# ELASTIX HIGH AVAILABILITY (HA) MODULE

**Authors: Alfio Muñoz, Alberto Santos**
**Version: 1.1**
**Date: March 28, 2016**

DTPQ

# Custom installation prior to HA module deployment

During this process Elastix will be installed and the required spaces will be left to be used by the partitions that will be created once the process of installation concludes.

You will need two servers, with at least one free partition in each equipment. In this installation two servers with 80GB disk will be used.

## First Server

Proceed normally with the installation until we reach the step where the assistant asks if you wish to partition the disk. A message will warn you that the disk is new and no partitions have been found. Also it will mention the disk will be initialized and all data will be erased.

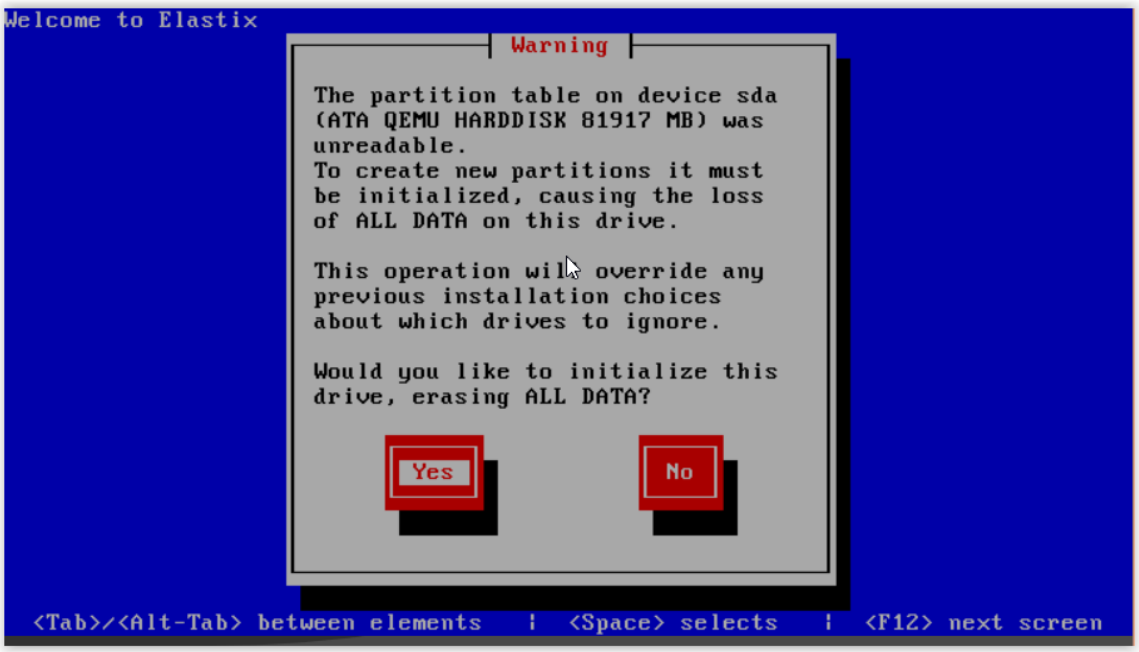


Select *"Yes"* and press ENTER.

In the next screen, select the option *"Create custom layout"*.

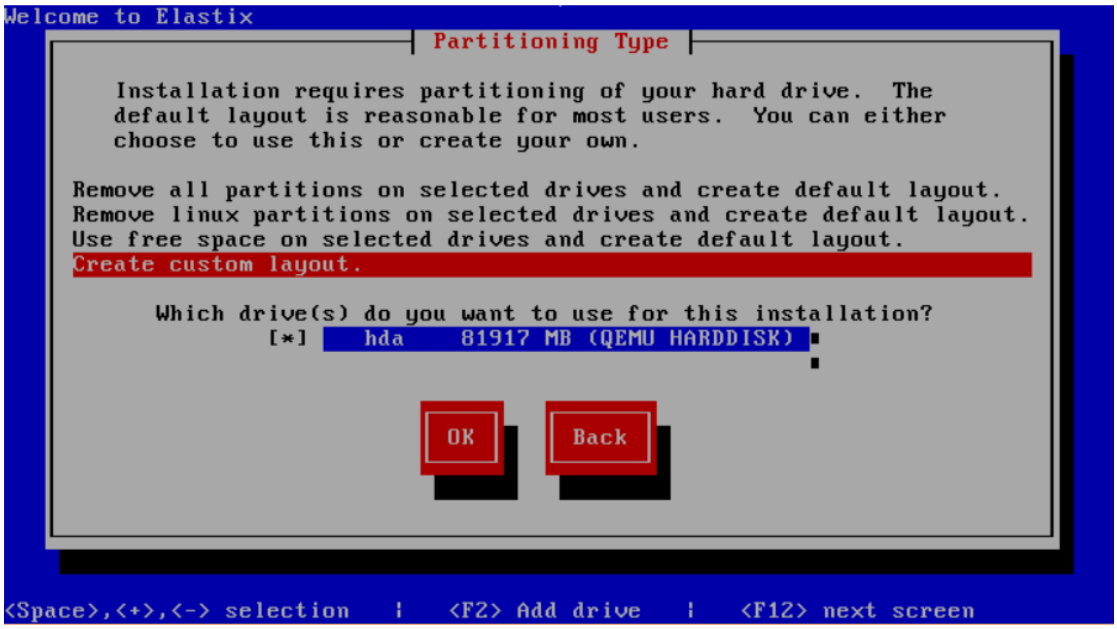


Select "OK" and press ENTER to continue to the Partitioning creation screen.

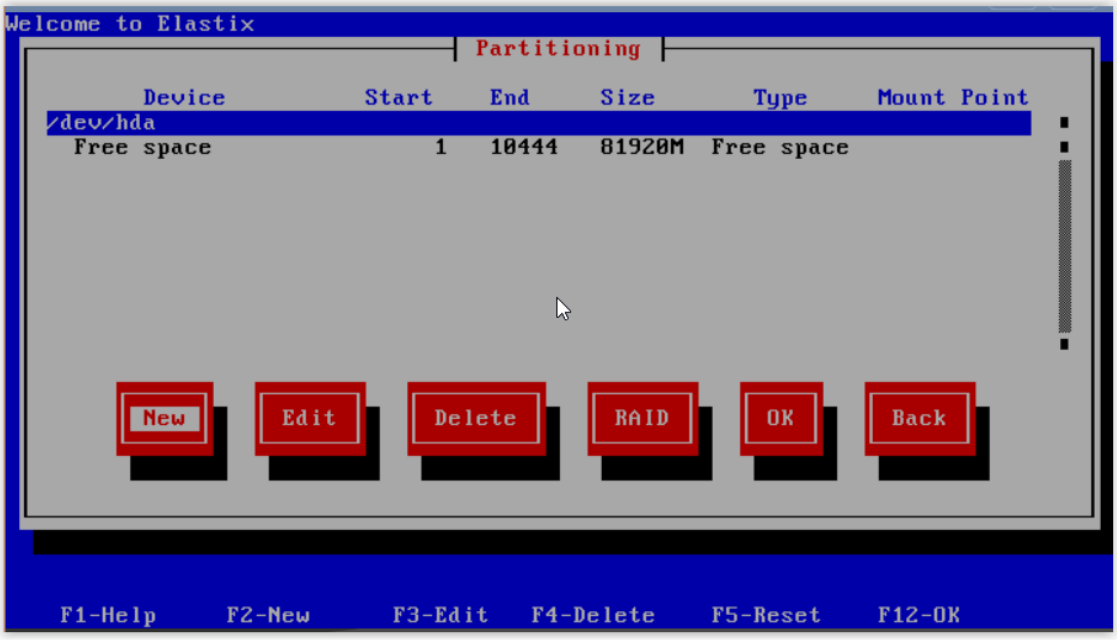

First you need to create the “boot” partition of the operating system. Select “New” and press ENTER.

This partition will be assigned with 101 MB, in the field “File System Type” select the option “ext3”.

In addition, make this partition primary by marking the option “*Force to be a primary partition”*.

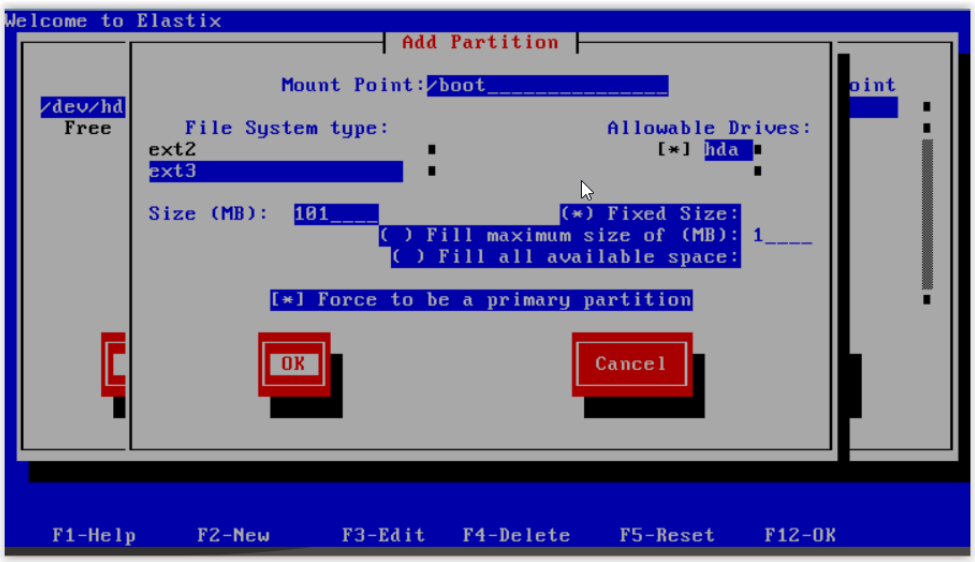


Select “OK” and press ENTER

Then, create the SWAP partition.

In the *“File System type”* field, select SWAP and assign 2048 MB to this partition. Usually this value is the double of the installed RAM memory in the server.

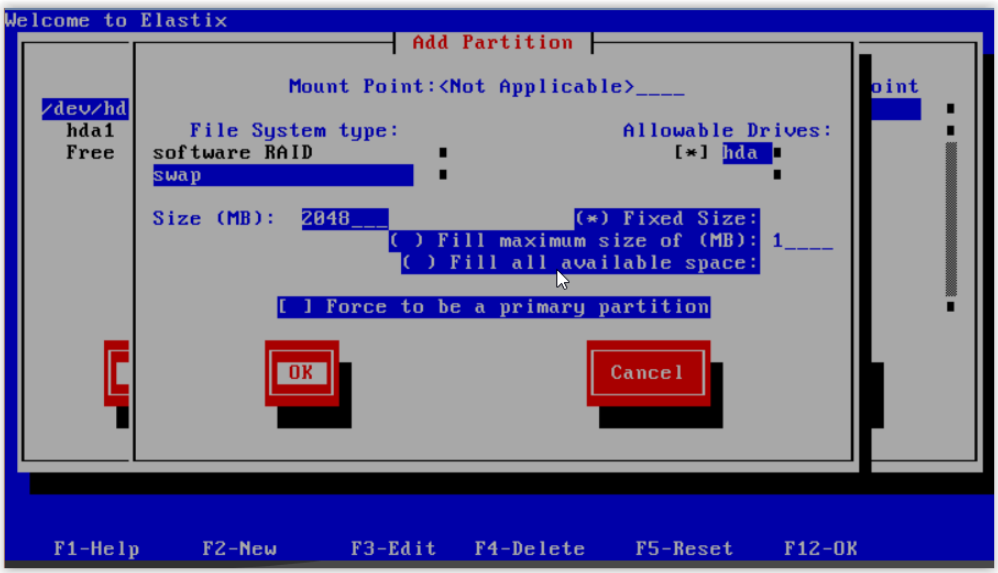

Select “OK” and press ENTER.

Now, create the ROOT partition (/).
Assign 45000 MB to this partition, leaving the rest of space free and without partitioning.

**Note:** When creating the root partition, free space must be left to create the partition where the replication will be mounted, the selected size will depend on the user but it is recommended that at least 60% of total space be destined for the root and the remaining space for the replication.

