

# DATA COMMUNICATION I

## 7,5 HP

---

### LAB I - NAME RESOLUTION

Jakob Ahlin (jakob.ahlin@his.se)

#### AIM

The lab aims to give the student:

- Understanding of the name resolution process in operating systems and applications
- Experience in configuring name servers
- Experience with tools to analyze network traffic
- Understanding of application layer network protocols
- Experience in configuring network related components in Unix operating system

#### REQUIREMENTS

To achieve a G on this assignment, your solution should be presented to one of the supervisors on one of the scheduled lab sessions, and you must turn in a report explaining your solutions.

A template for the lab report is present on the web site, and the report must be turned in before the deadline specified by the instructor.

The course web page is located at: <http://www.his.se/da120g>

# INDEX

<b>1</b>	<b>INTRODUCTION .....</b>	<b>1</b>
<b>2</b>	<b>LAB ENVIRONMENT SETUP .....</b>	<b>1</b>
2.1	SETTING UP THE VIRTUAL MACHINE HOST.....	2
2.2	BASIC OPERATING SYSTEM CONFIGURATION .....	2
2.2.1	<i>General requirements</i> .....	2
2.2.2	<i>Network configuration and hostnames</i> .....	2
2.2.3	<i>Updates</i> .....	3
<b>3</b>	<b>ASSIGNMENTS .....</b>	<b>3</b>
3.1	PART I - NAME RESOLUTION USING HOSTS FILES .....	3
3.2	PART II - DNS NAME RESOLUTION .....	5
3.2.1	<i>DNS infrastructre requirement</i> .....	5
3.2.2	<i>Resolver configuration</i> .....	5
3.2.3	<i>Verify your solution</i> .....	6
3.3	PART III – ADVANCED DNS CONFIGURATION .....	6
3.3.1	<i>Requirements</i> .....	6
3.3.2	<i>Verify your solution</i> .....	6
<b>4</b>	<b>GENERAL TROUBLESHOOTING AND HELP .....</b>	<b>7</b>
4.1	LOGGING.....	7
4.2	UTILITIES/TOOLS FOR TROUBLESHOOTING .....	7
4.3	MANAGING THE BIND DAEMON .....	8
4.4	NAME RESOLUTION CACHING .....	8
<b>5</b>	<b>LITERATURE .....</b>	<b>8</b>
<b>6</b>	<b>INFORMATION ON THE LAB ROOMS AND EQUIPMENT .....</b>	<b>9</b>
6.1	RULES .....	9
6.2	COMPUTERS AND HARDDRIVES .....	9
6.3	NETWORK TOPOLOGY AND LOGICAL ADRESSING .....	9

## BILAGOR

---

## INTRODUCTION

---

In this lab you will set up a name resolution infrastructure using two different strategies; hosts files and DNS. The aim is to gain understanding of the name lookup process in networks and operating systems and experience in configuring name servers and corresponding client operating system settings. You will be setting up your own DNS domain (a0Xlogin.local) with a subdomain (int.a0Xlogin.local), and do various configurations on this system.

You will be provided two physical machines for this lab. However, since this lab requires four systems in total, virtual machines will be used for all server operating systems. Debian Linux 5.0 will be used for the servers and Windows XP Professional for the client computer. .

Chapter 2 contains instructions on how to set up the lab environment, which includes installation and basic configuration of operating systems and the set up of a virtual machine hypervisor.

The assignment is split into the three following parts:

- I) Hosts files name resolution
- II) Basic DNS name resolution
- III) Advanced DNS configuration

After finishing each part, you are required to present your solutions to one of the supervisors before continuing onto the next part. Note that there may be several valid solution to each problem presented in the assignments. Often detailed requirements or instructions are omitted to allow for various solutions to be used. In your report, you will briefly motivate the selected solutions.

---

## 1 LAB ENVIRONMENT SETUP

---

This chapter covers the basic configuration of the operating systems to be set up for the lab assignment. As shown in Figure 1-1, two physical machines and three virtual machines will be used for this lab.

