

Oracle Cloud Infrastructure 2026 Observability Professional

Oracle 1z0-1111-25

Version Demo

Total Demo Questions: 6

Total Premium Questions: 69

Buy Premium PDF

<https://passqueen.com>

support@passqueen.com

Topic Break Down

Topic	No. of Questions
Topic 1, OCI Observability and Management Services	7
Topic 2, OCI Logging	9
Topic 3, OCI Monitoring	9
Topic 4, OCI Application Performance Monitoring	13
Topic 5, OCI Logging Analytics	14
Topic 6, OCI Stack Monitoring	6
Topic 7, OCI Operations Insights	9
Topic 8, Mix Questions	2
Total	69

QUESTION NO: 1

Which Management Agent group allows the agent to upload data to the discovery service?

- A. AgentUsersGrp
- B. StackMonitoringAdminGrp
- C. StackMonitoringViewerGrp
- D. Mgmt_agent_dynamic_group

ANSWER: D

Explanation:

Mgmt_agent_dynamic_group is correct because a Management Agent authenticates to OCI as a resource principal and must be included in an IAM dynamic group to receive permissions. A dynamic group is defined with matching rules for Management Agent resources, allowing OCI IAM policies to grant those agents access without creating individual user credentials. To upload discovery information, the dynamic group is granted the appropriate Discovery Service permission, typically through a policy such as `allow dynamic-group Mgmt_agent_dynamic_group to use discovery-service-family in tenancy`, with scope adjusted where applicable. The group itself identifies the agent resources; the attached IAM policy authorizes the upload operation. This model supports secure, scalable agent deployment because newly registered agents that satisfy the dynamic-group rule automatically receive the intended authorization. Oracle documents dynamic groups as the mechanism for assigning policies to OCI resource principals, including service-managed resources such as agents. See Oracle's [dynamic group documentation](#) and the [OCI IAM policy reference](#) for the resource-principal and policy syntax model.

QUESTION NO: 2

An organization must retain OCI API activity for longer than the OCI Audit service retention period. The security team wants a managed solution that continuously archives Audit events to an Object Storage bucket without writing custom polling code. Which TWO OCI capabilities should be used? (Choose two.)

- A. OCI Audit
- B. Service Connector Hub
- C. OCI Events Service
- D. Management Agent

ANSWER: A B

Explanation:

OCI Audit is correct because it automatically records supported OCI API calls and provides the activity events that must be retained. Audit records include details such as the principal, action, target resource, time, and response status.

Service Connector Hub is correct because it can create a managed delivery pipeline from Audit events to an Object Storage target. This supports long-term archival beyond the Audit service retention period. OCI Events can react to selected resource state changes, but it is not the service for continuously exporting the complete Audit event stream. A Management Agent collects telemetry from managed hosts and is not needed to archive OCI Audit records.

QUESTION NO: 3

Which of the following TWO are stored in a Log Source of Logging Analytics? (Choose two.)

- A. Which Parsers to use
- B. Where to store Log data

C. Where to find Logs

D. Which Management Agents to use

ANSWER: A C

Explanation:

In Oracle Logging Analytics, a log source defines the structure and collection characteristics for a particular kind of log data. It includes the parser or parsers used to interpret incoming log entries. Parsers identify record boundaries and extract meaningful fields, such as timestamps, severity values, messages, host information, and other attributes, so that the data can be indexed, queried, visualized, and used in alerts.

A log source also stores where to find logs. For file-based collection, this includes the log file location or file-name pattern that enables the Management Agent to locate the applicable log files on the monitored host. This source definition can then be associated with an entity to apply the appropriate collection configuration in the monitored environment. Together, the log location and parser configuration allow Logging Analytics to collect raw log content and transform it into structured, searchable log records consistently across matching entities.

Oracle documents log sources as the definitions that specify the log data to collect and the parser used to process it. See [Oracle Logging Analytics documentation](#) and the [Oracle Management Agents documentation](#) for the relationship between collection configuration and monitored resources.

