

Oracle AI for Fusion Applications

Questions and Answers



Oracle AI for Fusion Applications
Questions and Answers

G22516-14

Copyright © 2025, Oracle and/or its affiliates.

Author: Soniya Mathew

Contents

Get Help

i

1	Oracle AI for Fusion Applications Questions and Answers	1
	What's Oracle AI for Fusion Applications?	1
	Is Oracle AI in Fusion Applications the same as the OCI Generative AI offering?	1
	Can I enable or disable Fusion Applications features that use AI?	1
	Is there any operational impact if features with AI are enabled?	2
	What's the Service Level Agreement (SLA) and contingency plan for Oracle AI in Fusion Applications?	2
	What's the Oracle AI Platform for Fusion Applications?	2
	What are the technologies used in Oracle AI for Fusion Applications?	2
	What are the languages supported by Oracle AI for Fusion Applications?	3
	What's predictive AI?	3
	What's generative AI?	4
	What's a language-based agent?	4
	What's a context-aware agent?	4
	What's an action-taking agent?	4
	How are the models selected for use in Oracle AI for Fusion Applications?	5
	Is natural language processing (NLP) or a similar technology used when interacting with users?	5
	Does Oracle AI in Fusion Applications need calibration before using?	5
	How's unseen data, that wasn't used to calibrate or train the model, handled?	6
	How's user behavior, input, and feedback incorporated into the model?	6
	How is model drift handled?	6
	How do I override AI prompts for Oracle Fusion Cloud Applications?	7
	How do I resolve any issues in an AI prompt?	7
	What's a prompt?	8
	What's a prompt template?	8
	What's prompt engineering?	8
	What's the structure of a prompt?	8
	How do I create prompts for AI features?	9
	What are some common prompt techniques?	9
	What are some general best practices for creating prompts?	9
	Is there a central place where I can see and use all the AI agents I have access to?	10

How can I give users access to AI agents?

10

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 Support

You can get support at [My Oracle Support](#). For accessible support, visit [Oracle Accessibility Learning and Support](#).

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.

Learn About Accessibility

For information about Oracle's commitment to accessibility, visit the [Oracle Accessibility Program](#). Videos included in this guide are provided as a media alternative for text-based topics also available in this guide.

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.

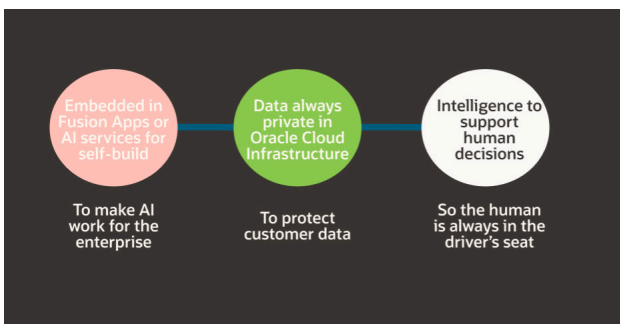
Thanks for helping us improve our user assistance!

1 Oracle AI for Fusion Applications Questions and Answers

What's Oracle AI for Fusion Applications?

Oracle AI for Fusion Applications refers to the AI functionality embedded within features in Oracle Fusion Cloud Applications. These features help you get business insights faster, automate routine tasks, access personalized recommendations, and improve business performance.

This image highlights the key focus areas of Oracle AI for Fusion Applications: enterprise integration, data privacy, and human-driven decision making.



The features are prebuilt to include AI and you don't need to have data science or machine learning expertise to use them. You can easily enable and set up these features, without any assistance from an external system integrator. Here's more information about the features that use [Oracle AI for Fusion Applications](#).

Is Oracle AI in Fusion Applications the same as the OCI Generative AI offering?

No, the Oracle AI Platform for Fusion Applications is built on top of OCI Generative AI. This platform provides additional functionality for features in Fusion Applications that use AI.

Can I enable or disable Fusion Applications features that use AI?

