[May 24, 2023 HPE2-N69 certification guide Q&A from Training Expert ValidBraindumps [Q10-Q33

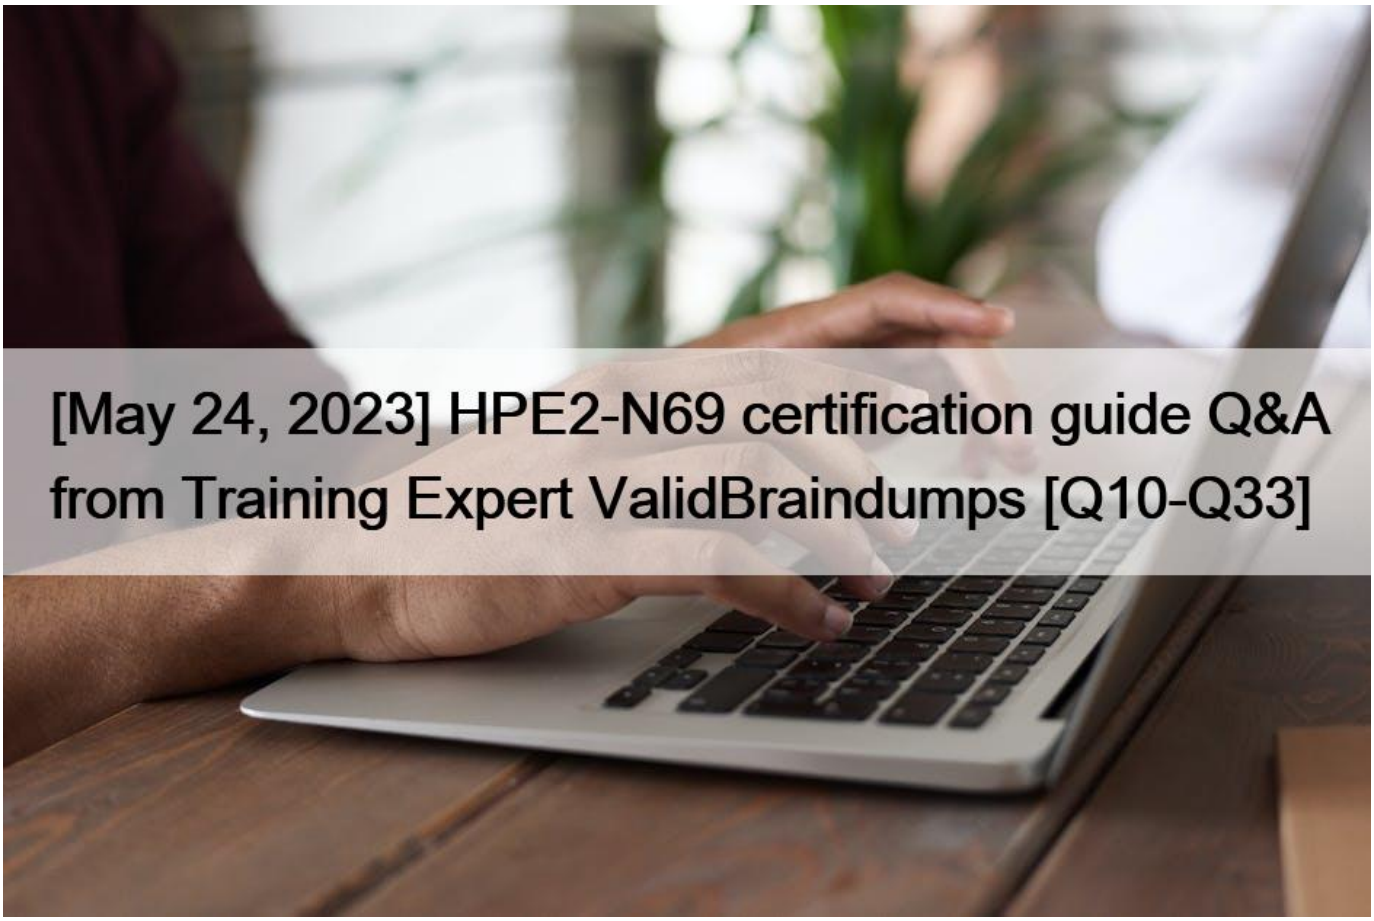

[May 24, 2023] HPE2-N69 certification guide Q&A from Training Expert ValidBraindumps
HPE2-N69 Certification Overview Latest HPE2-N69 PDF Dumps

