
Subject: [RFC][patch 3/12][CFQ-cgroup] Introduce cgroup subsystem

Posted by [Satoshi UCHIDA](#) on Thu, 03 Apr 2008 07:12:54 GMT

[View Forum Message](#) <> [Reply to Message](#)

This patch introduces a simple cgroup subsystem.

New cgroup subsystem is called cfq_cgroup.

Signed-off-by: Satoshi UCHIDA <uchida@ap.jp.nec.com>

```
block/Makefile          | 1 +
block/cfq-cgroup.c      | 57 +++++
include/linux/cgroup_subsys.h | 6 +
3 files changed, 64 insertions(+), 0 deletions(-)
create mode 100644 block/cfq-cgroup.c
```

```
diff --git a/block/Makefile b/block/Makefile
```

```
index 5a43c7d..ea07b46 100644
```

```
--- a/block/Makefile
```

```
+++ b/block/Makefile
```

```
@@ -11,6 +11,7 @@ obj-$(CONFIG_IOSCHED_NOOP) += noop-iosched.o
```

```
obj-$(CONFIG_IOSCHED_AS) += as-iosched.o
```

```
obj-$(CONFIG_IOSCHED_DEADLINE) += deadline-iosched.o
```

```
obj-$(CONFIG_IOSCHED_CFQ) += cfq-iosched.o
```

```
+obj-$(CONFIG_CGROUP_CFQ) += cfq-cgroup.o
```

```
obj-$(CONFIG_BLK_DEV_IO_TRACE) += blktrace.o
```

```
obj-$(CONFIG_BLOCK_COMPAT) += compat_ioctl.o
```

```
diff --git a/block/cfq-cgroup.c b/block/cfq-cgroup.c
```

```
new file mode 100644
```

```
index 0000000..de00a0d
```

```
--- /dev/null
```

```
+++ b/block/cfq-cgroup.c
```

```
@@ -0,0 +1,57 @@
```

```
+/*
```

```
+ * CFQ CGROUP disk scheduler.
```

```
+ *
```

```
+ * This program is a wrapper program that is
```

```
+ * extend CFQ disk scheduler for handling
```

```
+ * cgroup subsystem.
```

```
+ *
```

```
+ * This program is based on original CFQ code.
```

```
+ *
```

```
+ * Copyright (C) 2008 Satoshi UCHIDA <s-uchida@ap.jp.nec.com>
```

```
+ * and NEC Corp.
```

```
+ */
```

```
+
```

```
+#include <linux/blkdev.h>
```

```

+#include <linux/cgroup.h>
+#include <linux/cfq-iosched.h>
+
+struct cfq_cgroup {
+ struct cgroup_subsys_state css;
+};
+
+
+static inline struct cfq_cgroup *cgroup_to_cfq_cgroup(struct cgroup *cont)
+{
+ return container_of(cgroup_subsys_state(cont, cfq_subsys_id),
+ struct cfq_cgroup, css);
+}
+
+static struct cgroup_subsys_state *
+cfq_cgroup_create(struct cgroup_subsys *ss, struct cgroup *cont)
+{
+ struct cfq_cgroup *cfqc;
+
+ if (!capable(CAP_SYS_ADMIN))
+ return ERR_PTR(-EPERM);
+
+ if (!cgroup_is_descendant(cont))
+ return ERR_PTR(-EPERM);
+
+ cfqc = kzalloc(sizeof(struct cfq_cgroup), GFP_KERNEL);
+ if (unlikely(!cfqc))
+ return ERR_PTR(-ENOMEM);
+
+ return &cfqc->css;
+}
+
+static void cfq_cgroup_destroy(struct cgroup_subsys *ss, struct cgroup *cont)
+{
+ kfree(cgroup_to_cfq_cgroup(cont));
+}
+
+struct cgroup_subsys cfq_subsys = {
+ .name = "cfq",
+ .create = cfq_cgroup_create,
+ .destroy = cfq_cgroup_destroy,
+ .subsys_id = cfq_subsys_id,
+};
diff --git a/include/linux/cgroup_subsys.h b/include/linux/cgroup_subsys.h
index 1ddebfc..217ecb9 100644
--- a/include/linux/cgroup_subsys.h
+++ b/include/linux/cgroup_subsys.h
@@ -42,3 +42,9 @@ SUBSYS(mem_cgroup)

```

#endif

/* */

+

+#ifdef CONFIG_CGROUP_CFQ

+SUBSYS(cfq)

+#endif

+

+/* */

--

1.5.4.1

Containers mailing list

Containers@lists.linux-foundation.org

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