

Server Replacement

Cloud Execution Environment

OPERATING INSTRUCTIONS

Copyright

© Ericsson AB 2016–2018. All rights reserved. No part of this document may be reproduced in any form without the written permission of the copyright owner.

Disclaimer

The contents of this document are subject to revision without notice due to continued progress in methodology, design and manufacturing. Ericsson shall have no liability for any error or damage of any kind resulting from the use of this document.

Trademark List

All trademarks mentioned herein are the property of their respective owners. These are shown in the document Trademark Information.



Contents

| | | |
|----------|---|-----------|
| 1 | Introduction | 1 |
| 1.1 | Scope | 1 |
| 1.2 | Prerequisites | 1 |
| 1.2.1 | Documents | 1 |
| 1.2.2 | Tools | 2 |
| 1.2.3 | Data | 2 |
| 1.2.4 | Conditions | 3 |
| 2 | Procedure | 4 |
| 2.1 | Preparations for Server Removal | 6 |
| 2.2 | Hardware Replacement and Server Configuration | 9 |
| 2.2.1 | Replacement of HDS-Based System | 9 |
| 2.2.2 | vFuel Discovery for Unmanaged Servers | 10 |
| 2.3 | Executing Installation Command | 10 |
| 2.4 | Data Collection | 12 |
| 2.5 | Fuel Synchronization | 12 |
| 2.6 | Concluding Routine | 12 |
| 3 | ScaleIO MDM Cluster Member Replacement | 14 |
| 3.1 | Procedure for ScaleIO MDM Replacement without Standby MDM | 14 |
| 3.2 | Reproducing Original Role Distribution | 18 |
| | Reference List | 20 |





1 Introduction

The document provides instructions on how to replace a server in the Cloud Execution Environment (CEE). Throughout this document, compute and data host hardware are referred to as server.

1.1 Scope

This document describes how to replace a server in CEE.

The document is applicable for the replacement of the following servers:

- Compute host, including host containing the virtual Cloud Infrastructure Controller (vCIC) node
- ScaleIO Data Server (SDS)

Replacement of unmanaged ScaleIO nodes is out of the scope of this document. However, some changes to the Meta Data Manager (MDM) cluster and SDSs of unmanaged ScaleIO must be propagated in the CEE region manually. For more information, refer to the document [Runtime Configuration Guide](#).

1.2 Prerequisites

This section provides information on the documents, tools, and conditions that apply to the procedure.

1.2.1 Documents

Before starting this procedure, ensure that the following documents have been read and understood:

Note: This is a prerequisite for hardware replacement and is the responsibility of the Data Center owner.

- [Personal Health and Safety Information](#)
- [System Safety Information](#)

The following document is referred and used in this procedure:

- Refer to the [Limitations and Workarounds for Cloud Execution Environment \(CEE\)](#), Reference [3], for any limitation or workaround that applies to the procedures described in this document.



- Depending on the used hardware environment, use the relevant document from the ones referred in Section 2.2 on page 9. The documents contain further prerequisites.
- [Data Collection Guideline](#)

1.2.2 Tools

The following tools are needed:

- An Electrostatic Discharge (ESD) wrist strap (part number LYB 250 01/14)

Note: This is a prerequisite for hardware replacement and is the responsibility of the Data Center owner.

- A computer with the ability to do a Secure Shell (SSH) logon to the vCIC

1.2.3 Data

A site-specific IP and VLAN plan is required.

The address variables used in the document IP and VLAN plan are used throughout this document, and are summarized in the following table.

Table 1 IP Variables Based on IP and VLAN plan

| VLAN | Variable Name | Factory Default IP Address Allocation |
|--------------|------------------|---------------------------------------|
| fuel_ctrl_sp | <vfuel_(static)> | 192.168.0.11 |

Other site-specific data is listed in the following table:

Table 2 Other Variables Used

| Resource | Variable Name | Additional Information |
|--|-----------------|---|
| External IP address of the vCIC | <vcic_address> | |
| Personal username to the vCIC | <personal_user> | |
| Password for the personal username to the vCIC | | |
| Password for the root user on vFuel | | |
| Name of the host to be replaced | <hostname> | Hostnames are specified by the following scheme: compute-<shelf_number>-<blade_number> |



Table 2 Other Variables Used

| Resource | Variable Name | Additional Information |
|-----------------------|-----------------|---|
| Master MDM IP address | <master_mdm_ip> | Either of the following can be used: <ul style="list-style-type: none"> • IP address of the MDM Master • IP addresses of all active MDMs (including the Master MDM), separated with a comma “,” |
| SDN CLI password | | The password is specified by the <code>sdnc_admin_password</code> parameter in the <code>sdn</code> section of <code>config.yaml</code> . ⁽¹⁾ |

(1) For more information about password change, refer to the Cloud SDN Hardening Guideline, Reference [1].

