Red Hat OpenStack Services on OpenShift 18.0

Release notes

Release notes for the Red Hat OpenStack Services on OpenShift 18.0 release

Last Updated: 2025-12-17

Red Hat OpenStack Services on OpenShift 18.0 Release notes

Release notes for the Red Hat OpenStack Services on OpenShift 18.0 release

Legal Notice

Abstract

The release notes provide high-level coverage of the improvements and additions that have been implemented in Red Hat OpenStack Services on OpenShift 18.0 and document known problems in this release, as well as notable bug fixes, technology previews, deprecated functionality, and other details.

Table of Contents

PREFACE 10

PROVIDING FEEDBACK ON RED HAT DOCUMENTATION 11

CHAPTER 1. NEW AND ENHANCED FEATURES 12
1.1. NEW AND ENHANCED FEATURES IN 18.0.14 (FR4) 12
1.1.1. Compute 12
1.1.2. Data plane 12
1.1.3. Documentation 12
1.1.4. High availability 13
1.1.5. Migration 13
1.1.6. Networking 13
1.1.7. Security 14
1.1.8. Storage 14
1.1.9. Upgrades and updates 15
1.2. NEW AND ENHANCED FEATURES IN 18.0.10 (FR3) 16
1.2.1. Bare Metal Provisioning 16
1.2.2. Compute 16
1.2.3. Control plane 16
1.2.4. Dashboard 17
1.2.5. Networking 17
1.2.6. NFV 18
1.2.7. Observability 18
1.2.8. Security 18
1.2.9. Storage 19
1.2.10. Upgrades and updates 19
1.2.11. Resource optimization 19
1.3. NEW AND ENHANCED FEATURES IN 18.0.6 (FR2) 20
1.3.1. Bare Metal Provisioning 20
1.3.2. Control plane 20
1.3.3. High availability 20
1.3.4. Networking 21
1.3.5. Observability 21
1.3.6. Security 21
1.3.7. Storage 22
1.3.8. Upgrades and updates 22
1.4. NEW AND ENHANCED FEATURES IN 18.0.3 (FR1) 22
1.4.1. Distributed Compute nodes (DCN) 23
1.4.2. Networking 23
1.4.3. Observability 24
1.4.4. Upgrades and updates 24
1.5. NEW AND ENHANCED FEATURES IN 18.0 (GA) 24
1.5.1. Control plane 24
1.5.2. Dashboard 24
1.5.3. Data plane 25
1.5.4. Documentation 25
1.5.4.1. Documentation in progress 26
1.5.4.2. RHOCP feature documentation 26
1.5.4.3. Earlier documentation versions 27
1.5.5. High availability 27
1.5.6. Networking 27

1.5.7. NFV 27
1.5.8. Observability 27
1.5.9. Security 28
1.5.10. Storage 28

CHAPTER 2. RELEASE INFORMATION RHOSO 18.0 30
2.1. RELEASE INFORMATION RHOSO 18.0.15 30
2.1.1. Advisory list 30
2.1.2. Observability 30
2.1.2.1. Technology Previews 30
2.1.2.2. Known issues 30
2.1.3. Compute 30
2.1.3.1. Known issues 31
2.1.4. Data plane 31
2.1.4.1. Bug fixes 31
2.1.5. Hardware Provisioning 32
2.1.5.1. Known issues 32
2.1.6. Networking 32
2.1.6.1. Bug fixes 32
2.1.6.2. Known issues 32
2.1.7. Control plane 33
2.1.7.1. Known issues 33
2.1.8. Storage 33
2.1.8.1. Known issues 33
2.1.9. Optimize Service 34
2.1.9.1. Known issues 34
2.2. RELEASE INFORMATION RHOSO 18.0.14 34
2.2.1. Advisory list 34
2.2.2. Compute 34
2.2.2.1. New features 34
2.2.2.2. Bug fixes 35
2.2.2.3. Known issues 35
2.2.3. Networking 36
2.2.3.1. New features 36
2.2.3.2. Bug fixes 37
2.2.3.3. Technology Previews 37
2.2.4. Network Functions Virtualization 38
2.2.4.1. New features 38
2.2.4.2. Bug fixes 39
2.2.5. Control plane 39
2.2.5.1. New features 39
2.2.5.2. Known issues 40
2.2.6. High availability 40
2.2.6.1. New features 40
2.2.7. Security and hardening 41
2.2.7.1. New features 41
2.2.8. Storage 41
2.2.8.1. New features 41
2.2.8.2. Bug fixes 43
2.2.8.3. Known issues 43
2.2.9. Dashboard 44
2.2.9.1. New features 44
2.2.10. Optimize Service 44

2.2.10.1. New features 44
2.2.10.2. Bug fixes 45
2.2.10.3. Known issues 45
2.2.11. Migration 46
2.2.11.1. New features 46
2.2.12. Hardware Provisioning 46
2.2.12.1. Known issues 46
2.3. RELEASE INFORMATION RHOSO 18.0.13 47
2.3.1. Advisory list 47
2.3.2. Compute 47
2.3.2.1. Known issues 47
2.3.3. Control plane 47
2.3.3.1. Known issues 47
2.3.4. Storage 48
2.3.4.1. Known issues 48
2.3.5. Upgrades and updates 48
2.3.5.1. Deprecated functionality 48
2.3.6. Optimize Service 48
2.3.6.1. Known issues 48
2.4. RELEASE INFORMATION RHOSO 18.0.12 49
2.4.1. Advisory list 49
2.4.2. Control plane 49
2.4.2.1. Known issues 49
2.4.3. Storage 49
2.4.3.1. Known issues 49
2.5. RELEASE INFORMATION RHOSO 18.0.11 50
2.5.1. Advisory list 50
2.5.2. Compute 50
2.5.2.1. Bug fixes 50
2.5.2.2. Known issues 51
2.5.3. Control plane 51
2.5.3.1. Bug fixes 51
2.5.3.2. Known issues 51
2.5.4. Storage 52
2.5.4.1. Bug fixes 52
2.5.4.2. Known issues 52
2.5.5. Optimize Service 52
2.5.5.1. Known issues 52
2.6. RELEASE INFORMATION RHOSO 18.0.10 52
2.6.1. Advisory list 52
2.6.2. Observability 53
2.6.2.1. New features 53
2.6.2.2. Bug fixes 53
2.6.3. Compute 53
2.6.3.1. New features 53
2.6.3.2. Known issues 54
2.6.4. Hardware Provisioning 54
2.6.4.1. Bug fixes 54
2.6.5. Networking 55
2.6.5.1. New features 55
2.6.5.2. Bug fixes 55
2.6.5.3. Technology Previews 55
2.6.5.4. Deprecated functionality 56

2.6.6. Network Functions Virtualization 56
2.6.6.1. New features 56
2.6.6.2. Bug fixes 57
2.6.7. Control plane 58
2.6.7.1. New features 58
2.6.7.2. Known issues 59
2.6.8. High availability 59
2.6.8.1. Bug fixes 59
2.6.9. Storage 60
2.6.9.1. New features 60
2.6.9.2. Bug fixes 61
2.6.9.3. Technology Previews 61
2.6.9.4. Known issues 61
2.6.10. Upgrades and updates 62
2.6.10.1. New features 62
2.6.11. Optimize Service 62
2.6.11.1. New features 62
2.6.11.2. Bug fixes 62
2.6.11.3. Technology Previews 62
2.6.11.4. Known issues 63
2.7. RELEASE INFORMATION RHOSO 18.0.9 63
2.7.1. Advisory list 63
2.7.2. Compute 63
2.7.2.1. Known issues 63
2.7.3. Data plane 64
2.7.3.1. Bug fixes 64
2.7.4. Networking 64
2.7.4.1. Bug fixes 64
2.7.4.2. Known issues 64
2.7.5. Network Functions Virtualization 65
2.7.5.1. Bug fixes 65
2.7.6. Control plane 66
2.7.6.1. Known issues 66
2.7.7. High availability 66
2.7.7.1. New features 66
2.7.8. Storage 66
2.7.8.1. Known issues 66
2.8. RELEASE INFORMATION RHOSO 18.0.8 67
2.8.1. Advisory list 67
2.8.2. Compute 67
2.8.2.1. Known issues 67
2.8.3. Data plane 68
2.8.3.1. Bug fixes 68
2.8.3.2. Known issues 68
2.8.4. Hardware Provisioning 68
2.8.4.1. Bug fixes 68
2.8.5. Networking 68
2.8.5.1. Bug fixes 69
2.8.5.2. Known issues 69
2.8.6. Network Functions Virtualization 70
2.8.6.1. Known issues 70
2.8.7. Control plane 71
2.8.7.1. Bug fixes 71

2.8.7.2. Known issues 71
2.8.8. Security and hardening 71
2.8.8.1. Bug fixes 71
2.9. RELEASE INFORMATION RHOSO 18.0.7 71
2.9.1. Advisory list 72
2.9.2. Compute 72
2.9.2.1. Bug fixes 72
2.9.2.2. Known issues 72
2.9.3. Data plane 73
2.9.3.1. Known issues 73
2.9.4. Networking 73
2.9.4.1. Bug fixes 73
2.9.4.2. Known issues 73
2.9.5. Network Functions Virtualization 75
2.9.5.1. Bug fixes 75
2.9.5.2. Deprecated functionality 76
2.9.5.3. Known issues 76
2.9.6. Control plane 77
2.9.6.1. Bug fixes 77
2.9.6.2. Known issues 77
2.9.7. Security and hardening 77
2.9.7.1. Known issues 77
2.9.8. Optimize Service 77
2.9.8.1. Technology Previews 77
2.10. RELEASE INFORMATION RHOSO 18.0.6 78
2.10.1. Advisory list 78
2.10.2. Observability 78
2.10.2.1. New features 78
2.10.2.2. Bug fixes 80
2.10.3. Compute 81
2.10.3.1. New features 81
2.10.3.2. Bug fixes 82
2.10.3.3. Known issues 83
2.10.4. Data plane 84
2.10.4.1. New features 84
2.10.4.2. Bug fixes 84
2.10.5. Networking 84
2.10.5.1. Bug fixes 84
2.10.5.2. Technology Previews 85
2.10.5.3. Known issues 86
2.10.6. Network Functions Virtualization 88
2.10.6.1. Bug fixes 88
2.10.6.2. Technology Previews 89
2.10.6.3. Deprecated functionality 89
2.10.6.4. Known issues 89
2.10.7. Control plane 90
2.10.7.1. New features 90
2.10.7.2. Known issues 91
2.10.8. Storage 91
2.10.8.1. New features 91
2.10.8.2. Bug fixes 91
2.10.8.3. Known issues 92
2.10.9. Upgrades and updates 92

2.10.9.1. Known issues 92
2.11. RELEASE INFORMATION RHOSO 18.0.4 93
2.11.1. Advisory list 93
2.11.2. Compute 93
2.11.2.1. Bug fixes 93
2.11.2.2. Known issues 93
2.11.3. Data plane 94
2.11.3.1. Bug fixes 94
2.11.4. Hardware Provisioning 95
2.11.4.1. Bug fixes 95
2.11.5. Networking 95
2.11.5.1. Bug fixes 95
2.11.5.2. Known issues 96
2.11.6. Network Functions Virtualization 98
2.11.6.1. Known issues 98
2.11.7. Control plane 98
2.11.7.1. Known issues 98
2.11.8. Security and hardening 99
2.11.8.1. Bug fixes 99
2.11.9. Storage 99
2.11.9.1. Known issues 99
2.12. RELEASE INFORMATION RHOSO 18.0.3 100
2.12.1. Advisory list 100
2.12.2. Observability 101
2.12.2.1. New features 101
2.12.2.2. Technology Previews 101
2.12.3. Compute 101
2.12.3.1. New features 101
2.12.3.2. Bug fixes 102
2.12.3.3. Known issues 102
2.12.4. Data plane 103
2.12.4.1. New features 103
2.12.5. Networking 103
2.12.5.1. New features 103
2.12.5.2. Technology Previews 104
2.12.5.3. Known issues 104
2.12.6. Network Functions Virtualization 106
2.12.6.1. Known issues 106
2.12.7. Control plane 106
2.12.7.1. Known issues 106
2.12.8. High availability 106
2.12.8.1. Technology Previews 106
2.12.8.2. Known issues 107
2.12.9. Storage 107
2.12.9.1. New features 107
2.12.9.2. Known issues 108
2.12.10. Upgrades and updates 109
2.12.10.1. New features 109
2.12.10.2. Bug fixes 110
2.13. RELEASE INFORMATION RHOSO 18.0.2 110
2.13.1. Advisory list 110
2.13.2. Compute 111
2.13.2.1. Bug fixes 111

2.13.2.2. Known issues 111
2.13.3. Data plane 112
2.13.3.1. Bug fixes 112
2.13.4. Networking 112
2.13.4.1. Bug fixes 112
2.13.4.2. Known issues 113
2.13.5. Network Functions Virtualization 113
2.13.5.1. Known issues 113
2.13.6. Storage 113
2.13.6.1. Known issues 113
2.14. RELEASE INFORMATION RHOSO 18.0.1 114
2.14.1. Advisory list 114
2.14.2. Compute 114
2.14.2.1. Known issues 114
2.14.3. Data plane 115
2.14.3.1. Bug fixes 115
2.14.3.2. Known issues 115
2.14.4. Networking 116
2.14.4.1. New features 116
2.14.4.2. Bug fixes 116
2.14.4.3. Known issues 116
2.14.5. Network Functions Virtualization 117
2.14.5.1. Bug fixes 117
2.14.5.2. Known issues 117
2.14.6. Storage 117
2.14.6.1. Bug fixes 117
2.15. RELEASE INFORMATION RHOSO 18.0 GA 117
2.15.1. Advisory list 118
2.15.2. Observability 118
2.15.2.1. New features 118
2.15.3. Compute 119
2.15.3.1. New features 119
2.15.3.2. Bug fixes 122
2.15.3.3. Technology Previews 123
2.15.3.4. Removed functionality 123
2.15.3.5. Known issues 124
2.15.4. Data plane 125
2.15.4.1. Known issues 125
2.15.5. Hardware Provisioning 126
2.15.5.1. Bug fixes 126
2.15.6. Networking 126
2.15.6.1. New features 126
2.15.6.2. Bug fixes 127
2.15.6.3. Technology Previews 127
2.15.6.4. Known issues 127
2.15.7. Network Functions Virtualization 128
2.15.7.1. New features 128
2.15.7.2. Bug fixes 128
2.15.7.3. Technology Previews 129
2.15.7.4. Deprecated functionality 129
2.15.7.5. Known issues 129
2.15.8. High availability 130
2.15.8.1. New features 130

2.15.9. Storage 130
2.15.9.1. New features 130
2.15.9.2. Bug fixes 131
2.15.9.3. Technology Previews 131
2.15.9.4. Known issues 131
2.15.10. Dashboard 132
2.15.10.1. New features 132
2.15.10.2. Bug fixes 133
2.16. RELEASE INFORMATION RHOSO 18.0 BETA 133
2.16.1. Advisory list 133
2.16.2. Compute 133
2.16.2.1. New features 133
2.16.2.2. Deprecated functionality 134
2.16.2.3. Known issues 134
2.16.3. Networking 135
2.16.3.1. Known issues 135
2.16.4. Network Functions Virtualization 135
2.16.4.1. Known issues 135
2.16.5. Storage 135
2.16.5.1. Known issues 135
2.16.6. Release delivery 136
2.16.6.1. Removed functionality 136
2.16.7. Integration test suite 136
2.16.7.1. Known issues 136

PREFACE

The release notes provide high-level coverage of the improvements and additions that have been implemented in Red Hat OpenStack Services on OpenShift 18.0 and document known problems in this release, as well as notable bug fixes, technology previews, deprecated functionality, and other details.

PROVIDING FEEDBACK ON RED HAT DOCUMENTATION

We appreciate your input on our documentation. Tell us how we can make it better.

Providing documentation feedback in Jira

Use the Create Issue form to provide feedback on the documentation. The Jira issue will be created in the Red Hat OpenStack Platform Jira project, where you can track the progress of your feedback.

1. Ensure that you are logged in to Jira. If you do not have a Jira account, create an account to submit feedback.
2. Click the following link to open a the **Create Issue** page: Create Issue
3. Complete the **Summary** and **Description** fields. In the **Description** field, include the documentation URL, chapter or section number, and a detailed description of the issue. Do not modify any other fields in the form.
4. Click **Create**.

CHAPTER 1. NEW AND ENHANCED FEATURES

Review features that have been added to or significantly enhanced in Red Hat OpenStack Services on OpenShift (RHOSO).

RHOSO improves substantially over previous versions of Red Hat OpenStack Platform (RHOSP). The RHOSO control plane is natively hosted on the Red Hat OpenShift Container Platform (RHOCP) and the external RHEL-based data plane and workloads are managed with Ansible. This shift in architecture aligns with Red Hat's platform infrastructure strategy. You can future proof your existing investments by using RHOCP as a hosting platform for all of your infrastructure services.

For information about mapping RHOSO versions to OpenStack Operators and OpenStackVersion Custom Resources (CRs), see the Red Hat knowledge base article at https://access.redhat.com/articles/7125383.

RHOSP 17.1 is the last version of the product to use the director-based OpenStack on OpenStack form-factor for the control plane.

1.1. NEW AND ENHANCED FEATURES IN 18.0.14 (FR4)

Review features that have been added to or significantly enhanced in Red Hat OpenStack Services on OpenShift.

1.1.1. Compute

Erase all data on an NVMe device by using the NVMe cleanup agent

You can now deploy and configure the NVMe cleanup agent on data plane nodes to securely erase all data on the NVMe device before it is reallocated to the next instance.

1.1.2. Data plane

Deploy data plane nodes with Image Mode (bootc) images

Deploying data plane nodes using Image Mode (bootc) images is provided as a Technology Preview feature in this release. Technology Preview features are not fully supported by Red Hat. It should only be used for testing, and should not be deployed in a production environment.

1.1.3. Documentation

Documentation library restructured

The documentation library page was restructured to align better with the user life cycle and top-level jobs. The new structure includes the following enhancements:

- Validated Architectures were moved to a new category "Deploy a Validated Architecture environment".
- Several guides were moved to different categories to align with the actual user life cycle needs.
- The guide *Deploying RHOSO at scale* was renamed to *Planning a large-scale RHOSO deployment*.
- The new guide *Migrating VMs to a Red Hat OpenStack Services on Openshift deployment* was added.

1.1.4. High availability

Configuring authentication for the memcached service

Starting with RHOSO 18.0.14 (Feature Release 4), you can configure the cache maintained by the memcached service to require authentication to increase the security of your cloud by restricting access to the cached data of your cloud. For more information, see Configuring authentication for the memcached service in *Customizing the Red Hat OpenStack Services on OpenShift deployment* .

Configuring quorum queues for RabbitMQ in new deployments

Starting with RHOSO 18.0.14 (Feature Release 4), RabbitMQ supports the use of the **Quorum** queues for new RHOSO deployments. A **Quorum** queue is a durable, replicated queue based on the Raft consensus algorithm, providing increased data safety and high availability. For more information, see step 5 of Creating the control plane in *Deploying Red Hat OpenStack Services on OpenShift* .

1.1.5. Migration

Migrate VMs with VMware Migration Toolkit

In RHOSO 18.0.14 (Feature Release 4) and RHOSP 17.1, you can now migrate workloads from VMware to OpenStack using the VMware Migration Toolkit.

1.1.6. Networking

Observability metrics expanded from data plane nodes to data plane and control plane pods

The Prometheus OVS/OVN Exporter was previously available only on data plane nodes. Starting with RHOSO 18.0.14 (Feature Release 4), Prometheus OVS/OVN Exporter is also available on control plane pods. New groups of metrics are also included. For more information, see Network observability in *Managing networking resources*.

Firewall-as-a-Service (FWaaS) Technology Preview

In RHOSO 18.0.14 (Feature Release 4), you can test a Technology Preview of Firewall-as-a-Service (FWaaS). Do not use technology preview features in production environments. As more OpenStack-based clouds are adopted for multi-tenant applications, security remains a top priority. Network-level isolation and traffic control become critical, especially in public or hybrid cloud environments. Although security groups provide sufficient capability to specify security policy at a VM instance level or VM port level, it does not have support to specify policy at a network or router port level. FWaaS project provides this additional capability to specify the security policies at the router port level and enables specifying multiple policy rules within the same policy group and also supports application of L3 or L2 policy at the router port level. With the FWaaS Technology Preview, you can also test NGFW 3rd party plugins for integration with NGFW vendor solutions enabling firewall capabilities beyond the ACL level, including capabilities such as DPI, Malware protection, IPS, and IDP.

TAP-as-a-Service (TAPaas) Technology Preview

In RHOSO 18.0.14 (Feature Release 4), you can test a Technology Preview of TAP-as-a-Service (TAPaaS). Do not use technology preview features in production environments. As modern cloud infrastructure becomes increasingly complex and multi-tenant, observability and security monitoring have become foundational requirements for OpenStack operators. One key network diagnostic technique used in traditional and virtualized environments is port mirroring, which allows administrators to capture and analyze traffic flowing through a particular interface. Mirrored traffic can be re-directed to third party analytics tools and solutions hosted on a different or same host as the mirror port. Typically, the mirrored traffic is carried over overlay tunnels established between the source and destination of the mirror.

You can perform the following tasks with port mirroring:

- Security monitoring: Capture mirrored traffic for inspection by IDS/IPS tools.
- Performance analysis: Monitor bottlenecks, latency, and packet loss in real-time.
- Troubleshooting: Debug issues without logging into tenant VMs or affecting production traffic.
- Compliance auditing: Log and analyze data flows for regulatory purposes.
- Lawful intercept: In jurisdictions that require service providers to support legal requests for targeted surveillance, TAPaaS offers a programmable, isolated way to mirror traffic for specific endpoints without impacting other tenants or violating privacy constraints.

Port mirroring is available at OVS and OVN levels through a CLI interface, however, in highly dynamic, software-defined environments like OpenStack, traditional port mirroring does not scale well and does not offer the tenant-level abstraction and isolation. TAPaaS provides a Openstack integrated framework for scalable port mirroring in a multi-tenant shared environment maintaining the tenant isolation boundaries in Openstack deployments. TAPaaS is a Neutron extension that enables on-demand traffic mirroring for tenant or administrator purposes. It allows users to create TAP services that mirror traffic from one or more Neutron ports and redirect it to a TAP destination—often a virtual Network Packet Broker (NPB), intrusion detection system (IDS), or traffic analyzer instance.

Load-balancing service (Octavia) support for DCN deployments

In RHOSO 18.0.14, creating load balancers in availability zones (AZs) are now fully supported. For more information, see Creating availability zones for load balancing of network traffic at the edge in *Configuring load balancing as a service* .

DNS service (designate)

In RHOSO 18.0.14, the DNS service (designate) is now fully supported. For more information, see Configuring DNS as a service .

Dynamic routing with BGP support for IPv6 networks

In RHOSO 18.0.14, you can configure your dynamic routing environment using IPv6 networks. For more information, see Preparing RHOCP for BGP networks on RHOSO .

Avoiding taskflow interruptions by using flow resumption

In RHOSO 18.0.14, you can use Load-balancing service (octavia) flow resumption, which automatically reassigns the flow to an alternate controller if the original controller shuts down unexpectedly. For more information, see Avoiding taskflow interruptions by using flow resumption .

OVN provider driver for Load-balancing service (octavia) is now fully supported

In RHOSO 18.0.14, the OVN provider driver for the Load-balancing service is no longer a Technology Preview and is now fully supported. For more information, see Load-balancing service provider drivers .

1.1.7. Security

Multi-realm federation support

Starting with RHOSO 18.0.14 (Feature Release 4), you can configure RHOSO to allow users to log in to the OpenStack Dashboard by using single sign-on (SSO) and select from one of several external Identity Providers (IdPs). For more information, see Configuring multi-realm federated authentication in *Configuring security services*.

1.1.8. Storage

Notifications for events in the Block Storage service and Shared File Systems service

In RHOSO 18.0.14 (Feature Release 4), you can enable notifications in the Block Storage service (cinder) and Shared File System service (manila) by using the **notificationsBusInstance** parameter, allowing integration with either the existing RabbitMQ instance or a dedicated RabbitMQ instance.

Deployment of Object Storage service on data plane nodes

In RHOSO 18.0.14 (Feature Release 4), you can deploy the Object Storage service (swift) on external data plane nodes, improving scalability and performance for large storage clusters. By enabling DNS forwarding and creating an **OpenStackDataPlaneNodeSet** CR with specified properties, including disks for storage, you can customize the service configuration through additional **ConfigMap** or **Secret** CRs in the **OpenStackDataPlaneService** CR.

Shared File Systems service now supports transferring shares between tenants

In RHOSO 18.0.14 (Feature Release 4), the Shared File Systems service (manila) now supports transferring shares across projects. To ensure security and non-repudiation, a one-time transfer secret key is generated when you initiate a transfer. The key must be conveyed out-of-band so that a user in the recipient project can complete the transfer.

1.1.9. Upgrades and updates

Prevent minor update from proceeding when the custom container images have not been updated

This enhancement ensures correct version tracking and validation during minor updates by preventing the side effects and inconsistencies that result from custom container images not being updated when the target version is updated.
With this update, when a minor update is initiated by setting the **targetVersion**, the performance of the minor update is halted if the **customImages** version for the associated custom container images is not also updated. Users have the option to force the update if necessary.

Adopt RHOSP 17.1 Instance HA environments to RHOSO

Starting with RHOSO 18.0.14 (Feature Release 4), you can adopt Red Hat OpenStack Platform (RHOSP) 17.1 environments with Instance HA enabled to RHOSO 18.0. For more information about adopting Instance HA environments, see Preparing an Instance HA deployment for adoption and Enabling the high availability for Compute instances service in *Adopting a Red Hat OpenStack Platform 17.1 deployment*.

Shared File Systems service (manila) with CephFS through NFS adoption is fully supported

Adopting the Shared File Systems service (manila) with CephFS through NFS is now generally available. Previously, these adoption instructions were provided as a Technology Preview.
This enhancement allows you to migrate your existing Red Hat OpenStack Platform 17.1 deployment that uses CephFS through NFS as a back end for the Shared File Systems Service to RHOSO 18.0 with full support.

The adoption process includes:

- Creating a new clustered NFS Ganesha service managed directly on the Red Hat Ceph cluster
- Migrating export locations from the standalone Pacemaker-controlled **ceph-nfs** service to the new clustered service
- Decommissioning the previous standalone NFS service
  For more information, see Changes to CephFS through NFS and Creating an NFS Ganesha cluster in *Adopting a Red Hat OpenStack Platform 17.1 deployment* .

Full support for adopting environments that use iSCSI back ends for the Block Storage service (cinder)

Starting with RHOSO 18.0.14 (Feature Release 4), the procedure to adopt RHOSO 18.0 is fully supported for Red Hat OpenStack Platform 17.1 environments that use iSCSI as a back end for the Block Storage service (cinder). For more information, see Adopting the Block Storage service in *Adopting a Red Hat OpenStack Platform 17.1 deployment* .

Full support for adopting environments that use Block Storage service (cinder) back ends for the Image service (glance)

Starting with RHOSO 18.0.14 (Feature Release 4), RHOSO 18.0 adoption is fully supported for Red Hat OpenStack Platform 17.1 environments that use Block Storage service (cinder) as a back end for the Image service (glance). For more information, see Adopting the Image service that is deployed with a Block Storage service back end in *Adopting a Red Hat OpenStack Platform 17.1 deployment* .

1.2. NEW AND ENHANCED FEATURES IN 18.0.10 (FR3)

Review features that have been added to or significantly enhanced in Red Hat OpenStack Services on OpenShift.

1.2.1. Bare Metal Provisioning

Layer 2 network configuration using Networking Generic Switch in the Bare Metal Provisioning service (Technology Preview)

RHOSO 18.0.10 (Feature Release 3) introduces support for the configuration of L2 networks on non-provisioning NIC interfaces when using Baremetal as a Service (BMaaS) through the Bare Metal Provisioning service (ironic). This feature allows network configuration on switches by leveraging the networking-generic-switch Modular Layer 2 Neutron Mechanism driver.

1.2.2. Compute

PCI device tracking in the Placement service is now generally available

Previously, this feature was available as Technology Preview. You can use the Placement service to observe the PCI resource availability and usage across the whole cloud through the Placement API. The administrator can also reserve PCI devices for maintenance through the Placement API.

Configuration of notifications to the Telemetry service

Starting with RHOSO 18.0.10 (Feature Release 3), you can configure the Compute service (nova) to provide notifications to Telemetry services in your RHOSO environment.

Setting the maximum number of vGPUs that an SR-IOV NVIDIA GPU can create

Starting with RHOSO 18.0.10 (Feature Release 3), you can define the maximum number of vGPUs that a SR-IOV NVIDIA GPU can create.

Reserving One Time Use devices

Starting with RHOSO 18.0.10 (Feature Release 3), you can tag PCI devices as One Time Use (OTU) to reserve them for a single use by a single instance.

1.2.3. Control plane

Multiple RHOSO deployments on a single RHOCP cluster by using namespace separation

Starting with RHOSO 18.0.10 (Feature Release 3), you can deploy multiple RHOSO environments on a single RHOCP cluster by using namespace (project) isolation.

NOTE

Do not deploy multiple RHOSO environments on a single cluster with namespace separation in production. Multiple deployments are suitable only for development, staging, and testing environments.

Documentation: Guidance for deploying Red Hat OpenStack Services on OpenShift in a disconnected environment

RHOSO 18.0.10 (Feature Release 3) introduces documentation support for deploying Red Hat OpenStack Services on OpenShift (RHOSO) in a disconnected environment. For more information, see Deploying Red Hat OpenStack Services on OpenShift in a disconnected environment.

1.2.4. Dashboard

The horizon-operator creates an additional sidecar container for logging

Starting with RHOSO 18.0.10 (Feature Release 3), the Dashboard service **horizon-operator** implements a separate sidecar container to ensure the availability of logs for debugging. If you use a custom container image, you might need to rebuild your custom image when updating.

1.2.5. Networking

DNS service (designate) (Technology Preview)

With this technology preview, you can test the management of DNS records, names, and zones using the DNS service (designate). For more information, see *Configuring DNS as a service* .

Vertical scaling for load-balancing service (Octavia) instances (amphorae)

Starting with RHOSO 18.0.10 (Feature Release 3), RHOSO supports vertical scaling for load-balancing service instances. Users can scale-up their load balancers, increasing the CPU and RAM of the load-balancing instance, to improve performance and capacity. Vertically scaling increases the volume of network traffic processed. To scale-up a load balancer, use the appropriate load-balancing flavor when you create a load balancer. RHOSO ships with **amphora-4vcpus**, which creates an instance that contains 4 vCPUs, 4GB RAM, and 3GB of disk space. Your RHOSO administrator can create other custom load-balancing flavors that meet the load-balancing needs of your particular environment. For more information, see Creating Load-balancing service flavors in *Configuring load balancing as a service* .

Load-balancing service (Octavia) support for DCN deployments (Technology Preview)

With this technology preview, you can create load balancers in a distributed compute node (DCN) environment to increase traffic throughput and reduce latency. For more information, see Creating availability zones for load balancing of network traffic at the edge in *Configuring load balancing as a service*.

Load-balancing service (Octavia) TLS client authentication

Starting with RHOSO 18.0.10 (Feature Release 3), you can secure your web client communication with a load balancer by using two-way TLS authentication. For more information, see Creating a TLS-terminated HTTPS load balancer with client authentication in *Configuring load balancing as a service*.

