

Re: Initial 6.1 questions

Source: <http://unix.derkeiler.com/Mailing-Lists/FreeBSD/performance/2006-06/msg00010.html>

- *From:* Robert Watson <rwatson@xxxxxxxxxxxxx>
 - *Date:* Mon, 12 Jun 2006 21:02:51 +0100 (BST)
-

On Mon, 12 Jun 2006, Danial Thom wrote:

This is a design change that is in the process of being reconsidered. I expect that HZ will not be 1000 in 7.x, but can't tell you whether it will go back to 100, or some middle ground. There are a number of benefits to a higher HZ, not least is more accurate timing of some network timer events. Since I don't have my hands in the timer code, I can't speak to what the decision process here is, or when any change might happen, but I do expect to see some change.

Will anything break if I tweak this downward?

No, shouldn't do. I wouldn't go below 100 though, as things like process statistics, involuntary context switches, etc, are all affected.

Finally, there is a known performance problem involving loopback network traffic and preemption, which results in additional context switches. You may want to try disabling preemption and see if/how that impacts your numbers. There has been seen quite a bit of discussion of this problem, and I expect to see a solution for it in the near future. This problem does not manifest for remote traffic, only loopback traffic.

I'm sending this traffic from an external device, receiving on an em controller with blackhole set to 1. So I assume this loopback issue doesn't apply to this test?

The above comments only refer to traffic being sent over `if_loop` interfaces or certain other deferred work scenarios. Basically, deferring of work to the netisr from a user thread rather than an interrupt thread results in a premature context switch.

Robert N M Watson
Computer Laboratory
Universty of Cambridge

Re: Initial 6.1 questions

freebsd-performance@xxxxxxxxxxx mailing list

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

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