1.2.4 Conditions

Before starting this procedure, ensure that the following conditions are met:

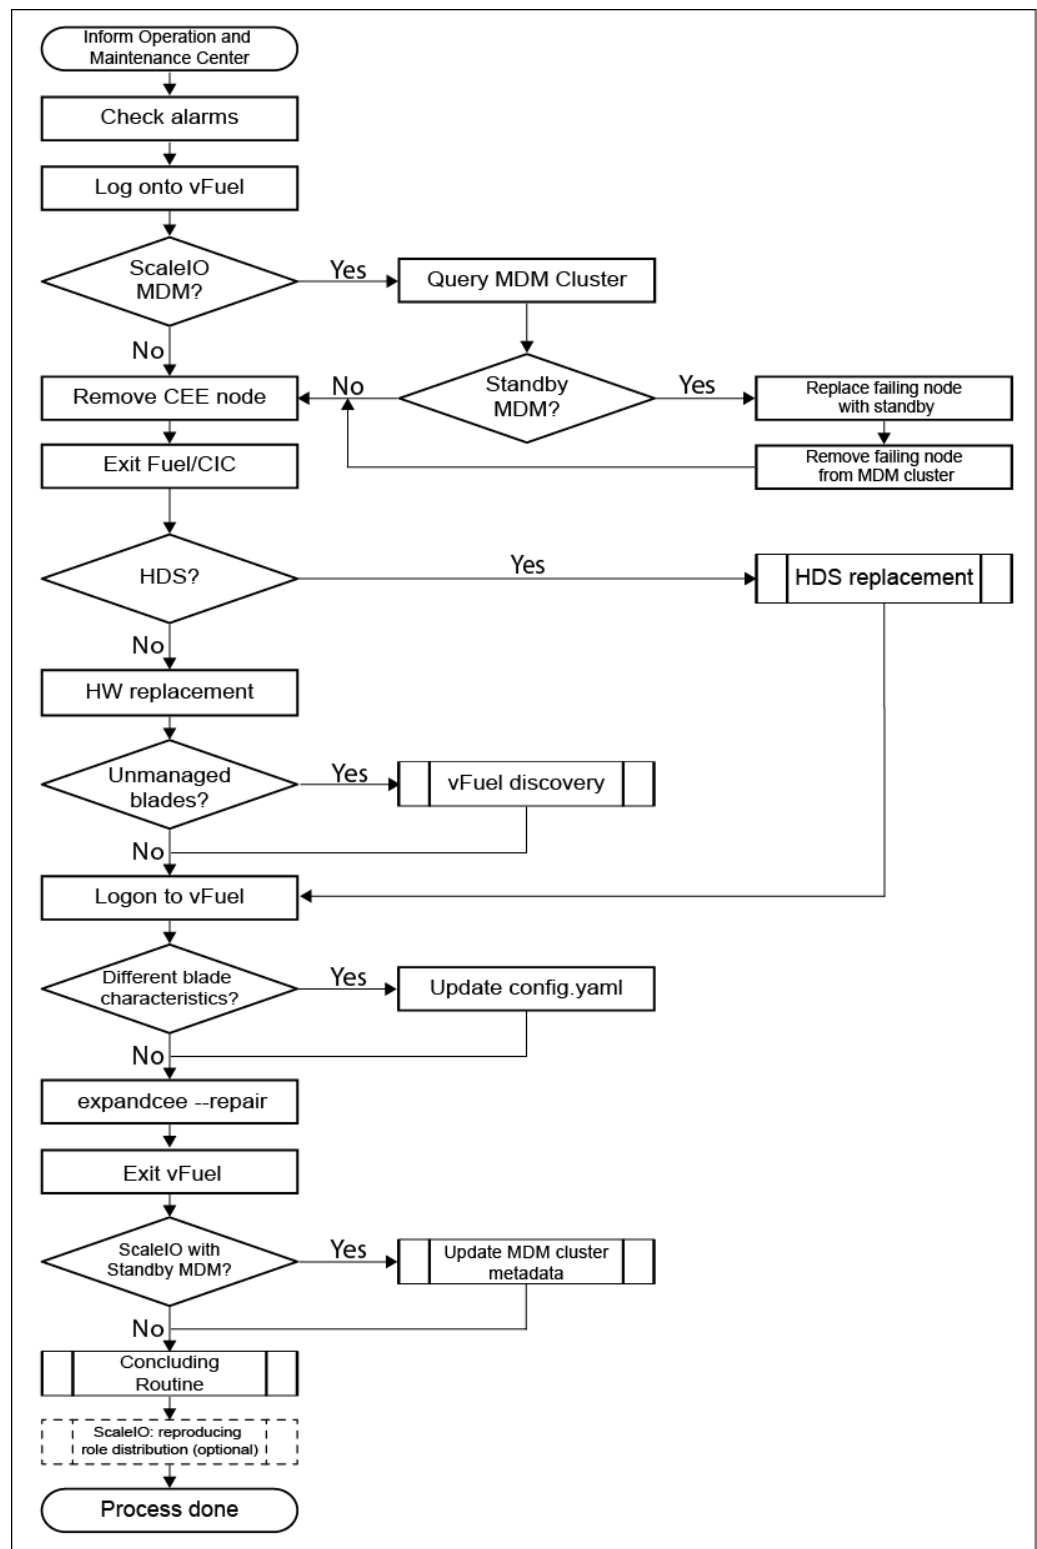
- A work order for the replacement is received or the document is referred from another procedure.
- A compute host failed alarm is active or an active distributed storage alarm points to server replacement, if faulty board is to be replaced.
- The new server is available, and it has been verified visually that it is undamaged.
- The IP addresses and credentials for SSH connections to the following devices are known. See also Section 1.2.3 on page 2.
 - vCIC
 - vFuel
 - ScaleIO MDM master, if ScaleIO MDM cluster member server is to be replaced
- If the `cscadm` password has been changed in ODL since deployment, it must be propagated in the CEE region. For more information, refer to the relevant section in the [Security User Guide](#).
- The data regarding dedicated ScaleIO servers in the configuration files corresponds to actual system parameters.
- Name of the host to be replaced is known. See also Section 1.2.3 on page 2.
- There is no active Fuel Failed alarm.
- All keys to the site are available and site access is granted.