BGP-EVPN support for provider network workloads without FDP support (Developer Preview)

Starting with RHOSO 18.0.10 (Feature Release 3), you can test a developer preview of BGP-EVPN support for provider network workloads without FDP support. Openstack provides a mature infrastructure platform for virtualized workload focussing on on-prem environments. With most of the Telco 4G workloads running on virtualized platforms and the expanding landscape of multiple sites and clusters, there is an imperative need for connectivity across the clusters that enables tenant workload deployment across multiple clusters. In addition to providing control plane and data

plane isolation in a shared environment, there is a need for multi-tenancy extending to the compute nodes. RHOSO 18 FR3 adds support for BGP-EVPN enabling multi-tenant, multi-VRF support with overlapping IP addresses for provider network workloads. The feature is available as developer preview in RHOSO 18 FR3 and is suitable for functional operation and testing in lab environments only.

Prometheus Exporter for OVN logical routers and logical switches

Starting with RHOSO 18.0.10 (Feature Release 3), you can use Prometheus Exporter for OVN logical routers and logical switches. Network observability requires metrics and KPIs to be available at the OVN layer, exposing packet statistics within the networking infrastructure orchestrated by OVN. RHOSO 18 FR3 adds support for monitoring metrics at the OVN layer (logical routers and switches) via prometheus exporter, allowing correlation between the top-level content management system (CMS), logical OVN, and physical representations of networking elements.

New OVN database synchronization tool to fix OVN load balancers

RHOSO 18.0.10 (Feature Release 3) introduces an OVN database synchronization tool to fix OVN load balancers that experience problems. The new tool, **octavia-ovn-db-sync-util**, is run on the command-line to synchronize the state of Load-balancing service (octavia) resources, with the OVN databases. For more information, see https://docs.redhat.com/en/documentation/red_hat_openstack_services_on_openshift/18.0/html/con maintain-lb-service_rhoso-lbaas#synch-lbs-ovn-provider_trbls-lbs

1.2.6. NFV

OVS-DPDK is now supported for all workloads

RHOSO 18.0.10 (Feature Release 3) introduces support for OVS-DPDK for all workloads. Previously, OVS-DPDK was only supported in NFV workloads.

TCP segmentation offloading (TSO) for OVS-DPDK is now generally available

Previously, TSO for OVS-DPDK was available as a technology preview. Now it is generally available. TSO offloads segmentation to NICs, freeing up host CPU resources and improving overall performance.

OVS-DPDK on networker nodes for acceleration of gateway traffic

RHOSO 18.0.10 (Feature Release 3) introduces support of DPDK-accelerated Open vSwitch (OVS-DPDK) on Networker nodes. The DPDK datapath provides lower latency and higher performance than the standard kernel OVS datapath. OVS-DPDK is a high-performance, user-space solution that bundles Open vSwitch with the Data Plane Development Kit (DPDK). This technology is designed to process packets quickly by running mostly in the user-space, allowing applications to directly handle packet processing to or from the Network Interface Card (NIC).

1.2.7. Observability

Database and Compute metrics available to Prometheus for telemetry data collection and storage

Starting with RHOSO 18.0.10 (Feature Release 3), the Telemetry service collects both database and Compute metrics and makes them available to Prometheus, enabling database telemetry and Compute node telemetry to be stored in the telemetry storage system.

1.2.8. Security

LDAP Support

RHOSO 18.0.10 (Feature Release 3) introduces support for connecting the Identity service (keystone) to LDAP for authentication.

Proteccio HSM support

In RHOSO 18.0.10 (Feature Release 3), the Key Manager service (barbican) supports the Proteccio HSM as a back end to store secrets.

1.2.9. Storage

Distributed zones with third-party storage

RHOSO 18.0.10 (Feature Release 3) introduces support for using certified third-party storage with distributed zones.

Adopting the Image service (glance) with an NFS storage back end

RHOSO 18.0.10 (Feature Release 3) introduces support for adopting the Image service from Red Hat OpenStack Platform (RHOSP) 17.1 with an NFS back end.

Improved parallel image upload performance with load distribution

Starting with RHOSO 18.0.10 (Feature Release 3), you can improve parallel image upload performance by using the **mod_wsgi** package to distribute the load across workers.

Image service (glance) notifications for events in image lifecycle

Starting with RHOSO 18.0.10 (Feature Release 3), you can enable notifications in the Image service by using the **notificationBusInstance** parameter, allowing integration with either the existing RabbitMQ instance or a dedicated one.

Adopting the Block Storage service (cinder) with an NFS storage back end

RHOSO 18.0.10 (Feature Release 3) introduces support for adopting the Block Storage service from Red Hat OpenStack Platform (RHOSP) 17.1 with an NFS back end.

Remote ring storage supports larger deployments of the Object Storage service (Technology Preview)

Starting with RHOSO 18.0.10 (Feature Release 3), you can use remotely stored rings to manage larger deployments of the Object storage service (swift).

CephFS file name added to CephFS share metadata

Starting with RHOSO 18.0.10 (Feature Release 3), you can check a CephFS file name when mounting a native CephFS share by viewing the **mount_options** metadata of the share. Starting with RHOSO 18.0.10 (Feature Release 3), you can check a CephFS file name when mounting a native CephFS share by viewing the **mount_options** metadata of the share.

Adopting the Shared File Systems service (manila) with a third-party back end

RHOSO 18.0.10 (Feature Release 3) introduces support for adopting the Shared File Systems service from Red Hat OpenStack Platform (RHOSP) 17.1 with a third-party back end, for example, NetApp or Dell.

1.2.10. Upgrades and updates

Granular package update workflow for RHOSO Compute nodes during the RHOSO update process (Technology Preview)

RHOSO 18.0.10 (Feature Release 3) introduces a mechanism to break down the update process for RHOSO Compute nodes running RHEL 9.4 into two distinct phases: updating OpenStack-related RPM packages and updating system-related RPM packages. By enabling this separation, operators gain finer control over the update process, reducing risks and simplifying troubleshooting in the event of issues.

1.2.11. Resource optimization

Optimize service (watcher) strategies for resource optimization (Technology Preview)

RHOSO 18.0.10 (Feature Release 3) introduces support for three new supported strategies in the

Optimize service: host maintenance, zone migration for instances, and workload balance. For more information about supported strategies to achieve resource optimization goals, see Sample Optimize service workflows in *Optimizing infrastructure resource utilization*.

1.3. NEW AND ENHANCED FEATURES IN 18.0.6 (FR2)

Review features that have been added to or significantly enhanced in Red Hat OpenStack Services on OpenShift.

1.3.1. Bare Metal Provisioning

RHOSO environment with a routed spine-leaf network topology

RHOSO 18.0.6 (Feature Release 2) introduces support for deploying a RHOSO environment with a routed spine-leaf network topology. For more information, see *Deploying a RHOSO environment with a routed spine-leaf network topology*.

1.3.2. Control plane

Streamlined RHOSO service Operators installation and initialization

RHOSO 18.0.6 (Feature Release 2) introduces a new initialization resource that streamlines the management of the RHOSO service Operators under a single Operator Lifecycle Manager (OLM) bundle. After you install the OpenStack Operator and before creating the control plane, you now create the new **OpenStack** initialization resource, which installs all the RHOSO service Operators.

Distributed zones

RHOSO 18.0.6 (Feature Release 2) introduces support for deploying a distributed control plane across multiple RHOCP cluster nodes that are located in distributed low latency L3 connected data centers.

Custom environment variables for the OpenStackClient pod

Starting with RHOSO 18.0.6 (Feature Release 2), you can customize the **OpenStackClient** pod environment variables to set the API version to use when connecting to the service API endpoints with the **openstackclient** CLI.

Multiple RHOSO deployments on a single RHOCP cluster using namespace separation (Technology Preview)

Starting with RHOSO 18.0.6 (Feature Release 2), you can test a technology preview of using namespace separation to deploy multiple RHOSO environments on a single RHOCP cluster. To deploy each RHOSO environment, create multiple isolated namespaces, then use the procedures in *Deploying Red Hat OpenStack Services on OpenShift*.

NOTE

Ensure that the NMState Operator on each host worker node provides for the multiple VLANs that are required to enable network isolation for each namespace.

1.3.3. High availability

Instance high availability

Starting with RHOSO 18.0.6 (Feature Release 2), you can use instance high availability (instance HA) to automatically evacuate and re-create instances on a different Compute node if a Compute node fails.

1.3.4. Networking

DNS service (Technology Preview)

Starting with RHOSO 18.0.6 (Feature Release 2), you can test a technology preview of the RHOSO DNS service (designate), a multi-tenant service that enables you to manage DNS records, names, and zones.

OVS-DPDK on networker nodes for OVN gateway acceleration

Starting with RHOSO 18.0.6 (Feature Release 2), you can enable OVS-DPDK on networker nodes for improved forwarding performance.

Support for nmstate provider in new greenfield deployments

Starting with RHOSO 18.0.6 (Feature Release 2), the **nmstate** provider is supported for new RHOSO deployments. The default **os-net-config** provider for new (greenfield) RHOSO deployments is **ifcfg**. For limitations and other details, see https://issues.redhat.com/browse/OSPRH-11309.

TCP segmentation offload for RHOSO environments with OVS-DPDK (Technology Preview)

Starting with RHOSO 18.0.6 (Feature Release 2), you can test a technology preview of TCP segmentation offload (TSO) for RHOSO environments with OVS-DPDK. For details, see *OVS-DPDK with TCP segmentation offload (Technology Preview)*.

1.3.5. Observability

Power consumption monitoring

Starting with RHOSO 18.0.6 (Feature Release 2), the visualization of IPMI power metrics is available in the dashboard. For more information, see https://issues.redhat.com/browse/OSPRH-10808.

Enhanced OpenStack Observability

Starting with RHOSO 18.0.6 (Feature Release 2), you can use the **openstack-network-exporter** to expose metrics from OVS or OVS-DPDK, OVN, and DPDK (PMD), and a dashboard has been added for these metrics.

Container health check

Starting with RHOSO 18.0.6 (Feature Release 2), you can use new metrics for monitoring the health of RHOSO services, including the following:

- **kube_pod_status_phase**
- **kube_pod_status_ready**
- **node_systemd_unit_state**
- **podman_container_state**
- **podman_container_health**
  You can use **kube_pod_status_phase** and **kube_pod_status_ready** to monitor control plane services. For more information, see https://issues.redhat.com/browse/OSPRH-1052.

1.3.6. Security

Key Manager (barbican) support for Luna

Starting with RHOSO 18.0.6 (Feature Release 2), when you install RHOSO, you have the option of using it with a Luna hardware security module (HSM). Using a hardware security module provides hardened protection for storing keys.

Identity service (keystone) support for Federation

RHOSO 18.0.6 (Feature Release 2) introduces Red Hat support for Red Hat Single Sign-On (RH-SSO) or Active Directory Federation Services (ADFS) as identity providers for RHOSO.

1.3.7. Storage

Integration with external Red Hat Ceph Storage (RHCS) clusters

Starting with RHOSO 18.0.6 (Feature Release 2), you can integrate RHOSO with external Red Hat Ceph Storage 8 clusters (as well as Red Hat Ceph Storage 7 clusters) to include Red Hat Ceph Storage capabilities with your deployment. Due to known issues, not all Red Hat Ceph Storage 8 functionality is supported. For more information about these issues, see the Known Issues section.

Image service (glance) support for S3 back end

Starting with RHOSO 18.0.6 (Feature Release 2), you can configure the Image service with an S3 back end.

1.3.8. Upgrades and updates

IMPORTANT

OpenStack Operator 18.0.6 now requires you to install a new initialization resource called **OpenStack**. You must create this resource when you update your RHOSO deployment from a version older than 18.0.6, or when you perform a new installation of 18.0.6. Also, if you deployed RHOSO 18.0.4 or earlier on RHOCP 4.16, you must create the OpenStack initialization resource before upgrading your RHOCP cluster to RHOCP 4.18.

RHOSO environments installed earlier than 18.0.6 have individual Operators, such as **horizon-operator**, **nova-operator**, and so on, in the **openstack-operators** namespace. Creation of the OpenStack resource automatically cleans up these unnecessary resources in the OpenShift environment. For more information about creating the OpenStack resource, see Installing the OpenStack Operator in *Deploying Red Hat OpenStack Services on OpenShift*.

Baremetal as a service (ironic) adoption from RHOSP 17.1 to RHOSO 18.0 (Technology Preview)

RHOSO 18.0.6 (Feature Release 2) introduces a technology preview of the ability to adopt Baremetal as a service (ironic) from RHOSP 17.1 to RHOSO 18.0. For details, see Adopting the Bare Metal Provisioning service in *Adopting a Red Hat OpenStack Platform 17.1 deployment* .

IPv6 stack adoption from RHOSP 17.1 to RHOSO 18.0 (Technology Preview)

Starting with RHOSO 18.0.6 (Feature Release 2), you can test a technology preview of configuring IPv6 networking for adoption. For more information, see *Adopting a Red Hat OpenStack Platform 17.1 deployment*.

Kernel live patching for RHOSO environments (Technology Preview)

Starting with RHOSO 18.0.6 (Feature Release 2), you can test a technology preview of kernel live patching support for RHOSO environments. With this feature, you can apply critical security updates and bug fixes to the kernel without requiring a system reboot. You cannot use this feature to apply custom live patches or third-party live patching solutions.

1.4. NEW AND ENHANCED FEATURES IN 18.0.3 (FR1)

Review features that have been added to or significantly enhanced in Red Hat OpenStack Services on OpenShift.

1.4.1. Distributed Compute nodes (DCN)

DCN with Red Hat Ceph storage

RHOSO 18.0.3 (Feature Release 1) introduces support for Distributed Compute Nodes (DCN) with persistent storage backed by Red Hat Ceph Storage.

1.4.2. Networking

Dynamic routing on data plane with FRR and BGP

RHOSO 18.0.3 (Feature Release 1) introduces support of Free Range Routing (FRR) border gateway protocol (BGP) to provide dynamic routing capabilities on the RHOSO data plane.

Limitations:

- If you use dynamic routing, you must also use distributed virtual routing (DVR).
- If you use dynamic routing, you also use dedicated networker nodes.
- You can not use dynamic routing in an IPv6 deployment or a deployment that uses the Load-balancing service (octavia).

Custom ML2 mechanism driver and SDN backend (Technology Preview)

RHOSO 18.0.3 (Feature Release 1) allows you to test integration of the Networking service (neutron) with a custom ML2 mechanism driver and software defined networking (SDN) back end components, instead of the default OVN mechanism driver and back end components. Do not use this feature in a production environment.

IPv6 metadata

RHOSO 18.0.3 (Feature Release 1) introduces support of the IPv6 metadata service.

NMstate provider for os-net-config (Development Preview)

RHOSO 18.0.3 (Feature Release 1) allows you to test a Development Preview of the NMstate provider for **os-net-config**. To test the NMstate provider, set **edpm_network_config_nmstate: true**. Do NOT use this Development Preview setting in a production environment.

Forwarding database (FDB) learning and aging controls

RHOSO 18.0.3 (Feature Release 1) introduces FDB learning and related FDB aging parameters. You can use FDB learning to prevent traffic flooding on ports that have port security disabled. Set **localnet_learn_fdb** to **true**.

Use the **fdb_age_threshold** parameter to set the maximum time (seconds) that the learned MACs stay in the FDB table. Use the **fdb_removal_limit** parameter to prevent OVN from removing a large number of FDB table entries at the same time.

Example configuration

```
apiVersion: core.openstack.org/v1beta1
kind: OpenStackControlPlane
metadata:
  name: unused
spec:
  neutron:
    template:
      customServiceConfig: |
        [ovn]
```

```
localnet_learn_fdb = true
fdb_age_threshold = 300
[ovn_nb_global]
fdb_removal_limit = 50
```

1.4.3. Observability

Power consumption monitoring (Technology Preview)

RHOSO 18.0.3 (Feature Release 1) introduces technology previews of power consumption monitoring capability for VM instances and virtual networking functions (VNFs).
See Jira Issue OSPRH-10006: Kepler Power Monitoring Metrics Visualization in RHOSO (Tech Preview) and Jira Issue OSPRH-46549: As a service provider I need a comprehensive dashboard that provides a power consumption matrix per VNF(Tech Preview).

RabbitMQ metrics dashboard

Starting in RHOSO 18.0.3 (Feature Release 1), RabbitMQ metrics are collected and stored in Prometheus. A new dashboard for displaying these metrics was added.

1.4.4. Upgrades and updates

Adoption from RHOSP 17.1

RHOSO 18.0.3 (Feature Release 1) introduces the ability to use the adoption mechanism to upgrade from RHOSP 17.1 to RHOSO 18.0 while minimizing impacts to your workloads.

1.5. NEW AND ENHANCED FEATURES IN 18.0 (GA)

Review features that have been added to or significantly enhanced in Red Hat OpenStack Services on OpenShift.

1.5.1. Control plane

Control plane deployed on Red Hat OpenShift Container Platform (RHOCP)

In RHOSO 18.0 (GA), the director-based undercloud is replaced by a control plane that is natively hosted on an RHOCP cluster and managed with the OpenStack Operator. The Red Hat OpenStack Services on OpenShift (RHOSO) control plane features include:

- Deployed in pods and governed by Kubernetes Operators.
- Deploys in minutes, consuming only a fraction of the CPU and RAM footprint required by earlier RHOSP releases.
- Takes advantage of native Kubernetes mechanisms for high availability.
- Features built-in monitoring based on RHOCP Observability.

1.5.2. Dashboard

Pinned CPUs

Starting with RHOSO 18.0 (GA), the OpenStack Dashboard service (horizon) shows how many pinned CPUs (pCPUs) are used and available to use in your environment.

1.5.3. Data plane

Ansible-managed data plane

In RHOSO 18.0 (GA), the director-deployed overcloud is replaced by a data plane driven by the OpenStack Operator and executed by Ansible. RHOSO data plane features include:

- The **OpenStackDataPlaneNodeSet** custom resource definition (CRD), which provides a highly parallel deployment model.
- Micro failure domains based on the **OpenStackDataPlaneNodeSet** CRD. If one or more node sets fail, the other node sets run to completion because there is no interdependency between node sets.
- Faster deployment times compared to previous RHOSP versions.
- Highly configurable data plane setup based on the **OpenStackDataPlaneNodeSet** and **OpenStackDataPlaneService** CRDs.

1.5.4. Documentation

In RHOSO 18.0 (GA), the documentation library has been restructured to align with the user lifecycle of RHOSO. Each guide incorporates content from one or more product areas that work together to cover end-to-end tasks. The titles are organized in categories for each stage in the user lifecycle of RHOSO.

The following categories are published with RHOSO 18.0:

Plan

Information about the release, requirements, and how to get started before deployment. This category includes the following guides:

- Release notes
- Planning your deployment
- Integrating partner content

Prepare, deploy, configure, test

Procedures for deploying an initial RHOSO environment, customizing the control plane and data plane, configuring validated architectures, storage, and testing the deployed environment. This category includes the following guides:

- Deploying Red Hat OpenStack Services on OpenShift
- Customizing the Red Hat OpenStack Services on OpenShift deployment
- Deploying a Network Functions Virtualization environment
- Deploying a hyper-converged infrastructure environment
- Configuring persistent storage
- Validating and troubleshooting the deployed cloud

Adopt and update

Information about performing minor updates to the latest maintenance release of RHOSO, and procedures for adopting a Red Hat OpenStack Platform 17.1 cloud. This category includes the following guides:

- Adopting a Red Hat OpenStack Platform 17.1 overcloud to a Red Hat OpenStack Services on OpenShift 18.0 data plane
- Updating your environment to the latest maintenance release

Customize and scale

Procedures for configuring and customizing specific components of the deployed environment. These procedures must be done before you start to operate the deployment. This category includes the following guides:

- Configuring the Compute service for instance creation
- Configuring data plane networking
- Configuring load balancing as a service
- Customizing persistent storage
- Configuring security services
- Auto-scaling for instances

Manage resources and maintain the cloud

Procedures that you can perform during ongoing operation of the RHOSO environment. This category includes the following guides:

- Maintaining the Red Hat OpenStack Services on OpenShift deployment
- Creating and managing instances
- Performing storage operations
- Performing security operations
- Managing networking resources
- Managing cloud resources with the Dashboard
- Monitoring high availability services

1.5.4.1. Documentation in progress

In RHOSO 18.0 (GA), the following titles are being reviewed and will be published asynchronously:

- Configuring the Bare Metal Provisioning service
- Configuring load balancing as a service (Technology Preview)

1.5.4.2. RHOCP feature documentation

Starting with RHOSO 18.0 (GA), features that are supported and managed natively in RHOCP are documented in the RHOCP documentation library. The RHOSO documentation includes links to relevant RHOCP documentation where needed.

1.5.4.3. Earlier documentation versions

The RHOSO documentation page shows documentation for version 18.0 and later. For earlier supported versions of RHOSP, see Product Documentation for Red Hat OpenStack Platform 17.1 .

1.5.5. High availability

High availability managed natively in RHOCP

Starting with RHOSO 18.0 (GA), RHOSO high availability (HA) uses RHOCP primitives instead of RHOSP services to manage failover and recovery deployment.

1.5.6. Networking

Egress QoS support at NIC level using DCB(Technology Preview)

Starting with RHOSO 18.0 (GA), egress quality of service (QoS) at the network interface controller (NIC) level uses the Data Center Bridging Capability Exchange (DCBX) protocol to configure egress QoS at the NIC level in the host. It triggers the configuration and provides the information directly from the top of rack (ToR) switch that peers with the host NIC. This capability, combined with egress QoS for OVS/OVN, enables end-to-end egress QoS.
This is a Technology Preview feature. A Technology Preview feature might not be fully implemented and tested. Some features might be absent, incomplete, or not work as expected.

For more information on this feature, see Feature Integration document - DCB for E2E QoS .

Configuring and deploying networking with Kubernetes NMState Operator and the RHEL NetworkManager service (Technology preview)

Starting with RHOSO 18.0 (GA), the RHOSO bare-metal network deployment uses **os-net-config** with a Kubernetes NMState Operator and NetworkManager back end. Therefore, administrators can use the Kubernetes NMState Operator, **nmstate**, and the RHEL NetworkManager CLI tool **nmcli** to configure and deploy networks on the data plane, instead of legacy **ifcfg** files and **network-init-scripts**.

1.5.7. NFV

Power optimization enhancements

RHOSO 18.0 (GA) features a Tuned power saving profile, **cpu-partitioning-powersave**. You can use this profile to improve CPU power consumption by shutting down idle CPU cores or associated sub-systems. Additionally, support for adaptive nano sleep enables power saving for low packet rates.

1.5.8. Observability

Enhanced Openstack Observability

- In RHOSO 18.0 (GA), enhanced dashboards provide unified observability with visualizations that are natively integrated into the RHOCP Observability UI. These include the **node_exporter** agent that exposes metrics to the Prometheus monitoring system.
- In RHOSO 18.0 (GA), the **node_exporter** agent replaces the **collectd** daemon, and Prometheus replaces the Time series database (Gnocchi).

Logging

In RHOSO 18.0 (GA), the OpenStack logging capability is significantly enhanced. You can now collect logs from the control plane and Compute nodes, and use RHOCP Logging to store them in-cluster via Loki log store or forward them off-cluster to an external log store. Logs that are stored in-cluster with Loki can be visualized in the RHOCP Observability UI console.

Service Telemetry Framework deprecation

The Observability product for previous versions of RHOSP is Service Telemetry Framework (STF). With the release of RHOSO 18.0 (GA), STF is deprecated and in maintenance mode. There are no feature enhancements for STF after STF 1.5.4, and STF status reaches end of life at the end of the RHOSP 17.1 lifecycle. Maintenance versions of STF will be released on new EUS versions of RHOCP until the end of the RHOSP 17.1 lifecycle.

1.5.9. Security

FIPS enabled by default

- Starting with RHOSO 18.0 (GA), Federal Information Processing Standard (FIPS) is enabled by default when RHOSO is installed on a FIPS enabled RHOCP cluster in new deployments.
- You do not enable or disable FIPS in your RHOSO configuration. You control the FIPS state in the underlying RHOCP cluster.

TLS-everywhere enabled by default

In RHOSO 18.0 (GA), after deployment, you can configure public services with your own certificates. You can deploy without TLS-everywhere and enable it later. You cannot disable TLS-everywhere after you enable it.

Secure RBAC enabled by default

The Secure Role-Based Access Control (RBAC) policy framework is enabled by default in RHOSO 18.0 (GA) deployments.

Key Manager (barbican) enabled by default

The Key Manager is enabled by default in RHOSO 18.0 (GA) deployments.

1.5.10. Storage

Integration with external Red Hat Ceph Storage (RHCS) 7 clusters

You can integrate RHOSO 18.0 (GA) with external RHCS 7 clusters to include RHCS capabilities with your deployment.

Distributed image import

RHOSO 18.0 (GA) introduces distributed image import for the Image service (glance). With this feature, you do not need to configure a shared staging area for different API workers to access images that are imported to the Image service. Now the API worker that owns the image data is the same API worker that performs the image import.

Block Storage service (cinder) backup and restore for thin volumes

Starting with RHOSO 18.0 (GA), the backup service for the Block Storage service service preserves sparseness when restoring a backup to a new volume. This feature ensures that restored volumes use the same amount of storage as the backed up volume. It does not apply to RBD backups, which use a different mechanism to preserve sparseness.

Support for RHCS RBD deferred deletion

RHOSO 18.0 (GA) introduces Block Storage service and Image service RBD deferred deletion, which improves flexibility in the way RBD snapshot dependencies are managed. With deferred deletion, you can delete a resource such as an image, volume, or snapshot even if there are active dependencies.

Shared File Systems service (manila) CephFS NFS driver with Ganesha Active/Active

In RHOSO 18.0 (GA), the CephFS-NFS driver for the Shared File Systems service consumes an active/active Ganesha cluster by default, improving both the scalability and high availability of the Ceph NFS service.

Unified OpenStack client parity with native Shared File Systems service client

Starting with RHOSO 18.0 (GA), the Shared File Systems service fully supports the **openstack** client command line interface.

CHAPTER 2. RELEASE INFORMATION RHOSO 18.0

These release notes highlight selected updates in some or all of the Red Hat Services on OpenShift (RHOSO) components. Consider these updates when you deploy this release of RHOSO. Each of the notes in this section refers to the Jira issue used to track the update. If the Jira issue security level is public, you can click the link to see the Jira issue. If the security level is restricted, the Jira issue ID does not have a link to the Jira issue.

2.1. RELEASE INFORMATION RHOSO 18.0.15

2.1.1. Advisory list

This release of Red Hat OpenStack Services on OpenShift (RHOSO) includes the following advisories:

RHBA-2025:23151

Release of components for RHOSO 18.0.15

RHBA-2025:23179

Release of containers for RHOSO 18.0.15

2.1.2. Observability

2.1.2.1. Technology Previews

This part provides a list of all Technology Previews available in Red Hat OpenStack Services on OpenShift 18.0.

For information on the scope of support for Technology Preview features, see Technology Preview Features - Scope of Support.

Chargeback and rating for RHOSO clouds with the Rating service (cloudkitty)

RHOSO administrators can evaluate the chargeback and rating capabilities of RHOSO clouds by enabling the Rating service (cloudkitty) in the Telemetry Operator.

The Rating service collects data from the RHOSO Telemetry service to generate cloud resource usage through rating reports. These reports can be consumed by external financial operations or billing systems. The Rating service does not provide a billing interface.

Jira:RHOSSTRAT-623

2.1.2.2. Known issues

This part describes known issues in Red Hat OpenStack Services on OpenShift 18.0.

Incomplete RabbitMQ cluster view

The problem occurs because a single RabbitMQ node is specifically selected in the scrape configuration, rather than the entire cluster. This leads to an incomplete RabbitMQ cluster view for clusters with multiple nodes, providing limited user insight. There is no workaround.

Jira:OSPRH-21020

2.1.3. Compute

2.1.3.1. Known issues

This part describes known issues in Red Hat OpenStack Services on OpenShift 18.0.

Potential scaling issue in PCI in Placement

If there are many similar child providers defined under the same root provider, the allocation candidate generation algorithm in the Placement service scales poorly with the default configuration of placement.

For example, if a Compute node has 8 or more child resource providers, each providing one resource, and an instance requests 8 or more such resources each in independent request groups, then without further optimization enabled, the GET **allocation_candidates** query takes too long to calculate and the scheduling of the instance will fail.

In this situation, make the following configuration changes in the **OpenStackControlPlane** CR:

```
spec:
placement:
    template:
      customServiceConfig: |
        [workarounds]
        optimize_for_wide_provider_trees = True
        [placement]
        max_allocation_candidates = 1000
        allocation_candidates_generation_strategy = breadth-first
```

Jira:OSPRH-20355

Compute service power management feature disabled by default

The Compute service (nova) power management feature is disabled by default. You can enable it with the following **nova-compute** configuration:

```
[libvirt]
cpu_power_management = true
cpu_power_management_strategy = governor
```

The default **cpu_power_management_strategy cpu_state** is currently unsupported. Restarting **nova-compute** causes all dedicated PCPUs on that host to be powered down, including those used by instances. If the **cpu_state** strategy is used, the CPUs of those instances become unpinned.

Jira:OSPRH-10772

Subset of image properties not being used in image properties weigher

The image properties weigher does not take the **os_version** or **os_admin_user** properties into account when calculating the raw weights when these properties are included in the configuration.

Jira:OSPRH-21740

2.1.4. Data plane

2.1.4.1. Bug fixes

This part describes bugs fixed in Red Hat OpenStack Services on OpenShift 18.0 that have a significant impact on users.

The edpm_bootstrap service no longer fails when a data plane node is rebooted

When a data plane node is rebooted, the **edpm_bootstrap** service no longer fails with the error **The conditional check 'not boot_file_entry_present' failed.**

Jira:OSPRH-21709

2.1.5. Hardware Provisioning

2.1.5.1. Known issues

This part describes known issues in Red Hat OpenStack Services on OpenShift 18.0.

Error replacing unprovisioned nodes

Red Hat OpenStack Services on Openshift (RHOSO) uses **metal3** for provisioning unprovisioned dataplane nodes. An error state occurs when you must replace a node where the **bootMacAddress** cannot be updated. The result is that the node is stuck in a state where it must be completely removed from the deployment and provisioned as if it is a new node. If the **automatedCleaningMode** attribute is set **automatedCleaningMode: disabled**, this error state does not occur.

Workaround: When provisioning unprovisioned dataplane nodes, ensure the **automatedCleaningMode** attribute is set to **automatedCleaningMode:disabled**.

Jira:OSPRH-20302

2.1.6. Networking

2.1.6.1. Bug fixes

This part describes bugs fixed in Red Hat OpenStack Services on OpenShift 18.0 that have a significant impact on users.

Bond interfaces can now be configured in RHOSO 18

Before this update, bond interfaces configured in the **OpenStackControlPlane** custom resource (CR) for OVN gateways failed. The **ovn-operator** would throw an error similar to the following:

```
…
Path:"" ERRORED: error configuring pod [openstack/ovn-controller-ovs-5rjnm] networking:
[openstack/ovn-controller-ovs-5rjnm/68204fb3-394d-4990-80d3-fc2e388c2ee3:datacentre]: error
adding container to network "datacentre": failed to move link invalid argument
': StdinData:
…
```

With this update, the problem has been resolved, and the **ovn-operator** correctly creates bond interfaces on the **ovn-controller** and **ovn-controller-ovs** pods.

