

## HP.HPE0-S59.v2026-07-10.q77

Exam Code:	HPE0-S59
Exam Name:	HPE Compute Solutions
Certification Provider:	HP
Free Question Number:	77
Version:	v2026-07-10
# of views:	144
# of Questions views:	770
<a href="https://www.freepdfdumps.com/HP.HPE0-S59.v2026-07-10.q77.html">https://www.freepdfdumps.com/HP.HPE0-S59.v2026-07-10.q77.html</a>	

### NEW QUESTION: 1

An organization manages its HPE ProLiant Gen11 servers, using iLO 6, in HPE Compute Ops Management.

What is one security risk that HPE Compute Ops Management can detect?

- A. The server ' s iLO allows SNMPv1 requests.
- B. The server ' s traffic patterns vary from normal.
- C. The server is exposed to attack due to the lack of antivirus software.
- D. The server has unencrypted data on its local drives.

**Answer: A (LEAVE A REPLY)**

HPE Compute Ops Management (COM) integrates deeply with the iLO 6 Security Dashboard to provide centralized security posture monitoring for a fleet of servers.

\* Security Risk Monitoring: COM monitors specific configuration parameters on the iLO management processor that are considered " risky " by modern security standards. One of the primary risks identified is the use of legacy, insecure protocols for management and alerting.

\* SNMPv1 Risk (Option A): SNMPv1 is an inherently insecure protocol because it sends community strings (passwords) in cleartext across the network. In the iLO 6 Security Dashboard, if the setting " SNMPv1 Request " is set to Enabled , iLO will flag this as an " At risk " security status. HPE Compute Ops Management surfaces these alerts in its cloud console, allowing administrators to identify servers where insecure management protocols are active.

\* Scope of Management: COM focuses on hardware, firmware, and management-level security . It evaluates settings such as:

- \* Minimum Password Length and complexity.
- \* Secure Boot status.
- \* Authentication requirements for host-based tools.
- \* Global Component Integrity (via SPDM).
- \* Protocol status (disabling IPMI/DCMI or SNMPv1).

Why other options are incorrect:

\* Option B: Monitoring for varying traffic patterns is typically a function of Network Detection and Response (NDR) or Intrusion Detection Systems (IDS) , such as those found in the Aruba networking portfolio. COM does not perform deep packet inspection or behavioral analysis of production network traffic.

\* Option C: Antivirus software is an Operating System (OS) level application. Since iLO and COM operate at the " below-the-OS " management layer, they do not have visibility into which software packages (like antivirus or EDR) are installed or running within the server ' s OS.

\* Option D: While iLO can manage Self-Encrypting Drives (SEDs) and display their encryption status, the presence of " unencrypted data " on a standard drive is an OS/Filesystem-level attribute. The security risks reported by the COM dashboard specifically target the management interface ' s vulnerabilities rather than the data residency state of the local storage.

### NEW QUESTION: 2

Your customer used to manage their HPE 3PAR arrays using HPE OneView. Now they have replaced these arrays with HPE Primera. How does managing the new arrays compare to managing their old arrays?

- A. Management procedures and features are the same for both types of the array.
- B. HPE OneView offers additional functionalities for HPE Primera like CPG management.
- C. HPE OneView offers limited support for HPE Primera because of the firmware differences.
- D. To manage HPE Primera using HPE OneView, a dedicated license is required.

**Answer: A (LEAVE A REPLY)**

HPE OneView is designed to provide a unified and consistent management experience across its entire software-defined storage portfolio. When a customer transitions from HPE 3PAR StoreServ to HPE Primera (or even newer Alletra 9000 systems), the core management workflows within HPE OneView remain essentially identical:

\* Storage Systems and Pools: In both cases, the storage system is added to the HPE OneView appliance.

Once connected, HPE OneView discovers the available storage pools, which are referred to as Common Provisioning Groups (CPGs) on both 3PAR and Primera.

\* Volume Management: The process of creating, provisioning, and managing volumes is the same.

Administrators use the same templates and server profiles to attach storage, regardless of the underlying hardware generation.

\* Automation and Templates: Because HPE OneView uses a software-defined approach, the " Logical " definitions of the infrastructure are decoupled from the physical hardware. This allows for a frictionless transition where existing automation scripts and REST API calls typically do not need to be modified when swapping a 3PAR array for a Primera array.

\* Licensing: Management of storage in HPE OneView does not require a different " type " of license when moving from 3PAR to Primera; the integration is built into the standard storage management capabilities of the appliance.

Why other options are incorrect:

\* Option B: While HPE Primera has many advanced internal features (like AI-driven optimizations via HPE InfoSight), these are handled at the array level. Within HPE OneView, CPG management and other basic provisioning features are functionally identical to those provided for 3PAR.

\* Option C: HPE OneView provides full support for HPE Primera. Far from being " limited, " Primera is a core platform for HPE's " Intelligent Storage " strategy.

\* Option D: There is no " dedicated " license required specifically for Primera management that differs from the existing 3PAR management model in OneView.

### NEW QUESTION: 3

What is a requirement for supervised training?

- A. Collecting performance metrics during the training process
- B. Using tools such as Prometheus to enhance visibility
- C. Having a team with multiple members collaborate across the ML/DL lifecycle
- D. Adding labels to training data

**Answer: D (LEAVE A REPLY)**

### NEW QUESTION: 4

For which interconnect type must a logical interconnect group be configured for only a single frame?

- A. HPE Virtual Connect SE 40 Gb F8 Module for Synergy
- B. Brocade 32GB Fibre Channel Switch Module for HPE Synergy
- C. HPE Virtual Connect SE 32 Gb FC Module for Synergy
- D. HPE Virtual Connect SE 100 Gb F32 Module for Synergy

**Answer: B (LEAVE A REPLY)**

In the HPE Synergy architecture, Logical Interconnect Groups (LIGs) define the connectivity patterns for the frames. The ability of an LIG to span multiple frames depends entirely on the interconnect technology being used:

\* Ethernet Virtual Connect (Options A & D): The HPE Virtual Connect SE 40 Gb F8 and 100 Gb F32 modules are designed for Multi-Frame Link Aggregation . Using a Master/Satellite architecture, these LIGs can span up to five Synergy frames, allowing a single logical management point for the entire fabric.

\* Fibre Channel Switches (Option B): The Brocade 32GB Fibre Channel Switch Module is a traditional " managed switch " persona. Unlike Virtual Connect, which acts as a fabric extender /aggregator, the Brocade module operates as a standalone switch within the frame ' s backplane. In HPE OneView, an LIG containing traditional Fibre Channel switches is strictly limited to a single frame . You cannot " stack " or span a Brocade switch LIG across multiple frames in the same way you do with Ethernet VC modules.

\* Virtual Connect FC (Option C): While Virtual Connect Fibre Channel modules also generally operate within a frame context for their uplinks, the question specifically targets the " Switch Module " (Brocade) which, as a standard industry switch, does not participate in the proprietary Synergy multi- frame horizontal stacking architecture.

#### **NEW QUESTION: 5**

An admin wants to experiment with HPE VM Essentials on a single node. What is one recommendation?

- A.** Manage the solution with the HPE VM Essentials Console, rather than HPE VM Essentials Manager (Morpheus)
- B.** Consolidate the management and compute networks
- C.** Use local data stores, rather than Ceph or other data store options
- D.** Use DHCP, rather than network pools, to assign IP addresses to VMs

**Answer: A (LEAVE A REPLY)**

#### **NEW QUESTION: 6**

Your customer is trying to grow a VMware duster using HPE OneView for VMware vCenter Server. They get an error and cannot proceed Which tool should you use to troubleshoot the problem?

- A.** HPE One View for VMware vCenter Server to verify if server profile template s property imported and registered
- B.** vCenter Server to verify if an OS Bund Plans property defined and has ESXi image uploaded
- C.** HPE OneView interface to import an existing cluster using the hypervisor cluster profile functionality.
- D.** ILO interlace of a given HPE Synergy compute node to verify whether this node is registered within HPE Oneview.

**Answer: A (LEAVE A REPLY)**

When encountering an error while growing a VMware cluster using HPE OneView for VMware vCenter Server, it is crucial to verify that the server profile template is properly imported and registered in HPE OneView. This ensures that the necessary configurations and settings are correctly applied to the new compute nodes being added to the cluster.

Reference: HPE OneView for VMware vCenter Server User Guide

#### **NEW QUESTION: 7**

What characteristic indicates that HPE Cray might be a better solution for a customer than HPE Private Cloud AI?

- A.** The customer Is focused on training many large and complex models.
- B.** The customer needs to use RAG with their inferencing workloads.
- C.** The customer needs to run AI workloads at the edge.
- D.** The customer is focused on deploying models in production.

**Answer: A (LEAVE A REPLY)**

#### **NEW QUESTION: 8**

In the past a customer experienced a networking issue where network ports went down and came back up in quick succession.

Which HPE Virtual Connect SE 100 GD F32 Module for HPE Synergy feature can be configured to disable such ports?

- A.** Pause flood protection

- B. sFlow
- C. Storm control
- D. Port flap protection

**Answer:** ([SHOW ANSWER](#))

Port flap protection is a feature of the HPE Virtual Connect SE 100Gb F32 Module for HPE Synergy that can be configured to disable ports experiencing frequent up and down transitions. This feature helps to stabilize the network by preventing the disruptive effects of port flapping.

Reference: HPE Virtual Connect SE 100Gb F32 Module for HPE Synergy User Guide

**NEW QUESTION: 9**

You configured a vVol datastore using HPE Storage integration Pack for VMware vCenter.

Which storage object should you check using SSMC to verify whether vVol datastore is configured properly?

- A. App Volume Set
- B. Storage Container
- C. Virtual Volume Set
- D. virtual Volume

**Answer:** ([SHOW ANSWER](#))

When you configure a vVol datastore using the HPE Storage Integration Pack for VMware vCenter, the storage object to check in SSMC (HPE StoreServ Management Console) is the Storage Container. The Storage Container is a logical storage entity that houses virtual volumes (vVols) and represents the vVol datastore in the storage system. Verifying the Storage Container ensures that the vVol datastore is properly configured and managed.

Reference: HPE 3PAR StoreServ Storage Concepts Guide

**NEW QUESTION: 10**

An HPE partner is selling an HPE Private Cloud AI solution.

For which need, not covered by the services always bundled with the solution, can an HPE partner offer a service?

- A. Installing the HPE AI Essentials software
- B. Connecting the HPE GreenLake File Storage to HPE AI Essentials
- C. Connecting HPE Private Cloud AI to the customer 's data center network
- D. Designing and setting up data pipelines for HPE Private Cloud AI

**Answer:** ([SHOW ANSWER](#))

**NEW QUESTION: 11**

Refer to the exhibit.

**Create Logical Interconnect Group** General ?

**General**

Name

**Logical Interconnect Group**

Using the selectors below, describe the logical interconnect group to be created and then click "Select interconnects" to see the bay and interconnect choices.

Interconnect type

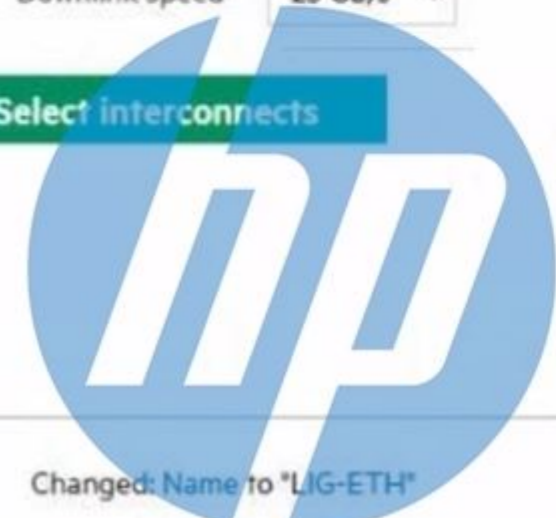
Enclosure count

Interconnect bay set

Redundancy

Downlink speed

[Select interconnects](#)



Changed: Name to "LIG-ETH"

Which statement about this logical interconnect group is true?

