
Subject: What is OpenVZ container scheduling granularity

Posted by [shule ney](#) on Tue, 20 Dec 2011 20:21:50 GMT

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

Hi all:

I'm eager to know what is OpenVZ container scheduling granularity, 1ms or something??? I really need information about this.

Subject: Re: What is OpenVZ container scheduling granularity

Posted by [Kirill Korotaev](#) on Wed, 21 Dec 2011 07:58:19 GMT

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

It's floating, depends on priorities. Plus more important for latency is not granularity, but preemptiveness.

Sent from my iPhonespam SPAMSPAM

On 21.12.2011, at 0:34, "shule ney" <neyshule@gmail.com> wrote:

> Hi all:

> I'm eager to know what is OpenVZ container scheduling granularity, 1ms or something??? I really need information about this.

> <ATT00001.c>

Subject: Re: What is OpenVZ container scheduling granularity

Posted by [shule ney](#) on Wed, 21 Dec 2011 15:50:01 GMT

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

Much thanks Kirill, I really appreciate your reply! My question is:

Suppose two containers exist on my machine which can use 0%-100% CPU, each of them has only one active process. If I sleep one container's process for 1us which makes this container has nothing to do, will the the container be scheduled off and the other container gets scheduled? Is 1us too small for container scheduling?? I want to know if this case is possible. Thanks very much.

2011/12/21 Kirill Korotaev <dev@parallels.com>

> It's floating, depends on priorities. Plus more important for latency is
> not granularity, but preemptiveness.

>

> Sent from my iPhonespam SPAMSPAM

>

> On 21.12.2011, at 0:34, "shule ney" <neyshule@gmail.com> wrote:

>

> > Hi all:
> > I'm eager to know what is OpenVZ container scheduling granularity, 1ms
> or something??? I really need information about this.
> > <ATT00001.c>
>

Subject: Re: What is OpenVZ container scheduling granularity
Posted by [Kirill Korotaev](#) on Wed, 21 Dec 2011 16:14:26 GMT
[View Forum Message](#) <> [Reply to Message](#)

if CPU has nothing to do since your app went to sleep (even for 1us), it will be rescheduled to another CPU. Just like for conventional tasks in Linux.

On Dec 21, 2011, at 19:50 , shule ney wrote:

> Much thanks Kirill, I really appreciate your reply! My question is:
> Suppose two containers exist on my machine which can use 0%-100% CPU, each of them has only one active process. If I sleep one container's process for 1us which makes this container has nothing to do, will the the container be scheduled off and the other container gets scheduled? Is 1us too small for container scheduling?? I want to know if this case is possible. Thanks very much.
>
> 2011/12/21 Kirill Korotaev <dev@parallels.com>
> It's floating, depends on priorities. Plus more important for latency is not granularity, but preemptiveness.
>
> Sent from my iPhonespam SPAMSPAM
>
> On 21.12.2011, at 0:34, "shule ney" <neyshule@gmail.com> wrote:
>
> > Hi all:
> > I'm eager to know what is OpenVZ container scheduling granularity, 1ms or something??? I really need information about this.
> > <ATT00001.c>
>

Subject: Re: What is OpenVZ container scheduling granularity
Posted by [MailingListe](#) on Wed, 21 Dec 2011 16:24:35 GMT
[View Forum Message](#) <> [Reply to Message](#)

Zitat von shule ney <neyshule@gmail.com>:

> Much thanks Kirill, I really appreciate your reply! My question is:
> Suppose two containers exist on my machine which can use 0%-100% CPU, each

> of them has only one active process. If I sleep one container's process for
> 1us which makes this container has nothing to do, will the the container be
> scheduled off and the other container gets scheduled? Is 1us too small for
> container scheduling?? I want to know if this case is possible. Thanks very
> much.

OpenVZ is shared Kernel so you must think of scheduling per
process(-group) like in a conventional Kernel, not per container.

