

Well-Prepared 1Z0-1084-25 New Learning Materials– Verified Exam Engine for 1Z0-1084-25: Oracle Cloud Infrastructure 2025 Developer Professional



P.S. Free & New 1Z0-1084-25 dumps are available on Google Drive shared by Itexamguide: https://drive.google.com/open?id=182epBhOPLrQuCYarPwvdSZ_N2zPso6sx

Want to crack the Oracle 1Z0-1084-25 certification test in record time? Look no further than Itexamguide! Our updated 1Z0-1084-25 Dumps questions are designed to help you prepare for the exam quickly and effectively. With study materials available in three different formats, you can choose the format that works best for you. Trust Itexamguide to help you pass the Oracle 1Z0-1084-25 Certification test with ease.

Compared with other education platform on the market, Itexamguide is more reliable and highly efficiently. It provide candidates who want to pass the 1Z0-1084-25 exam with high pass rate study materials, all customers have passed the exam in their first attempt. They all need 20-30 hours to learn on our website can pass the exam. 1Z0-1084-25 Exam Dump is really a high efficiently exam tool that can help you save much time and energy to do other things.

>> 1Z0-1084-25 New Learning Materials <<

1Z0-1084-25 Exam Engine, 1Z0-1084-25 100% Correct Answers

Our company is a professional certificate test materials provider, and we have rich experiences in providing exam materials. 1Z0-1084-25 exam materials are reliable, and we can help you pass the exam just one time. 1Z0-1084-25 exam dumps are also known as high pass rate, and the pas rate reaches 98.95%. We are pass guaranteed and money back guaranteed in case you fail to pass the exam. Moreover, we have free demo for 1Z0-1084-25 Exam Materials for you to have a general understanding of the product.

Oracle 1Z0-1084-25 Exam Syllabus Topics:

Topic	Details
Topic 1	<ul style="list-style-type: none"> Leveraging Serverless Technologies for Cloud Native Development: This section of the exam measures the skills of professionals in serverless development within OCI. It covers creating serverless applications using Oracle Functions, building API gateways for routing traffic, and integrating systems through OCI Streaming Service. Additionally, it explores event-driven architectures using OCI Event Service and how OCI Queue enables asynchronous messaging between microservices.

Topic 2	<ul style="list-style-type: none"> • Cloud Native Fundamentals: This section of the exam measures the skills of target audience and covers the essential principles of cloud-native development. It explains the core concepts, key pillars, and advantages of cloud-native applications. The section also focuses on microservices architecture, including its design methodology and how it supports scalable, distributed applications.
Topic 3	<ul style="list-style-type: none"> • Testing and Securing Cloud-Native Applications: This section focuses on testing strategies and security for cloud-native applications. It discusses different testing methodologies, securing sensitive information using OCI Vault, and implementing security measures to address cloud-native development challenges.
Topic 4	<ul style="list-style-type: none"> • Cloud Native Applications and Containerization: This section of the exam covers containerization technologies for cloud-native applications. It explains Docker architecture, its components, and the process of pulling and pushing container images using Oracle Cloud Infrastructure Registry (OCIR). It also explores container orchestration, deploying applications on Oracle Kubernetes Engine (OKE), and using OCI Service Mesh for Kubernetes deployments.
Topic 5	<ul style="list-style-type: none"> • Monitoring & Troubleshooting Cloud-Native Applications: This section of the exam focuses on monitoring and troubleshooting cloud-native applications. It covers using OCI Monitoring to track metrics, OCI Logging for managing logs and performing tasks related to monitoring, logging, and tracing for better observability and issue resolution.

Oracle Cloud Infrastructure 2025 Developer Professional Sample Questions (Q23-Q28):

NEW QUESTION # 23

Which technique is used for testing the entire user flow as well as the moving parts of a cloud native app, ensuring that there are no high-level discrepancies?

- A. Contract Testing
- B. Component Testing
- C. Unit Testing
- **D. End-to-end Testing**
- E. Integration Testing

Answer: D

Explanation:

End-to-end testing is a technique that involves checking the entire user flow as well as the moving parts of a cloud native app, ensuring that there are no high-level discrepancies³. End-to-end testing simulates real user scenarios and validates the functionality, performance, reliability, and security of the app from start to finish³. End-to-end testing has several benefits, such as³:

Comprehensive testing: You can test your app as a whole and verify that all the components work together as expected.

