

State of the Linux Kernel Security Subsystem

James Morris <james.l.morris@oracle.com>

LinuxCon Japan 2012, Yokohama

Introduction

- Who I am
 - Kernel security maintainer
 - Engineering manager
- Scope
 - Background
 - Discuss Linux-specific security
 - Ongoing developments

SECURITY

THE DISNEY STORE

STOP!
NO CLEARANCE
NO ENTRY!

STAR



Background

- Linux is a clone of Unix
- Inherits core security model
- DAC
 - Not sufficient for modern systems
 - Malware, bugs etc.
 - User manages own object security
 - Root user overrides security
 - Does not protect against many threats
- Linux kernel has many security extensions...

Linux Kernel Security Features

- Need to be retrofitted to existing design!
 - Constrained by that design
- Extensions of DAC
 - Access Control Lists (ACLs)
 - Posix Capabilities (privileges)
 - Process-based
 - File capabilities

Linux Kernel Security Features

- Namespaces
- **Seccomp** (“mode 2” coming in 3.5)
- Netfilter/IPtables
- Cryptographic subsystem
 - Ipsec
 - Disk encryption
 - dm-crypt
 - ecryptfs

Linux Kernel Security Features

- Mandatory Access Control (MAC)
 - SELinux
 - Smack
 - AppArmor
 - TOMOYO

Linux Kernel Security Features

- System Hardening
 - ASLR
 - NX
 - /dev/mem restrictions
 - Toolchain hardening
 - Yama LSM (3.4)
 - ptrace_scope (grsec)

Linux Kernel Security Features

- Audit
- Keys
- Integrity & platform security
 - IMA/EVM
 - TPM
 - TXT
 - VT-d
 - dm-verity

Integrity Management Architecture (IMA)

- Detects if files have been maliciously or accidentally altered
- Measures and stores file hashes in TPM
 - Remote attestation
 - Local validation
 - IMA appraisal (ongoing)
- Protect security attributes against offline attack (EVM)

Seccomp Mode 2

- General system call filtering
- Reduces attack surface of kernel
- Not a sandbox!
- BPF filters installed with
 - `prctl(PR_SET_SECCOMP, SECCOMP_MODE_FILTER, prog);`
- Action may be set to trap, kill, errno, trace, allow.
- Also, `PR_SET_NO_NEW_PRIVS`
 - Prevents privilege granting via `execve()`

Ongoing Work

- Security requirements also now being driven by mobile and virt
 - SE-Android
 - Tizen (Smack)
 - Svirt
- Integrity management a focus of current work
 - Signed modules
 - Trusted boot etc.
- Needs work:
 - Usability
 - System hardening

Conclusion

- Linux kernel security has significantly evolved beyond Unix DAC scheme
- Meets a *very* wide range of security requirements
- Security features are mainstream

Resources

- Linux Kernel Security Subsystem Wiki
 - kernsec.org
- LSM mailing list
- LWN security page
- Linux Security Summit
 - San Diego, USA, Aug 2012 with LinuxCon

Questions?