THE BCS PROFESSIONAL EXAMINATION

Professional Graduate Diploma

April 2003

EXAMINERS' REPORT

Network Information Systems

The number of candidates selecting this module continue to increase and this year the number taking the examination rose by 30%. Many of whom were well prepared and it is pleasing to report that the number achieving a pass standard was again around 80%.
A number of candidates attempted more than the three questions required, with some attempting all five. Only three answers are taken to give the result and candidates are advised to spend more time on their selected questions.

Many candidates however who failed to reach the pass standard had answered only one question well. Such candidates need to prepare themselves to answer questions from a wider selection of the published syllabus.

An indication is given below of the points expected; however any valid point which was relevant to the question received marks.

Question 1

The scope for open, configurable distributed systems is enhanced if the file service is structured as three components. These three components are, as shown in Figure 1, a flat file service, a directory service and a client module.

a) Explain the operation and the relationships between the various entities shown in the diagram.

(5 marks)

b) Define and discuss the division of responsibilities between the three components In your answer refer to unique file identifiers (UFIDs).

(10 marks)

c) Discuss the design issues involved in providing a range of services to address the requirements of clients with different goals with respect to these three components and the issue of fault tolerance.

(10 marks)

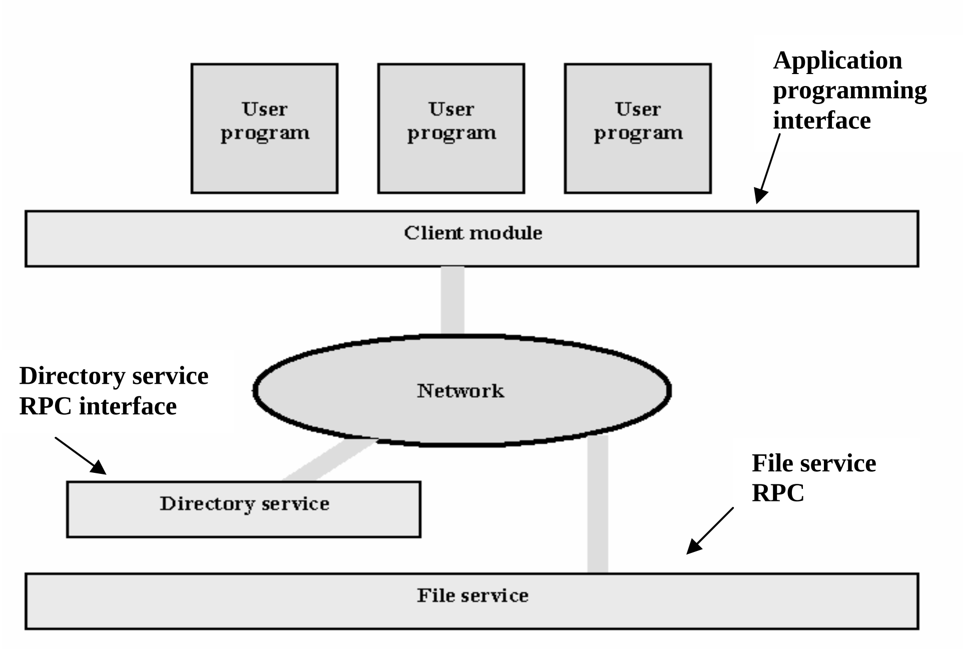

Figure 1: File service components

Answers were somewhat variable and reflected the difficulties some had in relating their prior knowledge and experience to the issues raised in the question. However, this was a reasonably popular question with some good answers.

Answer Pointers

Of the 36% of candidates who attempted this question 84% achieved a pass standard. The average mark was 13.

a)
Flat file service & directory service each export an interface for use by client programs and their RPC interfaces to provide a comprehensive set of operations for access to files. Client module integrates the flat file service and the directory service to provide a single program interface with operations on files similar to those found in conventional file systems. The design is open as different directory services can be used with a single flat file service to support different naming rules and directory structures. Different client modules can be used to implement different program interfaces to simulate the file operations of different operating systems and optimise performance to different workstation and server hardware configurations.

b)
Flat File service (FFS): concerned with implementing operations on files' contents. UFIDs are used to refer to files for flat file service operations. The division between file service & directory service is based on use of UFIDs, these are long integers unique to each file in a distributed system. The Flat file service generates & returns, to the requester, a new UFID when it receives a request to create a file.

