THE BRITISH COMPUTER SOCIETY

THE BCS PROFESSIONAL EXAMINATIONS
BCS Level 6 Professional Graduate Diploma in IT

NETWORK INFORMATION SYSTEMS

March 2015
Examiners' Report

General comments on candidates' performance

There was a variation in the standard of answers given to questions in section A of the paper.

Many candidates selected question 1 rather than questions 2 and 3, and answers were generally good.

Candidates are advised to pay attention to the marks allocated for each part of a question, as this is an indication of the complexity of the questions being asked. Although there is not necessarily any correlation between length of answers and marks awarded, it is unlikely that an answer for a question worth 25 marks will be complete if no more information is presented than for an answer to a question part worth six marks.
Some candidates presented very long answers to parts of questions worth only a few marks, and although typically these answers were correct, the same number of marks would be awarded for presenting the same information more succinctly.

SECTION A

A1. A company has an office based in London with a network that houses two web servers. One of these web servers is designed for external access, the other is an intranet web server designed only to be accessible to employees of the company. The company also has another office located in Paris, which is externally connected to the London office through a virtual private network routed over the public internet. This virtual private network must always be active, and not raised on demand by individual members of staff in the Paris office.

Both offices continue to use IP version 4 addresses, but the company only has a single public IP address range of 200.10.10.0/26. They require more nodes than the 62 allowed by this range, and so all nodes in the networks are numbered with private addresses from the 10.0.0.0/8 range.

Describe, with the aid of a diagram, a network topology that will interconnect the two offices and that will also allow public access to the external facing web server. You must clearly identify all routing or switching components required, and label IP addresses and subnets for at least one node in Paris, one node in London and all other identified components. Use Classless Internet Domain Routing (CIDR) notation to label the subnets. The diagram should also indicate how the routing of private addresses to the Internet is achieved, and the function of the relevant components should be described. **(25 marks)**

Answer Pointers

The private address range must be split into at least 2 parts. For instance, the London network could be 10.1.0.0/16 and Paris could use 10.2.0.0/16. These networks are then routed over a VPN that should be achieved in a router, e.g. an IPSec router. The answer does not require any specific VPN, but the candidate must recognise that there needs to be a point to point VPN over the routers through the public Internet. Both offices need a VPN end point that acts as a router, tunnelling encrypted data over the public internet. This VPN connection then handles the internal routing between subnets and there does not need to be any network address translation (NAT) at that point. **(10 marks)**

Outward connections from London do require NAT, and the Network Address Translator should be clearly described showing how one or more public IP addresses may be shared through the NAT, substituting the public addresses for the non externally routable private addresses. This is most concisely described by drawing up the NAT translation table showing the mappings of public IP addresses to private addresses and how the router rewrites the IP header with these addresses. Candidates will usually be more familiar with NPAT, sharing a single IP address. NPAT answers will be awarded most of the marks but an extra 2 marks are given to answers that recognise that the pool of 62 addresses allows a NAT solution without recourse to Port Address Translation. **(8 marks)**

The Web server must have a public address and may be located either on the external side of the NAT, or else within the NAT with a private address and a static mapping for one of the IP addresses. Both answers, if correctly described, carry full marks. **(4 marks)**

The intranet server should not have a public address and must be placed inside the privately addressed network with suitable IP address and subnet (e.g. 10.1.0.1/16). **(3 marks)**

Examiner's Comments

This was a popular question and generally answered very well, with some excellent answers. However, many network diagrams missed points that were asked for in the question. Attention to location of the web servers was important, and consideration of issues such as how the VPN should be set up. NAT is used very widely now, but many answers did not consider this aspect of the question. A few answers did not use CIDR notation, as was required in the question, either presenting subnet masks or not identifying subnets at all.