Yes. For generative AI, you just need to enable or disable features in the Setup and Maintenance work area and do some testing. For predictive AI, the model is trained on your data, so you also have some configuration steps.

Is there any operational impact if features with AI are enabled?

There's no significant impact. AI calculations are done in the AI platform and, for GenAI, the LLM – not within Oracle Fusion Cloud Applications.

What's the Service Level Agreement (SLA) and contingency plan for Oracle AI in Fusion Applications?

Oracle AI in Fusion Applications follows the same SLA as Oracle Fusion Cloud Applications. See the **Oracle Cloud Service Continuity Policy** section in the *Oracle Cloud Hosting and Delivery Policies* guide.

What's the Oracle AI Platform for Fusion Applications?

The Oracle AI Platform for Fusion Applications provides users with access to Oracle AI features, while ensuring superior performance and security.

Each organization gets their own securely partitioned instance of the platform within their local Oracle Cloud Infrastructure (OCI) region. This setup ensures high performance and scalability, while conforming to Oracle's stringent data security principles.

This platform is multitenant, which means that each organization's data and trained ML models (for predictive AI) are fully isolated. The platform accesses the pretrained generative AI models securely, and the input and output data from the models for the given use case is processed within a secure OCI environment, which might be outside of the local OCI region. There's no sharing or mixing of data between organizations, and your information won't be persisted or used to train models or to generate responses for others.

Finally, any AI-generated outcomes, predictions, or responses are securely sent from your platform instance to your applications using REST APIs, as part of your regular business workflows. The platform ensures that everything runs smoothly and securely from start to finish.

What are the technologies used in Oracle AI for Fusion Applications?

Oracle AI, in Fusion Applications, encompasses these types of AI technologies.

- *Predictive AI*
- *Generative AI*
 - *Language-based*
 - *Context-aware*
 - *Action-taking*

What are the languages supported by Oracle AI for Fusion Applications?

Language support might vary by release, product, use case, or feature. For example, only English is available for the Enterprise Performance Management features with AI.

Release	Languages Available
25B	English, Spanish, French, Portuguese, Arabic, German, Italian, Dutch, Japanese
25C, 25D	<p>You can interact with AI in the language you selected when signing in to Oracle Fusion Cloud Applications, or your preferred language as set in your preferences. However, the response language of the agent isn't yet automatically aligned with your sign-in or preferred language. Instead, the response language must be configured by the administrator during agent design, or you must explicitly request your desired language in your prompt.</p> <p>Note: For agent configuration, only English is supported.</p>

What's predictive AI?

Predictive AI identifies patterns in data and applies these insights to make predictions on new data.

It typically involves machine learning models tailored to specific use cases, such as data classification, forecasting, and anomaly detection. Predictive AI features use machine learning algorithms trained on each user's data, continuously learning and improving over time to deliver greater value and immediate benefits to users. Examples include summarizing a document, answering questions about your data, predictive cash forecasting, sales forecasting, project analysis, and so on.

What's generative AI?

Generative AI features use LLMs to create content, including text, images, videos, and even code.

Some common applications for text are assisted authoring, summarizing documents, and providing suggestions. For example, you can generate the reply for an email, and then review or edit the same before sending the reply.

Generative AI also includes various types of agents, which work together to help you create and interact with content in a dynamic and responsive way.

Type of agent	Description	Example
Language-based	Interacts with you using human languages.	Chatbots, virtual assistants.
Context-aware	Always on, sensing and reacting to data from your environment, even as the data changes.	If a user in your organization asks about their medical data, the agent can pull the relevant details directly from your organization's database.
Action-taking	Takes action automatically, with you guiding and approving in key moments.	If you request a vacation booking, the agent might suggest possible vacation dates based on your calendar, find the hotels with vacancy, and create a booking with your approval.

What's a language-based agent?

The language-based agent leverages the capabilities of LLMs to communicate with you, in a natural, human-like manner. For example, chatbots, virtual assistants, and so on.

What's a context-aware agent?

