Subject: Re: Container Test Campaign
Posted by Sam Vilain on Fri, 23 Jun 2006 03:40:49 GMT
View Forum Message <> Reply to Message

Marc E. Fiuczynski wrote:
> 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.

One area it would be interesting to see benchmarks for is the performance impact of filesystem unification and a lot of vservers - for instance, a system with 10 vservers, each running apache and actively serving pages, I'd expect to see more cache hits at the L2 and/or L3 CPU cache layers on account of the fact that, eg, C libraries are not being paged out to load in other (identical) C libraries.

My guess is that you just can't leverage that kind of benefit from a hypervisor approach, but I don't really know enough about how they work under the hood to be able to say.

Sam.

>> -----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;
>> dlezcano@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 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/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,