Subject: [PATCH] containers: define a namespace container subsystem
Posted by serue on Wed, 31 Jan 2007 06:12:04 GMT
View Forum Message <> Reply to Message

Hi,

Here's my next, much more satisfying attempt at doing namespace
tracking through a container subsystem. The commit log pretty
much details where I would expect to go with this.

It behaves pretty differently from other subsystems implemented
so far, which could either be seen as evidence that it doesn't
belong as a subsystem, or, more likely, that the container
subsystem approach is quite flexible. I particularly like the
implications of being able to mount both the ns_container subsystem
and a resource usage subsystem under the same hierarchy to get
automatic resource control on vservers or checkpointable jobs.

-serge

From: "Serge E. Hallyn" <serue@us.ibm.com>
Subject: [PATCH] containers: define a namespace container subsystem

Define a container subsystem to track namespace unshares.
So long as the ns_container subsystem is mounted in a
container hierarchy, every unshare (or clone with a new
mounts, uts, etc namespace) will create a new child container
and move the process into it.

The purpose of this is to eventually allow operations on
virtual servers such as kill or enter, and checkpoint and
kill checkpointable namespaces.

Manual entering of an ns_container (that is, using
echo > /../container_dir/tasks) is only allowed into a
container which is a child of the current. Currently
there is also the restriction that the new container must
not yet have any tasks. This will change soon, but will
of course require that the task being transitioned switch
to the nsproxy associated with the new container.

Manual moving of an *other* process into a new container
also requires that the process being moved be in the same
container as or a subcontainer of that of the process
requesting the move.

Creation of an ns_container is restricted to subcontainers
of the creating task's container.

Creation of an ns_container, and moving another task into
a ns_container, both require CAP_SYS_ADMIN. The mkdir
control may seem overly harsh, but since currently CAP_SYS_ADMIN
is also required to unshare, it results in no additional
limitations.

The next steps are (not necessarily in order):

1. allow rm -rf to kill all processes under a
   ns_container - with the intent of killing all
   processes in a virtual server

2. implement transitioning into a populated container,
   with the effect of setting the task's nsproxy to
   the one represented by the container.

3. define a file for each type of namespace in each
   ns_container, with the i_op->symlink() defined to
   allow creation of a new ns_container which references
   only some of the namespace pointers of an existing
   (child) container. All other namespaces will be
   taken from the existing process. In this way it
   is possible to enter just a network namespace of
   some vserver.

4. probably make containers mac-aware, that is add a
   ->security pointer, and LSM hooks at appropriate
   points so that, for instance, SELinux can control
   vserver kill and enters.

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

---

 include/linux/container.h | 5 ++
 init/Kconfig | 9 +++
 kernel/Makefile | 1
 kernel/container.c | 114 ++
 kernel/fork.c | 5 ++
 kernel/ns_container.c | 117 ++
 kernel/nsproxy.c | 8 +++
 7 files changed, 259 insertions(+), 0 deletions(-)
 create mode 100644 kernel/ns_container.c

