

Oracle.1Z0-1109-24.v2025-07-04.q23

Exam Code:	1z0-1109-24
Exam Name:	Oracle Cloud Infrastructure 2024 DevOps Professional
Certification Provider:	Oracle
Free Question Number:	23
Version:	v2025-07-04
# of views:	119
# of Questions views:	230
https://www.dumpsdb.com/dumps/Oracle/1Z0-1109-24/Oracle.1Z0-1109-24.v2025-07-04.q23	

NEW QUESTION: 1

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: 2

Which OCI DevOps project resource is responsible for defining the stages for compiling, testing, and running software applications before deployment?

- A. Deployment pipelines
- B. Build pipelines
- C. Development pipelines
- D. Integration pipelines

Answer: B (LEAVE A REPLY)

Build pipelines in OCI DevOps are responsible for defining the stages involved in compiling, testing, and running software applications. These pipelines automate the process of building the code, running unit tests, scanning for vulnerabilities, and packaging the software, all of which occur before deploying the application.

NEW QUESTION: 3

You're using Oracle Cloud Infrastructure (OCI) DevOps service to automate your software releases to release features more frequently and with fewer errors. While deploying an update to production, one of your deployment stages failed.

What action should you perform in your Deployment Pipeline?

- A. Use OCI DevOps Trigger and Rerun tool to avoid downtime.
- B. Automate back up and use the rerelease stage in the Deployment Pipeline.
- C. Add Rescue and Trigger stages to automatically trigger the failed deployment.
- D. Roll back the failed stage in the pipeline to the previous successful released version

Answer: D (LEAVE A REPLY)

When a deployment stage fails in a OCI DevOps Deployment Pipeline, the recommended approach is to roll back to the previous successful version to ensure that the production environment remains in a stable state. This approach minimizes disruption and ensures that your system continues to function properly while the issue with the failed deployment is investigated and fixed.

NEW QUESTION: 4

As a DevOps engineer working on a CI/CD pipeline for your company's application, you have completed code analysis, image scanning, and automated testing.

What is the next step to ensure a secure and reliable deployment?

- A. Add a traffic Shift stage to route the traffic between two sets of backend IPs.
- B. Add an invoke function stage to run code or custom logic in a serverless manner.
- C. Add a shell stage to run custom commands in the deployment pipeline.

D. Add an approval stage to pause the deployment for a specified duration for manual decision from the approver.

Answer: D (LEAVE A REPLY)

After completing code analysis, image scanning, and automated testing, the next step in the CI/CD pipeline should include a manual review to ensure that all necessary security and quality checks have been performed correctly. Adding an approval stage helps ensure that a secure and reliable deployment is achieved by requiring human verification and approval before proceeding with the deployment to production.

This step adds an extra layer of control to prevent unintended issues from moving forward without further review. It is a common practice in CI/CD pipelines to have an approval step, especially for critical deployments.

NEW QUESTION: 5

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: (SHOW ANSWER)

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: 6

A small company is moving to a DevOps framework to better accommodate their intermittent workloads, which are dynamic and irregular. They want to adopt a consumption-based pricing model.

Which Oracle Cloud Infrastructure service can be used as a target deployment environment?

- A. Virtual machine compute instance
- B. Oracle Kubernetes (OKE)
- C. Bare metal compute instance
- D. Functions

Answer: D (LEAVE A REPLY)

Oracle Cloud Infrastructure Functions is a serverless compute service that supports a consumption-based pricing model. This means that you are only charged for the compute

resources when your function is invoked. This is ideal for intermittent, dynamic, and irregular workloads since the company does not need to provision infrastructure in advance, and costs are directly tied to usage.

NEW QUESTION: 7

You're using Oracle Cloud Infrastructure (OCI) DevOps to deploy your application on an Oracle Container Engine for Kubernetes (OKE) environment. You push your code to the OCI Code Repository, add all the required stage and configure the build and deployment pipeline. When you run the build, you see "unable to clone the repository" error.

What could the configuration error be?

- A.** The Docker image used in the pipeline is incompatible with the OKE environment.
- B.** Dynamic Groups and OCI IAM policies to access the code repository are missing.
- C.** The OKE cluster is not configured to allow external access to the code repository.
- D.** CA bundle for Transport Layer Security (TLS) verification to download the build source during the build run is missing.

Answer: ([SHOW ANSWER](#))

The error "unable to clone the repository" typically indicates that there is an authentication or authorization issue preventing access to the OCI Code Repository. In OCI DevOps, the build pipeline must have the appropriate permissions to access the code repository.

