

# **Dell EMC Red Hat OpenStack NFV Solution**

**Release Notes**

**Version 6.0**



**Dell EMC Validated Solutions**

# Contents

---

- List of Tables..... 3
- Trademarks..... 4
- Notes, Cautions, and Warnings..... 5
  
- Chapter 1: Before You Begin..... 6
  - Related Information..... 7
  
- Chapter 2: Enhancements..... 8
  - Enhancements Listing..... 9
  
- Chapter 3: Fixes..... 10
  - Version 6.0 Fixes..... 11
  
- Chapter 4: Known Anomalies..... 12
  - Version 6.0 Known Anomalies..... 13
  
- Appendix A: Tempest Results Notes..... 21
  - Failing Tests..... 22
  
- Appendix B: References..... 25
  - To Learn More..... 26

# List of Tables

---

Table 1: Dell EMC Red Hat OpenStack Cloud Solution Fixes..... 11

Table 2: Dell EMC Red Hat OpenStack Cloud Solution Defects.....13

Table 3: Failing Tests.....22

## Trademarks

---

Copyright © 2014-2016 Dell EMC Inc. or its subsidiaries. All rights reserved.

Microsoft® and Windows® are registered trademarks of Microsoft Corporation in the United States and/or other countries.

Red Hat®, Red Hat Enterprise Linux®, and Ceph are trademarks or registered trademarks of Red Hat, Inc., registered in the U.S. and other countries. Linux® is the registered trademark of Linus Torvalds in the U.S. and other countries. Oracle® and Java® are registered trademarks of Oracle Corporation and/or its affiliates.

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

The Midokura® name and logo, as well as the MidoNet® name and logo, are registered trademarks of Midokura SARL.

## Notes, Cautions, and Warnings

---



A **Note** indicates important information that helps you make better use of your system.



A **Caution** indicates potential damage to hardware or loss of data if instructions are not followed.



A **Warning** indicates a potential for property damage, personal injury, or death.

This document is for informational purposes only and may contain typographical errors and technical inaccuracies. The content is provided as is, without express or implied warranties of any kind.

---

# Chapter 1

---

## Before You Begin

---

### Topics:

- [Related Information](#)

The Dell EMC Red Hat OpenStack Cloud Solution includes software from several parties. You should be aware of information that may impact your deployment.

## Related Information

Dell EMC recommends that, as a best practice, you become familiar with the following related information before deploying and using the Dell EMC Red Hat OpenStack Cloud Solution.



**Note:** The following list is not all-encompassing. It represents the major Solution components with which you should be familiar. Subscription access may be required.

- [Red Hat Enterprise Linux Server 7 Release Notes](#)
- [Red Hat OpenStack Platform 9 Release Notes](#)
- [Red Hat Ceph Storage 1.3.2 Release Notes](#)

---

# Chapter 2

---

## Enhancements

---

### Topics:

- [Enhancements Listing](#)

This topic describes enhancements for the Dell EMC Red Hat OpenStack NFV Solution, version 6.0.



## Enhancements Listing

Enhancements include:

- Red Hat OpenStack Platform version 9
- Added PowerEdge R730xd 24-drive variants as optional Compute and Storage nodes
- Automated deployment software
- Updated firmware for hardware components:
  - PowerEdge R630 and PowerEdge R730xd BIOS
  - iDRAC
  - Lifecycle Controller
  - RAID Controller
  - Power Supply
  - NICs



**Note:** See "*Tested BIOS and Firmware*" in the [Dell EMC Red Hat OpenStack Cloud Solution Hardware Deployment Guide](#) for specific firmware versions.

- Support for Cinder multiple concurrent backends
- Support for Dell Storage SC Series as a backend for Cinder
- Ephemeral and Cinder volume evacuation on a Red Hat Ceph Storage backend
- Improved Tempest results generation and analysis
- Support for node maintenance mode
- Support for Huge Pages
- Support for NUMA/CPU pinning optimization
- Support for Enablement of Security-Enhanced (SELinux) Protection Capability

---

# Chapter

# 3

---

## Fixes

---

### Topics:

- [Version 6.0 Fixes](#)

This topic describes issues in the Dell EMC Red Hat OpenStack Cloud Solution that have been corrected since the prior release.

