

---

Subject: Measuring and Adjusting CPU utilization

Posted by [Pradeep Padala](#) on Tue, 06 Jun 2006 18:03:37 GMT

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

---

Hi,

I am trying to measure the CPU utilization of the VZ containers, and change the cpu share dynamically. I have poured over most of the documentation, and looked at the code as well, and it seems like there's no utility that can directly show the current CPU utilization of a container (some thing like 30% of CPU). A search on the user list got me a message, where someone suggested using loadavg. However, It seems like the loadavg is not showing the proper utilization (or showing the total CPU utilization). This is what I am doing.

I setup a container with 1000 units limit (total CPU units: ~10000). I wrote a small `do {; }while(1);` loop and ran it in the container, now I do `cat /proc/loadavg` in both the container and on the host node. Since, the container is only using 1000 units, I should see something like 100% loadavg in the container, and 10% loadavg in the hostnode. But, I see 100% at both places. Am I doing something wrong? How do I get the current cpu utilization of a container?

Thanks,

--

Pradeep Padala

<http://ppadala.blogspot.com>

---

---

Subject: Re: Measuring and Adjusting CPU utilization

Posted by [kir](#) on Tue, 06 Jun 2006 23:17:15 GMT

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

---

Looks like you misunderstand the concept of cpunits. cpunits is not a hard limit, but just a suggestion, and a CPU time is shared proportionally to the values given. So, if you will have 9 VEs and the host system with cpunits set to 1000 for all of them, and run the loop in all of them, each VE will use 10% of the CPU time.

In case you will stop the loop running in 5 VEs so there will be 4 such VEs (plus the host system) left, each of them will use 20% of the CPU. So, all the CPU time is distributed between VEs which will need it, according with their proportional cpunits.

More to say, the concept of "total CPU units" is purely fiction, and is here just for the convenience. People do want to set CPU units in terms of processor's megahertz, and this is what cpunits does. But in fact it is not a megahertz but just a relative weights. I.e. all the cpunits

values are relative to each other, it doesn't matter what the actual numbers are -- what matters is a number given to a VE in relation to the sum of all cpuweights (which is expressed as "total CPU units" just for the convenience).

So, cpuunits, if you do not oversell them, are a CPU guarantee, not a limit. If you want CPU limit -- use cpulimit parameter.

Pradeep Padala wrote:

> Hi,  
>  
> I am trying to measure the CPU utilization of the VZ containers, and  
> change the cpu share dynamically. I have poured over most of the  
> documentation, and looked at the code as well, and it seems like  
> there's no utility that can directly show the current CPU utilization  
> of a container (some thing like 30% of CPU). A search on the user list  
> got me a message, where someone suggested using loadavg. However, It  
> seems like the loadavg is not showing the proper utilization (or  
> showing the total CPU utilization). This is what I am doing.  
>  
> I setup a container with 1000 units limit (total CPU units: ~10000). I  
> wrote a small do {; }while(1); loop and ran it in the container, now I  
> do cat /proc/loadavg in both the container and on the host node.  
> Since, the container is only using 1000 units, I should see something  
> like 100% loadavg in the container, and 10% loadavg in the hostnode.  
> But, I see 100% at both places. Am I doing something wrong? How do I  
> get the current cpu utilization of a container?  
>  
> Thanks,  
> --  
> Pradeep Padala  
> <http://ppadala.blogspot.com>  
>  
> -----  
>

---

Subject: Re: Measuring and Adjusting CPU utilization  
Posted by [Pradeep Padala](#) on Tue, 06 Jun 2006 23:27:32 GMT  
[View Forum Message](#) <> [Reply to Message](#)

Thanks for the explanation, but I DID use the --cpulimit parameter.

I set the limit to 1000 units, I run a while(1) loop, and I see a 100% loadavg on the host node. Shouldn't it be somewhere around 10% ?

Pradeep

On 6/6/06, Kir Kolyshkin <kir@openvz.org> wrote:

>  
> Looks like you misunderstand the concept of cpuunits. cpuunits is not a  
> hard limit, but just a suggestion, and a CPU time is shared  
> proportionally to the values given. So, if you will have 9 VEs and the  
> host system with cpuunits set to 1000 for all of them, and run the loop  
> in all of them, each VE will use 10% of the CPU time.  
>  
> In case you will stop the loop running in 5 VEs so there will be 4 such  
> VEs (plus the host system) left, each of them will use 20% of the CPU.  
> So, all the CPU time is distributed between VEs which will need it,  
> according with their proportional cpuunits.  
>  
> More to say, the concept of "total CPU units" is purely fiction, and is  
> here just for the convenience. People do want to set CPU units in terms  
> of processor's megahertz, and this is what cpuunits does. But in fact it  
> is not a megahertz but just a relative weights. I.e. all the cpuunits  
> values are relative to each other, it doesn't matter what the actual  
> numbers are -- what matters is a number given to a VE in relation to the  
> sum of all cpuweights (which is expressed as "total CPU units" just for  
> the convenience).  
>  
> So, cpuunits, if you do not oversell them, are a CPU guarantee, not a  
> limit. If you want CPU limit -- use cpulimit parameter.  
>  
> Pradeep Padala wrote:  
>  
>> Hi,  
>>  
>> I am trying to measure the CPU utilization of the VZ containers, and  
>> change the cpu share dynamically. I have poured over most of the  
>> documentation, and looked at the code as well, and it seems like  
>> there's no utility that can directly show the current CPU utilization  
>> of a container (some thing like 30% of CPU). A search on the user list  
>> got me a message, where someone suggested using loadavg. However, It  
>> seems like the loadavg is not showing the proper utilization (or  
>> showing the total CPU utilization). This is what I am doing.  
>>  
>> I setup a container with 1000 units limit (total CPU units: ~10000). I  
>> wrote a small do {; }while(1); loop and ran it in the container, now I  
>> do cat /proc/loadavg in both the container and on the host node.  
>> Since, the container is only using 1000 units, I should see something  
>> like 100% loadavg in the container, and 10% loadavg in the hostnode.  
>> But, I see 100% at both places. Am I doing something wrong? How do I  
>> get the current cpu utilization of a container?  
>>  
>> Thanks,

> > --  
> > Pradeep Padala  
> > <http://ppadala.blogspot.com>  
> >  
> > -----  
> >  
--  
Pradeep Padala  
<http://ppadala.blogspot.com>

---

---

Subject: Re: Measuring and Adjusting CPU utilization  
Posted by [kir](#) on Wed, 07 Jun 2006 00:12:43 GMT  
[View Forum Message](#) <> [Reply to Message](#)

---

CPU limit is in per cent units. I.e. if your server has a single CPU, use --cpulimit 10 to limit a VE to 10% of the CPU.

On a two-way SMP box max. value of cpulimit is 200. Say, if you want a VE to use no more than one CPUs, use --cpulimit 100.

Pradeep Padala wrote:

> Thanks for the explanation, but I DID use the --cpulimit parameter.  
>  
> I set the limit to 1000 units, I run a while(1) loop, and I see a 100%  
> loadavg on the host node. Shouldn't it be somewhere around 10% ?  
>  
> Pradeep  
>  
> On 6/6/06, \*Kir Kolyshkin\* <[kir@openvz.org](mailto:kir@openvz.org)> <<mailto:kir@openvz.org>>>  
> wrote:  
>  
> Looks like you misunderstand the concept of cpuunits. cpuunits is  
> not a  
> hard limit, but just a suggestion, and a CPU time is shared  
> proportionally to the values given. So, if you will have 9 VEs and the  
> host system with cpuunits set to 1000 for all of them, and run the  
> loop  
> in all of them, each VE will use 10% of the CPU time.  
>  
> In case you will stop the loop running in 5 VEs so there will be 4  
> such  
> VEs (plus the host system) left, each of them will use 20% of the CPU.  
> So, all the CPU time is distributed between VEs which will need it,  
> according with their proportional cpuunits.  
>  
> More to say, the concept of "total CPU units" is purely fiction,

> and is  
> here just for the convenience. People do want to set CPU units is  
> terms  
> of processor's megahertz, and this is what cpunits does. But in  
> fact it  
> is not a megahertz but just a relative weights. I.e. all the cpunits  
> values are relative to each other, it doesn't matter what the actual  
> numbers are -- what matters is a number given to a VE in relation  
> to the  
> sum of all cpuweights (which is expressed as "total CPU units"  
> just for  
> the convenience).

>  
> So, cpunits, if you do not oversell them, are a CPU guarantee, not a  
> limit. If you want CPU limit -- use cpulimit parameter.

>  
> Pradeep Padala wrote:

>  
> > Hi,  
> >  
> > I am trying to measure the CPU utilization of the VZ  
> containers, and  
> > change the cpu share dynamically. I have poured over most of the  
> > documentation, and looked at the code as well, and it seems like  
> > there's no utility that can directly show the current CPU  
> utilization  
> > of a container (some thing like 30% of CPU). A search on the  
> user list  
> > got me a message, where someone suggested using loadavg.  
> However, It  
> > seems like the loadavg is not showing the proper utilization (or  
> > showing the total CPU utilization). This is what I am doing.  
> >  
> > I setup a container with 1000 units limit (total CPU units:  
> ~10000). I  
> > wrote a small `do {; }while(1);` loop and ran it in the container,  
> now I  
> > do `cat /proc/loadavg` in both the container and on the host node.  
> > Since, the container is only using 1000 units, I should see  
> something  
> > like 100% loadavg in the container, and 10% loadavg in the hostnode.  
> > But, I see 100% at both places. Am I doing something wrong? How do I  
> > get the current cpu utilization of a container?  
> >  
> > Thanks,  
> > --  
> > Pradeep Padala  
> > <http://ppadala.blogspot.com>

