Hardware and Software
ORACLE
Engineered to Work Together

Oracle® Communications

Policy Management

Cloud Disaster Recovery 12.2

E82616-01
March 2017

CAUTION: In the event of a critical service situation, emergency response is offered by the Customer Access Support (CAS) main number at 1-800-223-1711 (toll-free in the US), or by calling the Oracle Support hotline for your local country from the list at *http://www.oracle.com/us/support/contact/index.html*. The emergency response provides immediate coverage, automatic escalation, and other features to ensure that the critical situation is resolved as rapidly as possible.

Contact Call the Oracle Customer Access Support Center at 1-800-223-1711 prior to executing this procedure to ensure that the proper recovery planning is performed.

Before disaster recovery, users must properly evaluate the outage scenario. This check ensures that the correct procedures are executed for the recovery.

****** WARNING *******

NOTE*:* DISASTER Recovery is an exercise that requires collaboration of multiple groups and is expected to be coordinated by the TAC prime. Based on TAC's assessment of Disaster, it may be necessary to deviate from the documented process.

EMAIL: support@oracle.com

Table of Contents

List of Tables

1. Introduction

1.1 *Purpose and Scope*

This document is a guide to describe procedures used to execute disaster recovery for Policy Management System, Release 12.2. This includes recovery of partial or a complete loss of one or more policy servers and policy components. This document provides step-by-step instructions to execute disaster recovery for Policy Management Systems. Executing this procedure also involves referring to and executing procedures in existing support documents.

1.2 *References*

[1] E82614-01 - Oracle Communications Policy Management Cloud Installation Guide 12.2

The above documents are available at the Oracle Help Center.

1.3 *Acronyms*

Acronym	Meaning
BIOS	Basic Input Output System
BOD	Bandwidth On Demand
CD	Compact Disk
ISO	The name *ISO* is taken from the ISO 9660 file system used with CD-ROM media, but an ISO image might also contain a UDF (ISO/IEC 13346) file system
c-Class	HP marketing term for their enterprise blade server platform
CMP	Configuration Management Platform
DR-CMP	Configuration Management Product for Disaster Recovery NOTE: It refers to the CMP on the secondary site
DVD	Digital Video Disc
GRUB	Grand Unified Boot loader
iLO	Integrated Lights-Out
IPM	Initial Product Manufacture – the process of installing TPD on a hardware platform
MPE	Multiprotocol Policy Engine
MRA	Multiprotocol Routing Agent
MA	Management Agent
OS	Operating System (e.g. TPD)
PM&C	Platform Management & Configuration
RMM	Remote Management Module
RMS	Rack Mount Server
SOL	Serial Over LAN
TPD	Tekelec Platform Distribution
TVOE	Tekelec Virtualization Operating Environment
FRU	Field Replaceable Unit
USB	Universal Serial Bus
VM	Virtual Machine

Table 1: Acronyms

1.4 *Logins and Passwords*

The standard configuration steps will configure standard passwords for root, admusr, admin, and some other standard logins referenced in this procedure. Please note that SSH to Policy servers as root user is restricted, but allowed using ‘admusr’ user. These passwords are not included in this document.

1.5 *Software Release Numbering*

This guide applies to all Policy Management versions 12.2.

1.6 *Terminology*

Base software	Base software includes deploying the VM image.
Failed server	A failed server in disaster recovery context refers to a server that has suffered partial or complete software and/or hardware failure to the extent that it cannot restart or be returned to normal operation and requires intrusive activities to re-install the software and/or hardware.
Perform initial configuration	The perform initial configuration put into the policy server through the platcfg utility that brings the server's network interface online and allows management and configuration from the CMP

Table 2. Terminology

2. General Description

The Policy Management disaster recovery procedure falls into two basic categories. It is primarily dependent on the state of the CMP VM instances:

- Recovery of one or more VM instance with at least one CMP VM instance intact
  - 1 or more CMP VM instances intact (this can include Geo-Redundant CMP(DR-CMP) VM instances)
  - 1 or more MPE/MRA/BOD/MA instances failed
- Recovery of the entire network from a total outage
  - No CMP instances are available (neither primary, nor secondary) and other MPE/MRA/BOD/MA VM instances will need to be recovered

The existence of Geo-redundant VM instances, including a geo-redundant CMP (DR-CMP) VM instance can mitigate massive outages by providing a running manager from which to synchronize new VM instances as they are restored.

No matter the number of VM instances involved in the outage, the key to the severity is the status of the CMP. The availability of regular system backups of the CMP are critical when all CMP VM instances are offline and must be restored.

Single node outage MRA/MPE/BOD/MA/CMP, with CMP VM instance available

The simplest case of recovery is to recover a single node of a cluster with one or both CMP VM instances intact. Each failed VM instance is recovered by:

- creating a new VM instance using as described in [1]
- performing the initial configuration of the VM instance manually or from a server backup file

After this recovery, the cluster will reform, and database replication from the active node of the cluster will recover the newly restored VM instance. This scenario can be used to recover one VM instance of a MRA/MPE/BOD/MA cluster or one VM instance of a CMP cluster. The SSH exchange keys with cluster mate from active CMP is also required.

Recovery of complete MRA/MPE/BOD/MA cluster, with CMP VM instance available

The failure of a complete cluster can be recovered by creating new VM instances for the failed cluster. Each failed VM instance is recovered by:

- creating a new VM instance using as described in [1]
- performing the initial configuration of the VM instance manually or from a server backup file

After this recovery, the CMP can push application level configuration to the newly restored cluster.

Recovery of the CMP Cluster when no geo-redundant CMP exists

The complete failure of all CMP VM instances when no geo-redundant CMP exists is recovered by:

- creating a new VM instance using as described in [1]
- performing the initial configuration of the CMP VM instances manually or from a server backup file

Once the cluster is available, completion of the recovery will require the use of a stored system backup in order to recover application level configuration including policies and configuration of the MPE/MRA/BOD/MA clusters in the network.

Recovery of the CMP Cluster when geo-redundant CMP (DR-CMP) is available

The availability of a geo-redundant CMP (DR-CMP) will simplify restoration of a failed CMP cluster. The geo-redundant CMP will be promoted to active primary, and the failed CMP VM instances are recovered by:

- creating a new VM instance using as described in [1]
- performing the initial configuration of the CMP VM instances manually or from a server backup file

Once the cluster is available, the primary running geo-redundant CMP will replicate databases to the replaced CMP cluster.

Complete Outage (All VM instances)

This is the worst case scenario where all the VM instances in the network have suffered complete failure, and no geo-redundant CMP is available. Each VM instance in the network is recovered by:

- creating a new VM instance using as described in [1]
- performing the initial configuration of the VM instance manually or from a server backup file

Once the VM instances are installed and available, completion of the recovery will require restoration of a stored system backup in order to recover the application level configuration including policies and configuration of the MPE/MRA/BOD/MA clusters in the network.

If no backup file is available, the only option is to rebuild the entire network from scratch in accordance with [1]. The network data must be reconstructed from whatever sources are available, including entering all data manually.

A note on performing the initial configuration of a VM instance:

The information required for initial configuration is not extensive, and may be readily available from customer site documents, or from the CMP's topology configuration. In some cases it can be easier to manually input the 'initial configuration' in platcfg than to try to load a server backup file into the newly installed hardware.

Needed initial configuration information:

- Hostname
- OAM real IP address and network mask
- OAM default router address
- NTP server
- DNS server (optional)
- DNS search (optional)
- Interface device (usually bond0)
- VLAN configuration for c-Class and Sun Netra systems.

Using the server backup file

When asked to restore from ‘server backup’, the platcfg utility will look in /var/camiant/backup/local-archive/serverbackup directory. If no files are in that directory, the box below will be presented.

You will have to enter the complete path and filename in order to restore from a file that is not in the /var/camiant/backup/local-archive/serverbackup directory.

Using the system restore file

When asked to restore from ‘system backup’, the platcfg utility will look in /var/camiant/backup/local-archive/systembackup directory. If no files are in that directory, the box below will be presented.

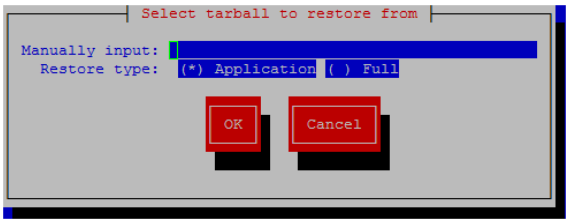

You will have to enter the complete path and filename in order to restore from a file that is not in the /var/camiant/backup/local-archive/systembackup directory.

3. Procedure Overview

This section lists the materials required to perform disaster recovery procedures and a general overview (disaster recovery strategy) of the procedure executed.

Disaster Recovery Strategy

Disaster recovery procedure execution is performed as part of a disaster recovery strategy with the basic steps listed below:

1. Evaluate failure conditions in the network and determine that normal operations cannot continue without disaster recovery procedures. This means the failure conditions in the network match one of the failure scenarios described in Recovery Scenarios
2. Evaluate the availability of server and system backup files for the servers that are to be restored.
3. Read and review the content in this document.
4. Determine whether a geo-redundant CMP(DR-CMP) is available
5. From the failure conditions, determine the Recovery Scenario and procedure to follow.
6. Execute appropriate recovery procedures.

Required materials

The following items are needed for disaster recovery:

1. A copy of this document and copies of all documents in the reference list.
2. Copy of all site surveys performed at the initial installation and network configuration of the customer's site. If the site surveys cannot be found, escalate this issue within Oracle CGBU Customer Service until the site survey documents can be located.
3. Policy 'System' backup file: electronic backup file (preferred) or hardcopy of all Policy system configuration and provisioning data.
4. Policy Application installation : OVA for CMP, MPE, MRA, BoD and MA of the target release.

Policy server backup

Backup of the policy server can be done either manually from platcfg, or on a schedule as configured in platcfg. There are 2 types of backup operations available; '*server backup'* and '*system backup'*:

- **Server Backup:** There is one Server Configuration backup for each server in the system. The server backup is a Back-up of the OS information unique to the server. Information includes hostname, IP Addresses, NTP, DNS, Static Route configuration. This operation create a Server Configuration Backup file, and should be executed on each of the server in the customer's network.

- **System Backup**: There is one Application Configuration backup for the entire Policy system. The system backup will gather PCRF configuration information that is unique to this system. Information such as: Topology, Policy(s), Feature Configuration. The system backup is executed only on the Active CMP at the primary site.

The availability of a recent system backup is critical to the restoration of the policy network when the CMP is not available.

4. Procedure Preparation

4.1 *Purpose and Scope*

Disaster recovery procedure execution is dependent on the failure conditions in the network. The severity of the failure determines the recovery scenario for the network. The first step is to evaluate the failure scenario and determine the procedure(s) that will be needed to restore operations. A series of procedures are included below that can be combined to recover one or more policy management nodes or clusters in the network.

Note: A failed VM instance in disaster recovery context refers to a VM instance that is no longer available to be restored. Examples of scenarios where this can happen are: host server failure and user deletion of VM instance.

The general steps recovering VM instances are:

1. Create a new VM as described in [1]
2. Perform the initial configuration of the VM or restore the initial configuration from a server backup file
3. Check NTP status after recovery
4. Check Active Alarms from GUI and both syscheck, alarmMgr --alarmStatus from CLI

4.2 *Recovery Scenarios*

4.2.1 Recovery Scenario 1 (Partial Cluster Outage with Primary CMP VM instance available)

A single CMP VM instance is capable of restoring the configuration database via replication to all MPE/MRA/BOD/MA servers, or to the other CMP node of a cluster. The major activities are summarized in the list below. Use this list to understand the recovery procedure summary. Do not use this list to execute the procedure. The actual procedures' detailed steps are in the Restore Procedures section. The major activities are summarized as follows:

- Recover Standby CMP VM instance (if necessary)
    - Create a new CMP VM instance
    - Initial Policy configuration is re-installed, either through the platcfg menu, or from the server backup file
    - The database is intact at the active CMP VM instance and will be replicated to the standby CMP VM instance.

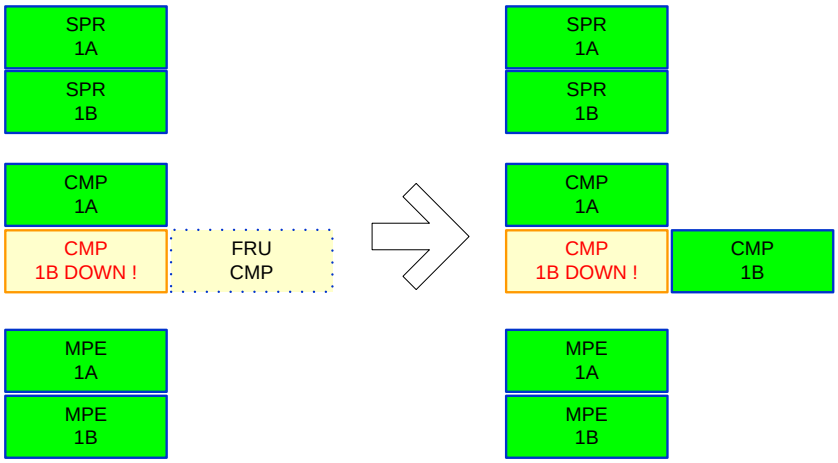

- Recover any failed MPE/MRA/BOD/MA servers by:
    - Create a new MPE/MRA/BOD/MA VM instance
    - Initial Policy configuration is re-installed, either through the platcfg menu, or from the server backup file
    - The configuration database is available at the active CMP VM instance and does not require restoration on the CMP. Configuration can be pushed from the CMP to the MPE/MRA/BOD/MA VMs using 're-apply configuration'

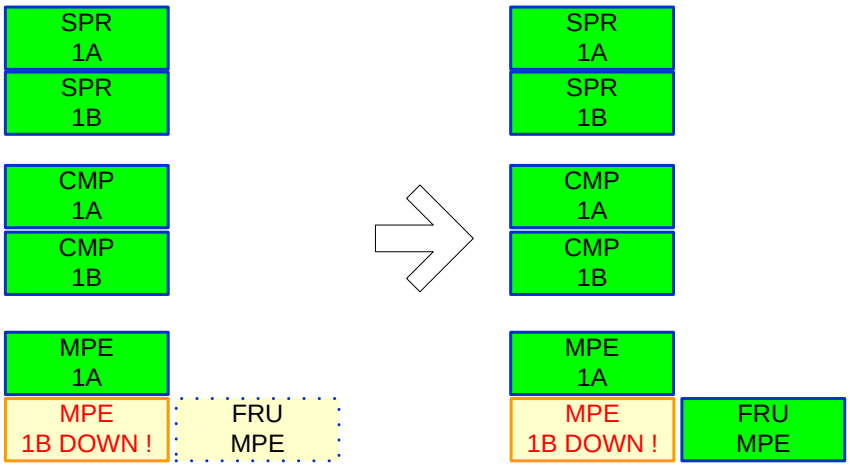

Follow the procedure below for detailed steps.

