Subject: Re: [Vserver] Re: Container Test Campaign
Posted by kir on Tue, 04 Jul 2006 12:19:02 GMT
View Forum Message <> Reply to Message

Clément,

Thanks for addressing my concerns! See comments below.

Clément Calmels wrote:
> Hi,
>
>
>> 1.1 It would be nice to run vmstat (say, vmstat 10) for the duration of
>> the tests, and put the vmstat output logs to the site.
>>
>
> Our benchmark framework allows us to use oprofile during test...
> couldn't it be better than vmstat?
>
Good idea.
>> Basically, the detailed description of a process would be nice to have,
>> in order to catch possible problems. There are a lot of tiny things
>> which are influencing the results. For example, in linux kernels 2.4
>> binding the NIC IRQ to a single CPU on an SMP system boosts network
>> performance by about 15%! Sure this is not relevant here, it's just an
>> example.
>>
>
> I agree. Actually, I always try to use 'default' configuration or
> installation but I will try to describe the tests in details.
>
>> 1.3 Would be nice to have diffs between different kernel configs.
>>
> The different configs used are available in the lxc site. You will
> notice that I used a minimal config file for most of the test, but for
> Openvz I had to use the one I found in the OpenVZ site because I faced
> kernel build error (some CONFIG_NET... issues).
We are trying to eliminate those, so a bug report would be nice.
> I think that the
> differences are more dealing with network stuff.
>
>> For example, the tbench test is probably failed to finish because it
>> hits the limits for privvmpages, tcpsndbuf and tcprcvbuf. I have
>> increased the limits for those parameters and the test was finished
>> successfully. Also, dbench test could hit the disk quota limit for a VE.
>> Some more info is available at http://wiki.openvz.org/Resource_management
>>
>

> I already used this page. I had to increase 'diskinodes' and 'diskspace'
> resources in order to run some test properly (the disk errors were more
> selfexplicit).
> I'm wondering why a default 'guest' creation implies some resources
> restrictions? Couldn't the resources be unlimited? I understand the need
> for resource management, but the default values look a little bit
> tiny...
>
The reason is security. A guest is untrusted by default, though sane
limits are applied. Same as ulimit which has some sane defaults (check
output of ulimit -a). Same as those kernel settings from /proc/sys --
should /proc/sys/fs/file-max be 'unlimited' by default?

In fact, those limits are taken from a sample configuration file during
"vzctl create" stage. Sample file is specified in global OpenVZ config
file (/etc/vz/vz.conf, parameter name is CONFIGFILE, default is to take
configuration from /etc/vz/conf/ve-vps.basic.conf-sample).

There are several ways to change that default configuration:

1. (globally) Put another sample config and specify it in /etc/vz/vz.conf
2. (globally) Edit the existing sample config
(/etc/vz/conf/ve-vps.basic.conf-sample)
3. (per VE) Specify another config during vzctl create stage, like this:
vzctl create VEID [--config name]
4. (per VE) Tune the specific parameters using vzctl set [--param value
...] --save
>
>> 2.2 For OpenVZ specifically, it would be nice to collect
>> /proc/user_beancounters output before and after the test.
>>
>
> For sure... I will take a look at how integrating it in our automatic
> test environment.
>
> Best regards,
>
>