Subject: Re: [Vserver] Re: Container Test Campaign
Posted by Herbert Poetzl on Sun, 02 Jul 2006 15:32:29 GMT
View Forum Message <> Reply to Message

> Hi,
>
> A first round about virtualisation benchmarks can be found here:
> http://lxc.sourceforge.net/bench/

very interesting results, tx ...

> These benchmarks run with vanilla kernels and the patched versions of
> well know virtualisation solutions: VServer and OpenVZ. Some benchs
> 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 produt) was added as it seems that the
> checkpoint/restart topic is close to the virtualisation's one (OpenVZ
> now provides a checkpoint/restart capability).

from the tests:
 "For benchs inside real 'guest' nodes (OpenVZ/VServer) you should
  take into account that the FS tested is not the 'host' node one's."

at least for Linux-VServer it should not be hard to avoid the
chroot/filesystem namespace part and have it run on the host fs.
a bind mount into the guest might do the trick too, if you need
help to accomplish that, just let me know ...

> For the moment, benchmarks only ran on xeon platform but we expect
> more architecture soon. Besides the 'classic' benchs used, more
> network oriented benchs 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,
Herbert

> Best regards,
>

> > 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)
> > * iozone (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/dotshowWe 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,

> >
> --

>
> _____
> Vserver mailing list
> Vserver@list.linux-vserver.org
> http://list.linux-vserver.org/mailman/listinfo/vserver