



Troubleshooting

- [Displaying Cisco NFVI Node Names and IP Addresses, page 1](#)
- [Verifying Cisco NFVI Node Interface Configurations, page 2](#)
- [Displaying Cisco NFVI Node Network Configuration Files, page 4](#)
- [Viewing Cisco NFVI Node Interface Bond Configuration Files, page 4](#)
- [Viewing Cisco NFVI Node Route Information, page 5](#)
- [Viewing Linux Network Namespace Route Information, page 6](#)
- [Prior to Remove Storage Operation, page 6](#)
- [Troubleshooting Cisco NFVI, page 7](#)
- [Management Node Recovery Scenarios, page 9](#)
- [Recovering Compute Node Scenario, page 17](#)
- [Running the Cisco VIM Technical Support Tool, page 19](#)
- [Tech-support configuration file, page 20](#)
- [Tech-support when servers are offline, page 23](#)
- [Disk-maintenance tool to manage physical drives, page 23](#)
- [OSD-maintenance tool, page 28](#)
- [Utility to resolve Cisco VIM Hardware Validation Failures, page 32](#)

Displaying Cisco NFVI Node Names and IP Addresses

Complete the following steps to display the Cisco NFVI node names and IP addresses.

-
- Step 1** Log into the Cisco NFVI build node.
- Step 2** Display the `openstack-configs/mercury_servers_info` file.

The info files displays node name and address similar to the following:

```
# more openstack-configs/mercury_servers_info
Total nodes: 5
Controller nodes: 3
+-----+-----+-----+-----+-----+
| Server          | CIMC          | Management    | Provision     | Tenant        | Storage       |
+-----+-----+-----+-----+-----+
| test-c-control-1 | 10.10.223.13 | 10.11.223.22 | 10.11.223.22 | 169.254.133.102 | None         |
|                 |               |               |               |                 |               |
| test-c-control-3 | 10.10.223.9  | 10.11.223.23 | 10.11.223.23 | 169.254.133.103 | None         |
|                 |               |               |               |                 |               |
| test-c-control-2 | 10.10.223.10 | 10.11.223.24 | 10.11.223.24 | 169.254.133.104 | None         |
|                 |               |               |               |                 |               |
+-----+-----+-----+-----+-----+
Compute nodes: 2
+-----+-----+-----+-----+-----+
| Server          | CIMC          | Management    | Provision     | Tenant        | Storage       |
+-----+-----+-----+-----+-----+
| test-c-compute-1 | 10.10.223.11 | 10.11.223.25 | 10.11.223.25 | 169.254.133.105 | None         |
|                 |               |               |               |                 |               |
| test-c-compute-2 | 10.10.223.12 | 10.11.223.26 | 10.11.223.26 | 169.254.133.106 | None         |
|                 |               |               |               |                 |               |
+-----+-----+-----+-----+-----+
```

Note During the Cisco NFVI deployment, SSH public keys for each node are added to `.../.ssh/authorized_keys`, so you should be able to log in from the build node into each of the Cisco NFVI nodes without passwords. If, for some reason you do need account information, see the `openstack-configs/secrets.yaml` file on the build node.

Verifying Cisco NFVI Node Interface Configurations

Complete the following steps to verify the interface configurations of Cisco NFVI nodes:

Step 1 SSH into the target node, for example, one of the Cisco VIM controllers:

```
[root@ j11-build-1~]# ssh root@j11-control-server-1
[root@j11-control-server-1 ~]#
```

Step 2 Enter the `ip a` command to get a list of all interfaces on the node:

```
[root@j11-control-server-1 ~]# ip a
1: lo: <LOOPBACK,UP,LOWER_UP> mtu 65536 qdisc noqueue state UNKNOWN
    link/loopback 00:00:00:00:00:00 brd 00:00:00:00:00:00
    inet 127.0.0.1/8 scope host lo
        valid_lft forever preferred_lft forever
2: enp8s0: <BROADCAST,MULTICAST,UP,LOWER_UP> mtu 1500 qdisc mq state UP qlen 1000
    link/ether 54:a2:74:7d:42:1d brd ff:ff:ff:ff:ff:ff
3: enp9s0: <BROADCAST,MULTICAST,UP,LOWER_UP> mtu 1500 qdisc mq state UP qlen 1000
    link/ether 54:a2:74:7d:42:1e brd ff:ff:ff:ff:ff:ff
4: mx0: <BROADCAST,MULTICAST,SLAVE,UP,LOWER_UP> mtu 1500 qdisc mq master mx state UP qlen 1000
    link/ether 54:a2:74:7d:42:21 brd ff:ff:ff:ff:ff:ff
5: mx1: <BROADCAST,MULTICAST,SLAVE,UP,LOWER_UP> mtu 1500 qdisc mq master mx state UP qlen 1000
```

```

link/ether 54:a2:74:7d:42:21 brd ff:ff:ff:ff:ff:ff
6: t0: <BROADCAST,MULTICAST,SLAVE,UP,LOWER_UP> mtu 1500 qdisc mq master t state UP qlen 1000
link/ether 54:a2:74:7d:42:23 brd ff:ff:ff:ff:ff:ff
7: t1: <BROADCAST,MULTICAST,SLAVE,UP,LOWER_UP> mtu 1500 qdisc mq master t state UP qlen 1000
link/ether 54:a2:74:7d:42:23 brd ff:ff:ff:ff:ff:ff
8: e0: <BROADCAST,MULTICAST,SLAVE,UP,LOWER_UP> mtu 1500 qdisc mq master e state UP qlen 1000
link/ether 54:a2:74:7d:42:25 brd ff:ff:ff:ff:ff:ff
9: e1: <BROADCAST,MULTICAST,SLAVE,UP,LOWER_UP> mtu 1500 qdisc mq master e state UP qlen 1000
link/ether 54:a2:74:7d:42:25 brd ff:ff:ff:ff:ff:ff
10: p0: <BROADCAST,MULTICAST,SLAVE,UP,LOWER_UP> mtu 1500 qdisc mq master p state UP qlen 1000
link/ether 54:a2:74:7d:42:27 brd ff:ff:ff:ff:ff:ff
11: p1: <BROADCAST,MULTICAST,SLAVE,UP,LOWER_UP> mtu 1500 qdisc mq master p state UP qlen 1000
link/ether 54:a2:74:7d:42:27 brd ff:ff:ff:ff:ff:ff
12: a0: <BROADCAST,MULTICAST,SLAVE,UP,LOWER_UP> mtu 1500 qdisc mq master a state UP qlen 1000
link/ether 54:a2:74:7d:42:29 brd ff:ff:ff:ff:ff:ff
13: a1: <BROADCAST,MULTICAST,SLAVE,UP,LOWER_UP> mtu 1500 qdisc mq master a state UP qlen 1000
link/ether 54:a2:74:7d:42:29 brd ff:ff:ff:ff:ff:ff
14: bond0: <BROADCAST,MULTICAST,MASTER> mtu 1500 qdisc noop state DOWN
link/ether 4a:2e:2a:9e:01:d1 brd ff:ff:ff:ff:ff:ff
15: a: <BROADCAST,MULTICAST,MASTER,UP,LOWER_UP> mtu 1500 qdisc noqueue master br_api state UP
link/ether 54:a2:74:7d:42:29 brd ff:ff:ff:ff:ff:ff
16: br_api: <BROADCAST,MULTICAST,UP,LOWER_UP> mtu 1500 qdisc noqueue state UP
link/ether 54:a2:74:7d:42:29 brd ff:ff:ff:ff:ff:ff
17: e: <BROADCAST,MULTICAST,MASTER,UP,LOWER_UP> mtu 1500 qdisc noqueue state UP
link/ether 54:a2:74:7d:42:25 brd ff:ff:ff:ff:ff:ff
18: mx: <BROADCAST,MULTICAST,MASTER,UP,LOWER_UP> mtu 1500 qdisc noqueue master br_mgmt state UP
link/ether 54:a2:74:7d:42:21 brd ff:ff:ff:ff:ff:ff
19: br_mgmt: <BROADCAST,MULTICAST,UP,LOWER_UP> mtu 1500 qdisc noqueue state UP
link/ether 54:a2:74:7d:42:21 brd ff:ff:ff:ff:ff:ff
inet 10.23.221.41/28 brd 10.23.221.47 scope global br_mgmt
    valid_lft forever preferred_lft forever
20: p: <BROADCAST,MULTICAST,MASTER,UP,LOWER_UP> mtu 1500 qdisc noqueue state UP
link/ether 54:a2:74:7d:42:27 brd ff:ff:ff:ff:ff:ff
21: t: <BROADCAST,MULTICAST,MASTER,UP,LOWER_UP> mtu 1500 qdisc noqueue state UP
link/ether 54:a2:74:7d:42:23 brd ff:ff:ff:ff:ff:ff
inet 17.16.3.8/24 brd 17.16.3.255 scope global t
    valid_lft forever preferred_lft forever
22: docker0: <NO-CARRIER,BROADCAST,MULTICAST,UP> mtu 1500 qdisc noqueue state DOWN
link/ether 02:42:70:f6:8b:da brd ff:ff:ff:ff:ff:ff
inet 172.17.42.1/16 scope global docker0
    valid_lft forever preferred_lft forever
24: mgmt-out@if23: <BROADCAST,MULTICAST,UP,LOWER_UP> mtu 1500 qdisc pfifo_fast master br_mgmt state
UP qlen 1000
link/ether 5a:73:51:af:e5:e7 brd ff:ff:ff:ff:ff:ff link-netnsid 0
26: api-out@if25: <BROADCAST,MULTICAST,UP,LOWER_UP> mtu 1500 qdisc pfifo_fast master br_api state
UP qlen 1000
link/ether 6a:a6:fd:70:01:f9 brd ff:ff:ff:ff:ff:ff link-netnsid 0

```

Displaying Cisco NFVI Node Network Configuration Files

Complete the following steps to view a Cisco NFVI node network configuration files:

Step 1 SSH into the target node, for example, one of the Cisco VIM controllers:

```
[root@ j11-build-1~]# ssh root@j11-control-server-1
[root@j11-control-server-1 ~]#
```

