

Hi Kirill,

Thanks for the feedback. Not sure whether you are referring to Clement's work or our paper. I'll assume you are referring to our paper.

> >From what I see just just after 1 minute check of your results:
>
> DBench:
> - different disk I/O schedulers are compared. This makes
> comparison useless (not virtualization technologies are
> compared itself).

Correct. As you can image, it is not easy to pick benchmarks. The intention was simply to repeat measurements presented by the Xen paper published in 2003 Symposium of Operating Systems Principles, and show how Vserver compares on those. The point of the paper is to show how "container based virtualization" compares to hypervisor based systems. We do not have the time nor expertise to show this with OpenVZ. It would be great if someone from the openvz community could step in and show how OpenVZ compares to Xen.

> - the fact that there is too much difference between measurements
> (e.g. vserver makes linux faster :lol:)

This is an interesting, odd, and likely laughable result. We just reported what we observed. It is quite possible that we made a mistake somewhere. However, I believe that the problem lies more with the dbench benchmark than with our setup. We did try to eliminate as many variables as possible. Please take a peek on the last page of the paper to see our discussion wrt normalizing the configurations. We are open to further suggestions to eliminate further variables.

> I also noticed that you do the measurements with different HZ settings.
> This influences the results as well...

Of course. My assumption is that it would negatively affect Vserver. Are you suggesting that it can positively affect the benchmark results to run at 1000HZ vs. 100HZ, as the Xen Domains are configured to do?

> BTW, do you plan to do functional testing in addition to performance?

Please clarify what you mean here? From what I gather, the main thing that Vserver lacks is the degree of network virtualization that OpenVZ supports. Is there anything else? From my perspective, the comparison will have to be with Xen/UML rather than, say, a contest between container based systems. I say this because it appears that the majority of the LKML community is believes that container-based systems don't add much above and beyond what Xen/UML/QEMU/VMware already offer today.

Best regards,
Marc

> -----Original Message-----

> From: Kirill Korotaev [mailto:dev@openvz.org]

> Sent: Monday, July 03, 2006 3:50 AM

> To: ? Calmels

> Cc: devel@openvz.org; vserver@list.linux-vserver.org;

> sam.vilain@catalyst.net.nz; serue@us.ibm.com; DLEZCANO@fr.ibm.com;

> mef@CS.Princeton.EDU

> Subject: Re: [Devel] Container Test Campaign

>

>

> >From what I see just just after 1 minute check of your results:

>

> DBench:

> - different disk I/O schedulers are compared. This makes

> comparison useless

> (not virtualization technologies are compared itself).

> - the fact that there is too much difference between measurements

> (e.g. vserver makes linux faster :lol:) makes me believe that

> you use large disk partiion,

> where data blocks allocation on the disk influence your results.

> To make these measurements correct the same partition with the

> size closest to the required

> max disk space should be used in all DBench tests.

>

> TBench:

> - when running inside VE, please, make sure, that

> /proc/user_beancounters doesn't show you

> resource allocation fails (failcnt column).

> Resource limits set by default can be just too small to finish

> your test case.

> And this doesn't mean your conclusion 'Concerning the results,

> obviously more isolation brings more overhead.'

> I'm really suprised to see such statements.

>

> I also noticed that you do the measurements with different HZ settings.

> This influences the results as well...

>

> BTW, do you plan to do functional testing in addition to performance?

>

> Thanks,

> Kirill

>

>

> > Hi,

> >

> > A first round about virtualisation benchmarks can be found here:
> > <http://lxc.sourceforge.net/bench/>
> > These benchmarks run with vanilla kernels and the patched versions of
> > well know virtualisation solutions: VServer and OpenVZ. Some benches also
> > run inside the virtual 'guest' but we ran into trouble trying to run
> > some of them... probably virtual 'guest' configuration issues... we will
> > trying to fix them...
> > The metacluster migration solution (formely a Meiosys company product)
> > was added as it seems that the checkpoint/restart topic is close to the
> > virtualisation's one (OpenVZ now provides a checkpoint/restart
> > capability).
> > For the moment, benchmarks only ran on xeon platform but we expect more
> > architecture soon. Besides the 'classic' benches used, more network
> > oriented benches will be added. Netpipe between two virtual 'guests' for
> > example. We hope we will be able to provide results concerning the
> > virtual 'guest' scalability, running several 'guest' at the same time.
> >
> > Best regards,
> >
> >
> > Le mercredi 07 juin 2006 à 16:20 +0200, Clement Calmels a écrit :
> >
> >>Hello !
> >>
> >>I'm part of a team of IBMers working on lightweight containers and we
> >>are going to start a new test campaign. Candidates are vserver,
> >>vserver context, namespaces (being pushed upstream), openvz, mcr (our
> >>simple container dedicated to migration) and eventually xen.
> >>
> >>We will focus on the performance overhead but we are also interested in
> >>checkpoint/restart and live migration. A last topic would be how well
> >>the
> >>resource managment criteria are met, but that's extra for the moment.
> >>
> >>We plan on measuring performance overhead by comparing the results on
> >>a vanilla kernel with a partial and with a complete virtual
> >>environment. By partial, we mean the patched kernel and a 'namespace'
> >>virtualisation.
> >>
> >>Test tools
> >>-----
> >>o For network performance :
> >>
> >> * netpipe (<http://www.scl.ameslab.gov/netpipe/>)
> >> * netperf (<http://www.netperf.org/netperf/NetperfPage.html>)
> >> * tbench (<http://samba.org/ftp/tridge/dbench/README>)
> >>
> >>o Filesystem :

> >>
> >> * dbench (<http://samba.org/ftp/tridge/dbench/README>)
> >> * iiozone (<http://www.iozone.org/>)
> >>
> >>o General
> >>
> >> * kernbench (<http://ck.kolivas.org/kernbench/>) stress cpu and
> >> filesystem through kernel compilation
> >> * More 'real world' application could be used, feel free to submit
> >> candidates...
> >>
> >>We have experience on C/R and migration so we'll start with our own
> >>scenario, migrating oracle under load. The load is generated by DOTS
> >>(<http://ltp.sourceforge.net/dotshow>)We ran into trouble trying
> to run sto.php).
> >>
> >>If you could provided us some material on what has already been done :
> >>URL, bench tools, scenarios. We'll try to compile them in. configuration
> >>hints and tuning are most welcome if they are reasonable.
> >>
> >>Results, tools, scenarios will be published on lxc.sf.net . We will
> >>set up the testing environment so as to be able to accept new
> >>versions, patches, test tools and rerun the all on demand. Results,
> >>tools, scenarios will be published on lxc.sf.net.
> >>
> >>thanks !
> >>
> >>Clement,
> >>
