

Introduction to Time Series:

It's everywhere, and you can harness it.

By Kirk Roybal



Roadmap

- 01 What is time-series data?
- 02 Visualization
- 03 Time series in PostgreSQL

Time-series data
**is recording every datapoint
as a new entry**

Time-series data
**is recording the change
of your world**

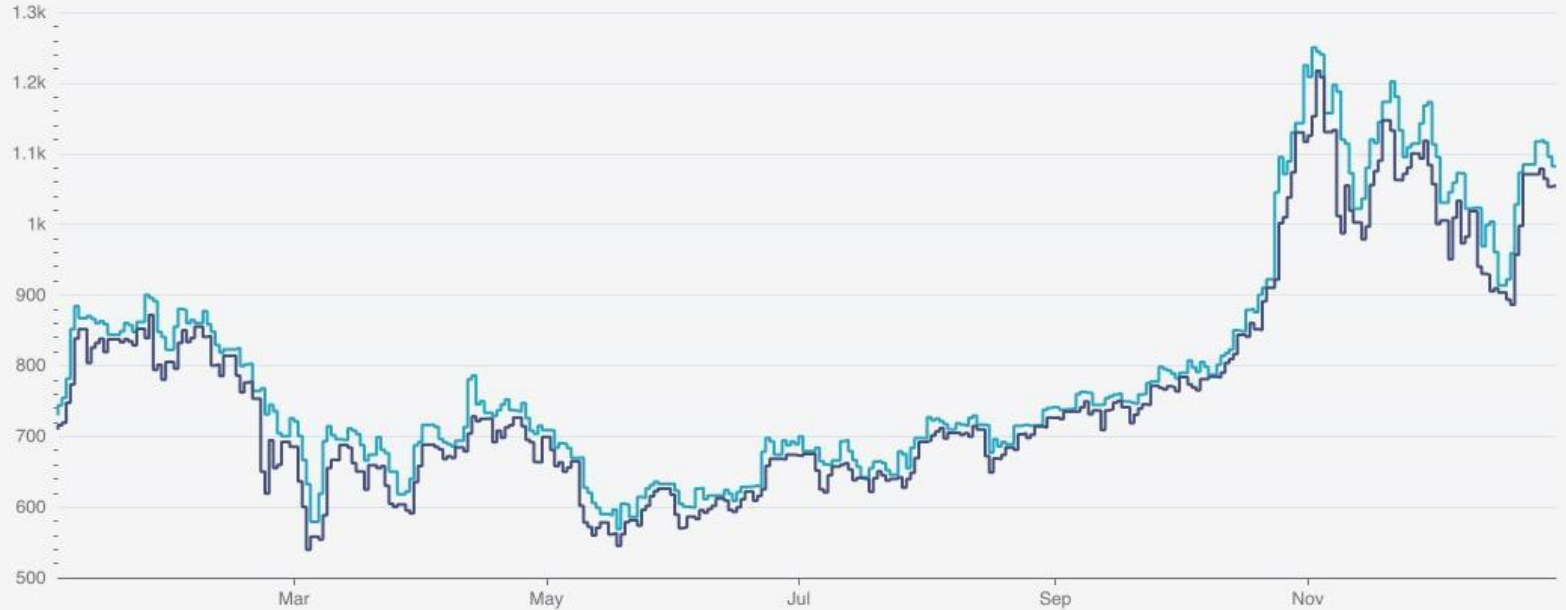
Time-series data
is everywhere



Daily Low vs High Price

2


Highest Price Lowest Price





Combined Ozone and PM2.5 Daily AQI Values

State College, PA



Source: U.S. EPA AirData <<https://www.epa.gov/air-data>>

Generated: January 28, 2022






New reported cases

All time


Last 90 days




Tests



Hospitalized



Deaths




New York Times: <https://www.nytimes.com/interactive/2021/us/covid-cases.html>



Possession team: KC

Play: (12:15) P.Mahomes pass deep middle to T.Hill for 73 yards, TOUCHDOWN. PENALTY on KC-T.Hill, Taunting, 15 yards, enforced between downs.



Data Intensive Applications

Existing databases aren't built for time-series

Relational Databases

Hard to scale efficiently, slower to query as data grows, storage is expensive, few time-series specific functions




