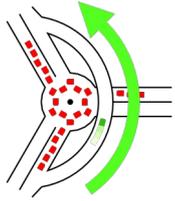


Overpass API

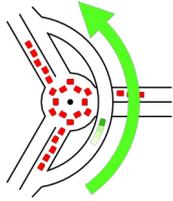
SotM-EU 2025

Roland Olbricht

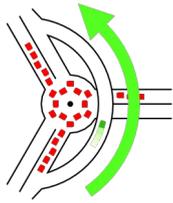


Overpass API

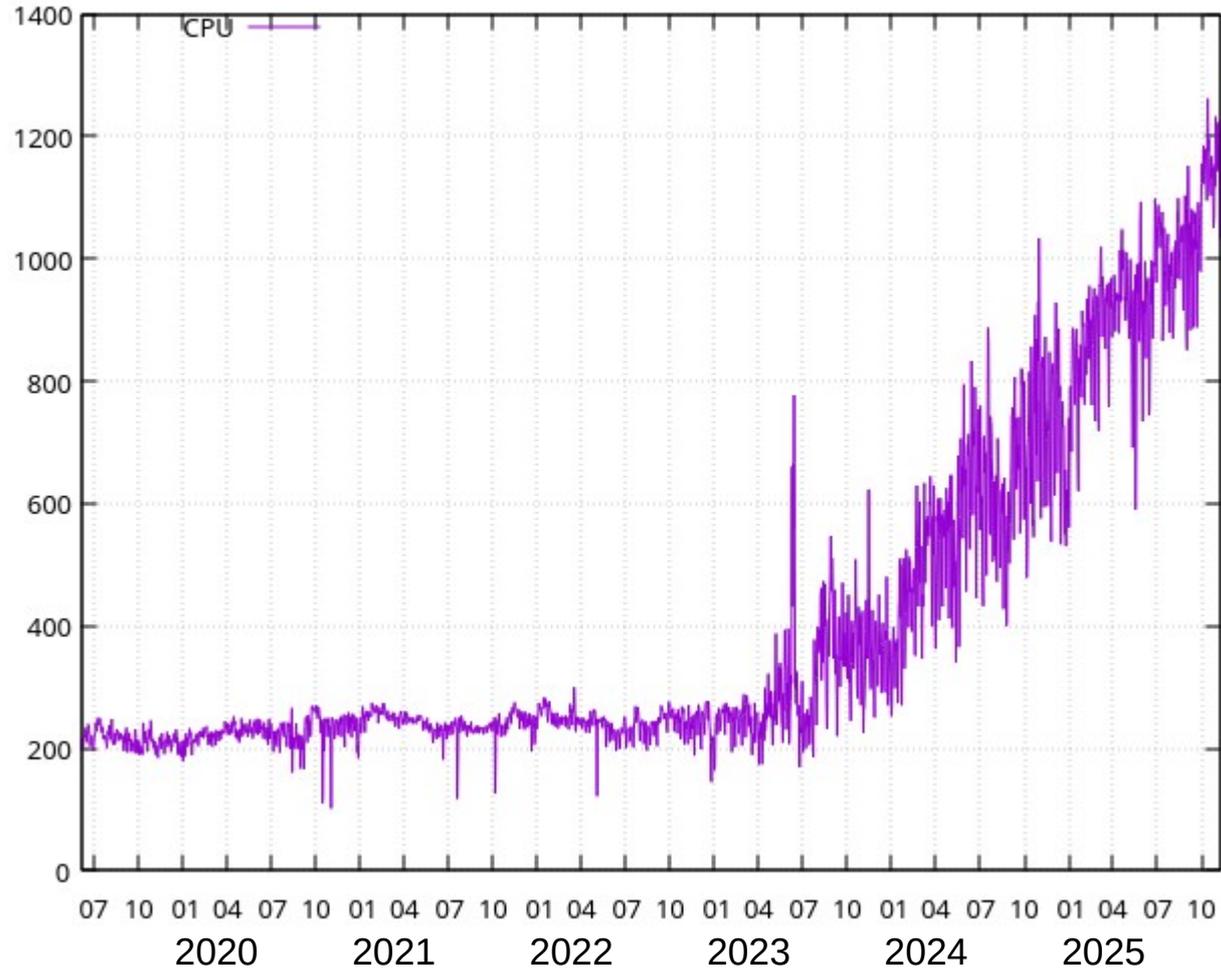
- 1. Quota System: Adaptions**
- 2. Considered Features**
- 3. Housekeeping: File Sizes**
- 4. Redaction: Adherence**

