

---

Subject: Re: Checkpoint/Restart mini-summit  
Posted by [serue](#) on Wed, 16 Jul 2008 15:15:30 GMT  
[View Forum Message](#) <> [Reply to Message](#)

---

Quoting Eric W. Biederman (ebiederm@xmission.com):

> Daniel Lezcano <dlezcano@fr.ibm.com> writes:

>

> > Hi all,

> >

> > Here is a proposition a more detailed agenda for the checkpoint/restart  
> > mini-summit. If everybody is ok with it, I will update the wiki.

> >

> > Comments are welcome :)

>

> A reading list is useful, even to help get some ideas circulating  
> before we get there.

>

> Ultimately the technical details will need to be resolve by  
> people discussing things and sending patches back and forth  
> on the mailing lists.

>

> I don't think a detailed agenda is going to get us anywhere.

> Especially not one focused on the implementation details.

Right, the whole point of Daniel including a 'reading list' was just so that we can avoid wasting time discussing existing implementations. So he wasn't suggesting that we would be discussing those in detail, in fact quite the opposite.

> I think we need to start by seeing what we can agree on. Certainly we  
> agree that checkpoint/restart needs to be part of the picture. What  
> are the problems that the linux community can solve with  
> checkpoint/restart.

>

> Then we need to talk about what kind of implementation we want to  
> merge into mainline. How do we sell it, and how do we implement  
> it without affecting long term maintainability.

>

> I think the granularity of our operations, and what state we  
> save is important. I don't think how we save it is important  
> unless it affects one of our requirements.

>

> As for the posix draft and the historical Cray & SGI implementations.  
> They were on the wrong track. The did not have namespace support  
> so they could not in general restore their checkpoints.

>

> There are also a lot of things you have failed to touch on, that  
> I'm not going to go into now.

>  
> With any luck the mini-summit before OLS will be the start of a  
> conversation that will go on all week, and continue on the mailing  
> lists.

Agreed.

This could be tough to pull off, but if we can walk out of there with a short focused list of coding todos with the intent of pumping out patches by the end of OLS, turning OLS into a bit of a hack-fest, that would imo be great.

(But then that's precisely what I like to do at conferences - sit by some wall and pick something completely new to code, while once in awhile getting up to chat or see a talk. Does that make me anti-social?)

> The real question is how do we coordinate our efforts to build a good  
> linux checkpoint/restart implementation.  
>  
> > \* Documentation  
> > \* Zap : [www.ncl.cs.columbia.edu/publications/userix2007\\_fordist.pdf](http://www.ncl.cs.columbia.edu/publications/userix2007_fordist.pdf)  
> > \* Metacluster : [lxc.sourceforge.net/doc/ols2006/lxc-ols2006.pdf](http://lxc.sourceforge.net/doc/ols2006/lxc-ols2006.pdf)  
> > \* OpenVZ : [http://wiki.openvz.org/Checkpointing\\_and\\_live\\_migration](http://wiki.openvz.org/Checkpointing_and_live_migration)  
> > \* Checkpoint/Restart technology :  
> > [http://en.wikipedia.org/wiki/Application\\_checkpointing](http://en.wikipedia.org/wiki/Application_checkpointing)  
> > \* Virtual Servers and Checkpoint/Restart in Mainstream Linux : Sigops  
> > document  
>  
> There is also the classic emacs undump.  
> The very simple vmadump from bproc.  
>  
> Eric  
> \_\_\_\_\_  
> Containers mailing list  
> Containers@lists.linux-foundation.org  
> <https://lists.linux-foundation.org/mailman/listinfo/containers>

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

---