
Subject: [PATCH 4/9] define group operations
Posted by [Srivatsa Vaddagiri](#) on Thu, 12 Apr 2007 17:56:08 GMT
[View Forum Message](#) <> [Reply to Message](#)

Define these operations for a task-group:

- create new group
- destroy existing group
- set bandwidth (quota) for a group
- get bandwidth (quota) of a group

Signed-off-by : Srivatsa Vaddagiri <vatsa@in.ibm.com>

1 file changed, 71 insertions(+)

```
diff -puN kernel/sched.c~grp_ops kernel/sched.c
--- linux-2.6.20/kernel/sched.c~grp_ops 2007-04-11 12:00:50.000000000 +0530
+++ linux-2.6.20-vatsa/kernel/sched.c 2007-04-12 11:07:17.000000000 +0530
@@ -7379,3 +7379,74 @@ void set_curr_task(int cpu, struct task_
 }

#endif
+
+#ifdef CONFIG_CPUMETER
+
+/* Allocate runqueue structures for the new task-group */
+static int sched_create_group(void)
+{
+    struct task_grp *tg;
+    struct task_grp_rq *tgrq;
+    int i;
+
+    tg = kzalloc(sizeof(*tg), GFP_KERNEL);
+    if (!tg)
+        return -ENOMEM;
+
+    tg->ticks = CPU_CONTROL_SHORT_WINDOW;
+    tg->long_ticks = NUM_LONG_TICKS;
+
+    for_each_possible_cpu(i) {
+        tgrq = kzalloc(sizeof(*tgrq), GFP_KERNEL);
+        if (!tgrq)
+            goto oom;
+    }
}
```

```

+ tg->rq[i] = tgrq;
+ task_grp_rq_init(tgrq, tg);
+ }
+
+ /* A later patch will make 'tg' accessible beyond this function */
+ return 0;
+oom:
+ while (i--)
+ kfree(tg->rq[i]);
+
+ kfree(tg);
+ return -ENOMEM;
+}
+
+/* Deallocate runqueue structures */
+static void sched_destroy_group(struct task_grp *tg)
+{
+ int i;
+
+ for_each_possible_cpu(i)
+ kfree(tg->rq[i]);
+
+ kfree(tg);
+}
+
+/* Assign quota to this group */
+static int sched_set_quota(struct task_grp *tg, int quota)
+{
+ int i;
+
+ tg->ticks = (quota * CPU_CONTROL_SHORT_WINDOW) / 100;
+
+ for_each_possible_cpu(i)
+ tg->rq[i]->ticks = tg->ticks;
+
+ return 0;
+}
+
+static inline int cpu_quota(struct task_grp *tg)
+{
+ return (tg->ticks * 100) / CPU_CONTROL_SHORT_WINDOW;
+}
+
+/* Return assigned quota for this group */
+static int sched_get_quota(struct task_grp *tg)
+{
+ return cpu_quota(tg);
+}

```

+
+##endif /* CONFIG_CPUMETER */

--

Regards,
vatsa

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