



OpenShift Container Platform 4.14

Release notes

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

OpenShift Container Platform 4.14 Release notes

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

Legal Notice

Copyright © Red Hat.

Except as otherwise noted below, the text of and illustrations in this documentation are licensed by Red Hat under the Creative Commons Attribution–Share Alike 3.0 Unported license . If you distribute this document or an adaptation of it, you must provide the URL for the original version.

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

Red Hat, the Red Hat logo, JBoss, Hibernate, and RHCE are trademarks or registered trademarks of Red Hat, LLC. or its subsidiaries in the United States and other countries.

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

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

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

All other trademarks are the property of their respective owners.

Abstract

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

Table of Contents

CHAPTER 1. OPENSIFT CONTAINER PLATFORM 4.14 RELEASE NOTES	11
1.1. ABOUT THIS RELEASE	11
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. RHCOS now uses RHEL 9.2	12
1.3.1.1.1. Considerations for upgrading to OpenShift Container Platform with RHEL 9.2	12
1.3.2. Installation and update	13
1.3.2.1. Installing a cluster on Amazon Web Services (AWS) by using a shared VPC	13
1.3.2.2. Enabling S3 bucket to be retained during cluster bootstrap on AWS	13
1.3.2.3. Installing a cluster on Microsoft Azure using a NAT gateway (Technology Preview)	13
1.3.2.4. Installing a cluster on Google Cloud Platform (GCP) using pd-balanced disk type	13
1.3.2.5. Optional capabilities in OpenShift Container Platform 4.14	13
1.3.2.6. Installing a cluster with Azure AD Workload Identity	13
1.3.2.7. User-defined tags for Microsoft Azure now generally available	13
1.3.2.8. Confidential VMs for Azure (Technology Preview)	13
1.3.2.9. Trusted launch for Azure (Technology Preview)	14
1.3.2.10. User-defined labels and tags for Google Cloud Platform (Technology Preview)	14
1.3.2.11. Installing an OpenShift Container Platform cluster on Microsoft Azure in a restricted network	14
1.3.2.12. Specifying installation disks using a by-path device alias	14
1.3.2.13. Applying existing AWS security groups to a cluster	14
1.3.2.14. Required administrator acknowledgment when updating from OpenShift Container Platform 4.13 to 4.14	14
1.3.2.15. Three-node cluster support for Nutanix	15
1.3.2.16. Installing a cluster on GCP using Confidential VMs is generally available	15
1.3.2.17. Root volume types parameter for RHOSP is now available	15
1.3.2.18. Static IP addresses for vSphere nodes	15
1.3.2.19. Additional validation for the Bare Metal Host CR	15
1.3.2.20. Install a cluster quickly in AWS Local Zones	16
1.3.2.21. Simplified installation and update experience for clusters with manually maintained cloud credentials	16
1.3.2.22. Quickly install RHCOS on vSphere hosts by using a pre-existing RHCOS image template	16
1.3.2.23. OpenShift Container Platform on 64-bit ARM	17
1.3.2.24. Using a custom RHCOS image for a Microsoft Azure cluster	17
1.3.2.25. Installing single-node OpenShift on cloud providers	17
1.3.3. Post-installation configuration	17
1.3.3.1. OpenShift Container Platform cluster with multi-architecture compute machines	18
1.3.4. Web console	18
1.3.4.1. Administrator Perspective	18
1.3.4.1.1. Dynamic plugin enhancements	18
1.3.4.1.2. Operating system based filtering in OperatorHub	18
1.3.4.1.3. Support for installing specific Operator versions in the web console	18
1.3.4.1.4. OperatorHub support for AWS STS	18
1.3.4.2. Developer Perspective	19
1.3.4.2.1. New quick starts	19
1.3.4.2.2. OpenShift Pipelines page improvements	19
1.3.5. OpenShift CLI (oc)	19
1.3.5.1. Supporting multi-arch OCI local images for catalogs with oc-mirror	19
1.3.5.2. Logging in to the CLI using a web browser	19
1.3.5.3. Enhancement to oc new-build	20

1.3.5.4. Enhancement to oc new-app	20
1.3.6. IBM Z and IBM LinuxONE	20
1.3.6.1. IBM Z and IBM LinuxONE notable enhancements	20
1.3.6.2. IBM Secure Execution	21
1.3.7. IBM Power	21
1.3.7.1. IBM Power notable enhancements	21
1.3.8. IBM Power, IBM Z, and IBM LinuxONE support matrix	21
1.3.9. Authentication and authorization	25
1.3.9.1. SCC preemption prevention	25
1.3.9.2. Pod security admission privileged namespaces	25
1.3.9.3. Pod security admission synchronization disabled on modified namespaces	25
1.3.9.4. OLM-based Operator support for AWS STS	26
1.3.9.5. Authentication Operator honors noProxy during connection checks	26
1.3.10. Networking	26
1.3.10.1. IPv6 as primary IP address family on vSphere dual-stack clusters	26
1.3.10.2. Multiple external gateway support for the OVN-Kubernetes network plugin	26
1.3.10.3. Ingress Node Firewall Operator is generally available	26
1.3.10.4. Dynamic use of non-reserved CPUs for OVS	26
1.3.10.5. Dual-stack configuration for multiple IP addresses	27
1.3.10.6. Exclude SR-IOV network topology for NUMA-aware scheduling	27
1.3.10.7. Update to HAProxy 2.6	27
1.3.10.8. Support for configuring the maximum length with sidecar logging in the Ingress Controller	27
1.3.10.9. NMstate Operator updated in console	27
1.3.10.10. OVN-Kubernetes network plugin support for IPsec on IBM Cloud	28
1.3.10.11. OVN-Kubernetes network plugin support for IPsec encryption of external traffic (Technology Preview)	28
1.3.10.12. Single-stack IPv6 support for Kubernetes NMstate	28
1.3.10.13. Egress service resource to manage egress traffic for pods behind a load balancer (Technology Preview)	28
1.3.10.14. VRF specification in MetalLB's BGPPeer resource (Technology Preview)	28
1.3.10.15. VRF specification in NMState's NodeNetworkConfigurationPolicy resource (Technology Preview)	28
1.3.10.16. Support for Broadcom BCM57504 is now GA	29
1.3.10.17. OVN-Kubernetes is available as a secondary network	29
1.3.10.18. Global IP forwarding is disabled on OVN-Kubernetes based cluster deployments	29
1.3.10.19. Admin Network Policy (Technology Preview)	29
1.3.10.20. MAC-VLAN, IP-VLAN, and VLAN subinterface creation for pods	29
1.3.10.21. Supporting all-multicast mode	29
1.3.10.22. Enhance network flexibility by using the TAP device plugin	30
1.3.10.23. Support for running rootless DPDK workloads with kernel access by using the TAP CNI plugin	30
1.3.10.24. Set or delete specific HTTP headers using an Ingress Controller or a Route object	30
1.3.10.25. Egress IPs on additional network interfaces is generally available	30
1.3.10.26. Parallel node draining during SR-IOV network policy updates	30
1.3.11. Registry	31
1.3.11.1. Optional Image Registry Operator	31
1.3.12. Storage	31
1.3.12.1. Support for OR logic in LVMS	31
1.3.12.2. Support for ext4 in LVMS	31
1.3.12.3. Standardized STS configuration workflow	31
1.3.12.4. Read Write Once Pod access mode (Technology Preview)	31
1.3.12.5. GCP Filestore storage CSI Driver Operator is generally available	31
1.3.12.6. Automatic CSI migration for VMware vSphere	32
1.3.12.7. Secrets Store CSI Driver Operator (Technology Preview)	32

1.3.12.8. Azure File supporting NFS is generally available	32
1.3.13. Oracle® Cloud Infrastructure	32
1.3.14. Operator lifecycle	32
1.3.14.1. Operator Lifecycle Manager (OLM) 1.0 (Technology Preview)	33
1.3.15. Operator development	33
1.3.15.1. Token authentication for Operators on cloud providers: AWS STS	34
1.3.15.2. Configuring Operator projects with support for multiple platforms	34
1.3.16. Builds	34
1.3.17. Machine Config Operator	34
1.3.17.1. Handling of registry certificate authorities	34
1.3.17.2. Additional metrics available in Prometheus	34
1.3.17.3. Support for offline Tang provisioning	34
1.3.17.4. Certificates are now handled by the Machine Config Daemon	35
1.3.18. Machine API	35
1.3.18.1. Support for control plane machine sets on Nutanix clusters	35
1.3.18.2. Support for control plane machine sets on RHOSP clusters	35
1.3.18.3. Support for assigning AWS machines to placement groups	35
1.3.18.4. Azure confidential VMs and trusted launch (Technology Preview)	36
1.3.19. Nodes	36
1.3.19.1. Descheduler resource limits for large clusters	36
1.3.19.2. Pod topology spread constraints matchLabelKeys parameter is now generally available	36
1.3.19.3. MaxUnavailableStatefulSet enabled (Technology Preview)	36
1.3.19.4. Pod disruption budget (PDB) unhealthy pod eviction policy	36
1.3.19.5. Linux Control Groups version 2 is now default	36
1.3.19.6. Cron job time zones general availability	37
1.3.20. Monitoring	37
1.3.20.1. Updates to monitoring stack components and dependencies	37
1.3.20.2. Changes to alerting rules	37
1.3.20.3. New option to create alerts based on core platform metrics	38
1.3.20.4. New option to specify resource limits for all monitoring components	38
1.3.20.5. New options to configure node-exporter collectors	38
1.3.20.6. New option to deploy monitoring web console plugin resources	39
1.3.21. Network Observability Operator	39
1.3.22. Scalability and performance	39
1.3.22.1. PAO must-gather image added to default must-gather image	39
1.3.22.2. Collecting data for the NUMA Resources Operator with the must-gather image of the Operator	39
1.3.22.3. Enabling more control over the C-states for each pod	40
1.3.22.4. Support for provisioning IPv6 spoke clusters from dual-stack hub clusters	40
1.3.22.5. Dual-stack networking for RHOSP clusters (Technology Preview)	40
1.3.22.6. Security group management for RHOSP clusters	40
1.3.22.7. Using custom CRs with PolicyGenTemplate CRs in the GitOps Zero Touch Provisioning (ZTP) pipeline	40
1.3.22.8. GitOps ZTP independence from managed cluster version	40
1.3.22.9. Pre-caching user-specified images with Topology Aware Lifecycle Manager	41
1.3.22.10. Disk cleaning option through SiteConfig and GitOps ZTP	41
1.3.22.11. Support for adding custom node labels in the SiteConfig CR through GitOps ZTP	41
1.3.22.12. Support for tuning etcd latency tolerances (Technology Preview)	41
1.3.22.13. Deprecation of fields in SiteConfig	41
1.3.23. Hosted control planes	41
1.3.23.1. Hosted control planes is Generally Available on bare metal and OpenShift Virtualization	41
1.3.23.2. Creating ARM NodePool objects on AWS hosted clusters (Technology Preview)	41
1.3.23.3. Hosted control planes on IBM Z (Technology Preview)	41
1.3.23.4. Hosted control planes on IBM Power (Technology Preview)	42

1.3.24. Insights Operator	42
1.3.24.1. On demand data gathering (Technology Preview)	42
1.3.24.2. Running gather operations as individual pods (Technology Preview)	42
1.4. NOTABLE TECHNICAL CHANGES	42
1.4.1. Cloud controller managers for additional cloud providers	42
1.4.2. Future restricted enforcement for pod security admission	42
1.4.3. Change in SSH key location	43
1.4.4. cert-manager Operator general availability	43
1.4.5. Improved scaling and stability with Open Virtual Network (OVN) Optimizations	43
1.4.6. Operator SDK 1.31.0	43
1.4.7. oc commands now default to storing and obtaining credentials from Podman configuration locations	44
1.4.8. Long-running pod requests are now recorded as CONNECT requests	44
1.4.9. Open Virtual Network Infrastructure Controller default range	44
1.5. DEPRECATED AND REMOVED FEATURES	44
1.5.1. Operator lifecycle and development deprecated and removed features	45
1.5.2. Images deprecated and removed features	45
1.5.3. Installation deprecated and removed features	45
1.5.4. Storage deprecated and removed features	46
1.5.5. Building applications deprecated and removed features	46
1.5.6. Multi-architecture deprecated and removed features	46
1.5.7. Networking deprecated and removed features	47
1.5.8. Node deprecated and removed features	47
1.5.9. OpenShift CLI (oc) deprecated and removed features	47
1.5.10. Workloads deprecated and removed features	47
1.5.11. Bare metal monitoring deprecated and removed features	48
1.5.12. Deprecated features	48
1.5.12.1. Deprecation of the OpenShift SDN network plugin	48
1.5.12.2. Service Binding Operator	48
1.5.12.3. DeploymentConfig resources are now deprecated	48
1.5.12.4. Operator-specific CatalogSource CRs used in GitOps ZTP are deprecated	48
1.5.12.5. The --cloud parameter for the oc adm release extract command	48
1.5.12.6. Red Hat Virtualization (RHV) as a host platform for OpenShift Container Platform	49
1.5.12.7. Using the REGISTRY_AUTH_PREFERENCE environment variable is now deprecated	49
1.5.12.8. operators.openshift.io/infrastructure-features annotations	49
1.5.13. Removed features	49
1.5.13.1. Bare Metal Event Relay Operator removed	49
1.5.13.2. Beta APIs removed from Kubernetes 1.27	49
1.5.13.3. Support for the LatencySensitive feature set is removed	50
1.5.13.4. oc registry login no longer stores credentials in Docker configuration file locations	50
1.6. BUG FIXES	50
1.6.1. API Server and Authentication	50
1.6.2. Bare Metal Hardware Provisioning	51
1.6.3. Cloud Compute	51
1.6.4. Cloud Credential Operator	52
1.6.5. Cluster Version Operator	52
1.6.6. Developer Console	53
1.6.7. etcd Cluster Operator	53
1.6.8. Installer	53
1.6.9. Machine Config Operator	57
1.6.10. Management Console	58
1.6.11. Monitoring	59
1.6.12. Networking	60
1.6.13. OpenShift CLI (oc)	62

1.6.14. Operator Lifecycle Manager (OLM)	63
1.6.15. OpenShift API server	64
1.6.16. Red Hat Enterprise Linux CoreOS (RHCOS)	64
1.6.17. Storage	64
1.7. TECHNOLOGY PREVIEW FEATURES	64
1.7.1. Networking Technology Preview features	65
1.7.2. Storage Technology Preview features	66
1.7.3. Installation Technology Preview features	67
1.7.4. Node Technology Preview features	68
1.7.5. Multi-Architecture Technology Preview features	68
1.7.6. Specialized hardware and driver enablement Technology Preview features	69
1.7.7. Scalability and performance Technology Preview features	69
1.7.8. Operator lifecycle and development Technology Preview features	70
1.7.9. Monitoring Technology Preview features	70
1.7.10. Hosted control plane Technology Preview features	71
1.7.11. Machine management Technology Preview features	71
1.7.12. Authentication and authorization Technology Preview features	72
1.7.13. Machine Config Operator Technology Preview features	72
1.8. KNOWN ISSUES	72
1.9. ASYNCHRONOUS ERRATA UPDATES	83
1.9.1. RHSA-2026:5107 - OpenShift Container Platform 4.14.63 image release, fixed issues, and security update	84
1.9.2. Enhancements	84
1.9.3. Fixed issues	84
1.9.3.1. Updating	85
1.9.4. RHSA-2026:2990 - OpenShift Container Platform 4.14.62 image release, fixed issues, and security update	85
1.9.5. Fixed issues	85
1.9.5.1. Updating	85
1.9.6. RHBA-2026:1003 - OpenShift Container Platform 4.14.61 bug fix advisory	85
1.9.6.1. Bug fixes	86
1.9.6.2. Updating	86
1.9.7. RHBA-2025:22733 - OpenShift Container Platform 4.14.60 bug fix and security update	86
1.9.7.1. Bug fixes	86
1.9.7.2. Updating	86
1.9.8. RHSA-2025:21331 - OpenShift Container Platform 4.14.59 bug fix and security update	86
1.9.8.1. Bug fixes	87
1.9.8.2. Updating	87
1.9.9. RHSA-2025:19058 - OpenShift Container Platform 4.14.58 bug fix and security update	87
1.9.9.1. Bug fixes	87
1.9.9.2. Updating	87
1.9.10. RHSA-2025:16165 - OpenShift Container Platform 4.14.57 bug fix and security update	87
1.9.10.1. Enhancement	88
1.9.10.2. Bug fixes	88
1.9.10.3. Updating	88
1.9.11. RHSA-2025:14855 - OpenShift Container Platform 4.14.56 bug fix and security update	88
1.9.11.1. Bug fixes	89
1.9.11.2. Updating	89
1.9.12. RHSA-2025:13289 - OpenShift Container Platform 4.14.55 bug fix and security update	89
1.9.12.1. Updating	89
1.9.13. RHSA-2025:11669 - OpenShift Container Platform 4.14.54 bug fix and security update	89
1.9.13.1. Bug fixes	89
1.9.13.2. Updating	90

1.9.14. RHSA-2025:9759 - OpenShift Container Platform 4.14.53 bug fix and security update	90
1.9.14.1. Bug fixes	90
1.9.14.2. Updating	91
1.9.15. RHSA-2025:7702 - OpenShift Container Platform 4.14.52 bug fix and security update	91
1.9.15.1. Bug fixes	91
1.9.15.2. Updating	92
1.9.16. RHSA-2025:4177 - OpenShift Container Platform 4.14.51 bug fix and security update	92
1.9.16.1. Bug fixes	92
1.9.16.2. Updating	93
1.9.17. RHSA-2025:3569 - OpenShift Container Platform 4.14.50 bug fix and security update	93
1.9.17.1. Bug fixes	93
1.9.17.2. Updating	93
1.9.18. RHSA-2025:2710 - OpenShift Container Platform 4.14.49 bug fix and security update	94
1.9.18.1. Bug fixes	94
1.9.18.2. Updating	94
1.9.19. RHSA-2025:1451 - OpenShift Container Platform 4.14.48 bug fix and security update	94
1.9.19.1. Bug fixes	94
1.9.19.2. Updating	95
1.9.20. RHSA-2025:0840 - OpenShift Container Platform 4.14.46 bug fix and security update	95
1.9.20.1. Bug fixes	95
1.9.20.2. Updating	95
1.9.21. RHSA-2025:0364 - OpenShift Container Platform 4.14.45 bug fix and security update	95
1.9.21.1. Bug fixes	96
1.9.21.2. Updating	96
1.9.22. RHSA-2025:0029 - OpenShift Container Platform 4.14.44 bug fix and security update	96
1.9.22.1. Bug fixes	97
1.9.22.2. Updating	97
1.9.23. RHSA-2024:11031 - OpenShift Container Platform 4.14.43 bug fix and security update	97
1.9.23.1. Bug fixes	97
1.9.23.2. Updating	98
1.9.24. RHSA-2024:10523 - OpenShift Container Platform 4.14.42 bug fix and security update	98
1.9.24.1. Bug fixes	98
1.9.24.2. Updating	98
1.9.25. RHSA-2024:9620 - OpenShift Container Platform 4.14.41 bug fix and security update	99
1.9.25.1. Bug fixes	99
1.9.25.2. Updating	99
1.9.26. RHSA-2024:8697 - OpenShift Container Platform 4.14.40 bug fix and security update	99
1.9.26.1. Bug fixes	99
1.9.26.2. Updating	100
1.9.27. RHSA-2024:8235 - OpenShift Container Platform 4.14.39 bug fix and security update	100
1.9.27.1. Enhancements	100
1.9.27.1.1. Using the Insights Operator to collect data from a Prometheus metric	100
1.9.27.2. Bug fixes	100
1.9.27.3. Updating	101
1.9.28. RHSA-2024:7184 - OpenShift Container Platform 4.14.38 bug fix and security update	101
1.9.28.1. Enhancements	101
1.9.28.1.1. Backporting configurable subnets	101
1.9.28.2. Bug fixes	102
1.9.28.3. Known issues	103
1.9.28.4. Updating	103
1.9.29. RHSA-2024:6689 - OpenShift Container Platform 4.14.37 bug fix and security update	103
1.9.29.1. Updating	103
1.9.30. RHSA-2024:6406 - OpenShift Container Platform 4.14.36 bug fix and security update	103

1.9.30.1. Enhancements	104
1.9.30.1.1. Updating the CENTOS 8 references to CENTOS 9	104
1.9.30.1.2. Collecting Ingress Controller certificates	104
1.9.30.2. Automatic leap seconds handling for PTP grandmaster clocks	104
1.9.30.3. Bug fixes	104
1.9.30.4. Updating	105
1.9.31. RHSA-2024:5433 - OpenShift Container Platform 4.14.35 bug fix and security update	105
1.9.31.1. Enhancements	105
1.9.31.1.1. Configuring Capacity Reservation by using machine sets	105
1.9.31.1.2. Updating to Kubernetes v1.27.16	105
1.9.31.2. Bug fixes	105
1.9.31.3. Updating	106
1.9.32. RHSA-2024:4960 - OpenShift Container Platform 4.14.34 bug fix and security update	106
1.9.32.1. Enhancements	106
1.9.32.1.1. Adding connectTimeout tuning option to Ingress Controller API	106
1.9.32.1.2. Collecting Prometheus and AlertManager resources with the Insights Operator	106
1.9.32.2. Bug fixes	107
1.9.32.3. Known issues	107
1.9.32.4. Updating	107
1.9.33. RHSA-2024:4479 - OpenShift Container Platform 4.14.33 bug fix and security update	108
1.9.33.1. Bug fixes	108
1.9.33.2. Updating	108
1.9.34. RHSA-2024:4329 - OpenShift Container Platform 4.14.32 bug fix and security updates	109
1.9.34.1. Enhancements	109
1.9.34.1.1. New custom resource definitions for pipeline plugin	109
1.9.34.1.2. Controller for defragmentating etcd	109
1.9.34.2. Bug fixes	109
1.9.34.3. Updating	110
1.9.35. RHSA-2024:4010 - OpenShift Container Platform 4.14.31 bug fix and security updates	111
1.9.35.1. Bug fixes	111
1.9.35.2. Updating	111
1.9.36. RHSA-2024:3881 - OpenShift Container Platform 4.14.30 bug fix and security updates	111
1.9.36.1. Enhancements	111
1.9.36.1.1. Allowing a colon character in pull secret passwords	111
1.9.36.1.2. Configuring the topology feature for the Cinder CSI Driver	111
1.9.36.2. Bug fixes	112
1.9.36.3. Updating	112
1.9.37. RHBA-2024:3697 - OpenShift Container Platform 4.14.29 bug fix and security updates	112
1.9.37.1. Bug fixes	113
1.9.37.2. Updating	113
1.9.38. RHSA-2024:3523 - OpenShift Container Platform 4.14.28 bug fix update and security update	113
1.9.38.1. Bug fixes	113
1.9.38.2. Updating	114
1.9.39. RHSA-2024:3331 - OpenShift Container Platform 4.14.27 bug fix update and security update	114
1.9.39.1. Bug fixes	114
1.9.39.2. Updating	115
1.9.40. RHSA-2024:2869 - OpenShift Container Platform 4.14.26 bug fix update and security update	115
1.9.40.1. Enhancements	115
1.9.40.1.1. OperatorHub filter renamed from FIPS Mode to Designed for FIPS	115
1.9.40.2. Bug fixes	115
1.9.40.3. Updating	116
1.9.41. RHBA-2024:2789 - OpenShift Container Platform 4.14.25 bug fix update	116
1.9.41.1. Bug fixes	116

1.9.41.2. Updating	117
1.9.42. RHSA-2024:2668 - OpenShift Container Platform 4.14.24 bug fix update and security update	117
1.9.42.1. Enhancements	117
1.9.42.1.1. IPv6 unsolicited neighbor advertisements now default on macvlan CNI plug-in	117
1.9.42.2. Bug fixes	117
1.9.42.3. Updating	118
1.9.43. RHBA-2024:2051 - OpenShift Container Platform 4.14.23 bug fix update and security update	118
1.9.43.1. Enhancements	118
1.9.43.1.1. Egress IP verification step for additional hops	118
1.9.43.1.2. New profile for RT kernel to drop unsupported parameters	118
1.9.43.1.3. Disable option for OLM default source	119
1.9.43.2. Bug fixes	119
1.9.43.3. Updating	119
1.9.44. RHSA-2024:1891 - OpenShift Container Platform 4.14.22 bug fix update and security update	119
1.9.44.1. Enhancements	119
1.9.44.1.1. Number of configured control plane replicas validated	120
1.9.44.2. Bug fixes	120
1.9.44.3. Updating	120
1.9.45. RHSA-2024:1765 - OpenShift Container Platform 4.14.21 bug fix update and security update	120
1.9.45.1. Bug fixes	120
1.9.45.2. Updating	120
1.9.46. RHSA-2024:1681 - OpenShift Container Platform 4.14.20 bug fix update and security update	121
1.9.46.1. Updating	121
1.9.47. RHBA-2024:1564 - OpenShift Container Platform 4.14.19 bug fix update and security update	121
1.9.47.1. Bug fixes	121
1.9.47.2. Updating	121
1.9.48. RHSA-2024:1458 - OpenShift Container Platform 4.14.18 bug fix update and security update	121
1.9.48.1. Bug fixes	122
1.9.48.2. Known issues	122
1.9.48.3. Updating	122
1.9.49. RHBA-2024:1260 - OpenShift Container Platform 4.14.17 bug fix update	122
1.9.49.1. Updating	122
1.9.50. RHBA-2024:1205 - OpenShift Container Platform 4.14.16 bug fix update	122
1.9.50.1. Updating	123
1.9.51. RHBA-2024:1046 - OpenShift Container Platform 4.14.15 bug fix update	123
1.9.51.1. Bug fixes	123
1.9.51.2. Updating	124
1.9.52. RHSA-2024:0941 - OpenShift Container Platform 4.14.14 bug fix and security update	124
1.9.52.1. Enhancements	124
1.9.52.1.1. Adding "eu-es" region support for IPI	124
1.9.52.2. Updating	124
1.9.53. RHSA-2024:0837 - OpenShift Container Platform 4.14.13 bug fix and security update	124
1.9.53.1. Bug fixes	125
1.9.53.2. Updating	125
1.9.54. RHSA-2024:0735 - OpenShift Container Platform 4.14.12 bug fix and security update	125
1.9.54.1. Features	125
1.9.54.1.1. Using dual Intel E810 Westport Channel NICs as grandmaster clock with the PTP Operator	125
1.9.54.2. Bug Fixes	125
1.9.54.3. Updating	126
1.9.55. RHSA-2024:0642 - OpenShift Container Platform 4.14.11 bug fix and security update	126
1.9.55.1. Features	126
1.9.55.1.1. Enabling configuration of whereabouts cron schedule	126
1.9.55.2. Bug fixes	126

1.9.55.3. Updating	127
1.9.56. RHSA-2024:0290 - OpenShift Container Platform 4.14.10 bug fix and security update	127
1.9.56.1. Bug fixes	127
1.9.56.2. Updating	127
1.9.57. RHSA-2024:0204 - OpenShift Container Platform 4.14.9 bug fix and security update	127
1.9.57.1. Bug fixes	127
1.9.57.2. Updating	128
1.9.58. RHSA-2024:0050 - OpenShift Container Platform 4.14.8 bug fix and security update	128
1.9.58.1. Features	128
1.9.58.2. Bug fixes	128
1.9.58.3. Updating	129
1.9.59. RHSA-2023:7831 - OpenShift Container Platform 4.14.7 bug fix and security update	129
1.9.59.1. Bug fixes	129
1.9.59.2. Updating	129
1.9.60. RHSA-2023:7682 - OpenShift Container Platform 4.14.6 bug fix and security update	130
1.9.60.1. Features	130
1.9.60.1.1. Using hardware-specific NIC features with the PTP Operator	130
1.9.60.1.2. Using GNSS timing synchronization for PTP grandmaster clocks	130
1.9.60.2. Bug fixes	130
1.9.60.3. Known issues	130
1.9.60.4. Updating	130
1.9.61. RHSA-2023:7599 - OpenShift Container Platform 4.14.5 bug fix and security update	131
1.9.61.1. Bug fixes	131
1.9.61.2. Updating	131
1.9.62. RHSA-2023:7470 - OpenShift Container Platform 4.14.4 bug fix and security update	131
1.9.62.1. Bug fixes	131
1.9.62.2. Updating	132
1.9.63. RHSA-2023:7315 - OpenShift Container Platform 4.14.3 bug fix and security update	132
1.9.63.1. Bug fixes	132
1.9.63.2. Updating	133
1.9.64. RHSA-2023:6837 - OpenShift Container Platform 4.14.2 bug fix and security update	133
1.9.64.1. Bug fixes	133
1.9.64.2. Known issue	134
1.9.64.3. Updating	134
1.9.65. RHBA-2023:6153 - OpenShift Container Platform 4.14.1 bug fix update	134
1.9.65.1. Updating	135
1.9.66. RHSA-2023:5006 - OpenShift Container Platform 4.14.0 image release, bug fix, and security update advisory	135

CHAPTER 1. OPENSIFT CONTAINER PLATFORM 4.14 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-2023:5006](#)) is now available. This release uses [Kubernetes 1.27](#) with CRI-O runtime. New features, changes, and known issues that pertain to OpenShift Container Platform 4.14 are included in this topic.

OpenShift Container Platform 4.14 clusters are available at <https://console.redhat.com/openshift>. With the Red Hat OpenShift Cluster Manager application for OpenShift Container Platform, you can deploy OpenShift Container Platform clusters to either on-premises or cloud environments.

OpenShift Container Platform 4.14 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 4.14. OpenShift Container Platform 4.14 is also supported on Red Hat Enterprise Linux CoreOS (RHCOS) 4.14. To understand RHEL versions used by RHCOS, see [RHEL Versions Utilized by Red Hat Enterprise Linux CoreOS \(RHCOS\) and OpenShift Container Platform](#) (Knowledgebase article).

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

For OpenShift Container Platform 4.12 on **x86_64** architecture, Red Hat has added a 6-month Extended Update Support (EUS) phase that extends the total available lifecycle from 18 months to 24 months. For OpenShift Container Platform 4.12 running on 64-bit ARM (**aarch64**), IBM Power® (**ppc64le**), and IBM Z® (**s390x**) architectures, the EUS lifecycle remains at 18 months.

Starting with OpenShift Container Platform 4.14, each EUS phase for even numbered releases on all supported architectures, including **x86_64**, 64-bit ARM (**aarch64**), IBM Power® (**ppc64le**), and IBM Z® (**s390x**) architectures, has a total available lifecycle of 24 months.

Starting with OpenShift Container Platform 4.14, Red Hat offers a 12-month additional EUS add-on, denoted as *Additional EUS Term 2*, that extends the total available lifecycle from 24 months to 36 months. The Additional EUS Term 2 is available on all architecture variants of OpenShift Container Platform.

For more information about this support, see the [Red Hat OpenShift Container Platform Life Cycle Policy](#).

Maintenance support ends for version 4.12 on 17 July 2024 and goes to extended update support phase. For more information, see the [Red Hat OpenShift Container Platform Life Cycle Policy](#).

Commencing with the 4.14 release, Red Hat is simplifying the administration and management of Red Hat shipped cluster Operators with the introduction of three new life cycle classifications; Platform Aligned, Platform Agnostic, and Rolling Stream. These life cycle classifications provide additional ease

and transparency for cluster administrators to understand the life cycle policies of each Operator and form cluster maintenance and upgrade plans with predictable support boundaries. For more information, see [OpenShift Operator Life Cycles](#).

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

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

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

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

1.3. NEW FEATURES AND ENHANCEMENTS

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

1.3.1. Red Hat Enterprise Linux CoreOS (RHCOS)

1.3.1.1. RHCOS now uses RHEL 9.2

RHCOS now uses Red Hat Enterprise Linux (RHEL) 9.2 packages in OpenShift Container Platform 4.14. These packages ensure that your OpenShift Container Platform instance receives the latest fixes, features, enhancements, hardware support, and driver updates. Excluded from this change, OpenShift Container Platform 4.12 is an Extended Update Support (EUS) release that will continue to use RHEL 8.6 EUS packages for the entirety of its lifecycle.

1.3.1.1.1. Considerations for upgrading to OpenShift Container Platform with RHEL 9.2

Because OpenShift Container Platform 4.14 now uses a RHEL 9.2 based RHCOS, consider the following before upgrading:

- Some component configuration options and services might have changed between RHEL 8.6 and RHEL 9.2, which means existing machine configuration files might no longer be valid.
- If you customized the default OpenSSH `/etc/ssh/sshd_config` server configuration file, you must update it according to this [Red Hat Knowledgebase article](#).
- RHEL 6 base image containers are not supported on RHCOS container hosts but are supported on RHEL 8 worker nodes. For more information, see the [Red Hat Container Compatibility matrix](#).
- Some device drivers have been deprecated, see the [RHEL documentation](#) for more information.

1.3.2. Installation and update

1.3.2.1. Installing a cluster on Amazon Web Services (AWS) by using a shared VPC

In OpenShift Container Platform 4.14, you can install a cluster on AWS that uses a shared Virtual Private Cloud (VPC), with a private hosted zone in a different account than the cluster. For more information, see [Installing a cluster on AWS into an existing VPC](#).

1.3.2.2. Enabling S3 bucket to be retained during cluster bootstrap on AWS

With this update, you can opt out of the automatic deletion of the S3 bucket during the cluster bootstrap on AWS. This option is useful when you have security policies that prevent the deletion of S3 buckets.

1.3.2.3. Installing a cluster on Microsoft Azure using a NAT gateway (Technology Preview)

In OpenShift Container Platform 4.14, you can install a cluster that uses a NAT gateway for outbound networking. This is available as a Technology Preview (TP). For more information, see [Additional Azure configuration parameters](#).

1.3.2.4. Installing a cluster on Google Cloud Platform (GCP) using pd-balanced disk type

In OpenShift Container Platform 4.14, you can install a cluster on GCP using the **pd-balanced** disk type. This disk type is only available for compute nodes and cannot be used for control plane nodes. For more information, see [Additional GCP configuration parameters](#).

1.3.2.5. Optional capabilities in OpenShift Container Platform 4.14

In OpenShift Container Platform 4.14, you can disable the **Build**, **DeploymentConfig**, **ImageRegistry** and **MachineAPI** capabilities during installation. You can only disable the **MachineAPI** capability if you install a cluster with user-provisioned infrastructure. For more information, see [Cluster capabilities](#).

1.3.2.6. Installing a cluster with Azure AD Workload Identity

During installation, you can now configure a Microsoft Azure cluster to use Azure AD Workload Identity. With Azure AD Workload Identity, cluster components use short-term security credentials that are managed outside the cluster.

For more information about the short-term credentials implementation for OpenShift Container Platform clusters on Azure, see [Azure AD Workload Identity](#).

To learn how to configure this credentials management strategy during installation, see [Configuring an Azure cluster to use short-term credentials](#).

1.3.2.7. User-defined tags for Microsoft Azure now generally available

The user-defined tags feature for Microsoft Azure was previously introduced as Technology Preview in OpenShift Container Platform 4.13 and is now generally available in OpenShift Container Platform 4.14. For more information, see [Configuring the user-defined tags for Azure](#).

1.3.2.8. Confidential VMs for Azure (Technology Preview)

You can enable confidential VMs when you install your cluster on Azure. You can use confidential computing to encrypt the virtual machine guest state storage during installation. This feature is in

Technology Preview due to known issues which are listed in the Known Issues section of this page. For more information, see [Enabling confidential VMs](#).

1.3.2.9. Trusted launch for Azure (Technology Preview)

You can enable trusted launch features when you install your cluster on Azure as a Technology Preview. These features include secure boot and virtualized Trusted Platform Modules. For more information, see [Enabling trusted launch for Azure VMs](#).

