PassQueen
a.k.a dumpsqueen.com

VCE Vblock System Deployment and Implementation – Virtualization

VCE 210-015

Version Demo

Total Demo Questions: 12

Total Premium Questions: 123

Buy Premium PDF

https://passqueen.com

support@passqueen.com

Topic Break Down

Topic	No. of Questions
Topic 1, Installation	24
Topic 2, Configuration	28
Topic 3, Administration	24
Topic 4, Troubleshooting	27
Total	123

QUESTION NO: 1

You want to reset hardware sensors in the vSphere client connected to vCenter.

Which statement describes how to accomplish this task?

A. From the Hosts and Clusters view select a host, select the Hardware Status tab, and select the Reset sensors button.

B. From the Hosts and Clusters view select a host, select the Summary tab, select the Reset sensors button.

C. Select a virtual machine on the host, install the VMware tools, and reboot the host.

D. Hardware status can only be reset from the command line.

ANSWER: A

Explanation:

The correct procedure is: “From the Hosts and Clusters view select a host, select the Hardware Status tab, and select the Reset sensors button.” In the legacy vSphere Client connected to vCenter Server, hardware health information for an ESXi host is exposed through the host’s Hardware Status view. This view obtains its sensor inventory and health readings from the host-management stack and the server’s hardware-management interfaces. The Reset Sensors action requests that the client refresh and reinitialize the displayed hardware-sensor status, which is useful when readings are stale or when a sensor alert remains visible after the underlying condition has been addressed. The action applies to the selected ESXi host, so the administrator must first select that host in Hosts and Clusters and then open its Hardware Status tab. No guest operating-system action, VMware Tools installation, or host reboot is required for this client-side sensor reset operation. VMware’s documentation describes hardware health monitoring as a host-management capability and identifies the Hardware Status interface as the place where host sensor information is viewed and managed. See the VMware vSphere Monitoring and Performance documentation and the Broadcom VMware Knowledge Base for VMware platform guidance.

QUESTION NO: 2

A user mentions that they are unable to migrate a virtual machine, but the user is able to power down and power on the virtual machine. The option to migrate is grayed out.

When inspecting the user's roles and permissions, which two settings would cause this problem? (Choose two.)

A. The user has read-only rights to the virtual machine.

B. The user is part of a user group that has the privilege to migrate virtual machines.

C. The user does not have administration privileges.

D. The user's permissions does not allow the privilege to migrate virtual machines.

E. The user's role lacks the Virtual machine > Inventory > Move privilege.

ANSWER: D E

Explanation:

VMware vSphere uses granular, task-based privileges rather than requiring a user to hold a full administrator role. Therefore, a user can legitimately be authorized to power a virtual machine off and on while still being unable to start a migration. The permission set must include the applicable migration privilege; without it, the client disables the migration action instead of allowing the workflow to begin. The relevant privilege for moving a running virtual machine with vMotion is *Resource* > *Migrate powered on virtual machine*. In addition, migration requires the inventory move privilege for the virtual machine, because the operation changes its placement in the vCenter inventory and destination compute environment. Missing either authorization can prevent the user from initiating migration, even when standard power operations remain available. These privileges are normally assigned through a custom role and must be granted at the appropriate inventory levels, including the virtual machine and the applicable source and destination resources. VMware’s role-based authorization model and privilege assignments are described in the vSphere Security documentation. VMware also documents vMotion operational requirements in the vSphere migration documentation.

QUESTION NO: 3

Which two statements are true regarding VMWare Update Manager? (Choose two)

A. Update Manager contains, by default, multiple default baselines.

B. After installing the Update Manager plug-in, you need to immediately create baselines and baseline groups.

C. Update Manager can only upgrade virtual machines.

D. Baselines and baseline groups can be created to upgrade hosts, virtual machines, and virtual appliances.

ANSWER: A D

Explanation:

Update Manager contains, by default, multiple default baselines. These predefined baselines give administrators an immediate starting point for assessing and remediating applicable ESXi host patches, including patch classifications supplied through VMware's patch repository. A baseline defines a desired compliance level against which an inventory object can be scanned. This lets administrators identify hosts that require remediation without first having to build every baseline manually.

