Red Hat Enterprise Linux OpenStack Platform 5 Technical Notes for RHEL7.1 Release

---

Technical Notes for Red Hat Enterprise Linux OpenStack Platform and supporting packages.

OpenStack Documentation Team

Red Hat Enterprise Linux OpenStack Platform 5 Technical Notes for RHEL7.1 Release

Technical Notes for Red Hat Enterprise Linux OpenStack Platform and supporting packages.

OpenStack Documentation Team
Red Hat Customer Content Services
rhos-docs@redhat.com

Legal Notice

Abstract

These Technical Notes are provided to supplement the information contained in the text of Red Hat Enterprise Linux OpenStack Platform errata advisories released via Red Hat Network.

Table of Contents

Chapter 1. Overview 3

Chapter 2. RHEA-2014:0854 — Red Hat Enterprise Linux OpenStack Platform ..... Enhancement - Identity 7
2.1. openstack-keystone 7
2.2. python-keystoneclient 10

Chapter 3. RHEA-2014:0853 — Red Hat Enterprise Linux OpenStack Platform ..... Enhancement - Compute 13
3.1. openstack-nova 13
3.2. python-novaclient 20

Chapter 4. RHEA-2014:0852 — Red Hat Enterprise Linux OpenStack Platform ..... Enhancement - Block Storage 21
4.1. openstack-cinder 21

Chapter 5. RHEA-2014:0851 — Red Hat Enterprise Linux OpenStack Platform ..... Enhancement - Image Service 27
5.1. openstack-glance 27

Chapter 6. RHEA-2014:0849 — Red Hat Enterprise Linux OpenStack Platform..... Enhancement - Orchestration 31
6.1. openstack-heat 31

Chapter 7. RHEA-2014:0848 — Red Hat Enterprise Linux OpenStack Platform..... Enhancement - Networking 33
7.1. openstack-neutron 33

Chapter 8. RHEA-2014:0855 — Red Hat Enterprise Linux OpenStack Platform ..... Enhancement - Dashboard 36
8.1. python-django-horizon 36
8.2. redhat-access-plugin-openstack 39

Chapter 9. RHEA-2014:0846 — Red Hat Enterprise Linux OpenStack Platform..... Enhancement - Packstack 40
9.1. openstack-packstack 40
9.2. openstack-puppet-modules 44

Chapter 10. RHBA-2014:0937 — Red Hat Enterprise Linux OpenStack Platform Bug . Fix and Enhancement Advisory 46
10.1. openstack-sahara 46
10.2. openstack-selinux 46
10.3. python-requests 47

Chapter 11. RHEA-2014:1003 — Red Hat Enterprise Linux OpenStack Platform ..... Enhancement Advisory 48
11.1. foreman 48
11.2. foreman-selinux 48
11.3. openstack-foreman-installer 48
11.4. openstack-puppet-modules 51
11.5. rhel-osp-installer 52
11.6. rubygem-staypuft 52

Chapter 12. RHBA-2014:1090 — Red Hat Enterprise Linux OpenStack Platform..... Enhancement Advisory 53
12.1. foreman 53
12.2. foreman-discovery-image 53

12.2. foreman-discovery-image 53
12.3. openstack-foreman-installer 53
12.4. rhel-osp-installer 55
12.5. rubygem-staypuft 57

Chapter 13. RHBA-2014:1324 — openstack-packstack and openstack-puppet-modules bug fix advisory 59
13.1. openstack-packstack 59

Chapter 14. RHBA-2014:1325 — Red Hat Enterprise Linux OpenStack Platform Bug Fix and Enhancement Advisory 61
14.1. ceph 61
14.2. galera 61
14.3. heat-cfntools 61
14.4. openstack-ceilometer 61
14.5. openstack-cinder 62
14.6. openstack-selinux 63

Chapter 15. RHBA-2015:0825 — Red Hat Enterprise Linux OpenStack Platform Bug Fix and Enhancement Advisory 64
15.1. openstack-sahara 64
15.2. openstack-selinux 64
15.3. openstack-utils 64
15.4. python-novaclient 65
15.5. python-sqlalchemy 65
15.6. rabbitmq-server 66

Chapter 16. RHSA-2015:0843 — Important: openstack-nova security, bug fix, and enhancement update 67
16.1. openstack-nova 67
16.2. vulnerability 70

Appendix A. Revision History 72

Chapter 1. Overview

These Technical Notes are provided to supplement the information contained in the text of Red Hat Enterprise Linux OpenStack Platform errata advisories released through Red Hat Network. If the text for an advisory's problem description is too lengthy to fit into the advisory itself, bug listings for that advisory are published as a chapter in this document.

The following table contains the list of errata advisories for this version.

Table 1.1. Errata Advisories

Release	Advisories
5.0.0	Issued in early July 2014: • Errata chapters: ◦ Chapter 2, *RHEA-2014:0854 — Red Hat Enterprise Linux OpenStack Platform Enhancement - Identity* ◦ Chapter 3, *RHEA-2014:0853 — Red Hat Enterprise Linux OpenStack Platform Enhancement - Compute* ◦ Chapter 4, *RHEA-2014:0852 — Red Hat Enterprise Linux OpenStack Platform Enhancement - Block Storage* ◦ Chapter 5, *RHEA-2014:0851 — Red Hat Enterprise Linux OpenStack Platform Enhancement - Image Service* ◦ Chapter 6, *RHEA-2014:0849 — Red Hat Enterprise Linux OpenStack Platform Enhancement - Orchestration* ◦ Chapter 7, *RHEA-2014:0848 — Red Hat Enterprise Linux OpenStack Platform Enhancement - Networking* ◦ Chapter 8, *RHEA-2014:0855 — Red Hat Enterprise Linux OpenStack Platform Enhancement - Dashboard* ◦ Chapter 9, *RHEA-2014:0846 — Red Hat Enterprise Linux OpenStack Platform Enhancement - Packstack* • Additional advisories include: ◦ RHEA-2014:0845 - Runtime Components. ◦ RHEA-2014:0847 - Object Storage. ◦ RHEA-2014:0850 - Telemetry Issued in late July 2014: • Errata chapters: ◦ Chapter 10, *RHBA-2014:0937 — Red Hat Enterprise Linux OpenStack Platform Bug Fix and Enhancement Advisory* ◦ Chapter 11, *RHEA-2014:1003 — Red Hat Enterprise Linux OpenStack Platform Enhancement Advisory* ◦ Chapter 12, *RHBA-2014:1090 — Red Hat Enterprise Linux OpenStack Platform Enhancement Advisory* • Additional advisories include: ◦ RHBA-2014:0930 - openstack-packstack and openstack-puppet-modules. ◦ RHBA-2014:0931 - openstack-keystone. ◦ RHBA-2014:0932 - openstack-cinder. ◦ RHBA-2014:0933 - openstack-glance and python-glanceclient. ◦ RHBA-2014:0934 - openstack-ceilometer and python-ceilometerclient. ◦ RHBA-2014:0935 - openstack-heat. ◦ RHBA-2014:0936 - openstack-neutron. ◦ RHEA-2014:0938 - Red Hat Enterprise Linux OpenStack Platform Enhancement - Ceph Client.

Release	Advisories
	RHSA-2014:0939 - Moderate: python-django-horizon security, bug fix, and enhancement update. RHSA-2014:0940 - Moderate: openstack-nova security and bug fix update. Issued in Septemner 2014: Advisories: RHBA-2014:1125 - PackStack and puppet-modules bug fix advisory. RHBA-2014:1127 - Compute bug fix update. RHBA-2014:1129 - Block Storage bug fix advisory. RHBA-2014:1131 - Image bug fix advisory. RHBA-2014:1135 - Orchestration bug fix advisory. RHBA-2014:1133 - Telemetry bug fix advisory. RHBA-2014:1116 - Red Hat Enterprise Linux OpenStack Platform Bug Fix and Enhancement Advisory. RHSA-2014:1121 - Low: Identity security and bug fix update. RHBA-2014:1137 - Dashboard bug fix update. RHSA-2014:1119 - Moderate: OpenStack Networking security, bug fix, and enhancement update. RHBA-2014:1138 - Red Hat Enterprise Linux OpenStack Platform Bug Fix Advisory. These September packages include rebases to 2014.1.2 for Compute, Block Storage, Image, Orchestration, Telemetry, Identity, Dashboard, OpenStack Networking and Data Processing (sahara - Technical Preview); and Open vSwitch package to 2.1.2
5.0.1	Errata Chapters: Chapter 13, *RHBA-2014:1324 — openstack-packstack and openstack-puppet-modules bug fix advisory* Chapter 14, *RHBA-2014:1325 — Red Hat Enterprise Linux OpenStack Platform Bug Fix and Enhancement Advisory* Additional advisories include: RHSA-2014:1335 - Moderate: python-django-horizon security and bug fix update. RHSA-2014:1337 - Moderate: openstack-glance security and bug fix update. RHBA-2014:1345 - openstack-heat bug fix advisory. RHBA-2014:1347 - openstack-keystone bug fix update.

Release	Advisories
5.0.2	Advisories include: - RHBA-2014:1770 - openstack-packstack and openstack-puppet-modules bug fix advisory. - RHBA-2014:1772 - openstack-heat bug fix advisory. - RHBA-2014:1775 - openstack-glance bug fix advisory. - RHBA-2014:1776 - openstack-ceilometer bug fix advisory. - RHBA-2014:1778 - python-django-horizon bug fix update. - RHBA-2014:1780 - Red Hat Enterprise Linux OpenStack Platform Bug Fix and Enhancement Advisory. - RHSA-2014:1782 - Important: openstack-nova security, bug fix, and enhancement update. - RHSA-2014:1784 - Moderate: python-keystoneclient security and bug fix update. - RHSA-2014:1786 - Moderate: openstack-neutron security, bug fix, and enhancement update. - RHSA-2014:1788 - Moderate: openstack-cinder security and bug fix update. - RHSA-2014:1790 - Important: openstack-keystone security and bug fix update. These packages include rebases to 2014.1.3 for the Block Storage, Compute, Dashboard, Identity, Image, Networking, Orchestration, and Telemetry services,
5.0.3	Advisories include: - RHBA-2014:1926 - openstack-cinder bug fix advisory. - RHBA-2014:1929 - python-django-horizon bug fix update. - RHBA-2014:1930 - openstack-heat bug fix advisory. - RHBA-2014:1933 - openstack-nova bug fix update. - RHBA-2014:1934 - openstack-packstack and openstack-puppet-modules bug fix advisory. - RHBA-2014:1935 - Red Hat Enterprise Linux OpenStack Platform Bug Fix and EnhancementAdvisory. - RHBA-2014:1936 - openstack-ceilometer bug fix advisory. - RHSA-2014:1939 - Low: openstack-trove security update. - RHSA-2014:1940 - Important: mariadb-galera security update. - RHSA-2014:1941 - Low: qemu-kvm-rhev security update. - RHSA-2014:1942 - Moderate: openstack-neutron security and bug fix update.

Release	Advisories
5.0.4	Errata chapters: - Chapter 15, *RHBA-2015:0825 — Red Hat Enterprise Linux OpenStack Platform Bug Fix and Enhancement Advisory* - Chapter 16, *RHSA-2015:0843 — Important: openstack-nova security, bug fix, and enhancement update* Additional advisories include: - RHBA-2015:0818 - openstack-cinder bug fix advisory. - RHBA-2015:0821 - openstack-ceilometer bug fix advisory. - RHBA-2015:0821 - openstack-keystone bug fix advisory. - RHBA-2015:0827 - openstack-heat bug fix advisory. - RHBA-2015:0829 - openstack-neutron bug fix advisory. - RHSA-2015:0831 - Important: openstack-packstack and openstack-puppet-modules update. - RHSA-2015:0834 - Moderate: novnc security update. - RHSA-2015:0835 - Moderate: openstack-swift security update. - RHSA-2015:0837 - Low: openstack-glance security and bug fix update. - RHSA-2015:0839 - Moderate: python-django-horizon and python-django-openstack-auth update. - RHSA-2015:0840 - Important: redhat-access-plugin security update. These packages include rebases to 2014.1.4 for the Block Storage, Compute, Dashboard, Identity, Image, OpenStack Networking, Orchestration, and Telemetry services.
5.0.5	Advisories: - RHBA-2015:1746 - openstack-packstack and openstack-puppet-modules bug fix advisory. - RHBA-2015:1748 - openstack-swift bug fix advisory. - RHBA-2015:1750 - python-django-horizon bug fix advisory. - RHBA-2015:1752 - openstack-keystone bug fix advisory. - RHBA-2015:1754 - openstack-neutron bug fix advisory. - RHBA-2015:1756 - openstack-nova bug fix advisory. - RHBA-2015:1758 - openstack-cinder bug fix advisory. - RHBA-2015:1760 - openstack-glance bug fix advisory. - RHBA-2015:1761 - openstack-heat bug fix advisory. - RHBA-2015:1762 - Red Hat Enterprise Linux OpenStack Platform Bug Fix and Enhancement advisory. - RHBA-2015:1765 - openstack-ceilometer bug fix advisory. - RHSA-2015:1767 - Moderate: python-django security update. - RHSA-2015:1769 - Low: libunwind security update.

