TestkingPDF

Try before you buy

Download a free sample of any of our exam questions and answers

Download Demo

- 24/7 customer support, Secure shopping site
- Free One year updates to match real exam scenarios
- If you failed your exam after buying our products we will refund the full amount back to you.

Quality and Value

TestkingPDF Practice Exams are written to the highest standards of technical accuracy, using only certified subject matter experts and published authors for development - no all dumps.

Tested and Approved

We are committed to the process of vendor and third party approvals. We believe professionals and executives alike deserve the confidence of quality coverage these authorizations provide.

Easy to Pass

If you prepare for the exams using our TestkingPDF testing engine, It is easy to succeed for all certifications in the first attempt. You don't have to deal with all dumps or any free torrent / rapidshare all stuff.

Try Before Buy

TestkingPDF offers free demo of each product. You can check out the interface, question quality and usability of our practice exams before you decide to buy.

http://www.testkingpdf.com

Valid test online & stable pass king & latest PDF dumps

Exam : **HPE0-V27**

Title : HPE Edge-to-Cloud Solutions

Vendor : HP

Version : DEMO

QUESTION NO: 1
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:
The HPE Apollo 6500 Gen10 Plus is a high-performance computing platform that supports up to 16 GPUs and 2 CPUs in a 4U chassis, providing an 8:1 GPU to CPU ratio. This platform is ideal for NLP applications that require massive parallel processing and high-speed data transfer between GPUs and CPUs. The HPE Apollo 6500 Gen10 Plus supports Nvidia L40 GPUs, which are optimized for NLP workloads and deliver up to 10x faster performance than previous generations. The HPE Apollo 6500 Gen10 Plus also offers flexible configuration options, advanced cooling and power efficiency, and comprehensive management and security features. References: HPE Apollo 6500 Gen10 Plus System , HPE Apollo 6500 Gen10 Plus Data sheet
, NVIDIA Teams With HPE to Take AI From Edge to Cloud

QUESTION NO: 2
What is the correct method to calculate incremental cash flow?
A. Substract projected decrease in revenue or increase in costs from the initial investment.
B. Divide the increase in revenue by the time needed to offset the IT investment.
C. Multiply the increase in revenue by the time needed to offset the IT investment.
D. Substract projected increase in revenue or decrease in costs from the initial investment.
Answer: D
Explanation:
Incremental cash flow is the difference between the cash flow of a project with the investment and the cash flow of the same project without the investment. It represents the net change in cash flow that results from making the investment. The correct method to calculate incremental cash flow is to substract projected increase in revenue or decrease in costs from the initial investment. This method captures the additional cash inflows and outflows that are attributable to the investment, and excludes any cash flows that are unrelated to the investment. Incremental cash flow can help you evaluate the profitability and feasibility of an IT investment by comparing it to the required rate of return or the payback period.
References: HPE Edge-to- Cloud Solutions - Self-Directed Lab , HPE Edge-to-Cloud Solution s - Official Certification Study Guide
, Incremental Cash Flow Definition , How to Calculate Incremental Cash Flow

QUESTION NO: 3
A customer has installed a new Fibre Channel SAN using a pair of 32Gbps B-Series switches. They have asked you to integrate their existing Synergy frame into this SAN. The

Synergy frame has a pair of Virtual Connect SE 100Gb F32 Modules configured for Ethernet only. There are enough unpopulated OSFP28 ports for you to be able to provide the 32 Gbps FC transceivers needed to facilitate the uplinks. What else needs to be included to facilitate the requirement?

A. Fibre Channel upgrade licenses

B. B-Series SANnav Management Software

C. Aruba Fabric Composer

D. HPE Compute Ops Management - Oneview Edition

Answer: A

Explanation:

The HPE Synergy Virtual Connect SE 100Gb F32 Module is a versatile interconnect that supports both Ethernet and Fibre Channel (FC) traffic. However, while the hardware ports are physically capable of supporting FC transceivers and traffic, the Fibre Channel functionality must be logically enabled via software licensing. If a Synergy frame was originally deployed in an " Ethernet only " configuration, it will not have the necessary licenses to process storage traffic. Therefore, to integrate the frame into a new B-Series FC SAN, the architect must include the appropriate Fibre Channel upgrade licenses for the Virtual Connect modules. These licenses activate the N_Port ID Virtualization (NPIV) features required to present the Synergy compute modules to the FC fabric. Without these licenses, the FC transceivers will not be initialized by the module even if they are physically installed.

