Subject: Where can we download the "vzstat" source code?
Posted by RapidVPS on Mon, 17 Apr 2006 05:46:33 GMT
View Forum Message <> Reply to Message

Hi, I downloaded "vzctl-3.0.0-4.tar.bz2" but could not find any mention of vzstat. Where is the source code to this app?

Rick

Subject: Re: Where can we download the "vzstat" source code?
Posted by dev on Mon, 17 Apr 2006 07:55:42 GMT
View Forum Message <> Reply to Message

vzstat is not available in OVZ.
what did you want to do with it? Maybe I can suggest some workaround?

Subject: Re: Where can we download the "vzstat" source code?
Posted by RapidVPS on Thu, 20 Apr 2006 03:26:48 GMT
View Forum Message <> Reply to Message

I am building some vps realtime CPU logging/monitoring tools. How can I determine the cpu usage of a given vps? the CPU value that is in vzstat.

Subject: Re: Where can we download the "vzstat" source code?
Posted by dev on Thu, 20 Apr 2006 06:03:28 GMT
View Forum Message <> Reply to Message

You can use /proc/vz/vestat file:

```
[root@tst ~]# cat /proc/vz/vestat
Version: 2.2
     VEID      user      nice     system     uptime                idle                strv             uptime
     used             maxlat             totlat   numsched
      100         0         0       5063   40598698       70123626936598                   0
35069066535737          8612439976              2314          737714316      53144
```

in this file for each VPS you have:
1. "user nice system uptime" in jiffies. i.e. divide by jiffies_per_second to get seconds (see below).
2. "idle strv uptime used" in cycles. much more accurate. seconds = used / cycles_per_jiffy / jiffies_per_seconds.
3. values of cycles_per_jiffy and jiffies_per_seconds can be read from /proc/vz/stats

BTW, while looking into this I found a small misprint in /proc/vz/vestat output code, so thank you very much

```
--- ./kernel/vecalls.c.usertm 2006-04-07 19:04:46.000000000 +0400
+++ ./kernel/vecalls.c 2006-04-20 10:05:24.000000000 +0400
@@ -2875,7 +2875,7 @@ static int vestat_seq_show(struct seq_fi
  idle_time = strv_time = used = 0;

  for (cpu = 0; cpu < NR_CPUS; cpu++) {
-  user_ve += VE_CPU_STATS(ve, cpu)->nice;
+  user_ve += VE_CPU_STATS(ve, cpu)->user;
   nice_ve += VE_CPU_STATS(ve, cpu)->nice;
   system_ve += VE_CPU_STATS(ve, cpu)->system;
   used += VE_CPU_STATS(ve, cpu)->used_time;
```

Subject: Re: Where can we download the "vzstat" source code?
Posted by RapidVPS on Fri, 21 Apr 2006 00:36:39 GMT
View Forum Message <> Reply to Message

Hi, I could not locate /proc/vz/stats. Can you verify?

```
#ls -al /proc/vz/
total 0
dr-xr-xr-x     3 root root 0 Apr 20 20:35 .
dr-xr-xr-x  2756 root root 0 Feb 24 13:59 ..
-r--------     1 root root 0 Apr 20 20:36 devperms
-r--------     1 root root 0 Apr 20 20:35 veinfo
-r--------     1 root root 0 Apr 20 20:36 veip
-r--------     1 root root 0 Apr 20 20:36 vestat
dr-x------     2 root root 0 Apr 20 20:36 vzaquota
-r--------     1 root root 0 Apr 20 20:36 vzquota
#uname -a
Linux sodium.rapidvps.net 2.6.8-022stab070.1-enterprise #1 SMP Mon Feb 20 19:31:28 MSK
2006 i686 i686 i386 GNU/Linux
```

Subject: Re: Where can we download the "vzstat" source code?
Posted by dev on Fri, 21 Apr 2006 08:43:53 GMT
View Forum Message <> Reply to Message

Oh, it's my poor memory.
cycles_per_jiffy = frequency_of_your_cpu/jiffies_per_second
jiffies_per_second = 1000

frequency_of_your_cpu can be read from /proc/meminfo, as:
cpu MHz : <frequency_of_your_cpu>

Subject: Re: Where can we download the "vzstat" source code?
Posted by KoolK on Sun, 28 Oct 2007 19:05:24 GMT
View Forum Message <> Reply to Message

Hi,

I am building tool for CPU, memory usage for VPS in OpenVZ. Can anybody help me and explain what do fields like uptime, used, maxlat, totlat, numsched in file /proc/vz/vestat signify?