## Version 6.0 Fixes

Subscriptions are required to view Red Hat Bugzilla and Dell EMC JIRA defect tracking systems.

**Table 1: Dell EMC Red Hat OpenStack Cloud Solution Fixes**

Issue	Component	Tracking Number
<p>When creating a medium-sized VM with a 1GB image, a message similar to the following appears:</p> <pre>Exception during message handling:   Timed out waiting for a reply to   message ID &lt;message_id&gt;</pre>	OpenStack, Messaging	Red Hat Bugzilla <a href="#">1341120</a>
Overcloud deployment fails if a pending Lifecycle Controller configuration job exists on a node.	OpenStack, Deployment	Dell EMC Jira <a href="#">3692</a>
<p>When a <code>swift list</code> command is executed at the CLI after containers are created in Horizon, the following error message appears:</p> <pre>/usr/lib/python2.7/site-packages/keystoneclient/service_catalog.py:196: UserWarning:   Providing attr without filter_value   to get_urls() is deprecated as   of the 1.7.0 release and may be   removed in the 2.0.0 release.   Either both should be provided or   neither should be provided. 'Providing attr without filter_value   to get_urls() is '</pre>	OpenStack, Red Hat Ceph Storage	Dell EMC JIRA <a href="#">5085</a>
After deploying the Overcloud, output from the <code>neutron net-list</code> command indicates that the tenant network is attached to VLAN 130.	OpenStack, Networking	Dell EMC JIRA <a href="#">5087</a>
Dell Storage SC Series local ephemeral volume instance fails with Nova evacuation.	OpenStack, Nova, Dell Storage SC Series	Dell EMC JIRA <a href="#">5095</a>
Live migration fails with timeout errors for Dell Storage SC Series volume instances with attached PS Series or SC Series volumes.	OpenStack, Storage	Dell EMC JIRA <a href="#">5228</a>
No ability to specify Overcloud nodes' domain names. The domain name defaults to <code>localdomain</code> .	OpenStack, Deployment	Dell EMC JIRA <a href="#">5238</a>
No ability to specify Overcloud nodes' static IP addresses.	OpenStack, Deployment	Dell EMC JIRA <a href="#">5244</a>
All shipped Overcloud images, used by Red Hat OpenStack Director to deploy OpenStack environments, have the same default root password. This could enable an attacker to gain root access to machines deployed by Red Hat OpenStack Director.	OpenStack, Red Hat OpenStack Director	Red Hat <a href="#">CVE-2016-4474</a>

---

# Chapter

# 4

---

## Known Anomalies

---

### Topics:

- [Version 6.0 Known Anomalies](#)

This section describes currently-known defects for the Dell EMC Red Hat OpenStack Cloud Solution, version 6.0. Workarounds are provided whenever possible.

## Version 6.0 Known Anomalies

Subscriptions are required to view Red Hat Bugzilla and Dell EMC JIRA defect tracking systems.