f4f9d31fa238b2ff804d35c6e1b5a0776f2bf303
diff --git a/include/linux/container.h b/include/linux/container.h
index 45db753..15b0446 100644

```
--- a/include/linux/container.h
+++ b/include/linux/container.h
@@ -207,6 +207,10 @@ static inline struct container_subsys_st
 }

 int container_path(const struct container *cont, char *buf, int buflen);
+int task_container_is_ancestor(struct task_struct *tsk, int subsys,
+       struct container *cont);
+
+int container_switch(struct task_struct *tsk);

 #else /* !CONFIG_CONTAINERS */

@@ -218,6 +222,7 @@ static inline void container_exit(struct

 static inline void container_lock(void) {}
 static inline void container_unlock(void) {}
+static inline void container_switch(struct task_struct *tsk) { }

 #endif /* !CONFIG_CONTAINERS */

diff --git a/init/Kconfig b/init/Kconfig
index ebaec57..c00b19c 100644
--- a/init/Kconfig
+++ b/init/Kconfig
@@ -297,6 +297,15 @@ config CONTAINER_CPUACCT
    Provides a simple Resource Controller for monitoring the
    total CPU consumed by the tasks in a container

+config CONTAINER_NS
+ bool "Namespace container subsystem"
+ select CONTAINERS
+ help
+   Provides a simple namespace container subsystem to
+   provide hierarchical naming of sets of namespaces,
+   for instance virtual servers and checkpoint/restart
+   jobs.
+
 config RELAY
  bool "Kernel->user space relay support (formerly relayfs)"
  help
diff --git a/kernel/Makefile b/kernel/Makefile
index feba860..6c73a5e 100644
--- a/kernel/Makefile
+++ b/kernel/Makefile
@@ -39,6 +39,7 @@ obj-$(CONFIG_COMPAT) += compat.o
 obj-$(CONFIG_CONTAINERS) += container.o
 obj-$(CONFIG_CPUSETS) += cpuset.o
```

```
 obj-$(CONFIG_CONTAINER_CPUACCT) += cpu_acct.o
+obj-$(CONFIG_CONTAINER_NS) += ns_container.o
 obj-$(CONFIG_IKCONFIG) += configs.o
 obj-$(CONFIG_STOP_MACHINE) += stop_machine.o
 obj-$(CONFIG_AUDIT) += audit.o auditfilter.o
diff --git a/kernel/container.c b/kernel/container.c
index 1f44351..97de43b 100644
--- a/kernel/container.c
+++ b/kernel/container.c
@@ -1558,6 +1558,120 @@ static int container_rmdir(struct inode
  return 0;
 }

+static int namecnt = 0;
+static DEFINE_SPINLOCK(namecnt_lock);
+
+static void get_unused_name(char *buf)
+{
+
+ spin_lock(&namecnt_lock);
+ memset(buf, 0, 20);
+ snprintf(buf, 20, "node%d", namecnt++);
+ spin_unlock(&namecnt_lock);
+}
+
+#ifndef CONFIG_CONTAINER_NS
+static int ns_container_subsys_idx = -1;
+#else
+extern int ns_container_subsys_idx;
+#endif
+
+static inline struct container *find_child_container(struct container *parent,
+      char *name)
+{
+ struct dentry *d = container_get_dentry(parent->dentry, name);
+ return d ? d->d_fsdata : NULL;
+}
+
+#define NS_CONT_MODE (S_IFDIR | S_IRUGO | S_IXUGO | S_IWUSR)
+int container_switch(struct task_struct *tsk)
+{
+ int h;
+ struct container *cur_cont, *new_cont;
+ char path[20];
+ struct qstr name;
+ struct dentry *dentry;
+ int ret;
+ char *pathbuf = NULL;
```

```
+ char buffer[20];
+
+ /* check if nsproxy subsys is registered */
+ if (ns_container_subsys_idx == -1)
+  return 0;
+
+ printk(KERN_NOTICE "%s: ns_container subsys registered\n", __FUNCTION__);
+ /* check if nsproxy subsys is mounted in some hierarchy */
