

---

Subject: Re: [ckrm-tech] [PATCH] BC: resource  
beancounters (v4) (added user memory)  
Posted by [Rohit Seth](#) on Thu, 14 Sep 2006 01:27:12 GMT  
[View Forum Message](#) <> [Reply to Message](#)

---

On Wed, 2006-09-13 at 15:24 -0700, Chandra Seetharaman wrote:  
> On Tue, 2006-09-12 at 18:33 -0700, Rohit Seth wrote:  
> > On Tue, 2006-09-12 at 18:13 -0700, Chandra Seetharaman wrote:  
> > > On Tue, 2006-09-12 at 17:43 -0700, Rohit Seth wrote:  
> > > <snip>  
> > >  
> > > > 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.  
> > >  
> > >  
> > > Yes, dynamic resource adjustments should be available. But, you can't  
> > > expect the sysadmin to sit around and keep tweaking the limits so as to  
> > > achieve the QoS he wants. (Even if you have an application sitting and  
> > > doing it, as I pointed in other email it may not be possible for  
> > > different scenarios).  
> > > >  
> > >  
> > > As said earlier, if strict QoS is desired then system should be  
> > > appropriately partitioned so that the sum of limits doesn't exceed  
> > > physical limit (that is cost of QoS). Let us first get at least that  
> > > much settled on and accepted in mainline before getting into these  
> > > esoteric features.  
> > >  
> > >  
> > > esoteric ?! Please look at the different operating system that provide  
> > > resource management and other resource management capability providers.  
> > > All of them have both guarantees and limits (they might call them  
> > > differently).  
> > >

Is this among the very first features you would like (to get absolutely  
right) before containers get in mm tree? Or is this something that can  
wait after the minimal infrastructure is in Andrew's tree and the code  
gets wider testing...And above all we have agreed upon user interface.

-rohit

---