

PRIVATE MODEL HOSTING. WORKFLOWS. EVALS.

# Production AI pipelines done properly.

Define your task. Build your pipeline. Measure and improve. Runs inside your own AWS account -- no data leaves your infrastructure.

## HOW IT WORKS

<p>01</p> <p><b>Inference API</b></p> <p>Async, typed, multi-modal. Text, image, audio, depth.</p>	<p>02</p> <p><b>Workflow execution</b></p> <p>Declarative YAML pipelines with parallel step execution.</p>	<p>03</p> <p><b>Evals</b></p> <p>Score any pipeline against labelled data. Improves with use.</p>
--------------------------------------------------------------------------------------------------------	------------------------------------------------------------------------------------------------------------	---------------------------------------------------------------------------------------------------

## WHAT RUNS ON IT

<p><b>Document Q&amp;A</b></p> <p>Extract, answer, and embed document content in one pipeline.</p> <p><a href="#">IMG2TXT</a> <a href="#">INSTRUCT</a></p>	<p><b>Accessibility audio</b></p> <p>Alt-text and audio descriptions in multiple languages from one image.</p> <p><a href="#">IMG2TXT</a> <a href="#">TTS</a></p>	<p><b>Site monitoring</b></p> <p>Embed, compare, describe changes, deliver audio report on schedule.</p> <p><a href="#">IMAGE-EMBED</a> <a href="#">INSTRUCT</a></p>
------------------------------------------------------------------------------------------------------------------------------------------------------------	-------------------------------------------------------------------------------------------------------------------------------------------------------------------	----------------------------------------------------------------------------------------------------------------------------------------------------------------------