View All 5V0-31.22 Actual Free Exam Questions Apr 02, 2024 Updated [Q18-Q35

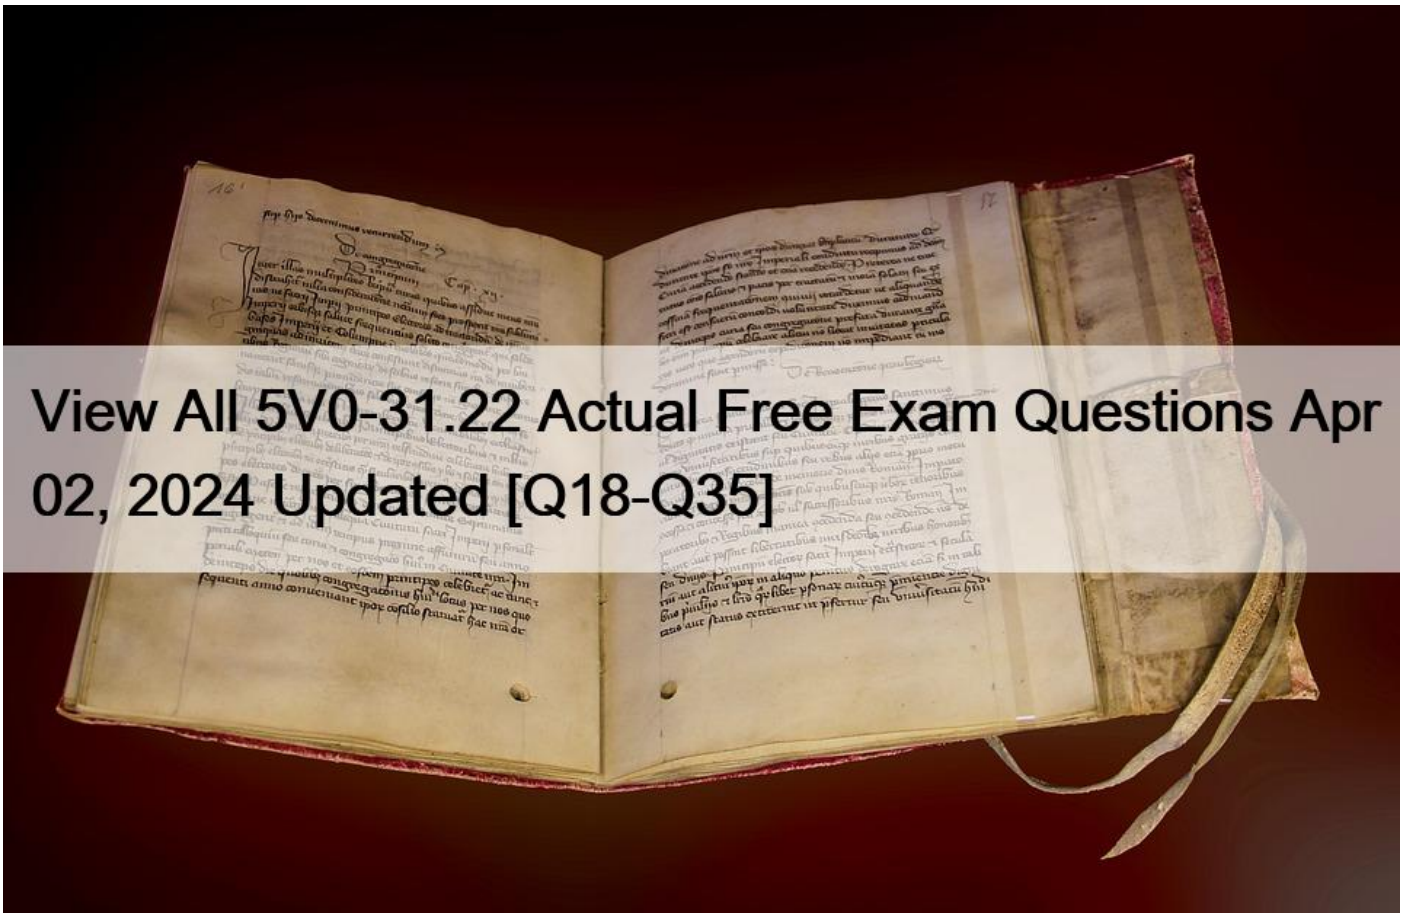

View All 5V0-31.22 Actual Free Exam Questions Apr 02, 2024 Updated
Pass Authentic VMware 5V0-31.22 with Free Practice Tests and Exam Dumps

