Troubleshooting Information

This chapter describes the common issues associated with SMI components and its resolution.

- Debugging OpenStack CSI Plugin for Kubernetes, on page 1
- ErrImagePull for System Kubernetes Images, on page 2
- Kubernetes Volumes Freezes on Node Shutdown, on page 3
- Log Forwarding - Missing or No Log, on page 4
- Show TAC Missing or No Information, on page 4
- Recovering Postgres Authentication, on page 5
- Network Connection Reset, on page 6
- Vulnerability Scan Reports a False Error Message During the Directory Listing, on page 7

Debugging OpenStack CSI Plugin for Kubernetes

This section describes how to debug the OpenStack CSI plugin for Kubernetes.

To debug the OpenStack CSI plugin for Kubernetes, validate the following:

- Connection to the OpenStack metadata service.
- OpenStack credentials.
- Cinder CSI label on the nodes.

Validating the Connection to the OpenStack Metadata Service

1. Validate the connection to the OpenStack metadata service using the following curl command.

```
ubuntu@cn2smi-controlplane1:~$ curl http://169.254.169.254/openstack
2012-08-10
2013-04-04
2013-10-17
2015-10-15
2016-06-30
2016-10-06
2017-02-22
latest
```

2. If the curl command fails to display any output, verify whether a proper route is there for 169.254.169.254.

```
Destination     Gateway         Genmask         Flags Metric Ref    Use Iface
0.0.0.0         172.16.181.1    0.0.0.0         UG    100    0        0 ens4
```

```
10.192.1.0      0.0.0.0         255.255.255.0   U     0      0        0 bond0.2400
169.254.169.254 172.16.180.1    255.255.255.255 UGH   100    0        0 ens3
169.254.169.254 172.16.181.1    255.255.255.255 UGH   100    0        0 ens4
```

Validating OpenStack Credentials

1. From the control plane node, debug the loaded OpenStack credentials using the following command.

   ```
   kubectl get secrets -n kube-system cloud-config -o json  | grep '"cloud.conf"'  | cut
   -f 4 -d \" | base64 --decode
   ```

   A sample output is shown below:

   ```
   [Global]
   username=core
   password="Csco@123"
   auth-url=http://10.81.68.210:5000//v3
   tenant-id=eb3c1fda05ab48faad1f94e015b84f63
   domain-id=default

   [BlockStorage]
   bs-version=auto
   ```

   Verify if the values specified in the following fields are accurate:

   - Username
   - Password
   - Auth URL
   - Tenant ID
   - Domain ID

   The OpenStack CSI plugin will fail to connect, if the specified values are inaccurate.

Validating the Cinder CSI Label on the Nodes

The following label must be configured on the nodes with openstack CSI controller plugin and node-plugin pods.

```
"topology.cinder.csi.openstack.org/zone":"nova"
```

For verifying the labels on the primary control plane node, use the following command:

```
kubectl get nodes --show-labels
```

ErrImagePull for System Kubernetes Images

This section describes how to resolve the ErrImagePull status in system Kuberenetes images.

Problem

The Kubernetes kubelet cmdlet attempts to purge unused disk images and pods from the system, when the image file system (or file system) is under pressure (disk space reaches lesser than 15%). This results in the system Kubernetes images being purged from the system and the pods end up in an "ErrImagePull" status.

A sample "ErrImagePull" status appearing for the kube-system pods are shown below:

```
kube-system     calico-node-lnz6g                                                 0/1
Init:ErrImagePull   0          41m    40.40.40.154    cn-svi-tb10-service2
kube-system     kube-proxy-zkk46                                                  0/1
ErrImagePull        0          51m    40.40.40.154    cn-svi-tb10-service2
smi-vips        keepalived-km4wx                                                  0/2
ErrImagePull        0          40m    40.40.40.154    cn-svi-tb10-service2
```

Resolution

You can restore the local images to resolve the "ErrImagePull" status in system Kubernetes images.

