Solaris Performance monitoring

Performance monitoring commands

The following are the system utilities you'll use to monitor the performance of your system:

- vmstat
- iostat
- uptime
- ps
- prstat
- sar

Finding the swap partition

Utilize the /etc/vfstab has the swap parition mounting information. Use the following to extract just the disk entry.

```
# grep "swap" /etc/vfstab | grep -v "^swap" | cut -f1
/dev/dsk/c0t0d0s1
```

or

Monitoring RAM and Swap

Examining how much memory is being used and how much is free has always been a source of confusion in the Solaris environment. Usually we can identify a memory shortfall by watching the system swap space usage. Remember that the Solaris operating system starts to use swap space when it runs out of physical memory, and we refer to this as paging. We can watch swap space usage by examining output from the vmstat command. Two indicators of a RAM shortage are the scan rate and swap device activity. Watch the 12th column (sr or scan rate) of information reported by vmstat in conjunction with I/O traffic displayed with the iostat -Pxn command. With the iostat -Pxn command, watch the swap partitions. The r/s and w/s columns might have high figures if a large amount of I/O is being generated through the file system and the page scanner needs to run to free up pages for I/O.

# vmsta	vmstat 60																
procs	memo			pa	age				C	lisk	2		fa	ults			
cpu																	
r b w	swap	free	re	mf	pi	ро	fr	de	sr	f0	s0	s1	s6	in	sy	CS	us
sy id																	
0 0 0	1374800	34881	63	13	1	0	0	0	10	0	0	0	0	306	91	147	1
1 98																	
0 0 0	1371600	33691	2 0	3	0	0	0	0	8	0	0	0	0	302	73	151	0
0 100																	
0 0 0	1371600	33691	2 0	3	0	0	0	0	14	0	0	0	0	303	72	150	0
0 100																	

With vmstat, use vmstat 60 to check memory usage every 60 seconds. Ignore the summary statistics on the first line. If page/sr exceeds 200 pages per second for an extended time, your system may be running short of physical memory. In the example I show above, the system is running between 8 and 14 pages per second, which is indicating an almost idle system.

```
# iostat -Pxn 60
                       extended device statistics
    r/s
            w/s
                   kr/s
                           kw/s wait actv wsvc_t asvc_t
                                                              γβ
                                                                  %b device
    1.0
            0.5
                    7.1
                            4.0
                                  0.0
                                        0.0
                                                0.2
                                                       11.8
                                                               0
                                                                    1 c0t0d0s0
    0.0
            0.0
                    0.0
                            0.0
                                  0.0
                                        0.0
                                                0.0
                                                       13.3
                                                               0
                                                                    0 c0t0d0s1
    0.0
            0.0
                    0.0
                            0.0
                                  0.0
                                        0.0
                                                0.0
                                                        0.0
                                                               0
                                                                    0 c0t0d0s2
                   10.4
    1.4
            0.1
                            1.5
                                  0.0
                                        0.0
                                                0.6
                                                        8.0
                                                               0
                                                                    0 c0t0d0s3
    0.8
            0.0
                    6.0
                            0.0
                                  0.0
                                        0.0
                                                0.1
                                                        2.4
                                                               0
                                                                    0 c0t0d0s4
    1.8
            0.0
                   13.5
                            0.1
                                  0.0
                                        0.0
                                                0.1
                                                        2.7
                                                               0
                                                                    0 c0t0d0s5
    0.2
            0.0
                    1.5
                            0.1
                                  0.0
                                        0.0
                                                0.1
                                                        4.2
                                                               0
                                                                    0 c0t0d0s6
   28.1
            2.4
                  204.4
                           24.2
                                  0.1
                                        0.2
                                                2.0
                                                        5.9
                                                               0
                                                                    5 c0t0d0s7
   16.6
            0.4
                  116.4
                            3.9
                                  0.0
                                        0.1
                                                0.7
                                                        3.5
                                                               0
                                                                    2 c0t9d0s0
    0.0
            0.0
                    0.0
                            0.0
                                  0.0
                                        0.0
                                                0.0
                                                        0.0
                                                               0
                                                                    0 c0t9d0s2
    0.0
           10.9
                    0.0
                          349.8
                                  0.0
                                        0.0
                                                0.0
                                                        3.9
                                                               0
                                                                    4 rmt/0
    0.0
            0.0
                    0.0
                             0.0
                                  0.0
                                        0.0
                                                0.0
                                                        2.5
                                                               0
                                                                    0
so6:vold(pid565)
                       extended device statistics
    r/s
                   kr/s
                           kw/s wait actv wsvc_t asvc_t
            w/s
                                                              8₩
                                                                  %b device
    0.0
            0.3
                    0.0
                            2.2
                                  0.0
                                        0.1
                                                0.0
                                                      185.8
                                                               0
                                                                    1 c0t0d0s0
    0.0
            0.0
                    0.0
                            0.0
                                  0.0
                                        0.0
                                                0.0
                                                        0.0
                                                               0
                                                                    0 c0t0d0s1
    0.0
            0.0
                    0.0
                            0.0
                                  0.0
                                        0.0
                                                0.0
                                                                    0 c0t0d0s2
                                                        0.0
                                                               0
    0.0
            0.7
                    0.3
                            5.3
                                  0.0
                                        0.1
                                                0.0
                                                     197.7
                                                               0
                                                                    1 c0t0d0s3
    0.0
            0.0
                    0.0
                            0.0
                                  0.0
                                        0.0
                                                0.0
                                                        0.0
                                                               0
                                                                    0 c0t0d0s4
    0.0
            0.0
                    0.0
                            0.0
                                  0.0
                                        0.0
                                                0.0
                                                        0.0
                                                               0
                                                                    0 c0t0d0s5
    0.0
            0.0
                    0.0
                            0.0
                                  0.0
                                        0.0
                                                0.0
                                                        0.0
                                                               0
                                                                    0 c0t0d0s6
   34.1
            2.2
                  505.9
                           16.4
                                  0.0
                                        0.2
                                                0.0
                                                        6.5
                                                               0
                                                                    7 c0t0d0s7
    2.6
          132.7
                   20.8 1276.2
                                  0.0
                                        2.9
                                                0.0
                                                       21.3
                                                               0 100 c0t9d0s0
```

