

Companion to article 6.4 of Block 6. Print on A4.

1. The two ideas

Family	Input	Output	Typical saving
Content-aware	Source video	Bit allocation	20-50% per title
Context-aware	Audience stats	Ladder rung set	30-50% storage

2. Three depths of content-aware encoding

Depth	Granularity	Tech	Pioneer	Saving
Per-title	Whole video	Convex hull, VMAF	Netflix 2015	20-50%
Per-scene	Per shot (4 s)	Lagrangian opt.	Netflix 2018	+15-30%
ROI	Per region	Face/object QP	Hikvision/Axis	+3-5 dB face

3. Which to add first - decision tree

- Step 1. Still shipping Apple HLS default ladder? -> add per-title first.
- Step 2. Catalogue has big complexity swings inside titles? -> add per-scene next.
- Step 3. Face / object reliably in frame (conferencing, surveillance)? -> add ROI.
- Step 4. More than 100k monthly viewers? -> add context-aware ladder design.
- Step 5. Re-measure audience mix quarterly and rebuild the ladder.

4. The four questions for any CAE vendor pitch

1. Content-aware or context-aware? Input is video frames or audience stats?
2. Per-title, per-scene, or ROI? Which granularity is the saving measured at?
3. What is the perceptual metric - VMAF, SSIMULACRA2, or PSNR? Never trust PSNR for CAE.
4. What is the baseline - Apple HLS default, a tuned fixed ladder, or per-title?

5. Ladder-design checklist (context-aware)

- Measured device class distribution (mobile / tablet / PC / TV).
- Bandwidth CDF from real player logs, not from a CDN estimate.
- Codec capability matrix per device class (H.264 / HEVC / AV1).
- Target rung count 4-6 (down from Apple-default 10-12).
- Rungs concentrated near audience-bandwidth modes.
- Feature-flag rollout with a fallback path to the old ladder.
- Quarterly re-measurement and ladder rebuild scheduled.

6. Rules of thumb

- Per-title saves money on your top 10% of titles by viewership.
- Per-scene is wasted on talking-head content with uniform complexity.
- ROI assumes the viewer's eye is in a known region - validate per vertical.
- Context-aware ladders only pay off above ~100k monthly viewers.
- AI helps content-aware (saliency, complexity), not context-aware (classical opt).