



**FACULTY OF SCIENCE**

**ACADEMY OF COMPUTER SCIENCE AND SOFTWARE  
ENGINEERING**

**MODULE** IT00297  
DATA COMMUNICATION  
**CAMPUS** APK  
**EXAM** 7 JUNE 2021

**DATE** 7 JUNE 2021

<b>ASSESSOR(S)</b>	DR. GJ VAN NIEKERK
<b>INTERNAL MODERATOR</b>	NONE
<b>EXTERNAL MODERATOR</b>	PROF. R VAN OLST
<b>DURATION</b> 2 HOURS 30 MINUTES	<b>MARKS</b> 100

---

**NUMBER OF PAGES: 9 (cover page included)**

**INSTRUCTIONS**  
OPEN BOOK/RESOURCE EXAM

**REQUIREMENTS**  
NONE

1. (3 marks) Support and/or criticize the following statement and give ample reason(s):

“Windows is the only operating system you’ll ever need!”

**Solution:** *(Highly opinionated question, looking for “best tool for the job” argument, but any valid argument worth 3 marks)*

Single OS might have been only needed solution decades back, but today most OSs have unique niche. Windows/macOS common on desktop environment; Linux/Unix common in server architectures, super-computers, cloud. Best tool for the job.

2. (2 marks) Linux Mint is a *fork* of Ubuntu, which is a *fork* of Debian, resulting in three different *distros*. Briefly explain the terms “fork” and “distro” in the context of the Linux ecosystem.

**Solution:** A distribution (or distro) is variation on an existing version of Linux, i.e. based on a parent version of Linux. ✓ A fork is often used to refer to the fact that a new development/maintenance path is taken from a specific point in the parent distro’s development. ✓

3. Consider each of the following commands and briefly explain what each does:

- (a) (2 marks) `$ mv myfile.txt ~/textfiles/`

**Solution:** Moves the file `myfile.txt` to the directory `textfiles` located under the current user’s home directory.

- (b) (2 marks) `$ tail -f /var/log/syslog | grep -i error`

**Solution:** Continuously checks the system logs for any errors (i.e. follow log), by searching for the keyword “error” (ignoring case).

- (c) (3 marks) `$ find ~ -regex .*mp[234] 2> /dev/null 1> songs`

**Solution:** Finds all mp2, mp3 and mp4 files under the current user’s home directory and write the results of this search to the `songs` file. All errors are ignored by sending to `/dev/null`.

- (d) (3 marks) `$ cat a1.txt a2.txt | sort > a3.txt`

**Solution:** Concatenates the content of a1.txt and a2.txt, sorts the results and writes to a3.txt.

- (e) (2 marks) `$ sudo rm -rf /home`

**Solution:** Forcefully removes the home directory of all users. The command is temporarily run with admin rights (given through sudo).

- (f) (2 marks) `$ sudo apt-get install ipcalc`

**Solution:** Installs the ipcalc tool on the computer. The command is temporarily run with admin rights (given through sudo).

4. Consider the following C-program that is compiled under Linux and answer the questions that follows:

```
#include <stdio.h>
#include <unistd.h>

int main{}
{
    int returnVal = fork();
    printf("Return value: %i\n.", returnVal);
    return 1;
}
```

- (a) (1 mark) Explain the purpose of the `fork()` system call.

**Solution:** Creates a duplicate of the current process, but with own PID and memory space. ✓

- (b) (6 marks) Give and explain the output that will be displayed when this program is executed under Linux.

**Solution:** (*Two marks for the output, four marks for the explanation.*)

Output will be:

Return value: <some positive integer>

Return value: 0

✓✓

The first line is the original parent process displaying the PID of the child process that was returned to it; ✓✓the second line is the newly created child process

that displays 0. ✓✓The ordering is arbitrary. Both processes then terminate, returning 1.

5. Consider the following script and answer the questions that follow:

```
1:  #!/bin/bash
2:  function show_usage {
3:      echo "Usage: $0 <textfile>"
4:      exit 1
5:  }
6:  if [ $# -ne 1 ] || [ ! -f $1 ]; then
7:      show_usage
8:  fi

9:  exec 0<$1
10: counter=1
11: while read line; do
12:     echo "$counter: $line"
13:     counter=$((counter+1))
14: done
```

- (a) (2 marks) Explain what this script can be used for.

