



**MYTHICS™**

**Lets build a AI based Knowledge Base**

**Erik Benner, VP Transformation , Mythics LLC**



# Agenda

Introductions  
What is a KB?  
Building the KB  
Q&A

Erik Benner  
VP Enterprise Transformation  
ebenner@mythics.com



@erik\_benner   
TalesFromTheDatacenter.com  
blog.mythics.com/posts



- ✓ 1st Cloud ACE Director
- ✓ Published Author
- ✓ Linux since 1992
- ✓ DB 12c ab 23 BETA
- ✓ Prelaunch ODA “comet”
- ✓ First Version of Oracle...7 in 1994
- ✓ OEM 13c since Product Launch
- ✓ Giving back to the community
  - ✓ OATUG EM for Apps SIG co-chair
  - ✓ OEM/OMC CAB Member
  - ✓ Quest Oracle Cloud SIG Leader
  - ✓ OATUG EM SIG Leader
  - ✓ OATUG Committee
  - ✓ ECO Conference Committee



One Company for ANY Oracle Need



**MYTHICS™**





# What is a KB?

What is a KB?

# What is a Knowledge Base?

A Knowledge Base (KB) is a centralized repository of structured and unstructured information that is used to store, organize, and retrieve knowledge. It helps users find answers, troubleshoot issues, and make informed decisions.

The challenge traditional KBs have, is finding the right answer to the question. Often the right answer comes from more than one source!

# Why build a Knowledge Base?

- Improves efficiency by reducing repetitive questions.
- Enhances customer support with self-service options.
- Facilitates knowledge sharing within organizations.
- Reduces operational costs by minimizing human intervention.

# Accessing the Knowledge Base?

OCI Chat Interface – great for troubleshooting and testing

REST API

- APEX Apps
- OCI Digital Assistant
- HTML5 embedded into existing apps

# Building the KB

Currently Requires the Chicago Region

You will need an Object Storage Bucket for the source files

Optionally, you can use OCI OpenSearch for a Oracle 23ai Database using a Vector Search.

A KB can have only a single type!

# Step 1 – What files can I use?

Files should be .pdf or .txt files.

You can easily convert web pages to .txt using w3m

```
w3m -dump https://talesfromthedatacenter.com/2025/01/lets-do-some-image-magic/ > /tmp/image.txt
```

For the demo, I dropped in a pdf file into a prebuild bucket.

# Navigate to Generative AI Agents

The screenshot shows the Oracle Cloud navigation interface. At the top, there is a search bar with the text "Search resources, services, documentation, and Marketplace" and a region selector set to "US Midwest (Chicago)". On the left, a vertical navigation menu lists various categories, with "Analytics & AI" highlighted. The main content area displays a grid of service categories: Analytics, Data Lake, Messaging, Machine Learning, AI Services, and Media Services. Under the "AI Services" category, "Generative AI Agents" is listed and highlighted with a red arrow.

Cloud

Search resources, services, documentation, and Marketplace

US Midwest (Chicago)

Search

Home

Compute

Storage

Networking

Oracle Database

Databases

**Analytics & AI**

Developer Services

Identity & Security

**Analytics & AI**

**Analytics**

Analytics Cloud

**Data Lake**

Big Data Service

Data Catalog

Data Integration

Data Flow

**Messaging**

Streaming

Connector Hub

**Machine Learning**

Data Science

Data Labeling

**AI Services**

Generative AI

Language

Speech

Vision

Document Understanding

Digital Assistant

Generative AI Agents

**Media Services**

Media Flow

Media Flow Jobs

Media Streams

# Create Agent

The screenshot shows the Oracle Cloud console interface for managing Generative AI Agents. The page title is "Agents in cloudfab Compartment". A red arrow points to the "Create agent" button. The page includes a search bar, navigation tabs (Overview, Agents, Knowledge Bases, Chat), and filter sections for Compartment (cloudlab) and State (Any state). The main content area shows a table with columns for Name, Lifecycle state, and Time created, but it is currently empty with the message "No items found." and "Showing 0 items < 1 of 1 >".

