Substituição de PCRF do servidor controlador UCS C240 M4

Contents

Introduction

Este documento descreve as etapas necessárias para substituir um servidor controlador com falha em uma configuração Ultra-M que hospeda as Funções de Rede Virtual (VNFs) do CPS.

Prerequisites

Backup

Em caso de recuperação, a Cisco recomenda fazer um backup do banco de dados OSPD (DB) com o uso destas etapas:

```
[root@director ~]# mysqldump --opt --all-databases > /root/undercloud-all-databases.sql
[root@director ~]# tar --xattrs -czf undercloud-backup-`date +%F`.tar.gz /root/undercloud-all-
databases.sql
/etc/my.cnf.d/server.cnf /var/lib/glance/images /srv/node /home/stack
tar: Removing leading `/' from member names
```

Verificação de status preliminar

Éimportante verificar o status atual do ambiente e dos serviços do OpenStack e garantir que ele esteja saudável antes de prosseguir com o procedimento de substituição. Ele pode ajudar a evitar complicações no momento do processo de substituição do controlador.

Etapa 1. Verifique o status do OpenStack e a lista de nós:

```
[stack@director ~]$ source stackrc
[stack@director ~]$ openstack stack list --nested
[stack@director ~]$ ironic node-list
[stack@director ~]$ nova list
```

Etapa 2. Verifique o status do Pacemaker nos controladores.

Faça login em um dos controladores ativos e verifique o status do pacemaker. Todos os serviços devem estar em execução nos controladores disponíveis e parados no controlador com falha.

```
[stack@pod1-controller-0 ~]# pcs status

<snip>
Online: [ pod1-controller-0 pod1-controller-1 ]
OFFLINE: [ pod1-controller-2 ]
Full list of resources:
ip-11.120.0.109 (ocf::heartbeat:IPaddr2): Started pod1-controller-0
ip-172.25.22.109 (ocf::heartbeat:IPaddr2): Started pod1-controller-1
ip-192.200.0.107 (ocf::heartbeat:IPaddr2): Started pod1-controller-0

Clone Set: haproxy-clone [haproxy]
Started: [ pod1-controller-0 pod1-controller-1 ]
Stopped: [ pod1-controller-2 ]

Master/Slave Set: galera-master [galera]
Masters: [ pod1-controller-0 pod1-controller-1 ]
Stopped: [ pod1-controller-2 ]
ip-11.120.0.110 (ocf::heartbeat:IPaddr2): Started pod1-controller-0
ip-11.119.0.110 (ocf::heartbeat:IPaddr2): Started pod1-controller-1

Clone Set: rabbitmq-clone [rabbitmq]
Started: [ pod1-controller-0 pod1-controller-1 ]
Stopped: [ pod1-controller-2 ]

Master/Slave Set: redis-master [redis]
Masters: [ pod1-controller-0 ]
Slaves: [ pod1-controller-1 ]
Stopped: [ pod1-controller-2 ]

ip-11.118.0.104 (ocf::heartbeat:IPaddr2): Started pod1-controller-1
openstack-cinder-volume (systemd:openstack-cinder-volume): Started pod1-controller-0

my-ipmilan-for-controller-6 (stonith:fence_ipmilan): Started pod1-controller-1
my-ipmilan-for-controller-4 (stonith:fence_ipmilan): Started pod1-controller-0
my-ipmilan-for-controller-7 (stonith:fence_ipmilan): Started pod1-controller-0

Failed Actions:
Daemon Status:

corosync: active/enabled
pacemaker: active/enabled
pcsd: active/enabled
```

Neste exemplo, Controller-2 está offline. Por conseguinte, será substituído. O controlador 0 e o controlador 1 estão operacionais e executando os serviços de cluster.

Etapa 3. Verifique o status de MariaDB nos controladores ativos.

```
[stack@director] nova list | grep control
| 4361358a-922f-49b5-89d4-247a50722f6d | pod1-controller-0 | ACTIVE | - | Running |
ctlplane=192.200.0.102 |
| d0f57f27-93a8-414f-b4d8-957de0d785fc | pod1-controller-1 | ACTIVE | - | Running |
ctlplane=192.200.0.110 |

[stack@director ~]$ for i in 192.200.0.102 192.200.0.110 ; do echo "*** $i ***" ; ssh heat-
admin@$i "sudo mysql --exec=\"SHOW STATUS LIKE 'wsrep_local_state_comment'\" ; sudo mysql --
exec=\"SHOW STATUS LIKE 'wsrep_cluster_size'\""; done
*** 192.200.0.152 ***
Variable_name     Value
wsrep_local_state_comment  Synced
Variable_name     Value
wsrep_cluster_size         2
*** 192.200.0.154 ***
Variable_name     Value
wsrep_local_state_comment  Synced
Variable_name     Value
wsrep_cluster_size         2
```

Verifique se essas linhas estão presentes para cada controlador ativo:

wsrep_local_state_comment: Sincronizado

wsrep_cluster_size: 2

Etapa 4. Verifique o status do Rabbitmq nos controladores ativos. O controlador com falha não deve aparecer na lista dos nós que são executados.

```
[heat-admin@pod1-controller-0 ~] sudo rabbitmqctl cluster_status
Cluster status of node 'rabbit@pod1-controller-0' ...
[{nodes,[{disc,['rabbit@pod1-controller-0','rabbit@pod1-controller-1',
               'rabbit@pod1-controller-2']}]},
{running_nodes,['rabbit@pod1-controller-1',
                'rabbit@pod1-controller-0']},
 {cluster_name,<<"rabbit@pod1-controller-2.localdomain">>},
 {partitions,[]},
 {alarms,[{'rabbit@pod1-controller-1',[]},
         {'rabbit@pod1-controller-0',[]}]}]


[heat-admin@pod1-controller-1 ~] sudo rabbitmqctl cluster_status
Cluster status of node 'rabbit@pod1-controller-1' ...
[{nodes,[{disc,['rabbit@pod1-controller-0','rabbit@pod1-controller-1',
               'rabbit@pod1-controller-2']}]},
 {running_nodes,['rabbit@pod1-controller-0',
                 'rabbit@pod1-controller-1']},
 {cluster_name,<<"rabbit@pod1-controller-2.localdomain">>},
 {partitions,[]},
 {alarms,[{'rabbit@pod1-controller-0',[]},
         {'rabbit@pod1-controller-1',[]}]}]
```

Etapa 5. Verifique se todos os serviços em nuvem estão no status carregado, ativo e em execução no nó OSP-D.

```
[stack@director ~]$ systemctl list-units "openstack*" "neutron*" "openvswitch*"

UNIT                                    LOAD   ACTIVE SUB     DESCRIPTION
```

```
neutron-dhcp-agent.service                 loaded active running OpenStack Neutron DHCP Agent
neutron-openvswitch-agent.service          loaded active running OpenStack Neutron Open vSwitch
Agent
neutron-ovs-cleanup.service                loaded active exited  OpenStack Neutron Open vSwitch
Cleanup Utility
neutron-server.service                     loaded active running OpenStack Neutron Server
openstack-aodh-evaluator.service           loaded active running OpenStack Alarm evaluator
service
openstack-aodh-listener.service            loaded active running OpenStack Alarm listener
service
openstack-aodh-notifier.service            loaded active running OpenStack Alarm notifier
service
openstack-ceilometer-central.service       loaded active running OpenStack ceilometer central
agent
openstack-ceilometer-collector.service     loaded active running OpenStack ceilometer collection
service
openstack-ceilometer-notification.service  loaded active running OpenStack ceilometer
notification agent
openstack-glance-api.service               loaded active running OpenStack Image Service (code-
named Glance) API server
openstack-glance-registry.service          loaded active running OpenStack Image Service (code-
named Glance) Registry server
openstack-heat-api-cfn.service             loaded active running Openstack Heat CFN-compatible
API Service
openstack-heat-api.service                 loaded active running OpenStack Heat API Service
openstack-heat-engine.service              loaded active running Openstack Heat Engine Service
openstack-ironic-api.service               loaded active running OpenStack Ironic API service
openstack-ironic-conductor.service         loaded active running OpenStack Ironic Conductor
service
openstack-ironic-inspector-dnsmasq.service loaded active running PXE boot dnsmasq service for
Ironic Inspector
openstack-ironic-inspector.service         loaded active running Hardware introspection service
for OpenStack Ironic
openstack-mistral-api.service              loaded active running Mistral API Server
openstack-mistral-engine.service           loaded active running Mistral Engine Server
openstack-mistral-executor.service         loaded active running Mistral Executor Server
openstack-nova-api.service                 loaded active running OpenStack Nova API Server
openstack-nova-cert.service                loaded active running OpenStack Nova Cert Server
openstack-nova-compute.service             loaded active running OpenStack Nova Compute Server
openstack-nova-conductor.service           loaded active running OpenStack Nova Conductor Server
openstack-nova-scheduler.service           loaded active running OpenStack Nova Scheduler Server
openstack-swift-account-reaper.service     loaded active running OpenStack Object Storage
(swift) - Account Reaper
openstack-swift-account.service            loaded active running OpenStack Object Storage
(swift) - Account Server
openstack-swift-container-updater.service  loaded active running OpenStack Object Storage
(swift) - Container Updater
openstack-swift-container.service          loaded active running OpenStack Object Storage
(swift) - Container Server
openstack-swift-object-updater.service     loaded active running OpenStack Object Storage
(swift) - Object Updater
openstack-swift-object.service             loaded active running OpenStack Object Storage
(swift) - Object Server
openstack-swift-proxy.service              loaded active running OpenStack Object Storage
(swift) - Proxy Server
openstack-zaqar.service                    loaded active running OpenStack Message Queuing
Service (code-named Zaqar) Server
openstack-zaqar@1.service                  loaded active running OpenStack Message Queuing
Service (code-named Zaqar) Server Instance 1
openvswitch.service                        loaded active exited  Open vSwitch

LOAD   = Reflects whether the unit definition was properly loaded.
ACTIVE = The high-level unit activation state, i.e. generalization of SUB.
SUB    = The low-level unit activation state, values depend on unit type.
```

```
37 loaded units listed. Pass --all to see loaded but inactive units, too.
To show all installed unit files use 'systemctl list-unit-files'.
```

