

HPE Edge-to-Cloud Solutions

HP HPE0-V27

Version Demo

Total Demo Questions: 20

Total Premium Questions: 204

Buy Premium PDF

<https://passqueen.com>

support@passqueen.com

Topic Break Down

Topic	No. of Questions
Topic 1, HPE Edge-to-Cloud Strategy	51
Topic 2, HPE GreenLake Cloud Platform	25
Topic 3, HPE Edge-to-Cloud Solutions	128
Total	204

QUESTION NO: 1

Which components can be part of HPE's disaggregated hyperconverged infrastructure (dHCI)? (Choose three.)

- A. HPE B-series switches
- B. HPE Alletra 6000 array with FC HBAs
- C. HPE ProLiant DX
- D. HPE Alletra 6000 array with iSCSI HBAs
- E. HPE Aruba switches
- F. HPE ProLiant DL

ANSWER: D E F

Explanation:

HPE Alletra 6000 array with iSCSI HBAs, HPE Aruba switches, and HPE ProLiant DL are the components that align with the HPE Alletra dHCI architecture. HPE dHCI deliberately separates compute from storage while integrating their lifecycle operations through VMware vCenter and HPE management tooling. The compute layer uses supported HPE ProLiant DL rack servers, which provide independently scalable CPU and memory resources for virtualized workloads. The shared-storage layer is supplied by an HPE Alletra 6000 system. Alletra 6000 is based on the HPE Nimble Storage architecture and provides block storage through iSCSI connectivity, allowing storage capacity and performance to scale independently of the server layer. Networking is provided through supported Ethernet switching, including HPE Aruba switches, to carry the iSCSI storage traffic and associated management connectivity. Together, these elements deliver HCI-like deployment and management simplicity without tying compute and storage expansion to fixed node pairs. HPE describes this as a disaggregated design in which customers can independently scale compute and storage while retaining a validated, integrated solution. See the [HPE Alletra dHCI product information](#) and HPE's [dHCI overview](#) for the architecture and its independently scalable compute, storage, and networking layers.

QUESTION NO: 2

Which HPE Aruba Networking capabilities support a zero trust approach for users and IoT devices connecting to a campus network? (Select two.)

- A. Identity-based access policies for network users and devices
- B. Device profiling to classify connected endpoints
- C. Unrestricted access for every device connected to the internal network
- D. A single shared credential for all employees and IoT devices
- E. Disabling authentication for devices at remote sites

ANSWER: A B

Explanation:

Identity-based access control is a core zero trust capability because access decisions should be based on verified user or device identity and assigned role, rather than assuming that a connection from inside the network is trusted. Device profiling also supports zero trust by helping IT identify connected IoT and unmanaged devices so that appropriate policies can be applied according to device type and behavior.

HPE Aruba Networking ClearPass provides network access control capabilities including authentication, profiling, and policy enforcement. These functions help organizations apply differentiated access policies for employees, guests, and IoT devices across wired and wireless networks. See [HPE Aruba Networking ClearPass](#) and [NIST Zero Trust Architecture](#).

QUESTION NO: 3

Your client needs a departmental Storage array to host a VDI workload at a remote office. The networking infrastructure is limited and the client has decided to connect the ESXi host servers with 12Gbps SAS.

Which HPE Storage product will meet their requirements?

- A. HPE Alletra 9060
- B. HPE Nimble AF20
- C. HPE MSA 2062
- D. HPE Alletra 6030

ANSWER: C

Explanation:

HPE MSA 2062 is the appropriate choice because it is an entry-level hybrid-flash storage array intended for cost-conscious small-to-midsize organizations, departmental deployments, and remote-office environments. Crucially for this design, the MSA 2062 platform is available with 12 Gb SAS host connectivity, enabling direct SAS attachment from supported ESXi hosts when an Ethernet or Fibre Channel SAN is not available. This meets the stated infrastructure constraint while providing shared storage suitable for a virtual-desktop deployment.

