



OpenShift Container Platform 3.11

Release Notes

Legal Notice

Copyright © 2019 Red Hat, Inc.

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

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

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

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

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

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

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

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

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

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

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

All other trademarks are the property of their respective owners.

Abstract

Table of Contents

CHAPTER 1. OVERVIEW	5
1.1. VERSIONING POLICY	5
CHAPTER 2. OPENSIFT CONTAINER PLATFORM 3.11 RELEASE NOTES	6
2.1. OVERVIEW	6
2.2. ABOUT THIS RELEASE	6
2.2.1. Major Changes Coming in Version 4.0	6
2.3. NEW FEATURES AND ENHANCEMENTS	7
2.3.1. Operators	8
2.3.1.1. Operator Lifecycle Manager (OLM) (Technology Preview)	8
2.3.1.2. Operator SDK	8
2.3.2. Brokers	8
2.3.2.1. OpenShift Container Platform Automation Broker Integration with Ansible Galaxy	8
2.3.2.2. Broker Support for Authenticated Registries	8
2.3.2.3. Service Catalog Namespaced Brokers	9
2.3.3. Installation and Upgrade	9
2.3.3.1. Checks for Expiring Certificates During Upgrade	9
2.3.3.2. Support for Ansible 2.6	9
2.3.3.3. Registry Auth Credentials Are Now Required	9
2.3.3.4. Customer Installations Are Now logged	10
2.3.4. Storage	10
2.3.4.1. OpenShift Container Storage	10
2.3.4.2. Container Storage Interface (Technology Preview)	10
2.3.4.3. Protection of Local Ephemeral Storage (Technology Preview)	10
2.3.4.4. Persistent Volume (PV) Provisioning Using OpenStack Manila (Technology Preview)	10
2.3.4.5. Persistent Volume (PV) Resize	10
2.3.4.6. Tenant-driven Storage Snapshotting (Technology Preview)	11
2.3.5. Scale	12
2.3.5.1. Cluster Limits	12
2.3.5.2. Scaling the Cluster Monitoring Operator	12
2.3.6. Metrics and Logging	12
2.3.6.1. Prometheus Cluster Monitoring	12
2.3.6.2. Elasticsearch 5 and Kibana 5	13
2.3.7. Developer Experience	13
2.3.7.1. CLI Plug-ins (Technology Preview)	13
2.3.7.2. Configure a Build Trigger Behavior without Triggering a Build Immediately	13
2.3.7.3. More Flexibility in Providing Configuration Options to Builds Using ConfigMaps	13
2.3.7.4. kubectl	13
2.3.8. Registry	13
2.3.8.1. Accessing and Configuring the Red Hat Registry	13
2.3.9. Quay	14
2.3.9.1. Red Hat Quay Registries	14
2.3.10. Networking	14
2.3.10.1. Improved OpenShift Container Platform and Red Hat OpenStack Integration with Kuryr	14
2.3.10.2. Router (HAProxy) Enhancements	14
2.3.10.3. HA Namespace-wide Egress IP	16
2.3.10.4. Fully-automatic Namespace-wide Egress IP	17
2.3.10.5. Configurable VXLAN Port	17
2.3.11. Master	17
2.3.11.1. Pod Priority and Preemption	17
2.3.11.2. The Descheduler (Technology Preview)	18

2.3.11.3. Podman (Technology Preview)	18
2.3.11.4. Node Problem Detector (Technology Preview)	18
2.3.11.5. Cluster Autoscaling (AWS Only)	19
2.3.12. Web Console	19
2.3.12.1. Cluster Administrator Console	19
2.3.12.2. Visibility into Nodes	20
2.3.12.3. Containers as a Service	20
2.3.12.4. Access Control Management	21
2.3.12.5. Cluster-wide Event Stream	21
2.3.13. Security	21
2.3.13.1. Control Sharing the PID Namespace Between Containers (Technology Preview)	21
2.3.13.2. GitHub Enterprise Added as Auth Provider	22
2.3.13.3. SSPI Connection Support on Microsoft Windows (Technology Preview)	22
2.3.14. Microservices	22
2.3.14.1. Red Hat OpenShift Service Mesh (Technology Preview)	22
2.4. NOTABLE TECHNICAL CHANGES	23
subjectaccessreviews.authorization.openshift.io and resourceaccessreviews.authorization.openshift.io Are Cluster-scoped Only	23
New SCC options	23
Removed oc deploy Command	23
Removed oc env and oc volume Commands	24
Removed the oc ex config patch Command	24
oc export Now Deprecated	24
oc types Now Deprecated	24
Pipeline Plug-in Is Deprecated	24
Logging: Elasticsearch 5	24
Hawkular Now Deprecated	24
New Registry Source for Red Hat images	24
New Storage Driver Recommendation	24
2.5. BUG FIXES	24
2.6. TECHNOLOGY PREVIEW FEATURES	28
2.7. KNOWN ISSUES	31
2.8. ASYNCHRONOUS ERRATA UPDATES	32
2.8.1. RHBA-2018:3537 - OpenShift Container Platform 3.11.43 Bug Fix and Enhancement Update	32
2.8.1.1. Bug Fixes	32
2.8.1.2. Upgrading	34
2.8.2. RHBA-2018:3743 - OpenShift Container Platform 3.11.51 Bug Fix and Enhancement Update	34
2.8.2.1. Upgrading	35
2.8.3. RHBA-2018:3688 - OpenShift Container Platform 3.11 Package Updates for IBM POWER	35
2.8.3.1. Upgrading	35
2.8.4. RHBA-2019:0024 - OpenShift Container Platform 3.11.59 Bug Fix and Enhancement Update	35
2.8.4.1. Bug Fixes	35
2.8.4.2. Enhancements	37
2.8.4.3. Upgrading	37
2.8.5. RHBA-2019:0096 - OpenShift Container Platform 3.11.69 Bug Fix and Enhancement Update	37
2.8.5.1. Bug Fixes	38
2.8.5.2. Enhancements	38
2.8.5.3. Upgrading	39
2.8.6. RHBA-2019:0326 - OpenShift Container Platform 3.11.82 Bug Fix Update	39
2.8.6.1. Bug Fixes	40
2.8.6.2. Upgrading	41
2.8.7. RHBA-2019:0407 - OpenShift Container Platform 3.11.88 Bug Fix Update	41
2.8.7.1. Upgrading	41

2.8.8. RHSA-2019:0739 - Important: OpenShift Container Platform 3.11 jenkins-2-plugins security update	41
2.8.8.1. Upgrading	41
2.8.9. RHBA-2019:0636 - OpenShift Container Platform 3.11.98 Bug Fix and Enhancement Update	41
2.8.9.1. Bug Fixes	42
2.8.9.2. Enhancements	44
2.8.9.3. Upgrading	44
2.8.10. RHBA-2019:0794 - OpenShift Container Platform 3.11.104 Bug Fix Update	44
2.8.10.1. Upgrading	44
2.8.11. RHBA-2019:1605 - OpenShift Container Platform 3.11.117 Bug Fix Update	44
2.8.11.1. Bug Fixes	44
2.8.11.2. Enhancements	47
2.8.11.3. Upgrading	47
2.8.12. RHSA-2019:1633 - Moderate: OpenShift Container Platform 3.11 atomic-openshift security update	47
2.8.12.1. Upgrading	47
2.8.13. RHBA-2019:1753 - OpenShift Container Platform 3.11.129 Bug Fix Update	47
2.8.13.1. Bug Fixes	48
2.8.13.2. Enhancements	49
2.8.13.3. Upgrading	49
2.8.14. RHBA-2019:2352 - OpenShift Container Platform 3.11.135 Bug Fix Update	49
2.8.14.1. Upgrading	49
2.8.15. RHBA-2019:2581 - OpenShift Container Platform 3.11.141 Bug Fix Update	49
2.8.15.1. Upgrading	49
2.8.16. RHBA-2019:2816 - OpenShift Container Platform 3.11.146 Bug Fix Update	50
2.8.16.1. Upgrading	50
CHAPTER 3. XPAAS RELEASE NOTES	51
CHAPTER 4. COMPARING WITH OPENSIFT ENTERPRISE 2	52
4.1. OVERVIEW	52
4.2. ARCHITECTURE CHANGES	52
4.3. APPLICATIONS	52
4.4. CARTRIDGES VERSUS IMAGES	53
4.5. BROKER VERSUS MASTER	54
4.6. DOMAIN VERSUS PROJECT	54

CHAPTER 1. OVERVIEW

The following release notes for OpenShift Container Platform 3.11 summarize all new features, major corrections from the previous version, and any known bugs upon general availability.

1.1. VERSIONING POLICY

OpenShift Container Platform provides strict backwards compatibility guarantees for all supported APIs, excluding alpha APIs (which may be changed without notice) and beta APIs (which may occasionally be changed in a non-backwards compatible manner).

The OpenShift Container Platform version must match between master and node hosts, excluding temporary mismatches during cluster upgrades. For example, in a 3.11 cluster, all masters must be 3.11 and all nodes must be 3.11. However, OpenShift Container Platform will continue to support older **oc** clients against newer servers. For example, a 3.5 **oc** will work against 3.4, 3.5, and 3.6 servers.

Changes of APIs for non-security related reasons will involve, at minimum, two minor releases (3.4 to 3.5 to 3.6, for example) to allow older **oc** to update. Using new capabilities may require newer **oc**. A 3.2 server may have additional capabilities that a 3.1 **oc** cannot use and a 3.2 **oc** may have additional capabilities that are not supported by a 3.1 server.

Table 1.1. Compatibility Matrix

	X.Y (oc Client)	X.Y+N ^[a] (oc Client)
X.Y (Server)	1	3
X.Y+N ^[a] (Server)	2	1
[a] Where N is a number greater than 1.		

- 1** Fully compatible.
- 2** **oc** client may not be able to access server features.
- 3** **oc** client may provide options and features that may not be compatible with the accessed server.

CHAPTER 2. OPENSIFT CONTAINER PLATFORM 3.11 RELEASE NOTES

2.1. OVERVIEW

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 and Kubernetes, OpenShift Container Platform provides a more secure and scalable multi-tenant 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.

2.2. ABOUT THIS RELEASE

Red Hat OpenShift Container Platform 3.11 ([RHBA-2018:2652](#)) is now available. This release is based on [OKD 3.11](#), and it uses Kubernetes 1.11. New features, changes, bug fixes, and known issues that pertain to OpenShift Container Platform 3.11 are included in this topic.

OpenShift Container Platform 3.11 is supported on Red Hat Enterprise Linux 7.4 and later with the latest packages from Extras, including CRI-O 1.11 and Docker 1.13. It is also supported on Atomic Host 7.5 and later.

OpenShift Container Platform 3.11 is supported on Red Hat Enterprise Linux 7 nodes running in Federal Information Processing Standards (FIPS) mode.

For initial installations, see the [Installing Clusters](#) documentation.

To upgrade to this release from a previous version, see the [Upgrading Clusters](#) documentation.



WARNING

In the initial release of OpenShift Container Platform version 3.11, downgrading does not completely restore your cluster to version 3.10. Do not downgrade.

If you need to downgrade, contact Red Hat support so they can help you determine the best course of action.

2.2.1. Major Changes Coming in Version 4.0

OpenShift Container Platform 3.11 is the last release in the 3.x stream. Large changes to the underlying architecture and installation process are coming in version 4.0, and many features will be deprecated.

Table 2.1. Features Deprecated in Version 4.0

Feature	Justification
Hawkular	Replaced by Prometheus monitoring.
Cassandra	Replaced by Prometheus monitoring.
Heapster	Replaced by Metrics-Server or Prometheus metrics adapter.
Atomic Host	Replaced by Red Hat CoreOS.
System containers	Replaced by Red Hat CoreOS.
projectatomic/docker-1.13 additional search registries	CRI-O is the default container runtime for 4.x on RHCOS and Red Hat Enterprise Linux.
oc adm diagnostics	Operator-based diagnostics.
oc adm registry	Replaced by the registry operator.
Custom Docker Build Strategy on Builder Pods	If you want to continue using custom builds, you must replace your Docker invocations with Podman and Buildah. The custom build strategy will not be removed, but the functionality will change significantly in OpenShift Container Platform 4.0.
Cockpit	Replaced by Quay.
Standalone Registry Installations	Quay is our enterprise container image registry.
DNSmasq	CoreDNS will be the default.
External etcd nodes	For 4.0, etcd is on the cluster always.
CloudForms OpenShift Provider and Podified CloudForms	Replaced by built-in management tooling.
Volume Provisioning via installer	Replaced by dynamic volumes or, if NFS is required, NFS provisioner.
blue-green-installation method	Ease of upgrade is a core value of 4.0.

Because of the extent of the changes in OpenShift Container Platform 4.0, the product documentation will also undergo significant changes, including the deprecation of large amounts of content. New content will be released based on the architectural changes and updated use cases.

2.3. NEW FEATURES AND ENHANCEMENTS

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

2.3.1. Operators

2.3.1.1. Operator Lifecycle Manager (OLM) (Technology Preview)

This feature is currently in [Technology Preview](#) and not for production workloads.

The OLM aids cluster administrators in installing, upgrading, and granting access to Operators running on their cluster:

- Includes a catalog of curated Operators, with the ability to load other Operators into the cluster
- Handles rolling updates of all Operators to new versions
- Supports role-based access control (RBAC) for certain teams to use certain Operators

See [Installing the Operator Framework](#) for more information.

2.3.1.2. Operator SDK

The Operator SDK is a development tool to jump-start building an Operator with generated code and a CLI to aid in building, testing, and publishing your Operator. The Operator SDK:

- Provides tools to get started quickly embedding application business logic into an Operator
- Saves you from doing the work to set up scaffolding to communicate with the Kubernetes API
- Helps run end-to-end tests of your logic on a local or remote cluster
- Is used by Couchbase, MongoDB, Redis and more

See [Getting started with the Operator SDK](#) in OKD documentation for more information and walkthroughs.

2.3.2. Brokers

Brokers mediate service requests in the Service Catalog. The goal is for you to initiate the request and for the system to fulfill the request in an automated fashion.

2.3.2.1. OpenShift Container Platform Automation Broker Integration with Ansible Galaxy

The Automation Broker manages applications defined in Ansible Playbook Bundles (APB). OpenShift Container Platform 3.11 includes support for discovering and running APB sources published to Ansible Galaxy from the OpenShift Container Platform Automation Broker.

See [OpenShift Automation Broker](#) for more information.