1.3.2.10. User-defined labels and tags for Google Cloud Platform (Technology Preview)

You can now configure user-defined labels and tags in Google Cloud Platform (GCP) for grouping resources and for managing resource access and cost. User-defined labels can be applied only to resources created with the OpenShift Container Platform installation program and its core components. User-defined tags can be applied only to resources created with the OpenShift Container Platform Image Registry Operator. For more information, see [Managing the user-defined labels and tags for GCP](#).

1.3.2.11. Installing an OpenShift Container Platform cluster on Microsoft Azure in a restricted network

In OpenShift Container Platform 4.14, you can install a cluster on Microsoft Azure in a restricted network for installer-provisioned infrastructure (IPI) and user-provisioned infrastructure (UPI). For IPI, you can create an internal mirror of the installation release content on an existing Azure Virtual Network (VNet). For UPI, you can install a cluster on Microsoft Azure by using infrastructure that you provide. For more information, see [Installing a cluster on Azure in a restricted network](#) and [Installing a cluster on Azure in a restricted network with user-provisioned infrastructure](#).

1.3.2.12. Specifying installation disks using a by-path device alias

You can now specify the installation disk using a by-path device alias, such as **deviceName: "/dev/disk/by-path/pci-0000:01:00.0-scsi-0:0:0:0"**, when you install a cluster on bare metal with installer-provisioned infrastructure. You can also specify this parameter during Agent-based installations. This type of disk alias persists across reboots. For more information, see [Configuring the install-config.yaml file for bare metal](#) or [About root device hints for Agent-based installations](#).

1.3.2.13. Applying existing AWS security groups to a cluster

By default, the installation program creates and attaches security groups to control plane and compute machines. The rules associated with the default security groups cannot be modified.

With OpenShift Container Platform 4.14, if you deploy a cluster to an existing Amazon Virtual Private Cloud (VPC), you can apply additional existing AWS security groups to control plane and compute machines. These security groups must be associated with the VPC that you are deploying the cluster to. Applying custom security groups can help you meet the security needs of your organization, in such cases where you must control the incoming or outgoing traffic of these machines. For more information, see [Applying existing AWS security groups to the cluster](#).

1.3.2.14. Required administrator acknowledgment when updating from OpenShift Container Platform 4.13 to 4.14

OpenShift Container Platform 4.14 uses Kubernetes 1.27, which removed a [deprecated API](#).

A cluster administrator must provide a manual acknowledgment before the cluster can be updated from

OpenShift Container Platform 4.13 to 4.14. This is to help prevent issues after updating to OpenShift Container Platform 4.14, 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 4.13 clusters require this administrator acknowledgment before they can be updated to OpenShift Container Platform 4.14.

For more information, see [Preparing to update to OpenShift Container Platform 4.14](#) .

1.3.2.15. Three-node cluster support for Nutanix

Deploying a three-node cluster is supported on Nutanix as of OpenShift Container Platform 4.14. This type of OpenShift Container Platform cluster is a more resource efficient cluster. It consists of only three control plane machines, which also act as compute machines. For more information, see [Installing a three-node cluster on Nutanix](#).

1.3.2.16. Installing a cluster on GCP using Confidential VMs is generally available

In OpenShift Container Platform 4.14, using Confidential VMs when installing your cluster is generally available. Confidential VMs are currently not supported on 64-bit ARM architectures. For more information, see [Enabling Confidential VMs](#).

1.3.2.17. Root volume types parameter for RHOSP is now available

You can now specify one or more root volume types in RHOSP, by using the **rootVolume.types** parameter. This parameter is available for both control plane and compute machines.

1.3.2.18. Static IP addresses for vSphere nodes

You can provision bootstrap, control plane, and compute nodes with static IP addresses in environments where Dynamic Host Configuration Protocol (DHCP) does not exist.



IMPORTANT

Static IP addresses for vSphere nodes is a Technology Preview feature only. Technology Preview features are not supported with Red Hat production service level agreements (SLAs) and might not be functionally complete. Red Hat does not recommend using them in production. These features provide early access to upcoming product features, enabling customers to test functionality and provide feedback during the development process.

For more information about the support scope of Red Hat Technology Preview features, see [Technology Preview Features Support Scope](#) .

After you have deployed your cluster to run nodes with static IP addresses, you can scale a machine to use one of these static IP addresses. Additionally, you can use a machine set to configure a machine to use one of the configured static IP addresses.

For more information, see the "Static IP addresses for vSphere nodes" section in the [Installing a cluster on vSphere](#) document.

1.3.2.19. Additional validation for the Bare Metal Host CR

The Bare Metal Host Custom Resource (CR) now contains the **ValidatingWebhooks** parameter. With this parameter, the Bare Metal Operator now catches any configuration errors before accepting the CR, and returns a message with the configuration errors to the user.

1.3.2.20. Install a cluster quickly in AWS Local Zones

For OpenShift Container Platform 4.14, you can quickly install a cluster on Amazon Web Services (AWS) to extend compute nodes to Local Zone locations. After you add zone names to the installation configuration file, the installation program fully automates the creation of required resources, network and compute, on each Local Zone. For more information, see [Intall a cluster quickly in AWS Local Zones](#) .

1.3.2.21. Simplified installation and update experience for clusters with manually maintained cloud credentials

This release includes changes that improve the experience of installing and updating clusters that use the Cloud Credential Operator (CCO) in manual mode for cloud provider authentication. The following parameters for the **oc adm release extract** command simplify the manual configuration of cloud credentials:

--included

Use this parameter to extract only the manifests that your specific cluster configuration needs. If you use cluster capabilities to disable one or more optional components, you are no longer required to delete the **CredentialsRequest** CRs for any disabled components before installing or updating a cluster.

In a future release, this parameter might make the CCO utility (**ccoctl**) **--enable-tech-preview** parameter unnecessary.

--install-config

Use this parameter to specify the location of the **install-config.yaml** file when installing a cluster. By referencing the **install-config.yaml** file, the extract command can determine aspects of the cluster configuration for the cluster that you are about to create. This parameter is not needed during a cluster update because **oc** can connect to the cluster to determine its configuration.


With this change, you are no longer required to specify the cloud platform you are installing on with the **--cloud** parameter. As a result, the **--cloud** parameter is deprecated starting in OpenShift Container Platform 4.14.

To understand how to use these parameters, see the installation procedure for your configuration and the procedures in [Preparing to update a cluster with manually maintained credentials](#) .

1.3.2.22. Quickly install RHCOS on vSphere hosts by using a pre-existing RHCOS image template

OpenShift Container Platform 4.14 includes a new VMware vSphere configuration parameter for use on installer-provisioned infrastructure: **template**. By using this parameter, you can now specify the absolute path to a pre-existing Red Hat Enterprise Linux CoreOS (RHCOS) image template or virtual machine in the installation configuration file. The installation program can then use the image template or virtual machine to quickly install RHCOS on vSphere hosts.

This installation method is an alternative to uploading an RHCOS image on vSphere hosts.



IMPORTANT

Before you set a path value for the **template** parameter, ensure that the default RHCOS boot image in the OpenShift Container Platform release matches the RHCOS image template or virtual machine version; otherwise, cluster installation might fail.

1.3.2.23. OpenShift Container Platform on 64-bit ARM

OpenShift Container Platform 4.14 is now supported on 64-bit ARM architecture-based Google Cloud Platform installer-provisioned and user-provisioned infrastructures. You can also now use the **oc mirror** CLI plug-in disconnected environments on 64-bit ARM clusters. For more information about instance availability and installation documentation, see [Supported installation methods for different platforms](#).

1.3.2.24. Using a custom RHCOS image for a Microsoft Azure cluster

By default, the installation program downloads and installs the Red Hat Enterprise Linux CoreOS (RHCOS) image that is used to boot control plane and compute machines. With this enhancement, you can now override the default behavior by modifying the installation configuration file (**install-config.yaml**) to specify a custom RHCOS image. Before you deploy the cluster, you can modify the following installation parameters:

- **compute.platorm.azure.osImage.publisher**
- **compute.platorm.azure.osImage.offer**
- **compute.platorm.azure.osImage.sku**
- **compute.platorm.azure.osImage.version**
- **controlPlane.platorm.azure.osImage.publisher**
- **controlPlane.platorm.azure.osImage.offer**
- **controlPlane.platorm.azure.osImage.sku**
- **controlPlane.platorm.azure.osImage.version**
- **platform.azure.defaultMachinePlatform.osImage.publisher**
- **platform.azure.defaultMachinePlatform.osImage.offer**
- **platform.azure.defaultMachinePlatform.osImage.sku**
- **platform.azure.defaultMachinePlatform.osImage.version**

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

1.3.2.25. Installing single-node OpenShift on cloud providers

OpenShift Container Platform 4.14 expands support for installing single-node OpenShift on cloud providers. Installation options for single-node OpenShift include Amazon Web Services (AWS), Google Cloud Platform (GCP), and Microsoft Azure. For more information about the supported platforms, see [Supported cloud providers for single node openshift](#).

1.3.3. Post-installation configuration

1.3.3.1. OpenShift Container Platform cluster with multi-architecture compute machines

OpenShift Container Platform 4.14 clusters with multi-architecture compute machines are now supported on Google Cloud Platform (GCP) as a Day 2 operation. OpenShift Container Platform clusters with multi-architecture compute machines on bare metal installations are now generally available. For more information on clusters with multi-architecture compute machines and supported platforms, see [About clusters with multi-architecture compute machines](#).

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. You can now perform the following actions:

- Narrow down the list of resources in a list view or search page with exact search capabilities. This action is useful when you have similarly named resources and the standard search functionality does not narrow down your search.
- Provide direct feedback about features and report a bug by clicking the **Help** button on the toolbar and clicking **Share Feedback** from the drop-down list.
- Display and hide tooltips in the YAML editor. Because the tooltips persist, you do not need to change a tooltip every time you navigate to a page.
- Configure the web terminal image for all users. For more information, see [Configuring the web terminal](#).

1.3.4.1.1. Dynamic plugin enhancements

With this update, you can add custom metric dashboards and extend the cluster's **Overview** page with the **QueryBrowser** extension. The OpenShift Container Platform release adds additional extension points, so you can add different types of modals, set the active namespace, provide custom error pages, and set proxy timeouts for your dynamic plugin.

For more information, see [Dynamic plugin reference](#) and **QueryBrowser** in the [OpenShift Container Platform console API](#).

1.3.4.1.2. Operating system based filtering in OperatorHub

With this update, Operators in OperatorHub are now filtered based on the operating systems of the nodes, because clusters can contain heterogenous nodes.

1.3.4.1.3. Support for installing specific Operator versions in the web console

With this update, you can now choose from a list of available versions for an Operator based on the selected channel on the **OperatorHub** page in the console. Additionally, you can view the metadata for that channel and version when available. When selecting an older version, a manual approval update strategy is required, otherwise the Operator immediately updates back to the latest version on the channel.

For more information, see [Installing a specific version of an Operator in the web console](#) .

1.3.4.1.4. OperatorHub support for AWS STS

With this release, OperatorHub detects when an Amazon Web Services (AWS) cluster is using the Security Token Service (STS). When detected, a "Cluster in STS Mode" notification displays with additional instructions before installing an Operator to ensure it runs correctly. The **Operator Installation** page is also modified to add the required **role ARN** field. For more information, see [Token authentication for Operators on cloud providers](#).

1.3.4.2. Developer Perspective

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

- Change the default timeout period for the web terminal for your current session. For more information, see [Configuring the web terminal timeout for a session](#).
- Test Serverless functions in the web console from the **Topology** view and the Serverless Service **List** and **Detail** pages, so that you can use a Serverless function with a CloudEvent or HTTP request.
- View status, start time, and duration of the latest build for **BuildConfigs** and Shipwright builds. You can also view this information on the **Details** page.

1.3.4.2.1. New quick starts

With this release, new quick starts exist where you can discover developer tools, such as installing the Cryostat Operator and getting started with JBoss EAP by using a helm chart.

1.3.4.2.2. OpenShift Pipelines page improvements

In OpenShift Container Platform 4.14, you can see the following navigation improvements on the **Pipelines** page:

- Autodetection of Pipelines as Code (PAC) in Git import flow.
- Serverless functions in the samples catalog.

1.3.5. OpenShift CLI (oc)

1.3.5.1. Supporting multi-arch OCI local images for catalogs with oc-mirror

With OpenShift Container Platform 4.14, oc-mirror supports multi-arch OCI local images for catalogs.

OCI layouts consist of an **index.json** file that identifies the images held within them on disk. This **index.json** file can reference any number of single or multi-arch images. However, oc-mirror only references a single image at a time in a given OCI layout. The image stored in the OCI layout can be a single-arch image, that is, an image manifest or a multi-arch image, that is, a manifest list.

The **ImageSetConfiguration** stores the OCI images. After processing the catalog, the catalog content adds new layers representing the content of all images in the layout. The ImageBuilder is modified to handle image updates for both single-arch and multi-arch images.

1.3.5.2. Logging in to the CLI using a web browser

With OpenShift Container Platform 4.14, a new **oc** command-line interface (CLI) flag, **--web** is now available for the **oc login** command.

With this enhancement, you can log in by using a web browser, so that you do not need to insert your access token into the command line.

For more information, see [Logging in to the OpenShift CLI using a web browser](#) .

1.3.5.3. Enhancement to `oc new-build`

A new `oc` CLI flag, `--import-mode`, has been added to the `oc new-build` command. With this enhancement, you can set the `--import-mode` flag to **Legacy** or **PreserverOriginal**, so that you trigger builds by using a single sub-manifest or all manifests.

1.3.5.4. Enhancement to `oc new-app`


A new `oc` CLI flag, `--import-mode`, has been added to the `oc new-app` command. With this enhancement, you can set the `--import-mode` flag to **Legacy** or **PreserverOriginal**, and then create new applications by using a single sub-manifest or all manifests.

For more information, see [Setting the import mode](#).

1.3.6. IBM Z and IBM LinuxONE

With this release, IBM Z[®] and IBM[®] LinuxONE are now compatible with OpenShift Container Platform 4.14. The installation can be performed with z/VM or Red Hat Enterprise Linux (RHEL) Kernel-based Virtual Machine (KVM). For installation instructions, see the following documentation:

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



IMPORTANT

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

1.3.6.1. IBM Z and IBM LinuxONE notable enhancements

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

The IBM Z[®] and IBM[®] LinuxONE release on OpenShift Container Platform 4.14 adds improvements and new capabilities to OpenShift Container Platform components and concepts.

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

- Assisted Installer with z/VM
- Installing on a single node
- Hosted control planes (Technology Preview)
- Multi-architecture compute nodes

- [oc-mirror plugin](#)

1.3.6.2. IBM Secure Execution

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

For installation instructions, see the following documentation:

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

1.3.7. IBM Power

IBM Power® is now compatible with OpenShift Container Platform 4.14. For installation instructions, see the following documentation:

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



IMPORTANT

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

1.3.7.1. IBM Power notable enhancements

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

The IBM Power® release on OpenShift Container Platform 4.14 adds improvements and new capabilities to OpenShift Container Platform components.

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

- IBM Power® Virtual Server Block CSI Driver Operator (Technology Preview)
- Installing on a single node
- Hosted control planes (Technology Preview)
- Multi-architecture compute nodes
- [oc-mirror plugin](#)

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

Table 1.1. OpenShift Container Platform features

Feature	IBM Power®	IBM Z® and IBM® LinuxONE
Alternate authentication providers	Supported	Supported
Automatic Device Discovery with Local Storage Operator	Unsupported	Supported

Feature	IBM Power®	IBM Z® and IBM® LinuxONE
Automatic repair of damaged machines with machine health checking	Unsupported	Unsupported
Cloud controller manager for IBM Cloud	Supported	Unsupported
Controlling overcommit and managing container density on nodes	Unsupported	Unsupported
Cron jobs	Supported	Supported
Descheduler	Supported	Supported
Egress IP	Supported	Supported
Encrypting data stored in etcd	Supported	Supported
FIPS cryptography	Supported	Supported
Helm	Supported	Supported
Horizontal pod autoscaling	Supported	Supported
IBM Secure Execution	Unsupported	Supported
IBM Power® Virtual Server Block CSI Driver Operator (Technology Preview)	Supported	Unsupported
Installer-provisioned Infrastructure Enablement for IBM Power® Virtual Server (Technology Preview)	Supported	Unsupported
Installing on a single node	Supported	Supported
IPv6	Supported	Supported
Monitoring for user-defined projects	Supported	Supported
Multi-architecture compute nodes	Supported	Supported
Multipathing	Supported	Supported
Network-Bound Disk Encryption - External Tang Server	Supported	Supported
Non-volatile memory express drives (NVMe)	Supported	Unsupported

Feature	IBM Power®	IBM Z® and IBM® LinuxONE
oc-mirror plugin	Supported	Supported
OpenShift CLI (oc) plugins	Supported	Supported
Operator API	Supported	Supported
OpenShift Virtualization	Unsupported	Unsupported
OVN-Kubernetes, including IPsec encryption	Supported	Supported
PodDisruptionBudget	Supported	Supported
Precision Time Protocol (PTP) hardware	Unsupported	Unsupported
Red Hat OpenShift Local	Unsupported	Unsupported
Scheduler profiles	Supported	Supported
Stream Control Transmission Protocol (SCTP)	Supported	Supported
Support for multiple network interfaces	Supported	Supported
Three-node cluster support	Supported	Supported
Topology Manager	Supported	Unsupported
z/VM Emulated FBA devices on SCSI disks	Unsupported	Supported
4K FCP block device	Supported	Supported

Table 1.2. Persistent storage options

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

Feature	IBM Power®	IBM Z® and IBM® LinuxONE
Persistent storage using Raw Block	Supported ^[1]	Supported ^{[1],[2]}
Persistent storage using EDEV/FBA	Supported ^[1]	Supported ^{[1],[2]}

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

Table 1.3. Operators

Feature	IBM Power®	IBM Z® and IBM® LinuxONE
Cluster Logging Operator	Supported	Supported
Cluster Resource Override Operator	Supported	Supported
Compliance Operator	Supported	Supported
File Integrity Operator	Supported	Supported
HyperShift Operator	Technology Preview	Technology Preview
Local Storage Operator	Supported	Supported
MetalLB Operator	Supported	Supported
Network Observability Operator	Supported	Supported
NFD Operator	Supported	Supported
NMState Operator	Supported	Supported
OpenShift Elasticsearch Operator	Supported	Supported
Vertical Pod Autoscaler Operator	Supported	Supported

Table 1.4. Multus CNI plugins

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

Table 1.5. CSI Volumes

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

1.3.9. Authentication and authorization

1.3.9.1. SCC preemption prevention

With this release, you can now require your workloads to use a specific security context constraint (SCC). By setting a specific SCC, you can prevent the SCC that you want from being preempted by another SCC in the cluster. For more information, see [Configuring a workload to require a specific SCC](#) .

1.3.9.2. Pod security admission privileged namespaces

With this release, the following system namespaces are always set to the **privileged** pod security admission profile:

- **default**
- **kube-public**
- **kube-system**

For more information, see [Privileged namespaces](#).

1.3.9.3. Pod security admission synchronization disabled on modified namespaces

With this release, if a user manually modifies a pod security admission label from the automatically labeled value on a label-synchronized namespace, synchronization is disabled for that label. Users can enable synchronization again, if necessary. For more information, see [Pod security admission synchronization namespace exclusions](#).

1.3.9.4. OLM-based Operator support for AWS STS

With this release, some Operators managed by Operator Lifecycle Manager (OLM) on Amazon Web Services (AWS) clusters can use the Cloud Credential Operator (CCO) in manual mode with the Security Token Service (STS). These Operators authenticate with limited-privilege, short-term credentials that are managed outside the cluster. For more information, see [Token authentication for Operators on cloud providers](#).

1.3.9.5. Authentication Operator honors noProxy during connection checks

With this release, if the **noProxy** field is set and the route is reachable without the cluster-wide proxy, the Authentication Operator will bypass the proxy and perform connection checks directly through the configured ingress route. Previously, the Authentication Operator always performed connection checks through the cluster-wide proxy, regardless of the **noProxy** setting. For more information, see [Configuring the cluster-wide proxy](#).

1.3.10. Networking

1.3.10.1. IPv6 as primary IP address family on vSphere dual-stack clusters

During cluster installation on vSphere, you can configure IPv6 as the primary IP address family on a dual-stack cluster. To enable this feature when installing a new cluster, specify an IPv6 address family before an IPv4 address family for the machine network, cluster network, service network, API VIPs, and ingress VIPs.

- Installer-provisioned infrastructure: [Deploying with dual-stack networking](#)
- User-provisioned infrastructure: [Network configuration parameters](#)

1.3.10.2. Multiple external gateway support for the OVN-Kubernetes network plugin

The OVN-Kubernetes network plugin supports defining additional default gateways for specific workloads. Both IPv4 and IPv6 address families are supported. You define each default gateway by using the **AdminPolicyBasedExternalRoute** object, in which you can specify two types of next hops, static and dynamic:

- Static next hop: One or more IP addresses of external gateways
- Dynamic next hop: A combination of pod and namespace selectors for pod selection, and a network attachment definition name previously associated with the selected pods.

The next hops that you define are scoped by a namespace selector that you specify. You can then use the external gateway for specific workloads that match the namespace selector.

For more information, refer to [Configure an external gateway through a secondary network interface](#) .

1.3.10.3. Ingress Node Firewall Operator is generally available

Ingress Node Firewall Operator was designated a Technology Preview feature in OpenShift Container Platform 4.12. With this release, Ingress Node Firewall Operator is generally available. You can now configure firewall rules at the node level. For more information, see [Ingress Node Firewall Operator](#).

1.3.10.4. Dynamic use of non-reserved CPUs for OVS

With this release, the Open vSwitch (OVS) networking stack can dynamically use non-reserved CPUs. This dynamic use of non-reserved CPUs occurs by default in nodes in a machine config pool that has a performance profile applied to it. The dynamic use of available, non-reserved CPUs maximizes compute resources for OVS and minimizes network latency for workloads during periods of high demand. OVS remains unable to dynamically use isolated CPUs assigned to containers in **Guaranteed** QoS pods. This separation avoids disruption to critical application workloads.



NOTE

When the Node Tuning Operator recognizes the performance conditions to activate the use of non-reserved CPUs, there is a several second delay while OVN-Kubernetes configures the CPU affinity alignment of OVS daemons running on the CPUs. During this window, if a **Guaranteed** QoS pod starts, it can experience a latency spike.

1.3.10.5. Dual-stack configuration for multiple IP addresses

In previous releases of the Whereabouts IPAM CNI plugin, only one IP address could be assigned per network interface.

Now, Whereabouts supports the assignment of an arbitrary number of IP addresses to support dual-stack IPv4/IPv6 functionality. See [Creating a configuration for assignment of dual-stack IP addresses dynamically](#).

1.3.10.6. Exclude SR-IOV network topology for NUMA-aware scheduling

With this release, you can exclude advertising the Non-Uniform Memory Access (NUMA) node for the SR-IOV network to the Topology Manager. By not advertising the NUMA node for the SR-IOV network, you can permit more flexible SR-IOV network deployments during NUMA-aware pod scheduling.

For example, in some scenarios, it is a priority to maximize CPU and memory resources for a pod on a single NUMA node. By not providing a hint to the Topology Manager about the NUMA node for the pod's SR-IOV network resource, the Topology Manager can deploy the SR-IOV network resource and the pod CPU and memory resources to different NUMA nodes. In earlier OpenShift Container Platform releases, the Topology Manager attempted to place all resources on the same NUMA node only.

For more information about this more flexible SR-IOV network deployment during NUMA-aware pod scheduling, see [Exclude the SR-IOV network topology for NUMA-aware scheduling](#).

1.3.10.7. Update to HAProxy 2.6

With this release, OpenShift Container Platform is updated to HAProxy 2.6.

1.3.10.8. Support for configuring the maximum length with sidecar logging in the Ingress Controller

Previously, the maximum length of the syslog message in the Ingress Controller was 1024 bytes. Now, the maximum value can be increased. For more information, see [Allow the Ingress Controller to modify the HAProxy log length when using a sidecar](#).

1.3.10.9. NMstate Operator updated in console

With this release, you can access the NMstate Operator and resources such as the **NodeNetworkState** (NNS), **NodeNetworkConfigurationPolicy** (NNCP), and **NodeNetworkConfigurationEnhancement** (NNCE) from the web console. In the **Administrator** perspective of the console from the **Networking**

page you can access NNCP, NNCE from the **NodeNetworkConfigurationPolicy** page, and NNS on the **NodeNetworkState** page. For more information about NMState resources and how to update them in the console, see [Updating node network configuration](#).

1.3.10.10. OVN-Kubernetes network plugin support for IPsec on IBM Cloud

IPsec is now supported on the IBM Cloud platform for clusters that use the OVN-Kubernetes network plugin, which is the default in OpenShift Container Platform 4.14. For more information, see [Configuring IPsec encryption](#).

1.3.10.11. OVN-Kubernetes network plugin support for IPsec encryption of external traffic (Technology Preview)

OpenShift Container Platform now supports encryption of external traffic, also known as *north-south traffic*. IPsec already supports encryption of network traffic between pods, known as *east-west traffic*. You can use both features in conjunction to provide full in-transit encryption for OpenShift Container Platform clusters. This is available as a Technology Preview feature.

To use this feature, you need to define an IPsec configuration tuned for your network infrastructure. For more information, refer to [Enabling IPsec encryption for external IPsec endpoints](#).

1.3.10.12. Single-stack IPv6 support for Kubernetes NMState

With this release, you can use Kubernetes NMState Operator in single-stack IPv6 clusters.

1.3.10.13. Egress service resource to manage egress traffic for pods behind a load balancer (Technology Preview)

With this update, you can use an **EgressService** custom resource (CR) to manage egress traffic for pods behind a load balancer service. This is available as a Technology Preview feature.

You can use the **EgressService** CR to manage egress traffic in the following ways:

- Assign the load balancer service's IP address as the source IP address of egress traffic for pods behind the load balancer service.
- Assign the egress traffic for pods behind a load balancer to a different network than the default node network.

For more information, see [Configuring an egress service](#).

1.3.10.14. VRF specification in MetalLB's BGPPeer resource (Technology Preview)

With this update, you can specify a Virtual Routing and Forwarding (VRF) instance in a **BGPPeer** custom resource. MetalLB can advertise services through the interfaces belonging to the VRF. This is available as a Technology Preview feature. For more information, see [Exposing a service through a network VRF](#).

1.3.10.15. VRF specification in NMState's NodeNetworkConfigurationPolicy resource (Technology Preview)

With this update, you can associate a Virtual Routing and Forwarding (VRF) instance with a network interface by using a **NodeNetworkConfigurationPolicy** custom resource. By associating a VRF instance with a network interface, you can support traffic isolation, independent routing decisions, and the logical

separation of network resources. This feature is available as a Technology Preview feature. For more information, see [Example: Network interface with a VRF instance node network configuration policy](#) .

1.3.10.16. Support for Broadcom BCM57504 is now GA

Support for the Broadcom BCM57504 network interface controller is now available for the SR-IOV Network Operator. For more information, see [Supported devices](#).

1.3.10.17. OVN-Kubernetes is available as a secondary network

With this release, the Red Hat OpenShift Networking OVN-Kubernetes network plugin allows the configuration of secondary network interfaces for pods. As a secondary network, OVN-Kubernetes supports both layer 2 switched and localnet switched topology networks. For more information about OVN-Kubernetes as a secondary network, see [Configuration for an OVN-Kubernetes additional network](#).

1.3.10.18. Global IP forwarding is disabled on OVN-Kubernetes based cluster deployments

Starting with this release, global IP address forwarding is disabled on OVN-Kubernetes based cluster deployments to prevent undesirable effects for cluster administrators with nodes acting as routers. OVN-Kubernetes now enables and restricts forwarding on a per-managed interface basis.

You can control IP forwarding for all traffic on OVN-Kubernetes managed interfaces by using the **gatewayConfig.ipForwarding** specification in the **Network** resource. Specify **Restricted** to forward all traffic related to OVN-Kubernetes only. Specify **Global** to allow forwarding of all IP traffic. For new installations, the default is **Restricted**. Clusters upgraded to 4.14 are not affected by this change; the IP forwarding behavior remains unchanged and continues to be globally enabled. For more information about enabling IP forwarding globally, see [Enabling IP forwarding globally](#) .

1.3.10.19. Admin Network Policy (Technology Preview)

Admin Network Policy is available as a Technology Preview feature. You can enable **AdminNetworkPolicy** and **BaselineAdminNetworkPolicy** resources, which are part of the Network Policy V2 API, in clusters running the OVN-Kubernetes CNI plugin. Cluster administrators can apply cluster-scoped policies and safeguards for an entire cluster before namespaces are created. Network administrators can secure clusters by enforcing network traffic controls that cannot be overridden by users. Network administrators can enforce optional baseline network traffic controls that can be overridden by users in the cluster, if necessary. Currently, these APIs support only expressing policies for intra-cluster traffic.

1.3.10.20. MAC-VLAN, IP-VLAN, and VLAN subinterface creation for pods

With this release, the ability to create a MAC-VLAN, IP-VLAN, and VLAN subinterface based on a master interface in a container namespace is generally available. You can use this feature to create the master interfaces as part of the pod network configuration in a separate network attachment definition. You can then base the VLAN, MACVLAN or IPVLAN on this interface without knowing the network configuration of the node. For more information, see [About configuring the master interface in the container network namespace](#).

1.3.10.21. Supporting all-multicast mode

OpenShift Container Platform release now supports configuring the all-multicast mode by using the tuning CNI plugin. This update eliminates the need to grant the **NET_ADMIN** capability to the pod's Security Context Constraints (SCC), enhancing security by minimizing potential vulnerabilities for your

Pods.

For further information about all-multicast mode, see [About all-multicast mode](#).

1.3.10.22. Enhance network flexibility by using the TAP device plugin

This release introduces a new Container Network Interface (CNI) network plugin type: the TAP device plugin. You can use this plugin to create TAP devices within containers, which enables user-space programs to handle network frames and act as an interface that receives frames from and that sends frames to user-space applications instead of through traditional network interfaces. For more information, see [Configuration for a TAP additional network](#).

1.3.10.23. Support for running rootless DPDK workloads with kernel access by using the TAP CNI plugin

In OpenShift Container Platform version 4.14 and later, DPDK applications that need to inject traffic to the kernel can run in non-privileged pods with the help of the TAP CNI plugin. For more information, see [Using the TAP CNI to run a rootless DPDK workload with kernel access](#).

1.3.10.24. Set or delete specific HTTP headers using an Ingress Controller or a Route object

Certain HTTP request and response headers can now be set or deleted either globally by using an Ingress Controller or for specific routes. You can set or delete the following headers:

- X-Frame-Options
- X-Cache-Info
- X-XSS-Protection
- X-Source
- X-SSL-Client-Cert
- X-Target
- Content-Location
- Content-Language

For more information, see [Setting or deleting HTTP request and response headers in an Ingress Controller](#) and [Setting or deleting HTTP request and response headers in a route](#).

1.3.10.25. Egress IPs on additional network interfaces is generally available

You can use egress IP addresses on additional network interfaces is generally available. This feature provides OpenShift Container Platform administrators with a greater level of control over networking aspects such as routing, addressing, segmentation, and security policies. You can also route workload traffic over specific network interfaces for purposes such as traffic segmentation or meeting specialized requirements.

For more information, see [Considerations for using an egress IP on additional network interfaces](#).

1.3.10.26. Parallel node draining during SR-IOV network policy updates

With this update, you can configure the SR-IOV Network Operator to drain nodes in parallel during network policy updates. The option to drain nodes in parallel enables faster rollouts of SR-IOV network configurations. You can use the **SriovNetworkPoolConfig** custom resource to configure parallel node draining and define the maximum number of nodes in the pool that the Operator can drain in parallel.

For more information, see [Configuring parallel node draining during SR-IOV network policy updates](#) .

1.3.11. Registry

1.3.11.1. Optional Image Registry Operator

With this release, the Image Registry Operator is now an optional component. This feature helps reduce the overall resources footprint of OpenShift Container Platform in Telco environments when the Image Registry Operator is not needed. For more information about disabling the Image Registry Operator, see [Selecting cluster capabilities](#).

1.3.12. Storage

1.3.12.1. Support for OR logic in LVMS

With this release, the logical volume manager (LVM) cluster custom resource (CR) provides **OR** logic in the **deviceSelector** setting. In previous releases, specifying the **paths** setting for device paths used **AND** logic only. With this release, you can also specify the **optionalPaths** setting, which supports **OR** logic. For more information, see the CR examples in [Persistent storage using logical volume manager storage](#).

1.3.12.2. Support for ext4 in LVMS

With this release, the logical volume manager (LVM) cluster custom resource (CR) provides support for the **ext4** filesystem with the **fstype** setting under **deviceClasses**. The default filesystem is **xfs**. For more information, see the CR examples in [Persistent storage using logical volume manager storage](#) .

1.3.12.3. Standardized STS configuration workflow

OpenShift Container Platform 4.14 provides a streamlined and standardized procedure to configure Security Token Service (STS) with the AWS Elastic File Storage (EFS) Container Storage Interface (CSI) Driver Operator.

For more information, see [Obtaining a role Amazon Resource Name for Security Token Service](#) .

1.3.12.4. Read Write Once Pod access mode (Technology Preview)

OpenShift Container Platform 4.14 introduces a new access mode for persistent volumes (PVs) and persistent volume claims (PVCs) called ReadWriteOncePod (RWOP), which can be used only in a single pod on a single node. This is compared to the existing ReadWriteOnce access mode where a PV or PVC can be used on a single node by many pods. This is available as a Technology Preview feature.

For more information, see [Access modes](#).

1.3.12.5. GCP Filestore storage CSI Driver Operator is generally available

OpenShift Container Platform is capable of provisioning persistent volumes (PVs) using the Container Storage Interface (CSI) driver for Google Compute Platform (GCP) Filestore Storage. The GCP

Filestore CSI Driver Operator was introduced in OpenShift Container Platform 4.12 with Technology Preview support. The GCP Filestore CSI Driver Operator is now generally available. For more information, see [Google Compute Platform Filestore CSI Driver Operator](#) .

1.3.12.6. Automatic CSI migration for VMware vSphere

The Automatic CSI migration for VMware vSphere feature automatically translates in-tree objects to their counterpart CSI representations and, ideally, must be completely transparent to users. Although storage class referencing to the in-tree storage plug-in continues to work, consider switching the default storage class to the CSI storage class.

In OpenShift Container Platform 4.14, CSI migration for vSphere is enabled by default under all circumstances and requires no action by an administrator.

If you are using vSphere in-tree persistent volumes (PVs) and want to upgrade from OpenShift Container Platform 4.12 or 4.13 to 4.14, update vSphere vCenter and ESXI host to 7.0 Update 3L or 8.0 Update 2, otherwise the OpenShift Container Platform upgrade is blocked. If you do not want to update vSphere, you can proceed with an OpenShift Container Platform upgrade by performing an administrator acknowledgment. However, with using the administrator acknowledgment, known issues can occur. Before proceeding with the administrator acknowledgement, carefully read the [Knowledge Base article](#).

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

1.3.12.7. Secrets Store CSI Driver Operator (Technology Preview)

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

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

1.3.12.8. Azure File supporting NFS is generally available

OpenShift Container Platform 4.14 supports Azure File Container Storage Interface (CSI) Driver Operator with Network File System (NFS) as generally available.