To restore the local images, execute the following command.

```
sudo docker load < /var/tmp/k8s-offline.tgz
```

Kubernetes Volumes Freezes on Node Shutdown

In the event of an unclean node shutdown, the detach and reattach logic of persistent volumes in Kubernetes are prone for issues. For more information on the issues, see https://github.com/kubernetes/kubernetes/issues/65392. This section describes the root cause of these issues and how to resolve these issues in the production system.

Problem

The following are some of the scenarios where the Kubernetes volumes freezes after an unclean node shutdown:

- When you shutdown a node that has a non-local persistent volumes (Cinder or vSphere) attached to it without using the cordon or drain approach recommended by Kubernetes.

```
ubuntu@cn2smi-controlplane1:~$ kubectl get nodes
NAME                    STATUS     ROLES           AGE   VERSION
cn2smi-controlplane1    Ready      control-plane   21h   v1.15.3
cn2smi-controlplane2    Ready      control-plane   21h   v1.15.3
cn2smi-controlplane3    Ready      control-plane   21h   v1.15.3
cn2smi-oam1             NotReady   <none>          14h   v1.15.3
```

- When the pods are in "creation" or "initialization" state waiting for the volumes to attach. The pods freeze approximately after thirty seconds.

```
ubuntu@cn2smi-controlplane1:~$ kubectl get pods -w -A -o wide| grep oam | grep bulk
cee-global      bulk-stats-68dc684d57-hqtdj                                       3/3
    Terminating         0          4m38s   192.200.7.68    cn2smi-oam1
<none>           <none>
cee-global      bulk-stats-68dc684d57-sdjgl                                       0/3
    ContainerCreating   0          15s     <none>          cn2smi-oam2
<none>           <none>
```

- The pods in the "ContainerCreating" state displays a "Multi-Attach error" event.

```
Events:
  Type     Reason              Age   From                     Message
  ----     ------              ----  ----                     -------
  Normal   Scheduled           76s   default-scheduler        Successfully assigned
cee-global/bulk-stats-68dc684d57-sdjgl to cn2smi-oam2
  Warning  FailedAttachVolume  76s   attachdetach-controller  Multi-Attach error for
volume "pvc-b01a4434-190f-4eec-b289-41ae990b0025" Volume is already used by pod(s)
bulk-stats-68dc684d57-hqtdj
ubuntu
```

Resolution

You can delete the "Terminating" node to resolve this issue. To delete the "Terminating" node, run the following command:

```
kubectl delete pod -n <namespace> <pod-name> --force --grace-period=0
```

Note You must use the `grace-period=0` and `force` options for deleting the node.

The pod reschedules approximately after seven minutes .

Log Forwarding - Missing or No Log

This section describes how to resolve the "missing logs" issue when Log Forwarding is enabled.

Problem

Log forwarding is enabled, but logs are missing or not seen on the external collector.

Resolution

1. Ping the Kubernetes node from the collection host to verify the network connectivity between each Kubernetes nodes and the collection server (where Fluentd or Fluentbit endpoint is hosted).

   ```
   ping [k8s-node-ip]
   ```

2. Verify whether the client machine's system clock is in synchronization with the Kubernetes node's clock. Any mismatch in the time will result in incorrect query in the front-end visualization tool (Kibana or Grafana), which uses the host clock as query parameter.
3. When Logs Forwarding is enabled, the Logs Forwarder dumps the entire JournalD entries - from the beginning - to the external collector. Also, the number of nodes on the deployment increases the amount of data available for processing. It may take a while for pushing out and processing the log entries before the frontend tools visualize it.

Show TAC Missing or No Information

This section describes how to resolve the failure of logs or Cores with the Show TAC feature.

Problem

When the Show TAC feature is executed, the logs or cores fail to show up in the Apache server.