Jira:OSPRH-18544

2.1.6.2. Known issues

This part describes known issues in Red Hat OpenStack Services on OpenShift 18.0.

Log failure when using neutron-fwaas driver and log API plugins in same neutron deployment

When deploying RHOSO with both the **neutron-fwaas** driver and log API plugins, end users might experience log creation failure. To work around this issue, do not use both the **neutron-fwaas** driver and log API plugins in the same neutron deployment.

Jira:OSPRH-21926

2.1.7. Control plane

2.1.7.1. Known issues

This part describes known issues in Red Hat OpenStack Services on OpenShift 18.0.

Control plane temporarily unavailable during minor update

During minor updates, the RHOSO control plane temporarily becomes unavailable. API requests might fail with HTTP error codes, such as error 500. Alternatively, the API requests might succeed but the underlying life cycle operation fails. For example, a virtual machine instance created with the **openstack server create** command during the minor update never reaches the **ACTIVE** state. The control plane outage is temporary and automatically recovers after the minor update is finished. The control plane outage does not affect the already running workload.

Workaround: To prevent this disruption, see the Red Hat Knowledgebase article How to enable mirrored queues in Red Hat Openstack Services on OpenShift.

Jira:OSPRH-10790

2.1.8. Storage

2.1.8.1. Known issues

This part describes known issues in Red Hat OpenStack Services on OpenShift 18.0.

Commands for quota usage failing with errors

In RHOSO 18, the **cinder-manage quota check** and **cinder-manage quota sync** commands fail when no **project-id** argument is specified, preventing accurate management of quota usage in the Block Storage service (cinder).

Workaround: There is currently no workaround for this issue.

Jira:OSPRH-19057

Image service uploads image that exceeds size quota before rejecting further uploads

When you upload an image to the Image service (glance) that is larger than the configured image size limit (**image_size_total**), the upload succeeds because the Image service does not verify the image size before upload. When the image is uploaded and stored, the Image service determines the image size, which might exceed the quota. However, the Image service rejects any subsequent uploads because the quota is exceeded.

Jira:OSPRH-20618

2.1.9. Optimize Service

2.1.9.1. Known issues

This part describes known issues in Red Hat OpenStack Services on OpenShift 18.0.

Volume migration failures in watcher

Currently, if you attempt to migrate a volume when a volume type has a **volume_backend_name** parameter value that does not match the destination pool **volume_backend_name** parameter value, an error is raised.

Workaround: Configure all volume types and cinder pools that will participate in volume migrations to a common value for the **volume_backend_name**.

Jira:OSPRH-20763

Workflow Engine does not revert actions for failed Action Plans

In RHOSO 18, the Optimize service (watcher) Engine does not automatically revert failed actions when an Action Plan fails, when configured to do so by enabling the **watcher_applier.rollback_when_actionplan_failed** configuration option.

Workaround: Manually revert each failed action in the Action Plan. To avoid the rollback, you can diagnose and fix the root cause of the failure and then run the Audit again to propose a new solution.

Jira:OSPRH-19735

2.2. RELEASE INFORMATION RHOSO 18.0.14

2.2.1. Advisory list

This release of Red Hat OpenStack Services on OpenShift (RHOSO) includes the following advisories:

RHBA-2025:20964

Release of components for RHOSO 18.0.14 (Feature Release 4)

RHSA-2025:21132

Release of containers for RHOSO 18.0.14

2.2.2. Compute

2.2.2.1. New features

This part describes new features and major enhancements introduced in Red Hat OpenStack Services on OpenShift 18.0.

Full support for adoption of Compute hosts with /var/lib/nova/instances on NFS

Previously, this feature was available as Technology Preview. With this release, you can migrate Compute hosts with **/var/lib/nova/instances** on NFS as part of the adoption of a RHOSP 17.1 deployment.

Jira:OSPRH-9175

2.2.2.2. Bug fixes

This part describes bugs fixed in Red Hat OpenStack Services on OpenShift 18.0 that have a significant impact on users.

Enhanced os-vif OVS plugin improves network performance on OVS interfaces

Previously, a bug fix was made to OVS that changed the application of the kernel's default QOS policy to OVS ports. This fix was applied to 17.1 but not to 18.0. As a result, a regression in the network configuration for OVS interfaces negatively impacted network performance of Openstack instances when using kernel OVS. With this update, the **os-vif** OVS plugin has been enhanced to improve network performance on OVS interfaces by using the linux-noop QOS policy by default. This can still be overridden by neutron QOS policies. To apply the update, restart or move the instance to recreate the port with a hard reboot, perform a detach followed by an attach operation, or perform a live migrate operation.

Jira:OSPRH-18532

2.2.2.3. Known issues

This part describes known issues in Red Hat OpenStack Services on OpenShift 18.0.

Compute service power management feature disabled by default

The Compute service (nova) power management feature is disabled by default. You can enable it with the following **nova-compute** configuration:

```
[libvirt]
cpu_power_management = true
cpu_power_management_strategy = governor
```

The default **cpu_power_management_strategy cpu_state** is currently unsupported. Restarting **nova-compute** causes all dedicated PCPUs on that host to be powered down, including those used by instances. If the **cpu_state** strategy is used, the CPUs of those instances become unpinned.

Jira:OSPRH-10772

Potential scaling issue in PCI in Placement

If there are many similar child providers defined under the same root provider, the allocation candidate generation algorithm in the Placement service scales poorly with the default configuration of placement.

For example, if a Compute node has 8 or more child resource providers, each providing one resource, and an instance requests 8 or more such resources each in independent request groups, then without further optimization enabled, the GET **allocation_candidates** query takes too long to calculate and the scheduling of the instance will fail.

In this situation, make the following configuration changes in the OpenStackControlPlane CR: spec: placement: template: customServiceConfig: | [workarounds] optimize_for_wide_provider_trees = True [placement] max_allocation_candidates = 1000 allocation_candidates_generation_strategy = breadth-first

Jira:OSPRH-20355

Subset of image properties not being used in image properties weigher

The image properties weigher does not take the **os_version** or **os_admin_user** properties into account when calculating the raw weights when these properties are included in the configuration.

Jira:OSPRH-21740

2.2.3. Networking

2.2.3.1. New features

This part describes new features and major enhancements introduced in Red Hat OpenStack Services on OpenShift 18.0.

DNS service (designate) now fully supported

In RHOSO 18.0.14, the DNS service (designate) is now fully supported. For more information, see *Configuring DNS as a service* .

Jira:RHOSSTRAT-109

Avoiding taskflow interruptions by using flow resumption

In RHOSO 18.0.14, the Load-balancing service (octavia) flow resumption, which automatically reassigns the flow to an alternate controller if the original controller shuts down unexpectedly. For more information, see Avoiding taskflow interruptions by using flow resumption .

Jira:RHOSSTRAT-554

Load-balancing service (octavia) support for DCN now fully supported

In RHOSO 18.0.14, creating load balancers in availability zones (AZs) are now fully supported. For more information, see Creating availability zones for load balancing of network traffic at the edge .

Jira:RHOSSTRAT-577

OVN provider driver for Load-balancing service (octavia) now fully supported

In RHOSO 18.0.14, the OVN provider driver for the Load-balancing service is no longer a Technology Preview and is now fully supported. For more information, see Load-balancing service provider drivers.

Jira:RHOSSTRAT-775

Collect OCP gateway, **ovn-controller**, and **ovs-vswitchd** metrics on data plane nodes and control plane pods

With this update, the Prometheus OVS/OVN exporter collects additional metrics related to OCP gateways, **ovn-controller**, and **ovs-vswitchd**. In addition, OVS/OVN metrics collection is now available on control plane pods. Previously, collection was only available on data plane nodes.

For more information, see Network observability.

Jira:RHOSSTRAT-847

Collect **northd** and RAFT metrics on data plane nodes and control plane pods

With this update, the Prometheus OVS/OVN exporter collects additional metrics related to **northd** and RAFT. In addition, OVS/OVN metrics collection is now available on control plane pods. Previously, collection was only available on data plane nodes. For more information, see Network observability

Jira:RHOSSTRAT-848

Dynamic routing with BGP support for IPv6 networks

In RHOSO feature release 4, you can configure your dynamic routing environment using IPv6 networks. For more information, see Preparing RHOCP for BGP networks on RHOSO .

Jira:RHOSSTRAT-972

2.2.3.2. Bug fixes

This part describes bugs fixed in Red Hat OpenStack Services on OpenShift 18.0 that have a significant impact on users.

Fixes load-balancer member re-enablement behavior

Before this update, disabled load-balancer members remained in an **ERROR** state even after being re-enabled, causing incorrect load balancer status reporting and potential traffic distribution issues. This problem was caused by the health monitor incorrectly reporting the disabled member status as **ERROR** instead of **OFFLINE**. This incorrect status prevented the disabled member from transitioning back to **ONLINE** when it was re-enabled. With this update, the health monitor correctly sets the disabled member status to **OFFLINE** and the problem has been resolved. Disabled members are now correctly marked as **OFFLINE** and the members automatically transition back to **ONLINE** when they are re-enabled.

Jira:OSPRH-20743

2.2.3.3. Technology Previews

This part provides a list of all Technology Previews available in Red Hat OpenStack Services on OpenShift 18.0.

For information on the scope of support for Technology Preview features, see Technology Preview Features - Scope of Support.

TAP-as-a-Service (TAPaas) [Technology Preview]

In this release, you can test a technology preview of TAPaaS.

TAPaaS provides a Openstack integrated framework for scalable port mirroring in a multi-tenant shared environment maintaining the tenant isolation boundaries in Openstack deployments. TAPaaS is a Neutron extension that enables on-demand traffic mirroring for tenant or administrator purposes. It allows users to create TAP services that mirror traffic from one or more Neutron ports and redirect it to a TAP destination—often a virtual Network Packet Broker (NPB), intrusion detection system (IDS), or traffic analyzer instance.

Jira:RHOSSTRAT-39

Firewall-as-a-Service (Technology Preview)

In RHOSO 18.0.14 (Feature Release 4), you can test a technology preview of Firewall-as-a-Service (FWaaS). Do not use technology preview features in production environments.

As more OpenStack-based clouds are adopted for multi-tenant applications, security remains a top priority. Network-level isolation and traffic control become critical, especially in public or hybrid cloud environments.

Although security groups provide sufficient capability to specify security policy at a VM instance level or VM port level, it does not have support to specify policy at a network or router port level.

FWaaS project provides this additional capability to specify the security policies at the router port level and enables specifying multiple policy rules within the same policy group and also supports application of L3 or L2 policy at the router port level.

FWaaS also provides support for NGFW 3rd party plugins for integration with NGFW vendor solutions enabling firewall capabilities beyond the ACL level. Features and capabilities such as DPI, Malware protection, IPS and IDP.

To enable the FWaaS service plugin, add **firewall_v2** to **service_plugins** in your control plane Custom Resource (CR) file as shown in the following example. The example includes other example services for context. These are not required for enabling FWaaS.

To configure the technology preview of FWaaS, add the following settings in your control plane CR:

```
customServiceConfig: |
  [DEFAULT]
  service_plugins = qos,ovn-router,trunk,segments,port_forwarding,log,firewall_v2

  [service_providers]
  service_provider =
FIREWALL_V2:fwaas_db:neutron_fwaas.services.firewall.service_drivers.ovn.firewall_l3_driver.OVNFwaasDriver:default
```

For FWaaS usage examples, see "Configure Firewall-as-a-Service v2" in *Firewall-as-a-Service (FWaaS) v3 scenario* [1].

[1] https://docs.openstack.org/neutron/latest/admin/fwaas-v2-scenario.html

Jira:RHOSSTRAT-499

Support Data plane adoption for BGP control plane (Technology Preview)

With this technology preview, you can upgrade in-place an OpenStack deployment from RHOSP 17.1 to RHOSO 18 by deploying a parallel control plane based on RHOSO 18 and then pointing the compute nodes to the new control plane.

Jira:RHOSSTRAT-902

2.2.4. Network Functions Virtualization

2.2.4.1. New features

This part describes new features and major enhancements introduced in Red Hat OpenStack Services on OpenShift 18.0.

Cleanup of obsolete network configurations

A NIC on a data plane node can be used for bare-metal provisioning and other initial configuration tasks. You can then use os-net-config to reconfigure the NIC for operational use in your deployment. In these cases, obsolete remnants of the initial configuration might cause IP address conflicts. To avoid such conflicts, you can use remove_config to clean up the obsolete configuration files from the initial configuration. For more information, see Cleaning up obsolete host network configurations .

Jira:RHOSSTRAT-660

2.2.4.2. Bug fixes

This part describes bugs fixed in Red Hat OpenStack Services on OpenShift 18.0 that have a significant impact on users.

The default os-net-config provider is nmstate

In previous RHOSO releases, Red Hat did not support **NMstate** as the **os-net-config provider**. It is now supported and the default configuration sets the **os-net-config** provider to **nmstate**.

The parameter is **edpm_network_config_nmstate**. The default value is **true**. If a specific limitation of the **nmstate** provider requires you to use the **ifcfg** provider, change **edpm_network_config_nmstate** to **false**.

For more information, see "The nmstate provider for os-net-config" in the guide Planning your deployment

Jira:OSPRH-11309

2.2.5. Control plane

2.2.5.1. New features

This part describes new features and major enhancements introduced in Red Hat OpenStack Services on OpenShift 18.0.

Prevent minor update from proceeding when the custom container images have not been updated

This enhancement ensures correct version tracking and validation during minor updates by preventing the side effects and inconsistencies that result from custom container images not being updated when the target version is updated.

With this update, when a minor update is initiated by setting the **targetVersion**, the performance of the minor update is halted if the **customImages** version for the associated custom container images is not also updated. Users have the option to force the update if necessary.

Jira:OSPRH-19183

Update rabbitmq-cluster-operator to v2.16.0

This enhancement updates the RabbitMQ service Operator, **rabbitmq-cluster-operator**, to version 2.16.0. With this update, the RabbitMQ clusters are restarted. If you want to control when the RabbitMQ clusters are updated, you can pause reconciliation before performing the update, then resume reconciliation when it is safe to update the RabbitMQ cluster.

Jira:RHOSSTRAT-914

The OpenStack Operator supports customization of the controller manager for service Operators

The OpenStack Operator initialization resource creates each {rhocp_long} service Operator with default CPU and memory resource limits, and default tolerations. This enhancement adds the ability to customize the configuration of the resource limits and tolerations of each {rhocp_long} service

Operator.

Jira:RHOSSTRAT-974

2.2.5.2. Known issues

This part describes known issues in Red Hat OpenStack Services on OpenShift 18.0.

Control plane temporarily unavailable during minor update

During minor updates, the RHOSO control plane temporarily becomes unavailable. API requests might fail with HTTP error codes, such as error 500. Alternatively, the API requests might succeed but the underlying life cycle operation fails. For example, a virtual machine instance created with the **openstack server create** command during the minor update never reaches the **ACTIVE** state. The control plane outage is temporary and automatically recovers after the minor update is finished. The control plane outage does not affect the already running workload.

Workaround: To prevent this disruption, see the Red Hat Knowledgebase article How to enable mirrored queues in Red Hat Openstack Services on OpenShift.

Jira:OSPRH-10790

2.2.6. High availability

2.2.6.1. New features

This part describes new features and major enhancements introduced in Red Hat OpenStack Services on OpenShift 18.0.

Configuring authentication for the memcached service

Starting with RHOSO 18.0.14 (Feature Release 4), you can configure the cache maintained by the memcached service to require authentication to increase the security of your cloud by restricting access to the cached data of your cloud. For more information, see Configuring authentication for the memcached service in Customizing the Red Hat OpenStack Services on OpenShift deployment.

Jira:RHOSSTRAT-288

Adopt RHOSP 17.1 Instance HA environments to RHOSO

Starting with RHOSO 18.0.14 (Feature Release 4), you can adopt Red Hat OpenStack Platform (RHOSP) 17.1 environments with Instance HA enabled to RHOSO 18.0. For more information about adopting Instance HA environments, see Preparing an Instance HA deployment for adoption and link:https://docs.redhat.com/en/documentation/red_hat_openstack_services_on_openshift/18.0/html/adc data-plane_adopt-control-plane#enabling-high-availability-for-instances_data-plane [Enabling the high availability for Compute instances service] in *Adopting a Red Hat OpenStack Platform 17.1 deployment*.

Jira:RHOSSTRAT-654

Configuring quorum queues for RabbitMQ in new deployments

Starting with RHOSO 18.0.14 (Feature Release 4), RabbitMQ supports the use of the **Quorum** queues for new RHOSO deployments. A **Quorum** queue is a durable, replicated queue based on the Raft consensus algorithm, providing increased data safety and high availability. For more information, see step 5 of link: Creating the control plane in the Deploying Red Hat OpenStack Services on OpenShift guide.

Jira:RHOSSTRAT-959

2.2.7. Security and hardening

2.2.7.1. New features

This part describes new features and major enhancements introduced in Red Hat OpenStack Services on OpenShift 18.0.

Multi-realm federation support

Starting with RHOSO 18.0.14 (Feature Release 4), you can configure RHOSO to allow users to log in to the OpenStack Dashboard by using single sign-on (SSO) and select from one of several external Identity Providers (IdPs). For more information, see Configuring multi-realm federated authentication in RHOSO in Configuring security services.

Jira:RHOSSTRAT-613

2.2.8. Storage

2.2.8.1. New features

This part describes new features and major enhancements introduced in Red Hat OpenStack Services on OpenShift 18.0.

Image service support for HTTPd customization

With this update, the Image service (glance) now supports the customization of HTTPd configuration files. You can use the **extraMounts** parameter to include and load a custom **httpd.conf** file.

Jira:OSPRH-20033

Shared File Systems service support for HTTPd customization

With this update, the Shared File Systems service (manila) now supports the customization of HTTPd configuration files. You can use the **extraMounts** parameter to include and load a custom **httpd.conf** file.

Jira:OSPRH-20335

Image service support for per-tenant quotas

With this update, the Image service (glance) supports per-tenant quotas for improved resource management in private clouds.

Jira:RHOSSTRAT-1004

Shared File Systems service (manila) with CephFS through NFS adoption is fully supported

Adopting the Shared File Systems service (manila) with CephFS through NFS is now generally available. Previously, these adoption instructions were provided as a Technology Preview.

This enhancement allows you to migrate your existing Red Hat OpenStack Platform 17.1 deployment that uses CephFS through NFS as a back end for the Shared File Systems Service to RHOSO 18.0 with full support.

The adoption process includes:

- Creating a new clustered NFS Ganesha service managed directly on the Red Hat Ceph cluster
- Migrating export locations from the standalone Pacemaker-controlled **ceph-nfs** service to the new clustered service
- Decommissioning the previous standalone NFS service

For more information, see Changes to CephFS through NFS and Creating an NFS Ganesha cluster in *Adopting a Red Hat OpenStack Platform 17.1 deployment* .

Jira:RHOSSTRAT-291

Full support for adopting environments that use iSCSI back ends for the Block Storage service (cinder)

Starting with RHOSO 18.0.14 (Feature Release 4), the procedure to adopt RHOSO 18.0 is fully supported for Red Hat OpenStack Platform 17.1 environments that use iSCSI as a back end for the Block Storage service (cinder). For more information, see Adopting the Block Storage service in *Adopting a Red Hat OpenStack Platform 17.1 deployment*.

Jira:RHOSSTRAT-507

Image service Cache API support

With this update, the Image service (glance) adds Cache API support. The Image service Cache API provides centralized management of cache nodes, eliminating the requirement for individual SSH connections. Key functionalities include listing cached images of a specific node, queuing images for caching, and deleting cached images or images queued for caching. With this update, administrators can call the API and integrate API calls into workflows and automated tasks.

Jira:RHOSSTRAT-606

Full support for adopting environments that use Block Storage service (cinder) back ends for the Image service (glance)

Starting with RHOSO 18.0.14 (Feature Release 4), RHOSO 18.0 adoption is fully supported for Red Hat OpenStack Platform 17.1 environments that use Block Storage service (cinder) as a back end for the Image service (glance). For more information, see Adopting the Image service that is deployed with a Block Storage service back end in *Adopting a Red Hat OpenStack Platform 17.1 deployment* .

Jira:RHOSSTRAT-670

Deployment of Object Storage service on data plane nodes

With this update, you can deploy the Object Storage service (swift) on external data plane nodes, improving scalability and performance for large storage clusters. By enabling DNS forwarding and creating an **OpenStackDataPlaneNodeSet** CR with specified properties, including disks for storage, you can customize the service configuration through additional **ConfigMap** or **Secret** CRs in the **OpenStackDataPlaneService** CR.

Jira:RHOSSTRAT-684

Shared File Systems service now supports transferring shares between tenants

The Shared File Systems service (manila) now supports transferring shares across projects. To ensure security and non-repudiation, a one-time transfer secret key is generated when you initiate a transfer. The key must be conveyed out-of-band so that a user in the recipient project can complete the transfer.

Jira:RHOSSTRAT-784

Integration of Shared File Systems service with Dell PowerScale

The Shared File Systems service (manila) now includes integration with Dell PowerScale storage systems (formerly Dell Isilon). The driver supports provisioning and managing the lifecycle of NFS and CIFS shared file systems, controlling client access to them, resizing them, creating snapshots, and cloning these snapshots into new shared file systems.

Jira:RHOSSTRAT-803

Notifications for events in the Block Storage service and Shared File Systems service

With this update, you can enable notifications in the Block Storage service (cinder) and Shared File System service (manila) by using the **notificationsBusInstance** parameter, allowing integration with either the existing RabbitMQ instance or a dedicated RabbitMQ instance.

Jira:RHOSSTRAT-876

2.2.8.2. Bug fixes

This part describes bugs fixed in Red Hat OpenStack Services on OpenShift 18.0 that have a significant impact on users.

Creation of snapshots for Block Storage service volumes succeeds in ONTAP FlexGroup pools

Before this update, the creation of snapshots for Block Storage service (cinder) volumes did not succeed in ONTAP FlexGroup pools. With this update, the creation of volume snapshots succeeds, ensuring data backup for end users.

Jira:OSPRH-19954

Image volume cache deletes PowerFlex volumes when the snapshot limit is reached

Before this update, when a PowerFlex volume in the image volume cache reached its snapshot limit, the Block Storage service (cinder) replaced the cache entry. However, the volume itself was not deleted, resulting in unusable volumes that consumed quota. Now, when a PowerFlex volume in the image volume cache reaches its snapshot limit, the Block Storage service replaces the cache entry and deletes the PowerFlex volume for the original cache entry.

Jira:OSPRH-20193

Improved clone deletion management

Before this update, NetApp ONTAP storage systems on version 9.13.1 or later rejected clone deletion requests for FlexVol snapshots and FlexClones if they were still processing. With this update, the NetApp ONTAP driver in the Shared File Systems service (manila) manages clone splitting operations before share or snapshot deletion, preventing deletion failures. Users can now delete Shared File Systems service snapshots and shares created from snapshots any time after creation without encountering clone deletion failures.

Jira:OSPRH-20328

2.2.8.3. Known issues

This part describes known issues in Red Hat OpenStack Services on OpenShift 18.0.

Commands for quota usage failing with errors

In RHOSO 18, the **cinder-manage quota check** and **cinder-manage quota sync** commands fail when no project-id argument is specified, preventing accurate management of quota usage in the Block Storage service (cinder).

Workaround: There is currently no workaround for this issue.

Jira:OSPRH-19057

Image service uploads image that exceeds size quota before rejecting further uploads

When you upload an image to the Image service (glance) that is larger than the configured image size limit (**image_size_total**), the upload succeeds because the Image service does not verify the image size before upload. When the image is uploaded and stored, the Image service determines the image size, which might exceed the quota. However, the Image service rejects any subsequent uploads because the quota is exceeded.

Jira:OSPRH-20618

2.2.9. Dashboard

2.2.9.1. New features

This part describes new features and major enhancements introduced in Red Hat OpenStack Services on OpenShift 18.0.

Customizing the Dashboard httpd timeout value

With this update, you can customize the httpd timeout value of your Dashboard **horizon-operator** by using extraMounts to load a httpd.conf file. For more information, see link: https://github.com/openstack-k8s-operators/horizon-operator/blob/main/config/samples/httpd-overrides/README.md

Jira:OSPRH-10961

Direct upload method and CORS enabled by default for image uploads to the Dashboard service

With this update, the **direct** upload method is the default upload method in the Dashboard service (horizon). The **direct** upload method requires Cross-Origin Resource Sharing (CORS) to be enabled in the Image service (glance), and CORS is now enabled by default in the **glance-operator**.

As a result, you can use the web UI in the Dashboard service to upload images that are greater than 1 GiB in size. The direct upload method sends the image directly from the web browser to the Image service instead of storing the image in the Dashboard service first and then sending it to the Image service.

Jira:RHOSSTRAT-917

2.2.10. Optimize Service

2.2.10.1. New features

This part describes new features and major enhancements introduced in Red Hat OpenStack Services on OpenShift 18.0.

The Optimize service (watcher) is now integrated into the OpenStack Operator

Before this update, the Optimize service (watcher) was provided as a technical preview that was installed with its own Operator. With this update, the Optimize service is now integrated into the OpenStack Operator and is fully supported.

> If you have installed the Optimize service as a technical preview, you must remove the Optimize service Operator and any custom resource of kind `Watcher` before updating your RHOSO deployment to 18.0.14 or later.

Jira:OSPRH-14852

Improve Optimize service accuracy with OpenStack services notifications

With this update, you can now enable the Optimize service (watcher) to receive notification updates from other OpenStack services, which improves Optimize service accuracy.

Jira:OSPRH-18674

Start and end time fields available for CONTINUOUS audits in the OpenStack Dashboard (horizon).

With this update, you can now set the start and end times when creating **CONTINUOUS** audits in the OpenStack Dashboard (horizon).

Jira:OSPRH-19133

Parameters field included in the OpenStack Dashboard (horizon).

With this update, you can now include parameters as JSON values when creating an Audit in the Openstack Dashboard.

Jira:OSPRH-20000

Filtered strategy selction in the OpenStack Dashboard (horizon)

With this update the OpenStack Dashboard now limits visible strategies to those applicable to the selected goal.

Jira:OSPRH-20148

2.2.10.2. Bug fixes

This part describes bugs fixed in Red Hat OpenStack Services on OpenShift 18.0 that have a significant impact on users.

Live migration fails after migrating the volume that an instance has attached

Before this update, legacy code in the Optimize service (watcher) related to the volume move operation could result in two instances having access to the same host block device, which could lead to corruption or incorrect data access between tenants after a volume was migrated.

This update resolves the data corruption and data access issues by removing legacy code and delegating volume migrations to the Block storage service (cinder).

Jira:OSPRH-19167

2.2.10.3. Known issues

This part describes known issues in Red Hat OpenStack Services on OpenShift 18.0.

Workflow Engine does not revert actions for failed Action Plans

In RHOSO 18, the Optimize service (watcher) Engine does not automatically revert failed actions when an Action Plan fails, when configured to do so by enabling the **watcher_applier.rollback_when_actionplan_failed** configuration option.

Workaround: Manually revert each failed action in the Action Plan. To avoid the rollback, you can diagnose and fix the root cause of the failure and then run the Audit again to propose a new solution.

Jira:OSPRH-19735

Volume migration failures in watcher

Currently, if you attempt to migrate a volume when a volume type has a **volume_backend_name** parameter value that does not match the destination pool **volume_backend_name** parameter value, an error is raised.

Workaround: Configure all volume types and cinder pools that will participate in volume migrations to a common value for the **volume_backend_name**.

Jira:OSPRH-20763

2.2.11. Migration

2.2.11.1. New features

This part describes new features and major enhancements introduced in Red Hat OpenStack Services on OpenShift 18.0.

Migrate VMs with VMware Migration Toolkit

In RHOSO 18.0.14 (Feature Release 4) and RHOSP 17.1, you can now migrate workloads from VMware to OpenStack using the VMware Migration Toolkit.

Jira:RHOSSTRAT-89

2.2.12. Hardware Provisioning

2.2.12.1. Known issues

This part describes known issues in Red Hat OpenStack Services on OpenShift 18.0.

Error replacing unprovisioned nodes

Red Hat OpenStack Services on Openshift (RHOSO) uses **metal3** for provisioning unprovisioned dataplane nodes. An error state occurs when you must replace a node where the **bootMacAddress** cannot be updated. The result is that the node is stuck in a state where it must be completely removed from the deployment and provisioned as if it is a new node. If the **automatedCleaningMode** attribute is set **automatedCleaningMode: disabled**, this error state does not occur.

Workaround: When provisioning unprovisioned dataplane nodes, ensure the **automatedCleaningMode** attribute is set to **automatedCleaningMode:disabled**.

Jira:OSPRH-20302

2.3. RELEASE INFORMATION RHOSO 18.0.13

2.3.1. Advisory list

This release of Red Hat OpenStack Services on OpenShift (RHOSO) includes the following advisory:

RHBA-2025:17561

Release of components for RHOSO 18.0.13

RHBA-2025:17990

Release of containers for RHOSO 18.0.13

2.3.2. Compute

2.3.2.1. Known issues

This part describes known issues in Red Hat OpenStack Services on OpenShift 18.0.

Compute service power management feature disabled by default

The Compute service (nova) power management feature is disabled by default. You can enable it with the following **nova-compute** configuration:

```
[libvirt]
cpu_power_management = true
cpu_power_management_strategy = governor
```

The default **cpu_power_management_strategy cpu_state** is currently unsupported. Restarting **nova-compute** causes all dedicated PCPUs on that host to be powered down, including those used by instances. If the **cpu_state** strategy is used, the CPUs of those instances become unpinned.

Jira:OSPRH-10772

2.3.3. Control plane

2.3.3.1. Known issues

This part describes known issues in Red Hat OpenStack Services on OpenShift 18.0.

Control plane temporarily unavailable during minor update

During minor updates, the RHOSO control plane temporarily becomes unavailable. API requests might fail with HTTP error codes, such as error 500. Alternatively, the API requests might succeed but the underlying life cycle operation fails. For example, a virtual machine instance created with the **openstack server create** command during the minor update never reaches the **ACTIVE** state. The control plane outage is temporary and automatically recovers after the minor update is finished. The control plane outage does not affect the already running workload.

Workaround: To prevent this disruption, see the Red Hat Knowledgebase article How to enable mirrored queues in Red Hat Openstack Services on OpenShift.

Jira:OSPRH-10790

2.3.4. Storage

2.3.4.1. Known issues

This part describes known issues in Red Hat OpenStack Services on OpenShift 18.0.

IPv6 export locations cannot be used with Shared File Systems service shares that have a CephFS-NFS back end

An issue with Red Hat Ceph Storage prevents the use of IPv6 export locations with Shared File Systems service (manila) shares that have a CephFS-NFS back end. Workaround: Currently, there is no workaround.

Jira:OSPRH-19498

2.3.5. Upgrades and updates

2.3.5.1. Deprecated functionality

This part provides an overview of functionality that has been *deprecated* in Red Hat OpenStack Services on OpenShift 18.0.

Deprecated functionality will likely not be supported in future major releases of this product and is not recommended for new deployments.

Deprecated the **update** service

RHOSO 18.0.10 (Feature Release 3) introduced a new update workflow that splits the OpenStack-related package updates and system-related package updates. With this new feature, the data plane **update** service is deprecated in favor of the split **update-services** and **update-system** services. The **update** service will be removed in a future release. Customers should transition to using the split update feature. For more information about the split update feature, see the Red Hat Knowledgebase article Performing a minor update of OpenStack service containers and RHEL RPMs separately .

