

Recovering Filesystems

Goal: Recover a LANforge system that boots into Emergency Mode

If your LANforge boots into Emergency Mode, your system is experiencing file-system corruption. Follow these instructions to check the filesystems. This process will require a monitor and keyboard or a serial cable connected to the LANforge.



i File system corruptions are caused by power-off events without properly shutting down the system

1. The emergency mode prompt looks similar to this:

Give root password for maintenance
(or type Control-D to continue):

```
[ OK ] Reached target NFS client services.
[ OK ] Reached target Remote File Systems (Pre).
[ OK ] Reached target Remote File Systems.
Starting Update: UTMP about System Runlevel Changes...
[ OK ] Started Update: UTMP about System Runlevel Changes.
[ OK ] Started Tell Plymouth To Write Out Runtime Data.
Give root password for maintenance
(or press Control-D to continue): ath10k_pci 0000:04:00.0: DANGER! You're overriding EEPROM-defined regulatory domain
ath10k_pci 0000:04:00.0: from: 0x0 to 0x9c (svc-ready-work)
ath10k_pci 0000:04:00.0: Your card was not certified to operate in the domain you chose.
ath10k_pci 0000:04:00.0: This might result in a violation of your local regulatory rules.
ath10k_pci 0000:04:00.0: Do not ever do this unless you really know what you are doing!
ath10k_pci 0000:05:00.0: DANGER! You're overriding EEPROM-defined regulatory domain
ath10k_pci 0000:05:00.0: from: 0x0 to 0x9c (svc-ready-work)
ath10k_pci 0000:05:00.0: Your card was not certified to operate in the domain you chose.
ath10k_pci 0000:05:00.0: This might result in a violation of your local regulatory rules.
ath10k_pci 0000:05:00.0: Do not ever do this unless you really know what you are doing!
```

A. The messages beforehand that begin with [**OK**] can be ignored.

B. Messages beginning with `ath10k_pci` can be ignored.

i We assume you'll press the **Enter** (**↵**) key when executing the commands in further examples.

3. Provide the root password: `lanforge`

4. The we next need to determine where our filesystems are mounted. You might have filesystems on partitions, or on LVM volumes. Use the command `mount | grep ext4` to find your partitions:

```
# mount | grep ext4
/dev/sda2 on / type ext4 (rw,relatime,nodev,nosuid)
/dev/sda4 on /home type ext4 (rw,relatime,nodev,nosuid)
/dev/sda1 on /boot type ext4 (rw,relatime,nodev,nosuid)
```

This example shows partitions.

i Note that inspecting `/etc/fstab` often does not show you device partitions, rather it shows you volume or partition UUIDs.

5. Example of filesystems on volumes:

```
# mount | grep ext4
```

```
/dev/mapper/fedora-root on / type ext4 (rw,relatime,nodev,noexec,nosuid)
/dev/mapper/fedora-home on /home type ext4 (rw,relatime,nodev,noexec,nosuid)
/dev/sda1 on /boot type ext4 (rw,relatime,nodev,noexec,nosuid)
```

Notice that `/boot` is typically not a LVM volume.

6. Issuing the fsck commands

7. For each of those **filesystems** in the partition example above, you will execute the below commands:

- A.

```
# fsck -fy /dev/sda2
```

 (for /)
- B.

```
# fsck -fy /dev/sda4
```

 (for /home)
- C.

```
# fsck -fy /dev/sda1
```

 (for /boot)

8. For the **volumes** example above, you will execute the below commands:

- A.

```
# fsck -fy /dev/mapper/fedora-root
```

 (for /)
- B.

```
# fsck -fy /dev/mapper/fedora-home
```

 (for /home)
- C.

```
# fsck -fy /dev/sda1
```

 (for /boot)

9. Configuring Filesystems Check on Reboot

10. The commands below involve setting the 'mount count' parameter to 1. Every time the filesystem is mounted, it will perform a check. The following examples ensure that full journaling is enabled on the filesystem. This ensures the most amount of recovery. And after that, modern versions of Fedora have metadata checksumming available. Older LANforge systems do not have this.

A. Filesystem Check on Every Mount

The filesystems do not need to be unmounted to set this parameter. Only the partition example is shown, the volumes example uses similar commands.

- A.

```
# touch /forcefsck
```

 (this technique works with any model of filesystem)
- B.

```
# tune2fs -c 1 /dev/sda2
```

 (for /)
- C.

```
# tune2fs -c 1 /dev/sda4
```

 (for /home)
- D.

```
# tune2fs -c 1 /dev/sda1
```

 (for /boot)

B. Tune the Filesystem for Full Journaling

The filesystems do not need to be unmounted to set this parameter. The commands for the volumes version is similar to the partition example below.

- A.

```
# tune2fs -o journal_data /dev/sda2
```

 (for /)
- B.

```
# tune2fs -o journal_data /dev/sda4
```

 (for /home)
- C.

```
# tune2fs -o journal_data /dev/sda1
```

 (for /boot)

- C. **Note:** Inspect `/etc/fstab` for conflicting mount options. The following example shows mount options applied in the `fstab` file:

```
$ mount | grep ext4
```

```
/dev/sda2 on / type ext4 (rw,relatime,nodelalloc)
/dev/sda4 on /home type ext4 (rw,relatime,nodelalloc)
/dev/sda1 on /boot type ext4 (rw,relatime,data=writeback)
```

```
$ grep data= /etc/fstab
```

```
UUID=1c1b4732-653f-47dd-a106-ae17cf5b12a9 /boot ext4 data=writeback 1 2
```

Notice the `fstab` entry for `/boot`? It has overridden the data journaling mode. Erase that setting from the `fstab` mount options.

D. Enable Metadata Checksumming (optional)

It is only practical to apply metadata checksumming in the below conditions. You will not be able to apply it to the root filesystem `/` because it cannot be unmounted when you boot the system.

- A. Fedora 27 or more recent
- B. `tune2fs` 1.43 or more recent
- C. `crc32c` or `libcrc32c` modules loaded
- D. you can unmount the filesystem.

- E. To Boot into Emergency Mode, follow these steps:

- A. Reboot system into Emergency Mode by appending the word `emergency` at the end of your grub boot option.

```
# umount -f /home
```

- B.

```
# umount -f /boot
```

- C.

```
# fsck -fD /dev/sda1
```

- D. (remember that's `/boot`)

```
# fsck -fD /dev/mapper/fedora-home
```

- E.

```
# tune2fs -O metadata_csum /dev/sda1
```

- F.

```
# tune2fs -O metadata_csum /dev/mapper/fedora-home
```

- G.

F. Disable Drive Write Caching (optional)

Drives typically ship with write caching enabled for a performance boost. SSDs have this as well. By turning it off you trade some performance for increased data safety. LANforge systems do not ship with write caching disabled. To make this setting enabled every boot, it needs to be added to `/etc/rc.local`. LANforge systems have `rc-local.service` enabled by default.

i Only disable write caching if your machine powers off without shutdown frequently

```
# vi /etc/rc.local
```

- A.

- B. Add the line `/sbin/hdparm -W0 /dev/sda`

- C. Save the file.

- D. Reboot, or issue the command for this booted session:

```
# /sbin/hdparm -W0 /dev/sda
```