Regards

Andreas

File Attachments

1) [smime.p7s](#), downloaded 440 times

Subject: Re: What is OpenVZ container scheduling granularity
Posted by [shule ney](#) on Wed, 21 Dec 2011 16:49:24 GMT

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

Thanks very much lst_hoe02, if I'm using kernel-2.6.24 which uses CFS
scheduler with nanosecond *granularity, processes *across* different
containers/groups should also have this * *granularity all right? *

2011/12/21 <lst_hoe02@kwsoft.de>

> Zitat von shule ney <neyshule@gmail.com>:
>
>
> Much thanks Kirill, I really appreciate your reply! My question is:
>> Suppose two containers exist on my machine which can use 0%-100% CPU,
>> each
>> of them has only one active process. If I sleep one container's process
>> for
>> 1us which makes this container has nothing to do, will the the container
>> be
>> scheduled off and the other container gets scheduled? Is 1us too small for
>> container scheduling?? I want to know if this case is possible. Thanks
>> very
>> much.
>>
>
> OpenVZ is shared Kernel so you must think of scheduling per
> process(-group) like in a conventional Kernel, not per container.
>
> Regards
>

> Andreas
>
>
>

Subject: Re: What is OpenVZ container scheduling granularity
Posted by [dowdle](#) on Wed, 21 Dec 2011 17:04:08 GMT
[View Forum Message](#) <> [Reply to Message](#)

Greetings,

----- Original Message -----

> Thanks very much lst_hoe02, if I'm using kernel-2.6.24 which uses CFS
> scheduler with nanosecond granularity, processes across different
> containers/groups should also have this granularity all right?

Just to clarify, the OpenVZ project has dropped the 2.6.24 branch. What you have may still be supported by Debian/Ubuntu perhaps?

For info about the various OpenVZ kernel branches, please see:

<http://wiki.openvz.org/Download/kernel>

Just saying.

TYL,

--

Scott Dowdle
704 Church Street
Belgrade, MT 59714
(406)388-0827 [home]
(406)994-3931 [work]

Subject: Re: What is OpenVZ container scheduling granularity
Posted by [shule ney](#) on Wed, 21 Dec 2011 17:18:37 GMT
[View Forum Message](#) <> [Reply to Message](#)

I'm using the old 2.6.24/2.6.24-ovz009.1. It works well for me.

2011/12/21 Scott Dowdle <dowdle@montanlinux.org>

> Greetings,

>

> ----- Original Message -----

> > Thanks very much lst_hoe02, if I'm using kernel-2.6.24 which uses CFS

> > scheduler with nanosecond granularity, processes across different
> > containers/groups should also have this granularity all right?
>
> Just to clarify, the OpenVZ project has dropped the 2.6.24 branch. What
> you have may still be supported by Debian/Ubuntu perhaps?
>
> For info about the various OpenVZ kernel branches, please see:
>
> <http://wiki.openvz.org/Download/kernel>
>
> Just saying.
>
> TYL,
> --
> Scott Dowdle
> 704 Church Street
> Belgrade, MT 59714
> (406)388-0827 [home]
> (406)994-3931 [work]

Subject: Re: What is OpenVZ container scheduling granularity
Posted by [MailingListe](#) on Thu, 22 Dec 2011 08:30:08 GMT

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

Zitat von shule ney <neyshule@gmail.com>:

> Thanks very much lst_hoe02, if I'm using kernel-2.6.24 which uses CFS
> scheduler with nanosecond *granularity, processes *across* different
> containers/groups should also have this * *granularity all right? *

To my knowledge OpenVZ does not alter the scheduler, but only
priorities. So yes it should work, but others may have more insight
then me.

Regards

Andreas

File Attachments

1) [smime.p7s](#), downloaded 407 times

Subject: Re: What is OpenVZ container scheduling granularity
Posted by [shule ney](#) on Thu, 22 Dec 2011 14:59:01 GMT

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

Much thanks for your help lst_hoe02[?]

2011/12/22 <lst_hoe02@kwsoft.de>

> Zitat von shule ney <neyshule@gmail.com>:
>
> Thanks very much lst_hoe02, if I'm using kernel-2.6.24 which uses CFS
>> scheduler with nanosecond *granularity, processes *across* different
>> containers/groups should also have this * *granularity all right? *
>>
>
> To my knowledge OpenVZ does not alter the scheduler, but only priorities.
> So yes it should work, but others may have more insight then me.
>
> Regards
>
> Andreas
>
>
>