Jira:OSPRH-20354

2.3.6. Optimize Service

2.3.6.1. Known issues

This part describes known issues in Red Hat OpenStack Services on OpenShift 18.0.

Volume migration failures in watcher

Currently, if you attempt to migrate a volume when a volume type has a **volume_backend_name** parameter value that does not match the destination pool **volume_backend_name** parameter value, an error is raised.

Workaround: Configure all volume types and cinder pools that will participate in volume migrations to a common value for the **volume_backend_name**.

Jira:OSPRH-20763

Workflow Engine does not revert actions for failed Action Plans

In RHOSO 18, the Optimize service (watcher) Engine does not automatically revert failed actions when an Action Plan fails, when configured to do so by enabling the **watcher_applier.rollback_when_actionplan_failed** configuration option.

Workaround: Manually revert each failed action in the Action Plan. To avoid the rollback, you can diagnose and fix the root cause of the failure and then run the Audit again to propose a new solution.

Jira:OSPRH-19735

2.4. RELEASE INFORMATION RHOSO 18.0.12

2.4.1. Advisory list

This release of Red Hat OpenStack Services on OpenShift (RHOSO) includes the following advisories:

RHBA-2025:15803

Release of containers for RHOSO 18.0.12

RHBA-2025:15804

Control plane Operators for RHOSO 18.0.12

RHBA-2025:15805

Release of components for RHOSO 18.0.12

RHBA-2025:15806

Data plane Operators for RHOSO 18.0.12

RHBA-2025:16120

Containers bug fix advisory for RHOSO 18.0.12

2.4.2. Control plane

2.4.2.1. Known issues

This part describes known issues in Red Hat OpenStack Services on OpenShift 18.0.

Control plane temporarily unavailable during minor update

During minor updates, the RHOSO control plane temporarily becomes unavailable. API requests might fail with HTTP error codes, such as error 500. Alternatively, the API requests might succeed but the underlying life cycle operation fails. For example, a virtual machine instance created with the **openstack server create** command during the minor update never reaches the **ACTIVE** state. The control plane outage is temporary and automatically recovers after the minor update is finished. The control plane outage does not affect the already running workload.

Workaround: To prevent this disruption, see the Red Hat Knowledgebase article How to enable mirrored queues in Red Hat Openstack Services on OpenShift.

Jira:OSPRH-10790

2.4.3. Storage

2.4.3.1. Known issues

This part describes known issues in Red Hat OpenStack Services on OpenShift 18.0.

IPv6 export locations cannot be used with Shared File Systems service shares that have a CephFS-NFS back end

An issue with Red Hat Ceph Storage prevents the use of IPv6 export locations with Shared File Systems service (manila) shares that have a CephFS-NFS back end. Workaround: Currently, there is no workaround.

Jira:OSPRH-19498

2.5. RELEASE INFORMATION RHOSO 18.0.11

2.5.1. Advisory list

This release of Red Hat OpenStack Services on OpenShift (RHOSO) includes the following advisories:

RHBA-2025:14763

Control plane Operators for RHOSO 18.0.11

RHBA-2025:14747

Data plane Operators for RHOSO 18.0.11

RHBA-2025:14762

Release of containers for RHOSO 18.0.11

RHBA-2025:14745

Release of components for RHOSO 18.0.11

2.5.2. Compute

2.5.2.1. Bug fixes

This part describes bugs fixed in Red Hat OpenStack Services on OpenShift 18.0 that have a significant impact on users.

Enhanced os-vif OVS plugin improves network performance on OVS interfaces

Previously, a bug fix was made to OVS that changed the application of the kernel's default QOS policy to OVS ports. This fix was applied to 17.1 but not to 18.0. As a result, a regression in the network configuration for OVS interfaces negatively impacted network performance of Openstack instances when using kernel OVS. With this update, the **os-vif** OVS plugin has been enhanced to improve network performance on OVS interfaces by using the linux-noop QOS policy by default. This can still be overridden by neutron QOS policies. To apply the update, restart or move the instance to recreate the port with a hard reboot, perform a detach followed by an attach operation, or perform a live migrate operation.

Jira:OSPRH-18532

Compute instance with ISO image boots correctly

Previously, a Compute instance with an ISO image booted via a block device instead of CD-ROM. This prevented the RHEL Kickstart installation from initiating from the CD-ROM. With this bug fix, Compute boots the instance from the ISO image correctly and via CD-ROM.

Jira:OSPRH-17569

2.5.2.2. Known issues

This part describes known issues in Red Hat OpenStack Services on OpenShift 18.0.

Compute service power management feature disabled by default

The Compute service (nova) power management feature is disabled by default. You can enable it with the following **nova-compute** configuration:

```
[libvirt]
cpu_power_management = true
cpu_power_management_strategy = governor
```

The default **cpu_power_management_strategy cpu_state** is currently unsupported. Restarting **nova-compute** causes all dedicated PCPUs on that host to be powered down, including those used by instances. If the **cpu_state** strategy is used, the CPUs of those instances become unpinned.

Jira:OSPRH-10772

2.5.3. Control plane

2.5.3.1. Bug fixes

This part describes bugs fixed in Red Hat OpenStack Services on OpenShift 18.0 that have a significant impact on users.

OpenStack Operator checks namespace field for upgrade of Operators

This update fixes an issue where upgrades from OpenStack Operator version 1.0.6 or earlier sometimes failed when OpenShift Lifecycle Manager (OLM) Operator resources contained data with no namespace field defined.

With this update, the OpenStack Operator checks that the namespace field is implemented for Operator references in the OpenStack controller and the OpenStack Service Operators upgrade is not affected.

Jira:OSPRH-17456

2.5.3.2. Known issues

This part describes known issues in Red Hat OpenStack Services on OpenShift 18.0.

Control plane temporarily unavailable during minor update

During minor updates, the RHOSO control plane temporarily becomes unavailable. API requests might fail with HTTP error codes, such as error 500. Alternatively, the API requests might succeed but the underlying life cycle operation fails. For example, a virtual machine instance created with the **openstack server create** command during the minor update never reaches the **ACTIVE** state. The control plane outage is temporary and automatically recovers after the minor update is finished. The control plane outage does not affect the already running workload.

Workaround: To prevent this disruption, see the Red Hat Knowledgebase article How to enable mirrored queues in Red Hat Openstack Services on OpenShift.

Jira:OSPRH-10790

2.5.4. Storage

2.5.4.1. Bug fixes

This part describes bugs fixed in Red Hat OpenStack Services on OpenShift 18.0 that have a significant impact on users.

Multipart image upload with S3 back end

Before this update, you had to use the image import workflow to upload multipart images if you had an S3 back end for the Image service (glance). With this update, you can set **s3_store_large_object_size** to **0** to force multipart upload when you create an image in the S3 back end from a Block Storage service (cinder) volume.

Jira:OSPRH-11018

2.5.4.2. Known issues

This part describes known issues in Red Hat OpenStack Services on OpenShift 18.0.

IPv6 export locations cannot be used with Shared File Systems service shares that have a CephFS-NFS back end

An issue with Red Hat Ceph Storage prevents the use of IPv6 export locations with Shared File Systems service (manila) shares that have a CephFS-NFS back end. Workaround: Currently, there is no workaround.

Jira:OSPRH-19498

2.5.5. Optimize Service

2.5.5.1. Known issues

This part describes known issues in Red Hat OpenStack Services on OpenShift 18.0.

Volume migration operations are technology preview

Volume migration operations that are a part of the zone migration strategy are provided as a technology preview only, and should not be used in production.

Jira:OSPRH-17639

2.6. RELEASE INFORMATION RHOSO 18.0.10

2.6.1. Advisory list

This release of Red Hat OpenStack Services on OpenShift (RHOSO) includes the following advisories:

RHBA-2025:12089

Release of components for RHOSO 18.0.10 (Feature Release 3)

RHBA-2025:12090

Release of containers for RHOSO 18.0.10 (Feature Release 3)

RHSA-2025:12091

Security release of Control plane Operators for RHOSO 18.0.10 (Feature Release 3)

RHBA-2025:12092

Data plane Operators for RHOSO 18.0.10 (Feature Release 3)

2.6.2. Observability

2.6.2.1. New features

This part describes new features and major enhancements introduced in Red Hat OpenStack Services on OpenShift 18.0.

The Telemetry service collects telemetry related to RHOSO database services

This enhancement implements a new exporter that enables observability of the MariaDB databases that run within RHOSO.

Jira:RHOSSTRAT-882

Compute metrics are available to Prometheus for telemetry data collection and storage

Telemetry data for Compute nodes is now collected directly from Prometheus rather than transiting the internal message bus, enabling storage of Compute node telemetry in the telemetry storage system.

[NOTE] You cannot collect Compute node metrics from Prometheus in an IPv6 environment.

Jira:RHOSSTRAT-905

2.6.2.2. Bug fixes

This part describes bugs fixed in Red Hat OpenStack Services on OpenShift 18.0 that have a significant impact on users.

Collection of telemetry data no longer disrupted by DNS search domains

This update fixes an issue where DNS search domains (**dns_search_domains**) shorter than 8 characters that appeared alphabetically before the control plane DNS domain caused disruption in the collection of telemetry data.

Jira:OSPRH-16162

2.6.3. Compute

2.6.3.1. New features

This part describes new features and major enhancements introduced in Red Hat OpenStack Services on OpenShift 18.0.

Enablement of Nova notifications in RHOSO

This update adds support for configuring a dedicated notifications message bus in the **nova-operator**. By setting the **notificationsBusInstance** in the Nova custom resource (CR), operators can specify an external RabbitMQ for emitting versioned and unversioned notifications.

The **[notification]** and **[oslo_messaging_notifications]** sections are rendered in **nova.conf**.

When **novaEnabledNotification** is set and a **transport_url** is provided via an OpenShift secret, **nova-compute** emits structured notifications to external systems, improving integration and observability in RHOSO environments.

To enable Nova notifications in RHOSO, you update the **OpenStackControlPlane** CR to add a new RabbitMQ instance and reference it in the Nova CR by using **notificationsBusInstance**. The **nova-operator** configures control plane services automatically.

For data plane updates, redeploy the data plane nodes.

Jira:OSPRH-16489

Support for data plane adoption of the source cloud with multiple nova cells

The cloud operator can now adopt an existing 17.1 multi-cell nova deployment with a common network from TripleO management to the new installer for RHOSO.

Jira:OSPRH-6548

2.6.3.2. Known issues

This part describes known issues in Red Hat OpenStack Services on OpenShift 18.0.

Compute service power management feature disabled by default

The Compute service (nova) power management feature is disabled by default. You can enable it with the following **nova-compute** configuration:

```
[libvirt]
cpu_power_management = true
cpu_power_management_strategy = governor
```

The default **cpu_power_management_strategy cpu_state** is currently unsupported. Restarting **nova-compute** causes all dedicated PCPUs on that host to be powered down, including those used by instances. If the **cpu_state** strategy is used, the CPUs of those instances become unpinned.

Jira:OSPRH-10772

2.6.4. Hardware Provisioning

2.6.4.1. Bug fixes

This part describes bugs fixed in Red Hat OpenStack Services on OpenShift 18.0 that have a significant impact on users.

Improved logging and error handling for cross-controller packet loss

Before this update, cross-controller packet loss could impact request handling by the python-networking-baremetal agent and prevent physical network mapping updates from occurring in the Networking service (neutron) for bare-metal nodes. With this update, there is additional logging and error handling so that the python-networking-baremetal provided service exits and the container can automatically restart if packet loss occurs. Physical network mappings for bare-metal nodes continue to to be updated if network interruptions for Controller nodes occur.

Jira:OSPRH-10799

Workflow operations persist through interruptions in connectivity

This update solves an issue in the Bare Metal Provisioning service (ironic) that caused the deployment process to loop and time out because of interruptions in connectivity while the deployment agent was starting. The issue occurred because only one attempt was made to evaluate if a RAM drive was recently booted. When this issue occurred, the bare metal nodes would fail to clean, deploy, or perform other workflow actions.

Jira:OSPRH-15883

2.6.5. Networking

2.6.5.1. New features

This part describes new features and major enhancements introduced in Red Hat OpenStack Services on OpenShift 18.0.

Adoption of combined Networker/Controller nodes

Adoption of RHOSP 17.1 environments that use combined Controller/Networker nodes are verified to work as documented in Adopting a Red Hat OpenStack Platform 17.1 deployment .

Jira:OSPRH-10714

2.6.5.2. Bug fixes

This part describes bugs fixed in Red Hat OpenStack Services on OpenShift 18.0 that have a significant impact on users.

North/south fragmentation fix

Before this update, OpenStack did not fragment north-south packets as expected when the external maximum transmission unit (MTU) was less than the internal MTU, which resulted in packets being dropped silently. With this update, fragmentation happens as expected, and packets are not dropped silently.

Jira:OSPRH-12695

Fix improves BGP recovery time

Before this update, disabling Bidirectional Forwarding Detection (BFD) in Free Range Routing (FRR) on RHEL 9.4 could cause traffic disruptions during error recovery.

With this release, you no longer need to disable BFD for BGP peers. Operating with BFD for BGP peers enhances BGP time recovery and minimizes traffic disruption.

Jira:OSPRH-15728

2.6.5.3. Technology Previews

This part provides a list of all Technology Previews available in Red Hat OpenStack Services on OpenShift 18.0.

For information on the scope of support for Technology Preview features, see Technology Preview Features - Scope of Support.

DNS as a service

With this technology preview, you can test the management of DNS records, names, and zones using the DNS service (designate). For more information, see *Configuring DNS as a service* .

Jira:RHOSSTRAT-441

Load-balancing service (octavia) support for DCN

With this technology preview, you can test creating load balancers in availability zones (AZs) to increase traffic throughput, reduce latency, and enhance security. For more information, see Creating availability zones for load balancing of network traffic at the edge.

Jira:RHOSSTRAT-528

Create Load-balancing service resources for a specific project

Load-balancing service (octavia) resources are created by default within a project (tenant) service. RHOSO 18.0.10 (Feature Release 3) introduces a technology preview of a new **TenantName** parameter for the Octavia Operator, which restricts the use of the resource to a specific project. RHOSO administrators can also change the domain of the project.

Jira:RHOSSTRAT-614

2.6.5.4. Deprecated functionality

This part provides an overview of functionality that has been *deprecated* in Red Hat OpenStack Services on OpenShift 18.0.

Deprecated functionality will likely not be supported in future major releases of this product and is not recommended for new deployments.

Deprecation of ovn-bgp-agent

Since RHOSO 18.0.10 (Feature Release 3), the OVN BGP Agent **ovn-bgp-agent** is deprecated. **ovn-bgp-agent** is the BGP integration component in RHOSO. An alternative BGP integration mechanism is scheduled for a future release. Until then, Red Hat will provide only bug fixes and support for this feature.

Jira:OSPRH-17130

2.6.6. Network Functions Virtualization

2.6.6.1. New features

This part describes new features and major enhancements introduced in Red Hat OpenStack Services on OpenShift 18.0.

TSO for OVS-DPDK promoted from technology preview to general availability

RHOSO 18.0.6 (Feature Release 2) introduced a technology preview of TCP segmentation offload (TSO) for RHOSO environments with OVS-DPDK.

As of RHOSO 18.0.10 (Feature Release 3), TCP segmentation offload (TSO) for RHOSO environments with OVS-DPDK is a general availability feature.

Jira:OSPRH-3885

2.6.6.2. Bug fixes

This part describes bugs fixed in Red Hat OpenStack Services on OpenShift 18.0 that have a significant impact on users.

Adoption no longer fails when physical function is attached to an instance

Previously, when the physical function (PF) was attached to the instance, if **os-net-config** was re-run, **os-net-config** could not find the SR-IOV PF in the host, and the adoption failed. With this release, the adoption does not fail.

Jira:OSPRH-12024

Fixes NetworkManager-dispatcher failures

Before this update, the **NetworkManager-dispatcher** service was blocked by SELinux permission denial, causing the service to fail when SELinux was enforced. With this release, NetworkManager has been updated to allow running the`NetworkManager-dispatcher` service with SELinux in enforcing mode. As a result, the **NetworkManager-dispatcher** service now runs with SELinux in enforcing mode, eliminating the failures.

Jira:OSPRH-13544

Data plane deployment no longer fails when using the nmstate provider to pre-provision Compute nodes over VLAN

Before this update, when pre-provisioning Compute nodes for communicating with the control plane over VLANs, theNetworkManager CLI (**nmcli**) connection was not always created with the proper interface name. This caused deployment failures.

With this release, the issue with the **nmstate** provider for handling vlan interfaces in pre-provisioned nodes has been resolved. As a result, data plane deployments using the **nmstate** provider succeeds.

Jira:OSPRH-16526

Fixes edpm_network_config_nonconfigured_cleanup parameter issue

The flag **edpm_network_config_nonconfigured_cleanup: true** was introduced as default in Feature Release 2 and caused some new deployments to fail.

With this update, appropriate use of the flag **edpm_network_config_nonconfigured_cleanup: true** no longer causes deployment failures.

You can now set **edpm_network_config_nonconfigured_cleanup: true** when you do the following configurations:

- Use unprovisioned or pre-provisioned nodes with a VLAN-tagged interface using either the ifcfg or nmstate provider.
- Have multiple dataplanes with separate namespaces and a tagged VLAN on the control plane interface.

Set **edpm_network_config_nonconfigured_cleanup: false** when you do the following configurations:

- Use unprovisioned or pre-provisioned physical interface with a flat network or bond using either the ifcfg or nmstate provider.

- Perform network updates or RHOSO minor updates.
- Perform a data plane adoption.
- Have multiple data planes with separate namespaces and a flat network on the control plane interface.

Jira:OSPRH-16537

Bandwidth limit now applied to instances with VLAN and flat ports with nmstate provider

Previously, in environments using the **os-net-config nmstate** provider, QoS bandwidth limit rules were not properly applied to the physical NIC attached to the **br-ex** bridge.

With this update, the QoS bandwidth limit rules are applied.

Jira:OSPRH-17551

Patch ports no longer cause network update failures

This update fixes an issue in environments with the nmstate provider where network updates failed on Compute nodes that hosted active instances with patch ports present in br-ex.

Jira:OSPRH-17622

2.6.7. Control plane

2.6.7.1. New features

This part describes new features and major enhancements introduced in Red Hat OpenStack Services on OpenShift 18.0.

Multiple RHOSO deployments on a single RHOCP cluster by using namespace separation

This feature enables you to deploy multiple RHOSO environments on a single RHOCP cluster by using namespace (project) isolation for development, staging and testing environments.

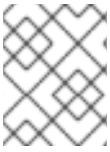

NOTE

Multiple RHOSO environments on a single cluster are not supported for production environments.

For more information, see Deploying multiple RHOSO environments on a single RHOCP cluster

Jira:RHOSSTRAT-612

Minimizing downtime during a minor update

During a minor update, the control plane services are updated concurrently. This enhancement isolates the **galera**, **rabbitmq**, **memcached**, and **keystone** services to perform their updates consecutively, in order, within the minor control plane services update phase.

Jira:RHOSSTRAT-871

Documentation: Updated "Installing the OpenStack Operator" procedure

The Installing the OpenStack Operator procedure has been updated to include changing the default automatic Operator update approvals to manual approvals. Using manual update approval enables RHOSO administrators to schedule RHOSO Operator updates.

Jira:RHOSSTRAT-890

2.6.7.2. Known issues

This part describes known issues in Red Hat OpenStack Services on OpenShift 18.0.

Control plane temporarily unavailable during minor update

During minor updates, the RHOSO control plane temporarily becomes unavailable. API requests might fail with HTTP error codes, such as error 500. Alternatively, the API requests might succeed but the underlying life cycle operation fails. For example, a virtual machine instance created with the **openstack server create** command during the minor update never reaches the **ACTIVE** state. The control plane outage is temporary and automatically recovers after the minor update is finished. The control plane outage does not affect the already running workload.

Workaround: To prevent this disruption, see the Red Hat Knowledgebase article How to enable mirrored queues in Red Hat Openstack Services on OpenShift.

Jira:OSPRH-10790

Upgrading the OpenStack Operator can fail due to Operators that are not namespaced

This update fixes an issue where upgrades from OpenStack Operator version 1.0.6 or earlier sometimes failed when OpenShift Lifecycle Manager (OLM) Operator resources contain data with no namespace field defined.

With this update, the OpenStack Operator checks that the namespace field is implemented for Operator references in the OpenStack controller and the OpenStack Service Operators upgrade is not affected.

Jira:OSPRH-17456

2.6.8. High availability

2.6.8.1. Bug fixes

This part describes bugs fixed in Red Hat OpenStack Services on OpenShift 18.0 that have a significant impact on users.

NodeName string updated in JSON tag for BGPConfiguration parameter

Before this update, the BGPConfiguration parameter **spec.frrNodeConfigurationSelector.nodename** had an inconsistency in its JSON tag where the NodeName string **json:"frrConfigurationNamespace,omitempty** was incorrect because **frrConfigurationNamespace** is a node name. With this update, the NodeName string in the JSON tag is correctly set as **json:"nodeName,omitempty"**. You can now configure the **frrNodeConfigurationSelector** by using the following spec:

```
frrNodeConfigurationSelector:
- nodeName: nodeA
  nodeSelector:
```

```
matchLabels:
  foo: bar
```

During an update to the fixed version, any node names that you previously specified by using the **frrConfigurationNamespace** JSON tag are removed and you must use the correct **nodeName** JSON tag to reconfigure your node names.

Jira:OSPRH-16550

2.6.9. Storage

2.6.9.1. New features

This part describes new features and major enhancements introduced in Red Hat OpenStack Services on OpenShift 18.0.

S3 driver for Image service (glance) has option to specify path to CA bundle

With this update, the S3 driver for the Image service has a new **s3_store_cacert** option that allows users to specify the path to a Certificate Authority (CA) bundle to use.

Jira:OSPRH-11189

Red Hat Ceph Storage 8 NFS is supported

Before this update, NFS was not supported when integrating with Red Hat Ceph Storage 8. With this update, NFS is now supported with Red Hat Ceph Storage 8 integrations.

Jira:OSPRH-14788

API token based authentication with the VAST Data storage driver in the Shared File Systems Service (manila)

With this update, cloud administrators can use either **vast_mgmt_user** and **vast_mgmt_password** or **vast_api_token** when configuring authentication in the Shared File Systems service for their VAST Data storage systems. API-based authentication is useful in RHOSO deployments if cloud administrators need an alternative to passwords when specifying VAST Data API users.

Jira:OSPRH-16085

Improved Fibre Channel performance when detaching a volume

With this update, there is improved Fibre Channel performance when detaching a volume because there is no longer a requirement to call the **lsscsi** command.

Jira:OSPRH-16924

Distributed zones with third-party storage

RHOSO 18.0.10 (Feature Release 3) supports the integration of third-party storage appliances within distributed zone environments. The NFS and Fiber Channel protocols in distributed zone environments are provided as a technology preview and are not yet recommended for production use.

Jira:RHOSSTRAT-259

Image service (glance) notifications for events in image lifecycle

With this update, you can enable notifications in the Image service by using the **notificationBusInstance** parameter, allowing integration with either the existing RabbitMQ instance or a dedicated one.

Jira:RHOSSTRAT-682

CephFS file name added to CephFS share metadata

With this update, you can check a CephFS file name when mounting a native CephFS share by viewing the **__mount_options** metadata of the share in the output of the following command:

```
$ openstack share show <share_id>
```

Jira:RHOSSTRAT-685

2.6.9.2. Bug fixes

This part describes bugs fixed in Red Hat OpenStack Services on OpenShift 18.0 that have a significant impact on users.

Improved reliability for Fibre Channel volume attachments

Before this update, Fibre Channel volume attachments failed intermittently with a **NoFibreChannelVolumeDeviceFound** error due to partial scanning of devices. With this update, a broader scan results in better discovery of devices and successful attach operations.

Jira:OSPRH-16737

2.6.9.3. Technology Previews

This part provides a list of all Technology Previews available in Red Hat OpenStack Services on OpenShift 18.0.

For information on the scope of support for Technology Preview features, see Technology Preview Features - Scope of Support.

Added options for customizing the Object Storage service (swift)

With this update, you can test two new options to customize deployments of the Object Storage service by using externally-managed rings. With this technology preview, you can now disable automatic ring management and spread large rings over multiple configmaps.

Jira:RHOSSTRAT-789

2.6.9.4. Known issues

This part describes known issues in Red Hat OpenStack Services on OpenShift 18.0.

Multipart image upload does not work with S3 back end

If you upload multipart images with an S3 back end, you must use the import workflow.

Jira:OSPRH-11018

Red Hat Ceph Storage 8 Object Gateway is not supported

The Red Hat Ceph Storage Object Gateway (RGW) is currently not supported when integrating with Red Hat Ceph Storage 8.

Workaround: There is no current workaround.

Jira:OSPRH-14789

2.6.10. Upgrades and updates

2.6.10.1. New features

This part describes new features and major enhancements introduced in Red Hat OpenStack Services on OpenShift 18.0.

Granular package update workflow

RHOSO 18.0.10 (Feature Release 3) introduces a feature to separate the update process for RHOSO EDPM nodes into two distinct phases: * updating OpenStack (containers & essential packages) and * updating the system (all packages).

This separation gives operators finer control over the update process, reducing risks and simplifying troubleshooting in the event of issues.

Jira:RHOSSTRAT-23

2.6.11. Optimize Service

2.6.11.1. New features

This part describes new features and major enhancements introduced in Red Hat OpenStack Services on OpenShift 18.0.

The dst_node parameter is now optional for the Zone migration strategy

Before this update, the implementation of the Zone migration strategy was affected by the **dst_node** parameter. Now the implementation is in line with API schema and the **dst_node** parameter is optional. If you do not specify a value for **dst_node**, the Nova scheduler selects an appropriate host automatically.

Jira:OSPRH-17179

2.6.11.2. Bug fixes

This part describes bugs fixed in Red Hat OpenStack Services on OpenShift 18.0 that have a significant impact on users.

Fix for RHOSO Watcher Action Plans status

This update fixes an issue where RHOSO Watcher did not correctly report the state of **Action Plans** after all **Actions** finished, for example reporting **SUCCESS** if some **Actions** actually finished with a state of **FAILED**.

Jira:OSPRH-17257

2.6.11.3. Technology Previews

This part provides a list of all Technology Previews available in Red Hat OpenStack Services on OpenShift 18.0.

For information on the scope of support for Technology Preview features, see Technology Preview Features - Scope of Support.

Support for new strategies for Optimize service (watcher)

RHOSO 18.0.10 (Feature Release 3) introduces support for three new supported strategies in the Optimize service: host maintenance, zone migration for instances, and workload balance. For more information about supported strategies to achieve resource optimization goals, see Sample Optimize service workflows in Optimizing infrastructure resource utilization.

Jira:RHOSSTRAT-237

2.6.11.4. Known issues

This part describes known issues in Red Hat OpenStack Services on OpenShift 18.0.

Volume migration operations are technology preview

Volume migration operations that are a part of the zone migration strategy are provided as a technology preview only, and should not be used in production.

Jira:OSPRH-17639

2.7. RELEASE INFORMATION RHOSO 18.0.9

2.7.1. Advisory list

This release of Red Hat OpenStack Services on OpenShift (RHOSO) includes the following advisories:

RHBA-2025:9211

Control plane Operators for RHOSO 18.0.9

RHBA-2025:9212

Data plane Operators for RHOSO 18.0.9

RHBA-2025:9213

Release of containers for RHOSO 18.0.9

RHBA-2025:9214

Release of components for RHOSO 18.0.9

2.7.2. Compute

2.7.2.1. Known issues

This part describes known issues in Red Hat OpenStack Services on OpenShift 18.0.

Compute service power management feature disabled by default

The Compute service (nova) power management feature is disabled by default. You can enable it with the following **nova-compute** configuration:

```
[libvirt]
cpu_power_management = true
cpu_power_management_strategy = governor
```

The default **cpu_power_management_strategy cpu_state** is currently unsupported. Restarting **nova-compute** causes all dedicated PCPUs on that host to be powered down, including ones used by instances. If the **cpu_state** strategy is used, the CPUs of those instances become unpinned.

Jira:OSPRH-10772

2.7.3. Data plane

2.7.3.1. Bug fixes

This part describes bugs fixed in Red Hat OpenStack Services on OpenShift 18.0 that have a significant impact on users.

The redhat service is restored to the default list of data plane services

Before this update, the **redhat** service was removed temporarily from the default list of data plane services, and users had to manually add the **redhat** service to the list of services in the **OpenStackDataPlaneNodeSet** CR. With this update, the **redhat** service is restored to the default list of data plane services.

Jira:OSPRH-15644

2.7.4. Networking

2.7.4.1. Bug fixes

This part describes bugs fixed in Red Hat OpenStack Services on OpenShift 18.0 that have a significant impact on users.

QoS information for VLAN or flat network ports persists through port updates

Any VLAN or flat network port with egress QoS policy rules (maximum and/or minimum bandwidth) stores this information in the **Logical_Switch_Port. options** dictionary. Before this update, any update on this port, from a port name change to a live migration, deleted this QoS information. With this update, the QoS information persists through port updates.

Jira:OSPRH-15457

2.7.4.2. Known issues

This part describes known issues in Red Hat OpenStack Services on OpenShift 18.0.

Packets silently dropped when external MTU is greater than internal MTU

RHOSO does not fragment north-south packets as expected when the external MTU is greater than the internal MTU. Instead, the ingress packets are dropped with no notification.

Also, fragmentation does not work on east/west traffic between tenant networks.

Until these issues are resolved, ensure that the external MTU settings are less than or equal to internal MTU settings, and that all MTU settings on east/west paths are equal.

Workaround:

Until these issues are resolved, perform the following steps to ensure that the external MTU settings are less than or equal to internal MTU settings, and that all MTU settings on east/west paths are equal.

1. Set **ovn_emit_need_to_frag** to **true**.
2. Set **global_physnet_mtu** to a size that is at least 58 bytes larger than the external network MTU, to accommodate the geneve tunnel encapsulation overhead.
3. Set **physical_network_mtus** value pairs to describe the MTU of each physical network.
4. Ensure that the MTU setting on every device on the external network is less than the internal MTU setting.
5. To apply the changes to an existing router, delete the router and re-create it.

Example

For example, suppose that the external network **datacentre** MTU is 1500.

- Enter the following neutron settings in your OpenStackControlPlane CR:

```
neutron:
    enabled: true
:
    template:
 :
      customServiceConfig: |
        [DEFAULT]
        global_physnet_mtu=1558
        [ml2]
        physical_network_mtus = ["datacentre:1500_{context}"]
        [ovn]
        ovn_emit_need_to_frag = true
```

  - Ensure that the MTU setting on every device on the external network is less than the internal MTU setting.
  - Ensure that all tenant networks that use the OVN router have the same MTU.
  - To apply the changes to an existing router, delete the router and re-create it.

Jira:OSPRH-12695

2.7.5. Network Functions Virtualization

2.7.5.1. Bug fixes

This part describes bugs fixed in Red Hat OpenStack Services on OpenShift 18.0 that have a significant impact on users.

Data plane deployment no longer fails when using the nmstate provider to pre-provision Compute nodes over VLAN

Before this update, when pre-provisioning Compute nodes for communicating with the control plane over VLANs, theNetworkManager CLI (**nmcli**) connection was not always created with the proper interface name. This caused deployment failures.