Directory service (DS): provides mapping between *text names* for files and their UFIDs. Upon file creation: client module must record the UFID and text name of each file, clients may then get a file's UFID by giving its text name to the directory service. Provides functions to generate & update directories & obtain UFIDs from directories,- it is a client to the flat file service, it's directory files are stored in files of the flat file service.. When a hierarchic file naming scheme is adopted (e.g. Unix) directories will hold references to other directories.

Client module (CM): It's an extension of the user package concept,- a single client module runs in each client computer to integrate and extend the operations of FFS & DS under one application program interface available to user-level programs in client computers. CM also holds information in relation to the network locations of the FFS & DS processors. CM also implements a cache of recently used file blocks at the client to enhance performance.

c)
FFS: designed to offer a simple, general purpose set of operations. File contain data (as items accessible for read/write any portion) & attributes as a single record to hold
i). file length, ii). timestamp, iii). type, iv). owner , v). access control list.
FFS maintains i - iii, and the DS maintains iv & v.

Fault Tolerance: RPC interface is designed as *idempotent (refer to the definition of* ***idempotent operations in*** *Distributed Systems: Concepts & Design, by Coulouris et al, 2nd ed, page 111)* operations ensuring duplicated requests do not result in invalid updates to files,- servers can be *stateless* to enable restart to restore service after failure without recovering previous state.

DS: creates & modifies entries in simple one-D directories, looks-up text names in directories to return UFID after checking user's authorization. Separation of DS from FS enables different DS to be designed for use with a single file service, each supporting a diff name syntax and access control regime,- e.g. Unix, MS/DOS, VMS. Translation from file name to UFID is a *stateless* substitute for the *open file* operating in non-distributed systems. DS also looks after access control where UFIDs take the role of *capabilities.*

CM: hides low level constructs such as UFIDs used in RPC interfaces of the FFS and DS from user level programs, i.e. emulating I/O functions of the host operating system in the client mode. CM locates files distributed in several nodes based on identity of the file group.

Question 2

a) **Describe the principles of public key encryption and show how public key encryption can be used to transmit confidential information between two parties. (8 marks)**

b) **Identify the problems associated with the integrity of public and secret keys. (4 marks)**

c) **A user wishes to download a file from a Web site with confidence that the data has not been tampered with. Describe how a message digest can be used to achieve this and identify the main problem associated with this mechanism. (9 marks)**

d) **Describe a digital certificate and indicate how a digital certificate can provide confidence to the browser of a secure Web site. (4 marks)**

Answer Pointers

Of the 91% of candidates who attempted this question 60% achieved a pass standard. The average mark was 11.

a) Public key encryption is asymmetric. Many wrong answers discussed symmetric key encryption. Two large numbers called keys are generated at the same time. One is designated the public key and the other the secret or private key. Many answers did not say that the public and private keys are generated together and that data encrypted with one can only be decrypted with the other.
The private key must be known only to the owner, the public key must really belong to the person you think it does.

b) Many failed to mention that the private key could fall into the wrong hands. A message digest is a statistically unique fixed length hash of some data. It is typically created using the MD5 algorithm and is a 128 bit hash represented by a hexadecimal string. Both the data and the hash are published. The data can then be downloaded and a new hash generated and compared with the published hash. The problem is that if the Web site is cracked into then both the data and the published hash can be modified. A digital signature is a hash which has been encrypted by the data owner's secret key. This cannot be modified without the secret key. This works if the public key is genuine and the secret key has not been compromised.

c) There was a common misconception that a message digest is some form of data transmission program. Few stated that a message digest is a hash which is published with the file to be downloaded.

d) A digital certificate contains the server identity and its public key which have a digital signature created by the secret key of a trusted third party called a certificate authority. Web browsers contain the public keys of certificate authorities. Digital certificates provide a mechanism for obtaining a server's public key with some confidence that it is genuine.