**Table 2: Dell EMC Red Hat OpenStack Cloud Solution Defects**

Issue	Component	Tracking No.	Workaround
<p>The boot order of 13G Servers, ordered with LOM daughter card, differs from that in prior server orders.</p> <p>The solution requires that the first 1G NIC be set to PXE boot. However, the NIC boot order as factory-configured is set to boot the first NIC, which is a 10G NIC.</p>	BIOS	N/A	<p>You can update the BIOS to set the boot order:</p> <ol style="list-style-type: none"> <li>1. Boot to System BIOS.</li> <li>2. Select <i>Device Settings</i>.</li> <li>3. Select <i>Integrated NIC 1 Port 1</i> (this is interface <i>em1</i>).</li> <li>4. Select <i>NIC Configuration</i>.</li> <li>5. Using the dropdown, change the boot protocol from <i>PXE</i> to <i>None</i>.</li> <li>6. In sequence, click on the <i>Back</i>, <i>Finish</i>, <i>Yes</i>, and <i>OK</i> buttons.</li> <li>7. Repeat steps 2-4 for <i>Integrated NIC 1 Port 3</i>.               <ol style="list-style-type: none"> <li>a. This time, change the boot protocol from <i>None</i> to <i>PXE</i>.</li> </ol> </li> <li>8. Reboot to System BIOS.</li> <li>9. Select <i>BIOS Boot Settings</i>.</li> <li>10. Select <i>Hard-Disk Drive Sequence</i>.</li> <li>11. Select <i>Integrated NIC 1 Port 3 Partition 1</i>.</li> <li>12. Use the + key to move the device to the top of the list.</li> <li>13. In sequence, click on the <i>OK</i>, <i>Back</i>, <i>Back</i>, <i>Finish</i>, <i>Yes</i>, and <i>OK</i> buttons.</li> <li>14. Reboot the system, which should now boot in the proper order.</li> </ol>
If local ephemeral storage is being used, <i>resize</i> operations do not currently function.	OpenStack, Compute Nodes	Red Hat Bugzilla <a href="#">975014</a>	None.
The Director Node takes a long time to shut down upon reboot, and may appear to hang.	OpenStack, Director Nodes	Red Hat Bugzilla <a href="#">1178497</a>	<p>On the Director Node, execute the following commands:</p> <ol style="list-style-type: none"> <li>1. <code>\$ sudo openstack-service \ stop</code></li> <li>2. <code>\$ sudo ip netns delete</code></li> <li>3. <code>\$ (ip netns)</code></li> <li>4. <code>\$ sudo poweroff</code></li> </ol>

Issue	Component	Tracking No.	Workaround
<p>Problem is <i>notification_driver</i> is not set in all required configuration files, and in Glance the <i>rabbit_hosts</i> may not be set when deploying the Overcloud.</p>	<p>OpenStack, Ceilometer</p>	<p>Red Hat Bugzilla <a href="#">1314732</a>, <a href="#">1316016</a></p>	<p><b>Controller Nodes</b></p> <ol style="list-style-type: none"> <li>1. Initiate <code>ssh</code> sessions to each Controller node.</li> <li>2. Execute the following commands: <pre>sudo -i grep rabbit_hosts /etc/nova/nova.conf</pre> <p>The above command will return something like "192.168.140.24,192.168.140.22".</p> </li> <li>3. Copy the value for step below, referred to as <code>RABBIT_HOSTS</code>.</li> <li>4. Open <i>glance-api.conf</i> and set the following attributes: <pre>vi /etc/glance/glance-api.conf rabbit_hosts = RABBIT_HOSTS notification_driver = messagingv2</pre> </li> <li>5. Save and close <i>glance-api</i>.</li> <li>6. Open <i>glance-registry.conf</i> and set the following attributes: <pre>vi /etc/glance/glance-registry.conf rabbit_hosts = RABBIT_HOSTS notification_driver = messagingv2</pre> </li> <li>7. Save and close <i>glance-registry.conf</i>.</li> <li>8. Open <i>cinder.conf</i> and set the following attribute: <pre>vi /etc/cinder/cinder.conf notification_driver = messagingv2</pre> </li> <li>9. Save and close <i>cinder.conf</i>.</li> <li>10. Restart the required services: <pre>systemctl restart openstack-glance-api.service openstack-glance-registry.service openstack-cinder-api.service openstack-cinder-scheduler.service</pre> </li> </ol>

