

004-LizardFS - Configuring Storage Node & Shadow Master

Notebook: <Inbox>

Created: 5/1/2018 9:19 PM

Updated: 5/8/2018 12:12 PM

Author: bdnog8@redskybd.com

URL: cherrytree0.38.4

PART # 01 : CONFIGURE SHADOW MASTER (BACKUP MASTER)

= Installation of LizardFS Packages

```
wget -O - http://packages.lizardfs.com/lizardfs.key | apt-key add -  
echo "deb [arch=amd64] http://packages.lizardfs.com/ubuntu/xenial xenial main" >  
/etc/apt/sources.list.d/lizardfs.list  
  
apt update  
apt install -y lizardfs-master lizardfs-chunkserver lizardfs-adm lizardfs-cgi lizardfs-cgiserv  
lizardfs-client lizardfs-common lizardfs-chunkserver
```

We have to modify Hosts file as follows.

Here make hostname according to your groupname following the bellow format

IP Address **groupX-nodeY**

// replace X with group-number & Y with node number as instructor say.

```
vim /etc/hosts
```

```
192.168.108.20 group1-node0 mfsmaster  
192.168.108.21 group1-nodel1  
192.168.108.22 group1-node2  
192.168.108.23 group1-node3
```

This is very important to follow the proper naming of the nodes.

For massive deployment you may use internal dns-server configured by unbound.

= Configuring Shadow Master Server (Shadow Master)

```
cd /etc/mfs  
cp mfsmaster.cfg.dist mfsmaster.cfg
```

```
vim mfsmaster.cfg
```

Now uncomment the following lines | OR | **Copy & Past** the following lines at the end of the file.

```
PERSONALITY = shadow  
WORKING_USER = mfs  
WORKING_GROUP = mfs
```

```
SYSLOG_IDENT = mfsmaster
NICE_LEVEL = -19
EXPORTS_FILENAME = /etc/mfs/mfsexports.cfg
TOPOLOGY_FILENAME = /etc/mfs/mfstopology.cfg
CUSTOM_GOALS_FILENAME = /etc/mfs/mfsgoals.cfg
DATA_PATH = /var/lib/mfs
AUTO_RECOVERY = 1
ENDANGERED_CHUNKS_PRIORITY = 1
MASTER_HOST = mfsmaster
MASTER_PORT = 9419
```

Here, change **false** to **true**

```
vim /etc/default/lizardfs-master
```

```
LIZARDFSMASTER_ENABLE=true
```

```
cp /var/lib/mfs/metadata.mfs.empty /var/lib/mfs/metadata.mfs
cp mfsexports.cfg.dist mfsexports.cfg
cp mfstopology.cfg.dist mfstopology.cfg
cp mfsgoals.cfg.dist mfsgoals.cfg
```

Change/uncomment parameter as follows in the file **mfstopology.cfg**

```
vim mfstopology.cfg
```

```
192.168.108.0/24          1
```

Modify the file **mfsgoals.cfg**

```
vim mfsgoals.cfg
```

```
1 1 : _
2 2 : _ _
3 3 : _ _ _
4 xor2 : $xor2
5 xor3 : $xor3
```

```
/etc/init.d/lizardfs-master restart
```

```
systemctl status lizardfs-master.service
```

Browse <http://192.168.108.20:9425/>

= Adding/Making the script to start the LizarFS service autostart at boot/reboot time.

```
vim /etc/network/if-up.d/rcinit.sh
```

```
#!/bin/bash  
/etc/init.d/lizardfs-master restart  
sleep 3  
/etc/init.d/lizardfs-cgiserv restart
```

```
chmod +x /etc/network/if-up.d/rcinit.sh  
echo 'post-up /etc/network/if-up.d/rcinit.sh' >> /etc/network/interfaces
```

Now reboot.

```
reboot
```

Now check from <http://192.168.108.20:9425>

PART # 02 : CONFIGURE CHUNK-SERVER (STORAGE NODE)

= Disk partition format for chunk server.

We have to allocate dedicated disks (HDD/SSD) or Disk Partition (For Practice), those will be the building-blocks of the LizardFS distributed file system.

Identify & format the HDD partition with **xfs** (please be careful to identify and format the partition).