Many missed the fact that a certificate is a mechanism for providing the public key of a web server with some confidence that it is genuine. An important fact is the certificate contains the server's public key which has been signed by a trusted third party.

Question 3

Very often the designer of a distributed system or application must consider issues that are largely unrelated to its distribution, such as software engineering techniques, human-computer interaction and algorithm design. However, you have been asked by your manager to restrict your attention to design issues that arise specifically from the distributed nature of systems. A full report covering these design issues would therefore consist of five sections to deal with: Naming, Communication, Software Structure, Workload Allocation and Consistency Maintenance.

You are required to produce a draft report, using suitable diagrams where necessary, covering only the section on Workload Allocation. Your report should be divided into the following 4 sections:

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)

Some thorough answers were evident which demonstrated that many candidates knew the material fairly well, based perhaps on actual experience and knowledge of syllabus material; some seemed to know the material well but were unsure to which section it belonged.
Of the 55% of candidates who attempted this question 78% achieved a pass standard.
The average mark was 13.

Answer Pointers

a) The Workstation-server model
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
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,
IW: as a fluctuating pool of extra computers, especially overnight, e.g. *worm programs* Xerox PARC used for animated graphics in parallel. Sprite operating system 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.
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 operating systems using shared memory as a hi-speed mechanism for inter-process communications and 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.

Question 4

a) **Describe the most direct mechanism by which an electronic mail message can Be sent between two users on separate workstations connected over the Internet using Mail User Agents (MUAs), Mail Delivery Agents (MDAs) and Mail Transfer Agents (MTAs). Give examples of actual software which can be used for each of MUAs, MDAs and MTAs. (15 Marks)**

b) **Describe how POP3 can be used to receive mail on a computer which is not permanently connected to the Internet. (5 Marks)**

c) **Define the terms *spoofing* and *spamming*. Explain how an MTA which is an open mail relay can be used by a spammer. (5 Marks)**

Few candidates obtained high marks for part a). Very few understood the role of the MDA which is to take a mail message from the receiver's MTA and place it in the recipient's mail box. Many answers discussed POP3 and message digests, both of which are irrelevant to this part.
In part b) most understood the POP3 mechanism.

And in part c) many discussed IP spoofing when the question was about mail spoofing - sending an email using someone else's address for the sender. Mail spoofing is the mechanism by which an open relay MTA can be made to emit spam. Very few got the open relay mechanism right.

Of the 74% of candidates who attempted this question 69% achieved a pass standard. The average mark was 12.

Answer Pointers

a) A message is composed in the MUA and passed to the sender's MTA. The MTA passes the message to the recipient's MTA. The MTA passes the message to an MDA which puts the message in the recipient's mail box. The recipient's MUA retrieves the message from the mail box for reading. MUA: pine, elm, netscape, mozilla, outlook, eurora. MTA:, sendmail, qmail, exim. MDA: procmail.

b) The message is received in the normal way by a mail server. The user's MUA can retrieve the mail message from the user's mail box on the mail server using the POP3 protocol. Many MUAs support POP3 directly, though there are programs which fetch all messages from a remote mail box into a local mail box.

c) Spoofing is the practise of sending an email which appears to have been sent by a third party. Spamming is the practice of sending large volumes of unsolicited email to a large number of different email addresses. An open relay MTA will forward email messages regardless of the sender or recipient. Spammers can send large volumes of mail through an open relay with any combination of sender or recipient addresses - even if the message generating machine has been blacklisted.

Question 5

a) **Replication is the maintenance of on-line copies of data and other resources. For example, the USENET system maintains a replica of each item posted to electronic bulletin boards across the Internet, each replica is held within, or close to, the various organizations that provide access to it.**

b) **Discuss the motivation for replication in terms of:**
performance enhancement, enhanced availability, and fault tolerance.

(9 marks)

The figure below shows a basic architectural model for the management of replicated data.