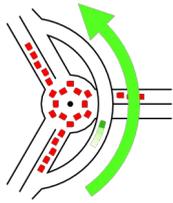


Quota System: Adaptions

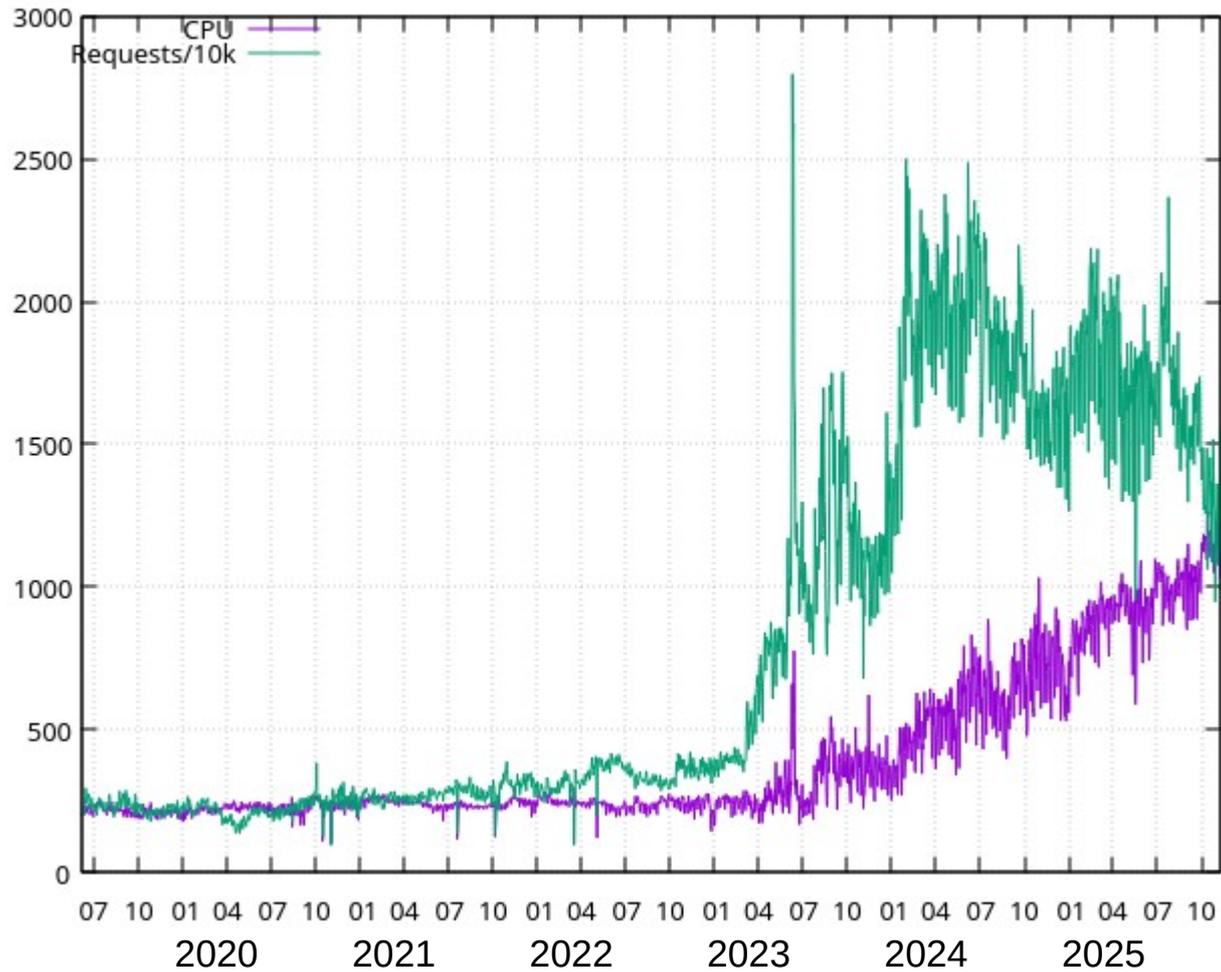


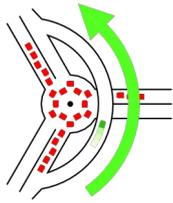
Quota System: History of Overpass API Usage



