

Introduction to Knowledge Graphs and Semantic Web Technologies

SPARQL Endpoints and Triple Stores

Olaf Hartig

olaf.hartig@liu.se

Acknowledgement: Some slides in this slide set are adaptations of slides of Olivier Curé (University of Paris-Est Marne la Vallée, France)


What is a SPARQL Endpoint?

- HTTP-based [Web service](#) that supports the [SPARQL protocol](#) (which is part of the family of W3C standards for SPARQL)
- Provides SPARQL-based access to a server-side dataset
 - [send](#) your SPARQL [query](#) → [receive](#) the query [result](#)
- Dataset may be
 - actual RDF data stored and maintained in a triple store*
 - a *virtual* RDF view of any other form of database (e.g., relational)

*DBMS for RDF data

RDF Virtualization Component

- Rewrites SPARQL queries to the source language (e.g., SQL)
- Relies on a mapping from the underlying database to RDF
 - R2RML is a W3C standard for mapping relational DBs to RDF
 - RML extends R2RML for other forms of source data (e.g., JSON)
- Example systems
 - Ontop <https://ontop-vkg.org/>
 - Morph-RDB <https://morph.oeg.fi.upm.es/tool/morph-rdb>
 - Sparqlify <http://aksw.org/Projects/Sparqlify.html>
 - D2R Server <http://d2rq.org/>



