

---

Subject: [PATCH 02/10] Task Containers(V11): Add tasks file interface

Posted by [Paul Menage](#) on Fri, 20 Jul 2007 18:31:54 GMT

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

---

This patch adds the per-directory "tasks" file for containerfs mounts; this allows the user to determine which tasks are members of a container by reading a container's "tasks", and to move a task into a container by writing its pid to its "tasks".

Signed-off-by: Paul Menage <[menage@google.com](mailto:menage@google.com)>

---

```
include/linux/container.h | 10 +
kernel/container.c       | 359 ++++++=====
2 files changed, 368 insertions(+), 1 deletion(-)
```

Index: container-2.6.22-rc6-mm1/include/linux/container.h

```
=====
--- container-2.6.22-rc6-mm1.orig/include/linux/container.h
+++ container-2.6.22-rc6-mm1/include/linux/container.h
@@ -144,6 +144,16 @@ int container_is_removed(const struct co
int container_path(const struct container *cont, char *buf, int buflen);

+int __container_task_count(const struct container *cont);
+static inline int container_task_count(const struct container *cont)
+{
+    int task_count;
+    rCU_read_lock();
+    task_count = __container_task_count(cont);
+    rCU_read_unlock();
+    return task_count;
+}
+
/* Return true if the container is a descendant of the current container */
int container_is_descendant(const struct container *cont);
```

Index: container-2.6.22-rc6-mm1/kernel/container.c

```
=====
--- container-2.6.22-rc6-mm1.orig/kernel/container.c
+++ container-2.6.22-rc6-mm1/kernel/container.c
@@ -40,7 +40,7 @@ @
#include <linux/magic.h>
#include <linux/spinlock.h>
#include <linux/string.h>
-
+#include <linux/sort.h>
#include <asm/atomic.h>
```

```

/* Generate an array of container subsystem pointers */
@@ -704,6 +704,127 @@ int container_path(const struct container
    return 0;
}

+/*
+ * Return the first subsystem attached to a container's hierarchy, and
+ * its subsystem id.
+ */
+
+static void get_first_subsys(const struct container *cont,
+    struct container_subsys_state **css, int *subsys_id)
+{
+    const struct containerfs_root *root = cont->root;
+    const struct container_subsys *test_ss;
+    BUG_ON(list_empty(&root->subsys_list));
+    test_ss = list_entry(root->subsys_list.next,
+        struct container_subsys, sibling);
+    if (css) {
+        *css = cont->subsys[test_ss->subsys_id];
+        BUG_ON(!*css);
+    }
+    if (subsys_id)
+        *subsys_id = test_ss->subsys_id;
+}
+
+/*
+ * Attach task 'tsk' to container 'cont'
+ *
+ * Call holding container_mutex. May take task_lock of
+ * the task 'pid' during call.
+ */
+
+static int attach_task(struct container *cont, struct task_struct *tsk)
+{
+    int retval = 0;
+    struct container_subsys *ss;
+    struct container *oldcont;
+    struct css_group *cg = &tsk->containers;
+    struct containerfs_root *root = cont->root;
+    int i;
+    int subsys_id;
+
+    get_first_subsys(cont, NULL, &subsys_id);
+
+    /* Nothing to do if the task is already in that container */
+    oldcont = task_container(tsk, subsys_id);
+    if (cont == oldcont)

```

```

+ return 0;
+
+ for_each_subsys(root, ss) {
+ if (ss->can_attach) {
+ retval = ss->can_attach(ss, cont, tsk);
+ if (retval) {
+ return retval;
+ }
+ }
+ }
+
+ task_lock(tsk);
+ if (tsk->flags & PF_EXITING) {
+ task_unlock(tsk);
+ return -ESRCH;
+ }
+ /* Update the css_group pointers for the subsystems in this
+ * hierarchy */
+ for (i = 0; i < CONTAINER_SUBSYS_COUNT; i++) {
+ if (root->subsys_bits & (1ull << i)) {
+ /* Subsystem is in this hierarchy. So we want
+ * the subsystem state from the new
+ * container. Transfer the refcount from the
+ * old to the new */
+ atomic_inc(&cont->count);
+ atomic_dec(&cg->subsys[i]->container->count);
+ rcu_assign_pointer(cg->subsys[i], cont->subsys[i]);
+ }
+ }
+ task_unlock(tsk);
+
+ for_each_subsys(root, ss) {
+ if (ss->attach) {
+ ss->attach(ss, cont, oldcont, tsk);
+ }
+ }
+
+ synchronize_rcu();
+ return 0;
+}
+
+/*
+ * Attach task with pid 'pid' to container 'cont'. Call with
+ * container_mutex, may take task_lock of task
+ */
+static int attach_task_by_pid(struct container *cont, char *pidbuf)
+{
+ pid_t pid;