To allow the build pipeline to clone the repository, you need to set up Dynamic Groups and OCI IAM policies that provide the necessary permissions for the build runner to access the code repository.

NEW QUESTION: 8

As a developer working on a software project to be deployed on OKE, you have created a Helm chart for your application. You want to deploy the chart using OCI DevOps service.

Which statement is true about deploying a Helm Chart using OCI DevOps service?

- A.** A single chart can be installed only once into the same cluster, and create single release.
- B.** The Helm charts must be located in the OCI Container Registry repository for deployment.
- C.** The values.yaml file is a generic file located in the OCI Container Registry. You must create a reference to this file.
- D.** Oracle Cloud Infrastructure (OCI) DevOps service supports deployment of Helm charts to Container Instances.

Answer: **B** ([LEAVE A REPLY](#))

OCI DevOps service can be used to deploy Helm charts for Kubernetes-based applications running on Oracle Kubernetes Engine (OKE). For deploying the Helm chart through the OCI DevOps service, the Helm chart must be stored in the OCI Container Registry. This allows the OCI DevOps service to easily access the Helm chart during the deployment process.

NEW QUESTION: 9

You are a DevOps engineer working on a project that requires you to push and pull Docker images to and from Oracle Cloud Infrastructure Registry (Container Registry) using Docker CLI. You have been given access to Container Registry and have installed Docker CLI on your local machine.

Which should you create and use to securely authenticate and store your Docker image in a private Docker registry in OCI?

- A. Auth Token
- B. JSON Web Token
- C. SSH Key Pair
- D. Master Encryption Key in OCI Vault

Answer: [\(SHOW ANSWER\)](#)

To authenticate with the Oracle Cloud Infrastructure Registry (Container Registry) when using the Docker CLI, you need to use an Auth Token. The Auth Token is created in the OCI console and acts as a password for the Docker login command, providing secure access to the container registry.

NEW QUESTION: 10

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: 11

Which is NOT a valid log category for the Oracle Cloud Infrastructure Logging service?

- A. Custom Logs
- B. Audit Logs
- C. Service Logs
- D. Execution Logs

Answer: [D \(LEAVE A REPLY\)](#)

The Oracle Cloud Infrastructure Logging service has several log categories that users can utilize:

Custom Logs: Users can create their own logs to collect application-specific data.

Audit Logs: These logs contain audit records that provide information about activities performed on resources, generated by OCI Audit service.

Service Logs: These are logs generated by OCI services, providing information about the functioning and performance of the services.

Execution Logs is not a valid log category in OCI Logging. There are no "execution logs" as a defined category for the Logging service.

NEW QUESTION: 12

Which two are prerequisites for creating a secret in Oracle Cloud Infrastructure Vault service? (Choose two.)

- A. You must first create a hash digest of the secret value.
- B. You must have the required permissions to create and manage secrets in the Vault service.
- C. You must have a Vault managed key to encrypt the secret.
- D. You must have an auth token to encrypt the secret.
- E. The user must create a compute instance to run the secret service.

Answer: B,C (LEAVE A REPLY)

You need the required permissions (such as policies allowing secret management) to create and manage secrets in Oracle Cloud Infrastructure (OCI) Vault service. These permissions are essential for performing operations such as creating, reading, and managing secrets.

Vault managed key is required to encrypt the secret before it is stored in the OCI Vault. The managed key acts as the encryption key for securing the secret, ensuring its confidentiality.

NEW QUESTION: 13

You are using the Oracle Cloud Infrastructure (OCI) DevOps service and you have successfully built and tested your software applications in your Build Pipeline. The resulting output needs to be stored in a container repository.

Which stage should you add next to your Build Pipeline?

- A. Trigger deployment
- B. Managed build
- C. Deliver artifacts
- D. Export packages

Answer: C (LEAVE A REPLY)

Step 1: Understanding the Requirement

The objective is to store the resulting build output from a Build Pipeline in a container repository.

In OCI DevOps, the build output is stored as an artifact, which can include Docker images or other build-generated files. To store these artifacts in a container repository, you need to explicitly deliver artifacts in the pipeline.

Step 2: Explanation of the Options

- A . Trigger deployment

This stage is used to trigger a deployment pipeline, which comes after the artifacts are already stored and prepared for deployment.

