Terminating in the name of eBPF

Shutting down Podman containers offending against Seccomp

Tuebix 2023

Lightning Talk by Cedric Casper

What's the goal?

Taking Seccomp security a bit further

- Podman containers can be started with Seccomp profiles
  - Nice way to reduce attack surface
- When a process inside the container offends against the Seccomp profile...
  - ...not only the process should be terminated...
  - ***...but also the whole container!***

Why?
Taking Seccomp security a bit further

- Currently only processes inside the container are blocked by Seccomp
  - Parent process is still running…
  - …and so, the container is still running.

- Malware or attacker inside the container could still try out other attack methods, allowed System Calls, overseen System Calls, …

- We wanted the container and therefore potential threats to be shut down for good.

One way of doing it
Using eBPF to spot Seccomp offences

- Writing a program, using Python and BCC [1]
- Tracing exited processes and filtering/checking:
  - How did it exit? (Checking for Signal 31 and abnormal exit)
  - Did it use Seccomp? (exited_process->seccomp.mode == 2 or 3)

- Getting info on the container:
  - Container owner UID = owner of the Seccomp violating process' PID namespace
  - PID namespace ID as container identifier

- Shutting down the container:
  - Sudo-ing into the user, listing Podman containers filtered by previously determined PID namespace ID
  - -> podman kill <container ID>

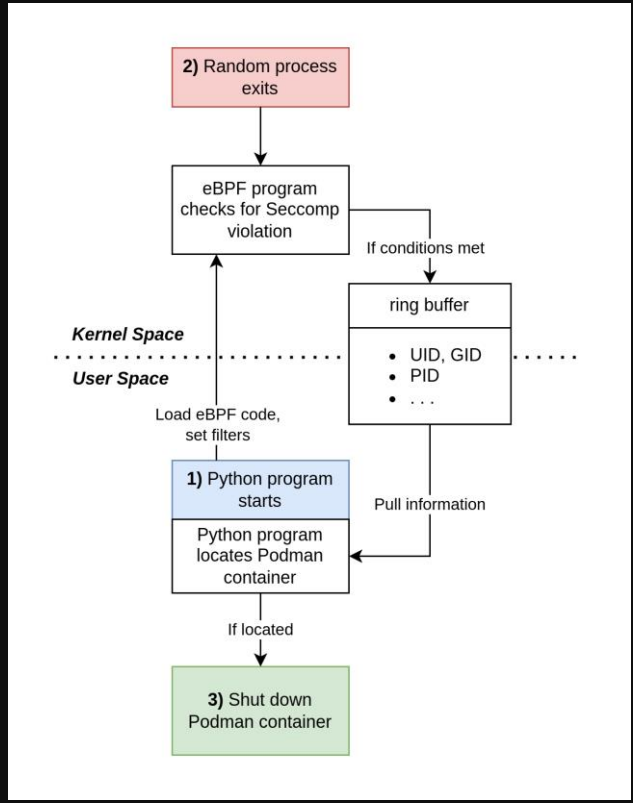

Program flow

[1] https://github.com/iovisor/bcc (Great source of inspiration for eBPF programs)

Result

- Bonus filters and infos:
  - Only shutdown container(s) of certain users or groups
  - Only shutdown certain container(s)
  - Get command that violated the Seccomp profile -> adjusting/debugging the Seccomp profile
  - And probably many more...

- After the Seccomp offence inside a container, that container is shut down
- -> Threat is reduced, the incident can be investigated

Example of a spotted Seccomp offence, followed by container shutdown

Code available on Github:
https://github.com/hashkeks/seccomSurv

Thank you!

For more information feel free to write me:

info@cedriccasper.com