadBalancer** type.

For more information, see [Network API ServiceSpec object](#)

1.3.11. Storage

1.3.11.1. Persistent storage using the GCP Filestore Driver Operator (Technology Preview)

OpenShift Container Platform is capable of provisioning persistent volumes (PVs) using the Container Storage Interface (CSI) driver for Google Compute Platform (GCP) Filestore. The GCP Filestore CSI Driver Operator that manages this driver is in Technology Preview.

For more information, see [GCP Filestore CSI Driver Operator](#).

1.3.11.2. Automatic CSI migration for AWS Elastic Block Storage auto migration is generally available

Starting with OpenShift Container Platform 4.8, automatic migration for in-tree volume plugins to their equivalent Container Storage Interface (CSI) drivers became available as a Technology Preview feature. Support for Amazon Web Services (AWS) Elastic Block Storage (EBS) was provided in this feature in OpenShift Container Platform 4.8, and OpenShift Container Platform 4.12 now supports automatic migration for AWS EBS as generally available. CSI migration for AWS EBS is now enabled by default and requires no action by an administrator.

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

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

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

1.3.11.3. Automatic CSI migration for GCP PD auto migration is generally available

Starting with OpenShift Container Platform 4.8, automatic migration for in-tree volume plugins to their equivalent Container Storage Interface (CSI) drivers became available as a Technology Preview feature. Support for Google Compute Engine Persistent Disk (GCP PD) was provided in this feature in OpenShift Container Platform 4.9, and OpenShift Container Platform 4.12 now supports automatic migration for GCP PD as generally available. CSI migration for GCP PD is now enabled by default and requires no action by an administrator.

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

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

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

1.3.11.4. Updating from OpenShift Container Platform 4.12 to 4.13 and later with vSphere in-tree PVs

Updates from OpenShift Container Platform 4.12 to 4.13 and from 4.13 to 4.14 are blocked if **all** of the following conditions are true:

- CSI migration is not already enabled
- OpenShift Container Platform is not running on vSphere 7.0u3L+ or 8.0u2+
- vSphere in-tree persistent volumes (PVs) are present

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

1.3.11.5. Storage capacity tracking for pod scheduling is generally available

This new feature exposes the currently available storage capacity using **CSIStorageCapacity** objects, and enhances scheduling of pods that use Container Storage Interface (CSI) volumes with late binding. Currently, the only OpenShift Container Platform storage type that supports this features is OpenShift Data Foundation.

1.3.11.6. VMware vSphere CSI topology is generally available

OpenShift Container Platform provides the ability to deploy OpenShift Container Platform for vSphere on different zones and regions, which allows you to deploy over multiple compute clusters, thus helping to avoid a single point of failure.

For more information, see [vSphere CSI topology](#).

1.3.11.7. Local ephemeral storage resource management is generally available

The local ephemeral storage resource management features is now generally available. With this feature, you can manage local ephemeral storage by specifying requests and limits.

For more information, see [Ephemeral storage management](#).

1.3.11.8. Volume populators (Technology Preview)

Volume populators use **datasource** to enable creating pre-populated volumes.

Volume population is currently enabled, and supported as a Technology Preview feature. However, OpenShift Container Platform does not ship with any volume populators.

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

1.3.11.9. VMware vSphere CSI Driver Operator requirements

For OpenShift Container Platform 4.12, VMWare vSphere Container Storage Interface (CSI) Driver Operator requires the following minimum components installed:

- VMware vSphere version 7.0 Update 2 or later, which includes version 8.0.
- vCenter 7.0 Update 2 or later, which includes version 8.0.
- Virtual machines of hardware version 15 or later
- No third-party CSI driver already installed in the cluster

If a third-party CSI driver is present in the cluster, OpenShift Container Platform does not overwrite it. The presence of a third-party CSI driver prevents OpenShift Container Platform from upgrading to OpenShift Container Platform 4.13 or later.

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

1.3.11.10. Azure File supporting NFS is generally available

OpenShift Container Platform 4.12 supports Azure File Container Storage Interface (CSI) Driver Operator with Network File System (NFS) as generally available.

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

1.3.12. Operator lifecycle

1.3.12.1. Platform Operators (Technology Preview)

Starting in OpenShift Container Platform 4.12, Operator Lifecycle Manager (OLM) introduces the *platform Operator* type as a Technology Preview feature. The platform Operator mechanism relies on resources from the RukPak component, also introduced in OpenShift Container Platform 4.12, to source and manage content.

A platform Operator is an OLM-based Operator that can be installed during or after an OpenShift Container Platform cluster's Day 0 operations and participates in the cluster's lifecycle. As a cluster administrator, you can use platform Operators to further customize your OpenShift Container Platform installation to meet your requirements and use cases.

For more information about platform Operators, see [Managing platform Operators](#). For more information about RukPak and its resources, see [Operator Framework packaging format](#).

1.3.12.2. Controlling where an Operator is installed

By default, when you install an Operator, OpenShift Container Platform randomly installs the Operator pod to one of your worker nodes.

In OpenShift Container Platform 4.12, you can control where an Operator pod is installed by adding affinity constraints to the Operator's **Subscription** object.

For more information, see [Controlling where an Operator is installed](#).

1.3.12.3. Pod security admission synchronization for user-created openshift-* namespaces

In OpenShift Container Platform 4.12, pod security admission synchronization is enabled by default if an Operator is installed in user-created namespaces that have an **openshift-** prefix. Synchronization is enabled after a cluster service version (CSV) is created in the namespace. The synchronized label inherits the permissions of the service accounts in the namespace.

For more information, see [Security context constraint synchronization with pod security standards](#).

1.3.13. Operator development

1.3.13.1. Configuring the security context of a catalog pod

You can configure the security context of a catalog pod by using the **--security-context-config** flag on the **run bundle** and **bundle-upgrade** subcommands. The flag enables seccomp profiles to comply with pod security admission. The flag accepts the values of **restricted** and **legacy**. If you do not specify a value, the seccomp profile defaults to **restricted**. If your catalog pod cannot run with restricted permissions, set the flag to **legacy**, as shown in the following example:

```
$ operator-sdk run bundle \  
  --security-context-config=legacy
```

1.3.13.2. Validating bundle manifests for APIs removed from Kubernetes 1.25

You can now check bundle manifests for deprecated APIs removed from Kubernetes 1.25 by using the Operator Framework suite of tests with the **bundle validate** subcommand.

For example:

```
$ operator-sdk bundle validate .<bundle_dir_or_image> \  
  --select-optional suite=operatorframework \  
  --optional-values=k8s-version=1.25
```

If your Operator requests permission to use any of the APIs removed from Kubernetes 1.25, the command displays a warning message.

If any of the API versions removed from Kubernetes 1.25 are included in your Operator's cluster service version (CSV), the command displays an error message.

See [Beta APIs removed from Kubernetes 1.25](#) and the [Operator SDK CLI reference](#) for more information.

1.3.14. Machine API

1.3.14.1. Control plane machine sets

OpenShift Container Platform 4.12 introduces control plane machine sets. Control plane machine sets provide management capabilities for control plane machines that are similar to what compute machine sets provide for compute machines. For more information, see [Managing control plane machines](#).

1.3.14.2. Specifying cluster autoscaler log level verbosity

OpenShift Container Platform now supports setting the log level verbosity of the cluster autoscaler by setting the **logVerbosity** parameter in the **ClusterAutoscaler** custom resource. For more information, see the [ClusterAutoscaler resource definition](#).

1.3.14.3. Enabling Azure boot diagnostics

OpenShift Container Platform now supports enabling boot diagnostics on Azure machines that your machine set creates. For more information, see "Enabling Azure boot diagnostics" for [compute machines](#) or [control plane machines](#).

1.3.15. Machine Config Operator

1.3.15.1. RHCOS image layering

Red Hat Enterprise Linux CoreOS (RHCOS) image layering allows you to add new images on top of the base RHCOS image. This layering does not modify the base RHCOS image. Instead, it creates a *custom layered image* that includes all RHCOS functionality and adds additional functionality to specific nodes in the cluster.

Currently, RHCOS image layering allows you to work with Customer Experience and Engagement (CEE) to obtain and apply Hotfix packages on top of your RHCOS image, based on the [Red Hat Hotfix policy](#). It is planned for future releases that you can use RHCOS image layering to incorporate third-party software packages such as Libreswan or numactl.

For more information, see [RHCOS image layering](#).

1.3.16. Nodes

1.3.16.1. Updating the interface-specific safe list (Technology Preview)

OpenShift Container Platform now supports updating the default interface-specific safe **sysctls**.

You can add or remove **sysctls** from the predefined list. When you add **sysctls**, they can be set across all nodes. Updating the interface-specific safe **sysctls** list is a Technology Preview feature only.

For more information, see [Updating the interface-specific safe sysctls list](#).

1.3.16.2. Cron job time zones (Technology Preview)

Setting a time zone for a cron job schedule is now offered as a [Technology Preview](#). If a time zone is not specified, the Kubernetes controller manager interprets the schedule relative to its local time zone.

For more information, see [Creating cron jobs](#).

1.3.16.3. Linux Control Group version 2 promoted to Technology Preview

OpenShift Container Platform support for [Linux Control Group version 2](#) (cgroup v2) has been promoted to Technology Preview. cgroup v2 is the next version of the kernel [control groups](#). cgroups v2 offers multiple improvements, including a unified hierarchy, safer sub-tree delegation, new features such as [Pressure Stall Information](#), and enhanced resource management and isolation. For more information, see [Enabling Linux Control Group version 2 \(cgroup v2\)](#).

1.3.16.4. crun container runtime (Technology Preview)

OpenShift Container Platform now supports the crun container runtime in Technology Preview. You can switch between the crun container runtime and the default container runtime as needed by using a **ContainerRuntimeConfig** custom resource (CR). For more information, see [About the container engine and container runtime](#).

1.3.16.5. Self Node Remediation Operator enhancements

OpenShift Container Platform now supports control plane fencing by the Self Node Remediation Operator. In the event of node failure, you can follow remediation strategies on both worker nodes and control plane nodes. For more information, see the [Workload Availability for Red Hat OpenShift](#) documentation.

1.3.16.6. Node Health Check Operator enhancements

OpenShift Container Platform now supports control plane fencing on the Node Health Check Operator. In the event of node failure, you can follow remediation strategies on both worker nodes and control plane nodes. For more information, see the [Workload Availability for Red Hat OpenShift](#) documentation.

The Node Health Check Operator now also includes a web console plugin for managing Node Health Checks. For more information, see the [Workload Availability for Red Hat OpenShift](#) documentation.

For installing or updating to the latest version of the Node Health Check Operator, use the **stable** subscription channel. For more information, see the [Workload Availability for Red Hat OpenShift](#) documentation.

1.3.17. Monitoring

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

1.3.17.1. Updates to monitoring stack components and dependencies

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

- kube-state-metrics to 2.6.0
- node-exporter to 1.4.0
- prom-label-proxy to 0.5.0
- Prometheus to 2.39.1

- prometheus-adapter to 0.10.0
- prometheus-operator to 0.60.1
- Thanos to 0.28.1

1.3.17.2. Changes to alerting rules



NOTE

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

- **New**
 - Added the **TelemeterClientFailures** alert, which triggers when a cluster tries and fails to submit Telemetry data at a certain rate over a period of time. The alert fires when the rate of failed requests reaches 20% of the total rate of requests within a 15-minute window.
- **Changed**
 - The **KubeAggregatedAPIDown** alert now waits 900 seconds rather than 300 seconds before sending a notification.
 - The **NodeClockNotSynchronising** and **NodeClockSkewDetected** alerts now only evaluate metrics from the **node-exporter** job.
 - The **NodeRAIDDegraded** and **NodeRAIDDiskFailure** alerts now include a device label filter to match only the value returned by **mmcblk.p.+|nvme.+|sd.+|vd.+|xvd.+|dm- .+|dasd.+**.
 - The **PrometheusHighQueryLoad** and **ThanosQueryOverload** alerts now also trigger when a high querying load exists on the query layer.

1.3.17.3. New option to specify pod topology spread constraints for monitoring components

You can now use pod topology spread constraints to control how Prometheus, Thanos Ruler, and Alertmanager pods are spread across a network topology when OpenShift Container Platform pods are deployed in multiple availability zones.

1.3.17.4. New option to improve data consistency for Prometheus Adapter

You can now configure an optional kubelet service monitor for Prometheus Adapter (PA) that improves data consistency across multiple autoscaling requests. Enabling this service monitor eliminates the possibility that two queries sent at the same time to PA might yield different results because the underlying PromQL queries executed by PA might be on different Prometheus servers.

1.3.17.5. Update to Alertmanager configuration for additional secret keys

With this release, if you configure an Alertmanager secret to hold additional keys and if the Alertmanager configuration references these keys as files (such as templates, TLS certificates, or tokens), your configuration settings must point to these keys by using an absolute path rather than a relative path. These keys are available under the **/etc/alertmanager/config** directory. In earlier releases of OpenShift Container Platform, you could use relative paths in your configuration to point to these keys because the Alertmanager configuration file was located in the same directory as the keys.



IMPORTANT

If you are upgrading to OpenShift Container Platform 4.12 and have specified relative paths for additional Alertmanager secret keys that are referenced as files, you must change these relative paths to absolute paths in your Alertmanager configuration. Otherwise, alert receivers that use the files will fail to deliver notifications.

1.3.18. New Network Observability Operator

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

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

1.3.19. Scalability and performance

1.3.19.1. Disabling realtime using workload hints removes Receive Packet Steering from the cluster

At the cluster level by default, a systemd service sets a Receive Packet Steering (RPS) mask for virtual network interfaces. The RPS mask routes interrupt requests from virtual network interfaces according to the list of reserved CPUs defined in the performance profile. At the container level, a **CRI-O** hook script also sets an RPS mask for all virtual network devices.

With this update, if you set **spec.workloadHints.realTime** in the performance profile to **False**, the system also disables both the systemd service and the **CRI-O** hook script which set the RPS mask. The system disables these RPS functions because RPS is typically relevant to use cases requiring low-latency, realtime workloads only.

To retain RPS functions even when you set **spec.workloadHints.realTime** to **False**, see the *RPS Settings* section of the Red Hat Knowledgebase solution [Performance addons operator advanced configuration](#).

For more information about configuring workload hints, see [Understanding workload hints](#).

1.3.19.2. Tuned profile

The **tuned** profile now defines the **fs.aio-max-nr sysctl** value by default, improving asynchronous I/O performance for default node profiles.

1.3.19.3. Support for new kernel features and options

The low latency tuning has been updated to use the latest kernel features and options. The fix for [2117780](#) introduced a new per-CPU **kthread**, **ktimers**. This thread must be pinned to the proper CPU cores. With this update, there is no functional change; the isolation of the workload is the same. For more information, see [2102450](#).

1.3.19.4. Power-saving configurations

In OpenShift Container Platform 4.12, by enabling C-states and OS-controlled P-states, you can use

different power-saving configurations for critical and non-critical workloads. You can apply the configurations through the new **perPodPowerManagement** workload hint, and the **cpu-states.crio.io** and **cpu-freq-governor.crio.io** CRI-O annotations. For more information about the feature, see [Power-saving configurations](#).

1.3.19.5. Expanding Single-node OpenShift clusters with worker nodes using GitOps ZTP (Technology Preview)

In OpenShift Container Platform 4.11, a feature allowing you to manually add worker nodes to single-node OpenShift clusters was introduced. This feature is now also available in GitOps ZTP.

For more information, see [Adding worker nodes to single-node OpenShift clusters with GitOps ZTP](#).

1.3.19.6. Factory-precaching-cli tool to reduce OpenShift Container Platform and Operator deployment times (Technology Preview)

In OpenShift Container Platform 4.12, you can use the factory-precaching-cli tool to pre-cache OpenShift Container Platform and Operator images on a server at the factory, and then you can include the pre-cached server to the site for deployment. For more information about the factory-precaching-cli tool, see [Pre-caching images for single-node OpenShift deployments](#).

1.3.19.7. Zero touch provisioning (ZTP) integration of the factory-precaching-cli tool (Technology Preview)

In OpenShift Container Platform 4.12, you can use the factory-precaching-cli tool in the GitOps ZTP workflow. For more information, see [Pre-caching images for single-node OpenShift deployments](#).

1.3.19.8. Node tuning in a hosted cluster (Technology Preview)

You can now configure OS-level tuning for nodes in a hosted cluster by using the Node Tuning Operator. To configure node tuning, you can create config maps in the management cluster that contain **Tuned** objects, and reference those config maps in your node pools. The tuning configuration that is defined in the **Tuned** objects is applied to the nodes in the node pool. For more information, see [Configuring node tuning in a hosted cluster](#).

1.3.19.9. Kernel module management Operator

The kernel module management (KMM) Operator replaces the Special Resource Operator (SRO). KMM includes the following features for connected environments only:

- Hub and spoke support for edge deployments
- Pre-flight checks for upgrade support
- Secure boot kernel module signing
- Must gather logs to assist with troubleshooting
- Binary firmware deployment

1.3.19.10. Hub and spoke cluster support (Technology Preview)

For hub and spoke deployments in an environment that can access the internet, you can use the kernel module management (KMM) Operator deployed in the hub cluster to manage the deployment of the required kernel modules to one or more managed clusters.

1.3.19.11. Topology Aware Lifecycle Manager (TALM)

Topology Aware Lifecycle Manager (TALM) now provides more detailed status information and messages, and redesigned conditions. You can use the **ClusterLabelSelector** field for greater flexibility in selecting clusters for update. You can use timeout settings to determine what happens if an update fails for a cluster, for example, skipping the failing cluster and continuing to upgrade other clusters, or stopping policy remediation for all clusters. For more information see [Topology Aware Lifecycle Manager for cluster updates](#).

1.3.19.12. Mount namespace encapsulation (Technology Preview)

Encapsulation is the process of moving all Kubernetes-specific mount points to an alternative namespace to reduce the visibility and performance impact of a large number of mount points in the default namespace. Previously, mount namespace encapsulation has been deployed transparently in OpenShift Container Platform specifically for Distributed Units (DUs) installed using GitOps ZTP. In OpenShift Container Platform v4.12, this functionality is now available as a configurable option.

A standard host operating system uses systemd to constantly scan all mount namespaces: both the standard Linux mounts and the numerous mounts that Kubernetes uses to operate. The current implementation of Kubelet and CRI-O both use the top-level namespace for all container and Kubelet mount points. Encapsulating these container-specific mount points in a private namespace reduces systemd overhead and enhances CPU performance. Encapsulation can also improve security, by storing Kubernetes-specific mount points in a location safe from inspection by unprivileged users.

For more information, see [Optimizing CPU usage with mount namespace encapsulation](#).

1.3.19.13. Changing the workload partitioning CPU set in single-node OpenShift clusters that are deployed with GitOps ZTP

You can configure the workload partitioning CPU set in single-node OpenShift clusters that you deploy with GitOps ZTP. To do this, you specify cluster management CPU resources with the **cpuset** field of the **SiteConfig** custom resource (CR) and the **reserved** field of the group **PolicyGenTemplate** CR. The value that you set for **cpuset** should match the value set in the cluster **PerformanceProfile** CR **.spec.cpu.reserved** field for workload partitioning.

For more information, see [Workload partitioning](#).

1.3.19.14. RHACM hub template functions now available for use with GitOps ZTP

Hub template functions are now available for use with GitOps ZTP using Red Hat Advanced Cluster Management (RHACM) and Topology Aware Lifecycle Manager (TALM). Hub-side cluster templates reduce the need to create separate policies for many clusters with similar configurations but with different values. For more information, see [Using hub templates in PolicyGenTemplate CRs](#).

1.3.19.15. ArgoCD managed cluster limits

RHACM uses **SiteConfig** CRs to generate the Day 1 managed cluster installation CRs for ArgoCD. Each ArgoCD application can manage a maximum of 300 **SiteConfig** CRs. For more information, see [Configuring the hub cluster with ArgoCD](#).

1.3.19.16. GitOps ZTP support for configuring policy compliance evaluation timeouts in PolicyGenTemplate CRs

In GitOps ZTP v4.11+, a default policy compliance evaluation timeout value is available for use in **PolicyGenTemplate** custom resources (CRs). This value specifies how long the related **ConfigurationPolicy** CR can be in a state of policy compliance or non-compliance before RHACM re-evaluates the applied cluster policies.

Optionally, you can now override the default evaluation intervals for all policies in **PolicyGenTemplate** CRs.

For more information, see [Configuring policy compliance evaluation timeouts for PolicyGenTemplate CRs](#).

1.3.19.17. Specifying the platform type for managed clusters

The Assisted Installer currently supports the following OpenShift Container Platform platforms:

- **BareMetal**
- **VSphere**
- **None**

Single-node OpenShift does not support **VSphere**.

1.3.19.18. Configuring the hub cluster to use unauthenticated registries

This release supports the use of unauthenticated registries when configuring the hub cluster. Registries that do not require authentication are listed under **spec.unauthenticatedRegistries** in the **AgentServiceConfig** resource. Any registry on this list is not required to have an entry in the pull secret used for the spoke cluster installation. **assisted-service** validates the pull secret by making sure it contains the authentication information for every image registry used for installation.

For more information, see [Configuring the hub cluster to use unauthenticated registries](#).

1.3.19.19. Ironic agent mirroring in disconnected GitOps ZTP installations

For disconnected installations using GitOps ZTP, if you are deploying OpenShift Container Platform version 4.11 or earlier to a spoke cluster with converged flow enabled, you must mirror the default Ironic agent image to the local image repository. The default Ironic agent images are the following:

- AMD64 Ironic agent image: **quay.io/openshift-release-dev/ocp-v4.0-art-dev@sha256:d3f1d4d3cd5fbcf1b9249dd71d01be4b901d337fdc5f8f66569eb71df4d9d446**
- AArch64 Ironic agent image: **quay.io/openshift-release-dev/ocp-v4.0-art-dev@sha256:cb0edf19fffc17f542a7efae76939b1e9757dc75782d4727fb0aa77ed5809b43**

For more information about mirroring images, see [Mirroring the OpenShift Container Platform image repository](#).

1.3.19.20. Configuring kernel arguments for the Discovery ISO by using GitOps ZTP

OpenShift Container Platform now supports specifying kernel arguments for the Discovery ISO in GitOps ZTP deployments. In both manual and automated GitOps ZTP deployments, the Discovery ISO

is part of the OpenShift Container Platform installation process on managed bare-metal hosts. You can now edit the **InfraEnv** resource to specify kernel arguments for the Discovery ISO. This is useful for cluster installations with specific environmental requirements. For example, you can define the **rd.net.timeout.carrier** kernel argument to help configure the cluster for static networking.

For more information about how to specify kernel arguments, see [Configuring kernel arguments for the Discovery ISO by using GitOps ZTP](#) and [Configuring kernel arguments for the Discovery ISO for manual installations by using GitOps ZTP](#).

1.3.19.21. Deploy heterogeneous spoke clusters from a hub cluster

With this update, you can create OpenShift Container Platform mixed-architecture clusters, also known as heterogeneous clusters, that feature hosts with both AMD64 and AArch64 CPU architectures. You can deploy a heterogeneous spoke cluster from a hub cluster managed by Red Hat Advanced Cluster Management (RHACM). To create a heterogeneous spoke cluster, add an AArch64 worker node to a deployed AMD64 cluster.

To add an AArch64 worker node to a deployed AMD64 cluster, you can specify the AArch64 architecture, the multi-architecture release image, and the operating system required for the node by using an **InfraEnv** custom resource (CR). You can then provision the AArch64 worker node to the AMD64 cluster by using the Assisted Installer API and the **InfraEnv** CR.

1.3.19.22. HTTP transport replaces AMQP for PTP and bare-metal events (Technology Preview)

HTTP is now the default transport in the PTP and bare-metal events infrastructure. AMQ Interconnect is end of life (EOL) from 30 June 2024.

For more information, see [About the PTP fast event notifications framework](#).

1.3.20. Insights Operator

1.3.20.1. Insights alerts

In OpenShift Container Platform 4.12, active Insights recommendations are now presented to the user as alerts. You can view and configure these alerts with Alertmanager.