Q18. Which three requirements are needed to provision an additional cluster within VMware Cloud Foundation? (Choose three.)
* Minimum of three hosts using vSAN as a principal storage
* Valid license specified in SDDC Manager
* All hosts in maintenance mode with no patches applied
* Minimum of three hosts using NFS, VMFS on FC, or vVols as a principal storage
* A new vCenter license to accommodate the additional cluster
* Minimum of five hosts for any type of principal storage

https://docs.vmware.com/en/VMware-Cloud-Foundation/5.0/vcf-admin/GUID-D3C55AA8-D4B9-49D4-A26F-7A713A141251.html

Q19. A VCF architect collected the following requirements when designing the expansion of a new VI Workload Domain with twenty four vSAN Ready nodes, each with a dual-port 25Gbps network interface card:

* Provide scalable high-performance networking with layer-3 termination at top-of-rack

* Protect workloads from switch/NIC/rack failure

* Provide isolation for DMZ workloads

* Provide at-least 25Gbps dedicated bandwidth to backup traffic

* Easily accept workloads on traditional VLAN-backed networks

* Fully-supported by VMware

Which three design considerations meet all of these requirements? (Choose three.)
* Two-node Edge Cluster with ECMP
* Spine and Leaf network topology with layer-3 at Spine
* Stretched Clustering
* Spine and Leaf network topology with layer-3 at top of rack
* Two-node Edge Cluster with BFD
* Core Aggregation network topology
Explanation

Option B: Spine and Leaf network topology with layer-3 at Spine – A spine and leaf network topology is designed for high scalability and performance, and layer-3 at the spine ensures that there is no single point of failure for the layer-3 termination. This meets several of the requirements, including scalable high-performance networking with layer-3 termination at top-of-rack, protecting workloads from switch/NIC/rack failure, and providing isolation for DMZ workloads.

Option D: Spine and Leaf network topology with layer-3 at top of rack – Similar to Option B, this topology also provides high scalability and performance, and layer-3 at the top of rack meets the requirement for layer-3 termination at top-of-rack.

Option F: Core Aggregation network topology – This topology provides a highly available, redundant core switch for aggregation and routing, which meets the requirement for protecting workloads from switch/NIC/rack failure.

Based on the given choices, the correct answers would be B, D, and F.

Sources: [1] Designing VMware Infrastructure Topology and Architecture; Authors: Russel Nolan, Eiad Al-Aqqad [2] Network Topology Considerations for VMware vSAN;https://docs.vmware.com/en/VMware-vSAN/7.0/com.vmware.vsan.networking.doc/GUID-1A901C10-48 Spine-Leaf Architecture:

Introduction;https://www.cisco.com/c/en/us/products/collateral/switches/nexus-9000-series-switches/datasheet-c

Q20. A VMware Cloud Foundation administrator is required to enable Workload Management (vSphere with Tanzu) on an existing workload domain cluster, which is currently licensed with a vSphere Enterprise Plus license.