0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0	0 c0t9d0s2
0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0	0 rmt/0
0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0	0
so6:vold(p	pid565)								

Looking at the output from iostat -Pxn, watch the row for the c0t0d0s1 (swap partition). If there are I/Os queued for the swap device (svc_t column), application paging is occurring. If there is significant, heavy I/O to the swap device, you might be experiencing a RAM shortage, and a RAM upgrade may be in order. Sometimes you can even hear the disk paging. How do you determine heavy swap I/O? As stated in the preceding section, any disk that is consistently more than 10% busy with svc_t above 30ms is getting heavy I/O. Also, compare the figures you have on a busy system against your baseline figures to determine if the disks are experiencing higher-than-average activity.

The system above is showing a some load against swap space but not significant for this system.

If there is insuffient physical memory, the system will be busy paging to the swap device and it will be unable to keep up with requests. This state is refered to as *thrashing*. This state is characterized by heavy I/O on the swap device and sluggish overall performance. In this state, the page scanner (the part of the kernel that handles swapping) can use up to 80% of CPU.

vmstat command

vmstat}} provides a convenient summary of system activity similar to command. When you run vmstat for the first time, the displayed result represents a summary of information since boot time. To obtain useful real-time statistics, run vmstat with a time step as follows:

```
vmstat 30
```

This tells vmstat to run every 30 seconds and to display the results on the screen as follows until you type Ctrl+C to interrupt the command:

# vmstat 30						
procs memo	ory	page	e	disk	faults	
cpu						
r b w swap	free re	mf pi p	o fr de	sr f0 s0 s1 s	6 in sy cs us	
sy id						
0 0 0 1374800	348800 3	13 1	0 0 0	0 0 0 0	0 306 91 147 1	
1 98						
0 0 0 1371744	337072 0	0 0	0 0 0	0 0 0 0	0 303 61 149 0	
1 99						
0 0 0 1371744	337072 2	9 0	0 0 0	0 0 0 0	0 303 68 157 0	