> >  
> > -----  
> >  
> >  
> >  
> >Users mailing list  
> >Users@openvz.org <mailto:Users@openvz.org>  
> >https://openvz.org/mailman/listinfo/users  
> <https://openvz.org/mailman/listinfo/users>  
> >  
> >  
>  
>  
> -----  
> Users mailing list  
> Users@openvz.org <mailto:Users@openvz.org>  
> https://openvz.org/mailman/listinfo/users  
> <https://openvz.org/mailman/listinfo/users>  
>  
>  
>  
>  
> --  
> Pradeep Padala  
> http://ppadala.blogspot.com  
>  
> -----  
>

---

Subject: Re: Measuring and Adjusting CPU utilization  
Posted by [Pradeep Padala](#) on Wed, 07 Jun 2006 00:33:23 GMT  
[View Forum Message](#) <> [Reply to Message](#)

---

On 6/6/06, Kir Kolyshkin <kir@openvz.org> wrote:  
>  
> CPU limit is in per cent units. I.e. if your server has a single CPU,  
> use --cpulimit 10 to limit a VE to 10% of the CPU.  
>  
> On a two-way SMP box max. value of cpulimit is 200. Say, if you want a  
> VE to use no more than one CPUs, use --cpulimit 100.

Thanks ! This solved the problem. However, I still see the loadavg at 100% (got from cat /proc/loadavg), but if I run top and see the CPU% for the one particular application (while(1) loop), it is limited to 10% CPU.

I can add up the CPU% for all the processes running in the container, and can get the current % utilization of the container. Is there a better way to

do this?

Pradeep

Pradeep Padala wrote:

```
>
>> Thanks for the explanation, but I DID use the --cpulimit parameter.
>>
>> I set the limit to 1000 units, I run a while(1) loop, and I see a 100%
>> loadavg on the host node. Shouldn't it be somewhere around 10% ?
>>
>> Pradeep
>>
>> On 6/6/06, *Kir Kolyshkin* <kir@openvz.org <mailto:kir@openvz.org>>
>> wrote:
>>
>> Looks like you misunderstand the concept of cpuunits. cpuunits is
>> not a
>> hard limit, but just a suggestion, and a CPU time is shared
>> proportionally to the values given. So, if you will have 9 VEs and
> the
>> host system with cpuunits set to 1000 for all of them, and run the
>> loop
>> in all of them, each VE will use 10% of the CPU time.
>>
>> In case you will stop the loop running in 5 VEs so there will be 4
>> such
>> VEs (plus the host system) left, each of them will use 20% of the
> CPU.
>> So, all the CPU time is distributed between VEs which will need it,
>> according with their proportional cpuunits.
>>
>> More to say, the concept of "total CPU units" is purely fiction,
>> and is
>> here just for the convenience. People do want to set CPU units is
>> terms
>> of processor's megahertz, and this is what cpuunits does. But in
>> fact it
>> is not a megahertz but just a relative weights. I.e. all the
> cpuunits
>> values are relative to each other, it doesn't matter what the actual
>> numbers are -- what matters is a number given to a VE in relation
>> to the
>> sum of all cpuweights (which is expressed as "total CPU units"
>> just for
>> the convenience).
>>
```

>> So, cpuunits, if you do not oversell them, are a CPU guarantee, not  
> a  
>> limit. If you want CPU limit -- use cpulimit parameter.  
>>  
>> Pradeep Padala wrote:  
>>  
>> > Hi,  
>> >  
>> > I am trying to measure the CPU utilization of the VZ  
>> containers, and  
>> > change the cpu share dynamically. I have poured over most of the  
>> > documentation, and looked at the code as well, and it seems like  
>> > there's no utility that can directly show the current CPU  
>> utilization  
>> > of a container (some thing like 30% of CPU). A search on the  
>> user list  
>> > got me a message, where someone suggested using loadavg.  
>> However, It  
>> > seems like the loadavg is not showing the proper utilization (or  
>> > showing the total CPU utilization). This is what I am doing.  
>> >  
>> > I setup a container with 1000 units limit (total CPU units:  
>> ~10000). I  
>> > wrote a small do {; }while(1); loop and ran it in the container,  
>> now I  
>> > do cat /proc/loadavg in both the container and on the host node.  
>> > Since, the container is only using 1000 units, I should see  
>> something  
>> > like 100% loadavg in the container, and 10% loadavg in the  
> hostnode.  
>> > But, I see 100% at both places. Am I doing something wrong? How do  
> I  
>> > get the current cpu utilization of a container?  
>> >  
>> > Thanks,  
>  
>

