

Time Series Database

A.A. 2021/22

Matteo Nardelli

Laurea Magistrale in Ingegneria Informatica - II anno

The reference Big Data stack

High-level Interfaces

Data Processing

Data Storage


Resource Management

Support / Integration

Time Series

- Time series data are measurements (or events) collected over time
 - Usually taken at equally spaced data points
 - Discrete-time measurements

Andamento demografico della popolazione residente in **Italia** dal 2001 al 2020. Grafici e statistiche su dati ISTAT al 31 dicembre di ogni anno.




Andamento della popolazione residente

ITALIA - Dati ISTAT al 31 dicembre di ogni anno - Elaborazione TUTTITALIA.IT

(*) post-censimento

CPU Utilization



Store Time Series

- Where to store data?
- Flat files
 - Limited utility for time series, data will outgrow them, access is inefficient
- Relational DBMS
 - Possible scalability issues
- NoSQL
 - Better scaling, efficient queries based on time range
 - The design can be challenging:
 - Row key as time series ID, column as time offset ?
 - Wide table stores data point-by-point
 - Hybrid design with wide tables and blob (aggregation of points)

Time	Time series ID	Value
15:51:00	101	0.01
15:51:03	102	1.16
15:52:07	101	0.04
15:52:11	101	0.08
15:53:17	103	4.18

RDBMS

Time series ID	Time-window start time	Columns are named by sample time offset within time window				
		+0	+3	+7	+11	+17
101	15:51:00	0.01				
101	15:52:00		0.04	0.08		
102	15:51:00		1.16			
103	15:53:00					4.18

NoSQL

Time Series Database (TSDB)

- Optimized for time series data
 - Key notion: time
 - Database optimized for handling time-stamped data
 - Examples: application performance monitoring, network data, sensor data, events, clicks, trades in a market, and many other types of analytics data.
- Similar to a key/value store:
 - Where the key is a **timestamp**
 - The value is the **measurement**, which can have multiple **fields**.
- A TSDB is optimized for measuring changes over time
 - Considers data lifecycle management
 - Optimizes data storage: storage, compression, data retention, and sharding
 - Optimizes data query: time-aware queries, data aggregation, large range scan of records

Time Series Database (TSDB)