**Solution:** Script accepts a text file as parameter ✓and outputs the file with line numbers.✓

- (b) (10 marks) Explain in detail how the script works by referring to the line numbers.

**Solution:** (1) The shebang indicating the interpreter for the script that follows; ✓(2) Defines a function called “show\_usage”; ✓(3) Outputs the name of the script followed by <textfile> to standard output; ✓(4) Exits the script with error code 1; ✓(6–8) If there is not exactly one parameter or it is not a regular file, the script calls the function “show\_usage” (which will also exit); ✓(9) Redirect the contents of the text file as standard input; ✓(10) Define variable counter and sets its value to 1. ✓(11) For every line in the file; ✓(12) Output the current contents of “counter” followed by the line; ✓(13) Increase counter by 1. ✓

- (c) (3 marks) The text file “names” contains people’s first names on separate lines. Assuming that the script above is called “number.sh”, the following command is executed:

```
$ ./number.sh names | sort -k 2 | tee newnames
```

Explain what this command will do (HINT: the `-k 2` parameter sorts using the second field, where fields are separated with white space).

**Solution:** Opens the text file and numbers the names it contains; output is then sorted using the second field (to ignore the line numbers) and ✓, finally, the result is stored in newnames (through tee command) and written to standard output. ✓✓

- (d) (2 marks) In order to execute the script without specifying bash, it needs to be executable. Give the Linux command to make it executable for the owner:

**Solution:** `chmod u+x number.sh` ✓ (octal notation: `chmod 744 number.sh` will also suffice)

6. Consider the network architecture depicted in the Figure 1. It consists of **five** subnets as shown (192.168.11.0/24 to 192.168.15.0/24). The subnets are interconnected with **six** routers, marked “Router 1” to “Router 6”. Only the last two octets for each IP is shown to avoid clutter (i.e. 192.168 is omitted for each). All nodes have a single network interface (eth0) with associated IP address.

Answer the questions that follow:

