Qapla: Ensuring policy compliance in database-backed systems

MAX-PLANCK-GESELLSCHAFT

Aastha Mehta, **Eslam Elnikety**, Deepak Garg, Peter Druschel

[early stage work in progress]

Max Planck Institute for Software Systems

1. Compliance failure in OSNs, search engines, healthcare providers, finance companies…

… causes loss of data, reputation, revenue

Goal: light-weight, application-independent, query-aware policy compliance layer

Challenges

- Sensitive data from individuals, many other sources
- Policies depend on data joins and transformations, time, location, principals
- Enforcement mangled with application code
- Application codebase dynamic

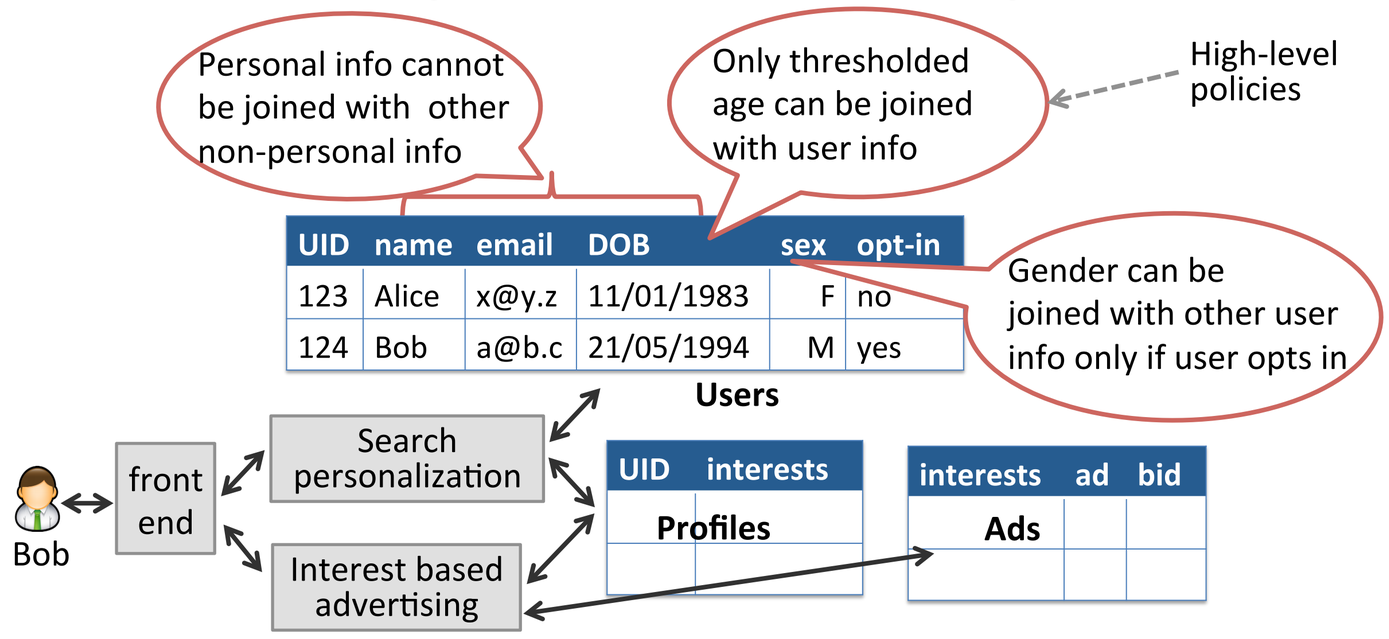

2. Qapla

Policies

- Declarative information flow conditions
- Attached to columns, rows, cells

Enforcement

- Coarse-grained policy tracking across queries
- Declassification based on SQL operators, system state (e.g. clock time)

Policy designer must reason about inference attacks, attacker knowledge

3. Policy restrictiveness proportional to information leak

Transforming data reveals no more info than input data
SELECT age > 30 FROM Users *vs.*
SELECT age FROM Users

→ Policy on transformed data **less restrictive than** source policy

Policy depends on past data transformations

Joining multiple data items may reveal more info than individual items
SELECT DOB, interests FROM
Users JOIN Profiles ON UID