With this release, the issue with the nmstate provider for handling vlan interfaces in pre-provisioned nodes has been resolved. As a result, data plane deployments using the nmstate provider succeeds.

Jira:OSPRH-16526

2.7.6. Control plane

2.7.6.1. Known issues

This part describes known issues in Red Hat OpenStack Services on OpenShift 18.0.

Control plane temporarily unavailable during minor update

During minor updates, the RHOSO control plane temporarily becomes unavailable. API requests might fail with HTTP error codes, such as error 500. Alternatively, the API requests might succeed but the underlying life cycle operation fails. For example, a virtual machine (VM) created with the **openstack server create** command during the minor update never reaches the **ACTIVE** state. The control plane outage is temporary and automatically recovers after the minor update is finished. The control plane outage does not affect the already running workload.

Workaround: To prevent this disruption, see the Red Hat Knowledgebase article How to enable mirrored queues in Red Hat Openstack Services on OpenShift.

Jira:OSPRH-10790

2.7.7. High availability

2.7.7.1. New features

This part describes new features and major enhancements introduced in Red Hat OpenStack Services on OpenShift 18.0.

The Instance HA service supports a new parameter

This enhancement adds the **TAGGED_AGGREGATES** parameter to the RHOSO high availability for Compute instances (Instance HA) service. By default, this parameter is set to **true**, so that the Instance HA service checks for tagged host aggregates. If you set this parameter to **false** then the Instance HA service does not check for tagged host aggregates and therefore will evacuate all the eligible Compute nodes.

Jira:OSPRH-16342

2.7.8. Storage

2.7.8.1. Known issues

This part describes known issues in Red Hat OpenStack Services on OpenShift 18.0.

Multipart image upload does not work with S3 back end

If you upload multipart images with an S3 back end, you must use the import workflow.

Jira:OSPRH-11018

Red Hat Ceph Storage 8 NFS is not supported

NFS is currently not supported when integrating with Red Hat Ceph Storage 8.

Workaround: There is no current workaround.

Jira:OSPRH-14788

Red Hat Ceph Storage 8 Object Gateway is not supported

The Red Hat Ceph Storage Object Gateway (RGW) is currently not supported when integrating with Red Hat Ceph Storage 8.

Workaround: There is no current workaround.

Jira:OSPRH-14789

2.8. RELEASE INFORMATION RHOSO 18.0.8

2.8.1. Advisory list

This release of Red Hat OpenStack Services on OpenShift (RHOSO) includes the following advisories:

RHBA-2025:8036
: Control plane Operators for RHOSO 18.0.8

RHBA-2025:8037
: Data plane Operators for RHOSO 18.0.8

RHBA-2025:8038
: Release of containers for RHOSO 18.0.8

RHBA-2025:8039
: Release of components for RHOSO 18.0.8

2.8.2. Compute

2.8.2.1. Known issues

This part describes known issues in Red Hat OpenStack Services on OpenShift 18.0.

Compute service power management feature disabled by default

The Compute service (nova) power management feature is disabled by default. You can enable it with the following **nova-compute** configuration:

```
[libvirt]
cpu_power_management = true
cpu_power_management_strategy = governor
```

The default **cpu_power_management_strategy cpu_state** is currently unsupported. Restarting **nova-compute** causes all dedicated PCPUs on that host to be powered down, including ones used by instances. If the **cpu_state** strategy is used, the CPUs of those instances become unpinned.

Jira:OSPRH-10772

2.8.3. Data plane

2.8.3.1. Bug fixes

This part describes bugs fixed in Red Hat OpenStack Services on OpenShift 18.0 that have a significant impact on users.

Default policy ensures nftables reload at the end of deployment

Before this update, iptables default tables were added to nftables to ensure backwards compatibility. However, there was a default ALLOW INPUT rule instead of a default DROP rule, and nftables were not reloaded at the end of the deployment. With this update, the correct rules are applied to ensure that nftables are reloaded at the end of the deployment.

Jira:OSPRH-15473

2.8.3.2. Known issues

This part describes known issues in Red Hat OpenStack Services on OpenShift 18.0.

Manually add the redhat service to the default list of data plane services

The **redhat** service has been removed temporarily from the default list of data plane services. As a result, when you attach subscriptions or repositories to Compute nodes and use the documented **rhc_*** parameters when creating the data plane secrets, the nodes are not registered and the data plane deployment fails.

Workaround: Override the services list in your **OpenStackDataPlaneNodeSet** CR, and ensure that you add the **redhat** service as the first service in the list. You can copy the default list shown in Data plane services in *Customizing the Red Hat OpenStack Services on OpenShift deployment* .

Jira:OSPRH-15644

2.8.4. Hardware Provisioning

2.8.4.1. Bug fixes

This part describes bugs fixed in Red Hat OpenStack Services on OpenShift 18.0 that have a significant impact on users.

Bare-metal data plane node with multi-path block storage boots after being provisioned

Before this update, pre-built whole disk images did not include the **device-mapper-multipath** package, which prevented the paired boot ramdisk from supporting multi-path block storage. This caused bare-metal nodes with multi-path block storage to fail to boot after deployment and instead be stuck in an emergency shell. With this update, pre-built whole disk images now include the **device-mapper-multipath** package and deployed bare-metal nodes no longer enter an emergency shell after being deployed.

Jira:OSPRH-15887

2.8.5. Networking

2.8.5.1. Bug fixes

This part describes bugs fixed in Red Hat OpenStack Services on OpenShift 18.0 that have a significant impact on users.

Logs now available for FRR service

Before this update, no logs were available for the Free Range Routing (FRR) service, which is deployed on the data plane nodes when RHOSO is configured to use Dynamic Routing with BGP. With this update, these logs are available.

Jira:OSPRH-10204

Legacy tripleo Networking services are removed after adoption

Before this update, there were legacy tripleo Networking service (neutron) services after the **edpm_tripleo_cleanup** task, which needed to be removed manually. These services were stopped after adoption, so the RHOSO services were not affected. With this update, all tripleo Networking services are removed after adoption on data plane nodes.

Jira:OSPRH-11323

2.8.5.2. Known issues

This part describes known issues in Red Hat OpenStack Services on OpenShift 18.0.

Packets silently dropped when external MTU is greater than internal MTU

RHOSO does not fragment north-south packets as expected when the external MTU is greater than the internal MTU. Instead, the ingress packets are dropped with no notification.

Also, fragmentation does not work on east/west traffic between tenant networks.

Until these issues are resolved, ensure that the external MTU settings are less than or equal to internal MTU settings, and that all MTU settings on east/west paths are equal.

Workaround:

Until these issues are resolved, perform the following steps to ensure that the external MTU settings are less than or equal to internal MTU settings, and that all MTU settings on east/west paths are equal.

1. Set **ovn_emit_need_to_frag** to **true**.
2. Set **global_physnet_mtu** to a size that is at least 58 bytes larger than the external network MTU, to accommodate the geneve tunnel encapsulation overhead.
3. Set **physical_network_mtus** value pairs to describe the MTU of each physical network.
4. Ensure that the MTU setting on every device on the external network is less than the internal MTU setting.
5. To apply the changes to an existing router, delete the router and re-create it.

Example

For example, suppose that the external network **datacentre** MTU is 1500.

- Enter the following neutron settings in your OpenStackControlPlane CR:

```
neutron:
   enabled: true
:
   template:
 :
     customServiceConfig: |
       [DEFAULT]
       global_physnet_mtu=1558
       [ml2]
       physical_network_mtus = ["datacentre:1500_{context}"]
       [ovn]
       ovn_emit_need_to_frag = true
```

- Ensure that the MTU setting on every device on the external network is less than the internal MTU setting.
- Ensure that all tenant networks that use the OVN router have the same MTU.
- To apply the changes to an existing router, delete the router and re-create it.

Jira:OSPRH-12695

Port updates delete QoS information for VLAN or flat network ports

Any VLAN or flat network port with egress QoS policy rules (maximum and/or minimum bandwidth) stores this information in the **Logical_Switch_Port. options** dictionary. Any update on this port, from a port name change to a live migration, will delete this QoS information.

Workaround: To restore the QoS information, you must remove the QoS policy for this port and set it again.

Jira:OSPRH-15457

2.8.6. Network Functions Virtualization

2.8.6.1. Known issues

This part describes known issues in Red Hat OpenStack Services on OpenShift 18.0.

Adoption fails when physical function is attached to a VM instance

When the physical function (PF) is attached to the instance, if **os-net-config** is re-run, **os-net-config** cannot find the SR-IOV PF in the host, and thus the deployment, update, or adoption fails.

Workaround: Before performing an adoption or network update, migrate the instances to another Compute host.

Jira:OSPRH-12024

NetworkManager-dispatcher scripts fail to run when SELinux is enabled

The **os-net-config** configuration tool uses **NetworkManager-dispatcher** scripts for driver bindings. When SELinux is enabled, these scripts fail to run, and the **os-net-config** network deployment fails.

Workaround: There is currently no workaround.

Jira:OSPRH-13544

2.8.7. Control plane

2.8.7.1. Bug fixes

This part describes bugs fixed in Red Hat OpenStack Services on OpenShift 18.0 that have a significant impact on users.

Failed service updates are being reflected accurately by the deployment status

Before this update, when updates to service configurations failed, the failure was not being reflected in the condition status of the deployment. Instead, the **Ready** condition showed as "True" because the new pods created by the update were not being considered when checking the deployment readiness. With this update, any new pods created during a configuration update are now considered when assessing deployment readiness. If rolling out new pods fails, then the deployment reflects that it is stuck in **Deployment in progress**.

Jira:OSPRH-14472

2.8.7.2. Known issues

This part describes known issues in Red Hat OpenStack Services on OpenShift 18.0.

Control plane temporarily unavailable during minor update

During minor updates, the RHOSO control plane temporarily becomes unavailable. API requests might fail with HTTP error codes, such as error 500. Alternatively, the API requests might succeed but the underlying life cycle operation fails. For example, a virtual machine (VM) created with the **openstack server create** command during the minor update never reaches the **ACTIVE** state. The control plane outage is temporary and automatically recovers after the minor update is finished. The control plane outage does not affect the already running workload.

Workaround: To prevent this disruption, see the Red Hat Knowledgebase article How to enable mirrored queues in Red Hat Openstack Services on OpenShift.

Jira:OSPRH-10790

2.8.8. Security and hardening

2.8.8.1. Bug fixes

This part describes bugs fixed in Red Hat OpenStack Services on OpenShift 18.0 that have a significant impact on users.

Generated CA bundle gets installed on data plane nodes

Before this update, the CA bundle that was generated by the RHOSO control plane was deployed on the data plane node for deployed or running services, but it did not get installed as the CA bundle on the data plane node itself. The CA bundle can include custom third-party CA files, for example, to access a satellite. With this update, the CA bundle gets installed on the data plane node.

Jira:OSPRH-14205

2.9. RELEASE INFORMATION RHOSO 18.0.7

Review the known issues, bug fixes, and other release notes for this release of Red Hat OpenStack Services on OpenShift.

RHOSO 18.0.7 introduces the Optimize service (watcher) to provide a flexible and scalable resource optimization service for multi-tenant RHOSO-based clouds. For more information about the Optimize service, see https://issues.redhat.com/browse/OSPRH-15037 and Optimizing infrastructure resource utilization.

2.9.1. Advisory list

This release of Red Hat OpenStack Services on OpenShift (RHOSO) includes the following advisories:

RHBA-2025:4083

Release of components for RHOSO 18.0.7

RHBA-2025:4084

Release of containers for RHOSO 18.0.7

RHBA-2025:4085

Data plane Operators for RHOSO 18.0.7

RHBA-2025:4086

Control plane Operators for RHOSO 18.0.7

2.9.2. Compute

2.9.2.1. Bug fixes

This part describes bugs fixed in Red Hat OpenStack Services on OpenShift 18.0 that have a significant impact on users.

Compute service fails a ready check for a deployment with an invalid configuration

Before this update, if the Compute service (nova) API raised a configuration error, it returned a 500 error once, and then continued to run with a broken configuration after a reload. This issue occurred because **mod_wsgi** reloaded the wsgi application into the same Python interpreter when an error was raised during application initialization. With this update, the Compute service has been modified to reraise the configuration error until the application can restart cleanly. Now, if you deploy with an invalid configuration, the Compute service API CR fails the **ready** check and updates the Status field in the OpenShift CR to prompt you to review the log files for the configuration error.

Jira:OSPRH-9737

2.9.2.2. Known issues

This part describes known issues in Red Hat OpenStack Services on OpenShift 18.0.

Compute service power management feature disabled by default

The Compute service (nova) power management feature is disabled by default. You can enable it with the following **nova-compute** configuration:

```
[libvirt]
cpu_power_management = true
cpu_power_management_strategy = governor
```

The default **cpu_power_management_strategy cpu_state** is currently unsupported. Restarting **nova-compute** causes all dedicated PCPUs on that host to be powered down, including ones used by instances. If the **cpu_state** strategy is used, those instances' CPUs will become unpinned.

Jira:OSPRH-10772

2.9.3. Data plane

2.9.3.1. Known issues

This part describes known issues in Red Hat OpenStack Services on OpenShift 18.0.

Manually add the redhat service to the default list of data plane services

The **redhat** service has been removed temporarily from the default list of data plane services. As a result, when you attach subscriptions or repositories to Compute nodes and use the documented **rhc_*** parameters when creating the data plane secrets, the nodes are not registered and the data plane deployment fails.

Workaround: Override the services list in your **OpenStackDataPlaneNodeSet** CR, and ensure that you add the **redhat** service as the first service in the list. You can copy the default list shown in Data plane services in *Customizing the Red Hat OpenStack Services on OpenShift deployment* .

Jira:OSPRH-15644

2.9.4. Networking

2.9.4.1. Bug fixes

This part describes bugs fixed in Red Hat OpenStack Services on OpenShift 18.0 that have a significant impact on users.

BFD now works as expected in RHOSO deployments with dynamic routing

Before this update, when you deployed RHOSO with Dynamic Routing with border gateway protocol (BGP), bi-directional forwarding (BFD) did not work as expected because there was no **nft** rule to permit BFD and BGP ports. With this update, an **nft** rule has been added and BFD works as expected:

```
BGP
    - 179 tcp
 BFD
    - 3784 udp
    - 3785 udp
    - 4784 udp
    - 49152 udp
    - 49153 udp
```

Jira:OSPRH-14536

2.9.4.2. Known issues

This part describes known issues in Red Hat OpenStack Services on OpenShift 18.0.

No logs available for FRR service

No logs are available for the FRR service, which is deployed on the data plane nodes when RHOSO is configured to use Dynamic Routing with BGP.

Workaround:

To obtain FRR logs after the **OpenstackDataplaneDeployment** is complete, perform the following actions on all the networker and Compute nodes that are running FRR:

1. Edit the **/var/lib/config-data/ansible-generated/frr/etc/frr/frr.conf`file and replace `log file** with **log file /var/log/frr/frr.log**.
2. Edit the **/var/lib/kolla/config_files/frr.json** and replace **sleep infinity** with **tail -f /var/log/frr/frr.log**.
3. Restart FRR: **systemctl restart edpm_frr**.

Jira:OSPRH-10204

Legacy tripleo Networking services (neutron) after adoption

After the **edpm_tripleo_cleanup** task, there are still legacy tripleo Networking service (neutron) services. These services are stopped after adoption, so the RHOSO services are not affected.

Workaround:

Perform the following steps to remove the legacy services manually:

- Check tripleo neutron services list: **systemctl list-unit-files --type service**
- Remove tripleo services from **/etc/systemd/system/**

Jira:OSPRH-11323

Packets silently dropped when external MTU is greater than internal MTU

RHOSO does not fragment north-south packets as expected when the external MTU is greater than the internal MTU. Instead, the ingress packets are dropped with no notification.

Also, fragmentation does not work on east/west traffic between tenant networks.

Until these issues are resolved, ensure that the external MTU settings are less than or equal to internal MTU settings, and that all MTU settings on east/west paths are equal.

Workaround:

Until these issues are resolved, perform the following steps to ensure that the external MTU settings are less than or equal to internal MTU settings, and that all MTU settings on east/west paths are equal.

1. Set **ovn_emit_need_to_frag** to **true**.
2. Set **global_physnet_mtu** to a size that is at least 58 bytes larger than the external network MTU, to accommodate the geneve tunnel encapsulation overhead.
3. Set **physical_network_mtus** value pairs to describe the MTU of each physical network.
4. Ensure that the MTU setting on every device on the external network is less than the internal MTU setting.

5. To apply the changes to an existing router, delete the router and re-create it.

Example

For example, suppose that the external network **datacentre** MTU is 1500.

- Enter the following neutron settings in your OpenStackControlPlane CR:

```
neutron:
   enabled: true
:
   template:
:
     customServiceConfig: |
       [DEFAULT]
       global_physnet_mtu=1558
       [ml2]
       physical_network_mtus = ["datacentre:1500_{context}"]
       [ovn]
       ovn_emit_need_to_frag = true
```

  - Ensure that the MTU setting on every device on the external network is less than the internal MTU setting.
  - Ensure that all tenant networks that use the OVN router have the same MTU.
  - To apply the changes to an existing router, delete the router and re-create it.

Jira:OSPRH-12695

Port updates delete QoS information for VLAN or flat network ports

Any VLAN or flat network port with egress QoS policy rules (maximum and/or minimum bandwidth) stores this information in the **Logical_Switch_Port. options** dictionary. Any update on this port, from a port name change to a live migration, will delete this QoS information.

Workaround: To restore the QoS information, you must remove the QoS policy for this port and set it again.

Jira:OSPRH-15457

2.9.5. Network Functions Virtualization

2.9.5.1. Bug fixes

This part describes bugs fixed in Red Hat OpenStack Services on OpenShift 18.0 that have a significant impact on users.

Fixes minor update failures starting with updates from 18.0.6

This update fixes a bug that causes minor update failures during updates from RHOSO 18.0.1 through 18.0.5 to 18.0.6 or later. The failure no longer occurs if you update from RHOSO 18.0.6 or later to any version.

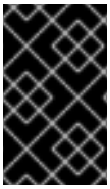

IMPORTANT

If you update from 18.0.1 through 18.0.5 to any version, the update fails because the **edpm_openstack_network_exporter.service** cannot be found. Before you perform these updates, you must perform the following workaround.

Workaround: Add the telemetry service to the **servicesOverride** field in the **openstack-edpm-update-services.yaml** file before you update the `OpenStackDataplaneService`custom resource. For example:

```
apiVersion: dataplane.openstack.org/v1beta1
kind: OpenStackDataPlaneDeployment
metadata:
  name: edpm-deployment-ipam-update-dataplane-services
spec:
  nodeSets:
    - openstack-edpm-ipam
  servicesOverride:
    - telemetry
    - update
```

Jira:OSPRH-14841

2.9.5.2. Deprecated functionality

This part provides an overview of functionality that has been *deprecated* in Red Hat OpenStack Services on OpenShift 18.0.

Deprecated functionality will likely not be supported in future major releases of this product and is not recommended for new deployments.

Deprecated the edpm_ovs_dpdk_lcore_list variable

You can stop using the **edpm_ovs_dpdk_lcore_list** Ansible variable in RHOSO deployments. Previously, it was used in nodeset CR definition files to enable OVS DPDK in data plane deployments in NFV environments. It is no longer required or supported, and its use now causes deployment errors.

Jira:OSPRH-14642

2.9.5.3. Known issues

This part describes known issues in Red Hat OpenStack Services on OpenShift 18.0.

Adoption fails when physical function is attached to a VM instance

When the physical function (PF) is attached to the instance, if **os-net-config** is re-run, **os-net-config** cannot find the SR-IOV PF in the host, and thus the deployment, update, or adoption fails.

Jira:OSPRH-12024

NetworkManager-dispatcher scripts fail to run when SELinux is enabled

The **os-net-config** configuration tool uses **NetworkManager-dispatcher** scripts for driver bindings. When SELinux is enabled, these scripts fail to run, and the **os-net-config** network deployment fails.

Workaround: There is currently no workaround.

Jira:OSPRH-13544

2.9.6. Control plane

2.9.6.1. Bug fixes

This part describes bugs fixed in Red Hat OpenStack Services on OpenShift 18.0 that have a significant impact on users.

TraceEnable parameter disabled by default in httpd configuration

Before this update, HTTP TRACE was enabled by default from the **OpenStackProvisionServer** CR which resulted in security scanners creating an alert. With this update, the **TraceEnable** parameter has been set to the value "off" by default in the httpd configuration.

Jira:OSPRH-14672

2.9.6.2. Known issues

This part describes known issues in Red Hat OpenStack Services on OpenShift 18.0.

Control plane temporarily unavailable during minor update

During minor updates, the RHOSO control plane temporarily becomes unavailable. API requests might fail with HTTP error codes, such as error 500. Alternatively, the API requests might succeed but the underlying life cycle operation fails. For example, a virtual machine (VM) created with the **openstack server create** command during the minor update never reaches the **ACTIVE** state. The control plane outage is temporary and automatically recovers after the minor update is finished. The control plane outage does not affect the already running workload.

Workaround: To prevent this disruption, see the Red Hat Knowledgebase article How to enable mirrored queues in Red Hat Openstack Services on OpenShift.

Jira:OSPRH-10790

2.9.7. Security and hardening

2.9.7.1. Known issues

This part describes known issues in Red Hat OpenStack Services on OpenShift 18.0.

Generated CA bundle does not get installed on data plane nodes

The CA bundle that is generated by the RHOSO control plane gets deployed on the data plane node for deployed or running services, but it does not get installed as the CA bundle on the data plane node itself. The CA bundle can include custom third-party CA files, for example, to access a satellite.
Workaround: There is currently no workaround.

Jira:OSPRH-14205

2.9.8. Optimize Service

2.9.8.1. Technology Previews

This part provides a list of all Technology Previews available in Red Hat OpenStack Services on OpenShift 18.0.

For information on the scope of support for Technology Preview features, see Technology Preview Features - Scope of Support.

Optimize service (watcher) for resource optimization

The Red Hat OpenStack Services on OpenShift (RHOSO) Optimize service (watcher) provides a flexible and scalable resource optimization service for multi-tenant RHOSO-based clouds. The Optimize service provides a framework to help you set and manage goals for infrastructure resource utilization.

The Optimize service is focused on helping users realize a wide range of infrastructure resource utilization goals toward reducing data center operating costs. It includes a metrics receiver, complex event processor and profiler, optimization processor, and an action plan applier.

This feature is currently being delivered in Technology Preview and supports a limited number of optimization strategies in this first release. For more information about the Optimize service, see https://docs.redhat.com/en/documentation/red_hat_openstack_services_on_openshift/18.0/html/optimiz

The Optimize service in RHOSO, which was released as a Technology Preview in RHOSO 18.0.6, is now functional as a Technology Preview for the supported strategies in 18.0.7.

Jira:OSPRH-15037

2.10. RELEASE INFORMATION RHOSO 18.0.6

Review the known issues, bug fixes, and other release notes for this release of Red Hat OpenStack Services on OpenShift.

2.10.1. Advisory list

This release of Red Hat OpenStack Services on OpenShift (RHOSO) includes the following advisories:

RHBA-2025:3029

Release of components for RHOSO 18.0.6 (Feature Release 2)

RHBA-2025:3030

Data plane Operators for RHOSO 18.0.6 (Feature Release 2)

RHBA-2025:3031

Release of Operators for RHOSO 18.0.6 (Feature Release 2)

RHBA-2025:3032

Control plane Operators for RHOSO 18.0.6 (Feature Release 2)

RHBA-2025:3033

Release of containers for RHOSO 18.0.6 (Feature Release 2)

2.10.2. Observability

2.10.2.1. New features

This part describes new features and major enhancements introduced in Red Hat OpenStack Services on OpenShift 18.0.

Improved metrics for RHOSO Observability

You can now use new metrics for monitoring the health of RHOSO services, including the following:

- **kube_pod_status_phase**
- **kube_pod_status_ready**
- **node_systemd_unit_state**
- **podman_container_state**
- **podman_container_health**

You can use the **kube_pod_status_phase** and **kube_pod_status_ready** to monitor control plane services.

- **kube_pod_status_phase** - The relevant parameter is **Phase**, with values of Pending, Running, Succeeded, Failed, or Unknown, and corresponding Boolean values of **1** or **0**.
- **kube_pod_status_ready** - This metric also has Boolean values, with **1** indicating that the pod has all the containers running and readiness probes succeeding, and **0** indicating that the pod has not all the containers running or that the readiness probe did not succeed.

You can use the **node_systemd_unit_state** to monitor the running state of data plane services.

- **node_systemd_unit_state ` - The relevant parameter is `State**, with values of activating, active, deactivating, failed, inactive, and corresponding Boolean values of **1** or **0**.

You can use the **podman_container_state** and **podman_container_health** to monitor the health of data plane containerized services.

- **podman_container_state** - This metric can have the following values: -1=unknown, 0=created, 1=initialized, 2=running, 3=stopped, 4=paused, 5=exited, 6=removing, 7=stopping.
- **podman_container_health** - This metric can have the following values: -1=unknown, 0=healthy, 1=unhealthy, 2=starting.

Jira:OSPRH-1052

An additional Ceilometer metric is available

You can now retrieve the **VMs: ceilometer_power_state** metric to indicate the **libvirt** power state.

Jira:OSPRH-10377

Additional VM metrics are available

You can now use the dashboards to view a dedicated VM Network Traffic Dashboard, as well as monitor the power state of VMs.

Jira:OSPRH-10481

Visualize hardware sensor metrics with Ceilometer IPMI

You can now use the dashboards to view the available IPMI sensor hardware metrics from your compute nodes.

Jira:OSPRH-10808

Kepler dashboard is more user-friendly (Technology Preview)

You can now view Compute nodes by human-readable hostnames, instead of with Compute service UUIDs.

Jira:OSPRH-11755

Improved compatibility between Telemetry Operator and OpenShift Logging

You can now use Telemetry Operator with OpenShift Logging versions 6.1 and newer.

Jira:OSPRH-12147

Prometheus connection information is exposed in a secret

Telemetry Operator now creates a secret with the internal Prometheus connection details. Other OpenShift services can use that secret as a service discovery mechanism to connect to Prometheus.

Jira:OSPRH-13223

2.10.2.2. Bug fixes

This part describes bugs fixed in Red Hat OpenStack Services on OpenShift 18.0 that have a significant impact on users.

Scrape Kepler metrics without manual intervention (Technology Preview)

Before this update, a firewall was not applied to Compute nodes, which allowed port 8888 to be open. However, access to port 8888 might be lost unexpectedly if a firewall is enabled. With this update, the ansible role checks whether the firewall is enabled and then opens port 8888. As a result, Prometheus can scrape Kepler metrics without manual intervention.

Jira:OSPRH-11066

Capture GPU metrics with Kepler (Technology Preview)

With this update to Red Hat OpenStack Services on OpenShift (RHOSO), you can now capture GPU metrics with Kepler.

Jira:OSPRH-11890

TLS errors caused node exporter scraping issues

This release of Red Hat OpenStack Services on OpenShift (RHOSO) fixes an issue with scraping metrics in specific dataplane configurations.

Jira:OSPRH-12359

Missing square brackets around IPv6 addresses

This release of Red Hat OpenStack Services on OpenShift (RHOSO) rectifies a potential issue with scraping data because of missing square brackets around IPv6 addresses.

Jira:OSPRH-13151

IPv6 connection refused with RabbitMQ metrics.

With this update to Red Hat OpenStack Services on OpenShift (RHOSO), the **RabbitMQ** metrics exporter now listens on the correct interface on the IPv6 Control Plane network.

Jira:OSPRH-13152

2.10.3. Compute

2.10.3.1. New features

This part describes new features and major enhancements introduced in Red Hat OpenStack Services on OpenShift 18.0.

Unified limits introduced

This update introduces unified limits to RHOSO 18.0. Unified limits is a modern quota system in which quota limits are stored centrally in the Identity service. Unified limits may be enabled by following the documented procedure.

Jira:OSPRH-9518

Verify systemd-container package installation on hypervisors

You can now verify that the **systemd-container** package is installed on the hypervisors before starting the final steps of the data plane adoption. You cannot adopt the source Red Hat OpenStack Platform cloud into Red Hat OpenStack Services on OpenShift (RHOSO) until the package is installed on all hypervisors.

Jira:OSPRH-10321

Periodic healing of Nova internal instance info cache is disabled

By default, **heal_instance_info_cache_interval** is now disabled to improve performance of the neutron API server by removing the load created by the nova-compute agent. This will not affect cache correctness, because it is updated during most VM operations.

Jira:OSPRH-10744

Adoption of nodes with hugepages now supported

With this update, data plane adoption now supports importing hypervisors with OSP workloads configured for use of hugepages.

Jira:OSPRH-12073

Nova enables the live migration of NVIDIA vGPU instances

Nova enables the live migration instances using vGPU resources between hosts if the target uses the same NVIDIA driver version and the same mediated types.

To live-migrate, operators need to modify the configuration for each of the hosts:

```
[libvirt]
live_migration_completion_timeout = 0
live_migration_downtime = 500000
live_migration_downtime_steps = 3
live_migration_downtime_delay = 3
```

Jira:OSPRH-13767

Topology support for Compute service (nova) and placement service

Implemented a new custom resource definition for scheduling RHOSO Nova and Placement services' pods based on TopologySpreadConstraints and Affinity/Anti-Affinity rules.

Jira:OSPRH-13785

Fixed update failure related to nova_statedir_ownership.py

Before this fix, updates from RHOSO18.0.3 to later releases failed with errors related to the missing nova_statedir_ownership.py script. With this fix, updates from RHOSO 18.0.3 to RHOSO 18.0.6 (Feature Release 2) do not generate these errors.

Jira:OSPRH-14506

2.10.3.2. Bug fixes

This part describes bugs fixed in Red Hat OpenStack Services on OpenShift 18.0 that have a significant impact on users.

Allocation candidates are spread between hosts in Placement GET /allocation_candidates query

In a deployment with wide and symmetric provider trees, for example, where there are multiple children providers under the same root having inventory from the same resource class, if the allocation candidate request asks for resources from those children resource providers in multiple request groups the number of possible allocation candidates grows rapidly. The Placement service generates these candidates fully before applying the limit parameter provided in the allocation candidate query. The Placement service takes excessive time and memory to generate this amount of allocation candidates and the client might time out.

To avoid request timeout or out of memory events, a new **[placement]max_allocation_candidates** configuration option is applied during the candidate generation process. By default, the **[placement]max_allocation_candidates** option is set to -1, which means there is no limit and which was the old behavior. Edit the value of this configuration option in affected deployments based on the memory available for the Placement service and the timeout setting of the clients. A suggested value is **100000**.

If the number of generated allocation candidates is limited by the **[placement]max_allocation_candidates** configuration option, you can get candidates from a limited set of root providers, for example, Compute nodes, as the Placement service uses a depth-first strategy, generating all candidates from the first root before considering the next one. To avoid this, use the **[placement]allocation_candidates_generation_strategy** configuration option, which has two possible values:

- **depth-first**: generates all candidates from the first viable root provider before moving to the next. This is the default and this triggers the legacy behavior.
- **breadth-first**: generates candidates from viable roots in a round-robin fashion, creating one candidate from each viable root before creating the second candidate from the first root. This is the possible new behavior.

In a deployment where **[placement]max_allocation_candidates** is configured to a positive number, set **[placement]allocation_candidates_generation_strategy** to **breadth-first**.

