

# **Dell EMC Ready Bundle for Red Hat OpenStack Platform**

**Release Notes**

**Version 10.0**



**Dell EMC Converged Platforms and Solutions**

# Contents

---

List of Tables.....	iii
Trademarks.....	4
Notes, Cautions, and Warnings.....	5
Chapter 1: Enhancements.....	6
Version 10.0 Enhancements.....	7
Chapter 2: Fixes.....	8
Version 10.0 Fixes.....	9
Chapter 3: Known Anomalies.....	11
Version 10.0 Known Anomalies.....	12
Appendix A: Tempest Results Notes.....	20
Failing Tests.....	21
Appendix B: References.....	23
To Learn More.....	24

# List of Tables

---

Table 1: Dell EMC Ready Bundle for Red Hat OpenStack Platform Fixes.....9

Table 2: Dell EMC Ready Bundle for Red Hat OpenStack Platform Defects..... 12

Table 3: Failing Tests..... 21

## Trademarks

---

Copyright © 2014-2017 Dell 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

---

## Enhancements

---


### Topics:

- [Version 10.0 Enhancements](#)

This chapter describes enhancements for the Dell EMC Ready Bundle for Red Hat OpenStack Platform, version 10.0.

## Version 10.0 Enhancements

Enhancements include:

- Support for the Dell EMC PowerEdge FX architecture:
  - Dell EMC PowerEdge FC630 as Compute and Controller nodes
  - Dell EMC PowerEdge FD332 with 2.5" or 1.8" drives
  - Requires a minimum of four (4) Dell EMC PowerEdge FX chassis
  - PowerEdge FX support for the Intel® Ethernet X710 Quad Port 10 GbE bNDC Converged NIC
  - PowerEdge FX support for the following Dell Networking switches, running Dell Networking OS 9:
    - S3048-ON
    - S4048-ON (used for 10 Gb connections only)
    - S6000-ON
    - S6010-ON
-  **Note:** Dell EMC PowerEdge FX-based clusters must be deployed using the Open Source Hardware Configuration Toolkit.
- Red Hat Ceph Storage version 2.0 support for both the Dell EMC PowerEdge R-Series and Dell EMC PowerEdge FX architectures
- Touch free Overcloud deployment, using Ironic support for full server management and installation, for the Dell EMC PowerEdge R-Series architecture only
- Documentation enhancements and reorganization:
  - Architecture Guides for both the Dell EMC PowerEdge R-Series and Dell EMC PowerEdge FX architectures
  - Hardware Deployment Guides for both the Dell EMC PowerEdge R-Series and Dell EMC PowerEdge FX architectures, with integrated Bills of Material
  - Combined automated and manual Software Deployment Guide for both the Dell EMC PowerEdge R-Series and Dell EMC PowerEdge FX architectures, with integrated Tempest Overcloud validation instructions

---

# Chapter

# 2

---

## Fixes

---

### Topics:

- [Version 10.0 Fixes](#)

This chapter describes issues in the Dell EMC Ready Bundle for Red Hat OpenStack Platform that have been corrected since the prior release.



## Version 10.0 Fixes

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

**Table 1: Dell EMC Ready Bundle for Red Hat OpenStack Platform Fixes**

Issue	Component	Tracking Number
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>
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>
Creating a snapshot of an empty Dell Storage SC Series Cinder volume fails.	OpenStack, Cinder, Dell Storage SC Series	Red Hat Bugzilla <a href="#">1369978</a>
After an Overcloud deployment, Ceilometer exceptions fill up the log files on the Controller nodes.	OpenStack, Ceilometer	Red Hat Bugzilla <a href="#">1337655</a>
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	Red Hat Bugzilla <a href="#">1384108</a>
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.	OpenStack, Host Evacuation	Red Hat Bugzilla <a href="#">1384110</a>
If Instance HA is installed in your environment, it cannot be removed in the Dell EMC Ready Bundle for Red Hat OpenStack Platform, version 10.0. This will affect removing a Compute Node if Instance HA was installed.	OpenStack, Instance HA	Dell EMC Jira <a href="#">CES-6590</a>
When configuring the Solution Admin Host, two networks external to the cluster are required: <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 Storage Console VM. 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="#">CES-6604</a>
Page 10, step 15 of the Dell EMC Ready Bundle for Red Hat OpenStack Platform <a href="#">Adding and Removing Compute and Storage Nodes Technical Guide</a> does not indicate how to extract the original invocation of the <code>deploy-overcloud.py</code> script.	OpenStack, Documentation	Dell EMC Jira <a href="#">CES-6649</a>

