Solutions for Fourth Quiz

COSC 6360

Fall 2014

First question

- What do we mean when we say that NFS client requests are: (2×10 pts)

  □ *self-contained*?

  □ *idempotent*?

Answer

- What do we mean when we say that NFS client requests are: (2×10 pts)

  ☐ *self-contained*?
  Each NFS request contains all information necessary to its completion.

  ☐ *idempotent*?
 Multiple executions of the same NFS request produce the same outcome as a single execution.

Second question

- Consider a distributed file system implementing **close-to-open consistency**. Assuming that
  - ☐ Alice opens the file at 9:30 AM, modifies it and closes it at 10:40 AM,
  - ☐ Bob opens the file at 10:00 AM, modifies it and closes it at 10:30 AM,
  - ☐ Carol opens the file at 10:25 AM, modifies it and closes it at 11:30 AM,
- Which of these three users would see his or her changes actually incorporated in the final version of the file? (10 pts)
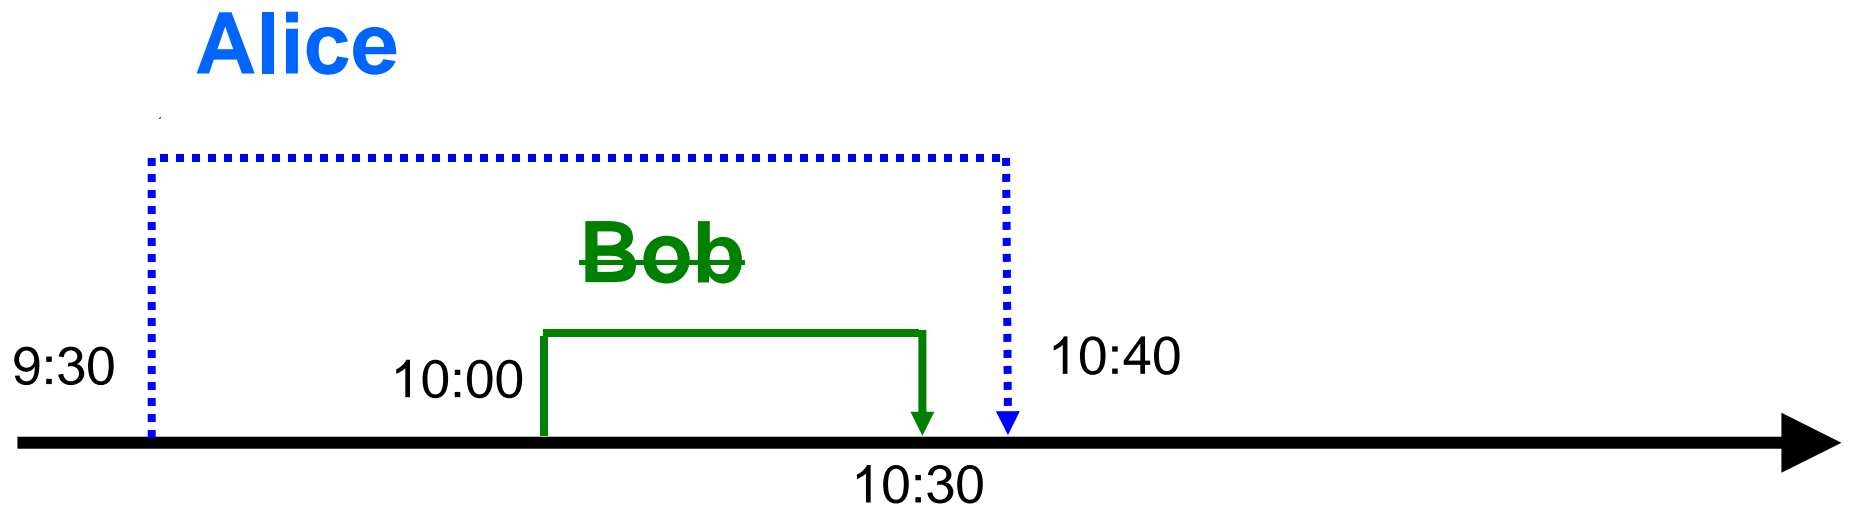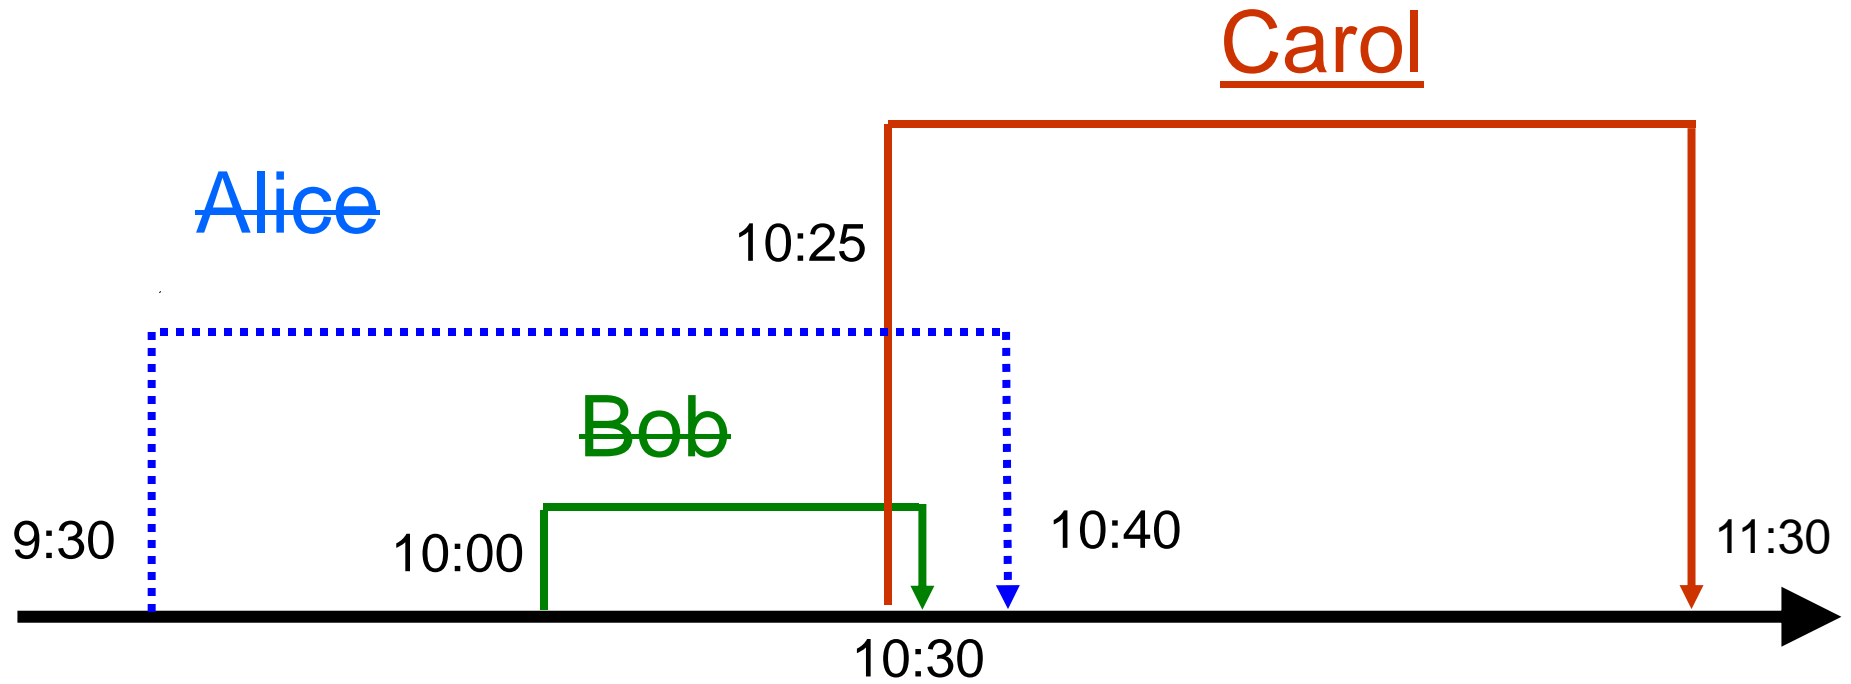
Answer

