

Apache Kafka Sydney Meetup

That looks weird!

Exploring Mastodon user behaviour with Kafka & DuckDB

 data-folks.masto.host/@saubury

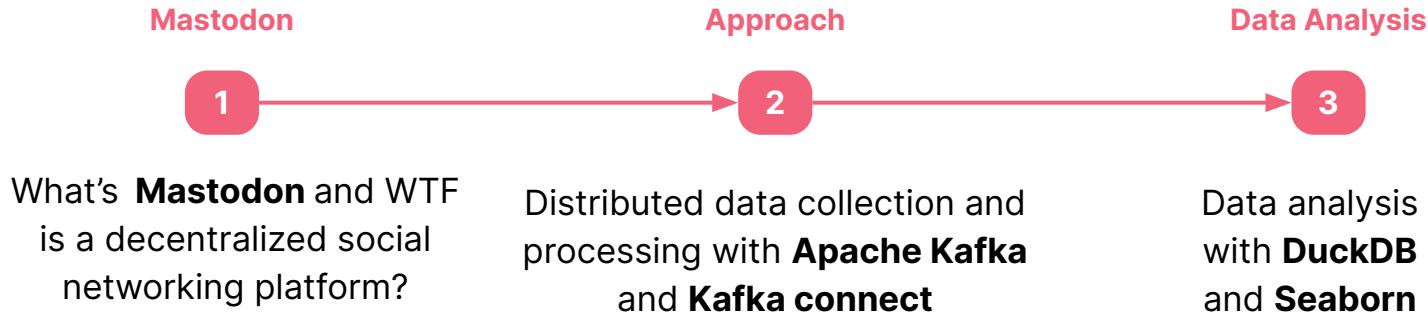
 [@SimonAubury](https://twitter.com/SimonAubury)

 [@saubury](https://github.com/saubury)



Mastodon user behaviour

What are we talking about today?



Simon Aubury

Principal Data Engineer

/thoughtworks

♥ Kafka enthusiast

🤗 Confluent Community Catalyst

🏡 Sydney, Australia



data-folks.masto.host/@saubury

Mastodon



data-folks.masto.host/@saubury



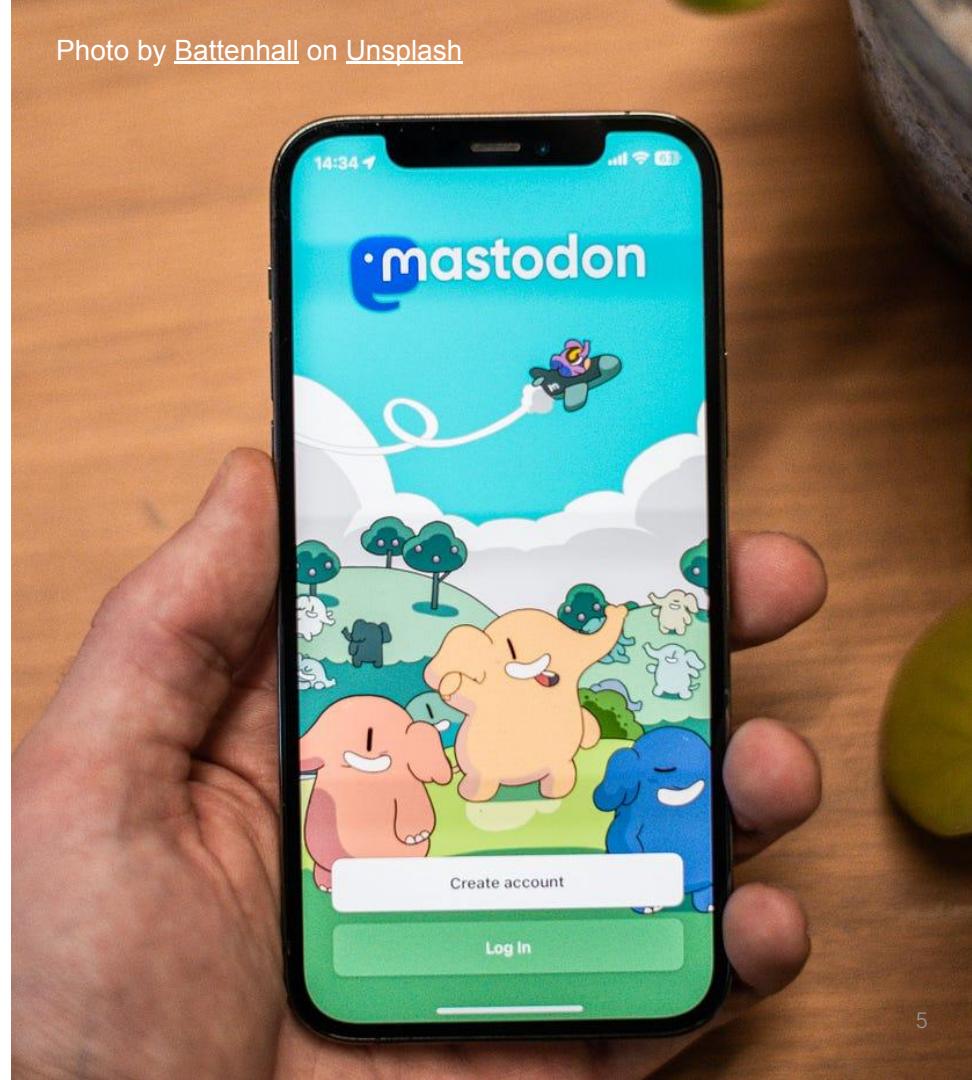
Mastodon

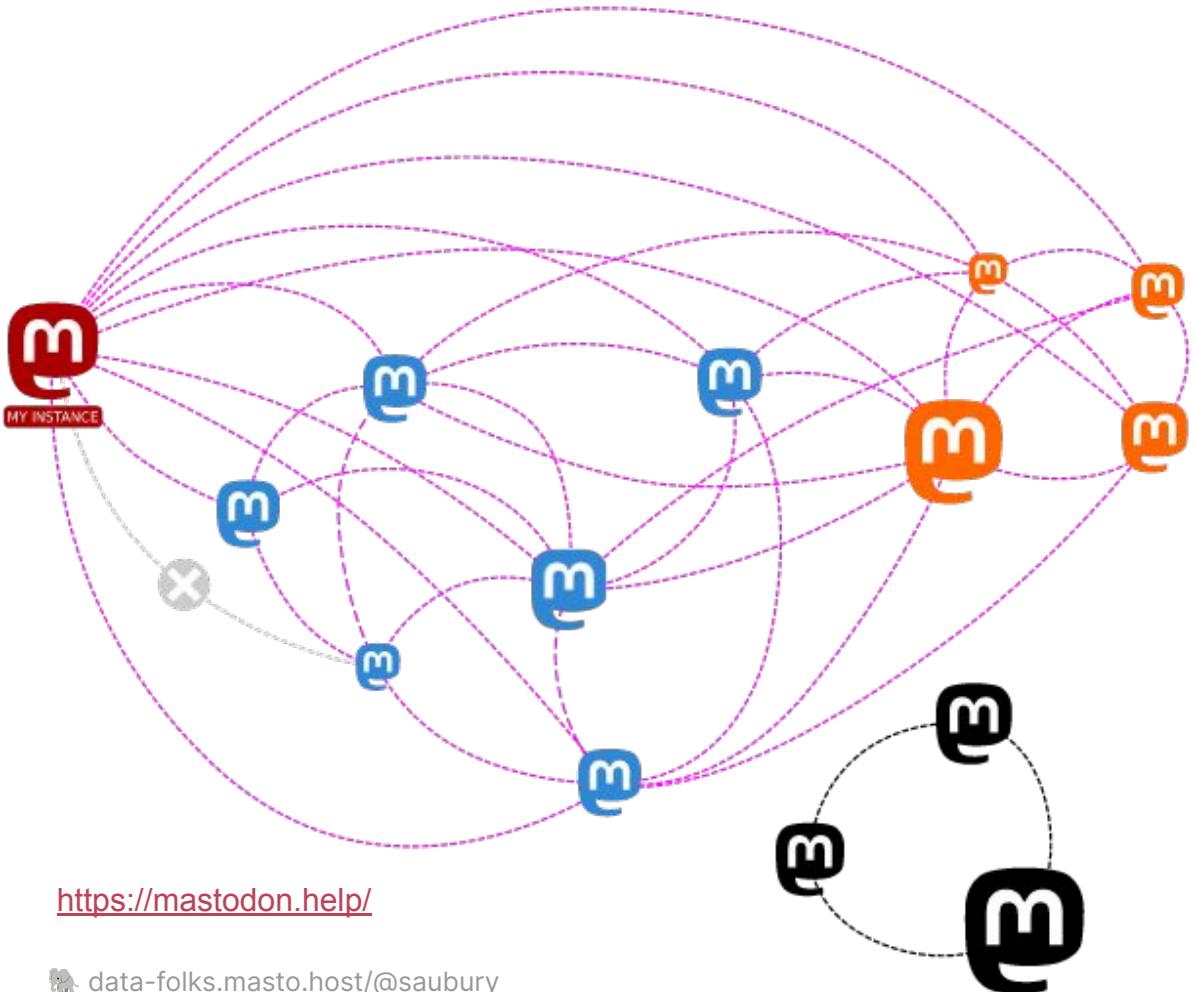
Mastodon is a *decentralized* social networking platform.

