

Mirantis Technical Bulletin 2019-005

April 3, 2019

Specific traffic packets with corrupted checksums may result in a kernel panic on a compute node

ISSUE

The issue occurs when the OpenContrail vRouter kernel module performs segmentation of specific traffic packets with corrupted lengths/checksums. The compute node stops responding with a kernel panic error and only a reboot restores it. Therefore, Mirantis highly recommends completing the steps described in this bulletin to prevent the compute nodes deadlock on your MCP cluster.

AFFECTS

The issue affects any OpenStack or Kubernetes compute node deployed on an MCP cluster prior to the MCP Release Version Q4'18 with the OpenContrail version 3.2 or 4.0 regardless of the OpenStack or Kubernetes version.

NOTE:

Only the vRouter kernel-based deployments are affected. The DPDK-based deployments are not affected.

SECURITY IMPACT

None

STEPS TO PREVENT

The following procedure describes the steps for the OpenStack-based compute nodes. To update the Kubernetes-based compute nodes, implement similar steps to isolate workloads and update packages.

PREREQUISITES

Before you proceed with the upgrade, complete the following prerequisite steps:

NOTE:

- Only nondestructive actions are applied during this stage.
- This procedure assumes that the hostname of a target node is `cmp00X`.
- This procedure assumes Keystone version 2. If your deployment is running Keystone version 3, change `/root/keystonerc` to `/root/keystonercv3` respectively.

1. Log in to the Salt Master node.
2. Verify that the affected version of the `contrail-vrouter-dkms` package is running.
For example:

```
salt 'cmp00X' cmd.run 'dpkg -l |grep contrail | egrep
"vrouter-dkms|vrouter-utils|vrouter-netns"; ls -l $(modinfo -n
vrouter)'
```

Example of system response:

```
contrail-vrouter-dkms_3.2~20190118065044-0
```

3. Obtain the FQDN used by OpenStack services that is `cmp00X.fqdn`:

```
salt 'ct101*' cmd.run './root/keystonerc; openstack host list | grep
cmp00X'
```

4. Verify that the OpenStack compute-related services are healthy:

```
salt 'cmp00X' cmd.run 'service nova-compute status; contrail-status'
salt 'ct101*' cmd.run './root/keystonerc; openstack compute service
list --host cmp00X.fqdn --long'
```

5. Verify that utilization of the disk space is less than 80%:

```
salt 'cmp00X' cmd.run 'df -h'
```

6. Verify that the expected amount of VMs is running:

```
salt 'ctl01*' cmd.run '. /root/keystonerc; openstack server list --host  
cmp00X.fqdn --all-projects'  
  
salt 'cmp00X' cmd.run 'virsh list | grep instance | wc -l'
```

7. Record the details of the compute node current state for the post-upgrade comparison:

```
salt 'cmp00X' cmd.run 'dpkg -l |grep contrail >>  
/root/pre-update-$(date +%F_%R)'  
  
salt 'cmp00X' cmd.run 'ls -l $(modinfo -n vrouter) >>  
/root/pre-update-$(date +%F_%R)'  
  
salt 'ctl01*' cmd.run '. /root/keystonerc; openstack server list --host  
cmp00X.fqdn --all-projects >> /root/pre-update-$(hostname -s)-$(date  
+%F_%R)'
```

8. Depending on the state of VMs running on a target compute node, choose from the following options:

WARNING: The options a and b assume actions on a compute node that can disrupt workload availability. The option c required if a compute node is already added to the OpenStack Nova cluster.

- a. If a compute node has running VMs to be live-migrated to different hosts, perform a smooth resource migration to minimize the workload downtime. Once done, disable the `nova-compute` service as described in the option c below.
- b. If a compute node has running VMs not to be live-migrated to different hosts, stop the VMs. Once done, disable the `nova-compute` service as described in the option c below.
- c. If a compute node does not have running VMs:
 - i. Log in to the Salt Master node.
 - ii. Disable the `nova-compute` service. For example:

```
salt 'cmp00X' cmd.run '. /root/keystonerc; openstack
compute service set --disable --disable-reason "Change
request <change_request_number>, upgrade of vrouter
cmp00x.fqdn compute"'
```

iii. Verify that the nova-compute is disabled:

```
salt 'ctl01*' cmd.run '. /root/keystonerc; openstack
compute service list --host cmp00X.fqdn --long'
```

Now, proceed with upgrading the packages.

