Oracle AI for Fusion Applications

How do I use AI Agent Studio?

Oracle AI for Fusion Applications
How do I use AI Agent Studio?

G36251-17

Author: Soniya Mathew

Contents

Get Help i

1 Overview 1

Overview of AI Agent Studio 1
Key Capabilities of AI Agent Studio 1
Components of AI Agent Studio 1

2 Access Requirements for AI Agent Studio 5

Access Requirements for AI Agent Studio 5
Provide Access to Configure AI Agents in all Products 7
Provide Access to Configure AI Agents in Oracle Fusion Cloud Human Capital Management 8
Provide Access to Configure AI Agents in Oracle Fusion Cloud Supply Chain & Manufacturing 12
Provide Access to Configure AI Agents in Oracle Fusion Cloud Procurement 17
Provide Access to Configure AI Agents in Oracle Fusion Cloud Financials 22
Provide Access to Configure AI Agents in Oracle Fusion Cloud CX 23
Provide Access to Configure AI Agents in Oracle Permitting and Licensing 23

3 Before You Begin 29

Choose How to Create AI Agents 29
Get Started with AI Agent Studio 30
Add External REST Tool 32
Add MCP Tool 33
Add Your LLM 33

4 Use Cases 35

Use Cases for Agents of Type Supervisor 35

5 Create AI Agents 43

Create AI Agents Using Preconfigured Templates 43
Create Custom AI Agents of Type Supervisor 46

Create Custom AI Agents of Type Workflow 48

6 Monitor and Evaluate AI Agents 63

Monitor and Evaluate AI Agents 63
Monitor Agents 65
Evaluate Agents 65
Compare Evaluation Runs 67

7 Migrate RAG Agents to AI Agent Studio 69

Migrate Document Tools of RAG Agents 69

8 Promote Published Agents 71

Migrate AI Agents from One Instance to Another 71

9 Create Channels in AI Agent Studio 73

Create Microsoft Teams Channel for AI Agent Studio 73
Create Slack Channel for AI Agent Studio 75

10 Allow External Access to Fusion Applications Agents 77

Enable Applications to Access Fusion Applications Agents 77

11 AI Agent Studio FAQs 81

When do I need a Custom AI Agent Subscription for AI Agent Studio? 81
Which large language models (LLMs) are currently used or supported for AI agents? 81
Can I use any agent in my agent team or workflow? 81
When should I use a single-agent flow vs a multiagent process? 81
Can I edit a preconfigured agent team? 82
How do I make agents respond faster? 82
How do I track the usage of AI Agents created in AI Agent Studio? 83
How can I best reduce the risk of my agent showing hallucinations? 83
How do I hide or restrict access to AI Agent Studio for specific users? 83
Can I rename the AI chat label to something other than Ask Oracle, or change the design of the chat panel? 83
What's the difference between editing a topic versus editing a system prompt? 84
What types of files can I upload as part of document tools in AI Agent Studio? 84

Why does the external REST tool only support HTTPS transactions? 84
Why is the Provider list in the LLM tab empty? 84
Why can't I see all agent teams in the Agent Teams tab of AI Agent Studio? 84
Why is my agent not fetching data from the document added to it? 85
Why is the Roles list in the Security tab empty? 85

Get Help

There are a number of ways to learn more about your product and interact with Oracle and other users.

Get Help in the Applications

Some application pages have help icons to give you access to contextual help. If you don't see any help icons on your page, click your user image or name in the global header and select Show Help Icons. If the page has contextual help, help icons will appear.

Get Training

Increase your knowledge of Oracle Cloud by taking courses at *Oracle University*.

Join Our Community

Use *Cloud Customer Connect* to get information from industry experts at Oracle and in the partner community. You can join forums to connect with other customers, post questions, suggest *ideas* for product enhancements, and watch events.

Share Your Feedback

We welcome your feedback about Oracle Applications user assistance. If you need clarification, find an error, or just want to tell us what you found helpful, we'd like to hear from you.

You can email your feedback to *oracle_fusion_applications_help_ww_grp@oracle.com*.

Thanks for helping us improve our user assistance!

1 Overview

Overview of AI Agent Studio

AI Agent Studio is a design-time environment that empowers you to create, configure, validate, and deploy AI agents to meet your organization's needs.

With AI Agent Studio, you can easily extend preconfigured agent templates, and even build new agents and multiagent flows. AI Agent Studio is fully integrated into Oracle Fusion Cloud Applications, providing secure and seamless access to the knowledge stores, tools, and APIs of Fusion Applications. This integration enables agents to be deployed directly into the flow, ensuring an efficient process.

Key Capabilities of AI Agent Studio

Feature	Description
Agent template libraries	Use templates and natural language prompts to create or fine tune agents for common business scenarios, such as opportunity-to-quote processing and shift scheduling.
Agent team orchestration	Configure multiple agents to collaborate on multistep processes, integrating user approvals where necessary.
Agent extensibility	Change and extend existing agents in Oracle Fusion Cloud Applications by incorporating new data sources, prompts, and APIs to fit specific business needs or industry requirements.
Native integration with Fusion Applications	Directly access APIs and tools in Fusion Applications, ensuring seamless deployment of agents without complex modifications.
Third-party system integration	Connect with external systems and collaborate with third-party agents for end-to-end automation with secure API support.
Trust and security framework	Automatically applies the security configurations, policies, and access controls of Fusion Applications, ensuring compliance with enterprise security standards.
Validation and testing tools	Use built-in tools to make your agents reliable, repeatable, easy to explain, and secure by verifying AI-driven flows before deployment.

Components of AI Agent Studio

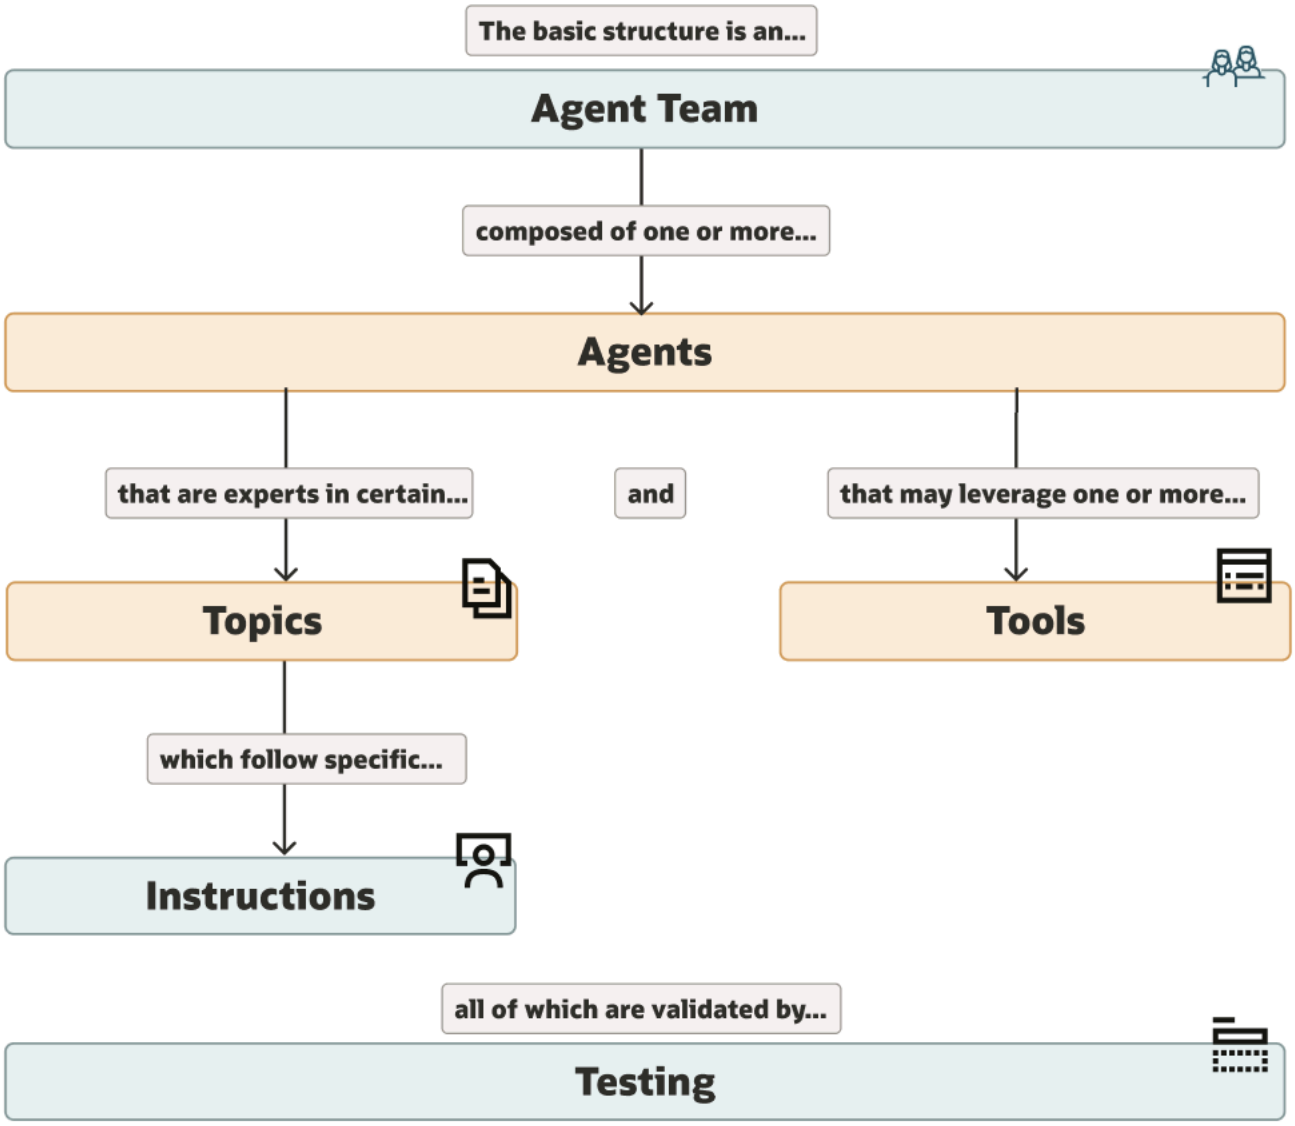

Component	What It Does	Example
Agent Team	Consists of a structured sequence of steps or actions that a single agent or a group of agents follows to accomplish a specific business task or answer a user query. The agent team is the component that can be deployed for use. • Consists of conversation logic, system integration, and user support flow. • Defines how the agent acts for a particular use case — what to do, when to do it, and how to respond based on user inputs or back-end system data.	A recruitment agent (that might consist of multiple agents) schedules interviews, screens resumes, calculates salary, and generates offers, based on policies and approvals.

Component	What It Does	Example
Agents	Leverages a large language model to reason, create action plans, and interact with users to gather information and take direction. On behalf of users, the AI agent can do tasks that enhance productivity, efficiency, and the overall user experience. An agent must be added to an agent team, so that it can be deployed for use. AI agents can be categorized into various types. • User-proxy agent: Acts on behalf of a business user to provide input to another agent or group of agents. It's sometimes referred to as a conversational agent. • Supervisor agent: Orchestrates the use of agents within an agentic flow. • Specialist or utility agent: Focuses on a specific role or expertise and can be skilled in using a particular tool. AI agents can also have one or more of these characteristics: • Persona-based agent: Represents a specific role, such as benefits administrator, customer service representative, and finance administrator. • Tool user: Uses technology-related tools, such as calculators, web search queries, document embedding, and calendar schedulers. • Task-oriented agent: Understands their assignment or task, as a single agent or as part of a multiagent flow.	A scheduling or calendar agent that manages your workday by following your instructions. It can accept new calendar invites and propose alternative times when needed.
Topics	Defines the areas of expertise through instructions that set the boundaries and constraints for agent conversations and abilities.	An employee benefits agent can contain topics such as Health Savings Account (HSA), retirement benefits, and stock plans.
Tools	Defines the additional utilities an agent can use to accomplish a task. One or more tools are assigned to agents, and they're reusable among agents.	• Calculator tool • Email tool • Business object tool • User query tool • Document retrieval tool for retrieval-augmented generation (RAG)
Instructions	Specifies the natural language rules that define the rules or conditions applied to a given topic. Instructions are part of the prompts that are sent to the underlying large language model. They can also contain guidelines and guardrails that set the parameters of an agent response.	Instructions for the payroll deduction topic: Make sure you've information regarding the number of dependents either by asking the user or querying the system. If you don't know the answer, don't make up a response.
Testing	Enables administrators to test the agent team design, ensure correct tool, topic, and instruction configurations, and validate reasoning and sources cited by the agent.	Provide example responses to a series of questions a user would likely ask an agent and details about how the agent arrived at its response.

2 Access Requirements for AI Agent Studio

Access Requirements for AI Agent Studio

You can give access to AI Agent Studio by assigning predefined duty roles to job roles. Also, make sure to complete these prerequisites:

- Enable security console to work with permission groups
- Run scheduled processes to import security data
- Run scheduled processes to get help from the integrated AI agent
- (Optional) Assign privilege to use external REST API tools
- (Optional) Assign permission groups for channels

Watch video

Enable Security Console to Work with Permission Groups

For the Security Console to work with permission groups and related objects, set the **Enable Security Console External Application Integration** (ORA_ASE_SAS_INTEGRATION_ENABLED) profile option at the site level.

1. In the Setup and Maintenance work area, search for the **Manage Administrator Profile Values** task using the search link in the panel.
2. Search for the profile option and set the value for the **Site** profile level to **Yes**.

Run Scheduled Processes to Import Security Data

To import resources from LDAP, and transfer the necessary information into the security tables of Fusion Applications, run these two scheduled processes sequentially.

1. Import Resource Application Security Data
2. Import User and Role Application Security Data

You must run the processes one after the other.

1. Go to **Navigator** > **Tools** > **Scheduled Processes**.
2. Click **Schedule New Process**.
3. Leave the type as **Job**.
4. Search for and select the process.
5. Submit the process.

Run Scheduled Processes to Get Help from the Integrated AI Agent

You can get answers to questions about developing AI agents from the AI help agent integrated into AI Agent Studio. Using this conversational agent, you can get answers to questions about existing agents, search for agents, tools, and topics using natural language, and receive AI powered suggestions for relevant resources to use in your agents.

To use the integrated AI help agent, run these two scheduled processes daily.

- Index AI Agent Studio Assistant Documents
- Index AI Agent Studio Assistant Objects and Attributes

Run each process using these steps:

1. Go to Scheduled Processes, and click **Schedule New Process**.
2. Leave the type as **Job**.
3. Search for and select the process.
4. Submit the process.

Assign Privilege to Use External REST API Tools

To create External REST API tools in AI Agent Studio, the Create and Edit Backends for Visual Builder Studio (ORA_FND_TRAP_PRIV) privilege must be added to the custom role assigned to the user. You can add this privilege while creating or editing a custom role.

1. Go to the Security Console.
2. To use a new custom role, create it. To use an existing custom role, search for the custom role and edit it.
3. Go to the Function Security Policies page and select **Add Function Security Policy**.
4. Add the Create and Edit Backends for Visual Builder Studio (ORA_FND_TRAP_PRIV) privilege to the role and save it.
5. Save the custom role and assign to the user.

Assign Permission Groups for Channels

To create channels from Credentials tab in AI Agent Studio, additional permission groups must be added to the custom role assigned to the user. You can add these permission groups to a duty role and assign the duty role while creating or editing a custom role.

1. Go to the Security Console and create a new duty role for permission groups.
2. Open the **Permission Groups** page and select **Add Permission Groups**.
3. Search for and add these permission groups:
    - create:ChannelManifest
    - create:ExternalChatCorrelation
    - delete:ChannelManifest
    - delete:ExternalChatCorrelation
    - read:ChannelManifest
    - read:ExternalChatCorrelation
    - update:ChannelManifest
    - update:ExternalChatCorrelation
4. Add security view for each permission group.
    a. Select the permission group added.
    b. In Details section, open Security Views tab.
    c. Select **Add Security Views** and add the **AllRowsAllFields** security view.
    d. Add the security view for all the permission groups.
5. Save the duty role and assign this duty role to the custom job role.

6. To use a new custom job role, create it. To use an existing custom job role, search for the custom job role and edit it.

   Note: Make sure to enable permission groups.

7. Go to the **Role Hierarchy** page. From the Roles and Permission Groups tab select **Add Role**.
8. Search for and add the duty role you've created.
9. Save the custom role and assign to the user.

Assign Predefined Duty Roles to Job Roles

Assign predefined product-specific duty roles to the appropriate job roles, and make sure permission groups are enabled. You can give people access to configure AI agents in all or specific products.

CAUTION: Using predefined roles might account for subscription consumption irrespective of whether you purchased the cloud service or not. See *Guidance for Assigning Predefined Roles*.

- *Provide Access to Configure AI Agents in all Products*
- *Provide Access to Configure AI Agents in Oracle Fusion Cloud Human Capital Management*
- *Provide Access to Configure AI Agents in Oracle Fusion Cloud Supply Chain & Manufacturing*
- *Provide Access to Configure AI Agents in Oracle Fusion Cloud Procurement*
- *Provide Access to Configure AI Agents in Oracle Permitting and Licensing*

Provide Access to Configure AI Agents in all Products

1. Go to the Security Console and create a new custom job role.

   Note: Make sure to enable permission groups.

2. On the Role Hierarchy page, open the Roles and Permission Groups tab and add these duty roles:
   - Fai Genai Agent CX Administrator Duty (ORA_DR_FAI_GENERATIVE_AI_AGENT_CX_ADMINISTRATOR_DUTY)
   - Fai Genai Agent FIN Administrator Duty (ORA_DR_FAI_GENERATIVE_AI_AGENT_FIN_ADMINISTRATOR_DUTY)
   - Fai Genai Agent GRC Administrator Duty (ORA_DR_FAI_GENERATIVE_AI_AGENT_GRC_ADMINISTRATOR_DUTY)
   - Fai Genai Agent HCM Administrator Duty (ORA_DR_FAI_GENERATIVE_AI_AGENT_HCM_ADMINISTRATOR_DUTY)
   - Fai Genai Agent PRC Administrator Duty (ORA_DR_FAI_GENERATIVE_AI_AGENT_PRC_ADMINISTRATOR_DUTY)
   - Fai Genai Agent PRJ Administrator Duty (ORA_DR_FAI_GENERATIVE_AI_AGENT_PRJ_ADMINISTRATOR_DUTY)
   - Fai Genai Agent PSC Administrator Duty (ORA_DR_FAI_GENERATIVE_AI_AGENT_PSC_ADMINISTRATOR_DUTY)

- Fai Genai Agent SCM Administrator Duty (ORA_DR_FAI_GENERATIVE_AI_AGENT_SCM_ADMINISTRATOR_DUTY)

3. Open the Roles and Privileges tab and add the Manage All Intelligent Agents (ORA_FAI_MANAGE_ALL_AI_AGENTS) role.
4. Save the custom role and assign it to users who want access.

Provide Access to Configure AI Agents in Oracle Fusion Cloud Human Capital Management

To give access to users without the Human Capital Management Application Administrator Job Role:

1. Go to the Security Console and create a new custom job role.

   Note: Make sure to enable permission groups.

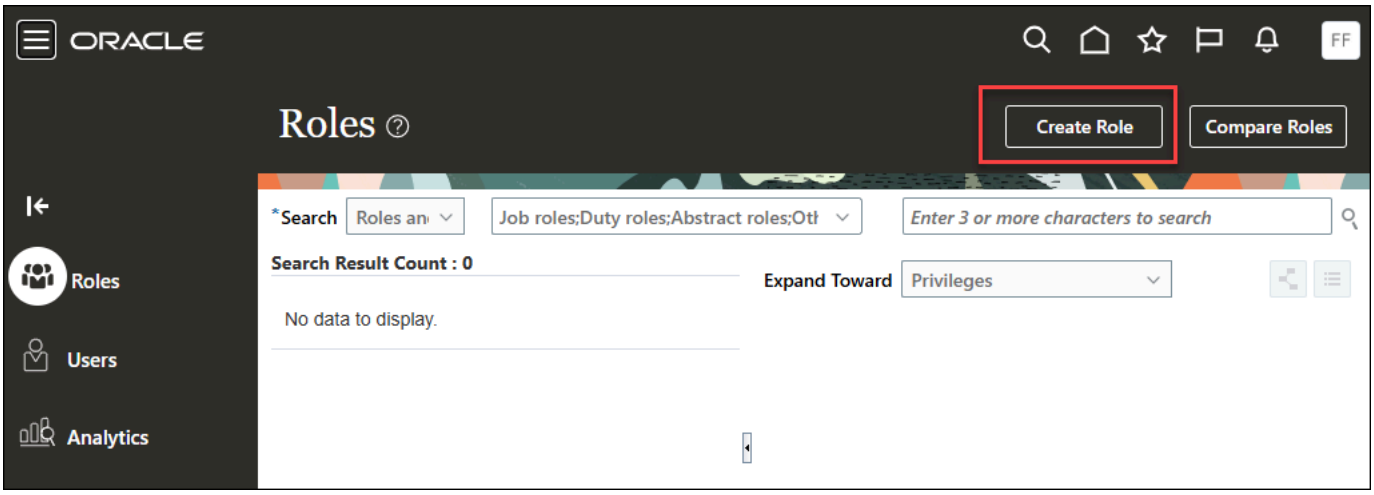

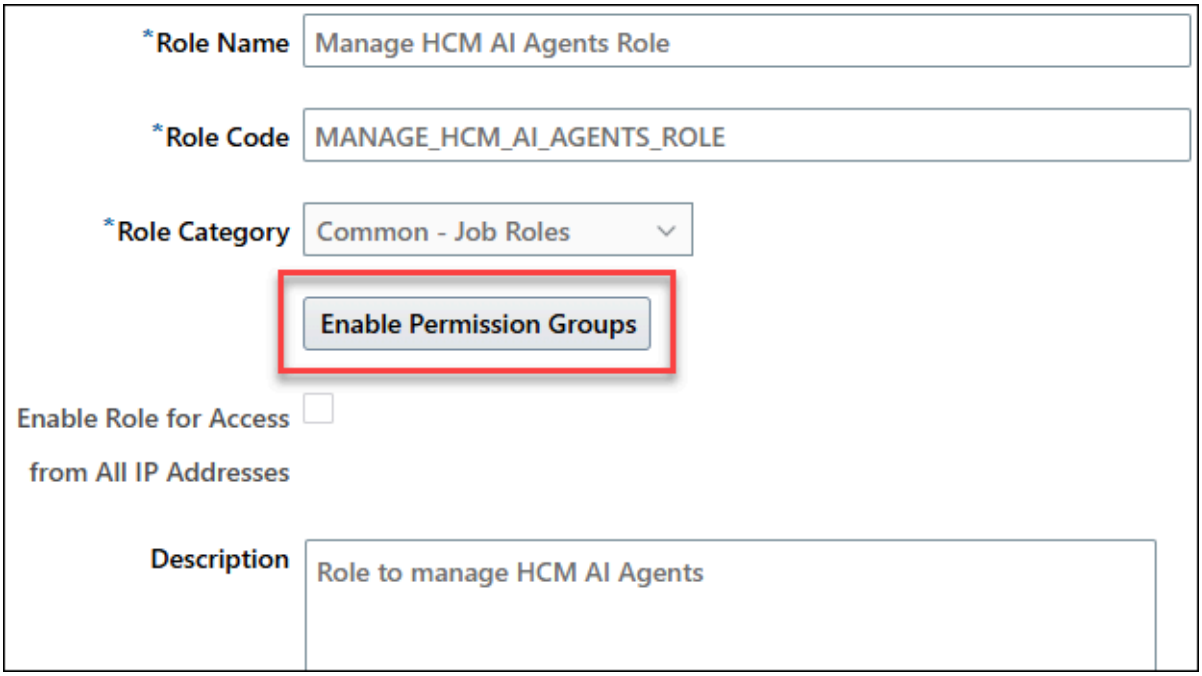

2. Go to the Role Hierarchy page.
    - Open the Roles and Privileges tab, and add the Manage HCM Intelligent Agent (ORA_HRC_HCM_AI_AGENT_MANAGEMENT_DUTY) duty role.

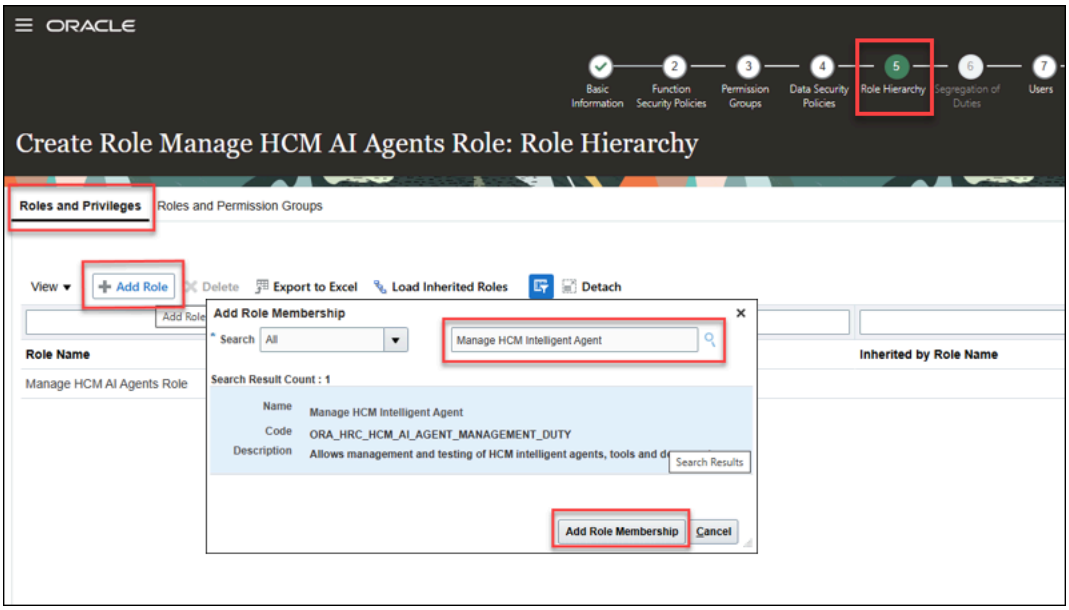

    - Open the Roles and Permission Groups tab, and add the Fai Genai Agent HCM Administrator Duty (ORA_DR_FAI_GENERATIVE_AI_AGENT_HCM_ADMINISTRATOR_DUTY) role.

