Oracle.1z0-1109-24.v2025-06-28.q18

Exam Code:	1z0-1109-24
Exam Name:	Oracle Cloud Infrastructure 2024 DevOps Professional
Certification Provider:	Oracle
Free Question Number:	18
Version:	v2025-06-28
# of views:	123
# of Questions views:	180
https://www.exam-tests.com/1z0-1109-24-exam/Oracle.1z0-1109-24.v2025-06-28.q18.html	

NEW QUESTION: 1
While adding variables to your build_spec.yaml file, you made a mistake that resulted in a failed build pipeline.
What is the error you could have made?
A. Defined parameters such as the ${VARIABLE_NAME} file and later assigned their values in the Parameters tab of the build pipeline
B. Defined variables as exportedVariables to make them available in subsequent stages of the same pipeline
C. Defined a field such as type: DOCKER_IMAGE in the outputArtifacts: section to specify the docker image produced by the Build stage
D. Used vaultVariable to hold the content of the vault secrets in Base64 format
Answer: A (LEAVE A REPLY)
In build_spec.yaml for OCI DevOps, variables must be correctly defined and used according to the specification. If you defined parameters using ${VARIABLE_NAME} and then attempted to assign their values through the Parameters tab in the build pipeline, you may have introduced an error. The syntax and usage of parameters must be properly defined in both the build_spec.yaml and the build pipeline interface.

NEW QUESTION: 2
The Kubernetes Master node serves as the central control plane for managing the cluster's resources and orchestrating workload deployment.
What are the primary responsibilities of the Kubernetes Master node?
A. The Master node is primarily tasked with executing application workloads and ensuring their availability within the cluster.
B. The Master node oversees the scheduling and allocation of cluster resources, ensuring optimal utilization across nodes.

C. The Master node monitors network traffic within the cluster, ensuring secure communication and efficient data transfer between nodes.
D. The Master node serves as a repository for storing container images, facilitating rapid deployment and scaling of applications.
Answer: B (LEAVE A REPLY)
The Kubernetes Master node is the central control plane responsible for managing the cluster's resources and orchestrating workload deployment. Its primary responsibilities include:
Scheduling: Deciding which workloads (pods) should run on which worker nodes, based on resource availability and scheduling constraints.
Resource Allocation: Allocating resources to workloads to ensure optimal utilization across nodes.
Cluster Management: Maintaining the desired state of applications, managing cluster events, and ensuring that all the components of the cluster are functioning properly.

NEW QUESTION: 3
How does the Oracle Cloud Infrastructure Container Engine for Kubernetes (OKE) Cluster Autoscaler determine when to create new nodes for an OKE cluster?
A. When the CPU or memory utilization crosses a configured threshold.
B. When the resource requests from pods exceed a configured threshold.
C. When the custom metrics from the services exceed a configured threshold.
D. When the rate of requests to the application crosses a configured threshold.
Answer: B (LEAVE A REPLY)
The OKE Cluster Autoscaler automatically adjusts the number of worker nodes in an OKE cluster based on the resource requests made by Kubernetes pods. When there are not enough resources available (e.g., CPU or memory) on existing nodes to accommodate pending pods, the Cluster Autoscaler will create new nodes to meet the resource demand.

NEW QUESTION: 4
What are the two items required to create a rule for the Oracle Cloud Infrastructure Events Service? (Choose two.)
A. Management Agent Cloud Service
B. Service Connector
C. Auth Token
D. Rule Conditions
E. Actions
Answer: D,E (LEAVE A REPLY)

NEW QUESTION: 5
A DevOps engineer is asked to access an Oracle Cloud Infrastructure Container Engine for Kubernetes (OKE) cluster to deploy new applications and manage existing ones.

Which two statements are true? (Choose two.)
A. To access the cluster using kubectl, you have to set up a Kubernetes configuration file for the cluster. The kubeconfig file by default is named config and stored in the $HOME/.kube directory.
B. When a cluster's Kubernetes API endpoint has a public IP address, you can access the cluster in Cloud Shell by setting up a kubeconfig file.
C. Generating an API signing key pair is not required while setting up cluster access using local machine if the public key is not already uploaded in the console.
D. The only available option when a cluster's Kubernetes API endpoint has a public IP address is to control the cluster locally using kubectl and the Kubernetes Dashboard.
E. To access the cluster using kubectl, you have to set up a Kubernetes manifest file for the cluster. The kubeconfig file by default is named config and stored in the $HOME/.manifest directory.
Answer: A,B (LEAVE A REPLY)
To access an OKE cluster using kubectl, you need to set up a Kubernetes configuration file (kubeconfig). By default, the kubeconfig file is named config and stored in the $HOME/.kube directory.
When a cluster's Kubernetes API endpoint has a public IP address, you can use Cloud Shell to access the cluster. Setting up a kubeconfig file is required to authenticate and manage the cluster.