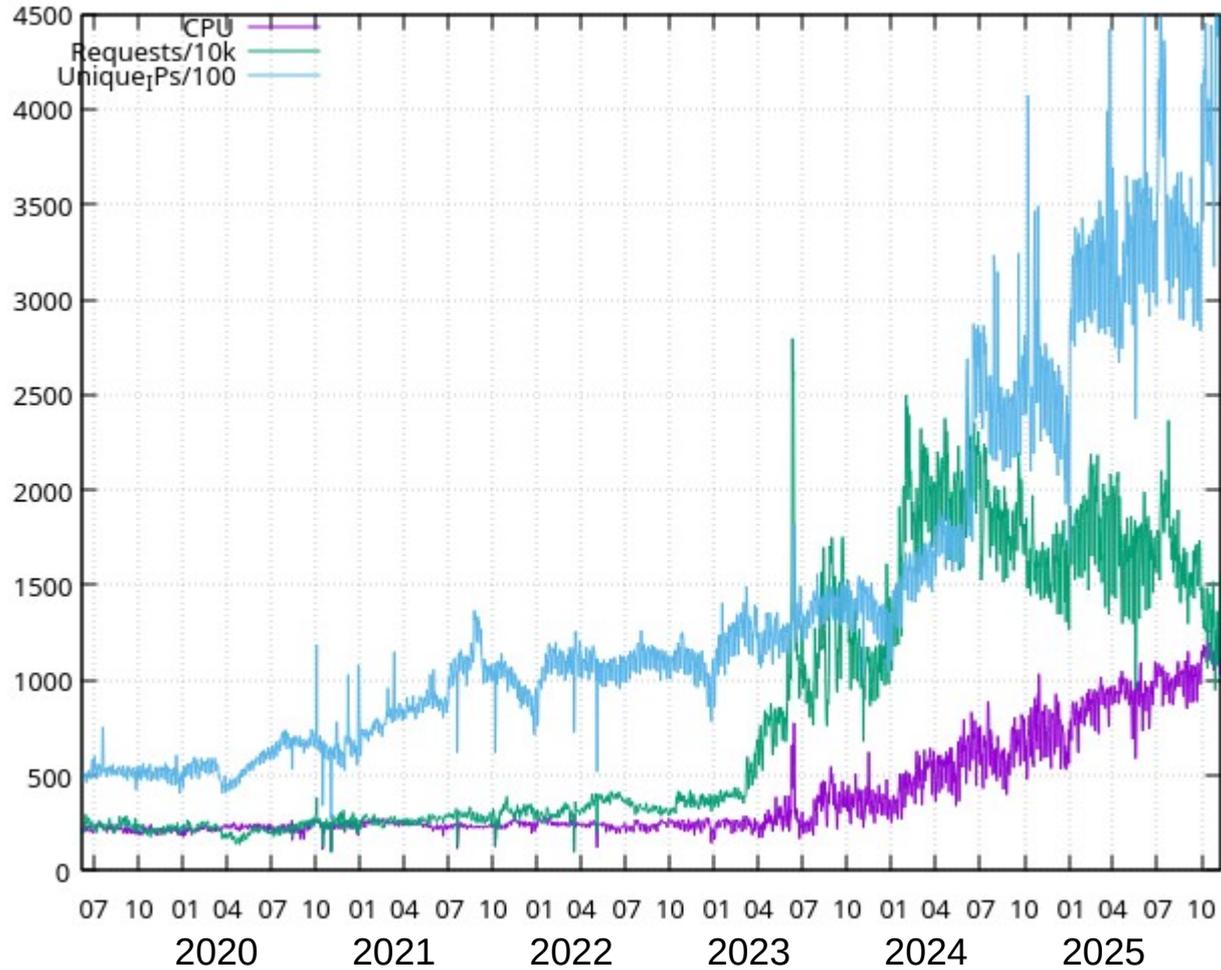


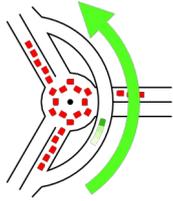
Quota System: History of Overpass API Usage





Quota System: History of Overpass API Usage





Quota System: Objectives

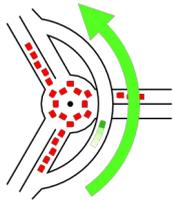
Augmented Diffs
(and all other
cool stuff)

vs.

