
Subject: RE: Container Test Campaign

Posted by [Clement Calmels](#) on Thu, 22 Jun 2006 16:33:19 GMT

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

Hi,

We have currently set up tests. Our goal is to provide a huge variety of test cases and measurements (to reduce imprecise plot). We started with same kind of microbenchmark: ltp, dbench, tbench and more complex benchmarks: kernel compilation. Tests are launched outside and inside a container. Besides we want to make tests with different number of container within a real node to get some clues on the different solutions's scalabilities.

After taking a glance at your paper, it seems we got same kind of results.

It would be a good idea to share our test setup. It seems there are easy ways to disadvantage one container solution against another one (Xen using loop device instead of a dedicated LVM partition for example). Or for example, the way we "share" a node between different containers could have some consequences on test results. I would prefer "fair" benchmarks.

In my opinion, a "fight" between the different container solutions would be as useless as a "Google Fight". But finding real world cases where a solution seems better than the others may result in more accurate conclusions.

Concerning the checkpoint/restart/migration topic, we (IBM) owned a solution called Metacluster. The main goal of Metacluster was the migration issue... but as a result it brought isolation in some areas (pid for example). We will use this solution and make some performance measures during the migration of well known application (Oracle under different workloads...). Openvz and Xen should be included in such benches.

Best regards,
Clement.

Le mercredi 21 juin 2006 à 15:25 -0400, Marc E. Fiuczynski a écrit :

> Hello Clement,

>

> Sorry for the late response, as I have been on vacation.

>

> We are interested in this test campaign. Our work so far has focused on
> performance, scalability, and isolation properties of vserver compared with
> xen. My guess is that you cc'd me due to the posting of our paper comparing
> vserver with xen (attached for those of you who have not seen it yet). In
> what way can be participate/contribute (i.e., where do we start)? We could
> share our test setup (except SpecWeb 99) that we used for our paper with

> everyone. Also, we'd appreciate if the folks participating in this test
> campaign could skim our paper and give us some feedback wrt the evaluation
> section and the appendix where we describe in reasonable the kernel vars,
> lvm partition setup, etc., we've used to eliminate differences between
> systems.

>
> Best regards,
> Marc

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

> > From: Clement Calmels [mailto:clement.calmels@fr.ibm.com]
> > Sent: Wednesday, June 07, 2006 10:20 AM
> > To: devel@openvz.org; vserver@list.linux-vserver.org
> > Cc: kir@openvz.org; dev@openvz.org; sam.vilain@catalyst.net.nz;
> > mef@CS.Princeton.EDU; clg@fr.ibm.com; serue@us.ibm.com;
> > haveblue@us.ibm.com; dlezcana@fr.ibm.com
> > Subject: Container Test Campaign

> >
> >
> > 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 management 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/dotshowto.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,