Issue	Component	Tracking No.	Workaround
(Continued) Problem is <i>notification_driver</i> is not set in all required configuration files, and in Glance the <i>rabbit_hosts</i> may not be set when deploying the Overcloud.	OpenStack, Ceilometer	Red Hat Bugzilla <a href="#">1314732</a> , <a href="#">1316016</a>	<p><b>Compute Nodes</b></p> <ol style="list-style-type: none"> <li>1. Initiate <code>ssh</code> sessions to each Compute node.</li> <li>2. Edit <i>nova.conf</i> and set the following attribute value: <pre>sudo -i vi /etc/nova/nova.conf notification_driver = messagingv2</pre> </li> <li>3. Save and close <i>nova.conf</i>.</li> <li>4. Restart the required services: <pre>systemctl restart openstack-ceilometer- compute.service openstack-nova- compute.service</pre> </li> </ol> <p><b>Storage Nodes</b></p> <ol style="list-style-type: none"> <li>1. SSH into each Storage node.</li> <li>2. Edit <i>nova.conf</i> and set the following attribute value: <pre>sudo -i vi /etc/cinder/ cinder.conf notification_driver = messagingv2</pre> </li> <li>3. Save and close <i>cinder.conf</i>.</li> <li>4. Restart the required service: <pre>systemctl restart openstack-cinder- volume.service</pre> </li> </ol>
Cold migrate of host does not migrate the instances from source host when using Horizon.	OpenStack, Horizon	Red Hat Bugzilla <a href="#">1245617</a>	None.
After an Overcloud deployment, Ceilometer exceptions fill up the log files on the Controller nodes.	OpenStack, Ceilometer	Red Hat Bugzilla <a href="#">1337655</a>	Rotate the Controller node log files upon completion of an Overcloud deployment.

Issue	Component	Tracking No.	Workaround
<p>Creating a snapshot of an empty Dell Storage SC Series Cinder volume fails. To reproduce:</p> <ol style="list-style-type: none"> <li>1. Create an empty SC Series volume.</li> <li>2. Try to create a snapshot of the empty volume.</li> <li>3. A new snapshot is listed with a status of <code>error</code> instead of <code>available</code>.</li> </ol>	OpenStack, Cinder, Dell Storage SC Series	Red Hat Bugzilla <a href="#">1369978</a>	Ensure that the Dell Storage SC Series Cinder volume contains data before attempting to take a snapshot of it.
<p>CLI host evacuation fails when using the <code>--on-shared-storage</code> option to the <code>nova host-evacuate</code> command. The <a href="#">Red Hat OpenStack Platform 9 Instances and Images Guide</a> incorrectly references this deprecated option.</p>	OpenStack, Host Evacuation	Red Hat Bugzilla <a href="#">1384110</a>	Do not use that deprecated CLI option when performing host evacuation.



Issue	Component	Tracking No.	Workaround
<p>On clusters with multi-backend storage with Red Hat Ceph Storage as one of the backends, Red Hat Ceph Storage should be the default block storage Cinder backend, but it is not.</p>	<p>OpenStack, Cinder, Red Hat Ceph Storage</p>	<p>Dell EMC Jira <a href="#">6155</a></p>	<p>To set Red Hat Ceph Storage as the default block storage Cinder backend:</p> <ol style="list-style-type: none"> <li>1. Initiate <code>ssh</code> sessions to each Controller node.</li> <li>2. Edit the <code>cinder.conf</code> file, and set the following attribute: <pre>vi /etc/cinder/ \   cinder.conf   default_volume_type = \   rbd_backend</pre> </li> <li>3. Save and close <code>cinder.conf</code>.</li> <li>4. Restart the required services: <pre>sudo pcs resource \   restart openstack-   cinder-api sudo pcs resource \   restart openstack-   cinder-volume sudo pcs resource \   restart openstack-   cinder-backup sudo pcs resource \   restart openstack-   cinder-scheduler</pre> </li> <li>5. Test by creating a volume, and verifying that the volume created is in a Red Hat Ceph Storage backend: <pre>cinder create -- \   display_name volumel 1 cinder show volumel</pre> </li> </ol>
<p>The automation code expects the number of nodes that get registered in Ironic to match the number of nodes in the stamp-specific initialization file. If the number of nodes does not match, then an error is declared.</p>	<p>OpenStack, Automation Deployment</p>	<p>Dell EMC Jira <a href="#">6220</a></p>	<p>When <code>use_custom_instack_json</code> is set to <code>True</code> in the stamp-specific initialization file, the user is responsible for ensuring the Ironic node info (specifically, the iDRAC address) precisely matches the initialization file.</p>

