NAPIER UNIVERSITY

SCHOOL OF COMPUTING

FIRST DIET (SEMESTER ONE) EXAMINATION

SESSION 2001-2002

MODULE: CO32010

NETWORK OPERATING SYSTEMS (SOLUTIONS)

DATE: **DURATION: 2 HOURS**

START TIME:

EXAMINER(S)

DR. W.BUCHANAN
MR. J.JACKSON

QUESTION PAPER DATA

Number of pages - 6
Number of questions - 5
Number of sections - ONE

INSTRUCTION TO CANDIDATES

Complete any three of the questions from five.

Module Organiser

Student Workload

Lectures/Tutorials	24 hours
Practicals/Project Work	12 hours
Examination	2 hours

Assessment

Examination	40% (Any 3 from 5)
Project	60%

Aims

The main learning outcomes are:

A1 Understand the basic concepts in implementing security methods, and develop practical solutions in a network operating system.
A2 Identify problems in distributed processing problems, and how they can be overcome.
A3 Develop an in-depth understanding of routing protocol methods, their associated drawbacks, and how they can be overcome.
A4 Contrast differing distributed files systems, and their associated characteristics.
A5 Develop an in-depth understanding of the mechanics of operating systems, such as process control and scheduling, and how these elements can operate over a network.

Module content

The areas covered are:

C1 **Process Control/scheduling**. Multitasking and threading.
C2 **Distributed Processing**. Interprocess communications, RPC. Deadlock.
C3 **Distributed File Systems**. NFS, NIS, XDR, Active Directory Structure.
C4 **Novell NetWare**. Protocols, NetWare 3, NDS.
C5 **Microsoft Window**s. Domain set-up, structure, peer-to-peer networks, server/workstation configuration.
C6 **Security**. Firewall configuration, Encryption methods, Public-key encryption.
C7 **Routing protocols**. Link-state, Distant vector, Routing protocol problems.

Exam/coursework coverage

The following table outlines how the examination and coursework covers the aims and content of the module.

	AIMS					CONTENT						
Q	**A1**	**A2**	**A3**	**A4**	**A5**	**C1**	**C2**	**C3**	**C4**	**C5**	**C6**	**C7**
1		✓	✓		✓		✓					✓
2	✓										✓	
3					✓			✓				
4	✓	✓		✓				✓			✓	
5		✓		✓	✓	✓	✓	✓	✓	✓		
C/W	✓			✓	✓	✓	✓	✓	✓	✓	✓	✓

1 (a) Define and explain the two major problems with distance-vector routing protocols, and the methods that can be used to overcome these problems. **(10)**

Part	Sample answer	Marking schedule
i	**Routing Loops.** These occur when slow convergence causes inconsistent routing entities when a new configuration occurs. In the case of the diagram in the notes, Network A becomes unavailable. Router V will report this to Router Y, which will then report to Router Z and Router X. Both Routers X and Z will stop routing to Network A. Unfortunately Router W still thinks it can reach Network A with 3 hops, thus Router Z will receive information that says that Router W can get to Network A in 3 hops, and that it is unreachable from Router Y. Thus Router Z updates its routing table so that Network A is reachable in 4 hops, and that the next router to the destination is Router W. Router Z will then send its updated information to Router Y which informs it that there is a path to Network A from Router Z to Router W, and so on. Router Y will then inform Router X, and so on. Thus any data packet which is destined for Network A will now loop around the loop from Router Z to Router W to Router X to Router Y to Router Z, and so on. [2]	[2] – Basic operation
	Counting to Infinity. Data packets can loop around forever, because of incorrect routing information. In this loop, the distance-vector of the hop count will increment each time the packet goes through a router. [2]	[2] – Basic operation
ii	**Setting infinity values.** The count-to-infinity will eventually resolve itself when the routers have counted to infinity (as infinity will be constrained with the maximum definable value), but while the network is counting to this value, the routing information will be incorrect. To reduce the time that it takes to get to this maximum, a maximum value is normally defined. In RIP this value is set at 16 hops for hop-count distance-vectors, thus the maximum number of hops that can occur is **15**. This leads to a problem in that a destination which has a distance of more than 15 hops is unreachable, as a value of 16 or more defines that the network is unreachable. [2]	[1] – Name [1] – Basic operation
	Split horizon. This method tries to overcome routing loops. With this routers do not update their routing table with information on a destination if they know that the network is already connected to the router (that is, the router knows more about the state of the network than any other router, as it connects to it). [2]	[1] – Name [1] – Basic operation
	Hold-Down Timers. This method overcomes the count-to-infinity problem. With a hold-time time, a router starts a hold-time timer when it receives an update from a neighbour indicating that a previously accessible network is now inaccessible. It also marks the route as inaccessible. [2]	[1] – Name [1] – Basic operation