Issue	Component	Tracking Number
The default gateway on the Controller nodes is on the Provisioning network. Some environments prefer it to be on the Public API network.	OpenStack, Networking	Dell EMC Jira <a href="#">CES-6823</a>
Accessing the Red Hat Storage Console returns an Internal Server Error with SELinux enforcing.	OpenStack, Red Hat Storage Console, SELinux	Dell EMC Jira <a href="#">CES-7151</a>
Currently, when the last HA router of a tenant is deleted, the tenant's HA network is not removed.	OpenStack, Networking	Launchpad <a href="#">1367157</a>

---

# Chapter

# 3

---

## Known Anomalies

---

### Topics:

- [Version 10.0 Known Anomalies](#)

This chapter describes currently-known defects for the Dell EMC Ready Bundle for Red Hat OpenStack Platform, version 10.0. Workarounds are provided whenever possible.

## Version 10.0 Known Anomalies

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

**Table 2: Dell EMC Ready Bundle for Red Hat OpenStack Platform 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>
Cold migrate of host does not migrate the instances from source host.	OpenStack, Host Migration	Red Hat Bugzilla <a href="#">1245617</a>	Configure SSH keys on the Compute nodes as documented <a href="#">here</a> .
If local ephemeral storage is being used, <code>nova resize</code> operations do not currently function.	OpenStack, Compute Nodes	Red Hat Bugzilla <a href="#">1267598</a>	Configure SSH keys on the Compute nodes as documented <a href="#">here</a> .

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.
Currently, instances will work on the first attempt to evacuate them from a node, but will go into an <code>ERROR</code> state if you attempt to evacuate them a second time.	OpenStack, Instance HA	Red Hat Bugzilla <a href="#">1422154</a> , <a href="#">1441368</a>	None.
<p>A design flaw was found in the Dell EMC Ready Bundle for Red Hat OpenStack Platform version 10.0 Red Hat OpenStack Platform Director's use of TripleO to enable <code>libvirt</code>-based live migration. See <a href="https://access.redhat.com/solutions/3022771">https://access.redhat.com/solutions/3022771</a> for more information.</p>	OpenStack, Migration	Red Hat Bugzilla <a href="#">1428240</a>	<ol style="list-style-type: none"> <li>1. Identify the files listed in the Red Hat bulletin.</li> <li>2. Unlock them following the instructions in <i>Appendix B</i> of the <u>Dell EMC Ready Bundle for Red Hat OpenStack Platform Software Deployment Guide</u>.</li> <li>3. Update these packages.</li> <li>4. Add them back to the version lock file.</li> </ol>

Issue	Component	Tracking No.	Workaround
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.	OpenStack, Cinder, Red Hat Ceph Storage	Dell EMC Jira <a href="#">CES-6155</a>	<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>
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.	OpenStack, Automation Deployment	Dell EMC Jira <a href="#">CES-6220</a>	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.
The sanity test creates a floating IP pool starting at <code>.2</code> . Many networks use <code>.1 - .10</code> and <code>245 - 254</code> for network gear (i.e., switches). The sanity test should use <code>.20 - .240</code> .	OpenStack, Deployment	Dell EMC Jira <a href="#">CES-6593</a>	None.