- A. The master modules are located in the interconnect bays 2 and 5
- B. For the given downlink speed, the maximum number of frames is reached
- C. Both master modules are located in the same HPE Synergy frame
- D. To use interconnect bay set 3, all compute nodes must be half-height

**Answer: C (LEAVE A REPLY)**

The exhibit shows the creation of a Logical Interconnect Group (LIG) using HPE Virtual Connect SE

100Gb F32 Modules . Here is the breakdown of why Option C is correct and why the others are not:

- \* Master Module Placement (Option C): In a redundant HPE Synergy configuration, the two " Master " (or Primary) interconnect modules must be installed in the same frame (referred to as the Master Frame). These modules are placed in the redundant bays of the chosen bay set (e.g., Bay 3 and Bay 6). The other frames in the logical enclosure (Enclosure count: 3) will house Interconnect Link Modules (ILMs) , which act as satellite modules to extend the fabric from the masters. This architecture ensures a single management point and a unified control plane for the entire multi-frame fabric.
- \* Bay Set Mapping (Option A): The exhibit explicitly shows the Interconnect bay set is set to 3 . In the HPE Synergy 12000 Frame, bay sets map to specific bays:
  - \* Bay Set 1: Bays 1 and 4
  - \* Bay Set 2: Bays 2 and 5
  - \* Bay Set 3: Bays 3 and 6 Since the LIG is configured for bay set 3, the master modules will be in bays 3 and 6 , making Option A incorrect.
- \* Frame Scalability (Option B): The HPE Virtual Connect SE 100Gb F32 Module supports a master /satellite architecture that can scale up to five (5) frames in a single fabric (1 Master Frame + 4 Satellite Frames). The exhibit shows an enclosure count of 3 , which is well within the supported limit and not the maximum for a 25 Gb/s downlink configuration.
- \* Compute Node Form Factor (Option D): Both half-height (e.g., Synergy 480) and full-height (e.g., Synergy 660) compute modules can use interconnect bay set 3. Half-height modules have three mezzanine slots, where slot 3 connects to bay set 3. Full-height modules have six mezzanine slots, where slot 3 and slot 6 can connect to these fabrics. There is no requirement that nodes must be half- height to utilize this fabric.

#### NEW QUESTION: 12

Which statement about port flap protection is true?

- A. It can disable link aggregation groups or stacking ports that have gone down and come up in quick succession.
- B. It allows the creation of LACP groups on the upstream switch that span multiple HPE Synergy Virtual Connect modules.
- C. It allows configuration of hardware Layer 2 switching behavior of multicast traffic to optimize network resource usage.
- D. It allows an administrator to suppress excessive inbound multicast, broadcast, and DLF packets.

**Answer: A (LEAVE A REPLY)**

#### NEW QUESTION: 13

Your customer used to manage their HPE 3PAR arrays using HPE OneView Now they have replaced these arrays HPE Primera.

How does managing the new arrays compare to managing their old arrays?

- A. To manage HPE Primera using HPE OneView. a dedicated license is required
- B. Management procedures and features are the same for both types of the array
- C. HPE OneView offers limited support for HPE Primera because of the firmware differences
- D. HPE OneView offers additional functionalities for HPE Primera like CPG management

**Answer: B (LEAVE A REPLY)**

#### NEW QUESTION: 14

Your customer wants to update the firmware for 36 HPE Synergy 480 Gen 10 compute modules with a server profile assigned from a server profile template.

Put the steps required to achieve this goal on the left into their correct order on the right.

Steps

- Select Firmware baseline within server profile template.
- Select Installation method (online or offline).
- Update all server profiles from the template.
- Upload SPP to HPE OneView repository.

Answer Area






Answer:

**Steps**

- Select Firmware baseline within server profile template.
- Select Installation method (online or offline).
- Update all server profiles from the template.
- Upload SPP to HPE OneView repository.

**Answer Area**

- Upload SPP to HPE OneView repository.
- Select Firmware baseline within server profile template.
- Select Installation method (online or offline).
- Update all server profiles from the template.

Explanation:

To update the firmware for the HPE Synergy 480 Gen 10 compute modules with a server profile assigned from a server profile template, the steps should be performed in the following order:

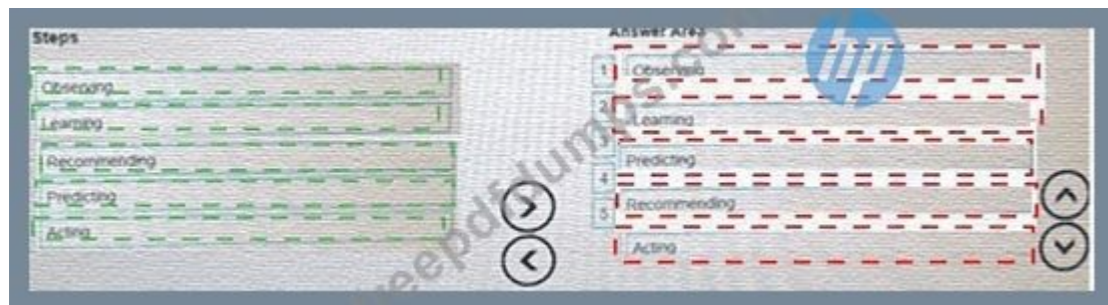
- \* Upload SPP to HPE OneView repository. (First, the Service Pack for ProLiant (SPP) containing the firmware updates must be uploaded to the HPE OneView repository.)
- \* Select Firmware baseline within server profile template. (Next, you need to select the desired firmware baseline in the server profile template.)
- \* Select installation method (online or offline). (Choose whether the firmware updates will be installed online or offline.)
- \* Update all server profiles from the template. (Finally, update all server profiles from the modified template to apply the firmware updates to the compute modules.)

**NEW QUESTION: 15**

Put the steps of the HPE infosight flow in the correct order.



**Answer:**



Explanation:

(observing, learning, predicting, recommending, acting)

- \* Observing: The first step involves collecting and monitoring data from the infrastructure. HPE InfoSight continuously gathers telemetry data from various components, including storage, compute, and networking devices.
- \* Learning: In this step, the collected data is analyzed to identify patterns and trends. Machine learning algorithms process the data to understand the normal behavior of the system and detect anomalies.
- \* Predicting: Based on the learned patterns, HPE InfoSight can predict potential issues before they occur. This predictive analysis helps in identifying areas that might cause problems in the future.
- \* Recommending: After predicting potential issues, HPE InfoSight provides actionable recommendations to prevent or resolve these issues. These recommendations are based on best practices and the vast amount of data analyzed by InfoSight.
- \* Acting: Finally, the recommendations are implemented to optimize the performance and reliability of the infrastructure. This step may involve configuring settings, applying updates, or other actions to maintain the health of the system.

Reference: HPE InfoSight Overview

**NEW QUESTION: 16**

A junior architect designed a solution based on HPE Synergy that must meet the following requirements:

- \* Support at least 30 HPE Synergy 480 Gen10 Plus compute modules
- \* Interconnect modules must provide MAC and WWN virtualization
- \* All HPE Synergy frames must formulate a single management entity
- \* The management appliance should not be a single point of failure
- \* Connectivity to FC array should be available to all compute nodes

The proposal consists of:

- \* 3 HPE Synergy 12000 frames
- \* 6 HPE Composer 2 modules
- \* 2 HPE Virtual Connect SE 100Gb F32 Module
- \* 4 HPE Synergy 50Gb Interconnect Link Module

Which statements about this proposed design are true? (Select two)

- A.** FC upgrade licenses for HPE Virtual Connect SE 100Gb F32 Module must be added.

- B. HPE Synergy 50Gb ILMs must be replaced with HPE Virtual Connect SE 100Gb F32 modules.
- C. Remove 4 HPE Composer 2 modules from the design, as they are not required.
- D. The selected interconnect modules do not provide the requested WWN virtualization.
- E. The current setup can support up to 18 HPE Synergy 480 Gen10 Plus compute modules.

**Answer: A,C (LEAVE A REPLY)**

This scenario tests your knowledge of HPE Synergy Master/Satellite architecture and Management Ring requirements. Based on the requirements provided, the design has two primary issues that must be addressed:

\* Management Redundancy (HPE Composer 2):

\* The requirement states that all frames must formulate a single management entity. In a Synergy environment, this is achieved by linking the Frame Link Modules (FLM) of the three frames together to form a management ring.

\* Within a single management ring (one Logical Enclosure), only two (2) HPE Synergy Composer 2 modules are required to provide a highly available, redundant management cluster.

\* The architect 's proposal included 6 Composers (2 per frame). This is incorrect for a single management entity design; therefore, 4 Composers are unnecessary and should be removed (Option C).

\* Fibre Channel Connectivity Licensing:

\* The HPE Virtual Connect SE 100Gb F32 Module is a converged interconnect. While it physically supports Fibre Channel (FC) and FCoE, the Fibre Channel uplink functionality is not enabled by default.

\* According to HPE QuickSpecs, a specific FC Upgrade License must be purchased and applied via HPE OneView to enable the use of the QSFP28 ports as native Fibre Channel uplinks to a SAN. Without these licenses, the customer cannot meet the requirement for connectivity to an FC array (Option A).

Analysis of other options:

\* Option B is incorrect: In a 3-frame setup, you only need one pair of " Master " modules (the F32s). The other frames are correctly equipped with Interconnect Link Modules (ILMs) which act as satellites to extend the Master ' s fabric.

\* Option D is incorrect: Virtual Connect technology is explicitly designed to provide MAC and WWN virtualization (often called " Server-Side Networking " ), allowing identities to be assigned to server profiles rather than being hardcoded to hardware.

\* Option E is incorrect: A 3-frame Synergy setup provides a total of 36 half-height slots (12 per frame).

Since the HPE Synergy 480 Gen10 Plus is a half-height module, the physical infrastructure can support up to 36 modules, easily meeting the " at least 30 " requirement. There is no architectural limitation in this specific F32/ILM combination that restricts it to only 18 modules.

**Valid HPE0-S59 Dumps** shared by Actual4test.com for Helping Passing HPE0-S59 Exam! Actual4test.com now offer the **newest HPE0-S59 exam dumps**, the Actual4test.com HPE0-S59 exam **questions have been updated** and **answers have been corrected** get the **newest** Actual4test.com HPE0-S59 dumps with Test Engine here: [https://www.actual4test.com/HPE0-S59\\_examcollection.html](https://www.actual4test.com/HPE0-S59_examcollection.html) (160 Q&As Dumps, **30%OFF Special Discount: Freepdfdumps**)

#### **NEW QUESTION: 17**

Your customer plans to deploy VMware ESXi 7.0 U2, and they are looking for a hardware platform that will allow them to use up to 16 CPU sockets Which HPE compute system meets the customer requirements?

- A. HPE ProLiant DL580 Gen10
- B. HPE ProLiant OL380 Gen 10 Plus
- C. HPE Synergy 480 Gen10 Plus
- D. HPE Superdome Flex system

**Answer: D (LEAVE A REPLY)**

The HPE Superdome Flex system is the only HPE compute system that supports up to 16 CPU sockets, making it suitable for customers who require a hardware platform with such high scalability. The Superdome Flex is designed for mission-critical workloads and provides exceptional scalability and performance.

Reference: HPE Superdome Flex QuickSpecs

**NEW QUESTION: 18**

An admin has found that an update is available for an HPE Private Cloud AI solution 's Private Cloud AI software. Before the admin can install the update, which step is required?

- A. Downloading the desired software to each node in the HPE Private Cloud AI solution
- B. Running a precheck on the desired new software catalog
- C. Updating the HPE GreenLake for File Storage software to a supported version for the desired software catalog
- D. Manually placing the AI optimized workers in maintenance mode

**Answer: B (LEAVE A REPLY)**

**NEW QUESTION: 19**

Your customer plans to deploy ESXi systems to HPE Synergy and HPE ProLiant servers using HPE OneView for VMware vCenter Server.

Which statement about HPE OneView for VMware vCenter Server functionality is true for this environment?