The MSA 2062 also includes all-flash capacity in the base configuration and supports automated tiering, helping it provide responsive performance for the boot, login, and steady-state I/O patterns commonly associated with VDI. Its dual-controller architecture, VMware integration, snapshots, and remote replication capabilities give a departmental implementation practical availability and data-protection features without requiring an enterprise-scale storage network. HPE positions the MSA family specifically for affordable shared storage and documents SAS as a supported host-interface choice for applicable MSA 2060-series configurations. Product-family details are available from [HPE MSA Storage](#), and platform specifications are documented in the [HPE MSA 2060 Storage QuickSpecs](#).

QUESTION NO: 4

Your customer has requested a proof of concept.

What should be included in your agreement with the customer? (Choose two.)

- A. production bill of materials
- B. success criteria
- C. estimated timeline
- D. production site survey
- E. 24/7 support

ANSWER: B C

Explanation:

“success criteria” and “estimated timeline” should be documented in a proof-of-concept agreement because they establish a shared, measurable basis for conducting and evaluating the engagement. Success criteria translate the customer’s intended business and technical outcomes into objective acceptance measures—for example, validating a required workload capability, performance target, manageability outcome, or operational benefit. Agreeing on these measures before implementation prevents ambiguity over whether the proof of concept achieved its purpose and provides a clear decision point for moving forward.

An estimated timeline defines the proof of concept’s expected duration, key activities, milestones, validation period, customer responsibilities, and review or closeout date. A time-bounded plan keeps the evaluation focused, ensures necessary resources are available when needed, and sets expectations for feedback and the final outcome. Together, these items align

the customer and solution provider on what will be tested, how success will be assessed, and when the evaluation will conclude. This follows HPE's consultative approach of identifying desired outcomes and validating a solution against agreed requirements before broader deployment. See HPE's [proof-of-concept overview](#) and its [consulting services approach](#).

QUESTION NO: 5

A retailer wants to improve the digital experience in its stores by connecting employee devices, guest devices, and IoT systems. Which HPE Aruba Networking capabilities should be included in the discussion? (Select two.)

- A. Secure wired and wireless connectivity
- B. Identity-based access control
- C. Replacing all local networking with unmanaged switches
- D. Providing identical unrestricted access to every device
- E. Eliminating network visibility for branch IT teams

ANSWER: A B

Explanation:

Secure wired and wireless connectivity should be included because reliable access is the foundation for employee mobility, guest services, point-of-sale systems, and connected operational devices. A modern network must provide consistent performance while supporting growing numbers of users and endpoints.

Identity-based access control should also be included because employees, guests, and IoT devices require different levels of network access. Applying policy according to identity and device role helps protect business systems while enabling the digital services the retailer wants to provide. See [HPE Aruba Networking](#) and [HPE Aruba Networking ClearPass](#).

QUESTION NO: 6

Your customer needs an 8:1 GPU to CPU ratio to support their natural language processing (NLP) application.

Which HPE platform meets their requirement?

- A. Apollo 6500 Gen10 Plus
- B. Synergy SY480 Gen10 Plus with 8 Bay PCIe Expansion
- C. ProLiant DL380 Gen10 Plus
- D. EdgeLine e910 Server Blades

ANSWER: A

Explanation:

Apollo 6500 Gen10 Plus is the appropriate platform because it is HPE's purpose-built, high-density GPU system for AI, deep learning, and high-performance accelerated-computing workloads. Its architecture supports configurations of up to eight double-wide GPU accelerators in a single 4U system, providing the required eight-GPU accelerated-computing density for an NLP deployment. This density is important for NLP training and inference because transformer-based language models benefit from substantial parallel GPU compute capacity and high-throughput GPU interconnect and I/O resources. The platform is designed to pair GPU-rich configurations with enterprise server management, memory, storage, and networking capabilities, allowing the customer to build a consolidated AI node rather than distributing the workload across lower-density servers. HPE's Apollo 6500 Gen10 Plus QuickSpecs identify support for up to eight GPUs and position the system for deep learning and AI workloads. NVIDIA also documents NLP as a core deep-learning use case for GPU-accelerated frameworks and models. See the [HPE Apollo 6500 Gen10 Plus System QuickSpecs](#) and [NVIDIA NeMo NLP documentation](#).

QUESTION NO: 7

Which of the following is necessary when selecting the appropriate HPE and 3rd party products and services for a traditional solution?

