
Subject: [Patch 01/05]- Containers: Documentation on using containers

Posted by [Rohit Seth](#) on Fri, 15 Sep 2006 01:38:33 GMT

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

This patch contains the Documentation for using containers.

Signed-off-by: Rohit Seth <rohitseth@google.com>

Documentation/containers.txt | 42 ++++++
1 files changed, 42 insertions(+)

--- linux-2.6.18-rc6-mm2.org/Documentation/containers.txt 1969-12-31 16:00:00.000000000 -0800

+++ linux-2.6.18-rc6-mm2.ctn/Documentation/containers.txt 2006-09-14 17:13:48.000000000
-0700

@@ -0,0 +1,42 @@

+This file contains information about how to use containers. Configfs support
+is needed in kernel as the container's user interface is through configfs. So
+first enable CONFIG_CONFIGFS_FS and CONFIG_CONTAINERS and recompile the kernel.

+

+1- Mount a configfs (for example):

+ mount -t configfs none /mnt/configfs

+ This will create a /mnt/configfs mount point.

+

+2- As the support of containers is built into kernel, so the mount point

+ /mnt/configfs will automatically contain a directory "containers"

+

+3- Create a container by name test_container

+ cd /mnt/configfs/containers

+ mkdir test_container

+

+All the current implemented attributes in the kernel will show up in the

+directory /configfs/containers/test_container

+

+4- Add a task to container

+ cd /mnt/configfs/containers/test_container

+ echo <pid> > addtask

+

+Now the <pid> and its subsequently forked children will belong to container

+test_container.

+

+5- Remove a task from container

+ echo <pid> rmtask

+

+6- Set a page limit for the container

+ echo some_number_of_pages > page_limit

+

+7- Read the id for the container

+ cat id

+
+8- Get the statistics for this container
+ cat num* (will print active pages, anon_pages, file_pages, num_files,
+ and num_task)
+ cat *hits (will print page_limit_hits and task_limit_hits: the number
+ of times container has gone over page_limit and task_limit)
+9- Freeing a container
+ cd /mnt/configs/containers/
+ rmdir test_container

Subject: Re: [Patch 01/05]- Containers: Documentation on using containers
Posted by [Rolf Eike Beer](#) on Fri, 15 Sep 2006 06:15:19 GMT
[View Forum Message](#) <> [Reply to Message](#)

Rohit Seth wrote:

> This patch contains the Documentation for using containers.

> +5- Remove a task from container

> + echo <pid> rmtask

echo <pid> > rmtask?

Please also give a short description what containers are for. From what I read here I can only guess it's about getting some statistics about a group of tasks.

Eike

Subject: Re: Re: [Patch 01/05]- Containers: Documentation on using containers
Posted by [kir](#) on Fri, 15 Sep 2006 10:12:40 GMT
[View Forum Message](#) <> [Reply to Message](#)

Rolf Eike Beer wrote:

> Please also give a short description what containers are for. From

> what I read

> here I can only guess it's about getting some statistics about a group of

> tasks.

>

Container is like FreeBSD Jail on steroids (and Jail is chroot() on steroids).

As a nice intro, I suggest you to read

http://en.wikipedia.org/wiki/Operating-system_level_virtualization

<http://en.wikipedia.org/wiki/OpenVZ>

Subject: Re: [Patch 01/05]- Containers: Documentation on using containers

Posted by [Rohit Seth](#) on Fri, 15 Sep 2006 16:45:25 GMT

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

On Fri, 2006-09-15 at 08:15 +0200, Rolf Eike Beer wrote:

> Rohit Seth wrote:

> > This patch contains the Documentation for using containers.

>

> > +5- Remove a task from container

> > + echo <pid> rmtask

>

> echo <pid> > rmtask?

>

rmtask is an attribute defined in test_container directory. So, first you have to cd into container directory (cd /mnt/configs/containers/test_container and then execute this command)

> Please also give a short description what containers are for. From what I read > here I can only guess it's about getting some statistics about a group of > tasks.

Containers allow different workloads to be run on the same platform with limits defined on per container basis. This basically allows a single platform to be (soft) partitioned among different workloads (each of which could be running many tasks). The limits could be amount of memory, number of tasks among other features. These two features are already implemented in the patch set that I posted. But it is possible to add other controllers like CPU that allows only finite amount of time to the processes belonging to a container.

