
Subject: Re: [PATCH (resubmit)][BRIDGE] Properly dereference the
br_should_route_hook

Posted by [paulmck](#) on Fri, 30 Nov 2007 01:25:01 GMT

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

On Fri, Nov 30, 2007 at 10:49:00AM +1100, Herbert Xu wrote:

> On Thu, Nov 29, 2007 at 06:36:50AM -0800, Paul E. McKenney wrote:

> >

> > That certainly is an interesting tradeoff... Save a memory barrier
> > when assigning NULL, but pay an extra test and branch in all cases.
> > Though it does make for a simpler rule -- just use rcu_assign_pointer()
> > in all cases. Of course, if almost all rcu_assign_pointer() executions
> > assign non-NULL pointers, the optimal strategy would be to leave the
> > implementation of rcu_assign_pointer() alone, and simply enforce use
> > of rcu_assign_pointer(), even if the pointer being assigned is NULL.

>

> I was thinking of something much simpler. If the second argument is
> constant and NULL, then skip the barrier. No run-time slow-down at
> all.

That certainly makes a lot of sense!!!

You have in mind something like the following?

```
#define rcu_assign_pointer(p, v) \
({ \
    if (!__builtin_constant_p(v) || \
        ((v) != NULL)) \
        smp_wmb(); \
    (p) = (v); \
})
```

If so, I will do some testing and submit a patch. Probably to Gautham's preemptible-RCU patchset to avoid gratuitously complicating his life, especially given that he very graciously agreed to take it over from me. We should be able to live with the overhead in the meantime. ;-)

Thanx, Paul

> > Although rcu_dereference() does a memory barrier only on Alpha, that of
> > rcu_assign_pointer() is needed on any machine that does not preserve store
> > order (Itanium, POWER, ARM, some MIPS boxes according to rumor, ...).

>

> Good point!

>

> Thanks,

> --

> Visit Openswan at <http://www.openswan.org/>

> Email: Herbert Xu ~{PmV>HI~} <herbert@gondor.apana.org.au>
> Home Page: <http://gondor.apana.org.au/~herbert/>
> PGP Key: <http://gondor.apana.org.au/~herbert/pubkey.txt>
