Transcoding-Farm Cost Worksheet

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

Managed cloud	Self-hosted (spot)	CPU vs GPU
Bills per minute of OUTPUT, so a 6-rung ladder is 6× the source minutes. ~$0.0075 SD to $0.06 UHD by tier. No idle cost.	Bills per machine-hour. Spot is ~75% off on-demand. You run the queue; chunk jobs so a 2-min reclaim re-runs.	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.**

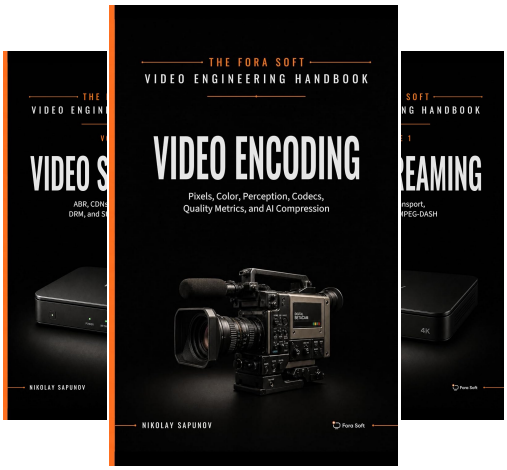
Estimate your monthly transcode bill, then compare E with H

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 × B	_____	what a managed service bills for
D · Managed $/output-min	_____	blend $0.0075 SD to $0.06 UHD
E · Managed bill = C × 60 × 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 × 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.

Companion to: 'The Transcoding Farm: Cloud, On-Prem, and the Cost of Compute'
www.forasoft.com/learn/ott/articles-ott/transcoding-farm-cloud-onprem-cost



GO DEEPER · THE 7-VOLUME HANDBOOK

Fora Soft Video Engineering Handbook

Seven volumes, 9,000+ pages: encoding, streaming, audio and AI for video products — from the pixel to production.

By Nikolay Sapunov, CEO at Fora Soft · Kindle & paperback on Amazon

Browse the handbook on Amazon →

This worksheet runs the managed-versus-self-hosted break-even for your own catalog. We wrote seven volumes on the stack behind that bill. Volume 1 covers encoders, codecs and rate control, Volume 3 bitrate ladders, CDNs and the economics of delivery, and Volume 2 packaging into HLS, DASH and CMAF.

WHAT YOU WILL FIND INSIDE

- Encoder and codec choices that set egress
- Economics of video delivery
- Bitrate ladders per title
- Packaging for HLS, DASH and CMAF

THE FULL HANDBOOK · 7 VOLUMES · 9,000+ PAGES · ALL ON AMAZON

Vol. 1 · Video Encoding — pixels and color to H.264, HEVC, AV1, quality metrics and AI
Vol. 2 · Video Streaming I — network transport, contribution and ingest, HLS, MPEG-DASH, CMAF
Vol. 3 · Video Streaming II — ABR, CDNs, players, WebRTC, DRM and the economics of delivery
Vol. 4 · Audio for Video I — PCM, codecs, loudness, accessibility, real-time voice processing
Vol. 5 · Audio for Video II — jitter buffers, lip sync, Atmos, quality metrics, neural codecs
Vol. 6 · AI for Video I — models and inference, computer vision, speech, search, generative video
Vol. 7 · AI for Video II — real-time pipelines, agents, serving, cost, compliance, verticals

Should your OTT catalog move from a managed encoder to a self-hosted farm?

We will model your catalog, ladder and views, compare managed and spot-based farms, and design a chunked, reclaim-safe pipeline if self-hosting wins. 250+ video products shipped since 2005.

Book a 30-min call → **WhatsApp →** **Email us →**