

Carnegie Mellon University in Qatar  
Distributed Systems

15-440 - Fall 2023

Mock Quiz 1

Name: \_\_\_\_\_

Andrew ID: \_\_\_\_\_

**Total time:** 45 minutes

**Instructions:**

- Write your answers in the spaces provided below the problems. If you make a mess, clearly indicate your final answers.
- This quiz has 8 questions over 5 pages, for a total of 25 points.
- Keep up with time.

Question	Points	Score
1	6	
2	3	
3	2	
4	2	
5	2	
6	4	
7	3	
8	3	
Total:	25	

6pts
------

1. Answer the following questions by selecting **True** or **False**:

- (a) (**True** / **False**) RPC allows passing parameters by value only.
- (b) (**True** / **False**) Marshalling and un-marshalling are performed within the stubs and skeletons of RPC.
- (c) (**True** / **False**) A workable distributed file system can be built without incorporating a naming service.
- (d) (**True** / **False**) The distributed Bellman-Ford algorithm uses essentially a flat naming service protocol for locating machines over the Internet.
- (e) (**True** / **False**) Blockchain uses an unstructured peer-to-peer architecture.
- (f) (**True** / **False**) Broadcasting is an effective naming service especially in WAN settings.
- (g) (**True** / **False**) An approach to locating mobile entities is to use forwarding pointers, whereby an entity moving from location A to location B can leave behind a server stub (or a skeleton) to its new location at B.
- (h) (**True** / **False**) No fault-tolerance measures need to be taken in RPC if it is layered on top of TCP.
- (i) (**True** / **False**) The at-least-once semantic in RPC can only be used with idempotent operations.
- (j) (**True** / **False**) Latency is a measure of throughput.
- (k) (**True** / **False**) If a server crashes before all the actions of a non-idempotent operation are done, the system has to redo all the actions that were executed before the crash.
- (l) (**True** / **False**) It is better to use synchronous RPC for operations like training a deep learning model (or what is referred to as batch processing).

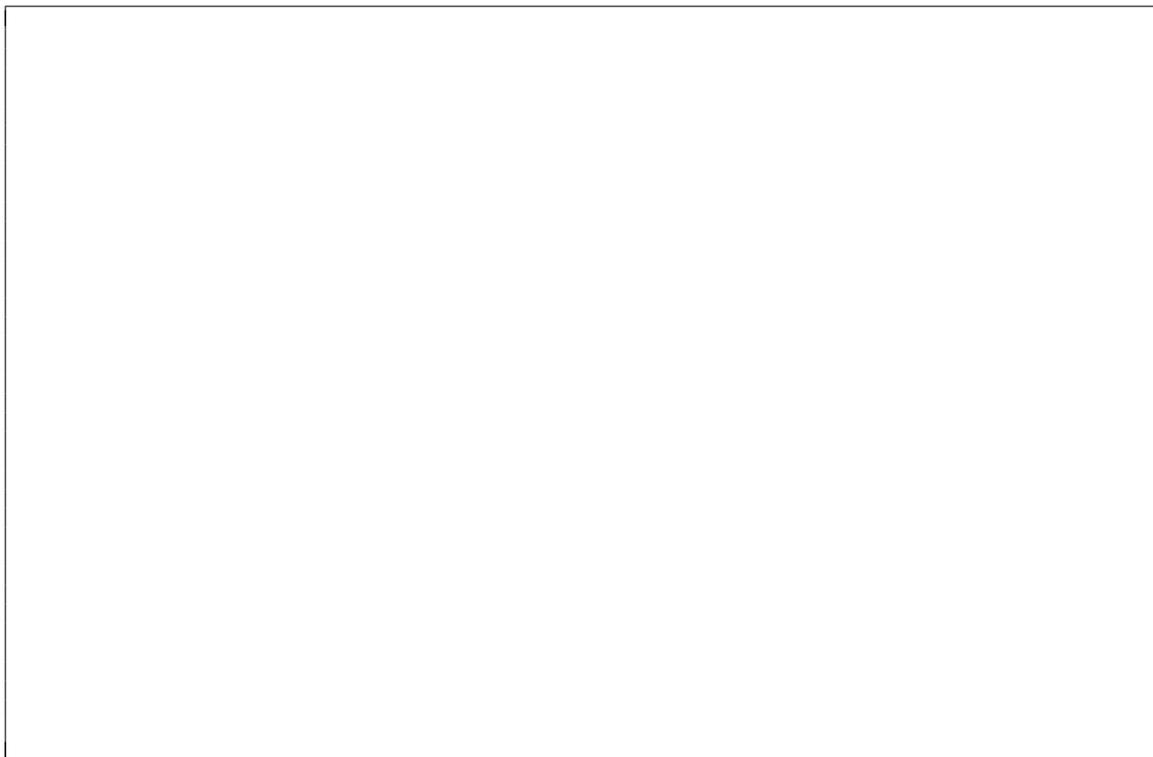
*Mock Quiz continues on the next page(s)*

