

## Admins - Anomalie #4236

### Reboot super long de calamus

02/02/2020 16:52 - Christian P. Momon

<b>Statut:</b>	Fermé	<b>Début:</b>	02/02/2020
<b>Priorité:</b>	Élevée	<b>Echéance:</b>	
<b>Assigné à:</b>	Christian P. Momon	<b>% réalisé:</b>	0%
<b>Catégorie:</b>		<b>Temps estimé:</b>	0.00 heure
<b>Version cible:</b>	Mai 2020	<b>Temps passé:</b>	0.00 heure
<b>Difficulté:</b>	2 Facile		
<b>Description</b>			
Le 02/02/2020, suite à un besoin d'upgrade, reboot de calamus : durée ~45 minutes. À noter que le teste de détection de fsck a répondu OK. Juste après, virola a rebootée dans un temps « habituel ».			
<b>Demandes liées:</b>			
Lié à Admins - Demande #4512: Ajouter la vérification des « Filesystem state ...		<b>Fermé</b>	<b>27/05/2020</b>

### Historique

#### #1 - 02/02/2020 16:55 - Christian P. Momon

- Statut changé de Nouveau à En cours de traitement
- Assigné à mis à Christian P. Momon
- Version cible changé de Backlog à Février 2020

Ci-après, le copier/coller de ma session ssh :

```
cpm@ocmstar (14:57:03) ~ 1 > sshapril root@calamus.april.org
Linux calamus.april.org 4.19.0-6-amd64 #1 SMP Debian 4.19.67-2+deb10u2 (2019-11-11) x86_64
```

```
The programs included with the Debian GNU/Linux system are free software;
the exact distribution terms for each program are described in the
individual files in /usr/share/doc/*/copyright.
```

```
Debian GNU/Linux comes with ABSOLUTELY NO WARRANTY, to the extent
permitted by applicable law.
```

```
Last login: Sun Feb  2 11:59:22 2020 from 91.160.139.56
```

```
(April) root@calamus:~# virsh list
```

Id	Name	State
1	adl	running
3	agir	running
5	bots	running
7	candidatsfr	running
12	lamp	running
15	pad	running
16	photos	running
17	pouet	running
18	republique-numerique	running
19	scm	running
21	webchat	running

```
(April) root@calamus:~# for disk in `mount -t ext4 | awk '{print $1}'`; do tune2fs -l $disk 2>/dev/null | /srv
/scripts/common/check_tune2fs.py ; done
```

```
Mount count and Fsck delay are OK
```

```
Mount count and Fsck delay are OK
```

```
Mount count and Fsck delay are OK
```

```
Mount count and Fsck delay are OK
```

```
Mount count and Fsck delay are OK
```

```
(April) root@calamus:~# uptime
```

```
13:58:02 up 19 days,  2:13,  1 user,  load average: 5.57, 5.86, 5.83
```

```
(April) root@calamus:~# apt-get update
```

```
Hit:1 http://apt.april.org/debian buster InRelease
```

```
Hit:2 http://apt.easter-eggs.com/debian buster-ee InRelease
```

```
Hit:3 http://ftp.fr.debian.org/debian buster InRelease
Hit:4 http://security.debian.org buster/updates InRelease
Reading package lists... Done
(April) root@calamus:~# apt-get dist-upgrade
Reading package lists... Done
Building dependency tree
Reading state information... Done
Calculating upgrade... Done
0 upgraded, 0 newly installed, 0 to remove and 0 not upgraded.
(April) root@calamus:~# cd /etc/libvirt/qemu
(April) root@calamus:/etc/libvirt/qemu[master u=# for host in $(ls *xml | sed -e 's/.xml//g') ; do virsh shut
down $host ; done; watch virsh list
Domain adl is being shutdown

error: Failed to shutdown domain admin
error: Requested operation is not valid: domain is not running

Domain agir is being shutdown

error: Failed to shutdown domain bastion
error: Requested operation is not valid: domain is not running

Domain bots is being shutdown

error: Failed to shutdown domain candidatsbe
error: Requested operation is not valid: domain is not running

Domain candidatsfr is being shutdown

error: Failed to shutdown domain cms-dev
error: Requested operation is not valid: domain is not running

error: Failed to shutdown domain dns
error: Requested operation is not valid: domain is not running

error: Failed to shutdown domain drupal6
error: Requested operation is not valid: domain is not running

error: Failed to shutdown domain dtc
error: Requested operation is not valid: domain is not running

Domain lamp is being shutdown

error: Failed to shutdown domain mail
error: Requested operation is not valid: domain is not running

error: Failed to shutdown domain mumble
error: Requested operation is not valid: domain is not running

Domain pad is being shutdown

Domain photos is being shutdown

Domain pouet is being shutdown

Domain republique-numerique is being shutdown

Domain scm is being shutdown

error: Failed to shutdown domain sympa
error: Requested operation is not valid: domain is not running

Domain webchat is being shutdown

(April) root@calamus:/etc/libvirt/qemu[master u=# umount /var/lib/libvirt/calamus
(April) root@calamus:/etc/libvirt/qemu[master u=# drbdadm disconnect virola
(April) root@calamus:/etc/libvirt/qemu[master u=# drbdadm secondary calamus
(April) root@calamus:/etc/libvirt/qemu[master u=# drbdadm disconnect calamus
(April) root@calamus:/etc/libvirt/qemu[master u=# reboot
===== ATTENTION !! =====

Vous êtes sur le point de rebooter calamus.

Il s'agit d'un hyperviseur qui requiert une attention particulière.
```