Context-aware agents answer questions based on your organization's data, which the LLM learns from. For example, if a user in your organization asks about their medical data, the agent can pull the relevant details directly from your organization's database.

What's an action-taking agent?

As the name suggests, an action-taking agent can do things on your behalf. These agents can autonomously complete tasks, use tools, and loop you in to review or approve. For example, if you request a vacation booking, the agent might suggest possible vacation dates based on your calendar, find the hotels with vacancy, and create a booking with your approval.

How are the models selected for use in Oracle AI for Fusion Applications?

An experienced data science team makes these decisions, which vary depending on the business problem at hand and the latest advancements in AI. But in general, the team considers these factors:

- Accuracy
- Coverage
- Processing intensity
- Response time
- Operational simplicity
- Ease of maintenance

Is natural language processing (NLP) or a similar technology used when interacting with users?

Yes, we use NLP, generative AI solutions, and machine learning tools to deliver responses to users. All responses are positioned as suggestions to keep users and operators fully in control.

Does Oracle AI in Fusion Applications need calibration before using?

Large language models (LLMs) for generative AI don't need to be trained further. For predictive AI, the base models are trained on your data to tailor them further for your use. These calibrated models are embedded into Oracle Fusion Cloud Applications, and they can be used without further training. Other than enabling the features with AI, usually you don't need to do additional implementation. At least, it would be very minimal.

How's unseen data, that wasn't used to calibrate or train the model, handled?

Fusion Applications features with AI are pretrained. They're robust without overfitting, and our algorithms gracefully handle unseen data. Some features learn from changes to your data so that the model can make more relevant suggestions.

How's user behavior, input, and feedback incorporated into the model?

Machine learning models incorporate user behavior and input by ingesting new or changed data to train the model every day. By continuously retraining, the model learns from use and delivers best results.

Note: Your organization's feedback data is isolated, and not incorporated into shared models.

Oracle also gathers telemetry data about product usage to incorporate user behavior and input, for example:

- We evaluate telemetry metrics to determine the data features and model configurations needed to optimize outputs. Telemetry data is also used to identify patterns that might lead to model or feature enhancements.
- Telemetry data is also used to identify changes in user behavior. We adapt model weightings to optimize outputs, based on the real-world data patterns that evolving user behavior displays.
- Telemetry data is also used to promote or relegate some outputs. For example, if a recommendation is liked, accepted, or used, it may be given a positive score, and if not, a negative one. We also monitor indicators that are more passive than clicking a Like button. For example, the number of changes to an autogenerated response, or the time spent on the task used to derive an implicit like or dislike.

The large language model (LLM) that Oracle Fusion Applications Cloud uses is a foundational model that's not further trained or fine tuned. But, we do monitor the LLM's performance by assessing the LLM's output and storing a score. We also record whether the user accepts or changes the LLM's output, which indicates the quality of the LLM's response from the user's perspective. We evaluate the model's performance when considering model upgrades.

How is model drift handled?

Model drift refers to the decline in the accuracy of a machine learning model over time. Oracle detects and mitigates model drift by monitoring model performance metrics, comparing model outputs with ground truth (verified data),

or human evaluation, and tracking evaluation metrics over time. To maintain performance, we also retrain models regularly with these practices:

- Updating the model with new data
- Reevaluating underlying assumptions
- Fine-tuning model parameters
- Applying transfer learning
- Adapting to evolving data patterns

Note: The exact processes vary according to product and feature.

How do I override AI prompts for Oracle Fusion Cloud Applications?

Use AI Configurator in Oracle HCM Experience Design Studio to override and enhance the default AI prompts. This configurator is available for all products in Fusion Applications.

Note: You must have the Access HCM Page Configurator (HRC_ACCESS_HCM_TRANSACTION_CONFIGURATOR_PRIV) privilege to access the design studio.

