
Subject: how to get the current HN CPU utilization by particular CT

Posted by [knavnd](#) on Thu, 12 May 2011 14:19:34 GMT

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

Hi!

Sorry if that question has been already asked and answered before but I couldn't find the reply.

I wonder what the proper way is to determine the current HN CPU utilization by particular container? I need to know what percentage from total hardware node CPU power a certain CT consumes. It is mentioned in [1] the CPUUNITS and CPULIMIT parameters but as far as I understand they are used to set some HN CPU usage limits (mini and max

I wonder if the following algorithm is the proper one.

1) If CPULIMIT is not set then a formula would be as below:

(CPUUNITS assigned to particular CT) x (summary of user and system CPU consumption shown by top inside the CT) / (Power of the node" produced by vzcpucheck).

For example,

```
ct_cpuunits=`vzcpucheck -v|grep 139`
```

```
ct_user_cpu=`vzctl exec 139 top -bn1|grep -i "cpu(s)"|cut -d: -f2|awk  
{print $1}`|cut -d% -f1
```

```
ct_sys_cpu=`vzctl exec 139 top -bn1|grep -i "cpu(s)"|cut -d: -f2|awk  
{print $2}`|cut -d% -f1
```

```
total_hw_cpu=`vzcpucheck |tail -1|cut -d: -f2|tr -d '[:space:]`
```

```
ct_cpu_usage = $ct_cpuunits * ($ct_user_cpu+$ct_sys_cpu) / $total_hw_cpu
```

2) if CPULIMIT is set then

(CPULIMIT assigned to particular CT) x (summary of user and system CPU consumption shown by top inside the CT)

Or there is a simpler way?

Thanks!

Nikolay.

[1] http://wiki.openvz.org/Resource_shortage#CPU

Subject: vz migrate standby mode

Posted by [Jean-Marc Pigeon](#) on Thu, 12 May 2011 15:56:50 GMT

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

Hello,

> Hello,
>
>
> I think it will be very convenient to have a way to
> prepare/clone a VPS ready to start at moment notice
> (standby mode), a kind of action ready backup.
>
> I think the best way to do that is to modify
> vzmigrate to add "--standby", such the vzmigrate
> will work in --keep_dst mode and doing the --online
> mode but NOT stopping the VPS source and keeping the
> destination VPS in stop mode.
>
> Such starting vzmigrate --standby lets say
> once a day, could keep a fresh enough VPS
> ready to take action in case of major
> trouble on production HOST.
>
> Is there a better way to achieve this?,
> is vzmigrate the best candidate to do it?
> if so, I'll try to add "--standby" to
> vzmigrate.
>
>
>
> --
> A bientôt
> =====
=====
> Jean-Marc Pigeon Internet: jmp@safe.ca
> SAFE Inc. Phone: (514) 493-4280
> Fax: (514) 493-1946
> Clement, 'a kiss solution' to get rid of SPAM (at last)
> Clement' Home base <"http://www.clement.safe.ca">
> =====
=====
>

Subject: Re: vzmigrate standby mode
Posted by [Jean-Marc Pigeon](#) on Fri, 13 May 2011 00:27:09 GMT
[View Forum Message](#) <> [Reply to Message](#)

Hello,

On Thu, 2011-05-12 at 20:12 -0400, Sterling Windmill wrote:

> I believe you can already do this with vzctl chkpnt and vzctl restore
> as documented here:
Sure enough... and I'll be using it, but I want
a simple script to be able to be run via cron.
IMHO, vzmigrate is a good candidate to add
a "--standby" capability (trying to write something now).

>
>
> http://wiki.openvz.org/Checkpointing_and_live_migration
>
> On Thu, May 12, 2011 at 11:56 AM, Jean-Marc Pigeon <jmp@safe.ca>
> wrote:

> Hello,

>
>
> I think it will be very convenient to have a way to
> prepare/clone a VPS ready to start at moment notice
> (standby mode), a kind of action ready backup.

>
> I think the best way to do that is to modify
> vzmigrate to add "--standby", such the vzmigrate
> will work in --keep_dst mode and doing the --online
> mode but NOT stopping the VPS source and keeping the
> destination VPS in stop mode.

>
> Such starting vzmigrate --standby lets say
> once a day, could keep a fresh enough VPS
> ready to take action in case of major
> trouble on production HOST.

>
> Is there a better way to achieve this?,
> is vzmigrate the best candidate to do it?
> if so, I'll try to add "--standby" to
> vzmigrate.