Chapter 2. RHEA-2014:0854 — Red Hat Enterprise Linux OpenStack Platform Enhancement - Identity

The bugs contained in this chapter are addressed by advisory RHEA-2014:0854. Further information about this advisory is available at https://rhn.redhat.com/errata/RHEA-2014-0854.html.

2.1. openstack-keystone

BZ#901955

Previously, attempting to start the Identity service when it was already running would throw a Python exception. So instead of reporting a useful error message, a raw exception with stack trace was displayed, which may not be very useful to the user.

This has been fixed and a more understandable error message is displayed if attempting to start the Identity service when it is already running.

BZ#908355

The SQL backend for Identity records tokens. It does not have a timeout, and it does not automatically remove tokens once they are recorded.

As a consequence, the SQL database can run out of storage space.

As a workaround, Identity now includes a command to remove tokens, namely 'keystone-manage token_flush'. This process should be scheduled to run regularly via cron. It is recommended that this command be run approximately once per minute.

BZ#970098

Previously, the service catalog used to return all endpoints, regardless of status. This meant that disabled endpoints were displayed as well.

Now, only enabled endpoints are returned by default.

BZ#1031214

Previously, using non-ASCII characters in names managed by Identity (i.e. non-English names), caused a server error and the request could not be completed.

Now, the UNICODE text is correctly encoded/decoded to/from UTF-8 during LDAP operations, and user names containing non-ASCII characters (i.e. non-English) are successfully stored and operated on in the LDAP backend.

BZ#1033190

Some errors, including those thrown from the database, are reported in an ASCII byte string format. This error message is included in some failures and so the error message would display the byte escaped output and UTF-8 values would be lost. As a consequence, UTF-8 characters would appear escaped, for example \\u2013 in messages instead of the correct UTF-8 character.

This has been fixed to ensure that UTF-8 characters that are present will be correctly rendered. Now the output displays full UTF-8 characters in error messages.

BZ#1041859

Previously, users could not update their own passwords using the V3 API, only administrators could update users' passwords using the V3 API.

Now that the V3 API is the default, (no longer the V2 API), users can update their own passwords too.

BZ#1041860

A region resource has been introduced to the Identity API for constructing a hierarchical container of groups of service endpoints.

Previously, service endpoints could refer to a region as an arbitrary string. Regions can now be explicitly defined and managed through the Identity API, which are then referred to when adding service endpoints. This allows more control over region management.

BZ#1041863

When requesting validation of an Identity token, the ability to opt-out of including the service catalog in the response is now available to the requestor.

Identity tokens sizes can get large due to the inclusion of the service catalog. When the service catalog is not needed, it may be desirable to get a token that omits the service catalog. This ability was previously available only when generating a new token, but is now also available when validating of existing tokens.

BZ#1041864

Previously, Identity only provided API error and exception messages in English. Providing translated messages would be more useful when the requestor is using a different locale.

Now, translated messages are provided if a translation exists for the requestor's locale.

BZ#1041865

A change was introduced to limit the number of results returned by list commands.

This was done to avoid problems when the number of results can be larger than available memory. For example, a user list on a system where there are millions of users would return all of them.

Now users can provide a configuration value (list_limit=<integer value>) on calls to the Identity API list functions. Lists that have values beyond those limits will return a 'truncated' value in the body of the results.

BZ#1041875

Audits are essential for maintaining security of a system, and especially applicable to Identity since Identity is a security-focused service.

OpenStack has adopted CADF as the format for audit events. The Identity service now emits CADF events upon Identity and Token operations.

BZ#1041930

Implementers of custom Identity extensions may need to perform tasks when Identity resources such as users or groups are created, updated, or deleted. The ability to register callbacks for these events has been added to allow for more complex custom Identity extensions.

BZ#1041959

Previously, Identity trusts allowed the trust to be used to issue tokens for an unlimited number of times as long as the trust was valid.

This new feature adds the ability to specify the exact number of times that a trust can be used to issue tokens, allowing for uses such as a one-time use trust.

BZ#1052807

The Identity service default token duration setting has been reduced to one hour.

The Identity service previously defaulted to a token duration of 24 hours, which would results in scalability problems due to a large number of tokens being persisted in Identity's token

database for tokens that are most likely not in use any more.

Now, the number of tokens persisted in Identity's token database will be vastly reduced compared to the previous default setting, resulting in improved scalability.

BZ#1055856

Identity previously emitted notifications for create, update, and delete operations for user, group, role, and project resources.

This has been extended to so that Identity also emits notifications for create, update, and delete operations for trust resources.

BZ#1056875

A new feature has been added that allows the Identity service's log messages to be translated based on the system locale.

Previously Identity only provided log messages in English. Now, translated messages will be provided if a translation exists for the system locale.

BZ#1059963

An enhancement has been made to use Oslo Messaging for notifications.

Previously, Identity used the incubated code, but messaging has graduated to a stand-alone library. Using the new library keeps Identity consistent with the rest of the OpenStack services in the way it handles notifications.

This ensures that Identity gets all bug fixes and feature enhancements.

BZ#1073011

Previously, the LDAP code in Identity was comparing attribute names by using string comparisons. Inconsistent capitalization caused the string comparisons to fail.

As a result, two values that should have matched would not match, and binding the LDAP query results to the Python variables would fail with the error "KeyError on user_ref['name']"

This has been fixed by doing attribute comparisons with all values forced to lowercase, so that attributes now match even if configuration values do not have consistent capitalization.

2.2. python-keystoneclient

BZ#1058577

Previously, the S3 middleware was in the keystone repository and package. Hence, in order to use the S3 middleware, the user had to either copy a specific file out, or pull in the whole keystone package. Neither of these were good solutions.

Now, the S3 middleware is in the keystone-client. A user can now deploy the middleware separately from the server, using just the client code.

BZ#1076869

Previously, keystoneclient was providing output in unicode instead of UTF-8 characters. Hence subsequent readers of the data were incorrectly trying to read ASCII data. This resulted in readers being unable to parse the data.

Now, keystoneclient always provides output in UTF-8 format.

BZ#1077182

Previously, the token revocation list was not being checked prior to the token cache when validating tokens in Identity.

As a result, unexpired tokens that were explicitly revoked would still be considered valid if they existed in the token cache. This caused a revocation cache time setting that was lower than the token cache time to be useless.

The revocation cache is now checked prior to validating tokens from the token cache. Now, tokens that have been explicitly revoked will be considered as invalid once the revocation list is retrieved, even if they exist in the token cache. The default revocation cache time is 10 seconds, which results in revoked tokens being unusable within 10 seconds.

BZ#1101713

The python-keystoneclient on RHEL 6 was deploying its Bash command completion file to /etc/profile.d/, instead of /etc/bash_completion.d/ because RHEL 6 did not include a bash-complation package. (Command completion expands command arguments after pressing TAB).

As a consequence, deploying Bash completion to /etc/profile.d/ might break a customer's Bash completion or cause problems when using other shells.

The bash-completion package is available in RHEL 7 and the python-keystoneclient package now deploys its Bash completion file to /etc/bash_completion.d/.

Note that to enable the keystone-client command completion, the bash-completion package must installed.

Chapter 3. RHEA-2014:0853 — Red Hat Enterprise Linux OpenStack Platform Enhancement - Compute

The bugs contained in this chapter are addressed by advisory RHEA-2014:0853. Further information about this advisory is available at https://rhn.redhat.com/errata/RHEA-2014-0853.html.

3.1. openstack-nova

BZ#966050

```
In the previous release, logging went to both syslog AND to file
if syslog was enabled in the Compute configuration. This double
logging was a result of the initscripts of the various components
including a '--logfile $logfile'.

The --logfile argument has now been removed from Compute service
scripts and
the default log file used. Logging is no longer doubled when
syslog is enabled.
Additionally, the default log files have been renamed to nova-
{service_name}.log (for example, from {log_dir}/compute.log to
{log_dir}/nova-compute.log).
```

BZ#978507

```
Compute users can now define a server group with an associated
policy; supported policies are affinity and anti-affinity.  When
servers are created, they can be associated with a server group.
When planning the new server, Compute enforces the affinity or
anti-affinity policy among all instances associated with this
group.
```

BZ#1024032

```
In OpenStack Compute, attaching volumes to instances is not
supported for device names like '/dev/hd*' which will cause the
bus to be defaulted to 'ide', which cannot be hot plugged.

Attaching volumes to running instances is only supported for
virtio, so the device name needs to be similar to '/dev/vd*'
```

BZ#1038668

```
Watchdog support has been added to the Libvirt driver. The
watchdog device used is "i6300esb", and is enabled by setting the
"hw_watchdog_action" property in the image properties or flavor
extra specifications ("extra_specs") to a value other than
"disabled".
```

Supported "hw_watchdog_action" property values, which specify the action for the watchdog device to take in the event of an instance failure, are "poweroff", "reset", "pause", and "none".

BZ#1040599

With this update, a Compute deployment which has a configured database slave (slave_connection) can send reads from periodic tasks to this slave. Periodic tasks are some of the most consistent load that any deployment will experience. To improve performance, these tasks can now be offloaded to database slaves.

BZ#1040985

In OpenStack Compute, the OS-DCF:diskConfig API attribute is no longer supported in V3 of the nova API.

BZ#1040993

The Compute service determines what action to take when instances are found to be running that were previously marked deleted based on the value of the "running_deleted_instance_action" configuration key. A new "shutdown" value has been added to the list of configurable actions.

Using this new value allows administrators to optionally keep instances found in this state for diagnostics while still releasing the run-time resources.

BZ#1041014

Compute services are now able to shutdown gracefully by disabling the processing of new requests when a service shutdown is requested but allowing requests already in process to complete before terminating.

BZ#1041017

Notifications are now generated when a Compute host is enabled, disabled, powered on, shut down, rebooted, put into maintenance mode, or taken out of maintenance mode.

BZ#1041018

The Compute API now exposes the hypervisor IP address, allowing it to be retrieved by administrators using the "nova hypervisor-show" command.

BZ#1041023

Notifications are now generated upon the creation and deletion of keypairs.

BZ#1041026

The Libvirt driver now allows instance configuration to use video drivers other than the default (cirros), so that different video driver models and amounts of video RAM can now be specified. These values are configured by setting the "hw_video_model" and "hw_video_ram" properties in the image metadata. Currently supported video-driver models are "vga", "cirrus", "vmvga", "xen", and "qxl".

BZ#1041031

With this update:
* Weights have been normalized in OpenStack Compute so that there is no need to articificially inflate multipliers. The maximum weight that a weigher puts for a node is 1.0, and the minimum is 0.0.
* A new multiplier option, 'offset_weight_multiplier' (nova.cells.weights.weight_offset.WeightOffsetWeigher), has been introduced.
 * Stacking flags for weighers have been introduced. Negative multipliers should not be using for stacking, but the weighers are still compatible (they issue a deprecation warning message).

BZ#1041038

The /etc/nova/nova.conf configuration file has been updated to ensure that all configuration groups in the file use descriptive names. A number of driver-specific flags, including those for the Libvirt driver, have also been moved to their own option groups.

BZ#1041051

The Compute service now uses the tenant identifier instead of the tenant name when authenticating with OpenStack Networking (neutron). This improves support for the OpenStack Identity API v3, which allows non-unique tenant names.

BZ#1041053

With this update, the API call for attaching volumes to instances in Compute (servers/<INSTANCE_UUID>/os-volume_attachments) now accepts two additional parameters in the body: disk_bus and device_type. If these parameters are specified, the libvirt driver attempts to honor them when attaching the volume. The following values are accepted:
* disk_bus: 'scsi' and 'virtio'
* device_type: 'disk', 'cdrom', 'floppy', and 'lun'

BZ#1041055

The V3 API admin_actions plugin has now been separated into logically separate plugins so that operators can enable subsets of the functionality currently present in the plugin.

BZ#1041067

VMware Compute drivers now support a virtual-machine diagnostics call. Diagnostics can be retrieved using the "nova diagnostics INSTANCE" command, where INSTANCE is replaced by an instance name or instance identifier.

BZ#1041084

Transient database-connection failures are now recovered automatically. There are a variety of circumstances which can cause a transient failure in database connection (for example, the restart or upgrade of the database, migration of VIP between an HA pair, or a network failure). Compute now catches these "db-has-gone-away" errors by automatically reconnecting and retrying the last operation in such a way that the caller is able to continue whatever operation was in progress. The user no longer has to abort long-running operations (such as 'nova boot' or 'glance image-create') just because of a momentary interruption in database connectivity.