Which action, if any, is required to complete this task?
* Add a license for vSphere with Tanzu with sufficient CPU capacity to the SDDC Manager inventory, and then assign the license to the cluster in SDDC Manager
* No action is required since SDDC Manager licenses include an entitlement for vSphere with Tanzu.
* No action is required since the vSphere Enterprise Plus license supports vSphere with Tanzu.
* Add a license for vSphere with Tanzu with sufficient CPU capacity to both the SDDC Manager and vCenter Server, and then assign the license to the cluster in vCenter Server

Explanation

https://docs.vmware.com/en/VMware-vSphere/7.0/vmware-vsphere-with-tanzu/GUID-9A190942-BDB1-4A19-B To enable Workload Management (vSphere with Tanzu) on an existing workload domain cluster, a license for vSphere with Tanzu with sufficient CPU capacity must be added to the SDDC Manager inventory, and then assigned to the cluster in SDDC Manager. This is because vSphere Enterprise Plus license does not include an entitlement for vSphere with Tanzu. Therefore, Option B and Option C are incorrect. Option D is also incorrect since adding the license to both the SDDC Manager and vCenter Server is not necessary to enable Workload Management.

Q21. An architect is designing networking for a developer-ready infrastructure on VMware Cloud Foundation.

During the discussion with the network team, a question comes up about the use of a routable CIDR range.

Which item uses this type of range?
* ClusterIP
* vSphere Pod
* Ingress
* Kubernetes services

Explanation

This is because an ingress is a Kubernetes resource that exposes HTTP and HTTPS routes from outside the cluster to services within the cluster1. An ingress can use a routable CIDR range to assign IP addresses to the ingress controllers that handle the traffic routing.

Q22. A VMware Cloud Foundation (VCF) architect is presented with a customer’s requirements for an architecture that needs to achieve:

* Network high availability across workloads in two data center locations.

* Maintain the least administrative overhead when performing day 2 operations.

* Decrease the RTO of both management plane and data plane when site-wide failure occurs Which VCF design consideration should the architect recommend?
* VCF with NSX-T Bridging
* VCF with NSX-T L2-VPN
* VCF with NSX-T Federation
* VCF with NSX-T Multi-Site

According to VMware Cloud Foundation Architecture Poster, VCF with NSX-T Federation provides network high availability across workloads in two data center locations by synchronizing network configuration and state across sites. It also simplifies day 2 operations by providing centralized management and policy enforcement across sites. It also reduces RTO by enabling fast failover of network services between sites.

Q23. An administrator has registered an external identity source in a consolidated architecture and would like to make sure that any subsequent workload domains can be accessed using the same identity sources.

How can this goal be achieved with VMware Cloud Foundation?
* By configuring IWA as an identity source
* By configuring LDAPS as an identity source
* By keeping the pre-configured defaults
* By replicating vSphere SSO configuration

Explanation

vSphere Single Sign-On (SSO) provides secure authentication and authorization services for VMware Cloud Foundation components, including vCenter Server and Platform Services Controller (PSC). In a consolidated architecture deployment of VMware Cloud Foundation, the vSphere SSO configuration is shared across all the workload domains.

To ensure that subsequent workload domains can use the same identity sources as an external identity source registered in a consolidated architecture, the administrator needs to replicate the vSphere SSO configuration.

This can be achieved by configuring the same identity sources for vSphere SSO across all the workload domains.

Configuring IWA (Integrated Windows Authentication) or LDAPS (Lightweight Directory Access Protocol over SSL) as an identity source is a part of configuring the vSphere SSO configuration for identity sources.

Keeping the pre-configured defaults does not guarantee that the subsequent workload domains will use the same identity sources as the external identity source registered in a consolidated architecture.

References:

* VMware Cloud Foundation Operations and Administration

Guide:https://docs.vmware.com/en/VMware-Cloud-Foundation/index.html

* VMware vSphere Security

Guide:https://docs.vmware.com/en/VMware-vSphere/7.0/vsphere-security-guide.pdf

* To ensure that any subsequent workload domains can be accessed using the same identity sources, it is necessary to replicate the vSphere SSO configuration across all the workload domains in a consolidated architecture deployment. This can be achieved by replicating the vSphere SSO configuration between the primary and additional SDDC Manager instances. This ensures that all the workload domains registered with the SDDC Manager will be able to consume resources and services from the same identity sources without any additional configuration in each individual workload domain.