References: HPE Virtual Connect SE 100Gb F32 Module QuickSpecs; HPE Synergy 12000 Frame Setup and Installation Guide.

QUESTION NO: 4

Your customer needs to deploy compute infrastructure in a harsh environment with inlet temperatures exceeding ASHRAE Class4, up to 55 degrees C.

Which HPE compute offering should you recommend?

A. HPE Apollo 6500 Gen10 Plus

B. HPE ProLiant DL380 Den10 Plus

C. HPE Edgeline 8000

D. HPE ProLiant DL380 Gen10 Plus Carrier Grade

Answer: C

Explanation:

HPE Edgeline 8000 is a ruggedized and modular edge system that delivers high performance, availability, and scalability for compute-intensive workloads in harsh environments. HPE Edgeline 8000 supports up to 16 blades with Intel Xeon Scalable processors, up to 1.5TB of memory per blade, and up to 48TB of storage per system. HPE Edgeline 8000 also offers integrated networking, security, and management features, as well as optional GPU and FPGA acceleration. HPE Edgeline 8000 can operate in temperatures ranging from -10°C to 55°C, exceeding the ASHRAE Class4 standard 1 2 . References: HPE Edgeline 8000 | HPE Store US
, HPE Edgeline 8000 - Data Sheet

QUESTION NO: 5

You are designing and sizing an HPE SMB multi-site hyper-converged infrastructure solution.

What tool can you use to aid in your design?

A. Solutions Wizard for SMB

B. SimpliVity Sizing Tool

C. CloudPhysics

D. NinjaSTARS

Answer: D

Explanation:

NinjaSTARS is a versatile sizing and design tool within the HPE portfolio that is essential for architecting hyper-converged infrastructure (HCI) solutions like HPE SimpliVity. When designing a multi-site solution for Small and Medium Business (SMB) customers, NinjaSTARS allows the architect to model the workload requirements across different locations. It provides critical insights into the expected data efficiency (deduplication and compression ratios), the number of nodes required per site to meet performance and availability goals, and the resulting hardware Bill of Materials (BOM). NinjaSTARS also validates the solution against HPE best practices, ensuring that the designed cluster has sufficient resources for high availability and workload failover between sites. While CloudPhysics is excellent for assessing current environments, NinjaSTARS is the dedicated tool for sizing the future HCI environment and ensuring it meets the specific hardware requirements for a SimpliVity deployment. References: HPE NinjaSTARS User Guide; HPE SimpliVity Sizing and Design.

QUESTION NO: 6

For each option, identify whether it is customizable for HPE GreenLake for Block Storage.

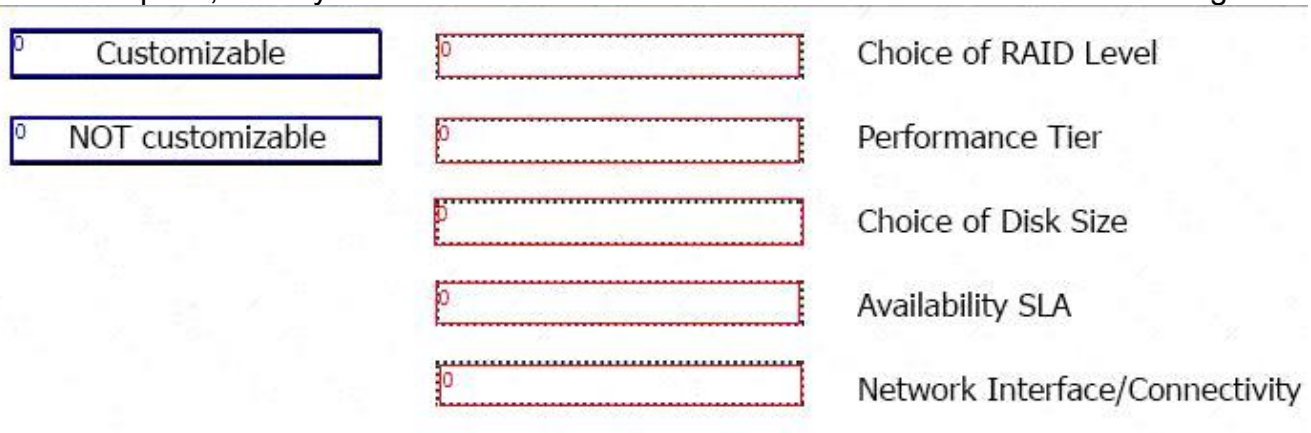

Answer:

