Subject: /dev/fd/3
Posted by Steve Hodges on Mon, 13 Aug 2007 17:33:01 GMT
View Forum Message <> Reply to Message

A program in a VE is complaining that it can't open /dev/fd/3

From reading, it appears that this is the equivalent of asking to open something with a file descriptor of 3. So I guess that the devices option on vzctl isn't going to help me here.

I also find that whilst I have made /dev/hdc available to my VE, I can't mount the cdrom. It tells me

mount: unknown filesystem type 'udf'
or
mount: unknown filesystem type 'iso9660'

depending on what I enter in the /etc/fstab.

My host OS is Debian etch, by VE is using a Debian etch template.

This has had me going around in circles for some time :-(

Steve

Subject: Re: /dev/fd/3
Posted by Gregor Mosheh on Mon, 13 Aug 2007 17:58:36 GMT
View Forum Message <> Reply to Message

> A program in a VE is complaining that it can't open /dev/fd/3
> From reading, it appears that this is the equivalent of asking to open
> something with a file descriptor of 3. So I guess that the devices
> option on vzctl isn't going to help me here.

Indeed. It sounds like a problem within the program. Something's causing a file descriptor to not be opened, and the program isn't performing sufficient error-checking. Cuz I know one can open file descriptors at will!

Or perhaps your VE has maxed out the number of FDs it's allowed to have? Check the /proc/user_beancounters and see if anything there is happening to shed light on whether the failure is VE-related.

> I also find that whilst I have made /dev/hdc available to my VE, I can't
> mount the cdrom. It tells me
> mount: unknown filesystem type 'udf'
> mount: unknown filesystem type 'iso9660'

For a CDROM, iso9660 is usually appropriate. Are you sure that that
module is already loaded in the hardware node? A lot of distros leave
iso9660 as a module and VEs cannot load kernel modules.

Do a "modprobe iso9660" in the HN, to ensure that the module is loaded.
Then try mounting the CD in the VE.

--
Gregor Mosheh / Greg Allensworth
System Administrator, HostGIS cartographic development & hosting services
http://www.HostGIS.com/

"Remember that no one cares if you can back up,
only if you can restore." - AMANDA

Subject: Re: /dev/fd/3
Posted by Gregor Mosheh on Mon, 13 Aug 2007 18:03:26 GMT
View Forum Message <> Reply to Message

> Indeed. It sounds like a problem within the program. Something's causing
> a file descriptor to not be opened, and the program isn't performing
> sufficient error-checking. Cuz I know one can open file descriptors at
> will!

To elaborate a bit more... The program "strace" is invaluable for
finding out what a program is doing (or failing to do). If you strace
your program's execution (and the failure is happening early enough)
then you can probably find out which fopen() is failing and why.

Like I said, this may be due to a VE issue (too many FDs open already,
out of memory while reading file, etc) or may be something about the
program itself.

--
Gregor Mosheh / Greg Allensworth
System Administrator, HostGIS cartographic development & hosting services
http://www.HostGIS.com/

"Remember that no one cares if you can back up,
only if you can restore." - AMANDA

Subject: Re: /dev/fd/3
Posted by Steve Hodges on Tue, 14 Aug 2007 08:17:29 GMT

On 14/08/2007 1:58 AM, Gregor Mosheh wrote:
>
> Or perhaps your VE has maxed out the number of FDs it's allowed to
> have? Check the /proc/user_beancounters and see if anything there is
> happening to shed light on whether the failure is VE-related.

Here is what I see, including the beancounters immediately after

Executing 'mkisofs -C 106736,113120 -M /dev/fd/3 -R -J -pad
-graft-points /2007-08-14__12-07-50=/home/archive/image | builtin_dd
of=/dev/hdc obs=32k seek=7070'
genisoimage: No such file or directory. Cannot open '/dev/fd/3'. Cannot
open SCSI driver.
genisoimage: No such file or directory. Unable to open previous session
image '/dev/fd/3'.
:-(mkisofs has failed: 2
***image written
***error writing to DVD. You may need another DVD
root@backup:/# cat /proc/user_beancounters
Version: 2.5
       uid resource   held   maxheld   barrier   limit
failcnt
     111: kmemsize   509969   691152   2752512
2936012   0
           lockedpages   0   8905   30000
30000   0
           privvmpages   1778   18792   100000
200000   0
           shmpages   0   0   8192
8192   0
           dummy   0   0   0
0   0
           numproc   10   14   65
65   0
           physspages   1112   9981   0
2147483647   0
           vmguarpages   0   0   6144
2147483647   0
           oomguarpages   1112   9981   6144
2147483647   0
           numtcpsock   3   3   80
80   0
           numflock   1   2   100
110   0
           numpty   1   1   16
16   0
           numsiginfo   0   2   256

```
256         0
        tcpsndbuf           0         0     319488
524288          0
        tcprcvbuf           0         0     319488
524288          0
        othersockbuf       2220       2904     132096
336896          0
        dgramrcvbuf          0         0     132096
132096          0
        numothersock         2         4        120
120         0
        dcachesize           0         0    1048576
1097728         0
        numfile            237       317       2048
2048         0
        dummy                0         0         0
0         0
        dummy                0         0         0
0         0
        dummy                0         0         0
0         0
        numiptent           10        10        128
128         0
```

the lines flagged with *** come from my script.

