

AWS State, Local, and Education Learning Days

Chicago



© 2024, Amazon Web Services, Inc. or its affiliates. All rights reserved.





Streamline grant proposal reviews with Amazon Bedrock Workshop

Archit Malpure (he/him)

Solutions Architect
Amazon Web Services,
malpurea@amazon.com

Mohan Alapati (he/him)

Sr. Solutions Architect
Amazon Web Services
maalapat@amazon.com



© 2024, Amazon Web Services, Inc. or its affiliates. All rights reserved.



Workshop Scenario: Administering Grant Programs

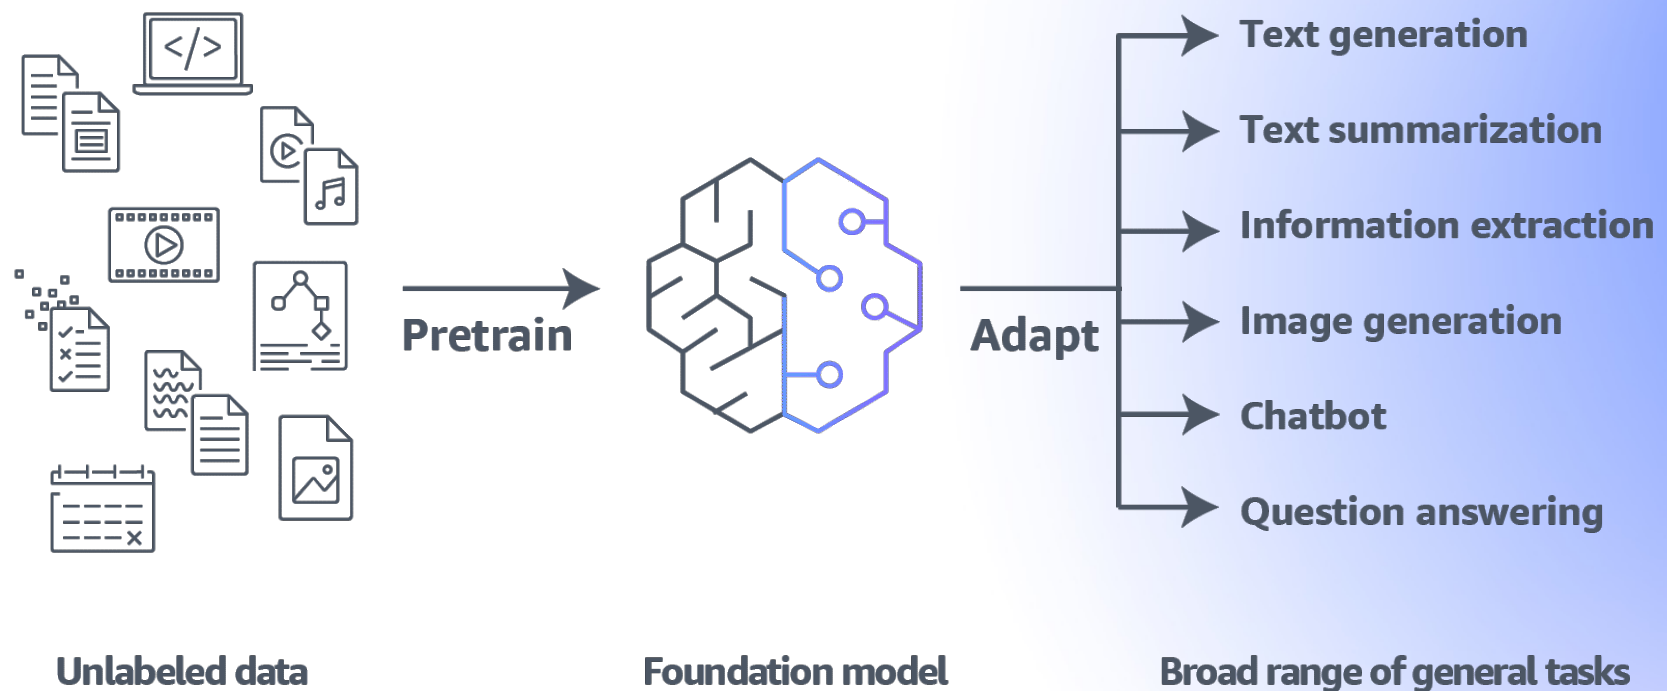


- Large amount of diverse proposals
- Limited resources to review and evaluate



- Automate manual review process
- Structure evaluations
- Incorporate diverse perspectives in a scalable manner

How does a foundation model function?



