

An Observatory for Internet Path Transparency

Mirja Kühlewind and Brian Trammell, ETH Zürich

Motivation: Efforts to design protocols and protocol extensions to work in the face of rampant middlebox interference in the Internet need quantitative data about the **prevalence** of different types of interference.

Problem: How to make this data available to networking research, engineering and operations communities?

Approach: build an open repository around a common data model for **comparable** and **repeatable** measurement of **observations** of these phenomena.


Observation: an assertion that at a given **time** along a given **path**, a given **condition** held:

$$O = \{t, p, c\}$$

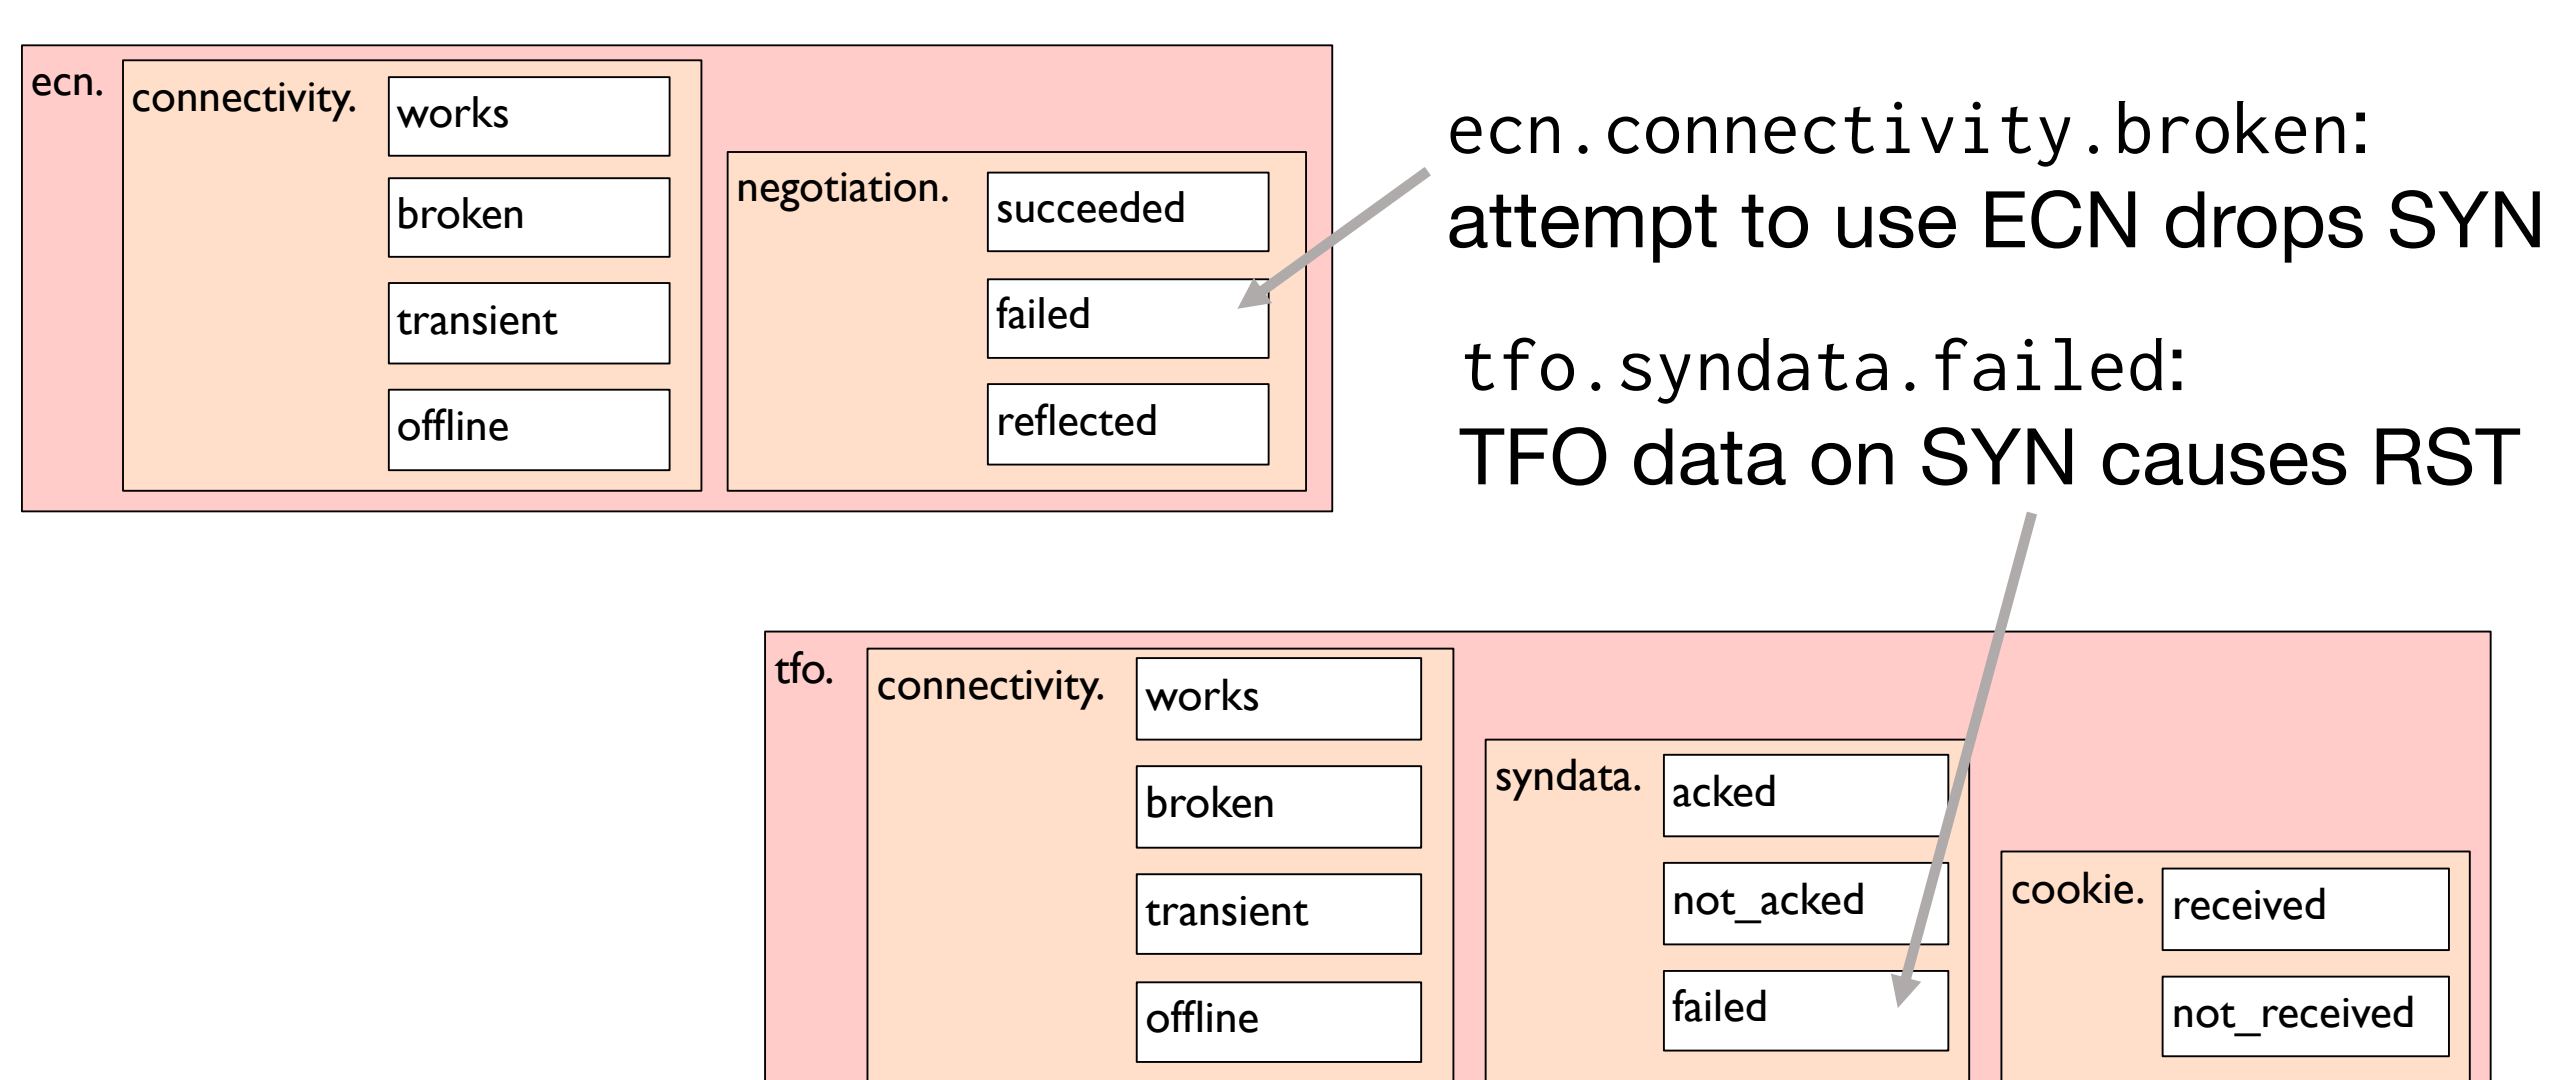
Path: sequence of elements (addresses, prefixes, BGP ASN, pseudonyms), allowing multi-resolution storage and correlation with topology (e.g. Tracebox)