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

1.3.13. Oracle® Cloud Infrastructure

You can now install an OpenShift Container Platform cluster on Oracle® Cloud Infrastructure (OCI) by using the Assisted installer or the Agent-based installer. To install an OpenShift Container Platform cluster on OCI, choose one of the following installation options:

- [Using the Assisted Installer to install a cluster on Oracle® Cloud Infrastructure \(OCI\)](#)
- [Using the Agent-based Installer to install a cluster on Oracle® Cloud Infrastructure \(OCI\)](#)

1.3.14. Operator lifecycle

1.3.14.1. Operator Lifecycle Manager (OLM) 1.0 (Technology Preview)

Operator Lifecycle Manager (OLM) has been included with OpenShift Container Platform 4 since its initial release. OpenShift Container Platform 4.14 introduces components for a next-generation iteration of OLM as a Technology Preview feature, known during this phase as *OLM 1.0*. This updated framework evolves many of the concepts that have been part of previous versions of OLM and adds new capabilities.

During this Technology Preview phase of OLM 1.0 in OpenShift Container Platform 4.14, administrators can explore the following features:

Fully declarative model that supports GitOps workflows

OLM 1.0 simplifies Operator management through two key APIs:

- A new **Operator** API, provided as **operators.operatorframework.io** by the new Operator Controller component, streamlines management of installed Operators by consolidating user-facing APIs into a single object. This empowers administrators and SREs to automate processes and define desired states by using GitOps principles.
- The **Catalog** API, provided by the new catalogd component, serves as the foundation for OLM 1.0, unpacking catalogs for on-cluster clients so that users can discover installable content, such as Operators and Kubernetes extensions. This provides increased visibility into all available Operator bundle versions, including their details, channels, and update edges.

For more information, see [Operator Controller](#) and [Catalogd](#).

Improved control over Operator updates

With improved insight into catalog content, administrators can specify target versions for installation and updates. This grants administrators more control over the target version of Operator updates. For more information, see [Installing an Operator from a catalog](#).

Flexible Operator packaging format

Administrators can use file-based catalogs to install and manage the following types of content:

- OLM-based Operators, similar to the existing OLM experience
- *Plain bundles*, which are static collections of arbitrary Kubernetes manifests

In addition, bundle size is no longer constrained by the etcd value size limit. For more information, see [Managing plain bundles in OLM 1.0](#).



NOTE

For OpenShift Container Platform 4.14, documented procedures for OLM 1.0 are CLI-based only. Alternatively, administrators can create and view related objects in the web console by using normal methods, such as the **Import YAML** and **Search** pages. However, the existing **OperatorHub** and **Installed Operators** pages do not yet display OLM 1.0 components.

For more information, see [About Operator Lifecycle Manager 1.0](#).

1.3.15. Operator development

1.3.15.1. Token authentication for Operators on cloud providers: AWS STS

With this release, Operators managed by Operator Lifecycle Manager (OLM) can support token authentication when running on Amazon Web Services (AWS) clusters that use the Security Token Service (STS). The Cloud Credential Operator (CCO) is updated to semi-automate provisioning certain limited-privilege, short-term credentials, provided that the Operator author has enabled their Operator to support AWS STS. For more information about enabling OLM-based Operators to support CCO-based workflows with AWS STS, see [Token authentication for Operators on cloud providers](#).

1.3.15.2. Configuring Operator projects with support for multiple platforms

With this release, Operator authors can configure their Operator projects with support for multiple architectures and operating systems, or *platforms*. Operator authors can configure support for multiple platforms by performing the following actions:

- Building a manifest list that specifies the platforms that the Operator supports.
- Setting the Operator's node affinity to support multi-architecture compute machines.

For more information, see [Configuring Operator projects for multi-platform support](#).

1.3.16. Builds

- With this update, the Source-to-Image (S2I) tool is now generally available in OpenShift Container Platform 4.14. You can use the S2I tool to build container images from source code, and transform application code into ready-to-deploy container images. This feature enhances the platform's ability to support reproducible containerized application development. For more information, see [Using Source-to-Image \(S2I\) tool](#).
- With this update, the Build CSI Volumes feature is now generally available in OpenShift Container Platform 4.14.

1.3.17. Machine Config Operator

1.3.17.1. Handling of registry certificate authorities

The Machine Config Operator now handles distributing certificate authorities for image registries. This change does not affect end users.

1.3.17.2. Additional metrics available in Prometheus

With this release, you can query additional metrics to more closely monitor the state of your machines and machine config pools.

For more information on how to use Prometheus, see [Viewing a list of available metrics](#).

1.3.17.3. Support for offline Tang provisioning

With this release, you can now provision an OpenShift Container Platform cluster with Tang-enforced, network-bound disk encryption (NBDE) by using Tang servers that are unreachable during first boot.

For more information, see [Configuring an encryption threshold](#) and [Configuring disk encryption and mirroring](#).

1.3.17.4. Certificates are now handled by the Machine Config Daemon

In previous OpenShift Container Platform versions, the MCO read and handled certificates directly from machine configuration files. This led to rotation issues and created unwanted situations, such as certificates getting stuck behind a paused machine config pool.

With this release, certificates are no longer templated from bootstrap into machine configuration files. Instead, they are put directly into the Ignition object, written onto a disk using the controller config, and handled by the Machine Config Daemon (MCD) during regular cluster operation. The certs are then visible by using the **ControllerConfig** resource.

The Machine Config Controller (MCC) holds the following certificate data:

- `/etc/kubernetes/kubelet-ca.crt`
- `/etc/kubernetes/static-pod-resources/configmaps/cloud-config/ca-bundle.pem`
- `/etc/pki/ca-trust/source/anchors/openshift-config-user-ca-bundle.crt`

The MCC also handles the image registry certificates and its associated user bundle certificate. This means that certificates are not bound by the machine config pool status and are more timely in their rotation. The previously listed CAs stored in machine configuration files are removed, and the templated files found during cluster installation no longer exist. For more information on how to access these certificates, see [Viewing and interacting with certificates](#).

1.3.18. Machine API

1.3.18.1. Support for control plane machine sets on Nutanix clusters

With this release, control plane machine sets are supported for Nutanix clusters. For more information, see [Getting started with the Control Plane Machine Set Operator](#).

1.3.18.2. Support for control plane machine sets on RHOSP clusters

With this release, control plane machine sets are supported for clusters that run on RHOSP.

For more information, see [Getting started with the Control Plane Machine Set Operator](#).



NOTE

For clusters that have root volume availability zones and are running on RHOSP that you upgrade to 4.14, you must converge control plane machines onto one server group before you can enable control plane machine sets. To make the required change, follow the instructions in [OpenShift on OpenStack with Availability Zones: Invalid Compute ServerGroup setup during OpenShift deployment](#).

For clusters that have compute zones configured with at least one zone and are running on RHOSP, which is upgradable to version 4.14, root volumes must now also be configured with at least one zone. If this configuration change does not occur, a control plane machine set cannot be generated for your cluster. To make the required change, follow the instructions in the related [OpenShift on OpenStack with compute Availability Zones: Missing rootVolume availability zone](#).

1.3.18.3. Support for assigning AWS machines to placement groups

With this release, you can configure a machine set to deploy machines within an existing AWS placement group. You can use this feature with Elastic Fabric Adapter (EFA) instances to improve network performance for machines within the specified placement group. You can use this feature with [compute](#) and [control plane](#) machine sets.

1.3.18.4. Azure confidential VMs and trusted launch (Technology Preview)

With this release, you can configure a machine set to deploy machines that use Azure confidential VMs, trusted launch, or both. These machines can use Unified Extensible Firmware Interface (UEFI) security features such as Secure Boot or a dedicated virtual Trusted Platform Module (vTPM) instance.

You can use this feature with [compute](#) and [control plane](#) machine sets.

1.3.19. Nodes

1.3.19.1. Descheduler resource limits for large clusters

With this release, the resource limits for the descheduler operand are removed. This enables the descheduler to be used for large clusters with many nodes and pods without failing due to out-of-memory errors.

1.3.19.2. Pod topology spread constraints `matchLabelKeys` parameter is now generally available

The **`matchLabelKeys`** parameter for configuring pod topology spread constraints is now generally available in OpenShift Container Platform 4.14. Previously, the parameter was available as a Technology Preview feature by enabling the **`TechPreviewNoUpgrade`** feature set. The **`matchLabelKeys`** parameter takes a list of pod label keys to select the pods to calculate spreading over.

For more information, see [Controlling pod placement by using pod topology spread constraints](#) .

1.3.19.3. `MaxUnavailableStatefulSet` enabled (Technology Preview)

With this release, the **`MaxUnavailableStatefulSet`** featureSet configuration parameter is available as Technology Preview feature by enabling the **`TechPreviewNoUpgrade`** feature set. You can now define the maximum number of **`StatefulSet`** pods that can be unavailable during updates, thereby reducing application downtime when upgrading.

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

1.3.19.4. Pod disruption budget (PDB) unhealthy pod eviction policy

With this release, specifying an unhealthy pod eviction policy for pod disruption budgets (PDBs) is Generally Available in OpenShift Container Platform and has been removed from the **`TechPreviewNoUpgrade`** featureSet. This helps evict malfunctioning applications during a node drain.

For more information, see [Specifying the eviction policy for unhealthy pods](#) .

1.3.19.5. Linux Control Groups version 2 is now default

Beginning with OpenShift Container Platform 4.14, new installs use Control Groups version 2 by default, also known as `cgroup v2`, `cgroup2`, or `cgroupsv2`. This enhancement includes many bug fixes, performance improvements, and the ability to integrate with new features. `cgroup v1` is still used in

upgraded clusters that have initial installation dates prior to OpenShift Container Platform 4.14. `cgroup v1` can still be used by changing the `cgroupMode` field in the `node.config` object to `v1`.

For more information, see [Configuring the Linux cgroup version on your nodes](#).

1.3.19.6. Cron job time zones general availability

Setting a time zone for a cron job schedule is now generally available. 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](#).

1.3.20. Monitoring


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

1.3.20.1. Updates to monitoring stack components and dependencies

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

- **kube-state-metrics** to 2.9.2
- **node-exporter** to 1.6.1
- **prom-label-proxy** to 0.7.0
- Prometheus to 2.46.0
- **prometheus-operator** to 0.67.1

1.3.20.2. Changes to alerting rules



NOTE

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

- **New**
 - Added the **KubeDeploymentRolloutStuck** alert to monitor if the rollout of a deployment has not progressed for 15 minutes.
 - Added the **NodeSystemSaturation** alert to monitor resource saturation on a node.
 - Added the **NodeSystemdServiceFailed** alert to monitor the systemd service on a node.
 - Added the **NodeMemoryMajorPagesFaults** alert to monitor major page faults on a node.
 - Added the **PrometheusSDRefreshFailure** alert to monitor failed Prometheus service discoveries.
- **Changed**
 - Modified the **KubeAggregatedAPIDown** alert and the **KubeAggregatedAPIErrors** alert to evaluate only metrics from the **apiserver** job.

- Modified the **KubeCPUOvercommit** alert to evaluate only metrics from the **kube-state-metrics** job.
- Modified the **NodeHighNumberContrackEntriesUsed**, **NodeNetworkReceiveErrs** and **NodeNetworkTransmitErrs** alerts to evaluate only metrics from the **node-exporter** job.
- **Removed**
 - Removed the **MultipleContainersOOMKilled** alert for not being actionable. Nodes under memory pressure are covered by other alerts.

1.3.20.3. New option to create alerts based on core platform metrics

With this release, administrators can create new alerting rules based on core platform metrics. You can now modify settings for existing platform alerting rules by adjusting thresholds and by changing labels. You can also define and add new custom alerting rules by constructing a query expression based on core platform metrics in the **openshift-monitoring** namespace. This feature was included as a Technology Preview feature in the OpenShift Container Platform 4.12 release, and the feature is now generally available in OpenShift Container Platform 4.14. For more information, see [Managing alerting rules for core platform monitoring](#).

1.3.20.4. New option to specify resource limits for all monitoring components

With this release, you can now specify resource requests and limits for all monitoring components, including the following:

- Alertmanager
- **kube-state-metrics**
- **monitoring-plugin**
- **node-exporter**
- **openshift-state-metrics**
- Prometheus
- Prometheus Adapter
- Prometheus Operator and its admission webhook service
- Telemeter Client
- Thanos Querier
- Thanos Ruler

In previous versions of OpenShift Container Platform, you could only set options for Prometheus, Alertmanager, Thanos Querier, and Thanos Ruler.

1.3.20.5. New options to configure node-exporter collectors

With this release, you can customize Cluster Monitoring Operator (CMO) config map settings for additional **node-exporter** collectors. The following **node-exporter** collectors are now optional, and you can enable or disable each one individually in the config map settings:

- **ksmd** collector
- **mountstats** collector
- **processes** collector
- **systemd** collector

In addition, you can now exclude network devices from the relevant collector configuration for the **netdev** and **netclass** collectors. You can also now use the **maxProcs** option to set the maximum number of processes that can run node-exporter.

1.3.20.6. New option to deploy monitoring web console plugin resources

With this release, the monitoring pages in the **Observe** section of the OpenShift Container Platform web console are deployed as a [dynamic plugin](#). With this change, the Cluster Monitoring Operator (CMO) is now the component that deploys the OpenShift Container Platform web console monitoring plugin resources. You can now use CMO settings to configure the following features of the console monitoring plugin resource:

- Node selectors
- Tolerations
- Topology spread constraints
- Resource requests
- Resource limits

1.3.21. Network Observability Operator

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

1.3.22. Scalability and performance

1.3.22.1. PAO must-gather image added to default must-gather image

With this release, the Performance Addon Operator (PAO) must-gather image is no longer required as an argument for the **must-gather** command to capture debugging data related to low-latency tuning. The functions of the PAO must-gather image are now under the default plugin image used by the **must-gather** command without any image arguments. For further information about gathering debugging information relating to low-latency tuning, see [Collecting low latency tuning debugging data for Red Hat Support](#).

1.3.22.2. Collecting data for the NUMA Resources Operator with the must-gather image of the Operator

In this release, the **must-gather** tool is updated to collect the data of the NUMA Resources Operator with the **must-gather** image of the Operator. For further information about gathering debugging information for the NUMA Resources Operator, see [Collecting NUMA Resources Operator data](#).

1.3.22.3. Enabling more control over the C-states for each pod

With this release, you have more control over the C-states for your pods. Now, instead of disabling C-states completely, you can specify a maximum latency in microseconds for C-states. You can configure this option in the **cpu-c-states.crio.io** annotation. This helps to optimize power savings in high-priority applications by enabling some of the shallower C-states instead of disabling them completely. For further information about controlling pod C-states, see [Disabling power saving mode for high priority pods](#).

1.3.22.4. Support for provisioning IPv6 spoke clusters from dual-stack hub clusters

With this update, you can provision IPv6 address spoke clusters from dual-stack hub clusters. In a zero touch provisioning (ZTP) environment, the HTTP server on the hub cluster that hosts the boot ISO now listens on both IPv4 and IPv6 networks. The provisioning service also checks the baseboard management controller (BMC) address scheme on the target spoke cluster and provides a matching URL for the installation media. These updates offer the ability to provision single-stack, IPv6 spoke clusters from a dual-stack hub cluster.

1.3.22.5. Dual-stack networking for RHOSP clusters (Technology Preview)

Dual-stack network configuration is now available for clusters that run on RHOSP. This is a Technology Preview feature. You can configure dual-stack networking during the deployment of a cluster on installer-provisioned infrastructure.

For more information, see [Configuring a cluster with dual-stack networking](#).

1.3.22.6. Security group management for RHOSP clusters

In OpenShift Container Platform 4.14, security for clusters that run on RHOSP is enhanced. By default, the OpenStack cloud provider now sets the **manage-security-groups** option for load balancers to **true**, ensuring that only node ports that are required for cluster operation are open. Previously, security groups for both compute and control plane machines were configured to open a wide range of node ports for all incoming traffic.

You can opt to use the previous configuration by setting the **manage-security-groups** option to **false** in the configuration of a load balancer and ensuring that the security group rules permit traffic from **0.0.0.0/0** on the node ports range 30000 through 32767.

For clusters that are upgraded to 4.14, you must manually remove permissive security group rules that open the deployment to all traffic. For example, you must remove a rule that permits traffic from **0.0.0.0/0** on the node ports range 30000 through 32767.

1.3.22.7. Using custom CRs with PolicyGenTemplate CRs in the GitOps Zero Touch Provisioning (ZTP) pipeline

You can now use GitOps ZTP to include custom CRs in addition to the base source CRs provided by the GitOps ZTP plugin in the **ztp-site-generate** container. For more information, see [Adding custom content to the GitOps ZTP pipeline](#).

1.3.22.8. GitOps ZTP independence from managed cluster version

You can now use GitOps ZTP to provision managed clusters that are running different versions of OpenShift Container Platform. This means that the hub cluster and the GitOps ZTP plugin version can be independent of the version of OpenShift Container Platform running on the managed clusters. For more information, see [Preparing the GitOps ZTP site configuration repository for version independence](#).

1.3.22.9. Pre-caching user-specified images with Topology Aware Lifecycle Manager

With this release, you can precache your application workload images before upgrading your applications on single-node OpenShift clusters with Topology Aware Lifecycle Manager. For more information, see [Pre-caching user-specified images with TALM on single-node OpenShift clusters](#).

1.3.22.10. Disk cleaning option through SiteConfig and GitOps ZTP

With this release, you can remove the partitioning table before installation by using the **automatedCleaningMode** field in the **SiteConfig** CR. For more information, see [Single-node OpenShift SiteConfig CR installation reference](#).

1.3.22.11. Support for adding custom node labels in the SiteConfig CR through GitOps ZTP

With this update, you can add the **nodeLabels** field in the **SiteConfig** CR to create custom roles for nodes in managed clusters. For more information about how to add custom labels, see [Deploying a managed cluster with SiteConfig and GitOps ZTP](#), [Generating GitOps ZTP installation and configuration CRs manually](#), and [single-node OpenShift SiteConfig CR installation reference](#).

1.3.22.12. Support for tuning etcd latency tolerances (Technology Preview)

With this release, you can set the control plane hardware speed to one of **"Standard"**, **"Slower"**, or the default, **""**, which allows the system to decide which speed to use. This is a Technology Preview feature. For more information, see [Setting tuning parameters for etcd](#).

1.3.22.13. Deprecation of fields in SiteConfig

With this update, the **apiVIP** and **ingressVIP** fields in the **SiteConfig** custom resource definition (CRD) are deprecated in favor of the plural forms, **apiVIPs** and **ingressVIPs**.

1.3.23. Hosted control planes

1.3.23.1. Hosted control planes is Generally Available on bare metal and OpenShift Virtualization

Hosted control planes for OpenShift Container Platform is now Generally Available on the bare-metal and OpenShift Virtualization platforms. Hosted control planes on AWS remains a Technology Preview feature.

1.3.23.2. Creating ARM NodePool objects on AWS hosted clusters (Technology Preview)

In this release, you can schedule application workloads on 64-bit ARM and AMD64 from the same hosted control plane. For more information, see [Creating ARM NodePool objects on AWS hosted clusters](#).

1.3.23.3. Hosted control planes on IBM Z (Technology Preview)

In this release, hosted control planes is available as a Technology Preview feature on IBM Z. For more information, see [Configuring the hosting cluster on 64-bit x84 bare metal for IBM Z compute nodes \(Technology Preview\)](#).

1.3.23.4. Hosted control planes on IBM Power (Technology Preview)

In this release, hosted control planes is available as a Technology Preview feature on IBM Power. For more information, see [Configuring the hosting cluster on a 64-bit x86 OpenShift Container Platform cluster to create hosted control planes for IBM Power compute nodes \(Technology Preview\)](#).

1.3.24. Insights Operator

1.3.24.1. On demand data gathering (Technology Preview)

In OpenShift Container Platform 4.14, Insights Operator can now run gather operations on demand. For more information about running gather operations on demand, see [Running an Insights Operator gather operation](#).

1.3.24.2. Running gather operations as individual pods (Technology Preview)

In OpenShift Container Platform 4.14 Technology Preview clusters, Insights Operator runs gather operations in individual pods. This supports the new on demand data gathering feature.

1.4. NOTABLE TECHNICAL CHANGES

OpenShift Container Platform 4.14 introduces the following notable technical changes.

1.4.1. Cloud controller managers for additional cloud providers

The Kubernetes community plans to deprecate the use of the Kubernetes controller manager to interact with underlying cloud platforms in favor of using cloud controller managers. As a result, there is no plan to add Kubernetes controller manager support for any new cloud platforms.

This release introduces the General Availability of using cloud controller managers for Amazon Web Services and Microsoft Azure.

To learn more about the cloud controller manager, see the [Kubernetes Cloud Controller Manager documentation](#).

To manage the cloud controller manager and cloud node manager deployments and lifecycles, use the Cluster Cloud Controller Manager Operator. For more information, see the [Cluster Cloud Controller Manager Operator](#) entry in the *Platform Operators reference*.

1.4.2. 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.

1.4.3. Change in SSH key location

OpenShift Container Platform 4.14 introduces a RHEL 9.2 based RHCOS. Before this update, SSH keys were located in **/home/core/.ssh/authorized_keys** on RHCOS. With this update, on RHEL 9.2 based RHCOS, SSH keys are located in **/home/core/.ssh/authorized_keys.d/ignition**.

If you customized the default OpenSSH **/etc/ssh/sshd_config** server configuration file, you must update it according to this [Red Hat Knowledgebase article](#).

1.4.4. cert-manager Operator general availability

The cert-manager Operator for Red Hat OpenShift 1.11 is now generally available in OpenShift Container Platform 4.14 and OpenShift Container Platform 4.13 and OpenShift Container Platform 4.12.

1.4.5. Improved scaling and stability with Open Virtual Network (OVN) Optimizations

OpenShift Container Platform 4.14 introduces an optimization of Open Virtual Network Kubernetes (OVN-K) in which its internal architecture was modified to reduce operational latency to remove barriers to scale and performance of the networking control plane. Network flow data is now localized to cluster nodes instead of centralizing information on the control plane. This reduces operational latency and reduces cluster-wide traffic between worker and control nodes. As a result, cluster networking scales linearly with node count, because additional networking capacity is added with each additional node, which optimizes larger clusters. Because network flow is localized on every node, RAFT leader election of control plane nodes is no longer needed, and a primary source of instability is removed. An additional benefit to localized network flow data is that the effect of node loss on networking is limited to the failed node and has no bearing on the rest of the cluster's networking, thereby making the cluster more resilient to failure scenarios. For more information, see [OVN-Kubernetes architecture](#).

1.4.6. Operator SDK 1.31.0

OpenShift Container Platform 4.14 supports Operator SDK 1.31.0. See [Installing the Operator SDK CLI](#) to install or update to this latest version.



NOTE

Operator SDK 1.31.0 supports Kubernetes 1.26.

If you have Operator projects that were previously created or maintained with Operator SDK 1.28.0, update your projects to keep compatibility with Operator SDK 1.31.0.

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

1.4.7. **oc** commands now default to storing and obtaining credentials from Podman configuration locations

Previously, OpenShift CLI (**oc**) commands that used the registry configuration, for example **oc adm release** or **oc image** commands, obtained credentials from Docker configuration file locations, such as `~/.docker/config.json`, first. If a registry entry could not be found in the Docker configuration locations, **oc** commands obtained the credentials from Podman configuration file locations, such as `${XDG_RUNTIME_DIR}/containers/auth.json`.

With this release, **oc** commands now default to obtaining the credentials from Podman configuration locations first. If a registry entry cannot be found in the Podman configuration locations, **oc** commands obtain the credentials from Docker configuration locations.

Additionally, the **oc registry login** command now stores credentials in the Podman configuration locations instead of the Docker configuration file locations.

1.4.8. Long-running pod requests are now recorded as **CONNECT** requests

In OpenShift Container Platform metrics, the pod **attach**, **exec**, **log**, **portforward**, and **proxy** requests are now recorded as **CONNECT** requests.

1.4.9. Open Virtual Network Infrastructure Controller default range

With this update, the Controller uses **100.88.0.0/16** as the default IP address range for the transit switch subnet. Do not use this IP range in your production infrastructure network. ([OCPBUGS-20261](#))

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 4.14, 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*

1.5.1. Operator lifecycle and development deprecated and removed features

Table 1.6. Operator lifecycle and development deprecated and removed tracker

Feature	4.12	4.13	4.14
SQLite database format for Operator catalogs	Deprecated	Deprecated	Deprecated
operators.openshift.io/infrastructure-features annotations	General Availability	General Availability	Deprecated

1.5.2. Images deprecated and removed features

Table 1.7. Images deprecated and removed tracker

Feature	4.12	4.13	4.14
ImageChangesInProgress condition for Cluster Samples Operator	Deprecated	Deprecated	Deprecated
MigrationInProgress condition for Cluster Samples Operator	Deprecated	Deprecated	Deprecated

1.5.3. Installation deprecated and removed features

Table 1.8. Installation deprecated and removed tracker

Feature	4.12	4.13	4.14
--cloud parameter for oc adm release extract	General Availability	General Availability	Deprecated
CoreDNS wildcard queries for the cluster.local domain	Deprecated	Removed	Removed
compute.platform.openstack.rootVolume.type for RHOSP	General Availability	General Availability	Deprecated
controlPlane.platform.openstack.rootVolume.type for RHOSP	General Availability	General Availability	Deprecated
ingressVIP and apiVIP settings in the install-config.yaml file for installer-provisioned infrastructure clusters	Deprecated	Deprecated	Deprecated
platform.gcp.licenses for Google Cloud Provider	Deprecated	Deprecated	Removed

Feature	4.12	4.13	4.14
VMware ESXi 7.0 Update 1 or earlier	General Availability	Removed ^[1]	Removed
vSphere 7.0 Update 1 or earlier	Deprecated	Removed ^[1]	Removed

1. For OpenShift Container Platform 4.14, you must install the OpenShift Container Platform cluster on a VMware vSphere version 7.0 Update 2 or later instance, including VMware vSphere version 8.0, that meets the requirements for the components that you use.

1.5.4. Storage deprecated and removed features

Table 1.9. Storage deprecated and removed tracker

Feature	4.12	4.13	4.14
Persistent storage using FlexVolume	Deprecated	Deprecated	Deprecated

1.5.5. Building applications deprecated and removed features

Table 1.10. Service Binding Operator deprecated and removed tracker

Feature	4.12	4.13	4.14
Service Binding Operator	General Availability	Deprecated	Deprecated

1.5.6. Multi-architecture deprecated and removed features

Table 1.11. Multi-architecture deprecated and removed tracker

Feature	4.12	4.13	4.14
IBM Power8 all models (ppc64le)	Deprecated	Removed	Removed
IBM Power® AC922 (ppc64le)	Deprecated	Removed	Removed
IBM Power® IC922 (ppc64le)	Deprecated	Removed	Removed
IBM Power® LC922 (ppc64le)	Deprecated	Removed	Removed
IBM z13 all models (s390x)	Deprecated	Removed	Removed
IBM® LinuxONE Emperor (s390x)	Deprecated	Removed	Removed

Feature	4.12	4.13	4.14
IBM® LinuxONE Rockhopper (s390x)	Deprecated	Removed	Removed
AMD64 (x86_64) v1 CPU	Deprecated	Removed	Removed

1.5.7. Networking deprecated and removed features

Table 1.12. Networking deprecated and removed tracker

Feature	4.12	4.13	4.14
Kuryr on RHOSP	Deprecated	Deprecated	Deprecated
OpenShift SDN network plugin	General Availability	General Availability	Deprecated

1.5.8. Node deprecated and removed features

Table 1.13. Node deprecated and removed tracker

Feature	4.12	4.13	4.14
ImageContentSourcePolicy (ICSP) objects	General Availability	Deprecated	Deprecated
Kubernetes topology label failure-domain.beta.kubernetes.io/zone	General Availability	Deprecated	Deprecated
Kubernetes topology label failure-domain.beta.kubernetes.io/region	General Availability	Deprecated	Deprecated

1.5.9. OpenShift CLI (oc) deprecated and removed features

Feature	4.12	4.13	4.14
--include-local-oci-catalogs parameter for oc-mirror	Not Available	General Availability	Removed
--use-oci-feature parameter for oc-mirror	General Availability	Deprecated	Removed

1.5.10. Workloads deprecated and removed features

Table 1.14. Workloads deprecated and removed tracker

Feature	4.12	4.13	4.14
DeploymentConfig objects	General Availability	General Availability	Deprecated

1.5.11. Bare metal monitoring deprecated and removed features

Table 1.15. Bare Metal Event Relay Operator tracker

Feature	4.14	4.15	4.16
Bare Metal Event Relay Operator	Removed	Removed	Removed

1.5.12. Deprecated features

1.5.12.1. Deprecation of the OpenShift SDN network plugin

OpenShift SDN CNI is deprecated as of OpenShift Container Platform 4.14. It is currently planned that the network plugin will not be an option for new installations in the next minor release of OpenShift Container Platform. In a subsequent future release, the OpenShift SDN network plugin is planned to be removed and no longer supported. Red Hat will provide bug fixes and support for this feature until removed, but this feature will no longer receive enhancements. As an alternative to OpenShift SDN CNI, you can use OVN Kubernetes CNI instead.

1.5.12.2. Service Binding Operator

The Service Binding Operator is deprecated and will be removed with the OpenShift Container Platform 4.16 release. Red Hat will provide critical bug fixes and support for this component during the current release lifecycle, but this component will no longer receive feature enhancements.

1.5.12.3. DeploymentConfig resources are now deprecated

As of OpenShift Container Platform 4.14, **DeploymentConfig** objects are deprecated.

DeploymentConfig objects are still supported, but are not recommended for new installations. Only security-related and critical issues will be fixed.

Instead, use **Deployment** objects or another alternative to provide declarative updates for pods.

1.5.12.4. Operator-specific CatalogSource CRs used in GitOps ZTP are deprecated

From OpenShift Container Platform 4.14, you must only use the **DefaultCatSrc.yaml** **CatalogSource** CR when updating Operators with Topology Aware Lifecycle Manager (TALM). All other **CatalogSource** CRs are deprecated and are planned to be removed in a future release. Red Hat will provide bug fixes and support for this feature during the current release lifecycle, but this feature will no longer receive enhancements and will be removed. For more information about **DefaultCatSrc** CR, see [Performing an Operator update](#).

1.5.12.5. The --cloud parameter for the oc adm release extract command

As of OpenShift Container Platform 4.14, the **--cloud** parameter for the **oc adm release extract** command is deprecated. The introduction of the **--included** and **--install-config** parameters make the **-cloud** parameter unnecessary.

For more information, see [Simplified installation and update experience for clusters with manually maintained cloud credentials](#).

1.5.12.6. Red Hat Virtualization (RHV) as a host platform for OpenShift Container Platform

Red Hat Virtualization (RHV) as a host platform for OpenShift Container Platform was deprecated and is no longer supported. This platform will be removed from OpenShift Container Platform in a future OpenShift Container Platform release.

1.5.12.7. Using the REGISTRY_AUTH_PREFERENCE environment variable is now deprecated

Using the **REGISTRY_AUTH_PREFERENCE** environment variable to specify your preferred location to obtain registry credentials for OpenShift CLI (**oc**) commands is now deprecated.

OpenShift CLI (**oc**) commands now default to obtaining the credentials from Podman configuration locations first, but will fall back to checking the deprecated Docker configuration file locations.

1.5.12.8. operators.openshift.io/infrastructure-features annotations

Starting in OpenShift Container Platform 4.14, the **operators.openshift.io/infrastructure-features** group of annotations are deprecated by the group of annotations with the **features.operators.openshift.io** namespace.



NOTE

Currently, the web console continues to support displaying and filtering based on the earlier annotations. However, because they are deprecated, this support will be removed from the web console in a future OpenShift Container Platform release, and therefore migration to the new annotations format is advised.

See [Deprecated infrastructure feature annotations](#) for the earlier group of annotations, and see [Infrastructure features annotations](#) for the latest group.

1.5.13. Removed features

1.5.13.1. Bare Metal Event Relay Operator removed

The Bare Metal Event Relay Operator was previously a Technology Preview feature and is now removed from OpenShift Container Platform. For complete lifecycle information for the Bare Metal Event Relay Operator, see [Product Life Cycles: Bare Metal Event Relay](#).

1.5.13.2. Beta APIs removed from Kubernetes 1.27

Kubernetes 1.27 removed the following deprecated API, 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 1.16. APIs removed from Kubernetes 1.27

Resource	Removed API	Migrate to
CSISStorageCapacity	storage.k8s.io/v1beta1	storage.k8s.io/v1

1.5.13.3. Support for the **LatencySensitive** feature set is removed

As of OpenShift Container Platform 4.14, support for the **LatencySensitive** feature set is removed.

1.5.13.4. **oc registry login** no longer stores credentials in Docker configuration file locations

As of OpenShift Container Platform 4.14, the **oc registry login** command no longer stores registry credentials in the Docker file locations, such as `~/.docker/config.json`. The **oc registry login** command now stores credentials in the Podman configuration file locations, such as `${XDG_RUNTIME_DIR}/containers/auth.json`.