2.3.2.2. Broker Support for Authenticated Registries

The Red Hat Container Catalog is moving from **registry.access.redhat.com** to **registry.redhat.io**. **registry.redhat.io** requires authentication for access to images and hosted content on OpenShift Container Platform.

OpenShift Container Platform 3.11 adds support for authenticated registries. The broker uses **cluster-wide** as the default setting for registry authentication credentials. You can define **oreg_auth_user** and **oreg_auth_password** in the inventory file to configure the credentials.

2.3.2.3. Service Catalog Namespaced Brokers

The Service Catalog added support for namespaced brokers in addition to the previous cluster-scoped behavior. This means you can register the broker with the service catalog as either a cluster-scoped **ClusterServiceBroker** or a namespace-scoped **ServiceBroker** kind. Depending on the broker's scope, its services and plans are available to the entire cluster or scoped to a specific namespace. When installing the broker, you can set the **kind** argument as **ServiceBroker** (namespace-specific) or **ClusterServiceBroker** (cluster-wide).

2.3.3. Installation and Upgrade

2.3.3.1. Checks for Expiring Certificates During Upgrade

In OpenShift Container Platform 3.11, **openshift_certificate_expiry_warning_days**, which indicates the amount of time the auto-generated certificates must be valid for an upgrade to proceed, is added.

Additionally, **openshift_certificate_expiry_fail_on_warn** is added, which determines whether the upgrade fails if the auto-generated certificates are not valid for the period specified by the **openshift_certificate_expiry_warning_days** parameter.

See [Configuring Your Inventory File](#) for more information.

2.3.3.2. Support for Ansible 2.6

openshift-ansible now requires Ansible 2.6 for both installation of OpenShift Container Platform 3.11 and upgrading from version 3.10.

The minimum version of Ansible required for OpenShift Container Platform 3.11 to run playbooks is now 2.6.x. On both master and node, use **subscription-manager** to enable the repositories that are necessary to install OpenShift Container Platform using Ansible 2.6. For example:

```
$ subscription-manager repos --enable="rhel-7-server-rpms" \
  --enable="rhel-7-server-extras-rpms" \
  --enable="rhel-7-server-ose-3.11-rpms" \
  --enable="rhel-7-server-ansible-2.6-rpms"
```

Ansible 2.7 is not yet supported.

2.3.3.3. Registry Auth Credentials Are Now Required

Registry auth credentials are now required for OpenShift Container Platform so that images and metadata can be pulled from an authenticated registry, registry.redhat.io.

Registry auth credentials are required prior to installing and upgrading when:

- **openshift_deployment_type** == 'openshift-enterprise'
- **oreg_url** == 'registry.redhat.io' or undefined

To configure authentication, **oreg_auth_user** and **oreg_auth_password** must be defined in the inventory file.

Pods can also be allowed to reference images from other secure registries.

See [Importing Images from Private Registries](#) for more information.

2.3.3.4. Customer Installations Are Now logged

Ansible configuration is now updated to ensure OpenShift Container Platform installations are logged by default.

The Ansible configuration parameter **log_path** is now defined. Users must be in the */usr/share/ansible/openshift-ansible* directory prior to running any playbooks.

2.3.4. Storage

2.3.4.1. OpenShift Container Storage

[OpenShift Container Storage](#) (OCS) provides software defined storage as a container for use with OpenShift Container Platform. Use OCS to define [persistent volumes](#) (PV) for use with your containers. ([BZ#1645358](#))

2.3.4.2. Container Storage Interface (Technology Preview)

This feature is currently in [Technology Preview](#) and not for production workloads.

CSI allows OpenShift Container Platform to consume storage from storage backends that implement the [CSI interface](#) as [persistent storage](#).

See [Persistent Storage Using Container Storage Interface \(CSI\)](#) for more information.

2.3.4.3. Protection of Local Ephemeral Storage (Technology Preview)

This feature is currently in [Technology Preview](#) and not for production workloads.

You can now control the use of the local ephemeral storage feature on your nodes. This helps prevent users from exhausting node local storage with their pods and other pods that happen to be on the same node.

This feature is disabled by default. If enabled, the OpenShift Container Platform cluster uses ephemeral storage to store information that does not need to persist after the cluster is destroyed.

See [Configuring Ephemeral Storage](#) for more information.

2.3.4.4. Persistent Volume (PV) Provisioning Using OpenStack Manila (Technology Preview)

This feature is currently in [Technology Preview](#) and not for production workloads.

OpenShift Container Platform is capable of provisioning PVs using the [OpenStack Manila](#) shared file system service.

See [Persistent Storage Using OpenStack Manila](#) for more information.

2.3.4.5. Persistent Volume (PV) Resize

You can expand PV claims online from OpenShift Container Platform for GlusterFS by creating a storage class with **allowVolumeExpansion** set to **true**, which causes the following to happen:

1. The PVC uses the storage class and submits a claim.
2. The PVC specifies a new increased size.
3. The underlying PV is resized.

Block storage volume types such as GCE-PD, AWS-EBS, Azure Disk, Cinder, and Ceph RBD typically require a file system expansion before the additional space of an expanded volume is usable by pods. Kubernetes takes care of this automatically whenever the pod or pods referencing your volume are restarted.

Network attached file systems, such as GlusterFS and Azure File, can be expanded without having to restart the referencing pod, as these systems do not require unique file system expansion.

See [Expanding Persistent Volumes](#) for more information.

2.3.4.6. Tenant-driven Storage Snapshotting (Technology Preview)

This feature is currently in [Technology Preview](#) and not for production workloads.

Tenants can now leverage the underlying storage technology backing the PV assigned to them to make a snapshot of their application data. Tenants can also now restore a given snapshot from the past to their current application.

You can use an external provisioner to access EBS, GCE pDisk, and hostPath. This Technology Preview feature has tested EBS and hostPath. The tenant must stop the pods and start them manually.

To use the external provisioner to access EBS and hostPath:

1. The administrator runs an external provisioner for the cluster. These are images from the Red Hat Container Catalog.
2. The tenant creates a PV claim and owns a PV from one of the supported storage solutions.
3. The administrator must create a new **StorageClass** in the cluster, for example:

```
kind: StorageClass
apiVersion: storage.k8s.io/v1
metadata:
  name: snapshot-promoter
provisioner: volumesnapshot.external-storage.k8s.io/snapshot-promoter
```

4. The tenant creates a snapshot of a PV claim named **gce-pvc**, and the resulting snapshot is **snapshot-demo**, for example:

```
$ oc create -f snapshot.yaml

apiVersion: volumesnapshot.external-storage.k8s.io/v1
kind: VolumeSnapshot
metadata:
  name: snapshot-demo
```

```
namespace: myns
spec:
  persistentVolumeClaimName: gce-pvc
```

5. The pod is restored to that snapshot, for example:

```
$ oc create -f restore.yaml
apiVersion: v1
kind: PersistentVolumeClaim
metadata:
  name: snapshot-pv-provisioning-demo
  annotations:
    snapshot.alpha.kubernetes.io/snapshot: snapshot-demo
spec:
  storageClassName: snapshot-promoter
```

2.3.5. Scale

2.3.5.1. Cluster Limits

Updated guidance around [Cluster Limits](#) for OpenShift Container Platform 3.11 is now available.

New recommended guidance for master

For large or dense clusters, the API server might get overloaded because of the default queries per second (QPS) limits. Edit `/etc/origin/master/master-config.yaml` and double or quadruple the QPS limits.

See [Recommended Practices for OpenShift Container Platform Master Hosts](#) for more information.

2.3.5.2. Scaling the Cluster Monitoring Operator

OpenShift Container Platform exposes metrics that can be collected and stored in backends by the [cluster-monitoring-operator](#). As an OpenShift Container Platform administrator, you can view system resources, containers, and component's metrics in one dashboard interface, Grafana.

In OpenShift Container Platform 3.11, the cluster monitoring operator installation is enabled by default as `node-role.kubernetes.io/infra=true` in your cluster. You can update this by setting `openshift_cluster_monitoring_operator_node_selector` in the inventory file of your customized node selector. Ensure there is an available node in your cluster to avoid unexpected failures.

See [Scaling Cluster Monitoring Operator](#) for capacity planning details.

2.3.6. Metrics and Logging

2.3.6.1. Prometheus Cluster Monitoring

Prometheus cluster monitoring is now fully supported in OpenShift Container Platform and deployed by default into an OpenShift Container Platform cluster.

- Query and plot cluster metrics collected by Prometheus.
- Receive notifications from pre-packaged alerts, enabling owners to take corrective actions and start troubleshooting problems.

- View pre-packaged Grafana dashboards for etcd, cluster state, and many other aspects of cluster health.

See [Configuring Prometheus Cluster Monitoring](#) for more information.

2.3.6.2. Elasticsearch 5 and Kibana 5

Elasticsearch 5 and Kibana 5 are now available. Kibana dashboards can be saved and shared between users. Elasticsearch 5 introduces better resource usage and performance and better resiliency.

Additionally, new numeric types, **half_float** and **scaled_float** are now added. There are now instant aggregations in Kibana 5, making it faster. There is also a new API that returns an explanation of why Elasticsearch shards are unassigned.

2.3.7. Developer Experience

2.3.7.1. CLI Plug-ins (Technology Preview)

This feature is currently in [Technology Preview](#) and not for production workloads.

Usually called *plug-ins* or *binary extensions*, this feature allows you to extend the default set of **oc** commands available and, therefore, allows you to perform new tasks.

See [Extending the CLI](#) for information on how to install and write extensions for the CLI.

2.3.7.2. Configure a Build Trigger Behavior without Triggering a Build Immediately

You can pause an image change trigger to allow multiple changes on the referenced image stream before a build is started. You can also set the **paused** attribute to **true** when initially adding an **ImageChangeTrigger** to a **BuildConfig** to prevent a build from being immediately triggered.

See [Triggering Builds](#) for more information.

2.3.7.3. More Flexibility in Providing Configuration Options to Builds Using ConfigMaps

In some scenarios, build operations require credentials or other configuration data to access dependent resources, but it is undesirable for that information to be placed in source control. You can define *input secrets* and *input ConfigMaps* for this purpose.

See [Build Inputs](#) for additional details.

2.3.7.4. kubectI

OpenShift Container Platform always shipped [kubectI](#) for Linux on the master's file system, but it is now available in the [oc client downloads](#).

2.3.8. Registry

2.3.8.1. Accessing and Configuring the Red Hat Registry

All container images available through the Red Hat Container Catalog are hosted on an image registry, **registry.access.redhat.com**. The Red Hat Container Catalog is moving from **registry.access.redhat.com** to **registry.redhat.io**. The new registry, **registry.redhat.io**, requires

authentication for access to images and hosted content on OpenShift Container Platform. Following the move to the new registry, the existing registry will be available for a period of time.

See [Authentication Enabled Red Hat Registry](#) for more information.

2.3.9. Quay

2.3.9.1. Red Hat Quay Registries

If you need an enterprise quality container image registry, Red Hat Quay is available both as a hosted service and as software you can install in your own data center or cloud environment. Advanced registry features in Red Hat Quay include geo-replication, image scanning, and the ability to roll back images. Visit the [Quay.io](#) site to set up your own hosted Quay registry account.

See [Container Registry](#) for more information.

2.3.10. Networking

2.3.10.1. Improved OpenShift Container Platform and Red Hat OpenStack Integration with Kuryr

See [Kuryr SDN Administration](#) and [Configuring Kuryr SDN](#) for best practices in OpenShift Container Platform and Red Hat OpenStack integration.

2.3.10.2. Router (HAProxy) Enhancements

The OpenShift Container Platform router is the most common way to get traffic into the cluster. The table below lists the OpenShift Container Platform router (HAProxy) enhancements for 3.11.

Table 2.2. Router (HAProxy) enhancements

Feature	Feature enhancements	Command syntax
HTTP/2	Implements HAProxy router HTTP/2 support (terminating at the router).	\$ oc set env dc/router ROUTER_ENABLE_HTTP2=true

Feature	Feature enhancements	Command syntax
Performance	Increases the number of threads that can be used by HAProxy to serve more routes.	<p>1. Scale down the default router and create a new router using two threads:</p> <pre data-bbox="1139 360 1439 667">\$ oc scale dc/router --replicas=0 \$ oc adm router myrouter --threads=2 --images='openshift3/ose-haproxy-router:v3.x'</pre> <p>2. Set a new thread count (for, example 7) for the HAProxy router:</p> <pre data-bbox="1139 824 1417 987">\$ oc set env dc/myrouter ROUTER_THREADS=7</pre>
Dynamic changes	Implements changes to the HAProxy router without requiring a full router reload.	<pre data-bbox="1034 1106 1422 1205">\$ oc set env dc/router ROUTER_HAPROXY_CONFIG_MANAGER=true</pre>
Client SSL/TLS cert validation	Enables mTLS for route support of older clients/services that do not support SNI, but where certificate verification is a requirement.	<pre data-bbox="1034 1263 1410 1435">\$ oc adm router myrouter --mutual-tls-auth=optional --mutual-tls-auth-ca=/root/ca.pem --images="\$image"</pre>

Feature	Feature enhancements	Command syntax
Logs captured by aggregated logging/EFK	Collects access logs so that Operators can see them.	<ol style="list-style-type: none"> Create a router with an rsyslog container: <pre>\$ oc adm router myrouter -- extended-logging -- images='xxxx'</pre> Set the log level: <pre>\$ oc set env dc/myrouter ROUTER_LOG_LEVEL=debug</pre> Check the access logs in the rsyslog container: <pre>\$ oc logs -f myrouter-x-xxxxx -c syslog</pre>

2.3.10.3. HA Namespace-wide Egress IP

Adding basic active/backup HA for project/namespace egress IPs now allows a namespace to have multiple egress IPs hosted on different cluster nodes.

To add basic active/backup HA to an existing project/namespace:

- Add two or more egress IPs to its **netnamespace**:

```
$ oc patch netnamespace myproject -p '{"egressIPs":["10.0.0.1","10.0.0.2"]}'
```

- Add the first egress IP to a node in the cluster:

```
# oc patch hostsubnet node1 -p '{"egressIPs":["10.0.0.1"]}'
```

- Add the second egress IP to a different node in the cluster:

```
# oc patch hostsubnet node2 -p '{"egressIPs":["10.0.0.2"]}'
```

The project/namespace uses the first listed egress IP by default (if available) until that node stops responding, upon which other nodes switch to using the next listed egress IP, and so on. This solution requires greater than or equal to two IPs.

If the original IP eventually comes back, the nodes switch back to using the original egress IP.

See [Enabling Static IPs for External Project Traffic](#) for more information.