1.3.20.2. Insights Operator data collection enhancements

In OpenShift Container Platform 4.12, the Insights Operator now collects the following metrics:

- **console_helm_uninstalls_total**
- **console_helm_upgrades_total**

1.3.21. Authentication and authorization

1.3.21.1. Application credentials on RHOSP

You can now specify [application credentials](#) in the **clouds.yaml** files of clusters that run on Red Hat OpenStack Platform (RHOSP). Application credentials are an alternative to embedding user account details in configuration files. As an example, see the following section of a **clouds.yaml** file that includes user account details:

```

clouds:
  openstack:
    auth:
      auth_url: https://127.0.0.1:13000
      password: thepassword
      project_domain_name: Default
      project_name: theprojectname
      user_domain_name: Default
      username: theusername
      region_name: regionOne

```

Compare that section to one that uses application credentials:

```

clouds:
  openstack:
    auth:
      auth_url: https://127.0.0.1:13000
      application_credential_id: '5dc185489adc4b0f854532e1af81ffe0'
      application_credential_secret:
'PDCTKans2bPBbaEqBLiT_lajG8e5J_nJB4kvQHjaAy6ufhod0ZI0NkNoBzjn_bWSYzk587ielGSIT11c4pV
ehA'
      auth_type: "v3applicationcredential"
      region_name: regionOne

```

To use application credentials with your cluster as a RHOSP administrator, create the credentials. Then, use them in a **clouds.yaml** file when you install a cluster. Alternatively, you can create the **clouds.yaml** file and rotate it into an existing cluster.

1.3.22. Hosted control planes (Technology Preview)

1.3.22.1. HyperShift API beta release now available

The default version for the **hypershift.openshift.io** API, which is the API for hosted control planes on OpenShift Container Platform, is now v1beta1. Currently, for an existing cluster, the move from alpha to beta is not supported.

1.3.22.2. Versioning for hosted control planes

With each major, minor, or patch version release of OpenShift Container Platform, the HyperShift Operator is released. The HyperShift command-line interface (CLI) is released as part of each HyperShift Operator release.

The **HostedCluster** and **NodePool** API resources are available in the beta version of the API and follow a similar policy to [OpenShift Container Platform](#) and [Kubernetes](#).

1.3.22.3. Backing up and restoring etcd on a hosted cluster

If you use hosted control planes on OpenShift Container Platform, you can back up and restore etcd by taking a snapshot of etcd and uploading it to a location where you can retrieve it later, such as an S3 bucket. Later, if needed, you can restore the snapshot. For more information, see [Backing up and restoring etcd on a hosted cluster](#).

1.3.22.4. Disaster recovery for a hosted cluster within an AWS region

In a situation where you need disaster recovery for a hosted cluster, you can recover the hosted cluster to the same region within AWS. For more information, see [Disaster recovery for a hosted cluster within an AWS region](#).

1.3.23. Red Hat Virtualization (RHV)

This release provides several updates to Red Hat Virtualization (RHV). With this release:

- The oVirt CSI driver logging was revised with new error messages to improve the clarity and readability of the logs.
- The cluster API provider automatically updates oVirt and Red Hat Virtualization (RHV) credentials when they are changed in OpenShift Container Platform.

1.4. NOTABLE TECHNICAL CHANGES

OpenShift Container Platform 4.12 introduces the following notable technical changes.

AWS Security Token Service regional endpoints

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

cert-manager Operator general availability

cert-manager Operator is generally available in OpenShift Container Platform 4.12.

Credentials requests directory parameter for deleting GCP resources with the Cloud Credential Operator utility

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

Future restricted enforcement for pod security admission

Currently, pod security violations are shown as warnings and logged in the audit logs, but do not cause the pod to be rejected.

Global restricted enforcement for pod security admission is currently planned for the next minor release of OpenShift Container Platform. When this restricted enforcement is enabled, pods with pod security violations will be rejected.

To prepare for this upcoming change, ensure that your workloads match the pod security admission profile that applies to them. Workloads that are not configured according to the enforced security standards defined globally or at the namespace level will be rejected. The **restricted-v2** SCC admits workloads according to the [Restricted](#) Kubernetes definition.

If you are receiving pod security violations, see the following resources:

- See [Identifying pod security violations](#) for information about how to find which workloads are causing pod security violations.
- See [Security context constraint synchronization with pod security standards](#) to understand when pod security admission label synchronization is performed. Pod security admission labels are not synchronized in certain situations, such as the following situations:
 - The workload is running in a system-created namespace that is prefixed with **openshift-**.
 - The workload is running on a pod that was created directly without a pod controller.

- If necessary, you can set a custom admission profile on the namespace or pod by setting the **pod-security.kubernetes.io/enforce** label.

Catalog sources and restricted pod security admission enforcement

Catalog sources built using the SQLite-based catalog format and a version of the **opm** CLI tool released before OpenShift Container Platform 4.11 cannot run under restricted pod security enforcement.

In OpenShift Container Platform 4.12, namespaces do not have restricted pod security enforcement by default and the default catalog source security mode is set to **legacy**.

If you do not want to run your SQLite-based catalog source pods under restricted pod security enforcement, you do not need to update your catalog source in OpenShift Container Platform 4.12. However, to ensure your catalog sources run in future OpenShift Container Platform releases, you must update your catalog sources to run under restricted pod security enforcement.

As a catalog author, you can enable compatibility with restricted pod security enforcement by completing either of the following actions:

- Migrate your catalog to the file-based catalog format.
- Update your catalog image with a version of the **opm** CLI tool released with OpenShift Container Platform 4.11 or later.

If you do not want to update your SQLite database catalog image or migrate your catalog to the file-based catalog format, you can configure your catalog to run with elevated permissions.

For more information, see [Catalog sources and pod security admission](#).

Operator SDK 1.25.4

OpenShift Container Platform 4.12 supports Operator SDK 1.25.4. See [Installing the Operator SDK CLI](#) to install or update to this latest version.



NOTE

Operator SDK 1.25.4 supports Kubernetes 1.25.

For more information, see [Beta APIs removed from Kubernetes 1.25](#) and [Validating bundle manifests for APIs removed from Kubernetes 1.25](#).

If you have Operator projects that were previously created or maintained with Operator SDK 1.22.2, update your projects to keep compatibility with Operator SDK 1.25.4.

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

LVM Operator is now called Logical Volume Manager Storage

The LVM Operator that was previously delivered with Red Hat OpenShift Data Foundation requires installation through the OpenShift Data Foundation. In OpenShift Container Platform v4.12, the LVM

Operator has been renamed *Logical Volume Manager Storage*. Now, you install it as a standalone Operator from the OpenShift Operator catalog. Logical Volume Manager Storage provides dynamic provisioning of block storage on a single, limited resources single-node OpenShift cluster.

End of support for RHOSP 16.1

OpenShift Container Platform no longer supports RHOSP 16.1 as a deployment target. See [OpenShift Container Platform on Red Hat OpenStack Platform Support Matrix](#) for complete details.

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 4.12, 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:

- *General Availability*
- *Deprecated*
- *Removed*

Operator deprecated and removed features

Table 1.2. Operator deprecated and removed tracker

Feature	4.10	4.11	4.12
SQLite database format for Operator catalogs	Deprecated	Deprecated	Deprecated

Images deprecated and removed features

Table 1.3. Images deprecated and removed tracker

Feature	4.10	4.11	4.12
ImageChangesInProgress condition for Cluster Samples Operator	Deprecated	Deprecated	Deprecated
MigrationInProgress condition for Cluster Samples Operator	Deprecated	Deprecated	Deprecated
Removal of Jenkins images from install payload	General Availability	Removed	Removed

Monitoring deprecated and removed features

Table 1.4. Monitoring deprecated and removed tracker

Feature	4.10	4.11	4.12
Grafana component in monitoring stack	Deprecated	Removed	Removed
Access to Prometheus and Grafana UIs in monitoring stack	Deprecated	Removed	Removed

Installation deprecated and removed features

Table 1.5. Installation deprecated and removed tracker

Feature	4.10	4.11	4.12
vSphere 6.x or earlier	Deprecated	Removed	Removed
vSphere 7.0 Update 1 or earlier	General Availability	Deprecated	Deprecated
VMware ESXi 6.x or earlier	Deprecated	Removed	Removed
VMware ESXi 7.0 Update 1 or earlier	General Availability	Deprecated	Deprecated
CoreDNS wildcard queries for the cluster.local domain	General Availability	General Availability	Deprecated
ingressVIP and apiVIP settings in the install-config.yaml file for installer-provisioned infrastructure clusters	General Availability	General Availability	Deprecated

Updating clusters deprecated and removed features

Table 1.6. Updating clusters deprecated and removed tracker

Feature	4.10	4.11	4.12
Virtual hardware version 13	Deprecated	Removed	Removed

Storage deprecated and removed features

Table 1.7. Storage deprecated and removed tracker

Feature	4.10	4.11	4.12
Snapshot.storage.k8s.io/v1beta1 API endpoint	Deprecated	Removed	Removed
Persistent storage using FlexVolume	Deprecated	Deprecated	Deprecated

Authentication and authorization deprecated and removed features

Table 1.8. Authentication and authorization deprecated and removed tracker

Feature	4.10	4.11	4.12
Automatic generation of service account token secrets	General Availability	Removed	Removed

Specialized hardware and driver enablement deprecated and removed features

Table 1.9. Specialized hardware and driver enablement deprecated and removed tracker

Feature	4.10	4.11	4.12
Special Resource Operator (SRO)	Technology Preview	Technology Preview	Removed

Multi-architecture deprecated and removed features

Table 1.10. Multi-architecture deprecated and removed tracker

Feature	4.10	4.11	4.12
IBM POWER8 all models (ppc64le)	General Availability	General Availability	Deprecated
IBM IBM POWER9 AC922 (ppc64le)	General Availability	General Availability	Deprecated
IBM IBM POWER9 IC922 (ppc64le)	General Availability	General Availability	Deprecated
IBM IBM POWER9 LC922 (ppc64le)	General Availability	General Availability	Deprecated
IBM z13 all models (s390x)	General Availability	General Availability	Deprecated
IBM LinuxONE Emperor (s390x)	General Availability	General Availability	Deprecated
IBM LinuxONE Rockhopper (s390x)	General Availability	General Availability	Deprecated
AMD64 (x86_64) v1 CPU	General Availability	General Availability	Deprecated

Networking deprecated and removed features

Table 1.11. Networking deprecated and removed tracker

Feature	4.10	4.11	4.12
Kuryr on RHOSP	General Availability	General Availability	Deprecated

Web console deprecated and removed features

Table 1.12. Web console deprecated and removed tracker

Feature	4.10	4.11	4.12
Multicluster console (Technology Preview)	REM	REM	REM

1.5.1. Deprecated features

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

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

1.5.1.2. Wildcard DNS queries for the **cluster.local** domain are deprecated

CoreDNS will stop supporting wildcard DNS queries for names under the **cluster.local** domain. These queries will resolve in OpenShift Container Platform 4.12 as they do in earlier versions, but support will be removed from a future OpenShift Container Platform release.

1.5.1.3. Specific hardware models on **ppc64le**, **s390x**, and **x86_64 v1** CPU architectures are deprecated

In OpenShift Container Platform 4.12, support for RHCOS functionality is deprecated for:

- IBM POWER8 all models (ppc64le)
- IBM POWER9 AC922 (ppc64le)
- IBM POWER9 IC922 (ppc64le)
- IBM POWER9 LC922 (ppc64le)
- IBM z13 all models (s390x)
- LinuxONE Emperor (s390x)
- LinuxONE Rockhopper (s390x)
- AMD64 (x86_64) v1 CPU

While these hardware models remain fully supported in OpenShift Container Platform 4.12, Red Hat recommends that you use later hardware models.

1.5.1.4. Kuryr support for clusters that run on RHOSP

In OpenShift Container Platform 4.12, support for Kuryr on clusters that run on RHOSP is deprecated. Support will be removed no earlier than OpenShift Container Platform 4.14.

1.5.2. Removed features

1.5.2.1. Beta APIs removed from Kubernetes 1.25

Kubernetes 1.25 removed the following deprecated APIs, so you must migrate manifests and API clients to use the appropriate API version. For more information about migrating removed APIs, see the [Kubernetes documentation](#).

Table 1.13. APIs removed from Kubernetes 1.25

Resource	Removed API	Migrate to	Notable changes
CronJob	batch/v1beta1	batch/v1	No
EndpointSlice	discovery.k8s.io/v1beta1	discovery.k8s.io/v1	Yes
Event	events.k8s.io/v1beta1	events.k8s.io/v1	Yes
HorizontalPodAutoscaler	autoscaling/v2beta1	autoscaling/v2	No
PodDisruptionBudget	policy/v1beta1	policy/v1	Yes
PodSecurityPolicy	policy/v1beta1	Pod Security Admission ^[1]	Yes
RuntimeClass	node.k8s.io/v1beta1	node.k8s.io/v1	No

1. For more information about pod security admission in OpenShift Container Platform, see [Understanding and managing pod security admission](#).

1.5.2.2. Empty file and stdout support for the oc registry login command

The **--registry-config** and **--to option** options for the **oc registry login** command now stop accepting empty files. These options continue to work with files that do not exist. The ability to write output to **-** (stdout) is also removed.

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

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

1.5.2.4. OpenShift CLI (oc) commands have been removed

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

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

1.5.2.5. Grafana component removed from monitoring stack

The Grafana component is no longer a part of the OpenShift Container Platform 4.12 monitoring stack. As an alternative, go to **Observe** → **Dashboards** in the OpenShift Container Platform web console to view monitoring dashboards.

1.5.2.6. Prometheus and Grafana user interface access removed from monitoring stack

Access to the third-party Prometheus and Grafana user interfaces have been removed from the OpenShift Container Platform 4.12 monitoring stack. As an alternative, click **Observe** in the OpenShift Container Platform web console to view alerting, metrics, dashboards, and metrics targets for monitoring components.

1.5.2.7. Support for virtual hardware version 13 is removed

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

1.5.2.8. Support for snapshot v1beta1 API endpoint is removed

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

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

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

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

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

1.5.2.11. Removal of Jenkins images from install payload

- OpenShift Container Platform 4.11 moves the "OpenShift Jenkins" and "OpenShift Agent Base" images to the **ocp-tools-4** repository at **registry.redhat.io** so that Red Hat can produce and update the images outside the OpenShift Container Platform lifecycle. Previously, these images were in the OpenShift Container Platform install payload and the **openshift4** repository at **registry.redhat.io**. For more information, see [OpenShift Jenkins](#).

- OpenShift Container Platform 4.11 removes "OpenShift Jenkins Maven" and "NodeJS Agent" images from its payload. Previously, OpenShift Container Platform 4.10 deprecated these images. Red Hat no longer produces these images, and they are not available from the **ocp-tools-4** repository at registry.redhat.io.

However, upgrading to OpenShift Container Platform 4.11 does not remove "OpenShift Jenkins Maven" and "NodeJS Agent" images from 4.10 and earlier releases. And Red Hat provides bug fixes and support for these images through the end of the 4.10 release lifecycle, in accordance with the [OpenShift Container Platform lifecycle policy](#).

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

1.5.3. Future Kubernetes API removals

The next minor release of OpenShift Container Platform is expected to use Kubernetes 1.26. Currently, Kubernetes 1.26 is scheduled to remove several deprecated APIs.

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

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

1.6. BUG FIXES

API Server and Authentication

- Previously, the Cluster Authentication Operator state was set to **progressing = false** after receiving a **workloadIsBeingUpdatedTooLong** error. At the same time, **degraded = false** was kept for the time of the **inertia** defined. Consequently, the shortened amount of progressing and increased time of degradation would create a situation where **progressing = false** and **degraded = false** were set prematurely. This caused inconsistent OpenShift CI tests because a healthy state was assumed, which was incorrect. This issue has been fixed by removing the **progressing = false** setting after the **workloadIsBeingUpdatedTooLong** error is returned. Now, because there is no **progressing = false** state, OpenShift CI tests are more consistent. ([BZ#2111842](#))

Bare Metal Hardware Provisioning

- In recent versions of server firmware the time between server operations has increased. This causes timeouts during installer-provisioned infrastructure installations when the OpenShift Container Platform installation program waits for a response from the Baseboard Management Controller (BMC). The new **python3-sushy** release increases the number of server side attempts to contact the BMC. This update accounts for the extended waiting time and avoids timeouts during installation. ([OCPBUGS-4097](#))
- Before this update, the Ironic provisioning service did not support Baseboard Management Controllers (BMC) that use weak eTags combined with strict eTag validation. By design, if the BMC provides a weak eTag, Ironic returns two eTags: the original eTag and the original eTag converted to the strong format for compatibility with BMC that do not support weak eTags. Although Ironic can send two eTags, BMC using strict eTag validation rejects such requests due to the presence of the second eTag. As a result, on some older server hardware, bare-metal provisioning failed with the following error: **HTTP 412 Precondition Failed**. In OpenShift Container Platform 4.12 and later, this behavior changes and Ironic no longer attempts to send two eTags in cases where a weak eTag is provided. Instead, if a Redfish request dependent on an

eTag fails with an eTag validation error, IroniC retries the request with known workarounds. This minimizes the risk of bare-metal provisioning failures on machines with strict eTag validation. ([OCPBUGS-3479](#))

- Before this update, when a Redfish system features a Settings URI, the IroniC provisioning service always attempts to use this URI to make changes to boot-related BIOS settings. However, bare-metal provisioning fails if the Baseboard Management Controller (BMC) features a Settings URI but does not support changing a particular BIOS setting by using this Settings URI. In OpenShift Container Platform 4.12 and later, if a system features a Settings URI, IroniC verifies that it can change a particular BIOS setting by using the Settings URI before proceeding. Otherwise, IroniC implements the change by using the System URI. This additional logic ensures that IroniC can apply boot-related BIOS setting changes and bare-metal provisioning can succeed. ([OCPBUGS-2052](#))

Builds

- By default, Buildah prints steps to the log file, including the contents of environment variables, which might include [build input secrets](#). Although you can use the **--quiet** build argument to suppress printing of those environment variables, this argument isn't available if you use the source-to-image (S2I) build strategy. The current release fixes this issue. To suppress printing of environment variables, set the **BUILDAH_QUIET** environment variable in your build configuration:

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

([BZ#2099991](#))

Cloud Compute

- Previously, instances were not set to respect the GCP infrastructure default option for automated restarts. As a result, instances could be created without using the infrastructure default for automatic restarts. This sometimes meant that instances were terminated in GCP but their associated machines were still listed in the **Running** state because they did not automatically restart. With this release, the code for passing the automatic restart option has been improved to better detect and pass on the default option selection from users. Instances now use the infrastructure default properly and are automatically restarted when the user requests the default functionality. ([OCPBUGS-4504](#))
- The **v1beta1** version of the **PodDisruptionBudget** object is now deprecated in Kubernetes. With this release, internal references to **v1beta1** are replaced with **v1**. This change is internal to the cluster autoscaler and does not require user action beyond the advice in the [Preparing to upgrade to OpenShift Container Platform 4.12](#) Red Hat Knowledgebase Article. ([OCPBUGS-1484](#))
- Previously, the GCP machine controller reconciled the state of machines every 10 hours. Other providers set this value to 10 minutes so that changes that happen outside of the Machine API system are detected within a short period. The longer reconciliation period for GCP could cause unexpected issues such as missing certificate signing requests (CSR) approvals due to an external IP address being added but not detected for an extended period. With this release, the GCP machine controller is updated to reconcile every 10 minutes to be consistent with other platforms and so that external changes are picked up sooner. ([OCPBUGS-4499](#))

- Previously, due to a deployment misconfiguration for the Cluster Machine Approver Operator, enabling the **TechPreviewNoUpgrade** feature set caused errors and sporadic Operator degradation. Because clusters with the **TechPreviewNoUpgrade** feature set enabled use two instances of the Cluster Machine Approver Operator and both deployments used the same set of ports, there was a conflict that lead to errors for single-node topology. With this release, the Cluster Machine Approver Operator deployment is updated to use a different set of ports for different deployments. ([OCPBUGS-2621](#))
- Previously, the scale from zero functionality in Azure relied on a statically compiled list of instance types mapping the name of the instance type to the number of CPUs and the amount of memory allocated to the instance type. This list grew stale over time. With this release, information about instance type sizes is dynamically gathered from the Azure API directly to prevent the list from becoming stale. ([OCPBUGS-2558](#))
- Previously, Machine API termination handler pods did not start on spot instances. As a result, pods that were running on tainted spot instances did not receive a termination signal if the instance was terminated. This could result in loss of data in workload applications. With this release, the Machine API termination handler deployment is modified to tolerate the taints and pods running on spot instances with taints now receive termination signals. ([OCPBUGS-1274](#))
- Previously, error messages for Azure clusters did not explain that it is not possible to create new machines with public IP addresses for a disconnected install that uses only the internal publish strategy. With this release, the error message is updated for improved clarity. ([OCPBUGS-519](#))
- Previously, the Cloud Controller Manager Operator did not check the **cloud-config** configuration file for AWS clusters. As a result, it was not possible to pass additional settings to the AWS cloud controller manager component by using the configuration file. With this release, the Cloud Controller Manager Operator checks the infrastructure resource and parses references to the **cloud-config** configuration file so that users can configure additional settings. ([BZ#2104373](#))
- Previously, when Azure added new instance types and enabled accelerated networking support on instance types that previously did not have it, the list of Azure instances in the machine controller became outdated. As a result, the machine controller could not create machines with instance types that did not previously support accelerated networking, even if they support this feature on Azure. With this release, the required instance type information is retrieved from Azure API before the machine is created to keep it up to date so the machine controller is able to create machines with new and updated instance types. This fix also applies to any instance types that are added in the future. ([BZ#2108647](#))
- Previously, the cluster autoscaler did not respect the AWS, IBM Cloud, and Alibaba Cloud topology labels for the CSI drivers when using the Cluster API provider. As a result, nodes with the topology label were not processed properly by the autoscaler when attempting to balance nodes during a scale-out event. With this release, the autoscaler's custom processors are updated so that it respects this label. The autoscaler can now balance similar node groups that are labeled by the AWS, IBM Cloud, or Alibaba CSI labels. ([BZ#2001027](#))
- Previously, Power VS cloud providers were not capable of fetching the machine IP address from a DHCP server. Changing the IP address did not update the node, which caused some inconsistencies, such as pending certificate signing requests. With this release, the Power VS cloud provider is updated to fetch the machine IP address from the DHCP server so that the IP addresses for the nodes are consistent with the machine IP address. ([BZ#2111474](#))
- Previously, machines created in early versions of OpenShift Container Platform with invalid configurations could not be deleted. With this release, the webhooks that prevent the creation of machines with invalid configurations no longer prevent the deletion of existing invalid

machines. Users can now successfully remove these machines from their cluster by manually removing the finalizers on these machines. ([BZ#2101736](#))