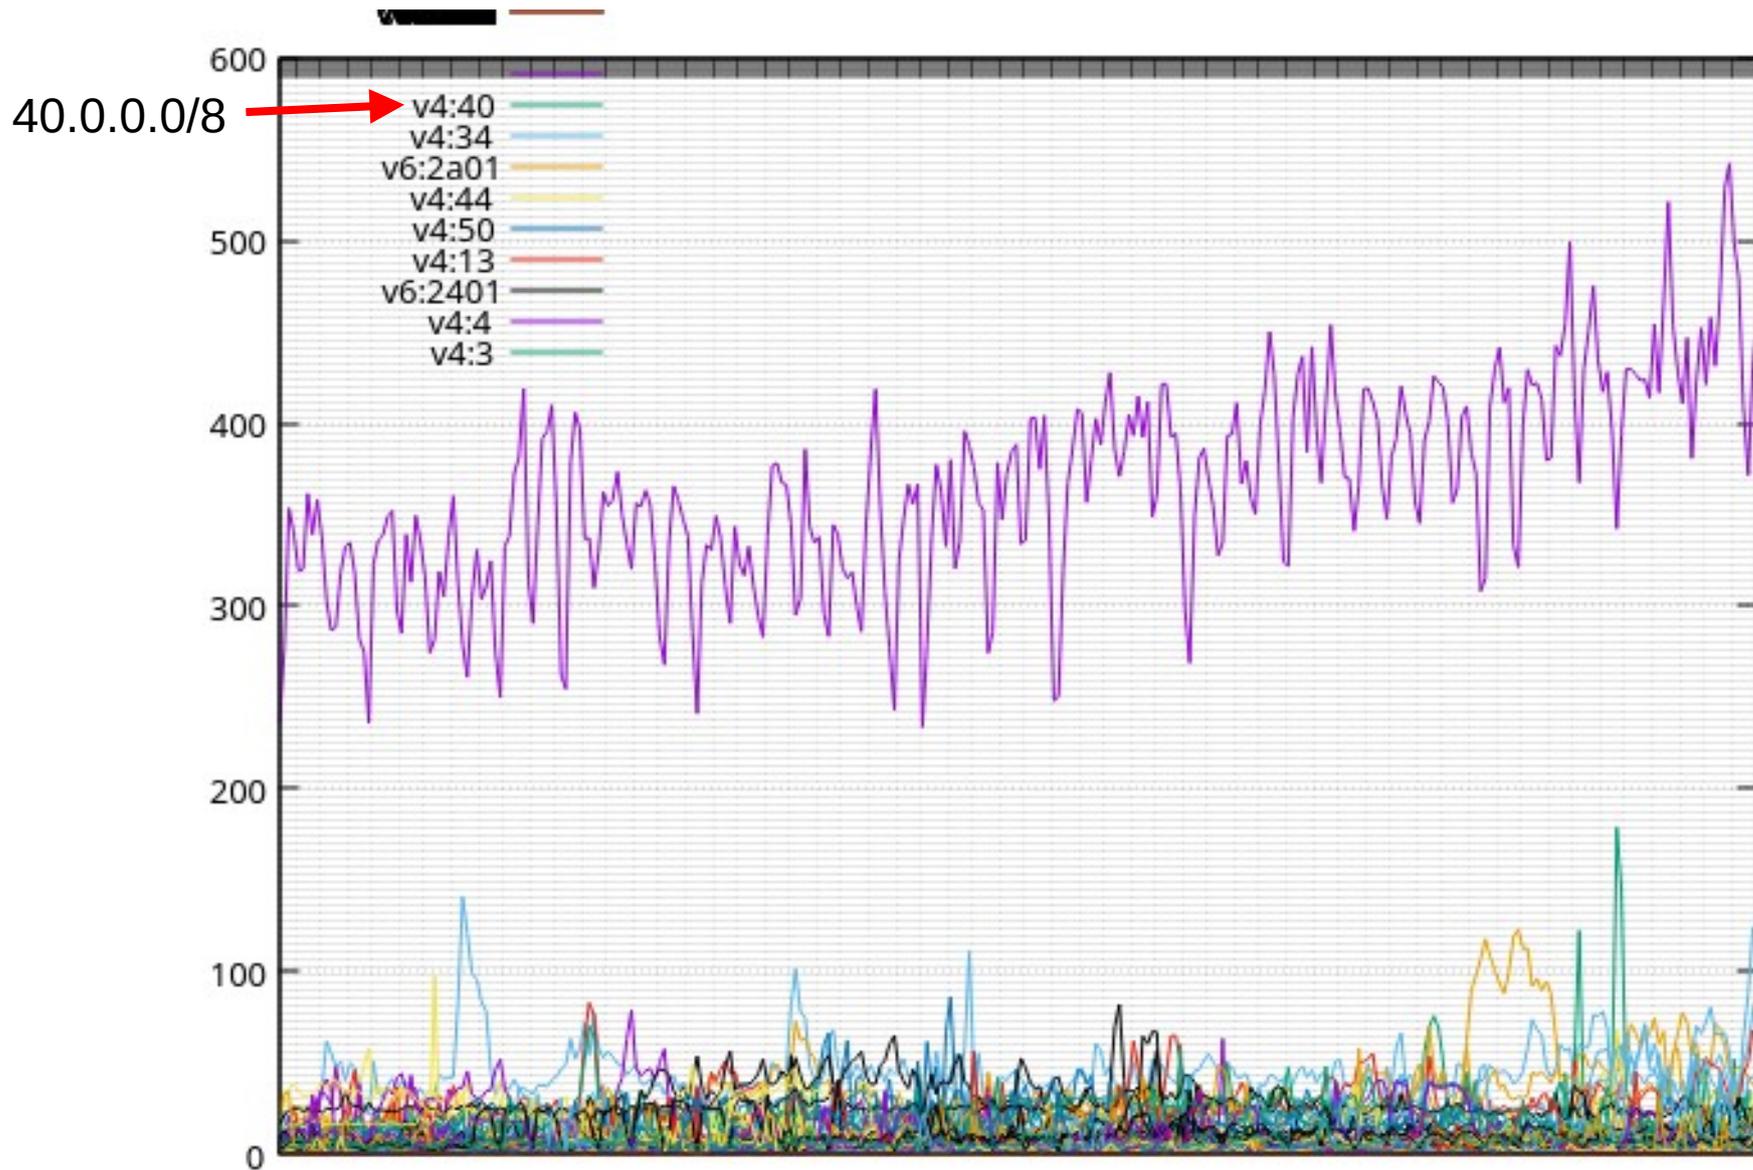
```
way(around:4,...);out;  
("Speedometer")
```