- A. Individual systems can be added to HPE OneView for VMware vCenter Server through an iLO management processor
- B. Only systems managed using HPE Oneview can use HPE Oneview for VMware vCenter Server functionalities
- C. HPE ProLiant servers need an HPE OneView for VMware vCenter Server license for full functionality
- D. HPE Synergy requires HPE Composer 2 for full functionality of HPE OneView for VMware vCenter Server

**Answer: B (LEAVE A REPLY)**

**NEW QUESTION: 20**

Your customer wants to compare HPE Superdome Flex with HPE Superdome Flex 280.

Which statement about these two systems is true?

- A. HPE Superdome Flex 280 can be managed using iLO5 management processor.
- B. HPE Superdome Flex must be in memory mode to support HPE Persistent Memory.
- C. HPE Superdome Flex requires Rack Management Controller to support multiple nPars.
- D. HPE Superdome Flex 280 can scale up to 16 sockets and 24TB of memory.

**Answer: A (LEAVE A REPLY)**

**NEW QUESTION: 21**

Your customer plans to deploy VMware vSAN using D3940 Storage Modules.

Which statement about this solution is true?

- A. The 12Gb SAS modules required for D3940 connectivity can only be installed in the first fabric.
- B. To use the storage controller that will access D3940 storage modules, a second CPU must be installed.
- C. For VMware vSAN deployment, all drives in the D3940 storage modules must be SSD drives.
- D. All drives in the D3940 storage modules must formulate a single logical drive with RAID5 configured

**Answer: C (LEAVE A REPLY)**

For VMware vSAN deployments, it is recommended that all drives in the D3940 storage modules be SSD drives to ensure optimal performance and efficiency. SSDs provide the necessary speed and reliability required by vSAN for handling storage-intensive operations and maintaining high performance.

Reference: VMware vSAN Hardware Quick Reference Guide

**NEW QUESTION: 22**

Which statement about HPE OneView milestone releases is true?

- A. A milestone release is available for selected partners that gives access to HPE OneView features which are not officially announced.
- B. A milestone release is a release which must be installed on HPE Composer to support current version of the Service Pack for Synergy.
- C. A milestone release is a release that is published when HPE releases new hardware components to immediately support new systems.
- D. A milestone release is a release with an enhanced update architecture that is a prerequisite prior to updating to a subsequent release.

**Answer:** [\(SHOW ANSWER\)](#)

In the lifecycle of HPE OneView software updates, a Milestone Release is a specific version of the appliance firmware that acts as a mandatory "stepping stone" or checkpoint in the upgrade path.

\* Prerequisite for Updates: Because HPE OneView undergoes significant architectural changes over time (such as database schema migrations, changes to the underlying operating system, or shifts in the container orchestration layer), you cannot always jump from a very old version directly to the latest version. A milestone release contains the necessary logic to transform the existing data and environment into a format that the subsequent versions require.

\* Update Path Enforcement: The HPE OneView update logic is designed to prevent administrators from skipping these critical versions. If an appliance is at a version significantly older than the target version, the update will fail or be blocked until the required milestone release is installed first.

\* Examples of Milestones: Historically, versions like HPE OneView 4.2, 5.2, 6.6, and 8.0 have served as milestones. For instance, an administrator on version 5.0 might be required to update to 5.2 (the milestone) before they are permitted to move to version 6.0 or higher.

\* Update Readiness Checker: HPE highly recommends using the Update Readiness Checker specifically when approaching a milestone release to ensure the appliance's internal health is sufficient to handle the intensive data migrations that typically occur during these transitions.

Why other options are incorrect:

\* Option A: This describes a "Beta" or "Early Access" program, not a milestone release.

\* Option B: While specific OneView versions are required for specific Service Packs for Synergy (SSP)

, the term "milestone" refers specifically to the appliance's internal update architecture and pathing, not its external hardware compatibility matrices.

\* Option C: While new hardware support is often included in major or minor releases, the primary definition of a milestone is its role as a prerequisite in the software update sequence.

**NEW QUESTION: 23**

One of the administrators added some connections within a server profile that was created using a template.

This server profile did not report inconsistency.

What is the most likely reason for this behavior?

- A. The server is in maintenance mode and does not report any issues.
- B. Consistency check for connections is set to Minimum or Not checked.
- C. The administrator removed FC connections that are not monitored by default.
- D. Connections defined in the server profile are not monitored for compliance.

**Answer:** [B \(LEAVE A REPLY\)](#)

**NEW QUESTION: 24**

Your customer added a new Ethernet network in HPE Oneview that is used to manage their HPE Synergy platform.

When the customer tries to connect one of the compute nodes to this network through the server profile, this network is not available.

What must be done to fix this issue?

- A. Check the Logical Enclosure settings and verify if the maximum number of networks allowed for this LE is reached.
- B. To make this network available to use in the server profile, the customer must enable the Smart Link feature within network definition.
- C. The newly created network must be added to the uplink set within the LIG, and Logical interconnect configuration must be updated from the LIG.

**D.** Check within the server profile which server hardware is used and perform the Refresh operation on this compute node to get the latest adapter information.

**Answer: C (LEAVE A REPLY)**

When a new Ethernet network is created in HPE OneView and is not available in the server profile, it is typically because the network has not been added to the uplink set within the Logical Interconnect Group (LIG). To resolve this issue, the network must be included in the uplink set configuration, and then the logical interconnect configuration must be updated from the LIG to propagate the changes.

Reference: HPE OneView Networking Guide

#### **NEW QUESTION: 25**

Your customer has an HPE Synergy platform deployed and managed using HPE OneView. The customer also has HPE Superdome Flex with RMC deployed. They plan to deploy a second nPar and HPE Superdome Flex

280. They plan to manage new components using HPE OneView.

What must be done to achieve this goal?

**A.** The customer must use IL05 to add HPE Superdome Flex systems to HPE OneView.

**B.** The customer should deploy an HPE OneView appliance for HPE Superdome Flex management.

**C.** The customer must add a second Rack Management Controller for HPE Superdome Flex 280.

**D.** The customer has to add HPE Superdome Flex to the HPE OneView used to manage HPE Synergy.

**Answer: D (LEAVE A REPLY)**

To manage HPE Superdome Flex systems, including the new Superdome Flex 280, using HPE OneView, the customer should deploy a dedicated HPE OneView appliance specifically for HPE Superdome Flex management. This approach ensures that the Superdome Flex systems are properly integrated and managed within the HPE OneView environment, allowing for consistent and centralized management alongside other HPE infrastructure.

Reference: HPE Superdome Flex Management with HPE OneView

#### **NEW QUESTION: 26**

Which tool can be used to evaluate the health of an HPE OneView appliance before proceeding with an appliance update?

**A.** HPE Composer Maintenance Console

**B.** HPE OneView Update Readiness Checker

**C.** HPE OneView Global Dashboard

**D.** HPE OneView Firmware Compliance Report

**Answer: B (LEAVE A REPLY)**

The HPE OneView Update Readiness Checker is a tool specifically designed to evaluate the health and readiness of an HPE OneView appliance before proceeding with an update. It checks for any potential issues that could affect the update process and provides recommendations for resolution.

Reference: HPE OneView Update Guide

#### **NEW QUESTION: 27**

Your customer is implementing an HPE Synergy solution based on three frames.

Click each of the ports that can be used to connect satellite modules.



Answer:



Explanation:



The QSFP+ ports (Q1 and Q2) on the HPE Virtual Connect SE 100Gb F32 Module for HPE Synergy can be used to connect satellite modules. These ports are designed to handle high-speed connectivity required for linking additional frames and expanding the Synergy infrastructure.

Reference: HPE Virtual Connect SE 100Gb F32 Module for Synergy

#### NEW QUESTION: 28

What is one benefit of NVIDIA AI Enterprise on HPE Private Cloud AI?

- A. Idle GPU-enabled workloads are cleaned up by NVIDIA AI Enterprise schedulers.
- B. All the necessary GPU operators and drivers are deployed for containerized AI applications to use GPUs.
- C. Users receive access to GPU models that are not publicly available.
- D. NVIDIA AI Enterprise sets up secure communications between all GPU-optimized workloads.

**Answer: B (LEAVE A REPLY)**

The integration of NVIDIA AI Enterprise (NVAIE) into the HPE Private Cloud AI stack is designed to remove the operational complexity of managing high-performance GPU hardware in a containerized environment.

\* Automation via GPU Operator: One of the most significant features of NVAIE is the NVIDIA GPU Operator . In a standard Kubernetes environment, administrators would typically need to manually install GPU drivers, container runtimes, and monitoring tools on every node. The GPU Operator automates this entire lifecycle. It detects the presence of NVIDIA GPUs and automatically deploys the necessary drivers, the NVIDIA Container Toolkit , and the Kubernetes Device Plugin .

\* Infrastructure Readiness: By including NVAIE, HPE ensures that the " Infrastructure Layer " is fully optimized for AI workloads. This means that as soon as an AI worker node (like an HPE ProLiant DL380a) is provisioned, the software stack is ready to pass GPU instructions from a containerized application directly to the hardware without manual intervention.

\* Consistency and Support: NVAIE provides a validated and supported path for these drivers and operators. This ensures that the versions of the drivers are compatible with the AI frameworks (like PyTorch or TensorFlow) and the specific version of Kubernetes running on the HPE Private Cloud AI, reducing " version hell " and ensuring enterprise-grade stability.

Why other options are incorrect:

\* Option A: While resource scheduling and orchestration (via tools like NVIDIA Run:ai, now part of NVAIE) can manage workload placement, the " cleanup of idle workloads " is typically a function of the Kubernetes scheduler or specific policy engines (like Kyverno), not the primary defining benefit of NVAIE itself.

\* Option C: NVIDIA AI Enterprise is a software platform . It does not provide access to " unreleased " or " non-public " hardware models; rather, it provides the software stack to run on commercially available NVIDIA GPUs like the H100, L40S, or B200.

\* Option D: Secure communication between workloads is usually handled by the Service Mesh (such as Istio , which is part of the HPE AI Essentials software layer) or networking operators, rather than NVAIE ' s primary role of GPU enablement.

**NEW QUESTION: 29**

Your customer has a logical enclosure configured using a single HPE Synergy frame with the following configuration:

\* Two HPE Virtual Connect SE 100Gb F32 Modules for Synergy

\* Two HPE Synergy Virtual Connect SE 32Gb FC Modules

\* Twelve HPE Synergy 480 Gen10 Plus compute nodes with appropriate mezzanine cards They plan to add a new HPE Synergy frame to an existing logical enclosure with the following configuration:

\* Two HPE Synergy Virtual Connect SE 32Gb FC Modules

\* Two HPE Synergy 20Gb Interconnect Link Modules

\* Eight HPE Synergy 480 Gen10 Plus compute nodes with appropriate mezzanine cards Which statement about the planned configuration change is true?

**A.** Four compute modules must be moved to the second frame to balance the configuration.

**B.** The two HPE Synergy 20Gb Interconnect Link Modules must be replaced with the 50Gb option.

**C.** A license for an extended logical enclosure must be added to HPE OneView through GUI.

**D.** The HPE Synergy Virtual Connect SE 32Gb FC Modules from both frames should be stacked.

**Answer: (SHOW ANSWER)**

The HPE Synergy Master/Satellite architecture is strictly defined by the compatibility between the " Master " (Primary) interconnect modules and the " Satellite " (Interconnect Link Modules - ILM) modules across multiple frames.

\* Fabric Generations: HPE Synergy fabrics are categorized into generations. The HPE Virtual Connect SE 100Gb F32 Module is a second-generation (Gen2) " Master " module. It is designed to work with the HPE Synergy 50Gb Interconnect Link Module as its satellite to extend the fabric to additional frames.

\* ILM Compatibility (The 20Gb vs 50Gb Issue): The 20Gb Interconnect Link Modules mentioned in the proposed new frame are first-generation (Gen1) satellite modules. They were designed to work with the HPE Virtual Connect SE 40Gb F8 Module (Gen1 Master).

\* According to HPE Synergy cabling and fabric rules, you cannot mix Gen1 satellites (20Gb ILM) with Gen2 masters (100Gb F32).

\* For the 100Gb F32 Master modules in Frame 1 to extend the Ethernet/FCoE fabric to Frame 2, Frame 2 must use the 50Gb Interconnect Link Modules .