(Typical) Prompt Structure

The diagram illustrates the typical prompt structure for the Amazon Titan XL v1.01 model. It shows four components on the left, each with a red arrow pointing to its corresponding part in a chat interface on the right.

- Instructions** (blue box) points to: "Act as an IT technical expert providing customer service. Consider the Context below to answer the user's questions with a friendly tone. Answer in English in 2 sentences or less providing instructions."
- Context** (dark blue box) points to: "Context: You work in the Support line of a technology company that commercializes Android smartphones. The user is calling because the phone is not charging"
- User Input** (light blue box) points to: "User: Hi, how can I fix my phone?"
- Output Indicator** (dark blue box) points to: "Assistant: I will try my best to assist you with this. Can you please tell me the model of your phone and what version of the Android operating system it is running?"

The chat interface also includes a header with the Amazon logo, a dropdown menu for "Amazon", and a dropdown menu for "Titan XL v1.01". Below the header, it states: "Powerful, general-purpose models pretrained on large datasets, Titan FMs are powerful, general-purpose models that can be used as-is or customized to perform s". At the bottom of the chat interface, there are icons for thumbs up, thumbs down, delete, and copy.

Instructions matter

Specificity, clarity, and persuasiveness are important!

Compare:

- “Tell me a story about cows”
- “Tell me a story about cows. It should be roughly 2000 words long and appropriate for American 5th graders. It should be entertaining, but with a moral message about the value of loyalty. Make it amazing and memorable.”

Neither humans nor LLMs can read minds, so you need to be clear and specific if you want something specific.

Many professions are prompt engineering for humans!

- Marketing, Education, Tech Writing, Law, Web Design



Amazon Bedrock

BROAD CHOICE OF MODELS

AI21labs

amazon

ANTHROPIC

cohere

Meta

MISTRAL
AI

stability.ai

Contextual answers,
summarization,
paraphrasing

Jamba 1.5 Large
Jamba 1.5 Mini
Jamba-Instruct
Jurassic-2 Ultra
Jurassic-2 Mid

Text summarization,
generation, Q&A,
search, image
generation

Amazon Titan
Text Premier
Amazon Titan
Text Lite
Amazon Titan
Text Express
Amazon Titan Text
Embeddings
Amazon Titan Text
Embeddings V2
Amazon Titan
Multimodal
Embeddings
Amazon Titan
Image Generator

Summarization,
complex reasoning,
writing, coding

Claude 3.5 Sonnet
Claude 3 Opus
Claude 3 Sonnet
Claude 3 Haiku
Claude 2.1
Claude 2
Claude Instant

Text generation,
search,
classification

Command
Command Light
Embed English
Embed Multilingual
Command R+
Command R

Q&A and reading
comprehension

Llama 3.2
Llama 3.1
Llama 3 8B
Llama 3 70B
Llama 2 13B
Llama 2 70B

Text summarization,
text classification,
text completion,
code generation, Q&A

Mistral Large 2 (24.07)
Mistral Large (24.02)
Mistral Small
Mixtral 8x7B
Mistral 7B

High-quality
images and art

Stable Image Ultra
Stable Diffusion 3
Large
Stable Image Core
Stable Diffusion XL1.0
Stable Diffusion
XL 0.8



© 2024, Amazon Web Services, Inc. or its affiliates. All rights reserved.

Prompt Guidance for Anthropic Claude

Being clear and direct

- Claude responds best to **clear and direct instructions**
- When in doubt, follow the **Golden Rule of Clear Prompting**: show your prompt to a friend and ask them if they can follow the instructions themselves and produce the exact result you're looking for



Example:

Human: Write a haiku about robots

Assistant: Here is a haiku about robots:

Metal bodies move
Circuits calculate tasks
Machines mimic life



Human: Write a haiku about robots. **Skip the preamble; go straight into the poem.**

Assistant: Metal bodies move
Circuits calculate tasks
Machines mimic life

Prompt Guidance for Anthropic Claude

Assigning roles (aka role prompting)

- Claude sometimes needs **context about what role it should inhabit**
- Assigning roles changes Claude's response in two ways:
 - **Improved accuracy** in certain situations (such as mathematics)
 - **Changed tone and demeanor** to match the specified role



Example:

Human: How is the U.S. President elected?