2.3.10.4. Fully-automatic Namespace-wide Egress IP

A fully-automatic HA option is now available. Projects/namespaces are automatically allocated a single egress IP on a node in the cluster, and that IP is automatically migrated from a failed node to a healthy node.

To enable the fully-automatic HA option:

1. Patch one of the cluster nodes with the **egressCIDRs**:

```
# oc patch hostsubnet node1 -p '{"egressCIDRs":["10.0.0.0/24"]}'
```

2. Create a project/namespace and add a single egress IP to its **netnamespace**:

```
# oc patch netnamespace myproject -p '{"egressIPs":["10.0.0.1"]}'
```

2.3.10.5. Configurable VXLAN Port

The OpenShift Container Platform SDN overlay VXLAN port is now configurable (default is **4789**). VMware modified the VXLAN port used in the VMware NSX SDN ($\geq v6.2.3$) from **8472** to **4789** to adhere to [RFC 7348](#).

When running the OpenShift Container Platform SDN overlay on top of VMware's NSX SDN underlay, there is a port conflict since both use the same VXLAN port (**4789**). With a configurable VXLAN port, users can choose the port configuration of the two products, used in combination, for their particular environment.

To configure the VXLAN port:

1. Modify the VXLAN port in *master-config.yaml* with the new port number (for example, **4889** instead of **4789**):

```
vxlanPort: 4889
```

2. Delete **clusternetwork** and restart the master API and controller:

```
$ oc delete clusternetwork default
$ master-restart api controllers
```

3. Restart all SDN pods in the **openshift-sdn** project:

```
$ oc delete pod -n openshift-sdn -l app=sdn
```

4. Allow the new port on the firewall on all nodes:

```
# iptables -i OS_FIREWALL_ALLOW -p udp -m state --state NEW -m udp --dport 4889 -j ACCEPT
```

2.3.11. Master

2.3.11.1. Pod Priority and Preemption

You can enable pod priority and preemption in your cluster. Pod priority indicates the importance of a

pod relative to other pods and queues the pods based on that priority. Pod preemption allows the cluster to evict, or preempt, lower-priority pods so that higher-priority pods can be scheduled if there is no available space on a suitable node. Pod priority also affects the scheduling order of pods and out-of-resource eviction ordering on the node.

See [Pod Priority and Preemption](#) for more information.

2.3.11.2. The Descheduler (Technology Preview)

This feature is currently in [Technology Preview](#) and not for production workloads.

The descheduler moves pods from less desirable nodes to new nodes. Pods can be moved for various reasons, such as:

- Some nodes are under- or over-utilized.
- The original scheduling decision does not hold true any more, as taints or labels are added to or removed from nodes, pod/node affinity requirements are not satisfied any more.
- Some nodes failed and their pods moved to other nodes.
- New nodes are added to clusters.

See [Descheduling](#) for more information.

2.3.11.3. Podman (Technology Preview)

This feature is currently in [Technology Preview](#) and not for production workloads.

Podman is a daemon-less CLI/API for running, managing, and debugging OCI containers and pods. It:

- Is fast and lightweight.
- Leverages runC.
- Provides a syntax for working with containers.
- Has remote management API via Varlink.
- Provides systemd integration and advanced namespace isolation.

For more information, see [Crioctl Vs Podman](#).

2.3.11.4. Node Problem Detector (Technology Preview)

This feature is currently in [Technology Preview](#) and not for production workloads.

The Node Problem Detector monitors the health of your nodes by finding specific problems and reporting these problems to the API server, where external controllers could take action. The Node Problem Detector is a daemon that runs on each node as a DaemonSet. The daemon tries to make the cluster aware of node level faults that should make the node not schedulable. When you start the Node Problem Detector, you tell it a port over which it should broadcast the issues it finds. The detector allows you to load sub-daemons to do the data collection. There are three as of today. Issues found by the problem daemon can be classified as **NodeCondition**.

The three problem daemons are:

- Kernel Monitor, which monitors the kernel log via journald and reports problems according to regex patterns.
- AbrtAdaptor, which monitors the node for kernel problems and application crashes from journald.
- CustomerPluginMonitor, which allows you to test for any condition and exit on a **0** or **1** should your condition not be met.

See [Node Problem Detector](#) for more information.

2.3.11.5. Cluster Autoscaling (AWS Only)

You can configure an auto-scaler on your OpenShift Container Platform cluster in Amazon Web Services (AWS) to provide elasticity for your application workload. The auto-scaler ensures that enough nodes are active to run your pods and that the number of active nodes is proportional to current demand.

See [Configuring the cluster auto-scaler in AWS](#) for more information.

2.3.12. Web Console

2.3.12.1. Cluster Administrator Console

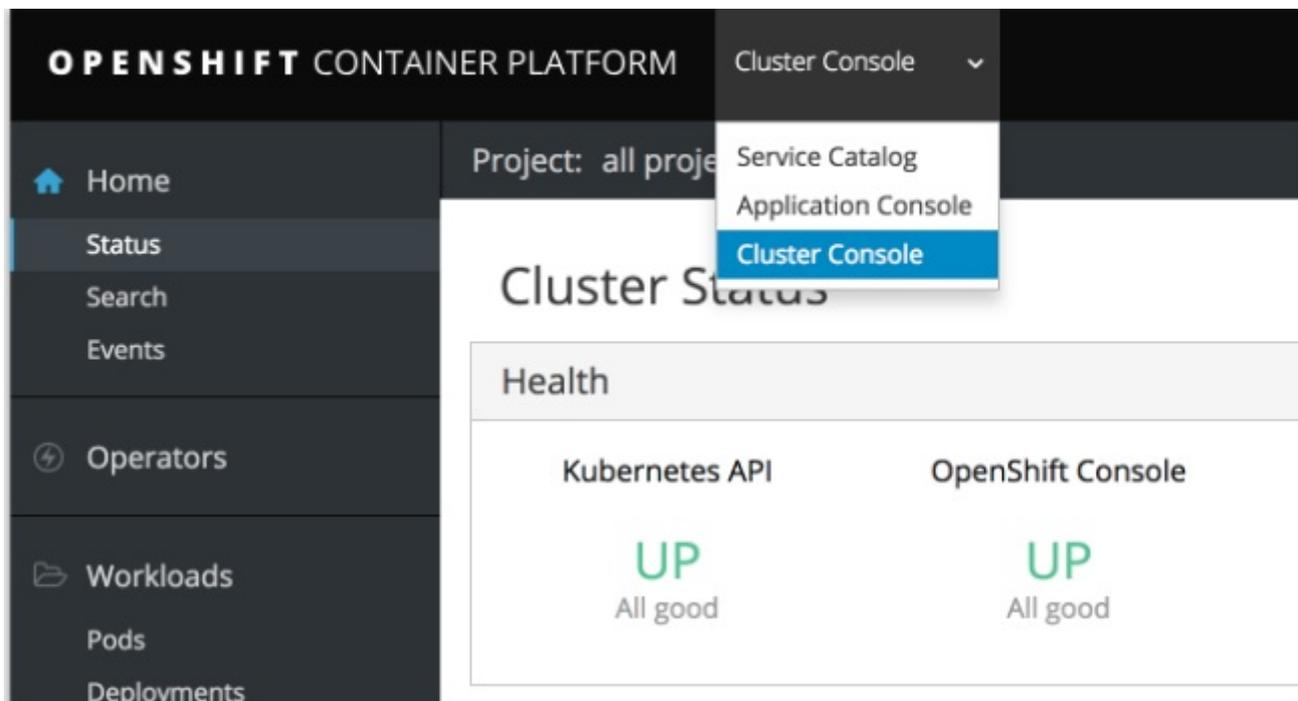
OpenShift Container Platform 3.11 introduces a cluster administrator console tailored toward application development and cluster administrator personas.

Users have a choice of experience based on their role or technical abilities, including:

- An administrator with Containers as a Service (CaaS) experience and with heavy exposure to Kubernetes.
- An application developer with Platform as a Service (PaaS) experience and standard OpenShift Container Platform UX.

Sessions are not shared across the consoles, but credentials are.

See [Configuring Your Inventory File](#) for details on configuring the cluster console.



2.3.12.2. Visibility into Nodes

OpenShift Container Platform now has an expanded ability to manage and troubleshoot cluster nodes, for example:

- Node status events are extremely helpful in diagnosing resource pressure and other failures.
- Runs **node-exporter** as a DaemonSet on all nodes, with a default set of scraped metrics from the **kube-state-metrics** project.
- Metrics are protected by RBAC.
- Those with **cluster-reader** access and above can view metrics.

2.3.12.3. Containers as a Service

You can view, edit, and delete the following Kubernetes objects:

- Networking
 - Routes and ingress
- Storage
 - PVs and PV claims
 - Storage classes
- Admin
 - Projects and namespaces
 - Nodes
 - Roles and RoleBindings

- CustomResourceDefinition (CRD)

2.3.12.4. Access Control Management

OpenShift Container Platform 3.11 includes visual management of the cluster's RBAC roles and RoleBindings, which allows you to:

- Find users and service accounts with a specific role.
- View cluster-wide or namespaced bindings.
- Visually audit a role's verbs and objects.

Project administrators can self-manage roles and bindings scoped to their namespace.

2.3.12.5. Cluster-wide Event Stream

The cluster-wide event stream provides the following ways to help debug events:

- All namespaces are accessible by anyone who can list the namespaces and events.
- Per-namespace is accessible for all project viewers.
- There is an option to filter by category and object type.

The screenshot shows the OpenShift Container Platform Cluster Console interface. The top navigation bar includes the OpenShift logo, the text 'CONTAINER PLATFORM', and 'Cluster Console'. A sidebar on the left contains navigation links for Home, Status, Catalog, Search, Events (selected), Operators, Workloads, Networking, Storage, Builds, Monitoring, and Administration. The main content area is titled 'Events' and shows a list of events for the 'all projects' namespace. The events are filtered by 'All Types' and 'All Categories'. The events list includes details such as the pod name, the source (kubelet), the time of occurrence, and the frequency of the event. The third event, 'reporting-operator-8556d966cf-mx57h', is highlighted with a red border and a red circle, indicating a readiness probe failure.

Pod Name	Source	Time	Frequency	Message
kafka-producer-6cb449bfc5-chw48	Generated from kubelet on ip-172-31-19-72.us-west-1.compute.internal	a few seconds ago	127011 times in the last 23 days	pulling image "scholzj/kafka-producer:latest"
kafka-consumer-7d87745d6b-x475q	Generated from kubelet on ip-172-31-19-72.us-west-1.compute.internal	a few seconds ago	126926 times in the last 23 days	Successfully pulled image "scholzj/kafka-consumer:latest"
reporting-operator-8556d966cf-mx57h	Generated from kubelet on ip-172-31-27-92.us-west-1.compute.internal	less than a minute ago	7985 times in the last 3 days	Readiness probe failed: Get https://10.131.1.240:8080/ready: dial tcp 10.131.1.240:8080: connect: connecti
my-cluster-zookeeper-0	Generated from kubelet on ip-172-31-22-146.us-west-1.compute.internal	a minute ago	2402 times in the last 3 days	Killing container with id docker://zookeeper:Need to kill Pod

2.3.13. Security

2.3.13.1. Control Sharing the PID Namespace Between Containers (Technology Preview)

This feature is currently in [Technology Preview](#) and not for production workloads.

You can use this feature to configure cooperating containers in a pod, such as a log handler sidecar container, or to troubleshoot container images that do not include debugging utilities like a shell, for example:

- The feature gate **PodShareProcessNamespace** is set to **false** by default.
- Set **feature-gates=PodShareProcessNamespace=true** in the API server, controllers, and kubelet.
- Restart the API server, controller, and node service.
- Create a pod with the specification of **shareProcessNamespace: true**.
- Run **oc create -f <pod spec file>**.

Caveats

When the PID namespace is shared between containers:

- Sidecar containers are not isolated.
- Environment variables are visible to all other processes.
- Any **kill all** semantics used within the process are broken.
- Any **exec** processes from other containers show up.

See [Expanding Persistent Volumes](#) for more information.

2.3.13.2. GitHub Enterprise Added as Auth Provider

GitHub Enterprise is now an auth provider. OAuth facilitates a token exchange flow between OpenShift Container Platform and GitHub or GitHub Enterprise. You can use the GitHub integration to connect to either GitHub or GitHub Enterprise. For GitHub Enterprise integrations, you must provide the **hostname** of your instance and can optionally provide a **ca** certificate bundle to use in requests to the server.

See [Configuring Authentication and User Agent](#) for more information.

2.3.13.3. SSPI Connection Support on Microsoft Windows (Technology Preview)

This feature is currently in [Technology Preview](#) and not for production workloads.

oc now supports the Security Support Provider Interface (SSPI) to allow for single sign-on (SSO) flows on Windows. If you use the request header identity provider with a GSSAPI-enabled proxy to connect an Active Directory server to OpenShift Container Platform, users can automatically authenticate to OpenShift Container Platform using the **oc** command line interface from a domain-joined Windows computer.

See [Configuring Authentication and User Agent](#) for more information.

2.3.14. Microservices

2.3.14.1. Red Hat OpenShift Service Mesh (Technology Preview)

This feature is currently in [Technology Preview](#) and not for production workloads.

Red Hat OpenShift Service Mesh is a platform that provides behavioral insights and operational control over the service mesh, providing a uniform way to connect, secure, and monitor microservice applications.

The term service mesh is often used to describe the network of microservices that make up applications based on a distributed microservice architecture and the interactions between those microservices. As a service mesh grows in size and complexity, it can become harder to understand and manage.

Based on the open source [Istio](#) project, Red Hat OpenShift Service Mesh layers transparently onto existing distributed applications, without requiring any changes in the service code.

See [Installing Red Hat OpenShift Service Mesh](#) for more information.

2.4. NOTABLE TECHNICAL CHANGES

OpenShift Container Platform 3.11 introduces the following notable technical changes.

subjectaccessreviews.authorization.openshift.io and resourceaccessreviews.authorization.openshift.io Are Cluster-scoped Only

subjectaccessreviews.authorization.openshift.io and *resourceaccessreviews.authorization.openshift.io* are now cluster-scoped only. If you need namespace-scoped requests, use *localsubjectaccessreviews.authorization.openshift.io* and *localresourceaccessreviews.authorization.openshift.io*.

New SCC options **No new privs flag**

Security Context Constraints have two new options to manage use of the (Docker) **no_new_privs** flag to prevent containers from gaining new privileges:

- The **AllowPrivilegeEscalation** flag gates whether or not a user is allowed to set the security context of a container.
- The **DefaultAllowPrivilegeEscalation** flag sets the default for the **allowPrivilegeEscalation** option.