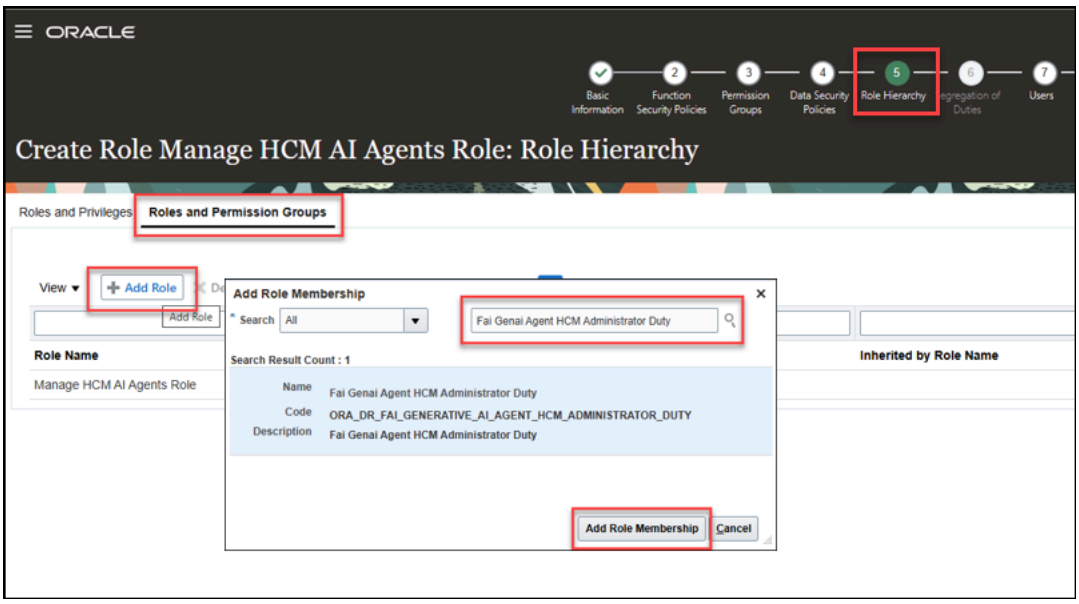

3. Save the custom role and assign to the appropriate job roles.

To give access to users with the Human Capital Management Application Administrator Job Role:

1. Go to the Security Console and create a new custom job role.

   Note: Make sure to enable permission groups.

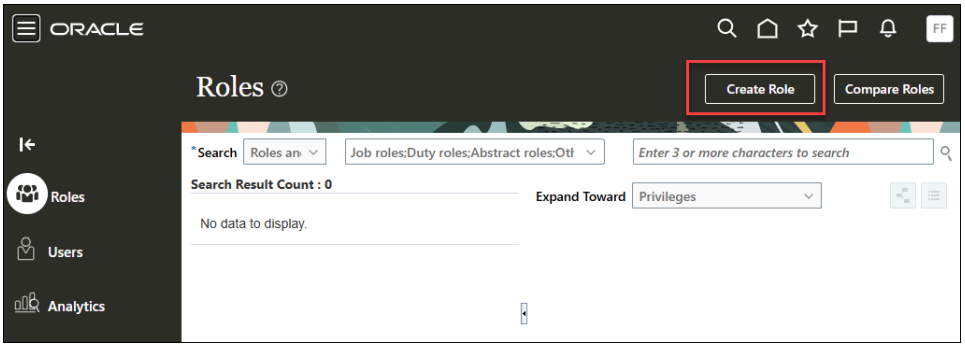

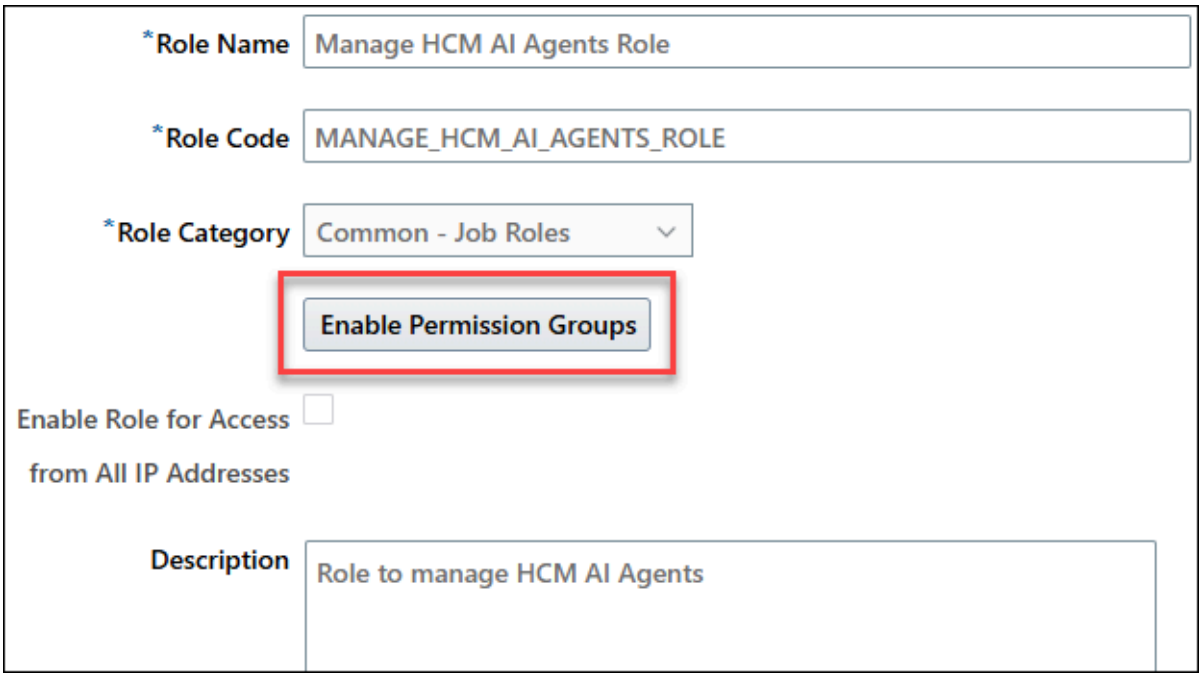

2. On the Role Hierarchy page, open the Roles and Permission Groups tab, and add the Fai Genai Agent HCM Administrator Duty (ORA_DR_FAI_GENERATIVE_AI_AGENT_HCM_ADMINISTRATOR_DUTY) role.

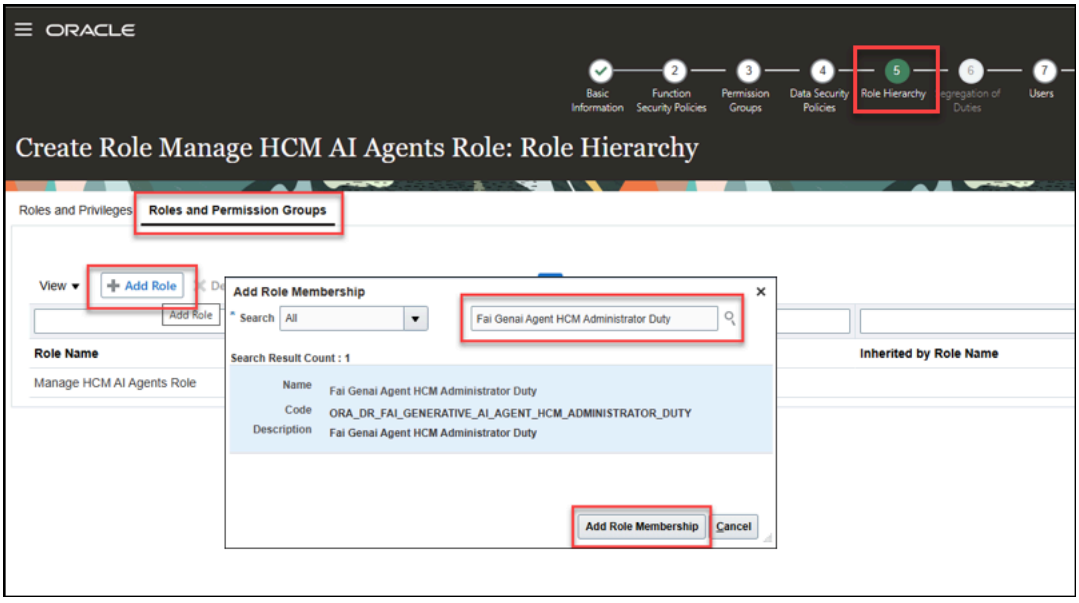

3. Save the custom role and assign to the appropriate job roles.

Provide Access to Configure AI Agents in Oracle Fusion Cloud Supply Chain & Manufacturing

To give access to users without the Supply Chain Application Administrator Job Role:

1. Go to the Security Console and create a new custom job role.

> **Note:** Make sure to enable permission groups.

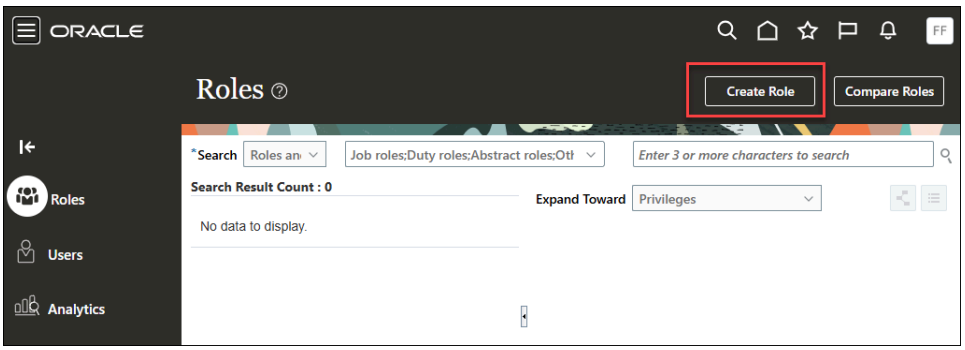

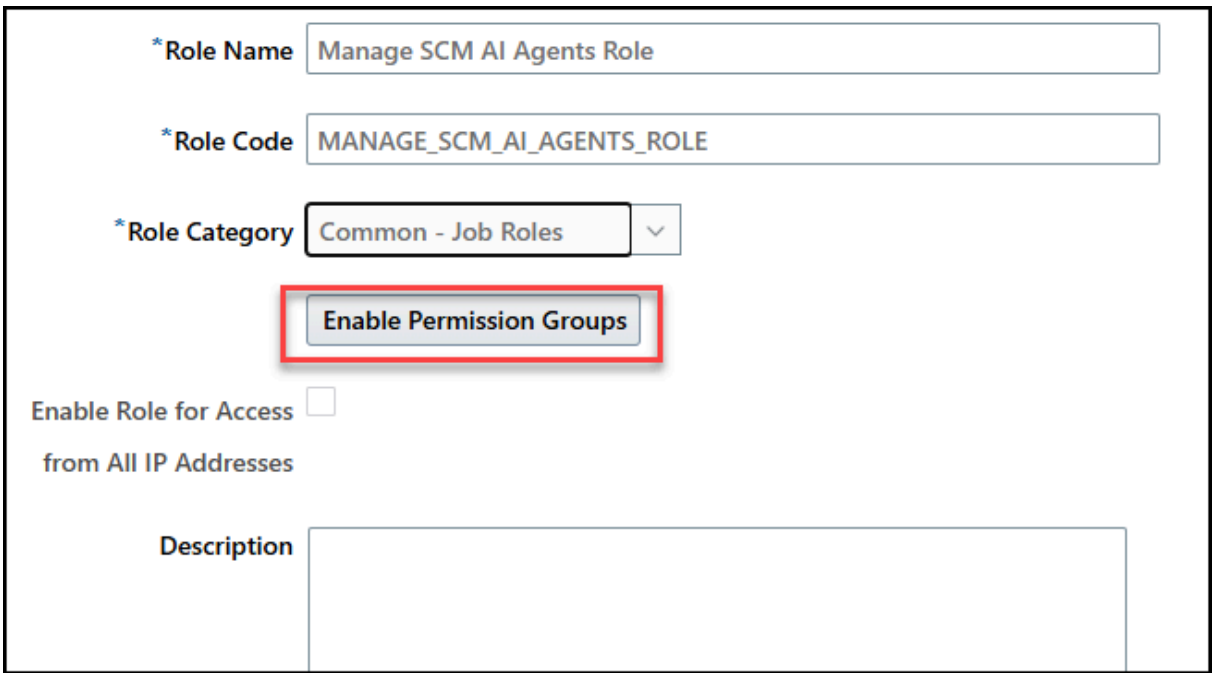
*Role Name
Manage SCM AI Agents Role
*Role Code
MANAGE_SCM_AI_AGENTS_ROLE
*Role Category
Common - Job Roles
Enable Permission Groups
Enable Role for Access from All IP Addresses
Description

2. On the Role Hierarchy page, open the Roles and Privileges tab and add these roles:
   - SCM Intelligent Agent Management Duty (ORA_RCS_SCM_AI_AGENT_MANAGEMENT_DUTY)

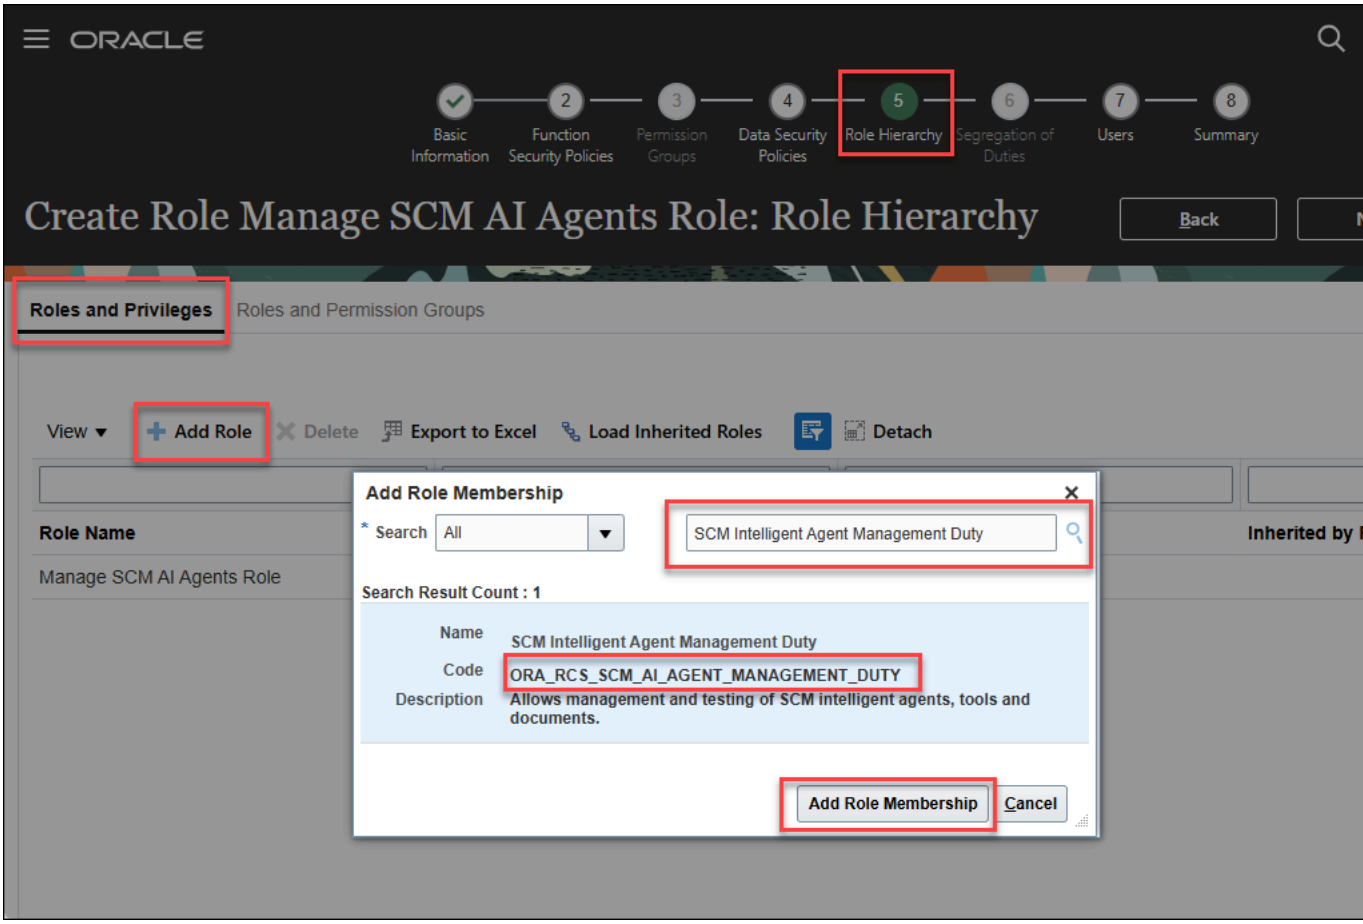

   - SCM Intelligent Agent Management Duty (ORA_RCS_SCM_AI_AGENT_MANAGEMENT_DUTY_HCM)

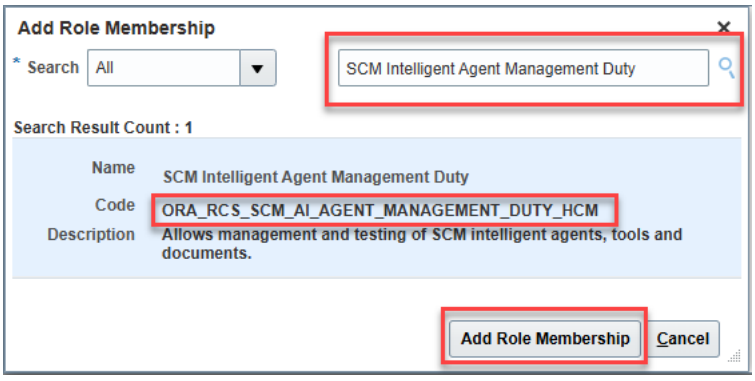

3. Open the Roles and Permission Groups tab and add the Fai Genai Agent SCM Administrator Duty (ORA_DR_FAI_GENERATIVE_AI_AGENT_SCM_ADMINISTRATOR_DUTY) duty role.

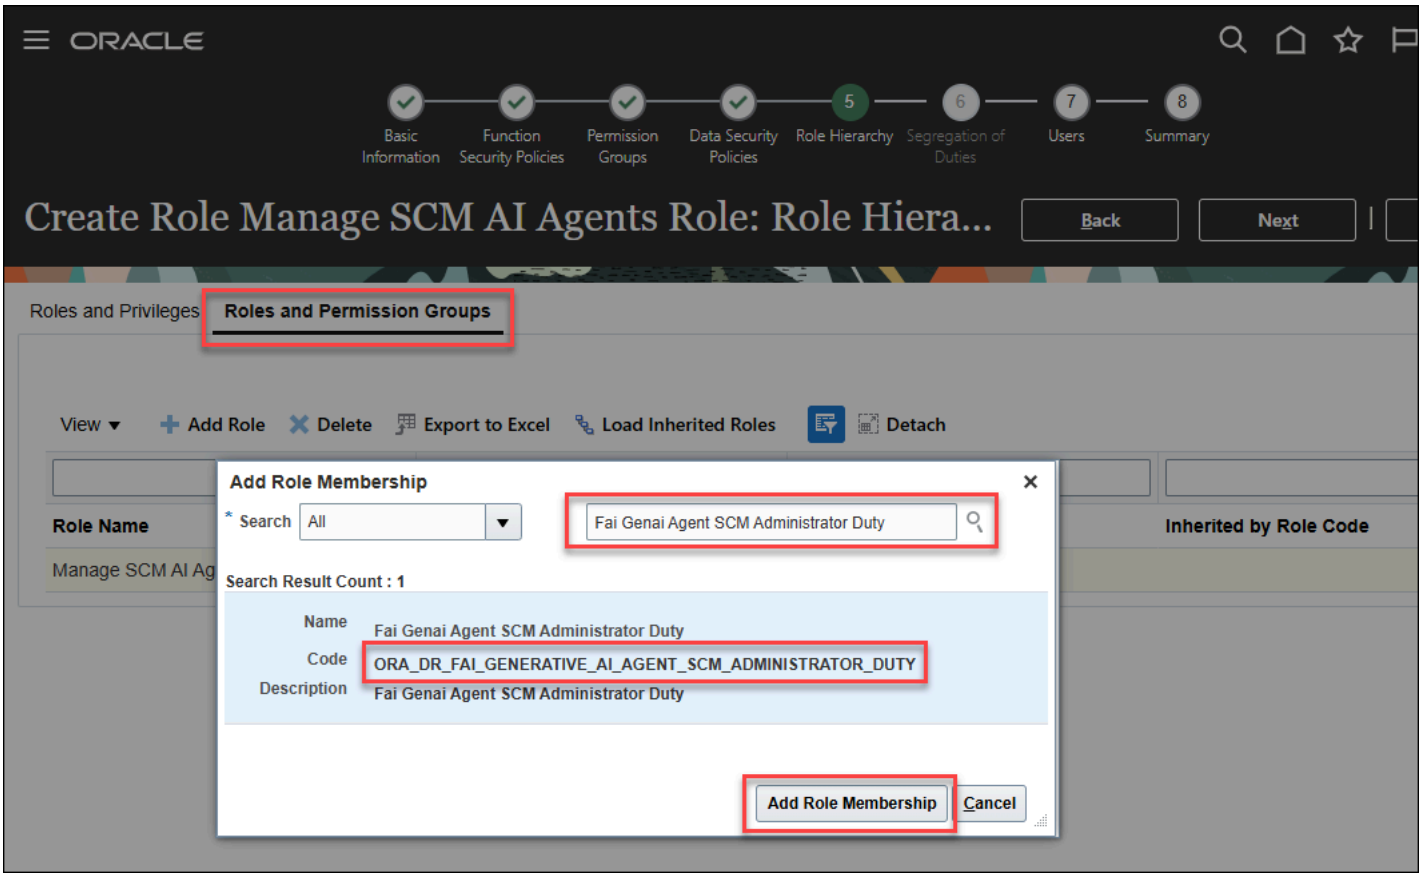

4. Save the custom role and assign this custom role to the appropriate job roles.

To give access to users with the Supply Chain Application Administrator Job Role:

1. Go to the Security Console.

2. Search for the Supply Chain Application Administrator (ORA_RCS_SUPPLY_CHAIN_APPLICATION_ADMINISTRATOR_JOB) job role, and edit it.

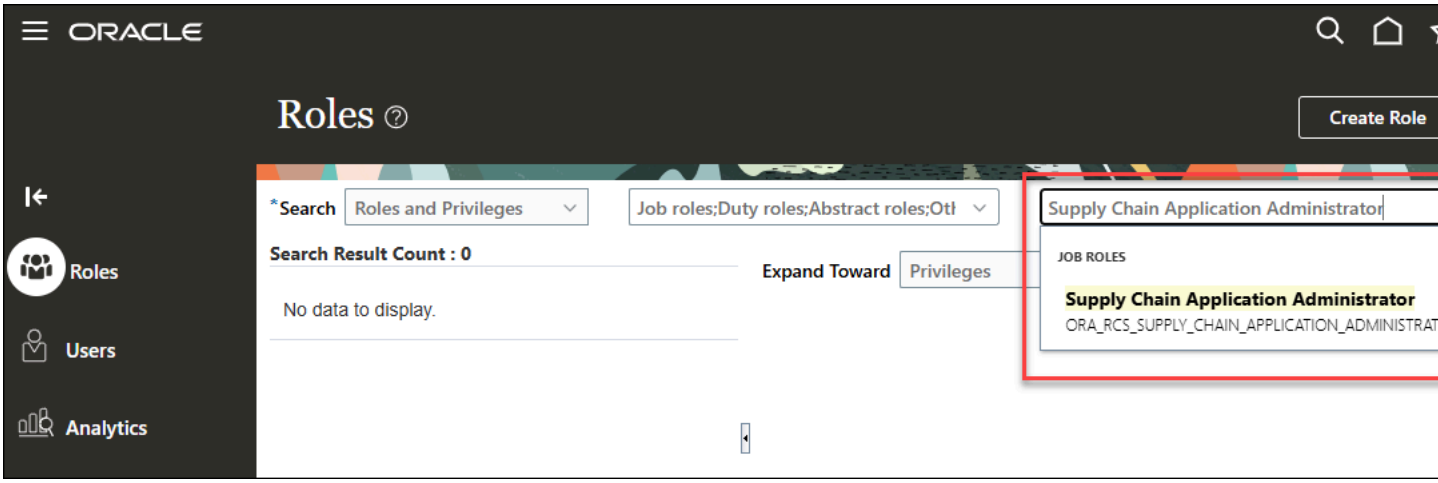

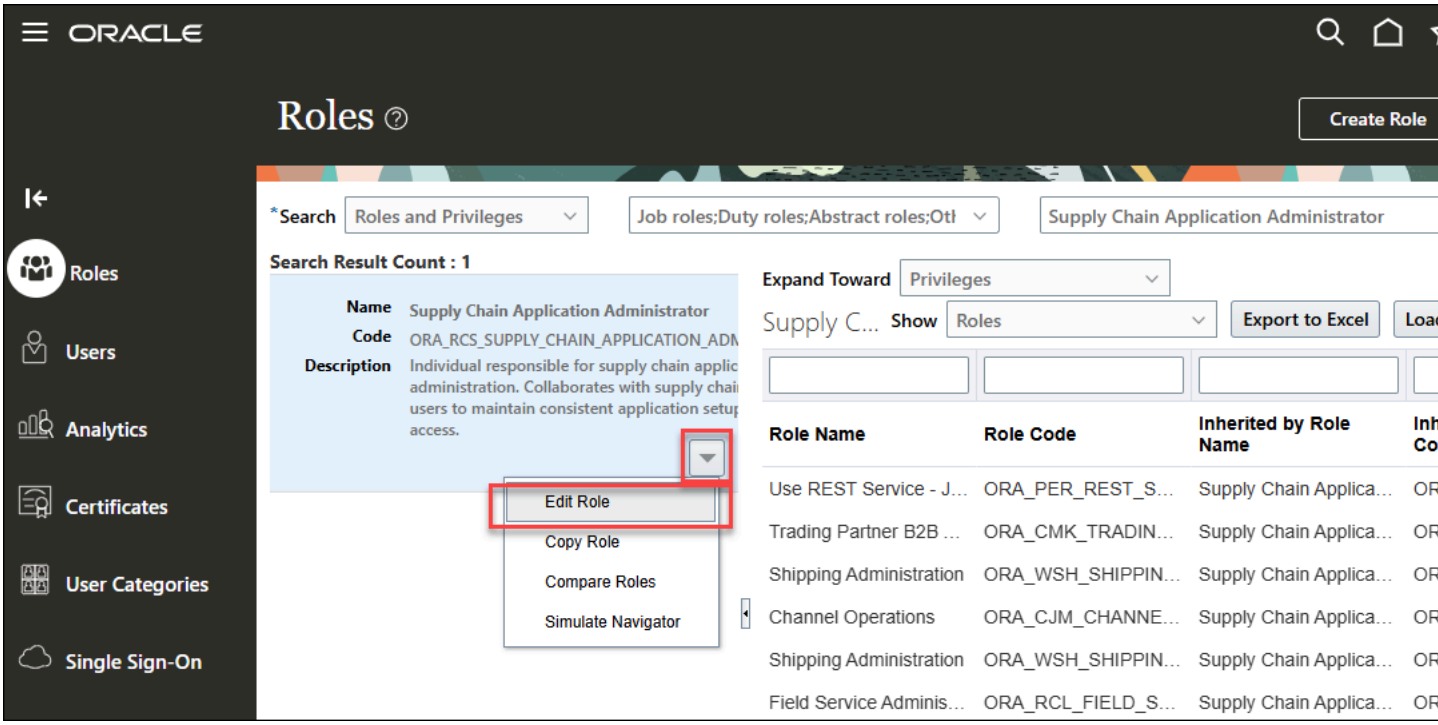

3. Enable permission groups and save the job role.

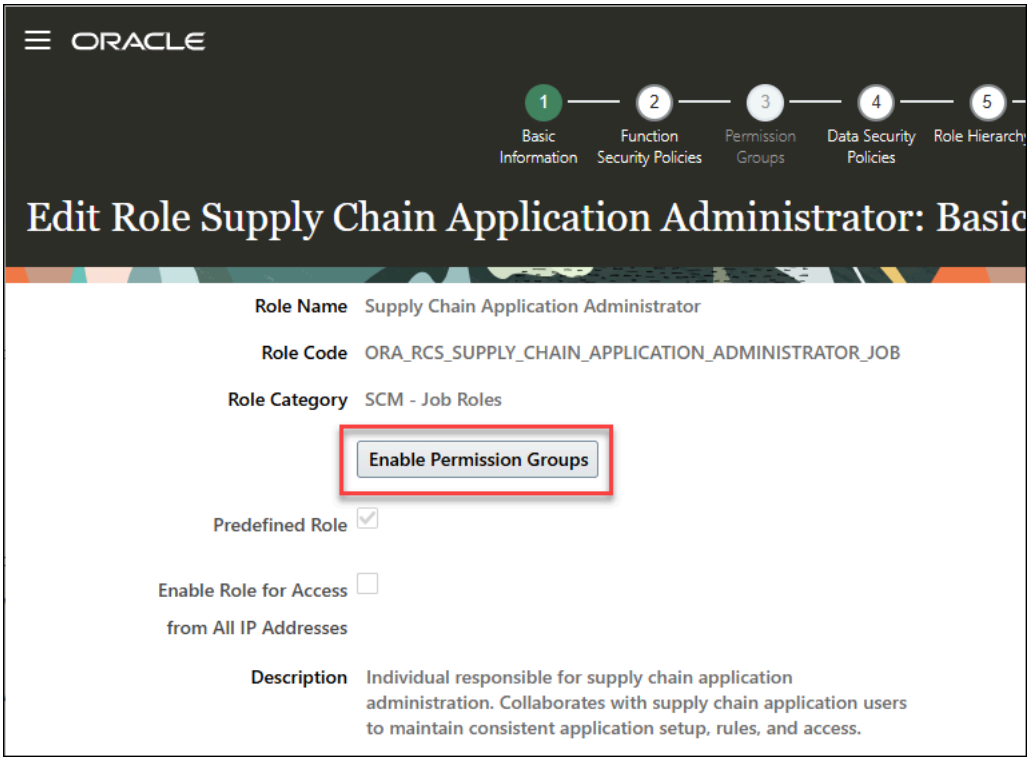

Provide Access to Configure AI Agents in Oracle Fusion Cloud Procurement

To give access to users without the Procurement Application Administrator Job Role:

1. Go to the Security Console and create a new custom job role.

   Note: Make sure to enable permission groups.