Desative a cerca no cluster do controlador

```
[root@pod1-controller-0 ~]# sudo pcs property set stonith-enabled=false
[root@pod1-controller-0 ~]# pcs property show

Cluster Properties:
 cluster-infrastructure: corosync
 cluster-name: tripleo_cluster
 dc-version: 1.1.15-11.el7_3.4-e174ec8
 have-watchdog: false
 last-lrm-refresh: 1510809585
 maintenance-mode: false
 redis_REPL_INFO: pod1-controller-0
stonith-enabled: false

Node Attributes:
 pod1-controller-0: rmq-node-attr-last-known-rabbitmq=rabbit@pod1-controller-0
 pod1-controller-1: rmq-node-attr-last-known-rabbitmq=rabbit@pod1-controller-1
 pod1-controller-2: rmq-node-attr-last-known-rabbitmq=rabbit@pod1-controller-2
```

Instale o novo nó do controlador

Etapa 1. As etapas para instalar um novo servidor UCS C240 M4 e as etapas de configuração inicial podem ser consultadas a partir do Guia de instalação e serviços do servidor Cisco UCS C240 M4

Etapa 2. Faça login no servidor usando o CIMC IP.

Etapa 3.Execute o upgrade do BIOS se o firmware não estiver de acordo com a versão recomendada usada anteriormente. As etapas para a atualização do BIOS são fornecidas aqui:

Guia de atualização do BIOS de servidor com montagem em rack Cisco UCS C-Series

Etapa 4.Verifique o status das unidades físicas. Ele deve ser **Não configurado como Bom**. Navegue até **Storage > Cisco 12G SAS Modular Raid Controller (SLOT-HBA) > Physical Drive Info (Armazenamento > Controlador RAID modular SAS Cisco 12G (SLOT-HBA) > Physical Drive Info (Informações da unidade física).**

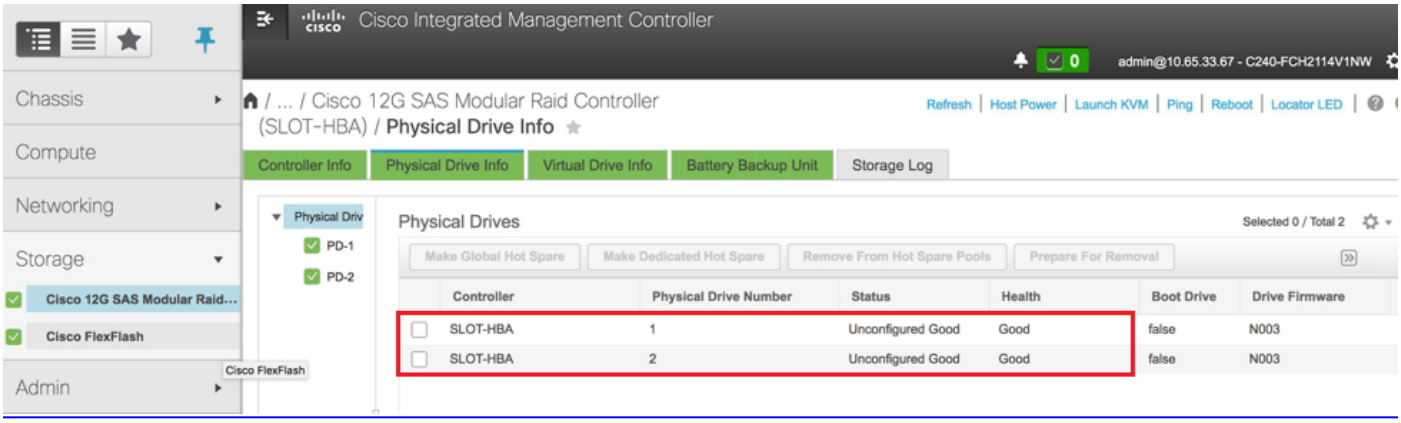

Etapa 5. Para criar uma unidade virtual a partir das unidades físicas com RAID Nível 1: navegue para **Storage > Cisco 12G SAS Modular Raid Controller (SLOT-HBA) > Controller Info > Create Virtual Drive from Unused Physical Drives**, como mostrado na imagem.

