Subject: Re: [PATCH 5/5] Add fair "control groups" scheduler
Posted by Randy Dunlap on Mon, 24 Sep 2007 16:58:15 GMT
View Forum Message <> Reply to Message

On Mon, 24 Sep 2007 22:11:59 +0530 Srivatsa Vaddagiri wrote:

> 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
> @@ -327,13 +327,6 @@ config FAIR_GROUP_SCHED
> This feature lets cpu scheduler recognize task groups and control cpu
> bandwidth allocation to such task groups.
>
> -config RESOURCE_COUNTERS
> - bool "Resource counters"
> - help
> - This option enables controller independent resource accounting
> - infrastructure that works with cgroups
> - depends on CGROUPS
> -
> choice
> depends on FAIR_GROUP_SCHED
> prompt "Basis for grouping tasks"
> @@ -345,8 +338,25 @@ choice
> This option will choose userid as the basis for grouping
> tasks, thus providing equal cpu bandwidth to each user.
>
> + config FAIR_CGROUP_SCHED
> + bool "Control groups"
> + depends on CGROUPS
> + help
> + This option allows you to create arbitrary task groups
> + using the "cgroup" psuedo filesystem and control
> + the cpu bandwidth allocated to each such task group.
> + Refer to Documentation/cgroups.txt for more information
> + on "cgroup" psuedo filesystem.

Too much indentation.

s/cpu/CPU/g please.

> +

> endchoice
>
> +config RESOURCE_COUNTERS
> + bool "Resource counters"
> + help
> +   This option enables controller independent resource accounting
> +   infrastructure that works with cgroups

Use tab + 2 spaces to indent help text.

> + depends on CGROUPS
> +
> config SYSFS_DEPRECATED
>   bool "Create deprecated sysfs files"
>   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