Lab assignment 3: Domain Name Service

Objective
To obtain practical experience regarding the role of DNS in a computer network and the interaction between users and DNS.

Introduction
In this lab we are going to investigate DNS through user commands in the terminal of a Linux distribution. In some of the assignments you need to look up arbitrary computers outside LTH network. Which computers are suitable and how can we obtain their names? A good idea is to choose your favorite web servers. E-mail entries can also be sources for names of computers.

How can we find the geographical location of a machine with a given name? In general, this is not trivial. In case the top domain indicates a country, such as .se, we know with some probability the country in which the server resides. On the other hand, computers in the .com, .edu or .org domains do not necessarily need to be in USA. In USA you can always select a machine at a well known university, such as MIT, Stanford, or Caltech.

Translation of a host name to its IP address is fundamental to most services in Internet. Thus, there is a requirement for checking how the translation is done explicitly. ***Dig***, ***host*** and ***nslookup*** are examples of free programs which explicitly contact a name server and ask for various information. In this lab students are supposed to use dig.

Before starting with the tasks in this lab each student should read the UNIX/Linux manual for ***dig***. The program has no interactive mode. One command will give one answer which contains all the details. A few things need to be clarified. In the answer there is a field starting with "flags:". If the answer is authoritative there will be the indication "aa" (authoritative answer) in this field. If the answer is not authoritative there will be no such indication. In the answer section, the Time To Live (TTL) for the entry is given in seconds.

Report
You shall follow all the following tasks in detail and demonstrate the commands, outcomes, answers and analysis of the logs from dig concisely. In each case just report the portion of the logs from dig that is relevant to the task and analyze it. Avoid pasting lengthy screen dumps and unprocessed command outputs. In order to pass the lab, each group of two students shall make sure that they have answered all questions in detail and correctly. Also you should follow the report structure introduced on the course home page. If you are required to resubmit the report you will have maximum 7 days from the time of the announcement to have an approved report.

Task Description:

1. How can we find the IP address of a computer from its name?

Let’s start with finding out the IP address of the computer, in the student’s computer room, that you are using. Use ***dig*** to do so. Repeat the task for some neighbor computers. Now let’s see what will happen if we try to lookup a nonexistent computer in the network using ***dig***. Finally look for the file: **/etc/hosts**. Describe the function of this file. Use a proper editor to study its content. When was the file created? Can you see the name of the computers you have searched for in it? Why?

2. Where is e-mail really sent?

As you know, names of the electronic mail servers are stored in the DNS database as MX records. Study the electronic mail servers (and their corresponding preference values) which are the real receivers of the emails sent to each of the LTH domains ***eit***, ***cs*** and ***control***. Describe the process of delivering an email to the address nobody@eit.lth.se.

3. What is the default DNS server used by your computer?

Use a proper editor to study the contents of the file **/etc/resolv.conf**. What does it contain? What is the type of the local name servers mentioned in this file? Can they give authoritative answers? How can you verify that? Use dig to verify your answers.

4. Using other DNS servers

Use dig to find the official authoritative name servers for LTH. We are interested to investigate if these servers support recursive queries. Describe how we can do that. Perform the test using dig.

5. Survival time for cached information

In this task we are going to investigate the survival time for cashed information in the name servers. We start by checking the translation for the web server at three Swedish universities namely Lund, Linköping and KTH using the default name server. Start with Lund university, www.lu.se, and look for the translation of its web server few times. Study the answers and figure out the time duration that the information will remain in cache at name servers. Note that TTL is counting down. Repeat the same task for Linköping university, www.liu.se. Afterwards repeat it for KTH, www.kth.se. Compare the TTL parameter for the name servers of these three universities. What is the consequence of such a choice?

As you have seen above, KTH has set a rather low value for the TTL of information. Try successive translation and see that TTL is counting down. What happens after TTL becomes zero? Repeat the process but, this time, use the official name server at KTH. Do you see any difference? If so, explain the reason.

6. Reverse look-up

There are cases that we are in need to recover the computer name corresponding to an IP address. This feature also exists in DNS. Translate your computer's IP address and some other computer addresses you have seen earlier, using the method described in the course book. Afterwards, do the same by just giving the **-x** flag and the IP-address to ***dig***. The program performs all the strange address manipulation for you. Don't be surprised if some reverse look-ups fail. The system administrators assign higher priority to forward look-ups since errors there means that some services will fail while reverse look-ups get less priority.

Good Luck!