

**THE BCS PROFESSIONAL EXAMINATION
Professional Graduate Diploma**

April 2007

EXAMINERS' REPORT

Network Information Systems

General

The number of candidates selecting this module continue to rise and those that selected the module this year appear to have taken the advice offered in the examiners report for 2006. Examiners were pleased to note that many candidates were very well prepared and examiners were able to award maximum or near maximum marks to all but one question. The overall pass rate improved from 66% to almost 80% this year.

Question 1

Please note: All sections of this question should be answered in the context of a local area network running TCP/IP over Ethernet.

- a) Discuss the functionality of a bridge, identifying at what layer of the ISO Open System Interconnection (OSI) network model it operates, what filtering can be performed, how it handles broadcasts. **(8 marks)**
- b) Discuss the functionality of a router, identifying at what layer of the OSI network model it operates, what filtering can be performed, how it handles broadcasts. **(8 marks)**
- c) Explain the need for routing protocols and describe the main differences between RIP and OSPF in terms of protocol type, ease of configuration and convergence speed. **(9 marks)**

Answer pointers

- a) A bridge connects two or more network segments at the data link layer of the OSI model. Traffic is forwarded through it from one network to another. It is capable of filtering using MAC (hardware) addresses. Broadcasts are retransmitted on all connected segments. (2 marks each)
- b) A router connects two or more network segments at the network layer of the OSI model. Traffic is forwarded through it from one network to another. It is capable of filtering using IP addresses. Broadcasts are not retransmitted on connected segments. (2 marks each)
- c) Routing protocols allow changes in network topology to be multicast to other routers to ensure that all routers have valid routes. If a network is connected or disconnected the connecting router informs others in its multicast group. (3 marks) RIP is a popular distance-vector protocol which is easy to configure but relatively slow to converge. (2 marks) OSPF is a link-state protocol which is quite complex to configure but converges quickly. (2 marks).

Examiner's comments

Most candidates attempted this question. The first two parts were answered well enabling over 80% to gain a pass. The third part on routing protocols was answered badly as the question had either not been read properly or had been misunderstood.

Most answers correctly described the function of bridges which operate at the datalink layer which can filter on MAC addresses and allow broadcasts.

Most answers correctly described the function of routers which operate at the datalink layer which can filter on IP addresses and do not allow broadcasts.

Routing protocols were misunderstood. They do not deliver datagrams. They are used to enable routers to communicate changes in network topology so that their routing tables can be updated dynamically. No answers mentioned that routing protocols are multicast.

Question 2

- a) There are seven protocol layers in the ISO Open System Interconnection (OSI) protocol. Explain the function of each layer and give examples of actual protocols used in practice. **(14 marks)**
- b) Inter-network protocols are overlaid on underlying networks as shown in **Figure 1** on the next page. The network interface layer accepts inter-network packets and converts them into packets suitable for transmission by the transport layer of a specific underlying network. The underlying network consists of the transport, network, data link and physical layers of all the real networks that constitute the inter-network.

Discuss the following relevant issues:

- i) Packet assembly **(3 marks)**
ii) Virtual circuit packet delivery **(4 marks)**
iii) Datagram packet delivery **(4 marks)**

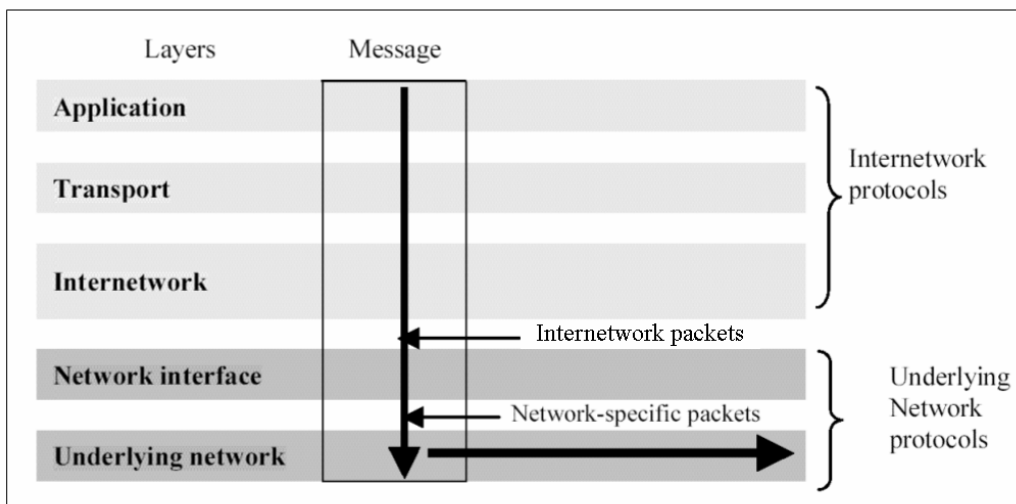


Figure 1: Inter-network Layers

Answer Pointers

a) Applications: protocols that are designed to meet the communications request of specific applications, e.g. FTP, Telnet, SMTP, X400 and X500.