For backward compatibility, the **AllowPrivilegeEscalation** flag defaults to **allowed**. If that behavior is not desired, this field can be used to default to **disallow**, while still permitting pods to request **allowPrivilegeEscalation** explicitly.

Forbidden and unsafe sysctls options

Security Context Constraints have two new options to control which sysctl options can be defined in a pod spec:

- The **forbiddenSysctls** option excludes specific sysctls.
- The **allowedUnsafeSysctls** option controls specific needs such as high performance or real-time application tuning.

All safe sysctls are enabled by default; all unsafe sysctls are disabled by default and must be manually allowed by the cluster administrator.

Removed oc deploy Command

The **oc deploy** command is deprecated in OpenShift Container Platform 3.7. The **oc rollout** command replaces this command.

Removed **oc env** and **oc volume** Commands

The deprecated **oc env** and **oc volume** commands are now removed. Use **oc set env** and **oc set volume** instead.

Removed the **oc ex config patch** Command

The **oc ex config patch** command will be removed in a future release, as the **oc patch** command replaces it.

oc export Now Deprecated

The **oc export** command is deprecated in OpenShift Container Platform 3.10. This command will be removed in a future release, as the **oc get --export** command replaces it.

oc types Now Deprecated

In OpenShift Container Platform 3.11, **oc types** is now deprecated. This command will be removed in a future release. Use the official documentation instead.

Pipeline Plug-in Is Deprecated

The OpenShift Container Platform Pipeline Plug-in is deprecated but continues to work with OpenShift Container Platform versions up to version 3.11. For later versions of OpenShift Container Platform, either use the **oc** binary directly from your Jenkins Pipelines or use the OpenShift Container Platform Client Plug-in.

Logging: Elasticsearch 5

Curator now works with Elasticsearch 5.

See [Aggregating Container Logs](#) for additional information.

Hawkular Now Deprecated

Hawkular is now deprecated and will be removed in a future release.

New Registry Source for Red Hat images

Instead of **registry.access.redhat.com**, OpenShift Container Platform now uses **registry.redhat.io** as the source of images for version 3.11. For access, **registry.redhat.io** requires credentials. See [Authentication Enabled Red Hat Registry](#) for more information.

New Storage Driver Recommendation

Red Hat strongly recommends [using the overlayFS storage driver instead of Device Mapper](#). For better performance, use overlayfs2 for Docker engine or overlayFS for CRI-O. Previously, we recommended using Device Mapper.

2.5. BUG FIXES

This release fixes bugs for the following components:

Builds

- ConfigMap Build Sources allows you to use ConfigMaps as a build source, which is transparent and easier to maintain than secrets. ConfigMaps can be injected into any OpenShift build. ([BZ#1540978](#))
- Information about out of memory (OOM) killed build pods is propagated to a build object. This information simplifies debugging and helps you discover what went wrong if appropriate failure reasons are described to the user. A build controller populates the status reason and message correctly when a build pod is OOM killed. ([BZ#1596440](#))

- The logic for updating the build status waited to update the log snippet containing the tail of the build log only ran after the build status changed to the failed state. The build would first transition to a failed state, then get updated again with the log snippet. This means code watching for the build to enter a failed state would not see the log snippet value populated initially. The code is now changed to populate the log snippet field when the build transitions to failed status, so the build update will contain both the failed state and the log snippet. Code that watches the build for a transition to the failed state will see the log snippet as part of the update that transitioned the build to failed, instead of seeing a subsequent update later. ([BZ#1596449](#))
- If a job used the **JenkinsPipelineStrategy** build strategy, the prune settings were ignored. As a result, setting **successfulBuildsHistoryLimit** and **failedBuildsHistoryLimit** did not correctly prune older jobs. The code has been changed to prune jobs properly. ([BZ#1543916](#))

Cloud Compute

- You can now configure NetworkManager for **dns=none** during installation. This configuration is commonly used when deploying OpenShift Container Platform on Microsoft Azure, but can also be useful in other scenarios. To configure this, set **openshift_node_dnsmasq_disable_network_manager_dns=true**. ([BZ#1535340](#))

Image

- Previously, because of improper handling of empty image stream updates, updates to an image stream that did not result in a change in tags resulted in a request to the image import API that included no content to be imported, which was invalid and lead to errors in the controller. Now, updates to the image stream that result in no new or updated tags that need to be imported will not result in an import API call. With this fix, invalid requests do not go to the import API, and no errors occur in the controller. ([BZ#1613979](#))
- Image pruning stopped on encountering any unexpected error while deleting blobs. In the case of an image deletion error, image pruning failed to remove any image object from etcd. Images are now being pruned concurrently in separated jobs. As a result, image pruning does not stop on a single unexpected blob deletion failure. ([BZ#1567657](#))

Installer

- When deploying to AWS, the **build_ami** play failed to clean **/var/lib/cloud**. An unclean **/var/lib/cloud** directory causes cloud-init to skip execution. Skipping execution causes a newly deployed node to fail to bootstrap and auto-register to OpenShift Container Platform. This bug fix cleans the **/var/lib/cloud** directory during **seal_ami** play. ([BZ#1599354](#))
- The installer now enables the router's extended route validation by default. This validation performs additional validation and sanitation of routes' TLS configuration and certificates. Extended route validation was added to the router in OpenShift Container Platform 3.3 and enhanced with certificate sanitation in OpenShift Container Platform 3.6. However, the installer did not previously enable extended route validation. There was initial concern that the validation might be too strict and reject valid routes and certificates, so it was disabled by default. But it has been determined to be safe to enable by default on new installs. As a result, extended route validation is enabled by default on new clusters. It can be disabled using by setting **openshift_hosted_router_extended_validation=False** in the Ansible inventory. Upgrading an existing cluster does **not** enable extended route validation. ([BZ#1542711](#))
- Without the fully defined **azure.conf** file when a load balancer service was requested through OpenShift Container Platform, the load balancer would never fully register and provide the external IP address. Now the **azure.conf**, with all the required variables, allows the load balancer to be deployed and provides the external IP address. ([BZ#1613546](#))

- To facilitate using CRI-O as the container-runtime for OpenShift Container Platform, update the *node-config.yaml* file with the correct endpoint settings. The **openshift_node_groups** defaults have been extended to include CRI-O variants for each of the existing default node groups. To use the CRI-O runtime for a group of compute nodes, use the following inventory variables:
 - **openshift_use_crio=True**
 - **openshift_node_group_name="node-config-compute-crio"**
Additionally, to deploy the Docker garbage collector, **docker gc**, the following variable must be set to **True**. This bug fix changes the previous variable default value from **True** to **False**:
 - **openshift_crio_enable_docker_gc=True** ([BZ#1615884](#))
- The *ansible.cfg* file distributed with **openshift-ansible** now sets a default log path of *~/openshift-ansible.log*. This ensures that logs are written in a predictable location by default. To use the distributed *ansible.cfg* file, you must first change directories to */usr/share/ansible/openshift-ansible* before running Ansible playbooks. This *ansible.cfg* file also sets other options meant to increase the performance and reliability of **openshift-ansible**. ([BZ#1458018](#))
- Installing Prometheus in a multi-zone or region cluster using dynamic storage provisioning causes the Prometheus pod to become unschedulable in some cases. The Prometheus pod requires three physical volumes: one for the Prometheus server, one for the Alertmanager, and one for the alert-buffer. In a multi-zone cluster with dynamic storage, it is possible that one or more of these volumes becomes allocated in a different zone than the others. This causes the Prometheus pod to become unschedulable due to each node in the cluster only able to access physical volumes in its own zone. Therefore, no node can run the Prometheus pod and access all three physical volumes. The recommended solution is to create a storage class which restricts volumes to a single zone using the **zone:** parameter, and assigning this storage class to the Prometheus volumes using the Ansible installer inventory variable, **openshift_prometheus_<COMPONENT>_storage_class=<zone_restricted_storage_class>**. With this workaround, all three volumes get created in the same zone or region, and the Prometheus pod is automatically scheduled to a node in the same zone. ([BZ#1554921](#))

Logging

- Previously, the **openshift-ansible installer** only supported **shared_ops** and **unique** as Kibana index methods. This bug fix allows users in a non-ops EFK cluster to share the default index in Kibana, to share queries, dashboards, and so on. ([BZ#1608984](#))
- As part of installing the ES5 stack, users need to create a *sysctl* file for the nodes that ES runs on. This bug fix evaluates which nodes/Ansible hosts to run the tasks against. ([BZ#1609138](#))
- Additional memory is required to support Prometheus metrics and retry queues to avoid periodic restarts from out-of-the-box memory. This bug fix increases out-of-the-box memory for Fluentd. As a result, Fluentd pods avoid out-of-the-box memory restarts. ([BZ#1590920](#))
- Fluentd will now reconnect to Elasticsearch every 100 operations by default. If one Elasticsearch starts before the others in the cluster, the load balancer in the Elasticsearch service will connect to that one and that one only, and so will all of the Fluentd connecting to Elasticsearch. With this enhancement, by having Fluentd reconnect periodically, the load balancer will be able to spread the load evenly among all of the Elasticsearch in the cluster. ([BZ#1489533](#))

- The rubygem ffi 1.9.25 reverted a patch, which allowed it to work on systems with SELinux **deny_execmem=1**. This causes Fluentd to crash. This bug fix reverts the patch reversion and, as a result, Fluentd does not crash when using SELinux **deny_execmem=1**. ([BZ#1628407](#))

Management Console

- The log viewer was not accounting for multi-line or partial line responses. If a response contained a multi-line message, it was appended and treated as a single line, causing the line numbers to be incorrect. Similarly, if a partial line were received, it would be treated as a full line, causing longer log lines sometimes to be split into multiple lines, again making the line count incorrect. This bug fix adds logic in the log viewer to account for multi-line and partial line responses. As a result, line numbers are now accurate. ([BZ#1607305](#))

Monitoring

- The **9100** port was blocked on all nodes by default. Prometheus could not scrape the **node_exporter** service running on the other nodes, which listens on port **9100**. This bug fix modifies the firewall configuration to allow incoming TCP traffic for the **9000 - 1000** port range. As a result, Prometheus can now scrape the **node_exporter** services. ([BZ#1563888](#))
- **node_exporter** starts with the **wifi** collector enabled by default. The **wifi** collector requires SELinux permissions that are not enabled, which causes AVC denials though it does not stop **node_exporter**. This bug fix ensures **node_exporter** starts with the **wifi** collector being explicitly disabled. As a result, SELinux no longer reports AVC denials. ([BZ#1593211](#))
- Uninstalling Prometheus currently deletes the entire **openshift-metrics** namespace. This has the potential to delete objects which have been created in the same namespace but are not part of the Prometheus installation. This bug fix changes the uninstall process to delete only the specific objects which were created by the Prometheus install and delete the namespace if there are no remaining objects, which allows Prometheus to be installed and uninstalled while sharing a namespace with other objects. ([BZ#1569400](#))

Pod

- Previously, a Kubernetes bug caused **kubectl drain** to stop when pods returned an error. With the [Kubernetes fix](#), the command no longer hangs if pods return an error. ([BZ#1586120](#))

Routing

- Because dnsmasq was exhausting the available file descriptors after the OpenShift Extended Conformance Tests and the Node Vertical Test, dnsmasq was hanging and new pods were not being created. A change to the code increases the maximum number of open file descriptors so the node can pass the tests. ([BZ#1608571](#))
- If 62 or more IP addresses are specified using an **haproxy.router.openshift.io/ip_whitelist** annotation on a route, the router will error due to exceeding the maximum parameters on the command (63). The router will not reload. The code was changed to use an overflow map if there are too many IPs in the whitelist annotation and pass the map to the HA-proxy ACL. ([BZ#1598738](#))
- By design, using a route with several services, when configuring a service with **set route-backend** set to **0**, the weight would drop all existing connections and associated end user connections. With this bug fix, a value of **0** means the server will not participate in load-balancing but will still accept persistent connections. ([BZ#1584701](#))
- Because the liveness and readiness probe could not differentiate between a pod that was alive and one that was ready, a router with **ROUTER_BIND_PORTS_AFTER_SYNC=true** was

reported as failed. This bug fix splits the liveness and readiness probe into separate probes, one for readiness and one for liveness. As a result, a router pod can be alive but not yet ready. ([BZ#1550007](#))

- When the HAproxy router contains a large number of routes (10,000 or more), the router will not pass the liveness and Readiness due to low performance, which kills the router repeatedly. The root cause of this issue is likely that a health check cannot be completed within the default readiness and liveness detection cycle. To prevent this problem, increase the interval of the probes. ([BZ#1595513](#))

Service Broker

- The deprovision process for Ansible Service Broker was not deleting secrets from the **openshift-ansible-service-broker** project. With this bug fix, the code was changed to delete all associated secrets upon Ansible Service Broker deprovisioning. ([BZ#1585951](#))
- Previously, the broker's reconciliation feature would delete its image references before getting the updated information from the registry, and there would be a period before the records appeared in the broker's data store while other jobs were still running. The reconciliation feature was redesigned to do an in-place update for items that have changed. For items removed from the registry, the broker deletes only those not already provisioned. It will also mark those items for deletion, which filters them out of the UI, preventing future provisions of those items. As a result, the broker's reconciliation feature makes provisioning and deprovisioning more resilient to registry changes. ([BZ#1577810](#))
- Previously, users would see an error message when an item was not found, even if it is normal not to be found. As a result, successful jobs might have an error message logged, causing the user concern that there might be a problem when there was none. The logging level of the message has now been changed from **error** to **debug**, because the message is still useful for debugging purposes, but not useful for a production installation, which usually has the level set to **info** or higher. As a result, users will not see an error message when the instance is not found unless there was an actual problem. ([BZ#1583587](#))
- If the cluster is not running or is not reachable, the **svcat version** command resulted in an error. The code has been changed to always report the client version, and if the server is reachable, it then reports the server version. ([BZ#1585127](#))
- In some scenarios, using the **svcat deprovision <service-instance-name> --wait** command sometimes resulted in the **svcat** command terminating with a panic error. When this happened, the **deprovision** command got executed, and the program then encountered a code bug when attempting to wait for the instance to be fully deprovisioned. This issue is now resolved. ([BZ#1595065](#))

Storage

- Previously, because the kubelet system containers could not write to the `/var/lib/iscsi` directory, iSCSI volumes could not be attached. Now, you can mount the host `/var/lib/iscsi` into the kubelet system container so that iSCSI volumes can be attached. ([BZ#1598271](#))

2.6. TECHNOLOGY PREVIEW FEATURES

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

[Technology Preview Features Support Scope](#)

In the table below, features marked **TP** indicate *Technology Preview* and features marked **GA** indicate *General Availability*.