- Users are members of a **specific Mastodon instance**
- Servers are capable of joining other servers to **form a federated social network**.

Great for independent communities

Difficult to get global themes





<https://mastodon.help/>

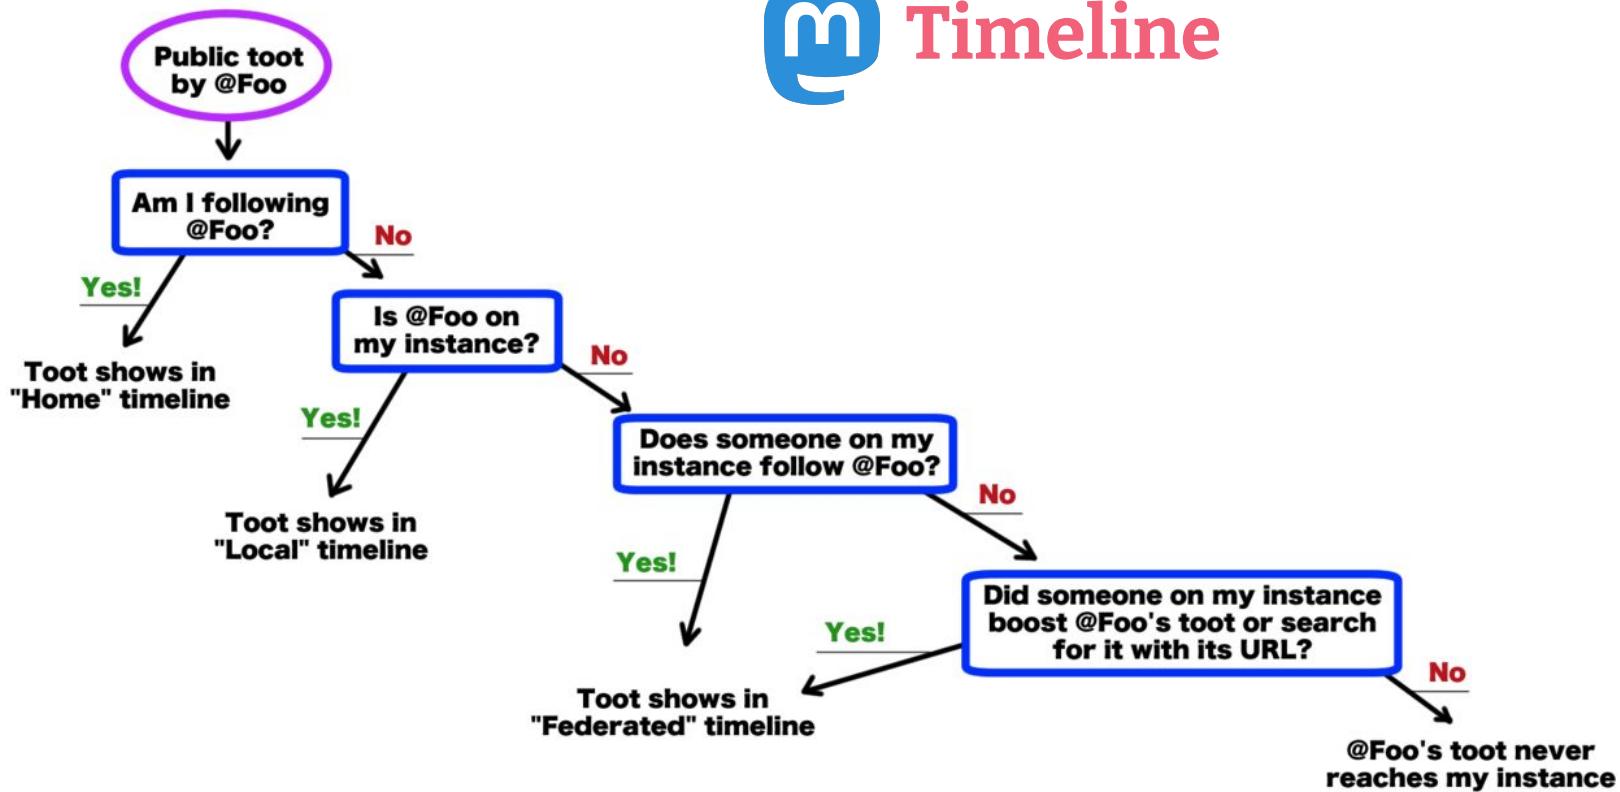
 data-folks.masto.host/@saubury

Mastodon is a galaxy of decentralized and independent networks called instances ...

Each one with its own website, rules and community.



Timeline



[Wikipedia](#)

Data collection



data-folks.masto.host/@saubury



Image by [Pexels](#) from [Pixabay](#)

