
Subject: Re: [ckrm-tech] [RFC][PATCH][2/4] Add RSS accounting and control
Posted by [Paul Menage](#) on Mon, 19 Feb 2007 12:09:38 GMT

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

On 2/19/07, Balbir Singh <balbir@in.ibm.com> wrote:

> >

> > More worrisome is the potential for use-after-free. What prevents the
> > pointer at mm->container from referring to freed memory after we're dropped
> > the lock?

> >

>

> The container cannot be freed unless all tasks holding references to it are
> gone,

... or have been moved to other containers. If you're not holding
task->alloc_lock or one of the container mutexes, there's nothing to
stop the task being moved to another container, and the container
being deleted.

If you're in an RCU section then you can guarantee that the container
(that you originally read from the task) and its subsystems at least
won't be deleted while you're accessing them, but for accounting like
this I suspect that's not enough, since you need to be adding to the
accounting stats on the correct container. I think you'll need to hold
mm->container_lock for the duration of memctl_update_rss()

Paul