+ rcu_read_lock();
+ h = rcu_dereference(subsys[ns_container_subsys_idx]->hierarchy);
+ rcu_read_unlock();
+ if (h == 0) {
+  /* do we mount the nsproxy subsys, or just skip
+   * creating a container?  I think we just skip
+   * it.
+   */
+  printk(KERN_NOTICE "%s: hierarchy (for idx %d) was 0\n", __FUNCTION__,
+   ns_container_subsys_idx);
+   return 0;
+ }
+ printk(KERN_NOTICE "%s: ns_container subsys not mounted.\n", __FUNCTION__);
+ cur_cont = tsk->container[h];
+ get_unused_name(path);
+
+ name.name = path;
+ name.len = strlen(name.name);
+ dentry = d_alloc(cur_cont->dentry, &name);
+ if (IS_ERR(dentry)) {
+  printk(KERN_NOTICE "%s: couldn't get dentry for current container\n", __FUNCTION__);
+  return PTR_ERR(dentry);
+ }
+ printk(KERN_NOTICE "%s: got dentry for current container\n", __FUNCTION__);
+ ret = vfs_mkdir(cur_cont->dentry->d_inode, dentry, NS_CONT_MODE);
+ dput(dentry);
+ if (ret) {
+  printk(KERN_NOTICE "%s: Couldn't make new directory (%d)\n", __FUNCTION__, ret);
+  return ret;
+ }
+ printk(KERN_NOTICE "%s: made new dir\n", __FUNCTION__);
+
+ new_cont = find_child_container(cur_cont, path);
+ if (!new_cont) {
+  printk(KERN_NOTICE "%s: Couldn't find new container\n", __FUNCTION__);
+  return -ENOMEM;
+ }
+ printk(KERN_NOTICE "%s: found new container\n", __FUNCTION__);
+
+ snprintf(buffer, 20, "%lu\n", (unsigned long)tsk->pid);
```

```
+ mutex_lock(&manage_mutex);
+ ret = attach_task(new_cont, buffer, &pathbuf);
+ mutex_unlock(&manage_mutex);
+
+ /* I dunno - do I *need* to call container_release_agent? */
+ // container_release_agent(cont->hierarchy, pathbuf);
+
+ return ret;
+}
+
+/*
+ * XXX - should probably be called with a lock
+ * isn't right now
+ */
+int task_container_is_ancestor(struct task_struct *tsk, int subsys,
+      struct container *cont)
+{
+ struct container *q = tsk->container[subsys];
+
+ if (!q)
+  return 1;
+
+ do {
+  cont = cont->parent;
+ } while (cont != cont->top_container && cont != q);
+
+ if (cont == q)
+  return 1;
+ return 0;
+}
+

 /**
  * container_init_early - initialize containers at system boot
diff --git a/kernel/fork.c b/kernel/fork.c
index 984fe31..cf06ff5 100644
--- a/kernel/fork.c
+++ b/kernel/fork.c
@@ -1661,6 +1661,10 @@ asmlinkage long sys_unshare(unsigned lon
    err = -ENOMEM;
    goto bad_unshare_cleanup_ipc;
   }
+
+  err = container_switch(current);
+  if (err)
+   goto bad_unshare_cleanup_dupns;
  }
```

```
  if (new_fs || new_ns || new_mm || new_fd || new_ulist ||
@@ -1715,6 +1719,7 @@ asmlinkage long sys_unshare(unsigned lon
   task_unlock(current);
  }

+bad_unshare_cleanup_dupns:
  if (new_nsproxy)
   put_nsproxy(new_nsproxy);

diff --git a/kernel/ns_container.c b/kernel/ns_container.c
new file mode 100644
index 0000000..fd2e92b
--- /dev/null
+++ b/kernel/ns_container.c
@@ -0,0 +1,117 @@
+/*
+ * ns_container.c - namespace container subsystem
+ *
+ * Copyright IBM, 2006
+ */
+
+#include <linux/module.h>
+#include <linux/container.h>
+#include <linux/fs.h>
+
+int ns_container_subsys_idx = -1;
+
+struct nscont {
+ struct container_subsys_state css;
+ spinlock_t lock;
+};
+
+static struct container_subsys ns_subsys;
