



**Free Questions for [HPE2-N69](#) by [go4braindumps](#)**

**Shared by [Gutierrez](#) on [22-07-2024](#)**

**For More Free Questions and Preparation Resources**

**[Check the Links on Last Page](#)**

# Question 1

---

## Question Type: MultipleChoice

---

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 "det deploy" to set up the cluster. What is one prerequisite?

### Options:

---

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

### Answer:

---

C

### Explanation:

---

In order to use the 'det 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 2

---

**Question Type:** MultipleChoice

---

A customer is using fair-share scheduling for an HPE Machine Learning Development Environment resource pool. What is one way that users can obtain relatively more resource slots for their important experiments?

### Options:

---

- A- Set the weight to a higher than default value.
- B- Set the weight to a lower than default value.
- C- Set the priority to a lower than default value.
- D- Set the priority to a higher than default value.

### Answer:

---

A

### **Explanation:**

---

Fair-share scheduling allocates resources to experiments based on the weight value of the resource pool. Increasing the weight value of a resource pool will result in more resource slots being allocated to it.

## **Question 3**

---

### **Question Type: MultipleChoice**

---

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?

### **Options:**

---

- 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

### **Explanation:**

---

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.

## Question 4

---

**Question Type: MultipleChoice**

---

What common challenge do ML teams face in implementing hyperparameter optimization (HPO)?

### Options:

---

- 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.

**Answer:**

---

C

**Explanation:**

---

Implementing hyperparameter optimization (HPO) manually can be time-consuming and demand a great deal of expertise. HPO is not a joint ML and IT Ops effort and it can be implemented on TensorFlow models, so these are not the primary challenges faced by ML teams. Additionally, ML teams often have access to large enough data sets to make HPO feasible and worthwhile.

## Question 5

---

**Question Type:** MultipleChoice

---

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

**Options:**

---

- 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.

**Answer:**

---

C

**Explanation:**

---

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.

## Question 6

---

**Question Type:** MultipleChoice

---



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?

### Options:

---

- 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.

### Answer:

---

A

### Explanation:

---

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 config'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.

## Question 7

---

**Question Type:** MultipleChoice

---

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?

### Options:

---

- 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.

**Answer:**

---

B

**Explanation:**

---

Make sure that the myfile.yaml file includes code for a PyTorchTrial or TFKerasTrial class. When creating an experiment with the det experiment create command, you need to specify the model\_def parameter to provide the code for the PyTorchTrial or TFKerasTrial class. This code should be specified in the myfile.yaml file, so make sure that the myfile.yaml file includes the code for the model you want to use.

## Question 8

---

**Question Type: MultipleChoice**

---

Where does TensorFlow fit in the ML/DL Lifecycle?

**Options:**

---

**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.

**Answer:**

---

B

**Explanation:**

---

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.

## Question 9

---

**Question Type:** MultipleChoice

---

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?

### Options:

---

- 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?

### Answer:

---

D

### Explanation:

---

A good discovery question to start a conversation about HPE Machine Learning Development Environment with an IT contact at a customer would be: 'What frustrations do you have with existing ML deployment and differencing solutions?' By understanding the customer's current challenges and frustrations, you can better determine how HPE's ML Development Environment could help to address those needs.

## Question 10

---

**Question Type:** MultipleChoice

---

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?

### Options:

---

- 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.

### Answer:

---

D

### Explanation:

---

The customer is a key target for an HPE Machine Learning Development solution, and you should continue the discussion. With 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 11

---

**Question Type:** MultipleChoice

---

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

### Options:

---

- 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

**Explanation:**

---

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 12

---

**Question Type: MultipleChoice**

---

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?

**Options:**

---

**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 l' 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.

**Answer:**

---

D

**Explanation:**

---

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.

**To Get Premium Files for HPE2-N69 Visit**

**<https://www.p2pexams.com/products/hpe2-n69>**

**For More Free Questions Visit**

**<https://www.p2pexams.com/hp/pdf/hpe2-n69>**

