2V0-71.23 Q&As

VMware Tanzu for Kubernetes Operations Professional

Pass VMware 2V0-71.23 Exam with 100% Guarantee

QUESTION 1

Which two statements describe Kubernetes observability characteristics? (Choose two.)

A. It provides network insight and detailed Kubernetes network topology view

B. Provides visibility into Kubernetes clusters for troubleshooting and impact assessment

C. It observes the code of the applications running in Kubernetes environment

D. Collects real-time metrics from all layers of Kubernetes

E. Automatically heals Kubernetes workloads after an issue has been observed

Correct Answer: BD

Kubernetes observability is the ability to monitor and analyze the performance, health, and behavior of Kubernetes clusters and applications. It provides visibility into Kubernetes clusters for troubleshooting and impact assessment, by collecting logs, events, traces, and alerts from various sources. It also collects real-time metrics from all layers of Kubernetes, such as nodes, pods, containers, services, and network policies, and displays them in dashboards and charts. Kubernetes observability helps administrators and developers to identify and resolve issues, optimize resource utilization, and ensure service quality and reliability. References: VMware Tanzu Observability Documentation, What is Kubernetes Observability?

QUESTION 2

Which Kubernetes object must be used to be able to upgrade a pod without disrupting services?

A. ReplicaSet

B. Service

C. Container

D. Deployment

Correct Answer: D

A Deployment is a Kubernetes object that allows you to perform a rolling update without disrupting services. A Deployment manages a ReplicaSet or a Pod and provides declarative updates for them. You can describe the desired state of

your application using a Deployment, and it will change the actual state to the desired state at acontrolled rate. A Deployment also allows you to roll back to a previous version if something goes wrong during the update14.

The other options are incorrect because:

A ReplicaSet is a Kubernetes object that ensures that a specified number of pod replicas are running at any given time. It does not provide any mechanism for updating or rolling back pods15.

A Service is a Kubernetes object that defines a logical set of pods and a policy to access them. It does not provide any mechanism for updating or rolling back pods16.

A Container is not a Kubernetes object, but rather a component of a Pod. A Pod is the smallest deployable unit of computing in Kubernetes. A Pod can contain one or more containers that share storage and network resources. A Pod does

not provide any mechanism for updating or rolling back itself or its containers17. References: Deployments, ReplicaSets, Services, Pods

QUESTION 3

What are two possible counts of control plane nodes in a Tanzu Kubernetes Grid Workload Cluster? (Choose two.)

A. 3

B. 5

C. 2

D. 0

E. 1

Correct Answer: AE

The control plane nodes are the nodes that run the Kubernetes control plane components, such as the API server, the scheduler, the controller manager, and etcd. The control plane nodes are responsible for managing the cluster state and orchestrating workload operations. The possible counts of control plane nodes in a Tanzu Kubernetes Grid workload cluster are 1 or 3. The control plane must have an odd number of nodes to ensure quorum and high availability. A single control plane node is suitable for development or testing purposes, while three control plane nodes are recommended for production clusters23. References: Deploy Workload Clusters - VMware Docs, Concepts and References - VMware Docs

QUESTION 4

Which two package management tools can be used to configure and install applications on Kubernetes? (Choose two.)

A. Grafana

B. Fluent bit

C. Carvel

D. Helm

E. Multus

Correct Answer: CD

Two package management tools that can be used to configure and install applications on Kubernetes are: Carvel. Carvel is a set of tools that provides a simple, composable, and flexible way to manage Kubernetes configuration, packaging, and deployment. Carvel includes tools such as kapp, which applies and tracks Kubernetes resources in a cluster; ytt, which allows templating YAML files; kbld, which builds and pushes images to registries; kpack, which automates image builds from source code; and vendir, which syncs files from different sources into a single directory. Carvel is integrated with VMware Tanzu Kubernetes Grid and can be used to deploy and manage applications on Tanzu clusters. Helm. Helm is a tool that helps users define, install, and upgrade complex Kubernetes applications using charts. Charts are packages of pre-configured Kubernetes resources that can be customized with values. Helm uses a client- server architecture with a command line tool called helm and an in-cluster component called Tiller. Helm can be used to deploy applications from the official Helm charts repository or from custom charts created by users or vendors. Helm is also integrated with VMware Tanzu Kubernetes Grid and can be used to deploy and manage applications on Tanzu clusters. References: : https://carvel.dev/ : https://docs.vmware.com/en/VMware-Tanzu-Kubernetes-Grid/1.6/vmware-tanzu-kubernetes-grid-16/GUID-tkg-carvel.html : https://helm.sh/ : https://docs.vmware.com/en/VMware-Tanzu-Kubernetes-Grid/1.6/vmware-tanzu- kubernetes-grid-16/GUID-tkg-helm.html