Issue	Component	Tracking No.	Workaround
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.	OpenStack, Instance HA	Dell EMC Jira <a href="#">CES-6684</a>	Dell EMC recommends that you do not enable Instance HA during automated installations.  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 Ready Bundle for Red Hat OpenStack Platform Software Deployment Guide</a> .
Restarting the Red Hat OpenStack Platform Director Node fails after <i>/tmp</i> is purged. Errors similar to the following appear:  <pre>error: Failed start domain director error: Failed to open file '/tmp/floppy-director.img'</pre> <p>The XML descriptions of both VMs define floppy drive configurations that reference a source file in <i>/tmp</i>. Once <i>/tmp</i> is cleaned out, the VMs can no longer start.</p>	OpenStack, Red Hat OpenStack Platform Director, Tempest	Dell EMC Jira <a href="#">CES-6770</a>	Remove the offending devices in the XML after installation:  <pre>&lt;disk type='file'   device='floppy'&gt; &lt;driver name='qemu'   type='raw' /&gt; &lt;source file='/tmp/floppy-director.img' /&gt; &lt;target dev='fda'   bus='fdc' /&gt; &lt;address type='drive'   controller='0' bus='0'   target='0' unit='0' /&gt;</pre>
Overcloud deployment fails due to Ironic on the Director Node locking one or more nodes, preventing them from advancing from the cleaning state to manageable.	OpenStack, Deployment	Dell EMC Jira <a href="#">CES-7619</a>	<ol style="list-style-type: none"> <li>Log onto the Director Node as the <i>root</i> user.</li> <li>Edit <i>/etc/ironic/ironic.conf</i>, replacing this line:   <pre>#node_locked_retry_attempts = 3</pre> with this line:   <pre>node_locked_retry_attempts = 15</pre> </li> <li>Execute the following command:   <pre># systemctl restart \ openstack-ironic-conductor</pre> </li> <li>Redeploy the Overcloud. <ul style="list-style-type: none"> <li>For automation installations, rerun <code>python deployer.py -s</code>, passing the <code>-overcloud_only</code> and <code>-skip_rhscn_vm</code> parameters.</li> <li>For manual installations, rerun <code>deploy-overcloud.py</code>.</li> </ul> </li> </ol>

Issue	Component	Tracking No.	Workaround
Currently, Dell EMC PowerEdge FX chassis blades are not assigned individual IP addresses via DHCP. Node discovery discovers the chassis' service tag and assigns it to all the blades, which are then all assigned the chassis' IP address.	Dell EMC PowerEdge FX, Networking	Dell EMC Jira <a href="#">CES-7643</a>	Assign static IP addresses to individual blades in a Dell EMC PowerEdge FX chassis.
Automation fails to delete Ironic nodes in the Overcloud that are in the <i>clean wait</i> state during redeployment.	OpenStack, Automation Deployment	Dell EMC Jira <a href="#">CES-7650</a>	<p>To work around the issue:</p> <ol style="list-style-type: none"> <li>1. Find the failed node by executing the following command:</li> </ol> <pre>\$ ironic node-list</pre> <ol style="list-style-type: none"> <li>2. Note the GUID of the failed node in the output.</li> <li>3. Delete the node by executing the following commands:</li> </ol> <pre>\$ ironic node-set-maintenance &lt;GUID&gt; False \$ ironic node-set-provision-state &lt;GUID&gt; manage \$ ironic node-delete &lt;GUID&gt;</pre>
When access to Dell Storage PS Series or SC Series backends is interrupted during cluster deployment, the deployment fails with a syntax error message that does not indicate an access failure.	OpenStack, Dell Storage	Dell EMC Jira <a href="#">CES-7689</a>	<p>None. To complete deployment:</p> <ol style="list-style-type: none"> <li>1. Ensure that access to the Dell Storage backends is functional.</li> <li>2. Re-deploy the cluster.</li> </ol>
<p>Cinder Ceph Storage is initially lost when using Local Ephemeral Storage. The <code>enable_rdb_backend</code> variable in the <code>settings_sample.ini</code> and <code>settings_sample-fx2.ini</code> files is now used to toggle both Local Ephemeral <b>and</b> Cinder Ceph Storage on/off:</p> <ul style="list-style-type: none"> <li>• <code>enable_rdb_backend=true</code> — Cluster uses Local Ephemeral Storage and Cinder Ceph Storage</li> <li>• <code>enable_rdb_backend=false</code> — Cluster uses Local Ephemeral Storage only.</li> </ul>	OpenStack, Cinder, Red Hat Ceph Storage	Dell EMC Jira <a href="#">CES-7727</a>	<p>To enable both Local Ephemeral Storage and Cinder Ceph Storage:</p> <ol style="list-style-type: none"> <li>1. Set <code>enable_rdb_backend=true</code> to configure Cinder Ceph Storage in the <code>settings_sample.ini</code> and <code>settings_sample-fx2.ini</code> files.</li> <li>2. Complete the automated installation.</li> <li>3. Set <code>Images_type=default</code> in the <code>/etc/nova/nova.conf</code> file on each Compute node to configure Local Ephemeral Storage.</li> <li>4. Run <code>sudo systemctl restart openstack-nova-compute</code> to restart the nova-compute service on each Compute node.</li> </ol>