BZ#1041103

The Compute API now exposes a mechanism for permanently removing decommissioned compute nodes. Previously, decommissioned nodes would continue to be listed even if the compute service had been disabled and the system re-provisioned. The removal functionality is provided by the "ExtendedServicesDelete" API extension.

BZ#1041118

A new scheduler filter, "AggregateImagePropertiesIsolation", has been introduced. The new filter schedules instances to hosts based on matching namespaced image properties with host aggregate properties. Hosts that do not belong to any host aggregate remain valid scheduling targets for instances based on all images.

The new Compute service configuration keys "aggregate_image_properties_isolation_namespace" and "aggregate_image_properties_isolation_separator" are used to determine which image properties are examined by the filter.

BZ#1045289

Support has been added for VMware instances which boot from an ISO image. Software licensing which prevents the distribution of modified software images is now supported. In particular, this enables compliant creation of Microsoft Windows instances. For more information, see: https://wiki.openstack.org/wiki/BootFromISO

BZ#1045805

The High Precision Event Timer (HPET) is now disabled for instances created using the Libvirt driver. The use of this option was found to lead to clock drift in Windows guests when under heavy load.

BZ#1052799

With this enhancement, the Libvirt compute driver now supports providing modified kernel arguments to booting compute instances from AMI images. Kernel arguments are retrieved from the "os_command_line" key in the image metadata (as stored in the Image service), if the key's value is provided; otherwise, the default kernel arguments are used.

BZ#1055529

OpenStack did not check previously whether the driver in use supports security groups. The VMware driver does not support security groups with flat networking, which resulted in the use of the feature resulting in an error.

OpenStack now checks to see whether the driver supports security groups. Attempting to use security groups with the VMware driver and flat networking now results in a warning rather than an error.

BZ#1055853

The Compute service now includes a caching scheduler driver to help improve scheduler performance. The caching scheduler uses existing facilities for applying scheduler filters and weights, and caches the resultant list of available hosts.

When a user request is passed to the caching scheduler, the driver attempts to schedule based on the list of cached hosts (and only uses non-cached information if the attempt fails).

BZ#1055855

The Libvirt Compute driver now supports adding a Virtio RNG device to compute instances to provide increased entropy. Virtio RNG is a paravirtual, random-number generation device, which allows the compute node to provide entropy to compute instances in order to fill their entropy pool.

The default entropy device used is /dev/random, however the use of a physical hardware RNG device attached to the host is also possible. The use of the Virtio RNG device is enabled using the hw_rng property in the metadata of the image used to build the instance.

BZ#1056381

File injection into VM images is now deprecated; instead, Red Hat

recommends that you use the ConfigDrive and metadata server facilities to modify guests at launch.

File injection is now disabled by default in OpenStack Compute. To enable file injection, modify the inject_key and inject_partition configuration keys in /etc/nova/nova.conf and restart the Compute services. Note that the file-injection mechanism will probably be disabled in a future release.

BZ#1058444

The Libvirt driver now supports using VirtIO SCSI (virtio-scsi) instead of VirtIO Block (virtio-blk) to provide block-device access for instances. Virtio SCSI is a para-virtualized SCSI controller device designed as a future successor to VirtIO Block, and provides improved scalability and performance.

BZ#1062815

To aid administrative troubleshooting, Guru Meditation reports are now sent by the Compute Service upon receipt of the SIGUSR1 signal. The report is sent to stderr and has the following sections:
* Package – Displays information about the package to which the process belongs, including version information.
* Threads – Displays stack traces and thread IDs for each of the threads within the process.
* Green Threads – Displays stack traces for each of the green threads within the process.
* Configuration – Lists all configuration options currently accessible through the CONF object for the current process.

BZ#1068691

Although the Compute API currently supports both XML and JSON formats, usage of the XML format is now deprecated. API support for XML will be retired in a future release.

BZ#1069430

Previously, instances running on a VMware host were allocated a VNC port in a manner which was not guaranteed to be unique. This meant it was possible for multiple instances to be given the same VNC port on the same VMware host. Accessing the console of any of these instances would give access to one of them, but which one was non-deterministic. With this update, the VNC port allocated is now guaranteed to be unique in a vCenter, and it is no longer possible to have a VNC port collision between two guests.

BZ#1080620

Previously, the VMware driver assumed that all clusters managed by a single vCenter were also managed by the same Compute. At startup, it checked that all VMs in the vCenter were known to it,

and killed those that were not. This meant that if multiple Computes managed clusters from the same vCenter, on startup Compute would kill all instances belonging to other Computes.

The VMware driver has now been updated to check only those instances in clusters it is actually managing. It is now possible to have multiple Computes managing clusters from the same vCenter.

BZ#1080621

Previously, instance names could be edited by vCenter administrators, and badly edited names could break the Compute administration of VMs. With this update, a vSphere metadata value is now used for the VM name (which is not easily edited). This makes Compute lookups far more robust.

BZ#1080622

Previously, if a VMware cluster contained a large number of datastores, the VMware driver would only consider those initially returned by vSphere when selecting a datastore for deployment. If there were more datastores than were initially returned, these were never considered, which meant that new datastores were not used.

With this fix, the VMware driver now checks all datastores in a cluster. A cluster may contain any number of datastores, and they will all be considered when deploying a new instance.

BZ#1100439

Previously, Compute's /etc/release file incorrectly listed the vendor as "Fedora Project" and the product as "OpenStack Nova", which then propagated into the SMBIOS for VMs. This meant that Satellite did not correctly recognize launched VMs, which meant that entitlements did not work correctly.

With this update, Compute's release information has been corrected to specify "Red Hat" and "OpenStack Compute", respectively. Satellite now correctly recognizes OpenStack VMs, and therefore entitlements now work as well.

BZ#1109182

An incorrect multiplier was previously used when resizing Ceph-backed instances (RBD image), resulting in Ceph volumes being 1024 times too large. With this update, Ceph volumes now use the flavor size correctly.

BZ#1110367

The live migration of Ceph-backed instances (RBD image) was previously unsupported. With this update, shared storage checks have been augmented so that Ceph-backed instances can now be migrated live.

3.2. python-novaclient

BZ#1101015

In the previous release, server group commands were not included in python-novaclient, and nova secgroup-* commands were not available. Server group commands have now been backported, so that nova secgroup-* commands are now supported.

Chapter 4. RHEA-2014:0852 — Red Hat Enterprise Linux OpenStack Platform Enhancement - Block Storage

The bugs contained in this chapter are addressed by advisory RHEA-2014:0852. Further information about this advisory is available at https://rhn.redhat.com/errata/RHEA-2014-0852.html.

4.1. openstack-cinder

BZ#984270

With this release, you can now modify a given volume's type. When modifying a volume's type, the Block Storage scheduler checks if the volume's current host can accept the new type: as in, the scheduler checks if the host passes the filters when using the specified type. If the current host is suitable, the volume's corresponding manager then calls the right driver to change the volume's type accordingly.

If the current host cannot accept the new type, or if the volume driver is unable to modify the volume's type, then you will need to migrate the volume in order to change the volume type. Specifically, you will need to create a new volume of the type you want, and then migrate the contents of original volume to this one.

BZ#985500

With this release, you can now set a volume to read-only access. This feature allows you to give multiple users shared, secure access to the same data.

BZ#1023384

In previous releases, if a source of a volume or image was larger than the destination, then cloning the source or image resulted in inconsistent errors, depending on the source type. This release applies several consistency fixes to the error messaging system of volume clones, as well as volume clones from:

* images,
* snapshots, and
* other volumes

In addition, with this release error messages now indicate the unit (for example, GB).

BZ#1030055

Previously, resizing an existing volume did not automatically send a request to tgt to update the volume's exported LUN. As a

result, the new size would not be displayed when the volume was attached to an instance.

To address this, volume LUNs are now exported when the volumes are attached to an instance (instead of during volume creation). This will ensure that any changes to the volume's size (before being attached) will be visible.

BZ#1041487

The Block Storage service now sends notifications for attach and detach events, allowing other openstack services (e.g. Telemetry) to listen for and display the results to the user.

With these notifications, a volume's status can now be updated automatically in the Telemetry service. This, in turn, allows an administrator to search samples for volume status history.

BZ#1041489

With this release, you can now export the metadata of a volume backup. This ability provides a more complete way to backup and restore a Cinder volume. Information like Glance metadata can now be included in volume backups.

BZ#1041491

Previously, the 'cinder absolute-limits' command only displayed the maximum usable limits of a tenant. With this release, the same command now returns a tenant's consumed resources as well, namely:

* totalVolumesUsed (as in, the total number of volumes used)
* totalGigabytesUsed

BZ#1041499

This update adds a new option for the Block Storage API service, namely osapi_volume_workers. This option allows you to specify the number of API service workers (or OS processes) to launch for openstack-cinder-api.

Typically, this option can be used to set the number of OS processes to the number of CPU cores/threads on a machine; doing so will greatly increase the number of API requests that can be handled per second. To set this option,
run the following command on the openstack-cinder-api host:

```
openstack-config --set /etc/cinder/cinder.conf \
DEFAULT osapi_volume_workers [X]
```

Replace [X] with the target number of OS processes you wish to set. By default, the Block Storage API service will still run in one process (that is, osapi_volume_workers is set to None).

BZ#1041638

The OpenStack Block Storage service now features a Fibre Channel Zone Manager. This allows OpenStack Block Storage to automatically manage fibre channel SAN zoning, making it easier to deploy a properly-configured Block Storage fibre channel setup.

For more information on how to use the Fibre Channel Zone Manager, refer to the Red Hat Enterprise Linux OpenStack 5 Configuration Reference Guide.

BZ#1041662

With this release, you can now export and import volume metadata. This, in conjunction with the ability to backup and restore volumes, now allows you to restore a volume even in the event of a catastrophic database failure.

In addition, volume backup metadata support also adds portability to volume backups. Now, exporting a volume backup's metadata allows you to restore the volume backup on a different Block Storage service (or even a different cloud service altogether).

For more information, refer to:

https://access.redhat.com/site/documentation/en-US/Red_Hat_Enterprise_Linux_OpenStack_Platform/5/html/Cloud_Administrator_Guide/section_manage-volumes.html#volume-backup-restore-export-import

BZ#1041693

A trial run of taskflow is implemented in cinder for only the create-volume task. The user should see no difference in behaviour.

With this implementation, cinder developers can now evaluate the usefulness and maintainability of taskflow actions without migrating the entire codebase.

BZ#1041696

The Chance and Simple schedulers are now deprecated. To replicate their behaviour, use the FilterScheduler instead. Both behaviours use the following setting in common:

```
    scheduler_driver =
cinder.scheduler.filter_scheduler.FilterScheduler
scheduler_default_filters = [‘AvailabilityZoneFilter’,
‘CapacityFilter’,

    ‘CapabilitiesFilter’]
```

To replicate the behaviour of the Chance scheduler, add the

following parameter:

    scheduler_default_weighers = 'ChanceWeigher'

To replicate the behavior of the Simple scheduler, use the following instead:

    scheduler_default_weighers = 'AllocatedCapacityWeigher'
    allocated_capacity_weight_multiplier = -1.0

BZ#1041709

The Block Storage LVM driver now supports the LIO iSCSI target as an iSCSI back-end for an OpenStack implementation on Red Hat Enterprise Linux 7. This is required due to the move from scsi-target-utils (tgtd) to LIO in Red Hat Enterprise Linux 7.

BZ#1043547

In previous releases, it was possible for a failure in the Block Storage volume driver initialization process to prevent the 'openstack-cinder-volume' service to fail at startup. Whenever this occurred in a multiple back-end environment, the 'openstack-cinder-volume' service would become inaccessible, and a failure in one volume driver could result in other volume drivers becoming unavailable.

With this update, the Block Storage service now marks uninitialized back-ends and disables requests to those back-ends. As a result, volume driver initialization failures now only affect the driver and not the entire 'openstack-cinder-volume' service.

BZ#1060685

This release corrects the default location of the 'memcached_servers' option in the default /etc/cinder/cinder.conf configuration file. In previous releases, this option was incorrectly listed under DEFAULT; the Block Storage service expects 'memcached_servers' to be in the [keystone_authtoken] section, where it now appears.

BZ#1065182

To help improve debugging, Block Storage service log messages now include OpenStack request IDs. These IDs provide an easy way to associate log items with requests from other services such as Compute.

BZ#1066035

Previously, as is with most OpenStack projects, Cinder used to rely on an RPC library coming from oslo-incubator. This library graduated from the incubator and moved into its own project,

oslo.messaging. This RFE tracks Cinder's adoption of oslo.messaging.

Although the messaging library kept backwards compatibility with older configuration options, it is highly recommended to upgrade the configuration files to use the `transport_url` and the new rpc options, where needed. The older configuration options are likely to be removed in future versions of the project.

BZ#1067230

This update provides Block Storage service administrators with the ability to remove any existing quotas set for a particular tenant.