Assistant: The president is not voted on directly by the people. Instead, each state gets a certain number of electors based on its population. The states with more people get more electors. When people vote on Election Day in November, they are actually voting for which candidate they want their state's electors to...



Human: **You are a kindergarten teacher.** How is the U.S. President elected?

Assistant: The president is the leader of our country. Every four years, we have a presidential election to choose a new president or re-elect the current president. During the election, people called candidates try to become president by getting votes. On Election Day in November, people go to voting places...

Prompt Guidance for Anthropic Claude

Using XML tags

- Disorganized prompts are hard for Claude to comprehend
- Just like section titles and headers help humans better follow information, **using delineators like XML tags `<></>` helps Claude understand the prompt's structure**

We recommend you use XML tags, as Claude has been especially trained on XML tags

Example:

Human: Hey Claude. Show up at 6AM because I say so. Make this email more polite.

Assistant: Dear Claude, I hope this message finds you well

Human: Hey Claude. **<email>**Show up at 6AM because I say so.**</email>** Make this email more polite.

Assistant: Good morning team, I hope you all had a restful weekend...

Prompt Guidance for Anthropic Claude

Formatting output & speaking for Claude

- You can get Claude to say exactly what you want by:
 - **Specifying the exact output format** you want
 - **Speaking for Claude** by writing the beginning of Claude's response for it (after "Assistant:")

Prompt

Human: Please write a haiku about a cat. Put it in <haiku> tags.
Assistant: <haiku>

Claude's
response

Curious little cat
Pouncing on strings and lasers
Naps in sunny spots
</haiku>

Thinking step by step

- Claude benefits from having time to think through tasks before executing
- Especially if a task is particularly complex, **tell Claude to think step by step before it answers**

Increases intelligence of responses but also increases latency by adding to the length of the output.

Prompt

Human: Are both the directors of Jaws and Casino Royale from the same country?
Can you answer the question by thinking through step by step?

Claude's
response

Human: Are both the directors of Jaws and Casino Royale from the same country?

Assistant: "Okay, let's think through this step-by-step:

1. The director of Jaws was Steven Spielberg. Steven Spielberg is American.
2. The director of Casino Royale (2006) was Martin Campbell. Martin Campbell is from New Zealand.
3. America and New Zealand are different countries.
Therefore, the directors of Jaws and Casino Royale are not from the same country, as Steven Spielberg is American and Martin Campbell is from New Zealand."

Using examples

- **Examples are probably the single most effective tool** for getting Claude to behave as desired
- Make sure to **give Claude examples of common edge cases.**
- Generally **more examples = more reliable responses** at the cost of latency and tokens

Prompt

Claude's response

Example:

Provide a quote from a book in the format of the below example

"When the reasoning mind is forced to confront the impossible again and again, it has no choice but to adapt."
— N.K. Jemisin, The Fifth Season
Author: N.K. Jemisin

"Some humans theorize that intelligent species go extinct before they can expand into outer space. If they're correct, then the hush of the night sky is the silence of the graveyard."
— Ted Chiang, Exhalation
Author: Ted Chiang

Sometimes, examples are so powerful that they don't even need to be accompanied by task instructions!

