Lead1Pass

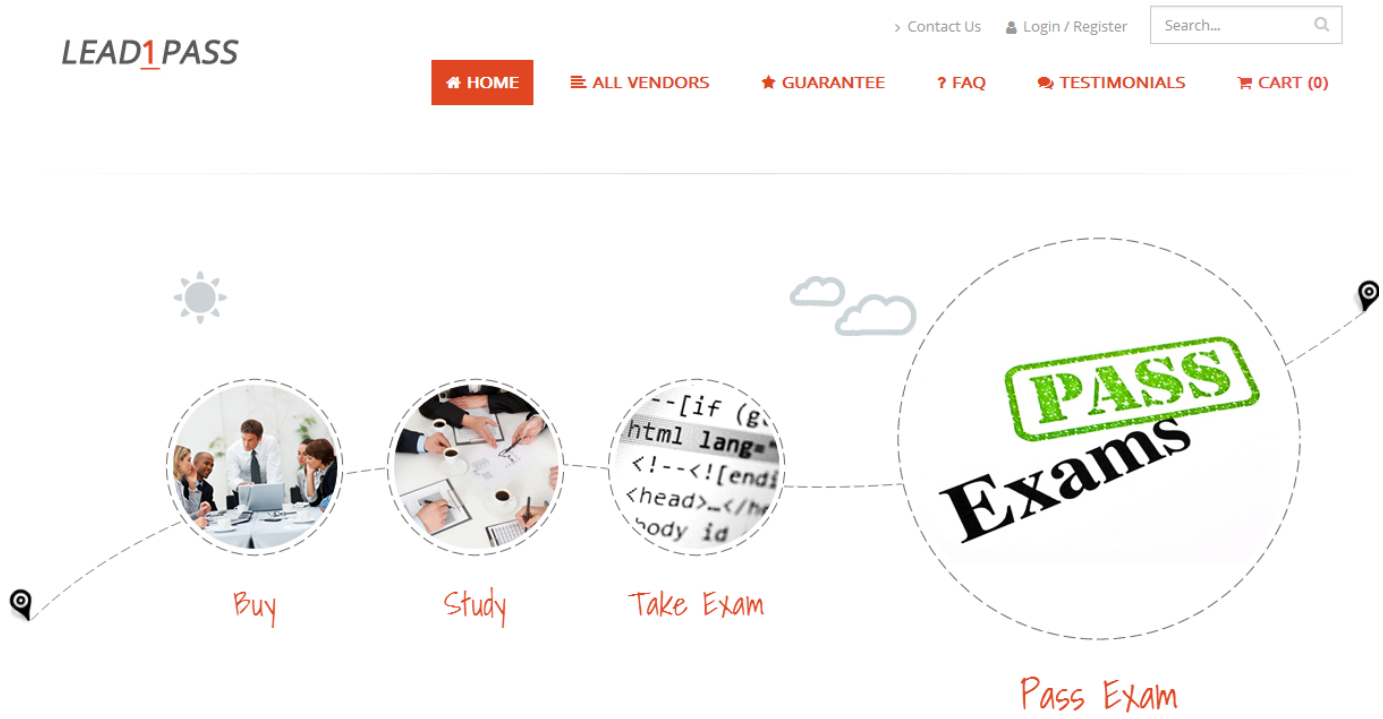

Try **PDF Demo** before you buy

Instant Download

After Payment, our system will send you the products you purchase in mailbox in a minute after payment. If not received within 2 hours, please contact us.

365 Days Free Updates

Free update is available within 365 days after your purchase. After 365 days, you will get 50% discounts for updating.

Money Back Guarantee

Full refund if you fail the corresponding exam in 60 days after purchasing. And Free get any another product.

Security & Privacy

We respect customer privacy. We use McAfee's security service to provide you with utmost security for your personal information & peace of mind.

http://www.lead1pass.com/

Latest Exam Guide & Learning Materials

Exam : **D-VXR-DY-01**

Title : Dell VxRail Deploy Exam

Vendor : EMC

Version : DEMO

NO.1 What is the minimum required network speed when using a vSAN 8 ESA with VxRail nodes?

A. 10 Gbps

B. 25 Gbps

C. 1 Gbps

D. 40 Gbps

Answer: B

Explanation:

When designing and deploying a Dell VxRail cluster utilizing VMware vSAN 8 Express Storage Architecture (ESA), network throughput requirements are significantly higher than traditional Original Storage Architecture (OSA) deployments. The minimum required network speed for all traffic-bearing network interfaces assigned to the vSAN ESA data path is strictly defined as 25 Gbps. This mandatory speed profile is dictated by the architectural modifications inherent to ESA, which replaces the traditional two-tier disk group model with a unified, high-performance single-tier storage pool built exclusively on NVMe-based flash media.

Because data paths are optimized for parallel disk accesses, activities such as inter-node replication, active write mirroring, and data erasure coding over the network domain can easily overwhelm lower-bandwidth infrastructure. Running vSAN ESA over 1 Gbps or 10 Gbps uplinks is structurally prohibited because network-induced latency would bottleneck the NVMe performance capabilities and trigger cluster instability during host failovers or data rebuild cycles. Therefore, 25 Gbps represents the absolute minimum speed requirement to pass pre-deployment hardware validation gates.

References: Dell VxRail Deploy Study Guide; vSAN ESA Hardware Prerequisites; Network Topology and Performance Requirements.

NO.2 An administrator physically installed the new nodes. Which two tasks must they complete before starting the deployment wizard? (Choose two.)

A. Upgrade the firmware to the appropriate version using iDRAC.

B. Validate the private management network VLAN.

C. Configure the iDRAC network settings.

D. Install the appropriate ESXi release using a VMware ISO.

Answer: B C

Explanation:

Before initiating the automated VxRail Deployment Wizard, specific network and administrative configuration steps must be executed immediately following the physical racking and cabling of the hardware nodes. First, the administrator must configure the iDRAC network settings (Choice C). This involves setting static IP addresses, updating default credentials, ensuring the interface is assigned to the dedicated management port, and enabling IPMI over LAN. Establishing the iDRAC network path allows deployment engineers to monitor hardware initialization states and perform out-of-band node validation tasks before the VxRail Manager cluster orchestrator takes control.

Second, the administrator must validate the private management network VLAN (Choice B) on the physical Top-of-Rack switches. VxRail uses an isolated VLAN-by default VLAN 3939-to run the Loudmouth discovery protocol using IPv6 multicast frames. If the switches are not properly tagged or trunked to allow this multicast domain traffic to pass between the nodes and the VxRail Manager, host discovery will fail.

Manually installing ESXi via a VMware ISO is unnecessary as nodes ship with a factory-installed

hypervisor image, and component firmware updates are natively orchestrated during or after initialization by the validated VxRail composite upgrade packages.
References: Dell VxRail Deploy Study Guide; Pre-Deployment Checklists and Node Preparation; Network Topology and VLAN Configuration.

NO.3 A customer will be using the Spanning Tree Protocol on Dell ToR switches in the VxRail environment. How should the switch ports that are connected to the VxRail nodes be configured?
A. Enable MLD snooping and querier.
B. Set the port mode to access.
C. Enable forwarding mode on the ports.
D. Set the port type to edge.
Answer: D
Explanation:
When implementing the Spanning Tree Protocol (STP) on Top-of-Rack (TOR) switches within a Dell VxRail environment, specific interface configurations are required to ensure uninterrupted host connectivity. Standard STP port states progress through listening and learning phases, which can introduce a propagation delay of up to 30 to 50 seconds before transitioning to a forwarding state. This delay can severely disrupt ESXi node discovery, initialization, and HA clustering during node reboots or link failover events, often causing the VxRail Deployment Wizard to fail.
To circumvent this, all switch ports directly interfacing with VxRail nodes must have their port type explicitly set to edge (Choice D). Configuring the port type as edge (the equivalent of PortFast in alternative switch operating systems) mandates that the interface bypass traditional STP convergence states and enter the forwarding state immediately upon link establishment. This prevents port-blocking cycles during critical boot sequences while preventing topology change notifications from destabilizing the local network tier. Setting ports to access mode is incorrect because VxRail mandates 802.1Q VLAN trunking to multiplex management, vSAN, and vMotion traffic over the same physical interfaces.
References: Dell VxRail Deploy Study Guide; ToR Switch Configuration Standards; Spanning Tree Protocol and Port Settings.

