

# Troubleshoot CNDP Cluster Manager HA Node Issue and Reinstall Force-VM Redeploy

## Contents

[Introduction](#)

[Prerequisites](#)

[Requirements](#)

[Components Used](#)

[Background Information](#)

[What is SMI Cluster Manager?](#)

[What is an Inception Server?](#)

[Problem](#)

[Procedure for the Maintenance](#)

[Identify Hosts](#)

[Identify Cluster Details from Inception Server](#)

[Remove the Virtual Drive to Clear the Operating System from the Server](#)

[Run Cluster Sync](#)

[Monitor the Cluster-sync Sync Logs](#)

[Verification](#)

## Introduction

This document describes the procedure to recover Cluster Manager from the inception server in the Cloud Native Deployment Platform (CNDP) setup.

## Prerequisites

### Requirements

Cisco recommends that you have knowledge of these topics:

- Cisco Subscriber Microservices Infrastructure (SMI)
- 5G CNDP or SMI-Bare-metal (BM) architecture
- Distributed Replicated Block Device (DRBD)

### Components Used

The information in this document is based on these software and hardware versions:

- SMI 2020.02.2.35
- Kubernetes v1.21.0

The information in this document was created from the devices in a specific lab environment. All of the devices used in this document started with a cleared (default) configuration. If your network is live, ensure that you understand the potential impact of any command.

# Background Information

## What is SMI Cluster Manager?

A cluster manager is a 2-node keepalived cluster used as the initial point for both control plane and user plane cluster deployment. It runs a single-node Kubernetes cluster and a set of PODs which are responsible for the entire cluster setup. Only the primary cluster manager is active and the secondary takes over only in case of a failure or brought down manually for maintenance.

## What is an Inception Server?

This node performs lifecycle management of the Cluster Manager (CM) that underlies and from here you can push Day0 Config.

This server is usually deployed region-wise or in the same data centre as the top-level orchestration function (for example NSO) and typically runs as a VM.

## Problem

The cluster manager is hosted in a 2-node cluster with Distributed Replicated Block Device (DRBD) and keepalived as Cluster Manager primary and Cluster Manager secondary. In this case, Cluster Manager secondary goes to power off state automatically while initialization/installation of OS in UCS, which indicates OS is corrupt.

```
cloud-user@POD-NAME-cm-primary:~$ drbd-overview status
0:data/0 WfConnection Primary/Unknown UpToDate/DUnknown /mnt/stateful_partition ext4 568G 369G 170G 69%
```

## Procedure for the Maintenance

This process helps to reinstall the OS on the CM server.

### Identify Hosts

Login to Cluster-Manager and identify hosts:

```
cloud-user@POD-NAME-cm-primary:~$ cat /etc/hosts | grep 'deployer-cm'
127.X.X.X POD-NAME-cm-primary POD-NAME-cm-primary
X.X.X.X POD-NAME-cm-primary
X.X.X.Y POD-NAME-cm-secondary
```

### Identify Cluster Details from Inception Server

Login to the Inception server and get into Deployer and verify the cluster name with hosts-IP from Cluster-Manager.

After successful login to the inception server, log in to the ops centre as shown here.

```
<#root>
```

```
user@
```

```
inception-server
```

```
: ~$
```

```
ssh -p 2022 admin@localhost
```

Verify Cluster Name from Cluster Manager SSH-IP (ssh-ip = Node SSH IP ip-address = ucs-server cimc ip-address).

```
<#root>
```

```
[
```

```
inception-server
```

```
] SMI Cluster Deployer#
```

```
show running-config clusters * nodes * k8s ssh-ip | select nodes * ssh-ip | select nodes * ucs-server cimc
```

NAME	NAME	SSH IP	SSH IP	IP ADDRESS
------	------	-----------	--------	------------

-----  
POD-NAME-deployer

```
cm-primary
```

```
-
```

```
X.X.X.X
```

10.X.X.X ---> Verify Name and SSH IP if Cluster is part of inception server SMI.

```
cm-secondary
```

```
-
```

```
X.X.X.Y
```

10.X.X.Y

Check the configuration for the target cluster.

```
<#root>
```

```
[inception-server]
```