```

```

+ struct task_struct *tsk;
+ int ret;
+
+ if (sscanf(pidbuf, "%d", &pid) != 1)
+ return -EIO;
+
+ if (pid) {
+ rCU_read_lock();
+ tsk = find_task_by_pid(pid);
+ if (!tsk || tsk->flags & PF_EXITING) {
+ rCU_read_unlock();
+ return -ESRCH;
+ }
+ get_task_struct(tsk);
+ rCU_read_unlock();
+
+ if ((current->euid) && (current->euid != tsk->uid)
+ && (current->euid != tsk->suid)) {
+ put_task_struct(tsk);
+ return -EACCES;
+ }
+ } else {
+ tsk = current;
+ get_task_struct(tsk);
+ }
+
+ ret = attach_task(cont, tsk);
+ put_task_struct(tsk);
+ return ret;
+}
+
/* The various types of files and directories in a container file system */

```

```

enum container_filetype {
@@ -712,6 +833,55 @@ enum container_filetype {
FILE_TASKLIST,
};

+static ssize_t container_common_file_write(struct container *cont,
+    struct cftype *cft,
+    struct file *file,
+    const char __user *userbuf,
+    size_t nbytes, loff_t *unused_ppos)
+{
+ enum container_filetype type = cft->private;
+ char *buffer;
+ int retval = 0;
+

```

```

+ if ( nbytes >= PATH_MAX)
+ return -E2BIG;
+
+ /* +1 for nul-terminator */
+ buffer = kmalloc(nbytes + 1, GFP_KERNEL);
+ if (buffer == NULL)
+ return -ENOMEM;
+
+ if (copy_from_user(buffer, userbuf, nbytes)) {
+ retval = -EFAULT;
+ goto out1;
+ }
+ buffer[nbytes] = 0; /* nul-terminate */
+
+ mutex_lock(&container_mutex);
+
+ if (container_is_removed(cont)) {
+ retval = -ENODEV;
+ goto out2;
+ }
+
+ switch (type) {
+ case FILE_TASKLIST:
+ retval = attach_task_by_pid(cont, buffer);
+ break;
+ default:
+ retval = -EINVAL;
+ goto out2;
+ }
+
+ if (retval == 0)
+ retval = nbytes;
+out2:
+ mutex_unlock(&container_mutex);
+out1:
+ kfree(buffer);
+ return retval;
+}

static ssize_t container_file_write(struct file *file, const char __user *buf,
    size_t nbytes, loff_t *ppos)
{
@@ -915,6 +1085,189 @@ int container_add_files(struct container
    return 0;
}

+/* Count the number of tasks in a container. Could be made more
+ * time-efficient but less space-efficient with more linked lists

```

```

+ * running through each container and the css_group structures that
+ * referenced it. Must be called with tasklist_lock held for read or
+ * write or in an rcu critical section.
+ */
+int __container_task_count(const struct container *cont)
+{
+ int count = 0;
+ struct task_struct *g, *p;
+ struct container_subsys_state *css;
+ int subsys_id;
+
+ get_first_subsys(cont, &css, &subsys_id);
+ do_each_thread(g, p) {
+ if (task_subsys_state(p, subsys_id) == css)
+ count++;
+ } while_each_thread(g, p);
+ return count;
+}
+
+/*
+ * Stuff for reading the 'tasks' file.
+ *
+ * Reading this file can return large amounts of data if a container has
+ * lots of attached tasks. So it may need several calls to read(),
+ * but we cannot guarantee that the information we produce is correct
+ * unless we produce it entirely atomically.
+ *
+ * Upon tasks file open(), a struct ctr_struct is allocated, that
+ * will have a pointer to an array (also allocated here). The struct
+ * ctr_struct is stored in file->private_data. Its resources will
+ * be freed by release() when the file is closed. The array is used
+ * to sprintf the PIDs and then used by read().
+ */
+struct ctr_struct {
+ char *buf;
+ int bufsz;
+};
+
+/*
+ * Load into 'pidarray' up to 'npids' of the tasks using container
+ * 'cont'. Return actual number of pids loaded. No need to
+ * task_lock(p) when reading out p->container, since we're in an RCU
+ * read section, so the css_group can't go away, and is
+ * immutable after creation.
+ */
+static int pid_array_load(pid_t *pidarray, int npids, struct container *cont)
+{
+ int n = 0;