NO.4 An implementation engineer wants to extract a VxRail cluster configuration report to validate the environment before deployment. Where can the configuration report be accessed on the VxRail Configuration Portal?
A. My Clusters
B. My Deployments
C. My Projects
D. My Final Checklists
Answer: C
Explanation:
The Dell VxRail Configuration Portal is a centralized cloud-based utility designed to streamline the planning, architecture layout, and validation stages of a VxRail cluster deployment. Within this portal, all configuration tasks, parameters, and structural variables are organized inside a hierarchical system grouped under the "My Projects" workspace.
When an implementation engineer completes a cluster design or needs to extract a comprehensive VxRail cluster configuration report to execute pre-deployment validation, they must navigate

specifically to the "My Projects" tab. Inside this dashboard, users can review the end-to-end parameters of their defined deployments, perform validation checks against known configuration constraints, and download the definitive configuration reports along with the automated deployment JSON file. This documentation acts as the single source of truth for both the onsite deployment wizard and network alignment verification. Other workspace options listed, such as "My Clusters" or "My Deployments," represent distinct operational scopes or monitoring segments that do not host the primary pre-initialization project configuration export tools.
References: Dell VxRail Deploy Study Guide; VxRail Configuration Portal User Guide; Pre-Deployment Validation Procedures.

NO.5 An administrator has deployed a VxRail dynamic node cluster using a VxRail-managed vCenter. Which task must they complete?

A. Assign a license to the vCenter server

B. Configure the appropriate storage policies.

C. Update the vCenter to the latest version using VAMI.

Answer: B

Explanation:
A VxRail dynamic node cluster functions as a compute-only topology designed to utilize external shared storage resources instead of local vSAN storage pools. During the automated initialization phase of a dynamic node cluster leveraging a VxRail-managed vCenter, the deployment framework configures ESXi hosts, internal networking, vSphere Distributed Switches, and lifecycle management components. However, because dynamic nodes do not contain local capacity drives to assemble a local vSAN datastore, a primary storage repository is not automatically created or allocated for workloads.
Consequently, as a primary post-deployment milestone, the administrator must explicitly map external storage targets-such as Dell PowerStore, PowerMax, Unity XT, or an adjacent cluster via vSAN HCI Mesh-and subsequently configure the appropriate VM storage policies. Defining these storage policies is an absolute requirement to dictate availability, performance characteristics, and structure placement rules for virtual machine workloads across the externally attached array fabrics. Alternative tasks, such as provisioning a vCenter license, are natively orchestrated during deployment, while modifying components via the VAMI does not address immediate datastore requirements.
References: Dell VxRail Deploy Study Guide; VxRail Dynamic Nodes Design and Planning; Post-Deployment Storage Configuration Procedures.

NO.6 What must be enabled on ToR switch ports that are connected to VxRail nodes to support multiple VLANs?

A. Trunk Mode

B. Access Mode

C. Link Aggregation

D. Rapid Spanning Tree

Answer: A

Explanation:
Within a standard Dell VxRail cluster topology, all physical interfaces on the Top-of-Rack (ToR) switches directly attached to the physical node interfaces must be explicitly configured in Trunk

Mode. This configuration is an absolute architectural requirement because a hyperconverged VxRail deployment shares a pair of high-bandwidth uplinks across a matrix of distinct virtual networks. These networks include ESXi Host Management, vSAN Storage, vMotion Live Migration, and various customer-facing production VM networks, which must remain structurally isolated from one another at Layer 2.
To maintain this network design while consolidating traffic over shared cabling, the environment relies heavily on explicit IEEE 802.1Q VLAN tagging. Setting the switch ports to Trunk Mode allows the physical infrastructure to accept, process, and pass traffic carrying multiple separate VLAN identifiers across the same interface links simultaneously. Conversely, configuring a port to Access Mode limits its transmission capability to a single untagged broadcast network domain, which completely breaks the multi-VLAN isolation required to discover and initialize a VxRail cluster.
References: Dell VxRail Deploy Study Guide; ToR Physical Switch Networking Requirements; IEEE 802.1Q VLAN and Trunking Standards.

NO.7 When will local updates be the preferred option to update a VxRail cluster?

A. When using a customer provided vCenter.

B. When working at a dark site.

C. When a proxy server is required.

Answer: B

Explanation:
Local updates involve downloading the standalone VxRail composite upgrade bundle manually from the Dell Support Portal onto an administrative workstation and uploading it directly to the VxRail Manager local filesystem. This approach is highly preferred and often mandatory when working within a "dark site"-an enterprise datacenter environment that is completely air-gapped, isolated, or restricted from external internet access due to strict corporate security regulations, defense protocols, or compliance policies.
Under normal operating conditions, an internet-connected VxRail cluster can leverage Secure Connect Gateway to automate file acquisitions directly from Dell software repositories. However, in dark sites where external connection hooks or proxy exceptions are completely prohibited, the local update capability ensures that critical lifecycle management (LCM) operations, security patching, and platform enhancements can proceed seamlessly. The composite package uploaded locally contains all required hypervisor binaries, hardware microcode, and component driver updates matching the validated baseline state, eliminating external dependencies.
References: Dell VxRail Deploy Study Guide; Lifecycle Management and Cluster Upgrades; Dark Site Operational Procedures.

NO.8 A deployment engineer runs these commands across all nodes in a new VxRail deployment, and VLAN ID
3940 is added to the ToR switches:
esxcli network vswitch standard portgroup set -p "Private Management Network" -v 3940 esxcli network vswitch standard portgroup set -p "Private VM Network" -v 3940 None of the nodes are discovered by the VxRail Manager. What is the cause of this issue?

A. The VLAN ID for node discovery must be set to 3939.

B. The Loudmouth service was not restarted.

C. Changing VLAN IDs is only supported via a JSON file uploaded during the Deployment Wizard.

D. The Private Management Network and the Private VM Network cannot have the same VLAN ID.

Answer: B

Explanation:

In a fresh Dell VxRail deployment, node discovery is orchestrated by the proprietary Loudmouth service, which natively broadcasts over the default IPv6 multicast address using VLAN 3939. If an implementation engineer explicitly alters the default port group configurations on the standard virtual switches across the ESXi hosts to utilize an alternative VLAN ID, such as 3940, the operational state of the discovery daemon is not dynamically updated.

Although the underlying esxcli commands successfully reconfigure the "Private Management Network" and

"Private VM Network" port group parameters within the hypervisor network layer, the active Loudmouth process continues to bind to the previous network state. Because the Loudmouth service was not restarted following the command executions, it fails to initialize its discovery listens and broadcasts over the newly designated VLAN 3940 interface. Consequently, the VxRail Manager cannot locate or establish communication with the nodes over the modified Top-of-Rack switch fabric. To rectify this discovery failure, the administrator must issue a service restart command directly to the Loudmouth daemon across all target nodes, ensuring that the background discovery process initializes cleanly on the updated VLAN assignment.

References: Dell VxRail Deploy Study Guide; Loudmouth Service Architecture; Node Discovery and Network Validation.

NO.9 Which VxRail configuration meets the VMware vSAN ESA requirements?

A. E-Series - 256 GB of Memory, 3 x NVMe Drives, 10 Gbps Networking

B. E-Series - 384 GB of Memory, 5 x NVMe Drives, 25 Gbps Networking

C. V-Series - 256 GB of Memory, 5 x NVMe Drives, 25 Gbps Networking

D. P-Series - 96 GB of Memory, 6 x NVMe Drives, 10 Gbps Networking

Answer: B

Explanation:

VMware Express Storage Architecture (ESA) introduces strict hardware baseline criteria to optimize the performance, data path efficiencies, and resilience profiles of its single-tier storage pool model. To prevent network transport bottleneck issues during intensive inter-node write replication and erasure coding tasks, vSAN ESA establishes a rigid minimum network speed threshold of 25 Gbps per uplink, which instantly disqualifies any 10 Gbps network configurations.

Furthermore, vSAN ESA requires a minimum density of 4 NVMe storage devices per contributing node to satisfy concurrent read/write streams and properly support RAID-5 or RAID-6 erasure coding layouts. From a compute perspective, the storage architecture demands substantial system memory allocation to track complex block metadata variations and manage high-performance native snapshot engines, defining a minimum host memory baseline of 384 GB for lowest-tier operational profiles. An E-Series configuration featuring 384 GB of memory, 5 x NVMe drives, and 25 Gbps networking successfully satisfies all architectural pre-checks.

Profiles with lower memory quantities or insufficient drive counts fail validation gates during first-run deployment planning.

References: Dell VxRail Deploy Study Guide; vSAN ESA Hardware Prerequisites; VxRail Node Choice and Planning.

NO.10 What are three optional components in a Dell VxRail node? (Choose three.)