(b) How do interior routing protocols differ from exterior routing protocols? What advantage(s) do exterior routing protocols have over interior routing protocols? (8)

Part	Sample answer	Marking schedule
i	Interior routing protocols, such as RIP and IGRP, build up routing tables based on the best metric value to get to a distant network [1]. Exterior routing protocols type route from one autonomous domain to another [1], and are mainly interested in connectivity, rather than best metrics [1]. An interior routing protocol typically sends the complete routing tables to neighbours [1], which is not good for network security [1]. An exterior routing protocol has a gateway, in which limits the external transmission of the network connections [1]. Exterior routing protocols also reduce the complexity of the Internet [1], as they split the high-level structure into domains [1].	[3]- Defining the interior and exterior protocols. [5]- Contrasting interior and exterior methods.

(c) Describe the four conditions which must occur for deadlock to occur, and give a practical example of this. (7)

Part	Sample answer	Marking schedule
i	• **Mutual exclusion condition.** This is where processes get exclusive control of required resources, and will not yield the resource to any other process.[1] • **Wait for condition.** This is where processes keep exclusive control of acquired resources while waiting for additional resources. [1] • **No pre-emption condition.** This is where resources cannot be removed from the processes which have gained them, until they have completed their access on them. [1] • **Circular wait condition.** This is a circular chain of processes on which each process holds one or more resources that are requested by the next process in the chain. [1] An example of exclusive access on devices is serial ports, and printers [1]. Thus an example would be when one process is writing to the serial port [1], but is printing out results, another process might want to gain access to the serial port, but currently has access to the printer [1].	[4] - Standard bookwork for the first part. [3] - Three marks awarded for the student thinking of their own example. An important factor is for the student to select resources which cannot be shared.

2 For the running configuration of a router given in Figure Q2:

(a) Show the programming steps to set the login password for a telnet session. (3)
(b) Determine the number of subnets that can connect to the network which connects to the `Ethernet0` connection, and also the number of hosts that can connect on each subnet. (8)
(c) Explain the ACL restrictions placed on the `Ethernet0` port. (5)
(d) Design an ACL for the `Ethernet0` port which blocks access to an FTP server (174.10.20.10) for all the addresses for the hosts from 180.70.1.128 to 180.70.1.254, but allows access for all the other hosts. (7)
(e) Explain the problems that an incorrectly set subnet mask on one of the ports would cause for the network. (2)

Total Marks [25]

Line no.	Router program
1	`version 12.0`
2	`service timestamps debug uptime`
3	`service timestamps log uptime`
4	`no service password-encryption`
5	`!`
6	`hostname my-router`
7	`!`
8	`enable secret 5 $1$op7P$LCHOURx5hc4Mns741ORvl/`
9	`!`
10	`ip subnet-zero`
11	`!`
12	`interface Ethernet0`
13	` ip address 180.2.1.130 255.255.255.192`
14	` ip access-group 104 in`
15	`!`
16	`interface Serial0`
17	` ip address 180.70.1.2 255.255.255.0`
18	` ip access-group 102 in`
19	` encapsulation ppp`
20	`!`
21	`router igrp 111`
22	` network 180.2.0.0`
23	` network 180.70.0.0`
24	`!`
26	`access-list 100 deny   ip host 180.2.1.134 host 180.70.1.1`
27	`access-list 100 permit ip any any`
28	`access-list 101 deny   tcp 180.2.1.128 0.0.0.63 host 180.70.1.1 eq www`
29	`access-list 101 permit ip any any`
30	`access-list 102 deny   tcp 180.2.1.128 0.0.0.63 180.70.1.0 0.0.0.255 eq www`
31	`access-list 102 permit ip any any`
32	`access-list 103 deny   ip 180.70.1.0 0.0.0.255 180.2.1.128 0.0.0.63`
33	`access-list 103 permit ip any any`
34	`access-list 104 deny   tcp 180.2.1.0 0.0.0.254 host 180.70.1.1 eq telnet`
35	`access-list 104 permit ip any any`
36	`!`
37	`line con 0`
38	` transport input none`
39	`line aux 0`
40	`line vty 0 4`
41	`!`

Figure Q2: Router program

Part	Sample answer	Marking schedule

a	`config t`
	`line vty 0 4`
	`password mypass`
	`login`
	`exit`
	`exit`
		[1] – config t
		[1] – line vty 0 4
		[0.5] – password mypass
		[0.5] – login

b	Subnet mask is `255.255.255.192` which is 255.255.1111 1111.1100 0000 [2] IP address is `180.2.1.130` This is a Class B address [1], thus there are 10 bits for the subnet [1], and 6 bits for the host [1]. Thus 2^{10}–2 subnets [1], and 2^{6}–2 for the number of hosts per subnet [1].	[8] See sample question.

