

Re: FW: 20TB Storage System

Source: <http://unix.derkeiler.com/Mailing-Lists/FreeBSD/performance/2003-09/0007.html>

From: Poul-Henning Kamp (*phk_at_phk.freebsd.dk*)

Date: 09/02/03

To: "Max Clark" <max.clark@media.net>

Date: Tue, 02 Sep 2003 22:01:44 +0200

In message <ILENIMHFIPIBHJLCDEHKKIEMEDCAA.max.clark@media.net>, "Max Clark" writes:

>Given the above:

>1) What would my expected IO be using vinum to stripe the storage enclosures
>detailed above?

That depends a lot on the applications I/O pattern, and I doubt a precise prediction is possible.

In particular the FibreChannel is hard to predict the throughput of because the various implementations seem to have each their own peculiar quirks performance wise.

On a SEAGATE ST318452 disks, I see sequential transfer rates at the outside rim of the disk of 58MB/sec. If I stripe two of them with CCD I get 107MB/sec.

CCD has a better performance than Vinum where they compare.

RAID-5 and striping a large number of disks does not scale linearly performance wise, in particular you *may* see your average access time drop somewhat, but there is by far no guarantee that it will be better than the individual drive.

>2) What is the maximum size of a filesystem that I can present to the host
>OS using vinum/ccd? Am I limited anywhere that I am not aware of?

Good question, I'm not sure we currently know the exact barrier.

>3) Could I put all 20TB on one system, or will I need two to sustain the IO
>required?

Spreading it will give you more I/O bandwidth.

--

Poul-Henning Kamp | UNIX since Zilog Zeus 3.20
phk@FreeBSD.ORG | TCP/IP since RFC 956

freebsd-performance: Re: FW: 20TB Storage System

FreeBSD committer | BSD since 4.3-tahoe
Never attribute to malice what can adequately be explained by incompetence.

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"