PRINCE OF SONGKLA UNIVERSITY

FACULTY OF ENGINEERING

Final Examination: Semester 2 **Academic Year:** 2008-2009

Date: February 25, 2009 **Time:** 09:00 – 12:00

Subject Number: 240-362 **Room:** R300

Subject Title: Internet Engineering

Name: _____

Student Number: _____ Signature: _____

Exam Duration: 3 hours

This paper has 16 pages (including this page).

- Write the answers in the spaces provided in the examination paper.
- Clearly write your student number in the space provided at the top of each page. Write your name, and sign, in the spaces provided on this cover page.
- There are 100 marks total for this exam. This will contribute 50% of the course total.

Authorised Materials:

- Anything the student can carry (except communication devices.)

Instructions to Students:

- Attempt all 9 questions.
- Anything illegible is incorrect.
- Show all calculations, not just the final result.
- Answer briefly where possible, essays are **not** required. There is no need to use all of the space provided for each answer!
- The marks allocated for each question are shown next to that question. There are 100 marks total for this examination.
- *Answer questions in English.* Good English is **not** required.

Question 1. *(10 marks)*

The list following gives a number of explanations for various networking events, scenarios, decisions, etc. Each entry in the list is numbered. On the following page there are a number of statements. Every one of those statements is false.

For each statement, choose the explanation from the list that best, in your opinion, explains why that statement is false. Write the number of the explanation you have chosen in the box appearing next to the false statement.

There are more possible explanations than statements, so obviously some of the explanations will not be used in answering this question. Further, it is possible that one, or more, of the explanations might be the correct answer to explain why more than one, or even all, of the statements are false.

(That is: you can use the same explanation more than once.)

The Explanations

1) ATM has a 48 octet cell size.
2) It is possible to manually configure everything.
3) A lost fragment wastes bandwidth transmitting the remaining fragments.
4) Too much state can cause busy routers to fail.
5) Tunnelling cases a reduction in available MTU.
6) Two IP addresses that differ only in the local parts should be connected to the same network.
7) Link local IPv6 addresses are all that is needed for communications when there is no router.
8) Only IPv6 has Router Advertisements.
9) TCP Sequence and Acknowledgment numbers cycle and continue forever.
10) Key distribution for symmetric key algorithms can be a difficult problem.
11) It is impossible to manually configure everything.
12) Lost packets help TCP adapt to network conditions.
13) An IPv6 optional header can be several hundreds of bytes long.
14) Any node can be a network management agent.
15) A certificate signs the private key of the organisation described.
16) Public Key security algorithms are complex and slow.
17) To access data from a MIB the Object Identifier of the leaf node, and an instance identifier, must be provided.
18) Routers never reply to an ICMP packet.
19) Packets smaller than the minimum required MTU never need fragmenting.
20) Anyone can act as a Certificate Authority.
21) The statement in the question is not false, it is true.

The Statements

(Write the number associated with the explanation you select from the list on the previous page in each box provided)

A) If there is no router, and no DHCPv6 server, an IPv6 node cannot operate, as it cannot obtain a prefix to configure its address. ☐

B) To digitally sign a document, file, or network packet, the node's private key is used to encrypt all of the data in the document, file, or packet. ☐

C) The minimum required MTU for IPv6 links should have been made 1500 bytes (octets), rather than 1280 as it was set, as these days no-one installs links with a smaller MTU than that provided by an Ethernet. ☐

D) DHCP is mandatory in all IPv4 networks as there is no other way to make sure every host gets a different IP address. ☐

E) The only way to obtain, or create, a digital certificate is to prove identity and pay much money to one of a few well known Certificate Authorities. ☐

F) Using Path MTU Discovery (PMTUD) is mandatory for all IPv6 communications as IPv6 routers are not permitted to fragment packets. ☐

G) Since every octet transferred over a TCP connection is given its own sequence number, and the sequence number is a 32 bit value, the maximum possible TCP transfer size is approximately 4 gigabytes (4 gigabytes is approximately 2^{32} bytes.) ☐

H) The **Integrated Services** model for providing Quality of Service is rapidly being deployed in the Internet. ☐

I) When an IPv6 node auto-configures its IPv6 address using an EUI-64 derived from its MAC address, and a prefix obtained from its local router, it is guaranteed that no other node can possibly have the same address. ☐

J) A **MIB** (Management Information Base) is a general purpose database implemented in network routers for network management and monitoring data. ☐

Question 2. *(6 marks)*

Which of the following three (3) statements about Internet routing are **TRUE** and which are **FALSE**?

