

Companion to article 6.5 of Block 6. Print on A4. Hand to procurement or run yourself.

1. The eight-step recipe

#	Step	What to lock down	Output
1	Corpus	6-12 clips: talking head, sports, cartoon, dark, nature	YUV 4:2:0 sources
2	Encoders	2-6 encoders, 1-2 presets each, locked	Preset list
3	Rate points	4-6 CRFs or fixed bitrates per encoder x clip	CRF/bitrate grid
4	Encode	Log wall-clock time + output bitrate	Encode logs
5	Measure VMAF	libvmaf, official model (vmaf_v0.6.1 or phone/4K)	Per-frame VMAF
6	Plot RD	Log-x bitrate, linear-y VMAF, one curve per encoder	RD curves
7	BD-rate	Bjontegaard cubic spline, per encoder pair, per clip	BD-rate %
8	Aggregate	Report mean AND per-clip range, not just the mean	Final number

2. Eight ways benchmarks lie - rule each one out

- [] One clip only -> always 6+. Per-clip BD-rate range tells you more than the mean.
- [] One preset -> match preset speeds across encoders before comparing quality.
- [] No quality control -> bitrate without VMAF is meaningless.
- [] Arithmetic-mean VMAF only -> also report harmonic mean + 1st percentile.
- [] Default encoder settings -> use published reference configs (tune=0, lookahead=120, ...).
- [] Wrong VMAF model -> 1080p TV (default) vs phone vs 4K shift VMAF by 3-5 points.
- [] Rate-control mode mismatch -> CRF vs CBR vs capped-CRF are not like-for-like.
- [] Cherry-picked aggregation -> median hides the worst case, mean hides variance.

3. The reporting table you must produce

Encoder	Preset	BD-rate vs ref	1080p fps	Watts/stream	Verdict
x264	medium	0% (ref)	120	~	Universal
x265	medium	-30...-35%	30	~	OTT today
SVT-AV1	preset 8	-40...-45%	25	~	Live AV1
SVT-AV1	preset 4	-55...-60%	10	~	Premium VOD
vvenc	medium	-55...-65%	2	~	Future-facing

4. VMAF reporting thresholds (apply both, not just one)

- Harmonic mean VMAF ≥ 95 -> top-rung perceptual quality.
- 1st percentile VMAF ≥ 89 -> 99% of frames clear the visible-artefact line.
- Standard deviation reported alongside, not hidden.
- Two encoders at same mean VMAF but different stddev look very different on TV.

5. The five questions for any encoder vendor

1. What corpus did you benchmark on? Was it AOMedia CTC, UVG, MSU, or 'internal'?
2. What is your BD-rate vs x264 medium and x265 medium - on MY content, not yours?
3. What is your 1080p encoded fps on a 16-core CPU at the preset you recommend?
4. Which VMAF model? Mean, harmonic mean, or percentile?
5. Will you run this benchmark live on my reference clip during the procurement call?