

DI/FCT/UNL
Mestrado Integrado em Engenharia Informática

Cloud Computing Systems
1st Semester, 2019/2020

Midterm Test (26/October/2019)

- 1) "One day, all (business) applications will run in (public) cloud platforms".
a) Give reasons why the previous sentence could become true.

- b) Give reasons why the previous sentence will never be true.

- 2) Vendor lock-in (i.e. the difficulty of changing of cloud provider) is a problem in the cloud computing market? Justify your answer.

Yes, because... / No, because... (select you answer)

3) App Service systems, such as Azure App Service, typically support scaling an application by scale out (run in more machines) and scale up (run in more powerful machines). Discuss the advantages/problems of each approach.

4) Azure Blob Storage supports three blob types. Explain the difference between block blobs and append blobs and give an example that is better supported with each type of blob.

5) Azure Cosmos DB supports different consistency models: strong consistency, bounded staleness, session, consistent prefix and eventual.

a) Assuming a geo-replicated deployment, what are the advantages of selecting a weaker consistency model (for an application).

b) "An application that uses the strong consistency level does not need to use the conditional update mechanism". State if you agree or disagree with this sentence and justify.

True, because... | False, because... (select the appropriate)

c) The recovery point objective of strong consistency is 0 for a deployment in multiple regions with a single master. Discuss what are the implications of this requirement for the implementation of the system and whether the implication could be relaxed by making the value of RPO larger.

Note: **Recovery point objective (RPO)** is the period of time for which updates might get lost in a failure.

- 6) Consider the following example of the use of Memcached (from the slides). Explain why the data is stored in Memcached with the key "userrow:" + **userid** and not only **userid**.

```
function get_foo(int userid) /* first try the cache */
    data = memcached_fetch("userrow:" + userid)
    if not data /* not found : request database */
        data = db_select("SELECT * FROM users WHERE userid = ?",
                        userid)
        /* then store in cache until next get */
        memcached_add("userrow:" + userid, data)
    return data
```

- 7) Consider the project you are developing in the course – a clone of reddit.
- a) Give two examples of data that you can cache in Azure Redis for Cache, the first in which there is no advantage of using Redis over Memcached and the other in which using Redis is better. Justify briefly.

b) For making your system available in the web, it would be necessary to develop a web interface – web interfaces typically consist of HTML pages, Javascript files and images.

For making these resources available to applications, they would have to be stored in one of the following alternatives:

- in the war files that is deployed to the App Service service, being served as files in a normal web site;
- on Azure blob storage;
- in Cosmos DB.

Discuss what would be the best alternative to store each of these resources. Justify.

