

DUMPSBOSS.

**Oracle Fusion AI Agent Studio Foundations
AssociateRel 1**

Oracle 1z0-1145-1

Version Demo

Total Demo Questions: 5

Total Premium Questions: 51

Buy Premium PDF

<https://dumpsboss.co>

support@dumpsboss.co

**support@dumpsboss.co
dumpsboss.co**

QUESTION NO: 1

A compliance agent must include an approved disclaimer in every response involving regulated product guidance, regardless of which compliance topic handles the request.

Where should the administrator define this instruction?

- A. In the system prompt
- B. In an individual topic only
- C. In the External REST tool configuration
- D. In the evaluation test data

ANSWER: A

Explanation:

The **system prompt** is correct because it establishes the agent-wide role, behavioral boundaries, and persistent response instructions. A requirement that applies across all regulated-product interactions should be defined as an overarching instruction rather than repeated only in selected scenarios.

Topics are useful for defining focused business areas and scenario-specific guidance. A tool enables an agent to retrieve information or perform an action, and evaluation test data is used to assess the configured behavior rather than define it.

QUESTION NO: 2

An HR manager wants a solution that allows managers to access employee information and initiate transactions like promotions within a single conversational interface.

When selecting a pre-built agent template for this purpose, which OOTB agent should be used?

- A. Pay Analyst Agent, which provides automated notifications for payroll cycle completions
- B. Personal and Employment Details Agent, which retrieves personal and employment details and can generate deep links for HR transactions
- C. Compensation Advisor, which manages the workflow for approving new benefits enrollments
- D. Leave and Absence Analyst, which creates and sends summary reports of team vacation usage to upper management

ANSWER: B

Explanation:

The **Personal and Employment Details Agent, which retrieves personal and employment details and can generate deep links for HR transactions** is the appropriate pre-built agent template. It is designed for conversational HR self-service scenarios in which a manager needs to obtain employee information, such as personal and employment-related details, without manually navigating through multiple application pages. Crucially, the agent can provide deep links into the relevant Fusion HCM transaction flows. This enables the conversation to move from answering an information request to starting an action, such as a promotion, in the appropriate secured HCM experience.

This matches the requirement for a single conversational entry point while retaining Fusion Applications security, role-based access, and the established transaction processes. The agent therefore serves as a practical manager-facing assistant: it can surface relevant worker context and guide the manager directly to the transaction needed to act on that context. Oracle positions AI Agent Studio as a framework for deploying embedded, business-process-aware agents in Fusion Applications; see [Oracle Fusion AI Agent Studio](#) and Oracle's [AI Agent Studio announcement](#).

QUESTION NO: 3

An HR agent can retrieve worker information and prepare proposed changes to employee records. The organization wants to reduce the information exposed to the agent and ensure that a responsible manager reviews sensitive changes before they are committed.

Which two design decisions should be used? Select two.

- A. Limit the agent to the business object fields required for the task
- B. Send an email notification after every sensitive update
- C. Add a human-in-the-loop approval before committing sensitive changes
- D. Allow the agent to retrieve all worker attributes for completeness

ANSWER: A C

Explanation:

Limiting the available business object fields to those required for the task applies data minimization. It reduces the amount of worker information that can be retrieved into the agent context or included in an output.

Adding a human-in-the-loop approval step ensures that an authorized manager reviews a sensitive proposed change before it is applied. Notifications sent after an update do not prevent an inappropriate update, and broadening access to all worker attributes conflicts with the requirement to minimize exposure.

QUESTION NO: 4

A manager-facing agent must retrieve current details from a Fusion employee record and prepare a proposed update to an approved Fusion business transaction. Because the update can affect employee data, a manager must authorize it before completion.

Which two capabilities should be included in the agent design? Select two.

- A. External REST tool for the Fusion employee record
- B. Business Object tool for the Fusion data and transaction
- C. Human-in-the-loop approval before the update is completed
- D. Email tool as the authorization mechanism

ANSWER: B C

Explanation:

A Business Object tool is appropriate for retrieving and working with business data and transactions in Fusion Applications. It gives the agent a controlled way to use the relevant Fusion business object as part of the workflow.

A human-in-the-loop approval step is appropriate because the proposed change affects sensitive employee information. It preserves an authorized business review before the transaction is completed. An External REST tool is intended for systems outside Fusion Applications, and an Email tool can notify users but does not provide the required Fusion data access or approval control.

QUESTION NO: 5

An employee asks an agent to explain the company's leave policy for a planned absence. The agent must provide grounded guidance from the available policy context, but it is not expected to plan a sequence of actions or execute a multi-step workflow.

Which type of agent is best suited to this requirement?

- A. Reasoning agent
- B. Agent team
- C. Answer Agent
- D. External REST tool

ANSWER: C

Explanation:

Answer Agent is correct because the primary requirement is to provide contextual guidance in response to a user question. This is appropriate when the desired outcome is an accurate, useful answer based on the agent's available business context.

A Reasoning agent is more appropriate when the task requires planning, optimizing, and executing a multi-step workflow. An agent team coordinates multiple related agents, while an External REST tool is an integration capability rather than an agent type.