Closing SolidLab Launching Trustflows

2026-01-19

Vlaamse overheid

DEPARTMENT OF ECONOMY SCIENCE & INNOVATION

Gefinancierd door de Europese Unie
NextGenerationEU

SolidLab highlights

The **technical** perspective

Pieter Colpaert (IDLab – UGent)

And

Vincent Naessens (DistriNet – KU Leuven)

Beatriz Esteves (IDLab – UGent)

Vincent Naessens
KU Leuven – Distrinet

Complex security challenges

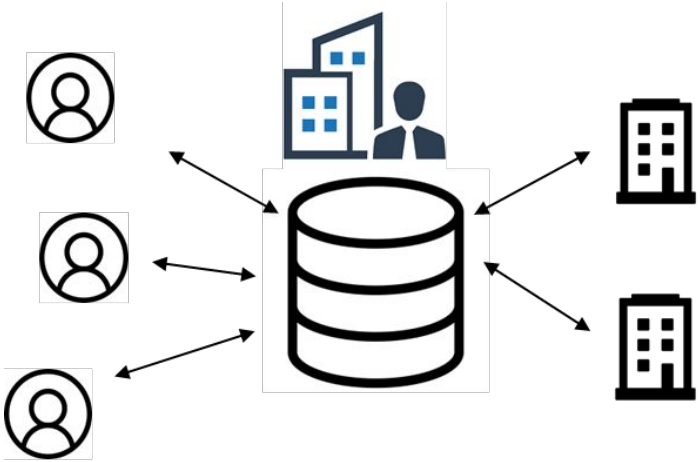

Massive amounts of data managed by Limited number of tech providers

Emerging amount of data breaches

- "data sources become honeypots"
- Stealing sensitive data || ransomware
- Impact can often not be undone
- *Data protection*

Semi-trusted platform providers

- "Honest-but-curious entities"
- Surveillance || prior knowledge
- Leaking data to intelligence agencies
- *Data sovereignty*

Solution space – data protection

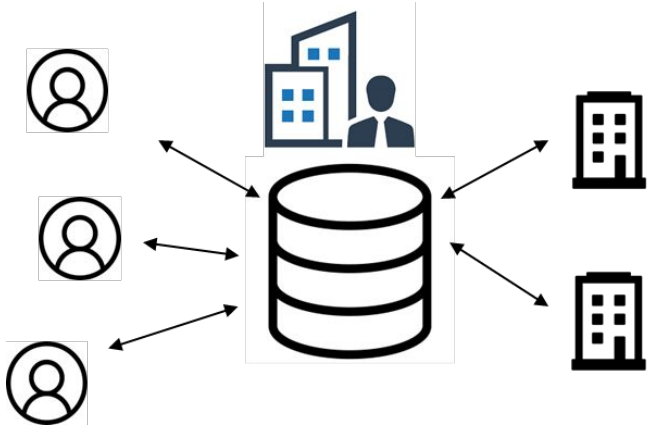

Controlled access to data (data at rest)

- Authentication and Access control
- Cryptographic technologies

Controlled processing (data in use)

- Trusted execution environments (TPM…)
- Processing on encrypted data

Secure communication (data in transit)

Solution space – controlling data flows

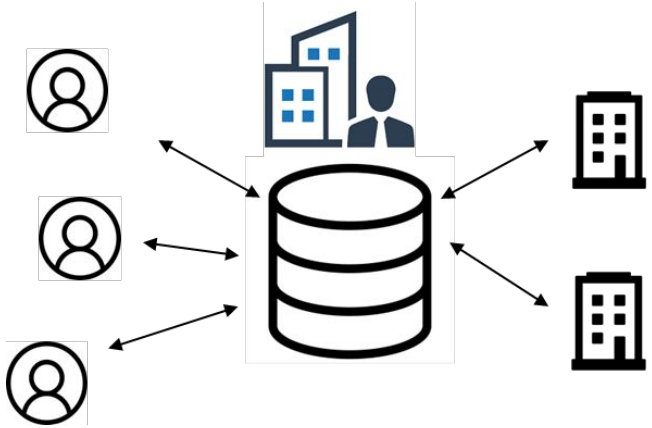

Advanced data collection strategies

- Local data management
- Pseudonymous data pods
- Decentralized data management

Controlling data retention

Transforming data over time

Controlled release of sensitive data

- Privacy-friendly queries (pull based)
- Dataset anonymizations (push based)