\* Fibre Channel Connectivity: Unlike the Ethernet fabric, the HPE Synergy Virtual Connect SE

32Gb FC Modules are not master/satellite modules. They are standalone SAN interconnects. Each frame that requires native Fibre Channel connectivity typically contains its own pair of FC modules.

They are managed within the same Logical Enclosure in HPE OneView, but they do not use " Interconnect Link Modules " or " stacking " in the way Ethernet masters do.

Why other options are incorrect:

\* Option A: There is no requirement to " balance " the number of compute modules across frames (12 in one, 8 in another) for a Logical Enclosure to function.

\* Option C: HPE OneView licensing for Synergy is per-node/per-frame (OneView Advanced). There is no specific " Extended Logical Enclosure " license required to add a second frame to an existing LE; the management infrastructure simply scales as frames are added.

\* Option D: " Stacking " is a term often used for traditional switches. In Synergy, the 32Gb FC modules provide individual uplinks to the SAN fabric. While they are part of the same logical management entity, the term " stacked " is technically inaccurate for how these modules are interconnected or managed across frames in this specific architecture.

#### NEW QUESTION: 30

What indicates that organization is more advanced than a beginner, but is still an early AI user?

A. The organization has a data scientist on staff but no models in production.

B. The organization has deployed some models but lacks standardization for their tools and processes.

C. The organization has identified use cases for AI, but does not yet have a team to work on the projects.

D. The organization has standardized AI development and deployment processes that they want to scale.

Answer: ([SHOW ANSWER](#))

#### NEW QUESTION: 31

What is one of the benefits of using HPE Composer 2?

A. HPE Composer 2 has 128GB of memory and 4 AMD CPUs to improve performance of the management system

B. HPE Composer 2 modules are required to manage HPE virtual Connect SE 100GD F32 Modules for Synergy

C. The administrator can access an HPE Synergy Composer 2 appliance remotely to perform First Time Setup

D. A pair of HPE Composer 2 modules can manage 42 frames, compared to 21 frames managed by HPE Composer

Answer: D ([LEAVE A REPLY](#))

One of the key benefits of using HPE Composer 2 is its enhanced scalability. A pair of HPE Composer 2 modules can manage up to 42 frames, which is double the capacity managed by the previous generation HPE Composer (which manages up to 21 frames). This increased capacity allows for more efficient and large-scale management of HPE Synergy environments.

Reference: HPE Synergy Composer 2 QuickSpecs

**Valid HPE0-S59 Dumps** shared by Actual4test.com for Helping Passing HPE0-S59 Exam! Actual4test.com now offer the **newest HPE0-S59 exam dumps**, the Actual4test.com HPE0-S59 exam **questions have been updated** and **answers have been corrected** get the **newest** Actual4test.com HPE0-S59 dumps with Test Engine here: [https://www.actual4test.com/HPE0-S59\\_examcollection.html](https://www.actual4test.com/HPE0-S59_examcollection.html) (160 Q&As Dumps, **30%OFF Special Discount: Freepdfdumps**)

#### NEW QUESTION: 32

You are setting up HPE VM Essentials on several hosts. The hosts communicate on VLAN 10, which is untagged on a link aggregation on each host.

A separate link aggregation on each host will carry traffic for the VMs. Each host also has two links in separate subnets for their storage traffic. All IP addressing uses IPv4.

The output for the ip a command on one of the hosts is shown below.

```

1: lo: <LOOPBACK,UP,LOWER_UP> mtu 65536 qdisc noqueue state UNKNOWN group default qlen 1000
   link/loopback 00:00:00:00:00:00 brd 00:00:00:00:00:00
   inet 127.0.0.1/8 scope host lo
       valid_lft forever preferred_lft forever
   inet6 ::1/128 scope host
       valid_lft forever preferred_lft forever
2: ens3f0: <BROADCAST,MULTICAST,SLAVE,UP,LOWER_UP> mtu 1500 qdisc mq master bond0 state UP group default qlen 10
   link/ether 46:b9:2e:25:e4:b5 brd ff:ff:ff:ff:ff:ff permaddr 92:68:47:c0:00:6a
3: ens3f1: <BROADCAST,MULTICAST,SLAVE,UP,LOWER_UP> mtu 1500 qdisc mq master bond0 state UP group default qlen 10
   link/ether 46:b9:2e:25:e4:b5 brd ff:ff:ff:ff:ff:ff permaddr 92:68:47:c0:00:6b
4: ens3f2: <BROADCAST,MULTICAST,SLAVE,UP,LOWER_UP> mtu 1500 qdisc mq master bond1 state UP group default qlen 10
   link/ether 02:2b:e2:c6:63:50 brd ff:ff:ff:ff:ff:ff permaddr 92:68:47:c0:00:6c
5: ens3f3: <BROADCAST,MULTICAST,SLAVE,UP,LOWER_UP> mtu 1500 qdisc mq master bond1 state UP group default qlen 10
   link/ether 02:2b:e2:c6:63:50 brd ff:ff:ff:ff:ff:ff permaddr 92:68:47:c0:00:6d
6: ens3f4: <BROADCAST,MULTICAST,SLAVE,UP,LOWER_UP> mtu 1500 qdisc noqueue state UP group default qlen 1000
   link/ether 02:2b:e2:c6:63:50 brd ff:ff:ff:ff:ff:ff permaddr 92:68:47:c0:00:6c
   inet 10.122.9.10/24 brd 10.122.9.255 scope global ens3f4
       valid_lft forever preferred_lft forever
7: ens3f5: <BROADCAST,MULTICAST,SLAVE,UP,LOWER_UP> mtu 1500 qdisc noqueue state UP group default qlen 1000
   link/ether 02:2b:e2:c6:63:50 brd ff:ff:ff:ff:ff:ff permaddr 92:68:47:c0:00:6d
   inet 10.122.10.10/24 brd 10.122.10.255 scope global ens3f5
       valid_lft forever preferred_lft forever
8: bond0: <BROADCAST,MULTICAST,MASTER,UP,LOWER_UP> mtu 1500 qdisc noqueue state UP group default qlen 1000
   link/ether 46:b9:2e:25:e4:b5 brd ff:ff:ff:ff:ff:ff
   inet 10.67.100.71/24 brd 10.67.100.255 scope global bond0
       valid_lft forever preferred_lft forever
   inet6 fe80::44b9:2eff:fe25:e4b5/64 scope link
       valid_lft forever preferred_lft forever
9: bond1: <BROADCAST,MULTICAST,MASTER,UP,LOWER_UP> mtu 1500 qdisc noqueue state UP group default qlen 1000
   link/ether 02:2b:e2:c6:63:50 brd ff:ff:ff:ff:ff:ff
   inet6 fe80::2b:e2ff:fec6:6350/64 scope link
       valid_lft forever preferred_lft forever
10: virbr0: <NO-CARRIER,BROADCAST,MULTICAST,UP> mtu 1500 qdisc noqueue state DOWN group default qlen 1000
   link/ether 52:54:00:8a:c2:db brd ff:ff:ff:ff:ff:ff
   inet 192.168.122.1/24 brd 192.168.122.255 scope global virbr0
       valid_lft forever preferred_lft forever

```

What should you specify for the management interface when installing the HPE VM Essentials Manager on this host?

- A. ens3f0
- B. bond0.10
- C. vlan10
- D. bond0

**Answer: D (LEAVE A REPLY)**

### NEW QUESTION: 33

Your customer plans to manage HPE Primera using HPE Oneview

Which storage management tasks can be completed using HPE OneView for this environment?

- A. CPG management
- B. Volume provisioning
- C. Cache management
- D. Replication setup

**Answer: B (LEAVE A REPLY)**

HPE OneView allows for several storage management tasks when managing HPE Primera, including volume provisioning. This task involves creating and managing storage volumes that can be allocated to compute nodes and applications within the environment. HPE OneView streamlines this process, providing a centralized management interface.

Reference: HPE OneView and HPE Primera Integration Guide

#### NEW QUESTION: 34

Which statement about HPE Superdome Flex is true?

- A. It requires at least 768GB of memory per chassis.
- B. It supports 2 to 8 sockets in 2-socket increments.
- C. It can support up to two nPars with external RMC.
- D. It cannot be managed using HPE OneView.

**Answer: A (LEAVE A REPLY)**

The HPE Superdome Flex (the 4-socket to 32-socket scale-up system) has specific architectural requirements that distinguish it from the smaller Superdome Flex 280:

\* Memory Configuration (Option A): The HPE Superdome Flex is designed for massive in-memory workloads. Each 4-socket chassis contains 48 DIMM slots (12 per processor). To maintain the required performance and RAS (Reliability, Availability, and Serviceability) features, the platform has a minimum memory requirement. According to the HPE QuickSpecs, the shared memory capacity starts at 768GB , which is achieved by populating all 48 slots with the minimum supported 16GB DIMMs ( $48 \times 16 \text{GB} = 768 \text{GB}$ ).

\* Scalability (Option B): This statement is incorrect for the standard Superdome Flex. The Superdome Flex 280 scales in 2-socket increments from 2 to 8 sockets. The original Superdome Flex scales in 4-socket increments (one full chassis at a time) up to a maximum of 32 sockets (8 chassis).

\* Hard Partitioning (Option C): Using an external Rack Management Controller (RMC) , the Superdome Flex can support significantly more than two partitions. It can support up to 16 nPars (hard partitions), allowing for massive workload consolidation with electrical isolation between partitions.

\* Management (Option D): This is incorrect. While the Superdome Flex initially focused on CLI management via the RMC, HPE OneView (versions 5.0 and later) fully supports the Superdome Flex for monitoring, inventory, and health management. More recent updates have introduced " Manage Mode " for template-based BIOS and firmware orchestration.

#### NEW QUESTION: 35

Your customer uses HPE OneView to manage their HPE Synergy environment. They plan to use it to manage a new set equipment that includes.

- 20 HPE ProLiant DL365 Gen10 Plus servers
- 20 HPE ProLiant DL380 Gen10 Plus servers
- 5 HPE Primera 650
- 5 HPE MSA 2062 Storage Array

Which statements about how OneView will work with this new equipment are true? (Select two)

- A. HPE ProLiant DL365 Gen10 Plus servers are not supported in HPE OneView
- B. HPE MSA 2062 Storage Array is not supported in HPE OneView
- C. HPE OneView replaces array management tools for both types of array
- D. HPE ProLiant DL380 Gen10 Plus servers have an HPE OneView license included
- E. HPE Primera 650 can be managed using HPE OneView.

**Answer: (SHOW ANSWER)**

\* HPE MSA 2062 Storage Array is not supported in HPE OneView: HPE OneView does not provide management capabilities for the HPE MSA series, which includes the MSA 2062. These arrays must be managed using their dedicated management tools.

\* HPE Primera 650 can be managed using HPE OneView: HPE OneView supports the management of HPE Primera storage arrays, including the Primera 650, allowing for integrated management and monitoring within the OneView environment.

Reference:

HPE OneView Support Matrix

**NEW QUESTION: 36**

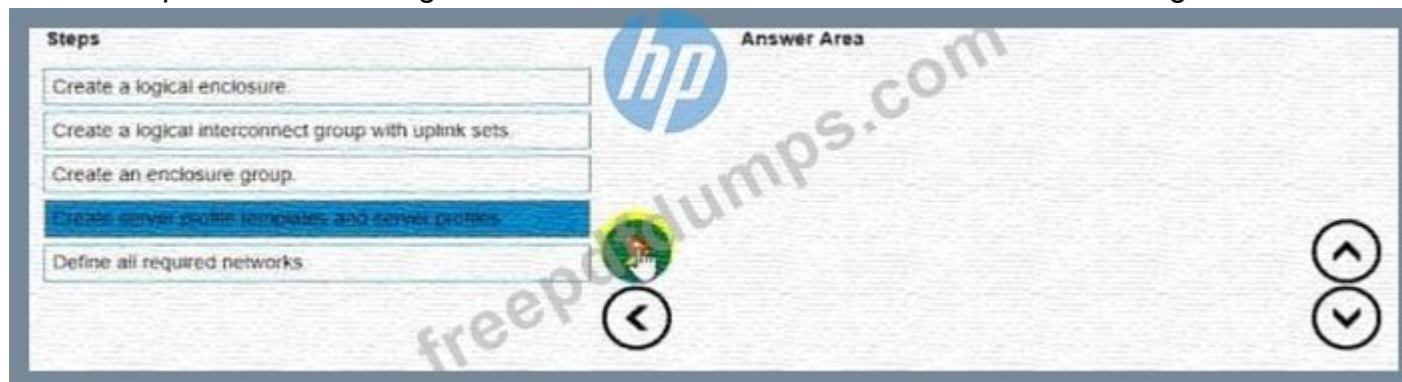
What technology do the HPE SN4000 Series (NVIDIA Spectrum) switches provide, which enables them to carry GPUDirect Storage (GDS)?