QUESTION 5

An administrator has a VMware Tanzu Kubernetes Grid management cluster named tanzu- mc0l which needs to be upgraded.

Which command can be used to upgrade this cluster?

A. kubectl management-cluster upgrade

B. tanzu mc upgrade

C. tanzu config use-context tanzu-mc01-admin@tanzu-mc01

D. kubectl tanzu-mc01 upgrade

Correct Answer: B

The tanzu mc upgrade command is used to upgrade a management cluster to a newer version of Tanzu Kubernetes Grid. The command requires the name of the management cluster as an argument, and optionally the version to upgrade to.

For example, toupgrade the management cluster named tanzu-mc01 to version v1.4.0, the command would be:

tanzu mc upgrade tanzu-mc01 --version v1.4.0

The other options are incorrect because:

kubectl management-cluster upgrade is not a valid command. The kubectl command is used to interact with Kubernetes clusters, not to upgrade them.

tanzu config use-context tanzu-mc01-admin@tanzu-mc01 is a command to switch the current context to the admin context of the management cluster named tanzu- mc01. It does not upgrade the cluster.

kubectl tanzu-mc01 upgrade is not a valid command. The kubectl command does not accept a cluster name as an argument, and there is no upgrade subcommand. References: VMware Tanzu for Kubernetes Operations Getting Started,

Upgrading Management Clusters

QUESTION 6

Which component must be installed upfront to deploy VMware Tanzu Kubernetes Grid management cluster?

A. Tanzu CLI

B. Cluster API

C. Kubeadm

D. External DNS

Correct Answer: A

The Tanzu CLI is a command-line tool that enables users to interact with VMware Tanzu products and services. It must be installed upfront to deploy VMware Tanzu Kubernetes Grid management cluster, as it provides commands to create, configure, scale, upgrade, and delete management clusters on different platforms. The Tanzu CLI also allows users to create workload clusters from the management cluster, and to perform various operations on both types of clusters. References: VMware Tanzu CLI Documentation, [Deploying Management Clusters with the Tanzu CLI]

QUESTION 7

An administrator will enable workload management in vSphere using NSX Advanced Load Balancer.

Which two components does the administrator need to prepare on NSX Advanced Load Balancer in advance? (Choose two.)

A. NSX Controller

B. Service Engine Group

C. Provide connectivity to NSX manaqer

D. NSX Advanced Load Balancer Controller

E. Avi Kubernetes Operator

Correct Answer: BD

To enable workload management in vSphere using NSX Advanced Load Balancer, an administrator needs to prepare two components on NSX Advanced Load Balancer in advance: the Service Engine Group and the NSX Advanced Load Balancer Controller1. The Service Engine Group is a logical group of Service Engines that share the same configuration and resources. A Service Engine is a virtual machine that handles the data plane operations of NSX Advanced Load Balancer, such as load balancing, health monitoring, SSL termination, and more2. The administrator needs to configure a Service Engine Group for each Supervisor Cluster that will use NSX Advanced Load Balancer as the load balancer provider1. The NSX Advanced Load Balancer Controller is a virtual machine that handles the control plane operations of NSX Advanced Load Balancer, such as configuration, analytics, orchestration, and management2. The administrator needs to deploy and configure the NSX Advanced Load Balancer Controller VM in the management network of the vSphere environment where workload management will be enabled1. The other options are incorrect because: The NSX Controller is not a component of NSX Advanced Load Balancer, but rather a component of NSX-T Data Center. The NSX Controller is a clustered virtual appliance that provides the control plane functions for logical switching and routing3. It is not required for enabling workload management in vSphere using NSX Advanced Load Balancer. Providing connectivity to NSX Manager is not a component of NSX Advanced Load Balancer, but rather a prerequisite for enabling workload management in vSphere using NSX-T Data Center. The NSX Manager is a virtual appliance that provides the management plane functions for NSX-T Data Center3. It is not required for enabling workload management in vSphere using NSX Advanced Load Balancer. The Avi Kubernetes Operator is not a component of NSX Advanced Load Balancer, but rather an optional tool that can be used to automate the installation and configuration of NSX Advanced Load Balancer on Kubernetes clusters4. It is not required for enabling workload management in vSphere using NSX Advanced Load Balancer.