Avant de procéder, veuillez svp ouvrir la page  
[https://admin.april.org/dokuwiki/doku.php?id=sysadm:redemarrer\\_calamus](https://admin.april.org/dokuwiki/doku.php?id=sysadm:redemarrer_calamus)  
et vérifier que vous suivez les points de la procédure  
de redémarrage.

En particulier, assurez vous que quelqu'un de qualifié est  
disponible pour intervenir physiquement en cas de besoin.

Pour continuer la procédure de reboot, saisissez RjU7

Saisissez le mot de passe : RjU7

W: molly-guard: SSH session detected!

Please type in hostname of the machine to reboot: calamus

Connection to calamus.april.org closed by remote host.

Connection to calamus.april.org closed.

```
[...]
cpm@ocmstar (15:16:50) ~ 8 > sshapril root@calamus.april.org
^[[1;2Bssh: connect to host calamus.april.org port 22: Connection timed out

cpm@ocmstar (15:49:02) ~ 23 > sshapril root@calamus.april.org
Linux calamus.april.org 4.19.0-6-amd64 #1 SMP Debian 4.19.67-2+deb10u2 (2019-11-11) x86_64
[...]
```

## #2 - 02/02/2020 16:56 - Christian P. Momon

Extrait des logs systèmes :

```
/var/log/messages :
Feb  2 14:04:42 calamus kernel: [1649977.633223] drbd calamus: Terminating drbd_a_calamus
Feb  2 14:04:42 calamus kernel: [1649977.756268] drbd calamus: Connection closed
Feb  2 14:04:42 calamus kernel: [1649977.756285] drbd calamus: conn( Disconnecting -> StandAlone )
Feb  2 14:04:42 calamus kernel: [1649977.756287] drbd calamus: receiver terminated
Feb  2 14:04:42 calamus kernel: [1649977.756294] drbd calamus: Terminating drbd_r_calamus
Feb  2 14:04:57 calamus kernel: [1649992.302489] block drbd1: disk( UpToDate -> Failed )
Feb  2 14:04:57 calamus kernel: [1649992.302625] block drbd1: 0 KB (0 bits) marked out-of-sync by on disk bit-
map.
Feb  2 14:04:57 calamus kernel: [1649992.302629] block drbd1: disk( Failed -> Diskless )
Feb  2 14:46:34 calamus kernel: [ 0.000000] Linux version 4.19.0-6-amd64 (debian-kernel@lists.debian.org) (
gcc version 8.3.0 (Debian 8.3.0-6)) #1 SMP Debian 4.19.67-2+
deb10u2 (2019-11-11)
Feb  2 14:46:34 calamus kernel: [ 0.000000] Command line: BOOT_IMAGE=/boot/vmlinuz-4.19.0-6-amd64 root=/dev
/mapper/vg_calamus-root ro quiet
Feb  2 14:46:34 calamus kernel: [ 0.000000] x86/fpu: Supporting XSAVE feature 0x001: 'x87 floating point re
gisters'
Feb  2 14:46:34 calamus kernel: [ 0.000000] x86/fpu: Supporting XSAVE feature 0x002: 'SSE registers'
Feb  2 14:46:34 calamus kernel: [ 0.000000] x86/fpu: Supporting XSAVE feature 0x004: 'AVX registers'
Feb  2 14:46:34 calamus kernel: [ 0.000000] x86/fpu: xstate_offset[2]: 576, xstate_sizes[2]: 256
Feb  2 14:46:34 calamus kernel: [ 0.000000] x86/fpu: Enabled xstate features 0x7, context size is 832 bytes
, using 'standard' format.
```

C'est très caractéristique d'un fsck.

