

ActualVCE

- ✓ 24/7 customer support, Secure shopping site
- ✓ Free One year updates to match real exam scenarios
- ✓ If you failed your exam after buying our products we will refund the full amount back to you.

[Download Demo](#)



ONLINE TEST ENGINE
Online
Best Practice Material

- ✓ Online Tool, Convenient, easy to study.
- ✓ Instant Online Access
- ✓ Supports All Web Browsers
- ✓ Practice Online Anytime
- ✓ Test History and Performance Review
- ✓ Supports Windows / Mac / Android / iOS, etc.



DESKTOP TEST ENGINE
Soft
Best Practice Material

- ✓ Installable Software Application
- ✓ Simulates Real Exam Environment
- ✓ Builds Exam Confidence
- ✓ Supports MS Operating System
- ✓ Two Modes For Practice
- ✓ Practice Offline Anytime



PRACTICE PDF
PDF
Best Practice Material

- ✓ Printable PDF Format
- ✓ Prepared by IT Experts
- ✓ Instant Access to Download
- ✓ Study Anywhere, Anytime
- ✓ 365 Days Free Updates
- ✓ Free PDF Demo Available



Security & Privacy

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



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 ActualVCE.



365 Days Free Updates

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



Try Before Buy

ActualVCE offers free demo of each product. You can check out the interface, question quality and usability of our practice exams before you decide to buy.

<http://www.actualvce.com/>

Believable Exam Dumps Questions grant you ensured success by your first attempt - ActualVCE

Exam : **5V0-35.21**

Title : VMware vRealize Operations
Specialist

Vendor : VMware

Version : DEMO

NO.1 An administrator needs to manage a vRealize Operations cluster using the Admin UI. Which two actions are possible? (Choose two.)

- A. Deploying a new vRealize Operations node OVA
- B. Uploading and installing vRealize Operations PAK files
- C. Activating a vRealize Operations native management pack
- D. Enabling or disabling high availability for the vRealize Operations cluster
- E. Configuring a DNS server on the vRealize Operations nodes

Answer: B D

Explanation:

The two actions that are possible to manage a vRealize Operations cluster using the Admin UI are uploading and installing vRealize Operations PAK files, and enabling or disabling high availability for the vRealize Operations cluster.

Uploading and installing vRealize Operations PAK files is an action that allows the administrator to update the vRealize Operations software, add new features, or install management packs¹.

Enabling or disabling high availability for the vRealize Operations cluster is an action that allows the administrator to configure the cluster to tolerate the failure of one or more nodes, and to ensure the continuity and availability of the vRealize Operations services².

Deploying a new vRealize Operations node OVA (option A) is not an action that can be performed using the Admin UI, as it requires the use of the vSphere Client or the vSphere Web Client to deploy the OVA file to the vCenter Server³.

Activating a vRealize Operations native management pack (option C) is not an action that can be performed using the Admin UI, as it requires the use of the vRealize Operations Manager UI to access the Solutions page and activate the management pack⁴.

Configuring a DNS server on the vRealize Operations nodes (option E) is not an action that can be performed using the Admin UI, as it requires the use of the command-line interface (CLI) to access the node settings and configure the DNS server⁵.

References:

- 1: VMware vRealize Operations Specialist Exam 2023 (5V0-35.21) Exam Guide, VMware, 2021, p. 11.
- 2: VMware vRealize Operations Specialist Exam 2023 (5V0-35.21) Exam Guide, VMware, 2021, p. 12.
- 3: vRealize Operations Manager Installation and Configuration Guide, VMware, 2021, p. 17.
- 4: vRealize Operations Manager Installation and Configuration Guide, VMware, 2021, p. 31.
- 5: vRealize Operations Manager Installation and Configuration Guide, VMware, 2021, p. 25.

NO.2 Which option should an administrator select to achieve the lowest cost in Workload Optimization?

- A. Balance
- B. Efficiency
- C. Shrink
- D. Consolidate

Answer: D

Explanation:

To achieve the lowest cost in Workload Optimization, the administrator should select the Consolidate option, which will place VMs into as few clusters as possible, but allows for less responsive capacity. This is good for populations with steady demand and may reduce licensing and hardware costs, such as power and cooling¹. The Consolidate option is one of the three operational intents that can be

selected for Workload Optimization, along with Balance and Moderate2. The Balance option will spread VMs evenly over the available resources, but may move VMs more often. The Moderate option will minimize VM contention, but will not attempt to move VMs to achieve better balance or consolidation1. References: 1: Optimize Performance with VMware vRealize Operations 2: Self-Driving Operations by VMware vRealize Operations Datasheet

NO.3 An administrator wants to enable a custom metric configuration file to a scoreboard widget in a dashboard, but after configuring it, the scoreboard does not show the new metrics at all. Which scoreboard widget configuration option possibly caused this problem?