Prompting Resource Portal

Anthropic Prompt Engineering Resource Guide

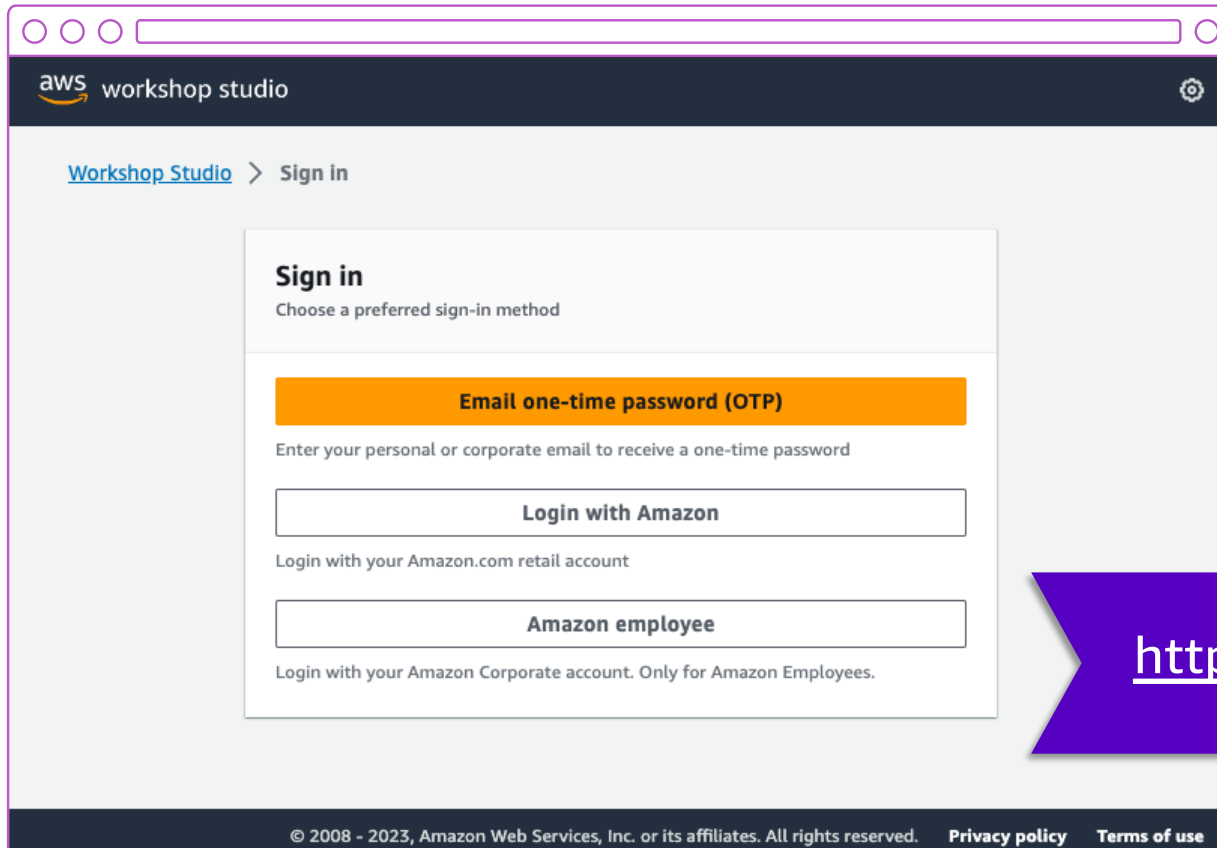
docs.anthropic.com/claude/docs

Contains:

Prompting tutorial, prompting guide, presentations, API information, sample code, Bedrock SDK, tools for experimenting with RAG and agents



Step 1: Sign in using your preferred method



The screenshot shows the AWS Workshop Studio sign-in interface. At the top, there's a dark header with the AWS logo and 'workshop studio' text. Below the header, a breadcrumb trail shows 'Workshop Studio' followed by a chevron and 'Sign in'. The main content area is a light gray box with a white border. Inside, the heading 'Sign in' is followed by the instruction 'Choose a preferred sign-in method'. There are three sign-in options, each in a white box with a thin border: 1. 'Email one-time password (OTP)' with an orange background, followed by the instruction 'Enter your personal or corporate email to receive a one-time password' and an empty input field. 2. 'Login with Amazon' with a white background, followed by the instruction 'Login with your Amazon.com retail account' and an empty input field. 3. 'Amazon employee' with a white background, followed by the instruction 'Login with your Amazon Corporate account. Only for Amazon Employees.' and an empty input field. At the bottom of the page, a dark footer contains the copyright notice '© 2008 - 2023, Amazon Web Services, Inc. or its affiliates. All rights reserved.', a 'Privacy policy' link, and a 'Terms of use' link.

aws workshop studio

Workshop Studio > Sign in

Sign in
Choose a preferred sign-in method

Email one-time password (OTP)
Enter your personal or corporate email to receive a one-time password

Login with Amazon
Login with your Amazon.com retail account

Amazon employee
Login with your Amazon Corporate account. Only for Amazon Employees.

© 2008 - 2023, Amazon Web Services, Inc. or its affiliates. All rights reserved. [Privacy policy](#) [Terms of use](#)

<https://catalog.workshops.aws/join>



© 2024, Amazon Web Services, Inc. or its affiliates. All rights reserved.