---

Subject: Re: Measuring and Adjusting CPU utilization

Posted by [dev](#) on Wed, 07 Jun 2006 08:20:05 GMT

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

---

> Thanks ! This solved the problem. However, I still see the loadavg at  
> 100% (got from cat /proc/loadavg), but if I run top and see the CPU% for  
> the one particular application (while(1) loop), it is limited to 10% CPU.  
This is also correct, since load average means 'average number of  
processes in running + uninterruptible state'. Since you busy loop is



constantly in running state, it is accounted as 1.0 in loadavg.

> I can add up the CPU% for all the processes running in the container,  
> and can get the current % utilization of the container. Is there a  
> better way to do this?

Sorry, didn't got what you mean... Probably you ask how to calculate the amount of CPU being consumed by VPS in %?

To do so, you need to monitor /proc/vz/vestat file.

[http://forum.openvz.org/index.php?t=msg&goto=2790&&a mp;srch=vestat#msg\\_2790](http://forum.openvz.org/index.php?t=msg&goto=2790&&a mp;srch=vestat#msg_2790)

Thanks,  
Kirill

> Pradeep

>  
>

> Pradeep Padala wrote:

>

> > Thanks for the explanation, but I DID use the --cpulimit parameter.

> >

> > I set the limit to 1000 units, I run a while(1) loop, and I see a  
> 100%

> > loadavg on the host node. Shouldn't it be somewhere around 10% ?

> >

> > Pradeep

> >

> > On 6/6/06, \*Kir Kolyshkin\* <kir@openvz.org

> <mailto:kir@openvz.org> <mailto:kir@openvz.org <mailto:kir@openvz.org>>

> > wrote:

> >

> > Looks like you misunderstand the concept of cpuunits.

> cpuunits is

> > not a

> > hard limit, but just a suggestion, and a CPU time is shared

> > proportionally to the values given. So, if you will have 9

> VEs and the

> > host system with cpuunits set to 1000 for all of them, and

> run the

> > loop

> > in all of them, each VE will use 10% of the CPU time.

> >

> > In case you will stop the loop running in 5 VEs so there will

> be 4

> > such

> > VEs (plus the host system) left, each of them will use 20% of  
> the CPU.

> > So, all the CPU time is distributed between VEs which will

> need it,

> > according with their proportional cpuunits.  
> >  
> > More to say, the concept of "total CPU units" is purely fiction,  
> > and is  
> > here just for the convenience. People do want to set CPU units in  
> > terms  
> > of processor's megahertz, and this is what cpuunits does. But in  
> > fact it  
> > is not a megahertz but just a relative weights. I.e. all the  
> cpuunits  
> > values are relative to each other, it doesn't matter what the  
> actual  
> > numbers are -- what matters is a number given to a VE in relation  
> > to the  
> > sum of all cpuweights (which is expressed as "total CPU units"  
> > just for  
> > the convenience).

> > So, cpuunits, if you do not oversell them, are a CPU  
> guarantee, not a  
> > limit. If you want CPU limit -- use cpulimit parameter.

> > Pradeep Padala wrote:

> >  
> > > Hi,  
> > >  
> > > I am trying to measure the CPU utilization of the VZ  
> > containers, and  
> > > change the cpu share dynamically. I have poured over most  
> of the  
> > > documentation, and looked at the code as well, and it seems  
> like  
> > > there's no utility that can directly show the current CPU  
> > utilization  
> > > of a container (some thing like 30% of CPU). A search on the  
> > user list  
> > > got me a message, where someone suggested using loadavg.  
> > However, It  
> > > seems like the loadavg is not showing the proper  
> utilization (or  
> > > showing the total CPU utilization). This is what I am doing.  
> > >  
> > > I setup a container with 1000 units limit (total CPU units:  
> > ~10000). I  
> > > wrote a small `do {; }while(1);` loop and ran it in the  
> container,  
> > now I  
> > > do `cat /proc/loadavg` in both the container and on the host

> node.  
> > > Since, the container is only using 1000 units, I should see  
> > something  
> > > like 100% loadavg in the container, and 10% loadavg in the  
> hostnode.  
> > > But, I see 100% at both places. Am I doing something wrong?  
> How do I  
> > > get the current cpu utilization of a container?  
> > >  
> > > Thanks,  
>  
>  
> -----  
>  
>

---