NEW QUESTION: 6
How can you scale a deployment named nodejs-deployment to have two replicas?
A. kubectl set replicas deployment nodejs-deployment --replicas=2
B. kubectl resize deployment nodejs-deployment --replicas=2
C. kubectl adjust deployment nodejs-deployment --replicas=2
D. kubectl scale deployment nodejs-deployment --replicas=2
Answer: D (LEAVE A REPLY)
The kubectl scale command is used to scale the number of replicas in a deployment. By specifying the --replicas flag, you define the desired number of replicas for the deployment.
(kubectl set replicas) is not the correct syntax for scaling a deployment.
(kubectl resize) is not a valid command for scaling a deployment.
(kubectl adjust) is also not a valid Kubernetes command.

NEW QUESTION: 7
As a cloud engineer, you are responsible for managing a Kubernetes cluster on the Oracle Cloud Infrastructure (OCI) platform for your organization. You are looking for ways to ensure reliable operations of Kubernetes at scale while minimizing the operational overhead of managing the worker node infrastructure.
Which cluster option is the best fit for your requirement?

A. Using OCI OKE managed nodes with cluster autoscalers to eliminate worker node infrastructure management
B. Using OCI OKE virtual nodes to eliminate worker node infrastructure management
C. Using Kubernetes cluster add-ons to automate worker node management
D. Creating and managing worker nodes using OCI compute instances
Answer: B (LEAVE A REPLY)
Step 1: Understanding the Requirement
The goal is to ensure reliable operations of Kubernetes at scale while minimizing the operational overhead of managing worker node infrastructure. In this context, a solution is needed that abstracts away the complexity of managing, scaling, and maintaining worker nodes.
Step 2: Explanation of the Options
A . Using OCI OKE managed nodes with cluster autoscalers
While this option provides managed node pools and uses cluster autoscalers to adjust resources based on demand, it still requires some level of management for the underlying worker nodes (e.g., patching, upgrading, monitoring).
Operational overhead: Moderate.
B . Using OCI OKE virtual nodes
Virtual nodes in OCI OKE are a serverless option for running Kubernetes pods. They remove the need to manage underlying worker nodes entirely.
OCI provisions resources dynamically, allowing scaling based purely on pod demand.
There's no need for node management, patching, or infrastructure planning, which perfectly aligns with the requirement to minimize operational overhead.
Operational overhead: Minimal.
Best Fit for This Scenario: Since the requirement emphasizes minimizing operational overhead, this is the ideal solution.
C . Using Kubernetes cluster add-ons to automate worker node management Kubernetes add-ons like Cluster Autoscaler or Node Problem Detector help in automating some aspects of worker node management. However, this still requires managing worker node infrastructure at the core level.
Operational overhead: Moderate to high.
D . Creating and managing worker nodes using OCI compute instances
This involves manually provisioning and managing compute instances for worker nodes, including scaling, patching, and troubleshooting.
Operational overhead: High.
Not Suitable for the Requirement: This option contradicts the goal of minimizing operational overhead.
Step 3: Why Virtual Nodes Are the Best Fit
Virtual Nodes in OCI OKE:
Virtual nodes provide serverless compute for Kubernetes pods, allowing users to run workloads without provisioning or managing worker node infrastructure.

Scaling: Pods are automatically scheduled, and the required infrastructure is dynamically provisioned behind the scenes.
Cost Efficiency: You only pay for the resources consumed by the running workloads.
Use Case Alignment: Eliminating the burden of worker node infrastructure management while ensuring Kubernetes reliability at scale.
Step 4: References and OCI Resources
OCI Documentation:
OCI Kubernetes Virtual Nodes
OCI Container Engine for Kubernetes Overview
Best Practices for Kubernetes on OCI:
Best Practices for OCI Kubernetes Clusters

NEW QUESTION: 8
As a DevOps engineer working on managing clusters on the OCI platform for your organization, which statement is true about managing cluster add-ons in OCI OKE Cluster?
A. When creating a new cluster, essential cluster add-ons cannot be disabled.
B. When enabling a cluster add-on, you cannot configure the add-on by specifying one or more key/value pairs to pass as arguments to the cluster add-on.
C. When creating a new cluster, essential cluster add-ons are set to manually update.
D. When you disable a cluster add-on using the console, the add-on is completely removed from the cluster.
Answer: A (LEAVE A REPLY)
Essential cluster add-ons are required for the basic functioning of the Kubernetes cluster and cannot be disabled during cluster creation. These add-ons provide necessary features such as core DNS, networking, and other critical functionalities for the cluster's operation.

