

Re: [ANNOUNCE] Minneapolis Cluster Summit, July 29-30

Source: <http://linux.derkeiler.com/Mailing-Lists/Kernel/2004-07/2356.html>

From: Lars Marowsky-Bree (*lmb_at_suse.de*)

Date: 07/12/04

Date: Mon, 12 Jul 2004 15:13:12 +0200
To: Arjan van de Ven <arjanv@redhat.com>

On 2004-07-12T14:01:27,
Arjan van de Ven <arjanv@redhat.com> said:

> *I'm not convinced that's a good idea, in that it exposes what is*
> *basically VM internals to userspace, which then would become a*
> *set-in-stone interface....*

But I'm also not a big fan of moving all HA relevant infrastructure into the kernel. Membership and DLM are the first ones; then follows messaging (and reliable and globally ordered messaging is somewhat complex – but if one node is slow, it will hurt global communication too, so...), next someone argues that a node always must be able to report which resources it holds and fence other nodes even under memory pressure, and there goes the cluster resource manager and fencing subsystem into the kernel too etc...

Where's the border?

And what can we do to make critical user-space infrastructure run reliably and with deterministic-enough & low latency instead of moving it all into the kernel?

Yes, the kernel solves these problems right now, but is that really the path we want to head down? Maybe it is, I'm not sure, afterall we also have the entire regular network stack in the kernel, but maybe also it is not.

Sincerely,
Lars Marowsky-Brée <lmb@suse.de>

--
High Availability & Clustering | ever tried. ever failed. no matter.
SUSE Labs, Research and Development | try again. fail again. fail better.
SUSE LINUX AG - A Novell company | -- Samuel Beckett

-
To unsubscribe from this list: send the line "unsubscribe linux-kernel" in the body of a message to majordomo@vger.kernel.org

Linux-Kernel: Re: [ANNOUNCE] Minneapolis Cluster Summit, July 29-30

More majordomo info at <http://vger.kernel.org/majordomo-info.html>
Please read the FAQ at <http://www.tux.org/lkml/>