

Confidential AI inference in practice: What's required and how to implement it

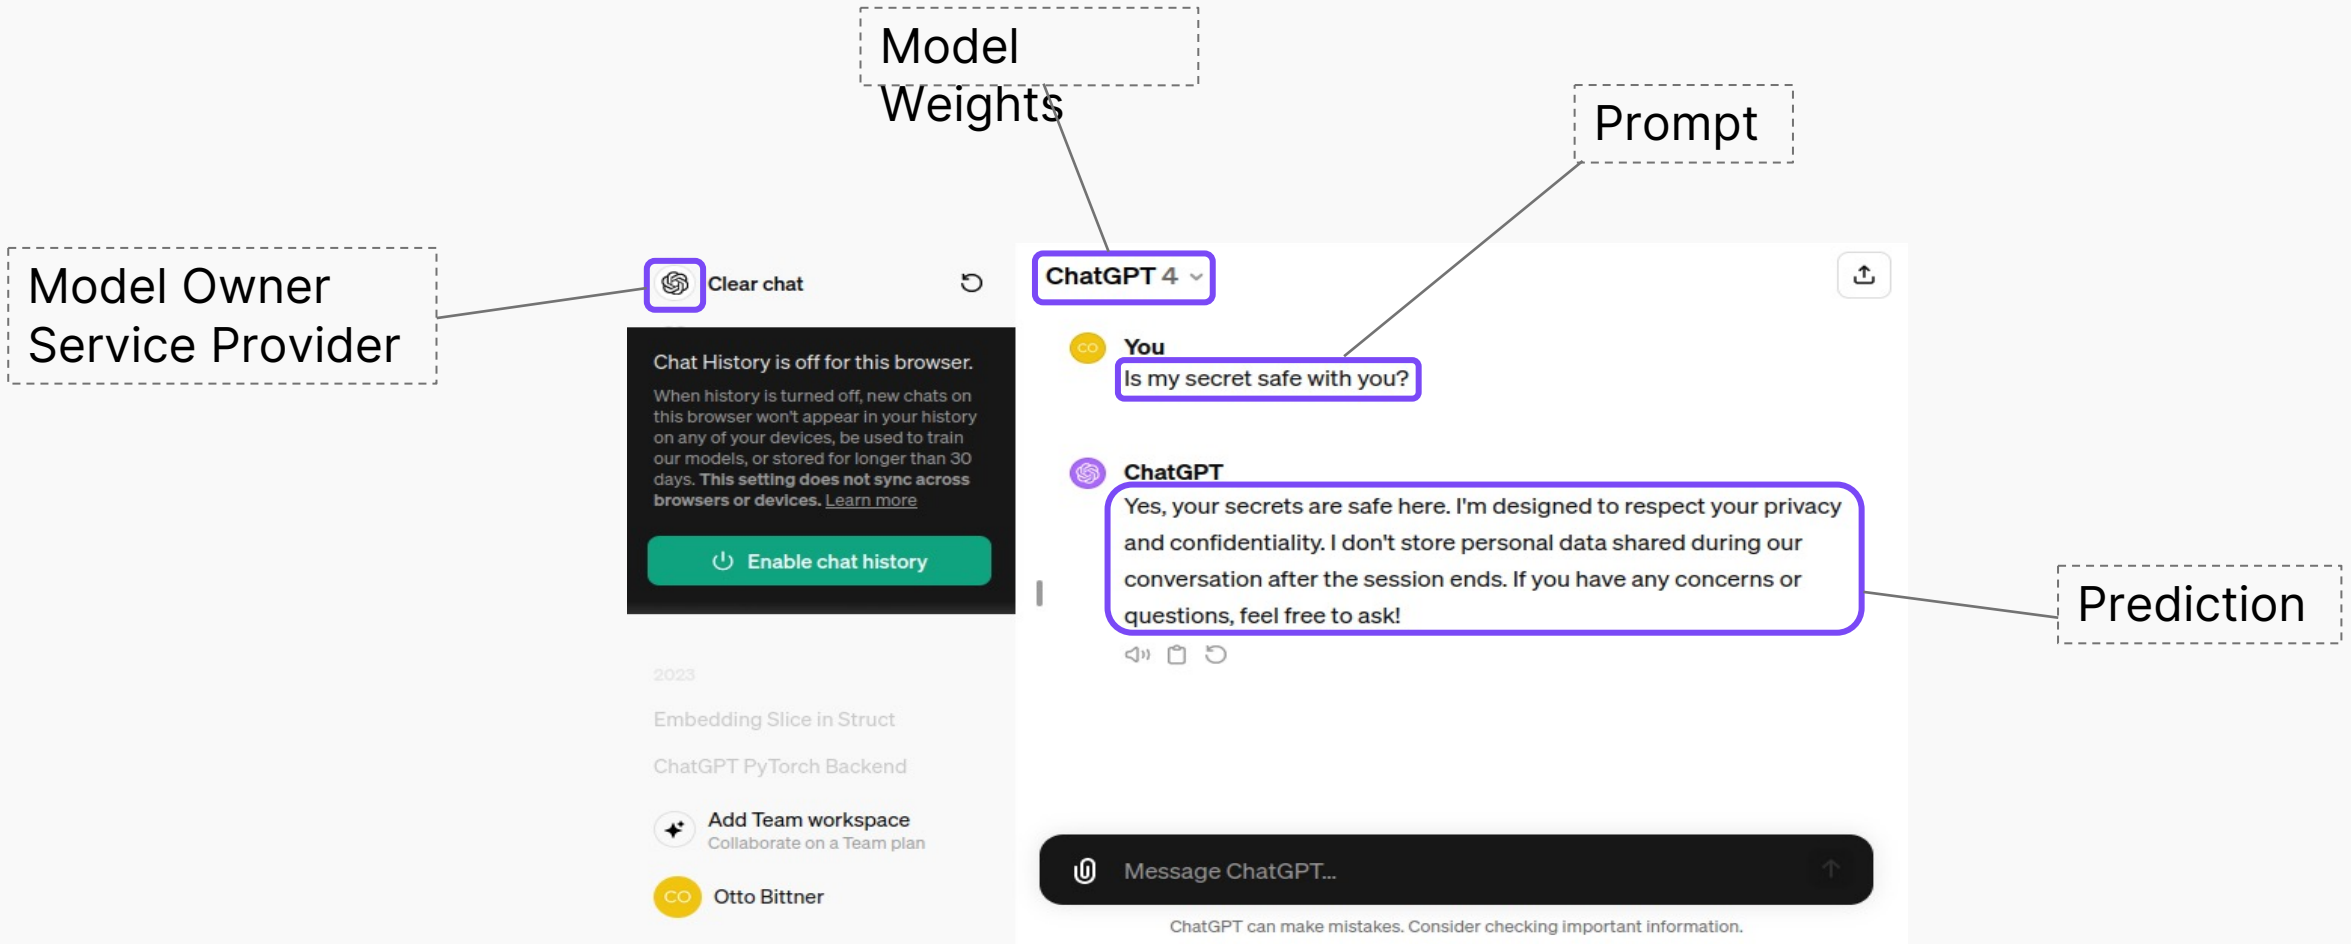
2024-03-17

Otto Bittner

Open Confidential Computing
Conference

01

Challenge



Service Provider

The screenshot shows the HuggingChat interface. On the left is a sidebar with navigation options: Login, Theme, Models (7), Assistants (New), Settings, and About & Privacy. The main chat area shows a conversation with the prompt "Is my secret safe with you?" and the response "Yes, your secret is safe with me. I am a machine learning model and do not have the ability to disclose any information. I will treat all input as confidential and use it only for the purpose of providing a response." Below the chat is a "Search web" toggle and an input field containing "Ask anything". At the bottom, the model name "mistralai/Mixtral-8x7B-Instruct-v0.1" is displayed, along with a disclaimer: "Generated content may be inaccurate or false." and a "Share this conversation" link.

