

Sans Digital ACM Configure Details

#1 - Be Patient.

- ❖ It takes time to sync setting between Nodes and adding Nodes.
- ❖ It takes time to create filesystem for LV/Share
- ❖ It takes time to add Public IP

#2 - Plan Ahead.

- ❖ Names for the NAS – Server Name ?
- ❖ NIC Assignment – Name and IP, which is the Public and which is for Synchronization ?
- ❖ Volume Cache – Write-Back or Write-Through ?

Steps to Configuring Cluster

Step-1: “System → General” Settings.

The screenshot shows the 'System > General' configuration page. The 'General' section includes fields for Server Name (NAS1), Server Description (!), Date (3 / 8 / 2012), Time (23 : 27 : 1), Time Zone (America/Los_Angeles), Language (English), and Admin Language (English). Below this is the 'Shutdown/Reboot' section with a 'Delay before shutdown' field set to 0 seconds and a 'Check filesystems on startup' checkbox. Buttons for 'Apply', 'Shutdown', and 'Reboot' are at the bottom.

(A): Set **Server Name**. In our example here, we name them as: “NAS0” and “NAS1”. or you can use name like “n0” / “n1” or “node0” / “node1”. Clustering starts the node from ‘node-0’, so try to make your naming scheme matche it to avoid confusion in the future.

(B) Set the correct Date, Time, and Time Zone. for joining ADS domain will need these to be accurate.

Step-2: NIC Port IP Settings

Go “**Network → TCP/IP**” page:

On **Node-0 (NAS0)**:

Assign the static IP to eth0, this IP is going to be used for the Public port for client access.

Assign the static IP to eth1, this IP is going to be used for Clustering Port / Data Replication between two Nodes.

On **Node-1 (NAS1)**:

Assign the static IP to eth0, this IP is going to be used for the Public port for client access.

Assign the static IP to eth1, this IP is going to be used for Clustering Port / Data Replication between two Nodes.

The image shows two screenshots of a network configuration interface. The top screenshot is for Node-0 / NAS0 and the bottom is for Node-1 / NAS1. Both show a table of Network Interface Cards (NICs) with columns for Port, Status, IP Address, MAC Address, Link, Speed, Duplex, MTU, and Control. Red and blue arrows indicate the intended use of each IP address.

Node-0 / NAS0

Port	Status	IP Address	MAC Address	Link	Speed	Duplex	MTU	Control
eth0	Enabled	192.168.2.8	00:1e:67:30:f6:d5	Up	1000Mb/s	Full	1500	Edit
eth1	Enabled	1.1.1.1	00:1e:67:30:f6:d4	Up	1000Mb/s	Full	1500	Edit

Node-1 / NAS1

Port	Status	IP Address	MAC Address	Link	Speed	Duplex	MTU	Control
eth0	Enabled	192.168.2.12	00:1e:67:30:e9:a7	Up	1000Mb/s	Full	1500	Edit
eth1	Enabled	1.1.1.2	00:1e:67:30:e9:a6	Up	1000Mb/s	Full	1500	Edit

To be used for Public port / client access (Red arrow pointing to eth0 IP)

To be used for Clustering Port / Replication (Blue arrow pointing to eth1 IP)

Step-3: DNS Setting

Go “ **Network** → **DNS**” page.

Set DNS Server info according to your actual DNS server.

If you have 2nd and/or 3rd DNS server, fill in the info too.

Both Nodes need to setup the DNS Server.

Remarks:

If your DNS Server is configured properly and working, you may not need to enter the DNS search path. But in the past, we have found many DNS servers out in the field were not properly configured. So, adding it here would be a like an insurance to guarantee no problem for later operations.

