
Subject: Re: Re: [RFC] Container mini-summit agenda for Sept 3, 2007

Posted by [Oren Laadan](#) on Fri, 31 Aug 2007 18:10:36 GMT

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

Kirill Korotaev wrote:

> Cedric Le Goater wrote:

>>>> Many of these were discussed in a recent Zap paper present in USENIX:

>>>> http://www.ncl.cs.columbia.edu/publications/usenix2007_fordist.pdf

>>>> The paper describes important design choices in Zap (but I'm biased ...).

>>>> I think it may serve as an appetizer for the discussion :P

>>> Thanks, I hope we all have time to read it.

>>

>> The abstract says :

>>

>> "...

>> Our results show checkpoint and restart times 3 to 55 times faster than

>> OpenVZ and 5 to 1100 times faster than Xen."

>>

>> I'm impressed ! :) When can we play it ?

>>

>> Thanks for the appetizer !

disclaimer(1): I'm one the other authors.

disclaimer(2): I hoped to focus on the technical issues in the paper rather than start a flame war...

> It is totally unfair to compare full virtualization solution such as OpenVZ

> with sync on VE stop (for quotas consistency) and which doesn't require shared storage for migration

> with POC which uses shared storage in the paper.

Indeed we didn't try to hack OpenVZ to get max performance, or isolate the components. For example, I suspect the long(er) restart times are mainly due to container setup time as opposed to the restoration of the processes. I even expressed this speculation during my talk.

Note the distinction between "migration" and "checkpoint-restart", as they are not the same. Generally if you have c/r you can migrate. The performance evaluation in the paper compared checkpoint and restart, not the migration from one machine to another.

BTW, Zap **does** sync selected files (in particular shared mapping of files). There are some optimization that move similar overhead out of the critical path to reduce the downtime.

> I'm not sure why author didn't pay attention to these HUGE differences in configuration...

> Maybe because 1100x times is so incredible :@)

It's always difficult to compare HW and OS virtualization approaches (in our terminology, HW approach - Xen, VMware, KVM etc .. they all work under the guest OS, while OS approach - OpenVZ, Vserver, Zap .. work with the OS, under the application). Both approaches have pros and cons. It's even hard to compare between OS virtualization (aka containers) approaches because they require different configs and provide different features. For example, Zap is not geared to provide performance isolation.

The quotes comparison highlights the advantages of working at the application (or container) level; For example if you want to run apache in an isolated environment and be able to c/r or migrate: you choose a system (linux, Zap, OpenVZ, Xen etc), and that implies some sort of configuration. Each method has goods and bads, be in in functionality, overhead, ease-of-use and so forth. In the specific case of running apache, this comparison gives an idea of the gain/loss in terms of checkpoint/restart performance.

Oren.

Containers mailing list
Containers@lists.linux-foundation.org
<https://lists.linux-foundation.org/mailman/listinfo/containers>