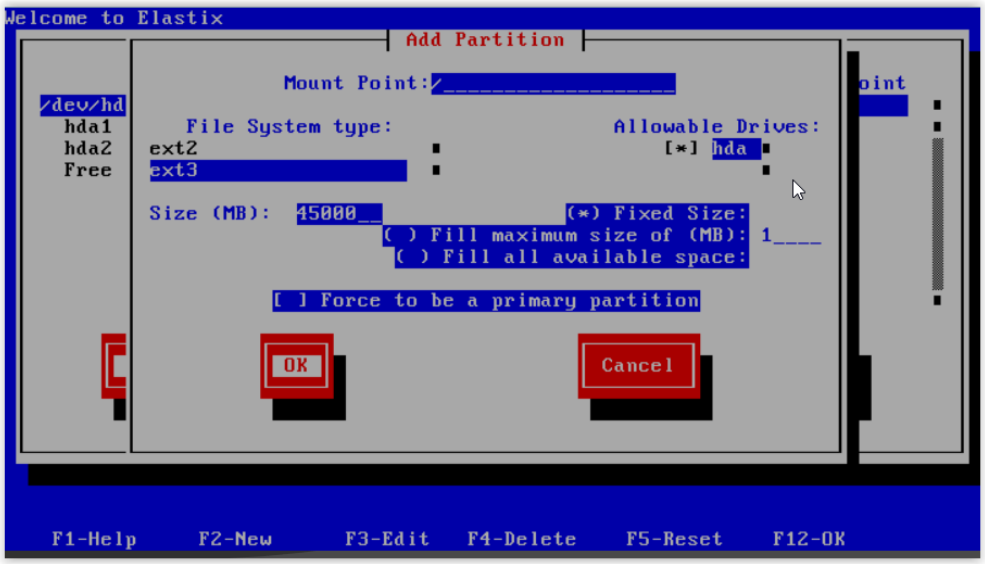


Select “OK” and press ENTER.

After completing this steps your partitions will look similar to the following example:

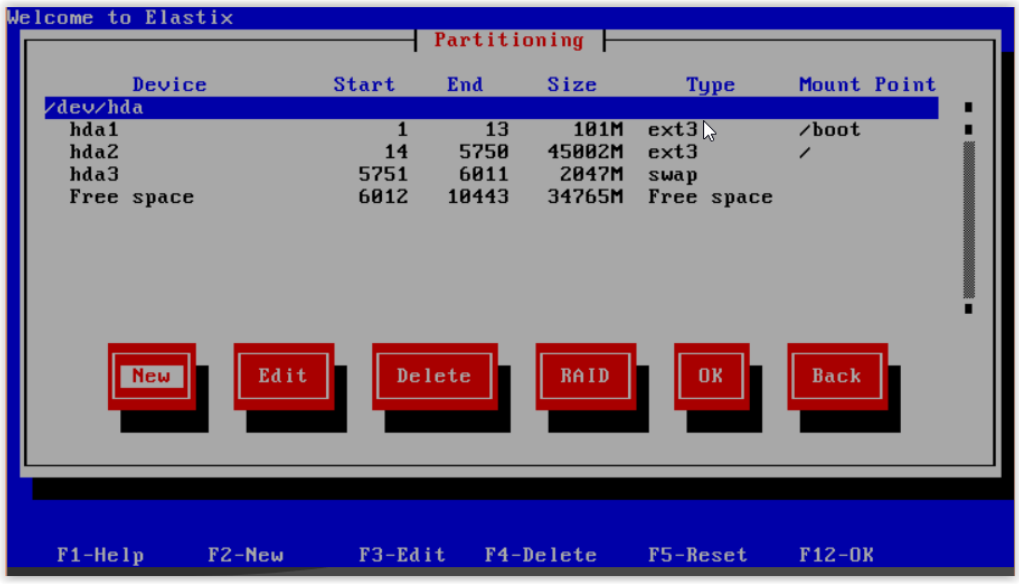

Once the disk partitioning is complete, proceed to the configuration of network adapters. Select “edit” and press ENTER.

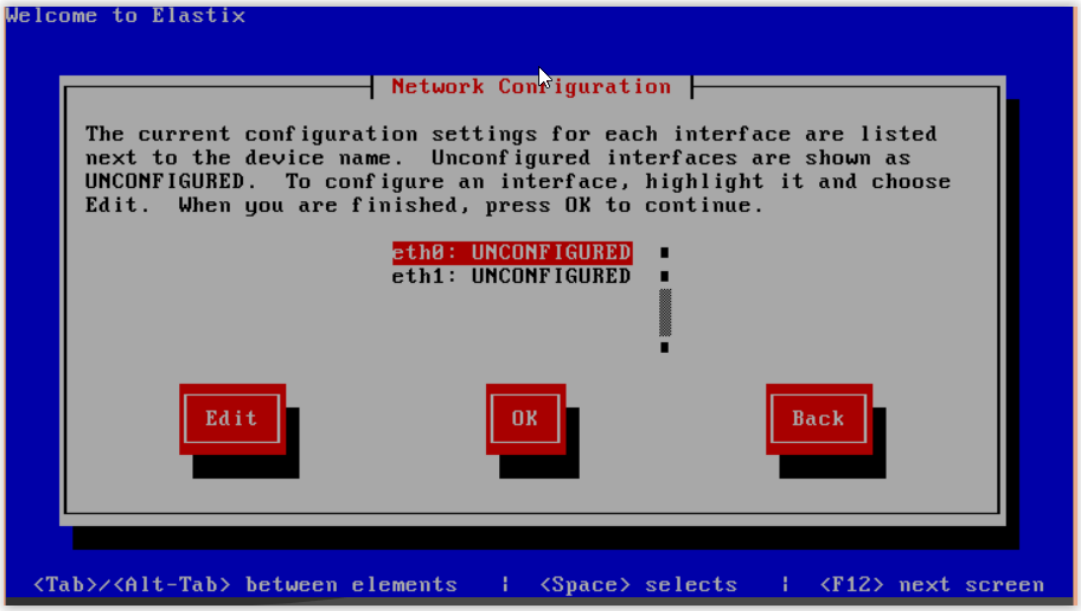


Disable the “Enable IPV6 support” option.

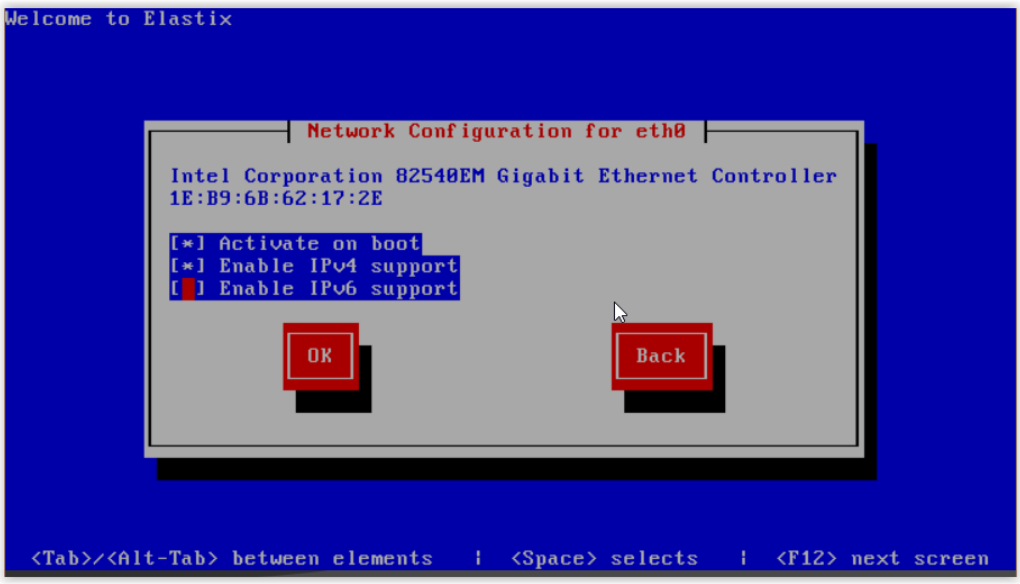


Then assign an IP address; in this example the address 10.0.0.220/24 has been used for the eth0 interface.

**Note:** It is recommended to have two network interfaces, one for the PBX traffic, and the other adapter dedicated to the communication between the two replicated servers.

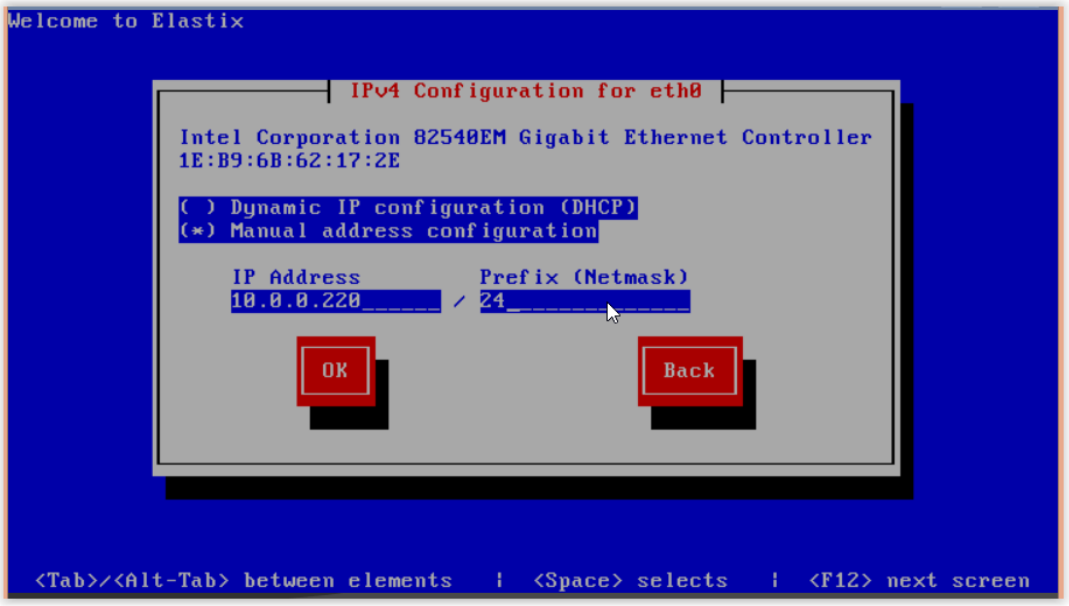


Select “OK” and press ENTER.

Then proceed to the second interface (eth1 in this case), select “Edit” and press ENTER.

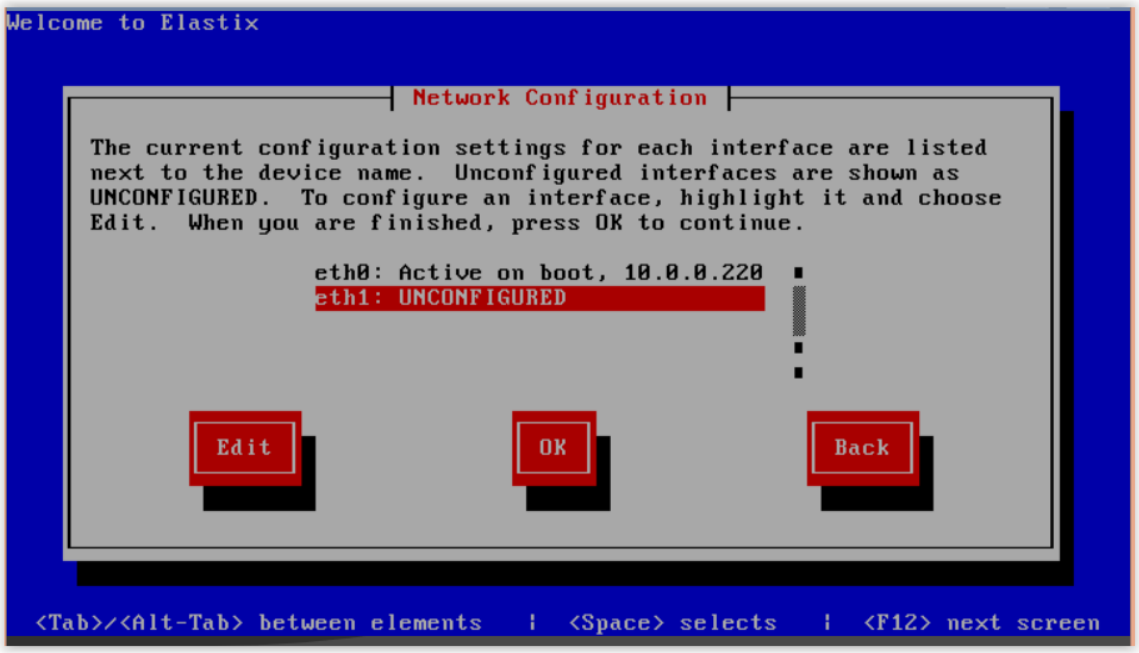


Select *“Activate on boot”* y and uncheck the “enable IPV6 support” option.