Jira:OSPRH-37

Instances with ephemeral storage on NFS share continue working after Compute service restart

Before this update, Compute service (nova) instances with ephemeral storage on NFS shares stopped working as soon as the containerized Compute agent service restarted on the hypervisor host.

With this update, Nova Compute service instances with ephemeral storage on NFS shares no longer stop working. The Nova Compute init container is triggered every time an Openstack Dataplane Deployment is created with the Nova EDPM service that is included into the linked Openstack Dataplane Nodeset, and corrects the permissions of the **/var/lib/nova/** directory contents on hypervisors.

Jira:OSPRH-10729

Fixed update failure related to nova_statedir_ownership.py

Before this fix, updates from RHOSO18.0.3 to later releases failed with errors related to the missing nova_statedir_ownership.py script. With this fix, updates from RHOSO 18.0.3 to RHOSO 18.0.6 (Feature Release 2) do not generate these errors.

Jira:OSPRH-13415

Fixed: Instances with ephemeral storage on NFS share stop working after Compute service restart

Before this update, Compute service (nova) instances with ephemeral storage on NFS shares stopped working as soon as the containerized Compute agent service restarted on the hypervisor host. That happened because of changed permissions of **/var/lib/nova/** instances.

This update fixes that bug.

Jira:OSPRH-14197

2.10.3.3. Known issues

This part describes known issues in Red Hat OpenStack Services on OpenShift 18.0.

Compute service power management feature disabled by default

The Compute service (nova) power management feature is disabled by default. You can enable it with the following **nova-compute** configuration:

```
[libvirt]
cpu_power_management = true
cpu_power_management_strategy = governor
```

The default **cpu_power_management_strategy cpu_state** is currently unsupported. Restarting nova-compute causes all dedicated PCPUs on that host to be powered down, including ones used by instances. If the **cpu_state** strategy is used, those instances' CPUs will become unpinned.

Jira:OSPRH-10772

Block Storage service (cinder) known issue

When you are using Red Hat Ceph Storage as the back end for the Block Storage service (cinder), then you might be unable to extend an attached encrypted volume. **Workaround:** detach the encrypted RBD volume, extend this volume and then reattach it.

Jira:OSPRH-14321

2.10.4. Data plane

2.10.4.1. New features

This part describes new features and major enhancements introduced in Red Hat OpenStack Services on OpenShift 18.0.

Replace faulty nodes without scaling the cloud

With this update, you have the option of replacing faulty nodes without scaling the cloud.

- For pre-provisioned nodes, set a new **ansibleHost** for the node in the **OpenStackDataPlaneNodeSet** CR.
- For provisioned nodes, delete the faulty bare metal host (BMH). The **OpenStackBaremetalSet** CR is reconciled to provision a new available BMH and reset the deployment status of the **OpenStackDataPlaneNodeSet**, prompting you to create a new **OpenStackDataPlaneDeployment** CR for deploying on the newly provisioned node.

You must still manually clean up the removed nodes as with scale-in.

Jira:OSPRH-13948

2.10.4.2. Bug fixes

This part describes bugs fixed in Red Hat OpenStack Services on OpenShift 18.0 that have a significant impact on users.

Documentation: Limitation added detailing maximum length of OpenStackDataPlaneNodeSet CR names

The description of the rules around naming `OpenStackDataPlaneNodeSet` CRs has been updated to include that the maximum length is 53 characters.

Jira:OSPRH-12327

2.10.5. Networking

2.10.5.1. Bug fixes

This part describes bugs fixed in Red Hat OpenStack Services on OpenShift 18.0 that have a significant impact on users.

Fixed default anti-affinity policy in the Load-balancing service

Before this update, the octavia operator did not enable anti-affinity when creating amphorae in active-standby topology. In some cases, the virtual machine were scheduled on the same compute node.

The release fixes this issue and ensures that anti-affinity is enabled.

Jira:OSPRH-10705

Corrected Load-balancing service provider network visibility

Before this update, end users could see the Load-balancing service provider network. Now the Load-balancing service provider network is only visible to the administrators.

Jira:OSPRH-12519

Deploy Load-balancing service in an offline cluster

Before this update, the container image URL of the **octavia-rsyslog** pod was hard-coded and could not be overridden, with the result that users could not deploy the Load-balancing service (octavia) in an offline cluster.

With this update, the container image URL can be overridden, and you can deploy the Load-balancing service offline.

Jira:OSPRH-13530

Fixed stability issue with Load-balancing service health manager in DCN mode

Before this update, when you ran Load-balancing service (octavia) health manager pods in DCN mode, pods were randomly restarted by the operator. With this update, the random restarts do not occur.

Jira:OSPRH-13951

Load-balancing service rsyslog endpoints no longer drop logs from remote area zones

Before this update, if you used the rsyslog service with DCN, the rsyslog pods dropped incoming rsyslog packets because the routes to the remote DCNs were missing. Now the packets are not dropped.

Jira:OSPRH-14085

2.10.5.2. Technology Previews

This part provides a list of all Technology Previews available in Red Hat OpenStack Services on OpenShift 18.0.

For information on the scope of support for Technology Preview features, see Technology Preview Features - Scope of Support.

Amphora Vertical Scaling (Threading/CPU pinning) (Technology Preview)

With this technology preview you can test an improved handling of Load-balancing service vertical scaling support for the amphora driver. The update in this technology preview use an additional amphora image that is specifically optimized for improved vertical scaling and an additional load balancer flavor that uses multiple vCPUs. This can help improve latency and throughput of the load balancer.

Jira:OSPRH-370

TLS client authentication with the Load-balancing service (Technology Preview)

This update includes a Technology Preview of TLS web client communication with a RHOSO Load-balancing service (octavia) TLS-terminated HTTPS load balancer by using certificates to establish two-way TLS authentication.

Jira:OSPRH-1375

Load-balancing service (octavia) support for Distributed zones feature (Technology Preview)

This update introduces a Technology Preview of Load-balancing service (octavia) availability zones (AZs) that enable project users to create load balancers in a Distributed zones environment to increase traffic throughput and reduce latency.

Jira:OSPRH-8336

Multiple Load-balancing service VIP addresses from same network

There are use-cases where a Load-balancing service in Octavia with the Amphora provider needs multiple VIP addresses allocated from the same Neutron network. With this technology preview, you can test the ability to specify additional subnet_id/ip_address pairs to associate with the VIP port. This enables scenarios such as having a load balancer with both IPv4 and IPv6 addresses simultaneously or being exposed to both public and private subnets.

Jira:OSPRH-14480

Improved TLS cipher and protocol support (Technology Preview)

This update introduces a Technology Preview of improved Load-balancing service (octavia) support for TLS cipher and protocol. You can now override the default cipher list with values that are more appropriate for your site, as well as use additional new features such as setting cipher and protocol lists for each listener.

Jira:OSPRH-14705

IPv6 load-balancing network [Technology preview]

You can now test a technology preview that uses IPv6 CIDRs for the load-balancing management network.

Jira:OSPRH-14811

2.10.5.3. Known issues

This part describes known issues in Red Hat OpenStack Services on OpenShift 18.0.

No logs available for FRR service

No logs are available for the FRR service, which is deployed on the data plane nodes when RHOSO is configured to use Dynamic Routing with BGP.

Workaround: To obtain FRR logs after the **OpenstackDataplaneDeployment** is complete, perform the following actions on all the networker and Compute nodes that are running FRR:

1. Edit the **/var/lib/config-data/ansible-generated/frr/etc/frr/frr.conf`file and replace `log file** with **log file /var/log/frr/frr.log**.
2. Edit the **/var/lib/kolla/config_files/frr.json** and replace **sleep infinity** with **tail -f /var/log/frr/frr.log**.
3. Restart FRR: **systemctl restart edpm_frr**.

Jira:OSPRH-10204

Legacy tripleo Networking services (neutron) after adoption

After the **edpm_tripleo_cleanup** task, there are still legacy tripleo Networking service (neutron) services. These services are stopped after adoption, so the RHOSO services are not affected.

Workaround: Perform the following steps to remove the legacy services manually:

- Check tripleo neutron services list: **systemctl list-unit-files --type service**
- Remove tripleo services from **/etc/systemd/system/**

Jira:OSPRH-11323

Packets silently dropped when external MTU is greater than internal MTU

RHOSO does not fragment north-south packets as expected when the external MTU is greater than the internal MTU. Instead, if the ingress packets are dropped with no notification.

Also, fragmentation does not work on east/west traffic between tenant networks.

Until these issues are resolved, ensure that the external MTU settings are less than or equal to internal MTU settings, and that all MTU settings on east/west paths are equal.

Procedure:

1. Set **ovn_emit_need_to_frag** to **true**.
2. Set **global_physnet_mtu** to a size that is at least 58 bytes larger than the external network MTU, to accommodate the geneve tunnel encapsulation overhead.
3. Set **physical_network_mtus** value pairs to describe the MTU of each physical network.
4. Ensure that the MTU setting on every device on the external network is less than the internal MTU setting.
5. To apply the changes to an existing router, delete the router and re-create it.

Example

For example, suppose that the external network **datacentre** MTU is 1500.

1. Enter the following neutron settings in your OpenStackControlPlane CR:

```
neutron:
   enabled: true
:
   template:
 :
     customServiceConfig: |
       [DEFAULT]
       global_physnet_mtu=1558
       [ml2]
       physical_network_mtus = ["datacentre:1500_{context}"]
       [ovn]
       ovn_emit_need_to_frag = true
```

2. Ensure that the MTU setting on on every device on the external network is less than the internal MTU setting.

3. Ensure that all tenant networks that use the OVN router have the same MTU.

4. To apply the changes to an existing router, delete the router and re-create it.

Jira:OSPRH-12695

BFD does not work as expected in RHOSO deployments with dynamic routing; workaround required

When you deploy RHOSO with Dynamic Routing with border gateway protocol (BGP), bi-directional forwarding (BFD) does not work as expected.

Workaround: Add NFT rules to the OpenstackDataplaneNoteSet CRs. There are two ways to do this. Choose one.

1. Disable BFD by setting **edpm_frr_bfd** to **false**.

2. Configure **edpm_nftables_user_rules** to allow BFD traffic:

```
edpm_nftables_user_rules: |
  - rule_name: 121 frr bgp port
    rule:
      proto: tcp
      dport:
       - 179
  - rule_name: 122 frr bfd ports
    rule:
      proto: udp
      dport:
       - 3784
       - 3785
       - 4784
       - 49152
       - 49153
      state: ["UNTRACKED_{context}"]
```

Jira:OSPRH-14536

QoS policies not enforced when only maximum bandwidth (egress) rules present, on ports in physical networks, when the physical interface is a bond

A port connected to a physical network (VLAN, flat), with maximum-bandwidth-only QoS rules and egress direction uses the physical network interface to enforce the QoS rule, via TC commands.

In previous versions, Neutron enforced the bandwidth limit rule using the OVN policer, regardless of the network type and rule direction.

Now, starting with RHOSO 18.0.6, if the environment uses a bond to connect the physical bridge to the physical network, there will be no QoS enforcement. For more information, see https://issues.redhat.com/browse/OSPRH-18010.

Jira:OSPRH-18011

2.10.6. Network Functions Virtualization

2.10.6.1. Bug fixes

This part describes bugs fixed in Red Hat OpenStack Services on OpenShift 18.0 that have a significant impact on users.

Change os-net-config provider to nmstate

In previous RHOSO releases, Red Hat did not support **NMstate** as the **os-net-config provider**. It is now supported but the default configuration sets the **os-net-config** provider to **ifcfg**.

The parameter is **edpm_network_config_nmstate**. The default value is **false**. Change it to **true** to use the **nmstate** provider unless a specific limitation of the **nmstate** provider requires you to use the **ifcfg** provider.

For more information, see "The nmstate provider for os-net-config" in the guide Planning your deployment

Jira:OSPRH-11309

2.10.6.2. Technology Previews

This part provides a list of all Technology Previews available in Red Hat OpenStack Services on OpenShift 18.0.

For information on the scope of support for Technology Preview features, see Technology Preview Features - Scope of Support.

TSO for OVS-DPDK (Technology Preview)

RHOSO 18.0.6 (Feature Release 2) introduces a technology preview of TCP segmentation offload (TSO) for RHOSO environments with OVS-DPDK.

For more information, see OVS-DPDK with TCP segmentation offload (Technology Preview) in *Deploying a network functions virtualization environment* (https://docs.redhat.com/en/documentation/red_hat_openstack_services_on_openshift/18.0/html/deploy ovs-dpdk-deploy_rhoso-nfv#ovsdpdk-tso_plndpdk-nfv).

Jira:OSPRH-3885

2.10.6.3. Deprecated functionality

This part provides an overview of functionality that has been *deprecated* in Red Hat OpenStack Services on OpenShift 18.0.

Deprecated functionality will likely not be supported in future major releases of this product and is not recommended for new deployments.

Deprecated the edpm_ovs_dpdk_lcore_list variable

Stop using the **edpm_ovs_dpdk_lcore_list** Ansible variable in RHOSO deployments. Previously, it was used in nodeset CR definition files to enable OVS DPDK in data plane deployments in NFV environments. It is no longer required or supported. Its use now causes deployment errors.

Jira:OSPRH-14642

2.10.6.4. Known issues

This part describes known issues in Red Hat OpenStack Services on OpenShift 18.0.

Adoption fails when physical function is attached to a VM instance

When the physical function (PF) is attached to the instance, if **os-net-config** is re-run, **os-net-config** cannot find the SR-IOV PF in the host, and thus the deployment/update/adoption fails.

Jira:OSPRH-12024

Cannot set no_turbo when using cpu-partitioning-powersave profile

Due to an issue with setting the **no_turbo** parameter in the kernel, tuned hangs and fails when using the cpu-partitioning-powersave profile.

Workaround: Downgrade tuned as part of the deployment to an older version by adding the following configuration to **edpm_bootstrap_command**:

```
...
edpm_bootstrap_command: |-
...
dnf downgrade tuned-2.24.0
…
```

Jira:OSPRH-14688

Requested service cannot be found during minor update

When updating the remaining services on the data plane, a minor update from 18.0.3 to 18.0.6 fails because the **edpm_openstack_network_exporter.service** cannot be found.

Workaround: Add the telemetry service to the **servicesOverride** field in the **openstack-edpm-update-services.yaml** file before you update the `OpenStackDataplaneService`custom resource. For example:

```
apiVersion: dataplane.openstack.org/v1beta1
kind: OpenStackDataPlaneDeployment
metadata:
  name: edpm-deployment-ipam-update-dataplane-services
spec:
  nodeSets:
    - openstack-edpm-ipam
  servicesOverride:
    - telemetry
    - update
```

Jira:OSPRH-14841

2.10.7. Control plane

2.10.7.1. New features

This part describes new features and major enhancements introduced in Red Hat OpenStack Services on OpenShift 18.0.

Manage service Operator installation under single OLM bundle

The OpenStack Operator no longer installs the multiple RHOSO service Operators individually. Instead, a new initialization resource manages the installation of the service Operators under a single Operator Lifecycle Manager (OLM) bundle. For more information about the new installation method, see

Installing and preparing the Operators.

Jira:OSPRH-11244

Custom environment variables for the OpenStackClient pod

You can set custom environment variables for the **OpenStackClient** pod.

Jira:OSPRH-13969

2.10.7.2. Known issues

This part describes known issues in Red Hat OpenStack Services on OpenShift 18.0.

Control plane temporarily unavailable during minor update

During the minor update to 18.0 Feature Release 1, the RHOSO control plane temporarily becomes unavailable. API requests might fail with HTTP error codes, such as error 500. Alternatively, the API requests might succeed but the underlying life cycle operation fails. For example, a virtual machine (VM) created with the **openstack server create** command during the minor update never reaches the **ACTIVE** state. The control plane outage is temporary and automatically recovers after the minor update is finished. The control plane outage does not affect the already running workload.

Jira:OSPRH-10790

2.10.8. Storage

2.10.8.1. New features

This part describes new features and major enhancements introduced in Red Hat OpenStack Services on OpenShift 18.0.

Enhanced Block Storage volume restoration on thinly provisioned back ends

This enhancement optimizes the process of restoring Block Storage volume backups on any thinly provisioned back end. Previously, when restoring a backup on a thinly provisioned back end, the full volume size was restored instead of only restoring the portion of the volume that was used. This caused unnecessary network traffic and greatly increased the time taken by the restoration process. This enhancement ensures that when restoring a volume on a thinly provisioned back end, only the portion of the volume that was used is restored.

Jira:OSPRH-1783

Red Hat Ceph Storage 8 support

This enhancement adds support for integration with external Red Hat Ceph Storage 8. Due to known issues, not all Red Hat Ceph Storage 8 functionality is supported. For more information about these issues, see the Known Issues section.

Jira:OSPRH-10661

2.10.8.2. Bug fixes

This part describes bugs fixed in Red Hat OpenStack Services on OpenShift 18.0 that have a significant impact on users.

ExtraMounts can process a per-instance propagation when the pod is prefixed with an arbitrary name

When **uniquePodNames** is **true**, every Cinder pod (and in general each component and service) is prefixed by a pseudo-random string. With this update, ExtraMounts can process a per-instance propagation when the pod is prefixed with an arbitrary name.

Jira:OSPRH-11240

2.10.8.3. Known issues

This part describes known issues in Red Hat OpenStack Services on OpenShift 18.0.

Multipart image upload does not work with S3 back end

If you upload multipart images with an S3 back end, you must use the import workflow.

Jira:OSPRH-11018

Red Hat Ceph Storage 8 NFS is not supported

In RHOSO 18.0.6, NFS is currently not supported when integrating with Red Hat Ceph Storage 8.

Workaround: There is no current workaround.

Jira:OSPRH-14788

Red Hat Ceph Storage 8 Object Gateway is not supported.

In RHOSO 18.0.6, the Red Hat Ceph Storage Object Gateway (RGW) is currently not supported when integrating with Red Hat Ceph Storage 8.

Workaround: There is no current workaround.

Jira:OSPRH-14789

2.10.9. Upgrades and updates

2.10.9.1. Known issues

This part describes known issues in Red Hat OpenStack Services on OpenShift 18.0.

Create instance of openstack during minor update

If you update your Red Hat OpenStack Services on OpenShift environment from any release before 18.0.6, you must create an instance of **openstack** after you update the **openstack-operator** to trigger the deployment of all operators. For example:

```
cat > openstack-init.yaml <<'EOF'
---
apiVersion: operator.openstack.org/v1beta1
kind: OpenStack
metadata:
  name: openstack
  namespace: openstack-operators
```

```
EOF

$ oc apply -f ./openstack-init.yaml
```

Jira:OSPRH-14826

2.11. RELEASE INFORMATION RHOSO 18.0.4

Review the known issues, bug fixes, and other release notes for this release of Red Hat OpenStack Services on OpenShift.

2.11.1. Advisory list

This release of Red Hat OpenStack Services on OpenShift (RHOSO) includes the following advisories:

RHBA-2025:0435

Release of components for RHOSO 18.0.4

RHBA-2025:0436

Release of containers for RHOSO 18.0.4

RHBA-2025:0437

Control plane Operators for RHOSO 18.0.4

RHBA-2025:0438

Data plane Operators for RHOSO 18.0.4

RHSA-2025:0439

Moderate: Red Hat OpenStack Platform 18.0.4 (openstack-ironic) security update

2.11.2. Compute

2.11.2.1. Bug fixes

This part describes bugs fixed in Red Hat OpenStack Services on OpenShift 18.0 that have a significant impact on users.

Boot instances from images with hardware architecture properties

Before this update, you could not boot an instance from an image that had **hw_architecture** or **hw_emulation_architecture** properties. With this update, you can boot instances from images that have **hw_architecture** and **hw_emulation_architecture** properties.

Jira:OSPRH-6215

2.11.2.2. Known issues

This part describes known issues in Red Hat OpenStack Services on OpenShift 18.0.

Instances with ephemeral storage on NFS share stop working after Compute service restart

Compute service (nova) instances with ephemeral storage on NFS shares stop working as soon as the containerized Compute agent service restarts on the hypervisor host. That happens because of changed permissions of **/var/lib/nova/** instances.

Workaround: Manually restore permissions to the original values and avoid the service restarts.

Jira:OSPRH-10729

Compute service power management feature disabled by default

The Compute service (nova) power management feature is disabled by default. You can enable it with the following **nova-compute** configuration:

```
[libvirt]
cpu_power_management = true
cpu_power_management_strategy = governor
```

The default **cpu_power_management_strategy cpu_state** is currently unsupported. Restarting nova-compute causes all dedicated PCPUs on that host to be powered down, including ones used by instances. If the **cpu_state** strategy is used, those instances' CPUs will become unpinned.

Jira:OSPRH-10772

Cold migration fails for server with swap in flavor for shared storage like NFS.

When the Compute service (nova) is enabled with shared storage, for example, NFS, cold migration fails if the instance uses the **FLAVOR_SWAP** flavor.

Jira:OSPRH-12784

2.11.3. Data plane

2.11.3.1. Bug fixes

This part describes bugs fixed in Red Hat OpenStack Services on OpenShift 18.0 that have a significant impact on users.

Ansible task writes the registries.conf file successfully in disconnected deployments

Before this update, the **registries.conf** file failed in disconnected deployments because the Ansible task attempted to use the **template** module and the raw string input as the **src**. With this update, the Ansible task writes the **registries.conf** file successfully because it uses the **ansible.builtin.copy** module with the **content** parameter.

Jira:OSPRH-11475

iscsi-starter.service disabled on EDPM nodes after a reboot

Before this update, the **iscsi.service** started after a reboot of EDPM nodes that run an instance with an iSCSI-backed volume, even though **iscsi.service** was not enabled in **edpm-ansible**. This issue occurred because the **iscsi-starter.service** was enabled in the EDPM node's image. With this update, the **iscsi-starter.service** is disabled on EDPM nodes to prevent the issue.

Jira:OSPRH-12372

Documentation: Specify a pool to register nodes if you have multiple Red Hat subscriptions

Before this update, there was an optional command missing to specify a pool when you registered nodes for your RHOSO deployment in the procedures for creating the **OpenStackDataPlaneNodeSet** CR with pre-provisioned nodes or unprovisioned nodes. The missing command could cause a registration error

when deploying the data plane or Compute nodes in RHOSO if you had multiple Red Hat subscriptions.

With this update, the optional command to specify a pool if you have multiple Red Hat subscriptions is included in these procedures.

Jira:OSPRH-12956

2.11.4. Hardware Provisioning

2.11.4.1. Bug fixes

This part describes bugs fixed in Red Hat OpenStack Services on OpenShift 18.0 that have a significant impact on users.

New **default_network_interface** parameter for Bare Metal service (ironic)

Before this update, if the Bare Metal service **network_interface** was not configured during ControlPlane deployment, RHOSO configured it as a no-op.

With this update, RHOSO set the **default_network_interface** parameter to a default value under **customServiceConfig** / **[DEFAULT]**

Jira:OSPRH-10697

2.11.5. Networking

2.11.5.1. Bug fixes

This part describes bugs fixed in Red Hat OpenStack Services on OpenShift 18.0 that have a significant impact on users.

Adds support of dynamic routing without DVR

Previously, you could not use dynamic routing on the data plane with Free Range Routing (FRR) and Border Gateway Protocol (BGP) unless you also used distributed virtual routing (DVR). Now you can use dynamic routing without enabling DVR.

Jira:OSPRH-8429

Update to latest RHOSP 17.1 version before adopting

When performing adoption of a source environment which is older than RHOSP 17.1.4, the workloads experience a prolonged network connectivity disruption. Make sure to update the source environment at least to RHOSP 17.1.4 before adopting.

Jira:OSPRH-10283

Corrected default value for **createDefaultLbMgmtNetwork** and **manageLbMgmtNetworks** when availability zones are defined

Before this update, **createDefaultLbMgmtNetwork** and **manageLbMgmtNetworks** were incorrectly set to **false** when availability zones were defined.

With this update, **createDefaultLbMgmtNetwork** and **manageLbMgmtNetworks** are set to **true** when availability zones are defined.

Jira:OSPRH-11092

Documentation: Set replicas: 3 for ovndbcluster fields

Previously, CR examples throughout the RHOSO documentation did not consistently represent that to support OVN database high availability, you must set **replicas: 3** in the **ovndbcluster-nb** and **ovndbcluster-sb** fields for every control plane CR that includes the OVN spec.

With this update, all CR examples now include the replicas requirement. The following example shows an excerpt from the CR example with the added section:

```
ovn:
  template:
    ovnDBCluster:
      ovndbcluster-nb:
        replicas: 3          <<----------
        dbType: NB
        storageRequest: 10G
        networkAttachment: internalapi
      ovndbcluster-sb:
        replcas: 3           <<----------
        dbType: SB
        storageRequest: 10G
        networkAttachment: internalapi
    ovnNorthd: {}
```

Jira:OSPRH-12462

2.11.5.2. Known issues

This part describes known issues in Red Hat OpenStack Services on OpenShift 18.0.

octavia-operator does not enable anti affinity

The Load-balancing service (octavia) does not currently configure anti-affinity settings in the Compute service (nova) to prevent Amphora VMs from being scheduled to the same Compute node.

Workaround: Add the relevant setting to the Load-balancing service by using the **customServiceConfig** parameter, as shown in the following example:

Example

```
# names will be dependent on the deployment
oc patch  -n openstack openstackcontrolplane openstack-galera-network-isolation --type=merge --
patch '
spec:
  octavia:
    template:
      octaviaHousekeeping:
        customServiceConfig: |
          [nova]
          enable_anti_affinity = true
      octaviaWorker:
        customServiceConfig: |
          [nova]
          enable_anti_affinity = true
```

```
      octaviaHealthManager:
        customServiceConfig: |
          [nova]
          enable_anti_affinity = true
'
```

Jira:OSPRH-10705

Adoption not supported with Load-balancing service (octavia) agents on networker nodes

Adoption of deployments that have Load-balancing service agents deployed on networker nodes is currently not supported.

Jira:OSPRH-10771

Legacy tripleo Networking services (neutron) after adoption

After the **edpm_tripleo_cleanup** task, there are still legacy tripleo Networking service (neutron) services. These services are stopped after adoption, so the RHOSO services are not affected.

Workaround: Perform the following steps to remove the legacy services manually:

- Check tripleo neutron services list: **systemctl list-unit-files --type service**
- Remove tripleo services from **/etc/systemd/system/**

Jira:OSPRH-11323

Packets silently dropped when external MTU is greater than internal MTU

RHOSO does not fragment north-south packets as expected when the external MTU is greater than the internal MTU. Instead, if the ingress packets are dropped with no notification.

Also, fragmentation does not work on east/west traffic between tenant networks.

Until these issues are resolved, ensure that the external MTU settings are less than or equal to internal MTU settings, and that all MTU settings on east/west paths are equal.

Procedure:

1. Set **ovn_emit_need_to_frag** to **true**.
2. Set **global_physnet_mtu** to a size that is at least 58 bytes larger than the external network MTU, to accommodate the geneve tunnel encapsulation overhead.
3. Set **physical_network_mtus** value pairs to describe the MTU of each physical network.
4. Ensure that the MTU setting on every device on the external network is less than the internal MTU setting.
5. To apply the changes to an existing router, delete the router and re-create it.

Example

For example, suppose that the external network **datacentre** MTU is 1500.

1. Enter the following neutron settings in your OpenStackControlPlane CR:

```
neutron:
    enabled: true
:
    template:
 :
      customServiceConfig: |
        [DEFAULT]
        global_physnet_mtu=1558
        [ml2]
        physical_network_mtus = ["datacentre:1500_{context}"]
        [ovn]
        ovn_emit_need_to_frag = true
```

2. Ensure that the MTU setting on every device on the external network is less than the internal MTU setting.

3. Ensure that all tenant networks that use the OVN router have the same MTU.

4. To apply the changes to an existing router, delete the router and re-create it.

Jira:OSPRH-12695

2.11.6. Network Functions Virtualization

2.11.6.1. Known issues

This part describes known issues in Red Hat OpenStack Services on OpenShift 18.0.

Do not use virtual functions (VF) for the RHOSO control plane interface

This RHOSO release does not support the use of VFs for the RHOSO control plane interface.

Jira:OSPRH-8882

Verify that the os-net-config provider is ifcfg for any production deployment

Red Hat does not currently support **NMstate** as the **os-net-config provider**. Ensure that you have the setting **edpm_network_config_nmstate: false**, which is the default. This ensures that your environment uses the **ifcfg** provider.

Jira:OSPRH-11309

2.11.7. Control plane

2.11.7.1. Known issues

This part describes known issues in Red Hat OpenStack Services on OpenShift 18.0.

Control plane temporarily unavailable during minor update

During the minor update to 18.0 Feature Release 1, the RHOSO control plane temporarily becomes unavailable. API requests might fail with HTTP error codes, such as error 500. Alternatively, the API requests might succeed but the underlying life cycle operation fails. For example, a virtual machine (VM)

created with the **openstack server create** command during the minor update never reaches the **ACTIVE** state. The control plane outage is temporary and automatically recovers after the minor update is finished. The control plane outage does not affect the already running workload.

Jira:OSPRH-10790

2.11.8. Security and hardening

2.11.8.1. Bug fixes

This part describes bugs fixed in Red Hat OpenStack Services on OpenShift 18.0 that have a significant impact on users.

Custom configuration support for Key Manager (barbican) services

Before this update, there was an issue in **common_types.go**, preventing the **customServiceConfig** field in the custom resource definition (CRD) for the Key Manager service from being applied correctly. With this update, the issue has been resolved, allowing custom configuration to be correctly generated and applied.

Jira:OSPRH-10935

2.11.9. Storage

2.11.9.1. Known issues

This part describes known issues in Red Hat OpenStack Services on OpenShift 18.0.

extraMounts propagation to instance does not work when uniquePodNames is true

When **uniquePodNames** is **true**, every Cinder pod (and in general each component and service) is prefixed by a pseudo-random string. This affects the per-instance propagation, because the legacy method, based on **strings.TrimPrefix**, is not valid anymore.

In a DCN deployment, propagate secrets to pods by matching the instance AZ name.

Example 1 results in pods with names that match az0 getting the secret ceph-conf-az-0, pods with names that match az1 getting the secret ceph-conf-az-0, and so on. Example 1 works for Glance pods but only works for Cinder pods if **uniquePodNames** is **false**.

Workaround: Set **uniquePodNames** to **false** as shown in Example 2, until this issue is resolved. The **uniquePodNames** setting is only required if the storage back end uses NFS.

Example 1

```
apiVersion: core.openstack.org/v1beta1
kind: OpenStackControlPlane
spec:
  extraMounts:
  - extraVol:
    - extraVolType: Ceph
      mounts:
      - mountPath: /etc/ceph
        name: ceph0
        readOnly: true
```

```
      propagation:
      - az0
      volumes:
      - name: ceph0
        projected:
          sources:
          - secret:
              name: ceph-conf-az-0
    - extraVolType: Ceph
      mounts:
      - mountPath: /etc/ceph
        name: ceph1
        readOnly: true
      propagation:
      - az1
      volumes:
      - name: ceph1
        projected:
          sources:
          - secret:
              name: ceph-conf-az-1
```

Example 2

```
apiVersion: core.openstack.org/v1beta1
kind: OpenStackControlPlane
<...>
spec:
  cinder:
    uniquePodNames: false    # workaround https://issues.redhat.com/browse/OSPRH-11240
    enabled: true
    apiOverride:
      <...>
```

Jira:OSPRH-11240

2.12. RELEASE INFORMATION RHOSO 18.0.3

Review the known issues, bug fixes, and other release notes for this release of Red Hat OpenStack Services on OpenShift.