- Previously, short DHCP lease times, caused by **NetworkManager** not being run as a daemon or in continuous mode, caused machines to become stuck during initial provisioning and never become nodes in the cluster. With this release, extra checks are added so that if a machine becomes stuck in this state it is deleted and recreated automatically. Machines that are affected by this network condition can become nodes after a reboot from the Machine API controller. ([BZ#2115090](#))
- Previously, when creating a new **Machine** resource using a machine profile that does not exist in IBM Cloud, the machines became stuck in the **Provisioning** phase. With this release, validation is added to the IBM Cloud Machine API provider to ensure that a machine profile exists, and machines with an invalid machine profile are rejected by the Machine API. ([BZ#2062579](#))
- Previously, the Machine API provider for AWS did not verify that the security group defined in the machine specification exists. Instead of returning an error in this case, it used a default security group, which should not be used for OpenShift Container Platform machines, and successfully created a machine without informing the user that the default group was used. With this release, the Machine API returns an error when users set either incorrect or empty security group names in the machine specification. ([BZ#2060068](#))
- Previously, the Machine API provider Azure did not treat user-provided values for instance types as case sensitive. This led to false-positive errors when instance types were correct but did not match the case. With this release, instance types are converted to the lowercase characters so that users get correct results without false-positive errors for mismatched case. ([BZ#2085390](#))
- Previously, there was no check for nil values in the annotations of a machine object before attempting to access the object. This situation was rare, but caused the machine controller to panic when reconciling the machine. With this release, nil values are checked and the machine controller is able to reconcile machines without annotations. ([BZ#2106733](#))
- Previously, the cluster autoscaler metrics for cluster CPU and memory usage would never reach, or exceed, the limits set by the **ClusterAutoscaler** resource. As a result, no alerts were fired when the cluster autoscaler could not scale due to resource limitations. With this release, a new metric called **cluster_autoscaler_skipped_scale_events_count** is added to the cluster autoscaler to more accurately detect when resource limits are reached or exceeded. Alerts will now fire when the cluster autoscaler is unable to scale the cluster up because it has reached the cluster resource limits. ([BZ#1997396](#))
- Previously, when the Machine API provider failed to fetch the machine IP address, it would not set the internal DNS name and the machine certificate signing requests were not automatically approved. With this release, the Power VS machine provider is updated to set the server name as the internal DNS name even when it fails to fetch the IP address. ([BZ#2111467](#))
- Previously, the Machine API vSphere machine controller set the **PowerOn** flag when cloning a VM. This created a **PowerOn** task that the machine controller was not aware of. If that **PowerOn** task failed, machines were stuck in the **Provisioned** phase but never powered on. With this release, the cloning sequence is altered to avoid the issue. Additionally, the machine controller now retries powering on the VM in case of failure and reports failures properly. ([BZ#2087981](#), [OCPBUGS-954](#))
- With this release, AWS security groups are tagged immediately instead of after creation. This means that fewer requests are sent to AWS and the required user privileges are lowered. ([BZ#2098054](#), [OCPBUGS-3094](#))

- Previously, a bug in the RHOSP legacy cloud provider resulted in a crash if certain RHOSP operations were attempted after authentication had failed. For example, shutting down a server causes the Kubernetes controller manager to fetch server information from RHOSP, which triggered this bug. As a result, if initial cloud authentication failed or was configured incorrectly, shutting down a server caused the Kubernetes controller manager to crash. With this release, the RHOSP legacy cloud provider is updated to not attempt any RHOSP API calls if it has not previously authenticated successfully. Now, shutting down a server with invalid cloud credentials no longer causes Kubernetes controller manager to crash. ([BZ#2102383](#))

Developer Console

- Previously, the **openshift-config** namespace was hard coded for the **HelmChartRepository** custom resource, which was the same namespace for the **ProjectHelmChartRepository** custom resource. This prevented users from adding private **ProjectHelmChartRepository** custom resources in their desired namespace. Consequently, users were unable to access secrets and configmaps in the **openshift-config** namespace. This update fixes the **ProjectHelmChartRepository** custom resource definition with a **namespace** field that can read the secret and configmaps from a namespace of choice by a user with the correct permissions. Additionally, the user can add secrets and configmaps to the accessible namespace, and they can add private Helm chart repositories in the namespace used the creation resources. ([BZ#2071792](#))

Image Registry

- Previously, the image trigger controller did not have permissions to change objects. Consequently, image trigger annotations did not work on some resources. This update creates a cluster role binding that provides the controller the required permissions to update objects according to annotations. ([BZ#2055620](#))
- Previously, the Image Registry Operator did not have a **progressing** condition for the **node-ca** daemon set and used **generation** from an incorrect object. Consequently, the **node-ca** daemon set could be marked as **degraded** while the Operator was still running. This update adds the **progressing** condition, which indicates that the installation is not complete. As a result, the Image Registry Operator successfully installs the **node-ca** daemon set, and the installer waits until it is fully deployed. ([BZ#2093440](#))

Installer

- Previously, the number of supported user-defined tags was 8, and reserved OpenShift Container Platform tags were 2 for AWS resources. With this release, the number of supported user-defined tags is now 25 and reserved OpenShift Container Platform tags are 25 for AWS resources. You can now add up to 25 user tags during installation. ([CFE#592](#))
- Previously, installing a cluster on Amazon Web Services started and then failed when the IAM administrative user was not assigned the **s3:GetBucketPolicy** permission. This update adds this policy to checklist that the installation program uses to ensure that all of the required permissions are assigned. As a result, the installation program now stops the installation with a warning that the IAM administrative user is missing the **s3:GetBucketPolicy** permission. ([BZ#2109388](#))
- Previously, installing a cluster on Microsoft Azure failed when the Azure DCasv5-series or DCadsv5-series of confidential VMs were specified as control plane nodes. With this update, the installation program now stops the installation with an error, which states that confidential VMs are not yet supported. ([BZ#2055247](#))

- Previously, gathering bootstrap logs was not possible until the control plane machines were running. With this update, gathering bootstrap logs now only requires that the bootstrap machine be available. ([BZ#2105341](#))
- Previously, if a cluster failed to install on Google Cloud Platform because the service account had insufficient permissions, the resulting error message did not mention this as the cause of the failure. This update improves the error message, which now instructs users to check the permissions that are assigned to the service account. ([BZ#2103236](#))
- Previously, when an installation on Google Cloud provider (GCP) failed because an invalid GCP region was specified, the resulting error message did not mention this as the cause of the failure. This update improves the error message, which now states the region is not valid. ([BZ#2102324](#))
- Previously, cluster installations using Hive could fail if Hive used an older version of the `install-config.yaml` file. This update allows the installation program to accept older versions of the **install-config.yaml** file provided by Hive. ([BZ#2098299](#))
- Previously, the installation program would incorrectly allow the **apiVIP** and **ingressVIP** parameters to use the same IPv6 address if they represented the address differently, such as listing the address in an abbreviated format. In this update, the installer correctly validates these two parameters regardless of their formatting, requiring separate IP addresses for each parameter. ([BZ#2103144](#))
- Previously, uninstalling a cluster using the installation program failed to delete all resources in clusters installed on GCP if the cluster name was more than 22 characters long. In this update, uninstalling a cluster using the installation program correctly locates and deletes all GCP cluster resources in cases of long cluster names. ([BZ#2076646](#))
- Previously, when installing a cluster on Red Hat OpenStack Platform (RHOSP) with multiple networks defined in the **machineNetwork** parameter, the installation program only created security group rules for the first network. With this update, the installation program creates security group rules for all networks defined in the **machineNetwork** so that users no longer need to manually edit security group rules after installation. ([BZ#2095323](#))
- Previously, users could manually set the API and Ingress virtual IP addresses to values that conflicted with the allocation pool of the DHCP server when installing a cluster on OpenStack. This could cause the DHCP server to assign one of the VIP addresses to a new machine, which would fail to start. In this update, the installation program validates the user-provided VIP addresses to ensure that they do not conflict with any DHCP pools. ([BZ#1944365](#))
- Previously, when installing a cluster on vSphere using a datacenter that is embedded inside a folder, the installation program could not locate the datacenter object, causing the installation to fail. In this update, the installation program can traverse the directory that contains the datacenter object, allowing the installation to succeed. ([BZ#2097691](#))
- Previously, when installing a cluster on Azure using arm64 architecture with installer-provisioned infrastructure, the image definition resource for **hyperVGeneration** V1 incorrectly had an architecture value of **x64**. With this update, the image definition resource for **hyperVGeneration** V1 has the correct architecture value of **Arm64**. ([OCBUGS-3639](#))
- Previously, when installing a cluster on VMware vSphere, the installation could fail if the user specified a user-defined folder in the **failureDomain** section of the **install-config.yaml** file. With this update, the installation program correctly validates user-defined folders in the **failureDomain** section of the **install-config.yaml** file. ([OCBUGS-3343](#))
- Previously, when destroying a partially deployed cluster after an installation failed on VMware

vSphere, some virtual machine folders were not destroyed. This error could occur in clusters configured with multiple vSphere datacenters or multiple vSphere clusters. With this update, all installer-provisioned infrastructure is correctly deleted when destroying a partially deployed cluster after an installation failure. ([OCPBUGS-1489](#))

- Previously, when installing a cluster on VMware vSphere, the installation failed if the user specified the **platform.vsphere.vcenters** parameter but did not specify the **platform.vsphere.failureDomains.topology.networks** parameter in the **install-config.yaml** file. With this update, the installation program alerts the user that the **platform.vsphere.failureDomains.topology.networks** field is required when specifying **platform.vsphere.vcenters**. ([OCPBUGS-1698](#))
- Previously, when installing a cluster on VMware vSphere, the installation failed if the user defined the **platform.vsphere.vcenters** and **platform.vsphere.failureDomains** parameters but did not define **platform.vsphere.defaultMachinePlatform.zones**, or **compute.platform.vsphere.zones** and **controlPlane.platform.vsphere.zones**. With this update, the installation program validates that the user has defined the **zones** parameter in multi-region or multi-zone deployments prior to installation. ([OCPBUGS-1490](#))

Kubernetes Controller Manager

- Previously, the Kubernetes Controller Manager Operator reported **degraded** on environments without a monitoring stack presence. With this update, the Kubernetes Controller Manager Operator skips checking the monitoring for cues about degradation when the monitoring stack is not present. ([BZ#2118286](#))
- With this update, Kubernetes Controller Manager alerts (**KubeControllerManagerDown**, **PodDisruptionBudgetAtLimit**, **PodDisruptionBudgetLimit**, and **GarbageCollectorSyncFailed**) have links to Github runbooks. The runbooks help users to understand debug these alerts. ([BZ#2001409](#))

Kubernetes Scheduler

- Previously, the secondary scheduler deployment was not deleted after a secondary scheduler custom resource was deleted. Consequently, the Secondary Schedule Operator and Operand were not fully uninstalled. With this update, the correct owner reference is set in the secondary scheduler custom resource so that it points to the secondary scheduler deployment. As a result, secondary scheduler deployments are deleted when the secondary scheduler custom resource is deleted. ([BZ#2100923](#))
- For the OpenShift Container Platform 4.12 release, the descheduler can now publish events to an API group because the release adds additional role-based access controls (RBAC) rules to the descheduler's profile. ([OCPBUGS-2330](#))

Machine Config Operator

- Previously, the Machine Config Operator (MCO) **ControllerConfig** resource, which contains important certificates, was only synced if the Operator's daemon sync succeeded. By design, unready nodes during a daemon sync prevent that daemon sync from succeeding, so unready nodes were indirectly preventing the **ControllerConfig** resource, and therefore those certificates, from syncing. This resulted in eventual cluster degradation when there were unready nodes due to inability to rotate the certificates contained in the **ControllerConfig** resource. With this release, the sync of the **ControllerConfig** resource is no longer dependent on the daemon sync succeeding, so the **ControllerConfig** resource now continues to sync if the

daemon sync fails. This means that unready nodes no longer prevent the **ControllerConfig** resource from syncing, so certificates continue to be updated even when there are unready nodes. ([BZ#2034883](#))

Management Console

- Previously, the **Operator details** page attempted to display multiple error messages, but the error message component can only display a single error message at a time. As a result, relevant error messages were not displayed. With this update, the **Operator details** page displays only the first error message so the user sees a relevant error. ([OCPBUGS-3927](#))
- Previously, the product name for Azure Red Hat OpenShift was incorrect in Customer Case Management (CCM). As a result, the console had to use the same incorrect product name to correctly populate the fields in CCM. Once the product name in CCM was updated, the console needed to be updated as well. With this update, the same, correct product name as CCM is correctly populated with the correct Azure product name when following the link from the console. ([OCPBUGS-869](#))
- Previously, when a plugin page resulted in an error, the error did not reset when navigating away from the error page, and the error persisted after navigating to a page that was not the cause of the error. With this update, the error state is reset to its default when a user navigates to a new page, and the error no longer persists after navigating to a new page. ([BZ#2117738](#), [OCPBUGS-523](#))
- Previously, the **View it here** link in the **Operator details** pane for installed Operators was incorrectly built when **All Namespaces** was selected. As a result, the link attempted to navigate to the **Operator details** page for a cluster service version (CSV) in **All Projects**, which is an invalid route. With this update, the **View it here** link to use the namespace where the CSV is installed now builds correctly and the link works as expected. ([OCPBUGS-184](#))
- Previously, line numbers with more than five digits resulted in a cosmetic issue where the line number overlaid the vertical divider between the line number and the line contents making it harder to read. With this update, the amount of space available for line numbers was increased to account for longer line numbers, and the line number no longer overlays the vertical divider. ([OCPBUGS-183](#))
- Previously, in the administrator perspective of the web console, the link to **Learn more about the OpenShift local update services** on the **Default update server** pop-up window in the **Cluster Settings** page produced a 404 error. With this update, the link works as expected. ([BZ#2098234](#))
- Previously, the **MatchExpression** component did not account for array-type values. As a result, only single values could be entered through forms using this component. With this update, the **MatchExpression** component accepts comma-separated values as an array. ([BZ#207690](#))
- Previously, there were redundant checks for the model resulting in tab reloading which occasionally resulted in a flickering of the tab contents where they rerendered. With this update, the redundant model check was removed, and the model is only checked once. As a result, the tab contents do not flicker and no longer rerender. ([BZ#2037329](#))
- Previously, when selecting the **edit** label from the action list on the OpenShift Dedicated node page, no response was elicited and a web hook error was returned. This issue has been fixed so that the error message is only returned when editing fails. ([BZ#2102098](#))
- Previously, if issues were pending, clicking on the **Insights** link would crash the page. As a workaround, you can wait for the variable to become **initialized** before clicking the **Insights** link. As a result, the Insights page will open as expected. ([BZ#2052662](#))

- Previously, when the **MachineConfigPool** resource was paused, the option to unpause said **Resume rollouts**. The wording has been updated so that it now says **Resume updates**. ([BZ#2094240](#))
- Previously, the wrong calculating method was used when counting master and worker nodes. With this update, the correct worker nodes are calculated when nodes have both the **master** and **worker** role. ([BZ#1951901](#))
- Previously, conflicting **react-router** routes for **ImageManifestVuln** resulted in attempts to render a details page for **ImageManifestVuln** with a **~new** name. Now, the container security plugin has been updated to remove conflicting routes and to ensure dynamic lists and details page extensions are used on the Operator details page. As a result, the console renders the correct create, list, and details pages for **ImageManifestVuln**. ([BZ#2080260](#))
- Previously, incomplete YAML was not synced was occasionally displayed to users. With this update, synced YAML always displays. ([BZ#2084453](#))
- Previously, when installing an Operator that required a custom resource (CR) to be created for use, the **Create resource** button could fail to install the CR because it was pointing to the incorrect namespace. With this update, the **Create resource** button works as expected. ([BZ#2094502](#))
- Previously, the **Cluster update** modal was not displaying errors properly. As a result, the **Cluster update** modal did not display or explain errors when they occurred. With this update, the **Cluster update** modal correctly display errors. ([BZ#2096350](#))

Monitoring

- Before this update, cluster administrators could not distinguish between a pod being not ready because of a scheduling issue and a pod being not ready because it could not be started by the kubelet. In both cases, the **KubePodNotReady** alert would fire. With this update, the **KubePodNotScheduled** alert now fires when a pod is not ready because of a scheduling issue, and the **KubePodNotReady** alert fires when a pod is not ready because it could not be started by the kubelet. ([OCPBUGS-4431](#))
- Before this update, **node_exporter** would report metrics about virtual network interfaces such as **tun** interfaces, **br** interfaces, and **ovn-k8s-mp** interfaces. With this update, metrics for these virtual interfaces are no longer collected, which decreases monitoring resource consumption. ([OCPBUGS-1321](#))
- Before this update, Alertmanager pod startup might time out because of slow DNS resolution, and the Alertmanager pods would not start. With this release, the timeout value has been increased to seven minutes, which prevents pod startup from timing out. ([BZ#2083226](#))
- Before this update, if Prometheus Operator failed to run or schedule Prometheus pods, the system provided no underlying reason for the failure. With this update, if Prometheus pods are not run or scheduled, the Cluster Monitoring Operator updates the **clusterOperator** monitoring status with a reason for the failure, which can be used to troubleshoot the underlying issue. ([BZ#2043518](#))
- Before this update, if you created an alert silence from the **Developer** perspective in the OpenShift Container Platform web console, external labels were included that did not match the alert. Therefore, the alert would not be silenced. With this update, external labels are now excluded when you create a silence in the **Developer** perspective so that newly created silences function as expected. ([BZ#2084504](#))

- Previously, if you enabled an instance of Alertmanager dedicated to user-defined projects, a misconfiguration could occur in certain circumstances, and you would not be informed that the user-defined project Alertmanager config map settings did not load for either the main instance of Alertmanager or the instance dedicated to user-defined projects. With this release, if this misconfiguration occurs, the Cluster Monitoring Operator now displays a message that informs you of the issue and provides resolution steps. ([BZ#2099939](#))
- Before this update, if the Cluster Monitoring Operator (CMO) failed to update Prometheus, the CMO did not verify whether a previous deployment was running and would report that cluster monitoring was unavailable even if one of the Prometheus pods was still running. With this update, the CMO now checks for running Prometheus pods in this situation and reports that cluster monitoring is unavailable only if no Prometheus pods are running. ([BZ#2039411](#))
- Before this update, if you configured OpsGenie as an alert receiver, a warning would appear in the log that **api_key** and **api_key_file** are mutually exclusive and that **api_key** takes precedence. This warning appeared even if you had not defined **api_key_file**. With this update, this warning only appears in the log if you have defined both **api_key** and **api_key_file**. ([BZ#2093892](#))
- Before this update the Telemeter Client (TC) only loaded new pull secrets when it was manually restarted. Therefore, if a pull secret had been changed or updated and the TC had not been restarted, the TC would fail to authenticate with the server. This update addresses the issue so that when the secret is rotated, the deployment is automatically restarted and uses the updated token to authenticate. ([BZ#2114721](#))

Networking

- Previously, routers that were in the terminating state delayed the **oc cp** command which would delay the **oc adm must-gather** command until the pod was terminated. With this update, a timeout for each issued **oc cp** command is set to prevent delaying the **must-gather** command from running. As a result, terminating pods no longer delay **must-gather** commands. ([BZ#2103283](#))
- Previously, an Ingress Controller could not be configured with both the **Private** endpoint publishing strategy type and PROXY protocol. With this update, users can now configure an Ingress Controller with both the **Private** endpoint publishing strategy type and PROXY protocol. ([BZ#2104481](#))
- Previously, the **routeSelector** parameter cleared the route status of the Ingress Controller prior to the router deployment. Because of this, the route status repopulated incorrectly. To avoid using stale data, route status detection has been updated to no longer rely on the Kubernetes object cache. Additionally, this update includes a fix to check the generation ID on route deployment to determine the route status. As a result, the route status is consistently cleared with a **routeSelector** update. ([BZ#2101878](#))
- Previously, a cluster that was upgraded from a version of OpenShift Container Platform earlier than 4.8 could have orphaned **Route** objects. This was caused by earlier versions of OpenShift Container Platform translating **Ingress** objects into **Route** objects irrespective of a given **Ingress** object's indicated **IngressClass**. With this update, an alert is sent to the cluster administrator about any orphaned Route objects still present in the cluster after Ingress-to-Route translation. This update also adds another alert that notifies the cluster administrator about any Ingress objects that do not specify an **IngressClass**. ([BZ#1962502](#))
- Previously, if a **configmap** that the router deployment depends on is not created, then the router deployment does not progress. With this update, the cluster Operator reports **ingress progressing=true** if the default ingress controller deployment is progressing. This results in

users debugging issues with the ingress controller by using the command **oc get co**. ([BZ#2066560](#))

- Previously, when an incorrectly created network policy was added to the OVN-Kubernetes cache, it would cause the OVN-Kubernetes leader to enter **crashloopbackoff** status. With this update, OVN-Kubernetes leader does not enter **crashloopbackoff** status by skipping deleting nil policies. ([BZ#2091238](#))
- Previously, recreating an EgressIP pod with the same namespace or name within 60 seconds of deleting an older one with the same namespace or name causes the wrong SNAT to be configured. As a result, packets could go out with nodeIP instead of EgressIP SNAT. With this update, traffic leaves the pod with EgressIP instead of nodeIP. ([BZ#2097243](#)).
- Previously, older Access Control Lists (ACL)s with **arp** produced **unexpectedly found multiple equivalent ACLs (arp v/s arp||nd)** errors due to a change in the ACL from **arp** to **arp || nd**. This prevented network policies from being created properly. With this update, older ACLs with just the **arp** match have been removed so that only ACLs with the new **arp || nd** match exist so that network policies can be created correctly and no errors will be observed on **ovnkube-master**. NOTE: This effects customers upgrading into 4.8.14, 4.9.32, 4.10.13 or higher from older versions. ([BZ#2095852](#)).
- With this update, CoreDNS has been updated to version 1.10.0, which is based on Kubernetes 1.25. This keeps both the CoreDNS version and OpenShift Container Platform 4.12, which is also based on Kubernetes 1.25, in alignment with one another. ([OCPBUGS-1731](#))
- With this update, the OpenShift Container Platform router now uses **k8s.io/client-go** version 1.25.2, which supports Kubernetes 1.25. This keeps both the **openshift-router** and OpenShift Container Platform 4.12, which is also based on Kubernetes 1.25, in alignment with one another. ([OCPBUGS-1730](#))
- With this update, the Ingress Operator now uses **k8s.io/client-go** version 1.25.2, which supports Kubernetes 1.25. This keeps both the Ingress Operator and OpenShift Container Platform 4.12, which is also based on Kubernetes 1.25, in alignment with one another. ([OCPBUGS-1554](#))
- Previously, the DNS Operator did not reconcile the **openshift-dns** namespace. Because OpenShift Container Platform 4.12 requires the **openshift-dns** namespace to have pod-security labels, this caused the namespace to be missing those labels upon cluster update. Without the pod-security labels, the pods failed to start. With this update, the DNS Operator now reconciles the **openshift-dns** namespace, and the pod-security labels are now present. As a result, pods start as expected. ([OCPBUGS-1549](#))
- Previously, the **ingresscontroller.spec.tuningOptions.reloadInterval** did not support decimal numerals as valid parameter values because the Ingress Operator internally converts the specified value into milliseconds, which was not a supported time unit. This prevented an Ingress Controller from being deleted. With this update, **ingresscontroller.spec.tuningOptions.reloadInterval** now supports decimal numerals and users can delete Ingress Controllers with **reloadInterval** parameter values which were previously unsupported. ([OCPBUGS-236](#))
- Previously, the Cluster DNS Operator used GO Kubernetes libraries that were based on Kubernetes 1.24 while OpenShift Container Platform 4.12 is based on Kubernetes 1.25. With this update, GO Kubernetes API is v1.25.2, which aligns the Cluster DNS Operator with OpenShift Container Platform 4.12 that uses Kubernetes 1.25 APIs. (link: [OCPBUGS-1558](#))
- Previously, setting the **disableNetworkDiagnostics** configuration to **true** did not persist when the **network-operator** pod was re-created. With this update, the **disableNetworkDiagnostics**

configuration property of network **operator.openshift.io/cluster** no longer resets to its default value after network operator restart. ([OCBUGS-392](#))

- Previously, **ovn-kubernetes** did not configure the correct MAC address of bonded interfaces in **br-ex** bridge. As a result, a node that uses bonding for the primary Kubernetes interface fails to join the cluster. With this update, **ovn-kubernetes** configures the correct MAC address of bonded interfaces in **br-ex** bridge, and nodes that use bonding for the primary Kubernetes interface successfully join the cluster. ([BZ2096413](#))
- Previously, when the Ingress Operator was configured to enable the use of mTLS, the Operator would not check if CRLs needed updating until some other event caused it to reconcile. As a result, CRLs used for mTLS could become out of date. With this update, the Ingress Operator now automatically reconciles when any CRL expires, and CRLs will be updated at the time specified by their **nextUpdate** field. ([BZ#2117524](#))

Node

- Previously, a symlinks error message was printed out as raw data instead of formatted as an error, making it difficult to understand. This fix formats the error message properly, so that it is easily understood. ([BZ#1977660](#))
- Previously, kubelet hard eviction thresholds were different from Kubernetes defaults when a performance profile was applied to a node. With this release, the defaults have been updated to match the expected Kubernetes defaults. ([OCBUGS-4362](#)).