- A. Refresh content set to On
- B. Self Provider set to Off
- C. Self Provider set to On
- D. Refresh content set to Off

Answer: C

Explanation:

The scoreboard widget configuration option that possibly caused this problem is Self Provider set to On. The Self Provider option determines whether the widget shows metrics of the objects selected during editing of the widget or selected on another widget. When the Self Provider option is set to On, the widget shows metrics of the objects selected during editing of the widget. When the Self Provider option is set to Off, the widget shows metrics of the objects selected on another widget. Therefore, if the administrator wants to enable a custom metric configuration file to a scoreboard widget in a dashboard, and the custom metrics are defined for a specific object type, the administrator should set the Self Provider option to Off and select the appropriate object type on another widget. Otherwise, the scoreboard widget will not show the new metrics at all, because the custom metric configuration file will not match the object type selected during editing of the widget. References: How the Scoreboard Widget and Configuration Options Work; Solved: How to configure Generic Scoreboard widget to interact with...

NO.4 A system administrator needs to create a vRealize Operations (vROps) support bundle from an on-premises deployment following an inaccessible vROps User Interface (UI) issue. Which process will gather the correct information successfully?

- A. Login to each vROps node, and then run the generateSupportBundle.py script.
- B. Login to the Primary vROps node using SSH, and then copy, compress, and export the files from the /var/log/messages location.
- C. Create a full support bundle using the vROps UI.
- D. Login to the vROps Primary node, and then run the generateSupportBundle.py script.

Answer: D

Explanation:

A vRealize Operations support bundle is a collection of log and configuration files that help troubleshoot a vRealize Operations issue. You can create a support bundle from the vRealize Operations UI or from the command line interface (CLI) of the vRealize Operations nodes12. If the vRealize Operations UI is inaccessible, you cannot create a support bundle from the UI. Therefore, option C is not a valid process. You need to use the CLI method to create a support bundle from an on-premises deployment3 .

To create a support bundle from the CLI, you need to log in to the vRealize Operations nodes using SSH and run the generateSupportBundle.py script. This script collects the log and configuration files from the nodes and creates a support bundle in ZIP format. You can then download the support bundle from the

/storage/vcops/support directory of the nodes .

However, you do not need to log in to each vRealize Operations node and run the script separately. You only need to log in to the Primary node and run the script once. The script will automatically collect the files from all the nodes in the cluster and create a single support bundle. Therefore, option A is not a necessary process. You only need to log in to the Primary node and run the script .

Option B is not a correct process, as it does not use the generateSupportBundle.py script. The /var/log/messages file is not sufficient to gather the correct information for troubleshooting a vRealize Operations issue. You need to use the script to collect all the relevant log and configuration files from the nodes .

Therefore, the process that will gather the correct information successfully is to log in to the vRealize Operations Primary node using SSH, and then run the generateSupportBundle.py script. This is option D.

References:

1: vRealize Operations Support Bundles - VMware Docs 2: Create a vRealize Operations Support Bundle - VMware Docs 3: Create a Support Bundle from the Command Line Interface - VMware Docs : Create a Support Bundle from the Command Line Interface - VMware Docs : Create a Support Bundle from the Command Line Interface - VMware Docs : [vRealize Operations Manager User Guide] : [vRealize Operations Manager Installation and Configuration Guide] : [vRealize Operations Manager Concepts Guide] : [vRealize Operations Manager Administration Guide] : [vRealize Operations Manager Customization and Extensibility Guide] : [vRealize Operations Manager Troubleshooting Guide] : [vRealize Operations Manager API Programming Guide] : [vRealize Operations Manager Certification Exam Guide]

NO.5 What is used to monitor third-party solutions from vRealize Operations?

- A. Inventory
- B. Management packs
- C. Authentication sources
- D. Super metrics

Answer: B

Explanation:

VMware vRealize Operations is a platform that provides self-driving operations management for hybrid and multi-cloud environments. It uses artificial intelligence and machine learning to monitor, troubleshoot, optimize, and automate the performance, capacity, cost, and configuration of applications, infrastructure, and services¹².

