Re: Diagnosing performance problems

Source: http://unix.derkeiler.com/Newsgroups/comp.sys.sun.admin/2003−12/0107.html

---

From: Vikas Agnihotri (*fornewsgroups_at_vikas.mailshell.com*)
Date: 12/04/03

Date: Thu, 04 Dec 2003 09:31:47 −0500

On Thu, 4 Dec 2003 10:17:29 +0000, Ric Davis <spam@fil.ion.ucl.ac.uk>
wrote:

> Vikas Agnihotri <fornewsgroups@vikas.mailshell.com> wrote:
>
>> 1. Out of the usual 3 suspects (cpu, memory, disk), which is the
>> bottleneck here? Looking at the top output, it seems like the iowait is
>> killing me here. But the disk svc times dont seem too bad. Yes, there
>> is a lot of IO going on, but doesnt seem like the disks are holding
>> things up.
>
> There are disks nearly 100% busy, and the CPUs are spending 40% of their
> time waiting for them. As long as the busy disks aren't busy with swap,
> I vote disk, or Oracle config.

Yes, I see that they are busy, but the asvc_t is not too high. These are
standard Sun 10000RPM SCSI disks, so a svc_t of around 30−40ms is
acceptable, right?

Also, you say, "as long as the busy disks arent busy with swap...". How
can I verify if they are busy with swap or not?

Another thing I dont quite understand is the memory usage. The way Oracle
works is when it starts up, it pre−allocates the entire memory configured
in its init files (called the SGA). This is set to 250M. All the processes
connecting to the database attach to this shared memory segment. NATTCH in
'ipcs −ma' shows me about 200 processes attaching.

But when I use the process monitoring utiltiies like top, prstat, ps −o
vsz, etc, they show each process using 250M for a grand total of some
ridiculous number like 100GB or something.

Dont these utilities understand the concept of shared memory segments and
adjust the numbers accordingly?

Thanks