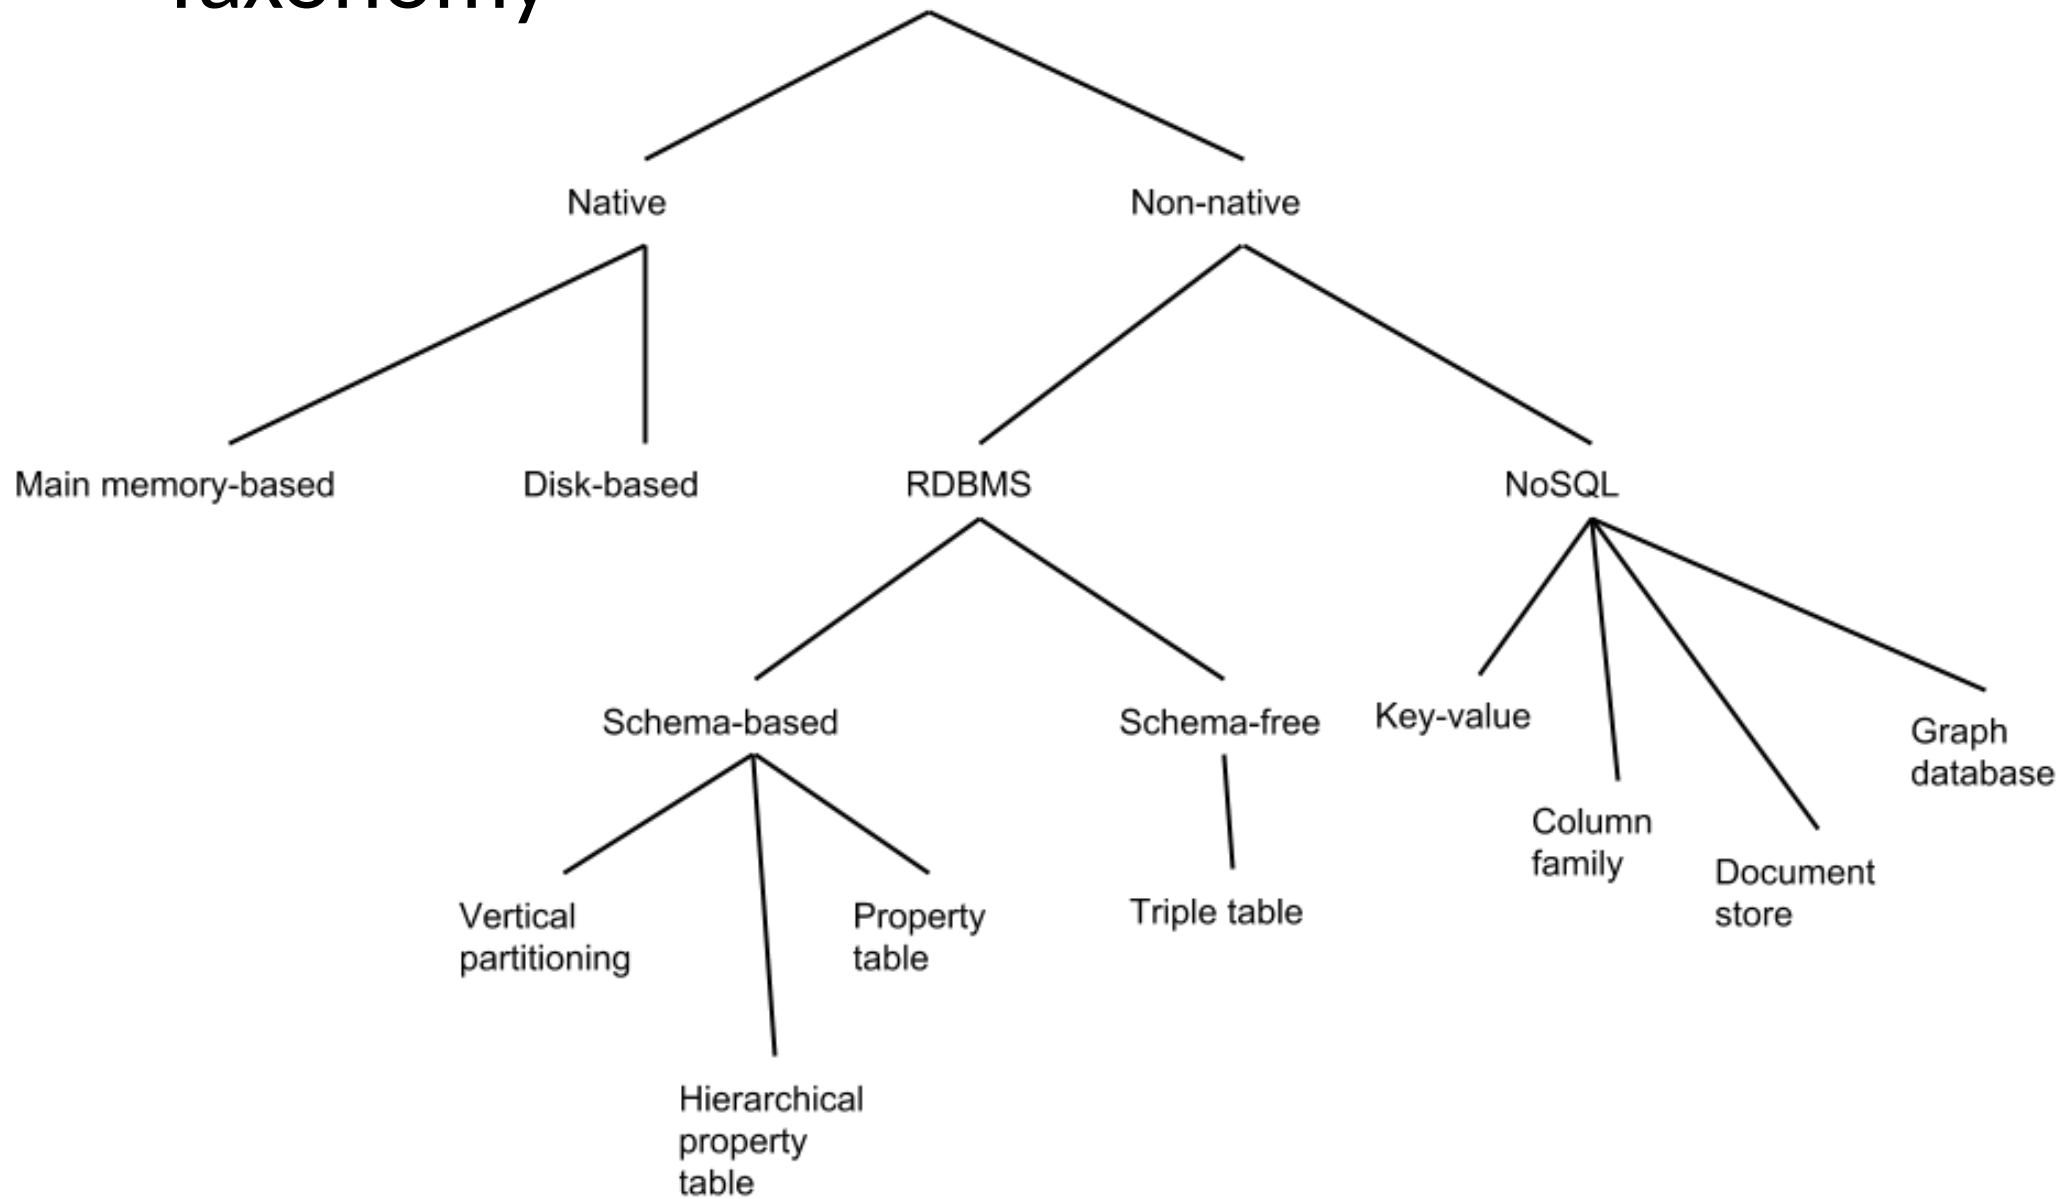
Classification of *RDF Triple Stores**

**Triple store* = DBMS for RDF data


RDF Storage

- RDF is a logical data model and, thus, does not impose any physical storage solution
- Existing triple stores are either
 - designed from scratch (“native”)
 - or
 - based on an existing DBMS
 - Relational model, e.g., PostgreSQL
 - NoSQL, e.g., Cassandra