PostgreSQL



Amazon RDS

Extend PostgreSQL for time-series

- Organize data into time-based chunks for faster ingest & access



looker

plotly

tableau

Grafana

MADlib

debezium

DATADOG

sensu

Works with all things PostgreSQL

OMNIDB



MODE


ZABBIX

kafka

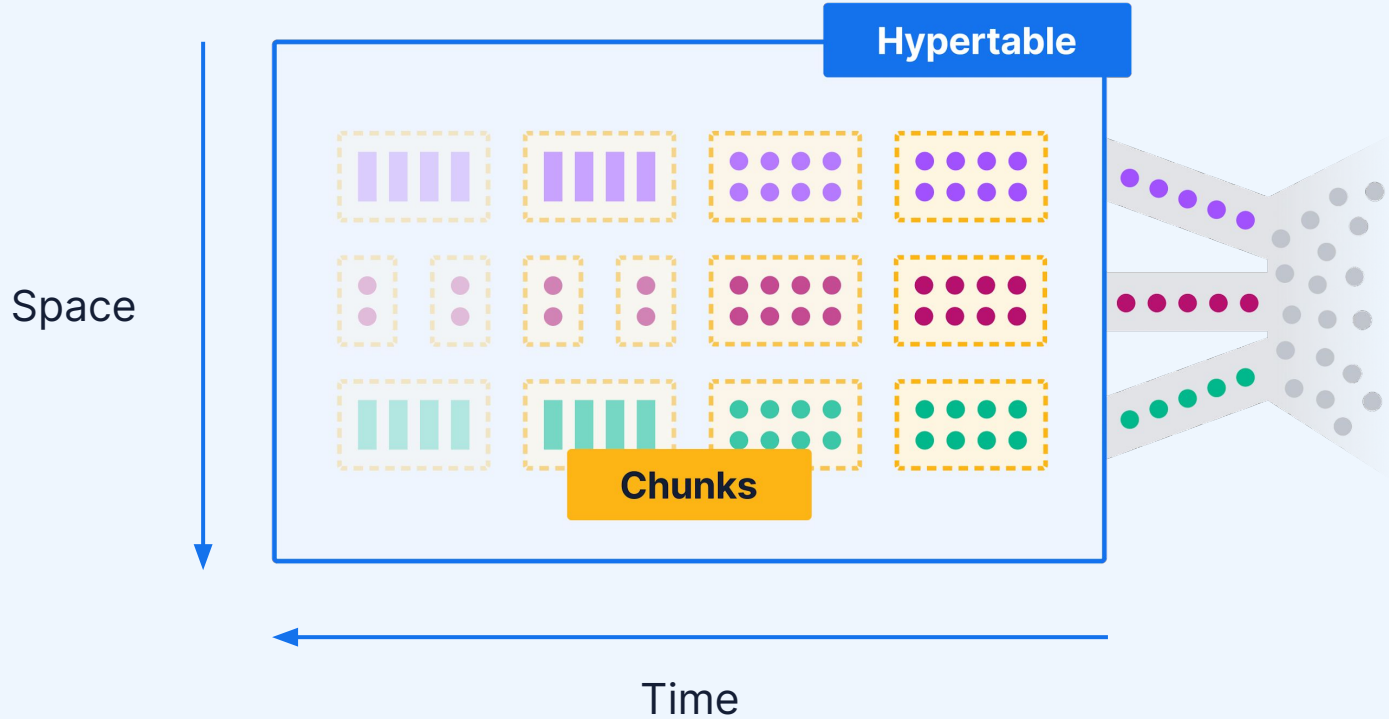
SQLAlchemy

CHARTIO

Time partitioning






Time and Space partitioning






Avoid querying chunks via constraint exclusion

```
SELECT time, device_id, temp FROM data  
WHERE time > now() - interval '24 hours'
```



Avoid global data structures and lookups

1. Select chunks **globally** via per-chunk constraints
2. Select data **locally** via per-chunk indexes
 - B-tree, Hash, GiST, SP-GiST, GIN and BRIN
 - Secondary and composite columns, UNIQUE* constraints




**Use partitions or you will
regret your life.**

Materialized Views



CREATE MATERIALIZED VIEW `one_hour_summary`
WITH ...

Oldest data  Hourly data



 Raw data



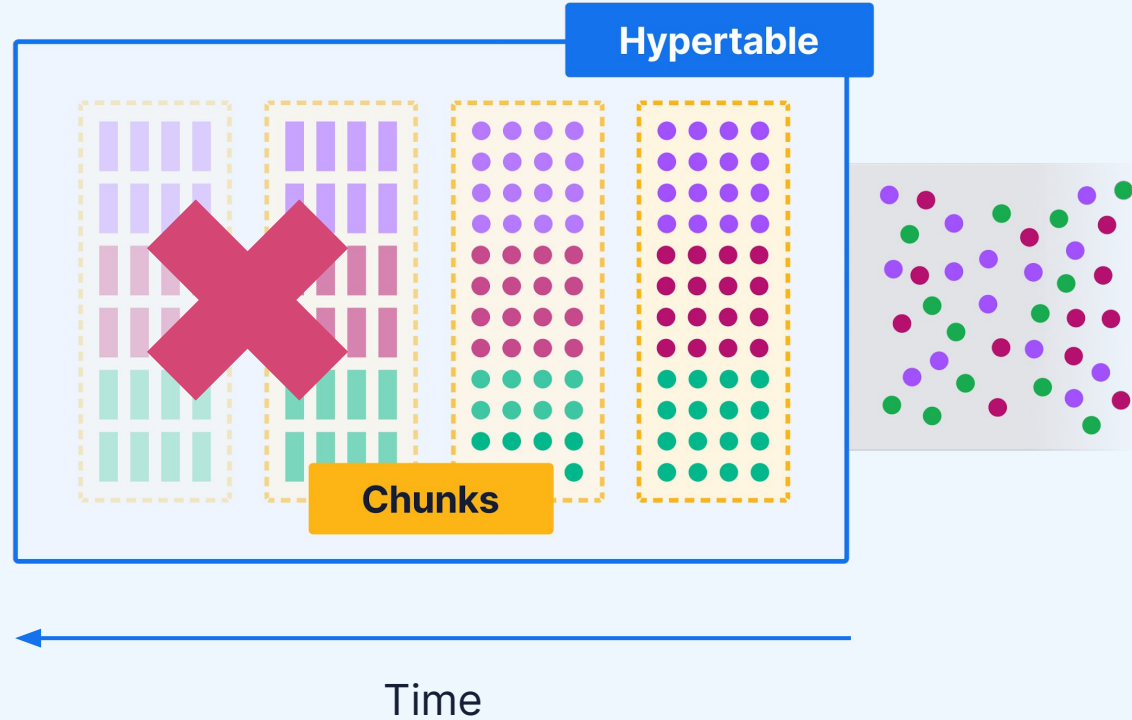


Schedule a refresh policy

```
REFRESH MATERIALIZED VIEW one_hour_summary;
```

Data Retention

Data retention





Try Timescale for free

<https://tsdb.co/try-timescale>

<https://docs.timescale.com>

<https://slack.timescale.com>