- Use Procedure 2: Restore standby CMP Node without server backup file
  Or Procedure 1: Restore standby CMP Node with server backup file to recover the second CMP node if necessary.
- Use Procedure 4: Restore single MPE/MRA/BOD/MA node without server backup file to recover MPE / MRA/BOD/MA nodes when one of the peers of the cluster is still available.
  Or Procedure 3: Restore single MPE/MRA/BOD/MA node with server backup file
- Use Procedure 5: Restoring complete cluster with the server backup files
  Or Procedure 6: Restoring complete cluster without the server backup to recover complete MPE / MRA/BOD clusters that have gone down.
- Use Procedure 7: Restoring CMP/MA cluster with system backup available files to recover first of 2 nodes in MA cluster
  Use Procedure 3: Restore single MPE/MRA/BOD/MA node with server backup file to recover the second node of MA cluster.

4.2.2 Recovery Scenario 2 (Partial Cluster Outage with geo-redundant CMP Server available)

For a partial outage with a geo-redundant CMP VM instance available, the secondary site CMP must be manually promoted to Primary status as the controlling CMP for the policy network. Then creation of new CMP VM instance and initial Policy configuration is needed. The now active CMP VM instance is capable of restoring the configuration database via replication to all MPE/MRA/BOD/MA servers, and to the other CMP cluster. The major activities are summarized in the list below. Use this list to understand the recovery procedure summary. Do not use this list to execute the procedure. The actual procedures' detailed steps are in the `Restore Procedures` section. The major activities are summarized as follows:

- Promote the geo-redundant CMP VM instance.
  - This step is done by logging into the OAM VIP address of the second site CMP cluster. Use procedure 8 below.

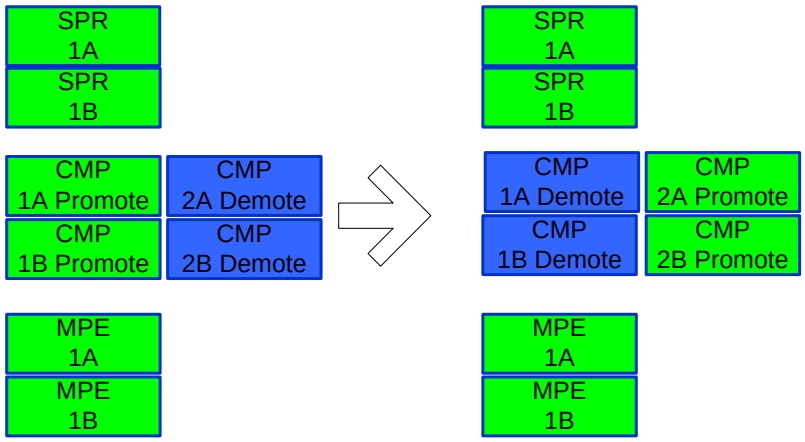

This would only need to be done if the Primary CMP cluster needs to be restored. If it's an MRA, MPE, BOD, MA, or secondary CMP cluster that needs to be restored, there is no need to promote the Geo CMP.

- Recover any failed MPE/MRA/BOD/MA VM instances by:
  - creating a new VM instance for the failed MPE/MRA/BOD/MA
  - Initial Policy configuration is re-installed, either through the platcfg menu, or from the server backup file
  - The configuration database is available at the active CMP VM instance and does not require restoration on the CMP. Configuration can be pushed from the CMP to the MPE/MRA/BOD/MA VM instances using 're-apply configuration'
- Recover other site CMP VM instance by:
  - creating a new CMP VM instance
  - Initial Policy configuration is re-installed, either through the platcfg menu, or from the server backup file.

The database of the active geo-redundant CMP VM instance will be replicated to the new CMP VM instance.

Follow the procedure below for detailed steps.

- Use Procedure 8: Promoting geo-redundant CMP cluster below to promote the geo-redundant CMP
- Use Procedure 4: Restore single MPE/MRA/BOD/MA node without server backup file to recover MPE / MRA / BOD / MA nodes when one of the peers of the cluster is still available.

  Or Procedure 3: Restore single MPE/MRA/BOD/MA node with server backup file
- Use Procedure 5: Restoring complete cluster with the server backup files

  Or Procedure 6: Restoring complete cluster without the server backup to recover complete MPE / MRA / BOD clusters that have gone down.
- Use Procedure 5: Restoring complete cluster with the server backup files

  Or Procedure 6: Restoring complete cluster without the server backup to recover the secondary site CMP. Recovery of the secondary site CMP can be left for late in the process because the now active CMP can handle all application level configuration as the network is brought back online.
- Use Procedure 7: Restoring CMP/MA cluster with system backup available files to recover first of 2 nodes in MA cluster

  Use Procedure 3: Restore single MPE/MRA/BOD/MA node with server backup file to recover the second node of MA cluster.

4.2.3 Recovery Scenario 3 (Full cluster outage of the CMP; geo-redundancy not available; other VM instances as needed)

For a full outage with a CMP VM instance unavailable, creation of new CMP VM instances is needed, then the recovery from system backup of the application configuration for the policy network. The first CMP VM instance is built and restored with the configuration database from a system backup. Replication of the restored database to a second rebuilt CMP node will form a CMP cluster. The major activities are summarized in the list below. Use this list to understand the recovery procedure summary. Do not use this list to execute the procedure. The actual procedures' detailed steps are in the `Restore Procedures section`. The major activities are summarized as follows:

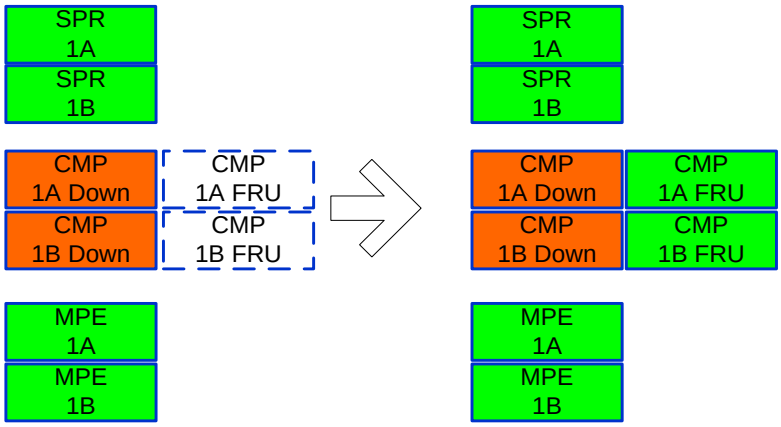

- Recover one Primary CMP VM instance (if necessary) by:
  - creating a new CMP VM instance
  - Recover the software.
  - Initial Policy configuration is re-installed, either through the platcfg menu, or from the server backup file.
  - The database of the CMP will be restored from a system backup provided by the customer.
  - If a system backup is not available, use customer site survey, and site installation documentation to restore application level configuration to the CMP. It is possible to use the data at the MPEs (that should still be good) to verify that the re-entered data on the CMPs matches the previous configuration that was in-use. Also, check with engineering team for possible approach to verify if the data at the operational MPEs matches the data that has been re-entered at the CMP after re-entering the Policies and other application level data to the CMP.
- Recover the second CMP VM instance by:
  - creating a new CMP VM instance
  - Initial Policy configuration is re-installed, either through the platcfg menu, or from the server backup file
  - The configuration database is available at the now active CMP VM instance and does not require restoration on the second CMP node. Configuration will be replicated when the two new CMP nodes form a cluster.
- Recover any failed MPE/MRA/BOD/MA VM instances by:
  - creating a new MPE/MRA/BOD/MA VM instance for the failed VM instance
  - Initial Policy configuration is re-installed, either through the platcfg menu, or from the server backup file

- The configuration database is available at the now active CMP VM instance and does not require restoration on the CMP. Configuration can be pushed from the CMP to the MPE/MRA/BOD/MA VM instances.

Follow the procedure below for detailed steps.

- Use Procedure 7: Restoring CMP/MA cluster with system backup available below to recover the first of 2 nodes in the CMP cluster.
- Use Procedure 2: Restore standby CMP Node below to recover the second node of the CMP cluster
- Use Procedure 4: Restore single MPE/MRA/BOD/MA node without server backup file to recover MPE/MRA/BOD/MA nodes when one of the peers of the cluster is still available.

  Or Procedure 3: Restore single MPE/MRA/BOD/MA node with server backup file
- Use Procedure 5: Restoring complete cluster with the server backup files

  Or Procedure 6: Restoring complete cluster without the server backup to recover complete MPE/MRA/BOD clusters that have gone down.
- Use Procedure 7: Restoring CMP/MA cluster with system backup available files to recover first of 2 nodes in MA cluster

  Use Procedure 3: Restore single MPE/MRA/BOD/MA node with server backup file to recover the second node of MA cluster.

5. Restore Procedures

5.2 *Procedure 1: Restore standby CMP Node with server backup file*

The purpose of this procedure is to replace one node of a CMP cluster. Restore initial Policy configuration from a server backup file, and then allow the new node to re-sync to the existing node to form a complete CMP cluster. In this example, initial Policy configuration is restored to the new nodes through the use of server backup files for each server to be restored.

Required resources:

- Host server identified for the new VM instance
- OVA file or equivalent (depending on hypervisor or NFV manager)
- *serverbackup*.ISO of the node to be replaced

Prerequisites:

- failed VM is no longer available (e.g. it has been removed from the hypervisor/NFV manager)
- a new VM has been created in accordance with [1].

S T E P #	This Procedure restores the standby CMP node when a server level backup is available. Check off (√) each step as it is completed. Boxes have been provided for this purpose under each step number. Should this procedure fail, contact the My Oracle Support (MOS) Customer Care Center and ask for assistance.	
1. ☐	Set the failed node to 'Forced Standby'	In the CMP GUI, navigate to: **Platform Setting → Topology Settings → All Clusters** 1. Determine the cluster with the failed node 2. Determine the failed node 3. Click the Modify Server-X for the failed node 4. Click the Forced Standby checkbox so that it is checked, then click Save 
2. ☐	Create the VM instance	Create the new VM instance in accordance with [1]
3. ☐	Load the ISO for server restore	Obtain the **serverbackup.iso** for the node to be restored. The server backup file should be copied via secure copy(pscp,scp, or WinSCP) to the following directory: **`/var/camiant/backup/local_archive/serverbackup`** ***Note:*** *Later in this procedure, the platcfg restore function check this directory and offer the user a convenient menu to choose from. The platcfg utility also allows the user to manually enter any mounted path on the server.*

4. ☐	Login via SSH to new node	SSH to the new VM instance:
		`# ssh admusr@<node_IP_Address>`
		`$ sudo su -`

5. ☐	Perform platcfg restore from SSH session to replacement node	Execute the following command: `# su - platcfg` From within the platcfg utility, navigate to: **Policy Configuration → Backup and Restore → Server Restore** Select the *serverbackup*.ISO that you just put on the system and hit OK, then 'Yes' to confirm. 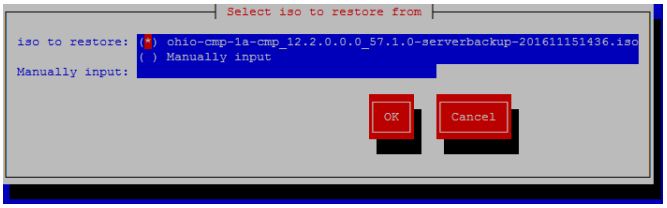 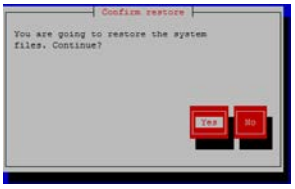
6. ☐	Verify the status	A window will pop-up, indicating restore operation was successful and will ask the user to press any key to exit. If it is not successful, retry the restore. If the second restore is not successful, stop and contact My Oracle Support or engineering team for assistance.

7. ☐	Perform Initial configuration	Choose Exit repeatedly until back to the Main Menu of the *platcfg* utility. While still within the *platcfg* utility, navigate to: **Policy Configuration → Verify Initial Configuration** 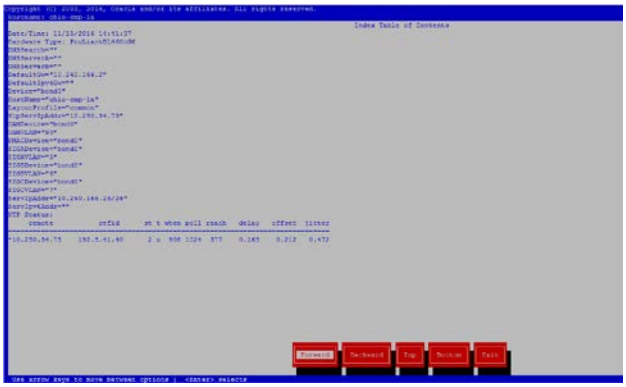 If the configuration does not exist, then navigate to '**Perform Initial Configuration'** and fill in the hostname, OAM IP and configuration as shown below: 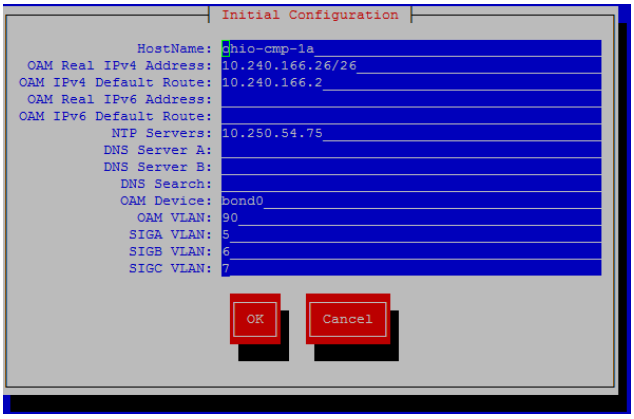 Ensure that your data is correct, and select 'Ok', then 'yes' to save and apply **`Exit platcfg`** Exit platcfg by selecting **Exit** from each platcfg menu until you are returned to the shell. 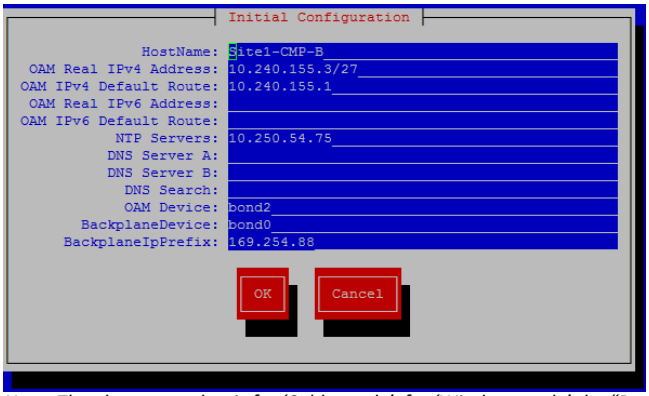 *Note: The above snapshot is for 'Cable mode', for 'Wireless mode' the "Backplane Device" and "Backplane IP Prefix" parameters will not exist.*
8. ☐	Reboot the server	Reboot: `# init 6` Allow the server time to reboot;

