NAME: _____ (First name first) **SCORE:** _____

COSC 6360 **QUIZ #3** **AUGUST 2, 2010**

Closed book. You can have with you one single-sided 8½ by 11 sheet of notes.

1. What steps must be taken by a ***log-structured file system*** to locate an i-node when it reboots after a ***normal shutdown***? ((3×5 points)

a) Access check-point area **to get address of the appropriate block of**

i-node map

b) Fetch that block to get address of the block containing i-node

2. Which are the main advantage and the main disadvantage of using journaling with ***asynchronous log updates*** compared to using journaling with ***synchronous log updates***?

a) ***Main advantage:*** (5 points)

Asynchronous log updates require fewer disk accesses and provide higher file system throughputs.

b) ***Main disadvantage:*** (5 points)

They do not guarantee the durability of metadata updates.

3. In Totem, what is the purpose of **guaranteed vector messages?** (5 points)

Guaranteed vector messages are part of the Totem multiple-ring protocol. They allow messages to be delivered to their destinations in a timely fashion without waiting for messages from a ring that remains silent.

4. In Kerberos, how does the ticket granting service (TGS)—and all other services—distinguish between valid authenticators and authenticators that are being ***replayed by intruders*****?** (2×5 points)

TGS will reject all authenticators whose timestamps are either too old or duplicates of a timestamp that it already received.

5. Consider a RAID level 5 disk array with four data blocks (b_0, b_1, b_2 and b_3) and one parity block p per stripe. What is the most efficient way to update block b_2 and leave the array in a consistent state? (4×5 points)

a) Read values of old block b_2 and old parity p.

b) Write new value of block b_2 and new parity p **= old parity p XOR nold b_2 XOR new b_2.**

6. Consider a file system using soft updates and assume that one directory block in its I/O buffer reflects the result of one file deletion and one file creation. Assuming that the i-nodes of the two files reside in two different i-node blocks, enumerate the steps the system will take to update the disk copies of the three blocks. (3×5 points)

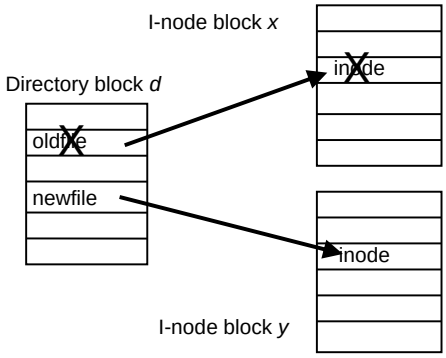

a) Write i-node block y.

b) Write directory block d.

c) Write i-node block x.

7. What is the main limitation of ***soft updates*?** (5 points)

They do not guarantee the durability of metadata updates.

8. Why do all NFS servers implement a ***write-through policy***? (5 points) What is the ***main drawback*** of that policy? (5 points) Which new feature of NFS v. 3 addresses that issue? (5 points) How? (5 points)

a) NFS servers implement a write-through policy in order to remain stateless.

b) The policy results in many more disk accesses than in a regular UNIX file sytem.

c) NFS now has safe asynchronous writes.

d) Safe asynchronous writes allow NFS servers to implement delayed writes as long as the client keeps a local copy of the blocks it sends to the server. A new commit primitive lets clients check with the server that it actually has written the data.