Step 2 List all of the network configuration files in the /etc/sysconfig/network-scripts directory, for example:

```
[root@j11-control-server-1 ~]# ls /etc/sysconfig/network-scripts/
ifcfg-a          ifcfg-enp15s0  ifcfg-mx0      ifdown-ib      ifup           ifup-ppp
ifcfg-a0         ifcfg-enp16s0  ifcfg-mx1      ifdown-ippv6   ifup-aliases  ifup-routes
ifcfg-a1         ifcfg-enp17s0  ifcfg-p        ifdown-ipv6    ifup-bnep     ifup-sit
ifcfg-br_api    ifcfg-enp18s0  ifcfg-p0       ifdown-isdn    ifup-eth      ifup-Team
ifcfg-br_mgmt   ifcfg-enp19s0  ifcfg-p1       ifdown-post    ifup-ib       ifup-TeamPort
ifcfg-e          ifcfg-enp20s0  ifcfg-t        ifdown-ppp     ifup-ippv6    ifup-tunnel
ifcfg-e0         ifcfg-enp21s0  ifcfg-t0       ifdown-routes  ifup-ipv6     ifup-wireless
ifcfg-e1         ifcfg-enp8s0   ifcfg-t1       ifdown-sit     ifup-isdn     init.ipv6-global
ifcfg-enp12s0   ifcfg-enp9s0   ifdown        ifdown-Team    ifup-plip     network-functions
ifcfg-enp13s0   ifcfg-lo       ifdown-bnep    ifdown-TeamPort ifup-plusb    network-functions-ipv6
ifcfg-enp14s0   ifcfg-mx       ifdown-eth     ifdown-tunnel  ifup-post
```

Viewing Cisco NFVI Node Interface Bond Configuration Files

Complete the following steps to view the Cisco NFVI node interface bond configuration files:

Step 1 SSH into the target node, for example, one of the Cisco VIM controllers:

```
[root@ j11-build-1~]# ssh root@j11-control-server-1
[root@j11-control-server-1 ~]#
```

Step 2 List all of the network bond configuration files in the /proc/net/bonding/ directory:

```
[root@j11-control-server-1 ~]# ls /proc/net/bonding/
a  bond0  e  mx  p  t
```

Step 3 To view more information about a particular bond configuration, enter:

```
[root@j11-control-server-1 ~]# more /proc/net/bonding/a
Ethernet Channel Bonding Driver: v3.7.1 (April 27, 2011)
```

```
Bonding Mode: load balancing (xor)
Transmit Hash Policy: layer3+4 (1)
MII Status: up
MII Polling Interval (ms): 100
Up Delay (ms): 0
Down Delay (ms): 0
```

```
Slave Interface: a0
MII Status: up
```

```

Speed: 10000 Mbps
Duplex: full
Link Failure Count: 1
Permanent HW addr: 54:a2:74:7d:42:29
Slave queue ID: 0

Slave Interface: a1
MII Status: up
Speed: 10000 Mbps
Duplex: full
Link Failure Count: 2
Permanent HW addr: 54:a2:74:7d:42:2a
Slave queue ID: 0
    
```

Viewing Cisco NFVI Node Route Information

Complete the following steps to view Cisco NFVI node route information. Note that this is not the HAProxy container running on the controller. The default gateway should point to the gateway on the management network using the br_mgmt bridge.

Step 1 SSH into the target node, for example, one of the Cisco VIM controllers:

```

[root@ j11-build-1~]# ssh root@j11-control-server-1
[root@j11-control-server-1 ~]#
    
```

Step 2 View the routing table (verify the default gateway) of the Cisco NFVI node:

```

[root@j11-control-server-1 ~]# route -n
Kernel IP routing table
Destination      Gateway          Genmask         Flags Metric Ref    Use Iface
0.0.0.0          10.23.221.33    0.0.0.0         UG    0      0      0 br_mgmt
10.23.221.32     0.0.0.0         255.255.255.240 U    0      0      0 br_mgmt
17.16.3.0        0.0.0.0         255.255.255.0   U    0      0      0 t
169.254.0.0      0.0.0.0         255.255.0.0     U    1016   0      0 br_api
169.254.0.0      0.0.0.0         255.255.0.0     U    1017   0      0 e
169.254.0.0      0.0.0.0         255.255.0.0     U    1019   0      0 br_mgmt
169.254.0.0      0.0.0.0         255.255.0.0     U    1020   0      0 p
169.254.0.0      0.0.0.0         255.255.0.0     U    1021   0      0 t
172.17.0.0       0.0.0.0         255.255.0.0     U    0      0      0 docker0
    
```

Viewing Linux Network Namespace Route Information

Complete the following steps to view the route information of the Linux network namespace that the HAProxy container uses on a Cisco NFVI controller node. The default gateway should point to the gateway on the API network using the API interface in the Linux network namespace.

Step 1 SSH into the target node, for example, one of the Cisco VIM controllers:

```
[root@ j11-build-1~]# ssh root@j11-control-server-1
[root@j11-control-server-1 ~]#
```

Step 2 Enter the `ip netns` command to find the name of the network namespace:

```
[root@j11-control-server-2 ~]# ip netns 17550 (id: 0)
```

Step 3 Enter the `ip netns exec` command to view the routing table (verify the default gateway) of the Linux network namespace:

```
[root@j11-control-server-2 ~]# ip netns exec 17550 route -n
Kernel IP routing table
Destination      Gateway          Genmask          Flags Metric Ref    Use Iface
0.0.0.0          172.29.86.1     0.0.0.0          UG    0      0      0 api
10.23.221.32    0.0.0.0         255.255.255.240 U    0      0      0 mgmt
172.29.86.0     0.0.0.0         255.255.255.0   U    0      0      0 api
```

Prior to Remove Storage Operation

Upon completion of the pod management operations such as `add-storage`, the operator needs to ensure that any subsequent operation such as `remove-storage` on the same storage node is done after accounting for all of the devices and their corresponding OSDs have been marked in the persistent crush map as shown in the output of the `ceph osd crush tree`.

Execute the following command on the storage node where a `remove-storage` pod operation is performed, to get a list of all the devices configured for ceph osds:

```
[root@storage-3 ~]$ df | grep -oh ceph-[0-9]*
[root@storage-3 ~]$ df | grep -oh ceph-[0-9]*
ceph-1
ceph-5
ceph-7
ceph-10
```

Login to any of the controller nodes and run the following commands within the `ceph mon` container:

```
$ cephmon
$ ceph osd crush tree
```

From the json output, locate the storage node to be removed and ensure all of the devices listed for ceph osds have corresponding osd entries for them by running the following commands:

```
{
  "id": -3,
  "name": "storage-3",
  "type": "host",
  "type_id": 1,
  "items": [
    {
      "id": 1,
      "name": "osd.1",
      "type": "osd",
      "type_id": 0,
```

```

        "crush_weight": 1.091095,
        "depth": 2
    },
    {
        "id": 5,
        "name": "osd.5",
        "type": "osd",
        "type_id": 0,
        "crush_weight": 1.091095,
        "depth": 2
    },
    {
        "id": 7,
        "name": "osd.7",
        "type": "osd",
        "type_id": 0,
        "crush_weight": 1.091095,
        "depth": 2
    },
    {
        "id": 10,
        "name": "osd.10",
        "type": "osd",
        "type_id": 0,
        "crush_weight": 1.091095,
        "depth": 2
    }
]
},

```

Troubleshooting Cisco NFVI

The following topics provide Cisco NFVI general troubleshooting procedures.

Container Download Problems

- 1 Check installer logs log file `/var/log/mercury/mercury_buildorchestration.log` for any build node orchestration failures including stuck "registry-Populate local registry". In some cases, the Docker container download from your management node might be slow.
- 2 Check the network connectivity between the management node and the remote registry in `defaults.yaml` on the management node (`grep "^registry:" openstack-configs/defaults.yaml`).
- 3 Verify valid remote registry credentials are defined in `setup_data.yaml` file.
- 4 A proxy server might be needed to pull the container images from remote registry. If a proxy is required, exclude all IP addresses for your setup including management node.

PXE Boot Problems

- 1 Check log file `/var/log/mercury/mercury_baremetal_install.log` and connect to failing node CIMC KVM console to find out more on PXE boot failure reason.
- 2 Ensure all validations (step 1) and (step 3) pass.
- 3 Check the kickstart file used in `setup_data.yaml` for controller, compute, and storage nodes matches with the hardware of corresponding nodes.
- 4 Check L2/L3 network connectivity between the failing node and the management node. Also, check the VPC configuration and port-channel status.
- 5 Check that the actual PXE boot order is not different from the configured boot-order.

- 6 Check that PXE (DHCP/TFTP) packets arrive at the management node by performing tcpdump on management interface and looking for UDP 67 or UDP 69 port.
- 7 Check that HTTP request reaches management node by performing tcpdump on management node management interface on TCP 80 and TCP 443 port. Also check the Docker logs for HTTP request from the failing node: `docker exec -it repo_mirror tail -f /var/log/httpd/access_log`.
- 8 Verify all nodes are running on supported Cisco NFVI CIMC firmware or above.
- 9 Verify supported VIC cards are installed (UCS C-series 1225/1227; UCS B-series: 1240/1280/1340/1380).
- 10 After redeploying when the overcloud nodes (controller or compute) do not get their boot order set correctly even though CIMC says the order is correct. In such instances step 4 of runner fails as PXE boot never completes in time. Manually enter the boot menu from the KVM console of the affected nodes and make sure the correct order (NIC interface) is chosen for the PXE boot. This is likely to clear the BIOS CMOS before PXE booting.
- 11 For Authorization Required error messages during bare metal (Step 4) with CIMC operations like hardware validations or cleaning up vNIC, check whether the maximum allowed simultaneous connection (4) are in use. This happens when the third-party application(s) monitoring CIMC does not properly close CIMC connections thereby using all 4 simultaneous connections and makes CiscoVIM installer to starve and not allow to login using xmlapi.

Cisco IMC Connection Problems during Bare Metal Installation

The cause may be Cisco IMC has too many connections, so the installer cannot connect to it. Clear the connections by logging into your Cisco IMC, going into the Admin->Sessions tab and clearing the connections.

API VIP Connection Problems

Verify the active HAProxy container is running in one of the controller nodes. On that controller within the HAProxy container namespace verify the IP address is assigned to the API interface. Also, verify that your ToR and the network infrastructure connecting your ToR is provisioned with API network segment VLAN.