0 100																	
0 0 0 1	371744	337072	0	0	0	0	0	0	0	0	0	0	0	303	58	149	0
0 100																	
0 0 0 1	371744	337072	0	0	0	0	0	0	0	0	0	0	0	303	58	148	0
1 99																	
0 0 0 1	371744	337072	2	9	0	0	0	0	0	0	0	0	0	303	102	150	0
0 99																	
0 0 0 1	371744	337072	0	0	0	0	0	0	0	0	0	0	0	302	60	153	0
0 100																	

vmstat fields

Field	Description
kthr/r	Run queue length.
kthr/b	Kernel threads blocked while waiting for I/O.
kthr/w	Idle processes that have been swapped.
memory/swap	Free, unreserved swap space (KB).
memory/free	Free memory (KB).
page/re	Pages reclaimed from the free list.
page/mf	Minor faults (page in memory but not mapped). If the page is still in memory, a minor fault remaps the page.
page/pi	Paged in from swap (KB/s). (When a page is brought back from the swap device, the process will stop execution and wait. This might affect performance.)
page/po	Paged out to swap (KB/s). The page has been written and freed.
page/fr	Freed or destroyed (KB/s). This column reports the activity of the page scanner.
page/de	Anticipated short-term memory shortfall (KB).
page/sr	Scan rate (pages). This number is not reported as a "rate" but as a total number of pages scanned.
disk/s#	Disk activity for disk # (disk operations per second).
faults/in	Interrupts per second.
faults/sy	System calls per second.
faults/cs	Context switches per second.
cpu/us	User CPU time (%).
cpu/sy	System (kernel) CPU time (%).
cpu/id	Idle + I/O wait CPU time (%).

iostat command

iostat output

Column	Description
r/s	Reads per second
w/s	Writes per second
Kr/s	Kilobytes read per second
Kw/s	Kilobytes written per second
Wait	Average number of transactions waiting for service (queue length)
actv	Average number of transactions actively being serviced (removed from the queue but not yet completed)
svc_t	Average service time, in milliseconds
%w	Percent of time there are transactions waiting for service (queue nonempty)
%b	Percent of time the disk is busy (transactions in progress)
device	The device name of the disk