1. *Create and activate a sandbox* that has the HCM Experience Design Studio tool in it.
2. On the sandbox bar before the global header, select **HCM Experience Design Studio** from the **Tools** menu.
3. Open **AI Configurator**.
4. Click **Override** for a seeded prompt, and edit the prompt. Here are a few things to know about prompts:
 - Lines beginning with multiple # signs are section headings.
 - Prompt variables are inside curly braces, for example, {variable_name}. They're integral to the prompts function.
 - Some prompts require responses to be in something that's not plain text, such as JSON or HTML. These instructions are essential to the business flow the prompt is embedded in. Changing the output or the output structure can cause potential issues in the business flow.
5. Preview and test your changes. After you save, you get a new overridden prompt, and the original seeded prompt remains as is.
6. *Publish the sandbox*.

When users access the page or process that triggers this prompt in the application, the overridden prompt will be processed instead of the seeded prompt. If you later delete the overridden prompt, the seeded prompt will be used.

How do I resolve any issues in an AI prompt?

Check if the prompt has been overridden. If so, delete the override and see if the issue is still there.

See [How do I override AI prompts for Oracle Fusion Cloud Applications?](#)

What's a prompt?

A prompt is a natural language text request sent to a generative AI model to get a response.

What's a prompt template?

Prompt templates convert user input and parameters into instructions for a language model, enabling it to generate responses. This approach helps to guide the model's response, allowing it to understand the context and produce relevant, coherent language-based output.

What's prompt engineering?

Prompt engineering is an iterative process of trial and error that creates prompts to generate the most relevant, accurate, and useful responses from large language models (LLMs). In prompt engineering, you continuously refine prompts until you get the desired outcomes from the AI model. It's a blend of art and science, focused on getting an LLM to perform as expected and intended.

What's the structure of a prompt?

Use these common elements to get the best results when you create your prompts.

Element	Description	Example
Instruction	The basic task you're asking the model to perform.	What do you want the model to do?
Context	Any extra information to help obtain a desired response.	A specific scenario, role, or situation.
Input data	Details of the task that reinforce the instruction.	Any extra data or variables.
Output format	The format of the desired response.	How do you want the response to look?

Tip: A prompt doesn't need to have all these elements, but responses can change significantly when these elements are added to or removed from the prompt.

How do I create prompts for AI features?

Use AI Agent Studio to create agentic workflows consisting of one or more agents. You can add prompts to the components (such as topics and tools) in your agents.

What are some common prompt techniques?

Technique	When to Use
Zero-shot prompting	When an LLM has enough existing knowledge to respond accurately without requiring any prior coaching. As training methods for LLMs improve, particularly with techniques like instruction tuning and reinforcement learning from human feedback (RLHF), these models can often generate desired responses without the need for any examples.
Few-shot prompting	When zero-shot prompting fails, you can provide examples or demonstrations directly in the prompt to guide the model.
Chain of thought (CoT) prompting	When an LLM is tasked with complex reasoning or problem-solving. This technique enhances the accuracy of the responses, increases transparency, and improves the LLM's ability to reason by encouraging it to provide a step-by-step explanation of its thought process. It reveals the model's reasoning to address situations where the model might give an incorrect answer and fail to recognize its mistake.
Prompt chaining	<p>When the LLM needs to do tasks with multiple stages, it's more efficient or exact to break the process down into smaller, manageable chunks. Unlike CoT, which is more focused on incremental reasoning, prompt chaining involves separating tasks into distinct steps, each of which can be handled by a separate prompt.</p> <p>This method links two prompts together. Instead of instructing the model to think step-by-step, you first ask it to analyze the situation and then use that analysis as input for a second prompt to obtain a final, straightforward answer.</p>

What are some general best practices for creating prompts?