$[IP_0, *, IP_8] \rightarrow [IP_0, IP_1, AS_3, *, AS_5, IP_6, *, IP_8]$

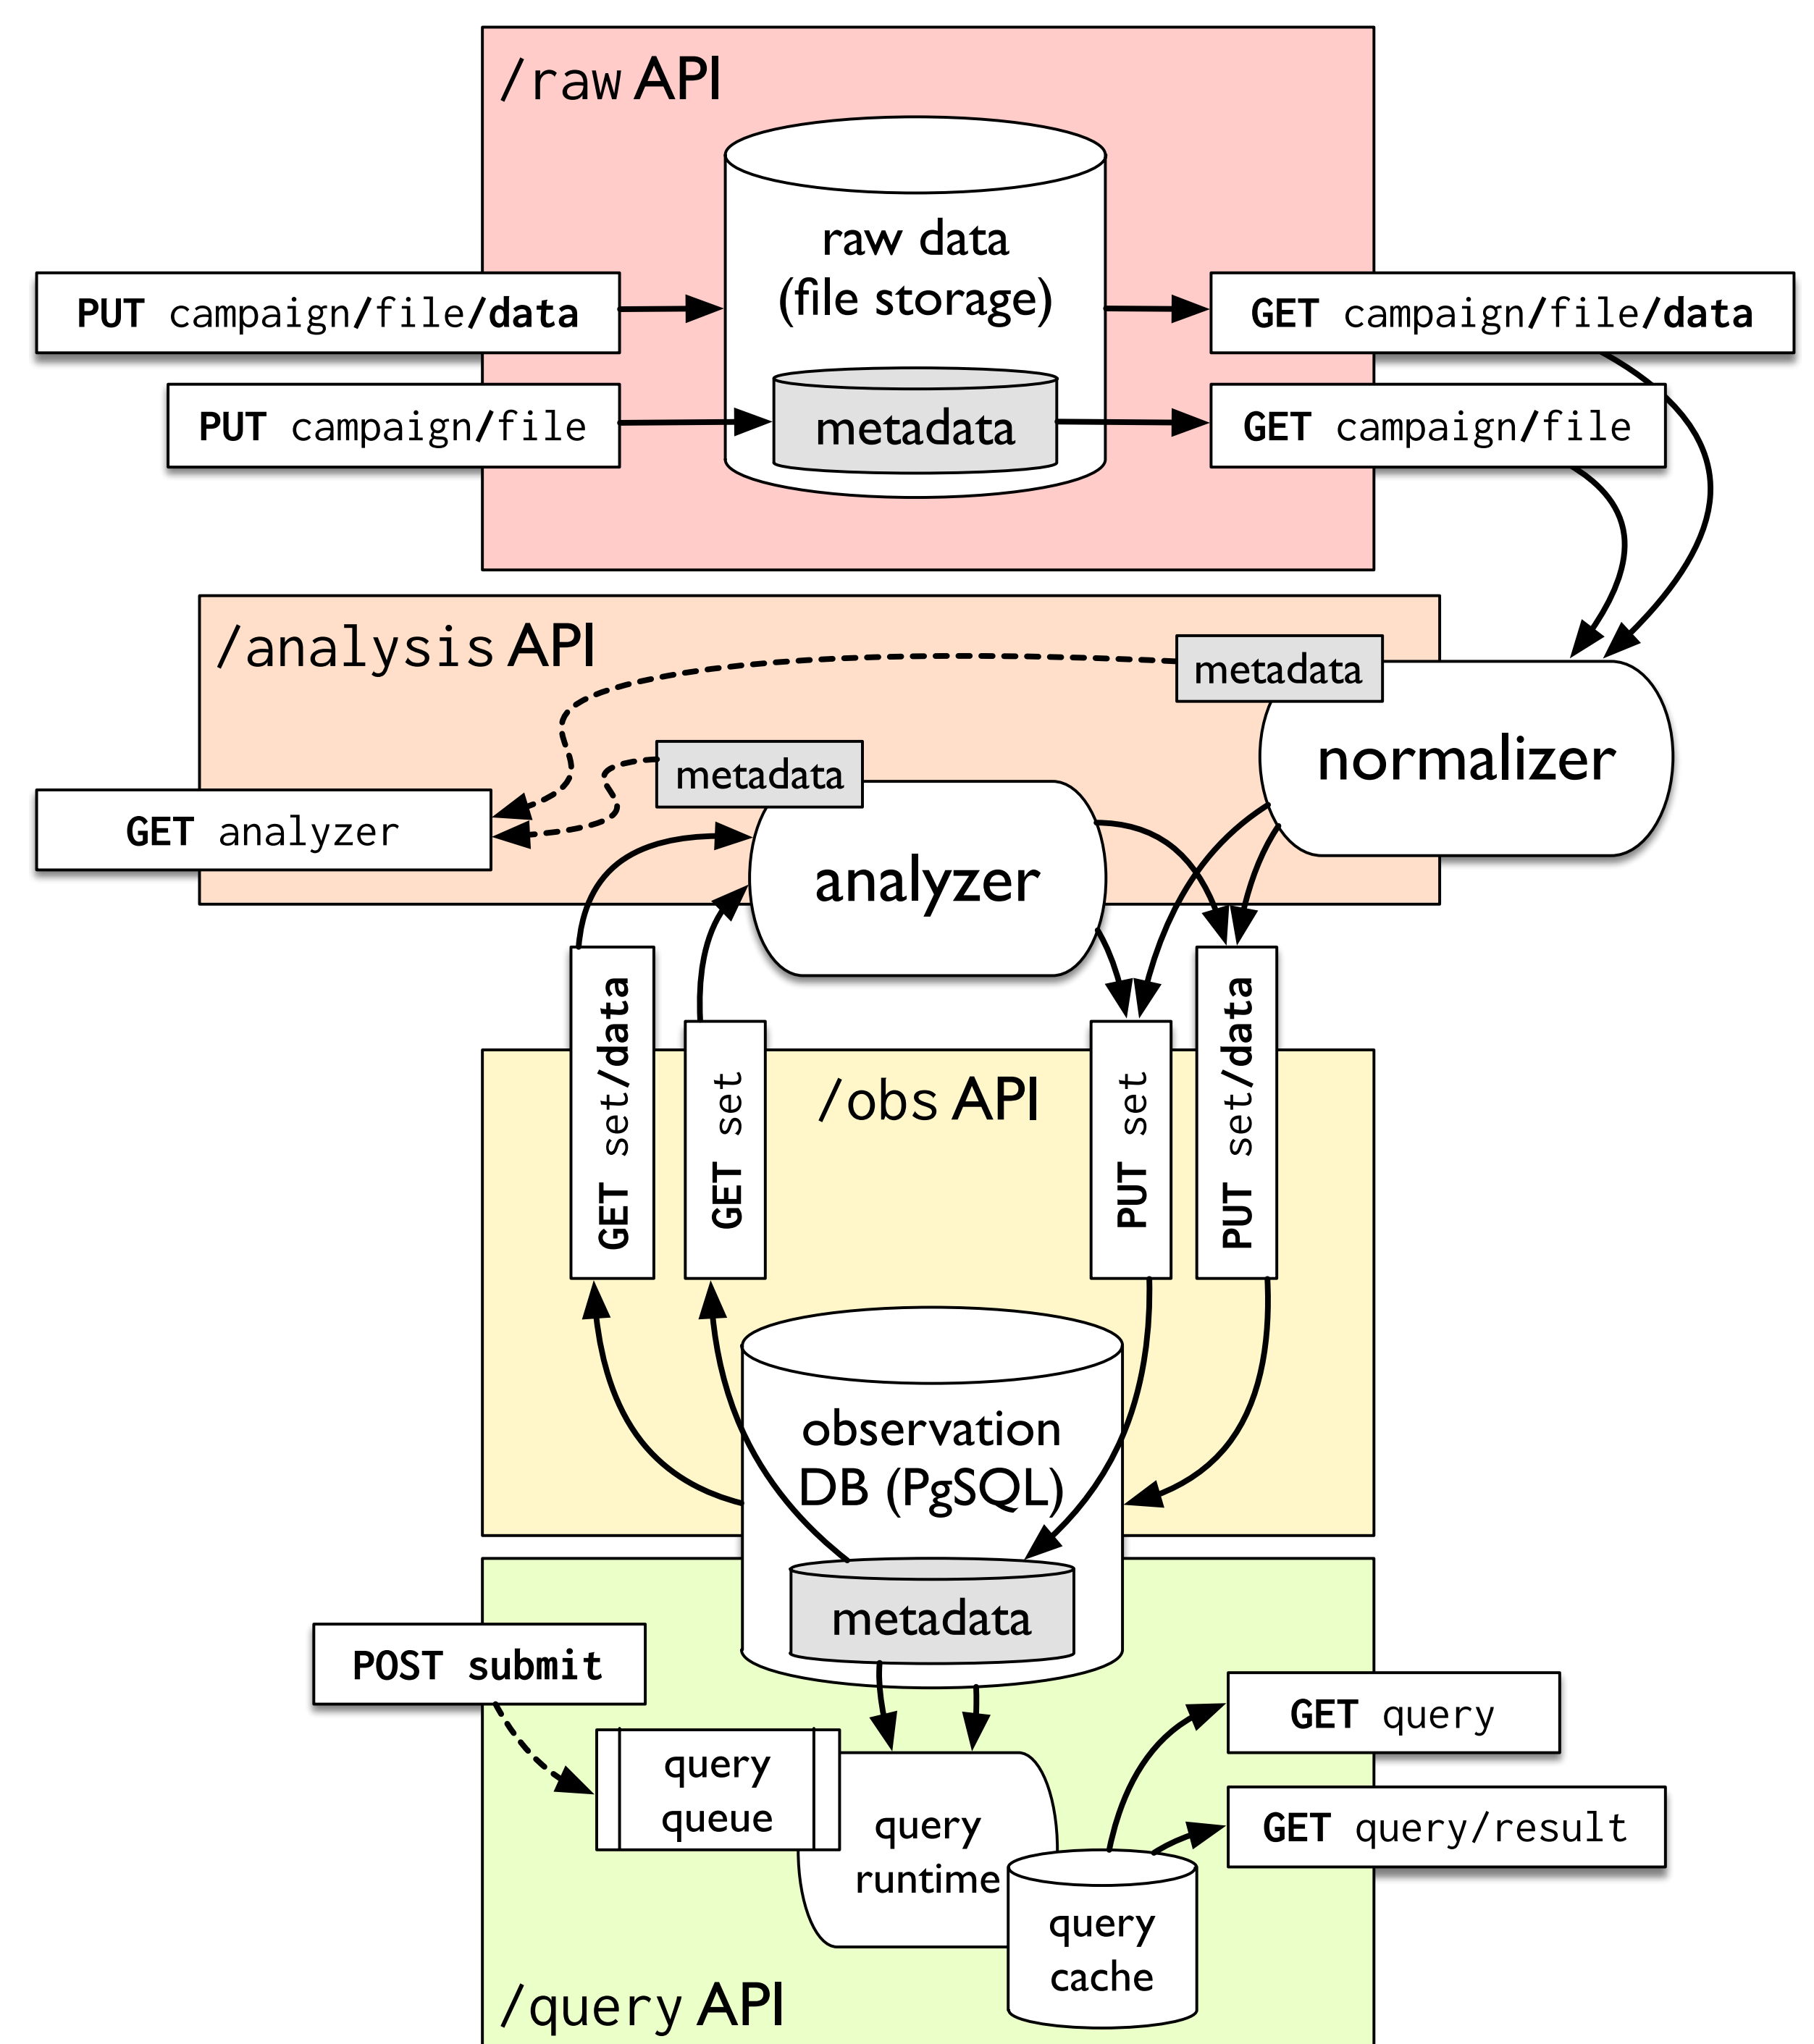
Condition: specific phenomenon observed on a path, defined in a structured namespace oriented to assign **states** to **aspects** of attempts to use a given protocol **feature**, fostering **comparability** of results.