- A. Fibre Channel over Ethernet (FCoE)
- B. iSCSI
- C. InfiniBand
- D. RDMA over Converged Ethernet (RoCE) v2

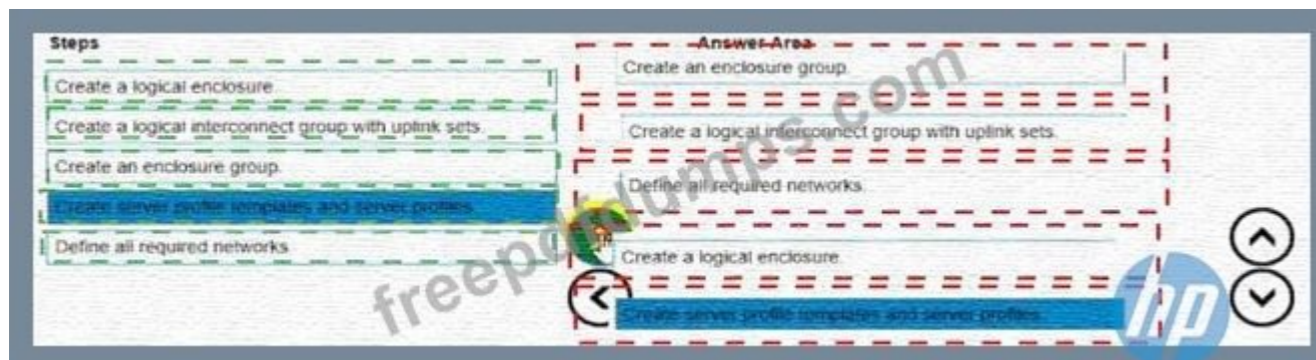
**Answer: D** ([LEAVE A REPLY](#))

**NEW QUESTION: 37**

You need to setup HPE Synergy frames after hardware components are discovered within HPE OneView. Put the steps to achieve this goal into the coned order in the answer area on the right.



**Answer:**



Explanation:

To set up HPE Synergy frames after the hardware components are discovered within HPE OneView, the following steps should be followed in the correct order:

- \* Create an enclosure group.
- \* Create a logical interconnect group with uplink sets.
- \* Define all required networks.
- \* Create a logical enclosure.
- \* Create server profile templates and server profiles.
- \* Create an enclosure group: This step involves defining the enclosure group, which includes the configurations for the enclosure such as firmware baselines and logical interconnect groups.
- \* Create a logical interconnect group with uplink sets: Next, set up the logical interconnect group, which includes the uplink sets that define the connectivity from the enclosures to the data center networks.
- \* Define all required networks: After setting up the logical interconnect group, define all the necessary networks that will be used within the HPE Synergy environment.
- \* Create a logical enclosure: Create a logical enclosure by combining the physical enclosures with the defined enclosure group, effectively binding the hardware configuration with the logical settings.

\* Create server profile templates and server profiles: Finally, create server profile templates and individual server profiles to apply the required configuration settings to the compute modules.  
Reference: HPE OneView User Guide

**NEW QUESTION: 38**

How does OpsRamp gather metrics across a customer ' s environment?

- A. OpsRamp agents must be installed on every server and storage device in the customer ' s environment.
- B. OpsRamp agents can be installed on supported Windows and Linux servers, and OpsRamp gateways can use protocols such as SNMP
- C. HPE compute and storage solutions automatically send telemetry information to global OpsRamp servers.
- D. Monitoring templates must be installed on each networking switch, determining what metrics are gathered and sent to the OpsRamp server.

Answer: B ([LEAVE A REPLY](#))

**NEW QUESTION: 39**

An administrator tries to open iLO interlace from HPE OneView Instead of being automatically logged in a login prompt displays Where should you start the troubleshooting process?

- A. At the iLO interface to verify whether the ILO Advanced license is properly applied
- B. At HPE OneView Interface to verify whether the server is In the maintenance mode
- C. At the iLO interface to verify whether the SSO certificate is removed from an iLO.
- D. At HPE OneView Interface to verify whether the iLO processor is disabled

Answer: ([SHOW ANSWER](#))

When the iLO interface prompts for a login instead of automatically logging in via HPE OneView, a common issue is that the Single Sign-On (SSO) certificate might have been removed or is not properly configured on the iLO. Verifying the presence and validity of the SSO certificate on the iLO interface is a critical step in troubleshooting this issue.

Reference: HPE OneView and iLO Integration Guide

**NEW QUESTION: 40**

Your customer plans to use HPE OneView for VMware vCenter Server to deploy ESXi systems. They created a server profile template that they want to use in HPE OneView. Put the steps required to configure HPE OneView for VMware vCenter Server on the left into their correct order on the right.

The screenshot shows a drag-and-drop interface with two columns: 'Steps' and 'Answer Area'. The 'Steps' column contains four items: 'Add HPE OneView credentials to the vCenter Server and create OS Deployment Plan', 'Add vCenter Server to HPE OneView for VMware vCenter Server', 'Deploy new ESXi systems using vCenter Server interface', and 'Upload HPE-friendly ESXi image'. The 'Answer Area' is currently empty.

Answer:

The screenshot shows the same drag-and-drop interface, but now the 'Answer Area' contains the steps in the correct order: 'Add HPE OneView credentials to the vCenter Server and create OS Deployment Plan', 'Add vCenter Server to HPE OneView for VMware vCenter Server', 'Upload HPE-friendly ESXi image', and 'Deploy new ESXi systems using vCenter Server interface'.

Explanation:

To configure HPE OneView for VMware vCenter Server to deploy ESXi systems, follow these steps in the correct order:

- \* Add HPE OneView credentials to the vCenter Server and create OS Deployment Plan.
- \* Add vCenter Server to HPE OneView for VMware vCenter Server.
- \* Upload HPE-friendly ESXi image.
- \* Deploy new ESXi systems using vCenter Server interface.
- \* Add HPE OneView credentials to the vCenter Server and create OS Deployment Plan: Integrate HPE OneView with vCenter Server by adding the necessary credentials and setting up the OS Deployment Plan.
- \* Add vCenter Server to HPE OneView for VMware vCenter Server: Ensure that vCenter Server is added to HPE OneView for VMware vCenter Server to enable communication and management.
- \* Upload HPE-friendly ESXi image: Upload an ESXi image that is optimized for HPE hardware to the vCenter Server for use in the deployment process.
- \* Deploy new ESXi systems using vCenter Server interface: Use the vCenter Server interface to deploy new ESXi systems based on the created server profile template and the uploaded ESXi image.

These steps ensure a seamless deployment and management of ESXi systems using HPE OneView integrated with VMware vCenter Server.

Reference: HPE OneView for VMware vCenter Server User Guide

#### **NEW QUESTION: 41**

You are troubleshooting an HPE OneView server profile that presents a critical state. The server profile was configured with a Link Aggregation Group. What should you check to fix the issue?

- A. if both FlexNICs are configured with the same speed
- B. If the logical enclosure is built on at least three HPE Synergy frames
- C. if both FlexNICs are connected to different networks
- D. if the logical enclosure is configured with redundant master modules

**Answer: A (LEAVE A REPLY)**

When troubleshooting a server profile in HPE OneView that is in a critical state due to a Link Aggregation Group (LAG), it is important to check if both FlexNICs are configured with the same speed. Inconsistent speeds can cause issues with link aggregation, leading to network instability and performance degradation.

Reference: HPE OneView Server Profile Troubleshooting Guide

#### **NEW QUESTION: 42**

Admins want to update the Mellanox kernel drivers running on the AI optimized workers in an HPE Private Cloud AI solution. What does HPE recommend?

- A. Using the OFED drivers downloaded from the HPE support site
- B. Updating the drivers as part of the HPE Private Cloud AI software catalog update
- C. Running the update from the HPE AI Essentials UI
- D. Updating the HPE Aruba Networking CX switches ' software first

**Answer: D (LEAVE A REPLY)**

#### **NEW QUESTION: 43**

Admins want to update the Mellanox kernel drivers running on the AI optimized workers in an HPE Private Cloud AI solution. What does HPE recommend?

- A. Running the update from the HPE AI Essentials UI
- B. Using the OFED drivers downloaded from the HPE support site
- C. Updating the drivers as part of the HPE Private Cloud AI software catalog update
- D. Updating the HPE Aruba Networking CX switches ' software first

**Answer: D (LEAVE A REPLY)**

#### **NEW QUESTION: 44**

Which challenge does distributed training with model parallelization address?

- A. Accelerating experimentation on servers with multi-core processors
- B. Avoiding drift by training multiple different models that check each other ' s results
- C. Fine-tuning models when data scientists are unsure which pretrained model will work best for their use case
- D. Training very large models that cannot fit on a single GPU

**Answer: D ([LEAVE A REPLY](#))**

#### **NEW QUESTION: 45**

What is a proper way to position HPE ProLiant DL145 servers?

- A. For providing NAS to a variety of clients, including ones running AI
- B. For fine-tuning pretrained AI models
- C. For adding Retrieval Augmented Generation (RAG) to pretrained LLMs
- D. For AI inferencing in physically challenging environments

**Answer: ([SHOW ANSWER](#))**

#### **NEW QUESTION: 46**

What is one component in an HPE Private Cloud AI solution that an admin needs to update separately from the HPE Private Cloud AI software catalog?

- A. The control plane nodes ' SPP
- B. The AI Essentials software
- C. The Red Hat Enterprise Linux OS running on the workers
- D. The HPE GreenLake for File Storage software

**Answer: ([SHOW ANSWER](#))**

HPE Private Cloud AI is designed as an integrated, full-stack solution, but it separates the lifecycle management of the " Compute/AI Stack " from the " Storage Stack. "

\* The Software Catalog: HPE provides a curated, tested, and versioned Software Catalog specifically for the Private Cloud AI environment. This catalog typically includes the firmware (SPP) for the worker nodes, the NVIDIA GPU drivers, the Kubernetes orchestration layer, and the HPE AI Essentials software (Option B). By updating the catalog, the admin ensures all these components remain compatible and supported as a single unit.

\* HPE GreenLake for File Storage (Option D): While this storage is a foundational component of the AI solution (providing the high-performance data lake), it is managed as a distinct service within the HPE GreenLake cloud platform. The software and firmware for the HPE Alletra Storage MP (which powers the file storage) follow their own release cycle and maintenance windows. Administrators must manage and trigger these updates through the HPE GreenLake for Storage portal, separate from the AI-specific software catalog orchestration.

\* Control Plane and Workers: The Control Plane nodes (Option A) and the Operating System on the workers (Option C) are generally included in the automated lifecycle managed by the Private Cloud AI solution ' s management plane to ensure the " turnkey " nature of the private cloud remains intact.

Key Technical Takeaway:

In a " Solution " context (like Private Cloud AI), components that are shared services or independent platforms (like a massive scale-out storage array) are often managed via their own dedicated lifecycle tools, even if they are sold as part of a single SKU/bundle.

**Valid HPE0-S59 Dumps** shared by Actual4test.com for Helping Passing HPE0-S59 Exam! Actual4test.com now offer the **newest HPE0-S59 exam dumps**, the Actual4test.com HPE0-S59 exam **questions have been updated** and **answers have been corrected** get the **newest** Actual4test.com HPE0-S59 dumps with Test Engine here: [https://www.actual4test.com/HPE0-S59\\_examcollection.html](https://www.actual4test.com/HPE0-S59_examcollection.html) (160 Q&As Dumps, **30%OFF** Special Discount: **Freepdfdumps**)

**NEW QUESTION: 47**

Which statement about HPE Apollo 2000 Gen10 Plus platform is true?

