



# MooseFS 2.0 Upgrade Guide

CORE TECHNOLOGY Development & Support Team

April 6, 2016

© 2014-2016

Piotr Robert Konopelko, CORE TECHNOLOGY Development & Support Team.  
All rights reserved.

v. 1.0.8

*Proofread by Agata Kruszona-Zawadzka.*

*Coordination & layout by Piotr Robert Konopelko.*

Please send corrections to [peter@mfs.io](mailto:peter@mfs.io).

# Contents

<b>1</b>	<b>Introduction</b>	<b>3</b>
1.1	Backing-up data and configuration files . . . . .	3
1.2	Repository . . . . .	4
1.2.1	Adding repository . . . . .	4
1.2.2	Repository branches . . . . .	4
<b>2</b>	<b>MooseFS 2.0 (CE and Pro) upgrade to newer version</b>	<b>5</b>
<b>3</b>	<b>MooseFS 1.6.27-5 to 2.0.x CE upgrade</b>	<b>6</b>
3.1	Making backup of data and configuration files . . . . .	6
3.2	Upgrading MooseFS master server: stopping, removing, installing, configuring and starting up . . . . .	6
3.2.1	Stopping . . . . .	6
3.2.2	Removing . . . . .	6
3.2.3	Installing new version of MooseFS master server . . . . .	7
3.2.4	Copying and comparing configuration files . . . . .	7
3.2.5	Starting up master server . . . . .	7
3.3	Upgrading the rest of MooseFS modules . . . . .	8
<b>4</b>	<b>MooseFS 2.0 CE to Pro HA migration</b>	<b>9</b>
4.1	MFS Master migration . . . . .	9
4.2	Upgrading chunkservers . . . . .	10
4.3	Upgrading clients (mounts) . . . . .	10
4.4	Adding master followers . . . . .	10
<b>5</b>	<b>MooseFS 2.0 Pro HA to CE migration</b>	<b>11</b>
5.1	Checking LEADER master . . . . .	11
5.2	Changing DNS entry . . . . .	11
5.3	Stopping FOLLOWERS . . . . .	11
5.4	Migrating master . . . . .	12
5.5	Migrating metaloggers . . . . .	12
5.6	Migrating chunkservers . . . . .	12
5.7	Migrating clients (mounts) . . . . .	13
<b>6</b>	<b>Adding master follower(s) server(s) procedure (Pro only)</b>	<b>14</b>
6.1	Installing 'extra masters' . . . . .	14
6.2	Adding DNS entries . . . . .	14

# Chapter 1

## Introduction

This book guides you through the process of upgrading and/or migrating MooseFS. It describes upgrading MooseFS to newer version in 2.0.x branch, upgrading from 1.6.27-5 to 2.0.x, migrating from Community Edition to Pro and migrating from Pro to Community Edition.

It is strongly advised to read the entire document before attempting to start the upgrade/migration process.

There are several prerequisites, discussed in the following paragraphs, that should be taken into consideration before upgrading/migrating MooseFS.

### 1.1 Backing-up data and configuration files

The most important thing is to **back-up all configuration, metadata and changelog files that contain critical MooseFS configuration. Losing those files may lead to data loss and a serious problem with recovering your files.**

If MooseFS was installed from officially supported packages or repository (version 2.0 and higher), the configuration files should be located in default configuration path i.e. `/etc/mfs` and metadata files should be located in `/var/lib/mfs`

If you installed MooseFS 1.6.27-5 from sources without any additional parameters to `configure` script, the configuration files and metadata should be located in `/usr/local/var/lib/mfs` or `/usr/local/etc/mfs`. Depending on parameters selected during compilation process, these paths may vary (e.g. it might be `/etc`, `/etc/mfs`, `/var/lib/mfs`). If you compiled MooseFS from sources and set different paths, you should refer to your site specific documentation to find out where MooseFS configuration files have been placed.

## 1.2 Repository

### 1.2.1 Adding repository

To install MooseFS 2.0 Pro or CE you need to add MooseFS Official Supported Repositories to your system. This process is described at <http://get.moosefs.com> (please select your distribution in menu on the left) or in paragraph 2.2 in document named *Installing MooseFS 2.0 Step by Step Tutorial*.

At this time there are repositories available for Ubuntu / Debian, RHEL / CentOS / Fedora, FreeBSD and MacOS X.

### 1.2.2 Repository branches

Our repository contains two branches: **stable** and **current**. Version from **stable** branch has been tested both in the production and in our test environment. Version from **current** branch – only in our test environment. MooseFS versions in these branches are upgraded automatically after finishing the tests.

At the time of writing this guide, **stable** branch contains version 2.0.39-1, and **current** branch contains version 2.0.40-1.

