

Basic routing lab

Laboration in data communication
OpenIPLab
Department of Information Technology, Uppsala University

Overview

This is a lab constructed as part of the OpenIPLab project.

Administration

Student 1

Name: _____

Email: _____

Personal number: _____

Student 2

Name: _____

Email: _____

Personal number: _____

Agreement

I/we have independently worked on the following assignment solution. In the case of two people, we have both taken part in creating the solution, according to the assignment specification.

Sign 1: _____

Sign 2: _____

General

Course instance (e.g. Datakom DV1): _____

Date: _____

Notes to the lab assistant:

Comments from the lab assistant:

Grade: _____

Sign: _____

1 About the lab

This lab will give you an insight in intra-domain routing. Intra-domain routing is routing performed within a single administrative system (AS), e.g. within a

company. It will concentrate on the network layer (i.e. layer 3 in the OSI model) and does not require any previous knowledge of other layers.

It is a “hands on” lab and you you will be using system that are used today in the Internet.

First you will set up your machine manually, by setting the IP address, netmask, etc. and discover why routing is necessary. You will be given the opportunity to gain knowledge about how to allocate address space in an efficient way and see in practice that your network nodes are able to communicate accordingly to the specification.

You will also be given some theoretical assignments, which might require you to read your excellent data communications book and do some actual thinking.

Enjoy the lab!

Objectives

- IP addressing
- Routing
- Practicing network tools and command usage

Reading instructions

Computer Networking - a top-down approach featuring the Internet (2nd edition)
Kurose, Ross

- 4.1-4.2 (routing)
- 4.4 (IPv4), especially 4.4.3 (packet format)

Theoretical and practical questions

There are a number of questions marked with an asterisk (*) which do not require you to use your computer. I.e., these questions can be done after the lab, in case you run out of time in the lab room.

2 Lab description

General

Background information

IP - the Internet Protocol.

One of the purposes of the Internet is to enable communication between computers. To be able to distinguish between computers/nodes in the Internet, all nodes

have unique IP addresses. An IP (version 4) address is four bytes long and is often written in dotted-decimal notation, e.g. 130.238.8.1.

One could also express the IP address in binary format.

10000010 10001010 00001000 00000001 is the binary representation of 130.238.8.1.

A portion of the IP address (the left part of the binary notation) will determine the network to which the node is connected. The rest of the IP address will determine a unique host on that IP network.

But, as you probably know, there are networks of different sizes, from two nodes to several thousands or even millions of nodes. How can you tell how big the network portion of the IP address is? That is determined by the network mask. In the above network it is determined that 254 computers could belong to this network. So to distinguish between 254 computers in this network we need 8 bits ($2^8 = 256$) of the address to be used for host identification. The right-most 8 bits are used for this purpose and the “rest” 24 bits are used to identify the network. We use something called the netmask to represent how big our network is, i.e., where the boundary is between the network and the host part of the IP address. All bits that are used for the networks are set to 1, the rest to 0.

In our example the netmask is:

11111111 11111111 11111111 00000000

or more commonly written 255.255.255.0.

“Internet - the network of networks”

The way Internet is designed and built is that it consists of many small LAN networks that are interconnected. This design requires some kind of mechanism to send data between networks. Traditionally, this has been done using some kind of forwarding mechanism, in the case of IP, IP forwarding. It’s basically a node which has connectivity to at least two networks, and whose task is to send traffic from one network to another (i.e. to forward data). A node with this task is called a router. In the simple case of two networks this is quite easy, but in more complex network topologies like the Internet, there has to be a systematic way to find a route from the source to the destination – typically by knowledge of which networks other routers are connected to. This can be formalized into a table, the routing table. Every machine on an IP network uses its own routing table to determine where to send packets; often these routing tables are quite simple on normal workstations, but tend to be more complex in routers on the core/backbone network.

IP routing and IP forwarding are two things that are often mixed up, and used interchangeably. IP routing is the process of updating and calculating the routing table used for IP forwarding. Often, a routing protocol such as Routing Information Protocol (RIP) or Open Shortest Path First (OSPF) is used to maintain the routing table. The actual packet forwarding, i.e. receiving a packet on one interface and sending it out on another interface (as defined by the routing table) is called IP

forwarding.

Basically, the motivation why routing/IP forwarding is needed is that one wants to interconnect several networks, such as in the Internet.

Routing in theory

The routing table consists of n rows of
<IP address, netmask, cost, next hop, interface>

For each row, calculate $IP(1) \text{ AND } mask(1)$, $IP(2) \text{ AND } mask(2)$, ..., $IP(n) \text{ AND } mask(n)$. The result will be the network to which the IP address belongs.

When a packet with destination address Y should be delivered, calculate $Y \text{ AND } mask(1)$, $Y \text{ AND } mask(2)$, ..., $Y \text{ AND } mask(n)$.

Now compare, $Y \text{ AND } mask(1)$ to $IP(1) \text{ AND } mask(1)$, $Y \text{ AND } mask(2)$ to $IP(2) \text{ AND } mask(2)$, ..., $Y \text{ AND } mask(n)$ to $IP(n) \text{ AND } mask(n)$.

