SOLUTIONS TO THE FOURTH 6360 QUIZ

Jehan-François Pâris
Fall 2016

First question

- NFS is said to use ***idempotent requests***.
  - What characterizes these requests?
  - What is the main advantage of the approach?

First question

- NFS is said to use ***idempotent requests***.
  - What characterizes these requests?
    - Multiple executions of any request produce the same result s a single execution
  - What is the main advantage of the approach?
    - When a server crashes, client just resends its requests until it gets answers from the rebooted server

Second question

- Consider an NFS file system that implements ***close-to-open consistency***. What should the system client do
  - When a user opens a file?
  - When a user close a file?

Second question

- Consider an NFS file system that implements ***close-to-open consistency***. What should the system client do
  - When a user opens a file?
    - It must check with the NFS server that any locally cached data are up-to-date
  - When a user close a file?
    - It must write all modified file data to the server

Third question

- Consider a distributed file system implementing close-to-open consistency.
- Assuming that
  - Alice opens a file at 9:30 AM, modifies it and closes it at 10:10 AM,
  - Bob opens the same file at 10:20 AM, modifies it and closes it at 10:30 AM,
  - Carol opens the same file at 10:25 AM, modifies it and closes it at 11:30 AM,

Third question

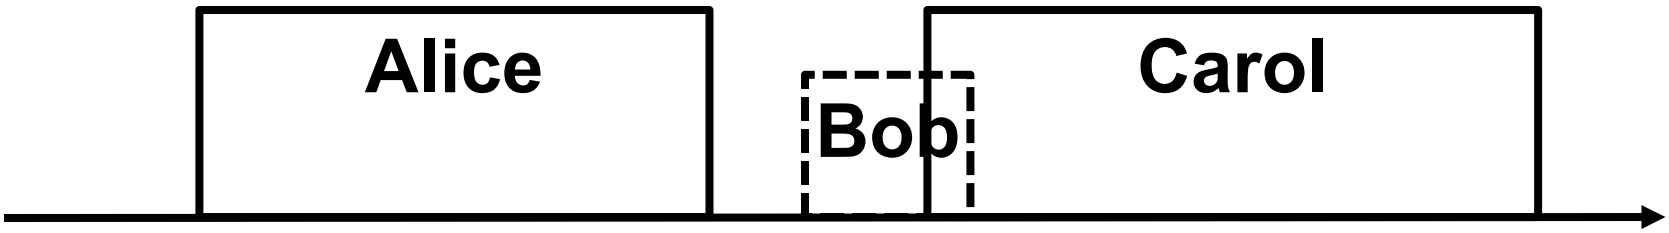

- Which of these three users would see his or her changes actually incorporated in the final version of the file?
  - Alice and Carol

Fourth question

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

Fourth question

- How do the Ceph metadata servers handle ***conflicting accesses*** by different clients to the same file?
  - When a Ceph MDS detects conflicting accesses by different clients to the same file
    - It revokes all caching and buffering permissions for that file
    - It forces ***synchronous I/O*** to the file

Fifth question

- How does FARSITE store ***users' secret keys***?

- Why?

Fifth question

- How does FARSITE store ***users' secret keys***?
  - User private keys are encrypted with a symmetric key derived from user password and stored in a globally-readable directory in Farsite
- Why?
  - Secret keys are too large to be memorized by users

Sixth question

- Assuming that we want to protect a FARSITE distributed file system against ***all double failures***,
  - What would be the ***minimum size*** for all your directory groups?
  - On how ***many hosts*** should the contents of your file be replicated?

Sixth question

- Assuming that we want to protect a FARSITE distributed file system against ***all double failures***,
  - What would be the ***minimum size*** for all your directory groups?
    - 3×2 + 1 = 7 hosts
  - On how ***many hosts*** should the contents of your file be replicated?
    - 2 + 1 = 3 hosts

Seventh question

- How does the LBFS file server ensure the atomicity of updates?

Seventh question

- How does the LBFS file server ensure the atomicity of updates?
  - The LBFS server ensures the ***atomicity of updates*** by ***writing them first into a temporary file***