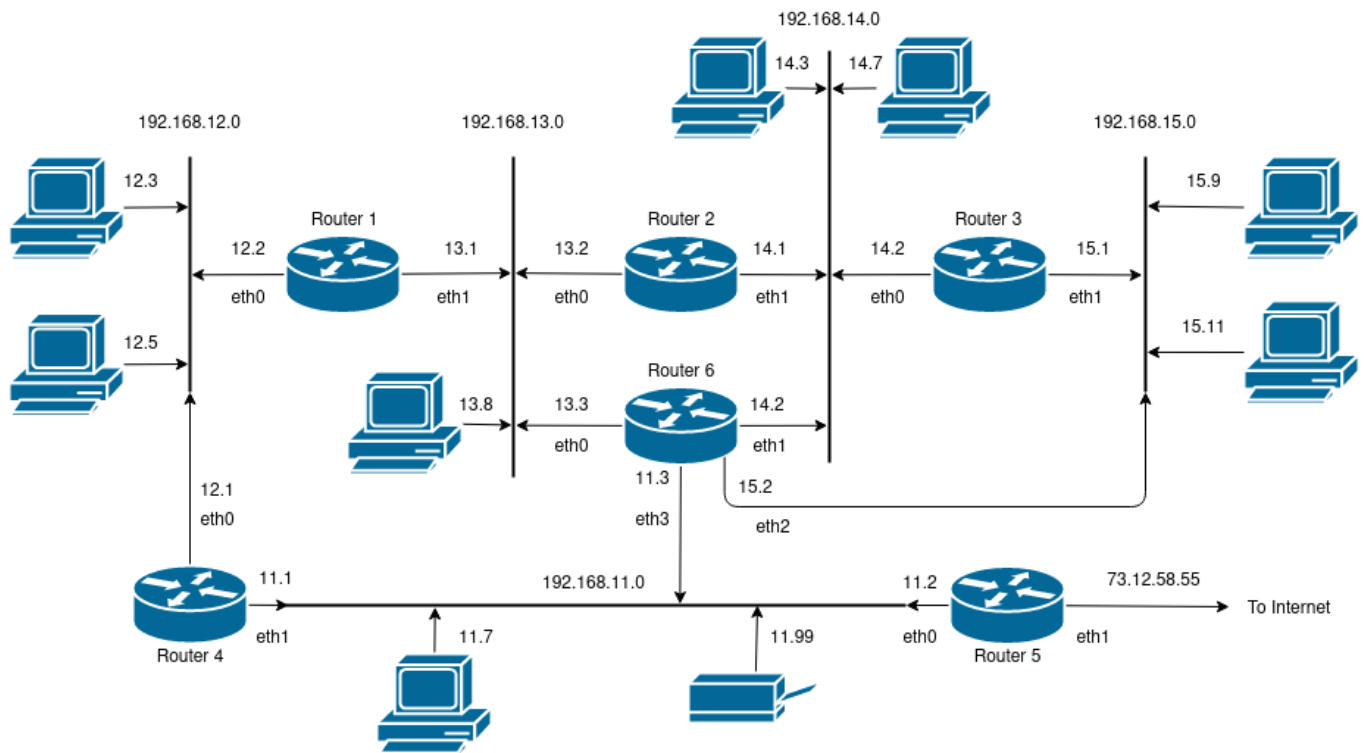


Figure 1: Network diagram

- (a) (4 marks) A junior assistant proposes to change the first three octets to “1.1.x” so that it is easier to remember (i.e. 192.168.11.0 should be 1.1.1.0 and 192.168.12.0 should be 1.1.2.0 and so on) . Will the network still work and is this a good idea? Motivate clearly.

**Solution:** Network should still work, ✓but is a bad idea. ✓The 192.168.x.x IP range is reserved for private networks as per RFC 1918. ✓Furthermore 1.1.1.1 is a known address for a DNS service, so this service will no longer be available from inside the network. ✓

- (b) (16 marks) Assuming hop count as efficiency metric, provide the minimal, yet optimal routing table for **Router 6** and **node 192.168.13.8**. Provide the value for fields: Destination, Gateway, Genmask and Iface.

**Solution:** (Two marks per row, subtract mark for each erroneous/dubious value, i.e. 2 incorrect values per row will result in 0)

(Should have been 18 marks, marked up to 16, negatively marked erroneous entries; also multiple routes are possibly and each student should be marked on merit, subtract marks for sub-optimal paths. Also, 192.168.14.2 is duplicated!! Worked around this error.)

Router 6 (max 10 marks):

Destination	Gateway	Genmask	Iface
192.168.11.0	0.0.0.0	255.255.255.0	eth3
192.168.12.0	192.168.13.1	255.255.255.0	eth0
192.168.13.0	0.0.0.0	255.255.255.0	eth0
192.168.14.0	0.0.0.0	255.255.255.0	eth1
192.168.15.0	0.0.0.0	255.255.255.0	eth2
0.0.0.0	192.168.11.2	0.0.0.0	eth3

Node 192.168.13.8 (max 6 marks):

Destination	Gateway	Genmask	Iface
192.168.12.0	192.168.13.1	255.255.255.0	eth0
192.168.13.0	0.0.0.0	255.255.255.0	eth0
0.0.0.0	192.168.13.3	0.0.0.0	eth0

- (c) Node 192.168.14.3 has the following naive single-line routing table:

Destination	Gateway	Genmask	Iface
0.0.0.0	192.168.14.2	0.0.0.0	eth0

- i. (2 marks) Will the node have Internet access? Motivate.

**Solution:** Yes, ✓ the default route will forward to Router 6 which will route the traffic correctly to R5 and to the external location. ✓

- ii. (2 marks) Explain the route that will be followed if the node communicates with 192.168.15.9.

**Solution:** Message will be sent to Router 6 ✓, which will forward the packet to 192.168.15.0. ✓

- iii. (4 marks) Assuming the ICMP redirect is configured, explain what will happen if the node sends a message to Node 192.168.14.7.