NAS > Network > DNS

DNS

Primary DNS	192	·	168	·	2	·	80
Secondary DNS	0	·	0	·	0	·	0
Tertiary DNS	0	·	0	·	0	·	0

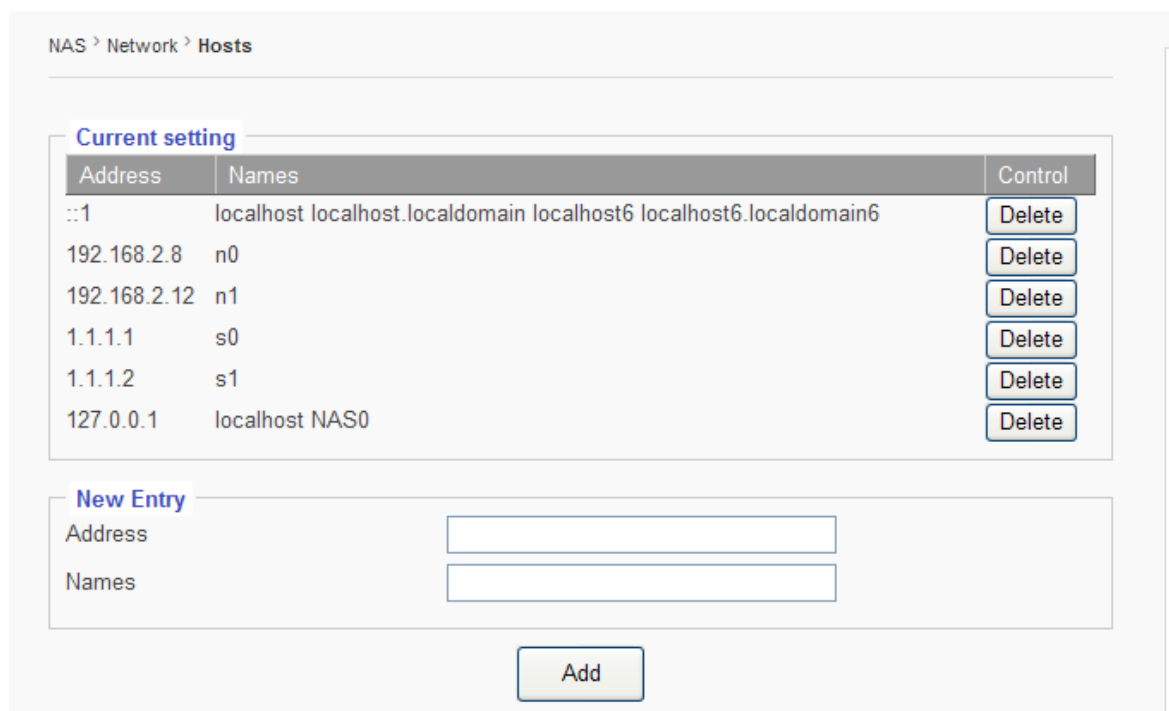
DNS search path

Step-4: Registering Hosts.

Go “**Network → Hosts**” page

Register the IP with associated host name.

The “**host name**” is defined by you for easy identification.



NAS > Network > Hosts

Current setting

Address	Names	Control
::1	localhost localhost.localdomain localhost6 localhost6.localdomain6	Delete
192.168.2.8	n0	Delete
192.168.2.12	n1	Delete
1.1.1.1	s0	Delete
1.1.1.2	s1	Delete
127.0.0.1	localhost NAS0	Delete

New Entry

Address

Names

Add

in our example here, we define:

n0 = 192.168.2.8 -- which is the eth0 on NAS0 (node0)

n1 = 192.168.2.12 – which is the eth0 on NAS1 (node1)

s0 = 1.1.1.1 – which is the eth1 on NAS0 to be the synchronization port.

s1 = 1.1.1.2 – which is the eth1 on NAS1 to be the synchronization port.

Both Nodes need to have the same hosts registrations.

If no hosts registration, when performing Clustering set up, the GUI page might turn Black for a long while before it comes back to the GUI. When it goes Black screen, it is not dead, just that the NAS needs time to find all the nodes.

Remarks:

This step of hosts registration can be skipped if you have a properly configured and working DNS server. But in the past, we have found many DNS servers out in the field

were not properly configured. So, registering hosts here would be a like an insurance to guarantee no problem for later operations.

Step-5: Configuring the Cluster → Adding Nodes

From NAS0/Node0, Go “Clustering → Configuration”, and give it a Name for the cluster and “Apply”

NAS > Clustering > Configuration

Nodes

Cluster Name

GUI displays the extended full page after the name and “Apply”

NAS > Clustering > Configuration

Nodes

Cluster Name

Node	IP Address	Status	Control
Add this Node first.			