HAProxy Services Downtime after Initial Installation or HA Failover

The HAProxy web interface can be accessed on TCP port 1936

```
http://<external_lb_vip_address>:1936/
Username: haproxy
Password: <HAPROXY_PASSWORD> from secrets.yaml file
```

After initial installation, the HAProxy web interface might report several OpenStack services with downtime depending upon when that OpenStack service was installed after HAProxy install. The counters are not synchronized between HAProxy active and standby. After HA proxy failover, the downtime timers might change based on the uptime of new active HAProxy container.

Management Node Problems

Service Commands

To identify all the services that are running, enter:

```
$ systemctl -a | grep docker | grep service
On controller ignore status of:
docker-neutronlb
On compute ignore status of:
docker-neutronlb, docker-keystone
```

To start a service on a host, enter:

```
$ systemctl start <service_name>
```

To stop a service on a host, enter:

```
$ systemctl stop <service_name>
```


To restart a service on a host, enter:

```
$ systemctl restart <service_name>
```

To check service status on a host, enter:


```
$ systemctl status <service_name>
```

Management Node Recovery Scenarios

The Cisco NFVI management node hosts the Cisco VIM Rest API service, Cobbler for PXE services, ELK for Logging/Kibana dashboard services and VMTP for cloud validation. Because the maintenance node currently does not have redundancy, understanding its points of failure and recovery scenarios is important. These are described in this topic.

The management node architecture includes a Cisco UCS C240 M4 server with dual CPU socket. It has a 1 Gbps on-board (LOM) NIC and a 10 Gbps Cisco VIC mLOM. HDDs are used in 8, 16, or 24 disk configurations. The following figure shows a high level maintenance node hardware and software architecture.

Figure 1: Cisco NFVI Management Node Architecture



Different management node hardware or software failures can cause Cisco NFVI service disruptions and outages. Some failed services can be recovered through manual intervention. In cases where the system is operational during a failure, double faults might not be recoverable. The following table lists different management node failure scenarios and their recovery options.

Table 1: Management Node Failure Scenarios

Scenario #	Failure/Trigger	Recoverable?	Operational Impact
1	Failure of 1 or 2 active HDD	Yes	No

Scenario #	Failure/Trigger	Recoverable?	Operational Impact
2	Simultaneous failure of more than 2 active HDD	No	Yes
3	Spare HDD failure: 4 spare for 24 HDD; or 2 spare for 8 HDD	Yes	No
4	Power outage/hard reboot	Yes	Yes
5	Graceful reboot	Yes	Yes
6	Docker daemon start failure	Yes	Yes
7	Service container (Cobbler, ELK) start failure	Yes	Yes
8	One link failure on bond interface	Yes	No
9	Two link failures on bond interface	Yes	Yes
10	REST API service failure	Yes	No
11	Graceful reboot with Cisco VIM Insight	Yes	Yes; CLI alternatives exist during reboot.
12	Power outage/hard reboot with Cisco VIM Insight	Yes	Yes
13	VIM Insight Container reinstallation	Yes	Yes; CLI alternatives exist during re-insight.
14	Cisco VIM Insight Container reboot	Yes	Yes; CLI alternatives exist during reboot.
15	Intel 1350 1Gbps LOM failure	Yes	Yes
16	Cisco VIC 1227 10 Gbps mLOM failure	Yes	Yes
17	DIMM memory failure	Yes	No
18	One CPU failure	Yes	No

Scenario 1: Failure of one or two active HDDs

The management node has either 8,16, or 24 HDDs. The HDDs are configured with RAID 6, which helps enable data redundancy and storage performance and overcomes any unforeseen HDD failures.

- When 8 HDDs are installed, 7 are active disks and one is spare disk.

- When 16 HDDs are installed, 14 are active disks and two are spare disks.
- When 24 HDDs are installed, 20 are active disks and four are spare disks.

With RAID 6 up to two simultaneous active HDD failures can occur. When an HDD fails, the system starts automatic recovery by moving the spare disk to active state and starts recovering and rebuilding the new active HDD. It takes approximately four hours to rebuild the new disk and move to synchronized state. During this operation, the system is completely functional and no impacts are seen. However, you must monitor the system to ensure that additional failures do not occur to enter into a double fault situation.

You can use the **storcli** commands to check the disk and RAID state as shown below:



Note Make sure the node is running with hardware RAID by checking the storcli output and comparing to the one preceding.

```
[root@mgmt-node ~]# /opt/MegaRAID/storcli/storcli64 /c0 show

<...snip...>

TOPOLOGY:
=====
-----
DG Arr Row EID:Slot DID Type State BT Size PDC PI SED DS3 FSpace TR
-----
0 - - - - RAID6 Optl N 4.087 TB dflt N N dflt N N
0 0 - - - RAID6 Optl N 4.087 TB dflt N N dflt N N <== RAID
6 in optimal state
0 0 0 252:1 1 DRIVE Onln N 837.258 GB dflt N N dflt - N
0 0 1 252:2 2 DRIVE Onln N 837.258 GB dflt N N dflt - N
0 0 2 252:3 3 DRIVE Onln N 930.390 GB dflt N N dflt - N
0 0 3 252:4 4 DRIVE Onln N 930.390 GB dflt N N dflt - N
0 0 4 252:5 5 DRIVE Onln N 930.390 GB dflt N N dflt - N
0 0 5 252:6 6 DRIVE Onln N 930.390 GB dflt N N dflt - N
0 0 6 252:7 7 DRIVE Onln N 930.390 GB dflt N N dflt - N
0 - - 252:8 8 DRIVE DHS - 930.390 GB - - - - - N

-----

<...snip...>

PD LIST:
=====
-----
EID:SlT DID State DG Size Intf Med SED PI SeSz Model Sp
-----
252:1 1 Onln 0 837.258 GB SAS HDD N N 512B ST900MM0006 U <== all disks
functioning
252:2 2 Onln 0 837.258 GB SAS HDD N N 512B ST900MM0006 U
252:3 3 Onln 0 930.390 GB SAS HDD N N 512B ST91000640SS U
252:4 4 Onln 0 930.390 GB SAS HDD N N 512B ST91000640SS U
252:5 5 Onln 0 930.390 GB SAS HDD N N 512B ST91000640SS U
252:6 6 Onln 0 930.390 GB SAS HDD N N 512B ST91000640SS U
252:7 7 Onln 0 930.390 GB SAS HDD N N 512B ST91000640SS U
252:8 8 DHS 0 930.390 GB SAS HDD N N 512B ST91000640SS D

-----

[root@mgmt-node ~]# /opt/MegaRAID/storcli/storcli64 /c0 show

<...snip...>

TOPOLOGY :
=====
-----
DG Arr Row EID:Slot DID Type State BT Size PDC PI SED DS3 FSpace TR
-----
```

```

0 - - - - RAID6 Pdgd N 4.087 TB dflt N N dflt N N <== RAID 6
in degraded state
0 0 - - - RAID6 Dgrd N 4.087 TB dflt N N dflt N N
0 0 0 252:8 8 DRIVE Rbld Y 930.390 GB dflt N N dflt - N
0 0 1 252:2 2 DRIVE Onln N 837.258 GB dflt N N dflt - N
0 0 2 252:3 3 DRIVE Onln N 930.390 GB dflt N N dflt - N
0 0 3 252:4 4 DRIVE Onln N 930.390 GB dflt N N dflt - N
0 0 4 252:5 5 DRIVE Onln N 930.390 GB dflt N N dflt - N
0 0 5 252:6 6 DRIVE Onln N 930.390 GB dflt N N dflt - N
0 0 6 252:7 7 DRIVE Onln N 930.390 GB dflt N N dflt - N
-----

```

<...snip...>

PD LIST :
=====

```

-----
EID:Slot DID State DG          Size Intf Med SED PI SeSz Model          Sp
-----
252:1      1 UGood - 837.258 GB SAS HDD N N 512B ST900MM0006      U <== active disk
in slot 1 disconnected from drive group 0
252:2      2 Onln 0 837.258 GB SAS HDD N N 512B ST900MM0006      U
252:3      3 Onln 0 930.390 GB SAS HDD N N 512B ST91000640SS      U
252:4      4 Onln 0 930.390 GB SAS HDD N N 512B ST91000640SS      U
252:5      5 Onln 0 930.390 GB SAS HDD N N 512B ST91000640SS      U
252:6      6 Onln 0 930.390 GB SAS HDD N N 512B ST91000640SS      U
252:7      7 Onln 0 930.390 GB SAS HDD N N 512B ST91000640SS      U
252:8      8 Rbld 0 930.390 GB SAS HDD N N 512B ST91000640SS      U <== spare disk
in slot 8 joined drive group 0 and in rebuilding state
-----

```

```

[root@mgmt-node ~]# /opt/MegaRAID/storcli/storcli64 /c0/e252/s8 show rebuild
Controller = 0
Status = Success
Description = Show Drive Rebuild Status Succeeded.

```

```

-----
Drive-ID      Progress% Status          Estimated Time Left
-----
/c0/e252/s8      20 In progress 2 Hours 28 Minutes <== spare disk in slot 8 rebuild
status
-----

```

To replace the failed disk and add it back as a spare:

```

[root@mgmt-node ~]# /opt/MegaRAID/storcli/storcli64 /c0/e252/s1 add hotsparedrive dg=0
Controller = 0
Status = Success
Description = Add Hot Spare Succeeded.

