
Subject: [RFC PATCH 3/4] namespace containers: add nsproxy to nscont struct
Posted by [serue](#) on Mon, 12 Feb 2007 22:23:00 GMT

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

From: "Serge E. Hallyn" <serue@us.ibm.com>

Subject: [RFC PATCH 3/4] namespace containers: add nsproxy to nscont struct

Each ns container is associated with an nsproxy. Add that nsproxy to the nscont struct, set it when a container is auto-created on clone/unshare, and inc/dec the nsproxy to account for each container referencing it.

Note that once the nscont->nsproxy is set, it will never change for the duration of the container's lifetime.

Signed-off-by: Serge E. Hallyn <serue@us.ibm.com>

kernel/container.c | 11 ++++++++
kernel/ns_container.c | 11 ++++++++
2 files changed, 22 insertions(+), 0 deletions(-)

e4dd14babf4fef1849cb0c5e797f8d176eb9a0a4
diff --git a/kernel/container.c b/kernel/container.c
index fc559ce..4608a12 100644

--- a/kernel/container.c

+++ b/kernel/container.c

```
@@ -1581,6 +1581,14 @@ static void get_unused_name(char *buf) {  
    sprintf(buf, "node%d", atomic_inc_return(&namecnt));  
}
```

```
+/* XXX need to create include/linux/ns_container.h and move this there */
```

```
+#ifdef CONFIG_CONTAINER_NS
```

```
+void ns_set_nsproxy_from_task(struct container *cont, struct task_struct *tsk);
```

```
+#else
```

```
+static inline void
```

```
+ns_set_nsproxy_from_task(struct container *cont, struct task_struct *tsk) { }
```

```
+#endif
```

```
+
```

```
/**
```

```
 * container_clone - duplicate the current container and move this
```

```
 * task into the new child
```

```
@@ -1669,6 +1677,9 @@ int container_clone(struct task_struct *  
    goto again;  
}
```

```
+/* mark the auto-created container with the new namespace list */
```

```

+ ns_set_nsproxy_from_task(child, tsk);
+
+ /* All seems fine. Finish by moving the task into the new container */
+ ret = attach_task(child, tsk);
+ mutex_unlock(&manage_mutex);
diff --git a/kernel/ns_container.c b/kernel/ns_container.c
index d60d4f5..2d5c578 100644
--- a/kernel/ns_container.c
+++ b/kernel/ns_container.c
@@ -11,6 +11,7 @@

struct nscont {
    struct container_subsys_state css;
+ struct nsproxy *nsproxy; /* never changes once set */
    spinlock_t lock;
};

@@ -81,10 +82,20 @@ static int ns_create(struct container_sub
    return 0;
}

+/* called from container_clone */
+void ns_set_nsproxy_from_task(struct container *cont, struct task_struct *tsk)
+{
+ struct nscont *ns = container_nscont(cont);
+ ns->nsproxy = tsk->nsproxy;
+ get_nsproxy(ns->nsproxy);
+}
+
+static void ns_destroy(struct container_subsys *ss,
+    struct container *cont)
+{
+ struct nscont *ns = container_nscont(cont);
+ if (ns->nsproxy)
+ put_nsproxy(ns->nsproxy);
+ kfree(ns);
+}

--
1.1.6

```
