Linux Security Modules

Nikos Kokkalis nikos.kokkalis@gmail.com

November 28, 2017

Linux Security Modules

- SELinux – applied to files (inodes)
- AppArmor – applied to file paths
- SMACK – applied to files, mostly used in embedded systems
- grsecurity – not mainlined, stable patches are no longer free
- TOMOYO – from NTT, current maintenance unclear
- Yama

Pros & Cons

- SELinux – very fine grained control, not going away
- AppArmor – feasible to write policies by hand, usable on FS w/out extended permissions, not going away
- SMACK – policies extremely easy to define
- grsecurity – handles lots of security issues beyond just MAC, stacks with other MAC implementations

Why AppArmor

- Enabled by default in one of the most common distributions, Ubuntu
- Relatively straightforward syntax
- No separate build process for new policies

What does AppArmor do?

Monitor and restrict:

- file access
- network access
- capabilities (chown, mknod, setuid, ...)
  - man 7 capabilities
- rlimit (aka ulimit)
- in general: restrict permissions

AppArmor Components

- Kernel
- Userspace Utilities
- Profiles

Example

```
   ~ sudo aa-status
apparmor module is loaded.
58 profiles are loaded.
25 profiles are in enforce mode.
   /usr/bin/evince
   /usr/bin/evince-previewer
   /usr/bin/evince-previewer//sanitized_helper
   /usr/bin/evince-thumbnailer
   /usr/bin/evince-thumbnailer//sanitized_helper
   /usr/bin/evince//sanitized_helper
   /usr/bin/irssi
   /usr/bin/pidgin
   /usr/bin/pidgin//launchpad_integration
   /usr/bin/pidgin//sanitized_helper
   /usr/bin/totem
   /usr/bin/totem-audio-preview
   /usr/bin/totem-video-thumbnailer
```

Creating Profile

- Start with aa-genprof (install apparmor-utils on Ubuntu)

```
# aa-genprof /usr/sbin/nginx
....
[(S)can system log for AppArmor events] / (F)inish
# #other window
# service nginx restart
```

- Do some operations then “S”can message logs, save and quit.. this will be inadequate
- Set new profile to complain mode with aa-complain and restart service
- Run tail -f and look for apparmor messages
- Edit profile to allow required permissions
- Repeat until warnings are gone
- Try aa-enforce and fix any further errors
- Important to run with exact same setup as production

More AppArmor Commands

- apparmor_parser – core executable
- aa-autodep – generate profile skeleton
- aa-cleanprof – remove superfluous rules, reorganize, and remove comments
- aa-decode: decodes hex strings in kernel log (if present)
- aa-logprof – interactively change profiles from system log warnings
- aa-mergeprof – merge profiles

Seccomp(-bpf)

- Seccomp denies access to all system calls except read, write, and exit.
- seccomp-bpf is an extension to seccomp that allows specifying a filter that is applied to every system call.

Grsecurity (+ PaX)

- a (large, 6.2M) patch against Linux kernel
- security oriented (obviously)
- started more than 14 years ago, pioneered multiple techniques
- includes multiple parts: -PaX protection against memory corruption bugs (sometimes called an HIPS)
- RBAC Role-based access control (not implemented as LSM)
- Other Generic hardening (memory, filesystem, network protections)

PaX

Major features

- NOEXEC segmentation or pagination-based implementation of NX bit (before it was available on CPUs)
- MPROTECT W X at the page level: forbid memory pages available both for write and execute, and completes NOEXEC
- KERNEXEC kernel equivalent of NOEXEC+MPROTECT; prevents injection and execution of foreign code to the kernel
- ASLR predates the Linux version (actually predates all version), and still an improvement
- UDEREF prevents the kernel dereferencing userland pointers, like SMEP/SMAP or PXN/PAN on steroid
- CONSTIFY constify structure containing only function pointers, using a gcc plugin

- GrSecurity/PaX is no longer offered as a free download and, as a result, almost every publicly available distribution has stopped providing it.
- KSPP

Yama does Discretionary Access Control of some kernel related functions, like defining if process tracing (ptrace) is allowed.

Firejail

Firejail is a project that uses Linux namespaces and seccomp-bpf to create a sandbox around Linux applications. (show a profile from /etc/firejail)

Other Security Mechanisms

- PaX
- seccomp(-bpf) – filtering for system calls
- capabilities (instead of setuid)
- sysctl settings
- systemtap (patching the kernel)
- compiler features (gcc stack-protector, etc.)
- namespaces
- cgroups