- 3pts 2. What is the difference between *layered* and *tiered* architectures? Give an example of when you would use a layered architecture, but not a tiered one?

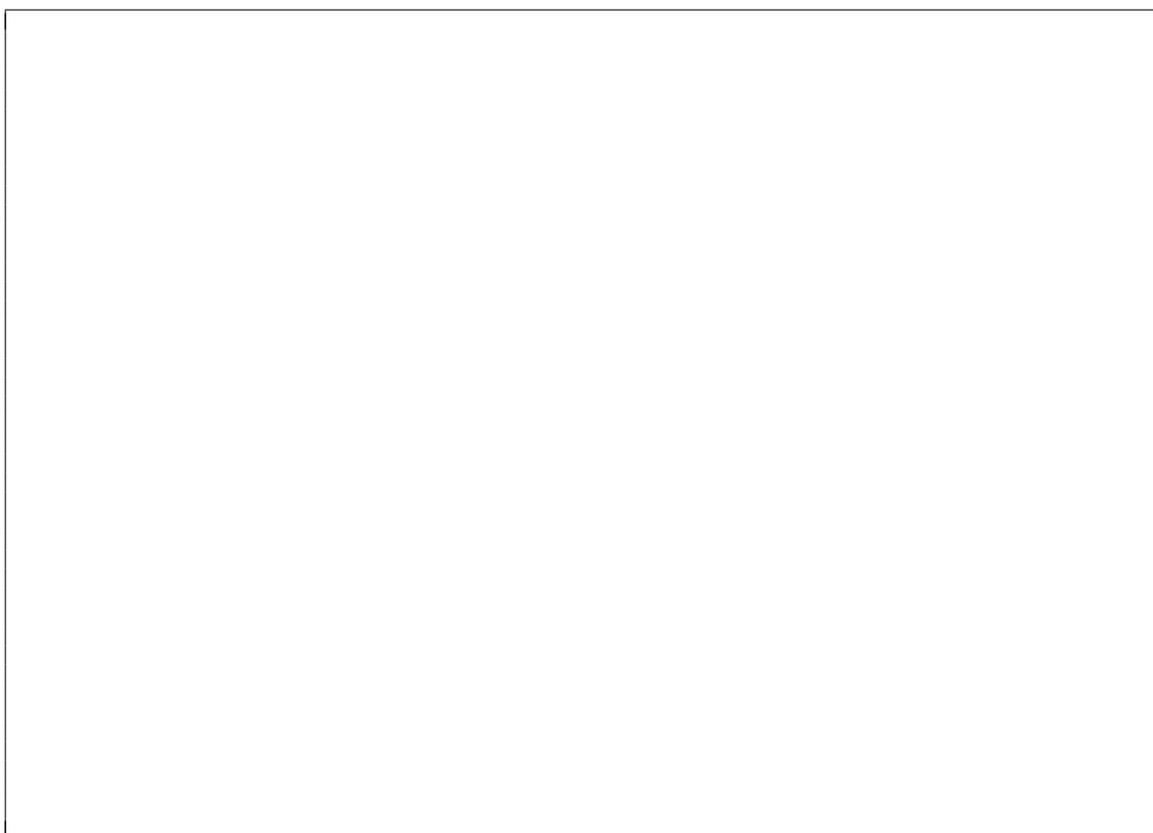
- 2pts 3. You have been asked to design and implement a distributed system for video streaming, which requires high Quality-of-Service. Would you use TCP or UDP for your middleware implementation? Justify your answer.

