
Subject: mysqld not running in all CT's
Posted by [plotinus](#) on Sat, 12 Nov 2011 10:13:03 GMT
[View Forum Message](#) <> [Reply to Message](#)

Hello all,

I'm running openVZ on centOS 5.7

uname -r

2.6.18-274.3.1.el5.028stab094.3

When I create new containers using a customized fedora v15 template, everything works ok.

But when I have a certain number of CT's running, mysqld can't start for the last started CT's. If I stop some containers then I can start mysqld and it works.

I enter in a container and see:

```
[root]# tail -100 /var/logs/mysqld.log
```

```
111112 10:22:50 mysqld_safe Starting mysqld daemon with databases from /var/lib/mysql
```

```
111112 10:22:50 [Note] Plugin 'FEDERATED' is disabled.
```

```
111112 10:22:50 InnoDB: The InnoDB memory heap is disabled
```

```
111112 10:22:50 InnoDB: Mutexes and rw_locks use GCC atomic builtins
```

```
111112 10:22:50 InnoDB: Compressed tables use zlib 1.2.5
```

```
111112 10:22:50 InnoDB: Using Linux native AIO
```

```
111112 10:22:50 InnoDB: Warning: io_setup() failed with EAGAIN. Will make 5 attempts before giving up.
```

```
InnoDB: Warning: io_setup() attempt 1 failed.
```

```
InnoDB: Warning: io_setup() attempt 2 failed.
```

```
InnoDB: Warning: io_setup() attempt 3 failed.
```

```
InnoDB: Warning: io_setup() attempt 4 failed.
```

```
InnoDB: Warning: io_setup() attempt 5 failed.
```

```
111112 10:22:52 InnoDB: Error: io_setup() failed with EAGAIN after 5 attempts.
```

```
InnoDB: You can disable Linux Native AIO by setting innodb_native_aio = off in my.cnf
```

```
111112 10:22:52 InnoDB: Initializing buffer pool, size = 128.0M
```

```
111112 10:22:52 InnoDB: Completed initialization of buffer pool
```

```
111112 10:22:52 - mysqld got signal 11 ;
```

This could be because you hit a bug. It is also possible that this binary

or one of the libraries it was linked against is corrupt, improperly built,

or misconfigured. This error can also be caused by malfunctioning hardware.

We will try our best to scrape up some info that will hopefully help diagnose

the problem, but since we have already crashed, something is definitely wrong

and this may fail.

```
key_buffer_size=8388608
```

```
read_buffer_size=131072
```

```
max_used_connections=0
```

```
max_threads=151
```

```
thread_count=0
```

```
connection_count=0
```

It is possible that mysqld could use up to

```
key_buffer_size + (read_buffer_size + sort_buffer_size)*max_threads = 337868 K
```

bytes of memory

Hope that's ok; if not, decrease some variables in the equation.

Thread pointer: 0x0

Attempting backtrace. You can use the following information to find out where mysqld died. If you see no messages after this, something went terribly wrong...

```
111112 10:22:52 mysqld_safe mysqld from pid file /var/run/mysqld/mysqld.pid ended
```

I get also this error:

```
111112 10:58:31 mysqld_safe Starting mysqld daemon with databases from /var/lib/mysql
```

```
111112 10:58:31 [Note] Plugin 'FEDERATED' is disabled.
```

```
111112 10:58:31 InnoDB: The InnoDB memory heap is disabled
```

```
111112 10:58:31 InnoDB: Mutexes and rw_locks use GCC atomic builtins
```

```
111112 10:58:31 InnoDB: Compressed tables use zlib 1.2.5
```

```
111112 10:58:31 InnoDB: Using Linux native AIO
```

```
111112 10:58:31 InnoDB: Warning: io_setup() failed with EAGAIN. Will make 5 attempts before giving up.
```

```
InnoDB: Warning: io_setup() attempt 1 failed.
```

```
InnoDB: Warning: io_setup() attempt 2 failed.
```

```
InnoDB: Warning: io_setup() attempt 3 failed.
```

```
InnoDB: Warning: io_setup() attempt 4 failed.
```

```
InnoDB: Warning: io_setup() attempt 5 failed.
```

```
111112 10:58:34 InnoDB: Error: io_setup() failed with EAGAIN after 5 attempts.
```

```
InnoDB: You can disable Linux Native AIO by setting innodb_native_aio = off in my.cnf
```

```
111112 10:58:34 InnoDB: Initializing buffer pool, size = 128.0M
```

```
111112 10:58:34 InnoDB: Completed initialization of buffer pool
```

```
111112 10:58:34 - mysqld got signal 11 ;
```

