

Name: _____ (First name first)

Score: _____

COSC 6360

QUIZ #5

DECEMBER 1, 2010

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

1. You want to build a Byzantine fault-tolerant system that can tolerate the failure of two entities. What is the **minimum** number of entities that you would need? (10 pts)

A minimum of 3x2 + 1 = 7 entities

2. **Pergamum**

- a) What equipment failures can be corrected by **intratome redundancy**? (10 pts)

Irrecoverable read errors.

- b) What would be the main drawback of a Pergamum system having plenty of **intratome redundancy** but no **intertome redundancy**? (10 pts)

It would not tolerate full disk failures.

- c) How are Pergamum tomes **powered**? (5 pts) Through conventional cables.

3. **Zyzyva**

- a) What is the main **disadvantage** of Byzantine fault-tolerant replication controls compared to protocols assuming fail-stop behaviors? (5 pts)

They are much slower AND/OR They require $3f + 1$ entities to tolerate f failures.

- b) Under which circumstances is Zyzyva **slower** than a traditional Byzantine fault-tolerant replication control protocol? (10 pts)

Anytime something goes wrong.

4. **FAWN**

- a) How is the FAWN **datastore** organized? (10 pts)?

As a log operating in append mode.

- b) Why? (10 pts)

Because flash memory performs sequential writes much faster than random writes.

- c) What is the purpose of allocating several randomly selected **virtual nodes** to each FAWN node? (10 pts)

To spread the workload of a failed physical node among several physical nodes.

5. **Ceph**

- a) What does Ceph when MDS detects conflicting accesses by different clients to the same file? (10 pts)?

It revokes all caching and buffering permissions and requires synchronous I/O to that file.

- b) What is the **main advantage** of this solution? (10 pts)?

It prevents inconsistent updates.