

Re: Old SUN NFS performance papers.

Source: <http://unix.derkeiler.com/Mailing-Lists/FreeBSD/performance/2004-01/0035.html>

From: Björn Grönvall (bg_at_sics.se)

Date: 01/27/04

Date: Tue, 27 Jan 2004 20:17:58 +0100
To: Robert Watson <rwatson@freebsd.org>

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.

Cheers,
Björn

--

Bjorn Gronvall (Björn Grönvall)
Swedish Institute of Computer Science
PO Box 1263, S-164 29 Kista, Sweden
Email: bg@sics.se, Phone +46 -8 633 15 25
Cellular +46 -70 768 06 35, Fax +46 -8 751 72 30

```
-----  
/-----/  
| Schroedingers |  
|           Cat  |  
|-----|  
-----
```

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

freebsd-performance: Re: Old SUN NFS performance papers.

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

- application/octet-stream attachment: [nfsclient.patch](#)