9. ☐	Verify basic network connectivity and server health.	From the newly installed VM, ping the OAM/XMI gateway. If the ping is not successful, verify all network settings match the old hardware configuration and reconfigure if needed. Contact My Oracle Support before proceeding if network ping tests still fail.
		`# ping <XMI or OAM gateway address>`
		
		Execute the `syscheck` command, ensuring that all tests return successfully. If errors are found, discontinue this procedure and contact My Oracle Support.
		`[root@ohio-cmp-1a ~]# syscheck`
		`Running modules in class disk...`
		`OK`
		`Running modules in class hardware...`
		`OK`
		`Running modules in class net...`
		`OK`
		`Running modules in class proc...`
		`OK`
		`Running modules in class system...`
		`OK`
		`LOG LOCATION: /var/TKLC/log/syscheck/fail_log`
		`[root@ohio-cmp-1a ~]#`

10. ☐	Remove 'Forced Standby' designation on current node.	In the CMP GUI, navigate to: **Platform Setting → Topology Settings → All Clusters → Current Cluster** 1. Modify for the server that has 'Forced Standby' 2. Clear the Forced Standby checkbox 3. Click Save  Accept the resulting pop-up by clicking OK: 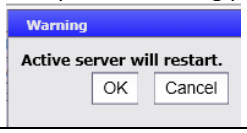
11. ☐	Verify cluster status	In the CMP GUI, navigate to: **Platform Setting → Topology Setting → All Clusters→ Current CMP Cluster** Monitor clustering of the new node to its peer, do not proceed until both nodes have a status of either 'active' or 'standby', and that there are no CMP related 'Active Alarms' as shown below. 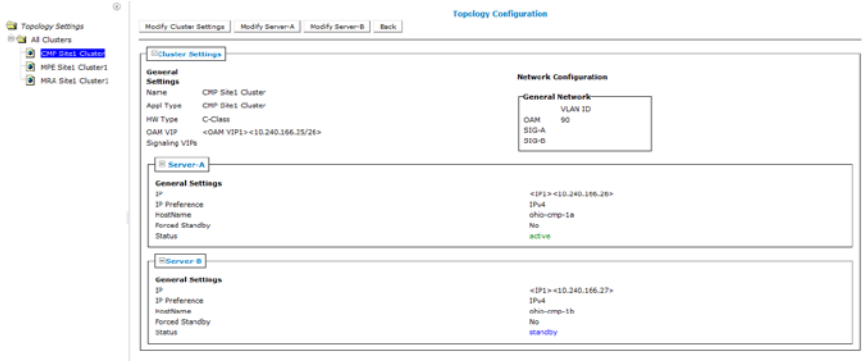

12. ☐	Alternative method to check replication status	You can also monitor the clustering of the new node from within the shell on the primary node with 'irepstat'. To do so, SSH to the Active node of the current cluster and execute the irepstat command: `# irepstat` Expected 'irepstat' output while waiting reconnection: 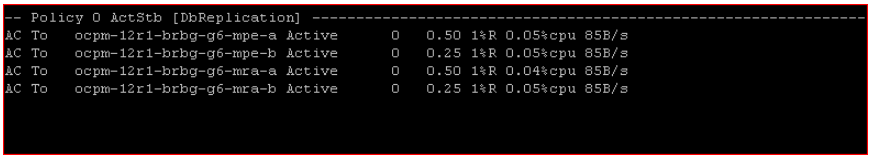 Expected 'irepstat' output after cluster has formed: 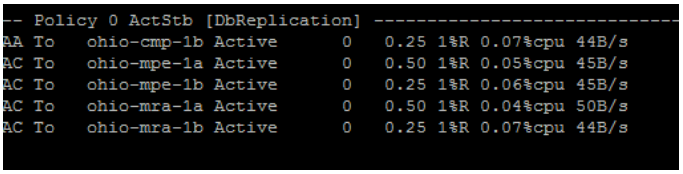
13. ☐	Exchange keys with cluster mate (This step needs to run from the active CMP)	Exchanging SSH keys Utility as `root`, please run `'/opt/camiant/bin/qpSSHKeyProv.pl -prov -user=root'`; as `admusr`, please run `'/usr/bin/sudo /opt/camiant/bin/qpSSHKeyProv.pl -prov'` 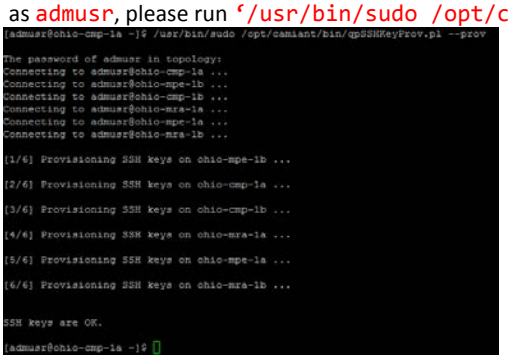 This procedure is completed

5.3 *Procedure 2: Restore standby CMP Node without server backup file*

The purpose of this procedure is to replace one node of a CMP cluster. Restore initial Policy configuration using platcfg's 'Perform Initial Configuration', and then allow the new node to re-sync to the existing node to form a complete CMP cluster. In this example, initial Policy configuration is restored to the new nodes through the use of platcfg's 'Perform Initial Configuration' menu for each server to be restored.

Required resources:

- Host server identified for the new VM instance
- OVA file or equivalent (depending on hypervisor or NFV manager)
- Node IP addresses, VLANs, NTP IP address, and hostname from CMP GUI

Prerequisites:

- failed VM is no longer available (e.g. it has been removed from the hypervisor/NFV manager)
- a new VM has been created in accordance with [1].

S T E P #	This Procedure restores the standby CMP node when a server level backup file is not available. Check off (√) each step as it is completed. Boxes have been provided for this purpose under each step number. Should this procedure fail, contact the My Oracle Support (MOS) Customer Care Center and ask for assistance.	
1. ☐	Set the failed node to 'Forced Standby'	In the CMP GUI, navigate to: **Platform Setting → Topology Setting → All Clusters** 1. Determine the cluster with the failed node 2. Determine the failed node 3. Click the Modify Server-X for the failed node 4. Click the Forced Standby checkbox so that it is checked, then click Save 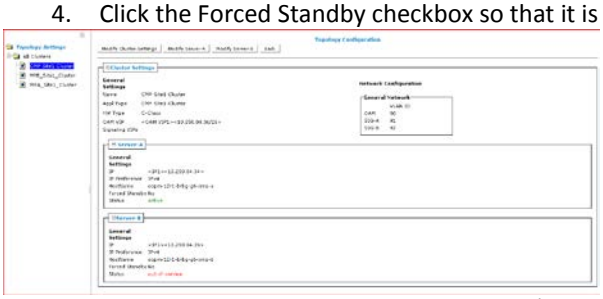 *Note: From the above screenshot, the Network Configuration/General Network(VLAN ID) will not appear for RMS (DL 360/ DL380) Hardware*
2. ☐	Create the VM instance	Create the new VM instance according to [1]

3. ☐	Login via SSH to new node	SSH session to the new VM instance:
		`# ssh admusr@<node_IP_Address>`
		`$ sudo su -`

4. ☐	Perform platcfg restore from SSH session to replacement node Perform Initial configuration	Execute the following command `# su - platcfg` From within the platcfg utility, navigate to: **Policy Configuration → Perform Initial Configuration** Enter the appropriate configuration details for this node, verify that entries are correct, and select 'OK' to continue. Accept the resulting popup that appears asking to apply the configuration. Once the operation is complete, select 'Exit' on the platcfg menu until you are dropped back to the shell. 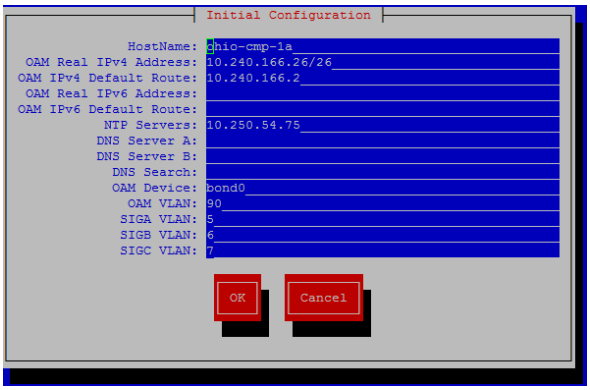 Ensure that configured data is correct, and select 'OK', then 'yes' to save and apply **Exit platcfg** Exit platcfg by selecting **Exit** from each platcfg menu until you are returned to the shell. 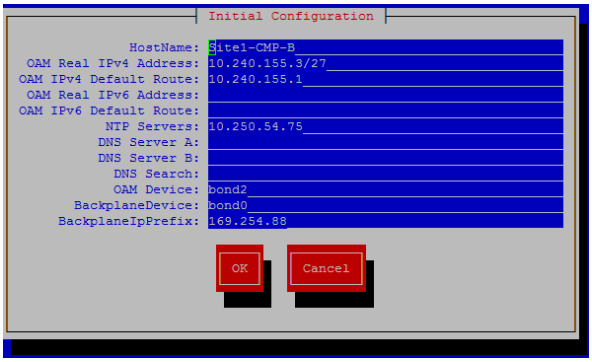 *Note: The above snapshot is for 'Cable mode', for 'Wireless mode' the "Backplane Device" and "Backplane IP Prefix" parameters will not exist.*
5. ☐	Reboot the server	Reboot: `# init 6` Allow the server time to reboot;

6. ☐	Verify basic network connectivity and server health.	From the newly installed VM, ping the OAM/XMI gateway. If the ping is not successful, verify all network settings match the old hardware configuration and reconfigure if needed. Contact My Oracle Support before proceeding if network ping tests still fail. `# ping <XMI or OAM gateway address>` Execute the syscheck command, ensuring that all tests return successfully. If errors are found, discontinue this procedure and contact My Oracle Support. 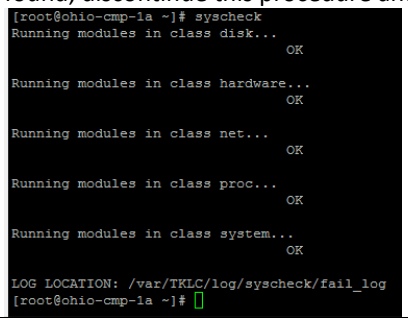
7. ☐	Remove 'Forced Standby' designation on current node.	In the CMP GUI, navigate to: **Platform Setting → Topology Setting → Current Cluster** 1. Modify for the server that has 'Forced standby' 2. Clear the Forced Standby checkbox 3. Click Save 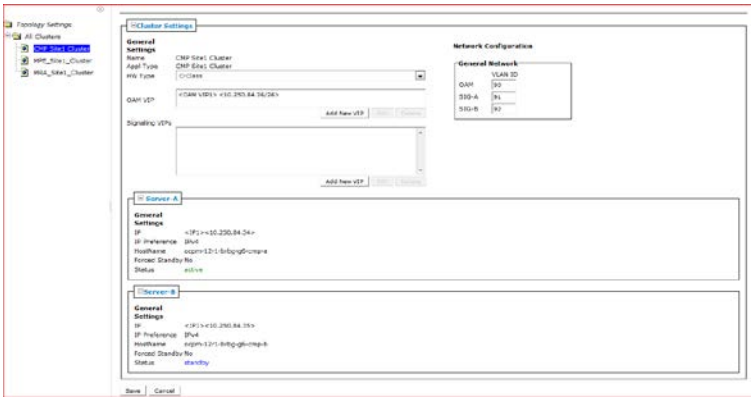 Accept the resulting pop-up by clicking OK: 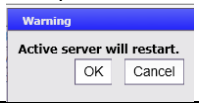
8. ☐	Verify cluster status	In the CMP GUI, navigate to: **Platform Setting → Topology Setting → All → Current CMP Cluster** Monitor clustering of the new node to its peer, do not proceed until both nodes have a status of either '*active*' or '*standby*', and that there are no CMP related 'Active Alarms' as shown below. 

9. ☐	Alternative method to check replication status	You can also monitor the clustering of the new node from within the shell on the primary node with 'irepstat'. To do so, SSH to the Active node of the current cluster and execute the irepstat command: `# irepstat` Expected 'irepstat' output while waiting reconnection: 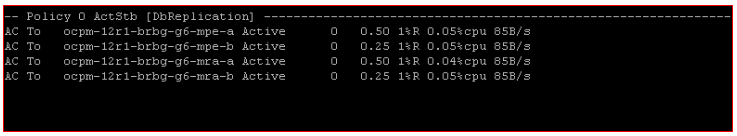 Expected 'irepstat' output after cluster has formed: 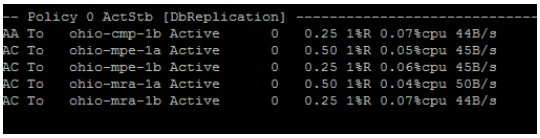
10. ☐	Exchange keys with cluster mate (This step needs to run from the active CMP)	Exchanging SSH keys Utility as `root`, please run `'/opt/camiant/bin/qpSSHKeyProv.pl --prov --user=root'`; as `admusr`, please run `'/usr/bin/sudo /opt/camiant/bin/qpSSHKeyProv.pl --prov';` 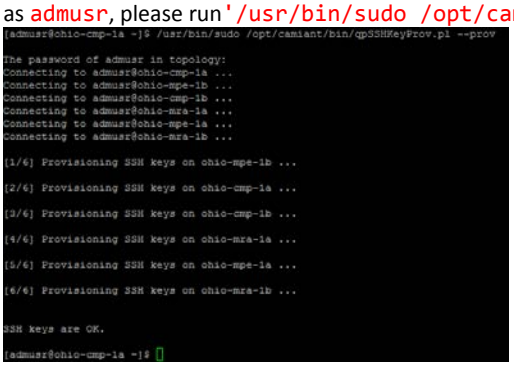 This procedure is completed.

