Ambient posix capabilities

Serge Hallyn

Canonical, Ltd

serge.hallyn@ubuntu.com

August 21, 2015

Overview

Capability recalculation

```
pB' = pB
pI' = pI
pP' = (pB & fP) | (fI & pI)
pE' = fE ? pP' : 0
```

Ambient capabilities

- What people thought pI should have been
- `pA` must be a subset of `pI` and `pP`
- Current ways of dropping privilege still work
  - pA cleared for setuid, fcaps, keepcaps

Ambient capabilities

- What people thought pI should have been
- pA must be a subset of pI and pP
- Current ways of dropping privilege still work
  - pA cleared for setuid, fcaps, keepcaps

```
pB' = pB
pA' = (fcaps|setxid) ? 0 : pA
pI' = pI
pP' = (pB & fP) | (fI & pI) | pA'
pE' = fE ? pP' : pA'
```

Namespaced file capabilities

- File caps cannot be set in user namespace
- Software in container must handle both file caps and setuid-root
- Less complexity $\rightarrow$ better
- Proposed:
  - sets of file capability sets
  - each tagged with userns root k_uid