THE BCS PROFESSIONAL EXAMINATIONS
Professional Graduate Diploma

April 2006

EXAMINERS' REPORT

Network Information Systems

General

Those that have selected the module over the last few years have demonstrated good knowledge and as a consequence a high proportion passed.

However this year seems to have attracted a large number of candidates who lacked even basic knowledge of the syllabus, were poorly prepared and as a consequence the pass rate has dropped. Fortunately a number of candidates were well prepared and able to give some excellent answers and they obtained marks of 84% or higher for each question.

Question 1

1. A process group is a collection of processes that co-operate towards a common goal or that consume one or more common streams of information. Group structures are defined according to the pattern of communication in which the members of a group are involved.

The figure below shows one such group structure, referred to as client-server group.

a) Explain how requests from clients are handled and the subsequent actions of the servers. (**9 marks)**

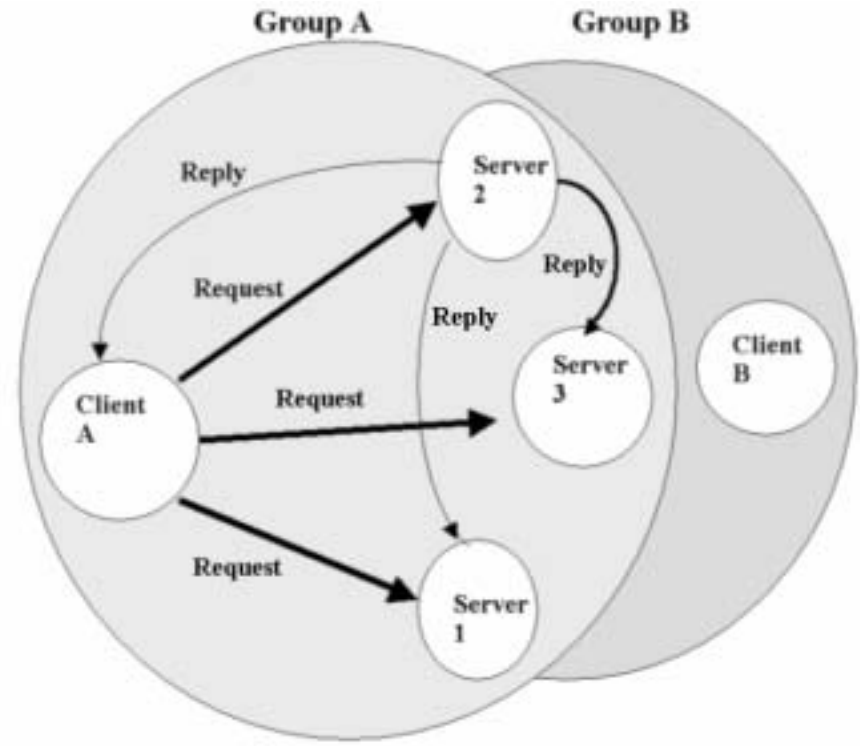

b) Using diagrams where appropriate, explain the principles of operation of the following group structures:

i) Peer group
ii) Server group
iii) Subscription group
iv) Hierarchical group

(16 marks)

Answer Pointers

a. Requests from clients are multicast to all members of the server group, but the server that processes the request from a particular client multicasts the reply to the client-server group, which contains all the server group members plus the client itself; there is only one client-server group for each client so the system has to support many overlapping groups for each service, and the other servers are able to update their state according to the results returned.

b)
i) Peer group: all comm. Is directed from processes within the group to the group,- suitable for processes implementing a multi-user editor (give diagram).
ii) Server group: a client multicasts its request to the group, and all members of the server group receive it, in simple case only one process need reply,- the one known to have the info by all group members, e.g. the oldest, or if ionfo is partitioned among servers only the relevant server process the request.
iii) Subscription group: sent the same info from the source, don't reply to messages but process info in their own app-specific way, e.g. trading room brokers workstations, central computers multicast their gathered data to such a group.
iv) Hierarchical group: avoid management overheads of very large group by dividing into sub-groups with one member joining a root group; multicast to root group whose members multicast to sub-groups,- can be extended to sub-sub-groups; advantages: parallel comm. And smaller vector timestamps for causal ordering; dis-ad: latency increase.

Examiner's Comments

This proved to be a reasonably popular question, attempted by around 60% of candidates who sat this module exam. Candidates seem to have difficulties in fully answering part a) of the question, but generally improved their overall mark by providing fuller and more accurate answers to part b). A few candidates obtained marks above 80% for this question. Answers varied somewhat, in both content and in depth, which represented the wide spectrum of practical experience of candidates. For the candidates who scored high marks, the solutions which they provided indicated first hand knowledge based on direct experience with the material. The majority, however, represented straightforward knowledge gleaned from the recommended text books. There were candidates with clear language difficulties in providing the answers, while others suffered from poor knowledge of the issues involved in the question.