+
+static inline struct nscont *container_nscont(struct container *cont)
+{
+ return container_of(container_subsys_state(cont, &ns_subsys),
+       struct nscont, css);
+}
+
+/*
+ * Rules:
+ *   1. you can only enter a container which is a child of your current
+ * container
+ *   2. you can only place another process into a container if
+ * a. you have CAP_SYS_ADMIN
+ * b. your container is an ancestor of tsk's destination container
+ *   (hence either you are in the same container as tsk, or in an
```

```
+ *    ancestor container thereof)
+ */
+int ns_can_attach(struct container_subsys *ss,
+        struct container *cont, struct task_struct *tsk)
+{
+ struct container *c;
+
+ if (current != tsk) {
+  int hierarchy;
+  if (!capable(CAP_SYS_ADMIN))
+   return -EPERM;
+
+  rcu_read_lock();
+  hierarchy = rcu_dereference(ns_subsys.hierarchy);
+  rcu_read_unlock();
+
+  if (!task_container_is_ancestor(current, hierarchy, cont))
+   return -EPERM;
+ }
+
+ if (atomic_read(&cont->count) != 0)
+  return -EPERM;
+
+ c = task_container(tsk, &ns_subsys);
+ if (c && c != cont->parent)
+  return -EPERM;
+
+ return 0;
+}
+
+/*
+ * Rules: you can only create a container if
+ * 1. you are capable(CAP_SYS_ADMIN)
+ * 2. the target container is a decendent of your own container
+ */
+static int ns_create(struct container_subsys *ss, struct container *cont)
+{
+ struct nscont *ns;
+ int hierarchy;
+
+ if (!capable(CAP_SYS_ADMIN))
+  return -EPERM;
+ rcu_read_lock();
+ hierarchy = rcu_dereference(ns_subsys.hierarchy);
+ rcu_read_unlock();
+ if (hierarchy && !task_container_is_ancestor(current, hierarchy, cont))
+  return -EPERM;
+
```

```
+ ns = kzalloc(sizeof(*ns), GFP_KERNEL);
+ if (!ns) return -ENOMEM;
+ spin_lock_init(&ns->lock);
+ cont->subsys[ns_subsys.subsys_id] = &ns->css;
+ return 0;
+}
+
+static void ns_destroy(struct container_subsys *ss,
+       struct container *cont)
+{
+ struct nscont *ns = container_nscont(cont);
+ kfree(ns);
+}
+
+static struct container_subsys ns_subsys = {
+ .name = "ns_container",
+ .create = ns_create,
+ .destroy  = ns_destroy,
+ .can_attach = ns_can_attach,
+ //.attach = ns_attach,
+ //.post_attach = ns_post_attach,
+ //.populate = ns_populate,
+ .subsys_id = -1,
+};
+
+int __init ns_init(void)
+{
+ int ret;
+
+ ret = container_register_subsys(&ns_subsys);
+ ns_container_subsys_idx = ns_subsys.subsys_id;
+
+ return ret < 0 ? ret : 0;
+}
+
+module_init(ns_init)
diff --git a/kernel/nsproxy.c b/kernel/nsproxy.c
index f5b9ee6..c4f41de 100644
--- a/kernel/nsproxy.c
+++ b/kernel/nsproxy.c
@@ -20,6 +20,7 @@
 #include <linux/mnt_namespace.h>
 #include <linux/utsname.h>
 #include <linux/pid_namespace.h>
+#include <linux/container.h>

 struct nsproxy init_nsproxy = INIT_NSPROXY(init_nsproxy);
```

```
@@ -115,11 +116,18 @@ int copy_namespaces(int flags, struct ta
  err = copy_pid_ns(flags, tsk);
  if (err)
   goto out_pid;
+
+ err = container_switch(tsk);
+ if (err)
+  goto out_container;

 out:
  put_nsproxy(old_ns);
  return err;

+out_container:
+ if (new_ns->pid_ns)
+  put_pid_ns(new_ns->pid_ns);
 out_pid:
  if (new_ns->ipc_ns)
   put_ipc_ns(new_ns->ipc_ns);
--
1.1.6
```