- A. It is equipped with HPE Persistent memory by default
- B. It can support a single processor only and up to 4 per chassis
- C. It is fully managed using HPE OneView including all connections
- D. It offers servers with AMD EPYC and Intel Xeon Scalable CPUs

**Answer:** [\(SHOW ANSWER\)](#)

The HPE Apollo 2000 Gen10 Plus platform offers flexibility in processor options, supporting both AMD EPYC and Intel Xeon Scalable CPUs. This allows customers to choose the best processor architecture based on their specific workload requirements and performance needs.

Reference: HPE Apollo 2000 Gen10 Plus QuickSpecs

**NEW QUESTION: 48**

You plan to configure a Link Aggregation Group (LAG) for two connections within a server profile for a compute node running an ESXi system.

What is required from the VMware side to enable LAG?

- A. Virtual Distributed Switch
- B. vSphere Standard license
- C. Load balancing cluster
- D. Dedicated VMKernel port

**Answer:** [A \(LEAVE A REPLY\)](#)

To configure a Link Aggregation Group (LAG) for two connections within a server profile for a compute node running an ESXi system, a Virtual Distributed Switch (VDS) is required from the VMware side. VDS allows for the configuration of LAGs and provides advanced network features such as centralized management and monitoring, which are essential for setting up and managing LAGs.

Reference: VMware vSphere Networking Guide

**NEW QUESTION: 49**

An admin is integrating HPE Compute Ops Management with Data Services Cloud Console to help automate connecting HPE servers to volumes hosted on HPE Alletra arrays. The admin is prompted to enter an API client ID and key.

Where should the admin obtain these?

- A. From the API credentials created from an account with Data Services Read Write access
- B. From the key created when the license for the HPE Alletra arrays was activated
- C. From API credentials created in an account with a Data Ops Manager Operator role
- D. From the key created when activating servers for HPE Compute Ops Management

**Answer:** [C \(LEAVE A REPLY\)](#)

**NEW QUESTION: 50**

Your customer has an HPE Synergy frame equipped with a D3940 Storage Module. Each of the compute nodes has the appropriate storage controller installed and two 12Gb SAS switches installed in the first fabric.

The customer reports that the compute nodes cannot access the storage module.

What will you verify first when troubleshooting this problem?

- A. If the D3940 module is imported in a managed state that allows volume provisioning.
- B. If the D3940 storage module is properly licensed through the HPE OneView interface.
- C. If the SAS logical interconnect is a part of the logical enclosure and has a healthy state.

D. If there is at least one iSCSI network configured that will provide access to the module.

Answer: ([SHOW ANSWER](#))

#### NEW QUESTION: 51

Which correctly describes the control plane of HPE Private Cloud AI?

- A. The control plane is distributed across all the worker nodes for redundancy.
- B. The Kubernetes-based control plane runs in HPE GreenLake cloud.
- C. Two worker nodes are elected to provide the control plane.
- D. Three HPE ProLiant servers host virtualized services to establish a redundant control plane.

Answer: B ([LEAVE A REPLY](#))

HPE Private Cloud AI is a key component of the HPE GreenLake for Private Cloud portfolio, specifically co-engineered with NVIDIA. The architecture is designed to provide a " cloud-like " experience on-premises.

To achieve this, HPE utilizes a distributed control plane model :

\* Management and Orchestration (The Control Plane): The management layer, which includes the Kubernetes orchestration, lifecycle management, and the user interface for provisioning AI workloads, is hosted in the HPE GreenLake cloud . This allows HPE to manage updates, monitoring, and security patches remotely as a managed service, reducing the operational burden on the customer.

\* The Data Plane (On-Premises): The actual compute power-consisting of HPE ProLiant servers (such as the DL380a or the newer Gen11/Gen12 NVIDIA-certified systems)-resides in the customer ' s data center. These are the " worker nodes " where the AI models are trained and inferred.

\* Connectivity: The on-premises infrastructure connects securely to the HPE GreenLake cloud control plane. While the compute and data stay local for performance, latency, and sovereignty reasons, the " logic " that dictates how those resources are sliced and managed stays in the cloud.

Why other options are incorrect:

\* Option A: Distributing the control plane across all worker nodes is a standard " vanilla " Kubernetes configuration but does not align with the " as-a-service " managed model of HPE GreenLake.

\* Option C: A two-node election would lack a " quorum, " making it unsuitable for high-availability control planes.

\* Option D: While some legacy or specific " Business Edition " private clouds used on-site management VMs, the Private Cloud AI architecture specifically leverages the HPE GreenLake cloud to provide a unified, scalable management experience across multiple locations.

#### NEW QUESTION: 52

You configured a virtual machine cluster using the HPE Storage Integration Pack for VMware vCenter.

Click the option in the SSMC main menu that you can use to verify whether a virtual storage operation is properly created on HPE Primera or 3PAR.

GENERAL	BLOCK PERSONA	STORAGE OPTIMIZATION	DATA PROTECTION	STORAGE SYSTEMS	FEDERATION & MIGRATION	SYSTEM REPORTER	SECURITY	VMWARE
Dashboard	Hosts	Adaptive Flash Cache	Remote Copy Configurations	Systems	Federation & Migration Configuration	Reports	Users	Storage Containers
Activity	Host Sets	Priority Optimization	Remote Copy Groups	Controller Nodes	Peer Motions	Threshold Alerts	LDAP	Virtual Machines
Schedules	Virtual Volumes		RMC Instances	Ports		Advanced Analytics	Roles	
Settings	App Volume Sets		Restore Points	Drive Enclosures			Connections	
	Virtual Volume Sets			Physical Drives			Domains	
	Common Provisioning Groups							
	Policies							

Answer:

GENERAL	BLOCK PERSONA	STORAGE OPTIMIZATION	DATA PROTECTION	STORAGE SYSTEMS	FEDERATION & MIGRATION	SYSTEM REPORTER	SECURITY	VMWARE
Dashboard	Hosts	Adaptive Flash Cache	Remote Copy Configurations	Systems	Federation & Migration Configuration	Reports	Users	Storage Containers
Activity	Host Sets	Priority Optimization	Remote Copy Groups	Controller Nodes	Peer Motions	Threshold Alerts	LDAP	Virtual Machines
Schedules	Virtual Volumes		RMC Instances	Ports		Advanced Analytics	Roles	
Settings	App Volume Sets		Restore Points	Drive Enclosures			Connections	
	Virtual Volume Sets			Physical Drives			Domains	
	Common Provisioning Groups							
	Policies							

Explanation:


To verify whether a virtual storage operation is properly created on HPE Primera or 3PAR using the HPE Storage Integration Pack for VMware vCenter, you would click on the "Virtualization" option in the SSMC (StoreServ Management Console) main menu. This option allows you to manage and monitor storage operations related to virtual environments, including VMware integration.

**NEW QUESTION: 53**

Click the port on the Frame Link Module which is used to include an HPE Synergy frame in the management ring.



Answer:



Explanation:

The port labeled " MGMT " (Management) in the upper right section of the HPE Frame Link Module (FLM) is the correct port used to include an HPE Synergy frame in the management ring. This port is crucial for linking the frame into the management network of the Synergy infrastructure.

**NEW QUESTION: 54**

A company manages their HPE ProLiant servers with HPE Compute Ops Management. Admins want to use this solution to deploy a Red Hat Enterprise Linux OS on a group of servers. What is one prerequisite?

- A. Admins must install HPE iSUT on the servers using iLO Intelligent Provisioning.
- B. Admins must ensure that the servers do not have any security risks found on them.
- C. Admins must obtain or create an Image that is set up for unintended, or automated, installation.
- D. Admins must prestage the ISO OS installation file in the servers ' iLO Repository.

**Answer: C (LEAVE A REPLY)**

#### **NEW QUESTION: 55**

What is a requirement for setting the CNSA security state on an HPE ProLiant server?

- A. The server uses iLO 7.
- B. The server has a certificate signed by a well-known CA installed on it.
- C. The server is managed by HPE Compute Ops Management.
- D. The server has an iLO Advanced license.

**Answer: D (LEAVE A REPLY)**

The Commercial National Security Algorithm (CNSA) security state (formerly known as Suite B mode) represents the highest level of cryptographic security for HPE ProLiant servers. Enabling this state enforces strict adherence to NSA-approved algorithms for all secure communications (including the web interface, SSH, and the RESTful API).

\* Licensing Requirement: The CNSA security state is classified by HPE as an Advanced Security feature . According to the HPE iLO 5, iLO 6, and iLO 7 Licensing Guides, the ability to configure and enable the CNSA security state is exclusively available with an HPE iLO Advanced license . Servers with only an iLO Standard license will not have the option to select this security state.

\* Prerequisites for Activation: Beyond the license, the CNSA security state can only be enabled if the FIPS (Federal Information Processing Standards) security state is already active. The process typically involves setting the server to FIPS mode, rebooting, and then elevating the security state to CNSA.

\* Architectural Support: While iLO 7 (Option A) introduces support for CNSA 2.0 algorithms, the CNSA security state itself has been available since HPE ProLiant Gen10 servers equipped with iLO 5. Therefore, using iLO 7 is not a requirement, as earlier iLO versions also support this state.

\* Certificate Nuance (Option B): While CNSA mode requires the use of highly secure certificates (specifically 384-bit ECDSA keys for SSL/TLS), having a certificate signed by a " well-known CA " (public CA) is not a prerequisite for enabling the state. Administrators can use internal CAs or even self- signed certificates, provided they meet the stringent CNSA cryptographic requirements.

Key Technical Takeaway:

If a customer requires their infrastructure to hold " Top Secret " classified data or comply with the most rigid government security standards, they must ensure their ProLiant fleet is equipped with iLO Advanced licenses to unlock the CNSA security state and the Silicon Root of Trust ' s most restrictive policies.

#### **NEW QUESTION: 56**

Your customer has the following HPE Synergy setup

- 3 HPE Synergy 12000 Frames
- 4 HPE Virtual Connect SE 100 Gb F32 Modules
- 2 HPE Synergy 50 GO Interconnect Link Modules
- 36 HPE Synergy 480 Gen10 Plus compute modules
- 2 HPE Composer 2 modules

Which statement about this setup is true?

- A. To use all compute nodes, the customer must create at least two logical enclosures
- B. To manage all frames, the customer must buy two additional HPE Composer 2 modules
- C. The customer must buy 36 HPE OneView Advanced licenses to manage compute modules
- D. The customer can add one D3940 storage Module to each of the frames

**Answer: A (LEAVE A REPLY)**

In an HPE Synergy environment, a logical enclosure represents a management domain that groups multiple Synergy frames and their associated compute, storage, and fabric resources. Given the setup with 3 HPE Synergy 12000 Frames and the specified components, the customer must create at least two logical enclosures to effectively manage all compute nodes and resources across multiple frames. This ensures proper resource allocation, management, and redundancy.

Reference: HPE Synergy Logical Enclosure Guide

**NEW QUESTION: 57**

What role does Prometheus play in HPE AI Essentials?

- A. It automates detecting drift in deployed models.
- B. It cleans up idle applications to make room for new applications.
- C. It gathers metrics and provides built-in alerting.
- D. It authenticates and secures communications between workloads.

**Answer: C (LEAVE A REPLY)**

In the HPE AI Essentials software stack-which provides the orchestration and management layer for AI workloads- Prometheus is the industry-standard component used for system observability.

\* Metric Collection (Scraping): Prometheus is responsible for "scraping" or collecting real-time numerical data (metrics) from across the environment. This includes hardware statistics from GPU-accelerated nodes (via the NVIDIA DCGM exporter) and performance data from Kubernetes pods.

\* Time-Series Database: It stores these metrics in a time-series format, allowing administrators to visualize performance over time and identify historical trends in resource consumption.