Prompt

Prediction

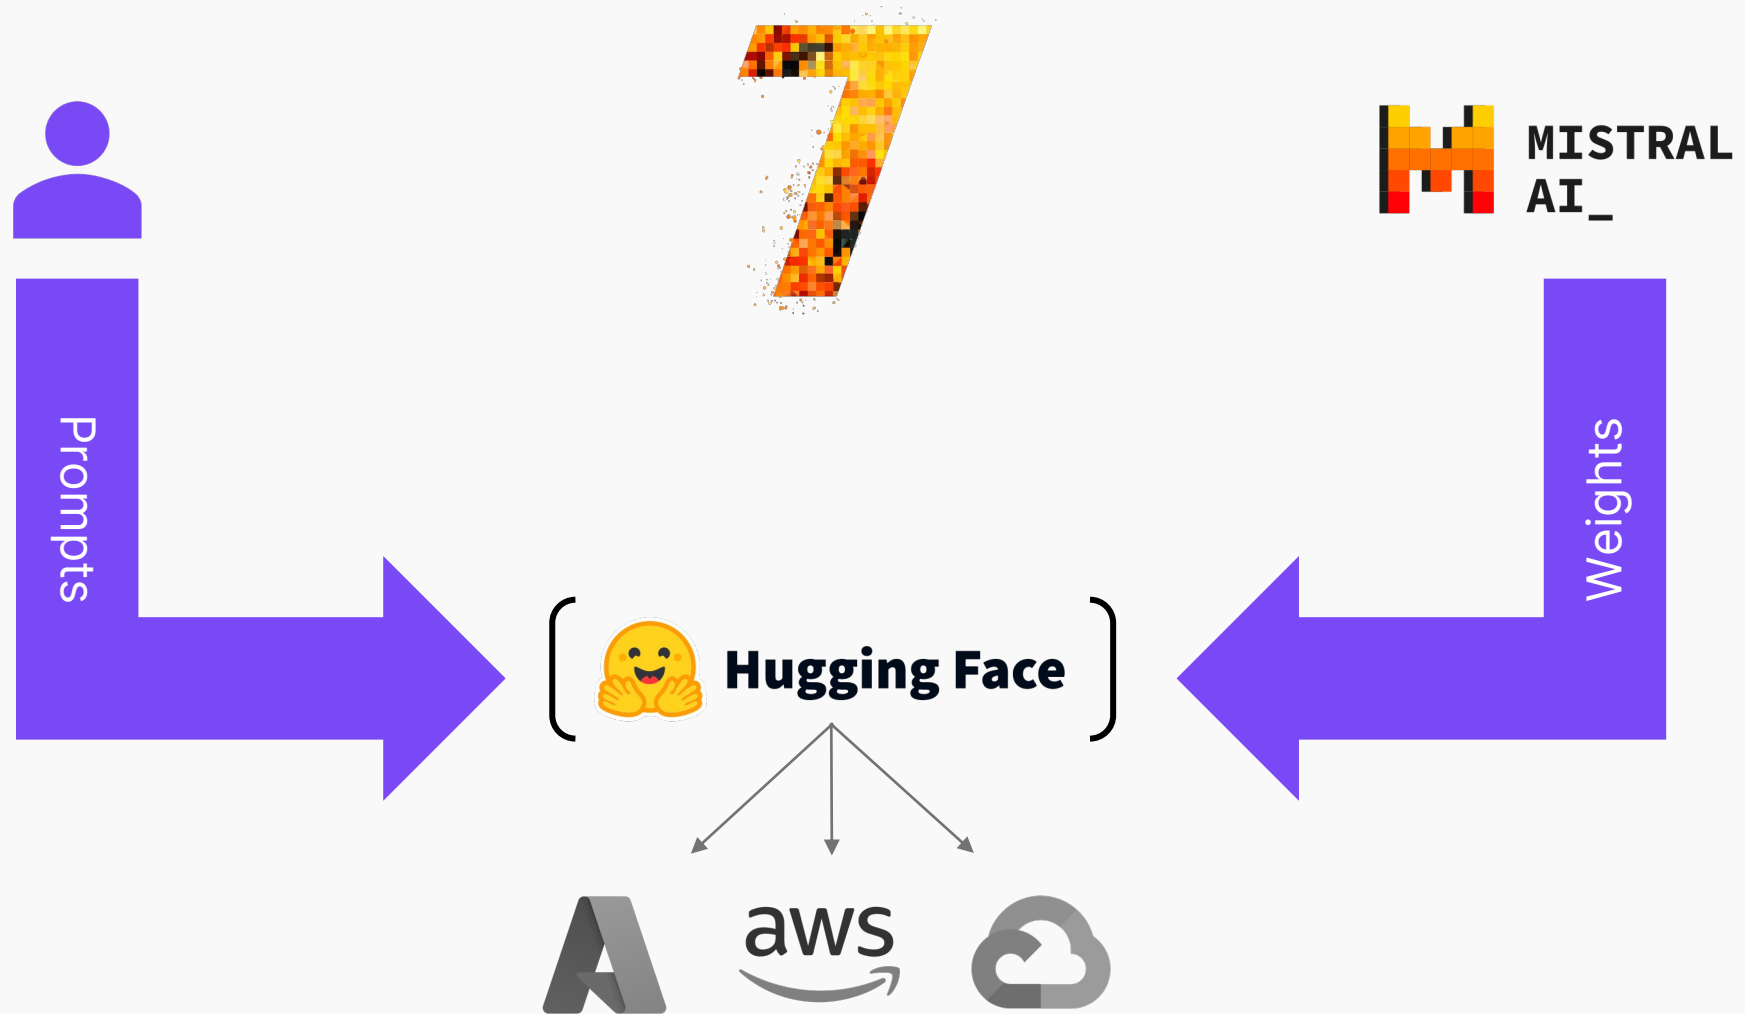
Model Weights

Model Owner

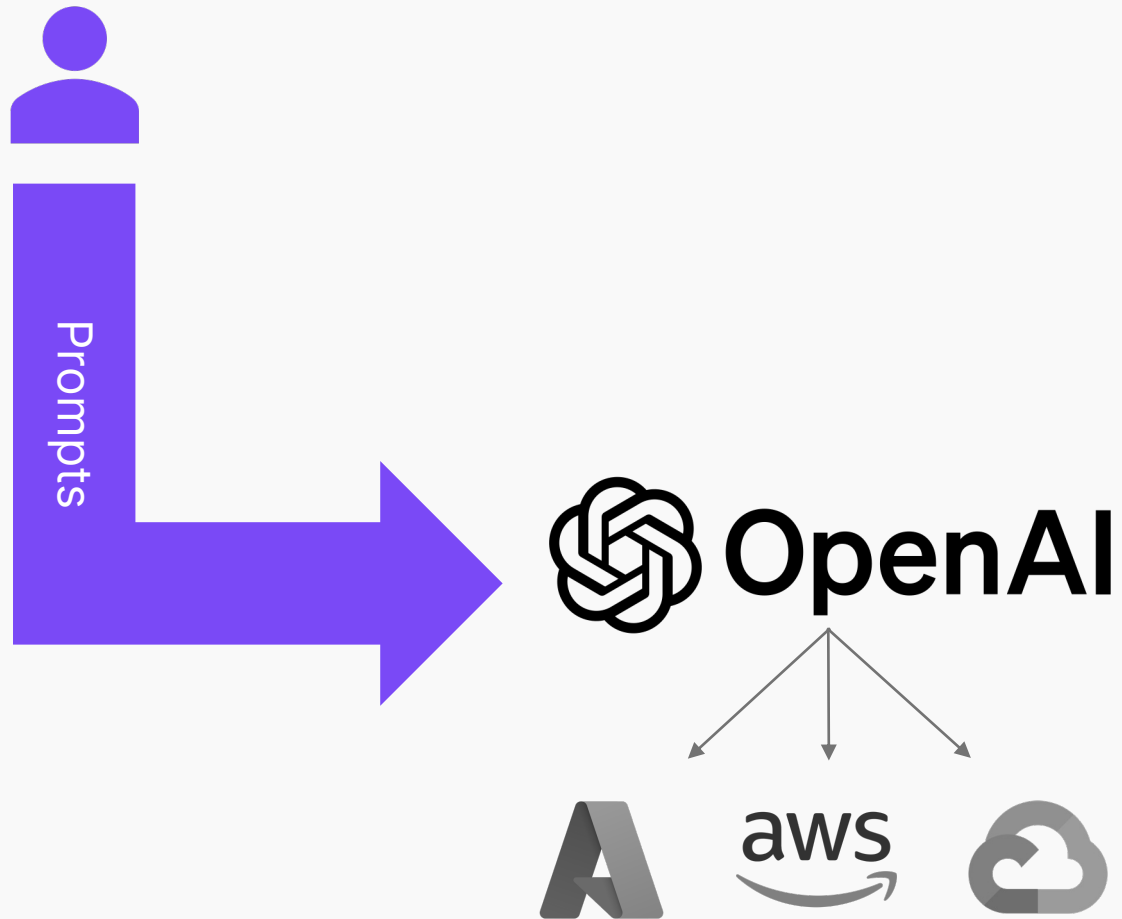


Challenge?