2 Procedure

This procedure describes how to replace a server.

An overview of the process is shown on the following flowchart:



Start the procedure with Section 2.1 on page 5.



2.1 Preparations for Server Removal

This section describes how to prepare for changing a server.

Note: In the case of multiple faulty servers, perform the server replacement procedure at the same time for all the faulty servers.

1. Inform the Operation and Maintenance Center (OMC) that work is in progress on the node with possible disturbance to the service.
2. Check current alarms to have a baseline for the same checks after compute host has been changed.
3. Log on to vFuel using `ssh` with the logon credentials given in the site documentation. For more information, refer to [CEE Connectivity User Guide](#).
4. If the server to be replaced runs ScaleIO MDM service, do the following:

Note: To check if the server to be replaced runs ScaleIO MDM service, refer to the `config.yaml`, or query the ScaleIO MDM cluster.

- a. Log on to any member of the MDM cluster.
- b. Check MDM cluster status by executing the following command:
`scli --mdm_ip <master_mdm_ip> --query_cluster`

An example of the output is the following:

```
root@scaleio-0-2:~# scli --mdm_ip 192.168.2.23, 192.168.2.25, 192.168.2.32 --query_cluster
Cluster:
  Mode: 5_node, State: Degraded, Active: 4/5, Replicas: 2/3
  Virtual IPs: N/A
Master MDM:
  Name: scaleio-0-3, ID: 0x34806ef91dc96031
  IPs: 192.168.17.23, 192.168.18.23, Management IPs: 192.168.2.23, Port: 9011,
  Virtual IP interfaces: N/A
  Version: 2.0.13000
Slave MDMs:
  Name: scaleio-0-1, ID: 0x5658fcdd39418640
  IPs: 192.168.17.25, 192.168.18.25, Management IPs: 192.168.2.25, Port: 9011,
  Virtual IP interfaces: N/A
  Status: Disconnected, Version: 2.0.13000
  Name: scaleio-0-2, ID: 0x188084867f3dfed2
  IPs: 192.168.17.29, 192.168.18.29, Management IPs: 192.168.2.32, Port: 9011,
  Virtual IP interfaces: N/A
  Status: Normal, Version: 2.0.13000
Tie-Breakers:
  Name: scaleio-0-4, ID: 0x2ef3c29e4c1e4374
  IPs: 192.168.17.24, 192.168.18.24, Port: 9011
  Status: Normal, Version: 2.0.13000
  Name: scaleio-0-5, ID: 0x54744e452be7e043
  IPs: 192.168.17.30, 192.168.18.30, Port: 9011
```



```
Status: Normal, Version: 2.0.13000
Standby MDMs:
Name: scaleio-0-7, ID: 0x26ee566356362451, Manager
  IPs: 192.168.17.31, Management IPs: 192.168.2.33, Port: 9011
Name: scaleio-0-6, ID: 0x13c925450656db74, Tie Breaker
  IPs: 192.168.17.32, Port: 9011
```