Example 1:
Collective anonymous release of pods for research purposes

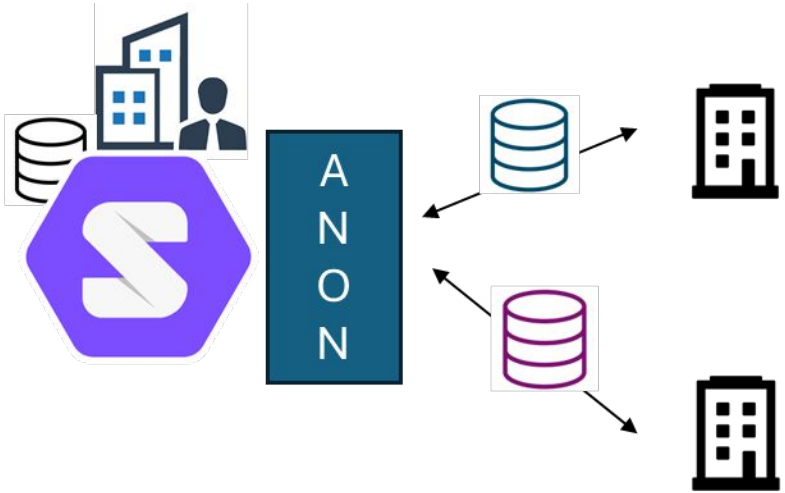

Non-perturbative techniques

- Generalization
- Sampling
- Outlier suppression

Perturbative techniques

- Randomization
- Data swapping
- Noise addition

Example 2:
Federated learning with Solid

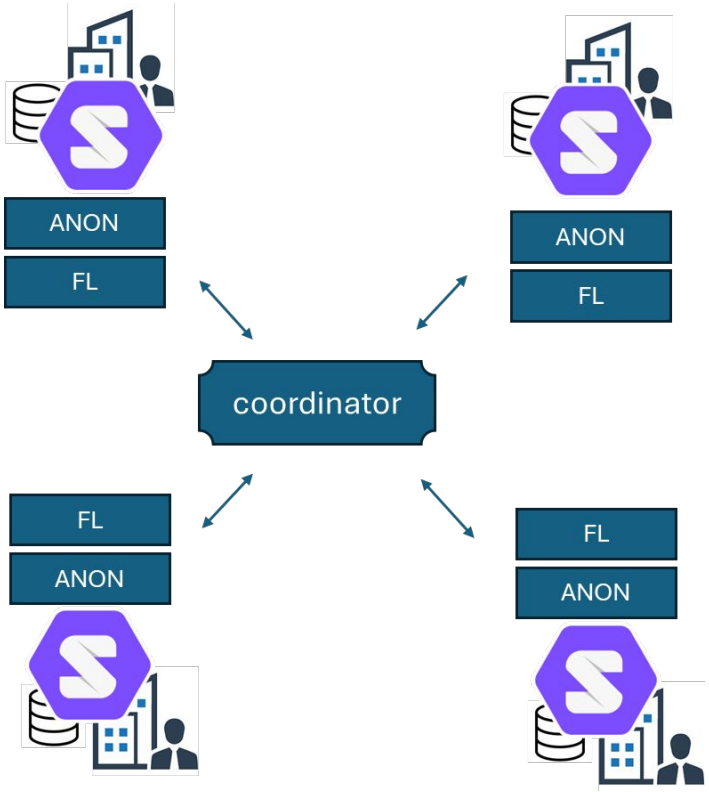

ANON
transforms data from multiple pods locally until a k-anonymity and l-diversity level is reached

FL
builds a local model based on anonymized data

Coordinator
builds a resulting model based on local models

We build **software**,

implementing **specifications** and **standards** we also contribute to.

We publish our **research** findings.

KPI **Contributed to 150+ projects**

KPI **Contributed to 50+ specs**

KPI **80+ publications**

Community Solid Server

https://github.com/CommunitySolidServer

A community storage implementation that became *the* reference implementation.

Allows to experiment with new features: ideal playground for **W3C LWS** and **W3C Solid**

The basis is stable and will be further maintained by IDLab.

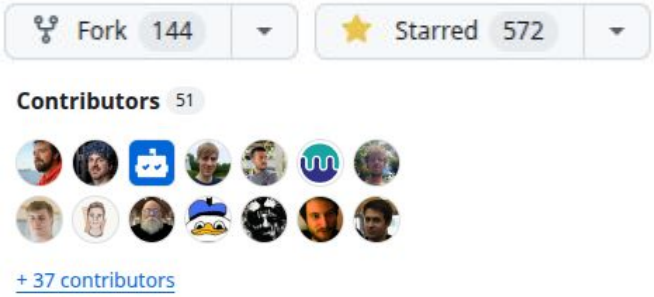

Comunica

https://comunica.dev

The client-side query engine was extended with Solid auth support.

It's a playground for **W3C SPARQL** and new data architecture paradigms.

The project will be further developed in follow-up projects.