A. Graphical Processing Units
B. Fibre Channel Host Bus Adapters
C. Smart Data Processing Units
D. Auxiliary Power Unit
E. InfiniBand Network Adapters
Answer: A B C
Explanation:
Dell VxRail nodes feature a highly modular hardware architecture that can be customized to align with specialized enterprise workloads, storage protocols, and advanced distributed networking topologies. While core compute elements like primary server CPUs and system memory are standard requirements, several acceleration and connectivity expansion options are classified as optional components depending on the specific node model.
First, Graphical Processing Units (GPUs) (Choice A) can be optionally integrated into performance-focused or graphics-ready profiles-such as the VxRail V-Series-to efficiently offload intensive graphic rendering, virtual desktop infrastructure (VDI), and artificial intelligence or machine learning computational workloads from the host CPU. Second, Fibre Channel Host Bus Adapters (FC HBAs) (Choice B) are available as an optional add-on card, providing dedicated hardware links to connect nodes to external SAN storage arrays, which is a mandatory configuration for VxRail dynamic node topologies. Third, modern iterations of the VxRail hardware family support Smart Data Processing Units (SmartDPUs) (Choice C) paired with the VMware vSphere Distributed Services Engine. These SmartDPUs offload foundational core networking and security microservices directly onto the smartNIC silicon, freeing up server CPU cycles for consumer application VMs.
References: Dell VxRail Deploy Study Guide; Hardware Components of the VxRail System; Node Models and Options.

NO.11 The default permissions of the vCenter server administrator account have been identified as a security concern. Due to this security issue, the customer cannot use the default administrator account for the VxRail deployment. What should the customer do to eliminate the concern?
A. Create a custom role with full permissions. Create a user and assign the new role.
B. Create a custom role with limited permissions. Add administrator@vsphere.local to the new role.
C. Create a custom role with limited permissions. Create a user and assign the new role.
D. Create a custom role with full permissions. Add administrator@vsphere.local to the new role.
Answer: C
Explanation:
Enterprise security standards often restrict the use of default administrative credentials, such as administrator@vsphere.local, due to auditing, accountability, and least-privilege enforcement policies. When preparing for a VxRail deployment with an external, customer-managed vCenter, it is not mandatory to utilize the unrestricted default administrator account. To eliminate security vulnerabilities while facilitating proper cluster initialization, the customer should create a custom role within vCenter configured with limited permissions, then provision a dedicated user account and assign it to this new role.
Dell specifies a precise, restricted matrix of privileges required by the VxRail Manager service account to manage host attachment, storage provisioning, network alignment, and lifecycle updates. By building a custom role that contains only these mandatory minimum privileges and binding it to a newly created non- default service user, the organization enforces a secure operational boundary.

Modifying the permissions of the default administrator@vsphere.local account or providing full administrative scope to a new user fails to satisfy least-privilege security requirements and can inadvertently destabilize global vCenter management operations.
References: Dell VxRail Deploy Study Guide; vCenter Server Choice and Planning; Security Roles and Privileges Configuration.

NO.12 Which three actions are Day 1 Operations for VxRail API requests? (Choose three.)
A. Monitoring the cluster initialization status
B. Configuring and deploying a new cluster
C. Expanding clusters
D. Upgrading software
E. Scanning for available hosts
Answer: A B E
Explanation:
The VxRail REST API classifies orchestration tasks into separate operational boundaries, explicitly distinguishing between Day 1 (initialization and deployment) and Day 2 (ongoing lifecycle management) requests. Day 1 operations are exclusively bound to the automated processes required to bring a fresh, unconfigured set of bare-metal nodes up to a fully realized hyperconverged cluster infrastructure. This includes executing a discovery scan across the physical network segment to locate candidate nodes, invoking the deployment engine to programmatically apply the system design specification, and programmatically monitoring the initialization build sequence.
Specifically, the corresponding VxRail API endpoints utilized during this architectural phase include queries to discover unconfigured nodes on the local fabric, POST /v1/system/initialize to launch the automated deployment, and GET /v1/system/initialize/status to track execution milestones.
Alternative operations, such as expanding an operational cluster by adding nodes or executing automated software runtime upgrades, are explicitly classified as Day 2 lifecycle management workflows. These Day 2 tasks require a pre-existing, fully initialized production environment and an active integration with the vCenter management layer.
References: Dell VxRail Deploy Study Guide; VxRail REST API Architecture and Reference; Day 1 Initialization and Cluster Creation.

NO.13 Refer to the exhibit.

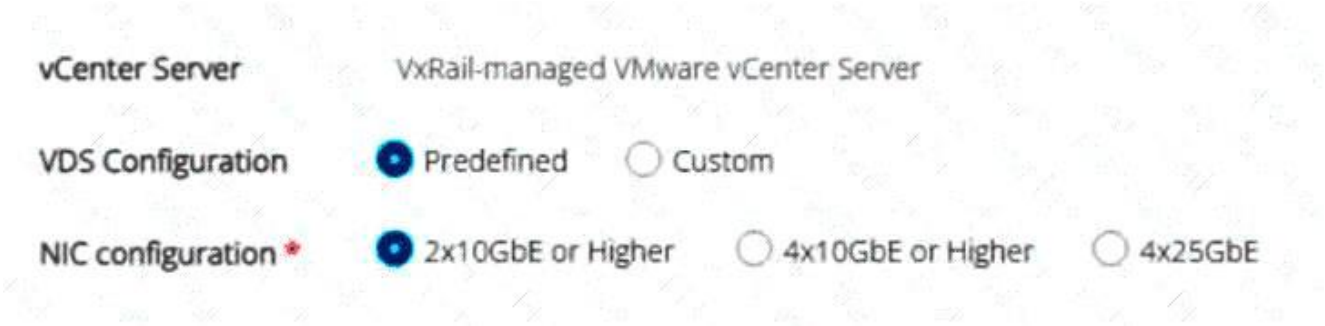

What teaming policy is applied to the vSAN port group?
A. Standby/Active
B. Active/Active
C. Unused/Active
Answer: A

Explanation:
In a VxRail deployment where the "Predefined" virtual distributed switch (VDS) configuration is selected alongside a "2x10GbE or Higher" dual-NIC layout, the automated installer applies deterministic network load- balancing and failover policies. Because a 2-NIC profile aggregates all foundational infrastructure traffic- including Management, vMotion, vSAN, and Virtual Machine networks-onto a single pair of physical uplinks (Uplink 1 and Uplink 2), logical segregation is critical to prevent contention and ensure data path reliability.
To achieve this, the predefined architecture utilizes an explicit failover ordering mechanism rather than a standard active/active load-balancing state across the uplinks. For the vSAN port group, which carries high- throughput storage synchronization traffic, the network design configures Uplink 1 as Standby and Uplink 2 as Active. This configuration maps directly to a Standby/Active setup.
Conversely, the Management port group operates in an inverse fashion, utilizing Uplink 1 as Active and Uplink 2 as Standby. This layout optimizes bandwidth usage across both physical connections during standard operations, while guaranteeing that full path redundancy remains maintained if an upstream switch link failover event occurs.
References: Dell VxRail Deploy Study Guide; Predefined VDS Network Topologies; Teaming and Failover Policies.

NO.14 What is a difference between compression and deduplication within a VxRail OSA cluster?
A. Compression is only applied on a per disk basis. Deduplication is applied to the entire disk group.
B. Compression requires significant resources to run. Deduplication requires marginal resources to run.
C. Compression requires reformatting of the disk group. Deduplication does not require reformatting of disk groups.
Answer: A
Explanation:
Within a VxRail cluster operating on the Original Storage Architecture (OSA), optimizing storage capacity is driven by software-defined data efficiency policies. Starting with vSphere 7.0 Update 1, vSAN introduced the option to enable "Compression Only" as a distinct toggle from the combined "Deduplication and Compression" feature set. A fundamental architectural distinction between these two mechanisms lies in their operational scope and data reduction boundaries.
When the compression-only feature is active, data reduction algorithms execute strictly on a per-disk basis.
This means blocks are evaluated and compressed locally before being written directly to that specific capacity drive, avoiding structural dependencies on adjacent drives or affecting the broader disk group cache-to- capacity relationship. In contrast, deduplication inherently requires a cross-device scope because it eliminates redundant data blocks across a larger logical boundary; hence, deduplication is applied to the entire disk group. For deduplication to occur, the hashing algorithms must track data signatures across all capacity drives tied to a specific caching drive within that given disk group. Consequently, deduplication carries a wider failure domain and distinct architectural requirements compared to localized per-disk compression profiles.
References: Dell VxRail Deploy Study Guide; vSAN OSA Space Efficiency Configurations; Storage Policy Design and Planning.

NO.15 Before performing a VxRail software upgrade, what report should be checked to ensure that no configuration drift exists due to manual updates?