If $Y \text{ AND } mask(x)$ is the same as $IP(x) \text{ AND } mask(x)$, the row is considered to be a match. If several rows match, use the row whose mask contains the most 1:s in the binary representation.

If no row matches, a default route (if it exists) is used. The default route is where all packets not captured by any other rule is sent to.

next hop is the link-layer address.

IP Networking in Linux

A substantial part of this lab will be about how to setup IP networking in Linux. There are some basic commands that you probably want to use, listed below.

In Linux the network interfaces are called ethX, where X is a number. The first network card is called eth0, the second eth1, etc.

ifconfig - to configure a network interface. Setup a network card, assign an IP address

ifconfig eth0 192.168.200.1 netmask 255.255.255.0 up

route - show / manipulate the IP routing table

route add -net 192.168.205.0 gw 192.168.200.1

ping - send ICMP ECHO_REQUEST packets to network hosts

ping 192.168.205.1

/proc/sys/net/ipv4/ip_forward - enable/disable IP forwarding

echo 1 > /proc/sys/net/ipv4/ip_forward - to enable IP forwarding

Use the man pages (**man <command>**) to retrieve additional information how to use these commands.

Running the system

- Login to the machine using the graphical login program. Enter the username **lab** and the password **lab**.
- Start a terminal window and read further...

You will be running multiple virtual Linux boxes on your computer using User-mode Linux (UML). Each machine will have its own xterm window and the systems are independent of each other. You will (if you setup the networking) be able to communicate with your other UML:s.

To start a UML: run an appropriate script file located in your home directory. E.g. **node_stockholm.sh**, **router_stockholm.sh**, **node_gothenburg.sh**, etc.

To login to your UML use the username **root** and the password **uml**.

Stage 1

Example scenario

A corporate business is located at two sites, Stockholm and Gothenburg. At each location there exists a LAN. Each such LAN can be used to communicate within the site, but not with the other site. A wire is setup between Stockholm and Gothenburg, but it's a leased line which costs per transferred amount of data. The simplest solution is to "extend" the LAN:s with a bridge so it can be seen as ONE big LAN. However, this might not be what we want, as traffic inside an office will have to travel to the other office, causing us to pay for this unnecessary traffic flowing on the leased line.

Another solution is to use routing/IP forwarding. This means that all packets that originate at the Stockholm office, which not are destined for Stockholm, will flow on the leased line to Gothenburg. So traffic inside an office will never leave that office. The way this is done is by using two IP networks.

A router is placed at each location, connected to both the local office network and the remote office network (via the leased line).

The Stockholm network is given the IP network address 10.1.0.0 and uses the netmask 255.255.0.0. This means that it can contain 2^{16} (-2 for technical reasons: one address is used to identify the network itself, and one address is used for broadcasts) nodes/computers. Similarly, the Gothenburg network is given the IP network address 10.2.0.0 and also uses the netmask 255.255.0.0.

Computers in the Stockholm network may have addresses ranging from 10.1.0.1 to 10.1.255.254. Recall that each octet of the IP address can take on 255 different values, so 10.1.353.3 is an example of an illegal address.

Similarly the Gothenburg network address range is 10.2.0.1 to 10.2.255.254.

By using the addressing scheme it is simple to know whether we are communicating within our network or not. If it has the same numbers in the first two bytes of the IP address, we know it's inside our own network. More specifically,

this is the purpose of the netmask. It decides how many bits that should be considered the network part and the host part, respectively. In our example, the netmask is 255.255.0.0. Converting it to binary format gives us 11111111 11111111 00000000 00000000. I.e. the first two bytes are all 1:s, meaning that the first two bytes are indicating the network part of the IP address.

Given an IP address and the netmask we can know which IP network that address belongs to.

E.g. a computer in Stockholm with IP address 10.1.10.10 belongs to network 10.1.0.0. We calculate this by taking the IP address and perform bitwise AND with the netmask. $10.1.10.10 \text{ AND } 255.255.0.0 = 10.1.0.0$.

OK, if we can distinguish which IP network an IP address belongs to, it should be easy to know which traffic we should send over the leased line. Right?

Yes.

In practice

In practice, how do you setup routing?

The first step is to setup IP addresses for all computers in the network. The next step is to setup routing at each node. This is done by stating which computers that we can communicate with directly, i.e., which computers that are located on the same LAN, and what to do with packets destined for other computers. This is often done by setting up routing rules, to tell which path/route the packet should take.

For a node in Stockholm it could look like this:

First we want a rule that says that packet in our network should be sent direct.

route add -net 10.1.0.0 netmask 255.255.0.0

Then for traffic to Gothenburg, all packets destined for 10.2.0.0 should go through our router which have connectivity to the 10.2.0.0 network via the leased line. However, in practice one almost always sets up a different kind of route, a default route.

The default route is used for any packets that do not match any previous rule.

route add -net 10.2.0.0 netmask 255.255.0.0 gw 10.1.0.1

10.1.0.1 is the IP address of the Stockholm router.