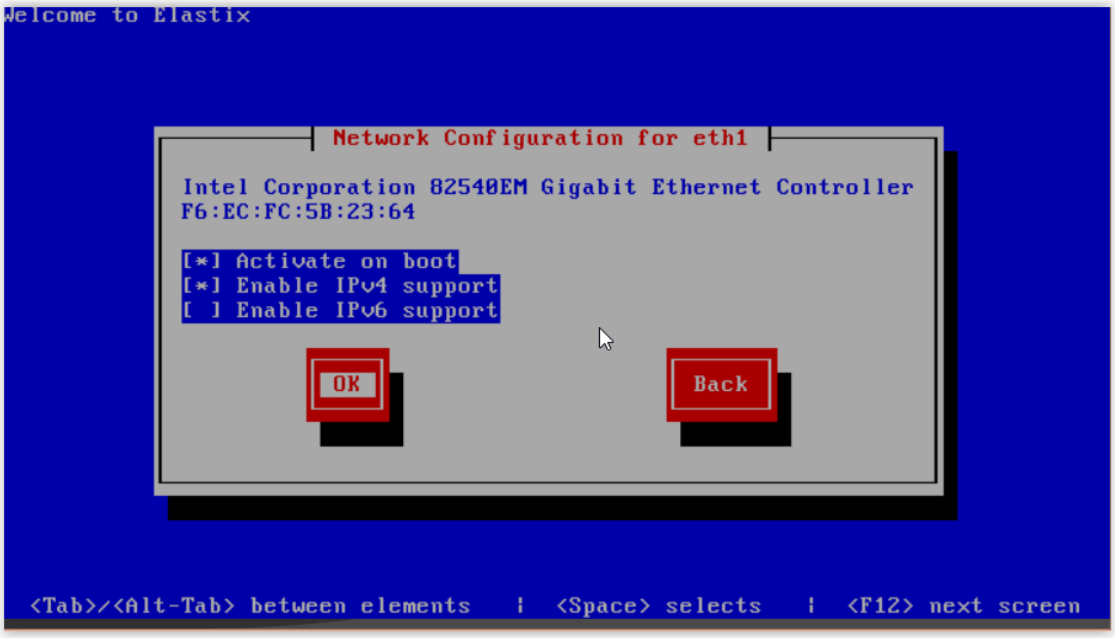


For this interface the IP address 192.168.25.220/24 has been used. It is valid for this example; you can choose another IP depending on your network configuration. With this interface our PBX will communicate to perform the replication transactions.

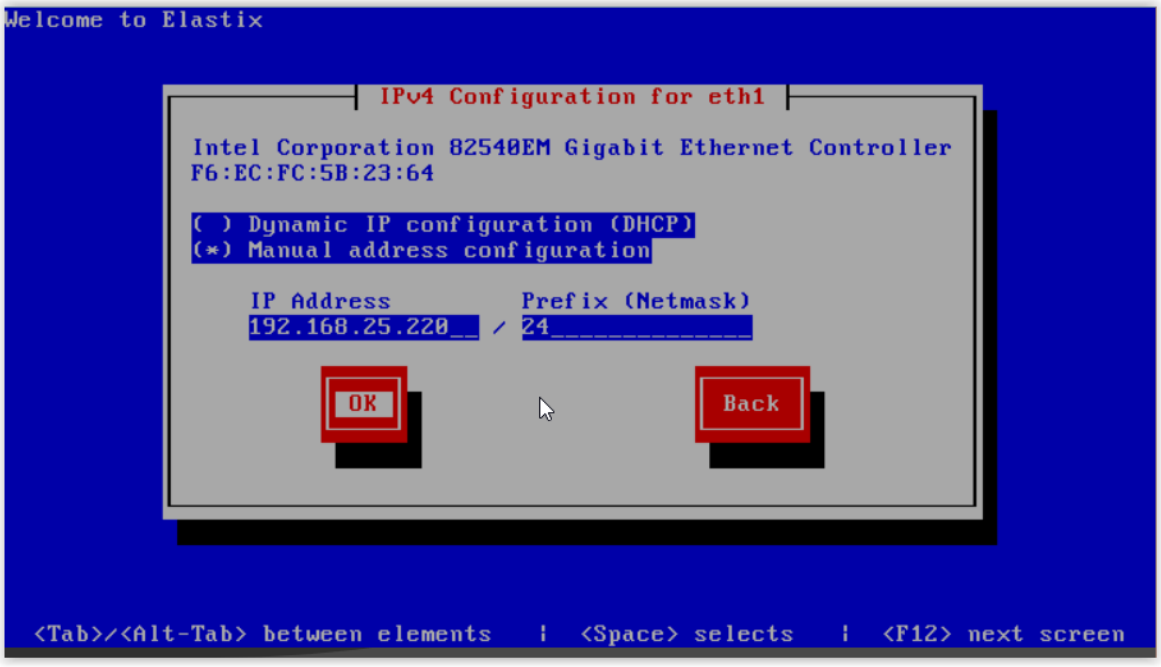


Select “OK” and press ENTER. Then proceed to assign a Gateway and DNS addresses.

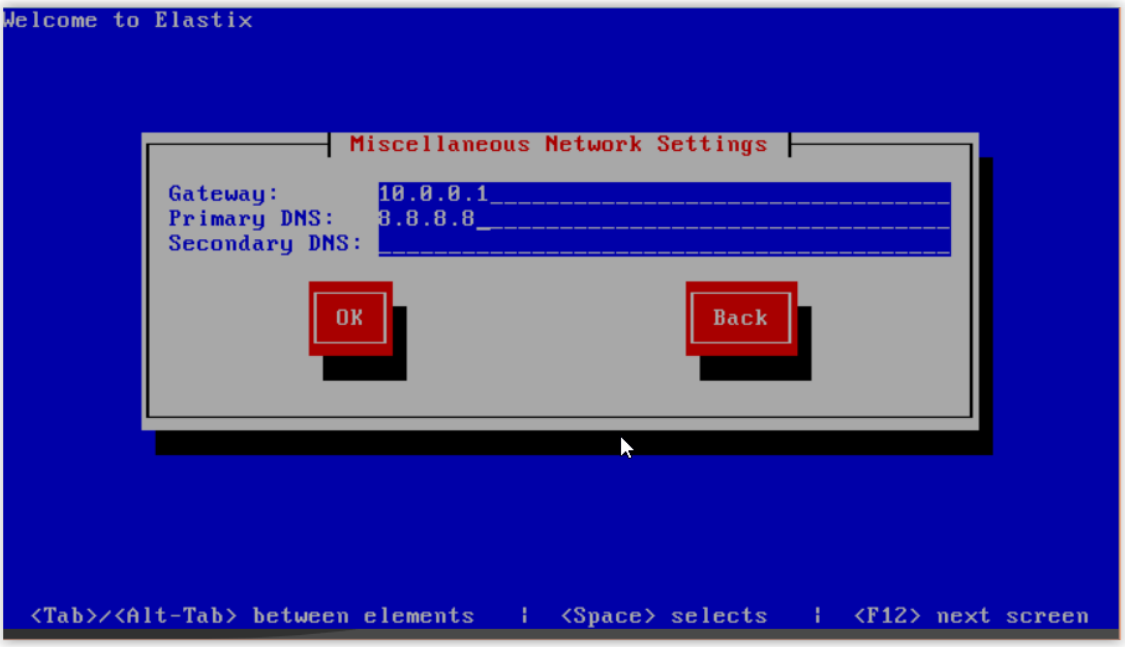


After that, the server host name must be defined, for this example we will use: **Elastix01.teconal.local**

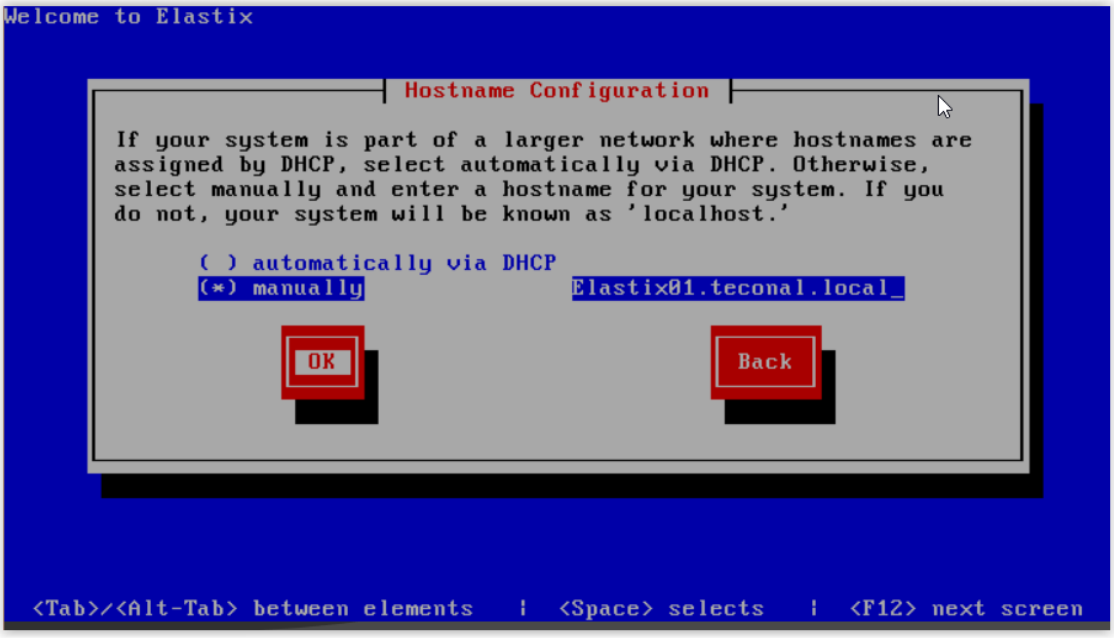


Select “OK” and press ENTER.

Then, select the time zone.

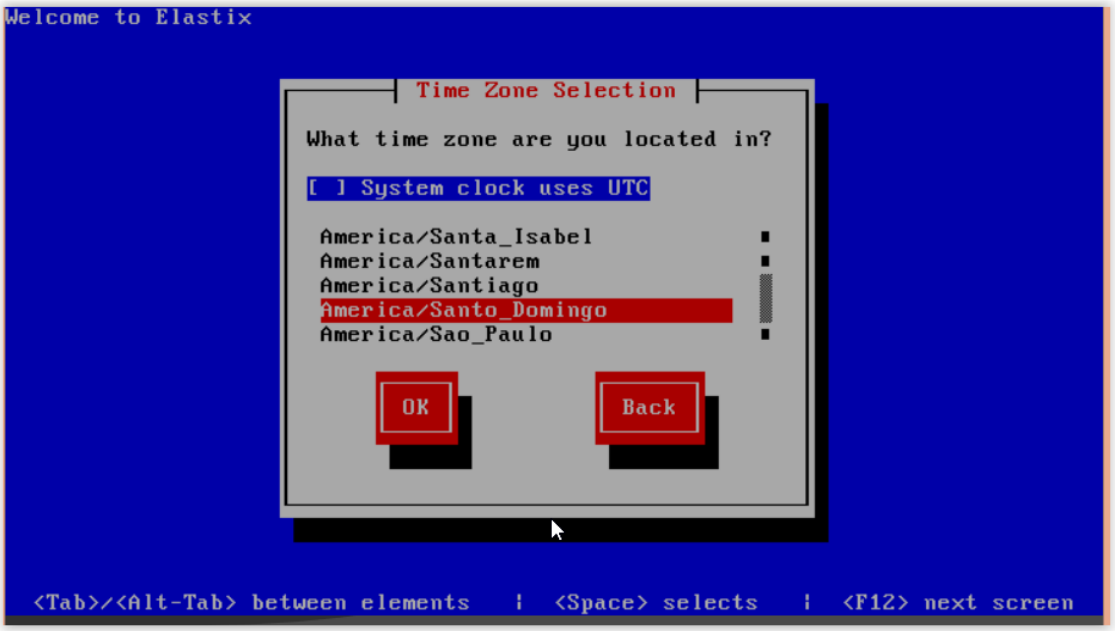


Select “OK” and press ENTER.

