Solutions to Midterm
Fall 2002 15-712

Please write your name on the following line (and on each page of the exam):

____

This midterm is open book and open notes.

You should write brief, precise, and legible answers. Rambling brain-dump answers may be punished twice — you run out of time writing them and we run out of patience reading them. :–)

If in doubt, list any assumptions you made about the question.

There are eight questions but we will only grade **six**. List the **six** questions we should grade:

____

Name: ______________________________

1. The Virtual Log paper shows how a disk with eager writing can reduce write service times by more than 90%. Two engineers are arguing about whether or not an FFS file system using soft updates would benefit from eager writing. Give one argument (in the form of a specific example workload) in favor of each position.

   Eager writing can speed up FFS with soft updates for workloads requiring synchronous writes. For example, online transaction processing applications mandate that log writes be committed to the non-volatile storage before reporting completion of a transaction. Similarly, mail servers use fsync() to make sure messages are stored before acknowledging them.

   Workloads with no explicitly synchronous writes, such as saving EMACS buffers, will get little to no benefit from eager writing. Soft updates makes metadata updates asynchronous. Eager writing may make the background writes faster, but one would expect this to offer little benefit to the applications.

Name: ____________________________________

2. The Emerald system provides flexible mechanisms for object placement and migration. Identify and explain two security dilemmas created by its flexibility. (Think about ways that an attacker could exploit it.)

An attacker can easily launch a DoS attack by spawning an object that consumes entire client's resources. Emerald's migration mechanism can then be used to move this object to anther system.

Another DoS would be to have numerous objects migrate repeatedly.

An object could be created that reads or writes various memory locations by fabricating pointers. Once migrated to another machine, such an object can violate confidentiality and integrity goals.

Name: ________________________________

3. Both Emerald and Abacus rely on an RPC style of communication between objects on different systems. For each, identify one pro and one con to switching to a streaming model of communication (like that in River).

Emerald

pro: a streaming model may help application writers better identify the most important inter-relationships between objects.

con: the streaming model may make it more difficult to identify when an object can be quiesced and moved.

Abacus

pro: less communication and more parallelism by letting the source push data at the sink.

con: less information about the data unit boundaries. The RPC call/return is convenient unit for tracking "per datum" resource requirements. Also, more difficult to identify clean quiescence points.

Name: ____________________________________

4. System crashes are "yucky" [Greg's wording, not Jiri's—Ed.]. Explain what data is lost (if any) and what recovery mechanism is used for each of these scenarios:

   (a) A crash of a system using an LFS file system.

   LFS will lose any data and metadata updates that occured in the active in-memory segment that hasn't been written out to the disk.

   LFS keeps a checkpoint that marks a consistent state of all data written to the disk. Any data modified after the checkpoint are "rolled forward" by examining segment summaries of segments written since last checkpoint. FS consistency is guaranteed before and after each log segment write.

   (b) A crash of a cache node in a hierarchical proxy cache like Harvest.

   No data will be lost since the node serves as a read-only cache.

   The Harvest cache node will not lose any data that are replicated in that node. When such a node fails, the "children" have to initiate recovery by contacting the "source" nodes in hierarchy chain to obtain data that were originally cached at the failed node. When the node restarts, it needs to reinsert itself into the hierarchy by contacting its parent and children.

Name: ____________________________________

5. The Abacus function placement approach dynamically splits an application object stack between client and server. Such partitioning has also been considered for wide-area computing. For each of Harvest and CFS, discuss the ramifications of their system architecture on the Abacus approach of dynamic function placement.

   The Harverst architecture would fit the Abacus model well, since its strict hierarchy represents a series of nodes between any client and server. The Abacus model could be extended to split the object stack at several points with zero-or-more parts on the client, each cache node along the way, and the server.

   The CFS architecture would be a poor match because it randomizes the location of each block across a multitude of servers. Extending Abacus to it would thus require replicating any "server side" object at many servers.

Name: ______________________________

6. xFS is designed to be scalable and to decouple file services from specific machines. But, it does so with globally-replicated metadata tables that must be coordinated when changes happen. CFS achieves similar goals with a very different mechanism. List two reasons that an xFS implementor might not want to immediately switch to the CFS mechanism.

   The lookup for data location in xFS is done by the globally replicated metadata table. Thus, one can contact simply lookup the location in the globally replicated metadata table and fetch the data directly. CFS, on the other hand, sends on the order of log(N) messages through a Chord ring to perform a lookup for the location of data. Thus, the CFS solution may be slower and involve more communication than the xFS implementation.

   CFS locates blocks by hash ID which tends to randomize their locations. xFS's design allows for locality via explicit assignment of locations.

Name: ____________________

7. Proxy caches are used by some companies (and countries) to control the inward flow of information. Describe three potential security dilemmas that this creates for the people using that proxy. (Think about what the proxy could do to you, if it wanted.)

   - *Loss of confidentiality and privacy for users. Since all the information between clients/users and the outside world passes through a centralized proxy, the proxy can determine who is sending what information.*
   - *This setup can make censorship very easy. It can selectively or totally deny service.*
   - *Data authenticity. Since the proxy is providing information to all the clients in the given domain, the data going through the proxy can be altered arbitrarily.*

Name: ___________________________________

8. AFS and xFS are both designed for scalability. Think about which system would be expected to scale better with the number of clients: a collection of 8 AFS servers or a collective of 8 xFS nodes exporting an NFS file service? Since systems answers are usually "it depends", explain when each would be better and why.

Compared to NFS, AFS offers more aggressive client caching and callbacks. Thus, AFS would be expected to perform better than NFS service on top of xFS when the load is distributed relatively evenly among the servers.

The xFS-based NFS service would be expected to scale better, thanks to automatically balanced load, if the AFS clients were overloading a single server with a popular volume.