5.4 *Procedure 3: Restore single MPE/MRA/BOD/MA node with server backup file*

The purpose of this procedure is to replace one node of a policy cluster. Restore initial Policy configuration from a server backup file, and then allow the new node to re-sync to the existing node to form a complete cluster. In this example, initial Policy configuration is restored to the new nodes through the use of server backup files for each server to be restored.

Required resources:

- Host server identified for the new VM instance
- OVA file or equivalent (depending on hypervisor or NFV manager)
- *serverbackup*.ISO of the node to be replaced

Prerequisites:

- failed VM is no longer available (e.g. it has been removed from the hypervisor/NFV manager)
- a new VM has been created in accordance with [1].

S T E P #	This procedure performs Restore single MPE/MRA/BOD/MA node with server backup file. Check off (√) each step as it is completed. Boxes have been provided for this purpose under each step number. Should this procedure fail, contact the My Oracle Support (MOS) Customer Care Center and ask for assistance.	
1. ☐	Set the failed node to 'Forced Standby'	In the CMP GUI, navigate to: **Platform Setting → Topology Setting → All Clusters** 1. Determine the cluster with the failed node 2. Determine the failed node 3. Click the Modify Server-X for the failed node 4. Click the Forced Standby checkbox so that it is checked, then click Save 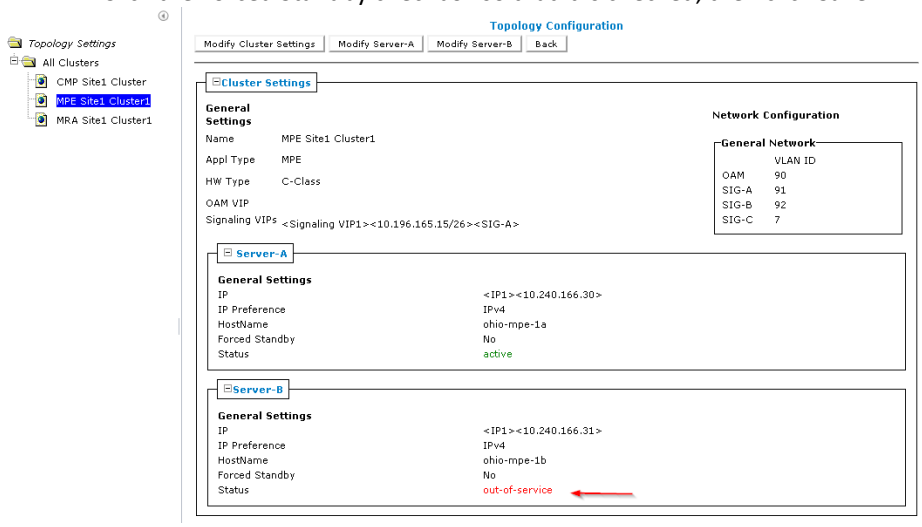
2. ☐	Create the VM instance	Create the VM instance according to [1]
3. ☐	Load the ISO for server backup	Obtain the *serverbackup.iso* for the node to be restored. The server backup file should be copied via secure copy(pscp,scp, or WinSCP) to the following directory: **`/var/camiant/backup/local_archive/serverbackup`** ***Note:*** *Later in this procedure, the platcfg restore function check this directory and offer the user a convenient menu to choose from. The platcfg utility also allows the user to manually enter any mounted path on the server.*

4. ☐	Login via SSH to new node	SSH to the new VM instance:
		`# ssh admusr@<node_IP_Address>`
		`$ sudo su -`

5. ☐	Perform platcfg restore from SSH session to replacement hardware	Execute the following command: `# su - platcfg` From within the platcfg utility, navigate to: **Policy Configuration → Backup and Restore → Server Restore** Select the *serverbackup*.ISO that you just put on the system and hit 'ok' – then 'yes' to confirm. 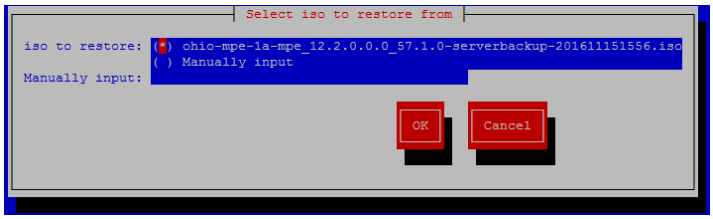 
6. ☐	Verify the status	A window will pop-up, indicating restore operation was successful and will ask the user to press any key to exit. If it is not successful, retry the restore. If the second restore is not successful, stop and contact My Oracle Support or engineering team for assistance.

7. ☐	Perform Initial configuration	Choose Exit repeatedly until back to the Main Menu of the platcfg utility. While still within the platcfg utility, navigate to:**Policy Configuration → Verify Initial Configuration** 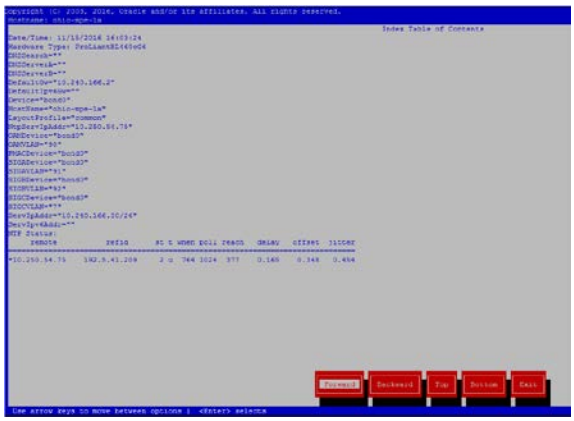 If the configuration does not exist, then navigate to **'Perform Initial Configuration'** and fill in initial configuration: hostname, OAM IP and NTP servers configurations as shown below: 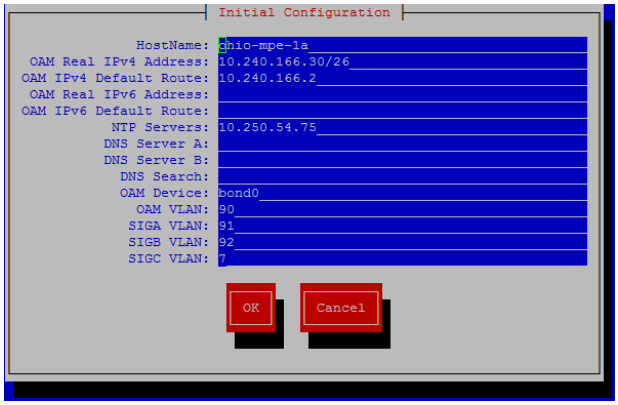 Ensure the configured data is correct, and select 'OK', then 'yes' to save and apply **Exit platcfg**. Exit platcfg by selecting Exit from each platcfg menu until you are returned to the shell. 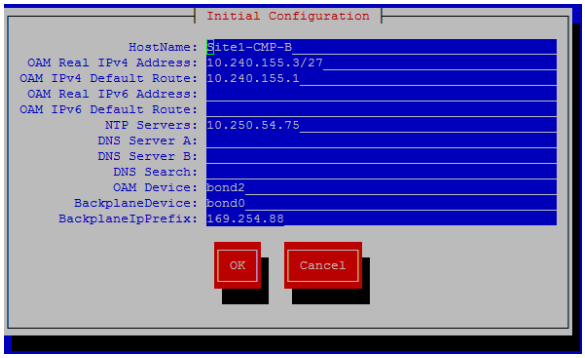 *Note: The above snapshot is for 'Cable mode', for 'Wireless mode' the "Backplane Device" and "Backplane IP Prefix" parameters will not exist.*	

8. ☐	Reboot the server	Reboot: `# init 6` Allow the server time to reboot;
9. ☐	Verify basic network connectivity and server health.	From the newly installed server, ping the OAM/XMI gateway. If the ping is not successful, verify all network settings match the old hardware configuration and reconfigure if needed. Contact My Oracle Support before proceeding if network ping tests still fail. `# ping <XMI or OAM gateway address>` Execute the *syscheck* command, ensuring that all tests return successfully. If errors are found, discontinue this procedure and contact My Oracle Support. 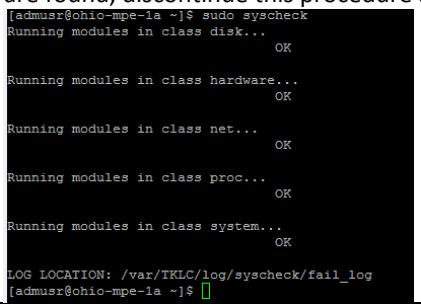
10. ☐	Remove 'Forced Standby' designation on current node.	In the CMP GUI, navigate to: **Platform Setting → Topology Settings → All Clusters → Current Cluster** 1. Modify for the server that has 'Forced Standby' 2. Clear the 'Forced Standby' checkbox 3. Click Save 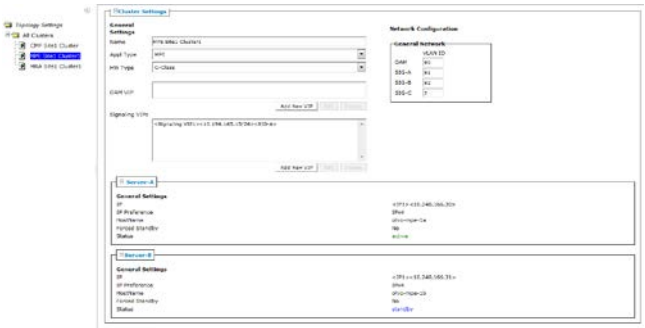 Accept the resulting pop-up by clicking OK: 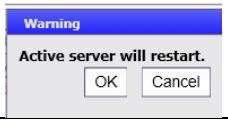

11. ☐	Check status	In the CMP GUI, depending on the type of the node, perform the following: If this is an MPE node, navigate to: **Policy Server → Configuration → All→*<Recovered MPE Cluster>*→ Reports Tab** If this is an MRA node, navigate to: **MRA → Configuration → All →*<Recovered MRA Cluster>*→ Reports Tab** If this is an BOD node, navigate to: **BOD → Configuration → All →*<Recovered BOD Cluster>*→ Reports Tab** If this is an MA node, navigate to: **MA →*<Recovered MA Cluster>*→Reports Tab** Monitor clustering of the new node to its peer, do not proceed until the Cluster Status returns from '*Degraded*' to '*On-line'* 

12. ☐	Alternative method to check replication status	You can also monitor the clustering of the new node from within the shell on the primary node with 'irepstat'. To do so, SSH to the Active node of the current cluster and execute the irepstat command:
		
		`# irepstat`
		Expected 'irepstat' output while waiting reconnection:
		`-- Policy 0 ActStb [DbReplication] ----------`
		`AC From ocpm-12r1-brbg-g6-cmp-a Active 0 0.25 ^0.04%cpu 45B/s A=me`
		
		Expected 'irepstat' output after cluster has formed:
		`-- Policy 0 ActStb [DbReplication] ----------`
		`AC From ocpm-12r1-brbg-g6-cmp-a Active 0 0.25 ^0.04%cpu 52B/s A=C2488.184`
		`CC From ocpm-12r1-brbg-g6-mpe-a Active 0 0.50 ^0.06 2.45%cpu 35B/s A=C2488.184`

13. ☐	Exchange keys with cluster mate(This step need to run from active CMP)	Exchanging SSH keys Utility as `root`, please run `'/opt/camiant/bin/qpSSHKeyProv.pl --prov --user=root'`; as `admusr`, please run`'/usr/bin/sudo /opt/camiant/bin/qpSSHKeyProv.pl --prov';` 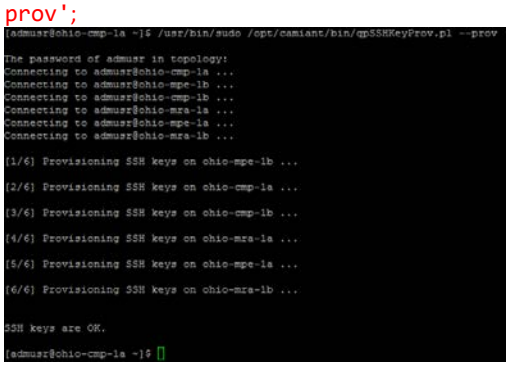 This procedure is completed.

5.5 *Procedure 4: Restore single MPE/MRA/BOD/MA node without server backup file*

The purpose of this procedure is to create a policy cluster from the replacement of one node of the cluster. The active primary node will then synchronize the newly installed node to complete the cluster. In this example, initial policy configuration is restored to the new node by manual entry.

Required resources:

- Host server identified for the new VM instance
- OVA file or equivalent (depending on hypervisor or NFV manager)
- Initial configuration information about the node to be restored:
  - OAM IP address, default gateway, NTP & SNMP server IP addresses
  - VLAN configuration information.

Hostname, OAM IP address, and VLAN configuration can be gleaned from:
Platform Setting → Topology Setting → <Cluster_Name>
NTP server configuration (and optionally DNS configuration can be gotten from platcfg of the running node)
Verify that routing is configured correctly i.e. XSI is default and any associated OAM routes are added.