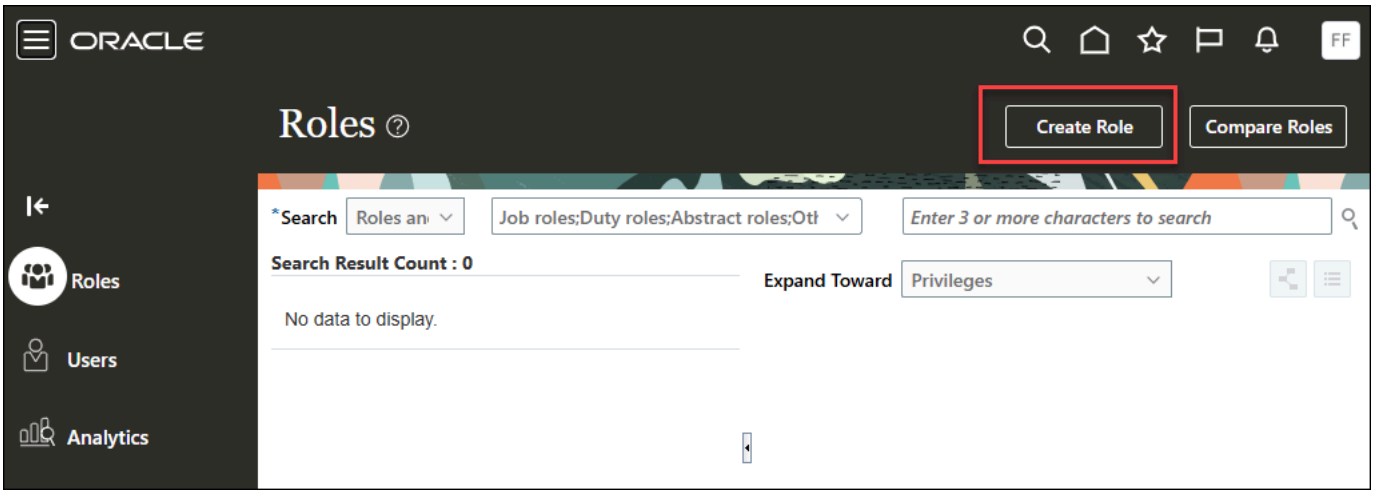

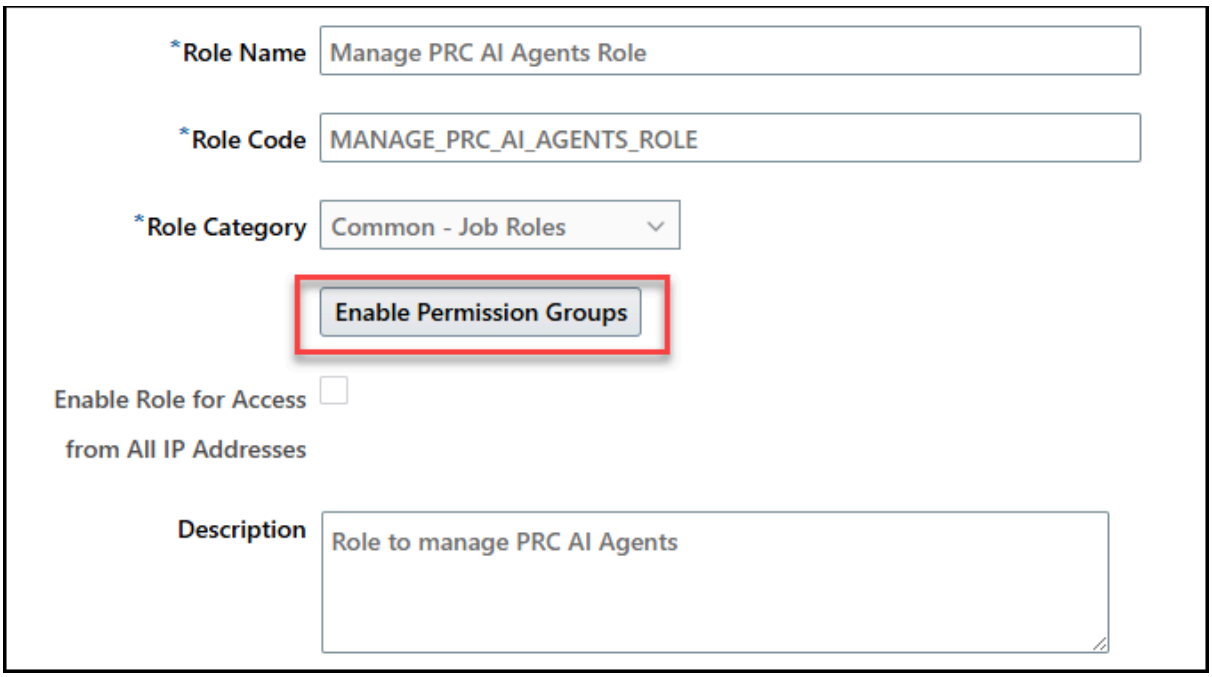

2. On the Role Hierarchy page, open the Roles and Privileges tab and add these roles:
   - PRC Intelligent Agent Management Duty (ORA_PO_PRC_AI_AGENT_MANAGEMENT_DUTY)

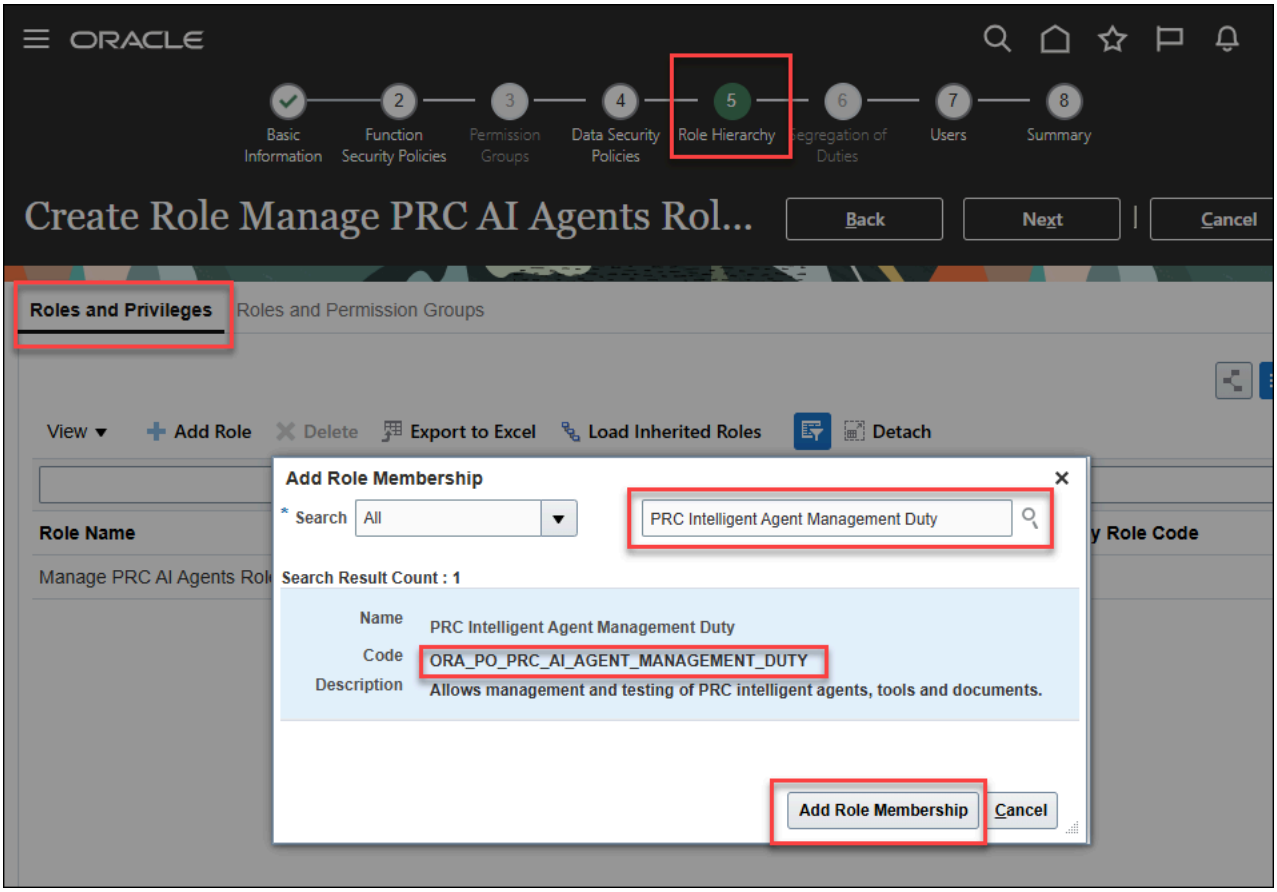

   - PRC Intelligent Agent Management Duty (ORA_PO_PRC_AI_AGENT_MANAGEMENT_DUTY_HCM)

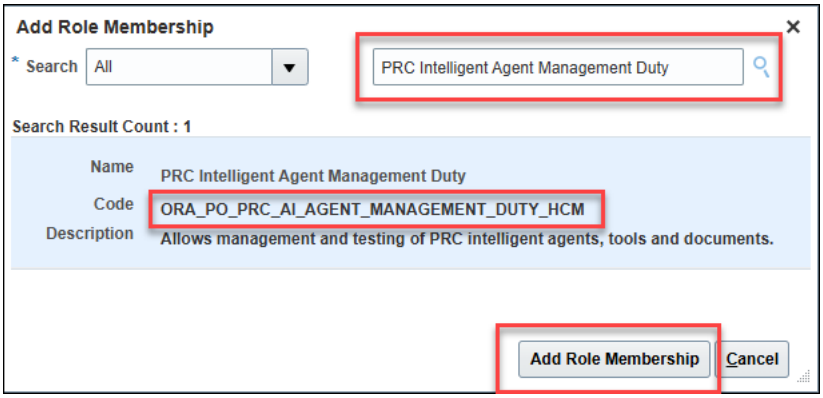

3. Open the Roles and Permission Groups tab and add the Fai Genai Agent PRC Administrator Duty (ORA_DR_FAI_GENERATIVE_AI_AGENT_PRC_ADMINISTRATOR_DUTY) duty role.

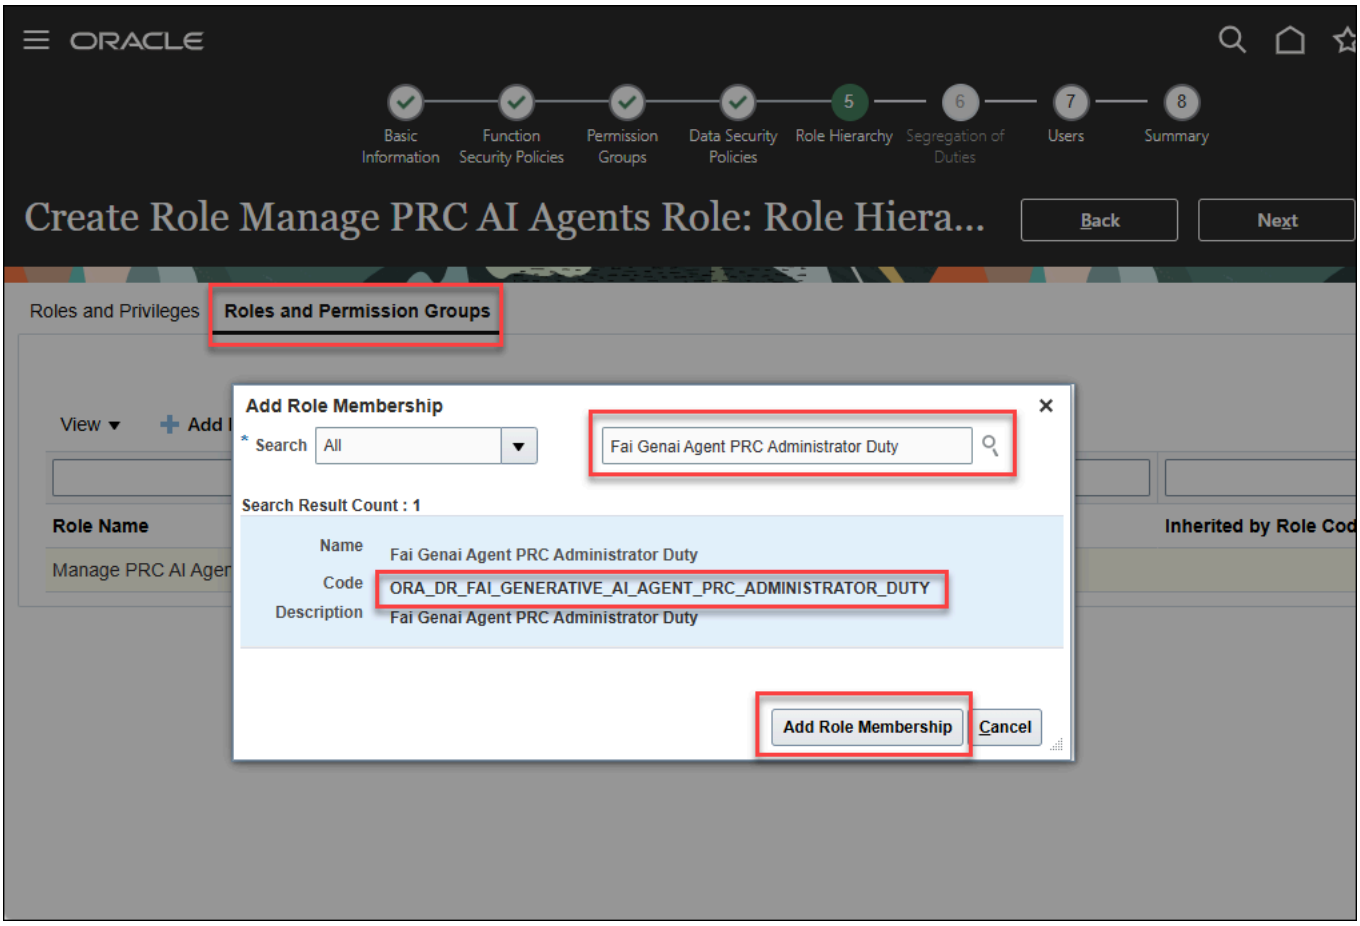

4. Save the custom role and assign this custom role to the appropriate job roles.

To give access to users with the Procurement Application Administrator Job Role:

1. Go to the Security Console.

2. Search for the Procurement Application Administrator (ORA_PO_PROCUREMENT_APPLICATION_ADMIN_JOB) job role, and edit it.

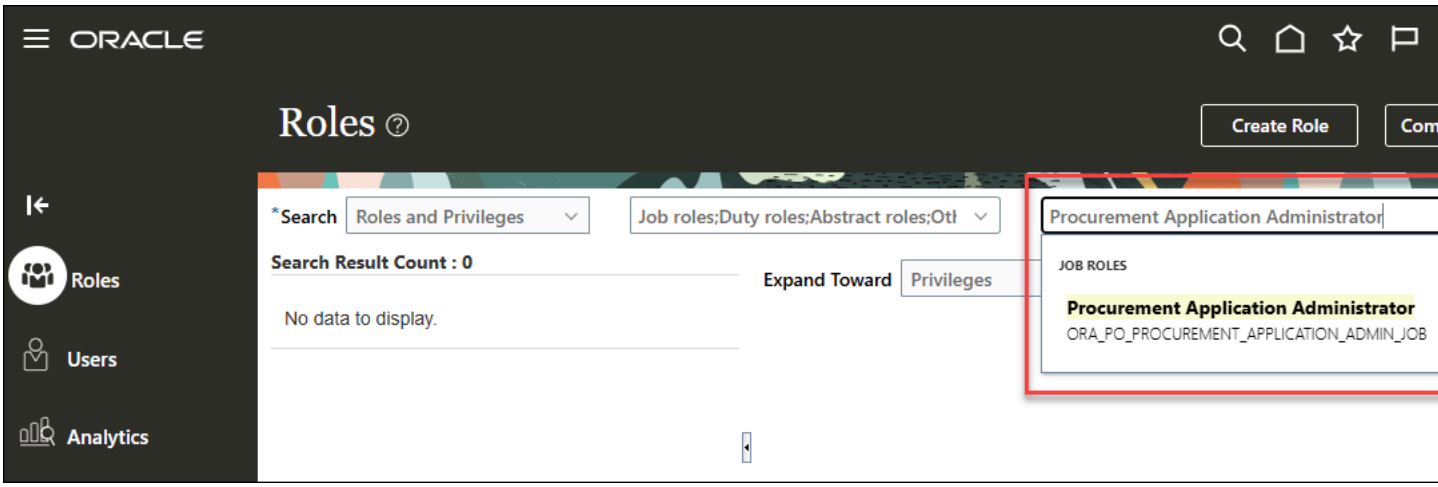

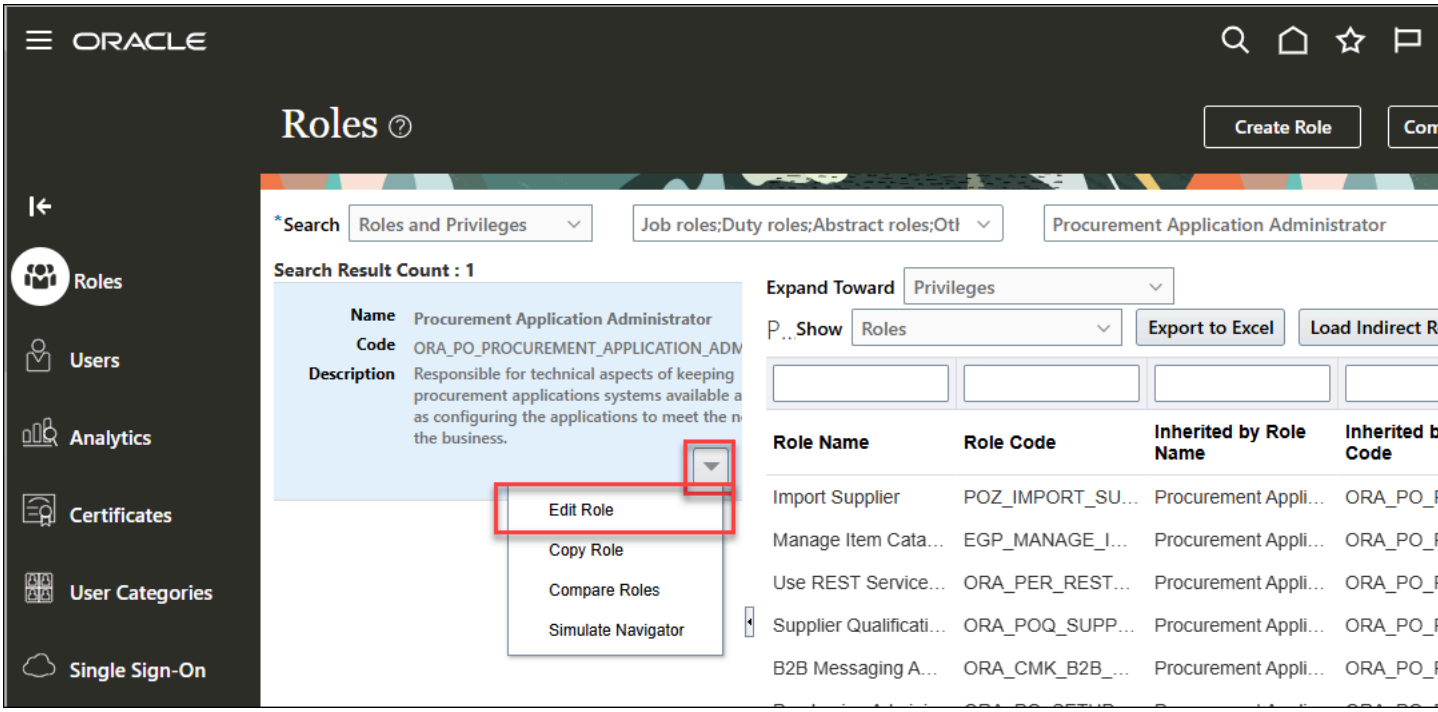

3. Enable permission groups and save the job role.

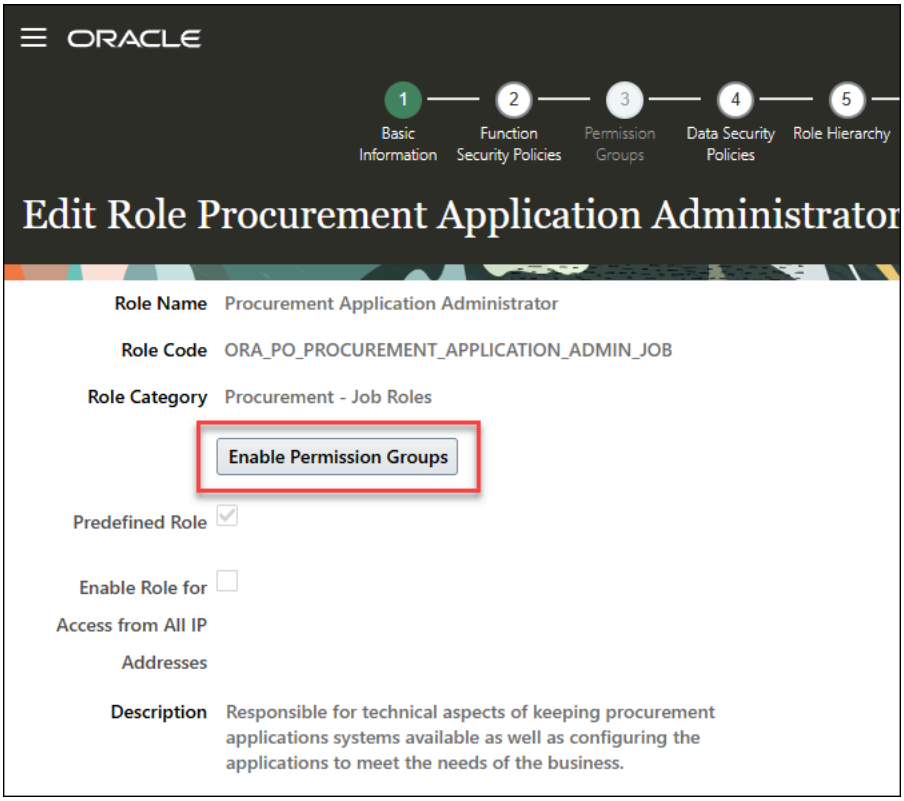

Provide Access to Configure AI Agents in Oracle Fusion Cloud Financials

1. Go to the Security Console and create a new custom job role.

   Note: Make sure to enable permission groups.

2. Go to the Role Hierarchy page.
   - Open the **Roles and Privileges** tab, and add the Manage Financials Intelligent Agent (ORA_FUN_MANAGE_FIN_AI_AGENT_HCM) and Manage Financials Intelligent Agent (ORA_FUN_MANAGE_FIN_AI_AGENT) duty roles.
   - Open the **Roles and Permission Groups** tab, and add the Fai Genai Agent FIN Administrator Duty (ORA_DR_FAI_GENERATIVE_AI_AGENT_FIN_ADMINISTRATOR_DUTY) duty role.
3. Save the custom role and assign it to the appropriate job roles.

Provide Access to Configure AI Agents in Oracle Fusion Cloud CX

To give access to users without the Sales Administrator Job Role:

1. Go to the Security Console and create a new custom job role.

   > **Note:** Make sure to enable permission groups.

2. On the Role Hierarchy page, open the Roles and Privileges tab and add these roles:
   - Manage CX AI Agents (ORA_ZCA_MANAGE_CX_AI_AGENTS)
   - Manage CX AI Agents (ORA_ZCA_MANAGE_CX_AI_AGENTS_HCM)
3. Open the Roles and Permission Groups tab and add the Fai Genai Agent CX Administrator Duty (ORA_DR_FAI_GENERATIVE_AI_AGENT_CX_ADMINISTRATOR_DUTY) duty role.
4. Save the custom role and assign this custom role to the appropriate job roles.

To give access to users with the Sales Administrator Job Role:

1. Go to the Security Console and search for the Sales Administrator (ORA_ZBS_SALES_ADMINISTRATOR_JOB) job role, and edit it.
2. Enable permission groups and save the job role.

Provide Access to Configure AI Agents in Oracle Permitting and Licensing

To give access to users without the PSC Application Administrator Job Role:

1. Go to the Security Console and create a new custom job role.

   Note: Make sure to enable permission groups.