There's a demo
– see minute madness

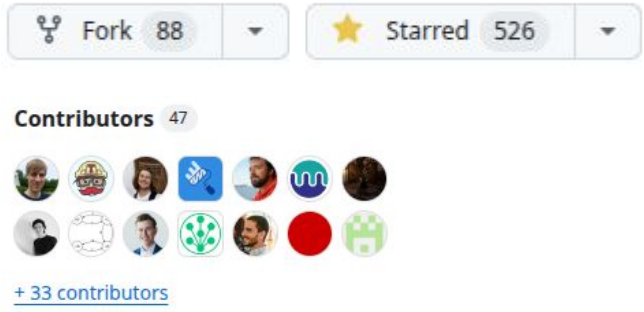

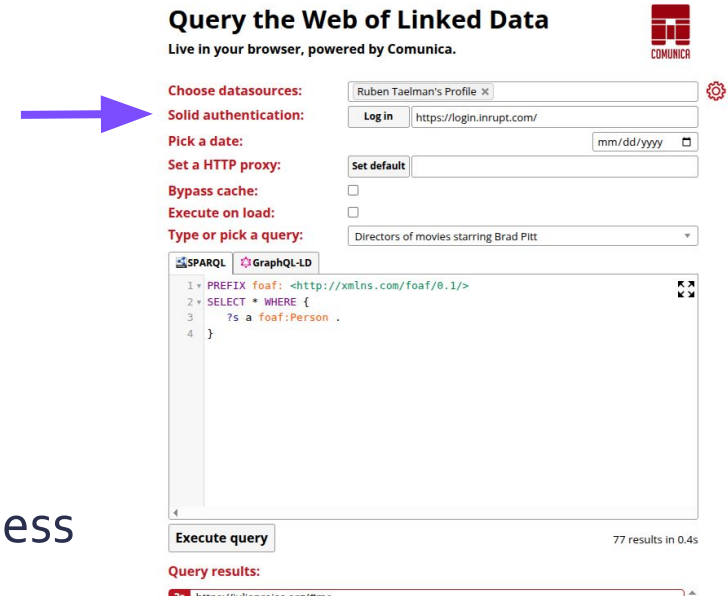

RML: mapping data to a Knowledge Graph

https://rml.io/

Allows to **transform** and share **existing data**,
from existing data management systems (RDB, API, ...),
using the Solid protocols.

Our implementations are a playground for **W3C Knowledge Graph Construction**

During SolidLab,

- Integration with Solid APIs and VC generation was introduced
- An algebraic foundation was formalized and implemented
- Feature extensions were researched
  - **incRML**: always up-to-date knowledge graphs
  - **Logical Views**: abstraction layer to merge any type of (deeply nested) source data

There's a demo
– see minute madness

EYE reasoner

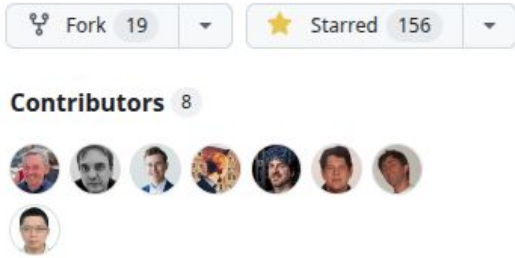

https://github.com/eyereasoner

Interoperable **logic** on Web-scale

Essential for:

- proving something is true
- knowledge engineering
- validation
- ...

It was applied in SHARCS, and is the basis for our work on a policy engine.

And much more

https://github.com/solidlabresearch

= a treasure trove for Solid developers for:

- Aggregating,
- stream reasoning,
- authorization with ODRL policies,
- etc.