Issue	Component	Tracking No.	Workaround
<p>The BIOS setting <i>Non RAID Disk Mode</i> can cause RAID configuration failures when set to <i>disabled</i>. As a result, nodes cannot be provisioned.</p>	BIOS	Dell EMC Jira <a href="#">CES-7739</a>	<p>On system boot:</p> <ol style="list-style-type: none"> <li>1. Press <b>[F2]</b> to invoke the BIOS System Setup console.</li> <li>2. Navigate to <i>Device Settings &gt; Integrated Raid Controller 1 &gt; Controller Management &gt; Advanced Controller properties &gt; Non RAID Disk Mode</i>.</li> <li>3. Change the setting from <i>disabled</i> to <i>enabled</i>.</li> <li>4. Continue the system boot process.</li> </ol> <p>Nodes should now be able to be provisioned.</p>
<p>The iDRAC console or Web session drops its connection during a Dell EMC Ready Bundle for Red Hat OpenStack Platform installation. This happens when the following commands are run:</p> <ul style="list-style-type: none"> <li>• Manual installations: — <code>config_idrac.py</code></li> <li>• Automation installations — <code>config_idracs.py</code></li> </ul>	iDRAC, Deployment	Dell EMC Jira <a href="#">CES-7740</a>	<p>Manual installations:</p> <ol style="list-style-type: none"> <li>1. After <code>config_idrac.py</code> completes its run:             <ol style="list-style-type: none"> <li>a. Restart the iDRACs.</li> <li>b. Reset the iDRAC consoles or Web sessions.</li> </ol> </li> </ol> <p>Automation installations:</p> <ol style="list-style-type: none"> <li>1. Monitor the <code>/auto_results</code> log on the SAH for "<code>config_idracs.py</code>".</li> <li>2. After <code>config_idracs.py</code> completes its run:             <ol style="list-style-type: none"> <li>a. Restart the iDRACs.</li> <li>b. Reset the iDRAC consoles or Web sessions.</li> </ol> </li> </ol>

Issue	Component	Tracking No.	Workaround
<p>Deploying the Dell EMC Ready Bundle for Red Hat OpenStack Platform with build 10.0.0-145 locked bits results in the following error:</p> <pre> ---&gt; Package firewalld.noarch 0:0.4.3.2-8.1.el7_3.3 will be an update  --&gt; Processing Dependency: firewalld-filesystem = 0.4.3.2-8.1.el7_3.3 for package: firewalld-0.4.3.2-8.1.el 7_3.3.noarch  --&gt; Processing Dependency: python-firewall = 0.4.3.2-8.1.el7_3.3 for package: firewalld-0.4.3.2-8.1.el 7_3.3.noarch  ---&gt; Package python- ply.noarch 0:3.4-10.el7 will be installed  --&gt; Finished Dependency Resolution  Error: Package: firewalld-0.4.3.2-8.1.el 7_3.3.noarch (rhel-7- server-rpms) Requires: firewalld- filesystem = 0.4.3.2-8.1.el 7_3.3 Removing: firewalld- filesystem-0.4.3.2-8.el 7.noarch (@anaconda/7.3) firewalld-filesystem = 0.4.3.2-8.el7 Updated By: firewalld- filesystem-0.4.3.2-8.1.el 7_3.2.noarch (rhel-7-server-rpms) firewalld-filesystem = 0.4.3.2-8.1.el7_3.2 </pre>	OpenStack, Deployment	Dell EMC Jira <a href="#">CES-7867</a>	<p>To work around the issue:</p> <ol style="list-style-type: none"> <li>1. Log onto the SAH node as <i>root</i>.</li> <li>2. Delete the <i>/root/director_vm.vlock</i> and the <i>/rhscon_vm.vlock</i> files.</li> <li>3. In a text editor of your choice, open the <i>/root/JetStream/data/vlock_files/director_vm.vlock</i> file.</li> <li>4. Add the following line (on a single line): <pre>0:firewalld-0.4.3.2-8.1.el7_3.2.*</pre> </li> <li>5. Save the file.</li> <li>6. Redeploy the cluster.</li> </ol>

