## The Semantic Web & Atomic Data

Joep Meindertsma, Declarative Amsterdam 2022-11-08

Slides: tiny.cc/atomicdata















# How did we get here?

- Developers have to invent their own API, because:
- The web is not machine readable
- The web is not read/write

(b<sub>ut it was tho)</sub>





# Is the web not machine readable?

- HTML pages are documents, not data
- Tim Berners-Lee wanted to fix this, too...
- ... with the **Semantic web** / RDF / Linked Data

### SCIENTIFIC AMERICAN.com

 Get Digital.
 > 12 digital issues
 > SCIENTIFIC
 DIGITAL

 Get More Science
 > Unlimited access 24 x7
 SUBSCRIBE TODAY >

#### May 17, 2001

The Semantic Web

#### A new form of Web content that is meaningful to computers will unleash a revolution of new possibilities

#### By Tim Berners-Lee, James Hendler and Ora Lassila

The entertainment system was belting out the Beatles' "We Can Work It Out" when the phone rang. When Pete answered, his phone turned the sound down by sending a message to all the other *local* devices that had a *volume control*. His sister, Lucy, was on the line from the doctor's office: "Mom needs to see a specialist and then has to have a series of physical therapy sessions. Biweekly or something. I'm going to have my agent set up the appointments." Pete immediately agreed to share the chauffeuring.



BY MIGUEL SALMERON

At the doctor's office, Lucy instructed her Semantic Web agent through her handheld Web browser. The agent promptly retrieved information about Mom's *prescribed treatment* from the doctor's agent, looked up several lists of *providers*, and checked for the ones *in-plan* for Mom's insurance within a 20-*mile radius* of her *home* and with a *rating* of *excellent* or *very good* on trusted rating services. It then began trying to find a match between available *appointment times* (supplied by the agents of individual providers through their Web sites) and Pete's and Lucy's busy schedules. (The emphasized keywords indicate terms whose semantics, or meaning, were defined for the agent through the Semantic Web.)

In a few minutes the agent presented them with a plan. Pete didn't like it—University Hospital was all the way across town from Mom's place, and he'd be driving back in the middle of rush hour. He set his own agent to redo the search with stricter preferences about *location* and *time*. Lucy's agent, having *complete trust* in Pete's agent in the context of the present task, automatically assisted by supplying access certificates and tod through

shortcuts to the data it had already sorted through.

Almost instantly the new plan was presented: a much closer clinic and earlier times—but there were two warning notes. First, Pete would have to reschedule a couple of his *less important* appointments. He checked



what they ware not a problem. The other was compating about the incurance company's list failing to include this provider under physical therapiets. "Service type and incurance plan status securely varified

### SCIENTIFIC AMERICAN.com

### Get Digital. Get More Science

> 12 digital issues
 > Archive with 150+ issues
 > Unlimited access 24 x7



May 17, 2001

**The Semantic Web** 

A new form of Web content that is meaningful to computers will unleash a revolution of new possibilities

#### By Tim Berners-Lee, James Hendler and Ora Lassila

The entertainment system was belting out the Beatles' "We Can Work It Out" when the phone rang. When Pete answered, his phone turned the sound down by sending a message to all the other *local* devices that had a *volume control*. His sister, Lucy, was on the line from the doctor's office: "Mom needs to see a specialist and then has to have a series of physical therapy sessions. Biweekly or something. I'm going to have my agent set up the appointments." Pete immediately agreed to share the chauffeuring.



At the doctor's office, Lucy instructed her Semantic Web agent through her handheld Web browser. The agent promptly retrieved information about Mom's prescribed treatment from the doctor's agent, looked up several lists of providers, and checked for the ones *in-plan* for Mom's insurance within a 20-mile radius of her home and with a rating of excellent or very good on trusted rating services. It then began trying to find a match between available appointment times (supplied by the agents of individual providers through their Web sites) and Pete's and Lucy's busy schedules. (The emphasized keywords indicate terms whose semantics, or meaning, were defined for the agent through the Semantic Web.)

n a few minutes the agent presented them with a plan. Pete didn't like —University Hospital was all the way across town from Mom's place, and e'd be driving back in the middle of rush hour. He set his own agent to ado the search with stricter preferences about *location* and *time*. Lucy's gent, having *complete trust* in Pete's agent in the context of the present ask, automatically assisted by supplying access certificates and eit through.

