

---

Subject: OLS Checkpoint/Restart BOF Summary  
Posted by [Dave Hansen](#) on Fri, 06 Jul 2007 16:35:26 GMT  
[View Forum Message](#) <> [Reply to Message](#)

---

First of all, thanks to everyone that attended the BOF. It was very productive.

First thing, does anyone have a small, working checkpoint/restart implementation that is mostly or all in userspace that is easy to build on and extend? If not, I'll probably continue hacking on the one that I have, or go try to steal something from one of the other projects.

One of the developers of bproc attended. They use an in-kernel mechanism called vmadump to migrate processes between nodes in a cluster. However, they only move processes in very particular states, such as with no sockets open.

Eric Focht from NEC attended. They're using a mechanism called KDDM: [http://www.linuxsymposium.org/2007/view\\_abstract.php?content\\_key=257](http://www.linuxsymposium.org/2007/view_abstract.php?content_key=257) to share kernel structures and user memory (among other things) across a cluster. This allows processes to be migrated, but requires heavy kernel modification?

Several OpenVZ/Virtuozzo developers attended.

Some Kerlabs developer attended.

Everyone seems to want to do as much as possible of checkpoint/restart in the kernel. Dave Hansen lobbied for doing as much in userspace as possible, at least until we've proved exactly what is really hard (or impossible) in userspace but easier in the kernel. That way we have some actual justification in order to get patches merged into mainline. Surprisingly, no one really disagreed with this sentiment.

Does anyone want to say something about networking developments?

-- Dave

---

Containers mailing list  
[Containers@lists.linux-foundation.org](mailto:Containers@lists.linux-foundation.org)  
<https://lists.linux-foundation.org/mailman/listinfo/containers>

---