> --

> A bientôt

> =====

=====

> Jean-Marc Pigeon Internet:
> jmp@safe.ca
> SAFE Inc. Phone:
> (514) 493-4280
> Fax: (514)
> 493-1946

> Clement, 'a kiss solution' to get rid of SPAM (at last)
> Clement' Home base <"http://www.clement.safe.ca">
> =====
=====

>
--
A bientôt

=====

=====

Jean-Marc Pigeon Internet: jmp@safe.ca
SAFE Inc. Phone: (514) 493-4280
Fax: (514) 493-1946
Clement, 'a kiss solution' to get rid of SPAM (at last)
Clement' Home base <"http://www.clement.safe.ca">

=====

=====

File Attachments

1) [smime.p7s](#), downloaded 487 times

Subject: Re: vzmigrate standby mode
Posted by [Aleksandar Ivanisevic](#) on Fri, 13 May 2011 09:54:44 GMT
[View Forum Message](#) <> [Reply to Message](#)

Jean-Marc Pigeon <jmp@safe.ca> writes:

> Hello,
>
>
> I think it will be very convenient to have a way to
> prepare/clone a VPS ready to start at moment notice
> (standby mode), a kind of action ready backup.

We do it with lsyncd <http://code.google.com/p/lsyncd/> that allows only changed files to be rsynced. Put /vz/private/veidt and /etc/vz/conf/veid.conf in it and you are all set.

Much easier on the machines, no cron necessary and the backups can be made really close to production.

[...]

Subject: Re: how to get the current HN CPU utilization by particular CT
Posted by [knewnd](#) on Fri, 13 May 2011 11:48:32 GMT
[View Forum Message](#) <> [Reply to Message](#)

knawnd@gmail.com wrote on 12/05/11 18:19:

> Hi!

>

> Sorry if that question has been already asked and answered before but

> I couldn't find the reply.

just have found two threads on the OpenVZ forum:

[1] <http://forum.openvz.org/index.php?t=msg&th=682>

[2] <http://forum.openvz.org/index.php?t=msg&th=479>

It seems that my first approach mentioned in my initial email in that thread is completely wrong. Trying to figure out the proper one. If someone can have already solved the same issue, please, share the tools or info with me.

Thanks.

Nikolay.

Subject: Re: how to get the current HN CPU utilization by particular CT
Posted by [Tim Small](#) on Fri, 13 May 2011 12:46:12 GMT

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

I can't remember how I derived it, but you might want to take a look at this munin plugin which I wrote a few months back:

<http://exchange.munin-monitoring.org/plugins/openvzcpu/details>

Cheers,

Tim.

--

South East Open Source Solutions Limited

Registered in England and Wales with company number 06134732.

Registered Office: 2 Powell Gardens, Redhill, Surrey, RH1 1TQ

VAT number: 900 6633 53 <http://seoss.co.uk/> +44-(0)1273-808309

Subject: Re: how to get the current HN CPU utilization by particular CT
Posted by [knawnd](#) on Fri, 13 May 2011 12:58:20 GMT

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

Tim, thanks a lot for reply and info!

Using all info I have now I will try to write what I need.

Thanks again!

Nikolay.

Tim Small wrote on 13/05/11 16:46:

> I can't remember how I derived it, but you might want to take a look at
> this munin plugin which I wrote a few months back:

>

> <http://exchange.munin-monitoring.org/plugins/openvzcpu/details>

>

> Cheers,

>

> Tim.

>

Subject: Re: Re: vzmigrate standby mode

Posted by [swindmill](#) on Fri, 13 May 2011 16:25:25 GMT

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

The problem with the lsyncd approach is it doesn't guarantee filesystem consistency. If you're running a database or any other application in the container that has open file handles to data that is constantly being manipulated, the data on the remote host isn't guaranteed to be consistent at any point in time. Obviously this can be partially worked around by creating application specific backups inside of the container using database dump tools, etc.

vzctl chkpnt and restore guarantees consistency in that the state of the container is dumped when the checkpoint is created and restored upon restoration of said checkpoint. As long as the dump file and unadulterated filesystem are available, restoration should result in a perfect copy of the container as of when it was checkpointed. That being said, checkpointing makes the container unavailable for a brief period of time and wouldn't be ideal as a means of taking backups of production systems on a regular basis.

Checkpointing for this purpose would probably only make sense for a container with data that doesn't change over time. A simple application server that relies on an external database for instance.

lsyncd is interesting though, I haven't heard of it before and will definitely be doing some research.