1.6. BUG FIXES

1.6.1. API Server and Authentication

- Previously, when creating a pod controller with a pod spec that would be mutated by security context constraints, users might get a warning that the pod did not meet the given namespace's pod security level. With this release, you no longer get a warning about pod security violations if the pod controller will create pods that do not violate pod security in that namespace. ([OCBUGS-7267](#))
- A **user:check-access** scoped token grants sufficient permissions to send a SelfSubjectAccessReview request. Previously, the cluster did not grant sufficient permissions to perform the access review unless the token also had the **user:full** scope or a role scope. With this release, the cluster authorizes a SelfSubjectAccessReview request as if it has either the full user's permissions or the permissions of the user's role set on the request in order to be able to perform the access review. ([OCBUGS-7415](#))
- Previously, the pod security admission controller required the **RoleBinding** object's **.subject[].namespace** field to be set when **.subjects[].kind** is set to **ServiceAccount** in order to successfully bind the service account to a role. With this release, the pod security admission controller uses the namespace of the **RoleBinding** object if the **.subject[].namespace** is not specified. ([OCBUGS-160](#))
- Previously, the **clientConfig** of all the webhooks of **ValidatingWebhookConfiguration** and **MutatingWebhookConfiguration** objects did not get a properly injected **caBundle** with the **service-ca** trust bundle. With this release, the **clientConfig** of all the webhooks of **ValidatingWebhookConfiguration** and **MutatingWebhookConfiguration** objects now get a properly injected **caBundle** with the **service-ca** trust bundle. ([OCBUGS-19318](#))
- Previously, kube-apiserver did not change to **Degraded=True** when an invalid secret name was specified for **servingCertificate** in **namedCertificates**. With this release, kube-apiserver now switches to **Degraded=True** and shows why the certificate was not accepted to allow for easier troubleshooting. ([OCBUGS-8404](#))
- Previously, observability dashboards used large queries to show data which caused frequent timeouts on clusters with a large number of nodes. With this release, observability dashboards use recording rules that are precalculated to ensure reliability on clusters with a large number of

nodes. ([OCPBUGS-3986](#))

1.6.2. Bare Metal Hardware Provisioning

- Previously, if the hostname of a bare-metal machine was not provided by either reverse DNS or DHCP, it would default to **localhost** during bare-metal cluster provisioning on installer-provisioned infrastructure. This issue caused Kubernetes node name conflicts and prevented the cluster from being deployed. Now, if the hostname is detected to be **localhost**, the provisioning agent sets the persistent hostname to the name of the **BareMetalHost** object. ([OCPBUGS-9072](#))

1.6.3. Cloud Compute

- Previously, the Machine API controller could not determine the zone of machines in vSphere clusters that use multiple zones. With this release, the zone lookup logic is based on the host of a VM and, as a result, machine objects indicate proper zones. ([OCPBUGS-7249](#))
- Previously, after the rotation of cloud credentials in the **clouds.yaml** file, the OpenStack machine API provider would need to be restarted in order to pick up the new cloud credentials. As a result, the ability of a machine set to scale to zero could be affected. With this change, cloud credentials are no longer cached, and the provider reads the corresponding secret freshly as needed. ([OCPBUGS-8687](#))
- Previously, some conditions during the startup process of the Cluster Autoscaler Operator caused a lock that prevented the Operator from successfully starting and marking itself as available. As a result, the cluster became degraded. The issue is resolved with this release. ([OCPBUGS-20038](#))
- Previously, the bootstrap credentials used to request client credentials for control plane nodes did not include the generic, all service accounts group. As a result, the cluster machine approver ignored certificate signing requests (CSRs) created during this phase. In certain conditions, this prevented approval of CSRs during bootstrap and caused the installation to fail. With this release, the bootstrap credential includes the groups that the cluster machine approver expects for a service account. This change allows the machine approver to take over from the bootstrap CSR approver earlier in the cluster lifecycle and should reduce bootstrap failures related to CSR approval. ([OCPBUGS-8349](#))
- Previously, if scaling the machines on a Nutanix cluster exceeded the available memory to complete the operation, machines would get stuck in the **Provisioning** state and could not be scaled up or down. The issue is resolved in this release. ([OCPBUGS-19731](#))
- Previously, for clusters on which the Control Plane Machine Set Operator is configured to use the **OnDelete** update strategy, cached information about machines caused the Operator to balance machines incorrectly and place them in an unexpected failure domain during reconciliation. With this release, the Operator refreshes this information immediately before creating new machines so that it correctly identifies the failure domains to place machines in. ([OCPBUGS-15338](#))
- Previously, the Control Plane Machine Set Operator used the **Infrastructure** object specification to determine the platform type for the cluster. For clusters upgraded from OpenShift Container Platform version 4.5 and earlier, this practice meant that the Operator could not correctly determine that a cluster was running on AWS, and therefore did not generate the **ControlPlaneMachineSet** custom resource (CR) as expected. With this release,

the Operator uses the status platform type, which is populated on all clusters independent of when they were created and is now able to generate the **ControlPlaneMachineSet** CR for all clusters. ([OCPBUGS-11389](#))

- Previously, machines created by a control plane machine set were considered ready once the underlying Machine API machine was running. With this release, the machine is not considered ready until the node linked to that machine is also ready. ([OCPBUGS-7989](#))
- Previously, the Control Plane Machine Set Operator prioritized failure domains alphabetically and moved machines from alphabetically later failure domains to alphabetically earlier failure domains, even if doing so did not improve the availability of the machines across the failure domains. With this release, the Operator is updated to prioritize failure domains that are present in the existing machines and to respect existing failure domains that provide better availability. ([OCPBUGS-7921](#))
- Previously, when a control plane machine on a vSphere cluster that uses a control plane machine set was deleted, sometimes two replacement machines were created. With this release, the control plane machine set no longer causes an extra machine to be created. ([OCPBUGS-7516](#))
- Previously, when the availability zone and subnet ID in a machine set were mismatched, a machine was created successfully by using the machine set specification with no indication to the user of the mismatch. Because the mismatched values can cause problems with some configurations, this occurrence might be visible as a warning message. With this release, a warning about the mismatch is logged. ([OCPBUGS-6882](#))
- Previously, when creating an OpenShift Container Platform cluster on Nutanix that uses Dynamic Host Configuration Protocol (DHCP) instead of an IP address management (IPAM) network configuration, the hostname of the VM was not set by DHCP. With this release, the VM hostname is set with values from the ignition configuration files. As a result, the issue is resolved for DHCP as well as other network configuration types. ([OCPBUGS-6727](#))
- Previously, multiple clusters could be created in the **openshift-cluster-api** namespace. This namespace must contain only one cluster. With this release, additional clusters cannot be created in this namespace. ([OCPBUGS-4147](#))
- Previously, clearing some parameters from the **providerSpec** field of a control plane machine set custom resource caused a loop of control plane machine deletion and creation. With this release, these parameters receive a default value if they are cleared or left empty, which resolves the issue. ([OCPBUGS-2960](#))


1.6.4. Cloud Credential Operator

- Previously, the Cloud Credential Operator utility (**ccoctl**) used an incorrect Amazon Resource Names (ARN) prefix for AWS GovCloud (US) and AWS China regions. The incorrect ARN prefix caused the **ccoctl aws create-all** command that is used to create AWS resources during installation to fail. This release updates the ARN prefixes to the correct values. ([OCPBUGS-13549](#))
- Previously, security changes to Amazon S3 buckets caused the Cloud Credential Operator utility (**ccoctl**) command that is used to create AWS resources during installation (**ccoctl aws create-all**) to fail. With this release, the **ccoctl** utility is updated to reflect the Amazon S3 security changes. ([OCPBUGS-11671](#))

1.6.5. Cluster Version Operator

- Previously, the Cluster Version Operator (CVO) did not reconcile **SecurityContextConstraints** resources as expected. The CVO now properly reconciles **SecurityContextConstraints** resources towards the state defined in the release image, reverting any unsupported modifications to them.
Users who want to upgrade from earlier OpenShift Container Platform versions and who operate workloads depending on modified system **SecurityContextConstraints** resources must follow the procedure in the [Knowledge Base article](#) to make sure their workloads are able to run without modified system **SecurityContextConstraint** resources. ([OCPBUGS-19465](#))
- Previously, the Cluster Version Operator did not prioritize likely targets when determining which conditional update risks to evaluate first. Now for conditional updates to which risks do not apply, these updates are available faster after Cluster Version Operator detection. ([OCPBUGS-5469](#))

1.6.6. Developer Console

- Previously, if you tried to edit a Helm chart repository in the **Developer** console by navigating to **Helm**, clicking the **Repositories** tab, then selecting **Edit HelmChartRepository** through the  menu for your Helm chart repository, an **Error** page displayed a **404: Page Not Found** error. This was caused by a component path that was not up to date. This issue is now fixed. ([OCPBUGS-14660](#))
- Previously, distinguishing between the types of samples listed in the **Samples** page was difficult. With this fix, you can easily identify the sample type from the badges displayed on the **Samples** page. ([OCPBUGS-7446](#))
- Previously on the Pipeline **Metrics** page, only four legends were visible for **TaskRun** duration charts. With this update, you can see all the legends present for the **TaskRun** duration charts. ([OCPBUGS-19878](#))
- Previously, an issue occurred when creating an application by using the **Import JAR** form in a disconnected cluster with the Cluster Samples Operator not installed. With this update, the **Import JAR** form from the **Add** page and the **Topology** page is hidden when the Java Builder Image is absent. ([OCPBUGS-15011](#))
- Previously, the Operator backed catalog did not show any catalog items if cluster service version (CSV) copies were disabled. With this fix, Operator backed catalogs are shown in every namespace even if CSV copies are disabled. ([OCPBUGS-14907](#))
- 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, the **Resource Type** section is moved to the **General** section. ([OCPBUGS-7395](#))

1.6.7. etcd Cluster Operator

- Previously, the **etcdctl** binary was cached on the local machine indefinitely, making updates to the binary impossible. The binary is now properly updated on every invocation of the **cluster-backup.sh** script. ([OCPBUGS-19499](#))

1.6.8. Installer

- Previously, if you did not specify a custom Red Hat Enterprise Linux CoreOS (RHCOS) Amazon Machine Image (AMI) when installing an AWS cluster to a supported secret partition, the

installation failed. With this update, the installation program validates that you have specified the ID of an RHCOS AMI in the installation configuration file before deploying the cluster. ([OCPBUGS-13636](#))

- Previously, the OpenShift Container Platform installation program did not find private hosted zones in the host project during installations on Google Cloud Platform (GCP) by using a shared VPC. With this update, the installation program checks for an existing private hosted zone in the host project and uses the private hosted zone if it exists. ([OCPBUGS-11736](#))
- Previously, if you configured user-defined outbound routing when installing a private Azure cluster, the cluster was incorrectly deployed with the default public load balancer. This behavior occurred when using the installer-provisioned infrastructure to install the cluster. With this update, the installation program no longer creates the public load balancer when user-defined routing is configured. ([OCPBUGS-9404](#))
- Previously, for clusters that run on RHOSP, in the deprovisioning phase of installation, the installer deleted object storage containers sequentially. This behavior caused slow and inefficient deletion of objects, especially with large containers. This problem occurred in part because image streams that use Swift containers accumulated objects over time. Now, bulk object deletion occurs concurrently with up to 3 calls to the RHOSP API, improving efficiency by handling a higher object count per call. This optimization speeds up resource cleanup during deprovisioning. ([OCPBUGS-9081](#))
- Previously, SSH access to bootstrap and cluster nodes failed when the bastion host ran in the same VPC network as the cluster nodes. Additionally, this configuration caused SSH access from the temporary bootstrap node to the cluster nodes to fail. These issues are now fixed by updating the IBM Cloud **SecurityGroupRules** to support SSH traffic between the temporary bootstrap node and cluster nodes, and to support SSH traffic from a bastion host to cluster nodes on the same VPC network. Log and debug information can be accurately collected for analysis during installer-provisioned infrastructure failure. ([OCPBUGS-8035](#))
- Previously, DNS records that the installation program created were not removed when uninstalling a private cluster. With this update, the installation program now correctly removes these DNS records. ([OCPBUGS-7973](#))
- Previously, a script provided in the documentation for checking invalid HTTPS certificates in the RHOSP API assumed a recent version of the RHOSP client. For users who did not have a recent version of the client, this script failed. Now, manual instructions are added to the documentation that users can follow to perform the check with any version of the client. ([OCPBUGS-7954](#))
- Previously, when defining static IP addresses in the **agent-config.yaml** or **nmstateconfig.yaml** files for the configuration of an Agent-based install, the configured static IP addresses might not have been configured during bootstrap. As a result, the host interfaces would choose an address through DHCP. With this update, timing issues are fixed to ensure that the configured static IP address is correctly applied to the host interface. ([OCPBUGS-16219](#))
- Previously, during an Agent-based installation, the certificates in the **AdditionalTrustBundle** field of the **install-config.yaml** file were only propagated to the final image when the **ImageContentSources** field was also set for mirroring. If mirroring was not set, the additional certificates were on the bootstrap but not the final image. This situation can cause issues when you have set up a proxy and want to add additional certificates as described in [Configuring the cluster-wide proxy during installation](#). With this update, these additional certificates are propagated to the final image whether or not the **ImageContentSources** field is also set. ([OCPBUGS-13535](#))

- Previously, the **openshift-install agent create** command did not return the help output when running an invalid command. With this update, the help output is now shown when you run an invalid **openshift-install agent create** command. ([OCPBUGS-10638](#))
- Previously, primary networks were not correctly set for generated machines that used Technology Preview failure domains. As a consequence, port targets with the ID **control-plane** were not set as the primary network on machines, which could cause installations that use Kuryr to function improperly. The field is now set to use the proper port target, if set. The primary network for generated machines is now set correctly, allowing installations that use Kuryr to complete. ([OCPBUGS-10570](#))
- Previously, when running the **openshift-install agent create image** command while using a **releasImage** that contained a digest, the command produced the following warning message: **WARNING The ImageContentSources configuration in install-config.yaml should have at least one source field matching the releasImage**. This message was produced every time, regardless of how **ImageContentSources** was configured, and could cause confusion. With this update, the warning message is only produced when **ImageContentSources** is legitimately not set to have at least one source field matching the release image. ([OCPBUGS-10207](#))
- Previously, when running the **openshift-install agent create image** command to generate a bootable ISO image, the command output provided a message indicating a successful generated image. This output message existed even if the Agent-based installer could not extract a base ISO image from the release image. With this update, the command output now produces an error message if the Agent-based Installer cannot locate the base ISO image, which might be indicative of an issue with **releasImage**. ([OCPBUGS-9949](#))
- Previously, shared VPC installations on GCP that used passthrough credentials mode could fail because the installation program used credentials from the default service account. With this update, you can specify another service account to use for node creation instead of the default. ([OCPBUGS-15421](#))
- Previously, if you defined more control plane nodes than compute nodes in either the **agent-config.yaml** or the **nmstateconfig.yaml** configuration file, you received a warning message. Now, if you specify this configuration in either file, you receive an error message, which indicates that compute nodes cannot exceed control plane nodes in either file. ([OCPBUGS-14877](#))
- Previously, an Agent-based installation would fail if a non-canonical IPv6 address was used for the **RendezvousIP** field in the **agent-config.yaml** file. Non-canonical IPv6 addresses contain leading zeros, for example, **2001:0db8:0000:0000:0000:0000:0000**. With this update, these valid addresses can now be used for the **RendezvousIP**. ([OCPBUGS-14121](#))
- Previously, the Operator cached the cloud credentials, which resulted in authentication issues when these credentials were rotated. Now, the Operator always uses the latest credentials. The Manila CSI Driver Operator now automatically creates an OpenShift storage class for each available Manila share type. As part of this operation, the Operator queries the Manila API. ([OCPBUGS-14049](#))
- Previously, when configuring the **install-config.yaml** file for use during an Agent-based installation, changing the **cpuPartitioning** field to a non-default value did not produce a warning to alert users that the field is ignored for Agent-based installations. With this update, changing the **cpuPartitioning** field causes a warning to users that the configuration does not impact the install. ([OCPBUGS-13662](#))
- Previously, installing an Azure cluster into an existing Azure Virtual Network (VNet) could fail because the installation program created a default network security group, which allowed traffic from **0.0.0.0**. The failure occurred when the existing VNet had the following rule enabled in the

tenant: **Rule: Network Security Groups shall not allow rule with 0.0.0.0/Any Source/Destination IP Addresses - Custom Deny**. With this fix, the installation program no longer creates the default network security group when installing a cluster into an existing VNet, and the installation succeeds. ([OCPBUGS-11796](#))

- During an installation, when the cluster status is **installing-pending-user-action**, the installation does not complete until the status is resolved. Previously, if you ran the **openshift-install agent wait-for bootstrap-complete** command, no indication existed of how to resolve the problem that caused this status. With this update, the command output provides a message indicating which actions must be taken to resolve the issue. ([OCPBUGS-4998](#))
For example, the **wait-for** output when an invalid boot disk is used is now as follows:

```
"level=info msg=Cluster has hosts requiring user input
level=debug msg=Host master-1 Expected the host to boot from disk, but it booted the
installation image - please reboot and fix boot order to boot from disk QEMU_HARDDISK
drive-scsi0-0-0-0 (sda, /dev/disk/by-path/pci-0000:03:00.0-scsi-0:0:0:0)
level=debug msg=Host master-2 Expected the host to boot from disk, but it booted the
installation image - please reboot and fix boot order to boot from disk QEMU_HARDDISK
drive-scsi0-0-0-0 (sda, /dev/disk/by-path/pci-0000:03:00.0-scsi-0:0:0:0)
level=info msg=cluster has stopped installing... working to recover installation"
```

- Previously, the **assisted-installer-controller** on the installed cluster would run continuously even after the cluster had completed installation. Because **assisted-service** runs on the bootstrap node and not on the cloud, and because the assisted-service goes offline after the bootstrap node reboots to join the cluster, the **assisted-installer-controller** was unable to communicate with assisted-service to post updates and upload logs and loops. Now, the **assisted-installer-controller** checks the cluster installation without using **assisted-service**, and exits when the cluster installation is complete. ([OCPBUGS-4240](#))
- Previously, installing a cluster to the AWS Commercial Cloud Services (C2S) **us-iso-east-1** region failed with an error message stating an **UnsupportedOperation**. With this fix, installing to this region now succeeds. ([OCPBUGS-2324](#))
- Previously, installations on AWS could fail because the installation program did not create the **cloud.conf** file with the necessary service endpoints in it. This led to the machine config operator creating an empty **cloud.conf** file that lacked the service endpoints, leading to an error. With this update, the installation program always creates the **cloud.conf** file so that the installation succeeds. ([OCPBUGS-20401](#))
- Previously, if you installed a cluster using the Agent-based installer and your pull secret had a null **auth** or **email** field, the installation would fail without providing a useful error. With this update, the **openshift-install agent wait-for install-complete** command validates your pull secret and notifies you if there are null fields. ([OCPBUGS-14405](#))
- Previously, the **create agent-config-template** command printed a line with **INFO** only, but no details about whether the command was successful and where the template file was written to. Now, if the command is successful, the command will print **INFO Created Agent Config Template in <path> directory**. ([OCPBUGS-13408](#))
- Previously, when a user specified the **vendor** hint in the **agent-config.yaml** file, the value was checked against the wrong field so that the hint would not match. With this update, the use of the **vendor** hint correctly selects a disk. ([OCPBUGS-13356](#))
- Previously, setting the **metadataService.authentication** field to **Required** when installing a

cluster on AWS did not configure the bootstrap VM to use IMDSv2 authentication. This could result in installations failing if you configured your AWS account to block IMDSv1 authentication. With this update, the **metadataService.authentication** field correctly configures the bootstrap VM to use IMDSv2 authentication when set to **Required**. ([OCPBUGS-12964](#))

- Previously, if you configured user-defined outbound routing when installing a private Azure cluster, the cluster was incorrectly deployed with the default public load balancer. This behavior occurred when using the installer-provisioned infrastructure to install the cluster. With this update, the installation program no longer creates the public load balancer when user-defined routing is configured. ([OCPBUGS-9404](#))
- Previously, the vSphere Terraform **vsphere_virtual_machine** resource did not include the **firmware** parameter. This issue caused the firmware of the VM to be set to **bios** by default instead of **efi**. Now, the resource includes the **firmware** parameter and sets **efi** as the default value for the parameter, so that the VM runs the Extensible Firmware Interface (EFI) instead of the basic input/output system (BIOS) interface. ([OCPBUGS-9378](#))
- Previously, for clusters that run on RHOSP, in the deprovisioning phase of installation, the installer deleted object storage containers sequentially. This behavior caused slow and inefficient deletion of objects, especially with large containers. This problem occurred in part because image streams that use Swift containers accumulated objects over time. Now, bulk object deletion now occurs concurrently with up to 3 calls to the RHOSP API, improving efficiency by handling a higher object count per call. This optimization speeds up resource cleanup during deprovisioning. ([OCPBUGS-9081](#))
- Previously, the installation program did not exit with an error if you installed a cluster on Azure by using disk encryption without providing a subscription ID. This caused the installation to begin, and then only to fail later on. With this update, the installation program requires you to specify a subscription ID for encrypted Azure installations and exits with an error if you do not provide one. ([OCPBUGS-8449](#))
- Previously, the Agent-based installer showed the results of secondary checks such as **ping** and **nslookup**, which can harmlessly fail even when the installation succeeds. This could result in errors being displayed despite the cluster installing successfully. With this update, secondary checks only display results if the primary installation checks fail, so that you can use the secondary checks to troubleshoot the failed installation. ([OCPBUGS-8390](#))
- Using an IPI **install-config** with the Agent-based Installer results in warning log messages showing the contents of any unused fields. Previously, these warnings printed sensitive information such as passwords. With this update, the warning messages for the credentials fields in the **vsphere** and **baremetal** platform sections have been changed to avoid logging any sensitive information. ([OCPBUGS-8203](#))
- Previously, clusters on Azure Stack Hub could not create new control plane nodes unless the nodes had custom disk sizes, because the default disk size could not be validated. With this update, the default disk size has been set to 128 GB and the installation program enforces user-specified disk size values between 128 and 1023 GB. ([OCPBUGS-6759](#))
- Previously, the installation program used port 80 to provide images to the Baseboard Management Controller (BMC) and the deployment agent when installing on bare metal with installer-provisioned infrastructure. This could present security concerns because many types of public traffic use port 80. With this update, the installation program uses port 6180 for this purpose. ([OCPBUGS-8509](#))

1.6.9. Machine Config Operator

- Previously, OpenShift Container Platform clusters that were installed on AWS used 4.1 boot images that were not able to scale up. This issue occurred because two systemd units, configured from Ignition then rendered and launched by the MCO during the initial boot of a new machine, have a dependency on the application Afterburn. Because OpenShift Container Platform 4.1 boot images do not contain Afterburn, this issue prevented new nodes from being able to join the cluster. Now, **systemd** units contain an additional check for Afterburn along with fallback code that does not rely on the presence of Afterburn. ([OCPBUGS-7559](#))

1.6.10. Management Console

- Previously, alerts loaded from non-Prometheus datasources such as logs. This caused the source of all alerts to be displayed always as **Prometheus**. With this update, alert sources are displayed correctly. ([OCPBUGS-9907](#))
- Previously, there was an issue with Patternfly 4 where you could not select or change the log component under the logs section of master node once a selection was already made. With this update, when you change to the log component from the log section of the master node, refresh the page to reload the default options. ([OCPBUGS-18727](#))
- Previously, an empty page was displayed when viewing route details on the **Metrics** tab of the **alertmanager-main** page. With this update, user privileges were updated so you can view the route details on the **Metrics** tab. ([OCPBUGS-15021](#))
- Previously, Red Hat OpenShift Service on AWS used custom branding and the favicon would disappear so no specific branding appeared when custom branding was being used. With this update, Red Hat OpenShift Service on AWS branding is now part of the branding API. ([OCPBUGS-14716](#))
- Previously, the OpenShift Container Platform web console did not render the monitoring **Dashboard** page when a proxy was expected. As a result, the websocket connection failed. With this update, the web console also detects proxy settings from environment variables. ([OCPBUGS-14550](#))
- Previously, if the **console.openshift.io/disable-operand delete: "true"** and **operator.openshift.io/uninstall-message: "some message"** annotations were used on an operator CSV, the uninstall instructions did not show up in the web console. With this update, the instructions to opt out of the installment are available. ([OCPBUGS-13782](#))
- Previously, the size on the **PersistentVolumeClaims** namespace **Details** page was incorrect. With this update, the Prometheus query on **PersistentVolumeClaims** namespace **Details** page includes the namespace label and the size is now correct. ([OCPBUGS-13208](#))
- Previously, after customizing the routes for console and downloads, the downloads route did not update in the **ConsoleCLIDownloads** link and pointed to the default downloads route. With this update, the **ConsoleCLIDownloads** link updates when the custom downloads route is set. ([OCPBUGS-12990](#))
- Previously, the print preview displayed incomplete topology information from the list view. With this update, a full list of resources is printed when they are longer than one page. ([OCPBUGS-11219](#))
- Previously, dynamic plugins that proxy to services with longer response times timed out at 30 seconds with a **504** error message. With this update, a 5-minute HAProxy timeout annotation was added to the console route to match the maximum timeout of most browsers. ([OCPBUGS-9917](#))

- Previously, the provided API page used the **displayName** of the provided API, but this value was not always set. As a result, the list was empty but you could still click all instances to get to the YAML of a new instance. With this update, if the **displayName** is not set, the list displays text. ([OCPBUGS-8682](#))
- Previously, the **CronJobs** table and details view did not have a **suspend** indication. With this update, **spec.suspend** was added to the list and details view for **CronJobs**. ([OCPBUGS-8299](#))
- Previously, when enabling a single plugin in the configuration of the console operator, the redeployed console fails. With this update, the list of plugins is now unique and pods run as expected. ([OCPBUGS-5059](#))
- Previously, after upgrading a plugin image, old plugin files were still requested. With this update, the **?cacheBuster=\${getRandomChars()}** query string was added when **plugin-entry.js** resources are requested. ([OCPBUGS-3495](#))

1.6.11. Monitoring

- Before this update, large amounts of CPU resources might be consumed during metrics scraping as a result of the way the **node-exporter** collected network interface information. This release fixes this issue by improving the performance of **node-exporter** when collecting network interface information, thereby resolving the issue with excessive CPU usage during metrics scraping. ([OCPBUGS-12714](#))
- Before this update, Thanos Querier failed to de-duplicate metrics by node roles. This update fixes the issue so that Thanos Querier now properly de-duplicates metrics by node roles. ([OCPBUGS-12525](#))
- Before this update, the **btrfs** collector of **node-exporter** was always enabled, which caused increased CPU usage because Red Hat Enterprise Linux (RHEL) does not support the **btrfs** storage format. With this update, the **btrfs** collector is now disabled, thereby resolving the issue. ([OCPBUGS-11434](#))
- Before this update, 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. ([OCPBUGS-10387](#))
- Before this update, the lack of a startup probe prevented the Prometheus Adapter pods from starting when the Kubernetes API had many custom resource definitions installed because the program initialization would take longer than what was allowed by the liveness probe. With this update, the Prometheus Adapter pods are now configured with a startup probe that waits five minutes before failing, thereby resolving the issue. ([OCPBUGS-7694](#))
- The **node_exporter** collector is meant to collect network interface metrics for physical interfaces only, but before this update, the **node-exporter** collector did not exclude Calico Virtual network interface controllers (NICs) when collecting these metrics. This update adds the **cali[a-f0-9]*** value to the **collector.netclass.ignored-devices** list to ensure that metrics are not collected for Calico Virtual NICs. ([OCPBUGS-7282](#))
- With this release, as a security measure, Cross Origin Resource Sharing (CORS) headers are now disabled by default for Thanos Querier. If you still need to use CORS headers, you can enable them by setting the value of the **enableCORS** parameter to **true** for the **ThanosQuerierConfig** resource. ([OCPBUGS-11889](#))

1.6.12. Networking

- Previously, when a client mutual TLS (mTLS) was configured on an ingress controller, and the certificate authority (CA) certificates in the CA bundle required more than 1 MB of certificate revocation lists (CRL) to be downloaded, the CRL config map could not be updated due to size limitations. Because of the missing CRLs, connections with valid client certificates might have been rejected with the following error: **unknown ca**.
With this update, CRLs are no longer placed in a config map, and the router now directly downloads CRLs. As a result, the CRL config map for each ingress controller no longer exists. CRLs are now downloaded directly and connections with valid client certificates are no longer rejected. ([OCPBUGS-6661](#))
- Previously, a non-compliant upstream DNS server that provided a UDP response larger than OpenShift Container Platform's specified buffer size of 512 bytes caused CoreDNS to throw an overflow error. Consequently, it would not provide a response to a DNS query.
With this update, users can now configure the **protocolStrategy** field on the **dnses.operator.openshift.io** custom resource (CR) to be **TCP**. With this field set to **TCP**, CoreDNS uses the TCP protocol for upstream requests and works around UDP overflow issues with non-compliant upstream DNS servers. ([OCPBUGS-6829](#))
- Previously, if cluster administrators configured an infra node using a taint with the **NoExecute** effect, the Ingress Operator's canary pods would not be scheduled on these infra nodes. After some time, the DaemonSet configuration would get overridden, and the pods would be terminated on the infra nodes.
With this release, the Ingress Operator now configures the canary DaemonSet to tolerate a **node-role.kubernetes.io/infra** node taint that specifies the **NoExecute** effect. As a result, canary pods are scheduled on infra nodes regardless of what effect has been specified. ([OCPBUGS-9274](#))
- Previously, when a client mutual TLS (mTLS) was configured on an ingress controller, if any of the client certificate authority (CA) certificates included a certificate revocation list (CRL) distribution point for a CRL issued by a different CA and that CRL expired, the mismatch between the distributing CA and the issuing CA caused the incorrect CRL to be downloaded. Consequently, the CRL bundle would be updated to contain an extra copy of the erroneously downloaded CRL, and the CRL that needed to be updated would be missing. Because of the missing CRL, connections with valid client certificates might have been rejected with the following error: **unknown ca**.
With this update, downloaded CRLs are now tracked by the CA that distributes them. When a CRL expires, the distributing CA's CRL distribution point is used to download an updated CRL. As a result, valid client certificates are no longer rejected. ([OCPBUGS-9464](#))
- Previously, when the Gateway API was enabled for Red Hat OpenShift Service Mesh, the Ingress Operator would fail to configure and would return the following error: **the spec.techPreview.controlPlaneMode field is not supported in version 2.4+; use spec.mode**. With this release, the Service Mesh **spec.techPreview.controlPlaneMode** API field in the **ServiceMeshControlPlane** custom resource (CR) has been replaced with **spec.mode**. As a result, the Ingress Operator is able to create a **ServiceMeshControlPlane** custom resource, and the Gateway API works properly. ([OCPBUGS-10714](#))
- Previously, when configuring DNS for Gateway API gateways, the Ingress Operator would attempt to create a DNS record for a gateway listener, even if the listener specified a hostname with a domain that was outside of the cluster's base domain. Consequently, the Ingress Operator attempted, and failed, to publish DNS records, and would return the following error: **failed to publish DNS record to zone**.
With this update, when creating a **DNSRecord** custom resource (CR) for a gateway listener, the

Ingress Operator now sets the **DNSRecord's** DNS management policy to **Unmanaged** if its domain is outside of the cluster's base domain. As a result, the Ingress Operator no longer attempts to publish records, and no longer logs the **failed to publish DNS record to zone** error. ([OCBUGS-10875](#))

- Previously, the **oc explain route.spec.tls.insecureEdgeTerminationPolicy** command documented the incorrect possible options that could be confusing to some users. With this release, the API documentation has been updated so that it shows the correct possible options for the **insecureEdgeTerminationPolicy** field. This is an API documentation fix only. ([OCBUGS-11393](#))
- Previously, a Cluster Network Operator controller monitored a broader set of resources than necessary, which resulted in its reconciler being triggered too often. Consequently, this increased the loads on both the Cluster Network Operator and the **kube-apiserver**. With this update, the Cluster Network Operator **allowlist** controller monitors its **cni-sysctl-allowlist** config map for changes. As a result, rather than being triggered when any config map is changed, the **allowlist** controller reconciler is only triggered when changes are made to the **cni-sysctl-allowlist** config map or the **default-cni-sysctl-allowlist** config map. As a result, Cluster Network Operator API requests and config map requests are reduced. ([OCBUGS-11565](#))
- **segfault** failures that were related to HaProxy have been resolved. Users should no longer receive these errors. ([OCBUGS-11595](#))
- Previously, CoreDNS terminated unexpectedly if a user created an **EndpointSlice** port without a port number. With this update, validation was added to CoreDNS to prevent it from unexpectedly terminated. ([OCBUGS-19805](#))
- Previously, the OpenShift router directed traffic to a route with a weight of **0** when it had only one back-end service. With this update, the router no longer sends traffic to routes with a single backend with weight **0**. ([OCBUGS-16623](#))
- Previously, the Ingress Operator created its canary route without specifying the **spec.subdomain** or the **spec.host** parameter on the route. Usually, this caused the API server to use the cluster's Ingress domain, which matches the domain of the default Ingress Controller, to set a default value for the **spec.host** parameter. However, if you configured the cluster by using the **appsDomain** option to set an alternative Ingress domain, the route host would have the alternative domain. Further, if you deleted the canary route, the route would be recreated with a domain that did not match the default Ingress Controller's domain, which would cause canary checks to fail. Now, the Ingress Controller specifies the **spec.subdomain** parameter when it creates the canary route. If you use the **appsDomain** option to configure your cluster and then delete the canary route, the canary checks do not fail. ([OCBUGS-16089](#))
- Previously, the Ingress Operator did not check status of DNS records in public hosted zones when updating the Operator status. This caused the Ingress Operator to report the DNS status as **Ready** when there could be errors in DNS records in public hosted zones. Now, the Ingress Operator checks the status of both public and private hosted zones, which fixes the issue. ([OCBUGS-15978](#))
- Previously, the CoreDNS **bufsize** setting was configured as 512 bytes. Now, the maximum size of the buffer for OpenShift Container Platform CoreDNS is 1232 bytes. This modification enhances DNS performance by reducing the occurrence of DNS truncations and retries. ([OCBUGS-15605](#))
- Previously, the Ingress Operator would specify the **spec.template.spec.hostNetwork: true**

parameter 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 port's **hostPort** field, 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-14995](#))

- Previously, the DNS Operator logged the **cluster-dns-operator startup has an error message: [controller-runtime] log.SetLogger(...) was never called, logs will not be displayed**: error message on startup, which could mislead users. Now, the error message is not displayed on startup. ([OCPBUGS-14395](#))
- 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-13190](#))
- Previously, transmission control protocol (TCP) connections were load balanced for all DNS. With this update, TCP connections are enabled to prefer local DNS endpoints. ([OCPBUGS-9985](#))
- Previously, for Intel E810 NICs, resetting a MAC address on an SR-IOV with a virtual function (VF) when a pod was deleted caused a failure. This resulted in a long delay when creating a pod with SR-IOV VF. With this update, the container network interface (CNI) does not fail fixing this issue. ([OCPBUGS-5892](#))
- 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](#))

1.6.13. OpenShift CLI (oc)

- Previously, container image references that have both tag and digest were not correctly interpreted by the oc-mirror plug-in and resulted in the following error:

```
"localhost:6000/cp/cpd/postgresql:13.7@sha256" is not a valid image reference: invalid reference format
```

This behavior has been fixed, and the references are now accepted and correctly mirrored. ([OCPBUGS-11840](#))

- Previously, you were receiving **401 - Unauthorized** error for registries where the number of path components exceeded the expected maximum path components. This issue is fixed by ensuring that the oc-mirror fails when the number of path components exceeds maximum path components. You can now set the maximum path components by using the flag **--max-nested-paths**, which accepts an integer value. By default, there is no limit to the maximum path components and is set to **0**. The generated **ImageContentSourcePolicy** will contain source and mirror references up to the repository level. ([OCPBUGS-8111](#), [OCPBUGS-11910](#), [OCPBUGS-11922](#))
- Previously, the oc-mirror flags **--short**, **-v**, and **--verbose** provided incorrect version information. You can now use the oc mirror **version** flag to know the correct version of oc-mirror. The oc-mirror flags **--short**, **-v**, and **--verbose** have been deprecated and will no longer

be supported. ([OCPBUGS-7845](#))

- Previously, mirroring from registry to disk would fail when several digests of an image were specified in the **imageSetConfig** without tags. The oc-mirror would add the default tag **latest** to the images. The issue is now fixed by using a truncated digest as the tag. ([OCPBUGS-2633](#))
- Previously, oc-mirror would incorrectly add the Operator catalog to **ImageContentSourcePolicy** specification. This is an unexpected behavior because the Operator catalog is directly used from the destination registry through **CatalogSource** resource. This bug is fixed by ensuring that the oc-mirror does not add the Operator catalog as an entry to **ImageContentSourcePolicy**. ([OCPBUGS-10051](#))
- Previously, mirroring images for Operators would fail when the registry domain name was not a part of the image reference. With this fix, the images are downloaded from **docker.io** if the registry domain name is not specified. ([OCPBUGS-10348](#))
- Previously, when both tag and digest were included in container image references, oc-mirror would incorrectly interpret it resulting in an **invalid reference format** error. This issue has been fixed and the images are successfully mirrored. ([OCPBUGS-11840](#))
- Previously, you could not create a **CatalogSource** resource if the name started with a number. With this fix, by default, the **CatalogSource** resource name is generated with the **cs-** prefix and is compliant with RFC 1035. ([OCPBUGS-13332](#))
- Previously, when using the **registries.conf** file, some images were not included in the mapping. With this bug fix, you can now see the images included in the mapping without any errors. ([OCPBUGS-13962](#))
- Previously, while using the insecure mirrors in the **registries.conf** file that is referenced in **--oci-registries-config** flag, oc-mirror tried to establish an HTTPS connection with the mirror registry. With this fix, you can configure oc-mirror to not use an HTTPS connection by specifying either **--source-skip-tls** or **--source-use-http** in the command line. ([OCPBUGS-14402](#))
- Previously, image mirroring would fail when you attempted to mirror OCI indexes by using oc-mirror plugins. With this fix, you can mirror OCI indexes by using oc-mirror plugins. ([OCPBUGS-15329](#))
- Previously, when mirroring several large catalogs on a low-bandwidth network, mirroring would be interrupted due to an expired authentication token resulting in an **HTTP 401 unauthorized** error. This issue is now fixed by refreshing the authentication tokens before starting the mirroring process of each catalog. ([OCPBUGS-20137](#))