Then, define the root password.

**Note:** For production environments please use a strong password to secure your server.

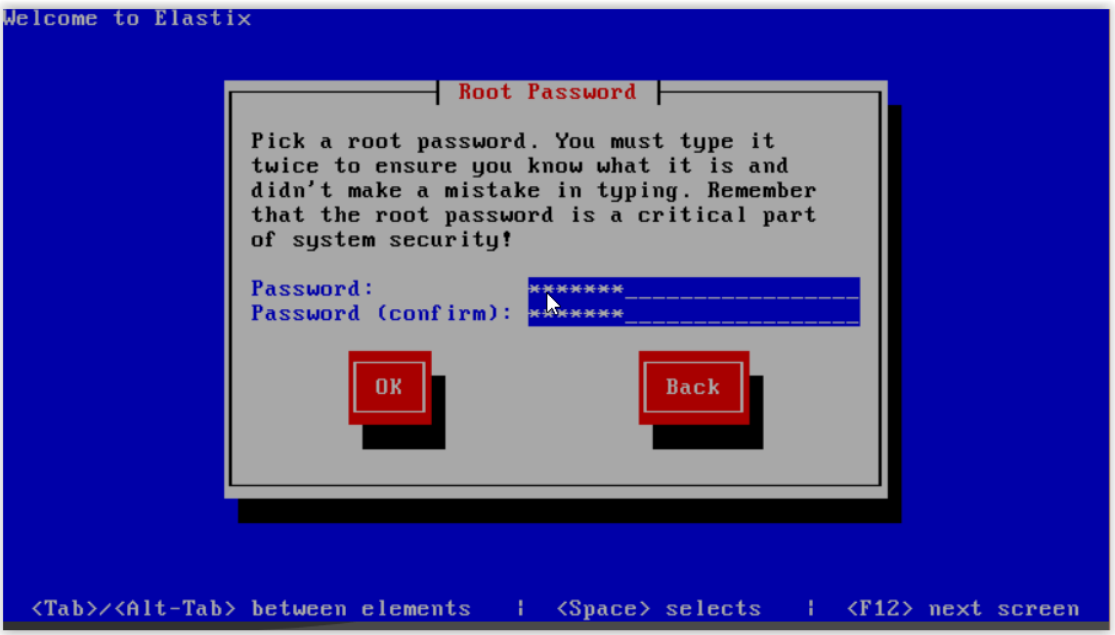


Select “OK” and press ENTER.

Now, the installation process will begin.

Once it finishes, the server will reboot and a screen will appear requesting the MySQL and the admin user from the GUI interface passwords.

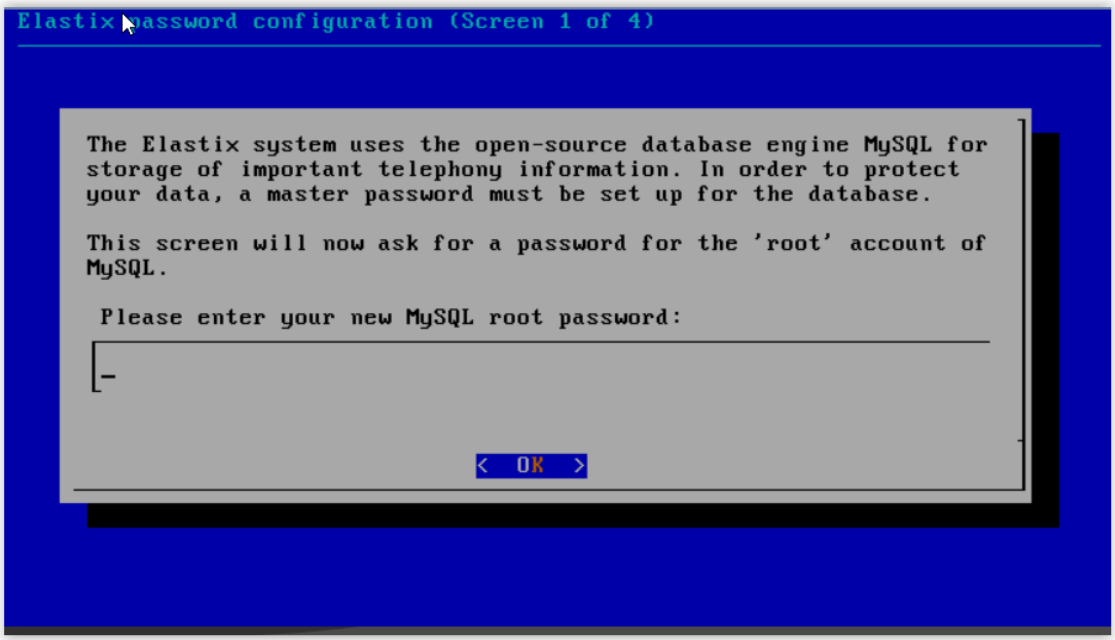


## Second server

Now, we will configure the second server for the replication, following the same steps as the first one. Consider that the IP addresses must be different from the first server; here the IP 10.0.0.221/24 is used for eth0.

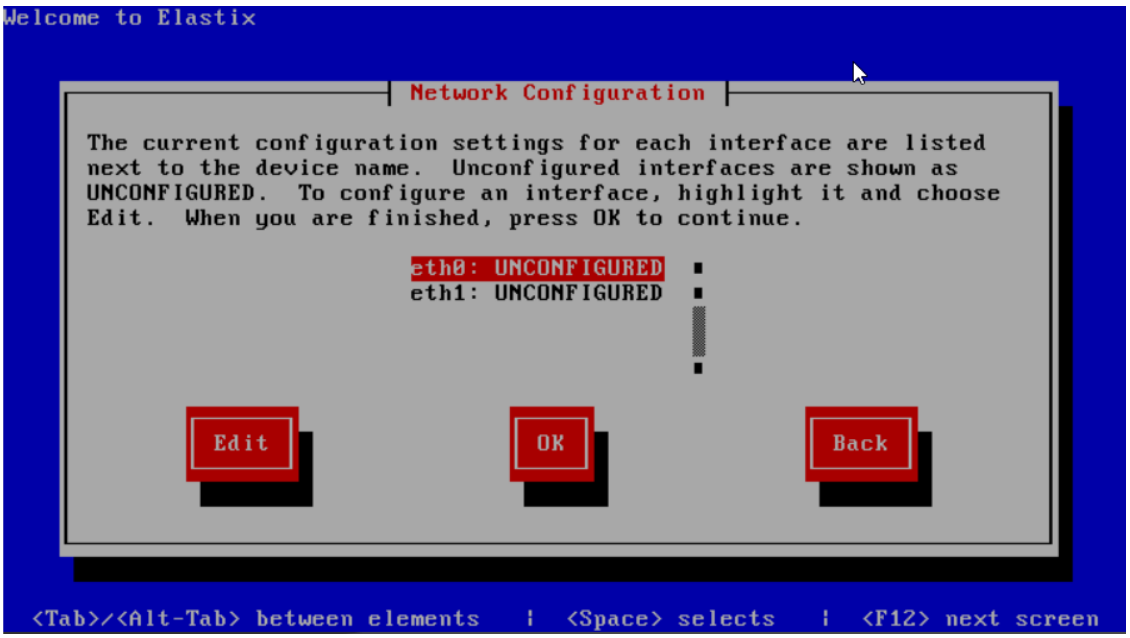

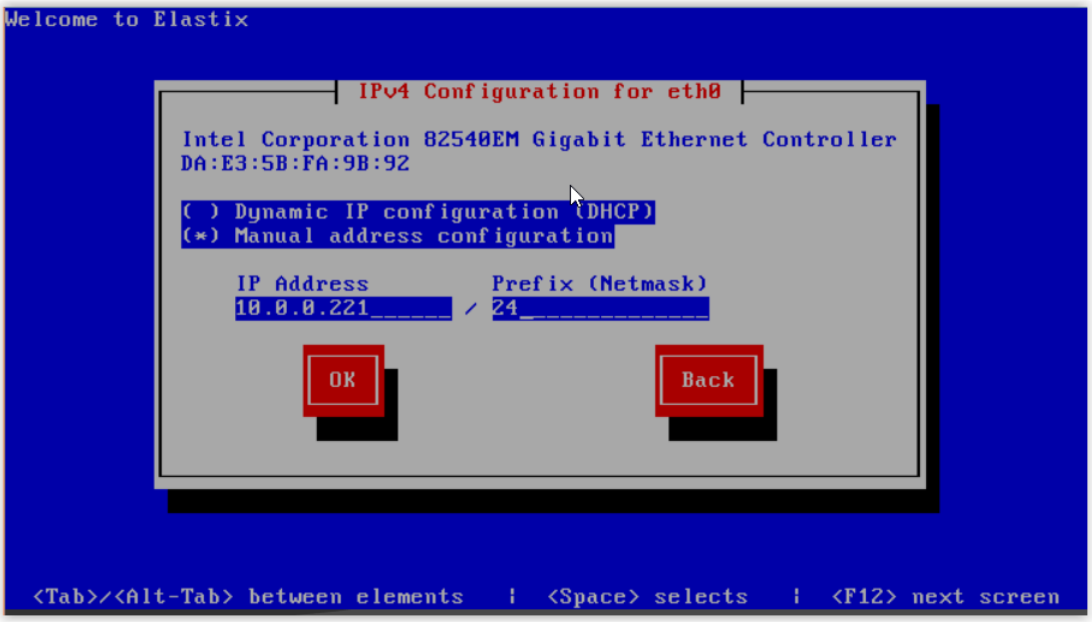


Use the same options for IPV6 support as in the first server

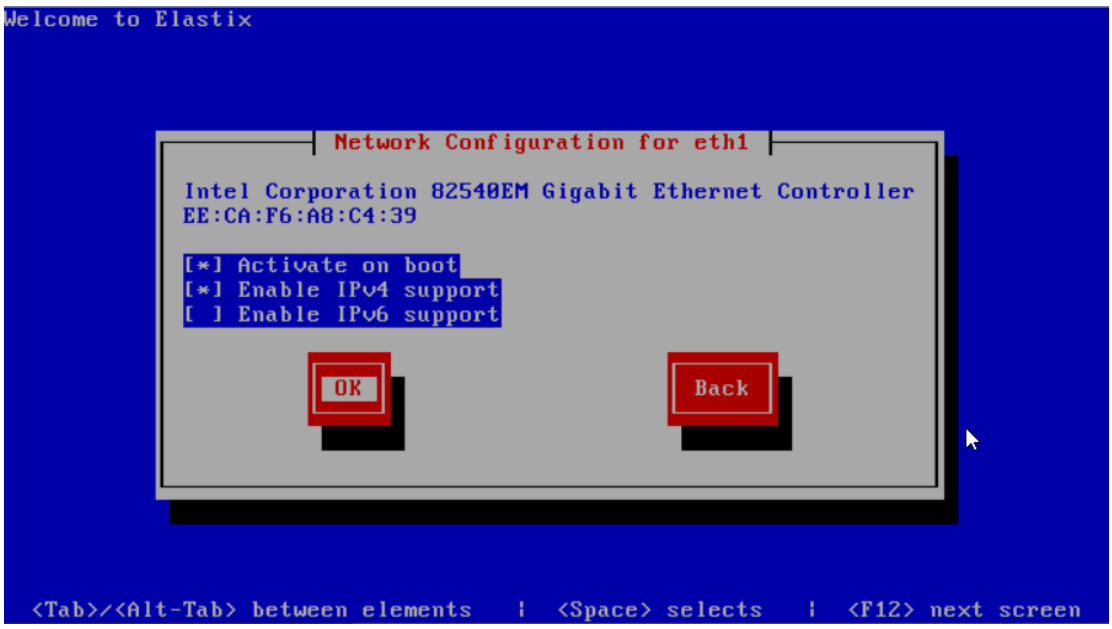


For eth1, the address in the same network segment of the first server will be used, 192.168.25.221/24.