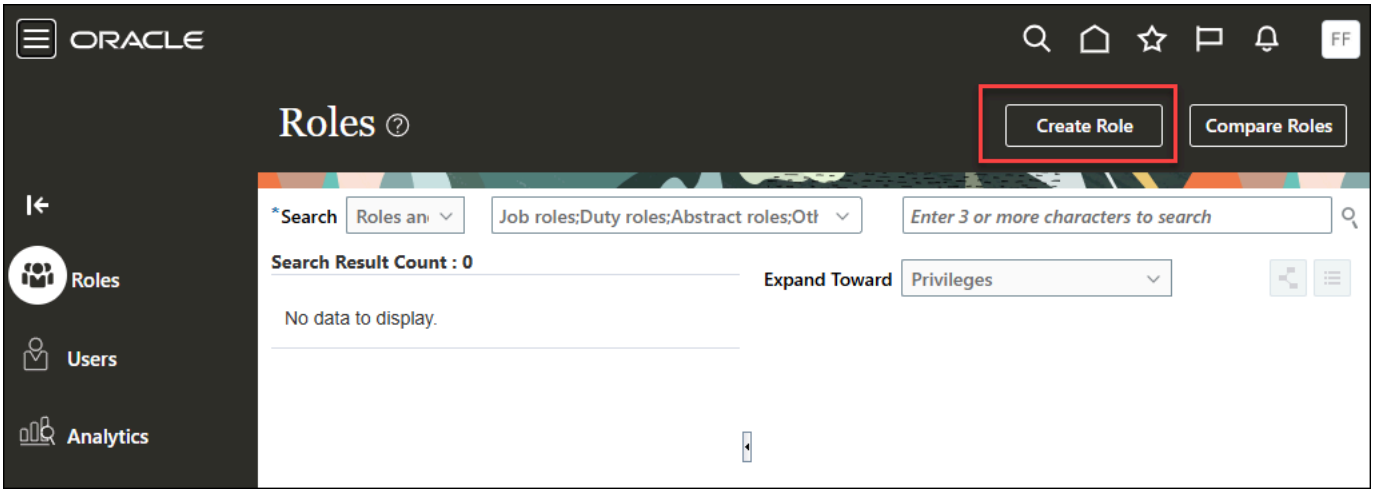

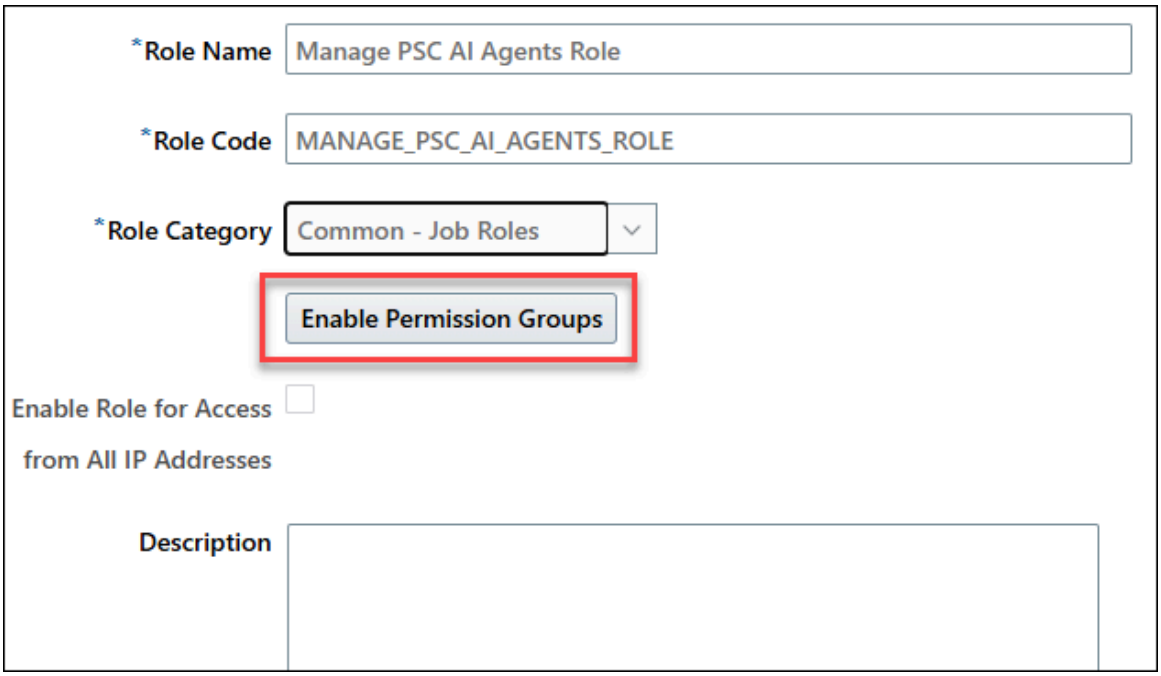

2. On the Role Hierarchy page, open the Roles and Privileges tab and add these roles:
   - PSC Intelligent Agent Management (ORA_PSC_AI_AGENT_MANAGEMENT_DUTY)

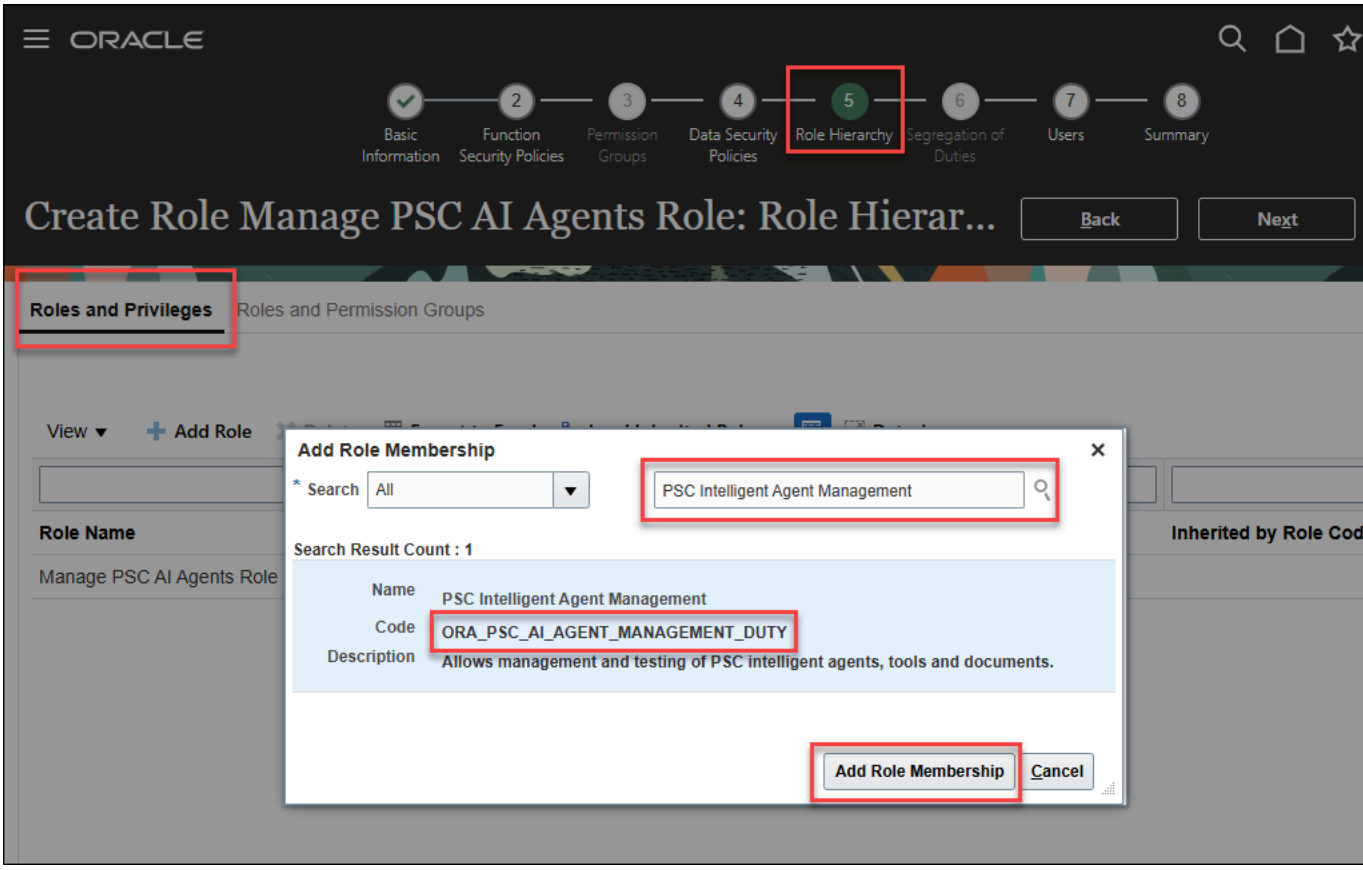

   - PSC Intelligent Agent Management (ORA_PSC_AI_AGENT_MANAGEMENT_DUTY_HCM)

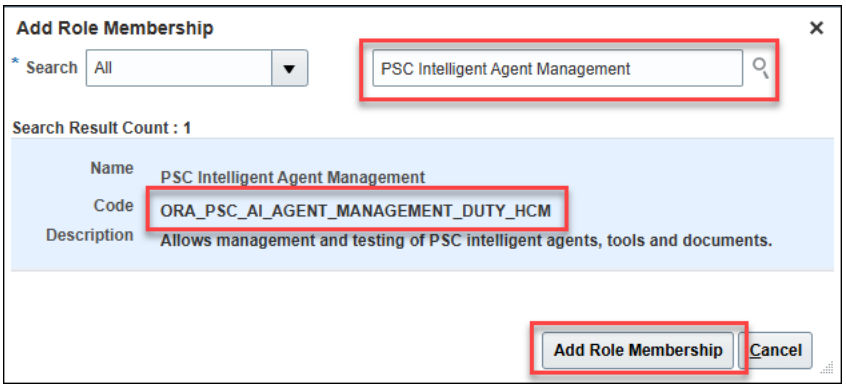

3. Open the Roles and Permission Groups tab and add the Fai Genai Agent PSC Administrator Duty (ORA_DR_FAI_GENERATIVE_AI_AGENT_PSC_ADMINISTRATOR_DUTY) duty role.

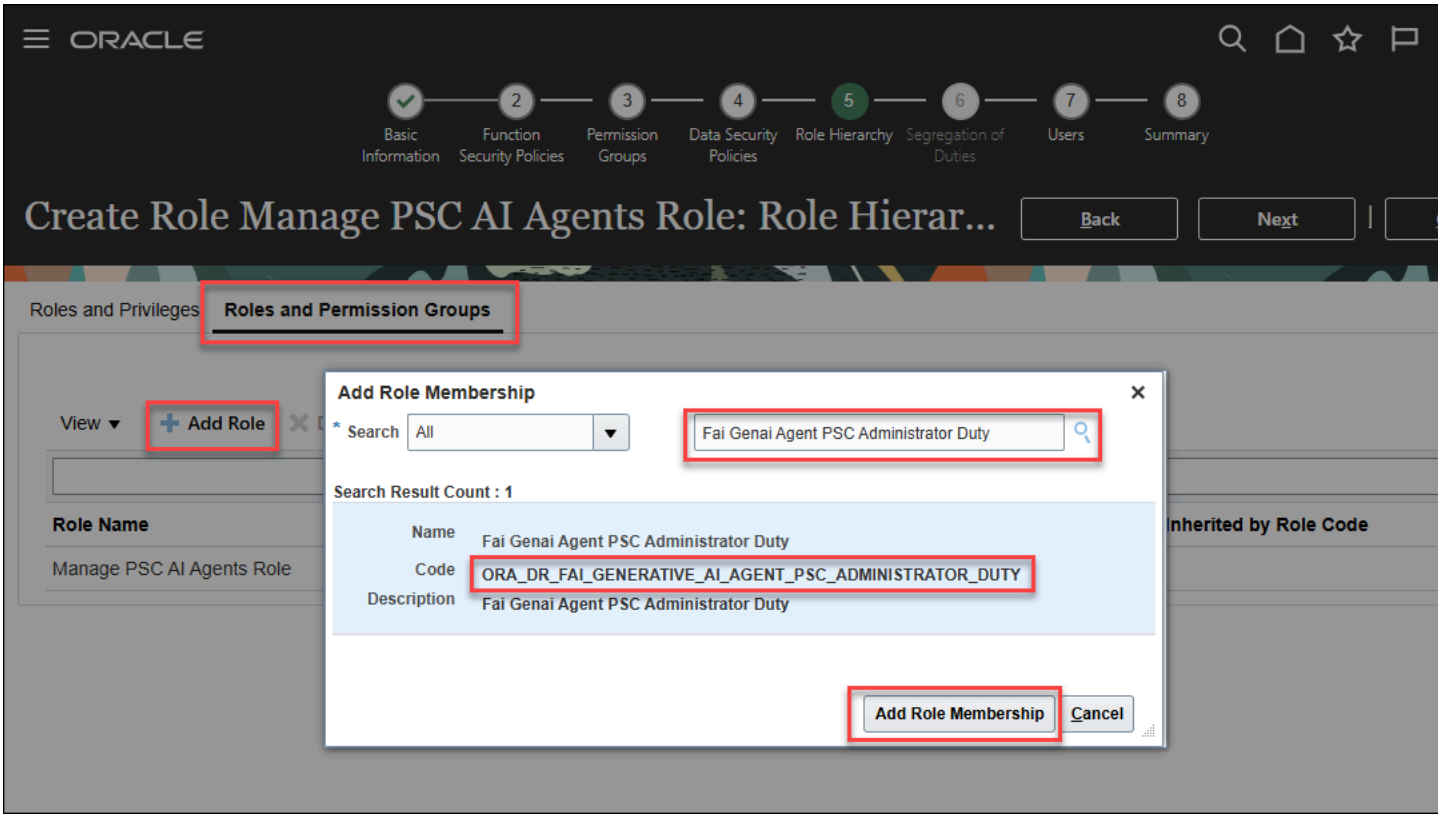

4. Save the custom role and assign this custom role to the appropriate job roles.

To give access to users with the PSC Application Administrator Job Role:

1. Go to the Security Console and create a new custom job role.

   Note: Make sure to enable permission groups.

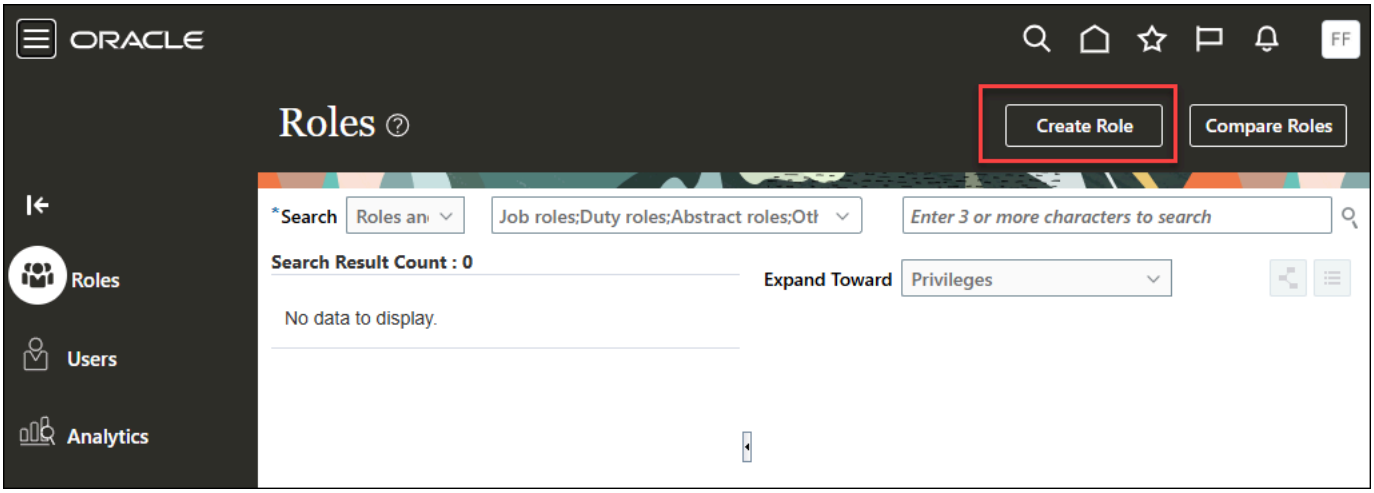

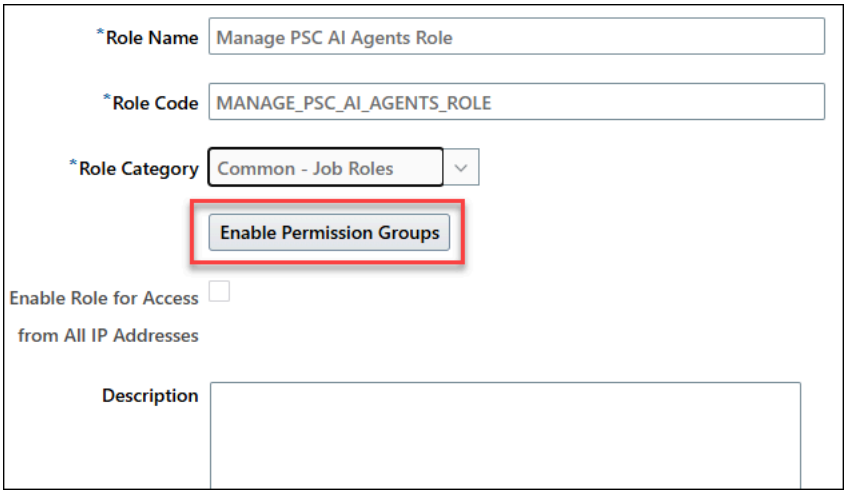

2. On the Role Hierarchy page, open the Roles and Permission Groups tab and add the Fai Genai Agent PSC Administrator Duty (ORA_DR_FAI_GENERATIVE_AI_AGENT_PSC_ADMINISTRATOR_DUTY) duty role.

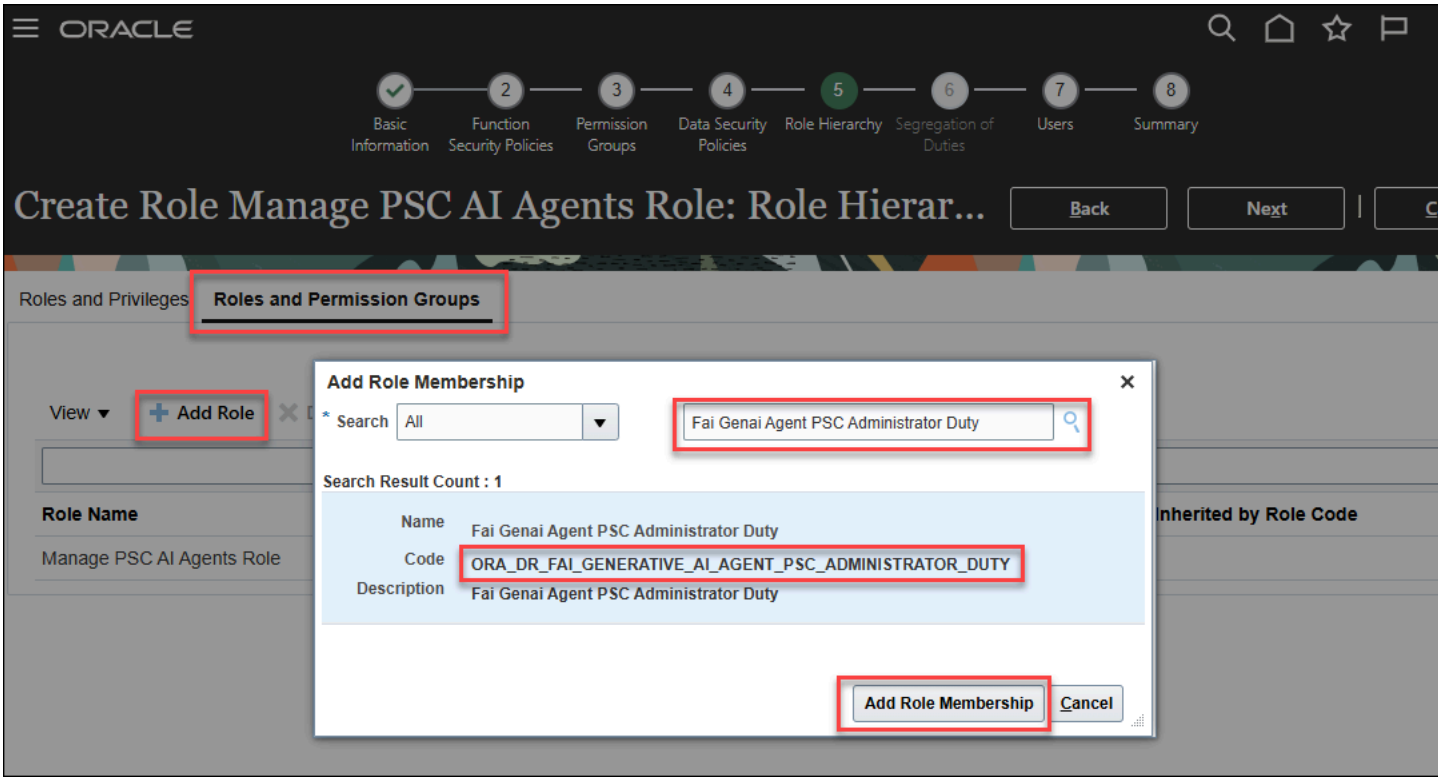

3. Save the custom role and assign this custom role to the appropriate job roles.

3 Before You Begin

Choose How to Create AI Agents

You can either use one of the preconfigured templates or create custom agents. If you want to create new agents or modify existing agents by adding additional tools and topics, a subscription is required. For more information, see *Fusion Custom AI Agents - Pricing FAQ (KB864473)*.

Here are some examples of how you choose between creating custom agents or using preconfigured templates, based on your use case.

Use Case	How to Create AI Agents
You want to turn on Benefits Agent, which allows employees to ask questions about their benefits, within AI Agent Studio. To ensure that this agent can understand the benefits specific to the organization, the documents specific to that organization must be uploaded.	Create an AI agent using a preconfigured agent team template. You can edit the existing agent, without adding any additional tools. **Note:** You can either use or copy a template and edit it. When you copy an agent template, you can automatically add a suffix to all the components to easily differentiate them.
You want to create a custom agent to help with onboarding by answering new hire questions and providing deep links to internal new hire resources if the agent is unable to answer. You can tailor the agent's topics and prompts to ensure they effectively address the specific needs of new hires.	Create a custom agent, and add the necessary tools and topics.
You want to create a Contract Assistant Agent to generate a new contract with auto-filled fields. This assistant will be provided with multiple contract templates to generate the new contract. The agent also reviews existing contracts to perform risk analysis, recommend changes and route the contract for final approvals. For these activities that need distinct expertise (template management, legal validation, recommendation, and routing), the administrator will need to build a supervisor agent that collaborates with other worker agents.	You can use either of these methods: • Create a team of custom agents, with one agent added as a supervisor to manage the others. • Use a preconfigured agent team template, and add more agents and a supervisor agent to it.

Use Case	How to Create AI Agents
You want to create a Supplier Quote Assistant agent to parse through quotes from multiple suppliers. The agent processes the quote document and maps the attributes. It then creates a draft requisition and submits for approval. For these activities, administrator will need to create a workflow agent that runs the tasks in the specified order.	Create an agent team of type Workflow, to run the predetermined set of tasks. Each node can perform a defined function, for example, extracting data, calling a business object, running an LLM, or sending an email in the order you define.

Get Started with AI Agent Studio

You can start with a preconfigured agent template or create your own agent team.

When using a preconfigured agent team template, the artifacts such as agents, tools, and topics, aren't directly editable. To change these artifacts, you create a copy of the artifact and add the copy to the agent team.

Note: Agents created directly within an agent team are directly editable within that team.

Here's a broad outline of the tasks involved in creating an agent team.

Task	Details
Define tools	To effectively define the tools required by an agent, you need to first identify the types of questions users might ask and then decide which tools the agent needs to answer those questions accurately. These are the available tools: • Oracle Fusion Cloud Applications business object: Business Object tools allow AI agents to retrieve, update, create, or delete business object records within Fusion Applications. Using this tool, agents can securely access information or call specific functions in the application. You can control what data the agents can access, by selecting the business objects and fields to use or ignore. AI agents adhere to the native security and role-based access controls of Fusion Applications, ensuring protection and privacy for your enterprise data. • Document tool: You can upload specific documents to be used by the AI agent, and provide natural language instructions on how the agent should use these documents. The agent can then search for information in the documents to provide a more exact answer to a user's question. • Email: This tool can access the email client to send emails that include summaries of interactions or details pulled from a knowledge store. • Deep link: A deep link will send a user directly to the part of Fusion Applications where they can update underlying information. For example, if the user moved and wants to update their home address in the HR system, a deep link can quickly route the user to the page for making that update. • External REST: You can connect to internal and external SaaS applications or public APIs using External REST tool. To connect to internal and third-party services, add the authorization

Task	Details
	parameters. For example, you can retrieve the real-time weather information for a specific location by creating an External REST tool that connects to the appropriate external API. • MCP: You can use the Model Context Protocol (MCP) tool and securely connect to external MCP servers, without building additional REST wrappers or plugin logic.
Define topics	Topics define the focus of the agent to a specific area of expertise. They are prompts that you can reuse across multiple systems or summarization prompts. For example, within a Benefits Administrator agent, we might define topics such as health policy coverage, vision policy coverage, and benefits enrollment. Use Topics to efficiently streamline your interactions with the agent. • Specify the instructions that help the agent decide which tools to use. • Enable the agent to better understand user intent by letting it identify and select the relevant topic based on the user's question. • Give each topic a clear, specific name, and include natural language instructions to ensure it's used correctly. You can reuse topics across agents.
Define credentials	To enable access to more services, you need to provide the necessary connection credentials. You can add credentials for this artifact: Custom LLM: In addition to the LLMs provided by Oracle, you can use other LLMs you've access to. You can add credentials for your LLM, and select it while creating your agent, node, or agent team.
Build new agents	Define the capabilities and scope of the agent. • Agent name • Product area the agent will work in • Natural language instructions to allow the agentic flow, or other agents, to understand the capabilities of this particular agent • Tools and topics the agent will need to use. In addition to the tools and topics you created, you can also add the predefined ones. In the Tools and Topics tabs, the predefined tools and topics are indicated by the icon.
Add a user (human) in the loop	If required, add an approval step for some actions that your AI agent will perform. The review step can be added at any point in the process for oversight and control over key actions, such as sending an email or updating a record.
Build the agent team	Create an agent team and add agents and other artifacts to it. Types of agent teams are: • Supervisor: A supervisor agent manages other agents and artifacts in the agentic flow. • Workflow: An agentic flow that does the tasks in a predetermined order. Agents and artifacts are added as nodes in the workflow, and each node performs a defined function, for example, extracting data, calling a business object, running an LLM, or sending an email. The node then passes its output to the next step.

Task	Details
	Note: You can also add predefined agents to your agent team. In the Agents tab, the predefined agents are indicated by the icon.
Test the agents	Make sure to test the agent before deploying to production. Ask a test query, and determine the accuracy and relevance of the agent's response. You can also see the instructions the agent is following, and the actions the agent has taken to arrive at the response.
Deploy the agent team	After defining and testing your agent team, you can deploy it directly from AI Agent Studio. • Embed the agent conversation chat experience into any website or application. • Trigger the agent from an external resource using Webhooks, or seamlessly embed the chat experience into HTML and React web pages.

Migrate RAG Agents to AI Agent Studio

If you've previously created any RAG agents in Fusion Applications, we recommend replacing your existing agent with a new one created in AI Agent Studio. For more information, see *Migrate Document Tools of RAG Agents*.

Add External REST Tool

You can enable your agents to securely connect to internal or external SaaS applications and public APIs by adding External REST tools in AI Agent Studio. Make sure you have a role with the Create and Edit Backends for Visual Builder Studio (ORA_FND_TRAP_PRIV) privilege assigned to it.