Resolution

1. Verify the status of the Show TAC process.

   ```
   tac-debug-pkg status
   ```

2. Verify the parameters and syntax.

- The TAC package creation time must be in *yyyy-mm-dd_hh:mm:ss* format.
- The parameters specified for filtering must match the microservices labeling.

3. In the target node, run the following command to verify if the targeted JournalD entries or Core files are within the specified time range.

```
journalctl
```

Note The first line of the output specifies the starting time of the current logs repository. Also, JournalID starts aging out old entries (log rotation) when the log size reaches the specified limit.

Similarly, for the Core files, run the following command to display the available Core files and its corresponding time stamp.

```
ls -al /var/lib/systemd/coredump
```

4. Verify if the JournalID is rate-limiting the entries in the target node. Filter for statements like the one shown in the following example:

```
journalctl | grep Suppressed
Suppressed <number> messages from /system.slice/...
```

Note JournalID is configured to allow 10000 messages per 10 seconds period.

Recovering Postgres Authentication

This section describes how to resolve recovering Postgres authentication issue.

Problem

The Postgres database is password protected. The passwords are stored as Kubernetes secrets. Also, the nodes in the Kubernetes cluster store the persisted data in the following location:

`/data/<cee-namespace>/data-postgres-x`.

You may encounter an authentication error while accessing Postgres database. A sample authentication error message is shown below:

```
2019-11-12 21:32:18.856 UTC [239] FATAL:  password authentication failed for user
"replica_user"
2019-11-12 21:32:18.856 UTC [239] DETAIL:  Password does not match for user "replica_user".
```

The following are the probable causes for the authentication error:

- The Kubernetes secrets are deleted.
- The `/data` folder is corrupted.

Resolution

When the Kubernetes secrets are deleted

When you remove the CEE Ops Center deployment from the SMI Cluster Manager, the CEE namespace is also deleted. This action in turn deletes the secrets from the Kubernetes, which restricts access to the old database. Even when you have the CEE Ops Center deployed again, you will only have a new secret created and cannot restore the deleted secret.

To resolve this issue:

- Delete the persisted database manually on the OAM nodes located here:

  `/data/<cee-namespace>/data-postgres-x`). or

- Upgrade the deployed software of the network function. To upgrade, perform the following:

  1. Edit the default helm repository the software uses through NETCONF/RESTCONF or CLI.

     `helm repository base-repos url` *`<url from the SMI deployer>`*

     Example:

     ```
     helm repository base-repos url https://charts.10.192.1.101.nip.io/cee-2019-09-30
     ```

Note The **show system status** command indicates the percentage of software upgrade process completed for the network function.

When the /data is corrupted

When the `/data` is corrupted, you must remove the `/data` to resolve the authentication error.

Network Connection Reset

This section describes how to resolve the network connection reset issue in Geographical Redundancy (GR) deployment.

Problem

During GR deployments, the SMI Cluster Manager is deployed on one or more remote sites. When the SMI Cluster Manager is deployed in at least two remote sites and if those deployed SMI Cluster Managers use a common *Vrrp-router ID* over a common VLAN/network with different IPs for running the traffic on SBI interface, the network connection resets with the following error:

Network error connection resets by peer

Resolution

The network connection resets because the Virtual IP (VIP) router ID is not unique across the two SMI Cluster Managers. You must ensure that the VIP router ID is unique while configuring the K8s clusters. For more details on configuring the cluster, see *Clusters Configuration* section in *UCC SMI Operations Guide* Guide.

Vulnerability Scan Reports a False Error Message During the Directory Listing

This section describes how to handle the false error that gets displayed during the vulnerability scan.

Problem

The Show TAC debug package interface requires the Director Listing feature configured. For debugging issues, the operator collects the TAC logs from the directory listing. However, the vulnerability scan incorrectly reports a security error during the listing.

Resolution

The error is invalid since the show-tac-url is secured through password protection, indicating no vulnerability threat associated with this operation.