Table 2.3. Technology Preview Tracker

Feature	OCP 3.9	OCP 3.10	OCP 3.11
Prometheus Cluster Monitoring	TP	TP	GA
Local Storage Persistent Volumes	TP	TP	TP
CRI-O for runtime pods	GA	GA* [a]	GA
Tenant Driven Snapshotting	TP	TP	TP
oc CLI Plug-ins	TP	TP	TP
Service Catalog	GA	GA	GA
Template Service Broker	GA	GA	GA
OpenShift Automation Broker	GA	GA	GA
Network Policy	GA	GA	GA
Service Catalog Initial Experience	GA	GA	GA
New Add Project Flow	GA	GA	GA
Search Catalog	GA	GA	GA
CFME Installer	GA	GA	GA
Cron Jobs	GA	GA	GA
Kubernetes Deployments	GA	GA	GA
StatefulSets	GA	GA	GA
Explicit Quota	GA	GA	GA
Mount Options		GA	GA

Feature	OCP 3.9	OCP 3.10	OCP 3.11
System Containers for Docker, CRI-O	Dropped	-	-
Installing from a System Container	GA	GA	GA
Hawkular Agent	-	-	-
Pod PreSets	-	-	-
experimental-qos-reserved	TP	TP	TP
Pod sysctls	TP	TP	TP
Central Audit	GA	GA	GA
Static IPs for External Project Traffic	GA	GA	GA
Template Completion Detection	GA	GA	GA
replicaSet	GA	GA	GA
Mux	TP	TP	TP
Clustered MongoDB Template	-	-	-
Clustered MySQL Template	-	-	-
Image Streams with Kubernetes Resources	GA	GA	GA
Device Manager	TP	GA	GA
Persistent Volume Resize	TP	TP	GA
Huge Pages	TP	GA	GA
CPU Manager	TP	GA	GA
Device Plug-ins	TP	GA	GA

Feature	OCP 3.9	OCP 3.10	OCP 3.11
syslog Output Plug-in for Fluentd	GA	GA	GA
Container Storage Interface (CSI)	-	TP	TP
Persistent Volume (PV) Provisioning Using OpenStack Manila	-	TP	TP
Node Problem Detector	-	TP	TP
Protection of Local Ephemeral Storage	-	TP	TP
Descheduler	-	TP	TP
Podman	-	TP	TP
Kuryr CNI Plug-in	-	TP	GA* [a]
Sharing Control of the PID Namespace	-	TP	TP
Cluster Administrator console	-	-	GA
Cluster Autoscaling (AWS Only)	-	-	GA
Operator Lifecycle Manager	-	-	TP
Red Hat OpenShift Service Mesh	-	-	TP
Multi-stage builds in Dockerfiles managed by the image builder	-	-	TP
[a] Features marked with * indicate delivery in a z-stream patch.			

2.7. KNOWN ISSUES

- Due to a change in the authentication for the Kibana web console, you must log back into the console after an upgrade and every 168 hours after initial login. The Kibana console has migrated to **oauth-proxy**. ([BZ#1614255](#))
- A Fluentd dependency on a systemd library is not releasing file handles. Therefore, the host eventually runs out of file handles. As a workaround, periodically recycle Fluentd to force the process to release unused file handles. See [Resolving Fluentd journald File Locking Issues](#) for more information on resolving this issue. ([BZ#1664744](#))

2.8. ASYNCHRONOUS ERRATA UPDATES

Security, bug fix, and enhancement updates for OpenShift Container Platform 3.11 are released as asynchronous errata through the Red Hat Network. All OpenShift Container Platform 3.11 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 via 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 3.11. Versioned asynchronous releases, for example with the form OpenShift Container Platform 3.11.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 [upgrading your cluster](#) properly.

2.8.1. RHBA-2018:3537 - OpenShift Container Platform 3.11.43 Bug Fix and Enhancement Update

Issued: 2018-11-19

OpenShift Container Platform release 3.11.43 is now available. The list of packages and bug fixes included in the update are documented in the [RHBA-2018:3537](#) advisory. The container images included in the update are provided by the [RHBA-2018:3536](#) advisory.

Space precluded documenting all of the bug fixes and enhancements for this release in the advisory. See the following sections for notes on upgrading and details on the bug fixes and enhancements included in this release.

2.8.1.1. Bug Fixes

- Log messages from a CRI-O pod could be split in the middle by nature. As a result, partial log messages were indexed in the Elasticsearch. The newer fluent-plugin-concat supports merging

the CRI-O style split messages into one, which is not available for the current fluentd (v0.12) that OpenShift Container Platform logging v3.11 uses. The functionality was backported to the fluentd v0.12. With this bug fix, the CRI-O style split log messages are merged back to the original full message. ([BZ#1552304](#))

- The event router intentionally generated duplicate event logs as to not lose them. The **elasticsearch_genid** plug-in is now extended to **elasticsearch_genid_ext** so that it takes the **alt_key** and **alt_tag**. If a log message has a tag matched the **alt_tag** value, it uses the **alt_key** value as the Elasticsearch primary key. You could specify a field, which is shared among the duplicate events to **alt_key**, which eliminates the duplicate events from the Elasticsearch. Sample filter using **elasticsearch_genid_ext**:

```
@type elasticsearch_genid_ext
hash_id_key viaq_msg_id
alt_key kubernetes.event.metadata.uid
alt_tags "#{ENV['GENID_ALT_TAG'] || 'kubernetes.var.log.containers.kube-eventrouter-
*.** kubernetes.journal.container._default_.kubernetes.event'}"
</filter>
```

With this bug fix, no duplicate event logs are indexed in Elasticsearch. ([BZ#1613722](#))

- The Netty dependency does not make efficient use of the heap. Therefore, Elasticsearch begins to fail on the network layer at a high logging volume. With this bug fix, the Netty recycler is disabled and Elasticsearch is more efficient in processing connections. ([BZ#1627086](#))
- The installer did not parameterize the configmap used by the Elasticsearch pods. The operations Elasticsearch pods used the configmap of the non-operations Elasticsearch pods. Parameterize the template used by the installer so that the pods use the **logging-es-ops** configmap. ([BZ#1627689](#))
- When using docker with the journald log driver, all container logs, including system and plain docker container logs, are logged to the journal, and read by fluentd. Consequently, fluentd does not know how to handle these non-Kubernetes container logs and throws exceptions. Treat non-Kubernetes container logs as logs from other system services (for example, send them to the operations index). Logs from non-Kubernetes containers are now indexed correctly and do not cause any errors. ([BZ#1632364](#))
- When using docker with log-driver journald, the setting in `/etc/sysconfig/docker` has changed to use **--log-driver** journald instead of **--log-driver=journald**. Fluentd cannot detect that journald is being used, so assumes **json-file**, and cannot read any Kubernetes metadata because it does not look for the journald **CONTAINER_NAME** field. This results in a lot of fluentd errors. Change the way Fluentd detects the docker log driver so that it looks for **--log-driver** journald in addition to **--log-driver=journald**. Fluentd can now detect the docker log driver, and can correctly process Kubernetes container logs. ([BZ#1632648](#))
- When fluentd is configured as the combination of collectors and MUX, event logs from the event were supposed to be processed by MUX, not by the collector for the both **MUX_CLIENT_MODE** maximal and minimal. This is because if an event log is formatted in the collector (and the event record is put under the Kubernetes key), the log is forwarded to MUX and passed to the k8s-meta plug-in there and the existing Kubernetes record is overwritten. It wiped out the event information from the log.
Fix 1: To avoid the replacement, if the log is from event router, the tag is rewritten to **#{tag}.raw** in `input-post-forward-mux.conf`, which makes the log treated in the **MUX_CLIENT_MODE=minimal** way.

Fix 2: There was another bug in Ansible. That is, the environment variable **TRANSFORM_EVENTS** was not set in MUX even if **openshift_logging_install_eventrouter** is set to **true**.

With these two bug fixes, the event logs are correctly logged when MUX is configured with **MUX_CLIENT_MODE=maximal** as well as minimal. ([BZ#1632895](#))

- In OpenShift Container Platform 3.10 and newer, the API server runs as a static pod and only mounted **/etc/origin/master** and **/var/lib/origin** inside that pod. CAs trusted by the host were not trusted by the API server. The API server pod definition now mounts **/etc/pki** into the pod. The API server now trusted all certificate authorities trusted by the host including those defined by the installer variable **openshift_additional_ca**. This can be used to import image streams from a registry verified by a private CA. ([BZ#1641657](#))
- The OSB Client Library used by the Service Catalog controller pod was not closing and freeing TCP connections used to communicate with brokers. Over a period of time, many TCP connections would remain open and eventually the communication between the Service Catalog controller and brokers would fail. Additionally, the pod would become unresponsive. Reuse the TCP connection when using the OSB Client Library. ([BZ#1641796](#))
- An unnecessarily short timeout resulted in a failure to reuse artifacts from a previous build when incremental builds were selected with S2I. This could occur when the size of the artifacts being reused was particularly large or the host system was running particularly slowly. Invalid artifacts could be used in a subsequent build, or artifacts would be recreated instead of reused resulting in performance degradation. With this bug fix, the timeout is increased to a sufficiently large value to avoid this problem. Artifact reuse should no longer timeout. ([BZ#1642350](#))
- The Automation Broker always created a network policy to give the transient namespace access to the target namespace. Adding a network policy to a namespace that does not have any other network policies in place causes the namespace to be locked down to the newly created policy. Before the network policy, everything was open and namespaces could communicate with each other. The Automation Broker now looks to see if there are any network policies in place for the target namespace. If there are none, the broker will not create a new network policy. The broker will assume that things are open enough to allow the transient namespace we create to communicate with the target namespace. The broker will still create a network policy giving the transient namespace access to the target namespace, if there are other network policies in place for the target namespace. This bug fix allows the broker to perform the APB actions without affecting existing services running on the target namespace. ([BZ#1643301](#))
- Previously, the cluster console in OpenShift Container Platform 3.11 would always show the value **0** for the crashlooping pods count on the cluster status page, even when there were crashlooping pods. The problem is now fixed and the count now accurately reflects the count for the selected projects. ([BZ#1643948](#))

2.8.1.2. Upgrading

To upgrade an existing OpenShift Container Platform 3.10 or 3.11 cluster to this latest release, see [Upgrade methods and strategies](#) for instructions.

2.8.2. RHBA-2018:3743 - OpenShift Container Platform 3.11.51 Bug Fix and Enhancement Update

Issued: 2018-12-12

OpenShift Container Platform release 3.11.51 is now available. The list of packages and bug fixes included in the update are documented in the [RHBA-2018:3743](#) advisory. The container images included in the update are provided by the [RHBA-2018:3745](#) advisory.

2.8.2.1. Upgrading

To upgrade an existing OpenShift Container Platform 3.10 or 3.11 cluster to this latest release, see [Upgrade methods and strategies](#) for instructions.

2.8.3. RHBA-2018:3688 - OpenShift Container Platform 3.11 Package Updates for IBM POWER

Issued: 2018-12-13

OpenShift Container Platform release 3.11 is now available with updates to packages for ppc64le. The list of packages and bug fixes included in the update are documented in the [RHBA-2018:3688](#) advisory.

2.8.3.1. Upgrading

To upgrade an existing OpenShift Container Platform 3.10 or 3.11 cluster to this latest release, see [Upgrade methods and strategies](#) for instructions.

2.8.4. RHBA-2019:0024 - OpenShift Container Platform 3.11.59 Bug Fix and Enhancement Update

Issued: 2019-01-10

OpenShift Container Platform release 3.11.59 is now available. The list of packages and bug fixes included in the update are documented in the [RHBA-2019:0024](#) advisory. The container images included in the update are provided by the [RHBA-2019:0023](#) advisory.

Space precluded documenting all of the bug fixes and enhancements for this release in the advisory. See the following sections for notes on upgrading and details on the bug fixes and enhancements included in this release.

2.8.4.1. Bug Fixes

- The openshift-ansible OpenStack playbook defaulted to the Kuryr-Kubernetes multi-pool driver, but that functionality was not merged on stable/queens kuryr-controller. This bug fix adds the option to select the pool driver to use for versions older than stable/queens. For newer versions, it will suffice with setting the **kuryr_openstack_pool_driver** to **multi** as described in the documentation. ([BZ#1573128](#))
- The Openshift Ansible installer did not check if any CNS are created before creating a security group. It would create a security group for CNS even when there were none created. The Openshift Ansible installer now checks that **openshift_openstack_num_cns** is greater than zero before creating a security group for CNS. CNS security groups are now only created when there is at least one CNS created. ([BZ#1613438](#))
- The ability to leave swap enabled is now removed and the **openshift_disable_swap** variable is deprecated. This variable was never publicly documented and was only used internally. Documentation has stated that system swap should be disabled since version 3.4. ([BZ#1623333](#))

