

BCS Higher Education Qualification

Diploma

May 2021

EXAMINERS' REPORT- Stephan Reiff-Marganec

Examination title – Advanced Database Management Systems

General comments

This exam covered the learning outcomes from the Database Systems syllabus.

Performance overall:

Overall performance was reasonable.

Where few marks were awarded candidates addressed only parts of the questions, did not include examples when asked and generally provided insufficient details. Candidates preferred questions that focused on 'terminology' and abstract concepts to those with 'practical' aspects; however they performed better in the latter where attempted.

Understanding of anything that is not strictly 'traditional databases' is weak.

I also noticed that candidates often chose to not answer questions – even a 'weak' answer can attract partial marks so some candidates are losing out by not answering questions.

Questions Report:

Q1	Comment
	<p>Around half of the candidates attempted this question and around half of the attempts achieved a pass mark.</p> <p>Overall understanding of BigData is relatively weak as shown by the results; When considering the 3 Vs, answers were often very simplistic (volume -> 'it's big', velocity -> 'need fast transport speed'). Answers should have considered more facets and provided more insight into the Vs. For part b) examples were not given or were very weak and very few candidates understood the point that the value comes from the insights gained. On part c) the answers focused very much on RDBMs and mostly stated that Hadoop is needed. NoSQL and solutions for unstructured and semi-structured data did not feature in many answers.</p>

Q2	Comment
	<p>This was not a popular question but of those who attempted it, the majority obtained pass marks. Part a) had some very good answers, but in the main, candidates generally explained what a trigger was and what the presented one did. For part b) some candidates wrote triggers, not queries and the 'self' (or recursive) relationship (a</p>

	manager is also an employee) caused problems for some. Answers to part c) were often very brief, not doing the question justice and again showing a lack of understanding of concepts beyond RDBMS.
--	---

Q3	Comment
	This was a very popular question with a high pass rate. Candidates did generally well on this question, there were minor mistakes in being sufficiently precise/ unambiguous in answers. Examples given in parts d) and e) also lacked crispness and precision, not allowing judgement of understanding.

Part B	
B4	This was a very popular question with a large majority of the candidates attempting it. Overall, a very satisfactory level of performance was achieved. Part a) generally had the most successful outcome in terms of marks showing candidates were well versed and had practised in creating an efficient way of processing a given query. Part c) produced a mixed response with a number of candidates unable to provide much in the way of an answer.
a)	The key to answering this part is to produce an efficient processing strategy on a complex SQL query that included multiple JOIN operations. Generally, most candidates were aware it was necessary to start by performing SELECT operations at the bottom of the query tree to reduce the number of rows processed in subsequent JOIN operations. Overall, this was the part in which the candidates gained the most success.
b)	In this part many candidates gave evidence of quite shallow knowledge with often vague answers. The key points were to state the advantages and the circumstances in which Denormalisation was beneficial. These are the main points that needed to be covered to gain full marks First, it is very important that prior to Denormalisation there must be a set of normalised tables to work with. Decisions can then be made on how to improve performance by selecting tables for denormalisation The circumstances in which denormalisation is beneficial are clearly in an OLAP (on line analytical processing) type of scenario rather than OLTP. Here, query processing predominates all other activities and is associated with topics such as data mining of big data sets.
c)	The CRUD acronym stands for CREATE READ UPDATE and DELETE. The relevant SQL statements are INSERT;SELECT;UPDATE;DELETE respectively. A number of candidates presented SQL operations rather than strings that would need to be embedded in the host programming language (for example, Java/PHP). Familiarity with an API and how connections to a database are made was necessary to appreciate what was actually required. There were many candidates who did not apply this knowledge and lost marks when standalone SQL code was submitted instead.
B5	Syllabus Coverage 5.1 : This was a fairly popular question with generally many good answers for all parts of the question. Again, as in previous years it was apparent that

	<p>many candidates failed to fully read and therefore understand the scenario. The scenario may have seemed familiar to most candidates but it was still very important to read it thoroughly before answering this question. To avoid duplicating answers particularly in parts a) and c) candidates are also advised to thoroughly read all parts of the question before writing any answers.</p>
a)	<p>(i) Answers describing full database replication needed to deal with the availability of data, problems associated with maintaining consistent copies of data and the update propagation problem</p> <p>(ii). Fragmentation. Apart from the fragmentation of the actual data, this type of distribution entails the fragmentation of the data dictionary, indexes and the overall impact on query performance</p> <p>Many candidates produced answers that lacked sufficient depth and clarity resulting in a loss of some of the marks.</p>
b)	<p>Most candidates applied reasonable knowledge on the different distribution strategies given a fairly well understood scenario.</p>
c)	<p>This part was answered reasonably. Candidates were able to distinguish between these two quite different distribution strategies.</p>