

Re: slowdown on internal harddisks

Source: <http://unix.derkeiler.com/Mailing-Lists/FreeBSD/performance/2004-01/0003.html>

From: Uwe Doering (gemin_i_at_geminix.org)

Date: 01/05/04

Date: Mon, 05 Jan 2004 09:48:06 +0100

To: freebsd-performance@freebsd.org

Cor Bosman wrote:

- > *Hi all, im running a large usenet setup (30+ servers, 20TB or so) and am*
- > *running into some trouble with one of the servers.*
- >
- > *It's a dual Xeon 2.6, 4GB memory, mirrored 2 x 17GB / on ADAPTEC 2015S (asr)*
- > *extra 70GB disk on the same asr controller, and 3 x 1.5 TB on*
- > *Qlogic ISP 2300 PCI FC-AL Adapter (isp) set up as 6 750GB partitions.*
- >
- > *This server pushes out about 300mbit of data on an Intel Pro 1000 (em).*
- > *Im also using a 1GB MFS which i need to support a very high amount of seeks*
- > *on a usenet history file (300 MB used).*
- >
- > *As the server starts doing lots of io on the Qlogic controller all disk writes*
- > *seem to drop to about 0. Also on the MFS, and on the internal asr controller.*
- > *Copying a file from any disk to any other, even to the unused extra 70GB*
- > *disk has a max throughput of about 50KB/s. Read performance is still ok, it's*
- > *just write(s) that are slow. (cp-ing to /dev/null is fast).*

Try increasing the sysctl(8) variables 'vfs.hirunningspace' and 'vfs.lorunningspace', for instance to '8388608' and '6291456', respectively. Lots of pending read operations at the same time tend to throttle write operations. This is by (kernel) design in order to give priority to read operations.

While the (lower) default values appear to work well with conventional "dumb" disk controllers we have found that they are insufficient for more sophisticated caching disk controllers. The effect of too low settings would be exactly what you're seeing: Good read response, but anemic write performance.

Maybe that tweak helps in your case, too.

Uwe

--

Uwe Doering | EscapeBox - Managed On-Demand UNIX Servers
gemin_i@geminix.org | <http://www.escapebox.net>

freebsd-performance: Re: slowdown on internal harddisks

freebsd-performance@freebsd.org mailing list

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

To unsubscribe, send any mail to "freebsd-performance-unsubscribe@freebsd.org"