BZ#1087886

In previous releases, deleting a snapshot of an attached volume also incorrectly deactivated the origin volume, rendering unusable by the instance. This made it necessary for the user to manually reactivate the volume.

This was caused by an incorrect routine in the commands used for volume snapshot deletion. With this release, this routine has been refined and correct, ensuring that attached volumes are no longer deactivated when any of their snapshots are deleted.

BZ#1096489

In previous releases, it was possible for the deletion of a volume's snapshot to prevent the volume from being attached to an instance. This was because the commands used for volume deletion did not contain any commands to deactivate (before snapshot deletion) and reactivate (after snapshot deletion) volumes. As such, when deleting a volume's snapshot, it was possible for the volume to be left with incorrect flags that would render it unattachable to an instance.

With this release, volume snapshot deletion now works with the right commands, ensuring that volumes are deactivated and reactivated as expected.

BZ#1103500

With this release, you can now use the 'retype' subcommand to set a volume's type after creating the volume with a type of 'None'.

BZ#1107733

Previously, cinder-rtstool incorrectly required /etc/iscsi/initiatorname.iscsi to be present in order to create a LUN/ACL/portal successfully. This should not have been required since the Block Storage service will create the required ACLs dynamically at attach time anyway.

With this update, cinder-rtstool no longer requires /etc/iscsi/initiatorname.iscsi to create a LUN/ACL/portal. As such, iscsi-initiator-utils no longer needs to be installed locally when using a remote Nova compute node.

Chapter 5. RHEA-2014:0851 — Red Hat Enterprise Linux OpenStack Platform Enhancement - Image Service

The bugs contained in this chapter are addressed by advisory RHEA-2014:0851. Further information about this advisory is available at https://rhn.redhat.com/errata/RHEA-2014-0851.html.

5.1. openstack-glance

BZ#985825

Previously, Glance would enable all stores by default, resulting in certain stores printing errors unless manually configured. This update restricts the stores enabled by default, only enabling those that work out of the box and leaving those that require manual configuration disabled unless explicity configured by users.

BZ#1027316

This update enhances the logic used to register multiple locations for the same image in Glance, preventing users from registering the same location more than once for a given image.

BZ#1030067

This update migrates the notifier used by Glance to that provided by oslo.messaging, enhancing the integration of the Glance notifier with that of other Red Hat OpenStack components and providing support for high availability while retaining backward compatibility by translating the former 'notifier_strategy' in oslo.messaging drivers.

BZ#1031689

Previously, as per most other OpenStack projects, Glance used to rely on an RPC library coming from oslo-incubator. This library graduated from the incubator and was moved into its own project, oslo.messaging. This request for enhancement tracks Glance's adoption of oslo.messaging.

Although the messaging library maintains backwards compatibility with older configuration options, upgrading the configuration files to use the `transport_url` and the new RPC options as necessary is highly recommended. The older configuration options are likely to be removed in future versions of the project.

BZ#1041737

Previously, it was not possible to mount all disks to a single

directory when configuring multiple NFS servers as a backend using the filesystem store. This was due to the filesystem store only allowing administrators to configure a single directory using the filesystem_store_datadir parameter in the glance-api.conf file.

While it is possible to use MHDDFS (a FUSE plug-in: https://romanrm.net/mhddfs), which mounts multiple NFS servers to a single directory, MHDDFS does not allow you to evenly store the data on all the disks. Another major drawback is that it is very difficult to know the number and type of images stored on a disk when one of the disks is broken because the Glance registry stores the location specified in the filesystem_store_datadir parameter.

This enhancement fixes the above issues by adding multi-filesystem support to the current filesystem store.

BZ#1041738

Multi-location support allows Glance to store locations pointing to copies of the same image data stored in different places. Previously, this feature did not allow Glance to choose the store from which to download an image. This update adds two strategies for selecting the location to send back to the client - the location from which to download the image. The two strategies are:

1. Location order: A simple round-robin that goes through all the available locations in order - as they were inserted.

2. Store type: Allows users to specify store preferences. A possible combination is: `http, file`. This tells Glance to try to download images from HTTP stores before downloading them from file stores.

BZ#1041747

This update incorporates common code for working with databases such as session management, connections, engines, models, migrations and other utilities.

BZ#1041820

This update adds support for making new requests to Swift to download the remainder of an image when downloading an image from Swift fails before the full image is transferred. The download operation is attempted a number of times based on the value set in the new configuration option 'swift_store_retry_get_count'. If the value of this key is set to '0', no attempts to retry downloading the remainder of the image are attempted.

BZ#1048174

Previously, the size of a Glance image would be incorrectly

reported under certain circumstances. This was caused by differences in the file size and the virtual size of the image.

The size of an image can refer to either the size of the file or the size of the actual image, which may not be the same in cases such as qcow2. This update splits the current size attribute into two separate attributes: image_size and file_size. The former refers to the real size of the image and the later the actual size of the uploaded file. In most cases, both fields will hold the same value. However, there are also many cases in which this value will differ.

This update considers both attributes important for an image, but not required for that image to exist. The value of the image_size attribute allows users to know the actual size of the image and how much space is needed to use that image, such as Cinder block allocation. The value of the file_size attribute is necessary to support quotas, CLI progress bars, rate-limits, and metering, etc.

While the actual image size could be included as part of the image properties or meta data, it is important to remember that the image_size attribute is a first citizen attribute in most external tools. A discrete attribute will ease the consumption of its value from Nova, Cinder and other tools that rely on the value of that attribute.

BZ#1058494

Previously, the owner of an image was a private property in Glance's image v2. This enhancement makes the owner a public property and sends it back to the client when information on an image is requested.

BZ#1070258

Previously, the 'glance image-create' operation would report that it had completed successfully under certain conditions even when it had failed. This was caused by the logic used to create images, whereby it was possible to create images that referenced a URI even when that URI was invalid. This update revises this logic so that it is no longer possible to create images with an invalid URI and so that an error message is provided on failure.

BZ#1084495

Previously, the Glance API would ask for the RabbitMQ driver during installation even when QPID was selected as the notification driver. This was caused by the logic used to determine the driver to load, whereby the Glance API would always try to get a transport from oslo.messaging assuming that the 'transport_url' option has been set, but because the default 'rpc_backend' is RabbitMQ, the Glance API would try to load the RabbitMQ driver. Furthermore, when 'kombu' is not installed and

'notifier_strategy' is set to QPID, the Glance API would fail due to trying to load the RabbitMQ driver before loading the QPID driver. This update revises this logic so that the correct driver is loaded.

BZ#1094675

Previously, attempting to perform the 'glance-manage db_sync' operation would fail without reporting any errors under certain conditions. This was caused by a requirement in the Icehouse release and above specifying that all tables in the database use the UTF-8 character set when certain Glance-related tables from previous releases used character sets other than UTF-8. This update adds the line 'db_enforce_mysql_charset = False' to the 'glance-api-dist.conf' file, disabling checking of the character set and making it possible for the 'glance-manage db_sync' operation to complete successfully under these conditions.

Chapter 6. RHEA-2014:0849 — Red Hat Enterprise Linux OpenStack Platform Enhancement - Orchestration

The bugs contained in this chapter are addressed by advisory RHEA-2014:0849. Further information about this advisory is available at https://rhn.redhat.com/errata/RHEA-2014-0849.html.

6.1. openstack-heat

BZ#1042150

With this release, heat_keystoneclient now uses the Identity service's v3 API exclusively. This API version allows non-admin users to perform autoscaling and use wait conditions.

BZ#1042153

With this release, you can now run multiple Orchestration engines simultaneously with just one database back-end. This feature adds horizontal scalability to the Orchestration service.

BZ#1042154

With this release, the OS::Neutron::SecurityGroup resource is now supported. This resource allows you to reliably specify security groups for the OpenStack Networking service. These security groups provide IP security for instances.

BZ#1103826

As is OpenStack Networking policy rule, non-admin users are forbidden from creating routers with the enable_snat value set. In previous releases, enable_snat was set by default in the Routers resource; this, in turn, prevented non-admin users from creating routers.

With this release, enable_snat is no longer set by default. Non-admin users should now be able to create routers using default resource settings.

BZ#1104709

The heat-keystone-setup-domain script has been ported to the heat-common package of Red Hat Enterprise Linux OpenStack Version 5 (Icehouse). This script is needed to provide domain creation support in the puppet-heat module.

BZ#1106858

In previous releases, a bug in the way QPID communicated with the

broker made it possible for a required routing key to be left out of the address for a direct producer. Whenever this occurred, instances could not be launched.

This release addresses the bug, thereby ensuring that routing keys are always specified in the address for a direct producer.

Chapter 7. RHEA-2014:0848 — Red Hat Enterprise Linux OpenStack Platform Enhancement - Networking

The bugs contained in this chapter are addressed by advisory RHEA-2014:0848. Further information about this advisory is available at https://rhn.redhat.com/errata/RHEA-2014-0848.html.

7.1. openstack-neutron

BZ#973750

```
The network namespace cleanup script 'neutron-netns-cleanup' now
works as expected.
'neutron-netns-cleanup' removes all unused network namespaces on
the network node:
neutron-netns-cleanup --config-file=/etc/neutron/neutron.conf --
config-file=/etc/neutron/dhcp_agent.ini

For example, creating a network and a router results in two
namespaces: One for the DHCP server and one for the virtual
router.
If the network or router is then deleted the namespaces are not
automatically cleaned up.
With this update, 'neutron-netns-cleanup' now removes these
superfluous namespaces.
```

BZ#985954

```
Previously, floating IP addresses were allocated by Networking,
but displayed by Compute. Synchronization was not always timely,
and resulted in a delay for the allocation to reflect in
Dashboard.
with this update, Dashboard receives floating IP allocation data
directly from Networking, resulting in faster reflection in the
Dashboard view.
```

BZ#1053727

```
With this update, floating IP addresses now have an operational
status: 'Active', 'Down', or 'Error'.
The operational status can be viewed with Dashboard and the
command-line interface.
```

BZ#1063583

```
This enhancement allows Networking to use additional CPU cores
for processing agent requests. This multiprocessing feature was
added to mitigate CPU contention issues.
Consequently, Networking is able to process more concurrent
```

requests.
To enable this feature, edit /etc/neutron/neutron.conf:
Uncomment the 'api_workers' and 'rpc_workers' options to commit any number of CPU cores.

BZ#1067211

This enhancement adds callbacks that notify Compute when a VIF has been plugged in and is ready on the host. In addition, Compute is notified when a floating IP address has been assigned or removed from the VIF.
This was added to prevent occurrences where instances started before the VIF was ready, which resulted in no IP address allocation received from the DHCP server.
Compute receives the notifications from Networking by default. Compute will not boot the instance if the notification has not been received; this will also apply if the VIF was plugged, but the notification was not received.
Networking now contains configuration values to notify Compute using the API that a VIF was plugged in, and that a floating IP address was assigned, updated, or removed.
The default configuration is to notify Compute, however the settings for the Compute API are not set by default; operators need to ensure that the Compute parameters are correctly configured.

BZ#1086004

Prior to this update, certain Qpid exceptions were not properly handled by the Qpid driver.
As a result, the Qpid connection would fail and stop processing subsequent messages.
With this update, all possible exceptions are handled to ensure the Qpid driver does not enter an unrecoverable failure loop. Consequently, Networking will continue to process Qpid messages, even after major exceptions occur.

BZ#1098121

Previously, 'neutron-vpn-agent' did not apply configuration options from the fwaas_driver.ini file, due to a missing argument in the service files. In addition, the legacy L3 agent did refer to the 'fwaas_driver.ini' file. This resulted in inconsistent L3 agent configuration procedures.
With this update, service files have been updated to ensure the configuration file is read on agent startup.
As a result, configuration in the fwaas_driver.ini file is now applied to neutron-vpn-agent.

BZ#1099261

Previously, policy profiles from N1KV VSM were not immediately visible in the Cisco N1KV Neutron plugin. This was due to delayed processing in the 'List events' API.
This update removes the 'List events' API from the Cisco N1kv

Neutron plugin. As a result, policy profiles from N1KV VSM are quickly updated in the Cisco N1KV Neutron plugin.

BZ#1102239

Previously, disttools configuration did not apply certain plug-in configuration files. As a result, these configuration files were not installed and packaged into Networking.
With this fix, disttools has been configured to install all plugin configuration files present in the Red Hat Enterprise Linux OpenStack Platform 5 (Icehouse) release.
Consequently, all plugin configuration files are now packaged into their corresponding plugin packages.

BZ#1106854

Previously, Networking would fail to reliably communicate with Qpid. This behavior was due to an incorrect message subject set in the Qpid layer used by Networking.
This update addresses this issue by setting a correct subject when sending a Qpid message. As a result, Networking now works reliably with the new Qpid server.

BZ#1108960

Previously, concurrent requests for network port and subnet management would result in database lock errors in the Cisco N1kv tables. This would result in the failure of the respective resource request.
With this update, the plug-in code has been revised to ensure that calls to the VSM are external to the database session.
Consequently, database locks are avoided on N1kv tables in the Cisco plug-in.