New Node IP Address Password

Service control

Select Service to Cluster SMB NFS AFP iSCSI FTP

Public IP Addresses

Public IP Addresses to Failover	Interface	Gateway	Owner	Control
Add Public IP Address	IP Address/Mask	Interface	Gateway (optional)	
	<input type="text"/> /24	eth0	<input type="text"/>	<input type="button" value="Add"/>

Now, add the **eth1 IP** of **NAS0/Node-0** itself.

This is the **s0** (**Synchronization / Replication** Port)

It might take a little bit of time before the GUI comes back from Black.

Be Patient

NAS > Clustering > Configuration

Nodes

Cluster Name

Node	IP Address	Status	Control
Add this Node first.			

New Node IP Address Password

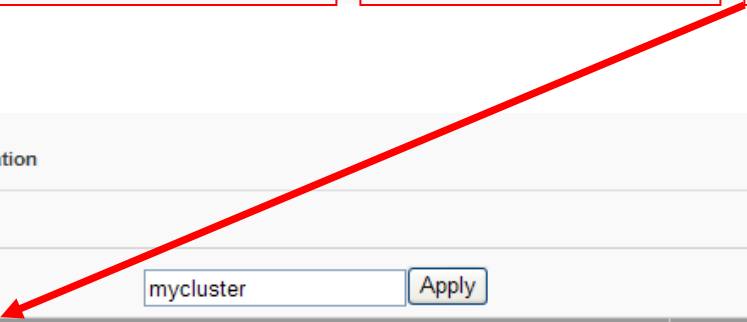
NAS > Clustering > Configuration

Nodes

Cluster Name

Node	IP Address	Status	Control
0	1.1.1.1	UNHEALTHY (THIS NODE) 2012/03/09 01:18:29.848041 [recovered:17079]: Resetting ban count to 0 for all nodes	<input type="button" value="Del"/> <input type="button" value="Stop"/>

New Node IP Address Password



Remarks: No need to worry about the message in the “**Status**” field at this point. It is normal.

Now, stay at the same window from **NAS0/Node-0**, let’s add the **eth1 IP** of **NAS1/Node-1**

Now, stay at the same window from **NAS0/Node-0**, let's add the **eth1 IP** of **NAS1/Node-1**
It might take a little bit of time before the GUI comes back from **Black**.

Be Patient

NAS > Clustering > Configuration

Nodes

Cluster Name

Node	IP Address	Status	Control
0	1.1.1.1	UNHEALTHY (THIS NODE) 2012/03/09 01:18:29.848041 [recoverd:17079]: Resetting ban count to 0 for all nodes	<input type="button" value="Del"/> <input type="button" value="Stop"/>

New Node IP Address Password

NAS > Clustering > Configuration

Nodes

Cluster Name

Node	IP Address	Status	Control
0	1.1.1.1	UNHEALTHY (THIS NODE) 2012/03/09 01:25:17.009596 [recoverd:21508]: server/ctdb_recoverd.c:2055 Reload nodes file from recovery daemon	<input type="button" value="Del"/> <input type="button" value="Stop"/>
1	1.1.1.2	UNHEALTHY 2012/03/09 01:24:31.317781 [recoverd: 6877]: server/ctdb_recoverd.c:3206 Current recmaster node 0 does not have CAP_RECMASTER, but we (node 1) have - force an election	<input type="button" value="Del"/> <input type="button" value="Stop"/>

Remarks: Again, no need to worry about the message in the “**Status**” field at this point. It is normal.

Step-6: Configuring the Cluster → Select Services

Now, stay at the same page “Clustering” on NAS0/Node0, and let’s enable **Services** that is to be “clustered” for **Failover**.