1. Go to **AI Agent Studio**.
2. From the Tools tab, click **Add**.
3. Add a new tool of type **External REST**.
4. Enter the name, code, and description for the tool, and select the appropriate family and product.
5. Select **Require Human Approval** to ensure that a person reviews and approves any action before the tool runs.
    a. Message: Enter instructions for the reviewer.
    b. Actions: Select the required action.
6. In the Authorization tab, add connection details for the third-party service.
    a. Instance URL: Enter the base domain of the external API, with no paths, parameters, or trailing slashes. For example, *https://api.weather.gov* is correct.
    b. Authentication: Define how your agents will authenticate with the external system.
    c. Description: Provide details about the tool, including what the tool does, and when the agent must call the tool.
7. In the Functions tab, provide endpoint details.
    a. Name: Assign a unique name to this endpoint. For example, **getWeatherByCoordinates**.
    b. Operation Type: Select the HTTP method required to process the REST call against the target endpoint.
    c. Resource Path: Specify the relative resource path to be appended to the Instance URL. Don't include the protocol or base domain. For example, **/points/{latitude},{longitude}** is correct.
    d. Description: Provide a clear summary of the endpoint's function and the specific task it performs within the workflow.

8. In the Parameters tab, define each dynamic variable in the path, including data type and purpose. For example, you can define latitude as the first parameter and longitude as the second.
9. In the Sample Queries tab, list example user queries that should trigger this tool.
10. In the Headers tab, specify any required metadata and security credentials for the external server.

The tool now appears in the Tools tab, and you can select it when creating agents, nodes, or agent teams.

Add MCP Tool

Using Model Context Protocol (MCP) tool, you can securely connect to external MCP servers, and use their capabilities inside agents and nodes, without building additional REST wrappers or plugin logic.

1. Go to **AI Agent Studio** and open the Tools tab.
2. Add a new tool of type **MCP**.
3. Enter the name, code, and description for the tool, and select the appropriate family and product.
4. Add MCP server connection details.
    - Instance URL: Enter the endpoint URL of the MCP server your agent will connect to.
    - Transport Type: Select the communication method:
        - Server Sent Events (SSE): Enables real-time, one-way streaming of data from a server to a client.
        - StreamableHTTP: Enables servers to independently handle multiple client connections using the HTTP POST and GET requests.
    - Credential Type: Select the authentication method required by the MCP server.
5. Update the details and create the tool.

Your tool now appears in the Tools tab, and you can select it when creating agents, nodes, or agent teams.

Add Your LLM

In addition to the LLMs provided by Oracle that are available for selection when creating an agent, node, or agent team, you can also use other LLMs by adding the credentials for it. These are the providers you can choose to add.

- OpenAI (through Microsoft Azure)
- Google Gemini (through Google Vertex AI)
- Google Gemini (through direct API)
- Anthropic (through Google Vertex AI)

To add the LLM, do these steps:

1. Go to **AI Agent Studio** and open the Credentials tab.
2. In the LLM tab, add the API key.
3. Add the details for the provider and save it.
    - Model: Choose the model you need from the available options.

- API Key: Add the API key for the selected model.
- Basic URL: Enter the URL to send LLM requests to.

Your LLM now appears in the LLM tab, and you can select it when creating agents, nodes, or agent teams.

4 Use Cases

Use Cases for Agents of Type Supervisor

Single Agent with Multiple Workers

In this pattern, a Benefits Advisor acts as a central agent, helping employees understand and make the most of their benefits packages. The agent coordinates several specialized workers to provide a personalized support for medical, dental, and vision coverage.

- The **Benefits Advisor** agent serves as the initial contact, helping the employee understand their benefits, and answers any preliminary questions.
- The **Policy Information** worker provides the employee with detailed information on available medical, dental, and vision benefits.
- The **Personal Enrollment** worker uses the employee's enrollment details to offer tailored advice on how they can optimize their benefits package.
- The **Eligibility and Assistance** worker answers any additional queries regarding eligibility, coverage, helping to resolve issues related to benefits.

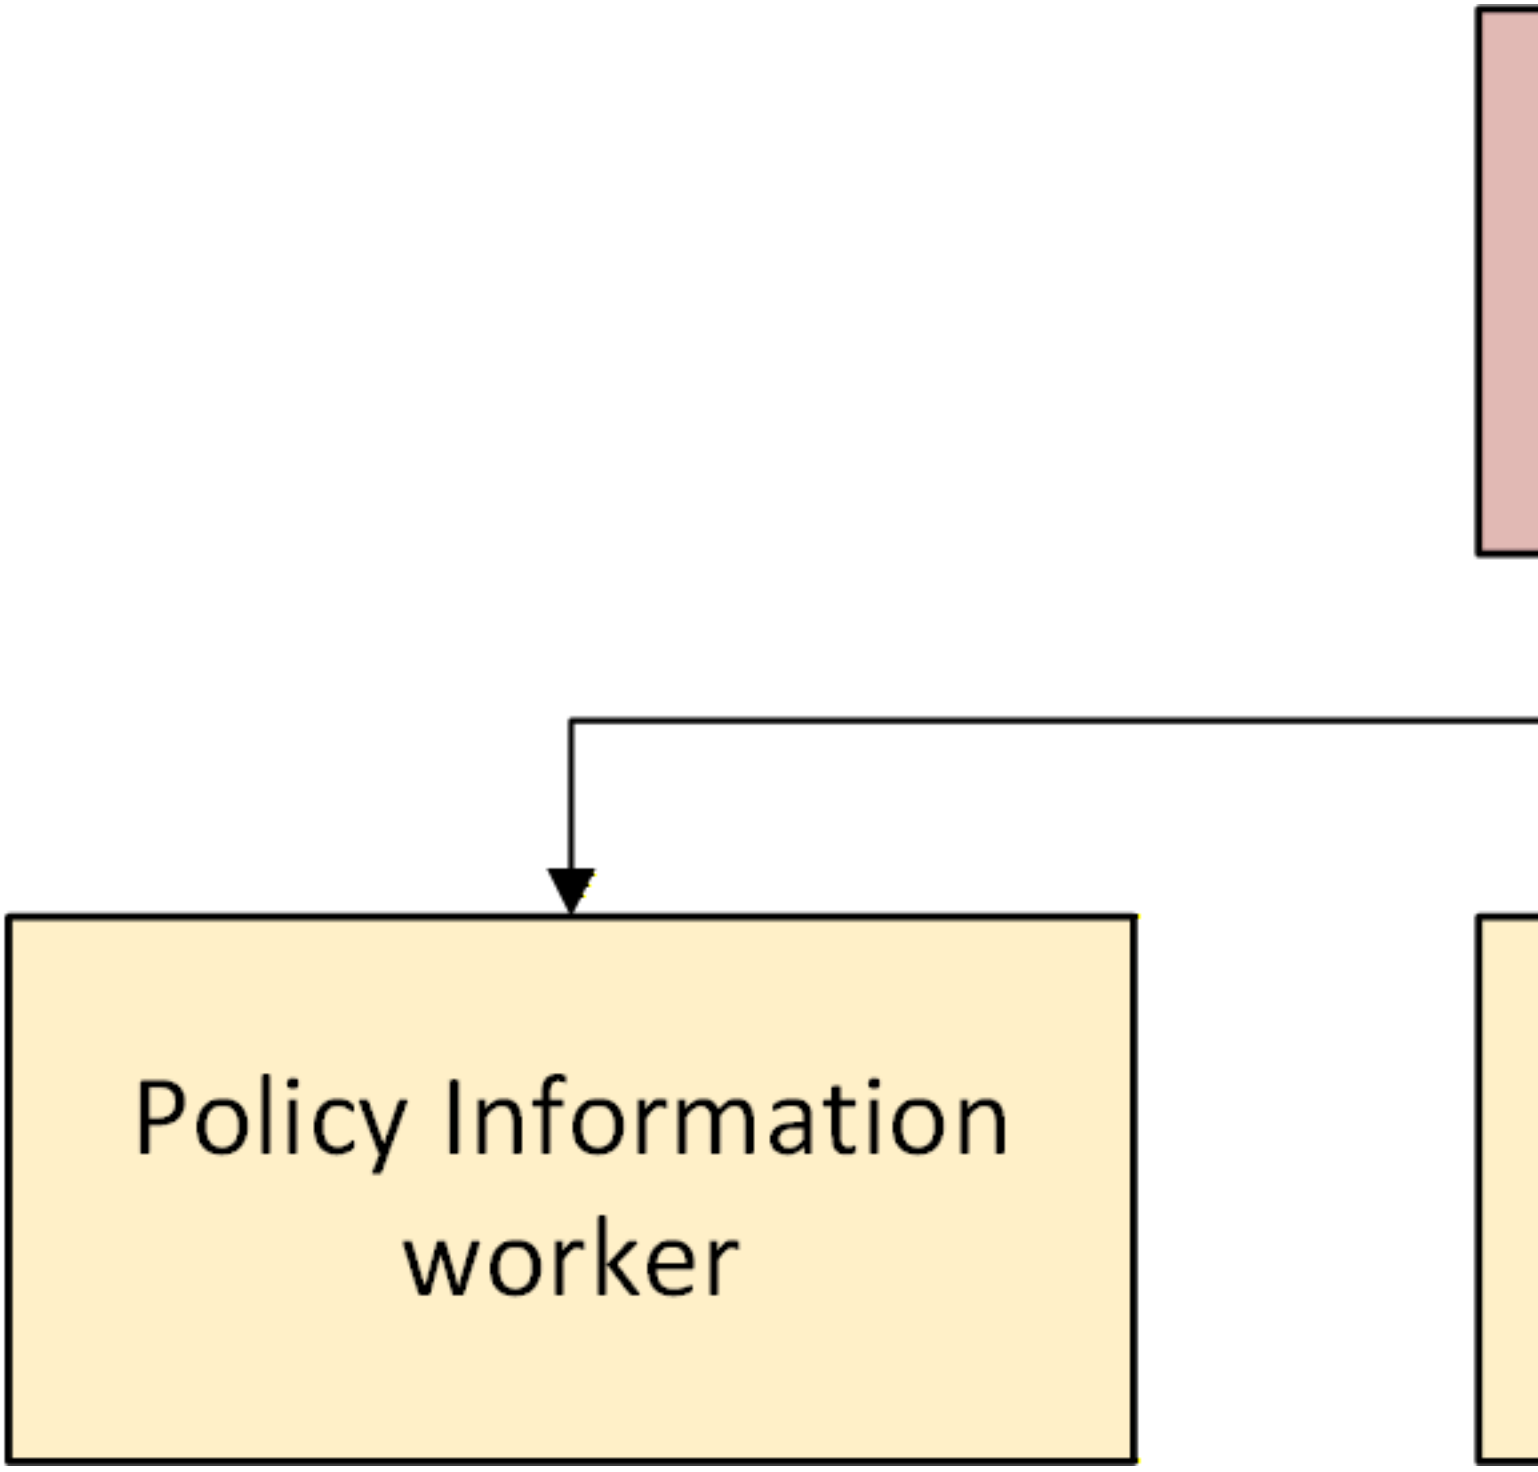

Single Agent with Multiple Tools and a Human in the Loop

This pattern shows the Benefits Advisor as a central agent that helps employees understand and optimize their benefits. This agent coordinates specialized workers to provide personalized insights on medical, dental, and vision coverage.

- The **Benefits Advisor** agent starts the interaction with a friendly introduction and asks for the employee's needs regarding benefits.
- The **Company Benefits Database** tool retrieves information on the company's benefits policies, eligibility, and plan details.
- The **Employee Enrollment Data** tool uses the employee's personal enrollment information to provide tailored advice, helping them understand and optimize their coverage.
- The **Benefits Optimizer** tool analyzes the employee's selections and offers personalized suggestions for improvement.
- The *Human-in-the-Loop* flag determines whether a user must provide explicit approval before the agent retrieves sensitive data, such as salary or personal information.

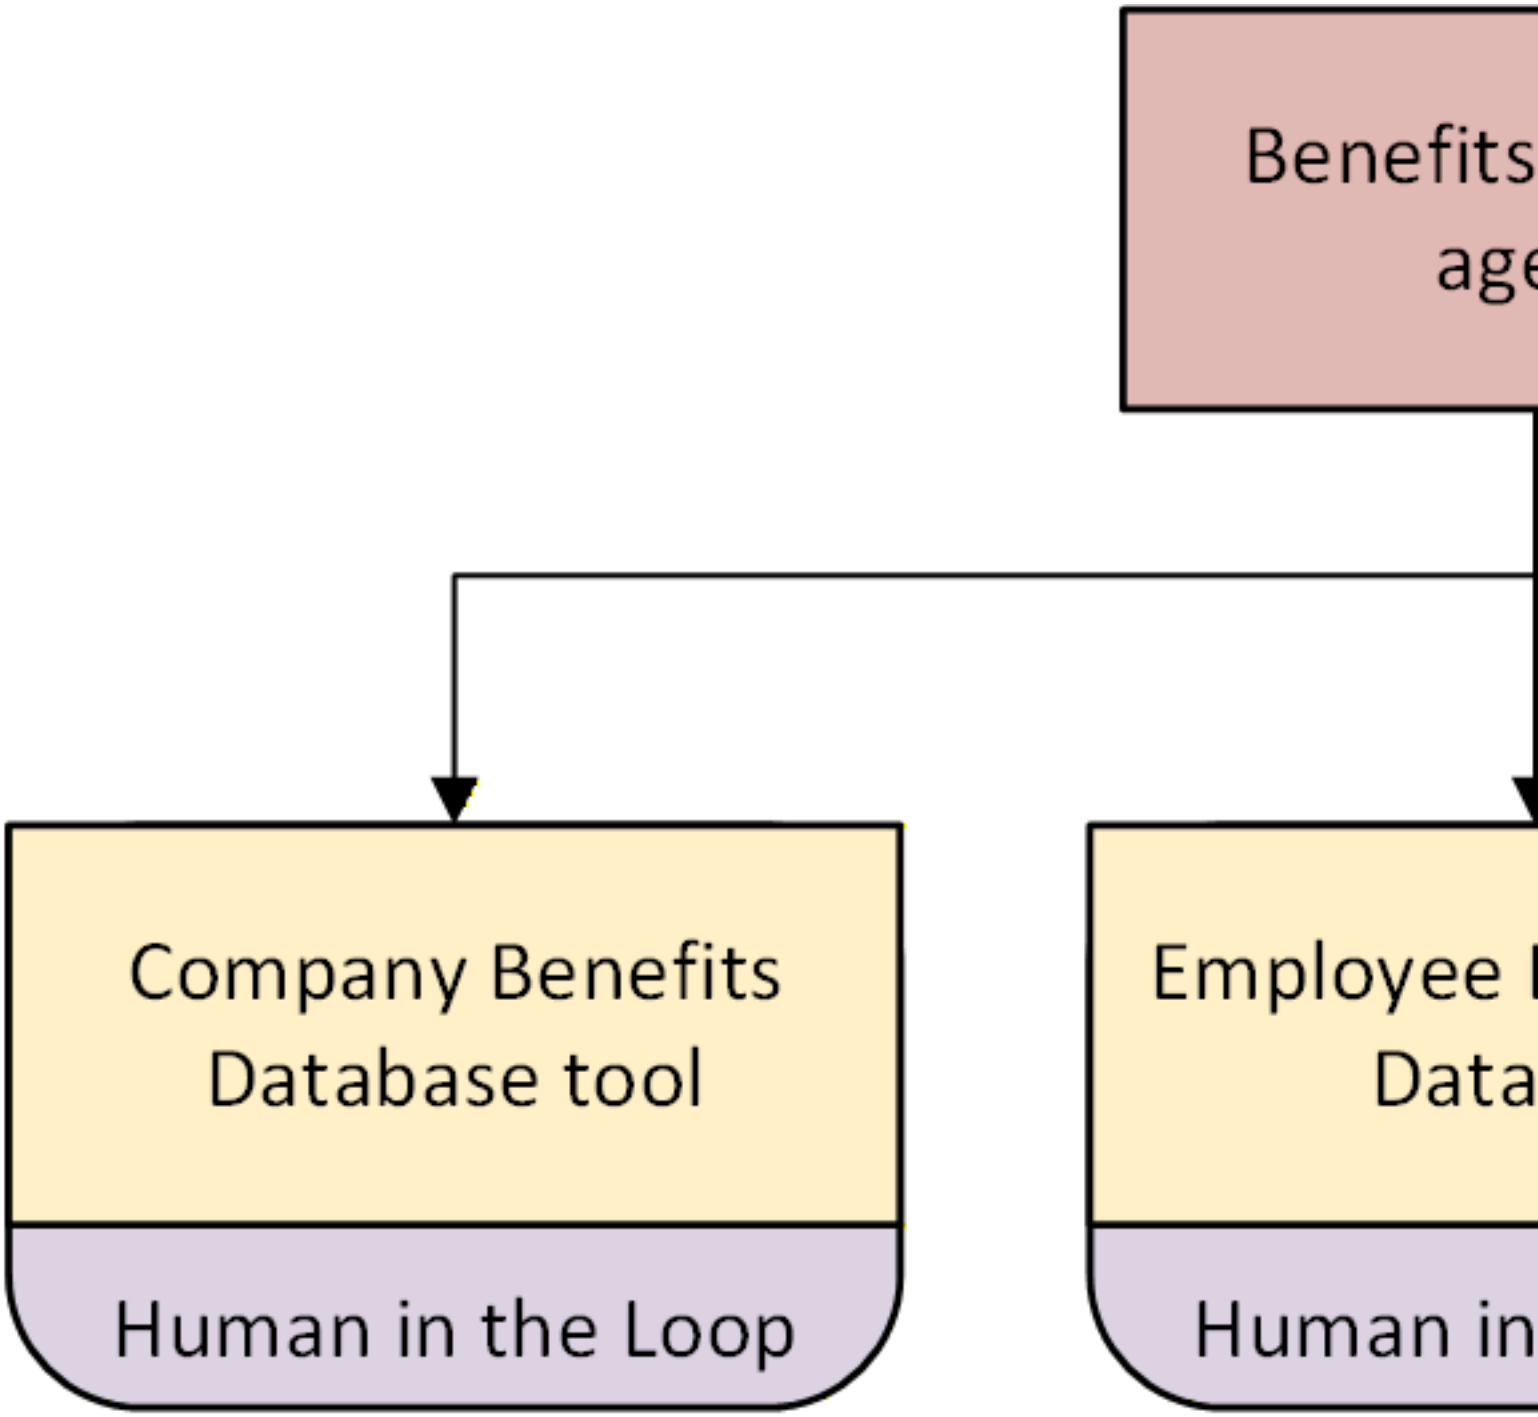

Multiple Agents With Multiple Tools and Delegations

In this pattern, the Central Onboarding Assistant agent acts as a supervisor, guiding new employees through their on-boarding process. The supervisor agent delegates tasks to specialized agents, each focused on a specific aspect of on-boarding, supported by dedicated workers.

Agent	Description
Policy and Compliance Agent	This agent helps employees understand and comply with company policies, and is supported by these workers. • **Policy Overview worker:** Shares key policies like leave, dress code, and workplace behavior. • **Compliance Training worker:** Provides access to mandatory training. For example, safety and social media policies. • **Document Management worker:** Ensures the employee completes and submits all required forms, such as tax forms and nondisclosure agreements.
Cultural Integration Agent	This agent helps employees adapt to the company's culture and connect with teams, and is supported by these workers. • **Company Values worker:** Explains the vision, mission, and goals of the organization. • **Community Introduction worker:** Connects employees with internal groups, events, and networking opportunities. • **Workplace Norms worker:** Provides guidance about unwritten rules, such as communication etiquette and meeting culture.
Tools and Resources Agent	This agent ensures that employees have access to tools, systems, and knowledge required for their job, and is supported by these workers. • **Account Setup worker:** Helps employees set up email, communication tools, and HR platforms. • **Resource Guide worker:** Provides access to handbooks, FAQs, and internal documentation. • **Technical Support worker:** Troubleshoots issues with tools or platforms.

Process

- The Central Onboarding Assistant agent acts as the entry point that gathers basic information about the employee's role and specific needs, and delegates tasks to specialized agents based on the query.
- The Policy and Compliance agent handles queries, or tasks, related to company rules and legal documentation, such as, "What's the process for requesting leave?"
- The Cultural Integration agent handles queries that focus on helping employees feel part of the team and aligned with the company values, such as, "How do I join company events?"
- The Tools and Resources agent handles queries about gaining access to systems and essential knowledge, such as, "How do I log into the time-tracking tool?"

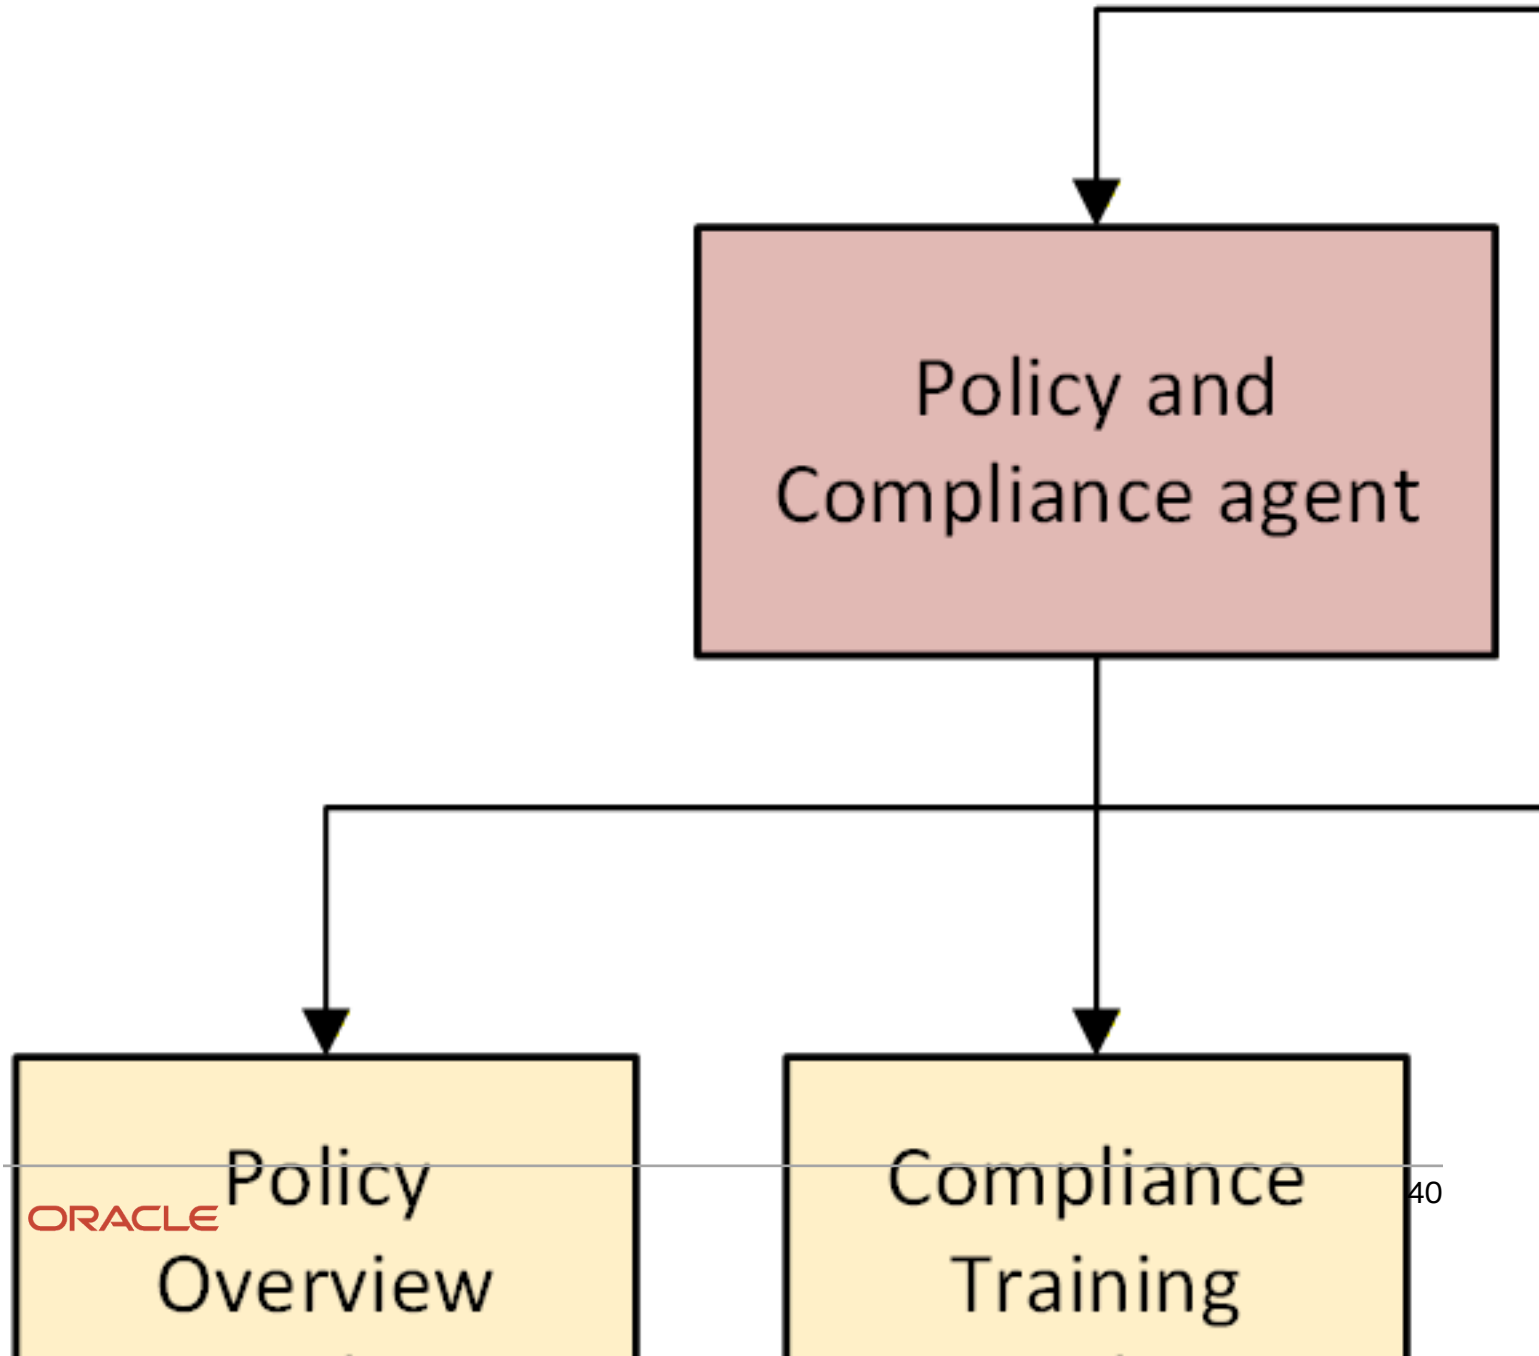
Policy and Compliance agent
Policy Overview
Compliance Training

Product-Specific Use Case Examples

Oracle Fusion Cloud CX

Agent Type	Key Characteristics	When to Use	Example
Customer service agent	• Interactive • Rule-based • Escalation-friendly • Customer-focused	• Repetitive, high-volume queries • Customer support • Query resolution	Chatbot for customer support agent answers FAQs and resolves issues.
Learning agent	• Adaptive • Feedback-driven • Continuously improves • Explainable	• Adaptive systems • Chatbot improvement • Demand forecasting refinement	Chatbot improvement agent learns from customer feedback to improve responses.

Oracle Fusion Cloud HCM

Agent Type	Key Characteristics	When to Use	Example
Copilots	• Collaborative • Proactive • Decision-making • Domain-specific	• Complex tasks requiring expertise • Draft documents • Input validation	Hiring copilot agent helps to build job descriptions and suggests interview questions.
Document referencing agent	• Knowledge-intensive • Accurate • Up-to-date • Query-driven	• Knowledge-intensive tasks • Answers to policy or FAQ questions • Compliance checks	Benefits agent answers employee questions about benefits.
Evaluating agent	• Assessment-focused • Criteria-based • Actionable insights • Feedback-driven	• Performance evaluation • Candidate ranking • Feedback analysis	Candidate evaluation agent ranks candidates based on interview performance.
Scheduling agent	• Task-oriented • Reactive • Limited decision-making • User-centric	• Routine task scheduling • Calendar management • Appointment booking	Interview scheduling agent coordinates availability between candidates and hiring managers.
Task automation agent	• Repetitive task handler • Efficient • Error-resistant	• Repetitive, high-volume tasks • Payroll processing	Payroll processing agent automates time sheet collection and payslip generation.