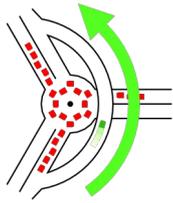
```
node[natural=water](poly:...);
```

„World’s free tile provider“
CPU time

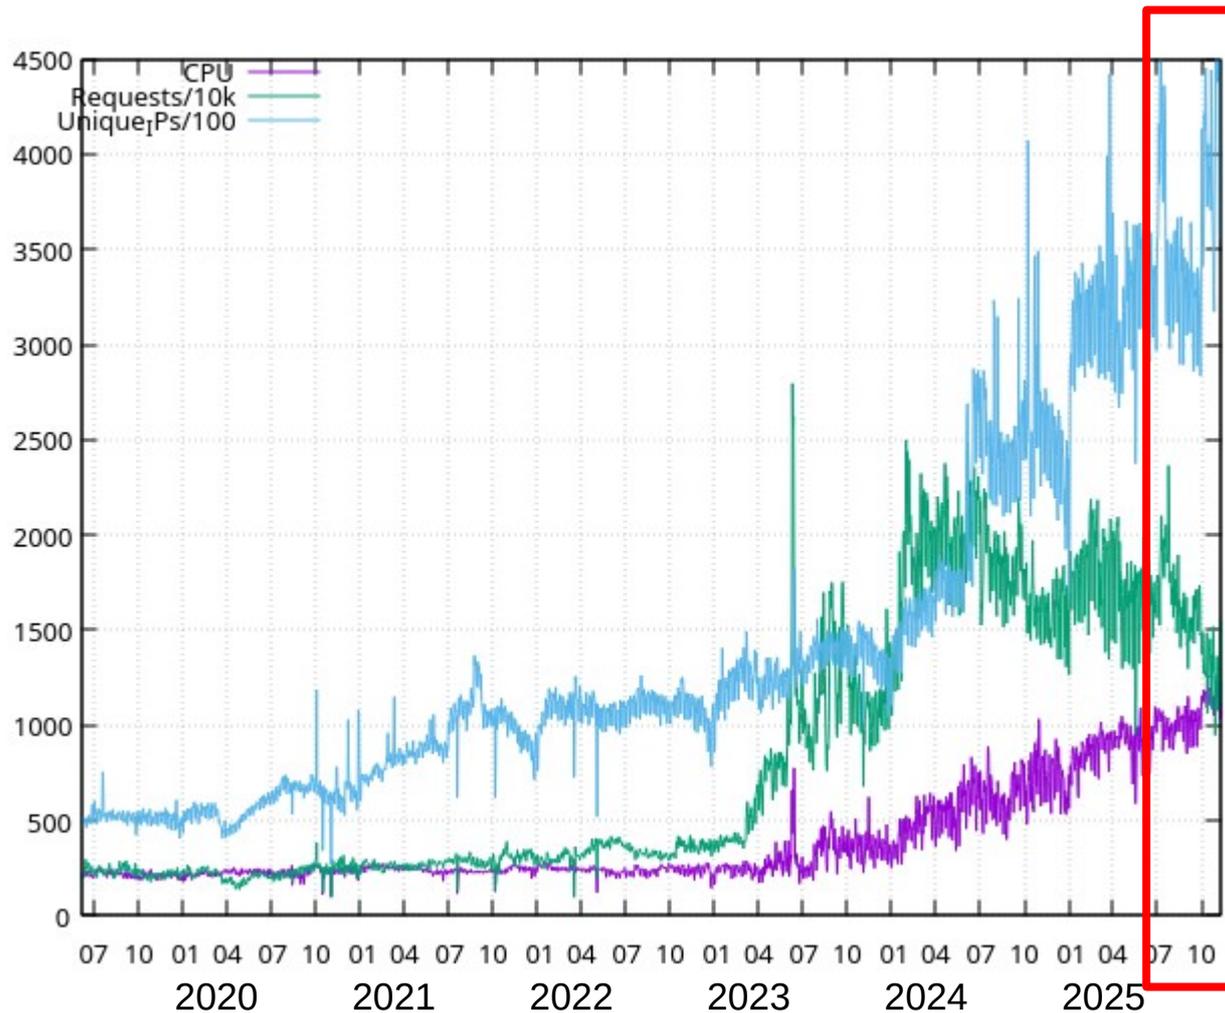


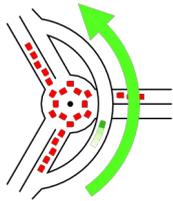
Quota System: Any Source Spiking?