NEW QUESTION: 9
A team wants to deploy artificial intelligence and machine learning workloads in their OCI Container Engine for Kubernetes (OKE) cluster. They prioritize strong isolation, cost-efficiency, and the ability to leverage serverless capabilities.
Which solution is best suited for their requirements?
A. Virtual nodes in OKE
B. Self-Managed Nodes in OKE
C. Managed nodes in OKE
D. Container Instances in OCI
Answer: A (LEAVE A REPLY)
Virtual nodes in OKE provide a serverless experience for deploying Kubernetes workloads, which means you do not have to manage or scale the underlying infrastructure. This solution is particularly cost-efficient because you only pay for the resources used by the pods, and it provides strong isolation for workloads.

Virtual nodes are well suited for AI/ML workloads as they allow users to easily scale compute resources without being constrained by the limits of individual worker nodes.

NEW QUESTION: 10
As a DevOps engineer at XYZ Corp, you are responsible for ensuring the smooth operation of high-traffic web applications hosted on Oracle Cloud Infrastructure (OCI). The web applications run on multiple OCI resources, including virtual machines, load balancers, and databases. Recently, users have reported failures while accessing one of the OCI-based web applications, and you suspect HTTP 5XX errors on the load balancer. You need to quickly identify and address this issue.
Which of the following statements can assist you in quickly identifying and monitoring the HTTP 5XX error rate on the load balancer and setting up notifications?
A. Use Custom Metrics of the Monitoring service to collect HTTP 5XX error rates from the load balancer and set up Service Connectors with third-party services such as PagerDuty or Slack.
B. Use Metrics and Alarms of the Monitoring service with Container Engine for Kubernetes (OKE) to monitor HTTP 5XX errors on Kubernetes resources and correlate them with other OCI resources.
C. Use Event Rules to detect HTTP 5XX errors on the load balancer and trigger automated actions using OCI Functions or API Gateway.
D. Use Metrics and Alarms of the Monitoring service to monitor the HTTP 5XX error rate on the load balancer and set up notifications with OCI Notifications.
Answer: B (LEAVE A REPLY)
The Monitoring service in OCI can be used to track metrics for various OCI resources, including load balancers. You can monitor specific metrics, such as HTTP 5XX error rates, to identify issues.
By using Alarms, you can set up thresholds for the HTTP 5XX error rate and receive notifications when the threshold is breached. The notifications can be configured through OCI Notifications, which allows integration with email, PagerDuty, Slack, and other channels.

NEW QUESTION: 11
What is the correct approach to upgrade an Oracle Container Engine for Kubernetes (OKE) Cluster to a newer version of Kubernetes?
A. Upgrade the control plane, then upgrade the node pools.
B. Initiate the automated upgrade process using the OCI Console, CLI, or API.
C. Upgrade the node pools one at a time, then once all node pools are upgraded, upgrade the control plane.
D. Initiate the control plane and node pool upgrades simultaneously.
Answer: A (LEAVE A REPLY)

The correct approach to upgrade an Oracle Container Engine for Kubernetes (OKE) cluster involves first upgrading the Kubernetes control plane, followed by upgrading the node pools. The control plane must be upgraded first to ensure compatibility with newer versions of Kubernetes, as node pools rely on the control plane for orchestration and management.
After upgrading the control plane, each node pool is upgraded to match the new Kubernetes version. This phased approach ensures the cluster remains in a stable state during the upgrade.

NEW QUESTION: 12
A DevOps team is deploying a new version of their application to their production environment using the Canary deployment strategy in the OCI DevOps service. They want to ensure that the production environment is not affected by any potential issues caused by the new version.
Which statement is true in regards to the Canary deployment strategy?
A. The Production stage in the Canary deployment strategy deploys the new version to the production environment without any manual approval.
B. The Canary deployment strategy only supports pipeline redeployment for OKE and not for instance group deployments.
C. The Invoke Function stage is an optional stage that can be used to validate the new version before moving to the production environment.
D. The Shift Traffic stage in the Canary deployment strategy shifts the entire 100% of the production traffic to the Canary environment.
Answer: C (LEAVE A REPLY)
In a Canary deployment strategy, a new version of an application is rolled out to a small subset of the production environment (the canary) first, while the majority of the production environment still runs the current stable version. This approach helps to detect potential issues before the new version is fully deployed.
The Invoke Function stage can be used as an optional step in the deployment pipeline to run additional validation (e.g., automated tests or other checks) on the new version before gradually shifting more traffic to it or moving to full production deployment.