User-centric testing: You can test your app from the user's perspective and ensure that it meets the user's needs and expectations.

Quality assurance: You can test your app in a realistic environment and identify any issues or defects before releasing it to the users.

NEW QUESTION # 24

A Docker image consists of one or more layers, each of which represents a Dockerfile instruction. The layers are stacked and each one is a delta of the changes from the previous layer. What permission is associated with these layers?

- A. write only
- B. movable
- C. write once
- **D. read only**
- E. read mostly

Answer: D

Explanation:

The correct answer is: "read only." The layers of a Docker image are read-only. Once a layer is created, it cannot be modified. Each

layer represents a Dockerfile instruction, and it is stacked on top of the previous layer, forming a stack of immutable layers. These layers are designed to be read-only to ensure consistency and integrity of the image. When a Docker image is built, each instruction in the Dockerfile creates a new layer. Each layer represents the changes made by that instruction relative to the previous layer. The layers are stacked on top of each other to form the complete image. This layer-based approach allows for efficient storage and distribution of Docker images. Because the layers are read-only, any changes or modifications to the image result in the creation of new layers rather than modifying the existing ones. This immutability ensures that each layer remains intact and preserves the integrity of the image. It also enables Docker's caching mechanism, where previously built layers can be reused if the corresponding instructions haven't changed, speeding up the image build process. The other options mentioned, such as "write only," "write once," "movable," and "read mostly," do not accurately describe the permission associated with Docker image layers. Docker image layers are specifically designed to be read-only.

NEW QUESTION # 25

You developed a microservices-based application that runs in an Oracle Cloud Infrastructure (OCI) Container Engine for Kubernetes (OKE) cluster. It has multiple endpoints that need to be exposed to the public internet. What is the most cost-effective way to expose multiple application endpoints without adding unnecessary complexity to the application?

- A. Create a separate load balancer instance for each service using the lowest 100 Mbps option.
- B. Use a ClusterIP service type in Kubernetes for each of your service endpoints using a load balancer to expose the endpoints.
- C. Use a NodePort service type in Kubernetes for each of your service endpoints using the node's public IP address to access the applications.
- **D. Deploy an Ingress Controller and use it to expose each endpoint with its own routing endpoint.**

Answer: D

Explanation:

An Ingress Controller is a Kubernetes resource that provides advanced routing and load balancing for your applications running on a Kubernetes cluster. An Ingress Controller allows you to define rules that specify how to route traffic to different services in your cluster based on the host name or path of the incoming request. By deploying an Ingress Controller and using it to expose multiple application endpoints, you can achieve the following benefits:

Cost-effectiveness: You only need to create one load balancer instance per cluster, instead of one per service, which reduces the cost of exposing your applications.

Simplicity: You only need to manage one set of routing rules for all your services, instead of configuring each service separately, which simplifies the application deployment and maintenance.

Flexibility: You can use different types of Ingress Controllers, such as NGINX or Traefik, that offer various features and customization options for your routing needs.

NEW QUESTION # 26

Your team has created a serverless application deployed in Oracle Functions. It uses a Python function leveraging the Oracle Cloud Infrastructure (OCI) Python SDK to stop any OCI compute instance that does not comply with your corporate security standards. Although there are three non-compliant OCI compute instances, when you invoke this function, none of the instances were stopped. With respect to this issue, which of the following is a valid troubleshooting strategy?

- A. Ensure that the application is deployed within the same OCI compartment as the instance, because you cannot enable function execution data from the OCI console.
- **B. Enable function logging in the OCI console, add some print statements in your function code, and then view the logs to troubleshoot.**
- C. Enable function remote debugging in the OCI console, and then use your favorite IDE to inspect the function running on Oracle Functions.
- D. Enable function tracing in the OCI console, and then go to the OCI Monitoring console to view the function stack trace.