BZ#1108962

Prior to this update, an internal server error was returned when selecting a network profile with the Cisco plug-in. This behavior was a result of erroneous validation of the subtype field.
With this fix, the subtype field is properly validated, and returns a meaningful error message on validation failure.
Consequently, no internal server error occurs when the subtype is missing from a request.

Chapter 8. RHEA-2014:0855 — Red Hat Enterprise Linux OpenStack Platform Enhancement - Dashboard

The bugs contained in this chapter are addressed by advisory RHEA-2014:0855. Further information about this advisory is available at https://rhn.redhat.com/errata/RHEA-2014-0855.html.

8.1. python-django-horizon

BZ#895586

Previously, when the status of an instance updated to 'Error', there was no feedback displayed on the UI. With this update, information about the error displayed for 'nova show <instance_id>' is displayed on the Instance Details UI page.

BZ#999393

A container that is not empty cannot be deleted. Previously, two contradictory messages, one for success and one for warning that a non-empty container cannot be deleted were displayed when a user tried to delete a non-empty container. With this fix, only the warning message is displayed.

BZ#1006736

Previously, live-migration for the Compute service was only possible using the command line interface. There was no UI option. With this enhancement, a new 'Live migrate' option is displayed in the menu for each instance, on the Instances page when logged in as an admin.

BZ#1012885

Previously, when a container was named with '%' character, the URL in the Dashboard would break, causing a 'Bad Request' page to be returned. As a result, Horizon would fail to create containers with the name containing '%' character. It would also fail to list, upload or download files from such a container.
With this bug fix, the URLs are quoted properly. Users can now create containers with '%' in the name using the Dashboard and also list, upload and downloads files from such a container.

BZ#1029719

Previously, unlimited quotas were represented in the code as a float ('inf') parameter. As a result, 'inf' was directly displayed in the strings related to the quotas on the Project overview page. This value is not translatable and caused issues in other languages.

With this update, 'inf' has been replaced by 'No Limit' and marked as translatable. The Overview now displays 'No Limit' and can be translated to other locales.

BZ#1033117

A new enhancement to create and edit Availability Zones using the Dashboard has been added with this release. This is a useful Compute service feature to manage groups of Compute nodes. Users can now manage Availability Zones, a property of the Host Aggregates, using the new Host Aggregates panel in the Admin Dashboard.

BZ#1033132

A new enhancement to create and edit Host Aggregates using the Dashboard has been added with this release. This is a useful Compute service feature to manage groups of Compute nodes. Users can now create, edit and delete host aggregates using the new Host Aggregates panel in the Admin Dashboard, as well as add and remove hosts from an aggregate.

BZ#1035790

If the Block Storage service endpoint was set to Cinder v2, Dashboard displayed a 500 error when trying to display volume-related pages. With this update, Cinder v2 is supported and the OPENSTACK_API_VERSIONS dictionary in the local_settings file can now take a 'volume' attribute that can be set to either 1 or 2, depending on the Block Storage version the administrator wants to use.

BZ#1041965

With this update, Role Based Access Control (RBAC) support has been added for Block Storage service. After copying the Block Storage policy.json file to /etc/openstack_dashboard/cinder_policy.json, some action such as Delete, etc. are displayed only when allowed by the policy.

BZ#1041976

With this update, a new look and feel was introduced for the dashboard featuring a horizontal instead of vertical navigation bar. This significantly differs from the new accordion navigation bar used upstream.

BZ#1041977

To add features or new components to the Dashboard, changes to the local_settings file were necessary.
With this update, a plugin architecture for horizon is introduced. As a result, users can now place a config file in the

/openstack_dashboard/enabled and restart the web server.
For more information, see
http://docs.openstack.org/developer/horizon/topics/settings.html#pluggable-settings-for-dashboards.

BZ#1041981

With this update, support for public container is added. This allows the user to set ACL as 'read' for the container to either public or private. This can be set at the container creation time or can also be updated at a later time by the user.

BZ#1042060

Pseudo-folders are similar to folders in your desktop operating system. They are virtual collections defined by a common prefix on the object's name.

With this update, users can create pseudo directories using the Create Pseudo-folder button without uploading an object.

BZ#1043717

With this update, Block Storage has the ability to 'extend' (that is, expand or resize) a volume using the Dashboard.

BZ#1056388

With this update, when logged in as an admin, a new tab titled Host Aggregates is available. This can be used to create and manage host aggregates and availability zones.

BZ#1056878

With this enhancement, you can create a new volume as a copy of an existing volume using the Create Volume option on the Dashboard.

BZ#1057830

With this update, Role Based Access Control (RBAC) support has been added for Image service. It is now possible to configure access to images via the
/etc/openstack_dashboard/glance_policy.json file.

BZ#1057831

With this update, the functionality to add and modify an object or its information was added to the Dashboard.

BZ#1063585

With this update, when logged in as an admin, Resource Usage tab

in the System Panel has a new Daily Report tab. You can select a period and generate a report per Project. This is useful to know how much resources various projects are using across all services.

BZ#1067217

With this update, a new disk configuration option is added to the Dashboard. In the Launch an Instance window, there is a new Advanced Options tab for disk partitioning. You can select either Automatic or Manual option from the dropdown list.

8.2. redhat-access-plugin-openstack

BZ#1091367

With this update, the Red Hat Access Plugin for Red Hat Enterprise Linux OpenStack PLatform now uses the new modular Red Hat Access Angular UI framework.

Chapter 9. RHEA-2014:0846 — Red Hat Enterprise Linux OpenStack Platform Enhancement - Packstack

The bugs contained in this chapter are addressed by advisory RHEA-2014:0846. Further information about this advisory is available at https://rhn.redhat.com/errata/RHEA-2014-0846.html.

9.1. openstack-packstack

BZ#895042

This update adds the ability to set the password for the keystone_admin user when running PackStack in interactive mode.

BZ#914648

This update introduces support for generating and distributing SSH keys to provide support for migrating instances via connections using QEMU and SSH to remote hypervisors. This update also ensures QEMU migration ports are open for Nova compute nodes and that Nova passes 'no_verify' to libvirt, making it possible for Nova to access compute nodes correctly when resizing instances.

BZ#1072070

Previously, the PackStack all-in-one installation process would fail under certain conditions when more than one Swift disk was specified. This was caused by the logic used to deploy firewall rules during the swift.pp puppet run, whereby the rules would be deployed for each device rather than for each host, resulting in duplicate entries. Now, the logic used to deploy firewall rules has been revised so that those rules are only deployed once for each host, making it possible to specify more than one Swift disk.

BZ#1072268

This update adds support for automatically mapping GRE and VXLAN network interfaces with the bridge during the PackStack all-in-one installation process.

BZ#1084461

Previously, the cinder-backup service would not be started correctly under certain conditions after the PackStack all-in-one installation process completed. This was caused by the logic used to set up the cinder-backup service, whereby there was no dependency in the puppet used to install the Cinder component, resulting in the puppet attempting to start the cinder-backup

service prior to running the 'cinder-manage db_sync' operation. Now, this logic has been revised so that the cinder-backup service is always started after the 'cinder-manage db_sync' operation is executed.

BZ#1096154

Previously, PackStack would fail under certain circumstances when attempting to install Nagios. This was caused by PackStack attempting to install the nagios-plugins-nrpe package, which was not present in the OpenStack repositories, resulting in a fatal error. Now, the logic used to install this package has been revised so that PackStack will attempt to install the nagios-plugins-nrpe package if that package is available, or install the monitoring-plugins-nrpe package if the nagios-plugins-nrpe package is not available, allowing PackStack to install Nagios successfully.

BZ#1097306

Previously, the PackStack all-in-one installation process would time out and fail under certain conditions when QPID was specified as the messaging service. This was caused by the logic used to interact with QPID, whereby PackStack would specify a deprecated option for QPID during installation of Glance and would specify a different location for the qpidd.conf file for Red Hat Enterprise Linux 7.0. Now, this logic has been revised to update the deprecated option and specify the same location for the qpidd.conf file for all versions of Red Hat Enterprise Linux, allowing the all-in-one installation process to complete all affected steps successfully.

BZ#1098716

This update adds the ability for PackStack to install an L3 metering agent. Installation of the metering agent is controlled by a new command line option (--os-neutron-metering-agent-install) and a new Boolean in PackStack answer files (CONFIG_NEUTRON_METERING_AGENT_INSTALL), which install the metering agent on all nodes running L3 agents when enabled.

BZ#1100993

Previously, virtual machines in a Red Hat OpenStack environment configured using the PackStack all-in-one installation process would have no network connectivity under certain circumstances. This was caused by the logic used in the PackStack all-in-one installation process, whereby the VXLAN and GRE ports required for virtual machine network connectivity would not be opened in the firewall. Now, this logic has been revised so that these ports are opened correctly during the PackStack all-in-one installation process, making it possible for virtual machines to communicate over the network without having to manually configure the firewall.

BZ#1101134

This update provides corrections to the default values of several keys in answer files generated by PackStack, changing value of the default CONFIG_NEUTRON_L2_PLUGIN to 'ml2' and switch default segregation type (the NEUTRON_OVS_PLUGIN and NEUTRON_ML2_PLUGIN keys) to 'vxlan'.

BZ#1103695

Previously, PackStack would fail under certain circumstances when attempting to install Nagios. This was caused by PackStack attempting to install the nagios-plugins-nrpe package, which was not present in the OpenStack repositories, resulting in a fatal error. Now, the logic used to install this package has been revised so that PackStack will attempt to install the nagios-plugins-nrpe package if that package is available, or install the monitoring-plugins-nrpe package if the nagios-plugins-nrpe package is not available, allowing PackStack to install Nagios successfully.

BZ#1105166

Previously, the PackStack all-in-one installation process would fail during the Neutron puppet run due to a syntax error. This would occur when the character ':' was included in the name of a network interface, which the PackStack all-in-one installation process was not able to parse correctly. This update provides support for the characters '.', '-', and ':' in network interface names, making it possible for the PackStack all-in-one installation process to correctly configure network interfaces with names that contain these characters.

BZ#1106394

Previously, Red Hat OpenStack environments deployed using the PackStack all-in-one installation process would not be able to perform autoscaling under certain conditions. This was caused by the default value of the 'keystone_ec2_uri' configuration key being written to the the heat.conf file during the installation process, which would prevent Heat from authenticating against Keystone correctly. With this update, the correct value required by Heat is now explicitly specified and written to the heat.conf file, making it possible for Heat to authenticate correctly against Keystone.

BZ#1106512

Previously, the PackStack all-in-one installation process would fail during the Cinder puppet run under certain conditions when configured to use a Gluster mount. This was caused by the logic used to handle volume declarations, whereby iSCSI was defined by

default and only one volume declaration could be included in the manifest. Now, the iSCSI definition has been moved to a separate template, allowing the PackStack all-in-one installation process to complete successfully when Gluster has been specified.

BZ#1109250

This update ensures that the Audit daemon is installed and running on each node (Controller, Network, and Compute nodes).

BZ#1109308

This update adds availability of the openstack-selinux package for Red Hat Enterprise Linux 7.0 in Red Hat OpenStack environments.

BZ#1109362

This update adds support for restarting libvirtd during the PackStack all-in-one installation process, ensuring that all filters loaded during installation are correctly defined at the end of the installation process.

BZ#1112019

This update resolves an issue in which the Packstack all-in-one installation process would open TCP ports 67 and 68 for DHCP instead of UDP ports 67 and 68.

BZ#1114121

Previously, there was a race condition when shutting down the firewalld service and starting the iptables service. While it was intended that shutting down the firewalld service before calling the puppet firewall class would resolve this issue, puppet resource ordering with before does not work with classes. Now, the puppet iptables service is used for ordering, making it possible to shut down the firewalld service and start the iptables service successfully under these conditions.

BZ#1114261

This update fixes an issue in which LBaaS agents would be installed on the controller instead of the network nodes when installed using the PackStack all-in-one installation process and a network node was configured independently from the controller. Now, LBaaS agents are installed on the network nodes.

BZ#1114262

Previously, the Swift proxy service would fail to start in Red Hat OpenStack environments installed using the PackStack all-in-one installation process. This was caused by the Swift proxy

service attempting to read filters in the proxy-server.conf using names different to those actually written to that file. Now, the names of those filters have been standardized so that the Swift proxy service starts successfully under these conditions.

BZ#1114930

Previously, the PackStack all-in-one installation process would fail under certain conditions when attempting to restart the libvirtd service. This would occur during the Nova puppet run when the PackStack all-in-one installation process was run on a controller node on which Compute had not been deployed. Now, the PackStack all-in-one installation only attempts to restart the libvirtd service when Nova has been deployed on a node on which Compute has also been deployed, allowing the process to complete successfully under these conditions.