Thanks in advance

Subject: Re: Where can we download the "vzstat" source code?
Posted by dev on Mon, 29 Oct 2007 08:31:17 GMT
View Forum Message <> Reply to Message

it's quite easy:

[root@dhcp0-95 ~]# cat /proc/vz/vestat
Version: 2.2
VEID user nice system uptime idle
strv uptime used maxlat totlat numsched
100 1634 0 1712 343011435
1277332114762816 0 638679179771038 16971458021 4921
5633592426 145722
1 157635 0 29103 488618806
1818929923633445 0 909796661898187 512196984005 4858
50362877628 2524559
101 759 0 1844 488627361
1819537209544684 0 909812591751777 33663331702 4529
10257230389 713559
103 765 0 1468 488628230
1819533866071061 0 909814208459884 28408677122 5117
11686660006 712016
105 4522 0 5088 488628270
1819572969105220 0 909814283772065 22621645151 4949
7557632985 125350

- column 1 "VEID": VE id
- column 2 "user", 3 "nice", 4 "system", 5 "uptime" - corresponding std user/nice/system/uptime values in jiffies to standalone linux /proc/stat. Note, there is no "idle" time here, since it can't be calculated this way.

the next group comes in cycles units:

- column 6 "idle" - idle time, 7 "strv" - not used, 8 "uptime" - uptime in cycles, 9 "used" - used time by VE across all CPUs in cycles

the next group is scheduling latency statistics in cycles:
- column 10 "maxlat" - max latency in cycles meaning how long VE process has to wait before it actually got CPU time.
- column 11/12 "totlat/numsched", i.e. divide 11 on 12 to get average scheduling latency.

you can also read:

[root@dhcp0-95 ~]# cat /proc/vz/stats
Version: 2.6
cycles_per_jiffy: 1860831
jiffies_per_second: 1000

to understand how many cycles and jiffies are per second.

Subject: Re: Where can we download the "vzstat" source code?
Posted by Valmont on Mon, 29 Oct 2007 11:07:07 GMT
View Forum Message <> Reply to Message

Impresses. It is easy to trace cpu using

Subject: Re: Where can we download the "vzstat" source code?
Posted by KoolK on Sat, 03 Nov 2007 06:28:15 GMT
View Forum Message <> Reply to Message

Hi Dev,

Thank you very much for help...I have few doubts...

(1) Do the two uptime fields in /proc/vz/vestat denotes same (time since last boot) or different values(if so which values) because I tried to convert each field in file to sec by formulas give in post above..

first uptime in file:

uptime in sec= uptime_in_jiffies/jiffies_per_sec;

second uptime in file:
uptime in sec= uptime_in_cycle/cpu_freq

(cpu_freq is in MHZ..so no need to divide by jiffies_per_sec*jiffies_per_sec...I tried with this also still both values do not match)

the values are not same.

(2) I ran several infinite processes in two different VEs each...still the values of Maxlat, Totlat, Numsched remain 0. Am I doing any mistake?

thanks in advance...

Subject: Re: Where can we download the "vzstat" source code?
Posted by rickb on Sat, 03 Nov 2007 09:45:23 GMT
View Forum Message <> Reply to Message

Hi, I thought these fields were only avaialble in the virtuozzo kernel.

virtuozzo 3.x 2.6.9
Quote:
[root@yttrium ~]# cat /proc/vz/vestat
Version: 2.2

VEID user nice system uptime idle strv uptime used maxlat totlat numsched
1030 34518869 562190 10999333 783842709 9045444692949996 0 2402735226105746 141451539991400 31786364 249070126464914 69045720
2086 29900400 0 4506271 1636856393 19450221162219816 0 5017502202734454 108719002917852 2942466202 345587087247930 57824594
1024 21191653 0 23079329 2060298740 24614175256611440 0 6315492003670710 141218633354376 55571736 187231807734994 75413693
1033 22317199 0 72997502 2744987204 12939562299845728 0 8414286920957352 1083430833802996 162172532 22732088731566506 3584853667
9889220 52074253 4596608 18967106 3807321208 34535853929424632 0 11670689390962514 289382121038592 5121172724 19637723631512872 191177819
[root@yttrium ~]# uname -a
Linux yttrium 2.6.9-023stab033.7-enterprise #1 SMP Wed Nov 15 14:58:00 MSK 2006 i686 i686 i386 GNU/Linux
[root@yttrium ~]#

openvz 2.6.18 el5 kernel:

Quote:
[root@beast ~]# more /proc/vz/vestat
Version: 2.2

VEID	user	nice	system	uptime	idle	strv	uptime
used	maxlat	totlat	numsched				
9889998	133305	119	32588	18161899	337438033162426	0	
42371308510342	397498699339	0	0	0			
9890597	3578	59165	216691	20955483	385816188773967	0	
48888681988797	675360951251	0	0	0			
9890697	1082186	151	167382	21137597	384817445703089	0	
49313548979135	2953753939612	0	0	0			
9891687	98727	395	100033	59602854	1088224694685604	0	
139052138699873	525936556194	0	0	0			
9890885	879421	144	32052757	70982874	577295962439301	0	
165601472128687	80575127114055	0	0	0			
9890873	15481	338	10649	76233353	1390464816241648	0	
177850721787158	74312085540	0	0	0			
9890825	325128	330	192029	91927040	1699110331139417	0	
214463745852951	1260942674159	0	0	0			
9891890	67485	625	42338	121264030	2162101569813360	0	
282906292312196	264202327921	0	0	0			
9890978	278071	538	298084	146878626	2668734525714551	0	
342666792897541	1481079142459	0	0	0			
9890715	311897	635	187855	147589429	2735871461013208	0	
344325080204707	1291547097603	0	0	0			

Subject: Re: Where can we download the "vzstat" source code?
Posted by dev on Mon, 05 Nov 2007 18:24:56 GMT
View Forum Message <> Reply to Message

they should be the same, but the uptime in cycles should give more precise values.
P.S. However, there are rare cases when it doen't work in case your SMP system have non-syncronized cycles accross cpus.

Subject: Re: Where can we download the "vzstat" source code?
Posted by KoolK on Wed, 07 Nov 2007 23:30:10 GMT
View Forum Message <> Reply to Message

Hi,

My second doubt is still unsolved.I created two VEs and then executed infinite process with sleep of 5 seconds. In VE I could run 40 instances of infinite processes in background. After that system failed to run the process as there was no memory to allocate.

Still fields in /proc/ve/vestat maxlat, totlat, numsched remains 0. According to my understanding the latency should increase as system load increases. I want to use this to show system performance for eg to show values in red if average latency in last scanning period has increased. My plan is to show it in graphs.

So please anybody can tell why maxlat, totlat and numsched remain 0?

Thanks in advance

Subject: Re: Where can we download the "vzstat" source code?
Posted by dev on Thu, 08 Nov 2007 07:36:25 GMT
View Forum Message <> Reply to Message

what do you mean by "executed infinite process with sleep of 5 seconds"?

if process sleeps - it doesn't take CPU and doesn't affect latency.
as you can see in my example output of /proc/vz/vestat latencies were not 0. Do you have zeros there in your case?

Subject: Re: Where can we download the "vzstat" source code?
Posted by KoolK on Thu, 08 Nov 2007 08:59:55 GMT
View Forum Message <> Reply to Message

Hi,

by infinite process I mean following file

```
#include<stdio.h>

int main()
{

 for(;;);
 return 0;
}
```

so I run 40 processes of such kind with some having nice value 19, after that system says memory insufficient to spawn new process.

Still in this case /proc/vz/vestat shows 0 for maxlat, totallat n numsched. even in rickB's example in previous post, the file has 0 values.

this is my top command output in VE 1001 after running 37 instances of infinite process

top - 11:58:15 up 2:22, 0 users, load average: 33.96, 26.71, 13.17
Tasks: 51 total, 36 running, 15 sleeping, 0 stopped, 0 zombie
Cpu(s): 99.3% us, 0.3% sy, 0.3% ni, 0.0% id, 0.0% wa, 0.0% hi, 0.0% si
Mem: 505200k total, 495700k used, 9500k free, 39208k buffers
Swap: 1084348k total, 0k used, 1084348k free, 187308k cached

