Professional 2V0-13.24 Latest Exam Pdf & Passing 2V0-13.24 Exam is No More a Challenging Task

BTW, DOWNLOAD part of Exam4Tests 2V0-13.24 dumps from Cloud Storage: https://drive.google.com/open?id=1CQL6uHzfSmYI6RUI_aKTmukR_jchaQS_

Each of the Exam4Tests VMware 2V0-13.24 exam dumps formats excels in its way and carries actual VMware Cloud Foundation 5.2 Architect (2V0-13.24) exam questions for optimal preparation. All of these VMware Cloud Foundation 5.2 Architect (2V0-13.24) practice question formats are easy to use and extremely convenient such that even newbies find them simple.

If you have limited budget, and also need complete value package, why not try our Exam4Tests's 2V0-13.24 exam training materials. It is easy to understand with reasonable price and high accuracy. It's suitable for all kinds of learners. If you choose Exam4Tests' 2V0-13.24 Exam Training materials, you will get one year free renewable service.

>> 2V0-13.24 Latest Exam Pdf <<

2V0-13.24 Latest Exam Testking & Reliable 2V0-13.24 Exam Pdf

To advance your career, take the VMware Cloud Foundation 5.2 Architect exam. Your VMware demonstrates your commitment to lifelong learning. Passing the VMware Cloud Foundation 5.2 Architect exam in one sitting is not a walk in the park. The VMware 2V0-13.24 exam preparation process takes a lot of time and effort. You have to put time and money into passing the VMware Cloud Foundation 5.2 Architect exam. The best method to reap the rewards of your investment in becoming an expert is by using VMware 2V0-13.24 Exam Questions. Additionally, you can confidently study for the 2V0-13.24 exam.Passing an VMware Cloud Foundation 5.2 Architect exam on the first attempt can be stressful, but VMware 2V0-13.24 exam questions can help manage stress and allow you to perform at your best.

VMware 2V0-13.24 Exam Syllabus Topics:

Topic	Details
Topic 1	• Plan and Design the VMware by Broadcom Solution: This section of the exam measures the skills of VMware administrators. It involves gathering and analyzing business objectives and requirements to create a conceptual model. Additionally, it covers the creation of VMware Cloud Foundation logical and physical designs. This includes prerequisites and design decisions related to Network Infrastructure, VCF Management Domain, VCF Workload Domain, VCF Edge Cluster, VCF Cloud Automation, and VCF Cloud Operations. Designs should consider availability within and across availability zones, manageability (Lifecycle Management, Scalability, Capacity Management), performance, recoverability (BCDR strategies), and security for VCF Management Components and Workloads. Workload mobility, consumption, and monitoring strategies are also addressed in this section.

Topic 2	• IT Architectures, Technologies, Standards: This section of the exam measures the skills of enterprise architects and solution architects and focuses on the fundamentals of IT architectures, technologies, and standards. It covers differentiating between business and technical requirements, understanding conceptual models, and logical and physical designs, and recognizing the distinctions between requirements, assumptions, constraints, and risks. Also included are availability, manageability, performance, recoverability, and security (AMPRS), developing risk mitigation strategies, documenting design decisions, and creating design validation strategies.
Topic 3	• Troubleshoot and Optimize the VMware by Broadcom Solution: This section has NO TESTABLE OBJECTIVES in this version of the exam.
Topic 4	• VMware by Broadcom Solution: This section of the exam measures the skills of cloud architects and infrastructure engineers and focuses on understanding the architecture of VMware by Broadcom solution. Candidates should be able to differentiate between various VMware Cloud Foundation architecture options based on different scenarios.
Topic 5	• Install, Configure, and Administrate the VMware by Broadcom Solution: This section has NO TESTABLE OBJECTIVES in this version of the exam.

VMware Cloud Foundation 5.2 Architect Sample Questions (Q36-Q41):

NEW QUESTION # 36
A VMware Cloud Foundation multi-AZ (Availability Zone) design mandates that:
All availability zones must operate independently of each other.
The availability SLA must adhere to no less than 99.9%.
What would be the three design decisions that would help satisfy those requirements? (Choose three.)

