NAME: _____ (FIRST NAME FIRST) **SCORE**: _____

COSC 6360 FINAL EXAMINATION DECEMBER 12, 2008

THIS EXAM IS CLOSED BOOK. YOU CAN HAVE ONE SHEET (I.E., TWO PAGES) OF NOTES. UH EXPELS CHEATERS

1. Mark all the properties that apply to the following distributed file systems. (5 points per correct line)

Distributed file system	***NFS***	***AFS***	***Coda***
Does not ***trust*** its clients		X	X
Allows clients to operate in ***disconnected mode***			X
Supports ***diskless clients***	X		
Uses ***callbacks***		X	X
Uses ***replicated servers***			X
Uses ***stateless servers***	X		

2. Consider a Coda file system having two copies ***A*** and ***B*** of the same file ***X***. In which state are these two copies if their LSID's and CVV's respectively are:

$\mathbf{LSID_A = 7435}$ $\mathbf{LSID_B = 4987}$
$\mathbf{CVV_A = \{5 , 4\}}$ $\mathbf{CVV_B = \{4 , 4\}}$

(5 points)

Replica A dominates replica B and has the most recent version of the file. _____

3. Consider a diskless client trying to access a file named "**`/usr/joe/6360/paper.doc`**" that is stored on its NFS server.

a) Assuming that the client already has a handle for its root directory, how many ***lookup()*** requests will it issue? (5 points)

Four _____ requests

b) What does NFS do to speed up these requests? (5 points)

NSF caches recently used file handles._____

4. What is the meaning of the following Swift label? (5 points)

`{ alice ← bob }`

alice trusts the information and believes only bob should modify it. _____

T:___/50

5. Mark all the properties that apply to the following distributed file systems. (5 points per correct line)

Distributed file system	***LBFS***	***BlueFS***	***FARSITE***
Uses ***Byzantine agreement*** for directory updates.			X
Uses ***leases*** to implement close-to-open consistency	X		
Vulnerable to ***Sybil attacks***			X
Aggregates writes to save energy		X	

6. What are the three exchanges of messages that occur during the gracious execution of the Zyzzyva Byzantine fault-tolerant protocol (3 × 5 points)

a) Primary replica _____ send(s) a message to secondary replicas. _____

b) Secondary replicas _____ send(s) a message to client. _____

c) Client _____ send(s) a message to all replicas. _____

7. What is the purpose of the ***CRUSH function*** in the Ceph distributed file system? (5 points) Why is it publicly available? (5 points) How does it simplify the design of the Ceph metadata server cluster? (5 points)

(**Hint:** Since the question has three parts, we should have three separate paragraphs in the answer.)

The CRUSH function is a pseudo-random function that lets clients assign placement groups to object storage devices.

CRUSH is publicly available to all Ceph clients because they need to execute it to locate their files.

CRUSH simplifies the design of the Ceph metadata server cluster by eliminating the need of storing object addresses in the cluster.