System Maintenance

System Maintenance

Use the platform user login application command line to perform maintenance - either by SSH or from the VM console command line. The password is configured during installation, and can be changed using **system password**. If the password is lost, it must be reset. To do this, please contact Support.

A local home directory is available to the user and can be managed by the user with standard Unix commands: **ls**, **cp**, **mv**, **rm**, **less**, **grep** and **scp**. The user is not permitted to view directories or run commands outside the home directory.

Typing **help** provides information on available commands. Entering the command without further parameters will provide additional information on that command. Refer to the output of the help command for more details.

The following brief command summary can be used to maintain the system.

Command	Description
health	Provides a short system summary indicative of system health.
app status	Checks that all applications are running correctly.
app start	Restarts all services.
diag	Lists the available diagnostic commands.
network <interface> <ip> <netmask> <gateway>	Reconfigures the network interface, for example: **network eth0 192.168.1.2 255.255.255.0 192.168.2.1**
system restart	Reboots the system.
system mount	Mounts any ISO images copied to the system so that the applications within it are visible.

Command	Description
app list	Displays the available software installation packages.
app install <application>	Installs an application.
system provision	Re-provisions the system to ensure that the system components are configured to work together.
voss cleardown	Removes all database transactions and reinitialises the database. All current data is cleared.
cluster status	Displays the status of all nodes in the cluster (if a clustering topology is used).

The command **cluster provision** (and the standalone equivalent command **system provision**) needs to be carried out after the topology of the cluster is changed (for example, if nodes are added or deleted; or when network addresses are changed).

Cluster Failure Scenarios

The status of the cluster can be displayed from the command-line on any node using the command:

cluster status

The system can automatically signal email and/or SNMP events in the event that a node is found to be down.

Refer to the diagrams in the section on deployments.

Loss of an Application role

The Web Proxy will keep directing traffic to alternate Application role servers. There is no downtime.

Loss of a Web Proxy

Communication via the lost Web Proxy will fail, unless some another loadbalancing infrastructure is in place (DNS, external loadbalancer, VIP technology). The node can be installed as a HA pair so that the VMware infrastructure will restore the node if it fails. Downtime takes place while updating the DNS entry or returning the Web Proxy to service. For continued service, traffic can be directed to an alternate Web Proxy or directly to an Application node if available. Traffic can be directed manually (i.e. network elements must be configured to forward traffic to the alternate Web Proxy).

Loss of a Database role

If the primary Database service is lost, the system will automatically revert to the secondary Database. The primary and secondary database nodes can be configured via the CLI using **database weight <ip> <weight>**. For example, the primary can be configured with a weight of 50, and the secondary with a weight of 20. If both the primary and the secondary Database servers are lost, the remaining Database servers will vote to elect a new primary Database server. There is downtime (usually no more than a few seconds) during election and failover, with a possible loss of data in transit (a single transaction). The GUI web-frontend transaction status can be queried to determine if any transactions failed. The downtime for a Primary to Secondary failover is significantly less and the risk of data loss likewise reduced. A full election (with higher downtime and risk) is therefore limited only to cases of severe outages where it is unavoidable. Although any values can be used, for 4 database nodes the weights: 4, 3, 2, 1 is recommended.

Loss of a site

Unified and Database nodes have database roles. The status of the roles can be displayed using **cluster status**. If 50% or more of the database roles are down, then there is insufficient availability for the cluster to function as is. Either additional role servers must be added, or the nodes with down roles must be removed from the cluster and the cluster needs to be reprovisioned. If there is insufficient (less than 50% means the system is down) Database role availability, manual intervention is required to reprovision the system – downtime is dependent on the size of the cluster. Refer to the Operations Guide for details on DR Failover. Database role availability can be increased by adding Database roles, providing greater probability of automatic failover. To delete a failed node and replace it with a new one if database primary is for example lost: The node can be deleted using **cluster del <ip>**. Additional nodes can be deployed and added to the cluster with **cluster add <ip>**. The database weights can be adjusted using **database weight <ip> <weight>**. Finally, the cluster can be reprovisioned with **cluster provision**.

The console output below shows examples of these commands.

The cluster status:

```
platform@cpt-bld2-cluster-01:~$ cluster status


Data Centre: jhb
     application : cpt-bld2-cluster-04[172.29.21.243]
                   cpt-bld2-cluster-03[172.29.21.242]

        webproxy : cpt-bld2-cluster-06[172.29.21.245]
                   cpt-bld2-cluster-04[172.29.21.243]
                   cpt-bld2-cluster-03[172.29.21.242]

        database : cpt-bld2-cluster-04[172.29.21.243]
                   cpt-bld2-cluster-03[172.29.21.242]


Data Centre: cpt
     application : cpt-bld2-cluster-02[172.29.21.241]
                   cpt-bld2-cluster-01[172.29.21.240] (services down)

        webproxy : cpt-bld2-cluster-05[172.29.21.244]
                   cpt-bld2-cluster-02[172.29.21.241]
                   cpt-bld2-cluster-01[172.29.21.240] (services down)

        database : cpt-bld2-cluster-02[172.29.21.241]
                   cpt-bld2-cluster-01[172.29.21.240] (services down)
```

Deleting a node:

```
platform@cpt-bld2-cluster-01:~$ cluster del 172.29.21.245
You are about to delete a host from the cluster. Do you wish to continue? y
```

```
Cluster successfully deleted node 172.29.21.245

Please run 'cluster provision' to reprovision the services in the cluster

Please note that the remote host may still be part of the database clustering
and should either be shut down or reprovisioned as a single node BEFORE this
cluster is reprovisioned
You have new mail in /var/mail/platform
```

Adding a node:

```
platform@cpt-bld2-cluster-01:~$ cluster add 172.29.21.245

Cluster successfully invited node 172.29.21.245

Please run 'cluster provision' to provision the services in the cluster
```

Database weights: listing and adding

```
platform@cpt-bld2-cluster-01:~$ database weight list
    172.29.21.240:
        weight: 5
    172.29.21.241:
        weight: 3
    172.29.21.243:
        weight: 2
    172.29.21.244:
        weight: 1

platform@cpt-bld2-cluster-01:~$ database weight 172.29.21.240 10
    172.29.21.240:
        weight: 10
    172.29.21.241:
        weight: 3
    172.29.21.243:
        weight: 2
    172.29.21.244:
        weight: 1
```