task

Almost instantly the new plan was presented: a much closer clinic and earlier times—but there version notes. First, Pate would have to reschedule a couple of his less important appointment



what they ware \_\_\_\_\_\_ and a problem. The other was compatible about the incurance company's list failing to include this provider under physical therapists: "Service two and incurance plan status securely varifier





## So... we made it?

## Not really... Monopoly, not interoperability

# URLs can make data linked and interoperable

- Why don't we use them for, like, everything?
- URLs should represent **things**, not (HTML) strings!
- So Tim Berners-Lee came up with RDF!

# Let's upgrade a piece of information from sentence to RDF

## Sentence

"Joep is born in Baarn on the 20th on january, 1991."

Not machine readable

## Table

Name	City	Birth Date
Joep	Baarn	02-01-1991

**1.** Machine readable (query, sort, filter, serialize)

#### 2. Still ambiguous

a. Who's Joep? Where's Baarn? How to interpret date?

# Links

<u>Name</u>	<b>BirthPlace</b>	<b>BirthDate</b>
Joep	Baarn	1991-01-20

- 1. Machine readable
- 2. Unambiguous (even in a global context, because URLs are global)
- **3. Browseable** (we can open links and learn more)
- 4. Standardized (predicates can specify how they should be used)

# **RDF Triples**subjectpredicateobject

<<u>https://example.com/joep</u>> <<u>https://schema.org/birthDate</u>> "1991-01-20"xsd:date <<u>https://example.com/joep</u>> <<u>https://schema.org/birthPlace</u>> <<u>https://example.com/baarn</u>>





# We (Argu) went all-in on RDF

Because we believe in the merits of the semantic web

argu

Over

Blog

Help

### Alle functionaliteiten in huis

Argu is een krachtig en flexibel platform waarmee je snel online bent. Mis je iets? Dan bouwen wij precies dat wat jij nodig hebt.

#### Gebruiksvriendelijk Toegankelijke participatie voor iedereen

#### Eenvoudig registratieproces

Gebruikers kunnen alle content zien en hoeven pas te registreren nádat ze ergens op hebben gestemd. Liever een besloten omgeving? Dan kan je mensen uitnodigen via een veilige stemcode.

- > Gebruiksvriendelijk platform dat op alle apparaten werkt
- > Look and feel aanpassen naar huisstijl
- > Voldoet aan WCAG standaarden



# Things we had to do

- 1. Make RDF work with **React (link-redux)**
- 2. Make RDF work with **Rails (<u>linked\_rails</u>)**
- 3. Create a lot of **ontologies**
- 4. Invent a new RDF serialization format (hextuples)
- 5. Dynamic **forms** in RDF (powered by SHACL)
- 6. Full-text search

Read more @ ontola.io/blog/full-stack-linked-data/

## And the more we built with RDF...

...the less hopeful I became

## Why is the cloud not mostly Linked Data?

# RDF is not attractive enough for developers

## Why is the cloud not mostly Linked Data?

JSON > RDF

# What makes RDF less attractive than JSON?

It's not *just* the lack of tools & libraries

# 1. No native arrays

- Collections, Sequences, Bags... It's confusing!
- Linked lists in RDF are fundamentally complex
- Works different in all serialization formats

Read more @ <u>ontola.io/blog/ordered-data-in-rdf/</u>

## 2. No subject-predicate uniqueness

- A subject can have *many* triples with the *same* predicate, each with different datatypes
- Hard to store in Maps (and maps are easy and fast!)
- Hard to query => bad Developer Experience (compared with JSON: thing.property)

## **3. Blank nodes**

- Necessary, but are hard to store & query
- Name collisions
- A high cost for clients to deal with

## 4. Named graphs

- Add an extra identifier that is required to filter and find a set of triples
- No consensus on best practice, lots of hacky usages

## 5. It's hard to select a specific value

1 <example:joep> <schema:birthDate> "1991-01-20"^^xsd:date <example:someNamedGraph> 2 <example:joep> <schema:birthDate> <example:birthDateObject> <example:someOtherNamedGraph> 3 <example:joep> <schema:birthDate> "20th of januari 1991"@en <example:someNamedGraph> 4 <example:joep> <schema:birthDate> "20 januari 1991"@nl <example:someNamedGraph> 5 <example:joep> <schema:birthDate> "2000-02-30"^^xsd:date <example:someNamedGraph>

wish: joep.birthDate = Date

reality: someNamedGraph.joep.birthdate.filter(b.datatype == date)[0]

