Subject: Re: [ckrm-tech] [PATCH] BC: resource beancounters (v4) (added user memory)
Posted by Rohit Seth on Wed, 13 Sep 2006 00:43:23 GMT
View Forum Message <> Reply to Message

On Tue, 2006-09-12 at 17:02 -0700, Chandra Seetharaman wrote:
> On Tue, 2006-09-12 at 10:22 -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.
>
> It won't be a complete solution, as the user won't be able to
> - set both guarantee and limit for a resource group
> - use limit on some and guarantee on some
> - optimize the usage of available resources

I think, if we have some of the dynamic resource limit adjustments
possible then some of the above functionality could be achieved. And I
think that could be a good start point.

-rohit