- A. Designing the solution based on customer requirements
- B. Validating the customer requirements
- C. Architecting the solution
- D. Documenting customer intent

ANSWER: B

Explanation:

Validating the customer requirements is necessary because product and service selection must be based on confirmed, complete requirements rather than assumptions. For a traditional solution, validation establishes the customer's actual business outcomes, workloads, capacity and performance needs, availability targets, physical-site constraints, compatibility dependencies, budget, and support expectations. Those validated inputs allow the solution professional to choose appropriately sized HPE infrastructure, software, support services, and any required third-party components. It also provides a defensible basis for configuration, quotation, and the resulting statement of work, helping ensure that the proposed solution can be delivered and supported as intended. HPE's solution approach emphasizes understanding customer requirements and matching technology to required outcomes before finalizing a solution configuration. HPE describes its portfolio and services as being designed around customer needs and outcomes; see [HPE Solutions](#) and [HPE Consulting Services](#). Requirement validation is therefore the essential activity that makes an appropriate selection possible.

QUESTION NO: 8 - (SIMULATION)

Match the solution from left to right the benefits.

ANSWER: See the explanation for the answer

Explanation:

Match the solutions to these benefits:

HPE InfoSight → **AI-driven predictive analytics and performance optimization** (identifies performance issues/opportunities from workload patterns and provides recommendations).

HPE GreenLake Central → **Cost control and cloud-spend optimization** (tracks consumption, budgets, invoices, and cloud costs).

HPE Data Services Cloud Console → **Intent-based provisioning** (provisions data services according to workload type, capacity, and access requirements).

QUESTION NO: 9

When assessing the potential impact of an HPE Edge-to-Cloud solution on an organization's existing infrastructure, which of the following is a key metric to collect and analyze?

- A. Customer churn rate
- B. Number of product returns
- C. Mean time between failures
- D. Employee satisfaction score

ANSWER: C

Explanation:

Mean time between failures is correct because it is a core infrastructure reliability metric that helps establish the current environment's operational baseline before an Edge-to-Cloud solution is introduced. MTBF measures the average operating time between repairable failures. Collecting it across relevant servers, storage, network components, and supporting systems helps an organization quantify service reliability, identify infrastructure areas with recurring disruptions, and estimate the operational benefit that modernization, automation, predictive analytics, or managed cloud services may deliver.

For an Edge-to-Cloud assessment, understanding reliability is important alongside capacity, performance, availability, and operational-management data. A documented baseline enables the customer and solution team to compare conditions before and after implementation, assess risk to dependent workloads, and plan maintenance or resiliency improvements based on evidence. HPE's Edge-to-Cloud platform is designed to provide a cloud operating experience across distributed environments, where observability and operational insights support informed infrastructure decisions. HPE InfoSight also uses predictive analytics and operational telemetry to help identify and prevent infrastructure issues. See [HPE GreenLake cloud platform](#) and [HPE InfoSight](#).

QUESTION NO: 10

A customer is defining success criteria for an HPE GreenLake consumption-based solution. Which factors should be included when planning the capacity and consumption model? (Select two.)

- A. Expected workload growth and usage patterns
- B. Service-level and performance requirements
- C. The color standard used for the data center racks
- D. The number of vendors invited to the initial proposal meeting
- E. The age of the organization's desktop productivity software

ANSWER: A B**Explanation:**

Expected workload growth and usage patterns must be considered because capacity needs can change as application demand, data volumes, and user populations evolve. Understanding normal utilization, peak demand, and growth expectations helps establish an appropriate capacity plan.

Service-level and performance requirements must also be considered. Requirements for availability, response time, recovery, and workload performance influence the architecture, the amount of capacity required, and the operational model. A consumption-based solution should align technology and costs to actual business demand while continuing to meet required service outcomes. See [HPE GreenLake](#) and [HPE: What is cloud computing?](#).

QUESTION NO: 11

A customer has already virtualized much of the data center, but some workloads remain on non-virtualized servers. The customer does not have a strategy for moving to cloud. However, the customer is interested in moving to an IT as a Service (ITaaS) approach in the data center.

Which HPE solution should you target for this customer?