Challenge



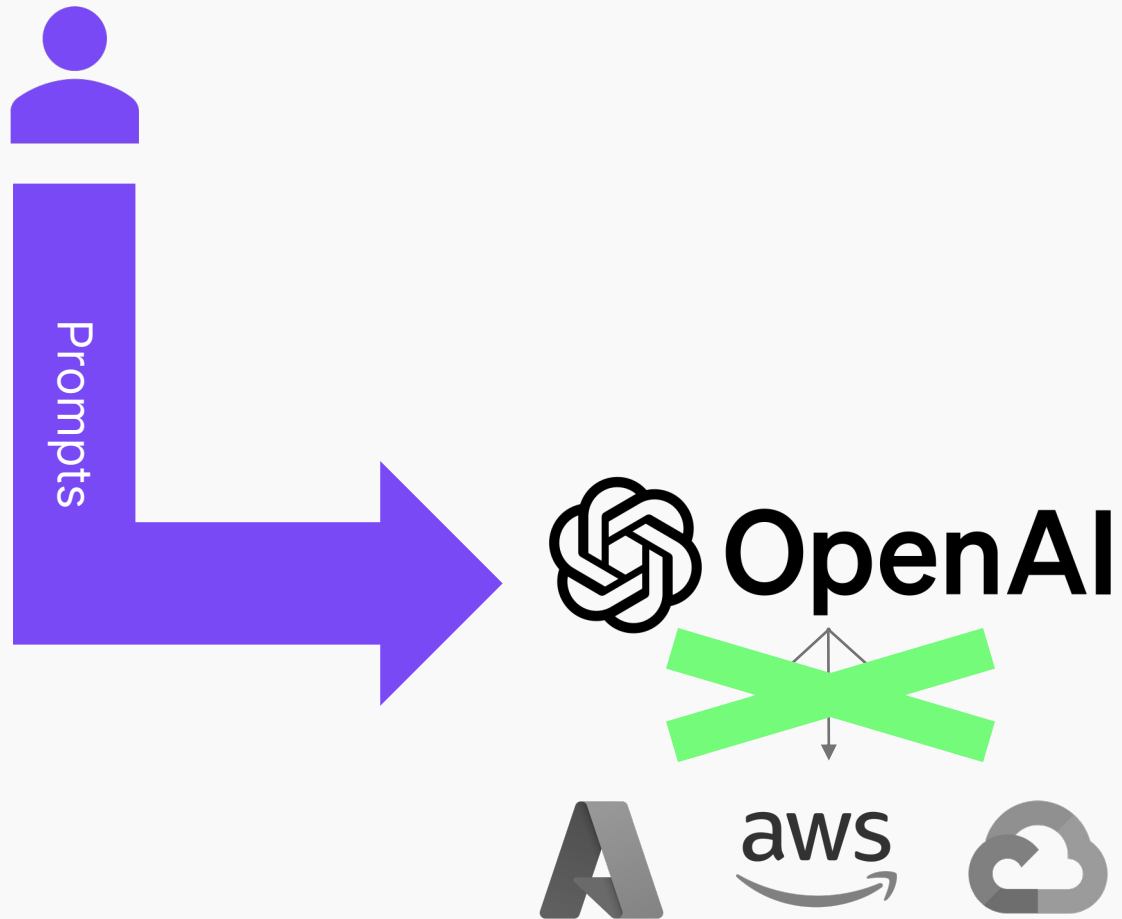
Challenge



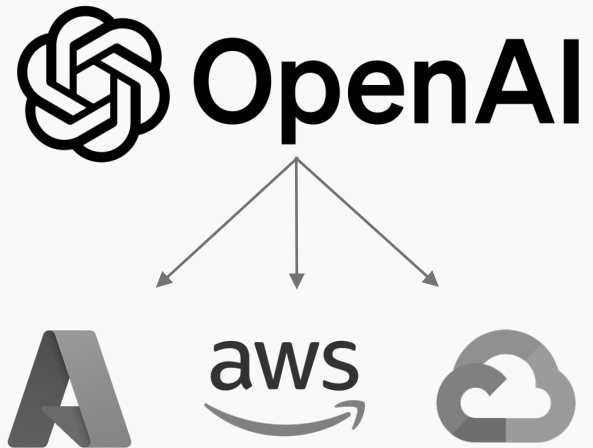
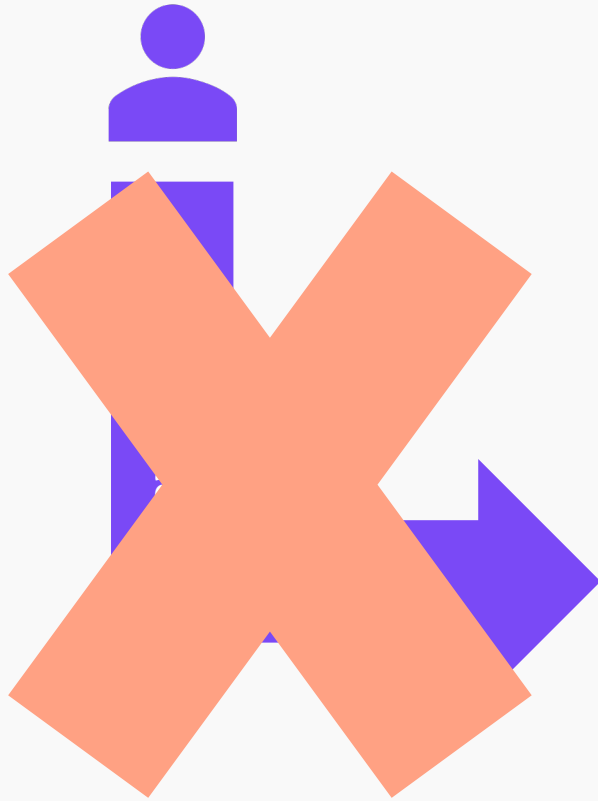
02

