

High Availability

MayaNAS provides high-availability operation using Linux Heartbeat mechanism in

- **Active-Passive**
One resource group is specified with the cluster id number and owned by the active primary node. Only in the event of failure or manual event the resources will be moved
- **Active-Active**
There will be two resource groups specified with different cluster id numbers, and each node gets to own one resource group. Only in the event of failure or manual event the surviving node will own both the resources.

MayaNAS provides high-availability operation in both

- Shared Storage as in dual-path disk arrays, NVMe-oF array, cloud elastic block storage, etc..
- Shared nothing storage requires some form of synchronous mirroring setup:
 - P2P NVMe-oF setup with zpool mirror or Raid Group mirror-1. Provided with MayaScale setup.
 - DRBD synchronous replication . See [Replication Management](#)

Planning for cluster setup

1. Install proper licenses that enables the following optional features
 - Failover option
 - Snapshot option
 - Replication option
2. Select unique number between 1 ..255 for cluster id, and use that value for all resources to be managed as cluster resource. It is recommended to defer any storage pool or volume creation till fail-over configuration is finished.
3. For Google Cloud attach persistent disks to the primary node such that the device names and disk resource name are the same. This is used by HA scripts to fence and forcibly take over disk resources.



gcloud

```
gcloud compute instances attach-disk mayanas-hal --disk=data-1 --device-name data-1
```

4. The node names should match the `uname -n` output and resolve to IP address correctly either with proper `/etc/hosts` entries or DNS. Make sure it does not resolve to localhost 127.0.0.1 address.
5. Decide on the virtual IP address. This is the floating IP address for the currently active server as decided by cluster heartbeat mechanism.



gcloud

For google cloud define secondary ranges which you can use for virtual IP purpose

```
gcloud compute networks subnets update default --add-secondary-ranges range1=10.9.0.0/24
```

6. Optionally you would need ping node ip address (usually the gateway ip address)
7. Creating volume groups or zfs pool require no previous magic header to be present. If needed erase any previous information on the disks with command **wipefs**

Setup using Web GUI

Click **Failover Management** from sidebar menu and then click **New** from Failover Status tab. This opens up the **Setup Failover** dialog as shown below

Failover Management

Setup Failover

New Failover Configuration

Provide failover server and communication parameters.

Create Failover type: ☐ Active/Passive ☐ Active/Active

Primary:	mayanas-ha1	Secondary:	mayanas-ha2
Resource ID:	5	Resource ID:	
Virtual IP address:	10.9.0.100	Virtual IP address:	
Heartbeat Link1:	eth0 <> 10.138.0.8	Heartbeat Link1:	eth0 <> 10.138.0.7
Heartbeat Link2:	eth0 <>	Heartbeat Link2:	eth0 <>
Ping node:			

☐ Advanced Parameters:

Finish

Cancel

1. Select one of the **Create Failover type**. The above shows setup for **Active/Passive** which requires only one **Resource ID** and one **Virtual IP address**.
2. Enter the **Primary** node name. It should match the `uname -n` output on that server. No IP address allowed.
3. Enter the **Secondary** node name and hit **Tab** key. This should result in **Heartbeat Link1**: to be populated with the discovered network interfaces on the secondary.
4. Enter the unique **Resource ID** number to tag the cluster resources. In MayaNAS resources are VGs, zpools, LUN mappings, NFS shares, replication volumes.
5. Enter the **Virtual IP address** that will be floated by the HA resource scripts.
6. Define **Heartbeat Link1** by selecting the network interface and then enter **IP address** of the **peer**.
7. Similarly for secondary side define **Heartbeat Link1** by selecting the network interface and then enter **IP address** of the **peer**.
8. If you have additional network interfaces you may strengthen heartbeat mechanism by filling **Heartbeat Link2** fields.
9. Enter IP address for **Ping node** that is always available. Usually it will be the gateway router address for on-prem setup.
10. Click **Finish** to save the cluster configuration. This will finish cluster setup on Primary and Secondary node also.
11. Successful completion will show Failover Status of **stopped** as shown below.
12. Without starting the cluster services, you may proceed with storage configuration and finish tasks such as creating storage pool, initial volumes, and mappings. Be sure to use the same resource ID number as cluster id for all those operations.

Failover Management

Failover Status

Failover Status on mayanas-ha1

New Standby Takeover **Start** Stop

Node List: mayanas-ha1 mayanas-ha2

Failover Service: **stopped**

Virtual IP Address:

Resource ID:

Auto Failback: off

Heartbeat Links: eth0 10.138.0.8

Heartbeat Parameters: keepalive:1
deadtime:10
warntime:5
initdead:60

Node Status:

13. After that you may start the cluster services by clicking **Start**. This step has to be performed on Primary and on Secondary independently.
14. You can check the status of the cluster services by clicking on **Node Status**. Below it shows that heartbeat links are good and both the nodes are up and running with one node owning the resources.

Failover Management

Failover Status

Failover Status on mayanas-ha1

New

Standby

Takeover

Start

Stop

Node List:

mayanas-ha1 mayanas-ha2

Failover Service:

running

Virtual IP Address:

10.9.0.100

Resource ID:

5

Auto Failback:

off

Heartbeat Links:

eth0 10.138.0.8

Heartbeat Parameters:

keepalive:1

deadtime:10

warmtime:5

initdead:60

Node Status:

mayanas-ha2	mayanas-ha1
Virtual IP address	Virtual IP address 10.9.0.100
Resource ID	Resource ID 5
Role	Role active
Status	Status active
Heartbeat Status eth0 up	Heartbeat Status eth0 up
Resource ownership local	Resource ownership local

15. Congratulations! This completes the initial cluster setup.
16. From primary node click on **Standby** and watch the peer to take ownership of resources.
17. Or from secondary node click on **Takeover** to take ownership of resources from primary node
18. To completely stop the cluster services click on **Stop**