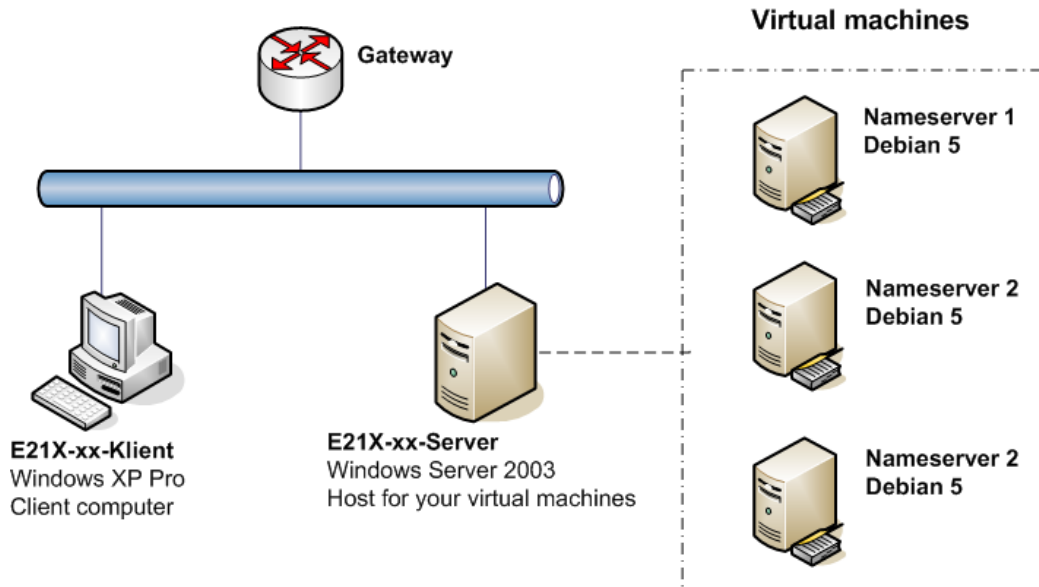


Figure 1-1. Lab topology overview

## 1.1 SETTING UP THE VIRTUAL MACHINE HOST

The virtual machines will run on top of Windows Server 2003 R2 (64-bit edition) using VMware Server version 1.010. You will use the computer labeled ”-Server” for this purpose. You will find the CDs for Windows Server 2003 on in the disc case mounted on the computer labeled ”-klient”. You will find the serial number for this product on a label attached to your computer labeled ”-server”. VMware Server can be fetched from the following URL:

<http://10.0.252.100/VMware-server-installer-1.0.10-203137.zip>

Use the following serial number for VMWare:

930DJ-YAV6Y-1C7C2-421VN

When creating the virtual machines, make sure that the network interfaces of your virtual machines get “bridged” to the physical network interface connected to your lab network.

## 1.2 BASIC OPERATING SYSTEM CONFIGURATION

For the client computer, Windows XP will be used. In addition to acting a resolver (client) for your name resolutions infrastructure, this computer will also be used to search information (e.g read documentation) needed to solve the assignments. On the name servers, run in virtual machines, Debian Linux 5.0 will be used.

The password of all user accounts should be “Syp9393”. Note that the “S” is a capital letter.

### 1.2.1 GENERAL REQUIREMENTS

A general requirement is to use the English language on all operating system installations, but have the locale settings (timezone etc) set to Sweden. The Swedish keyboard layout should also be selected. How you partition your harddrives is up to you, and should not have an impact on the lab assignments. However, make sure that Windows XP has a system partition of at least 20 GB. The Debian systems need at least 4 GB.

### 1.2.2 NETWORK CONFIGURATION AND HOSTNAMES

Each student is assigned a computer group which is isolated in a private subnetwork containing 254 usable addresses (10.<room>.<group>.1-254). Table 1-1 lists the IP addresses and hostnames to be used on each system. When domain name is asked for, use your login name here on the university followed by “.local”, for example ”a04jakah.local”. The gateway must be set to

the first usable address in your subnetwork. Also, configure all systems to use the 10.0.252.100 name server. This setting will however be changed when you set up your own DNS servers later on.