- A. Make sure the recovery VLAN for the infrastructure management components has access to both AZ(s)
- B. Choose two distant AZ(s) and consider each AZ the DR for the other
- C. Configure a non-routable separate recovery VLAN for the infrastructure management components within each AZ
- D. Make sure all configuration backups are replicated between the selected AZ(s)
- E. Choose two close proximity AZ(s) and configure a stretched management workload domain
- F. Configure array-based replication between the selected AZ(s) for the management domain

Answer: C,D,F

Explanation:
This scenario involves a VCF multi-AZ design where AZs must operate independently (no shared dependencies) and achieve a 99.9% availability SLA (allowing ~8.76 hours of downtime annually). The design decisions must ensure resilience, fault isolation, and recovery capabilities across AZs.
Requirement Analysis:
Independent AZ operation:Each AZ must function standalone, with no single point of failure or dependency across AZs.
99.9% availability:The design must minimize downtime through redundancy, replication, and recovery mechanisms.
Option Analysis:
A: Configure array-based replication between the selected AZ(s) for the management domain:Array- based replication (e.g., vSphere Replication or SAN replication) for the management domain (vCenter, NSX Manager, SDDC Manager) ensures that critical management VMs are duplicated across AZs. If one AZ fails, the other can take over with minimal downtime, supporting independent operation and high availability. The VCF 5.2 Design Guide recommends replication for multi-AZ deployments to meet SLAs, as it provides a recovery point objective (RPO) near zero. This option enhances availability and is correct.
B: Make sure all configuration backups are replicated between the selected AZ(s):Replicating configuration backups (e.g., SDDC Manager backups, NSX configurations) ensures that each AZ has access to recovery data. If an AZ's management components fail, the other AZ can restore operations independently using its local backup copy. This supports the independence requirement and reduces downtime (contributing to 99.9% SLA) by enabling quick recovery. The VCF Administration Guide emphasizes backup replication for multi-AZ resilience, making this option correct.
C: Make sure the recovery VLAN for the infrastructure management components has access to both AZ (s):A recovery VLAN spanning both AZs implies a shared network dependency. If this VLAN fails (e.g., due to a network outage), both AZs could be impacted, violating the independence requirement. Multi-AZ designs in VCF favor isolated networks per AZ to avoid cross-AZ single points of failure. The VCF Design Guide advises against shared VLANs for critical components in independent AZ setups. This option undermines the requirements and is incorrect.

D: Choose two distant AZ(s) and consider each AZ the DR for the other:Distant AZs (e.g., separate data centers) with mutual DR (disaster recovery) roles enhance geographic fault tolerance. However, "operate independently" in VCF typically means each AZ can run workloads standalone, not that one is a passive DR site. Distant AZs introduce latency, complicating synchronous replication needed for 99.9% availability, and may rely on shared management, conflicting with independence. The VCF Multi-AZ Guide focuses on active- active AZs, not DR-centric designs, making this less suitable.
E: Choose two close proximity AZ(s) and configure a stretched management workload domain:A stretched management domain (e.g., using vSAN stretched cluster) spans AZs with synchronous replication, ensuring high availability. However, this creates a dependency: both AZs share the same vCenter and management stack, so a failure (e.g., vCenter outage) could affect both, violating independence. The VCF 5.2 Design Guide notes stretched clusters are for single logical domains, not independent AZs. This option contradicts the requirement and is incorrect.
F: Configure a non-routable separate recovery VLAN for the infrastructure management components within each AZ:A non-routable, AZ-specific recovery VLAN isolates management recovery traffic (e.g., for vMotion, backups) within each AZ. This ensures that each AZ's management components operate independently, with no cross-AZ network reliance. If one AZ's network fails, the other remains unaffected, supporting the SLA through fault isolation. The VCF Multi-AZ Design Guide recommends separate, isolated networks per AZ for resilience, making this option correct.
Conclusion:The three design decisions areConfigure array-based replication between the selected AZ(s) for the management domain (A),Make sure all configuration backups are replicated between the selected AZ(s) (B), andConfigure a non-routable separate recovery VLAN for the infrastructure management components within each AZ (F). These ensure independent operation and meet the 99.9% SLA through replication and isolation.
References:
VMware Cloud Foundation 5.2 Design Guide (Section: Multi-AZ Design)
VMware Cloud Foundation 5.2 Administration Guide (Section: Backup and Recovery) VMware Cloud Foundation Multi-AZ Deployment Guide (Section: Networking) VMware vSphere 8.0 Update 3 Documentation (Section: vSAN Stretched Clusters)