Generative AI Agents

Agents *in cloudfab* Compartment

Generative AI agents connect to your data sources, retrieve data, and augment model responses with relevant, grounded information to generate more accurate and relevant responses. [Learn more about working with agents.](#)




Create agent

Name	Lifecycle state	Time created
No items found.		

Showing 0 items < 1 of 1 >



# Basic Info

☰ Cloud  ▼ US Midwest (Chicago)   


## Create agent

- 1 Basic information
- 2 Add tool
- 3 Setup agent endpoint
- 4 Review and create

### Basic information

Create an agent and optionally add a knowledge base to connect the agent to. You can also add a knowledge base to the agent after you create the agent.

Name

Compartment    
benneraced2022 (root)/cloudlab

Description *Optional*

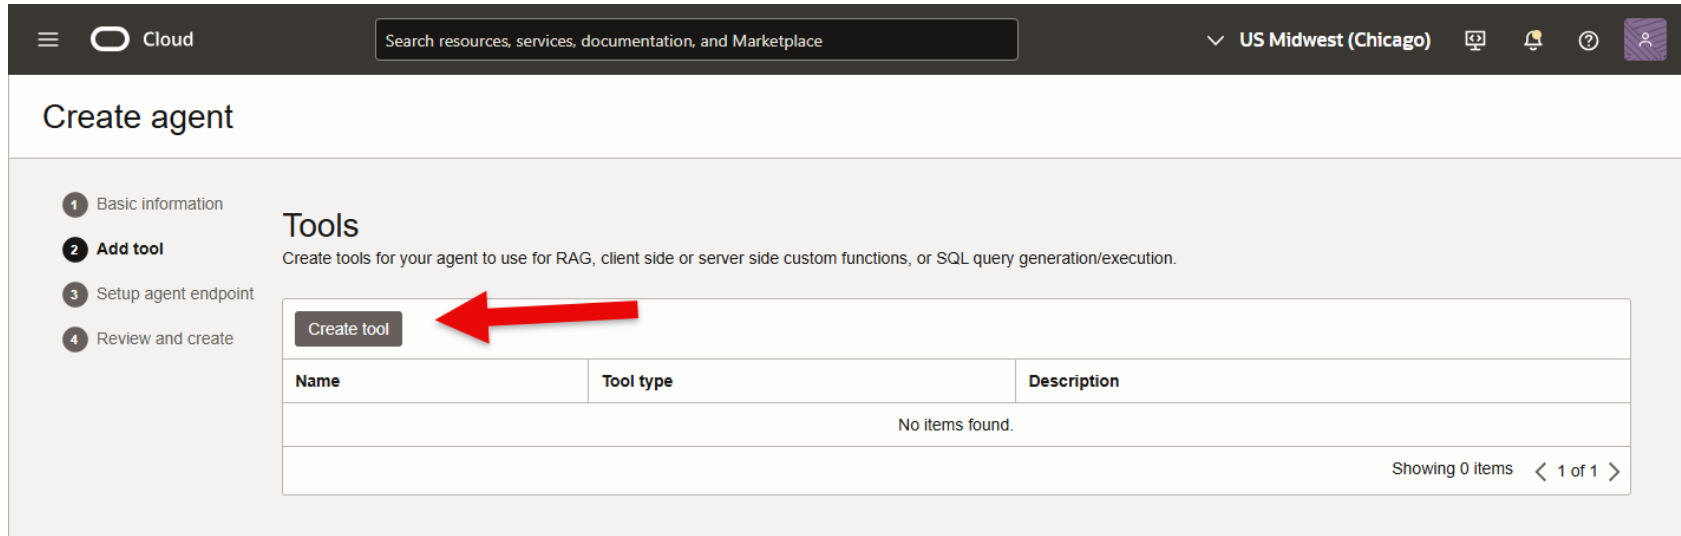
Welcome message *Optional* ⓘ

Routing instructions *Optional* ⓘ

[Show tagging options](#)



# We need a tool



The screenshot shows the 'Create agent' page in the Cloud console. The page has a dark header with a search bar and navigation icons. The main content area is titled 'Create agent' and contains a sidebar with four steps: 1. Basic information, 2. Add tool (highlighted), 3. Setup agent endpoint, and 4. Review and create. The 'Tools' section is active, displaying a 'Create tool' button with a red arrow pointing to it. Below the button is a table with columns for Name, Tool type, and Description. The table is currently empty, showing 'No items found.' and 'Showing 0 items < 1 of 1 >'.