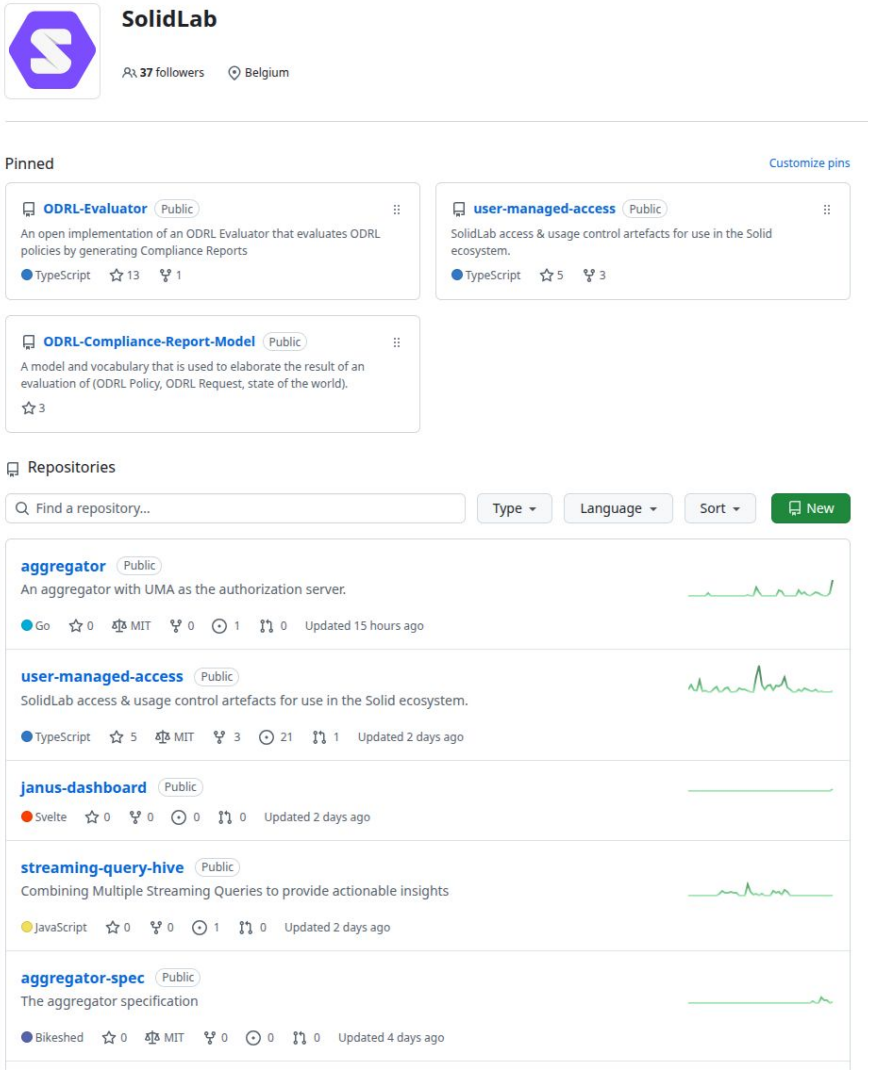

Impact starts with a specification

RDF, SPARQL 1.2, RML, Triple Pattern Fragments, the Solid specifications, Linked Web Storage (LWS), Verifiable credentials, RDF Stream Processing, Linked Data Event Streams (SEMIC), SHACL, ODRL, Framework for ODRL Rule Compliance through Evaluation (FORCE), Data Privacy Vocabulary, DCAT-AP Feeds (SEMIC), Various OSLO Flanders specifications (Flanders), CPSV-AP, RDF Time Series Snippets, TREE hypermedia, Event Notifications, Shape Indexes, the function ontology, Incremental RML, Authorization for Dataspaces, Yarrrml, Notation3, Time Functions, DCAT-AP, …

We maintain a full list of specifications we build, contribute to, or endorse at https://spec-catalogue.knows.idlab.ugent.be/

Towards enterprise: Kvasir

https://kvasir.pages.ilabt.imec.be/kvasir-server/what-is-kvasir.html

A microservices architecture for high throughput writing
preparing data for specialized read interfaces.

Seeks compliance with the relevant parts of S3, Solid and (soon) Linked Web Storage.

Decouples authorization from storage
→ *bring your own authorization server*

Specialized read interfaces for:

- GraphQL-based subsets
- Time series
- Assets
- …

Solid separates data from apps

2021

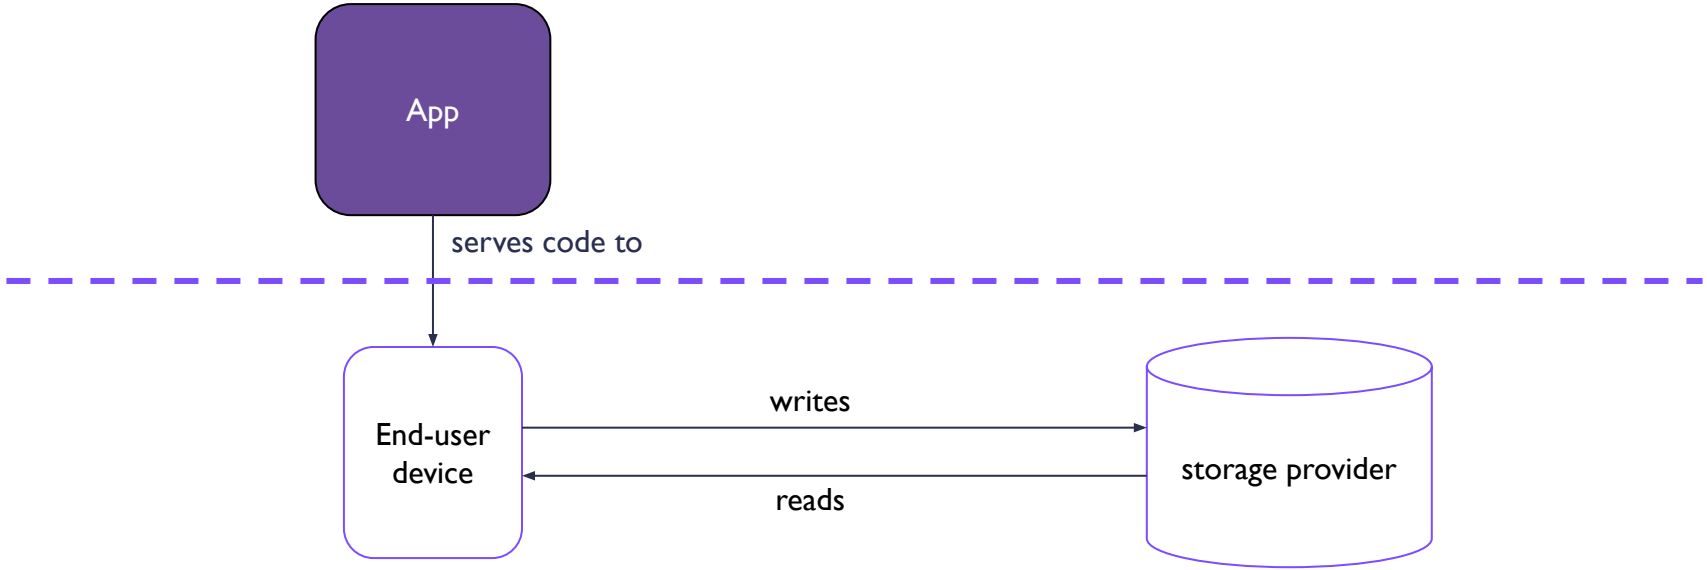

New opportunities for

1. redefining our relation to app builders
2. simplifying legal compliance and putting the user *in control*

We need more than just storing a file

2022

What's in a Pod?

