


2023 Updated Verified Pass HPE2-N69 Exam - Real Questions & Answers [Q11-Q35]



2023 Updated Verified Pass HPE2-N69 Exam - Real Questions and Answers
Dumps Moneyack Guarantee - HPE2-N69 Dumps Approved Dumps

QUESTION 11

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?

- * Kubernetes
- * PyTorch
- * Terralorm
- * Docker

Before running the `del deploy`; command to set up the cluster, you must first install Docker on the machine. Docker is a containerization platform that is used to run applications in an isolated environment. It is necessary to have Docker installed before running the `del deploy`; command to set up the cluster for the open source Determined AI on a local machine.

QUESTION 12

What are the mechanics of how a model trains?

- * Decides which algorithm can best meet the use case for the application in question
- * Adjusts the model's parameter weights such that the model can better perform its tasks
- * Tests how accurately the model performs on a wide array of real world data
- * Detects data drift or concept drift that might compromise the ML model's performance

QUESTION 13

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?

- * Streaming requires just one bucket, while downloading requires many.
- * The trial can more quickly start up and begin training the model.
- * The trial can better separate training and validation data.
- * Setting up streaming is easier than setting up downloading.

QUESTION 14

ML engineers are defining a convolutional neural network (CNN) model but they are not sure how many filters to use in each convolutional layer. What can help them address this concern?

- * Using hyperparameter optimization (HPO)
- * Distributing the training across multiple CPUs
- * Using a variable learning rate
- * Training the model on multiple epochs

Hyperparameter optimization is a process of tuning the hyperparameters of a machine learning model, such as the number of filters in a convolutional neural network (CNN) model, to determine the best combination of hyperparameters that will result in the best model performance. HPO techniques are used to automatically find the optimal hyperparameter values, which can greatly increase the accuracy and performance of the model.

QUESTION 15

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?

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

Validation loss is a metric used to measure how well the model is performing on unseen data. It is calculated by taking the difference between the predicted values and the actual values. The lower the validation loss, the better the model's performance on new data.

QUESTION 16

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

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

QUESTION 17

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?

- * Under searcher, set max_concurrent_trails to cap the number of trials run at once by this

experiment.

- * Under `searcher`; set `divisor` to 2 to reduce the share of the resource slots that the experiment receives.
- * Set the `scheduling_unit`; to cap the number of resource slots used at once by this experiment.
- * Under `resources`.- set `priority` to I to reduce the share of the resource slots that the experiment receives.

QUESTION 18

You are meeting with a customer who has several DL models deployed. They want 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?

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

The customer's dedicated IT staff, the customer is ready to deploy an on-premise GPU cluster with at least 14 CPUs. The HPE Machine Learning Development Environment is a comprehensive solution that provides the tools and technologies required to develop, manage, and deploy ML models. It includes a distributed training framework, an orchestration layer, a powerful development environment, and an integrated MLOps platform. With this solution, the customer can expand their ML/DL projects and scale up their team.

QUESTION 19

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?

- * Streaming requires just one bucket, while downloading requires many.
- * The trial can better separate training and validation data.
- * Setting up streaming is easier than setting up downloading.
- * The trial can more quickly start up and begin training the model.

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 20

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?

- * In the experiment config that you used, look for the `bucket` field under `hyperparameters`; This is the UUID for checkpoints.
- * Use the `experiment download -top-n I` command, referencing the experiment ID.
- * In the Web UI, go to the Task page and click the checkpoint task that has the experiment ID.
- * Look for a `determined-checkpoint` bucket within Amazon S3, referencing your experiment ID.

HPE Machine Learning Development Environment uses Amazon S3 to store checkpoints. To find the location of the best checkpoint created during an experiment, you need to look for a `determined-checkpoint` bucket within Amazon S3, referencing your experiment ID. This bucket will contain all of the checkpoints that were created during the experiment.

QUESTION 21

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 GPUs.

What should you determine about this customer?

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

QUESTION 22

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?

- * Training
- * GPU Jupyter Notebook
- * Validation
- * CPU-only Jupyter Notebook

Pool 1 has two agents, each with 100 aux containers, and a total of 0 slots. This means that the cluster is configured to run CPU-only workloads, such as running a CPU-only Jupyter Notebook. Training, GPU Jupyter Notebook, and validation workloads cannot be run on this cluster due to the lack of GPU resources.