The exam consists of 60 questions, which the candidate must answer in 90 minutes. The questions are multiple-choice and cover a range of topics related to HPE Cray AI Development Environment. The exam is available in English and Japanese, and the passing score is 70%. Candidates who pass the exam will receive the HPE ATP - AI Solutions V2 certification.

NEW QUESTION 10

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 tor 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 “det experiment download -top-n I” command, referencing the experiment ID.
* In the Web Ul, 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.

NEW QUESTION 11

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 mat the experiment receives.
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.

NEW QUESTION 12

What type of interconnect does HPE Machine learning Development System use for high-speed, agent-to-agent communications?
* Remote Direct Memory Access (RDMA) overconverged Ethernet (RoCE)
* Slingshot
* InfiniBand
* Data Center Bridging (OCB)-enabled Ethernet
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.

NEW QUESTION 13

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.

NEW QUESTION 14

ML engineers are defining a convolutional neural network (CNN) model bur 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 late
* Training the model on multiple epochs

NEW QUESTION 15

At what FQDN (or IP address) do users access the WebUI Tor 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.

NEW QUESTION 16

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

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

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 17

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

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

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.

NEW QUESTION 18

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

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

The best way to explain how HPE Machine Learning Development Environment addresses this pain point is to mention that the solution can take periodic checkpoints during the training process and automatically restart failed training from the latest checkpoint. This ensures that in case of a hardware failure, the engineers will not lose their work and training can be resumed from the last successful checkpoint.

NEW QUESTION 19

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

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 20

What is a benefit of HPE Machine Learning Development Environment mat tends to resonate with executives?
* It uses a centralized training architecture that is highly efficient.
* It helps DL projects complete faster for a faster ROI.
* It helps companies deploy models and generate revenue.
* It automatically cleans up data to create better end results.

NEW QUESTION 21

At what FQDN (or IP address) do users access the WebUI Tor 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

NEW QUESTION 22

What are the mechanics of now 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 of content drift that might compromise the ML model’s performance

This is done by running the model through a training loop, where the model is fed data and the parameter weights are adjusted based on the results of the model’s performance on the data. For example, if the model is a neural network, the weights of the connections between the neurons are adjusted based on the results of the model’s performance on the data. This process is repeated until the model performs better on the data, at which point the model is considered trained.

NEW QUESTION 23

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.

NEW QUESTION 24

What is a reason to use the best tit 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

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 25

A customer mentions that the ML team wants to avoid overfitting models. What does this mean?
* The team wants to avoid wasting resources on training models with poorly selected hyperparameters.
* The team wants to spend less time on creating the code tor models and more time training models.
* The team wants to avoid training models to the point where they perform less well on new data.
* The team wants to spend less time figuring out which CPUs are available for training models.

Overfitting occurs when a model is trained too closely on the training data, leading to a model that performs very well on the training data but poorly on new data. This is because the model has been trained too closely to the training data, and so cannot generalize the patterns it has learned to new data. To avoid overfitting, the ML team needs to ensure that their models are not overly trained on the training data and that they have enough generalization capacity to be able to perform well on new data.

NEW QUESTION 26

You are meeting with a customer how 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 start. 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 Al.
* 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.

NEW QUESTION 27

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.

NEW QUESTION 28

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?
* profiling: enabled: false
* hyperparameters; optimizer:none
* searcher: name: single
* resources: slots_per_trial: 1

NEW QUESTION 29

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

NEW QUESTION 30

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.

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 31

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 that setting up downloading.

NEW QUESTION 32

What is a reason to use the best tit 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

NEW QUESTION 33

The 10 agents in “my-compute-poor nave 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 “resourcespertrial” for the 10 agents in “my-compute-poor” 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.