This could be because you hit a bug. It is also possible that this binary

or one of the libraries it was linked against is corrupt, improperly built,

or misconfigured. This error can also be caused by malfunctioning hardware.

```
111112 10:58:34 mysqld_safe mysqld from pid file /var/run/mysqld/mysqld.pid ended
```

Here is the /etc/my.cnf file of a container:

```
[mysqld]
```

```
datadir=/var/lib/mysql
```

```
socket=/var/lib/mysql/mysql.sock
```

```
user=mysql
```

```
# Disabling symbolic-links is recommended to prevent assorted security risks
```

```
symbolic-links=0
```

```
[mysqld_safe]
```

```
log-error=/var/log/mysqld.log
```

```
pid-file=/var/run/mysqld/mysqld.pid
```

Here are the UBC:

```
[root]# cat /proc/user_beancounters
```

```
Version: 2.5
```

uid	resource	held	maxheld	barrier	limit	failcnt
113:	kmemsize	2526590	3377739	168442484		185286732

```

0
lockedpages          0          0          8224          8224          0
privvmpages         11502         78323         1381448         1519592
0
shmpages            128           128           138144          138144
0
dummy               0             0             0               0               0
numproc             14            28            4112            4112            0
physpages           4268          10532          0 9223372036854775807
0
vmguarpages         0             0             1381448 9223372036854775807
0
oomguarpages        4268          10532          1381448 9223372036854775807
0
numtcpsock          4             4             4112            4112            0
numflock            4             5             1000            1100            0
numpty              1             1             411             411             0
numsiginfo          0             4             1024            1024            0
tcpsndbuf           70016         70016         39304742         56147494
0
tcprcvbuf           65536         65536         39304742         56147494
0
othersockbuf        2328          12200          19652371         36495123
0
dgramrcvbuf         0             8472          19652371         19652371
0
numothersock        2             5             4112            4112            0
dcachesize          231933        255408         36792419         37896192
0
numfile             519           778           65792           65792            0
dummy               0             0             0               0               0
dummy               0             0             0               0               0
dummy               0             0             0               0               0
numiptent           24            24            100             100             0

```

SELINUX is disabled on the hn.

I've found the fantstic vzsplint and used it for partitioning the resources, but the problem still persists.

How do you solve it?

I'm reading some articles but I don't identify where lies the problem.

Please can anyone help?

Thanks in advance!!!

P

Subject: Re: mysqld not running in all CT's
 Posted by [plotinus](#) on Sun, 04 Dec 2011 22:10:17 GMT

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

Now I've got 70 running containers, but only on 24 of them mysqld can run. I start them randomly, and mysqld runs only on the 24 first containers that I boot.

I've changed ubc parms, but no success.

Can please anybody give me any hint? I dont know where lies the problem: mysql? ubc parms? ulimit?

Thanks in advance,

P

Subject: Re: mysqld not running in all CT's

Posted by [plotinus](#) on Mon, 05 Dec 2011 10:47:26 GMT

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

I've read lots of posts.

Some of they say "recompile kernel", some others "unload modules".

How can I know what is the right way to follow?

I can't identify where is the problem.

Total lost and becoming crazy.

Subject: More info: it has to do with the template

Posted by [plotinus](#) on Mon, 05 Dec 2011 22:02:36 GMT

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

I've used the same template for all containers: the standard fedora 15 x86_64, with some changes: mysql, and security issues.

Now, I've used the standard ubuntu 11.04 template plus mysql and it works!

Does anybody know why the fedora template doesn't work?

I've copied the ubuntu's my.conf file into the fedora template, but it doesn't work.

Thanks in advance,

P

Subject: Re: mysqld not running in all CT's

Posted by [ergosteur](#) on Tue, 26 Jun 2012 14:28:42 GMT

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

Did you ever figure this out? I have the same issue, but running Proxmox (Debian 2.6.32-7-pve) with the default Proxmox Debian 6.0 template and MySQL 5.5 from DotDeb.