Q24. During a VCF design workshop, the architect gathered the following customer requirements:

* There must be two environments: PROD and DEV.

* PROD and DEV should be administratively separated.

* PROD will use two different hardware server types, and DEV will only use one hardware server type.

* The VCF infrastructure design should be flexible and scalable as much as possible How many NSX local managers in total will be provisioned after deploying the full VCF infrastructure?
* 6
* 3
* 12
* 9
Explanation

According to the VMware documentation, each NSX-T Local Manager is associated with a vCenter Server, and each NSX-T Local

Manager can manage up to three vCenters. In a VCF deployment with two environments (PROD and DEV) and two different hardware server types in PROD, there would be a total of three vCenter Servers. Therefore, a total of three NSX-T Local Managers would be provisioned to manage the three vCenter Servers.

Q25. A VMware Cloud Foundation (VCF) architect is presented with a customer’s requirements for an architecture that needs to achieve:

* Network high availability across workloads in two data center locations.

* Maintain the least administrative overhead when performing day 2 operations.

* Decrease the RTO of both management plane and data plane when site-wide failure occurs Which VCF design consideration should the architect recommend?
* VCF with NSX-T Bridging
* VCF with NSX-T L2-VPN
* VCF with NSX-T Federation
* VCF with NSX-T Multi-Site
Explanation

According to VMware Cloud Foundation Architecture Poster, VCF with NSX-T Federation provides network high availability across workloads in two data center locations bysynchronizing network configuration and state across sites. It also simplifies day 2 operations by providing centralized management and policy enforcement across sites. It also reduces RTO by enabling fast failover of network services between sites.

Q26. Which two options can be used to create a new VMware Cloud Foundation VI workload domain? (Choose two.)
* SDDC Manager Ul
* PowerCLI
* Cloud Builder Ul
* vCenter Ul
* REST API
The SDDC Manager UI provides a single point of control for managing and monitoring your VMware Cloud Foundation instance and for provisioning workload domains. You use the navigation bar to move between the main areas of the user interface 1. The SDDC Manager UI provides an integrated view of the physical and virtual infrastructure and centralized access to manage the physical and logical resources 2.

The REST API can also be used to create a new VI workload domain using VMware Cloud Foundation. The VMware Cloud Foundation API Reference Guide provides information on available operations 3.

Q27. What is the correct upgrade order for VMware Cloud Foundation (VCF) components for a VI workload domain with a stretched vSAN cluster?
* 1. NSX-T

2. vCenter Server

3. ESXi hosts

4. vSAN Witness host
* 1. vCenter Server

2. NSX-T

3. ESXi hosts

4. vSAN Witness host
* 1. vSAN witness host

2. NSX-T

3. ESXi hosts

4. vCenter Server
* 1. vSAN witness host

2. NSX-T

3. vCenter Server

4. ESXi hosts
https://docs.vmware.com/en/VMware-Cloud-Foundation/4.5/vcf-lifecycle/GUID-3B41CF79-C721-4AFC-A263-0672143DF41E.html

Q28. A VMware administrator is tasked to upgrade a VMware Cloud Foundation (VCF) environment that is running on DellEMC PowerEdge servers.

During the ESXi software upgrade for the VI Workload Domain hosts, the administrator receives an error stating that the correct storage driver is not available, although the storage adapters are enabled in the BIOS.

Which action should the administrator take to fix this issue?
* Use the Dell EMC customized image for the ESXi build in the VCF bill of materials.
* Upgrade the storage adapter firmware to the latest version.
* Use the image for the ESXi build in the VCF bill of materials.
* Upgrade the BIOS firmware to the latest version.
when upgrading ESXi software on Dell EMC PowerEdge servers using SDDC Manager Lifecycle Management (LCM), you must use a Dell EMC customized image that contains drivers for specific hardware components such as storage adapters.

https://docs.vmware.com/en/VMware-vSphere/6.7/com.vmware.vsphere.virtualsan.doc/GUID-08911FD3-2462-4C1C-AE81-0D4DBC8F7990.html

Q29. An administrator has registered an external identity source in a consolidated architecture and would like to make sure that any subsequent workload domains can be accessed using the same identity sources.