1.6.14. Operator Lifecycle Manager (OLM)

- Before this update, Operator Lifecycle Manager (OLM) could cause failed installations due to initialization errors when the API server was busy. This update fixes the issue by adding a one-minute-retry interval for initialization errors. ([OCPBUGS-13128](#))
- Before this update, a race condition occurred if custom catalogs used the same names as the default Red Hat catalogs in a disconnected environment. If the default Red Hat catalogs were disabled, the catalogs were created at start and deleted after the OperatorHub custom resource (CR) was reconciled. As a result, the custom catalogs were deleted along with the default Red Hat catalogs. With this update, the OperatorHub CR is reconciled before any catalogs are deleted, preventing the race condition. ([OCPBUGS-9357](#))

- Before this update, the channels of some Operators were displayed on OperatorHub in a random order. With this update, Operator channels are displayed in lexicographical order. ([OCPBUGS-7910](#))
- Before this update, registry pods were not drained gracefully by the autoscaler if the controller flag was not set to true in the owner references file. With this update, the controller flag is set to true and draining nodes no longer requires a forceful shutdown. ([OCPBUGS-7431](#))
- Before this update, **collect-profiles** pods caused regular spikes of CPU usage due to the way certificates were generated. With this update, certificates are generated daily, the loading of the certificate is optimized, and CPU usage is lower. ([OCPBUGS-1684](#))

1.6.15. OpenShift API server

- Previously, the **metadata.namespace** field would be automatically populated in update and patch requests to the **projects** resource. As a result, the affected requests would generate spurious validation errors. With this release, the **projects** resource is no longer automatically populated. ([OCPBUGS-8232](#))

1.6.16. Red Hat Enterprise Linux CoreOS (RHCOS)

- Previously, pods in OpenShift Container Platform that access block persistent volume claims (PVC) storage with logical volume manager (LVM) metadata could get stuck when terminating. This is because the same LVM devices were active both inside the container and on the host. An example of this occurred when running a virtual machine inside a pod that used OpenShift Virtualization that in turn used LVM for the virtual machine. With this update, RHCOS by default only attempts to setup and access devices that are in the **/etc/lvm/devices/system.devices** file. This prevents contentious access to the LVM devices inside the virtual machine guests. ([OCPBUGS-5223](#))
- Previously, pods were stuck in the **ContainerCreating** state on Google Cloud Platform (GCP) Confidential Computing instances, which caused a volume mount failure. This fix adds support for the Persistent Disk storage type for Confidential Computing instances in Google Cloud Platform, which can be used as persistent volumes in OpenShift Container Platform. As a result, pods are able to enter a **Running** state and volumes can be mounted. ([OCPBUGS-7582](#))

1.6.17. Storage

- Previously, when the cluster-wide proxy is enabled on IBM Cloud® clusters, there was a failure to provision volumes. ([OCPBUGS-18142](#))
- The **vsphereStorageDriver** field of the Storage Operator object has been deprecated. This field was used to opt in to CSI migration on OpenShift Container Platform 4.13 vSphere clusters, but it has no effect on OpenShift Container Platform 4.14 and newer clusters. ([OCPBUGS-13914](#))

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*

1.7.1. Networking Technology Preview features

Table 1.17. Networking Technology Preview tracker

Feature	4.12	4.13	4.14
PTP dual NIC hardware configured as boundary clock	Technology Preview	General Availability	General Availability
Egress IPs on additional network interfaces	Not Available	Not Available	General Availability
Intel E810 Westport Channel NIC as PTP grandmaster clock	Not Available	Technology Preview	Technology Preview
Dual Intel E810 Westport Channel NICs as PTP grandmaster clock	Not Available	Not Available	Technology Preview
Ingress Node Firewall Operator	Technology Preview	Technology Preview	General Availability
Advertise using L2 mode the MetalLB service from a subset of nodes, using a specific pool of IP addresses	Technology Preview	Technology Preview	Technology Preview
Multi-network policies for SR-IOV networks	Technology Preview	Technology Preview	Technology Preview
OVN-Kubernetes network plugin as secondary network	Not Available	Technology Preview	General Availability
Updating the interface-specific safe sysctls list	Technology Preview	Technology Preview	Technology Preview
MT2892 Family [ConnectX-6 Dx] SR-IOV support	Technology Preview	General Availability	General Availability
MT2894 Family [ConnectX-6 Lx] SR-IOV support	Technology Preview	General Availability	General Availability
MT42822 BlueField-2 in ConnectX-6 NIC mode SR-IOV support	Technology Preview	General Availability	General Availability

Feature	4.12	4.13	4.14
Silicom STS Family SR-IOV support	Technology Preview	General Availability	General Availability
MT2892 Family [ConnectX-6 Dx] OvS Hardware Offload support	Technology Preview	General Availability	General Availability
MT2894 Family [ConnectX-6 Lx] OvS Hardware Offload support	Technology Preview	General Availability	General Availability
MT42822 BlueField-2 in ConnectX-6 NIC mode OvS Hardware Offload support	Technology Preview	General Availability	General Availability
Switching Bluefield-2 from DPU to NIC	Technology Preview	General Availability	General Availability
Intel E810-XXVDA4T	Not Available	General Availability	General Availability
Egress service custom resource	Not Available	Not Available	Technology Preview
VRF specification in BGPPeer custom resource	Not Available	Not Available	Technology Preview
VRF specification in NodeNetworkConfigurationPolicy custom resource	Not Available	Not Available	Technology Preview
Admin Network Policy (AdminNetworkPolicy)	Not Available	Not Available	Technology Preview
IPsec external traffic (north-south)	Not Available	Not Available	Technology Preview

1.7.2. Storage Technology Preview features

Table 1.18. Storage Technology Preview tracker

Feature	4.12	4.13	4.14
Automatic device discovery and provisioning with Local Storage Operator	Technology Preview	Technology Preview	Technology Preview
Google Filestore CSI Driver Operator	Technology Preview	Technology Preview	General Availability

Feature	4.12	4.13	4.14
CSI automatic migration (Azure file, VMware vSphere)	Technology Preview	General Availability	General Availability
CSI inline ephemeral volumes	Technology Preview	General Availability	General Availability
IBM Power® Virtual Server Block CSI Driver Operator	Not Available	Technology Preview	Technology Preview
NFS support for Azure File CSI Operator Driver	Generally Available	Generally Available	Generally Available
Read Write Once Pod access mode	Not available	Not available	Technology Preview
Build CSI Volumes in OpenShift Builds	Technology Preview	Technology Preview	General Availability
Shared Resources CSI Driver in OpenShift Builds	Technology Preview	Technology Preview	Technology Preview
Secrets Store CSI Driver Operator	Not available	Not available	Technology Preview

1.7.3. Installation Technology Preview features

Table 1.19. Installation Technology Preview tracker

Feature	4.12	4.13	4.14
Adding kernel modules to nodes with kvc	Technology Preview	Technology Preview	Technology Preview
Azure Tagging	Not Available	Technology Preview	General Availability
Enabling NIC partitioning for SR-IOV devices	Not Available	Technology Preview	Technology Preview
GCP Confidential VMs	Not Available	Technology Preview	General Availability
User-defined labels and tags for Google Cloud Platform (GCP)	Not Available	Not Available	Technology Preview
Installing a cluster on Alibaba Cloud by using installer-provisioned infrastructure	Technology Preview	Technology Preview	Technology Preview

Feature	4.12	4.13	4.14
Mount shared entitlements in BuildConfigs in RHEL	Technology Preview	Technology Preview	Technology Preview
Multi-architecture compute machines	Technology Preview	General Availability	General Availability
AWS Outposts platform	Technology Preview	Technology Preview	Technology Preview
Installing OpenShift Container Platform on Oracle® Cloud Infrastructure (OCI) with VMs	N/A	N/A	General Availability
Installing OpenShift Container Platform on Oracle® Cloud Infrastructure (OCI) on bare metal	N/A	Developer Preview	Developer Preview
Selectable Cluster Inventory	Technology Preview	Technology Preview	Technology Preview
Static IP addresses with vSphere (IPI only)	Not Available	Not Available	Technology Preview

1.7.4. Node Technology Preview features

Table 1.20. Nodes Technology Preview tracker

Feature	4.12	4.13	4.14
Linux Control Group version 2 (cgroup v2)	Technology Preview	General Availability	General Availability
crun container runtime	Technology Preview	General Availability	General Availability
Cron job time zones	Technology Preview	Technology Preview	General Availability
MaxUnavailableStatefulSet featureset	Not Available	Not Available	Technology Preview

1.7.5. Multi-Architecture Technology Preview features

Table 1.21. Multi-Architecture Technology Preview tracker

Feature	4.12	4.13	4.14
IBM Secure Execution on IBM Z® and IBM® LinuxONE	Technology Preview	General Availability	General Availability
IBM Power® Virtual Server using installer-provisioned infrastructure	Not Available	Technology Preview	Technology Preview
kdump on arm64 architecture	Technology Preview	Technology Preview	Technology Preview
kdump on s390x architecture	Technology Preview	Technology Preview	Technology Preview
kdump on ppc64le architecture	Technology Preview	Technology Preview	Technology Preview

1.7.6. Specialized hardware and driver enablement Technology Preview features

Table 1.22. Specialized hardware and driver enablement Technology Preview tracker

Feature	4.12	4.13	4.14
Driver Toolkit	General Availability	General Availability	General Availability
Hub and spoke cluster support	Technology Preview	General Availability	General Availability

1.7.7. Scalability and performance Technology Preview features

Table 1.23. Scalability and performance Technology Preview tracker

Feature	4.12	4.13	4.14
Tuning etcd latency tolerances	Not Available	Not Available	Technology Preview
Hyperthreading-aware CPU manager policy	Technology Preview	Technology Preview	Technology Preview
Node Observability Operator	Technology Preview	Technology Preview	Technology Preview
factory-precaching-cli tool	Not Available	Technology Preview	Technology Preview

Feature	4.12	4.13	4.14
Single-node OpenShift cluster expansion with worker nodes	Technology Preview	General Availability	General Availability
Topology Aware Lifecycle Manager (TALM)	Technology Preview	General Availability	General Availability
Mount namespace encapsulation	Not Available	Technology Preview	Technology Preview
NUMA-aware scheduling with NUMA Resources Operator	Technology Preview	General Availability	General Availability
HTTP transport replaces AMQP for PTP and bare-metal events	Not Available	Technology Preview	Technology Preview
Workload partitioning for three-node clusters and standard clusters	Not Available	Technology Preview	General Availability

1.7.8. Operator lifecycle and development Technology Preview features

Table 1.24. Operator lifecycle and development Technology Preview tracker

Feature	4.12	4.13	4.14
Operator Lifecycle Manager (OLM) v1	Not Available	Not Available	Technology Preview
RukPak	Technology Preview	Technology Preview	Technology Preview
Platform Operators	Technology Preview	Technology Preview	Technology Preview
Hybrid Helm Operator	Technology Preview	Technology Preview	Technology Preview
Java-based Operator	Technology Preview	Technology Preview	Technology Preview

1.7.9. Monitoring Technology Preview features

Table 1.25. Monitoring Technology Preview tracker

Feature	4.12	4.13	4.14
---------	------	------	------

Feature	4.12	4.13	4.14
Alerting rules based on platform monitoring metrics	Technology Preview	Technology Preview	General Availability
Metrics Collection Profiles	Not Available	Technology Preview	Technology Preview

1.7.10. Hosted control plane Technology Preview features

Table 1.26. Hosted control plane Technology Preview tracker

Feature	4.12	4.13	4.14
Hosted control planes for OpenShift Container Platform on Amazon Web Services (AWS)	Technology Preview	Technology Preview	Technology Preview
Hosted control planes for OpenShift Container Platform on bare metal	Technology Preview	Technology Preview	General Availability
Hosted control planes for OpenShift Container Platform on OpenShift Virtualization	Not Available	Technology Preview	General Availability
Hosted control planes for an ARM64 OpenShift Container Platform cluster on AWS	Not available	Technology Preview	Technology Preview
Hosted control planes for OpenShift Container Platform on IBM Power	Not Available	Not Available	Technology Preview
Hosted control planes for OpenShift Container Platform on IBM Z	Not Available	Not Available	Technology Preview

1.7.11. Machine management Technology Preview features

Table 1.27. Machine management Technology Preview tracker

Feature	4.12	4.13	4.14
Managing machines with the Cluster API for Amazon Web Services	Technology Preview	Technology Preview	Technology Preview
Managing machines with the Cluster API for Google Cloud	Technology Preview	Technology Preview	Technology Preview
Managing machines with the Cluster API for IBM Power® Virtual Server	Technology Preview	Technology Preview	Technology Preview

Feature	4.12	4.13	4.14
Cloud controller manager for Alibaba Cloud	Technology Preview	Technology Preview	Technology Preview
Cloud controller manager for Amazon Web Services	Technology Preview	Technology Preview	General Availability
Cloud controller manager for Google Cloud Platform	Technology Preview	Technology Preview	Technology Preview
Cloud controller manager for IBM Cloud Power VS	Not Available	Technology Preview	Technology Preview
Cloud controller manager for Microsoft Azure	Technology Preview	Technology Preview	General Availability
Cloud controller manager for Nutanix	Technology Preview	General Availability	General Availability
Cloud controller manager for VMware vSphere	Technology Preview	General Availability	General Availability

1.7.12. Authentication and authorization Technology Preview features

Table 1.28. Authentication and authorization Technology Preview tracker

Feature	4.12	4.13	4.14
Pod security admission restricted enforcement	Technology Preview	Technology Preview	Technology Preview

1.7.13. Machine Config Operator Technology Preview features

Table 1.29. Machine Config Operator Technology Preview tracker

Feature	4.12	4.13	4.14
Red Hat Enterprise Linux CoreOS (RHCOS) image layering	Technology Preview	General Availability	General Availability

1.8. KNOWN ISSUES

- A regression in the behaviour of **libreswan** caused some nodes with IPsec enabled to lose communication with pods on other nodes in the same cluster. To resolve this issue, consider disabling IPsec for your cluster. ([OCPBUGS-42952](#))
- 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 4.14, you can either revoke or continue to allow unauthenticated access. Unless there is a specific need for unauthenticated access, you should revoke it. If you do continue to allow unauthenticated access, be aware of the increased risks.



WARNING

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

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

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

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

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

([BZ#1821771](#))

- The **oc annotate** command does not work for LDAP group names that contain an equal sign (=), because the command uses the equal sign as a delimiter between the annotation name and value. As a workaround, use **oc patch** or **oc edit** to add the annotation. ([BZ#1917280](#))
- If the installation program cannot get all of the projects that are associated with the Google Cloud Platform (GCP) service account, the installation fails with a **context deadline exceeded** error message.

This behavior occurs when the following conditions are met:

- The service account has access to an excessive number of projects.
- The installation program is run with one of the following commands:

- **openshift-install create install-config**

Error message

```
FATAL failed to fetch Install Config: failed to fetch dependency of "Install Config":
failed to fetch dependency of "Base Domain": failed to generate asset "Platform":
failed to get projects: context deadline exceeded
```

- **openshift-install create cluster** without an existing installation configuration file (**install-config.yaml**)

Error message

```
FATAL failed to fetch Metadata: failed to fetch dependency of "Metadata": failed to
fetch dependency of "Cluster ID": failed to fetch dependency of "Install Config": failed
to fetch dependency of "Base Domain": failed to generate asset "Platform": failed to
get projects: context deadline exceeded
```

- **openshift-install create manifests** with or without an existing installation configuration file

Error message

```
ERROR failed to fetch Master Machines: failed to load asset "Install Config": failed to
create install config: platform.gcp.project: Internal error: context deadline exceeded
```

As a workaround, if you have an installation configuration file, update it with a specific project id to use (**platform.gcp.projectID**). Otherwise, manually create an installation configuration file, and enter a specific project id. Run the installation program again, specifying the file. ([OCPBUGS-15238](#))

- Booting fails in a large compute node. ([OCPBUGS-20075](#))
- When you deploy a cluster with a Network Type of **OVNKubernetes** on IBM Power®, compute nodes might reboot because of a kernel stack overflow. As a workaround, you can deploy the cluster with a Network Type of **OpenShiftSDN**. ([RHEL-3901](#))
- The following known issue applies to users who updated their OpenShift Container Platform deployment to an early access version of 4.14 using release candidate 3 or 4: After the introduction of the node identify feature, some pods that were running as root are updated to run unprivileged. For users who updated to an early access version of OpenShift Container Platform 4.14, attempting to upgrade to the official version of 4.14 might not progress. In this scenario, the Network Operator reports the following state, indicating an issue with the update: **DaemonSet "/openshift-network-node-identity/network-node-identity" update is rolling.**

As a workaround, you can delete all pods in the **openshift-network-node-identity** namespace by running the following command: **oc delete --force=true -n openshift-network-node-identity --all pods**. After running this command, the update continues.

For more information about early access, [candidate-4.14 channel](#).

- Currently, users cannot modify the **interface-specific** safe `sysctl` list by updating the **cni-sysctl-allowlist** config map in the **openshift-multus** namespace. As a workaround, you can modify, either manually or with a DaemonSet, the `/etc/cni/tuning/allowlist.conf` file on the node or nodes. ([OCPBUGS-11046](#))
- In OpenShift Container Platform 4.14, all nodes use Linux control group version 2 (cgroup v2) for internal resource management in alignment with the default RHEL 9 configuration. However, if you apply a performance profile in your cluster, the low-latency tuning features associated with the performance profile do not support cgroup v2.
As a result, if you apply a performance profile, all nodes in the cluster reboot to switch back to the cgroup v1 configuration. This reboot includes control plane nodes and worker nodes that were not targeted by the performance profile.

To revert all nodes in the cluster to the cgroups v2 configuration, you must edit the **Node** resource. For more information, see [Configuring Linux cgroup v2](#). You cannot revert the cluster to the cgroups v2 configuration by removing the last performance profile. ([OCPBUGS-16976](#))

- AWS **M4** and **C4** instances might fail to boot properly in clusters installed using OpenShift Container Platform 4.14. There is no current workaround. ([OCPBUGS-17154](#))
- There is a known issue in this release which prevents installing a cluster on Alibaba Cloud by using installer-provisioned infrastructure. Installing a cluster on Alibaba Cloud is a Technology Preview feature in this release. ([OCPBUGS-20552](#))
- For clusters that have root volume availability zones and are running on RHOSP that you upgrade to 4.14, you must converge control plane machines onto one server group before you can enable control plane machine sets. To make the required change, follow the instructions in the [knowledge base article](#). ([OCPBUGS-13300](#))
- For clusters that have compute zones configured with at least one zone and are running on RHOSP, which is upgradable to version 4.14, root volumes must now also be configured with at least one zone. If this configuration change does not occur, a control plane machine set cannot be generated for your cluster. To make the required change, follow the instructions in the [knowledge base article](#). ([OCPBUGS-15997](#))
- Currently, an error might occur when deleting a pod that uses an SR-IOV network device. This error is caused by a change in RHEL 9 where the previous name of a network interface is added to its alternative names list when it is renamed. As a consequence, when a pod attached to an SR-IOV virtual function (VF) is deleted, the VF returns to the pool with a new unexpected name, such as **dev69**, instead of its original name, such as **ensf0v2**. Although this error is not severe, the Multus and SR-IOV logs might show the error while the system recovers on its own. Deleting the pod might take a few seconds longer due to this error. ([OCPBUGS-11281](#), [OCPBUGS-18822](#), [RHEL-5988](#))
- Starting from **RHEL 5.14.0-284.28.1.el9_2**, if you configure a SR-IOV virtual function with a specific MAC address, configuration errors might occur in the i40e driver. Consequently, Intel 7xx Series NICs might have connectivity issues. As a workaround, avoid specifying MAC addresses in the **metadata.annotations** field in the Pod resource. Instead, use the default address that the driver assigns to the virtual function. ([RHEL-7168](#), [OCPBUGS-19536](#), [OCPBUGS-19407](#), [OCPBUGS-18873](#))
- Currently, defining a **sysctl** value for a setting with a slash in its name, such as for bond devices, in the **profile** field of a **Tuned** resource might not work. Values with a slash in the **sysctl** option name are not mapped correctly to the **/proc** filesystem. As a workaround, create a

MachineConfig resource that places a configuration file with the required values in the `/etc/sysctl.d` node directory. ([RHEL-3707](#))

- Currently, due to an issue with Kubernetes, the CPU Manager is unable to return CPU resources from the last pod admitted to a node to the pool of available CPU resources. These resources can be allocated if a subsequent pod is admitted to the node. However, this in turn becomes the last pod, and again, the CPU manager cannot return the resources of this pod to the available pool.
This issue affects the CPU load balancing features because these features depend on the CPU Manager releasing CPUs to the available pool. Consequently, non-guaranteed pods might run with a reduced number of CPUs. As a workaround, schedule a pod with a **best-effort** CPU Manager policy on the affected node. This pod will be the last pod admitted, ensuring that resources are properly released to the available pool. ([OCPBUGS-17792](#))
- Currently, the Machine Config Operator (MCO) might apply an incorrect cgroup version argument for custom pools because of how the MCO handles machine configurations for worker pools and custom pools. As a consequence, nodes in the custom pool might have an incorrect cgroup kernel argument, resulting in unpredictable behavior. As a workaround, specify the cgroup version kernel arguments for worker and control plane pools only. ([OCPBUGS-19352](#))
- Currently, due to a race condition between the application of a **udev** rule on physical network devices and the application of the default requests per second (RPS) mask to all network devices, some physical network devices might feature the wrong RPS mask configuration. As a consequence, a performance degradation might affect the physical network devices with the wrong RPS mask configuration. It is anticipated that an upcoming z-stream release will include a fix for this issue. ([OCPBUGS-21845](#))
- Broadcom network interface controllers in legacy Single Root I/O Virtualization (SR-IOV) do not support quality of service (QoS) and tag protocol identifier (TPID) settings for the SRIOV VLAN. This affects Broadcom BCM57414, Broadcom BCM57508, and Broadcom BCM57504. ([RHEL-9881](#))
- When you create a hosted cluster in an environment that uses the dual-stack network, you might encounter the following DNS-related issues:
 - **CrashLoopBackOff** state in the **service-ca-operator** pod: When the pod tries to reach the Kubernetes API server through the hosted control plane, the pod cannot reach the server because the data plane proxy in the **kube-system** namespace cannot resolve the request. This issue occurs because in the HAProxy setup, the front end uses an IP address and the back end uses a DNS name that the pod cannot resolve.
 - Pods stuck in **ContainerCreating** state: This issue occurs because the **openshift-service-ca-operator** cannot generate the **metrics-tls** secret that the DNS pods need for DNS resolution. As a result, the pods cannot resolve the Kubernetes API server.

To resolve those issues, configure the DNS server settings by following the guidelines in [Configuring DNS for a dual stack network](#) . ([OCPBUGS-22753](#), [OCPBUGS-23234](#))

- In hosted control planes for OpenShift Container Platform, the following Operators and components are not tested ([OCPSTRAT-605](#)):
 - Performance Addon Operator
 - OpenShift sandboxed containers
 - Red Hat OpenShift GitOps

- Red Hat OpenShift Service Mesh
- Red Hat OpenShift Pipelines
- Red Hat OpenShift Dev Spaces
- Red Hat's single sign-on technology
- The web terminal in the OpenShift Container Platform web console
- Migration toolkit for applications
- In hosted control planes for OpenShift Container Platform, installing the File Integrity Operator on a hosted cluster fails. ([OCPBUGS-3410](#))
- In hosted control planes for OpenShift Container Platform, the Vertical Pod Autoscaler Operator fails to install on a hosted cluster. ([PODAUTO-65](#))
- In hosted control planes for OpenShift Container Platform, on the bare metal and OpenShift Virtualization platforms, the auto-repair function is disabled. ([OCPBUGS-20028](#))
- In hosted control planes for OpenShift Container Platform, using the Secrets Store CSI Driver Operator with AWS Secrets Manager or AWS Systems Manager Parameter Store is not supported. ([OCPBUGS-18711](#))
- In hosted control planes for OpenShift Container Platform, the **default**, **kube-system**, and **kube-public** namespaces are not properly excluded from pod security admission. ([OCPBUGS-22379](#))
- In hosted control planes on OpenShift Virtualization, worker nodes might lose network connectivity after a restart. ([OCPBUGS-23208](#))
- In hosted control planes for OpenShift Container Platform, the HyperShift Operator extracts the release metadata only once during Operator initialization. When you make changes in the management cluster or create a hosted cluster, the HyperShift Operator does not refresh the release metadata. As a workaround, restart the HyperShift Operator by deleting its pod deployment. ([OCPBUGS-29110](#))
- In hosted control planes for OpenShift Container Platform, when you create the custom resource definition (CRD) for **ImageDigestMirrorSet** and **ImageContentSourcePolicy** objects at the same time in a disconnected environment, the HyperShift Operator creates the object only for the **ImageDigestMirrorSet** CRD, ignoring the **ImageContentSourcePolicy** CRD. As a workaround, copy the **ImageContentSourcePolicies** object configuration in the **ImageDigestMirrorSet** CRD. ([OCPBUGS-29466](#))
- In hosted control planes for OpenShift Container Platform, when creating a hosted cluster in a disconnected environment, if you do not set the **hypershift.openshift.io/control-plane-operator-image** annotation explicitly in the **HostedCluster** resource, the hosted cluster deployment fails with an error. ([OCPBUGS-29494](#))
- Agent-based installations on vSphere will fail due to a failure to remove node taint, which causes the installation to be stuck in a pending state. Single-node OpenShift clusters are not impacted. You can work around this issue by running the following command to manually remove the node taint:

```
$ oc adm taint nodes <node_name>
node.cloudprovider.kubernetes.io/uninitialized:NoSchedule-
```

- ([OCPBUGS-20049](#))
- There is a known issue with using Azure confidential virtual machines, which is a Technology Preview feature in this release. Configuring a cluster to encrypt the managed disk and the Azure VM Guest State (VMGS) blob with a platform-managed key (PMK) or a customer-managed key (CMK) is unsupported. To avoid this issue, only enable encryption of the VMGS blob by setting the value of the **securityEncryptionType** parameter to **VMGuestStateOnly**.
([OCPBUGS-18379](#))
- There is a known issue with using Azure confidential virtual machines, which is a Technology Preview feature in this release. Installing a cluster configured to use this feature fails because the control plane provisioning process times out after 30 minutes.
If this occurs, you can run the **openshift-install create cluster** command a second time to complete the installation.

To avoid this issue, you can enable confidential VMs on an existing cluster by using machine sets.
([OCPBUGS-18488](#))
- When you run hosted control planes for OpenShift Container Platform on a bare-metal platform, if a worker node fails, another node is not automatically added to the hosted cluster, even when other agents are available. As a workaround, manually delete the machine that is associated with the failed worker node. ([MGMT-15939](#))
- Since the source catalog bundles an architecture specific **opm** binary, you must run the mirroring from that architecture. For instance if you are mirroring a ppc64le catalog, you must run **oc-mirror** from a system that runs on the ppc64le architecture. ([OCPBUGS-22264](#))
- If more than one OpenShift Container Platform group points to the same LDAP group, only one OpenShift Container Platform group is synced. The **oc adm groups sync** command prints a warning when multiple groups point to the same LDAP group, indicating that only a single group is eligible for mapping. ([OCPBUGS-11123](#))
- Installation fails when installing OpenShift Container Platform with the **bootMode** set to **UEFISecureBoot** on a node where Secure Boot is disabled. Subsequent attempts to install OpenShift Container Platform with Secure Boot enabled will proceed normally. ([OCPBUGS-19884](#))
- In OpenShift Container Platform 4.14, a **MachineConfig** object with Ignition version 3.4 might fail scans of the **api-collector** pods with **CrashLoopBackOff** errors, causing the Compliance Operator to not work as expected. ([OCPBUGS-18025](#))
- In OpenShift Container Platform 4.14, assigning an IPv6 egress IP to a network interface that is not the primary network interface is unsupported. This is a known issue and will be fixed in a future version of OpenShift Container Platform. ([OCPBUGS-17637](#))
- Run Once Duration Override Operator (RODOO) cannot be installed on clusters managed by the HyperShift Operator. ([OCPBUGS-17533](#))
- When you run CNF latency tests on an OpenShift Container Platform cluster, the **oslat** test can sometimes return results greater than 20 microseconds. This results in an **oslat** test failure. ([RHEL-9279](#))
- When you use **preempt-rt** patches with the realtime kernel and you update the SMP affinity of a network interrupt, the corresponding IRQ thread does not immediately receive the update. Instead, the update takes effect when the next interrupt is received, and the thread is

subsequently migrated to the correct core. ([RHEL-9148](#))

- 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 20 μ s, 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 20 μ s for over 99.999999% of the samples. ([RHELPLAN-138733](#))
- The global navigation satellite system (GNSS) module in an Intel Westport Channel e810 NIC that is configured as a grandmaster clock (T-GM) can report the GPS **FIX** state and the GNSS offset between the GNSS module and the GNSS constellation satellites.
The current T-GM implementation does not use the **ubxtool** CLI to probe the **ublox** module for reading the GNSS offset and GPS **FIX** values. Instead, it uses the **gpsd** service to read the GPS **FIX** information. This is because the current implementation of the **ubxtool** CLI takes 2 seconds to receive a response, and with every call, it increases CPU usage threefold. ([OCPBUGS-17422](#))
- In a PTP grandmaster clock clocked sourced from GNSS, when the GNSS signal is lost, the Digital Phase Locked Loop (DPLL) clock state can change in 2 ways: it can transition to unlocked, or it can enter a holdover state. Currently, the driver transitions the DPLL state to unlocked by default. An upstream change is currently being developed to handle the holdover state functionality and to configure which state machine handling is used. ([RHELPLAN-164754](#))
- The DPLL subsystem and DPLL support is not currently enabled in the Intel Westport Channel e810 NIC ice driver. ([RHELPLAN-165955](#))
- The current grandmaster clock (T-GM) implementation has a single NMEA sentence generator sourced from the GNSS without a backup NMEA sentence generator. If NMEA sentences are lost on their way to the e810 NIC, the T-GM cannot synchronize the devices in the network synchronization chain and the PTP Operator reports an error. A proposed fix is to report a **FREERUN** event when the NMEA string is lost. ([OCPBUGS-19838](#))
- Currently, due to differences in setting up a container's cgroup hierarchy, containers that use the **crun** OCI runtime along with a **PerformanceProfile** configuration encounter performance degradation. As a workaround, use the **runc** OCI container runtime. Although the **runc** container runtime has lower performance during container startup, shutdown operations, and **exec** probes, the **crun** and **runc** container runtimes are functionally identical. It is anticipated that an upcoming z-stream release will include a fix for this issue. ([OCPBUGS-20492](#))
- There is a known issue after enabling and disabling IPsec during runtime that causes the cluster to be in an unhealthy state with the error message: **an unknown error has occurred: MultipleErrors**. ([OCPBUGS-19408](#))
- Creating pods with Microsoft Azure File NFS volumes that are scheduled to the control plane node causes the mount to be denied.
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. ([OCPBUGS-18581](#))
- For clusters that run on RHOSP 17.1 and use network function virtualization (NFV), a known issue in RHOSP prevents successful cluster deployment. There is no workaround for this issue. Contact Red Hat Support to request a hotfix. ([BZ#2228643](#))
- There is no support for Kuryr installations on RHOSP 17.1.
- Currently, the update to HAProxy version 2.6.13 in OpenShift Container Platform 4.14 causes an increase in P99 latency for re-encrypt traffic. This is observed when the volume of ingress

traffic puts the HAProxy component of the **IngressController** custom resource (CR) under a considerable load. The latency increase does not affect overall throughput, which remains consistent.

The default **IngressController** CR is configured with 4 HAProxy threads. If you experience elevated P99 latencies during high ingress traffic conditions, specifically with re-encrypt traffic, it's recommended to increase the number of HAProxy threads to reduce latency. ([OCPBUGS-18936](#))

- For Single-node OpenShift on 4.14 and Google Cloud Platform (GCP), there is a known issue with the Cloud Network Config Controller (CNCC) entering a **CrashLoopBackOff** state. This occurs at initialization time when the CNCC tries to reach the GCP internal load balancer address and the resulting hairpin traffic is not correctly prevented in OVN-Kubernetes shared gateway mode on GCP causing it to get dropped. Cluster Network Operator will show a **Progressing=true** status in such case. Currently, there is no workaround for this issue. ([OCPBUGS-20554](#))
- On an Single-node OpenShift that has guaranteed CPUs and where Interrupt Request (IRQ) load balancing is disabled, large latency spikes can occur at container startup. ([OCPBUGS-22901](#))
- When deploying an application that has a large number of pods, some of which have CPU limits configured, the deployment can fail. The workaround is to re-deploy the application. ([RHEL-7232](#))
- On an Single-node OpenShift that has disabled capabilities, the **openshift-controller-manager-operator** may continuously restart. As a workaround, enable the build capability or manually create the **builds.config.openshift.io** CRD.

Perform the following steps to manually create the **builds.config.openshift.io** CRD:

1. Run the following command to extract the release manifests:

```
$ oc adm release extract --to manifests
```

2. Search for **builds.config.openshift.io** within the **manifests** directory and subdirectories:

```
$ grep -r builds.config.openshift.io manifests
```

Expected output

```
manifests/0000_10_openshift-controller-manager-operator_01_build.crd.yaml: name:
builds.config.openshift.io
```

3. Apply the configuration specified in the **0000_10_openshift-controller-manager-operator_01_build.crd.yaml**:

```
$ oc apply -f manifests/0000_10_openshift-controller-manager-
operator_01_build.crd.yaml
```

([OCPBUGS-21778](#))

- There is a known issue that prevents installing a cluster on or updating a cluster to this version of OpenShift Container Platform on Microsoft Azure Stack Hub. For more details and a workaround, see the information in this [Red Hat Knowledgebase article](#). ([OCPBUGS-20548](#))

- There is a known issue with Microsoft Azure clusters that use Azure AD Workload Identity in versions of OpenShift Container Platform 4.14 prior to version 4.14.2. A recent change to the default security settings for new Azure storage accounts in the **eastus** region prevents the installation of clusters that use Azure AD Workload Identity in that region. Other regions do not seem to be impacted at this time, but might be impacted in the future. This issue is resolved in OpenShift Container Platform [4.14.2](#).

To work around this issue, manually create a storage account that allows public access before running **ccoctl azure create-all** in the procedure [Configuring an Azure cluster to use short-term credentials](#).

Perform the following steps:

1. Create a resource group for the storage account by running the following Azure CLI command:

```
$ az group create --name <oidc_resource_group_name> --location <azure_region>
```

2. Create a storage account that allows public access by running the following Azure CLI command:

```
$ az storage account create --name <storage_account_name> --resource-group <oidc_resource_group_name> --location <azure_region> --sku Standard_LRS --kind StorageV2 --allow-blob-public-access true
```

3. When you use the **ccoctl** tool to process all **CredentialsRequest** objects by running the following command, you must specify the resources created in the previous steps.

```
$ ccoctl azure create-all \
  --name=<azure_infra_name> \
  --output-dir=<ccoctl_output_dir> \
  --region=<azure_region> \
  --subscription-id=<azure_subscription_id> \
  --credentials-requests-dir=<path_to_credentials_requests_directory> \
  --dnszone-resource-group-name=<azure_dns_zone_resource_group_name> \
  --tenant-id=<azure_tenant_id> \
  --storage-account-name=<storage_account_name> \
  --oidc-resource-group-name=<oidc_resource_group-name>
```

([OCPBUGS-22651](#))

- 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"] 1 2 3 4 5
    inspect.metal3.io: disabled
    bmac.agent-install.openshift.io/hostname: <fqdn> 6
    bmac.agent-install.openshift.io/role: <role> 7

  generation: 1
  name: openshift-worker-0
  namespace: mynamespace
spec:
  automatedCleaningMode: disabled
  bmc:
    address: idrac-virtualmedia://<bmc_ip>/redfish/v1/Systems/System.Embedded.1 8
    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 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.