OpenShift CLI (oc)

- The OpenShift Container Platform 4.12 release fixes an issue with entering a debug session on a target node when the target namespace lacks the appropriate security level. This caused the **oc** CLI to prompt you with a pod security error message. If the existing namespace does not contain the appropriate security levels, OpenShift Container Platform now creates a temporary namespace when you enter **oc** debug mode on a target node. ([OCBUGS-852](#))
- Previously, on macOS arm64 architecture, the **oc** binary needed to be signed manually. As a result, the **oc** binary did not work as expected. This update implements a self-signing binary for **oc** mimicking. As a result, the **oc** binary on macOS arm64 architectures works properly. ([BZ#2059125](#))
- Previously, **must-gather** was trying to collect resources that were not present on the server. Consequently, **must-gather** would print error messages. Now, before collecting resources, **must-gather** checks whether the resource exists. As a result, **must-gather** no longer prints an error when it fails to collect non-existing resources on the server. ([BZ#2095708](#))
- The OpenShift Container Platform 4.12 release updates the **oc-mirror** library, so that the library supports multi-arch platform images. This means that you can choose from a wider selection of architectures, such as **arm64**, when mirroring a platform release payload. ([OCBUGS-617](#))

Operator Lifecycle Manager (OLM)

- Before the OpenShift Container Platform 4.12 release, the **package-server-manager** controller would not revert any changes made to a **package-server** cluster service version (CSV), because of an issue with the **on-cluster** function. These persistent changes might impact how an Operator starts in a cluster. For OpenShift Container Platform 4.12, the **package-server-manager** controller always rebuilds a **package-server** CSV to its original state, so that no modifications to the CSV persist after a cluster upgrade operation. The **on-cluster** function no longer controls the state of a **package-server** CSV. ([OCBUGS-867](#))

- Previously, Operator Lifecycle Manager (OLM) would attempt to update namespaces to apply a label, even if the label was present on the namespace. Consequently, the update requests increased the workload in API and etcd services. With this update, OLM compares existing labels against the expected labels on a namespace before issuing an update. As a result, OLM no longer attempts to make unnecessary update requests on namespaces. ([BZ#2105045](#))
- Previously, Operator Lifecycle Manager (OLM) would prevent minor cluster upgrades that should not be blocked based on a miscalculation of the **ClusterVersion** custom resource's **spec.DesiredVersion** field. With this update, OLM no longer prevents cluster upgrades when the upgrade should be supported. ([BZ#2097557](#))
- Previously, the reconciler would update a resource's annotation without making a copy of the resource. This caused an error that would terminate the reconciler process. With this update, the reconciler no longer stops due the error. ([BZ#2105045](#))
- The **package-server-manifest** (PSM) is a controller that ensures that the correct **package-server** Cluster Service Version (CSV) is installed on a cluster. Previously, changes to the **package-server** CSV were not being reverted because of a logical error in the reconcile function in which an on-cluster object could influence the expected object. Users could modify the **package-server** CSV and the changes would not be reverted. Additionally, cluster upgrades would not update the YAML for the **package-server** CSV. With this update, the expected version of the CSV is now always built from scratch, which removes the ability for an on-cluster object to influence the expected values. As a result, the PSM now reverts any attempts to modify the **package-server** CSV, and cluster upgrades now deploy the expected **package-server** CSV. ([OCBUGS-858](#))
- Previously, OLM would upgrade an Operator according to the Operator's CRD status. A CRD lists component references in an order defined by the group/version/kind (GVK) identifier. Operators that share the same components might cause the GVK to change the component listings for an Operator, and this can cause the OLM to require more system resources to continuously update the status of a CRD. With this update, the Operator Lifecycle Manager (OLM) now upgrades an Operator according to the Operator's component references. A change to the custom resource definition (CRD) status of an Operator does not impact the OLM Operator upgrade process. ([OCBUGS-3795](#))

Operator SDK

- With this update, you can now set the security context for the registry pod by including the **securityContext** configuration field in the pod specification. This will apply the security context for all containers in the pod. The **securityContext** field also defines the pod's privileges. ([BZ#2091864](#))

File Integrity Operator

- Previously, the File Integrity Operator deployed templates using the **openshift-file-integrity** namespace in the permissions for the Operator. When the Operator attempted to create objects in the namespace, it would fail due to permission issues. With this release, the deployment resources used by OLM are updated to use the correct namespace, fixing the permission issues so that users can install and use the operator in non-default namespaces. ([BZ#2104897](#))
- Previously, underlying dependencies of the File Integrity Operator changed how alerts and notifications were handled, and the Operator didn't send metrics as a result. With this release the Operator ensures that the metrics endpoint is correct and reachable on startup. ([BZ#2115821](#))

- Previously, alerts issued by the File Integrity Operator did not set a namespace. This made it difficult to understand where the alert was coming from, or what component was responsible for issuing it. With this release, the Operator includes the namespace it was installed into in the alert, making it easier to narrow down what component needs attention. ([BZ#2101393](#))
- Previously, the File Integrity Operator did not properly handle modifying alerts during an upgrade. As a result, alerts did not include the namespace in which the Operator was installed. With this release, the Operator includes the namespace it was installed into in the alert, making it easier to narrow down what component needs attention. ([BZ#2112394](#))
- Previously, service account ownership for the File Integrity Operator regressed due to underlying OLM updates, and updates from 0.1.24 to 0.1.29 were broken. With this update, the Operator defaults to upgrading to 0.1.30. ([BZ#2109153](#))
- Previously, the File Integrity Operator daemon used the **ClusterRoles** parameter instead of the **Roles** parameter for a recent permission change. As a result, OLM could not update the Operator. With this release, the Operator daemon reverts to using the **Roles** parameter and updates from older versions to version 0.1.29 are successful. ([BZ#2108475](#))

Compliance Operator

- Previously, the Compliance Operator used an old version of the Operator SDK, which is a dependency for building Operators. This caused alerts about deprecated Kubernetes functionality used by the Operator SDK. With this release, the Compliance Operator is updated to version 0.1.55, which includes an updated version of the Operator SDK. ([BZ#2098581](#))
- Previously, applying automatic remediation for the **rhcos4-high-master-sysctl-kernel-yama-ptrace-scope** and **rhcos4-sysctl-kernel-core-pattern** rules resulted in subsequent failures of those rules in scan results, even though they were remediated. The issue is fixed in this release. ([BZ#2094382](#))
- Previously, the Compliance Operator hard coded notifications to the default namespace. As a result, notifications from the Operator would not appear if the Operator was installed in a different namespace. This issue is fixed in this release. ([BZ#2060726](#))
- Previously, the Compliance Operator failed to fetch API resources when parsing machine configurations without Ignition specifications. This caused the **api-check-pods** check to crash loop. With this release, the Compliance Operator is updated to gracefully handle machine configuration pools without Ignition specifications. ([BZ#2117268](#))
- Previously, the Compliance Operator held machine configurations in a stuck state because it could not determine the relationship between machine configurations and kubelet configurations. This was due to incorrect assumptions about machine configuration names. With this release, the Compliance Operator is able to determine if a kubelet configuration is a subset of a machine configuration. ([BZ#2102511](#))

OpenShift API server

- Previously, adding a member could remove previous members from a group. As a result, the user lost group privileges. With this release, the dependencies were bumped and users no longer lose group privileges. ([OCPBUGS-533](#))

Red Hat Enterprise Linux CoreOS (RHCOS)

- Previously, updating to Podman 4.0 prevented users from using custom images with toolbox containers on RHCOS. This fix updates the toolbox library code to account for the new Podman behavior, so users can now use custom images with toolbox on RHCOS as expected.

([BZ#2048789](#))

- Previously, the **podman exec** command did not work well with nested containers. Users encountered this issue when accessing a node using the **oc debug** command and then running a container with the **toolbox** command. Because of this, users were unable to reuse toolboxes on RHCOS. This fix updates the toolbox library code to account for this behavior, so users can now reuse toolboxes on RHCOS. ([BZ#1915537](#))
- With this update, running the **toolbox** command now checks for updates to the default image before launching the container. This improves security and provides users with the latest bug fixes. ([BZ#2049591](#))
- Previously, updating to Podman 4.0 prevented users from running the **toolbox** command on RHCOS. This fix updates the toolbox library code to account for the new Podman behavior, so users can now run **toolbox** on RHCOS as expected. ([BZ#2093040](#))
- Previously, custom SELinux policy modules were not properly supported by **rpm-ostree**, so they were not updated along with the rest of the system upon update. This would surface as failures in unrelated components. Pending SELinux userspace improvements landing in a future OpenShift Container Platform release, this update provides a workaround to RHCOS that will rebuild and reload the SELinux policy during boot as needed. ([OCPBUGS-595](#))

Scalability and performance

- The tuned profile has been modified to assign the same priority as **ksoftirqd** and **rcuc** to the newly introduced per-CPU kthreads (**ktimers**) added in a recent Red Hat Enterprise Linux (RHEL) kernel patch. For more information, see [OCPBUGS-3475](#), [BZ#2117780](#), and [BZ#2122220](#).
- Previously, restarts of the **tuned** service caused improper reset of the **irqbalance** configuration, leading to IRQ operation being served again on the isolated CPUs, therefore violating the isolation guarantees. With this fix, the **irqbalance** service configuration is properly preserved across **tuned** service restarts (explicit or caused by bugs), therefore preserving the CPU isolation guarantees with respect to IRQ serving. ([OCPBUGS-585](#))
- Previously, when the tuned daemon was restarted out of order as part of the cluster Node Tuning Operator, the CPU affinity of interrupt handlers was reset and the tuning was compromised. With this fix, the **irqbalance** plugin in tuned is disabled, and OpenShift Container Platform now relies on the logic and interaction between **CRI-O** and **irqbalance**.([BZ#2105123](#))
- Previously, a low latency hook script executing for every new **veth** device took too long when the node was under load. The resultant accumulated delays during pod start events caused the rollout time for **kube-apiserver** to be slow and sometimes exceed the 5-minute rollout timeout. With this fix, the container start time should be shorter and within the 5-minute threshold. ([BZ#2109965](#)).
- Previously, the **oslat** control thread was colocated with one of the test threads, which caused latency spikes in the measurements. With this fix, the **oslat** runner now reserves one CPU for the control thread, meaning the test uses one less CPU for running the busy threads. ([BZ#2051443](#))
- Latency measurement tools, also known as **oslat**, **cyclictest**, and **hwlatdetect**, now run on completely isolated CPUs without the helper process running in the background that might cause latency spikes, therefore providing more accurate latency measurements. ([OCPBUGS-2618](#))

- Previously, although the reference **PolicyGenTemplate** for **group-du-sno-ranGen.yaml** includes two **StorageClass** entries, the generated policy included only one. With this update, the generated policy now includes both policies. ([BZ#2049306](#)).

Storage

- Previously, checks for generic ephemeral volumes failed. With this update, checks for expandable volumes now include generic ephemeral volumes. ([BZ#2082773](#))
- Previously, if more than one secret was present for vSphere, the vSphere CSI Operator randomly picked a secret and sometimes caused the Operator to restart. With this update, a warning appears when there is more than one secret on the vCenter CSI Operator. ([BZ#2108473](#))
- Previously, OpenShift Container Platform detached a volume when a Container Storage Interface (CSI) driver was not able to unmount the volume from a node. Detaching a volume without unmount is not allowed by CSI specifications and drivers could enter an **undocumented** state. With this update, CSI drivers are detached before unmounting only on unhealthy nodes preventing the **undocumented** state. ([BZ#2049306](#))
- Previously, there were missing annotations on the Manila CSI Driver Operator's VolumeSnapshotClass. Consequently, the Manila CSI snapshotter could not locate secrets, and could not create snapshots with the default VolumeSnapshotClass. This update fixes the issue so that secret names and namespaces are included in the default VolumeSnapshotClass. As a result, users can now create snapshots in the Manila CSI Driver Operator using the default VolumeSnapshotClass. ([BZ#2057637](#))
- Users can now opt into using the experimental VHD feature on Azure File. To opt in, users must specify the **fstype** parameter in a storage class and enable it with **--enable-vhd=true**. If **fstype** is used and the feature is not set to **true**, the volumes will fail to provision. To opt out of using the VHD feature, remove the **fstype** parameter from your storage class. ([BZ#2080449](#))
- Previously, if more than one secret was present for vSphere, the vSphere CSI Operator randomly picked a secret and sometimes caused the Operator to restart. With this update, a warning appears when there is more than one secret on the vCenter CSI Operator. ([BZ#2108473](#))

Web console (Developer perspective)

- Previously, the users could not deselect a Git secret in add and edit forms. As a result, the resources had to be recreated. This fix resolves the issue by adding the option to choose **No Secret** in the select secret option list. As a result, the users can easily select, deselect, or detach any attached secrets. ([BZ#2089221](#))
- In OpenShift Container Platform 4.9, when it is minimal or no data in the **Developer Perspective**, most of the monitoring charts or graphs (CPU consumption, memory usage, and bandwidth) show a range of -1 to 1. However, none of these values can ever go below zero. This will be resolved in a future release. ([BZ#1904106](#))
- Before this update, users could not silence alerts in the **Developer** perspective in the OpenShift Container Platform web console when a user-defined Alertmanager service was deployed because the web console would forward the request to the platform Alertmanager service in the **openshift-monitoring** namespace. With this update, when you view the **Developer** perspective in the web console and try to silence an alert, the request is forwarded to the correct Alertmanager service. ([OCPBUGS-1789](#))

- Previously, there was a known issue in the **Add Helm Chart Repositories** form to extend the Developer Catalog of a project. The **Quick Start** guides shows that you can add the **ProjectHelmChartRepository** CR in the required namespace whereas it does not mention that to perform this you need permission from the kubeadmin. This issue was resolved with **Quickstart** mentioning the correct steps to create **ProjectHelmChartRepository** CR. ([BZ#2057306](#))

1.7. TECHNOLOGY PREVIEW FEATURES

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

Technology Preview Features Support Scope

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

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

Networking Technology Preview features

Table 1.14. Networking Technology Preview tracker

Feature	4.10	4.11	4.12
PTP single NIC hardware configured as boundary clock	Technology Preview	General Availability	General Availability
PTP dual NIC hardware configured as boundary clock	Not Available	Technology Preview	Technology Preview
PTP events with boundary clock	Technology Preview	General Availability	General Availability
HTTP transport replaces AMQP for PTP and bare-metal events	Not Available	Not Available	General Availability
Pod-level bonding for secondary networks	General Availability	General Availability	General Availability
External DNS Operator	Technology Preview	General Availability	General Availability
AWS Load Balancer Operator	Not Available	Technology Preview	General Availability

Feature	4.10	4.11	4.12
Ingress Node Firewall Operator	Not Available	Not Available	Technology Preview
Advertise using BGP mode the MetalLB service from a subset of nodes, using a specific pool of IP addresses	Not Available	Technology Preview	General Availability
Advertise using L2 mode the MetalLB service from a subset of nodes, using a specific pool of IP addresses	Not Available	Technology Preview	Technology Preview
Multi-network policies for SR-IOV networks	Not Available	Not Available	Technology Preview
Updating the interface-specific safe sysctls list	Not Available	Not Available	Technology Preview
MT2892 Family [ConnectX-6 Dx] SR-IOV support	Not Available	Not Available	Technology Preview
MT2894 Family [ConnectX-6 Lx] SR-IOV support	Not Available	Not Available	Technology Preview
MT42822 BlueField-2 in ConnectX-6 NIC mode SR-IOV support	Not Available	Not Available	Technology Preview
Silicom STS Family SR-IOV support	Not Available	Not Available	Technology Preview
MT2892 Family [ConnectX-6 Dx] OvS Hardware Offload support	Not Available	Not Available	Technology Preview
MT2894 Family [ConnectX-6 Lx] OvS Hardware Offload support	Not Available	Not Available	Technology Preview
MT42822 BlueField-2 in ConnectX-6 NIC mode OvS Hardware Offload support	Not Available	Not Available	Technology Preview
Switching Bluefield-2 from DPU to NIC	Not available	Not available	Technology Preview

Storage Technology Preview features

Table 1.15. Storage Technology Preview tracker

Feature	4.10	4.11	4.12
Shared Resources CSI Driver and Build CSI Volumes in OpenShift Builds	Technology Preview	Technology Preview	Technology Preview

Feature	4.10	4.11	4.12
CSI volume expansion	Technology Preview	General Availability	General Availability
CSI Azure File Driver Operator	Technology Preview	General Availability	General Availability
CSI Google Filestore Driver Operator	Not Available	Not Available	Technology Preview
CSI automatic migration (Azure file, VMware vSphere)	Technology Preview	Technology Preview	Technology Preview
CSI automatic migration (Azure Disk, OpenStack Cinder)	Technology Preview	General Availability	General Availability
CSI automatic migration (AWS EBS, GCP disk)	Technology Preview	Technology Preview	General Availability
CSI inline ephemeral volumes	Technology Preview	Technology Preview	Technology Preview
CSI generic ephemeral volumes	Not Available	General Availability	General Availability
Shared Resource CSI Driver	Technology Preview	Technology Preview	Technology Preview
CSI Google Filestore Driver Operator	Not Available	Not Available	Technology Preview
Automatic device discovery and provisioning with Local Storage Operator	Technology Preview	Technology Preview	Technology Preview
NFS support for Azure File CSI Operator Driver	Not Available	Not Available	Generally Available

Installation Technology Preview features

Table 1.16. Installation Technology Preview tracker

Feature	4.10	4.11	4.12
Adding kernel modules to nodes with kvc	Technology Preview	Technology Preview	Technology Preview

Feature	4.10	4.11	4.12
IBM Cloud VPC clusters	Technology Preview	Technology Preview	General Availability
Selectable Cluster Inventory	Technology Preview	Technology Preview	Technology Preview
Multi-architecture compute machines	Not Available	Technology Preview	Technology Preview
Disconnected mirroring with the oc-mirror CLI plugin	Technology Preview	General Availability	General Availability
Mount shared entitlements in BuildConfigs in RHEL	Technology Preview	Technology Preview	Technology Preview
Agent-based OpenShift Container Platform Installer	Not Available	Not Available	General Availability
AWS Outposts platform	Not Available	Not Available	Technology Preview
Installing a cluster on Alibaba Cloud using installer-provisioned infrastructure	Technology Preview	Technology Preview	Technology Preview

Node Technology Preview features

Table 1.17. Nodes Technology Preview tracker

Feature	4.10	4.11	4.12
Non-preempting priority classes	Technology Preview	General Availability	General Availability
Node Health Check Operator	Technology Preview	General Availability	General Availability
Linux Control Group version 2 (cgroup v2)	Not Available	Not Available	Technology Preview
crun container runtime	Not Available	Not Available	Technology Preview

Multi-Architecture Technology Preview features

Table 1.18. Multi-Architecture Technology Preview tracker

Feature	4.10	4.11	4.12
kdump on x86_64 architecture	Technology Preview	General Availability	General Availability
kdump on arm64 architecture	Not Available	Technology Preview	Technology Preview
kdump on s390x architecture	Technology Preview	Technology Preview	Technology Preview
kdump on ppc64le architecture	Technology Preview	Technology Preview	Technology Preview
IBM Secure Execution on IBM Z and LinuxONE	Not Available	Not Available	Technology Preview

Serverless Technology Preview features

Table 1.19. Serverless Technology Preview tracker

Feature	4.10	4.11	4.12
Serverless functions	Technology Preview	Technology Preview	Technology Preview

Specialized hardware and driver enablement Technology Preview features

Table 1.20. Specialized hardware and driver enablement Technology Preview tracker

Feature	4.10	4.11	4.12
Driver Toolkit	Technology Preview	Technology Preview	General Availability
Special Resource Operator (SRO)	Technology Preview	Technology Preview	Not Available
Hub and spoke cluster support	Not Available	Not Available	Technology Preview

Web console Technology Preview features

Table 1.21. Web console Technology Preview tracker

Feature	4.10	4.11	4.12
---------	------	------	------

Feature	4.10	4.11	4.12
Dynamic Plugins	Technology Preview	Technology Preview	General Availability

Scalability and performance Technology Preview features

Table 1.22. Scalability and performance Technology Preview tracker

Feature	4.10	4.11	4.12
Hyperthreading-aware CPU manager policy	Technology Preview	Technology Preview	Technology Preview
Node Observability Operator	Not Available	Technology Preview	Technology Preview
factory-precaching-cli tool	Not Available	Not Available	Technology Preview
Adding worker nodes to Single-node OpenShift clusters with GitOps ZTP	Not Available	Not Available	Technology Preview
Topology Aware Lifecycle Manager (TALM)	Technology Preview	Technology Preview	General Availability
Mount namespace encapsulation	Not Available	Not Available	Technology Preview
NUMA-aware scheduling with NUMA Resources Operator	Technology Preview	Technology Preview	General Availability

Operator Technology Preview features

Table 1.23. Operator Technology Preview tracker

Feature	4.10	4.11	4.12
Hybrid Helm Operator	Technology Preview	Technology Preview	Technology Preview
Java-based Operator	Not Available	Technology Preview	Technology Preview
Node Observability Operator	Not Available	Not Available	Technology Preview

Feature	4.10	4.11	4.12
Network Observability Operator	Supported	Supported	General Availability
Platform Operators	Not Available	Not Available	Technology Preview
RukPak	Not Available	Not Available	Technology Preview
cert-manager Operator	Technology Preview	Technology Preview	General Availability

Monitoring Technology Preview features

Table 1.24. Monitoring Technology Preview tracker

Feature	4.10	4.11	4.12
Alert routing for user-defined projects monitoring	Technology Preview	General Availability	General Availability
Alerting rules based on platform monitoring metrics	Not Available	Technology Preview	Technology Preview

Red Hat OpenStack Platform (RHOSP) Technology Preview features

Table 1.25. RHOSP Technology Preview tracker

Feature	4.10	4.11	4.12
Support for RHOSP DCN	Technology Preview	Technology Preview	Technology Preview
Support for external cloud providers for clusters on RHOSP	Technology Preview	Technology Preview	General Availability
OVS hardware offloading for clusters on RHOSP	Technology Preview	General Availability	General Availability

Architecture Technology Preview features

Table 1.26. Architecture Technology Preview tracker

Feature	4.10	4.11	4.12
Hosted control planes for OpenShift Container Platform on bare metal	Not Available	Not Available	Technology Preview
Hosted control planes for OpenShift Container Platform on Amazon Web Services (AWS)	Not Available	Technology Preview	Technology Preview

Machine management Technology Preview features

Table 1.27. Machine management Technology Preview tracker

Feature	4.10	4.11	4.12
Managing machines with the Cluster API for Amazon Web Services	Not Available	Technology Preview	Technology Preview
Managing machines with the Cluster API for Google Cloud Platform	Not Available	Technology Preview	Technology Preview
Cron job time zones	Not Available	Not Available	Technology Preview
Cloud controller manager for Alibaba Cloud	Technology Preview	Technology Preview	Technology Preview
Cloud controller manager for Amazon Web Services	Technology Preview	Technology Preview	Technology Preview
Cloud controller manager for Google Cloud Platform	Technology Preview	Technology Preview	Technology Preview
Cloud controller manager for IBM Cloud	Technology Preview	Technology Preview	General Availability
Cloud controller manager for Microsoft Azure	Technology Preview	Technology Preview	Technology Preview
Cloud controller manager for Red Hat OpenStack Platform (RHOSP)	Technology Preview	Technology Preview	General Availability
Cloud controller manager for VMware vSphere	Technology Preview	Technology Preview	Technology Preview
Custom Metrics Autoscaler Operator	Not Available	Technology Preview	Technology Preview

Authentication and authorization Technology Preview features

Table 1.28. Authentication and authorization Technology Preview tracker

Table 1.28. Authentication and authorization Technology Preview tracker

Feature	4.10	4.11	4.12
Pod security admission restricted enforcement	Not Available	Not Available	Technology Preview

1.8. KNOWN ISSUES

- In OpenShift Container Platform 4.1, anonymous users could access discovery endpoints. Later releases revoked this access to reduce the possible attack surface for security exploits because some discovery endpoints are forwarded to aggregated API servers. However, unauthenticated access is preserved in upgraded clusters so that existing use cases are not broken. If you are a cluster administrator for a cluster that has been upgraded from OpenShift Container Platform 4.1 to 4.12, you can either revoke or continue to allow unauthenticated access. Unless there is a specific need for unauthenticated access, you should revoke it. If you do continue to allow unauthenticated access, be aware of the increased risks.



