

CS423UG Homework 6 Solution

December 9, 2005

1. Remote Procedure Call (RPC) is a protocol that allows a computer program running on one host to cause code to be executed on another host without the programmer needing to explicitly code for this. (excerpt from Wikipedia: http://en.wikipedia.org/wiki/Remote_procedure_call)

Since RPC middleware marshals and unmarshals parameters and results, and the middleware format is platform-independent, an RPC between a Sun and a Windows machine will work.

2. The original wired telephone network is circuit-switched, while the IP telephony service is packet-switched. In circuit-switched networks, the resources remain dedicated to the circuit during the entire transfer and the entire message follows the same path, which guarantees a high quality communications. On the other hand, packet-switched networks are better for larger scale, as messages are broken into small data packets that seek out the most efficient route as circuits become available.

3.
 - Name space

While the name space of NFS is per-client linkage, AFS employs global name space.

- Location transparency

If a directory is moved from one server to another, client must remount in NFS. On the other hand, in AFS, only the volume location database on the servers needs to be updated.

- Caching

While the client cache location is main memory in NFS, AFS stores client cache in local disk.

- Scalability

While NFS is intended for only small number of clients and servers, AFS can scale to large systems by reducing server load and clustering

- Security

While server trusts client in NFS, security mechanisms have been implemented in AFS, such as access control lists (ACL) and Kerberos authentication between client and server, etc.

There might be other differences, as described in the textbook.

4.
 - Adapt (mostly decrease) the value of TTL as long as the query returns a good result.
 - Cache the original query parameters for a while (and use for similar subsequent queries).
 - Cache query results, and propagate files back along reverse route, caching at intermediate nodes.
5. SETI@home is a grid computing project by Space Sciences Laboratory, at University of California, Berkeley, using Internet-connected computers. This project aims at analyzing data incoming from the Arecibo radio telescope, searching for possible evidence of radio transmissions from extraterrestrial intelligence. Each user can download a software from SETI@home, which makes use of wasted processor power to analyze data, divided as 0.35MB work units sent from SETI@home server. Since its launch in May 1999, aggregate computing time has reached over two million years, and number of users has reached five million. Recently, SETI@home has mostly transferred its computing software to Berkeley Open Infrastructure for Network Computing (BOINC), which allows users to contribute to other grid computing projects running on the BOINC platform.

6. From security perspective, it is good to erase each block when a file is removed. Otherwise, attackers would be able to access such blocks by tampering hard disks or using buffer overflow attack even after a file has been deleted. However, there is a trade-off between security and performance. In this case, file delete operation will require an additional cost to erase each block. This can be problematic, as it will take much time to erase a big file. Also, once a file is erased, it will be lost for good.