- 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-17895](#))

- There is a known issue in this release preventing access to the Web Terminal Operator after it has been installed. This issue will be fixed in a future OpenShift Container Platform release. ([OCPBUGS-14463](#))
- There is a known issue when upgrading the Cluster Network Operator from version 4.13 to 4.14. The conversion to the new OVN Kubernetes interconnect multi-zone architecture can lead to packet drops that cause brief network outages. The impact is on East/West traffic in local gateway mode with **routingViaHost** set to **true**. ([OCPBUGS-38891](#))
- There is a known issue with HAProxy in OpenShift Container Platform 4.14 involving applications that send invalid **Transfer-Encoding** headers. This issue results in losing external access to pods when they expose a route that sends these invalid headers. For more details, see the information in this [Red Hat Knowledgebase article](#). ([OCPBUGS-43095](#))

1.9. ASYNCHRONOUS ERRATA UPDATES

Security, bug fix, and enhancement updates for OpenShift Container Platform 4.14 are released as asynchronous errata through the Red Hat Network. All OpenShift Container Platform 4.14 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 4.14. Versioned asynchronous releases, for example with the form OpenShift Container Platform 4.14.z, will be detailed in subsections. In addition, releases in which the errata text cannot fit in the space provided by the advisory will be detailed in subsections that follow.



IMPORTANT

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

1.9.1. RHSA-2026:5107 - OpenShift Container Platform 4.14.63 image release, fixed issues, and security update

Issued: 26 March 2026

OpenShift Container Platform release 4.14.63 is now available. The list of bug fixes that are included in the update is documented in the [RHSA-2026:5107](#) advisory. The RPM packages that are included in the update are provided by the [RHSA-2026:5086](#) 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 4.14.63 --pullspecs
```

1.9.2. Enhancements

- Before this update, the self-signed loopback certificate for the Kubernetes API Server expired after one year. With this release, the expiration date of the certificate is extended to three years. ([OCPBUGS-59473](#))

1.9.3. Fixed issues

- Before this update, when you configured a Windows VM on which the default PowerShell execution policy that was set to **Restricted**, the Windows Machine Config Operator (WMCO) could not execute the necessary commands. As a consequence, the Windows node was not

properly configured. With this release, the WMCO bypasses the execution policy when running PowerShell commands on a Windows VM. ([OCPBUGS-37611](#))

- Before this update, aggregated API servers on OpenShift Container Platform used in-memory loopback certificates that were valid for only one year. With this release, these servers now use in-memory loopback certificates that are valid for three years. ([OCPBUGS-74922](#))
- Before this update, service account creation with the **openshift/local/google** provider in earlier OpenShift Container Platform versions caused inconsistent results during service account creation. As a consequence, errors occurred during the apply process, which affected cluster provisioning in Google Cloud. With this release, the issue is resolved by upgrading Terraform and the provider. As a result, service account creation issues are resolved, which reduces inconsistencies during cluster installation in Google Cloud. ([OCPBUGS-76930](#))

1.9.3.1. Updating

To update an OpenShift Container Platform 4.14 cluster to this latest release, see [Updating a cluster using the CLI](#).

1.9.4. RHSA-2026:2990 - OpenShift Container Platform 4.14.62 image release, fixed issues, and security update

Issued: 26 February 2026

OpenShift Container Platform release 4.14.62 is now available. The list of bug fixes that are included in the update is documented in the [RHSA-2026:2990](#) advisory. The RPM packages that are included in the update are provided by the [RHSA-2026:2973](#) 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 4.14.62 --pullspecs
```

1.9.5. Fixed issues

There are no notable fixed issues in this release.

1.9.5.1. Updating

To update an OpenShift Container Platform 4.14 cluster to this latest release, see [Updating a cluster using the CLI](#).

1.9.6. RHBA-2026:1003 - OpenShift Container Platform 4.14.61 bug fix advisory

Issued: 29 January 2026

OpenShift Container Platform release 4.14.61, which includes security updates, is now available. The list of bug fixes that are included in the update is documented in the [RHBA-2026:1003](#) advisory. The RPM packages that are included in the update are provided by the [RHSA-2026:0995](#) 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 4.14.61 --pullspecs
```

1.9.6.1. Bug fixes

There are no notable fixed issues in this release.

1.9.6.2. Updating

To update an existing OpenShift Container Platform 4.14 cluster to this latest release, see [Updating a cluster using the CLI](#).

1.9.7. RHBA-2025:22733 - OpenShift Container Platform 4.14.60 bug fix and security update

Issued: 11 December 2025

OpenShift Container Platform release 4.14.60, which includes security updates, is now available. The list of bug fixes that are included in the update is documented in the [RHBA-2025:22733](#) advisory. The RPM packages that are included in the update are provided by the [RHBA-2025:22731](#) 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 4.14.60 --pullspecs
```

1.9.7.1. Bug fixes

The following bug is fixed with this release:

- Before this update, an unsafe process file system in a container prevented a process from starting, and caused network issues. With this release, the unsafe issue in a container is fixed, and the pods start. As a result, the pods do not stay in the initial startup state, and the Border Gateway Protocol (BGP) operation continues. ([OCPBUGS-58892](#))

1.9.7.2. Updating

To update an existing OpenShift Container Platform 4.14 cluster to this latest release, see [Updating a cluster using the CLI](#).

1.9.8. RHSA-2025:21331 - OpenShift Container Platform 4.14.59 bug fix and security update

Issued: 20 November 2025

OpenShift Container Platform release 4.14.59, which includes security updates, is now available. The list of bug fixes that are included in the update is documented in the [RHSA-2025:21331](#) advisory. The RPM packages that are included in the update are provided by the [RHSA-2025:21328](#) 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 4.14.59 --pullspecs
```

1.9.8.1. Bug fixes

The following bug is fixed with this release:

- Before this update, the Machine API Controller (MAPI) in a hosted control planes management cluster had a memory leak, and caused extreme memory use on the primary node. As a consequence, the cluster performance was affected. With this release, the memory leak in the MAPI controller is fixed, excessive memory use on primary nodes is reduced, and the cluster stability is improved. ([OCPBUGS-65531](#))

1.9.8.2. Updating

To update an existing OpenShift Container Platform 4.14 cluster to this latest release, see [Updating a cluster using the CLI](#).

1.9.9. RHSA-2025:19058 - OpenShift Container Platform 4.14.58 bug fix and security update

Issued: 30 October 2025

OpenShift Container Platform release 4.14.58, which includes security updates, is now available. The list of bug fixes that are included in the update is documented in the [RHSA-2025:19058](#) advisory. The RPM packages that are included in the update are provided by the [RHBA-2025:19040](#) 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 4.14.58 --pullspecs
```

1.9.9.1. Bug fixes

- Before this release, the **properties.storageProfile.[osDiskImage/dataDiskImages].source.Id** field was supported for Microsoft Azure. With this release, use the **properties.storageProfile.osDiskImage.source.storageAccountId** field instead of the **properties.storageProfile.[osDiskImage/dataDiskImages].source.Id** field. ([OCPBUGS-62814](#))

1.9.9.2. Updating

To update an existing OpenShift Container Platform 4.14 cluster to this latest release, see [Updating a cluster using the CLI](#).

1.9.10. RHSA-2025:16165 - OpenShift Container Platform 4.14.57 bug fix and security update

Issued: 25 September 2025

OpenShift Container Platform release 4.14.57, which includes security updates, is now available. The list of bug fixes that are included in the update is documented in the [RHSA-2025:16165](#) advisory. The RPM packages that are included in the update are provided by the [RHBA-2025:16163](#) 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 4.14.57 --pullspecs
```

1.9.10.1. Enhancement

- In OpenShift Container Platform 4.14.57, Operator Lifecycle Manage (OLM) allows Operators to include network policy manifests in their resource bundles. These tailored network policies protect against data leaks and harden against many attack vectors on OpenShift Container Platform clusters.

TIP

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

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

Update to OpenShift Container Platform 4.14.57 or later to enable OLM support for tailored network policies.

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

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

1.9.10.2. Bug fixes

- Before this update, cluster installations failed due to a missing application programming interface (API) endpoint certificate in affected versions. This occurrence caused unresponsive certificate issues and installation problems. With this release, certificate issues during a Hosted control planes cluster installation are resolved, and the Advanced Cluster Management (ACM) agent does not stall. ([OCPBUGS-61176](#))

1.9.10.3. Updating

To update an existing OpenShift Container Platform 4.14 cluster to this latest release, see [Updating a cluster using the CLI](#).

1.9.11. RHSA-2025:14855 - OpenShift Container Platform 4.14.56 bug fix and security update

Issued: 04 September 2025

OpenShift Container Platform release 4.14.56, which includes security updates, is now available. The list of bug fixes that are included in the update is documented in the [RHSA-2025:14855](#) advisory. There is no RPM package for this update.

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 4.14.56 --pullspecs
```

1.9.11.1. Bug fixes

- Before this update, the sidecar of the **openshift-ntp** pod crashed when it stopped during a restart. As a consequence, the clock class metrics were unavailable. With this release, the sidecar for the **openshift-ntp** pod does not stop during a restart. As a result, the clock class metrics are available. ([OCPBUGS-59233](#))

1.9.11.2. Updating

To update an existing OpenShift Container Platform 4.14 cluster to this latest release, see [Updating a cluster using the CLI](#).

1.9.12. RHSA-2025:13289 - OpenShift Container Platform 4.14.55 bug fix and security update

Issued: 14 August 2025

OpenShift Container Platform release 4.14.55, which includes security updates, is now available. The list of bug fixes that are included in the update is documented in the [RHSA-2025:13289](#) advisory. There are no RPM packages for this update.

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 4.14.55 --pullspecs
```

1.9.12.1. Updating

To update an existing OpenShift Container Platform 4.14 cluster to this latest release, see [Updating a cluster using the CLI](#).

1.9.13. RHSA-2025:11669 - OpenShift Container Platform 4.14.54 bug fix and security update

Issued: 31 July 2025

OpenShift Container Platform release 4.14.54, which includes security updates, is now available. The list of bug fixes that are included in the update is documented in the [RHSA-2025:11669](#) advisory. The RPM packages that are included in the update are provided by the [RHBA-2025:11670](#) 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 4.14.54 --pullspecs
```

1.9.13.1. Bug fixes

- Before this update, when a port entered a faulty state, Precision Time Protocol (PTP) pods did not mask the initial summary metrics. As a result, a new interface appeared with a value that remained unchanged, displaying incorrect or stale data. With this update, PTP pods mask and alias the summary metrics correctly, ensuring that interface data is updated as expected. ([OCPBUGS-55309](#))

1.9.13.2. Updating

To update an existing OpenShift Container Platform 4.14 cluster to this latest release, see [Updating a cluster using the CLI](#).

1.9.14. RHSA-2025:9759 - OpenShift Container Platform 4.14.53 bug fix and security update

Issued: 02 July 2025

OpenShift Container Platform release 4.14.53, which includes security updates, is now available. The list of bug fixes that are included in the update is documented in the [RHSA-2025:9759](#) advisory. The RPM packages that are included in the update are provided by the [RHBA-2025:9760](#) 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 4.14.53 --pullspecs
```

1.9.14.1. Bug fixes

- Previously, when users added a custom certificate with a Subject Alternative Name (SAN) that conflicted with the Kubernetes API server (KAS) hostname defined in the **hc.spec.services.servicePublishingStrategy** parameter, the KAS certificate was not included in the new payload generation. This caused an issue with certificate validation for new nodes attempting to join the hosted control planes cluster. With this release, a new validation is implemented to proactively identify and alert users about conflicts ensuring this bug no longer occurs. ([OCPBUGS-57321](#))
- Previously, the incorrect handling of proxy environment variables led to errors with external binaries and failed builds due to improperly formatted proxy settings. With this release, the issue is resolved by excluding proxy variables from the build process when they are not configured. This fix ensures that builds do not fail because of these variables. ([OCPBUGS-56951](#))
- Previously, deploying OpenShift Container Platform clusters with Hosted control planes and a **Kyverno** policy engine for validation resulted in the **Konnectivity** service failing to proxy API pod requests to the validating webhook. This prevented the creation of additional groups within the cluster. With this release, the **Konnectivity** service proxy issue is resolved, which enables the successful creation of additional groups by users. ([OCPBUGS-55936](#))
- Previously, during the Time-Grandmaster (T-GM) holdover, if the global navigation satellite system (GNSS) source was reacquired, the system would incorrectly announce a **T-GM-STATUS** of **S2** before the DPLL (Digital Phase-Locked Loop) had achieved a locked state. This premature announcement compromised T-GM stability and led to inaccurate status reports following GNSS reacquisition. With this release, the DPLL verifies the T-GM status

against the GNSS source availability before declaring an **S2** state. This change improves T-GM stability and ensures status announcements are sent after the GNSS reacquisition. ([OCPBUGS-55467](#))

1.9.14.2. Updating

To update an existing OpenShift Container Platform 4.14 cluster to this latest release, see [Updating a cluster using the CLI](#).

1.9.15. RHSA-2025:7702 - OpenShift Container Platform 4.14.52 bug fix and security update

Issued: 21 May 2025

OpenShift Container Platform release 4.14.52, which includes security updates, is now available. The list of bug fixes that are included in the update is documented in the [RHSA-2025:7702](#) advisory. The RPM packages that are included in the update are provided by the [RHBA-2025:7704](#) 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 4.14.52 --pullspecs
```

1.9.15.1. Bug fixes

- Previously, an Operator would incorrectly appear twice in the installed Operators list if the current project matched its default namespace and the copied CSVs were disabled in the Operator Lifecycle Manager (OLM). With this release, the Operator displays only once. ([OCPBUGS-55942](#))
- Previously, OpenShift Container Platform release 4.14.52 and later releases remained on the stable-4.14 channel instead of switching to the eus-4.14 channel. This caused 'VersionNotFound' errors when attempting to retrieve update information. Because the version was not found in the stable-4.14 channel, the recommended updates were not available. With this release, the installer directs 4.14.52 and later releases to the Extended Update Support (EUS) channel, ensuring the releases receive the recommended update notices, and the **VersonNotFound** error no longer displays. ([OCPBUGS-55193](#))
- 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-54868](#))
- Previously, cluster nodes repeatedly lost communication due to improper remote port binding by Open Virtual Network (OVN)-Kubernetes. This affected pod communication across nodes. With this release, the remote port binding functionality has been updated to be handled by OVN directly, improving the reliability of cluster node communication. ([OCPBUGS-48522](#))
- Previously, when you installed a cluster on IBM Cloud® into an existing Virtual Private Cloud (VPC), the installation program retrieved an unsupported VPC region. If you tried to install into a supported VPC region that followed the unsupported VPC region alphabetically, the installation program crashed. With this release, the installation program ignores any VPC regions that are not fully available during resource lookups. ([OCPBUGS-48196](#))

1.9.15.2. Updating

To update an existing OpenShift Container Platform 4.14 cluster to this latest release, see [Updating a cluster using the CLI](#).

1.9.16. RHSA-2025:4177 - OpenShift Container Platform 4.14.51 bug fix and security update

Issued: 30 April 2025

OpenShift Container Platform release 4.14.51, which includes security updates, is now available. The list of bug fixes that are included in the update is documented in the [RHSA-2025:4177](#) advisory. There are no RPM packages for this update.

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 4.14.51 --pullspecs
```

1.9.16.1. Bug fixes

- Previously, an update to the IBM Cloud® Cloud Internet Services (CIS) implementation impacted the upstream Terraform plugin. If you attempted to create an external-facing cluster on IBM Cloud®, 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-54264](#)).

- Previously, the **Observe** section 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, network observability, and others 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-53437](#))
- Previously, on a **condition.Status**, the ignition-server controller overloaded the Kubernetes agent server (KAS) by updating that condition with the same message in every reconcile loop. With this release, the controller checks the message and validates whether it is the existing message so that the KAS is not overloaded. ([OCPBUGS-53433](#))
- Previously, a custom Security Context Constraint (SCC) impacted pods that the Cluster Version Operator generated from receiving a cluster version upgrade. With this release, OpenShift Container Platform now sets a default SCC to each pod, so that any custom SCC created does not impact a pod. ([OCPBUGS-50592](#))
- Previously, a User Datagram Protocol (UDP) packet that was larger than the maximum transmission unit (MTU) value set for the cluster, could not be sent to the endpoint of the packet using a service. With this release, the pod IP address is used instead of the service IP

address regardless of the packet size so that the UDP packet can be sent to the endpoint. ([OCPBUGS-50584](#))

1.9.16.2. Updating

To update an existing OpenShift Container Platform 4.14 cluster to this latest release, see [Updating a cluster using the CLI](#).

1.9.17. RHSA-2025:3569 - OpenShift Container Platform 4.14.50 bug fix and security update

Issued: 09 April 2025

OpenShift Container Platform release 4.14.50, which includes security updates, is now available. The list of bug fixes that are included in the update is documented in the [RHSA-2025:3569](#) 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 4.14.50 --pullspecs
```

1.9.17.1. Bug fixes

- Previously, installing a cluster on VMware vSphere required specifying the full path to the vSphere datastore. With this release, the installation program accepts full paths and relative paths to the datastore. ([OCPBUGS-54260](#))
- Previously, the **Cluster Settings** page would not properly render during a cluster update if the **ClusterVersion** did not receive a **Completed** update. With this release, the **Cluster Setting** page properly renders even if the **ClusterVersion** has not received a **Completed** update. ([OCPBUGS-54167](#))
- Previously, incorrect addresses were being passed to the Kubernetes EndpointSlice on a cluster. This issue prevented the installation of the MetalLB Operator on an Agent-based cluster in an IPv6 disconnected environment. With this release, a fix modified the address evaluation method. Red Hat Marketplace pods can now successfully connect to the cluster API server so that the installation of MetalLB Operator and handling of ingress traffic in IPv6 disconnected environments can occur. ([OCPBUGS-53314](#))
- Previously, a code migration operation failed to process external labels correctly on the **Home > Overview > Status** pane on the **Administrator perspective** of the web console. These external labels are required to prevent silenced alert notifications from getting added to **Status** pane. Because the **Status** pane did not handle external labels correctly, the pane provided links to **Alert detail** pages but generated a "no matching alerts found" message when you clicked on a link. With this release, the **Status** pane accepts external labels so clicking on an alert links to the correct **Alert detail** page. ([OCPBUGS-51118](#))

1.9.17.2. Updating

To update an existing OpenShift Container Platform 4.14 cluster to this latest release, see [Updating a cluster using the CLI](#).

1.9.18. RHSA-2025:2710 - OpenShift Container Platform 4.14.49 bug fix and security update

Issued: 19 March 2025

OpenShift Container Platform release 4.14.49, which includes security updates, is now available. The list of bug fixes that are included in the update is documented in the [RHSA-2025:2710](#) advisory. The RPM packages that are included in the update are provided by the [RHSA-2025:2712](#) 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 4.14.49 --pullspecs
```

1.9.18.1. Bug fixes

- Previously, after a compute node was removed from the cluster, provisioning of a **PersistentVolume** (PV) resource failed with a message that stated: **The object <virtual machine ID> has already been deleted or has not been completely created.** With this release, a fix ensures that when a compute node is removed, this removal does not impact provisioning of a PV resource. ([OCPBUGS-51045](#))
- Previously, requests to the **deploymentconfig/scale** subresource would fail when there was an admission webhook matching the request. With this release, the issue is resolved and requests to the **deploymentconfig/scale** subresource succeed. ([OCPBUGS-50477](#))
- Previously, when you used the **Form View** to edit **Deployment** or **DeploymentConfig** API objects on the OpenShift Container Platform web console, duplicate **ImagePullSecrets** parameters existed in the YAML configuration for either object. With this release, a fix ensures that duplicate **ImagePullSecrets** parameters do not get automatically added for either object. ([OCPBUGS-49753](#))

1.9.18.2. Updating

To update an existing OpenShift Container Platform 4.14 cluster to this latest release, see [Updating a cluster using the CLI](#).

1.9.19. RHSA-2025:1451 - OpenShift Container Platform 4.14.48 bug fix and security update

Issued: 19 February 2025

OpenShift Container Platform release 4.14.48, which includes security updates, is now available. The list of bug fixes that are included in the update is documented in the [RHSA-2025:1451](#) advisory. The RPM packages that are included in the update are provided by the [RHSA-2025:1453](#) 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 4.14.48 --pullspecs
```

1.9.19.1. Bug fixes

- Previously, an invalid or unreachable identity provider (IDP) blocked updates to hosted control planes. With this release, the **ValidIDPConfiguration** condition in the **HostedCluster** object now reports any IDP errors so that these errors do not block updates to hosted control planes. ([OCPBUGS-49405](#))

1.9.19.2. Updating

To update an existing OpenShift Container Platform 4.14 cluster to this latest release, see [Updating a cluster using the CLI](#).

1.9.20. RHSA-2025:0840 - OpenShift Container Platform 4.14.46 bug fix and security update

Issued: 06 February 2025

OpenShift Container Platform release 4.14.46, which includes security updates, is now available. The list of bug fixes that are included in the update is documented in the [RHSA-2025:0840](#) advisory. The RPM packages that are included in the update are provided by the [RHSA-2025:0842](#) 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 4.14.46 --pullspecs
```

1.9.20.1. Bug fixes

- Previously, crun failed to stop a container if you opened a terminal session and then disconnected from it. With this release, the issue is resolved. ([OCPBUGS-48752](#))
- Previously, you could not remove a "finally" pipeline task from the **edit Pipeline** form if you created a pipeline with only one "finally" task. With this release, you can remove the "finally" task from the **edit Pipeline** form and the issue is resolved. ([OCPBUGS-46603](#))

1.9.20.2. Updating

To update an existing OpenShift Container Platform 4.14 cluster to this latest release, see [Updating a cluster using the CLI](#).

1.9.21. RHSA-2025:0364 - OpenShift Container Platform 4.14.45 bug fix and security update

Issued: 22 January 2025

OpenShift Container Platform release 4.14.45, which includes security updates, is now available. The list of bug fixes that are included in the update is documented in the [RHSA-2025:0364](#) advisory. The RPM packages that are included in the update are provided by the [RHBA-2025:0367](#) 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 4.14.45 --pullspecs
```

1.9.21.1. Bug fixes

- Previously, when using the **SiteConfig** custom resource (CR) to delete a cluster or a node, the **BareMetalHost** CR could get stuck in the **Deprovisioning** state. This bug has been fixed to ensure that the order deletion is correct. The fix requires Red Hat OpenShift GitOps 1.13 or later version. ([OCPBUGS-48339](#))
- Previously, a machine controller failed to save the VMware vSphere task ID of an instance template clone operation. This caused the machine to go into the **Provisioning** state and to power off. With this release, the VMware vSphere machine controller can detect and recover from this state. ([OCPBUGS-48245](#))
- Previously, during node reboots, especially during update operations, the node that interacts with the rebooting machine entered a **Ready=Unknown** state for a short amount of time. This situation caused the Control Plane Machine Set Operator to enter an **UnavailableReplicas** condition and then an **Available=false** state. The **Available=false** state triggers alerts that demand urgent action, but in this case, intervention was only required for a short period of time until the node rebooted. With this release, a grace period for node unreadiness is provided where 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-48211](#))
- Previously the algorithm used to calculate the priority for machine removal equated machines over a certain age as equal priority to machines that were preferred for removal. With this release, the priority of unmarked machines sorted by age is reduced so that they never conflict with machines explicitly marked for removal. The algorithm is also updated to ensure that age ordering is guaranteed for machines up to ten years old. ([OCPBUGS-47659](#))
- Previously, when you upgraded an OpenShift Container Platform cluster from 4.14 to 4.15, the **vCenterCluster** parameter was not populated with a value in the **use-connection-form.ts** configuration file. As a result, the VMware vSphere GUI did not display VMware vSphere vCenter information. With this release, an update to the **Infrastructure** custom resource (CR) ensures it has that the GUI checks the **cloud-provider-config** config map for the **vCenterCluster** value. ([OCPBUGS-45323](#))

1.9.21.2. Updating

To update an existing OpenShift Container Platform 4.14 cluster to this latest release, see [Updating a cluster using the CLI](#).

1.9.22. RHSA-2025:0029 - OpenShift Container Platform 4.14.44 bug fix and security update

Issued: 09 January 2025

OpenShift Container Platform release 4.14.44, which includes security updates, is now available. The list of bug fixes that are included in the update is documented in the [RHSA-2025:0029](#) advisory. The RPM packages that are included in the update are provided by the [RHBA-2025:0032](#) 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 4.14.44 --pullspecs
```

1.9.22.1. Bug fixes

- Previously, the Single-Root I/O Virtualization (SR-IOV) Operator did not expire the acquired lease during the Operator's shutdown operation. This impacted a new instance of the Operator, because the new instance had to wait for the lease to expire before the new instance was ready. With this release, an update to the Operator shutdown logic ensures that the Operator expires the lease when the Operator is shutting down. ([OCPBUGS-44726](#))
- Previously, when you attempted to use Operator Lifecycle Manager (OLM) to upgrade an Operator, the upgrade was blocked and an **error validating existing CRs against new CRD's schema** message was generated. An issue existed with OLM, whereby OLM erroneously identified incompatibility issues with the validation of an existing custom resources definition (CRD) of the the new Operator version. With this release, the validation is corrected so that Operator upgrades are no longer blocked. ([OCPBUGS-46595](#))
- Previously, the **aws-sdk-go-v2** software development kit (SDK) failed to authenticate an **AssumeRoleWithWebIdentity** API operation on an Amazon Web Services (AWS) Security Token Service (STS) cluster. With this release, **pod-identity-webhook** includes a default region so that the authenticate issue no longer exists. ([OCPBUGS-46487](#))
- Previously, when you used the Agent-based Installer to install a cluster on a node that had an incorrect date, the cluster installation failed. With this release, a patch is applied to the live ISO time synchronization that exists in the Agent-based Installer. The patch fixes the date issue and configures the **/etc/chrony.conf** file with the list of additional Network Time Protocol (NTP) servers. You can now set any of these additional NTP servers in the **agent-config.yaml** without experiencing cluster installation issues. ([OCPBUGS-45464](#))
- Previously, there was an error when adding parameters to the Pipeline when the resource field was added to the payload, and as a result resources were deprecated. With this update, the resource fields have been removed from the payload, and you can add parameters to the Pipeline without getting an error. ([OCPBUGS-39368](#))

1.9.22.2. Updating

To update an existing OpenShift Container Platform 4.14 cluster to this latest release, see [Updating a cluster using the CLI](#).

1.9.23. RHSA-2024:11031 - OpenShift Container Platform 4.14.43 bug fix and security update

Issued: 18 December 2024

OpenShift Container Platform release 4.14.43, which includes security updates, is now available. The list of bug fixes that are included in the update is documented in the [RHSA-2024:11031](#) advisory. The RPM packages that are included in the update are provided by the [RHBA-2024:11034](#) 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 4.14.43 --pullspecs
```

1.9.23.1. Bug fixes

- Previously, a trailing period was not permitted in a Kubernetes object name per the DNS 1123

sub-domain name standard. If you configured your AWS DHCP option set with a custom domain name containing a trailing period, the logic for extracting the hostname of EC2 instances and turning it into Kubelet node names would not trim the hostnames for trailing periods. With this release, the logic is updated to trim trailing periods, and trailing periods are permitted within the domain name in the DHCP option set. ([OCPBUGS-46057](#))

- Previously, hosted control planes-based clusters were unable to authenticate through the **oc login** command. The web browser displayed an error when it attempted to retrieve the token after selecting **Display Token**. With this release, **cloud.ibm.com** and other cloud-based endpoints are no longer proxied and authentication is successful. ([OCPBUGS-44279](#))

1.9.23.2. Updating

To update an existing OpenShift Container Platform 4.14 cluster to this latest release, see [Updating a cluster using the CLI](#).

1.9.24. RHSA-2024:10523 - OpenShift Container Platform 4.14.42 bug fix and security update

Issued: 05 December 2024

OpenShift Container Platform release 4.14.42, which includes security updates, is now available. The list of bug fixes that are included in the update is documented in the [RHSA-2024:10523](#) advisory. The RPM packages that are included in the update are provided by the [RHBA-2024:10526](#) 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 4.14.42 --pullspecs
```

1.9.24.1. Bug fixes

- Previously, the approval mechanism for certificate signing requests (CSRs) failed because the node name and internal DNS entry for a CSR did not match in terms of character case differences. With this release, an update to the approval mechanism for CSRs skips case-sensitive checks so that a CSR with a matching node name and internal DNS entry does not fail the check because of character case differences. ([OCPBUGS-44774](#))
- Previously, when the Cluster Version Operator (CVO) pod restarted while it was initializing the synchronization work, the Operator interrupted the guard of the blocked upgrade request. The blocked request was unexpectedly accepted. With this release, the guard of the blocked upgrade request continues after the CVO restarts. ([OCPBUGS-44704](#))
- Previously, when the Cluster Resource Override Operator was unable to completely deploy its operand controller, the Operator would restart the process. Each time the Operator attempted the deployment process, the Operator created a new set of secrets. This resulted in a large number of secrets created in the namespace where the Cluster Resource Override Operator was deployed. With this release, the fixed version correctly processes the service account annotations and only one set of secrets is created. ([OCPBUGS-44435](#))

1.9.24.2. Updating

To update an existing OpenShift Container Platform 4.14 cluster to this latest release, see [Updating a cluster using the CLI](#).

1.9.25. RHSA-2024:9620 - OpenShift Container Platform 4.14.41 bug fix and security update

Issued: 20 November 2024

OpenShift Container Platform release 4.14.41, which includes security updates, is now available. The list of bug fixes that are included in the update is documented in the [RHSA-2024:9620](#) advisory. The RPM packages that are included in the update are provided by the [RHSA-2024:9623](#) 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 4.14.41 --pullspecs
```

1.9.25.1. Bug fixes

- Previously, the Machine Config Operator (MCO) on the vSphere **resolv-prepender** script used systemd directives that were not compatible with old boot image versions of OpenShift Container Platform 4. With this release, OpenShift Container Platform nodes are compatible with old boot images with one of the following solutions: scaling with a boot image 4.13 or later by using manual intervention, or upgrading to a release with this fix. ([OCPBUGS-42111](#))

1.9.25.2. Updating

To update an existing OpenShift Container Platform 4.14 cluster to this latest release, see [Updating a cluster using the CLI](#).

1.9.26. RHSA-2024:8697 - OpenShift Container Platform 4.14.40 bug fix and security update

Issued: 07 November 2024

OpenShift Container Platform release 4.14.40, which includes security updates, is now available. The list of bug fixes that are included in the update is documented in the [RHSA-2024:8697](#) advisory. The RPM packages that are included in the update are provided by the [RHSA-2024:8700](#) 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 4.14.40 --pullspecs
```

1.9.26.1. Bug fixes

- Previously, the Cloud Credential Operator (CCO) for a cluster installed on Google Cloud entered a **Degraded=True** if the cluster was not using manual mode credentials and did not have the **backupdr** API enabled. With this release, an update ensures the cluster does not enter a degraded state when the cluster is configured with this environment. ([OCPBUGS-43821](#))

- Previously, when you attempted to use the **oc import-image** command to import an image in a hosted control planes cluster, the command failed because of access issues with a private image registry. With this release, an update to **openshift-apiserver** pods in a hosted control planes cluster resolves names that use the data plane so that the **oc import-image** command now works as expected with private image registries. ([OCPBUGS-43468](#))
- Previously for hosted control planes, a cluster that used mirroring release images might result in existing node pools to use the hosted cluster's operating system version instead of the **NodePool** version. With this release, a fix ensures that node pools use their own versions. ([OCPBUGS-43368](#))
- Previously, when you used the **must-gather** tool, a Multus Container Network Interface (CNI) log file, **multus.log**, was stored in a node's file system. This situation caused the tool to generate unnecessary debug pods in a node. With this release, the Multus CNI no longer creates a **multus.log** file, and instead uses a CNI plugin pattern to inspect any logs for Multus DaemonSet pods in the **openshift-multus** namespace. ([OCPBUGS-43058](#))
- Previously, when you configured the image registry to use an Microsoft Azure storage account that was located in a resource group other than the cluster's resource group, the Image Registry Operator would become degraded. This occurred because of a validation error. With this release, an update to the Operator allows for authentication only by using a storage account key. Validation of other authentication requirements is not required. ([OCPBUGS-42935](#))

1.9.26.2. Updating

To update an existing OpenShift Container Platform 4.14 cluster to this latest release, see [Updating a cluster using the CLI](#).

1.9.27. RHSA-2024:8235 - OpenShift Container Platform 4.14.39 bug fix and security update

Issued: 23 October 2024

OpenShift Container Platform release 4.14.39, which includes security updates, is now available. The list of bug fixes that are included in the update is documented in the [RHSA-2024:8235](#) advisory. The RPM packages that are included in the update are provided by the [RHSA-2024:8238](#) 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 4.14.39 --pullspecs
```

1.9.27.1. Enhancements

The following enhancement is included in this z-stream release:

1.9.27.1.1. Using the Insights Operator to collect data from a Prometheus metric

- The Insights Operator (IO) now collects data from the **haproxy_exporter_server_threshold** metric of Prometheus. ([OCPBUGS-41918](#))

1.9.27.2. Bug fixes

- Previously, if a connection on port 9637 for Windows nodes was refused, the Kubelet Service Monitor threw a **target down** alert because CRI-O does not run on Windows nodes. With this release, Windows nodes are excluded from the Kubelet Service Monitor. ([OCPBUGS-42603](#))
- Previously, when the Node Tuning Operator (NTO) was configured using the **PerformanceProfiles** specification, it created an **ocp-tuned-one-shot systemd** service which ran prior to kubelet and blocked NTO execution. This prevented Podman from fetching the image. With this release, support for cluster-wide proxy environment variables defined in **/etc/mco/proxy.env** is available. This allows Podman to pull NTO images into environments that need to use the http(s) proxy for out-of-cluster connections. ([OCPBUGS-42567](#))
- 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, the issue is resolved. ([OCPBUGS-41246](#))
- Previously, the inspection process failed if special or invalid characters were present in the serial number of a block device and could not escape the **lsblk** command. With this release, the issue is resolved. ([OCPBUGS-39019](#))
- Previously, after installing the Pipelines Operator, users could still create the Deployment before the Pipeline template became available. With this update, the **Create** button on the **Import from Git** page is disabled if a pipeline template is unavailable for the selected resource. ([OCPBUGS-37353](#))
- Previously, a node registration issue prevented you from using Redfish Virtual Media to add an xFusion bare metal node to your cluster. The issue occurred because the hardware was not fully compliant with Redfish. With this release, you can add xFusion bare-metal nodes to your cluster. ([OCPBUGS-32266](#))

1.9.27.3. Updating

To update an existing OpenShift Container Platform 4.14 cluster to this latest release, see [Updating a cluster using the CLI](#).