Prerequisites:

- failed VM is no longer available (e.g. it has been removed from the hypervisor/NFV manager)
- a new VM has been created in accordance with [1]

S T E P #	This Procedure performs Restore single MPE/MRA/BOD/MA node without server backup file Check off (√) each step as it is completed. Boxes have been provided for this purpose under each step number. Should this procedure fail, contact the My Oracle Support (MOS) Customer Care Center and ask for assistance.	
1. ☐	Set the failed node to 'Forced Standby'	In the CMP GUI, navigate to: **Platform Setting → Topology Setting → All Clusters** 1. Determine the cluster with the failed node 2. Determine the failed node Note: It is possible that for a Geo-Redudant Topology, the server C will be a failed node 3. Click the Modify Server-X for the failed node 4. Click the Forced Standby checkbox so that it is checked, then click Save 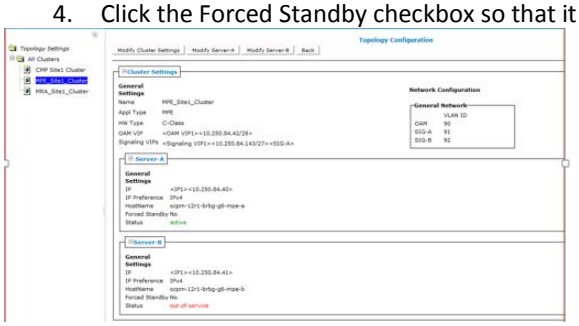 *Server-C (spare): In a Geo-Redundant Topology* 
2. ☐	Create the VM instance	Create the VM instance according to [1]

3. ☐	Login via SSH to new node	SSH session to the new VM instance:
		`# ssh admusr@<node_IP_Address>`
		`$ sudo su -`

4. ☐	Perform 'Initial Policy Configuration' from within platcfg utility on newly installed node	Execute the following command `# su - platcfg` From within the platcfg utility, navigate to: **Policy Configuration → Perform Initial Configuration**. Enter the configuration details from the node being replaced: 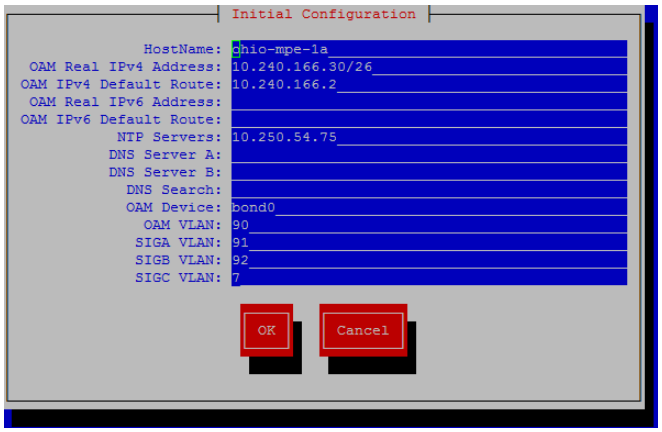 Once the server details are entered and verified for correctness select 'Ok'. A menu will appear asking if the new settings should be applied, select 'YES' and allow the operation to complete. No specific message is given when the operation is successful, but an error will appear if it was not completed. In this case, review the settings from the 'Perform Initial Configuration screen again, if all appears as expected, contact My Oracle Support before proceeding. Exit platcfg by selecting Exit from each platcfg menu until you are returned to the shell. 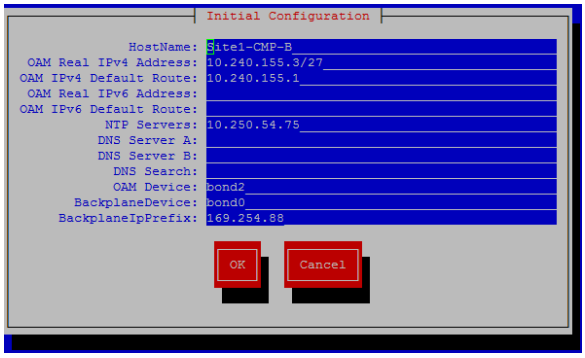 *Note: The above snapshot is for 'Cable mode', for 'Wireless mode' the "Backplane Device" and "Backplane IP Prefix" parameters will not exist*.
5. ☐	Reboot the server	Reboot: `# init 6` Allow the server time to reboot;

6. ☐	Verify basic network connectivity and server health.	From the newly installed server, ping the OAM/XMI gateway. If the ping is not successful, verify all network settings match the old blade configuration and reconfigure if needed. Contact My Oracle Support before proceeding if network ping tests still fail. `# ping <XMI or OAM gateway address>` Execute the syscheck command, ensuring that all tests return successfully. If errors are found, discontinue this procedure and contact My Oracle Support. 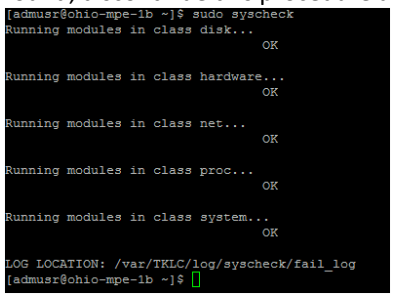	
7. ☐	Remove 'Forced Standby' designation on current blade.	In the CMP GUI, navigate to: **Platform Setting → Topology Setting → Current Cluster** 1. Modify for the server that has 'Forced Standby' 2. Clear the 'Forced Standby' checkbox 3. Click Save 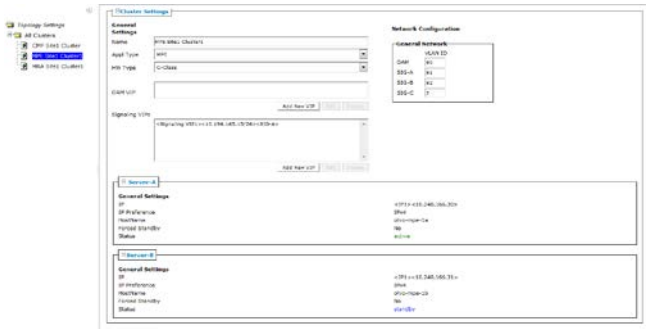 Accept the resulting pop-up by clicking OK: 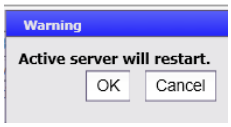	

8. ☐	Check status	In the CMP GUI, depending on the type of the blade, perform the following: If this is an MPE node, navigate to: **Policy Server → Configuration → All→*<Recovered MPE Cluster>*→ Reports Tab** If this is an MRA node, navigate to: **MRA → Configuration → All →*<Recovered MRA Cluster>*→ Reports Tab** If this is an BOD node, navigate to: **BOD → Configuration → All →*<Recovered BOD Cluster>*→ Reports Tab** If this is an MA node, navigate to: **MA →*<Recovered MA Cluster>*→Reports Tab** Monitor clustering of the new blade to its peer, do not proceed until the Cluster Status returns from '*Degraded*' to '*On-line'* 

9. ☐	Alternative method to check replication status	You can also monitor the clustering of the new blade from within the shell on the primary node with 'irepstat'. To do so, SSH to the Active node of the current cluster and execute the irepstat command:
		`# irepstat`
		Expected 'irepstat' output while waiting reconnection:
		 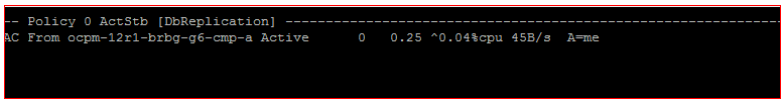
		`-- Policy 0 ActStb [DbReplication] ----------`
		`AC From ocpm-12r1-brbg-g6-cmp-a Active 0 0.25 ^0.04%cpu 45B/s A=me`
		
		Expected 'irepstat' output after cluster has formed:
		 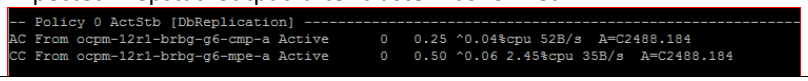
		`-- Policy 0 ActStb [DbReplication] ----------`
		`AC From ocpm-12r1-brbg-g6-cmp-a Active 0 0.25 ^0.04%cpu 52B/s A=C2488.184`
		`CC From ocpm-12r1-brbg-g6-mpe-a Active 0 0.50 ^0.06 2.45%cpu 35B/s A=C2488.184`

10. ☐	Exchange keys with cluster mate (This step needs to run from the active CMP)	Exchanging SSH keys Utility
		as `root`, please run `'/opt/camiant/bin/qpSSHKeyProv.pl --prov --user=root'`;
		as `admusr`, please run `'/usr/bin/sudo /opt/camiant/bin/qpSSHKeyProv.pl --prov';`
		 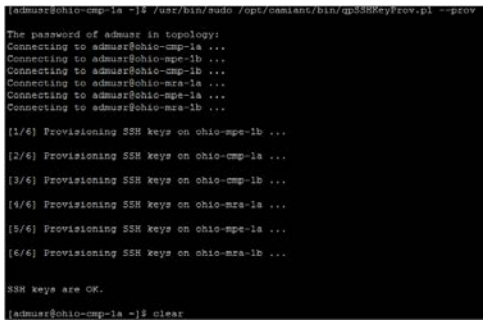
		`[admusr@ohio-cmp-1a ~]$ /usr/bin/sudo /opt/camiant/bin/qpSSHKeyProv.pl --prov`
		`The password of admusr in topology:`
		`Connecting to admusr@ohio-cmp-1a ...`
		`Connecting to admusr@ohio-mpe-1b ...`
		`Connecting to admusr@ohio-cmp-1b ...`
		`Connecting to admusr@ohio-mra-1a ...`
		`Connecting to admusr@ohio-mpe-1a ...`
		`Connecting to admusr@ohio-mra-1b ...`
		`[1/6] Provisioning SSH keys on ohio-mpe-1b ...`
		`[2/6] Provisioning SSH keys on ohio-cmp-1a ...`
		`[3/6] Provisioning SSH keys on ohio-cmp-1b ...`
		`[4/6] Provisioning SSH keys on ohio-mra-1a ...`
		`[5/6] Provisioning SSH keys on ohio-mpe-1a ...`
		`[6/6] Provisioning SSH keys on ohio-mra-1b ...`
		`SSH keys are OK.`
		`[admusr@ohio-cmp-1a ~]$ clear`
		
		This procedure is completed.

5.6 *Procedure 5: Restoring complete cluster with the server backup files*

The purpose of this procedure is to create policy cluster VM instances, then restore application level configuration by pushing that configuration from the active CMP. In this example, initial Policy configuration is restored to the new blades through the use of server backup files for each server to be restored.

Required resources:

- Host server(s) identified for the new VM instances
- OVA file or equivalent (depending on hypervisor or NFV manager)
- *serverbackup*.iso of the blade to be replaced

Prerequisites:

- failed VMs are no longer available (e.g. they have been removed from the hypervisor/NFV manager)
- new VM instances have been created in accordance with [1].
    - Note: In case it is a CMP Cluster that is being rebuilt, restore application data either from system backup or manually if no backup available.

S T E P #	This Procedure performs Restoring complete cluster with the server backup files Check off (√) each step as it is completed. Boxes have been provided for this purpose under each step number. Should this procedure fail, contact the My Oracle Support (MOS) Customer Care Center and ask for assistance.	
1. ☐	Create the VM instance	Create the VM instance according to [1]
2. ☐	Load the ISO to restore 1st server of the cluster	Obtain the *serverbackup.iso* for the blade to be restored. The server backup file should be copied via secure copy(pscp,scp, or WinSCP) to the following directory: **/var/camiant/backup/local_archive/serverbackup** ***Note:*** *Later in this procedure, the platcfg restore function check this directory and offer the user a convenient menu to choose from. The platcfg utility also allows the user to manually enter any mounted path on the server.*

3. ☐	SSH to replacement VM instance	SSH to the new VM instance:
		`# ssh admusr@<node_IP_Address>`
		`$ sudo su -`

4. ☐	Perform platcfg restore from SSH session to replacement VM instance	Execute the following command `# su - platcfg` From within the platcfg utility, navigate to: **Policy Configuration → Backup and Restore → Server Restore** Select the *serverbackup*.iso that you just put on the system and hit ok – then 'yes' to confirm. 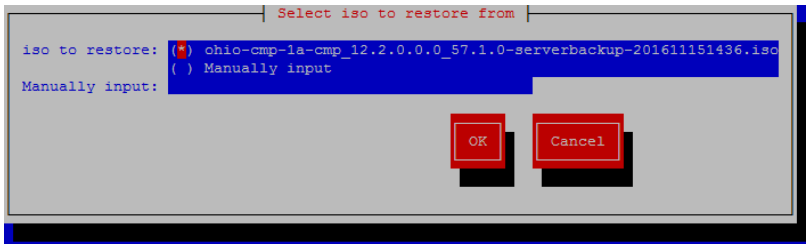
5. ☐	Verify the status	A window will pop-up, indicating restore operation was successful and will ask the user to press any key to exit. If it is not successful, retry the restore. If the second restore is not successful, stop and contact My Oracle Support or engineering team for assistance.

6. ☐	Verify Initial configuration	Choose Exit repeatedly until back to the Main Menu of the platcfg utility. While still within the platcfg utility, navigate to:**Policy Configuration → Verify Initial Configuration**  If the configuration does not exists, then navigate to **'Perform Initial Configuration'** and fill in initial configuration: hostname, OAM IP and NTP servers configurations as shown below: 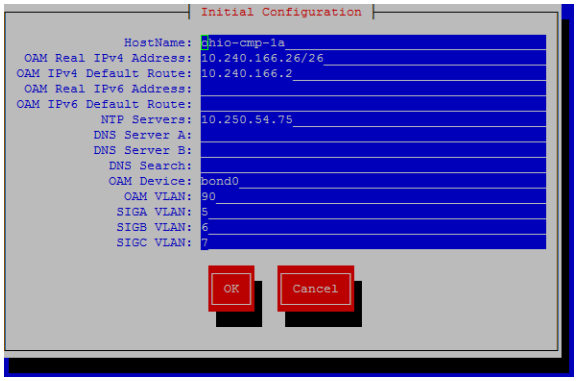 Ensure that your data is correct, and select 'Ok', then 'yes' to save and apply **Exit platcfg**: Exit platcfg by selecting Exit from each platcfg menu until you are returned to the shell. 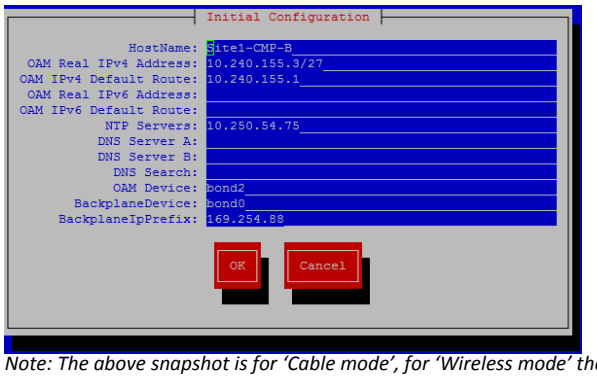 *Note: The above snapshot is for 'Cable mode', for 'Wireless mode' the "Backplane Device" and "Backplane IP Prefix" parameters will not exist.*
7. ☐	Reboot the server	Reboot: `# init 6` Allow the server time to reboot;

