Subject: possible container states/statuses
Posted by knawnd on Mon, 16 May 2011 13:43:23 GMT
View Forum Message <> Reply to Message

Hi!

I wonder what the possible container states/statuses are and how they can be detected?
I read about 'status' option in 'man vzctl' which is below:

==========
status
Shows a container status. This is a line with five or six words, separated by spaces.

First word is literally CTID.

Second word is the numeric CT ID.

Third word is showing whether this container exists or not, it can be either exist or deleted.

Fourth word is showing the status of the container filesystem, it can be either mounted or unmounted.

Fifth word shows if the container is running, it can be either running or down.

Sixth word, if exists, is suspended. It appears if both a container and its dump file exist (see chkpnt).
[...]
==========

So I tried to summarize possible CT states as the following ('+' means states is possible and '-' it is impossible):
exist mounted down (-)
exist mounted running (+)
exist unmounted down (+)
exist unmounted running (-)
exist unmounted down suspended (+)
deleted unmounted down (+)
deleted mounted down (-)
deleted mounted running (-)
deleted unmounted down (-)

Could someone confirms these are all possible container states, please?
Can container for some reason crash? If it can then what the 'vzctl status' shows for it?

Thanks,
Nikolay.

Subject: Re: possible container states/statuses
Posted by kir on Mon, 16 May 2011 18:48:02 GMT
View Forum Message <> Reply to Message

Generally speaking, all the states (words) are independent, with a few
exceptions (mostly obvious and logical):

1. Running CT can't be unmounted.
2. Deleted CT can't be mounted or running or suspended (or so I think, might
be wrong here — if you manually delete running CT's configuration file it
might show as deleted mounted running).
3. (Contrary to common thinking) running CT can be suspended!
On May 16, 2011 5:47 PM, <knawnd@gmail.com> wrote:

Subject: Re: possible container states/statuses
Posted by KriS83 on Mon, 16 May 2011 19:18:25 GMT
View Forum Message <> Reply to Message

Hi,

Am 16.05.2011 15:43, schrieb knawnd@gmail.com:

> ==========
>
> So I tried to summarize possible CT states as the following ('+' means
> states is possible and '-' it is impossible):
> exist mounted down (-)

I'd think this one is actually possible?
A CT can exist, not run but be mounted.

Regards,
Kristian

Subject: Re: possible container states/statuses
Posted by knawnd on Tue, 17 May 2011 05:50:13 GMT
View Forum Message <> Reply to Message

Thanks everyone who replied!

Nikolay.

Kristian Marcroft wrote on 16/05/11 23:18:
> Hi,
>
> Am 16.05.2011 15:43, schrieb knawnd@gmail.com:
>
>> ==========
>>
>> So I tried to summarize possible CT states as the following ('+' means
>> states is possible and '-' it is impossible):
>> exist mounted down (-)
> I'd think this one is actually possible?
> A CT can exist, not run but be mounted.
>
> Regards,
> Kristian

Subject: Re: possible container states/statuses
Posted by knawnd on Tue, 17 May 2011 06:13:14 GMT
View Forum Message <> Reply to Message

Kir Kolyshkin wrote on 16/05/11 22:48:
>
> Generally speaking, all the states (words) are independent, with a few
> exceptions (mostly obvious and logical):
>
> 1. Running CT can't be unmounted.
> 2. Deleted CT can't be mounted or running or suspended (or so I think,
> might be wrong here — if you manually delete running CT's
> configuration file it might show as deleted mounted running).
> 3. (Contrary to common thinking) running CT can be suspended!
>
Does your last sentence mean "CT can be in running and suspended state"
i.e. 'vzctl status' has to show 'exist mounted running suspended'?
I've just tried to reproduce that state following [1] guide. So I did:
$ vzctl create 103 --ostemplate centos-5-x86
Creating container private area (centos-5-x86)
Performing postcreate actions
Container private area was created

$ vzctl start 103
Starting container ...
Container is mounted
Setting CPU units: 1000
Container start in progress...

```
$ vzlist
      CTID      NPROC STATUS    IP_ADDR         HOSTNAME
       103          6 running   -               -

$ vzctl chkpnt 103 --suspend
Setting up checkpoint...
        suspend...
        get context...
Checkpointing completed succesfully

$ vzctl status 103
CTID 103 exist mounted running

$ vzlist
      CTID      NPROC STATUS    IP_ADDR         HOSTNAME
       103         12 running   -               -

$vzctl chkpnt 103 --dump --dumpfile /vz/dump/dump.103
Setting up checkpoint...
        join context..
        dump...
Checkpointing completed succesfully

$ vzlist
      CTID      NPROC STATUS    IP_ADDR         HOSTNAME
       103         12 running   -               -

$ vzctl status 103
CTID 103 exist mounted running
```

As one can see 'vzctl status' doesn't show running CT as suspended. So what is the proper way to detect that state (running and suspended CT)?

Software versions used:

```
$ rpm -q vzctl
vzctl-3.0.26.3-1

$ uname -rsvpi
Linux 2.6.18-238.9.1.el5.028stab089.1 #1 SMP Thu Apr 14 14:34:26 MSD
2011 i686 i386
```

[1] http://wiki.openvz.org/Checkpointing_and_live_migration

> On May 16, 2011 5:47 PM, <knawnd@gmail.com <mailto:knawnd@gmail.com>>
> wrote:
>

>

Subject: Re: possible container states/statuses
Posted by knawnd on Tue, 17 May 2011 06:48:40 GMT
View Forum Message <> Reply to Message

> $ vzctl status 103
> CTID 103 exist mounted running
>
> As one can see 'vzctl status' doesn't show running CT as suspended. So
> what is the proper way to detect that state (running and suspended CT)?
I guess because 'suspended' is missing from the output of 'vzctl status' command for such CT the attempt to stop a CT in such state (running and suspended) has to be properly traced. For example, 'vzctl stop <CT_ID>' has to print something like "CT is running and suspended. Please, resume it by executing 'vzctl chkpnt <CTID> --resume' command in order to stop it". Right now the action 'vzctl stop <CTID>' on running and suspended CT ends with "Unable to stop container: operation timed out" what is expected.

Regards,
Nikolay.