

HP.HPE2-N69.v2023-07-19.q27

Exam Code:	HPE2-N69
Exam Name:	Using HPE Cray AI Development Environment
Certification Provider:	HP
Free Question Number:	27
Version:	v2023-07-19
# of views:	571
# of Questions views:	270
https://www.exam-tests.com/HPE2-N69-exam/HP.HPE2-N69.v2023-07-19.q27.html	

NEW QUESTION: 1

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 does not play a role during the forward pass of data through the ANN, but it helps to optimize during the backward pass.
- D. It receives and weighs inputs from the preceding layer and produces outputs for the next layer.

Answer: C (LEAVE A REPLY)

NEW QUESTION: 2

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

Answer: (SHOW ANSWER)

To train a model on HPE Machine Learning Development Environment without implementing hyper parameter optimization (HPO), you need to set the "optimizer" field to "none" in the hyperparameters section of the experiment config. This will instruct the ML engine to not use any hyperparameter optimization when training the model.

NEW QUESTION: 3

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

Answer: (SHOW ANSWER)

NEW QUESTION: 4

You are helping a customer start to implement hyper parameter optimization (HPO) with HPE Machine learning Development Environment. An ML engineer is putting together an experiment config file with the desired Adaptive A5HA settings. The engineer asks you questions, such as how many trials will be trained on the max length and what the min length for all trials will be.

What should you explain?

- A. The engineer should access the HPE Machine Learning Development online calculator and input the mode, max_trials, max_length, divisor, and max_runs.
- B. The engineer should run the "det preview-search" command, referencing the experiment config.
- C. The engineer should upload the experiment config to the HPE Machine Learning Development Environment WebUI and view the graph of the experiment plan.
- D. The engineer should run a preliminary experiment with one tenth the desired number of max trials, assess the results, and then run the full experiment.

Answer: (SHOW ANSWER)

NEW QUESTION: 5

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

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

Answer: D (LEAVE A REPLY)

One key target vertical for HPE Machine Learning Development solutions is Manufacturing. Manufacturing businesses are using machine learning to automate processes, reduce costs, and improve safety and quality control. HPE ML solutions provide the tools and technologies to help manufacturers develop and deploy ML models in their production environments, enabling them to optimize and automate their operations.

NEW QUESTION: 6

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

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

Answer: B ([LEAVE A REPLY](#))

NEW QUESTION: 7

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 trained through multiple training processes implemented by different team members
- C. Models defined with Apache Spark rather than MapReduce
- D. Models that are trained through unsupervised, rather than supervised, training

Answer: ([SHOW ANSWER](#))

NEW QUESTION: 8

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

Answer: B ([LEAVE A REPLY](#))

NEW QUESTION: 9

A customer has been expanding its deep learning (DL) projects 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)

Answer: D ([LEAVE A REPLY](#))

The HPE Machine Learning Development Environment specifically addresses Complex and time-consuming hyperparameter optimization (HPO). HPO is a process used to

identify the most effective set of hyperparameters for a given machine learning model. HPE's ML Development Environment provides a suite of tools that allow users to quickly and easily design and deploy deep learning models, as well as optimize their hyperparameters to get the best results.

NEW QUESTION: 10

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?

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

Answer: B (LEAVE A REPLY)

NEW QUESTION: 11

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.

Answer: C (LEAVE A REPLY)

HPE ProLiant DL325 servers play an important role in the HPE Machine Learning Development System. They are used to host the management software such as the Conductor and HPCM, and they also run non-distributed training workloads that do not require GPUs. They can also be used to run validation and checkpoint workloads.

NEW QUESTION: 12

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 trial fails, and the ML engineer must manually restart it from the latest checkpoint using the WebUI.

D. The conductor reschedules the trial on another available GPU in the pool, and the trial restarts from the latest checkpoint.

Answer: D (LEAVE A REPLY)

If a GPU fails during a trial running on a resource pool on HPE Machine Learning Development Environment, the conductor will reschedule the trial on another available GPU in the pool, and the trial will restart from the latest checkpoint. The trial will not fail, and the ML engineer will not have to manually restart it from the latest checkpoint using the WebUI.

NEW QUESTION: 13

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

Answer: D (LEAVE A REPLY)

The OS supported for setting up a simple demo cluster for HPE Machine Learning Development Environment for the open source Determined on a local machine is Red Hat 7-based Linux. Red Hat 7-based Linux is an open source operating system that is used extensively in enterprise applications. It provides a stable and secure platform for running applications and is suitable for use in a demo cluster.