8. ☐	Verify basic network connectivity and server health.	From the newly installed VM instance, ping the OAM/XMI gateway. If the ping is not successful, verify all network settings match the old blade configuration and reconfigure if needed. Contact My Oracle Support before proceeding if network ping tests still fail. `# ping <XMI or OAM gateway address>` Execute the syscheck command, ensuring that all tests return successfully. If errors are found, discontinue this procedure and contact My Oracle Support. 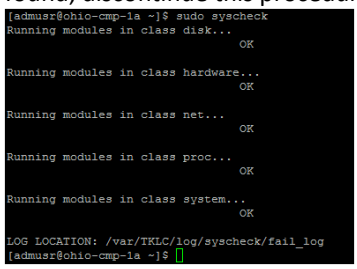	
9. ☐	Check status	In the CMP GUI, navigate to: **Platform Setting → Topology Setting → All Clusters** Check system tab for the cluster. If the Status field indicates '*Config Mismatch*', click the '***Reapply Configuration***' button and wait for the '*Config Mismatch*' designation to disappear. If it does not, contact My Oracle Support before proceeding. 	
10. ☐	Set 'Forced Standby' designation on cluster node that is still 'out-of-service'.	In the CMP GUI, navigate to: **Platform Setting → Topology Setting → Current Cluster** 1. Modify for the server that has a status of '*out-of-service*' 2. Check the Forced Standby checkbox 3. Click Save 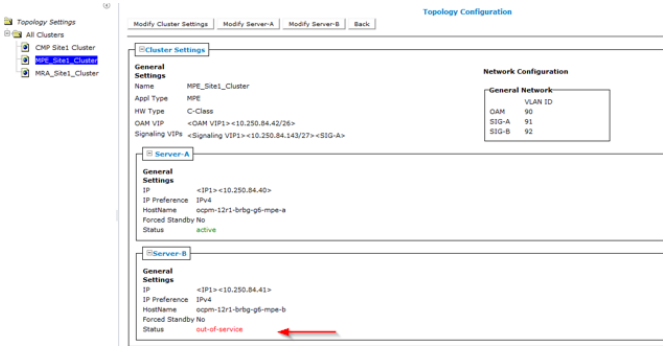 Accept the resulting pop-up by clicking OK: 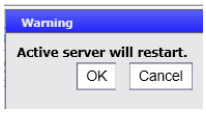	
11.	Create the VM instance	Create the VM instance according to [1]	

12. ☐	Load the ISO to restore 2nd server of the cluster	Obtain the *serverbackup.iso* for the blade to be restored. The server backup file should be copied via secure copy(pscp,scp, or WinSCP) to the following directory: **/var/camiant/backup/local_archive/serverbackup** ***Note:*** *Later in this procedure, the platcfg restore function check this directory and offer the user a convenient menu to choose from. The platcfg utility also allows the user to manually enter any mounted path on the server.*

13. ☐	SSH to replacement VM instance	SSH to the new VM instance:
		`# ssh admusr@<node_IP_Address>`
		`$ sudo su -`

14. ☐	Perform platcfg restore from SSH session to replacement VM instance	Execute the following command `# su – platcfg` From within the platcfg utility, navigate to: **Policy Configuration → Backup and Restore → Server Restore** Select the *serverbackup*.iso that you just put on the system and hit ok – then 'yes' to confirm. 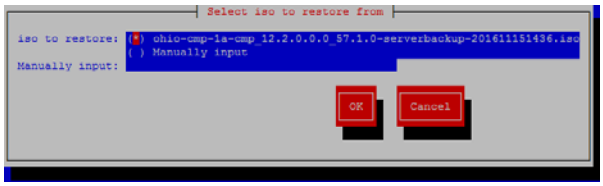
15. ☐	Verify the status	A window will pop-up, indicating restore operation was successful and will ask the user to press any key to exit. If it is not successful, retry the restore. If the second restore is not successful, stop and contact My Oracle Support or engineering team for assistance.

16. ☐	Verify Initial configuration	Choose Exit repeatedly until back to the Main Menu of the platcfg utility. While still within the platcfg utility, navigate to:**Policy Configuration → Verify Initial Configuration** 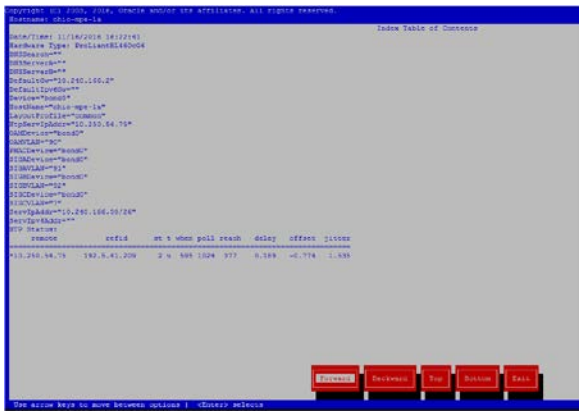 If the configuration does not exists, then navigate to **'Perform Initial Configuration'** and fill in initial configuration: hostname, OAM IP and NTP servers configurations as shown below: 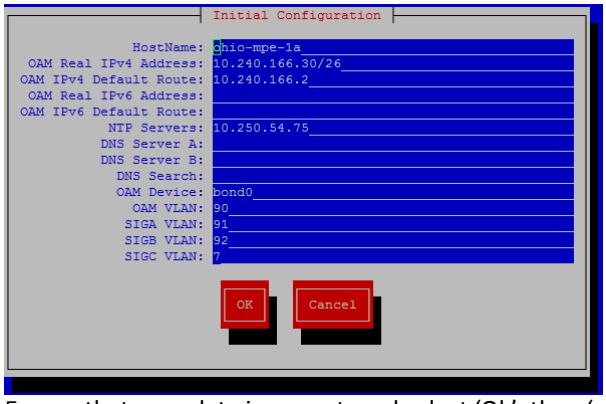 Ensure that your data is correct, and select 'Ok', then 'yes' to save and apply **Exit platcfg** Exit platcfg by selecting Exit from each platcfg menu until you are returned to the shell. 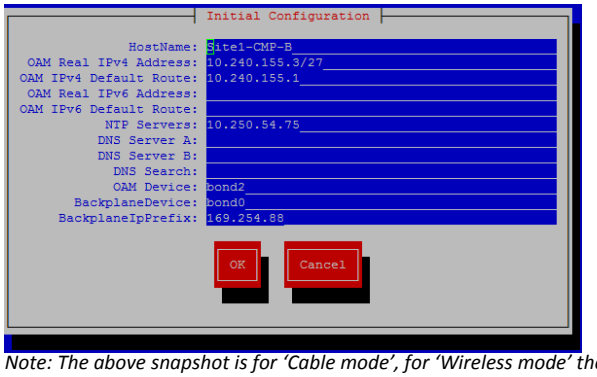 *Note: The above snapshot is for 'Cable mode', for 'Wireless mode' the "Backplane Device" and "Backplane IP Prefix" parameters will not exist.*
17. ☐	Reboot the server	Reboot: `# init 6` Allow the server time to reboot;

18. ☐	Remove 'Forced Standby' designation on current blade.	In the CMP GUI, navigate to: **Platform Setting → Topology Settings → Current Cluster** 1. Modify for the server that has 'Forced Standby' 2. Clear the 'Forced Standby' checkbox 3. Click Save 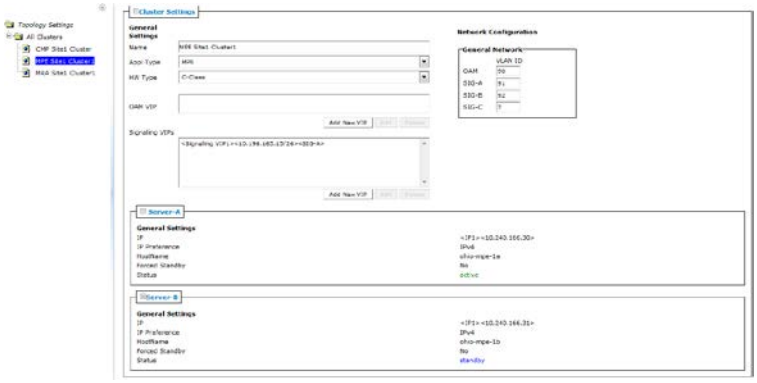 Accept the resulting pop-up by clicking OK: 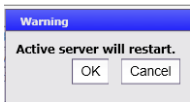
19. ☐	Check status	In the CMP GUI, depending on the type of the blade, perform the following: If this is an MPE node, navigate to: **Policy Server → Configuration → All→*<Recovered MPE Cluster>*→ Reports Tab** If this is an MRA node, navigate to: **MRA → Configuration → All →*<Recovered MRA Cluster>*→ Reports Tab** If this is an BOD node, navigate to: **BOD → Configuration → All →*<Recovered BOD Cluster>*→ Reports Tab** If this is an MA node, navigate to: **MA →*<Recovered MA Cluster>*→Reports Tab** Check CMP cluster status (as indicated in the previous step), navigate to: **Platform Setting → Topology Setting → Current CMP Cluster** Monitor clustering of the new blade to its peer, do not proceed until the Cluster Status returns from '*Degraded*' to '*On-line*' 

20. ☐	Alternative method to check replication status	You can also monitor the clustering of the new blade from within the shell on the primary node with 'irepstat'. To do so, SSH to the Active node of the current cluster and execute the irepstat command: # irepstat Expected 'irepstat' output while waiting reconnection: -- Policy 0 ActStb [DbReplication] ---------- AC To ocpm-12r1-brbg-g6-mpe-a Active 0 0.50 1%R 0.05%cpu 85B/s AC To ocpm-12r1-brbg-g6-mpe-b Active 0 0.25 1%R 0.05%cpu 85B/s AC To ocpm-12r1-brbg-g6-mra-a Active 0 0.50 1%R 0.04%cpu 85B/s AC To ocpm-12r1-brbg-g6-mra-b Active 0 0.25 1%R 0.05%cpu 85B/s Expected 'irepstat' output after cluster has formed: -- Policy 0 ActStb [DbReplication] ---------- AA To ohio-cmp-1b Active 0 0.25 1%R 0.07%cpu 79B/s AC To ohio-mpe-1a Active 0 0.50 1%R 0.05%cpu 65B/s AC To ohio-mpe-1b Active 0 0.25 1%R 0.07%cpu 78B/s AC To ohio-mra-1a Active 0 0.50 1%R 0.05%cpu 65B/s AC To ohio-mra-1b Active 0 0.25 1%R 0.07%cpu 79B/s
21. ☐	Exchange keys with cluster mate (This step needs to run from the active CMP)	Exchanging SSH keys Utility as root, please run '/opt/camiant/bin/qpSSHKeyProv.pl --prov --user=root'; as admusr, please run '/usr/bin/sudo /opt/camiant/bin/qpSSHKeyProv.pl --prov'; [admusr@ohio-cmp-1a ~]$ /usr/bin/sudo /opt/camiant/bin/qpSSHKeyProv.pl --prov The password of admusr in topology: Connecting to admusr@ohio-cmp-1a ... Connecting to admusr@ohio-mpe-1b ... Connecting to admusr@ohio-cmp-1b ... Connecting to admusr@ohio-mra-1a ... Connecting to admusr@ohio-mpe-1a ... Connecting to admusr@ohio-mra-1b ... [1/6] Provisioning SSH keys on ohio-mpe-1b ... [2/6] Provisioning SSH keys on ohio-cmp-1a ... [3/6] Provisioning SSH keys on ohio-cmp-1b ... [4/6] Provisioning SSH keys on ohio-mra-1a ... [5/6] Provisioning SSH keys on ohio-mpe-1a ... [6/6] Provisioning SSH keys on ohio-mra-1b ... SSH keys are OK. [admusr@ohio-cmp-1a ~]$ This procedure is completed.

5.7 *Procedure 6: Restoring complete cluster without the server backup*

The purpose of this procedure is to restore a policy cluster without the server backup file. The active primary blade will then synchronize the newly installed blade to complete the cluster. In this example, initial Policy configuration is restored to the new blade by manual entry.

Required resources:

- Host server(s) identified for the new VM instances
- OVA file or equivalent (depending on hypervisor or NFV manager)
- Initial configuration information about the blade to be restored:
    - OAM blade Ip address, default gateway, ntp server ip address
    - Vlan configuration information.

Hostname, OAM IP address, and VLAN configuration can be gleaned from:
Platform Setting → Topology Setting → <Cluster_Name>
NTP server configuration (and optionally DNS configuration can be gotten from platcfg of the running blade)
Verify that routing is configured correctly i.e. XSI is default and any associated OAM routes are added.

Prerequisites:

- failed VMs are no longer available (e.g. they has been removed from the hypervisor/NFV manager)
- new VM instances have been created in accordance with [1]
- Install application software – CMP, MPE, MRA, BOD or MA
    - Note: In case it is a CMP Cluster that is being rebuilt, restore application data either from system backup or manually if no backup available.