```

```

[root@mgmt-node ~]# /opt/MegaRAID/storcli/storcli64 /c0 show

```

<...snip...>

TOPOLOGY :
=====

```

-----
DG Arr Row EID:Slot DID Type   State BT          Size PDC  PI SED DS3  FSpace TR
-----
0 - - - - RAID6 Pdgd N 4.087 TB dflt N N dflt N N
0 0 - - - RAID6 Dgrd N 4.087 TB dflt N N dflt N N
0 0 0 252:8 8 DRIVE Rbld Y 930.390 GB dflt N N dflt - N
0 0 1 252:2 2 DRIVE Onln N 837.258 GB dflt N N dflt - N
0 0 2 252:3 3 DRIVE Onln N 930.390 GB dflt N N dflt - N
0 0 3 252:4 4 DRIVE Onln N 930.390 GB dflt N N dflt - N
0 0 4 252:5 5 DRIVE Onln N 930.390 GB dflt N N dflt - N
0 0 5 252:6 6 DRIVE Onln N 930.390 GB dflt N N dflt - N
0 0 6 252:7 7 DRIVE Onln N 930.390 GB dflt N N dflt - N
0 - - - 252:1 1 DRIVE DHS - 837.258 GB - - - - - N

```

```
-----
<...snip...>
```

```
PD LIST :
=====
```

```
-----
EID:Slit DID State DG          Size Intf Med SED PI SeSz Model          Sp
-----
252:1      1 DHS   0  837.258 GB SAS  HDD N   N  512B ST900MM0006      U   <== replacement
  disk added back as spare
252:2      2 Onln  0  837.258 GB SAS  HDD N   N  512B ST900MM0006      U
252:3      3 Onln  0  930.390 GB SAS  HDD N   N  512B ST91000640SS      U
252:4      4 Onln  0  930.390 GB SAS  HDD N   N  512B ST91000640SS      U
252:5      5 Onln  0  930.390 GB SAS  HDD N   N  512B ST91000640SS      U
252:6      6 Onln  0  930.390 GB SAS  HDD N   N  512B ST91000640SS      U
252:7      7 Onln  0  930.390 GB SAS  HDD N   N  512B ST91000640SS      U
252:8      8 Rbld  0  930.390 GB SAS  HDD N   N  512B ST91000640SS      U
-----
```

Scenario 2: Simultaneous failure of more than two active HDDs

If more than two HDD failures occur at the same time, the management node goes into an unrecoverable failure state because RAID 6 allows for recovery of up to two simultaneous HDD failures. To recover the management node, reinstall the operating system.

Scenario 3: Spare HDD failure

When the management node has 24 HDDs, four are designated as spares. Failure of any of the disks does not impact the RAID or system functionality. Cisco recommends replacing these disks when they fail (see the steps in Scenario 1) to serve as standby disks and so when an active disk fails, an auto-rebuild is triggered.

Scenario 4: Power outage/hard reboot

If a power outage or hard system reboot occurs, the system will boot up and come back to operational state. Services running on management node during down time will be disrupted. See the steps in Scenario 9 for the list of commands to check the services status after recovery.

Scenario 5: System reboot

If a graceful system reboot occurs, the system will boot up and come back to operational state. Services running on management node during down time will be disrupted. See the steps in Scenario 9 for the list of commands to check the services status after recovery.

Scenario 6: Docker daemon start failure

The management node runs the services using Docker containers. If the Docker daemon fails to come up, it causes services such as ELK, Cobbler and VMTP to go into down state. You can use the **systemctl** command to check the status of the Docker daemon, for example:

```
# systemctl status docker
docker.service - Docker Application Container Engine
Loaded: loaded (/usr/lib/systemd/system/docker.service; enabled; vendor preset: disabled)
Active: active (running) since Mon 2016-08-22 00:33:43 CEST; 21h ago
Docs: http://docs.docker.com
Main PID: 16728 (docker)
```

If the Docker daemon is in down state, use the **systemctl restart docker** command to restart the Docker service. Run the commands listed in Scenario 9 to verify that all the Docker services are active.

Scenario 7: Service container (Cobbler, ELK) start failure

As described in Scenario 8, all the services run as Docker containers on the management node. To find all services running as containers, use the **docker ps -a** command. If any services are in Exit state, use the **systemctl** command and **grep** for Docker to find the exact service name, for example:

```
# systemctl | grep docker- | awk '{print $1}'
docker-cobbler-tftp.service
docker-cobbler-web.service
```

```
docker-cobbler.service
docker-container-registry.service
docker-elasticsearch.service
docker-kibana.service
docker-logstash.service
docker-vmtp.service
```

If any services need restarting, use the **systemctl** command. For example, to restart a Kibana service:

```
# systemctl restart docker-kibana.service
```

Scenario 8: One link failure on the bond Interface

The management node is set up with two different networks: br_api and br_mgmt. The br_api interface is the external. It is used for accessing outside services such as the Cisco VIM REST API, Kibana and Cobbler. The br_mgmt interface is internal. It is used for provisioning and to provide management connectivity to all OpenStack nodes (control, compute and storage). Each network has two ports that are bonded to provide redundancy. If one port fails, the system will remain completely functional through the other port. If a port fails, check for physical network connectivity and/or remote switch configuration to debug the underlying cause of the link failure.

Scenario 9: Two link failures on the bond Interface

As described in Scenario 10, each network is configured with two ports. If both ports are down, the system is not reachable and management node services could be disrupted. After the ports are back up, the system is fully operational. Check the physical network connectivity and/or the remote switch configuration to debug the underlying link failure cause.

Scenario 10: REST API service failure

The management node runs the REST API service for Cisco VIM clients to reach the server. If the REST service is down, Cisco VIM clients cannot reach the server to trigger any server operations. However, with the exception of the REST service, other management node services remain operational.

To verify the management node REST services are fully operational, use the following command to check that the httpd and mercury-restapi services are in active and running state:

```
# systemctl status httpd
httpd.service - The Apache HTTP Server
  Loaded: loaded (/usr/lib/systemd/system/httpd.service; enabled; vendor preset: disabled)
  Active: active (running) since Mon 2016-08-22 00:22:10 CEST; 22h ago

# systemctl status mercury-restapi.service
mercury-restapi.service - Mercury Restapi
  Loaded: loaded (/usr/lib/systemd/system/mercury-restapi.service; enabled; vendor preset: disabled)
  Active: active (running) since Mon 2016-08-22 00:20:18 CEST; 22h ago
```

A tool is also provided so that you can check the REST API server status and the location of the directory it is running from. To execute run the following command:

```
# cd installer-<tagid>/tools
# ./restapi.py -a status
Status of the REST API Server: active (running) since Thu 2016-08-18 09:15:39 UTC; 9h ago
REST API launch directory: /root/installer-<tagid>/
```

Confirm the server status is active and check that the restapi launch directory matches the directory where the installation was launched. The restapi tool also provides the options to launch, tear down, and reset password for the restapi server as shown below:

```
# ./restapi.py -h
usage: restapi.py [-h] --action ACTION [--yes] [--verbose]
REST API setup helper

optional arguments:
  -h, --help            show this help message and exit
```

```

--action ACTION, -a ACTION
    setup - Install and Start the REST API server.
    teardown - Stop and Uninstall the REST API
server.
    restart - Restart the REST API server.
    regenerate-password - Regenerate the password for
REST API server.
    reset-password - Reset the REST API password with
user given password.
    status - Check the status of the REST API server
--yes, -y      Skip the dialog. Yes to the action.
--verbose, -v  Perform the action in verbose mode.

```

If the REST API server is not running, execute **ciscovim** to show the following error message:

```

# cd installer-<tagid>/
# ciscovim -setupfile ~/Save/<setup_data.yaml> run

```

If the installer directory or the REST API state is not correct or points to an incorrect REST API launch directory, go to the installer-<tagid>/tools directory and execute:

```
# ./restapi.py -action setup
```

To confirm that the REST API server state and launch directory is correct run the following command:

```
# ./restapi.py -action status
```

Scenario 11: Graceful reboot with Cisco VIM Insight

Cisco VIM Insight runs as a container on the management node. After a graceful reboot of the management node, the VIM Insight and its associated database containers will come up. So there is no impact on recovery.

Scenario 12: Power outage or hard reboot with VIM Insight

The Cisco VIM Insight container will come up automatically following a power outage or hard reset of the management node.

Scenario 13: Cisco VIM Insight reinstallation

If the management node which is running the Cisco VIM Insight fails and cannot come up, you must uninstall and reinstall the Cisco VIM Insight. After the VM Insight container comes up, add the relevant bootstrap steps as listed in the install guide to register the pod. VIM Insight then automatically detects the installer status and reflects the current status appropriately.

To clean up and reinstall Cisco VIM Insight run the following command:

```

# cd /root/installer-<tagid>/insight/
# ./bootstrap_insight.py -a uninstall -o standalone -f </root/insight_setup_data.yaml>

```

Scenario 14: VIM Insight Container reboot

On Reboot of the VIM Insight container, services will continue to work as it is.

Scenario 15: Intel (I350) 1Gbps LOM failure

The management node is set up with an Intel (I350) 1 Gbps LOM for API connectivity. Two 1 Gbps ports are bonded to provide connectivity redundancy. No operational impact occurs if one of these ports goes down. However, if both ports fail, or the LOM network adapter fails, the system cannot be reached through the API IP address. If this occurs you must replace the server because the LOM is connected to the system motherboard. To recover the management node with a new server, complete the following steps. Make sure the new management node hardware profile matches the existing server and the Cisco IMC IP address is assigned.

- 1 Shut down the existing management node.
- 2 Unplug the power from the existing and new management nodes.
- 3 Remove all HDDs from existing management node and install them in the same slots of the new management node.
- 4 Plug in the power to the new management node, but do not boot the node.

- 5 Verify the configured boot order is set to boot from local HDD.
- 6 Verify the Cisco NFVI management VLAN is configured on the Cisco VIC interfaces.
- 7 Boot the management node for the operating system to start.

After the management node is up, the management node bond interface will be down due to the incorrect MAC address on the ifcfg files. It will point to old node network card MAC address.
- 8 Update the MAC address on the ifcfg files under `/etc/sysconfig/network-scripts`.
- 9 Reboot the management node.

It will come up and be fully operational. All interfaces should be in an up state and be reachable.
- 10 Verify that Kibana and Cobbler dashboards are accessible.
- 11 Verify the Rest API services are up. See Scenario 15 for any recovery steps.

Scenario 16: Cisco VIC 1227 10Gbps mLOM failure

The management node is configured with a Cisco VIC 1227 dual port 10 Gbps mLOM adapter for connectivity to the other Cisco NFVI nodes. Two 10 Gbps ports are bonded to provide connectivity redundancy. If one of the 10 Gbps ports goes down, no operational impact occurs. However, if both Cisco VIC 10 Gbps ports fail, the system goes into an unreachable state on the management network. If this occurs, you must replace the VIC network adapters. Otherwise pod management and the Fluentd forwarding service will be disrupted. If you replace a Cisco VIC, update the management and provisioning VLAN for the VIC interfaces using Cisco IMC and update the MAC address in the interfaces under `/etc/sysconfig/network-scripts` interface configuration file.

