
Subject: How to reset the failcount (failcnt) in /dev/user_beancounters ?

Posted by [bjmg](#) on Wed, 15 Mar 2006 11:33:37 GMT

[View Forum Message](#) <> [Reply to Message](#)

Hi,

How can I reaset the failcount in /dev/user_beancounters ?

In previous cases I just restarted the hardware node to be sure that everything is in a defined state.

Because this time I have to do this on a production system I don't want to reboot the hardware node. I hope there is another way to do this without even restarting any of the VPSs. If so, please explain it to me (or just say it is described in document X).

Bernhard

Subject: Re: How to reset the failcount (failcnt) in /dev/user_beancounters ?

Posted by [kir](#) on Wed, 15 Mar 2006 12:05:09 GMT

[View Forum Message](#) <> [Reply to Message](#)

AFAIK there is no way to reset failcnt in /proc/user_beancounters without restarting an appropriate VPS. This is actually a valid behaviour since there can be many different programs reading those values, and they should not expect those values to suddenly go down to zero.

So what you should do is to collect the current values and later check against them. For the 'problematic' VPSs I usually employ a cron job which periodically cats /proc/user_beancounters to a file, together with a timestamp.

A simple shell script can also be written to mail a sysadmin about "bad" changes in /proc/user_beancounters.

Subject: Re: How to reset the failcount (failcnt) in /dev/user_beancounters ?

Posted by [bjmg](#) on Wed, 15 Mar 2006 12:17:25 GMT

[View Forum Message](#) <> [Reply to Message](#)

A simple restart of the VPS does not work. The failcounts are still on the same level. Anyway I'll just write a script that reports changes on failcounts to me.

Bernhard

Subject: Re: How to reset the failcount (failcnt) in /dev/user_beancounters ?

Posted by [kir](#) on Wed, 15 Mar 2006 12:56:00 GMT

[View Forum Message](#) <> [Reply to Message](#)

Quote:A simple restart of the VPS does not work.

That tells that there is some UBC leak (incorrect counting). What kernel are you using? Do you have some strange messages from the kernel? Can you post the output of /proc/user_beancounters for the given VPS when this VPS is stopped (some 'usage' counters should be > 0, thus failcnt is not being reset).

Subject: Re: How to reset the failcount (failcnt) in /dev/user_beancounters ?

Posted by [bjmg](#) on Wed, 15 Mar 2006 13:35:58 GMT

[View Forum Message](#) <> [Reply to Message](#)

Here is the data:

user_beancounters (before stop of VPS 22114):

22114: kmemsize	8006290	12162801	131972572	145169828	20533210
lockedpages	0	0	6442	6442	0
privvmpages	84085	112592	144989	159487	0
shmpages	671	1978	14498	14498	0
dummy	0	0	0	0	0
numproc	113	183	5332	5332	24
physpages	37958	65609	0	2147483647	0
vmguarpages	0	0	144989	2147483647	0
oomguarpages	39682	66342	144989	2147483647	0
numtcpsock	30	127	5332	5332	1347
numflock	13	31	1000	1100	0
numpty	2	5	512	512	0
numsiginfo	0	50	1024	1024	0
tcpsndbuf	31192	739696	22150986	43990858	1549
tcprcvbuf	2228	300780	22150986	43990858	0
othersockbuf	187152	430392	11075492	32915364	0
dgramrcvbuf	0	169328	11075492	11075492	0
numothersock	153	193	5332	5332	0
dcachesize	357099	466733	28811184	29675520	0
numfile	1084	1593	51520	51520	0
dummy	0	0	0	0	0
dummy	0	0	0	0	0
dummy	0	0	0	0	0
numiptent	143	143	200	200	0

user_beancounters (after stop of VPS 22114)

22114: kmemsize	1177	12162801	131972572	145169828	20533210
lockedpages	0	0	6442	6442	0
privvmpages	0	112592	144989	159487	0
shmpages	0	1978	14498	14498	0
dummy	0	0	0	0	0
numproc	0	183	5332	5332	24

physpages	0	65609	0	2147483647	0
vmguarpages	0	0	144989	2147483647	0
oomguarpages	14	66342	144989	2147483647	0
numtcpsock	0	127	5332	5332	1347
numflock	0	31	1000	1100	0
numpty	0	5	512	512	0
numsiginfo	0	50	1024	1024	0
tcpsndbuf	0	739696	22150986	43990858	1549
tcprcvbuf	0	300780	22150986	43990858	0
othersockbuf	0	430392	11075492	32915364	0
dgramrcvbuf	0	169328	11075492	11075492	0
numothersock	0	193	5332	5332	0
dcachesize	0	466733	28811184	29675520	0
numfile	0	1593	51520	51520	0
dummy	0	0	0	0	0
dummy	0	0	0	0	0
dummy	0	0	0	0	0
numiptent	0	143	200	200	0

user_beancounters (after start of VPS 22114):