Agent Type	Key Characteristics	When to Use	Example
	• Trigger-based	• Form generation	

Oracle Fusion Cloud SCM

Agent Type	Key Characteristics	When to Use	Example
Collaborative workflow agent	• Multi-stakeholder • Real-time coordination • Role-based task allocation • Communication-driven	• Complex processes requiring collaboration • Supplier collaboration • Crisis management	Supplier collaboration agent coordinates with procurement, supplier, and warehouse bots to work together.
Decision-making agent	• Data-driven • Recommendation-focused • Transparent • Adaptive	• Data-driven tasks • Supplier evaluation • Candidate shortlisting	Supplier evaluation agent evaluates suppliers based on cost and quality.
Goal-seeking agent	• Optimization-focused • Real-time adjustments • Objective-driven • Explainability	• Resource allocation • Route planning • Workforce scheduling	Route optimization agent plans optimal delivery routes.
Monitoring agent	• Proactive • Alert-driven • Data-intensive • Response-oriented	• Proactive management • Inventory monitoring • Social media tracking	Inventory monitoring agent tracks in-stock levels and alerts for replenishment.

5 Create AI Agents

Create AI Agents Using Preconfigured Templates

You can create and deploy AI agents using preconfigured templates available in these two tabs:

- AI Agents: Access the preconfigured templates built and provided by Oracle. When using a preconfigured agent team template, the artifacts such as agents, tools, and topics aren't directly editable. To change these artifacts, you create a copy of the artifact and add the copy to the agent team.
- Marketplace: Access the preconfigured templates built and provided by Oracle partners. View template details using the ⓘ icon. To create a copy of any template, click **Create**.

To use preconfigured agent team templates from the AI Agents tab:

1. Go to **AI Agent Studio**.
2. Select **Use Template** from the required agent team and provide details for the new agent team.

   > **Tip:** To automatically add a suffix to your agent team, choose **Copy Template** instead of **Use Template**.

   Details Tab

Field	Description
Family	Select the family to which this agent team belongs.
Product	Select the product within the family to which this agent team belongs.

Field	Description
Maximum Interactions	Indicate the number of times an agent within this agent team can interact with the topics and tools assigned to it.

LLM Tab

Field	Description
Provider	Select the LLM for your agent.

Security Tab

Field	Description
Add	Select the roles which will have access to this agent team.

Questions Tab

Field	Description
Starter Questions	Enter initial questions for the agent team.
Follow-up Questions	Enable this to indicate that the agent team can ask follow-up questions based on the user's conversation history.
Prompt	Enter the prompt to be used for the follow-up questions. For example: **using {chatHistory} generate 3 follow up questions in json format. JSON Schema format mentioned below. Remove the ``json markdown from the output. Here is the JSON Schema format the output should adhere to:** **[{"question": "<put first question generated here>"}, {"question": "<put second question generated here>"}, {"question": "<put third question generated here>"}]**

Field	Description
Insert Expression	Add additional variables to the prompt. For example, to add the current system date to your prompt, select the **Current Date Time** option.

Chat Experience Tab

You can allow users to upload attachments while interacting with agents. Users can upload up to five files, with a total combined size limit of 50 MB. These file types are supported: PDF (tagged or scanned), TXT, DOCX, XLSX, PNG, and JPEG.

Field	Description
Enable file upload	Select this option to allow users to upload files from local storage.
Enable third-party upload	Select this option to allow users to upload files from connected cloud storage accounts. To enable this functionality, add credentials for at least one provider by navigating to the Credentials tab in AI Agent Studio and completing the Chat Experience setup.

> **Note:** To process attachments, make sure you add the runtime processing tool, **MultiFileProcessor**, to your agent or to the relevant agent team.

Output Tab

Define the overall structure of the agent's output using JSON schema, to specify the exact output.

Field	Description
Specification Mode	Select this mode to directly modify the JSON schema for the output.
Simple Mode	Select this mode to define the output values and types. The corresponding JSON schema will be generated automatically and displayed in the specification mode for any further changes.

3. Continue to edit and add details for the agent team. If you want to edit the artifacts in the agent team including agents, tools, and topics, you can create a copy of the artifact and edit it. After editing, remove the original artifact from the agent team, and then add the copy you created. See *Can I edit a preconfigured agent team?*

   > **Note:** The agents and other artifacts included within an agent template are optimized to provide the best usage of the agent team. We recommend not to change the basic functionality of the artifacts because that might impact your agent team's performance.

4. If needed, use ▷ to test the agent team. For any required fine-tuning, you can edit the agent team using ⚙.
5. Publish your agent team.

Users can view the published agentic flows from the AI Agents page. To open this page, add **agent-explore** to the end of the URL for AI Agent Studio. For example, **https://example.com/myApp/redwood/human-resources/ai-studio/agent-explore**.

Create Custom AI Agents of Type Supervisor

The supervisor agent orchestrates and plans tasks for your agent team in Supervisor flows. It directs other agents on how to interact and generate responses. You can configure supervisor agents as a single agent or within a multiple agent pattern. To see some of the common supervisor agent patterns, go to *Use Cases for Agents of Type Supervisor*.

Tip: For recommendations about improving your agent's efficiency, see *How do I make agents respond faster?*.

1. Go to **Navigator** > **Tools** > **AI Agent Studio**.
2. From the Tools tab, add the required tools.

   For example, to create an HR benefits administrator agent that can answer questions related to medical, vision, dental, retirement, and stock plans, these are some of the tools needed.
   - Document tools, with the organization's health and financial benefits documents.

     Note: Before you add any document tools to an agent, make sure you've done these things:
     - Set the status of the document tool as **Ready to Publish**.
     - Run the **Process Agent Documents** scheduled process from the Scheduled Processes work area.
     - After the scheduled process completes, set the status of the document tool as **Published**.
   - Business object tools, to fetch the employee enrollments data.
   - Calculator tool, to check balance amounts and percentages.
3. From the Topics tab, add the required topics. Make sure you include instructions about these key areas.
   - What the topic is about
   - Tools to use with the topic
   - Examples of possible questions
   - Any guidelines and guardrails
4. From the Agents tab, add an agent.
   - In **Maximum Interactions** field, specify the number of times the agent can interact with the topics and tools assigned to it.
   - Describe the persona and role of the agent, including the tone to use.
   - Add the prompt for the agent, and select any variables from **Insert Expression**.
5. Assign the required tools and topics to the agent, and create it.

6. From the **Agent Teams** tab, add an agent team and provide details for the new agent team.
    - **Details Tab**

Field	Description
Family	Select the family to which this agent team belongs.
Product	Select the product within the family to which this agent team belongs.
Type	Select **Supervisor** to create an agent team where one agent acts as supervisor and manages other agents and artifacts.
Maximum Interactions	Indicate the number of times an agent within this agent team can interact with the topics and tools assigned to it.

    - **LLM Tab**

Field	Description
Provider	Select the LLM for your agent team.

    - **Security Tab**

Field	Description
Add	Select the roles which will have access to this agent team.

    - **Questions Tab**

Field	Description
Starter Questions	Enter initial questions for the agent team.
Follow-up Questions	Enable this to indicate that the agent team can ask follow-up questions based on the user's conversation history.
Prompt	Enter the prompt to be used for the follow-up questions. For example: using {chatHistory} generate 3 follow up questions in json format. JSON Schema format mentioned below. Remove the ``json markdown from the output. Here's the JSON Schema format the output should adhere to: [{"question": "<put first question generated here>"}, {"question": "<put second question generated here>"}, {"question": "<put third question generated here>"}]
Insert Expression	Add additional variables to the prompt. For example, to add the current system date to your prompt, select the Current Date Time option.

- **Chat Experience Tab**

  You can allow users to upload attachments while interacting with agents. Users can upload up to five files, with a total combined size limit of 50 MB. These file types are supported: PDF (tagged or scanned), TXT, DOCX, XLSX, PNG, and JPEG.

Field	Description
Enable file upload	Select this option to allow users to upload files from local storage.
Enable third-party upload	Select this option to allow users to upload files from connected cloud storage accounts. To enable this functionality, add credentials for at least one provider by navigating to the Credentials tab in AI Agent Studio and completing the Chat Experience setup.

  > **Note:** To process attachments, make sure you add the runtime processing tool, **Chat Attachments Reader**, to the relevant agent.

- **Output Tab**

Field	Description
Specification Mode	Select this mode to directly modify the JSON schema for the output.
Simple Mode	Select this mode to define the output values and types. The corresponding JSON schema will be generated automatically and displayed in the specification mode for any further changes.

7. Select **Create**.
8. Select and click **New Supervisor Agent**.
9. Enter details for the supervisor agent and create it. This supervisor agent is specific only to the agent team in which it's created, and can't be reused.
10. Select and add any existing agents as worker agents, or create worker agents.

    When you create worker agents from within the agent team, those agents are available only for this agent team.
11. Add any needed artifacts such as tools and topics.
12. If needed, use to test the agent team. For any required fine-tuning, you can edit the agent team using .
13. Publish your agent team.

Users can view the published agents from the AI Agents page. To open this page, add **agent-explore** to the end of the URL for AI Agent Studio. For example, **https://example.com/myApp/redwood/human-resources/ai-studio/agent-explore**.

Create Custom AI Agents of Type Workflow

Agent teams of type Workflow consist of a sequence of nodes used for deterministic, rule-based orchestration of tasks in which every step is preconfigured. Workflows are ideal for scenarios where compliance, repeatability, and governance

are essential. The workflow tasks are represented as a connected sequence of nodes. Each node within the workflow performs a defined function, such as extracting data, calling a business object function, running an LLM, or sending an email. The output from each node is seamlessly passed to the subsequent step, ensuring a controlled and efficient process flow.

When designing your workflow, keep these process breakdown steps in mind:

- Model the manual workflow. Start from a high-level process and break it down into smaller, automatable steps.
- Identify task dependencies and group nodes by type. For example, distinguish between backend logic tasks and those driven by LLMs. Doing so helps to streamline execution and minimize unnecessary LLM usage.
- For each task, define the node type, required inputs, expected outputs, and any transformation logic needed.
- For recommendations about improving your agent's efficiency, see *How do I make agents respond faster?*

Let’s look at an example to illustrate how a workflow operates in practice. Suppose an employee needs to purchase a product or service and contacts a supplier to request a quotation. The supplier responds with a quotation, which might vary in format and content, but typically includes key information such as the seller’s details, a list of items, and the total quoted price. The quotation is processed internally, and a requisition order is created.

Here's an example workflow for this scenario, created in AI Agent Studio.

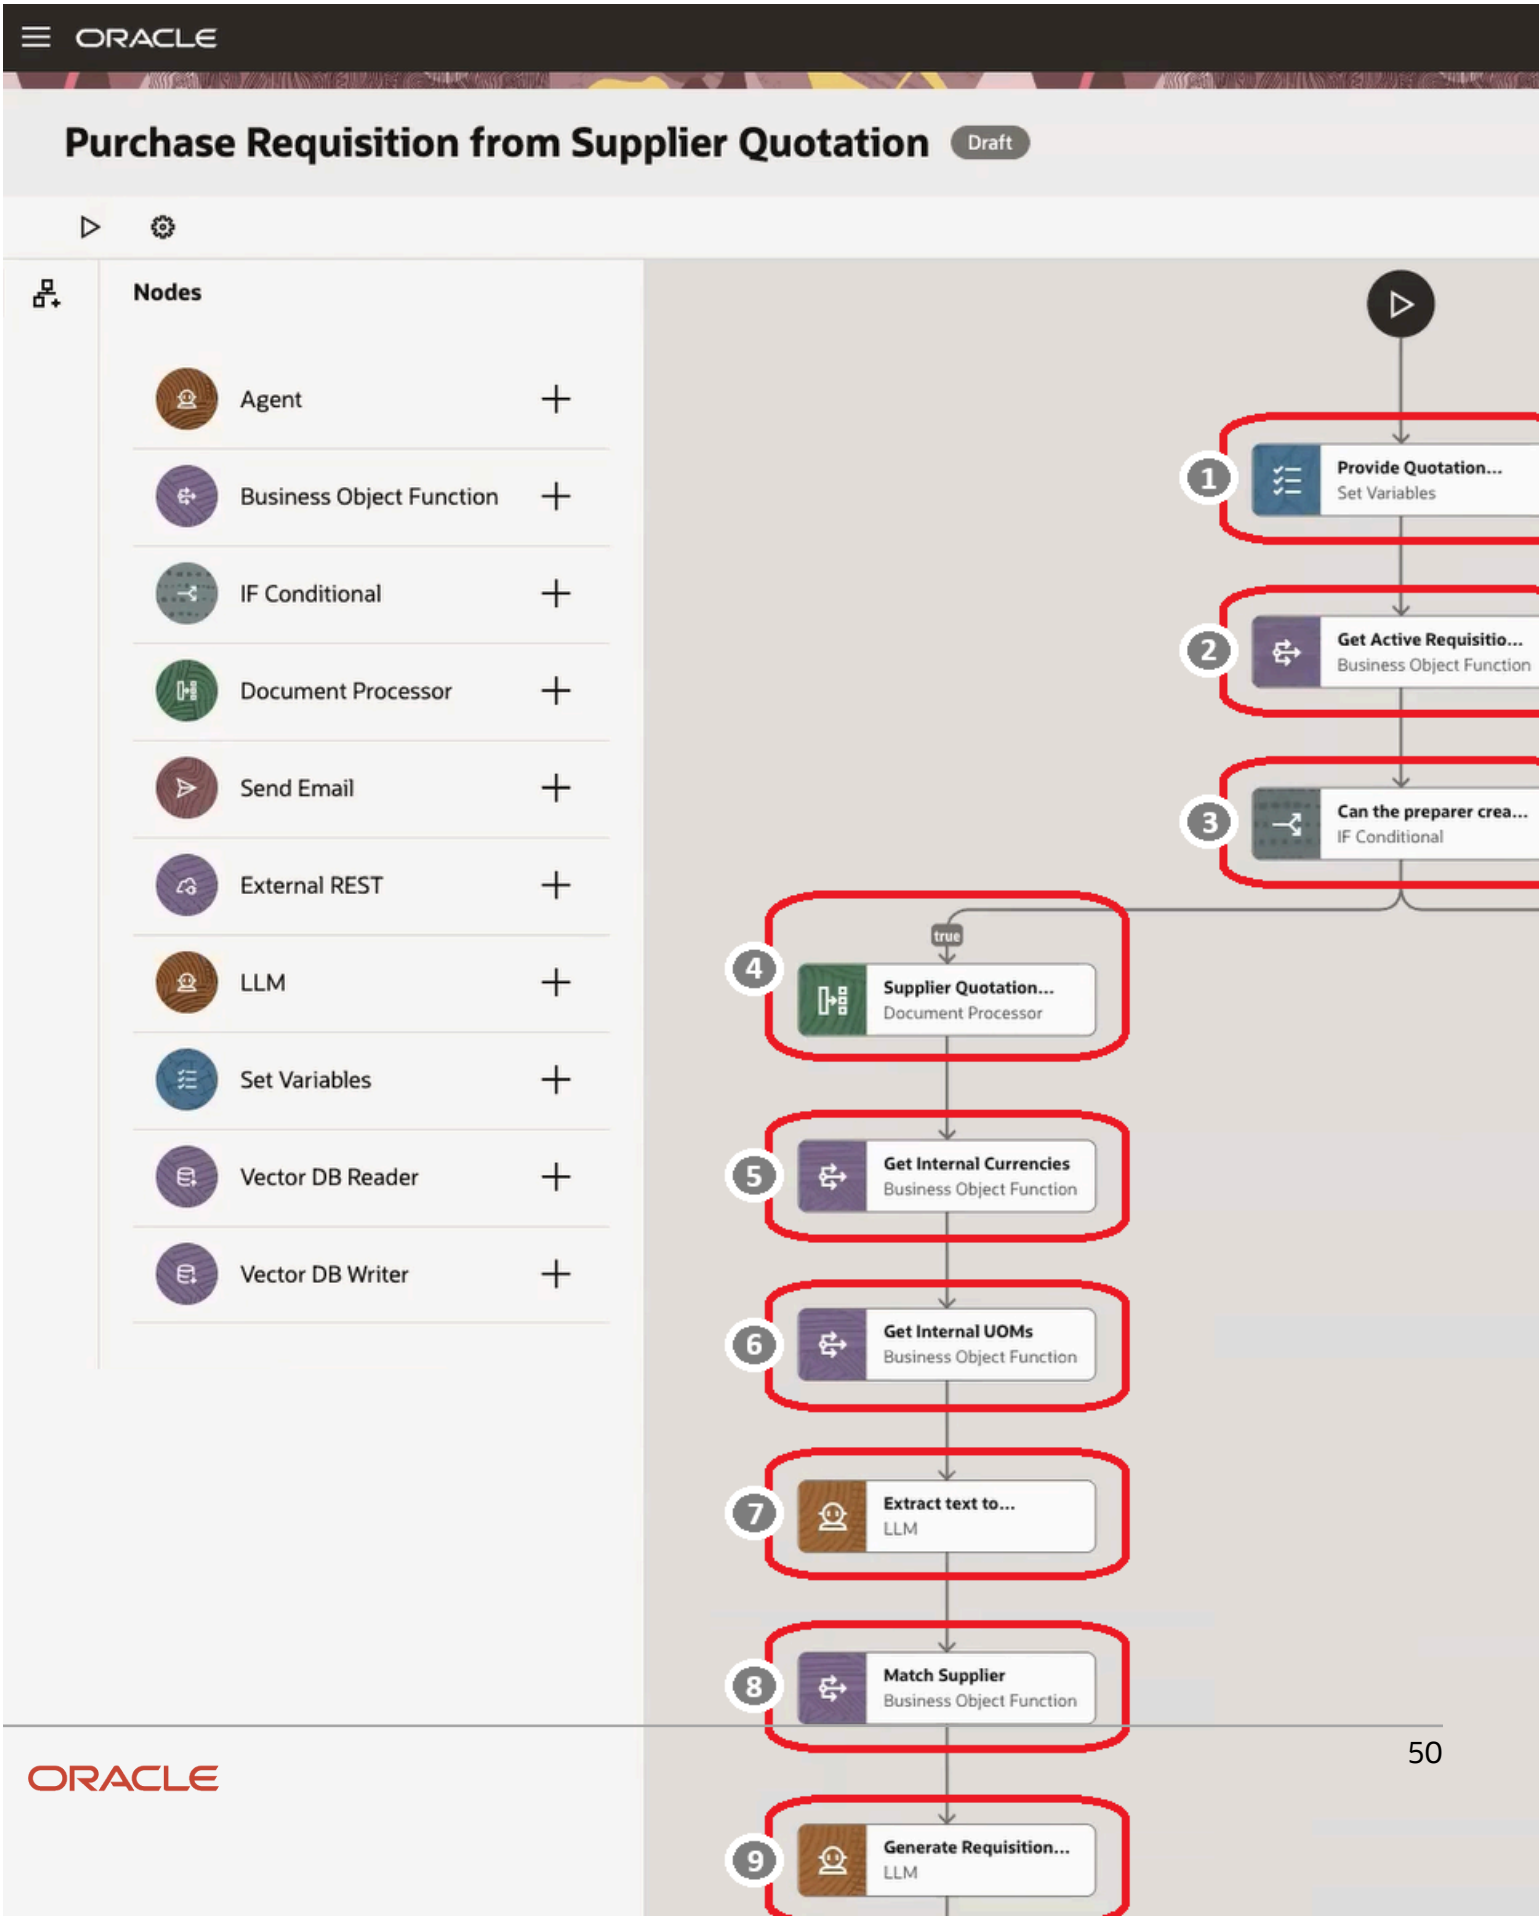
ORACLE
Purchase Requisition from Supplier Quotation
Draft
Nodes
Agent
Business Object Function
IF Conditional
Document Processor
Send Email
External REST
LLM
Set Variables
Vector DB Reader
Vector DB Writer
1
Provide Quotation...
Set Variables
2
Get Active Requisitio...
Business Object Function
3
Can the preparer crea...
IF Conditional
true
4
Supplier Quotation...
Document Processor
5
Get Internal Currencies
Business Object Function
6
Get Internal UOMs
Business Object Function
7
Extract text to...
LLM
8
Match Supplier
Business Object Function
9
Generate Requisition...
LLM

Callout Number	Node Name	Node Type	Description
1	Provide Quotation	Set variables	Sets the values for the input variables.
2	Get Active Requisition	Business Object Function	Retrieves the default preferences of the requester.
3	Can the preparer create a requisition?	IF conditional	Validates whether the user has permission to create a requisition. • If true, the workflow continues to the subsequent nodes. • If false, the workflow indicates that a requisition wasn't created, sends a failur
4	Supplier Quotation	Document processor	Downloads the document from the WebCenter Content server and parses the cont
5	Get Internal Currencies	Business Object Function	Pulls the valid internal currency value.
6	Get Internal UOM	Business Object Function	Pulls the valid internal units of measure value.
7	Extract text to LLM	LLM	Uses LLM to convert the unstructured text from upstream nodes to a structured JS
8	Match supplier	Business Object Function	Checks and validates the supplier information extracted earlier.
9	Generate requisition payload	LLM	Maps all gathered data from the document and database to generate a valid post r requisition payload.
10	Create purchase requisition	Business Object Function	Invokes the API to create the draft requisition. After this call, the requisition is save
11	Is requisition created?	IF conditional	Checks if the requisition is successfully created. • If true, the workflow proceeds to the success confirmation node. • If false, the workflow proceeds to the send failure email node.
12	Success confirmation	LLM	Displays a message indicating the successful creation of the requisition.
13	Send failure email	Send email	Sends an email to the user with failure details.

Here are the steps to create your own workflow.

1. Go to **AI Agent Studio**.
2. From the **Agent Teams** tab, add an agent team and provide details for the new agent team.

   Details Tab

Field	Description
Family	Select the family to which this agent team belongs.
Product	Select the product within the family to which this agent team belongs.

Field	Description
Type	Select **Workflow**.

LLM Tab

Field	Description
Provider	Select the LLM for your agent team.

Security Tab

Field	Description
Add	Select the roles which will have access to this agent team.

Triggers Tab

Create triggers to initiate the workflow. Triggers define the conditions or events that start the workflow and determine when and why it begins.

Field	Description
Type	Select the data type of the trigger.

Variables Tab

Define workflow-level variables, making them accessible to all nodes within the workflow. Variables are ideal for storing IDs, constants, or any values you need to share and use across different parts of the workflow.

Field	Description
Type	Select the data type of the variable.

Chat Experience Tab

You can allow users to upload attachments while interacting with agents. Users can upload up to five files, with a total combined size limit of 50 MB. These file types are supported: PDF (tagged or scanned), TXT, DOCX, XLSX, PNG, and JPEG.

Field	Description
Enable file upload	Select this option to allow users to upload files from local storage.

Field	Description
Enable third party file upload	Select this option to allow users to upload files from connected cloud storage accounts. To enable this functionality, add credentials for at least one provider by navigating to the Credentials tab in AI Agent Studio and completing the Chat Experience setup.

Note: To process attachments, make sure you add the runtime processing tool, **Runtime File Processor**, to the relevant node.

Output Tab

Define the overall structure of the workflow's output using JSON schema, to specify exactly what you will receive when the workflow is completed. If no schema is specified, the output type defaults to that of the individual node. When a schema is provided, the system validates the outputs against the schema and enables auto-complete for those fields.

Field	Description
Specification Mode	Select this mode to directly modify the JSON schema for the output.
Simple Mode	Use this mode to define output values and types. The corresponding JSON schema will be generated automatically and displayed in the specification mode for you to edit.

Error Handling Tab

Specify where to send an email when a workflow instance encounters a permanent error. You can also use context expressions in these fields for additional clarity.

3. Click **Create**.

4. Add nodes to the workflow by hovering over the line between the start and end of the workflow. Click ⊕ to choose the type of node you want to add.

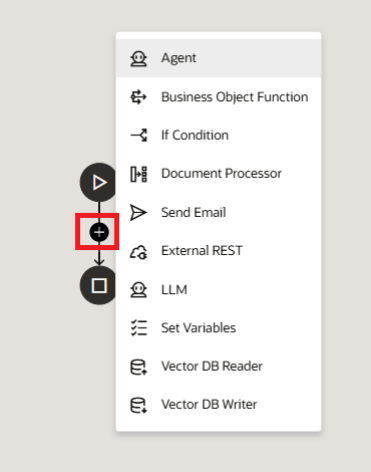

Tips

- Connect nodes in sequence to control how inputs are received, processed, and forwarded throughout the workflow.
- You can insert new nodes before, after, or between existing nodes as needed.
- For more information about Vector nodes, see *Vector Write and Read Nodes* .
- Once created, you can't rearrange the nodes. To change a node's position, delete it and create it again in the desired spot.

5. If needed, use ▷ to test the agent team. For any required fine-tuning, you can edit the agent team using ⚙ .
6. Publish the agent team.

   Users can view the published agentic workflows from the AI Agents page. To open this page, add agent-explore to the end of the URL for AI Agent Studio. For example, `https://example.com/myApp/redwood/human-resources/ai-studio/agent-explore`.

Vector Write and Read Nodes

To automate business processes effectively, your AI workflows require access to high-quality, well-structured, and reusable knowledge. You can use Vector Write and Read nodes in AI Agent Studio to store knowledge as embeddings and retrieve it intelligently with semantic search and metadata filters. By leveraging these nodes, you can design workflows that capture, manage, and reuse critical business knowledge across processes, thereby improving accuracy, scalability, and reliability in an enterprise environment.

Here's what you need to know to effectively use Vector Write and Read nodes in your workflow agents, along with some best practices for organizing and managing enterprise knowledge for optimal retrieval.

Vector Write Node

Vector write nodes store high-value knowledge as embeddings for future semantic retrieval. These nodes are also referred to as vector store nodes.

When you create a node, you can configure all values as expressions that resolve dynamically at runtime. Here's an example of the node creation window. Each numbered field is described in detail in the following table.

