

Re: amrd disk performance drop after running under high load

Re: amrd disk performance drop after running under high load

Source: <http://unix.derkeiler.com/Mailing-Lists/FreeBSD/hackers/2007-10/msg00133.html>

- *From:* Krassimir Slavchev <krassi@xxxxxxxxxxxx>
 - *Date:* Tue, 16 Oct 2007 11:51:33 +0300
-

-----BEGIN PGP SIGNED MESSAGE-----

Hash: SHA1

Alexey Popov wrote:

Hi.

Kris Kennaway wrote:

After some time of running under high load disk performance become extremely poor. At that periods 'systat -vm 1' shows something like this:

What does "high load" mean? You need to explain the system workload more.

This web service is similiar to YouTube. This server is video store. I have around 200G of *.flv (flash video) files on the server.

I run lighttpd as a web server. Disk load is usually around 50%, network output 100Mbit/s, 100 simultaneous connections. CPU is mostly idle.

As you can see it is a trivial service – sending files to network via HTTP.

```
Disks amrd0
KB/t 85.39
tps 5
MB/s 0.38
% busy 99
```

Re: amrd disk performance drop after running under high load

Apart of all, I tried to make mutex profiling and here's the results
(sorted by the total number of acquisitions):

Bad case:

```
102 223514 273977 0 14689 1651568
/usr/src/sys/vm/uma_core.c:2349 (512)
950 263099 273968 0 15004 14427
/usr/src/sys/vm/uma_core.c:2450 (512)
108 150422 175840 0 10978 22988519
/usr/src/sys/vm/uma_core.c:1888
(mbuf)
```

> Here you can see that high UMA activity happens in periods of low disk
> performance. But I'm not sure whether this is a root of the
problem, not
> a consequence.

The extremely high contention there does seem to say you have a mbuf
starvation problem and not a disk problem. I don't know why this
would be happening off-hand.

But there's no mbuf shortage in `netstat -m`.

What else can I try to track down the source of the problem?

Can you also provide more details about the system hardware and
configuration?

This is Dell 2850 2 x Xeon 3.2, 4Gb RAM, 6x300Gb SCSI RAID5. I'll attach
details.

With best regards,
Alexey Popov

freebsd-stable@xxxxxxxxxxxxx mailing list
<http://lists.freebsd.org/mailman/listinfo/freebsd-stable>
To unsubscribe, send any mail to "freebsd-stable-unsubscribe@xxxxxxxxxxxxx"

last pid: 11008; load averages: 0.07, 0.10, 0.08 up 47+08:32:50
11:46:15
38 processes: 1 running, 37 sleeping

Re: amrd disk performance drop after running under high load

Re: amrd disk performance drop after running under high load

Mem: 46M Active, 3443M Inact, 246M Wired, 144M Cache, 208M Buf, 5596K Free
Swap: 2048M Total, 4K Used, 2048M Free

```
PID USERNAME THR PRI NICE SIZE RES STATE C TIME WCPU COMMAND
56386 root 1 4 0 19856K 10000K kqread 1 115:19 2.88% lighttpd
636 root 1 96 0 18292K 4212K select 0 25:39 0.00% snmpd
784 root 1 96 0 19668K 2072K select 1 2:31 0.00% sshd
680 root 1 96 0 7732K 1384K select 0 1:59 0.00% ntpd
1540 root 1 96 0 35092K 6496K select 0 1:30 0.00% httpd
769 root 4 20 0 14148K 2632K kserel 0 1:04 0.00% bacula-fd
755 root 1 96 0 3852K 1060K select 1 0:22 0.00% master
568 root 1 96 0 3648K 908K select 0 0:18 0.00% syslogd
80663 root 1 8 0 3688K 1016K nanslp 1 0:05 0.00% cron
760 postfix 1 96 0 3944K 1160K select 0 0:04 0.00% qmgr
89776 www 1 20 0 35180K 6684K lockf 0 0:04 0.00% httpd
89763 www 1 20 0 35180K 6684K lockf 0 0:04 0.00% httpd
89774 www 1 20 0 35180K 6684K lockf 0 0:04 0.00% httpd
89775 www 1 96 0 35180K 6684K select 0 0:04 0.00% httpd
699 root 1 20 0 7732K 1388K pause 0 0:03 0.00% ntpd
484 root 1 4 0 652K 220K select 0 0:00 0.00% devd
10904 llp 1 96 0 30616K 3564K select 0 0:00 0.00% sshd
10915 root 1 20 0 3912K 2340K pause 1 0:00 0.00% csh
```

You run apache with mod_perl or php too. How many clients handle this apache server? Also in this light load you have locked files! Check script execution times (/server-status may be useful).

When you have hight load check swap usage and haw many processes are in lockf state.

Best Regards

-----BEGIN PGP SIGNATURE-----

Version: GnuPG v1.4.7 (FreeBSD)

iD8DBQFHfHuVxJBWvpalMpkRAnqTAJ9FgURNk98dtD0HYX6xIz17R6sLpQCgh5nJ
XBtfOyzJJbkjzVzSF/WfmHc=
=oTHZ

-----END PGP SIGNATURE-----

freebsd-hackers@xxxxxxxxxxx mailing list

<http://lists.freebsd.org/mailman/listinfo/freebsd-hackers>

To unsubscribe, send any mail to "freebsd-hackers-unsubscribe@xxxxxxxxxxx"

Re: amrd disk performance drop after running under high load