Cloud

Search resources, services, documentation, and Marketplace

US Midwest (Chicago)

## Create agent

- 1 Basic information
- 2 Add tool**
- 3 Setup agent endpoint
- 4 Review and create

### Tools

Create tools for your agent to use for RAG, client side or server side custom functions, or SQL query generation/execution.

**Create tool**

Name	Tool type	Description
No items found.		

Showing 0 items < 1 of 1 >




# More than one Tool!

The screenshot displays a user interface for creating tools. At the top, there is a navigation bar with a hamburger menu, a 'Cloud' label, a search bar containing 'Search resources, services, documentation, and Marketplace', and a region selector set to 'US Midwest (Chicago)'. Below the navigation bar, the main heading is 'Create tool'. Underneath, a sub-heading reads 'Add prebuilt tools or define custom tools for your agent'. Three tool options are presented in a grid:

- RAG**: Retrieval-Augmented Generation combines retrieval of information from knowledge bases with text generation for more accurate and contextually relevant responses.
- SQL**: Converts natural language queries into SQL statements for seamless database interactions.
- Custom tool**: Define tools for function calls (client execution) or API endpoints (agent execution).

A red arrow points from the top right towards the RAG option. At the bottom left of the tool options, there is a link that says 'Show tagging options'.


# RAG KB

Cloud  US Midwest (Chicago)   

## Create tool

### RAG Configuration

Name

Description 

Custom instructions

### Add knowledge bases

Compartment

benneraced2022 (root)/cloudlab

<input type="checkbox"/>	Name	OCID	Lifecycle state	Time updated
No items found.				

0 selected Showing 0 items < Page 1 >

Terms of Use and Privacy [Cookie Preferences](#) Copyright © 2025, Oracle and/or its affiliates. All rights reserved. [What's this?](#)  Redwood preview



# Feed me Data

documentation, and Marketplace US Midwest (Chicago) [Icons]

## New knowledge base

Create a knowledge base that either imports content into a retrieval-augmented generation (RAG) system or uses your own data store.

Name

Compartment

benneraced2022 (root)/cloudlab

Description *Optional*

Data store type

Enable hybrid search ⓘ

### Data sources

Name	Type	Bucket name
No items found.		

Showing 0 items < 1 of 1 >

# It's all in the bukcket

documentation, and Marketplace

US Midwest (Chicago)

## Specify data source

**i** Information ingested into the knowledge base is securely stored and used as reference context when the Agent generates responses. The ingestion process does not attempt to identify, filter, or remove personal information from the input.

Name  
Sample Data for the KB

Description *Optional*

Enable multi-modal parsing **i**

### Data bucket

Select bucket in **cloudlab** [\(Change connection\)](#)

AI\_Agent\_Bucket

Select directories and/or individual files inside the Object Storage bucket name specified above. Content indexing is supported for the following file types: PDF, TXT

Select all in bucket **i**

Object prefixes

Time updated

<input type="checkbox"/> Drivers_Licensing_Handbook_web.pdf	Wed, 02 Apr 2025 03:34:16 GMT
---	-------------------------------

0 selected Showing 1 item < Page 1 >

[Add object prefixes manually.](#)

Create Cancel

Copyright © 2025, Oracle and/or its affiliates. All rights reserved. [What's this?](#) | Redwood preview



# Create the tool

Cloud Search resources, services, documentation, and Marketplace US Midwest (Chicago)

**Create tool**

Demo RAG

Description ⓘ

Custom instructions

Custom instruction is a guideline message that gives clear, specific task directions and can change overall chat behaviour and conversation style. Example: answer in a professional tone

**Add knowledge bases**

Compartment

cloudlab

benneraced2022 (root)/cloudlab

Create knowledge base

<input type="checkbox"/>	Name	OCID	Lifecycle state	Time updated
<input type="checkbox"/>	<a href="#">Demo KB</a>	...pgz7gmoaza <a href="#">Show</a> <a href="#">Copy</a>	● Creating	Wed, 02 Apr 2025 03:35:16 GMT