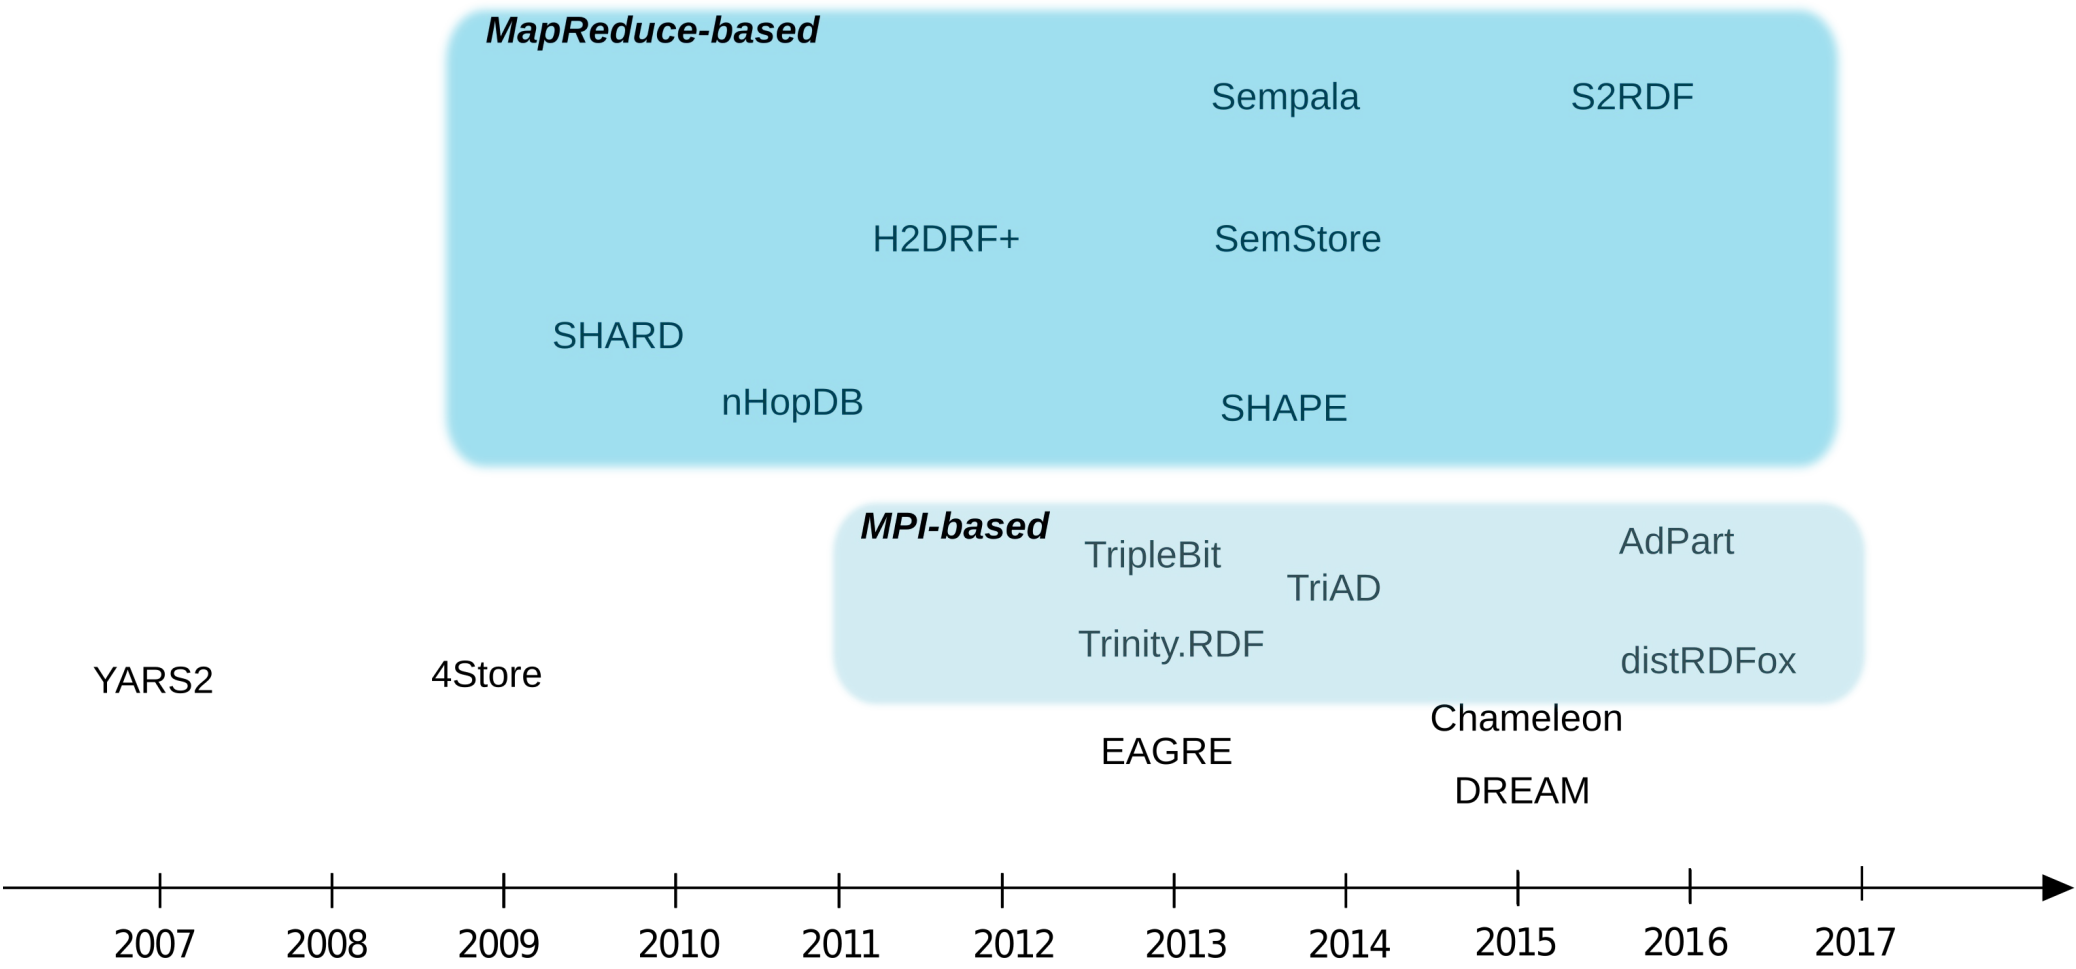
Taxonomy



Timeline of Triple Store Proposals



Prototypes of Distributed Triple Stores



by @oliviercure (feb. 2017)

Production-Ready Triple Stores

Overview

Name

Allegrograph

Blazegraph

GraphDB

MarkLogic

Oracle

Stardog

Virtuoso



AllegroGraph

Franz Inc.



GraphDB by ontotext

ORACLE[®]



Transactions with ACID Properties

Name
Allegrograph ✓
Blazegraph ✓
GraphDB ✓
MarkLogic ✓
Oracle ✓
Stardog ✓
Virtuoso ✓

- **Atomicity**: a transaction (TA) is an atomic unit of processing; it is either performed in its entirety or not performed at all
- **Consistency preservation**: a correct execution of a TA must take the DB from one consistent state to another
- **Isolation**: even if TAs are executing concurrently, they should appear to be executed in isolation; that is, their final effect should be as if each TA was executed alone from start to end
- **Durability**: once a TA is committed, its changes applied to the database must never be lost due to subsequent failure

Cluster Setups

Name

Allegrograph ✓

Blazegraph ✓

GraphDB ✓

MarkLogic ✓

Oracle ✓

Stardog ✓

Virtuoso ✓

- Replication: mostly master-slave, some master-master
- Partitioning: range, hash