c	`access-list 104 deny   tcp 180.2.1.0 0.0.0.254 host 180.70.1.1 eq telnet`
	`access-list 104 permit ip any any`
	
	This denies all hosts with an even IP address [1] from the 180.2.1.0 subnet [1] to the 180.70.1.1 telnet server [1], on the incoming port [1]. Everything else is accepted [1].
		[5] See sample question.

d	`access-list 105 deny   tcp 180.70.1.0 0.0.0.128 host 174.10.20.10 eq ftp` `access-list 105 permit ip any any`	[1] – Basic statement [3] – 0.0.0.128 [1] – Explain 128 part [1] – 174.10.20.10 eq ftp [1] – permit ip any any
e	The network part would be confused with the host part, and vice versa.	[2]

3 (a) Figure Q3.1 shows the result of running the command `ps -ef` on the Unix system `mars`.

(i) By examining the process list determine whether this system is a network file system *client* or *server*, or both. (3)

(ii) List the processes from the list which are critical to the operation of the network file system and describe the main function of each one. (5)

(b) Describe the sequence of client/server operations required to carry out a remote procedure call, for example writing data to a remotely mounted network file system. (6)

(c) Figure Q3.2 shows an outline directory structure for three networked Unix systems (`mercury`, `venus` and `mars`) each is both NFS server and client. It is desired that any one of three users (`anne`, `bob` or `colin`) should be able to login to any of the three systems and see their own files using the same pathname. In other words the perceived location of each users' files is not dependent on the machine they are using. Detail the configuration actions required to achieve this. (11)

Line no.	UNIX processes
1	`mars:~ > ps -ef`
2	`UID   PID  PPID  C    STIME TTY      TIME CMD`
3	`root     0     0  0   Oct 15 ?        0:01 sched`
4	`root     1     0  0   Oct 15 ?        0:15 /etc/init -`
5	`root     2     0  0   Oct 15 ?        0:00 pageout`
6	`root     3     0  0   Oct 15 ?        5:54 fsflush`
7	`root   205     1  0   Oct 15 ?        0:00 /usr/lib/utmpd`
8	`root   107     1  0   Oct 15 ?        0:07 /usr/sbin/in.routed -q`
9	`root   246     1  0   Oct 15 ?        0:00 /usr/lib/saf/sac -t 300`
10	`root    49     1  0   Oct 15 ?        0:00 /usr/lib/devfsadm/devfseventd`
11	`root    51     1  0   Oct 15 ?        0:00 /usr/lib/devfsadm/devfsadmd`
12	`root   113     1  0   Oct 15 ?        0:00 /usr/sbin/rpcbind`
13	`root   173     1  0   Oct 15 ?        0:02 /usr/sbin/syslogd`
14	`root   115     1  0   Oct 15 ?        0:01 /usr/sbin/keyserv`
15	`root   200     1  0   Oct 15 ?        0:00 /usr/lib/power/powerd`
16	`root   162     1  0   Oct 15 ?        9:35 /usr/lib/autofs/automountd`
17	`root   209     1  0   Oct 15 ?        0:01 /usr/sbin/vold`
18	`root   156     1  0   Oct 15 ?        0:03 /usr/sbin/inetd -s -t`
19	`root   151     1  0   Oct 15 ?        0:00 /usr/lib/nfs/lockd`
20	`root   192     1  0   Oct 15 ?        0:17 /usr/sbin/nscd`
21	`daemon   150     1  0   Oct 15 ?        0:00 /usr/lib/nfs/statd`
22	`root   249   246  0   Oct 15 ?        0:00 /usr/lib/saf/ttymon`
23	`root   239     1  0   Oct 15 ?        0:00 /usr/dt/bin/dtlogin -daemon`
24	`root   250   230  0   Oct 15 ?        1:57 mibiisa -r`
26	`root   377   156  0   Oct 15 ?        0:00 in.telnetd`
27	`root 13040 13014  1 14:07:56 pts/0    0:00 ps -ef`
28	`root  1710   156  0   Oct 16 ?        0:00 in.telnetd`
29	`root 13006   156  0 14:03:25 ?        0:00 in.telnetd`
30	`root 25389   156  0   Oct 16 ?        0:00 in.telnetd`
31	`root 26096   156  0   Oct 16 ?        0:00 in.telnetd`
32	`root  3212   156  0   Oct 16 ?        0:01 sadmind`
33	`jim  13014 13006  0 14:03:35 pts/0    0:00 -tcsh`