1.9.28. RHSA-2024:7184 - OpenShift Container Platform 4.14.38 bug fix and security update

Issued: 3 October 2024

OpenShift Container Platform release 4.14.38, which includes security updates, is now available. The list of bug fixes that are included in the update is documented in the [RHSA-2024:7184](#) advisory. The RPM packages that are included in the update are provided by the [RHSA-2024:7187](#) 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 4.14.38 --pullspecs
```

1.9.28.1. Enhancements

The following enhancement is included in this z-stream release:

1.9.28.1.1. Backporting configurable subnets

- This release includes configurable subnets that you can use to configure the masquerade. You can also use configurable subnets to join and transition subnets to prevent overlaps with IP addresses that are used in your local infrastructure. ([OCPBUGS-38440](#))

1.9.28.2. Bug fixes

- Previously, an **AdditionalTrustedCA** field specified in the Hosted Cluster image configuration was not reconciled into the **openshift-config** namespace as expected and the component was unavailable. With this release, the issue is resolved. ([OCPBUGS-42184](#))
- Previously, Operator Lifecycle Manager (OLM) catalog source pods did not recover from node failure if the **registryPoll** field was **none**. With this release, OLM **CatalogSource** registry pods recover from cluster node failures and the issue is resolved. ([OCPBUGS-42150](#))
- Previously, if you created a hosted cluster by using a proxy to make the cluster reach a control plane from a compute node, the compute node would be unavailable to the cluster. With this release, the proxy settings are updated for the node so that the node can use a proxy to communicate with the control plane successfully. ([OCPBUGS-42021](#))
- Previously, when the Operator Lifecycle Manager (OLM) evaluated a potential upgrade, it 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. ([OCPBUGS-42017](#))
- Previously, the OpenShift Container Platform web console **Topology** view could only display a maximum of 100 nodes. If you attempted to view more than 100 nodes, the web console would output a **Loading is taking longer than expected.** error message. With this release, the **MAX_NODES_LIMIT** parameter for the web console is set to **200** so that the web console can display a maximum of 200 nodes. ([OCPBUGS-41581](#))
- Previously, when you configured a hosted cluster to use an identity provider (IdP) that has either an **http** or **https** endpoint, the IdP hostname did not resolve when sent through the proxy. With this release, a DNS lookup operation checks the IdP before IdP traffic is sent through a proxy, so that IdPs with hostnames can only be resolved by the data plane and verified by the Control Plane Operator (CPO). ([OCPBUGS-41374](#))
- Previously, the Cloud Credential Operator (CCO) would produce an error when starting or restarting a large number of secrets in the cluster. The CCO attempted to fetch the secrets at the same time. With this release, CCO fetches the secrets in batches of 100 and the issue is resolved. ([OCPBUGS-41236](#))
- Previously, the grace period for a node to become ready was not aligned with upstream behavior. Occasionally, the grace period caused a node to cycle between **Ready** and **Not ready** states. With this release, the issue is fixed so that the grace period does not cause a node to cycle between the two states. ([OCPBUGS-39378](#))
- Previously, the **openvswitch** service used older cluster configurations after a cluster upgrade and this caused the **openvswitch** service to stop. With this release, the **openvswitch** service is now restarted after a cluster upgrade so that the service uses the newer cluster configurations. ([OCPBUGS-39192](#))
- Previously, proxying for Operators that run in the control plane of a hosted cluster was performed through proxy settings on the Connectivity agent pod that runs in the data plane. As a consequence, it was not possible to distinguish whether proxying was needed based on application protocol.
For parity with OpenShift Container Platform, IdP communication through HTTPS or HTTP

protocols must be proxied, but LDAP communication must not be proxied. This type of proxying also ignores **NO_PROXY** entries that rely on host names because by the time traffic reaches the Konnectivity agent. This means that only the destination IP address is available. With this release, in hosted clusters, proxy is invoked in the control plane through **konnectivity-https-proxy** and **konnectivity-socks5-proxy**, so that proxying traffic is stopped from the Konnectivity agent. As a result, traffic that is destined for LDAP servers is no longer proxied. Other HTTPS or HTTPS traffic is proxied correctly. The **NO_PROXY** setting is honored when you specify hostnames. ([OCPBUGS-38066](#))

- Previously, proxying for LDAP communication occurred in the Konnectivity agent. By the time traffic reached Konnectivity, its protocol and hostname were no longer available. Consequently, proxying was not done correctly for the OAUTH server pod. It did not distinguish between protocols that require proxying (HTTP or HTTPS) and protocols that do not (LDAP). In addition, it did not honor the **no_proxy** variable that is configured in the **HostedCluster.spec.configuration.proxy** spec.

With this release, you can configure the proxy on the Konnectivity sidecar of the OAUTH server so that traffic is routed appropriately, honoring your **no_proxy** settings. As a result, the OAUTH server can communicate properly with identity providers when a proxy is configured for the hosted cluster. ([OCPBUGS-38060](#))

1.9.28.3. Known issues

- Deploying a self-managed private hosted cluster on AWS fails because the **bootstrap-kubeconfig** file uses an incorrect **api-server** port. As a result, the AWS instances are provisioned but cannot join the hosted cluster as nodes. ([OCPBUGS-42221](#))

1.9.28.4. Updating

To update an existing OpenShift Container Platform 4.14 cluster to this latest release, see [Updating a cluster using the CLI](#).

1.9.29. RHSA-2024:6689 - OpenShift Container Platform 4.14.37 bug fix and security update

Issued: 19 September 2024

OpenShift Container Platform release 4.14.37, which includes security updates, is now available. The list of bug fixes that are included in the update is documented in the [RHSA-2024:6689](#) advisory. There are no RPM packages for this update.

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 4.14.37 --pullspecs
```

1.9.29.1. Updating

To update an existing OpenShift Container Platform 4.14 cluster to this latest release, see [Updating a cluster using the CLI](#).

1.9.30. RHSA-2024:6406 - OpenShift Container Platform 4.14.36 bug fix and security update

Issued: 11 September 2024

OpenShift Container Platform release 4.14.36, which includes security updates, is now available. The list of bug fixes that are included in the update is documented in the [RHSA-2024:6406](#) advisory. The RPM packages that are included in the update are provided by the [RHSA-2024:6412](#) 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 4.14.36 --pullspecs
```

1.9.30.1. Enhancements

The following enhancements are included in this z-stream release:

1.9.30.1.1. Updating the CENTOS 8 references to CENTOS 9

- CENTOS 8 has recently ended its lifecycle. This release updates the CENTOS 8 references to CENTOS 9. ([OCPBUGS-39160](#))

1.9.30.1.2. Collecting Ingress Controller certificates

- The Insights Operator now collects information about all Ingress Controller certificates (NotBefore and NotAfter dates). It aggregates the data into a JSON file in the path 'aggregated/ingress_controllers_certs.json'. ([OCPBUGS-37673](#))

1.9.30.2. Automatic leap seconds handling for PTP grandmaster clocks

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

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

For more information, see [Configuring dynamic leap seconds handling for PTP grandmaster clocks](#) .

1.9.30.3. Bug fixes

- Previously, if a virtual machine (VM) was deleted and the network interface controller (NIC) still existed for that VM, the Microsoft Azure VM verification check crashed. With this release, the verification check gracefully processes the issue without crashing. ([OCPBUGS-39413](#))
- Previously, the **spec.noProxy** field from the cluster-wide proxy was not considered when the Cluster Monitoring Operator (CMO) configured proxy capabilities for a Prometheus remote write endpoint. With this release, the CMO no longer configures proxy capabilities for any remote-write endpoints that have an URL that bypasses the proxy according to the **noProxy** field. ([OCPBUGS-39176](#))
- Previously, when running **oc logs -f <pod>**, the logs would not output anything after the log file was rotated. With this release, kubelet outputs a log file after the file is rotated and the issue is resolved. ([OCPBUGS-38959](#))

- Previously, the OpenShift Container Platform web console failed to restart a bare-metal node. This release fixes this issue so that you can restart a bare metal node by using the OpenShift Container Platform web console. ([OCPBUGS-38053](#))
- Previously, when the Cloud Credential Operator (CCO) checked if passthrough mode permissions were correct, the CCO sometimes received a response from the Google Cloud API about an invalid permission for a project. This bug caused the CCO to enter a degraded state that, in turn, impacted the installation of the cluster. With this release, the CCO checks specifically for this error so that it diagnoses it separately without impacting the installation of the cluster. ([OCPBUGS-37823](#))
- Previously, the TuneD profile could unnecessarily reload an additional time after a custom resource (CR) update. With this release, the TuneD object has been removed and the TuneD profiles are carried directly in the TuneD profile Kubernetes objects. As a result, the issue has been resolved. ([OCPBUGS-37754](#))

1.9.30.4. Updating

To update an existing OpenShift Container Platform 4.14 cluster to this latest release, see [Updating a cluster using the CLI](#).

1.9.31. RHSA-2024:5433 - OpenShift Container Platform 4.14.35 bug fix and security update

Issued: 22 August 2024

OpenShift Container Platform release 4.14.35, which includes security updates, is now available. The list of bug fixes that are included in the update is documented in the [RHSA-2024:5433](#) advisory. The RPM packages that are included in the update are provided by the [RHSA-2024:5436](#) 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 4.14.35 --pullspecs
```

1.9.31.1. Enhancements

The following enhancements are included in this z-stream release:

1.9.31.1.1. Configuring Capacity Reservation by using machine sets

- OpenShift Container Platform release 4.14.35 introduces support for on-demand Capacity Reservation with Capacity Reservation groups on Microsoft Azure clusters. For more information, see *Configuring Capacity Reservation by using machine sets* for [compute](#) or [control plane](#) machine sets. ([OCPCLLOUD-1646](#))

1.9.31.1.2. Updating to Kubernetes v1.27.16

- This release contains the updates resulting from an update to Kubernetes v1.27.16. ([OCPBUGS-37623](#))

1.9.31.2. Bug fixes

- Previously, the **OVNKubernetesNorthdInactive** alert did not fire as expected. With this release, the issue is resolved. ([OCPBUGS-38073](#))
- Previously, a machine config pool (MCP) with a higher **maxUnavailable** value than the amount of unavailable nodes would cause cordoned nodes to receive updates if they were in a certain position on the node list. With this release, a fix ensures that cordoned nodes are not added to a queue to receive updates. ([OCPBUGS-37738](#))
- Previously, **HostedClusterConfigOperator** did not delete the **ImageDigestMirrorSet** (IDMS) object after a user removed the **ImageContentSources** field from the **HostedCluster** object. This caused the IDMS object to remain in the **HostedCluster** object. With this release, **HostedClusterConfigOperator** removes all IDMS resources in the **HostedCluster** object so that this issue no longer exists. ([OCPBUGS-37175](#))
- Previously, when the default gateway of a node was set to **vlan** and multiple network manager connection had the same name, the node would fail as it could not configure the default OVN-Kubernetes bridge. With this release, the **configure-ovs.sh** shell script includes an **nmcli connection show uuid** command that retrieves the correct network manager connection if many connections with the same name exist. ([OCPBUGS-33590](#))

1.9.31.3. Updating

To update an existing OpenShift Container Platform 4.14 cluster to this latest release, see [Updating a cluster using the CLI](#).

1.9.32. RHSA-2024:4960 - OpenShift Container Platform 4.14.34 bug fix and security update

Issued: 7 August 2024

OpenShift Container Platform release 4.14.34, which includes security updates, is now available. The list of bug fixes that are included in the update is documented in the [RHSA-2024:4960](#) advisory. The RPM packages that are included in the update are provided by the [RHSA-2024:4963](#) 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 4.14.34 --pullspecs
```

1.9.32.1. Enhancements

1.9.32.1.1. Adding connectTimeout tuning option to Ingress Controller API

- The IngressController API is updated with a new tuning option, **ingresscontroller.spec.tuningOptions.connectTimeout**, which defines how long the router waits for a response when establishing a connection to a back-end server. ([OCPBUGS-36555](#))

1.9.32.1.2. Collecting Prometheus and AlertManager resources with the Insights Operator

- The Insights Operator now collects the **Prometheus** and **AlertManager** resources outside of the **openshift-monitoring** namespace. ([OCPBUGS-36380](#))

1.9.32.2. Bug fixes

- Previously, AWS HyperShift clusters leveraged their VPC's primary CIDR range to generate security group rules on the data plane. As a result, installing AWS HyperShift clusters into an AWS VPC with multiple CIDR ranges could cause the generated security group rules to be insufficient. With this update, security group rules are generated based on the provided Machine CIDR range instead to resolve this issue. ([OCPBUGS-36159](#))
- Previously, when a pod specified a node selector with no matching nodes, the Kubernetes scheduler panicked with the following error: "Observed a panic: integer divide by zero". With this release, the issue in the Kubernetes scheduler code base is resolved, and the Kubernetes scheduler no longer panics when a pod specifies a node selector with no matching nodes. ([OCPBUGS-36397](#))
- Previously, installing an Operator could sometimes fail if the same Operator had been previously installed and uninstalled. This was due to a caching issue. With this release, the Operator Lifecycle Manager (OLM) is updated to install the Operator correctly in this scenario and the issue is resolved. ([OCPBUGS-36452](#))
- Previously, the Ingress Operator could not successfully update the canary route because the Operator did not have permission to update **spec.host** or **spec.subdomain** on an existing route. With this release, the required permission is added to the cluster role for the Operator's service account and the Ingress Operator can update the canary route. ([OCPBUGS-36467](#))
- Previously, when the OVN-Kubernetes setting for routing-via-host was set to the default value of shared gateway mode, OVN-Kubernetes did not correctly handle traffic streams that mixed non-fragmented and fragmented packets from the IP layer on cluster ingress. With this release, OVN-Kubernetes correctly reassembles and handles external traffic IP packet fragments when ingressing and the issue is resolved. ([OCPBUGS-36554](#))
- Previously, in Amazon Web Services (AWS) security token service (STS), Cloud Credential Operator (CCO) checked **awsSTSRoleARN** in **CredentialsRequest** to create the secret. When **awsSTSRoleARN** was not present, CCO logged an error. With this release, CCO no longer logs the error and the issue is resolved. ([OCPBUGS-36716](#))
- Previously, the Open vSwitch (OVS) pinning procedure set the CPU affinity of the main thread, but other CPU threads did not pick up this affinity if they had already been created. As a consequence, some OVS threads did not run on the correct CPU set and could interfere with the performance of pods with a Quality of Service (QoS) class of **Guaranteed**. With this update, the OVS pinning procedure updates the affinity of each OVS thread, ensuring that all OVS threads run on the correct CPU set. ([OCPBUGS-37197](#))

1.9.32.3. Known issues

- On clusters with the SR-IOV Network Operator installed and configured, pods with a secondary interface of SR-IOV VF fail with the error message: **SRIOV-CNI failed to configure VF "failed to set vf 0 vlan configuration - id 0, qos 0 and proto 802.1q: invalid argument"**. To resolve this issue, upgrade the SR-IOV Network Operator before upgrading OpenShift Container Platform. This ensures that the fix to this issue is included in the SR-IOV Network Operator. ([OCPBUGS-38091](#))

1.9.32.4. Updating

To update an existing OpenShift Container Platform 4.14 cluster to this latest release, see [Updating a cluster using the CLI](#).

1.9.33. RHSA-2024:4479 - OpenShift Container Platform 4.14.33 bug fix and security update

Issued: 17 July 2024

OpenShift Container Platform release 4.14.33, which includes security updates, is now available. The list of bug fixes that are included in the update is documented in the [RHSA-2024:4479](#) advisory. There are no RPM packages in this update.

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 4.14.33 --pullspecs
```

1.9.33.1. Bug fixes

- Previously, nodes that were started by using 4.1 and 4.2 boot images for OpenShift Container Platform got stuck during provisioning because the **machine-config-daemon-firstboot.service** had incompatible machine-config-daemon binary code. With this release, the binary is updated and the issue is resolved. ([OCPBUGS-36776](#))
- Previously, a change of dependency targets was introduced in OpenShift Container Platform 4.14 that prevented disconnected ARO installations from scaling up new nodes after they upgraded to affected versions. With this release, disconnected ARO installs can scale up new nodes after upgrading to OpenShift Container Platform 4.14. ([OCPBUGS-36593](#))
- Previously, if a new deployment was completed at the OSTree level on a host that was identical to the current deployment but on a different stateroot, the OSTree saw them as equal. This behavior incorrectly prevented the boot loader from updating when **set-default** was invoked, as OSTree did not identify the two stateroots as a differentiation factor for deployments. With this release, the OSTree logic is modified to analyze the stateroots and OSTree can properly set the default deployment to a new deployment with different stateroots. ([OCPBUGS-36437](#))
- Previously, the **HighOverallControlPlaneCPU** alert triggered warnings based on criteria for multi-node clusters with high availability. As a result, misleading alerts were triggered in single-node OpenShift clusters because the configuration did not match the environment criteria. This update refines the alert logic to use single-node OpenShift-specific queries and thresholds and account for workload partitioning settings. As a result, CPU utilization alerts in single-node OpenShift clusters are accurate and relevant to single-node configurations. ([OCPBUGS-31354](#))
- Previously, the DNS Operator did not verify that a cluster had ready nodes with available CPU in at least two availability zones, and the DNS daemon set did not use surge for rolling updates. As a result, clusters in which all nodes were in the same availability zone repeatedly emitted **TopologyAwareHintsDisabled** events for the cluster DNS service. With this release, the **TopologyAwareHintsDisabled** events are no longer emitted on clusters that do not have nodes in multiple availability zones, and the issue is resolved. ([OCPBUGS-5943](#))

1.9.33.2. Updating

To update an existing OpenShift Container Platform 4.14 cluster to this latest release, see [Updating a cluster using the CLI](#).

1.9.34. RHSA-2024:4329 - OpenShift Container Platform 4.14.32 bug fix and security updates

Issued: 11 July 2024

OpenShift Container Platform release 4.14.32, which includes security updates, is now available. The list of bug fixes that are included in the update is documented in the [RHSA-2024:4329](#) advisory. The RPM packages that are included in the update are provided by the [RHBA-2024:4332](#) 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 4.14.32 --pullspecs
```

1.9.34.1. Enhancements

The following enhancements are included in this z-stream release:

1.9.34.1.1. New custom resource definitions for pipeline plugin

This release updates the pipelines plugin to support the latest Pipeline Trigger API version for the custom resource definitions (CRDs) **ClusterTriggerBinding**, **TriggerTemplate**, and **EventListener**. ([OCPBUGS-35723](#))

1.9.34.1.2. Controller for defragmentating etcd

This release introduces a controller to defragment etcd for hosted clusters on Hypershift. ([OCPBUGS-35723](#))

1.9.34.2. Bug fixes

- Previously, the **alertmanager-trusted-ca-bundle ConfigMap** was not injected into the user-defined Alertmanager container, which prevented the verification of HTTPS web servers receiving alert notifications. With this update, the trusted CA bundle **ConfigMap** is mounted into the Alertmanager container at the `/etc/pki/ca-trust/extracted/pem/tls-ca-bundle.pem` path. ([OCPBUGS-36416](#))
- Previously, for clusters updated from older versions of OpenShift Container Platform, enabling **kdump** on an OVN-enabled cluster sometimes prevented the node from rejoining the cluster or returning to the **Ready** state. This fix removes problematic stale data from older OpenShift Container Platform versions and ensures this kind of stale data is always cleaned up. The node can now start correctly and rejoin the cluster. ([OCPBUGS-36356](#))
- Previously, a bug in **growpart** caused the device to lock which, prevented the LUKS-encrypted devices from opening. The system would fail to boot successfully. With this release, **growpart** was removed from the process and the the system will boot successfully. ([OCPBUGS-35989](#))
- Previously, in user provisioned infrastructure clusters that were updated from older versions, **failureDomains** might be missing in Infrastructure objects which caused certain checks to fail. With this release, a fallback **failureDomains** is synthesized from **cloudConfig** if none are available in **infrastructures.config.openshift.io**. ([OCPBUGS-35913](#))
- Previously, when installing a cluster on VMware vSphere, the installation failed if an ESXi host

was in maintenance mode because the installation program could not retrieve version information from the host. With this update, the installation program does not attempt to retrieve version information from ESXi hosts that are in maintenance mode, allowing the installation to proceed. ([OCPBUGS-35827](#))

- Previously, a bug in systemd might cause the **coreos-multipath-trigger.service** unit to hang forever. The system would fail to finish booting. With this release, systemd was removed and the boot is successful. ([OCPBUGS-35750](#))
- Previously, registry overrides configured by a cluster administrator on the management side applied to non-relevant data-plane components. With this release, registry overrides no longer apply to these components. ([OCPBUGS-35549](#)).
- Previously, the OpenShift Cluster Manager container did not have the right TLS Certificates. As a result, image streams could not be used in disconnected deployments. With this update, the TLS Certificates are added as projected volumes. ([OCPBUGS-35482](#))
- Previously, in disconnected environments, the HyperShift Operator ignored registry overrides. As a consequence, changes to node pools were ignored, and node pools encountered errors. With this update, the metadata inspector works as expected during the HyperShift Operator reconciliation, and the override images are properly populated. ([OCPBUGS-35401](#))
- Previously, the secrets-store CSI driver on Hypershift was failing to mount secrets due to an issue with the Hypershift CLI. With this release, the driver is able to mount volumes and the issue has been resolved. ([OCPBUGS-35183](#))
- Previously, the reduction of the network queue did not work as expected for inverted rules such as **!ens0**. This happened because the exclamation mark symbol was duplicated in the generated tuned profile. With this release, the duplication no longer occurs so that inverted rules apply as intended. ([OCPBUGS-35012](#))
- Previously, a race condition meant the Kubelet might report a false error relating to a failed resize from the volume resizer. With this release, the race condition has been fixed and the false error is no longer printed. ([OCPBUGS-33964](#))
- Previously, escaped strings were not handled properly when editing a vSphere connection, causing broken vSphere configuration. With this update, the escape strings work as expected and the vSphere configuration no longer breaks. ([OCPBUGS-33942](#))
- Previously, when the default gateway of a node was set to **vlan** and multiple network manager connections had the same name, the node would fail as it could not configure the default OVN-Kubernetes bridge. With this release, the **configure-ovs.sh** shell script includes an **nmcli connection show uuid** command that retrieves the correct network manager connection if many connections with the same name exist. ([OCPBUGS-33590](#))
- Previously, when manually restarting the **kubelet** service on a node, some state files were deleted after an assumed node reboot, which led to kubelet resetting the CPU Manager state. After the state reset, the CPU Manager computes new CPU assignments to running workloads. As a result, the new and initial **cpuset** configuration might differ. With this update, the **cpuset** configuration is correctly restored after a kubelet restart. ([OCPBUGS-32472](#))

1.9.34.3. Updating

To update an existing OpenShift Container Platform 4.14 cluster to this latest release, see [Updating a cluster using the CLI](#).

1.9.35. RHSA-2024:4010 - OpenShift Container Platform 4.14.31 bug fix and security updates

Issued: 26 June 2024

OpenShift Container Platform release 4.14.31, which includes security updates, is now available. The list of bug fixes that are included in the update is documented in the [RHSA-2024:4010](#) advisory. The RPM packages that are included in the update are provided by the [RHBA-2024:4013](#) 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 4.14.31 --pullspecs
```

1.9.35.1. Bug fixes

- Previously, the logic for fetching the names of nodes from the host was not accounting for multiple values and terminated unexpectedly when it returned multiple values for the name that contained a space. With this release, the logic is updated to use only the first returned hostname as the node name and the issue has been resolved. ([OCPBUGS-34716](#))

1.9.35.2. Updating

To update an existing OpenShift Container Platform 4.14 cluster to this latest release, see [Updating a cluster using the CLI](#).

1.9.36. RHSA-2024:3881 - OpenShift Container Platform 4.14.30 bug fix and security updates

Issued: 19 June 2024

OpenShift Container Platform release 4.14.30, which includes security updates, is now available. The list of bug fixes that are included in the update is documented in the [RHSA-2024:3881](#) advisory. The RPM packages that are included in the update are provided by the [RHSA-2024:3918](#) 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 4.14.30 --pullspecs
```

1.9.36.1. Enhancements

The following enhancements are included in this z-stream release:

1.9.36.1.1. Allowing a colon character in pull secret passwords

This release adds the ability to include a colon character in your pull secret password for the OpenShift Container Platform Assisted Installer. ([OCPBUGS-35034](#))

1.9.36.1.2. Configuring the topology feature for the Cinder CSI Driver

Previously, the topology feature for the Cinder CSI Driver was always active, because you could not disable the topology feature. With this release, the topology feature is based on the compute and storage availability zones and you can disable the topology feature. ([OCPBUGS-34792](#))

1.9.36.2. Bug fixes

- Previously, the OpenShift Container Platform Assisted Installer reported installed SATA SDDs as removable and did not use any of them as installation targets. With this release, removable disks are eligible for installation and the issue has been resolved. ([OCPBUGS-35085](#))
- Previously, an Amazon Web Services (AWS) policy issue prevented the Cluster API Provider AWS from retrieving the necessary domain information. As a consequence, installing an AWS hosted cluster with a custom domain failed. With this update, the policy issue is resolved. ([OCPBUGS-34856](#))
- Previously, when you deleted a cluster or BareMetal Host (BMH), a **PreprovImage** image was created during cluster deletion. As a result, the cluster resources were stuck. With this release, an exception is created for the powering off before the delete stage and the issue has been resolved. ([OCPBUGS-34814](#)).
- Previously, when you enabled the **virtualHostedStyle** parameter with **regionEndpoint** in the Image Registry Operator configuration, the image registry ignored **virtualHostedStyle** and failed to start. This release drops the use of **virtualHostedStyle** and uses **ForcePathStyle** instead which fixes the issue. ([OCPBUGS-34668](#))
- Previously, when upgrading to OpenShift Container Platform 4.15.11 from 4.15.8, the **metal3-ironic** and **metal3-ironic-inspector** pods would fail due to an installation failure related to FIPS mode enablement. With this release, the issue has been resolved. ([OCPBUGS-34657](#)).
- Previously, the OVN-Kubernetes network plugin could not send gratuitous Address Resolution Protocol (ARP) requests to peers to inform them about the medium access control (MAC) address of a new node, because some situations caused the transfer of the Egress IP address from one node to another to fail. This resulted in failover issues. With this release, the OVN-Kubernetes network plugin correctly informs peers about the medium access control (MAC) address of a new node, without causing a failover issue. ([OCPBUGS-34570](#))
- Previously, if you used the **wait-for-ceo** command during the bootstrapping process, the command would not report the error messages in case of a failure. With this release, the command reports the error messages in the **bootkube** script for you to view. ([OCPBUGS-34495](#))
- Previously, the installation program did not support the **ca-west-1** Amazon Web Services (AWS) region. With this release, the **ca-west-1** region is supported and the issue has been resolved. ([OCPBUGS-34024](#))

1.9.36.3. Updating

To update an existing OpenShift Container Platform 4.14 cluster to this latest release, see [Updating a cluster using the CLI](#).

1.9.37. RHBA-2024:3697 - OpenShift Container Platform 4.14.29 bug fix and security updates

Issued: 13 June 2024

OpenShift Container Platform release 4.14.29, which includes security updates, is now available. The list of bug fixes that are included in the update is documented in the [RHBA-2024:3697](#) advisory. The RPM packages that are included in the update are provided by the [RHSA-2024:3700](#) 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 4.14.29 --pullspecs
```

1.9.37.1. Bug fixes

- Previously, for OpenShift Container Platform deployments on Red Hat OpenStack Platform (RHOSP), the **MachineSet** object did not correctly apply the value for the **Port Security** parameter. This meant that the **port_security_enabled** parameter in the RHOSP server's port had an unexpected value. With this release, the **MachineSet** object applies the **port_security_enabled** flag as expected. ([OCPBUGS-32428](#))
- Previously, when an **IngressController** object was configured with client SSL/TLS, but did not have the **clientca-configmap** finalizer, the Ingress Operator tried to add the finalizer without checking whether the **IngressController** object was marked for deletion. Consequently, if an **IngressController** was configured with client SSL/TLS and was subsequently deleted, the Operator would correctly remove the finalizer. The Operator would then repeatedly, and erroneously, try and fail to update the **IngressController** object to add the finalizer back, resulting in error messages in the Operator's logs. With this update, the Ingress Operator no longer adds the **clientca-configmap** finalizer to an **IngressController** object that is marked for deletion. As a result, the Ingress Operator no longer tries to perform erroneous updates and no longer logs the associated errors. ([OCPBUGS-34410](#))

1.9.37.2. Updating

To update an existing OpenShift Container Platform 4.14 cluster to this latest release, see [Updating a cluster using the CLI](#).

1.9.38. RHSA-2024:3523 - OpenShift Container Platform 4.14.28 bug fix update and security update

Issued: 10 June 2024

OpenShift Container Platform release 4.14.28, which includes security updates, is now available. The list of bug fixes that are included in the update is documented in the [RHSA-2024:3523](#) advisory. The RPM packages that are included in the update are provided by the [RHBA-2024:3526](#) 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 4.14.28 --pullspecs
```

1.9.38.1. Bug fixes

- Previously, the HyperShift Operator was not using the **RegistryOverrides** mechanism to inspect the image from the internal registry. With this release, the metadata inspector works as

expected during the HyperShift Operator reconciliation, and the **OverrideImages** are properly populated. ([OCPBUGS-33844](#))

- Previously, the recycler pods for Hosted Control Planes (HCPs) were not starting in disconnected environments. With this release, the HCP recycler-pod image now points to the OpenShift Container Platform payload reference and the issue has been resolved. ([OCPBUGS-33843](#))
- Previously, the information from **imageRegistryOverrides** was only extracted once on the HyperShift Operator initialization and did not refresh. With this release, the HyperShift Operator retrieves the new **ImageContentSourcePolicy** files from the management cluster and adds them to the HyperShift Operator and Control Plane Operator in every reconciliation loop. ([OCPBUGS-33713](#))
- 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. ([OCPBUGS-33321](#))

1.9.38.2. Updating

To update an existing OpenShift Container Platform 4.14 cluster to this latest release, see [Updating a cluster using the CLI](#).

1.9.39. RHSA-2024:3331 - OpenShift Container Platform 4.14.27 bug fix update and security update

Issued: 30 May 2024

OpenShift Container Platform release 4.14.27, which includes security updates, is now available. The list of bug fixes that are included in the update is documented in the [RHSA-2024:3331](#) advisory. The RPM packages that are included in the update are provided by the [RHBA-2024:3335](#) 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 4.14.27 --pullspecs
```

1.9.39.1. Bug fixes

- 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-33537](#))
- Previously, the **Topology** view in the OpenShift Container Platform web console did not show the visual connector between a virtual machine (VM) node and other non-VM components. With this release, the visual connector shows interaction activity of a component. ([OCPBUGS-33640](#))

- Previously, a logo in the masthead element of the OpenShift Container Platform web console could grow beyond 60 pixels in height. This caused the masthead to increase in height. With this release, the masthead logo is constrained to a **max-height** of 60 pixels. ([OCPBUGS-33635](#))

1.9.39.2. Updating

To update an existing OpenShift Container Platform 4.14 cluster to this latest release, see [Updating a cluster using the CLI](#).

1.9.40. RHSA-2024:2869 - OpenShift Container Platform 4.14.26 bug fix update and security update

Issued: 23 May 2024

OpenShift Container Platform release 4.14.26, which includes security updates, is now available. The list of bug fixes that are included in the update is documented in the [RHSA-2024:2869](#) advisory. The RPM packages that are included in the update are provided by the [RHBA-2024:2873](#) 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 4.14.26 --pullspecs
```

1.9.40.1. Enhancements

The following enhancement is included in this z-stream release:

1.9.40.1.1. OperatorHub filter renamed from FIPS Mode to Designed for FIPS

- Previously, OperatorHub included a filter named **FIPS Mode**. With this release, that filter is named **Designed for FIPS**. ([OCPBUGS-33110](#))

1.9.40.2. Bug fixes

- Previously, when the **ContainerRuntimeConfig** resource was created as an extra manifest for single-node OpenShift Container Platform installation, the bootstrap failed with the following error message: "more than one ContainerRuntimeConfig found that matches MCP labels". With this release, the incorrect processing of the **ContainerRuntimeConfig** resource is fixed and the issue has been resolved. ([OCPBUGS-30153](#))
- 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-32319](#))
- Previously, for OpenShift Container Platform deployments on Red Hat OpenStack Platform (RHOSP), the **MachineSet** object did not correctly apply the value for the **Port Security** parameter. With this release, the **MachineSet** object applies the **port_security_enabled** flag as expected. ([OCPBUGS-32428](#))
- Previously, static Persistent Volumes in Azure File on Workload Identity clusters could not be configured due to an issue with the driver. With this release, the issue has been resolved and static Persistent Volumes mount correctly. ([OCPBUGS-33039](#))

- 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-33389](#))
- Previously, the Ironic Python Agent (IPA) failed when trying to wipe disks because it expected the wrong byte sector size, which caused the node provisioning to fail. With this release, the IPA checks the disk sector size and node provisioning succeeds. ([OCPBUGS-33452](#))
- Previously, attempting to remove an alternate service when editing a Route by using the form view did not remove the alternate service from the Route. With this update, the alternate service is removed and the issue has been resolved. ([OCPBUGS-33462](#))
- Previously, the **vsphere-problem-detector** Operator could not connect to vCenter because the operator did not have the HTTP(S) proxy configured. With this release, the **vsphere-problem-detector** Operator uses the same HTTP(S) proxy as the rest of the cluster and the issue has been resolved. ([OCPBUGS-33467](#))

1.9.40.3. Updating

To update an existing OpenShift Container Platform 4.14 cluster to this latest release, see [Updating a cluster using the CLI](#).

1.9.41. RHBA-2024:2789 - OpenShift Container Platform 4.14.25 bug fix update

Issued: 16 May 2024

OpenShift Container Platform release 4.14.25 is now available. The list of bug fixes that are included in the update is documented in the [RHBA-2024:2789](#) advisory. The RPM packages that are included in the update are provided by the [RHBA-2024:2792](#) 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 4.14.25 --pullspecs
```

1.9.41.1. 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-32482](#))
- Previously, timeout values larger than what the Go programming language 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, it is capped at the maximum that HAProxy can parse. As a result, issues are no longer caused for HAProxy. ([OCPBUGS-30773](#))
- Previously, when users imported image stream tags, **ImageContentSourcePolicy** (ICSP) could not co-exist with **ImageDigestMirrorSet** (IDMS) and **ImageTagMirrorSet** (ITMS). OpenShift Container Platform ignored any IDMS/ITMS created by the user and favored ICSP. With this release, the image stream tags can co-exist because importing image stream tags now respect IDMS/ITMS when ICSP is also present. ([OCPBUGS-31509](#))

- Previously, after you performed a Control Plane Only update on your OpenShift Container Platform cluster that involved pausing and unpausing the machine config pool, two reboot operations occurred after the unpauses operation. This additional reboot was not expected and was caused by the performance profile controller being reconciled against an older **MachineConfig** object that is listed in the **MachineConfigPool** object. With this release, the performance profile controller reconciles against the latest **MachineConfig** object that is listed in the **MachineConfigPool** object so that the extra reboot does not occur. ([OCPBUGS-32980](#))
- Previously, a kernel regression that was introduced in OpenShift Container Platform 4.14.14 caused kernel issues, such as nodes crashing and rebooting, in nodes that mounted to CephFS storage. In this release, the regression issue is fixed so that the kernel regression issue no longer occurs. ([OCPBUGS-33251](#))
- Previously, **ovs-if-br-ex.nmconnection.*** files caused the failure of **ovs-configuration.service**, which resulted in nodes being moved to the **NotReady** state. With this release, **ovs-if-br-ex.nmconnection.*** files are removed from **/etc/NetworkManager/system-connections**, so that this issue no longer exists. ([OCPBUGS-32341](#))

1.9.41.2. Updating

To update an existing OpenShift Container Platform 4.14 cluster to this latest release, see [Updating a cluster using the CLI](#).