This was a single question, worth 25 marks, and so it is important to understand the scenario being presented and to attempt to produce a network map that would be full and workable. Many maps provided adequate information, but unable to gain marks for not considering the detail demanded in the scenario (particularly detail regarding NAT, the web server and VPN)

A2.

a) With reference to a packet switched network describe the concepts of :-

i) Traffic Contracts.
Clearly identify why these contracts are necessary and what the effect would be of having no such contracts. **(6 marks)**

ii) Traffic Shaping.
Include a suitable traffic shaping algorithm. **(7 marks)**

iii) Traffic Policing.
Include in your answer how the impact of traffic policing on networks with small frames or cells, such as Asynchronous Transfer Mode, may have severe and unintended consequences when carrying TCP/IP traffic, and how this may be alleviated. **(6 marks)**

b) Discuss the extent to which traffic shaping is supported in TCP/IP, and how such shaping may be avoided in some cases. **(6 marks)**

Answer Pointers

Part a) has 3 sections. (i) looks at traffic contracts where the packet switched network is often owned and managed by a service provider who supports multiple customers. These customers agree upon and pay for a certain level of bandwidth and performance from the service provider over that WAN. This agreement becomes the basis of the traffic contract, which defines the traffic parameters and the QoS that is negotiated for each virtual connection for that user on the network. Without such contracts, the shared packet switched network would become a free for all with no guaranties of service delivery. For instance, video conference traffic, which may need a limited bandwidth but with minimal delay could find insufficient bandwidth, and increased latencies at times of heavy usage. Customers with high bandwidth requirements could not pay for, nor have a guaranty of receiving more bandwidth than those with smaller requirements. **(6 marks)**

Part a (ii) then moves to traffic shaping, which may be used to enforce the contract and asks about a suitable algorithm. This calls for a description of either the leaky bucket algorithm or token bucket algorithm, or one of the variations on these. The algorithm should be adequately described to show how traffic shaping is achieved. **(7 marks)**

Part a (iii) asks candidates to be aware of or notice that dropping a single cell, where a cell is much shorter than a full sized packet, as is the case in IP over ATM, can cause a whole packet loss. If the algorithm for dropping packets just randomly drops occasional packets, it may heavily impact on overall network performance. Thus, when dropping cells, it is more efficient to drop a range of cells from the first cell dropped to the end of the transmission of the current packet. Two of the marks are reserved for knowledge of relative sizes of the ATM cells and IP packets (52 byte cells with 48 byte payloads in ATM compared to typically 1500 byte IP packets) **(6 marks)**

Part b asks about the use of differentiated services through the old QOS field in the IP header. Although the QOS field was not widely used, RFC 2474 provided a means to transform this into a differentiated services field to implement a means for classifying traffic and managing it for traffic contract purposes. Candidates are not expected to remember RFC numbers, but rather should understand how the DiffServ field in an IP header may be used for this. Candidates should also recognise that this is not enforceable across all networks, and different networks may treat information in the DiffServ header differently to that intended in the source network. Moreover, differentiated services can be circumvented through encrypted tunnelling of packets. Candidates may also mention traffic policing based on examination of following headers (e.g., looking at TCP port numbers in the TCP header). Although not intended by this question, candidates will be given most of the credit for this if they make the other points above, that the policing will be handled differently in other networks and can be circumvented through tunnelling encrypted packets. However, two marks are reserved for understanding of the DiffServ field. **(6 marks)**

Examiner's Comments

Although this question was much less popular than question 1, there were some very good answers, especially over part a) where some candidates presented good knowledge of traffic shaping algorithms such as the leaky bucket algorithm. Sadly, averages on the question were pulled down by candidates who started to answer it but did not provide any written answer to the questions.

Part b) of the question probed how well traffic shaping might be supported in TCP/IP, and at this point very few candidates knew of the DiffServ field in the IP header, but could pick up some of the marks for a discussion of QoS.