References: Install and Configure the NSX Advanced Load Balancer for vSphere with Tanzu with NSX, NSX Advanced Load Balancer Architecture, NSX-T Data Center Architecture, Avi Kubernetes Operator

QUESTION 8

An administrator was requested to create a pod with two interfaces to separate the application and management traffic for security reasons.

Which two packages have to be installed in VMware Tanzu Kubernetes Grid cluster to satisfy the requirement? (Choose two.)

A. multus

B. external-dns

C. cert-manager

D. qrafana

E. contour

Correct Answer: AE

Multus is an open-source container network interface plugin for Kubernetes that enables attaching multiple network interfaces to pods. Contour is an open-source Kubernetes ingress controller that provides dynamic configuration updates and makes use of the Envoy proxy as a data plane. By installing these two packages in a VMware Tanzu Kubernetes Grid cluster, an administrator can create a pod with two interfaces and use Contour to route the application and management traffic to different networks. The other options are incorrect because: external-dns is a package that synchronizes exposed Kubernetes services and ingresses with DNS providers. It does not provide multiple interfaces for pods. cert-manager is a package that automates the management and issuance of TLS certificates from various sources. It does not provide multiple interfaces for pods. qrafana is not a valid package name. The correct spelling is Grafana, which is a package that provides visualization and analytics for metrics collected by Prometheus. It does not provide multiple interfaces for pods. References: Install Multus and Whereabouts for Container Networking, Install Contour for Ingress

QUESTION 9

What two steps are required to visualize API connectivity and enable API protection in VMware Tanzu Service Mesh? (Choose two.)

A. Activate API Discovery for the Global Namespace

B. Create API Security Policy for the Global Namespace

C. Enable Threat Detection Policy for the Global Namespace

D. Set a Distributed Firewall policy for the Global Namespace

E. Create an Autoscaling policy for API for the Global Namespace

Correct Answer: AB

To visualize API connectivity and enable API protection in VMware Tanzu Service Mesh, the administrator needs to perform two steps: Activate API Discovery for the Global Namespace. This allows Tanzu Service Mesh to automatically discover the APIs signatures between microservices running inside or outside the mesh. API Discovery creates a custom API schema for each API that is close to OpenAPI spec 3.0. Tanzu Service Mesh graph renders the detected APIs in the Enforcing mode by default, which means that any new API is considered as a violated API unless accepted by the administrator1 Create API Security Policy for the Global Namespace. This allows the administrator to block or allow layer 4 and layer 7 traffic, as well as create granular policies that provide API and data segmentation, OWASP 10 attack defense, schema validation, geofencing, data compliance, and egress controls. The administrator can create the API Security policy through the Tanzu Service Mesh Console UI or by using the Tanzu Service Mesh API Explorer2 References: 1: https:// docs.vmware.com/en/VMware-Tanzu-Service-Mesh/services/tanzu- service-mesh-enterprise/GUID-E6FB9FB3-FDB3-4D2B-B5CB-614608EEF537.html 2: https://docs.vmware.com/en/VMware-Tanzu-Service-Mesh/services/tanzu-service-mesh- enterprise/GUID-5B635420-3BD2-4EC1-B67E-2015F991A91C.html

QUESTION 10

Which two are valid options for obtaining kubectl config file in Tanzu Kubernetes environment? (Choose two.)

