

Cellular Network Infrastructure - Feature #4107

Start systemd services as non-root user

07/15/2019 06:56 AM - osmith

Status: New	Start date: 07/15/2019
Priority: High	Due date:
Assignee:	% Done: 0%
Category:	
Target version:	
Spec Reference:	
Description laforge wrote in OS#3369 : Ideally, as far as possible, we should start them as non-root user (which may require changes to our systemd service files, etc. in the individual git repos - but that is fine!). Starting them as non-root will also means that any writes to unintended directories like '/' will be discovered as they then would make the program start fail.	
Related issues:	
Related to Cellular Network Infrastructure - Bug #3369: no automatic testing ...	Resolved 06/29/2018
Related to Cellular Network Infrastructure - Bug #4821: Update working dir in...	New 10/20/2020

History

#1 - 07/15/2019 06:56 AM - osmith

- Related to Bug #3369: no automatic testing of Debian/Ubuntu packages added

#2 - 12/01/2019 09:38 AM - laforge

- Priority changed from Normal to Low

#3 - 10/01/2020 02:42 PM - laforge

Programs like osmo-msc, osmo-sgsn, osmo-cbc, osmo-smlc, osmo-hlr have no real time requirements or special needs in terms of raw networks sockets or tun devices. All of those should be executed as normal, non-privileged user from the start. This could be done via the systemd unit files. This could be done via the systemd unit files, or explicitly inside the osmocom programs via a privilege dropping approach.

the only processes that need special privileges are (AFAICT):

- osmo-gbproxy requires CAP_NET_RAW if IPPROTO_GTP sockets are required for FR/GRE/IP
- osmo-trx, osmo-bts, osmo-pcu requires CAP_SYS_NICE if SCHED_RR is to be used per command line argument (and is not done by e.g. systemd before starting it)
- osmo-ggsn requires CAP_NET_ADMIN for setting up the gtp0/tun0 devices (unless this is done externally before starting it)
- any program requires CAP_SYS_NICE if it uses the relatively new libosmocom/src/vty/cpu_sched_vty.c code to have user-configured scheduling

For those above, we basically have three possible strategies:

- at least drop all privileges except those we really ever need in the specific program (CAP_NET_RAW / CAP_NET_ADMIN / CAP_SYS_NICE). We can first constrain the permitted capabilities using cap_set_flag, then use prctl(PR_SET_KEEPCAPS, 1L) to keep capabilities while changing from root to non-root, and then change the user ID / group ID. <https://stackoverflow.com/a/13186076> has a nice example
- if it is sufficient to perform those privileged operations once on start-up, we could even drop those capabilities after performing the operations like creating netdev, binding socket, changing scheduler policy. This would mean that no subsequent changes can be made later on.

#4 - 10/01/2020 02:42 PM - laforge

- Assignee deleted (osmith)

- Priority changed from Low to High

#5 - 10/20/2020 03:09 PM - keith

- Related to Bug #4821: Update working dir in systemd unit files added