State of SELinux

Paul Moore

September 2017

Kernel Changes

Obligatory Container Slide

- Individual file labeling for cgroup, cgroup2, and tracefs
  - Allows for labels unique to each container, enabling SELinux enforced separation for these filesystems
- Context mount tmpfs, ramfs, and devpts in non-init namespaces
  - Allows containers running under SELinux to mount these filesystems inside the container

“SELinux <heart emoji> KSPP”

- LSM/SELinux hooks marked read-only after boot
  - Prevents hijacking/bypass of SELinux kernel hooks
- Various internal structures “constified”
  - Prevents malicious tampering of data structures

New Stuff We Can’t Claim As Container Related

- Access controls for prlimit(2)
  - Control access to resource limits of other processes
- Access controls for mmap(2)
  - Enable rights revocation by preventing direct memory access to resources
- Access controls for Infiniband/RDMA
  - Add SELinux access controls to Infiniband traffic
- Display SELinux policy capabilities during policy load
  - Helpful when trying to determine supported capabilities

Old Stuff We Finally Got Working Correctly

- Expand the number of socket related object classes
  - All the visible socket address families are supported
  - Eliminate the generic “socket” object class
- Shuffle the DAC_OVERRIDE / DAC_READ_SEARCH capabilty checks
  - Only check for the capabilities which are needed
- Enable domain transitions under NNP and nosuid
  - New SELinux permission to enable transitions
  - Plays well with the new systemd NNP protections

Fun Stuff I Learned Using gitdm

- Change summary
  - 102 changesets
  - 37 developers
  - 1407 lines added
  - 613 lines removed

- Top 10 developers by lines changed

1)	Daniel Jurgens	584 (37.6%)
2)	Stephen Smalley	564 (36.3%)
3)	Andreas Gruenbacher	89 (5.7%)
4)	Markus Elfring	84 (5.4%)
5)	Scott Mayhew	33 (2.1%)
6)	Junil Lee	27 (1.7%)
7)	Florian Westphal	24 (1.5%)
8)	Kees Cook	20 (1.3%)
9)	Gary Tierney	19 (1.2%)
10)	Matthias Kaehlcke	17 (1.1%)

Userspace and Policy Changes

Something For Everyone ...

- New genhomedircon template additions
  - Enables greater policy flexibility
- Support for ioctl(2) xperms in policy modules
  - Enables the ioctl whitelisting in modular policy
- Generate CIL/policy.conf from binary policy
  - Enables inspection of packaged or loaded policy
- Improved attribute handling
  - Better performance, memory footprint
- Improved libsemanage’s relinking
  - Better performance due to less policy relinking

… Even Your Favorite Distribution

- Migrated to setools4
  - setools3 was deprecated/unsupported
- Improved support for Python 3
- Added support for PCRE2 in libselinux
  - Support building with PCRE1 or PCRE2
- Split policycoreutils into individual components
  - Easier for distributions to package and ship
  - Similar to existing Fedora / RHEL packaging

More Fun With gitdm

- Change summary
  - 588 changesets
  - 37 developers
  - 26655 lines added
  - 14816 lines removed

- Top 10 developers by lines changed

1)	James Carter	10757 (29.9%)
2)	Stephen Smalley	9927 (27.6%)
3)	Daniel Jurgens	5439 (15.1%)
4)	Jason Zaman	3237 (9.0%)
5)	Richard Haines	1487 (4.1%)
6)	Nicolas Iooss	1304 (3.6%)
7)	Janis Danisevskis	957 (2.7%)
8)	William Roberts	495 (1.4%)
9)	Jeff Vander Stoep	323 (0.9%)
10)	Alan Jenkins	309 (0.9%)

SEAndroid

Stats We Can Brag About

- ~92% of devices with some SELinux protection
  - KitKat (v4.4) and above
  - Up from 80% at LSS 2016
- ~77% of devices with full SELinux protection
  - Lollipop (v5.0) and above
  - Up from 50% at LSS 2016

Proof It All Really Works

- “Honey, I Shrunk the Attack Surface” (click for URL)
  - Nick Kralevich, Google at Black Hat 2017

- Security improvements in Android
- Demonstrated effectiveness of SEAndroid against published attacks (CVEs)
- Impact of SEAndroid on the exploit market

More Information

You Should Take A Picture Of This Slide

- Kernel
  - git://git.kernel.org/pub/scm/linux/kernel/git/pcmoore/selinux.git
- Userspace / Tests
  - https://github.com/SELinuxProject
- Reference Policy
  - https://github.com/TresysTechnology/refpolicy
- Mailing List
  - https://www.nsa.gov/what-we-do/research/selinux/mailing-list.shtml
- Me
  - @securepaul
  - paul@paul-moore.com