Rethinking Audit

Paul Moore
Red Hat
August 2015

The Problems

- Event record strings are generated in the kernel
  - Lots of vsnprintf() calls and memory copies
  - Escaping “untrusted” strings in the kernel
- Event records are larger than needed
  - Overhead when writing to userspace, storage, remote log servers
- Event records are not easily parsed for searching, translation, etc.
  - Records are poorly formatted quasi name/value pair strings
  - Record formatting is the responsibility of the caller, not audit

One Possible Solution

- Remove the formatting burden from individual kernel subsystems
  - Replace audit_log_format() with data specific APIs:
  - audit_log_int(AUDIT_FIELD_RES, result)
  - audit_log_string(AUDIT_FIELD_SUBJ, lsm_subj_string)
  - audit_log_creds(cred)
- Support the generation of audit records in binary format
  - Move the name/value pair idea to Netlink attributes
  - Leverage attribute nesting to minimize records for each event

Thoughts on the Attribute Solution

- If nothing else, there is value in removing audit_log_format()
  - Well behaved subsystems don't need to worry about formatting
  - Easier to spot event record abuses
  - Enables any future work we want to do regarding record formats
- BC: “Backwards Compatible” or “Badly Cursed”?
  - The string format will need to be supported for years to come
- Supporting the audit multicast listeners
  - Continue to generate string based records just for multicast?
  - Modify string record format, e.g. proper name/value pair format?