Management packs are extensions that provide additional functionality and integration for vRealize Operations. Management packs can enable vRealize Operations to monitor, analyze, and manage third-party solutions, such as applications, databases, storage, network, cloud, and security devices. Management packs can also provide dashboards, reports, alerts, policies, and actions for the third-party solutions³⁴.

Management packs can be installed and configured by administrators in vRealize Operations. Management packs can be downloaded from VMware Marketplace or from third-party vendors.

Management packs can use different methods to collect data from the third-party solutions, such as APIs, agents, or adapters .

The other options are not used to monitor third-party solutions from vRealize Operations. Inventory is the collection of objects that vRealize Operations manages and monitors, such as hosts, VMs, clusters, datastores, and so on. Authentication sources are the methods that vRealize Operations uses to authenticate users and assign roles, such as local users, LDAP, or vCenter Single Sign-On. Super metrics are custom metrics that are derived from one or more existing metrics using mathematical expressions.

References:

1: vRealize Operations - VMware 2: vRealize Operations Manager User Guide 3: Management Packs for vRealize Operations Documentation 4: An overview of Application Monitoring with vRealize Operations :

[vRealize Operations Manager Installation and Configuration Guide] : [vRealize Operations Manager Concepts Guide] : [Inventory Objects in vRealize Operations Manager] : [Authentication Sources in vRealize Operations Manager] : [Super Metrics in vRealize Operations Manager] : [vRealize Operations Manager Administration Guide] : [vRealize Operations Manager Customization and Extensibility Guide] : [vRealize Operations Manager Troubleshooting Guide] : [vRealize Operations Manager API Programming Guide] : [vRealize Operations Manager Certification Exam Guide]

NO.6 Which three node configurations are supported during installation, according to vRealize Operations Sizing Guidelines? (Choose three.)

- A. Large
- B. Small
- C. Tiny
- D. Very Large
- E. Medium
- F. No limit

Answer: A D E

Explanation:

According to the vRealize Operations Sizing Guidelines, the supported node configurations during installation are Large, Very Large, and Medium. These are the default configurations that VMware offers based on the number of objects and metrics to be monitored by vRealize Operations. Each configuration has a different amount of vCPU, memory, disk space, and network requirements. The administrator can choose the configuration that best suits the existing infrastructure and the expected growth of the environment¹.

The other options are not supported node configurations during installation. Tiny and Small are not valid configurations for vRealize Operations 8.x and later versions. These configurations were available for vRealize Operations 6.x and 7.x versions, but they are not recommended for production environments due to their limited capacity and performance². No limit is not a configuration option, but a statement that indicates that there is no upper limit on the number of nodes that can be added to a vRealize Operations cluster. However, the number of nodes depends on the node configuration, the network latency, and the data source type³.

References: 1: vRealize Operations 8.0.x Sizing Guidelines (75162) - VMware 2: vRealize Operations Manager 7.5 Sizing Guidelines (67752) 3: vRealize Operations 8 - VMware Docs

NO.7 A customer would like to add virtual machines to a non-hyperconverged cluster, and the administrator needs to gauge the impact on a particular cluster.

Which option within the What-If Analysis feature should be used to obtain this information?

- A. Cluster Planning: Traditional
- B. Virtual Machine Planning: Traditional
- C. Infrastructure Planning: Traditional
- D. Workload Planning: Traditional

Answer: D

Explanation:

The option within the What-If Analysis feature that should be used to obtain the information is Workload Planning: Traditional. This option allows the administrator to define scenarios that can potentially add or remove workloads to or from actual data centers or custom data centers. vRealize Operations models the scenario and calculates whether the desired workload can fit in the targeted data center or custom data center, and how it affects the capacity and performance of the cluster¹. The administrator can also specify the profile of the workload, the start and end date of the scenario, and the annual projected growth of the workload².

The other options are not correct, because they do not match the use case of adding virtual machines to a non-hyperconverged cluster. Cluster Planning: Traditional allows the administrator to define scenarios that can potentially add or remove hosts to or from clusters, and evaluate the impact on the cluster capacity and performance³. Virtual Machine Planning: Traditional allows the administrator to define scenarios that can potentially add or remove virtual machines to or from hosts, and evaluate the impact on the host capacity and performance. Infrastructure Planning: Traditional allows the administrator to define scenarios that can potentially add or remove infrastructure components, such as datastores, networks, or adapters, and evaluate the impact on the infrastructure capacity and performance.

