Pelton: Privacy-Compliant Storage For Web Applications By Construction

Abstract.

Data privacy laws like the EU's GDPR grant users new rights to their data, such as the right to request access and deletion. Manual compliance with these requests is error-prone, and imposes costly burdens, especially on smaller organizations, as non-compliance risks steep fines.

Pelton is a new, MySQL-compatible database that complies with privacy laws by construction. The key idea is to make the data ownership and sharing semantics explicit in the storage system. This requires Pelton to capture and enforce applications' complex data ownership and sharing semantics, but in exchange simplifies privacy compliance. Using a small set of schema annotations, Pelton infers storage organization, procedures for data retrieval and deletion, and reports compliance errors if an application risks violating the GDPR.

Annotation	Example
DATA SUBJECT	CREATE TABLE users (...) DATA SUBJECT
$T_A(x)$ OWNED BY $T_B(y)$	stories(author_id) OWNED BY users(id)
$T_A(x)$ OWNS $T_B(y)$	share(file_id) OWNS files(id)
$T_A(x)$ ACCESSED BY $T_B(y)$	orders(uid) ACCESSED BY auth_user(id)
$T_A(x)$ ACCESSES $T_B(y)$	member(group_id) ACCESSES groups(id)
ON DEL $T_A(x)$ ANON $T_A(...)$	ON DEL orders(uid) ANON orders(name, address)
ON GET $T_A(x)$ ANON $T_A(...)$	ON GET chat(receiver) ANON chat(sender_name)

Figure 1: A list of schema annotations that Pelton makes available to developers

```
CREATE TABLE member (
  id INT PRIMARY KEY,
  uid INT NOT NULL OWNED BY user(id),
  gid INT NOT NULL OWNS group(id)
);
CREATE TABLE share (
  id INT PRIMARY KEY,
  uid_owner INT NOT NULL OWNED BY user(id),
  share_with INT ACCESSED BY user(id),
  share_with_group INT ACCESSED BY group(id), ...
);
```

Figure 2: An example of using Pelton's annotations in a partial schema for OwnCloud, one of the applications we tested Pelton with

We built a prototype of Pelton and evaluated its expressivity and performance. Pelton successfully expresses the data sharing semantics of real web applications, and guides developers to getting privacy compliance right. Pelton also matches or exceeds the performance of existing storage systems, at the cost of a modest increase in state size.