\* Alerting Framework: Prometheus includes a built-in alerting engine. Administrators can define specific thresholds (e.g., if a GPU temperature exceeds a certain limit or if a training job stalls). When these conditions are met, Prometheus generates an alert and forwards it to the Alertmanager for notification.

\* Infrastructure Health: By providing a unified view of the cluster's health, Prometheus ensures that the AI platform remains stable and that bottlenecks are identified before they impact model development.

**NEW QUESTION: 58**

Your customer has an HPE Synergy frame equipped with a D3940 Storage Module. Each of the compute nodes has an appropriate storage controller installed and two 12Gb SAS switches installed in the first fabric.

The customer reports that compute nodes cannot access the storage module. What should you verify first when troubleshooting this problem?

- A. If the D3940 module is imported in managed state that allows volume provisioning.
- B. If there is a logical interconnect created based on the 12Gb SAS switches.
- C. If the D3940 storage module is properly licensed through HPE OneView interface.
- D. If there is at least one iSCSI network configured that will provide access to the module.

**Answer: B (LEAVE A REPLY)**

When compute nodes cannot access the D3940 Storage Module in an HPE Synergy environment, the first thing to check is whether a logical interconnect has been created based on the 12Gb SAS switches.

Logical interconnects are essential for defining the connection and communication pathways between the compute nodes and the storage module. Without a proper logical interconnect configuration, the compute nodes will not be able to access the storage module.

Reference: HPE Synergy Configuration and Management Guide

**NEW QUESTION: 59**

Which statement about manual orchestration during the firmware update is true?

- A. It is supported only for logical enclosures with at least three frames
- B. It provides an ability to select and update one side of the Interconnect topology at a time
- C. It provides the ability to install an older firmware version than currently installed in the interconnect
- D. It allows all the interconnect modules to be activated manually at the same time.

**Answer: B (LEAVE A REPLY)**

Manual orchestration during a firmware update allows administrators to control the update process more precisely. Specifically, it enables the selection and update of one side of the interconnect topology at a time.

This capability is particularly useful in maintaining network availability and minimizing downtime during the update process, as one side can be updated and verified before proceeding to the other side.

Reference: HPE OneView Firmware Update Guide

**NEW QUESTION: 60**

An HPE partner is managing an HPE Private Cloud AI solution 's software as a service for the customer.

Which role or roles in HPE GreenLake cloud must the customer assign the partner?

- A. Private Cloud AI Administrator and Data Ops Manager Administrator
- B. Private Cloud AI Administrator only
- C. Private Cloud AI Cloud Administrator and Data Ops Manager Operator
- D. Private Cloud AI Cloud Administrator only

**Answer: A (LEAVE A REPLY)**

**NEW QUESTION: 61**

Admins want to use HPE Compute Ops Management to apply a custom BIOS/Workload profile to several servers. What is a required part of the setup?

- A. They must ensure that the servers ' BIOS settings are first set to factory defaults.
- B. They must use the HPE Compute Ops Management API to create the custom profile.
- C. They must use iLO to configure the servers to boot in optimized UEFI mode.
- D. They must ensure that the servers ' OS does not attempt to control any of the settings that the custom profile controls.

**Answer: (SHOW ANSWER)**

HPE Compute Ops Management (COM) , part of the HPE GreenLake cloud platform, allows for centralized management of BIOS settings and Workload Profiles across a fleet of servers. When implementing custom profiles, the interaction between the hardware (BIOS) and the Operating System (OS) is a critical architectural consideration.

\* OS vs. BIOS Conflict: Modern operating systems (such as Windows Server, VMware ESXi, and various Linux distributions) have built-in power management and performance scaling features. If the OS is configured to manage " P-States " or " C-States " (standard processor power/performance states), it will override the settings defined in the BIOS Workload Profile.

\* The " Required Setup " (Option D): For a BIOS/Workload profile to be effective, the administrator must ensure that the OS Power Management is set to a mode that defers to the BIOS (often called " BIOS Match " or " Static High Performance " within the OS settings). If the OS attempts to control these settings, the benefits of the custom profile applied via COM may be negated, or the server may exhibit inconsistent performance.

\* HPE Workload Profiles: These are pre-defined sets of BIOS settings optimized for specific tasks (e.g., Virtualization - Max Performance, Low Latency, etc.). When applying a Custom profile through COM, you are essentially fine-tuning these variables. Ensuring the OS is not " fighting " the BIOS for control is a prerequisite for a stable, optimized deployment.

Why other options are incorrect:

\* Option A: While starting from factory defaults can provide a clean baseline, it is not a technical requirement. The profile applied via COM will overwrite existing settings regardless of their current state.

\* Option B: Custom profiles can be easily created using the HPE GreenLake/COM GUI . While the API is available for automation, it is not a " required part of the setup. "

\* Option C: While most modern features require UEFI, specifically configuring " Optimized UEFI mode

" via the local iLO interface is not the primary requirement for applying a COM profile; COM itself is designed to handle these configurations remotely once the server is onboarded.

**Valid HPE0-S59 Dumps** shared by Actual4test.com for Helping Passing HPE0-S59 Exam! Actual4test.com now offer the **newest HPE0-S59 exam dumps**, the Actual4test.com HPE0-S59 exam **questions have been updated and answers have been corrected** get the **newest** Actual4test.com HPE0-S59 dumps with Test Engine here: [https://www.actual4test.com/HPE0-S59\\_examcollection.html](https://www.actual4test.com/HPE0-S59_examcollection.html) (160 Q&As Dumps, **30%OFF Special Discount: Freepdfdumps**)

**NEW QUESTION: 62**

What is one reason that organizations might prefer HPE Machine Learning Inference Software to KServe?

- A. HPE Machine Learning Inference Software supports autoscaling while KServe does not.
- B. HPE Machine Learning Inference Software 's batching and versioning capabilities exceed KServe ' s.
- C. HPE Machine Learning Inference Software is designed to run on Kubemetes, making it easy to deploy.
- D. HPE Machine Learning Inference Software ' s API makes it easy to automate model deployment.

**Answer: B (LEAVE A REPLY)**

**NEW QUESTION: 63**

What is one benefit of the infrastructure that underlies HPE GreenLake for File Storage within HPE Private Cloud AI?

- A. It is cost-effective and resilient, based on NVMe drives in HPE ProLiant servers.
- B. It provides 100% data availability.
- C. It is based on local drives on the AI-optimized nodes to accelerate data access.
- D. It is based on Lustre for massive scalability.

**Answer: C (LEAVE A REPLY)**

HPE GreenLake for File Storage is a key component of the HPE Private Cloud AI architecture, designed to provide the high-throughput, low-latency performance required for AI workloads.

\* Hardware Foundation: The infrastructure underlying this service is based on the HPE Alletra Storage MP platform. These modular " MP " (Multi-Protocol) nodes are built using standard HPE ProLiant Gen11 hardware components. By leveraging the proven ProLiant ecosystem, HPE ensures the solution is both cost-effective (due to global supply chain efficiencies) and highly resilient (utilizing enterprise-grade power supplies, fans, and management via iLO).

\* NVMe Performance: To handle the massive data requirements of AI model training and inference, the infrastructure utilizes all-NVMe drives. This provides the parallelism and speed necessary to feed hungry GPUs without the bottlenecks associated with traditional SAS or SATA storage.

\* Disaggregated Architecture: Unlike legacy systems, this infrastructure separates performance (compute nodes) from capacity (JBOFs/Storage shelves), allowing the environment to scale efficiently as data grows.

Why other options are incorrect:

\* Option B: While Alletra MP Block storage offers a 100% availability guarantee, the File persona is designed for extreme durability and high availability but is typically marketed with 99.9999% (six nines) availability rather than the absolute 100% block storage claim.

\* Option C: HPE Private Cloud AI uses shared storage (Alletra MP) rather than local drives on the worker nodes. This ensures data persistence and allows multiple GPU nodes to access the same datasets simultaneously.

\* Option D: HPE GreenLake for File Storage is powered by the VAST Data software stack, not Lustre.

While Lustre is common in classic HPC, the VAST-based HPE solution is preferred for modern enterprise AI due to its ease of management and superior data reduction capabilities.

**NEW QUESTION: 64**

Which compute node parameters are captured within server hardware type?

- A. installed operating system
- B. Size of the memory Installed
- C. Number of the CPUs

D. Mezzanine card configuration

Answer: D ([LEAVE A REPLY](#))

**NEW QUESTION: 65**

Which statement about the HPE Synergy Image Streamer is true?

- A. Image Streamer is used to deploy OS images to the internal disks of compute modules.
- B. Image Streamer can manage firmware and driver baselines across multiple frames.
- C. Image Streamer is responsible for the configuration of network settings across Synergy frames.
- D. Image Streamer can manage OS images, firmware, and drivers across multiple frames.

Answer: D ([LEAVE A REPLY](#))

**NEW QUESTION: 66**

Your customer plans to deploy HPE OneView for VMware vCenter Server together with HPE Storage Integration Pack for VMware vCenter.

The customer wants to use them to manage HPE Synergy Gen10 compute modules, HPE ProLiant Gen10 servers, and an MSA array.

Which statement about compatibility of the existing environment with the planned software components is true?

- A. HPE OneView for VMware vCenter Server does not support standalone HPE ProLiant servers.
- B. Not all of HPE Storage Integration Pack for VMware vCenter features are supported with MSA arrays.
- C. HPE ProLiant Gen10 systems require HPE OneView for VMware vCenter server licenses for each node.
- D. Both HPE Synergy Gen10 full-height compute modules plugins will require a license.

Answer: B ([LEAVE A REPLY](#))

**NEW QUESTION: 67**

Match HPE SimpliVity term with the definition.

Term	Definition
Arbiter	Non-hyperconverged servers that can exist in an HPE SimpliVity environment and provide additional compute capacity while consuming the storage provided by the hyperconverged nodes.
Compute node	Provides centralized management and events end-point for vCenter.
Management Virtual Controller	Integrates with VMware Distributed Resources Scheduler to ensure optimal placement of the VM compute resources.
Intelligent Workload Optimizer	Facilitates communication between nodes and resolves state conflicts to ensure service continuity.

Answer:

Term	Definition
Arbiter	Facilitates communication between nodes and resolves state conflicts to ensure service continuity.
Compute node	Non-hyperconverged servers that can exist in an HPE SimpliVity environment and provide additional compute capacity while consuming the storage provided by the hyperconverged nodes.
Management Virtual Controller	Provides centralized management and events end-point for vCenter.
Intelligent Workload Optimizer	Integrates with VMware Distributed Resources Scheduler to ensure optimal placement of the VM compute resources. Facilitates communication between nodes and resolves state conflicts to ensure service continuity.

Explanation:

Here are the correct matches for each HPE SimpliVity term with their respective definitions:

- \* Arbiter Definition: Facilitates communication between nodes and resolves state conflicts to ensure service continuity.
- \* Compute node Definition: Non-hyperconverged servers that can exist in an HPE SimpliVity environment and provide additional compute capacity while consuming the storage provided by the hyperconverged nodes.
- \* Management Virtual Controller Definition: Provides centralized management and events end-point for vCenter.
- \* Intelligent Workload Optimizer Definition: Integrates with VMware Distributed Resources Scheduler to ensure optimal placement of the VM compute resources.
- \* Arbiter: This component is essential for ensuring that the HPE SimpliVity nodes communicate properly and that any state conflicts are resolved, which is critical for maintaining service continuity.
- \* Compute node: These are non-hyperconverged servers that add additional compute capacity to the HPE SimpliVity environment. They utilize the storage resources provided by the hyperconverged nodes.
- \* Management Virtual Controller: This controller is responsible for centralized management and acts as the events end-point for vCenter, allowing for streamlined management of the SimpliVity environment.
- \* Intelligent Workload Optimizer: This tool integrates with VMware 's Distributed Resources Scheduler (DRS) to ensure that virtual machine (VM) compute resources are placed optimally across the infrastructure.

Reference: HPE SimpliVity 380 Data Sheet

### NEW QUESTION: 68

Match HPE Superdome Flex system with its specific feature.