0 selected Showing 1 item < Page 1 >

1 Knowledge bases selected, Maximum 5 allowed *Read-only*

Selected knowledge bases are: Demo KB

[Show tagging options](#)

**Create tool**

Successfully initiated ingestion job creation



# Next, up... endpoints

## Create agent

- 1 Basic information
- 2 **Add tool**
- 3 Setup agent endpoint
- 4 Review and create

### Tools

Create tools for your agent to use for RAG, client side or server side custom functions, or SQL query generation/execution.

Create tool

Name	Tool type	Description	
Demo RAG	RAG_TOOL_CONFIG	Demo system	⋮

Showing 1 item < 1 of 1 >

# The devil is in the details, or lack there of!

Cloud Search resources, services, documentation, and Marketplace US Midwest (Chicago)

## Create agent

- 1 Basic information
- 2 Add tool
- 3 **Setup agent endpoint**
- 4 Review and create

### Setup agent endpoint

Configure settings for Human-in-the-loop and guardrails. Additional endpoint setup options can be configured on the endpoint creation page.

Automatically create an endpoint for this agent

Enable human in the loop ⓘ

Disabled

#### Guardrails

Specify guardrails for content moderation, prompt injection (PI) protection and personally identifiable information (PII) protection.

##### Content moderation

Choose whether to apply the content moderation feature to input and output. Content moderation aims to remove toxic or violent content from responses. It is recommended to use content moderation.

Input

<b>Disable</b> Keeps the input prompt as is, without aiming to remove any toxic or violent content from the input. ✓	<b>Block</b> Aims to detect and remove content that it finds toxic or violent from the input (recommended).	<b>Inform</b> Does not filter explicit content from the input, but aims to notify if it finds certain content to be toxic or violent from the input.
---	--	---

Output

<b>Disable</b> Outputs explicit content, without aiming to remove toxic or violent content from the response. ✓	<b>Block</b> Aims to detect and remove content that it finds toxic or violent from the response (recommended).	<b>Inform</b> Outputs explicit content and aims to notify if it finds toxic or violent content in the response.
--	---	--

Previous Next Cancel

Terms of Use and Privacy Cookie Preferences Copyright © 2025, Oracle and/or its affiliates. All rights reserved. What's this? | Redwood preview

# Lets protect folks with guard rails!

## **Content moderation**

Content moderation aims to remove toxic or violent content from responses. It is recommended to use content moderation.




## **Prompt injection (PI) protection**

Choose whether to help prevent malicious inputs from manipulating system behaviour/responses or not. It is recommended to use prompt injection protection.

## **Personally identifiable information (PII) protection**

Choose whether to help remove personally identifiable data or not. Disable for unrestricted data exposure.

# Double Check

Cloud  US Midwest (Chicago)   

## Create agent

- 1 Basic information
- 2 Add tool
- 3 Setup agent endpoint
- 4 **Review and create**

### Basic information

**Name:** Demo Agent  
**Compartment:** ...52aycgwtya [Show](#) [Copy](#)  
**Description:** A demo of an AI Agen  
**Welcome message:** I am the Oracle of all things. Ask me a question  
**Routing instructions:**

### Tools

Name	Tool type	Description
Demo RAG	RAG_TOOL_CONFIG	Demo system

Showing 1 item < 1 of 1 >

### Setup agent endpoint

**Enable human in the loop:** Disabled  
**Content moderation on input:** disable  
**Content moderation on output:** disable  
**Prompt injection (PI) protection:** disable  
**Personally identifiable information (PII) protection on input:** disable  
**Personally identifiable information (PII) protection on output:** disable

[Previous](#) [Next](#) [Create agent](#) [Cancel](#)



# Be patient!

Cloud  US Midwest (Chicago)