System ▾ Network ▾ Storage ▾ Account ▾ Services ▾ Features ▾ Clu

NAS > Clustering > Configuration

Nodes

Cluster Name

Node	IP Address	Status	Control
0	1.1.1.1	OK (THIS NODE)	<input type="button" value="Del"/> <input type="button" value="Stop"/>
1	1.1.1.2	OK	<input type="button" value="Del"/> <input type="button" value="Stop"/>

New Node IP Address Password

Service control

Select Service to Cluster SMB NFS AFP iSCSI FTP

Public IP Addresses

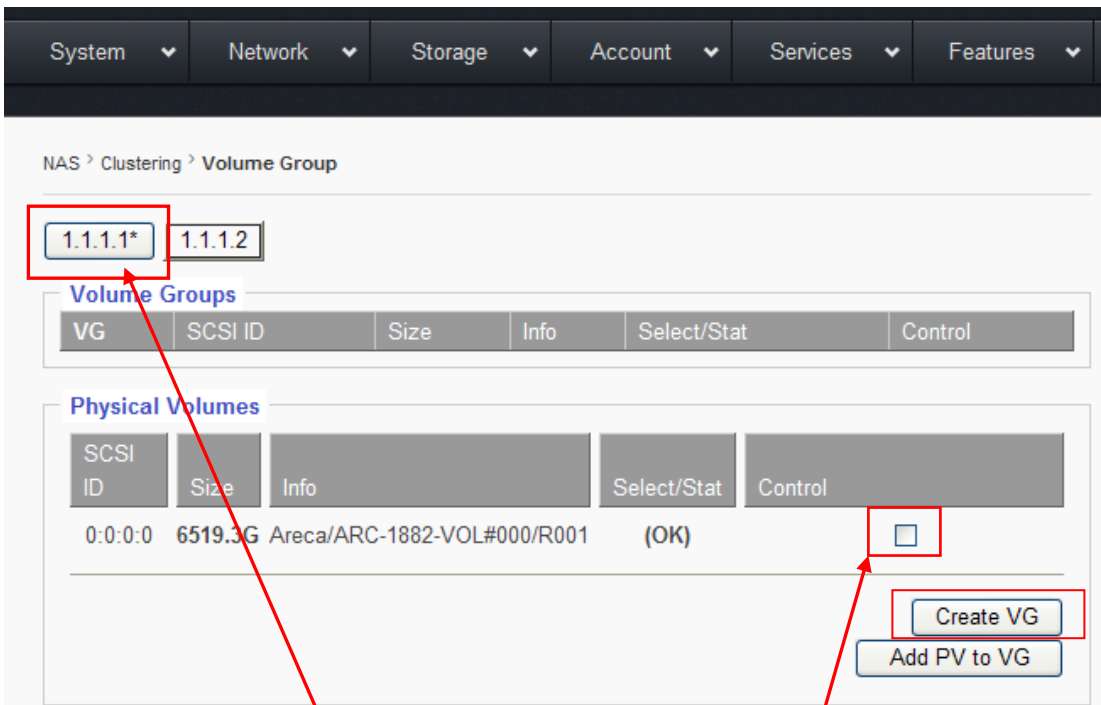
Public IP Addresses to Failover	Interface	Gateway	Owner	Control
Add Public IP Address	IP Address/Mask	Interface	Gateway (optional)	
	<input type="text"/> /24	eth0 ▾	<input type="text"/>	<input type="button" value="Add"/>

Remarks:

For **Cluster mode**, all the **Services** are to be controlled in the **Clustering** page here for **turn-on/turn-off** instead of the pages in “**Services**” top menu.

Step-7: Configuring the Cluster → Create VGs for Cluster on Both Nodes

Stay with **NAS0/Node0**, go “**Clustering → Volume Group**”



In a single page, the GUI displays both Nodes' Volume Group info. The current select node has an “*” after the IP to indicate it is the selected one. In the above picture, IP: **1.1.1.1*** is the current selected node.

No, select the **Physical Volume** by check-marking the “**Control**” box, then click “**Create VG**” button to create Volume Group from the PV.