Question 2

2. *a)* Explain the difference between a GET and a POST HTTP request to a Web server and within your explanation clearly identify when POST should be used in preference to GET. **(10 marks)**

b) Explain why HTTP is stateless and describe the mechanisms which can be used to implement the following sessions on a Web site:

i) log in

ii) update profile

iii) log out **(9 marks)**

c) Identify the mechanisms which can be used to restrict access to a Web site to authorised users only. **(6 marks)**

Answer Pointers

a) An HTTP GET request encodes any additional information by appending it to the end of the request URL in the form ?name1=values1&name2=value2. GET should only be used for small volumes of additional information. The HTTP RFC stipulates that GET requests should be read only and not result in any changes on the server. An HTTP POST request encodes any additional information in the HTTP request body (after the headers and blank line). POST requests can send large volumes of additional information. Server content can be updated as a result of a POST operation. (10 marks)

b) Each HTTP request is processed independently by a Web server. In normal operation a Web server cannot distinguish between different users as the HTTP protocol is machine to machine. A login mechanism can be implemented using dynamic Web content such as CGI, PHP, ASP, JSP. A session identifier is created on successful login authentication. This session identifier is past back and forth between browser and server either in a cookie or by encoding it into the URL or using an HTML form hidden control. The dynamic Web content uses the session id to provide state. The session id is deleted on logout. (9 marks)

c) Access to a Web site can be restricted by hostname/ip address, basic or digest HTTP authentications, a more sophisticated authentication mechanism implemented in dynamic Web content. (6 marks).

Examiner's Comments

This was the least popular question and the answers provided were very poor, resulting in a very low pass rate. The information required from candidates is given above.

Many candidates did not have the basic understanding that a GET request encodes the data in the URL and a POST request encodes data in the HTTP body.

Many didn't mention the use of cookies, session variables and URL Encoding, to implement state.

Question 3

3. You are required to produce a draft report identifying the design issues that arise specifically from the distributed nature of systems. However as a full report would cover five sections: Naming, Communication, Software Structure, Workload Allocation and Consistency Maintenance, your report should only consider Workload Allocation.

Your report, using suitable diagrams where necessary, should be divided into the following four parts:

a) the Workstation-server model **(6 marks)**

b) the processor-pool model **(6 marks)**

c) use of idle workstations **(6 marks)**

d) shared-memory multiprocessors **(7 marks)**

Answer Pointers

a) The Workstation-server model
(6 marks)

WSM: Diagram of workstations, PCs, file servers, & other misc. servers for login, print, etc, connected via LAN & to a WAN gateway. Processor power and memory capacity of a work station determines the size of the largest task done by it's user. 'Processor cycles is put near the user' especially for very interactive applications,- very effective in the WSM. But WSM does not optimize use of processing & memory resources, doesn't enable a single user with a large processing/memory- requirement computing task to obtain additional resources.

b) The processor-pool model
(6 marks)

PPM: Diagram of workstations, X terminals & other servers, plus an integrated processor pool, connected via LAN & to WAN gateway. Processor pool: collection of low cost cpu/memory/network interface; each pool has independent network connection. Processors are allocated to processes for their lifetime, resulting in sharing processing at the grain of the whole process. A user with task > 1 process can have more power than one WS can offer, e.g. compiling multi segment C program. X-terminals & Window systems, e.g. X-11 client-server. Ameba. Pool processors allocated dynamically. Others: lan-9 and Clouds.

c) Use of idle workstations,
(6 marks)

IW: as a fluctuating pool of extra computers, especially overnight, e.g. *worm programs* Xerox PARC used for animated graphics in parallel. Sprite op sys for distributed systems,- target workstation is chosen transparently by the system; supports *process migration*, i.e. relocation of an executing program from one machine to another,- a remotely executing program can be migrated safely to its 'home' machine where it can continue with it's execution when a user logs on or starts to use its workstation more heavily.

d) Shared-memory multiprocessors.
(7 marks)

SMM: Diagram of shared memory with several programs each connected to a cache memory & processor. Commonly used as server machines in open distributed systems to give performance for relatively low hardware & software costs. Several independent processors each able to execute a separate program. Number deployed 2-64 determined by engineering costs and

performance overheads, need efficient interfaces to shared memory & large cache memory with each processor. Support is integrated into a distributed op sys using shared mem as hi-speed mech for inter-process comms. & allocating processors to system & user tasks. Can service many client requests in parallel thus reducing bottlenecks in distributed systems. Multiprocessor workstation are also now widely available for demanding processing applications e.g. VR and multimedia.

Examiner's Comments

Attempted by nearly two thirds of the candidates this question proved to be quite popular. Only 13 candidates failed to score a pass mark, representing an 82% pass rate. There were some excellent answers obtaining marks in the 75-85% range. Most candidates gave good answers for parts a) and b) but somewhat lost momentum when getting to the latter parts of the question. The high scoring candidates showed a good knowledge of the subject, in terms of both theoretical principles and the practical implications of the issues. Low scoring candidates had difficulty distinguishing between the four models, for example one or two confused shared memory with idle workstation and workstation-server model.