**Stable** branch is default and you don't need to make any changes in default URL:

```
http://ppa.moosefs.com/stable/.
```

If you want to use **current** branch, you just need to replace **stable** with **current** after `http://ppa.moosefs.com/` and before `apt`, `yum`, `freebsd` or `osx`, so URL will look like:

```
http://ppa.moosefs.com/current/[rest of url]
```

It is also possible to use version number instead of "branch" if you want to upgrade to a specific version of MooseFS (e.g. 2.0.40-1):

```
http://ppa.moosefs.com/2.0.40/[rest of url]
```

If you want to use this option, please remember you need to manually change version number on each server to the selected one before doing an upgrade.

## Chapter 2

# MooseFS 2.0 (CE and Pro) upgrade to newer version

This chapter describes general MooseFS Upgrade process (e.g. from version 2.0.34 to 2.0.35). Upgrading to newer version is very simple, but **there is one important thing**: the order of upgrade. We only support the following upgrade order:

1. Upgrading all MooseFS master follower(s) (one-by-one) (Pro only)
2. Upgrading MooseFS leader master and MooseFS tools (cgi & cgiserv, cli, supervisor (Pro only))
3. Upgrading all MooseFS metaloggers (one-by-one)
4. Upgrading all MooseFS chunkservers (one-by-one)
5. Upgrading all MooseFS clients (one-by-one)

On every server upgrading comes down to running two commands:

```
# apt-get update  
# apt-get upgrade
```

The second command listed above will upgrade all packages installed on your system. If you want to upgrade only one package (e.g. `moosefs-pro-master`), please run e.g.:

```
# apt-get install moosefs-pro-master
```

– this will upgrade only the specified package to the newest version.

## Chapter 3

# MooseFS 1.6.27-5 to 2.0.x CE upgrade

**Notice:** It is not possible to roll-back to older (1.6.27-5) version after this upgrade!

### 3.1 Making backup of data and configuration files

Now you **need to** back-up configuration files in a safe place, as described in **Paragraph 1.1** above. The backup should be done on **all servers**, i.e. configuration and metadata from **each** server should be backed-up.

### 3.2 Upgrading MooseFS master server: stopping, removing, installing, configuring and starting up

#### 3.2.1 Stopping

In order to stop `mfsmaster`, run the following command:

```
# mfsmaster stop
```

#### 3.2.2 Removing

- If you have installed MooseFS 1.6.27-5 from packages:  
# `apt-get purge --purge mfs-master`
- or:  
# `dpkg -P mfs-master`

- Otherwise, if MooseFS has been compiled and installed from sources, you have to remove all MooseFS binaries (`mfsmaster`, `mfsctgi`, ...) (most likely) from `/usr/sbin`, `/usr/local/bin` or `/usr/local/sbin`.

### 3.2.3 Installing new version of MooseFS master server

The next step to do is installing MooseFS 2.0 CE master server. For example on Debian/Ubuntu you can do it by running the following command:

```
# apt-get install moosefs-master
```

### 3.2.4 Copying and comparing configuration files

It is assumed, that old configuration is in `/etc` (it may exist in `/etc`, `/usr/local/etc`, or `/usr/local/etc/mfs` depending on the way MooseFS has been installed, so make sure where your configuration files are and make changes in following commands). Metadata should exist in `/var/lib/mfs`. If metadata files are in different location, please move them to `/var/lib/mfs`.

- Next step is to move old configuration files:

```
# mv /etc/mfsmaster.cfg /etc/mfs/mfsmaster.cfg.old
# mv /etc/mfsexports.cfg /etc/mfs/mfsexports.cfg.old
# mv /etc/mfstopology.cfg /etc/mfs/mfstopology.cfg.old
```

And copy examples as new configuration files:

```
# cp -i /etc/mfs/mfsmaster.cfg.dist /etc/mfs/mfsmaster.cfg
# cp -i /etc/mfs/mfsexports.cfg.dist /etc/mfs/mfsexports.cfg
# cp -i /etc/mfs/mfstopology.cfg.dist /etc/mfs/mfstopology.cfg
```

- The last step before starting up the new master server is to **compare new configuration files with old ones** and **make changes in new configuration**. In this step you need to rewrite modified parameters from old file to new one:

```
# vim -d /etc/mfs/mfsmaster.cfg /etc/mfs/mfsmaster.cfg.old
```

**Perform this step for every configuration files.**

### 3.2.5 Starting up master server

The last step is to start up the `mfsmaster` process.