Best regards,
Sterling

On Fri, May 13, 2011 at 5:54 AM, Aleksandar Ivanisevic <aleksandar@ivanisevic.de> wrote:

> Jean-Marc Pigeon <jmp@safe.ca> writes:

>
> > Hello,
> >
> >
> > I think it will be very convenient to have a way to
> > prepare/clone a VPS ready to start at moment notice
> > (standby mode), a kind of action ready backup.
>
> We do it with lsyncd <http://code.google.com/p/lsyncd/> that allows only
> changed files to be rsynced. Put /vz/private/veidt and
> /etc/vz/conf/veid.conf in it and you are all set.
>
> Much easier on the machines, no cron necessary and the backups can be
> made really close to production.
>
> [...]
>
>

Subject: Re: Re: vzmigrate standby mode
Posted by [Jean-Marc Pigeon](#) on Fri, 13 May 2011 18:28:11 GMT
[View Forum Message](#) <> [Reply to Message](#)

On Fri, 2011-05-13 at 12:25 -0400, Sterling Windmill wrote:
> The problem with the lsyncd approach is it doesn't guarantee
> filesystem consistency. If you're running a database or any other
[...]
Agree with you, difficult to be an effective solution with
a real server running database, filesystem consistency
could be a real issue
> .
>
>
> vzctl chkpnt and restore guarantees consistency in that the state of
> the container is dumped when the checkpoint is created and restored
[...]
> period of time and wouldn't be ideal as a means of taking backups of
> production systems on a regular basis.
Using vzmigrate as starting point, I wrote a small
shell, with a small VPS (1.2 Gig Bytes) as test bench,
the preliminary result are (so fare) encouraging.
"standby" is done between 2 HOST on local network

Syncing private
Live migrating VPS...
Resuming local VPS

Times:

Second rsync: 0.507628 sec

Quota Generation: 0.00138903 sec

Suspended time: 0.631432 sec

Total time: 6.47325 sec

less than one second of suspend is good enough to me,
It is true the data-base within this test VPS
is not a very active one.

Keep in mind, there is many task (data dump transfer,
reload, ...) which can be done on the "standby"
VPS while the production one already resume task.

>

>

> Checkpointing for this purpose would probably only make sense for a
> container with data that doesn't change over time. A simple
> application server that relies on an external database for instance.

Agree with you.

Critical issue is the "suspend" time, if production
VPS is in very "volatile" data context, "suspend"
time could be a real issue.

>

>

> Lsyncd is interesting though, I haven't heard of it before and will
> definitely be doing some research.

Agreed

>

>

> Best regards,
> Sterling

>

>

> On Fri, May 13, 2011 at 5:54 AM, Aleksandar Ivanisevic

> <aleksandar@ivanisevic.de> wrote:

> Jean-Marc Pigeon <jmp@safe.ca> writes:

>

> > Hello,

> >

> >

> > I think it will be very convenient to have a way to
> > prepare/clone a VPS ready to start at moment notice
> > (standby mode), a kind of action ready backup.

>

>

> We do it with Lsyncd <http://code.google.com/p/lsyncd/> that

> allows only
> changed files to be rsynced. Put /vz/private/veidt and
> /etc/vz/conf/veid.conf in it and you are all set.
>
> Much easier on the machines, no cron necessary and the backups
> can be
> made really close to production.
>
> [...]

--
A bientôt

=====
=====
Jean-Marc Pigeon Internet: jmp@safe.ca
SAFE Inc. Phone: (514) 493-4280
Fax: (514) 493-1946
Clement, 'a kiss solution' to get rid of SPAM (at last)
Clement' Home base <"http://www.clement.safe.ca">
=====
=====

File Attachments

1) [smime.p7s](#), downloaded 460 times

Subject: Re: how to get the current HN CPU utilization by particular CT
Posted by [k Nawnd](#) on Sat, 14 May 2011 08:27:22 GMT
[View Forum Message](#) <> [Reply to Message](#)

As far as I understand in order to determine HN CPU load produced by certain CT it is enough to extract 'user', 'system' and 'uptime' values from /proc/vz/veidstat twicely. Then the formula is
$$ct_cpu_usage = 100\% * [(user_time_2 + system_time_2) - (user_time_1 + system_time_1)] / (uptime_2 - uptime_1)$$

I wonder if the formula above is correct and if it is then if it is valid for both cases: when cpulimit is set and unset? I guess it should be valid in both cases.

The second question is if there are any plans to implement HN CPU load produced by particular CT in some of vz* tools (e.g. in vzcpucheck providing some additional option and print one more column required info about HN CPU utilization by each container)?

Regards,

Nikolay.

knawnd@gmail.com wrote on 13/05/11 16:58:

> Tim, thanks a lot for reply and info!

>

> Using all info I have now I will try to write what I need.

>

> Thanks again!

> Nikolay.