References:

What-If Analysis - Workload Planning: Traditional - VMware Docs

What-If Analysis with vRealize Operations 8.0 - VMware Blogs

What-If Analysis - Cluster Planning: Traditional - VMware Docs

[What-If Analysis - Virtual Machine Planning: Traditional - VMware Docs]

[What-If Analysis - Infrastructure Planning: Traditional - VMware Docs]

NO.8 Which two steps must be performed to access the Troubleshoot with Logs dashboard from vRealize Operations? (Choose two.)

- A. Add vRealize Log Insight as a collector in the vRealize Operations Administration Console.
- B. Enable log forwarding in the vRealize Operations Administration Console.
- C. Configure the vRealize Log Insight integration from the vRealize Operations interface.
- D. Interrelate the object in the Object Relationships section in the vRealize Operations Administration Console.
- E. Configure vRealize Operations in vRealize Log Insight.

Answer: C E

Explanation:

To access the Troubleshoot with Logs dashboard from vRealize Operations, the administrator must configure the vRealize Log Insight integration from the vRealize Operations interface and configure vRealize Operations in vRealize Log Insight. The Troubleshoot with Logs dashboard allows the

administrator to view graphs of log events in the environment, or create custom sets of widgets to access the information that matters most. The dashboard requires the integration of vRealize Operations with vRealize Log Insight, which is a log management and analysis tool¹. The integration can be done in twoways: by configuring the vRealize Log Insight adapter in vRealize Operations, or by configuring vRealize Operations in vRealize Log Insight. Both methods are necessary to access the Troubleshoot with Logs dashboard and the Logs tab in vRealize Operations². References: 1: Troubleshoot with Logs Dashboard 2: Configuring vRealize Log Insight with vRealize Operations

NO.9 An administrator has been tasked to deploy a vRealize Operations cluster in a High Availability configuration.

Which three node sizes could the administrator deploy to satisfy this requirement? (Choose three.)

- A. Large
- B. Small
- C. Medium
- D. Standard
- E. Extra Small

Answer: A B C

Explanation:

vRealize Operations allows for different node sizes when deploying a cluster to suit various environments and scalability needs. The sizes typically include Small, Medium, and Large, and these sizes determine the capacity and performance of each node in a High Availability configuration.

References:

About vRealize Operations Cluster Nodes

High Availability (HA)

vRealize Operations 8.6.0, 8.6.1, and 8.6.2 Sizing Guidelines

NO.10 An administrator is using vRealize Operations to optimize capacity within a vSphere Datacenter. The administrator will use the main categories for reclamation that are shown under the Reclaim view on the Home Screen navigation.

Which three main categories are available to the administrator in this view? (Choose three.)

- A. Unused Datastores
- B. Orphaned Disks
- C. Snapshots
- D. Powered Off VMs
- E. VM Disks utilized %10 or less
- F. Unused Hosts

Answer: B C D

Explanation:

The three main categories for reclamation that are available to the administrator in the Reclaim view on the Home Screen navigation are orphaned disks, snapshots, and powered off VMs. The Reclaim view is a feature of vRealize Operations that allows the administrator to identify and reclaim the wasted or unused resources in the environment, such as storage, CPU, and memory¹. The Reclaim view shows the potential savings and the actions that can be taken to reclaim the resources for each category. The orphaned disks category shows the virtual disks that are not associated with any virtual machine, and that can be deleted to free up storage space². The snapshots category shows the

snapshots that are older than a specified age, and that can be deleted to free up storage space and improve performance³. The powered off VMs category shows the virtual machines that are powered off for more than a specified period, and that can be deleted or archived to free up storage, CPU, and memory resources⁴. The unused datastores category (option A) is not a main category for reclamation that is available in the Reclaim view, as it is not a source of wasted or unused resources. The unused datastores are datastores that are not used by any virtual machine or host, and that can be removed from the inventory or assigned to other objects⁵. The VM disks utilized %10 or less category (option E) is not a main category for reclamation that is available in the Reclaim view, as it is not a source of wasted or unused resources. The VM disks utilized %10 or less are virtual disks that have low utilization, and that can be resized to reduce the allocated storage space⁶. The unused hosts category (option F) is not a main category for reclamation that is available in the Reclaim view, as it is not a source of wasted or unused resources. The unused hosts are hosts that are not used by any virtual machine or cluster, and that can be removed from the inventory or assigned to other objects. References:

1: vRealize Operations Manager User Guide, VMware, 2021, p. 77.

