

Homework 1

Data Science and Database Technology

Data analysts of the National Association of Italian Museums are interested in analyzing the average revenue per ticket. In particular, they would like the analyses to address the following features.

A museum has a unique name, and it is located in a specific city. The province and region are also stored. The same city can host different museums. Each museum belongs to a specific category (e.g., “Art”, “Historic sites”, “Natural history”).

A museum may have some additional services available for its public. The systems record which services are available for each museum. Examples of additional services are “guided tours”, “audio guides”, “wardrobe”, “café”, “Wi-Fi”. The number of additional services is 10 and their complete list is known.

The tickets sold by each museum are recorded. There are 3 different types of tickets: “Full Revenue”, “Reduced-student” (for students from 14 to 24 years old) and “Reduced-junior” (for young people less than 14 years old).

The systems also store how the ticket is purchased. A ticket can be purchased in three modalities: online, in authorized ticket offices, or directly at the entrance of the museum.

The analyses must be carried out considering the date, month, bimester, trimester, semester, year, if the date is a working day or a holiday, and time slot of the ticket validity date. The time slot is stored in 3 ranges of 4-hour blocks (08:00-12:00, 12:01-16:00, 16:01-20:00).

The company is interested in the statistics on the average revenue per ticket. The analysis must be carried on based on:

- museum name, museum category, city, province, region
- museum services
- Ticket type (full-Revenue, reduced-student, reduced-junior)
- Purchase modality (online, in authorized ticket offices, or at the museum entrance)
- Ticket validity date, month, bimester, trimester, semester, working day, and time slot

Homework tasks

1. Design the data warehouse to address the specifications and to efficiently answer to all the provided frequent queries. Draw the conceptual schema of the data warehouse and the logical schema (fact and dimension tables).

2. Write the following frequent queries using the extended SQL language.
 - 2.1. Separately for each ticket type and for each month (of the ticket validity), analyze: the average daily revenue, the cumulative revenue from the beginning of the year, the percentage of tickets related to the considered ticket type over the total number of tickets of the month

 - 2.2. Considering the ticket of 2021. Separately for each museum and ticket type analyze: the average revenue for a ticket, the percentage of revenue over the total revenue for the corresponding museum category, assign a rank to the museum, for each ticket type, according to the total number of tickets in decreasing order.

3. Create and update a materialized view with **CREATE MATERIALIZED VIEW** and **CREATE MATERIALIZED VIEW LOG** in ORACLE

Consider the following frequent queries of interest:

- Analyze the average monthly revenue related to each ticket type and for each semester.
- Separately for each ticket type and for each month analyze the cumulative revenue from the beginning of the year.
- Considering only the tickets purchased online, separately for each ticket type and for each month analyze the total number of tickets, the total revenue and the average revenue.
- Separately for each ticket type and for each month analyze the total number of tickets, the total revenue and the average revenue for year 2021.
- Analyze the percentage of tickets related to each ticket type and month over the total number of tickets of the month for each ticket type.

- 3.1. Define a materialized view with CREATE MATERIALIZED VIEW useful to reduce the response time of the reported frequent queries.
 - 3.2. Define the materialized view logs with CREATE MATERIALIZED VIEW LOG for each table where you deem it necessary. For which tables is it useful to keep track of logs? Identify *all* and *only* the necessary tables. Furthermore, for each table identify *all* and *only* the attributes for which it is necessary to keep track of the variations.
 - 3.3. Specify which operations (e.g. INSERT on a specific table) cause an update of the defined MATERIALIZED VIEW
4. Update and management of views via Trigger
- Assuming that the CREATE MATERIALIZED VIEW command is not available, create the materialized views defined in the previous exercise and define the update procedure starting from changes on the fact table created by means of a trigger
- 4.1 Create the structure of the materialized view with CREATE TABLE VM1 (...)
 - 4.2 Specify an example of statement to populate the VM1 table with the necessary records using the statement

```
INSERT INTO VM1 (...)  
( SELECT ...  
  ... )
```
 - 4.3 Write the triggers necessary to propagate the changes (insertion of a new record) made in the FACTS table to the materialized view VM1.
 - 4.4 Specify which operations (e.g. INSERT) trigger the trigger created in 4.3.