A3.

a) Demonstrate your understanding of a Remote Procedure Call (RPC), indicating how the client/server architecture may communicate with one another. Illustrate your answer with a diagram. **(7 marks)**

b) Clearly identify the difference between idempotent and non-idempotent operations. **(4 marks)**

c) Explain the difference between synchronous and asynchronous RPC calls. **(4 marks)**

d) Explain what is meant by "at least once" and "at most once" operations in RPC semantics, and how would these relate to idempotent and non-idempotent operations. **(5 marks)**

e) Remote Procedure Calls can be wrapped in XML-RPC or SOAP as an alternative to traditional Open Network Computing Remote Procedure Calls. List advantages and disadvantages of wrapping the calls in an XML schema. **(5 marks)**

Answer Pointers

Part a - Remote procedure calls are inter process communications that allow subroutines to be executed in a remote address space, either on the same node or on other network nodes. Local code is compiled against a stub that handles the packaging and dispatch of a message to the remote (server) process and deals with the reply but does not implement the subroutine itself. The server process provides the subroutine code as well as stubs for handling the messages from the client. A diagram should show client and server processes running on remote nodes (and optionally on separate address spaces in one node), connected by a network link (or loopback interface in the second case) over which messages are passed. Marks are awarded for a correct diagram, locating stubs, and correctly describing the message passing between processes as well as a clear representation of separate address spaces or nodes. **(7 marks)**

Part b - idempotent operations are those that have the same effect no matter how many times they are performed. Examples are reading a block of a file or setting an account balance to a particular value. Non-idempotent operations are the opposite. Their effects depend on how many times they are performed. Incrementing a count is non-idempotent, whereas reading a block of a (non volatile) file would be idempotent. **(4 marks)**

Part c – In synchronous calls, when the server is processing a call, the client is blocked (it waits until the server has finished processing before resuming execution). Asynchronous calls are non blocking. The client is not blocked until completion of the call. **(4 marks)**

Part d – At least once operations guarantee that a remote-procedure call (RPC) is performed by the server at least once, even in the face of server failures. A client might retransmit the RPC request repeatedly until a reply is received to achieve this, and the call may be handled more than once (as the server may fail after performing the operation but before sending the reply to the client, causing a repeat call to be sent and the server will handle the repeat call too). It is guaranteed to be handled at least once.

At most once operations will be handled no more than once by the server, even in the face of server failures. This can be accomplished by keeping track of identifying session numbers on the server that the clients must include in the RPC requests. The server increments the session number when it reboots and refuses to service RPC requests from the previous session (or epoch). Client's time-out when a reply is not received.

At least once RPC calls require idempotent operations since the RPC may be performed multiple times, and non idempotent operations would cause an inconsistent state. At most once RPC operations can use either idempotent or non idempotent operations, and a client will detect failure to process calls and decide how to proceed. **(5 marks)**

Part e – RPC requires stubs to be compiled into code (usually using an interface description language or IDL). It is these stubs that handle the message packaging and message passing to RPC services on this or a remote node. The RPC services are typically client and server processes that either open TCP connections or handle inter process communication via the UDP protocol over TCP/IP, but other implementations exist. These communications over specific ports, in the typical case, may be blocked by firewalls or be vectors for attack of a node, and where nodes already offer a web service, it was recognised that the transport service for the messages could be carried over existing client/server HTTP transport. Other existing transports can also be used by SOAP, e.g. SMTP messages.

Thus, wrapping the messages passed in the inter process communication of RPC inside industry standard XML with an appropriate schema could reduce complexity and improve security, and allowed for code that did not need to be compiled against the stubs if the messages were passed by API provided subroutines.

The key disadvantages of XML solutions surround verbosity of the XML messages, where XML data is wrapped in an XML schema that significantly increases the size of the RPC messages. Where these are transmitted across high latency links, and particularly where there is limited bandwidth, the extra verbosity can impact upon performance of the code using the XML based RPC solutions. **(5 marks)**

Examiner's Comments

This was another question that was not popular with candidates, although it covers a core area of the syllabus. Moreover, there is evidence that a number of candidates who attempted the question did not seem to understand the meaning of terms such as synchronous and asynchronous calls in the specific context of the question (i.e. RPC) and thus gave descriptions of synchronous communications in general.

