
Subject: Capabilities

Posted by [Benny Amorsen](#) on Wed, 14 Mar 2007 09:16:35 GMT

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

The kernel kernel-smp-2.6.18-ovz028test018.1 (and others before it) has CONFIG_SECURITY and therefore CONFIG_SECURITY_CAPABILITIES turned off. Unfortunately, quagga (zebra, actually) needs capabilities in at least Fedora Core.

Would it be possible to enable CONFIG_SECURITY and make CONFIG_SECURITY_CAPABILITIES a module in the standard OpenVZ kernel? Does OpenVZ fundamentally not work with capabilities?

/Benny

Subject: Re: Capabilities

Posted by [dev](#) on Wed, 14 Mar 2007 09:32:11 GMT

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

Benny,

Not sure what you mean, cause std Linux capabilities do work fine. Is there any clue on what is required by zebra? any messages? exact problem description?

Thanks,
Kirill

> The kernel kernel-smp-2.6.18-ovz028test018.1 (and others before it)
> has CONFIG_SECURITY and therefore CONFIG_SECURITY_CAPABILITIES turned
> off. Unfortunately, quagga (zebra, actually) needs capabilities in at least
> Fedora Core.
>
> Would it be possible to enable CONFIG_SECURITY and make
> CONFIG_SECURITY_CAPABILITIES a module in the standard OpenVZ kernel?
> Does OpenVZ fundamentally not work with capabilities?
>
>
> /Benny
>
>

Subject: Re: Capabilities

Posted by [Benny Amorsen](#) on Wed, 14 Mar 2007 09:35:27 GMT

>>>> "BA" == Benny Amorsen <benny+usenet@amorsen.dk> writes:

BA> Would it be possible to enable CONFIG_SECURITY and make
BA> CONFIG_SECURITY_CAPABILITIES a module in the standard OpenVZ
BA> kernel? Does OpenVZ fundamentally not work with capabilities?

Ah, a quick test gave me this result:

```
kernel/sched.c: In function '__activate_task':
kernel/sched.c:1565: warning: implicit declaration of function 've_stop_idle'
kernel/sched.c:1565: error: 've' undeclared (first use in this function)
kernel/sched.c:1565: error: (Each undeclared identifier is reported only once
kernel/sched.c:1565: error: for each function it appears in.)
kernel/sched.c: In function 'deactivate_task':
kernel/sched.c:1718: error: 'pcpu' undeclared (first use in this function)
kernel/sched.c:1735: warning: implicit declaration of function 've_strt_idle'
kernel/sched.c:1735: error: 've' undeclared (first use in this function)
kernel/sched.c:1735: error: 'cpu' undeclared (first use in this function)
kernel/sched.c: In function 'pull_task':
kernel/sched.c:3037: error: 've' undeclared (first use in this function)
kernel/sched.c: In function 'vsched_add_vcpu':
kernel/sched.c:8218: warning: implicit declaration of function 'VE_CPU_STATS'
kernel/sched.c:8218: error: 'struct fairsched_node' has no member named 'owner_env'
kernel/sched.c:8219: error: invalid application of 'sizeof' to incomplete type 'struct ve_cpu_stats'
kernel/sched.c:8219: warning: passing argument 1 of 'memset' makes pointer from integer without
a cast
kernel/sched.c:8221: error: 'struct fairsched_node' has no member named 'owner_env'
make[1]: *** [kernel/sched.o] Error 1
make: *** [kernel] Error 2
```

So I guess it isn't trivial to enable.

/Benny

Subject: Re: Capabilities

Posted by [Benny Amorsen](#) on Wed, 14 Mar 2007 09:40:49 GMT

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

```
[root@router01 /]# service zebra start
Starting zebra: Nothing to flush.
privs_init: initial cap_set_proc failed
```

[FAILED]

This is with kernel-smp-2.6.18-ovz028test018.1 and quagga-0.98.6-2.1

from Fedora Core 6.

```
$ fgrep CONFIG_SECURITY /lib/modules/2.6.18-ovz028test018.1-smp/build/.config
# CONFIG_SECURITY is not set
```

/Benny

Subject: Re: Re: Capabilities

Posted by [dev](#) on Wed, 14 Mar 2007 09:49:44 GMT

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

Benny Amorsen wrote:

>>>>>"BA" == Benny Amorsen <benny+usenet@amorsen.dk> writes:

>

>

> BA> Would it be possible to enable CONFIG_SECURITY and make
> BA> CONFIG_SECURITY_CAPABILITIES a module in the standard OpenVZ
> BA> kernel? Does OpenVZ fundamentally not work with capabilities?