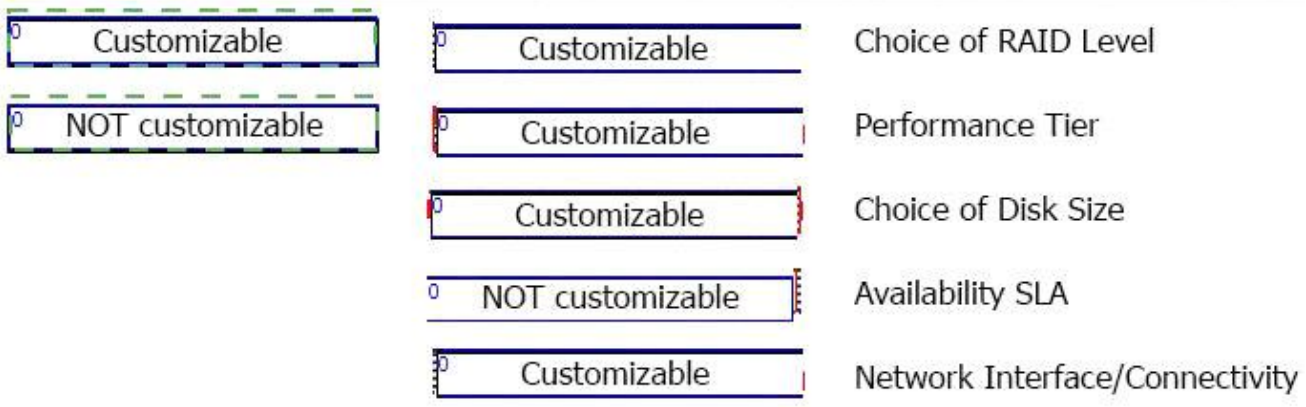

Explanation:
According to the HPE GreenLake for Block Storage configuration overview , the following options are customizable for HPE GreenLake for Block Storage:
* Choice of RAID Level: Customizable. You can choose the RAID level for each virtual pool of storage based on your performance and availability requirements. The supported RAID levels are RAID 1, RAID 5, RAID 6, and RAID 10 1 .
* Performance Tier: Customizable. You can choose the performance tier for each virtual pool of storage based on your workload characteristics and service level objectives. The available performance tiers are Performance, Business Critical, and Mission Critical 1 .
* Choice of Disk Size: Customizable. You can choose the disk size for each virtual pool of storage based on your capacity and density needs. The supported disk sizes are 1.92 TB, 3.84 TB, 7.68 TB, and 15.36 TB for SSDs, and 1.6 TB and 3.2 TB for SCM 1 .
* Availability SLA: NOT customizable. HPE GreenLake for Block Storage offers a built-in 100% data availability guarantee for mission-critical environments. This means that HPE will ensure that your data is always accessible and protected, and will compensate you for any downtime or data loss that may occur 2 .
* Network Interface/Connectivity: Customizable. You can choose the network interface and connectivity for each storage system based on your host and network requirements. The supported network interfaces are Fibre Channel and iSCSI, and the supported connectivity options are direct attach, switch attach, and SAN 1 .

QUESTION NO: 7
A customer wants to implement a solution that doesn't limit any future operational expenditures.
Which type of solution should you choose and why?
A. A hybrid solution because moving workloads into the cloud always reduce total cost.
B. A hybrid solution as the expenditure model is simple.
C. A Cloud solution with a pay-per-use model with on-demand scaling of resources.
D. A traditional solution works best because it takes no capital expenditure to deploy.
Answer: C
Explanation:
* A cloud solution with a pay-per-use model allows the customer to pay only for the resources they consume, such as compute, storage, network, or software services 1 .
* A cloud solution with a pay-per-use model also enables the customer to scale up or down their resources on-demand, depending on their workload and performance needs 1 .
* A cloud solution with a pay-per-use model can help the customer reduce their operational expenditures (OPEX) by eliminating the need for upfront capital investments, maintenance costs, and overprovisioning of resources 1 .
* A hybrid solution, which combines cloud and on-premises resources, may not always reduce the total cost of ownership, as it depends on the workload characteristics, the cloud pricing model, and the integration and management complexity 2 .
* A hybrid solution may also have a more complicated expenditure model, as it involves both OPEX and CAPEX, and requires careful planning and optimization of the resource allocation and utilization 2 .
* A traditional solution, which relies on on-premises hardware and software, requires a high

capital expenditure (CAPEX) to deploy, as well as ongoing maintenance, upgrade, and support costs 3 .
* A traditional solution also limits the customer's flexibility and agility, as they have to deal with fixed and finite resources, and longer provisioning and deployment cycles 3 .
The references are:
1 : Cloud Computing Pricing Models 2 : Hybrid Cloud Cost Optimization 3 : Cloud vs On-Premise Software Comparison

QUESTION NO: 8

Refer to the following exhibit.