A. Baseline Report on the Lifecycle Manager Baselines tab
B. Compliance Report on the VxRail Updates Compliance tab
C. Advisory Report on the VxRail Updates Internet Upgrades tab
D. Baseline Report on the Lifecycle Manager Updates tab
Answer: B
Explanation:
Maintaining configuration integrity across a VxRail cluster is imperative prior to executing automated lifecycle management (LCM) operations. Before initiating a comprehensive VxRail software upgrade, deployment engineers and administrators must evaluate the environment using the Compliance Report found specifically on the VxRail Updates Compliance tab (Choice B). This specialized report actively scans the entire hyperconverged infrastructure cluster, cross-referencing the active software, driver, and component firmware states against the definitive target VxRail Continuously Validated State baseline.
The core purpose of this evaluation is to proactively identify any instances of configuration drift. Configuration drift typically occurs when an administrator manually introduces unvalidated updates, installs standalone VIBs, or applies out-of-band firmware updates directly via the ESXi command line or iDRAC, bypassing the unified VxRail Manager interface. If drift is detected, the upgrade engine flags the non- compliant nodes, allowing engineers to remediate or reconcile the variations before launching the orchestration script, thereby preventing mid-upgrade installation failures or cluster state validation timeouts.
References: Dell VxRail Deploy Study Guide; VxRail Lifecycle Management Lifecycle Updates; Pre-Upgrade Health Check and Troubleshooting.

NO.16 Which three settings must be verified when using an existing VDS during VxRail setup? (Choose three.)
A. Names of VDS
B. Details of each port group
C. Permissions on the VDS
D. System traffic shares
E. Details of uplinks
Answer: A B E
Explanation:
When deploying a VxRail cluster utilizing an existing customer-managed Virtual Distributed Switch (VDS), the deployment framework shifts networking orchestration from automated switch generation to target environment integration. To guarantee seamless mapping of cluster network components, three primary settings must be rigorously verified within the VxRail configuration workflow: the exact Names of the VDS, the precise details of each port group, and the detailed configuration of the uplinks.
The name of the VDS must be explicitly matched because the initialization wizard relies on strict string matching to locate the target switch fabric inside the vCenter inventory. Furthermore, the details of each port group-including management, vMotion, vSAN, and VM networks-must be validated to ensure correct naming syntax and corresponding VLAN tagging structures are intact. Finally, the details of the uplinks must be verified, particularly the number of assigned active physical adapters and any pre-existing Link Aggregation Group (LAG) topologies. If mismatches exist among these structural properties, the wizard cannot successfully bind the virtual network interfaces of the

incoming ESXi hosts to the pre-established switch layers, precipitating initialization failure. Advanced switch policy variables like system traffic shares or specific object permissions are not structural binding metrics validated directly in the setup interface.
References: Dell VxRail Deploy Study Guide; Customer-Managed VDS Requirements; Advanced Network Configuration and VDS Integration.

NO.17 What is a difference between vSAN Data-in-Transit Encryption and vSAN Data-At-Rest Encryption?
A. Data-In-Transit Encryption: Does not require a Key Management Server.
Data-At-Rest Encryption: Requires a Key Management Server.
B. Data-In-Transit Encryption: Key Management Server must reside on the vSAN cluster.
Data-At-Rest Encryption: Key Management Server can reside on the vSAN cluster.
C. Data-In-Transit Encryption: Encrypts data after it is written to the datastore.
Data-At-Rest Encryption: Encrypts data as it moves between hosts.
D. Data-In-Transit Encryption: Encrypts traffic between VMs and the VxRail nodes.
Data-At-Rest Encryption: Encrypts traffic between VxRail nodes only.
Answer: A
Explanation:
Within a VMware vSAN-backed VxRail environment, cryptographic features are distinctly divided into Data- at-Rest Encryption and Data-in-Transit Encryption, each leveraging entirely different key exchange architectures. vSAN Data-at-Rest Encryption is implemented to safeguard data blocks residing within cache and capacity drives against physical media theft or unauthorized offline modifications. Because of its persistent state, Data-at-Rest Encryption strictly requires an integrated Key Provider solution-either an external, KMIP-compliant Key Management Server (KMS) or a VMware Native Key Provider-to generate, distribute, and protect the root Key Encryption Keys (KEKs) .
In contrast, vSAN Data-in-Transit Encryption secures the active data packets migrating across the dedicated local vSAN transport network between cluster hosts. Data-in-Transit Encryption does not require a Key Management Server or any external Key Provider infrastructure. Instead, the participating ESXi hosts within the cluster leverage internal secure tokens to dynamically generate, rotate, and pass ephemeral session keys directly between one another as peer connections form. Other answers incorrectly define the cryptographic boundaries of data traffic or specify arbitrary rules for KMS placement that contradict core vSAN storage security profiles.
References: Dell VxRail Deploy Study Guide; vSAN Security Architecture; Data Encryption Best Practices.

NO.18 What is the minimum and maximum number of nodes that can be selected to create a standard VxRail cluster in the deployment wizard?
A. 2 and 12
B. 2 and 32
C. 3 and 6
D. 3 and 64
Answer: D
Explanation:
During the architectural phase and the subsequent execution of the VxRail Deployment Wizard, strict

boundary constraints govern the node scaling of a standard VxRail cluster. The minimum and maximum number of nodes that can be selected to create a standard VxRail cluster are 3 and 64 nodes, respectively. A standard hyperconverged cluster requires a baseline minimum of three nodes during initial Day 1 deployment to establish a fully functional vSAN datastore that satisfies the default primary level of failures to tolerate (FTT=1) using a standard RAID-1 mirroring configuration.
While specialized deployment profiles like 2-node clusters exist, they require an external witness appliance and are classified as a distinct topology rather than a standard cluster path. On the upper boundary, the VxRail deployment framework scales in full alignment with VMware vSphere and vSAN core software limitations, allowing a single logical cluster to expand up to a maximum of 64 physical ESXi hosts. Adhering to these minimum and maximum boundaries ensures long-term operational stability and complete alignment with validated support matrices.
References: Dell VxRail Deploy Study Guide; Cluster Scaling Requirements; Deployment Wizard Configuration Boundaries.

NO.19 What are three prerequisites to deploy an additional VxRail cluster on an existing VxRail-managed vCenter Server? (Choose three.)

A. Existing SSO domain

B. Reserved IP addresses on the vCenter Server network

C. Preconfigured data center in vCenter

D. Preconfigured Management VLAN with ID 3939

E. Preconfigured forward and reverse DNS entries

Answer: A C E

Explanation:
Deploying a secondary or subsequent VxRail cluster onto an active, pre-existing VxRail-managed vCenter Server creates a multi-cluster architecture overseen by a single centralized instance. To ensure a successful deployment via the configuration portal and wizard, three rigid organizational and environmental prerequisites must be satisfied. First, the automated deployment wizard must join the incoming infrastructure directly to the existing Single Sign-On (SSO) domain (Choice A) established during the primary cluster's deployment, matching security and identity tokens across the management boundary.
Second, an administrative anchor point must exist prior to launching initialization; therefore, a preconfigured data center object (Choice C) must be explicitly defined within the vCenter inventory to house the new cluster and its associated resource groups. Third, network-level name resolution mapping must be established in advance, which requires preconfigured forward and reverse DNS entries (Choice E) on the customer's enterprise DNS server for the new VxRail Manager and all incoming ESXi host management interfaces.
Because the vCenter Server is already running and active on the network, reserved IP addresses for the vCenter itself are irrelevant, and management VLAN tracking applies to node-specific targets rather than global vCenter definitions.
References: Dell VxRail Deploy Study Guide; Multi-Cluster Deployment Prerequisites; vCenter Integration and Configuration.

NO.20 Which three actions cause a host "Incompatible" message to be displayed when expanding an OSA VxRail cluster? (Choose three.)

A. Adding a host with VxRail code 7.0.411 to a cluster with VxRail code 8.0.200

B. Adding a host with an AMD CPU to a cluster with Intel CPUs.
C. Adding a host with all NVMe SSDs to a cluster with all SAS SSDs
D. Adding a host with all SSDs to a cluster with HDDs
E. Adding a host with only 1 GbE NICs to a cluster with 10 GbE NICs
Answer: A B D
Explanation:
When expanding an existing VxRail Original Storage Architecture (OSA) cluster, the VxRail Manager performs stringent pre-validation checks to ensure that the incoming node conforms to the operational and architectural boundaries of the active cluster. Three distinct violations will immediately trigger a host
"Incompatible" status message within the expansion wizard. First, a major software version mismatch (Choice A), such as attempting to inject a node running VxRail code 7.0.411 into a cluster operating on 8.0.200, is strictly blocked because all nodes must share the same major release line to maintain lifecycle management integrity.
Second, a CPU vendor mismatch (Choice B), where an AMD-based node is introduced into an Intel-powered cluster, is fundamentally incompatible due to disparate processor instruction sets that break VMware vMotion requirements and vSAN cluster homogeneity. Third, mixing storage tiers (Choice D), specifically attempting to join an all-flash node (all SSDs) to a hybrid node cluster (utilizing HDDs for capacity), violates core vSAN deployment design configurations. While mixing drive interfaces like NVMe and SAS SSDs within an all- flash cluster is structurally acceptable under specific constraints, mixing drive types or processor architectures introduces deviations that fail validation.
References: Dell VxRail Deploy Study Guide; Cluster Expansion Guidelines and Restrictions; Host Compatibility Validation.

