Subject: What vzctl --features nfs:on ... do?
Posted by piavlo on Wed, 20 Aug 2008 10:55:37 GMT
View Forum Message <> Reply to Message

Hi,

I'd like to know exactly what does the vzctl --features nfs:on ... do? When i should be using it? If i have a VE with nfs client, while HN does not have any nfs client/server, is this nfs:on feature useful for me?

ps. does VE has nfs4 client support?

THanks

Subject: Re: What vzctl --features nfs:on ... do?
Posted by piavlo on Sat, 23 Aug 2008 21:42:36 GMT
View Forum Message <> Reply to Message

bump

Subject: Re: What vzctl --features nfs:on ... do?
Posted by maratrus on Mon, 25 Aug 2008 11:38:22 GMT
View Forum Message <> Reply to Message

Hi,

each process has a "struct ve_struct" associated with it (current->ve_task_info.exec_env), this struct contains different information about the VE for example it contains veid field. It also contains "__u64 features" field - it is responsible for the VE features.
When any VE starts the appropriate ve_struct is created.

So, when you set "nfs feature:on" vzctl put appropriate string inside VEID.conf file and when you start your VE this information goes to features field of the "struct ve_struct".

During the mount process inside VE nfs_get_sb function if called. It checks that VE_FEATURE_NFS must be set (if (!ve_is_super(ve) && !(get_exec_env()->features & VE_FEATURE_NFS))) otherwise it returns -ENODEV. So, to be able to mount nfs inside VE we have to set nfs feature for this VE.

Quote:
ps. does VE has nfs4 client support?

Unfortunately, not yet.

Subject: Re: What vzctl --features nfs:on ... do?
Posted by piavlo on Mon, 25 Aug 2008 14:54:46 GMT
View Forum Message <> Reply to Message

Quote:So, to be able to mount nfs inside VE we have to set nfs feature for this VE. Strange since i've never set this feature explicitly to ON, still i'm able to mount nfs shares from VE. Maybe it is ON by default anyway?

Subject: Re: What vzctl --features nfs:on ... do?
Posted by maratrus on Mon, 25 Aug 2008 15:02:58 GMT
View Forum Message <> Reply to Message

Hi,

try to increase the log level (/etc/vz/vz.conf let's set to 10)
and then start your VE.
vzctl must output something like this:

Quote:
Set features mask */*

then you can find out if nfs_feature is set.

Subject: Re: What vzctl --features nfs:on ... do?
Posted by piavlo on Mon, 25 Aug 2008 21:09:04 GMT
View Forum Message <> Reply to Message

I'm getting the following outputSet features mask 0000000000000000/0000000000000000 What does it mean?

Subject: Re: What vzctl --features nfs:on ... do?
Posted by maratrus on Tue, 26 Aug 2008 08:34:51 GMT
View Forum Message <> Reply to Message

Hi,

what kernel do you use, what vzctl version do you use?
Is this issue affected all VEs, i.e. the way to reproduce it - just create any VE?

Subject: Re: What vzctl --features nfs:on ... do?
Posted by piavlo on Mon, 01 Sep 2008 09:42:23 GMT
View Forum Message <> Reply to Message

I use several 2.6.18 kernels currently:
Kernel 2.6.18 028stab056.1
Kernel 2.6.18 028stab053.14
Kernel RHEL5 028stab057.2

And i haveSet features mask 0000000000000000/0000000000000000 for all VEs with all these kernels.

What actually does this zero mask mean?

Subject: Re: What vzctl --features nfs:on ... do?
Posted by maratrus on Mon, 01 Sep 2008 10:45:14 GMT
View Forum Message <> Reply to Message

Hi,

you didn't provide with vzctl version you use.

Quote:
What actually does this zero mask mean?

As you can see the mask is shown as two numbers separated by "/" symbol.
Let's name the first number as "feature_mask", the second one as "known_mask". What do these numbers mean?
The first number represent the features you want to be enabled/disabled or you don't bother about.
The features you want to be enabled are the features that are in configuration file for particular VE and marked as :on.
The features you want to be disabled are the features that are in configuration file for particular VE and marked as :off.
The features you don't bother about are the features that don't be mentioned in conf file.

