

Re: [ANN] Userspace M-on-N threading model implementation. Alpha release.

Re: [ANN] Userspace M-on-N threading model implementation. Alpha release.

Source: <http://linux.derkeiler.com/Mailing-Lists/Kernel/2007-01/msg07967.html>

- *From:* Evgeniy Polyakov <johnpol@xxxxxxxxxxxxx>
 - *Date:* Tue, 30 Jan 2007 13:43:04 +0300
-

On Tue, Jan 30, 2007 at 11:24:51AM +0100, Samuel Thibault (samuel.thibault@xxxxxxxxxxxxx) wrote:

Evgeniy Polyakov, le Tue 30 Jan 2007 12:53:16 +0300, a écrit :

You may want to have a look at some existing implementations:

I saw most of them.

As far as I recall, only PTL (is not shown here) has preemptible scheduler. NTL has it too, but is based on different approach.

Marcel has a preemptible scheduler too, based on signals and setjmp/longjmp.

Do some documentation and benchmarks exist for that library – site seems to only described environment is was created for? How does blocking problem solved?

Samuel

--

Evgeniy Polyakov

-

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

More majordomo info at <http://vger.kernel.org/majordomo-info.html>

Please read the FAQ at <http://www.tux.org/lkml/>