

Future Storage Layer

Roman Chyla and the ADS Team

ADS Users Group Meeting, 15-16 Nov. 2021





The retreat of Russia.
Swebach Bernard
Edouard (1800-1870)
The retreat of Russia in
1812 (oil on canvas 1;
26 x 1; 93) 1838
Museum of Fine Arts
Besancon.

<https://www.gettyimages.com/detail/news-photo/swebach-bernard-edouard-the-retreat-of-russia-in-1812-1838-news-photo/1048362648>



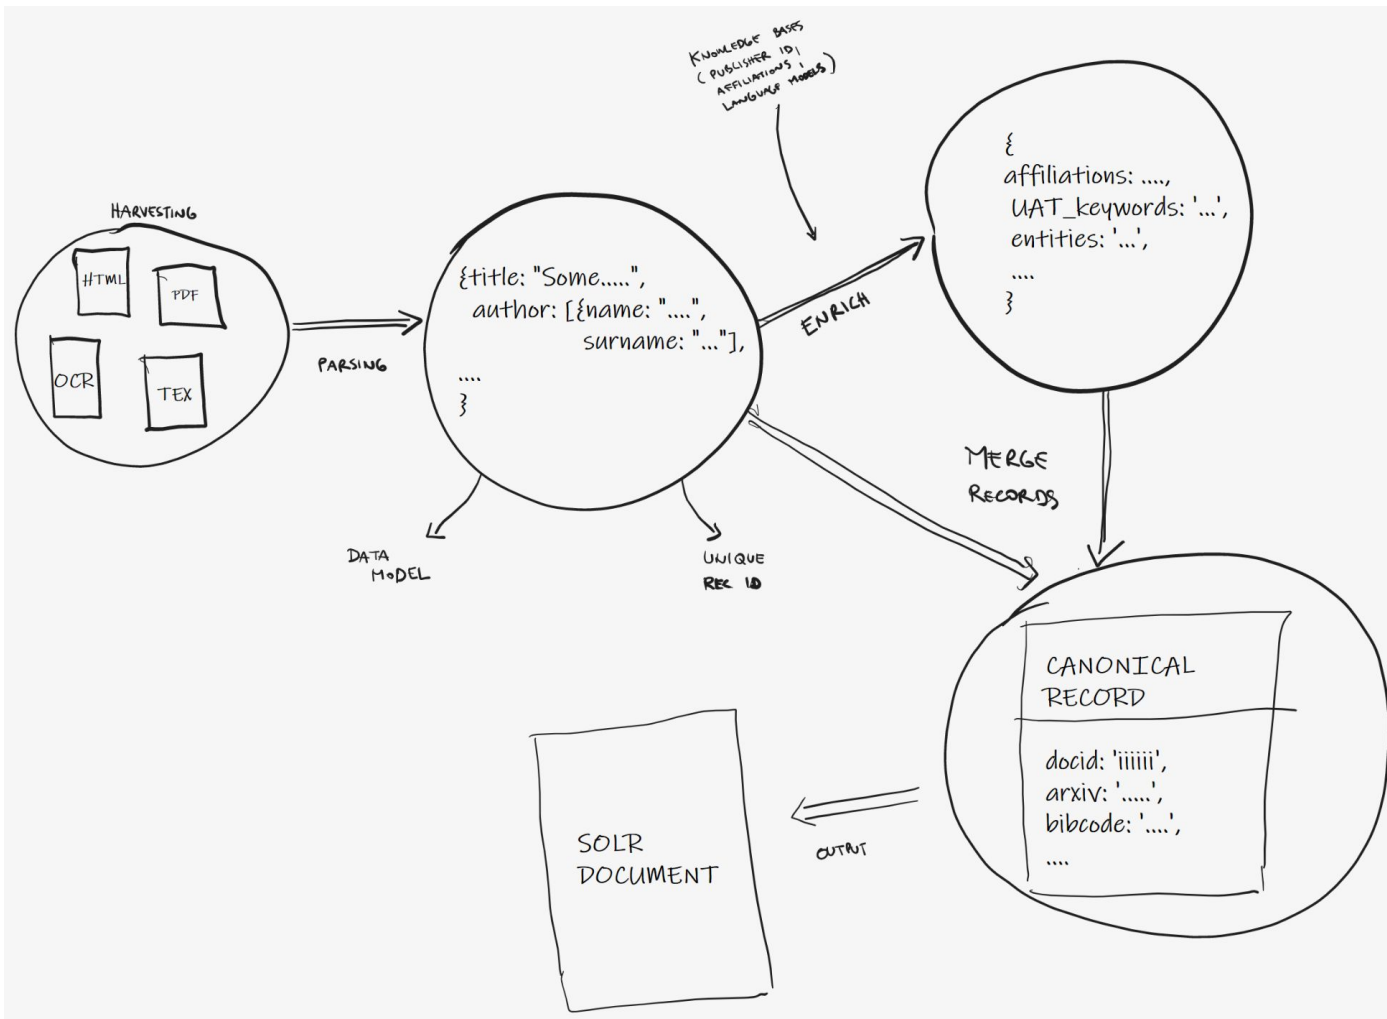
*March of Suvorov
through the Alps by
Vasily Surikov.*

The problem

- Technologically newer parts of ADS
 - lots of attention, gradual improvements
 - steady and stable
- Content curated by ADS
 - linear growth
 - 2M docs/5 yrs
 - doubled in 8(?) years
- Possible expansion
 - At least 10x more of **everything**
 - will also accentuate deficiencies
 - however: speed alone is not what informs the design
 - *Massive change in quantity is change in **quality***

The ~~problem~~ exciting opportunity

- Connect all separate data
 - One road connecting two cities, that's something
 - But network of roads: that changes how people think
- Remove (mental) roadblocks
 - Access to computation/data via API
 - First internally, in next phase to end-users
 - Researchers will be able to run experiments on data ADS curates
- Possibly redesign API
 - Pipe dream for quite some time



Reqs

- Access to **all** data ADS curates
- Via unified interface (API)
 - internally, not exposed to public
- **Highly available + scalable + failure tolerant**
 - 144TB (for data - not counting logs, telemetry)

How?

- Exploration ongoing
 - number of unknowns decreasing
 - evaluated multiple distributed storage technologies
 - Cassandra
 - Hadoop/HBase
 - Yugabyte (and Postgres)
 - need to 'marry' storage and computation
 - Hadoop vs Spark vs Kafka streaming...
 - schema vs schemaless
 - thrift/protocol buffers and polyglot bindings (unified access)
 - need to assemble cost projections
 - bare metal
 - cloud and vendor products

And then put it together...

