THE BCS PROFESSIONAL EXAMINATIONS
BCS Level 5 Diploma in IT

October 2007

EXAMINERS' REPORT

COMPUTER NETWORKS

General Comments

There was a wide range of quality in the answers provided by candidates. There were some excellent answers to all questions, but there were also some quite poor answers. There was some evidence that students had missed key words in questions and thus had given a general answer to the topic rather than addressing the precise scenario as given in the question.

Question 1

a) A total of 10 computers (C1 to C10) are connected to a LAN switching hub to form a LAN. Assuming that the LAN switching hub has just been switched on and C1 transmits a LAN frame destined to C5 and then C5 replies with a frame destined to C1, determine which ports these frames are transmitted over and how the switching hub is able to learn the port numbers on which both C1 and C5 are located.

(12 marks)

b) What is meant by the term 'Virtual LAN' ?

(4 marks)

c) How could the LAN switching hub described in part (a) be configured to support two virtual LANs; one with a membership of C1 to C5 and the other with a membership of C6 to C10?

(9 marks)

Answer Pointers

a) When the switching hub is first switched on then it does not know the address of any computer connected to it or the ports on which they reside. So when C1 transmits a frame destined to C5, the switching hub must transmit this frame on ALL ports, except the one from which it was received. However, the switching hub is able to read the source address of this frame and therefore learns the port number on which C1 is located. When C5 receives the frame, it will respond with a frame of its own. This frame will be destined to C1, however, the switching hub now knows the port on which C1 is located and therefore the frame is only transmitted on this port. In receiving this frame from C5, the switching hub has now learned the port number on which C5 is located. Hence, any further frames between C1 and C5 will be constrained to their respective ports.

b) A virtual LAN is a logical, rather than a physical, separation of nodes. In this way, it appears as though a group of nodes within the same vLAN are located on their own private LAN and physically separated from any other nodes.

c) There are two ways that this can be achieved:

Port-based vLANs

The vLAN routing table will be configured such that the port numbers on which C1-C5 are connected are defined as vLAN 1 and the port numbers on which C6-C10 are connected are defined as vLAN2.

MAC address based vLANs

The vLAN routing table will be configured with the MAC addresses of C1-C5 being defined as vLAN1 and the MAC addresses of C6-C10 being defined as vLAN 2.

Examiner's Guidance Notes

This question was tackled by only about 30% of the candidates. A small number of candidates submitted perfect, or near perfect answers. Unfortunately, this was matched by a number of very poor answers. For part a), of the question, many candidates ignored the fact that the switching hub has just been turned on and was thus devoid of knowledge. Switching hubs take their decision based purely on information associated with MAC addresses, however, many candidates gave answers that indicated that the students thought that the switching hubs behaved more like a router, which is of course false. Many candidates referred to the use of the ARP protocol which was irrelevant. For part b), some students gave good answers, whereas others confused the concept of a VLAN with a VPN. For part c), several answers did not fully consider the scenario.

Question 2

a) Explain what is meant by the term 'protocol layering'.

(6 marks)

b) By means of a protocol layer diagram based on the ISO seven layer reference model, show how data is transmitted from one computer to another over a network and clearly indicate on the diagram what is meant by a 'peer to peer protocol.

(14 marks)

c) What functions are performed by the Transport layer within the ISO seven layer model?

(5 marks)

Answer Pointers

a) If one considers a complete communication system from user application through to the physical media over which the information is transmitted, you are faced with a very complex problem involving many different protocol functions. As a consequence, it is now generally accepted that overall protocol functionality is broken down into a series of simpler protocols or layers. In this way, each layer provides a finite functionality that builds upon the services offered by the layer beneath and offers its service to the layer above.

b)