1.9.42. RHSA-2024:2668 - OpenShift Container Platform 4.14.24 bug fix update and security update

Issued: 9 May 2024

OpenShift Container Platform release 4.14.24, which includes security updates, is now available. The list of bug fixes that are included in the update is documented in the [RHSA-2024:2668](#) advisory. The RPM packages that are included in the update are provided by the [RHSA-2024:2672](#) 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 4.14.24 --pullspecs
```

1.9.42.1. Enhancements

The following enhancements are included in this z-stream release:

1.9.42.1.1. IPv6 unsolicited neighbor advertisements now default on macvlan CNI plug-in

- Pods created using the macvlan CNI plug-in, where the IP address management CNI plug-in has assigned IP addresses, now send IPv6 unsolicited neighbor advertisements by default onto the network. This notifies hosts of the new pod's MAC address for a particular IP address to refresh IPv6 neighbor caches. ([OCPBUGS-33066](#))

1.9.42.2. Bug fixes

- Previously, when a cluster was installed using proxy and the proxy information contained escape characters in the format **%XX** the installation would fail. With this release, the issue has been fixed. ([OCPBUGS-33010](#))

- Previously, in hosted control planes for OpenShift Container Platform, when you created the custom resource definition (CRD) for **ImageDigestMirrorSet** and **ImageContentSourcePolicy** objects at the same time in a disconnected environment, the HyperShift Operator created the object only for the **ImageDigestMirrorSet** CRD, ignoring the **ImageContentSourcePolicy** CRD. With this release, the HyperShift Operator creates objects for the **ImageDigestMirrorSet** and **ImageContentSourcePolicy** CRDs. ([OCPBUGS-32471](#))
- Previously, nodes of paused MachineConfigPools might be incorrectly unpaused when performing a cluster update. With this update, nodes of paused MachineConfigPools correctly stay paused when performing a cluster update. ([OCPBUGS-32168](#))
- Previously, the image registry did not support Amazon Web Services (AWS) region **ca-west-1**. With this release, the image registry can now be deployed in this region. ([OCPBUGS-31857](#))
- Previously, Terraform would create the compute server group with the policy set for the control plane. As a consequence, the **serverGroupPolicy** property of the **install-config.yaml** file was ignored for the compute server group. With this release, the server group policy set in the **install-config.yaml** file for the compute machine pool is correctly applied at install-time in the Terraform flow. ([OCPBUGS-31756](#))

1.9.42.3. Updating

To update an existing OpenShift Container Platform 4.14 cluster to this latest release, see [Updating a cluster using the CLI](#).

1.9.43. RHBA-2024:2051 - OpenShift Container Platform 4.14.23 bug fix update and security update

Issued: 2 May 2024

OpenShift Container Platform release 4.14.23, which includes security updates, is now available. The list of bug fixes that are included in the update is documented in the [RHBA-2024:2051](#) advisory. The RPM packages that are included in the update are provided by the [RHSA-2024:2054](#) 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 4.14.23 --pullspecs
```

1.9.43.1. Enhancements

The following enhancements are included in this z-stream release:

1.9.43.1.1. Egress IP verification step for additional hops

- Previously, if egress IPs were hosted by any interface other than the primary interface, there was no verification to determine if a next hop is required. With this release, the IP will inspect the main routing table and determine if a next hop is required. ([OCPBUGS-31854](#))

1.9.43.1.2. New profile for RT kernel to drop unsupported parameters

- Previously, the **net.core.busy_read**, **net.core.busy_poll**, and **kernel.numa_balancing sysctl** parameters did not exist within the RT kernel and were therefore unsupported. With this release,

the **openshift-node-performance-rt** profile is added and included if an RT kernel is detected, which drops the unsupported kernel parameters before they are applied. ([OCPBUGS-31905](#))

1.9.43.1.3. Disable option for OLM default source

- Previously, there was no way to disable the Operator Lifecycle Manager (OLM) default source in a disconnected situation. With this release, the **OperatorHubSpec** field is integrated into the **hostedcluster.Spec.Configuration** API to facilitate disabling and enabling default sources during creation. The CLI also includes a flag for this functionality. ([OCPBUGS-32221](#))

1.9.43.2. Bug fixes

- Previously, the Node Tuning Operator (NTO) checked if there were profiles that share the same priority, regardless of their associated node. The process was for the NTO to first collect the profiles, check for priority conflicts, and then filter for the associated node. As a result, if multiple performance profiles were present on two different nodes, false priority warnings were dumped into the logs. With this release, the steps of this process have been changed so that the NTO filters the associated nodes first and then checks for priority conflicts. ([OCPBUGS-31735](#))
- Previously, there was a fundamental issue that prevented egress IPv6 from working for Elastic IPs (EIPs) with multi-network interface controllers (NICs). With this release, the issue has been resolved. ([OCPBUGS-31853](#))
- Previously, certain HTTP clients caused Ingress traffic to degrade after upgrading to OpenShift Container Platform 4.14 when closed idle connections were erroneously reused. With this release, the issue has been resolved. ([OCPBUGS-32437](#))
- Previously, the image registry's Azure path fix job incorrectly required the presence of client and tenant IDs to function, which caused valid configurations to produce validation errors. With this release, a check to account for key-in connection to missing client and tenant IDs is added. ([OCPBUGS-32450](#))

1.9.43.3. Updating

To update an existing OpenShift Container Platform 4.14 cluster to this latest release, see [Updating a cluster using the CLI](#).

1.9.44. RHSA-2024:1891 - OpenShift Container Platform 4.14.22 bug fix update and security update

Issued: 25 April 2024

OpenShift Container Platform release 4.14.22, which includes security updates, is now available. The list of bug fixes that are included in the update is documented in the [RHSA-2024:1891](#) advisory. The RPM packages that are included in the update are provided by the [RHSA-2024:1897](#) 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 4.14.22 --pullspecs
```

1.9.44.1. Enhancements

1.9.44.1.1. Number of configured control plane replicas validated

- Previously, you could set the number of control plane replicas to an invalid value, such as **2**. With this release, a validation is added to prevent any misconfiguration of the control plane replicas at the ISO generation time. ([OCPBUGS-31885](#))

1.9.44.2. Bug fixes

- Previously, the **network-tools** image, which is a debugging tool, included the Wireshark network protocol analyzer. Wireshark had a dependency on the **gstreamer1** package, and this package has specific licensing requirements. With this release, the **gstreamer1** package is removed from the **network-tools** image and the image now includes the **wireshark-cli** package. ([OCPBUGS-31862](#))
- Previously, an external neighbor could change its Media Access Control (MAC) address while the cluster was shutting down or hibernating. Although a Gratuitous Address Resolution Protocol (GARP) was supposed to inform the other neighbors about this change, the cluster did not process the GARP. After restarting the cluster, the neighbor might no longer be available from the OVN-Kubernetes cluster network because the outdated MAC address was being used. With this release, an update enables an aging mechanism so that a neighbor's MAC address is updated regularly every 300 seconds. ([OCPBUGS-11710](#))

1.9.44.3. Updating

To update an existing OpenShift Container Platform 4.14 cluster to this latest release, see [Updating a cluster using the CLI](#).

1.9.45. RHSA-2024:1765 - OpenShift Container Platform 4.14.21 bug fix update and security update

Issued: 18 April 2024

OpenShift Container Platform release 4.14.21, which includes security updates, is now available. The list of bug fixes that are included in the update is documented in the [RHSA-2024:1765](#) advisory. The RPM packages that are included in the update are provided by the [RHBA-2024:1768](#) 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 4.14.21 --pullspecs
```

1.9.45.1. Bug fixes

- Previously, the console backend proxy server was sending operand list requests to the public API server endpoint. This caused Certificate Authority (CA) issues under some circumstances. With this release, the proxy configuration was updated to point to the internal API server endpoint which fixed this issue. ([OCPBUGS-29783](#))

1.9.45.2. Updating

To update an existing OpenShift Container Platform 4.14 cluster to this latest release, see [Updating a cluster using the CLI](#).

1.9.46. RHSA-2024:1681 - OpenShift Container Platform 4.14.20 bug fix update and security update

Issued: 8 April 2024

OpenShift Container Platform release 4.14.20, which includes security updates, is now available. The list of bug fixes that are included in the update is documented in the [RHSA-2024:1681](#) 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 4.14.20 --pullspecs
```

1.9.46.1. Updating

To update an existing OpenShift Container Platform 4.14 cluster to this latest release, see [Updating a cluster using the CLI](#).

1.9.47. RHBA-2024:1564 - OpenShift Container Platform 4.14.19 bug fix update and security update

Issued: 3 April 2024

OpenShift Container Platform release 4.14.19, which includes security updates, is now available. The list of bug fixes that are included in the update is documented in the [RHBA-2024:1564](#) advisory. The RPM packages that are included in the update are provided by the [RHSA-2024:1567](#) 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 4.14.19 --pullspecs
```

1.9.47.1. Bug fixes

- Previously, pods without an IP in their status failed to trigger a new reconciliation loop when being processed by the Admin Policy Based (APP) controller, which caused the logic that adds their configuration to the north-bound DB to go missing. With this release, pods without an IP in their status field continue to be processed by the controller on each event change until their IP field is populated and the controller can complete the reconciliation loop. ([OCPBUGS-29342](#))

1.9.47.2. Updating

To update an existing OpenShift Container Platform 4.14 cluster to this latest release, see [Updating a cluster using the CLI](#).

1.9.48. RHSA-2024:1458 - OpenShift Container Platform 4.14.18 bug fix update and security update

Issued: 27 March 2024

OpenShift Container Platform release 4.14.18, which includes security updates, is now available. The list of bug fixes that are included in the update is documented in the [RHSA-2024:1458](#) advisory. The RPM packages that are included in the update are provided by the [RHSA-2024:1461](#) 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 4.14.18 --pullspecs
```

1.9.48.1. Bug fixes

- Previously, under certain conditions the installation program would fail with the error message: **unexpected end of JSON input**. With this release, the error message is clarified and suggests users set the **serviceAccount** field in the **install-config.yaml** configuration file to fix the issue. ([OCPBUGS-30027](#))

1.9.48.2. Known issues

- Currently, providing a performance profile as an extra manifest when installing a OpenShift Container Platform cluster is not supported. ([OCPBUGS-18640](#))

1.9.48.3. Updating

To update an existing OpenShift Container Platform 4.14 cluster to this latest release, see [Updating a cluster using the CLI](#).

1.9.49. RHBA-2024:1260 - OpenShift Container Platform 4.14.17 bug fix update

Issued: 20 March 2024

OpenShift Container Platform release 4.14.17 is now available. The list of bug fixes that are included in the update is documented in the [RHBA-2024:1260](#) advisory. The RPM packages that are included in the update are provided by the [RHBA-2024:1263](#) 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 4.14.17 --pullspecs
```

1.9.49.1. Updating

To update an existing OpenShift Container Platform 4.14 cluster to this latest release, see [Updating a cluster using the CLI](#).

1.9.50. RHBA-2024:1205 - OpenShift Container Platform 4.14.16 bug fix update

Issued: 13 March 2024

OpenShift Container Platform release 4.14.16 is now available. The list of bug fixes that are included in the update is documented in the [RHBA-2024:1205](#) advisory. The RPM packages that are included in the update are provided by the [RHBA-2024:1208](#) 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 4.14.16 --pullspecs
```

1.9.50.1. Updating

To update an existing OpenShift Container Platform 4.14 cluster to this latest release, see [Updating a cluster using the CLI](#).

1.9.51. RHBA-2024:1046 - OpenShift Container Platform 4.14.15 bug fix update

Issued: 4 March 2024


OpenShift Container Platform release 4.14.15 is now available. The list of bug fixes that are included in the update is documented in the [RHBA-2024:1046](#) advisory. The RPM packages that are included in the update are provided by the [RHBA-2024:1049](#) 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 4.14.15 --pullspecs
```

1.9.51.1. Bug fixes

- Previously, the **manila-csi-driver-controller-metrics** service had empty endpoints because of 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-23443](#))
- Previously, an Amazon Web Services (AWS) code that provided image credentials was removed from the kubelet in OpenShift Container Platform 4.14. Consequently, pulling images from Amazon Elastic Container Registry (ECR) failed without a specified pull secret, because the kubelet could no longer authenticate itself and pass credentials to the container runtime. With this update, a separate credential provider has been configured, which is now responsible for providing ECR credentials for the kubelet. As a result, the kubelet can now pull private images from ECR. ([OCPBUGS-29630](#))
- Previously, the OpenShift Container Platform 4.14 release introduced a change that gave users the perception that their images were lost when updating from OpenShift Container Platform version 4.13 to 4.14. A change to the default internal registry caused the registry to use an incorrect path when using the Microsoft Azure object storage. With this release, the correct path is used and a job has been added to the registry operator that moves any blobs pushed to the registry that used the wrong storage path into the correct storage path, which effectively merges the two distinct storage paths into a single path. ([OCPBUGS-29604](#))



NOTE

This fix does not work on Azure Stack Hub. For Azure Stack Hub users who used OpenShift Container Platform versions 4.14.0 through to 4.14.13 when upgrading to 4.14.14 and later versions will need to complete manual steps to move their object blobs to the correct storage path. See the [Red Hat Knowledgebase article](#).

- Previously, machine sets that ran on Microsoft Azure regions with no availability zone support always created **AvailabilitySets** objects for Spot instances. This operation caused Spot instances to fail because the instances did not support availability sets. Now, machine sets do not create **AvailabilitySets** objects for Spot instances that operate in non-zonal configured regions. ([OCPBUGS-29152](#))

1.9.51.2. Updating

To update an existing OpenShift Container Platform 4.14 cluster to this latest release, see [Updating a cluster using the CLI](#).

1.9.52. RHSA-2024:0941 - OpenShift Container Platform 4.14.14 bug fix and security update

Issued: 28 February 2024

OpenShift Container Platform release 4.14.14, which includes security updates, is now available. The list of bug fixes that are included in the update is documented in the [RHSA-2024:0941](#) advisory. The RPM packages that are included in the update are provided by the [RHSA-2024:0944](#) 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 4.14.14 --pullspecs
```

1.9.52.1. Enhancements

The following enhancements are included in this z-stream release:

1.9.52.1.1. Adding "eu-es" region support for IPI

- Previously, the installation program failed to install a cluster on IBM Cloud VPC for the "eu-es" region, although it is supported. With this update, the installation program successfully installs a cluster on IBM Cloud VPC for the "eu-es" region. ([OCPBUGS-19398](#))

1.9.52.2. Updating

To update an existing OpenShift Container Platform 4.14 cluster to this latest release, see [Updating a cluster using the CLI](#).

1.9.53. RHSA-2024:0837 - OpenShift Container Platform 4.14.13 bug fix and security update

Issued: 21 February 2024

OpenShift Container Platform release 4.14.13, which includes security updates, is now available. The list of bug fixes that are included in the update is documented in the [RHSA-2024:0837](#) advisory. The RPM packages that are included in the update are provided by the [RHBA-2024:0840](#) 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 4.14.13 --pullspecs
```

1.9.53.1. Bug fixes

- Previously, Kubelet was running with an incorrect **unconfined_service_t** label, which caused an error related to SELinux. With this release, the issue is resolved and kubelet runs with the **kubelet_exec_t** label. ([OCBUGS-22270](#))

1.9.53.2. Updating

To update an existing OpenShift Container Platform 4.14 cluster to this latest release, see [Updating a cluster using the CLI](#).

1.9.54. RHSA-2024:0735 - OpenShift Container Platform 4.14.12 bug fix and security update

Issued: 13 February 2024

OpenShift Container Platform release 4.14.12, which includes security updates, is now available. The list of bug fixes that are included in the update is documented in the [RHSA-2024:0735](#) advisory. The RPM packages that are included in the update are provided by the [RHBA-2024:0738](#) 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 4.14.12 --pullspecs
```

1.9.54.1. Features

The following feature is included in this z-stream release:

1.9.54.1.1. Using dual Intel E810 Westport Channel NICs as grandmaster clock with the PTP Operator

- You can now configure **linuxptp** services as grandmaster clock (T-GM) for dual Intel E810 Westport Channel NICs by creating a **PtpConfig** custom resource (CR) that configures both NICs. The host system clock is synchronized from the NIC that is connected to the GNSS time source. The second NIC is synced to the 1PPS timing output provided by the NIC that is connected to GNSS. For more information see, [Configuring linuxptp services as a grandmaster clock for dual E810 Westport Channel NICs](#). ([RHBA-2024:0734](#))

1.9.54.2. Bug Fixes

- Previously, the release-to-channel strategy and **oc-mirror** behavior caused an error with the selective mirroring of packages. When selectively mirroring the most recent (and hence default) channel for the package and a new release introduced a new channel, the current default channel became invalid and the automatic assignment of the new default channel failed. With this release, the issue has been resolved. You can now define a **defaultChannel** field in the **ImageSetConfig** CR that overrides the **currentDefault** channel. ([OCBUGS-28871](#))
- Previously, the CPU limits from the EFS CSI driver container had the potential to cause performance degradation. With this release, the CPU limits from the EFS CSI driver container have been removed. ([OCBUGS-28823](#))

- Previously, when using the **routingViaHost** mode, access to the **ExternalTrafficPolicy=Local** load balancer services broke. With this release, the issue has been resolved. ([OCPBUGS-28789](#))
- Previously, when a HostedCluster was deployed and a user defined a KAS **AdvertiseAddress**, it conflicted with the current deployment, overlapping with the other networks like Service, Cluster or Machine network, which caused a deployment failure. With this release, network validations for **AdvertiseAddress** have been added. ([OCPBUGS-20547](#))

1.9.54.3. Updating

To update an existing OpenShift Container Platform 4.14 cluster to this latest release, see [Updating a cluster using the CLI](#).

1.9.55. RHSA-2024:0642 - OpenShift Container Platform 4.14.11 bug fix and security update

Issued: 7 February 2024

OpenShift Container Platform release 4.14.11, which includes security updates, is now available. The list of bug fixes that are included in the update is documented in the [RHSA-2024:0642](#) advisory. The RPM packages that are included in the update are provided by the [RHBA-2024:0645](#) 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 4.14.11 --pullspecs
```

1.9.55.1. Features

The following feature is included in this z-stream release:

1.9.55.1.1. Enabling configuration of whereabouts cron schedule

- The Whereabouts reconciliation schedule was hard-coded to run once per day and could not be reconfigured. With this release, a **ConfigMap** has enabled the configuration of the whereabouts cron schedule. For more information, see [Configuring the Whereabouts IP reconciler schedule](#).

1.9.55.2. Bug fixes

- Previously, updating OpenShift Container Platform could lead to DNS queries failing due to upstream returning a payload larger than 512 bytes for non-EDNS queries using CoreDNS 1.10.1. With this release, clusters with a non-compliant upstream will retry with TCP upon overflow errors which will prevent disruption of function when updating. ([OCPBUGS-28200](#))
- Previously, the **node.env** file would be overwritten on every restart due to a typo in an environment variable. With this release, edits to the **node.env** will persist after a restart. ([OCPBUGS-27362](#))
- Previously, **container_t** could not access Direct Rendering Infrastructure (DRI) devices. With this release, the policy has been updated so **container_t** can now access DRI devices and GPU devices exposed by a device plug-in by default. ([OCPBUGS-27275](#))
- Previously, pods assigned an IP address 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 plug-in issue associated with the IP allocation after a node force reboot is resolved.([OCPBUGS-26553](#))

1.9.55.3. Updating

To update an existing OpenShift Container Platform 4.14 cluster to this latest release, see [Updating a cluster using the CLI](#).

1.9.56. RHSA-2024:0290 - OpenShift Container Platform 4.14.10 bug fix and security update

Issued: 23 January 2024

OpenShift Container Platform release 4.14.10, which includes security updates, is now available. The list of bug fixes that are included in the update is documented in the [RHSA-2024:0290](#) advisory. The RPM packages that are included in the update are provided by the [RHSA-2024:0293](#) 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 4.14.10 --pullspecs
```

1.9.56.1. Bug fixes

- Previously, when the Cloud Credential Operator (CCO) was in default mode, CCO used an incorrect client for root credential queries. The CCO failed to find the intended secret and wrongly reported a **credsremoved** mode in the **cco_credentials_mode** metric. With this release, the CCO now uses the correct client so to ensure accurate reporting of the **cco_credentials_mode** metric. ([OCPBUGS-26512](#))

1.9.56.2. Updating

To update an existing OpenShift Container Platform 4.14 cluster to this latest release, see [Updating a cluster using the CLI](#).

1.9.57. RHSA-2024:0204 - OpenShift Container Platform 4.14.9 bug fix and security update

Issued: 17 January 2024

OpenShift Container Platform release 4.14.9, which includes security updates, is now available. The list of bug fixes that are included in the update is documented in the [RHSA-2024:0204](#) advisory. The RPM packages that are included in the update are provided by the [RHSA-2024:0207](#) 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 4.14.9 --pullspecs
```

1.9.57.1. Bug fixes

- Previously, The Cluster Version Operator (CVO) continually retrieves update recommendations and evaluates known conditional update risks against the current cluster state. CVO changes caused failing risk evaluations to block the CVO from fetching new update recommendations. This bug caused the CVO to fail to notice the update recommendation service serving an improved risk declaration. With this release, the CVO continues to poll the update recommendation service, regardless of whether update risks are being successfully evaluated or not. ([OCPBUGS-26207](#))
- Previously, The use of eus-* channels for mirroring releases was causing a failure in mirroring with the oc-mirror plugin. This was due to the fact that the oc-mirror plugin did not acknowledge that eus-* channels are even-numbered only. With this release users of the oc-mirror plugin should be able to use eus-* channels for mirroring releases. ([OCPBUGS-26065](#))

1.9.57.2. Updating

To update an existing OpenShift Container Platform 4.14 cluster to this latest release, see [Updating a cluster using the CLI](#).

1.9.58. RHSA-2024:0050 - OpenShift Container Platform 4.14.8 bug fix and security update

Issued: 9 January 2024

OpenShift Container Platform release 4.14.8, which includes security updates, is now available. The list of bug fixes that are included in the update is documented in the [RHSA-2024:0050](#) advisory. The RPM packages that are included in the update are provided by the [RHBA-2024:0053](#) 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 4.14.8 --pullspecs
```

1.9.58.1. Features

The following features are included in this z-stream release:

- With this release, Telemetry data is collected from clusters where pod security admission violations are occurring. The data being collected is whether the offending namespace is a Kubernetes system namespace, an OpenShift Container Platform system namespace, or a custom namespace. This data collection will help Red Hat evaluate customer cluster readiness for pod security admission global restricted enforcement in the future. For more information about pod security admission, see [Understanding and managing pod security admission](#) . ([OCPBUGS-25384](#))
- Previously, customers could not leverage OpenShift's Azure Identity capability for short-lived authentication tokens with layered products. With this release, OLM-managed operators have increased security by enabling this support. ([OCPBUGS-25275](#))

1.9.58.2. Bug fixes

- Previously, installation would fail when installing OpenShift Container Platform with the **bootMode** set to **UEFISecureBoot** on a node where Secure Boot is disabled. With this release, subsequent attempts to install OpenShift Container Platform with Secure Boot enabled

proceed normally. ([OCPBUGS-19884](#))

- Previously, the installer would fail to destroy a cluster when using regional PDs on the Google Cloud Platform. With this release, the replicated zones are found and the disks are deleted properly. ([OCPBUGS-22770](#))
- Previously, if the **additionalSecurityGroupIDs** field was not specified in the control plane nodes, the **additionalSecurityGroupIDs** in the **defaultMachinePlatform** stanza was not used. With this release, if the **additionalSecurityGroupIDs** field is not specified in the control plane nodes, the **additionalSecurityGroupIDs** in the **defaultMachinePlatform** stanza is now used. ([OCPBUGS-22771](#))

1.9.58.3. Updating

To update an existing OpenShift Container Platform 4.14 cluster to this latest release, see [Updating a cluster using the CLI](#).

1.9.59. RHSA-2023:7831 - OpenShift Container Platform 4.14.7 bug fix and security update

Issued: 3 January 2024

OpenShift Container Platform release 4.14.7, which includes security updates, is now available. The list of bug fixes that are included in the update is documented in the [RHSA-2023:7831](#) advisory. The RPM packages that are included in the update are provided by the [RHBA-2023:7834](#) 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 4.14.7 --pullspecs
```

1.9.59.1. Bug fixes

- Previously, when restarting IPsec pod, it killed existing policies. With this release, the IPsec service is also restarted, which reinstates existing policies and solves the issue. ([OCPBUGS-24633](#))
- Previously, updating a control plane machine set custom resource to reference an invalid resource, such as an invalid network name or image, created a control plane machine that would become stuck in the provisioning state and could not be deleted. With this release, this issue has been resolved. ([OCPBUGS-23202](#))
- Previously, applying a Performance Profile caused the tuned profile to report a **DEGRADED** condition. This was because the generated tuned profile was trying to set a second sysctl value. With this release, the sysctl value is no longer set by tuned, instead it is only set by the **sysctl.d** file. ([OCPBUGS-25305](#))

1.9.59.2. Updating

To update an existing OpenShift Container Platform 4.14 cluster to this latest release, see [Updating a cluster using the CLI](#).

1.9.60. RHSA-2023:7682 - OpenShift Container Platform 4.14.6 bug fix and security update

Issued: 12 December 2023

OpenShift Container Platform release 4.14.6, which includes security updates, is now available. The list of bug fixes that are included in the update is documented in the [RHSA-2023:7682](#) advisory. The RPM packages that are included in the update are provided by the [RHBA-2023:7685](#) 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 4.14.6 --pullspecs
```

1.9.60.1. Features

The following PTP features are included in this z-stream release:

1.9.60.1.1. Using hardware-specific NIC features with the PTP Operator

- A new PTP Operator hardware plugin is available that allows you to use hardware-specific features for supported NICs with the PTP Operator. Currently the Intel Westport channel E810 NIC is supported. For more information see, [E810 hardware configuration reference](#).

1.9.60.1.2. Using GNSS timing synchronization for PTP grandmaster clocks

- The PTP Operator now supports receiving precision PTP timing from Global Navigation Satellite System (GNSS) sources connected to grandmaster clocks (T-GM). For more information see, [Configuring linuxptp services as a grandmaster clock](#).

1.9.60.2. Bug fixes

- Previously, when you deployed IPv6-only hosts from a dual-stack cluster, an issue prevented the Baseboard Management Controller (BMC) from receiving the correct callback URL. Instead, the BMC received an IPv4 URL. With this update, the issue no longer occurs because the IP version of the URL depends on the IP version of the BMC address. ([OCPBUGS-23903](#))
- Previously, on a Single-node OpenShift that has guaranteed CPUs and where Interrupt Request (IRQ) load balancing is disabled, large latency spikes could occur at container startup. With this update, the issue no longer occurs. ([OCPBUGS-22901](#)) ([OCPBUGS-24281](#))

1.9.60.3. Known issues

- For PTP timing synchronization, DPLL phase offset monitoring is required to fully determine the grandmaster clock (T-GM) state. This is currently absent in the in-tree ice driver DPLL API, which creates a blind spot for determining the grandmaster state.

1.9.60.4. Updating

To update an existing OpenShift Container Platform 4.14 cluster to this latest release, see [Updating a cluster using the CLI](#).

1.9.61. RHSA-2023:7599 - OpenShift Container Platform 4.14.5 bug fix and security update

Issued: 5 December 2023

OpenShift Container Platform release 4.14.5, which includes security updates, is now available. The list of bug fixes that are included in the update is documented in the [RHSA-2023:7599](#) advisory. The RPM packages that are included in the update are provided by the [RHBA-2023:7603](#) 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 4.14.5 --pullspecs
```

1.9.61.1. Bug fixes

- Previously, when the build capability is not enabled, ConfigObserver controller would fail when trying to sync with build informer. With this release, ConfigObserver starts successfully when build capability is not enabled. ([OCPBUGS-23490](#)) ([OCPBUGS-21778](#))
- Previously, the Cloud Credential Operator (CCO) did not support updating the VMware vSphere root secret (**vsphere-creds**) in the **kube-system** namespace. This prevented the component secrets from synchronizing correctly. With this release, the CCO supports updating the vSphere root secret and resets the secret data when synchronized. ([OCPBUGS-23426](#))
- Previously, when deploying an application that has a large number of pods, some of which have CPU limits configured, the deployment can fail. With this update, the issue no longer occurs. ([RHEL-7232](#))

1.9.61.2. Updating

To update an existing OpenShift Container Platform 4.14 cluster to this latest release, see [Updating a cluster using the CLI](#).

1.9.62. RHSA-2023:7470 - OpenShift Container Platform 4.14.4 bug fix and security update

Issued: 29 November 2023

OpenShift Container Platform release 4.14.4, which includes security updates, is now available. The list of bug fixes that are included in the update is documented in the [RHSA-2023:7470](#) advisory. The RPM packages that are included in the update are provided by the [RHSA-2023:7473](#) 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 4.14.4 --pullspecs
```

1.9.62.1. Bug fixes

- Previously, when you specified Key Management Service (KMS) encryption keys in the **kmsKeyARN** section of the **install-config.yaml** configuration file for installing a cluster on

Amazon Web Services (AWS), permission roles were not added during the cluster installation operation. With this update, after you specify the keys in the configuration file, an additional set of keys are added to the cluster so that the cluster successfully installs. If you specify the **credentialsMode** parameter in the configuration file, all KMS encryption keys are ignored. ([OCPBUGS-22774](#))

1.9.62.2. Updating

To update an existing OpenShift Container Platform 4.14 cluster to this latest release, see [Updating a cluster using the CLI](#).

1.9.63. RHSA-2023:7315 - OpenShift Container Platform 4.14.3 bug fix and security update

Issued: 21 November 2023

OpenShift Container Platform release 4.14.3, which includes security updates, is now available. The list of bug fixes that are included in the update is documented in the [RHSA-2023:7315](#) advisory. The RPM packages that are included in the update are provided by the [RHBA-2023:7321](#) 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 4.14.3 --pullspecs
```

1.9.63.1. Bug fixes

- Previously, **ImageDigestMirrorSet** (IDMS) and **ImageTagMirrorSet** (ITMS) objects could not be used if there were any **ImageContentSourcePolicy** (ICSP) objects in the cluster. As a result, to use IDMS or ITMS objects, you needed to delete any ICSP objects in the cluster, which required a cluster reboot. With this release, ICSP, IDMS, and ITMS objects now function in the same cluster at the same time. As a result, you can use any or all of the three types of objects to configure repository mirroring after the cluster is installed. For more information, see [Understanding image registry repository mirroring](#). ([RHIBMCS-185](#))



IMPORTANT

Using an ICSP object to configure repository mirroring is a deprecated feature. Deprecated functionality is still included in OpenShift Container Platform and continues to be supported. However, it might be removed in a future release of this product. Because it is deprecated functionality, avoid using it for new deployments.

- Previously, **LoadBalancer** services were not created for a deployment if a node contained an additional port with the **enable_port_security** parameter set to **false**. With this release, **LoadBalancer** services are created for a deployment that contains additional ports with this setting. ([OCPBUGS-22974](#))
- Previously, a **ClusterAutoscaler** resource would go into a **CrashBackoff** loop for nodes configured with the Container Storage Interface (CSI) implementation. This release updated dependencies so that a **ClusterAutoscaler** resource no longer goes into a **CrashBackoff** for nodes configured in this way. ([OCPBUGS-23270](#))

1.9.63.2. Updating

To update an existing OpenShift Container Platform 4.14 cluster to this latest release, see [Updating a cluster using the CLI](#).

1.9.64. RHSA-2023:6837 - OpenShift Container Platform 4.14.2 bug fix and security update

Issued: 15 November 2023

OpenShift Container Platform release 4.14.2, which includes security updates, is now available. The list of bug fixes that are included in the update is documented in the [RHSA-2023:6837](#) advisory. The RPM packages that are included in the update are provided by the [RHSA-2023:6840](#) 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 4.14.2 --pullspecs
```

1.9.64.1. Bug fixes

- Previously, a change to the default security settings for new Microsoft Azure storage accounts in the **eastus** region prevented the installation of OpenShift Container Platform clusters that use Azure AD Workload Identity in that region. The issue has been resolved in this release. ([OCPBUGS-22651](#))
- Previously, the Docker build deployment would fail because the inline Dockerfile hooks did not preserve the modification time of the file that was copied. With this release, the '-p' flag was added to the **cp** command when copying artifacts between containers to preserve timestamps. ([OCPBUGS-23006](#))
- 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 many OpenShift Container Platform (OCP) clusters, this could lead to API limits being reached causing 429 errors when attempting to create new clusters. With this release, the time between calls is increased from 5 minutes to 20 minutes. ([OCPBUGS-22127](#))
- Previously, an Azure Managed Identity role was omitted from the cloud-controller-manager (CCM) service account, which meant that CCM could not properly manage **Service type LoadBalancers** in environments deployed to existing VNets with the private publishing method. With this release, the missing role was added to the CCO utility (**ccoctl**) so it is possible for Azure Managed Identity installations into existing VNet with private publishing. ([OCPBUGS-21926](#))
- This patch enables egress IP for Azure setups that use outbound rules to achieve outbound connectivity. An architectural constraint in Azure prevents the secondary IP acting as egress IP from having outbound connectivity in such setups. This means that matching pods will have no outbound connectivity to the internet, but will be able to reach external servers in the infrastructure network, which is the intended use case for egress IP. ([OCPBUGS-21785](#))
- Previously, when MetalLB Operator's controller restarted while having an IP assigned and unassigned load balancer services, it restarted with an empty internal state, which can break workloads. With this release, MetalLB's controller is modified to first process the services that already have an assigned IP. ([OCPBUGS-16267](#))

- Previously, if you created an Operator group with same name as a cluster role used by OpenShift Container Platform or Kubernetes resources, Operator Lifecycle Manager (OLM) would overwrite the cluster roles. With this fix, if you create an Operator group that conflicts with a cluster role used by OpenShift Container Platform or Kubernetes, OLM generates a unique cluster role name using the following syntax:

Naming syntax

```
olm.og.<operator_group_name>.<admin_edit_or_view>-<hash_value>
```

OLM generates unique names only for Operator groups that conflict with a set of defined cluster roles used by OpenShift Container Platform and Kubernetes. You must ensure that the name of your Operator group does not conflict with other cluster roles that exist on the cluster.

OLM generates unique names for Operator groups that conflict with the following cluster roles:

- **aggregate-olm**
- **alert-routing**
- **cluster**
- **cluster-monitoring**
- **monitoring**
- **monitoring-rules**
- **packagemanifests-v1**
- **registry**
- **storage**
([OCPBUGS-19789](#))

1.9.64.2. Known issue

- There is a known issue in this release which prevents installing a cluster on Alibaba Cloud using installer-provisioned infrastructure. Installing a cluster on Alibaba Cloud is a Technology Preview feature in this release. ([OCPBUGS-20552](#))

1.9.64.3. Updating

To update an existing OpenShift Container Platform 4.14 cluster to this latest release, see [Updating a cluster using the CLI](#).

1.9.65. RHBA-2023:6153 - OpenShift Container Platform 4.14.1 bug fix update

Issued: 1 November 2023

OpenShift Container Platform release 4.14.1 is now available. The list of bug fixes that are included in the update is documented in the [RHBA-2023:6153](#) advisory. The RPM packages that are included in the update are provided by the [RHBA-2023:6152](#) 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 4.14.1 --pullspecs
```

1.9.65.1. Updating

To update an existing OpenShift Container Platform 4.14 cluster to this latest release, see [Updating a cluster using the CLI](#).

1.9.66. RHSA-2023:5006 - OpenShift Container Platform 4.14.0 image release, bug fix, and security update advisory

Issued: 31 October 2023

OpenShift Container Platform release 4.14.0, which includes security updates, is now available. The list of bug fixes that are included in the update is documented in the [RHSA-2023:5006](#) advisory. The RPM packages that are included in the update are provided by the [RHSA-2023:5009](#) advisory. The list of security updates that are included in the update are documented in the [RHSA-2023:6143](#) 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 4.14.0 --pullspecs
```