Step 2: Enter the event access code

aws workshop studio

Workshop Studio > Join event

Step 1
Enter event access code

Step 2
Review and join

Enter event access code

Event access code

Event access code
A 12 digit code that was given to you for this event

abcd-012345-ef

Cancel Next

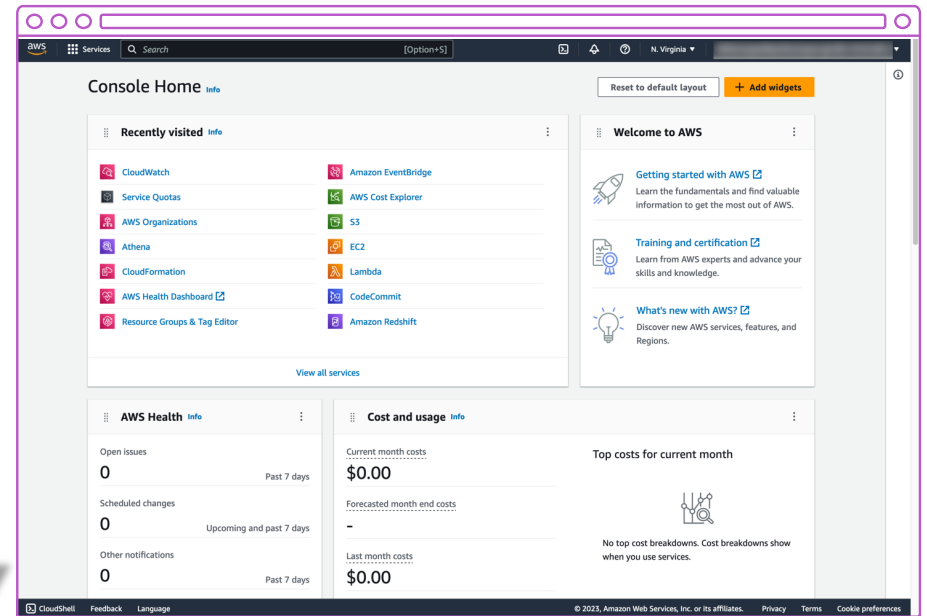
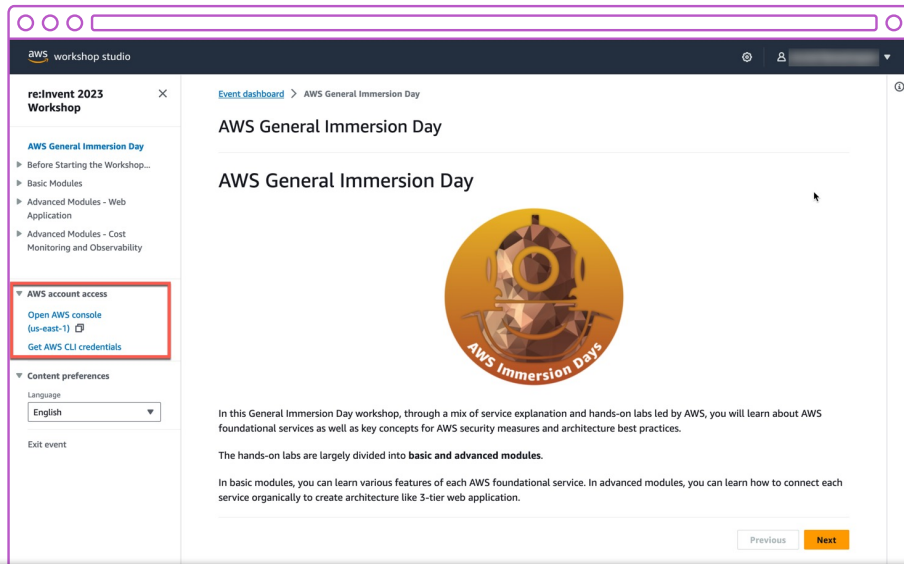
841a-08fe05-45

<https://catalog.workshops.aws/join>

© 2008 - 2023, Amazon Web Services, Inc. or its affiliates. All rights reserved. [Privacy policy](#) [Terms of use](#)



Step 3: Access AWS account



Access the AWS Management Console or generate AWS CLI credentials as needed



© 2024, Amazon Web Services, Inc. or its affiliates. All rights reserved.

Additional resources



Bedrock Docs

AWS documentation for
Amazon Bedrock



© 2024, Amazon Web Services, Inc. or its affiliates. All rights reserved.



Thank you!

Archit Malpure

Solutions Architect
Amazon Web Services,
malpurea@amazon.com

Mohan Alapati

Sr. Solutions Architect
Amazon Web Services
maalapat@amazon.com

Please complete the survey
for this session



Track Name: Artificial intelligence and machine
learning track

Session Name: Streamline
grant proposal reviews with
Amazon Bedrock



© 2024, Amazon Web Services, Inc. or its affiliates. All rights reserved.