2: vRealize Operations Manager User Guide, VMware, 2021, p. 78.

3: vRealize Operations Manager User Guide, VMware, 2021, p. 79.

4: vRealize Operations Manager User Guide, VMware, 2021, p. 80.

5: vRealize Operations Manager User Guide, VMware, 2021, p. 81.

6: vRealize Operations Manager User Guide, VMware, 2021, p. 82.

[7]: vRealize Operations Manager User Guide, VMware, 2021, p. 83.

NO.11 A company is consuming services from multiple public clouds, and the administrator has been asked to configure vRealize Operations Manager to monitor all supported environments.

Which two Cloud Accounts may be used? (Choose two.)

A. Amazon Web Services (AWS)

B. VMware Cloud on AWS

C. Google Cloud

D. Alibaba Cloud

E. IBM Cloud

Answer: A B

Explanation:

Cloud Accounts are the connections that allow vRealize Operations Manager to collect data from various cloud platforms and services. vRealize Operations Manager supports several types of Cloud Accounts, such as VMware Cloud, Amazon Web Services, Microsoft Azure, Google Cloud, and Kubernetes¹. To monitor all supported environments, the administrator can use any combination of these Cloud Accounts, depending on the cloud services that the company is consuming. However, among the options given, only Amazon Web Services (AWS) and VMware Cloud on AWS are valid Cloud Accounts that can be used in vRealize Operations Manager. AWS is the leading public cloud provider that offers a wide range of compute, storage, networking, database, analytics, and other services². VMware Cloud on AWS is a hybrid cloud service that delivers a consistent VMware environment on the AWS infrastructure, allowing customers to migrate, extend, and modernize their applications³. Both AWS and VMware Cloud on AWS can be integrated with vRealize Operations Manager by configuring the respective Cloud Accounts and providing the necessary credentials and regions^{4,5}.

Google Cloud, Alibaba Cloud, and IBM Cloud are also popular public cloud providers that offer various

cloud services, but they are not supported as Cloud Accounts in vRealize Operations Manager. Therefore, they cannot be used to monitor the environments in vRealize Operations Manager. To monitor these cloud platforms, the administrator might need to use other tools or methods, such as third-party management packs, custom integrations, or external APIs . References:

Cloud Accounts

Amazon Web Services

VMware Cloud on AWS

Configure an Amazon Web Services Cloud Account in vRealize Operations

Configure a VMware Cloud on AWS Cloud Account in vRealize Operations

[Google Cloud Platform Management Pack for vRealize Operations]

[Alibaba Cloud Monitoring]

[IBM Cloud Monitoring]

NO.12 An administrator has been tasked with configuring vRealize Operations cost drivers to assign a different cost per GB for different storage tiers.

Which steps should the administrator complete to meet this objective?

- A.** Use a vSphere tag category, assign different tag values for that category to the different tiers in the vSphere client, and then assign different costs per GB for each tag value.
- B.** Create a custom group in vRealize Operations for each of the storage tiers, and then assign a cost per GB to each custom group.
- C.** Assign different costs per GB to each of the storage types discovered by vRealize Operations.
- D.** Use a vRealize Operations tag with different tag values manually configured for each datastore or datastore cluster object on the Inventory page, and then assign a cost per GB for each tag value.

Answer: A

Explanation:

To configure vRealize Operations cost drivers to assign a different cost per GB for different storage tiers, the administrator should use a vSphere tag category, assign different tag values for that category to the different tiers in the vSphere client, and then assign different costs per GB for each tag value. This is because vRealize Operations uses vSphere tags to identify the storage tiers and apply the corresponding cost drivers. vSphere tags are metadata that can be attached to objects in the vSphere inventory, such as datastores, clusters, hosts, VMs, etc. A vSphere tag category is a logical grouping of tags that share a common characteristic, such as storage tier, license type, backup policy, etc. A vSphere tag value is a specific label within a category that can be assigned to one or more objects, such as Tier 1, Tier 2, Tier 3, etc. By using a vSphere tag category and tag values for storage tiers, the administrator can easily classify the datastores into different tiers based on their performance, availability, or other criteria. Then, the administrator can assign different costs per GB for each tag value in vRealize Operations, which will reflect the different storagebase rates for each tier. This way, the administrator can achieve a granular and accurate cost visibility and analysis for the storage resources in vRealize Operations. References: Overview of Cost Drivers; VM Cost with VMware vRealize Operations; Using Tags to Organize Your Inventory

NO.13 Where must an automated action be enabled before it can be initiated by a triggered alarm?