BZ#1115163

This update resolves an issue in which the CONFIG_PROVISION_ALL_IN_ONE_OVS_BRIDGE configuration option would always resolve to true. This was caused by a change in the logic used to produce the value against which this configuration option was tested, whereby a test result of false was unreacheable. Now, this logic has been revised so that testing this configuration key returns the correct value.

9.2. openstack-puppet-modules

BZ#1017210

This update provides support for installing MariaDB on Red Hat Enterprise Linux 7 using the puppet-mysql module.

BZ#1093949

Previously, policy files would not be loaded due to the POLICY_FILES override in local_settings. This update removes the override, allowing all default policy files to be loaded.

BZ#1095279

Previously, network connectivity would be lost under certain conditions when running the PackStack all-in-one installation process. This was caused by puppet-neutron not moving the IP address from the network interface controller to the bridge when the bridge is created. Now, the bridge is created before running puppet-neutron, and the IP address is moved to the bridge as follows:

```
ip addr del <ip>/24 dev <nic>
ip addr add <ip>/24 dev br-ex
```

With this update, network connectivity is no longer lost during the puppet-neutron run due to the IP address not being moved to the bridge.

BZ#1109445

Previously, the Puppet modules used to manage OpenStack contained an erroneous configuration for the neutron-ovs-cleanup service. In environments with periodic Puppet runs (such as those deployed using Foreman), this would disable Neutron networking. This change ensures that the neutron-ovs-cleanup service only runs at boot, rather than at every Puppet run.

BZ#1109890

This update increases the default value of the maximum number of simultaneous connections that can be made to MariaDB to 1024 when deployed using the PackStack all-in-one installation process, enhancing support for systems with a large number of cores.

BZ#1110281

With this update, the neutron-server service is now restarted after ML2 configuration changes, ensuring those changes are applied.

BZ#1114583

This update resolves an issue that would break existing Neutron namespaces when the network service was restarted on a Neutron network node.

Chapter 10. RHBA-2014:0937 — Red Hat Enterprise Linux OpenStack Platform Bug Fix and Enhancement Advisory

The bugs contained in this chapter are addressed by advisory RHBA-2014:0937. Further information about this advisory is available at https://rhn.redhat.com/errata/RHBA-2014-0937.html.

10.1. openstack-sahara

BZ#1117903

```
This rebase package includes a number of notable enhancements and
fixes under version 2014.1.1:

* Fixed validation of node group templates using Nova networking
with auto assignment of floating IPs.
* Fixed occasional failure to launch EDP jobs on transient
clusters.
* Fixed issue with transient cluster creation failing due to
incorrect tenant id.
* Updated architecture diagram and IRC channel name in
documentation.
```

10.2. openstack-selinux

BZ#1116755

```
In previous releases, using LBaaS would generate an SELinux AVC
denial in the audit log. This was because the Networking service
was not explicitly allowed to connect to sockets with the
haproxy_t label. Note that even when this occurred, the SELinux
AVC denial did not prevent LBaaS from functioning properly.

With this release, the Networking service is now explicitly
allowed to connect to sockets with the haproxy_t label, thereby
preventing any further SELinux AVC denial logs.
```

BZ#1117301

```
In the previous release, SELinux in enforcing mode prevented live
migration by not allowing SSH to look at Compute. Consequently,
migrations failed and AVC messages were generated (for example,
'Live Migration failure: operation failed: Failed to connect to
remote libvirt URI qemu+ssh<snip>').
With this update, SSH is now allowed to look at files with the
nova_var_lib_t label, and migrations succeed (without the
appearance of AVCs).
```

BZ#1119151

In the previous release, SELinux prevented Ceph Storage from connecting to an unreserved port. As a result, Ceph was unable to receive traffic due to it being unable to connect to this port.

The Image Service is now allowed to connect to all TCP ports; Ceph is allowed to connect to the unreserved port and receive incoming traffic.

BZ#1119400

Previously, SELinux in enforcing mode prevented image creation in the Image Service (glance) when using Ceph storage (RADOS block devices). This meant that image creation failed, and SELinux generated AVC messages.
With this update, the Image Service can now write to memory with the same label, so that image creation succeeds and no AVC messages are output.

BZ#1119845

In the previous release, SELinux in enforcing mode blocked the attachment of block storage using 'nova volume-attach'. As a result, Compute failed to attach block storage.

With this update, the svirt process in SELinux has been updated and can now write to memory with the same label; Compute's 'nova volume-attach' now succeeds without being blocked by SELinux.

10.3. python-requests

BZ#1115794

An incorrect string length calculation in python-requests caused images larger than 5GB to not be uploaded.

The string length calculation in python-requests has been corrected, so that images larger than 5GB may now be uploaded.

Chapter 11. RHEA-2014:1003 — Red Hat Enterprise Linux OpenStack Platform Enhancement Advisory

The bugs contained in this chapter are addressed by advisory RHEA-2014:1003. Further information about this advisory is available at https://rhn.redhat.com/errata/RHEA-2014-1003.html.

11.1. foreman

BZ#1039011

When the installer or other applications used Foreman's API and an error occurred, the locale to render error messages was not initialized.

As a consequence error messages in a random, non-English locale would be shown to the user.

This has been fixed and the locale is now initialized when using the API, defaulting to US English.

Now, when an error occurs between the API and the installer, error messages will be logged in English.

11.2. foreman-selinux

BZ#1043887

Previously, OpenStack Foreman installer operations resulted in SELinux denials.
This update amends the Foreman SELinux policy. Consequently, denials are no longer generated during installation.

BZ#1123279

A post-install scriplet in the foreman-selinux package was issuing errors during uninstall. This was caused by the Elasticsearch port (9200-9300) not being removed properly before unloading the SELinux policy.

This resulted in the "yum uninstall" transaction being canceled, leaving Foreman in an uninstallable state.

Now, the scriplet has been fixed to remove ports prior to unloading the policy. As a result Foreman now uninstalls cleanly.

11.3. openstack-foreman-installer

BZ#1050203

This enhancement adds support for NFS backends in the 'Block

This enhancement adds support for NFS backends in the 'Block Storage' host group.
As a result, NFS backends for Block Storage (Cinder) are available in the 'Block Storage' host group.

BZ#1064958

This update adds support for NFS backends on non-HA controller host groups.
Consequently, NFS backends for Block Storage are available on host groups 'Controller (Neutron)' and 'Controller (Nova Network)'.

BZ#1091536

This enhancement adds support for Dell EqualLogic as a backend for Block Storage (Cinder).
As a result, Dell EqualLogic is available as a Block Storage backend on the controller host groups: 'HA All In One Controller', 'Controller (Neutron)', and 'Controller (Nova Network)'.

BZ#1094385

The Block Storage (Cinder) service now supports multiple concurrent back-ends. To enable this feature, set the following Host Group parameters to 'true':

* cinder_multiple_backends
* cinder_backend_[backend type] (where [backend_type] is each back-end type you wish to enable)

You will also need to set values specific to the selected back-ends.

BZ#1100369

Prior to this update, a wait condition was present in a Puppet manifest for an event that does not occur when fencing is disabled.
This resulted in an error when running the Puppet agent in deployments where fencing was disabled (typically proof-of-concept).
With this update, the wait condition has been changed to an event that always occurs during deployment, regardless of fencing status.
As a result, the Puppet agent runs successfully in conditions where fencing has been disabled.

BZ#1104093

Modular Layer 2 (ML2) and VXLAN are now the default neutron networking driver/type for Red Hat Enterprise Linux OpenStack Platform 5.

ML2 is the new Networking core plug-in introduced in OpenStack's Havana release. Superseding the previous model of singular plug-ins, ML2's modular design enables the concurrent operation of mixed network technologies. Controllers should now be configured with ML2 as the OpenStack Networking (neutron) driver by default, and the default tenant network type is VXLAN.

The monolithic Open vSwitch and linuxbridge plug-ins have been deprecated and will be removed in a future release; their functionality has instead been re-implemented as ML2 mechanisms.

BZ#1104219

This new feature ensures that the Compute scheduler runs in Active/Active mode for high availability (HA) by default.

This is required because Active/Active mode in HA for the Compute scheduler allows better scaling than the Active/Passive mode.

Now, Compute scheduler runs Active/Active by default. If you wish to make it Active/Passive, simply set scheduler_host_subset_size=1.

BZ#1105218

This enhancement enables support for multiple instances of the Dell EqualLogic backend for Block Storage (Cinder).
As a result, all Dell EqualLogic parameters are now arrays instead of single values, with the exception of 'cinder_backend_eqlx' which remains a true/false switch to enable/disable the EqualLogic backend.
In addition, all EqualLogic parameter arrays are expected to have the same number of elements: the first elements of each array form properties of one instance of the EqualLogic backend, while the second elements of each array form properties of another instance. This pattern continues for subsequent elements.

BZ#1109311

This update adds availability of the openstack-selinux package for Red Hat Enterprise Linux 7.0 in Red Hat Enterprise Linux OpenStack Platform environments.

BZ#1109329

In the previous release, the Telemetry (ceilometer) notification agent was not started on the controller node. As a result, the Telemetry API did not function correctly. Setting have now been updated to start the notification agent, and the Telemetry API now works as expected.

BZ#1110504

Prior to this update, the 'nfs-utils' package was not installed

on controller nodes when NFS was selected as the Block Storage backend.
Consequently, NFS shares designated as Block Storage backends would fail to mount.
This update addresses this issue by installing the 'nfs-utils' package if NFS is selected as the Block Storage backend.
As a result, NFS backends now mount as expected.

BZ#1111158

The Telemetry alarm evaluator and notifier services were not installed on the controller node in a Staypuft-driven deployment.

As a consequence Telemetry alarms were not evaluated against their respective thresholds, so Orchestration's autoscaling events were not triggered.

This has been fixed and the missing services were added as dependencies.

Now Telemetry alarms are evaluated against their respective thresholds.

BZ#1113294

Block Storage's configuration including db_sync happened simultaneously across nodes.

Due to multiple simultaneous cinder-manage db_sync processes, some tables ended up in unexpected states for some of the processes, causing error messages.

To avoid this problem, the first node that runs cinder-manage db_sync now completes before the other HA nodes attempt to run it.

This avoids tables ending up in incorrect states.

11.4. openstack-puppet-modules

BZ#1124027

The vswitch Puppet module provides a new implementation for the vs_port and vs_bridge Puppet resource providers.

1) vs_port

When a physical network interface (port) is associated to an Openvswitch bridge, both the bridge and the physical interface need a network configuration file to be created for the configuration to be resilient. For instance, if 'eth0' is attached to 'br-ex', two files, respectively 'ifcfg-br-ex' and 'ifcfg-eth0' will be created in '/etc/sysconfig/network-scripts/' directory.

```
However, in order to addresses various network configuration
cases (virtual interface, bonding, vlan, etc.) and therefore
potential problems, a new approach was needed.

The new implementation still creates the 'ifcfg' file for the
port (physical interface) the same way, but it uses the physical
interface definitions, if the link is active, to create the
'ifcfg' bridge configuration itself, then it adds the necessary
Openvswitch network entries.

The vs_port (ovs_redhat) provider includes the following
technical changes:
  - Inheritance from the default provider (OVS) class
  - A library to handle ifcfg content
  - Automatic behavior replaces keep_ip and sleep parameters -
Puppet-neutron never implemented them so the change is
transparent
  - Requires Puppet 2.7.8+: using commands instead of
optional_commands

B. vs_bridge

The vs_bridge/ovs_redhat.rb has been removed. The puppet resource
bridge feature relies on the default OVS provider.
```

11.5. rhel-osp-installer

BZ#1116877

```
In a non-HA deployment of OpenStack on Red Hat Enterprise Linux
6.5 that uses LVM as the backing store for Block Storage, the
disk partitioning creates a 500MB boot and 100GB LVM set for
root. The rest of the space is used for a physical volume that is
used for a cinder-volumes volume group that will be use by Block
Storage to create volumes.
```

11.6. rubygem-staypuft

BZ#1103276

```
This update disables Ruby 1.8.7 garbage collection when running
the Red Hat Enterprise Linux OpenStack Platform Installer. Ruby
1.8.7 garbage collection contained a bug that could cause a
segmentation fault in Puppet when running the the Red Hat
Enterprise Linux OpenStack Platform Installer, ultimately
preventing Puppet from completing the installation process.
```

Chapter 12. RHBA-2014:1090 — Red Hat Enterprise Linux OpenStack Platform Enhancement Advisory

The bugs contained in this chapter are addressed by advisory RHBA-2014:1090. Further information about this advisory is available at https://rhn.redhat.com/errata/RHBA-2014-1090.html.

12.1. foreman

BZ#979266

When Foreman provisions hosts, a root password is now requested, instead of being filled in with a default which was insecure.

Users now need to specify a password either during host creation or by providing a default password hash under settings, ensuring the security of provisioned hosts.