extended device statistics r/s w/s kr/s kw/s wait actv wsvc_t asvc_t %w %b device 33.4 3.1 242.9 29.8 0.1 0.2 1.7 6.0 0 5 c0t000 0.1 0.0 0.2 0.0 0.0 0.3 3.0 0 0 c0t600 16.6 0.4 116.3 4.1 0.0 0.1 0.7 3.5 0 2 c0t900 0.0 10.9 0.0 349.7 0.0 0.0 0.0 3.9 0 4 rmt/0 0.0 0.0 0.0 0.0 0.0 0.0 2.5 0 0 sof:vold(pid565) extended device statistics 1 c0t040 1.6 0.4 22.8 2.2 0.0 0.0 5.4 1 c0t040 0.0 0.0 0.0 0.0 0.0 0.0 0.0 0.0 0.0 0.0 0.0 0.0 0.0 0.0 0.0 0.0 <	# iostat	t -xn 5									
r/s w/s kr/s kw/s wait actv wsvc_t asvc_t %w %b device 33.4 3.1 242.9 29.8 0.1 0.2 1.7 6.0 0 5 c0t0d0 0.1 0.0 0.2 0.0 0.0 0.0 0.3 3.0 0 0 c0t6d0 16.6 0.4 116.3 4.1 0.0 0.1 0.7 3.5 0 2 c0t9d0 0.0 10.9 0.0 349.7 0.0 0.0 0.0 3.9 0 4 rmt/0 0.0			ex	tended	devi	ce sta	atistic	5			
33.4 3.1 242.9 29.8 0.1 0.2 1.7 6.0 0 5 c0t0d0 0.1 0.0 0.2 0.0 0.0 0.0 0.3 3.0 0 0 c0t6d0 16.6 0.4 116.3 4.1 0.0 0.1 0.7 3.5 0 2 c0t9d0 0.0 10.9 0.0 349.7 0.0 0.0 0.0 3.9 0 4 rmt/0 0.0 0.0 0.0 0.0 0.0 0.0 0.0 2.5 0 0 so6:vold(pid565) extended device statistics r/s w/s kr/s kw/s wait actv wsvc_t asvc_t %w %b device 1.6 0.4 22.8 2.2 0.0 0.0 0.0 5.4 0 1 c0t0d0 0.0 0.0 0.0 0.0 0.0 0.0 0.0 0.0 0.0 0 0 c0t6d0 0.0 0.0 0.0 0.0 0.0 0.0 0.0 0.0 0.0 0 0 c0t6d0 0.0 0.0 0.0 0.0 0.0 0.0 0.0 0.0 0.0 0.0	r/s	w/s	kr/s	kw/s	wait	actv	wsvc_t	asvc_t	%₩	%b	device
0.1 0.0 0.2 0.0 0.0 0.0 0.3 3.0 0 0 c0t6d0 16.6 0.4 116.3 4.1 0.0 0.1 0.7 3.5 0 2 c0t9d0 0.0 10.9 0.0 349.7 0.0 0.0 0.0 3.9 0 4 rmt/0 0.0 0.0 0.0 0.0 0.0 0.0 0.0 2.5 0 0 so6:vold(pid565) r/s w/s kr/s kw/s wait actv wsvc_t asvc_t %w %b device 1.6 0.4 22.8 2.2 0.0 0.0 0.0 5.4 0 1 c0t0d0 0.0 0.0 0.0 0.0 0.0 0.0 0.0 0.0 0 0 0 c0t6d0 0.0 0.0 0.0 0.0 0.0 0.0 0.0 0.0 0 0 0 c0t6d0 0.0 0.0 0.0 0.0 0.0 0.0 0.0 0.0 0.0 0 0 c0t6d0 0.0 0.0 0.0 0.0 0.0 0.0 0.0 0.0 0.0 0 0 c0t6d0 0.0 0.0 0.0 0.0 0.0 0.0 0.0 0.0 0.0 0 0 0 c0t6d0 0.0 0.0 0.0 0.0 0.0 0.0 0.0 0.0 0.0 0.0	33.4	3.1	242.9	29.8	0.1	0.2	1.7	6.0	0	5	c0t0d0
16.6 0.4 116.3 4.1 0.0 0.1 0.7 3.5 0 2 c0t9d0 0.0 10.9 0.0 349.7 0.0 0.0 0.0 3.9 0 4 rmt/0 0.0 0.0 0.0 0.0 0.0 0.0 0.0 2.5 0 0 so6:vold(pid565) extended device statistics r/s w/s kr/s kw/s wait actv wsvc_t asvc_t %w %b device 1.6 0.4 22.8 2.2 0.0 0.0 0.0 5.4 0 1 c0t0d0 0.0 0.0 0.0 0.0 0.0 0.0 0.0 0.0 0.0 0 0 c0t6d0 0.0 0.0 0.0 0.0 0.0 0.0 0.0 0.0 0.0 0 0 c0t6d0 0.0 0.0 0.0 0.0 0.0 0.0 0.0 0.0 0.0 0 0 c0t6d0 0.0 0.0 0.0 0.0 0.0 0.0 0.0 0.0 0.0 0 0 c0t6d0 0.0 0.0 0.0 0.0 0.0 0.0 0.0 0.0 0.0 0 0 c0t6d0 0.0 0.0 0.0 0.0 0.0 0.0 0.0 0.0 0.0 0 0 c0t6d0 0.0 0.0 0.0 0.0 0.0 0.0 0.0 0.0 0.0 0.	0.1	0.0	0.2	0.0	0.0	0.0	0.3	3.0	0	0	c0t6d0
0.0 10.9 0.0 349.7 0.0 0.0 0.0 3.9 0 4 rmt/0 0.0 0.0 0.0 0.0 0.0 0.0 0.0 0.0 2.5 0 0 so6:vold(pid565) r/s w/s kr/s kw/s wait actv wsvc_t asvc_t %w %b device 1.6 0.4 22.8 2.2 0.0 0.0 0.0 5.4 0 1 c0t0d0 0.0 0.0 0.0 0.0 0.0 0.0 0.0 0.0 0.0 0 0 0 c0t6d0 0.0 0.0 0.0 0.0 0.0 0.0 0.0 0.0 0.0 0 0 0 c0t9d0 0.0 0.0 0.0 0.0 0.0 0.0 0.0 0.0 0.0 0 0 0 rmt/0 0.0 0.0 0.0 0.0 0.0 0.0 0.0 0.0 0.0 0.0	16.6	0.4	116.3	4.1	0.0	0.1	0.7	3.5	0	2	c0t9d0
0.0 0.0 0.0 0.0 0.0 0.0 0.0 2.5 0 0 so6:vold(pid565) extended device statistics r/s w/s kr/s kw/s wait actv wsvc_t asvc_t %w %b device 1.6 0.4 22.8 2.2 0.0 0.0 0.0 5.4 0 1 c0t0d0 0.0 0.0 0.0 0.0 0.0 0.0 0.0 0.0 0.0 0 0 0 c0t6d0 0.0 0.0 0.0 0.0 0.0 0.0 0.0 0.0 0.0 0 0 c0t6d0 0.0 0.0 0.0 0.0 0.0 0.0 0.0 0.0 0.0 0 0 c0t6d0 0.0 0.0 0.0 0.0 0.0 0.0 0.0 0.0 0.0 0 0 c0t6d0 so6:vold(pid565) extended device statistics r/s w/s kr/s kw/s wait actv wsvc_t asvc_t %w %b device 0.0 1.2 0.0 8.4 0.0 0.0 0.0 0.0 0.0 0.0 0 0 c0t6d0	0.0	10.9	0.0	349.7	0.0	0.0	0.0	3.9	0	4	rmt/0