Baselines and baseline groups can be created to upgrade hosts, virtual machines, and virtual appliances. Update Manager uses different baseline types for its supported inventory objects and lifecycle tasks. Administrators can use host patch, extension, and upgrade baselines for ESXi hosts; virtual-machine upgrade baselines for VMware Tools and virtual hardware; and virtual-appliance upgrade baselines for supported appliances. Baseline groups combine compatible baselines so that related compliance checks and remediation actions can be managed together. This design supports a consistent, policy-based maintenance workflow: attach the relevant baseline or group, scan to determine compliance, and remediate the applicable objects during an approved maintenance window.

For VMware's Update Manager baseline concepts and supported remediation workflows, see the VMware vSphere Update Manager Administration documentation and the Broadcom VMware vSphere technical documentation portal.

QUESTION NO: 4

You power on a Vblock System 320 and look at the Maps tab of an ESXi server in vCenter's Hosts And Clusters view and see an ESXi server with a red "x" over it.

What is causing the ESX server problem?

A. The server has crashed.

B. The server is an unsuitable candidate for migration.

C. The data store is out of space.

D. The server is rebooting.

ANSWER: B

Explanation:

The server is an unsuitable candidate for migration. In the vCenter Maps view, the red "x" overlay is a migration-compatibility indicator: vCenter has determined that the ESXi host cannot currently be selected as a valid destination for a virtual-machine migration. This status is based on the compatibility checks vCenter and Distributed Resource Scheduler perform before allowing or recommending a migration. Those checks include whether the destination has access to the required virtual-machine networks and datastores, satisfies processor compatibility requirements, and meets applicable host or virtual-machine constraints. In a Vblock environment, the condition should be investigated as a configuration and compatibility issue affecting migration eligibility, rather than treating the overlay itself as a host hardware-state diagnosis. An administrator should review the migration compatibility details and the host's configured networking, storage presentation, cluster settings, and CPU compatibility configuration to identify the unmet requirement. VMware's resource-management documentation describes how DRS evaluates hosts as migration destinations and only migrates workloads where compatibility requirements are met. See the vSphere Resource Management documentation and the vCenter Server configuration documentation.

QUESTION NO: 5

A new VLAN has been configured on your Nexus 1000V for VM data traffic on VLAN 202 using the following commands.

vlan 202

name datavlan-202

port-profile type vethernet VM-DATA-202

vmware port-group

switchport mode access

switchport access vlan 202

no shutdown

max-ports 512

state enabled

After deploying a virtual machine on the new VLAN, connectivity outside of this network does not appear to work.

Which action will resolve the problem?

A. Apply the system VLAN attribute.

B. Use the switchport mode trunk command.

C. Add VLAN to the uplink port-profile configuration.

D. Increase default port count.

ANSWER: C

Explanation:

Add VLAN to the uplink port-profile configuration. The virtual Ethernet port profile correctly associates the VM's vNIC with access VLAN 202, but traffic can leave the ESXi host only through the Nexus 1000V uplink interfaces. Therefore, VLAN 202 must also be permitted on the uplink port profile that connects the host to the physical network. In a typical Nexus 1000V design, uplinks operate as trunks and their port profile defines the VLANs that are carried between the virtual switch and the upstream physical switch. Adding VLAN 202 to the allowed VLAN list on that uplink profile makes frames from the VM's access port available to the external network on the appropriate VLAN. The corresponding upstream physical-switch trunk must likewise carry VLAN 202, and VLAN 202 must be present where routing or gateway services for that subnet are provided. This creates an end-to-end Layer 2 path from the VM port group through the host uplink to the external network. Cisco's Nexus 1000V documentation describes configuring uplink port profiles and VLAN trunk membership; see the Cisco Nexus 1000V installation and configuration guides. For the underlying VLAN trunking behavior, see Cisco's VLAN trunk configuration guidance.

QUESTION NO: 6

Click the Exhibit button.

