Transcoding-Farm Cost Worksheet

Run the managed-vs-self-hosted break-even for your own catalog before you commit a pipeline.

Managed cloud

Bills per minute of OUTPUT, so a 6-rung ladder is 6x the source minutes. ~$0.0075 SD to $0.06 UHD by tier. No idle cost.

Self-hosted (spot)

Bills per machine-hour. Spot is ~75% off on-demand. You run the queue; chunk jobs so a 2-min reclaim re-runs.

CPU vs GPU

CPU compresses smaller (lower egress on every view). GPU is faster and denser. Head to CPU, long tail to GPU.

Encode cost is paid once. Egress is paid on every view. Let demand decide which one you minimize per title.

Estimate your monthly transcode bill, then compare E with H

Step	Your number	Formula / note
A · Catalog-hours / month	_____	new hours you encode each month
B · Rungs in the ladder	_____	renditions per title (e.g. 6)
C · Output-hours = A x B	_____	what a managed service bills for
D · Managed $/output-min	_____	blend $0.0075 SD to $0.06 UHD
E · Managed bill = C x 60 x D	_____	per month; zero fixed cost
F · Machine-hours (self-host)	_____	~A if one ladder per machine-hour
G · Spot $/machine-hour	_____	~$0.10 on the spot market
H · Self-host = F x G + ops	_____	add engineering / queue ops time
Verdict = lower of E and H	_____	self-host wins past the break-even

Before you commit the pipeline

- ☐ Output, not input. You multiplied source hours by ladder rungs BEFORE applying the per-minute managed rate.
- ☐ Break-even checked. You compared the managed bill (E) against self-host plus ops (H), not encode price alone.
- ☐ Spot resilience. Jobs are chunked so a two-minute spot reclaim re-runs the chunk instead of failing the title.
- ☐ Demand routing. Popular titles get byte-efficient CPU encodes; the long tail gets fast GPU throughput.
- ☐ Re-encode budgeted. Adding a codec (AV1) or a device tier means encoding the whole catalog again.
- ☐ Live budgeted apart. Any 24/7 channel is a per-channel-hour bill (~$550-$1,250/mo HD), separate from VOD.