Name: _____

CS 6/73201 **Final Exam** Advanced OS

Monday 12 December 2005

1. **Explain how the SUN NFS virtual file system co-exists with and interacts with the UNIX file system. (15 points)**

   The NFS virtual file system sits "above" the UNIX file system. System calls to access a file are made to NFS, which determines whether the file being accessed is local or remote. If local, the system call is passed on to the UNIX file system. If remote, the RPC-based NFS protocol is used to access the remote machine, where the system call is passed on to the UNIX file system on that machine.

2. **Two semantic models used in distributed file systems are UNIX semantics and session semantics.**

   a. **Briefly define each of these models. (10 points)**

      UNIX semantics — result of a write goes immediately to the file, so reads always return the latest value written

      Session semantics — result of a write goes to local storage, and only goes to the file when the session/file is closed. (Further, files are cached when opened, and modified files are flushed back to the server when closed, so any client using a previously-cached copy continues to do so even if the file has been updated on the server.)

   b. **If the file data is cached, which cache modification policy is most appropriate for each of these semantic models and why? (6 points)**

      UNIX semantics — write-through, so the new value is immediately written to the server

      Session semantics — write-on-close (variation of write-back / delayed-write), so the new value is written to the server only when the file is closed

3. **Two process migration mechanisms are total freeze and pre-transfer.**

   a. **Briefly define each of these mechanisms. (10 points)**

      Total freeze — stop execution of the process while the address space is being transferred

      Pre-transfer — continue executing the process while the address space is being transferred, then stop execution while the remaining modified pages are being transferred

b. Which has less freeze time? Explain. (6 points)

Pre-transfer, since the process is frozen only during transfer of the remaining modified pages instead being frozen during transfer of the entire address space (as in total freeze).

4. Three sender-initiated load distribution algorithms were studied by Eager, Lazowska, and Zohorjan, which differ only in their location policy. Would the use of broadcast or multicast (broadcast to a specific subset of notes) improve any of these three location policies or would that only decrease performance? Explain. (13 points)

It wouldn't help with the *random* policy, but it could make the polling in the *threshold* policy more efficient if multiple polls are necessary (though at the cost of requiring more other processors to listen to the poll, which would be particularly detrimental in a heavily loaded system). Similarly, it could get the request for queue length to the processors more quickly in the *shortest* policy.

5. Why do receiver-initiated load distribution algorithms tend to be stable at high system loads, whereas sender-initiated load distribution algorithms are unstable at high system loads? (15 points)

At high system loads, most nodes are senders, so receiver-initiated algorithms can easily find a sender without increasing the system load appreciably, whereas sender-initiated algorithms are unlikely to find a sender so they waste their time and the listener's time with fruitless polling that can push the system into instability.

6. Define strict memory consistency, and explain why it is impossible to achieve in a distributed system with shared memory. (15 points)

In strict memory consistency, the value returned by a read operation is always the same as the value written by the most recent write operation. Implementing strict memory consistency would require global time, which is not possible in a distributed system.

7. How does a polyalphabetic cipher such as the vigenere cipher improve on the Caesar cipher? (10 points)

In the Caesar cipher the statistical frequencies of characters are not concealed and the key can be found by exhaustive search. In polyalphabetic ciphers such as the vigenere cipher the key is longer, which smoothes out the statistical frequencies to make analysis harder.

8. Define private key cryptography, and explain the key distribution problem. (15 points)

In private key cryptosystems (e.g., DES), both the encryption and decryption keys are kept private to the communicating parties (as opposed to public key cryptosystems, where the encoding key is made public). Thus before communication can take place, the keys must be passed from one of the parties to the other using some secure channel.

9. What is an embedded system? (10 points)

An embedded system is a combination of hardware and software dedicated to perform a specific function, generally as part of a larger system.