


CSPAI Sample Questions | CSPAI Exam Demo



P.S. Free & New CSPAI dumps are available on Google Drive shared by ITCertMagic: <https://drive.google.com/open?id=1aUSSUbaOQSswWasunt5OpJBVA6v-IYbm>

They can print these real Certified Security Professional in Artificial Intelligence (CSPAI) questions to save them as paper notes. And you can also use the Certified Security Professional in Artificial Intelligence (CSPAI) PDF on smart devices like smartphones, laptops, and tablets. The second one is the web-based Certified Security Professional in Artificial Intelligence (CSPAI) practice exam which can be accessed through the browsers like Firefox, Safari, and Google Chrome.

SISA CSPAI Exam Syllabus Topics:

Topic	Details
Topic 1	• 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 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	• Securing AI Models and Data: This section of the exam measures skills of the Cybersecurity Risk Manager and focuses on the protection of AI models and the data they consume or generate. Topics include adversarial attacks, data poisoning, model theft, and encryption techniques that help secure the AI lifecycle.

Topic 4	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.
---------	--

>> CSPAI Sample Questions <<

CSPAI Exam Demo & Latest CSPAI Exam Registration

Achieving a good score on the SISA CSPAI exam on the first attempt is a common goal for many candidates. However, some believe that studying good Certified Security Professional in Artificial Intelligence (CSPAI) materials isn't necessary. This notion, however, is far from true. The right preparation material for the CSPAI Exam is critical for success, and failing to find the most up-to-date SISA CSPAI materials can lead to a wasted effort and expense.

SISA Certified Security Professional in Artificial Intelligence Sample Questions (Q25-Q30):

NEW QUESTION # 25

What is a key benefit of using GenAI for security analytics?

- A. Predicting future threats through pattern recognition in large datasets.
- B. Limiting analysis to historical data only.
- C. Increasing data silos to protect information.
- D. Reducing the use of analytics tools to save costs.

Answer: A

Explanation:

GenAI revolutionizes security analytics by mining massive datasets for patterns, predicting emerging threats like zero-day attacks through generative modeling. It synthesizes insights from disparate sources, enabling proactive defenses and anomaly detection with high precision. This foresight allows organizations to allocate resources effectively, preventing breaches before they occur. In practice, it integrates with SIEM systems for enhanced threat hunting. The benefit lies in transforming reactive security into predictive, bolstering posture against sophisticated adversaries. Exact extract: "A key benefit of GenAI in security analytics is predicting future threats via pattern recognition, improving proactive security measures." (Reference: Cyber Security for AI by SISA Study Guide, Section on Predictive Analytics with GenAI, Page 220-223).

NEW QUESTION # 26

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. Freezing the majority of model parameters and only updating a small subset relevant to the task
- C. Implementing multiple independent models for each specific task instead of fine tuning a single model
- D. Training the model from scratch on the target task to achieve optimal performance.

Answer: B

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 # 27

During the development of AI technologies, how did the shift from rule-based systems to machine learning models impact the efficiency of automated tasks?

- A. Enhanced the precision and relevance of automated outputs with reduced manual tuning.
- B. Improved scalability and performance in handling diverse and evolving data.
- C. Increased system complexity and the requirement for specialized knowledge,
- **D. Enabled more dynamic decision-making and adaptability with minimal manual intervention**

Answer: D

Explanation:

The transition from rigid rule-based systems, which rely on predefined logic and struggle with variability, to machine learning models introduced data-driven learning, allowing systems to adapt dynamically to new patterns with less human oversight. This shift boosted efficiency in automated tasks by enabling real-time adjustments, such as in spam detection where ML models evolve with threats, unlike static rules. It minimized manual rule updates, fostering scalability and handling complex, unstructured data effectively. However, it introduced challenges like interpretability needs. In GenAI evolution, this paved the way for advanced models like Transformers, impacting sectors by automating nuanced decisions. Exact extract: "The shift enabled more dynamic decision-making and adaptability with minimal manual intervention, significantly improving the efficiency of automated tasks." (Reference: Cyber Security for AI by SISA Study Guide, Section on AI Evolution and Impacts, Page 20-23).

NEW QUESTION # 28

In a scenario where Open-Source LLMs are being used to create a virtual assistant, what would be the most effective way to ensure the assistant is continuously improving its interactions without constant retraining?

- **A. Implementing reinforcement learning from human feedback (RLHF) to refine responses based on user input.**
- B. Reducing the amount of feedback integrated to speed up deployment.
- C. Shifting the assistant to a completely rule-based system to avoid reliance on user feedback.
- D. Training a larger proprietary model to replace the open-source LLM

Answer: A

Explanation:

For continuous improvement in open-source LLM-based virtual assistants, RLHF integrates human evaluations to align model outputs with preferences, iteratively refining behavior without full retraining. This method uses reward models trained on feedback to guide policy optimization, enhancing interaction quality over time. It addresses limitations like initial biases or suboptimal responses by leveraging real-world user inputs, making the system adaptive and efficient. Unlike full retraining, RLHF is parameter-efficient and scalable, ideal for production environments. Security benefits include monitoring feedback for adversarial attempts. Exact extract: "Implementing RLHF allows continuous refinement of the assistant's interactions based on user feedback, avoiding the need for constant full retraining while improving performance." (Reference: Cyber Security for AI by SISA Study Guide, Section on AI Improvement Techniques in SDLC, Page 85-88).

NEW QUESTION # 29

Fine-tuning an LLM on a single task involves adjusting model parameters to specialize in a particular domain. What is the primary challenge associated with fine tuning for a single task compared to multi task fine tuning?

- A. Single-task fine-tuning tends to degrade the model's performance on the original tasks it was trained on.
- B. Single-task fine-tuning introduces more complexity in managing different versions of the model compared to multi-task fine-tuning.
- **C. Single-task fine-tuning is less effective in generalizing to new, unseen tasks compared to multi-task fine-tuning.**
- D. Single-task fine-tuning requires significantly more data to achieve comparable performance to multi-task fine tuning.

Answer: C

Explanation:

Single-task fine-tuning specializes the LLM but risks overfitting, limiting generalization to novel tasks unlike multi-task approaches that promote transfer learning across domains. This challenge requires careful regularization in SDLC to balance specificity and versatility, often needing more resources for version management. Exact extract: "Single-task fine-tuning is less effective in generalizing to new tasks compared to multi-task fine-tuning." (Reference: Cyber Security for AI by SISA Study Guide, Section on

• • • • •

CSPA Exam Demo: <https://www.itcertmagic.com/SISA/real-CSPA-exam-prep-dumps.html>

- BTW, DOWNLOAD part of ITCertMagic CSPAI dumps from Cloud Storage: <https://drive.google.com/open?id=1aUSSUbaOQSswWasunt5OpJBVA6v-IYbm>