CASE STUDY

AI Property Valuation Engine

Full property analysis and a defensible price prediction in under two minutes, for less than a quarter.

THE WORK

I built an end-to-end AI engine that does the work of a manual comparative market analysis in a fraction of the time and cost. The pipeline runs computer vision over every listing photo to read quality and condition, automatically selects and ranks the best comparable sales, scores the location from geospatial signals, and factors in live market conditions. It then predicts the actual sale price and days on market with a confidence range.

Crucially, it explains its reasoning in plain English, so the output is trusted rather than just delivered. To keep it affordable at scale, expensive analysis is cached and reused, payloads stay lean to control token cost, and an appraisal-style review step sanity-checks the model's own output and corrects errors before they reach the user.

TECHNOLOGIES

Vellum · Claude · GPT-4o Vision · Gemini Vision · AWS Rekognition · Pinecone · Supabase · Retool · Python

WHAT I TOOK AWAY

- Cost discipline is the real engineering. Lean payloads and cached analysis are what let an AI pipeline scale affordably.
- Explainability is the product. The number is a commodity; the defensible “why” is the differentiator.
- A step that reviews the model's own output catches errors before they ever reach the client.

THE REPORT IT PRODUCES

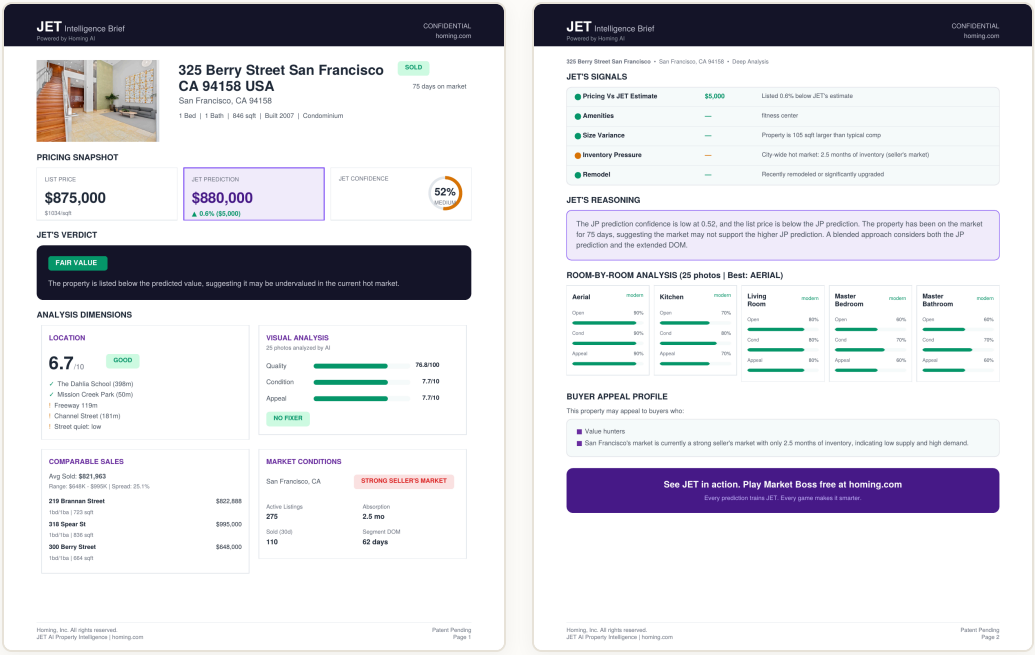

The JET Intelligence Brief, front and back — the full report the engine produces for one property.

WATCH IT RUN

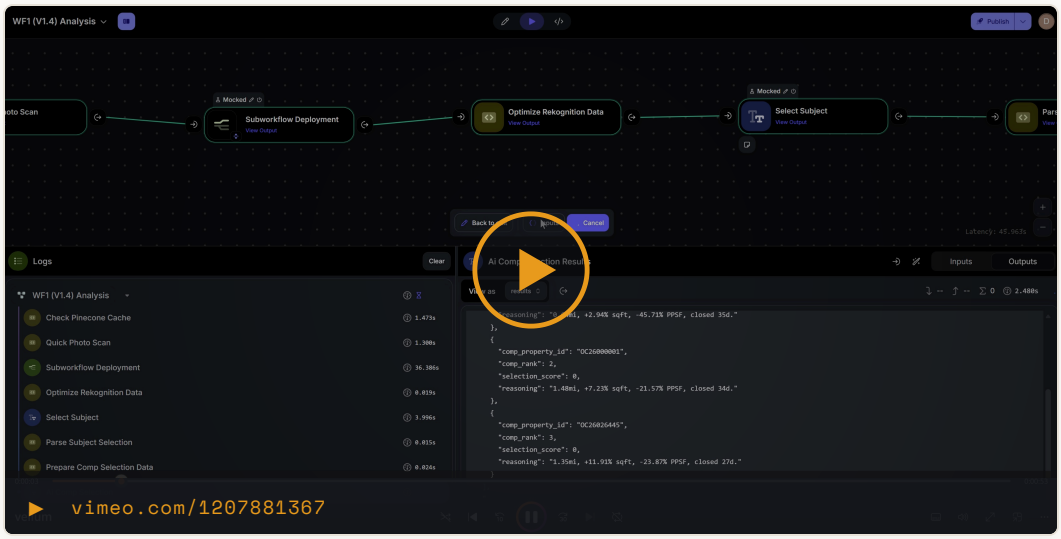

The Vellum pipeline running live: per-node timings, comp selection, and reasoning output.