## #3 - 02/02/2020 17:00 - Christian P. Momon

Vérification des valeurs utilisées par la commande de détection d'un fsck prochain, réalisé après le reboot super long de calamus :

```
(April) root@calamus:/var/log# for disk in `mount -t ext4 | awk '{print $1}'`; do tune2fs -l $disk 2>/dev/null
; done |egrep -i "(Last checked|interval|mount count)"
Mount count:          75
Maximum mount count: -1
Last checked:        Mon Mar  7 16:24:15 2016
Check interval:      0 (<none>)
Mount count:          75
Maximum mount count: -1
Last checked:        Mon Mar  7 16:24:18 2016
Check interval:      0 (<none>)
Mount count:          75
Maximum mount count: -1
Last checked:        Mon Mar  7 16:24:18 2016
Check interval:      0 (<none>)
Mount count:          72
```

```
Maximum mount count:      -1
Last checked:              Mon Mar  7 19:54:05 2016
Check interval:           0 (<none>)
```

#### #4 - 02/02/2020 17:01 - Christian P. Momon

Valeurs avant le reboot de virola :

```
cpm@ocmstar (15:55:33) ~ 24 > sshapril root@virola.april.org
Linux virola.april.org 4.19.0-6-amd64 #1 SMP Debian 4.19.67-2+deb10u2 (2019-11-11) x86_64
[...]
(April) root@virola:/etc/libvirt/qemu[master$ u=# for disk in `mount -t ext4 | awk '{print $1}'`; do tune2fs
-l $disk 2>/dev/null ; done |grep -i "(Last checked|interval|mount count)"
Mount count:              99
Maximum mount count:      -1
Last checked:              Mon Mar  7 16:05:58 2016
Check interval:           0 (<none>)
Mount count:              99
Maximum mount count:      -1
Last checked:              Mon Mar  7 16:06:04 2016
Check interval:           0 (<none>)
Mount count:              99
Maximum mount count:      -1
Last checked:              Mon Mar  7 16:06:03 2016
Check interval:           0 (<none>)
Mount count:              96
Maximum mount count:      -1
Last checked:              Mon Mar  7 19:52:45 2016
Check interval:           0 (<none>)
Mount count:              78
Maximum mount count:      -1
Last checked:              Mon Mar  7 19:57:03 2016
Check interval:           0 (<none>)
```

Valeurs après le reboot de virola :

```
Last login: Sun Feb  2 15:12:47 2020 from 91.160.139.56
(April) root@virola:~# for disk in `mount -t ext4 | awk '{print $1}'`; do tune2fs -l $disk 2>/dev/null ; done
|grep -i "(Last checked|interval|mount count)"
Mount count:              100
Maximum mount count:      -1
Last checked:              Mon Mar  7 16:05:58 2016
Check interval:           0 (<none>)
Mount count:              100
Maximum mount count:      -1
Last checked:              Mon Mar  7 16:06:03 2016
Check interval:           0 (<none>)
Mount count:              100
Maximum mount count:      -1
Last checked:              Mon Mar  7 16:06:04 2016
Check interval:           0 (<none>)
Mount count:              97
Maximum mount count:      -1
Last checked:              Mon Mar  7 19:52:45 2016
Check interval:           0 (<none>)
Mount count:              79
Maximum mount count:      -1
Last checked:              Mon Mar  7 19:57:03 2016
Check interval:           0 (<none>)
```

#### #5 - 02/02/2020 17:05 - Christian P. Momon

- *Priorité changé de Normale à Élevée*

À en croire les valeurs suivantes, la détection ne peut pas fonctionner.

```
Mount count:              100
Maximum mount count:      -1
Last checked:              Mon Mar  7 16:05:58 2016
Check interval:           0 (<none>)
```

Questions :

- pourquoi un fsck s'est-il déclenché quand même ?
- Le champs « Last checked » est-il le bon ?
- Que faut-il modifier dans la commande de détection ?
- Faut-il forcer un fsck préventif ?
- Des sondes existent-elles et alors faut-il les modifier ? Ou faut-il en créer ?