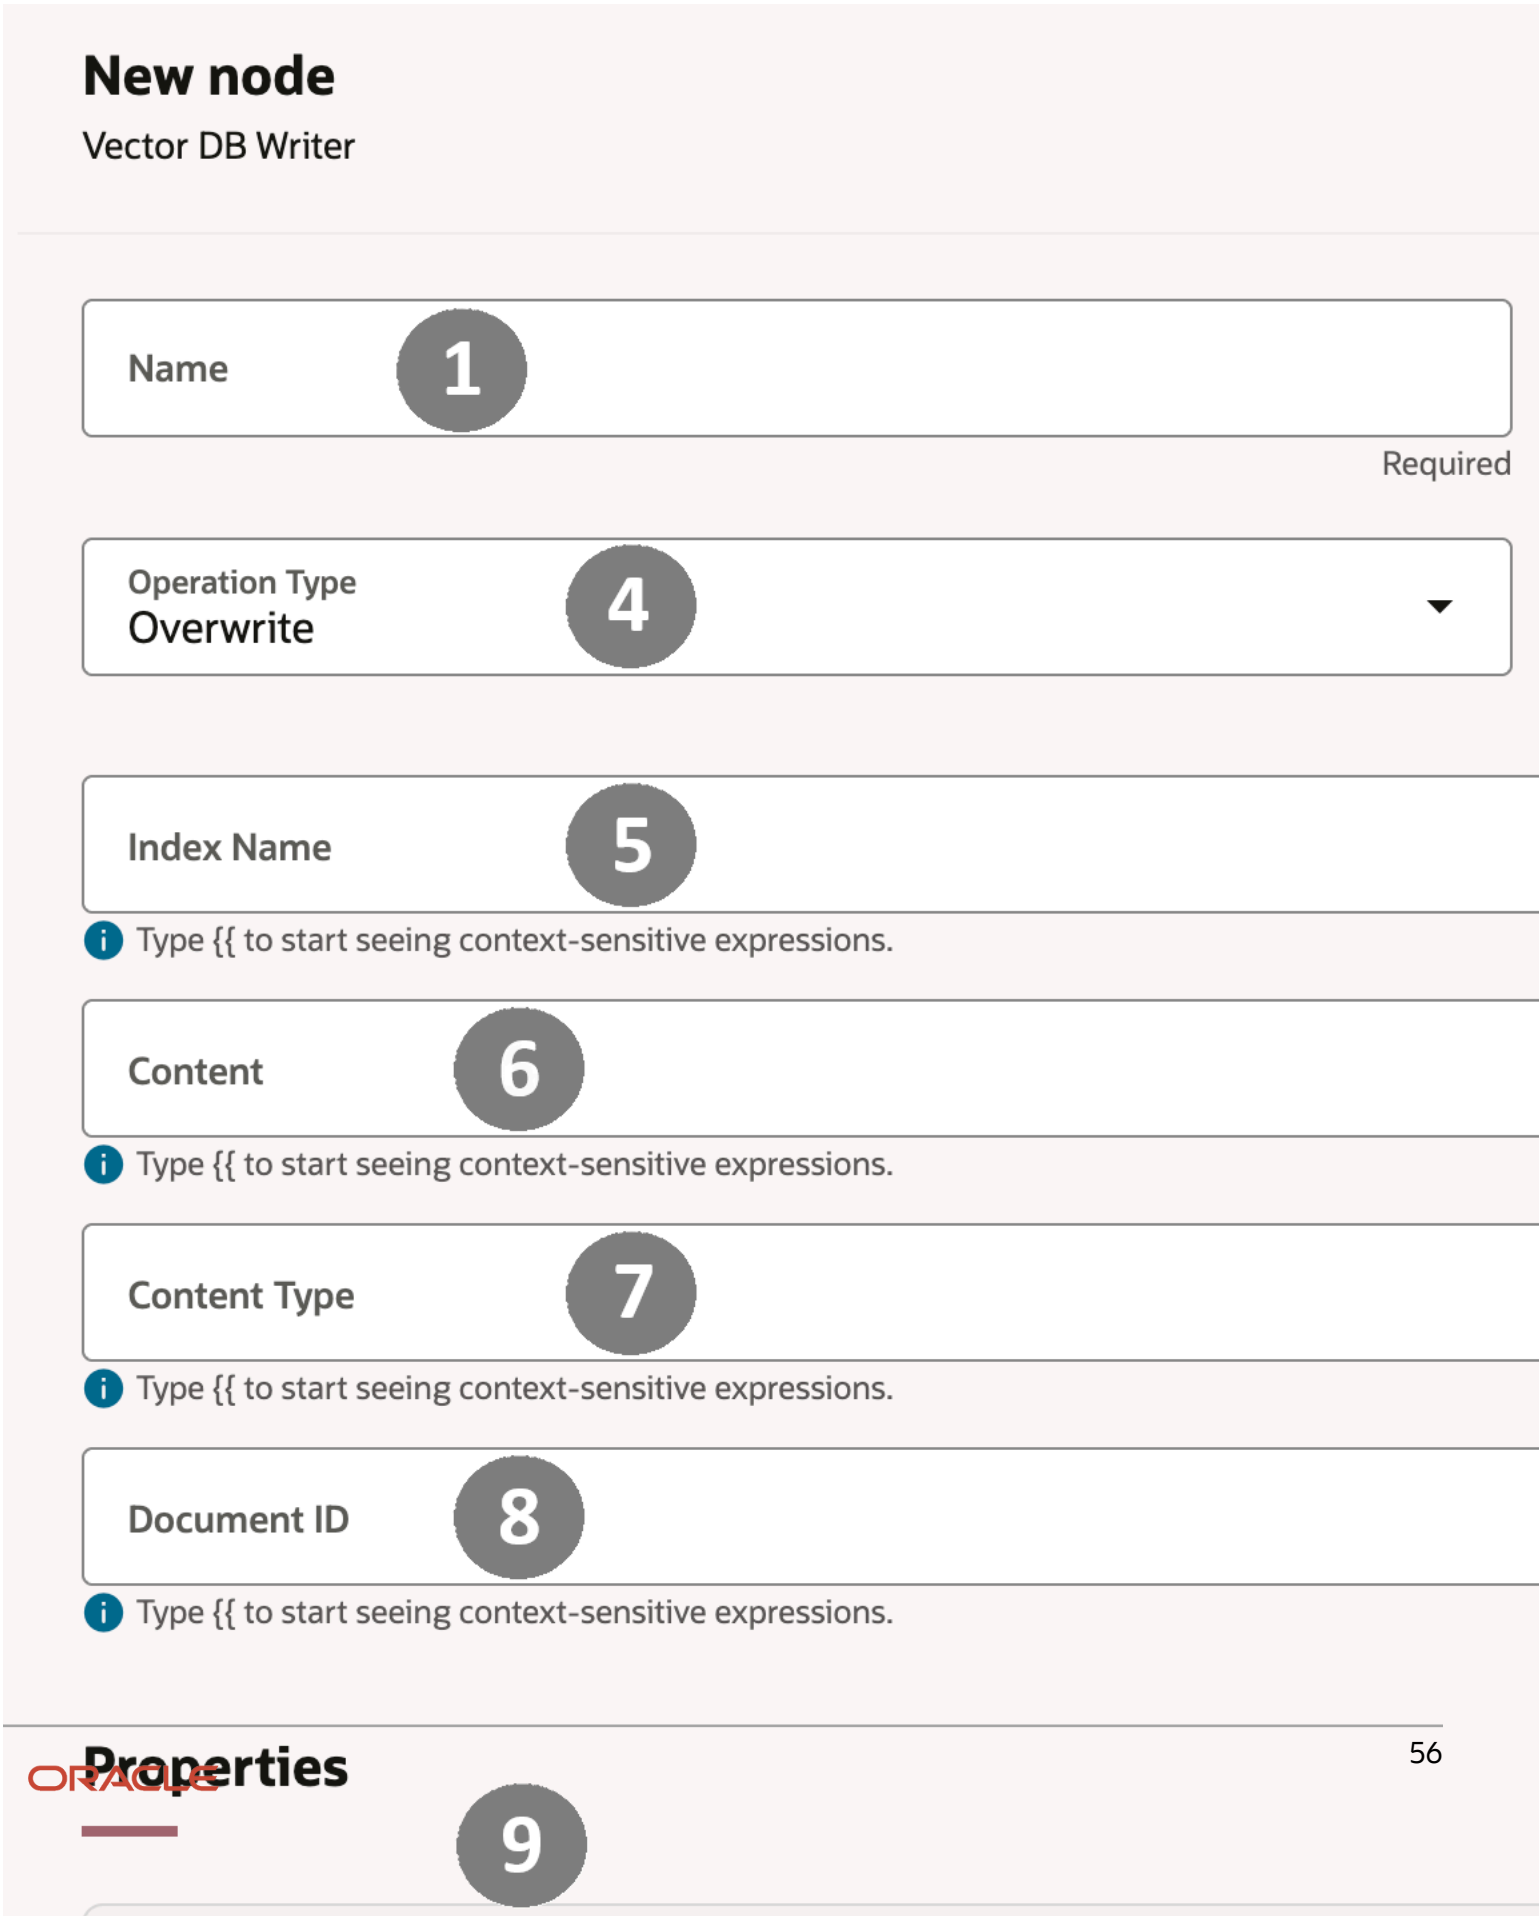
New node
Vector DB Writer
Name
1
Required
Operation Type
Overwrite
4
Index Name
5
Type {{ to start seeing context-sensitive expressions.
Content
6
Type {{ to start seeing context-sensitive expressions.
Content Type
7
Type {{ to start seeing context-sensitive expressions.
Document ID
8
Type {{ to start seeing context-sensitive expressions.
Properties
9

Callout Number	Field	Description	Example
1	Name	Design-time name for the node	Use descriptive names, like `WriteResolutionSummaryToVectorDB`, so the workflow is easy to understand.
2	Code	Programmatic identifier in the workflow schema	This field is auto-generated but can be changed to a user-defined one, such as `write_resolution_vector`. Use lowercase and underscores.
3	Error Handler	Fallback path on failure	Select an error branch or a dedicated handler node to manage failures cleanly.
4	Operation Type	Method of how the document is written to the index	Use `INSERT` for new entries, `OVERWRITE` to replace content, `UPSERT` to safely update or create, or `DELETE` to remove entries.
5	Index Name	Name of the vector index to write to	Choose an existing index or specify a new one, such as `support_ticket_summaries` or `product_docs_index`.
6	Content	Textual data to embed	Summarize this in a clean and structured manner. Avoid raw logs. Instead, use LLM-generated summaries, extracted facts, or curated knowledge.
7	Content Type	Type of content being indexed or embedded	Typical values include `json` or `text`.
8	Document ID	Unique identifier for this record	Use stable identifiers, such as `ticket_1123` or `customer_450_profile`.
9	Properties	Additional optional metadata key-values	`{objectId:"a12345", region:"NA", severity:"High"}`

Best Practices for Building Vector Write Nodes

A vector store works best when you write clear, meaningful, metadata-rich knowledge while avoiding noise, duplication, and topic drift.

Best Practice	Description
Store only public, reusable knowledge	Exclude sensitive, permission-controlled, or private information, for example, personally identifiable information (PII), credentials, confidential documents, or any information requiring permission checks.
Focus on durable, high-value content	Store information that will have long-term value for future workflows, such as case resolutions, structured summaries, or validated insights. Avoid writing noisy, one-off, or ephemeral content, like raw chat logs or temporary instructions.
Clean and normalize before writing	Always standardize content, summarize documents, deduplicate overlaps, remove irrelevant details, and attach key metadata.

Best Practice	Description
Keep data up to date	Vector writes must reflect the latest truth. Stale entries decrease accuracy and lead to incorrect answers. • Set auto-refresh triggers to update vectors whenever business objects change (for example, case closures, policy updates, PO resolutions). • Use OVERWRITE or UPSERT to keep summaries up to date. • Delete outdated or unused entries to prevent stale data from accumulating. • Maintain freshness so workflows rely on accurate, up-to-date knowledge.
Follow smart update practices	Update vectors in a controlled and intentional manner. Don't create new entries for each update, or store conflicting or outdated data. • Avoid creating new entries for every update. Refine existing data instead of duplicating. • Use UPSERT to update and improve existing summaries. • Use OVERWRITE for significant changes or new, higher-quality content. • Prune redundant or conflicting entries to keep the vector store clean. • Establish clear versioning strategies for policies, product versions, and configuration changes to prevent fragmentation.
Always include rich metadata	Tag with business object IDs, product numbers, version information, and other searchable attributes. Use a consistent metadata schema across agents and indexes.
Prevent data pollution	Check for existing knowledge to avoid duplicates and inconsistent tags. Regularly remove low-value content.
Exclude business object data	Duplicating this data in vector memory creates noise, redundancy, and version drift.

Vector Read Node

Vector read nodes retrieve the most relevant knowledge using semantic similarity and metadata filters.

Here's an overview of what the node creation window looks like. When creating a node, all values can be configured as expressions that will be resolved at runtime.

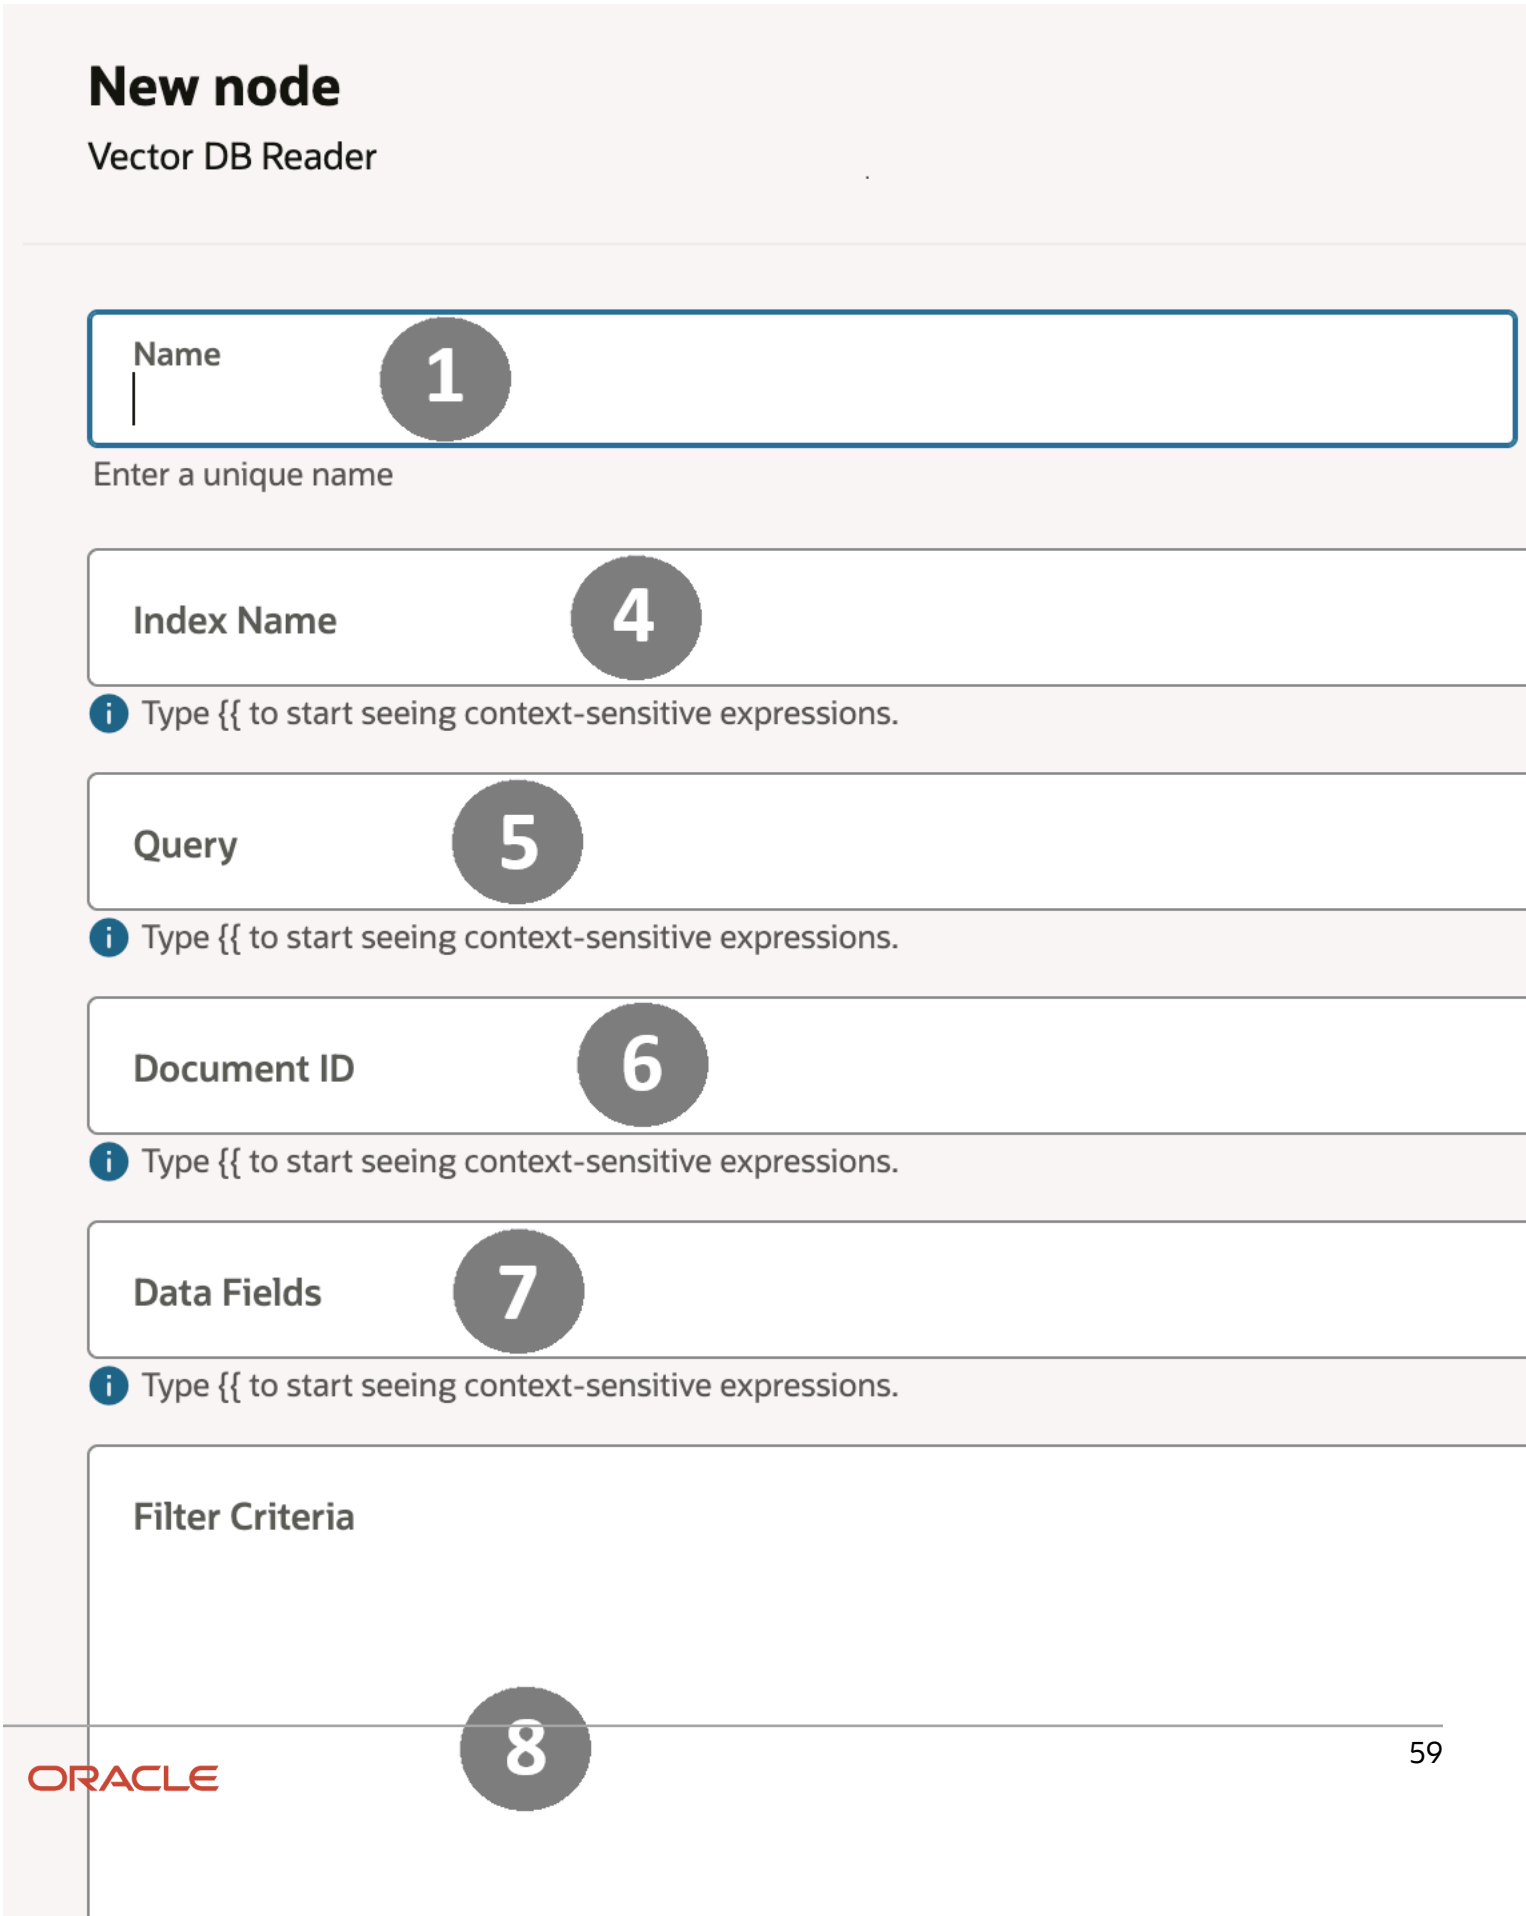
New node
Vector DB Reader
Name
1
Enter a unique name
Index Name
4
Type {{ to start seeing context-sensitive expressions.
Query
5
Type {{ to start seeing context-sensitive expressions.
Document ID
6
Type {{ to start seeing context-sensitive expressions.
Data Fields
7
Type {{ to start seeing context-sensitive expressions.
Filter Criteria
8

Callout Number	Field	Description	Example
1	Name	Name of the node	Use clear names like `RetrieveTicketContextFromVectorDB` to keep the workflow readable.
2	Code	Internal programmatic identifier	This field is auto-generated but can be changed to a user-defined one, like `read_ticket_context_vector`.
3	Error Handler	Defines what occurs on failure	Route to an error handler or fallback logic to avoid empty or invalid agent responses.
4	Index Name	Vector index to search	Choose the same index used by the writer, such as `support_ticket_summaries` or `employee_profile_index`.
5	Query	Natural language search query	Avoid vague queries and ask for intent-specific info, such as, “What troubleshooting steps were taken?”
6	Document ID	Unique identifier for this record.	Specify this when you want details tied to a known record, such as `ticket_12345_summary`.
7	Data Fields	Metadata fields to return	Specify an array of strings that will be used in the filter.
8	Filter Criteria	Logical filters applied before ranking	Specify this field to enable high precision and to constrain results from a retrieval, such as `product = Payroll`, `region = US`, or `severity >= High`.
9	Maximum Results	Maximum number of ranked results returned	Specify an integer value for the max number of results to be displayed.

Best Practices for Building Vector Read Nodes

A vector read node works best when you write clear, specific queries that leverage metadata filters to ensure semantically accurate responses, while preventing topic drift and hallucinations.

Best Practice	Description
Write clear, intent-driven queries	Ensure queries are specific and aligned with the exact workflow or agent task. For example, use precise prompts like "What resolved similar issues?" instead of vague ones like "Help me with this."
Use metadata filters for precision	Apply filters (for example, entity type, business object ID, and product number) to target only relevant content.
Validate retrieved results before using them	Check metadata and context before trusting or using results in workflows. This validation prevents incorrect answers, silent workflow failures, and LLM hallucinations grounded in bad evidence.

Best Practice	Description
Set an appropriate `maxResults` value	Set the `maxResults`field to return three to five, or more results. Don't depend on a single query result.
Add graceful fallback logic	Implement backup steps, such as business objects or API lookups, to handle cases where queries return no results.
Avoid broad or irrelevant retrieval	Scope queries to just what's needed. Skip irrelevant or noisy indexes to boost performance and accuracy of retrieved results.

6 Monitor and Evaluate AI Agents

Monitor and Evaluate AI Agents

Monitor and gain insights into how your AI agents are performing, and also evaluate the agents for accuracy. You can also track the interactions with your agents, understand real-world usage patterns, identify common errors, and measure overall performance.

- Monitoring: Monitoring tracks performance and provides insight on how your agents behave in production. Monitor agents to ensure that your quality bars for response time and token counts are maintained over time. You can also see any errors logged here.
- Evaluation: Evaluate agents before you deploy them, to ensure that they're ready for production. Test your agents for response correctness, response time, and token usage to meet your quality standards. You can also check the quality of answers generated through the document tool to assess how effectively agents utilize the retrieved context from the retrieval-augmented generation (RAG) metrics. After making any changes to your agent, or after a model update, rerun evaluations to confirm that your agent continues to perform as expected. This proactive approach helps you maintain high-quality experiences for your users.

This table summarizes some key metrics, their descriptions, and their availability for monitoring or evaluation.

Metric	Description	Available to Evaluate	Available to Monitor
Error Rate	Percentage of user sessions that ended in an error.	Yes	Yes
Error Count	Total number of errors recorded.	Yes	Yes
Session Count	Total number of conversations initiated with agents.	Yes	Yes
P99 Latency	The maximum wait time in milliseconds for 99% of users, revealing any areas where you should review and optimize the prompts or structure of the agent.	Yes	Yes
P50 Latency	The maximum wait time in milliseconds for 50% of users, helping identify performance issues. You can view this metric in the details of the monitoring or evaluation results.	Yes	Yes

Metric	Description	Available to Evaluate	Available to Monitor
Total Tokens	Cumulative number of tokens used by all agents.	Yes	Yes
Input Token Count	Total tokens sent to the LLM for requests. This includes system prompts, user messages, retrieved or context data, chat history, and tool or function definitions.	No	Yes
Output Token Count	Total tokens generated by the LLM for requests sent to it.	Yes	Yes
Median Correctness	The 50th percentile of correctness scores across evaluation runs. Each score (0–1) is computed by comparing the agent's answer to the reference answer provided in the evaluation set.	Yes	No
Session Count	The number of unique conversational sessions between a user and an AI agent. One session can include multiple messages or evaluation runs.	Yes	Yes
Groundedness	Alignment of the generated answer with the retrieved source content, indicating how faithfully the response reflects the supporting information.	Yes	No
Answer Relevance	Degree to which the answer directly addresses the user's question, measuring how fully and precisely the content covers the required subject matter.	Yes	No
Context Relevance	Quality and appropriateness of the retrieved information, assessing whether it meets the necessary standards to be considered reliable evidence.	Yes	No

Prerequisites

Aggregate the metrics that are displayed on the Monitoring and Evaluation tab in AI Agent Studio.

1. Go to **Navigator** > **Tools** > **Scheduled Processes**.
2. In Scheduled Processes, click **Schedule New Process**.
3. Leave the type as **Job**.

4. Search for and select **Aggregate AI Agent Usage and Metrics**.
5. Run the **Aggregate AI Agent Usage and Metrics** scheduled process.

   You can schedule this process to run on a recurring basis, for example, once a day.

The process aggregates the metrics that are displayed in the Monitoring and Evaluation tab of AI Agent Studio.

Monitor Agents

1. Go to **Navigator** > **Tools** > **AI Agent Studio**.
2. From the Monitoring and Evaluation tab, go to the Monitoring tab.
   - The Monitoring subtab displays aggregated metrics of all the agent runs over the selected time frame.
   - All supervisor agent runs are monitored, including agents in draft status.
3. Select the agent to view additional details.

   Each row represents a single session and displays the number of turns (back-and-forth messages), the session's completion status (successful or error), and the number of tokens used.
4. Select any session in the list to open the detailed trace view. This view displays a step-by-step timeline of the entire conversation, showing exactly which tools were called, the duration of each step, and the metrics for each step.

Evaluate Agents

Use evaluation sets to assess your agents' performance. An evaluation set contains one or more test questions, the expected agent responses, and the metrics to be measured. Evaluation sets are specific to each agent, and an agent can have multiple evaluation sets.

