

2009 Paper 9 Question 2

Advanced Systems Topics

- (a) Distributed storage systems can typically be divided into *network attached storage* (NAS) and *storage area networks* (SAN).
- (i) Describe with the aid of a diagram the operation of a typical NAS system. Use as an example the access of a file by a client system. [3 marks]
 - (ii) Describe with the aid of a diagram the operation of a typical SAN system. Use as an example the access of a file by a client system. [3 marks]
 - (iii) Which would be more suitable for a high-performance database system? Justify your answer. [2 marks]
- (b) Database systems often use *indexes* in order to accelerate certain operations.
- (i) What exactly is an index used for? [1 mark]
 - (ii) Sequential indexes can be either sparse or dense. Give *two* advantages of sparse indexes and *two* advantages of dense indexes. [4 marks]
 - (iii) The B+ -tree is a commonly used data-structure for implementing indexes. Sketch the structure of a B+ -tree, and describe how lookup and insertion occur. [4 marks]
 - (iv) A related data-structure is the B-tree. What are the differences between B-trees and B+ -trees? [2 marks]
 - (v) Why are B+ -trees typically preferred over B-trees in database systems? [1 mark]