_____
Containers mailing list
Containers@lists.osdl.org
https://lists.osdl.org/mailman/listinfo/containers

Subject: Re: [PATCH] containers: define a namespace container subsystem
Posted by serue on Wed, 31 Jan 2007 17:19:10 GMT
View Forum Message <> Reply to Message

Quoting Serge E. Hallyn (serue@us.ibm.com):
> Hi,
>
> Here's my next, much more satisfying attempt at doing namespace
> tracking through a container subsystem. The commit log pretty
> much details where I would expect to go with this.

How to use this:
 . mount -t container -ons_container ns_container /containers
 . (unshare a namespace)
 . cat /proc/self/containers
   You'll see something like
     ns_container:/node0
   Unshare again, and you'll be in
     ns_container:/node0/node1

thanks,

-serge

_____

Containers mailing list
Containers@lists.osdl.org
https://lists.osdl.org/mailman/listinfo/containers

Subject: Re: [PATCH] containers: define a namespace container subsystem
Posted by Cedric Le Goater on Fri, 02 Feb 2007 16:19:52 GMT
View Forum Message <> Reply to Message

> The next steps are (not necessarily in order):
>
> 1. allow rm -rf to kill all processes under a
> ns_container - with the intent of killing all
> processes in a virtual server
>
> 2. implement transitioning into a populated container,
> with the effect of setting the task's nsproxy to
> the one represented by the container.
>
> 3. define a file for each type of namespace in each

could that file be a directory exposing some critical data
from each namespace ?

I would imagine the network devices for the net namespace
and be able to interact with them (Daniel ?). the task list
for the pid namespace, etc.

> ns_container, with the i_op->symlink() defined to
> allow creation of a new ns_container which references
> only some of the namespace pointers of an existing
> (child) container. All other namespaces will be
> taken from the existing process. In this way it
> is possible to enter just a network namespace of
> some vserver.
> 4. probably make containers mac-aware, that is add a
> ->security pointer, and LSM hooks at appropriate
> points so that, for instance, SELinux can control
> vserver kill and enters.
>

_____

Containers mailing list
Containers@lists.osdl.org
https://lists.osdl.org/mailman/listinfo/containers

Quoting Cedric Le Goater (clg@fr.ibm.com):
>
> > The next steps are (not necessarily in order):
> >
> > 1. allow rm -rf to kill all processes under a
> > ns_container - with the intent of killing all
> > processes in a virtual server
> >
> > 2. implement transitioning into a populated container,
> > with the effect of setting the task's nsproxy to
> > the one represented by the container.
> >
> > 3. define a file for each type of namespace in each
>
> could that file be a directory exposing some critical data
> from each namespace ?

it probably could be, but that might be confusing since subcontainers
are also directories. Would just putting the data into the namespace
files suffice? This isn't sysfs so no 1-value-per-file restrictions...

> I would imagine the network devices for the net namespace
> and be able to interact with them (Daniel ?). the task list
> for the pid namespace, etc.

Well the tasklist will already be in the 'tasks' file created by the
containers code :)

But actually, making them directories might actually be easier, because
iirc cftypes only have f_ops right now, whereas dirs already have i_ops,
so doing the symlink magic should be easier that way.

Great idea! :)

thanks,
-serge

> > ns_container, with the i_op->symlink() defined to
> > allow creation of a new ns_container which references
> > only some of the namespace pointers of an existing
> > (child) container. All other namespaces will be
> > taken from the existing process. In this way it
> > is possible to enter just a network namespace of
> > some vserver.
> > 4. probably make containers mac-aware, that is add a

> > ->security pointer, and LSM hooks at appropriate
> > points so that, for instance, SELinux can control
> > vserver kill and enters.
> >

_____

Containers mailing list
Containers@lists.osdl.org
https://lists.osdl.org/mailman/listinfo/containers

Subject: Re: [PATCH] containers: define a namespace container subsystem
Posted by Paul Menage on Fri, 02 Feb 2007 22:25:25 GMT
View Forum Message <> Reply to Message

On 1/30/07, Serge E. Hallyn <serue@us.ibm.com> wrote:
>
> It behaves pretty differently from other subsystems implemented
> so far, which could either be seen as evidence that it doesn't
> belong as a subsystem, or, more likely, that the container
> subsystem approach is quite flexible.