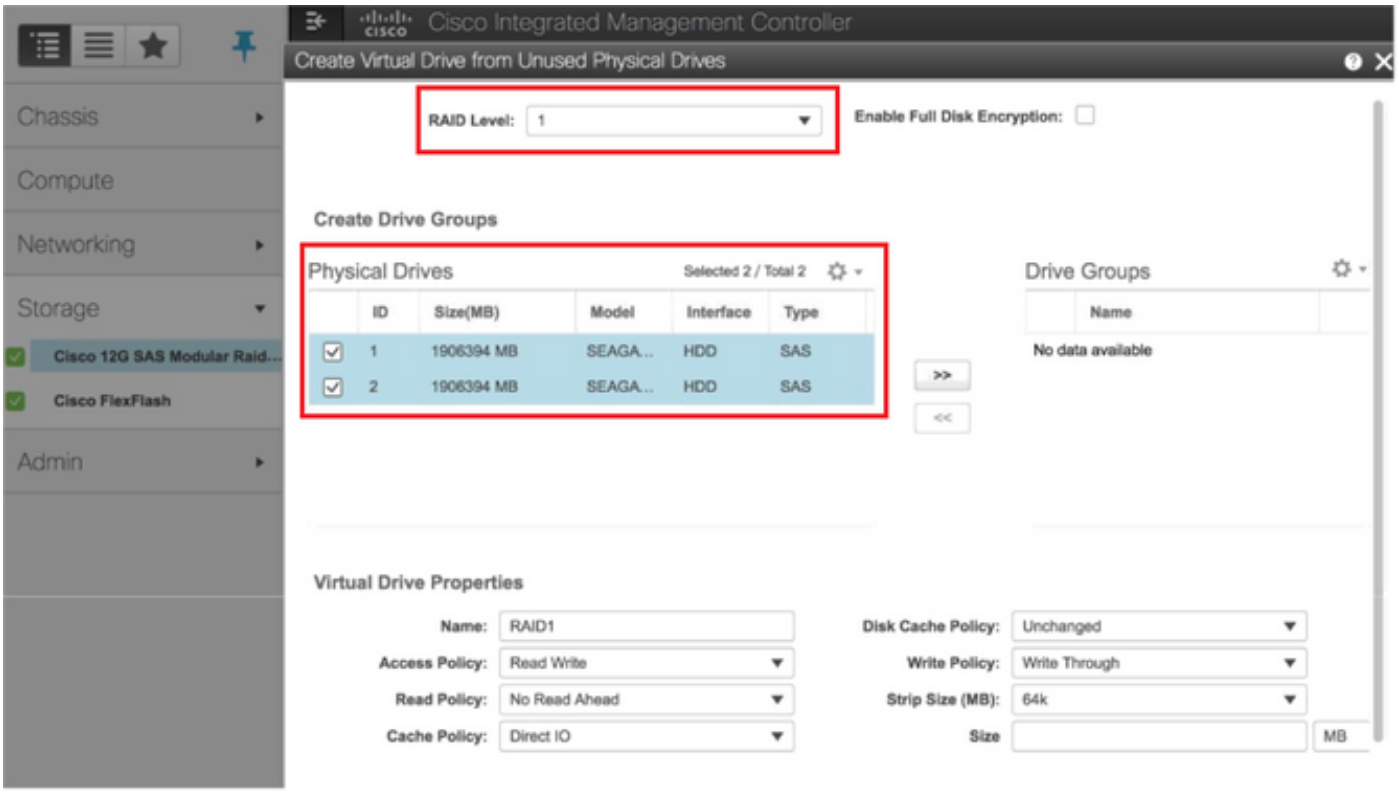

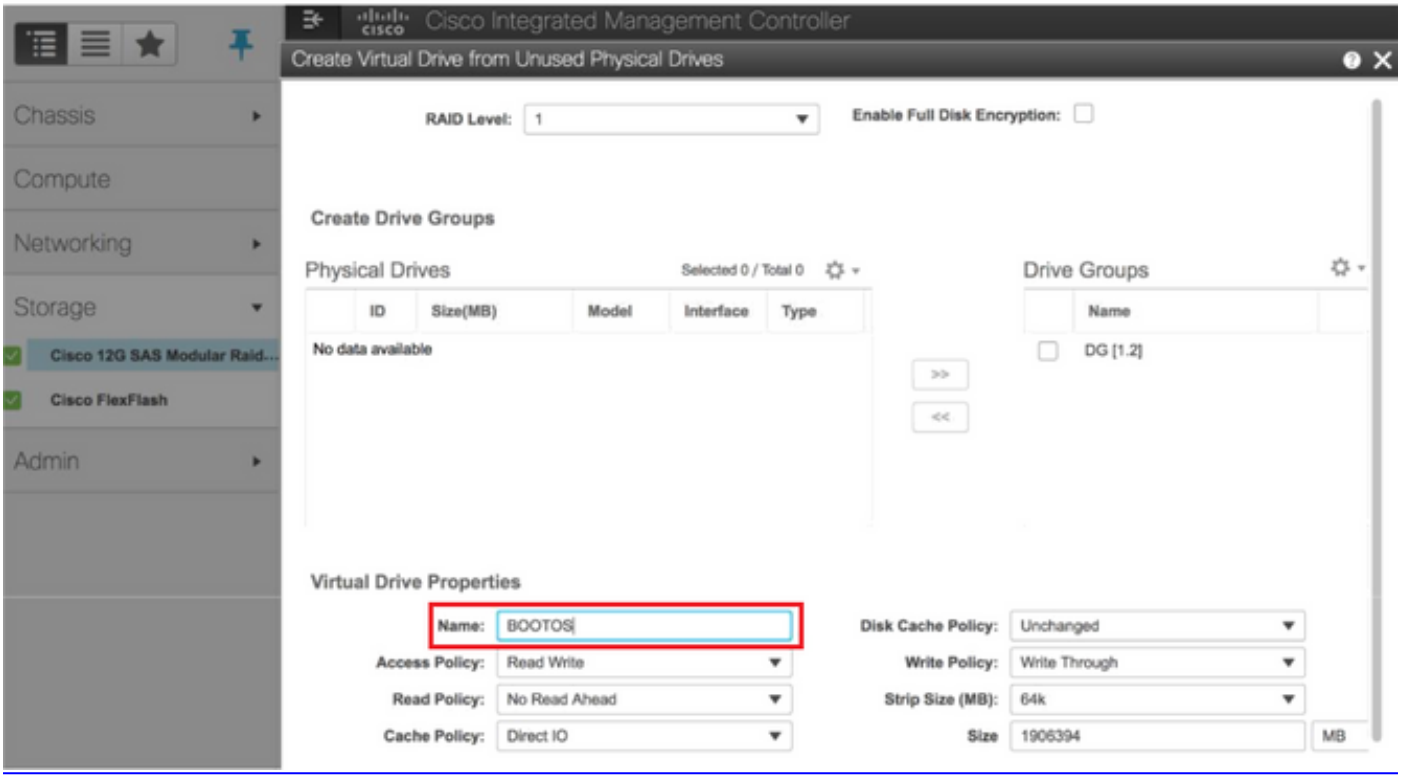

- Selecione o VD e configure **Set as Boot Drive (Definir como unidade de inicialização):**

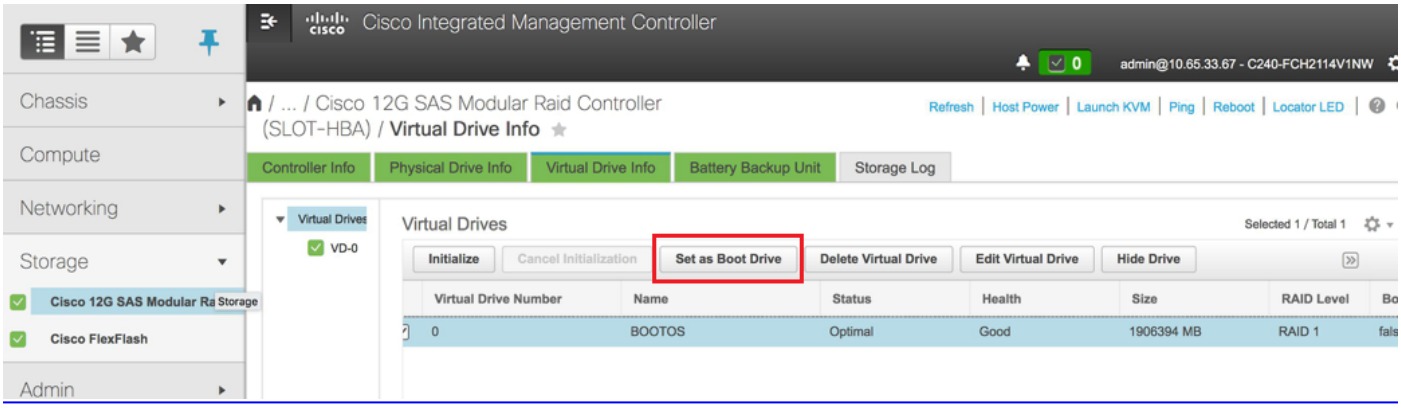

Etapa 6. Para habilitar o IPMI na LAN, navegue até **Admin > Communication Services > Communication Services**.

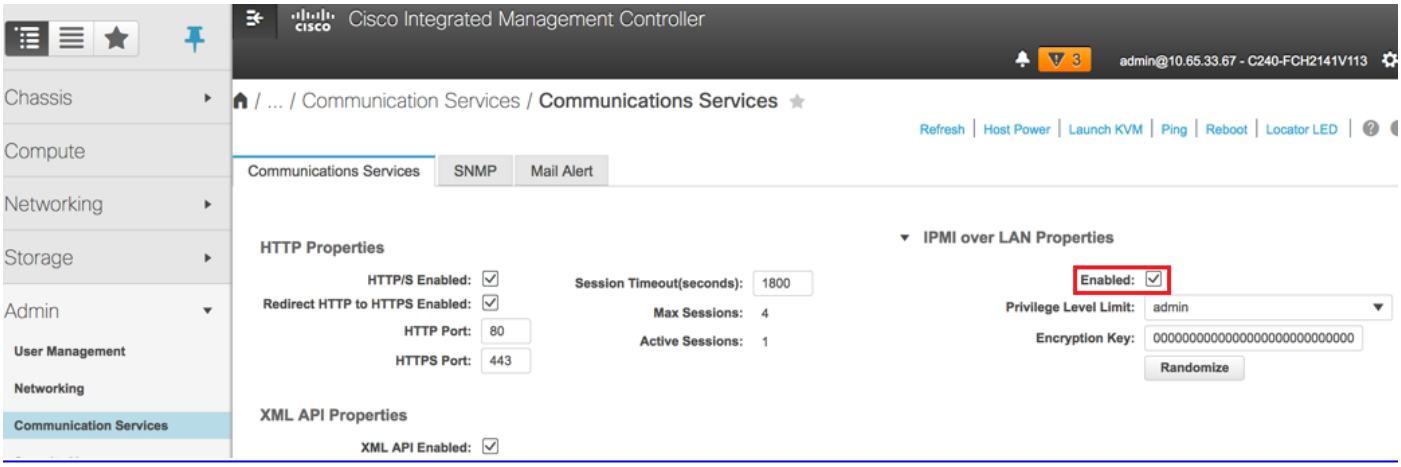

Passo 7. Para desabilitar o hyperthreading, navegue até **Compute > BIOS > Configure BIOS > Advanced > Processor Configuration**, como mostrado na imagem.