so6:vold(pid565) extended device statistics r/s kr/s kw/s wait actv wsvc_t asvc_t %w %b device 1.6 0.4 22.8 2.2 0.0 0.0 0.0 1 c0t0d0 0.0 0.0 0.0 0.0 0.0 0.0 0.0 0.0 0 0 c0t0d0 0 0.0 0.0 0.0 0.0 0.0 0.0 0.0 0 0 c0t0d0 0 c0t0d0 0 c0t0d0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	2.5	0	0	
r/s w/s kr/s kw/s wait actv wsvc_t asvc_t %w %b device 1.6 0.4 22.8 2.2 0.0 0.0 0.0 5.4 0 1 c0t0dd 0.0 <	so6:vol	d(pid565)								
r/s w/s kr/s kw/s wait actv wsvc_t asvc_t %w %b device 1.6 0.4 22.8 2.2 0.0 0.0 0.0 5.4 0 1 c0t0d0 0.0 0.0 0.0 0.0 0.0 0.0 0.0 0.0 0 0 c0t0d0 0.0 0.0 0.0 0.0 0.0 0.0 0.0 0.0 0 c0t6d0 0.0 0.0 0.0 0.0 0.0 0.0 0.0 0 c0t6d0 0.0 0.0 0.0 0.0 0.0 0.0 0.0 0 c0t6d0 0.0 0.0 0.0 0.0 0.0 0.0 0.0 0 0 c0t9d0 0.0 0.0 0.0 0.0 0.0 0.0 0 0 c0t9d0 softward 0.0 0.0 0.0 0.0 0.0 0 0 c c softward 0.0 0.0 0.0 0.0 0.0 0 c <td></td> <td></td> <td>ex</td> <td>tended</td> <td>devi</td> <td>ce sta</td> <td>atistic</td> <td>5</td> <td></td> <td></td> <td></td>			ex	tended	devi	ce sta	atistic	5			
1.6 0.4 22.8 2.2 0.0 0.0 0.0 5.4 0 1 c0t0d0 0.0 <	r/s	w/s	kr/s	kw/s	wait	actv	wsvc_t	asvc_t	₩	%b	device
0.0 0.0 0.0 0.0 0.0 0.0 0.0 0.0 0.0 0 0 0 c0t6d0 0.0 0.0 0.0 0.0 0.0 0.0 0.0 0.0 0.0 0 0 c0t9d0 0.0 0.0 0.0 0.0 0.0 0.0 0.0 0.0 0 0 0 rmt/0 0.0 0.0 0.0 0.0 0.0 0.0 0.0 0.0 0 0 0 rmt/0 so6:vold(pid565) r/s w/s kr/s kw/s wait actv wsvc_t asvc_t %w %b device 0.0 1.2 0.0 8.4 0.0 0.0 0.0 0.0 6.1 0 1 c0t0d0 0.0 0.0 0.0 0.0 0.0 0.0 0.0 0.0 0 0 0 c0t6d0	1.6	0.4	22.8	2.2	0.0	0.0	0.0	5.4	0	1	c0t0d0
0.0 0.0 0.0 0.0 0.0 0.0 0.0 0.0 0.0 0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0	0	c0t6d0
0.0 0.0 0.0 0.0 0.0 0.0 0.0 0.0 0.0 0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0	0	c0t9d0
0.0 0.0 0.0 0.0 0.0 0.0 0.0 0.0 0.0 0 0 0 so6:vold(pid565) r/s w/s kr/s kw/s wait actv wsvc_t asvc_t %w %b device 0.0 1.2 0.0 8.4 0.0 0.0 0.0 6.1 0 1 c0t0d0 0.0 0.0 0.0 0.0 0.0 0.0 0.0 0.0 0 0 0 c0t6d0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0	0	rmt/0
so6:vold(pid565) extended device statistics r/s w/s kr/s kw/s wait actv wsvc_t asvc_t %w %b device 0.0 1.2 0.0 8.4 0.0 0.0 0.0 6.1 0 1 c0t0d0 0.0 0.0 0.0 0.0 0.0 0.0 0.0 0.0 0 0 c0t6d0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0	0	
extended device statistics r/s w/s kr/s kw/s wait actv wsvc_t asvc_t %w %b device 0.0 1.2 0.0 8.4 0.0 0.0 0.0 6.1 0 1 c0t0d0 0.0 0.0 0.0 0.0 0.0 0.0 0.0 0.0 0 0 c0t6d0	so6:vol	d(pid565)								
r/s w/s kr/s kw/s wait actv wsvc_t asvc_t %w %b device 0.0 1.2 0.0 8.4 0.0 0.0 0.0 6.1 0 1 c0t0d0 0.0 0.0 0.0 0.0 0.0 0.0 0.0 0.0 0 c0t6d0			ex	tended	devi	ce sta	atistic	S			
0.0 1.2 0.0 8.4 0.0 0.0 0.0 6.1 0 1 c0t0d0 0.0 0.0 0.0 0.0 0.0 0.0 0.0 0.0 0.0 0.0 0.0 0.0 0.0 0.0 0 0 c0t6d0	r/s	w/s	kr/s	kw/s	wait	actv	wsvc_t	asvc_t	8₩	۶b	device
0.0 0.0 0.0 0.0 0.0 0.0 0.0 0.0 0 c0t6d0	0.0	1.2	0.0	8.4	0.0	0.0	0.0	6.1	0	1	c0t0d0
	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0	0	c0t6d0
0.0 0.0 0.0 0.0 0.0 0.0 0.0 0.0 0 c0t9d0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0	0	c0t9d0
0.0 0.0 0.0 0.0 0.0 0.0 0.0 0.0 0.0 0 rmt/0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0	0	rmt/0