NO.21 Refer to the exhibit.

Which selection in the VxRail deployment wizard leads to the displayed locked option?
A. Customer-managed vCenter Server
B. vSAN 2-Node Cluster
C. Stretched Cluster
D. VxRail-managed vCenter Server
Answer: A
Explanation:
During the configuration phase of the VxRail Deployment Wizard, choices made within the foundational global settings dynamically govern subsequent system requirements and user interface behaviors. When an implementation engineer selects the option to deploy the cluster using a Customer-managed VMware vCenter Server, the system architecture shifts name resolution responsibilities entirely to the customer's pre-existing enterprise environment.
Because an external, customer-managed vCenter Server already relies on established enterprise datacenter infrastructure to resolve hostnames, lookup records, and management components, the internal VxRail Manager DNS service cannot be used to anchor the deployment. Consequently,

choosing a Customer- managed vCenter Server forces the wizard to automatically lock and restrict the DNS Server configuration field to "External," graying out the "Internal (VxRail Manager Service)" alternative option entirely. This built- in validation guardrail ensures that the incoming VxRail nodes and management components are integrated within the exact same authoritative external DNS servers handling the external management plane, preventing name resolution mismatches or deployment initialization failures. Other deployment options, such as using a standard VxRail-managed internal vCenter, leave this choice unlocked, permitting either internal or external placement.
References: Dell VxRail Deploy Study Guide; VxRail Deployment Wizard Configurations; vCenter Server Choice and Planning.

NO.22 An administrator wants to use the /rest/vxm/v5/system API to retrieve current cluster configuration information. Which HTTP method must be used for this API call?
A. GET
B. POST
C. PUT
D. DELETE
E. PATCH
Answer: A
Explanation:
The VxRail REST API platform adheres strictly to Representational State Transfer (REST) design principles, leveraging standardized HTTP methods to dictate operations against infrastructure resources. When an administrator intends to interact with the /rest/vxm/v5/system endpoint specifically to retrieve the current cluster configuration parameters, operational state, and component inventory details, they must invoke the GET HTTP method.
In RESTful web services, the GET verb is universally defined as a safe, idempotent read operation used exclusively to fetch a representation of a specified resource without altering its underlying state. Alternative HTTP methods are mapped to destructive or state-changing actions within the VxRail Manager API gateway.
For instance, POST is reserved for executing initialization workflows or generating new system objects, PUT and PATCH are used to modify existing resource configurations or update target properties, and DELETE handles object removals. Because retrieving configuration data is purely an information readout task, GET is the correct operational method.
References: Dell VxRail Deploy Study Guide; VxRail REST API Architecture and Reference; API Command Structures and Methods.

NO.23 Which two types of key provider are supported in vSAN ESA using a VxRail-managed vCenter Server?
(Choose two.)
A. Standard Key Provider with TPM 1.2 protected hosts
B. Native Key Provider with TPM 1.2 protected hosts
C. Native Key Provider
D. Standard Key Provider
Answer: C D
Explanation:

When deploying VMware vSAN Express Storage Architecture (ESA) within a Dell VxRail cluster environment managed by an internal, VxRail-managed vCenter Server, enforcing data security through cryptographic measures is fully supported. To enable vSAN Data-at-Rest Encryption, the platform requires an active key management infrastructure to handle keys securely. The architecture supports two primary types of key providers: the Native Key Provider (NKP) and the Standard Key Provider.
The Native Key Provider is an embedded, software-defined feature within vSphere that generates and manages keys internally, eliminating the requirement to deploy or license an external, third-party Key Management Server (KMS) system. The Standard Key Provider allows the internal vCenter Server to securely interface with external, KMIP-compliant enterprise Key Management Servers to manage key distribution.
Options mentioning Trusted Platform Module (TPM) 1.2 are invalid because modern vSAN ESA implementations demand advanced cryptographic sealing mechanisms that strictly require TPM 2.0 on the host hardware if hardware-backed protection is layered over the key provider architecture; TPM 1.2 fails validation checks.
References: Dell VxRail Deploy Study Guide; vSAN ESA Security and Encryption; Key Provider Configuration and Management.

NO.24 A VxRail Cluster has just been deployed. Use the VxRail simulator to determine the Service Tag, Model, and ESXJ IP Address of the first node - vcluster730.esx01.edu.local.
Note: It is necessary to dose (x) the simulator window before you can select a response to this question.