WARNING

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

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

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

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

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

(BZ#1821771)

- Intermittently, an IBM Cloud VPC cluster might fail to install because some worker machines do not start. Rather, these worker machines remain in the **Provisioned** phase. There is a workaround for this issue. From the host where you performed the initial installation, delete the failed machines and run the installation program again.

- Verify that the status of the internal application load balancer (ALB) for the master API server is **active**.

- Identify the cluster's infrastructure ID by running the following command:

```
$ oc get infrastructure/cluster -ojson | jq -r '.status.infrastructureName'
```

- Log into the IBM Cloud account for your cluster and target the correct region for your cluster.

- Verify that the internal ALB status is **active** by running the following command:

```
$ ibmcloud is lb <cluster_ID>-kubernetes-api-private --output json | jq -r '.provisioning_status'
```

- Identify the machines that are in the **Provisioned** phase by running the following command:

```
$ oc get machine -n openshift-machine-api
```

Example output

```
NAME                                PHASE    TYPE    REGION  ZONE    AGE
example-public-1-x4gpn-master-0     Running  bx2-4x16  us-east  us-east-1  23h
example-public-1-x4gpn-master-1     Running  bx2-4x16  us-east  us-east-2  23h
example-public-1-x4gpn-master-2     Running  bx2-4x16  us-east  us-east-3  23h
example-public-1-x4gpn-worker-1-xqzzm Running  bx2-4x16  us-east  us-east-1  22h
example-public-1-x4gpn-worker-2-vg9w6 Provisioned bx2-4x16  us-east  us-east-2  22h
example-public-1-x4gpn-worker-3-2f7zd Provisioned bx2-4x16  us-east  us-east-3  22h
```

- Delete each failed machine by running the following command:

```
$ oc delete machine <name_of_machine> -n openshift-machine-api
```

- Wait for the deleted worker machines to be replaced, which can take up to 10 minutes.

- Verify that the new worker machines are in the **Running** phase by running the following command:

```
$ oc get machine -n openshift-machine-api
```

Example output

```
NAME                                PHASE    TYPE    REGION  ZONE    AGE
```

```

example-public-1-x4gpn-master-0    Running  bx2-4x16  us-east  us-east-1  23h
example-public-1-x4gpn-master-1    Running  bx2-4x16  us-east  us-east-2  23h
example-public-1-x4gpn-master-2    Running  bx2-4x16  us-east  us-east-3  23h
example-public-1-x4gpn-worker-1-xqzzm Running  bx2-4x16  us-east  us-east-1  23h
example-public-1-x4gpn-worker-2-mnlsz Running  bx2-4x16  us-east  us-east-2
8m2s
example-public-1-x4gpn-worker-3-7nz4q Running  bx2-4x16  us-east  us-east-3
7m24s

```

6. Complete the installation by running the following command. Running the installation program again ensures that the cluster's **kubeconfig** is initialized properly:

```
$ ./openshift-install wait-for install-complete
```

([OCPBUGS#1327](#))

- The **oc annotate** command does not work for LDAP group names that contain an equal sign (=), because the command uses the equal sign as a delimiter between the annotation name and value. As a workaround, use **oc patch** or **oc edit** to add the annotation. ([BZ#1917280](#))
- Due to the inclusion of old images in some image indexes, running **oc adm catalog mirror** and **oc image mirror** might result in the following error: **error: unable to retrieve source image**. As a temporary workaround, you can use the **--skip-missing** option to bypass the error and continue downloading the image index. For more information, see [Service Mesh Operator mirroring failed](#).
- When using the egress IP address feature in OpenShift Container Platform on RHOSP, you can assign a floating IP address to a reservation port to have a predictable SNAT address for egress traffic. The floating IP address association must be created by the same user that installed the OpenShift Container Platform cluster. Otherwise any delete or move operation for the egress IP address hangs indefinitely because of insufficient privileges. When this issue occurs, a user with sufficient privileges must manually unset the floating IP address association to resolve the issue. ([OCPBUGS-4902](#))
- There is a known issue with Nutanix installation where the installation fails if you use 4096-bit certificates with Prism Central 2022.x. Instead, use 2048-bit certificates. ([KCS](#))
- Deleting the bidirectional forwarding detection (BFD) profile and removing the **bfdProfile** added to the border gateway protocol (BGP) peer resource does not disable the BFD. Instead, the BGP peer starts using the default BFD profile. To disable BFD from a BGP peer resource, delete the BGP peer configuration and recreate it without a BFD profile. ([BZ#2050824](#))
- Due to an unresolved metadata API issue, you cannot install clusters that use bare-metal workers on RHOSP 16.1. Clusters on RHOSP 16.2 are not impacted by this issue. ([BZ#2033953](#))
- The **loadBalancerSourceRanges** attribute is not supported, and is therefore ignored, in load-balancer type services in clusters that run on RHOSP and use the OVN Octavia provider. There is no workaround for this issue. ([OCPBUGS-2789](#))
- After a catalog source update, it takes time for OLM to update the subscription status. This can mean that the status of the subscription policy may continue to show as compliant when Topology Aware Lifecycle Manager (TALM) decides whether remediation is needed. As a result the operator specified in the subscription policy does not get upgraded. As a workaround, include a **status** field in the **spec** section of the catalog source policy as follows:

```
metadata:
```

```

name: redhat-operators-disconnected
spec:
  displayName: disconnected-redhat-operators
  image: registry.example.com:5000/disconnected-redhat-operators/disconnected-redhat-operator-index:v4.11
status:
  connectionState:
    lastObservedState: READY

```

This mitigates the delay for OLM to pull the new index image and get the pod ready, reducing the time between completion of catalog source policy remediation and the update of the subscription status. If the issue persists and the subscription policy status update is still late you can apply another **ClusterGroupUpdate** CR with the same subscription policy, or an identical **ClusterGroupUpdate** CR with a different name. ([OCBUGS-2813](#))

- TALM skips remediating a policy if all selected clusters are compliant when the **ClusterGroupUpdate** CR is started. The update of operators with a modified catalog source policy and a subscription policy in the same **ClusterGroupUpdate** CR does not complete. The subscription policy is skipped as it is still compliant until the catalog source change is enforced. As a workaround, add the following change to one CR in the **common-subscription** policy, for example:

```

metadata.annotations.upgrade: "1"

```

This makes the policy non-compliant prior to the start of the **ClusterGroupUpdate** CR. ([OCBUGS-2812](#))

- On a single-node OpenShift instance, rebooting without draining the node to remove all the running pods can cause issues with workload container recovery. After the reboot, the workload restarts before all the device plugins are ready, resulting in resources not being available or the workload running on the wrong NUMA node. The workaround is to restart the workload pods when all the device plugins have re-registered themselves during the reboot recovery procedure. ([OCBUGS-2180](#))
- The default **dataset_comparison** is currently **ieee1588**. The recommended **dataset_comparison** is **G.8275.x**. It is planned to be fixed in a future version of OpenShift Container Platform. In the short term, you can manually update the ptp configuration to include the recommended **dataset_comparison**. ([OCBUGS-2336](#))
- The default **step_threshold** is 0.0. The recommended **step_threshold** is 2.0. It is planned to be fixed in a future version of OpenShift Container Platform. In the short term, you can manually update the ptp configuration to include the recommended **step_threshold**. ([OCBUGS-3005](#))
- The **BMCEventSubscription** CR fails to create a Redfish subscription for a spoke cluster in an ACM-deployed multi-cluster environment, where the metal3 service is only running on a hub cluster. The workaround is to create the subscription by calling the Redfish API directly, for example, by running the following command:

```

curl -X POST -i --insecure -u "<BMC_username>:<BMC_password>"
https://<BMC_IP>/redfish/v1/EventService/Subscriptions \
-H 'Content-Type: application/json' \
--data-raw '{
  "Protocol": "Redfish",
  "Context": "any string is valid",
  "Destination": "https://hw-event-proxy-openshift-bare-metal-

```

```
events.apps.example.com/webhook",  
  "EventTypes": ["Alert"]  
}'
```

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

- When using the GitOps ZTP pipeline to install a single-node OpenShift cluster in a disconnected environment, there should be two **CatalogSource** CRs applied in the cluster. One of the **CatalogSource** CRs gets deleted following multiple node reboots. As a workaround, you can change the default names, such as **certified-operators** and **redhat-operators**, of the catalog sources. ([OCPBUGSM-46245](#))
- If an invalid subscription channel is specified in the subscription policy that is used to perform a cluster upgrade, the Topology Aware Lifecycle Manager indicates a successful upgrade right after the policy is enforced because the **Subscription** state remains **AtLatestKnown**. ([OCPBUGSM-43618](#))
- The **SiteConfig** disk partition definition fails when applied to multiple nodes in a cluster. When a **SiteConfig** CR is used to provision a compact cluster, creating a valid **diskPartition** config on multiple nodes fails with a Kustomize plugin error. ([OCPBUGSM-44403](#))
- If secure boot is currently disabled and you try to enable it using ZTP, the cluster installation does not start. When secure boot is enabled through ZTP, the boot options are configured before the virtual CD is attached. Therefore, the first boot from the existing hard disk has the secure boot turned on. The cluster installation gets stuck because the system never boots from the CD. ([OCPBUGSM-45085](#))
- Using Red Hat Advanced Cluster Management (RHACM), spoke cluster deployments on Dell PowerEdge R640 servers are blocked when the virtual media does not disconnect the ISO in the iDRAC console after writing the image to the disk. As a workaround, disconnect the ISO manually through the Virtual Media tab in the iDRAC console. ([OCPBUGSM-45884](#))
- Low-latency applications that rely on high-resolution timers to wake up their threads might experience higher wake up latencies than expected. Although the expected wake up latency is under 20us, latencies exceeding this can occasionally be seen when running the `cyclictest` tool for long durations (24 hours or more). Testing has shown that wake up latencies are under 20us for over 99.999999% of the samples. ([RHELPLAN-138733](#))
- A Chapman Beach NIC from Intel must be installed in a bifurcated PCIe slot to ensure that both ports are visible. A limitation also exists in the current devlink tooling in RHEL 8.6 which prevents the configuration of 2 ports in the bifurcated PCIe slot. ([RHELPLAN-142458](#))
- Disabling an SR-IOV VF when a port goes down can cause a 3-4 second delay with Intel NICs. ([RHELPLAN-126931](#))
- When using Intel NICs, IPV6 traffic stops when an SR-IOV VF is assigned an IPV6 address. ([RHELPLAN-137741](#))
- When using VLAN strip offloading, the offload flag (**ol_flag**) is not consistently set correctly with the `iavf` driver. ([RHELPLAN-141240](#))
- A deadlock can occur if an allocation fails during a configuration change with the `ice` driver. ([RHELPLAN-130855](#))

- SR-IOV VFs send GARP packets with the wrong MAC address when using Intel NICs. ([RHELPLAN-140971](#))
- When using the GitOps ZTP method of managing clusters and deleting a cluster which has not completed installation, the cleanup of the cluster namespace on the hub cluster might hang indefinitely. To complete the namespace deletion, remove the **baremetalhost.metal3.io** finalizer from two CRs in the cluster namespace:
 1. Remove the finalizer from the secret that is pointed to by the BareMetalHost CR **.spec.bmc.credentialsName**.
 2. Remove the finalizer from the **BareMetalHost** CR. When these finalizers are removed the namespace termination completes within a few seconds. ([OCPBUGS-3029](#))
- The addition of a new feature in OCP 4.12 that enables UDP GRO also causes all veth devices to have one RX queue per available CPU (previously each veth had one queue). Those queues are dynamically configured by OVN and there is no synchronization between latency tuning and this queue creation. The latency tuning logic monitors the veth NIC creation events and starts configuring the RPS queue cpu masks before all the queues are properly created. This means that some of the RPS queue masks are not configured. Since not all NIC queues are configured properly there is a chance of latency spikes in a real-time application that uses timing-sensitive cpus for communicating with services in other containers. Applications that do not use kernel networking stack are not affected. ([OCPBUGS-4194](#))
- Platform Operator and RukPak known issues:
 - Deleting a platform Operator results in a cascading deletion of the underlying resources. This cascading deletion logic can only delete resources that are defined in the Operator Lifecycle Manager-based (OLM) Operator's bundle format. In the case that a platform Operator creates resources that are defined outside of that bundle format, then the platform Operator is responsible for handling this cleanup interaction. This behavior can be observed when installing the cert-manager Operator as a platform Operator, and then removing it. The expected behavior is that a namespace is left behind that the cert-manager Operator created.
 - The platform Operators manager does not have any logic that compares the current and desired state of the cluster-scoped **BundleDeployment** resource it is managing. This leaves the possibility for a user who has sufficient role-based access control (RBAC) to manually modify that underlying **BundleDeployment** resource and can lead to situations where users can escalate their permissions to the **cluster-admin** role. By default, you should limit access to this resource to a small number of users that explicitly require access. The only supported client for the **BundleDeployment** resource during this Technology Preview release is the platform Operators manager component.
 - OLM's Marketplace component is an optional cluster capability that can be disabled. This has implications during the Technology Preview release because platform Operators are currently only sourced from the **redhat-operators** catalog source that is managed by the Marketplace component. As a workaround, a cluster administrator can create this catalog source manually.
 - The RukPak provisioner implementations do not have the ability to inspect the health or state of the resources that they are managing. This has implications for surfacing the generated **BundleDeployment** resource state to the **PlatformOperator** resource that owns it. If a **registry+v1** bundle contains manifests that can be successfully applied to the

cluster, but will fail at runtime, such as a **Deployment** object referencing a non-existent image, the result is a successful status being reflected in individual **PlatformOperator** and **BundleDeployment** resources.

- Cluster administrators configuring **PlatformOperator** resources before cluster creation cannot easily determine the desired package name without leveraging an existing cluster or relying on documented examples. There is currently no validation logic that ensures an individually configured **PlatformOperator** resource will be able to successfully roll out to the cluster.
- When using the Technology Preview OCI feature with the `oc-mirror` CLI plugin, the mirrored catalog embeds all of the Operator bundles, instead of filtering only on those specified in the image set configuration file. ([OCPBUGS-5085](#))
- There is currently a known issue when you run the Agent-based OpenShift Container Platform Installer to generate an ISO image from a directory where the previous release was used for ISO image generation. An error message is displayed with the release version not matching. As a workaround, create and use a new directory. ([OCPBUGS#5159](#))
- The defined capabilities in the **install-config.yaml** file are not applied in the Agent-based OpenShift Container Platform installation. Currently, there is no workaround. ([OCPBUGS#5129](#))
- Fully populated load balancers on RHOSP that are created with the OVN driver can contain pools that are stuck in a pending creation status. This issue can cause problems for clusters that are deployed on RHOSP. To resolve the issue, update your RHOSP packages. ([BZ#2042976](#))
- Bulk load-balancer member updates on RHOSP can return a 500 code in response to **PUT** requests. This issue can cause problems for clusters that are deployed on RHOSP. To resolve the issue, update your RHOSP packages. ([BZ#2100135](#))
- Clusters that use external cloud providers can fail to retrieve updated credentials after rotation. The following platforms are affected:
 - Alibaba Cloud
 - IBM Cloud VPC
 - IBM Power
 - OpenShift Virtualization
 - RHOSP

As a workaround, restart **openshift-cloud-controller-manager** pods by running the following command:

```
$ oc delete pods --all -n openshift-cloud-controller-manager
```

([OCPBUGS-5036](#))

- There is a known issue when **cloud-provider-openstack** tries to create health monitors on OVN load balancers by using the API to create fully populated load balancers. These health monitors become stuck in a **PENDING_CREATE** status. After their deletion, associated load balancers are stuck in a **PENDING_UPDATE** status. There is no workaround. ([BZ#2143732](#))

- Due to a known issue, to use stateful IPv6 networks with cluster that run on RHOSP, you must include **ip=dhcp,dhcpv6** in the kernel arguments of [worker nodes](#). ([OCBUGS-2104](#))
- It is not possible to create a macvlan on the physical function (PF) when a virtual function (VF) already exists. This issue affects the Intel E810 NIC. ([BZ#2120585](#))
- There is currently a known issue when manually configuring IPv6 addresses and routes on an IPv4 OpenShift Container Platform cluster. When converting to a dual-stack cluster, newly created pods remain in the **ContainerCreating** status. Currently, there is no workaround. This issue is planned to be addressed in a future OpenShift Container Platform release. ([OCBUGS-4411](#))
- When an OVN cluster installed on IBM Public Cloud has more than 60 worker nodes, simultaneously creating 2000 or more services and route objects can cause pods created at the same time to remain in the **ContainerCreating** status. If this problem occurs, entering the **oc describe pod <podname>** command shows events with the following warning:
FailedCreatePodSandBox...failed to configure pod interface: timed out waiting for OVS port binding (ovn-installed). There is currently no workaround for this issue. ([OCBUGS-3470](#))
- When a control plane machine is replaced on a cluster that uses the OVN-Kubernetes network provider, the pods related to OVN-Kubernetes might not start on the replacement machine. When this occurs, the lack of networking on the new machine prevents etcd from allowing it to replace the old machine. As a result, the cluster is stuck in this state and might become degraded. This behavior can occur when the control plane is replaced manually or by the control plane machine set.
There is currently no workaround to resolve this issue if encountered. To avoid this issue, [disable the control plane machine set](#) and do not replace control plane machines manually if your cluster uses the OVN-Kubernetes network provider. ([OCBUGS-5306](#))
- If a cluster that was deployed through ZTP has policies that do not become compliant, and no **ClusterGroupUpdates** object is present, you must restart the TALM pods. Restarting TALM creates the proper **ClusterGroupUpdates** object, which enforces the policy compliance. ([OCBUGS-4065](#))
- Currently, a certificate compliance issue, specifically outputted as **x509: certificate is not standards compliant**, exists when you run the installation program on macOS for the purposes of installing an OpenShift Container Platform cluster on VMware vSphere. This issue relates to a known issue with the **golang** compiler in that the compiler does not recognize newly supported macOS certificate standards. No workaround exists for this issue. ([OSDOCS-5694](#))
- Currently, when using a persistent volume (PV) that contains a very large number of files, the pod might not start or can take an excessive amount of time to start. For more information, see this [knowledge base article](#). ([BZ1987112](#))
- Creating pods with Azure File NFS volumes that are scheduled to the control plane node causes the mount to be denied. ([OCBUGS-18581](#))
To work around this issue: If your control plane nodes are schedulable, and the pods can run on worker nodes, use **nodeSelector** or Affinity to schedule the pod in worker nodes.
- When installing an OpenShift Container Platform cluster with static IP addressing and Tang encryption, nodes start without network settings. This condition prevents nodes from accessing the Tang server, causing installation to fail. To address this condition, you must set the network settings for each node as **ip** installer arguments.
 1. For installer-provisioned infrastructure, before installation provide the network settings as

ip installer arguments for each node by executing the following steps.

- a. Create the manifests.
- b. For each node, modify the **BareMetalHost** custom resource with annotations to include the network settings. For example:

```
$ cd ~/clusterconfigs/openshift
$ vim openshift-worker-0.yaml
```

```
apiVersion: metal3.io/v1alpha1
kind: BareMetalHost
metadata:
  annotations:
    bmac.agent-install.openshift.io/installer-args: ["--append-karg", "ip=<static_ip>::
<gateway>:<netmask>:<hostname_1>:<interface>:none", "--save-partindex", "1", "n"]
    inspect.metal3.io: disabled
    bmac.agent-install.openshift.io/hostname: <fqdn>
    bmac.agent-install.openshift.io/role: <role>
  generation: 1
  name: openshift-worker-0
  namespace: mynamespace
spec:
  automatedCleaningMode: disabled
  bmc:
    address: idrac-virtualmedia://<bmc_ip>/redfish/v1/Systems/System.Embedded.
    credentialsName: bmc-secret-openshift-worker-0
    disableCertificateVerification: true
  bootMACAddress: 94:6D:AE:AB:EE:E8
  bootMode: "UEFI"
  rootDeviceHints:
    deviceName: /dev/sda
```

For the **ip** settings, replace:

- 1 **<static_ip>** with the static IP address for the node, for example, **192.168.1.100**
- 2 **<gateway>** with the IP address of your network's gateway, for example, **192.168.1.1**
- 3 **<netmask>** with the network mask, for example, **255.255.255.0**
- 4 **<hostname_1>** with the node's hostname, for example, **node1.example.com**
- 5 **<interface>** with the name of the network interface, for example, **eth0**
- 6 **<fqdn>** with the fully qualified domain name of the node
- 7 **<role>** with **worker** or **master** to reflect the node's role
- 8 **<bmc_ip>** with the BMC IP address and the protocol and path of the BMC, as needed.

- c. Save the file to the **clusterconfigs/openshift** directory.
 - d. Create the cluster.
2. When installing with the Assisted Installer, before installation modify each node's installer arguments using the API to append the network settings as **ip** installer arguments. For example:

```
$ curl https://api.openshift.com/api/assisted-install/v2/infra-
envs/${infra_env_id}/hosts/${host_id}/installer-args \
-X PATCH \
-H "Authorization: Bearer ${API_TOKEN}" \
-H "Content-Type: application/json" \
-d '
{
  "args": [
    "--append-karg",
    "ip=<static_ip>:<gateway>:<netmask>:<hostname_1>:<interface>:none", 1 2
  ] 3 4 5
  "--save-partindex",
  "1",
  "-n"
}
'| jq
```

For the previous network settings, replace:

- 1 <static_ip> with the static IP address for the node, for example, **192.168.1.100**
- 2 <gateway> with the IP address of your network's gateway, for example, **192.168.1.1**
- 3 <netmask> with the network mask, for example, **255.255.255.0**
- 4 <hostname_1> with the node's hostname, for example, **node1.example.com**
- 5 <interface> with the name of the network interface, for example, **eth0**.

Contact Red Hat Support for additional details and assistance.

([OCPBUGS-23119](#))

1.9. ASYNCHRONOUS ERRATA UPDATES


Security, bug fix, and enhancement updates for OpenShift Container Platform 4.12 are released as asynchronous errata through the Red Hat Network. All OpenShift Container Platform 4.12 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 4.12. Versioned asynchronous releases, for example with the form OpenShift Container Platform 4.12.z, will be detailed in subsections. In addition, releases in which the errata text cannot fit in the space provided by the advisory will be detailed in subsections that follow.

**IMPORTANT**

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

1.9.1. RHSA-2022:7399 - OpenShift Container Platform 4.12.0 image release, bug fix, and security update advisory

Issued: 17 January 2023

OpenShift Container Platform release 4.12.0, which includes security updates, is now available. The list of bug fixes that are included in the update is documented in the [RHSA-2022:7399](#) advisory. The RPM packages that are included in the update are provided by the [RHSA-2022:7398](#) advisory.

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

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

```
$ oc adm release info 4.12.0 --pullspecs
```

1.9.1.1. Features

1.9.1.1.1. General availability of pod-level bonding for secondary networks

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

1.9.2. RHSA-2023:0449 - OpenShift Container Platform 4.12.1 bug fix and security update

Issued: 30 January 2023

OpenShift Container Platform release 4.12.1, which includes security updates, is now available. The list of bug fixes that are included in the update is documented in the [RHSA-2023:0449](#) advisory. The RPM packages that are included in the update are provided by the [RHBA-2023:0448](#) advisory.