A knowledge graph interpretation for the Solid ecosystem

Ruben Dedecker[1] Wout Slabbinck[1] Jesse Wright[2] Patrick Hochstenbach[1] Pieter Colpaert[1] Ruben Verborgh[1]

[1]IDLab, Department of Electronics and Information Systems, Ghent University – imec
[2]Australian National University, College of Engineering & Computer Science

In reply to: *QuWeDa 2022 call for papers*

The Solid vision aims to make data independent of applications through technical specifications, which detail how to publish and consume permissioned data across multiple autonomous locations called *"pods"*. The current document-centric interpretation of Solid, wherein a pod is a single hierarchy of Linked Data documents, cannot fully realize this independence. Applications are left to define their own APIs within the Solid Protocol, leading to fundamental interoperability problems and the need for associated workarounds. The long-term vision for Solid is confounded with the concrete HTTP interface to pods today, leading to a narrower solution space to address core issues. We examine the mismatch between the vision and its prevalent document-centric interpretation, and propose a reconciliatory graph-centric interpretation wherein a pod is a hybrid, contextualized knowledge graph. In this article, we contrast the existing and proposed interpretations in terms of how they support the Solid vision. We argue that the graph-centric interpretation can improve pod access through different Web APIs that act as views into the knowledge graph. We show how the latter interpretation provides improved opportunities for storage, publication, and querying of decentralized data in more flexible and sustainable ways. These insights are crucial to reduce the dependency of Solid apps on implicit API semantics and local assumptions about the shape and organization of data and the resulting performance. The suggested broader interpretation can guide Solid through its evolution into a heterogeneous yet interoperable ecosystem that better supports the diverging read/write data access patterns of different use cases.

What's in a Pod? a knowledge graph interpretation for the Solid ecosystem
In *6th Workshop on Storing, Querying and Benchmarking Knowledge Graphs (QuWeDa) at ISWC 2022* (Vol. 3279, pp. 81-96). CEUR.

Towards an **evolving trust relation**

2024

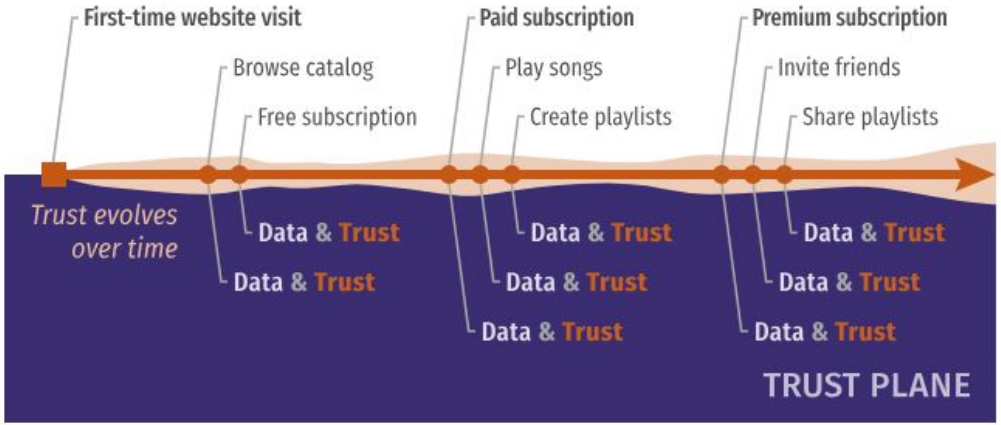

Read the full blog post by Ruben Verborgh:
https://ruben.verborgh.org/blog/2024/10/15/trust-takes-time/

