BCS Higher Education Qualification

Profession Graduate Diploma

April 2022

EXAMINERS' REPORT

Advanced Database Management Systems

General comments

Questions Report:

Qu.	Comment
A1	A1 The question covered the syllabus areas of database recovery, integrity and security. It was attempted by nearly half of the candidates with some excellent answers but a large number of superficial attempts, often addressing only some of the sub questions.
a)	Most candidates were able to define a transaction; the notion of change of state could have been clearer in some answers.
b)	While most candidates were able to identify what the two concepts were, there was a lack of detail in the explanation
c)	The concepts of authentication and authorisation have some very precise meanings as shown in the sample answer; answers provided by candidates often confused the two and either explained them the wrong way round or mixed aspects of each together, leading to answers that identified both as almost the same.
d)	The concept of database auditing caused many candidates problems; a good number of answers described in detail the idea of using a database for financial auditing, but the ask here was to look at auditing the database itself and not how databases can be used in other audits. Where candidates answered in the right domain, they frequently failed to describe how a database audit could be planned and executed (i.e. they did not provide a strategy/ technique).
A2	A2 This question covered concurrency control. It was attempted by over half of the candidates.
a)	Almost all candidates are able to define ACID and describe how each feeds into database correctness. The isolation property can reduce throughput, as it requires blocking of records to operate correctly. This was identified by many candidates, but without any explanation why. Some identified one of the other properties without giving a clear explanation as to why they thought that applied.
b)	Relatively clear and well argued answers from most students for the first part; however some focused away from the database problem and discussed the business matter of 'overbooking flights'. Answers to subquestion ii) were less convincing with many students not showing clarity of the key concepts and relatively few students being able to show clear examples of how the concepts manifest themselves.
c)	No additional comments
A3	The question covered the syllabus topic on referential integrity and SQL DDL. The most popular question in Part a) of the paper, with three quarters of candidates attempting it. However, a combination of not answering sub questions at all and not understanding the various NFs resulted in a low pass rate.
a)	Most candidates identified the repeating groups correctly; some managed to identify a subset of the attributes involved.

b)	Normalisation as a concept is understood by most candidates, however very often there is less clarity in the understanding of what 1NF and 2NF entail, compared to 3NF. Consequently, candidates often arrive at something close to the right answer but not necessarily attributing it at the right level. 1st NF in particular removes repeating groups – it does not introduce any 'new relations' or new tables.
c)	c) + d) Many candidates did not answer parts c) and d). Those that did generally answered well, but could often have provided some more detail in their answers.
B4	Syllabus 4.2 (Web services) 3.2 (New database applications and architectures). This was the least popular question on the paper. The overall performance was also weak, reflecting a lack of knowledge on this topic.
a) / b)	The diagram provided in this part of the question described a typical web services architecture. It was expected that this would drive candidates to describe in more detail the configuration and role of the various components that were presented. Unfortunately most candidates failed to describe in any significant detail how the various components interacted and achieved the main aim of providing seamless machine to machine integration when transporting data and services over the internet. It is recommended that candidates address this question by going through the various stages that are needed to handle the transport/communication of data from its source (the client) to the database and back again to the source. Very few candidates understood the processing involved and how the various components differed from a traditional n-tier (client application database) server architecture.
c)	A small number of candidates had knowledge of JSON largely seen as a more portable and lightweight data format than XML. Candidates are recommended to acquaint themselves with recent developments in the general area of NoSQL databases.
B5	Syllabus 3.2 New database applications and architectures: e.g. Data Warehousing. This question was reasonably popular. The overall performance was good and it was pleasing to observe that candidates have good knowledge in this topic and were able to apply their knowledge to the supplied scenario.
a)	Most candidates realised the importance of using a data warehouse to in effect combine many sources of historical data into one managed entity. The best answers expressed some excellent analysis for example that a data warehouse approach would consolidate and reinforce Keebo's marketing and management information system and that a data warehouse is extremely good at handling Keebo's legacy data and systems in a coherent way.
b)	Candidates' answers were good at following up their answers to part a) by addressing the concepts and techniques available in a data warehouse environment, in particular the staging of the ETL (Extract translate Load) process. The best answers were from candidates who went through each stage with examples of what data is extracted ; how data is transformed and loaded into a suitable form (such as star schema).
c)	Following on from part b) candidates were required to apply their knowledge of data models that a data warehouse supports. The highest marks were from candidates that produced some form of schema that related to the scenario of selling products. A number of candidates confused a snowflake schema with a star schema and lost marks as a result. This might explain why many candidates could not explain the final subpart of part c) in describing what is multidimensional modelling, which allows for example timeframes to be broken down to months then to weeks and days if necessary.