

Companion to article 4.13 of Block 4. Print on A4.

## 1. The five modes at a glance

Mode	Holds fixed	Lets vary	Use when
CBR	bitrate	quality	live ingest, conferencing, narrow pipe
VBR (2-pass)	average bitrate	instantaneous bitrate	VOD where target avg matters most
CRF	perceptual quality	bitrate	archival, mastering, no bandwidth cap
ABR (rate-control)	running avg bitrate	instantaneous bitrate	single-pass VOD when you skip 2-pass
Capped CRF	quality + ceiling	bitrate below ceiling	modern streaming, ABR rungs, default

## 2. FFmpeg recipes (one line each, x264/x265/AV1)

CBR x264 (live ingest): `-c:v libx264 -preset veryfast -b:v 4M -maxrate 4M -minrate 4M -bufsize 4M -g 60`  
 2-pass VBR x264 (VOD): `-c:v libx264 -b:v 4M -pass 1/2 -an -f mp4 /dev/null then -pass 2 ... out.mp4`  
 Plain CRF x264 (archive): `-c:v libx264 -preset slow -crf 22`  
 Capped CRF x264 (streaming): `-c:v libx264 -preset medium -crf 22 -maxrate 5M -bufsize 10M`  
 Capped CRF x265 (streaming): `-c:v libx265 -preset medium -crf 26 -x265-params vbv-maxrate=5000:vbv-buffersize=10000`  
 Capped CRF SVT-AV1 (streaming): `-c:v libsvtav1 -preset 6 -crf 30 -maxrate 5M -bufsize 10M`

## 3. VBVB buffer sizing rules

Application	bufsize / maxrate ratio	Why
WebRTC / conferencing	0.25 - 0.5	any buffering shows as audio lag / echo
Live broadcast (Twitch / YouTube)	1.0	low latency, predictable ingest
VOD streaming (HLS / DASH)	2.0	comfortable headroom across scene boundaries
Broadcast contribution (DVB / ATSC)	4.0 - 8.0	high quality, encoder reorders aggressively
Archive / mastering	infinite (no -maxrate)	buffer is irrelevant; CRF dominates

## 4. Decision tree - which mode

Is it live? yes -> latency budget? under 500 ms -> CBR + tiny VBVB (WebRTC). 1-3 s -> CBR @ ingest (Twitch / YouTube).  
 Is it VOD? yes -> ABR ladder? yes -> Netflix scale? yes -> per-shot capped CRF. no -> per-title capped CRF.  
 Single VOD file (not ABR)? yes -> CRF (no cap). Archive grade -> x264 CRF 18 / x265 CRF 22 / SVT-AV1 CRF 26.  
 Surveillance recorder? yes -> capped CRF with low cap, tuned to disk-day budget.

## 5. Common mistakes (and the fix)

bufsize == maxrate -> use 2x for VOD, 1x for live; never less unless you have a latency reason.  
 CRF for an ABR rung with no cap -> use capped CRF with maxrate = rung nameplate.  
 Two-pass on a live stream -> impossible; use capped CRF or CBR for live.  
 Comparing presets at iso-CRF -> always compare at iso-bitrate or iso-VMAF.  
 Setting -b:v and -crf together -> they compete; -crf wins. Pick one or the other.

## 6. What the big players run (2026 snapshot)

Netflix - per-shot capped CRF, VMAF-tuned, ~30% savings vs per-title (Dynamic Optimizer).  
 YouTube - per-title capped CRF for VOD; CBR for live, 2-second keyframes.  
 Twitch - CBR end-to-end, 2-second keyframes, no per-title.  
 Disney+ / Hulu / HBO - per-title capped CRF on AWS Elemental / Bitmovin, x265 or AV1.  
 Zoom / Teams - CBR over WebRTC, ~250-500 ms VBVB, SVC layers on top.