### 3.3 Upgrading the rest of MooseFS modules

Now you should upgrade the rest of MooseFS modules. It should be done analogously to MooseFS master server upgrade, in following order:

- Make sure you have backed-up configuration and metadata files
- Stop MooseFS module process
- Remove MooseFS module binaries
- Install version 2.0.x of MooseFS module
- Copy old, **compare with new** and **make proper changes to new configuration files** and make sure metadata and changelogs files are in `/var/lib/mfs` (on metaloggers)
- Start module's process

The proper modules upgrade order is:

- MooseFS metalogger(s) (one-by-one)
- MooseFS chunkservers (one-by-one)
- MooseFS clients (mounts) (one-by-one)

## Chapter 4

# MooseFS 2.0 CE to Pro HA migration

### Notes:

- MooseFS PRO **depends on DNS**, so make sure that your local DNS server is running and resolving name for `mfsmaster` (or other chosen host name) properly.
- In Pro version it is not possible to use `/etc/hosts` file to resolve `mfsmaster` (or other chosen name). After making sure your DNS is working properly, you must remove entries pointing to master host from `/etc/hosts`.
- Before you start the migration it is good to know that during the packages uninstall process files like: `/etc/default/moosefs-master` and `/etc/init.d/moosefs-master` will be removed. Configuration and metadata files will not be deleted during migration, but good idea is to backup them first.

### 4.1 MFS Master migration

- Stop MooseFS CE master: `mfsmaster -c /etc/mfs/mfsmaster.cfg stop`
- Uninstall `moosefs-master`, `moosefs-cli`, `moosefs-cgiserv`, `moosefs-cgi`, `moosefs-supervisor` packages
- Install `moosefs-pro-master`, `moosefs-pro-supervisor`, `moosefs-pro-cgi`, `moosefs-pro-cgiserv`, `moosefs-pro-cli` packages
- Restart MooseFS cgi server `mfs-cgiserv restart`
- Place `mfslicence.bin` file in `/etc/mfs` folder
- Start `mfsmaster` pro version `mfsmaster -c /etc/mfs/mfsmaster.cfg start`

## 4.2 Upgrading chunkservers

It is important to perform chunkservers migration before adding master FOLLOWER in HA configuration. Only pro chunkservers are capable of switching to new LEADER master during the failure.

Chunkservers can be updated one by one on a running system **only** if the minimum goal of all files in the system is at least **2**. Otherwise you need to stop the master before performing any and all chunkservers updates.

If in `/etc/default/moosefs-chunkserver` file value `MFSCHUNKSERVER_ENABLE` is set to `true` then during uninstall proces `mfschunkserver` will automatically stop and then start. If not, you need to restart it manually.

- Stop chunkserver `mfschunkserver stop`
- Uninstall `moosefs-chunkserver` from system
- Install `moosefs-pro-chunkserver`
- Start chunkserver `mfschunkserver start`

## 4.3 Upgrading clients (mounts)

- Umount mfs mountpoints: `umount /mnt/mfs`
- Remove `moosefs-client` from system
- Install `moosefs-pro-client`
- Mount MooseFS: `mfsmount /mnt/mfs`

Alternatively you can add this entry to `/etc/fstab` to mount MooseFS automatically when system starts:

- For Linux: `mfsmount /mnt/mfs fuse defaults 0 0`
- For FreeBSD: `mfsmount_magic /mnt/mfs moosefs rw,mfsmaster=mfsmaster, mountprog=/usr/local/bin/mfsmount,late 0 0`

## 4.4 Adding master followers

This process is described below in **Chapter 6: Adding master follower(s) servers(s) procedure** To add master followers, jump to Chapter 6 now and follow instructions placed there.

## Chapter 5

# MooseFS 2.0 Pro HA to CE migration

MooseFS 2.0 is fully switchable between CE and PRO version. So all your data are safe and there is no reason to be afraid to test PRO version and then switch to CE version. Remember that the most important thing in migration is to preserve the order in these steps:

### 5.1 Checking LEADER master

If your LEADER master is not running on selected (i.e. the one, that will become the CE master) hardware machine, switch it using `mfssupervisor` command. For example your LEADER master is running on IP 192.168.0.2 machine and your CE master needs to be on IP 192.168.0.1 machine.

```
mfssupervisor -l 192.168.0.1
```

Now wait a few seconds for chunkservers and clients to switch connection to new master LEADER.

### 5.2 Changing DNS entry

When your master LEADER is on selected machine, it is time to remove all 'extra masters' – FOLLOWERS entries from your DNS configuration. Leave only one mfsmaster IP. Now restart DNS service. Remember that propagation of new DNS entries can take some time and it depends on your TTL DNS configuration.

### 5.3 Stopping FOLLOWERS

Next step in migration is to stop all 'extra masters' – FOLLOWERS in your cluster. On all master FOLLOWERS machines stop mfsmaster process

```
mfsmaster -c /etc/mfs/mfsmaster.cfg stop
```

## 5.4 Migrating master

This step will stop master for installation time. Whole MFS cluster will be down!