>

> Tim Small wrote on 13/05/11 16:46:

>> I can't remember how I derived it, but you might want to take a look at

>> this munin plugin which I wrote a few months back:

>>

>> <http://exchange.munin-monitoring.org/plugins/openvzcpu/details>

>>

>> Cheers,

>>

>> Tim.

>>

Subject: Re: Re: vzmigrate standby mode

Posted by [Jean-Marc Pigeon](#) on Mon, 16 May 2011 12:58:47 GMT

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

Hello,

Standby mode is working fine, I defined 2 scripts "vzstandby" and "vzonair". I tested with one VPS with a postgresql database application, sampling data every minute and reporting via a web interface.

It is "funny" (but expected) to see last data sampling vanishing from web page, when stopping production VPS and "vzonair" backup VPS, beside the system appearing just dropped in time, everything seems to be working as usual.

The backup process suspend time was the most critical to me. Test show, while the whole backup process can be up to one minute long, the suspend time is less than a second, which is quite acceptable.

vzstandby generate a standby VPS without quota, as dumping quota seems to be "suspend time" expensive. IMHO, to have a backup VPS working without quota is not a real issue.

(quota can be restarted, when restarting the backup VPS at a more convenient time).

I can provide vzstandby and vzonair to list if some of us think it could be of interest. (sure there is space for improvements...).

--

A bientôt

=====

=====

Jean-Marc Pigeon
SAFE Inc.

Internet: jmp@safe.ca
Phone: (514) 493-4280

Fax: (514) 493-1946

Clement, 'a kiss solution' to get rid of SPAM (at last)
Clement' Home base <"http://www.clement.safe.ca">

=====

=====

File Attachments

1) [smime.p7s](#), downloaded 490 times

Subject: Re: how to get the current HN CPU utilization by particular CT
Posted by [Tim Small](#) on Mon, 16 May 2011 13:59:25 GMT

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

On 14/05/11 09:27, knawnd@gmail.com wrote:

> As far as I understand in order to determine HN CPU load produced by
> certain CT it is enough to extract 'user', 'system' and 'uptime'
> values from /proc/vz/vestat twicely. Then the formula is
> $ct_cpu_usage = 100\% * [(user_time_2 + system_time_2) - (user_time_1 +$
> $system_time_1)] / (uptime_2 - uptime_1)$

Can't really remember I'm afraid, but that sounds plausible, I think the figures collected from /proc are a total amount of CPU time used for each container, so to determine usage, you have to take two readings and divide by elapsed time. The figure you get is the total across all CPU cores.

>

> The second question is if there are any plans to implement HN CPU load
> produced by particular CT in some of vz* tools (e.g. in vzcpucheck
> providing some additional option and print one more column required
> info about HN CPU utilization by each container)?

Dunno. Email the developers and see if they'd be willing to accept a

patch in principle?

If not (or in any case) - I wonder if a "dstat" plugin might be a good solution?

Actually I just looked at the dstat man page, and there seems to be a --vz-cpu option already....

Tim.

--

South East Open Source Solutions Limited
Registered in England and Wales with company number 06134732.
Registered Office: 2 Powell Gardens, Redhill, Surrey, RH1 1TQ
VAT number: 900 6633 53 <http://seoss.co.uk/> +44-(0)1273-808309

Subject: Re: how to get the current HN CPU utilization by particular CT
Posted by [knawnd](#) on Tue, 17 May 2011 07:02:50 GMT

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

Tim Small wrote on 16/05/11 17:59:

> On 14/05/11 09:27, knawnd@gmail.com wrote:

>> As far as I understand in order to determine HN CPU load produced by
>> certain CT it is enough to extract 'user', 'system' and 'uptime'
>> values from /proc/vz/vestat twice. Then the formula is

>> $ct_cpu_usage = 100\% * [(user_time_2 + system_time_2) - (user_time_1 + system_time_1)] / (uptime_2 - uptime_1)$

> Can't really remember I'm afraid, but that sounds plausible, I think the
> figures collected from /proc are a total amount of CPU time used for
> each container, so to determine usage, you have to take two readings and
> divide by elapsed time. The figure you get is the total across all CPU
> cores.

that's exactly what I need for now.

>> The second question is if there are any plans to implement HN CPU load
>> produced by particular CT in some of vz* tools (e.g. in vzcpucheck
>> providing some additional option and print one more column required
>> info about HN CPU utilization by each container)?

> Dunno. Email the developers and see if they'd be willing to accept a
> patch in principle?

I guess (and hope) OpenVZ developers are also reading that list. So it would be interesting and useful to know their opinions on that.

> If not (or in any case) - I wonder if a "dstat" plugin might be a good
> solution?

