

## 1 - The feature menu (group by job, then pick features)

<b>Prepare the catalog</b>	Per-title encoding, upscaling, thumbnails, chapters. Once at ingest.
<b>Find &amp; understand</b>	Captions, dubbing, translation, summaries, search. Mostly ingest.
<b>Keep it safe</b>	Moderation, AI-content disclosure, age rating. The gate before publish.
<b>Grow watch time</b>	Recommendations, personalized thumbnails, ad targeting. Continuous.

## 2 - When does it run? (the decision that drives cost)

Once at ingest ... encoding, captions, dubbing, moderation. Paid once, reused forever.  
Per view ..... on-the-fly summary, clip, search. Cost multiplied by your audience.  
Continuous ..... recommendations, ad targeting. Always-on serving you operate.  
Rule: push everything you can to ingest - the single biggest OTT AI cost lever.

## 3 - Build the platform: rent vs assemble vs build

Rent managed platform . ~80% of features in 8-10 wks. AI & data live with the vendor.  
Assemble hyperscaler .. cloud media + AI APIs. Control for a small video team; pay per use.  
Build open source+APIs 14-22 wks. Data stays inside; 30-50% cheaper at scale; you run it.

## 4 - The cost math (ingest-once beats per-view)

1,000-hr catalog x \$1/hr: once at ingest = \$1,000 total, forever.  
Same feature per view at 100k views/mo = \$100,000/mo. Same output, 100x the bill.

## 5 - Compliance gate - every asset passes it at ingest (context, not legal advice)

- Moderate every upload: AI does the first pass (~1/30-1/100 of human cost); humans judge edge cases.
- Disclose AI-generated or AI-altered video - EU AI Act Article 50 makes labelling an obligation.
- Attach a C2PA content credential to synthetic or AI-edited assets where you can.
- Add captions for accessibility - AI caption generation makes this cheap; many markets require it.
- Run the whole gate at ingest, before publish - fixing a live title costs far more.