Support for other Data Models (besides RDF)

Name
Allegrograph
Blazegraph ✓
GraphDB ✓
MarkLogic ✓
Oracle ✓
Stardog ✓
Virtuoso ✓

- **Relational Model** (with SQL)
 - Virtuoso, Oracle
- **XML** (with XQuery)
 - MarkLogic, Virtuoso
- **Document Model**
 - MarkLogic
- **Property Graphs** (with Gremlin)
 - Blazegraph, GraphDB, Stardog

Licenses

Name

Allegrograph

Blazegraph

GraphDB

MarkLogic

Oracle

Stardog

Virtuoso

- Most of these systems have a free-to-use edition, some even have a feature-limited free software version (open source)
- All have commercial editions

Full-Text Search Support

Name	Full-text search
Allegrograph	Integrated + Solr
Blazegraph	Integrated + Solr
GraphDB	Integrated + Solr + ElasticSearch (enterp.)
MarkLogic	Integrated
Oracle	Integrated
Stardog	Integrated + Lucene
Virtuoso	Integrated

Cloud Readiness

Name	Full-text search	Cloud-ready
Allegrograph	Integrated + Solr	AMI
Blazegraph	Integrated + Solr	AMI
GraphDB	Integrated + Solr + ElasticSearch (enterp.)	AMI
MarkLogic	Integrated	AMI
Oracle	Integrated	
Stardog	Integrated + Lucene	AMI
Virtuoso	Integrated	AMI

AMI: Amazon Machine Image

Other, cloud-native options:



Automated Reasoning in Triple Stores

Automated Reasoning?

- Definition of properties and classes in RDF vocabularies based on an ontology language such as RDFS or OWL (covered later)
- Allows for inferring additional data from existing data that uses these properties and classes
- Example:
 - assume `foaf:knows` is defined as a symmetric relationship
 - then, given the RDF triple (`ex:olaf`, `foaf:knows`, `ex:eva`), we can infer the triple (`ex:eva`, `foaf:knows`, `ex:olaf`)
- Example:
 - assume `ex:Man` is defined as a subclass of `foaf:Person`
 - then, given the triple (`ex:olaf`, `rdf:type`, `ex:Man`), we can infer (`ex:olaf`, `rdf:type`, `foaf:Person`)

Approach 1: *Materialization*

a.k.a. *forward reasoning or closure*

- Idea: make explicit all inferences in the triple store
- Pros:
 - efficient query processing (no reasoning at query runtime)
- Cons:
 - slow data loading
 - data volume expansion
 - tricky update management

Approach 2: Query Rewriting

a.k.a. *backward reasoning* or *query reformulation*

- Idea: reformulate the original query such that all answers can be retrieved
- Pros:
 - no preprocessing overhead
 - no expansion of stored data volume
 - easy update management
- Cons:
 - slower query processing due to cost of reasoning at runtime