Question 4

4. *a)* Describe the principles of public key cryptography. Explain what makes public key encryption secure when using the RSA algorithm. **(7 marks)**

b) Describe mechanisms by which public keys can be exchanged which avoid a "Man in the middle Attack". **(8 marks)**

c) Explain how a digital certificate is used as:
i) a server certificate
ii) a client certificate **(10 marks)**

Answer Pointers

a) Public key encryption works on the principle that each party has two keys, a private key which is kept secret, and a public key which is widely published. A message encrypted with the public key can only be decrypted with the private key. Conversely a message encrypted with the private key can only be decrypted with the public key. In RSA each key consists of a modulus which is the product of two large primes and an exponent. The security of the algorithm depends on the fact that it impossible to compute the other key in a reasonable amount of time. It is theoretically possible to decrypt a message which has been encrypted using the public key with the public key, but it requires factorising the modulus which will take a long time. Security does depend on good key management. (7 marks).

b) Public keys can be exchanged physically on removable media. A Key server can be used. Having public keys signed by a trusted third party using a digital signature. The public key can be encapsulated in a digital certificate. (8 marks).

c) A digital certificate is a public key, information about the owner of the key and a digital signature from a trusted third party called a certificate authority. The main purpose of a certificate is that a party can be sent a public key and have a mechanism for verifying that the public key genuinely belongs to the sender. This is done by validating the certificate authority's digital signature using the certificate authority's public key which is itself stored in a certificate. A server certificate is used to send a server's public key to a client usually for the purpose of key exchange for use with establishing SSL and TLS connections. A client certificate can be used to authenticate a client to a server. The client certificate is digitally signed by the server's private key (the server is a certificate authority). If the server receives a correctly signed client certificate and the client has

the corresponding private key to encrypt and decrypt transmissions, then the client is genuine and has authenticated itself. (10 marks).

Examiners Comments

Almost all candidates attempted this question. However the question was answered very badly by many candidates and only half of the candidates reached a pass standard. In general candidates appeared confused as to the actual sequence of coding and decoding.

Marks were lost for not stating that encoding with either a public or a private key can only be decoded with the other key. Many candidates just stated one direction of coding. Very few stated that RSA technique is secure because it is computationally very difficult, but not impossible, to decrypt with one key.

A lot of answers mentioned key sizes of 56 bits when RSA uses at least 1024 bits. None of the solutions mentioned key exchange algorithms such as Diffie Hellman. Very few mentioned exchanging keys in person.

In part c) many answers correctly identified what a certificate was and how it can be used as a server certificate. Few stated that a client certificate validates a client's public key which can then be used to send a challenge for authentication.

Question 5

5. *a)* With reference to Wide Area Networks explain the terms:

i) circuit switching

ii) packet switching

ii) message switching.

Also explain why message switching suffers from several weaknesses related to message length. **(11 marks)**

b) With the aid of an example and a diagram verify the following statement; 'A message gets to its destination faster when sent within packets'. **(8 marks)**

c) Interface Message Processors (IMPs) are used in international Wide Area Networks. Produce a schematic design for such a network, making reference to the length of packets and messages as parameters of data communications between hosts and IMPs. **(6 marks)**

Answer Pointers

a. Circuit switching establishes a physical circuit between nodes. Message switching uses special routing comm. SW to route message (of un-specifiable length) along appropriate paths to destination nodes. Packet switching divides message into standard size packets. Explanation: 3 marks each. Message length problems: link and node-storage monopolized during transmission, until positively acknowledged by destination, and may overflow intermediate node capacity, 2 marks. Total 11 marks.

b. Packets are interleaved or 'pipelined' along links between intermediate nodes to enable nodes to send out packets while still receiving later packets. Two diagrams are expected showing i) a whole message being sent (3 marks) and ii) when the same message is divided into a number of packets (3 marks). Explanation of time saved , 2 marks. Total 8 marks.

c. IMPs need 2 type of interface: to host and to other IMPs to form the network. No.s of hosts per IMP is limited by it's speed and storage, similarly no. of other IMPs it is connected to. Messages from host to IMP are broken down into packets, 'packetizing' is shared between

host & IMP, e.g. host divides into 1 K bytes and IMP breaks further into 1 K bits. Suitable diagram to show a number of hosts & IMPS connected together and message/packet sizes flowing between. 4 marks for explanation and 2 for diagram. Total 6 marks.

Examiner's Comments

Most of those who attempted this question showed reasonable comprehension of the issues involved and associated syllabus material. As usual the candidates' answers represented their personal knowledge based on their individual practical experience and the extent of understanding the underlying theoretical principles. The depth and length of answers were quite varied, some demonstrating an excellent grasp of the subject matter and related surrounding issues. Those who scored high clearly enjoyed attempting this question and answered all parts of the question. Most of those attempted this question managed to explain the 3 types of switching in part a), but nearly half then failed to provide a satisfactory answer to part b) which meant they were unable to obtain marks.

Similarly, most candidates seem to find part c) difficult to answer fully and again were unable to gain marks. However, it is pleasing that some candidates obtained marks above 75% for this question.