2.12.1. Advisory list

This release of Red Hat OpenStack Services on OpenShift (RHOSO) includes the following advisories:

RHBA-2024:9480

Release of components for RHOSO 18.0.3 (Feature Release 1)

RHSA-2024:9481

Moderate: Red Hat OpenStack Platform 18.0.3 (python-django) security update

RHBA-2024:9482

Release of containers for RHOSO 18.0.3 (Feature Release 1)

RHBA-2024:9483

Data plane Operators for RHOSO 18.0.3 (Feature Release 1)

RHBA-2024:9484

Release of operators for RHOSO 18.0.3 (Feature Release 1)

RHSA-2024:9485

Important: Control plane Operators for RHOSO 18.0.3 (Feature Release 1) security update

RHBA-2024:9486

Control plane Operators for RHOSO 18.0.3 (Feature Release 1)

2.12.2. Observability

2.12.2.1. New features

This part describes new features and major enhancements introduced in Red Hat OpenStack Services on OpenShift 18.0.

RabbitMQ metrics now in Prometheus.

With this update, RabbitMQ metrics are collected and stored in Prometheus. A new dashboard for displaying these metrics was added.

Jira:OSPRH-7610

Autoscaling improvements

Autoscaling has been updated to use the server_group metadata. This improves the stability of the autoscaling feature. For more information, see Autoscaling for instances

Jira:OSPRH-9202

2.12.2.2. Technology Previews

This part provides a list of all Technology Previews available in Red Hat OpenStack Services on OpenShift 18.0.

For information on the scope of support for Technology Preview features, see Technology Preview Features - Scope of Support.

VM power usage monitoring (Technology Preview)

With the integration of the kepler component, youi can expose the power usage of VM instances in a dashboard.

Jira:OSPRH-10006

2.12.3. Compute

2.12.3.1. New features

This part describes new features and major enhancements introduced in Red Hat OpenStack Services on OpenShift 18.0.

vGPUs enablement

This update introduces enhancements for mdev and vGPU.

Jira:OSPRH-63

2.12.3.2. Bug fixes

This part describes bugs fixed in Red Hat OpenStack Services on OpenShift 18.0 that have a significant impact on users.

NUMA resource tracking works correctly

With this release, a bug that causes NUMA resource tracking issues has been fixed. Previously, Libvirt reported all powered down CPUs on NUMA node 0 instead of on the correct NUMA node. Now, Nova caches the correct CPU topology before powering down any CPUs, fixing the resource tracking issues.

Jira:OSPRH-8712

2.12.3.3. Known issues

This part describes known issues in Red Hat OpenStack Services on OpenShift 18.0.

Setting hw-architecture or architecture on Image service (glance) image does not work as expected

In RHOSO 18.0, the image metadata prefilter is enabled by default. RHOSO does not support emulation of non-native architectures. As part of the introduction of emulation support upstream, the image metadata prefilter was enhanced to support the scheduling of instances based on the declared VM architecture, for example, **hw_architecture=x86_64**.

When nova was enhanced to support emulating non-native architecture by using image properties, a bug was introduced, because the native architecture was not reported as a trait by the virt driver.

Therefore, by default, support for setting **hw_architecture** or **architecture** on an image was rendered inoperable.

Workaround: To mitigate this bug, perform one of the following tasks:

- Unset the **architecture**/**hw_architecture** image property. RHOSO supports only one architecture, x86_64. There is no valid use case that requires this to be set for an RHOSO cloud, so all hosts will be x86_64.
- Disable the image metadata prefilter in the **CustomServiceConfig** section of the nova scheduler:

```
[scheduler]
image_metadata_prefilter=false
```

Jira:OSPRH-6215

Instances with ephemeral storage on NFS share stop working after Compute service restart

Compute service (nova) instances with ephemeral storage on NFS shares stop working as soon as the containerized Compute agent service restarts on the hypervisor host. That happens because of changed permissions of /var/lib/nova/instances.

Workaround: Manually restore permissions to the original values and avoid the service restarts.

Jira:OSPRH-10729

Compute service power management feature disabled by default

The Compute service (nova) power management feature is disabled by default. You can enable it with the following **nova-compute** configuration:

```
[libvirt]
cpu_power_management = true
cpu_power_management_strategy = governor
```

The default **cpu_power_management_strategy cpu_state** is currently unsupported. Restarting nova-compute causes all dedicated PCPUs on that host to be powered down, including ones used by instances. if the **cpu_state** strategy is used, those instances' CPUs will become unpinned.

Jira:OSPRH-10772

2.12.4. Data plane

2.12.4.1. New features

This part describes new features and major enhancements introduced in Red Hat OpenStack Services on OpenShift 18.0.

OpenStackAnsibleEE custom resource replaced with functionality in openstack-operator

Enhancement: The OpenStackAnsibleEE custom resource has been removed along with the openstack-ansibleee-operator. This functionality has been integrated into the openstack-operator to allow the direct creation of Kubernetes jobs without the unnecessary abstraction provided by the additional operator and associated custom resource.

Reason: The additional abstraction was unnecessary. This change reduces the amount of code that we need to maintain, along with reducing the number of CRD's and operators running in the cluster.

Result: Users can expect that there will no longer be any OpenStackAnsibleEE resources created when they deploy dataplane nodes. Instead, they will just see Kubernetes Jobs.

Existing OpenStackAnsibleEE resources will remain in the cluster for posterity, or if users no longer require them for historical reference, they can be deleted. Documentation is provided to cleanup unnecessary resources and operators.

Jira:OSPRH-7650

2.12.5. Networking

2.12.5.1. New features

This part describes new features and major enhancements introduced in Red Hat OpenStack Services on OpenShift 18.0.

Dynamic routing on data plane with FRR and BGP

This update introduces support of Free Range Routing (FRR) border gateway protocol (BGP) to provide dynamic routing capabilities on the RHOSO data plane.

Limitations:

- If you use dynamic routing, you must also use distributed virtual routing (DVR).
- If you use dynamic routing, you also use dedicated networker nodes.
- You can not use dynamic routing in an IPv6 deployment or a deployment that uses the Load-balancing service (octavia).

Jira:OSPRH-9298

2.12.5.2. Technology Previews

This part provides a list of all Technology Previews available in Red Hat OpenStack Services on OpenShift 18.0.

For information on the scope of support for Technology Preview features, see Technology Preview Features - Scope of Support.

Custom ML2 mechanism driver and SDN back end support (Technology Preview)

This update introduces a Technology Preview of the ability to integrate the Networking service (neutron) with a custom ML2 mechanism driver and software defined networking (SDN) back end components, instead of the default OVN mechanism driver and back end components.

Jira:OSPRH-3678

2.12.5.3. Known issues

This part describes known issues in Red Hat OpenStack Services on OpenShift 18.0.

Update to latest RHOSP 17.1 version before adopting

When performing adoption of a source environment which is older than RHOSP 17.1.4, the workloads experience a prolonged network connectivity disruption. Make sure to update the source environment at least to RHOSP 17.1.4 before adopting.

Jira:OSPRH-10283

octavia-operator does not enable anti affinity

The Load-balancing service (octavia) does not currently configure anti-affinity settings in Nova to prevent amphora VMs from being scheduled to the same compute node.

Workaround: Add the relevant setting to octavia through customConfig, as shown in the following example:

Example

```
# names will be dependent on the deployment
oc patch  -n openstack openstackcontrolplane openstack-galera-network-isolation --type=merge --
patch '
spec:
  octavia:
    template:
      octaviaHousekeeping:
        customServiceConfig: |
```

```
        [nova]
        enable_anti_affinity = true
    octaviaWorker:
      customServiceConfig: |
        [nova]
        enable_anti_affinity = true
    octaviaHealthManager:
      customServiceConfig: |
        [nova]
        enable_anti_affinity = true
'
```

Jira:OSPRH-10705

Adoption not supported with Load-balancing service (octavia) agents on networker nodes

Adoption of deployments that have Load-balancing service agents deployed on networker nodes is currently not supported.

Jira:OSPRH-10771

createDefaultLbMgmtNetwork and **manageLbMgmtNetworks** set to **false** when availability zones are defined

When setting a list of availability zones in the Octavia CR (in **spec.lbMgmtNetwork.availabilityZones**), the default values of the **spec.lbMgmtNetwork.createDefaultLbMgmtNetwork** and **spec.lbMgmtNetwork.manageLbMgmtNetworks** settings are incorrectly reset to **false**.

Workaround: When setting **availabilityZones** to a non-empty list in **spec.lbMgmtNetwork**, explicity set **createDefaultLbMgmtNetwork** and **manageLbMgmtNetworks to `true**.

Jira:OSPRH-11092

Adoption of combined Controller/Networker nodes not verified

Red Hat has not verified a process for adoption of a RHOSP 17.1 environment where Controller and Networker roles are composed together on Controller nodes. If your RHOSP 17.1 environment does use combined Controller/Networker roles on the Controller nodes, the documented adoption process will not produce the expected results.

Adoption of RHOSP 17.1 environments that use dedicated Networker nodes has been verified to work as documented.

Jira:OSPRH-11301

Legacy tripleo Networking services (neutron) after adoption

After edpm_tripleo_cleanup task, there are still legacy tripleo Networking service (neutron) services. These services are stopped after adoption, so the RHOSO services are not affected.

Workaround: Perform the following steps to remove the legacy services manually:

```
- Check tripleo neutron services list: systemctl list-unit-files --type service
- Remove tripleo services from: /etc/systemd/system/
```

Jira:OSPRH-11323

2.12.6. Network Functions Virtualization

2.12.6.1. Known issues

This part describes known issues in Red Hat OpenStack Services on OpenShift 18.0.

Do not use virtual functions (VF) for the RHOSO control plane interface

This RHOSO release does not support use of VFs for the RHOSO control plane interface.

Jira:OSPRH-8882

Verify that the **os-net-config** provider `ifcfg`for any production deployment

Red Hat does not currently support **NMstate** as the **os-net-config provider**.

Ensure that you have the setting **edpm_network_config_nmstate: false**, which is the default. This ensures that your environment uses the **ifcfg** provider.

Jira:OSPRH-11309

2.12.7. Control plane

2.12.7.1. Known issues

This part describes known issues in Red Hat OpenStack Services on OpenShift 18.0.

Control plane temporarily unavailable during minor update

During the minor update to 18.0 Feature Release 1, the Red Hat OpenStack Platform control plane temporarily becomes unavailable. API requests might fail with HTTP error codes, such as error 500. Alternatively, the API requests might succeed but the underlying life cycle operation fails. For example, a virtual machine (VM) created with the **openstack server create** command during the minor update never reaches the **ACTIVE** state. The control plane outage is temporary and automatically recovers after the minor update is finished. The control plane outage does not affect the already running workload.

Jira:OSPRH-10790

2.12.8. High availability

2.12.8.1. Technology Previews

This part provides a list of all Technology Previews available in Red Hat OpenStack Services on OpenShift 18.0.

For information on the scope of support for Technology Preview features, see Technology Preview Features - Scope of Support.

Instance high availability

RHOSO 18.0.3 (Feature Release 1) introduces a technology preview of instance high availability (instance HA). With instance HA, RHOSO can automatically evacuate and re-create instances on a different Compute node when a Compute node fails.

To use the instance HA technology preview in a test environment, see https://access.redhat.com/articles/7094761.

Do not use this technology preview in a production environment.

Jira:OSPRH-9902

2.12.8.2. Known issues

This part describes known issues in Red Hat OpenStack Services on OpenShift 18.0.

Possible database error before adoption

You might see a database error for system table **mysql.proc** if you run **mysqlcheck** before adopting the OSP database.

```
[...]
mysql.plugin                       OK
mysql.proc                         Needs upgrade
mysql.procs_priv                     OK
[...]
```

This error message is harmless and results from a system table's redo log that was not replicated correctly when the galera cluster was bootstrapped.

Workaround: You can remove the error by repairing the **mysql.proc** system table:

Example command

```
oc run mariadb-client ${MARIADB_CLIENT_ANNOTATIONS} -q --image ${MARIADB_IMAGE} -i --
rm --restart=Never -- \
    mysql -h $SOURCE_MARIADB_IP -u root -p"$SOURCE_DB_ROOT_PASSWORD" -e "repair
table mysql.proc;"
```

Example output

```
+------------+--------+----------+---------------------------------+
| Table      | Op     | Msg_type | Msg_text                        |
+------------+--------+----------+---------------------------------+
| mysql.proc | repair | info     | Running zerofill on moved table |
| mysql.proc | repair | status   | OK                              |
+------------+--------+----------+---------------------------------+
```

The table and its redo log are fixed and replicated across all galera nodes. Re-run the mysqlcheck and continue the adoption procedure.

Jira:OSPRH-10783

2.12.9. Storage

2.12.9.1. New features

This part describes new features and major enhancements introduced in Red Hat OpenStack Services on OpenShift 18.0.

Decommission a manilaShare back end

You can now decommission a **manilaShare** back end from RHOSO. When you delete a **manila-share**, a clean-up job runs to clean up the service list for the Shared File Systems service (manila). The output of the **openstack share pool list** command does not reflect storage pool changes. To update and display the latest statistics, you must restart the scheduler service. Perform the restart during scheduled downtime because it causes a minor disruption.

Jira:OSPRH-1099

Rebuilding volume-backed server

This release adds the support to rebuild a volume-backed server with the same or different image.

Jira:OSPRH-1391

Migrate director-deployed Ceph cluster to external Ceph cluster

With this update, after adoption from RHOSP 17.1 to a RHOSO 18.0 data plane, you can migrate a Director deployed Ceph cluster and turn it into an external Ceph cluster. The Ceph daemons deployed on the Controller nodes are migrated to a set of target nodes.

Jira:OSPRH-8369

Shared File Systems service (manila) support for VAST Data Platform

The Shared File Systems service now includes a storage driver to support VAST Data Platform. The driver allows provisioning and management of NFS shares and point-in-time backups through snapshots.

Jira:OSPRH-8821

Block Storage service (cinder) volume deletion

With this release, the Block Storage service RBD driver takes advantage of recent Ceph developments to allow RBD volumes to meet normal volume deletion expectations.

In previous releases, when the Block Storage service used an RBD (Ceph) volume back end, it was not always possible to delete a volume.

Jira:OSPRH-9477

2.12.9.2. Known issues

This part describes known issues in Red Hat OpenStack Services on OpenShift 18.0.

extraMounts propagation to instance does not work when uniquePodNames is true

When **uniquePodNames** is **true**, every Cinder Pod (and in general each component and service) is prefixed by a pseudo-random string. This affects the per-instance propagation, because the legacy method, based on **strings.TrimPrefix**, is not valid anymore.

In a DCN deployment, Red Hat recommends propagating secrets to pods by matching the instance AZ name.

Example 1 results in pods whose names match az0 getting the secret ceph-conf-az-0, pods whose names match az1 getting the secret ceph-conf-az-0, and so on. Example 1 works for Glance pods but only works for Cinder pods if **uniquePodNames** is **false**.

Workaround: Set **uniquePodNames** to false as shown in Example 2, until this bug is resolved. The **uniquePodNames** setting is only needed if the storage backend uses NFS.

Example 1

```
apiVersion: core.openstack.org/v1beta1
kind: OpenStackControlPlane
spec:
  extraMounts:
  - extraVol:
    - extraVolType: Ceph
      mounts:
      - mountPath: /etc/ceph
        name: ceph0
        readOnly: true
      propagation:
      - az0
      volumes:
      - name: ceph0
        projected:
          sources:
          - secret:
              name: ceph-conf-az-0
    - extraVolType: Ceph
      mounts:
      - mountPath: /etc/ceph
        name: ceph1
        readOnly: true
      propagation:
      - az1
      volumes:
      - name: ceph1
        projected:
          sources:
          - secret:
              name: ceph-conf-az-1
```

Example 2

```
apiVersion: core.openstack.org/v1beta1
kind: OpenStackControlPlane
<...>
spec:
  cinder:
    uniquePodNames: false    # workaround https://issues.redhat.com/browse/OSPRH-11240
    enabled: true
    apiOverride:
      <...>
```

Jira:OSPRH-11240

2.12.10. Upgrades and updates

2.12.10.1. New features

This part describes new features and major enhancements introduced in Red Hat OpenStack Services on OpenShift 18.0.

os-diff tool identifies differences between source 17.1 and adopted RHOSO environments

RHOSO ships the **os-diff tool**, which can help the operator find differences between the source RHOSP 17.1 environment configuration and the adopted RHOSO environment configuration.

Jira:OSPRH-1490

Baremetal adoption

You can now adopt baremetal RHOSP 17.1 environments into RHOSO environments.

Jira:OSPRH-2428

Adoption roll-back

You can now roll back a failed adoption of a RHOSP 17.1 control plane.

Jira:OSPRH-7817

2.12.10.2. Bug fixes

This part describes bugs fixed in Red Hat OpenStack Services on OpenShift 18.0 that have a significant impact on users.

Erroneous error messages no longer generated

Previously, the Keystone services and endpoints clean up step in the adoption procedure generated false errors when some services weren't deployed in the source cloud. The false errors are not generated anymore.

Jira:OSPRH-10174

2.13. RELEASE INFORMATION RHOSO 18.0.2

Review the known issues, bug fixes, and other release notes for this release of Red Hat OpenStack Services on OpenShift.

2.13.1. Advisory list

This release of Red Hat OpenStack Services on OpenShift (RHOSO) includes the following advisories:

RHBA-2024:8151

Release of containers for RHOSO 18.0.2

RHBA-2024:8152

Release of components for RHOSO 18.0.2

RHBA-2024:8153

Control plane Operators for RHOSO 18.02

RHBA-2024:8154

Data plane Operators for RHOSO 18.0.2

RHBA-2024:8155

Release of components for RHOSO 18.0.2

2.13.2. Compute

2.13.2.1. Bug fixes

This part describes bugs fixed in Red Hat OpenStack Services on OpenShift 18.0 that have a significant impact on users.

Fix for instances created before OpenStack Victoria

In OpenStack Victoria, the instance_numa_topology object was extended to enable mix cpus (pinned and unpinned cpus) in the same instance. Object conversion code was added to handle upgrades but did not account for flavors that have either **hw:mem_page_size** or **hw:numa_nodes** set with **hw:cpu_policy** not set to dedicated

As a result instances created before the Victoria release could not be started after an upgrade to Victoria.

With this update, non-pinned numa instances can be managed after an FFU from 16.2.

Jira:OSPRH-10183

2.13.2.2. Known issues

This part describes known issues in Red Hat OpenStack Services on OpenShift 18.0.

Setting hw-architecture or architecture on Image service (glance) image does not work as expected

In RHOSO 18.0, the image metadata prefilter is enabled by default. RHOSO does not support emulation of non-native architectures. As part of the introduction of emulation support upstream, the image metadata prefilter was enhanced to support the scheduling of instances based on the declared VM architecture, for example, **hw_architecture=x86_64**.

When nova was enhanced to support emulating non-native architecture by using image properties, a bug was introduced, because the native architecture was not reported as a trait by the virt driver.

Therefore, by default, support for setting **hw_architecture** or **architecture** on an image was rendered inoperable.

Workaround: To mitigate this bug, perform one of the following tasks:

- Unset the **architecture**/**hw_architecture** image property. RHOSO supports only one architecture, x86_64. There is no valid use case that requires this to be set for an RHOSO cloud, so all hosts will be x86_64.
- Disable the image metadata prefilter in the **CustomServiceConfig** section of the nova scheduler:

```
[scheduler]
image_metadata_prefilter=false
```

Jira:OSPRH-6215

Compute service power management feature disabled by default

The Compute service (nova) power management feature is disabled by default. You can enable it with the following **nova-compute** configuration:

```
[libvirt]
cpu_power_management = true
cpu_power_management_strategy = governor
```

The default **cpu_power_management_strategy cpu_state** is not supported at the moment due to a bug that causes NUMA resource tracking issues, as all disabled CPUs are reported on NUMA node 0 instead of on the correct NUMA node.

Jira:OSPRH-8712

QEMU process failure

A paused instance that uses local storage cannot be live migrated more than once. The second migration causes the QEMU process to crash and nova puts the instance to ERROR state.

Workaround: if feasible, unpause the instance temporarily, then pause it again before the second live migration.

It is not always feasible to unpause an instance. For example, suppose the instance uses a multi-attach cinder volume, and pause is used to limit the access to that volume to a single instance while the other is kept in paused state. In this case, unpausing the instance is not a feasible workaround.

Jira:OSPRH-8699

2.13.3. Data plane

2.13.3.1. Bug fixes

This part describes bugs fixed in Red Hat OpenStack Services on OpenShift 18.0 that have a significant impact on users.

The value for edpm_kernel_hugepages is reliably set on the kernel command line.

Before this update, the value for **edpm_kernel_hugepages** could be missing from the kernel commandline due to an an error in an ansible role that configures it. With this update, this problem is resolved, and no work arounds are required.

Jira:OSPRH-10007

2.13.4. Networking

2.13.4.1. Bug fixes

This part describes bugs fixed in Red Hat OpenStack Services on OpenShift 18.0 that have a significant impact on users.

Metadata rate-limiting feature

This update fixes a bug that prevented successful use of metadata rate-limiting. Metadata rate limiting is now available.

Jira:OSPRH-9569

2.13.4.2. Known issues

This part describes known issues in Red Hat OpenStack Services on OpenShift 18.0.

Router deletion problem and workaround

After an update to RHOSO 18.0.2, you cannot delete a pre-existing router as expected.

The following error is displayed in the CLI:

> Internal Server Error: The server has either erred or is incapable of performing the requested operation.

Also, the Neutron API logs include the following exception message:

> Could not find a service provider that supports distributed=False and ha=False

Workaround: Manually create a database register. In a SQL CLI:

```
$ use ovs_neutron;
$ insert into providerresourceassociations (provider_name, resource_id) values ("ovn", "
<router_id>");
```

Jira:OSPRH-10537

2.13.5. Network Functions Virtualization

2.13.5.1. Known issues

This part describes known issues in Red Hat OpenStack Services on OpenShift 18.0.

Do not use virtual functions (VF) for the RHOSO control plane interface

This RHOSO release does not support use of VFs for the RHOSO control plane interface.

Jira:OSPRH-8882

2.13.6. Storage

2.13.6.1. Known issues

This part describes known issues in Red Hat OpenStack Services on OpenShift 18.0.

OpenStack command output does not account for storage pool changes in the Shared File Systems service (manila)

The **openstack share pool list** command output does not account for storage pool changes, for example, changes to pool characteristics on back end storage systems, or removal of existing pools from the deployment. Provisioning operations are not affected by this issue. **Workaround:** Restart the scheduler service to reflect the latest statistics. Perform the restart during scheduled downtime because it causes a minor disruption.

Jira:OSPRH-1099

2.14. RELEASE INFORMATION RHOSO 18.0.1

Review the known issues, bug fixes, and other release notes for this release of Red Hat OpenStack Services on OpenShift.

2.14.1. Advisory list

This release of Red Hat OpenStack Services on OpenShift (RHOSO) includes the following advisories:

RHBA-2024:6773

Release of containers for RHOSO 18.0.1

RHBA-2024:6774

Release of components for RHOSO 18.0.1

RHBA-2024:6775

Moderate: Red Hat OpenStack Platform 18.0 (python-webob) security update

RHBA-2024:6776

Control plane Operators for RHOSO 18.0.1

RHBA-2024:6777

Data plane Operators for RHOSO 18.0.1

RHBA-2024:6778

Data plane Operators for RHOSO 18.0.1

2.14.2. Compute

2.14.2.1. Known issues

This part describes known issues in Red Hat OpenStack Services on OpenShift 18.0.

Setting hw-architecture or architecture on Image service (glance) image does not work as expected

In RHOSO 18.0, the image metadata prefilter is enabled by default. RHOSO does not support emulation of non-native architectures. As part of the introduction of emulation support upstream, the image metadata prefilter was enhanced to support the scheduling of instances based on the declared VM architecture, for example, **hw_architecture=x86_64**.

When nova was enhanced to support emulating non-native architecture by using image properties, a bug was introduced, because the native architecture was not reported as a trait by the virt driver.

Therefore, by default, support for setting **hw_architecture** or **architecture** on an image was rendered inoperable.

Workaround: To mitigate this bug, perform one of the following tasks:

- Unset the **architecture**/**hw_architecture** image property. RHOSO supports only one architecture, x86_64. There is no valid use case that requires this to be set for an RHOSO cloud, so all hosts will be x86_64.
- Disable the image metadata prefilter in the **CustomServiceConfig** section of the nova scheduler:

```
[scheduler]
image_metadata_prefilter=false
```

Jira:OSPRH-6215

Compute service power management feature disabled by default

The Compute service (nova) power management feature is disabled by default. You can enable it with the following **nova-compute** configuration:

```
[libvirt]
cpu_power_management = true
cpu_power_management_strategy = governor
```

The default **cpu_power_management_strategy cpu_state** is not supported at the moment due to a bug that causes NUMA resource tracking issues, as all disabled CPUs are reported on NUMA node 0 instead of on the correct NUMA node.

Jira:OSPRH-8712

QEMU process failure

A paused instance that uses local storage cannot be live migrated more than once. The second migration causes the QEMU process to crash and nova puts the instance to ERROR state.

Workaround: if feasible, unpause the instance temporarily, then pause it again before the second live migration.

It is not always feasible to unpause an instance. For example, suppose the instance uses a multi-attach cinder volume, and pause is used to limit the access to that volume to a single instance while the other is kept in paused state. In this case, unpausing the instance is not a feasible workaround.

Jira:OSPRH-8699

2.14.3. Data plane

2.14.3.1. Bug fixes

This part describes bugs fixed in Red Hat OpenStack Services on OpenShift 18.0 that have a significant impact on users.

Using the download-cache service no longer prevents Podman from pulling images for data plane deployment

Before this bug fix, if you included **download-cache** service in **spec.services** of the **OpenStackDataPlaneNodeSet**, a bug prevented Podman from pulling container images that are required by the data plane deployment.

With this bug fix, you can include **download-cache** service in **spec.services** of the **OpenStackDataPlaneNodeSet** and doing so does not prevent Podman from pulling the required container images.

Jira:OSPRH-9500

2.14.3.2. Known issues

This part describes known issues in Red Hat OpenStack Services on OpenShift 18.0.

Set edpm_kernel_args variable if you configure the Ansible variable edpm_kernel_hugepages

To configure the Ansible variable **edpm_kernel_hugepages** in the **ansibleVars** section of an **OpenStackDataPlaneNodeSet** CR, you must also set the **edpm_kernel_args** variable. If you do not need to configure **edpm_kernel_args** with a particular value, then set it to an empty string:

```
edpm_kernel_args: ""
```

Jira:OSPRH-10007

2.14.4. Networking

2.14.4.1. New features

This part describes new features and major enhancements introduced in Red Hat OpenStack Services on OpenShift 18.0.

Support for security group logging on Compute nodes

With this update, when security group logging is enabled, RHOSO writes logs to the data plane node that hosts the project instance. In the **/var/log/messages** file, each log entry contains the string, **acl_log**.

Jira:OSPRH-9248

2.14.4.2. Bug fixes

This part describes bugs fixed in Red Hat OpenStack Services on OpenShift 18.0 that have a significant impact on users.

Fixed delay between the oc patch command and update of OVN databases

Before this update, custom configuration settings applied with the **oc patch** command did not affect the Networking service (neutron) OVN databases until 10 minutes passed.

This update eliminates the delay.

Jira:OSPRH-9035

MAC_Binding aging functionality added back in RHOSO 18.0.1

The MAC_Binding aging functionality that was added in RHOSP 17.1.2 was missing from 18.0 GA. This update to RHOSO 18.0.1 adds it back.

Jira:OSPRH-8716

2.14.4.3. Known issues

This part describes known issues in Red Hat OpenStack Services on OpenShift 18.0.

Delayed OVN database update after oc patch command

Any custom configuration settings applied with the **oc patch** command do not affect the Networking service OVN databases until 10 minutes have passed.

Workaround: After you replace old pods by using the **oc patch** command, use the **oc delete pod** command to delete the new neutron pods.

The pod deletion forces a new configuration to be set without the delay issue.

Jira:OSPRH-7998

Metadata rate-limiting feature

Metadata rate-limiting is not available in RHOSO 18.0.1. A fix is in progress.

Jira:OSPRH-9569

2.14.5. Network Functions Virtualization

2.14.5.1. Bug fixes

This part describes bugs fixed in Red Hat OpenStack Services on OpenShift 18.0 that have a significant impact on users.

DPDK bonds are now validated in os-net-config

Previously, when OVS or DPDK bonds were configured with a single port, no error was reported despite the ovs bridge not being in the right state. With this update **os-net-config** reports an error if the bond has a single interface.

Jira:OSPRH-9307

2.14.5.2. Known issues

This part describes known issues in Red Hat OpenStack Services on OpenShift 18.0.

Do not use virtual functions (VF) for the RHOSO control plane interface

This RHOSO release does not support use of VFs for the RHOSO control plane interface.

Jira:OSPRH-8882

2.14.6. Storage

2.14.6.1. Bug fixes

This part describes bugs fixed in Red Hat OpenStack Services on OpenShift 18.0 that have a significant impact on users.

Image import no longer remains in importing state after conversion with ISO image format

Before this update, when you used image conversion with the ISO image format, the image import operation remained in an "importing" state.

Now the image import operation does not remain in an "importing" state.

Jira:OSPRH-8580

2.15. RELEASE INFORMATION RHOSO 18.0 GA

Review the known issues, bug fixes, and other release notes for this release of Red Hat OpenStack Services on OpenShift.

2.15.1. Advisory list

This release of Red Hat OpenStack Services on OpenShift (RHOSO) includes the following advisories:

RHEA-2024:5245

Release of components for RHOSO 18.0

RHEA-2024:5246

Release of containers for RHOSO 18.0

RHEA-2024:5247

Data plane Operators for RHOSO 18.0

RHEA-2024:5248

Control plane Operators for RHOSO 18.0

RHEA-2024:5249

Release of components for RHOSO 18.0

2.15.2. Observability

2.15.2.1. New features

This part describes new features and major enhancements introduced in Red Hat OpenStack Services on OpenShift 18.0.

Deploy metric storage with Telemetry Operator

The Telemetry Operator now supports deploying and operating Prometheus by using the **cluster-observability-operator** through a MonitoringStack resource.

Jira:OSPRH-1896

Expanded interaction with metrics and alarms

You can now use the **openstack metric** and **openstack alarm** commands in the OpenStack CLI to interact with metrics and alarms. These commands are useful for troubleshooting.

Jira:OSPRH-2892

Ceilometer uses TCP publisher to expose data for Prometheus

Ceilometer can now use the TCP publisher to publish metric data to sg-core, which exposes them for scraping by Prometheus.

Jira:OSPRH-2957

Prometheus replaces Gnocchi for metrics storage and metrics-based autoscaling

In RHOSO 18.0, Prometheus replaces Gnocchi for metrics and metrics-based autoscaling.

Jira:OSPRH-3057

Compute node log collection

RHOSO uses the Cluster Logging Operator (**cluster-logging-operator**) to collect and centrally store logs from OpenStack Compute nodes.

Jira:OSPRH-802

Graphing dashboards for OpenStack metrics

The Red Hat OpenShift Container Platform (RHOCP) console UI now provides graphing dashboards for OpenStack Metrics.