Transmit data in a network representation independent of individual Computer rep's; encryption also done here, e.g. XDR, ANS.1

Session: for setting up process communications and error recovery, not needed for connectionless communications.

Transport: lowest message level, connection-oriented or connectionless, e.g. TCP, UDP

Network: data packet between computers, e.g. X25, IP.

Data Link: does the error free transmission of packets for directly connected computers, in WANs between pairs of PSE & PSEs and hosts, e.g. HDLC,, Ethernet: CSMA/CD

Physical: for circuit/HW, e.g. X21, Ethernet: base-band & signalling. (14 marks)

b) Packet Assembly: In the internet suit of protocols, the IP protocol is a 'network layer' protocol; the MTU packet for IP packets is unusually large- 64 Kbytes including packet leader & the data field. This means UDP transport layer data-grams seldom need to be subdivided before being put in IP packets, although this may need to be sub-divided by the network interface layer, e.g. to fit into Ethernet packets. (3 marks)

Virtual circuit packet delivery: Each network layer packet contains a virtual circuit number; it needs to contain the destination address for identification. Packets are routed at each switch by ref to the circuit no.. Packets are also checked & acknowledged at each step along the route. On arrival at destination, they are passed to the transport layer in a format that includes a channel identifier in a connection-oriented service and sender's address in a connectionless service. (4 marks)

Datagram packet delivery: Each network level packet contains the network addresses of the source and destination PCs. At the source, and at each switch along the route to the destination, the destination address is used to find the next step along the route, using pre-defined routing tables held in each switch. Tables are modified subject to network faults or loading. (4 marks)

Examiner's comments

This question proved to be the most popular choice with nearly 94% of the candidates attempting the question. The pass rate was also very encouraging, with 109 candidates passing it, just over 83% of those who attempted it. The average mark for this question was around 57%, or about 14 out of 25. The answers varied widely in quality, detail and length, with quite a number of candidates demonstrating knowledge based on first hand practical experience with computer network layers.

Question 3

- a) Explain the structure and contents of an X.509 digital certificate. **(8 marks)**
- b) Explain the main function of a certificate and discuss the role of a certificate authority in achieving this functionality. **(8 marks)**
- c) Using an example explain clearly how a certificate is used to authenticate its owner when using Secure Sockets Layer (SSL). **(9 marks)**

Answer pointers

- a) Certificate: Version, Serial Number, Signature Algorithm,
Issuer: Certificate authority,
Validity dates,
Subject: Owner identity
Subject Public Key Info:
Public Key Algorithm: rsaEncryption
RSA Public Key: (1024 bit)
Signature Algorithm: md5WithRSAEncryption
Digital signature. **(8 marks, 1 per major section)**
- b) The main function of a certificate is to allow a public key to be obtained with some assurance that it is genuine. A certificate authority is a trusted third party who's own certificate is readily available. The CA creates a digital signature of other certificates using its private key. The public key in the CA certificate can be used to verify the digital signature.
(2 marks each = 8 marks).
- c) A certificate is sent to the server when the SSL connection is established. The certificate is validated using its digital signature and the public key from its CA certificate. A symmetric session key is generated by the server. The session key is encrypted using the public key from the certificate. This is sent to the client and decrypted with the private key of the certificate owner. The session key is used to exchange data. **(1 mark each = 7 marks)**

Examiner's comments

Around 75% of candidates attempted this question. The first two parts were answered well. The third part on SSL was often answered badly due to a lack of understanding of SSL.

The most important point is that a certificate encapsulates a public key and had a third party CA digital signature to verify that the public key is genuine. Few stated that the CA public key, which is required to validate the digital signature, is itself stored in a certificate which is contained within web browsers. There was some confusion about the role of a CA. Clients do not request server certificates from a CA web site. The server obtains a certificate from a CA which is then sent to clients as a means of delivering a public key.