Make sure that the server to be replaced has Status: Disconnected.

Note: If the original distribution of roles must be reproduced after the server replacement, save the printout of the command for a later step in the procedure.

- c. If no standby MDM is available with the same role as the server to be replaced (Tie-Breaker or Manager), continue with Step 6.
- d. If a Standby MDM is available with the same role as the server to be replaced (Tie-Breaker or Manager), activate the Standby MDM or Tie-Breaker (TB) and remove the failing MDM or TB from the MDM cluster by executing one of the following commands:

— If the node to be replaced is an MDM:
`scli --replace_cluster_mdm --remove_slave_mdm_name =>`
`<failing_mdm_server_name> --add_slave_mdm_name=>`
`<standby_mdm_name>`

Note: If the CLI does not reside on the MDM, the `--mdm_ip <master_mdm_ip>` parameter must be added to the command.

— If the node to be replaced is a TB:
`scli --replace_cluster_mdm --remove_tb_name=>`
`<failing_mdm_server_name> --add_tb_name=>`
`<standby_mdm_name>`

Note: If the CLI does not reside on the MDM, the `--mdm_ip <master_mdm_ip>` parameter must be added to the command.

- e. Remove the replaced node from the system by executing the following command:
`scli --remove_standby_mdm --remove=>`
`mdm_name <node_name>`
 where `<node_name>` corresponds to the name of the node to be replaced.

5. If the CEE is configured with Tight Integrated SDN, remove the transport zone. Do the following:



- a. Log onto any of the vCICs using SSH. For more information, refer to the [CEE CLI Guide](#).
- b. Start the CSC CLI by executing the following command:
`/opt/sdnc/opendaylight/comcli/runCli.sh`
- c. Identify the DPN ID of the server to be replaced by executing the following command:

```
cli>display all-dpns |grep <hostname>
```

Note: The SDN CLI password is required. For more information, see Table 2.

Record the DPN ID of the server to be replaced.

- d. Collect data about the server to be replaced by executing the following command:

```
cli>display tep-show-config |grep <host_dpn_id>
```

where <host_dpn_id> corresponds to the DPN ID of the server to be replaced.

From the printout, record the following data for the server to be replaced:

| Parameter | Corresponds to variable |
|---------------|-------------------------|
| portName | <port_name> |
| vlanid | <vlan_id> |
| ipaddress | <ip_address> |
| cidr | <subnet_mask> |
| gatewayip | <gateway_ip_address> |
| transportZone | <transport_zone> |

- e. Delete the TEP by executing the following commands:

```
cli> exec tep-delete <host_dpn_id> =>  
<port_name> <vlan_id> =>  
<ip_address> <subnet_mask> =>  
<gateway_ip_address> <transport_zone>=>  
cli> exec tep-commit
```

where the variables correspond to the DPN ID of the server to be replaced and the data collected in Step d.

6. Remove the node by issuing the following command on the vFuel node:

```
removeceenode --name <hostname>
```



If the removal fails, stop the process and consult the next level of support. Further actions are outside the scope of this instruction.

7. Log out from vFuel:

exit

8. Log out from vCIC:

exit

9. Continue with Section 2.2 on page 9.

2.2 Hardware Replacement and Server Configuration

Do the following:

1. Depending on the used hardware infrastructure, refer to the relevant instruction indicated below and perform the steps provided for hardware installation of the server including BIOS configuration:
 - BSP-based systems: CPI documentation of the Blade Server Platform (BSP). Use section Replace Device Board in the instruction Manage Blade, Reference [2].
 - Dell-based systems:
 - Dell PowerEdge R630 HW Installation
 - Dell PowerEdge R630 Server Configuration
 - Dell PowerEdge R630 Server BIOS Configuration for CEE
 - HDS-based systems: see Section 2.2.1 on page 9
 - Other hardware: documentation provided by the manufacturer
2. In case of unmanaged servers, continue with Section 2.2.2 on page 9. Else, continue with Section 2.3 on page 10.

