

Re: SATA 3ware RAID review...sort of.

Source: <http://unix.derkeiler.com/Mailing-Lists/FreeBSD/isp/2005-02/0005.html>

From: Simon (*simon_at_optinet.com*)

Date: 02/04/05

To: "freebsd-isp@freebsd.org" <freebsd-isp@freebsd.org>, "Gerald" <gcoon@inch.com>
Date: Fri, 04 Feb 2005 17:59:13 -0500

While your review is helpful, I hope you realize that server's throughput doesn't indicate the load you are putting on the card or the harddrives. If you were serving large movie files, you could quickly fill up both of your 100mbps links with a single 5.4k RPM ATA drive, given you have some RAM so FreeBSD could cache the data being served. 30Mbps of bandwidth is merely ~3mBps of sustained disk transfer. Most of this data could be cached by FreeBSD without you even knowing. It could also be cached by the memory on the control and harddrives.

What would be very useful is if you could mention what kind of files are being served (their size), how many hits the server serves every second. Are they mostly the same files or completely random, and so on... how often is the server writing to disk, is it creating many random files? how big? how often? there is just so many things involved, you can't merely post your hardware and say you are pushing 30mbps of bandwidth. I have a 8x250gb using PATA and 3ware 8port 8600 series card which can do sustained reads of over 100mBps which would translate into ~1000mbps of bandwidth. However, this doesn't indicate that it would be able to serve 50,000,000 small, mostly random files, a day like a similar server using SCSI could. Don't forget, it's one thing when you just read, but completely different when you read and write, especially with RAID5.

-Simon

On Fri, 4 Feb 2005 17:23:22 -0500 (EST), Gerald wrote:

>*Details of a web server upgrade using SATA RAID5 drives.*

>

>*Old Machine:*

>*FreeBSD 4.8 (CUSTOM), Dual Pentium III 850s, 1 GB RAM*

>*ADAPTEC 3200S FW Rev. 370F*

>*External disk case for SCSI SCA 10k RPM drives.*

>

>*Problem: Not enough memory. No way to upgrade RAM on existing MB.*

>*Solution: Kill soon to be hard drive issue and memory problem at same*

>*time with new machine. (Long story about "hard drive issue"s on old*

>*machine.)*

freebsd-isp: Re: SATA 3ware RAID review...sort of.

>
>After shopping around we decided to try the SATA RAID setup for this
>web server. Upon advice from this list, I went for the 3ware 9000 series
>controller.
>
>New Machine:
>Freebsd 5.3 (GENERIC), Dual Xeon 3.06GHz, 4GB RAM
>3ware 9500S8CH
>Disks: (5x) 7200RPM 8MB cache SATA150 RAID5'ed
>
>Some information about this server: It's totally dedicated to serving
>web pages for one site. Site traffic during down months has a 95th
>percentile of about 15 Mb. During Feb and Sep it jumps to 25 Mb and
>this month we expect 30 Mb 95th percentile. For those that do transfer
>numbers, the last week of September 04 saw 895.96 gigabytes total
>transferred and 127.99 gigabytes average per day. (7 days)
>
>On the old server I watched systat -v 1 for about 30 minutes recently
>to see how busy the individual disks were during normal load. Most
>partitions bounced around from 20% busy to 75% busy. This was a SCSI
>setup though so I was concerned about moving over to SATA for something
>that uses disk I/O so much.
>
>The new setup is 1 TB of usable space. The new setup allows all of the
>data to be on one partition that is RAID5'ed. I switched to the new
>server Wed night at midnight and peak traffic hit today. They went from
>15 mb to 30 Mb between 6AM and 8AM. (INCH has redundant 100 Mb upstream
>connections.)
>
>Systat and top now show the machine taking the load in stride.
>LA: 0.94 0.52 0.58
>Disk usage: hovers around 50% (40-55% predominantly)
>CPU: stays about 80% Idle
>
>I obviously don't have the apache defaults to serve up this many pages,
>but I wanted to share the SATA success on a high usage web server. At
>least from my perspective you can easily do 30 Mb of web traffic on a
>properly configured RAID5 SATA system. If I had to guess, you could make
>it to somewhere in the neighborhood of 75-100 Mb of traffic before you
>>would run in to I/O problems configured the way I have it. I wouldn't do
>that much on it though since a drive failure would put your processes
>waiting on the Disk I/O.
>
>Problems related to drives:
>
>The twa driver would not let me create more than I think 7 partitions.
>I'm pretty sure it was 7. I've slept since then. It tried to create the
>like "/dev/...i" as /dev/X literally. I just reworked my FS layout to
>work around this. This might just be the 5.3 installer though. I don't
>know. I didn't spend too much time on this.
>

freebsd-isp: Re: SATA 3ware RAID review...sort of.

>Concessions:

>- There are many many changes between FreeBSD 4.8 and 5.3. I'm

>benefiting from quite a few improvements in that upgrade.

>

>- UFS -> UFS2 probably plays quite a bit in to the above information.

>

>- If I could have upgraded the RAM on the old machine, I could have made

>do with a memory upgrade for an indeterminate amount of time. The memory

>is what I needed the most.

>

>- I've not tested (even for the fun of it) a drive failure and the

>ensuing load should such a situation happen. I feel comfortable that I

>have enough breathe room with the current load on the disks.

>

>- I did stress test the disks with bonnie++ from the ports and it seemed

>to do well.

>

>- You can squeeze even more out of this setup with 10k RPM drives, but I

>had to do 7200s to keep within customer's budget and it's handling the

>load quite well.

>

>- YMMV, This information comes with no warranty either expressed or

>implied...etc etc. Contact me if you want a server setup the same way.

>

>Gerald Coon

>System Administrator

>Internet Channel

>

>

freebsd-isp@freebsd.org mailing list

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

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

>

freebsd-isp@freebsd.org mailing list

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

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