
Subject: git inside VPS fails on mmap
Posted by [pollux](#) on Wed, 28 May 2008 07:44:41 GMT
[View Forum Message](#) <> [Reply to Message](#)

Hi,

I am trying to setup a git server inside a VPS. The git repository has fairly large objects (it's a local kernel tree).

The following error happens on git clone (using the git protocol):

```
2008-05-28_07:11:14.41523 [32100] Request upload-pack for '/git/linux-2.6.git'  
2008-05-28_07:11:14.41641 fatal: packfile  
./objects/pack/pack-95d32a83b9918b18223858e7349d5b64585b2a51.pack cannot be mapped.
```

Running stace on the git servers shows indeed a mmap failure with -ENOMEM:

```
[pid 32195] open("./objects/pack/pack-95d32a83b9918b18223858e7349d5b64585b2a51.pack",  
O_RDONLY) = 5  
[pid 32195] fstat(5, {st_mode=S_IFREG|0444, st_size=155244017, ...}) = 0  
[pid 32195] mmap(NULL, 155244017, PROT_READ, MAP_PRIVATE, 5, 0) = -1 ENOMEM  
(Cannot allocate memory)
```

Now what is weird is that I have no quota exceeded in /proc/user_beancounters (last columns is 0 everywhere). I have tried raising all quota to 9223372036854775807, no changes.

Any idea ?

Thanks

Host environment: 2.6.18-028stab053 x86_64
Guest OS: Debian Etch x86_64

Subject: Re: git inside VPS fails on mmap
Posted by [maratrus](#) on Wed, 28 May 2008 09:59:02 GMT
[View Forum Message](#) <> [Reply to Message](#)

Hello,

It might have been several reasons:

1. You had physically lack of memory on your HN.
2. The virtual address space was come to the end.
3. You have exceeded the number of map areas your process was allowed. (sysctl

vm.max_map_count parameter). Information about map areas can be obtained via /proc/PID/maps
4. There was not appropriate "hole" in address space of particular process to place there additional map area.

Subject: Re: git inside VPS fails on mmap
Posted by [pollux](#) on Wed, 28 May 2008 10:16:38 GMT
[View Forum Message](#) <> [Reply to Message](#)

1. The HN has 4 GB of RAM, + 4 of swap. It does not seem to have problems

```
# free
      total    used    free   shared  buffers   cached
Mem:   4041736  3972848    68888        0   433104   848136
-/+ buffers/cache:  2691608  1350128
Swap:   4194296      336  4193960
```

2. That, I don't know how to check. The host is 64-bits, so I assume there is no limitation on 4GB.

3.

```
# sysctl vm.max_map_count
vm.max_map_count = 65536
```

(same on the HN and on the guest)

The proc values shows no large number of maps:

```
# cat /proc/`pidof git-daemon`/maps | wc -l
16
```

Strace shows that git-daemon adds 27 mmmaps before failing.

I tried raising vm.max_map_count to an arbitrary high number, no changes.

4. I don't know how to check that.

I would add that the HN is running 34 servers, maybe it is too big for its capacity ? AFAICT, there is no sign of saturation (cpu load is 0.14, There is some free RAM, etc.)

Thanks for your help

Subject: Re: git inside VPS fails on mmap
Posted by [maratrus](#) on Wed, 28 May 2008 10:55:43 GMT
[View Forum Message](#) <> [Reply to Message](#)

Hi,

could you please show "/proc/user_beancounters" from inside the VE where is the issue occurs.
And also "sysctl -a | grep overcommit"
Thanks.

Subject: Re: git inside VPS fails on mmap
Posted by [pollux](#) on Wed, 28 May 2008 11:31:06 GMT
[View Forum Message](#) <> [Reply to Message](#)

sure:

```
git:~# sysctl -a | grep overcommit
error: "Operation not permitted" reading key "net.ipv4.route.flush"
vm.overcommit_ratio = 50
vm.overcommit_memory = 0
git:~# cat /proc/user_beancounters
Version: 2.5
```

uid	resource	held	maxheld	barrier	limit	failcnt
40:	kmemsize	3480487	4136968	9223372036854775807		
9223372036854775807		0				
	lockedpages	0	0	9223372036854775807		
9223372036854775807		0				
	privvmpages	6561	18731	9223372036854775807		
9223372036854775807		0				
	shmpages	1424	2080	9223372036854775807		
9223372036854775807		0				
	dummy	0	0	0	0	0
	numproc	21	29	9223372036854775807		
9223372036854775807		0				
	physpages	3210	4990	9223372036854775807		
9223372036854775807		0				
	vmguarpages	0	0	9223372036854775807		
9223372036854775807		0				
	oomguarpages	3210	4990	9223372036854775807		
9223372036854775807		0				
	numtcpsock	7	23	9223372036854775807		
9223372036854775807		0				
	numflock	4	5	9223372036854775807	9223372036854775807	
0						
	numpty	2	2	9223372036854775807	9223372036854775807	
0						

numsiginfo	0	7	9223372036854775807	9223372036854775807
0				
tcpsndbuf	50512	342104	9223372036854775807	
9223372036854775807	0			
tcprcvbuf	2544	67984	9223372036854775807	
9223372036854775807	0			
othersockbuf	4592	12224	9223372036854775807	
9223372036854775807	0			
dgramrcvbuf	0	8440	9223372036854775807	
9223372036854775807	0			
numothersock	3	10	9223372036854775807	
9223372036854775807	0			
dcachesize	0	0	9223372036854775807	9223372036854775807
0				
numfile	478	564	9223372036854775807	
9223372036854775807	0			
dummy	0	0	0	0
dummy	0	0	0	0
dummy	0	0	0	0
numiptent	10	10	9223372036854775807	
9223372036854775807	0			