- A. a high performance computing (HPO solution with bundled management services
- B. an HPE hybrid cloud that integrates with leading third-party clouds
- C. an enterprise platform for a Hadoop big data ecosystem
- D. HPE Hybrid IT software for automating Infrastructure provisioning

ANSWER: D

Explanation:

HPE Hybrid IT software for automating Infrastructure provisioning is the appropriate target because the customer's immediate objective is to operate its existing, mixed data-center environment as an internal service. An ITaaS model depends on standardized infrastructure, policy-based automation, repeatable provisioning workflows, and a unified operational view. HPE infrastructure-management software, including HPE OneView capabilities, is designed to automate the lifecycle of physical compute, storage, and fabric resources through template-based infrastructure definitions and API-driven operations. This lets IT deliver resources more quickly and consistently while reducing manual configuration effort.

This approach is especially relevant where virtualized workloads coexist with workloads that must remain on physical servers. Rather than requiring an immediate public-cloud migration, automated Hybrid IT management modernizes the on-premises operating model first. It establishes the operational foundation for self-service-like delivery, governance, capacity control, and service-oriented IT processes, while preserving support for traditional infrastructure. The customer can later extend this foundation toward cloud or hybrid-cloud services when its strategy and workload requirements are defined. HPE describes OneView as infrastructure-management software that automates compute, storage, and networking management; see [HPE OneView](#). HPE's broader edge-to-cloud platform direction is described at [HPE GreenLake cloud platform](#).

QUESTION NO: 12

When evaluating an edge-to-cloud solution for a customer with critical applications, which availability and resiliency requirements should be documented? (Select two.)

- A. Recovery time objective
- B. Recovery point objective
- C. The preferred logo placement on user dashboards
- D. The number of project-status meetings per month
- E. The age of the oldest endpoint in the organization

ANSWER: A B**Explanation:**

Recovery time objectives should be documented because they define how quickly an application or service must be restored after a disruption. This requirement influences the recovery architecture, operational procedures, and level of supporting infrastructure needed.

Recovery point objectives should also be documented because they define the maximum acceptable amount of data loss, measured in time. Together, RTO and RPO provide measurable requirements for data protection, replication, backup, and disaster-recovery design. See [HPE Data Protection](#) and [HPE: What is disaster recovery?](#).

QUESTION NO: 13

What information should you gather during a physical site survey? (Choose two.)

- A. existing VLANs
- B. datacenter dimensions
- C. current patch levels
- D. existing disk capacity
- E. fire suppression capabilities

ANSWER: B E**Explanation:**

A physical site survey establishes whether the intended environment can safely and practically support the proposed infrastructure. Gathering **datacenter dimensions** provides the measurements needed to validate equipment placement, rack layouts, aisle clearances, delivery paths, floor loading considerations, and space for future growth. Accurate dimensions also support cooling-layout planning, including hot- and cold-aisle containment and the location of power and network infrastructure. Gathering **fire suppression capabilities** identifies the facility's installed fire-detection and suppression arrangements, such as pre-action sprinklers or clean-agent systems. This is necessary to assess environmental risk, coordinate equipment and rack placement with safety systems, and ensure that the deployment aligns with the site's operational and compliance requirements. HPE's data-center planning guidance treats facilities characteristics—including physical space, power, cooling, and environmental safeguards—as core inputs to solution design. Similarly, HPE Aruba site-survey guidance distinguishes physical inspection of the site and its constraints from later logical assessment of the existing network configuration. These two items therefore belong in the physical-survey discovery record used before finalizing the implementation design. See [HPE Data Center Facilities Design Guide](#) and [HPE Aruba Central site-survey documentation](#).

QUESTION NO: 14

Your customer wants to implement a software defined storage solution. They require the lowest latency possible in the back-end. Which technology should you recommend to address the customer ' s concern?

