

# Temporary 'lock-up' under heavy write, MegaRAID RAID-5

*Source:* <http://linux.derkeiler.com/Mailing-Lists/Debian/2005-11/1145.html>

---

**From:** Dave Ewart (*davee\_at\_ceu.ox.ac.uk*)

**Date:** 11/09/05

Date: Wed, 9 Nov 2005 10:46:24 +0000

To: debian-amd64@lists.debian.org, debian-user@lists.debian.org

System: 4-way Opteron, generic Debian Sarge AMD64  
RAID controller: LSI Logic MegaRAID 320-1, 64MB cache  
RAID config: Three 146GB 15K SCSI/320 disks, RAID-5  
Kernel: 2.6.14 SMP, includes megaraid driver

The above system is incredibly fast under almost all conditions, except when writing very large files (say, 100s of MB, or even GB). When writing such files, the system effectively locks-up for many seconds – typically, for as long as it takes to finish writing/flushing the file to disk. This lock-up affects all other processes: local text editor sessions, workstations with /home NFS-mounted, web server stops serving. (I guess all the affected processes are those which are contending for disk write access, actually). In particular, the workstations which have /home NFS-mounted experience a \*workstation\* hang (if trying to write) during the \*server\* disk flush, which is very frustrating. Given that a 'write' may simply involve updating a web browser history stored in /home, this is an extremely serious problem.

Example while system is idle, out of work hours: while creating a 1GB file (copying an existing file, already cached in RAM), 'vmstat 3' shows the following:

```
procs -----memory-----swap-----io-----system-- -----cpu-----
r b swpd free buff cache si so bi bo in cs us sy id wa
0 0 0 4280268 55888 3366484 0 0 0 0 260 49 0 0 100 0
1 1 0 3777432 56364 3847652 0 0 0 6419 319 55 0 23 77 0
1 4 0 3258408 56852 4342476 0 0 0 10243 407 45 0 25 75 0
0 3 0 3070296 57028 4520868 0 0 0 8856 403 99 0 9 75 16
0 4 0 3068152 57044 4520852 0 0 5 9561 417 153 0 1 72 27
0 3 0 3069316 57044 4520852 0 0 0 10240 429 144 0 0 75 25
0 3 0 3069356 57044 4520852 0 0 0 10219 411 85 0 0 75 25
0 3 0 3069368 57044 4520852 0 0 0 8876 391 78 0 0 75 25
[...]
0 2 0 3077856 57044 4520852 0 0 0 9557 409 44 0 0 75 25
0 2 0 3077856 57044 4520852 0 0 0 8875 384 41 0 0 75 25
```

Debian–User: Temporary 'lock–up' under heavy write, MegaRAID RAID–5

0 1 0 3097748 57044 4520852 0 0 0 7704 421 42 0 2 73 25  
0 0 0 3100096 57048 4520848 0 0 0 56 259 20 0 0 99 1  
0 0 0 3100112 57052 4520844 0 0 0 552 362 32 0 0 97 3  
0 0 0 3100