Scenario 17: DIMM memory failure

The management node is set up with multiple DIMM memory across different slots. Failure of one or memory modules could cause the system to go into unstable state, depending on how many DIMM memory failures occur. DIMM memory failures are standard system failures like any other Linux system server. If a DIMM memory fails, replace the memory module(s) as soon as possible to keep the system in stable state.

Scenario 18: One CPU failure

Cisco NFVI management nodes have dual core Intel CPUs (CPU1 and CPU2). If one CPU fails, the system remains operational. However, always replace failed CPU modules immediately. CPU failures are standard system failures like any other Linux system server. If a CPU fails, replace it immediately to keep the system in stable state.

Recovering Compute Node Scenario

The Cisco NFVI Compute node hosts the OpenStack services to provide processing, network, and storage resources to run instances. The node architecture includes a Cisco UCS C220 M4 server with dual CPU socket, 10 Gbps Cisco VIC mLOM, and two HDDs in RAID 1 configuration.

Failure of one active HDD

With RAID 1, data are mirrored and can allow up to one active HDD failure. When a HDD fails, the node is still functional with no impacts. However, the data are no longer mirrored and losing another HDD will result in unrecoverable and operational downtime. The failed disk should be replaced soon as it takes approximately two hours to rebuild the new disk and move to synchronized state.

To check the disk and RAID state run the `storcli` commands as follows:



Note

Make sure that the node is running with hardware RAID by checking the storcli output and comparing to the one below.

```
[root@compute-node ~]# /opt/MegaRAID/storcli/storcli64 /c0 show
<...snip...>
TOPOLOGY :
=====
-----
DG Arr Row EID:Slot DID Type State BT Size PDC PI SED DS3 FSpace TR
-----
0 - - - - RAID1 Opt1 N 837.258 GB dflt N N dflt N N <== RAID 1 in
optimal state
0 0 - - - RAID1 Opt1 N 837.258 GB dflt N N dflt N N
0 0 0 252:2 9 DRIVE Onln N 837.258 GB dflt N N dflt - N
0 0 1 252:3 11 DRIVE Onln N 837.258 GB dflt N N dflt - N
-----
```

```
<...snip...>
Physical Drives = 2
PD LIST :
=====
-----
EID:SlT DID State DG Size Intf Med SED PI SeSz Model Sp
-----
252:2 9 Onln 0 837.258 GB SAS HDD N N 512B ST900MM0006 U <== all disks
functioning
252:3 11 Onln 0 837.258 GB SAS HDD N N 512B ST900MM0006 U
-----
```

```
[root@compute-node ~]# /opt/MegaRAID/storcli/storcli64 /c0 show
<...snip...>
TOPOLOGY :
=====
-----
DG Arr Row EID:Slot DID Type State BT Size PDC PI SED DS3 FSpace TR
-----
0 - - - - RAID1 Dgrd N 837.258 GB dflt N N dflt N N <== RAID 1 in
degraded state.
0 0 - - - RAID1 Dgrd N 837.258 GB dflt N N dflt N N
0 0 0 - - DRIVE Msng - 837.258 GB - - - - N
0 0 1 252:3 11 DRIVE Onln N 837.258 GB dflt N N dflt - N
-----
```

```
<...snip...>
PD LIST :
=====
-----
EID:SlT DID State DG Size Intf Med SED PI SeSz Model Sp
-----
252:2 9 UGood - 837.258 GB SAS HDD N N 512B ST900MM0006 U <== active disk
in slot 2 disconnected from drive group 0
252:3 11 Onln 0 837.258 GB SAS HDD N N 512B ST900MM0006 U
-----
```

To replace the failed disk and add it back as a spare run the following command:

```
[root@compute-node ~]# /opt/MegaRAID/storcli/storcli64 /c0/e252/s2 add hotsparedrive dg=0
Controller = 0
Status = Success
Description = Add Hot Spare Succeeded.
```

```
[root@compute-node ~]# /opt/MegaRAID/storcli/storcli64 /c0 show
<...snip...>
TOPOLOGY :
=====
-----
DG Arr Row EID:Slot DID Type State BT Size PDC PI SED DS3 FSpace TR
-----
0 - - - - RAID1 Dgrd N 837.258 GB dflt N N dflt N N
0 0 - - - RAID1 Dgrd N 837.258 GB dflt N N dflt N N
0 0 0 252:2 9 DRIVE Rbld Y 837.258 GB dflt N N dflt - N
0 0 1 252:3 11 DRIVE Onln N 837.258 GB dflt N N dflt - N
-----

<...snip...>
PD LIST :
=====
-----
EID:SlT DID State DG Size Intf Med SED PI SeSz Model Sp
-----
252:2 9 Rbld 0 837.258 GB SAS HDD N N 512B ST900MM0006 U <== replacement
disk in slot 2 joined device group 0 and in rebuilding state
252:3 11 Onln 0 837.258 GB SAS HDD N N 512B ST900MM0006 U
-----

[root@compute-node ~]# /opt/MegaRAID/storcli/storcli64 /c0/e252/s2 show rebuild
Controller = 0
Status = Success
Description = Show Drive Rebuild Status Succeeded.

-----
Drive-ID Progress% Status Estimated Time Left
-----
/c0/e252/s2 10 In progress 1 Hours 9 Minutes <== replacement disk in slot 2 rebuild
status
-----
```

Running the Cisco VIM Technical Support Tool

Cisco VIM includes a tech-support tool that you can use to gather Cisco VIM information to help solve issues working with Cisco Technical Support. The tech-support tool can be extended to execute custom scripts. It can be called after runner is executed at least once. The tech-support tool uses a configuration file that specifies what information to collect. The configuration file is located in the following location:
/root/openstack-configs/tech-support/tech_support_cfg.yaml.

The tech-support tool keeps track of the point where the Cisco VIM installer has executed and collects the output of files or commands indicated by the configuration file. For example, if the installer fails at Step 3 (VALIDATION), the tech-support will provide information listed in the configuration file up to Step 3 (included). You can override this default behavior by adding the --stage option to the command.

The tech-support script is located at the management node /root/installer-{tag-id}/tech-support directory. To run it after the runner execution, enter the following command:

```
./tech-support/tech_support.py
```

The command creates a compressed tar file containing all the information it gathered. The file location is displayed in the console at the end of the execution. You need not have to execute the command with any options. However, if you want to override any default behavior, you can use the following options:

```
/tech_support.py --help
Usage: tech_support.py [options]

Tech-support collects information about your cloud

Options:
  -h, --help                show this help message and exit
  --stage=STAGE             specify the stage where installer left off
  --config-file=CFG_FILE   specify alternate configuration file name
  --tmp-dir=TMP_DIR        specify alternate temporary directory name
  --file-size=TAIL_SIZE    specify max size (in KB) of each file collected
  --host-list=HOST_LIST    List (comma separated) of the hostnames of the servers
                           to collect info from
  --ip-list=IP_LIST        List (comma separated) of the IPv4 of the hosts to
                           collect info from
  --exclude-mgmt-node      specify if mgmt node info needs to be excluded
```

Where:

- **stage**—tells at which state the installer left off. The possible values are: `INPUT_VALIDATION`, `BUILDNODE_ORCHESTRATION`, `VALIDATION`, `BAREMETAL_INSTALL`, `COMMON_SETUP`, `CEPH`, `ORCHESTRATION` or `VMTP`
- **config-file**—Provides the path for a specific configuration file. Make sure that your syntax is correct. Look at the default `/root/tech-support/openstack-configs/tech_support_cfg.yaml` file as an example on how to create a new config-file or modify the default file.
- **tmp-dir**—Provides the path to a temp directory tech-support can use to create the compressed tar file. The tech-support tool provides the infrastructure to execute standard Linux commands from packages included in the Cisco VIM installation. This infrastructure is extensible and you can add commands, files, or custom bash/Python scripts into the appropriate configuration file sections for the tool to collect the output of those commands/scripts. (See the README section at the beginning of the file for more details on how to do this.)
- **file-size**—Is an integer that specifies (in KB) the maximum file size that tech-support will capture and tail the file if needed. By default, this value is set to 10 MB. For example, if no file-size option is provided and the tech-support needs to collect `/var/log/mercury/data.log` and the `data.log` is more than 10 MB, tech-support will get the last 10 MB from `/var/log/mercury/data.log`.
- **host-list**: Provides the list of hosts one wants to collect from the tech-support through hostname defaults, to all hosts.
- **ip-list**: Provides the list of hosts one wants to collect the tech-support through management IP, defaults to all hosts.
- **exclude-mgmt-node**: It is an option not to collect tech-support from management node.

Tech-support configuration file

Cisco VIM tech-support is a utility tool designed to collect the VIM pod logs which help users to debug the issues offline. The administrator uses the tech-support configuration files to provide the list of commands or configuration files. The tech support tool of the Cisco VIM gathers list of commands or configuration files for the offline diagnostic or debugging purposes.

By default the tech-support configuration file is located at the `/root/openstack-configs/tech-support/tech_support_cfg.yaml` file. Alternatively, you can use a different one by specifying the `-config-file` option. The syntax of this configuration file must be as follows:

The tech-support configuration file section is divided into eight sections which corresponds to each of the installer stages:

- INPUT_VALIDATION
- BUILDNODE_ORCHESTRATION
- VALIDATION
- BAREMETAL_INSTALL
- COMMON_SETUP
- CEPH
- ORCHESTRATION
- VMTP

Inside each of these eight sections, there are tags divided on hierarchical levels. At the first level, the tag indicates the host(s) or path on which the command(s) run and from where the file(s) can be collected. The possible tags are as follows:

- - HOSTS_MANAGEMENT: Run in Management node only
- - HOSTS_CONTROL: Run in all the Control nodes
- - HOSTS_COMPUTE: Run in all the Compute nodes
- - HOSTS_STORAGE: Run in all the Storage nodes
- - HOSTS_COMMON: Run in all the Compute and Control nodes
- - HOSTS_ALL: Run in all the Compute, Control and Storage nodes

**Note**

In any of these eight sections, if HOSTS tag is not specified then no information is collected for that stage.