Digital omnibus

General Data Protection Regulation

INSPIRE

2025

Data Governance Act

Data Act

eIDAS

EU Data Strategy

Interoperable Europe Act

Open Data Directive

AI Act

Legal compliance as an enabler of superior datatech

A new team on legaltech

By Beatriz Esteves

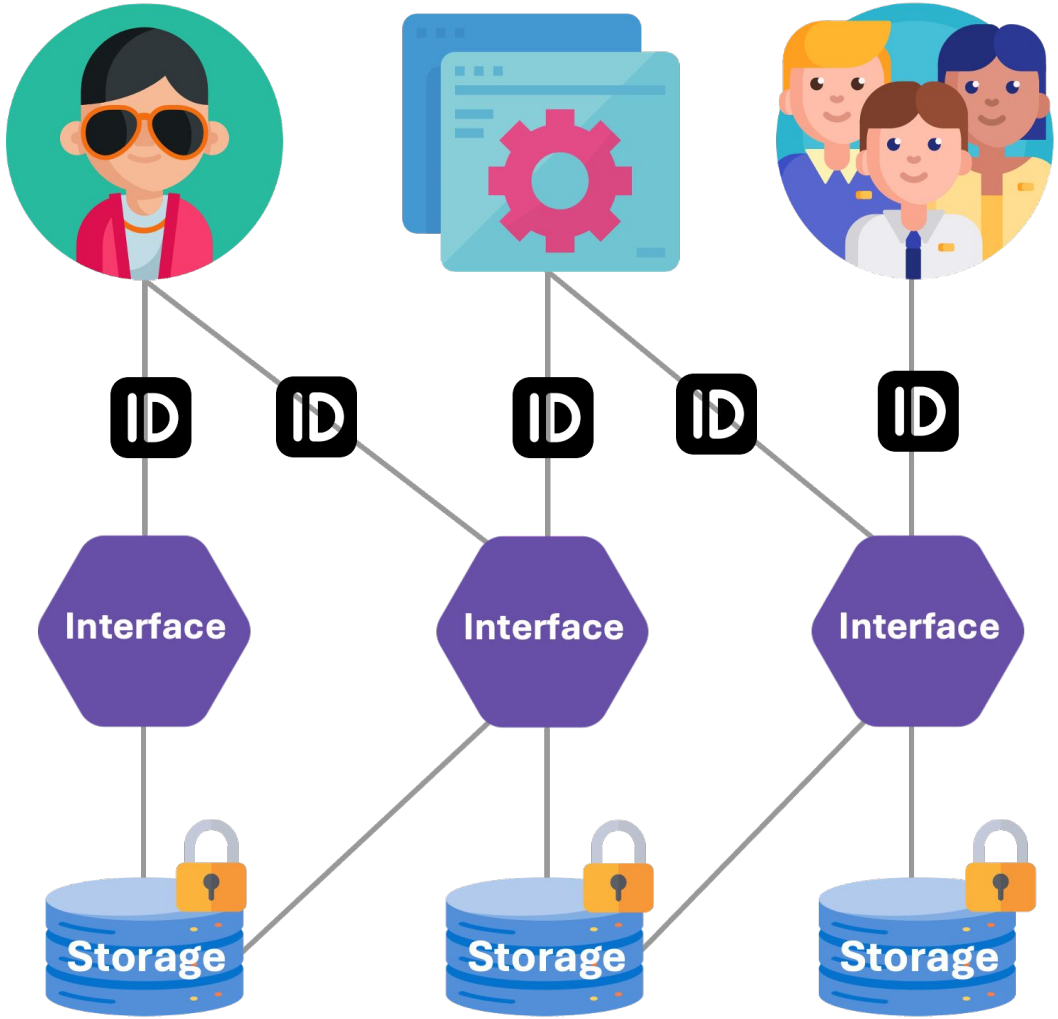
ID
ID
ID
ID
ID
Interface
Interface
Interface
Storage
Storage
Storage