- An incorrect **etcdctl** command was used during etcd backup for system containers, causing the etcd backup to fail during upgrade. The etcd system container is now identified correctly. The upgrade succeeds with etcd in the system container. ([BZ#1625534](#))
- During etcd scaleup, facts about the etcd cluster are required in order to add new hosts. The necessary tasks are now added to ensure those facts are set before configuring new hosts and, therefore, allow the scale-up to complete as expected. ([BZ#1628201](#))
- The default log format for audit was set to **json**. The audit log was always printed using JSON format. You can now set the log format as specified in the **master-config.yaml** file. The audit log now contains values per the configured log format. ([BZ#1632155](#))
- **sync daemonset** did not run on all nodes. The pgrade failed, as some nodes did not have an annotation set. With this bug fix, **sync daemonset** now tolerates all taints and runs on all nodes and the upgrade succeeds. ([BZ#1635462](#))
- **sync daemonset** did not wait a sufficient amount of time for nodes to restart. The sync DS verification task failed, as nodes did not come up in time. A number of retries was increased and the install or upgrade now succeeds. ([BZ#1636914](#))
- A deployment would take longer than some of the infrastructure or API server-related timeouts. Long-running deployments would fail. The deployer is now fixed to tolerate long running deployments by re-establishing the watch. ([BZ#1638140](#))
- Ansible 2.7.0 changed the way variables were passed to roles. Some roles did not have necessary variables set, resulting in a failed installation. The required Ansible version is now set to 2.6.5 and the installation succeeds. ([BZ#1638699](#))
- Node, pod, and control-plane images were not pre-pulled when CRI-O was used. Tasks timed out, as they included pull time. Images are now pre-pulled when Docker and CRI-O are used and the installation succeeds. ([BZ#1639201](#))
- The scale-up playbooks, when used in conjunction with Calico, did not properly configure the Calico certificate paths causing them to fail. The playbooks have been updated to ensure that master scale-up with Calico works properly. ([BZ#1644416](#))
- In some cases, CRI-O was restarted before verifying that the image pre-pull was finished. Images were not pre-pulled. Now, CRI-O is restarted before image pre-pull begins and installation succeeds. ([BZ#1647288](#))
- The CA was not copied to the master config directory when GitHub Enterprise was used as a identity provider. The API server failed to start without a CA. New variables, **openshift_master_github_ca** and **openshift_master_github_ca_file**, were introduced to set the GitHub Enterprise CA and installation now succeeds. ([BZ#1647793](#))
- The **curator** image was built with the wrong version of the **python-elasticsearch** package and the **curator** image would not start. Use the correct version of the **python-elasticsearch** package to build the **curator** image and the **curator** image works as expected. ([BZ#1648453](#))
- There was improper evaluation of a user's Kibana index. A minor upgrade in server version caused an error when the expected configuration object was not as expected. Its reation was skipped due to the existence of kibana index. Remove a user's Kiana index, evaluate the stored version against the Kibana version, and recreate the configuration object if necessary. With this bug fix, users will no longer see the error. ([BZ#1652224](#))
- Egress IP-related iptables rules were not recreated if they were deleted. If a user restarted **firewalld** or **iptables.service** on a node that hosted egress IPs, then those egress IPs would stop

working. Traffic that should have used the egress IP would use the node's normal IP instead. Egress IP iptables rules are now recreated if they are removed. Egress IPs now work reliably. ([BZ#1653380](#))

- A bug in earlier releases of **cluster-logging** introduced Kibana index-patterns where the title was not properly replaced and was left with the placeholder of **'\$TITLE\$'**. As a result, the user sees a permission error of **no permissions for [indices:data/read/field_caps]** Remove all index-patterns that have the bad data, either by upgrading or running:

```
$ oc exec -c elasticsearch -n $NS $pod --es_util \
--query=".kibana.*/_delete_by_query?pretty" -d \
"{\"query\":{\"match\":{\"title\":\"*TITLE*\"}}}"
```

With this bug fix, the permission error is no longer generated. ([BZ#1656086](#))

2.8.4.2. Enhancements

- A new playbook was added to cleanup etcd2 data. If the cluster was upgraded from OpenShift Container Platform 3.5, it might still carry etcd2 data and use up space. The new playbook safely removes etcd2 data. ([BZ#1514487](#))
- A new multi-pool driver is added to Kuryr-Kubernetes to support hybrid environments where some nodes are bare metal while others are running inside VMs, therefore having different pod VIF drivers (e.g., **neutron** and **nested-vlan**). To make use of this new feature, the available configuration mappings for the different pools and **pod_vif** drivers need to be specified in the **kuryr.conf** configmap. In addition, the nodes must be annotated with the correct information about the **pod_vif** to be used. Otherwise, the default one is used. ([BZ#1553070](#))
- Scale out Ansible playbooks for the OpenStack deployed clusters are now added. When installing OpenShift on top of OpenStack with the OpenStack provisioning playbooks (**playbooks/openstack/openshift-cluster/provision_install.yml**), scaling the cluster out required several manual steps such as writing the inventory by hand and running two extra playbooks. This was more brittle, required more complex documentation, and did not match the initial deployment experience. To scale out OpenShift on OpenStack, you can now change the desired number of nodes and run one of the following playbooks (depending on whether you want to scale the worker or master nodes):

```
playbooks/openstack/openshift-cluster/node-scaleup.yml
playbooks/openstack/openshift-cluster/master-scaleup.yml
```

([BZ#1566424](#))

- Define the recreate strategy timeout for Elasticsearch. There are examples on AWS OpenShift clusters where rollout of new Elasticsearch pods fail because the cluster is having issues attaching storage. Defining a long recreate timeout allows the the cluster more time to attach storage to the new pod. Elasticsearch pods have more time to restart and experience fewer rollbacks. ([BZ#1655675](#))

2.8.4.3. Upgrading

To upgrade an existing OpenShift Container Platform 3.10 or 3.11 cluster to this latest release, see [Upgrade methods and strategies](#) for instructions.

2.8.5. RHBA-2019:0096 - OpenShift Container Platform 3.11.69 Bug Fix and Enhancement Update

Issued: 2019-01-31

OpenShift Container Platform release 3.11.69 is now available. The list of packages and bug fixes included in the update are documented in the [RHBA-2019:0096](#) advisory. The container images included in the update are provided by the [RHBA-2019:0097](#) advisory.

Space precluded documenting all of the bug fixes and enhancements for this release in the advisory. See the following sections for notes on upgrading and details on the bug fixes and enhancements included in this release.

2.8.5.1. Bug Fixes

- The location of the master proxy API changed. Since the MetricsApiProxy diagnostic uses this endpoint, it broke. The diagnostic was updated to look at the correct endpoint and it should now work as expected. ([BZ#1632983](#))
- Pods would not schedule because they did not have free ports. This issue is now resolved. ([BZ#1647674](#))
- Bootstrap v3.3.5 contains a Cross-Site Scripting (XSS) vulnerability. The management console does not allow user input to be displayed via a data-target attribute. Upgrade Bootstrap to v3.4.0, which fixes the vulnerability. With this bug fix, the management console is no longer at risk of possible exploit via the Cross-Site Scripting (XSS) vulnerability in Bootstrap v3.3.5. ([BZ#1656438](#))
- Improper error checking ignored errors from object creation during template instantiation. Template instances would report successful instantiation when some objects in the template failed to be created. Errors on creation are now properly checked and the template instance will report failure if any object within it cannot be created. ([BZ#1662339](#))
- The rsync package was removed from the registry image, so rsync cannot be used to backup content from the registry container. The rsync package is now added back to the image and can now be used. ([BZ#1664853](#))

2.8.5.2. Enhancements

- This enhancement ensures that OpenShift-on-OpenStack playbook execution will fail at the prerequisites check if the public net ID is not configured when the Kuryr SDN is used. ([BZ#1579414](#))
- You can now control the assignment of floating IP addresses for OpenStack cloud provisioning. The playbook responsible for creating the OpenStack virtual servers would always associate a floating IP address with each virtual machine (each OpenShift node). This had two negative implications:
 1. The OpenShift cluster size was limited by the number of floating IPs available to the OpenStack user.
 2. All OpenShift nodes were directly accessible from the outside, increasing the potential attack surface.
A role-based control over which nodes get floating IPs and which do not is now introduced. This is controlled by the following inventory variables:
- **openshift_openstack_master_floating_ip**
- **openshift_openstack_infra_floating_ip**

- **openshift_openstack_compute_floating_ip**

- **openshift_openstack_load_balancer_floating_ip**

They are all boolean and all default to **true**. This allows for use cases such as:

- A cluster where all the master and infra nodes have floating IPs but the compute nodes do not.
- A cluster where none of the nodes have floating IPs, but the load balancers do (so OpenShift is used through the load balancers, but none of the nodes are directly accessible).

If some of the nodes do not have floating IPs (by setting **openshift_openstack_compute_floating_ip = false**), the openshift-ansible playbooks must be run from inside the node network. This is because a server without a floating IP is only accessible from the network it is in. A common way to do this is to pre-create the node network and subnet, create a "bastion" host in it, and run Ansible there:

```
$ openstack network create openshift
$ openstack subnet create --subnet-range 192.168.0.0/24 --dns-nameserver 10.20.30.40 --
network openshift openshift
$ openstack router create openshift-router
$ openstack router set --external-gateway public openshift-router
$ openstack router add subnet openshift-router openshift
$ openstack server create --wait --image RHEL7 --flavor m1.medium --key-name openshift --
network openshift bastion
$ openstack floating ip create public
$ openstack server add floating ip bastion 172.24.4.10
$ ping 172.24.4.10
$ ssh cloud-user@172.24.4.10
```

Then, install openshift-ansible and add the following to the inventory (*inventory/group_vars/all.yml*):

```
openshift_openstack_node_network_name: openshift
openshift_openstack_router_name: openshift-router
openshift_openstack_node_subnet_name: openshift
openshift_openstack_master_floating_ip: false
openshift_openstack_infra_floating_ip: false
openshift_openstack_compute_floating_ip: false
openshift_openstack_load_balancer_floating_ip: false
```

([BZ#1611839](#))

2.8.5.3. Upgrading

To upgrade an existing OpenShift Container Platform 3.10 or 3.11 cluster to this latest release, see [Upgrade methods and strategies](#) for instructions.

2.8.6. RHBA-2019:0326 - OpenShift Container Platform 3.11.82 Bug Fix Update

Issued: 2019-02-20

OpenShift Container Platform release 3.11.82 is now available. The list of packages and bug fixes included in the update are documented in the [RHBA-2019:0326](#) advisory. The container images included in the update are provided by the [RHBA-2019:0327](#) advisory.

Space precluded documenting all of the bug fixes and enhancements for this release in the advisory. See the following sections for notes on upgrading and details on the bug fixes and enhancements included in this release.

2.8.6.1. Bug Fixes

- All Docker related packages are not removed during the uninstall process. Docker is not re-installed properly during installation, causing Docker CLI tasks to fail. With this bug fix, all related Docker packages to uninstall are now added. Re-installation succeeds after running the uninstall playbook. ([BZ#1635254](#))
- Polling of quotas resulted in undesirable toast notifications. Now, quota polling errors are suppressed and users no longer see these notifications. ([BZ#1651090](#))
- Previously, running the install playbook multiple times with no changes to the cluster console configuration could cause the cluster console login to stop working. The underlying problem has been fixed, and now running the playbook more than once will correctly roll out a new console deployment. This problem can be worked around without the installer fix by manually deleting the console pods using the command:

```
$ oc delete --all pods -n openshift-console
```

([BZ#1651632](#))

- Certain certificate expiry check playbooks did not call properly initialization functions resulting in an error. Those playbooks have been updated to avoid this problem. ([BZ#1655183](#))
- The OpenShift SDN/OVS DaemonSets were upgraded during control plane upgrades with an **updateStrategy** of **RollingUpdate**, an upgrade of the pods in the entire cluster was performed. This caused unexpected network and application outages on nodes. This bug changed the **updateStrategy** for SDN/OVS pods to **OnDelete** in the template, affecting only new installations. Control plane upgrade tasks were added to modify SDN/OVS daemonsets to use **OnDelete updateStrategy**. Node upgrade tasks were added to delete all SDN/OVS pods while nodes are drained. Network outages for nodes should only occur during the node upgrade when nodes are drained. ([BZ#1657019](#))
- Previously, the 3.11 admin console did not correctly display whether a storage class was the default storage class, as it was checking an out-of-date annotation value. The admin console has been updated to use the **storageclass.kubernetes.io/is-default-class=true** annotation, and service classes are now properly marked as default when that value is set. ([BZ#1659976](#))
- A change introduced in Kubernetes 1.11 affected nodes with many IP addresses in **vSphere** deployments. Under vSphere, a node hosting several **Egress IPs** or **Router HA** addresses would sporadically lose IP addresses and start using one of the other ones, causing networking problems. Now, if a **node IP** is specified in the node configuration, it will be used correctly, regardless of how many other IP addresses are assigned to the node. ([BZ#1666820](#))
- A type error in the OpenStack code prevented installation on OpenShift nodes without floating IP addresses. This error has been corrected, and installation proceeds as expected. ([BZ#1667270](#))
- Certain certificate expiry check playbooks did not call initialization functions properly, resulting in an error. Those playbooks have been updated to avoid this issue. ([BZ#1667618](#))
- The cluster role **system:image-pruner** was required for all DELETE requests to the registry. As a result, the regular client could not cancel its uploads, and the **S3 multipart** uploads were

accumulating. Now, the cluster role **system:image-pruner** will accept DELETE requests for uploads from clients who are allowed to write into them. ([BZ#1668412](#))

- If the specified router certificate, key, or CA did not end with a new line character, the router deployment would fail. A new line is now appended to each of the input files ensuring this problem doesn't occur. ([BZ#1668970](#))
- The **volume-config.yaml** was not copied to `/etc/origin/node`. As a result, volume quotas were not observed, so local storage size was not limited. Now, the **volume-config.yaml** is copied to `/etc/origin/node`. Volume quotas are observed and local storage size is limited by setting **openshift_node_local_quota_per_fsgroup** in the inventory. ([BZ#1669555](#))
- **oc** image mirror failed with error **tag: unexpected end of JSON input** when attempting to mirror images from Red Hat registry. This was a result of commits from a dependency were dropped from the product build. The commits have been re-introduced, and the command can now parse the output successfully, as well as mirror from the Red Hat registry. ([BZ#1670551](#))

2.8.6.2. Upgrading

To upgrade an existing OpenShift Container Platform 3.10 or 3.11 cluster to this latest release, see [Upgrade methods and strategies](#) for instructions.

2.8.7. RHBA-2019:0407 - OpenShift Container Platform 3.11.88 Bug Fix Update

Issued: 2019-03-14

OpenShift Container Platform release 3.11.88 is now available. The list of packages and bug fixes included in the update are documented in the [RHBA-2019:0407](#) advisory. The container images included in the update are provided by the [RHBA-2019:0406](#) advisory.

With this release, [Kuryr](#) is now moved out of [Technology Preview](#) and now generally available.

2.8.7.1. Upgrading

To upgrade an existing OpenShift Container Platform 3.10 or 3.11 cluster to this latest release, see [Upgrade methods and strategies](#) for instructions.

2.8.8. RHSA-2019:0739 - Important: OpenShift Container Platform 3.11 jenkins-2-plugins security update

Issued: 2019-04-10

An update for jenkins-2-plugin is now available for OpenShift Container Platform 3.11. Details of the update are documented in the [RHSA-2019:0739](#) advisory.

2.8.8.1. Upgrading

To upgrade an existing OpenShift Container Platform 3.10 or 3.11 cluster to this latest release, see [Upgrade methods and strategies](#) for instructions.

2.8.9. RHBA-2019:0636 - OpenShift Container Platform 3.11.98 Bug Fix and Enhancement Update

Issued: 2019-04-11

OpenShift Container Platform release 3.11.98 is now available. The list of packages and bug fixes included in the update are documented in the [RHBA-2019:0636](#) advisory. The container images included in the update are provided by the [RHBA-2019:0637](#) advisory.

