

The Tester's Handbook: 1Z0-1084-25 Online Test Engine



BTW, DOWNLOAD part of Actual4Cert 1Z0-1084-25 dumps from Cloud Storage: <https://drive.google.com/open?id=1uCR7lRnO9g298MdvIITr3Wpvn1uTyxy>

In the course of your study, the test engine of 1Z0-1084-25 actual exam will be convenient to strengthen the weaknesses in the learning process. This can be used as an alternative to the process of sorting out the wrong questions of 1Z0-1084-25 learning guide in peacetime learning, which not only help you save time, but also makes you more focused in the follow-up learning process with our 1Z0-1084-25 learning materials.

Oracle 1Z0-1084-25 Exam Syllabus Topics:

| Topic | Details |
|---------|--|
| Topic 1 | <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 2 | <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 3 | <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 4 | <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. |
| Topic 5 | <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. |

Latest 1Z0-1084-25 Exam Question - 1Z0-1084-25 Top Questions

You don't have to worry about your problems on our 1Z0-1084-25 exam questions are too much or too simple. Our staff will give you a smile and then answer them carefully. All we do is just want you to concentrate on learning on our 1Z0-1084-25 study guide! Let other things go to us. And as long as you focus on our 1Z0-1084-25 Training Materials, we believe you will pass for sure for our 1Z0-1084-25 practice braindumps are always the latest and valid for all of our customers.

Oracle Cloud Infrastructure 2025 Developer Professional Sample Questions (Q95-Q100):

NEW QUESTION # 95

You have two microservices, A and B running in production. Service A relies on APIs from service B. You want to test changes to service A without deploying all of its dependencies, which includes service B. Which approach should you take to test service A?

- A. Test using a previous test version of service B.
- B. This is not possible because service B is a dependency.
- **C. Test using an API mock of service B.**
- D. Test using the current production version of service B.

Answer: C

Explanation:

The correct answer is: Test using an API mock of service B. To test service A without deploying all of its dependencies, including service B, you can use an API mock of service B. An API mock is a simulated version of the API that mimics the behavior of the actual service. By using an API mock, you can isolate the testing of service A and simulate the responses and behavior of service B's APIs. With an API mock, you can define the expected responses and behavior of service B's APIs, allowing you to test the integration between service A and the mocked version of service B. This approach enables you to verify the functionality of service A without relying on the availability or changes in the actual service B. By decoupling the dependencies and using an API mock, you can perform independent testing of service A, ensuring its functionality in isolation.

NEW QUESTION # 96

Which TWO are required to access the Oracle Cloud Infrastructure (OCI) Container Engine for Kubernetes (OKE) cluster from the kubectl CLI? (Choose two.)

- **A. A configured OCI API signing key pair.**
- B. An SSH key pair with the public key added to the cluster worker nodes.
- C. Tiller enabled on the OKE cluster.
- **D. Install and configure the OCI CLI.**
- E. OCI Identity and Access Management (IAM) Auth Token.

Answer: A,D

Explanation:

The correct options are: A configured OCI API signing key pair: The API signing key pair is used for authentication and authorization to access OCI resources, including the OKE cluster. The private key should be configured on your local machine to authenticate API requests. An SSH key pair with the public key added to the cluster worker nodes: This is required for secure SSH access to the worker nodes in the OKE cluster. You need to generate an SSH key pair and add the public key to the cluster's worker node pool during cluster creation or update. Therefore, the correct options are having a configured OCI API signing key pair and an SSH key pair with the public key added to the cluster worker nodes.

NEW QUESTION # 97

You have been asked to update an OKE cluster to a network configuration that has the least attack surface while the deployed applications are still directly available for access from the Internet. Which is a valid OKE cluster network configuration that meets this requirement? (Choose the best answer.)