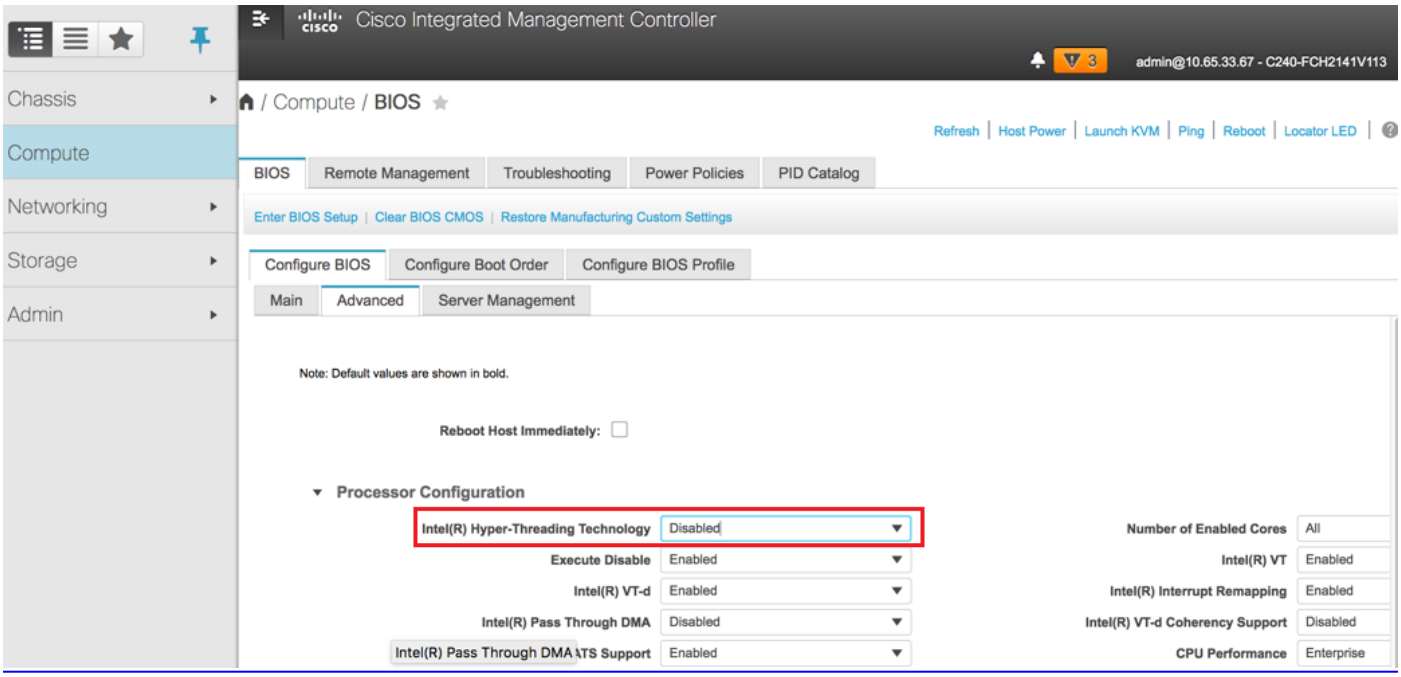

Note: A imagem é mostrada aqui e as etapas de configuração mencionadas nesta seção referem-se à versão de firmware 3.0(3e) e pode haver pequenas variações se você trabalhar em outras versões.

Substituição do nó do controlador na nuvem

Esta seção aborda as etapas necessárias para substituir o controlador com falha pelo novo na nuvem. Para isso, o script **Deployment.sh** usado para ativar a pilha seria reutilizado. No momento da implantação, na fase ControllerNodesPostDeployment, a atualização falharia devido a algumas limitações nos módulos Puppet. A intervenção manual é necessária antes de reiniciar o script de implantação.

Prepare-se para remover o nó do controlador com falha

Etapa 1. Identifique o índice do controlador com falha. O índice é o sufixo numérico no nome do controlador na saída da lista do servidor OpenStack. Neste exemplo, o índice é 2:

```
[stack@director ~]$ nova list | grep controller
| 5813a47e-af27-4fb9-8560-75decd3347b4 | pod1-controller-0  | ACTIVE | -          | Running
| ctlplane=192.200.0.152 |
| 457f023f-d077-45c9-bbea-dd32017d9708 | pod1-controller-1  | ACTIVE | -          | Running
| ctlplane=192.200.0.154 |
| d13bb207-473a-4e42-a1e7-05316935ed65 | pod1-controller-2  | ACTIVE | -          | Running
| ctlplane=192.200.0.151 |
```

Etapa 2. Crie um arquivo Yaml **~templates/remove-controller.yaml** que defina o nó a ser excluído. Use o índice encontrado na etapa anterior para a entrada na lista de recursos.

```
[stack@director ~]$ cat templates/remove-controller.yaml

parameters:
 ControllerRemovalPolicies:
   [{'resource_list': ['2']}]

parameter_defaults:
 CorosyncSettleTries: 5
```

Etapa 3. Faça uma cópia do script de implantação usado para instalar a nuvem geral e insira uma linha para incluir o arquivo **remove-controller.yaml** criado anteriormente.

```
[stack@director ~]$ cp deploy.sh deploy-removeController.sh
[stack@director ~]$ cat deploy-removeController.sh
time openstack overcloud deploy --templates \
-r ~/custom-templates/custom-roles.yaml \
-e /home/stack/templates/remove-controller.yaml \
-e /usr/share/openstack-tripleo-heat-templates/environments/puppet-pacemaker.yaml \
-e /usr/share/openstack-tripleo-heat-templates/environments/network-isolation.yaml \
-e /usr/share/openstack-tripleo-heat-templates/environments/storage-environment.yaml \
-e /usr/share/openstack-tripleo-heat-templates/environments/neutron-sriov.yaml \
-e ~/custom-templates/network.yaml \
-e ~/custom-templates/ceph.yaml \
-e ~/custom-templates/compute.yaml \
-e ~/custom-templates/layout-removeController.yaml \
-e ~/custom-templates/rabbitmq.yaml \
--stack pod1 \
--debug \
--log-file overcloudDeploy_$(date +%m_%d_%y__%H_%M_%S).log \
--neutron-flat-networks phys_pcie1_0,phys_pcie1_1,phys_pcie4_0,phys_pcie4_1 \
--neutron-network-vlan-ranges datacentre:101:200 \
--neutron-disable-tunneling \
```

```
--verbose --timeout 180
```

Etapa 4. Identifique o ID do controlador a ser substituído, usando os comandos mencionados aqui e mova-o para o modo de manutenção.

```
[stack@director ~]$ nova list | grep controller

| 5813a47e-af27-4fb9-8560-75decd3347b4 | pod1-controller-0  | ACTIVE | -          | Running
| ctlplane=192.200.0.152 |

| 457f023f-d077-45c9-bbea-dd32017d9708 | pod1-controller-1  | ACTIVE | -          | Running
| ctlplane=192.200.0.154 |

| d13bb207-473a-4e42-a1e7-05316935ed65 | pod1-controller-2  | ACTIVE | -          | Running
| ctlplane=192.200.0.151 |


[stack@director ~]$ openstack baremetal node list | grep d13bb207-473a-4e42-a1e7-05316935ed65

| e7c32170-c7d1-4023-b356-e98564a9b85b | None | d13bb207-473a-4e42-a1e7-05316935ed65 | power
off   | active             | False       |


[stack@b10-ospd ~]$ openstack baremetal node maintenance set e7c32170-c7d1-4023-b356-
e98564a9b85b


[stack@director~]$ openstack baremetal node list | grep True

| e7c32170-c7d1-4023-b356-e98564a9b85b | None | d13bb207-473a-4e42-a1e7-05316935ed65 | power
off   | active             | True        |
```

Etapa 5. Para garantir que o DB seja executado no momento do procedimento de substituição, remova o Galera do controle do marca-passo e execute esse comando em um dos controladores ativos.

```
[root@pod1-controller-0 ~]# sudo pcs resource unmanage galera
[root@pod1-controller-0 ~]# sudo pcs status

Cluster name: tripleo_cluster
Stack: corosync
Current DC: pod1-controller-0 (version 1.1.15-11.el7_3.4-e174ec8) - partition with quorum
Last updated: Thu Nov 16 16:51:18 2017                  Last change: Thu Nov 16 16:51:12 2017
by root via crm_resource on pod1-controller-0
3 nodes and 22 resources configured
Online: [ pod1-controller-0 pod1-controller-1 ]
OFFLINE: [ pod1-controller-2 ]

Full list of resources:

 ip-11.120.0.109         (ocf::heartbeat:IPaddr2):        Started pod1-controller-0
 ip-172.25.22.109        (ocf::heartbeat:IPaddr2):        Started pod1-controller-1
 ip-192.200.0.107        (ocf::heartbeat:IPaddr2):        Started pod1-controller-0

 Clone Set: haproxy-clone [haproxy]
    Started: [ pod1-controller-0 pod1-controller-1 ]
    Stopped: [ pod1-controller-2 ]

Master/Slave Set: galera-master [galera] (unmanaged)
    galera         (ocf::heartbeat:galera):          Master pod1-controller-0 (unmanaged)
    galera         (ocf::heartbeat:galera):          Master pod1-controller-1 (unmanaged)
```

```
    Stopped: [ pod1-controller-2 ]
 ip-11.120.0.110         (ocf::heartbeat:IPaddr2):       Started pod1-controller-0
 ip-11.119.0.110         (ocf::heartbeat:IPaddr2):       Started pod1-controller-1

<snip>
```

Prepare-se para adicionar um novo nó de controlador

Etapa 1. Crie um arquivo **controllerRMA.json** com apenas os novos detalhes do controlador. Verifique se o número de índice no novo controlador não foi usado antes. Normalmente, incremente para o próximo número de controlador mais alto.