Goals

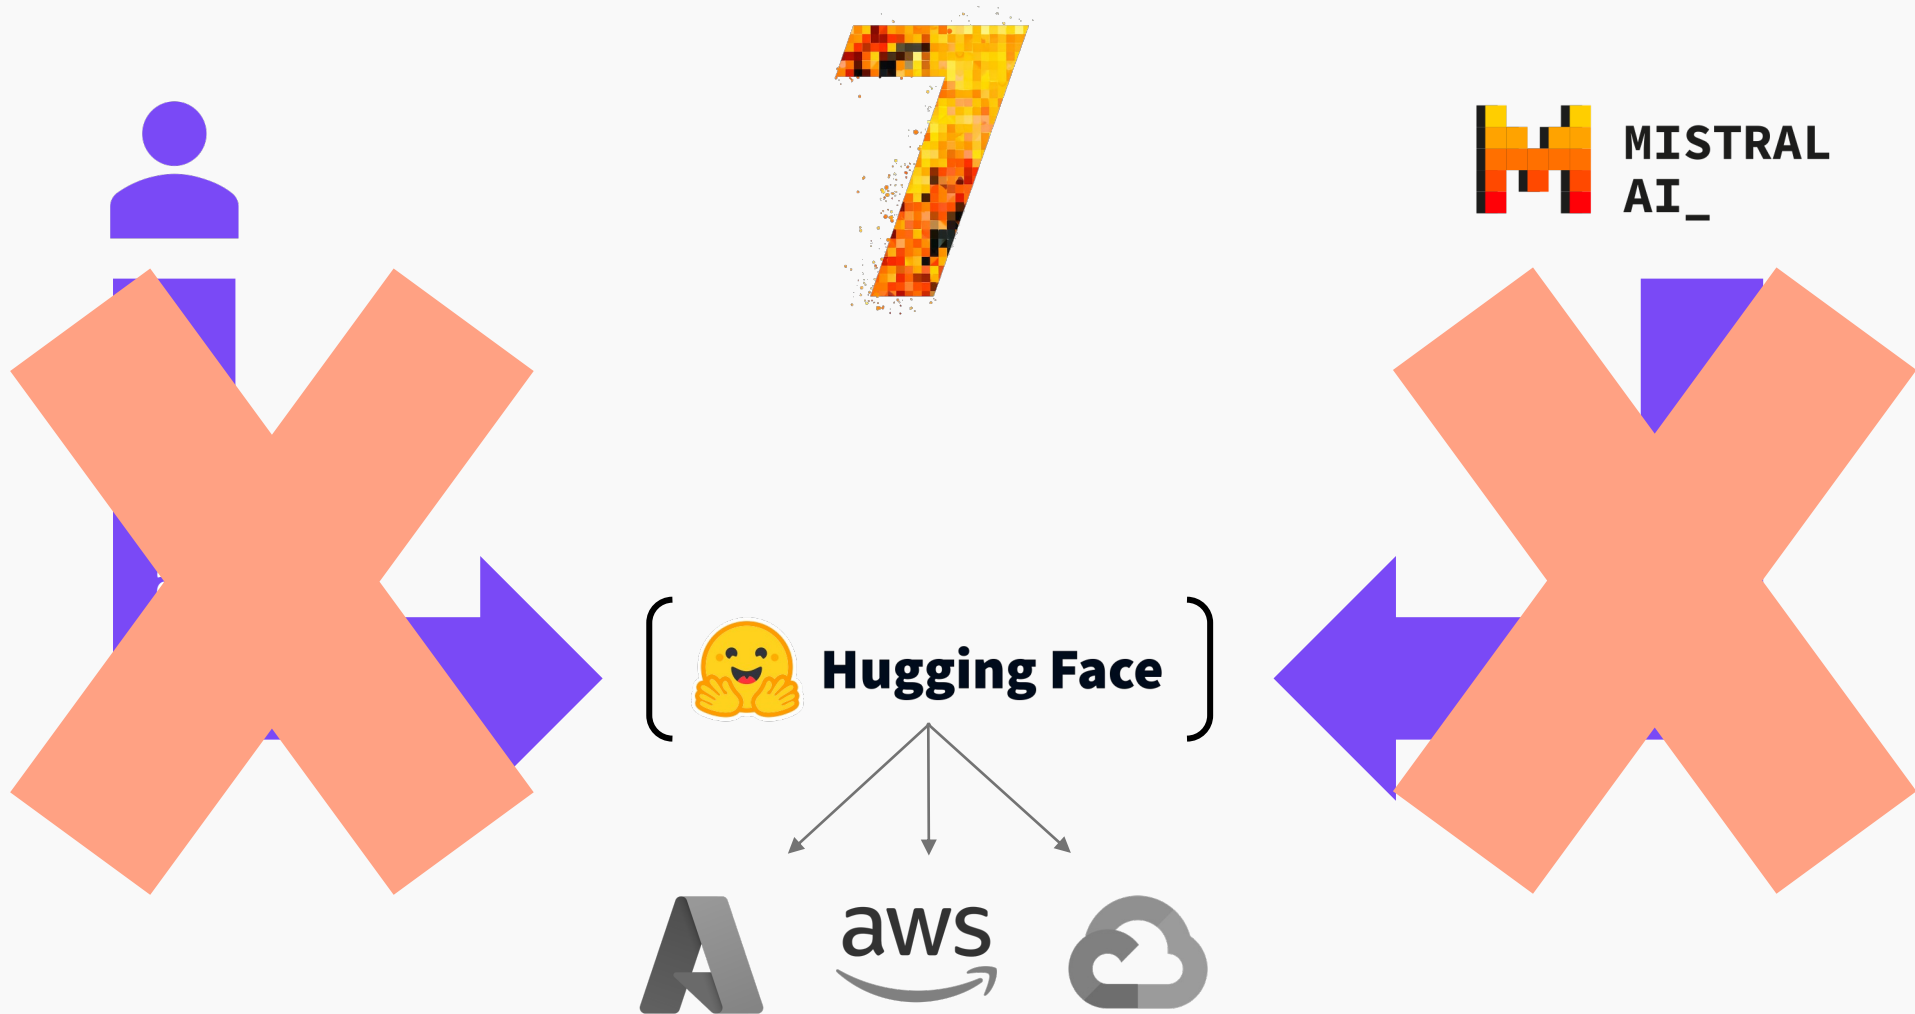
Goal #1



Goal #2



Goal #3



Goals - specific

- A **service** provider should have neither access to **prompts** nor **weights**

- An **infrastructure** provider should have neither access to **prompts** nor **weights**