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

```
$ oc adm release info 4.12.1 --pullspecs
```

1.9.2.1. Bug fixes

- Previously, due to a wrong check in the OpenStack cloud provider, the load balancers were populated with External IP addresses when all of the Octavia load balancers were created. This increased the time for the load balancers to be handled. With this update, load balancers are still created sequentially and External IP addresses are populated one-by-one. ([OCBUGS-5403](#))
- Previously, the **cluster-image-registry-operator** would default to using persistent volume claim (PVC) when it failed to reach Swift. With this update, failure to connect to Red Hat OpenStack Platform (RHOSP) API or other incidental failures cause the **cluster-image-registry-operator** to retry the probe. During the retry, the default to PVC only occurs if the RHOSP catalog is correctly found, and it does not contain object storage; or alternatively, if RHOSP catalog is there and the current user does not have permission to list containers. ([OCBUGS-5154](#))

1.9.2.2. Updating

To update an existing OpenShift Container Platform 4.12 cluster to this latest release, see [Updating a cluster using the CLI](#).

1.9.3. RHSA-2023:0569 - OpenShift Container Platform 4.12.2 bug fix and security update

Issued: 7 February 2023

OpenShift Container Platform release 4.12.2, which includes security updates, is now available. The list of bug fixes that are included in the update is documented in the [RHSA-2023:0569](#) advisory. The RPM packages that are included in the update are provided by the [RHBA-2023:0568](#) advisory.

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

```
$ oc adm release info 4.12.2 --pullspecs
```

1.9.3.1. Updating

To update an existing OpenShift Container Platform 4.12 cluster to this latest release, see [Updating a cluster using the CLI](#).

1.9.4. RHSA-2023:0728 - OpenShift Container Platform 4.12.3 bug fix and security update

Issued: 16 February 2023

OpenShift Container Platform release 4.12.3, which includes security updates, is now available. The list of bug fixes that are included in the update is documented in the [RHSA-2023:0728](#) advisory. The RPM packages that are included in the update are provided by the [RHSA-2023:0727](#) advisory.

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

```
$ oc adm release info 4.12.3 --pullspecs
```

1.9.4.1. Bug fixes

- Previously, when a control plane machine was replaced on a cluster that used the OVN-

Kubernetes network provider, the pods related to OVN-Kubernetes sometimes did not start on the replacement machine, and prevented etcd from allowing it to replace the old machine. With this update, pods related to OVN-Kubernetes start in the replacement machine as expected. ([OCPBUGS-6494](#))

1.9.4.2. Updating

To update an existing OpenShift Container Platform 4.12 cluster to this latest release, see [Updating a cluster using the CLI](#).

1.9.5. RHSA-2023:0769 - OpenShift Container Platform 4.12.4 bug fix and security update

Issued: 20 February 2023

OpenShift Container Platform release 4.12.4, which includes security updates, is now available. The list of bug fixes that are included in the update is documented in the [RHSA-2023:0769](#) advisory. The RPM packages that are included in the update are provided by the [RHBA-2023:0768](#) advisory.

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

```
$ oc adm release info 4.12.4 --pullspecs
```

1.9.5.1. Updating

To update an existing OpenShift Container Platform 4.12 cluster to this latest release, see [Updating a cluster using the CLI](#).

1.9.6. RHSA-2023:0890 - OpenShift Container Platform 4.12.5 bug fix and security update

Issued: 28 February 2023

OpenShift Container Platform release 4.12.5, which includes security updates, is now available. The list of bug fixes that are included in the update is documented in the [RHSA-2023:0890](#) advisory. The RPM packages that are included in the update are provided by the [RHBA-2023:0889](#) advisory.

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

```
$ oc adm release info 4.12.5 --pullspecs
```

1.9.6.1. Bug fixes

- Previously, in the repositories list, you could see the **PipelineRuns** only when the status was **Succeeded** or **Failed** but not when the status was **Running**. With this fix, when the **PipelineRuns** is triggered, you can see it in the repositories list with the status **Running**. ([OCPBUGS-6816](#))
- Previously, when creating a **Secret**, the **Start Pipeline** model created an invalid JSON value, As a result, the Secret was unusable and the **PipelineRun** could fail. With this fix, the **Start Pipeline** model creates a valid JSON value for the **Secret**. Now, you can create valid Secrets while starting a Pipeline. ([OCPBUGS-6671](#))

- Previously, when a **BindableKinds** resource did not have a status, the web console crashed, fetching and showing the same data in a loop. With this fix, you can set the **BindableKinds** resource status array to `[]`, expecting it to exist without a status field. As a result, the web browser or the application does not crash. ([OCPBUGS-4072](#))
- Previously, the associated webhook `<kn-service-name>-github-webhook-secret` did not delete when deleting a Knative (**kn**) service from OpenShift Container Platform. With this fix, all the associated webhook secrets are deleted. Now, you can create a Knative (**kn**) service with the same name as the deleted one. ([OCPBUGS-7437](#))

1.9.6.2. Updating

To update an existing OpenShift Container Platform 4.12 cluster to this latest release, see [Updating a cluster using the CLI](#).

1.9.7. RHSA-2023:1034 - OpenShift Container Platform 4.12.6 bug fix and security update

Issued: 7 March 2023

OpenShift Container Platform release 4.12.6, which includes security updates, is now available. The list of bug fixes that are included in the update is documented in the [RHBA-2023:1034](#) advisory. The RPM packages that are included in the update are provided by the [RHSA-2023:1033](#) advisory.

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

```
$ oc adm release info 4.12.6 --pullspecs
```

1.9.7.1. Updating

To update an existing OpenShift Container Platform 4.12 cluster to this latest release, see [Updating a cluster using the CLI](#).

1.9.8. RHBA-2023:1163 - OpenShift Container Platform 4.12.7 bug fix update

Issued: 13 March 2023

OpenShift Container Platform release 4.12.7 is now available. The list of bug fixes that are included in the update is documented in the [RHBA-2023:1163](#) advisory. The RPM packages that are included in the update are provided by the [RHBA-2023:1162](#) advisory.

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

```
$ oc adm release info 4.12.7 --pullspecs
```

1.9.8.1. Updating

To update an existing OpenShift Container Platform 4.12 cluster to this latest release, see [Updating a cluster using the CLI](#).

1.9.9. RHBA-2023:1269 - OpenShift Container Platform 4.12.8 bug fix and security update

Issued: 21 March 2023

OpenShift Container Platform release 4.12.8, which includes security updates, is now available. The list of bug fixes that are included in the update is documented in the [RHBA-2023:1269](#) advisory. The RPM packages that are included in the update are provided by the [RHSA-2023:1268](#) advisory.

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

```
$ oc adm release info 4.12.8 --pullspecs
```

1.9.9.1. Updating

To update an existing OpenShift Container Platform 4.12 cluster to this latest release, see [Updating a cluster using the CLI](#).

1.9.10. RHSA-2023:1409 – OpenShift Container Platform 4.12.9 bug fix and security update

Issued: 27 March 2023

OpenShift Container Platform release 4.12.9, which includes security updates, is now available. The list of bug fixes that are included in the update is documented in the [RHSA-2023:1409](#) advisory. The RPM packages that are included in the update are provided by the [RHSA-2023:1408](#) advisory.

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

```
$ oc adm release info 4.12.9 --pullspecs
```

1.9.10.1. Bug fixes

- Previously, validation was not preventing users from installing a GCP cluster into a shared VPC if they did not enable the Technology Preview feature gate. Therefore, you could install a cluster into a shared VPC without enabling the Technology Preview feature gate. This release added a feature gate validation to 4.12 so you must enable **featureSet: TechPreviewNoUpgrade** to install a GCP cluster into a shared VPC. ([OCPBUGS-7469](#))
- Previously, MTU migration configuration would sometimes be cleaned up before the migration was complete causing the migration to fail. This release ensures that the MTU migration is preserved while migration is in progress so that the migration can complete successfully. ([OCPBUGS-7445](#))

1.9.10.2. Updating

To update an existing OpenShift Container Platform 4.12 cluster to this latest release, see [Updating a cluster using the CLI](#).

1.9.11. RHBA-2023:1508 – OpenShift Container Platform 4.12.10 bug fix update

Issued: 3 April 2023

OpenShift Container Platform release 4.12.10 is now available. The list of bug fixes that are included in the update is documented in the [RHBA-2023:1508](#) advisory. The RPM packages that are included in the update are provided by the [RHBA-2023:1507](#) advisory.

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

```
$ oc adm release info 4.12.10 --pullspecs
```

1.9.11.1. Updating

To update an existing OpenShift Container Platform 4.12 cluster to this latest release, see [Updating a cluster using the CLI](#).

1.9.12. RHSA-2023:1645 - OpenShift Container Platform 4.12.11 bug fix and security update

Issued: 11 April 2023

OpenShift Container Platform release 4.12.11, which includes security updates, is now available. The list of bug fixes that are included in the update is documented in the [RHBA-2023:1645](#) advisory. The RPM packages that are included in the update are provided by the [RHBA-2023:1644](#) advisory.

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

```
$ oc adm release info 4.12.11 --pullspecs
```

1.9.12.1. Features

1.9.12.1.1. New flag for the oc-mirror plugin: --max-nested-paths

With this update, you can now use the **--max-nested-paths** flag for the oc-mirror plugin to specify the maximum number of nested paths for destination registries that limit nested paths. The default is **2**.

1.9.12.1.2. New flag for the oc-mirror plugin: --skip-pruning

With this update, you can now use the **--skip-pruning** flag for the oc-mirror plugin to disable automatic pruning of images from the target mirror registry.

1.9.12.2. Bug fixes

- Previously, the **openshift-install agent create cluster-manifests** command required a non-empty list of **imageContentSources** in the **install-config.yaml** file. If no image content sources were supplied, the command generated the error **failed to write asset (Mirror Registries Config) to disk: failed to write file: open .: is a directory**. With this update, the command works whether or not the **imageContentSources** section of **install-config.yaml** file contains anything. ([OCPBUGS-8384](#))
- Previously, the OpenStack Machine API provider had to be restarted so that new cloud credentials were used in the event of a rotation of the OpenStack **clouds.yaml** file. Consequently, the ability of a MachineSet to scale to zero was affected. With this update, cloud credentials are no longer cached and the OpenStack Machine API provider reads the corresponding secret on demand. ([OCPBUGS-10603](#))

1.9.12.3. Updating

To update an existing OpenShift Container Platform 4.12 cluster to this latest release, see [Updating a cluster using the CLI](#).

1.9.13. RHBA-2023:1734 - OpenShift Container Platform 4.12.12 bug fix

Issued: 13 April 2023

OpenShift Container Platform release 4.12.12 is now available. The list of bug fixes that are included in the update is documented in the [RHBA-2023:1734](#) advisory. There are no RPM packages for this update.

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

```
$ oc adm release info 4.12.12 --pullspecs
```

1.9.13.1. Updating

All OpenShift Container Platform 4.12 users are advised that the only defect fixed in this release is limited to install time; therefore, there is no need to update previously installed clusters to this version.

1.9.14. RHBA-2023:1750 - OpenShift Container Platform 4.12.13 bug fix update

Issued: 19 April 2023

OpenShift Container Platform release 4.12.13 is now available. The list of bug fixes that are included in the update is documented in the [RHBA-2023:1750](#) advisory. The RPM packages that are included in the update are provided by the [RHBA-2023:1749](#) advisory.

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

```
$ oc adm release info 4.12.13 --pullspecs
```

1.9.14.1. Features

1.9.14.1.1. Pod security admission restricted enforcement (Technology Preview)

With this release, pod security admission restricted *enforcement* is available as a Technology Preview feature by enabling the **TechPreviewNoUpgrade** feature set. If you enable the **TechPreviewNoUpgrade** feature set, pods are rejected if they violate pod security standards, instead of only logging a warning.



NOTE

Pod security admission restricted enforcement is only activated if you enable the **TechPreviewNoUpgrade** feature set after your OpenShift Container Platform cluster is installed. It is not activated if you enable the **TechPreviewNoUpgrade** feature set during cluster installation.

For more information, see [Understanding feature gates](#).

1.9.14.2. Updating

To update an existing OpenShift Container Platform 4.12 cluster to this latest release, see [Updating a cluster using the CLI](#).

1.9.15. RHBA-2023:1858 - OpenShift Container Platform 4.12.14 bug fix update

Issued: 24 April 2023

OpenShift Container Platform release 4.12.14 is now available. The list of bug fixes that are included in the update is documented in the [RHBA-2023:1858](#) advisory. The RPM packages that are included in the update are provided by the [RHBA-2023:1857](#) advisory.

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

```
$ oc adm release info 4.12.14 --pullspecs
```

1.9.15.1. Features

1.9.15.1.1. Cloud provider OpenStack is updated to 1.25

With this release, Cloud Provider Red Hat OpenStack Platform (RHOSP) is updated to 1.25.5. The update includes the addition of an annotation for real load balancer IP addresses and the global source for **math/rand** packages are seeded in **main.go**.

1.9.15.2. Updating

To update an existing OpenShift Container Platform 4.12 cluster to this latest release, see [Updating a cluster using the CLI](#).

1.9.16. RHBA-2023:2037 - OpenShift Container Platform 4.12.15 bug fix update

Issued: 3 May 2023

OpenShift Container Platform release 4.12.15 is now available. The list of bug fixes that are included in the update is documented in the [RHBA-2023:2037](#) advisory. The RPM packages that are included in the update are provided by the [RHBA-2023:2036](#) advisory.

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

```
$ oc adm release info 4.12.15 --pullspecs
```

1.9.16.1. Bug fixes

- Previously, the Cluster Network Operator (CNO) configuration ignored Kuryr's maximum transmission unit (MTU) settings when using the Red Hat OpenStack Platform (RHOSP) Networking service (neutron) component to create a network for OpenShift services. CNO would create a network in Neutron with the wrong MTU property, and this action could cause incompatibility issues among network components. With this update, the CNO does not ignore the Kuryr MTU setting when creating the network for services. You can then use the network to host OpenShift services. ([OCPBUGS-4896](#))

1.9.16.2. Updating

To update an existing OpenShift Container Platform 4.12 cluster to this latest release, see [Updating a cluster using the CLI](#).

1.9.17. RHSA-2023:2110 - OpenShift Container Platform 4.12.16 bug fix and security update

Issued: 10 May 2023

OpenShift Container Platform release 4.12.16, which includes security updates, is now available. The list of bug fixes that are included in the update is documented in the [RHSA-2023:2110](#) advisory. The RPM packages that are included in the update are provided by the [RHBA-2023:2109](#) advisory.

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

```
$ oc adm release info 4.12.16 --pullspecs
```

1.9.17.1. Bug fixes

- Previously, in the **Import from Git** and **Deploy Image** flows, the **Resource Type** section was moved to **Advanced** section. As a result, it was difficult to identify the type of resource created. With this fix, **Resource Type** section is moved to the **General** section. ([OCPBUGS-7395](#))

1.9.17.2. Updating

To update an existing OpenShift Container Platform 4.12 cluster to this latest release, see [Updating a cluster using the CLI](#).

1.9.18. RHBA-2023:2699 - OpenShift Container Platform 4.12.17 bug fix update

Issued: 18 May 2023

OpenShift Container Platform release 4.12.17 is now available. The list of bug fixes that are included in the update is documented in the [RHBA-2023:2699](#) advisory. The RPM packages that are included in the update are provided by the [RHBA-2023:2698](#) advisory.

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

```
$ oc adm release info 4.12.17 --pullspecs
```

1.9.18.1. Bug fixes

- Previously, you used the edit form for creating **ConfigMaps**, **Secrets**, **Deployments**, and **DeploymentConfigs**. For **BuildConfigs**, you used the edit form only for editing. With this fix, you can use the edit form for creating **BuildConfigs** too. ([OCPBUGS-9336](#))

1.9.18.2. Updating

To update an existing OpenShift Container Platform 4.12 cluster to this latest release, see [Updating a cluster using the CLI](#).

1.9.19. RHBA-2023:3208 - OpenShift Container Platform 4.12.18 bug fix update

Issued: 23 May 2023

OpenShift Container Platform release 4.12.18 is now available. The list of bug fixes that are included in the update is documented in the [RHBA-2023:3208](#) advisory. The RPM packages that are included in the update are provided by the [RHBA-2023:3207](#) advisory.

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

```
$ oc adm release info 4.12.18 --pullspecs
```

1.9.19.1. Bug fixes

- Previously, the **Samples** page in the OpenShift Container Platform did not allow distinguishing between the types of samples listed. With this fix, you can identify the sample from the badges displayed on the **Samples** page. ([OCBUGS-7446](#))
- Previously, when viewing resource consumption for a specific pod, graphs displaying **CPU usage** and **Memory Usage** metrics were stacked even though these metrics are static values, which should be displayed as a static line across the graph. With this update, OpenShift Container Platform correctly displays the values for **CPU Usage** and **Memory Usage** in the monitoring dashboard. ([OCBUGS-5353](#))

1.9.19.2. Updating

To update an existing OpenShift Container Platform 4.12 cluster to this latest release, see [Updating a cluster using the CLI](#).

1.9.20. RHSA-2023:3287 - OpenShift Container Platform 4.12.19 bug fix and security update

Issued: 31 May 2023

OpenShift Container Platform release 4.12.19 is now available. The list of bug fixes that are included in the update is documented in the [RHSA-2023:3287](#) advisory. The RPM packages that are included in the update are provided by the [RHBA-2023:3286](#) advisory.

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

```
$ oc adm release info 4.12.19 --pullspecs
```

1.9.20.1. Updating

To update an existing OpenShift Container Platform 4.12 cluster to this latest release, see [Updating a cluster using the CLI](#).

1.9.21. RHSA-2023:3410 - OpenShift Container Platform 4.12.20 bug fix update

Issued: 7 June 2023

OpenShift Container Platform release 4.12.20 is now available. The list of bug fixes that are included in the update is documented in the [RHSA-2023:3410](#) advisory. The RPM packages that are included in the update are provided by the [RHSA-2023:3409](#) advisory.

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

```
$ oc adm release info 4.12.20 --pullspecs
```

1.9.21.1. Bug fixes

- Previously, mirroring from a registry to a disk by using an image set configuration file that specifies several digests of the same image, without tags, caused an error because the oc-mirror plugin added a default tag **latest** to all the images (digests). With this update, the oc-mirror plugin now uses a truncated digest of the image, which eliminates the error. ([OCPBUGS-13432](#))

1.9.21.2. Updating

To update an existing OpenShift Container Platform 4.12 cluster to this latest release, see [Updating a cluster using the CLI](#).

1.9.22. RHBA-2023:3546 - OpenShift Container Platform 4.12.21 bug fix and security update

Issued: 14 June 2023

OpenShift Container Platform release 4.12.21, which includes security updates, is now available. The list of bug fixes that are included in the update is documented in the [RHBA-2023:3546](#) advisory. The RPM packages that are included in the update are provided by the [RHSA-2023:3545](#) advisory.

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

```
$ oc adm release info 4.12.21 --pullspecs
```

1.9.22.1. Bug fixes

- Previously on single-node OpenShift, in case of node reboot there was a race condition that could result in admission of application pods requesting devices on the node even if devices were unhealthy or unavailable to be allocated. This resulted in runtime failures when the application tried to access devices. With this update, the resources requested by the pod are only allocated if the device plugin has registered itself to kubelet and healthy devices are present on the node to be allocated.
If these conditions are not met, the pod can fail at admission with **UnexpectedAdmissionError** error, which is an expected behavior. If the application pod is part of deployments, in case of failure subsequent pods are created and ultimately successfully run when devices are suitable to be allocated. ([OCPBUGS-14437](#))

1.9.22.2. Updating

To update an existing OpenShift Container Platform 4.12 cluster to this latest release, see [Updating a cluster using the CLI](#).

1.9.23. RHSA-2023:3615 - OpenShift Container Platform 4.12.22 bug fix and security update

Issued: 26 June 2023

OpenShift Container Platform release 4.12.22, which includes security updates, is now available. The list of bug fixes that are included in the update is documented in the [RHSA-2023:3615](#) advisory. The RPM packages that are included in the update are provided by the [RHSA-2023:3613](#) advisory.

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

```
$ oc adm release info 4.12.22 --pullspecs
```

1.9.23.1. Bug fixes

- Previously, client TLS (mTLS) was configured on an Ingress Controller, and the certificate authority (CA) in the client CA bundle required more than 1MB of certificate revocation list (CRLs) to be downloaded. The CRL **ConfigMap** object size limitations prevented updates from occurring. As a result of the missing CRLs, connections with valid client certificates may have been rejected with the error **unknown ca**. With this update, the CRL **ConfigMap** for each Ingress Controller no longer exists; instead, each router pod directly downloads CRLs, ensuring connections with valid client certificates are no longer rejected. ([OCPBUGS-14454](#))
- Previously, because client TLS (mTLS) was configured on an Ingress Controller, mismatches between the distributing certificate authority (CA) and the issuing CA caused the incorrect certificate revocation list (CRL) to be downloaded. As a result, the incorrect CRL was downloaded instead of the correct CRL, causing connections with valid client certificates to be rejected with the error message **unknown ca**. With this update, downloaded CRLs are now tracked by the CA that distributes them. This ensures that valid client certificates are no longer rejected. ([OCPBUGS-14455](#))

1.9.23.2. Updating

To update an existing OpenShift Container Platform 4.12 cluster to this latest release, see [Updating a cluster using the CLI](#).

1.9.24. RHSA-2023:3925 - OpenShift Container Platform 4.12.23 bug fix and security update

Issued: 6 July 2023

OpenShift Container Platform release 4.12.23, which includes security updates, is now available. This update includes a Red Hat security bulletin for customers who run OpenShift Container Platform in FIPS mode. For more information, see [RHSA-2023:001](#).

The list of bug fixes that are included in the update is documented in the [RHSA-2023:3925](#) advisory. The RPM packages that are included in the update are provided by the [RHSA-2023:3924](#) advisory.

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

```
$ oc adm release info 4.12.23 --pullspecs
```

1.9.24.1. Updating

To update an existing OpenShift Container Platform 4.12 cluster to this latest release, see [Updating a cluster using the CLI](#).

1.9.25. RHBA-2023:3977 - OpenShift Container Platform 4.12.24 bug fix and security update

Issued: 12 July 2023

OpenShift Container Platform release 4.12.24, which includes security updates, is now available. The list of bug fixes that are included in the update is documented in the [RHBA-2023:3977](#) advisory. The RPM packages that are included in the update are provided by the [RHSA-2023:3976](#) advisory.

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

```
$ oc adm release info 4.12.24 --pullspecs
```

1.9.25.1. Features

1.9.25.2. NUMA-aware scheduling with the NUMA Resources Operator is generally available

NUMA-aware scheduling with the NUMA Resources Operator was previously introduced as a Technology Preview in OpenShift Container Platform 4.10. It is now generally available in OpenShift Container Platform version 4.12.24 and later.

The NUMA Resources Operator deploys a NUMA-aware secondary scheduler that makes scheduling decisions for workloads based on a complete picture of available NUMA zones in clusters. This enhanced NUMA-aware scheduling ensures that latency-sensitive workloads are processed in a single NUMA zone for maximum efficiency and performance.

This update adds the following features:

- Fine-tuning of API polling for NUMA resource reports.
- Configuration options at the node group level for the node topology exporter.



NOTE

NUMA-aware scheduling with the NUMA Resources Operator is not yet available on single-node OpenShift.

For more information, see [Scheduling NUMA-aware workloads](#).

1.9.25.3. Updating

To update an existing OpenShift Container Platform 4.12 cluster to this latest release, see [Updating a cluster using the CLI](#).

1.9.26. RHBA-2023:4048 - OpenShift Container Platform 4.12.25 bug fix update

Issued: 19 July 2023

OpenShift Container Platform release 4.12.25 is now available. The list of bug fixes that are included in the update is documented in the [RHBA-2023:4048](#) advisory. The RPM packages that are included in the update are provided by the [RHBA-2023:4047](#) advisory.

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

```
$ oc adm release info 4.12.25 --pullspecs
```

1.9.26.1. Updating

To update an existing OpenShift Container Platform 4.12 cluster to this latest release, see [Updating a cluster using the CLI](#).