Architecture



Data collection



Mastodon instance #2

Mastodon instance #3



Data processing



Kafka topic



Notebook



DuckDB



Parquet

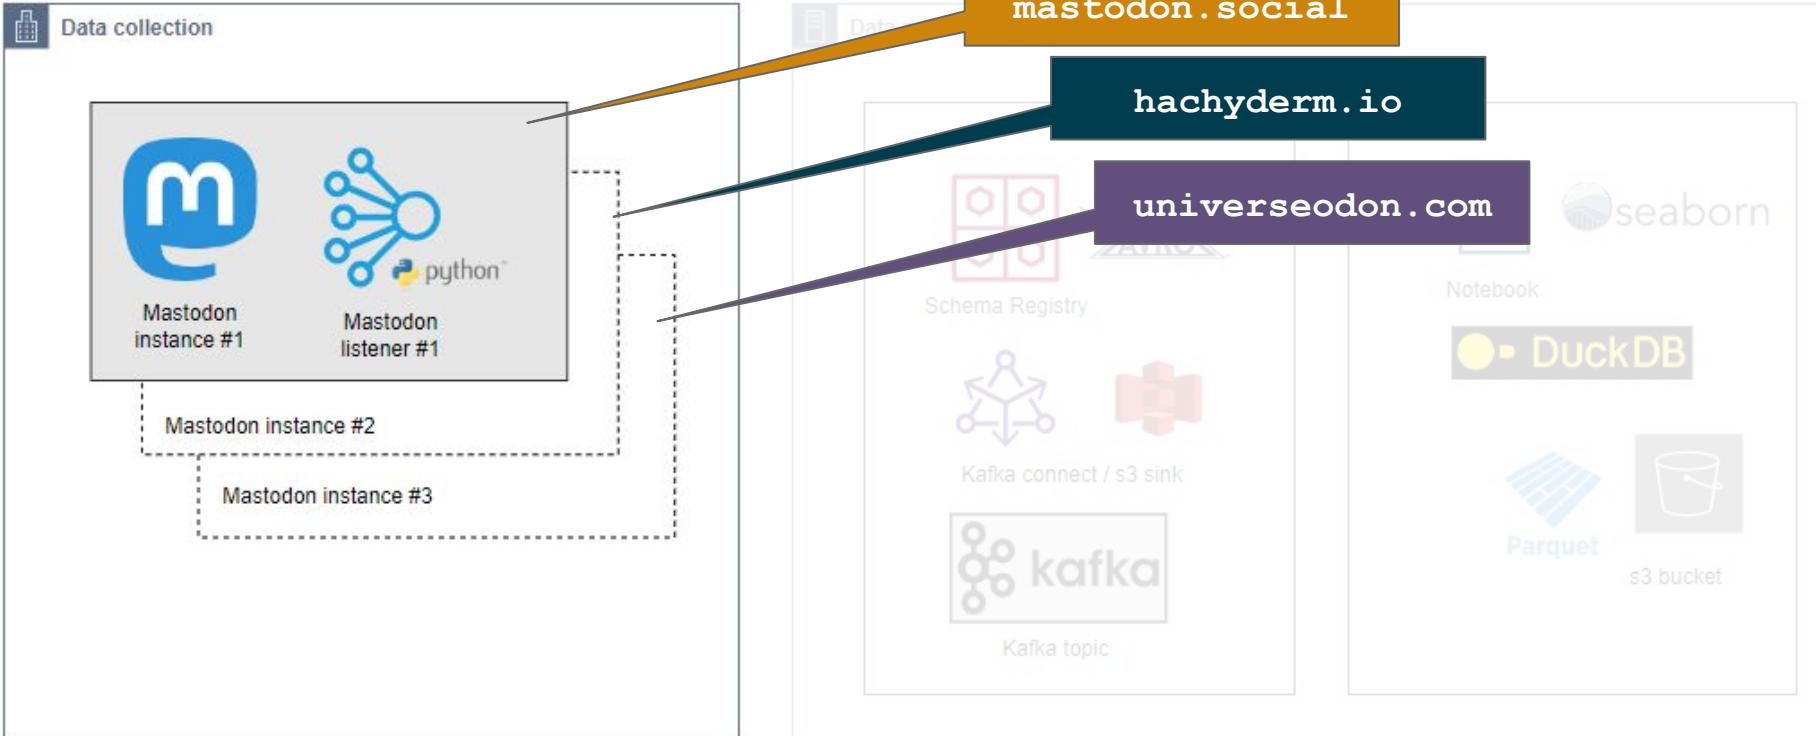


s3 bucket



data-folks.masto.host/@saubury

Data collection



Mastodon Listener

Kafka producer

```
# Kafka AVRO producer
def kafka_producer():
    producer_config = {
        'bootstrap.servers': 'localhost:9092',
        'schema.registry.url': 'http://localhost:8081',
        'broker.address.family': 'v4'
    }
    value_schema = avro.load('avro/mastodon-topic-value.avsc')
    producer = AvroProducer(producer_config, default_value_schema=value_schema)
```

AVRO serializer

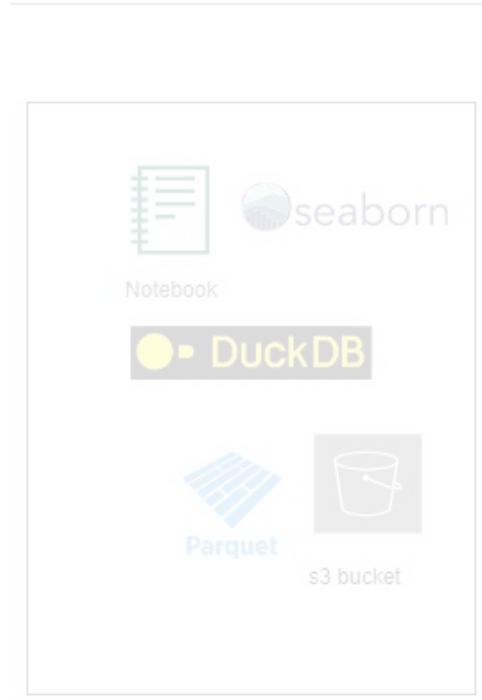
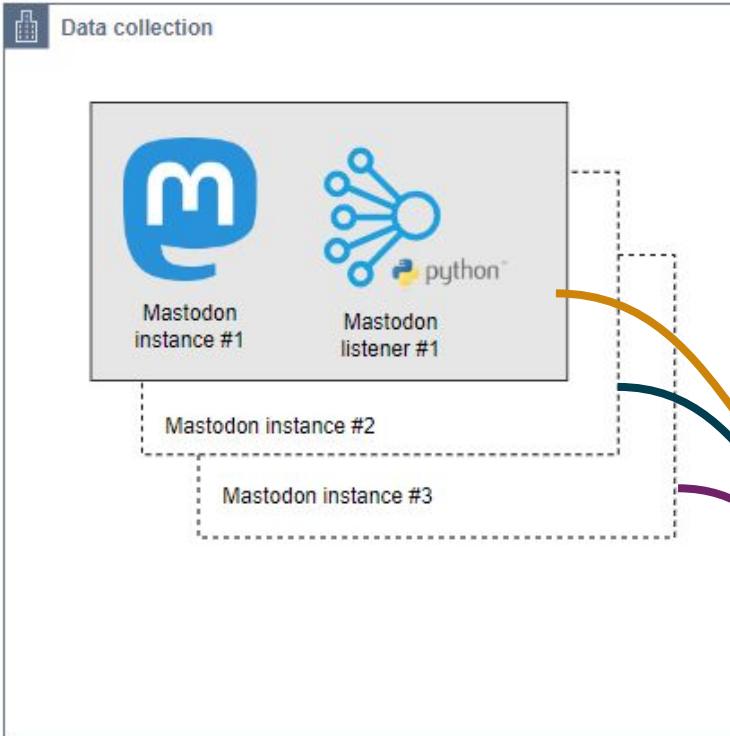
Listener

```
# Listener for Mastodon events
class Listener(mastodon.StreamListener):
    def on_update(self, status):
        m_text = BeautifulSoup(status.content, 'html.parser').text
        producer.produce(topic = topic_name, value = value_dict)
```

Servers

```
mastodon = Mastodon('https://mastodon.social')
mastodon.stream_public(Listener())
```