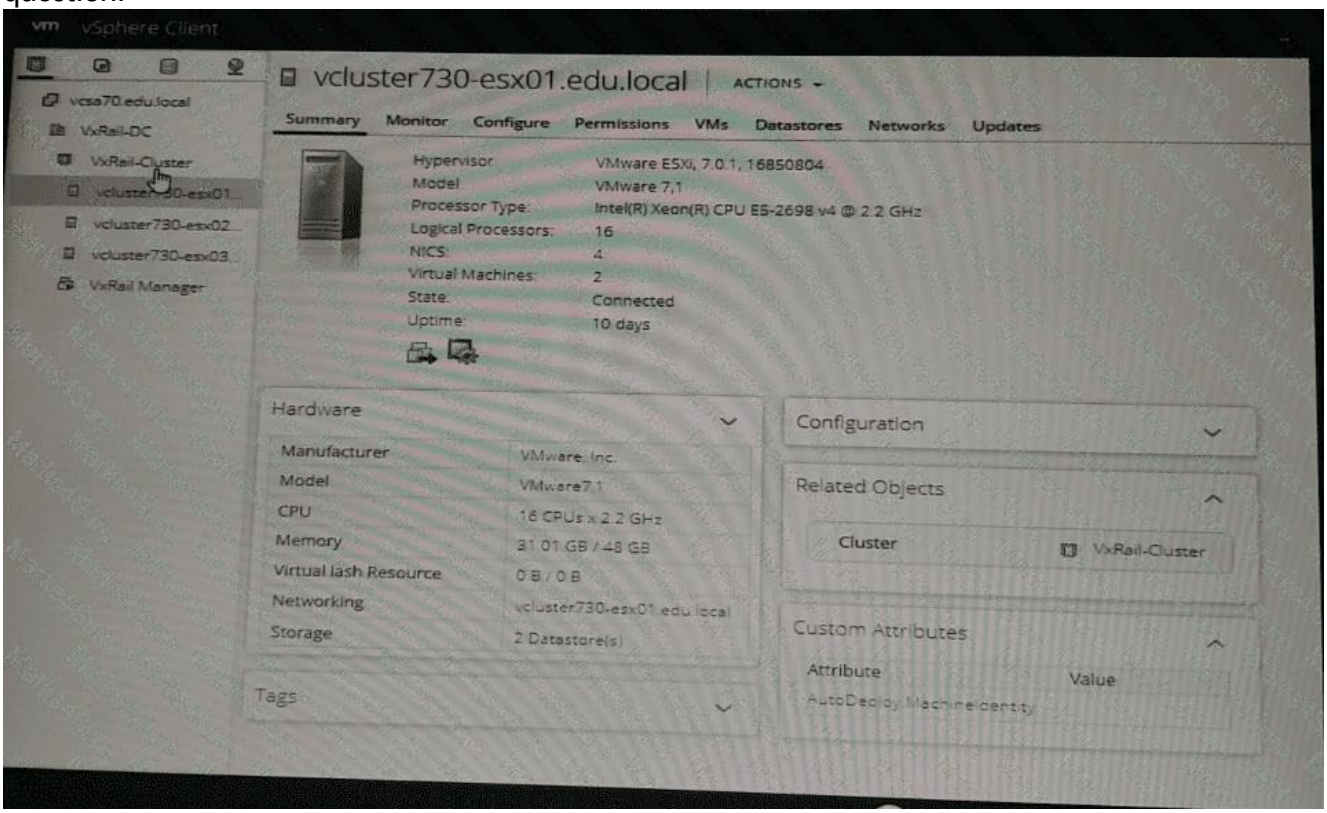

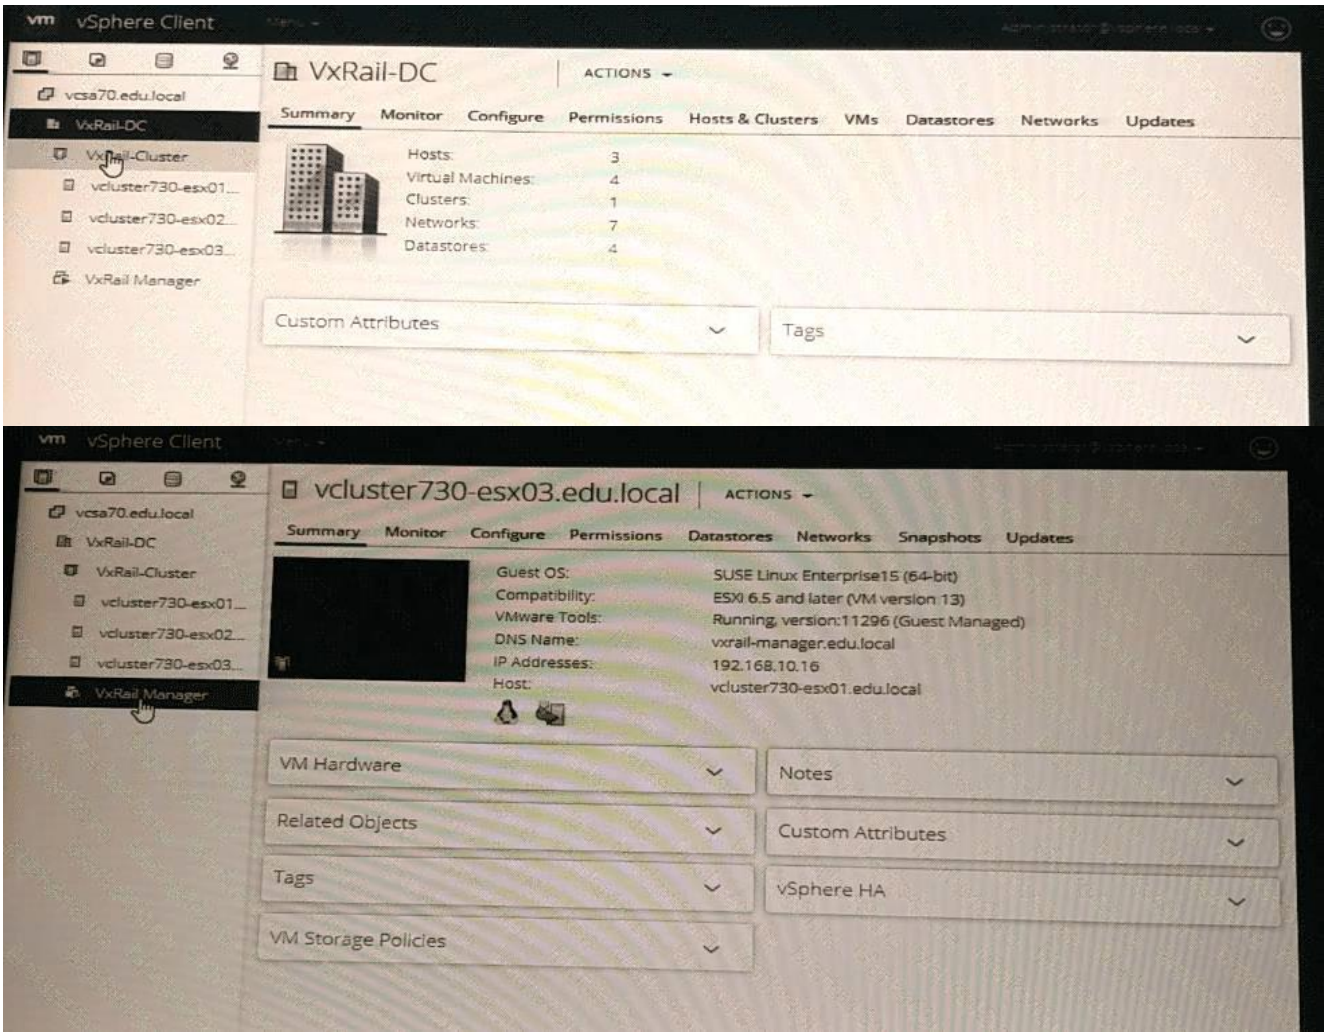

A. Option A
B. Option B
C. Option C
D. Option D
Answer: A

NO.25 Refer to the exhibit.

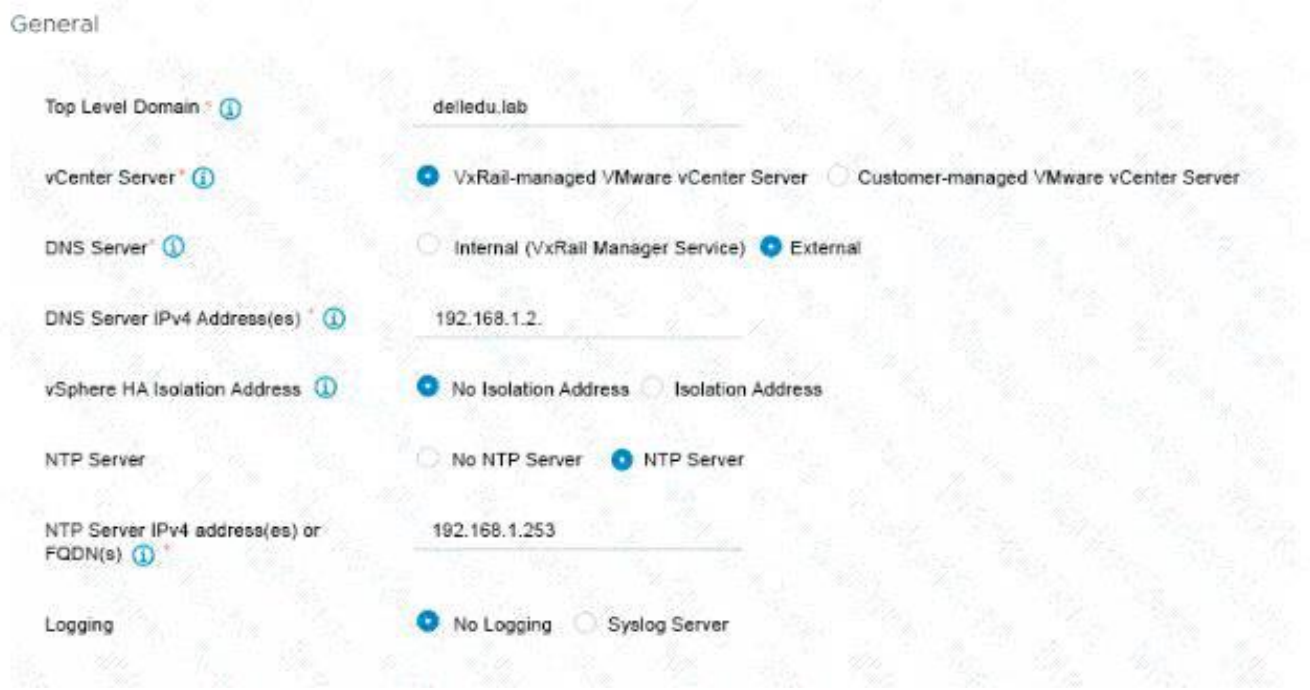

Which component is included in future VxRail Manager updates when the deployment of a VxRail Cluster is completed using the Global Settings shown?

A. DNS Server

B. NTP Server

C. vCenter Server

D. Windows Server

Answer: C

Explanation:

The exhibit illustrates the "General" settings section within the VxRail Deployment Wizard, where foundational infrastructure parameters are defined. Crucially, the radio button for "VxRail-managed VMware vCenter Server" is explicitly selected. In a VxRail deployment topology, choosing an internal, VxRail- managed vCenter server fundamentally changes how lifecycle management (LCM) operations are orchestrated for the cluster.

When a cluster is deployed with a VxRail-managed vCenter, the vCenter Server virtual appliance is tightly coupled with the VxRail Manager platform. This deep architectural integration guarantees that all future patch cycles, security definitions, and major platform upgrades for the vCenter Server are natively included within the unified VxRail composite update bundles. When an administrator applies an update package through the VxRail Manager dashboard, the orchestration engine automatically conducts pre-checks and applies the software versions to both the VxRail Manager and the vCenter Server in a fully automated, validated sequence. Conversely, infrastructure services configured as "External," such as the designated DNS or NTP servers, reside outside the VxRail ecosystem and are not targeted by VxRail lifecycle payloads. Therefore, the vCenter Server is the specific component captured in future update packages.