For each of the hosts mentioned above there will be a second level tag which specifies where to run the command. The possible values of those tags are as follows:

- - SERVER_FILES: Path(s) to the file(s) that tech-support needs to collect.
- - SERVER_COMMANDS: Command(s) or script name(s) which need to be executed directly on the desired server. The command(s) need to be already included in the \$PATH. For the scripts, please refer to the Custom Scripts paragraph below.
- - CONTAINERS: Indicates the tech-support tool that the command(s) need to be executed and the files to be gathered from inside a container. See the following steps for more specific information of what can be added in this section.

In the CONTAINERS section, indicate the path in which container the commands should be executed or gathered from. This is done with a `<container_name>` tag (Shown below are example of where to get the string for the `<container_name>` tag):

- `all_containers`: Execute inside all containers (regardless of the state).
- `<container_name>`: This must be the name of a container and it indicates in which container to run the command or gather the information. This will run commands inside the container only if the mentioned container is up (as we cannot run commands on dead containers). Example of how to get the container name:
 - Execute **docker ps** and get the name (without any numbers) of the last column of output **docker ps -a**.

For example:

```
CONTAINER ID   IMAGE                                COMMAND <snip>   NAMES
81bc4e54cbfb  <registry>/vmtpl:4263              /bin/bash"      vmtpl_4263
```

The tech-support runs the linux commands on the server (from packages included in RHEL7.3). Add the name of the commands under the `SERVER_COMMANDS` section of the configuration file to run the commands.

However, if the administrator wants to add a custom bash or python script to be executed in some set of servers in the cloud. In this case the user just needs to add the script into the custom-scripts directory on the current directory path (`/root/openstack-configs/tech-support/`) and add the script name into the corresponding `SERVER_COMMANDS` section.

The tech-support tool will scp the script(s) included in the custom-scripts directory into the appropriate cloud nodes where it will be executed (as# indicated in this config file) and capture the output (stdout and stderr) and add it to the collection of files collected by the tech-support tool. It is assumed that the scripts are self-standing and independent and needs no external input.

Following is an example of a custom tech-support configuration file. This is just an example of what information the tech-support tool will gather if given the following configuration file:

```
COMMON_SETUP:
HOSTS_ALL: # All compute, control and storage hosts
SERVER_FILES:
- /usr/lib/docker-storage-setup
SERVER_COMMANDS:
- docker info
- my_script.sh
CONTAINERS:
all_containers: #execute in all containers (even if they are in down state)
CONTAINER_COMMANDS:
- docker inspect
- docker logs
logstash:
CONTAINER_FILES:
- /var/log/
CONTAINER_COMMANDS:
- ls -l
```

Given this example of configuration, and assuming that the installer ended in at least the `COMMON_SETUP` state, the tech-support tool will run under all OpenStack nodes (Compute, Control and Storage) and it will:

- Gather (if exists) the contents of `/usr/lib/docker-storage-setup` file.
- Run **docker info** command and collect the output.
- Run **my_script.sh** and collect the output. The **my_script.sh** is an example of a bash script which the user previously added to the `/root/openstack-configs/tech-support/custom-scripts` directory.
- Collect the output of `docker inspect` and `docker logs` for all containers.
- Collect the files in `/var/log` inside the `logstash` container (if there is container with that name). This is equivalent to running the following command (where `/tmp` indicates a temporary location where the tech-support tool gathers all the information): **docker cp logstash_{tag}:/var/log/ /tmp**.

- Collect the output of the command `docker exec logstash_{{tag}}: ls -l`.

Tech-support when servers are offline

If one or more cloud nodes are not reachable, you will not be able to collect the information from the servers.. In this case, you can connect through the KVM console into those servers and run the local tech-support tool.

Step 1 To run the local tech-support tool run the following command:

```
/root/tech_support_offline
```

Step 2 Cisco VIM tech_support_offline will collect the logs and other troubleshooting output from the server and put it in the location of that server:

```
/root/tech_support
```

Note Then after the server is reachable again, the user can just use the Cisco VIM tech-support tool which will collect all files under the /root/tech-support/ directory among with the rest of the VIM logs which will be used to debug any issue offline.

Disk-maintenance tool to manage physical drives

In VIM 2.2 you can use the disk-maintenance tool to check the status of all physical drives that are present in running and operational nodes in the following roles -

- Management
- Control (all or specific nodes)
- Compute (all or specific nodes)

This provides the information about the current status of the physical drives - if they are in Online, Offline, Rebuilding, Unconfigured Good or JBOD states if all disks are ok. If not, the disks that have gone bad are displayed with the slot number and server information, that needs to be replaced. When multiple disks needs to be replaced, we recommend you to execute the remove/add of the node.

- Physically remove and insert a new disk before attempting to replace.
- For smooth operation, wipe out disk before attempting replace operations.
- Call Cisco TAC if any issue is hit. Do not reattempt.



Note Make sure that each node is running with hardware RAID , the steps for which can be found in the section titled Recovering Compute Node Scenario. Refer to step 15 of the section "Upgrading Cisco VIM Software Using a USB" on how to move the pod from hardware RAID to software RAID.

To check the status of the Disk Management run the following commands:

```
[root@cork-mgmt]# ciscovim help diskmgmt
usage: ciscovim diskmgmt [--server <node1,node2,...>] [-y]
      (--check_disks <disk_check_server_roles> | --replace_disks
<disk_replace_server_roles>)
```

HDD maintenance helper

Optional arguments:

- server <node1,node2,...> List of specific control/compute hostnames
- y, --yes Yes option to perform the action
- check_disks <disk_check_server_roles>
 - all,management,control,compute - Run on all servers and/or on management server and/or on all controller servers and/or on all compute servers.
- replace_disks <disk_replace_server_roles>
 - all,management,control,compute - Run on all servers and/or on management server and/or on all controller servers and/or on all compute servers.

```
[root@EMC-Testbed ~]# ciscovim diskmgmt --check_disks all
```

Perform the action. Continue (Y/N)Y

Message	Host	Role	Server	State
Raid Health Status	EMC-Testbed	management	localhost	Optimal
	i13-20	control	15.0.0.7	Optimal
	i13-22	control	15.0.0.5	Optimal
	i13-21	control	15.0.0.8	Optimal
	i13-23	compute	15.0.0.6	Optimal
	i13-24	compute	15.0.0.10	Optimal
VD Health Status	EMC-Testbed	management	localhost	Optimal
	i13-20	control	15.0.0.7	Optimal
	i13-22	control	15.0.0.5	Optimal
	i13-21	control	15.0.0.8	Optimal
	i13-23	compute	15.0.0.6	Optimal
	i13-24	compute	15.0.0.10	Optimal
RAID Level and Type	EMC-Testbed	management	localhost	Type - HW; Level - RAID10
	i13-20	control	15.0.0.7	Type - HW; Level - RAID10
	i13-22	control	15.0.0.5	Type - HW; Level - RAID10

		i13-21	control	15.0.0.8	Type - HW; Level - RAID10	
		i13-23	compute	15.0.0.6	Type - HW; Level - RAID10	
		i13-24	compute	15.0.0.10	Type - HW; Level - RAID1	
Number of Physical Disks	EMC-Testbed	management	localhost	4		
		i13-20	control	15.0.0.7	4	
		i13-22	control	15.0.0.5	5	
		i13-21	control	15.0.0.8	4	
		i13-23	compute	15.0.0.6	5	
		i13-24	compute	15.0.0.10	6	
Number of Virtual Disks	EMC-Testbed	management	localhost	1		
		i13-20	control	15.0.0.7	1	
		i13-22	control	15.0.0.5	1	
		i13-21	control	15.0.0.8	1	
		i13-23	compute	15.0.0.6	1	
		i13-24	compute	15.0.0.10	1	
Boot Drive Disk Media-Type	EMC-Testbed	management	localhost	HDD		
		i13-20	control	15.0.0.7	HDD	
		i13-22	control	15.0.0.5	HDD	
		i13-21	control	15.0.0.8	HDD	
		i13-23	compute	15.0.0.6	HDD	
		i13-24	compute	15.0.0.10	SSD	
Spare disks available	i13-22	control	15.0.0.5	Slot - 5; State - DHS		
	i13-23	compute	15.0.0.6	Slot - 5; State - DHS		

State Keys:

DHS-Dedicated Hot Spare|UGood-Unconfigured Good|GHS-Global Hotspare
 UBad-Unconfigured Bad|Onln-Online|Offln-Offline
 Rbld-Rebuilding|JBOD-Just a Bunch Of Disks

```
[root@EMC-Testbed ~]# ciscovim diskmgmt --check_disks compute
```

Perform the action. Continue (Y/N)Y

Message	Host	Role	Server	State

Raid Health Status	i13-23	compute	15.0.0.6	Optimal
	i13-24	compute	15.0.0.10	Optimal
VD Health Status	i13-23	compute	15.0.0.6	Optimal
	i13-24	compute	15.0.0.10	Optimal
RAID Level and Type	i13-23	compute	15.0.0.6	Type - HW; Level - RAID10
	i13-24	compute	15.0.0.10	Type - HW; Level - RAID1
Number of Physical Disks	i13-23	compute	15.0.0.6	5
	i13-24	compute	15.0.0.10	6

Number of Virtual Disks	i13-23	compute	15.0.0.6	1
	i13-24	compute	15.0.0.10	1
Boot Drive Disk Media-Type	i13-23	compute	15.0.0.6	HDD
	i13-24	compute	15.0.0.10	SSD
Spare disks available	i13-23	compute	15.0.0.6	Slot - 5; State - DHS

```
[root@EMC-Testbed ~]# ciscovim diskmgmt --check_disks control --server i13-22
```

Perform the action. Continue (Y/N)Y

Message	Host	Role	Server	State
Raid Health Status	i13-22	control	15.0.0.5	Optimal
VD Health Status	i13-22	control	15.0.0.5	Optimal
RAID Level and Type	i13-22	control	15.0.0.5	Type - HW; Level - RAID10
Number of Physical Disks	i13-22	control	15.0.0.5	5
Number of Virtual Disks	i13-22	control	15.0.0.5	1
Boot Drive Disk Media-Type	i13-22	control	15.0.0.5	HDD
Spare disks available	i13-22	control	15.0.0.5	Slot - 5; State - DHS

```
[root@EMC-Testbed ~]# ciscovim diskmgmt --check_disks control --server i13-24
```