Space precluded documenting all of the bug fixes and enhancements for this release in the advisory. See the following sections for notes on upgrading and details on the bug fixes and enhancements included in this release.

2.8.9.1. Bug Fixes

- Administrative users were not able to access the cluster endpoints because permissions were defined incorrectly. Now, the correct permissions have been defined, and administrative users can use the `_cat` endpoints. ([BZ#1548640](#))
- Image garbage collection failed to remove an image correctly if it has only one tag but more than one repository associated with the image. This has now been resolved and garbage collection completes successfully. ([BZ#1647348](#))
- The **docker** registry Health Check would fail if the bucket was empty on AWS S3 environments, returning a **PathNotFound** message. Now, **PathNotFound** is treated as a success and Health Check works as expected for empty buckets. ([BZ#1655641](#))
- Playbooks ran a check to see if images existed on the disk with specific version tags, but did not ensure the version on the disk was up-to-date to the tagged image in the repo, resulting in skipping the z-stream image pulls, and z-stream upgrades would fail. Now, the on-disk check has been removed, and image pulls are efficient so that there is no need to check whether the image exists on the disk prior to downloading. ([BZ#1658387](#))
- Health Check playbooks would fail at checking **Elasticsearch** because the `exec` call would not specify a container. The call failed because the output included incorrectly formatted JSON text. Now, the target container is included in the `exec` call and the Health Check succeeds. ([BZ#1660956](#))
- An error in **glusterfs** pod mount points prevented the use of **gluster-block**. As a result, the provisioner would fail to create devices. The mount points have now been updated and the provisioning process succeeds as expected. ([BZ#1662312](#))
- The **openshift-ansible** package was incorrectly checking if a value in the **etcd-servers-overrides** was a valid path. Some values were considered invalid by the **openshift-ansible-3.11.51-2.git.0.51c90a3.el7.noarch** package. Now, **etcd-servers-overrides** does not contain paths, and is ignored during path checks. ([BZ#1666491](#))
- **etcd** non-master host nodes were excluded from upgrades. Now, **etcd** host nodes are able to be upgraded. ([BZ#1668317](#))
- The Ansible variable **openshift_master_image_policy_allowed_registries_for_import** was incorrectly parsed, causing a corrupted **master-config.yaml** file. Now, the **openshift_master_image_policy_allowed_registries_for_import** variable is correctly parsed and a simple registry image policy can be set as expected. ([BZ#1670473](#))
- The playbooks and manual configuration steps to redeploy router certificates were replaced with service serving certificates secret. This would overwrite or miss the router wild certificates secret, causing certificate errors due to incorrect certificates redeployed. Now, the playbooks and manual redeployment steps do not overwrite router certificates secret. The router certificates are redeployed based on the specified sub domain or customer certificates. ([BZ#1672011](#))

- The **ImageStream** used in the **BuildConfig** editor did not have edit properties, causing runtime errors in the **BuildConfig** editor. Now, the editor is initializing tags and objects, even if **ImageStream** in the **BuildConfig** is missing or if the user does not have the correct permissions to use it. ([BZ#1672904](#))
- Master pods did not match time zones with worker nodes, which led to errors in logging timestamps. Now, the host's timezone configuration is mounted into the control plane pods. ([BZ#1674170](#))
- When a cluster was installed, the user name in the loopback kubeconfig is the same as the host name of the master. Now, the variable in the playbook is changed to a different value. ([BZ#1675133](#))
- The Ansible Health Check playbook failed when checking the **curator** status. This occurred because the Health Check assumed **curator** is a **DeploymentConfig** instead of a **cronjob**, resulting in a failed check. Now, Health Check properly evaluates for a **cronjob** instead of a **DeploymentConfig**. ([BZ#1676720](#))
- Some namespaces would be missing from **oc get projects** if more than 1,000 projects were listed. Now, all items correctly appear when looking at large resource lists. ([BZ#1677545](#))
- High network latency existed between **Kibana** and **Elasticsearch** due to either network issues or under-allocated memory for **Elasticsearch**. As a result, **Kibana** would be unusable because of a gateway timeout. Now, changes are backported from **Kibana** version 6, which allows modification to the ping timeout. Administrators are not able to override the default **pingTimeout** of 3000ms by setting the **ELASTICSEARCH_REQUESTTIMEOUT** environment variable. **Kibana** is functional until the underlying network issues or under-allocated memory conditions can be resolved. ([BZ#1679159](#))
- The **defaultIndex** in the **Kibana** config entry was null, causing the seeding process to fail and the user was presented with a white screen. Now, the **defaultIndex** value is evaluated and returns to the default screen if there is a null value. The **Kibana** seeding process completes successfully. ([BZ#1679613](#))
- Previously, the upgrade process for **CRI-O** would attempt to stop **docker** on nodes that had been configured to only run **CRI-O**, resulting in playbook failures. Now, the playbook does not stop **docker** on nodes that are configured only for **CRI-O** operation, ensuring successful upgrades. ([BZ#1685072](#))
- Using **MERGE_JSON_LOG=true** would create fields in the record that would cause syntax violations or create too many fields in **Elasticsearch**, causing severe performance problems. Now, users who experience these problems can tune **fluentd** to accommodate their log record fields without errors or **Elasticsearch** performance degradation. ([BZ#1685243](#))
- The SSL and TLS service uses Diffie-Hellman groups with insufficient strength (a key size less than 2048 bytes). As a result, the keys are more vulnerable. Now, the key strength has been increased and certificates are more secure. ([BZ#1685618](#))
- The **fluentd** daemonset did not include a **tolerate everything** toleration. If a node became tainted, the **fluentd** pod would get evicted. Now, a **tolerate everything** toleration has been added, and **fluentd** pods do not get evicted. ([BZ#1685970](#))
- Upgrade playbooks ran several **oc** commands that used resource aliases that may not be immediately available after a restart or other reasons. Now, the **oc** suite of commands uses the fully qualified resource name to avoid potential failure. ([BZ#1686590](#))
- The files that implemented log rotation functionality were not copied to the correct **fluentd**

directory. As a result, logs were not being rotated. Now, the container build has been changed to inspect the **fluentd** gem to find out where to install the files. The files that implement log rotation are copied to the correct directory for **fluentd** usage. ([BZ#1686941](#))

2.8.9.2. Enhancements

- The command **oc label --list** is now added, and now shows the resource and name of all the labels. ([BZ#1268877](#))
- This enhancement allows the AWS cloud provider to parse additional endpoint configuration and customization of both core Kubernetes and cluster autoscaler environments. AWS now allows custom and private regions, which do not follow the conventions of their public cloud endpoints. OpenShift Container Platform deployments were limited to the public AWS cloud regions only, and this limited the adoption of the product in these scenarios. Additional configuration elements can be added to the **aws.conf** file and will be honored by OpenShift Container Platform as well as the **cluster-autoscaler** to ensure the correct cloud endpoints are used to automatically provision EBS volumes, load balancers, and EC2 instances. ([BZ#1644084](#))

2.8.9.3. Upgrading

To upgrade an existing OpenShift Container Platform 3.10 or 3.11 cluster to this latest release, see [Upgrade methods and strategies](#) for instructions.

2.8.10. RHBA-2019:0794 - OpenShift Container Platform 3.11.104 Bug Fix Update

Issued: 2019-06-06

OpenShift Container Platform release 3.11.104 is now available. The list of packages and bug fixes included in the update are documented in the [RHBA-2019:0794](#) advisory. The container images included in the update are provided by the [RHBA-2019:0795](#) advisory.

2.8.10.1. Upgrading

To upgrade an existing OpenShift Container Platform 3.10 or 3.11 cluster to this latest release, use the automated upgrade playbook. See [Performing Automated In-place Cluster Upgrades](#) for instructions.

2.8.11. RHBA-2019:1605 - OpenShift Container Platform 3.11.117 Bug Fix Update

Issued: 2019-06-26

OpenShift Container Platform release 3.11.117 is now available. The list of packages and bug fixes included in the update are documented in the [RHBA-2019:1605](#) advisory. The container images included in the update are provided by the [RHBA-2019:1606](#) advisory.

2.8.11.1. Bug Fixes

- The **oc create route dry-run -o yaml** command would not output a route object. This has been resolved by implementing the printing of the route object to the command line. ([BZ#1418021](#))
- Some **.operations** index projects were given a value of **default openshift-**. This has now been changed to **kube-system**. ([BZ#1571190](#))

- On a director-deployed OpenShift environment, the GlusterFS playbooks auto-generate a new heketi secret key for each run. As a result of this, operations such as scale out or configuration changes on CNS deployments fail. As a workaround, complete the following steps:

1. Post-deployment, retrieve the heketi secret key. Use this command on one of the master nodes:

```
$ sudo oc get secret heketi-storage-admin-secret --namespace glusterfs -o json | jq -r
.data.key | base64 -d
```

2. In an environment file, set the following parameters to that value:

```
openshift_storage_glusterfs_heketi_admin_key
openshift_storage_glusterfs_registry_heketi_admin_key
```

As a result of this workaround, operations such as scale out or configuration changes on CNS deployments work as long as the parameters were manually extracted. ([BZ#1640382](#))

- When a new CA was generated, the certificates on the nodes were not updated and would not become ready. Now, the redeploy-certificates playbook will copy the certificates and join nodes. Nodes no longer go to a **NotReady** state when replacing the CA. ([BZ#1652746](#))
- The `oc_adm_router` Ansible module allowed edits to add duplicate environment variables to the router `DeploymentConfig`. An Ansible inventory file that specified edits to the router `DeploymentConfig` that added duplicate environment variables could produce a `DeploymentConfig` with unpredictable behavior. If an edit appends an environment variable to the router `DeploymentConfig`, and a variable by that name already exists, the `oc_adm_router` module now deletes the old variable. Using an Ansible inventory file to append environment variables to the router `DeploymentConfig` now has predictable behavior and allows users to override default environment variable settings. ([BZ#1656487](#))
- A playbook which redeployed master certificates did not update web console secrets, causing the web console to fail to start. Now, web console secrets are recreated when the master certificate redeployment playbook is run. ([BZ#1667063](#))
- The logging playbooks did not work with Ansible 2.7. The **include_role** and **import_role** behavior changed between versions 2.6 and 2.7, which caused issues with logging. As a result, errors with "-ops" suffixes would appear even when not deploying with the ops cluster. To resolve this, use **include_role** instead of **import_role** in logging playbooks and roles. The logging Ansible code works on both Ansible 2.6 and Ansible 2.7. ([BZ#1671315](#))
- Undesired DNS IP addresses were selected by the OpenShift service if multiple network cards were present. As a result, DNS requests failed to work from pods. Now, there are sane defaults present for DNS and it follows a similar pattern used by kubelet to fetch routable node IP addresses. ([BZ#1680059](#))
- Initialization during upgrades was slow. Sanity checks were using inefficient code to validate host variables. This code has been updated and host variables are now stored in the class. As a result, the host variables are not being copied on every check. The sanity checks and initialization during upgrades takes less time to complete. ([BZ#1682924](#))
- The **oreg_url** variable would not function correctly on disconnected installs using Satellite because the `etcd` image could not perform pulls on disconnected installs. Now, guidance and examples have been added to specify the `etcd` image URL issuing **osm_etcd_image** in the associated documentation. ([BZ#1689796](#))

- If a build pod was evicted, the build reported a **GenericBuildFailure**. Determining the cause of build failures was difficult as a result. Now a new failure reason, **BuildPodEvicted**, has been added. ([BZ#1690066](#))
- Nodes would sometimes panic due to cadvisor index reporting out of range errors. This has now been resolved by a backporting of kube code. ([BZ#1691023](#))
- ElasticSearch could not be monitored with Prometheus because the **oauth-proxy** was not passing a user's token. Now, the token is exchanged to ElasticSearch and users with proper roles can retrieve metrics in Prometheus. ([BZ#1695903](#))
- Deploying nodes would fail in the **setup_dns.yaml** playbook during multi-node setup. This was resolved by fixing the host name that was passed to the **add_host** function. Now, multi-node setup proceeds as expected. ([BZ#1698922](#))
- Upgrading between minor versions would fail because several OpenShift variables were not used during the upgrade process. Now, **api_port** and other apiserver-related variables are read during the upgrade process and upgrades complete successfully. ([BZ#1699696](#))
- ElasticSearch would fail to start due to invalid certificate dates if hosts had non-UTC timezones. When OpenShift nodes' timezone is not set to UTC, the current non-UTC timestamp is used for the **NotBefore** checking. If the timezone is ahead of UTC, the **NotBefore** checking would fail. Now, regardless of the nodes' timezone, the UTC timestamp is set to the start date in the certificates and failures are not reported due to non-UTC timestamps. ([BZ#1702544](#))
- **CustomResourceDefinition** errors were presented in a confusing manner that made troubleshooting difficult. Now, the CRD error messages have been clarified to assist in troubleshooting CRD errors. ([BZ#1702693](#))
- There was a missing **@** for an instance variable in the Fluentd remote syslog plugin code. In some cases, systemd-journald logged errant values. This resulted in rsyslog forwarding failures. Now, the variable has been corrected and remote logging completes successfully. ([BZ#1703904](#))
- Long running Jenkins agents and slave pods would experience defunct process errors, causing a high number of processes to appear in process listings until the pod is terminated. Now, **dumb-init** is deployed to clean up these defunct processes. ([BZ#1707448](#))
- The environment variable **JOURNAL_READ_FROM_HEAD** was set to an empty string. This caused the default value of **read_from_head** for the journald input to be true. When Fluentd starts up for the first time on a node, it reads in the entire journal. This could result in hours of delays for system messages to show up in ElasticSearch and Kibana. Now, Fluentd will check if the value is set and is not empty, or will use the default value of false. Fluentd will read from the tail of the journal when it starts on a new node. ([BZ#1707524](#))
- The script **99-origin-dns.sh** had a debug flag set to enabled, which would log debug level messages by default. This has been resolved and debug is now set to false. ([BZ#1707799](#))
- Kubernetes pod templates were removed at random. This was because the OpenShift Jenkins Sync plugin confused ImagesStreams and ConfigMaps with the same name while processing them. An event for one type could delete the pod template created for another type. The plugin has been modified to keep track of which API object type created the pod template of a given name. ([BZ#1709626](#))
- The **openshift_set_node_ip** variable was deprecated, but still included in inventory example files. This has now been removed from example files and code for the **openshift_set_node_ip** variable has been cleaned up. ([BZ#1712488](#))

- Previously, the web console could show an incorrect "Scaling to..." value for stateful sets in the project overview under some conditions. The stateful set desired replicas value now correctly updates in the web console project overview. ([BZ#1713211](#))
- Previously, a service would not correctly show up in the project overview when it selected the **DeploymentConfig** label that is automatically set for pods created by a deployment config. Now, it correctly show services that select the **DeploymentConfig** label on the overview. ([BZ#1717028](#))
- The cluster autoscaler did not have the **clusterrole** permission to evict pods and nodes would not be automatically scaled as a result. Now, eviction permissions have been added to the autoscaler cluster role. Pods can be evicted and nodes can be scaled down. ([BZ#1718458](#))
- If a pod using an egress IP tried to contact an external host that was not responding, the egress IP monitoring code may have mistakenly interpreted that as meaning that the node hosting the egress IP was not responding. High-availability egress IPs may have been switched from one node to another spuriously. The monitoring code now distinguishes the case of "egress node not responding" from "final destination not responding". High-availability egress IPs will not be switched between nodes unnecessarily. ([BZ#1718542](#))
- Refactoring of **openshift_facts** caused the MTU to be improperly set. Hosts could not communicate on networks with non-default MTU settings. The **openshift_facts.py** script was updated to properly detect and set the MTU for the host environment. Hosts now can properly communicate on networks with non-default MTU. ([BZ#1720581](#))

2.8.11.2. Enhancements

- The Cisco ACI CNI plugin is now available. ([BZ#1708552](#))
- You can now use an Ansible playbook to perform a certificate rotation for the EFK stack without needing to run the install/upgrade playbook. This playbook deletes the current certificate files, generates new EFK certificates, updates certificate secrets, and restarts Elasticsearch and Kibana. ([BZ#1710424](#))

2.8.11.3. Upgrading

To upgrade an existing OpenShift Container Platform 3.10 or 3.11 cluster to this latest release, see [Upgrade methods and strategies](#) for instructions.