References: Dell VxRail Deploy Study Guide; VxRail Deployment Wizard Configurations; Lifecycle Management and Component Bundles.

NO.26 When following Dell best practices, what are two options for correctly installing VxRail nodes in a rack considering serial number order? (Choose two.)

A. Start from the bottom of the rack, ordered by serial number from highest to lowest.

B. Start from the bottom of the rack, ordered by serial number from lowest to highest.
C. Start from the upper portion of the rack, ordered by serial number from lowest to highest.
D. Start from the upper portion of the rack, ordered by serial number from highest to lowest.
Answer: B D
Explanation:
Dell best practices for physical rack deployment dictate specific structural and sequential guidelines when installing multiple VxRail nodes within an enterprise server rack. These methodologies ensure physical stability, optimize cooling efficiency, and streamline subsequent logical node mapping. The two approved physical installation configurations based on server serial numbers are Choice B and Choice D.
Under Choice B, engineers start installation from the bottom portion of the designated rack space, ordering the nodes by serial number from lowest to highest. This means the host with the lowest serial number (Node
1) is secured in the lowest physical position, with subsequent nodes stacked sequentially upward. Conversely, under Choice D, if the deployment architecture commands utilizing an upper portion of the rack first, the nodes are installed starting from the upper boundary ordered by serial number from highest to lowest. This specific configuration ensures that the lowest serial number (Node 1) remains positioned at the top of that specific node block, allowing the physical layout to maintain a logical sequence that maps cleanly to automated node discovery tools and management interfaces.
References: Dell VxRail Deploy Study Guide; Hardware Installation and Maintenance Procedures; Node Physical Racking Guidelines.

NO.27 What is a consideration when implementing an existing customer-managed VDS on a VxRail cluster?
A. Supports only two NIC uplinks.
B. Requires multiple switches to share each distributed port group.
C. Supports only four NIC uplinks.
D. Requires each LAG uplink to exist in the vCenter before implementing.
Answer: D
Explanation:
Utilizing a customer-managed Virtual Distributed Switch (VDS) allows enterprise organizations to pre-configure network settings, security policies, and custom port groups within their external vCenter before initiating the VxRail Deployment Wizard. When choosing this deployment pathway, particularly when Link Aggregation Control Protocol (LACP) is selected for network redundancy and load balancing, specific sequencing prerequisites must be followed.
If the architecture dictates the use of an LACP configuration to bundle physical network adapters, each Link Aggregation Group (LAG) uplink structure must be manually created and exist within the customer-managed VDS inside vCenter prior to running the deployment process. The VxRail Deployment Wizard cannot dynamically generate or inject new LAG configurations into an existing, externally managed VDS. During automated initialization, the VxRail Manager discovers these predefined LAG uplinks and maps the designated physical network interface cards (pNICs) of the incoming VxRail nodes to them. This ensures that high-bandwidth traffic such as vSAN and vMotion matches the upstream physical switch configurations immediately upon node discovery and initialization.
References: Dell VxRail Deploy Study Guide; Customer-Managed VDS Requirements; Advanced

Network Configuration and LACP Integration.

NO.28 Use the VxRail simulator to explore the system. What is the presented configuration? Note: It fs necessary to close (x) the simulator window before you can select a response to this question.

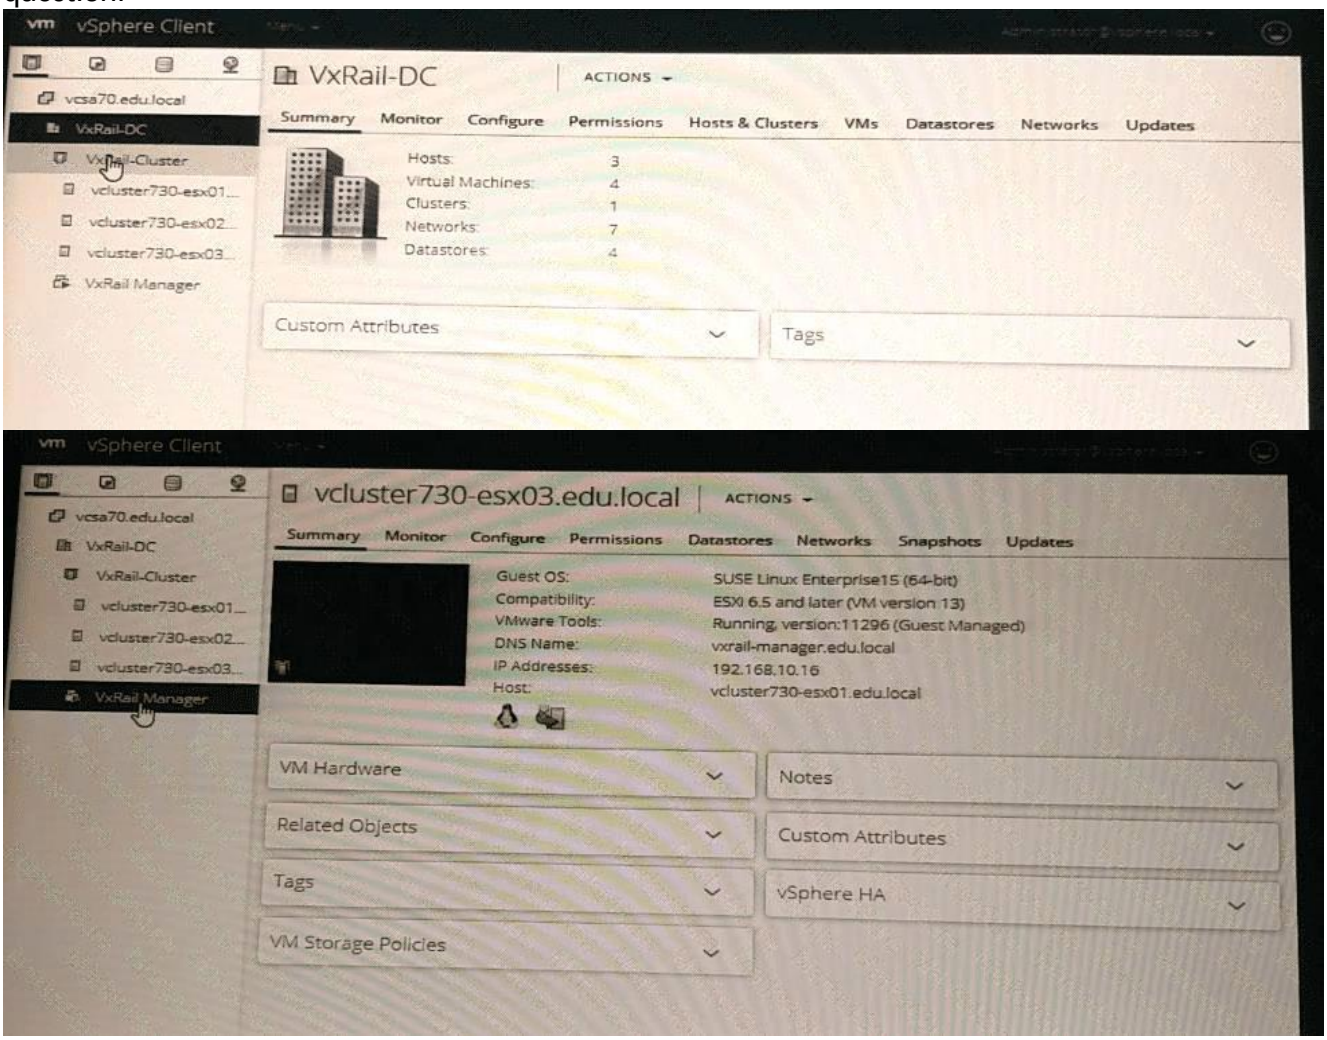

A. Internal vCenter and external DNS
B. Internal vCenter and Internal DNS
C. External vCenter and external DNS
D. External vCenter and internal DNS
Answer: A

NO.29 When following Dell best practices during an implementation, where do you assign the VxRail Manager permanent IP address?
A. All nodes
B. Any node
C. Primary node
Answer: C
Explanation:
When preparing for a Dell VxRail cluster implementation, establishing network connectivity and identity for the management plane is a key foundational step. During the automated initialization

process orchestrated by the VxRail Deployment Wizard or via a pre-configured configuration JSON template, the permanent management IP address designated for the VxRail Manager virtual appliance is assigned to the primary node.
The primary node-traditionally the first physical host in the rack sequence (Node 1)-acts as the bootstrap anchor point for the entire deployment workflow. The VxRail Manager virtual machine is initially stood up, configured, and run locally on this primary node's storage and compute resources before the cluster is fully formed and resources are aggregated into a shared vSAN datastore. Once the initial node setup completes and the hyperconverged cluster environment is established, the VxRail Manager VM participates in standard vSphere High Availability (HA) rules and can migrate across other nodes as needed, but its definitive configuration and initial IP binding are bound directly to the primary node profile during Day 1 tasks.
Assigning the permanent IP addresses randomly or globally across all nodes is incorrect, as only the specific management host entity receives the dedicated identifier.
References: Dell VxRail Deploy Study Guide; VxRail Architecture Overview; Day 1 Deployment and Initialization Procedures.

