
Subject: Re: Side effects of enabling CAP_SYS_TIME inside VE

Posted by [vaverin](#) on Thu, 20 Sep 2007 13:07:22 GMT

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

Alex,

some action inside kernel (for example like setting system time) checks process permissions. However kernel do not checks process or user names, it checks process's capability for this purpose.

(you can read man 7 capabilities for more details)

Ususally capabilites are inherited from parent to children but parent process is able to restrict the children's permissions too.

By default cap_sys_time capability is not allowed inside VE, therefore even VE root cannot change system time.

As far as I understand ntpd tries to set cap_sys_time capability for its working thread and fails because it does not have this capability.

When cap_sys_time is allowed inside VE this operation finished with success, unprivileged working thread gets the capability allowing it to change the system time.

thank you,
Vasily Averin