NEW QUESTION: 14

You are helping a customer start to implement hyper parameter optimization (HPO) with HPE Machine Learning Development Environment. An ML engineer is putting together an experiment config file with the desired Adaptive A5HA settings. The engineer asks you questions, such as how many trials will be trained on the max length and what the min length for all trials will be.

What should you explain?

A. The engineer should run the "det preview-search" command, referencing the experiment config.

B. The engineer should access the HPE Machine Learning Development online calculator and input the mode, max_trials, max_length, divisor, and max_runs.

C. The engineer should upload the experiment config to the HPE Machine Learning Development Environment WebUI and view the graph of the experiment plan.

D. The engineer should run a preliminary experiment with one tenth the desired number of max trials, assess the results, and then run the full experiment.

Answer: B (LEAVE A REPLY)

The engineer should specify the number of trials to train on the max length and the minimum length for all trials in the experiment config file. For example, if the engineer wants to run 10 trials with a max length of 10, the config file should look something like this:

```
{  
"mode": "A5HA",  
"max_trials": 10,  
"max_length": 10,  
"min_length": 1,  
"divisor": 2,  
"max_runs": 1  
}
```

Once the config file is complete, the engineer should upload it to the HPE Machine Learning Development Environment WebUI and view the graph of the experiment plan. This will allow the engineer to see how the Adaptive A5HA settings will affect the experiment. After that, the engineer can run the experiment and assess the results.

NEW QUESTION: 15

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

Answer: (SHOW ANSWER)

NEW QUESTION: 16

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?

- A. Using hyperparameter optimization (HPO)
- B. Distributing the training across multiple CPUs
- C. Using a variable learning rate
- D. Training the model on multiple epochs

Answer: A (LEAVE A REPLY)

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.

Valid HPE2-N69 Dumps shared by BraindumpsPass.com for Helping Passing HPE2-N69 Exam! BraindumpsPass.com now offer the **newest HPE2-N69 exam dumps**, the BraindumpsPass.com HPE2-N69 exam **questions have been updated** and **answers have been corrected** get the **newest** BraindumpsPass.com HPE2-N69 dumps with Test Engine here: <https://www.braindumps.com/HP/HPE2-N69-practice-exam-dumps.html> (42 Q&As Dumps, **40%OFF Special Discount: Exam-Tests**)

NEW QUESTION: 17

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

Answer: D (LEAVE A REPLY)

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.

NEW QUESTION: 18

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?

- A. Trial 1 is allowed to finish. Then Trial 3 is scheduled.
- B. Trial 2 is scheduled on 8 of the slots. Then, after 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 (LEAVE A REPLY)

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.

NEW QUESTION: 19

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

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

Answer: B (LEAVE A REPLY)

NEW QUESTION: 20

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

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

Answer: B (LEAVE A REPLY)

NEW QUESTION: 21

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.** Adaptive ASHA tries multiple exploration/exploitation tradeoffs by running multiple Instances of ASHA.
- C.** ASHA selects hyperparameter configs entirely at random while Adaptive ASHA clones higher-performing configs.
- D.** Adaptive ASHA can train more trials in certain amount of time, as compared to ASHA.

Answer: B (LEAVE A REPLY)

NEW QUESTION: 22

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.

Answer: (SHOW ANSWER)

HPE Machine Learning Development Environment is designed to deliver results more quickly than traditional methods, allowing companies to get a return on their investment sooner and benefit from their DL projects faster. This tends to be a benefit that resonates with executives, as it can help them realize their goals more quickly and efficiently.

NEW QUESTION: 23

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

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

Answer: D (LEAVE A REPLY)

NEW QUESTION: 24

What are the mechanics of how a model trains?

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

Answer: D (LEAVE A REPLY)

NEW QUESTION: 25

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

Answer: D (LEAVE A REPLY)

The best fit policy on an HPE Machine Learning Development Environment resource pool ensures that the highest priority experiments obtain access to more resources, while still ensuring that all experiments receive their fair share. This allows you to make the most of your resources and prioritize the experiments that are most important to you.

NEW QUESTION: 26

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

Answer: D (LEAVE A REPLY)

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.

NEW QUESTION: 27

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 (LEAVE A REPLY)

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.

Valid HPE2-N69 Dumps shared by BraindumpsPass.com for Helping Passing HPE2-N69 Exam! BraindumpsPass.com now offer the **newest HPE2-N69 exam dumps**, the BraindumpsPass.com HPE2-N69 exam **questions have been updated** and **answers have been corrected** get the **newest** BraindumpsPass.com HPE2-N69 dumps with Test Engine here: <https://www.braindumpsPass.com/HP/HPE2-N69-practice-exam-dumps.html> (42 Q&As Dumps, **40%OFF Special Discount: Exam-Tests**)