
Subject: How to determine a container from the filesystem?
Posted by [Brad Alexander](#) on Fri, 13 Apr 2012 20:07:58 GMT
[View Forum Message](#) <> [Reply to Message](#)

I just found out through the proxmox-ve forums that running ntp on a container is considered a Bad Thing. So I am reworking my puppet installation to disable ntp on the containers...But I was trying to figure out a foolproof way of looking on the machine and determining if it is a container or not. The only thing I have found so far is that /proc/mtrr exists on the physical servers, but not on the containers. Is this a viable way to make this determination or is there a better way?

Thanks,
--b

Subject: Re: How to determine a container from the filesystem?
Posted by [fruitwerks](#) on Fri, 13 Apr 2012 22:24:03 GMT
[View Forum Message](#) <> [Reply to Message](#)

You could run ifconfig and grep for venetX:X or by mac address (all zero) unless you have changed that specifically. I physical machine should not have venetX:X, simply venetX. This may be distribution dependent though, I am not sure.

- C

On Fri, Apr 13, 2012 at 8:07 PM, Brad Alexander <storm16@gmail.com> wrote:

> I just found out through the proxmox-ve forums that running ntp on a
> container is considered a Bad Thing. So I am reworking my puppet
> installation to disable ntp on the containers...But I was trying to
> figure out a foolproof way of looking on the machine and determining
> if it is a container or not. The only thing I have found so far is
> that /proc/mtrr exists on the physical servers, but not on the
> containers. Is this a viable way to make this determination or is
> there a better way?
>
> Thanks,
> --b
--

Any use, dissemination, distribution, posting on Internet bulletin boards, disclosure or copying of this e-mail or any information contained herein by or to anyone other than the intended recipient(s) is strictly prohibited. Use of this content for any other purpose is a violation of International

Subject: Re: How to determine a container from the filesystem?

Posted by [efball](#) on Fri, 13 Apr 2012 22:33:21 GMT

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

On Fri, Apr 13, 2012 at 04:07:58PM -0400, Brad Alexander wrote:

> I just found out through the proxmox-ve forums that running ntp on a
> container is considered a Bad Thing. So I am reworking my puppet
> installation to disable ntp on the containers...But I was trying to
> figure out a foolproof way of looking on the machine and determining
> if it is a container or not. The only thing I have found so far is
> that /proc/mtrr exists on the physical servers, but not on the
> containers. Is this a viable way to make this determination or is
> there a better way?

if [-d /proc/vz -a ! -d /proc/bc]

/proc/vz - always exists if OpenVZ kernel is running (inside
and outside container) /proc/bc - exists on node, but not
inside container.

--

E Frank Ball efball@efball.com

Subject: Re: How to determine a container from the filesystem?

Posted by [jjs - mainphrame](#) on Fri, 13 Apr 2012 22:36:22 GMT

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

Unfortunately that won't work if you are using only bridged networking - I
don't have any venet devices on my servers.

My host has only lo, ethx, brx, and vethnxx devices, and the containers
have only lo and ethx devices.

The puppet "factor" program is able to figure out if a machine is a vz
CT or a vz host, but I haven't looked into how it does it.

Joe

On Fri, Apr 13, 2012 at 3:24 PM, Corey Carpenter <fruitwerks@gmail.com>wrote:

> You could run ifconfig and grep for venetX:X or by mac address (all zero)

> unless you have changed that specifically. I physical machine should not
> have venetX:X, simply venetX. This may be distribution dependent though, I
> am not sure.
>
> - C
>
>
> On Fri, Apr 13, 2012 at 8:07 PM, Brad Alexander <storm16@gmail.com> wrote:
>
>> I just found out through the proxmox-ve forums that running ntp on a
>> container is considered a Bad Thing. So I am reworking my puppet
>> installation to disable ntp on the containers...But I was trying to
>> figure out a foolproof way of looking on the machine and determining
>> if it is a container or not. The only thing I have found so far is
>> that /proc/mtrr exists on the physical servers, but not on the
>> containers. Is this a viable way to make this determination or is
>> there a better way?
>>
>> Thanks,
>> --b
> --
>
> _____

> Any use, dissemination, distribution, posting on Internet bulletin boards,
> disclosure or copying of this e-mail or any information contained herein by
> or to anyone other than the intended recipient(s) is strictly prohibited.
> Use of this content for any other purpose is a violation of International
> Copyright Laws.
>
>

Subject: Re: How to determine a container from the filesystem?