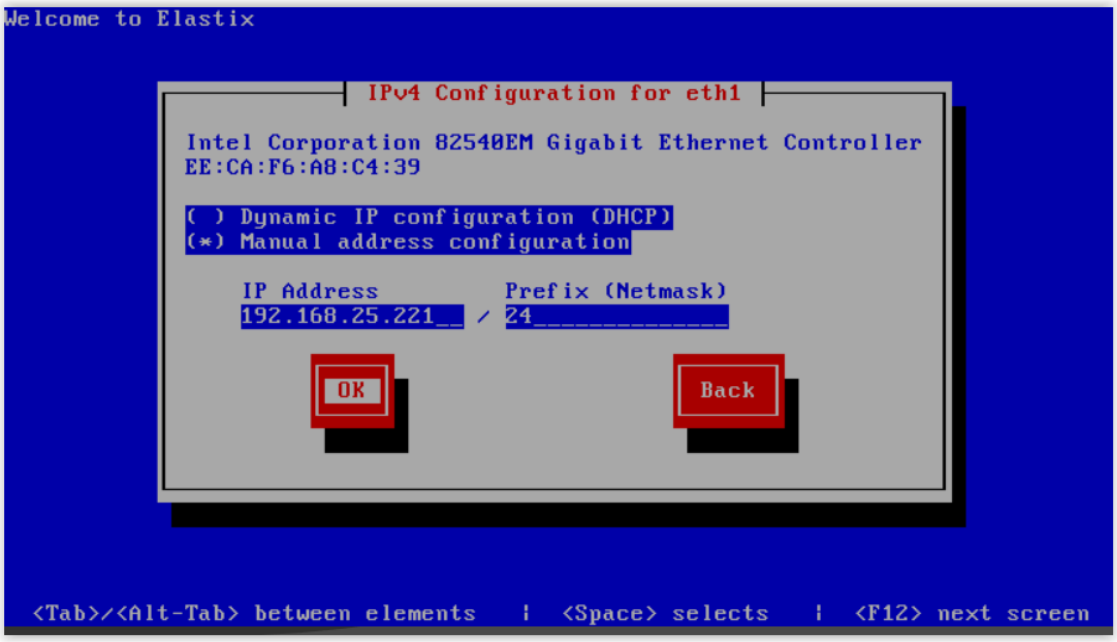


Assign Gateway and DNS addresses and continue.

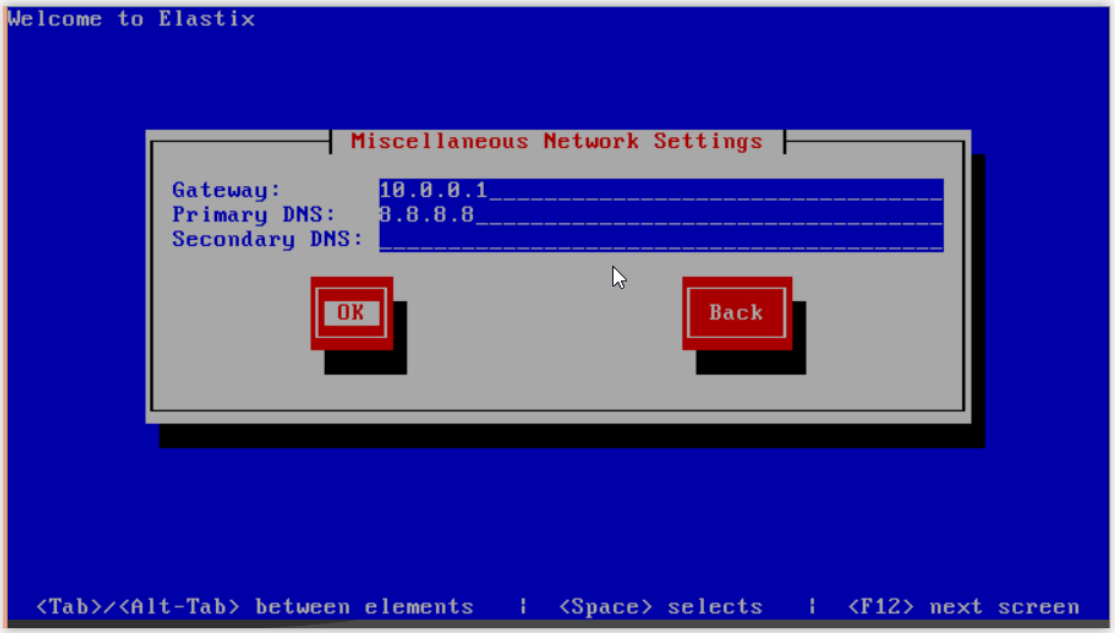


Define the hostname, it is recommended that it keeps similarity with the firs server's name: **Elastix02.teconal.local**

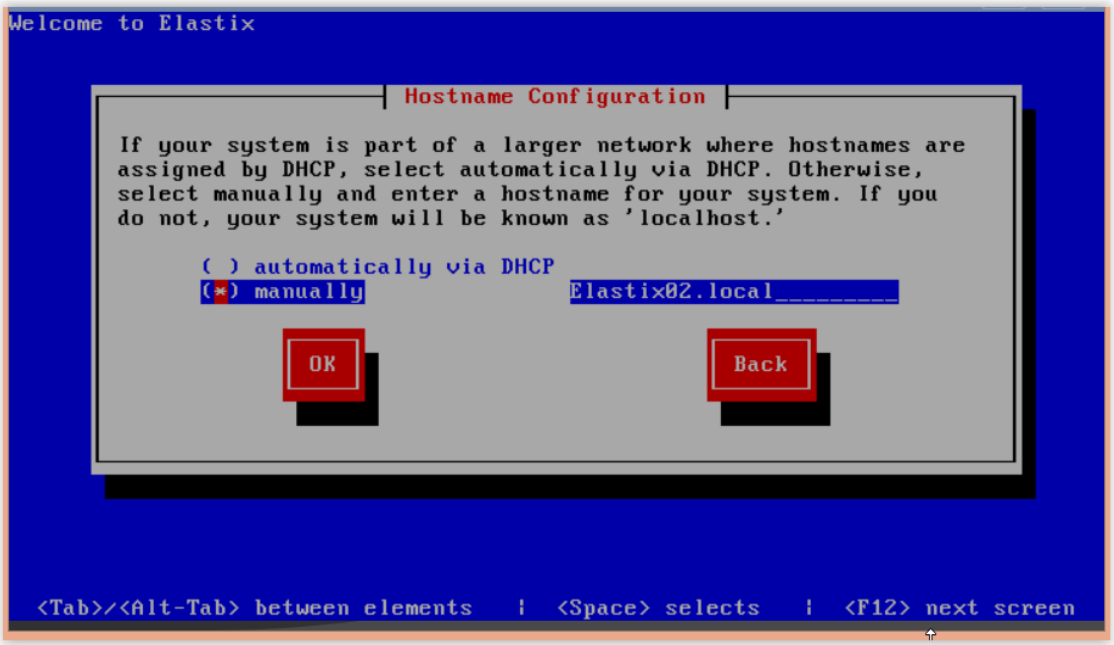


**Note:** It is necessary that the **same root, mysql and admin GUI passwords** are established in both servers.

## Disk partitioning

These steps will create the partitions for the replication. Execute the command `fdisk -l`, the following partitions will appear:

```
[root@Elastix01 ~]# fdisk -l

Disk /dev/hda: 85.8 GB, 85899345920 bytes
255 heads, 63 sectors/track, 10443 cylinders
Units = cylinders of 16065 * 512 = 8225280 bytes

   Device Boot      Start         End      Blocks   Id  System
/dev/hda1   *           1          13      104391   83  Linux
/dev/hda2              14        5750    46082452+  83  Linux
/dev/hda3            5751        6011     2096482+  82  Linux swap / Solaris
```

Now, execute the command: `fdisk /dev/hda`

**Note:** Execute these commands in both servers.

```
[root@Elastix01 ~]# fdisk /dev/hda

The number of cylinders for this disk is set to 10443.
There is nothing wrong with that, but this is larger than 1024,
and could in certain setups cause problems with:
1) software that runs at boot time (e.g., old versions of LILO)
2) booting and partitioning software from other OSs
   (e.g., DOS FDISK, OS/2 FDISK)

Command (m for help):
```

First a new partition will be created, proceed with the following commands after the `Command (m for help)` line:

1. Add a new partition with (n)
2. Select type as *primary* (p)
3. Partition number (4)

```
Command (m for help): n
Command action
   e   extended
   p   primary partition (1-4)
p
Selected partition 4
First cylinder (6012-10443, default 6012):
Using default value 6012
Last cylinder or +size or +sizeM or +sizeK (6012-10443, default 10443):
Using default value 10443
```

4. Press ENTER until the `Command (m for help)` appears again (accept all options by default)
5. Press *"t"* to change the ID of the partition system
6. Press *"4"* to select the partition number
7. Write 83 as the hex code to use *HEX83* partition type
8. Press *"w"* to save changes

```
Command (m for help): t
Partition number (1-4): 4
Hex code (type L to list codes): 83

Command (m for help): w
The partition table has been altered!

Calling ioctl() to re-read partition table.

WARNING: Re-reading the partition table failed with error 16: Device or resource busy.
The kernel still uses the old table.
The new table will be used at the next reboot.
Syncing disks.
[root@Elastix02 ~]#
```

A warning message will appear saying that the partition table re-reading failed, this is normal, reboot both servers after finishing to apply changes.

Now execute `fdisk -l` to verify all partitions and if they were successfully created. A new partition *hda4* will appear.

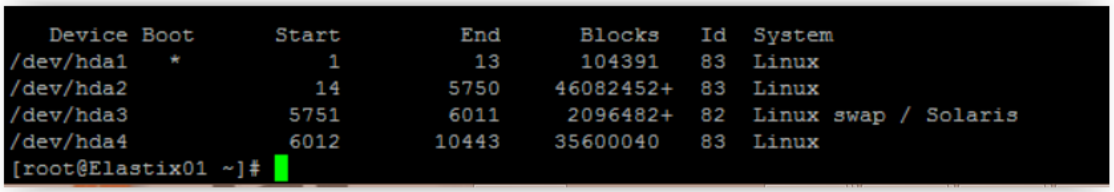


## HA module installation

Now, install the addon of Elastix High Availability by using the `yum` command. Do this in both servers. Execute:

```
yum install elastix-highAvailability
```

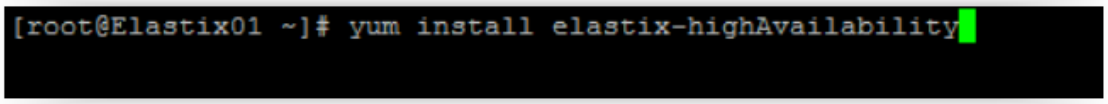


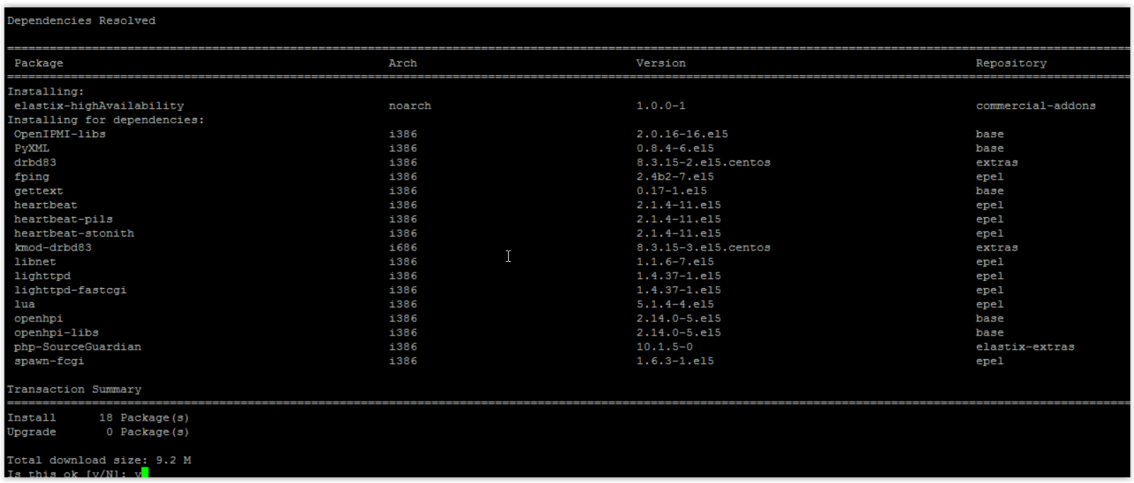


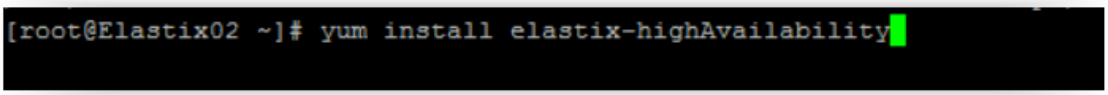



Once the module has been installed, proceed to the web administration interface of the Elastix first server (master) by going to: `PBX` → `High Availability`