→ Policy on joined data **more restrictive than** conjunction of source policies

Policy is a function of the input policies and join operation

4. Policy specification

Declarative policies
Finite state machines with merge operations

Users

UID	name	DOB
123	Alice	11/03/1983
124	Bob	24/12/1994

Profiles

UID	interests

Intended policy

<u>Within SQL</u>

A. Only thresholded age can be extracted (by provider)

B. Only thresholded age can be joined with user info (by analysts)

Qapla policy state machine

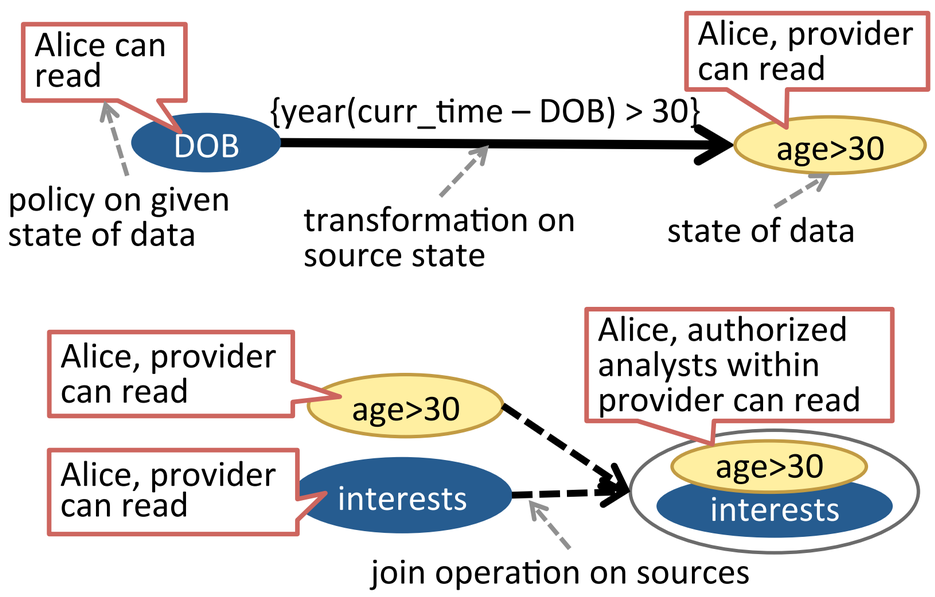

{Example analyst query allowed:
SELECT interests FROM Profiles JOIN Users on UID WHERE age > 30}

<u>Outside SQL</u>

C. Joined data is made public after end of 2025 (for research)

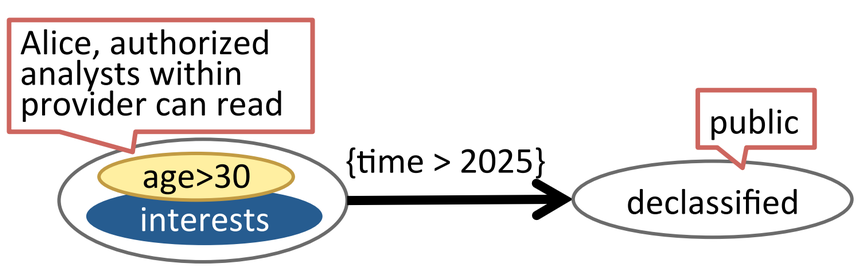

5. Enforcement

Policy tracking within SQL

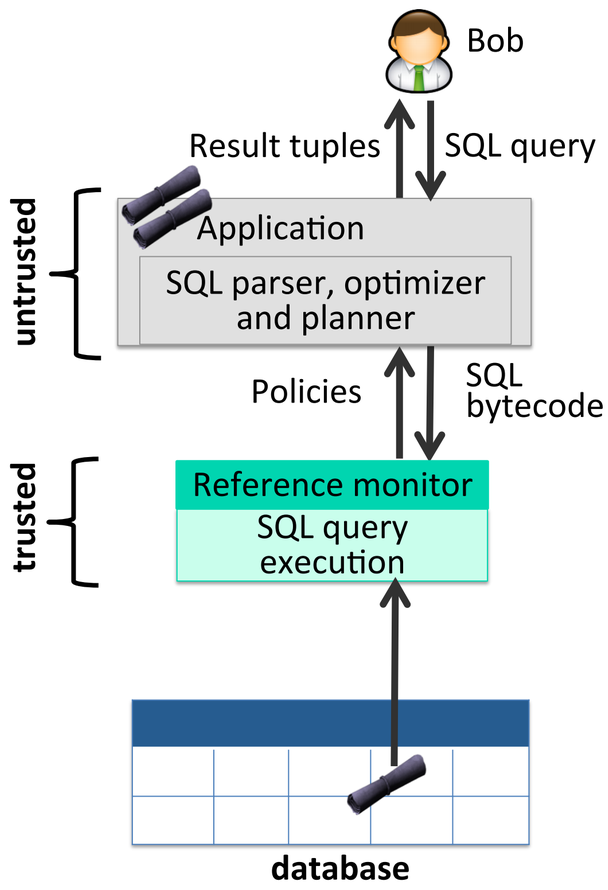

Qapla reference monitor

- Static analysis of compiled SQL query
- Track row accesses in query execution