

Oracle

1Z0-1111-25 Exam

Oracle Cloud Infrastructure 2025 Observability Professional

**Questions & Answers
Demo**

Version: 4.0

Question: 1

You are working on a project to automate the deployment of Oracle Cloud Infrastructure (OCI) compute instances that are pre-configured with web services. As part of the deployment workflow, you also need to create a corresponding OCI object storage bucket bearing the same name as that of the compute instance. Which of these two options can help you achieve this requirement? (Choose two.)

- A. Cloud Agent Plugin for the compute instance
- B. Service Connector Hub
- C. Oracle Functions
- D. OCI CLI command, `oci os bucket create auto`
- E. Events Service

Answer: B, C

Explanation:

To automate the creation of an OCI Object Storage bucket with the same name as a compute instance during deployment, you need a mechanism to detect the instance creation event and trigger an action to create the bucket. Two OCI services that can achieve this are Service Connector Hub and Oracle Functions, used in conjunction with the Events Service.

Service Connector Hub (B): This service acts as a cloud message bus that facilitates data movement between OCI services. You can configure a service connector with the Events Service as the source

(to detect compute instance creation events, e.g., `com.oraclecloud.computeapi.launchinstance.end`) and Oracle Functions as the target. The service connector filters and routes the event to trigger a function.

Oracle Functions (C): This is a serverless platform that allows you to write and execute code in response to events. You can create a function that retrieves the compute instance name from the event payload and uses the OCI SDK or API to create an Object Storage bucket with the same name.

Why not A, D, or E alone?

Cloud Agent Plugin (A): This is used for monitoring and managing compute instances but does not directly support bucket creation automation.

OCI CLI command (D): The command `oci os bucket create auto` is not a valid OCI CLI command (`oci os bucket create` is valid but requires manual invocation or scripting, not event-driven automation).

Events Service (E): While critical for detecting instance creation, it alone cannot execute the logic to create a bucket—it needs a target like Functions or Notifications.

This solution leverages the event-driven architecture of OCI, combining Events Service (implicitly used with Service Connector Hub) and Oracle Functions for execution.

Reference: OCI Events Service, Service Connector Hub, Oracle Functions

Question: 2

What happens in Stack Monitoring after Management Agents are set up and resources are discovered?

- A. Metric data is immediately collected
- B. Alarm rules will trigger when resources are down or performance thresholds are crossed
- C. Management Agents discover resources that are running locally on the instance
- D. OCI Notifications send email notifications

Answer: A

Explanation:

In OCI Stack Monitoring, once Management Agents are deployed and resources (e.g., databases, applications) are discovered, the immediate next step is the collection of metric data.

Metric data is immediately collected (A): Management Agents are lightweight processes that continuously collect performance and health metrics from discovered resources (e.g., CPU usage, memory utilization) and send them to OCI services like Monitoring or Stack Monitoring. This data becomes available for visualization and analysis right after discovery.

Why not B, C, or D?

Alarm rules (B): Alarms are configured separately in the OCI Monitoring service and only trigger after metric data is collected and thresholds are breached—not an immediate post-discovery action.

Resource discovery (C): Discovery happens before this stage, as the question assumes resources are already discovered. Agents don't rediscover resources post-setup.

Notifications (D): Notifications require separate configuration (e.g., via the Notifications service) and are not an automatic outcome of agent setup and discovery.

This aligns with Stack Monitoring's purpose of providing real-time visibility into resource performance.

Reference: Stack Monitoring Overview, Management Agent

Question: 3

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

- A. Management Agent Cloud Service
- B. Actions
- C. Rule Conditions
- D. Install Key
- E. Service Connector

Answer: B, C

Explanation:

To create a rule in the OCI Events Service, you need to define what triggers the rule and what happens when it's triggered. The two required components are:

Actions (B): These specify the tasks to perform when an event matches the rule (e.g., invoking a function, sending a notification, or streaming to a service). Without an action, the rule has no effect.

Rule Conditions (C): These define the criteria for matching events (e.g., event type like `com.oraclecloud.computeapi.launchinstance.end` or resource attributes). Conditions filter which events trigger the rule.

Why not A, D, or E?

Management Agent Cloud Service (A): This is unrelated to Events Service rules; it's for monitoring resources.

Install Key (D): This is used for agent installation, not event rules.

Service Connector (E): While it can work with Events Service, it's a separate service and not a required component of an event rule itself.

These two elements form the core of an OCI Events Service rule, enabling event-driven automation.

Reference: OCI Events Service Rules

Question: 4

Choose two FluentD scenarios that apply when using continuous log collection with client-side processing. (Choose two.)

A. Managing apps/services which push logs to Object Storage

B. Comprehensive monitoring for OKE/Kubernetes

C. Monitoring systems that are not currently supported by Management Agent

D. Log Source

Answer: A, B

Explanation:

FluentD is an open-source data collector used for continuous log collection with client-side processing in OCI Logging. Two applicable scenarios are:

Managing apps/services which push logs to Object Storage (A): FluentD can be configured to collect logs from applications or services (e.g., Oracle Functions) that write logs to Object Storage buckets. It processes these logs client-side and forwards them to OCI Logging or Logging Analytics.

Comprehensive monitoring for OKE/Kubernetes (B): FluentD is widely used in Kubernetes environments like Oracle Container Engine for Kubernetes (OKE) to collect logs from pods, containers, and nodes. It processes these logs locally before sending them to OCI services for analysis.

Why not C or D?

Monitoring unsupported systems (C): While possible, this is not a primary FluentD scenario in OCI—it's more about extending Management Agent capabilities.

Log Source (D): This is a component of Logging Analytics, not a FluentD scenario.

FluentD's flexibility makes it ideal for these use cases in OCI's observability ecosystem.

Reference: FluentD with OCI Logging, OKE Logging

Question: 5

Which of the following is not a key interaction element in the Log Explorer UI of Logging Analytics?

- A. Fields Panel
- B. Time Picker
- C. Scope Filter
- D. Dashboard

Answer: D

Explanation:

The Log Explorer UI in OCI Logging Analytics includes four key interaction elements: Fields Panel, Time Picker, Scope Filter, and Results Panel. These allow users to search, filter, and analyze logs interactively.

Dashboard (D): This is not part of the Log Explorer UI. Dashboards are separate visualizations in Logging Analytics for summarizing data, not an interactive element of the Log Explorer.

Why A, B, and C are key elements:

Fields Panel (A): Displays log fields for filtering and analysis.

Time Picker (B): Sets the time range for log queries.

Scope Filter (C): Defines the scope (e.g., compartments, log groups) of the log search.

Reference: Log Explorer UI