the relevant lines of the script are:

```
# create a path on the dvd
writedate=$(date +%Y-%m-%d__%H-%M-%S)

# burn files to dvd (either new or append) in timestamped directory
# might be a good idea to log this fact too.
err=
case "$dvdstatus" in
  blank)
    growisofs -Z /dev/hdc -R -J -pad -graft-points
"/${writedate}=${image}" || err=Y
    ;;

  appendable)
    growisofs -M /dev/hdc -R -J -pad -graft-points
"/${writedate}=${image}" || err=Y
    ;;

  *)
    echo Internal error deciding how to write to dvd
    exit 1
```

```
    ;;

esac
echo image written

# if error then send email and ask for new dvd (reporting error too)
if [[ "$err" == "Y" ]]; then
  echo error writing to DVD.  You may need another DVD
  eject
  exit 1
fi
```

(clearly the script isn't complete yet :-)

>
>
>> I also find that whilst I have made /dev/hdc available to my VE, I
>> can't mount the cdrom. It tells me
>> mount: unknown filesystem type 'udf'
>> mount: unknown filesystem type 'iso9660'
>
> For a CDROM, iso9660 is usually appropriate. Are you sure that that
> module is already loaded in the hardware node? A lot of distros leave
> iso9660 as a module and VEs cannot load kernel modules.
>
> Do a "modprobe iso9660" in the HN, to ensure that the module is
> loaded. Then try mounting the CD in the VE.

modprobe iso9660 seems to work on the hardware node, however lsmod |
grep iso shown only
isofs   21348 0

(which my googling suggests may actually me the right thing to see.)

nevertheless, I can
mount -t iso9660 /dev/hdc /media/cdrom0
whereapon I get
mount: block device /dev/hdc is write-protected, mounting read-only
and df, in part, reports
/dev/hdc   458   458   0 100% /media/cdrom0
/etc/fstab on the HN contains:
/dev/hdc /media/cdrom0 udf,iso9660 user,noauto 0 0

I then umount /dev/hdc and go into the ve (which has access to hdc via
vzctl set $1 --devnodes hdc:rw --save)
it's /etc/fstab looks like this:
UNCONFIGURED FSTAB FOR BASE SYSTEM
/dev/hdc /media/cdrom0 udf,iso9660 user,noauto 0 0

if I do this:
mount -t iso9660 /dev/hdc /media/cdrom0
I get this:
mount: unknown filesystem type 'iso9660'

>
>

Subject: Re: /dev/fd/3
Posted by TheWiseOne on Tue, 14 Aug 2007 11:30:42 GMT
View Forum Message <> Reply to Message

The same bug exists in Virtuozzo, "ln -s /proc/self/fd /dev/fd" appears
to be the fix.

Reference: http://www.tektonic.net/forum/showthread.php?t=1936

Steve Hodges wrote:
> On 14/08/2007 1:58 AM, Gregor Mosheh wrote:
>>
>> Or perhaps your VE has maxed out the number of FDs it's allowed to
>> have? Check the /proc/user_beancounters and see if anything there is
>> happening to shed light on whether the failure is VE-related.
>
> Here is what I see, including the beancounters immediately after
>
> Executing 'mkisofs -C 106736,113120 -M /dev/fd/3 -R -J -pad
> -graft-points /2007-08-14__12-07-50=/home/archive/image | builtin_dd
> of=/dev/hdc obs=32k seek=7070'
> genisoimage: No such file or directory. Cannot open '/dev/fd/3'.
> Cannot open SCSI driver.
> genisoimage: No such file or directory. Unable to open previous
> session image '/dev/fd/3'.
> :-(mkisofs has failed: 2
> ***image written
> ***error writing to DVD. You may need another DVD
> root@backup:/# cat /proc/user_beancounters
> Version: 2.5
>   uid resource   held   maxheld   barrier   limit
> failcnt
>   111: kmemsize   509969   691152   2752512
> 2936012   0
>   lockedpages   0   8905   30000
> 30000   0
>   privvmpages   1778   18792   100000