*Table 1-1. IP addressing scheme*

System	IP address	Hostname <sup>1</sup>
Windows Server 2003 (virtual machine host)	10.<roomnr>.<group nr>.2	vmhost
Windows XP Client	10.<roomnr>.<group nr>.5	client
Name Server 1	10.<roomnr>.<group nr>.11	ns1
Name Server 2	10.<roomnr>.<group nr>.12	ns2
Name Server 3	10.<roomnr>.<group nr>.13	ns3

To be able to easily administer the Debian systems without using the VMWare Console, an SSH (Secure SHell) service must be installed on the debian systems. Install this package using this command:

```
# apt-get install openssh-server
```

On both the Windows XP and Windows Server 2003 machines, install the SSH client Putty, which can be downloaded from the following URL:

```
http://www.putty.org/
```

Set up one "stored session" in Putty for each name server, to be able to easily connect to the servers. Shortcuts should be created on the desktop on both Windows machines, used to connect to each name server with Putty using the stored sessions.

### 1.2.3 UPDATES

Install the latest security updates on your Windows machine using Windows Update, and the latest updates for your Debian machines using APT.

---

## 2 ASSIGNMENTS

---

The assignment is split into three parts. One for name resolution using hosts files, one for basic name resolution using DNS and a third for a more advanced DNS configuration. Ask your supervisor to let you present your solutions after each part is done, before you continue on to the next one.

In chapter 3, general help for troubleshooting tips are provided to aid you in the practical configuration of the components to be set up.

### 2.1 PART I - NAME RESOLUTION USING HOSTS FILES

A simple form of name resolution in an IP network is to store mappings between IP addresses and host names in plain text host files. The mappings stored using this approach will only have local significance, i.e. a mapping done on one particular system will not be recognized by other systems.

---

<sup>1</sup> Hostname is sometimes called "computer name" in Windows

Find out where the hosts files are stored in Windows and Debian respectively, and set up each system with host name mappings in Table 2-1.

*Table 2-1. Host name and IP address mappings*

IP address	Hostname
10.<roomnr>.<group nr>.2	vmhost.<loginname>.local.
10.<roomnr>.<group nr>.5	client.<loginname>.local.
10.<roomnr>.<group nr>.11	ns1.<loginname>.local.
10.<roomnr>.<group nr>.12	ns2.<loginname>.local.
10.<roomnr>.<group nr>.13	ns3.<loginname>.local.
10.0.252.201	www.<loginname>.local.
10.0.252.201	webmail.<loginname>.local.
10.0.252.100	fileserver.<loginname>.local.

All systems should be able to contact all four systems using the hostnames defined in Table 2-1. Use a tool such as ping to verify connectivity to all systems from all systems using their hostnames. Now, change the stored sessions in Putty to use the hostnames of the systems instead of their IP addresses and make sure it still works to connect to all of them.

**\*\*\* Present your solution to a supervisor before you continue onto part II \*\*\***

## 2.2 PART II - DNS NAME RESOLUTION

Using hosts files will work fine with a small numbers of systems, where manual replication is not a big deal. However, in a larger network such as the Internet, this approach becomes practically impossible since the name database becomes too large to store and synchronize between all systems on the Internet.

In this part, you will replace the simple name resolution infrastructure using hosts files with a DNS infrastructure. The mappings from Table 2-1 currently stored in local hosts files will now be transferred to a central DNS server. Only one name server, the one called "ns1", will be used for now. The other two name servers will first be used in part III, where more advanced DNS configuration will be performed. In addition to answering name resolutions requests for your "<login>.local" domain, your DNS server should also be able to resolve Internet domain names, for example "google.com". A reverse lookup zone should also be created representing the IP address range of your computer group, to be able to resolve IP addresses into hostnames.

Before you start, save one of the hosts files from at least one of your systems, and then remove all the hosts entries from all the hosts files on all systems. The reason to save one hosts file is that you may need it for the lab report you will be writing.

### 2.2.1 DNS INFRASTRUCTURE REQUIREMENT

The name server software Bind (version 9) will be used for this lab. Install Bind on "ns1" using APT, from the Debian package named "bind9". It is now up to you to configure Bind, given the following requirements:

- From all machines, it should be possible to resolve all hostnames from Table 2-1 to IP addresses (forward name resolution)
- Where two hostnames point at the same IP address, that IP address should only have to be specified once in a zone (file).
- It should also be possible to resolve the same IP addresses into hostnames (reverse/inverse name resolution)
- All machines should also be able to resolve Internet domain names through your DNS server

For this to work, you need to create two zones, and a number of resource records of different types.