Currently this patch set is only tracking user memory (both file based and anonymous). The memory handler is currently deactivating pages belonging to a container that has gone over the limit. Even though this allows containers to go over board their limits but 1- once they are over the limit then they run in degraded manner and 2- if there is any memory pressure then the (extra) pages belonging to this container are the prime candidates for swapping (for example). The statistics that are shown in each container directory are the current values of each resource consumption.

Please let me know if you need any more specific information about the patch set.

-rohit

Subject: Re: Re: [Patch 01/05]- Containers: Documentation on using containers
Posted by [Rohit Seth](#) on Fri, 15 Sep 2006 16:46:07 GMT
[View Forum Message](#) <> [Reply to Message](#)

On Fri, 2006-09-15 at 14:12 +0400, Kir Kolyshkin wrote:

> Rolf Eike Beer wrote:
> > Please also give a short description what containers are for. From
> > what I read
> > here I can only guess it's about gettings some statistics about a group of
> > tasks.
> >
> Container is like FreeBSD Jail on steroids (and Jail is chroot() on
> steroids).

Nice choice of words!

-rohit

Subject: Re: [Patch 01/05]- Containers: Documentation on using containers
Posted by [B.Steinbrink](#) on Fri, 15 Sep 2006 17:44:12 GMT
[View Forum Message](#) <> [Reply to Message](#)

On 2006.09.15 09:45:25 -0700, Rohit Seth wrote:

> On Fri, 2006-09-15 at 08:15 +0200, Rolf Eike Beer wrote:
> > Rohit Seth wrote:
> > > This patch contains the Documentation for using containers.
> >
> > > +5- Remove a task from container
> > > + echo <pid> rmtask
> >
> > echo <pid> > rmtask?
> >
>
> rmtask is an attribute defined in test_container directory. So, first
> you have to cd into container directory
> (cd /mnt/configs/containers/test_container and then execute this
> command)

The > is missing in the patch though, guess that's what Rolf wanted to point out ;)

>
> > Please also give a short description what containers are for. From what I read
> > here I can only guess it's about gettings some statistics about a group of
> > tasks.
>

> Containers allow different workloads to be run on the same platform with
> limits defined on per container basis. This basically allows a single
> platform to be (soft) partitioned among different workloads (each of
> which could be running many tasks). The limits could be amount of
> memory, number of tasks among other features. These two features are
> already implemented in the patch set that I posted. But it is possible
> to add other controllers like CPU that allows only finite amount of time
> to the processes belonging to a container.
>
> Currently this patch set is only tracking user memory (both file based
> and anonymous). The memory handler is currently deactivating pages
> belonging to a container that has gone over the limit. Even though this
> allows containers to go over board their limits but 1- once they are
> over the limit then they run in degraded manner and 2- if there is any
> memory pressure then the (extra) pages belonging to this container are
> the prime candidates for swapping (for example). The statistics that
> are shown in each container directory are the current values of each
> resource consumption.
>
> Please let me know if you need any more specific information about the
> patch set.

A general description for the containers needs to go into Documentation/
along with the usage example, so that potential users know what to do
with it. The above seems fine, just remove the "this patch" references.

Subject: Re: [Patch 01/05]- Containers: Documentation on using containers
Posted by [Rohit Seth](#) on Fri, 15 Sep 2006 18:19:14 GMT
[View Forum Message](#) <> [Reply to Message](#)

On Fri, 2006-09-15 at 19:44 +0200, Björn Steinbrink wrote:
> On 2006.09.15 09:45:25 -0700, Rohit Seth wrote:
> > On Fri, 2006-09-15 at 08:15 +0200, Rolf Eike Beer wrote:
> > > Rohit Seth wrote:
> > > > This patch contains the Documentation for using containers.
> > >
> > > > +5- Remove a task from container
> > > > + echo <pid> rmtask
> > >
> > > echo <pid> > rmtask?
> > >
> >
> > rmtask is an attribute defined in test_container directory. So, first
> > you have to cd into container directory
> > (cd /mnt/configs/containers/test_container and then execute this

> > command)

>

> The > is missing in the patch though, guess that's what Rolf wanted to
> point out ;)

>

Aha. Will fix it.