Perform the action. Continue (Y/N)Y

Message	Host	Role	Server	State
Raid Health Status	i13-24	compute	15.0.0.10	Need Attention
VD Health Status	i13-24	compute	15.0.0.10	Degraded
RAID Level and Type	i13-24	compute	15.0.0.10	Type - HW; Level - RAID1
Number of Physical Disks	i13-24	compute	15.0.0.10	6
Number of Virtual Disks	i13-24	compute	15.0.0.10	1
Boot Drive Disk Media-Type	i13-24	compute	15.0.0.10	SSD
BAD DISKS/ NEEDS CHECKING	i13-24	compute	15.0.0.10	Slot - 2; State - Offln

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

```
[root@EMC-Testbed ~]# ciscovim diskmgmt --check_disks control --server i13-21
```

```
Perform the action. Continue (Y/N)Y
```

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

Message	Host	Role	Server	State
Raid Health Status	i13-21	control	15.0.0.8	Need Attention
VD Health Status	i13-21	control	15.0.0.8	Degraded
RAID Level and Type	i13-21	control	15.0.0.8	Type - HW; Level - RAID10
Number of Physical Disks	i13-21	control	15.0.0.8	4
Number of Virtual Disks	i13-21	control	15.0.0.8	1
Boot Drive Disk Media-Type	i13-21	control	15.0.0.8	HDD
BAD DISKS/ NEEDS CHECKING	i13-21	control	15.0.0.8	Slot - 2; State - Offln

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

```
[root@EMC-Testbed ~]# ciscovim diskmgmt --check_disks control --server i13-21
```

```
Perform the action. Continue (Y/N)Y
```

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

Message	Host	Role	Server	State
Raid Health Status	i13-21	control	15.0.0.8	Need Attention
VD Health Status	i13-21	control	15.0.0.8	Degraded
RAID Level and Type	i13-21	control	15.0.0.8	Type - HW; Level - RAID10
Number of Physical Disks	i13-21	control	15.0.0.8	4
Number of Virtual Disks	i13-21	control	15.0.0.8	1
Boot Drive Disk Media-Type	i13-21	control	15.0.0.8	HDD
Disks currently Rebuilding	i13-21	control	15.0.0.8	Slot - 2; State - Rbld
DISK READY TO ADD AS SPARE	i13-21	control	15.0.0.8	Slot - 3; State - UGood

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

```
[root@EMC-Testbed ~]# ciscovim diskmgmt --replace_disks control --server i13-21
```

```

Perform the action. Continue (Y/N)Y
+-----+-----+-----+-----+-----+
| Message          | Host   | Role   | Server  | State          |
+-----+-----+-----+-----+-----+
| DISK(S) ADDED ONLINE | i13-21 | control | 15.0.0.8 | Slot - 3; State - Success |
|                   |        |         |          |                |
+-----+-----+-----+-----+-----+
State Keys:
DHS-Dedicated Hot Spare|UGood-Unconfigured Good|GHS-Global Hotspare
UBad-Unconfigured Bad|Onln-Online|Offln-Offline
Rbld-Rebuilding|JBOD-Just a Bunch Of Disks

```

OSD-maintenance tool

In VIM 2.2 you can use the `osd-maintenance` tool to check the status of all OSDs that are present in running and operational block storage nodes. OSD maintenance tool gives you the detailed information about the current status of the OSDs - if they are Up or Down, in addition to what HDD corresponds to which OSD, including the slot number and server hostname.

- If a down OSD is discovered after `check_osds` is performed, run cluster recovery and recheck.
- If still down, wait 30 minutes before attempting replace - time for `ceph-mon` to sync
- Physically remove and insert a new disk before attempting replace.
- For smooth operation, please wipe out disk before attempting replace operations.
- Need a dedicated journal SSD for each storage server where `osdmgmt` is attempted.
- Only allowed to replace one OSD at a time. Space out each replace OSD by 30 minutes - time for `ceph-mon` to sync.
- Call TAC if any issue is hit. Do not reattempt.

Examples of usage of this tool:

To check the status of the OSDs run the following commands:

```

[root@cork-mgmt]# ciscovim help osdmgmt
usage: ciscovim osdmgmt [--server <node1,node2,...>] [--detail]
        [--osd <osd_name>] [-y] (--check_osds | --replace_osd)

```

OSD maintenance helper

Optional arguments:

```

--server <node1,node2,...>   List of specific block_storage hostnames
--detail                     Display full OSD details
--osd <osd_name>           Name of down OSD to replace. Eg. 'osd.xx'

```

```
-y, --yes                Yes option to perform the action
--check_osds            Check status of OSDs on storage servers
--replace_osd          Replace a sepcific down OSD
```

```
[root@cork-mgmt tools]# ciscovim osdmgmt --check_osds
```

Perform the action. Continue (Y/N)Y

```
+-----+-----+-----+-----+-----+
| Message          | Host                | Role          | Server  | State    |
+-----+-----+-----+-----+-----+
| Overall OSD Status | cork-store-1.cisco.com | block_storage | 8.8.8.11 | Optimal  |
|                   | cork-store-3.cisco.com | block_storage | 8.8.8.20 | Optimal  |
|                   | cork-store-6.cisco.com | block_storage | 8.8.8.13 | Needs Attention |
|                   | cork-store-5.cisco.com | block_storage | 8.8.8.19 | Optimal  |
|                   | cork-store-4.cisco.com | block_storage | 8.8.8.17 | Optimal  |
|                   |                       |               |         |         |
| Number of OSDs    | cork-store-1.cisco.com | block_storage | 8.8.8.11 | 4        |
|                   | cork-store-3.cisco.com | block_storage | 8.8.8.20 | 8        |
|                   | cork-store-6.cisco.com | block_storage | 8.8.8.13 | 8        |
|                   | cork-store-5.cisco.com | block_storage | 8.8.8.19 | 8        |
|                   | cork-store-4.cisco.com | block_storage | 8.8.8.17 | 8        |
+-----+-----+-----+-----+-----+
| Host              | OSDs  | Status | ID | HDD Slot | Path      | Mount   | Journal |
+-----+-----+-----+-----+-----+
| cork-store-6.cisco.com | osd.33 | down   | 33 | 6 (UGood) | Missing  | Missing | Missing |
+-----+-----+-----+-----+-----+
```

```
[root@cork-mgmt tools]# ciscovim osdmgmt --check_osds --detail
```

Perform the action. Continue (Y/N)Y

```
+-----+-----+-----+-----+-----+
| Message          | Host                | Role          | Server  | State    |
+-----+-----+-----+-----+-----+
| Overall OSD Status | cork-store-1.cisco.com | block_storage | 8.8.8.11 | Optimal  |
|                   | cork-store-3.cisco.com | block_storage | 8.8.8.20 | Optimal  |
|                   | cork-store-5.cisco.com | block_storage | 8.8.8.19 | Optimal  |
|                   | cork-store-6.cisco.com | block_storage | 8.8.8.13 | Needs Attention |
|                   | cork-store-4.cisco.com | block_storage | 8.8.8.17 | Optimal  |
|                   |                       |               |         |         |
| Number of OSDs    | cork-store-1.cisco.com | block_storage | 8.8.8.11 | 4        |
|                   | cork-store-3.cisco.com | block_storage | 8.8.8.20 | 8        |
|                   | cork-store-5.cisco.com | block_storage | 8.8.8.19 | 8        |
|                   | cork-store-6.cisco.com | block_storage | 8.8.8.13 | 8        |
|                   | cork-store-4.cisco.com | block_storage | 8.8.8.17 | 8        |
+-----+-----+-----+-----+-----+
| Host              | OSDs  | Status | ID | HDD Slot | Path      | Mount   |
| Journal          |
+-----+-----+-----+-----+-----+
```

```

| cork-store-1.cisco.com | osd.0 | up | 0 | 5 | /dev/sde1 | /var/lib/ceph/osd/ceph-0
| /dev/sda4 |
| | osd.1 | up | 1 | 6 | /dev/sdf1 | /var/lib/ceph/osd/ceph-1
| /dev/sdb4 |
| | osd.2 | up | 2 | 7 | /dev/sg1 | /var/lib/ceph/osd/ceph-2
| /dev/sdc4 |
| | osd.7 | up | 7 | 8 | /dev/sdh1 | /var/lib/ceph/osd/ceph-7
| /dev/sdd4 |
| | | | | | |
| | | | | | |
| cork-store-3.cisco.com | osd.3 | up | 3 | 5 (JBOD) | /dev/sdc1 | /var/lib/ceph/osd/ceph-3
| /dev/sdk1 |
| | osd.8 | up | 8 | 4 (JBOD) | /dev/sdd1 | /var/lib/ceph/osd/ceph-8
| /dev/sdl1 |
| | osd.12 | up | 12 | 3 (JBOD) | /dev/sde1 | /var/lib/ceph/osd/ceph-12
| /dev/sdk2 |
| | osd.16 | up | 16 | 8 (JBOD) | /dev/sdf1 | /var/lib/ceph/osd/ceph-16
| /dev/sdl2 |
| | osd.20 | up | 20 | 9 (JBOD) | /dev/sg1 | /var/lib/ceph/osd/ceph-20
| /dev/sdk3 |
| | osd.24 | up | 24 | 7 (JBOD) | /dev/sdh1 | /var/lib/ceph/osd/ceph-24
| /dev/sdl3 |
| | osd.28 | up | 28 | 10 (JBOD) | /dev/sdi1 | /var/lib/ceph/osd/ceph-28
| /dev/sdk4 |
| | osd.32 | up | 32 | 6 (JBOD) | /dev/sdj1 | /var/lib/ceph/osd/ceph-32
| /dev/sdl4 |
| | | | | | |
| | | | | | |
| cork-store-5.cisco.com | osd.5 | up | 5 | 8 (JBOD) | /dev/sdc1 | /var/lib/ceph/osd/ceph-5
| /dev/sdk1 |
| | osd.11 | up | 11 | 3 (JBOD) | /dev/sdd1 | /var/lib/ceph/osd/ceph-11
| /dev/sdl1 |
| | osd.14 | up | 14 | 10 (JBOD) | /dev/sde1 | /var/lib/ceph/osd/ceph-14
| /dev/sdk2 |
| | osd.19 | up | 19 | 5 (JBOD) | /dev/sdf1 | /var/lib/ceph/osd/ceph-19
| /dev/sdl2 |
| | osd.22 | up | 22 | 7 (JBOD) | /dev/sg1 | /var/lib/ceph/osd/ceph-22
| /dev/sdk3 |
| | osd.25 | up | 25 | 4 (JBOD) | /dev/sdh1 | /var/lib/ceph/osd/ceph-25
| /dev/sdl3 |
| | osd.30 | up | 30 | 6 (JBOD) | /dev/sdi1 | /var/lib/ceph/osd/ceph-30
| /dev/sdk4 |
| | osd.34 | up | 34 | 9 (JBOD) | /dev/sdj1 | /var/lib/ceph/osd/ceph-34
| /dev/sdl4 |
| | | | | | |
| | | | | | |
| cork-store-6.cisco.com | osd.6 | up | 6 | 9 (JBOD) | /dev/sdc1 | /var/lib/ceph/osd/ceph-6
| /dev/sdk1 |
| | osd.9 | up | 9 | 7 (JBOD) | /dev/sdd1 | /var/lib/ceph/osd/ceph-9
| /dev/sdl1 |
| | osd.13 | up | 13 | 3 (JBOD) | /dev/sde1 | /var/lib/ceph/osd/ceph-13
| /dev/sdk2 |
| | osd.17 | up | 17 | 10 (JBOD) | /dev/sdf1 | /var/lib/ceph/osd/ceph-17
| /dev/sdl2 |