For literature on configuring Bind, see the references presented in chapter *4Literature*. Bind, and DNS servers in general, are complicated software to configure. Hence, a lot of things can go wrong. A general tip is to check the syslog main log file for messages when something is not working, located at the following path:

```
/var/log/syslog
```

To list the last lines in a text file (such as a log file), it is recommended to use the "tail" command.

### 2.2.2 RESOLVER CONFIGURATION

Point all your machines to use this name server for DNS name resolutions. On the Debian systems, the DNS server is specified in "/etc/resolv.conf" and on the Windows systems, this is specified on the TCP/IP properties of the network interface (card) in use. Also, make sure that the DNS suffix on each machine is set to your domain name. This tells your system to append your domain name (<login>.local) when you query a non-fully qualified domain name (e.g client rather than client.login.local). For example, this makes it possible to do this:

```
# ping client
reply from 10.218.4.5 (client.a04jakah.local): time=32ms
```

### 2.2.3 VERIFY YOUR SOLUTION

Use the "nslookup" utility on the client to resolve hostnames from your domain and Internet domain names. Also perform reverse lookup on IP addresses in your lab network using the same utility.

Perform the same tests on all your Debian systems. The "nslookup" utility exist on your Debian systems as well. However, the more powerful (and with better output) "dig" utility is recommended to use here.

Now, also verify that hostnames can be derived from knowing an IP address (reverse lookup). Both nslookup and dig can make such queries.

For the oral presentation of your solution, create a simple shell script that makes DNS queries for all relevant resource records.

**\*\*\* Present your solution to a supervisor before you continue onto part III \*\*\***

## 2.3 PART III – ADVANCED DNS CONFIGURATION

In this part, you will extend your DNS infrastructure. For an organization, a common naming strategy is to create a sub domain of the organization's "external" domain, to use for all internal hosts. In this part, a sub domain for your "<login>.local" domain, called "int.<login>.local" will be created, and a number of records will be set up in this domain. Two hostnames for the XP client and the file server, will be "moved down" to this new sub domain. Hence, their fully qualified domain names will be "client.int.<login>.local" and "fileserver.int.<login>.local".

In addition to creating a new sub domain, your solution should provide redundancy for this zone in case of a DNS server failure; two DNS server (ns2 and ns3) will be set up to be authoritative for the zone. Given a number of requirements presented below, it is up to you to design a working solution.

### 2.3.1 REQUIREMENTS

A general requirement is that all of your systems should be able to resolve all host names of your old domain, the new sub domain and also Internet domain names.

Another requirement is that the name servers authoritative for your new sub domain should not answer queries for other domains, except for possibly future sub domains of this one (e.g sub.int.<login>.local).

If one of your two servers authoritative for the sub domain is turned off, all systems should still be able to resolve all names in that zone. However, it does not have to be possible to change zone data if one server is down. Note that you must only provide redundancy for your sub domain, not the parent domain you created in part II.

### 2.3.2 VERIFY YOUR SOLUTION

Make sure that all your systems can make name resolutions for all records stored in your two domains, and Internet domains. Also, it must also be possible to lookup names from the sub domain with at least one of the "ns2" and "ns3" name servers running. Zone transfers are usually performed using either AXFR or IXFR queries. For troubleshooting, these queries can be performed using DNS utilities such as dig and nslookup.

Tip: For zone replication to work as intended, keep in mind that you need be careful in how you set up (and maintain!) your SOA record.

---

### 3 GENERAL TROUBLESHOOTING AND HELP

---

Here are a number of general tips that may help you with the lab assignment.

- Before you make big changes to (or replace) a configuration file, make a backup of the file!
- Read log files when something goes wrong, and find ways to monitor log files efficiently
- Make sure name caches are emptied when troubleshooting name resolution

#### 3.1 LOGGING

Bind is logging events using “syslog”, a common logging daemon in Unix. If for example Bind does not start up because of a syntax error, nothing is usually put out in the screen. Instead, the error messages and details can be found in the syslog main file located at “/var/log/syslog”.