Issue	Component	Tracking No.	Workaround
<p>The following files contain an outdated package:</p> <ul style="list-style-type: none"> <li><code>/root/director_vm.vlock</code></li> <li><code>/rhscon_vm.vlock</code></li> </ul> <p>This package:</p> <ul style="list-style-type: none"> <li><code>0:python-crypto-2.6.1-1.el7ost.*</code></li> </ul> <p>Has been replaced by:</p> <ul style="list-style-type: none"> <li><code>0:python2-crypto-2.6.1-1.el7ost.*</code></li> </ul>	OpenStack, Deployment	Dell EMC Jira <a href="#">CES-7962</a>	<p>To work around the issue:</p> <ol style="list-style-type: none"> <li>1. Log onto the Director Node as root.</li> <li>2. Delete the <code>/root/director_vm.vlock</code> and the <code>/rhscon_vm.vlock</code> files.</li> <li>3. In a text editor of your choice, open the <code>/root/JetStream/data/vlock_files/director_vm.vlock</code> file.</li> <li>4. Search for the line that contains <code>0:python-crypto-2.6.1-1.el7ost.*</code>.</li> <li>5. Add the following line under it: <pre>0:python2-crypto-2.6.1-1.el7ost.*</pre> </li> <li>6. Edit the <code>/root/JetStream/data/vlock_files/rhscon_vm.vlock</code> file to add the same package.</li> <li>7. Save both files.</li> <li>8. Redeploy the cluster.</li> </ol>
<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.
In the Horizon Instance launcher, when selecting Image or Image snapshot as boot source, the <i>Create New Volume</i> option is set to <b>Yes</b> by default. This will initially create the instance via an image, but then store it on a volume.	Horizon, Ephemeral Storage	Launchpad <a href="#">1678109</a>	<p>To create a true ephemeral image using Horizon:</p> <ol style="list-style-type: none"> <li>1. Click on Launch Instance</li> <li>2. Select the following options: <ol style="list-style-type: none"> <li>a. <i>Select Boot Source</i> = <b>Image</b></li> <li>b. <i>Create New Volume</i> = <b>No</b></li> </ol> </li> </ol>

---

## Appendix

# A

---

### Tempest Results Notes

---

**Topics:**

- [Failing Tests](#)

This appendix describes known Tempest failures for the Dell EMC Ready Bundle for Red Hat OpenStack Platform, version 10.0.

## 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 Ready Bundle for Red Hat OpenStack Platform, thus the following backup related tests will fail.
<pre>tempest.api.volume.admin.test_volumes_backup.VolumesBackupsV1Test.* tempest.api.volume.admin.test_volumes_backup.VolumesBackupsV2Test.*</pre>
Encrypted volumes are not configured in the Dell EMC Ready Bundle for Red Hat OpenStack Platform 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>
Neutron Distributed Virtual Router (DVR) has production support in the Dell EMC Ready Bundle for Red Hat OpenStack Platform, albeit with some limited functionality. The following tests may fail depending on the Neutron configuration.
<pre>neutron.tests.tempest.api.admin.test_routers_ dvr.RoutersTestDVR.test_centralized_router_update_to_dvr  neutron.tests.tempest.api.test_routers.DvrRoutersTest.test_ convert_centralized_router</pre>
<b>Telemetry</b>

**Failing Tests**

Telemetry installation is now configured to use Gnocchi in Dell EMC Ready Bundle for Red Hat OpenStack Platform. Please use the Gnocchi API available on the metric endpoint to retrieve data. Therefore Ceilometer API tests may fail.

```
ceilometer.tests.tempest.api.test_telemetry_  
notification_api.TelemetryNotificationAPITest.test_check_  
glance_v1_notifications
```

```
ceilometer.tests.tempest.api.test_telemetry_  
notification_api.TelemetryNotificationAPITest.test_check_  
glance_v2_notifications
```

```
ceilometer.tests.tempest.api.test_telemetry_  
notification_api.TelemetryNotificationAPITest.test_check_  
nova_notification
```

---

# Appendix

# B

---

## References

---

### Topics:

- [To Learn More](#)

Additional information can be obtained at <http://www.dell.com/en-us/work/learn/openstack-cloud> or by e-mailing [openstack@dell.com](mailto:openstack@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 Ready Bundle for Red Hat OpenStack Platform visit <http://www.dell.com/learn/us/en/04/solutions/red-hat-openstack>.

Copyright © 2014-2017 Dell 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 Inc.