System ▾ Network ▾ Storage ▾ Account ▾ Services ▾ Features ▾

NAS > Storage > Volume Group

1.1.1.1* 1.1.1.2

Now VG0 is created for Node0 / 1.1.1.1

Volume Groups

VG	SCSI ID	Size	Info	Select/Stat	Control
VG0		6518.3G		<input type="checkbox"/> (OK)	Delete
	0:0:0:0	6518.3G	Areca/ARC-1882-VOL#000/R001	(OK)	

Physical Volumes

SCSI ID	Size	Info	Select/Stat	Control
---------	------	------	-------------	---------

Create VG Add PV to VG

System ▾ Network ▾ Storage ▾ Account ▾ Services ▾ Features ▾

NAS > Storage > Volume Group

1.1.1.1* 1.1.1.2

Click on it

Volume Groups

VG	SCSI ID	Size	Info	Select/Stat	Control
VG0		6518.3G		<input type="checkbox"/> (OK)	Delete
	0:0:0:0	6518.3G	Areca/ARC-1882-VOL#000/R001	(OK)	

Physical Volumes

SCSI ID	Size	Info	Select/Stat	Control
---------	------	------	-------------	---------

Create VG Add PV to VG

Now, let's switch to Node1 / 1.1.1.2 and do the VG creation.
 To switch to Node1 / 1.1.1.2, just click on the tag: **1.1.1.2**
 and notice the IP 1.1.1.2 now has the "*" to indicate it is selected.
 Just do the same thing to create VG for Node1 / 1.1.1.2

NAS > Storage > Volume Group

1.1.1.1 | 1.1.1.2*

Volume Groups

VG	SCSI ID	Size	Info	Select/Stat	Control

Physical Volumes

SCSI ID	Size	Info	Select/Stat	Control
0:0:0:0	6519.3G	Areca/ARC-1880-VOL#000/R001	(OK)	<input type="checkbox"/>

Create VG | Add PV to VG

Warnings:

VGs MUST be the same size from both Nodes !

Step-8: Configuring the Cluster → Create Clustered LV

Go “Clustering → Storage” page to create Clustered Logical Volume (LV)

NAS > Clustering > Storage

Cluster Volumes

Volume Name	Type	Status	Bricks	Select/Control
-------------	------	--------	--------	----------------

Bricks

Node	Brick Name	Type	Size/Usage	Select
------	------------	------	------------	--------

Create Bricks at LV page.

Timeout

Create Cluster Volumes

Volume Type File System iSCSI

Bricks Use Existing Bricks Mirrors Write Through

Select Volume Group

Volume Name

Volume Description

Volume Size (MB)

Warnings:

There are many functions and features in this page.

But for now, we are only supporting what are listed and guided below.

Volume Type:

File-System: when selected, the LV is using **EXT4** as file system, and what you are created is to be a network share.

iSCSI: when selected, it is going to be an iSCSI target volume. *For now, the iSCSI for clustering is still under debugging.*

Bricks: For now we do not support any of its options here except the **Write Through**. **Write Through** – by default, all LVs (network share or iSCSI target) are created to be **Write-Back** mode. By check-marking this “Write-Through” option, the LV is created to be **Write-Through** mode. *The LV mode is not changeable or interchangeable once a LV is created.*

“**Write-Back**” mode gives better Write performance, while it have the risk of losing data in the cache if any thing happens, such as power outage, system locked up, etc.

On the other hand, “**Write-Through**” gives better Data Safety/Integrity, but the Write performance is significantly slower.

Disclaimer: Which mode to be used is purely up to the users based on their own judgment for the nature of their applications. We are not responsible for data lost whatsoever if the user picks one method over the other. We are simply laying out the facts here, and we are not suggesting one method over the other.

Here are some basic numbers:

Over 1Gb Ethernet

“Write Back” = ~ 80 ~ 90+MB/sec sequential Write.

“Write Through” = ~ 36MB/sec sequential Write.

Over 10Gb Ethernet

“Write Back” = 500+MB/sec sequential Write, or whatever the internal RAID system permits.

“Write Through” = ~ 100+MB/sec sequential Write, regardless how fast the internal RAID system may go.

Select Volume Group.

Here it displays what VG is available to create the LV from. If there is more than 1 VG, scroll down the window the select the desired VG.

Now, after select the VG, give the LV a ‘**name**’ and assign the size to it, then click “**Create**” button.

NAS > Clustering > Storage

Cluster Volumes

Volume Name	Type	Status	Bricks	Select/Control
-------------	------	--------	--------	----------------

Bricks

Node	Brick Name	Type	Size/Usage	Select
------	------------	------	------------	--------

Create Bricks at LV page.