- Ranking of TSDBs on DB-Engines

☐ include secondary database models

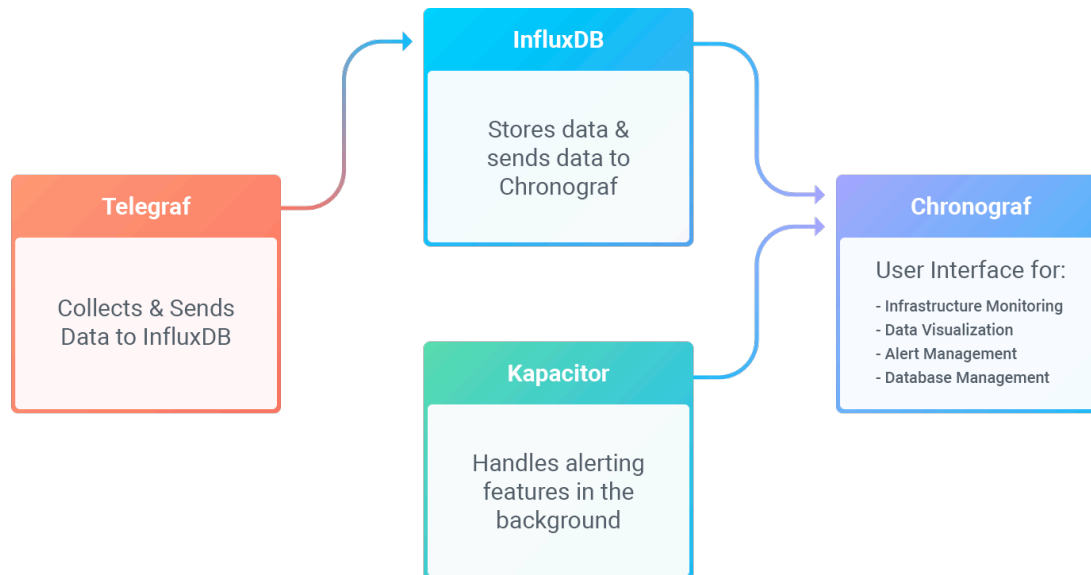
39 systems in ranking, April 2022

Rank			DBMS	Database Model	Score		
Apr 2022	Mar 2022	Apr 2021			Apr 2022	Mar 2022	Apr 2021
1.	1.	1.	InfluxDB	Time Series, Multi-model	30.02	+0.33	+3.47
2.	2.	2.	Kdb+	Time Series, Multi-model	8.78	-0.24	+0.92
3.	3.	3.	Prometheus	Time Series	6.31	-0.01	+0.58
4.	4.	4.	Graphite	Time Series	5.36	-0.11	+0.83
5.	5.	5.	TimescaleDB	Time Series, Multi-model	4.56	+0.10	+1.79
6.	6.	7.	Apache Druid	Multi-model	3.18	-0.08	+0.55
7.	7.	6.	RRDtool	Time Series	2.58	-0.03	-0.12
8.	8.	8.	OpenTSDB	Time Series	1.82	+0.03	+0.06
9.	9.	11.	DolphinDB	Time Series, Multi-model	1.62	+0.09	+0.78
10.	10.	9.	Fauna	Multi-model	1.42	+0.07	-0.10
11.	11.	10.	GridDB	Time Series, Multi-model	1.28	-0.07	+0.30
12.	12.	16.	QuestDB	Time Series, Multi-model	1.15	0.00	+0.75
13.	13.	14.	Amazon Timestream	Time Series	0.98	-0.04	+0.36
14.	14.		TDengine	Time Series, Multi-model	0.87	+0.11	
15.	15.	12.	eXtremeDB	Multi-model	0.69	-0.01	-0.06
16.	16.	13.	KairosDB	Time Series	0.65	+0.01	-0.08
17.	17.	21.	VictoriaMetrics	Time Series	0.57	-0.05	+0.25
18.	19.	19.	IBM Db2 Event Store	Multi-model	0.53	+0.01	+0.17
19.	18.	15.	Raima Database Manager	Multi-model	0.51	-0.06	-0.01
20.	20.	25.	Apache IoTDB	Time Series	0.41	+0.01	+0.23

<https://db-engines.com/en/ranking/time+series+dbms>

InfluxDB

- Natively built to manage time series data
- InfluxDB is part of an ecosystem that supports:
 - Collection: Telegraf
 - Storage: InfluxDB
 - Monitoring/Processing: Kapacitor
 - Visualization: Chronograf
 - Alerting



InfluxDB

- Data model:

measurement-name tag-set field-set timestamp

- Measurement: string (indexed)
- Tag-set: key-value pairs (indexed, only string allowed)
- Field-set: key-value pairs (values can be of numbers, booleans or strings)
- Timestamp: can have second, millisecond, microsecond, or nanosecond precision

Example: `cpu host=serverA,region=uswest idle=23,user=42,system=12 1464623548s`

- Data compression depends on the level of precision
- Data storage (on disk):
 - Data is organized in a columnar style format
 - Contiguous blocks of time are set for the measurement, tagset, field.
 - Each field is organized sequentially on disk for blocks of time, which make calculating aggregates on a single field a very fast operation.
 - There is no limit to the number of tags and fields that can be used.

InfluxDB

- InfluxDB creates a shard for each block of time, whose size depends on the retention policy:
 - describes how long to keep data (retention)
 - how many copies (replication factor)
 - the time range covered by shard groups (shard group duration)
- Each shard maps to an underlying database with
 - A WAL file:
 - Write-optimized storage format (durable writes, but not easily queryable)
 - Writes to the WAL are appended to segments of a fixed size.
 - A TSM file:
 - Contains sorted, compressed series data.
 - Read-only files that are memory mapped.