- A. HPE Aruba CX Ethernet switch family
- B. HPE M-series Ethernet switch family
- C. HPE B-series FC switch family
- D. HPE C-series FC switch family

ANSWER: B

Explanation:

HPE M-series Ethernet switch family is the appropriate recommendation because it is designed for performance-sensitive data-center, AI, HPC, and storage fabrics in which extremely low, predictable latency is essential. HPE M-series switches, based on NVIDIA networking technology, support high-speed Ethernet and capabilities used by modern scale-out software-defined storage designs, including lossless Ethernet features and RDMA over Converged Ethernet (RoCE). RoCE enables direct memory-to-memory data movement across the Ethernet fabric, reducing host CPU overhead and helping storage nodes exchange replication, rebuild, and client data with very low delay. This is particularly valuable when the storage back-end carries east-west traffic between clustered nodes, where network latency directly affects application and storage response times. The M-series portfolio is engineered for high throughput, low latency, and congestion-aware data-center fabrics, making it a strong fit when minimizing back-end latency is the customer's explicit priority. HPE describes the portfolio and its high-performance networking use cases at [HPE Networking M-series](#). Additional HPE storage-networking guidance is available at [HPE Storage Networking](#).

QUESTION NO: 15

What is one benefit of HPE Solution Sales Enablement Tool (SSET)?

- A. It helps you to quickly set up meetings with HPE GreenLake and HPE Financial Services personnel.
- B. It helps you identify which customers will be good candidates for HPE GreenLake.
- C. HPE SSET is web-based and can be accessed over the network from a supported web browser.
- D. HPE SSET is a Windows-based application that does not require internet connectivity.

ANSWER: D

Explanation:

HPE SSET is a Windows-based application that does not require internet connectivity. This is a practical benefit for solution sellers and presales teams because it lets them work with the tool locally when network access is unavailable, restricted, or

unreliable. SSET supports the creation of solution configurations and related sales outputs, so offline availability helps maintain productivity during customer visits, travel, or work in secured environments where external connectivity may not be permitted. The locally installed application model also provides a consistent working environment rather than requiring users to depend on browser compatibility and continuous access to a web service. When connectivity is available, users can still obtain applicable product and content updates through the normal SSET update process, but an active internet connection is not a prerequisite for using the installed tool itself. This characteristic aligns with HPE's positioning of SSET as a sales-enablement application used to develop solution proposals and configurations. See HPE's [compute and solution portfolio information](#) and the [HPE Support Center](#) for HPE software documentation and update resources.

QUESTION NO: 16

Which traditional use cases can be handled by HPE GreenLake for Block? (Choose two.)

- A. Containerized AI workload utilizing Microsoft Azure and an Apollo 4000.
- B. Highly secure, air-gapped NFS share from an MSA 2062.
- C. Collocated VMFS datastore mounted from an Alletra 9080 and protected by Zerto.
- D. Mainframe volumes provided by HPE XP8 storage.
- E. On-premises mission critical Oracle database backed by an Alletra 6050.

ANSWER: C E

Explanation:

HPE GreenLake for Block is designed to deliver enterprise block storage as a service for conventional, performance-sensitive workloads that run in customer data centers. "Collocated VMFS datastore mounted from an Alletra 9080 and protected by Zerto" is an appropriate use case because VMware environments commonly consume block volumes as VMFS datastores, while an HPE Alletra 9000-class platform supplies the availability and performance expected for consolidated virtual-machine workloads. Data protection through Zerto complements this architecture by providing continuous data protection and disaster-recovery workflows for virtualized applications.

"On-premises mission critical Oracle database backed by an Alletra 6050" is also appropriate. Oracle databases depend on predictable latency, IOPS, resilient storage, and enterprise data services. HPE Alletra 6000 storage is intended for business-critical applications and provides the block-storage capabilities needed for database deployments, including availability-oriented architecture and nondisruptive operations. GreenLake consumption adds a cloud operating experience and flexible capacity model while allowing these traditional applications and their data to remain on premises. HPE describes GreenLake for Block Storage as a service built for mission-critical workloads, and its Alletra portfolio specifically positions the 6000 and 9000 families for business-critical and mission-critical block-storage needs. See [HPE GreenLake for Block Storage](#) and [HPE Alletra storage](#).

QUESTION NO: 17

You are generating a customer HPE GreenLake proposal for a customer.

Select the items that are mandatory when submitting the initial proposal to HPE for quoting.

(Choose two.)

- A. Start Bill of Materials
- B. End Bill of Materials
- C. Signed statement of Work
- D. Completed Order Checklist
- E. Credit Check Form

ANSWER: A D

Explanation:

For an initial HPE GreenLake proposal submission, **Start Bill of Materials** and **Completed Order Checklist** provide the baseline information HPE needs to create an accurate quote and begin the commercial review. The Start Bill of Materials identifies the proposed solution configuration, including the relevant infrastructure and software components, capacity quantities, and GreenLake service parameters. This establishes the starting consumption environment on which pricing and the proposed service are based.