UPGRADE THE PACKAGES

1. Log in to the Salt Master node.
2. Download the .tar archive from [Artifactory](#) depending on the MCP Release Version and the OpenContrail version installed on the target compute nodes:

```
salt 'cmp00X' cmd.run 'wget
http://artifactory.mirantis.com/fixes/fixes-<opencontrail_version>-<mcp_r
elease_version>.tar.gz /root/'
```

For example, if your cluster is based on OpenContrail 3.2 and MCP Release Version Q2'18 with the Build ID 2018.8.0, the link for the archive is as follows:

```
salt 'cmp00X' cmd.run 'wget
http://artifactory.mirantis.com/fixes/fixes-3.2-q2.tar.gz /root/'
```

3. Disable the alerts for the target compute node that will require packages update.
4. Log in to the target compute node.
5. Change to the /root/ directory.
6. Extract the content of the archive:

```
# gzip -d fixes-<opencontrail_version>-<mcp_release_version>.tar.gz
# tar -xvf fixes-<opencontrail_version>-<mcp_release_version>.tar.gz
# cd fixes-<opencontrail_version>-<mcp_release_version>.tar.gz
```

Each archive contains three .deb packages and one .pdf file with a list of instructions.

The .deb packages are as follows:

- contrail-vrouter-dkms: DKMS package to compile and install vRouter kernel module with the required fix
- contrail-vrouter-utils: dependency package
- python-opencontrail-vrouter-netns: dependency package

7. Install the fix packages from the archive:

```
# dpkg -i python-opencontrail-vrouter-netns*.deb
# dpkg -i contrail-vrouter-utils*.deb
# dpkg -i contrail-vrouter-dkms*.deb
```

8. Reboot the target compute node:

```
# reboot
```

9. Log in to the Salt Master node.

10. Verify that the target compute node is up and running:

```
salt 'cmp00X' test.ping
```

11. Verify that the packages have been updated and have the new timestamp:

```
salt 'cmp00X' cmd.run 'dpkg -l |grep contrail | egrep
"vrouter-dkms|vrouter-utils|vrouter-netns"; ls -l $(modinfo -n vrouter) '
```

12. Verify that the Nova and OpenContrail services are healthy:

```
salt 'cmp00X' cmd.run 'service nova-compute status; contrail-status'

salt 'ctl01*' cmd.run './root/keystonerc; openstack compute service
list --host cmp00X.fqdn --long'
```

13. Repeat the steps 3-12 on the remaining compute nodes.

Now, proceed to the post-upgrade activities.

POST-UPGRADE ACTIVITIES

Once you complete the upgrade procedure, verify that the compute nodes agents and services are up and running and add them back to scheduling.

1. Log in to the Salt Master node.
2. Verify that the following MCP cluster operations on the target compute node are successful:
 - a. Provisioning of VM
 - b. Access from VM to external network and metadata
 - c. FIP association and reachability
 - d. Volume attaching and writing to it
3. Enable the `nova-compute` service:

```
salt 'ctl01*' cmd.run '. /root/keystonerc; openstack compute service set --enable cmp00X nova-compute'
```

4. Verify that the `nova-compute` is enabled:

```
salt 'ctl01*' cmd.run '. /root/keystonerc; openstack compute service list --host cmp00X.fqdn --long'
```

5. Enable the alerts for the target compute node.
6. Stop and start the `nova-compute` service:

```
salt 'cmp00X' cmd.run 'service nova-compute stop; sleep 15; service nova-compute start'
```

7. In the Prometheus web UI, verify that the `NovaServiceDown` alert has fired.
8. Remove virtual resources created for testing purposes.
9. Remove the temporary files. For example:

```
salt 'cmp00X' cmd.run 'rm /root/pre-update-$(date +%F)*'  
  
salt 'cmp00X' cmd.run 'rm -r /root/fixes-3.2-q3*'
```

10. Repeat the steps 2-9 on the remaining compute nodes.

WARNING: Since the APT repository may not yet contain the fixed packages described in this article, during further repository updates if required, you may be asked to downgrade the OpenContrail packages updated during this procedure. In this case, Mirantis recommends keeping the existing version of the OpenContrail packages not to override the fix described in this article.

STEPS TO REVERT THE PATCH

Not recommended

REFERENCES

[1] <https://review.opencontrail.org/#/c/47462/>

[2] <http://artifactory.mirantis.com/fixes/>