

Streaming Latency Budget — Cheat Sheet

Glass-to-glass = camera lens to viewer's display. Sum of seven contributors.

Version 2026-05-20

The 7 Contributors

1. Capture & ISP	10 - 200 ms	Sensor exposure, image-signal processor, audio capture, USB/HDMI grabber.
2. Encoder	50 - 400 ms	Lookahead + B-frames dominate; LL profile cuts both.
3. Packager	50 ms - 4 s	HLS/DASH segments add seconds; CMAF chunks ~200-500 ms; WebRTC has none.
4. Contribution	20 - 200 ms	Encoder to origin: RTMP, SRT, RIST, WHIP.
5. CDN edge	20 - 200 ms	Warm edge ~40 ms; cache miss ~150 ms; chunked-transfer support matters.
6. Player buffer	0.1 - 30 s	Safety reserve; almost always the largest term.
7. Decoder + display	30 - 100 ms	Hardware decode + frame queue + panel refresh (16.7 ms @ 60 Hz).

2026 Protocol Latency Floors

Protocol	Floor	Typical	Condition
HLS (RFC 8216, 6-s)	18 s	20 - 30 s	Any CDN, defaults
Classic DASH	12 s	15 - 25 s	2-4 s segments
LL-HLS	1.5 s	2 - 5 s	Chunked transfer; tuned player
LL-DASH	1.5 s	2 - 5 s	Chunked CMAF; tuned player
HESP	0.4 s	0.6 - 1.2 s	HESP-aware origin and player
WebRTC	0.2 s	0.3 - 0.8 s	SFU or P2P; tuned jitter buffer
Media over QUIC	0.2 s	0.4 - 1 s	Draft, May 2026

Worked Example


Stadium in London to viewer in Berlin, 1080p sports feed.

Classic HLS, 6-s segments

Camera & ISP	80 ms
Encoder, 10-frame LA	333 ms
Packager, 6 s	6,000 ms
Contribution	100 ms
CDN warm edge	40 ms
Player buffer 3 x 6	18,000 ms
Decoder + display	50 ms
Total	24.6 s

WebRTC, default jitter buffer

Camera & ISP	80 ms
Encoder, no lookahead	100 ms
No packager (RTP)	0 ms
Contribution RTP/SRTP	80 ms
SFU forward	30 ms
Jitter buffer	100 ms
Decoder + display	30 ms
Total	0.42 s



GO DEEPER · VOLUME 2 OF 7

Video Streaming: A Complete Guide to Video Delivery

From Network Protocols and Transport to Ingest, HLS, and MPEG-DASH

By Nikolay Sapunov, CEO at Fora Soft · 1,124 pages · Kindle & paperback on Amazon

Read it on Amazon →

amazon.com/dp/B0HH8F4R82

This sheet adds up the seven glass-to-glass contributors and compares protocol floors from HLS to WebRTC and MoQ. Volume 2 is built around this budget: live, VOD and near-live modes, the latency budget hop by hop, contribution over RTMP, SRT and RIST, and HLS, DASH and CMAF segment delivery.

WHAT YOU WILL FIND INSIDE

- Glass-to-glass latency budget, hop by hop
- Contribution over RTMP, SRT and RIST
- HLS, MPEG-DASH and CMAF segment timing
- QUIC and HTTP/3 in delivery

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

Need to cut glass-to-glass latency without losing stability?

We will measure each contributor in your pipeline and show which protocol and settings reach your latency target. 250+ video products shipped since 2005.

Book a 30-min call →

WhatsApp →

Email us →