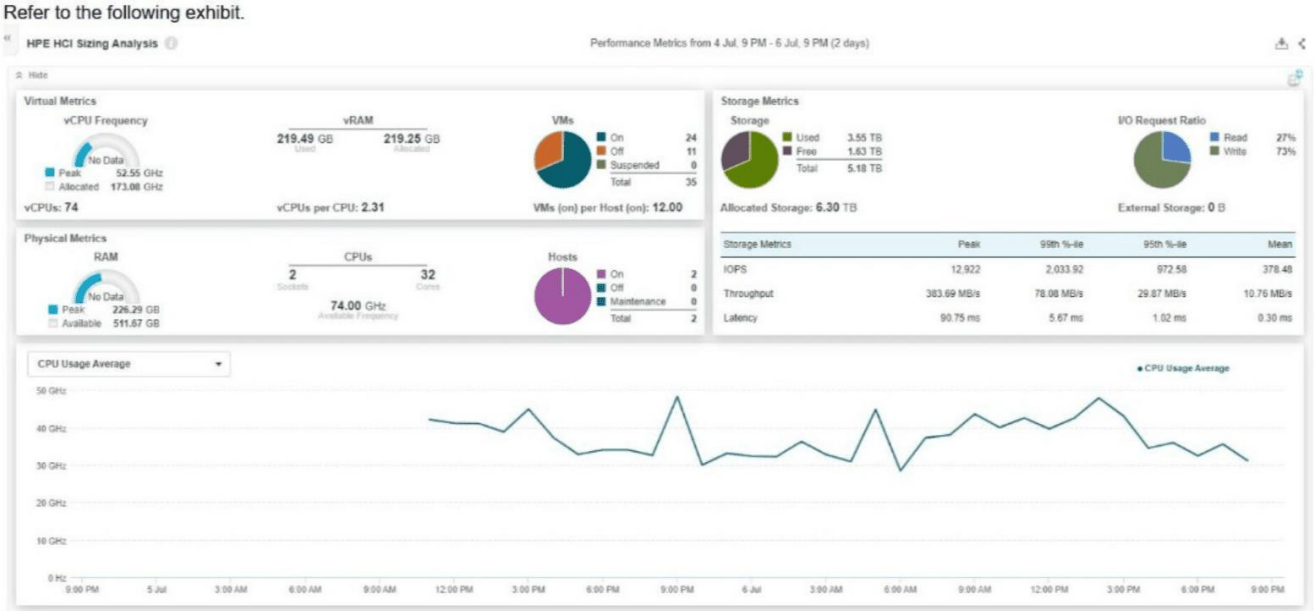

Which statement about the CloudPhysics report shown (vCenter Summary) is true?

A. The cluster has four total CPUs.

B. Each host has 256GB of RAM.

C. The cluster has enough resources to run all VMs during a double host failure.

D. Each host has 6.3TB of direct attached storage.

Answer: B

Explanation:
To verify the host-level specifications in the CloudPhysics report, one must look at the " Hosts " and " Memory " cards. The report indicates that there are " 2 Clustered Hosts " and " 0 Standalone Hosts, " for a total of 2 hosts. In the Memory card, the " Physical " RAM for the entire environment is listed as 511.75 GB.
By dividing the total physical RAM by the number of hosts (511.75 / 2), the result is approximately 255.875 GB. This confirms that each host is equipped with 256GB of physical RAM. This calculation is a standard part of the " Analyze " phase in the HPE architecting process to ensure the current hardware profile is understood before recommending a replacement or expansion. Statements about CPU counts or HA failure capacity often require deeper drill-down reports beyond this summary view.
References: HPE CloudPhysics vCenter Summary Report Metrics; HPE Edge-to-Cloud Solutions Course H61Z2S.

QUESTION NO: 9

Your customer wants to renew their outdated VMware environment. You have analyzed their environment and interviewed key stakeholders. Based on your findings, you have designed a new solution which you will present to the customer. How should you start your presentation?