The latter - the container system is meant to be able to supportmore
than just resource controllers.

```
> +                                       char *name)
> +{
> +       struct dentry *d = container_get_dentry(parent->dentry, name);
> +       return d ? d->d_fsdata : NULL;
> +}
> +
> +#define NS_CONT_MODE (S_IFDIR | S_IRUGO | S_IXUGO | S_IWUSR)
> +int container_switch(struct task_struct *tsk)
> +{
> +       int h;
> +       struct container *cur_cont, *new_cont;
> +       char path[20];
> +       struct qstr name;
> +       struct dentry *dentry;
> +       int ret;
> +       char *pathbuf = NULL;
> +       char buffer[20];
> +
> +       /* check if nsproxy subsys is registered */
> +       if (ns_container_subsys_idx == -1)
> +               return 0;
> +
> +       printk(KERN_NOTICE "%s: ns_container subsys registered\n", __FUNCTION__);
> +       /* check if nsproxy subsys is mounted in some hierarchy */
```

> +	rcu_read_lock();
> +	h = rcu_dereference(subsys[ns_container_subsys_idx]->hierarchy);
> +	rcu_read_unlock();
> +	if (h == 0) {
> +		/* do we mount the nsproxy subsys, or just skip
> +		 * creating a container? I think we just skip
> +		 * it.

I'd say that we should try to create a fresh hierarchy with just the nsproxy subsystem on it. Otherwise if someone tries to mount the nsproxy subsystem later, we end up with some namespaces with no container.

This looks great - I'll incorporate it or something like it in my next patch set.

Cheers,

Paul

_____
Containers mailing list
Containers@lists.osdl.org
https://lists.osdl.org/mailman/listinfo/containers

Subject: Re: [PATCH] containers: define a namespace container subsystem
Posted by serue on Fri, 02 Feb 2007 23:35:09 GMT
View Forum Message <> Reply to Message

Quoting Paul Menage (menage@google.com):
> On 1/30/07, Serge E. Hallyn <serue@us.ibm.com> wrote:
> >
> >It behaves pretty differently from other subsystems implemented
> >so far, which could either be seen as evidence that it doesn't
> >belong as a subsystem, or, more likely, that the container
> >subsystem approach is quite flexible.
>
> The latter - the container system is meant to be able to supportmore
> than just resource controllers.
>
> >+ char *name)
> >+{
> >+	struct dentry *d = container_get_dentry(parent->dentry, name);
> >+	return d ? d->d_fsdata : NULL;
> >+}
> >+
> >+#define NS_CONT_MODE (S_IFDIR | S_IRUGO | S_IXUGO | S_IWUSR)
> >+int container_switch(struct task_struct *tsk)

```
> >+{
> >+	int h;
> >+	struct container *cur_cont, *new_cont;
> >+	char path[20];
> >+	struct qstr name;
> >+	struct dentry *dentry;
> >+	int ret;
> >+	char *pathbuf = NULL;
> >+	char buffer[20];
> >+
> >+	/* check if nsproxy subsys is registered */
> >+	if (ns_container_subsys_idx == -1)
> >+		return 0;
> >+
> >+	printk(KERN_NOTICE "%s: ns_container subsys registered\n",
> >__FUNCTION__);
> >+	/* check if nsproxy subsys is mounted in some hierarchy */
> >+	rcu_read_lock();
> >+	h = rcu_dereference(subsys[ns_container_subsys_idx]->hierarchy);
> >+	rcu_read_unlock();
> >+	if (h == 0) {
> >+		/* do we mount the nsproxy subsys, or just skip
> >+		 * creating a container?  I think we just skip
> >+		 * it.
>
> I'd say that we should try to create a fresh hierarchy with just the
> nsproxy subsystem on it. Otherwise if someone tries to mount the
> nsproxy subsystem later, we end up with some namespaces with no
> container.
```

Yes, but if we automatically create a fresh hierarchy with just the nsproxy subsystem on it, then if you do any unsharing during boot or login, which with pam_namespace on LSPP systems is very possible, then you'll never be able to manually mount a hierarchy with an nsproxy subsystem and a resource controller in the same hierarchy.

