Re: about ufs filesystem io performance!

Source: http://unix.derkeiler.com/Mailing−Lists/FreeBSD/performance/2006−05/msg00095.html

- *From*: Eric Anderson <anderson@xxxxxxxxxxxx>
- *Date*: Tue, 30 May 2006 07:40:54 −0500

etalk etalk wrote:

> From: Eric Anderson <anderson@xxxxxxxxxxxx>
> To: etalk etalk <yanyuejin2004@xxxxxxxxxxx>
> CC: freebsd−fs@xxxxxxxxxxx, freebsd−performance@xxxxxxxxxxx
> Subject: Re: about ufs filesystem io performance!
> Date: Thu, 25 May 2006 07:46:44 −0500
>
> etalk etalk wrote:
>
> > 5.3 vs 6.0 The test tool is Iozone3_257, and the test command is ?/iozone −A −f /mnt/tmpfile.test −g 1g −n 1m −q 8k −y 2k −R −b outfile−Af.xls ?(http://www.iozone.org/src/current/). We ran all the tests on the same PC with 2.4 GHz Pentium CPU and 512M main memory. Figure1~Figure5 show the results of the file system performance comparison between Bsd5.3¯ UFS2 and Bsd6.0¯ UFS2 when testing with different file system (local, sync, async, softupdate, sync+softupdate).
> > According to the figures, our conclusion is: On all kinds of file systems, the write, rewrite, read and reread performance of the two is almost same and we cant say that Bsd6.0 make a improvement on file system IO performance.
> > http://blog.csdn.net/minerboyIo/Gallery/204114.aspx
> > linux2.6.11 vs bsd 5.3 The test tool is Iozone3_257, and the test command is ?/iozone −A −f /mnt/tmpfile.test −g 1g −n 4m −q 8k −y 2k −R −b outfile−Af.xls ?(http://www.iozone.org/src/current/). We ran all the tests on the same PC with 2.4 GHz Pentium CPU and 512M main memory, Figure1, Figure2, Figure3 show the results of the file system performance comparison between Bsd¯ UFS2 and Linux?Ext3 (the Linux kernel version is 2.6.11, and the Bsd kernel version is 5.3) when testing with sync, async and local (Bsd using softupdate) file system. According to the figures, our conclusion is: a.On local file system and async file system, Fedora4¯ write and rewrite is much faster than Bsd5.3¯ (about 5−10 times). b.On all kinds of file systems,

the read and reread performance of FreeBsd5.3 is about 50%−90% lower than that of Fedora4. c.On sync file system, Bsd5.3 writes several times faster than Fedora4 does and rewrites over two hundred times faster than Fedora4 does. http://blog.csdn.net/minerboyIo/Gallery/204107.aspx

You don't report the type of disks you are using, or anything about the storage. For the first test, I'd think that it's possible that you were hitting hardware performance bottlenecks before actually testing the filesystem performance.

Also, what are the 2,4,8 numbers referencing? How many times did you run the tests?

Eric

−−
−−
Eric Anderson Sr. Systems Administrator Centaur Technology
Anything that works is better than anything that doesn't.
−−

Thanks for your reply!
My disk is Maxtor 2F040L0 with 40GB capacity , 5400 rpm and <12ms ave seek time, and the 2,4,8 is the block size when doing the writes and reads, We do the tests two times,and the results is same,One of the results is in the attachment!

According to the conclusions,our puzzles is :

a. Why the write and rewrite performance of FreeBsd5.3 is so lower than that of Fedora4 in async system or in local system? Can we improve the performace by tuning the FreeBsd5.3's kernel or by making some modifition to the kernel of FreeBsd5.3 in the file vfs_bio.c?

b. Is Bsd6.0 make improvement in file system io performance when comparing to Bsd5.3?

I am eager to have your reply!

Best Regards

etalk

I'm no expert, but the drive you used is (in my opinion) insufficient to fully test filesystem's. You really need to remove as much hardware from the bottleneck path as possible. Preferably using something like a Gigabyte

i−Ram device, or possibly memory disks.

FreeBSD 6.1+ has many many filesystem improvements, so I would highly suggest re−running the benchmarks on that, and also on the newer 2.6 Linux kernels. I also suggest more than 2 runs − something like 5−7 would be a good start I think.

You might also do some tests on very large (100Tb or bigger) and very small (30mb or smaller) filesystem sizes.

Eric

−−
−−−
Eric Anderson Sr. Systems Administrator Centaur Technology
Anything that works is better than anything that doesn't.
−−−