Authorisation separated from storage

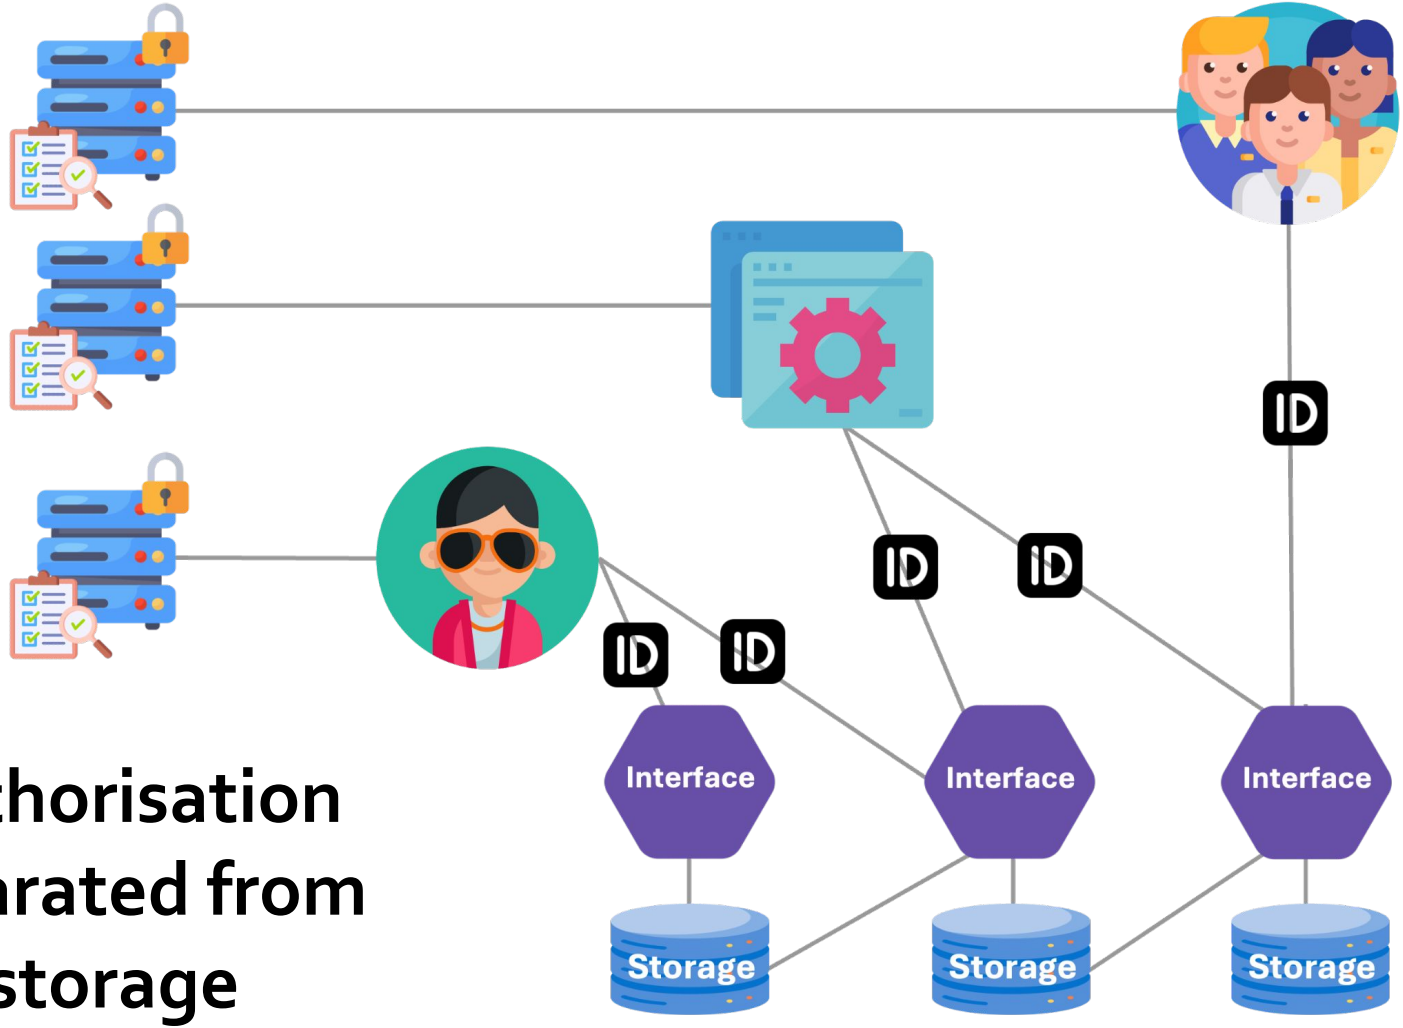

From Access Control…

Involved entities

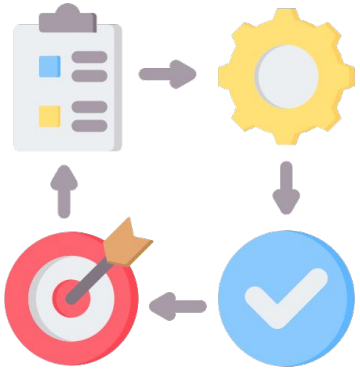

Involved action(s)

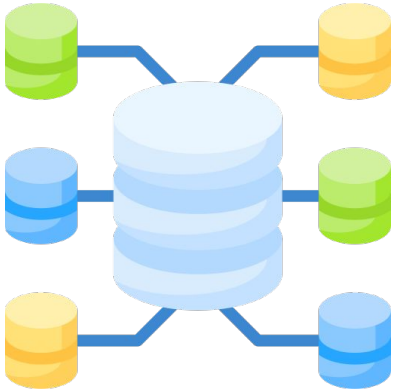

Involved resource(s)

… To Usage Control

Involved entities

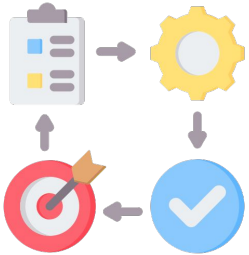

Involved action(s)

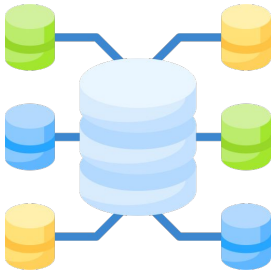

Involved resource(s)

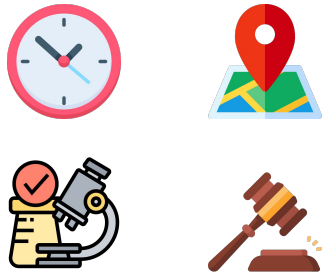

Time
Location
Purpose
Legal ground

Standards for Regulatory Compliance

ISO/IEC TS 27560:2023

Privacy technologies — Consent record information structure

HealthDCAT-AP Release 5

IEEE 7012-2025

IEEE Approved Draft Standard for Machine Readable Personal Privacy Terms

data privacy vocabulary

How to interoperably enforce ODRL policies?