Streaming - Kafka



Kafka - CCC

bulwarkonline en "No shortage of Republicans ha
dnddeutsch de #LorianStudios hat den Release
RSMockimon en b10 EWE: UN biodiversity sumi
realTuckFrumper en Trump allies demand purge of J
AgentStrangeTV en Why does Miles Morales run war
ellebgomersall en Odious nonsense. Labour cannot
PaulKlee en Magdalene before the conversio
liquid_sunroy de Neln, ich gucke nicht die erst
ShaneMcCarthy en @Elonmusk Thank you fake Elon!
gmcpoth en You know that a government is
WillyCoyote69 en Eigenlijk wil ik knuffelen, ma
cgl unknown Esta noticia me causa una gran
loutre en going to school only 6h or 7h
yontou zh 将军上的评论都好好笑呀
tugu ja 「世の中には性的消費するために腰高い格好してる人もいるし性
Caladin en bro some of y'all's toots need
LucianFreud en The Butcher's Daughter, 2000 #
iembot_gyx en #Mewether #MEwx841 ACUS11 KWN
RedditGoneWild en I work with heavy equipment #
PROSPERwick en #ASA294 / #MDF56: Squawk 3311,
taniel en Chicago has created new electe
RSMockimon en b10 EWE: Rook and swift added
toonces en My kid: You know that expressi
crashkrispy en Today's accomplishments: dropp
karasu_sue ja @dqx_online Lemmyへの投稿例:https://
[]

raye en Signage at the dialysis clinic
stigusa en (then again, that would be com
Independent en Trump ordered to give evidence
hoffm en Do we report the height and we
shanaout en It's a niche question so yes,
begamit en ARTICLE I. NAME OF ORGANIZATIO
mporables en "Situated discretion is the gl
jeffsheets en I asked ChatGPT if people trus
sivation en Finally got all of garmentcras
dclarehorris en With more snow on the way.. how
uvideos en t.me/ukraina_novosti/38662 --
NZPriest en @Kirysssdal Thank you for the
wolfe en i wish someone would love me a
techbolt en #Waffle98 5/
wynch in ALL BLACK People Should Be Aw
bulwarkonline en In California, the wine and we
malcontato pt Vizualizaci um compa reclamando
bulwarkonline en "No shortage of Republicans ha
dnddeutsch de #LorianStudios hat den Release
ellebgomersall en Odious nonsense. Labour cannot
PaulKlee en Magdalene before the conversio
loutre en going to school only 6h or 7h
Caladin en bro some of y'all's toots need
NormanRockwell en Babysitter, 1927 #regionalism
toonces en My kid: You know that expressi
[]

Shoro fr Vehicle emissions may cause av
gratefuldead en Jail deaths hit 20-year high,
Malik de @mark I am not sure about my S
TsukikoshiYoshitoshi en Tsukikoshi Yoshitoshi, Appearing
emsc en #Earthquake (#isismo) M3.0 strl
SavedFox en New pet, who is this? https://
darkosubotica en So, a while back I've manage t
jikodesu en Sushi restaurants in Japan hav
davemark en #butterflyfly
stigusa en (then again, that would be com
Gondolthim es Veré el primer cap de Carnaval
ckeskrisy en Somewhere over Nevada with #Mu
mporables en "Situated discretion is the gl
NZPriest en @Kirysssdal Thank you for the
news in ASX edges up as nervy Wall Str
jonutens en The Republic of Estonia celeb
raspau en Has anyone asked #chatGTP to w
eritfactCollection en Göbekli Tepe is the oldest tem
wing en New local transit survey! A hi
bulwarkonline en In California, the wine and we
bulwarkonline en "No shortage of Republicans ha
PaulKlee en Magdalene before the conversio
feijoo en Doing the #cancer shuffle toda
WillyCoyote69 en Eigenlijk wil ik knuffelen, ma
toonces en My kid: You know that expressi
[]

bulwarkonline en "No shortage of Republicans ha
dnddeutsch de #LorianStudios hat den Release
RSMockimon en b10 EWE: UN biodiversity sumi
realTuckFrumper en Trump allies demand purge of J
AgentStrangeTV en Why does Miles Morales run war
ellebgomersall en Odious nonsense. Labour cannot
PaulKlee en Magdalene before the conversio
liquid_sunroy de Neln, ich gucke nicht die erst
ShaneMcCarthy en @Elonmusk Thank you fake Elon!
gmcpoth en You know that a government is
WillyCoyote69 en Eigenlijk wil ik knuffelen, ma
cgl unknown Esta noticia me causa una gran
loutre en going to school only 6h or 7h
yontou zh 将军上の评论都好好笑呀
tugu ja 「世の中には性的消費のために腰高い格好してる人もいるし性
Caladin en bro some of y'all's toots need
LucianFreud en The Butcher's Daughter, 2000 #
iembot_gyx en #Mewether #MEwx841 ACUS11 KWN
RedditGoneWild en I work with heavy equipment #
PROSPERwick en #ASA294 / #MDF56: Squawk 3311,
taniel en Chicago has created new electe
RSMockimon en b10 EWE: Rook and swift added
toonces en My kid: You know that expressi
crashkrispy en Today's accomplishments: dropp
karasu_sue ja @dqx_online Lemmyへの投稿例:https://
[]



Authenticating

Mastodon.py

You generally *don't* need your application to be authorized to listen to a mastodon feed.



The docs recommend to **not** authenticate

The screenshot shows a web browser displaying the Mastodon.py documentation at mastodonpy.readthedocs.io/en/stable/01_general.html. The page title is "Mastodon.py" and the sub-section is "01_general". A sidebar on the left lists "General information" (Rate limiting, Pagination, IDs and unpacking, Versioning, A brief note on block lists), "Return values", and "Error handling". Below that is a section titled "API METHODS" containing "App registration, authentication and preferences", "Statuses, media and polls", "Accounts, relationships and lists", "Reading data: Timelines", and "Instance-wide data and search". The main content area starts with a note about advanced polling and then discusses pagination. A callout box highlights a specific note: "If your application requires many hits to endpoints that are available without logging in, do consider using Mastodon.py without authenticating to get the full per-IP limit." The "Pagination" section is expanded, showing its definition and examples.

advanced one and allows you to just poll in a loop without ever sleeping at all yourself. It is for applications that would rather just pretend there is no such thing as a rate limit and are fine with sometimes not being very interactive.

In addition to the per-user limit, there is a per-IP limit of 7500 requests per 5 minute time slot, and tighter limits on logins. Mastodon.py does not make any effort to respect these.

If your application requires many hits to endpoints that are available without logging in, do consider using Mastodon.py without authenticating to get the full per-IP limit.

Pagination ¶

Many of Mastodon's API endpoints are paginated. What this means is that if you request data from them, you might not get all the data at once - instead, you might only get the first few results.

All endpoints that are paginated have four parameters: `since_id`, `max_id`, `min_id` and `limit`. `since_id` allows you to specify the smallest id you want in the returned data, but you will still always get the newest data, so if there are too many statuses between the newest one and `since_id`, some will not be returned. `min_id`, on the other hand, gives you statuses with that minimum id and newer, starting at the given id. `max_id`, similarly, allows you to specify the largest id you want. By specifying either `min_id` or `max_id` (generally, only one, not both, though specifying both is supported starting with Mastodon