In log files, new events are generally appended to the end of the log file. The Unix command “tail” is therefore handy when reading log files, since it prints the last few lines of a given file.

#### 3.2 UTILITIES/TOOLS FOR TROUBLESHOOTING

A number of utilities that may be of help are listed in Table 3-1. For Unix commands, there is usually a man page to describe the syntax and usage of the command. For Windows commands, the “/?” may display a help text on how to use the command.

*Table 3-1. Tools for troubleshooting*

<b>Tool</b>	<b>Description</b>
Nslookup (Windows/Unix)	A tool used to query DNS servers. Exist both in Windows and Unix systems.
Dig (Unix)	Dig has the same purpose as nslookup but is more advanced and produce a more verbose output. It may also be installed on Windows systems, but is not present by default
Getent (hosts) (Unix)	Lookup of a hostname through the operating system. Will translate between hostnames and IP addresses, independent of source (hosts file or DNS).
Ping (Windows/Unix)	Test connectivity to a host using ICMP
Ipconfig (Windows)	Configure and view network related settings in Windows
Ifconfig (Unix)	Configure and view network settings on a network interface
Tcpdump (Unix)	Can be used to capture all traffic send and received on an interface, to a file.
Wireshark (Windows/Unix)	A graphical tool to capture network traffic sent and received on an interface and to analyze it. It is also possible to analyze the contents of a traffic capture file created by tcpdump using this tool. For example, if you want to capture all traffic send and received from one of your name servers, it is possible to capture it using tcpdump, and transfer it to a Windows machine using Wireshark.
Arp (Windows/Unix)	Display the arp cache, and add/remove entries from it

### 3.3 MANAGING THE BIND DAEMON

Rndc (remote name daemon control) is a protocol, and also a utility, to manage a DNS server. The “rndc” command can be used to e.g start stop and reload Bind. A parameter handy to know is “flush”, which empties the current name resolution cache of Bind:

```
# rndc flush
```

### 3.4 NAME RESOLUTION CACHING

When troubleshooting name resolution, cached lookups (at various levels) may cause trouble. To clear the resolver cache on a Windows system, use the “flushdns” argument to “ipconfig”. Unix/Linux systems most often do not have common name resolution cache in the operating system level. However, individual applications may have caching mechanisms. Closing and starting up for example a web browser may clear its cached entries.

DNS servers, including Bind, also cache name resolution entries. To clear all cached entries of your Bind daemon, use the “flush” argument to RNDNC.

---

## 4 LITERATURE

---

### Computer Networking and the Internet (course literature)

It is **highly** recommended to read section 8.2 on DNS in the course literature “Computer Networking and the Internet”.

### CCNA Network Exploration 4.0 – Network fundamentals (course literature)

Chapter 3.3.1 offers an introduction to DNS.

### The official Bind howto

The official Bind howto contains reference documentation for Bind (configuration/zone file syntax, general help etc). It can be found on the following URL:

```
http://www.bind9.net/manuals
```

### DNS HOWTO

This how-to document contains guides on how to configure Bind in practice (rather than just reference documentation), and also has quite a good explanation of how DNS works. It is quite old, and some of the syntax is not valid to the version of Bind you are using. Hence, don’t “cut and paste” complete example configuration, but rather see it as a good guide/reference for understanding of Bind.

```
http://www.langfeldt.net/DNS-HOWTO/BIND-9/
```

### Man pages

Unix Man pages can be of help for syntax and general use of commands (for example dig), but there are also man pages for many configuration files, such as “named.conf”.

Examples:

```
root@ns1:~# man named.conf
```

---

## 5 INFORMATION ON THE LAB ROOMS AND EQUIPMENT

---

The lab sessions will be held in E210, E212, E214 and E218. To get access to these rooms you need a security badge and a personal code. This badge, and security code, may be picked up at the janitors in the A-building during their opening hours. Remember that you are responsible when you are in the lab rooms, so do not let people without badges inside. You are expected to work by your own, on unscheduled time, in the lab premises to finish the assignment in time.