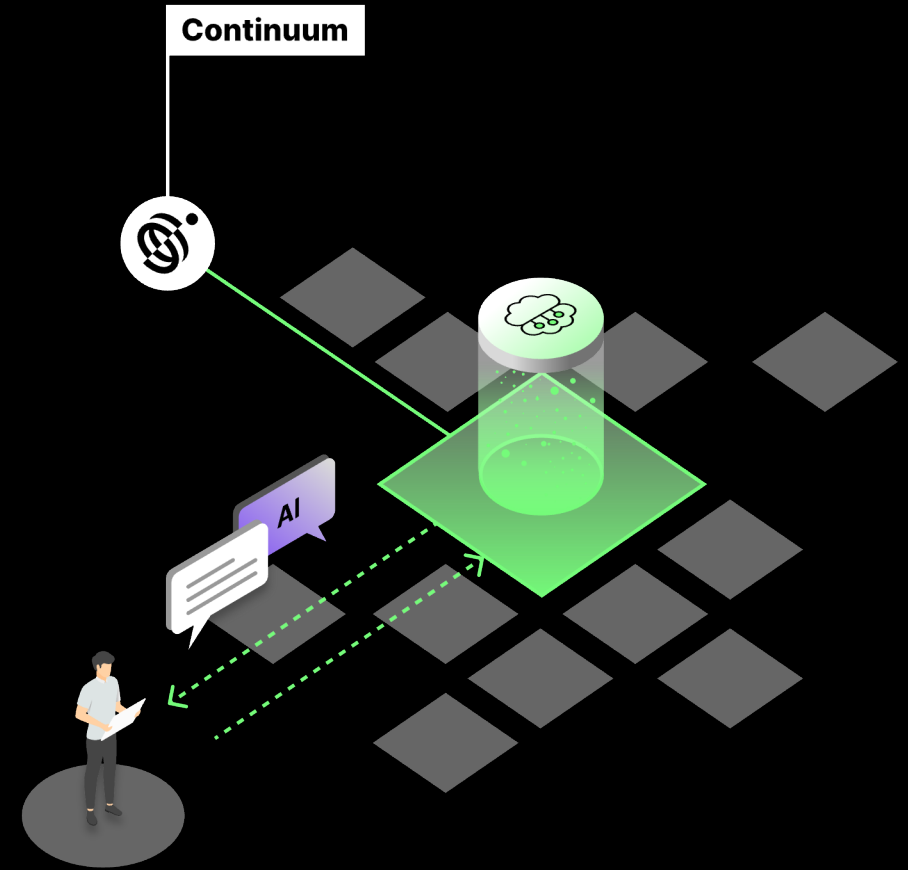
03

Solution



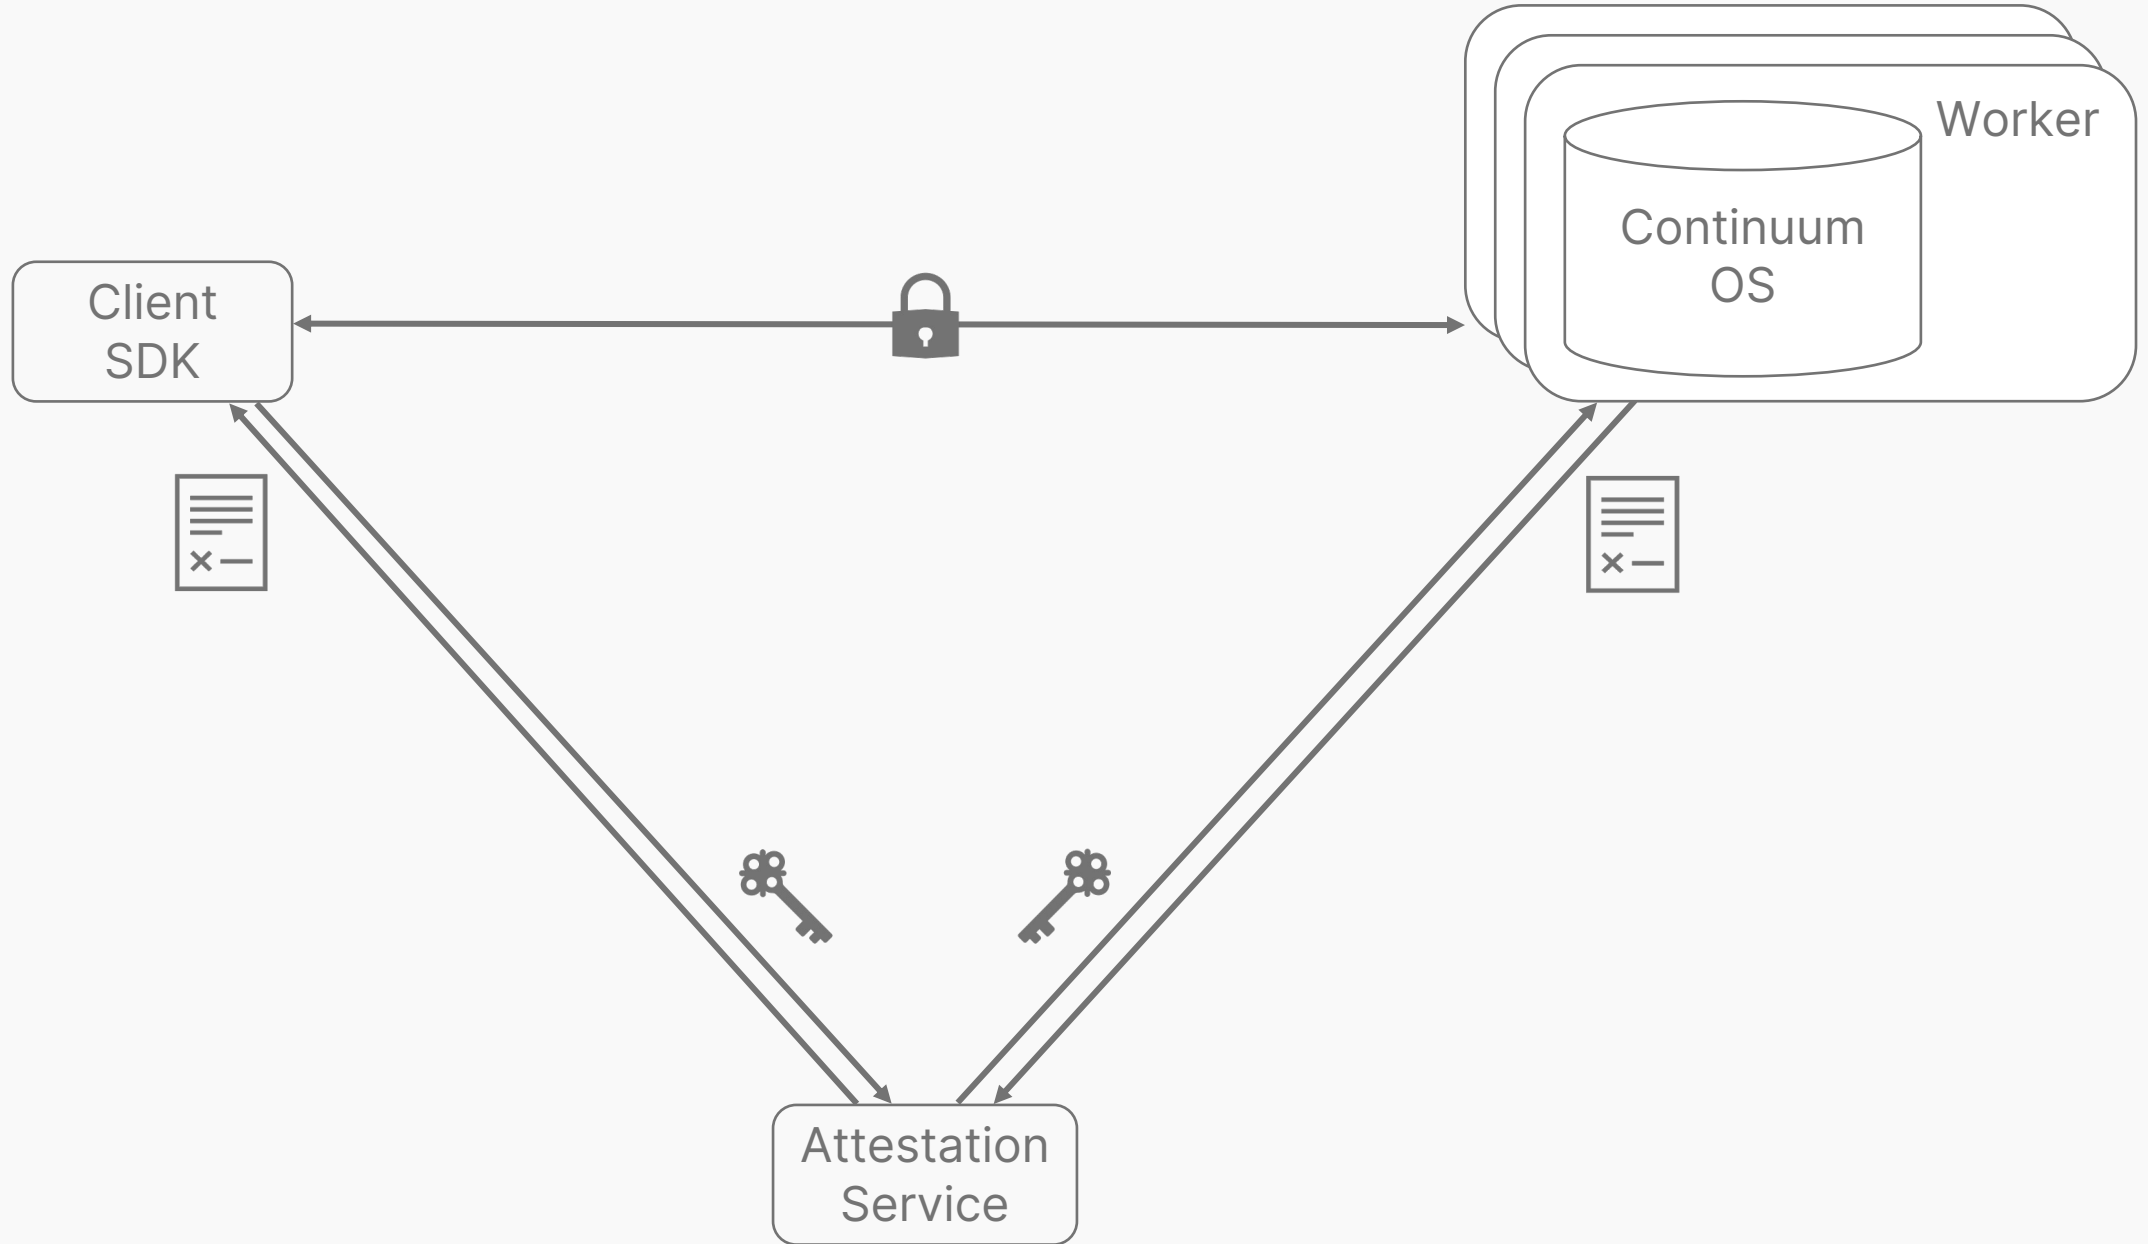
Continuum

Continuum

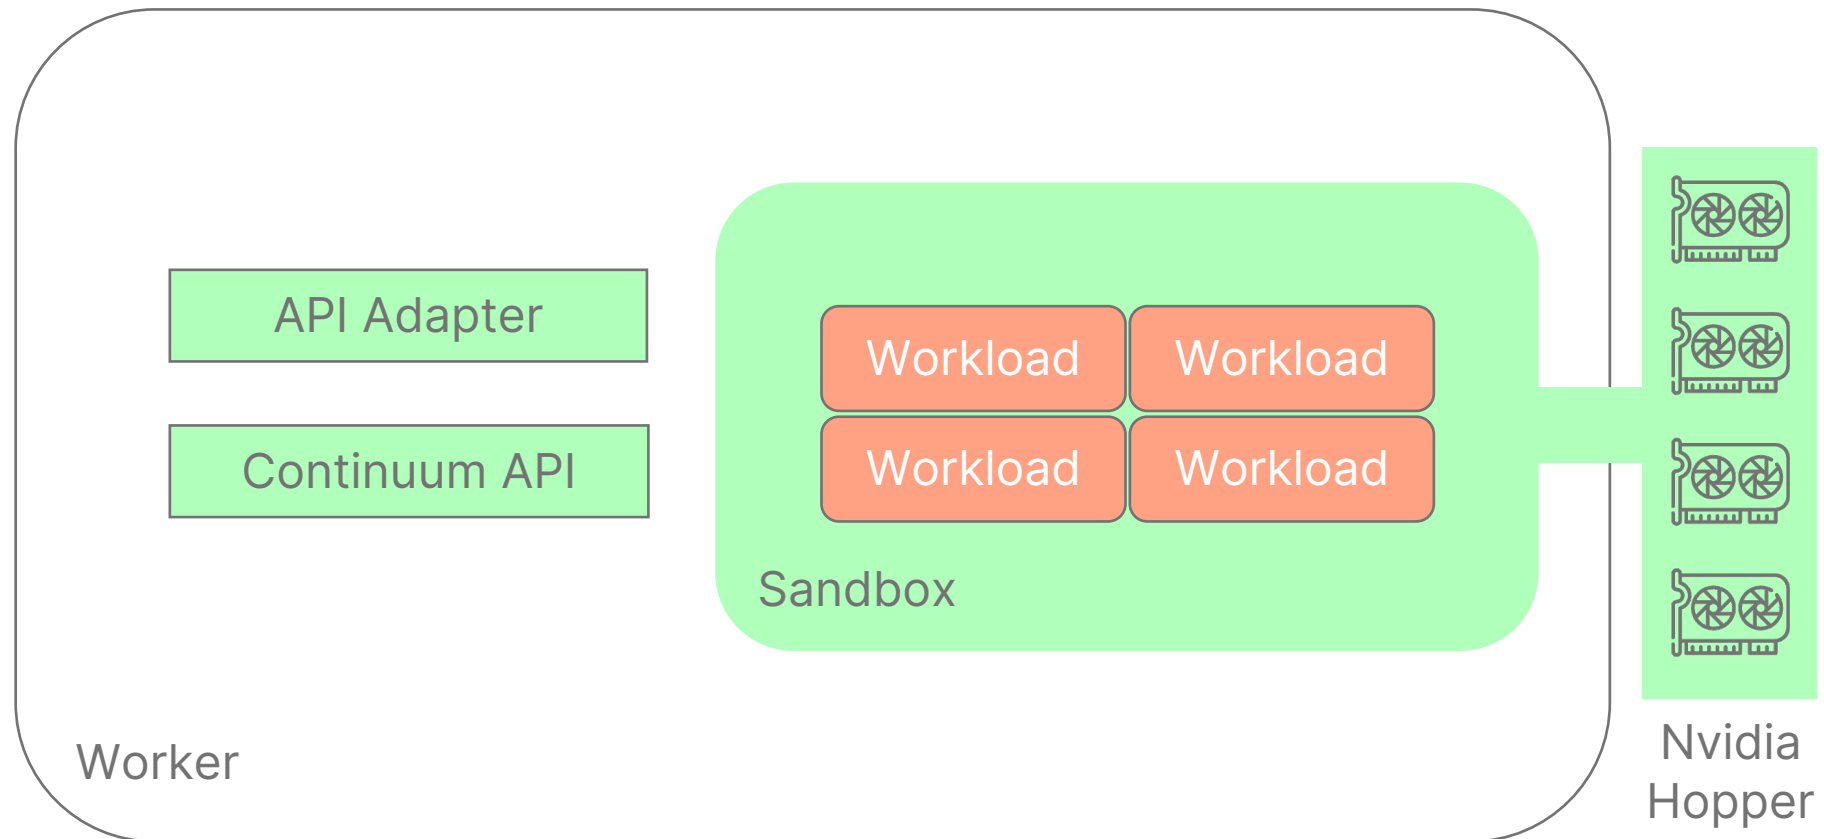


04

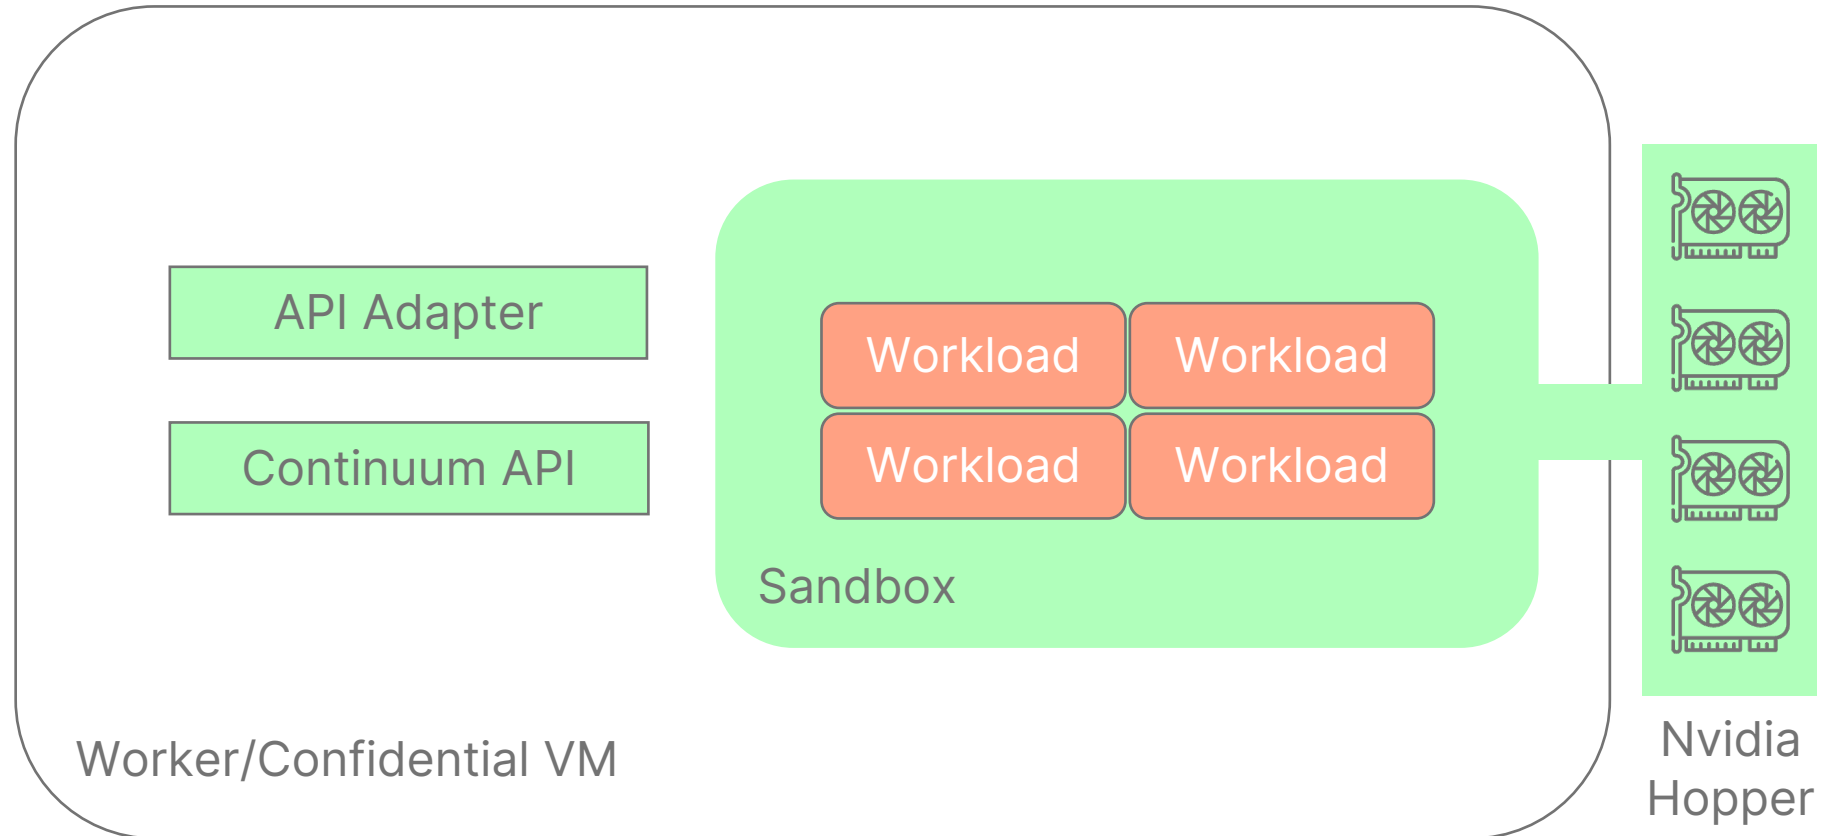
Implementation