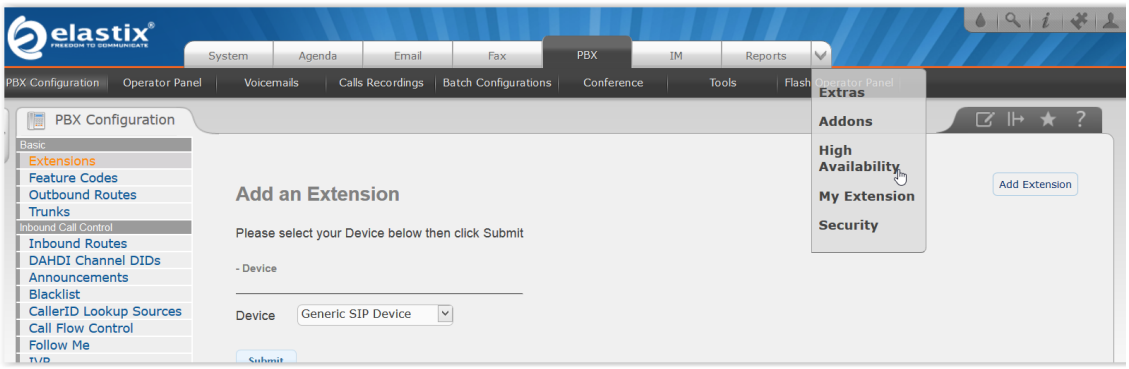


Once there, a page will appear requesting for the valid license key, purchased at http://store.palosanto.com/index.php/elastix-addons/elastix-highavailability.html

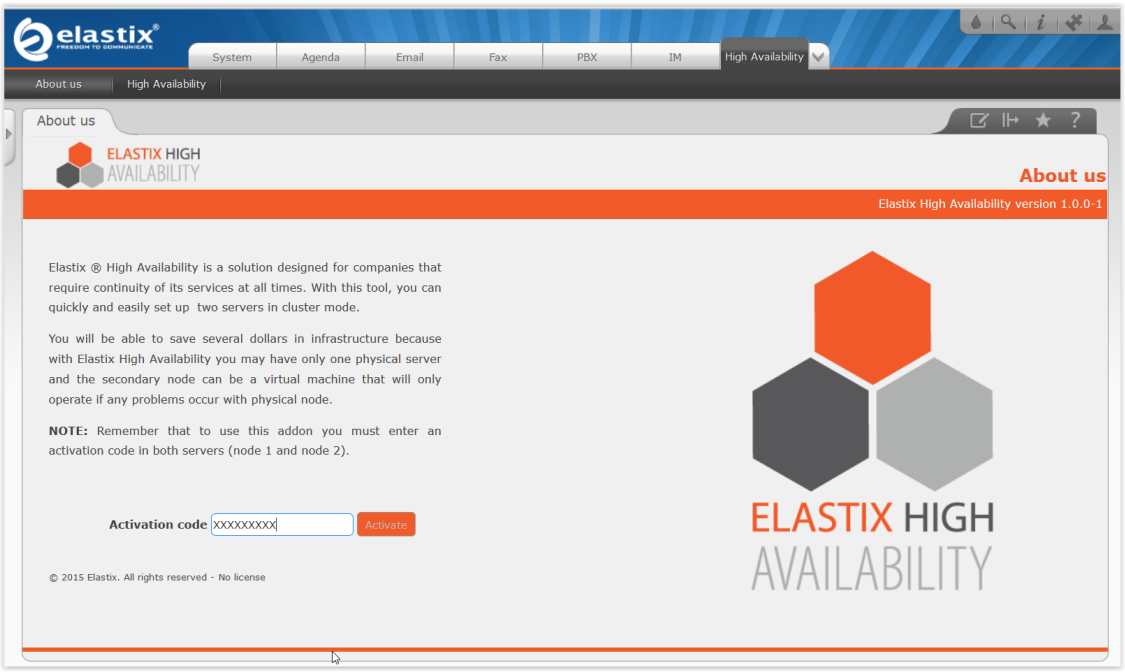


**Note:** This step must be performed on both servers, with a valid license acquired for each server (node).

Once the licenses have been approved, you can begin to configure the addons, by clicking in the sub-menu *"High Availability"*.

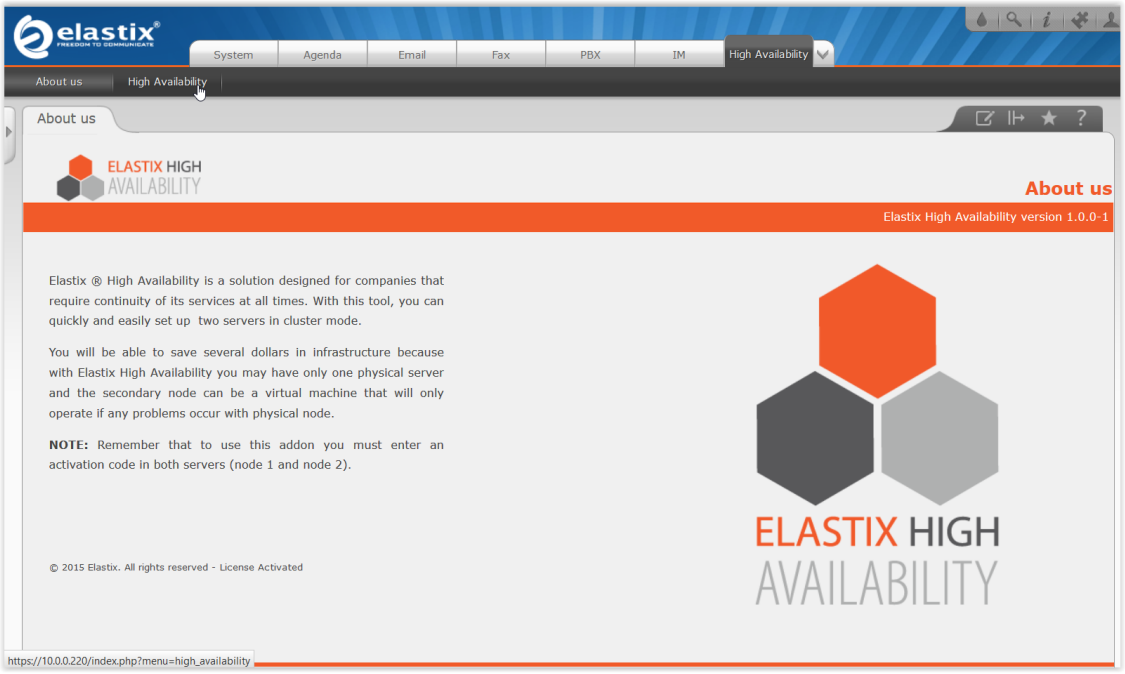


When entering the first time, the module will request the IP address of the second server (*secondary node* or *slave server*) and for its root user password.

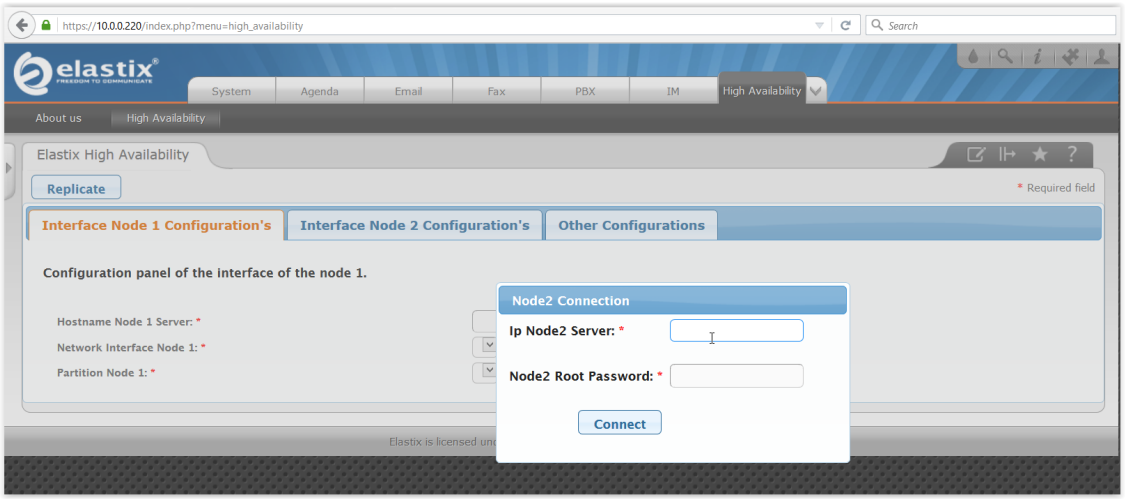


**Note:** In this example the scheme uses 2 network interfaces per server, so the IP that needs to be entered here is the one that will be used for the replication transmissions in the cluster.

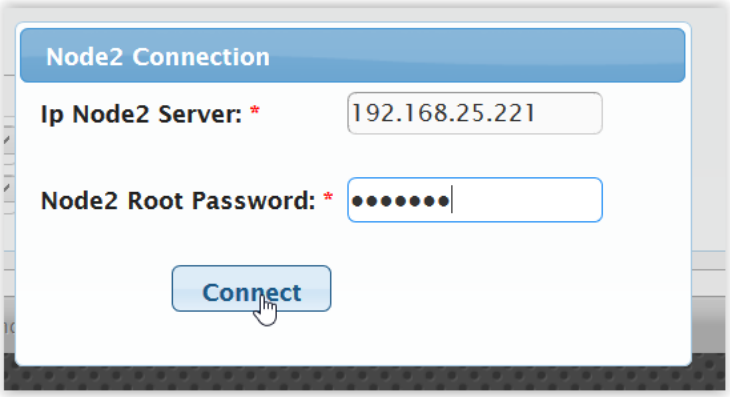


Click on *"Connect"* and wait until the process completes, this can take 10 seconds to 4 minutes.

Then, a window will appear where you will configure all the data in the primary node (or master server): in the *"Interface Node 1 Configuration's"* menu, enter the following information:

- Hostname Node 1 server: Elastix01.teconal.local*
- Network Interface Node 1: 10.0.0.220/24*
- Partition node 1: /dev/hda4*

***Note:** These data is related to the example depicted in this manual, you may see differences in some information depending on the options you entered.