Part three was about how SSL works. Most answers correctly stated that the client obtains and validates the server's public key using a certificate. Few stated that the certificate is used to encrypt a session key negotiation with the server. The session key is used for subsequent data transfers.

Question 4.

- a) Distributed processes often need to coordinate their activities. For example, if a collection of processes share a single or a collection of resources managed by a server, then often mutual exclusion is required to prevent interference and ensure consistency when accessing resources.

Explain the requirements for mutual exclusion in terms of safety, liveness and ordering.

(9 marks)

- b) The simplest way to achieve distributed mutual exclusion is to employ a server that grants permission to enter a critical section, commonly referred to as 'the central server algorithm'. **Figure 2** on the next page depicts such a server managing a mutual exclusion token for a set of processes.

Explain the structure and operation of this algorithm.

(16 marks)




Figure 2

Answer pointers

- a). Safety: only one process inside a critical section (CS) at a time, use of semaphores in Wait & Signal operations. (3 marks)

Liveness: a process requesting entry to CS is eventually granted access when CS is no longer occupied. (3 marks)

Ordering: access to CS is granted to processes in the time order of the request. (3 marks)

b).

-A process sends out a request to the server & awaits a reply, the reply is a 'token' permitting entry to the CS.

-If no other process has the token the server replies immediately granting the token, otherwise server doesn't reply but queues the request.

-On exiting the CS a message is sent to the server returning the token.

-If the queue is not empty, then server chooses the oldest process in the queue, removes it & replies to it, the chosen process then hold the token.

-In the figure, P2's request has been appended to the queue which already has P4's request.

-P3 exits CS and the server removes P4's entry and grants permission to P4 by replying to it.

-P1 has no need to enter CS.

(16 marks)

Examiner's comments

Only about half of candidates attempted this question, perhaps put off by the conceptual/theoretical bias of the question, which may give the impression of remoteness from the practical hands-on issues encountered in the daily management and operation of Network Information Systems. Never the less, the pass rate was very encouraging at 75% with the highest mark being 22 out of 25. Some of the high scoring candidates clearly enjoyed the challenge posed by the question which allowed them the opportunity to draw on both their experience and knowledge of the subject.

Question 5

- a) Explain why there is a need for server side dynamic content generation on Web servers. **(4 marks)**
- b) For each section below, explain the details of the technology and give examples of appropriate and inappropriate uses of the technology.
- i) Common Gateway Interface (CGI) **(7 marks)**
 - ii) Server parsed HTML **(7 marks)**
 - iii) Java Servlets and Java Server Pages (JSP) **(7 marks)**

Answer pointers.

a) HTTP is stateless so dynamic content is needed to create and maintain sessions. There is a need to create Web pages from information stored on the server such as in a database. (2 marks each = 4 marks)

b) CGI describes the interface between a Web server and an external program. It allows access to HTTP request headers and information from forms or encoded in the URL. It allows HTTP response headers to be set and content to be returned. Languages C/C++, Perl, Python, shell. Advantages: can be written quickly using scripting languages. Disadvantages: performance overhead and security risks of running external program. Need to print all HTML by the program. (1 mark each = 7 marks)

c) Server parsed HTML uses special tags embedded in normal HTML documents. The files are processed by a program or a Web server module which process the special tags and either remove them or replace them with generated content to produce a new document. Languages: SHTML, PHP, ASP, TAL. Advantages: HTML is written as normal and not written by a program. Disadvantages: can't decouple HTML and program. Overhead and administration of pre-processing documents. (1 mark each = 7 marks)

d) A Java Servlet is a Java object which runs continuously. A method is called on the object each time an HTTP request is directed to the servlet. The servlet generates content in a similar way to CGI. A JSP is a server parsed HTML document which gets converted to a servlet on the first access. Servlets and JSPs require a Web application server such as Tomcat. Advantages: use of Java and its APIs. Disadvantage: need for application server and the URL recognition to direct requests. Overhead of a Java JVM. (1 mark each = 7 marks)

Examiner's comments

Very few candidates attempted this question and very few reached a pass standard. Many candidates appeared to guess their answers rather than demonstrate knowledge.

The important point was that dynamic content uses programs on the server to generate HTML and other web content.

CGI is an interface between a web server and an external program which generates content.

Server parsed HTML is HTML with special tags embedded which enable additional content to be generated when the document is processed on the server.

Java servlets and JSPs are java programs which run in an application server to which content requests are redirected by a web server.