2.2.1 Replacement of HDS-Based System

Do the following:

1. Order the replacement of the faulty server from the Data Center owner.
2. Make sure that the new server is assigned to the CEE vPOD, and that the UUID of the new server is available.
3. Continue with Section 2.3 on page 10.



2.2.2 vFuel Discovery for Unmanaged Servers

In the case of unmanaged servers, before the execution of the installation command, the new server needs to be discovered in vFuel. After the new server has been configured, do the following:

1. Set the boot device to PXE.
2. Force restart the server.
3. Log on to vFuel using `ssh`. For more information, refer to [CEE Connectivity User Guide](#).
4. Check from vFuel if the new server has been discovered using the command `fue1 node`. Depending on the result, do one of the following:
 - If the new server is discovered in vFuel, continue with Section 2.3 on page 10.
 - If the new server is not discovered by vFuel after 10 minutes, repeat Step 1 and Step 2.
 - If the new server is not discovered by vFuel after multiple attempts, continue with Section 2.4 on page 12.

2.3 Executing Installation Command

Do the following:

1. Log on to vFuel using `ssh`. For more information, refer to [CEE Connectivity User Guide](#).
2. If the replacement server has different characteristics from those of the old server (for example, if GEP5 blade is exchanged to GEP7 or GEP7L), update the server parameter in the `/mnt/cee_config/config.yaml` file. For more information, refer to the [platform-related System Dimensioning Guide](#) and the [Configuration File Guide](#).

Note: In the case of HDS system, the `<server.uuid>` of the replaced server must be updated.

In the case of a system with unmanaged servers, `<left.mac.address>` and `<right.mac.address>` of the replaced server must be updated.

Note: CSS CPU reservation mode reconfiguration (changing the `css_mode` parameter in the `config.yaml`) can be performed at this point. For details, refer to the [Runtime Configuration Guide](#)

3. Issue the following command to start the installation:

```
expandcee --repair
```



Note: Verification for the newly replaced server is also performed by the command.

Note: Simultaneous repair of a compute server hosting a vCIC and one or more computes previously added using the region expansion procedure (refer to [Region Expansion](#)) is not possible.

4. Do the relevant action:

— If the installation ends without issues, continue with Step 5.

— If the command execution fails, continue with Section 2.4 on page 12.

Note: On BSP platforms, Nova and Neutron commands in some cases return the following error after the installation:

```
message": "<html><body><h1>504 Gateway Time-out</h1>
The server didn't respond in time
```

In this case, log on to one of the vCICs and restart the Neutron server:

```
crm resource restart neutron-server
```

Verify that the Neutron server successfully restarted:

```
crm resource status neutron-server
```

5. Log out from vFuel:

```
exit
```

6. If the CEE is configured with Tight Integrated SDN, update ARP entries of neighbors by doing the following:

a. Log on to the replaced server using SSH. For more information, refer to the [CEE CLI Guide](#).

b. List the IP address of the vNIC on the `br-sdnc-sbi` interface by executing the following command:

```
ifconfig br-sdnc-sbi | grep 'inet addr' | cut -d: -f2 |⇒
awk '{print $1}'
```

From the printout, record the IP address listed as `inet addr:`.

c. Update the ARP entries of the neighbors by executing the following command:

```
arping -U <vnic_ip_address> -I br-sdnc-sbi -c 10
```

where `<vnic_ip_address>` corresponds to the IP address recorded in substep b.

7. Do the relevant action:



- If the replaced server runs ScaleIO MDM service without Standby MDM, continue with Section 3.1 on page 14.
- Else, continue with Section 2.5 on page 12.

2.4 Data Collection

1. Collect the console printout.
2. Log out from vFuel, if applicable:
`exit`
3. Collect all logs as described in the [Data Collection Guideline](#).
4. Consult the next level of maintenance support. Further actions are outside the scope of this instruction.

2.5 Fuel Synchronization

To synchronize Fuel VM, follow the instructions in [Fuel Synchronization](#).

Continue with Section 2.6 on page 12.