Jira:OSPRH-824

2.15.3. Compute

2.15.3.1. New features

This part describes new features and major enhancements introduced in Red Hat OpenStack Services on OpenShift 18.0.

The compute service now supports native Secure RBAC

In RHOSP 17.1 secure role-based access control was implemented using custom policy. In RHOSO-18.0.0 this is implemented using nova native support for SRBAC. As a result all OpenStack deployments support the ADMIN, MEMBER and READER roles by default.

Jira:OSPRH-1505

Setting the hostname of the Compute service (nova) instance by using the Compute service API microversions 2.90 and 2.94

This enhancement enables you to set the hostname of the Compute service (nova) instance by using the Compute service API microversions 2.90 and 2.94 that are now included in the 18.0 release of RHOSO.

API microversion 2.90 enables you to specify an optional hostname when creating, updating, or rebuilding an instance. This is a short name (without periods), and it appears in the metadata available to the guest OS, either through the metadata API or on the configuration drive. If installed and configured in the guest, **cloud-init** uses this optional hostname to set the guest hostname.

API microversion 2.94 extends microversion 2.90 by enabling you to specify fully qualified domain names (FQDN) wherever you specify the hostname. When using an FQDN as the instance hostname, you must set the **[api]dhcp_domain** configuration option to the empty string in order for the correct FQDN to appear in the hostname field in the metadata API.

Jira:OSPRH-17

Manage dedicated CPU power state

You can now configure the **nova-compute** service to manage dedicated CPU power state by setting [libvirt]cpu_power_management to True.

This feature requires the Compute service to be set with [compute]cpu_dedicated_set. With that setting, all dedicated CPUs are powered down until they are used by an instance. They are powered up when an instance using them is booted. If power management is configured but [compute]cpu_dedicated_set isn't set, then the compute service will not start.

By default, the power strategy offlines CPUs when powering down and onlines the CPUs on powering up, but another strategy is possible. Set [libvirt]cpu_power_management_strategy=governor to instead

use governors, and use [libvirt]cpu_power_governor_low [libvirt]cpu_power_governor_high to direct which governors to use in online and offline mode (performance and powersave).

Jira:OSPRH-18

Evacuate to STOPPED with v2.95

Starting with the v2.95 micro version, any evacuated instance will be stopped at the destination. Operators can still continue using the previous behaviour by selecting a microversion below v2.95. Prior to v2.95, if the VM was active prior to the evacuation, it was restored to the active state following a failed evacuation. If the workload encountered I/O corruption as a result of the hypervisor outage, this could potentially make recovery effort harder or cause further issues if the workload was a clustered application that tolerated the failure of a single VM. For this reason, it is considered safer to always evacuate to Stopped and allow the tenant to decide how to recover the VM.

Jira:OSPRH-184

Compute service hostname change

If you start the Compute service (nova) and your Compute host detects a name change, you must know the reason for the change of the host names. When you resolve the issue, you must restart the Compute service.

Jira:OSPRH-20

Create a neutron port without an IP address if the port requires only L2 network connectivity

You can now create an instance with a **non-deferred** port that has no fixed IP address if the network back end has L2 connectivity.

In previous releases of RHOSP, all neutron ports were required to have an IP address. The IP address assignment could be immediate (default) or deferred for L3 routed networks. In RHOSO 18.0, that requirement has been removed. You can now create a neutron port without an IP address if the port requires only L2 network connectivity.

To use this feature, set **ip_allocation = 'none'** on the neutron port before passing it to nova to use when creating a VM instance or attaching the port to an existing instance.

Jira:OSPRH-57

New enlightenments to the libvirt XML for Windows guests in RHOSO 18.0.0

This update adds the following enlightenments to the libvirt XML for Windows guests:

- vpindex
- runtime
- synic
- reset
- frequencies
- tlbflush
- ipi

This adds to the list of existing enlightenments:

- relaxed
- vapic
- spinlocks retries
- vendor_id spoofing

Jira:OSPRH-58

New default for managing instances on NUMA nodes

In RHOSP 17.1.4, the default was to pack instances on NUMA nodes.

In RHOSO 18.0, the default has been changed to balance instances across NUMA nodes. To change the default, and pack instances on NUMA nodes, set

```
[compute]
packing_host_numa_cells_allocation_strategy = True
```

in both the scheduler and compute node nova.conf

Jira:OSPRH-59

Rebuild a volume-backed instance with a different image

This update adds the ability to rebuild a volume-backed instance from a different image.

Before this update, you could only rebuild a volume-backed instance from the original image in the boot volume.

Now you can rebuild the instance after you have reimaged the boot volume on the cinder side.

This feature requires API microversion 2.93 or later.

Jira:OSPRH-66

Archive 'task_log' database records

This enhancement adds the **--task-log** option to the **nova-manage db archive_deleted_rows** CLI. When you use the **--task-log** option, the **task_log** table records get archived while archiving the database. This option is the default in the nova-operator database purge cron job. Previously, there was no method to delete the **task_log** table without manual database modification.

You can use the **--task-log** option with the **--before** option for records that are older than a specified **<date>**. The **updated_at** field is compared to the specifed **<date>** to determine the age of a **task_log** record for archival.

If you configure **nova-compute** with **[DEFAULT]instance_usage_audit = True**, the **task_log** database table maintains an audit log of **--task-log** use.

Jira:OSPRH-68

Support for virtual IOMMU device

The Libvirt driver can add a virtual IOMMU device to guests. This capability applies to x86 hosts that use

the Q35 machine type. To enable the capability, provide the **hw:viommu_model** extra spec or equivalent image metadata property **hw_viommu_model**.The following values are supported: **intel**, **smmuv3**, **virtio**, **auto**. The default value is **auto**, which automatically selects **virtio**.

NOTE

Due to the possible overhead introduced with vIOMMU, enable this capability only for required workloads.

Jira:OSPRH-69

More options for the server unshelve command

With this update, new options are added to the **server unshelve** command in RHOSO 18.0.0.

The **--host** option allows administrators to specify a destination host. The **--no-availability-zone** option allows administrators to specify the availability zone. Both options require the server to be in the **SHELVED_OFFLOADED** state and the Compute API version to be **2.91** or greater.

Jira:OSPRH-74

Support for the bochs libvirt video model

This release adds the ability to use the **bochs** libvirt video model. The **bochs** libvirt video model is a legacy-free video model that is best suited for UEFI guests. In some cases, it can be usable for BIOS guests, such as when the guest does not depend on direct VGA hardware access.

Jira:OSPRH-76

Schedule archival and purge of deleted rows from Compute service (nova) cells

The nova-operator now schedules a periodic job for each Compute service (nova) cell to archive and purge the deleted rows from the cell database. The frequency of the job and the age of the database rows to archive and purge can be fine tuned in the **{{OpenStackControlPlane.spec.nova.template.cellTemplates[].dbPurge}}** structure for each cell in the cellTemplates.

Jira:OSPRH-86

2.15.3.2. Bug fixes

This part describes bugs fixed in Red Hat OpenStack Services on OpenShift 18.0 that have a significant impact on users.

Migrating paused instance no longer generates error messages

Before this update, live migration of a paused instance with live_migration_permit_post_copy=True in nova.conf caused the libvirt driver to erroneously generate error messages similar to [1].

Now the error message is not generated when you live migrate a paused instance with live_migration_permit_post_copy=True.

[1] Error message example: "Live Migration failure: argument unsupported: post-copy migration is not supported with non-live or paused migration: libvirt.libvirtError: argument unsupported: post-copy migration is not supported with non-live or paused migration."

Jira:OSPRH-41

No network block device (NBD) live migration with TLS enabled

In RHOSO 18.0 Beta, a bug prevents you from using a network block device (NBD) to live migrate storage between Compute nodes with TLS enabled. See https://issues.redhat.com/browse/OSPRH-6931.

This has now been resolved and live migration with TLS enabled is supported with local storage.

Jira:OSPRH-6740

Cannot delete instance when cpu_power_managment is set to true

In the RHOSO 18.0.0 beta release, a known issue was discovered preventing the deletion of an instance shortly after it was created if power management was enabled.

This has now been fixed in the RHOSO 18.0.0 GA release

Jira:OSPRH-7103

2.15.3.3. Technology Previews

This part provides a list of all Technology Previews available in Red Hat OpenStack Services on OpenShift 18.0.

For information on the scope of support for Technology Preview features, see Example.

Technology preview of PCI device tracking in Placement service

RHOSO 18.0.0 introduces a technology preview of the ability to track PCI devices in the OpenStack Placement service.

Tracking PCI devices in the Placement service enables you to use granular quotas on PCI devices when combined with the Unified Limits Technology Preview.

PCI tracking in the Placement service is disabled by default and is limited to flavor-based PCI passthrough. Support for the Networking service (neutron) SRIOV ports is not implemented, but is required before this feature is fully supported.

Jira:OSPRH-19

Use of Identity service (Keystone) unified limits in the Compute service (nova)

This RHOSO release supports Identity service unified limits in the Compute service. Unified limits centralize management of resource quota limits in the Identity service (Keystone) and enable flexibility for users to manage quota limits for any Compute service resource being tracked in the Placement service.

Jira:OSPRH-70

2.15.3.4. Removed functionality

This part provides an overview of functionality that has been *removed* in Red Hat OpenStack Services on OpenShift 18.0.

Removed functionality is no longer supported in this product and is not recommended for new deployments.

Keypair generation removed from RHOSO 18

Keypair generation was deprecated in RHOSP 17 and has been removed from RHOSO 18. Now you need to precreate the keypair by the SSH command line tool **ssh-keygen** and then pass the public key to the nova API.

Jira:OSPRH-67

i440fx PC machine type no longer tested or supported

In RHOSP 17, the i440fx PC machine type, pc-i440fx, was deprecated and Q35 became the default machine type for x86_64.

In RHOSP 18, the i440fx PC machine type is no longer tested or supported.

The i440fx PC machine type is still available for use under a support exception for legacy applications that cannot function with the Q35 machine type. If you have such a workload, contact Red Hat support to request a support exception.

With the removal of support for the i440fx PC machine type from RHOSP, you cannot use pc-i440fx to certify VNFs or third-party integrations. You must use the Q35 machine type.

Jira:OSPRH-7373

Unsupported: vDPA and Hardware offload OVS are unsupported

Hardware offload OVS consists of processing network traffic in hardware with the kernel swtichdev and tcflower protocols.

vDPA extends Hardware offload OVS by providing a vendor-neutral virtio net interface to the guest, decoupling the workload from the specifics fo the host hardware instead of presenting a vendor-specific virtual function.

Both Hardware offload OVS and vDPA are unsupported in RHOSO 18.0 with no upgrade path available for existing users.

At this time there is no plan to reintroduce this functionality or continue to invest in new features related to vdpa or hardware offloaded ovs.

If you have a business requirement for these removed features, please reach out to Red Hat support or your partner and Technical Account Manager so that Red Hat can reassess the demand for these features for a future RHOSO release.

Jira:OSPRH-7829

2.15.3.5. Known issues

This part describes known issues in Red Hat OpenStack Services on OpenShift 18.0.

Setting hw-architecture or architecture on Image service (glance) image does not work as expected

In RHOSO 18.0, the image metadata prefilter is enabled by default. RHOSO does not support emulation of non-native architectures. As part of the introduction of emulation support upstream, the image metadata prefilter was enhanced to support the scheduling of instances based on the declared VM architecture, for example **hw_architecture=x86_64**.

When nova was enhanced to support emulating non-native architecture via image properties, a bug was introduced, because the native architecture was not reported as a trait by the virt driver.

Therefore, by default, support for setting **hw_architecture** or **architecture** on an image was rendered inoperable.

To mitigate this bug, you have two choices:

- Unset the **architecture**/**hw_architecture** image property. RHOSO supports only one architecture, x86_64. There is no valid use case that requires this to be set for an RHOSO cloud, so all hosts will be x86_64.
- Disable the image metadata prefilter in the **CustomServiceConfig** section of the nova scheduler:

```
[scheduler]
image_metadata_prefilter=false
```

Jira:OSPRH-6215

QEMU process failure

A paused instance that uses local storage cannot be live migrated more than once. The second migration causes the QEMU process to crash and nova puts the instance to ERROR state.

Workaround: if feasible, unpause the instance temporarily then pause it again before the second live migration.

It is not always feasible to unpause an instance. For example, suppose the instance uses a multi-attach cinder volume, and pause is used to limit the access to that volume to a single instance while the other is kept in paused state. In this case, unpausing the instance is not a feasible workaround.

Jira:OSPRH-8699

Compute service power management feature disabled by default

The Compute service (nova) power management feature is disabled by default. You can enable it with the following nova-compute configuration.

```
[libvirt]
cpu_power_management = true
cpu_power_management_strategy = governor
```

The default cpu_power_management_strategy cpu_state is not supported at the moment due to a bug that causes NUMA resource tracking issues, because all disabled CPUs are reported on NUMA node 0 instead of on the correct NUMA node.

Jira:OSPRH-8712

2.15.4. Data plane

2.15.4.1. Known issues

This part describes known issues in Red Hat OpenStack Services on OpenShift 18.0.

Using the **download-cache** service prevents Podman from pulling images for data plane deployment

Do not list the **download-cache** service in spec.services of the **OpenStackDataPlaneNodeSet**. If you list **download-cache** in **OpenStackDataPlaneNodeSet**, Podman can not pull the container images required by the data plane deployment.

Workaround: Omit the **download-cache** service from the default services list in **OpenStackDataPlaneNodeSet**.

Jira:OSPRH-9500

2.15.5. Hardware Provisioning

2.15.5.1. Bug fixes

This part describes bugs fixed in Red Hat OpenStack Services on OpenShift 18.0 that have a significant impact on users.

Increased EFI partition size

Before RHOSP 17.1.4, the EFI partition size of an overcloud node was 16MB. With this update, the image used for provisioned EDPM nodes now has an EFI partition size of 200MB to align with RHEL and to accommodate firmware upgrades.

Jira:OSPRH-6691

2.15.6. Networking

2.15.6.1. New features

This part describes new features and major enhancements introduced in Red Hat OpenStack Services on OpenShift 18.0.

Octavia Operator availability zones

The Octavia Management network created and managed by the Octavia operator requires that the OpenStack routers and networks are scheduled on the OVN controller on the OpenShift worker nodes.

If the OpenStack Networking Service (neutron) is configured with non-default availability zones, the OVN controller pod on the OpenShift worker and Octavia must be configured with the same availability zone.

Example:

```
ovn:
  template:
    ovnController:
      external-ids:
        availability-zones:
          - zone1
octavia:
  template:
    lbMgmtNetwork:
      availabilityZones:
        zone1
```

-

Jira:OSPRH-6901

2.15.6.2. Bug fixes

This part describes bugs fixed in Red Hat OpenStack Services on OpenShift 18.0 that have a significant impact on users.

OVN pod no longer goes into loop due to NIC Mapping

When using a large number of NIC mappings, OVN could go into a creation loop. This is now fixed

Jira:OSPRH-7480

2.15.6.3. Technology Previews

This part provides a list of all Technology Previews available in Red Hat OpenStack Services on OpenShift 18.0.

For information on the scope of support for Technology Preview features, see Technology Preview Features - Scope of Support.

QoS minimum bandwidth policy (technology preview)

In RHOSO 18.0.0, a technology preview is available for the Networking service (neutron) for QoS minimum bandwidth for placement reporting and scheduling.

Jira:OSPRH-507

Load-balancing service (Octavia) support of multiple VIP addresses

This update adds a technology preview of support for multiple VIP addresses allocated from the same Neutron network for the Load-balancing service.

You can now specify additional subnet_id/ip_address pairs for the same VIP port. This makes it possible to configure the Load-balancing service with both IPv4 and IPv6 exposed to both public and private subnets.

Jira:OSPRH-2154

2.15.6.4. Known issues

This part describes known issues in Red Hat OpenStack Services on OpenShift 18.0.

Delayed OVN database update after `oc patch` command

Any custom configuration settings applied with 'oc patch ...' command do not affect neutron ovn databases until 10 minutes have passed.

Workaround: After you replace old pods using the **oc patch ...** command, delete the new neutron pod(s) manually using **oc delete pod ...** command.

The pod deletion forces a new configuration to be set without the delay issue.

Jira:OSPRH-7998

MAC_Binding aging functionality missing in RHOSO 18.0.0

The MAC_Binding aging functionality that was added in RHOSP 17.1.2 is missing from 18.0 GA. A fix is in progress.

Jira:OSPRH-8716

10-minute delay between 'oc patch` command and update of OVN databases

Custom configuration settings applied with the 'oc patch' command do not affect the Networking service (neutron) OVN databases until 10 minutes have passed.

Workaround: After the old Networking service pods are replaced new pods after an 'oc patch' command operation, delete the new Networking service pods manually using the 'oc delete pod' command.

This deletion forces a new configuration to be set without the delay issue.

Jira:OSPRH-9035

Metadata rate-limiting feature

Metadata rate-limiting is not available in RHOSO 18.0.0. A fix is in progress.

Jira:OSPRH-9569

2.15.7. Network Functions Virtualization

2.15.7.1. New features

This part describes new features and major enhancements introduced in Red Hat OpenStack Services on OpenShift 18.0.

AMD CPU powersave profiles

A power save profile, cpu-partitioning-powersave, was introduced in Red Hat Enterprise Linux 9 (RHEL 9), and made available in Red Hat OpenStack Platform (RHOSP) 17.1.3.

This TuneD profile is the base building block for saving power in NFV environments. RHOSO 18.0 adds cpu-partitioning-powersave support for AMD CPUs.

Jira:OSPRH-2268

2.15.7.2. Bug fixes

This part describes bugs fixed in Red Hat OpenStack Services on OpenShift 18.0 that have a significant impact on users.

Physical function (PF) MAC address now matches between VM instances and SR-IOV physical functions (PFs)

This update fixes a bug that caused a PF MAC address mismatch between VM instances and SR-IOV PFs (Networking service ports with **vnic-type** set to **direct-physical**).

In the RHOSO 18.0 Beta release, a bug in the Compute service (nova) prevented the MAC address of SR-IOV PFs from being updated correctly when attached to a VM instance.

Now the MAC address of the PF is set on the corresponding neutron port.

Jira:OSPRH-7085

2.15.7.3. Technology Previews

This part provides a list of all Technology Previews available in Red Hat OpenStack Services on OpenShift 18.0.

For information on the scope of support for Technology Preview features, see Example.

In RHOSO 18.0, a technology preview is available for the nmstate provider back-end in os-net-config.

This technology preview of nmstate and NIC hardware offload has known issues that make it unsuitable for production use. For production, use the **openstack-network-scripts** package rather than nmstate and NetworkManager.

There is a production-ready native nmstate mode you can select during installation, but network configuration, which must be provided in nmstate format, is not backwards-compatible with templates from TripleO. It also lacks certain features that os-net-config provides, such as NIC name mapping or DSCP configuration.

Jira:OSPRH-2273

Data Center Bridge (DCB)-based QoS settings technology preview

Specific to port/interface, DCB-based QoS settings are now available as a technology preview as part of the **os-net-config** tool's network configuration template. For more information, see this knowledge base article: https://access.redhat.com/articles/7062865

Jira:OSPRH-2889

2.15.7.4. Deprecated functionality

This part provides an overview of functionality that has been *deprecated* in Red Hat OpenStack Services on OpenShift 18.0.

Deprecated functionality will likely not be supported in future major releases of this product and is not recommended for new deployments.

TimeMaster service is deprecated in RHOSO 18.0

In RHOSO 18.0, support for the TimeMaster service is deprecated. Bug fixes and support are provided through the end of the RHOSO 18.0 lifecycle, but no new feature enhancements will be made.

Jira:OSPRH-8244

2.15.7.5. Known issues

This part describes known issues in Red Hat OpenStack Services on OpenShift 18.0.

Do not use virtual functions (VF) for the RHOSO control plane interface

This RHOSO release does not support use of VFs for the RHOSO control plane interface.

Jira:OSPRH-8882

Bonds require minimum of two interfaces

If you configure an OVS or DPDK bond, always configure at least two interfaces. Bonds with only a single interface do not function as expected.

Jira:OSPRH-9307

2.15.8. High availability

2.15.8.1. New features

This part describes new features and major enhancements introduced in Red Hat OpenStack Services on OpenShift 18.0.

Password rotation

This update introduces the ability to generate and rotate OpenStack database passwords.

Jira:OSPRH-92

2.15.9. Storage

2.15.9.1. New features

This part describes new features and major enhancements introduced in Red Hat OpenStack Services on OpenShift 18.0.

Shared File Systems support for scalable CephFS-NFS

The Shared File Systems service (manila) now supports a scalable CephFS-NFS service. In earlier releases of Red Hat OpenStack Platform, only active/passive high-availability that was orchestrated with Director, using Pacemaker/Corosync, was supported. With this release, deployers can create active/active clusters of CephFS-NFS and integrate these clusters with the Shared File Systems service for improved scalability and high availability for NFS workloads.

Jira:OSPRH-1024

Block Storage service (cinder) volume deletion

With this release, the Block Storage service RBD driver takes advantage of recent Ceph developments to allow RBD volumes to meet normal volume deletion expectations.

In previous releases, when the Block Storage service used an RBD (Ceph) volume back end, it was not always possible to delete a volume.

Jira:OSPRH-1777

project_id in API URLs now optional

You are no longer required to include **project_id** in Block Storage service (cinder) API URLs.

Jira:OSPRH-1787

Dell PowerStore storage systems driver

A new share driver has been added to support Dell PowerStore storage systems with the Shared File Systems service (Manila) service.

Jira:OSPRH-4425

Dell PowerFlex storage systems driver

A new share driver has been added to support Dell PowerFlex storage systems with the Shared File Systems service (Manila) service.

Jira:OSPRH-4426

openstack-must-gather SOS report support

You can now collect diagnostic information about your RHOSO deployment using the openstack-must-gather.

You can retrieve SOS reports for both the RHOCP control plane and RHOSO data plane nodes using a single command, and options are available to dump specific information related to a particular deployed service.

Jira:OSPRH-866

2.15.9.2. Bug fixes

This part describes bugs fixed in Red Hat OpenStack Services on OpenShift 18.0 that have a significant impact on users.

Key Manager service configuration fix enables Image service image signing and verification

With this fix, the Image service (glance) is automatically configured to interact with the Key Manager service (barbican), and you can now perform encrypted image signing and verification.

Jira:OSPRH-7155

Fixed faulty share creation in the NetApp ONTAP driver when using SVM scoped accounts

Due to a faulty kerberos enablement check upon shares creation, the NetApp ONTAP driver failed to create shares when configured with SVM scoped accounts. A fix has been committed to openstack-manila and shares creation should work smoothly.

Jira:OSPRH-8044

2.15.9.3. Technology Previews

This part provides a list of all Technology Previews available in Red Hat OpenStack Services on OpenShift 18.0.

For information on the scope of support for Technology Preview features, see Technology Preview Features - Scope of Support.

Deployment and scale of Object Storage service

This feature allows for the deployment and scale of Object Storage service (swift) data on data plane nodes. This release of the feature is a technology preview.

Jira:OSPRH-1307

2.15.9.4. Known issues

This part describes known issues in Red Hat OpenStack Services on OpenShift 18.0.

RGW does not pass certain Tempest object storage metadata tests

Red Hat OpenStack Services on OpenShift 18.0 supports Red Hat Ceph Storage 7. Red Hat Ceph Storage 7 RGW does not pass certain Tempest object storage metadata tests as tracked by the following Jiras:

https://issues.redhat.com/browse/RHCEPH-6708https://issues.redhat.com/browse/RHCEPH-9119https://issues.redhat.com/browse/RHCEPH-9122https://issues.redhat.com/browse/RHCEPH-4654

Jira:OSPRH-7464

Image import remains in importing state after conversion with ISO image format

When you use image conversion with the ISO image format, the image import operation remains in an "importing" state.

> *Workaround:* If your deployment supports uploading images in ISO format, you can use the `image-create` command to upload ISO images as shown in the following example (instead of using image conversion with the `image-create-via-import` command).

Example:

```
glance image-create \
--name <iso_image> \
--disk-format iso \
--container-format bare \
--file <my_file.iso>
```

- Replace **<iso_image>** with the name of your image.
- Replace **<my_file.iso>** with the file name for your image.

Jira:OSPRH-8580

2.15.10. Dashboard

2.15.10.1. New features

This part describes new features and major enhancements introduced in Red Hat OpenStack Services on OpenShift 18.0.

Hypervisor status now includes vCPU and pCPU information

Before this update, pCPU usage was excluded from the hypervisor status in the Dashboard service (horizon) even if the **cpu_dedicated_set** configuration option was set in the **nova.conf** file. This enhancement uses the Placement API to display information about vCPUs and pCPUs. You can view vCPU and pCPU usage diagrams under the **Resource Providers Summary** and find more information on vCPUs and pCPUs on the new **Resource provider** tab at the Hypervisors panel.

Jira:OSPRH-1516

With this update, you can now customize the OpenStack Dashboard (horizon) container.

The customization can be performed by using the extra mounts feature to add or change files inside of the Dashboard container.

Jira:OSPRH-5644

TLS everywhere in RHOSO Dashboard Operator

With this update, the RHOSO Dashboard (horizon) Operator automatically configures TLS-related configuration settings.

These settings include certificates and response headers when appropriate, including the secure cookies and HSTS headers for serving over HTTPS.

Jira:OSPRH-5882

2.15.10.2. Bug fixes

This part describes bugs fixed in Red Hat OpenStack Services on OpenShift 18.0 that have a significant impact on users.

Host spoofing protective measure

Before this update, the hosts configuration option was not populated with the minimum hosts necessary to protect against host spoofing.

With this update, the hosts configuration option is now correctly populated.

Jira:OSPRH-5832

Dashboard service operators now include HSTS header

Before this update, HSTS was only enabled in Django through the Dashboard service (horizon) application. However, user HTTPS sessions were going through the OpenShift route, where HSTS was disabled. With this update, HSTS is enabled on the OpenShift route.

Jira:OSPRH-7367

2.16. RELEASE INFORMATION RHOSO 18.0 BETA

2.16.1. Advisory list

This release of Red Hat OpenStack Services on OpenShift (RHOSO) includes the following advisories:

RHEA-2024:3646

RHOSO 18.0 Beta container images, data plane 1.0 Beta

RHEA-2024:3647

RHOSO 18.0 Beta container images, control plane 1.0 Beta

RHEA-2024:3648

RHOSO 18.0 Beta service container images

RHEA-2024:3649

RHOSO 18.0 Beta packages

2.16.2. Compute

2.16.2.1. New features

This part describes new features and major enhancements introduced in Red Hat OpenStack Services on OpenShift 18.0.

You can schedule archival and purge of deleted rows from Compute service (nova) cells

The nova-operator now schedules a periodic job for each Compute service (nova) cell to archive and purge the deleted rows from the cell database. The frequency of the job and the age of the database rows to archive and purge can be fine tuned in the **{{OpenStackControlPlane.spec.nova.template.cellTemplates[].dbPurge}}** structure for each cell in the cellTemplates.

Jira:OSPRH-86

2.16.2.2. Deprecated functionality

This part provides an overview of functionality that has been *deprecated* in Red Hat OpenStack Services on OpenShift 18.0.

Deprecated functionality will likely not be supported in future major releases of this product and is not recommended for new deployments.

i440fx PC machine type no longer tested or supported

In RHOSP 17, the i440fx PC machine type, pc-i440fx, was deprecated and Q35 became the default machine type for x86_64.

In RHOSP 18, the i440fx PC machine type is no longer tested or supported.

The i440fx PC machine type is still available for use under a support exception for legacy applications that cannot function with the Q35 machine type. If you have such a workload, contact Red Hat support to request a support exception.

With the removal of support for the i440fx PC machine type from RHOSP, you cannot use pc-i440fx to certify VNFs or third-party integrations. You must use the Q35 machine type.

Jira:OSPRH-7373

2.16.2.3. Known issues

This part describes known issues in Red Hat OpenStack Services on OpenShift 18.0.

No network block device (NBD) live migration with TLS enabled

In RHOSO 18.0 Beta, a bug prevents you from using a network block device (NBD) to live migrate storage between Compute nodes with TLS enabled. See https://issues.redhat.com/browse/OSPRH-6931.

This issue only affects storage migration when TLS is enabled. You can live migrate storage with TLS not enabled.

Jira:OSPRH-6740

Do not mix NUMA and non-NUMA instances on same Compute host

Instances without a NUMA topology should not coexist with NUMA instances on the same host.

Jira:OSPRH-83

Cannot delete instance when **cpu_power_managment** is set to **true**

When an instance is first started and the host core state is changed there is a short time period where it cannot be updated again. During this period instance deletion can fail. if this happens a second delete attempt should succeed after a short delay of a few seconds.

Jira:OSPRH-7103

2.16.3. Networking

2.16.3.1. Known issues

This part describes known issues in Red Hat OpenStack Services on OpenShift 18.0.

OVN pod goes into loop due to NIC Mapping

When using a large number of NIC mappings, OVN might go into a creation loop.

Jira:OSPRH-7480

2.16.4. Network Functions Virtualization

2.16.4.1. Known issues

This part describes known issues in Red Hat OpenStack Services on OpenShift 18.0.

Listing physical function (PF) ports using neutron might show the wrong MAC

Lists of PF ports might show the wrong MAC.

Jira:OSPRH-7085

2.16.5. Storage

2.16.5.1. Known issues

This part describes known issues in Red Hat OpenStack Services on OpenShift 18.0.

Image uploads might fail if a multipathing path for Block Storage service (cinder) volumes is offline

If you use multipath for Block storage service volumes, and you have configured the Block Storage service as the back end for the Image service (glance), image uploads might fail if one of the paths goes offline.

Jira:OSPRH-7393

RGW does not pass certain Tempest object storage metadata tests

Red Hat OpenStack Services on OpenShift 18.0 supports Red Hat Ceph Storage 7. Red Hat Ceph Storage 7 RGW does not pass certain Tempest object storage metadata tests as tracked by the following Jiras:

https://issues.redhat.com/browse/RHCEPH-6708https://issues.redhat.com/browse/RHCEPH-9119https://issues.redhat.com/browse/RHCEPH-9122https://issues.redhat.com/browse/RHCEPH-4654

Jira:OSPRH-7464

Missing Barbican configuration in the Image service (glance)

The Image service is not automatically configured to interact with Key Manager (barbican), and encrypted image signing and verification fails due to the missing configuration.

Jira:OSPRH-7155

2.16.6. Release delivery

2.16.6.1. Removed functionality

This part provides an overview of functionality that has been *removed* in Red Hat OpenStack Services on OpenShift 18.0.

Removed functionality is no longer supported in this product and is not recommended for new deployments.

Removal of **snmp** and **snmpd**

The **snmp** service and **snmpd** daemon are removed in RHOSO 18.0.

Jira:OSPRH-2960

2.16.7. Integration test suite

2.16.7.1. Known issues

This part describes known issues in Red Hat OpenStack Services on OpenShift 18.0.

Tempest test-operator does not work with LVMS storage class

When the test-operator is used to run Tempest, it requests a “ReadWriteMany” PersistentVolumeClaim (PVC) which the LVMS storage class does not support. This causes the tempest-test pod to become stuck in the **pending** state.

Workaround: Use the test-operator with a storage class supporting **ReadWriteMany** PVCs. The test-operator should work with a **ReadWriteOnce** PVC so the fixed version will no longer request a **ReadWriteMany** PVC.

Jira:OSPRH-7062