

Re: RH 8; intermittend very high load

Source: <http://linux.derkeiler.com/Newsgroups/linux.redhat/2004-04/0739.html>

From: Stephan Wonzak (wonzak_at_Uni-Koeln.DE)

Date: 04/28/04

Date: Wed, 28 Apr 2004 11:56:35 +0200

Tommy Reynolds wrote:

>
> *On Tue, 27 Apr 2004 21:50:37 +0200, Stephan Wonzak wrote:*
>
>> *To me this sounds like something is wedged in 'D'-state; maybe an I/O*
>> *issue. The 2.4 VM is well-known for ... 'creative' behaviour.*
>>
>> *Look at the output of "ps -ealf | grep ' D '" and see if there are any*
>> *stuck processes (during the time of high load, of course :-). Also*
>> *looking at /proc/meminfo and /proc/slabinfo might provide clues.*
>
> *Those two "/proc" entries are good places to look for memory*
> *thrashing and/or swap issues.*
>
> *I don't get the connection with a process simply waiting in*
> *un-interruptable sleep causing a process scheduling backlog. As*
> *long as a process is on the un-interruptable wait queue the kernel*
> *doesn't even see the process anymore during schedule time; it does*
> *not get woken up for signal delivery or semaphore releases.*

The OP said he had a load of 4 while the CPUs being more or less idle. Usually this is an indication that a process is stuck in kernel space somewhere. This does count towards the load average while actual CPU consumption is zero. Whether this is the right thing to do was the subject of a much discussion :-)

Anyway, in my experience this is usually caused by a process waiting on IO; stuck NFS is a prime candidate for this (just try an 'ls' on a NFS mount while the server is not responding. Result: Load 1, CPU 0.0%)

> *Can you enlightenment about this?*

I hope my point is a bit clearer now.

> *No quibble about earlier 2.4 VM being buggy,*

:~)

> *but after about 2.4.14 this smoothed out pretty well.*

linux.redhat: Re: RH 8; intermittend very high load

Yes, the overhaul of the 2.4 VM in the 2.4.10 timeframe was needed, but is still is not too good for some loads. Especially on big memory machines (...as I learned the hard way.... *sigh*) Some of the problems have been sorted out, but we still have not reached a really stable environment on our big machines.

> *HEY! ORIGINAL POSTER! What does the 'uname -r' program say?*

Yes, and the output of /proc/meminfo and /proc/slabinfo would be helpful, too!

--

C U, Stephan

"I haven't lost my mind; I know exactly where I left it."
"The meaning of my life is to make me crazy"