Best Practice	Description
Start simple and iterate	Prompt engineering is an iterative process that starts with a simple prompt, evaluates the AI's response, and then adds more details to get better results.
Experiment with the elements of prompt structures	Responses are highly sensitive to small changes in prompt structure. To improve your results, experiment with these simple structural elements to see how they affect the model's response: <ul style="list-style-type: none">• Word order: Changing the order of words can generate different responses.• Line spacing: Adding or removing line breaks can affect the output.• Placement: Information placed later in the prompt might hold more weight with some models.
Use detailed commands	Prompts are effective when you use commands to instruct the model for the exact task you're asking the model to achieve, such as Write , Classify , Summarize , Translate , Order , and so on.
Be specific	Specify the instruction and task you want the model to perform. Descriptive and detailed prompts usually generate better results. Being specific is crucial when aiming for a particular result or desired style of generating responses.
Be mindful of length	Aim for a balance between simplicity and detail. Ensure all information is relevant to the task and avoid unnecessary details.
Specify the output	Specify the exact output format needed for Oracle Fusion Cloud Applications.
Avoid using personally identifiable information (PII)	PII shouldn't be used for privacy and security reasons.
Test your prompts	Experiment with various phrasings or versions of the same prompt. Include both <i>happy</i> and <i>sad</i> case data (inputs designed to challenge the model) to ensure it stays on topic.

Is there a central place where I can see and use all the AI agents I have access to?

Explore the available agents from the AI Agents page. To open this page, add `/fscmUI/redwood/human-resources/ai-studio/agent-explore` to the end of the URL for your application environment.

For example, <https://example.com/fscmUI/redwood/human-resources/ai-studio/agent-explore>.

How can I give users access to AI agents?


You can provide users access to explore and interact with AI agents.


The following instructions apply to these products:

- Oracle Fusion Cloud Human Capital Management
- Oracle Fusion Cloud Procurement
- Oracle Fusion Cloud Sales

- Oracle Fusion Cloud Service
- Oracle Fusion Cloud Supply Chain & Manufacturing - For users without the SCM job roles


Here's what you do:

1. Set the profile options to enable external application integration for Security Console and the VBCS Progressive Web Application user interface.
 - a. In the Setup and Maintenance work area, search for the **Manage Administrator Profile Values** task using the search link in the  panel.
 - b. Search for the **Enable Security Console External Application Integration** (ORA_ASE_SAS_INTEGRATION_ENABLED) profile option and set the value for the **Site** profile level to **Yes**.
 - c. Search for the **Enable VBCS Progressive Web Application User Interface** (ORA_HCM_VBCS_PWA_ENABLED) profile option and set the value for the **Site** profile level to **Y**.
2. Create a custom role.
 - a. Go to **Navigator > Tools > Security Console**, and create a new custom job role.


Note: Make sure to enable permission groups.
 - b. On the Function Security Policies page, click **Add Function Security Policy** and add the Access Intelligent Agent Chat (HRC_ACCESS_AI_AGENT_CHAT_PRIV) privilege.
 - c. On the Role Hierarchy page, open the **Roles and Permission Groups** tab and add the Fai Genai Agent Runtime Duty (ORA_DR_FAI_GENERATIVE_AI_AGENT_RUNTIME_DUTY) duty role.
 - d. Save the custom role.
3. Assign the custom role to users.
 - a. In the Security Console work area, open the **Users** tab and select the user you want to assign the custom role to.
 - b. On the User Account Details page, edit the user account and add the role.
4. Associate the custom role with an agent team.
 - a. Go to **Navigator > Tools > AI Agent Studio**.
 - b. Open the **Agent Teams** tab and edit the agent team that you want to give users access to.
 - c. Click  and open the **Security** tab.
 - d. Add the custom role.
 - e. Apply your changes, and update and publish your agent team.

Users with the custom role can now access the agents associated with that role.