How can this goal be achieved with VMware Cloud Foundation?
* By configuring IWA as an identity source
* By configuring LDAPS as an identity source
* By keeping the pre-configured defaults
* By replicating vSphere SSO configuration
vSphere Single Sign-On (SSO) provides secure authentication and authorization services for VMware Cloud Foundation components, including vCenter Server and Platform Services Controller (PSC). In a consolidated architecture deployment of VMware Cloud Foundation, the vSphere SSO configuration is shared across all the workload domains.

To ensure that subsequent workload domains can use the same identity sources as an external identity source registered in a consolidated architecture, the administrator needs to replicate the vSphere SSO configuration. This can be achieved by configuring the same identity sources for vSphere SSO across all the workload domains.

Configuring IWA (Integrated Windows Authentication) or LDAPS (Lightweight Directory Access Protocol over SSL) as an identity source is a part of configuring the vSphere SSO configuration for identity sources.

Keeping the pre-configured defaults does not guarantee that the subsequent workload domains will use the same identity sources as the external identity source registered in a consolidated architecture.

Reference:

VMware Cloud Foundation Operations and Administration Guide:
https://docs.vmware.com/en/VMware-Cloud-Foundation/index.html VMware vSphere Security Guide:
https://docs.vmware.com/en/VMware-vSphere/7.0/vsphere-security-guide.pdf To ensure that any subsequent workload domains can be accessed using the same identity sources, it is necessary to replicate the vSphere SSO configuration across all the workload domains in a consolidated architecture deployment. This can be achieved by replicating the vSphere SSO configuration between the primary and additional SDDC Manager instances. This ensures that all the workload domains registered with the SDDC Manager will be able to consume resources and services from the same identity sources without any additional configuration in each individual workload domain.

VMware Cloud Foundation Administration Guide https://docs.vmware.com/en/VMware-Cloud-Foundation/index.html

Q30. Which license is required to enable Workload Management on VMware Cloud Foundation?
* VMware vSphere Evaluation
* VMware vSphere Standard
* VMware vSphere Enterprise Plus
* VMware Tanzu Basic

Explanation

A Tanzu Basic license is required to enable Workload Management on VMware Cloud Foundation. Once enabled, the Supervisor Cluster must be assigned a Tanzu license before the 60-day evaluation period expires.

This license can be added to the license inventory of vSphere if a valid Tanzu Edition license is available.

Q31. Which order of steps should an administrator use to replace a failed host in a stretched cluster?
* Decommission the failed host.

2. Remove the host using cluster APIs.

3. Add the newly commissioned host to the cluster using cluster APIs. 4 Commission the new host with the correct network.
* 1 Remove the host using cluster APIs.

2. Decommission the failed host.

3. Commission the new host with the correct network.

4. Add the newly commissioned host to the cluster using cluster APIs.
* Remove the host using cluster APIs

2. Decommission the failed host.

3. Add the newly commissioned host to the cluster using cluster APIs.

4. Commission the new host with the correct network
D.

1 Decommission the failed host

2. Remove the host using cluster APIs.

3. Commission the new host with the correct network.

4. Add the newly commissioned host to the cluster using cluster APIs.

Explanation:

This is because according to VMware documentation1, these are the steps to replace a failed host in a stretched cluster:

Run the compact cluster API to remove any stale data from vSAN.

Decommission the host to be removed using SDDC Manager UI or API.

Commission the replacement host to the same network pool as the removed host using SDDC Manager UI or API.

Add the newly commissioned host to the cluster using SDDC Manager UI or API.

According to the VMware documentation, when replacing a failed host in a stretched cluster, the first step is to decommission the failed host. This should be followed by removing the host using cluster APIs, commissioning the new host with the correct network, and then adding the newly commissioned host to the cluster using cluster APIs.