Data storage



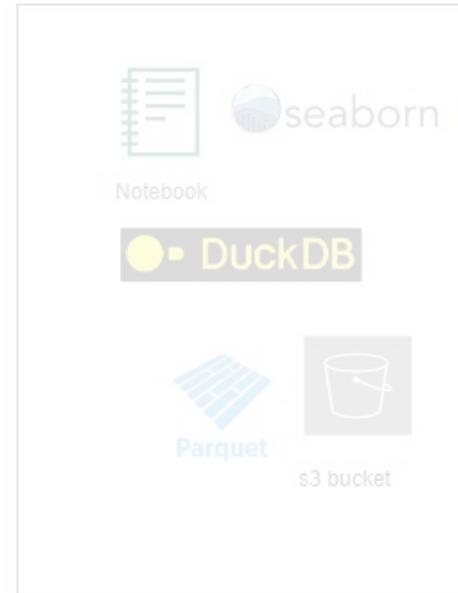
data-folks.masto.host/@saubury



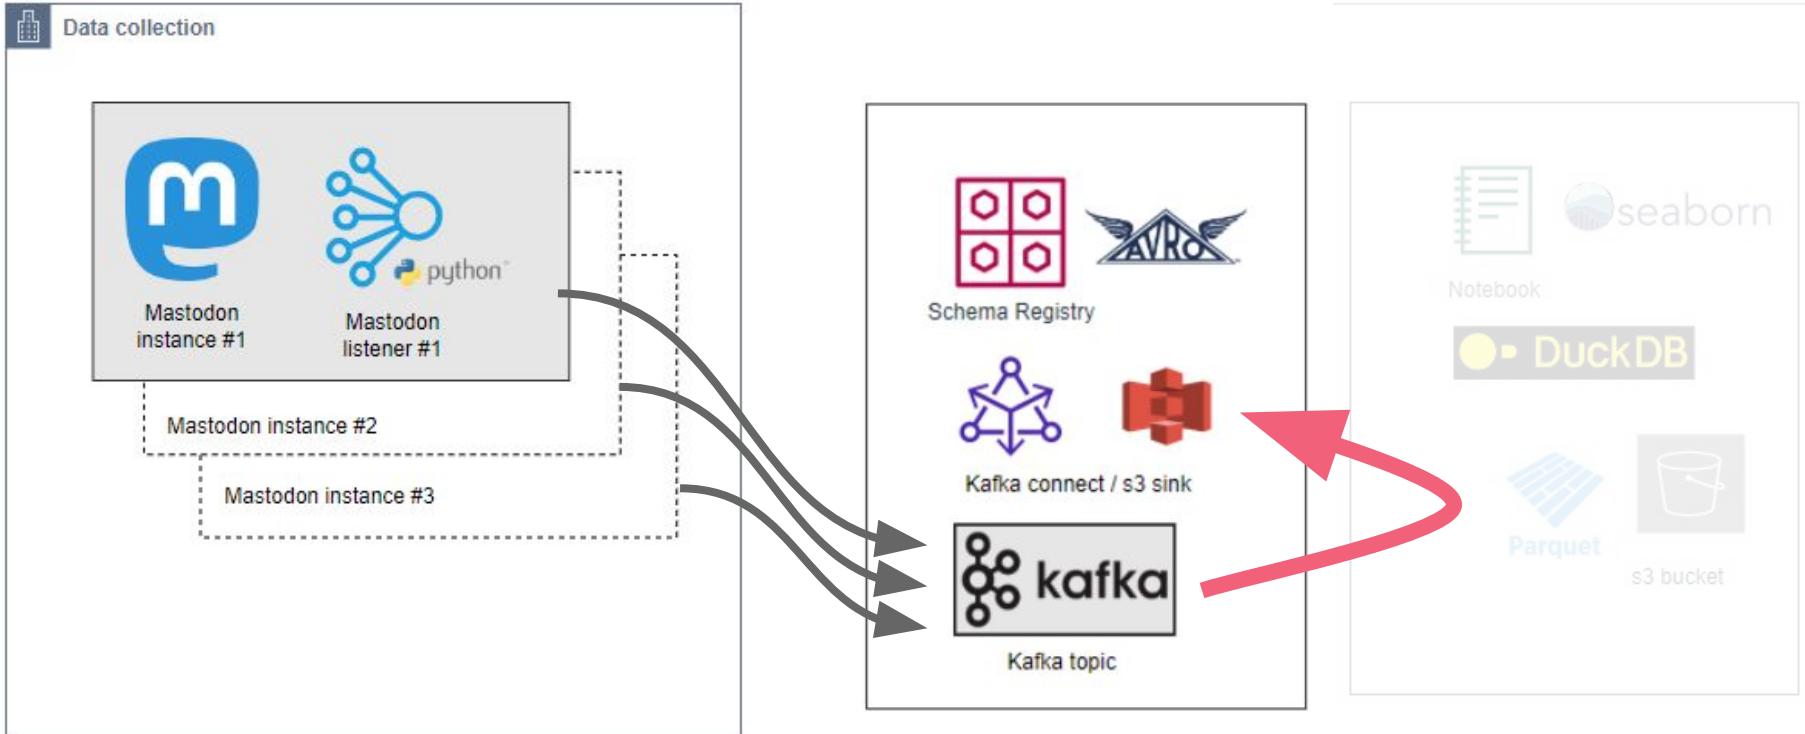
Image by [Pexels](#) from [Pixabay](#)

Data storage

Data collection



Store to s3

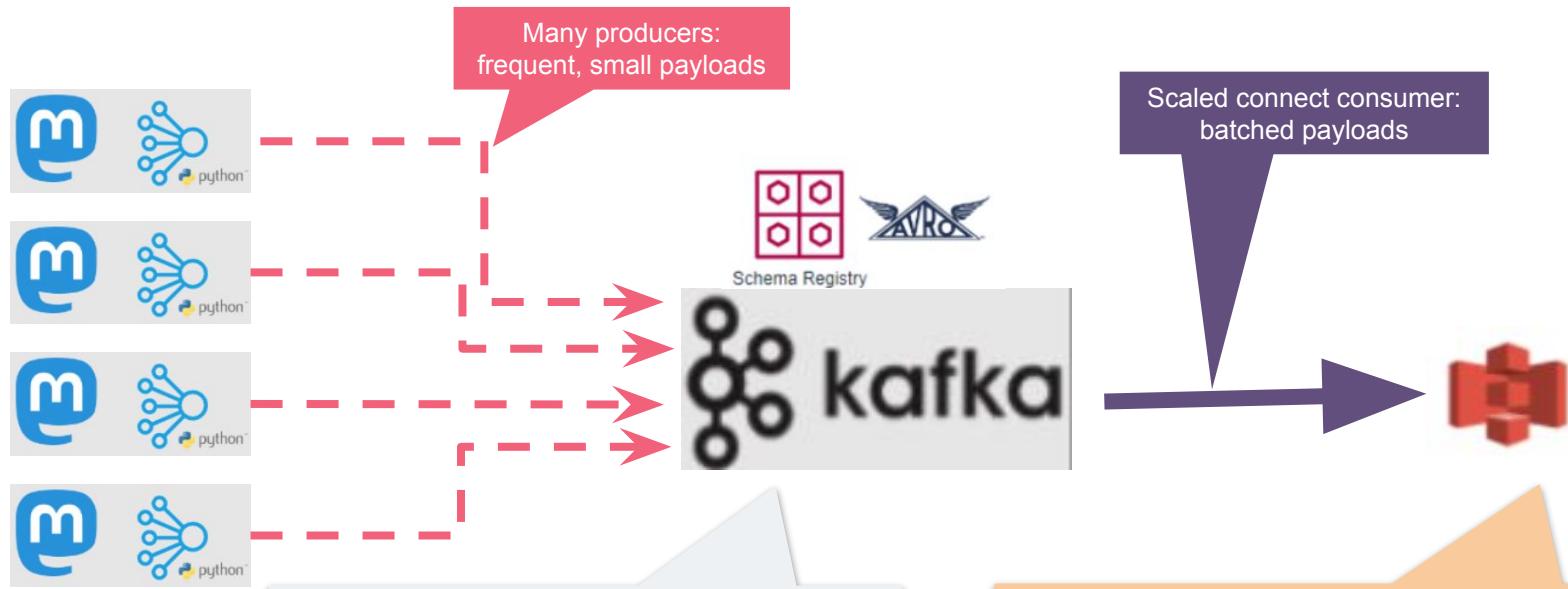


Kafka to s3

Kafka connect / s3 sink

```
{  
  "Connector name": "mastodon-sink-s3-aws",  
  "S3 sink class": "io.confluent.connect.s3.S3SinkConnector",  
  "Kafka topic": "mastodon-topic",  
  "Write as parquet": "true",  
  "S3 bucket details": {  
    "name": "2023mastodon",  
    "access.key.id": "XXXXXXXXXXXXXX",  
    "secret.access.key": "XXXXXXXXXXXXXXXXXXXXXX",  
    "region": "us-east-1",  
    "storage.class": "io.confluent.connect.s3.storage.S3Storage"  
  }  
}
```