Timeout

Create Cluster Volumes

Volume Type File System iSCSI

Bricks Use Existing Bricks Mirrors Write Through

Select Volume Group

Volume Name

Volume Description

Volume Size (MB)

Remarks:

▪ **If you get an error at this point:**

If you get an error now at this point, it is most likely; you forgot to create another VG in one of the Nodes. Please go back the pages to check if both Nodes have equal number of VGs created with equal size.

▪ **Time needed to Create EXT4 network share / LV**

EXT4 file system takes time to create, and it is almost linearly proportional to the size of the LV. In this example, 6.6TB takes about 13 minutes and 18 seconds to finish.

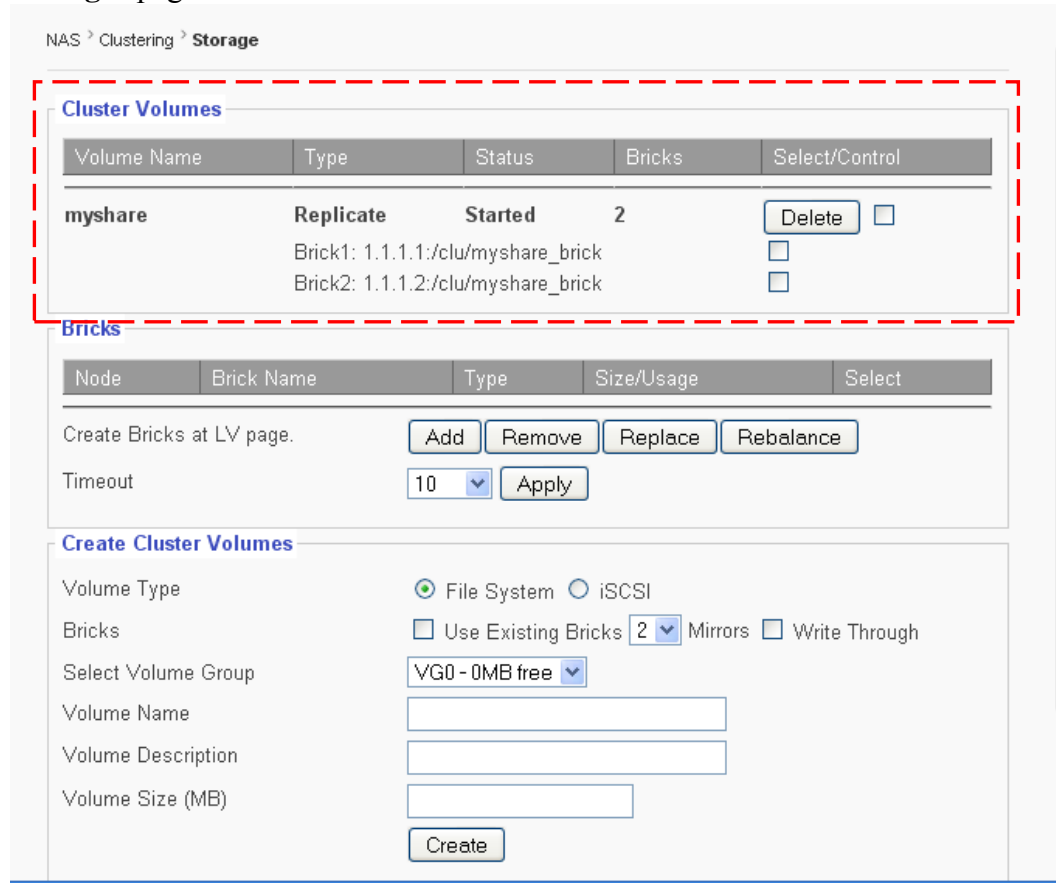
(**6.6TB = ~ 00:13:18**)

During the **EXT4** file system creation, the GUI screen is black and blank, but no need to panic; it will come back after the **EXT4** file system creation is done.

▪ **Time needed to Create iSCSI Target LV**

iSCSI target volume also takes time to create too, but we have found it not to be linearly proportional to the size of the iSCSI target.

This is how it looks after the EXT4 network share / LV created in the “**Clustering → Storage**” page.



Step-9: Configuring the Cluster → Creating Public IPs

Go “Clustering → Configuration” page.

