LSS 2017: linux-integrity subsystem update

Mimi Zohar

Linux Integrity Subsystem Status Update

- Review of IMA goals
- New IMA features and other changes
- TPM related work
- Possible IMA specific “fixes” for TPM performance issues
- Namespacing IMA
- Summary

IMA: Extending secure and trusted boot to the OS

Integrity Measurement Architecture: Linux kernel integrity subsystem whose goal is to detect if files have been accidentally or maliciously altered, both remotely and locally, appraise a file's measurement or signature, and enforce local file integrity.

- IMA-measurement (Linux 2.6.30): "**trusted boot"**
  - Hardware TPM signs list of file hashes, attests to third party
  - Complete Trusted Computing Base (TCB) verification does not scale
- IMA-appraisal (Linux 3.7): "**secure boot"**
  - Verifies signatures on files
  - Scalable prevention, but does not do attestation
- IMA ima-sig template (Linux 3.13) "**combined model**"
  - Attestation of hashes and signatures
- All IMA modes are controlled by a single **policy** file

Linux Integrity Subsystem Ecosystem: IMA-measurement & IMA-appraisal

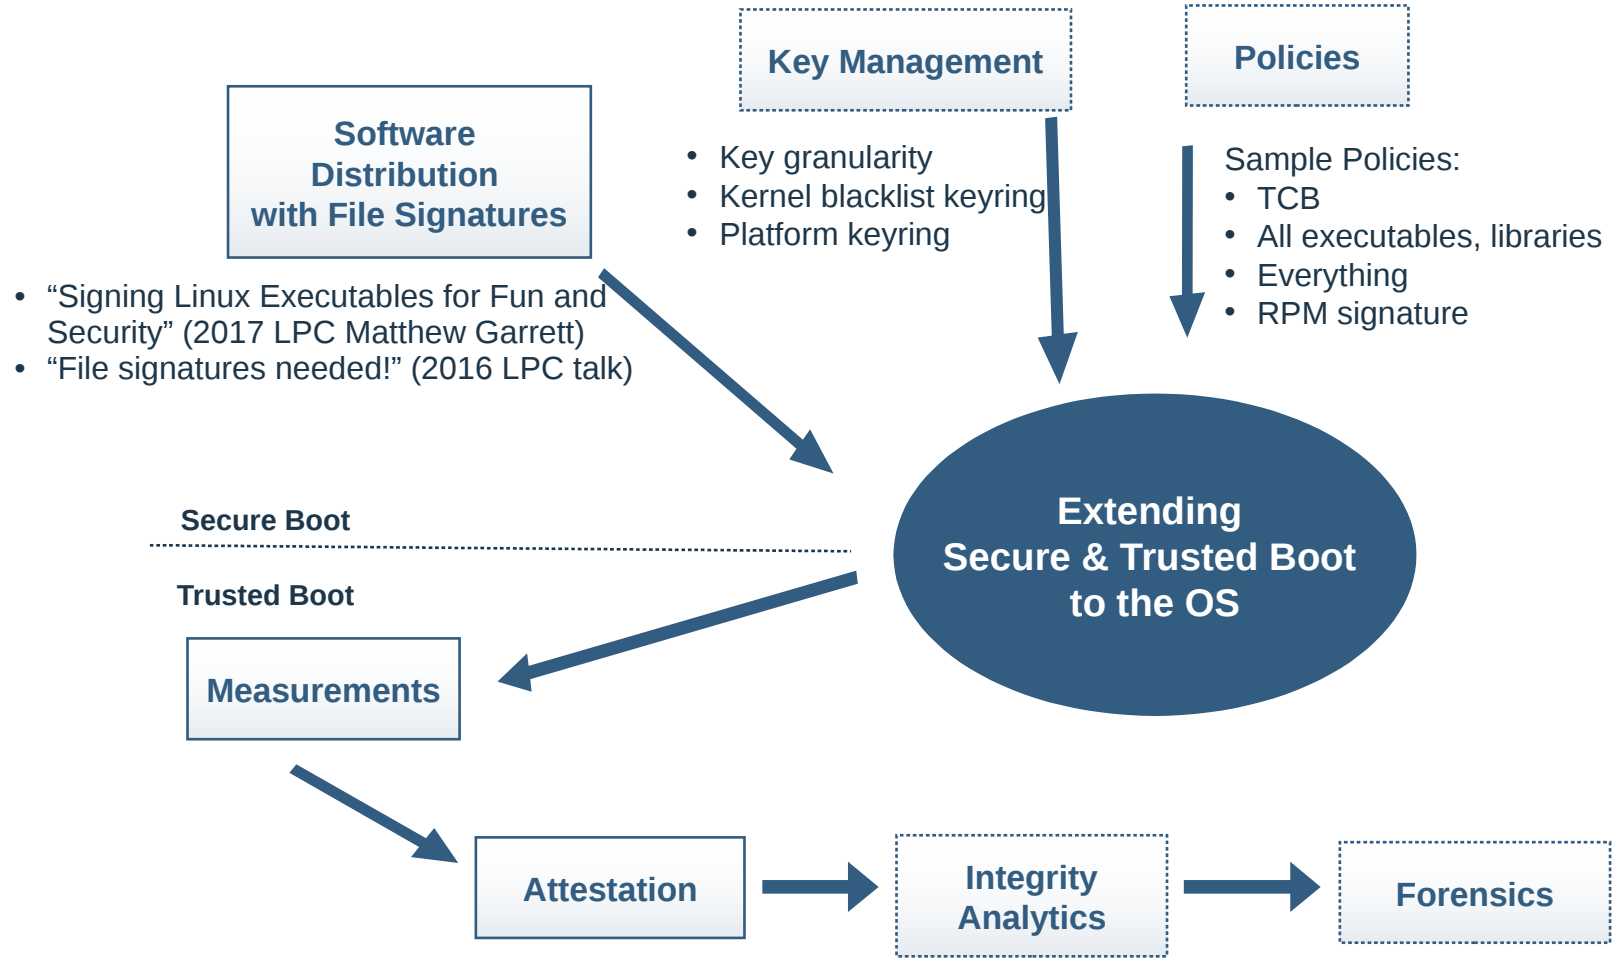

Linux Integrity Subsystem Ecosystem: IMA-audit

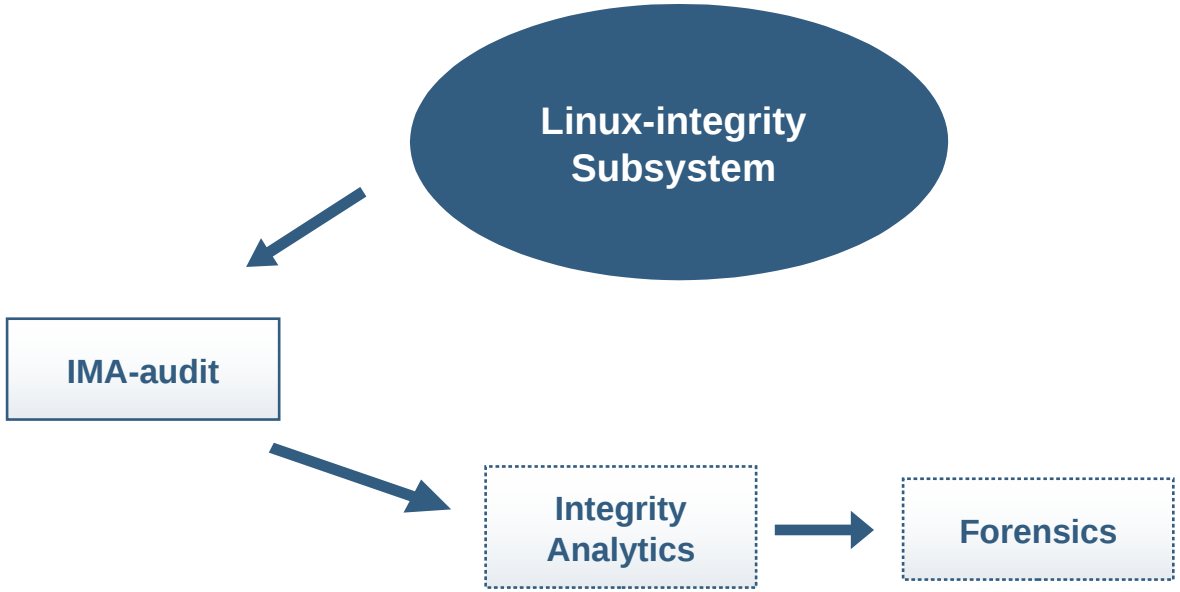

- IMA-audit (linux-3.10) audit logs file measurements – Peter Moody (while at Google)
- Can be used to augment existing integrity analytics and forsenics tools (https://www.fireeye.com/blog/threatresearch/2016/11/extending_linux_exec.html)

