Linux Kernel Security Update

LinuxCon Japan
Tokyo, 2016

James Morris
james.l.morris@oracle.com

DRAFT VERSION

Introduction

Who am I?

- Kernel security subsystem maintainer
- Mainline Linux kernel development @ Oracle

Introduction

Scope of Talk

- Linux kernel security background
- Overview of Linux kernel security subsystem
- Developments since ~2013
- Current and future challenges

Linux Kernel Security Overview: DAC

- Core security model is Discretionary Access Control (DAC)
- Inherited from Unix, designed in late 1960s
- Insufficient for modern security threats

Linux Kernel Security Overview: Extensions

- Posix ACLs
- Capabilities (privileges)
- Namespaces
- Seccomp
- Audit

Linux Kernel Security Overview: Extensions (cont'd)

- Netfilter
  - Iptables
- Cryptography API
  - IPSec
  - Disk encryption
  - Key management
- Mandatory Access Control (MAC)
  - SELinux
  - Smack
  - Apparmor

Linux Kernel Security Overview: Extensions (cont'd)

- Integrity Management
  - IMA, EVM
- Kernel self protection (KSP)
- Platform Security
  - Hardware features
    - TPM, SGX etc.

“Recent” Changes

- Scope: since May 2013
  - Kernel releases v3.10 to v3.19
  - and v4.0 (April 2015) to v4.6 (current)

Updates: LSM API

- kernel_fw_from_file hook (v3.17)
- Generalized security module stacking (v4.2)
  - Simple manual stacking previously allowed
  - Now: any number of smaller LSMs can be stacked on top of a major (“monolithic”) LSM
  - e.g. SELinux + YAMA + Capabilities, but not SELinux + TOMOYO + AppArmor.

Updates: Capabilities

- Ambient capabilities (v4.3)

Updates: SELinux

- Labeled NFS, and SELinux support (v3.11)
- Android Binder IPC support (v4.0)
- Full Netlink coverage (v4.1)
- Performance improvements (v4.1)
- Fine grained ioctl coverage (v4.3)
- Export validatetrans decisions to userspace (v4.6)

Updates: Smack

- Support modification of existing rules (v3.10)
- Local IPv6 network labeling (v3.11), IPv6 host labeling (v4.3)
- Bring-up (“permissive”) access mode (v3.18)
  - Allows logging only of policy violations
- Netfilter support (v4.0)
- Multiple label MAC bypass via onlycap (v4.2)

Updates: AppArmor

- Add interfaces to report profile and namespace information (v3.12)
- SHA1 hash reporting of loaded profiles (v3.12)
- Profile introspection (v3.12)
- Allow any profile to be set to unconfined (v3.12)

Updates: Integrity Subsystem

- IMA support for x.509 signed policy (v3.19)
- Custom IMA templates via kernel command line (v3.19)
- Integration of TPM 2.0 authorization policies with kernel keys, allow hash algorithm selection (v4.5)
- EVM support for x.509 kernel certificates (v4.5)
- Measurement & appraisal of IMA policy (v4.6)
- Support for kernexec image & initramfs (v4.6)

Updates: Platform Security

- Intel Memory Protection Extensions (MPX) (v3.19)
- Intel Trusted Execution Engine (TXE) (v3.15)
- TPM 2.0 chip support (v4.0)
- Intel Memory Protection Keys (MPK) (v4.5)

Updates: Audit

- Add support for auditing by executable file, rather than just PID (v4.3)
- Add ioctl device and command info to LSM audit data (v4.3)

Updates: Kernel Self Protection

- Kernel Address Sanitizer (KASan) (v4.0)
- Always enable RODATA checking (v4.6)
- KASLR
  - x86 (v3.14)
  - arm64 (v4.6)
- Page zero-poisoning (v4.6)

- See also: http://kernsec.org/wiki/index.php/Feature_List

Updates: Seccomp

- Integrate with BPF JIT (v3.16)
- Set filter across all threads (v3.17)
- seccomp() syscall (v3.17)
- SECCOMP_FILTER_FLAG_TSYNC (v3.17)
- ptrace options for suspend/resume (v4.3)
- ppc support (v4.3)
- Dump seccomp filters via ptrace (v4.4)

Updates: Crypto API Users

- ext4 filesystem encryption (v4.1)
- Kernel module signing (v4.3)
- MACsec/IEEE 802.1AE (v4.6)

Current / Evolving Work

- Kernel Self Protection Project
- Platform security
  - ARM
  - TBA

Future Work

Resources

- Linux Security Module mailing list
  - http://vger.kernel.org/vger-lists.html#linux-security-module
- Linux Security Summit (Aug 2016, Toronto)
  - http://events.linuxfoundation.org/events/linux-security-summit
- Kernel Self Protection Project
  - http://kernsec.org/wiki/index.php/Kernel_Self_Protection_Project
- LWN Security
  - http://lwn.net/Security