Free Questions for C1000-172

Shared by Banks on 04-10-2024

For More Free Questions and Preparation Resources

Question 1

What is an advantage of choosing IBM Cloud VMware for regulated workloads?

Options:

A- Offers a secure-by-default architecture

B- Offers the ability to bring your own VMware license

C- IBM offers to manage the guest OS

D- Offers the ability to bring your own AIX license

Answer:

A

Explanation:

IBM Cloud VMware offers a 'secure-by-default' architecture, which is particularly beneficial for regulated workloads. This means that security best practices are integrated into the default configuration, providing enhanced security controls, compliance readiness, and

data protection measures from the outset.

Why Secure-by-Default is Advantageous: It simplifies compliance for regulated workloads by ensuring that security is baked into the infrastructure, reducing the effort required to meet regulatory requirements.

Comparison with Other Options:

B (Bring your own VMware license): This is a feature, but not the main advantage for regulated workloads.

C (Manage the guest OS): Not specific to the security or regulatory compliance needs.

D (Bring your own AIX license): Irrelevant for VMware workloads.

IBM Cloud VMware Solutions

IBM Cloud Architect Exam Study Guide

Question 2

Question Type: MultipleChoice

What are the basic components in an event-driven architecture?

Options:

A- Event Store, Switch, and Event Consumer

B- Event Consumer, Switch, and Event Producer

C- Event Producer, Router, and Event Consumer

D- Event Producer, Router, and Event Store

Answer:

D

Explanation:

In an event-driven architecture, the basic components are:

Event Producer: The source that generates events, such as a service or application emitting a message whenever a significant change or action occurs.

Router (sometimes called an Event Router or Event Bus): Routes the event to the appropriate consumer(s). The router can handle complex event processing, filtering, and transformation before delivering the event.

Event Store: A component that persists events, making them available for future analysis, auditing, or replaying if needed.

Why These Components are Core:

These components are essential to ensure that events are generated, routed to the correct consumers, and stored for traceability and recovery purposes. This architecture is crucial for building scalable, decoupled, and responsive systems.

Comparison with Other Options:

A, B, and C: Do not provide a complete representation of all three core components needed in an event-driven architecture.

IBM Cloud Event-driven Architecture

IBM Cloud Architect Exam Study Guide

Question 3

A data analyst working for a retail company uses IBM Cloud Watson Discovery service to analyze client dat

a. The company's customer support team has been experiencing an increase in customer complaints regarding delayed deliveries. The analyst has been asked to identify the root cause of this issue.

Which feature of Watson Discovery would be most helpful in this situation?

Options:

A- AI-powered search engine

B- Natural language processing

C- User behavior tracking

D- Pattern and trend analysis

Answer:

D

Explanation:

The feature of Watson Discovery most helpful in identifying the root cause of customer complaints regarding delayed deliveries is Pattern and Trend Analysis.

Pattern and Trend Analysis in Watson Discovery: This feature helps to identify recurring themes, trends, or patterns within large datasets. For a data analyst investigating customer complaints, pattern and trend analysis can help reveal common factors or trends contributing to the delays, such as issues with specific delivery locations, times, or processes.

Use Case for Analyzing Customer Complaints: By analyzing patterns and trends in the text data from customer complaints, the analyst can pinpoint potential bottlenecks or root causes leading to delays in deliveries.

Reference from IBM Cloud Professional Architect Materials:

IBM's Watson Discovery documentation highlights its ability to perform pattern and trend analysis to extract meaningful insights from unstructured data.

Other options are incorrect:

A . AI-powered search engine is for retrieving relevant documents or information.

B . Natural language processing processes text but does not focus on identifying patterns or trends.

C . User behavior tracking is not relevant in this context.

Question 4

Microservices need to be deployed to a platform where container communication can be orchestrated and customized. Which IBM Cloud solution should be selected to meet this requirement?

Options:

A- IBM Cloud Kubernetes Service

B- IBM Cloud Code Engine

C- IBM Cloud for VMware Solutions

D- IBM Cloud Functions

Answer:

A

Explanation:

To deploy microservices where container communication can be orchestrated and customized, the best IBM Cloud solution is the IBM Cloud Kubernetes Service.