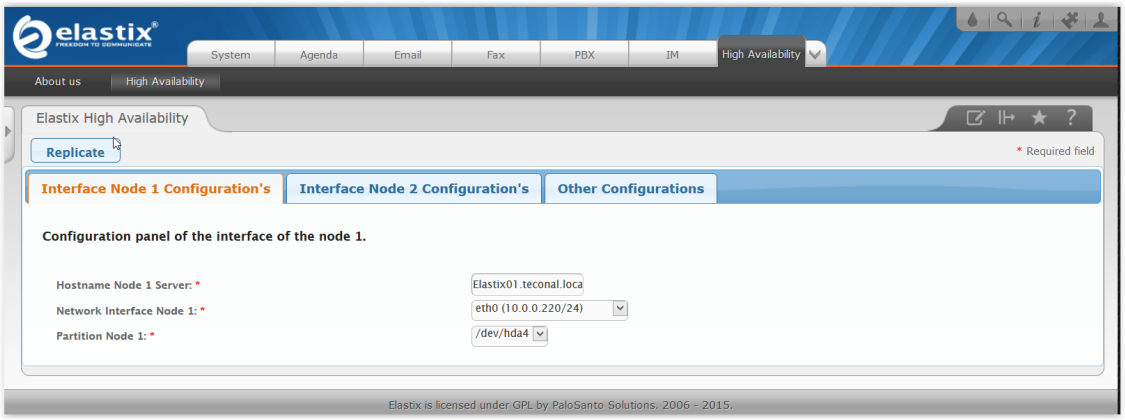


If you don't remember the sever hostname, you can execute the command `hostname` in the command prompt:

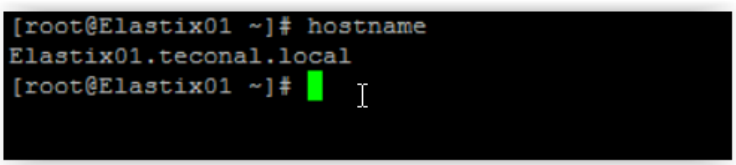

Then proceed to the secondary server (*Interface Node 2 Configuration)* and fill the information similar as you did in server 1.

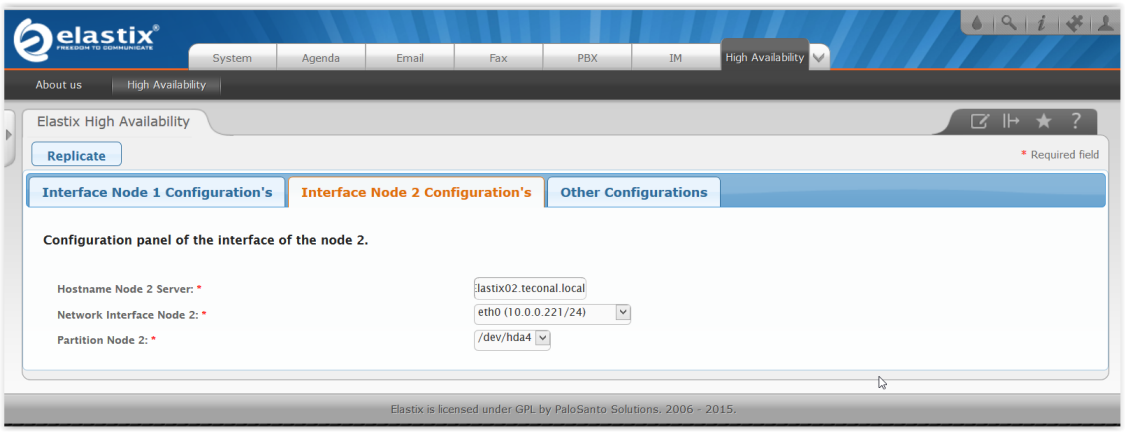


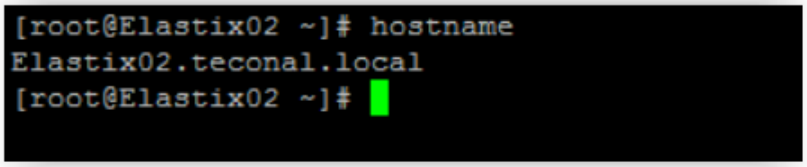


**Note:** In case a different hostname is entered than the one configured in the server, this one will be changed for the new entered hostname.

Once the information of the interfaces in both nodes has been entered, proceed to *"Other configurations"* to configure the following parameters:

- ***Ip Floating*:** This is the IP of the Cluster. This IP will be shared between both servers to deliver services in the network (PBX, SIP, mail, etc). The primary server is the one that answers the requests, if by any reason this one fails, the secondary server will answer the requests on the floating IP.
- ***Network Interface for ip floating in Node 1*:** This is the interface that will assume the floating IP on the primary server, which will reply to requests. It has to be in the same network segment that the floating IP.
- ***Network Interface for ip floating in Node 2*:** This is the interface that will assume the floating IP, it works like in the primary server.
- ***Auth Key*:** It is the key that will be used as authentication between the servers in the cluster. It can be any key.

In addition, there is an option to replicate other directories that may be important for the server administrator depending on the use this Elastix will have, for example, the directory */tftpboot.*

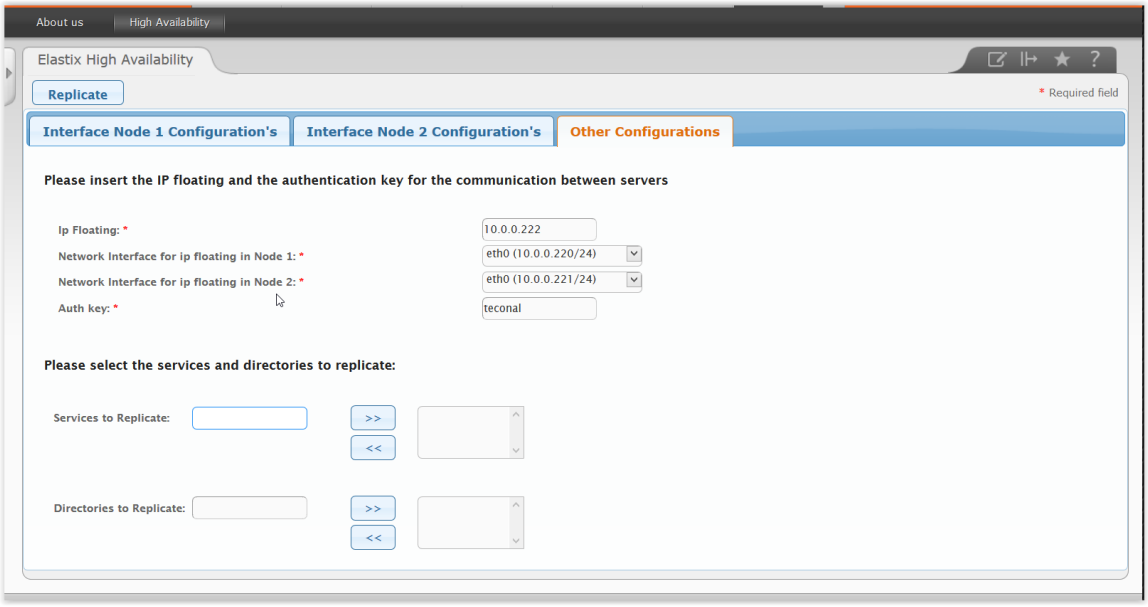


**Note:** By default, the services that are replicated are mysql, apache y asterisk; along with their directories. Add all additional services that you wish to replicate.

Once the configurations have been done, click on *"Replicate"* - button at the upper left part above the menus -.

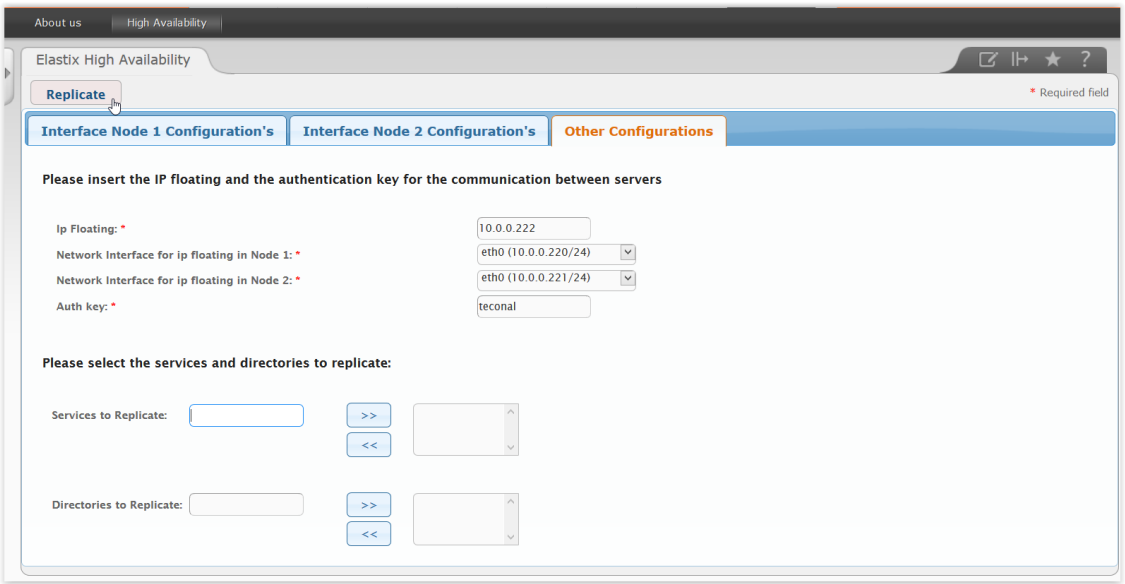


After clicking, a message will warn us that the process is irreversible and that the replication can last several minutes.

Click "OK" to proceed with the replication process and configurations.

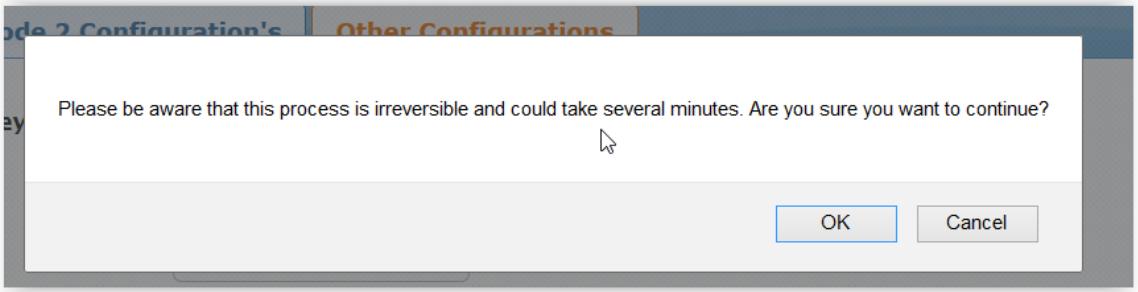


The replication time will vary depending on the total size of the disk and the partitions that were selected. For example, for virtual machines with a small drive size (16GB), the time of replication is approximately 10 minutes. For disks of 500GB, replication time may take as much as 24 hours if the network interface dedicated to the replication is not adequate.

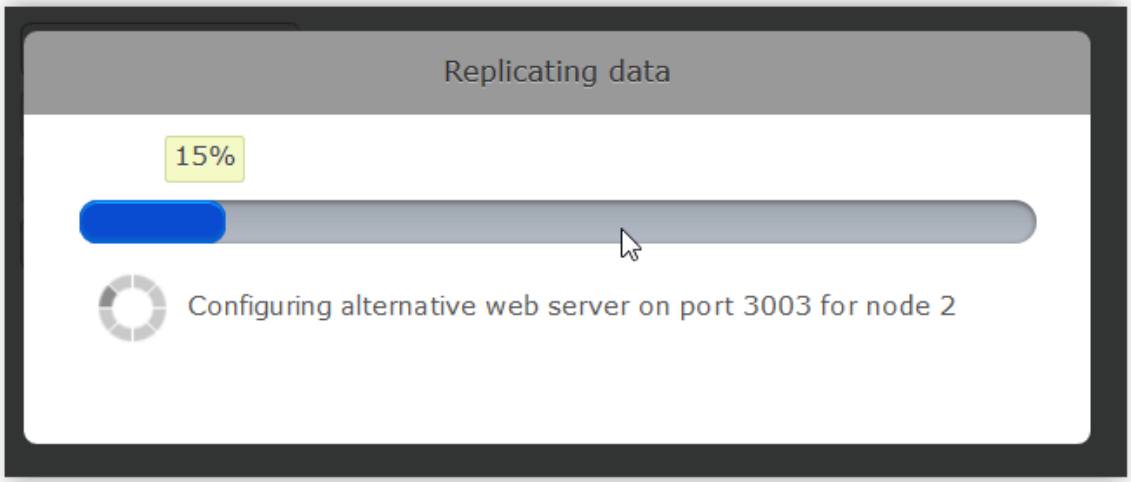


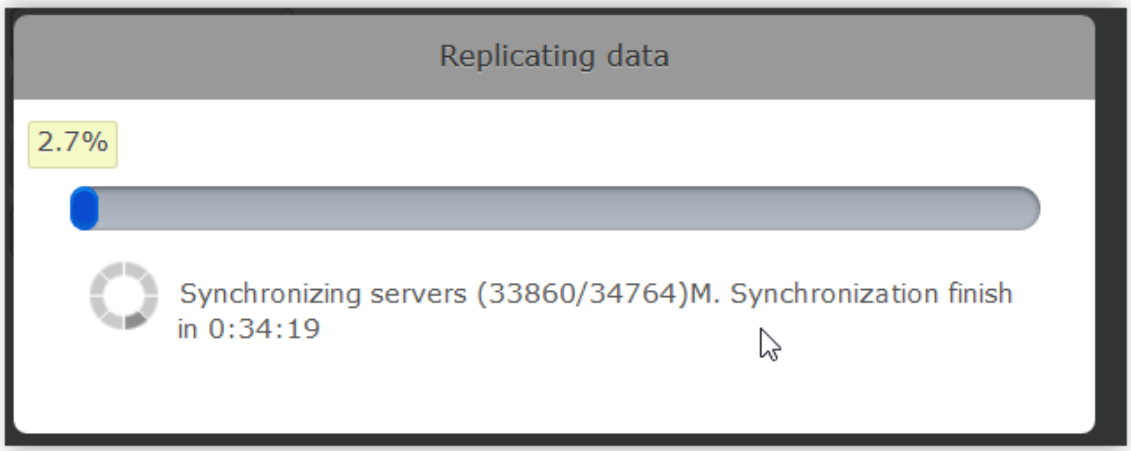


As a test you can execute the following command in the prompt:

```
cat /proc/drbd
```