Subject: Re: git inside VPS fails on mmap
 Posted by [maratrus](#) on Wed, 28 May 2008 12:12:12 GMT
[View Forum Message](#) <> [Reply to Message](#)

Hello,

1. Please show "ulimit -a" from inside the VE.
2. I was looking through RHEL based kernel to find out from where we can obtain ENONEM and where have you got your kernel "2.6.18-028stab053"?

Subject: Re: git inside VPS fails on mmap
 Posted by [pollux](#) on Wed, 28 May 2008 12:18:38 GMT
[View Forum Message](#) <> [Reply to Message](#)

```
# ulimit -a
Linux git.inl.fr 2.6.18-028stab053 #1 SMP Fri Feb 15 16:06:31 CET 2008 x86_64 GNU/Linux
```

The HN is a Debian Etch, with a kernel compiled locally. I have uploaded the config file here:
<http://pastebin.ca/1032013>

Subject: Re: git inside VPS fails on mmap
Posted by [maratrus](#) on Wed, 28 May 2008 12:22:22 GMT
[View Forum Message](#) <> [Reply to Message](#)

Hi,

Quote:

```
# ulimit -a  
Linux git.inl.fr 2.6.18-028stab053 #1 SMP Fri Feb 15 16:06:31 CET 2008 x86_64 GNU/Linux
```

Very weird output. Are you sure that is was "ulimit -a" (from inside the VE) command may be "uname -a".

Subject: Re: git inside VPS fails on mmap
Posted by [pollux](#) on Wed, 28 May 2008 12:26:51 GMT
[View Forum Message](#) <> [Reply to Message](#)

Haha, I have been caught

```
git:~# uname -a  
Linux git.inl.fr 2.6.18-028stab053 #1 SMP Fri Feb 15 16:06:31 CET 2008 x86_64 GNU/Linux  
git:~# ulimit -a  
core file size          (blocks, -c) 0  
data seg size           (kbytes, -d) unlimited  
max nice                (-e) 0  
file size               (blocks, -f) unlimited  
pending signals         (-i) unlimited  
max locked memory       (kbytes, -l) unlimited  
max memory size         (kbytes, -m) unlimited  
open files              (-n) 1024  
pipe size               (512 bytes, -p) 8  
POSIX message queues    (bytes, -q) unlimited  
max rt priority         (-r) 0  
stack size              (kbytes, -s) 8192  
cpu time                (seconds, -t) unlimited  
max user processes      (-u) unlimited  
virtual memory          (kbytes, -v) unlimited  
file locks              (-x) unlimited
```

Subject: Re: git inside VPS fails on mmap

Posted by [maratrus](#) on Wed, 28 May 2008 13:01:45 GMT

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

Hi,

Quote:

with a kernel compiled locally.

How can I get kernel src like yours? Please, describe how did you build your kernel tree? I want to look at it because I've got stuck.

Subject: Re: git inside VPS fails on mmap

Posted by [pollux](#) on Wed, 28 May 2008 13:14:10 GMT

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

Well, I am using linux-2.6.18.tar.bz2 and patch-ovz028stab053.4-combined.gz.

So the steps I use to build the kernel are:

1. untar kernel, cd linux-2.6.18, apply ovz patch
2. copy config-2.6.18-028stab053 .config; make oldconfig
- 3 build the kernel and install. I generally use the debian way:

fakeroot make-kpkg --revision=ovz.02 --initrd --config menuconfig binary-arch modules_image

to produce a .deb, but running make && make install should do the same.

Subject: Re: git inside VPS fails on mmap

Posted by [maratrus](#) on Wed, 28 May 2008 14:45:23 GMT

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

Hi,

Can we play a little at that node or it's impossible?

I want to put some debug messages in do_mmap_pgoff() function to find out where is error occurred and (error == -ENOMEM) (or you can do it by yourself).

Or can I get tarball with problem VE may be I will be able to start it on my node and reproduce the issue. But if it is possible could you please also attach appropriate steps to start git inside VE and steps how were you able to reproduce the problem. Thank you very much for the patient!

Subject: Re: git inside VPS fails on mmap [solved]

Posted by [pollux](#) on Fri, 30 May 2008 20:39:34 GMT

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

Thanks to maratrus, I have found the solution ! It was indeed related to git, and not to OpenVZ.

Debian uses runit (<http://smarden.org/runit/>) to run git-daemon, and creates a default run file:

```
git:~# cat /etc/sv/git-daemon/run
#!/bin/sh
exec 2>&1
echo 'git-daemon starting.'
exec chpst -m64000000 \
  git-daemon --verbose --base-path=/var/cache /var/cache/git
```

There is indeed a hardcoded (stupid) limitation. Removing it solved the problem.

Pierre