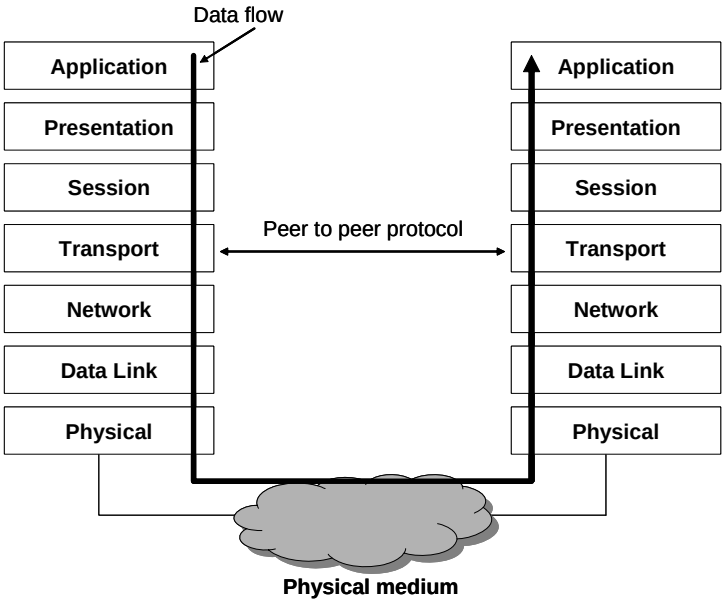

c) Transport Layer

- responsible for end to end communications and present only in the end-stations, not the network
- responsible for accepting the quality of service from the network layer and translating this into the quality of service required by the application
- where the network layer provides limited quality of service, as in the case of IP, then the transport layer must provide these functions as in the case of TCP
- provides application level multiplexing functions

Examiner's Guidance Notes

This question was tackled by over 90% of the candidates. No candidates submitted perfect answers, but having said that, the quality of many answers was very good, a high proportion being deemed of pass quality. Part a) of this question was perhaps the most poorly answered. Candidates should look at our answer outline and marking scheme. Part b) was in general answered well. Many candidates are obviously well aware of the concepts represented by the ISO OSI model. The answers for part c) did indicate some misunderstanding. Some candidates did not indicate that the transport layer operates in an end-to-end manner (rather than hop by hop).

Question 3

a) What functions are performed by a router?

(8 marks)

b) Explain the basic principles behind the operation of the Open Shortest Path (OSPF) routing protocol.

(12 marks)

c) How are routers able to prioritise traffic and provide a different quality of service to different traffic flows?

(5 marks)

Answer Pointers

a) A router is responsible for three key functions. These are traffic routing where the addressing information contained within a received packet/datagram is examined to determine how it should be forwarded. This requires a router to maintain a routing table that allows it to determine the next hop (outgoing port) on which to send a received packet/datagram. This table is constructed through the use of a routing protocol that operates between routers and from which routes through the network can be determined and classified.

The second function is that of quality of service control through packet classification/prioritisation. In today's networks it is increasingly important for traffic to be differentiated and a different quality of service offered to higher priority traffic. Routers use fields within received packets/datagrams – such as addresses/port numbers/protocol fields – to classify each packet or traffic flow. These are then assigned to different output queues from which traffic is transmitted based upon a scheduling algorithm. High priority queues would expect to use more of the available bandwidth on the outgoing port.

The third and final function is providing connectivity of different network technologies. A router operates at layer 3 of the ISO model and as such, is independent of the underlying hardware of the network. In this way a router can connect, for example, a Token Ring LAN to a CSMA/CD LAN, or a CSMA/CD LAN to a wide area network, such as ISDN.

b) OSPF is a link state database routing algorithm.

Each router determines those routers that are adjacent to it and calculates a cost for each such link. This cost can reflect the delay currently experienced on that link, the bandwidth available, the technology and event the financial cost of using the link. The knowledge of these adjacent routers and their associated costs are stored within a link state database.

Using a protocol, each router communicates with those adjacent to it and sends a copy of its link state database. To prevent excessive traffic being generated, it is normal for one router to communicate on behalf of a group of others. So, for each subnetwork, one router will be designated and become responsible for passing on a copy of each router's link state database to the next subnetwork.