```

/dev/sdk3	osd.21	up	21	4 (JBOD)	/dev/sdg1	/var/lib/ceph/osd/ceph-21
/dev/sdl3	osd.26	up	26	5 (JBOD)	/dev/sdh1	/var/lib/ceph/osd/ceph-26
/dev/sdk4	osd.29	up	29	8 (JBOD)	/dev/sdi1	/var/lib/ceph/osd/ceph-29
Missing	osd.33	down	33	6 (UGood)	Missing	Missing
cork-store-4.cisco.com	osd.4	up	4	6 (JBOD)	/dev/sdb1	/var/lib/ceph/osd/ceph-4
/dev/sdj1	osd.10	up	10	7 (JBOD)	/dev/sdc1	/var/lib/ceph/osd/ceph-10
/dev/sdk1	osd.15	up	15	4 (JBOD)	/dev/sdd1	/var/lib/ceph/osd/ceph-15
/dev/sdj2	osd.18	up	18	8 (JBOD)	/dev/sde1	/var/lib/ceph/osd/ceph-18
/dev/sdk2	osd.23	up	23	10 (JBOD)	/dev/sdf1	/var/lib/ceph/osd/ceph-23
/dev/sdj3	osd.27	up	27	9 (JBOD)	/dev/sdg1	/var/lib/ceph/osd/ceph-27
/dev/sdk3	osd.31	up	31	3 (JBOD)	/dev/sdh1	/var/lib/ceph/osd/ceph-31
/dev/sdj4	osd.35	up	35	5 (JBOD)	/dev/sdi1	/var/lib/ceph/osd/ceph-35
/dev/sdk4						

```
[root@cork-mgmt tools]# ciscovim osdmgmt --check_osds --server cork-store-4.cisco.com
```

Perform the action. Continue (Y/N)Y

Message	Host	Role	Server	State
Overall OSD Status	cork-store-4.cisco.com	block_storage	8.8.8.17	Optimal
Number of OSDs	cork-store-4.cisco.com	block_storage	8.8.8.17	8

```
[root@cork-mgmt tools]# ciscovim osdmgmt --check_osds --server cork-store-4.cisco.com --detail
```

Perform the action. Continue (Y/N)Y

Message	Host	Role	Server	State
Overall OSD Status	cork-store-5.cisco.com	block_storage	8.8.8.19	Optimal
Number of OSDs	cork-store-5.cisco.com	block_storage	8.8.8.19	8

```

| Host          | OSDs  | Status | ID | HDD Slot | Path          | Mount
| Journal      |
+-----+-----+-----+-----+-----+-----+-----+
| cork-store-5.cisco.com | osd.5  | up    | 5  | 8 (JBOD) | /dev/sdc1    | /var/lib/ceph/osd/ceph-5
| /dev/sdk1    |
|              | osd.11 | up    | 11 | 3 (JBOD) | /dev/sdd1    | /var/lib/ceph/osd/ceph-11
| /dev/sdl1    |
|              | osd.14 | up    | 14 | 10 (JBOD) | /dev/sde1    | /var/lib/ceph/osd/ceph-14
| /dev/sdk2    |
|              | osd.19 | up    | 19 | 5 (JBOD) | /dev/sdf1    | /var/lib/ceph/osd/ceph-19
| /dev/sdl2    |
|              | osd.22 | up    | 22 | 7 (JBOD) | /dev/sdg1    | /var/lib/ceph/osd/ceph-22
| /dev/sdk3    |
|              | osd.25 | up    | 25 | 4 (JBOD) | /dev/sdh1    | /var/lib/ceph/osd/ceph-25
| /dev/sdl3    |
|              | osd.30 | up    | 30 | 6 (JBOD) | /dev/sdi1    | /var/lib/ceph/osd/ceph-30
| /dev/sdk4    |
|              | osd.34 | up    | 34 | 9 (JBOD) | /dev/sdj1    | /var/lib/ceph/osd/ceph-34
| /dev/sdl4    |
+-----+-----+-----+-----+-----+-----+-----+

[root@cork-mgmt ~]# ciscovim osdmgmt --replace_osd --server cork-store-6.cisco.com --osd osd.33

Perform the action. Continue (Y/N)Y

Successfully deleted, removed and replaced OSD osd.33 from server cork-store-6.cisco.com
+-----+-----+-----+-----+-----+-----+-----+
| Host | New device UUID | New OSD.id | HDD Slot | Path | Mount | Journal |
+-----+-----+-----+-----+-----+-----+-----+
| cork-store-6.cisco.com | UUID=297e4aac-ad1e-4755-9e90-9e3b79c86824 | 33 | 6 | /dev/sdj1 | /var/lib/ceph/osd/ceph-33 | /dev/sdl4 |
+-----+-----+-----+-----+-----+-----+-----+

```

Utility to resolve Cisco VIM Hardware Validation Failures

The Cisco VIM Hardware Validation utility tool is used to perform hardware validation during the installation of UCS C-series servers. It captures the user and environmental hardware validation errors that occur during the installation process. The tool enables you to fix these errors based on the inputs you provide in the Command Line Interface (CLI). It validates the updated configurations to verify if the changes are applied properly. After the error is resolved, you can resume the installation from the point of failure.



Note

You cannot use the utility for the following tasks:

- Configuring BIOS settings for the B-series pods.
- Upgrading or changing the firmware version.
- Resolving hardware failures other than lom, hba, flexflash, pcie_slot, power and vnic_pxe_boot.

Command Usage

To capture the list of failures that can be resolved by using the utility, go to the install directory and execute the help command as shown below:

```
# cd <installer-id>/clouddeploy
# python hw_validations.py -help .
```

The following shows the output of the help command.

```
usage: hw_validations.py [-h] [--resolve-failures RESOLVE_FAILURES]
[--validate VALIDATE_OF] [-y] [--host HOSTS]
[--file SETUP_FILE_LOCATION]
UCS Hardware Validations
optional arguments:
-h, --help show this help message and exit
--resolve-failures RESOLVE_FAILURES, -rf RESOLVE_FAILURES
    all - Fix all the failures.
    lom - Fix LOM port(s) status failures.
    hba - Fix HBA port status failures.
    flexflash - Fix Flexflash failures.
    pcie_slot - Fix PCIe slot status failures.
    power - Fix Power failures.
    vnic_pxe_boot - Fix Vnic PXE_Boot statusfailures
-y, -yes
--host HOSTS Comma separated list of hostnames
--file SETUP_FILE_LOCATION, -f SETUP_FILE_LOCATION
    Provide a valid 'setup_data.yaml' file
```

Command Syntax

```
hw_validations.py [-h] [--resolve-failures RESOLVE_FAILURES] [--validate VALIDATE_OF] [-y]
[--host HOSTS] [--file SETUP_FILE_LOCATION]
```

The following table provides the description of the parameters of the command.

Optional	Description
[-h], --help	Provides detailed information about the command.
[--resolve-failures RESOLVE_FAILURES], -rf RESOLVE_FAILURES	Enables you to specify the failure that you want to resolve. The optional arguments are as follows:
[-y]	Yes

Optional	Description
[--host HOSTS]	Enables you to specify the hostname of the server for which you want to resolve failures. You cannot specify ip address or cimc ip address of servers as arguments. You can specify a list of hostnames as comma separated arguments. If the -host option is not specified, the failures of all the servers that are specified in the setup_data.yaml file are resolved.
[--file SETUP_FILE_LOCATION] [-f SETUP_FILE_LOCATION]	Enables you to specify the name of a setup_data.yaml file.

Examples of Command Usage

The following table provides the commonly used commands along with their examples.

Purpose	Syntax	Example
To resolve all failures.	<code>python hw_validations.py --resolve-failures all -y</code>	<code>python hw_validations.py --resolve-failures all -y</code>
To simultaneously resolve one or more failures.	<code>python hw_validations.py --resolve-failures <failure-1>,<failure-2> -y</code>	To resolve the lom and hba status failures: <code>python hw_validations.py --resolve-failures lom,hba -y</code>
To resolve the errors by using the setup_data.yaml file.	<code>python hw_validations.py --resolve-failures <failure-1>,<failure-2> -y --file <location-of-yaml file></code>	To resolve the LOM status failures by using ~/save/setup_data.yaml file: <code>python hw_validations.py --resolve-failures lom,hba -y --file ~/save/setup_data.yaml</code>
To resolve failures on a particular server as specified in the setup_data.yaml file by using the -- host option.	<code>python hw_validations.py --resolve-failures <failure-1> -y --host <name of host server-1>,<name of host server-2></code>	To resolve the PCIe slot failures on hiccup-controller-1 server as specified in the setup_data.yaml: <code>python hw_validations.py --resolve-failures pcie_slot -y --host hiccup-controller-1</code>