Re: Old SUN NFS performance papers.

Source: http://unix.derkeiler.com/Mailing−Lists/FreeBSD/performance/2004−01/0036.html

---

From: Kris Kennaway (*kris_at_obsecurity.org*)
Date: 01/27/04

```
Date: Tue, 27 Jan 2004 13:04:45 −0800
To: Bj?rn Gr?nvall <bg@sics.se>
```

On Tue, Jan 27, 2004 at 08:17:58PM +0100, Bj?rn Gr?nvall wrote:
> On Sat, 24 Jan 2004 21:14:51 −0500 (EST)
> Robert Watson <rwatson@freebsd.org> wrote:
>
> > I haven't done much benchmarking on NFS lately, but something worth
> > remembering is that people have spent a lot of time researching and
> > optimizing TCP for a variety of connection types, whereas the NFS code
> > basically has a static implementation of RPC backoff and flow control that
> > hasn't evolved much.
>
> One reason that FreeBSD users experience poor NFSv3/TCP performance is
> that the defaults for rsize and wsize is unusually small, only
> 8k. Solaris and HP−UX defaults to 32k for a good reason. I guess TCP
> simply needs a little bit more data to chew on to be efficient.
>
> I tested this on 5.2−CURRENT and found that large file read
> performance went up from 56Mbit/s to 80Mbit/s, an improvement by 43%.
>
> I have written a patch that makes FreeBSD use the same defaults as
> Solaris and HP−UX. Note that with NFSv3 there is no risk associated
> with specifying to large values for [rw]size. The server automatically
> limits these values in the fsinfo rpc. Patch is attached.

If no−one picks this up in the next few days, can you please send−pr
it so it does not get lost?

Kris

- application/pgp−signature attachment: stored