Or by using default router/gateway:

route add default gw 10.1.0.1

Similarly, at Gothenburg, the rules will be:

route add -net 10.2.0.0 netmask 255.255.0.0

route add default gw 10.2.0.1

where 10.2.0.1 is the IP address of the Gothenburg router.

Now we should be able to communicate with both networks, and only traffic destined for the other network goes over the expensive and slow leased line.

Assignment 1

The company recently created a new office in Uppsala, which needs to be connected to the other offices in Stockholm and Gothenburg. Uppsala is connected to Stockholm via a leased line, but is not connected to Gothenburg. The office in Uppsala is small, so it will only use 250 computers.

Task: Allocate a suitable IP network and the smallest possible netmask, and show what changes you will have to do in the routing tables at all places (including Stockholm and Gothenburg).

IP network:

Netmask:

Routing table changes
in Stockholm:

in Gothenburg:

in Uppsala:

Assignment 2

The company is offered a free fibre using Gigabit technology between Uppsala and Gothenburg. This means that all traffic preferably should flow over this fibre rather than the expensive and slow leased line. However, we don't want to remove the leased line altogether, in case the fibre line goes down. (The fibre line is rather unreliable, but what do you expect when you get things for free?)

There is something called "metrics" which come in handy in situations like this.

Task: State the changes in the routing tables. Try to change the metric and see if you get the wanted result.

IP network:

Netmask:

Routing table changes
in Stockholm:

in Gothenburg:

in Uppsala:

Assignment 3

Finally the price for Internet connectivity has dropped enough that the company can afford to get an Internet connection! However, to save money, there is only one connection in Stockholm. Based on your recently acquired knowledge of routing, you should have a solution to how both offices can share the Internet connection.

Task: State your solution and the changes in the routing tables.

Hint: default gateway/router

You can ignore the fact that the IP addresses used in this scenario are private. If you don't understand what this means, don't care.

Answer :

Routing table changes

in Stockholm:

in Gothenburg:

in Uppsala:

Assignment 4

Your task is to design a network given the specification below. You should allocate IP addresses in an efficient way (i.e. not wasting addresses - remember that there are only 2^{32} IPv4 addresses in the world). State IP network addresses, netmasks and routing tables. Draw a picture of your network, with lines representing the links, and indicate the IP addresses of the gateways.

Networks: A B C D E

Connectivity between: A-B, A-D, B-E, C-D

Population per network: (excluding extra routers)

A 160

B 200

C 4097

D 10

E 127

Stage 2 - the theoretical part

These are theoretical questions which will probably require you to read the data communication book to be able to answer them. They might require some serious thinking, but do NOT require any computer access. So this part of the lab could be skipped during the lab session.

Assignment 1 - routing protocol design

Gary Scott Malkin said:

“Distance-vector protocols operate by locally distributing global information. Link state protocols operate by globally distributing local information.”

It actually captures the whole idea of those types of routing protocols in two sentences. Give a more exhaustive description of the two protocol types to convince your lab instructor that you have understood what Gary Scott Malkin said.

Assignment 2 - security

Consider using a Virtual Private Network (VPN) to interconnect the offices, i.e. let the traffic flow over the Internet. Obviously, one wants to address security issues here, as confidentiality is important if business secrets are to be distributed over the Internet. One solution to this is to use IPSec, which uses encryption to increase the security. Is this enough?

Compare the usage of a dedicated line for communication with the usage of a VPN over the Internet. Which security benefits and drawbacks do the different solutions imply?

Assignment 3 - applied security

In this assignment you will discover how vital security is to routing protocols.

You want to capture your teacher’s private IRC messages sent from his/her computer. The teacher’s computer is located at 10.100.1.20/24 and the IRC server is at 192.168.200.42/24. There are two routers in between, 10.100.1.1 and 192.168.200.1. One of the routers is located in the server room at the IT department, whereas the other one is located at IT-stöd. These routers are connected using Gigabit Ethernet (copper) through a switch outside the server room at the IT department.

One dark night, you happen to be outside the server room, and the wiring closet (“korskopplingskåpet”) is open - so you can access the switch that connects the two routers. As always, you have your data communication book and your laptop in your backpack, and you are curious about what fun stuff you can do in this tempting situation. Perhaps you could spy on your teacher’s IRC session, without anyone knowing it?

Outline a description of how this can be performed. What equipment is needed?

You happen to know that the sysadmin is really nice guy who isn’t using any of the security mechanisms in RIP.

Inform the sysadmin about how he could improve the security.

3 The report

You should answer all of the questions above. Also make sure that you covered all subquestions.

You may write in English or in Swedish.

Hand in this printed out lab with your answers filled in (in a legible style) to the lab assistant's pigeonhole. This should be done before the date announced at the lab page for your course instance.

4 Marking guidelines

Laboratory results are allocated according to the one of the following grades.

Godkänd	G	You have answered all the related questions in a satisfactory manner.
Komplettering	K	The questions need further work in order to address shortcomings. The lab assistant will contact you with information on how this is done. The code or questions need further work in order to address shortcomings or implementation bugs revealed during the tests associated with marking
Underkänd	U	Failed. Questions have not been answered to the satisfaction of the marker, and the time limit (generally one year from the date when the work was due) for handing in material has expired.