Issue	Component	Tracking No.	Workaround
If the ARP cache on a network device (router or switch) does not update correctly, intermittent connectivity failures to floating IP (FIP) addresses may result.	OpenStack, Networking	Dell EMC Jira <a href="#">6560</a>	<ol style="list-style-type: none"> <li>1. Connect to your network device.</li> <li>2. Clear the ARP cache.</li> <li>3. Perform one or both of the following: <ol style="list-style-type: none"> <li>a. Set the ARP cache timer for the VLAN associated to the FIPs down to five (5) minutes.</li> <li>b. Increase the size of the pool for your FIPs, so that a FIP is not reused as often.</li> </ol> </li> </ol>
If Instance HA is installed in your environment, it cannot be removed in the Dell EMC Red Hat OpenStack Cloud Solution, version 6.0. This will affect removing a Compute Node if Instance HA was installed.	OpenStack, Instance HA	Dell EMC Jira <a href="#">6590</a>	See <i>Removing a Compute Node</i> in the Dell EMC Red Hat OpenStack Cloud Solution Technical Guide, <a href="#">Adding and Removing Compute and Storage Nodes</a> , for instructions to remove a Compute node if Instance HA is installed.
The sanity test creates a floating IP pool starting at .2. Many networks use .1 - .10 and 245 - 254 for network gear (i.e., switches). The sanity test should use .20 - .240.	OpenStack, Deployment	Dell EMC Jira <a href="#">6593</a>	None.
<p>When configuring the Solution Admin Host, two networks external to the cluster are required:</p> <ul style="list-style-type: none"> <li>• External Network for Management with at least 3 IP addresses; for the SAH, Director Node, and Red Hat Ceph Storage Admin Node. This is the default route.</li> <li>• Public API network used for the RESTful API in OpenStack.</li> </ul>	OpenStack, Networking	Dell EMC Jira <a href="#">6604</a>	The External Network for Tenants (Floating IP Network) does not need to be brought to the Director Node, but can be used for the External Network. You must then ensure that when creating the External Network for Tenants (Floating IP), there are no overlapping IP addresses.
If Instance HA is installed in your environment, updates and upgrades will fail for the Dell EMC Red Hat OpenStack Cloud Solution, version 6.0.	OpenStack, Instance HA	Dell EMC Jira <a href="#">6643</a>	None at this time. This issue is currently under investigation. If updates and upgrades are important to your organization, do not enable Instance HA.

Issue	Component	Tracking No.	Workaround
<p>Page 10, step 15 of the Dell EMC Red Hat OpenStack Cloud Solution <a href="#">Adding and Removing Compute and Storage Nodes Technical Guide</a>, Version 6.0, instructs the user to re-invoke the <code>deploy-overcloud.py</code> script.</p> <p>Several notes advise using the same options as were used initially.</p> <p>If the initial deployment was manual, then the user is responsible for tracking that initial invocation.</p> <p>If the initial deployment was done with automation, then the user had no contact with the automated original invocation of this script.</p> <p>No documentation of the record of that invocation formulation is currently provided.</p>	OpenStack, Documentation	Dell EMC Jira <a href="#">6649</a>	<p>The original invocation of the <code>deploy-overcloud.py</code> script, made by the deployment automation, is stored by the root user on the SAH node in the most recent log file in the <code>/auto_results</code> directory.</p> <p>The exact invocation can be extracted with the following command:</p> <pre>[root@sah ~]# grep \   deploy-overcloud.py \   \$(ls -tr /auto_results/\   deployer.log.*   tail \   -1)   tail -1   sed -e \   's/ ^.*running : //'</pre>
<p>Due to differing system node types, a cluster sometimes needs more time to stabilize the nodes. Therefore, enabling Instance HA during an automated installation may fail.</p>	OpenStack, Instance HA	Dell EMC Jira <a href="#">6684</a>	<p>Dell EMC recommends that you do not enable Instance HA during automated installations.</p> <p>Instead, proceed with the enablement of Instance HA, following the procedure in the <i>Scripted HA Installation</i> section of the <a href="#">Dell EMC Red Hat OpenStack Cloud Solution Deployment Guide</a>.</p>
<p>Currently, when the last HA router of a tenant is deleted, the tenant's HA network is not removed.</p>	OpenStack, Networking	Launchpad <a href="#">1367157</a>	None.
<p>Currently, while using the Horizon GUI, if you edit a file on your disk that is stored in a container, an error similar to the following is displayed when you attempt to upload the latest version of that file with the <i>Update object</i> control:</p> <pre>Object with the name   already exists</pre> <p>This same operation completes successfully when using the CLI.</p>	Horizon, Red Hat Ceph Storage	Launchpad <a href="#">1572635</a>	Use the CLI to perform this operation, instead of the Horizon GUI.