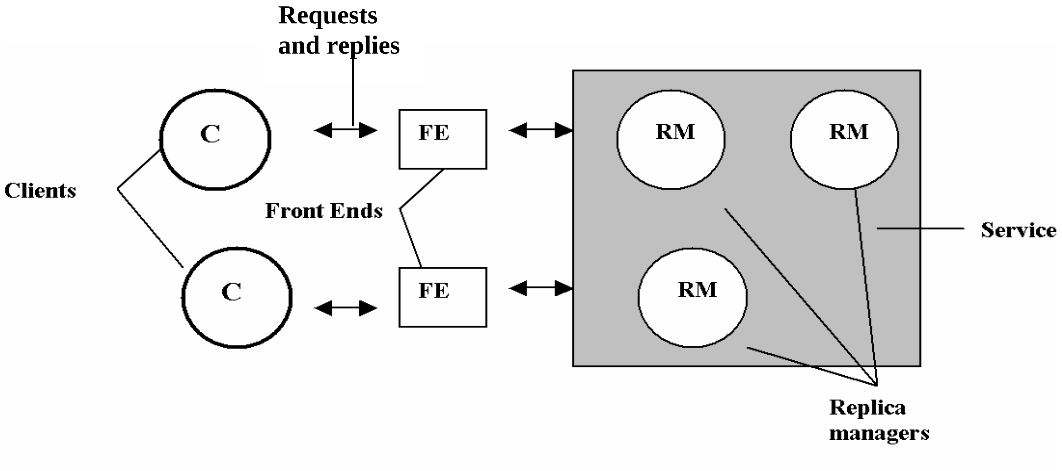

Figure 2. A basic model for the management of replicated data

i) **Briefly explain the operation of the model.**

(3 marks)

ii) **If the clients and replica managers are separate processes, the model becomes what is commonly referred to as the 'gossip' model. With the aid of a modified version of the above diagram explain the gossip model.**

(5 marks)

iii) **To enhance availability, the 'primary' copy model is often used. Modify the above diagram and explain the operation of this modified model. Suggest an architecture for a shared editor which might be used in a multi-user collaboration environment and referred to as 'groupware'.**

(8 marks)

Most candidates who attempted this question managed to give a reasonable answer to some parts of the question. Of the 44% of candidates who attempted this question 77% achieved a pass standard. The average mark was 13

Answer Pointers

a)
Performance enhancement: Data shared between a large client community is not held at a single server but distributed among many servers, each provides to a smaller community of users close to it.
Enhanced availability: client SW can access an alternative server when default server fails. If n servers have probability p of failing then enhanced availability is 1 - p^n . Much better than caching, which doesn't always hold files in their entirety,- Coda filing system is an exception.

Fault tolerance: if each server processes every client request in parallel it is possible to guarantee correct request processing should a server fails. It includes real time guarantees against arbitrary (Byzantine) failures, e.g. stock market, rocket engine calculations.

b)
i) Clients each make a series of requests first handled by FE’s to communicate by message passing with one or more of the replica managers. FE can be a user package executed in each client or a user process,- trade-off: efficiency & sharinf info.

ii) Diagram: 2 way communication between all RM’s which exchange gossip messages periodically to convey the updates (news) they receive. FE normally talks to one RM for each operation; alternatively, it can talk to more than one RM.
FE’s propagate their vector timestamps whenever clients communicate directly.

iii) Diagram: all FE’s talk with the same primary server (PS) when updating a particular data item; PS is accessed for all update requests,- PS propagates the updates to the other, slave servers (SS). FE’s may read the item from a slave. If PS fails then one of the SS is promoted to PS, e.g. in Sun Net Info Service, e.g. ‘Yellow Pages’ or special telephone directories with only business details, passwords which change infrequently are updated at a master server & propagated to SS.

Shared editor, diagram: each editor talks to all others, each circle is a combined client, FE and RM, one per user. Each holds a replica of the overall document state; only one class of process & it performs roles of client & RM; a FE module is used to hide replication from other modules that access the shared document state.