Exemplo: O anterior mais alto foi o Controller-2, então crie o Controller-3.

> **Note**: Lembre-se do formato json.

```
[stack@director ~]$ cat controllerRMA.json
{
   "nodes": [
       {
           "mac": [
               <MAC_ADDRESS>
           ],
           "capabilities": "node:controller-3,boot_option:local",
           "cpu": "24",
           "memory": "256000",
           "disk": "3000",
           "arch": "x86_64",
           "pm_type": "pxe_ipmitool",
           "pm_user": "admin",
           "pm_password": "<PASSWORD>",
           "pm_addr": "<CIMC_IP>"
       }
   ]
}
```

Etapa 2. Importe o novo nó com o uso do arquivo json criado na etapa anterior.

```
[stack@director ~]$ openstack baremetal import --json controllerRMA.json

Started Mistral Workflow. Execution ID: 67989c8b-1225-48fe-ba52-3a45f366e7a0

Successfully registered node UUID 048ccb59-89df-4f40-82f5-3d90d37ac7dd

Started Mistral Workflow. Execution ID: c6711b5f-fa97-4c86-8de5-b6bc7013b398

Successfully set all nodes to available.

[stack@director ~]$ openstack baremetal node list | grep available

| 048ccb59-89df-4f40-82f5-3d90d37ac7dd | None | None                                    | power
off   | available          | False
```

Etapa 3. Defina o nó para gerenciar o estado.

```
[stack@director ~]$ openstack baremetal node manage 048ccb59-89df-4f40-82f5-3d90d37ac7dd
[stack@director ~]$ openstack baremetal node list | grep off
| 048ccb59-89df-4f40-82f5-3d90d37ac7dd | None | None | power off | manageable | False |
```

Etapa 4. Façam introspecção.

```
[stack@director ~]$ openstack overcloud node introspect 048ccb59-89df-4f40-82f5-3d90d37ac7dd --
provide
Started Mistral Workflow. Execution ID: f73fb275-c90e-45cc-952b-bfc25b9b5727
Waiting for introspection to finish...
Successfully introspected all nodes.
Introspection completed.
Started Mistral Workflow. Execution ID: a892b456-eb15-4c06-b37e-5bc3f6c37c65
Successfully set all nodes to available

[stack@director ~]$ openstack baremetal node list | grep available
| 048ccb59-89df-4f40-82f5-3d90d37ac7dd | None | None                                    | power
off   | available          | False       |
```

Etapa 5. Marque o nó disponível com as novas propriedades do controlador. Certifique-se de usar a ID da controladora designada para a nova controladora, conforme usada no arquivo controllerRMA.json.

```
[stack@director ~]$ openstack baremetal node set --property capabilities='node:controller-
3,profile:control,boot_option:local' 048ccb59-89df-4f40-82f5-3d90d37ac7dd
```

Etapa 6. No script de implantação, há um modelo personalizado chamado layout.yaml que, entre outras coisas, especifica quais endereços IP são atribuídos aos controladores para as várias interfaces. Em uma nova pilha, há 3 endereços definidos para Controller-0, Controller-1 e Controller-2. Ao adicionar um novo controlador, certifique-se de adicionar um próximo endereço IP em sequência para cada sub-rede.

```
ControllerIPs:
internal_api:
- 11.120.0.10
- 11.120.0.11
- 11.120.0.12
 - 11.120.0.13
tenant:
- 11.117.0.10
- 11.117.0.11
- 11.117.0.12
 - 11.117.0.13
storage:
- 11.118.0.10
- 11.118.0.11
- 11.118.0.12
 - 11.118.0.13
storage_mgmt:
- 11.119.0.10
- 11.119.0.11
- 11.119.0.12
 - 11.119.0.13
```

Passo 7. Agora execute o comando de distribuição-remoção.sh que foi criado anteriormente, para remover o nó antigo e adicionar o novo nó.

> Note: Espera-se que esta etapa falhe em ControllerNodesDeployment_Step1. Nesse ponto, é necessária uma intervenção manual.

```
[stack@b10-ospd ~]$ ./deploy-addController.sh
```

```
START with options: [u'overcloud', u'deploy', u'--templates', u'-r', u'/home/stack/custom-
templates/custom-roles.yaml', u'-e', u'/usr/share/openstack-tripleo-heat-
templates/environments/puppet-pacemaker.yaml', u'-e', u'/usr/share/openstack-tripleo-heat-
templates/environments/network-isolation.yaml', u'-e', u'/usr/share/openstack-tripleo-heat-
templates/environments/storage-environment.yaml', u'-e', u'/usr/share/openstack-tripleo-heat-
templates/environments/neutron-sriov.yaml', u'-e', u'/home/stack/custom-templates/network.yaml',
u'-e', u'/home/stack/custom-templates/ceph.yaml', u'-e', u'/home/stack/custom-
templates/compute.yaml', u'-e', u'/home/stack/custom-templates/layout-removeController.yaml',
u'-e', u'/home/stack/custom-templates/rabbitmq.yaml', u'--stack', u'newtonoc', u'--debug', u'--
log-file', u'overcloudDeploy_11_15_17__07_46_35.log', u'--neutron-flat-networks',
u'phys_pcie1_0,phys_pcie1_1,phys_pcie4_0,phys_pcie4_1', u'--neutron-network-vlan-ranges',
u'datacentre:101:200', u'--neutron-disable-tunneling', u'--verbose', u'--timeout', u'180']
 :
DeploymentError: Heat Stack update failed
END return value: 1

real     42m1.525s
user     0m3.043s
sys      0m0.614s
```

O progresso/status da implantação pode ser monitorado com estes comandos:

```
[stack@director~]$ openstack stack list --nested | grep -iv complete

+--------------------------------------+-------------------------------------------------------
---------------------------------------------------------------------------------------
-----------------+-----------------+----------------------+----------------------+-------------
------------------------+

| ID                                   | Stack
Name
                                                                  | Stack Status    | Creation
Time        | Updated Time         | Parent                               |

+--------------------------------------+-------------------------------------------------------
---------------------------------------------------------------------------------------
-----------------+-----------------+----------------------+----------------------+-------------
------------------------+

| c1e338f2-877e-4817-93b4-9a3f0c0b3d37 | pod1-AllNodesDeploySteps-5psegydpwxij-
ComputeDeployment_Step1-swnuzjixac43
                                                                      | UPDATE_FAILED   |
2017-10-08T14:06:07Z | 2017-11-16T18:09:43Z | e90f00ef-2499-4ec3-90b4-d7def6e97c47 |

| 1db4fef4-45d3-4125-bd96-2cc3297a69ff | pod1-AllNodesDeploySteps-5psegydpwxij-
ControllerDeployment_Step1-
hmn3hpruubcn
   | UPDATE_FAILED   | 2017-10-08T14:03:05Z | 2017-11-16T18:12:12Z | e90f00ef-2499-4ec3-90b4-
d7def6e97c47 |

| e90f00ef-2499-4ec3-90b4-d7def6e97c47 | pod1-AllNodesDeploySteps-
5psegydpwxij
                                          | UPDATE_FAILED   | 2017-10-08T13:59:25Z | 2017-11-
16T18:09:25Z | 6c4b604a-55a4-4a19-9141-28c844816c0d |

| 6c4b604a-55a4-4a19-9141-28c844816c0d |
pod1
                                                                  | UPDATE_FAILED   | 2017-10-
08T12:37:11Z | 2017-11-16T17:35:35Z | None                                 |

+--------------------------------------+-------------------------------------------------------
---------------------------------------------------------------------------------------
-----------------+-----------------+----------------------+----------------------+-------------
```

```
-----------------------+
```

Intervenção manual

Etapa 1. No servidor OSP-D, execute o comando OpenStack server list para listar os controladores disponíveis. O controlador recém-adicionado deve aparecer na lista.

```
[stack@director ~]$ openstack server list | grep controller
| 3e6c3db8-ba24-48d9-b0e8-1e8a2eb8b5ff | pod1-controller-3  | ACTIVE | ctlplane=192.200.0.103 |
overcloud-full |
| 457f023f-d077-45c9-bbea-dd32017d9708 | pod1-controller-1  | ACTIVE | ctlplane=192.200.0.154 |
overcloud-full |
| 5813a47e-af27-4fb9-8560-75decd3347b4 | pod1-controller-0  | ACTIVE | ctlplane=192.200.0.152 |
overcloud-full |
```

Etapa 2. Conecte-se a um dos controladores ativos (não ao controlador recém-adicionado) e examine o arquivo **/etc/corosync/corosycn.conf**. Localize a **lista** que atribui um **nó** a cada controlador. Localize a entrada do nó com falha e anote seu **nó:**

```
[root@pod1-controller-0 ~]# cat /etc/corosync/corosync.conf
totem {
   version: 2
   secauth: off
   cluster_name: tripleo_cluster
   transport: udpu
   token: 10000
}

nodelist {
   node {
       ring0_addr: pod1-controller-0
       nodeid: 5
   }
   node {
       ring0_addr: pod1-controller-1
       nodeid: 7
   }
   node {
       ring0_addr: pod1-controller-2
       nodeid: 8
   }
}
```