Alice

9:30

10:40

Answer

Alice

~~Bob~~

9:30

10:00

10:30

10:40
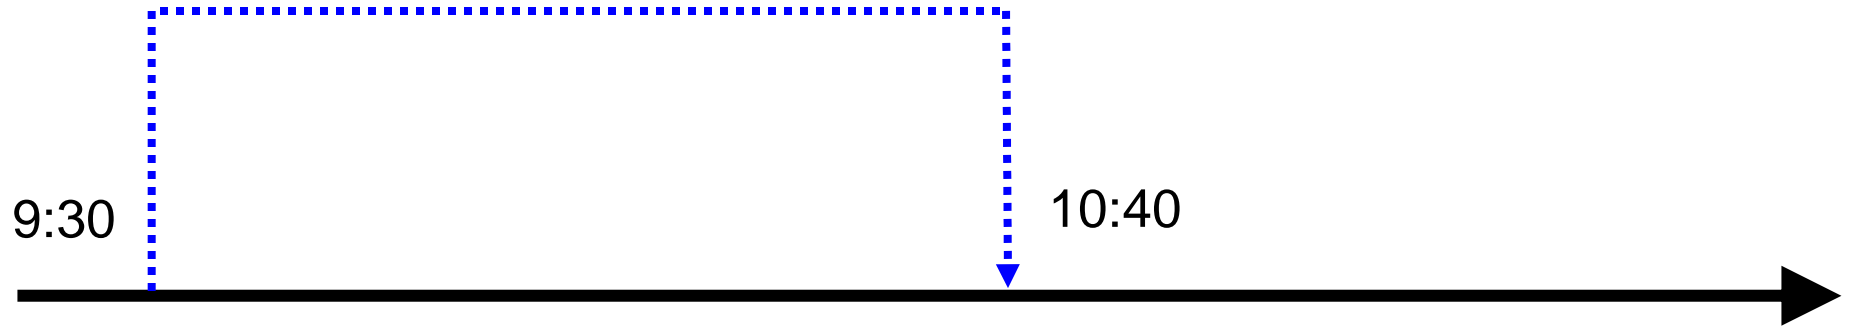
Answer

Third question

- Explain why NSF *safe asynchronous writes* mitigate the need for adding *non-volatile RAM* to the server. (10 pts)

Answer

- Explain why NSF *safe asynchronous writes* mitigate the need for adding *non-volatile RAM* to the server. (10 pts)

  - Safe asynchronous writes allow NSF servers to complete write requests as soon as it has received updated blocks without having to write them to disk first.

  - Otherwise the server would either have to write them first on disk or to store them in an NVRAM buffer.

Fourth question

- Why are BlueFS concepts such as *dynamic storage hierarchy* and *aggregating writes* much less important in today's smart phones, tablets and iPads that they were for the portable devices of ten years ago? (10 pts)

Answer

- Why are BlueFS concepts such as *dynamic storage hierarchy* and *aggregating writes* much less important in today's smart phones, tablets and iPads that they were for the portable devices of ten years ago? (10 pts)

  - Because these techniques were primarily designed for devices incorporating magnetic disks.

Fifth question

- Consider a hypothetical version of NSF that would use *leases* to control access to its files. (2×10 pts)
  - ☐ Would the server of this file system still be *stateless*?

  - ☐ Would that make it *less robust* than NSF?

Answer

- Consider a hypothetical version of NSF that would use *leases* to control access to its files. (2×10 pts)
  - ☐ Would the server of this file system still be *stateless*?
    - *No*, because the server would have to keep track of which clients hold which leases.
  - ☐ Would that make it *less robust* than NSF?
    - *No*, as long as the leases remain short-lived.

Sixth question

- How do the Ceph metadata servers handle **conflicting accesses** by different clients to the same file? (10 points)

Answer

- How do the Ceph metadata servers handle *conflicting accesses* by different clients to the same file? (10 points)

  - They revoke all caching and buffering permissions for that file and require *synchronous I/O* to the file.

Seventh question

- How can **Sybil attacks** defeat FARSITE security? (10 pts)

Answer

- How can *Sybil attacks* defeat FARSITE security? (10 pts)

  - A Sybil attack will let a few rogue hosts pretend to be many and possibly gain the necessary quorums in one or more directory services or hold all the replicas of some files.

Eighth question

- Assuming that we want to protect a FARSITE distributed file system against a **single host failure**,
  - What would be the **minimum size** for all your directory groups? (5 pts)

  - On how **many hosts** should the contents of your files be replicated? (5 pts)

Answer

- Assuming that we want to protect a FARSITE distributed file system against a ***single host failure***,

  - What would be the ***minimum size*** for all your directory groups? (5 pts)
    - **3 + 1 = 4 hosts**

  - On how ***many hosts*** should the contents of your files be replicated? (5 pts)
    - **1 + 1 = 2 hosts**