2.6 Concluding Routine

Do the following:

1. Check that there are no additional active alarms. If new active alarms are found, act on them according to the relevant Alarm OPI.
2. Check that there are no alerts created during the repair process. If any such alerts are present, act on them according to the relevant OPI.
3. If the configuration files have been updated, it is recommended to perform a backup of the configuration files for disaster recovery purposes. For more information, refer to [Disaster Recovery](#).
4. If the replaced server is a ScaleIO MDM and the original distribution of roles must be reproduced, continue with Section 3.2 on page 18.
5. If the replaced server is a vCIC host, perform the following procedures:
 - Check Pacemaker (vCIC state and cluster resource state), refer to section [Check Pacemaker - vCIC State and Cluster Resource State in Health Check Procedure](#).
 - Check Nova services, refer to section [Check Nova Services in Health Check Procedure](#).



- Check RabbitMQ cluster status, refer to section [Check RabbitMQ Cluster Status in Health Check Procedure](#).
 - Check OpenStack components, refer to section [Check OpenStack Components in Health Check Procedure](#).
6. If the CEE is configured with tightly integrated SDN, do the following:
 - a. Wait for a few minutes.
 - b. Log on to any of the vCICs using SSH. For more information, refer to the [CEE CLI Guide](#).
 - c. Start the CSC CLI by executing the following command:
`/opt/sdnc/opendaylight/comcli/runCli.sh`
 - d. Check DNP status and Transport Zone for the replaced server by executing the following command:
`display tep-show-config`
 7. Log out from the vCIC.

`exit`
 8. If the CEE is configured with tightly integrated SDN, do the following:
 - a. Log on to the replaced server using SSH. For more information, refer to the [CEE CLI Guide](#).
 - b. Verify that the VXLAN tunnels are created by executing the following command:
`ovs-vsctl show`
 - c. Log out of the server.

`exit`
 9. Collect all tools and equipment.
 10. Report that the server has been replaced.
 11. Handle the removed unit according to company procedures regarding repair and data security.

Note: Sensitive data may be present on the disk.
 12. Do any remaining actions according to the work order, if applicable.
 13. The job is completed.



3 ScaleIO MDM Cluster Member Replacement

Note: If the CLI does not reside on the MDM, the `--mdm_ip <master_mdm_ip>` parameter must be added to every CLI command described in this section.

3.1 Procedure for ScaleIO MDM Replacement without Standby MDM

This section describes additional steps that must be executed if no standby MDM is available.

Do the following when directed here from Section 2.3 on page 10:

1. Check MDM cluster status by executing the following command:
`scli --mdm_ip <master_mdm_ip> --query_cluster`

An example of the printout is the following:

```
root@scaleio-0-2:~# scli --mdm_ip 192.168.2.23, 192.168.2.25
, 192.168.2.32 --query_cluster
Cluster:
  Mode: 5_node, State: Degraded, Active: 4/5, Replicas: 2/3
  Virtual IPs: N/A
Master MDM:
  Name: scaleio-0-3, ID: 0x34806ef91dc96031
  IPs: 192.168.17.23, 192.168.18.23, Management IPs: 192.168.2.23, Port: 9011,
  Virtual IP interfaces: N/A
  Version: 2.0.13000
Slave MDMs:
  Name: scaleio-0-1, ID: 0x5658fcdd39418640
  IPs: 192.168.17.25, 192.168.18.25, Management IPs: 192.168.2.25, Port: 9011,
  Virtual IP interfaces: N/A
  Status: Disconnected, Version: 2.0.13000
  Name: scaleio-0-2, ID: 0x188084867f3dfed2
  IPs: 192.168.17.29, 192.168.18.29, Management IPs: 192.168.2.32, Port: 9011,
  Virtual IP interfaces: N/A
  Status: Normal, Version: 2.0.13000
Tie-Breakers:
  Name: scaleio-0-4, ID: 0x2ef3c29e4c1e4374
  IPs: 192.168.17.24, 192.168.18.24, Port: 9011
  Status: Normal, Version: 2.0.13000
  Name: scaleio-0-5, ID: 0x54744e452be7e043
  IPs: 192.168.17.30, 192.168.18.30, Port: 9011
  Status: Normal, Version: 2.0.13000
```

The replaced server is still listed in the printout. The expected status of the replaced MDM is Disconnected.

Note: From the printout, record the name and IP addresses of the replaced server and an MDM or a TB, based on the following:

If the replaced node is an MDM, select a **TB**.
If the replaced node is a TB, select a **Slave MDM**.

This data is required later during the process.



