

Seminar: Linked Data and the Semantic Web
Authors: Lukas Kubelka and Benjamin Meyjohann

The desirable target state would be an integrated data ecosystem, enabling organizations to make informed decisions based on accurate and up-to-date information. A key aspect of this challenge is the integration of Identity and Access Management (IAM), which in legacy systems is typically performed separately for each data source.

For that we utilize the Ontology-based Data Access framework and provide R2RML mappings that enable this transition. Next to the mappings of the actual data, our main contribution lies in the mapping of Role-based Access Control (RBAC) policies to Web Access Control (WAC) ACL resources.

For the practical implementation of our theoretical framework, we selected PostgreSQL as our relational database for its open-source nature and widespread adoption. Ontop was chosen as the virtual knowledge graph system for its ability to provide a SPARQL endpoint for just-in-time data transformation, interfacing seamlessly with both PostgreSQL and the Solid protocol. The Solid



The Solid server configuration is critical for routing requests to the appropriate data sources. We use `RegexRule` configurations to direct all RDF graph store queries to the SPARQL endpoint and other requests to the server's file storage. This ensures that requests for RDF data are processed through the virtualization layer, maintaining real-time data fidelity and applying the necessary access controls as defined by the R2RML mappings as shown in Fig. 2. Additionally, the R2RML mappings account for Solid's requirement of meta-resources for each resource accessed.

This method integrates PostgreSQL databases with the capabilities of Ontop, providing an accessible solution for exposing legacy relational data as linked data within the Solid ecosystem.

By employing virtualization over materialization, it significantly reduces storage demands, guarantees real-time synchronization between legacy and linked data, and ensures accessibility by embedding the data in the Solid ecosystem, making it well-suited especially for step-by-step adaptation from legacy relational data to linked data when running both systems in parallel.

Nonetheless, there are still issues open for future research. These include writing and updating operations at the linked data level synchronizing to the legacy relational data. Furthermore, we expect a script for generation of the R2RML mappings to be the next big step towards even more accessibility.

Solid
Community
Server requests
ACL resources
using 'acl'

This will be dynamically replaced with the WebID from the SQL query for each table.

This is dynamically replaced by the access mode from the SQL query for each table. Currently, only 'Read' can be used.

'{table_name}' is dynamically replaced by the table name

SQL query used by mappings for metadata resources that are required by Solid Community Server

Metadata mapping for

Solid Community Server always requests metadata

Metadata mapping for the ACL resource

```

TablesAclMetadataMap a rr:TriplesMap;
rr:logicalTable :RDFGraphStoreView;
rr:subjectMap [
  rr:template "http://localhost:3000/rdf-graph-
               store/{table_name}/.acl";
  rr:graphMap [
    rr:template "meta:http://localhost:3000/rdf-graph-

```

Fig. 2: The most important B2BMI Mappings