>

> Actually I just looked at the dstat man page, and there seems to be a

> --vz-cpu option already....

Thanks for pointing that out! I guess dstat is also can be used. But the implementation of HN CPU utilization by each CT in vz-native tools is much more preferable way for me since there wouldn't be a need to install additional third-party software and it would be simpler to configure sudoers file.

Regards,
Nikolay.

>

> Tim.

>

Subject: Re: vzmigrate standby mode

Posted by [Aleksandar Ivanisevic](#) on Wed, 18 May 2011 10:54:42 GMT

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

Sterling Windmill <sterling@ampx.net> writes:

[...]

> vzctl chkpnt and restore guarantees consistency in that the state of the
> container is dumped when the checkpoint is created and restored upon
> restoration of said checkpoint. As long as the dump file and unadulterated
> filesystem are available, restoration should result in a perfect copy of the
> container as of when it was checkpointed. That being said, checkpointing
> makes the container unavailable for a brief period of time and wouldn't be
> ideal as a means of taking backups of production systems on a
> regular basis.

Yes, sorry, forgot to mention that. Unfortunately there is no easy solution for applications like databases that constantly change big files. You simply have to handle them separately, either by replication on a database level or by putting the disk in DRBD or SAN.

--

Ti si arogantan, prepotentan i peglaš vlastitu frustraciju. -- Ivan
Tišljar, hr.comp.os.linux

Subject: Re: Re: vzmigrate standby mode

Posted by [Tim Small](#) on Wed, 18 May 2011 11:30:51 GMT

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

On 18/05/11 11:54, Aleksandar Ivanisevic wrote:

> Sterling Windmill <sterling@ampx.net> writes:

>
> [...]
>
>
>> vzctl chkpnt and restore guarantees consistency in that the state of the
>> container is dumped when the checkpoint is created and restored upon
>> restoration of said checkpoint. As long as the dump file and unadulterated
>> filesystem are available, restoration should result in a perfect copy of the
>> container as of when it was checkpointed. That being said, checkpointing
>> makes the container unavailable for a brief period of time and wouldn't be
>> ideal as a means of taking backups of production systems on a
>> regular basis.
>>
> Yes, sorry, forgot to mention that. Unfortunately there is no easy
> solution for applications like databases that constantly change big
> files. You simply have to handle them separately, either by
> replication on a database level or by putting the disk in DRBD or SAN.
>

We use a hacked version of mylvmbbackup to backup an entire container. Each container lives on its own logical volume, and the process calls into the logical volume to ask the database (mysql in this case) to make its data-on-disk consistent. At this point, an LVM snapshot is taken, then mysql is told it can carry on writing to disk. The LVM snapshot is then fscked and mounted on a different mountpoint, once mounted, the contents are rsynced to the standby machine, and the lvm snapshot is removed.

I'm guessing that this could be combined with a vzctl chkpnt, but I haven't looked into that.

Tim.

--

South East Open Source Solutions Limited
Registered in England and Wales with company number 06134732.
Registered Office: 2 Powell Gardens, Redhill, Surrey, RH1 1TQ
VAT number: 900 6633 53 <http://seoss.co.uk/> +44-(0)1273-808309

Subject: Re: Re: vzmigrate standby mode
Posted by [Tim Small](#) on Wed, 18 May 2011 11:49:03 GMT
[View Forum Message](#) <> [Reply to Message](#)

On 18/05/11 12:30, Tim Small wrote:

>
> We use a hacked version of mylvmbbackup to backup an entire container.

> Each container lives on its own logical volume, and the process calls
> into the logical volume to ask the database (mysql in this case) to make
> its data-on-disk consistent. At this point, and LVM snapshot is taken,
> then mysql is told it can carry on writing to disk. The LVM snapshot is
> then fscked and mounted on a different mountpoint, once mounted, the
> contents are rsynced to the standby machine, and the lvm snapshot is
> removed.
>

I should add that this is done at a relatively quiet time on the server
(if it's a server with lots of writes), as it has the side effect of
turning a single write on the containers storage into two writes and a
read (plus associated disk seeks) whilst the LVM snapshot exists.

Whilst the lvm snapshot is being created (< 1 second), writes to the
database block (reads continue as normal, however), then continue once
mysql gets the UNLOCK TABLES command. I don't know if there is
equivalent functionality available in other databases such as PostgreSQL
etc.

Cheers,

Tim.

--

South East Open Source Solutions Limited
Registered in England and Wales with company number 06134732.
Registered Office: 2 Powell Gardens, Redhill, Surrey, RH1 1TQ
VAT number: 900 6633 53 <http://seoss.co.uk/> +44-(0)1273-808309
