
Subject: MySQL performance inside VE

Posted by [openfauzett](#) on Sat, 13 Dec 2008 14:37:04 GMT

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

Hi OpenVZ-Folks,

on a decent hardware we are facing serious performance problems and after checking for all we could think of, the last guess is, that open VZ could be responsible for the issues. I'd appreciate any comments on the problem described below.

The Problem:

A php-based website is to be delivered by a decent machine described below. The experimentally determined limit for the machine is at 10 requests per second, which is way below our expectations. At this rate, all cpu resources are used up by mysql, which runs in its own openvz-instance. After continued stressing, the box becomes unusable and recovers slowly if the load alleviates. We have thoroughly examined the system and are still not sure what causes this behavior.

All this leads to the following question: Could open VZ be the cause?

Here are further details:

Hardware used:

The box powering the whole setup is a modern server equipped with two quad-core Opteron 2354 processors and 8GB RAM built in. Data is stored on 8 serial attached SCSI-disks forming a RAID6. All Network-adapters are gigabit speed.

Operating System:

The installed OS is Centos 5.2. The kernel used is a rebuild of 2.6.18-92.1.1.el5.028stab057 with drbd7 disabled. However, drbd is currently not used at all. The disk is partitioned in two pieces, one of which holds a lvm-group for the system and with the other one holding a second, large lvm-group for the openVZ instances. The filesystem used is ext3.

OpenVZ setup:

Each VE is located on its own LV, quotas are not used or enforced. System resources are split equally between all VEs. For networking, the veth way is used with bridges connecting the VEs to physical networks where necessary.

VEs:

Mysql(ID 105), webserver(107) and fastcgi-workers(108,109) are located in own VEs each. Management systems are also put in own VEs so that in total 8 VEs are running on the host. This gives plenty of resources to each of the VEs. The fastcgi-workers share a nfs-mount from the host.

MySQL:

The mysql-server mostly uses the config shipped with mysql for very large setups. If necessary I can post it. There are about 80 queries being executed by the database per request on the website. There is also a slave server which runs on a different machine and replicates the local dbs.

/proc/fairsched2:

Among the symptoms are the following values in /proc/fairsched2:

Version: 2.7 debug

id	weight	rate	run	cpus	flg	ready	start_tag	value	delay
2147483647	6	0	0	1	1	..	1	130843172784826	0
101	5	0	0	0	..	0	130843135560820	0	0
105	5	2048	0	0	L.	0	130843154252806	329155854336	
1597319284612829									
108	5	3072	0	0	L.	0	130843135708425	329155854336	
1505901203181114									
109	5	3072	0	0	L.	0	130843150223017	329155854336	
1505901210530697									
107	5	1024	0	0	L.	0	130843139046280	329155854336	
1608674694367369									
10	5	1024	0	0	L.	0	130841315590750	329155854336	
503020793403475									
231	5	1024	0	0	L.	0	130843138592665	329155854336	
474103127903629									

I don't know why, but on all other systems I have running with openVZ, the last column (delay) is always zero.

What we already checked:

The network throughput between VEs is approx 500mbit/sec.

The system is not swapping, there's enough free ram.

The database can handle approx 10000 SELECTs per second (1 row by id from a small table).

During simulated load, the database uses all resources at about 1500 queries per second. Those seem to be significantly more complicated and some include JOINS.

Thanks for your comments.

Regards,

Lars
