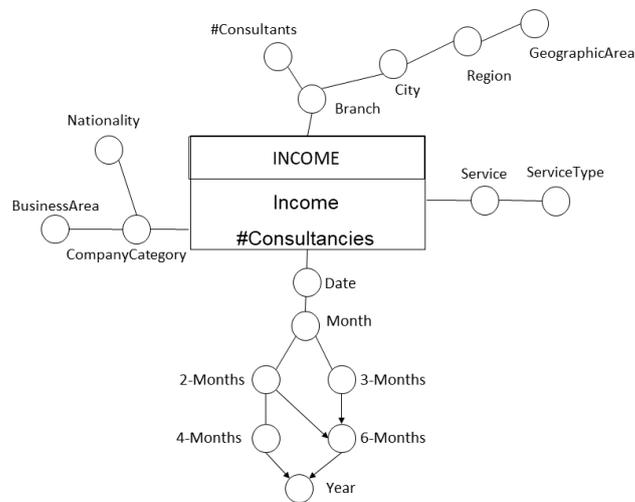


Exercise on materialized views

A data mart is characterized by the following conceptual schema and corresponding logical schema



INCOME (BranchID, ServiceID, CompanyCategoryID, TimeID, #Consultancies, Income)

SERVICE (ServiceID, Service, ServiceType)

TIME (TimeID, Date, Month, 2M, 3M, 4M, Semester, Year)

CONSULTANTS_BRANCH (BranchID, Branch, City, Region, GeographicArea, #Consultants)

COMPANY (CompanyCategoryID, CompanyCategory, BusinessArea, Nationality)

Point 1

Given the above logic schema, consider the following queries of interest:

- For each pair (type of services, semester), show the total income and the total number of consultancies carried out by the consultants of the branches located in the Italian region Lombardia.
- Considering only Italian and German companies, show for each pair (Region of consultants' branch, Service), the total income and number of consultancies, separately for each year.
- Considering only the income of years 2017, 2018, and 2019, for each pair (Type of services, Nationality of the company), show the six-months income and the average six-months income per consulting.

Considering the above information, define a materialized view with CREATE MATERIALIZED VIEW, to reduce the response time of queries (a) to (c) above. In particular, specify the SQL query associated with Block A in the following statement:

```
CREATE MATERIALIZED VIEW ViewIncome
BUILD IMMEDIATE
REFRESH FAST ON COMMIT
AS
    Block A
```

Point 2

Specify the minimal combination of attributes that constitutes an identifier for the materialized ViewIncome view.

Point 3

Assume that the ViewIncome materialized view is managed **without using** the CREATE MATERIALIZED VIEW statement. The derived table that implements the materialized view is defined by the following SQL statement

```
CREATE TABLE ViewIncome
(
  Service          VARCHAR(20),
  ServiceType      VARCHAR(20),
  Nationality      VARCHAR(20),
  Semester         VARCHAR(20),
  Year             VARCHAR(20),
  Region           VARCHAR(20),
  TotIncome        INTEGER
                  CHECK (TotIncome IS NOT NULL and TotIncome >0),
  TotNumConsulting INTEGER
                  CHECK (TotNumConsulting IS NOT NULL and TotNumConsulting >0),
  PRIMARY KEY (Nationality, Semester, Region, Service)
)
```

Write the INSERT statement for the first data loading into the ViewIncome materialized view, using the INSERT (SELECT ...) statement.

Point 4

Assume that the updating of the materialized view is done using triggers. Write the trigger to propagate the changes to the ViewIncome materialized view when a new record is inserted in the INCOME fact table.

Point 5

Assume that the updating of the materialized view is done using triggers. Write the trigger to propagate the update of the ServiceType attribute in the SERVICE table to the ViewIncome materialized view. Note: Consider the Type 1 mode for time management.

Point 6

The VM1 materialized view is built with the following SQL statement

```
CREATE MATERIALIZED VIEW
BUILD IMMEDIATE
REFRESH FAST ON DEMAND
ENABLE QUERY REWRITE
AS <Block A>
```

For the definition of Block A see Point 1.

Write the instructions that define the MATERIALIZED VIEW LOG in Oracle needed for the automatic FAST update of the ViewIncome materialized view. Indicate **all and only** the necessary logs and within each log definition indicate **all and only** the necessary attributes.