PID	USER	PR	NI	VIRT	RES	SHR	S	%CPU	%MEM	TIME+	COMMAND
22084	root	25	0	1384	232	184	R	3.7	0.0	0:12.62	infinite
22036	root	25	0	1384	232	184	R	3.3	0.0	0:13.89	infinite
22037	root	25	0	1384	232	184	R	3.3	0.0	0:13.19	infinite
22038	root	25	0	1384	232	184	R	3.3	0.0	0:13.14	infinite
22039	root	25	0	1388	236	184	R	3.3	0.0	0:13.12	infinite
22040	root	25	0	1384	232	184	R	3.3	0.0	0:13.10	infinite
22043	root	25	0	1384	232	184	R	3.3	0.0	0:13.09	infinite
22044	root	25	0	1384	232	184	R	3.3	0.0	0:12.89	infinite
22046	root	25	0	1384	232	184	R	3.3	0.0	0:12.92	infinite
22049	root	25	0	1388	236	184	R	3.3	0.0	0:12.90	infinite
22050	root	25	0	1388	236	184	R	3.3	0.0	0:12.89	infinite
22051	root	25	0	1384	232	184	R	3.3	0.0	0:13.59	infinite
22052	root	25	0	1388	236	184	R	3.3	0.0	0:12.89	infinite
22065	root	25	0	1384	232	184	R	3.3	0.0	0:11.80	infinite
22066	root	25	0	1384	232	184	R	3.3	0.0	0:12.59	infinite
22067	root	25	0	1384	232	184	R	3.3	0.0	0:11.89	infinite
22068	root	25	0	1384	232	184	R	3.3	0.0	0:11.84	infinite

But still output of /proc/vz/vestat is -
[root@khyati-laptop home]# cat /proc/vz/vestat
Version: 2.2

VEID	user	nice	system	uptime	idle	strv	uptime
used	maxlat	totlat	numsched				
1001	990058	931	1922	8628377	2320326584544	0	
3988923534912	1559444504772	0	0	0			

Subject: Re: Where can we download the "vzstat" source code?
Posted by gurtaj on Thu, 11 Sep 2008 01:55:03 GMT
View Forum Message <> Reply to Message

Hello. I've found this message while looking for information about how to calculate the CPU usage percentage of a container (yeah... one more trying to do the same).

The fact is that I'm having some troubles, so I'd appreciate any help.

I started a container (# 109) and captured the following values:

VEID	user	nice	system	uptime	idle	strv	uptime	used	maxlat	totlat	numsched
109	9	0	30	225316	4494871181910	0	4494982391820	793267650	0	0	0

So, let's start with the calculations:

- if jiffies_per_second = 1000 (as indicated by dev), the first uptime value equals 225.316 seconds.

Anyway, I was trying to match both uptime values, the one in jiffies and the one in cycles.

As far as I understand, the second value should be divided by cycles_per_jiffy:

CPU: Intel Core Duo T2450 @ 2.00 GHz
CPU MHz: 1995.115

cycles_per_jiffy = 1995115 / 1000 = 1995.115
(which seems pretty small compared with the one given by dev as example of his processor, which I guess it was a 1.86 GHz)

Now, if I devide the second uptime value by the previously obtained cycles_per_jiffy:

seconds = uptime / cycles_per_jiffy = 4494982391820 / 1995.115 = 22527270349023.399653654

... OK, the value starts similar, but they are quite different.

So... what I'm understanding wrong?

Thanks for the help!

Subject: Re: Where can we download the "vzstat" source code?
Posted by gurtaj on Sun, 14 Sep 2008 19:22:58 GMT
View Forum Message <> Reply to Message

Anybody around?

I kept playing with the containers in this Ubuntu system.

I started a container, obtained the first uptime value, waited 60 seconds and obtained the uptime value again.

First measure: uptime = 2033
60 seconds later: uptime = 8030

The difference is 6000 jiffies (in 60 seconds), so there are 100 jiffies per second.

Subject: Re: Where can we download the "vzstat" source code?
Posted by gurtaj on Sun, 14 Sep 2008 19:43:43 GMT
View Forum Message <> Reply to Message

Following on the previous message...

The second uptime value (on cycles) in those two measures was:

First measure: uptime = 40566822150
60 seconds later: uptime = 160205154630

Dividing by the CPU Hz (should get seconds) obtained from /proc/cpuinfo, 1995067:

First measure: uptime = 20333.563810138
60 seconds later: uptime = 80300.638840701

As it can be seen, both uptime values from the first measure are similar, but differ on orders of magnitude.

uptime in jiffies = 2033 ==> 20.33 seconds
uptime in cycles = 40566822150
uptime in cycles/frequency ==> 20333.563810138 seconds

So, where am I missing the point?

Subject: Re: Where can we download the "vzstat" source code?
Posted by disaster on Tue, 23 Sep 2008 12:01:40 GMT
View Forum Message <> Reply to Message

any news on this topic? Or any opensource solutions for viewing CPU Usage?

Subject: Re: Where can we download the "vzstat" source code?
Posted by gurtaj on Tue, 23 Sep 2008 15:01:47 GMT
View Forum Message <> Reply to Message

No news around here. I haven't made sense of those numbers yet.