- Stop mfsmaster on LEADER machine  
`mfsmaster -c /etc/mfs/mfsmaster.cfg`
- Remove `moosefs-pro-*` packages from system
- Install `moosefs-:` master, cgi, cgiserv, cli package
- Restart cgiserv (`mfs cgiserver restart`)
- Start mfsmaster CE version `mfsmaster -c /etc/mfs/mfsmaster.cfg start`

Now your master is running in CE version and CGI and cli are not displaying LEADER – FOLLOWER informations in tables.

## 5.5 Migrating metaloggers

Installation of metaloggers in HA cluster is not necessary, but in CE configuration is strongly recommended option. Without metaloggers, your metadata files are saved on one physical machine only. Running metalogger on another machine is the best solution for metadata backup. If metalogger was running in your PRO version perform these steps to migrate:

- Stop metalogger `mfsmetallogger -c /etc/mfs/mfsmetallogger.cfg stop`
- Remove `moosefs-pro-metalogger` from system
- Install `moosefs-metalogger`
- Start metalogger `mfsmetallogger -c /etc/mfs/mfsmetallogger.cfg start`

If metalogger was not running in your PRO version, install one now – refer to chapter **3.4** of **”Installing MooseFS 2.0 Step by Step Tutorial”**

## 5.6 Migrating chunkservers

Now it’s time to perform chunkserver migration. It can be done on running system – if your minimal goal is set to 2. If not, we suggest stopping master server before starting chunkservers migration. If you have set goal 2, on whole MFS then update can be done with One-By-One method and without stopping the cluster.

- Stop chunkserver `mfschunkserver -c /etc/mfs/mfschunkserver.cfg stop`
- Remove `moosefs-pro-chunkserver` from system
- Install `moosefs-chunkserver`

- Start chunkserver `mfschunkserver -c /etc/mfs/mfschunkserver.cfg start`

Repeat these steps for every chunk server.

## 5.7 Migrating clients (mounts)

This step is similar to chunkserver migration. First you need to unmount mfs and then uninstall `mfsclient` from system.

- Umount all mount points for MooseFS
- Uninstall `moosefs-pro-client`
- Install `moosefs-client`
- Mount MooseFS: `mfsmount /mnt/mfs`

Alternatively you can add this entry to `/etc/fstab` to mount MooseFS automatically when system starts:

- For Linux: `mfsmount /mnt/mfs fuse defaults 0 0`
- For FreeBSD: `mfsmount_magic /mnt/mfs moosefs rw,mfsmaster=mfsmaster, mountprog=/usr/local/bin/mfsmount,late 0 0`

## Chapter 6

# Adding master follower(s) server(s) procedure (Pro only)

### 6.1 Installing 'extra masters'

Please remember that **all masters servers need to have the same configuration and licence files** and identical or very similar hardware configuration. For example if your current MASTER have 64GB RAM and 256GB HDD, other masters need to have at least the same (or greater) amount of RAM and HDD space.

Let's start by adding 'extra master' in your cluster:

- Install `moosefs-pro-master`, `moosefs-pro-supervisor`, `moosefs-pro-cli`
- Copy configuration from `mfsmaster`  

```
scp root@mfsmaster: "/etc/mfs/mfsmaster.cfg /etc/mfs/mfsexports.cfg /etc/mfs/mfstopology.cfg /etc/mfs/mfslicence.bin" /etc/mfs
```
- First start of Follower Master Server should be performed with `-e` option to let it download metadata from Leader Master Server. You can achieve this by issuing the following command:  

```
mfsmaster -e -c /etc/mfs/mfsmaster.cfg start
```

Your new master is FOLLOWER and is downloading metadata from LEADER master but it is not visible in CGI and CLI tables. Now it's time to add DNS entry for the new master server.

### 6.2 Adding DNS entries

To start working with HA and 'extra masters' – FOLLOWERS in your cluster you need to point all masters IP to one name in DNS.

For example in your configuration there will be three masters and your entries should look similar to these ones:

```
mfsmaster IN A 192.168.0.1;
```

```
mfsmaster IN A 192.168.0.2;  
mfsmaster IN A 192.168.0.3;
```

```
mfsmaster1 IN A 192.168.0.1;  
mfsmaster2 IN A 192.168.0.2;  
mfsmaster3 IN A 192.168.0.3;
```

Now is time to restart DNS service and wait for it to propagate changes in your network. This step can take some time depending on TTL settings. The whole system will be fully-HA after DNS propagation.

You can test if DNS propagation is done in your LAN. Just run `host mfsmaster` command on several machines in LAN. If you see three addresses near hostname it means that DNS cache is updated and now MFS is running in HA configuration.

Now check CGI or use CLI (`mfscli -SIM`) to see if your LEADER and FOLLOWERS are visible.