

HP.HPE2-N69.by.Polyn.31q

Number: HPE2-N69
Passing Score: 800
Time Limit: 120
File Version: 6.0

Exam Code: HPE2-N69



Exam A

QUESTION 1

What common challenge do ML teams face 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.

Correct Answer: A

Section:

QUESTION 2

The 10 agents in "my-compute-pool" have 8 GPUs each, you want to change an experiment config to run on multiple GPUs at once. What is a valid setting for "resources_per_trial"?

- A. 10
- B. 24
- C. 12
- D. 20

Correct Answer: A

Section:

QUESTION 3

You are meeting with a customer, and MUDL engineers express frustration about losing work flow to hardware failures. What should you explain about how HPE Machine Learning Development Environment addresses this pain point?

- A. The solution automatically mirrors the training process on redundant agents, which take over if an issue occurs.
- B. The solution continuously monitors agent hardware and sends out proactive alerts before failed hardware causes training to fail.
- C. The conductor and each of the agents are deployed in an active-standby model, which protects in case of hardware issues.
- D. The solution can take periodic checkpoints during the training process and automatically restart failed training from the latest checkpoint.

Correct Answer: A

Section:

QUESTION 4

At what FQDN (or IP address) do users access the WebUI for an HPE Machine Learning Development cluster?

- A. Any of the agent's in a compute pool
- B. A virtual one assigned to the cluster
- C. The conductor's
- D. Any of the agent's in an aux pool



Correct Answer: D

Section:

QUESTION 5

What is the role of a hidden layer in an artificial neural network (ANN)?

- A. It is responsible for passively reformatting data for use in the ANN.
- B. It is responsible for making the final decision about how to label a record, based on weighted input from preceding layers.
- C. It receives and weighs inputs from the preceding layer and produces outputs for the next layer.
- D. It does not play a role during the forward pass of data through the ANN, but it helps to optimize during the backward pass.

Correct Answer: D

Section:

QUESTION 6

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 than setting up downloading.

