

2010 Paper 8 Question 7

Distributed Systems

- (a) The goal is to design a naming system that, along with a storage service, allows authorisation software to retrieve role-based access control data.
- (i) Suppose that the authorisation software is presented with an authenticated principal. To make its authorisation decision, what information will it need to retrieve from storage and therefore what sorts of objects should be in the naming system? [3 marks]
 - (ii) Should names be pure or impure? Justify your claim. [2 marks]
 - (iii) How might impure names be structured? [2 marks]
 - (iv) It is likely that the authorisation software will need to understand principals from several organisations. How should names reflect this? [2 marks]
- (b) Suppose that it is early 2005 and you are designing a web site allowing registered users to upload videos and anybody to download and watch them. You expect and hope that the site will be popular, so you use replication to handle the anticipated request rate. (Assume that, unlike a certain mainstream site whose name you can probably guess, no comments are associated with the videos.)
- (i) What are strong and weak consistency? Describe the advantages and disadvantages of using each for this application, paying particular attention to any assumptions you must make about the users' behaviour. [4 marks]
 - (ii) You are asked to modify your site so that it can be used to store and annotate video evidence for law enforcement. This requires using strong consistency; why? [2 marks]
 - (iii) Suppose that you use nine replicas to meet demand and quorums to ensure strong consistency. What are the constraints on the sizes of read and write quorums? Give an example of quorum processing involving a write, two reads, then a write. Use the notation v_i to indicate version i of the data and suppose that at the start of your example, each replica contains version 3. Make sure that your diagram is easy to understand and that you explain any notation that you use. [5 marks]