```

```

+ struct task_struct *g, *p;
+ struct container_subsys_state *css;
+ int subsys_id;
+
+ get_first_subsys(cont, &css, &subsys_id);
+ rCU_read_lock();
+ do_each_thread(g, p) {
+   if (task_subsys_state(p, subsys_id) == css) {
+     pidarray[n++] = pid_nr(task_pid(p));
+     if (unlikely(n == npids))
+       goto array_full;
+   }
+ } while_each_thread(g, p);
+
+array_full:
+ rCU_read_unlock();
+ return n;
+}
+
+static int cmppid(const void *a, const void *b)
+{
+  return *(pid_t *)a - *(pid_t *)b;
+}
+
+/*
+ * Convert array 'a' of 'npids' pid_t's to a string of newline separated
+ * decimal pids in 'buf'.  Don't write more than 'sz' chars, but return
+ * count 'cnt' of how many chars would be written if buf were large enough.
+ */
+static int pid_array_to_buf(char *buf, int sz, pid_t *a, int npids)
+{
+  int cnt = 0;
+  int i;
+
+  for (i = 0; i < npids; i++)
+    cnt += snprintf(buf + cnt, max(sz - cnt, 0), "%d\n", a[i]);
+  return cnt;
+}
+
+/*
+ * Handle an open on 'tasks' file.  Prepare a buffer listing the
+ * process id's of tasks currently attached to the container being opened.
+ *
+ * Does not require any specific container mutexes, and does not take any.
+ */
+static int container_tasks_open(struct inode *unused, struct file *file)
+{
+  struct container *cont = __d_cont(file->f_dentry->d_parent);

```

```

+ struct ctr_struct *ctr;
+ pid_t *pidarray;
+ int npids;
+ char c;
+
+ if (!(file->f_mode & FMODE_READ))
+   return 0;
+
+ ctr = kmalloc(sizeof(*ctr), GFP_KERNEL);
+ if (!ctr)
+   goto err0;
+
+ /*
+  * If container gets more users after we read count, we won't have
+  * enough space - tough. This race is indistinguishable to the
+  * caller from the case that the additional container users didn't
+  * show up until sometime later on.
+ */
+ npids = container_task_count(cont);
+ if (npids) {
+   pidarray = kmalloc(npids * sizeof(pid_t), GFP_KERNEL);
+   if (!pidarray)
+     goto err1;
+
+   npids = pid_array_load(pidarray, npids, cont);
+   sort(pidarray, npids, sizeof(pid_t), cmppid, NULL);
+
+   /* Call pid_array_to_buf() twice, first just to get bufsz */
+   ctr->bufsz = pid_array_to_buf(&c, sizeof(c), pidarray, npids) + 1;
+   ctr->buf = kmalloc(ctr->bufsz, GFP_KERNEL);
+   if (!ctr->buf)
+     goto err2;
+   ctr->bufsz = pid_array_to_buf(ctr->buf, ctr->bufsz, pidarray, npids);
+
+   kfree(pidarray);
+ } else {
+   ctr->buf = 0;
+   ctr->bufsz = 0;
+ }
+ file->private_data = ctr;
+ return 0;
+
+err2:
+ kfree(pidarray);
+err1:
+ kfree(ctr);
+err0:
+ return -ENOMEM;

```

```

+}
+
+static ssize_t container_tasks_read(struct container *cont,
+    struct cftype *cft,
+    struct file *file, char __user *buf,
+    size_t nbytes, loff_t *ppos)
+{
+    struct ctr_struct *ctr = file->private_data;
+
+    return simple_read_from_buffer(buf, nbytes, ppos, ctr->buf, ctr->bufsz);
+}
+
+static int container_tasks_release(struct inode *unused_inode,
+    struct file *file)
+{
+    struct ctr_struct *ctr;
+
+    if (file->f_mode & FMODE_READ) {
+        ctr = file->private_data;
+        kfree(ctr->buf);
+        kfree(ctr);
+    }
+    return 0;
+}
+
+/*
+ * for the common functions, 'private' gives the type of file
+ */
+static struct cftype cft_tasks = {
+    .name = "tasks",
+    .open = container_tasks_open,
+    .read = container_tasks_read,
+    .write = container_common_file_write,
+    .release = container_tasks_release,
+    .private = FILE_TASKLIST,
+};
+
+static int container_populate_dir(struct container *cont)
{
    int err;
@@ -923,6 +1276,10 @@ static int container_populate_dir(struct
/* First clear out any existing files */
container_clear_directory(cont->dentry);

+    err = container_add_file(cont, NULL, &cft_tasks);
+    if (err < 0)
+        return err;
+

```

```
for_each_subsys(cont->root, ss) {  
    if (ss->populate && (err = ss->populate(ss, cont)) < 0)  
        return err;  
  
--
```

---