Posted by [Martin Dobrev](#) on Fri, 13 Apr 2012 22:39:28 GMT

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

Better way to do it is to look for /proc/user_beancounters. If it exists then it's a distro with OpenVZ kernel installation. In it there is a info about different parameters of the container (if you look into it inside the container) or containers (if checked from the HN). Container 0: is the HN, so if you have it listed in the file then you run outside of the container.

Martin Dobrev

Sent from iPhonespam SPAMSPAM 4

On 14.04.2012, at 01:24, Corey Carpenter <fruitwerks@gmail.com> wrote:

> You could run ifconfig and grep for venetX:X or by mac address (all zero) unless you have changed that specifically. I physical machine should not have venetX:X, simply venetX. This may be distribution dependent though, I am not sure.

>

> - C

>

> On Fri, Apr 13, 2012 at 8:07 PM, Brad Alexander <storm16@gmail.com> wrote:

> I just found out through the proxmox-ve forums that running ntp on a

> container is considered a Bad Thing. So I am reworking my puppet

> installation to disable ntp on the containers...But I was trying to

> figure out a foolproof way of looking on the machine and determining

> if it is a container or not. The only thing I have found so far is

> that /proc/mtrr exists on the physical servers, but not on the

> containers. Is this a viable way to make this determination or is

> there a better way?

>

> Thanks,

> --b

> --

>

> Any use, dissemination, distribution, posting on Internet bulletin boards, disclosure or copying of this e-mail or any information contained herein by or to anyone other than the intended recipient(s) is strictly prohibited. Use of this content for any other purpose is a violation of International Copyright Laws.

>

Subject: Re: How to determine a container from the filesystem?

Posted by [Brad Alexander](#) on Sat, 14 Apr 2012 01:40:09 GMT

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

Thanks Joe!

I didn't realize that facter had that level of detail. For my opevz hosts, virtual => openvzhn, but the containers all have virtual => openvzve...

--b

On Fri, Apr 13, 2012 at 6:36 PM, jjs - mainphrame <jjs@mainphrame.com> wrote:

> Unfortunately that won't work if you are using only bridged networking - I

> don't have any venet devices on my servers.

>

> My host has only lo, ethx, brx, and vethnxx devices, and the containers

> have only lo and ethx devices.

>

> The puppet "facter" program is able to figure out if a machine is a vz CT

> or a vz host, but I haven't looked into how it does it.

>

> Joe

>

>

> On Fri, Apr 13, 2012 at 3:24 PM, Corey Carpenter <fruitwerks@gmail.com>

> wrote:

>>

>> You could run ifconfig and grep for venetX:X or by mac address (all zero)

>> unless you have changed that specifically. I physical machine should not

>> have venetX:X, simply venetX. This may be distribution dependent though, I

>> am not sure.

>>

>> - C

>>

>>

>> On Fri, Apr 13, 2012 at 8:07 PM, Brad Alexander <storm16@gmail.com> wrote:

>>>

>>> I just found out through the proxmox-ve forums that running ntp on a

>>> container is considered a Bad Thing. So I am reworking my puppet

>>> installation to disable ntp on the containers...But I was trying to

>>> figure out a foolproof way of looking on the machine and determining

>>> if it is a container or not. The only thing I have found so far is

>>> that /proc/mtrr exists on the physical servers, but not on the

>>> containers. Is this a viable way to make this determination or is

>>> there a better way?

>>>

>>> Thanks,

>>> --b

>> --

>>

>>

>> Any use, dissemination, distribution, posting on Internet bulletin boards,

>> disclosure or copying of this e-mail or any information contained herein by

>> or to anyone other than the intended recipient(s) is strictly prohibited.

>> Use of this content for any other purpose is a violation of International

>> Copyright Laws.

>>

>>

Subject: Re: How to determine a container from the filesystem?

Posted by [jjs - mainphrame](#) on Sat, 14 Apr 2012 03:19:06 GMT

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

That's another good tip.

Joe

On Fri, Apr 13, 2012 at 3:39 PM, Martin Dobrev <martin@dobrev.eu> wrote:

> Better way to do it is to look for /proc/user_beancounters. If it exists
> then it's a distro with OpenVZ kernel installation. In it there is a info
> about different parameters of the container (if you look into it inside the
> container) or containers (if checked from the HN). Container 0: is the HN,
> so if you have it listed in the file then you run outside of the container.

>
> Martin Dobrev

>
> Sent from iPhonespam SPAMSPAM 4

> On 14.04.2012, at 01:24, Corey Carpenter <fruitwerks@gmail.com> wrote:

>
> You could run ifconfig and grep for venetX:X or by mac address (all zero)
> unless you have changed that specifically. I physical machine should not
> have venetX:X, simply venetX. This may be distribution dependent though, I
> am not sure.