NEW QUESTION # 37
Which design decision should be prioritized when creating a VMware Cloud Foundation logical design for a network infrastructure?
Response:

- A. The specific models of firewalls and load balancers
- B. The exact placement of storage devices
- C. The physical layout of servers and switches
- D. The logical segmentation of networks and VLANs

Answer: D

NEW QUESTION # 38
An architect is designing a new VMware Cloud Foundation (VCF) solution. During the discovery workshops, the customer explained that the solution will initially be used to host a single business application and some internal management tooling. The customer provided the following background information:
The business application consists of two virtual machines.
The business application is sensitive to changes in its storage I/O.
The business application must be available during the company's business hours of 9 AM - 5 PM on weekdays.
The architect has made the following design decisions in response to the customer's requirements and the additional information provided during discovery:
The solution will use the VCF consolidated architecture model.
A single cluster will be created, consisting of six ESXi hosts.
Which design decision should the architect include in the design to mitigate the risk of impacting the business application?

- A. Perform ESXi host maintenance activities outside of the stated business hours.
- B. Implement FTT=6 for the business application virtual machines.
- C. Replace the vSAN shared storage exclusively with an All-Flash Fibre Channel shared storage solution.
- D. Use resource pools to apply CPU and memory reservations on the business application virtual machines.
- E. Use Anti-Affinity Distributed Resource Scheduler (DRS) rules on the business application virtual machines.

Answer: A

Explanation:
The VCF 5.2 design must ensure the business application (two VMs) remains available during business hours (9 AM - 5 PM weekdays) and is protected from storage I/O disruptions in a consolidated architecture with a single six-host cluster using vSAN.

The goal is to mitigate risks to the application's performance and availability. Let's evaluate each option:
Option A: Use resource pools to apply CPU and memory reservations on the business application virtual machinesResource pools with reservations ensure CPU and memory availability, which could help performance. However, the application's sensitivity is tostorage I/O, not CPU/memory, and the availability requirement (business hours) isn't directly addressed by reservations. While useful, this doesn't fully mitigate the primary risks identified, making it less optimal.
Option B: Implement FTT=6 for the business application virtual machinesThis is incorrect and infeasible. In vSAN, Failures to Tolerate (FTT) defines the number of host or disk failures a storage object can withstand, with a maximum FTT dependent on cluster size. FTT=6 requires at least 13 hosts (2n+1 where n=6), but the cluster has only six hosts, supporting a maximum FTT=2 (RAID-5/6). Even if feasible, FTT addresses data redundancy, not runtime availability or I/O sensitivity during business hours, making this irrelevant to the stated risks.
Option C: Perform ESXi host maintenance activities outside of the stated business hoursThis is the correct answer. In a vSAN-based VCF cluster, ESXi host maintenance (e.g., patching, reboots) triggers data resyncs and VM migrations (via vMotion), which can impact storage I/O performance and potentially cause brief disruptions. The application's sensitivity to storage I/O and its availability requirement (9 AM - 5 PM weekdays) mean maintenance during business hours poses a risk. Scheduling maintenance outside these hours (e.g., nights or weekends) mitigates this by ensuring uninterrupted I/O performance and availability during critical times, directly addressing the customer's needs.
Option D: Replace the vSAN shared storage exclusively with an All-Flash Fibre Channel shared storage solutionThis is incorrect. While an All-Flash Fibre Channel array might offer better I/O performance, VCF's consolidated architecture relies on vSAN as the primary storage for management and workload domains.
Replacing vSAN entirely contradicts the chosen architecture and introduces unnecessarycomplexity and cost.
The sensitivity to storage I/O changes doesn't justify abandoning vSAN, especially since All-Flash vSAN could meet performance needs if properly tuned.
Option E: Use Anti-Affinity Distributed Resource Scheduler (DRS) rules on the business application virtual machinesAnti-Affinity DRS rules ensure the two VMs run on separate hosts, improving availability by avoiding a single host failure impacting both. While this mitigates some risk, it doesn't address storage I/O sensitivity (a vSAN-wide concern) or guarantee availability during business hours if maintenance occurs. It's a partial solution but less effective than scheduling maintenance outside business hours.
Conclusion:The best design decision is toperform ESXi host maintenance activities outside of the stated business hours(Option C). This directly mitigates the risk of storage I/O disruptions and ensures availability during 9 AM - 5 PM weekdays, aligning with the customer's requirements in the VCF 5.2 consolidated architecture.
References:
VMware Cloud Foundation 5.2 Architecture and Deployment Guide (Section: Consolidated Architecture Design) VMware vSAN 7.0U3 Planning and Deployment Guide (integrated in VCF 5.2): Maintenance Mode Considerations VMware Cloud Foundation 5.2 Planning and Preparation Guide (Section: Availability and Performance Design)

