

## The four failure modes

- 1 Audio priming** AAC adds silent samples (2112 Apple / 2048 FDK / 1024 FFmpeg). fMP4 edit list trims it; MPEG-TS cannot.
- 2 Rendition edits** If some bitrate levels carry an edit/CTS others don't, video shifts vs audio on a quality switch.
- 3 Ad insertion** HLS EXT-X-DISCONTINUITY or new DASH Period resets the clock. New content = new priming = new offset.
- 4 Segment drift** Sample duration shorter than advertised leaves a hole per segment; audio drifts behind over long playback.

## The DASH offset + the tolerance

- DASH offset**  $MSE.timestampOffset = Period@start - Period@presentationTimeOffset$
- The rule**  $presentationTimeOffset = \text{earliest presentation time of the period's first segment, in the track timescale.}$
- Tolerance** Keep audio within +90 ms lead / -185 ms lag of video (ITU-R BT.1359-1).

## Symptom to cause: five checks

- Small constant lag from second one, worse on MSE players than Safari = audio priming in TS or stripped edit list.
- Fine in main content, breaks at the first ad and worsens each break = discontinuity or period-offset error.
- Imperceptible at start, obvious 20 min in = segment-duration rounding accumulating into drift.
- Fine on fast Wi-Fi, breaks only when the network forces a quality switch = inconsistent edits across renditions.
- Encode once with CMAF (ISO/IEC 23000-19): one timeline, edit lists survive, no HLS-vs-DASH divergence.