IBM Cloud Kubernetes Service: This service provides a managed Kubernetes environment to orchestrate and manage containerized applications. It supports microservices architecture by providing tools for managing container communication, scaling, deployment, and load balancing. Kubernetes allows for custom networking policies, service meshes, and other tools necessary for microservices communication and orchestration.

Microservices and Kubernetes: Kubernetes is designed to manage and scale microservices. It can handle the orchestration of complex, interdependent services and provides features such as service discovery, traffic routing, load balancing, and monitoring.

Reference from IBM Cloud Professional Architect Materials:

IBM documentation on IBM Cloud Kubernetes Service confirms its suitability for managing microservices architecture with container orchestration capabilities.

Other options are incorrect:

B . IBM Cloud Code Engine is a serverless platform, not specifically designed for microservices orchestration.

C . IBM Cloud for VMware Solutions is used for running VMware workloads in the cloud.

D . IBM Cloud Functions is used for running serverless functions, not for orchestrating containers.

Question 5

Question Type: **MultipleChoice**

Which two use cases for IBM Cloud Flow logs can be supported by a single collector?

Options:

A- Collect data for a single network interface in a virtual server instance

B- Collect data for all network interfaces in an IBM Cloud Classic infrastructure

C- Collect data for all network interfaces in an IBM Cloud region

D- Collect data for all network interfaces in an IBM Cloud account

E- Collect data for all network interfaces in a subnet

Answer:
A, E

Explanation:
A single collector in IBM Cloud Flow logs can support the following two use cases:

Collect Data for a Single Network Interface in a Virtual Server Instance: IBM Cloud Flow logs can collect network traffic data specific to a single network interface, which allows detailed monitoring and analysis of traffic patterns and security incidents for that specific instance.

Collect Data for All Network Interfaces in a Subnet: A single collector can also be configured to gather data from all network interfaces in a given subnet. This provides a comprehensive view of the network traffic within that subnet, useful for monitoring, troubleshooting, and security analysis.

Reference from IBM Cloud Professional Architect Materials:

IBM documentation on IBM Cloud Flow Logs supports these use cases, showing how collectors can be set up to monitor specific network interfaces or entire subnets.

Other options are incorrect:

B . Collect data for all network interfaces in an IBM Cloud Classic infrastructure and C. Collect data for all network interfaces in an IBM Cloud region are too broad for a single collector.

D . Collect data for all network interfaces in an IBM Cloud account would require multiple collectors.

Therefore, the correct answers are A. Collect data for a single network interface in a virtual server instance and E. Collect data for all network interfaces in a subnet.

Question 6

Which two are the private registry options available to IBM Cloud OpenShift clusters?

Options:

A- GitHub

B- IBM Cloud Container Registry

C- Docker Hub

D- The internal registry setup at cluster creation time

E- Red Hat Quay

Answer:

B, D

Explanation:

The two private registry options available to IBM Cloud OpenShift clusters are IBM Cloud Container Registry and the internal registry setup at cluster creation time.

IBM Cloud Container Registry: This is a private image registry that integrates with IBM Cloud OpenShift. It allows users to store and manage Docker images in a private repository securely. This option is specifically designed to work seamlessly with Kubernetes and OpenShift clusters in IBM Cloud.

Internal Registry at Cluster Creation Time: IBM Cloud OpenShift also supports an internal image registry that is set up during cluster creation. This internal registry allows for the storage and management of container images directly within the OpenShift cluster, providing a secure and private environment for images that are not meant to be publicly accessible.

Reference from IBM Cloud Professional Architect Materials:

According to IBM Cloud OpenShift documentation on Configuring Private Registries, these two options are available for IBM Cloud OpenShift clusters.

Other options are incorrect:

A . GitHub and C. Docker Hub are public registries.

E . Red Hat Quay is another private registry solution but is not directly listed as an option specific to IBM Cloud OpenShift clusters.

Therefore, the correct answers are B. IBM Cloud Container Registry and D. The internal registry setup at cluster creation time.

Question 7

The client requires 99.99% uptime. Why would a Citrix NetScaler help in this case?

Options:

A- Predictability

B- Scalability

C- Reliability

D- Availability

Answer:

D

Explanation:

A Citrix NetScaler (now known as Citrix ADC) is designed to improve the availability of applications by providing advanced load balancing, traffic management, and redundancy features. It ensures that applications remain accessible even in the event of server failures or high traffic volumes.