(Write **T** *or* **F** *in each box provided)*

A) Routers running a Bellman-Ford (or Distance Vector) routing protocol know to which router(s) every network in the system is connected. ☐

Why?

_____

_____

_____

B) In Link-State routing (Dijkstra's Shortest Path First algorithm, or similar) every router must obtain a LSA (Link State Advertisement) from every other router. ☐

Why?

_____

_____

_____

C) For Path Vector routing (such as BGP), all Autonomous Systems (AS) must implement a common **internal gateway protocol** *(local routing protocol)* so that consistent forwarding can be achieved. ☐

Why?

_____

_____

_____

Question 3. *(10 marks)*

Which of the following five (5) statements about multicast are correct (or true), and which are incorrect (or false)? For each incorrect (or false) statement, briefly explain why the statement is incorrect.

A) A sender of multicast packets must join the multicast group before sending packets to it.

B) Multicast is not suitable for TCP sessions.

C) Multicast packets are transmitted to a multicast router on the sender's link (Ethernet) to be forwarded to members of the group.

D) A multicast router will always know the identities of all members of multicast groups on links connected to it.

E) Multicast routers can forget which links should be **pruned** for a particular group without harming the correctness of the protocol.

Question 4. *(5 marks)*

For each of the following, is it **True** or **False** that the capability described can be provided using only Quality of Service (QoS) mechanisms?

(Write **T** *or* **F** *in each box provided)*

A) Lower packet delay. ☐

B) Ordered packet delivery. ☐

C) Less variation in the intervals between packet arrival times (that is, less jitter). ☐

D) Higher probability of packet loss. ☐

E) Reliable data transfer. ☐

Question 5. *(30 marks)*

Examine the network in the diagram:

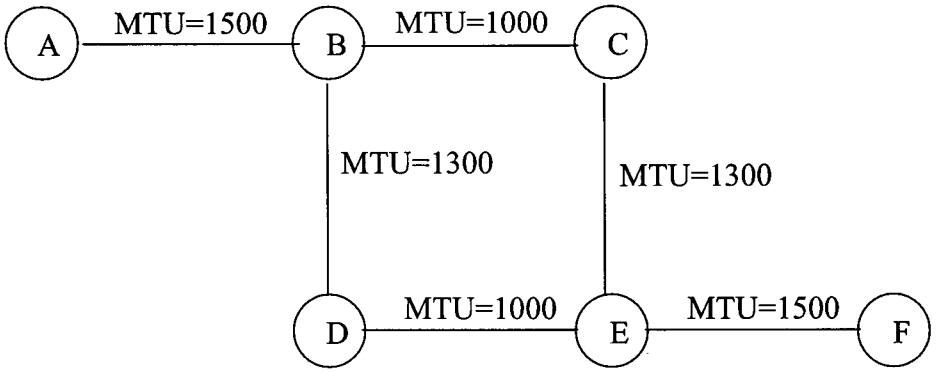

The MTUs of the various links are as indicated in the diagram. Do not assume that any node has knowledge of the MTU of any link to which it is not directly connected.

The path from A to F can be A-B-C-E-F or A-B-D-E-F the routing costs of each path are equal.

Node B is configured to send consecutive packets to the same destination that can be reached via equal cost paths, using each of the possible paths in order, then repeating.

In this example, for this question, that means that when packets arrive at B. for destination F, one packet will be sent to C, the next to D, then to C again, then D, repeating forever. At the time of this question, the next link to be used by B for packets to F is the link to C.