Standardised

Not standardised

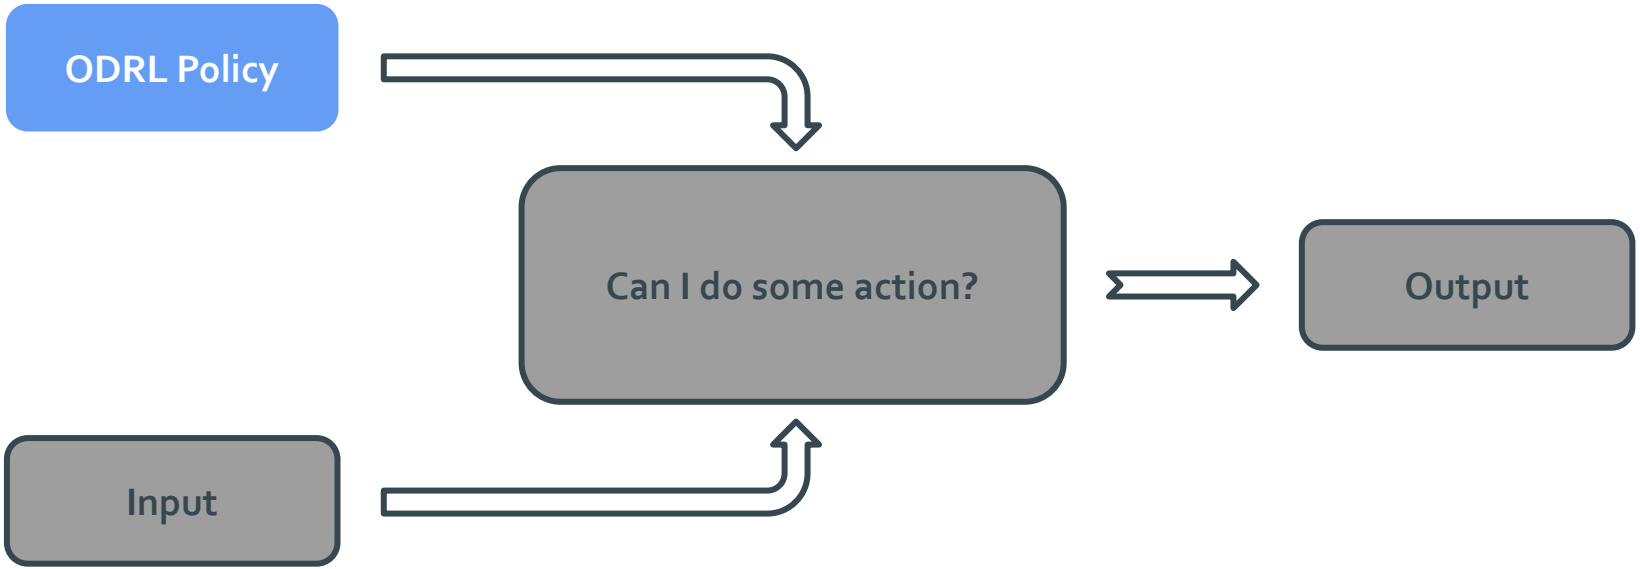

How to interoperably enforce ODRL policies?

Standardised

SolidLab

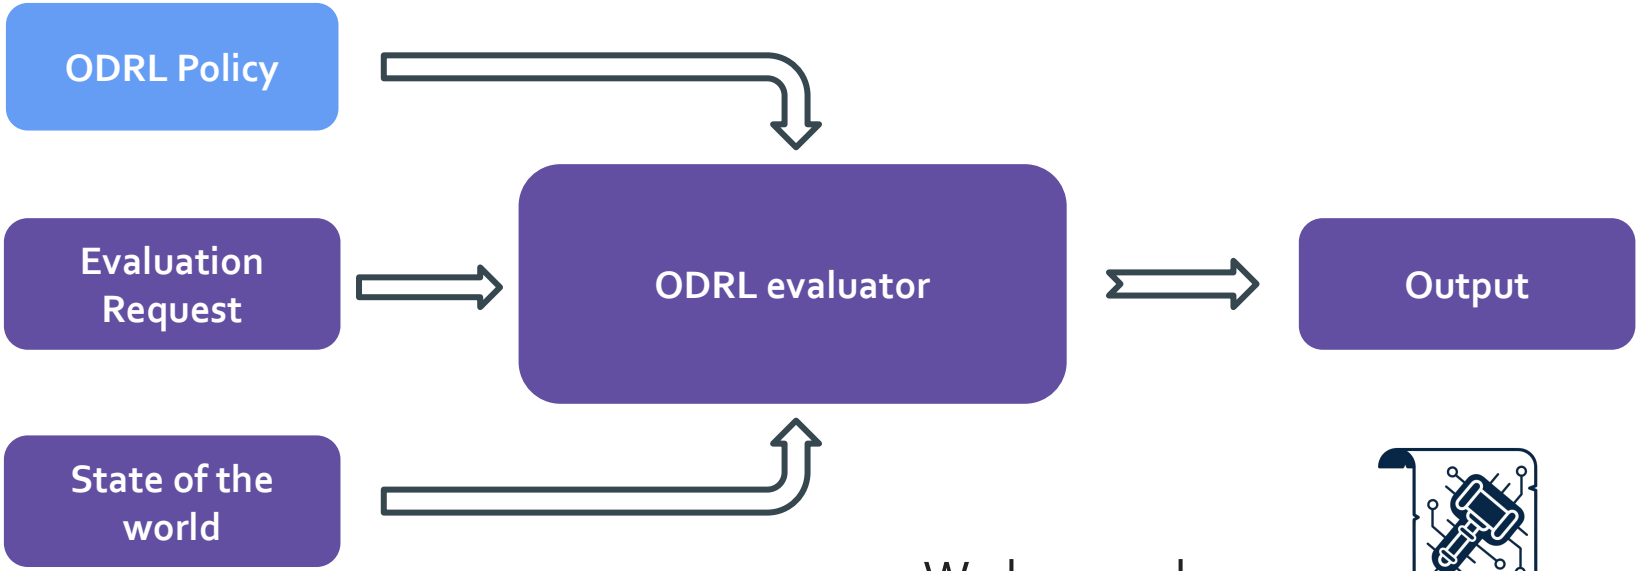

We have a demo …

More on the minute madness 😉

2026

Towards Trustflows

More about that at 16:00