

Re: What is the PREEMPTION option good for?

Re: What is the PREEMPTION option good for?

Source: <http://unix.derkeiler.com/Mailing-Lists/FreeBSD/arch/2006-11/msg00080.html>

- *From:* Bruce Evans <bde@xxxxxxxxxxxx>
 - *Date:* Fri, 24 Nov 2006 11:10:00 +1100 (EST)
-

On Thu, 23 Nov 2006, Ivan Voras wrote:

Ivan Voras wrote:

Bruce Evans wrote:

Most of the difference is caused by pgzero becoming too active with PREEMPTION.

Don't know about the other things but I've noticed pagezero is suspiciously active on heavy loaded SMP web servers (even complained on @stable a long time ago). I'll try disabling PREEMPTION and see how it goes.

Ok, I couldn't run extensive tests because people were waiting to use the machine, so this should be considered anecdotal evidence. On a simple benchmark that repeatedly (for 1 minute) and concurrently (target=50 concurrent requests) hits a dynamic web page on a development machine (2 proc true SMP), the performance goes up from ~85 requests/sec. to ~105 requests/s by disabling PREEMPTION. This improvement looks suspiciously high to me, but I don't think I'll be going back :) pagezero is now not noticable in 'top' output.

I (or someone) will make pgzero not depend on PREEMPTION.

I found what PREEMPTION is good for -- it is to not completely lose to RELENG_4 for interrupt latency. Without PREEMPTION, hitting caps lock results in the keyboard interrupt handler busy-waiting for 3-5 msec while programming the LED. The keyboard driver has always done this stupidly, but it doesn't matter much if interrupt priorities work right. Without PREEMPTION, interrupt priorities don't work right -- higher priority interrupt handlers don't run while the very low priority keyboard interrupt handler is running. This results in several thousand

Re: What is the PREEMPTION option good for?

packets per second being dropped in my network blast tests. I'm testing bge mainly on a !SMP system with an rx load of 240kpps. My bge has an rx ring size of 1024, so drops shouldn't occur unless the latency is > 4mS, but it is apparently much higher than 3-5mS since I lose 1-3k packets when I hit caps lock. PREEMPTION is less needed for SMP systems since it takes something like 2 kernel threads busy-waiting for too long at the same time to give very large interrupt latencies.

Bruce

freebsd-arch@xxxxxxxxxxx mailing list

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

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