- **A. Private subnets for nodes and the Kubernetes API endpoint; public subnets for load balancers**
- B. Private subnet for the Kubernetes API endpoint; public subnets for nodes and load balancers
- C. Private subnets for nodes; public subnets for the Kubernetes API endpoint and load balancers

- D. Private subnets for nodes, the Kubernetes API endpoint, and load balancers

Answer: A

Explanation:

The valid OKE cluster network configuration that meets the requirement of having the least attack surface while still allowing direct access to the deployed applications from the Internet is: Private subnets for nodes and the Kubernetes API endpoint; public subnets for load balancers. By placing the nodes and the Kubernetes API endpoint in private subnets, they are not directly accessible from the Internet, reducing the attack surface. The load balancers, on the other hand, are placed in public subnets, allowing them to be accessed from the Internet and serve as the entry point for accessing the deployed applications. This configuration ensures that the critical components of the cluster, such as the nodes and the API endpoint, are protected within the private network, while still providing accessibility to the applications through the load balancers. It helps to enhance security by limiting direct access to the internal components of the cluster while maintaining the availability of the deployed applications.

NEW QUESTION # 98

You plan to implement logging in your services that will run in Oracle Cloud Infrastructure (OCI) Container Engine for Kubernetes (OKE). Which statement describes the appropriate logging approach?

- A. All services log to standard output only.
- B. All services log to an external logging system.
- C. All serviceAAs log to a shared log file.
- D. Each service logs to its own log file.

Answer: A

Explanation:

The appropriate logging approach for services running in Oracle Cloud Infrastructure (OCI) Container Engine for Kubernetes (OKE) is: "All services log to standard output only." When running services in a containerized environment like OKE, it is recommended to follow the Twelve-Factor App methodology, which suggests treating logs as event streams. According to this methodology, services should write their log events to standard output (stdout) instead of writing to log files. By logging to standard output, the container runtime (such as Kubernetes) can collect and aggregate the logs generated by the services. These logs can then be accessed and managed through the container runtime's logging infrastructure. Logging to standard output offers several advantages in a containerized environment: Simplicity and consistency: Standardizing on logging to stdout ensures a consistent approach across different services, making it easier to manage and analyze logs. Log aggregation: The container runtime can collect the logs from all the running containers and provide centralized log management, allowing you to access and search logs from different services in one place. Scalability: Since logs are written to stdout, they can be easily handled by the container runtime's log management system, which can scale to handle large volumes of log data. Separation of concerns: By logging to stdout, the responsibility of managing log files and their rotation is shifted to the container runtime, allowing the services to focus on their core functionality. While it is possible to log to log files or external logging systems, the recommended approach in a containerized environment like OKE is to log to standard output and leverage the logging infrastructure provided by the container runtime.

NEW QUESTION # 99

Which TWO statements are correct regarding Docker images and containers? (Choose two.)

- A. If multiple containers are spawned from the same image, then they all use the same copy of that image in memory.
- B. Only one container can be spawned from a given image at a time.
- C. An image is a collection of immutable layers whereas a container is a running instance of an image.
- D. Writing and building a new Dockerfile is the only way you can create new Docker images.
- E. A container can exist without an image but an image cannot exist without a container.

Answer: A,C

Explanation:

Which TWO statements are correct regarding Docker images and containers? (Choose two.) Writing and building a new Dockerfile is the only way you can create new Docker images. A container can exist without an image but an image cannot exist without a container. Only one container can be spawned from a given image at a time. If multiple containers are spawned from the same image, then they all use the same copy of that image in memory. (Correct) An image is a collection of immutable layers whereas a container is a running instance of an image. (Correct) ChatGPT The correct statements regarding Docker images and containers are: If multiple containers are spawned from the same image, then they all use the same copy of that image in memory. This statement is

BONUS!!! Download part of Actual4Cert 1Z0-1084-25 dumps for free: <https://drive.google.com/open?>

id=1uCR7lRnO9g298MdvllTrB3Wpvn1uTyxy