Aspect and state definition is feature-specific, based on examination of measurement data; e.g.:

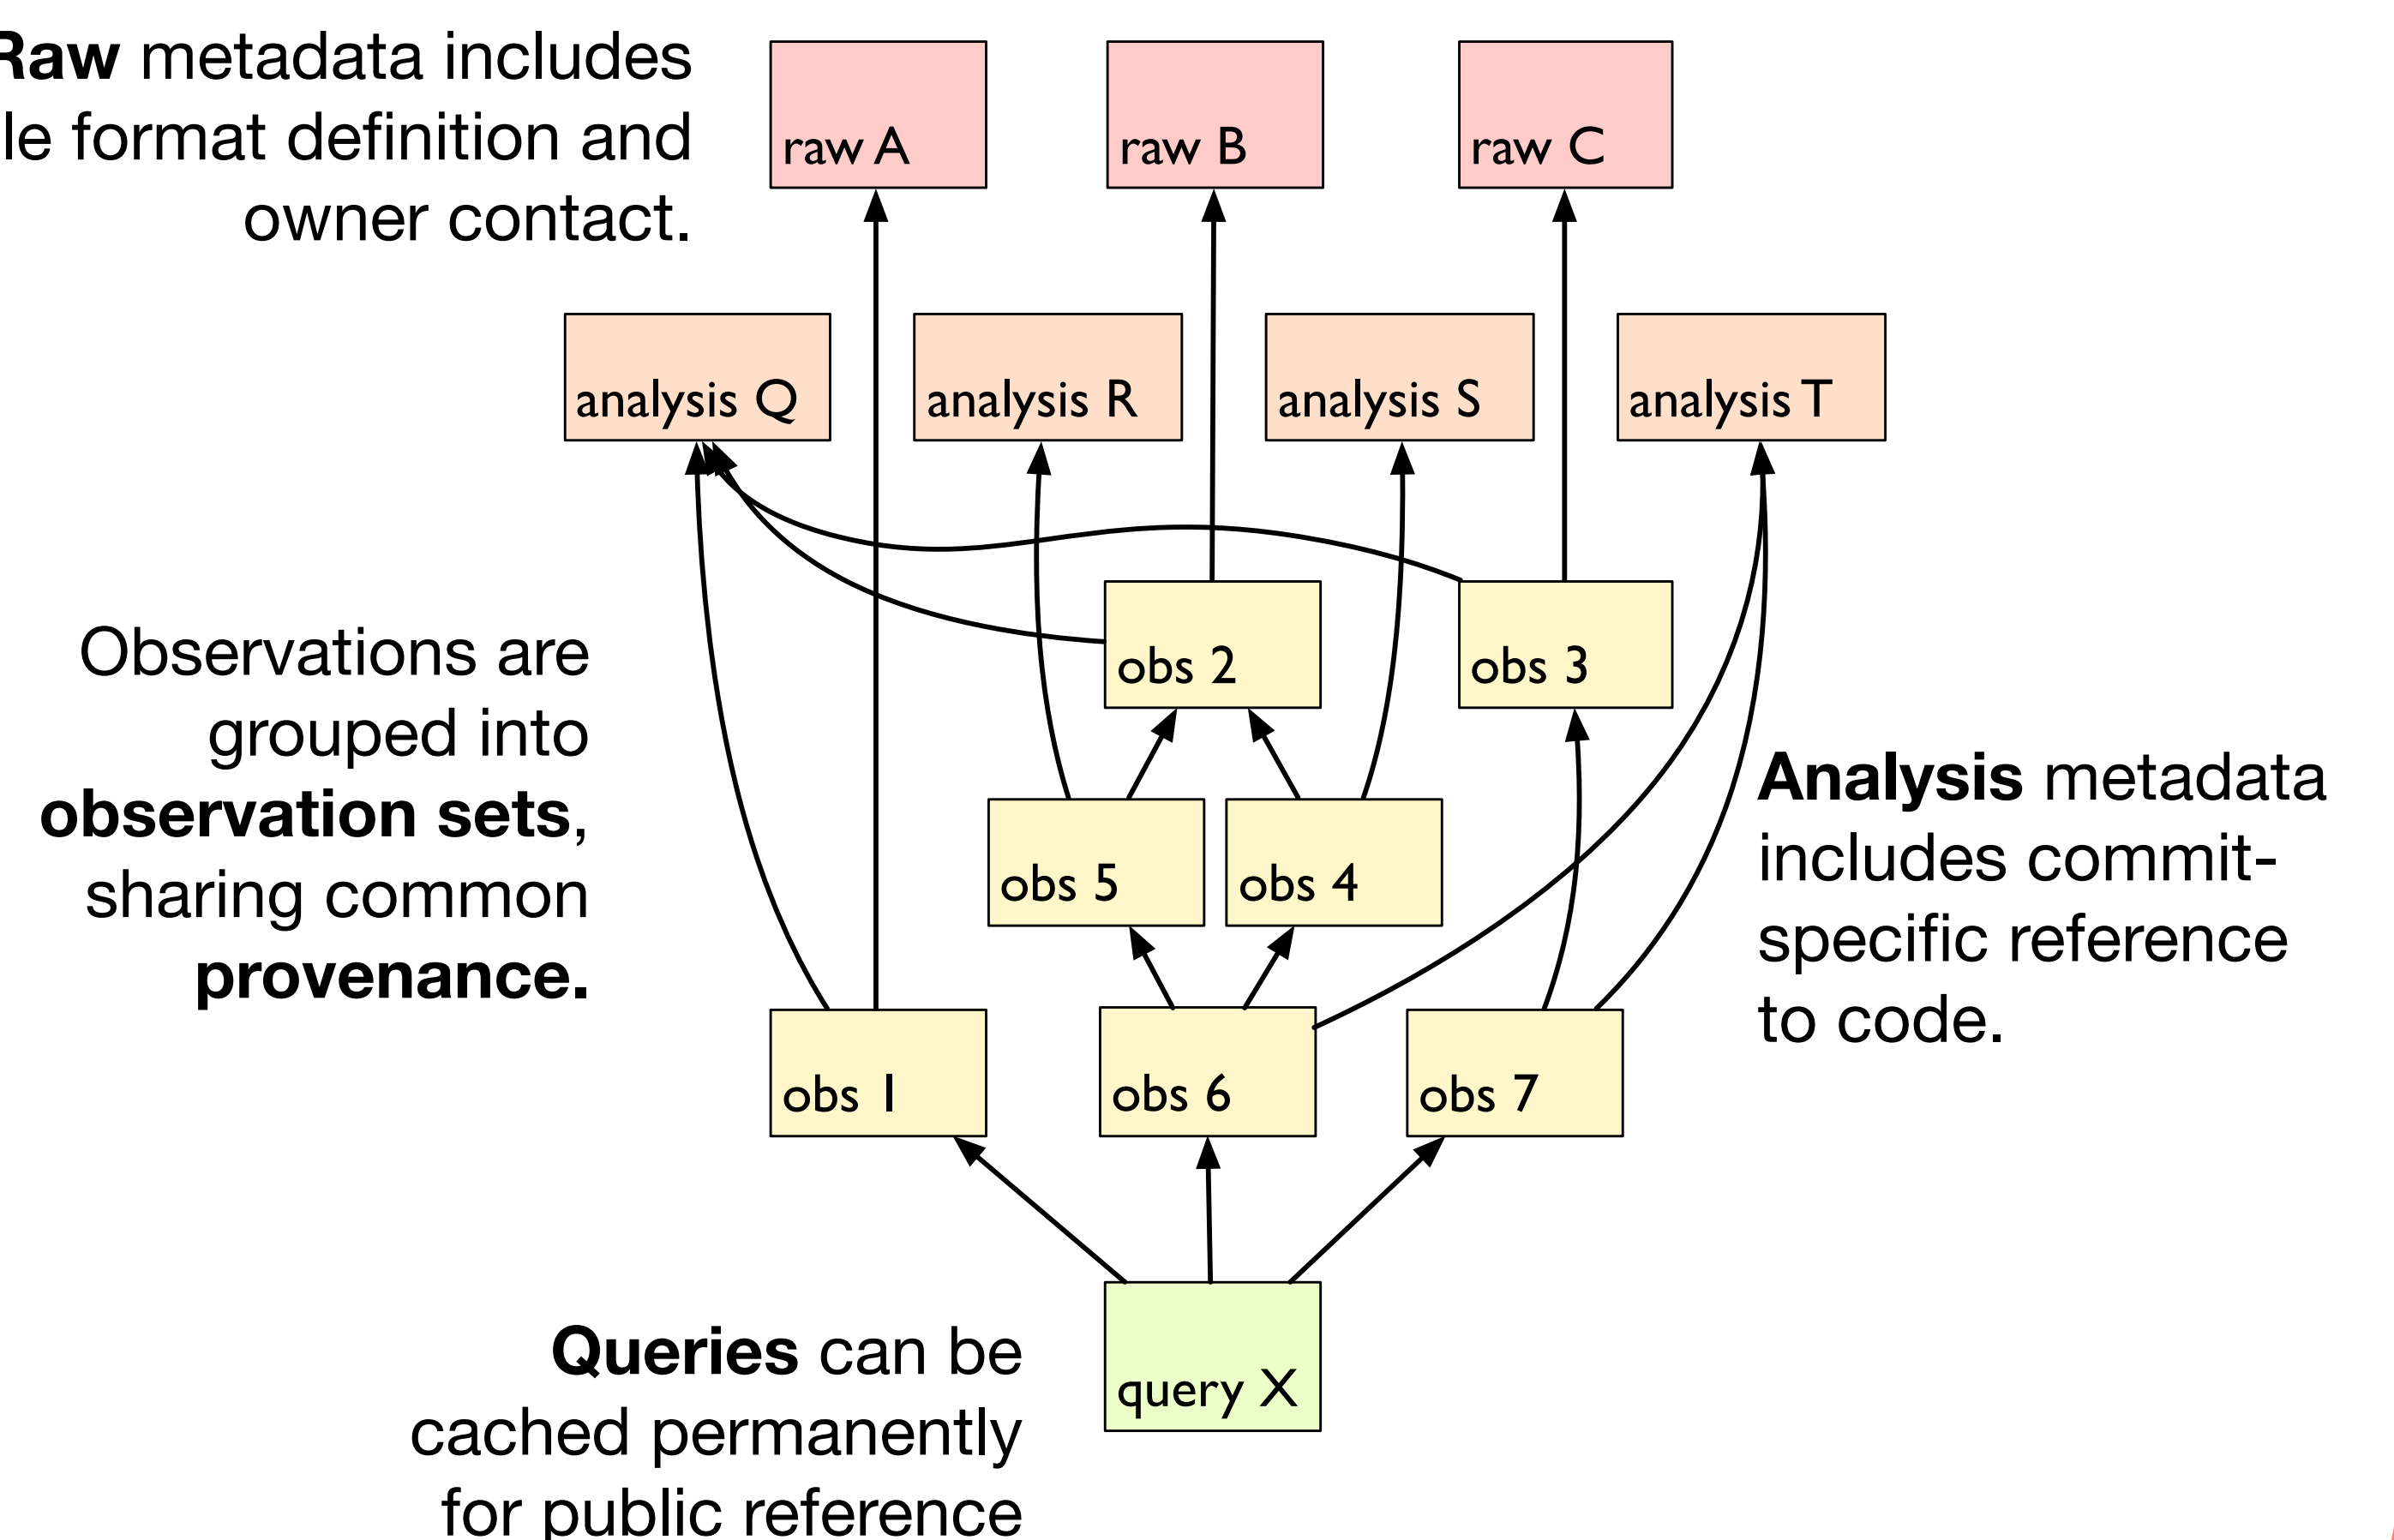


The PTO is implemented as a RESTful API, storing raw data files organized into campaigns, normalizing these into a queryable observation database.



Every object stored in the observatory, including queries, keeps its *provenance*, including arbitrary metadata, fostering measurement **repeatability**. Following links yields all antecedents for any observation or aggregate.

Raw metadata includes file format definition and owner contact.



Deploying soon: <https://observatory.mami-project.eu/> — code: <https://github.com/mami-project/pto3-go>

measurement

architecture

experimentation