2. Update the MDM cluster metadata with the new UUID of the replaced server by manually removing the node and adding it again to the cluster. Do the following:
 - a. Switch the MDM cluster to three-node cluster by removing the replaced server and an MDM or TB, according to the data recorded in Step 1. Execute the following command:


```
scli --switch_cluster_mode --cluster_mode⇒
3_node--remove_tb_name <tb_name>⇒
--remove_slave_mdm_name <mdm_name>
```
 - b. Check MDM cluster status by executing the following command:


```
scli --mdm_ip <master_mdm_ip> --query_cluster
```

An example of the printout is the following:

```
root@scaleio-0-2:~# scli --mdm_ip 192.168.2.23, 192.168.2.25,
192.168.2.32 --query_cluster
Cluster:
  Mode: 3_node, State: Normal, Active: 3/3, Replicas: 2/2
  Virtual IPs: N/A
Master MDM:
  Name: scaleio-0-3, ID: 0x34806ef91dc96031
  IPs: 192.168.17.23, 192.168.18.23, Management IPs: 192.168.2.23, Port: 9011,
  Virtual IP interfaces: N/A
  Version: 2.0.13000
Slave MDMs:
  Name: scaleio-0-2, ID: 0x188084867f3dfed2
  IPs: 192.168.17.29, 192.168.18.29, Management IPs: 192.168.2.32, Port: 9011,
  Virtual IP interfaces: N/A
  Status: Normal, Version: 2.0.13000
Tie-Breakers:
  Name: scaleio-0-4, ID: 0x2ef3c29e4c1e4374
  IPs: 192.168.17.24, 192.168.18.24, Port: 9011
  Status: Normal, Version: 2.0.13000
Standby MDMs:
  Name: scaleio-0-1, ID: 0x5658fcdd39418640, Manager
  IPs: 192.168.17.25, 192.168.18.25, Management IPs: 192.168.2.25, Port: 9011
  Name: scaleio-0-5, ID: 0x54744e452be7e043, Tie Breaker
  IPs: 192.168.17.30, 192.168.18.30, Port: 9011
```

- c. Remove the replaced node from the system by executing the following command:


```
scli --remove_standby_mdm --remove_⇒
mdm_name <node_name>
```

 where <node_name> corresponds to the name of the replaced node from the data recorded in Step 1.
 - d. Check MDM cluster status by executing the following command:


```
scli --mdm_ip <master_mdm_ip> --query_cluster
```

An example of the printout is the following:



```

root@scaleio-0-2:~# scli --mdm_ip 192.168.21.21,192.168.20.23,19
2.168.20.24 --query_cluster
Cluster:
  Mode: 3_node, State: Normal, Active: 3/3, Replicas: 2/2
  Virtual IPs: N/A
Master MDM:
  Name: scaleio-0-3, ID: 0x34806ef91dc96031
  IPs: 192.168.17.23, 192.168.18.23, Management IPs: 192.168.2.23, Port: 9011, Virtual
  N/A
  Version: 2.0.13000
Slave MDMs:
  Name: scaleio-0-2, ID: 0x188084867f3dfed2
  IPs: 192.168.17.29, 192.168.18.29, Management IPs: 192.168.2.32, Port: 9011, Virtual
  N/A
  Status: Normal, Version: 2.0.13000
Tie-Breakers:
  Name: scaleio-0-4, ID: 0x2ef3c29e4c1e4374
  IPs: 192.168.17.24, 192.168.18.24, Port: 9011
  Status: Normal, Version: 2.0.13000
Standby MDMs:
  Name: scaleio-0-5, ID: 0x54744e452be7e043, Tie Breaker
  IPs: 192.168.17.30, 192.168.18.30, Port: 9011

```

- e. Add the replaced node back to the system by executing the following command:

```

scli --add_standby_mdm --new_mdm_ip⇒
  <mdm_ip> --mdm_role <role> --new_mdm⇒
management_ip <mdm_management_ip> ⇒
--new_mdm_name <mdm_name> --approve_certificate

```

Use the following values for variables:

- <mdm_ip> corresponds to the IP address or addresses of the removed node, according to the data recorded in Step 1. If the node has multiple addresses, provide all addresses, separated by a comma “,”.
- <mdm_management_ip> corresponds to the management IP address of the removed node, according to the data recorded in Step 1.
- <role> is manager or tie-breaker, according to the role of the removed node.

