
Subject: Re: [PATCH 1/1, v6] cgroup/freezer: add per freezer duty ratio control
Posted by [Vaidyanathan Srinivas](#) on Mon, 14 Feb 2011 18:03:54 GMT
[View Forum Message](#) <> [Reply to Message](#)

* Arjan van de Ven <arjan@linux.intel.com> [2011-02-09 19:06:15]:

> On 2/9/2011 7:04 PM, Matt Helsley wrote:
> >On Tue, Feb 08, 2011 at 05:05:41PM -0800, jacob.jun.pan@linux.intel.com wrote:
> >>From: Jacob Pan<jacob.jun.pan@linux.intel.com>
> >>
> >>Freezer subsystem is used to manage batch jobs which can start
> >>stop at the same time. However, sometime it is desirable to let
> >>the kernel manage the freezer state automatically with a given
> >>duty ratio.
> >>For example, if we want to reduce the time that backgroup apps
> >>are allowed to run we can put them into a freezer subsystem and
> >>set the kernel to turn them THAWED/FROZEN at given duty ratio.
> >>
> >>This patch introduces two file nodes under cgroup
> >>freezer.duty_ratio_pct and freezer.period_sec
> >>
> >>Usage example: set period to be 5 seconds and frozen duty ratio 90%
> >>[root@localhost aoa]# echo 90> freezer.duty_ratio_pct
> >>[root@localhost aoa]# echo 5000> freezer.period_ms
> >I kept wondering how this was useful when we've got the "cpu" subsystem
> >because for some reason "duty cycle" made me think this was a scheduling
> >policy knob. In fact, I'm pretty sure it is -- it just happens to
> >sometimes reduce power consumption.
> >
> >Have you tried using the cpu cgroup subsystem's share to see if it can
> >have a similar effect?
>
> does the cpu cgroup system work on a 20 to 30 second time window?
> the objective is to have the CPU idle, without wakeups, for that long...
> (to save power)

This is an interesting idea to force idle. The cpu cgroup will maintain resource ratio but will not restrict runtime of a cgroup if there is nothing else to run in the system.

CFS hardlimits (<http://lwn.net/Articles/368685/>) can do something like this but will need to be tuned for long intervals. On multi cpu system, synchronising the idle times across cpus has been the key challenge that reduces the power saving benefits.

Does this technique provide good power savings for a specific usecase/workload or platform?

--Vaidy

Containers mailing list

Containers@lists.linux-foundation.org

<https://lists.linux-foundation.org/mailman/listinfo/containers>