Issue	Component	Tracking No.	Workaround
The scripts used to enable HugePages, and NUMA/CPU pinning Optimizations sometimes do not wait for the last Compute node in the solution to reboot, when used with the action=remove option.	OpenStack Compute nodes, enable_hugepages script, enable_numa script.	N/A	None
The NUMA scripts fail with the error "openstack-nova-scheduler.service cancelled".	enable_numa script.	N/A	Ensure that the services are manually started before retrying the script.
The SR-IOV Script_pass instance does not execute if any of the links are down.	SR-IOV	N/A	Ensure that the links are up before executing the script.
If the SR-IOV Bond is assigned an existing IP address the script executes successfully.	SR-IOV Bonds	N/A	Use different IP addresses when creating the SR-IOV bonds.
The SR-IOV configuration on the instances clears after the nodes are rebooted.	SR-IOV	N/A	Rerun the SR-IOV script on the Instance.

---

# Appendix

# A

---

## Tempest Results Notes

---

### Topics:

- [Failing Tests](#)

This appendix describes known Tempest failures for the Dell EMC Red Hat OpenStack Cloud Solution, version 6.0.

## Failing Tests

**Table 3: Failing Tests**

Failing Tests
<b>Cinder</b>
These tests are invalid if backends defined in <i>cinder.conf</i> do not declare <i>vendor_name</i> or <i>storage_protocol</i> attributes. The backend will be valid and fully functional, but the test will fail and should be ignored in such cases.
<pre>tempest.api.volume.admin.test_volume_types.VolumeTypesV1Test.test_ volume_crud_with_volume_type_and_extra_specs  tempest.api.volume.admin.test_ volume_types.VolumeTypesV2Test.test_volume_crud_with_volume_type_and_ extra_specs</pre>
Cinder backup service is not deployed and configured in the Dell EMC Red Hat OpenStack Cloud Solution, thus the following backup related tests will fail.
<pre>tempest.api.volume.admin.test_volumes_ backup.VolumesBackupsV1Test.test_volume_backup_create_get_ detailed_list_restore_delete  tempest.api.volume.admin.test_volumes_ backup.VolumesBackupsV2Test.test_volume_backup_create_get_ detailed_list_restore_delete</pre>
Encrypted volumes are not configured in the Dell EMC Red Hat OpenStack Cloud Solution by default, thus the following encrypted volume related tests will fail.
<pre>tempest.scenario.test_encrypted_cinder_ volumes.TestEncryptedCinderVolumes.test_encrypted_cinder_volumes_cryptsetup  tempest.scenario.test_encrypted_cinder_ volumes.TestEncryptedCinderVolumes.test_encrypted_cinder_volumes_luks</pre>
<b>Neutron</b>
Neturon Distributed Virtual Router (DVR) is currently a technical preview only in the Dell EMC Red Hat OpenStack Cloud Solution. The following tests may fail depending on the Neutron configuration.
<pre>tempest.api.network.admin.test_routers_dvr.RoutersTestDVR.test_ centralized_router_update_to_dvr  tempest.api.network.admin.test_routers_dvr.RoutersTestDVR.test_ distributed_router_creation  tempest.api.network.test_routers.DvrRoutersTest.test_ convert_centralized_router  tempest.api.network.test_routers.DvrRoutersTest.test_ create_distributed_router</pre>