1.9.27. RHBA-2023:4221 - OpenShift Container Platform 4.12.26 bug fix update

Issued: 26 July 2023

OpenShift Container Platform release 4.12.26 is now available. The list of bug fixes that are included in the update is documented in the [RHBA-2023:4221](#) advisory. There are no RPM packages for this update.

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

```
$ oc adm release info 4.12.26 --pullspecs
```

1.9.27.1. Updating

To update an existing OpenShift Container Platform 4.12 cluster to this latest release, see [Updating a cluster using the CLI](#).

1.9.28. RHBA-2023:4319 - OpenShift Container Platform 4.12.27 bug fix update

Issued: 2 August 2023

OpenShift Container Platform release 4.12.27 is now available. The list of bug fixes that are included in the update is documented in the [RHBA-2023:4319](#) advisory. The RPM packages that are included in the update are provided by the [RHBA-2023:4322](#) advisory.

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

```
$ oc adm release info 4.12.27 --pullspecs
```

1.9.28.1. Updating

To update an existing OpenShift Container Platform 4.12 cluster to this latest release, see [Updating a cluster using the CLI](#).

1.9.29. RHBA-2023:4440 - OpenShift Container Platform 4.12.28 bug fix update

Issued: 9 August 2023

OpenShift Container Platform release 4.12.28 is now available. The list of bug fixes that are included in the update is documented in the [RHBA-2023:4440](#) advisory. The RPM packages that are included in the update are provided by the [RHBA-2023:4443](#) advisory.

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

```
$ oc adm release info 4.12.28 --pullspecs
```

1.9.29.1. Updating

To update an existing OpenShift Container Platform 4.12 cluster to this latest release, see [Updating a cluster using the CLI](#).

1.9.30. RHBA-2023:4608 - OpenShift Container Platform 4.12.29 bug fix update

Issued: 16 August 2023

OpenShift Container Platform release 4.12.29 is now available. The list of bug fixes that are included in the update is documented in the [RHBA-2023:4608](#) advisory. The RPM packages that are included in the update are provided by the [RHBA-2023:4611](#) advisory.

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

```
$ oc adm release info 4.12.29 --pullspecs
```

1.9.30.1. Bug fixes

- Previously, the Ingress Operator did not include an Amazon Web Services (AWS) permission in its cloud credentials request. This impacted the management of domain name system (DNS) records in the Commercial Cloud Services (C2S) **us-iso-east-1** and the Secret Commercial Cloud Services (SC2S) **us-isob-east-1** AWS Regions. If you installed an OpenShift Container Platform cluster in a C2S or an SC2S AWS Region, the Ingress Operator failed to publish DNS records for the Route 53 service and you received an error message similar to the following example:

```
The DNS provider failed to ensure the record: failed to find hosted zone for record: failed to
get tagged resources: AccessDenied: User: [...] is not authorized to perform:
route53:ListTagsForResources on resource: [...]
```

With this update, the Ingress Operator's cloud credentials request includes the **route53:ListTagsForResources** permission, so that the Operator can publish DNS records in the C2S and SC2S AWS Regions for the Route 53 service. ([OCPBUGS-15467](#))

1.9.30.2. Updating

To update an existing OpenShift Container Platform 4.12 cluster to this latest release, see [Updating a cluster using the CLI](#).

1.9.31. RHSA-2023:4671 - OpenShift Container Platform 4.12.30 bug fix update

Issued: 23 August 2023

OpenShift Container Platform release 4.12.30, which includes security updates, is now available. The list of bug fixes that are included in the update is documented in the [RHSA-2023:4671](#) advisory. The RPM packages that are included in the update are provided by the [RHSA-2023:4674](#) advisory.

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

```
$ oc adm release info 4.12.30 --pullspecs
```

1.9.31.1. Updating

To update an existing OpenShift Container Platform 4.12 cluster to this latest release, see [Updating a cluster using the CLI](#).

1.9.32. RHBA-2023:4756 - OpenShift Container Platform 4.12.31 bug fix update

Issued: 31 August 2023

OpenShift Container Platform release 4.12.31 is now available. The list of bug fixes that are included in the update is documented in the [RHBA-2023:4756](#) advisory. The RPM packages that are included in the update are provided by the [RHBA-2023:4759](#) advisory.

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

```
$ oc adm release info 4.12.31 --pullspecs
```

1.9.32.1. Updating

To update an existing OpenShift Container Platform 4.12 cluster to this latest release, see [Updating a cluster using the CLI](#).

1.9.33. RHBA-2023:4900 - OpenShift Container Platform 4.12.32 bug fix update

Issued: 6 September 2023

OpenShift Container Platform release 4.12.32 is now available. The list of bug fixes that are included in the update is documented in the [RHBA-2023:4900](#) advisory. The RPM packages that are included in the update are provided by the [RHBA-2023:4903](#) advisory.

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

```
$ oc adm release info 4.12.32 --pullspecs
```

1.9.33.1. Bug fix

- Previously, an issue was observed in OpenShift Container Platform with some pods getting stuck in the **terminating** state. This affected the reconciliation loop of the allowlist controller, which resulted in unwanted retries that caused the creation of multiple pods. With this update, the allowlist controller only inspects pods that belong to the current daemon set. As a result, retries no longer occur when one or more pods are not ready. ([OCPBUGS-16019](#))

1.9.33.2. Updating

To update an existing OpenShift Container Platform 4.12 cluster to this latest release, see [Updating a cluster using the CLI](#).

1.9.34. RHBA-2023:5016 - OpenShift Container Platform 4.12.33 bug fix update

Issued: 12 September 2023

OpenShift Container Platform release 4.12.33 is now available. The list of bug fixes that are included in the update is documented in the [RHBA-2023:5016](#) advisory. The RPM packages that are included in the update are provided by the [RHBA-2023:5018](#) advisory.

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

```
$ oc adm release info 4.12.33 --pullspecs
```

1.9.34.1. Updating

To update an existing OpenShift Container Platform 4.12 cluster to this latest release, see [Updating a cluster using the CLI](#).

1.9.35. RHBA-2023:5151 - OpenShift Container Platform 4.12.34 bug fix update

Issued: 20 September 2023

OpenShift Container Platform release 4.12.34 is now available. The list of bug fixes that are included in the update is documented in the [RHBA-2023:5151](#) advisory. There are no RPM packages for this release.

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

```
$ oc adm release info 4.12.34 --pullspecs
```

1.9.35.1. Bug fixes

- Previously, a non-compliant upstream DNS server that provided a UDP response larger than OpenShift Container Platform specified bufsize of 512 bytes, caused an overflow error in CoreDNS in which a response to a DNS query was not given. With this update, users can configure the **protocolStrategy** field on the **dnsmas.operator.openshift.io** custom resource to be "TCP". This resolves issues with non-compliant upstream DNS servers. ([OCPBUGS-15251](#))
- Previously, the OpenShift Container Platform Router directed traffic to a route with a weight of **0** when it had only one back end. With this update, the router will not send traffic to routes with a single back end with weight **0**. ([OCPBUGS-18639](#))
- Previously, the cloud credentials used in Manila CSI Driver Operator were cached, resulting in authentication issues if these credentials were rotated. With this update, this issue is resolved. ([OCPBUGS-18475](#))

1.9.35.2. Updating

To update an existing OpenShift Container Platform 4.12 cluster to this latest release, see [Updating a cluster using the CLI](#).

1.9.36. RHBA-2023:5321 - OpenShift Container Platform 4.12.35 bug fix update

Issued: 27 September 2023

OpenShift Container Platform release 4.12.35 is now available. The list of bug fixes that are included in the update is documented in the [RHBA-2023:5321](#) advisory. The RPM packages that are included in the update are provided by the [RHBA-2023:5323](#) advisory.

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

```
$ oc adm release info 4.12.35 --pullspecs
```

1.9.36.1. Updating

To update an existing OpenShift Container Platform 4.12 cluster to this latest release, see [Updating a cluster using the CLI](#).

1.9.37. RHSA-2023:5390 - OpenShift Container Platform 4.12.36 bug fix and security update

Issued: 4 October 2023

OpenShift Container Platform release 4.12.36, which includes security updates, is now available. The list of bug fixes that are included in the update is documented in the [RHSA-2023:5390](#) advisory. The RPM packages that are included in the update are provided by the [RHBA-2023:5392](#) advisory.

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

```
$ oc adm release info 4.12.36 --pullspecs
```

1.9.37.1. Updating

To update an existing OpenShift Container Platform 4.12 cluster to this latest release, see [Updating a cluster using the CLI](#).

1.9.38. RHBA-2023:5450 - OpenShift Container Platform 4.12.37 bug fix update

Issued: 11 October 2023

OpenShift Container Platform release 4.12.37 is now available. The list of bug fixes that are included in the update is documented in the [RHBA-2023:5450](#) advisory. The RPM packages that are included in the update are provided by the [RHBA-2023:5452](#) advisory.

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

```
$ oc adm release info 4.12.37 --pullspecs
```

1.9.38.1. Updating

To update an existing OpenShift Container Platform 4.12 cluster to this latest release, see [Updating a cluster using the CLI](#).

1.9.39. RHSA-2023:5677 - OpenShift Container Platform 4.12.39 bug fix and security update

Issued: 18 October 2023

OpenShift Container Platform release 4.12.39, which includes security updates, is now available. The list of bug fixes that are included in the update is documented in the [RHSA-2023:5677](#) advisory. The RPM packages that are included in the update are provided by the [RHSA-2023:5679](#) advisory.

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

```
$ oc adm release info 4.12.39 --pullspecs
```

1.9.39.1. Bug fixes

- Previously, CoreDNS would crash if an EndpointSlice port was created without a port number. With this update, validation was added to CoreDNS so it will no longer crash in this situation. ([OCBUGS-20144](#))
- Previously, large clusters were slow to attach volumes through **cinder-csi-driver**. With this update, **cinder-csi-driver** is updated with slow volume attachment when the number of Cinder volumes in the project exceed 1000. ([OCBUGS-20124](#))

1.9.39.2. Updating

To update an existing OpenShift Container Platform 4.12 cluster to this latest release, see [Updating a cluster using the CLI](#).

1.9.40. RHSA-2023:5896 - OpenShift Container Platform 4.12.40 bug fix and security update

Issued: 25 October 2023

OpenShift Container Platform release 4.12.40, which includes security updates, is now available. The list of bug fixes that are included in the update is documented in the [RHSA-2023:5896](#) advisory. The RPM packages that are included in the update are provided by the [RHBA-2023:5898](#) advisory.

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

```
$ oc adm release info 4.12.40 --pullspecs
```

1.9.40.1. Updating

To update an existing OpenShift Container Platform 4.12 cluster to this latest release, see [Updating a cluster using the CLI](#).

1.9.41. RHSA-2023:6126 - OpenShift Container Platform 4.12.41 bug fix and security update

Issued: 2 November 2023

OpenShift Container Platform release 4.12.41, which includes security updates, is now available. The list of bug fixes that are included in the update is documented in the [RHSA-2023:6126](#) advisory. The RPM packages that are included in the update are provided by the [RHSA-2023:6128](#) advisory.

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

```
$ oc adm release info 4.12.41 --pullspecs
```

1.9.41.1. Updating

To update an existing OpenShift Container Platform 4.12 cluster to this latest release, see [Updating a cluster using the CLI](#).

1.9.42. RHSA-2023:6276 - OpenShift Container Platform 4.12.42 bug fix and security update

Issued: 8 November 2023

OpenShift Container Platform release 4.12.42, which includes security updates, is now available. The list of bug fixes that are included in the update is documented in the [RHSA-2023:6276](#) advisory. The RPM packages that are included in the update are provided by the [RHBA-2023:6278](#) advisory.

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

```
$ oc adm release info 4.12.42 --pullspecs
```

1.9.42.1. Feature

1.9.42.1.1. APIServer.config.openshift.io is now tracked by Insights Operator

After running the Insights Operator, a new file is now available in the archive in the path **config/apiserver.json** with the information about the audit profile for **APIServer.config.openshift.io**.

Access to audit profiles help you to understand what audit policy is common practice, what profiles are most commonly used, what differences there are between industries, and what kind of customization is applied.

1.9.42.2. Bug fixes

- Previously, the Cluster Version Operator (CVO) did not reconcile **SecurityContextConstraints** (SCC) resources as expected. The CVO now properly reconciles the **volumes** field in the **SecurityContextConstraints** resources towards the state defined in the release image. User modifications to system SCC resources are tolerated. For more information about how SCC resources can impact updating, see [Resolving Detected modified SecurityContextConstraints update gate before upgrading to 4.14](#). ([OCPBUGS-22198](#))
- Previously, a large number of **ClusterServiceVersion** (CSV) resources on startup caused a pod running the Node Tuning Operator (NTO) to restart and loop, which resulted in an error. With this update, the issue is fixed. ([OCPBUGS-21837](#))

1.9.42.3. Updating

To update an existing OpenShift Container Platform 4.12 cluster to this latest release, see [Updating a cluster using the CLI](#).

1.9.43. RHSA-2023:6842 - OpenShift Container Platform 4.12.43 bug fix and security update

Issued: 16 November 2023

OpenShift Container Platform release 4.12.43, which includes security updates, is now available. The list of bug fixes that are included in the update is documented in the [RHSA-2023:6842](#) advisory. The RPM packages that are included in the update are provided by the [RHBA-2023:6844](#) advisory.

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

```
$ oc adm release info 4.12.43 --pullspecs
```

1.9.43.1. Updating

To update an existing OpenShift Container Platform 4.12 cluster to this latest release, see [Updating a cluster using the CLI](#).

1.9.44. RHSA-2023:6894 - OpenShift Container Platform 4.12.44 bug fix and security update

Issued: 21 November 2023

OpenShift Container Platform release 4.12.44, which includes security updates, is now available. The list of bug fixes that are included in the update is documented in the [RHSA-2023:6894](#) advisory. The RPM packages that are included in the update are provided by the [RHBA-2023:6896](#) advisory.

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

```
$ oc adm release info 4.12.44 --pullspecs
```

1.9.44.1. Updating

To update an existing OpenShift Container Platform 4.12 cluster to this latest release, see [Updating a cluster using the CLI](#).

1.9.45. RHSA-2023:7608 - OpenShift Container Platform 4.12.45 bug fix and security update

Issued: 6 December 2023

OpenShift Container Platform release 4.12.45, which includes security updates, is now available. The list of bug fixes that are included in the update is documented in the [RHSA-2023:7608](#) advisory. The RPM packages that are included in the update are provided by the [RHSA-2023:7610](#) advisory.

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

```
$ oc adm release info 4.12.45 --pullspecs
```

1.9.45.1. Bug fixes

- Previously, using the cluster autoscaler with nodes that have CSI storage would cause the cluster autoscaler pods to enter in a **CrashLoopBackoff** status. With this release, you can use the cluster autoscaler with nodes that have CSI storage successfully. ([OCPBUGS-23274](#))
- Previously, you could not assign an egress IP to the egress node on an Azure private cluster. With this release, egress IP is enabled for Azure private clusters that use outbound rules to achieve outbound connectivity. ([OCPBUGS-22949](#))
- Previously, there was no suitable virtual media device for Cisco UCS Blade. With this release, you can use Redfish virtual media to provision Cisco UCS hardware. ([OCPBUGS-19064](#))

1.9.45.2. Updating

To update an existing OpenShift Container Platform 4.12 cluster to this latest release, see [Updating a cluster using the CLI](#).

1.9.46. RHSA-2023:7823 - OpenShift Container Platform 4.12.46 bug fix and security update

Issued: 4 January 2024

OpenShift Container Platform release 4.12.46, which includes security updates, is now available. The list of bug fixes that are included in the update is documented in the [RHSA-2023:7823](#) advisory. The RPM packages that are included in the update are provided by the [RHBA-2023:7825](#) advisory.

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

```
$ oc adm release info 4.12.46 --pullspecs
```

1.9.46.1. Bug fixes

- Previously, the Image Registry Operator made API calls to the Storage Account List endpoint as part of obtaining access keys every 5 minutes. In projects with several OpenShift Container Platform clusters, this could lead to API rate limits being reached, which could result in several HTTP errors when attempting to create new clusters. With this release, the time between calls is increased from 5 minutes to 20 minutes. ([OCPBUGS-22125](#))

1.9.46.2. Updating

To update an existing OpenShift Container Platform 4.12 cluster to this latest release, see [Updating a cluster using the CLI](#).

1.9.47. RHSA-2024:0198 - OpenShift Container Platform 4.12.47 bug fix and security update

Issued: 17 January 2024

OpenShift Container Platform release 4.12.47, which includes security updates, is now available. The list of bug fixes that are included in the update is documented in the [RHSA-2024:0198](#) advisory. The RPM packages that are included in the update are provided by the [RHSA-2024:0200](#) advisory.

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

```
$ oc adm release info 4.12.47 --pullspecs
```

1.9.47.1. Bug fixes

- Previously, the **spec.storage.deviceClasses.thinPoolConfig.overprovisionRatio** value on a Logical Volume Manager Storage (LVMS) cluster custom resource could only be set to a minimum of **2**. With this release, the **spec.storage.deviceClasses.thinPoolConfig.overprovisionRatio** value can now be set to as low as **1**, which disables overprovisioning. ([OCPBUGS-24480](#))

1.9.47.2. Updating

To update an existing OpenShift Container Platform 4.12 cluster to this latest release, see [Updating a cluster using the CLI](#).

1.9.48. RHSA-2024:0485 - OpenShift Container Platform 4.12.48 bug fix and security update

Issued: 31 January 2024

OpenShift Container Platform release 4.12.48, which includes security updates, is now available. The list of bug fixes that are included in the update is documented in the [RHSA-2024:0485](#) advisory. The RPM packages that are included in the update are provided by the [RHSA-2024:0489](#) advisory.

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

```
$ oc adm release info 4.12.48 --pullspecs
```

1.9.48.1. Updating

To update an existing OpenShift Container Platform 4.12 cluster to this latest release, see [Updating a cluster using the CLI](#).

1.9.49. RHSA-2024:0664 - OpenShift Container Platform 4.12.49 bug fix and security update

Issued: 9 February 2024

OpenShift Container Platform release 4.12.49, which includes security updates, is now available. The list of bug fixes that are included in the update is documented in the [RHSA-2024:0664](#) advisory. The RPM packages that are included in the update are provided by the [RHSA-2024:0666](#) advisory.

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

```
$ oc adm release info 4.12.49 --pullspecs
```

1.9.49.1. Bug fixes

- Previously, pods assigned an IP from the pool created by the Whereabouts CNI plugin were getting stuck in **ContainerCreating** state after a node force reboot. With this release, the Whereabouts CNI plugin issue associated with the IP allocation after a node force reboot is resolved. ([OCPBUGS-16008](#))

1.9.49.2. Updating

To update an existing OpenShift Container Platform 4.12 cluster to this latest release, see [Updating a cluster using the CLI](#).

1.9.50. RHSA-2024:0833 - OpenShift Container Platform 4.12.50 bug fix and security update

Issued: 21 February 2024

OpenShift Container Platform release 4.12.50, which includes security updates, is now available. The list of bug fixes that are included in the update is documented in the [RHSA-2024:0833](#) advisory. The RPM packages that are included in the update are provided by the [RHBA-2024:0835](#) advisory.

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

```
-
```

```
$ oc adm release info 4.12.50 --pullspecs
```

1.9.50.1. Bug fixes

- Previously, CPU limits applied on the Amazon Elastic File System (EFS) Container Storage Interface (CSI) driver container caused performance degradation issues for I/O operations to EFS volumes. Now, the CPU limits for the EFS CSI driver are removed so the performance degradation issue no longer occurs. ([OCPBUGS-29066](#))

1.9.50.2. Updating

To update an existing OpenShift Container Platform 4.12 cluster to this latest release, see [Updating a cluster using the CLI](#).

1.9.51. RHSA-2024:1052 - OpenShift Container Platform 4.12.51 bug fix and security update

Issued: 6 March 2024

OpenShift Container Platform release 4.12.51, which includes security updates, is now available. The list of bug fixes that are included in the update is documented in the [RHSA-2024:1052](#) advisory. The RPM packages that are included in the update are provided by the [RHBA-2024:1054](#) advisory.

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

```
$ oc adm release info 4.12.51 --pullspecs
```

1.9.51.1. Bug fixes

- Previously, when the most recent and default channels were selectively mirrored, and a new release introduced a new channel, the current default channel became invalid. This caused the automatic assignment of the new default channel to fail. With this release, you can now define a **defaultChannel** field in the **ImageSetConfig** custom resource (CR) that overrides the **currentDefault** channel. ([OCPBUGS-29232](#))
- Previously, the **compat-openssl10** package was included in the Red Hat Enterprise Linux CoreOS (RHCOS). This package did not meet Common Vulnerabilities and Exposures (CVE) remediation requirements for Federal Risk and Authorization Management Program (FedRAMP). With this release, **compat-openssl10** has been removed from the RHCOS. As a result, security scanners will no longer identify potential common vulnerabilities and exposures (CVEs) in this package. Any binary running on the host RHCOS requiring Red Hat Enterprise Linux (RHEL) OpenSSL compatibility must be upgraded to support RHEL8 OpenSSL. ([OCPBUGS-22928](#))

1.9.51.2. Updating

To update an existing OpenShift Container Platform 4.12 cluster to this latest release, see [Updating a cluster using the CLI](#).

1.9.52. RHSA-2024:1265 - OpenShift Container Platform 4.12.53 bug fix update

Issued: 20 March 2024

OpenShift Container Platform release 4.12.53 is now available. The list of bug fixes that are included in the update is documented in the [RHSA-2024:1265](#) advisory. The RPM packages that are included in the update are provided by the [RHSA-2024:1267](#) advisory.

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

```
$ oc adm release info 4.12.53 --pullspecs
```

1.9.52.1. Updating

To update an existing OpenShift Container Platform 4.12 cluster to this latest release, see [Updating a cluster using the CLI](#).

1.9.53. RHSA-2024:1572 - OpenShift Container Platform 4.12.54 bug fix and security update

Issued: 3 April 2024

OpenShift Container Platform release 4.12.54, which includes security updates, is now available. The list of bug fixes that are included in the update is documented in the [RHSA-2024:1572](#) advisory. The RPM packages that are included in the update are provided by the [RHSA-2024:1574](#) advisory.

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

```
$ oc adm release info 4.12.54 --pullspecs
```

1.9.53.1. Updating

To update an existing OpenShift Container Platform 4.12 cluster to this latest release, see [Updating a cluster using the CLI](#).

1.9.54. RHSA-2024:1679 - OpenShift Container Platform 4.12.55 bug fix and security update

Issued: 8 April 2024

OpenShift Container Platform release 4.12.55, which includes security updates, is now available. The list of bug fixes that are included in the update is documented in the [RHSA-2024:1679](#) advisory. There are no RPM packages for this update.

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

```
$ oc adm release info 4.12.55 --pullspecs
```

1.9.54.1. Bug fixes

- Previously, the **manila-csi-driver-controller-metrics** service had empty endpoints due to an incorrect name for the app selector. With this release the app selector name is changed to **openstack-manila-csi** and the issue is fixed. ([OCPBUGS-30295](#))

1.9.54.2. Updating

To update an existing OpenShift Container Platform 4.12 cluster to this latest release, see [Updating a cluster using the CLI](#).

1.9.55. RHSA-2024:1896 – OpenShift Container Platform 4.12.56 bug fix and security update

Issued: 25 April 2024

OpenShift Container Platform release 4.12.56, which includes security updates, is now available. The list of bug fixes that are included in the update is documented in the [RHSA-2024:1896](#) advisory. The RPM packages that are included in the update are provided by the [RHSA-2024:1899](#) advisory.

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

```
$ oc adm release info 4.12.56 --pullspecs
```

1.9.55.1. Bug fixes

- Previously, two components (**tuned** and **irqbalanced**) were modifying the **irq** CPU affinity simultaneously, which caused issues. With this release, the **irqbalanced** component is the only component that configures the interrupt affinity and the issues are resolved. ([OCPBUGS-32205](#))

1.9.55.2. Updating

To update an existing OpenShift Container Platform 4.12 cluster to this latest release, see [Updating a cluster using the CLI](#).

