Examcollection CSPAI Dumps Torrent & CSPAI Exam Papers

Answer: C

NEW QUESTION # 22
What is a function of Cisco Intersight?

- A. Cisco Intersight is a great tool for core management and complication
- B. Cisco Intersight includes a RESTful API to provide full programmability
- C. Cisco Intersight provides Platform as a Service management
- D. Cisco Intersight automates many tasks including RBAC and reporting of cases to TAC

Answer: D

NEW QUESTION # 23

Cisco - Reliable 700-755 - Examcollection Cisco Small Business Technical Overview Dumps Torrent

If you hope to get a job with opportunity of promotion, it will be the best choice chance for you to choose the CSPAI study question from our company. Because our CSPAI study materials have the enough ability to help you improve yourself and make you more excellent than other people. The CSPAI Learning Materials from our company have helped a lot of people get the certification and achieve their dreams. And you also have the opportunity to contact with the CSPAI test guide from our company.

SISA CSPAI Exam Syllabus Topics:

Topic	Details
Topic 1	• Improving SDLC Efficiency Using Gen AI: This section of the exam measures skills of the AI Security Analyst and explores how generative AI can be used to streamline the software development life cycle. It emphasizes using AI for code generation, vulnerability identification, and faster remediation, all while ensuring secure development practices.
Topic 2	• Models for Assessing Gen AI Risk: This section of the exam measures skills of the Cybersecurity Risk Manager and deals with frameworks and models used to evaluate risks associated with deploying generative AI. It includes methods for identifying, quantifying, and mitigating risks from both technical and governance perspectives.

Topic 3	• Using Gen AI for Improving the Security Posture: This section of the exam measures skills of the Cybersecurity Risk Manager and focuses on how Gen AI tools can strengthen an organization's overall security posture. It includes insights on how automation, predictive analysis, and intelligent threat detection can be used to enhance cyber resilience and operational defense.
Topic 4	• AIMS and Privacy Standards: ISO 42001 and ISO 27563: This section of the exam measures skills of the AI Security Analyst and addresses international standards related to AI management systems and privacy. It reviews compliance expectations, data governance frameworks, and how these standards help align AI implementation with global privacy and security regulations.
Topic 5	• Evolution of Gen AI and Its Impact: This section of the exam measures skills of the AI Security Analyst and covers how generative AI has evolved over time and the implications of this evolution for cybersecurity. It focuses on understanding the broader impact of Gen AI technologies on security operations, threat landscapes, and risk management strategies.

CSPAI Exam Papers, CSPAI Valid Test Papers

Sometimes choice is greater than important. Good choice may do more with less. If you still worry about your exam, our SISA CSPAI braindump materials will be your right choice. Our exam braindumps materials have high pass rate. Most candidates purchase our products and will pass exam certainly. If you want to fail exam and feel depressed, our SISA CSPAI braindump materials can help you pass exam one-shot.

SISA Certified Security Professional in Artificial Intelligence Sample Questions (Q16-Q21):

NEW QUESTION # 16
When integrating LLMs using a Prompting Technique, what is a significant challenge in achieving consistent performance across diverse applications?

- A. Handling the security concerns that arise from dynamically generated prompts
- B. The need for optimizing prompt templates to ensure generalization across different contexts.
- C. Overcoming the lack of transparency in understanding how the LLM interprets varying prompt structures.
- D. Reducing latency in generating responses to meet real-time application requirements.

Answer: B

Explanation:
Prompting techniques in LLM integration, such as zero-shot or few-shot prompting, face challenges in consistency due to the need for meticulously optimized templates that generalize across tasks. Variations in prompt phrasing can lead to unpredictable outputs, requiring iterative engineering to balance specificity and flexibility, especially in diverse domains like legal or medical apps. This optimization involves A/B testing, semantic alignment, and incorporating chain-of-thought to enhance reasoning, but it demands expertise and time in SDLC phases. Unlike latency issues, which are hardware-related, prompt optimization directly affects performance reliability. Security overlaps, as poor prompts might expose vulnerabilities, but the core challenge is generalization. Efficient SDLC uses automated prompt tuning tools to streamline this, reducing development overhead while maintaining efficacy. Exact extract: "A significant challenge is optimizing prompt templates to ensure generalization across different contexts, crucial for consistent LLM performance in varied applications." (Reference: Cyber Security for AI by SISA Study Guide, Section on Prompting in SDLC, Page 100-103).

NEW QUESTION # 17
In a time-series prediction task, how does an RNN effectively model sequential data?

- A. By processing each time step independently, optimizing the model's performance over time.
- B. By storing only the most recent time step, ensuring efficient memory usage for real-time predictions
- C. By focusing on the overall sequence structure rather than individual time steps for a more holistic approach.