mastodon-topic

- Overview
- Messages
- Schema
- Configuration

Producers
Bytes in/sec: 1.02K

Consumers
Bytes out/sec: 2.01K

Message fields

- topic
- partition
- offset
- timestamp
- timestampType
- headers
- key
- value

+ Produce a new message to this topic

Value Header Key

```

1  {
2   "id": 189793213642298359,
3   "account": "saubury",
4   "uri": "https://data-folks.masto.host/@saubury/18979321354757539",
5   "base_url": "https://mastodon.social",
6   "uri_normalized": "https://mastodon.social/@saubury/18979321354757539",
7   "favorited": 0,
8   "username": "saubury",
9   "url": "https://mastodon.social/@saubury/18979321354757539",
10  "tagged": 0,
11  "character": 59,
12  "reposts": 0,
13  "mentions": 0,
14  "text": "% GitHub just rolled out support for Mastodon profiles! % 😊"
}

```

data-folks.masto.host/@saubury

Amazon S3 > Buckets > 2023mastodon > topics/ > mastodon-topic/ > partition=0/

partition=0/

Objects (999+)

To enable sorting in the table below, use the search to reduce the size of the list to 999 objects or fewer.

Name	Type	Last modified	Size
masodon-topic=0+000000000000.json.parquet	parquet	February 10, 2023, 15:30:45 UTC<11:00	189.7 KB
masodon-topic=0+000000000001.json.parquet	parquet	February 10, 2023, 15:31:01 UTC<11:00	185.5 KB
masodon-topic=0+000000000002.json.parquet	parquet	February 10, 2023, 15:40:13 UTC<11:00	177.2 KB
masodon-topic=0+000000000003.json.parquet	parquet	February 10, 2023, 15:45:18 UTC<11:00	184.3 KB
masodon-topic=0+000000000004.json.parquet	parquet	February 10, 2023, 15:57:29 UTC<11:00	186.5 KB
masodon-topic=0+000000000005.json.parquet	parquet	February 10, 2023, 16:04:27 UTC<11:00	185.5 KB
masodon-topic=0+000000000006.json.parquet	parquet	February 10, 2023, 16:12:38 UTC<11:00	190.5 KB
masodon-topic=0+000000000007.json.parquet	parquet	February 10, 2023, 16:12:20 UTC<11:00	197.0 KB
masodon-topic=0+000000000008.json.parquet	parquet	February 10, 2023, 16:29:56 UTC<11:00	185.4 KB
masodon-topic=0+000000000009.json.parquet	parquet	February 10, 2023, 16:36:33 UTC<11:00	182.5 KB
masodon-topic=0+000000000010.json.parquet	parquet	February 10, 2023, 16:46:24 UTC<11:00	185.5 KB

DuckDB



data-folks.masto.host/@saubury

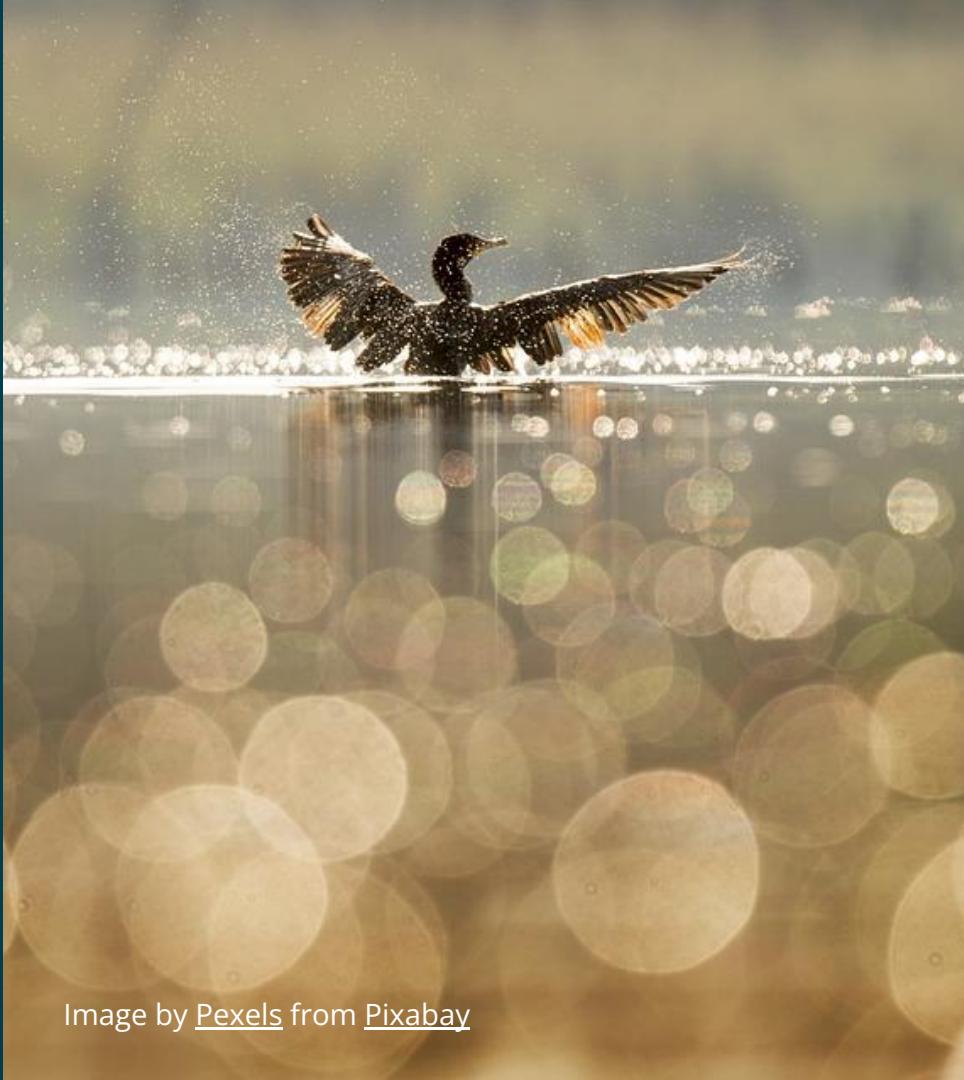
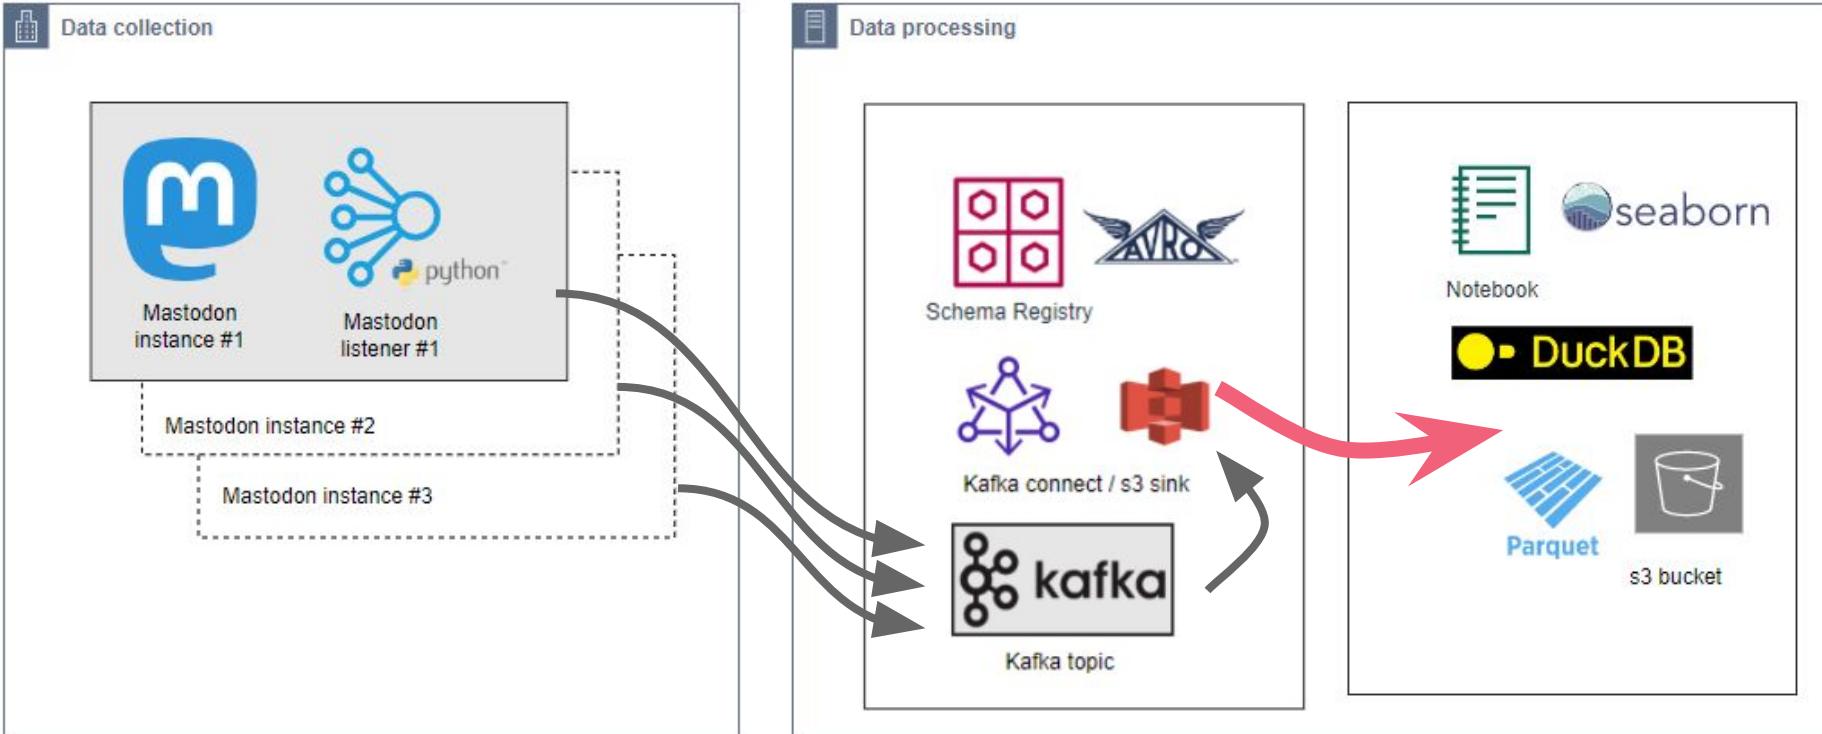


