Exam : **LLM-ACP**

Title : Alibaba Cloud LLM Engineer (Professional)

Version : DEMO

1.Which practice best reduces the risk of a fine-tuned model memorizing sensitive information?
A. Remove or mask sensitive data before training
B. Add more sensitive examples to improve recall
C. Disable all dataset review steps
D. Publish the training data for easier debugging
Answer: A

2.Which practices improve production reliability for LLM applications? (Choose TWO answers)
A. Continuous monitoring of inference performance
B. Controlled deployment and rollback procedures
C. Storing model artifacts only on local laptops
D. Disabling health checks to reduce overhead
E. Deploying updates directly to production without testing
Answer: AB

3.A user asks a question that is not answered by the retrieved documents.
What should a well-designed RAG application do?
A. State that the available context does not contain the answer
B. Invent a likely answer based on general assumptions
C. Always cite the most similar unrelated passage
D. Disable retrieval for all future questions
Answer: A

4.An LLM application sometimes exceeds the model context limit when users submit long documents.
Which design is most appropriate?
A. Split or summarize input before sending relevant context to the model
B. Send the full document repeatedly until the request succeeds
C. Remove all user instructions from the prompt permanently
D. Disable validation and accept incomplete responses
Answer: A

5.Which practices help evaluate prompt quality? (Choose TWO answers)
A. Test prompts against representative user inputs
B. Compare outputs against expected quality criteria
C. Evaluate only a single successful example
D. Ignore edge cases and unsafe requests
E. Change prompts in production without review
Answer: AB