Not applicable: This stage is downstream of storing artifacts and is used for deploying software, not for saving the build output to a repository.

B . Managed build

The managed build stage is where you compile, test, and package the application. This has already been completed successfully according to the question.

Not applicable: The question specifies that the build has been completed, so this stage is not relevant at this point.

C . Deliver artifacts

The Deliver Artifacts stage in OCI DevOps pipelines is designed to store the output of the build process in an artifact repository, such as:

OCI Container Registry (OCIR) for Docker images.

Artifact Registry for build artifacts like binaries or JAR files.

Applicable and Correct answer: This is the correct next step for storing the resulting D . Export packages This is not a standard OCI DevOps pipeline stage. It may be relevant in other contexts but is not related to OCI DevOps for storing build artifacts.

Step 3: Key Concepts of "Deliver Artifacts" in OCI DevOps

Purpose: Save build outputs (artifacts) to an artifact repository.

Artifact Types: Includes Docker container images, binaries, JAR files, or other build outputs.

Repositories Supported:

OCI Container Registry (OCIR)

OCI Artifact Registry

Configuration:

Specify the artifact source (build stage output).

Define the destination repository (e.g., OCIR).

Step 4: References and OCI Resources

OCI DevOps Build Pipelines:

Build Pipeline Documentation

Deliver Artifacts Stage

OCI Container Registry (OCIR):

OCI Container Registry Overview

OCI Artifact Registry:

OCI Artifact Registry Overview

NEW QUESTION: 14

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: 15

Which command creates the docker registry secret required in the application manifests for OKE to pull images from Oracle Cloud Infrastructure Registry?

A.

```
kubectl create secret docker-registry <secret-name> --docker-server=  
<region-key>.ocir.io --docker-username=<oci-username> --docker-  
password='<oci-auth-token>' --docker-email=<email-address>
```

B.

```
kubectl create passwd docker-registry <secret-name> --docker-server=  
<region-key>.oke.io --docker-username=<tenancy-namespace>/<oci-  
username> --docker-password='<oci-auth-token>' --docker-email=<email-  
address>
```

C.

```
kubectl create passwd docker-registry <secret-name> --docker-server=  
<region-key>.oke.io --docker-username=<oci-username> --docker-  
password='<oci-auth-token>' --docker-email=<email-address>
```

D.

```
kubectl create secret docker-registry <secret-name> --docker-server=  
<region-key>.ocir.io --docker-username=<tenancy-namespace>/<oci-  
username> --docker-password='<oci-auth-token>' --docker-email=<email-  
address>
```

Answer: D (LEAVE A REPLY)

To create a Docker registry secret to pull images from the Oracle Cloud Infrastructure Registry (OCIR), you need to specify the correct parameters such as the region key, namespace, OCI username, and OCI authentication token.

Chosen command is correct because:

The `kubectl create secret docker-registry` command creates a Docker registry secret.

The `--docker-server=<region-key>.ocir.io` specifies the correct endpoint for OCIR.

The `--docker-username=<tenancy-namespace>/<oci-username>` provides both the tenancy namespace and the OCI username, which is the required format for authentication with OCIR. The `--docker-password='<oci-auth-token>'` specifies the OCI auth token, which acts as a password for authentication.

The `--docker-email=<email-address>` is also included.

The other commands have errors, such as missing tenancy namespace or using incorrect flags (`passwd` instead of `secret`).

NEW QUESTION: 16

How do OCI DevOps Deployment Pipelines reduce risk and complexity of production applications?

- A. By reducing change-driven errors introduced by manual deployments
- B. By scaling builds with service-managed build runners
- C. By working with existing Git repositories and CI systems
- D. By eliminating downtime of production applications

Answer: A (LEAVE A REPLY)

OCI DevOps Deployment Pipelines automate the process of deploying applications to production environments. By using automated, repeatable deployment processes, they help reduce the risk of change-driven errors, which are often introduced during manual deployments. This automation reduces human errors and ensures consistency across environments, thus minimizing complexity and risk in production.

Valid 1z0-1109-24 Dumps shared by TrainingQuiz.com for Helping Passing 1z0-1109-24 Exam! TrainingQuiz.com now offer the **newest 1z0-1109-24 exam dumps**, the TrainingQuiz.com 1z0-1109-24 exam **questions have been updated** and **answers have been corrected** get the **newest** TrainingQuiz.com 1z0-1109-24 dumps with Test Engine here: <https://www.trainingquiz.com/1z0-1109-24-practice-quiz.html> (52 Q&As Dumps, **40%OFF**

Special Discount: Exam-Tests)

NEW QUESTION: 17

What are the two items required to create a rule for the Oracle Cloud Infrastructure Events Service? (Choose two.)