New features and other changes

- Carrying the measurement list across kexec (Linux-4.10) - Thiago Bauermann (IBM LTC Core Kernel Team)
- Embedding IMA more deeply into the VFS layer
  - New integrity_read file operation method - Christoph Hellwig
  - Reporting i_version status
- modsig: appended signature support - Thiago Bauermann
- platform keyring: using “UEFI“ keys for verifying the kernel image - Nayna Jain (IBM LTC)

IMA: measurements, TPM performance & memory usage

- Spectrum of what “needs” to be measured
- TPM performance & memory impact
- What are we trying to detect?
- Need some sample policies
- Usecase:
    - ~1,500 measurements from boot to logon prompt
    - ~5,500 additional measurements to application start
    - ~63,000 application measurements (pre-measure)

TPM related work

- Extend/truncate SHA1 hash to extend TPM 2.0 enabled banks (linux-4.11) : Nayna Jain (IBM LTC)
- IMA hash agile format: Roberto Sassu (Huawei)

- Performance degradation: non-cascading wheel timer (linux-4.8)
  - Nuvoton driver: replace msleep() with usleep_range()
  - Nuvoton driver: removing unnecessary msleep(): Nayna Jain (IBM LTC)
  - tpm_msleep() wrapper: Attack Hazma (HP)
- Other performance improvements: Ken Goldman & the TCG device driver wg
  - Send/receive entire command, not just burst_count size (Nayna Jain)
  - Reduce long poll timeouts (Nayna Jain)
  - Other performance improvements ???

Possible IMA specific “fixes” for TPM performance issues

- IMA digest white lists: Roberto Sassu (Huawei)
  - Single white list measurement, fewer TPM extends
  - Only “unknown” measurements in measurement list
  - Issues:
    - Loss of file access information
    - Measurements hidden behind white list measurement
    - Integrity guarantee change: measurement list + possible white list files
- Queueing and batching TPM extends
  - Integrity guarantees unchanged
  - Maximum file measurement extend is 2x TPM roundtrip
  - Requires new measurement list template field or record type

Namespacing IMA*

- Which aspects (eg. ima-measurement, ima-appraisal, ima-audit)?
- What are the goals?
- What are the issues/concerns?

* Based on initial PoC IMA namespacing discussions by Yuqiong Feng, David Safford, Dimitrios Pendarakis, w/fixes by Stefan Berger, and rebased by Mehmet Kayaalp.

Namespacing IMA-measurement

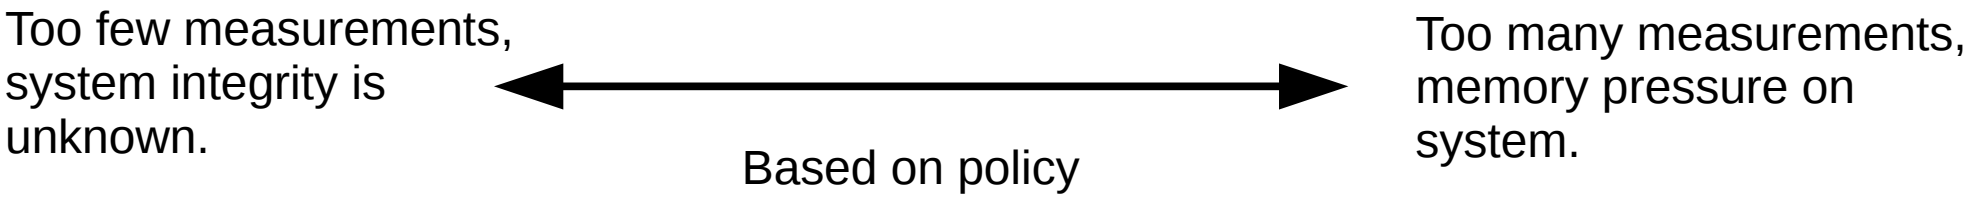

- Too many or is it too few measurements?
- Containers come and go ...
- Where do we store the measurements (eg. native vs namespace, or both)?
- Who defines the namespace measurement policy - the container owner or the container orchestration tools? (Requires securityfs namespace support.)
- Is there an initial “builtin” namespace measurement policy?
- Can root in the namespace replace the policy?
- Measurements with per namespace info result in more measurements.
- Minimize amount of namespace specific info in native list?

Namespacing IMA-appraisal

- Enforce file integrity based on namespace keys and keyrings
  - With Mat Martineau's "Make keyring link restrictions accessible from userspace" patch set, we can simulate a kernel trusted keyring.
  - Trusted dot prefixed keyrings are "special". How "special" are they?
- Who defines the namespace appraisal policy - the container owner, the container orchestration tools, or the native system?
- Is there a default appraisal policy, similar to the builtin "appraise_tcb" policy?