If you are seeing svc_t (service time) values of more than 30ms on disks that are in use (more than 10% busy), the end user will see noticeably sluggish performance.

If a disk is more than 60% busy over sustained periods of time, this can also indicate overuse of that resource.

iostat sometimes reports excessive svc_t (service time) readings for disks that are inactive. This is because fsflush (a kernel activity) tries to keep the data in memory and on the disk up-to-date. Because many writes are specified over a very short period of time to random parts of the disk, a queue forms briefly, and the average service time goes up. svc_t should only be taken seriously on a disk that is showing 5% or more activity.

The "wait" time reported by iostat refers to time spent by a process while waiting for block device (such as disk) I/O to finish. If iostat consistently reports %w > 5, the disk subsystem is too busy.

#	ios	stat i	10														
	tt	ty		sd0			sd6			sd8			st5				
C]	pu																
1	cin	tout	kps	tps	serv	us	sy										
W	t id	d															
	2	100	273	36	8	0	0	3	120	17	4	350	11	4	1	2	
5	92																
	1	144	111	14	5	0	0	0	0	0	0	0	0	0	8	7	
3	82																
	0	8	5	1	7	0	0	0	0	0	0	0	0	0	2	1	
0	96																
	0	8	34	5	69	0	0	0	0	0	11	0	0	0	0	1	
0	98																
	0	8	6	1	6	0	0	0	0	0	0	0	0	0	0	1	
0	98																
	0	8	5	1	6	0	0	0	0	0	0	0	0	0	0	1	
0	98																
	0	8	14	2	25	0	0	0	0	0	0	0	0	0	0	2	
1	97																

uptime

The uptime command can be used to give you a rough estimate of the system load. The uptime command prints the current time, the length of time that the system has been up, and the average number of jobs in the run queue over the last 1, 5, and 15 minutes.

uptime
3:01pm up 65 day(s), 21:04, 21 users, load average: 0.08, 0.07, 0.09

Look at the load average numbers. The load average is the sum of the run queue length and the number of jobs currently running on CPUs. In short, it's a rough estimate of CPU usage. Notice the figures, showing averages over the last 1, 5, and 15 minutes.