QUESTION 23

You want to set up a simple demo cluster for HPE Machine Learning Development Environment (or the open source Determined AI) on Amazon Web Services (AWS). You plan to use `hpe-dl-deploy` to set up the cluster. What is one prerequisite?

- * installing the NVIDIA Container Toolkit on your local machine
- * Manually creating the AWS EC2 instance with a PostgreSQL database
- * Recording the name of a valid AWS EC2 keypair
- * Adding Amazon Elastic Kubernetes Services (EKS) to your AWS account

In order to use the `hpe-dl-deploy` command to set up a cluster for HPE Machine Learning Development Environment (or the open source Determined AI) on Amazon Web Services (AWS), you will need to have a valid AWS EC2 keypair. The keypair

will authenticate your access to the cluster and allow you to securely access the cluster once it is set up.

QUESTION 24

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

- * it downloads datasets for training.
- * It uploads model checkpoints.
- * It validates trained models.
- * It ensures experiment metadata is stored.

QUESTION 25

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

- * Automated user provisioning
- * Pipeline-based data management
- * Distributed training
- * Automated hyperparameter optimization (HPO)

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 26

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`?

- * 10
- * 24
- * 12
- * 20

The valid setting for `resources_per_trial` for the 10 agents in `my-compute-pool` with 8 GPUs each would be 20, as this would be the total number of GPUs available across all 10 agents. This setting would allow the experiment config to run on multiple GPUs at once.

QUESTION 27

An HPE Machine Learning Development Environment resource pool uses priority scheduling with preemption disabled. Currently Experiment 1 Trial 1 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: 1 trial (Trial 3) that needs 24 slots; priority 1

What happens?

- * Trial 1 is allowed to finish. Then Trial 3 is scheduled.
- * Trial 2 is scheduled on 8 of the slots. Then, after Trial 1 has finished, it receives 16 more slots.
- * Trial 1 is allowed to finish. Then Trial 2 is scheduled.
- * Trial 3 is scheduled on 8 of the slots. Then, after Trial 1 has finished, it receives 16 more slots.

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 28

ML engineers are defining a convolutional neural network (CNN) model but they are not sure how many filters to use in each convolutional layer. What can help them address this concern?

- * Using hyperparameter optimization (HPO)
- * Distributing the training across multiple CPUs
- * Using a variable learning rate
- * Training the model on multiple epochs

QUESTION 29

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

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

The WebUI for an HPE Machine Learning Development cluster can be accessed at the FQDN or IP address of the conductor. The conductor is responsible for managing the cluster and providing access to the WebUI.

QUESTION 30

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 A100 GPUs and other experiments on servers with only 2 NVIDIA T4 GPUs. What should you recommend?

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

QUESTION 31

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

- * Experiment tracking
- * Model Inference
- * Distributed training
- * Premium dedicated support

The benefit of HPE Machine Learning Development Environment beyond open source Determined AI is Distributed Training. Distributed training allows multiple machines to train a single model in parallel, greatly increasing the speed and efficiency of the training process. HPE ML Development Environment provides tools and support for distributed training, allowing users to make the most of their resources and quickly train their models.

QUESTION 32

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

- * The number of servers in the cluster
- * The number of agent GPUs

- * The number of processor cores on agents
- * The number of processor cores on all servers in the cluster

The license count for the HPE Machine Learning Development Environment solution would be based on the number of processor cores on all servers in the cluster. This includes all servers in the cluster, regardless of whether they are running agents or not. Each processor core in the cluster requires a license and these licenses can be purchased in packs of 2, 4, 8, and 16.

QUESTION 33

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

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

QUESTION 34

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?

- * HP-UX v11i
- * Windows Server 2016 or above
- * Windows 10 or above
- * Red Hat 7-based Linux

QUESTION 35

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 I?

- * Training
- * GPU Jupyter Notebook
- * Validation
- * CPU-only Jupyter Notebook

Updated PDF (New 2023) Actual HP HPE2-N69 Exam Questions:

<https://www.examcollectionpass.com/HP/HPE2-N69-practice-exam-dumps.html>