

HP HPE0-S62

HPE Compute Architect

**Thank You for Downloading HPE0-S62 Updated Exam
Questions**

<https://certs4sale.com/hp/hpe0-s62-pdf-exam-dumps>

Question 1. (Single Select)

A customer's HPE Synergy estate has grown, and the operations team has started applying firmware and platform settings to compute modules one at a time. The customer asks the architect how the design is supposed to keep a group of modules consistent.

What should the architect describe?

- A: Re-stream a new operating-system image to each compute module in the group
- B: Edit each server profile individually and record the change in a configuration register
- C: Change the server profile template and let the derived profiles follow
- D: Replace the compute modules with matching hardware

Correct Answer: C

Question 2. (Single Select)

An architect is completing the design for a customer whose installation is scheduled for the following quarter.

Why does the architect gather addressing, naming, credential, and upstream network details as part of the design rather than leaving them to the delivery team?

- A: They are customer facts nobody can discover on site during the install window
- B: They determine the physical rack layout and the routing of power feeds into the cabinet
- C: They are needed to confirm that the configuration the architect proposed is orderable
- D: They form the acceptance criteria the customer signs against once the platform is running

Correct Answer: A

Question 3. (Single Select)

A customer asks an architect to simply state how many accelerators they need for HPE Private Cloud AI. The architect wants to reach a defensible configuration size rather than a component count.

What should the architect establish first?

- A: The rack space, power and cooling capacity the customer's facility can make available today
- B: The number of data scientists, analysts, and application teams who will be given accounts, and the order in which they are onboarded
- C: The volume of data currently held in the source systems, the archives, and the file shares that would be connected to the platform
- D: The workload profile: training or inference, model size, concurrency and latency target

Correct Answer: D

Question 4. (Single Select)

A customer's data scientists say they can produce an accurate model in a notebook within weeks, but that months of argument with the platform team follow before anyone outside the data science group can use it. Nobody disputes the model's accuracy.

Which stage of the AI workload lifecycle does this pain sit in?

- A: Data preparation -- the inputs are not being assembled in a repeatable way.
- B: Deployment and serving -- moving a trained model into a form others can consume.
- C: Training -- the available compute is not letting the model converge quickly enough.
- D: Monitoring -- nobody is observing how the model behaves once it is answering real requests.

Correct Answer: B

Question 5. (Single Select)

An application team currently serves each model behind its own hand-built inference service. Every new model has meant new serving code, new tuning work and a new interface for the consuming application.

Which benefit of NVIDIA NIMs directly addresses that pattern?

A: Every model is retrained against the customer's own data before the application is allowed to call it

B: Every model is chosen automatically for the application according to the prompt it receives

C: Every model is fronted by the same standard API, so the application does not change per model

D: Every model arrives with a serving stack, so the team tunes it to its own hardware and traffic

Correct Answer: C

Question 6. (Single Select)

Two customers describe the same application serving a comparable user population. One loses trading revenue for every minute the application is unavailable; the other can run the process manually for a working day.

Which requirement does that difference establish for the architect?

A: The processor and memory capacity each estate will need to be sized to

B: The management layer each customer's operations team should be asked to adopt

C: The physical location in which each customer's equipment must be installed

D: The level of redundancy and recovery the design must carry

Correct Answer: D

Question 7. (Single Select)

A customer new to composable infrastructure is being walked through the parts of an HPE Synergy frame and asks what the Composer contributes.

What is HPE Synergy's Composer for?

- A: Streaming operating-system images and boot volumes to compute modules from a central repository
- B: Linking several frames together and holding them in one managed ring, so they form a single management domain
- C: Composing and managing server profiles for the compute modules
- D: Carrying the data-center network into the frame

Correct Answer: C

Question 8. (Single Select)

A mid-sized customer wants an internal assistant over their own policy documents. They employ two application developers and no ML engineers, and they have decided not to hire. Their security team requires that the documents never leave the customer's own facility.

Which positioning follows from this assessment?

- A: A hosted service consumed over the public internet by the two developers.
- B: An on-premises platform the customer consumes rather than assembles.
- C: A components-based build the two developers integrate themselves over the coming year.
- D: An HPC cluster with a batch scheduler through which the developers submit jobs.

Correct Answer: B

Question 9. (Single Select)

During a design review a customer reports that their fine-tuning jobs spend long periods waiting for training data to reach the accelerators, while utilization on the accelerated servers stays low.

Which layer of HPE Private Cloud AI is this design conversation about?

- A: The storage and data path, which has to feed the accelerators fast enough to keep them busy
- B: The accelerated compute layer, which should be moved to a larger configuration size
- C: The management plane, which drives updates and support insight from the HPE GreenLake cloud
- D: The software layer, whose scheduling decides which project's jobs run at any moment

Correct Answer: A

Question 10. (Single Select)

A customer re-purposes the same HPE Synergy compute modules between a seasonal analytics run and a test estate several times a year. Each turn requires a module to come up under a different operating system, and the customer will not accept a rebuild from installation media each time.

Which component should the architect present as delivering this?

- A: Composer, reapplying a revised server profile template to the module
- B: Image Streamer, delivering a different image to the module
- C: Frame Link Module, widening the management domain across both frames
- D: Composable storage modules, reallocating direct-attached disk to the module

Correct Answer: B

THANK YOU FOR DOWNLOADING HPE0-S62 UPDATED EXAM QUESTIONS

Note: Thanks For Trying The Demo Of Our HPE0-S62 Exam Product

Visit Our Site to Purchase the Full Set of Actual HPE0-S62 Exam Questions With Answers.

Money Back Guarantee



Click The Link Below

<https://certs4sale.com/hp/hpe0-s62-pdf-exam-dumps>