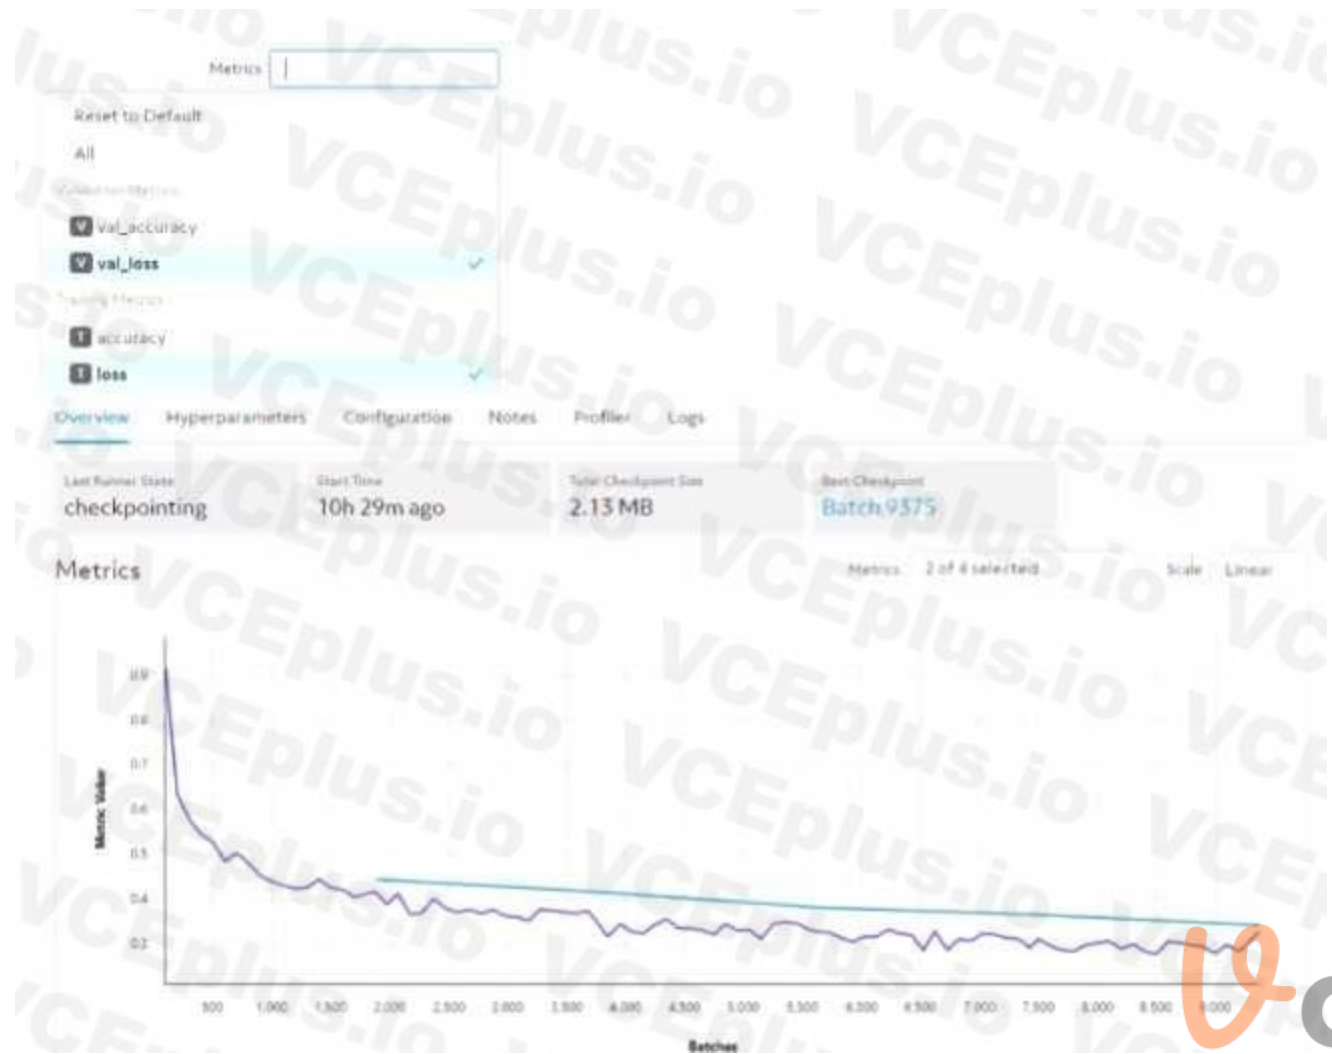
Correct Answer: C

Section:

QUESTION 7

Refer to the exhibit.





You are demonstrating HPE Machine Learning Development Environment, and you show details about an experiment, as shown in the exhibits. The customer asks about what "validation loss" means. What should you respond?

- A. Validation refers to testing how well the current model performs on new data; the lower the loss the better the performance.
- B. Validation refers to an assessment of how efficient the model code is; the lower the loss the lower the demand on GPU memory resources.
- C. Validation loss refers to the loss detected during the backward pass of training, while training loss refers to loss during the forward pass.
- D. Validation loss is metadata that indicates how many updates were lost between the conductor and agents.

Correct Answer: C

Section:

QUESTION 8

An ml engineer wants to train a model on HPE Machine Learning Development Environment without implementing hyper parameter optimization (HPO). What experiment config fields configure this behavior?

- A. profiling: enabled: false
- B. hyperparameters; optimizer:none
- C. searcher: name: single
- D. resources: slots_per_trial: 1

Correct Answer: D

Section:

QUESTION 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)

Correct Answer: B

Section:

QUESTION 10

A customer has Men expanding its deep learning (DO prefects and is confronting several challenges. Which of these challenges does HPE Machine Learning Development Environment specifically address?

- A. Time-consuming data collection
- B. Complex model deployment processes
- C. Complex and time-consuming data cleansing process
- D. Complex and time-consuming hyperparameter optimization (HPO)

Correct Answer: C

Section:

QUESTION 11

You want to set up a simple demo cluster for HPE Machine Learning Development Environment for the open source Determined all on a local machine. Which OS Is supported?

- A. HP-UX v11i
- B. Windows Server 2016 or above
- C. Windows 10 or above
- D. Red Hat 7-based Linux

Correct Answer: A

Section:

QUESTION 12

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 l to reduce the share of the resource slots mat the experiment receives.

Correct Answer: A

Section:

QUESTION 13

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

- A. Experiment tracking
- B. Model Inferencing
- C. Distributed training
- D. Premium dedicated support

Correct Answer: A

Section:

QUESTION 14

A customer is deploying HPE Machine learning Development Environment on on-prem infrastructure. The customer wants to run some experiments on servers with 8 NVIDIA A too GPUs and other experiments on servers with only Z

NVIDIA T4 GPUs. What should you recommend?

- A. Letting the conductor automatically determine which servers to use for each experiment, based on the number of resource slots required
- B. Deploying two HPE Machine Learning Development Environment clusters, one for each server type
- C. Deploying servers with 8 GPUs as agents and using the conductor to run experiments that require only 2 GPUs
- D. Establishing multiple compute resource pools on the cluster, one for servers of each type

Correct Answer: D

Section:

QUESTION 15

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 higherperforming 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 by running multiple Instances of ASHA.

Correct Answer: D

Section:

QUESTION 16

What is a benefit of HPE Machine Learning Development Environment that tends to resonate with executives?

- A. It uses a centralized training architecture that is highly efficient.
- B. It helps DL projects complete faster for a faster ROI.
- C. It helps companies deploy models and generate revenue.
- D. It automatically cleans up data to create better end results.

Correct Answer: B

Section:

QUESTION 17

Your cluster uses Amazon S3 to store checkpoints. You ran an experiment on an HPE Machine Learning Development Environment cluster, you want to find the location for the best checkpoint created during the experiment. What can you do?