Etapa 3. Faça login em cada um dos controladores ativos. Remova o nó com falha e reinicie o serviço. Nesse caso, remova **pod1-controller-2**. Não execute esta ação no controlador recém-adicionado.

```
[root@pod1-controller-0 ~]# sudo pcs cluster localnode remove pod1-controller-2
pod1-controller-2: successfully removed!
[root@pod1-controller-0 ~]# sudo pcs cluster reload corosync
Corosync reloaded

[root@pod1-controller-1 ~]# sudo pcs cluster localnode remove pod1-controller-2
pod1-controller-2: successfully removed!
[root@pod1-controller-1 ~]# sudo pcs cluster reload corosync
Corosync reloaded
```

Etapa 4. Execute esse comando de um dos controladores ativos para excluir o nó com falha do cluster.

```
[root@pod1-controller-0 ~]# sudo crm_node -R pod1-controller-2 --force
```

Etapa 5. Execute este comando de um dos controladores ativos para excluir o nó com falha do cluster **rabbitmq**.

```
[root@pod1-controller-0 ~]# sudo rabbitmqctl forget_cluster_node rabbit@pod1-controller-2
Removing node 'rabbit@newtonoc-controller-2' from cluster ...
```

Etapa 6. Exclua o nó com falha do MongoDB. Para fazer isso, você precisa encontrar o nó Mongo ativo. Use **netstat** para encontrar o endereço IP do host.

```
[root@pod1-controller-0 ~]# sudo netstat -tulnp | grep 27017
tcp        0      0 11.120.0.10:27017       0.0.0.0:*               LISTEN
219577/mongod
```

Passo 7. Faça login no nó e verifique se ele é o mestre com o uso do endereço IP e do número de porta do comando anterior.

```
[heat-admin@pod1-controller-0 ~]$ echo "db.isMaster()" | mongo --host 11.120.0.10:27017
MongoDB shell version: 2.6.11
connecting to: 11.120.0.10:27017/test
{
        "setName" : "tripleo",
        "setVersion" : 9,
        "ismaster" : true,
        "secondary" : false,
        "hosts" : [
                "11.120.0.10:27017",
                "11.120.0.12:27017",
                "11.120.0.11:27017"
        ],
        "primary" : "11.120.0.10:27017",
        "me" : "11.120.0.10:27017",
        "electionId" : ObjectId("5a0d2661218cb0238b582fb1"),
        "maxBsonObjectSize" : 16777216,
        "maxMessageSizeBytes" : 48000000,
        "maxWriteBatchSize" : 1000,
        "localTime" : ISODate("2017-11-16T18:36:34.473Z"),
        "maxWireVersion" : 2,
        "minWireVersion" : 0,
        "ok" : 1
}
```

Se o nó não for o mestre, faça login no outro controlador ativo e execute a mesma etapa.

Etapa 8. Do mestre, liste os nós disponíveis com o uso do comando **rs.status()**. Localize o nó antigo/sem resposta e identifique o nome do nó mongo.

```
[root@pod1-controller-0 ~]# mongo --host 11.120.0.10
MongoDB shell version: 2.6.11
connecting to: 11.120.0.10:27017/test
<snip>
tripleo:PRIMARY> rs.status()
{
        "set" : "tripleo",
        "date" : ISODate("2017-11-14T13:27:14Z"),
```

```
        "myState" : 1,
        "members" : [
                {
                        "_id" : 0,
                        "name" : "11.120.0.10:27017",
                        "health" : 1,
                        "state" : 1,
                        "stateStr" : "PRIMARY",
                        "uptime" : 418347,
                        "optime" : Timestamp(1510666033, 1),
                        "optimeDate" : ISODate("2017-11-14T13:27:13Z"),
                        "electionTime" : Timestamp(1510247693, 1),
                        "electionDate" : ISODate("2017-11-09T17:14:53Z"),
                        "self" : true
                },
                {
                        "_id" : 2,
                        "name" : "11.120.0.12:27017",
                        "health" : 1,
                        "state" : 2,
                        "stateStr" : "SECONDARY",
                        "uptime" : 418347,
                        "optime" : Timestamp(1510666033, 1),
                        "optimeDate" : ISODate("2017-11-14T13:27:13Z"),
                        "lastHeartbeat" : ISODate("2017-11-14T13:27:13Z"),
                        "lastHeartbeatRecv" : ISODate("2017-11-14T13:27:13Z"),
                        "pingMs" : 0,
                        "syncingTo" : "11.120.0.10:27017"
                },
                {
                        "_id" : 3,
                        "name" : "11.120.0.11:27017
                        "health" : 0,
                        "state" : 8,
                        "stateStr" : "(not reachable/healthy)",
                        "uptime" : 0,
                        "optime" : Timestamp(1510610580, 1),
                        "optimeDate" : ISODate("2017-11-13T22:03:00Z"),
                        "lastHeartbeat" : ISODate("2017-11-14T13:27:10Z"),
                        "lastHeartbeatRecv" : ISODate("2017-11-13T22:03:01Z"),
                        "pingMs" : 0,
                        "syncingTo" : "11.120.0.10:27017"
                }
        ],
        "ok" : 1
}
```

Etapa 9. Do mestre, exclua o nó com falha com o uso do comando **rs.remove**. Alguns erros são vistos quando você executa este comando, mas verifique o status mais uma vez para descobrir que o nó foi removido:

```
[root@pod1-controller-0 ~]$ mongo --host 11.120.0.10
<snip>
tripleo:PRIMARY> rs.remove('11.120.0.12:27017')
2017-11-16T18:41:04.999+0000 DBClientCursor::init call() failed
2017-11-16T18:41:05.000+0000 Error: error doing query: failed at src/mongo/shell/query.js:81
2017-11-16T18:41:05.001+0000 trying reconnect to 11.120.0.10:27017 (11.120.0.10) failed
2017-11-16T18:41:05.003+0000 reconnect 11.120.0.10:27017 (11.120.0.10) ok

tripleo:PRIMARY> rs.status()
{
        "set" : "tripleo",
        "date" : ISODate("2017-11-16T18:44:11Z"),
```

```
        "myState" : 1,
        "members" : [
                {
                        "_id" : 3,
                        "name" : "11.120.0.11:27017",
                        "health" : 1,
                        "state" : 2,
                        "stateStr" : "SECONDARY",
                        "uptime" : 187,
                        "optime" : Timestamp(1510857848, 3),
                        "optimeDate" : ISODate("2017-11-16T18:44:08Z"),
                        "lastHeartbeat" : ISODate("2017-11-16T18:44:11Z"),
                        "lastHeartbeatRecv" : ISODate("2017-11-16T18:44:09Z"),
                        "pingMs" : 0,
                        "syncingTo" : "11.120.0.10:27017"
                },
                {
                        "_id" : 4,
                        "name" : "11.120.0.10:27017",
                        "health" : 1,
                        "state" : 1,
                        "stateStr" : "PRIMARY",
                        "uptime" : 89820,
                        "optime" : Timestamp(1510857848, 3),
                        "optimeDate" : ISODate("2017-11-16T18:44:08Z"),
                        "electionTime" : Timestamp(1510811232, 1),
                        "electionDate" : ISODate("2017-11-16T05:47:12Z"),
                        "self" : true
                }
        ],
        "ok" : 1
}
tripleo:PRIMARY> exit
bye
```

Etapa 10. Execute esse comando para atualizar a lista de nós de controlador ativos. Inclua o novo nó de controlador nesta lista.

```
[root@pod1-controller-0 ~]# sudo pcs resource update galera wsrep_cluster_address=gcomm://pod1-controller-0,pod1-controller-1,pod1-controller-2
```

Etapa 11. Copie esses arquivos de um controlador que já existe para o novo controlador:

- /etc/sysconfig/clustercheck
- /root/.my.cnf

```
On existing controller:

[root@pod1-controller-0 ~]# scp /etc/sysconfig/clustercheck stack@192.200.0.1:/tmp/.
[root@pod1-controller-0 ~]# scp /root/.my.cnf stack@192.200.0.1:/tmp/my.cnf

On new controller:

[root@pod1-controller-3 ~]# cd /etc/sysconfig

[root@pod1-controller-3 sysconfig]# scp stack@192.200.0.1:/tmp/clustercheck .

[root@pod1-controller-3 sysconfig]# cd /root

[root@pod1-controller-3 ~]# scp stack@192.200.0.1:/tmp/my.cnf .my.cnf
```

Etapa 12. Execute o comando **cluster node add** de um dos controladores já existentes.

```
[root@pod1-controller-1 ~]# sudo pcs cluster node add pod1-controller-3

Disabling SBD service...
pod1-controller-3: sbd disabled
pod1-controller-0: Corosync updated
pod1-controller-1: Corosync updated

Setting up corosync...
pod1-controller-3: Succeeded
Synchronizing pcsd certificates on nodes pod1-controller-3...
pod1-controller-3: Success

Restarting pcsd on the nodes in order to reload the certificates...
pod1-controller-3: Success
```