Whereas having the topmost container contain a bunch of nsproxies really isn't a problem I don't think.

> This looks great - I'll incorporate it or something like it in my next
> patch set.

Great. Note though that I just found a little buglet - the following patch is needed for the !CONFIG_CONTAINERS case :)

thanks,
-serge

Subject: [PATCH] fix !config_containers compile

when container_switch was changed from returning void to int,
the !config_containers inline version was not updated.

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

---

 include/linux/container.h | 2 +-
 1 files changed, 1 insertions(+), 1 deletions(-)

6adc0fbab23e7f971750c3eb6c244d1008a636c0
diff --git a/include/linux/container.h b/include/linux/container.h
index 15b0446..a9075ac 100644
--- a/include/linux/container.h
+++ b/include/linux/container.h
@@ -222,7 +222,7 @@ static inline void container_exit(struct

 static inline void container_lock(void) {}
 static inline void container_unlock(void) {}
-static inline void container_switch(struct task_struct *tsk) { }
+static inline int container_switch(struct task_struct *tsk) { }

 #endif /* !CONFIG_CONTAINERS */

--
1.1.6

_____
Containers mailing list
Containers@lists.osdl.org
https://lists.osdl.org/mailman/listinfo/containers

Subject: Re: [PATCH] containers: define a namespace container subsystem
Posted by serue on Sat, 03 Feb 2007 05:37:27 GMT
View Forum Message <> Reply to Message

Quoting Serge E. Hallyn (serue@us.ibm.com):
> Great. Note though that I just found a little buglet - the following
> patch is needed for the !CONFIG_CONTAINERS case :)
>
> thanks,
> -serge
>
> Subject: [PATCH] fix !config_containers compile
>
> when container_switch was changed from returning void to int,

> the !config_containers inline version was not updated.
>
> Signed-off-by: Serge E. Hallyn <serue@us.ibm.com>
>
> ---
>
> include/linux/container.h | 2 +-
> 1 files changed, 1 insertions(+), 1 deletions(-)
>
> 6adc0fbab23e7f971750c3eb6c244d1008a636c0
> diff --git a/include/linux/container.h b/include/linux/container.h
> index 15b0446..a9075ac 100644
> --- a/include/linux/container.h
> +++ b/include/linux/container.h
> @@ -222,7 +222,7 @@ static inline void container_exit(struct
>
> static inline void container_lock(void) {}
> static inline void container_unlock(void) {}
> -static inline void container_switch(struct task_struct *tsk) { }
> +static inline int container_switch(struct task_struct *tsk) { }
>
> #endif /* !CONFIG_CONTAINERS */

Egads, I blame that one on Friday. Let's try the following instead...

-serge

Subject: [PATCH] fic !config_containers compile

when container_switch was changed from returning void to int,
the !config_containers inline version was not updated.

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

---

 include/linux/container.h | 2 +-
 1 files changed, 1 insertions(+), 1 deletions(-)

c588c6f4604c80f26b799ebd317726cba8696688
diff --git a/include/linux/container.h b/include/linux/container.h
index 15b0446..294e996 100644
--- a/include/linux/container.h
+++ b/include/linux/container.h
@@ -222,7 +222,7 @@ static inline void container_exit(struct

 static inline void container_lock(void) {}
 static inline void container_unlock(void) {}

-static inline void container_switch(struct task_struct *tsk) { }
+static inline int container_switch(struct task_struct *tsk) { return 0; }

 #endif /* !CONFIG_CONTAINERS */

--
1.1.6

_____
Containers mailing list
Containers@lists.osdl.org
https://lists.osdl.org/mailman/listinfo/containers

Subject: Re: [PATCH] containers: define a namespace container subsystem
Posted by Cedric Le Goater on Mon, 05 Feb 2007 12:21:52 GMT
View Forum Message <> Reply to Message