Failing Tests
Listing virtual interfaces is not supported in the Dell EMC Red Hat OpenStack Cloud Solution.
<pre>tempest.api.compute.servers.test_virtual_ interfaces.VirtualInterfacesTestJSON.test_list_virtual_interfaces</pre>
Nova
EC2 support has been deprecated in the Dell EMC Red Hat OpenStack Cloud Solution, and the nova-cert service not running on the overcloud controller nodes, therefore the following tests will fail.
<pre>tempest.api.compute.certificates.test_ certificates.CertificatesV2TestJSON.test_create_root_certificate  tempest.api.compute.certificates.test_ certificates.CertificatesV2TestJSON.test_get_root_certificate</pre>
Even with Firewall as a Service (FWaaS) disabled in your overcloud deployment you may see the teardown phase of the FWaaS fail, this failure can be safely ignored.
<pre>tearDownClass (neutron_fwaas.tests.tempest_ plugin.tests.scenario.test_fwaas.TestFWaaS)</pre>
Resize
The Resize feature is not currently supported by the solution, and is a known issue. See the following defect:
<ul style="list-style-type: none"> <li><a href="https://bugzilla.redhat.com/show_bug.cgi?id=1221776">https://bugzilla.redhat.com/show_bug.cgi?id=1221776</a></li> </ul>
<pre>tempest.api.compute.admin.test_migrations.MigrationsAdminTest.test_list_ migrations_in_flavor_resize_situation  tempest.api.compute.servers.test_ delete_server.DeleteServersTestJSON.test_delete_server_while_in_verify_ resize_state  tempest.api.compute.servers.test_disk_config.ServerDiskConfigTestJSON. test_resize_server_from_auto_to_manual  tempest.api.compute.servers.test_disk_config.ServerDiskConfigTestJSON.test_ resize_server_from_manual_to_auto  tempest.api.compute.servers.test_server_actions.ServerActionsTestJSON.test_ resize_server_confirm  tempest.api.compute.servers.test_server_actions.ServerActionsTestJSON.test_ resize_server_confirm_from_stopped  tempest.api.compute.servers.test_server_actions.ServerActionsTestJSON.test_ resize_server_revert  tempest.scenario.test_network_advanced_server_ ops.TestNetworkAdvancedServerOps.test_server_connectivity_resize</pre>
Telemetry

**Failing Tests**

Gnocchi is only available as a tech preview in the Dell EMC Red Hat OpenStack Cloud Solution, and Gnocchi related tests may fail.

```
gabbi.suitemaker.tempest.scenario.gnocchi.test_  
live_post_some_measures_to_the_metric_on_instance.test_request  
  
gabbi.suitemaker.tempest.scenario.gnocchi.test_  
live_get_instance_measures_with_poll.test_request
```



---

# Appendix

# B

---

## References

---

### Topics:

- [To Learn More](#)

Additional information can be obtained at <http://www.dell.com/nfv> or by e-mailing [nfv@dell.com](mailto:nfv@dell.com).

If you need additional services or implementation help, please contact your Dell EMC sales representative.

## To Learn More

For more information on the Dell EMC Red Hat OpenStack Cloud Solution visit <http://www.dell.com/learn/us/en/04/solutions/red-hat-openstack>.

Copyright © 2014-2016 Dell EMC Inc. or its subsidiaries. All rights reserved. Trademarks and trade names may be used in this document to refer to either the entities claiming the marks and names or their products. Specifications are correct at date of publication but are subject to availability or change without notice at any time. Dell EMC and its affiliates cannot be responsible for errors or omissions in typography or photography. Dell EMC's Terms and Conditions of Sales and Service apply and are available on request. Dell EMC service offerings do not affect consumer's statutory rights.

Dell EMC, the DELL EMC logo, the DELL EMC badge, and PowerEdge are trademarks of Dell EMC Inc.