- Assumptions:
  - Support for multiple security xattrs (on a per "user" basis).
  - Permit root in the namespace to write security xattrs.
- Concerns:
  - xattr file system size limitations

Namespacing IMA-audit

- Audit log messages containing the file hash on a per namespace basis
- Ability to correlate containers with namespace ids
- Integrate with integrity analytics & forensic software

- No other kernel changes required

Staging IMA namespacing

1) IMA-audit: define architecture
- Safely initializing/freeing IMA namespaces
- Extend integrity audit messages with namespacing info

2) IMA-measurement: introduce securityfs files
- Define per namespace policy
- Defind per namespace measurement list

Concerns:
- Prevent securityfs information leakage
- Hierarchical measurements – too many/few measurements

3) IMA-appraisal: introduce keyrings and xattrs
- Define per namespace keyrings
- Support for multiple security xattrs (on a per “user” basis?)
- Permit root in the namespace to write security xattrs

Concerns:
- xattr file system size limitation

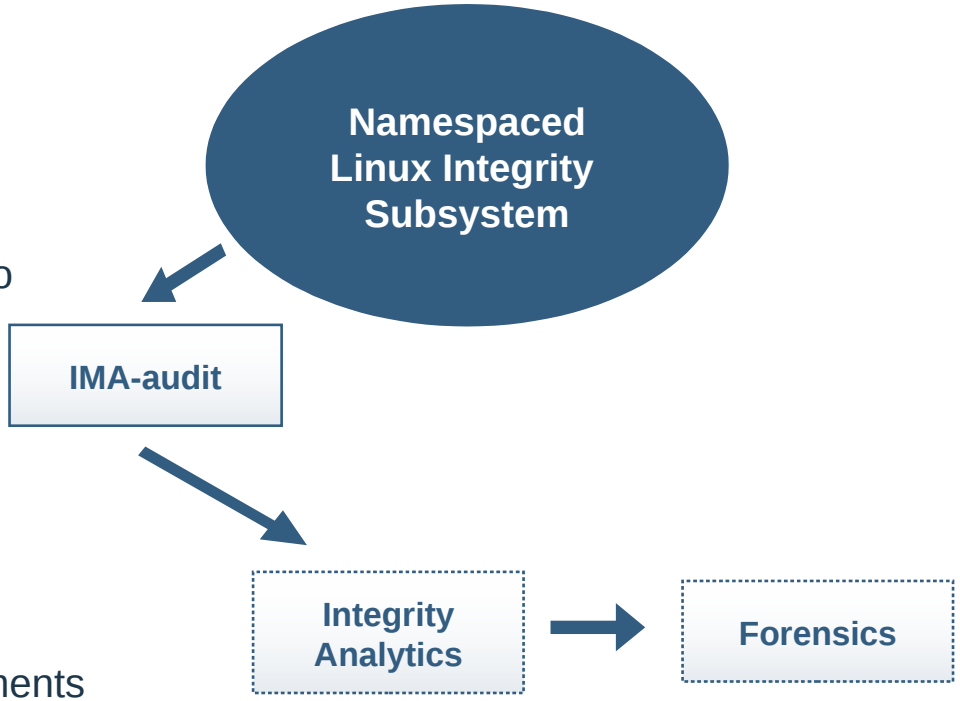

Linux Integrity Subsystem: summary

- Need to simplify usage
  - Understanding which filesystems support IMA and their level of support
  - Define sample sane policies for different use cases
  - Signining files: from distro's, firmware vendors, package maintainers
- Performance improvements
  - TPM related work
  - Possible IMA specific "fixes" for TPM performance issues
- Continuing to add support for new features
  - modsig - appended signature support
  - Platform keyring – using UEFI keys for verifying the kernel image
  - IMA hash agile measurement list
- Staging IMA Namespacing: ima-audit, ima_measurement, ima_appraisal

Questions?

Current and future work

- Appended signature support -Thiago Bauermann (IBM LTC Core Kernel Team)
- Using the platform keyring (UEFI keys) for validating file signatures - Nayna Jain (IBM LTC)
- IMA hash agile measurement list - Roberto Sassu (Huawei)
- TPM performance
- Simplify usage: sample policies, updating documentation
- Key management: black lists, revocation, resetting IMA cache status
- Namespacing IMA (ima-audit, ima-measurement, ima-audit)
- initramfs: CPIO extended attribute support
- Closing measurement gaps: eBPF, ?
- Directory protection support
- Re-work the IMA measurement list memory allocation