- A. In the experiment config that you used, look for the "bucket" field under "hyperparameters." This is the UUID for checkpoints.
- B. Use the "det experiment download -top-n 1" command, referencing the experiment ID.
- C. In the Web UI, go to the Task page and click the checkpoint task that has the experiment ID.
- D. Look for a "determined-checkpoint/" bucket within Amazon S3, referencing your experiment ID.

Correct Answer: B

Section:

QUESTION 18

What is a reason to use the best fit policy on an HPE Machine Learning Development Environment resource pool?

- A. Ensuring that all experiments receive their fair share of resources
- B. Minimizing costs in a cloud environment
- C. Equally distributing utilization across multiple agents
- D. Ensuring that the highest priority experiments obtain access to more resources

Correct Answer: B

Section:

QUESTION 19

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.



Correct Answer: B

Section:

QUESTION 20

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

Correct Answer: D

Section:

QUESTION 21

An ML engineer is running experiments on HPE Machine Learning Development Environment. The engineer notices all of the checkpoints for a trial except one disappear after the trial ends. The engineer wants to Keep more of these checkpoints. What can you recommend?

- A. Adjusting how many of the latest and best checkpoints are saved in the experiment config's checkpoint storage settings.

- B. Monitoring ongoing trials In the WebUI and clicking checkpoint nags to auto-save the desired checkpoints.
- C. Double-checking that the checkpoint storage location is operating under 90% of total capacity.
- D. Adjusting the checkpoint storage settings to save checkpoints to a shared file system instead of cloud storage.

Correct Answer: D

Section:

QUESTION 22

You are proposing an HPE Machine Learning Development Environment solution for a customer. On what do you base the license count?

- A. The number of servers in the cluster
- B. The number of agent GPUs
- C. The number of processor cores on agents
- D. The number of processor cores on all servers in the cluster

Correct Answer: B

Section:

QUESTION 23

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?



Correct Answer: A

Section:

QUESTION 24

What role do HPE ProLiant DL325 servers play in HPE Machine Learning Development System?

- A. They run validation and checkpoint workloads.
- B. They run training workloads that do not require GPUs.
- C. They host management software such as the conductor and HPCM.
- D. They run non-distributed training workloads.

Correct Answer: C

Section:

QUESTION 25

What is one key target vertical (or HPE Machine Learning Development solutions)?

- A. Hospitality
- B. K-12education
- C. Retail
- D. Manufacturing

Correct Answer: D

Section:

QUESTION 26

You are meeting with a customer who has several DL models deployed. Out wants to expand the projects.

The ML/DL team is growing from 5 members to 7 members. To support the growing team, the customer has assigned 2 dedicated IT staff. The customer is trying to put together an on-prem GPU cluster with at least 14 CPUs.

What should you determine about this customer?

- A. The customer is not ready for an HPE Machine Learning Development solution, but you could recommend open-source Determined AI.
- B. The customer is not ready for an HPE Machine Learning Development solution. Out you could recommend an educational HPE Pointnext ASPS workshop.
- C. The customer is a key target for HPE Machine Learning Development Environment, but not HPE Machine Learning Development System.
- D. The customer is a key target for an HPE Machine Learning Development solution, and you should continue the discussion.

Correct Answer: A

Section:

QUESTION 27

Where does TensorFlow fit in the ML/DL Lifecycle?

- A. it helps engineers use a language like Python to code and train DL models.
- B. it provides pipelines to manage the complete lifecycle.
- C. It is primarily used to transport trained models to a deployment environment.
- D. It adds system and GPU monitoring to the training process.

Correct Answer: A

Section:

QUESTION 28

You want to set up a simple demo Ouster for HPE Machine Learning Development Environment for the open source Determined AI) on a local machine. You plan to use "del deploy" to set up the cluster. What software must be installed on the machine before you run that command?

- A. Kubernetes
- B. PyTorch
- C. Terraform
- D. Docker

Correct Answer: A

Section:

QUESTION 29

An HPE Machine Learning Development Environment cluster has this resource pool:

Name: pool 1

Location: On-prem

Agents: 2

Aux containers per agent: 100

Total slots: 0

Which type of workload can run in pool 1?



- A. Training
- B. GPU Jupyter Notebook
- C. Validation
- D. CPU-only Jupyter Notebook

Correct Answer: D

Section:

QUESTION 30

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 fails, and the ML engineer must restart it manually by re-running the experiment.
- B. The conductor 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.

Correct Answer: C

Section:

QUESTION 31

You are in a directory on your machine with your experiment config file and your model code. You enter this command: `det experiment create myfile.yaml` You receive this error: `det experiment create: error: the following arguments are required: model_def` What should you do?

- A. Re-enter the command with "-m" in which is the code filename.
- B. Make sure that the myfile.yaml file includes code for a PyTorchTrial or TFKerasTrial class.
- C. Re-enter the command with a period (.) at the end.
- D. Make sure that you have already logged into the cluster with the "det login" command.

Correct Answer: B

Section:

