

IT-DUMPS Q&A

Accurate study guides, High passing rate!
IT-dumps provides update free of charge in one year!

Exam : **HPE2-N69**

Title : Using HPE AI and Machine Learning

Version : DEMO

1.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

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

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

Answer: A

3.At what FQDN (or IP address) do users access the WebUI Tor an HPE Machine Learning Development cluster?

- A. Any of the agent's in a compute pool
- B. A virtual one assigned to the cluster
- C. The conductor's
- D. Any of the agent's in an aux pool

Answer: C

4.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. 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 I" 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: B

5.A customer mentions that the ML team wants to avoid overfitting models.

What does this mean?

- A. The team wants to avoid wasting resources on training models with poorly selected hyperparameters.
- B. The team wants to spend less time on creating the code tor models and more time training models.
- C. The team wants to avoid training models to the point where they perform less well on new data.
- D. The team wants to spend less time figuring out which CPUs are available for training models.

Answer: C