The Completed Order Checklist supplies the required order and customer administration details so that HPE can validate the request, identify the appropriate contracting and billing entities, and confirm that the quotation can proceed through the GreenLake ordering workflow. Together, these documents connect the technical configuration to the commercial and operational information needed for a quote.

This aligns with HPE GreenLake's consumption-service model, in which the solution is sized for an initial installed capacity and delivered under a defined service and commercial arrangement. HPE's overview of the model is available at [HPE GreenLake](#), and the HPE GreenLake platform documentation is available through [HPE Support Center](#).

QUESTION NO: 18

What is a sign that a business is ready to modernize their IT infrastructure?

- A. Executives accept that modernizing the infrastructure will take several years to complete.
- B. The company wants to enhance customers' experience and gather data to uncover insights about customers.
- C. Executives understand the need to automate IT in order to implement new technologies.
- D. IT managers want to maintain control of the IT infrastructure by ensuring that traditional IT processes remain intact.

ANSWER: C

Explanation:

Executives understand the need to automate IT in order to implement new technologies is the clearest sign of organizational readiness for infrastructure modernization. Modernization is not only a hardware-refresh initiative; it requires the business to adopt more agile operating practices that reduce manual effort, standardize repeatable work, and enable IT teams to provision and manage resources consistently. Executive recognition of automation's role indicates that leadership appreciates the operational changes needed to support cloud-like experiences, hybrid environments, and rapid delivery of new digital services.

Automation is a foundational modernization capability because it supports infrastructure-as-code approaches, policy-based operations, streamlined lifecycle management, and more efficient use of IT staff. With leadership support, the organization can align technology investments, processes, and skills around measurable business outcomes rather than treating modernization as an isolated technical project. HPE describes IT modernization as transforming IT environments and operations to improve agility, efficiency, and innovation, while HPE GreenLake emphasizes a cloud operating model across edge, data center, and cloud environments. These principles make executive commitment to IT automation a strong readiness indicator. See [HPE: What is IT modernization?](#) and [HPE GreenLake cloud platform](#).

QUESTION NO: 19 - (SIMULATION)

Match each solution to the appropriate customer.



ANSWER: See the explanation for the answer

Explanation:

Match the solutions as follows:

Customers looking to unlock agility and collapse data management silos → Alletra

Customers looking to run apps and collapse infrastructure management silos → SimpliVity

Customers looking to run apps and build a resilient enterprise → Zerto

Alletra provides cloud-managed data services that reduce data-management silos. SimpliVity is HPE hyperconverged infrastructure, simplifying infrastructure operations for applications. Zerto provides continuous data protection and disaster recovery to improve enterprise resilience.

QUESTION NO: 20

Which key benefit of HPE Synergy helps to distinguish HPE from the competition?

- A. the ability to integrate IoT devices into the network with minimal security risks
- B. the ability to dynamically deploy location-based service apps to customer mobile devices
- C. the ability to support virtualized workloads from leading vendors such as VMware
- D. the ability to independently scale compute and storage resources, and redefine them dynamically

ANSWER: D

Explanation:

HPE Synergy distinguishes itself through a composable-infrastructure approach that lets organizations treat physical compute, storage, and fabric capacity as fluid resources rather than as permanently fixed server configurations. The ability to independently scale compute and storage resources, and redefine them dynamically, is central to this model. For example, an administrator can compose the appropriate infrastructure profile for a workload, allocate compute capacity separately from storage capacity, and then recompose those resources when application requirements change. This reduces overprovisioning, improves utilization, and enables faster deployment of traditional, virtualized, and cloud-native workloads.

HPE Synergy combines pooled infrastructure with software-defined management through HPE OneView. Templates and profiles enable infrastructure to be configured consistently through an API-driven process, helping IT teams deploy and adapt environments at cloud-like speed while retaining the performance and control of on-premises infrastructure. The HPE Synergy platform overview describes this as composable infrastructure designed to accelerate application delivery and operational agility. See [HPE Synergy](#) and [HPE OneView](#).