1. Create an evaluation set for an agent.
    a. Go to **AI Agent Studio**.
    b. Select the Monitoring and Evaluation tab.

       All evaluations run on the agents are displayed in the Evaluation tab.
    c. To create an evaluation set, click **Manage Evaluations** and select  .
    d. Enter a name, code, and description for the evaluation set, and select the agent team to be evaluated.
    e. Choose the run mode.

       Sequential: Runs the questions in the exact order you define them. Use this if one question depends on the context of the previous one.

       Random: Runs the questions in a random order.
    f. Select **Enable document tool evaluation** if you want to view RAG metrics in the evaluation report.
    g. From the Questions tab, add common questions that users are likely to ask the agent and the answers you would like the agent to deliver. Ensure both the questions and answers are concise, user-friendly, and reflect best practices.

       You can either add your questions and the expected answers individually, or upload a CSV file with the questions in the first column and the expected answers in the second column.
    h. From the **Metrics** tab, edit each metric to set the pass and fail criteria. For example, to indicate that the test fails if the correctness score is less than 0.7, choose **<** as the threshold condition and enter 0.7 as the threshold value.
    i. Select **Create** to save the evaluation set.
2. Run the evaluation set.
    a. On the Manage Evaluations page, select the **Initiate Evaluation Run** action for your evaluation set.
    b. Choose the version of the agent team to evaluate and run the evaluation.
3. Analyze the results.
    a. Click the evaluation set to view the Evaluation Runs page.
    b. Select the evaluation run and select the **View Run Results** action.

Tab	Information Displayed
Response performance	- Comparison of the expected response versus the actual response from the agent for each question, along with the metrics for each question in the evaluation. - Trace provides information about the detailed timeline for each question in the evaluation.
Correctness	Detailed breakdown of the correctness score. The LLM provides an initial score and feedback, and you can add your own feedback for record keeping in the **Correctness Score by Human** column.

Compare Evaluation Runs

You can see a side-by-side comparison of two different runs of the same evaluation, and easily spot regressions or improvements in latency, correctness, and token usage. Doing so, you can understand how an agent's performance changes over time, especially after you've made modifications.

1. From the Evaluation tab, select the evaluation.
2. Select any two runs and click **Compare**.
    - The Summary tab displays a high-level overview of the performance differences between the runs.
    - The Details tab provides a granular, question-by-question breakdown of the runs. For each question in the evaluation set, you can directly compare the actual response from Run 1 against the actual response from Run 2. You can also compare the specific latency, tokens used, and trace links for each question, making it easy to pinpoint exactly where and why performance or accuracy has changed.

7 Migrate RAG Agents to AI Agent Studio

Migrate Document Tools of RAG Agents

If you've created any RAG agents in Oracle Fusion Cloud Applications previously, we recommend that you replace your existing agent with an agent you create in AI Agent Studio.

You can migrate the Document tool you created for your existing agent to AI Agent Studio.

1. Go to the **Configure RAG Agents** work area.

   > **Tip:** You can use global search to get to the work area.

2. From the Tools tab, click  to migrate your tool.
3. Enter a unique tool name and code. That will help you easily find your migrated tool in AI Agent Studio.

After migrating your tool, create a new agent in AI Agent Studio using an appropriate template, add your migrated tool, and publish your agent. After testing and verifying your new agent, you can delete the original agent in the **Configure RAG Agents** work area.

8 Promote Published Agents

Migrate AI Agents from One Instance to Another

You can migrate an individual agent team from one environment to another using the import and export options available in AI Agent Studio. For example, you can export from the test or staging environment and import into the production environment.

After importing the agent team, you must publish it in the new environment to make it available for use. Also, if any of your agents use the following tools, make sure you complete the corresponding post-import tasks for each tool:

- Document tool: Upload and republish the document to ensure its content is correctly processed in the new environment.
- Email tool: Reconfigure the email tool to make sure necessary alerts are generated properly in the new environment.
- External REST tool: Reconfigure authentication settings for external REST API tools.

To migrate an agent team:

1. From the source environment, go to **AI Agent Studio**.
2. From the Agent Teams tab, locate the agent team you want to export and click the icon.
3. Enter the export details and click **Export** to download the JSON file to your local system.
4. From the target environment, go to **AI Agent Studio** and select the Agents tab.
5. Click **Import** and upload the previously downloaded JSON file.
6. Use the **Draft** button to edit the imported agent team.
7. Publish the agent team.

9 Create Channels in AI Agent Studio

Create Microsoft Teams Channel for AI Agent Studio

You can integrate Microsoft Teams with AI Agent Studio, and enable users to access Oracle AI agents from Microsoft Teams. The integration preserves user identity and data security across Oracle Fusion Cloud Applications and Microsoft accounts. After integration, users can select **Ask Oracle** app in Microsoft Teams and interact with the AI agents available to them.

Integration with Microsoft Teams involves configuring several Microsoft 365 services, including Microsoft Entra (Azure Active Directory), Microsoft Bot Framework, and the Teams App Manifest. These components are closely interdependent and must be configured correctly across environments and tenants. Proper coordination and version management are essential for identity flows, manifest packaging, channel registrations, permission scoping, and endpoint exposure.

Here are the steps to set up Ask Oracle in Microsoft Teams:

1. Register Azure bot in Microsoft Entra (Azure Active Directory). Ensure to save the bot ID, password and tenant ID after registration.
2. Create a manifest file for Microsoft Teams in AI Agent Studio, and download it.
3. Upload the manifest file downloaded from AI Agent Studio in Microsoft Teams.

Register Azure Bot in Microsoft Entra

1. As an Azure/Exchange administrator, sign in to the *Microsoft Azure portal*.
2. Click **Create a resource** and search for **Azure Bot**.
3. Create the bot by adding these details:
    - Enter a unique bot handle.
    - Choose your active subscription.
    - Select an existing resource group or create one. For **Type of App**, select **Single Tenant**.
    - Select **Create new Microsoft App ID** as **Creation Type**.
4. Click next and create the bot.
5. After deployment is complete, click **Go to resource**.
6. Enter the URL of the bot service in **Messaging endpoint**. The URL will be of this pattern:

    ```
    https://<pod name>/api/fusion-ai/orchestrator/collabTools/v1/teams-message
    ```

    Replace <pod name> with the base URL of your Oracle application environment.

    > **Note:** Make a note of the App ID and Tenant ID because you'll need them when creating the manifest file in AI Agent Studio.

7. Generate a client secret:
    a. Click **Manage Password** next to **Microsoft App ID** and create a new client secret.
    b. Provide a description, for example, Oracle Agents Secret, and select a desired expiry date.

> **Note:** Make a note of the client secret because you'll need it when creating the manifest file in AI Agent Studio.

8. Configure Application Permissions:
    a. From the left panel, open **API Permissions**.
    b. Go to **Add Permission** > **Microsoft Graph** > **Application Permissions**.
    c. In the **User** folder, search for **User.ReadBasic.All** and add it.
    d. Click **Grant admin consent** for the added permission.
9. Configure Delegate Permissions:
    a. From the left panel, open **API Permissions**.
    b. Go to **Add Permission** > **Microsoft Graph** > **Delegate Permissions**.
    c. In the **User** folder, search for **User.ReadBasic.All** and add it.
    d. Click **Grant admin consent** for the added permission.
10. Enable Microsoft Teams Channel:
    a. Go back to your bot resource, and open **Channels** from the left panel.
    b. Select the **Microsoft Teams** channel, agree to the terms and conditions, and click **Apply**.

Create Manifest File in AI Agent Studio

Prerequisite: Make sure that your role has the necessary permission groups. For more information, see Assign Permission Groups to Use Channels in Credentials in *Access Requirements for AI Agent Studio*.

1. In **AI Agent Studio**, go to the Credentials tab, and then open the Channels tab.
2. Configure **Microsoft Teams**.
3. Enter details for the channel manifest file.
    - Short Manifest Name: Enter the short name for the app as displayed in Microsoft Teams. The default value is **Ask Oracle**.
    - Full Manifest Name: Enter the full name for the app as displayed in Microsoft Teams. For example, Ask Oracle Fusion Agent.
    - Manifest Code: Enter a unique ID to be used in AI Agent Studio (not part of the Microsoft Teams schema). For example, b3e4a289-7cbb-43d3-9a4c-d52bc874e830
    - Version: Select the manifest schema version.
    - Manifest Schema Version: Select the Microsoft schema version that's to be used for the manifest file.
    - Oracle Instance URL: Enter the URL for the Oracle tenant for Fusion Applications.

      Format is

      ```
      https://<fully qualified pod name>:<optional port>
      ```

      Create separate manifests for each tenant.
    - Bot ID or App ID: Enter the App ID that was generated during the Azure bot setup.
    - Tenant: Enter the Tenant ID that was generated during Azure bot setup.
    - Secret: Enter the client secret that was generated during the Azure bot setup.
    - Advanced Settings: List of valid domains or websites the bot can access or redirect to. Include all API or identity-related endpoints.

      > **Note:** You must add token.botframework.com with the fully qualified pod name.
4. Create the manifest file.

5. Select **Edit** and download the manifest file, to be used in Microsoft Teams.

Upload Bot Artifacts in Microsoft Teams Console

1. As an administrator, go to *Microsoft Teams Admin Center*.
2. Open **Manage Apps** and select **Upload new app** from the **Actions** menu.
3. Upload the manifest file downloaded from AI Agent Studio.
4. Publish the app. Your app now appears in the list of available apps.
5. In **Manage Apps**, go to **Setup Policies** > **Org Wide Policy** and add the app you published.

Your app is now available to all users included in Org Wide Policy.

Create Slack Channel for AI Agent Studio

You can integrate Slack with AI Agent Studio, and enable users to access Oracle AI agents from Slack. After integration, users in the workspace can select the **Ask Oracle** app in Slack and interact with the AI agents available to them.

The setup process involves these high-level steps:

1. Create the Slack channel manifest in AI Agent Studio.
2. Create the Ask Oracle app in *api.slack.com/apps* and add the manifest file.
3. Update the manifest file created in AI Agent Studio and add the credentials to the Bot configuration section.
4. Add the Ask Oracle app in Slack.

Create Manifest File in AI Agent Studio

Before you start, make sure your role has the necessary permission groups. For more information, see **Assign Permission Groups to Use Channels** in *Access Requirements for AI Agent Studio*.

1. In AI Agent Studio, go to the Credentials tab, and then open the Channels tab.
2. Configure Slack.
3. Enter details for the channel manifest file.
    - Manifest Short Name: Enter the short name for the app. The default value is Ask Oracle.
    - Manifest Full Name: Enter the full name for the app. For example, Ask Oracle Fusion Agent.
    - Version: Enter the manifest file version.
    - Manifest Schema Version: Select the schema version that's to be used for the manifest file.
4. Create the manifest file.
5. Copy the manifest details from the Manifest.json section. You'll need this information when creating the app on *api.slack.com/apps*.

Create Ask Oracle App in Slack

1. Go to *api.slack.com/apps* and select **Create New App**.
2. Select **From a manifest**, choose the workspace, and click **Next**.
3. In the JSON tab, paste the manifest details copied from the manifest file created in AI Agent Studio.

4. Click **Next** and create the app.
5. Select the app you created and go to **Settings** > **Basic Information**.
6. Note down these credentials for later use:
    - App ID
    - Client ID
    - Client Secret
    - Signing Secret
7. Go to **Settings** > **Install App** and install the app.
8. Make a note of the **Bot User OAuth Token** value.

Update Manifest File in AI Agent Studio

1. In AI Agent Studio, go to the Credentials tab, and then open the Channels tab.
2. Edit the Slack manifest file.
3. In the Bot configuration section, enter the credentials you saved from the Slack application setup.
    - Bot ID or App ID
    - Client ID
    - Secret
    - Token
    - Signing Secret
4. Update the manifest file.

Add the Ask Oracle App in Slack

1. Open the Slack application and select **Add app**.
2. Add Ask Oracle.

Users in the workspace can now interact with the Oracle AI agents they've access to.

10 Allow External Access to Fusion Applications Agents

Enable Applications to Access Fusion Applications Agents

You can access the agents of Oracle Fusion Cloud Applications from other Oracle applications and from external applications, using the `/invokeAsync` API.

Access from Other Oracle Applications

To use `/invokeAsync` API, you must authenticate through Oracle Cloud Infrastructure Identity and Access Management (OCI IAM) using an OAuth 2.0 bearer token. For agent operations to function, integration must be bidirectional. Both systems must securely access each other's services.

As an administrator, you create the confidential application in Fusion Applications IAM, representing the other Oracle application service. Then create a matching confidential application in the other application service, representing the AI Agent Studio's service. These two applications form the trusted link for the two-way communication.

Set Up Two-Way Trust Between Applications

1. Create a confidential application in Fusion Applications IAM:

   This application acts as a representation of the external Oracle application in Fusion Applications IAM. After creating the confidential application, enable access to the AI Agent Studio service so the application can call the `/invokeAsync` API exposed by this service.
   - Identify the identity domain (authorization server) where you'll create your confidential application.
   - Specify OAuth grant types. For example, a 2-legged OAuth.
   - Define the access scope.
   - Generate the access token and use it to make REST API calls.

   For more information, see *Configure OAuth Using the Fusion Applications Identity Domain*.
2. Create a confidential application in the external Oracle application's IAM:

   Similar to the previous confidential application you created, create another confidential application that acts as a representation of the Fusion Application service in the external Oracle application's IAM. After creating the confidential application, configure it to allow AI Agent Studio to call the appropriate service endpoints.

   For more information, see *Configure OAuth Using the Fusion Applications Identity Domain*.
3. Register the confidential application created in the external Oracle application's IAM (from step 2) in AI Agent Studio:
   a. Go to **AI Agent Studio** and open the Credentials tab.
   b. Go to the Data Source Applications tab and add the details of the external confidential application.
      - Base URL: Enter the base URL of the external API

- IDCS URL: Enter the URL of the Fusion Applications IAM, to enable AI Agent Studio to get the OAuth token
- Scope: Specify what access needs to be given whenever that OAuth token is retrieved.
- Public and Private keys: Enter the keys for signing authentication requests. Used by IDCS to verify that the request is from a trusted and authorized source.

4. Create a business object to access the external Oracle application:
    a. In **AI Agent Studio**, open the Business Object tab.
    b. Add a new business object with the **Resource Type** as **Other Data Source Application**.
    c. From the **Oracle Data Source Application** list, select the data source application that you've registered in step 3.

Access Flow During Runtime

At runtime, the external Oracle application obtains an access token from Fusion Applications using the configured credentials. The application then uses this token to call the `/invokeAsync` API in AI Agent Studio, enabling access to the agent team.

This illustration shows the interactions between Warehouse Management System (WMS), which is an Oracle application outside Fusion Applications, Fusion Applications IAM, AI Agent Studio, and WMS IAM.

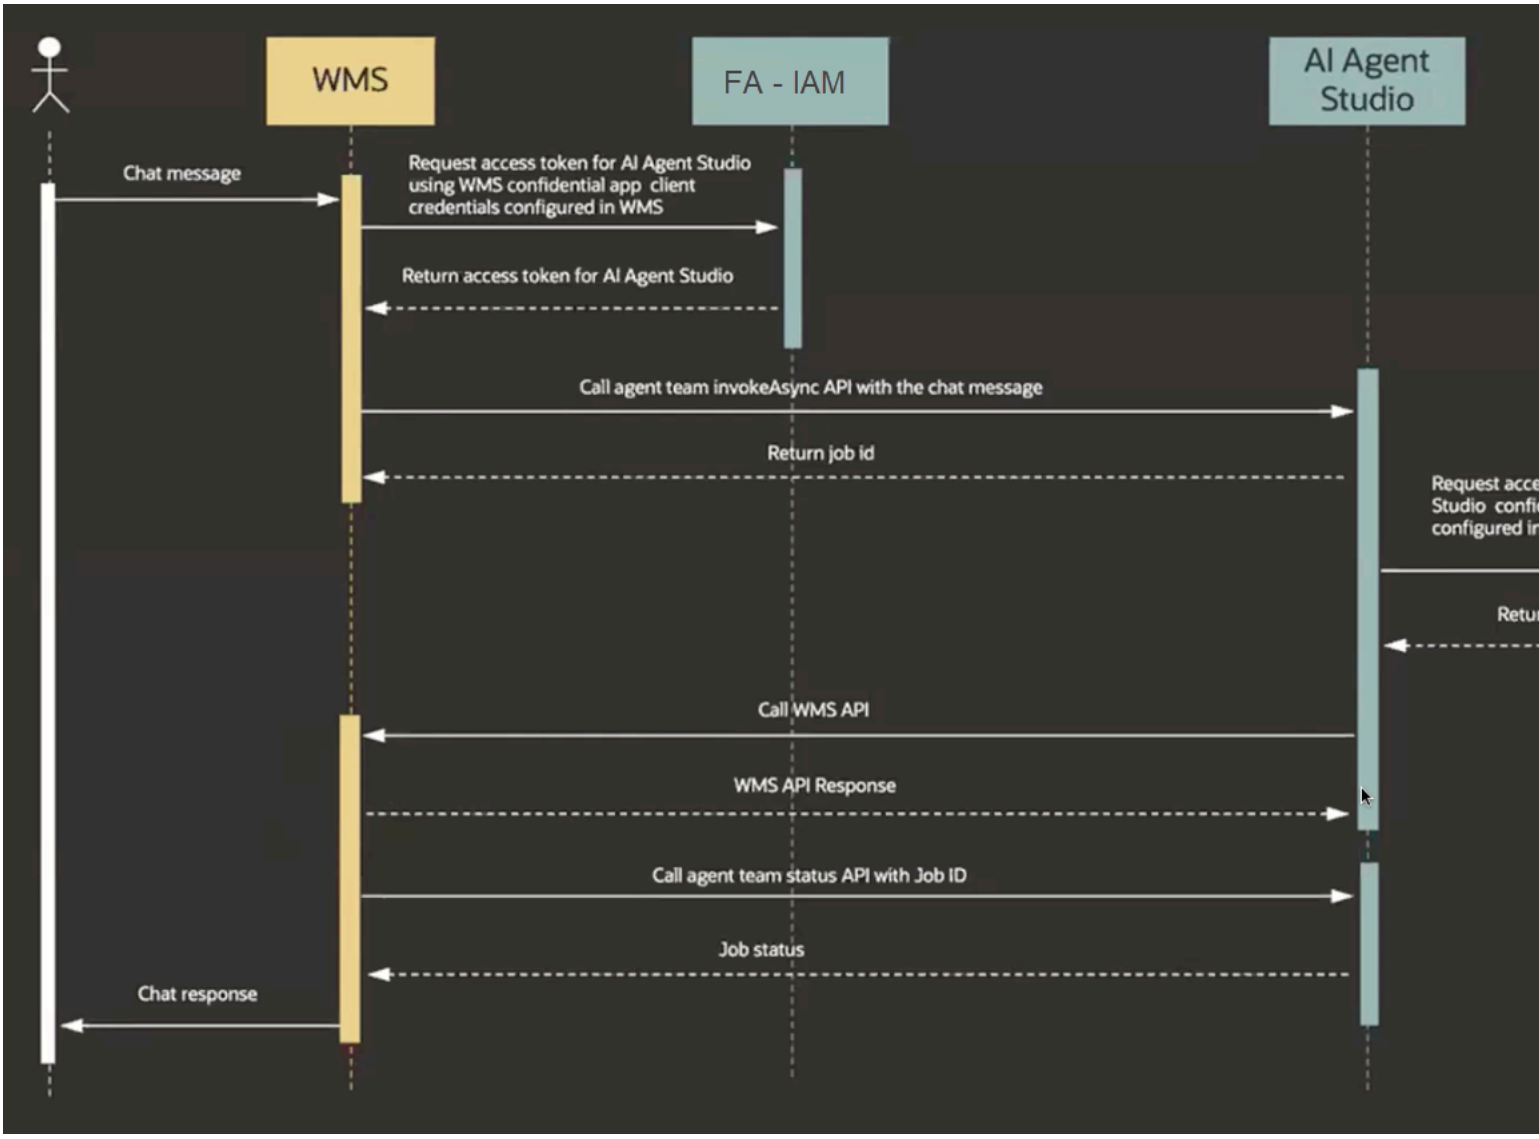

- The user interacts with the WMS application and sends a chat message.
- WMS requests an access token from Fusion Applications IAM using its configured credentials.
- WMS calls the `/invokeAsync` API in AI Agent Studio, passing the user's chat message and access token.
- AI Agent Studio returns a Job ID to WMS, which WMS uses to track progress and get the response.
- To call WMS API, AI Agent Studio requests an access token from WMS IAM.
- AI Agent Studio calls the WMS API using the access token.
- WMS returns the API response and requests for the job status.
- AI Agent Studio returns the job status.
- Final chat response is processed and delivered to the user.

Access from External Applications

Use `/invokeAsync` API to access the agents of Fusion Applications. For more information, see *Agent Team REST Endpoints*.

11 AI Agent Studio FAQs

When do I need a Custom AI Agent Subscription for AI Agent Studio?

AI Agent Studio is included with your Oracle Fusion SaaS Cloud subscription and comes with a variety of ready-to-use agent templates at no additional cost. You may use and make minor changes to these templates, such as uploading documents, adjusting display fields, or editing prompts, without requiring a Custom AI Agent subscription. However, a subscription is required if you want to create entirely new agents, significantly modify existing templates (for example, by adding new integrations, actions, or changing the agent's overall purpose), use third-party or marketplace agents, or select large language models (LLMs) that aren't provided by Oracle. Additional fees may also apply if your usage of premium LLMs exceeds your default token allocation. For more information, see *Fusion Custom AI Agents - Pricing FAQ (KB864473)*.

Which large language models (LLMs) are currently used or supported for AI agents?

Oracle Fusion Cloud Applications support several LLMs from providers such as, Cohere, Meta, and OpenAI, for powering AI features including content authoring, summarization, and interactive Q&A. For the most current list of supported models, see *What Large Language Models (LLM) are in use or supported by Fusion SaaS Applications? (KB796229)*.

Can I use any agent in my agent team or workflow?

No, you can't use just any agent in your agent team or workflow. You can use predefined agents and reuse agents created from the Agents tab in AI Agent Studio in other agent teams and workflows. However, if you create an agent directly within an agent team or workflow, that agent can't be reused and will only be available in the specific agent team or workflow where it was created.

When should I use a single-agent flow vs a multiagent process?

- Use a single-agent process for straightforward business requirements, such as retrieving answers from a single policy document.

- Opt for a multiagent process when dealing with more complex scenarios. For example, when multiple documents need to be parsed, or when the answers vary based on the user asking the question. In these cases, multiple agents are used, and the user's question is directed to the appropriate worker agent for processing.

Can I edit a preconfigured agent team?

You can't directly modify the components in a preconfigured agent team template. However, you can create copies of individual components (such as topics), edit them, and then add the edited versions to your agent team after removing the original components. For example, to edit a topic that's part of a preconfigured agent team template you're using, do these steps.

1. Go to **AI Agent Studio**, and select the preconfigured agent team template you need, either using **Copy Template** or **Use Template**.
2. From the selected template, remove the topic you want to edit.
3. Go to the Topics tab, and search for the topic you removed.
4. Click  and create a copy of this topic.
5. Edit the copied topic as needed and add it back to your agent team.
6. Save and publish your agent team.

How do I make agents respond faster?

AI agent response time is directly linked to the combined length of input and output text, measured in tokens for each transaction. For example, 100 tokens is roughly 75 words, though this ratio varies by model, writing style, and language. By being deliberate with your instructions, you ensure that your agents reason clearly and avoid unnecessary processing.

Here are some recommendations to help your agents respond faster.

- Edit summarization prompts to include only essential instructions relevant to your specific use case. The default prompts for supervisor and workflow agents are broadly written to cover many scenarios, so remove any general or redundant sections to streamline processing and improve response speed.
- Minimize the use of input and output tokens with concise prompts. Use retrieval-augmented generation (RAG) instead of overloading the context window, and set specific output length limits.
- Specify response length to avoid overly verbose answers.
- Use smaller, specialized agents that work in parallel and cache static instructions.
- Include only relevant information. Remove unnecessary details from context and documents to streamline processing.
- Consider a multi-agent approach, where specialized agents handle distinct tasks, such as research, coding, or Q&A, instead of loading all information into a single prompt.

By carefully selecting relevant context and breaking complex tasks into smaller, focused agents, you can improve both the speed and efficiency of your AI agents.

How do I track the usage of AI Agents created in AI Agent Studio?

You can use the Monitoring and Evaluation tab within the AI Agent Studio to track real-time interactions, analyze usage patterns, and identify errors to improve agent accuracy and performance. For more information, see *Monitor and Evaluate AI Agents*.

How can I best reduce the risk of my agent showing hallucinations?

Large language models (LLMs) can unintentionally introduce extra details when they lack sufficient information.

To mitigate this issue, consider these strategies:

- Supply the necessary knowledge: Integrate relevant sources such as business objects, REST APIs, or RAG documents, ensuring the LLM has access to accurate and up-to-date information.
- Use explicit prompt instructions: Clearly instruct the LLM not to add or infer information beyond what's provided, reducing the risk of embellishment or fabrication.

How do I hide or restrict access to AI Agent Studio for specific users?

To restrict a user's access to AI Agent Studio, remove any related custom roles from the user's account.

1. Go to the Security Console, open the Users tab and search for the specific user.
2. From the search results, select the user, and edit the user account details.
3. In the Roles section of the Edit User Account page, select the custom roles associated with AI Agent Studio, and remove them using the × icon. These roles were added to the user account when access was provided. For a list of these roles, see *Provide Access to Configure AI Agents in all Products*.
4. Save your changes.

Can I rename the AI chat label to something other than Ask Oracle, or change the design of the chat panel?

No, you can't change the background image or rename 'Ask Oracle' within the active chat window or AI Agent Studio preview.

What's the difference between editing a topic versus editing a system prompt?

When you update a topic, the changes are available across agents because you can reuse topics across agents. When you're updating a system prompt, the changes are available only for the agent being edited.

Tip: We recommend updating topics instead of prompts, so that you can reuse the topics.

What types of files can I upload as part of document tools in AI Agent Studio?

You can upload PDF, TXT, and Markdown files.

Why does the external REST tool only support HTTPS transactions?

The tool is limited to only HTTPS to maintain data privacy in transit.

Why is the Provider list in the LLM tab empty?

If the Provider list is empty when creating an agent team, it's likely due to missing permissions. To resolve this, make sure the required permission groups with the **AllRowsAllFields** security view are added to the custom role assigned to the user, and run the necessary scheduled processes to synchronize security data. For more information, see *LLM Provider LOV Empty when Creating a new Agent in Ai Agent Studio (KB863374)*.

Why can't I see all agent teams in the Agent Teams tab of AI Agent Studio?

That's because only published agent teams are displayed by default. Use the **Draft** button to view agents that aren't published yet.

Why is my agent not fetching data from the document added to it?

Check if you've marked your document as published. If not, publish your document. To do so, from the Tools tab in AI Agent Studio, set the status of your document as **Ready to Publish**, and then run the **Process Agent Documents** scheduled process. After the scheduled process completes, from the Tools tab, select your document, and set its status as **Published**.

Also remember, the best practice is to create one Document tool per logical document. For example:

- Java Resumes Tool - One logical document containing all Java resumes.
- Python Resumes Tool - One logical document containing all Python resumes.

Why is the Roles list in the Security tab empty?

That's probably because permission groups aren't enabled for your duty role. Security administrators can enable permission groups for the duty role assigned to you, in Security Console. For details, see *Access Requirements for AI Agent Studio*.