NEW QUESTION: 13
What are the two items required to create a rule for the Oracle Cloud Infrastructure Events Service? (Choose two.)
A. Rule Conditions
B. Management Agent Cloud Service
C. Auth Token
D. Actions
E. Service Connector
Answer: (SHOW ANSWER)

Oracle Cloud Infrastructure (OCI) Events Service allows you to create rules to automatically respond to changes in your OCI environment. To create a rule, you need the following items:
Rule Conditions: Conditions define the events that will trigger the rule. This specifies what event (e.g., instance creation, object storage bucket modification) will trigger an action.
Actions: Actions define what should happen when a condition is met. For example, the action might be to trigger an OCI Function or send a notification.

NEW QUESTION: 14
Your team is working on a project to deploy a microservices-based application on a cloud platform using Terraform. Each microservice has specific configurations and dependencies, and you want to ensure modularity, reusability, and consistency across deployments.
Which Terraform feature would you use to achieve these objectives efficiently?
A. Terraform Providers
B. Terraform Workspaces
C. Terraform Variables
D. Terraform Modules
Answer: D (LEAVE A REPLY)
Terraform Modules are used to organize and group related configuration resources into reusable components. By using modules, you can achieve modularity, reusability, and consistency across different deployments, making it easier to manage complex infrastructure setups.
For a microservices-based application, where each microservice has specific configurations and dependencies, modules allow you to define the infrastructure for each microservice in a modular way. This helps to maintain clean, reusable code and ensures consistency across deployments.

NEW QUESTION: 15
As a DevOps engineer working on containerizing a microservices-based application to be hosted on OCI Cloud platforms, which step can help ensure that the container images have not been modified after being pushed to Oracle Cloud Infrastructure Registry (OCIR)?
A. Scanning the image upon ingestion and comparing the image size for changes
B. Enabling scanning of container images stored in OCI Registry
C. Deploying a manifest to the Kubernetes cluster that references the container image and its unique hash
D. Signing the image using the Container Registry CLI and creating an image signature that associates the image with the master encryption key and key version in the Vault service
Answer: D (LEAVE A REPLY)

To ensure that container images have not been modified after being pushed to the Oracle Cloud Infrastructure Registry (OCIR), you should sign the image. This involves using the Container Registry CLI to create a digital signature for the image, which associates the image with a master encryption key and key version stored in the OCI Vault service. This signature can then be verified at the time of deployment, ensuring that the image has not been tampered with since it was signed.

NEW QUESTION: 16
As a DevOps Engineer, you are tasked with explaining the key concepts of Terraform to a new team member. You want to ensure they understand the fundamental concepts of Terraform.
Which of the following best describes the purpose of Terraform variables?
A. Terraform variables are used to manage the life cycle of Terraform resources.
B. Terraform variables are used to define input values for Terraform configurations, allowing for customization and reuse of infrastructure code.
C. Terraform variables are used to output the final state of the infrastructure after deployment.
D. Terraform variables are used to define the structure and organization of Terraform configuration files.
Answer: B (LEAVE A REPLY)
Terraform variables are used to define input values for Terraform configurations. They allow users to customize infrastructure deployments by providing different values without modifying the configuration files themselves. Variables help in creating reusable infrastructure code, making it easy to maintain and adjust the infrastructure setup according to different environments or needs.

NEW QUESTION: 17
Why is it important to extract output artifacts from the Oracle Cloud Infrastructure (OCI) DevOps build pipeline and store them in an Artifact Registry repository?
A. All artifacts are permanently stored in the build pipeline. Extracting just the ones required for deployment tells the deployment pipeline which artifacts to use.

B. Storing build artifacts in registries helps the deployment pipeline differentiate output artifacts created by the build pipeline from artifacts copied from a Git repository.
C. Deliver Artifacts is a required stage of the build pipeline, and the entire pipeline won't work if it is not included in order to extract artifacts after the Managed Build stage.
D. Output artifacts aren't permanent. If they are to be used in the Deliver Artifacts stage, they need to be exported as output artifacts to a registry.
Answer: (SHOW ANSWER)
In OCI DevOps Build Pipeline, the output artifacts generated during the build are temporary and will be discarded unless explicitly stored in a persistent location. By extracting and storing these artifacts in an Artifact Registry (such as OCI Artifact Registry or OCI Container Registry), you ensure that they are available for subsequent stages, such as deployment.
Storing artifacts in a registry provides a persistent location where the deployment pipeline can access them, ensuring the artifacts are available for reliable deployment.

NEW QUESTION: 18
Observability helps understand system behavior, but it has limitations.
Which of the following tasks cannot be accomplished solely through observability?
A. Optimizing resource utilization
B. Automating software deployments
C. Identifying infrastructure bottlenecks
D. Identifying resource consumption
Answer: (SHOW ANSWER)
Observability refers to the ability to understand a system's internal state by observing its outputs. It involves metrics, logs, and tracing to provide insights into the system's behavior, which helps in monitoring and identifying issues.