Answer:

Explanation:

Here are the correct matches for each HPE Superdome Flex system with its respective feature:

- \* HPE Superdome Flex
  - \* Support for 32 sockets and 48TB of memory
  - \* Support for nPar with 16 sockets
  - \* HPE Superdome Flex 280
  - \* Maximum of 24TB memory
  - \* Minimum of 768 GB of memory
  - \* HPE Superdome Flex:
    - \* Support for 32 sockets and 48TB of memory: The HPE Superdome Flex system is designed for extremely large-scale workloads and can support configurations with up to 32 CPU sockets and 48TB of memory.
    - \* Support for nPar with 16 sockets: This system supports nPartitions (nPars) with configurations that can include up to 16 CPU sockets, providing flexible partitioning and resource allocation.
  - \* HPE Superdome Flex 280:
    - \* Maximum of 24TB memory: The HPE Superdome Flex 280 supports a maximum memory capacity of 24TB, suitable for large but not as massive workloads compared to the full Superdome Flex.
    - \* Minimum of 768 GB of memory: The minimum memory configuration for the HPE Superdome Flex 280 is 768 GB, ensuring a base level of performance and capacity for smaller deployments.
- Reference: HPE Superdome Flex QuickSpecs

**NEW QUESTION: 69**

A company needs to provide object storage to applications running on their HPE servers. Which HPE solution should you recommend?

- A. HPE Alletra MP X10000
- B. HPE Solutions with Weka
- C. HPE Solutions with Qumulo
- D. HPE Alletra MP B10000

**Answer: A (LEAVE A REPLY)**

The HPE Alletra Storage MP X10000 is the specific HPE solution designed to deliver high-performance, scale-out object storage . As HPE transitioned its storage portfolio to the modular " MP " (Multi-Protocol) hardware platform, it introduced distinct software-defined personas for different data types.

- \* HPE Alletra MP X10000 (The Object Specialist): This model is optimized specifically for unstructured data and massive-scale object workloads. It provides a cloud-native, S3-compatible interface, making it ideal for modern application development, active data lakes, and AI training sets. It leverages a disaggregated architecture that allows capacity and performance to scale independently.
- \* HPE Alletra MP B10000: The " B " series is designated for Block and File services. While it is a versatile platform providing mission-critical performance with 100% availability guarantees, it is not the primary recommendation when the requirement is specifically for an object storage service.
- \* HPE Solutions with Weka/Qumulo: While HPE partners with Weka and Qumulo for high- performance file systems that can support object protocols (like S3), these are typically categorized as " File-first " solutions or specialized third-party integrations. The X10000 is the first-party, flagship HPE Alletra branded solution for dedicated high-performance object storage.

Key Technical Advantages of X10000 for Object Storage:

- \* S3 Native API: Provides the standard interface required by modern applications for object access.
- \* Exabyte Scalability: Designed to handle massive growth in unstructured data without the performance bottlenecks found in traditional NAS.
- \* Cloud-Managed: Fully integrated into the HPE GreenLake cloud console for unified management alongside block and file services.

**NEW QUESTION: 70**

Your customer used to manage their HPE 3PAR arrays using HPE OneView Now they have replaced these arrays HPE Primera.

How does managing the new arrays compare to managing their old arrays?

- A. HPE OneView offers additional functionalities for HPE Primera like CPG management
- B. HPE OneView offers limited support for HPE Primera because of the firmware differences
- C. To manage HPE Primera using HPE OneView. a dedicated license is required

**D.** Management procedures and features are the same for both types of the array

**Answer:** [\(SHOW ANSWER\)](#)

When managing HPE Primera arrays using HPE OneView, the management procedures and features are largely the same as those for HPE 3PAR arrays. HPE OneView provides a consistent management experience across both storage platforms, allowing customers to easily transition from HPE 3PAR to HPE Primera without significant changes to their management workflows.

Reference: HPE OneView Storage Management Guide

**NEW QUESTION: 71**

Your customer experienced some problems caused by outdated HPE Superdome Flex firmware. Which update method should they use to avoid these issues in the future?

- A.** HPE OneView method for I/O firmware update procedure
- B.** RMC CLI to update I/O and HPE Persistent Memory firmware
- C.** HPE OneView IT HPE Persistent Memory firmware must be updated.
- D.** HPE SUM if HPE Persistent Memory firmware must be updated

**Answer:** **D** [\(LEAVE A REPLY\)](#)

HPE System Update Manager (SUM) is the recommended method for updating firmware on HPE Superdome Flex systems, especially when HPE Persistent Memory firmware needs to be updated. HPE SUM ensures that all firmware components are updated correctly and consistently, reducing the risk of compatibility issues.

Reference: HPE Superdome Flex Firmware Update Guide

**NEW QUESTION: 72**

A logical interconnect group can span on multiple HPE Synergy frames for which interconnect type?

- A.** Brocade 32GB Fibre Channel Switch Module for HPE Synergy
- B.** HPE virtual Connect SE 32 Gb FC Module for Synergy
- C.** HPE Virtual Connect SE 100 Gb F32 Module for Synergy
- D.** HPE Synergy 12 Gb SAS Connection Module


**Answer:** **C** [\(LEAVE A REPLY\)](#)

A logical interconnect group in HPE Synergy is used to define a consistent network configuration across multiple frames. The HPE Virtual Connect SE 100 Gb F32 Module for Synergy allows for logical interconnect groups to span across multiple frames. This is because the 100 Gb module supports high-speed connectivity and the necessary infrastructure to maintain consistent network configurations over multiple frames, which is essential for scalable and flexible Synergy environments.

Reference: HPE Synergy Logical Interconnect Groups

**NEW QUESTION: 73**

Refer to the exhibit.



A company has an HPE Synergy system and an HPE Alletra array, which are connected as shown in the exhibit. What is one recommended setting for the FC networks defined on HPE Synergy?

- A. Setting each network up as a direct-attach network
- B. Setting the network type on each network as Fibre Channel over Ethernet (FCoE)
- C. Assigning at least one port on each redundant VC FC module to the uplink set that carries each network
- D. Associating each network with a SAN Manager

**Answer: C** ([LEAVE A REPLY](#))

#### NEW QUESTION: 74

Your customer plans to deploy HPE OneView for VMware vCenter Server together with HPE Storage Integration Pack for VMware vCenter. The customer wants to use them to manage HPE Synergy Gen10 compute modules, HPE ProLiant Gen 10 servers and an MSA array.

Which statement about compatibility of the existing environment with the planned software components is true?

- A. Both HPE Synergy Gen10 full-height compute module plugins will require a license
- B. HPE OneView for VMware vCenter Server does not support standalone HPE ProLiant servers
- C. To use the required plugins HPE ProLiant Gen10 systems must be managed using HPE OneView
- D. HPE Storage integration Pack for VMware vCenter does not support MSA arrays

**Answer: (SHOW ANSWER)**

#### NEW QUESTION: 75

Your customer plans to add four HPE Synergy frames to an existing management ring. All installed frames are equipped with two 2-port FLM modules, while new frames will be equipped with 4-port FLM modules.

Which statement about mixing different FLM modules is true?

- A. Existing 2-port FLM modules must be replaced with 4-port FLM modules
- B. A management ring can contain mixed frame link module configurations
- C. A management ring with mixed FLM modules cannot have more than 12 frames
- D. Mixing different FLM modules is allowed if all of them have the same firmware version

**Answer: (SHOW ANSWER)**

In HPE Synergy, a management ring can indeed contain mixed frame link module (FLM) configurations. This means you can have a mix of 2-port and 4-port FLM modules within the same management ring without needing to replace the existing 2-port modules. This flexibility allows for gradual upgrades and scalability within the management ring.

Reference: HPE Synergy Management Guide

**NEW QUESTION: 76**

Match each HPE OneView network type with its definition.

**Answer:**

**Explanation:**

Here are the correct matches for each HPE OneView network type with their respective definitions:

- \* Internal network Definition: A network that does not utilize any uplink ports and is used for communication between servers in the same frame.
- \* Tagged network Definition: A network in which any tagged packets are dropped and forwarding is done by a MAC address.
- \* Tunnel network Definition: A network with a dedicated set of uplink ports used to pass a group of VLANs without changing the VLAN tags.
- \* Untagged network Definition: A network that contains multiple networks on a single uplink set that allows your customer to share uplinks with those networks.
- \* Internal network: This network type facilitates communication between servers within the same frame without utilizing uplink ports, providing internal data exchange.
- \* Tagged network: This network type drops any packets that are tagged and forwards packets based on their MAC address, ensuring only untagged traffic is processed.
- \* Tunnel network: This type of network uses a specific set of uplink ports to pass VLAN groups while preserving the VLAN tags, useful for maintaining VLAN configurations across different network segments.
- \* Untagged network: This network type enables the sharing of uplinks among multiple networks by placing them on a single uplink set, useful for optimizing network resource usage.

Reference: HPE OneView Networking Guide

**Valid HPE0-S59 Dumps** shared by Actual4test.com for Helping Passing HPE0-S59 Exam! Actual4test.com now offer the **newest HPE0-S59 exam dumps**, the Actual4test.com HPE0-S59 exam **questions have been updated and answers have been corrected** get the **newest** Actual4test.com HPE0-S59 dumps with Test Engine here: [https://www.actual4test.com/HPE0-S59\\_examcollection.html](https://www.actual4test.com/HPE0-S59_examcollection.html) (160 Q&As Dumps, **30%OFF Special Discount: Freepdfdumps**)

**NEW QUESTION: 77**

An HPE partner is working with a customer, who says that they lack a large enough IT staff to stay ahead of help desk requests. What should you explain that OpsRamp provides?

- A. Unified provisioning of server and storage resources in a hybrid environment
- B. An overview of HPE solutions, support contracts, and registered support cases
- C. AI-powered help in reducing alert noise and identifying root causes
- D. Automated generation of cases in HPE support

**Answer: C (LEAVE A REPLY)**

OpsRamp (now an HPE company) is a cloud-based IT Operations Management (ITOM) platform that utilizes Artificial Intelligence for IT Operations (AIOps) to simplify the management of complex, hybrid environments.

\* Solving the Staffing Challenge: For customers with small IT teams, the primary bottleneck is often " alert fatigue " -a constant flood of low-level notifications and help desk tickets that mask critical issues.

OpsRamp addresses this by automating the " first response " layers of IT operations.

\* Alert Noise Reduction: OpsRamp ' s AIOps engine uses machine learning to deduplicate and suppress redundant alerts. For example, if a network switch goes down, it might trigger 100 alerts from connected servers. OpsRamp identifies the commonality and presents the admin with a single " Inferred Incident " rather than 100 separate tickets.

\* Root Cause Analysis (RCA): By analyzing topology maps and event sequences, OpsRamp can automatically pinpoint the probable root cause of a service disruption. This significantly reduces the " Mean Time to Repair " (MTTR) because staff members do not have to spend hours manually correlating logs to find where the failure started.

\* Automation and Remediation: Beyond just detection, OpsRamp provides Runbook Automation , allowing a small staff to create policy-driven workflows that automatically fix common issues (like restarting a service or clearing a full disk) without human intervention.

Why other options are incorrect:

\* Option A: This describes the core functionality of HPE OneView or the provisioning aspects of HPE GreenLake , which focuses on the " Day 0/Day 1 " deployment of resources.

\* Option B: This refers to the My Insights and Software Dashboard features found in the HPE Support Center (often part of an HPE Tech Care contract).

\* Option D: While OpsRamp can integrate with ITSM tools to create tickets, the automated generation of HPE support cases specifically for hardware failures is a primary feature of HPE InfoSight and HPE Support Center (Auto-case) .

**Valid HPE0-S59 Dumps** shared by Actual4test.com for Helping Passing HPE0-S59 Exam! Actual4test.com now offer the **newest HPE0-S59 exam dumps**, the Actual4test.com HPE0-S59 exam **questions have been updated and answers have been corrected** get the **newest** Actual4test.com HPE0-S59 dumps with Test Engine here: [https://www.actual4test.com/HPE0-S59\\_examcollection.html](https://www.actual4test.com/HPE0-S59_examcollection.html) (160 Q&As Dumps, **30%OFF Special Discount: Freepdfdumps**)