EXAMTOPICS

-Expert Verified, Online, **Free**.



CERTIFICATION TEST

-CertificationTest.net - Cheap & Quality Resources With Best Support

Question #1 *Topic 1*

How much time is usually required for pre-training a model?

A. More than 1 month

B. 3 weeks

C. 2 weeks

D. 1 week

Suggested Answer: *B*

Currently there are no comments in this discussion, be the first to comment!

Question #2 *Topic 1*

Which of the following statements are true about the features of the baseline solution for the hardware foundation of Huawei's Center Training and Inference Computing Platform? (Choose all that apply.)

A. Third-party AI platforms are not supported.

B. Computing and storage collaboration, fast inference speed

C. Collaborative computing and network, improving the performance of the computing infrastructure.

D. Cloud-edge collaboration, building an elastic resource pool

Suggested Answer: *B, C, D*

Currently there are no comments in this discussion, be the first to comment!

Question #3 *Topic 1*

On which plane do read and write operations of sample data between compute nodes and storage nodes occur?

A. Management plane

B. Sample plane

C. Service plane

D. Parameter plane

Suggested Answer: *B*

Currently there are no comments in this discussion, be the first to comment!

Question #4 *Topic 1*

What parameter scale of large models in MoE scenarios does Ascend multi-machine inference support for deployment?

A. 1 billion

B. 10 billion

C. 100 billion

D. 1 trillion

Suggested Answer: *D*

Currently there are no comments in this discussion, be the first to comment!

Question #5 *Topic 1*

Which of the following computing capabilities has extended from the device-side to the server market and has become a new option in the field of general computing?

A. ARM

B. x86

C. MIPS

D. RISC-V

Suggested Answer: *A*

Currently there are no comments in this discussion, be the first to comment!

Question #6 *Topic 1*

In Huawei AI solution, the computing platform reads and writes data in the storage system through the storage service plane.

A. TRUE

B. FALSE

Suggested Answer: *A*

Currently there are no comments in this discussion, be the first to comment!

Question #7 *Topic 1*

What are the main features of the liquid cooling solution in new AI equipment room scenarios?

A. Improved computing power per unit of electricity consumed

B. 1 week go-live

C. PUE=1.6

D. Low Capex investment

Suggested Answer: *A*

Currently there are no comments in this discussion, be the first to comment!

Question #8 *Topic 1*

Which capabilities can Huawei AI Storage provide during the AI model training process?

A. Supports hot and cold data tiering

B. Multi-protocol interworking, supporting multiple types of data

C. Supports high-performance read and write for hybrid workloads, and supports high-speed data loading and CKPT saving.

D. Support input of long sequence requirements

Suggested Answer: *C*

Currently there are no comments in this discussion, be the first to comment!

Question #9 *Topic 1*

System high availability means maximizing fault-free running time and minimizing fault recovery time.

A. TRUE

B. FALSE

Suggested Answer: *A*

Currently there are no comments in this discussion, be the first to comment!

Question #10 *Topic 1*

Which of the following statements are correct about large language models and capability emergence? (Choose all that apply.)

A. When the number of parameters, data volume, and training intensity reach a critical value, the model capability will experience a qualitative leap.

B. Emergent capabilities enable large language models to achieve stronger general intelligence.

C. Once emergent phenomena occur, the model loses its fundamental understanding capabilities.

D. The model capability improves only slightly as the model size increases; it does not suddenly become more powerful.

Suggested Answer: *A, B*

Currently there are no comments in this discussion, be the first to comment!

Question #11 *Topic 1*

Which of the following features is not included in the Ascend AI Training Solution designed for large-scale computing power clusters?

A. Highly ease of use

B. High performance

C. High cost

D. High availability

Suggested Answer: *C*

Currently there are no comments in this discussion, be the first to comment!

Question #12 *Topic 1*

The CCAE Cluster Operation and Management System transforms cluster management from “single-domain mode” to “centralized governance”, achieving full visibility and controllability of training job paths across the entire system.

A. TRUE

B. FALSE

Suggested Answer: *A*

Currently there are no comments in this discussion, be the first to comment!

Question #13 *Topic 1*

The rapid development of new models and new operators relies on user-friendly tools and an open ecosystem.

A. TRUE

B. FALSE

Suggested Answer: *A*

Currently there are no comments in this discussion, be the first to comment!

Question #14 *Topic 1*

Huawei fully opens up the four-layer architecture covering operators, frameworks, acceleration libraries and models.

A. TRUE

B. FALSE

Suggested Answer: *B*

Currently there are no comments in this discussion, be the first to comment!

Question #15 *Topic 1*

Large scale does not equate to high computing power. What challenges are faced in the construction of high-computing-power clusters? (Choose all that apply.)

A. Cluster scale

B. Development Ease of Use

C. Cluster Availability

D. Cluster performance

Suggested Answer: *A, B, C, D*

Currently there are no comments in this discussion, be the first to comment!

Question #16 *Topic 1*

What networking mode is usually used for large-scale AI clusters to support the training of 100-billion or trillion parameter models?

A. 2-level box-based networking

B. 3-level box networking

C. 2-level box networking

D. 3-level box-based networking

Suggested Answer: *C*

Currently there are no comments in this discussion, be the first to comment!

Question #17 *Topic 1*

Ascend training and inference use the same architecture, but the training and inference operator capabilities are not compatible.

A. TRUE

B. FALSE

Suggested Answer: *B*

Currently there are no comments in this discussion, be the first to comment!

Question #18 *Topic 1*

Which of the following statements about the Atlas 800I A2 inference server are true? (Choose all that apply.)

A. The overall inference throughput is higher, and the user experience of latency is better.

B. Recommended for non-latency-sensitive scenarios

C. 8 GPUs per server, full-meshed high bandwidth

D. Supports models of hundreds of billions of parameters and long-sequence inference.

Suggested Answer: *A, C, D*

Currently there are no comments in this discussion, be the first to comment!

Question #19 *Topic 1*

Which of the following statements about the Atlas 300I Duo inference card are true? (Choose all that apply.)

A. Standard cards deliver a balanced cost-performance ratio

B. The 280T INT8 high computing power is suitable for computing power-sensitive scenarios (such as text-to-image generation), improving user experience.

C. 96 GB large memory, enabling deployment of larger parameter models and support for more concurrent users, thereby reducing construction costs

D. Recommended for latency-sensitive scenarios

Suggested Answer: *A, B, C, D*

Currently there are no comments in this discussion, be the first to comment!

Question #20 *Topic 1*

Ascend enables users to develop custom operators, efficiently meeting operator compatibility requirements in heterogeneous migration scenarios.

A. TRUE

B. FALSE

Suggested Answer: *A*

Currently there are no comments in this discussion, be the first to comment!

Question #21 *Topic 1*

What are the dimensions and form factor of the Atlas 800I A2 inference server?

A. 1U

B. 2U

C. 4U

D. 8U

Suggested Answer: *C*

Currently there are no comments in this discussion, be the first to comment!

Question #22 *Topic 1*

Latency, accuracy, throughput, and ease of use are the core metrics for building a leading AI inference system.

A. TRUE

B. FALSE

Suggested Answer: *A*

Currently there are no comments in this discussion, be the first to comment!

Question #23 *Topic 1*

The development trend of foundation model inference is an increase in inference overhead and a rise in the cost per token.

A. TRUE

B. FALSE

Suggested Answer: *B*

Currently there are no comments in this discussion, be the first to comment!