As the protocol completes a cycle, so every router will have received a copy of the link state database from every other router. From these it is possible to construct a complete map of the overall network topology and all of the associated costs. A routing algorithm – Dijkstrar's algorithm – is then able to compute the shortest path between any given two points. It is these shortest paths that routers use in order to send traffic through the network.

c) This is achieved by firstly identifying each traffic flow when it arrives at the router. A traffic classifier is able to examine the addressing information, or higher layer protocol information, such as port numbers, to determine the traffic flow and hence, derived its quality of service requirements. Secondly, having classified the traffic the router is able to prioritise traffic flows by placing them into different output queues and then scheduling transmission from the queues in some defined order. For example, high priority queue data is always sent before lower priority queue data

Examiner's guidance notes

This question was tackled by over 90% of candidates. The quality of the answers presented varied widely. A small number of candidates submitted very good answers, but there were a large number of poor quality answers. Many of the answers to part a) neglected important issues. Only a small number of candidates mentioned the Quality of Service issue and only a small number mentioned the ability of a router to connect networks that used different technologies. It was also clear that many candidates falsely believed that routing tables were calculated every time a new packet arrived. While packet forwarding decisions are taken for every packet, the calculation of the routing tables takes place MUCH less often. For part b), many answers confused the behaviour of link state protocols with the behaviour of distance vector protocols. Part c) was in general answered poorly. Many candidates ignored the essential role played by the existence of multiple queues within a router that offer QoS. Traffic classification is important. Based on the way traffic is classified, a router will then cause it to enter the appropriate queue. To provide the appropriate QoS, a router then has to use an appropriate algorithm to decide how to take traffic from the queues and transmit it onwards.

Question 4

a) What does it mean to say that IP offers a best-effort unreliable delivery service? Include in your answer how datagrams may arrive out of order.

(6 marks)

b) Explain how TCP builds a reliable delivery service, using an unreliable service (IP).

(12 marks)

c) Give examples of the type of applications that typically use UDP in preference to TCP and explain why these applications work better using an unreliable protocol for the delivery mechanism.

(7 marks)

Answer Pointers

a) On an IP network, data is encapsulated in datagrams. These datagrams each contain the address of the receiver. They are routed independently so may arrive out of order. If they are lost on the way, there is no check that they have been lost.

b) There are two aspects to the answer:

i) setting up a virtual connection between the two TCP endpoints – 3 way handshake etc.

ii) numbering of TCP segments so that the receiving station can know if segments are arriving out of order (and re-order them, before passing them to the application) or if segments have got lost.

c) Real time applications such as video-conferencing work better using UDP because it is better for some datagrams to be lost, than to lose time in retransmission (which would impair the quality of the transmission more than the loss of some datagrams)

Examiner's Guidance Notes

Generally this question was done well. A substantial portion of the students did understand why datagrams might arrive out of order (due to each one being routed individually). And most had a fair idea that IP was an unreliable service and could give a reasonable explanation of why that was. But there was a significant proportion of students who thought that TCP set up 'virtual paths' through the network – i.e. they failed to understand that TCP works over IP and has to use the unreliable service that IP offers. The proper term is virtual connection – and that does not mean that the datagrams all follow the same route. As before they are routed independently. This confusion lost marks for a lot of students. Some students had the idea that UDP is used primarily for request/ response protocols – meaning, presumably, applications that access a database, or perhaps the Web. Generally these applications use TCP. On the other hand many students did understand that UDP was best adapted to applications like video conferencing where the loss of some packets is not critical and speed and continuity of delivery is the main problem.

Question 5

a) Illustrate the structure of a cell in an ATM (Asynchonous Transfer Mode) network. You may use a diagram to illustrate your answer.

(5 marks)

b) Explain how ATM combines the advantages of circuit switching and packet switching.

(12 marks)

c) In the context of ATM explain the difference between a Permanent Virtual Circuit and a Switched Virtual Circuit and the advantages and disadvantages of each.