NAS > Clustering > Configuration

Nodes

Cluster Name

Node	IP Address	Status	Control
0	1.1.1.1	OK (THIS NODE)	<input type="button" value="Del"/> <input type="button" value="Stop"/>
1	1.1.1.2	OK	<input type="button" value="Del"/> <input type="button" value="Stop"/>

New Node IP Address Password

Service control

Select Service to Cluster SMB NFS AFP iSCSI FTP

Public IP Addresses

IP Address	Subnet Mask	Interface	Gateway	Owner	Control
<input type="text"/>	<input type="text" value="/24"/>	<input type="text" value="eth0"/>	<input type="text"/>		<input type="button" value="Add"/>

Annotations:

- Enter the IP address here.
- Define subnet mask.
- Define which NIC port to serve data to clients

Do NOT click Refresh button on the Web Browser !!!
It will take a little bit time for GUI to Come Back, Just Be Patient

NAS > Clustering > Configuration

Nodes

Cluster Name

Node	IP Address	Status	Control
0	1.1.1.1	OK (THIS NODE)	<input type="button" value="Del"/> <input type="button" value="Stop"/>
1	1.1.1.2	OK	<input type="button" value="Del"/> <input type="button" value="Stop"/>

New Node IP Address Password

Service control

Select Service to Cluster SMB NFS AFP iSCSI FTP

Public IP Addresses

Public IP Addresses to Failover	Interface	Gateway	Owner	Control
192.168.2.66/24	eth0		Node 0	<input type="button" value="Del"/>

Add Public IP Address

IP Address/Mask	Interface	Gateway (optional)
<input type="text"/> /24	<input type="text" value="eth0"/>	<input type="text"/>

NAS > Clustering > Configuration

Nodes

Cluster Name:

Node	IP Address	Status	Control
0	1.1.1.1	OK (THIS NODE)	<input type="button" value="Del"/> <input type="button" value="Stop"/>
1	1.1.1.2	OK	<input type="button" value="Del"/> <input type="button" value="Stop"/>

New Node IP Address: Password:

Service control

Select Service to Cluster: SMB NFS AFP iSCSI FTP

Public IP Addresses

Public IP Addresses to Failover	Interface	Gateway	Owner	Control
192.168.2.66/24	eth0		Node 0	<input type="button" value="Del"/>

Add Public IP Address

IP Address/Mask	Interface	Gateway (optional)
<input type="text" value=""/> /24	<input type="text" value="eth0"/>	<input type="text" value=""/>

“**Public IP Address**” – it is the service IP that all network clients will be mapped to and accessing from/to. The “/xx” defines the subnet mask. “24” = 255.255.255.0

“**Interface**” – it is the NIC port on each node that selected to be the actual network port that will be serving data underneath the virtual “**Public IP**”. Both Nodes should be using the same **ethX / nicX**; and these ethX/nicX should have been assigned with a static IP in the beginning before configuring the **Cluster**.

“**Owner**” – that means, the **Node-0** (1.1.1.1 / 192.168.2.8) will own the **Public IP** in normal operation. All Write operations will write to the **Node-0 / eth0** first, then, from **Node-0**, data will synchronize to **Node-1** via the dedicated ports 1.1.1.1 and 1.1.1.2.

For all the Read operations, clients will read from **Node-0** via eth0 only.

When **Node-0** Fails, **Node-1** will automatically become the owner of the **Public IP**, and operations will failover to **Node-1**.

This configuration is called: “**Active / Standby**” clustering configuration. The **Node-0** is **Active** and **Node-1** is “**Standby**”.

NAS > Clustering > Configuration

Nodes

Cluster Name

Node	IP Address	Status	Control
0	1.1.1.1	OK (THIS NODE)	<input type="button" value="Del"/> <input type="button" value="Stop"/>
1	1.1.1.2	OK	<input type="button" value="Del"/> <input type="button" value="Stop"/>

New Node IP Address Password

Service control

Select Service to Cluster SMB NFS AFP iSCSI FTP

Public IP Addresses

Public IP Addresses to Failover	Interface	Gateway	Owner	Control
192.168.2.66/24	eth0		Node 1	<input type="button" value="Del"/>
192.168.2.99/24	eth0		Node 0	<input type="button" value="Del"/>

Add Public IP Address

IP Address/Mask	Interface	Gateway (optional)
<input type="text"/> /24	<input type="text" value="eth0"/>	<input type="text"/>

Now, notice there are two **Public IPs** and each has a different owner, one by **Node-0** and the other by **Node-1**.