22114: kmemsize	6021807	12162801	131972572	145169828	20533210
lockedpages	0	0	6442	6442	0
privvmpages	67938	112592	144989	159487	0
shmpages	26	1978	14498	14498	0
dummy	0	0	0	0	0
numproc	87	183	5332	5332	24
physpages	24228	65609	0	2147483647	0
vmguarpages	0	0	144989	2147483647	0
oomguarpages	24242	66342	144989	2147483647	0
numtcpsock	20	127	5332	5332	1347
numflock	16	31	1000	1100	0
numpty	0	5	512	512	0
numsiginfo	0	50	1024	1024	0
tcpsndbuf	33420	739696	22150986	43990858	1549
tcpvbuf	0	300780	22150986	43990858	0
othersockbuf	158188	430392	11075492	32915364	0
dgramrcvbuf	0	169328	11075492	11075492	0
numothersock	138	193	5332	5332	0
dcachesize	345619	466733	28811184	29675520	0
numfile	912	1593	51520	51520	0
dummy	0	0	0	0	0
dummy	0	0	0	0	0
dummy	0	0	0	0	0
numiptent	143	143	200	200	0

uname:

Linux master.mediadesign-gruen.de 2.6.8-022stab070.1-enterprise #1 SMP Mon Feb 20 19:31:28
MSK 2006 i686 athlon i386 GNU/Linux

This is an excerpt of dmesg output:
(not only from that stop - start event but all relevant infos I found)

VPS: 22114: started
Unable to load interpreter
Unable to load interpreter
Unable to load interpreter
VPS: 22114: stopped
...
VPS: 22114: started
Fatal resource shortage: numiptent, UB 22114.
Fatal resource shortage: numiptent, UB 22114.
Fatal resource shortage: numiptent, UB 22114.
Fatal resource shortage: numiptent, UB 22114.
VPS: 22114: stopped
...
VPS: 22114: started
Out of socket memory
Out of socket memory
Out of socket memory
Out of socket memory
Out of socket memory
VPS: 22114: stopped
...
VPS: 22114: started
TCP: time wait bucket table overflow
TCP: time wait bucket table overflow
TCP: time wait bucket table overflow
TCP: time wait bucket table overflow
VPS: 22114: stopped
...
...
TCP: time wait bucket table overflow
TCP: time wait bucket table overflow
TCP: time wait bucket table overflow
TCP: time wait bucket table overflow
TCP: too many of orphaned sockets
TCP: too many of orphaned sockets
TCP: too many of orphaned sockets
TCP: too many of orphaned sockets
TCP: too many of orphaned sockets
TCP: too many of orphaned sockets

TCP: too many of orphaned sockets
VPS: 22114: stopped
VPS: 22114: started
[no problem memory problems since that point - but still high failcount]
...
VPS: 22114: stopped
VPS: 22114: started
[this is the restart I did for getting the infos above]

I hope this helps debugging it.
I can also provide a psaux output (with VPS 22114 stoped).

Bernhard

Subject: Re: How to reset the failcount (failcnt) in /dev/user_beancounters ?
Posted by [kir](#) on Wed, 15 Mar 2006 14:23:34 GMT
[View Forum Message](#) <> [Reply to Message](#)

Thanks for the response!

Apparently, as a neighbour kernel hacker just explained me, this is not a bug.

The problem is tcp time wait buckets are still there after VPS is stopped, thus we are having >0 kmemsize. They should go away after a timeout (5 minutes or so), after that kmemsize will be set (zeroed) and fail counters should be reset to zero.

'oomguarpages > 0' is known bug; the good thing is fail counters will be reset even with oomguarpages > 0.

So, you have two options

- (1) stop VPS and wait 5 minutes or so
 - (2) check the difference of failcnts, not their absolute values.
-

Subject: Re: How to reset the failcount (failcnt) in /dev/user_beancounters ?
Posted by [kir](#) on Thu, 13 Jul 2006 08:03:01 GMT
[View Forum Message](#) <> [Reply to Message](#)

The info from this thread has been summarized at wiki: UBC failcnt reset.