**Solution:** Message will reach Router 6 ✓, which will place packet back on 192.168.14.0 where 14.7 will pick it up. ✓ It will send a ICMP redirect to 14.3 ✓ to update its routing table to send messages destined for 14.7 directly to 14.7 and not to the router. 14.3 can then update the routing table or ignore. ✓

- (d) (4 marks) A webserver is installed on network 192.168.12.0 with IP 192.168.12.3 (as shown). The node is correctly configured to accept external connections (in fact, any node on any of the sub-networks can access the pages on the webserver), however, it is completely inaccessible from outside the network. Explain the reason for this and propose workarounds.

**Solution:** The NAT firewall translates internal IP addresses to the external address and vice versa. From the outside only 73.12.58.55 is visible. There is no way to address the webserver directly as it only has a RFC1918 private address. ✓✓ Port forwarding can be used on the firewall to redirect connections 76.211.216.12:80 to the internal webserver. ✓✓

7. (6 marks) A legacy client/server application communicates insecurely using a TCP connection. The legacy application listens on static IP 76.118.12.20 on port 9900. Briefly explain how communication can be made temporarily secure using “ssh”. Give any commands that needs to be executed.

**Solution: (Corrected ssh command)** Setting up a secure tunnel will solve the problem. ✓ The `ssh` utility needs to be installed on both the client and the server machine. ✓ The tunnel is set up by running the following on the client:  
`ssh -L 9900:76.118.12.20:9900 user@76.118.12.20.` ✓✓✓ The client needs to connect to `localhost:9900` in the usual manner. The traffic will then be redirected over the secure tunnel. ✓

8. (10 marks) You are near a water hole in a nature reserve. Wild life frequent the water hole at various times throughout the day. You have access to a good camera, connected to a Raspberry PI which in turn is 3G enabled. Although the Internet is reliable and responsive, it's relatively slow. The camera can be controlled via the command line, using the following command which will take a photo and save it in the current directory as web-cam-shot.jpg:

```
$ fswebcam -r 1024x768 --jpeg 85 -D 1 web-cam-shot.jpg
```

The main problem you are facing is that taking constant photographs (say 1 every 10 seconds) takes up too much storage and searching through them is time consuming. Using the technology described above, come up with a solution to make the most of this setup. Clearly describe any software that would need to install.

**Solution:** *A very open-ended question that should be evaluated on merit.*

An outline given below:

Installation of a webserver like Apache hosting simple website; configuring cron to take pictures periodically (or a python script that takes photos every x seconds); connecting a simple motion sensor to trigger the taking of a photo through the gpio interface; using an online service to identify animals could be considered; otherwise using technology like “opencv” to identify animals or changes in consecutive images; notifications through email or other mechanism such as “Pushover” or “Pushsafer”; publish positively identified images to web site.

9. (4 marks) Give **two** reasons for and **two** reason against using a cloud-based solution as oppose to an on-premise solution.

**Solution:** *(Multiple possibilities, should be evaluated on own merit)*

- (+) For implementations that needs to scale on-demand, automatically based on time-of-day.
- (+) Easier collaboration between users/vendors.
- (+) Out-of-the-box security.
- (+) Simplified application hosting.
- (+) Rapid implementation and time-to-value.
- (-) Some industries are highly regulated and does not allow data to leave the country.

- (-) High security requirements.
- (-) Environments with poor/no Internet connectivity.
- (-) If the current environment works well then moving for the sake of moving is not a good idea.

10. (5 marks) The 7-layer OSI model is based on a technique called “enveloping” or “encapsulation”. Briefly discuss what enveloping is within the OSI model context and two advantages of this technique.

**Solution:** Protocols are organized in a stack. As the message passes down the stack, the entire packet is placed into a new packet with a new header (the previous message becoming the entire payload/body of the next protocol). ✓✓ After transmission, the protocol moves upwards through the stack and each protocol removes its header information and passes the message on. From the perspective of the protocol, it communicates with the same protocol on the other side being agnostic to protocols above or below the stack. ✓✓

The primary advantage is the separation of concern (each protocol typically responsible for one aspect) which allows protocols to be swapped in/out based based on the needs of the environment. ✓