COSC 6360 FINAL EXAMINATION DECEMBER 13, 2002

FIRST PART AND GRADUATE PART OF THE QUALIFYING EXAMINATION

This exam is closed book. You can have two sheets of self-prepared notes. UH expels cheaters.

1. ***Remote Procedure Calls:***

Most distributed systems use a remote procedure call facility to let clients communicate with servers. These facilities were usually encoding data using some predefined format—such as Sun's External Data Representation or XDR—and using specific ports to interprocess communication.

The recent years have seen several proposals using XML as external data representation and one of them—SOAP—can run on any http server and does not require any specific ports.

What are the advantages and disadvantages of these new protocols in terms of portability, efficiency, ease of deployment and security?

Hint: Consider how these proposals will interact with existing firewalls.

2. Long term Authenticity of Documents

Most if not all digital signature schemes are based on the property that some specific operations like guessing a person's private key knowing her private key and more generally computing the inverse of a one-way function are not feasible in the current state of the technology. The solution is very satisfactory for short-term contracts that are to expire in a few years.

They are however many long-lived contracts, think only of mortgages, whose duration is such that we can reasonably expect that it will become computationally feasible to tamper with their contents during their duration. The same applies to official documents, which tend to be declassified after many years.

How would you guarantee—or reaffirm —over time the authenticity of the original of these documents?

Hint: The best solutions should make no assumptions about the speed of change of the computing technology.

COSC 6360 FINAL EXAMINATION DECEMBER 13, 2002

SECOND PART

1. What is a *replay?* (5 points) What does Kerberos to allow servers to distinguish relays from authentic messages? (10 points)

2. What is ***sequential sharing*** in Sprite? What could Sprite do to eliminate it? (5 points) What would be the main disadvantage of doing so? (5 points)

3. Which previous versions of a file does Elephant keep and why? (10 points)

4. What are the main advantage and the main disadvantage of the AFS ***callback*** mechanism? (10 points)

5. Why does HARP use ***uninterruptible power supplies*** for its servers? (5 points)

6. Does the ***small write problem*** applies to RAID-3 organizations and why? (5 points)