1.9.56. RHSA-2024:2782 – OpenShift Container Platform 4.12.57 bug fix and security update

Issued: 16 May 2024

OpenShift Container Platform release 4.12.57, which includes security updates, is now available. The list of bug fixes that are included in the update is documented in the [RHSA-2024:2782](#) advisory. The RPM packages that are included in the update are provided by the [RHSA-2024:2784](#) advisory.

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

```
$ oc adm release info 4.12.57 --pullspecs
```

1.9.56.1. Bug fixes

- Previously, the **clear-irqbalance-banned-cpus.sh** script set an empty value for **IRQBALANCE_BANNED_CPUS** in the **/etc/sysconfig/irqbalance** pod annotation. As a result, the IRQs only balanced over the reserved CPUs. With this release, the **clear-irqbalance-banned-cpus.sh** script sets the banned mask to zeros on startup and the issue has been resolved. ([OCPBUGS-31442](#))
- Previously, a kernel regression introduced in OpenShift Container Platform versions 4.15.0, 4.14.14, 4.13.36, and 4.12.54 led to potential kernel panics in nodes and mounted CephFS volumes. In this release, the regression issue is fixed so that the kernel regression issue no longer occurs. ([OCPBUGS-33253](#))

1.9.56.2. Updating

To update an existing OpenShift Container Platform 4.12 cluster to this latest release, see [Updating a cluster using the CLI](#).

1.9.57. RHSA-2024:3349 - OpenShift Container Platform 4.12.58 bug fix and security update

Issued: 30 May 2024

OpenShift Container Platform release 4.12.58, which includes security updates, is now available. The list of bug fixes that are included in the update is documented in the [RHSA-2024:3349](#) advisory. The RPM packages that are included in the update are provided by the [RHSA-2024:3351](#) advisory.

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

```
$ oc adm release info 4.12.58 --pullspecs
```

1.9.57.1. Bug fixes

- Previously, some container processes created by using the **exec** command persisted even when CRI-O stopped the container. Consequently, lingering processes led to tracking issues, causing process leaks and defunct statuses. With this release, CRI-O tracks the **exec** calls processed for a container and ensures that the processes created as part of the **exec** calls are terminated when the container is stopped. ([OCPBUGS-33175](#))
- Previously, an issue with NodePort traffic-forwarding caused the Transmission Control Protocol (TCP) traffic to be directed to pods under a terminating state. With this release, the endpoints selection logic fully implements **KEP-1669 ProxyTerminatingEndpoints** and the issue has been resolved. ([OCPBUGS-33422](#))
- Previously, the load balancing algorithm did not differentiate between active and inactive services when determining weights, and it employed the **random** algorithm excessively in environments with many inactive services or environments routing backends with weight 0. This led to increased memory usage and a higher risk of excessive memory consumption. With this release, changes are made to optimize traffic direction towards active services only and prevent unnecessary use of the **random** algorithm with higher weights, reducing the potential for excessive memory consumption. ([OCPBUGS-33517](#))
- Previously, the load-balancing algorithm had flaws that led to increased memory usage and a higher risk of excessive memory consumption. With this release, the service filtering logic for load-balancing is updated and the issue has been resolved. ([OCPBUGS-33778](#))

1.9.57.2. Known issue

- Sometimes, the Console Operator status can get stuck in a failed state when the Operator is removed. To work around this issue, patch the Console controller to reset all conditions to the default state when the Console Operator is removed. For example, log in to the cluster and run the following command:

```
$ oc patch console.operator.openshift.io/cluster --type='merge' -p='{"status": {"conditions": null}}'
```

([OCPBUGS-33386](#))

1.9.57.3. Updating

To update an existing OpenShift Container Platform 4.12 cluster to this latest release, see [Updating a cluster using the CLI](#).

1.9.58. RHSA-2024:3713 - OpenShift Container Platform 4.12.59 bug fix and security update

Issued: 12 June 2024

OpenShift Container Platform release 4.12.59, which includes security updates, is now available. The list of bug fixes that are included in the update is documented in the [RHSA-2024:3713](#) advisory. The RPM packages that are included in the update are provided by the [RHSA-2024:3715](#) advisory.

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

```
$ oc adm release info 4.12.59 --pullspecs
```

1.9.58.1. Bug fixes

- Previously, the Ingress Operator would specify the **spec.template.spec.hostNetwork: true** on a router deployment without specifying the **spec.template.spec.containers[].ports[].hostPort**. This caused the API server to set a default value for each filed port's **hostPort**, which the Ingress Operator would then detect as an external update and attempt to revert it. Now, the Ingress Operator no longer incorrectly performs these updates. ([OCPBUGS-34888](#))
- Previously, the Ingress Operator was leaving the **spec.internalTrafficPolicy**, **spec.ipFamilies**, and **spec.ipFamilyPolicy** fields unspecified for **NodePort** and **ClusterIP** type services. The API would then set default values for these fields, which the Ingress Operator would try to revert. With this update, the Ingress Operator specifies an initial value and fixes the error caused by API default values. ([OCPBUGS-34757](#))
- Previously, if you configured an OpenShift Container Platform cluster with a high number of internal services or user-managed load balancer IP addresses, you experienced a delayed startup time for the OVN-Kubernetes service. This delay occurred when the OVN-Kubernetes service attempted to install **iptables** rules on a node. With this release, the OVN-Kubernetes service can process a large number of services in a few seconds. Additionally, you can access a new log to view the status of installing **iptables** rules on a node. ([OCPBUGS-34273](#))
- Previously, the Ingress Operator did not specify **ingressClass.spec.parameters.scope**, while the Ingress Class API object specifies type **cluster** by default. This caused unnecessary updates to all Ingress Classes when the Operator starts. With this update, the Ingress Operator specifies **ingressClass.spec.parameters.scope** of type **cluster**. ([OCPBUGS-34110](#))

1.9.58.2. Updating

To update an existing OpenShift Container Platform 4.12 cluster to this latest release, see [Updating a cluster using the CLI](#).

1.9.59. RHSA-2024:4006 - OpenShift Container Platform 4.12.60 bug fix and security update

Issued: 27 June 2024

OpenShift Container Platform release 4.12.60, which includes security updates, is now available. The list of bug fixes that are included in the update is documented in the [RHSA-2024:4006](#) advisory. The RPM packages that are included in the update are provided by the [RHSA-2024:4008](#) advisory.

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

```
$ oc adm release info 4.12.60 --pullspecs
```

1.9.59.1. Bug fixes

- Previously, for the **cluster:capacity_cpu_cores:sum** metric, nodes with the `infra`` role but not **master** role were not assigned a value of **infra** for the **label_node_role_kubernetes_io** label. With this update, nodes with the **infra** role but not **master** role are now correctly labeled as **infra** for this metric. ([OCBUGS-35558](#))
- Previously, when an Ingress Controller was configured with client SSL/TLS, but did not have the **clientca-configmap** finalizer, the Ingress Operator would try to add the finalizer without checking whether the Ingress Controller was marked for deletion. Consequently, if an Ingress Controller was configured with client SSL/TLS and was subsequently deleted, the Operator would correctly remove the finalizer. It would then repeatedly try and fail to update the IngressController to add the finalizer back, resulting in error messages in the Operator's logs. With this update, the Ingress Operator does not add the **clientca-configmap** finalizer to an Ingress Controller that is marked for deletion. As a result, the Ingress Operator does not attempt incorrect updates and does not log the associated errors. ([OCBUGS-35027](#))
- Previously, timeout values larger than what Golang could parse were not properly validated. Consequently, timeout values larger than what HAProxy could parse caused issues with HAProxy. With this update, if the timeout specifies a value larger than what can be parsed, the value is capped at the maximum value that HAProxy can parse. As a result, issues are no longer caused by HAProxy. ([OCBUGS-33432](#))

1.9.59.2. Updating

To update an existing OpenShift Container Platform 4.12 cluster to this latest release, see [Updating a cluster using the CLI](#).

1.9.60. RHSA-2024:4677 - OpenShift Container Platform 4.12.61 bug fix and security update

Issued: 25 July 2024

OpenShift Container Platform release 4.12.61, which includes security updates, is now available. The list of bug fixes that are included in the update is documented in the [RHSA-2024:4677](#) advisory. The RPM packages that are included in the update are provided by the [RHBA-2024:4679](#) advisory.

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

```
$ oc adm release info 4.12.61 --pullspecs
```

1.9.60.1. Bug fixes

- Previously, the **mapi_instance_create_failed** alert metric did not fire when there was an error for the Accelerated Networking feature on Microsoft Azure clusters. This release adds the

missing alert so that clusters with Accelerated Networking enabled generate alerts when required. ([OCBUGS-5235](#))

- Previously, the **wait-for-ceo** command that is used during the bootstrap operation to verify etcd rollout did not report errors for some failure modes. With this release, a code fix ensures that error messages for these failure modes get reported. ([OCBUGS-35501](#))
- Previously, the Helm Plugin index view did not display the same number of charts as the Helm CLI if the chart names varied. With this release, the Helm catalog now looks for **charts.openshift.io/name** and **charts.openshift.io/provider** so that all versions are grouped together in a single catalog title. ([OCBUGS-34933](#))

1.9.60.2. Enhancements

- Previously, the installation program failed to install a cluster on IBM Cloud VPC on the "eu-es" region, although it is supported. With this update, the installation program successfully installs a cluster on IBM Cloud VPC on the "eu-es" region. ([OCBUGS-22981](#))

1.9.60.3. Updating

To update an existing OpenShift Container Platform 4.12 cluster to this latest release, see [Updating a cluster using the CLI](#).

1.9.61. RHSA-2024:5200 - OpenShift Container Platform 4.12.63 bug fix and security update

Issued: 19 August 2024

OpenShift Container Platform release 4.12.63, which includes security updates, is now available. The list of bug fixes that are included in the update is documented in the [RHSA-2024:5200](#) advisory. The RPM packages that are included in the update are provided by the [RHSA-2024:5202](#) advisory.

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

```
$ oc adm release info 4.12.63 --pullspecs
```

1.9.61.1. Updating

To update an existing OpenShift Container Platform 4.12 cluster to this latest release, see [Updating a cluster using the CLI](#).

1.9.62. RHSA-2024:5808 - OpenShift Container Platform 4.12.64 bug fix and security update

Issued: 29 August 2024

OpenShift Container Platform release 4.12.64, which includes security updates, is now available. The list of bug fixes that are included in the update is documented in the [RHSA-2024:5808](#) advisory. The RPM packages that are included in the update are provided by the [RHSA-2024:5810](#) advisory.

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

```
$ oc adm release info 4.12.64 --pullspecs
```

1.9.62.1. Updating

To update an existing OpenShift Container Platform 4.12 cluster to this latest release, see [Updating a cluster using the CLI](#).

1.9.63. RHSA-2024:6642 - OpenShift Container Platform 4.12.65 bug fix and security update

Issued: 18 September 2024

OpenShift Container Platform release 4.12.65, which includes security updates, is now available. The list of bug fixes that are included in the update is documented in the [RHSA-2024:6642](#) advisory. The RPM packages that are included in the update are provided by the [RHBA-2024:6644](#) advisory.

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

```
$ oc adm release info 4.12.65 --pullspecs
```

1.9.63.1. Bug fixes

- Previously, when you received an **OVNKubernetesNorthdInactive** alert, you could not view the associated runbook. With this release, the runbook is added so you can reference it to resolve an **OVNKubernetesNorthdInactive** alert. ([OCPBUGS-38905](#))

1.9.63.2. Updating

To update an existing OpenShift Container Platform 4.12 cluster to this latest release, see [Updating a cluster using the CLI](#).

1.9.64. RHSA-2024:6705 - OpenShift Container Platform 4.12.66 bug fix and security update

Issued: 19 September 2024

OpenShift Container Platform release 4.12.66, which includes security updates, is now available. The list of bug fixes that are included in the update is documented in the [RHSA-2024:6705](#) advisory. There are no RPM packages for this release.

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

```
$ oc adm release info 4.12.66 --pullspecs
```

1.9.64.1. Updating

To update an existing OpenShift Container Platform 4.12 cluster to this latest release, see [Updating a cluster using the CLI](#).

1.9.65. RHSA-2024:7590 - OpenShift Container Platform 4.12.67 bug fix and security update

Issued: 09 October 2024

OpenShift Container Platform release 4.12.67, which includes security updates, is now available. The list of bug fixes that are included in the update is documented in the [RHSA-2024:7590](#) advisory. The RPM packages that are included in the update are provided by the [RHBA-2024:7592](#) advisory.

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

```
$ oc adm release info 4.12.67 --pullspecs
```

1.9.65.1. Bug fixes

- Previously, when the Operator Lifecycle Manager (OLM) evaluated a potential upgrade, the Operator used the dynamic client list for all custom resource (CR) instances in the cluster. Clusters with a large number of CRs could experience timeouts from the **apiserver** and stranded upgrades. With this release, the issue is resolved. ([OCBUGS-42161](#))
- Previously, the proxy service for the web console plugin handled non-200 response codes as error responses. This, in turn, caused browser caching issues. With this release, the proxy service is fixed so that it does not handle these responses as errors. ([OCBUGS-41600](#))

1.9.65.2. Updating

To update an existing OpenShift Container Platform 4.12 cluster to this latest release, see [Updating a cluster using the CLI](#).

1.9.66. RHSA-2024:8692 - OpenShift Container Platform 4.12.68 bug fix and security update

Issued: 07 November 2024

OpenShift Container Platform release 4.12.68, which includes security updates, is now available. The list of bug fixes that are included in the update is documented in the [RHSA-2024:8692](#) advisory. The RPM packages that are included in the update are provided by the [RHSA-2024:8694](#) advisory.

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

```
$ oc adm release info 4.12.68 --pullspecs
```

1.9.66.1. Bug fixes

- Previously, an upstream change in Google Cloud Platform (GCP) caused the Cloud Credential Operator (CCO) to degrade. With this release, the CCO is no longer degraded, and the issue is resolved. ([OCBUGS-43872](#))
- Previously, when updating to OpenShift Container Platform version 4.14, HAProxy 2.6 enforced strict RFC 7230 compliance and rejected requests with multiple **Transfer-Encoding** headers. Duplicate **Transfer-Encoding** headers were configured at the application level, so the requests resulted in **502 Bad Gateway** errors and service disruptions. With this release, cluster administrators can use a procedure to proactively detect applications that would send duplicate **Transfer-Encoding** headers before updating their clusters. This allows administrators to mitigate the issue in advance and prevents service disruption. ([OCBUGS-43703](#))
- Previously, a group ID was not added to the **/etc/group** within the container when the **spec.securityContext.runAsGroup** attribute was set in the pod resource. With this release, this issue is resolved. ([OCBUGS-41248](#))

- Previously, stale data prevented the node of an updated, OVN-enabled cluster from rejoining the cluster and returning to the **Ready** state. This fix removes problematic stale data from older versions of OpenShift Container Platform and resolves the issue. ([OCPBUGS-38382](#))
- Previously, when opening a targeted link in a new tab, an error displayed the following message: **Cannot read properties of undefined**. This issue was the result of a missing object validation check. With this release, the object validation check is added, and the new tab of the target detail displays correctly. ([OCPBUGS-33519](#))

1.9.66.2. Updating

To update an existing OpenShift Container Platform 4.12 cluster to this latest release, see [Updating a cluster using the CLI](#).

1.9.67. RHBA-2024:8996 - OpenShift Container Platform 4.12.69 bug fix

Issued: 14 November 2024

OpenShift Container Platform release 4.12.69 is now available. The list of bug fixes that are included in the update is documented in the [RHBA-2024:8996](#) advisory. The RPM packages that are included in the update are provided by the [RHBA-2024:8998](#) advisory.

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

```
$ oc adm release info 4.12.69 --pullspecs
```

1.9.67.1. Updating

To update an existing OpenShift Container Platform 4.12 cluster to this latest release, see [Updating a cluster using the CLI](#).

1.9.68. RHBA-2024:10533 - OpenShift Container Platform 4.12.70 bug fix and security update

Issued: 05 December 2024

OpenShift Container Platform release 4.12.70 is now available. The list of bug fixes that are included in the update is documented in the [RHBA-2024:10533](#) advisory. The RPM packages that are included in the update are provided by the [RHSA-2024:10535](#) advisory.

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

```
$ oc adm release info 4.12.70 --pullspecs
```

1.9.68.1. Updating

To update an existing OpenShift Container Platform 4.12 cluster to this latest release, see [Updating a cluster using the CLI](#).

1.9.69. RHSA-2025:0014 - OpenShift Container Platform 4.12.71 bug fix and security update

Issued: 09 January 2025

OpenShift Container Platform release 4.12.71 is now available. The list of bug fixes that are included in the update is documented in the [RHSA-2025:0014](#) advisory. The RPM packages that are included in the update are provided by the [RHBA-2025:0016](#) advisory.

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

```
$ oc adm release info 4.12.71 --pullspecs
```

1.9.69.1. Updating

To update an existing OpenShift Container Platform 4.12 cluster to this latest release, see [Updating a cluster using the CLI](#).

1.9.70. RHSA-2025:0832 - OpenShift Container Platform 4.12.72 bug fix and security update

Issued: 06 February 2025

OpenShift Container Platform release 4.12.72 is now available. The list of bug fixes that are included in the update is documented in the [RHSA-2025:0832](#) advisory. The RPM packages that are included in the update are provided by the [RHSA-2025:0834](#) advisory.

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

```
$ oc adm release info 4.12.72 --pullspecs
```

1.9.70.1. Bug fixes

- Previously, during node reboots, especially during update operations, the node that interacted with the rebooting machine entered a **Ready=Unknown** state. This caused the Control Plane Machine Set Operator to enter an **UnavailableReplicas** condition and an **Available=false** state, which triggered alerts that demanded urgent action. However, manual intervention was not necessary because the condition was resolved when the node rebooted. With this release, a grace period for node unreadiness is provided. If a node enters an unready state, the Control Plane Machine Set Operator does not instantly enter an **UnavailableReplicas** condition or an **Available=false** state. ([OCPBUGS-48325](#)).

1.9.70.2. Updating

To update an existing OpenShift Container Platform 4.12 cluster to this latest release, see [Updating a cluster using the CLI](#).

1.9.71. RHSA-2025:1242 - OpenShift Container Platform 4.12.73 bug fix and security update

Issued: 13 February 2025

OpenShift Container Platform release 4.12.73 is now available. The list of bug fixes that are included in the update is documented in the [RHSA-2025:1242](#) advisory. The RPM packages that are included in the update are provided by the [RHBA-2025:1244](#) advisory.

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

```
$ oc adm release info 4.12.73 --pullspecs
```

1.9.71.1. Bug fixes

- Previously, pods could get stuck when generating **FailedMount** errors. These errors caused nodes to need additional reboots and for the Network File System (NFS) volume mount to stay in a pending state. With this release, a kernel update fixes the issue so that pods no longer get stuck because nodes no longer need to be rebooted to clear the **FailedMount** errors. ([OCBUGS-49398](#))

1.9.71.2. Updating

To update an existing OpenShift Container Platform 4.12 cluster to this latest release, see [Updating a cluster using the CLI](#).

1.9.72. RHSA-2025:2441 - OpenShift Container Platform 4.12.74 bug fix and security update

Issued: 13 March 2025

OpenShift Container Platform release 4.12.74 is now available. The list of bug fixes that are included in the update is documented in the [RHSA-2025:2441](#) advisory. The RPM packages that are included in the update are provided by the [RHSA-2025:2443](#) advisory.

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

```
$ oc adm release info 4.12.74 --pullspecs
```

1.9.72.1. Updating

To update an existing OpenShift Container Platform 4.12 cluster to this latest release, see [Updating a cluster using the CLI](#).

1.9.73. RHSA-2025:3573 - OpenShift Container Platform 4.12.75 bug fix and security update

Issued: 10 April 2025

OpenShift Container Platform release 4.12.75 is now available. The list of bug fixes that are included in the update is documented in the [RHSA-2025:3573](#) advisory. The RPM packages that are included in the update are provided by the [RHBA-2025:3575](#) advisory.

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

```
$ oc adm release info 4.12.75 --pullspecs
```

1.9.73.1. Updating

To update an existing OpenShift Container Platform 4.12 cluster to this latest release, see [Updating a cluster using the CLI](#).

1.9.74. RHSA-2025:4409 - OpenShift Container Platform 4.12.76 bug fix and security update

Issued: 08 May 2025

OpenShift Container Platform release 4.12.76 is now available. The list of bug fixes that are included in the update is documented in the [RHSA-2025:4409](#) advisory. The RPM packages that are included in the update are provided by the [RHBA-2025:4411](#) advisory.

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

```
$ oc adm release info 4.12.76 --pullspecs
```

1.9.74.1. Bug fixes

- Previously, the **Observe** sections on the web console did not show items contributed from plugins unless certain flags related to the monitoring were set. These flags prevented other plugins, such as logging, Distributed Tracing Platform, and network observability, 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-55210](#))

1.9.74.2. Updating

To update an existing OpenShift Container Platform 4.12 cluster to this latest release, see [Updating a cluster using the CLI](#).

1.9.75. RHSA-2025:8303 - OpenShift Container Platform 4.12.77 bug fix and security update

Issued: 05 June 2025

OpenShift Container Platform release 4.12.77 is now available. The list of bug fixes that are included in the update is documented in the [RHSA-2025:8303](#) advisory. The RPM packages that are included in the update are provided by the [RHBA-2025:8304](#) advisory.

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

```
$ oc adm release info 4.12.77 --pullspecs
```

1.9.75.1. Bug fixes

- Previously, when the **openshift-host-network' namespace was modified by the user or by an upgrade, the network policy did not correctly set the VXLAN virtual network ID parameter, `VNID, to 0**. With this release, the VNID parameter is correctly set after the namespace is modified. ([OCPBUGS-56130](#))
- 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](#))

1.9.75.2. Updating

To update an existing OpenShift Container Platform 4.12 cluster to this latest release, see [Updating a cluster using the CLI](#).

1.9.76. RHSA-2025:10270 - OpenShift Container Platform 4.12.78 bug fix and security update

Issued: 10 July 2025

OpenShift Container Platform release 4.12.78 is now available. The list of bug fixes that are included in the update is documented in the [RHSA-2025:10270](#) advisory. The RPM packages that are included in the update are provided by the [RHSA-2025:10271](#) advisory.

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

```
$ oc adm release info 4.12.78 --pullspecs
```

1.9.76.1. Bug fixes

- Previously, when you injected proxy environment variables in a build container, they replaced the default proxy settings and caused build failures due to external binary compatibility issues. Also, if no default variable was set, a null value was replaced, which caused applications to break. With this release, a check for a null proxy environment variable is added to the process and the injected proxy variables are added only if they are defined in the default variables and if the variable is not equal to a null value. ([OCPBUGS-57171](#))

1.9.76.2. Updating

To update an existing OpenShift Container Platform 4.12 cluster to this latest release, see [Updating a cluster using the CLI](#).

1.9.77. RHSA-2025:12323 - OpenShift Container Platform 4.12.79 bug fix and security update

Issued: 07 August 2025

OpenShift Container Platform release 4.12.79 is now available. The list of bug fixes that are included in the update is documented in the [RHSA-2025:12323](#) advisory. The RPM packages that are included in the update are provided by the [RHBA-2025:12324](#) 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 4.12.79 --pullspecs
```

1.9.77.1. Updating

To update an existing OpenShift Container Platform 4.12 cluster to this latest release, see [Updating a cluster using the CLI](#).

1.9.78. RHSA-2025:15332 - OpenShift Container Platform 4.12.80 bug fix and security update

Issued: 11 September 2025

OpenShift Container Platform release 4.12.80 is now available. The list of bug fixes that are included in the update is documented in the [RHSA-2025:15332](#) advisory. The RPM packages that are included in the update are provided by the [RHBA-2025:15307](#) 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 4.12.80 --pullspecs
```

1.9.78.1. Updating

To update an existing OpenShift Container Platform 4.12 cluster to this latest release, see [Updating a cluster using the CLI](#).