Practice Exam Questions

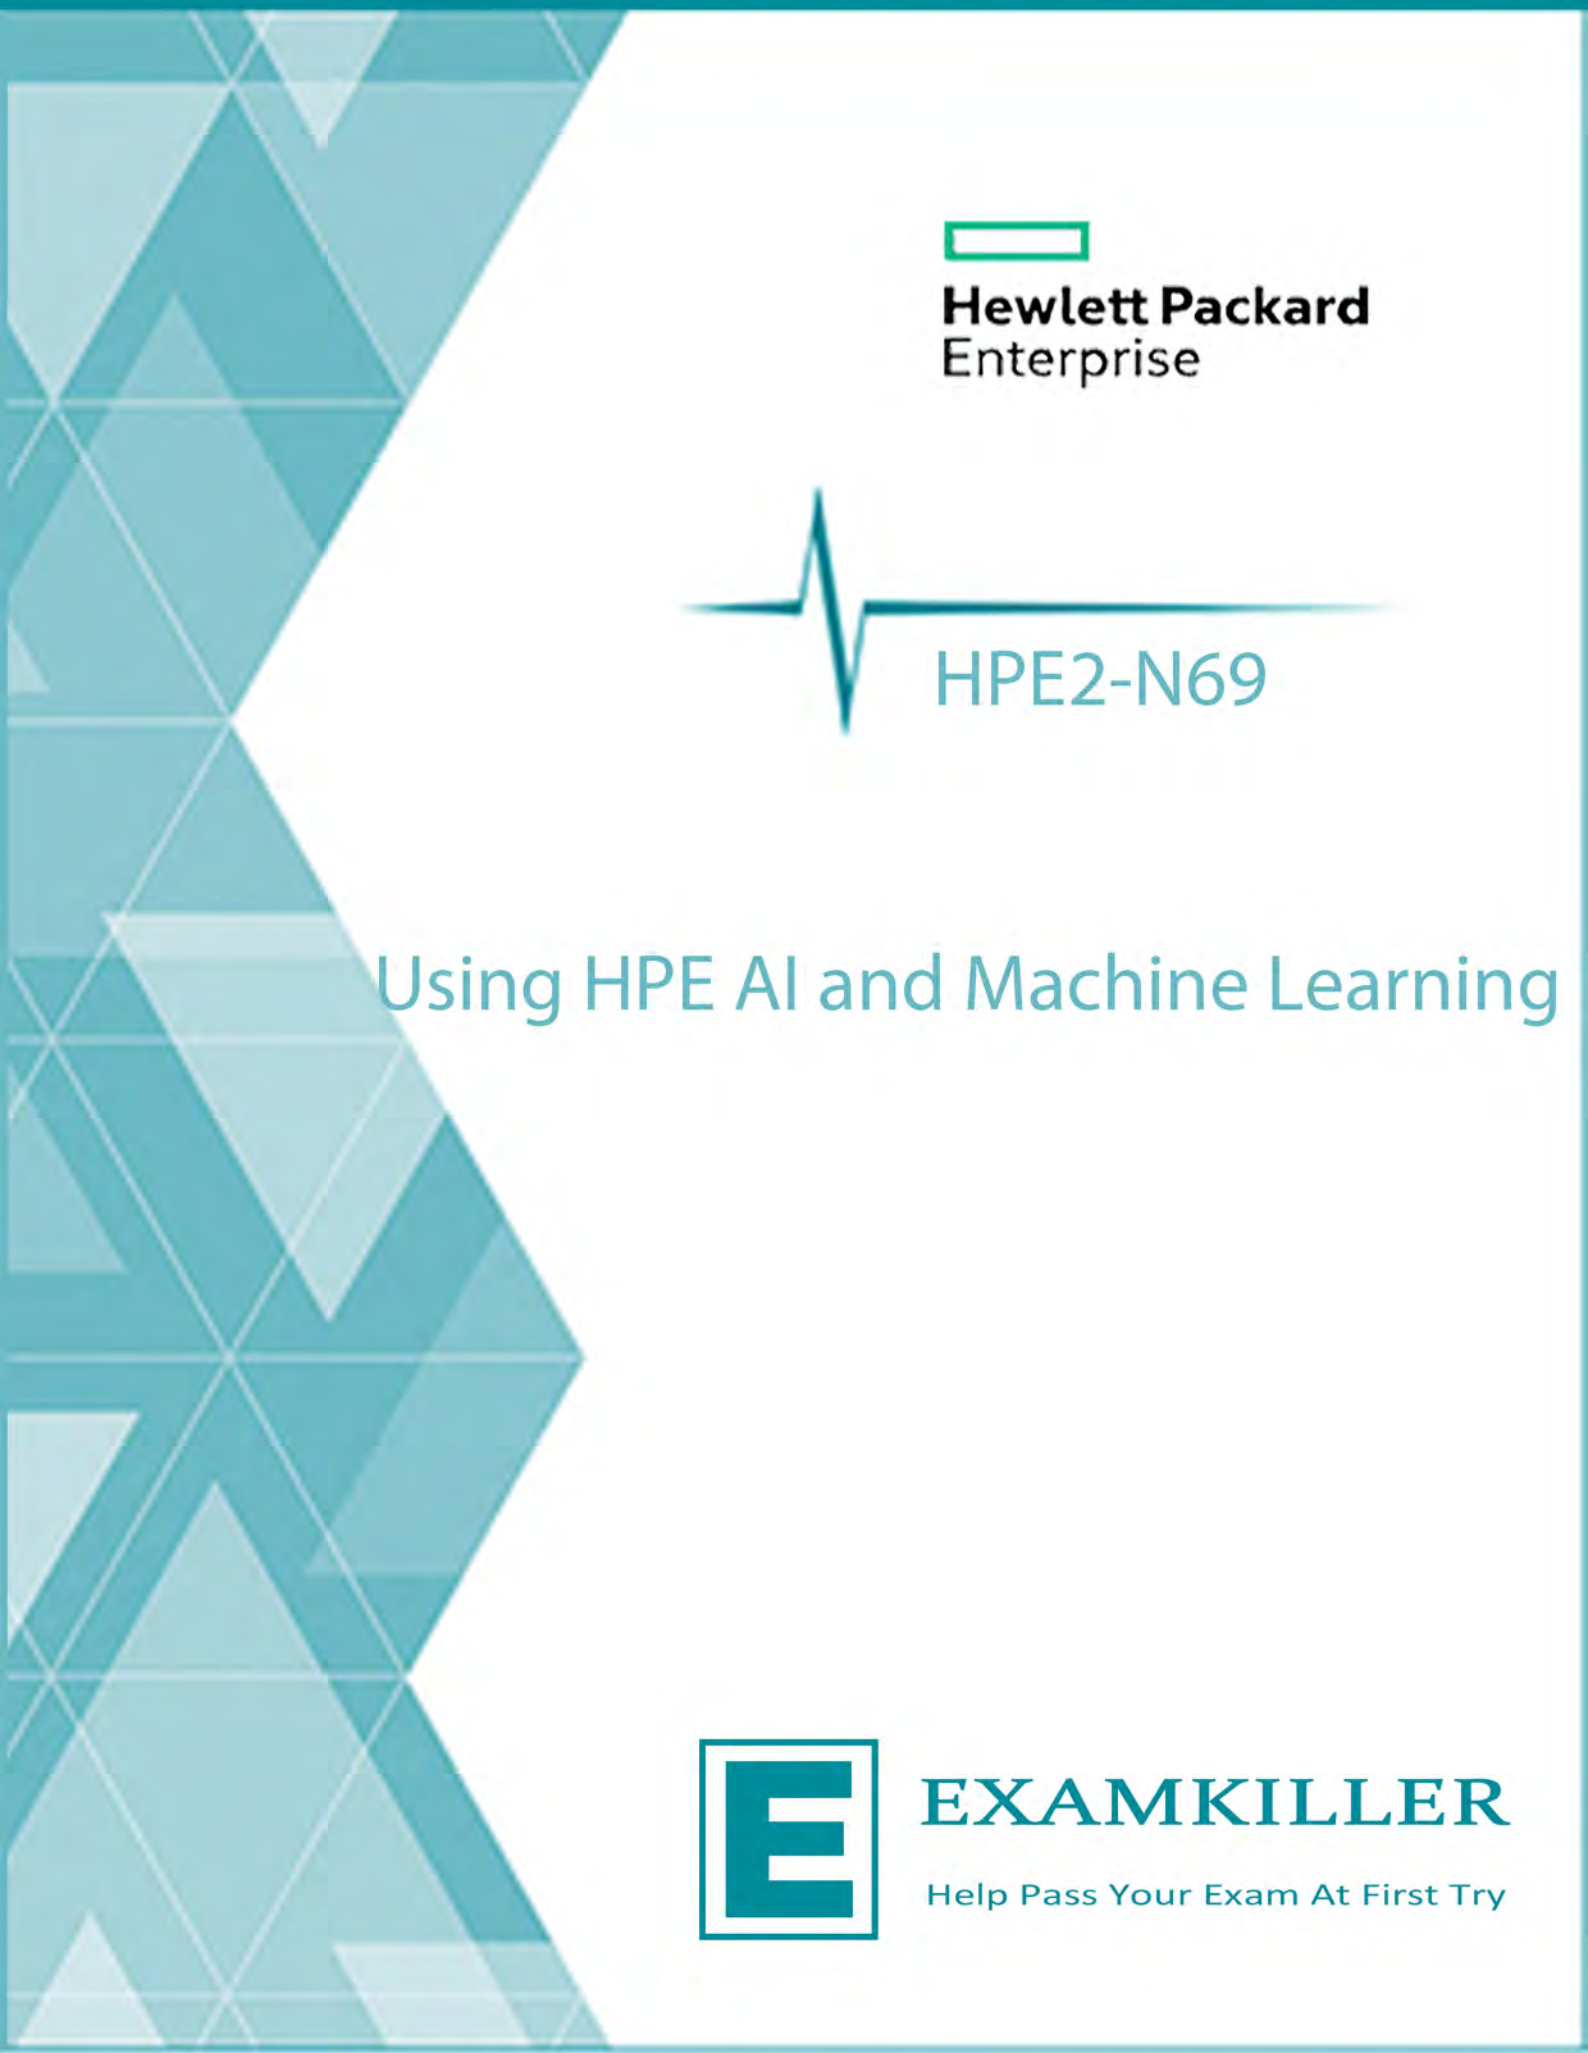

HPE2-N69

Using HPE AI and Machine Learning

EXAMKILLER

Help Pass Your Exam At First Try

HP

Exam HPE2-N69

Using HPE AI and Machine Learning

Version: 4.0

[Total Questions: 40]

Question No : 1

The ML engineer wants to run an Adaptive ASHA experiment with hundreds of trials. The engineer knows that several other experiments will be running on the same resource pool, and wants to avoid taking up too large a share of resources. What can the engineer do in the experiment config file to help support this goal?

A. Under "searcher," set "max_concurrent_trails" to cap the number of trials run at once by this experiment.
B. Under "searcher," set "divisor- to 2 to reduce the share of the resource slots that the experiment receives.
C. Set the "scheduling_unit" to cap the number of resource slots used at once by this experiment.
D. Under "resources.- set 'priority to I to reduce the share of the resource slots mat the experiment receives.

Answer: A

Explanation: The ML engineer can set "maxconcurrenttrials" under "searcher" in the experiment config file to cap the number of trials run at once by this experiment. This will help ensure that the experiment does not take up too large a share of resources, allowing other experiments to also run concurrently.