Image by [Pexels](#) from [Pixabay](#)

Parquet in DuckDB





DuckDB is free / open-source

https://duckdb.org/why_duckdb.html



RDBMS / SQL / OLAP - Extensive SQL with a large function library, window functions etc. ACID transactional



No external dependencies - compiled into two files, embedded within a host process. DuckDB uses **PostgreSQL's SQL parser**, Google's RE2 regular expression engine and SQLite's shell

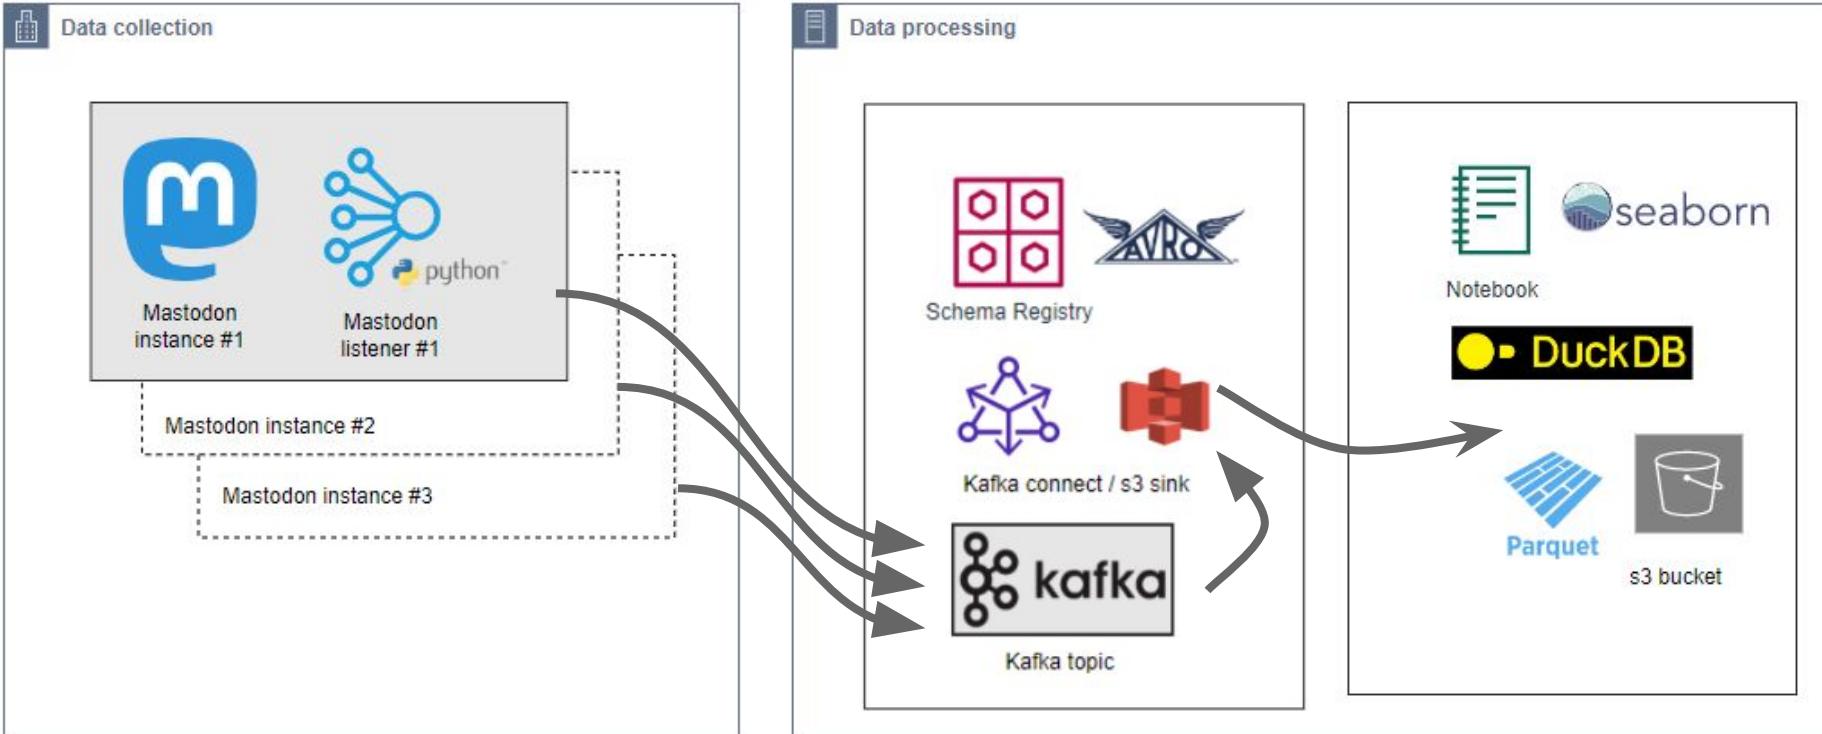


Process foreign data without copying, run queries directly on Pandas data, Python and R, APIs for Java, C, C++