Whenever any node needs to fragment a packet, for the purposes of this question, it always makes as many maximum sized fragments as it possibly can, starting with the smallest offset and working forwards, ending with a possibly small fragment containing the last of the original packet's data. It transmits the fragments it made in order of increasing offset.

Node A desires to send a 4000 byte packet to node F. Node A does not use Path MTU Discovery (PMTUD), and so sends all packets with the DF bit set to 0.

The relevant fields of the IP header of the packet node A wants to send are shown in the figure:

Vers=4 HdrLen=5	IP Version 4, Header Length 5 (words)
Pkt_Len= 4000	Total packet length (including header) 4000 octets
ID= 12120	Packet Identifier value 12120
M=0 D=0 R=0	More Flag reset, Don't Frag Flag reset, Reserved Flag reset
Offset = 0	Fragmentation offset (units of 8 octets) 0

Other fields of the packet header are not relevant to this question.

The source address will be **A**, and the destination address will be **F**.

The protocol might be anything (TCP, UDP, ICMP, or any other meaningful value.)

The TOS field is probably not used, and set to zero, but any value is possible.

When the packet is created at node A, the TTL might be set to any value bigger than 4. It will then be altered as packet passes through the network. Its value at any point is not relevant to this question.

The header checksum will be calculated as required by the other header fields, adjusted as necessary as the TTL alters, and you can assume is always correct. Its actual value is irrelevant to this question.

In the following diagram, complete the missing information (that is, **packet length**, **M**, **D** and **R** flags, and **offset**) to show the packet or packets that arrive at B in the order that they should arrive.

Vers=4 HdrLen=5	Vers=4 HdrLen=5	Vers=4 HdrLen=5
Pkt_Len=	Pkt_Len=	Pkt_Len=
ID= 12120	ID= 12120	ID= 12120
M= D= R=	M= D= R=	M= D= R=
Offset =	Offset =	Offset =
Vers=4 HdrLen=5	Vers=4 HdrLen=5	Vers=4 HdrLen=5
Pkt_Len=	Pkt_Len=	Pkt_Len=
ID= 12120	ID= 12120	ID= 12120
M= D= R=	M= D= R=	M= D= R=
Offset =	Offset =	Offset =

Note that there is no requirement (or expectation) that all fragment headers provided will be needed. You may leave some blank, or use them to correct errors you make – simply cross out any header that is not to be considered part of your answer.

Assume no packets are lost anywhere, and that this one original packet from A is the only packet in the network at the time (there is no other network traffic at all.)

The first packet to arrive at B will be transmitted towards F via node C. In the following diagram, show the packet or packets that will arrive at C as a result of the first arriving packet at B.

Vers=4 HdrLen=5
Pkt_Len=
ID= 12120
M= D= R=
Offset =

Vers=4 HdrLen=5
Pkt_Len=
ID= 12120
M= D= R=
Offset =

Vers=4 HdrLen=5
Pkt_Len=
ID= 12120
M= D= R=
Offset =

Because on the path from C to F, the MTU only increases, no further fragmentation will be needed, so the fragments that arrive at C should also arrive at F, unaltered (excluding the TTL and checksum of course, which are not relevant to this question, and not included in our header pieces).

The third packet to arrive at B (if such a packet exists) will also be forwarded via C. Show what arrives at C from this packet (or nothing if there is no third packet).

Vers=4 HdrLen=5
Pkt_Len=
ID= 12120
M= D= R=
Offset =

Vers=4 HdrLen=5
Pkt_Len=
ID= 12120
M= D= R=
Offset =

Vers=4 HdrLen=5
Pkt_Len=
ID= 12120
M= D= R=
Offset =

The second packet to arrive at B, will be transmitted towards F via node D. In the following diagram, show the packet or packets that will arrive at D.