Architecture



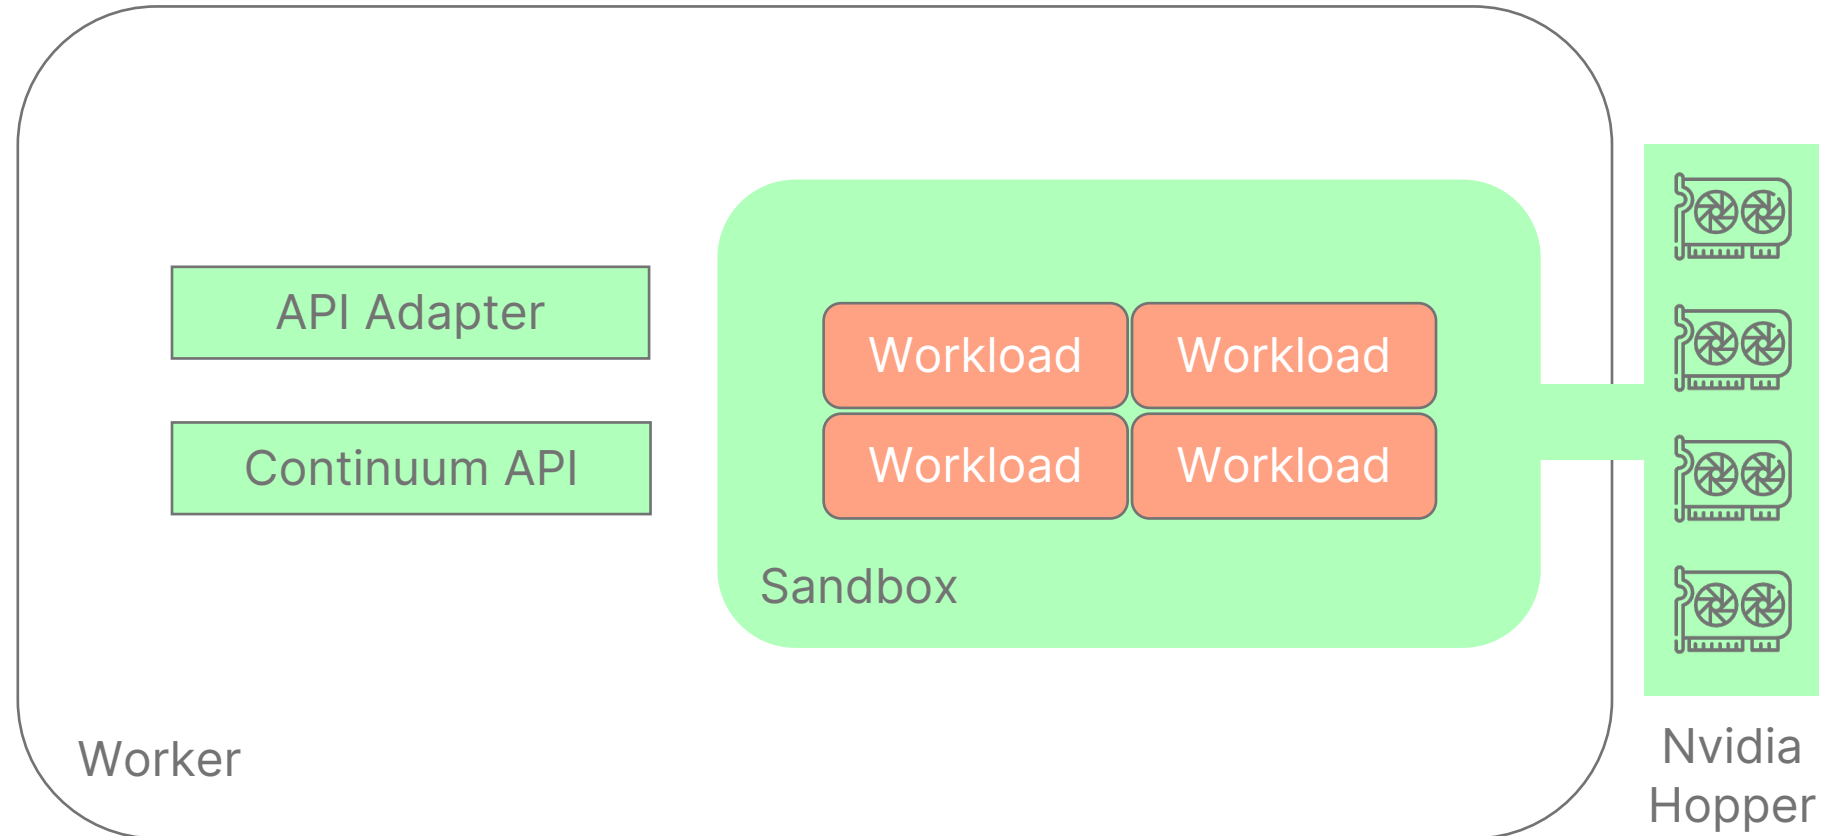
Architecture



Architecture

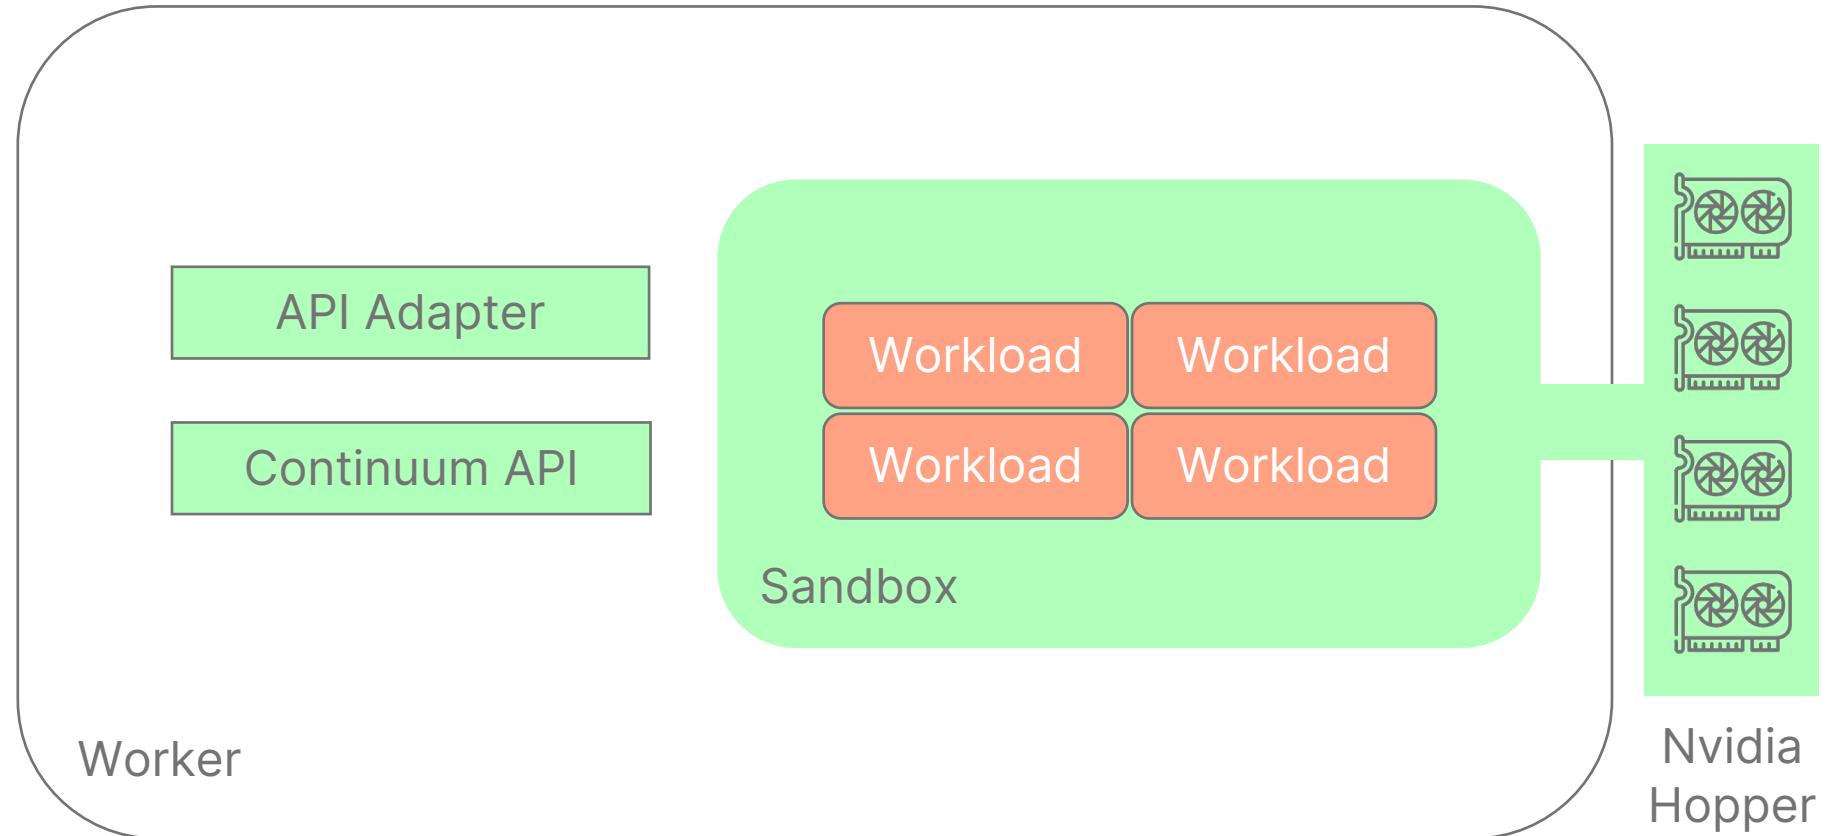
Sandbox prevents:

- Network access
- Persistent storage access



Architecture

API Adapter provides:
- Transparent encryption

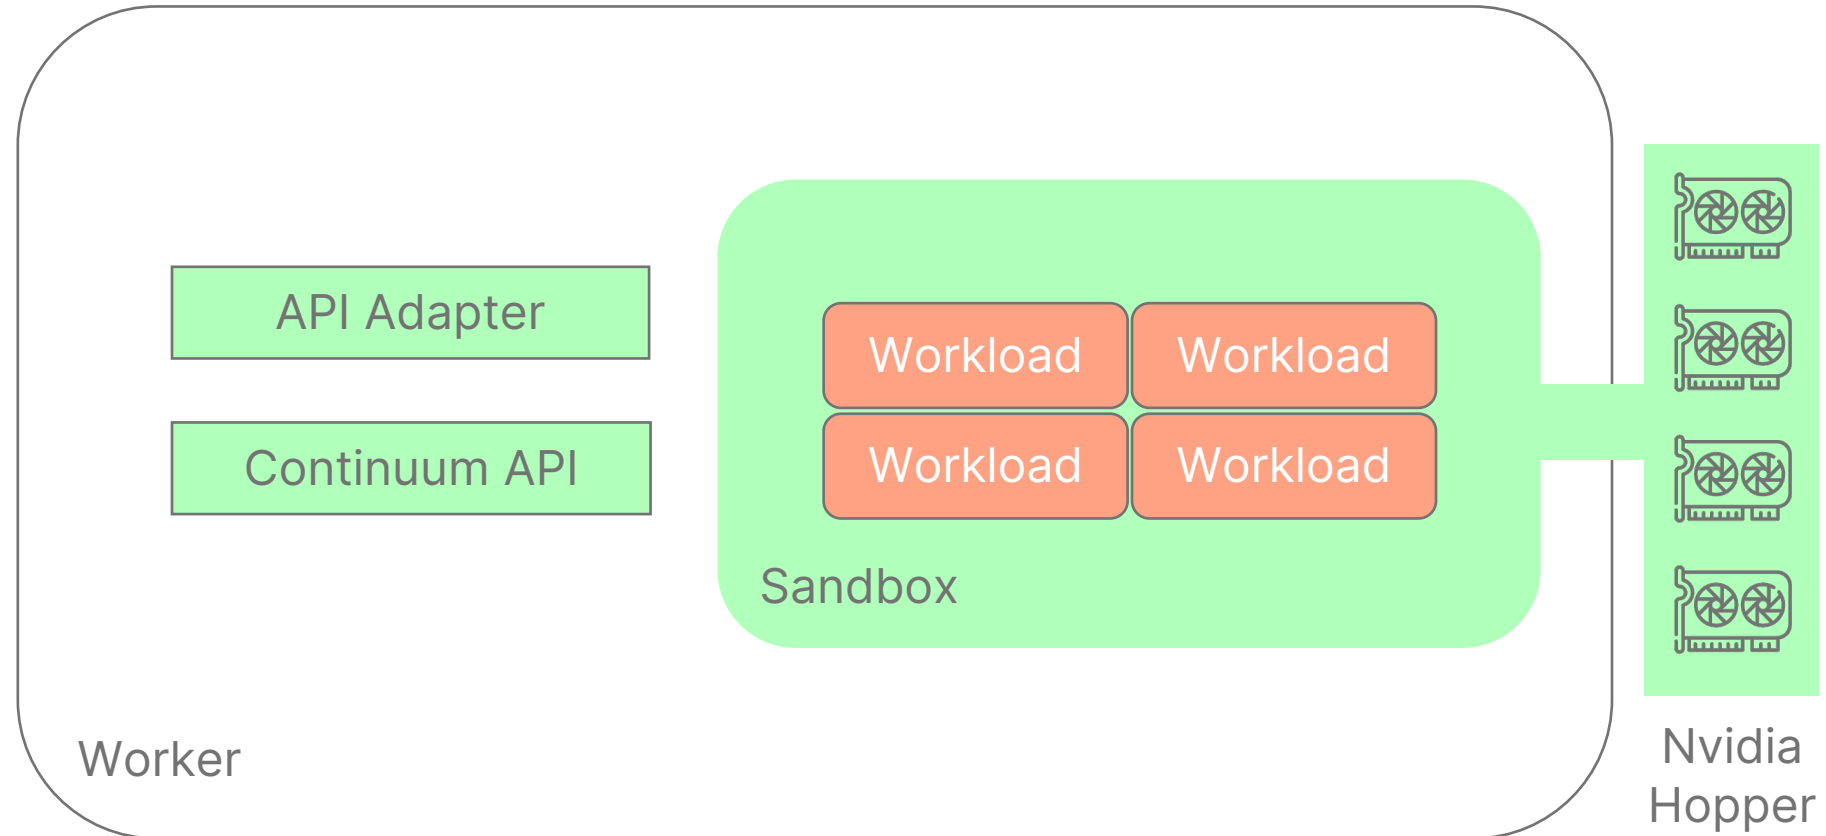


Nvidia
Hopper

Architecture

Continuum API provides:

- Deployment API
- Secret management
- Remote attestation



Nvidia
Hopper

05

Demo

Upload Attestation Service Manifest & Secrets

Attestation Service Manifest

 no file selected

Encryption Secrets

 no file selected

06

Outlook

Outlook

- Working with early partners for evaluation
- More Client SDKs
- Open source release in Q2/24.



Thank you!