Figure Q3.1: UNIX processes

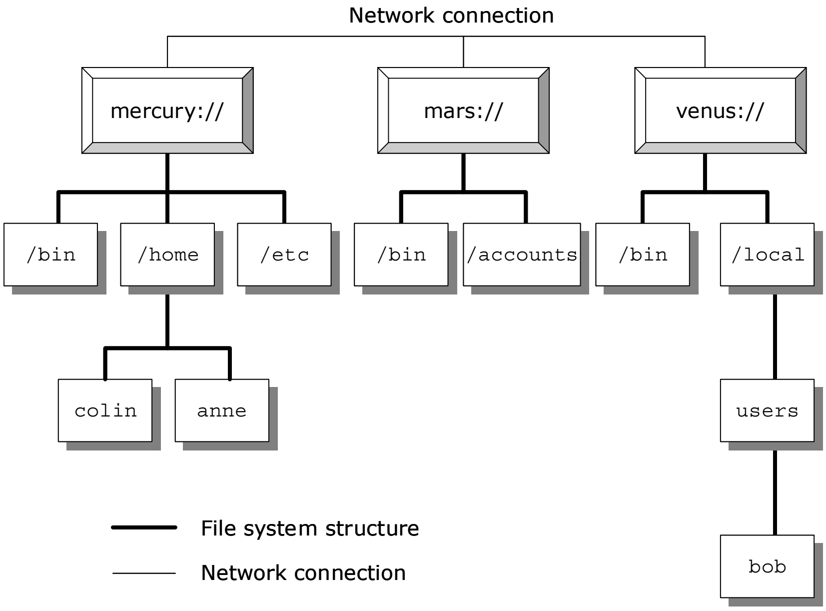

Figure Q3.2: UNIX processes

4 (a) Describe the role and benefits of the NIS (network information service). (8)

(b) Prior to the transmission of a message across a network, explain why it can be beneficial to encrypt it using:

(i) The sender's private key (5).
(ii) The recipient's public key (5).

(c) Using the RSA private key pair E,n (5,21) show how the trivial encrypted message 16 can be decrypted back into is original value. (7)

Total Marks [25]

5 (a) Explain the operation of NFS, and reasons for its architecture. How does NFS allow different data representations to be represented, and how does the client run processes on the server. (10)

Part	Sample answer	Marking schedule
i	NFS uses a client/server architecture where a computer can act as an NFS client, an NFS server or both [1]. An NFS client makes requests to access data and files on servers; the server then makes that specific resource available to the client. [1] NFS servers are passive and stateless. [1] They wait for requests from clients and they maintain no information on the client. One advantage of servers being stateless is that it is possible to reboot servers without adverse consequences to the client. [1] External data representation (XDR) is a universal data representation used by all nodes [1]. It provides a common data representation if applications are to run transparently on a heterogeneous network or if data is to be shared among heterogeneous systems [1]. Each node translates machine-dependent data formats to XDR format when sending and translating data [1]. It is XDR that enables heterogeneous nodes and operating systems to communicate with each other over the network. [1] Remote Procedure Call (RPC) provides the ability for clients to transparently execute procedures on remote systems of the network. NFS services run on top of the RPC. [2]	[10] - Mostly bookwork, but some discussion on the reasons for the NFS architecture.

(b) Outline a major weakness of the Windows NT domain structure that is overcome in NDS and NFS. (5)

Part	Sample answer	Marking schedule
i	The NT domain structure does not allow for a hierarchical structure [1], as nodes within a domain are mounted onto the top level (\\) [1]. This makes it difficult is find resources [1]. NDS and NFS mount devices onto a hierarchical structure [1], which creates a global file system [1].	[5] - Students should be able to identify the weakness in the domain structure, and contrast it with NDS and NFS.

(c) An operating system is currently running four processes (A-D), which each require a specified amount of memory of memory to complete. The total amount of memory on the system is 100 MB, and the total current requirements for the processes are:

Process A requires 70 MB of memory.
Process B requires 50 MB of memory.
Process C requires 40 MB of memory.
Process D requires 20 MB of memory.

The current allocation is:

Process	**Current allocation**
A	50 MB
B	20 MB
C	10 MB
D	10 MB

Show how deadlock could occur in this case, and how the Banker's algorithm could overcome this. What is the major drawback of the Banker's algorithm? (10)

Part	Sample answer	Marking schedule
i	Current allocation has only 10MB of resource allocation left [1]. If any of the processes A-C requires any more memory then none of the processes can complete [1]. The Bankers algorithm would not allow the allocation of memory to the processes, apart from D, as none of the processes could complete. [1] A possible solution would be: Process D gets the remaining 10MB it needs [1], and then releases 20MB [1]. Only Process A would then be allocated the 20MB of memory [1], as the other two processes could not run to completion [1]. When complete it would release 70MB [1], of which any of the other processes can use to complete [1]. The major drawback is that the algorithm requires that each process requires to define the maximum amount of resource that it requires [1].	[10] - Students should carefully identify the key decisions that would be made in allowing all the processes to complete.