By Falko Timme
Published: 2009-01-19 18:18

Splitting Resources Evenly Between OpenVZ VMs With vzsplit

Version 1.0
Author: Falko Timme <ft [at] falkotimme [dot] com>
Last edited 12/15/2008

This short guide shows how you can split your OpenVZ host resources evenly between multiple virtual machines with the help of vzsplit. vzsplit generates a sample container configuration file with a certain set of system resource control parameters that you can then apply to your virtual machines.

I do not issue any guarantee that this will work for you!

Using vzsplit

Take a look at

```
man vzsplit
```

to learn more about its usage.

Now let's assume we want to run 5 virtual machines on the OpenVZ host. We call vzsplit as follows:

```
vzsplit -n 5 -f max-limits
```

`-n 5` specifies the number of virtual machines.

```
server1:~# vzsplit -n 5 -f max-limits
```

```
The optimal swap space size is 2022 Mb, twice bigger than the RAM size

WARNING: Recommended minimal size of partition holding /var/lib/vz/private/ is 20Gb!
  Config /etc/vz/conf/ve-max-limits.conf-sample was created
  server1:~#
```

As you see, this has created the file */etc/vz/conf/ve-max-limits.conf-sample* which contains the max. values of the configuration settings for each of the 5 virtual machines. Take a look at that file:

```
vi /etc/vz/conf/ve-max-limits.conf-sample
```

```
# Configuration file generated by vzsplit for 5 containers
# on HN with total amount of physical mem 1011 Mb
# low memory 883 Mb, swap size 2047 Mb, Max treads 8000
# Resourse commit level 0:
# Free resource distribution. Any parameters may be increased
# Primary parameters
NUMPROC="1600:1600"
AVNUMPROC="452:452"
NUMTCPSOCK="1600:1600"
NUMOTHERSOCK="1600:1600"
VMGUARPAGES="135985:2147483647"

# Secondary parameters
KMEMSIZE="37052743:40758017"
TCPSNDBUF="5797314:12350914"
TCPRCVBUF="5797314:12350914"
OTHERSOCKBUF="2898657:9452257"
DGRAMRCVBUF="2898657:2898657"
OOMGUARPAGES="135985:2147483647"
PRIVVMPAGES="155342:170876"
```

```
# Auxiliary parameters
LOCKEDPAGES="1809:1809"
SHMPAGES="15534:15534"
PHYSPAGES="0:2147483647"
NUMFILE="14464:14464"
NUMFLOCK="1000:1100"
NUMPTY="160:160"
NUMSIGINFO="1024:1024"
DCACHESIZE="8088605:8331264"
NUMIPTENT="200:200"
DISKSPACE="117286:129015"
DISKINODES="110065:121072"
CPUUNITS="8587"
```

As you see, this file does not contain VM-specific settings such as hostname, IP address, name servers, etc.which makes sense because we don't want to overwrite these settings.

Now let's say we want to apply these settings to our VM with the ID *101*. This is how we do it:

```
vzctl set 101 --applyconfig max-limits --save
```

Afterwards, you can start/restart the VM 101:

```
vzctl start 101
```

You might see something like this:

```
server1:~# vzctl start 101
  Starting container ...
```

```
  vzquota : (warning) block_hard_limit [129115] < block_current_usage [139856]
  Container is mounted
  Adding IP address(es): 192.168.0.101
  bash: line 402: printf: write error: Disk quota exceeded
  bash: line 416: printf: write error: Disk quota exceeded
  bash: line 421: echo: write error: Disk quota exceeded
  bash: line 447: printf: write error: Disk quota exceeded
  Setting CPU units: 8587
  Configure meminfo: 155342
  Set hostname: test.example.com
  awk: close failed on file /dev/stdout (Disk quota exceeded)
ERROR: Can't change file /etc/hosts
vzquota : (warning) block_hard_limit [129015] < block_current_usage [139856]
Container start in progress...
server1:~#
```

For some reason vzsplit has calculated a wrong disk size. You can correct that as follows (this will set a disk size with a soft limit of 10GB and a hard limit of 11GB - adjust this to your needs):

```
vzctl set 101 --diskspace "10000000:11000000" --save
```

Then restart the VM:

```
vzctl restart 101
```

You can take a look at the new settings for the VM as follows:

```
vzctl exec 101 cat /proc/user_beancounters
```

```
server1:/etc/vz/conf# vzctl exec 101 cat /proc/user_beancounters
```

```
Version: 2.5
       uid  resource           held    maxheld    barrier      limit    failcnt
      101:  kmemsize         340278     580635   37052743   40758017          0
            lockedpages           0          0       1809       1809          0
            privvmpages         796       1655     155342     170876          0
            shmpages              0          0      15534      15534          0
            dummy                 0          0          0          0          0
            numproc               5          9       1600       1600          0
            physspages          486       1060          0 2147483647          0
            vmguarpages           0          0     135985 2147483647          0
            oomguarpages        486       1060     135985 2147483647          0
            numtcpsock            1          2       1600       1600          0
            numflock              1          2       1000       1100          0
            numpty                0          1        160        160          0
            numsiginfo            0          2       1024       1024          0
            tcpsndbuf          8928          0    5797314   12350914          0
            tcprcvbuf         16384          0    5797314   12350914          0
            othersockbuf       2232       3624    2898657    9452257          0
            dgramrcvbuf           0          0    2898657    2898657          0
            numothersock          1          3       1600       1600          0
            dcachesize            0          0    8088605    8331264          0
            numfile             129        193      14464      14464          0
            dummy                 0          0          0          0          0
            dummy                 0          0          0          0          0
            dummy                 0          0          0          0          0
            numiptent            10         10        200        200          0
server1:/etc/vz/conf#
```

The `failcnt` column is very important, it should contain only zeros; if it doesn't, this means that the vm needs more resources than are currently allocated to the VM.

To find out about the memory allocated to the VM, run

```
vzctl exec 101 free
```

```
server1:~# vzctl exec 101 free
             total       used       free     shared    buffers     cached
Mem:        621368       3268     618100          0          0          0
-/+ buffers/cache:       3268     618100
Swap:            0          0          0
server1:~#
```

Links

- OpenVZ: **http://openvz.org**