12.2. foreman-discovery-image

BZ#1108512

Previously, when trying to discover a VM with two network interfaces, one in the discovery network, and one in an external network with DHCP and DNS, the VM failsed to send facts to Foreman, and did not get discovered by Foreman Discovery.

With this release, Foreman Discovery image is changed so that DNS is configured for the network interface that is specified by BOOTIF parameter. This is a required parameter and therefore IPAPPEND 2 option must be present in the kernel command line. As a result of this fix, the network interface that is configured from boot using LAN is used to configure DNS.

12.3. openstack-foreman-installer

BZ#1104856

This enhancement adds support for CEPH RBD as a backend for Block Storage (Cinder) and Image Service (Glance).
Ceph RBD was previously only available for the 'HA All In One Controller' host group. This support has now been extended to non-HA controller host groups.

BZ#1122701

Previously, on the RHEL 7 hosts, firewalld was not disabled early during puppet run, which resulted in the default iptables rules (for example, rejection of traffic which was not explicitly allowed) not being present.

With this update, on the RHEL 7 hosts, firewalld is removed and iptables is started early in the puppet run, before Puppet attempts to create any custom iptables rules. As a result, default iptables rules (including rejection of traffic which was not explicitly allowed) are now present.

BZ#1123300

Previously, configuration parameter 'check' was missing the 'interval' option defined in HA Proxy which resulted in longer recovery times in case of failure.

With this update, a new configuration option 'check inter 1s' is added to /etc/haproxy/haproxy.cfg file and as a result, service recover as expected after failures.

BZ#1123301

The nova-metadata parameter was not enabled for HA deployments, so it did not function as expected.

This fix enables service in the haproxy configuration file and binds listener to the correct address. As a result, nova-metadata service runs as expected for a HA deployment.

BZ#1123312

Previously, an appropriately long start time had not been assigned for Galera (MariaDB) service under the Pacemaker control which lead to a false error condition as Galera did not start up within the allocated window.

With this update, the start timeout has been increased to 300s and as a result, Pacemaker is able to start up Galera under systemd.

BZ#1123314

Previously, openstack-heat-engine monitor interval parameter value was set too low for Galera (MariaDB) for Pacemaker to identify that Orchestration service was running. This caused the Orchestration service to restart as random intervals.

With this bug fix, the monitor interval has been increased to 60s and as a result, Orchestration service is not unnecessarily restarted by Pacemaker.

BZ#1123318

Previously, the max_retries parameter in various configuration files was not set to -1. As a result, an OpenStack service would 'give up' rather than retry if the initial database connection failed.

With this update, the value for max_retries parameter is set to -1 in the OpenStack configuration files, so services will retry connecting to the database service if the initial attempt fails, allowing for a more robust setup.

BZ#1127887

Previously, when puppet utility ran for the first time with include_galera parameter set to 'true' and other 'include' parameters (such as include_keystone) set to false, the 'rsync' package was not installed in Galera and the puppet utility would fail on the non-bootstrap Galera nodes.

With this update, 'rsync' package is added as a requirement for all Galera nodes and the puppet utility no longer fails on the non-bootstrap nodes and Galera starts as expected.

BZ#1128457

Previously, horizon was not configured to allow requests for all FQDNs. As a result, OpenStack Dashboard would display 'Not Found' for the HA and non-HA controller node if the FQDN in the web request was not explicitly allowed in the Apache configuration.

This update allows Apache to serve all requests to the Dashboard, regardless of the web address specified in the user's web request (URL) and Dashboard loads without errors.

BZ#1130304

Previously, the nfs-utils package was not installed before a Image sevice or Block Storage service NFS export was attempted to be mounted by a HA controller node and so the NFS mount operation failed.

This update ensures that the nfs-utils package is installed before any NFS mounts are made. As a result, NFS mount now succeed during the initial puppet run on the HA controller nodes, whether or not the nfs-utils package has already been installed before puppet is executed.

12.4. rhel-osp-installer

BZ#1117019

When an administrator wants to use single or multiple external networks with VLAN ID, the dashboard interface requests the VLAN mappings information and sets them in the configuration file. As a result, when using the CLI, the user needs to use the same information, else the external network will fail. If the VLAN ID needs to be modified, then the user needs to update the parameter in Foreman.

As a fix to this issue:
1. An empty mapping for 'network_vlan_ranges' parameter as "physnet-external" is set as follows:
[("physnet-tenants:#{self.tenant_vlan_ranges}" if self.vlan_segmentation?), ("physnet-external")].compact

2. In the Dashboard, '[ ] Use VLAN ID _____ for External Network' is removed.

This fix allows the administrator to create an external network to be configured on any VLAN+VID, or a flat external network without VLAN.

BZ#1122753

In a non-HA deployment of OpenStack on Red Hat Enterprise Linux 6.5 that uses LVM as the backing store for Block Storage, the disk partitioning creates a 500MB boot and will initially split the remaining space equally between physical volumes for the root and cinder-volumes. The PV for root will cap at 100G and the rest of the space is allocated for cinder-volumes.

BZ#1123463

Previously, an initial puppet run during provisioning could result in puppet errors. The initial puppet run is only done to exchange certificates and errors in other puppet actions are not fatal. As a result, this would lead to erroneous reports of failed deployments.

This has been fixed by ignoring errors in the first puppet run. Now, deployments no longer fail erroneously.

BZ#1123492

In some deployment scenarios, puppet configures a NIC as part of a bridge. As a consequence, if NetworkManager is running, this change causes the puppet agent to terminate when the NIC being changed is the one in use.

This has been fixed by disabling NetworkManager, so now puppet-runs no longer get killed mid-run as a result of a valid configuration change.

BZ#1124545

Previously, the provisioning wizard did not properly save the values for skipping subscription seeding or provisioning repositories. As a result, when running the provisioning wizard in the non-interactive mode without the values, the wizard would throw an 'Unprocessable Entity' stacktrace.

With this fix, you can update the answer file according to the following steps:

1. Add the following values to the :custom: section on the answer file:
    :skip_subscription_seeding: (true|false)
    :skip_repo_path: (true|false)
2. Save the answer file properly and retrieve these values from the provisioning wizard when run interactively.
As a result, the provisioning wizard will no longer throw an error.

BZ#1125136

An ordering issue in the puppet classes was not waiting for firewalld to completely shut down before starting iptables.
As a result, iptables would be started too soon and the firewalld process would kill it.

The ordering has been fixed and now puppet waits for firewalld to stop completely before starting iptables.

BZ#1126219

A new feature has been added which gives users the ability to specify proxy information for subscription-manager details.

This was added because many environments require a proxy between the host and CDN/RHN.

Now, users can specify proxy information in the installer.

BZ#1126982

Previously, both 'biosdevname' and 'net.ifnames' were enabled by default, resulting in some hosts booting with 'biosdevname' NICs and other using 'net.ifnames'.

With this update, 'biosdevname' has been disabled and all hosts now use the 'net.ifnames' (for example, ens8).

12.5. rubygem-staypuft

BZ#1120426

Foreman discovery plugin generated a hostname from the host's MAC address. If the MAC address is all numeric, the resulting host name was all numeric. But all numeric hostnames are a violation of internet protocol and caused the glibc resolver to fail.

As a workaround, change the hostnames of the machines you discover to include at least one non-numeric character and as a result, the DNS resolution succeeds.

BZ#1124494

Currently, there is no validation to check if a network interface is correct. As a result, issues could occur later in a deployment.

As a workaround, use only lower case to enter the interface names. For example, use eth0 instead of ETH0.

BZ#1127297

When changing layout for a Deployment with hosts already added (i.e. non-HA to HA, Neutron to Nova networking, etc), any hostgroups that must be deleted to make the layout change will have any assigned hosts removed first. Changing from HA to non-HA (or vice versa) will delete the Controller hostgroup and create a new one. Changing from Neutron to Nova (or vice versa) will delete the Compute hostgroup (and Controller too for non-HA). Changing from Non-HA Neutron to HA or to Nova networking will delete the Neutron Network Node hostgroup.

Chapter 13. RHBA-2014:1324 — openstack-packstack and openstack-puppet-modules bug fix advisory

The bugs contained in this chapter are addressed by advisory RHBA-2014:1324. Further information about this advisory is available at https://rhn.redhat.com/errata/RHBA-2014-1324.html.

13.1. openstack-packstack

BZ#1053734

With this enhancement, PackStack now consistently performs the installation of the sos, sos-plugins-openstack and rhos-collector packages on all hosts.

BZ#1098765

This Packstack enhancement adds FireWall as a Service (FWaaS) when Openstack Networking (neutron) is selected.
FWaaS is optional and not activated by default. Activate FWaaS using either the command line, or the answer file.

BZ#1103148

This PackStack enhancement adds Load Balance as a Service (LBaaS) when Openstack Networking (neutron) is selected.
LBaaS is optional and not activated by default. Activate LBaaS using either the command line, or the answer file.

BZ#1108155

With this enhancement, PackStack configures the Orchestration service (heat) to use trusts by default. For the Orchestration service to work with trusts, each user must have a role for delegation (by default, this role is 'heat_stack_owner').

BZ#1128303

This Packstack update adds parameter validation for CONFIG_SWIFT_STORAGES, which accepts the same values and deprecates CONFIG_SWIFT_STORAGE_HOSTS. Values of CONFIG_SWIFT_STORAGES must be a comma-separated list of paths to devices, for example: /path/to/device

BZ#1134069

New PackStack parameters have been added for specifying a HTTP proxy for 'subscription-manager'. The new parameters are:

* CONFIG_RH_PROXY
* CONFIG_RH_PROXY_PORT
* CONFIG_RH_PROXY_USER

```
* CONFIG_RH_PROXY_PW

Note that the parameters CONFIG_RH_USER and CONFIG_RH_PW have to
be completed in order for PackStack to successfully set a HTTP
proxy on hosts.
```

Chapter 14. RHBA-2014:1325 — Red Hat Enterprise Linux OpenStack Platform Bug Fix and Enhancement Advisory

The bugs contained in this chapter are addressed by advisory RHBA-2014:1325. Further information about this advisory is available at https://rhn.redhat.com/errata/RHBA-2014-1325.html.

14.1. ceph

BZ#1133145

```
The Ceph client has been rebased to the ceph-0.80.5-1
version. Fixes include:
* When a computer with CephFS that was mounted using ceph-
fuse woke up from suspend, the /mnt/ceph mount did not
respond (that is, `df -h /mnt/ceph` never returned). The
sleep-recover and client-sleep functions have been fixed so
that this no longer occurs.
```

14.2. galera

BZ#1090604

```
In the past, if Galera was configured to use SSL, the client
certificate was not provided during the SSL handshakes when an
IST (incremental state transfer) occurred. As a consequence the
IST operation failed, causing the joining node to crash.

This has been fixed and now Galera provides the client
certificate as part of SSL handshake, so IST operations succeed
when SSL is enabled.
```

14.3. heat-cfntools

BZ#1138865

```
The packaging for heat-cfntools was missing a few dependencies
required for some of the external commands called in the
cfn_helper script. Wget specifically has been replaced by curl to
maintain functionality while reducing a required dependency.

These dependencies are now added to the package.
```

14.4. openstack-ceilometer

BZ#1085998

There was an internal error in the python-qpid library which Telemetry would fail to handle gracefully and Qpid communication would be broken.

This has been fixed so that Telemetry gracefully handles the failure and restarts Qpid communication.

Now, Telemetry services recover after an internal error in the python-qpid library.

BZ#1120990

Previously, Telemetry HTTP alarm notifications did not include a Content-Type HTTP header. As a consequence, the handler of the notification was unable to determine the content type of the notification.

This has been fixed by including a Content-Type HTTP header in the HTTP alarm notification, set to 'json'.

Notification handlers can now correctly determine the content type of the notification.

BZ#1123376

Previously, if Telemetry encountered a failed connection to a message broker, re-connection attempts kept failing as well. This was due to the Telemetry service trying to reconnect to the same failing message broker, even if there were several hosts configured.

This has been fixed by making the reconnect() implementation select the next broker in the list. As a result, when several broker hosts are provided, it will try the next one in the list at every connection attempt.

This means that non-failure reconnect attempts will also switch from the current broker to the next in the list. Hence, users should not rely on any particular order when using brokers from the list.

14.5. openstack-cinder

BZ#1026202

With this enhancement, Block Storage backup support has been added to the NFS driver.

BZ#1102340

Previously, OpenStack Block Storage (cinder) would retry any vSphere API call which failed, not just those which failed due to authentication or temporary network issues.

Additionally, the code to automatically recreate a session could fail spuriously under some circumstances. This meant that any vSphere API call that failed would be called multiple times, which increased the load on the vSphere server. Also, API failures could result in a failure to re-establish a new session, causing all subsequent API calls to fail.
With this update, vSphere API calls are only retried for transient network or session errors, and session recreation is independently retried for connection errors. This reduces the load on the vSphere server, and sessions are reliably recreated if they expire or fail.