S T E P #	This Procedure performs Restoring complete cluster without the server backup Check off (√) each step as it is completed. Boxes have been provided for this purpose under each step number. Should this procedure fail, contact the My Oracle Support (MOS) and ask for assistance.	
1. ☐	Create the VM	Create the VM instance according to [1]

2. ☐	Login via SSH to replacement VM instance	SSH to the new VM instance:
		`# ssh admusr@<node_IP_Address>`
		`$ sudo su -`

3. ☐	Perform 'Initial Policy Configuration' from within platcfg utility on newly installed VM instance	Execute the following command `# su - platcfg` From within the platcfg utility, navigate to: **Policy Configuration → Perform Initial Configuration** Enter the relevant configuration details from the blade being replaced: 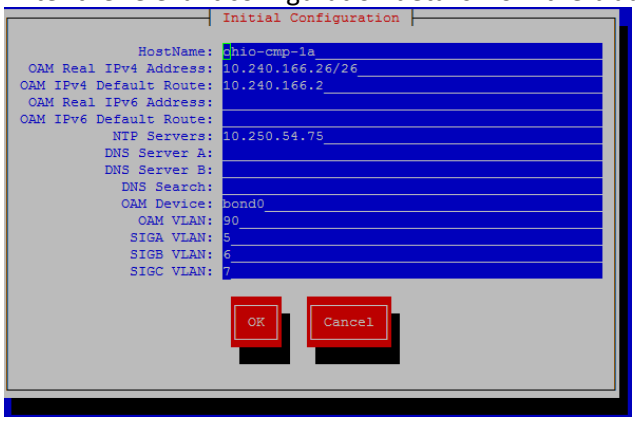 Once the server details are entered and verified for correctness select 'Ok'. A menu will appear asking if the new settings should be applied, select 'YES' and allow the operation to complete. No specific message is given when the operation is successful, but an error will appear if it was not completed. In this case, review the settings from the 'Perform Initial Configuration screen again, if all appears as expected, contact My Oracle Support before proceeding. Exit platcfg by selecting Exit from each platcfg menu until you are returned to the shell. 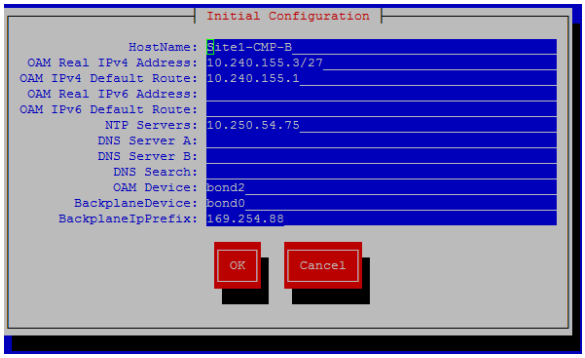 *Note: The above snapshot is for 'Cable mode', for 'Wireless mode' the "Backplane Device" and "Backplane IP Prefix" parameters will not exist.*

4. ☐	Reboot the server	Reboot: `# init 6` Allow the server time to reboot;
5. ☐	Verify basic network connectivity and server health.	From the newly installed VM, ping the OAM/XMI gateway. If the ping is not successful, verify all network settings match the old blade configuration and reconfigure if needed. Contact My Oracle Support before proceeding if network ping tests still fail. `# ping <XMI or OAM gateway address>` Execute the syscheck command, ensuring that all tests return successfully. If errors are found, discontinue this procedure and contact My Oracle Support. 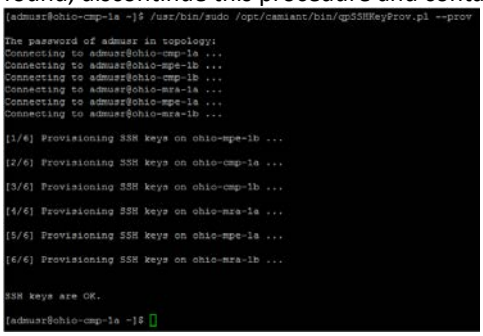
6. ☐	Check status	In the CMP GUI, depending on the type of the blade, perform the following: If this is an MPE node, navigate to: **Policy Server → Configuration → All→*<Recovered MPE Cluster>*→ Reports Tab** If this is an MRA node, navigate to: **MRA → Configuration → All →*<Recovered MRA Cluster>*→ Reports Tab** If this is an BOD node, navigate to: **BOD → Configuration → All →*<Recovered BOD Cluster>*→ Reports Tab** If this is an MA node, navigate to: **MA →*<Recovered MA Cluster>*→Reports Tab** Monitor clustering of the new blade to its peer, do not proceed until the Cluster Status returns from 'Off-line' to 'Degraded'. *Off-line*  *Degraded* 
7. ☐	Create the VM	Create the VM according to [1]

8. ☐	Login via SSH to second node of the current cluster	SSH to the new VM instance:
		`# ssh admusr@<node_IP_Address>`
		`$ sudo su -`	

9. ☐	Perform 'Initial Policy Configuration' from within platcfg utility on second node of cluster	Execute the following command `# su - platcfg` From within the platcfg utility, navigate to: **Policy Configuration → Initial Configuration** Enter the relevant details from the blade being replaced: 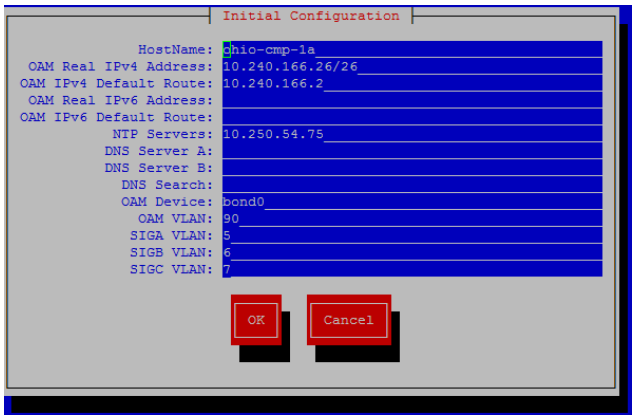 Once the server details are entered and verified for correctness select 'Ok'. A menu will appear asking if the new settings should be applied, select 'YES' and allow the operation to complete. No specific message is given when the operation is successful, but an error will appear if it was not completed. In this case, review the settings from the 'Perform Initial Configuration screen again, if all appears as expected, contact My Oracle Support before proceeding. **Exit platcfg** by selecting Exit from each platcfg menu until you are returned to the shell. 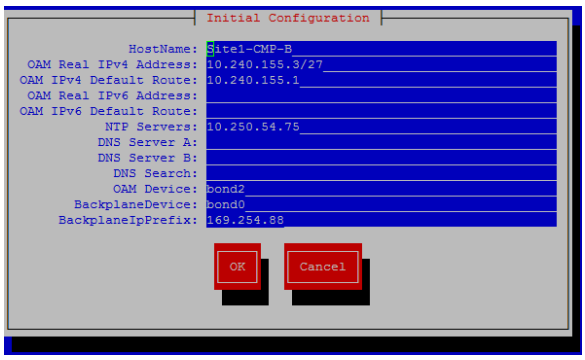 *Note: The above snapshot is for 'Cable mode', for 'Wireless mode' the "Backplane Device" and "Backplane IP Prefix" parameters will not exist.*	
10. ☐	Reboot the server	Reboot: `# init 6` Allow the server time to reboot;	

11. ☐	Verify basic network connectivity and server health.	From the newly installed VM, ping the OAM/XMI gateway. If the ping is not successful, verify all network settings match the old blade configuration and reconfigure if needed. Contact My Oracle Support before proceeding if network ping tests still fail.
		`# ping <XMI or OAM gateway address>`
		
		Execute the syscheck command, ensuring that all tests return successfully. If errors are found, discontinue this procedure and contact My Oracle Support.
		`[admusr@ohio-cmp-1a ~]$ sudo syscheck`
		`Running modules in class disk...`
		`OK`
		`Running modules in class hardware...`
		`OK`
		`Running modules in class net...`
		`OK`
		`Running modules in class proc...`
		`OK`
		`Running modules in class system...`
		`OK`
		`LOG LOCATION: /var/TKLC/log/syscheck/fail_log`
		`[admusr@ohio-cmp-1a ~]$`

12. ☐	Check status	In the CMP GUI, depending on the type of the blade, perform the following: If this is an MPE node, navigate to: **Policy Server → Configuration → All→*<Recovered MPE Cluster>*→ Reports Tab** If this is an MRA node, navigate to: **MRA → Configuration → All →*<Recovered MRA Cluster>*→ Reports Tab** If this is an BOD node, navigate to: **BOD → Configuration → All →*<Recovered BOD Cluster>*→ Reports Tab** If this is an MA node, navigate to: **MA →*<Recovered MA Cluster>*→Reports Tab** Monitor clustering of the new blade to its peer, do not proceed until the Cluster Status returns from 'Degraded' to 'On-line' *MPE:*  *MRA:* 

13. ☐	Alternative method to check replication status	You can also monitor the clustering of the new blade from within the shell on the primary node with 'irepstat'. To do so, SSH to the Active node of the current cluster and execute the irepstat command: # irepstat Expected 'irepstat' output while waiting reconnection: 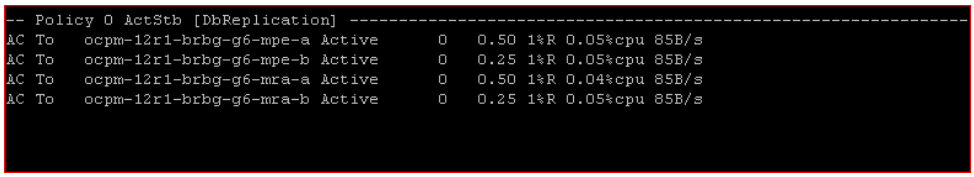 Expected 'irepstat' output after cluster has formed: 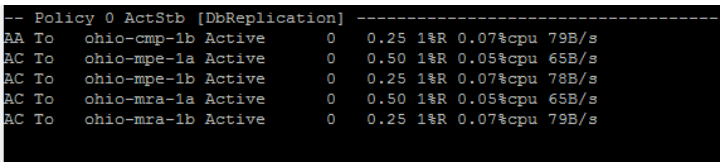
14. ☐	Exchange keys with cluster mate (This step needs to run from the active CMP)	Exchanging SSH keys Utility as root, please run '/opt/camiant/bin/qpSSHKeyProv.pl --prov --user=root'; as admusr, please run'/usr/bin/sudo /opt/camiant/bin/qpSSHKeyProv.pl --prov'; 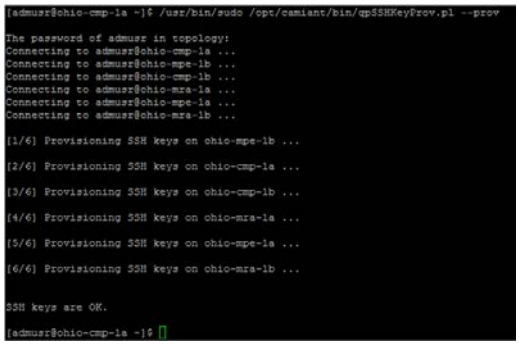 This procedure is completed.

5.8 *Procedure 7: Restoring CMP/MA cluster with system backup available*

The purpose of this procedure is to re-create a CMP with the application level configuration of the policy network that can be used to re-create the policy network that is to be recovered. Once a CMP is online, all other VM instances of the policy network can be re-created using the above procedures and then their application level configuration restored from this CMP. In the case of a massive outage that includes the CMP, at least one of the CMP VM instances should be restored first.

Required resources:

- Host server(s) identified for the new VM instances
- OVA file or equivalent (depending on hypervisor or NFV manager)
- Recent System backup file.
- Initial configuration information about the blade to be restored:
    - OAM IP address, default gateway, NTP & SNMP server IP addresses
    - VLAN configuration information.

Hostname, OAM IP address, and VLAN configuration can be gleaned from:
Platform Setting → Topology Setting → <Cluster_Name>
NTP server configuration (and optionally DNS configuration can be gotten from platcfg of the running blade)
Verify that routing is configured correctly i.e. XSI is default and any associated OAM routes are added.

Prerequisites:

- failed VMs are no longer available (e.g. they have been removed from the hypervisor/NFV manager)
- new VM instances have been created in accordance with [1].

S T E P #	This Procedure performs Restoring CMP cluster with system backup available Check off (√) each step as it is completed. Boxes have been provided for this purpose under each step number. Should this procedure fail, contact the My Oracle Support (MOS) Customer Care Center and ask for assistance.	
1. ☐	Create the VM	Create the VM according to [1]

2. ☐	Login via SSH to new VM instance	SSH to the new VM instance:
		`# ssh admusr@<node_IP_Address>`
		`$ sudo su -`

3. ☐	Perform 'Initial Policy Configuration' from within platcfg utility on newly installed VM instance	Execute the following command `# su – platcfg` From within the platcfg utility, navigate to: **Policy Configuration → Perform Initial Configuration** Enter the relevant details from the blade being replaced: 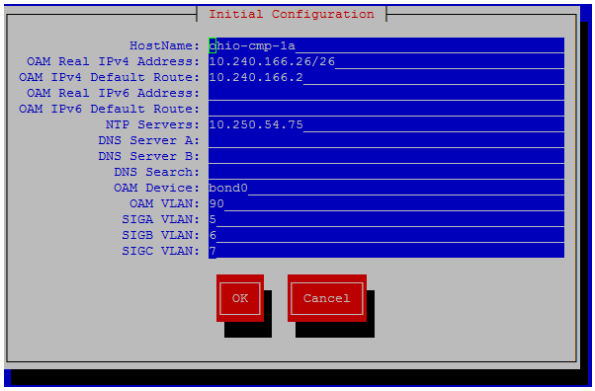 Once the server details are entered and verified for correctness select 'Ok'. A menu will appear asking if the new settings should be applied, select 'YES' and allow the operation to complete. No specific message is given when the operation is successful, but an error will appear if it was not completed. In this case, review the settings from the 'Perform Initial Configuration screen again, if all appears as expected, contact My Oracle Support before proceeding. Exit platcfg by selecting Exit from each platcfg menu until you are returned to the shell. 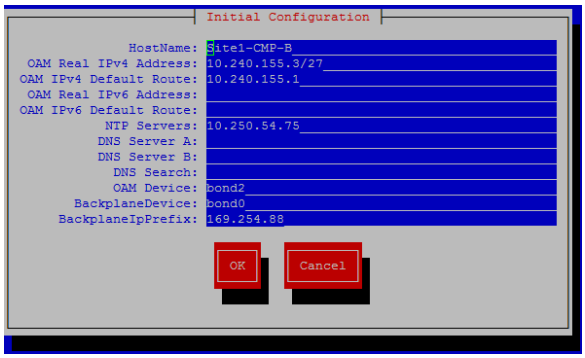 *Note: The above snapshot is for 'Cable mode', for 'Wireless mode' the "Backplane Device" and "Backplane IP Prefix" parameters will not exist.*

4. ☐	Reboot the server	Reboot: `# init 6` Allow the server time to reboot;
5. ☐	Verify basic network connectivity and server health.	From the newly installed VM, ping the OAM/XMI gateway. If the ping is not successful, verify all network settings match the old blade configuration and reconfigure if needed. Contact My Oracle Support before proceeding if network ping tests still fail. `# ping <XMI or OAM gateway address>` Execute the syscheck command, ensuring that all tests return successfully. If errors are found, discontinue this procedure and contact My Oracle Support. 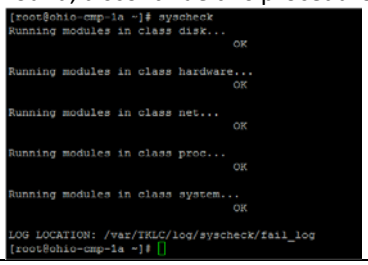
6. ☐	Load the system backup file for server restore	The system backup file contains the database information that makes up the application level configuration of the policy network. Without that backup, the application configuration will have to be restored either through the platcfg menu, or from the server backup file from site documentation. If the system backup file is available, put a copy of the file on the newly constructed CMP VM instance into the: via secure copy (pscp scp, or WinSCP). **/var/camiant/backup/local_archive/systembackup/**

7. ☐	Perform platcfg restore from SSH session to replacement VM instance	Execute the following command `# su - platcfg` From within the platcfg utility, navigate to: **Policy Configuration → Backup and Restore → System Restore** A message will appear prompting confirmation to restore even though this node is not recognized as the active member. This behavior is expected, continue by selecting 'NO'. 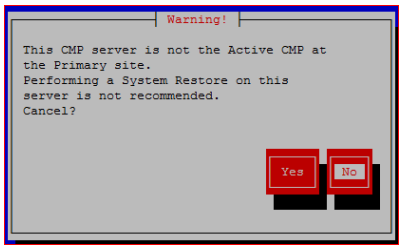 Then a screen will appear asking to select the file to restore from. If the file was copied correctly in the previous step, it will be shown here as an option, otherwise select 'Manually Input', and Select 'Full' and then select OK to proceed. 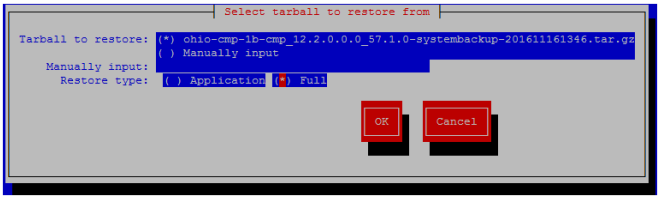 ***Note**: "Full" will also restore Comcol data, But "Application" will exclude Comcol.*	
8. ☐	Verify the status	A window will pop-up, indicating restore operation was successful and will ask the user to press any key to exit. If it is not successful, retry the restore. If the second restore is not successful, stop and contact My Oracle Support or engineering team for assistance.	