2.8.12. RHSA-2019:1633 - Moderate: OpenShift Container Platform 3.11 atomic-openshift security update

Issued: 2019-06-27

An update for atomic-openshift is now available for OpenShift Container Platform 3.11. Details of the update are documented in the [RHSA-2019:1633](#) advisory.

2.8.12.1. Upgrading

To upgrade an existing OpenShift Container Platform 3.10 or 3.11 cluster to this latest release, use the automated upgrade playbook. See [Performing Automated In-place Cluster Upgrades](#) for instructions.

2.8.13. RHBA-2019:1753 - OpenShift Container Platform 3.11.129 Bug Fix Update

Issued: 2019-07-23

OpenShift Container Platform release 3.11.129 is now available. The list of packages and bug fixes included in the update are documented in the [RHBA-2019:1753](#) advisory. The container images included in the update are provided by the [RHBA-2019:1754](#) advisory.

2.8.13.1. Bug Fixes

- In OpenShift on Azure environments, conditional arguments were missing that would result in incorrect kubelet node names in certain cases. The conditionals to set **nodeName** in **node-config** were added, and now kubelet names can be set as required. ([BZ#1656983](#))
- Health check playbooks would assume **Curator** was a **deploymentconfig** instead of a **cronjob**, and would fail the check because the resource type had changed. Now, the health check playbook properly evaluates for a **cronjob** instead of a **deploymentconfig**. ([BZ#1676720](#))
- Some OpenShift Container Platform installations would fail because the **selinux** check was occurring in the **openshift_node** role instead of the **init** role. Now, the **selinux** check occurs earlier in the installation process and is completed successfully. ([BZ#1710020](#))
- Access to the **ElasticSearch** root URL was denied from a project's pod in OpenShift Container Platform 3.11 instances that had been upgraded from version 3.10. This was due to strenuous permissions that denied non-administrative users access to the root endpoints. Now, permissions have been changed so that all users are able to access the root endpoints. ([BZ#1710868](#))
- **ElasticSearch** metrics were unavailable in the **Prometheus** role. Now, the **Prometheus** role has been enabled access to monitor all **ElasticSearch** indices. ([BZ#1712423](#))
- **ImageStreams** would fail if not using a hosted managed registry due to an unset **referencePolicy** field. Now, the dictionary has been changed to read and modify the **referencePolicy** as needed, and **ImageStreams** can be used without a hosted managed registry. ([BZ#1712496](#))
- The **templateinstance** controller did not properly manage cluster level objects in its create path, and as a result failed to create projects specified in templates. Now, the **templateinstance** controller determines if the objects in its create path and passes correct values in secrets through namespaces. The **templateinstance** can now create projects as defined in templates. ([BZ#1713982](#))
- Redeployment of certificates did not recreate the **ansible-service-broker** pod secrets, causing the service catalog to fail. A new playbook has been created to support updating the certificates. ([BZ#1715322](#))
- The IPv4 dictionary was recently modified and MTU was set incorrectly as a result. This IPv4 conditional has been removed, and now MTU is established correctly. ([BZ#1719362](#))
- The **pom.xml** of some of the OpenShift Jenkins plugins had **http://** references instead of **https://** references for some of its build time dependencies, and dependency downloads would occur over **http** instead of the **https** protocol. The **pom.xml** references have now been corrected and dependency downloads only occur using the **https** protocol. ([BZ#1719477](#))
- The readiness probe for **ElasticSearch curl** commands used **NSS**, which bloated the dentry cache. This would cause **ElasticSearch** to become unresponsive. To resolve this, set the **NSS_SDB_USE_CACHE=no** flag in the readiness probe to work around the dentry cache bloating. ([BZ#1720479](#))

- Previously, the web console showed a misleading warning that metrics might not be configured for horizontal pod autoscalers when only the metrics server had been set up. The warning has been removed. ([BZ#1721428](#))
- Previously, the **image-signature-import** controller would only import up to three signatures, but the registry would often have more than three signatures. This would cause importing signatures to fail. The limit of signatures has been increased, and importing signatures from **registry.redhat.io** completes successfully. ([BZ#1722581](#))
- The **prerequisites** playbook would fail because default values were not loaded correctly, causing sanity checks to fail. A step to run **openshift_facts** has been added to load all the default values, and sanity checks complete successfully. ([BZ#1724718](#))
- **Kibana** would present a blank page or timeout if a large number of projects were creating too many calls to the **ElasticSearch** cluster, resulting in the timeout before a response is returned. Now, API calls are cached and processing is more efficient, reducing the opportunity for page timeouts. ([BZ#1726433](#))

2.8.13.2. Enhancements

- The service catalog did not have a **redeploy-certificate** playbook. The certificates for the service catalog need to be rotated like other components of OpenShift Container Platform, and a playbook has now been created for the service catalog. ([BZ#1702401](#))

2.8.13.3. Upgrading

To upgrade an existing OpenShift Container Platform 3.10 or 3.11 cluster to this latest release, see [Upgrade methods and strategies](#) for instructions.

2.8.14. RHBA-2019:2352 - OpenShift Container Platform 3.11.135 Bug Fix Update

Issued: 2019-08-13

OpenShift Container Platform release 3.11.135 is now available. The list of packages and bug fixes included in the update are documented in the [RHBA-2019:2352](#) advisory. The container images included in the update are provided by the [RHBA-2019:2353](#) advisory.

2.8.14.1. Upgrading

To upgrade an existing OpenShift Container Platform 3.10 or 3.11 cluster to this latest release, see [Upgrade methods and strategies](#) for instructions.

2.8.15. RHBA-2019:2581 - OpenShift Container Platform 3.11.141 Bug Fix Update

Issued: 2019-09-03

OpenShift Container Platform release 3.11.141 is now available. The list of packages and bug fixes included in the update are documented in the [RHBA-2019:2581](#) advisory. The container images included in the update are provided by the [RHBA-2019:2580](#) advisory.

2.8.15.1. Upgrading

To upgrade an existing OpenShift Container Platform 3.10 or 3.11 cluster to this latest release, see [Upgrade methods and strategies](#) for instructions.

2.8.16. RHBA-2019:2816 - OpenShift Container Platform 3.11.146 Bug Fix Update

Issued: 2019-09-23

OpenShift Container Platform release 3.11.146 is now available. The list of packages and bug fixes included in the update are documented in the [RHBA-2019:2816](#) advisory. The container images included in the update are provided by the [RHBA-2019:2824](#) advisory.

2.8.16.1. Upgrading

To upgrade an existing OpenShift Container Platform 3.10 or 3.11 cluster to this latest release, see [Upgrade methods and strategies](#) for instructions.

CHAPTER 3. XPAAS RELEASE NOTES

The release notes for xPaaS docs have migrated to their own book on the [Red Hat customer portal](#).

CHAPTER 4. COMPARING WITH OPENSHIFT ENTERPRISE 2

4.1. OVERVIEW

OpenShift Container Platform 3 is based on the OpenShift version 3 (v3) architecture, which is very different product than OpenShift version 2 (v2). Many of the same terms from OpenShift v2 are used in v3, and the same functions are performed, but the terminology can be different, and behind the scenes things may be happening very differently. Still, OpenShift remains an application platform.

This topic discusses these differences in detail, in an effort to help OpenShift users in the transition from OpenShift v2 to OpenShift v3.

4.2. ARCHITECTURE CHANGES

Gears Versus Containers

Gears were a core component of OpenShift v2. Technologies such as kernel namespaces, cGroups, and SELinux helped deliver a highly-scalable, secure, containerized application platform to OpenShift users. Gears themselves were a form of container technology.

OpenShift v3 takes the gears idea to the next level. It uses Docker as the next evolution of the v2 container technology. This container architecture is at the core of OpenShift v3.

Kubernetes

As applications in OpenShift v2 typically used multiple gears, applications on OpenShift v3 will expectedly use multiple containers. In OpenShift v2, gear orchestration, scheduling, and placement was handled by the OpenShift broker host. OpenShift v3 integrates Kubernetes into the master host to drive container orchestration.

4.3. APPLICATIONS

Applications are still the focal point of OpenShift. In OpenShift v2, an application was a single unit, consisting of one web framework of no more than one cartridge type. For example, an application could have one PHP and one MySQL, but it could not have one Ruby, one PHP, and two MySQLs. It also could not be a database cartridge, such as MySQL, by itself.

This limited scoping for applications meant that OpenShift performed seamless linking for all components within an application using environment variables. For example, every web framework knew how to connect to MySQL using the **OPENSHIFT_MYSQL_DB_HOST** and **OPENSHIFT_MYSQL_DB_PORT** variables. However, this linking was limited to within an application, and only worked within cartridges designed to work together. There was nothing to help link across application components, such as sharing a MySQL instance across two applications.

While most other PaaS limit themselves to web frameworks and rely on external services for other types of components, OpenShift v3 makes even more application topologies possible and manageable.

OpenShift v3 uses the term "application" as a concept that links services together. You can have as many components as you desire, contained and flexibly linked within a [project](#), and, optionally, labeled to provide grouping or structure. This updated model allows for a standalone MySQL instance, or one shared between JBoss components.

Flexible linking means you can link any two arbitrary components together. As long as one component can export environment variables and the second component can consume values from those

environment variables, and with potential variable name transformation, you can link together any two components without having to change the images they are based on. So, the best containerized implementation of your desired database and web framework can be consumed directly rather than you having to fork them both and rework them to be compatible.

This means you can build anything on OpenShift. And that is OpenShift's primary aim: to be a container-based platform that lets you build entire applications in a repeatable lifecycle.

4.4. CARTRIDGES VERSUS IMAGES

In OpenShift v3, an [image](#) has replaced OpenShift v2's concept of a cartridge.

Cartridges in OpenShift v2 were the focal point for building applications. Each cartridge provided the required libraries, source code, build mechanisms, connection logic, and routing logic along with a preconfigured environment to run the components of your applications.

However, cartridges came with disadvantages. With cartridges, there was no clear distinction between the developer content and the cartridge content, and you did not have ownership of the home directory on each gear of your application. Also, cartridges were not the best distribution mechanism for large binaries. While you could use external dependencies from within cartridges, doing so would lose the benefits of encapsulation.

From a packaging perspective, an image performs more tasks than a cartridge, and provides better encapsulation and flexibility. However, cartridges also included logic for building, deploying, and routing, which do not exist in images. In OpenShift v3, these additional needs are met by [Source-to-Image \(S2I\)](#) and [configuring the template](#).

Dependencies

In OpenShift v2, cartridge dependencies were defined with **Configure-Order** or **Requires** in a cartridge manifest. OpenShift v3 uses a declarative model where [pods](#) bring themselves in line with a predefined state. Explicit dependencies that are applied are done at runtime rather than just install time ordering.

For example, you might require another service to be available before you start. Such a dependency check is always applicable and not just when you create the two components. Thus, pushing dependency checks into runtime enables the system to stay healthy over time.

Collection

Whereas cartridges in OpenShift v2 were colocated within gears, [images](#) in OpenShift v3 are mapped 1:1 with [containers](#), which use [pods](#) as their colocation mechanism.

Source Code

In OpenShift v2, applications were required to have at least one web framework with one Git repository. In OpenShift v3, you can choose which images are built from source and that source can be located outside of OpenShift itself. Because the source is disconnected from the images, the choice of image and source are distinct operations with source being optional.

Build

In OpenShift v2, builds occurred in application gears. This meant downtime for non-scaled applications due to resource constraints. In v3, [builds](#) happen in separate containers. Also, OpenShift v2 build results used rsync to synchronize gears. In v3, build results are first committed as an immutable image and published to an internal registry. That image is then available to launch on any of the nodes in the cluster, or available to rollback to at a future date.

Routing

In OpenShift v2, you had to choose up front as to whether your application was scalable, and whether the routing layer for your application was enabled for high availability (HA). In OpenShift v3, [routes](#) are first-class objects that are HA-capable simply by scaling up your application component to two or more replicas. There is never a need to recreate your application or change its DNS entry.

The routes themselves are disconnected from images. Previously, cartridges defined a default set of routes and you could add additional aliases to your applications. With OpenShift v3, you can use templates to set up any number of routes for an image. These routes let you modify the scheme, host, and paths exposed as desired, with no distinction between system routes and user aliases.

4.5. BROKER VERSUS MASTER

A [master](#) in OpenShift v3 is similar to a broker host in OpenShift v2. However, the MongoDB and ActiveMQ layers used by the broker in OpenShift v2 are no longer necessary, because **etcd** is typically installed with each master host.

4.6. DOMAIN VERSUS PROJECT

A [project](#) is essentially a v2 domain.