

## **Tutorial 9**

### **Distributed File Systems, WWW, Distributed Coordination**

1. Consider fault tolerance of the file systems covered in lectures. If you had to implement a high-availability solution, which would you use, independent of other criteria like performance?
2. You have a project where a specific requirement is that if a file is visible to a user, it should have the same name (including path) wherever any user can access it. Which of the following meet this criterion? Explain your answer.
  - a. NFS
  - b. Coda
3. Fault tolerance and performance are always at odds in distributed file systems. True or false? Think of cases where they are or are not in conflict.
4. To what extent are NFS and Coda stateless? Comment generally on the desirability of statelessness and why it may reasonably be adopted or abandoned.
5. Does NFS implement entry consistency?
6. Would it make sense to associate a replication strategy with each Web document separately, as opposed to using one or only a few global strategies?
7. Is coordination a logical thing to split from processing? What else is there which is specifically an organizational or management aspect of distributing computing which would be logical to split off from processes and threads?
8. What type of coordination model would you classify the message-queuing systems discussed in the lecture on communication?
9. Why JavaSpaces has a scalability problem?