Vers=4 HdrLen=5	Vers=4 HdrLen=5	Vers=4 HdrLen=5
Pkt_Len=	Pkt_Len=	Pkt_Len=
ID= 12120	ID= 12120	ID= 12120
M= D= R=	M= D= R=	M= D= R=
Offset =	Offset =	Offset =

Then for all of those packets (the ones that arrive from B at D shown just above), show the packet, or packets, that will later arrive at node E (and then unchanged at node F):

Vers=4 HdrLen=5	Vers=4 HdrLen=5	Vers=4 HdrLen=5
Pkt_Len=	Pkt_Len=	Pkt_Len=
ID= 12120	ID= 12120	ID= 12120
M= D= R=	M= D= R=	M= D= R=
Offset =	Offset =	Offset =

Vers=4 HdrLen=5	Vers=4 HdrLen=5	Vers=4 HdrLen=5
Pkt_Len=	Pkt_Len=	Pkt_Len=
ID= 12120	ID= 12120	ID= 12120
M= D= R=	M= D= R=	M= D= R=
Offset =	Offset =	Offset =

The fourth packet to arrive at B, if there is any such packet, will be transmitted towards F via node D. If there are any such packets, in the following diagram, show the packet or packets that will arrive at D.

Vers=4 HdrLen=5	Vers=4 HdrLen=5	Vers=4 HdrLen=5
Pkt_Len=	Pkt_Len=	Pkt_Len=
ID= 12120	ID= 12120	ID= 12120
M= D= R=	M= D= R=	M= D= R=
Offset =	Offset =	Offset =

Then for all of those packets (that you have shown above), show the packet, or packets, that will later arrive at node E (and then unchanged at node F):

Vers=4 HdrLen=5	Vers=4 HdrLen=5	Vers=4 HdrLen=5
Pkt_Len=	Pkt_Len=	Pkt_Len=
ID= 12120	ID= 12120	ID= 12120
M= D= R=	M= D= R=	M= D= R=
Offset =	Offset =	Offset =

Vers=4 HdrLen=5	Vers=4 HdrLen=5	Vers=4 HdrLen=5
Pkt_Len=	Pkt_Len=	Pkt_Len=
ID= 12120	ID= 12120	ID= 12120
M= D= R=	M= D= R=	M= D= R=
Offset =	Offset =	Offset =

Question 6. *(10 marks)*

Explain how a playout buffer can help reduce the effects of jitter upon a real time application.

Indicate what is the cost (in terms of application or network performance, or apparent performance) of using a playout buffer.

Is a playout buffer ever useful for applications using TCP? Why, or why not?

Question 7. *(4 marks)*

What does the UDP protocol add to the protocol stack that is not already provided by the IP protocol, over which UDP runs?

Question 8. *(10 marks)*

On one TCP connection, between port 1123 on host A, and port 8097 on host B, host A sends to B

SYN=1, ACK=0, RST=0, Seq=12345, Ack=X, Src port=1123, Dst port=8097

at about the same time as host B sends to A

SYN=1, ACK=0, RST=0, Seq=9012, Ack=X, Src port=8097, Dst port=1123

When B receives the first of those, and A receives the second, each will send a packet to the other in reply.

Those packets will be *(fill in the values in the boxes)*:

From A to B, Src port=1123, Dst port=8097,

SYN= ☐

ACK= ☐

RST= ☐

Seq= ☐

Ack= ☐

From B to A, Src port=8097, Dst port=1123,

SYN= ☐

ACK= ☐

RST= ☐

Seq= ☐

Ack= ☐

Question 9. *(15 marks)*

A web (www) server has a certificate signed by a certificate authority (**CA**) that is known and trusted by all web browsers (accept that as a fact).

A) A web web browser connects to that server attempting to make a secure connection.

The server will send its certificate to the browser.

Upon receiving the certificate the web browser will check it.

How?

_____

_____

_____

B) After verifying the certificate (and succeeding) the browser will now know:

_____

and

_____

C) The server will also have sent a list of symmetric key algorithms to the browser.

Why?

_____

_____

D) With this list the browser can select one algorithm and generate a secret key for the selected algorithm.

It then sends that key to the server.

How is that done securely?

_____

_____