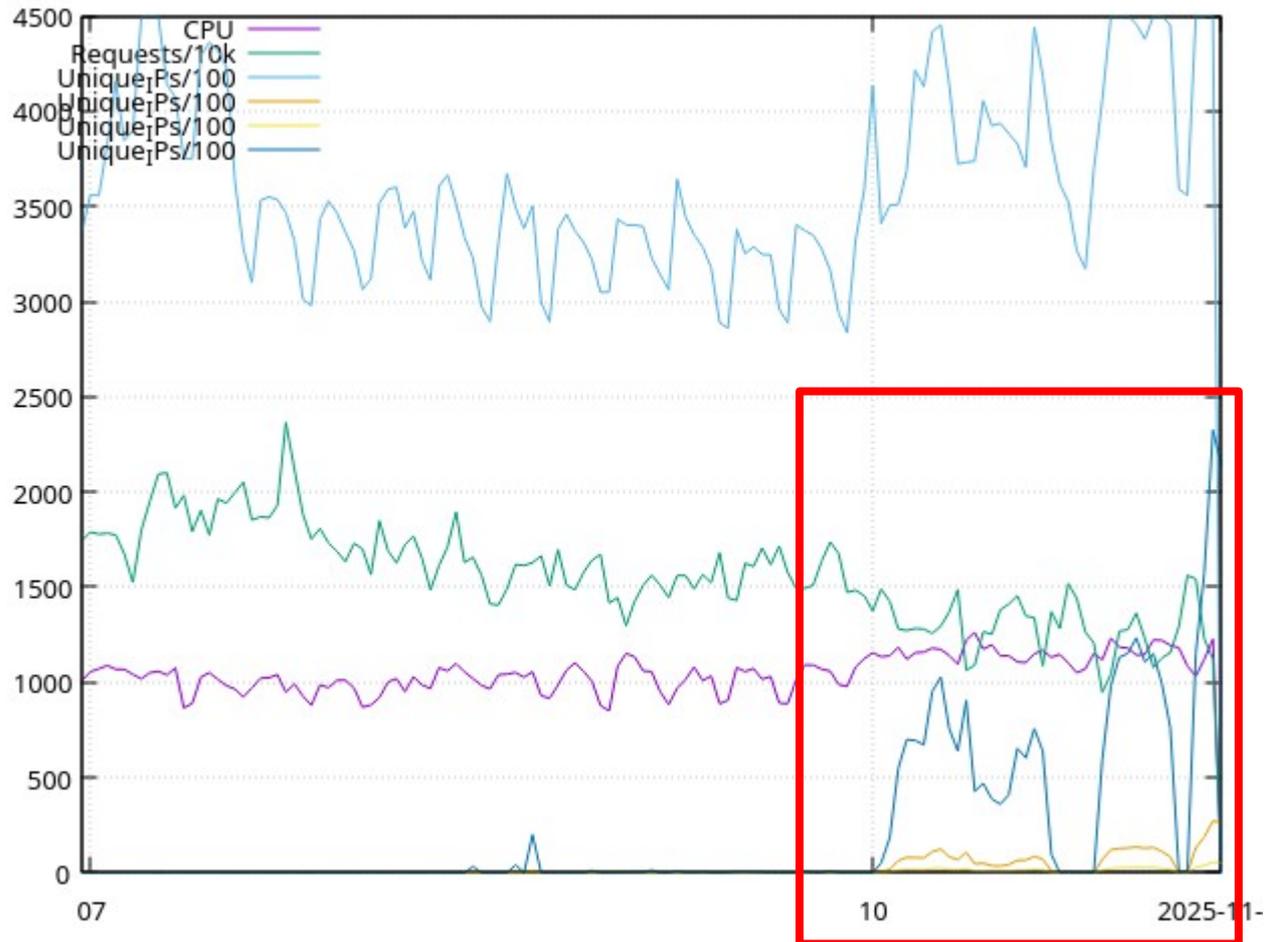


Quota System: History of Overpass API Usage



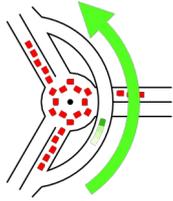


Quota System: Any Source Spiking?