14.6. openstack-selinux

BZ#1135637

Previously, when running puppet as a systemd service, SELinux would deny its services. As a result this caused rsync to fail repeatedly.

This has been fixed by giving rsync the ability to relabel a few directories and files so it has access rights to them.

Now, rsync runs successfully with SELinux in enforcing mode.

Chapter 15. RHBA-2015:0825 — Red Hat Enterprise Linux OpenStack Platform Bug Fix and Enhancement Advisory

The bugs contained in this chapter are addressed by advisory RHBA-2015:0825. Further information about this advisory is available at https://rhn.redhat.com/errata/RHBA-2015-0825.html.

15.1. openstack-sahara

BZ#1101516

Previously, SQLite database was created by a user who ran the database management script, resulting in Sahara being unable to read the default database without changing ownership of the database.

With this update, the file is not touched and the ownership is assigned to Sahara (for only the default file location). As a result, Sahara now has access to its database in the default flow.

BZ#1163420

Previously, the log directory permissions for Sahara was set to 755, resulting in the Sahara service not conforming to the Red Hat log security standards.

With this update, the directory permissions are modified to 750, thus, conforming to the Red Hat log security standards.

15.2. openstack-selinux

BZ#1149975

Previously, SELinux prevented the nova scheduler from searching directories labeled 'cert_t', resulting in SELinux causing Compute to fail.

With this update, an 'allow' rule has been created to give the nova scheduler permission to search the 'cert_t' directories. As a result, Compute service functions normally.

15.3. openstack-utils

BZ#1133920

If an existing haproxy process was already running before installing and running LBaaS (Load-Balancing-as-a-Service),

attempting to start LBaaS will fail. This typically happens when upgrading to Red Hat Enterprise Linux OpenStack Platform 5 with an existing LBaaS service.

To work around this, you will have to kill the running haproxy process and restart the LBaaS agent:

```
# kill $(pgrep haproxy)
# service neutron-lbaas-agent restart
```

15.4. python-novaclient

BZ#1139413

With this enhancement, a new command 'service-delete' has been added to the nova client to allow disabling services through the nova CLI as opposed to manually editing the nova-services table.

BZ#1147958

Previously, the 'nova list' was inefficient and took very long to complete as the number of instances increased.

With this update, 'nova list' command code has been optimized and uses server-side filtering, resulting in faster response.

15.5. python-sqlalchemy

BZ#1121796

Previously, an improvement to the connection pool such that new connections could be made concurrently, made it so that the 'init on first connect' routine of a SQLAlchemy dialect would not have been completed if concurrent routines proceeded at the same time. As a result, when a SQLAlchemy engine was first used, operations which relied on the state acquired during initial startup could fail, as this information would not have been completed.

To resolve this issue, with this update, 'mutexing' was added to the event system which handles the initial dialect startup phase, so that connection attempts are again serialized, but only when the engine first starts up.

BZ#1121798

Previously, the MySQL-Python DBAPI was observed under some circumstances using the ProgrammingError exception class to report on the 'command out of sync' errors, which is considered to be the case where a connection need to be thrown away; the SQLAlchemy dialect only expected this error to be emitted within the OperationalError class. As a result, in some cases a MySQL-Python connection that became corrupt would not signal to the

SQLAlchemy engine that the pool of connections should be disposed, leading the engine not being able to proceed with new operations.

With this update, the error handling scheme of the MySQL-Python dialect is modified to expect either the OperationalError or ProgrammingError exception class when testing for this particular class of error. As a result, the SQLAlchemy engine/connection pool now correctly disposes off its connections when a MySQL-Python ProgrammingError delivers the 'command out of sync' error code.

15.6. rabbitmq-server

BZ#1126680

Previously, the rabbitmq-plugins command was not available in the default path. As a result, trying to run rabbitmq-plugins command would result in a 'Command Not Found' error.

With this update, the rabbitmq-plugins command is added to the default path and it executes as expected.

Chapter 16. RHSA-2015:0843 — Important: openstack-nova security, bug fix, and enhancement update

The bugs contained in this chapter are addressed by advisory RHSA-2015:0843. Further information about this advisory is available at https://rhn.redhat.com/errata/RHSA-2015-0843.html.

16.1. openstack-nova

BZ#1134992

Previously, Nova would default to cold snapshots of instances. As a result, instances needed to be shutdown before the snapshot was taken.

With this enhancement, Nova now uses live snapshots by default. Instances remain powered on during snapshots, and the process is unnoticeable to the user.

BZ#1151114

The ephemeral disk format (for example, 'ext4' or 'xfs') was previously not taken into account if using the nova boot command ([--ephemeral size=<size>[,format=<format>]]). Instead, the default format was used. With this update, Compute's libvirt driver now uses the specified format.

BZ#1151150

A previous overly restrictive ban on live migration of vfat config drives, and the incorrect handling of config drives with RBD storage, meant that the live migration of instances with config drives was not supported.

With this update, vfat config drives can now be live migrated, and config drive persistence is handled appropriately with RBD storage. This means that live migration is now possible when using vfat config drives with storage either local to the compute node or remote with RBD storage. (To use vfat config drives, set config_drive_format in /etc/nova/nova.conf to 'vfat'.)

BZ#1170212

The arguments to the multipath rescan command used by the libvirt volume driver were incorrectly formatted, which led to the rescan failing. The multipath command was invoked with two arguments '-' and 'r', instead of a single argument '-r'. With this update, the command-line argument is correctly formatted, and the multipath command can now succeed.

BZ#1174422

Previously, the evacuate function did not consider RBD storage as shared and the evacuate procedure failed with RBD-backed instances. With this fix, RBD storage is now marked as shared, and the evacuate function handles the shared storage attribute and therefore now operates on RBD.

BZ#1180600

When Compute is configured to only set up VNC/SPICE servers on a specific network interface, the host's IP address is recorded in the libvirt guest XML. Previously, if the guest was migrated to a different host, the IP address of the source host remained in the guest XML and the guest failed to launch on the target host because the IP address was incorrect.

With this update, the libvirt guest XML is now updated during migration to refer to the IP address of the target host. Migration can be performed for guests, even when the VNC/SPICE servers are configured to only bind to the IP address of a specific network interface.

BZ#1188355

When using the command "nova host-evacuate" with the option "on-shared-storage", the instance was evacuated, but the guest was rebuilt using the original image. With this fix, the Compute API has been updated so that a rebuild is no longer requested after evacuating the instance. As a result, when evacuating an instance with shared storage, the instance is now moved to the new compute host, and the root disk is not rebuilt.

BZ#1191696

Previously, the Compute service incorrectly handled exceptions when migrating instances between different OpenStack versions. This meant that an instance migrated from an older version would appear to hang forever in the migrating state. With this update, the exception for a 'forbidden version' is now handled correctly, and migrations are properly disallowed.

BZ#1199106

The Compute service has been rebased to version: 2014.1.4

Important fixes and enhancements include:

* Security fixes. The websocket proxy of the Compute service console now verify the origin HTTP header to block cross-site attacks. CVE 2013-2255: Local CAs are now verified by default. By default, SSL certificate verifications are disabled. A new attestation_insecure_ssl option was added to enable verification by setting the option to False.

* Block device mapping retries are now configurable, with two new

```
configuration
  options: block_device_allocate_retries (the number of block
device
    mapping retries) and block_device_allocate_retries_interval
(the time interval
    between consecutive retries).

* Two new configuration options have been added to control keep-
alive and client connection
  timeout: wsgi_keep_alive option (default=True),
client_socket_timeout option (default=0).

* Fixed issue with the Compute service not doing Image service
server certificate validation.

* Fixed instance root-disk size restriction with QCOW2 images.

* Fixed the initialization sequence of nova-compute service to
handle binding   failures of virtual interface. Failures are now
logged when nova-compute starts. Before, nova-compute failed to
start.

* Set a check for minimum disk and RAM when booting from a
volume. Previously,
  the minimum attributes were ignored.

* Fixed a multipath iSCSI sessions issue when connecting or
disconnecting a
  volume.

* Fixed a race condition in the creation of security groups.

* Fixed the resource tracking and now updates the number of
instance during delete
  instance.

* Fixed a Compute service evacuate issue with RDB.

* Fixed nova-compute start issue after evacuate.

* Fixed denial-of-service issue in instance-list IP filter.

* Now retry on closing of LUKS-encrypted volume in case device is
busy.

* Now share OpenStack Networking admin authentication tokens
resizing.

* Fixed a bug in cell management which prevented the start of the
nova-cells service.

* Fixed instance cross-AZ check when attaching volumes.

* Now ignore errors when deleting non-existing VIFs so that
instances are not left in the state of "Error(deleting)".
```

BZ#1199142

The previous setting of 'iscsi_use_multipath=true' in nova.conf meant that detaching a multipath iSCSI volume killed all iSCSI volumes visible from the Compute service's compute node.

There are two types of iSCSI multipath devices. One which shares the same IQN between multiple portals, and the other which uses different IQNs on different portals. With this update, connect_volume() now identifies the type by checking iscsiadm (the output is the IQN is used by multiple portals), and then connecting to the correct targets using connect_to_iscsi_portal().

BZ#1201851

When providing an affinity of the anti-affinity server group for two or more VMs, Compute previously only checked when the instance was booted but not with instance migration. This meant that two VMs in the same anti-affinity group could end up on the same compute host. With this update, the groups policy of a migrating instance is now checked, so that the policy is not violated when migrating.

BZ#1205024

Because the Compute service did not include a UUID in the XML it sent, it would always get a random UUID generated, which caused failures when re-defining an existing network filter (name and UUID must match). As a result, the Compute service would fail to start up and fail to migrate if there was an existing guest running.

With this update, libvirt is now queried to see whether the nwfilter already exists. If the filter exists, Compute now extracts the UUID from its XML and uses it when re-defining the filter so that the service can successfully start up.

BZ#1205806

Previously, the Compute service used incorrect size checks against the QCOW2 virtual size, which meant that the root-disk size for a flavor was not honored. With this fix, size checks have been reinstated, and instances are again restricted to the root-disk size of a flavor.

16.2. vulnerability

BZ#1154890

A flaw was found in the OpenStack Compute (nova) VMWare driver, which could allow an authenticated user to delete an instance while it was in the resize state, causing the instance to remain

on the back end. A malicious user could use this flaw to cause a denial of service by exhausting all available resources on the system.

BZ#1154951

A denial of service flaw was found in the way OpenStack Compute (nova) looked up VM instances based on an IP address filter. An attacker with sufficient privileges on an OpenStack installation with a large amount of VMs could use this flaw to cause the main nova process to block for an extended amount of time.

BZ#1190112

It was discovered that the OpenStack Compute (nova) console websocket did not correctly verify the origin header. An attacker could use this flaw to conduct a cross-site websocket hijack attack. Note that only Compute setups with VNC or SPICE enabled were affected by this flaw.

Appendix A. Revision History

Revision 5.0.5-1 Tue November 17 2015 Don Domingo
Updated to reflect 7.2 minimum requirement.

Revision 5.0.5-0 Thu September 10 2015 Andrew Dahms
Updated overview to include 5.0.5 advisories.

Revision 5.0.4-0 Thu April 16 2015 Summer Long
BZ#1211637 - Updated overview to include 5.0.4 advisories. Chapters added for RHBA-2015:0825 and RHSA-2015:0843.
BZ#1210112 - Updated to reflect 7.1 minimum requirement.

Revision 5.0.0-16 Tue December 2 2014 Summer Long
Updated overview to include 5.0.3 advisories.

Revision 5.0.0-15 Mon November 3 2014 Summer Long
Updated overview to include 5.0.2 advisories.

Revision 5.0.0-13 Tue September 30 2014 Summer Long
Updated overview to include 5.0.1 advisories. Chapters added for RHBA-2014:1324 and RHBA-2014:1325.

Revision 5.0.0-10 Wed September 3 2014 Bruce Reeler
Updated to include 5.0.0 advisories released in Sept. 2014.

Revision 5.0.0-9 Fri August 22 2014 Bruce Reeler
Added chapter for advisory RHBA-2014:1090 for the wizard-based installer in release 5.0.0.

Revision 5.0.0-8 Mon August 18 2014 Bruce Reeler
Corrected links to RHEA-2014:1003.xml and updated advisory which now includes openstack-puppet-modules BZ#1124027.

Revision 5.0.0-7 Mon August 5 2014 Bruce Reeler
Added https://rhn.redhat.com/errata/RHEA-2014-1003.html advisory.

Revision 5.0.0-6 Fri July 25 2014 Bruce Reeler
Final revision for Red Hat Enterprise Linux OpenStack Platform 5.0.0 on RHEL 7.

Revision 5.0.0-5 Tue July 8 2014 Bruce Reeler
Final revision for Red Hat Enterprise Linux OpenStack Platform 5.0.

Revision 5.0.0-2 Thu June 26 2014 Summer Long
Initial draft-Preface removed.