- A. Rule Conditions
- B. Actions
- C. Management Agent Cloud Service
- D. Auth Token
- E. Service Connector

Answer: A,B (LEAVE A REPLY)

NEW QUESTION: 18

As a DevOps engineer working with OCI DevOps, you are managing artifacts for a microservices application.

Based on your understanding of working with DevOps projects and artifacts, which statement is true?

- A. Artifacts can be used directly by OCI DevOps without the need for them to be located or mirrored in an OCI Artifact or Container registry.
- B. In the build pipeline, to store the Managed Build stage outputArtifacts. you need an OCI Object storage.
- C. It is recommended to make artifacts immutable to prevent any modifications after they are uploaded.
- D. Once created, the artifact's name, type, and source cannot be modified.

Answer: C (LEAVE A REPLY)

Making artifacts immutable ensures that the build artifacts are not altered after being published. This is a best practice to maintain the integrity and consistency of the artifacts used in deployments, preventing unintentional changes that could introduce issues during subsequent deployment stages.

NEW QUESTION: 19

Which of the following is NOT considered a DevOps resource in the context of the OCI DevOps project service?

- A. Environments
- B. API integrations
- C. Build pipelines
- D. Code repositories

Answer: B (LEAVE A REPLY)

In the context of OCI DevOps project service, the key DevOps resources include:

Environments: These are used to represent deployment targets, such as Kubernetes clusters, compute instances, or functions.

Build pipelines: These automate the building, testing, and packaging of software.

Code repositories: These store source code for the application being built, providing version control and collaboration.

API integrations are not directly considered a DevOps resource in OCI DevOps projects. Instead, they facilitate communication and integration with other tools or services but do not represent a core component of DevOps projects like environments, build pipelines, or code repositories.

NEW QUESTION: 20

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: 21

As a DevOps engineer at XYZ Corp, you have been assigned the task of setting up a new OKE (Oracle Kubernetes Engine) cluster to manage the organization's Kubernetes applications hosted on Oracle Cloud Infrastructure (OCI). Your goal is to ensure a smooth and efficient process while preparing for the cluster creation.

Which of the following statements is false regarding the preparation process for setting up a new OKE cluster?

- A. Container Engine for Kubernetes cannot utilize existing network resources for the creation of the new cluster.
- B. Container Engine for Kubernetes automatically creates and configures new network resources for the new cluster.
- C. It is necessary to ensure sufficient quota on different resource types in your OCI tenancy for the cluster setup.
- D. Access to an Oracle Cloud Infrastructure tenancy is required to set up the new OKE cluster.

Answer: (SHOW ANSWER)

This statement is false because Container Engine for Kubernetes (OKE) can utilize existing network resources such as Virtual Cloud Networks (VCNs), subnets, security lists, and route tables for the creation of a new cluster. You can either use pre-existing network resources or let OKE create new network resources automatically.

NEW QUESTION: 22

A software development team is working on a prototype using Node.js and MongoDB as their programming language and database, respectively. They need to develop and test isolated web applications or RESTful APIs. They are looking for a simple, quick, and secure way to run containers without managing any servers.

Which OCI service is best suited for this use case?

- A. OCI Functions
- B. Container Instances
- C. OCI DevOps Project

D. Compute Instances

Answer: (SHOW ANSWER)

OCI Container Instances provide a serverless way to run containers in Oracle Cloud Infrastructure (OCI). They allow you to run containers without having to manage any underlying servers or clusters. This is ideal for development teams that need a simple, quick, and secure way to develop and test isolated web applications or RESTful APIs.

NEW QUESTION: 23

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: (SHOW ANSWER)

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.

Valid 1z0-1109-24 Dumps shared by TrainingQuiz.com for Helping Passing 1z0-1109-24 Exam! TrainingQuiz.com now offer the **newest 1z0-1109-24 exam dumps**, the TrainingQuiz.com 1z0-1109-24 exam **questions have been updated** and **answers have been corrected** get the **newest** TrainingQuiz.com 1z0-1109-24 dumps with Test Engine here: <https://www.trainingquiz.com/1z0-1109-24-practice-quiz.html> (52 Q&As Dumps, **40%OFF** Special Discount: **Exam-Tests**)