# 6. Lack of consensus on usage

- Many resources (most notably properties) do not resolve (data dumps, SPARQL / TPF endpoints, but not as resources!)
- Many **different serialization formats** used no requirements on which is needed, which leads to fat clients with many parsers
- Highly **fragmented documentation** with a high focus on semantics instead of pragmatics

## 7. No type safety

- Any predicate can have **any datatype** (string / bool / etc.)
- SHACL / SHEX are **optional** (and kind of complicated)

## What I'd like to have...

- **Decentralized** nature of RDF (lots of URLs)
- Easy for devs like JSON
- **Type-safe** like typescript

## Take a (very) strict subset of RDF!

# Atomic Data



### Atomic Data is

- Resolvable RDF
- Idiomatic JSON
- Fully Type-safe

## How?

### With strict **Properties**!

# Atomic **Properties**

"@id": "https://atomicdata.dev/properties/description",

- "https://atomicdata.dev/properties/datatype": "https://atomicdata.dev/datatypes/markdown",
- "https://atomicdata.dev/properties/description": "A textual description of something.",
- https://atomicdata.dev/properties/isA": [

https://atomicdata.dev/classes/Property"

"https://atomicdata.dev/properties/shortname": "description"

• **JSON-AD** as serialization

··],

- **@id** must resolve to the actual JSON-AD resource
- Every **key** is a **Property**
- Properties have **datatypes**, triples can't set their own!
- **<u>Shortnames</u>** can be used for easy mapping (to simple JSON)
- **Descriptions** can be used for semantic meaning
- See <u>docs.atomicdata.dev</u>

## Atomic <u>Classes</u>

- E.g. **BlogPost** or **Agent**
- Describe which Properties are *required* and *optional*
- Can be used to **validate data** and **show Forms**
- Difference from RDF: schema tightly coupled to Class

## **Atomic Data Core**

JSON-AD, Properties, Datatypes, Classes

## **Atomic Data Extended**

Authorization, authentication, collections, more

# Static data is relatively easy

But data *changes over time*, how can we deal with that?

# What can we gain by standardizing *Changes*?

- Versioning, history, global undo
- Verifiability for checking authorship, auditing
- **P2P** sharing
- So basically **Git**, but for structured data
- A read-write web! Collaboration everywhere

## **Atomic Commits**

- <u>Commits</u> (changes) are resources, too, with their own URLs
- **PK Cryptography** is used for signing every single change
- So every change becomes **reversible** and **verifiable**
- And can therefore be shared **peer to peer**

## **Atomic Agents**

- Agents are users with their own URLs
- Agents have a **public key** used for signing Commits
- They are also used for read / write **authorization**
- Authentication by **<u>signing</u>** every HTTP GET request

## And More

- <u>Endpoints</u> (machine readable API docs)
- <u>Collections</u> (queries for sorting, filtering, pages)
- <u>Invitations</u> (easy sign up + redirect)
- <u>Websockets</u> (real time collaboration)
- File upload / download

# A spec needs an implementation

- **<u>Atomic-server</u>** (graph database, written in Rust)
  - Fast: <1ms responses
  - 14mb binary, no runtime dependencies
  - Full-text search, indexed queries, sorting, filtering, pagination
  - Authorization, authentication, attachments, versioning
- **<u>Atomic-data-browser</u>** (GUI for viewing / editing)
  - Browse data, dynamic forms, tables, search, endpoints, collaborative documents

## Time for a demo!

Follow along at atomicdata.dev

### **Atomic Data**

## **Getting started**

- <u>@tomic/lib</u> for js, typescript, node projects
- <u>@tomic/react</u> for web GUIs (<u>codesandbox template</u>)
- <u>atomic lib</u> (rust library, powers server)
- Everything mentioned is MIT licensed, so fully free & open source
- Check out the <u>Discord</u> and <u>Docs</u>!

## **Questions?**

Joep Meindertsma, Declarative Amsterdam 2022-11-08

Slides: tiny.cc/atomicdata