Suppose here our disk partition is /dev/sda3; (Please choose your appropriate partition)

```
mkfs.xfs /dev/sda3  
mkdir /mnt/disk1
```

Now mount the partition in **/mnt/disk1**;

Here we will mount the disks/partition via **/etc/fstab** file, so that while reboot the node, it mount automatically.

```
vim /etc/fstab
```

```
/dev/sda3    /mnt/disk1    xfs    defaults,noatime,nodiratime 0    0
```

Run the mount command;

```
mount -a
```

See the mount status;

```
df -h
```

Now change the ownership of the mounted disks

```
chown -R mfs:mfs /mnt/disk1
```

= Chunk Server Configuration:

```
cd /etc/mfs
cp mfschunkserver.cfg.dist mfschunkserver.cfg
cp mfshdd.cfg.dist mfshdd.cfg
```

```
vim mfschunkserver.cfg
```

Now uncomment the following lines | OR | Copy & Past the following lines at the end of the file.

```
WORKING_USER = mfs
WORKING_GROUP = mfs
SYSLOG_IDENT = mfschunkserver
NICE_LEVEL = -19
DATA_PATH = /var/lib/mfs
MASTER_HOST = mfsmaster
MASTER_PORT = 9420
HDD_CONF_FILENAME = /etc/mfs/mfshdd.cfg
```

```
vim mfshdd.cfg
```

Add the following line;

```
/mnt/disk1
```

```
vim /etc/default/lizardfs-chunkserver
```

Change false to true

```
=true
```

Now restart the service

```
/etc/init.d/lizardfs-chunkserver restart
```

= Mounting the clustered storage.

```
cp mfsmount.cfg.dist mfsmount.cfg
```

```
vim mfsmount.cfg
```

Uncomment the following line

```
/mnt/mfs
```

```
mkdir /mnt/mfs  
chown -R mfs:mfs /mnt/mfs  
mfsmount
```

Now set the default replication factor to 2, so that anything kept in the clustered storage it has two copies of those to survive the failure of any clustered node/disks.

```
mfssetgoal 2 /mnt/mfs
```

```
df -h ; to see its mounted or not
```

Modify the file [/etc/network/if-up.d/rcinit.sh](#) and add the following lines at the end of the file

```
vim /etc/network/if-up.d/rcinit.sh
```

```
sleep 2  
/etc/init.d/lizardfs-chunkserver restart  
sleep 2  
mfsmount
```

Now reboot and check after reboot that everything is working properly or not.

```
reboot
```

PART # 01 : CONFIGURE SHADOW MASTER (BACKUP MASTER)

= Installation of LizardFS Packages

```
wget -O - http://packages.lizardfs.com/lizardfs.key | apt-key add -  
echo "deb [arch=amd64] http://packages.lizardfs.com/ubuntu/xenial xenial main" >  
/etc/apt/sources.list.d/lizardfs.list  
  
apt update  
apt install -y lizardfs-master lizardfs-chunkserver lizardfs-adm lizardfs-cgi lizardfs-cgiserv  
lizardfs-client lizardfs-common lizardfs-chunkserver
```

We have to modify Hosts file as follows.

Here make hostname according to your groupname following the bellow format

IP Address **groupX-nodeY**

// replace X with group-number & Y with node number as instructor say.

```
vim /etc/hosts
```

```
192.168.108.20 group1-node0 mfsmaster
192.168.108.21 group1-node1
192.168.108.22 group1-node2
192.168.108.23 group1-node3
```

This is very important to follow the proper naming of the nodes.
For massive deployment you may use internal dns-server configured by unbound.

= Configuring Shadow Master Server (Shadow Master)

```
cd /etc/mfs
cp mfsmaster.cfg.dist mfsmaster.cfg
```

```
vim mfsmaster.cfg
```

Now uncomment the following lines | OR | **Copy & Past** the following lines at the end of the file.

```
PERSONALITY = shadow
WORKING_USER = mfs
WORKING_GROUP = mfs
SYSLOG_IDENT = mfsmaster
NICE_LEVEL = -19
EXPORTS_FILENAME = /etc/mfs/mfsexports.cfg
TOPOLOGY_FILENAME = /etc/mfs/mfstopology.cfg
CUSTOM_GOALS_FILENAME = /etc/mfs/mfsgoals.cfg
DATA_PATH = /var/lib/mfs
AUTO_RECOVERY = 1
ENDANGERED_CHUNKS_PRIORITY = 1
MASTER_HOST = mfsmaster
MASTER_PORT = 9419
```

Here, change **false** to **true**

```
vim /etc/default/lizardfs-master
```

```
LIZARDFSMASTER_ENABLE=true
```

```
cp /var/lib/mfs/metadata.mfs.empty /var/lib/mfs/metadata.mfs
cp mfsexports.cfg.dist mfsexports.cfg
cp mfstopology.cfg.dist mfstopology.cfg
cp mfsgoals.cfg.dist mfsgoals.cfg
```