Provide Access to the SCM Agents for Users with SCM Job Roles

1. Set the **Enable VBCS Progressive Web Application User Interface** (ORA_HCM_VBCS_PWA_ENABLED) profile option.
 - a. In the Setup and Maintenance work area, search for the **Manage Administrator Profile Values** task using the search link in the  panel.
 - b. Search for the profile option and set the value for the **Site** profile level to **Y**.
2. In the Security Console work area, search for any of these roles to edit:
 - ORA_CJM_CHANNEL_CLAIMS_MANAGER_JOB
 - ORA_CJM_CUSTOMER_CHANNEL_PROGRAM_MANAGER_JOB
 - ORA_CMF_FISCAL_DOCUMENT_SPECIALIST_JOB
 - ORA_CMF_RECEIVING_SPECIALIST_JOB

- ORA_CMK_B2B_ADMINISTRATOR_ABSTRACT
- ORA_CMK_TRADING_PARTNER_B2B_ADMINISTRATOR_ABSTRACT
- ORA_CSE_ASSET_ADMINISTRATOR_JOB
- ORA_CST_COST_ACCOUNTANT_JOB
- ORA_DOO_ORDER_ADMINISTRATOR_JOB
- ORA_DOO_ORDER_MANAGER_JOB
- ORA_DOS_SUPPLY_CHAIN_OPERATIONS_MANAGER_JOB
- ORA_EGI_PRODUCT_DATA_STEWARD_JOB
- ORA_EGP_PRODUCT_MANAGER_JOB
- ORA_ENQ_PRODUCT_RECALL_MANAGER_JOB
- ORA_ENQ_QUALITY_ENGINEER_JOB
- ORA_FOM_ORDER_ENTRY_SPECIALIST_JOB
- ORA_FOS_SUPPLY_CHAIN_CONTROLLER_JOB
- ORA_INV_HEALTHCARE_INVENTORY_SPECIALIST_JOB
- ORA_INV_INVENTORY_MANAGER_JOB
- ORA_INV_WAREHOUSE_MANAGER_JOB
- ORA_INV_WAREHOUSE_OPERATOR_JOB
- ORA_MNT_MAINTENANCE_MANAGER_JOB
- ORA_MNT_MAINTENANCE_TECHNICIAN_JOB
- ORA_MSC_BACKLOG_MANAGER
- ORA_MSC_DEMAND_PLANNER
- ORA_MSC_DEMAND_AND_SUPPLY_PLANNER
- ORA_MSC_MATERIALS_PLANNER
- ORA_MSC_ORDER_PROMISING_MANAGER
- ORA_MSC_PRODUCTION_SCHEDULER
- ORA_MSC_REPLENISHMENT_PLANNER
- ORA_MSC_SALES_AND_OPERATIONS_PLANNER
- ORA_MSC_SUPPLY_CHAIN_PLANNER
- ORA_MSC_SUPPLY_CHAIN_PLANNING_APPLICATION_ADMINISTRATOR
- ORA_QP_PRICING_ADMINISTRATOR_JOB
- ORA_QP_PRICING_ANALYST_JOB
- ORA_QP_PRICING_MANAGER_JOB
- ORA_RCL_DEPOT_REPAIR_MANAGER_JOB
- ORA_RCL_FIELD_SERVICE_ADMINISTRATOR_JOB
- ORA_RCS_SUPPLY_CHAIN_APPLICATION_ADMINISTRATOR_JOB
- ORA_RCV_RECEIVING_AGENT_JOB
- ORA_SCH_PURCHASE_PRICING_MANAGER_JOB

- ORA_VCS_SUPPLY_CHAIN_COLLABORATION_PLANNER_JOB
 - ORA_WIE_PRODUCTION_OPERATOR_JOB
 - ORA_WIE_PRODUCTION_SUPERVISOR_JOB
 - ORA_WIS_MANUFACTURING_ENGINEER_JOB
 - ORA_WSH_SHIPPING_AGENT_JOB
 - ORA_WSH_SHIPPING_MANAGER_JOB
 - a. Enable permission groups.
3. Associate the edited job roles with the agents.
- a. Go to **Navigator > Tools > AI Agent Studio**.
 - b. Open the **Agent Teams** tab and edit the required agent team.
 - c. Click  and open the **Security** tab.
 - d. Add the job roles.
 - e. Apply your changes, and update and publish your agent team.

Users with any of the edited job roles can now access the agents associated with that role.