A. With a report from Cloud Physics to explain your findings.

B. With a detailed Bill of Material (BOM) for all solution components.

C. With a schedule for a comprehensive on-site POC.

D. With a completed HPE Smart CID document.

Answer: A

Explanation:

When presenting a designed solution to a customer, it is essential to ground the proposal in data-driven evidence that validates the architect ' s conclusions. Starting with a CloudPhysics report is the most effective way to begin because it provides an objective baseline of the customer's current environment. It highlights existing performance bottlenecks, resource overprovisioning, and actual consumption patterns. By presenting these findings first, the architect builds credibility and ensures the customer understands " why " the new solution was designed a certain way before getting into technical specifications or pricing (BOM). This approach aligns the presentation with the customer ' s specific pain points and business requirements discovered during the analysis phase. Starting with a BOM or a POC schedule is premature until the customer agrees on the underlying problems identified in the discovery phase.

References: HPE Edge-to-Cloud Solutions Presentation Best Practices; HPE CloudPhysics Assessment Strategy.

QUESTION NO: 10

You are assessing an environment of 100 servers with mixed workloads on VMware and bare metal. Which tool should you use for this assessment?

A. iLO 5

B. CloudPhysics

C. SAF Collector

D. RVTools

Answer: C

Explanation:

While CloudPhysics is a powerful tool for analyzing virtualized VMware environments, it does not natively gather the same level of granular performance data from " bare metal " (non-virtualized) servers. For a large- scale assessment of 100 servers with a " mixed " workload profile, the Storage Assessment Foundry (SAF) Collector is the more appropriate tool. The SAF Collector is designed to discover and gather performance metadata across a wide variety of infrastructures, including various hypervisors and physical operating systems. It provides a holistic view of the environment ' s actual resource consumption, which is critical for designing a consolidated edge-to-cloud solution. By using SAF Collector, the architect can ensure that the resource requirements for both the virtual machines and the physical servers are accurately captured and factored into the final solution design. References: HPE Storage Assessment Foundry (SAF) Collector Guide; HPE Discovery and Assessment Strategy.

QUESTION NO: 11

For each use case, identify whether it is a traditional solution or a hybrid solution.

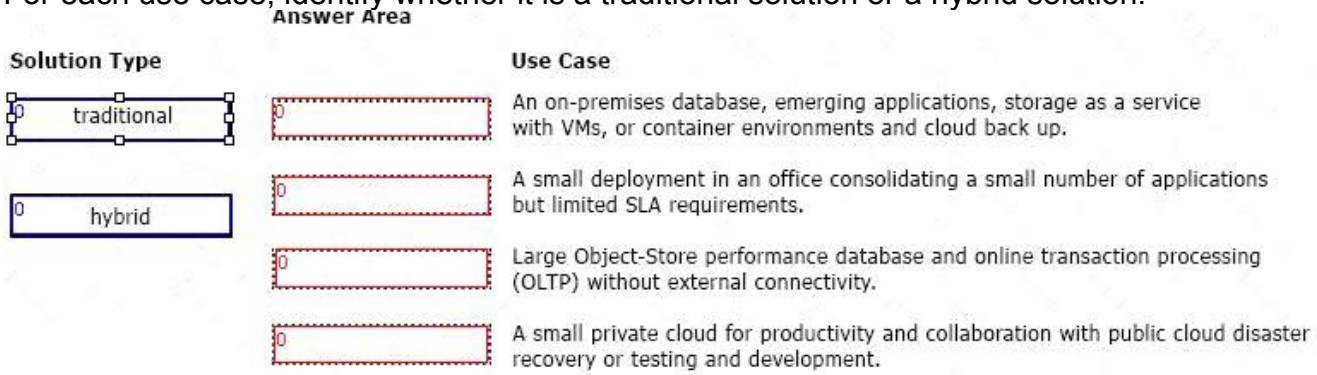

Answer:

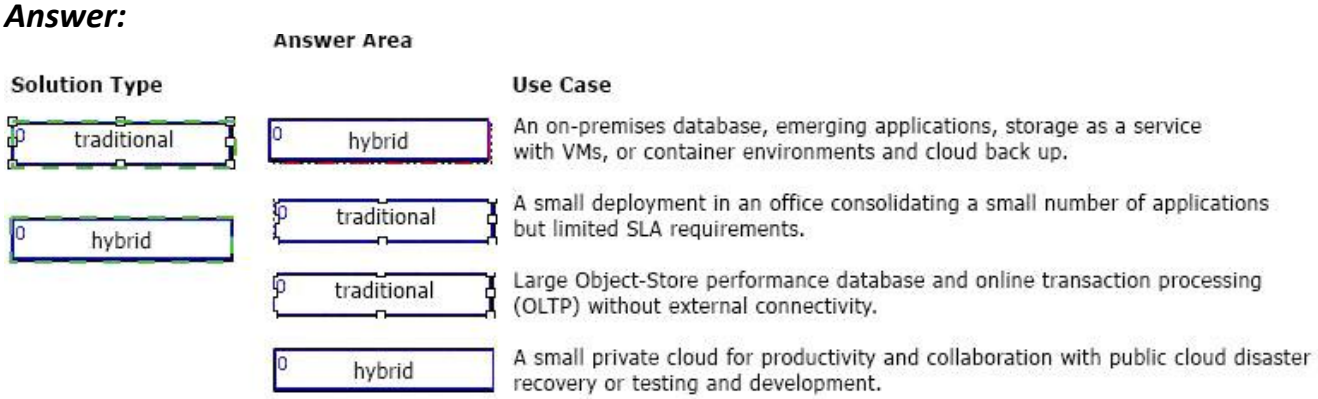

Explanation:

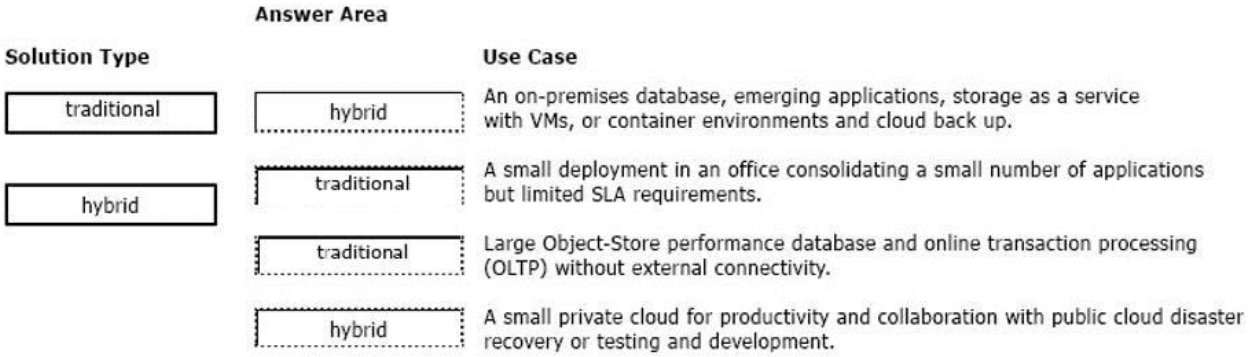

According to the HPE Edge-to-Cloud Adoption Framework , page 5, a traditional solution is a cloud deployment model where the cloud infrastructure is provisioned for exclusive use by a single organization comprising multiple consumers (e.g., business units). It may be owned, managed, and operated by the organization, a third party, or some combination of them, and it may exist on or off premises. A traditional solution offers the organization more control, security, and customization over their cloud resources, but it also requires more investment, maintenance, and expertise.

A hybrid solution is a cloud deployment model where the cloud infrastructure is a composition of two or more distinct cloud infrastructures (private, community, or public) that remain unique entities, but are bound together by standardized or proprietary technology that enables data and application portability (e.g., cloud bursting for load balancing between clouds). A hybrid solution offers the consumers more choice, agility, and innovation over their cloud resources, but it also requires more integration, management, and complexity.

Based on these definitions, the following use cases are hybrid solutions:

* An on-premises database, emerging applications, storage as a service with VMs, or container environments and cloud back up. = Hybrid solution. This use case involves a

combination of on- premises and cloud resources, and requires data and application portability between them.
* A small private cloud for productivity and collaboration with public cloud disaster recovery or testing and development. = Hybrid solution. This use case involves a combination of private and public cloud resources, and requires data and application portability between them.
The following use cases are traditional solutions:
* A small deployment in an office consolidating a small number of applications but limited SLA requirements. = Traditional solution. This use case involves a single organization using a cloud infrastructure for exclusive use, and does not require data and application portability to other clouds.
* Large Object-Store performance database and online transaction processing (OLTP) without external connectivity. = Traditional solution. This use case involves a single organization using a cloud infrastructure for exclusive use, and does not require data and application portability to other clouds