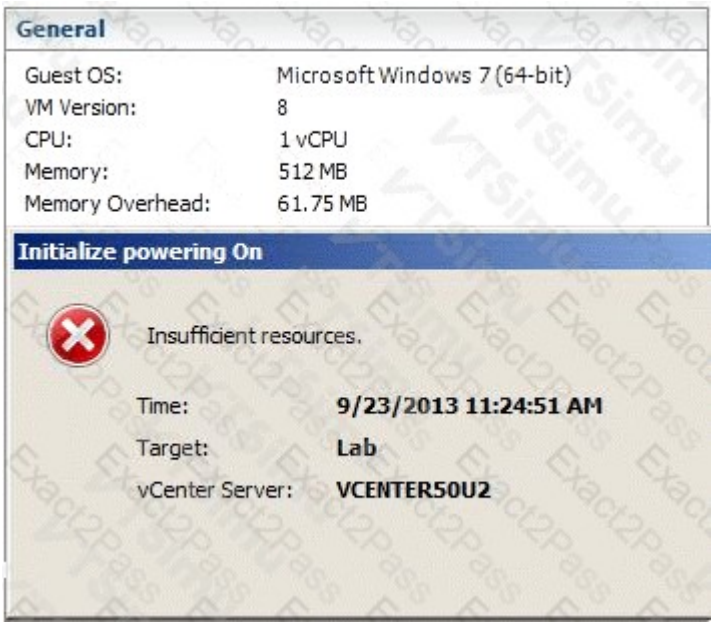

Referring to the exhibit, which two actions would allow powering on the virtual machine? (Choose two.)

A. Increase the memory reservation on the virtual machine.

B. Decrease the memory reservation on the virtual machine.

C. Add more memory to the ESXi host.

D. Increase the amount of memory allocated to the virtual machine.

ANSWER: B C

Explanation:

Decreasing the memory reservation on the virtual machine can allow it to power on because a reservation is a guaranteed amount of physical host memory that vSphere must be able to commit when the virtual machine starts. If the host does not have sufficient unreserved memory to meet that guarantee, lowering the reservation reduces the amount that must be immediately available and may satisfy the power-on admission check. Adding more memory to the ESXi host also resolves the capacity constraint by increasing the physical RAM from which ESXi can meet the virtual machine's reservation and other active memory commitments. These actions directly address an insufficient-memory condition at power-on while preserving vSphere's requirement that configured reservations be backed by available host resources. VMware documents that memory reservations guarantee physical memory to a virtual machine and are considered during resource allocation and admission-control decisions. See the vSphere Resource Management documentation and VMware's memory resource troubleshooting guidance.

QUESTION NO: 7

Which three components are part of the Vblock System Release Certification Matrix? (Choose three.)

A. VMware vSphere hypervisor ESXi

B. VMware vSphere Server

C. VMware vCloud Director

D. VMware vSphere Client

E. VMware vCenter Operations Manager

ANSWER: A B D

Explanation:

The Vblock System Release Certification Matrix identifies the interoperable, validated versions of the infrastructure software that make up a Vblock release. VMware vSphere hypervisor ESXi is included because it is the bare-metal virtualization layer installed on the Vblock compute hosts. VMware vSphere Server is included as the vCenter Server management component for the vSphere environment; it provides centralized inventory, host management, and operational control for the ESXi hosts. VMware vSphere Client is also included because the client release must be compatible with the vCenter Server and vSphere environment it administers. Together, these components represent the core VMware vSphere stack that must be version-aligned in a certified Vblock configuration. The matrix is important because Vblock certification is based on testing the complete integrated release, rather than treating individual infrastructure components as independently interchangeable. VMware's vSphere documentation describes ESXi and vCenter Server as core vSphere components and identifies the vSphere Client as the management interface: VMware vSphere documentation. Dell Support provides Vblock product documentation and release-related materials: Dell Support.

QUESTION NO: 8

While reviewing the Vblock System 320 configuration, you verify the physical NIC mode is configured for failover for a teamed vSwitch.

Which mode would be used to configure the NICs?

A. active

B. passive

C. promiscuous

D. unused

ANSWER: A

Explanation:

active is correct because physical uplinks that are intended to carry traffic as part of a teamed vSwitch are assigned to the active uplink group. In VMware vSphere NIC teaming, active physical adapters are the uplinks currently eligible to forward and receive traffic for the vSwitch or port group. The teaming policy monitors the availability of those uplinks and, when a failure is detected, preserves network connectivity by moving traffic according to the configured failover order. This is the fundamental behavior required when validating a Vblock System 320 vSwitch configuration that uses NIC teaming for uplink resiliency.

The active setting therefore identifies the NICs participating in normal production forwarding while allowing the vSwitch failover mechanism to react to a link, switch, or adapter failure. This configuration is commonly paired with multiple physical paths so the virtual networking layer has redundant connectivity available. VMware documents the role of active adapters and failover order in its vSphere networking and NIC-teaming guidance: Understanding network failover and load balancing in vSphere and vSphere Networking documentation.

QUESTION NO: 9

When installing the virtual Ethernet module (VEM) for a Nexus 1000V, which two software installation methods would be used for ESXi host configuration? (Choose two.)