NEW QUESTION # 39
A customer has stated the following requirements for Aria Automation within their VCF implementation:
Users must have access to specific resources based on their company organization.
Developers must only be able to provision to the Development environment.
Production workloads can be placed on DMZ or Production clusters.
What two design decisions must be implemented to satisfy these requirements? (Choose two.)

- A. Users' access to resources will be controlled by tenant membership.
- B. Separate tenants will be configured for Development and Production.
- C. Users' access to resources will be controlled by project membership.
- D. Separate cloud zones will be configured for Development and Production.

Answer: C,D

Explanation:
In VMware Cloud Foundation (VCF) 5.2, Aria Automation (formerly vRealize Automation) manages resource provisioning and access control. The requirements involve role-based access, environment isolation, and workload placement flexibility. Let's analyze each option:
Option A: Separate tenants will be configured for Development and ProductionAria Automation in VCF 5.2 operates as a single-tenant application by default, integrated with SDDC Manager and vCenter. Multi- tenancy (separate tenants) is an advanced configuration typically used for service providers, not standard VCF private cloud designs. TheVMware Aria Automation Installation Guidenotes that multi-tenancy adds complexity and isn't required for environment segregation within a single organization. Instead, projects and cloud zones handle these needs, making this unnecessary.
Option B: Users' access to resources will be controlled by tenant membershipTenant membership applies in multi-tenant setups, where users are assigned to distinct tenants (e.g., Dev vs. Prod). Since VCF 5.2 typically uses a single tenant, and the requirements can be met with projects (group-based access), this isn't a must-have decision. TheVCF 5.2 Architectural Guidefavors project-

based access over tenant separation for organizational control, rendering this optional.
Option C: Users' access to resources will be controlled by project membershipProjects in Aria Automation group users and define their access to resources (e.g., cloud zones, policies). To meet the first requirement (access based on company organization) and the second (developers provisioning only to Development), projects can restrict developers to a "Dev" project linked to a Development cloud zone, while other teams (e.g., ops) access Production/DMZ via separate projects. TheVMware Aria Automation Administration Guideconfirms projects as the primary mechanism for role-based access in VCF, making this a required decision.
Option D: Separate cloud zones will be configured for Development and ProductionCloud zones in Aria Automation map to vSphere clusters or resource pools (e.g., Development, Production, DMZ clusters). To satisfy the second requirement (developers limited to Development) and the third (Production workloads on DMZ or Production clusters), separate cloud zones ensure environment isolation and placement flexibility.
TheVCF 5.2 Architectural Guidemandates cloud zones for workload segregation, tying them to projects for access control, making this essential.
Conclusion:
C: Project membership enforces user access per organization and restricts developers to Development, meeting the first two requirements.
D: Separate cloud zones isolate Development from Production/DMZ, enabling precise workload placement per the third requirement.These decisions align with Aria Automation's design in VCF 5.2.References:
VMware Cloud Foundation 5.2 Architectural Guide(docs.vmware.com): Aria Automation Design and Cloud Zones.
VMware Aria Automation Administration Guide(docs.vmware.com): Projects and Access Control.
VMware Aria Automation Installation Guide(docs.vmware.com): Tenancy Options in VCF.

