

CentOS Not Booting

Apr 2021

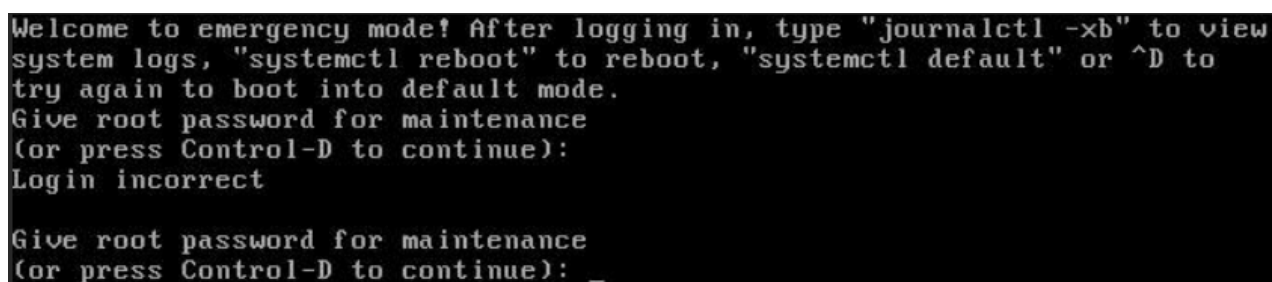
aka Centos Emergency Mode.

Introduction

This has been listed under IPHE, not because it is specific to IPHE (because it is not) but because it was on an IPHE system that I first observed this phenomenon.

After a server was rebooted (and I cannot say whether this was a reboot from a command line, a server crash, someone pulling out the power cables, a power cut etc) the end result is that the server will not fully boot up, instead we end up in '**emergency mode**'.

When logging in to the server from command KVM or from BMC (or whatever remote management is present) we can see the following tale screen:



```
Welcome to emergency mode! After logging in, type "journalctl -xb" to view
system logs, "systemctl reboot" to reboot, "systemctl default" or ^D to
try again to boot into default mode.
Give root password for maintenance
(or press Control-D to continue):
Login incorrect

Give root password for maintenance
(or press Control-D to continue): _
```

The Cause

To move past this page you have to enter **Control-D** to get to the login prompt and complete the boot up process. This is a hassle to do every time there is a reboot. A clue of what to look at is in the text, the command **journalctl -xb** can be executed once you have logged in.

On my system this file was 30,888 lines long, so hundreds of pages long. After searching through it though I did find this section:

```
-- The start-up result is done.
Apr 20 17:04:10 k8master1 kernel: power_meter ACPI000D:00: Found ACPI power meter.
Apr 20 17:04:10 k8master1 systemd-fsck[19721]: /dev/mapper/vg_main-lv_var: Inodes that were part of a
corrupted orphan linked list found.
Apr 20 17:04:10 k8master1 systemd-fsck[19721]: /dev/mapper/vg_main-lv_var: UNEXPECTED INCONSISTENCY; RUN
fsck MANUALLY.
Apr 20 17:04:10 k8master1 systemd-fsck[19721]: (i.e., without -a or -p options)
Apr 20 17:04:10 k8master1 systemd-fsck[19721]: fsck failed with error code 4.
Apr 20 17:04:10 k8master1 systemd-fsck[19721]: Running request emergency.target/start/replace
```

```
Apr 20 17:04:10 k8master1 systemd[1]: Started File System Check on /dev/mapper/vg_main-lv_var.
-- Subject: Unit systemd-fsck@dev-mapper-vg_main\x2dlv_var.service has finished start-up
```

The important parts are:

```
Apr 20 17:04:10 k8master1 systemd-fsck[19721]: /dev/mapper/vg_main-lv_var: UNEXPECTED INCONSISTENCY; RUN fsck MANUALLY.
```

The partition **vg_main-lv_var** has an issue, and Linux wants us to run **fsck**. **fsck** is a file check tool, however it will not run on mounted file systems, and the only way to unmount them is if you boot from say a USB stick.

The other important part is:

```
Apr 20 17:04:10 k8master1 systemd-fsck[19721]: /dev/mapper/vg_main-lv_var: Inodes that were part of a corrupted orphan linked list found.
```

So **/dev/mapper/vg_main-lv_var** has an **Inode** issue. An **Inode** is a reference that points to each file on the disk, there is an **Inode** for each file, and for empty space also. The **Inodes** should all line up nicely, but in the case of a server crash or sudden power loss, the process writing to the disk may not have time to update all the **Inodes**, and so a discrepancy occurs on the disk that the system notices.

Each time the server boots, this discrepancy causes the server to go in to **Emergency Mode**. On a disk that is being written to a lot, this is not an uncommon occurrence, there are two choices, run the **fsck**, which may not be possible (especially remotely) or tell the system to not perform the check at boot, thus avoiding going in to **Emergency Mode** while rebooting.

The Remedy

We need to tell the system to not perform the **fsck** at boot, and to do this we need to edit the file **fstab** in the **etc** directory (this is for CentOS).

Open the file in a text editor by using **vi /etc/fstab**

```
#
# /etc/fstab
# Created by anaconda on Thu Nov  5 18:45:24 2020
#
# Accessible filesystems, by reference, are maintained under '/dev/disk'
# See man pages fstab(5), findfs(8), mount(8) and/or blkid(8) for more info
#
/dev/mapper/vg_main-lv_root / ext4 defaults 1 1
UUID=ce253ff9-86ee-480b-87b3-7efa803bedda /boot ext4 defaults 1 2
/dev/mapper/vg_main-lv_opt /opt ext4 defaults 1 2
/dev/mapper/vg_main-lv_var /var ext4 defaults 1 1
/dev/mapper/vg_main-lv_applogs /applogs ext4 defaults 1 2
```

Above we see the contents of the **fstab** file (your own **fstab** will differ) and we know that the **journalctl -xe** was complaining about **/dev/mapper/vg_main-lv_var**

Each entry in the **fstab** file has six fields:

1	2	3	4	5 6
/dev/mapper/vg_main-lv_var	/var	ext4	defaults	1 1

It is the last number we want to change, and we will change this to a zero, this will mark this partition (and only this partition) to ignore **fsck** at boot.

So we will change:

```
/dev/mapper/vg_main-lv_var /var ext4 defaults 1 1
to
/dev/mapper/vg_main-lv_var /var ext4 defaults 1 0
```

Save this file, and now you should be able to do a test reboot, and the system should boot normally to user login page rather than emergency mode.

From:

<http://cameraangle.co.uk/> - WalkerWiki - wiki.alanwalker.uk

Permanent link:

http://cameraangle.co.uk/doku.php?id=centos_not_booting

Last update: **2023/03/09 22:35**