Answer: B

Explanation:

The valid troubleshooting strategy in this scenario is to enable function logging in the OCI console, add some print statements in your function code, and then view the logs to troubleshoot. Enabling function logging allows you to capture and store logs generated by your function during its execution. By adding print statements or log statements in your function code, you can output relevant information and debug messages to the logs. This helps you understand the execution flow, identify any errors or issues, and gather more information about the function's behavior. To troubleshoot the issue of the Python function not stopping the non-compliant OCI

compute instances, you can follow these steps: Enable function logging in the OCI console: Enable logging for your function to ensure that logs are captured during its execution. Modify your function code: Add relevant print statements or log statements at key points in your code to output debug information or verify the execution flow. For example, you can print the instance details that are being evaluated for compliance. Invoke the function: Trigger the function execution either through an event or manually. View the logs: Access the function logs in the OCI console or retrieve them programmatically. Look for the expected print statements or log entries that indicate the status of each instance and the decisions made by the function. By reviewing the logs, you can analyze the output and identify any issues or discrepancies. It can help you determine if the function is correctly evaluating the compliance criteria, retrieving the instance details, or making the necessary API calls to stop the instances. You may need to adjust your code logic or investigate further based on the information provided in the logs. Enabling function remote debugging is not a suitable strategy in this case because it is primarily used for inspecting and debugging the function code during development, rather than troubleshooting issues in a deployed function. Enabling function tracing can provide insights into the execution flow and performance of the function but may not directly address the issue of the instances not being stopped. Ensuring that the application is deployed within the same OCI compartment as the instance is not directly related to troubleshooting the issue with the non-compliant instances. It is a consideration for access and permissions but does not provide specific insights into the problem at hand. Remember to refer to the Oracle Functions documentation and consult the official resources for detailed instructions and best practices on troubleshooting and monitoring Oracle Functions.

NEW QUESTION # 27

You need to push a new Docker container image to a repository in the Oracle Cloud Infrastructure (OCI) Registry. Which mechanism must you use to provide authentication?

- A. Generate an Auth Token to complete the authentication via the OCI CLI.
- B. Generate an API signing key to complete the authentication via the Docker CLI.
- C. Generate an API signing key to complete the authentication via the OCI CLI.
- **D. Generate an Auth Token to complete the authentication via the Docker CLI.**

Answer: D

Explanation:

To push a new Docker container image to a repository in OCI Registry, you need to use an Auth Token to complete the authentication via the Docker CLI. An Auth Token is a secure, auto-generated password that you can use to authenticate with OCI services such as OCI Registry. You can generate an Auth Token in the Console by following these steps:

In the top-right corner of the Console, open the Profile menu and then click User settings to view the details.

On the Auth Tokens page, click Generate Token.

Enter a friendly description for the auth token. Avoid entering confidential information.

Click Generate Token. The new auth token is displayed.

Copy the auth token immediately to a secure location from where you can retrieve it later, because you won't see the auth token again in the Console.

Close the Generate Token dialog. After generating an Auth Token, you need to log in to OCI Registry by entering `docker login <region-key>.ocir.io` in a terminal window on the client machine running Docker, where `<region-key>` corresponds to the key for the OCI Registry region you're using. When prompted for a username, enter your username in the format `<tenancy-namespace>/<username>`, where `<tenancy-namespace>` is the auto-generated Object Storage namespace string of your tenancy. When prompted for a password, enter the Auth Token you copied earlier.

NEW QUESTION # 28

.....

After the client pay successfully they could receive the mails about 1Z0-1084-25 guide questions our system sends by which you can download our test bank and use our study materials in 5-10 minutes. The mail provides the links and after the client click on them the client can log in and gain the 1Z0-1084-25 Study Materials to learn. For the client the time is limited and very important and our product satisfies the client's needs to download and use our 1Z0-1084-25 practice engine immediately.

1Z0-1084-25 Exam Engine: https://www.itexamguide.com/1Z0-1084-25_braindumps.html

- Lab 1Z0-1084-25 Questions 1Z0-1084-25 Valid Exam Book 1Z0-1084-25 Exam Questions Fee Open website [www.pass4test.com] and search for **> 1Z0-1084-25** for free download 1Z0-1084-25 Valid Dump
- Valid 1Z0-1084-25 Exam Papers New 1Z0-1084-25 Exam Testking Free Sample 1Z0-1084-25 Questions The page for free download of 1Z0-1084-25 on www.pdfvce.com will open immediately Braindump 1Z0-1084-25 Pdf