>

> Ah, a quick test gave me this result:

>

> kernel/sched.c: In function '__activate_task':
> kernel/sched.c:1565: warning: implicit declaration of function 've_stop_idle'
> kernel/sched.c:1565: error: 've' undeclared (first use in this function)
> kernel/sched.c:1565: error: (Each undeclared identifier is reported only once
> kernel/sched.c:1565: error: for each function it appears in.)
> kernel/sched.c: In function 'deactivate_task':
> kernel/sched.c:1718: error: 'pcpu' undeclared (first use in this function)
> kernel/sched.c:1735: warning: implicit declaration of function 've_strt_idle'
> kernel/sched.c:1735: error: 've' undeclared (first use in this function)
> kernel/sched.c:1735: error: 'cpu' undeclared (first use in this function)
> kernel/sched.c: In function 'pull_task':
> kernel/sched.c:3037: error: 've' undeclared (first use in this function)
> kernel/sched.c: In function 'vsched_add_vcpu':
> kernel/sched.c:8218: warning: implicit declaration of function 'VE_CPU_STATS'
> kernel/sched.c:8218: error: 'struct fairsched_node' has no member named 'owner_env'
> kernel/sched.c:8219: error: invalid application of 'sizeof' to incomplete type 'struct ve_cpu_stats'
> kernel/sched.c:8219: warning: passing argument 1 of 'memset' makes pointer from integer
without a cast
> kernel/sched.c:8221: error: 'struct fairsched_node' has no member named 'owner_env'
> make[1]: *** [kernel/sched.o] Error 1
> make: *** [kernel] Error 2
>

> So I guess it isn't trivial to enable.

No, it looks like you used non-OVZ config and disabled OVZ options
like CONFIG_FAIRSCHEM and CONFIG_SCHED_VCPU. So compilation failed.

So this one is not related to zebra/capabilities.

Thanks,
Kirill

Subject: Re: Re: Capabilities
Posted by [dev](#) on Wed, 14 Mar 2007 09:57:41 GMT
[View Forum Message](#) <> [Reply to Message](#)

Benny,

from what I see after googling:

1. this is not related to SECURITY. it is related to std Linux capabilities.
2. Looks like quagga tries to setup more capabilities than it have,
i.e. extend it's capabilities. That's why it fails.
It is possible to patch zebra to find out why.
3. AFAIK it is possible to configure it not to do so.

Thanks,
Kirill

```
> [root@router01 /]# service zebra start
> Starting zebra: Nothing to flush.
> privs_init: initial cap_set_proc failed
>                                     [FAILED]
>
> This is with kernel-smp-2.6.18-ovz028test018.1 and quagga-0.98.6-2.1
> from Fedora Core 6.
>
> $ fgrep CONFIG_SECURITY /lib/modules/2.6.18-ovz028test018.1-smp/build/.config
> # CONFIG_SECURITY is not set
>
>
> /Benny
>
>
```

Subject: Re: Capabilities
Posted by [Benny Amorsen](#) on Wed, 14 Mar 2007 10:02:11 GMT
[View Forum Message](#) <> [Reply to Message](#)

>>>>> "KK" == Kirill Korotaev <dev@sw.ru> writes:

KK> No, it looks like you used non-OVZ config and disabled OVZ options

KK> like CONFIG_FAIRSCHEM and CONFIG_SCHED_VCPU. So compilation
KK> failed. So this one is not related to zebra/capabilities.

```
$ egrep '(CONFIG_FAIRSCHEM|CONFIG_SCHED_VCPU)'  
kernel-2.6.18-x86_64-smp.config.ovz.ba  
CONFIG_SCHED_VCPU=y  
CONFIG_FAIRSCHEM=y
```

```
$ diff -u SOURCES/kernel-2.6.18-x86_64-smp.config.ovz  
SOURCES/kernel-2.6.18-x86_64-smp.config.ovz.ba  
--- SOURCES/kernel-2.6.18-x86_64-smp.config.ovz 2007-02-27 13:27:56.000000000 +0100  
+++ SOURCES/kernel-2.6.18-x86_64-smp.config.ovz.ba 2007-03-14 10:19:24.000000000  
+0100  
@@ -1631,7 +1631,12 @@  
# Security options  
#  
# CONFIG_KEYS is not set  
-# CONFIG_SECURITY is not set  
+CONFIG_SECURITY=y  
+CONFIG_SECURITY_NETWORK=y  
+# CONFIG_SECURITY_NETWORK_XFRM is not set  
+CONFIG_SECURITY_CAPABILITIES=y  
+# CONFIG_SECURITY_ROOTPLUG is not set  
+# CONFIG_SECURITY_SECLVL is not set  
  
#  
# Cryptographic options
```

Here kernel-2.6.18-x86_64-smp.config.ovz is the file provided in the
OpenVZ source RPM, and kernel-2.6.18-x86_64-smp.config.ovz.ba is my
modified file. Compilation works fine with
kernel-2.6.18-x86_64-smp.config.ovz, but if I change the SPEC file to
use kernel-2.6.18-x86_64-smp.config.ovz.ba, compilation fails.

