


2023 HPE2-N69 Dumps PDF - HPE2-N69 Real Exam Questions Answers [Q13-Q37]



2023 HPE2-N69 Dumps PDF - HPE2-N69 Real Exam Questions Answers
Valid HPE2-N69 Test Answers & HP HPE2-N69 Exam PDF

NO.13 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.

NO.14 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.

NO.15 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.

NO.16 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?

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

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

The best recommendation for an ML engineer running experiments on HPE Machine Learning Development Environment to keep more of the checkpoints is to adjust the experiment configuration's checkpoint storage settings to save more of the latest and best checkpoints. This can be done by monitoring ongoing trials in the WebUI and clicking checkpoint flags to auto-save the desired checkpoints. Additionally, the engineer should double-check that the checkpoint storage location is operating under 90% of total capacity to ensure that enough capacity is available to store the checkpoints. Finally, they can adjust the checkpoint storage settings to save checkpoints to a shared file system instead of cloud storage if desired.

NO.17 What common challenge do ML teams face in implementing hyperparameter optimization (HPO)?

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

NO.18 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?

- * 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.

NO.19 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.

NO.20 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.

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.

NO.21 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 fails, and the ML engineer must restart it manually by re-running the experiment.
- * The conductor 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.

NO.22 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?

- * Adjusting how many of the latest and best checkpoints are saved in the experiment config's checkpoint storage settings.
- * Monitoring ongoing trials in the WebUI and clicking checkpoint nags to auto-save the desired checkpoints.
- * Double-checking that the checkpoint storage location is operating under 90% of total capacity.
- * Adjusting the checkpoint storage settings to save checkpoints to a shared file system instead of cloud storage.

NO.23 Where does TensorFlow fit in the ML/DL Lifecycle?

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

TensorFlow provides pipelines to manage the complete lifecycle of ML/DL models, from data ingestion to model training, evaluation, and deployment. It helps engineers use a language like Python to code and train DL models, and it also adds system and GPU monitoring to the training process. Additionally, it can be used to transport trained models to a deployment environment.

NO.24 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 fails, and the ML engineer must restart it manually by re-running the experiment.
- * The conductor reschedules the trial on another available GPU in the pool, and the trial restarts from the state of the latest training workload.
- * The trial fails, and the ML engineer must manually restart it from the latest checkpoint using the WebUI.
- * The conductor reschedules the trial on another available GPU in the pool, and the trial restarts from the latest checkpoint.

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.

NO.25 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 40

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.

NO.26 What is the role of a hidden layer in an artificial neural network (ANN)?

- * It is responsible for passively reformatting data for use in the ANN.

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

NO.27 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.

NO.28 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

NO.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

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.

NO.30 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 `det 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.

NO.31 What is one key target vertical (or HPE Machine Learning Development solutions)?

- * Hospitality
- * K-12 education
- * Retail
- * Manufacturing

NO.32 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.

NO.33 Compared to Asynchronous Successive Halving Algorithm (ASHA), what is an advantage of Adaptive ASHA?

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

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.

NO.34 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.

NO.35 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

NO.36 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; 1 trial (Trial 3) that needs 24 slots; priority I

What happens?

- * Trial I is allowed to finish. Then Trial 3 is scheduled.
- * Trial 2 is scheduled on 8 of the slots. Then, alter 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.

NO.37 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.

HPE2-N69 Exam Dumps - PDF Questions and Testing Engine: <https://www.vceprep.com/HPE2-N69-latest-vce-prep.html>