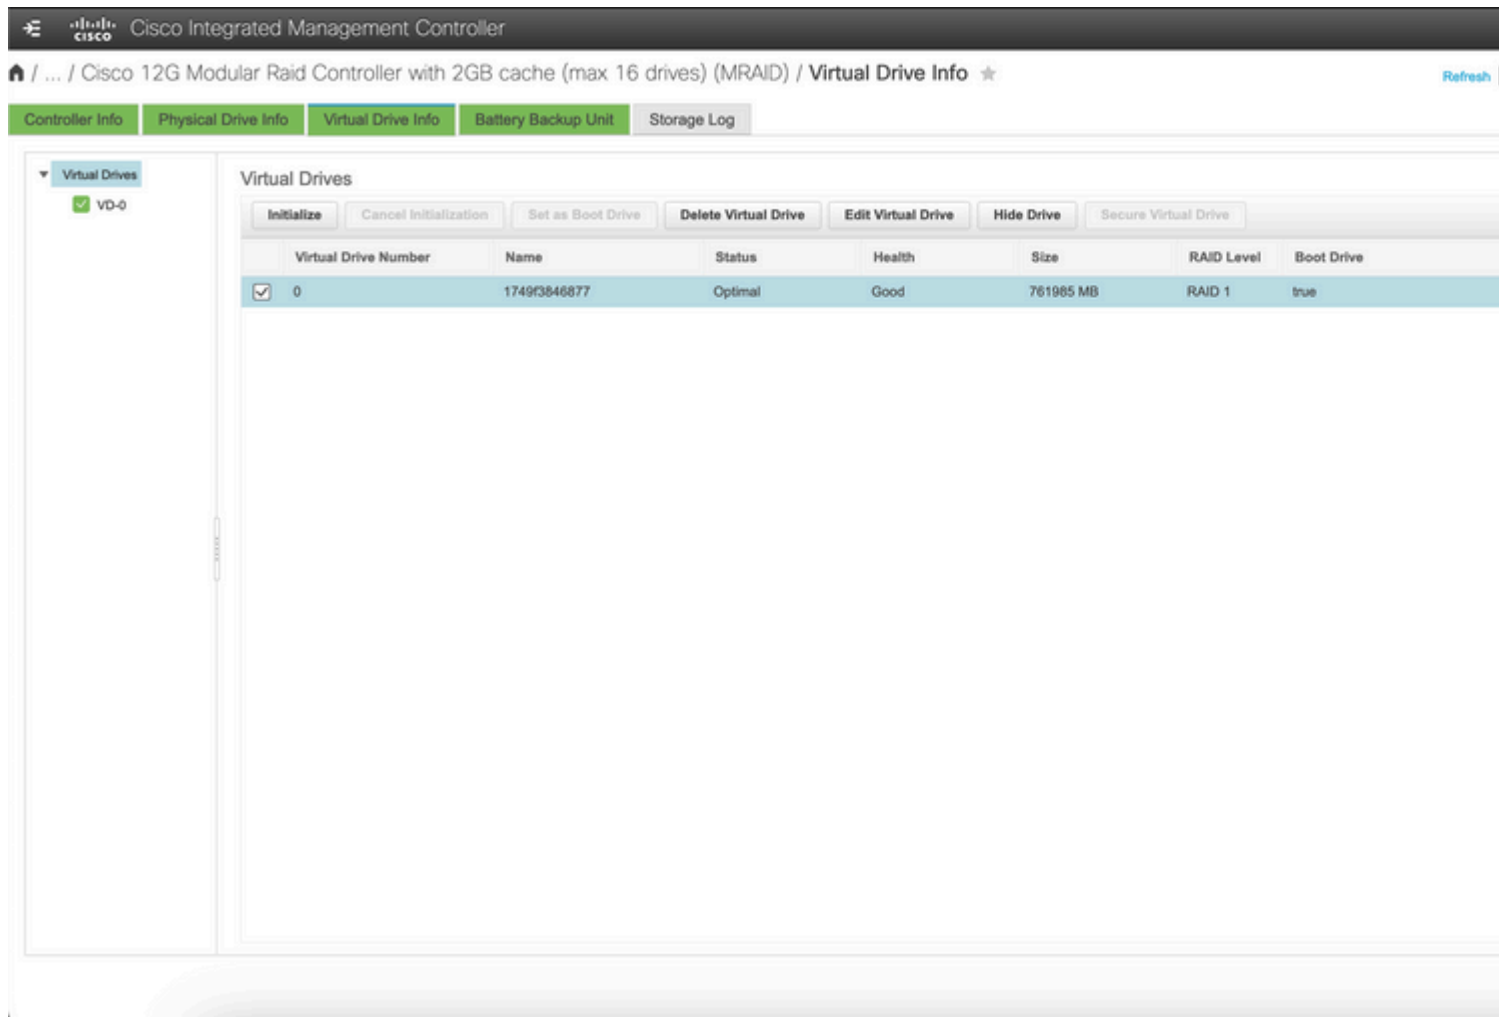
```
SMI Cluster Deployer#
```

```
show running-config clusters POD-NAME-deployer
```

## Remove the Virtual Drive to Clear the Operating System from the Server

Connect to the CIMC of the affected host and clear the boot drive and delete the virtual drive (VD).