QUESTION NO: 4

There are several ways to reduce Logging Analytics noise. Select the TWO options that apply. (Choose two.)

A. Use histogram records

B. Use specific keywords

C. Use parsed logs search

D. Use time-picker to limit the volume of logs

ANSWER: C D

Explanation:

Use parsed logs search is correct because Logging Analytics extracts fields from log content during parsing, allowing searches to focus on meaningful attributes such as severity, source, entity, message type, or other defined fields. Filtering on these structured fields narrows broad log results to records relevant to the investigation, improving the signal-to-noise ratio and making patterns or operational issues easier to identify.

Use time-picker to limit the volume of logs is also correct because the time picker constrains a search to the incident window or operational period of interest. Log repositories can contain very large volumes of records, and limiting the search to an appropriate recent or historical interval reduces unrelated events returned by the query. This supports faster investigation and more focused visualizations while preserving the context needed to analyze an event. Oracle documents Logging Analytics search capabilities, including using parsed fields and selecting the relevant search time range, in its [Logging Analytics: Search Logs documentation](#). Additional product concepts and operational guidance are available in the [Oracle Logging Analytics documentation](#).

QUESTION NO: 5

Which is the correct monitoring query that will monitor the CPU utilization threshold including an alarm?

A. `CpuUtilization[1m](shape "VM.Standard.E4.Flex").max()`

B. `(CpuUtilization[1m].max() > 80).grouping().sum()`

C. `CpuUtilization[1m].max()`

D. `cpuutilization[1m].max().grouping().sum()`

ANSWER: B

Explanation:

`(CpuUtilization[1m].max() > 80).grouping().sum()` is the correct query because it evaluates CPU utilization as a threshold-based condition suitable for alarm monitoring. `CpuUtilization[1m]` selects the CPU utilization metric using one-minute intervals, while `.max()` evaluates the maximum observed value for each interval and metric stream. Comparing that result with `> 80` converts the metric evaluation into a condition representing CPU utilization greater than 80 percent. The parentheses ensure that this comparison is evaluated before the subsequent aggregation.

The `.grouping().sum()` portion then aggregates the resulting condition across the applicable metric streams. This is useful when the alarm must account for multiple resources or dimension values rather than only one individually selected stream. The alarm can use the resulting aggregated value together with its configured alarm rule to detect and notify on the defined CPU threshold condition. Oracle Monitoring Query Language supports metric interval selection, statistical functions, comparison operators, and grouping/aggregation operations for metric evaluation. See Oracle's documentation on [querying metrics with MQL](#) and [creating and managing alarms](#).

QUESTION NO: 6

Which of the following capabilities does the performance management feature of Database Management Services offer to a managed database?

- A. Automatically invokes full stats gathering of objects to improve performance of regressed SQLs
- B. Visualizes and performs trend analysis from AWR data to detect issues using AWR Explorer
- C. Dynamically modifies database initialization parameters to improve performance

ANSWER: B

Explanation:

Visualizes and performs trend analysis from AWR data to detect issues using AWR Explorer is correct because Oracle Cloud Infrastructure Database Management provides AWR Explorer as a performance-management capability for managed Oracle databases. Automatic Workload Repository (AWR) data contains historical database workload and performance statistics, including information that helps administrators identify resource-intensive SQL, waits, load patterns, and performance changes over time. AWR Explorer presents this collected historical information in an interactive form, allowing a database administrator to select and compare AWR snapshots across time periods. This makes it possible to investigate trends, recognize regressions, and isolate when a performance problem began or changed.

The feature is designed to support evidence-based diagnosis using the database's recorded workload history rather than requiring administrators to manually assemble performance data from separate sources. Its value is especially clear for intermittent or historical issues, because the relevant workload statistics can be examined for the period in which the issue occurred. Oracle documents AWR Explorer as part of Database Management's tools for analyzing Oracle Database performance and AWR reports. See the [Oracle Database Management documentation for AWR Explorer](#) and the [Database Management overview](#).