The Streaming Pipeline — 9 Stages, End to End

Map of every stage between camera and screen, with typical protocol and latency contribution. 2026.

Fora Soft Learn

The nine stages

Stage	What it does	Latency added	Typical protocol	Common failure
1. Capture	Sensor → raw frame buffer	**+16–33 ms**	—	Bad focus / exposure
2. Encode	Compress with H.264/HEVC/AV1, 2 s GOP	**+~200 ms**	—	Wrong preset / GOP drift
3. Contribution	Ingest to cloud	**+400 ms–4 s**	WHIP / SRT / RTMP	Link drop / uplink starve
4. Transcode	Build bitrate ladder, 5 rungs	**+~500 ms**	—	Software on CPU vs GPU
5. Package	CMAF segments + manifest	**+~300 ms**	CMAF / fMP4	GOP–segment misalign
6. Origin	HTTP storage front	**+20–100 ms**	HTTPS	Window expiry / signed URL
7. CDN	Shield → mid-tier → edge	**+10–500 ms**	HTTPS / HTTP/3	Cache-key mismatch
8. Last-mile	Wi-Fi / mobile / fibre / satellite	**+5–100 ms RTT**	TCP / QUIC	Bufferbloat / handover
9. Player	Fetch, decode, paint	**+2 000–18 000 ms**	HLS / DASH / WebRTC	Default buffer too deep

Three typical configurations and their glass-to-glass total

Configuration	Segment / chunk	Total glass-to-glass	Typical use
Classic HLS	6 s seg, 18 s hold-back	**18–30 s**	OTT VOD-like live
Low-latency HLS / DASH	2 s seg, 200 ms chunks	**2–5 s**	Live sports, live shopping
WebRTC (WHEP / SFU)	no segments, paced RTP	**0.2–0.5 s**	Auctions, telemedicine

Key principle — the player buffer dominates HTTP-based latency

Apple HLS Authoring Spec requires the player buffer be at least 3 × segment duration before starting playback.

A 6 s segment forces an 18 s floor; a 2 s segment with LL-HLS partials drops to 2–5 s; WebRTC removes the buffer entirely.

Optimise the player config first if the goal is low-latency. The earlier stages combined rarely add more than a second.