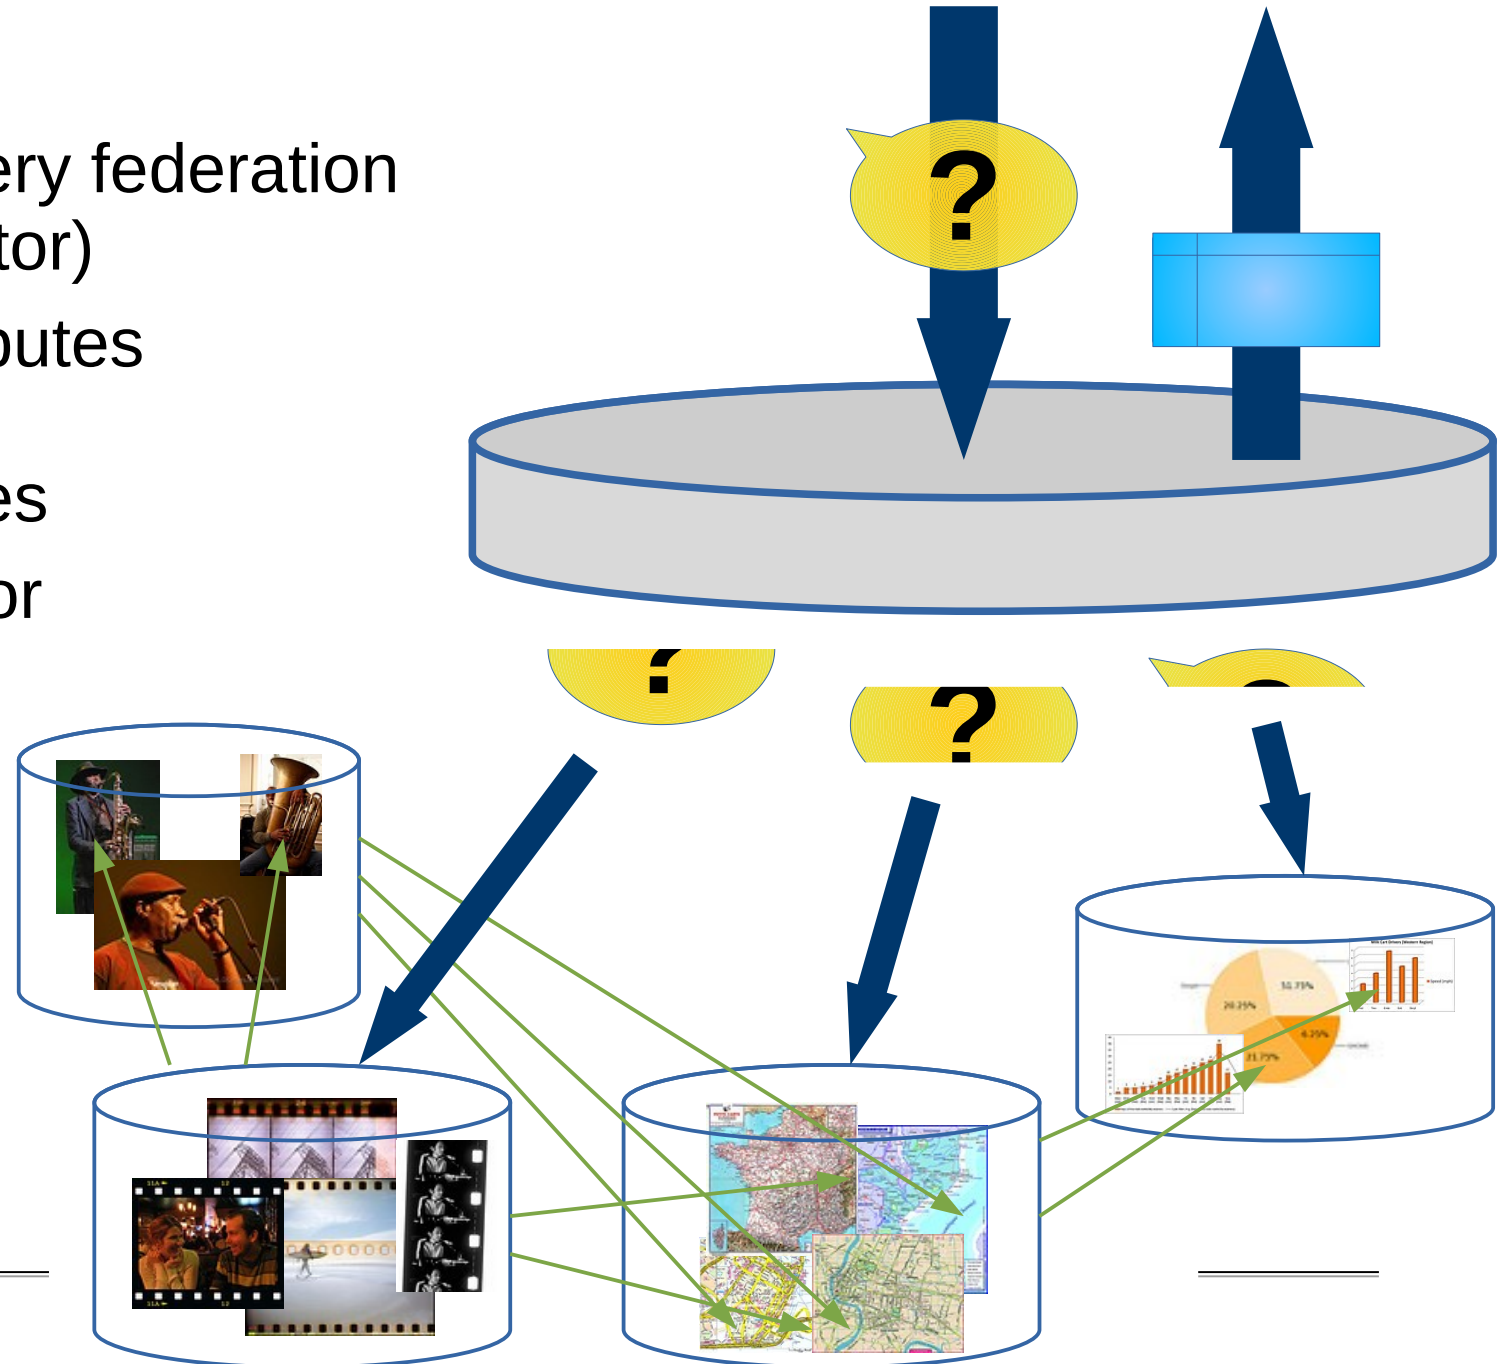
Reasoning in the Production-Ready Systems

