

BCS Higher Education Qualification

Profession Graduate Diploma

September 2019

EXAMINERS' REPORT

Advanced Database Management Systems

Question number: A1
Syllabus area:
Total marks allocated: 25
Examiners' Guidance Notes
<p>This question was not very popular, attempted by around a quarter of candidates. The overall performance was fairly good though with 60% of candidates obtaining a pass mark. Most candidates could differentiate the data models presented in part a) though the lack of depth in many answers was particularly noticeable, The best answers described real world examples that illustrated and demonstrated knowledge of the concepts of GIS (Geographical Information systems) that exhibit both spatial and graphical modelling concepts</p> <p>Part b) was generally less well answered with many candidates giving stock answers rather than opening up discussion of the comparisons of the two types of data model, At this level more discussion is expected to obtain marks when explicitly required in the question,</p>

Question number: A2
Syllabus area:
Total marks allocated: 25
Examiners' Guidance Notes
<p>This was a very popular question, attempted by a large majority of candidates. The overall performance was good with a high average mark of 13.5. Around 71% of candidates obtained a pass mark. This question required more factual and definitive answers. For Part a) candidates generally relied on illustrative examples of transaction schedules in their answer, The few candidates who produced answers without example schedules generally gained few marks. Part b) again required factual knowledge with most candidates unfamiliar with the purpose of Intent locks, Part c) produced mainly stock text book answers on 2PL,. It was surprising though that having demonstrated an understanding of 2PL, about half of candidates could not correctly determine whether the schedule provided (in Part d) actually complied with this protocol. In part e) it was surprising to find quite a lot of candidates unable to explain the full range of trade-offs because they were unfamiliar with the full range of locks that can be applied to database objects.</p>

Question number: A3

Syllabus area:

Total marks allocated: 25

Examiners' Guidance Notes

<p>This question was fairly popular, attempted by around half of the candidates. The question provided a good test of applying knowledge of fundamental techniques such as indexing, security and data integrity to given problems or scenarios. Parts a) and b) fully tested candidates knowledge in this regard. As expected, the overall performance was reasonable and in line with expectations, producing an average mark of 11.7. Just over half of candidates obtained a pass mark. Many candidates seem to take the term data integrity in part c) out of context. Answers that focused on features like check constraints, entity and referential integrity gained few marks because this was not what was asked.</p>

Question number: B4

Syllabus area:

Total marks allocated: 25

Examiners' Guidance Notes

<p>Most candidates attempted this question. Answers differentiating OLTP vs DataWarehouse were generally good, capturing the fundamental criteria and exploring details of the operational differences. The ETL process was generally explained well with candidates capturing the main purpose of the phases. Some candidates moved aspects of the T phase to E (especially conversion and cleaning), and some could not really explain the details of the purpose of the Transform phase focusing on the single aspect of mapping a specific datatype to an equal one (a common example was M=Male and F=Female). Candidates struggled with the concept of multiple inheritance (i.e. an object inheriting aspects from 2 parents simultaneously), quite often confusing it with inheritance over a number of generations. This clearly then meant that the concept of ambiguity was almost meaningless for them.</p>

Question number: B5

Syllabus area:

Total marks allocated: 25

Examiners' Guidance Notes

<p>Considering that this question was generally of a bookwork type, where concepts had to be understood and explained rather than applied, the answers were disappointing. Many candidates did not answer all of the sub questions. Extremely few candidates could explain that an outer join is a child-parent type of relation. Many did not appear to know what a trigger was, and even fewer could provide a coded example. The concept of a star schema was generally explained at some level, with candidates highlighting the key features and drawing a diagram. The fundamental concept of data replication was often explained, but at a very superficial level. Query optimization was discussed in detail using some query trees by many of the candidates. However a significant number of candidates did not answer this part.</p>
--