MySQL 5.1 works fine, but as soon as I update to 5.5 I can only start 24 CTs. the 25th fails with

```
Jun 26 10:20:26 myhost mysqld_safe: Starting mysqld daemon with databases from /var/lib/mysql
Jun 26 10:20:26 myhost mysqld: 120626 10:20:26 [Note] Plugin 'FEDERATED' is disabled.
```

Jun 26 10:20:26 myhost mysqld: 120626 10:20:26 InnoDB: The InnoDB memory heap is disabled
Jun 26 10:20:26 myhost mysqld: 120626 10:20:26 InnoDB: Mutexes and rw_locks use GCC atomic builtins
Jun 26 10:20:26 myhost mysqld: 120626 10:20:26 InnoDB: Compressed tables use zlib 1.2.3.4
Jun 26 10:20:26 myhost mysqld: 120626 10:20:26 InnoDB: Using Linux native AIO
Jun 26 10:20:26 myhost mysqld: 120626 10:20:26 InnoDB: Warning: io_setup() failed with EAGAIN. Will make 5 attempts before giving up.
Jun 26 10:20:26 myhost mysqld: InnoDB: Warning: io_setup() attempt 1 failed.
Jun 26 10:20:27 myhost mysqld: InnoDB: Warning: io_setup() attempt 2 failed.
Jun 26 10:20:27 myhost mysqld: InnoDB: Warning: io_setup() attempt 3 failed.
Jun 26 10:20:28 myhost mysqld: InnoDB: Warning: io_setup() attempt 4 failed.
Jun 26 10:20:28 myhost mysqld: InnoDB: Warning: io_setup() attempt 5 failed.
Jun 26 10:20:29 myhost mysqld: 120626 10:20:29 InnoDB: Error: io_setup() failed with EAGAIN after 5 attempts.
Jun 26 10:20:29 myhost mysqld: InnoDB: You can disable Linux Native AIO by setting innodb_use_native_aio = 0 in my.cnf
Jun 26 10:20:29 myhost mysqld: 120626 10:20:29 InnoDB: Fatal error: cannot initialize AIO sub-system
Jun 26 10:20:29 myhost mysqld: 120626 10:20:29 [ERROR] Plugin 'InnoDB' init function returned error.
Jun 26 10:20:29 myhost mysqld: 120626 10:20:29 [ERROR] Plugin 'InnoDB' registration as a STORAGE ENGINE failed.
Jun 26 10:20:29 myhost mysqld: 120626 10:20:29 [ERROR] Unknown/unsupported storage engine: InnoDB
Jun 26 10:20:29 myhost mysqld: 120626 10:20:29 [ERROR] Aborting
Jun 26 10:20:29 myhost mysqld:
Jun 26 10:20:29 myhost mysqld: 120626 10:20:29 [Note] /usr/sbin/mysqld: Shutdown complete
Jun 26 10:20:29 myhost mysqld:
Jun 26 10:20:29 myhost mysqld_safe: mysqld from pid file /var/run/mysqld/mysqld.pid ended

Subject: Re: mysqld not running in all CT's
Posted by [plotinus](#) on Wed, 27 Jun 2012 18:45:21 GMT
[View Forum Message](#) <> [Reply to Message](#)

What a mistery!
I forgot about fedora, and now with ubuntu 11.04 server everything goes fine, since months ago.
I'm afraid I can't help you, I'm sorry.

Subject: Re: mysqld not running in all CT's
Posted by [RedKrieg](#) on Tue, 07 Aug 2012 22:36:42 GMT
[View Forum Message](#) <> [Reply to Message](#)

I ran into the same issue and was able to work around it. There are two solutions I found.

Disable AIO in all containers' my.cnf files by adding "innodb_use_native_aio=0"
Increase the global limit for the server in "/proc/sys/fs/aio-max-nr"

The first solution is simple and causes MySQL to revert to the previous 5.1 style behavior. Unfortunately if you're offering service to end-users, they may remove this setting or set up their own MySQL instance.

The second solution is more nuanced, but probably better if you can't rely on the containers' configurations to be static or if you want the enhanced performance of Asynchronous I/O.

On my CentOS5 based Virtuozzo node, the default value of aio-max-nr is 65536. A default MySQL 5.5 install will use 2661 of these, leading to the maximum of 24 started MySQL instances with the setting enabled. You should tune the value to be greater than max_instances * 2661 to ensure that you don't run into this issue.

In my case, I opted for the first solution because I didn't have time to test the effects of having this many AIO tasks happening in parallel on a single hardware node and I understand the 5.1 behavior very well for large tasks. If you opt to go with option 2 on a large install, I'd be interested in hearing how it worked out for you (as I'm sure many others reading this thread are as well).

All the docs I could find related to this:

docs.oracle.com/cd/E17952_01/refman-5.5-en/innodb-performance-aio-linux.html