Etapa 13. Faça login em cada controlador e visualize o arquivo **/etc/corosync/corosync.conf**. Verifique se o novo controlador está listado e se o **nó** atribuído a esse controlador é o próximo número na sequência que não foi usado anteriormente. Assegure-se de que essa alteração seja feita em todos os três controladores.

```
[root@pod1-controller-1 ~]# cat /etc/corosync/corosync.conf
totem {
   version: 2
   secauth: off
   cluster_name: tripleo_cluster
   transport: udpu
   token: 10000
}
nodelist {
   node {
       ring0_addr: pod1-controller-0
       nodeid: 5
   }
   node {
       ring0_addr: pod1-controller-1
       nodeid: 7
   }
   node {
       ring0_addr: pod1-controller-3
        nodeid: 6
   }
}
quorum {
   provider: corosync_votequorum
}
logging {
   to_logfile: yes
   logfile: /var/log/cluster/corosync.log
   to_syslog: yes
}
```

Por exemplo, **/etc/corosync/corosync.conf** após modificação:

```
totem {
version: 2
secauth: off
cluster_name: tripleo_cluster
```

```
transport: udpu
token: 10000
}
nodelist {
   node {
       ring0_addr: pod1-controller-0
       nodeid: 5
   }
   node {
       ring0_addr: pod1-controller-1
       nodeid: 7
   }
   node {
       ring0_addr: pod1-controller-3
       nodeid: 9
   }
}
quorum {
   provider: corosync_votequorum
}
logging {
   to_logfile: yes
   logfile: /var/log/cluster/corosync.log
   to_syslog: yes
}
```

Etapa 14. Reinicie a **sincronização corporativa** nos controladores ativos. Não inicie a **sincronização corporativa** no novo controlador.

```
[root@pod1-controller-0 ~]# sudo pcs cluster reload corosync
[root@pod1-controller-1 ~]# sudo pcs cluster reload corosync
```

Etapa 15. Inicie o novo nó controlador a partir de um dos controladores em ação.

```
[root@pod1-controller-1 ~]# sudo pcs cluster start pod1-controller-3
```

Etapa 16. Reinicie Galera de um dos controladores.

```
[root@pod1-controller-1 ~]#  sudo pcs cluster start pod1-controller-3

pod1-controller-0: Starting Cluster...

[root@pod1-controller-1 ~]# sudo pcs resource cleanup galera
Cleaning up galera:0 on pod1-controller-0, removing fail-count-galera
Cleaning up galera:0 on pod1-controller-1, removing fail-count-galera
Cleaning up galera:0 on pod1-controller-3, removing fail-count-galera
 * The configuration prevents the cluster from stopping or starting 'galera-master' (unmanaged)

Waiting for 3 replies from the CRMd... OK

[root@pod1-controller-1 ~]#
[root@pod1-controller-1 ~]# sudo pcs resource manage galera
```

Etapa 17. O cluster está no modo de manutenção. Desative o modo de manutenção para que os serviços sejam iniciados.

```
[root@pod1-controller-2 ~]# sudo pcs property set maintenance-mode=false --wait
```

Etapa 18. Verifique o status dos PCs do Galera até que todos os 3 controladores sejam listados como mestres no Galera.

Note: Para configurações grandes, pode levar algum tempo para sincronizar DBs.

```
[root@pod1-controller-1 ~]# sudo pcs status | grep galera -A1

Master/Slave Set: galera-master [galera]
     Masters: [ pod1-controller-0 pod1-controller-1 pod1-controller-3 ]
```

Etapa 19. Mude o cluster para o modo de manutenção.

```
[root@pod1-controller-1~]# sudo pcs property set maintenance-mode=true --wait

[root@pod1-controller-1 ~]# pcs cluster status
Cluster Status:
 Stack: corosync
 Current DC: pod1-controller-0 (version 1.1.15-11.el7_3.4-e174ec8) - partition with quorum
 Last updated: Thu Nov 16 19:17:01 2017                  Last change: Thu Nov 16 19:16:48 2017
by root via cibadmin on pod1-controller-1
              *** Resource management is DISABLED ***
  The cluster will not attempt to start, stop or recover services

PCSD Status:
 pod1-controller-3: Online
 pod1-controller-0: Online
 pod1-controller-1: Online
```

Etapa 20. Execute novamente o script de implantação executado anteriormente. Desta vez, deve ser bem sucedido.

```
[stack@director ~]$ ./deploy-addController.sh
START with options: [u'overcloud', u'deploy', u'--templates', u'-r', u'/home/stack/custom-
templates/custom-roles.yaml', u'-e', u'/usr/share/openstack-tripleo-heat-
templates/environments/puppet-pacemaker.yaml', u'-e', u'/usr/share/openstack-tripleo-heat-
templates/environments/network-isolation.yaml', u'-e', u'/usr/share/openstack-tripleo-heat-
templates/environments/storage-environment.yaml', u'-e', u'/usr/share/openstack-tripleo-heat-
templates/environments/neutron-sriov.yaml', u'-e', u'/home/stack/custom-templates/network.yaml',
u'-e', u'/home/stack/custom-templates/ceph.yaml', u'-e', u'/home/stack/custom-
templates/compute.yaml', u'-e', u'/home/stack/custom-templates/layout-removeController.yaml',
u'--stack', u'newtonoc', u'--debug', u'--log-file', u'overcloudDeploy_11_14_17__13_53_12.log',
u'--neutron-flat-networks', u'phys_pcie1_0,phys_pcie1_1,phys_pcie4_0,phys_pcie4_1', u'--neutron-
network-vlan-ranges', u'datacentre:101:200', u'--neutron-disable-tunneling', u'--verbose', u'--
timeout', u'180']
options: Namespace(access_key='', access_secret='***', access_token='***',
access_token_endpoint='', access_token_type='', aodh_endpoint='', auth_type='',
auth_url='https://192.200.0.2:13000/v2.0', authorization_code='', cacert=None, cert='',
client_id='', client_secret='***', cloud='', consumer_key='', consumer_secret='***', debug=True,
default_domain='default', default_domain_id='', default_domain_name='', deferred_help=False,
discovery_endpoint='', domain_id='', domain_name='', endpoint='', identity_provider='',
identity_provider_url='', insecure=None, inspector_api_version='1', inspector_url=None,
interface='', key='', log_file=u'overcloudDeploy_11_14_17__13_53_12.log', murano_url='',
old_profile=None, openid_scope='', os_alarming_api_version='2',
os_application_catalog_api_version='1', os_baremetal_api_version='1.15', os_beta_command=False,
os_compute_api_version='', os_container_infra_api_version='1',
os_data_processing_api_version='1.1', os_data_processing_url='', os_dns_api_version='2',
os_identity_api_version='', os_image_api_version='1', os_key_manager_api_version='1',
os_metrics_api_version='1', os_network_api_version='', os_object_api_version='',
os_orchestration_api_version='1', os_project_id=None, os_project_name=None,
os_queues_api_version='2', os_tripleoclient_api_version='1', os_volume_api_version='',
os_workflow_api_version='2', passcode='', password='***', profile=None, project_domain_id='',
project_domain_name='', project_id='', project_name='admin', protocol='', redirect_uri='',
region_name='', roles='', timing=False, token='***', trust_id='', url='', user='',
```

```
user_domain_id='', user_domain_name='', user_id='', username='admin', verbose_level=3,
verify=None)
Auth plugin password selected

Starting new HTTPS connection (1): 192.200.0.2
"POST /v2/action_executions HTTP/1.1" 201 1696
HTTP POST https://192.200.0.2:13989/v2/action_executions 201
Overcloud Endpoint: http://172.25.22.109:5000/v2.0
Overcloud Deployed
clean_up DeployOvercloud:
END return value: 0

real     54m17.197s
user     0m3.421s
sys       0m0.670s
```

Verifique os serviços em nuvem no controlador

Verifique se todos os serviços gerenciados são executados corretamente nos nós do controlador.

```
[heat-admin@pod1-controller-2 ~]$ sudo pcs status
```

Finalizar os roteadores do agente L3

Verifique os roteadores para garantir que os agentes L3 estejam hospedados corretamente. Certifique-se de originar o arquivo overcloudrc ao executar essa verificação.

Etapa 1. Localize o nome do roteador.

```
[stack@director~]$ source corerc
[stack@director ~]$ neutron router-list

+--------------------------------------+------+-----------------------------------------------
-----------------+-------------+------+

| id                                   | name |
external_gateway_info                                          | distributed | ha   |

+--------------------------------------+------+-----------------------------------------------
-----------------+-------------+------+

| d814dc9d-2b2f-496f-8c25-24911e464d02 | main | {"network_id": "18c4250c-e402-428c-87d6-
a955157d50b5",            | False       | True |
```

Neste exemplo, o nome do roteador é main.