- 2pts 4. What is the weakest RPC semantic (i.e., *exactly-once*, *at-most-once*, or *at-least-once*) that you would suggest for requesting a taxi through *Uber*. Discuss.

- 2pts 5. Why are marshaling and unmarshaling important in exchanging data between communicating entities in a distributed system?

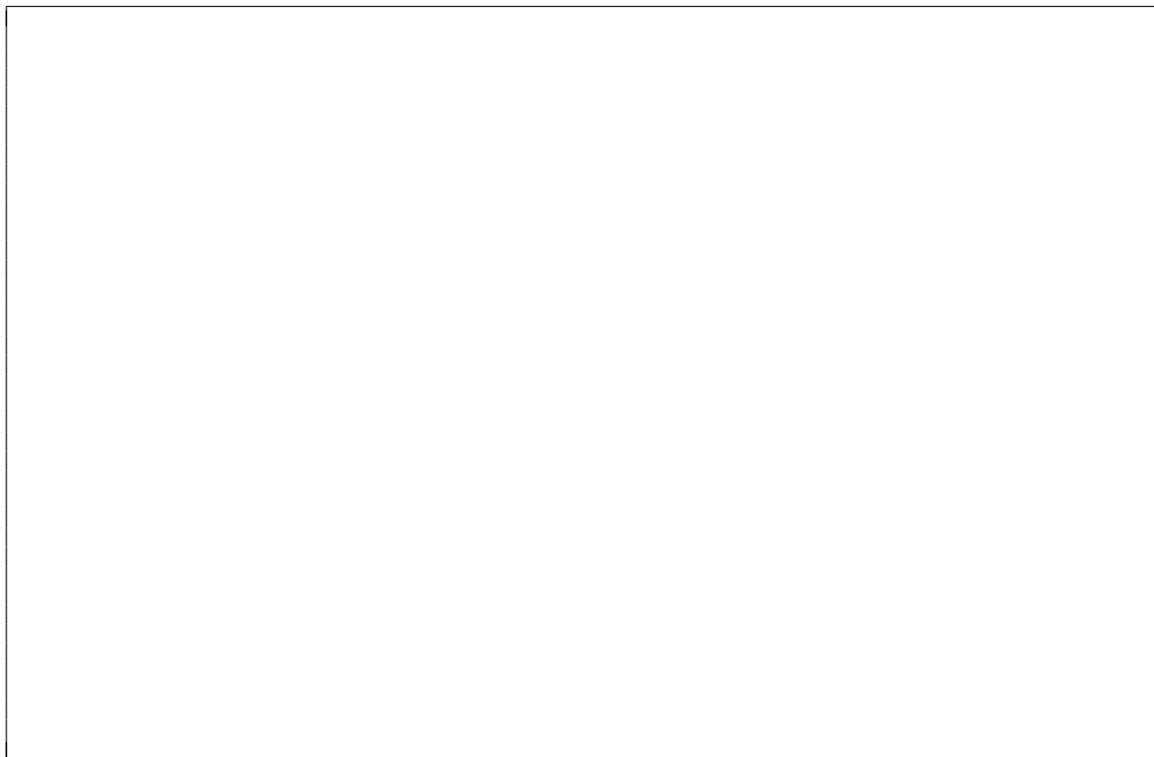


- 4pts 6. What are the three major problems that blockchain tries to solve? Discuss briefly how it solves only *one* of them?



3pts

7. If a mobile computer is to remain accessible to clients when it moves across LANs, it must retain a single IP number. However, IP routing is subnet-based. Subnets are at fixed locations, and the correct routing of packets to them depends upon their positions on the network. Discuss a way of how location transparency can be achieved in such an environment (i.e., IP communication continues normally when a mobile computer moves between subnets at different locations).



3pts

8. How to ensure the durability of non-idempotent operations?