With this command you can see which is the primary server and the status of the data replication.

```
[root@Elastix01 ~]# cat /proc/drbd
version: 8.3.15 (api:88/proto:86-97)
GIT-hash: 0ce4d235fc02b5c53c1c52c53433d11a694eab8c build by mockbuild@builder17.centos.org, 2013-03-27 16:04:08
 0: cs:SyncSource ro:Primary/Secondary ds:UpToDate/Inconsistent C r-----
    ns:1658112 nr:0 dw:0 dr:1665792 al:0 bm:100 lo:7 pe:20 ua:64 ap:0 ep:1 wo:b oos:33943364
        [>....................] sync'ed:  4.7% (33144/34764)M
        finish: 0:20:26 speed: 27,664 (21,224) K/sec
[root@Elastix01 ~]#
```

```
[root@Elastix02 ~]# cat /proc/drbd
version: 8.3.15 (api:88/proto:86-97)
GIT-hash: 0ce4d235fc02b5c53c1c52c53433d11a694eab8c build by mockbuild@builder17.centos.org, 2013-03-27 16:04:08
 0: cs:SyncTarget ro:Secondary/Primary ds:Inconsistent/UpToDate C r-----
    ns:0 nr:2465280 dw:2464640 dr:0 al:0 bm:150 lo:6 pe:7492 ua:5 ap:0 ep:1 wo:b oos:33134276
        [>....................] sync'ed:  7.0% (32356/34764)M
        finish: 0:18:33 speed: 29,740 (23,696) want: 102,400 K/sec
[root@Elastix02 ~]#
```

After data replication is completed, the HA module will stop the services in the secondary server, as the cluster is the responsible for starting the services in that server.

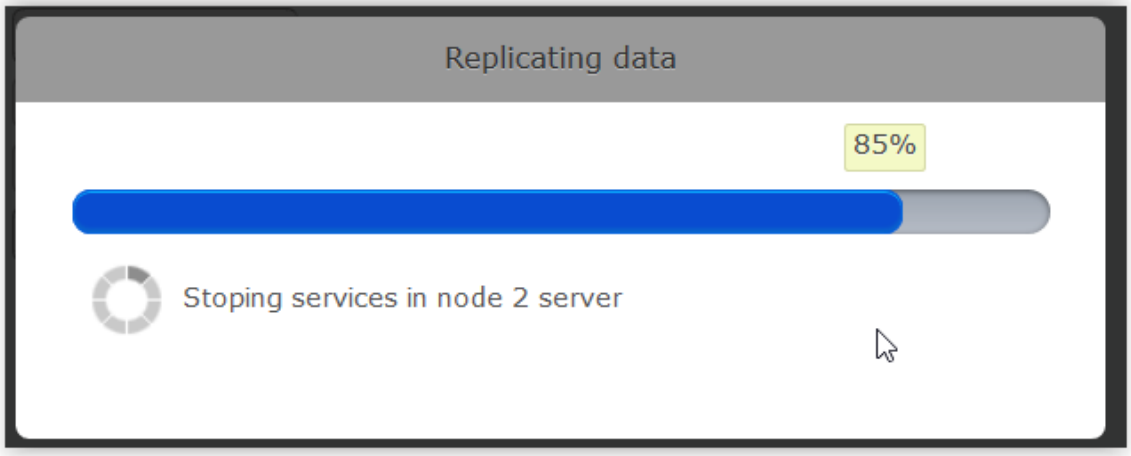


After the process is complete, you can go to the *"Replication Status"* menu and review the process and the master server.

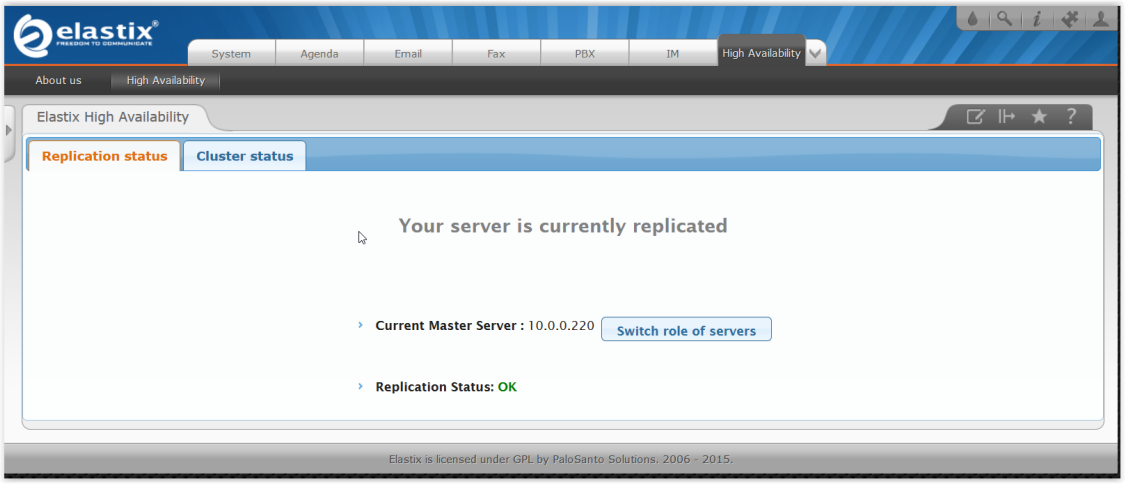


In the *"Cluster status"* menu you can see information about the cluster and useful information as the free space in the partition used for the cluster; and the route used to mount that partition.

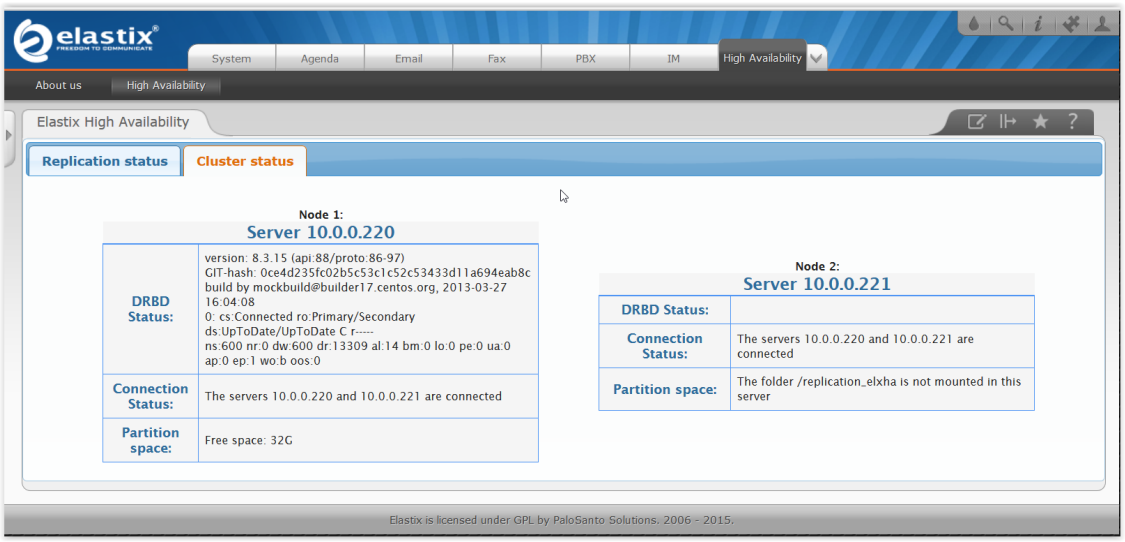


If you want to do an additional verification, execute the command:

- `cd /replication_elxha`
- `ls`

```
[root@Elastix01 ~]# cd /replication_elxha
[root@Elastix01 replication_elxha]# ls
etc  lost+found  usr  var
[root@Elastix01 replication_elxha]# cd /etc/asterisk
[root@Elastix01 asterisk]#
```

With this you can ping the floating IP to verify everything is working correctly.

```
Microsoft Windows [Version 6.3.9600]
(c) 2013 Microsoft Corporation. All rights reserved.

C:\WINDOWS\system32>ping 10.0.0.222

Pinging 10.0.0.222 with 32 bytes of data:
Reply from 10.0.0.222: bytes=32 time=5ms TTL=64
Reply from 10.0.0.222: bytes=32 time=1ms TTL=64
Reply from 10.0.0.222: bytes=32 time=1ms TTL=64
Reply from 10.0.0.222: bytes=32 time=1ms TTL=64

Ping statistics for 10.0.0.222:
    Packets: Sent = 4, Received = 4, Lost = 0 (0% loss),
Approximate round trip times in milli-seconds:
    Minimum = 1ms, Maximum = 5ms, Average = 2ms

C:\WINDOWS\system32>
```

The Elastix HA module uses replication mechanisms of the DRBD with heartbeat. The inner workings of it exceed the scope of this manual.

In case of any issue, both servers assume a secondary role, an alternative interface can be used to revert this by entering:

http://server_ip_address:3003/elxha/index.php  the user and password will be required.

User: admin
Password: the same you entered as root pass for secondary node.

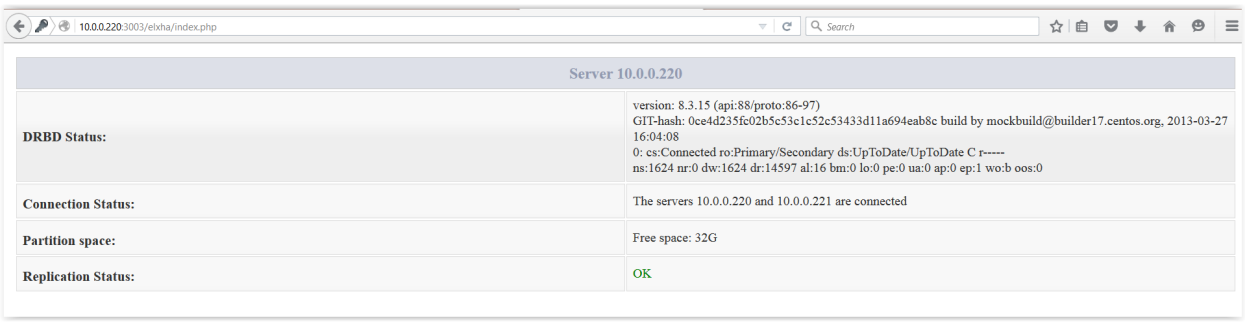


Now the module is successfully configured and your cluster has been deployed!

If you have any question regarding the installation you can send an email to: addons@elastix.com, we'll be glad to assist you during this process.