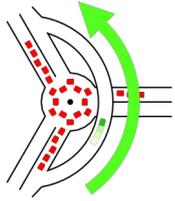
Pokemon-
Go-Cheaters

90k human users vs. **180k from AWS/Azure**

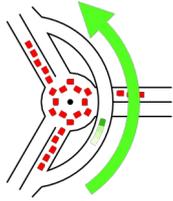


Quota System: Remedies

	IP addresses	User Agent	OSM account
Ensure contact is possible	hardly	mostly	yes
Fair resource sharing	degraded	no	yes
Enforcement	automatic	manual	automatic
Reuse for data protection	no	no	yes



Considered Features



Considered Features: Overview

Augmented Diffs:
little used

Achavi much more

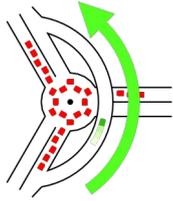
June 2024 vandalism:
Heavy undo tool?

from community:
rather **Redaction**

„All Deleted Objects“?

User experience disaster:
Moved? Split? Retagged?
Moved-deleted? Repurposed?

rather:
Tag Timeline

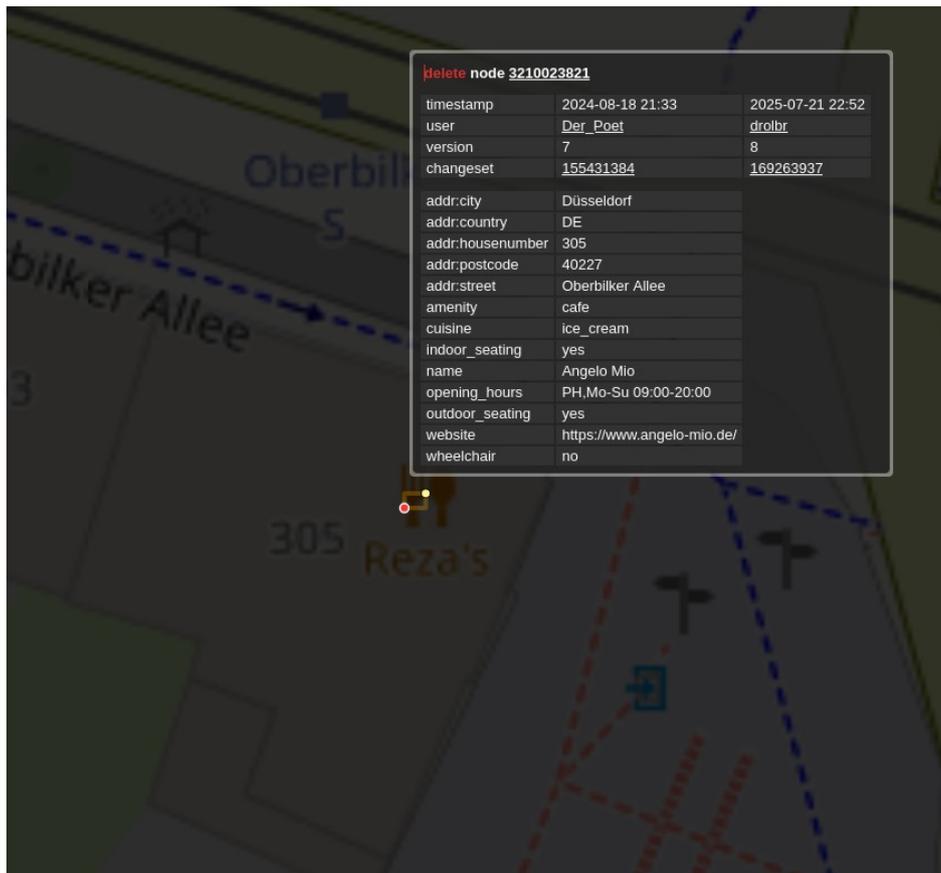


Considered Features: Tag Timeline

Has there ever been a restaurant
in this Bbox? ice cream parlor

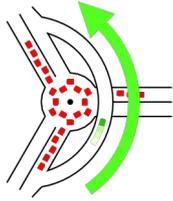
[https://overpass-api.de/achavi/
?changeset=169263937](https://overpass-api.de/achavi/?changeset=169263937)

```
[out:csv(timestamp,elements)];  
timeline(tag_kv,"cuisine","ice_cream");  
out;
```

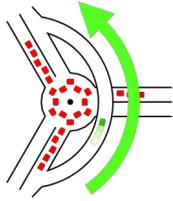


```
timestamp  elements  
2014-11-28T23:38:38Z  
2017-12-26T14:24:10Z  
2019-07-13T16:14:46Z  
2019-07-21T08:18:50Z      n3210023821  
2021-07-31T14:49:59Z      n3210023821  
2023-02-09T08:12:51Z      n3210023821  
2024-08-18T19:33:14Z      n3210023821  
2025-07-21T20:52:13Z
```

