Subject: Re: [ckrm-tech] [PATCH] BC: resource beancounters (v4) (added user memory)
Posted by Rohit Seth on Tue, 12 Sep 2006 18:02:18 GMT
View Forum Message <> Reply to Message

On Tue, 2006-09-12 at 23:10 +0530, Srivatsa Vaddagiri wrote:
> On Tue, Sep 12, 2006 at 10:22:32AM -0700, Rohit Seth wrote:
> > On Tue, 2006-09-12 at 16:14 +0530, Srivatsa Vaddagiri wrote:
> > > On Mon, Sep 11, 2006 at 12:10:31PM -0700, Rohit Seth wrote:
> > > > It seems that a single notion of limit should suffice, and that limit
> > > > should more be treated as something beyond which that resource
> > > > consumption in the container will be throttled/not_allowed.
> > >
> > > The big question is : are containers/RG allowed to use *upto* their
> > > limit always? In other words, will you typically setup limits such that
> > > sum of all limits = max resource capacity?
> > >
> >
> > If a user is really interested in ensuring that all scheduled jobs (or
> > containers) get what they have asked for (guarantees) then making the
> > sum of all container limits equal to total system limit is the right
> > thing to do.
> >
> > > If it is setup like that, then what you are considering as limit is
> > > actually guar no?
> > >
> > Right.  And if we do it like this then it is up to sysadmin to configure
> > the thing right without adding additional logic in kernel.
>
> Perhaps calling it as "limit" in confusing then (otoh it may go down well
> with Linus!).

It is similar to what ulimit generally implies.  It will definitely help getting container code in mainline if we can minimize the changes in mainline but still providing some meaningful functionality.

CKRM is excellent piece of code but I think its complexity is what making it sit out side the mainline for such a long period of time.

>  I perhaps agree we need to go with one for now (in the
> interest of making some progress), but we probably will come back to
> this at a later point.

Indeed if there are still some requirements.

-rohit