#### #6 - 02/02/2020 18:39 - François Poulain

Tu peux aussi avoir un soucis sur le volume, indétecté par la sonde, mais qui pousse au fsck via le paramètre pass du fstab, non ?

#### #7 - 26/02/2020 22:24 - Quentin Gibeaux

- Version cible changé de Février 2020 à Mars 2020

#### #8 - 25/03/2020 21:36 - Quentin Gibeaux

- Version cible changé de Mars 2020 à Avril 2020

#### #9 - 27/04/2020 00:13 - Christian P. Momon

François Poulain a écrit :

Tu peux aussi avoir un soucis sur le volume, indétecté par la sonde, mais qui pousse au fsck via le paramètre pass du fstab, non ?

Extrait du man tune2fs :

```
A filesystem error detected by the kernel will still force an fsck on the next reboot, but it may already be too late to prevent data loss at that point.
```

Donc oui. À la question « **pourquoi un fsck s'est-il déclenché quand même ?** », on va donc dire que c'est parce que le noyau a détecté un problème sur le disque.

#### #10 - 27/04/2020 00:59 - Christian P. Momon

Pour information, à la question « pourquoi les valeurs sont-elles désactivée », en fait, il semble que ce soit le comportement par défaut :

```
(April) root@virola:~# grep period /etc/mke2fs.conf
    enable_periodic_fsck = 0
```

En cherchant encore, ça date depuis 2011 :

```
From: Eric Sandeen <sandeen@redhat.com>
Date: Thu, 17 Feb 2011 15:55:15 -0600
Subject: [PATCH] e2fsprogs: turn off enforced fsck intervals by default
```

```
The forced fsck often comes at unexpected and inopportune moments, and even enterprise customers are often caught by surprise when this happens. Because a filesystem with an error condition will be marked as requiring fsck anyway, I submit that the time-based and mount-based checks are not particularly useful, and that administrators can schedule fscks on their own time, or tune2fs the enforced intervals if they so choose. This patch disables the intervals by default, and I've added a new mkfs.conf option to turn on the old behavior of random, unexpected, time-consuming fscks at boot time.
```

#### #11 - 29/04/2020 21:22 - Quentin Gibeaux

- Version cible changé de Avril 2020 à Mai 2020

#### #12 - 27/05/2020 01:37 - Christian P. Momon

Autre indication trouvée dans le man de tune2fs :

```
Mount-count-dependent checking is disabled by default to avoid unanticipated long reboots while e2fsck does its work. [...]
```

```
If you are using journaling on your filesystem, your filesystem will never be marked dirty, so it will not normally be checked.
```

Donc, il est normal que *Check interval* et *Maximum mount count* soient désactivés.

#### #13 - 27/05/2020 18:40 - Christian P. Momon

Malgré le fait que *Check interval* et *Maximum mount count* soient désactivés, en cherchant encore, on trouve que le superblock ext4 contient un champ dédié ([https://ext4.wiki.kernel.org/index.php/Ext4\\_Disk\\_Layout](https://ext4.wiki.kernel.org/index.php/Ext4_Disk_Layout)) :

```
s_state File system state. Valid values are:
0x0001   Cleanly unmounted
0x0002   Errors detected
0x0004   Orphans being recovered
```

Ce champ est modifiable via la commande *debugfs*, extrait du man :

```
dirty [-clean]
Mark the filesystem as dirty, so that the superblocks will be written on exit.  Additionally, clear the superblock's valid flag, or set it if -clean is specified.
```

Donc quand le man de *tune2fs* dit « A filesystem error detected by the kernel will still force an fsck on the next reboot », on peut penser que le noyau modifie ce flag et que c'est ce flag qui déclenche un fsck.

Le man de *e2fsck* nous le confirme :

```
pour les systèmes de fichiers journalisés, e2fsck applique le journal puis quitte, à moins que le superbloc n'indique que d'autres vérifications sont nécessaires.
```

#### #14 - 27/05/2020 18:52 - Christian P. Momon

- Statut changé de *En cours de traitement* à *Résolu*

On va dire que c'est ça et donc le ticket peut être passé à résolu.

Suggestion : ajouter une détection d'état du superblock avant le reboot ([#4512](#)).

#### #15 - 27/05/2020 18:53 - Christian P. Momon

- Lié à *Demande #4512: Ajouter la vérification des « Filesystem state » avant le reboot d'une pm ajouté*

#### #16 - 27/05/2020 21:54 - Quentin Gibeaux

- Statut changé de *Résolu* à *Fermé*