Triple store	Materialization	Query rewriting
Allegrograph	OWLRL	RDFS++, Prolog
Blazegraph	RDFS, OWL Lite	
GraphDB	RDFS, OWL Horst, OWLRL, OWLQL	
MarkLogic		RDFS, RDFS++, OWL Horst
Oracle	RDFS, OWLRL, OWLQL	
Stardog	All OWL2	
Virtuoso		RDFS++

Federated Query Processing

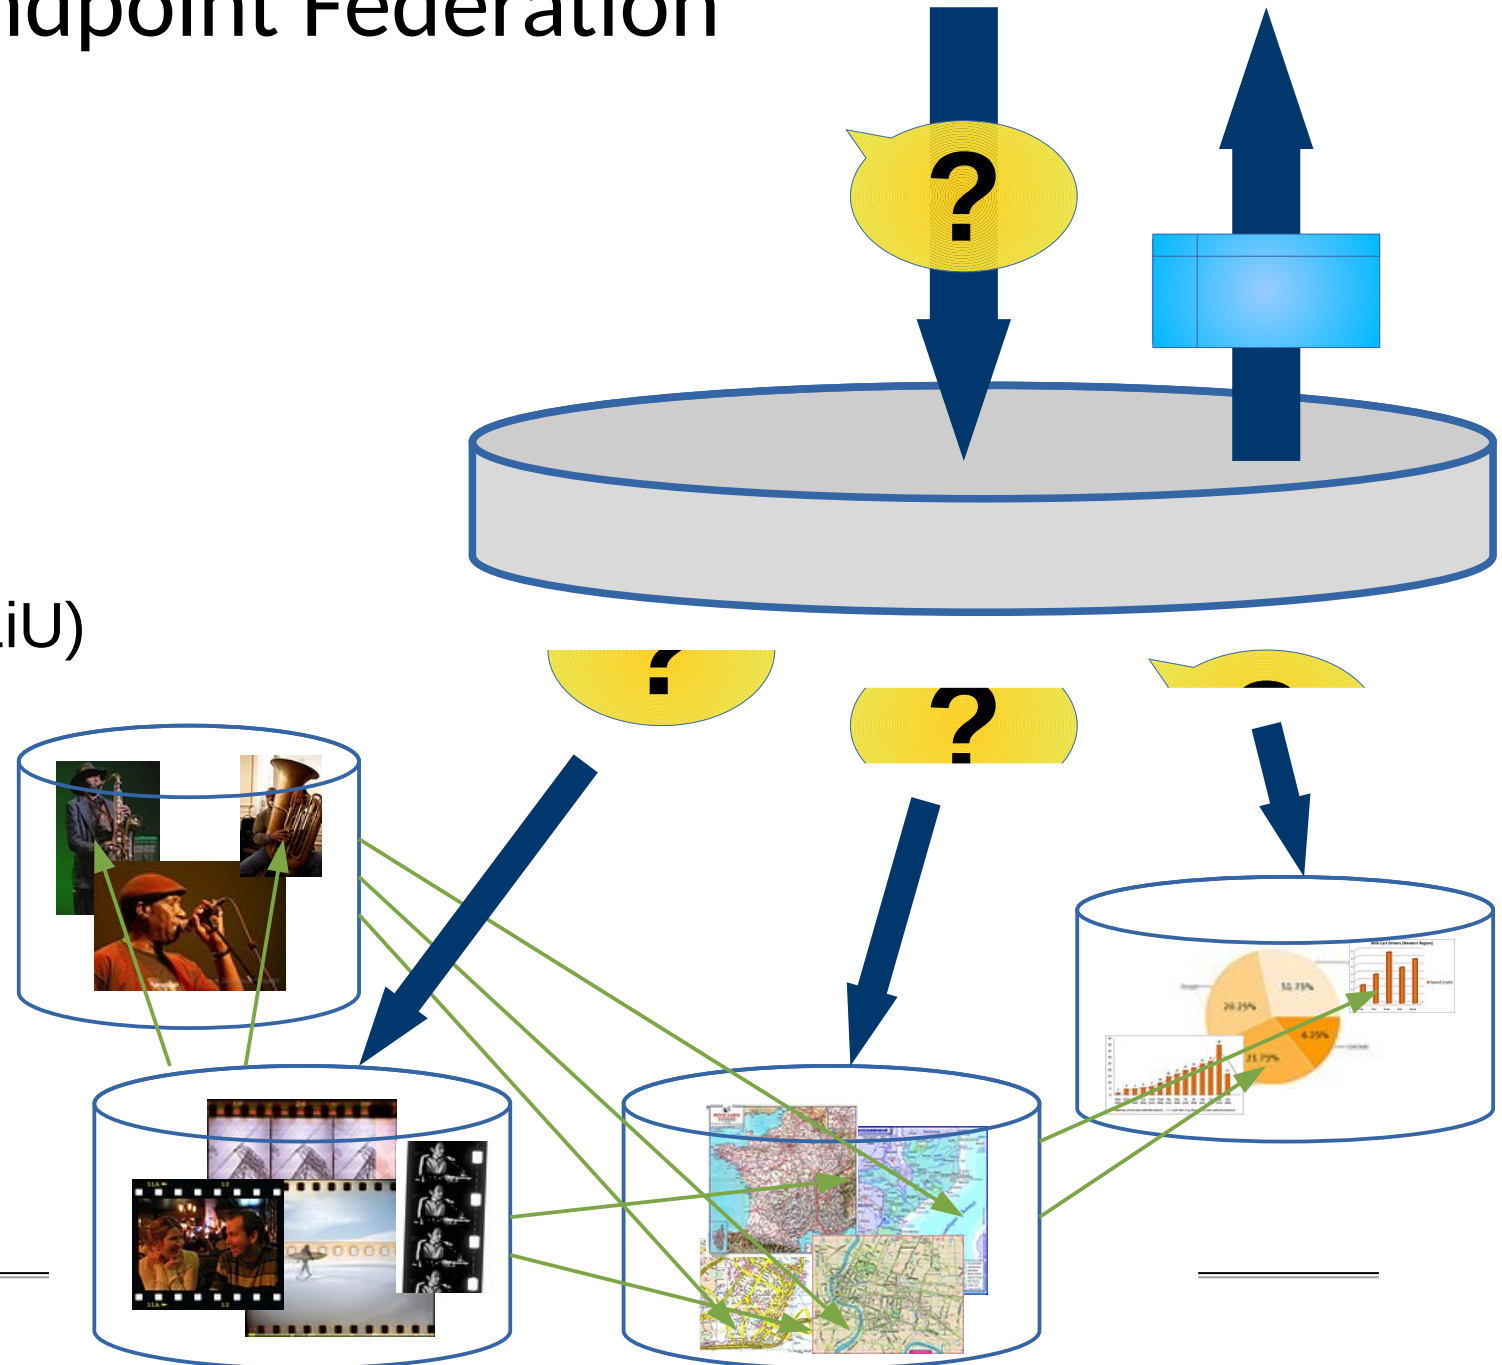
Idea

- Querying a query federation service (mediator)
- Mediator distributes sub-queries to relevant sources
- Finally, mediator combines sub-results



SPARQL Endpoint Federation

- Prototypes:
 - FedX
 - SPLENDID
 - ANAPSID
 - CostFed
 - HeFQUIN (@LiU)
 - etc.



SPARQL 1.1 Federation Extension

- SERVICE pattern in SPARQL 1.1
- Explicitly specify query patterns whose execution must be distributed to a remote SPARQL endpoint

```
SELECT ?v ?ve WHERE
{
  SERVICE <http://volcanos.example.org/query> {
    ?v rdf:type umbel-sc:Volcano ;
    p:location dbpedia:Italy .
  }
  SERVICE <http://volcano-eruptions.org/sparql> {
    ?v p:lastEruption ?ve .
  }
}
```

www.liu.se