

Re: "Load Balancing": How Busy are the servers?

## Re: "Load Balancing": How Busy are the servers?

---

*Source:* <http://unix.derkeiler.com/Mailing-Lists/FreeBSD/questions/2006-01/msg00069.html>

---

- *From:* "Marc G. Fournier" <[scrappy@xxxxxxx](mailto:scrappy@xxxxxxx)>
  - *Date:* Mon, 2 Jan 2006 01:40:38 -0400 (AST)
- 

I just installed cacti, which seems fairly useful for 'long term views' of how a server is doing ... now I have to figure out what SNMP MIBs related to all of the "important things" :(

On Sun, 1 Jan 2006, Francisco Reyes wrote:

Marc G. Fournier writes:

For all the technology, I was kinda hoping for some 'scientific formula' :)

There are..

Now, I really hate to ask, but how do you use vmstat to get a feel for how busy the disk subsystem is?

For me, reading "Absolute BSD" by Michael Lucas was very helpfull. In particular Chapter 18, System performance.

The three columns I look at are for vmstat "r" and "b" on the left, and "fault".

"r" shows how many processes are waiting for CPU, "b" shows how many processes are waiting for disk. The fault column(s) show how badly your system is accesing swap.

## Re: "Load Balancing": How Busy are the servers?

Quick example:

```
r b w
2 5 0
1 5 0
2 4 0
2 5 0
3 4 0
1 5 0
1 5 0
```

That's from my home machine as I am doing some backups. The machine at this point is more disk bound than CPU bound with 4 to 5 disk operations at any point in time waiting for disk access

I am also falling behind in CPU, but not as bad.

On the far right of vmsat you also have CPU stats.. in my case the vmstat from the above lines showed 70% to 90% iddle which confirmed I was disk bound at that point. The fault column show you how actively you are using swap. The lines above had between 30 and 200 approximately. If you look at swapinfo and you have a large amount of swap in use and then you see a high number in vmstat for fault, the machine is short on RAM for the load you have on it.

So far in my experience nothing hurts a machine as badly as hitting swap (given that you have adequate CPU/disks). Once you start to hit swap heavily you need to do something (if you can...) such as moving services to another machine or putting in more memory.

Instead of looking for fixed number I think that relative figures are more important.. like looking at your machines at their lowest usage and then at their busiest.. or at spikes.. If at slow times of activity the machines are already falling behind on "b", "r" on vmstat.. then that machine is overloaded.

One possible quick way to start benchmarking your machines, until you can do something better is to capture snapshots of vmstat every 15 to 30 minutes and take a look.. perhaps even write a short script to summarize it. On my list of things to do.. is to do a simple setup of that nature.. just because it would be easy to setup and can provide very valuable information until you setup something more feature rich.

## Re: "Load Balancing": How Busy are the servers?

"top" in 5.X branch and up is also very usefull. If you hit "m" it shows you disk processes so you can see what programs are doing the most I/O.

One thing to watch out for in top when using 'm' is if you see all low numbers ( hit 'o' to sort and then type 'total').. is that you may have lots of programs doing little I/O, but their combined load is a problem for your disk subsystem... like having 200+ IMAP connections. Each single IMAP connection may not be doing more than a handfull of transactions per second, but all of them combined can give a disk subsystem a pretty good workout.

The load averages from 'w' are also good figures to do comparative tests. I started to wokr on a script (but needs more work) that dumps 'w' and 'vmstat' .. next have to work on parsing them and giving summaries. In particular one wants to know peak times.. since that is the best time to determine if the machine can handle it's load.. and more importantly spikes. If a machine is usually under 2.. and it spikes at 5+.. that machine is possibly able to do "normal" loads, but may not be able to handle spikes in traffic (ie a customer doing a mailing list, or a site just got press.. and there are a larger number than usual of people going to their URL).

I still thinkg I have MUCH, MUCH to learn.. but I would be glad to expand on anything mentioned above.. or anything else. Ultimately each machine/company is unique enough that absolute numbers from other people (ie what is a good value for 'r' and 'b' to be around most of the time) may be less important than learning what are the different figures for your different machines under "normal" operation.

-----

Marc G. Fournier

Hub.Org Networking Services (<http://www.hub.org>)

Email: [scrappy@xxxxxxxx](mailto:scrappy@xxxxxxxx)

Yahoo!: [yscrappy](mailto:yscrappy)

ICQ: 7615664

---

[freebsd-questions@xxxxxxxxxxxxx](mailto:freebsd-questions@xxxxxxxxxxxxx) mailing list

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

To unsubscribe, send any mail to "[freebsd-questions-unsubscribe@xxxxxxxxxxxxx](mailto:freebsd-questions-unsubscribe@xxxxxxxxxxxxx)"