

28.04.2025 Week 10 exercises: Key-value stores and CAP theorem

Exercise 1:

| What is the CAP theorem?

Exercise 2:

| What is the definition of consistency in the CAP theorem?

Exercise 3:

| Describe the simplest system you can think of that provides consistency and partition tolerance, in a distributed setting.

Hint: It is a very dumb system.

Exercise 4:

| Describe the simplest system you can think of that provides availability and partition tolerance, in a distributed setting.

Exercise 5:

| Explain the difference between the C in ACID and the C in CAP.

Exercise 6:

Among the following propositions, identify which of the CAP theorem guarantees is/are violated.

- You matched on Tinder with your soul mate. Whenever you try to send a message, the application tells you to retry.
- You create a slack channel to discuss about the next PUBG gaming session you will have with your Swiss and American friends. After a while, you only see messages from your Swiss friends. Worried, you call Bob who lives in the US, and learn that he only sees messages from your American friends.
- You open Facebook messenger on both your phone and computer and start talking with Mark Zuckerberg about privacy issues. You do not see the same content on your phone and on your laptop, making the conversation really confusing. All your private data is sold to Cambridge Analytica.
- You want to find a torrent of the latest Smurfs movie. The website keeps telling you that there is no matching result. The same thing happens for any other movie you try to find. Time to get a streaming account.
- You play an online video game. Suddenly, everybody freezes, and none of the players is able to move. Apparently, a cat played with the plug of one of the servers in the data-center, disconnecting it from the other machines.

Exercise 7:

| What are Quorums? How do they relate to consistency levels in Cassandra?

Exercise 8:

| How does Cassandra compare to MySQL in terms of speed? Why?

Exercise 9:

| To ensure serializability (*i.e.*, avoid overlapping of writes or writes and reads) in quorum-based protocols, one must obtain the permission of a quorum in order to read or write a data item. Concerning quorums, answer the following questions:

- (a) What property of majorities make them required in write quorums (*i.e.*, $W > \frac{N}{2}$)?
- (b) Why $R + W$ must be strictly greater than N ?
- (c) What are the implications of choosing between $\langle R = \lceil \frac{N}{2} \rceil, W = \lceil \frac{N+1}{2} \rceil \rangle$ and $\langle R = 1, W = N \rangle$?
where R is the read quorum, W is the write quorum and N is total number of replicas in a distributed database.