> 200000 0
> shmpages 0 0 8192
> 8192 0
> dummy 0 0 0
> 0 0
> numproc 10 14 65
> 65 0
> physpages 1112 9981 0
> 2147483647 0
> vmguarpages 0 0 6144
> 2147483647 0
> oomguarpages 1112 9981 6144
> 2147483647 0
> numtcpsock 3 3 80
> 80 0
> numflock 1 2 100
> 110 0
> numpty 1 1 16
> 16 0
> numsiginfo 0 2 256
> 256 0
> tcpsndbuf 0 0 319488
> 524288 0
> tcprcvbuf 0 0 319488
> 524288 0
> othersockbuf 2220 2904 132096
> 336896 0
> dgramrcvbuf 0 0 132096
> 132096 0
> numothersock 2 4 120
> 120 0
> dcachesize 0 0 1048576
> 1097728 0
> numfile 237 317 2048
> 2048 0
> dummy 0 0 0
> 0 0
> dummy 0 0 0
> 0 0
> dummy 0 0 0
> 0 0
> numiptent 10 10 128
> 128 0
>
> the lines flagged with *** come from my script.
>
> the relevant lines of the script are:
>

> # create a path on the dvd
> writedate=$(date +%Y-%m-%d__%H-%M-%S)
>
> # burn files to dvd (either new or append) in timestamped directory
> # might be a good idea to log this fact too.
> err=
> case "$dvdstatus" in
> blank)
> growisofs -Z /dev/hdc -R -J -pad -graft-points
> "/${writedate}=${image}" || err=Y
> ;;
>
> appendable)
> growisofs -M /dev/hdc -R -J -pad -graft-points
> "/${writedate}=${image}" || err=Y
> ;;
>
> *)
> echo Internal error deciding how to write to dvd
> exit 1
> ;;
>
> esac
> echo image written
>
> # if error then send email and ask for new dvd (reporting error too)
> if [["$err" == "Y"]]; then
> echo error writing to DVD. You may need another DVD
> eject
> exit 1
> fi
>
> (clearly the script isn't complete yet :-)
>
>>
>>
>>> I also find that whilst I have made /dev/hdc available to my VE, I
>>> can't mount the cdrom. It tells me
>>> mount: unknown filesystem type 'udf'
>>> mount: unknown filesystem type 'iso9660'
>>
>> For a CDROM, iso9660 is usually appropriate. Are you sure that that
>> module is already loaded in the hardware node? A lot of distros leave
>> iso9660 as a module and VEs cannot load kernel modules.
>>
>> Do a "modprobe iso9660" in the HN, to ensure that the module is
>> loaded. Then try mounting the CD in the VE.
>

> modprobe iso9660 seems to work on the hardware node, however lsmod |
> grep iso shown only
> isofs 21348 0
>
> (which my googling suggests may actually me the right thing to see.)
>
> nevertheless, I can
> mount -t iso9660 /dev/hdc /media/cdrom0
> whereapon I get
> mount: block device /dev/hdc is write-protected, mounting read-only
> and df, in part, reports
> /dev/hdc 458 458 0 100% /media/cdrom0
> /etc/fstab on the HN contains:
> /dev/hdc /media/cdrom0 udf,iso9660 user,noauto 0 0
>
> I then umount /dev/hdc and go into the ve (which has access to hdc via
> vzctl set $1 --devnodes hdc:rw --save)
> it's /etc/fstab looks like this:
> # UNCONFIGURED FSTAB FOR BASE SYSTEM
> /dev/hdc /media/cdrom0 udf,iso9660 user,noauto 0 0
> if I do this:
> mount -t iso9660 /dev/hdc /media/cdrom0
> I get this:
> mount: unknown filesystem type 'iso9660'
>
>
>
>>
>>

Subject: Re: /dev/fd/3
Posted by Steve Hodges on Tue, 14 Aug 2007 14:40:44 GMT
View Forum Message <> Reply to Message

thank you

The really strange thing is that I'm damn sure I had this working at
least once before. I wonder if I managed to snag a different version of
the template when I rebuilt the server?

Anyway, that's fixed now :-)

It still won't mount the dvd, but that's not a huge issue at present.

Steve

On 14/08/2007 7:30 PM, Matt Ayres wrote:

> The same bug exists in Virtuozzo, "ln -s /proc/self/fd /dev/fd"
> appears to be the fix.
>
> Reference: http://www.tektonic.net/forum/showthread.php?t=1936
>
> Steve Hodges wrote:
>> On 14/08/2007 1:58 AM, Gregor Mosheh wrote:
>>>
>>> Or perhaps your VE has maxed out the number of FDs it's allowed to
>>> have? Check the /proc/user_beancounters and see if anything there is
>>> happening to shed light on whether the failure is VE-related.
>>
>> Here is what I see, including the beancounters immediately after
>>
>> Executing 'mkisofs -C 106736,113120 -M /dev/fd/3 -R -J -pad
>> -graft-points /2007-08-14__12-07-50=/home/archive/image | builtin_dd
>> of=/dev/hdc obs=32k seek=7070'
>> genisoimage: No such file or directory. Cannot open '/dev/fd/3'.
>> Cannot open SCSI driver.
>> genisoimage: No such file or directory. Unable to open previous
>> session image '/dev/fd/3'.
>> :-(mkisofs has failed: 2
>> ***image written
>> ***error writing to DVD. You may need another DVD
>> root@backup:/# cat /proc/user_beancounters
>> Version: 2.5
>>   uid resource   held   maxheld   barrier
>> limit   failcnt
>>   111: kmemsize   509969   691152   2752512
>> 2936012   0
>>   lockedpages   0   8905   30000
>> 30000   0
>>   privvmpages   1778   18792   100000
>> 200000   0
>>   shmpages   0   0   8192
>> 8192   0
>>   dummy   0   0   0
>> 0   0
>>   numproc   10   14   65
>> 65   0
>>   physpages   1112   9981   0
>> 2147483647   0
>>   vmguarpages   0   0   6144
>> 2147483647   0
>>   oomguarpages   1112   9981   6144
>> 2147483647   0
>>   numtcpsock   3   3   80
>> 80   0

>> numflock 1 2 100
>> 110 0
>> numpty 1 1 16
>> 16 0
>> numsiginfo 0 2 256
>> 256 0
>> tcpsndbuf 0 0 319488
>> 524288 0
>> tcprcvbuf 0 0 319488
>> 524288 0
>> othersockbuf 2220 2904 132096
>> 336896 0
>> dgramrcvbuf 0 0 132096
>> 132096 0
>> numothersock 2 4 120
>> 120 0
>> dcachesize 0 0 1048576
>> 1097728 0
>> numfile 237 317 2048
>> 2048 0
>> dummy 0 0 0
>> 0 0
>> dummy 0 0 0
>> 0 0
>> dummy 0 0 0
>> 0 0
>> numiptent 10 10 128
>> 128 0
>>
>> the lines flagged with *** come from my script.
>>
>> the relevant lines of the script are:
>>
>> # create a path on the dvd
>> writedate=$(date +%Y-%m-%d__%H-%M-%S)
>>
>> # burn files to dvd (either new or append) in timestamped directory
>> # might be a good idea to log this fact too.
>> err=
>> case "$dvdstatus" in
>> blank)
>> growisofs -Z /dev/hdc -R -J -pad -graft-points
>> "/${writedate}=${image}" || err=Y
>> ;;
>>
>> appendable)
>> growisofs -M /dev/hdc -R -J -pad -graft-points
>> "/${writedate}=${image}" || err=Y

>>   ;;
>>
>> *)
>>   echo Internal error deciding how to write to dvd
>>   exit 1
>>   ;;
>>
>> esac
>> echo image written
>>
>> # if error then send email and ask for new dvd (reporting error too)
>> if [["$err" == "Y"]]; then
>> echo error writing to DVD. You may need another DVD
>> eject
>> exit 1
>> fi
>>
>> (clearly the script isn't complete yet :-)
>>
>>>
>>>
>>>> I also find that whilst I have made /dev/hdc available to my VE, I
>>>> can't mount the cdrom. It tells me
>>>> mount: unknown filesystem type 'udf'
>>>> mount: unknown filesystem type 'iso9660'
>>>
>>> For a CDROM, iso9660 is usually appropriate. Are you sure that that
>>> module is already loaded in the hardware node? A lot of distros
>>> leave iso9660 as a module and VEs cannot load kernel modules.
>>>
>>> Do a "modprobe iso9660" in the HN, to ensure that the module is
>>> loaded. Then try mounting the CD in the VE.
>>
>> modprobe iso9660 seems to work on the hardware node, however lsmod |
>> grep iso shown only
>> isofs   21348 0
>>
>> (which my googling suggests may actually me the right thing to see.)
>>
>> nevertheless, I can
>> mount -t iso9660 /dev/hdc /media/cdrom0
>> whereapon I get
>> mount: block device /dev/hdc is write-protected, mounting read-only
>> and df, in part, reports
>> /dev/hdc   458   458   0 100% /media/cdrom0
>> /etc/fstab on the HN contains:
>> /dev/hdc /media/cdrom0 udf,iso9660 user,noauto 0 0
>>

>> I then umount /dev/hdc and go into the ve (which has access to hdc
>> via vzctl set $1 --devnodes hdc:rw --save)
>> it's /etc/fstab looks like this:
>> # UNCONFIGURED FSTAB FOR BASE SYSTEM
>> /dev/hdc /media/cdrom0 udf,iso9660 user,noauto 0 0
>> if I do this:
>> mount -t iso9660 /dev/hdc /media/cdrom0
>> I get this:
>> mount: unknown filesystem type 'iso9660'
>>
>>
>>
>>>
>>>