

**Big Data Mangement - A.A. 2013/14**  
Written Exam of **January 28, 2014**

**Exercise 1**

Describe one-, two-, and three-layer architectures for data warehousing and expose the main advantages and drawbacks of each such architecture.

**Exercise 2**

(a) Write an RDF/RDFS model representing the following statements about URIs City, Train, ESTrain, DepartsFrom, ArrivesIn, DepartsAt, xsd:integer, Rome, Milan, ES730.

1. City, Train, and ESTrain are classes;
2. ESTrain is a subclass of Train;
3. DepartsFrom, ArrivesIn, and DepartsAt are properties;
4. the domain of DepartsFrom, ArrivesIn and DepartsAt is Train;
5. the range of DepartsFrom and ArrivesIn is City;
6. the range of DepartsAt is xsd:integer;
7. Rome and Milan are Cities;
8. ES730 is an ESTrain that DepartsFrom Rome at 14:30 and ArrivesIn Milan.

(b) Write SPARQL queries corresponding to the following requests: (b1) “return all trains from Rome to Milan”; (b2) “return all trains that depart from Rome after 13:00”.

**Exercise 3**

Describe the column family data model and make it clear why column family stores are classified as aggregate databases. Provide a small example together your description.

**Exercise 4**

Provide the definition of “eventually consistent” and explain the importance of this notion in a distributed environment.