A. Install with the VMware Update Manager (VUM).

B. Install the VIB remotely using the esxcli.

C. Install the VEM from the OVF template.

D. Install the VIB using the Nexus 1000V Web interface.

ANSWER: A B

Explanation:

The supported ESXi-host deployment approaches for the Cisco Nexus 1000V Virtual Ethernet Module are installation through VMware Update Manager and remote installation of the VIB with `esxcli`. VMware Update Manager provides a centralized, policy-driven mechanism for importing the VEM software package into a repository, attaching the relevant baseline to selected ESXi hosts, remediating those hosts, and tracking deployment compliance. This is particularly useful when the VEM must be applied consistently across multiple hosts in a vSphere cluster.

Remote VIB installation with `esxcli` is also a standard host-level method. An administrator can place the VEM VIB or offline bundle at an accessible location and use the ESXi command-line software installation workflow to install it on a target host. This method is appropriate for direct administration, scripted rollout, or environments where Update Manager is not being used. In either case, the VEM is installed as ESXi software and integrates the host networking with the Nexus 1000V distributed switch architecture. VMware's guidance for VIB deployment using ESXCLI is available in the Broadcom Knowledge Base; Cisco Nexus 1000V product documentation and resources are available from Cisco Support.

QUESTION NO: 10

An ESXi host is running a licensed version of PowerPathVE. A single path to a PowerPathVE-controlled device is lost.

Which two statements are true? (Choose two.)

A. The time required to restore a failed path on a heavily loaded system is less than the time required on a lightly loaded system.

B. After the path has been repaired, PowerPathVE will automatically restore the path and resume IO down the path.

C. The only way to restore a failed path is with the rpowermt restore CLI command.

D. The periodic testing of live and dead paths is low priority for PowerPathVE.

ANSWER: B D

Explanation:

After the path has been repaired, PowerPathVE will automatically restore the path and resume IO down the path. PowerPathVE continuously manages the available paths to a device and provides automated path failover and recovery. When a previously unavailable path becomes accessible and is validated for use, PowerPathVE can return it to the active path set without requiring an administrator to manually restore it. This behavior maintains availability during a transient path outage while allowing the repaired path to again contribute to load balancing and redundancy.

The periodic testing of live and dead paths is low priority for PowerPathVE. Path-health checking is deliberately performed as background work so that its probes do not compete significantly with production I/O. Its execution and the observed recovery timing can therefore be influenced by host activity and I/O workload; the essential point is that PowerPathVE performs ongoing, lower-priority path testing to discover restored connectivity and support automatic reintegration. Dell describes PowerPath/VE as multipathing software that provides automated path management for virtualized environments. VMware also documents that storage-path management relies on path state and path selection to maintain access to LUNs. See Dell PowerPath and VMware vSphere Storage documentation.

QUESTION NO: 12

While using the vSphere Client, you are experiencing high latency between issuing a search and obtaining a response from one of the four vCenter servers configured for LinkedMode. When using the vSphere Web Client, the inventory tree does not load properly and often shows objects out of date.

How would you improve the troubleshooting process

A. Update the logging level in vCenter to verbose, and restart the vCenter services.

B. Update the logging level for the Inventory Service to TRACE, and restart the Inventory Service.

C. Update the ESXi logging level to verbose, and restart the host.

D. Update the logging level for the Inventory Service to verbose, update the vCenter service to verbose, and restart all the services.

ANSWER: B

Explanation:

Update the logging level for the Inventory Service to TRACE, and restart the Inventory Service. The symptoms center on inventory retrieval, search responsiveness, and stale or incomplete inventory presentation in the vSphere Web Client. In the vSphere versions that used the separate Inventory Service, this service maintained an inventory model and supplied inventory-related data to the Web Client. Raising its log level to TRACE produces the most detailed diagnostic records for requests, processing activity, and failures affecting inventory synchronization or query handling. Restarting the Inventory Service applies the revised logging configuration and starts a clean diagnostic interval, making it easier to correlate a user search or inventory-tree refresh with the resulting service-side events. This focused approach is especially useful when only one vCenter Server in a Linked Mode group exhibits the issue, because it captures evidence at the affected inventory component without unnecessarily broadening logging across unrelated hosts or services. Review the generated Inventory Service logs alongside the timestamps of the slow searches and stale-object observations to identify communication, inventory-update, or processing delays. VMware's vCenter Server troubleshooting documentation provides the framework for collecting and reviewing vCenter diagnostic logs: vSphere Troubleshooting Documentation. Additional VMware knowledge resources are available through the Broadcom VMware Knowledge Base.