- D. By using hidden states to retain context from prior time steps, allowing it to capture dependencies across the sequence.

Answer: D

Explanation:
RNNs model sequential data in time-series tasks by maintaining hidden states that propagate information across time steps, capturing temporal dependencies like trends or seasonality. This memory mechanism allows RNNs to learn from past data, unlike independent processing or holistic approaches, though they face gradient issues for long sequences. Exact extract: "RNNs use hidden states to retain context from prior time steps, effectively capturing dependencies in sequential data for time-series tasks." (Reference: Cyber Security for AI by SISA Study Guide, Section on RNN Architectures, Page 40-43).

NEW QUESTION # 18
In line with the US Executive Order on AI, a company's AI application has encountered a security vulnerability. What should be prioritized to align with the order's expectations?

- A. Immediate public disclosure of the vulnerability.
- B. Ignoring the vulnerability if it does not affect core functionalities.
- C. Implementing a rapid response to address and remediate the vulnerability, followed by a review of security practices.
- D. Halting all AI projects until a full investigation is complete.

Answer: C

Explanation:
The US Executive Order on AI emphasizes proactive risk management and robust security to ensure safe AI deployment. When a vulnerability is detected, rapid response to remediate it, coupled with a thorough review of security practices, aligns with these mandates by minimizing harm and preventing recurrence. This approach involves patching the issue, assessing root causes, and updating protocols to strengthen defenses, ensuring compliance with standards like ISO 42001, which prioritizes risk mitigation in AI systems. Public disclosure, while important, is secondary to remediation to avoid premature exposure, and halting projects is overly disruptive unless risks are critical. Ignoring vulnerabilities contradicts responsible AI principles, risking regulatory penalties and trust erosion. This strategy fosters accountability and aligns with governance frameworks for secure AI operations. Exact extract: "Addressing vulnerabilities promptly through remediation and reviewing security practices is prioritized to meet the US Executive Order's expectations for safe and secure AI systems." (Reference: Cyber Security for AI by SISA Study Guide, Section on AI Governance and US EO Compliance, Page 165-168).

NEW QUESTION # 19
When deploying LLMs in production, what is a common strategy for parameter-efficient fine-tuning?

- A. Using external reinforcement learning to adjust the model's parameters dynamically.
- B. Training the model from scratch on the target task to achieve optimal performance.
- C. Implementing multiple independent models for each specific task instead of fine tuning a single model
- D. Freezing the majority of model parameters and only updating a small subset relevant to the task

Answer: D

Explanation:
Parameter-efficient fine-tuning (PEFT) strategies, like LoRA or adapters, freeze most pretrained parameters and train only lightweight modules, reducing computational costs while adapting to new tasks. This preserves general knowledge, prevents catastrophic forgetting, and enables quick deployments in resource-constrained settings. For LLMs, it's crucial for efficiency in production, allowing specialization without retraining billions of parameters. Security-wise, it minimizes exposure to new data risks. Exact extract: "A common strategy is freezing the majority of model parameters and updating only a small task-relevant subset, ensuring efficiency in fine-tuning for production deployment." (Reference: Cyber Security for AI by SISA Study Guide, Section on Efficient Fine-Tuning in SDLC, Page 90-92).

NEW QUESTION # 20
In transformer models, how does the attention mechanism improve model performance compared to RNNs?

- A. By processing each input independently, ensuring the model captures all aspects of the sequence equally.
- B. By enhancing the model's ability to process data in parallel, ensuring faster training without compromising context.

- C. By enabling the model to attend to both nearby and distant words simultaneously, improving its understanding of long-term dependencies
- D. By dynamically assigning importance to every word in the sequence, enabling the model to focus on relevant parts of the input.

Answer: C

Explanation:
Transformer models leverage self-attention to process entire sequences concurrently, unlike RNNs, which handle inputs sequentially and struggle with long-range dependencies due to vanishing gradients. By computing attention scores across all words, Transformers capture both local and global contexts, enabling better modeling of relationships in tasks like translation or summarization. For example, in a long sentence, attention links distant pronouns to their subjects, improving coherence. This contrasts with RNNs' sequential limitations, which hinder capturing far-apart dependencies. While parallelism (option C) aids efficiency, the core improvement lies in dependency modeling, not just speed. Exact extract: "The attention mechanism enables Transformers to attend to nearby and distant words simultaneously, significantly improving long-term dependency understanding over RNNs." (Reference: Cyber Security for AI by SISA Study Guide, Section on Transformer vs. RNN Architectures, Page 50-53).

NEW QUESTION # 21
.....