NEW QUESTION # 40
An architect is designing a VMware Cloud Foundation (VCF)-based private cloud solution for a customer that will include two physical locations. The customer has stated the following requirement:
All management tooling must be resilient at the component level within a single site.
When considering the design decisions for VMware Aria Suite components, what should the Architect document to meet the stated requirement?

- A. The solution will deploy Aria Suite Lifecycle Manager in a high availability configuration.
- B. The solution will implement an external load balancer for Aria Operations Cloud Proxies.
- C. The solution will deploy three Aria Automation appliances in a clustered configuration.
- D. The solution will configure the VCF Workload domain in a stretched topology across two locations.

Answer: C

Explanation:
The requirement specifies that management tooling must be resilient at the component level within a single site, meaning each site's management components (e.g., VMware Aria Suite) must withstand individual failures without relying on the other site. Let's evaluate each option in the context of VCF 5.2 and Aria Suite:
Option A: The solution will implement an external load balancer for Aria Operations Cloud Proxies Aria Operations Cloud Proxies collect data for monitoring and don't inherently require an external load balancer for resiliency within a site. TheVMware Aria Operations Administration Guideindicates that proxies are lightweight and typically deployed per cluster, with resiliency achieved via multiple proxies, not load balancing. This doesn't directly address component-level resiliency for the broader Aria Suite management tools.
Option B: The solution will configure the VCF Workload domain in a stretched topology across two locationsA stretched topology extends a workload domain across two sites for site-level resiliency (e.g., disaster recovery), not component-level resiliency within a single site. TheVCF 5.2 Architectural Guidenotes that stretched clusters rely on cross-site failover, which contradicts the requirement for single-site resilience, making this irrelevant to management tooling within one site.
Option C: The solution will deploy three Aria Automation appliances in a clustered configuration VMware Aria Automation (formerly vRealize Automation) supports a clustered deployment with three appliances (primary, replica, and failover) to ensure high availability within a site. TheVMware Aria Automation Installation Guideconfirms that this configuration provides component-level resiliency by allowing the cluster to tolerate individual appliance failures without service disruption. In VCF, Aria Automation is a key management tool, and this design meets the requirement for single-site resilience.
Option D: The solution will deploy Aria Suite Lifecycle Manager in a high availability configuration Aria Suite Lifecycle Manager (LCM) manages the lifecycle of Aria components but isn't deployed in a clustered HA configuration itself in VCF 5.2-it's a single appliance with backup/restore options. TheVCF
5.2 Administration Guidenotes that LCM resiliency is typically achieved via infrastructure HA (e.g., vSphere HA), not native clustering, making this less directly aligned with component-level resiliency compared to Aria Automation clustering.
Conclusion:Option C best meets the requirement by ensuring Aria Automation, a critical management tool, is resilient at the component level within a single site through clustering, aligning with VCF and Aria Suite best practices.References:

VMware Cloud Foundation 5.2 Architectural Guide(docs.vmware.com): Management Component Design.
VMware Aria Automation Installation Guide(docs.vmware.com): Clustered Configuration for HA.
VMware Aria Suite Lifecycle Administration Guide(docs.vmware.com): LCM Deployment Options.

NEW QUESTION # 41
......