What is a high load average?

It depends on your system. If you've been keeping an eye on the load average, you'll know what is a good average and what is a bad average based on the history of the system. Normally, I would say a load average of 3 or less is good, but I've seen systems with a load average of 5 in which performance is still good. Different system configurations behave differently under the same load averages.

How to Display Virtual Memory Statistics (vmstat)

The following example shows the vmstat display of statistics gathered at five-second intervals.

\$ vmst	\$ vmstat 5																	
procs	me	emory				page					dis	sk		fau	ults		срі	1
r b w	swap	free	re	mf	pi	ро	fr	de	sr	f0	s3			in	sy	CS	us	sy
id																		
0 0 8	28312	668	0	9	2	0	1	0	0	0	1	0	0	10	61	82	1	2
97																		
0 0 3	31940	248	0	10	20	0	26	0	27	0	4	0	0	53	189	191	6	6
88																		
0 0 3	32080	288	3	19	49	6	26	0	15	0	9	0	0	75	415	277	6	15
79																		
0 0 3	32080	256	0	26	20	6	21	0	12	1	6	0	0	163	110	138	1	3
96																		
0 1 3	32060	256	3	45	52	28	61	0	27	5	12	0	0	195	191	223	7	11
82																		
0 0 3	32056	260	0	1	0	0	0	0	0	0	0	0	0	4	52	84	0	1
99																		

Category	Field Name	Description
procs		Reports the following states:
	r	The number of kernel threads in the dispatch queue
	b	Blocked kernel threads waiting for resources

	W	Swapped out LWPs waiting for processing resources to finish
memory		Reports on usage of real and virtual memory:
	swap	Available swap space
	free	Size of the free list
page		Reports on page faults and paging activity, in units per second:
	re	Pages reclaimed
	mf	Minor and major faults
	pi	Kbytes paged in
	ро	Kbytes paged out
	fr	Kbytes freed
	de	Anticipated memory needed by recently swapped-in processes
	sr	Pages scanned by page daemon (not currently in use). If sr does not equal zero, the page daemon has been running.
disk		Reports the number of disk operations per second, showing data on up to four disks
faults		Reports the trap/interrupt rates (per second):
	in	Interrupts per second
	sy	System calls per second
	cs	CPU context switch rate
cpu		Reports on the use of CPU time:
	us	User time
	sy	System time
	id	Idle time

How to Display System Event Information

Run vmstat -s to show the total of various system events that have taken place since the system was last booted.

```
0 swap ins
0 swap outs
0 pages swapped in
0 pages swapped out
409376480 total address trans. faults taken
3075036 page ins
2601555 page outs
```

```
3812452 pages paged in
  6525552 pages paged out
11007609 total reclaims
 10927650 reclaims from free list
        0 micro (hat) faults
409376480 minor (as) faults
  2957386 major faults
102738273 copy-on-write faults
 61711047 zero fill page faults
1002562077 pages examined by the clock daemon
     7881 revolutions of the clock hand
16716370 pages freed by the clock daemon
  4999048 forks
  1138206 vforks
  5747009 execs
741660225 cpu context switches
736047593 device interrupts
528054538 traps
2496638575 system calls
430283487 total name lookups (cache hits 95%)
    81727 toolong
 10484677 user cpu
  9528364 system cpu
443762786 idle cpu
 16281790 wait cpu
```

How to Display Swapping Statistics

Run vmstat -S to show swapping statistics.

procs	memory				pa	age				(dis	ζ		f			
cpu																	
r b w	swap	free	si	SO	pi	ро	fr	de	sr	m1	mЗ	m4	m5	in	sy	CS	us
sy id																	
0 0 0	8512	888	0	0	12	21	55	0	417	1	0	0	0	206	1040	308	2
2 96																	

si Average number of LWPs swapped in per secondso Number of whole processes swapped out

How to Display Disk Utilization Information (iostat)

You can display disk activity information by using the iostat command with a time interval. The following example shows disk statistics gathered every five seconds.