An example of the command is the following:

```

scli --add_standby_mdm --new_mdm⇒
_ip 192.168.21.23 --mdm_role manager ⇒
--new_mdm_management_ip 192.168.20.23 ⇒
--new_mdm_name scaleio-0-1 --approve_certificate

```



- f. Confirm that the new node has received an ID by executing the following command:

```
scli --mdm_ip <master_mdm_ip> --query_cluster
```

An example of the printout is the following:

```
root@scaleio-0-2:~# scli --mdm_ip 192.168.21.21,192.168.20.23,19
2.168.20.24 --query_cluster
Cluster:
  Mode: 3_node, State: Normal, Active: 3/3, Replicas: 2/2
  Virtual IPs: N/A
Master MDM:
  scaleio-0-3, ID: 0x34806ef91dc96031
  IPs: 192.168.17.23, 192.168.18.23, Management IPs: 192.168.2.23, Port: 9011,
  Virtual IP interfaces: N/A
  Version: 2.0.13000
Slave MDMs:
  scaleio-0-2, ID: 0x188084867f3dfed2
  IPs: 192.168.17.29, 192.168.18.29, Management IPs: 192.168.2.32, Port: 9011,
  Virtual IP interfaces: N/A
  Status: Normal, Version: 2.0.13000
Tie-Breakers:
  Name: scaleio-0-4, ID: 0x2ef3c29e4c1e4374
  IPs: 192.168.17.24, 192.168.18.24, Port: 9011
  Status: Normal, Version: 2.0.13000
Standby MDMs:
  Name: scaleio-0-1, ID: 0x5658fcdd39418640 ,Manager
  IPs: 192.168.17.25, 192.168.18.25, Management IPs: 192.168.2.25, Port: 9011
  Name: scaleio-0-5, ID: 0x54744e452be7e043, Tie Breaker
  IPs: 192.168.17.30, 192.168.18.30, Port: 9011
```

In the printout, the replaced server is now listed with the UUID of the new server as a standby MDM.

- g. Switch MDM cluster back to five-node mode by activating the Standby MDM and the TB. Execute the following command:

```
scli --switch_cluster_mode =>
--cluster_mode 5_node --add_slave_mdm_name =>
<mdm_name> --add_tb_name <tb_name>
```

where <mdm_name> and <tb_name> correspond to the following:

- The name of the node added in Step e
- Data recorded in Step 1

An example of the command is the following:

```
scli --switch_cluster_node --cluster_mode=>
5_node --add_slave_mdm_name scaleio-0-1=>
--add_tb_name scaleio-0-5
```



3. Continue with Section 2.5 on page 12.

Note: After replacement, the new server is assigned Slave MDM role. Additional manual MDM cluster management can be required. For more information, refer to Section 3.2 on page 18.

3.2 Reproducing Original Role Distribution

As the failure and the replacement procedure can result in changes in MDM cluster member roles, the following procedures can be required:

- Replacing Slave MDM or TB with standby MDM
- Switching Master MDM role

Replacing Members of MDM Cluster

If the original distribution of Slave MDM or TB roles must be reproduced in the MDM Cluster, execute one of the following commands:

- To replace a Slave MDM with a Standby MDM:

```
scli --replace_cluster_mdm --remove_slave_mdm_name=>  
  <node_a_name> --add_slave_mdm_name =>  
  <node_b_name>
```
- To replace a TB with a Standby MDM:

```
scli --replace_cluster_mdm --remove_tb_name =>  
  <node_a_name> --add_tb_name =>  
  <node_b_name>
```

Use the following values for the variables:

- <node_a_name> corresponds to the name of the node that took over the role of the failed node before server replacement. Use the name of the standby MDM from Step d in Section 2.1 on page 5.
- <node_b_name> corresponds to the name of the replaced node from the data recorded in Step 1 Section 3.1 on page 14.

For more information, refer to the [Dell EMC ScaleIO Version 2.x CLI Reference Guide](#).

Switching MDM Ownership

Note: A prerequisite to this procedure is that the node intended to take on the Master MDM role is already a slave MDM.

If the failed node was the master MDM, and it required that the new server assumes the Master MDM role after replacement, execute the following command:

where <mdm_name> corresponds to the name of the replaced server according to the data saved in Step b in Section 2.1 on page 5.



For more information, refer to the Dell EMC ScaleIO Version 2.x CLI Reference Guide.



Reference List

- [1] Cloud SDN Hardening Guideline, 2/1553-AXD 101 08/6-V1
- [2] Manage Blade, 53/1543-APR 901 0549/1
- [3] Limitations and Workarounds for Cloud Execution Environment (CEE) 6.6, 5/109 21-AZE 102 01/5-12 Uen