Where candidates understood RPC, they provided good, and in a few cases, excellent answers. As with question 2, however, there were many candidates who left sections blank, and this was perhaps more telling in this question where each part tended to follow on from the previous part. A good answer to part b), for instance, guaranteed a good answer to part d), because understanding what an idempotent operation is foundational to part d).

The question also suggested illustrating the answer to part a) with a diagram. Some candidates did not provide this, and although they were not marked down for the lack of a diagram, they generally did not manage to convey the nature of client/server RPC communication without the diagram.

Diagrams are a means to quickly convey a lot of information, and where an answer can usefully be illustrated with a diagram, it is usually a good idea to do so.

B4

a) Briefly describe the security concepts of Confidentiality, Authentication, Data Integrity and Availability. **(4 Marks)**

b) All security threats are divided into four broad classes: **Disclosure, Disruption, Deception** and **Usurpation**.

Briefly discuss these threats and give an example of each. **(8 Marks)**

c) Clearly explain the digital signature concept and identify which, AND how, the security concepts listed in a) above are addressed by a digital signature. **(6 Marks)**

d) Explain how the BS7799 (or ISO 17799) standard relates to computer security. Relate your answer to the Plan-Do-Check-Act (PDCA) model in Information Security Management System (ISMS). **(7 Marks)**

Answer Pointers:

a)

Confidentiality is the concealment of information or resources. The need for keeping information secret arises from the use of computers in sensitive fields such as government and industry. For example, military and civilian institutions in the government often restrict access to information to those who need that information. **(1 Mark)**

Authentication is the process of verification of the server or the user to confirm that either is who they say they are and that they have the necessary privileges to access or change data. **(1 Mark)**

Data Integrity is the process of ensuring that the data are correct and have not been modified from their original state. **(1 Mark)**

Availability is the ability to use the information or resource desired by the users worldwide. The aspect of availability that is relevant to security is that someone may deliberately arrange to deny access to data or to a service by making it unavailable. **(1 Mark)**

b)
Disclosure is unauthorized access to information. **Snooping**, the unauthorized interception of information is a form of disclosure. **(2 Marks)**

Deception is acceptance of false data. **Spoofing**, an impersonation of one entity by another is a form of both deception and usurpation. Denial of receipt, a false denial that an entity received some information or message, is a form of deception. **(2 Marks)**

Disruption is interruption or prevention of correct operation. Denial of Service attacks on routers fall into this category. **(2 Marks)**

Usurpation is unauthorized control of some part of a system. Long-term inhibition of service is a form of usurpation, although it is often used with other mechanisms to deceive. The attacker prevents a server from providing a service. The denial may occur at the source by preventing the server from obtaining the resources needed to perform its function or through obtaining full control of the server. **(2 Marks)**

c) A digital signature is a mathematical scheme for demonstrating the authenticity of a digital message or document. **(2 Marks)**

A valid digital signature gives a recipient reason to believe that the message was created by a known sender, such that the sender cannot deny having sent the message – authentication (**(2 Marks)** - and that the message was not altered in transit – integrity **(2 Marks)**.

d) Initially developed from BS7799-1, ISO 17799 is an international standard that sets out the requirements of good practice for Information Security Management. **(3 Marks)**

An important aspect of ISO 27001 (ref BS 7799-2:2002) is that of the Plan-Do-Check-Act (PDCA) model, which must be applied to the ISMS. This is an approach to developing, implementing and improving the effectiveness of an organization's ISMS. **(4 Marks)**