### 5.1 RULES

The following rules apply in the lab rooms:

1. You are not allowed to bring or eat food or drink fluids in the labs. There are tables for eating both outside in the hall as well next to the cafeteria on floor 1.
2. When you are finished for the day you are expected to turn off your computer, put your harddrives into the shelf and bring all your personal stuff with you. No personal belongings may be left in the rooms.

### 5.2 COMPUTERS AND HARDDRIVES

The computers within the lab rooms are grouped together in pairs. Each group consists of a computer labeled “E-*{room nr}*-*{group nr}*-Server” and “E-*{room nr}*-*{group nr}*-Klient”. To exemplify, with room E218 and computer group 4 the label would look like this: “E218-4-Server” and “E218-4-Klient”.

To make it possible to use the lab rooms for more than one course at a time a system with removable harddrives is used. The removable disks are numbered in pairs with the markings of Sxxx and Kxxx, where xxx is the number of the removable harddrive, for example K054. The numbering of the harddrives is not related to the numbering of the computers. You will get a pair of harddrives at your disposal for the lab, one for the server and one for the client. The harddrives can be found on a shelf in the lab room, and they should be put back there after you are done for the day unless your instructor gives you other instructions. The computer designated as server in each computer group always has two network interface cards installed, while the computers designated as client only have one (except in room E216 where all computers have two network cards).

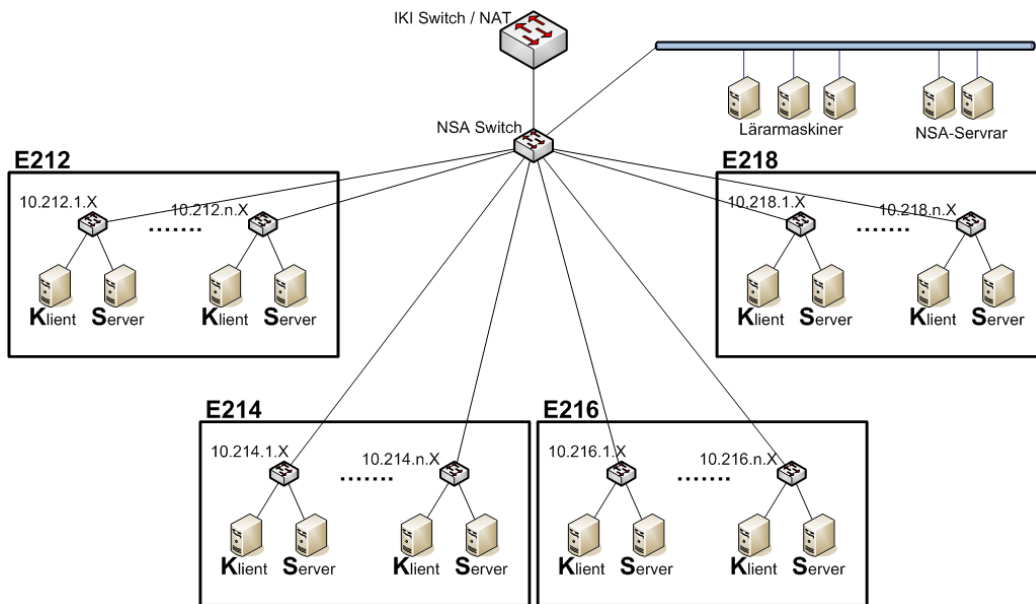
### 5.3 NETWORK TOPOLOGY AND LOGICAL ADDRESSING

Every computer group within the lab rooms are placed in their own network, where the IP settings are based on the name of the computer group. The following IP-settings are used (where room number corresponds to the number in a rooms name, that is for room E212 the room number is 212, and the group number corresponds to the number of the computer group). Tabell 5-1 gives an overview of the IP-addresses and the Figur 1 shows an overview of the network topology.

*Tabell 5-1. IP adressing rules*

Usable address range	10.<sal>.<gnr>.2-254 (e.g 10.212.7.2) Note that the first address (.1) is used for the gateway
Subnetmask	255.255.255.0
Gateway	10.<room>.<group>.1
Name Server (DNS)	10.0.252.100

Note that in this assignment you will set up your own name servers and not directly use the one specified in the table above.



Figur 1. Lab topology overview