
Subject: Re: [PATCH -mm] Hook up group scheduler with control groups
Posted by [Randy Dunlap](#) on Thu, 27 Sep 2007 17:40:05 GMT
[View Forum Message](#) <> [Reply to Message](#)

On Thu, 27 Sep 2007 20:36:24 +0530 Srivatsa Vaddagiri wrote:

> Andrew,
> This is a resend of the patch I had sent earlier at:
>
> <http://marc.info/?l=linux-kernel&m=119065506607858>
>
> This patch enables group cpu scheduler feature to work with control
> groups.
>
> Could you include this in -mm for folks to test it?

Hi :)

Uh, a few of my previous comments weren't fixed... (below)

> --
>
> Enable "cgroup" (formerly containers) based fair group scheduling.
> This will let administrator create arbitrary groups of tasks (using
> "cgroup" pseudo filesystem) and control their cpu bandwidth usage.
>
> Signed-off-by : Srivatsa Vaddagiri <vatsa@linux.vnet.ibm.com>
> Signed-off-by : Dhaval Giani <dhaval@linux.vnet.ibm.com>
>
> ---
> include/linux/cgroup_subsys.h | 6 ++
> init/Kconfig | 24 +++++---
> kernel/sched.c | 122 +++
> 3 files changed, 145 insertions(+), 7 deletions(-)
>

> Index: current/init/Kconfig
> =====
> --- current.orig/init/Kconfig
> +++ current/init/Kconfig
> @@ -319,6 +319,13 @@ config CPUSETS
>
> Say N if unsure.
>
> +config RESOURCE_COUNTERS
> + bool "Resource counters"
> + help

> + This option enables controller independent resource accounting
> + infrastructure that works with cgroups

Use tab + 2 spaces consistently for help text indentation.
End that sentence with a ".".

> + depends on CGROUPS
> +
> config FAIR_GROUP_SCHED
> bool "Fair group CPU scheduler"
> default y

~Randy
Phaedrus says that Quality is about caring.

Containers mailing list
Containers@lists.linux-foundation.org
<https://lists.linux-foundation.org/mailman/listinfo/containers>