InfluxDB: Query and Manipulation Language

- **Flux** is InfluxData's functional data scripting language designed for querying, analyzing, and acting on data
- Flux supports multiple data source types, including:
 - Time series databases (such as InfluxDB)
 - Relational SQL databases (such as MySQL and PostgreSQL)
 - CSV
- Flux unifies code for querying, processing, writing, and acting on data into a single syntax.

InfluxDB: Flux

Like treating water, a Flux query does the following:

1. Retrieves a specified amount of data from a source.
2. Filters data based on time or column values.
3. Processes and shapes data into expected results.
4. Returns the result.

```
from(bucket: "example-bucket")           // — Source
  |> range(start: -1d)                     // — Filter on time
  |> filter(fn: (r) => r._field == "foo") // — Filter on column values
  |> group(columns: ["sensorID"])          // — Shape
  |> mean()                               // — Process
```

InfluxDB: Flux

```
from(bucket: "example-bucket")           // — Source
  |> range(start: -1d)                     // — Filter on time
  |> filter(fn: (r) => r._field == "foo") // — Filter on column values
  |> group(columns: ["sensorID"])          // — Shape
  |> mean()                               // — Process
```

- `from()` to retrieve data from the data source.
- pipe-forward operator (`|>`) to send the output of each function to the next function as input.
- `range()`, `filter()`, or both to filter data based on column values.
- `mean()` to calculate the average of values returned from the data source.
- `yield()` to yield results to the user.

InfluxDB: The Flux Data Model

The Flux data model comprises the following:

- Stream of tables: returned by data sources
- Table: collection of columns partitioned by group key
- Column: collection of values of the same basic type that contains one value for each row.
- Row: a collection of associated column values.
- Group key: key-value pairs, where each key represents a column name and each value represents the column value included in the table.
 - All rows in a table contain the same values in group keys. All tables in a stream of tables have a unique group key
 - e.g., each group key represents a table containing data for a unique location

```
[_measurement: "production", facility: "us-midwest", _field: "apq"]  
[_measurement: "production", facility: "eu-central", _field: "apq"]  
[_measurement: "production", facility: "ap-east", _field: "apq"]
```

InfluxDB: Flux

- The majority of basic Flux queries include the following steps: Source, Filter, Shape, Process
- Source:
 - Flux input functions retrieve data from a data source.
 - All input functions return a stream of tables.
 - Flux supports multiple data sources including, time series databases (such as InfluxDB and Prometheus), relational databases (such as MySQL and PostgreSQL), CSV, and more.
- Filter:
 - Filter functions iterate over and evaluate each input row to see if it matches specified conditions.
 - `range()`: filter data based on time.
 - `filter()`: use a predicate function (fn) to filter data based on column values.

<https://docs.influxdata.com/flux/v0.x/function-types/>

<https://docs.influxdata.com/flux/v0.x/get-started/syntax-basics/>

InfluxDB: Flux

- Shape data:
 - Queries may require to change the structure of data
 - Functions that reshape data include:
 - `group()`: modify group keys
 - `window()`: modify `_start` and `_stop` values of rows to group data by time
 - `pivot()`: pivot column values into rows
 - `drop()`: drop specific columns
 - `keep()`: keep scientific columns and drop all others
- Process:
 - Aggregate data: into a single row (e.g., `count()`, `mean()`, `sum()`, `quantile()`)
 - Select specific data points: return specific rows from each input table.
 - e.g., `distinct()`, `first()`, `last()`, `min()`, `max()`, `limit()`, `top()`, `unique()`
 - Rewrite rows: transform values (e.g., maths operations, process strings, add new columns)
 - e.g., `map()`, which by default drops not explicitly mapped columns, the `with` operator updates a column if it already exists and includes all existing columns: `map(fn: (r) => ({ r with newColumn: r._value * 2 }))`
 - Send notifications