9. ☐	Verify Initial configuration	Choose Exit repeatedly until back to the Main Menu of the platcfg utility. While still within the platcfg utility, navigate to: **Policy Configuration → Verify Initial Configuration**  Ensure that your data is correct, if configuration is not there, then navigate to '**Perform Initial Configuration**' and fill in the hostname, OAM IP etc as shown below: 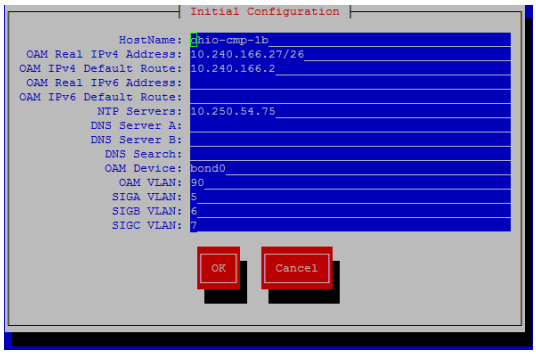 Select 'OK', then 'YES' to save and apply Once the server details are entered and verified for correctness select 'OK'. A menu will appear asking if the new settings should be applied, select 'YES' and allow the operation to complete. No specific message is given when the operation is successful, but an error will appear if it was not completed. In this case, review the settings from the 'Perform Initial Configuration screen again, if all appears as expected, contact My Oracle Support before proceeding. Exit platcfg by selecting Exit from each platcfg menu until you are returned to the shell. 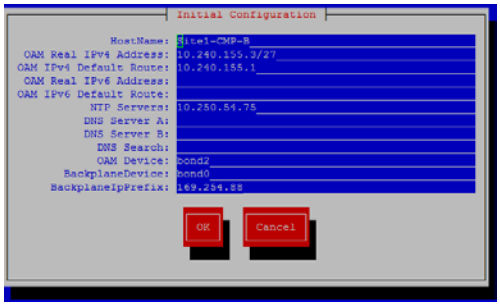 *Note: The above snapshot is for 'Cable mode', for 'Wireless mode' the "Backplane Device" and "Backplane IP Prefix" parameters will not exist.*
10. ☐	Reboot the server	Reboot. **# init 6** Allow the server time to reboot;

11. ☐	Verify basic network connectivity and server health.	From the newly installed VM, ping the OAM/XMI gateway. If the ping is not successful, verify all network settings match the old hardware configuration and reconfigure if needed. Contact My Oracle Support before proceeding if network ping tests still fail. `# ping <XMI or OAM gateway address>` Execute the syscheck command, ensuring that all tests return successfully. If errors are found, discontinue this procedure and contact My Oracle Support. 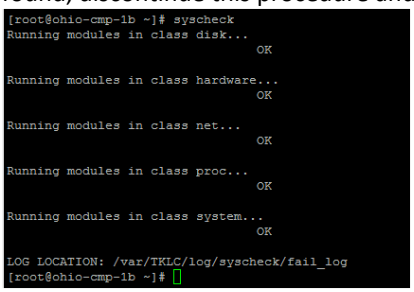
12. ☐	Exchange keys with cluster mate (This step needs to run from the active CMP)	Exchanging SSH keys Utility as root, please run `'/opt/camiant/bin/qpSSHKeyProv.pl --prov --user=root'`; as admusr, please run `'/usr/bin/sudo /opt/camiant/bin/qpSSHKeyProv.pl --prov'` 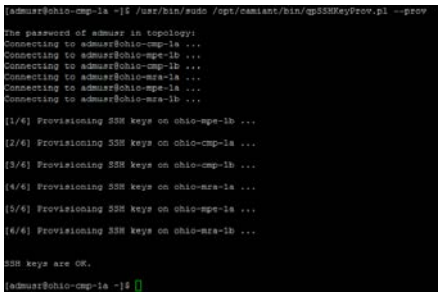
13. ☐	Check status	In the CMP GUI, navigate to: **Platform Setting → Topology Setting → All Clusters** When the server has returned to online status, log into the GUI on the OAM virtual IP address • Verify to the best of your abilities that the new manager has configuration for the MPE clusters in the network (whether those clusters are online or not) • Verify other application configuration properties as you are able. Once one CMP is in place, the other node of the CMP cluster can replaced with the procedures above, and any other clusters or individual nodes that need replacement can be handled with the above procedures. This procedure is completed.

5.9 *Procedure 8: Promoting geo-redundant CMP cluster*

This procedure is used to bring a geo-redundant secondary active CMP online before beginning restoration of other policy clusters in the network. Once a CMP is online, all other servers of the policy network can be re-created using the above procedures and then their application level configuration restored from this CMP.

S T E P #		
	This Procedure performs Promoting geo-redundant CMP cluster Check off (√) each step as it is completed. Boxes have been provided for this purpose under each step number. Should this procedure fail, contact the My Oracle Support (MOS) Customer Care Center and ask for assistance.	
1. ☐	Access to the system	Log into the GUI on the OAM VIP of the geo-redundant CMP.
2. ☐	Check status	In the CMP GUI, navigate to: **Platform Setting → Topology Setting → All Clusters** You will be warned that you are not on the primary cluster of the policy network. The secondary server has limited functionality. 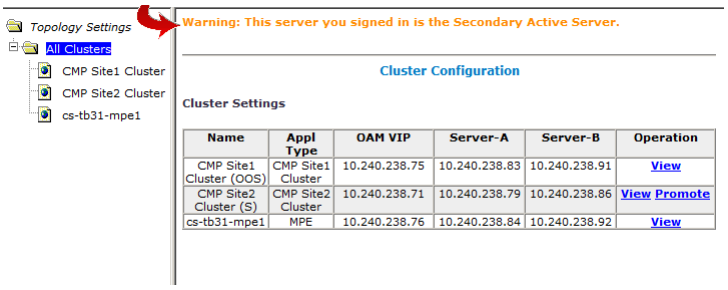
3. ☐	Verify basic network connectivity and server health.	From the active VM of site 2 CMP('Promote' server), ping the OAM/XMI gateway. If the ping is not successful, verify all network settings match the old hardware configuration and reconfigure if needed. Contact My Oracle Support before proceeding if network ping tests still fail. `# ping <XMI or OAM gateway address>` Execute the syscheck command, ensuring that all tests return successfully. If errors are found, discontinue this procedure and contact My Oracle Support. 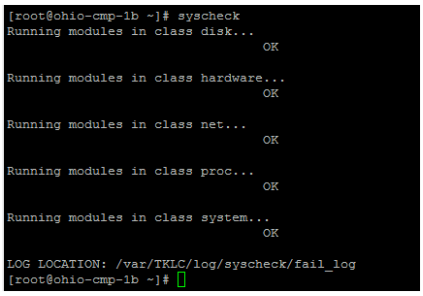

4. ☐	Promote secondary CMP cluster	In the CMP GUI, navigate to: **Platform Setting → Topology Setting → All Clusters** Select 'Promote' on the secondary server. Accept the resulting pop-up by clicking 'OK'. 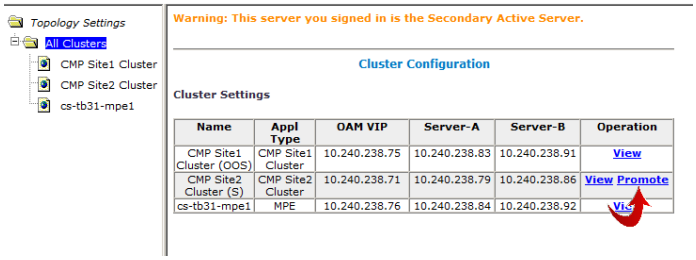 You should see a message appear above the 'Cluster Configuration' header indicating successful promotion (see example below). If not, retry the operation and/or contact My Oracle Support. 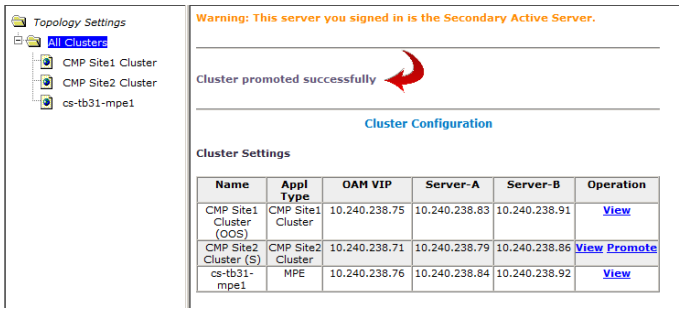
5. ☐	Logout of the CMP GUI	Logout of the CMP GUI by clicking the 'Logout' link or closing the browser window.
6. ☐	Verify operation via CMP GUI	Relogin to the CMP GUI using the VIP of CMP Site2 In the CMP GUI, navigate to: **Platform Setting → Topology Setting → All Clusters** Ensure all clusters are performing as expected. Follow procedures listed in this document to bring other failed servers/clusters back online.

7. ☐	SSH to active node of newly promoted cluster	SSH to the active node of the newly promoted cluster
		
		`# ssh admusr@<node_IP_Address>`
		`$ sudo su -`

8. ☐	Verify irepstat output shows expected status	Execute the irepstat command to verify that cluster replication is 'Active'. If not 'Active' after 5 minutes, check the CMP GUI for any active alarms. # irepstat 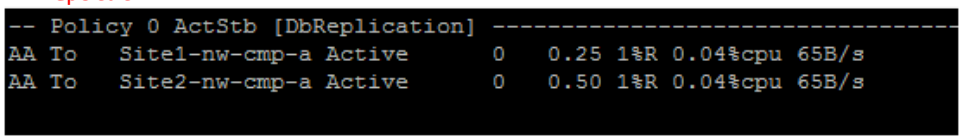 The status of all clusters except known failed servers should have a status of 'Active' as in the above snapshot. Otherwise if any of the replication paths show 'DownConnecting' as in the snapshot below contact My Oracle Support. The example shown below shows our installation with servers 'cs-tb31-cmp2-a' and 'cs-tb31-cmp2-b' failed, while all other cluster replication is working properly. 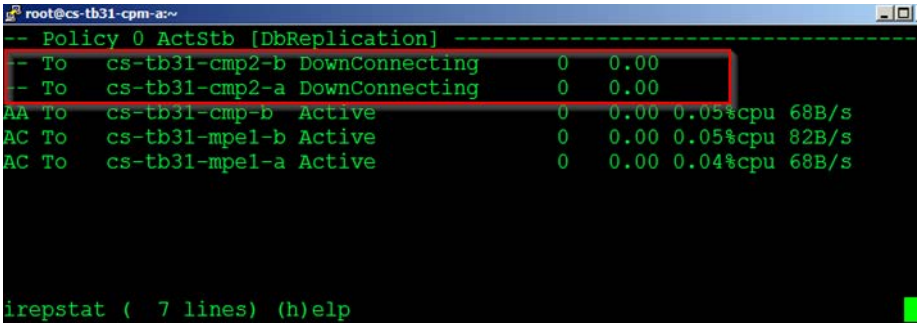
9. ☐	Rebuild failed CMP cluster	Refer to Procedure 6: Restoring complete cluster without the server backup to rebuild failed CMP cluster. This procedure is completed.

6. Contact Oracle

Disaster recovery activity may require real-time assessment by Oracle Engineering in order to determine the best course of action. Customers are instructed to contact the Oracle Customer Access Support for assistance if an enclosure FRU is requested.

My Oracle Support (MOS)

MOS is available 24 hours a day, 7 days a week, 365 days a year:
Web portal (preferred option): My Oracle Support (MOS) at https://support.oracle.com/
Phone: **+1.800.223.1711** (toll-free in the US),
Or retrieve your local hotline from Oracle Global Customer Support Center at http://www.oracle.com/support/contact.html
Make the following selections on the Support telephone menu:
Select **2** for New Service Request Then select **3** for Hardware, Networking, and Solaris Operating System Support

Then: either

4.2.3. select **1** for **Technical Issues**,
When talking to the agent, please indicate that you are an existing Oracle customer.

Note: Oracle support personnel performing installations or upgrades on a customer site must obtain the customer Support Identification (SI) number prior to seeking assistance.

OR

4.2.4. Select **2** for **Non-Technical Issues**, for example, for My Oracle Support (MOS) registration.
When talking to the agent, mention that you are a Oracle Customer new to MOS.

Emergency Response

In the event of a critical service situation, emergency response is offered by the Customer Access Support (CAS) main number at **1-800-223-1711** (toll-free in the US), or by calling the Oracle Support hotline for your local country from the list at *http://www.oracle.com/us/support/contact/index.html*. The emergency response provides immediate coverage, automatic escalation, and other features to ensure that the critical situation is resolved as rapidly as possible.

A critical situation is defined as a problem with the installed equipment that severely affects service, traffic, or maintenance capabilities, and requires immediate corrective action. Critical situations affect service and/or system operation resulting in one or several of these situations:

- A total system failure that results in loss of all transaction processing capability
- Significant reduction in system capacity or traffic handling capability
- Loss of the system's ability to perform automatic system reconfiguration
- Inability to restart a processor or the system
- Corruption of system databases that requires service affecting corrective actions
- Loss of access for maintenance or recovery operations
- Loss of the system ability to provide any required critical or major trouble notification

Any other problem severely affecting service, capacity/traffic, billing, and maintenance capabilities may be defined as critical by prior discussion and agreement with Oracle.