So, there are 4 variants:
1. The feature is set in feature_mask number and set in known_mask number (1-1)
This means that we know about this feature and want it to be enabled. The kernel must enable this feature for particular VE.
2. (0-1)
This means that we know about this feature and want it to be disabled. The kernel must disable this feature for particular VE.
3. (0-0)
This means that feature is not mentioned in conf file and kernel must set this feature to default value.
4. (1-0)
doesn't make any sence. This situation shouldn't be in real life.

So, we have third situation. I looked through the kernel and haven't been able to find that NFS feature should be enabled as default value.

Does your situation is easily reproducible, i.e. the scenario
- vzctl creat ....
- vzctl start ....
- mount -t nfs ....
always works and you are able to mount nfs inside VE?

Subject: Re: What vzctl --features nfs:on ... do?
Posted by piavlo on Mon, 01 Sep 2008 13:31:39 GMT
View Forum Message <> Reply to Message

vzctl-3.0.22

Yes mount -t nfs... always works from inside of VE
while there is mask 0000000000000000/0000000000000000

I just did vzctl set 111 --features nfs:on --save
and restarted the VE and now the mask is 0000000000000002/0000000000000002

And then i also add vzctl set 111 --features sysfs:on --save the mask is
0000000000000003/0000000000000003

So AFAI understand sysfs is least significant bit and nfs is second significant bit. And mask is in decimal or hexadicimal format.

Subject: Re: What vzctl --features nfs:on ... do?
Posted by piavlo on Mon, 08 Sep 2008 09:24:04 GMT
View Forum Message <> Reply to Message

So could some one tell me if the 0000000000000000/0000000000000000
mask is ok or not, considering that both nfs and sysfs is available from VE with
0000000000000000/0000000000000000 mask?

Subject: Re: What vzctl --features nfs:on ... do?
Posted by maratrus on Wed, 10 Sep 2008 13:41:14 GMT
View Forum Message <> Reply to Message

Hi,

I'm terribly sorry for delay.
I'm not able to do anything concerning this problem right now.
I think I can compile test kernel in a week.

Subject: Re: What vzctl --features nfs:on ... do?
Posted by maratrus on Wed, 17 Sep 2008 12:17:26 GMT
View Forum Message <> Reply to Message

Hello,

no wonder that VE_FEATURE_SYSFS is available indide VE because default set of features includes it. But it's very strange from my point of view that NFS if available inside container.
Could you possibly try to compile the latest 028stab057.2 rhel5 based kernel with the following debug patch. It only shows VEs features during the environment creation and nfs mounting.
So, I'd like you to do the following:
- vzctl stop VE_ID
- vzctl start VE_ID
- mount -t nfs ip:dir dir (from inside the VE)
And then please show the dmesg output.
Thank you very much for the patient.

File Attachments

1) diff-nfs-features-debug, downloaded 298 times

Subject: Re: What vzctl --features nfs:on ... do?
Posted by piavlo on Thu, 18 Sep 2008 09:22:53 GMT
View Forum Message <> Reply to Message

After some testing I can say that it does work properly on 028stab057.2 rhel5 based kernel but not on
028stab053.14 vanilla kernel

Ve with FEATURES="nfs:off" can still mount nfs shares on vanilla kernel in spite of that the fetures mask reported in vzctl.log is correct 0000000000000000/0000000000000002
While on rhel5 kernel i get mount.nfs: No such device
trying to mount nfs share.

Subject: Re: What vzctl --features nfs:on ... do?
Posted by maratrus on Thu, 18 Sep 2008 11:12:35 GMT
View Forum Message <> Reply to Message

Hi,

here is the reason - I've looked through rhel based kernel.
No doubt it's a bug:
http://bugzilla.openvz.org/show_bug.cgi?id=1018
Thank you very much for your great patient.