# iostat 5																	
		tty			md1			md3			md4			md5			
C]	сри																
1	cin t	cout	kps	tps	serv	kps	tps	serv	kps	tps	serv	kps	tps	serv	us	sy	
w	t id																
	0	2	10	1	28	2	0	22	0	0	0	1	0	10	2	2	
3	92																
	0	47	58	7	39	16	2	34	0	0	0	0	0	0	0	2	
1	19 78																
	0	16	0	0	0	0	0	0	0	0	0	0	0	0	0	1	
0	98																
	0	16	0	0	0	0	0	0	0	0	0	0	0	0	0	0	
1	99																
	0	16	2	0	22	0	0	0	0	0	0	0	0	0	2	3	
1	95																
	0	24	0	0	0	0	0	0	0	0	0	0	0	0	0	1	
1	98																
1	99 0 95 0 98	16 24	2 0	0	22 0	0	0	0	0	0	0	0	0	0	2 0	3	

For Each	Field Name	Description
Terminal		
	tin	Number of characters in the terminal input queue
	tout	Number of characters in the terminal output queue
Disk		
	bps	Blocks per second
	tps	Transactions per second
	serv	Average service time, in milliseconds
CPU		
	us	In user mode
	sy	In system mode
	wt	Waiting for I/O
	id	Idle

How to Display Extended Disk Statistics

Run iostat -xtc to get extended disk statistics. This command displays a line of output for each disk.

cpu													
device	r/s	w/s	kr/s	kw/s	wait	actv	svc_t	%₩	%b	tin	tout	us	sy
wt id													
md1	0.4	0.9	3.6	6.9	0.0	0.0	27.7	1	1	0	2	2	2
3 92													
md3	0.1	0.2	1.0	1.3	0.0	0.0	21.7	0	0				
md4	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0	0				
md5	0.0	0.0	0.7	0.0	0.0	0.0	9.9	0	0				
md8	0.8	0.3	6.7	14.2	0.0	0.0	13.1	0	1				
md10	0.2	0.9	1.8	6.8	0.0	0.0	15.5	0	1				
md11	0.2	0.9	1.8	6.8	0.0	0.0	14.8	0	1				
md30	0.0	0.2	0.5	1.3	0.0	0.0	11.4	0	0				
md31	0.0	0.2	0.5	1.3	0.0	0.0	10.2	0	0				
md40	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0	0				
md41	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0	0				
md50	0.0	0.0	0.4	0.0	0.0	0.0	9.4	0	0				
md51	0.0	0.0	0.4	0.0	0.0	0.0	7.3	0	0				
md80	0.4	0.3	3.3	14.2	0.0	0.0	10.3	0	0				
md81	0.4	0.3	3.3	14.2	0.0	0.0	11.7	0	1				
sd0	0.6	2.1	6.0	22.8	0.0	0.0	16.3	0	3				
sd1	0.6	2.1	6.0	22.8	0.0	0.0	15.2	0	2				

Field Name	Description
r/s	Reads per second
w/s	Writes per second
Kr/s	Kbytes read per second
Kw/s	Kbytes written per second
wait	Average number of transactions waiting for service (queue length)
actv	Average number of transactions actively being serviced
svc_t	Average service time, in milliseconds
%w	Percentage of time the queue is not empty
%b	Percentage of time the disk is busy

How to Check CPU Utilization (sar)

Display CPU utilization with the sar -u command. (The sar command without any options is equivalent to sar -u.) At any given moment, the processor is either busy or idle. When busy, the processor is in either user or system mode. When idle, the processor is either waiting for I/O completion or "sitting still" with no work to do.

Measure CPU utilization during 5 secs one time.

sar -u 5 1

Measure CPU utilization during 60 secs 1440 times and write result in file sar.log.

sar -u -o sar.log 60 1440

To later review disk and tape activity from that period:

sar -d -f sar.log

Field Name	Description
%sys	Lists the percentage of time that the processor is in system mode
%user	Lists the percentage of time that the processor is in user mode
%wio	Lists the percentage of time the processor is idle and waiting for I/O completion
%idle	Lists the percentage of time the processor is idle and is not waiting for I/O

NOTE: A high %wio generally means a disk slowdown has occurred.

Copyright 2005 - J. Michael McGarrah

last edited 2005-09-27 12:09:19 by MichaelMcGarrah