Question No : 2

Compared to Asynchronous Successive Halving Algorithm (ASHA), what is an advantage of Adaptive ASHA?

A. Adaptive ASHA can handle hyperparameters related to neural architecture while ASHA cannot.
B. ASHA selects hyperparameter configs entirely at random while Adaptive ASHA clones higher-performing configs.
C. Adaptive ASHA can train more trials in certain amount of time, as compared to ASHA.
D. Adaptive ASHA tries multiple exploration/exploitation tradeoffs oy running multiple Instances of ASHA.

Answer: B

Explanation: Adaptive ASHA is an enhanced version of ASHA that uses a reinforcement learning approach to select hyperparameter configurations. This allows Adaptive ASHA to select higher-performing configs and clone those configurations, allowing for better performance than ASHA.

Question No : 3

A company has an HPE Machine Learning Development Environment cluster. The ML engineers store training and validation data sets in Google Cloud Storage (GCS). What is an advantage of streaming the data during a trial, as opposed to downloading the data?

A. Streaming requires just one bucket, while downloading requires many.
B. The trial can more quickly start up and begin training the model.
C. The trial can better separate training and validation data.
D. Setting up streaming is easier that setting up downloading.

Answer: B

Explanation: Streaming the data during a trial allows the data to be processed more quickly, as it does not need to be downloaded onto the cluster before training can begin. This means that the trial can start up faster and the model can begin training more quickly.

Question No : 4

What common challenge do ML teams lace in implementing hyperparameter optimization (HPO)?

A. HPO is a joint ml and IT Ops effort, and engineers lack deep enough integration with the IT team.
B. They cannot implement HPO on TensorFlow models, so they must move their models to a new framework.
C. Implementing HPO manually can be time-consuming and demand a great deal of expertise.
D. ML teams struggle to find large enough data sets to make HPO feasible and worthwhile.

Answer: C

Explanation: Implementing hyperparameter optimization (HPO) manually can be time-consuming and demand a great deal of expertise. HPO is not a joint ML and IT Ops effort and it can be implemented on TensorFlow models, so these are not the primary challenges faced by ML teams. Additionally, ML teams often have access to large enough data sets to make HPO feasible and worthwhile.

Question No : 5

You want to open the conversation about HPE Machine Learning Development Environment with an IT contact at a customer. What can be a good discovery question?