Change/uncomment parameter as follows in the file **mfstopology.cfg**

```
vim mfstopology.cfg
```

Modify the file `mfsgoals.cfg`

```
vim mfsgoals.cfg
```

```
1 1 : _  
2 2 : _ _  
3 3 : _ _ _  
4 xor2 : $xor2  
5 xor3 : $xor3
```

```
/etc/init.d/lizardfs-master restart
```

```
systemctl status lizardfs-master.service
```

Browse <http://192.168.108.20:9425/>

= Adding/Making the script to start the LizarFS service autostart at boot/reboot time.

```
vim /etc/network/if-up.d/rcinit.sh
```

```
#!/bin/bash  
/etc/init.d/lizardfs-master restart  
sleep 3  
/etc/init.d/lizardfs-cgiserv restart
```

```
chmod +x /etc/network/if-up.d/rcinit.sh  
echo 'post-up /etc/network/if-up.d/rcinit.sh' >> /etc/network/interfaces
```

Now reboot.

```
reboot
```

Now check from <http://192.168.108.20:9425>

PART # 02 : CONFIGURE CHUNK-SERVER (STORAGE NODE)

= Disk partition format for chunk server.

We have to allocate dedicated disks (HDD/SSD) or Disk Partition (For Practice), those will be the building-blocks of the LizardFS distributed file system.

Identify & format the HDD partition with **xfs** (please be careful to identify and format the partition).

Suppose here our disk partition is `/dev/sda3`; (Please choose your appropriate partition)

```
mkfs.xfs /dev/sda3
mkdir /mnt/disk1
```

Now mount the partition in `/mnt/disk1`;

Here we will mount the disks/partition via `/etc/fstab` file, so that while reboot the node, it mount automatically.

```
vim /etc/fstab
```

```
/dev/sda3    /mnt/disk1    xfs    defaults,noatime,nodiratime 0    0
```

Run the mount command;

```
mount -a
```

See the mount status;

```
df -h
```

Now change the ownership of the mounted disks

```
chown -R mfs:mfs /mnt/disk1
```

= Chunk Server Configuration:

```
cd /etc/mfs
cp mfschunkserver.cfg.dist mfschunkserver.cfg
cp mfshdd.cfg.dist mfshdd.cfg
```

```
vim mfschunkserver.cfg
```

Now uncomment the following lines | OR | Copy & Past the following lines at the end of the file.

```
WORKING_USER = mfs
WORKING_GROUP = mfs
SYSLOG_IDENT = mfschunkserver
NICE_LEVEL = -19
DATA_PATH = /var/lib/mfs
MASTER_HOST = mfsmaster
MASTER_PORT = 9420
HDD_CONF_FILENAME = /etc/mfs/mfshdd.cfg
```



```
vim mfsbdd.cfg
```

Add the following line;

```
/mnt/disk1
```

```
vim /etc/default/lizardfs-chunkserver
```

Change false to true

```
=true
```

Now restart the service

```
/etc/init.d/lizardfs-chunkserver restart
```

= Mounting the clustered storage.

```
cp mfsmount.cfg.dist mfsmount.cfg
```

```
vim mfsmount.cfg
```

Uncomment the following line

```
/mnt/mfs
```

```
mkdir /mnt/mfs  
chown -R mfs:mfs /mnt/mfs  
mfsmount
```

Now set the default replication factor to 2, so that anything kept in the clustered storage it has two copies of those to survive the failure of any clustered node/disks.

```
mfsssetgoal 2 /mnt/mfs
```

```
df -h ; to see its mounted or not
```

Modify the file /etc/network/if-up.d/rcinit.sh and add the following lines at the end of the file

```
vim /etc/network/if-up.d/rcinit.sh
```

```
sleep 2  
/etc/init.d/lizardfs-chunkserver restart  
sleep 2  
mfsmount
```

Now reboot and check after reboot that everything is working properly or not.

```
reboot
```

LizardFS -Add Cluster Storage in Virt-Manager

```
mkdir /mnt/mfs/iso
mkdir /mnt/mfs/kvm

cd /mnt/mfs/iso/
wget -c http://192.168.108.8/iso/ubuntu-16.04.3-server-amd64.iso
```

- Add mfs directory in virt-manager.
- Install ubuntu-server 16.04.3.
- Please Use the VM name as follows;

naming-format:- **groupX-nodeY-vm1**

Example for **group1-node1-vm1**