A. Use the command tubeccl vsphere login

B. Download from vSphere Ul

C. Download on the Supervisor Cluster Webpage

D. Access from vCenter Server Appliance Management Interface

E. Access from VMware Tanzu Mission Control

Correct Answer: BE

Two valid options for obtaining kubectl config file in Tanzu Kubernetes environment are:

Download from vSphere UI: For Tanzu Kubernetes clusters that are deployed on vSphere with Tanzu, you can download the kubeconfig file from the vSphere UI by selecting the cluster and clicking on the Download kubeconfig button1. This

file contains the credentials and connection information for the cluster, which you can use to access it with kubectl1.

Access from VMware Tanzu Mission Control: For Tanzu Kubernetes clusters that are attached or provisioned by VMware Tanzu Mission Control, you can access the kubeconfig file from the Tanzu Mission Control console by selecting the

cluster and clicking on the Access this cluster button2. This will generate a YAML file that you can download and use to connect to the cluster with kubectl2. References: Download a Kubeconfig File for a Tanzu Kubernetes Cluster - VMware Docs, Connect to a Managed Cluster with kubectl - VMware Docs

QUESTION 11

What is the role of the Tanzu Kubernetes Grid Service?

A. It provides declarative, Kubernetes-style APIs for cluster creation, configuration, and management.

B. It provides a declarative, Kubernetes-style API for management of VMs and associated vSphere resources.

C. It provisions an extension inside the Kubernetes cluster to validate user authentication tokens.

D. It provisions Kubernetes clusters that integrate with the underlying vSphere Namespace resources and Supervisor Services.

Correct Answer: D

The role of the Tanzu Kubernetes Grid Service is to provision Kubernetes clusters that integrate with the underlying vSphere Namespace resources and Supervisor Services. The Tanzu Kubernetes Grid Service is a component of vSphere with Tanzu that provides self-service lifecycle management of Tanzu Kubernetes clusters3. A Tanzu Kubernetes cluster is an opinionated installation of Kubernetes that runs on top of the Supervisor Cluster and inherits its capabilities, such as storage integration, pod networking, load balancing, authentication, and authorization4. The Tanzu Kubernetes Grid Service exposes three layers of controllers to manage the lifecycle of a Tanzu Kubernetes cluster: Cluster API, Virtual Machine Service, and Tanzu Kubernetes Release Service3. References: Tanzu Kubernetes Grid Service Architecture - VMware Docs, What Is a Tanzu Kubernetes Cluster? - VMware Docs

QUESTION 12

A Tanzu Mission Control administrator would like to enforce the following container controls:

1.

Only allows container images that match the specified names or tags.

2.

Ensure that the container image is not tampered with.

Which type of policy can be used?

A. Access

B. Security

C. Image Security

D. Image Registry

E. Network

Correct Answer: C

The type of policy that can be used to enforce the container controls is image security. Image security policies allow users to define rules for validating container images before they are deployed on clusters. Users can specify image names, tags, signatures, or digests to whitelist or blacklist images based on their source and integrity. Users can also enable or disable image scanning for vulnerabilities and configure the severity threshold for admission decisions. References: Image Security Policy - VMware Docs, Image Policy - VMware Docs

QUESTION 13

Which L7 ingress mode leverages the integration between NSX Advanced Load Balancer and Antrea?