A. How long does it currently take for a DL training to run the backward pass?
B. How much do you understand about building ML and DL models?
C. How much time do you spend managing the ML infrastructure?
D. What frustrations do you have with existing ML deployment and differencing solutions?

Answer: D

Explanation: A good discovery question to start a conversation about HPE Machine Learning Development Environment with an IT contact at a customer would be: "What frustrations do you have with existing ML deployment and differencing solutions?" By understanding the customer's current challenges and frustrations, you can better determine how HPE's ML Development Environment could help to address those needs.

Question No : 6

What type of interconnect does HPE Machine learning Development System use for high-speed, agent-to-agent communications?

A. Remote Direct Memory Access (RDMA) overconverged Ethernet (RoCE)
B. Slingshot
C. InfiniBand
D. Data Center Bridging (OCB)-enabled Ethernet

Answer: A

Explanation: HPE Machine Learning Development System uses Remote Direct Memory Access (RDMA) overconverged Ethernet (RoCE) for high-speed, agent-to-agent communications. This technology allows data to be transferred directly between agents without the need for copying, which results in improved performance and reduced latency.

Question No : 7

What is one of the responsibilities of the conductor of an HPE Machine Learning Development Environment cluster?

A. it downloads datasets for training.
B. It uploads model checkpoints.
C. It validates trained models.
D. It ensures experiment metadata is stored.

Answer: D

Explanation: The conductor of an HPE Machine Learning Development Environment

cluster is responsible for ensuring that all experiment metadata is stored and accessible. This includes tracking experiment runs, storing configuration parameters, and ensuring results are stored for future reference.

Question No : 8

What distinguishes deep learning (DL) from other forms of machine learning (ML)?

A. Models based on neural networks with interconnected layers of nodes, including multiple hidden layers
B. Models defined with Apache Spark rather than MapReduce
C. Models that are trained through unsupervised, rather than supervised, training
D. Models trained through multiple training processes implemented by different team members

Answer: A

Explanation: Models based on neural networks with interconnected layers of nodes, including multiple hidden layers. Deep learning (DL) is a type of machine learning (ML) that uses models based on neural networks with interconnected layers of nodes, including multiple hidden layers. This is what distinguishes it from other forms of ML, which typically use simpler models with fewer layers. The multiple layers of DL models enable them to learn complex patterns and features from the data, allowing for more accurate and powerful predictions.

Question No : 9

What is a benefit of HPE Machine Learning Development Environment, beyond open source Determined AI?

A. Automated user provisioning
B. Pipeline-based data management
C. Distributed training
D. Automated hyperparameter optimization (HPO)

Answer: D

Explanation: One of the main benefits of HPE Machine Learning Development Environment is its ability to automate the process of hyperparameter optimization (HPO). HPO is a process of automatically tuning the hyperparameters of a model during training, which can greatly improve a model's performance. HPE ML DE provides automated HPO,

making the process of tuning and optimizing the model much easier and more efficient.

Question No : 10

An HPE Machine Learning Development Environment resource pool uses priority scheduling with preemption disabled. Currently Experiment 1 Trial I is using 32 of the pool's 40 total slots; it has priority 42. Users then run two more experiments:

• Experiment 2:1 trial (Trial 2) that needs 24 slots; priority 50

• Experiment 3; I trial (Trial 3) that needs 24 slots; priority I

What happens?

A. Trial I is allowed to finish. Then Trial 3 is scheduled.
B. Trial 2 is scheduled on 8 of the slots. Then, alter Trial 1 has finished, it receives 16 more slots.
C. Trial 1 is allowed to finish. Then Trial 2 is scheduled.
D. Trial 3 is scheduled on 8 of the slots. Then, after Trial 1 has finished, it receives 16 more slots.

Answer: D

Explanation: Trial 3 is scheduled on 8 of the slots. Then, after Trial 1 has finished, it receives 16 more slots. This is because priority scheduling is used in the HPE Machine Learning Development Environment resource pool, which means higher priority tasks will be given priority over lower priority tasks. As such, Trial 3 with priority 1 will be given priority over Trial 2 with priority 50.

Question No : 11

A trial is running on a GPU slot within a resource pool on HPE Machine Learning Development Environment. That GPU fails. What happens next?

A. The trial tails, and the ML engineer must restart it manually by re-running the experiment.
B. The concluded reschedules the trial on another available GPU in the pool, and the trial restarts from the state of the latest training workload.
C. The conductor reschedules the trial on another available GPU in the pool, and the trial restarts from the latest checkpoint.
D. The trial fails, and the ML engineer must manually restart it from the latest checkpoint using the WebUI.