- a) CIMC > Storage > Cisco 12G Modular Raid Controller > Storage Log > Clear Boot Drive
- b) CIMC > Storage > Cisco 12G Modular Raid Controller > Virtual drive > Select the virtual drive > Delete



## Run Cluster Sync

Run default cluster sync for Cluster-Manager from the inception server.

```
<#root>
```

```
[inception-server] SMI Cluster Deployer# clusters POD-NAME-deployer actions sync run debug true
```

```
This will run sync. Are you sure? [no,yes]
```

```
yes
```

```
message accepted
```

```
[inception-server] SMI Cluster Deployer#
```

If default cluster-sync fails, perform cluster-sync with force-vm redeploy option for complete re-install (Cluster-sync activity can take ~45-55 mins to complete, it depends on the number of nodes hosted on the cluster)

```
<#root>
```

```
[  
i  
[inception-server] SMI Cluster Deployer  
#  
clusters POD-NAME-deployer actions sync run debug true force-vm-redeploy true
```

```
This will run sync. Are you sure? [no,yes]
```

```
yes
```

```
message accepted
```

```
[inception-server] SMI Cluster Deployer#
```

## Monitor the Cluster-sync Sync Logs

```
<#root>
```

```
[  
inception-server  
] SMI Cluster Deployer#  
monitor sync-logs POD-NAME-deployer
```

```
2023-02-23 10:15:07.548 DEBUG cluster_sync.POD-NAME: Cluster name: POD-NAME  
2023-02-23 10:15:07.548 DEBUG cluster_sync.POD-NAME:
```

```
Force VM Redeploy: true
```

```
2023-02-23 10:15:07.549 DEBUG cluster_sync.POD-NAME: Force partition Redeploy: false  
2023-02-23 10:15:07.549 DEBUG cluster_sync.POD-NAME: reset_k8s_nodes: false  
2023-02-23 10:15:07.549 DEBUG cluster_sync.POD-NAME: purge_data_disks: false  
2023-02-23 10:15:07.549 DEBUG cluster_sync.POD-NAME: upgrade_strategy: auto  
2023-02-23 10:15:07.549 DEBUG cluster_sync.POD-NAME: sync_phase: all  
2023-02-23 10:15:07.549 DEBUG cluster_sync.POD-NAME: debug: true
```

```
...  
...  
...
```

The server is re-provisioned and installed by successful cluster-sync.

```
<#root>
```

```
PLAY RECAP *****
```

```
cm-primary : ok=535 changed=250 unreachable=0
```

```
failed=0
```

```
skipped=832 rescued=0 ignored=0
```

```
cm-secondary : ok=299 changed=166 unreachable=0
```

```
failed=0
```

```
skipped=627 rescued=0 ignored=0
```

```
localhost : ok=59 changed=8 unreachable=0
```

```
failed=0
```

```
skipped=18 rescued=0 ignored=0
```

```
Thursday 23 February 2023 13:17:24 +0000 (0:00:00.109)
```

```
0:56:20.544
```

```
****. ---> ~56 mins to complete cluster sync
```

```
=====
```

```
2023-02-23 13:17:24.539 DEBUG cluster_sync.POD-NAME:
```

```
Cluster sync successful
```

```
2023-02-23 13:17:24.546 DEBUG cluster_sync.POD-NAME: Ansible sync done
```

```
2023-02-23 13:17:24.546 INFO cluster_sync.POD-NAME: _sync finished. Opening lock
```

## Verification

Check affected Cluster Manager is reachable and DRBD overview of the Primary and Secondary Cluster Managers is in UpToDate status.

```
cloud-user@POD-NAME-cm-primary:~$ ping X.X.X.Y
```

```
PING X.X.X.Y (X.X.X.Y) 56(84) bytes of data.
```

```
64 bytes from X.X.X.Y: icmp_seq=1 ttl=64 time=0.221 ms
```

```
64 bytes from X.X.X.Y: icmp_seq=2 ttl=64 time=0.165 ms
```

```
64 bytes from X.X.X.Y: icmp_seq=3 ttl=64 time=0.151 ms
```

```
64 bytes from X.X.X.Y: icmp_seq=4 ttl=64 time=0.154 ms
```

```
64 bytes from X.X.X.Y: icmp_seq=5 ttl=64 time=0.172 ms
```

```
64 bytes from X.X.X.Y: icmp_seq=6 ttl=64 time=0.165 ms
```

```
64 bytes from X.X.X.Y: icmp_seq=7 ttl=64 time=0.174 ms
```

```
--- X.X.X.Y ping statistics ---
```

```
7 packets transmitted, 7 received, 0% packet loss, time 6150ms
```

```
rtt min/avg/max/mdev = 0.151/0.171/0.221/0.026 ms
```

```
cloud-user@POD-NAME-cm-primary:~$ drbd-overview status
0:data/0 Connected Primary/Secondary UpToDate/UpToDate /mnt/stateful_partition ext4 568G 17G 523G 4%
```

The affected cluster manager is installed and re-provisioned to the network successfully.