A. L7 ingress in NodePort mode

B. L7 ingress in ClusterIP mode

C. L7 ingress in NodePortLocal mode

D. L7 ingress in NodeIntegration mode

Correct Answer: C

L7 ingress in NodePortLocal mode is an ingress mode that leverages the integration between NSX Advanced Load Balancer and Antrea. NSX Advanced Load Balancer (NSX ALB) is a solution that provides L4 and L7 load balancing and ingress control for Kubernetes clusters5. Antrea is a Kubernetes networking solution that implements the Container Network Interface (CNI) specification and uses Open vSwitch (OVS) as the data plane6. In NodePortLocal mode, the ingress backend service must be ClusterIP mode, and Antrea assigns a unique port on each node for each pod that serves as a backend for the service. The network traffic is routed from the client to the NSX ALB Service Engine (SE), and then directly to the pods without going through the nodes or kube-proxy. This mode reduces network latency and improves performance by avoiding extra hops7. The following diagram illustrates how the network traffic is routed in NodePortLocal mode: !NodePortLocal mode diagram The other options are incorrect because: L7 ingress in NodePort mode is an ingress mode that does not leverage the integration between NSX ALB and Antrea. In this mode, the ingress backend service must be NodePort mode, and the network traffic is routed from the client to the NSX ALB SE, and then to the cluster nodes, before it reaches the pods. The NSX ALB SE routes the traffic to the nodes, and kube-proxy helps route the traffic from the nodes to the target pods. This mode requires an extra hop for kube-proxy to route traffic from node to pod7. L7 ingress in ClusterIP mode is an ingress mode that does not leverage the integration between NSX ALB and Antrea. In this mode, the ingress backend service must be ClusterIP mode, and Antrea assigns a virtual IP (VIP) for each service. The network traffic is routed from the client to the NSX ALB SE, and then to one of the VIPs assigned by Antrea, before it reaches the pods. The NSX ALB SE routes the traffic to one of the VIPs, and kube-proxy helps route the traffic from the VIPs to the target pods. This mode requires an extra hop for kube-proxy to route traffic from VIPs to pod7. L7 ingress in NodeIntegration mode is not a valid ingress mode for NSX ALB. References: NSX Advanced Load Balancer, Antrea, NSX ALB as L7 Ingress Controller

QUESTION 14

Which statement correctly describes the Cluster API?

A. It is a specialized toolset to bring declarative, Kubernetes-style APIs to cluster creation, configuration, and management in the Kubernetes ecosystem.

B. It enables pod networking and enforces network Kubernetes policies.

C. It is responsible for scanning language-specific packages in container images, such as Java, Python, Go, and others.

D. It is a native Kubernetes certificate management controller that adds certificates and certificate issuers as resource types in Kubernetes clusters.

Correct Answer: A

The statement that correctly describes the Cluster API is that it is a specialized toolset to bring declarative, Kubernetes-style APIs to cluster creation, configuration, and management in theKubernetes ecosystem. Cluster API is a Kubernetes sub-project that provides declarative APIs and tooling to simplify provisioning, upgrading, and operating multiple Kubernetes clusters5. Cluster API uses a set of custom resource definitions (CRDs) to represent clusters, machines, and other objects. Cluster API also relies on providers to implement the logic for interacting with different infrastructure platforms5. References: Introduction - The Cluster API Book

QUESTION 15

What are three capabilities of VMware Aria Operations for Applications (formerly known as Tanzu Observability)? (Choose three.)

A. Create Alerts

B. Set Application Container security policy

C. Set Service Level Objectives

D. Create Kubernetes Clusters

E. Create Charts and Dashboards

F. Create Queries

Correct Answer: AEF

VMware Aria Operations for Applications (formerly known as Tanzu Observability) is a unified observability platform that provides full-stack visibility using metrics, traces, and logs across distributed applications, application services, container services, and multi-cloud environments. Some of the capabilities of VMware Aria Operations for Applications are: Create alerts: Users can monitor for certain behaviors and get smart notifications based on query conditions. Users can create alerts independently or directly from charts, and use advanced and accurate alerting powered by AI/analytics and query language1. Create charts and dashboards: Users can visualize their data based on query results in various chart types (such as line plot, point plot, table, pie chart, etc.) and organize them in dashboards. Users can also interact with charts and dashboards in real time, such as zoom in, zoom out, change the time window, change the focus, and so on1. Create queries: Users can use the Wavefront Query Language (WQL) to extract the information they need from their data. Users can use the Chart Builder for easy query creation or the Query Editor for advanced query editing. Users can also use functions, operators, variables, macros, and expressions to manipulate their data1. References: VMware Aria Operations for Applications Documentation, Unified Observability Platform by VMware Aria Operations for Applications