Talk: Thursday 1st Session on Storage at 9:50 am

Guardat: Enforcing data policies at the storage layer

Anjo Vahldiek-Oberwagner[1], Eslam Elnikety[1], Aastha Mehta[1], Deepak Garg[1], Peter Druschel[1], Rodrigo Rodrigues[2], Johannes Gehrke[3,4], Ansley Post[5]

[1]MPI-SWS [2]NOVA LINCS/Nova University of Lisbon [3]Microsoft [4]Cornell [5]Google

Max Planck Institute for Software Systems

1. Problem

Complex storage systems threaten data integrity and confidentiality.

- **Bugs**, security **vulnerabilities**, **viruses**, operator errors, misconfiguration
- **Protection** implicit in application code or configuration **spread** across software layers
- Lack of **transparency/accountability** in third-party storage

Examples:

- A virus overwriting a system binary with an infected binary
- A file system bug falsely allowing access to private data

2. Guardat

Storage layer mediates all I/O, enforces user-defined data policy per file and certifies its state.

Key Idea:

- Data owner, provider, system designer and legislators provide per-file **policy**
- Storage device **intercepts** I/O and **enforces** the policy
- Storage device **certifies**
  - its own properties (type, reliability, etc.)
  - current policies associated with stored files

Benefits:

- Safety net against viruses, bugs, or operator error
- Centralized & declarative policy specification

3. Policy

Access rules in declarative policy language, conditioned on:

Identity

External facts (e.g. time)

HW/SW configuration

File state (content, size, …)

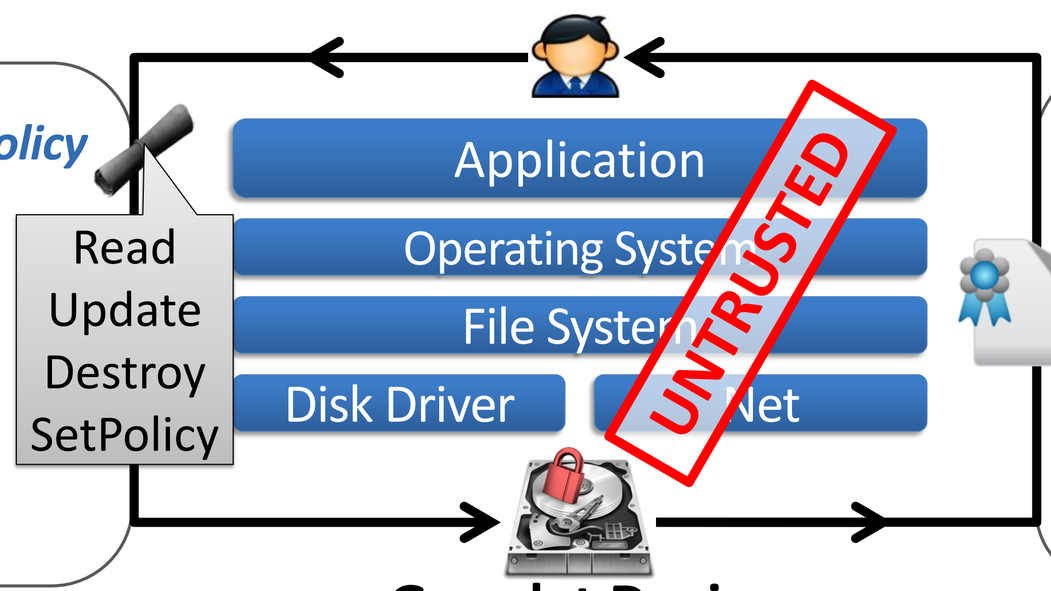

4. Certificate

Guardat attests file state & policy.

Full name

Associated policy

Size & hash

Physical layout

Device properties

5. Guardat Controller

A controller (e.g. disk/RAID firmware, or extra microcontroller) that provides security primitives.

- **Trusted firmware** with secure updates
- **Declarative policy language**
- **Cryptographic support** (embedded key, …)
- **Protocols for secure channel** establishment with remote applications

6. Data Protection Examples

Data confidentiality, integrity & accountability guarantees depend only on Guardat integrity.

Integrity:

- Append-only files
- Protected executables
- Mandatory access logging

Update: fileNameIs(f) ^ K_{Vendor} signs nextVersion(f, nH, cH) ^ f hasHash cH ^ f willHaveHash nH

Read: sessionIs(Bob_pk)

Confidentiality:

- Read private files using an authenticated secure session

7. Evaluation

iSCSI SAN prototype with moderate overhead.

- Latency: HDD < 1%, SSD < 18% (random accesses < 2x)
- Throughput: HDD & SSD < 2%