Extensions - Httpfs, s3, parquet, json, FTS, geospatial

Parquet in DuckDB



DuckDB SQL

Start DuckDB

Install extensions

Set s3 endpoint

Select parquet

CLI / macOS

```
./duckdb
```

```
INSTALL httpfs;  
LOAD httpfs;
```

```
set s3_access_key_id='XXXXXX';  
set s3_secret_access_key='XXXXXXXXXXXXXXXXXXXX';  
set s3_region='us-east-1';
```

```
select *  
from  
read_parquet('s3://mastodon/mastodon-topic/partition=0/*');
```



DuckDB SQL

1.6 million

```
create table mastodon_toot
as
select m_id, created_at, app, url, username, mastodon_text
from read_parquet('s3://mastodon-topic/partition=0/*');
```

SQL



Data Analysis



data-folks.masto.host/@saubury



Photo by [Isaac Smith](#) on [Unsplash](#)

Daily Mastodon Usage

```
select strftime(created_tz, '%Y/%m/%d %a')
, count(*) as "Num toots"
, count(distinct(username)) as "Num users"
, count(distinct(from_instance)) as "Num urls"
, mode(case when bot='False' then username end)
, mode(case when bot='True' then username end)
, mode(base_url) as "Most freq host"
from mastodon_toot
group by 1 order by 1;
```

[SQL / Notebook](#)

200,000 toots a day
from 50,000 users,
4,000 servers

**mastodon.social was the
most popular host**

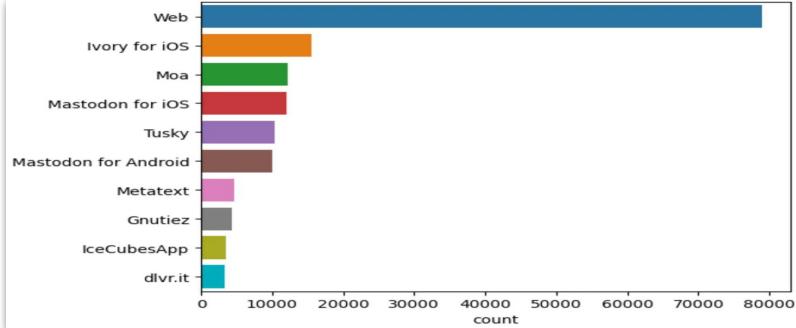
News organisations are
the biggest generator of
content

Created day	Num toots	Num users	Num urls	Most freq non-bot	Most freq bot	Most freq host
2023/02/03 Fri	17880	8537	1524	gnutiez	nieuws	https://mastodon.social
2023/02/04 Sat	210646	54006	4562	gnutiez	cnexnews	https://mastodon.social
2023/02/05 Sun	191391	49241	4310	IzumiHal	ua	https://mastodon.social
2023/02/06 Mon	41632	17846	2255	gnutiez	nieuws	https://mastodon.social
2023/02/07 Tue	99097	30701	3350	gnutiez	cnexnews	https://mastodon.social
2023/02/08 Wed	188503	49649	4372	gnutiez	cnexnews	https://mastodon.social
2023/02/09 Thu	166096	48532	4227	worldeconomicfella	cnexnews	https://mastodon.social
2023/02/10 Fri	207877	54230	4608	gnutiez	cnexnews	https://mastodon.social

Mastodon App Landscape

```
%%sql  
mastodon_app_df <<  
    select *  
    from mastodon_toot  
    where app is not null  
    and app <> ''  
    and bot='False';  
  
sns.countplot(data=mastodon_app_df, y="app")
```

SQL / Notebook



Mastodon application landscape is rapidly changing

Web usage is the preferred client, with mobile apps like Ivory, Moa, Tusky, and the Mastodon app

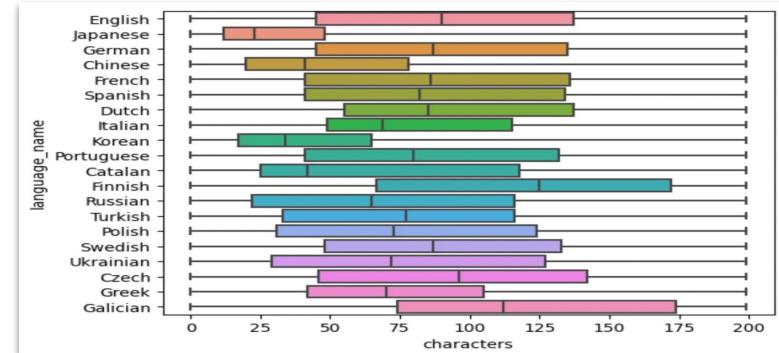
Generally, the app attribute does not federate



Toot Length by Language Usage

```
%%sql  
mastodon_lang_df <<  
    select *  
    from mastodon_toot  
    where language not in ('unknown');  
  
sns.boxplot(data=mastodon_lang_df,  
x="characters", y="language_name", orient="h")
```

SQL / Notebook



Chinese, Japanese, and Korean toots are shorter than English

Galicia and Finish messages are longer

Perhaps logographic languages (like Mandarin) convey more with fewer characters?

Longest toot ...

```
%%sql
with cte as
(
    select username
    , mastodon_text
    , len(mastodon_text) as len_toot
    , max(len(mastodon_text)) over () as max_len_toot
    from mastodon_toot
)
select username, mastodon_text
from cte
where len_toot = max_len_toot;
```

SQL / Notebook

68,991 characters



Marsey the Cat



Surprising statistics

Quick statistics from the data collected over **ten days**
(February 3 to February 12)

	1,622,149 Mastodon toots seen
	142,877 unique Mastodon users
	8,309 unique Mastodon instances, 131 languages seen
	Shortest toot is 0 characters, average toot length is 151 characters, and longest toot is 68,991 characters
	All toots 245,245,677 characters (over 1.6 million toots) in DuckDB's memory only 745.5MB
	Time it takes to calculate the above statistics in a single SQL query is 0.7 seconds

*NOT SO-*Surprising discoveries

Kafka - great for scale & resilience

Kafka ecosystem - flexible data distribution

DuckDB - super flexible tool for EDA

The internet is [fun | weird]



Thanks / questions?

👉 [Medium Blog](#)

🛠 [Code](#)

Ⓜ️ [data-folks.masto.host/@saubury](#)

🐦 [@SimonAubury](#)

🐘 [data-folks.masto.host/@saubury](#)



DATAENGBYTES KEY DATES



MARCH 20TH | *CFP Open*

MAY 31ST | *CFP Closed*

AUGUST 21ST | *Perth Conference*

AUGUST 25TH | *Sydney Conference*

AUGUST 29TH | *Brisbane Conference*

AUGUST 31ST | *Melbourne Conference*

🔗 [DataEngBytes 2023](#)