Solves also **splitting tracking!**

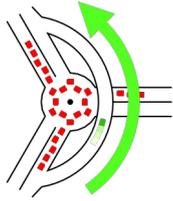


Housekeeping: File Sizes



File Sizes: OpenStreetMap context

planet...bz2	2012-11-09	24G	
planet...bz2	2025-11-08	156G	6-fold
history...bz2	2013-02-08	40G	
history...bz2	2025-11-08	240G	6-fold



File Sizes: Overpass API Database

652 GB
snapshot of
2025-11-08

842 GB
working set

ZeLonewolf's backup:

1. clone from dev.openstreetmap.de over network to disk
2. copy disk to backup disk to prepare for power outage
3. restore in case of power outage

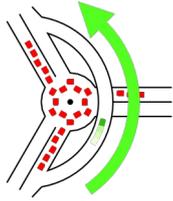
Recovery time for 670 GB (with 50 MB/s): 13400 s ~ 4 hrs
4 hrs x 3 passes

Cloning Trouble:

1. read from updating database
2. write to snapshot
3. read from snapshot for every download

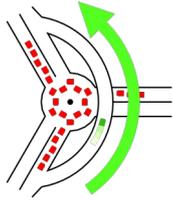
Other things happen in parallel!

Result: cloning takes **more than 24 hours**



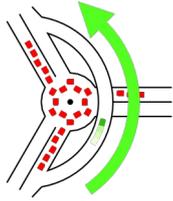
File Sizes: Breakdown

	nodes	ways	rels	sum	running
Geo data	95.4 GB	87.0 GB	1.7 GB	184.1 GB	184.1 GB
User data	77.9 GB	10.2 GB	0.2 GB	88.4 GB	272.4 GB
Attic data	77.0 GB	73.4 GB	8.3 GB	158.6 GB	431.1 GB
Speedup data	169.2 GB	48.9 GB	2.2 GB	220.6 GB	651.6 GB



File Sizes: Breakdown

	nodes	ways	rels	sum	running
Geo data	95.4 GB	87.0 GB	1.7 GB	184.1 GB	184.1 GB
User data	77.9 GB	10.2 GB	0.2 GB	88.4 GB	272.4 GB
Attic data	77.0 GB	73.4 GB	8.3 GB	158.6 GB	431.1 GB
Speedup data	169.2 GB	48.9 GB	2.2 GB	220.6 GB	651.6 GB



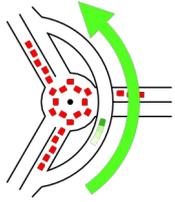
File Sizes: Improvements

Compress user data
via **Id and User Assumption**

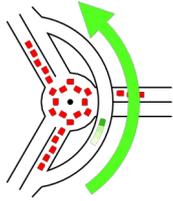
In February: **629 GB ~> 544 GB**

Replicate instead of Clone
non-essential files

In February: **544 GB ~> 462 GB**



Redaction: Adherence



Redaction: Contested Data

Peer review can result in one of the following outcomes:

Confirmed
invisible

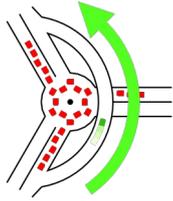
Controversial
altered
"deleted"
visible on request
(old version)

Hazardous
not visible
at all



Redaction

Slur, sb.else's passwords, data breach, copyright infringes



Redaction: Strategy for Overpass API, part 1

Redacted elements are

- **not visible through main API**
- but copies may exist at data consumers (e.g. snapshot export once a year)
- no (sincere) promise possible: rather **best effort**

OpenStreetMap shall not claim that this trap road exists.

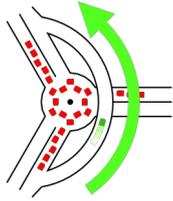
Shall OpenStreetMap claim that it never claimed that trap road existed?

We are able to not claim it has been in the database implies
We are unable to claim it not has been in the database.

However:

- neither OpenStreetMap nor Overpass API wants to spread slur

~> **Any documentation must set clear it is best effort**



Redaction: Strategy for Overpass API, part 2

S/t redaction resulted in relations w/o members (late 2012)

Strategic:

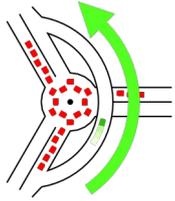
Declare **all attic** and meta **non-public** (i.e. project internal) **(done)**
Subsequently, no unredact feature is needed.

Tactic:

Overpass API will withhold output (and not the computation).
Access to redacted elements via setting "[redacted:yes]" possible,
OSM users only (maybe restrict even further).

Operational:

Redacted **element versions** are marked in the database.
The required redaction feed is created from screening full history file.



Summary

Quota via **OSM users**

Redaction via **OSM users** and feed

Size reduction is important housekeeping.

Hopefully resume feature dev at some point.