Secure Systems Group, Aalto University

Yizhou Ye and Lachlan J. Gunn

Semi-Automated Integrity Policy Analysis

Goal

- Preserve the integrity of important mutable files on a system.
    - Configuration, data, etc. that need to change at run-time
- Allow normal developers (not just security experts) to protect their applications' resources

Integrity policies

- Critical resources marked as high-integrity.
- Untrusted entities are marked as low-integrity.
- The classical integrity model Biba aims to preserve this integrity by enforcing:
  - No read down. High-integrity subjects can't read low-integrity objects.
  - No write up. Low-integrity subjects can't modify high-integrity objects.
- Biba model is too inflexible in practice (no good for e.g. web services).
- This has led to the development of the Clark-Wilson policy that introduces the concept of filters through which high-integrity subjects are allowed to read low-integrity objects.

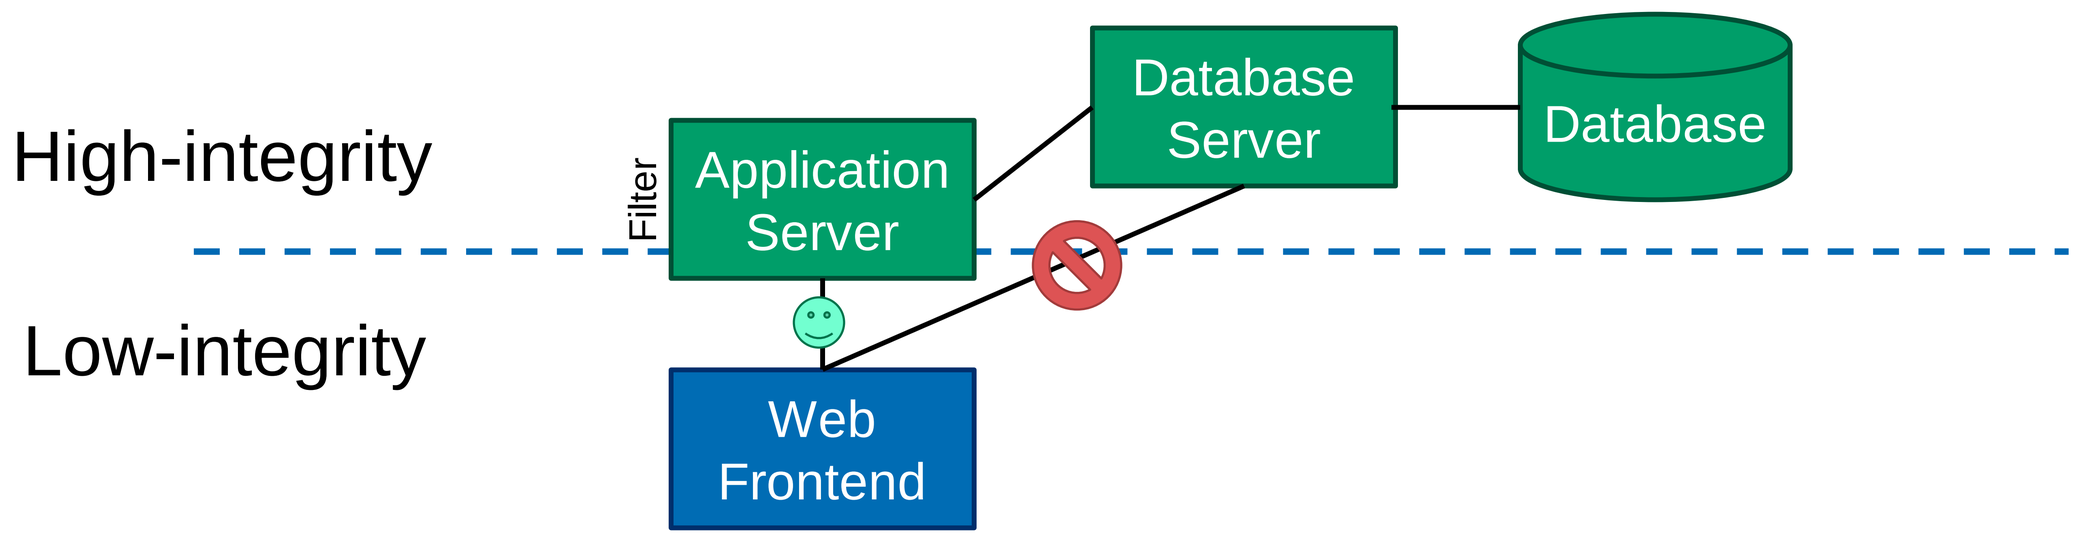

Approach

- Use an easily-comprehensible Clark-Wilson policy to protect important resources.
- Use existing SELinux policy to identify information flows.
- [Semi-]automatically generate a Clark-Wilson policy from these information flows

Work so far

- Reviewed information-flow analysis tools.
- Developed code complexity estimator to guide Clark-Wilson policy formulation.

Aalto University

HUAWEI