

DUMPSQUEEN

Using HPE AI and Machine Learning

HP HPE2-N69

Version Demo

Total Demo Questions: 5

Total Premium Questions: 40

Buy Premium PDF

<https://dumpsqueen.com>

support@dumpsqueen.com

dumpsqueen.com

QUESTION NO: 1

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 A too GPUs and other experiments on servers with only Z NVIDIA T4 GPUs. What should you recommend?

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

ANSWER: D

Explanation:

By establishing multiple compute resource pools on the cluster, you can ensure that the correct servers are used for each experiment, depending on the number of GPUs required. This will help ensure that the experiments are run on the servers with the correct resources without having to manually assign each experiment to the appropriate server.

QUESTION NO: 2

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

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

ANSWER: D

Explanation:

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.

QUESTION NO: 3

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?

- 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 NO: 4

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

Explanation:

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 NO: 5

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?

- 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.