Subject: Re: Using threads on OpenVZ, and memory allocation versus memory usage

Posted by hvdkamer on Wed, 09 May 2007 10:47:39 GMT

View Forum Message <> Reply to Message

The above is a proof of concept program, but may be you like something from real live? On the internet you can find many people who have trouble installing Apache 2 mpm worker on Virtuozzo/OpenVZ based VE's with a privvmpages of around 512 MiB. The same applied on my virtual servers which are running on OpenVZ. So I investigated and found the following:

[1] very minimalistic server, only 4 processes (init, cron, syslogd and OpenSSH)

```
            privvmpages        663        972     131072     147456          0
            physpages          486        739          0   (2^63-1)          0
```

As you can see, hardly any memory is being used, and physpages is almost the same as privvmpages. Now I install Apache 2 mpm worker (Debian package). On a normal server which isn't stripped down like mine, you run into problems. On this stripped down server and 512 MiB burstable memory it is just possible. After a reboot, I see the following:

```
            privvmpages     112933     122415     131072     147456          0
            physpages         1573      12562          0   (2^63-1)          0
```

As you can see, there is a huge difference between physpages and privvmpages. You can't do anything with this server, because it is almost out om memory. Starting aptitude for example crashes. When I do a pmap on the four instances (two are control forks, two start in the deafult configuration 25 threads each) I see the following:

```
ve35:/# pmap 26120
26120:   /usr/sbin/apache2 -k start -DSSL
...
 total     7224K
ve35:/# pmap 26121
26121:   /usr/sbin/apache2 -k start -DSSL
...
 total     7016K
ve35:/# pmap 26122
26122:   /usr/sbin/apache2 -k start -DSSL
...
405dd000   8188K rwx--    [ anon ]
40ddc000      4K -----    [ anon ]
  => 27 times
...
```

```
 total   228444K
ve35:/# pmap 26127
26127:   /usr/sbin/apache2 -k start -DSSL
...
405dd000   8188K rwx--    [ anon ]
40ddc000      4K -----    [ anon ]
  => 27 times
...
 total   228444K
```

So every thread uses a buffer of 8 MiB. Most of the times it doesn't use thsi. On a normal machine with 512 MiB available RAM this is not a problem. You can start aptitude and a lot more programs. The only solutions is configuring Apache to use less default waiting threads. In my opinion 50 is way to much for most servers, but this is the default. And according to physpages and a free on the hardware node, it is not a real problem that this huge amount disappears in buffers...

This strange behavior between a server with real RAM and a OpenVZ server with the same amount is seen in more situations. I even asked about this in this forum but got no answer. I now know what is the problem and most of the times I can work around it. But is always a discussion with clients, because they se it working on a normal server.