Q32. What is required as part of enabling the Harbor Image Registry?
* Storage Policy
* Tanzu Enabled Cluster
* Access Control
* Resource Limits
Explanation

This is because according to Dell documentation , to enable the Harbor Image Registry, you need to select the VM Storage Policy that will be used to store the images.

As part of enabling the Harbor Image Registry in VMware Cloud Foundation, a storage policy needs to be defined to specify the storage requirements for the registry. The storage policy should define the storage characteristics for the datastores where the registry will be deployed, including the redundancy level, disk type, and disk space. This is documented in the VMware documentation titled “Enabling Harbor Image Registry in Workload Domains.”

Q33. A vSphere administrator is tasked with enabling Workload Management on a VMware Cloud Foundation Workload Domain.

Which three components are configured as part of the Supervisor Cluster control plane after this task is completed? (Choose three.)
* Tanzu Kubernetes Grid Service

* kubectl-vSphere
* Kubernetes Grid Orchestrator
* Spherelet
* Kubernetes Mission Control
* Container Runtime Executive

This is because according to VMware documentation2, these are some of the components that are configured as part of the Supervisor Cluster control plane after enabling Workload Management on a VMware Cloud Foundation Workload Domain:

Tanzu Kubernetes Grid Service: This service enables you to create and manage Tanzu Kubernetes clusters on vSphere with Tanzu.

Kubernetes Grid Orchestrator: This component manages the lifecycle of Tanzu Kubernetes clusters on vSphere with Tanzu.

Spherelet: This component runs on each ESXi host and acts as a kubelet agent that communicates with the Supervisor Cluster control plane.

Q34. VCF design workshops were conducted, and the architect collected the following customer requirements for the newly planned VCF infrastructure:

* The new VCF infrastructure must target two zones: DEV/UAT and DMZ.

* The security team would like to have full management and network isolation between these two zones

* 12 hosts have been ordered for the solution.

* DEV/UAT workloads must comply with an erasure coding vSAN storage policy with the ability to tolerate the failure of two hosts.

Which workload domain sizing will be required to achieve these requirements?

* 12-hosts workload domain for both zones, having a 4-hosts DEV cluster a 4-hosts UAT cluster, and a

4-hosts DMZ cluster

* 12-hosts workload domain for both zones, having an 8-hosts DEV/UAT cluster, and a 4-hosts DMZ cluster
* 8-hosts DEV/UAT workload domain, having a 4-hosts DEV cluster and a 4-hosts UAT cluster, in addition to a 4-hosts DMZ workload domain, having a 4-hosts DMZ cluster
* 8-hosts DEV/UAT workload domain, having an 8-hosts DEV/UAT cluster, and a 4-hosts DMZ workload domain, having a 4-hosts DMZ cluster

Explanation

erasure coding vSAN storage policy with the ability to tolerate two host failures requires at least six fault domains (hosts) in a cluster. Therefore, an 8-hosts cluster can meet this requirement for DEV/UAT workloads.

Additionally, creating separate workload domains for DEV/UAT and DMZ can provide full management and network isolation between these two zones.

https://docs.vmware.com/en/VMware-vSphere/6.7/com.vmware.vsphere.virtualsan.doc/GUID-AD408FA8-5898

Q35. A developer is deploying pods with Persistent Volumes (PV) on vSphere with Tanzu. Which component determines the datastore that the PV will be placed on?

* CNS-CSI
* Hostd
* Spherelet

* SPBM
Explanation

This is because according to VMware documentation34, vSphere with Tanzu uses storage policies to integrate with shared datastores available in your environment, including VMFS, NFS, vSAN, or vVols datastores. The storage policies represent datastores and manage the storage placement of such objects as persistent volumes (PVs). Storage Policy Based Management (SPBM) is a framework that provides a single unified control plane across different types of datastores and enables administrators to define policies based on storage capabilities and requirements5.