>
> - C

>
> On Fri, Apr 13, 2012 at 8:07 PM, Brad Alexander <storm16@gmail.com> wrote:

>
>> I just found out through the proxmox-ve forums that running ntp on a
>> container is considered a Bad Thing. So I am reworking my puppet
>> installation to disable ntp on the containers...But I was trying to
>> figure out a foolproof way of looking on the machine and determining
>> if it is a container or not. The only thing I have found so far is
>> that /proc/mtrr exists on the physical servers, but not on the
>> containers. Is this a viable way to make this determination or is
>> there a better way?

>>
>> Thanks,
>> --b

> --

>

> _____

> Any use, dissemination, distribution, posting on Internet bulletin boards,
> disclosure or copying of this e-mail or any information contained herein by
> or to anyone other than the intended recipient(s) is strictly prohibited.
> Use of this content for any other purpose is a violation of International
> Copyright Laws.

>

Subject: Re: How to determine a container from the filesystem?

Posted by [kir](#) on Tue, 17 Apr 2012 07:29:23 GMT

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

On 04/14/2012 12:07 AM, Brad Alexander wrote:

> I just found out through the proxmox-ve forums that running ntp on a
> container is considered a Bad Thing.

Not necessarily. In fact, it's a good thing to run ntpd inside a container, it's just you need to

1. Have only ONE container doing that.
2. Grant that container sys_time capability, so it will be able to set system time.

This is because time is not virtualized, ie all the containers share the same time (because indeed there's only one time -- time zones of course can be different).

> So I am reworking my puppet
> installation to disable ntp on the containers...But I was trying to
> figure out a foolproof way of looking on the machine and determining
> if it is a container or not. The only thing I have found so far is
> that /proc/mtrr exists on the physical servers, but not on the
> containers. Is this a viable way to make this determination or is
> there a better way?

Solutions provided here in this thread by E Frank Ball and Martin Dobrev are both good.

Subject: Re: How to determine a container from the filesystem?

Posted by [Brad Alexander](#) on Tue, 17 Apr 2012 11:07:43 GMT

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

Thanks Kir.

On Tue, Apr 17, 2012 at 3:29 AM, Kir Kolyshkin <kir@openvz.org> wrote:

> On 04/14/2012 12:07 AM, Brad Alexander wrote:

>>

>> I just found out through the proxmox-ve forums that running ntp on a
>> container is considered a Bad Thing.

>

>

> Not necessarily. In fact, it's a good thing to run ntpd inside a container,
> it's just you need to

>

> 1. Have only ONE container doing that.

So that one container can be Container 0 (the HN)?

> 2. Grant that container sys_time capability, so it will be able to set
> system time.

Perhaps I misunderstood the sys_time flag, it was my understanding that it was better to turn off ntp on the containers, make sure it is on in container 0 (the hardware node), then turn on sys_time on the remaining containers.

> This is because time is not virtualized, ie all the containers share the
> same time (because indeed there's only one time -- time zones of course can
> be different).

Thanks,
--b

Subject: Re: How to determine a container from the filesystem?

Posted by [kir](#) on Tue, 17 Apr 2012 11:14:21 GMT

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

On 04/17/2012 03:07 PM, Brad Alexander wrote:

> Thanks Kir.

>

> On Tue, Apr 17, 2012 at 3:29 AM, Kir Kolyshkin<kir@openvz.org> wrote:

>> On 04/14/2012 12:07 AM, Brad Alexander wrote:

>>> I just found out through the proxmox-ve forums that running ntp on a

>>> container is considered a Bad Thing.

>>

>> Not necessarily. In fact, it's a good thing to run ntpd inside a container,

>> it's just you need to

>>

>> 1. Have only ONE container doing that.

> So that one container can be Container 0 (the HN)?

Yes, but from the privilege separation perspective it might make sense to have a dedicated container for that, so you don't clog HN with all sorts of services and daemons.

>

>> 2. Grant that container sys_time capability, so it will be able to set

>> system time.

> Perhaps I misunderstood the sys_time flag, it was my understanding

> that it was better to turn off ntp on the containers

Right, it doesn't make sense to run ntpd in more than one container (or HN).

> , make sure it is
> on in container 0 (the hardware node)

Right. Or any other `_single_` container.

> , then turn on `sys_time` on the
> remaining containers.

Ughm. That way, root user of any of those container can change system time (and affect other users of CTs on the same HN).

>
>> This is because time is not virtualized, ie all the containers share the
>> same time (because indeed there's only one time -- time zones of course can
>> be different).
> Thanks,
> --b