/Benny

Subject: Re: Re: Capabilities

Posted by [dev](#) on Wed, 14 Mar 2007 10:10:26 GMT

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

Benny Amorsen wrote:

>>>>>"KK" == Kirill Korotaev <dev@sw.ru> writes:

>

```

>
> KK> No, it looks like you used non-OVZ config and disabled OVZ options
> KK> like CONFIG_FAIRSCHEM and CONFIG_SCHED_VCPU. So compilation
> KK> failed. So this one is not related to zebra/capabilities.
>
> $ egrep '(CONFIG_FAIRSCHEM|CONFIG_SCHED_VCPU)'
kernel-2.6.18-x86_64-smp.config.ovz.ba
> CONFIG_SCHED_VCPU=y
> CONFIG_FAIRSCHEM=y
>
> $ diff -u SOURCES/kernel-2.6.18-x86_64-smp.config.ovz
SOURCES/kernel-2.6.18-x86_64-smp.config.ovz.ba
> --- SOURCES/kernel-2.6.18-x86_64-smp.config.ovz 2007-02-27 13:27:56.000000000 +0100
> +++ SOURCES/kernel-2.6.18-x86_64-smp.config.ovz.ba 2007-03-14 10:19:24.000000000
+0100
> @@ -1631,7 +1631,12 @@
> # Security options
> #
> # CONFIG_KEYS is not set
> -# CONFIG_SECURITY is not set
> +CONFIG_SECURITY=y
> +CONFIG_SECURITY_NETWORK=y
> +# CONFIG_SECURITY_NETWORK_XFRM is not set
> +CONFIG_SECURITY_CAPABILITIES=y
> +# CONFIG_SECURITY_ROOTPLUG is not set
> +# CONFIG_SECURITY_SECLVL is not set
Ahh... I forget that CONFIG_SECURITY automatically disables CONFIG_VE due to:
kernel/Kconfig.openvz:
config VE
    bool "Virtual Environment support"
    default y
    depends on !SECURITY <<<< this!
So if you remove this line it should be better.

```

Thanks,
Kirill

Subject: Re: Capabilities

Posted by [Benny Amorsen](#) on Wed, 14 Mar 2007 10:20:25 GMT

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

>>>>> "KK" == Kirill Korotaev <dev@sw.ru> writes:

```

KK> Ahh... I forget that CONFIG_SECURITY automatically disables
KK> CONFIG_VE due to: kernel/Kconfig.openvz: config VE bool "Virtual
KK> Environment support" default y depends on !SECURITY <<<< this! So
KK> if you remove this line it should be better.

```

It seems a bit dangerous to remove a depends line which must have been put there deliberately by OpenVZ developers. I guess I will have to try to patch Quagga instead.

/Benny

Subject: Re: Capabilities (solved, user stupidity)
Posted by [Benny Amorsen](#) on Wed, 14 Mar 2007 10:37:41 GMT
[View Forum Message](#) <> [Reply to Message](#)

The solution is to simply:

```
for CAP in net_admin net_raw sys_admin; do vzctl set 114 --capability ${CAP}:on --save ; done
```

It was <http://forum.openvz.org/index.php?t=msg&goto=4214&> which got me on the right track.

/Benny

Subject: Re: Re: Capabilities (solved, user stupidity)
Posted by [kir](#) on Wed, 14 Mar 2007 14:26:53 GMT
[View Forum Message](#) <> [Reply to Message](#)

Benny,

Would be cool if you'd add this piece of sacred knowledge :) to our wiki
-- <http://wiki.openvz.org/>
The title would be "Quagga in VE" or "Zebra in VE" or smth like that.

Benny Amorsen wrote:

```
> The solution is to simply:  
>  
> for CAP in net_admin net_raw sys_admin; do vzctl set 114 --capability ${CAP}:on --save ; done  
>  
> It was http://forum.openvz.org/index.php?t=msg&goto=4214& which got me  
> on the right track.  
>
```