(8 marks)

Answer Pointers

a) Either the answer can be given in words only, or else a diagram can be used to illustrate the answer as follows:

<table>
<tr><td colspan="3">VPI</td></tr>
<tr><td>VPI</td><td colspan="2">VCI</td></tr>
<tr><td colspan="3">VCI</td></tr>
<tr><td>VCI</td><td>PT</td><td>CLP</td></tr>
<tr><td colspan="3">HEC</td></tr>
</table>

b) ATM offers a connection-oriented service (although it can also offer connectionless service if required). In the case of the connection-oriented service a virtual circuit is set up before transmission starts. The advantages are guaranteed capacity and constant (i.e. predictable) transmission delay. The advantages of packet switching are flexibility and efficiency for intermittent traffic. ATM can offer this as well because it is asynchronous.

c) A permanent virtual circuit is similar to a leased line and endures until it is taken down manually. The disadvantage is that it has to be set up manually in the first place and the connection that is made is static. The advantage is that the availability of the connection is guaranteed and there are no call setup procedures that need to take place before data is transmitted. A switched virtual circuit is like a telephone call, and remains in use only while data is being transferred. This gives flexibility. The call setup procedures can be done automatically by a suitable networking device and do not require manual intervention. But extra time and overhead are taken up in this setting up of the circuit.

Examiner's Guidance Notes

Most students gave the right diagram or textual description of an ATM cell.
Quite a lot of students reviewed the whole question of packet switching and circuit switching a review the advantages and disadvantages of both. Marks were given for doing this, since this approach responded to the sense of the question. However, very few students were able to apply the argument to ATM because the point usually stressed about ATM is the fact that it uses circuit switching. On that basis alone it is, of course, difficult to see how ATM can combine the advantages of circuit switching and packet switching. A great many students were able to explain the difference between a Permanent Virtual Circuit and a Switched Virtual Circuit and a fair proportion of them could explain quite exactly the advantages and disadvantages of each.

Question 6

a) If a physical network is based on (a) optical fibre, (b) shielded copper wire or (c) uses wireless transmission, explain which is the least likely to produce errors and which is most likely to produce errors and why?

(5 marks)

b) Explain the difference between error detection and error correction. Which approach requires more information to be sent, in addition to the original data?

(4 marks)

c) Explain an approach to burst error correction.

(8 marks)

d) A 7-bit ASCII character is encoded, using the Hamming code and is transmitted to a receiver. If the bit pattern 11000010101 is received:

Show how the receiving station checks for an error. Determine if there is any error, and show how the error can be corrected.
What was the original ASCII character?
Determine the code efficiency of the encoder.

(8 marks)

Answer Pointers

a) A wireless transmission is most likely to incur errors, because it is subject to many potential types of interference. Shielded copper wire is protected against interference by the shielding, but not perfectly. It can still be subject to electrical interference. Optical fibre is not affected by electrical interference, and so is the most reliable and can be used in environments where there is a great deal of interference.

b) Error detection is necessary is both cases. Error detection only detects errors and does not directly try to reconstruct the original information (retransmission may to used to get the information again. Error detection requires less information to be sent. Error correction carries more information and the system uses that information to reconstruct the original data from the received data.

c) The use of Hamming codes can be extended to deal with burst errors by working out a series of code words then sending all the first bits from each codeword, then all the second bits and so on. Thus, if a single burst error occurs, only one bit from each codeword will (hopefully) be corrupted, and the codewords can be dealt with as usual.

d) Working is as follows:

11 = 1011

10 = 1010

5 = 0101

3 = 0011

1 = 0001

TOTAL IS: 0110 indicating that there is an error in the bit 6.

The corrected codeword is 11000110101

The ASCII character is 1100111

The code efficiency of the encoder is 7/11

Examiner's Guidance Notes

There was quite a wide variation in terms of what students thought would be most and least likely to suffer from errors from (a) optical fibre, (b) shielded copper wire and (c) wireless transmission. Amongst those who did answer correctly, there was a tendency just to single out the two 'extreme' cases – the best and worst – and to forget to say anything at all about shielded copper wire. Since they had to explain why one or another system was more or less likely to produce errors, the question required comment on all 3 systems. But the marks for this part were quite good, in any case.

Very few students were able to suggest an approach to burst error correction. If they explained clearly what a burst error actually was then a mark was given. If they gave a proposal for how a burst error could be detected then a mark was given. If they suggested the Hamming code should be used then marks were given, since that is basically the right approach. Virtually no students could explain how the Hamming code could be adapted to deal with burst errors.

A fair proportion of the students both understood how the Hamming code works and were able to do the necessary work of calculation to determine what was the error in the received transmission – and therefore to correct it. If errors were made in the calculations carry through marking was done to ensure that marks were not lost more than once for the same error.