Why Citrix NetScaler Improves Availability: It distributes incoming traffic across multiple servers to prevent overload on a single server, thus maintaining application uptime. It also provides failover capabilities, ensuring that if one server fails, traffic is rerouted to other healthy servers.

Comparison with Other Options:

A (Predictability): While Citrix NetScaler can improve predictability, its main function is related to availability.

B (Scalability): Citrix NetScaler can aid in scalability, but this is not the primary benefit for achieving 99.99% uptime.

C (Reliability): While reliability is an aspect, the key feature directly supporting 99.99% uptime is availability.

Citrix ADC Documentation

IBM Cloud Load Balancer Services

IBM Cloud Architect Exam Study Guide

Question 8

A company is using Watson Assistant to improve customer support. The client has integrated the virtual agent technology to handle common inquiries and provide quick assistance.

How can Watson Assistant benefit the company in terms of customer support?

Options:

A- Reduce costs per contact

B- Remove the requirement for all human agents

C- Provide live agents who can speak multiple languages

D- Check insurance claim details for non-native English speakers

Answer:

A

Explanation:

Watson Assistant helps companies reduce costs per contact by handling common inquiries and providing quick assistance through automation. By using a virtual agent, companies can reduce the number of queries that require human intervention, thereby lowering the

overall cost associated with customer support operations.

How Watson Assistant Reduces Costs: Watson Assistant automates responses to frequently asked questions and common tasks, freeing up human agents to handle more complex inquiries. This leads to lower staffing requirements and operational costs.

Comparison with Other Options:

B (Remove the requirement for all human agents): Watson Assistant reduces but does not eliminate the need for human agents.

C (Provide live agents who can speak multiple languages): Not a function of Watson Assistant; it is a virtual agent technology.

D (Check insurance claim details for non-native English speakers): Watson Assistant can handle multilingual conversations, but the primary benefit is cost reduction.

IBM Watson Assistant Overview

IBM Cloud Architect Exam Study Guide

Question 9

Which IBM Cloud service can be used to store and analyze data collected to identify the root causes of slowness?

Options:

A- Databases for etcd

B- Databases for Redis

C- Databases for Elasticsearch

D- Cloudant

Answer:

C

Explanation:

Elasticsearch is a powerful search and analytics engine commonly used to store, search, and analyze large volumes of data in real time. IBM Cloud's 'Databases for Elasticsearch' service is specifically designed for this purpose, allowing users to ingest, search, and analyze log data or other large datasets to identify root causes of slowness or performance issues.

Why Databases for Elasticsearch? Elasticsearch is ideal for storing and analyzing logs due to its ability to quickly index and query large amounts of data. It provides powerful search capabilities, aggregations, and visualization tools that are well-suited for identifying performance bottlenecks and troubleshooting.

Comparison with Other Options:

A (Databases for etcd): Not used for analytics or identifying performance issues.

B (Databases for Redis): Primarily an in-memory data store, not suitable for complex data analysis.

D (Cloudant): A NoSQL database service optimized for web and mobile applications, but not primarily for analyzing log data.

IBM Cloud Databases for Elasticsearch

IBM Cloud Architect Exam Study Guide

Question 10

Question Type: **MultipleChoice**

Which two search capabilities are available in IBM Cloud Log Analysis?

Options:

A- Simple search

B- SQL search

C- Wildcard search

D- Field search

E- Multi search

Answer:

A, C

Explanation:

IBM Cloud Log Analysis provides several search capabilities to help users analyze log data efficiently:

Simple Search: This allows users to perform straightforward keyword searches across log data. It helps find specific entries or patterns without needing complex syntax or operators.

Wildcard Search: This feature allows users to perform searches using wildcards, enabling more flexible pattern matching. It is particularly useful when searching for entries that match a certain pattern or partial data.

Comparison with Other Options:

B (SQL search): Not a feature of IBM Cloud Log Analysis.

D (Field search): IBM Cloud Log Analysis supports field-specific searches, but this option was not one of the correct ones.

E (Multi search): Not explicitly a feature of IBM Cloud Log Analysis.

IBM Cloud Log Analysis Documentation

IBM Cloud Architect Exam Study Guide