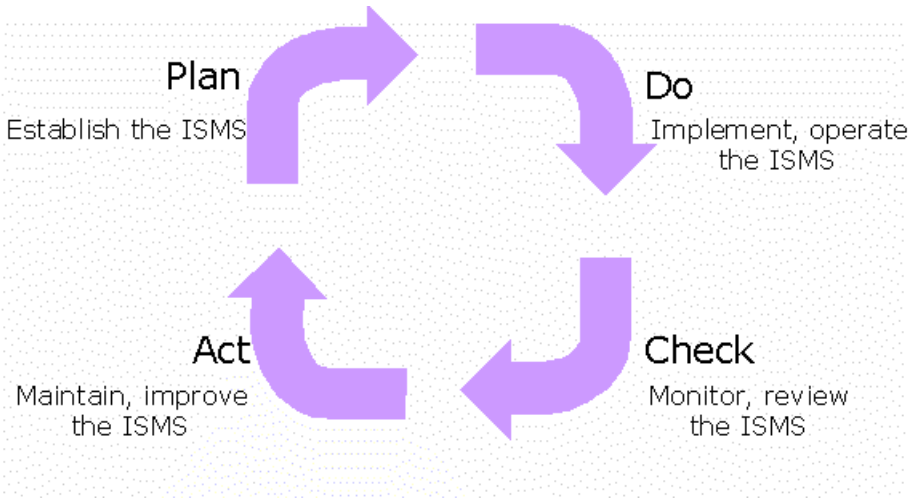

Examiner's Comments

This question covered security issues and was trying to cover the basic concepts of maintaining the data privacy and secure at the same time. It is matter of staying safe on-line irrespectively whether you program web-interface or you are a general user trying to see bank details on the Internet. These concepts need to be instilled in every user nowadays to prevent identity theft or other imminent threats.

Most of the students attempting the question managed to score a pass.

B5. A company provides gardening services and has a number of useful web applications to enable a customer to find their nearest garden centre and then to order products and services. The company wants to modernise and move to web services.

a) By providing a clear explanation for each term, demonstrate an understanding of the following terms: XML, SOAP, WSDL, and UDDI.
(10 Marks)

b) With the help of a diagram describe and justify the mechanism of working with web services on the garden centre company web site.
(15 Marks)

Answer Pointers

a)

XML is Extensible Markup Language.

XML solves the problem of data independence. You use it to describe data, and also to map that data into and out of any application or programming language.
(2 Marks)

WSDL is Web Services Description Language)

You use this XML-based language to create a description of an underlying application. It is this description that turns an application into a Web service, by acting as the interface between the underlying application and other Web-enabled applications. **(3 Marks)**

SOAP stands for Simple Object Access Protocol.

SOAP is the core communications protocol for the Web, and most Web services use this protocol to talk to each other. SOAP is an XML format for Web services requests.
(2 Marks)

UDDI is Universal Description, Discovery and Integration (UDDI) registry

A UDDI registry provides a way to publish and discover information about the Web services that are available within your organization. It does the same job for Web services that a business telephone directory does for business addresses and telephone numbers. **(3 Marks)**

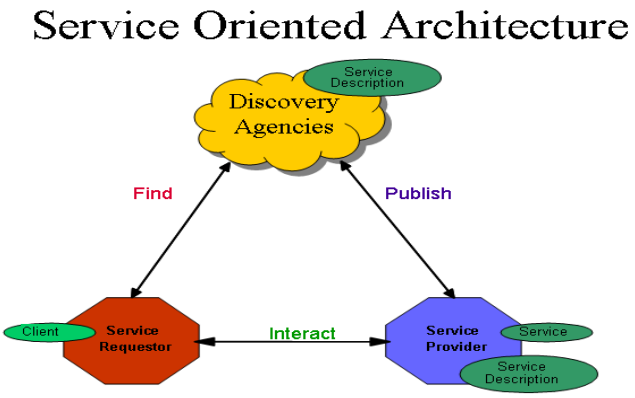

b)

The diagram or similar - **(7 Marks)**

Explanation of the diagram - **(8 Marks)**

The figure above illustrates the basic Web services architecture, in which a service requestor and service provider interact, based on the service's description information published by the provider and discovered by the requester through some form of discovery agency. The explanation must include references to XML, WSDL, SOAP and UDDI.

Examiner's Comment

This question was about the latest open standards developments. The technologies described in the question all use these standards and this gives an advantage to the companies using them. The standards allow better integration of the legacy technologies and opens up the services.