Serge E. Hallyn wrote:
> Quoting Cedric Le Goater (clg@fr.ibm.com):
>>> The next steps are (not necessarily in order):
>>>
>>> 1. allow rm -rf to kill all processes under a
>>>   ns_container - with the intent of killing all
>>>   processes in a virtual server
>>>
>>> 2. implement transitioning into a populated container,
>>>   with the effect of setting the task's nsproxy to
>>>   the one represented by the container.
>>>
>>> 3. define a file for each type of namespace in each
>> could that file be a directory exposing some critical data
>> from each namespace ?
>
> it probably could be, but that might be confusing since subcontainers
> are also directories. Would just putting the data into the namespace
> files suffice? This isn't sysfs so no 1-value-per-file restrictions...

ok.

Would it be reasonable to use such a file to expose or hide network
interfaces in an l3 network namespace ?

what would be nice now is to rebase Paul's patchset on next -mm and
see how we interact with it and the namespaces ? I already did such a
merge a while ago but there was no connections between the features.
We need to come to that point.

I'll try again when andrew releases and include your patch, serge.

regards,

C.

_____

Containers mailing list
Containers@lists.osdl.org
https://lists.osdl.org/mailman/listinfo/containers

Subject: Re: [PATCH] containers: define a namespace container subsystem
Posted by Cedric Le Goater on Wed, 07 Feb 2007 16:29:25 GMT
View Forum Message <> Reply to Message

> what would be nice now is to rebase Paul's patchset on next -mm and
> see how we interact with it and the namespaces ? I already did such a
> merge a while ago but there was no connections between the features.
> We need to come to that point.

Paul, are you planning on rebasing your patchset on 2.6.20 ?

thanks,

C.

_____

Containers mailing list
Containers@lists.osdl.org
https://lists.osdl.org/mailman/listinfo/containers

Subject: Re: [PATCH] containers: define a namespace container subsystem
Posted by Paul Menage on Wed, 07 Feb 2007 17:06:25 GMT
View Forum Message <> Reply to Message

On 2/7/07, Cedric Le Goater <clg@fr.ibm.com> wrote:
>
> > what would be nice now is to rebase Paul's patchset on next -mm and
> > see how we interact with it and the namespaces ? I already did such a
> > merge a while ago but there was no connections between the features.
> > We need to come to that point.
>
> Paul, are you planning on rebasing your patchset on 2.6.20 ?
>

Yes, I did this last night - there were very few changes required to
go from -rc1 to 2.6.20. I'll be sending them out to the list today.
The changes include having just a single aggregated pointer per task,

and hopefully with a version of Serge's container_switch() function
incorporated.

Paul

_____
Containers mailing list
Containers@lists.osdl.org
https://lists.osdl.org/mailman/listinfo/containers

Subject: Re: [PATCH] containers: define a namespace container subsystem
Posted by Cedric Le Goater on Wed, 07 Feb 2007 17:41:44 GMT
View Forum Message <> Reply to Message

Paul Menage wrote:
> On 2/7/07, Cedric Le Goater <clg@fr.ibm.com> wrote:
>>
>> > what would be nice now is to rebase Paul's patchset on next -mm and
>> > see how we interact with it and the namespaces ? I already did such a
>> > merge a while ago but there was no connections between the features.
>> > We need to come to that point.
>>
>> Paul, are you planning on rebasing your patchset on 2.6.20 ?
>>
>
> Yes, I did this last night - there were very few changes required to
> go from -rc1 to 2.6.20. I'll be sending them out to the list today.
> The changes include having just a single aggregated pointer per task,
> and hopefully with a version of Serge's container_switch() function
> incorporated.

thanks, i'll include it the next -lxc patchset i'm maintaining on -mm.

C.

_____
Containers mailing list
Containers@lists.osdl.org
https://lists.osdl.org/mailman/listinfo/containers
