Subject: nscd on VE
Posted by edel on Sat, 19 Jan 2008 03:17:17 GMT
View Forum Message <> Reply to Message

hello great people, in this great society

i have problem running nscd on top of cos ve. from the logfile, it says "30624: Failed to drop capabilities", so i thought i have to enable some caps to run nscd.

here's the relevant info from strace:

Quote:...
chmod("/var/run/nscd/socket", 0666) = 0
listen(12, 128) = 0
getuid32() = 0
prctl(0x8, 0x1, 0x8c42050, 0xa84f63, 0x8c43258) = 0
capset(0x19980330, 0, {CAP_SETGID|CAP_SETUID|0x20000000,
CAP_SETGID|CAP_SETUID|0x20000000, 0}) = -1 EPERM (Operation not permitted)
write(3, "30490: Failed to drop capabiliti"..., 35) = 35
write(2, "nscd: ", 6) = 6
write(2, "cap_set_proc failed", 19) = 19
write(2, "\n", 1) = 1
exit_group(1) = ?

i did try:
vzctl set 1502 --capability setuid:on --save
vzctl set 1502 --capability setgid:on --save
...and reboot.

but it does not succed. i also tried enabling all but setpcap but to no avail still nscd failed to start.

thanks in advance.

--edel

Subject: Re: nscd on VE
Posted by edel on Fri, 25 Jan 2008 05:16:15 GMT
View Forum Message <> Reply to Message

hi,

anyone can help me here? thank you.

--edel

Subject: Re: nscd on VE
Posted by maratrus on Sat, 02 Feb 2008 13:25:19 GMT
View Forum Message <> Reply to Message

Hi,

I'm very sorry for delay.

As we can see from strace output setuid and setgid capabilities are not sufficient. We need also CAP_AUDIT_WRITE.

From src/include/linux/capability.h (openvz kernel tree) inside VE we should to set SETVEID capability and "man vzctl" allows us to do this. I tried this but suddenly found that this capability was not actually set. The reason is in kernel's do_env_create() function which is called when we start our VPS. It displaces setveid capability otherwise this VPS becomes the serious hole in security.

So I can suggest you two ways:
1. If it is possible you can use the elder version of nscd. I noticed that it doesn't try to set this capabilities. (I experimented with VPS based on Fedora 4).
2. You can compile nscd by yourself. We easily can find the relevant piece of code and can try to comment unnecessary string.

Subject: Re: nscd on VE
Posted by edel on Sat, 02 Feb 2008 23:59:35 GMT
View Forum Message <> Reply to Message

hello maratrus,

thank you very much for the reply.

i was already thnking of recompiling nscd before, however i couldnt find the source for centos5. and honestly, i didnt tried hard enough. since it's a nss, im thinking its part of glibc. not sure. maybe, when i have the time, i'll look into it.

i also tried setting setting up nsscache[1], but didnt get it to work.

for now, here's my workaround (for the benefit of others):

1. on HN, setup nss_ldap and nscd as usual
2. HN's /var/run/nscd is on tmpfs (1mb)
3. HN's /var/db/nscd is on 10mb loopback-mounted image
4. mount (--bind) HN's /var/{run,db}/nscd to /vz/root/${VEID}/var/{run,db}/nscd (via $VEID.mount)

reason for #2/3 is that i dont want to give HN's /var partition open/fully accessible to VEs.

and it appears working.

i got this idea from the mysql/shared hosting form the openvz wiki.

openvz is great!

thanks.

--edel

[1] http://code.google.com/p/nsscache/