> >

> > > Please also give a short description what containers are for. From what I read
> > > here I can only guess it's about getting some statistics about a group of
> > > tasks.

> >

> > Containers allow different workloads to be run on the same platform with
> > limits defined on per container basis. This basically allows a single
> > platform to be (soft) partitioned among different workloads (each of
> > which could be running many tasks). The limits could be amount of
> > memory, number of tasks among other features. These two features are
> > already implemented in the patch set that I posted. But it is possible
> > to add other controllers like CPU that allows only finite amount of time
> > to the processes belonging to a container.

> >

> > Currently this patch set is only tracking user memory (both file based
> > and anonymous). The memory handler is currently deactivating pages
> > belonging to a container that has gone over the limit. Even though this
> > allows containers to go over board their limits but 1- once they are
> > over the limit then they run in degraded manner and 2- if there is any
> > memory pressure then the (extra) pages belonging to this container are
> > the prime candidates for swapping (for example). The statistics that
> > are shown in each container directory are the current values of each
> > resource consumption.

> >

> > Please let me know if you need any more specific information about the
> > patch set.

>

> A general description for the containers needs to go into Documentation/
> along with the usage example, so that potential users know what to do
> with it. The above seems fine, just remove the "this patch" references.

>

Okay.

Thanks,
-rohit

Subject: Re: [Patch 01/05]- Containers: Documentation on using containers
Posted by [Valdis.Kletnieks](#) on Tue, 19 Sep 2006 15:32:19 GMT

On Thu, 14 Sep 2006 18:38:33 PDT, Rohit Seth said:

(Sorry for the late reply...)

```
> --- linux-2.6.18-rc6-mm2.org/Documentation/containers.txt 1969-12-31 16:00:00.000000000
-0800
> +++ linux-2.6.18-rc6-mm2.ctn/Documentation/containers.txt 2006-09-14 17:13:48.000000000
-0700
```

```
> +5- Remove a task from container
> + echo <pid> rmtask
```

```
echo <pid> > rmtask
```

is what I think was intended. Also, I'm not sure <pid> is the best meta-syntax - anybody got a better idea?

```
> +9- Freeing a container
> + cd /mnt/configfs/containers/
> + rmdir test_container
```

What happens if you try to remove a container that still has active tasks? Are you relying on the VFS to catch the 'non-empty directory'? (of course, 'rm -r' has predictable semantics here).

I see support for "add a task" and "remove a task", but none for listing the current tasks in the container?

Subject: Re: [Patch 01/05]- Containers: Documentation on using containers
Posted by [Rohit Seth](#) on Tue, 19 Sep 2006 16:30:24 GMT

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

On Tue, 2006-09-19 at 11:32 -0400, Valdis.Kletnieks@vt.edu wrote:

> On Thu, 14 Sep 2006 18:38:33 PDT, Rohit Seth said:

```
>
> (Sorry for the late reply...)
```

```
>
>
> > --- linux-2.6.18-rc6-mm2.org/Documentation/containers.txt 1969-12-31 16:00:00.000000000
-0800
> > +++ linux-2.6.18-rc6-mm2.ctn/Documentation/containers.txt 2006-09-14 17:13:48.000000000
-0700
>
> > +5- Remove a task from container
```

> > + echo <pid> rmtask
>
> echo <pid> > rmtask
>
> is what I think was intended. Also, I'm not sure <pid> is the best
> meta-syntax - anybody got a better idea?
>

Fixed that.

> > +9- Freeing a container
> > + cd /mnt/configfs/containers/
> > + rmdir test_container
>
> What happens if you try to remove a container that still has active tasks? Are
> you relying on the VFS to catch the 'non-empty directory'? (of course, 'rm -r'
> has predictable semantics here).
>

When a rmdir operation is done on a directory, then all tasks have to
move out of container (before we the rmdir operation is complete). This
removal of tasks from a container is done as part of
container_remove_tasks.

> I see support for "add a task" and "remove a task", but none for listing
> the current tasks in the container?

This is one of the TODO items that I've listed. Will be there soon.

thanks,
-rohit