## Generative AI Agents

- Overview
- Agents**
- Knowledge Bases
- Chat

List scope

Compartment: cloudlab  
benneraced2022 (root)/cloudlab

Filters

State: Any state

Tag filters: [add](#) | [clear](#)  
no tag filters applied

### Agents in cloudlab Compartment

Generative AI agents connect to your data sources, retrieve data, and augment model responses with relevant, grounded information to generate more accurate and relevant responses. [Learn more about working with agents.](#)

[Create agent](#)

Name	Lifecycle state	Time created	
<a href="#">Demo Agent</a>	● Creating	Wed, 02 Apr 2025 03:39:03 GMT	⋮

Showing 1 item < 1 of 1 >



# Test with Chat

Generative AI Agents

Overview

Agents

Knowledge Bases

Chat

Agent compartment

cloudlab

benneraced2022 (root)/cloudlab

Agent endpoint compartment

cloudlab

benneraced2022 (root)/cloudlab

## Chat

Agent

Agent endpoint

Demo Agent

kendpoint

Close traces

Chatting with [Demo Agent](#)

I am the Oracle of all things. Ask me a question

Type a message...

Submit

Reset chat session

Traces

Metadata filters

Terms of Use and Privacy

Cookie Preferences

Copyright © 2025, Oracle and/or its affiliates. All rights reserved.

What's this? | Redwood preview



# Troubleshooting Chats

The screenshot displays the Oracle Cloud AI Agents chat interface. The top navigation bar includes the 'Cloud' logo, a search bar, and the region 'US Midwest (Chicago)'. The left sidebar contains navigation options: Overview, Agents, Knowledge Bases, Chat, Agent compartment (cloudlab), and Agent endpoint compartment (cloudlab). The main chat area is titled 'Chat' and shows a conversation with 'Demo Agent'. The user's message is 'I am the Oracle of all things. Ask me a question'. The agent's response is 'what are driver points?'. The agent's reply explains that driver points are used by the DMV to identify and control problem drivers, based on the seriousness of the violation. It also provides a citation for 'Drivers\_Licensing\_Handbook\_web.pdf' with details on object storage path, document ID, page numbers, and source text. A 'View citations' section is visible, and a 'Traces' panel on the right shows the search string and generated text. The bottom of the interface includes a message input field, 'Submit' and 'Reset chat session' buttons, and footer information.

Cloud

Search resources, services, documentation, and Marketplace

US Midwest (Chicago)

Generative AI Agents

Chat

Agent

Agent endpoint

Demo Agent

kbendpoint

Close traces

Chatting with **Demo Agent**

I am the Oracle of all things. Ask me a question

what are driver points?

Driver points are used by the DMV to identify and control problem drivers. Points are assigned to various in-state and out-of-state moving traffic violations depending on the seriousness of the violation. These points remain on your driving record for two years from the conviction date. When you accumulate 12 points on your record, your driver's license is subject to suspension.

**View citations**

1. **Title:** [Drivers\\_Licensing\\_Handbook\\_web.pdf](#)

**Object storage path:** ...ket/o/Drivers\_Licensing\_Handbook\_web.pdf [Show](#) [Copy](#).

**Document ID:** ...ok\_web.pdf [Show](#) [Copy](#).

**Page numbers:** 27, 28

**Source text:**

Third and subsequent offenses = \$300 and three (3) points18

Use of an electronic communications device to report a fire, traffic crash, serious road hazard, or emergency to the appropriate authorities...

[Show more](#)

**Traces**

2025-04-02T10:11:08.020Z: Retrieval request

**Search string:** "what are driver points?"

> **Sources retrieved**

2025-04-02T10:11:09.545Z: Generation

**Generated text:**

Driver points are used by the DMV to identify and control problem drivers. Points are assigned to various in-state and out-of-state moving traffic violations depending on the seriousness of the violat...

[Show more](#)

> **Metadata filters**

Type a message...

Submit

Reset chat session

Terms of Use and Privacy

Cookie Preferences

Copyright © 2025, Oracle and/or its affiliates. All rights reserved.

What's this? | Redwood preview



# DEMO TIME



Erik Benner

@erik\_benner

TalesFromTheDatacenter.com

[blog.mythics.com/posts](http://blog.mythics.com/posts)

[ebenner@mythics.com](mailto:ebenner@mythics.com)