NO.30 What are three planning considerations when using a customer-supplied vCenter Server? (Choose three.)

A. Includes a standard vCenter license.

B. Requires an external DNS server.

C. Requires a separate vCenter license.

D. Uses customer-defined procedures for vCenter software upgrades.

E. Uses VxRail Lifecycle Management for vCenter software upgrades.

Answer: B C D

Explanation:
When an organization opts to deploy a VxRail cluster using a customer-supplied (external) vCenter Server instead of the embedded VxRail-managed virtual appliance, specific architectural and licensing boundaries must be accommodated during the design and planning phase. First, this model dictates that the customer must provide a separate, valid VMware vCenter Server license (Choice C), as the automated, bundled license included with the embedded deployment model is non-transferable to external management entities.
Second, because the infrastructure services are completely decoupled, the environment requires a fully functional, external DNS server (Choice B) with pre-established forward and reverse lookup records for all node management components, ESXi hosts, and the incoming VxRail Manager instance prior to initialization.
Third, the long-term maintenance paradigm shifts completely; the cluster uses customer-defined procedures for vCenter software upgrades (Choice D). This means the customer is independently responsible for patching and upgrading the external vCenter Server instance using standard VMware utilities like the vCenter Server Appliance Management Interface (VAMI) before initiating updates on the VxRail cluster components, effectively separating the external management platform from the automated VxRail Manager lifecycle management (LCM) framework.
References: Dell VxRail Deploy Study Guide; vCenter Server Choice and Planning; External Infrastructure Prerequisites.

NO.31 What is a software component of the VxRail HCI system software?

A. VMware Cloud Foundation

B. VMware Tanzu

C. Automation and Orchestration Services

D. VxRail Manager

Answer: D

Explanation:

The Dell VxRail HCI System Software is a proprietary suite of integrated management, automation, and orchestration capabilities designed to simplify the operational lifecycle of hyperconverged infrastructure.

Within this architectural ecosystem, VxRail Manager serves as the core software virtual appliance that contains and delivers these software capabilities.

VxRail Manager integrates directly with VMware vCenter Server to extend standard virtual environment oversight, offering an embedded control plane that manages automated hardware discovery, node additions, drive replacements, and non-disruptive software upgrades. While solutions like VMware Cloud Foundation or VMware Tanzu can be layered onto or integrated with the VxRail platform to provide cloud operating models and container orchestration, they are standalone VMware product portfolios rather than native software components of the core VxRail HCI system software suite itself. The automation and orchestration capabilities represent functional attributes delivered directly by the primary software component, which is the VxRail Manager appliance. Therefore, VxRail Manager represents the explicit internal engine driving the system software layer.

References: Dell VxRail Deploy Study Guide; VxRail HCI System Software Overview; VxRail Manager Architecture.

NO.32 An administrator is adding a node to an existing VxRail 8.0.100 vSAN ESA cluster. Upon starting the Add VxRail Hosts wizard, the node is listed as incompatible. Why is the node showing as incompatible?

A. The new node is running VxRail 8.0.000 software.

B. The ports on the ToR switch are not configured with the appropriate VLANs.

C. The network switch ports are cabled differently than the existing nodes.

D. The new node is a different model than the existing nodes.

Answer: A

Explanation:

When expanding an existing VxRail cluster utilizing the Express Storage Architecture (ESA), strict baseline validation rules govern node addition. A critical prerequisite for the automated "Add VxRail Hosts" wizard is the alignment of the VxRail software release train. Attempting to add a node running VxRail 8.0.000 code into an active cluster operating on a newer release like 8.0.100 will immediately flag the host as

"Incompatible" within the manager interface.

The automated cluster expansion framework mandates that the incoming host operate at a codebase version fully compatible with or exactly matching the existing cluster's active patch level to prevent lifecycle management vulnerabilities and operational drift. If a version mismatch exists where the target host operates on an older or disparate software baseline, the system cannot orchestrate uniform component updates or properly map vSAN ESA storage pools across the nodes. To resolve this incompatibility, the engineer must update the standalone node's runtime software environment to match the specific 8.0.100 cluster baseline before re-initiating the host expansion workflow. Other

infrastructure variations, such as switch tagging adjustments or cabling layouts, will not clear this software compatibility gate.
References: Dell VxRail Deploy Study Guide; Cluster Expansion Guidelines and Restrictions; Host Compatibility Validation.

NO.33 A VxRail API request has been performed, and the user account does not have vCenter privileges to perform the task. What code is returned?
A. 500
B. 207
C. 403
D. 200
E. 401
Answer: C
Explanation:
The VxRail REST API uses standard HTTP status code conventions to communicate the outcome of programmatic management commands. When a REST API request is submitted to the VxRail Manager appliance, the request undergoes a strict authorization verification step against the vCenter Server Single Sign- On (SSO) and Role-Based Access Control (RBAC) frameworks to ensure the calling entity possesses the required privileges.
If the user account successfully completes authentication (proving identity) but lacks the specific administrative permissions within vCenter to execute that particular infrastructure task-such as adding a node or altering cluster network properties-the VxRail API gateway intercepts the operation and returns an HTTP status code 403 Forbidden. This response signifies that while the server understands who the user is, the account does not possess the structural clearance or functional rights to execute the requested payload. In contrast, an HTTP 401 Unauthorized status is returned if the request fails authentication completely due to invalid or missing credentials. Status code 500 signifies a general internal server error, making 403 the definitive response for privilege omissions.
References: Dell VxRail Deploy Study Guide; VxRail REST API Security Architecture; HTTP Status Codes and Error Handling.

NO.34 An administrator used the script python vxm_backup_restore.py -b -l to back up the VxRail Manager. What is the result from this script?
A. Backup of the VxRail Manager and a list of backup copies
B. Back up of the VxRail Manager and existing services
C. Backup of the VxRail Manager and logs
D. Backup of the VxRail Manager only
Answer: A
Explanation:
The vxm_backup_restore.py script is an internal system tool designed to execute file-based backups, scheduling modifications, and recovery tasks for the VxRail Manager management appliance. When an administrator issues the command utilizing both the -b and -l flags simultaneously, the script evaluates and applies both parameter arguments in a sequential execution pipeline.
The primary -b argument initializes a manual, file-based backup immediately, compressing the core database records, identity tokens, and environment parameters into a secure target archive file pushed to the vCenter vSAN datastore storage directory. Simultaneously, the trailing -l parameter

commands the execution framework to poll the active archive path to retrieve and list the files. This outputs a detailed, chronological summary of all stored backup generations, and validation states directly to the terminal console. Running this flag combination ensures a fresh configuration snapshot is successfully committed while returning a clear inventory view of all current backup copies for operational confirmation.
References: Dell VxRail Deploy Study Guide; VxRail Manager Backup and Recovery; System Maintenance and File-Based CLI Tools.

NO.35 Which vSAN feature requires an all-flash cluster?
A. Data-at-rest encryption
B. Data-in-transit encryption
C. Storage policy-based management
D. Deduplication and compression
Answer: D
Explanation:
Within the VMware vSAN framework embedded inside the Dell VxRail infrastructure, certain storage efficiency features are structurally bound to the underlying media types. Specifically, deduplication and compression (Choice D) is a core space-saving technology that fundamentally requires an all-flash cluster deployment when utilizing the Original Storage Architecture (OSA). In an all-flash configuration, vSAN executes deduplication and compression at the block level during the destaging phase, as data transitions from the caching tier to the capacity tier.
Because these hashing and block-matching algorithms impose intensive mathematical processing overhead and demand rapid, random read/write lookups to maintain throughput without introducing system latency, VMware restricts this capability entirely to solid-state drives across both storage tiers. Attempting to run deduplication and compression on a hybrid node configuration utilizing spinning magnetic disks for capacity is structurally barred because traditional HDDs cannot support the high-performance random I/O required for metadata manipulation. In contrast, security and orchestration features such as data-at-rest encryption, data-in- transit encryption, and storage policy -based management (SPBM) are software-defined storage abstractions that operate uniformly across both hybrid and all-flash cluster environments.
References: Dell VxRail Deploy Study Guide; vSAN Storage Components and Capabilities; Space Efficiency Principles.