

Re: Old SUN NFS performance papers.

Source: <http://unix.derkeiler.com/Mailing-Lists/FreeBSD/performance/2004-01/0033.html>

From: Dominic Mitchell (*dom_at_happygiraffe.net*)

Date: 01/25/04

Date: Sun, 25 Jan 2004 11:04:25 +0000
To: Robert Watson <rwatson@FreeBSD.ORG>

On Sat, Jan 24, 2004 at 09:14:51PM -0500, Robert Watson 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. TCP is aware of things like the pathwise-mtu to the
> server and adapts, whereas UDP just loses packets due to fragmentation,
> especially if you are using larger block sizes. Please do post your
> discoveries on performance@, and perhaps we could build an NFS performance
> tuning section in the FreeBSD Handbook (or if there's not that much
> content, add it to the FAQ)?*

I'm just playing with this... The first thing to note (probably) is to check that you can ping your server with a similiar size packet to the one you're using. I realised that my network isn't as robust as I thought it was very quickly yesterday, when pinging my server with an 8k packet. I was seeing 70% packet loss. The default ping showed no problems at all.

The reason I mention it is that I'd been playing with NFS tuning because I had been seeing lockups. But the fault really lies at a lower level than NFS, it appears.

-Dom

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"