Linux kernel contributions by ANSSI

Yves-Alexis Perez
ANSSI
LSS-EU 2018

Introduction

Who am I?

Yves-Alexis Perez

ANSSI head of software and hardware architectures lab

- OS security (Linux, micro-kernels)
- mobile and embedded device security
- CLIP OS[a] development

Debian developer

- security team
- Linux kernel packaging
- Xfce, strongSwan, imobiledevice packaging

Mostly interested in low-level security and hardening

a. `https://www.clip-os.org`

ANSSI

Missions

- not an intelligence agency
- protect French administration & industry networks
- defensive only

Labs

- pool of expertise on relevant area: cryptography, network, OS, ...
- research & development
- academic publications
- free-software contributions

French government & free software

User

- Large user of free software
- Linux distributions widely used accross the whole government
- cost reason, sometime about control

Contributor

- some large projects (SPIP, ...)
- DINSIC maintains a repository list[8]
- 2016 Digital Republic law[11] introduced open-data by default
- including code

Current situation with Linux

Linux usage quite diverse

- Linux servers from various distribution
- Linux-based appliances (firewalls, IDS, VPN...)
- some Linux workstations

ANSSI can't secure everything downstream

- upstream work benefits more people
- Linux distributions and Linux kernel
- contribute where possible

Past and current contributions

Past and current contributions / Documentation

Documentation

Administrators and products developers

- Linux distribution hardening guide[5]
  French
- Linux kernel hardening guide
  French (to be released)

Integrators and products developers

- CLIP OS 4 security architecture[16]
  French
- CLIP OS 5 documentation[4]
  English

Past and current contributions / CLIP OS

CLIP OS 4

ANSSI internal operating system

- Gentoo-based Linux distribution
- developed in-house since 2005
- targeted at the French government
- hardened and multi-level security
- includes specific kernel hardening
- recently released as free software[2]

CLIP OS 4 feedback

Hardened system in production since 10+ years

User

Not that much complains

- *"Where is MS Office?"*
- even on invasive security measures
- technical users miss scripting complete $W \oplus X$ policy

Developer

Complex beast

- distribution maintenance
- diversion from upstream
- complex toolchain and SDK
- isolated network

CLIP OS 5[1]

Same basis as CLIP OS 4

- suitable for managed networks
- hardened Linux distribution
- non-technical users (office work)
- multi-level

Different choices

- open development from the start
- upstream relevant code directly
- share with the community
- modular design to facilitate forks

State of the art implementation of ANSSI recommendations

CLIP OS 5 kernel[3]

Objectives

- provide isolation primitives to userspace
- maintain trust in hardware resources
- guarantee kernel self protection

Mainline hardening

- KSPP recommendations[10]
- minimize system (attack surface)
- security rather than performance

Security by default for all choices

Kernel patches

Some out of tree patches

- linux-hardened[7]
- lockdown[9]
- stackleak[12]

Testbed for hardening features

- provide some real-life feedback
- hopefully help inclusion into mainline
- minimize differences with mainline

Past and current contributions / Landlock

Landlock[15]

Status

- development by Mickaël Salaün from ANSSI
- submission process ongoing
- last presentation at LSS-NA 2018[14]

Specs

- unprivileged sandboxing
- define security policy directly from the application
- similar to seccomp-bpf but not limited to syscalls
- generic kernel objects access control
- implemented as a LSM

Planned involvement

Planned involvement / CLIP OS 4 upstreaming

O_MAYEXEC (from clip-patches[6])

- new flag for open(2)
- enforce and extend W⊕X policy to scripts
- currently requires patching interpreters

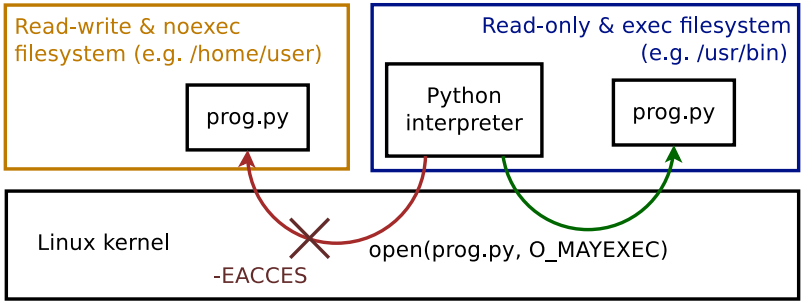

O_MAYEXEC upstreaming

Some re-architecturing needed

Options

- replace the hardcoded policy with a system-wide runtime configuration (sysctl)
- enable more fine-grained policies by making this property available to LSMs

Timeline

- first internal drafts ongoing
- first RFC by the end of the year
- integration into CLIP OS 5 in parallel

VServer[13]

Description

- large external patch (700kB)
- provide unified containers infrastructure
- upstream has no Linux upstreaming plans
- some bits could still be interesting

VServer as LSM

- VServer has the concept of XID (container ID)
- can be used for access control decisions
- VServer as LSM might be interesting
- plans and timeline are still unsure

Conclusion

Community feedback

On current work

- is it helpful?
- is it enough?

What can we do next?

- specific items to work on
- specific tasks

Questions

?

References I

ANSSI.
The CLIP OS project.
URL: https://clip-os.org/en/.

ANSSI.
CLIP OS 4.
URL: https://github.com/clipos-archive.

ANSSI.
CLIP OS kernel documentation.
URL: https://docs.clip-os.org/clipos/kernel.html.

ANSSI.
CLIP OS project documentation.
URL: https://docs.clip-os.org/.

ANSSI.
Recommandations de configuration d'un système GNU/Linux.
URL: https://www.ssi.gouv.fr/uploads/2015/10/NP_Linux_Configuration.pdf.

References II

ANSSI.
Set of Linux patches for CLIP OS.
URL: https://github.com/clipos-archive/src_platform_clip-patches.

Levente 'anthraxx' Polyak.
linux-hardened.
URL: https://github.com/anthraxx/linux-hardened.

Etalab.
Inventaire des codes sources des organismes publics.
URL: https://github.com/etalab/inventaire-codes-sources-organismes-publics.

Justin Forbes.
lockdown.
URL: https://git.kernel.org/pub/scm/linux/kernel/git/jforbes/linux.git/log/?h=lockdown.

References III

KSPP.
KSPP recommended settings.
URL: https://kernsec.org/wiki/index.php/Kernel_Self_Protection_Project/Recommended_Settings.

Axelle Lemaire.
Digital republic law.
URL: https://www.republique-numerique.fr/pages/in-english.

Alexander Popov.
stackleak.
URL: https://git.kernel.org/pub/scm/linux/kernel/git/kees/linux.git/log/?h=for-next/gcc-plugin/stackleak.

VServer project.
VServer project .
URL: http://linux-vserver.org/.

References IV

Mickaël Salaün.
How to safely restrict access to files in a programmatic way with landlock?
URL: https://landlock.io/talks/2018-08-27_landlock-lss.pdf.

Mickaël Salaün.
Landlock: programmatic access control.
URL: https://landlock.io/.

Vincent Strubel.
Documentation CLIP - Architecture de sécurité.
URL: https://github.com/clipos-archive/clipos4_doc/blob/master/developpeur/1002_Architecture_Securite_1.2.pdf.