

Introduction to Networking

Homework 1

Handed out: Monday January 13, 2003

Due back: Wednesday, January 29, 2002 (at the start of class)

Notes: To be done individually

1. Textbook problem 1-3
2. Textbook problem 1-5
3. Textbook problem 1-8
4. Textbook problem 1-15 – also record and comment on the **paths** that you found. Note how paths and delay can change over time. You will want to use Dylan, or your own machine for this. The TLAB firewall doesn't allow traceroutes.
5. Textbook problem 2-1
6. Textbook problem 2-6
7. Textbook problem 2-8
8. Look at <http://www.omg.org>. Write a one paragraph summary of what CORBA is. Avoiding marketing speak.
9. Look at <http://java.sun.com>. Write a one paragraph summary of what Java RMI is. Again, avoid marketing speak.
10. Look at <http://www.xml.com>. Write a one paragraph summary of what XML is. Again, avoid marketing speak.
11. Look at <http://www.openldap.org/>. Write a one paragraph summary of what LDAP is. Again, avoid marketing speak.
12. Look at <http://www.gridforum.org>. Write a one paragraph summary of what Grid computing is. Again, avoid marketing speak.

Extra Credit Problem

Akamai (<http://www.akamai.com>, stock symbol AKAM) is a leading provider of web caching services. An Akamai client can serve its own content or have Akamai's global system of web caches serve it. For example, the client may want to serve the text and advertisement portions of its web site in order to keep track of the number of times each page or advertisement is visited, but not want to be burdened with serving images. The client can choose which objects Akamai will cache simply by "akamizing" their URLs. If you view the source of the CNN home page (<http://www.cnn.com>) or many other high-volume web sites, you will find many akamized URLs. When your browser fetches an akamized URL, Akamai selects a cache that is near you and redirects your request (and your future requests for akamized URLs from that client) to that cache. It does all this by cleverly exploiting the DNS system. Your task in this problem is to learn and describe how this is done. You are free to read Akamai's web pages, but you must give your answers in your own words and they must be accompanied by an example. You may find the programs "nslookup" and "dig" helpful both to explore what Akamai does and in creating your example. Your description should answer the following questions:

- a. How does Akamai discover where the browser's host is?
- b. How does Akamai redirect the browser's request to the cache?
- c. Why does Akamai scale? (Why aren't all requests processed in same costly manner?)