Hands-on InfluxDB (Docker image)

InfluxDB with Docker

- We use the official standalone InfluxDB image

```
$ docker pull influxdb:2.0
```

- We can now create an instance of InfluxDB

```
$ docker run -p 8086:8086 \  
  -v $PWD:/var/lib/influxdb2 \  
  influxdb:2.0
```

- Chronograf can be reached at <http://localhost:8086>

Loading Sample Data


```
import "influxdata/influxdb/sample"

sample.data(set: "airSensor")
  |> to(
    org: "example-org",
    bucket: "example-bucket"
  )
```

_time	_value	_field	_measurement	sensor_id
2022-04-09 17:05:49 GMT+2	0.5108571955231641	co	airSensors	TLM0100
2022-04-09 17:05:59 GMT+2	0.5164828607775828	co	airSensors	TLM0100
2022-04-09 17:06:09 GMT+2	0.5086872079247073	co	airSensors	TLM0100
2022-04-09 17:06:19 GMT+2	0.5019140187838028	co	airSensors	TLM0100
2022-04-09 17:06:29 GMT+2	0.5061828461616715	co	airSensors	TLM0100
2022-04-09 17:06:39 GMT+2	0.5156546512843109	co	airSensors	TLM0100
2022-04-09 17:06:49 GMT+2	0.5149306168772608	co	airSensors	TLM0100
2022-04-09 17:06:59 GMT+2	0.5254830148241758	co	airSensors	TLM0100
2022-04-09 17:07:09 GMT+2	0.519322216384614	co	airSensors	TLM0100


Get Raw Temperature Data

```
from(bucket: "example-bucket")
|> range(start: v.timeRangeStart, stop: v.timeRangeStop)
|> filter(fn: (r) => r["_measurement"] == "airSensors")
|> filter(fn: (r) => r["_field"] == "temperature")
|> yield(name: "raw values")
```



Get Sensor's Temperature Data

```
from(bucket: "example-bucket")
|> range(start: v.timeRangeStart, stop: v.timeRangeStop)
|> filter(fn: (r) => r["_measurement"] == "airSensors")
|> filter(fn: (r) => r["_field"] == "temperature")
|> filter(fn: (r) => r["sensor_id"] == "TLM0100")
|> yield(name: "raw values")
```



Get Average Sensor's Temperature Data

```
from(bucket: "example-bucket")
|> range(start: v.timeRangeStart, stop: v.timeRangeStop)
|> filter(fn: (r) => r["_measurement"] == "airSensors")
|> filter(fn: (r) => r["_field"] == "temperature")
|> filter(fn: (r) => r["sensor_id"] == "TLM0100")
|> group(columns: ["sensorID"])
|> mean()
```

_start	_stop	_value	sensorID
2022-04-09 12:26:52 GMT+2	2022-04-09 18:26:52 GMT+2	71.20310898758906	

Since InfluxDB groups data by series, `mean()` returns a table for **each unique** `sensor_id` containing a single row with the average value in the `_value` column.

Here, since we query for a single sensor, we could have omitted the `group()` operation; however, if we are interested in computing the average value across all sensors, we really need to `group()` correctly data.

Get Average Temperature Data

```
from(bucket: "example-bucket")
|> range(start: v.timeRangeStart, stop: v.timeRangeStop)
|> filter(fn: (r) => r["_measurement"] == "airSensors")
|> filter(fn: (r) => r["_field"] == "temperature")
|> filter(fn: (r) => r["sensor_id"] == "TLM0100")
|> group(columns: ["sensorID"])
|> mean()
```

_start	_stop	_value	sensorID
2022-04-09 12:26:52 GMT+2	2022-04-09 18:26:52 GMT+2	71.20310898758906	

Not filtering on a single sensor_id, group() allows to compute the mean value across all sensors

_start	_stop	_value
2022-04-09 12:31:07 GMT+2	2022-04-09 18:31:07 GMT+2	72.86225604580238