- 192.168.2.66 – owned by Node-1
- 192.168.2.99 – owned by Node-0

This configuration is called “**Active/Active**” clustering configuration, it achieves the purpose of “**Static Load-Balancing**” while it is also the **Mirroring** and **Failover** clustering. Now, you can assign some of the network clients to access one **Public IP (192.168.2.66)**, and some other clients to access the other **Public IP (192.168.2.99)**.

In **Active/Active** mode, if clients are accessing both Nodes, the Write performance is better than writing to a single Node, and the Read performance is nearly double of reading from a single Node, because now clients are indeed reading from two Nodes.

If users prefer and know how to configure a **Round-Robin DNS Server**, and all network clients can access a single **Virtual IP** provided by the **DNS Server**, but actually accessing these two **Public IPs** in a **Round-Robin Load-Balancing** method.

Disclaimer: How to setup a Round-Robin DNS Server is not our responsibility. Users are the sole party taking on the risk and responsibility for setting /configuring / maintaining the DNS Server.

Step-10: Creating Local Users / Groups

If local users are to be used, now create the local user accounts. Go “**Account – Local User**” page, and create local users.

Only need to create Local User on one Node.
All local users accounts created from a node in an existing cluster will be synced to all other nodes in the cluster.

Also create Local User Groups here if necessary.

Step-11: Join ADS for ADS Users/Groups

If ADS is to be used, now join the ADS from each node. Go “**Account – ADS**” page, fill in necessary info and perform the join action.

***Each Node Must perform the same Join action.** ADS users/groups will not be synchronized between Nodes in the cluster.*

Step-12: Services – NFS for LVs

By default, all and any LV/share, when created, is automatically enabled for SMB/CIFS access. But LV/share is not automatically enabled for NFS access. So if a LV/share is intended to allow NFS client accessing, then, user needs to manually turn on the NFS service for the LV/share.

“Clustering → Configuration” → NFS Service = Enable/Disable NFS Service for the Node/Cluster.
“Services → NFS” = Enable/Disable NFS access to the LV/Share

Go “**Services → NFS**” page, click on the “**Enable**” button, then click on the “**Edit**” button to edit the access permission for the LV/share.

NAS > Services > NFS

Service Control

Enable Service ?

Number of Daemons

Shares

Path	Options	Status	Control
/exports/myshare	public(root_squash,rw,sync)	Disabled	<input type="button" value="Enable"/> <input type="button" value="Edit"/> <input type="button" value="Delete"/>

NFS service for LV/share and permission setting on one Node in the cluster will get synced to all Nodes in the cluster.

Step-13: SMB/CIFS & Access-Control – for LVs and Folders.

Go “**Services → SMB/CIFS**” page to assign basic permission to the LVs(Shares).

Go “**Services → Access Control**” page to assign detail permission to folders / subfolders.

Settings done on one Node will be synced to all Nodes in the cluster.

Here is the Summary for what syncs and what don't sync:

Pre-Condition: In An Existing Cluster	
Services / Actions	Synced to All Nodes in a Cluster
Service Control – Selected Services to be Clustered: SMB/NFS/AFP/iSCSI/FTP (in clustering -> configuration page)	Yes, will synced to all nodes
Create Cluster Volume – (in clustering -> storage page)	Yes, will synced to all nodes
Create Public IP – (in clustering -> configuration page)	Yes, will synced to all nodes
Create Local User / Group	Yes, will synced to all nodes
SMB/CIFS Enable/Disable for LV/share	Yes, will synced to all nodes
SMB/CIFS Permission Setting	Yes, will synced to all nodes
NFS Enable/Disable for LV/share	Yes, will synced to all nodes
NFS Permission Setting	Yes, will synced to all nodes
Access Control – Permission Setting	Yes, will synced to all nodes
Join ADS	No! need to join from each node
Creating VG	No! need to create from each node

Now, you are ready to have some fun, configuring and testing the Failover.