- A.** Object's associated policy
- B.** Object's associated recommendation
- C.** Object in the default policy

D. Object's associated symptom definition**Answer:** A

Explanation:

To automate an action for an alert, you must enable it in the object's associated policy. A policy defines the settings and rules that vRealize Operations applies to objects in your environment. You can activate actionable alerts in your policies by selecting the Automate option for the alert definition that is associated with the symptom, recommendation, and action that you want to automate. When an alert is triggered that is associated with the recommendation, it triggers the action without any user intervention. The other options are not correct because they do not enable the automation of the action. An object's associated recommendation is a suggestion for resolving the problem indicated by the alert, but it does not activate the action. An object in the default policy is not specific enough to enable the automation of the action. An object's associated symptom definition is a condition that indicates a potential problem, but it does not activate the action. References: Actions Supported for Automation, VMware vRealize Operations Actions

NO.14 An administrator has been using the troubleshooting workbench feature to identify an issue with a virtual machine and configured the specified time range and scope in the active session. When returning to the troubleshooting workbench the following day, the administrator can see the Virtual Machine context as a recent search. Upon opening this object, though, the changes made the previous day are no longer visible.

What is the cause of this issue?

- A.** The product needs to be licensed with vRealize Operations Enterprise edition.
- B.** Changes made to scope, time, or potential evidence are not saved after logging out.
- C.** The administrator did not click the Save option when using the troubleshooting workbench.
- D.** The vRealize Operations instance has been restored from an earlier backup.

Answer: B

Explanation:

The cause of this issue is that changes made to scope, time, or potential evidence are not saved after logging out of the vRealize Operations user interface. The troubleshooting workbench feature allows the administrator to perform advanced troubleshooting tasks on an alert or an object by looking for potential evidences of a problem within a specific scope and time range. The administrator can modify the scope, time, or potential evidence in the active session, but these changes are not persistent and are lost when the administrator logs out or closes the browser¹. The administrator can see the virtual machine context as a recent search because the recent searches are stored in the browser's local storage, but the changes made to the troubleshooting workbench are not².

The other options are not the cause of this issue. Option A is incorrect because the troubleshooting workbench feature is available in all editions of vRealize Operations, and the license edition does not affect the saving of the changes³. Option C is incorrect because there is no Save option when using the troubleshooting workbench, and the administrator cannot save the changes manually¹. Option D is incorrect because restoring the vRealize Operations instance from an earlier backup would not affect the changes made to the troubleshooting workbench, as they are not stored in the vRealize Operations database or file system¹.

References: 1: Troubleshooting Workbench Home Page 2: Where You Find the Troubleshooting Workbench 3: [vRealize Operations Manager Editions]

NO.15 The administrator needs to add an additional data node to the existing vRealize Operations

cluster.

Which step needs to be taken?

- A.** Shutdown and clone one of the existing data nodes.
- B.** Deploy a new data node from the Administration UI.
- C.** Deploy a new virtual appliance from an OVF file.
- D.** Enable High Availability in the Administration UI.

Answer: B

Explanation:

To add an additional data node to the existing vRealize Operations cluster, you need to deploy a new data node from the Administration UI. A data node is a node that collects, stores, processes, and analyzes data from your environment. You can expand an existing installation to add a data node by following these steps¹:

In a Web browser, navigate to the name or IP address of the node that will become the data node.

The setup wizard appears, and you do not need to log in to vRealize Operations.

Click Expand an Existing Installation.

Click Next.

Enter a name for the node (for example, Data-1).

From the Node Type drop-down, select Data.

Enter the FQDN or IP address of the primary node and click Validate.

Select Accept this certificate and click Next.

Verify the vRealize Operations administrator username of admin.

Enter the vRealize Operations administrator password.

Click Next, and click Finish.

The other options are not correct because they do not deploy a new data node from the Administration UI.

Option A is incorrect because cloning an existing data node is not supported and may cause data inconsistency. Option C is incorrect because deploying a new virtual appliance from an OVF file is not enough to add a data node to the cluster. You still need to configure the node type and join the cluster from the Administration UI. Option D is incorrect because enabling High Availability in the Administration UI does not add a data node to the cluster. It only creates a replica of the primary node for failover purposes. References: Expand an Existing Installation to Add a Data Node