CS 5523 Lecture 28: Review

- *Discuss Farsite Paper*
- *Questions from Chapter 8*
- *Questions from Chapter 9*

Discussion Questions from CDK Chapter 8

[CDK 8.1]

Why is there no open or close operation in the interface to the flat file service or the directory service. What are the differences between our directory service Lookup operation and the UNIX open?

Discussion Questions from CDK Chapter 8

[CDK 8.4]

Why should UFIDs be unique across all possible file systems? How is uniqueness for UFIDs ensured?

Discussion Questions from CDK Chapter 8

[CDK 8.5]

To what extent does Sun NFS deviate from one-copy file update semantics? Construct a scenario in which two user-level processes sharing a file would operate correctly in a single UNIX host but would observe inconsistencies when running in different hosts.

Discussion Questions from CDK Chapter 8

[CDK 8.6]

Sun NFS aims to support heterogeneous distributed systems by the provision of an operating system-independent file service. What are the key decisions that the implementer of an NFS server for an operating system other than UNIX would have to take? What constraints should an underlying filing system obey to be suitable for the implementation of NFS servers?

Discussion Questions from CDK Chapter 8

[CDK 8.7]

What data must the NFS client module hold on behalf of each user-level process?

Discussion Questions from CDK Chapter 8

[CDK 8.14]

How does AFS gain control when an open or close system call referring to a file in the shared file space is issued by a client?

Discussion Questions from CDK Chapter 8

[CDK 8.15]

Compare the update semantics of UNIX when accessing local files with those of NFS and AFS. Under what circumstances might clients become aware of the differences?

Discussion Questions from CDK Chapter 8

[CDK 8.16]

How does AFS deal with the risk that callback messages may be lost?

Discussion questions from CDK Chapter 9

CDK [9.2]

Discuss the problems raised by the use of aliases in a name service, and indicate how, if at all, these may be overcome.

Discussion questions from CDK Chapter 9

CDK [9.3]

Explain why iterative navigation is necessary in a name service in which different name spaces are partially integrated, such as the file naming scheme provided by NFS.

Discussion questions from CDK Chapter 9

CDK [9.4]

Describe the problem of unbound names in multicast navigation. What is implied by the installation of a server for responding to lookups of unbound names?

Discussion questions from CDK Chapter 9

CDK [9.8]

Why do DNS root servers hold entries for two-level names such as yahoo.com and purdue.com, rather than one-level names such as edu and com?

Discussion questions from CDK Chapter 9

CDK [9.9]

Which other name server addresses do DNS name servers hold by default, and why?

Discussion questions from CDK Chapter 9

CDK [9.12]

The Jini lookup service matches service offers to client request based on attributes or on Java typing. Explain with examples the difference between these two methods of matching. What is the advantage of allowing both sorts of matching?

Discussion questions from CDK Chapter 9

CDK [9.16]

Discuss the potential advantages and drawbacks in the use of a X.500 directory service in place of DNS and the Internet mail delivery programs. Sketch the design of a mail delivery system for an internet work in which all mail users and mail hosts are registered in an X.500 database.