Etapa 2. Liste todos os agentes L3 para localizar o UUID do nó com falha e o novo nó.

```
[stack@director ~]$ neutron agent-list | grep "neutron-l3-agent"

| 70242f5c-43ab-4355-abd6-9277f92e4ce6 | L3 agent           | pod1-controller-0.localdomain  |
nova              | :-)   | True           | neutron-l3-agent          |
| 8d2ffbcb-b6ff-42cd-b5b8-da31d8da8a40 | L3 agent           | pod1-controller-2.localdomain  |
nova              | xxx   | True           | neutron-l3-agent          |
| a410a491-e271-4938-8a43-458084ffe15d | L3 agent           | pod1-controller-3.localdomain  |
nova              | :-)   | True           | neutron-l3-agent          |
| cb4bc1ad-ac50-42e9-ae69-8a256d375136 | L3 agent           | pod1-controller-1.localdomain  |
```

```
nova              | :-)   | True          | neutron-l3-agent          |
```

Etapa 3. Neste exemplo, o agente L3 que corresponde a **pod1-controller-2.localdomain** deve ser removido do roteador e o que corresponde a **pod1-controller-3.localdomain** deve ser adicionado ao roteador.

```
[stack@director ~]$ neutron l3-agent-router-remove 8d2ffbcb-b6ff-42cd-b5b8-da31d8da8a40  main

Removed router main from L3 agent

[stack@director ~]$ neutron l3-agent-router-add a410a491-e271-4938-8a43-458084ffe15d main

Added router main to L3 agent
```

Etapa 4. Verifique a lista atualizada de agentes L3.

```
[stack@director ~]$ neutron l3-agent-list-hosting-router main

+--------------------------------------+-----------------------------------+----------------+---
----+----------+
| id                                   | host                              | admin_state_up |
alive | ha_state |
+--------------------------------------+-----------------------------------+----------------+---
----+----------+

| 70242f5c-43ab-4355-abd6-9277f92e4ce6 | pod1-controller-0.localdomain | True           | :-)
| standby  |
| a410a491-e271-4938-8a43-458084ffe15d | pod1-controller-3.localdomain | True           | :-)
| standby  |
| cb4bc1ad-ac50-42e9-ae69-8a256d375136 | pod1-controller-1.localdomain | True           | :-)
| active   |
+--------------------------------------+-----------------------------------+----------------+---
----+----------+
```

Etapa 5. Liste todos os serviços que são executados no nó do controlador removido e remova-os.

```
[stack@director ~]$ neutron  agent-list | grep controller-2

| 877314c2-3c8d-4666-a6ec-69513e83042d | Metadata agent     | pod1-controller-2.localdomain
|                   | xxx   | True           | neutron-metadata-agent    |
| 8d2ffbcb-b6ff-42cd-b5b8-da31d8da8a40 | L3 agent           | pod1-controller-2.localdomain  |
nova              | xxx   | True           | neutron-l3-agent          |
| 911c43a5-df3a-49ec-99ed-1d722821ec20 | DHCP agent         | pod1-controller-2.localdomain  |
nova              | xxx   | True           | neutron-dhcp-agent        |
| a58a3dd3-4cdc-48d4-ab34-612a6cd72768 | Open vSwitch agent | pod1-controller-2.localdomain
|                   | xxx   | True           | neutron-openvswitch-agent |

[stack@director ~]$ neutron agent-delete 877314c2-3c8d-4666-a6ec-69513e83042d
Deleted agent(s): 877314c2-3c8d-4666-a6ec-69513e83042d
[stack@director ~]$ neutron agent-delete 8d2ffbcb-b6ff-42cd-b5b8-da31d8da8a40
Deleted agent(s): 8d2ffbcb-b6ff-42cd-b5b8-da31d8da8a40
[stack@director ~]$ neutron agent-delete 911c43a5-df3a-49ec-99ed-1d722821ec20
Deleted agent(s): 911c43a5-df3a-49ec-99ed-1d722821ec20
[stack@director ~]$ neutron agent-delete a58a3dd3-4cdc-48d4-ab34-612a6cd72768
Deleted agent(s): a58a3dd3-4cdc-48d4-ab34-612a6cd72768

[stack@director ~]$ neutron agent-list | grep controller-2
[stack@director ~]$
```

Finalizar serviços de computação

Etapa 1. Verifique os itens **da nova lista de serviços** deixados do nó removido e exclua-os.

```
[stack@director ~]$ nova service-list | grep controller-2

| 615 | nova-consoleauth | pod1-controller-2.localdomain  | internal          | enabled | down
| 2017-11-16T16:08:14.000000 | -               |
| 618 | nova-scheduler   | pod1-controller-2.localdomain  | internal          | enabled | down
| 2017-11-16T16:08:13.000000 | -               |
| 621 | nova-conductor   | pod1-controller-2.localdomain  | internal          | enabled | down
| 2017-11-16T16:08:14.000000 | -

[stack@director ~]$ nova service-delete 615
[stack@director ~]$ nova service-delete 618
[stack@director ~]$ nova service-delete 621

stack@director ~]$ nova service-list | grep controller-2
```

Etapa 2. Certifique-se de que o processo de **console** seja executado em todos os controladores ou reinicie-o com o uso deste comando: **reinicialização do recurso pcs openstack-nova-consoleauth:**

```
[stack@director ~]$ nova service-list | grep consoleauth

| 601 | nova-consoleauth | pod1-controller-0.localdomain  | internal          | enabled | up
| 2017-11-16T21:00:10.000000 | -               |
| 608 | nova-consoleauth | pod1-controller-1.localdomain  | internal          | enabled | up
| 2017-11-16T21:00:13.000000 | -               |
| 622 | nova-consoleauth | pod1-controller-3.localdomain  | internal          | enabled | up
| 2017-11-16T21:00:13.000000 | -
```

Reinicie a cerca nos nós da controladora

Etapa 1. Verifique todos os controladores da rota IP para a nuvem inferior 192.0.0.0/8

```
[root@pod1-controller-3 ~]# ip route
default via 172.25.22.1 dev vlan101
11.117.0.0/24 dev vlan17  proto kernel  scope link  src 11.117.0.12
11.118.0.0/24 dev vlan18  proto kernel  scope link  src 11.118.0.12
11.119.0.0/24 dev vlan19  proto kernel  scope link  src 11.119.0.12
11.120.0.0/24 dev vlan20  proto kernel  scope link  src 11.120.0.12
169.254.169.254 via 192.200.0.1 dev eno1
172.25.22.0/24 dev vlan101  proto kernel  scope link  src 172.25.22.102
192.0.0.0/8 dev eno1  proto kernel  scope link  src 192.200.0.103
```

Etapa 2. Verifique a configuração **contínua** atual. Remova qualquer referência ao nó antigo do controlador.

```
[root@pod1-controller-3 ~]# sudo pcs stonith show --full
 Resource: my-ipmilan-for-controller-6 (class=stonith type=fence_ipmilan)
 Attributes: pcmk_host_list=pod1-controller-1 ipaddr=192.100.0.1 login=admin
passwd=Csco@123Starent lanplus=1
 Operations: monitor interval=60s (my-ipmilan-for-controller-6-monitor-interval-60s)
 Resource: my-ipmilan-for-controller-4 (class=stonith type=fence_ipmilan)
 Attributes: pcmk_host_list=pod1-controller-0 ipaddr=192.100.0.14 login=admin
passwd=Csco@123Starent lanplus=1
 Operations: monitor interval=60s (my-ipmilan-for-controller-4-monitor-interval-60s)
Resource: my-ipmilan-for-controller-7 (class=stonith type=fence_ipmilan)
 Attributes: pcmk_host_list=pod1-controller-2 ipaddr=192.100.0.15 login=admin
```

```
passwd=Csco@123Starent lanplus=1
 Operations: monitor interval=60s (my-ipmilan-for-controller-7-monitor-interval-60s)

[root@pod1-controller-3 ~]# pcs stonith delete my-ipmilan-for-controller-7
Attempting to stop: my-ipmilan-for-controller-7...Stopped
```

Etapa 3. Adicione a configuração confiável para um novo controlador.

```
[root@pod1-controller-3 ~]sudo pcs stonith create my-ipmilan-for-controller-8 fence_ipmilan
pcmk_host_list=pod1-controller-3 ipaddr=<CIMC_IP> login=admin passwd=<PASSWORD> lanplus=1 op
monitor interval=60s
```

Etapa 4. Reinicie a vedação de qualquer controlador e verifique o status.

```
[root@pod1-controller-1 ~]# sudo pcs property set stonith-enabled=true
[root@pod1-controller-3 ~]# pcs status

<snip>
my-ipmilan-for-controller-1 (stonith:fence_ipmilan): Started pod1-controller-3
my-ipmilan-for-controller-0 (stonith:fence_ipmilan): Started pod1-controller-3
my-ipmilan-for-controller-3 (stonith:fence_ipmilan): Started pod1-controller-3
```