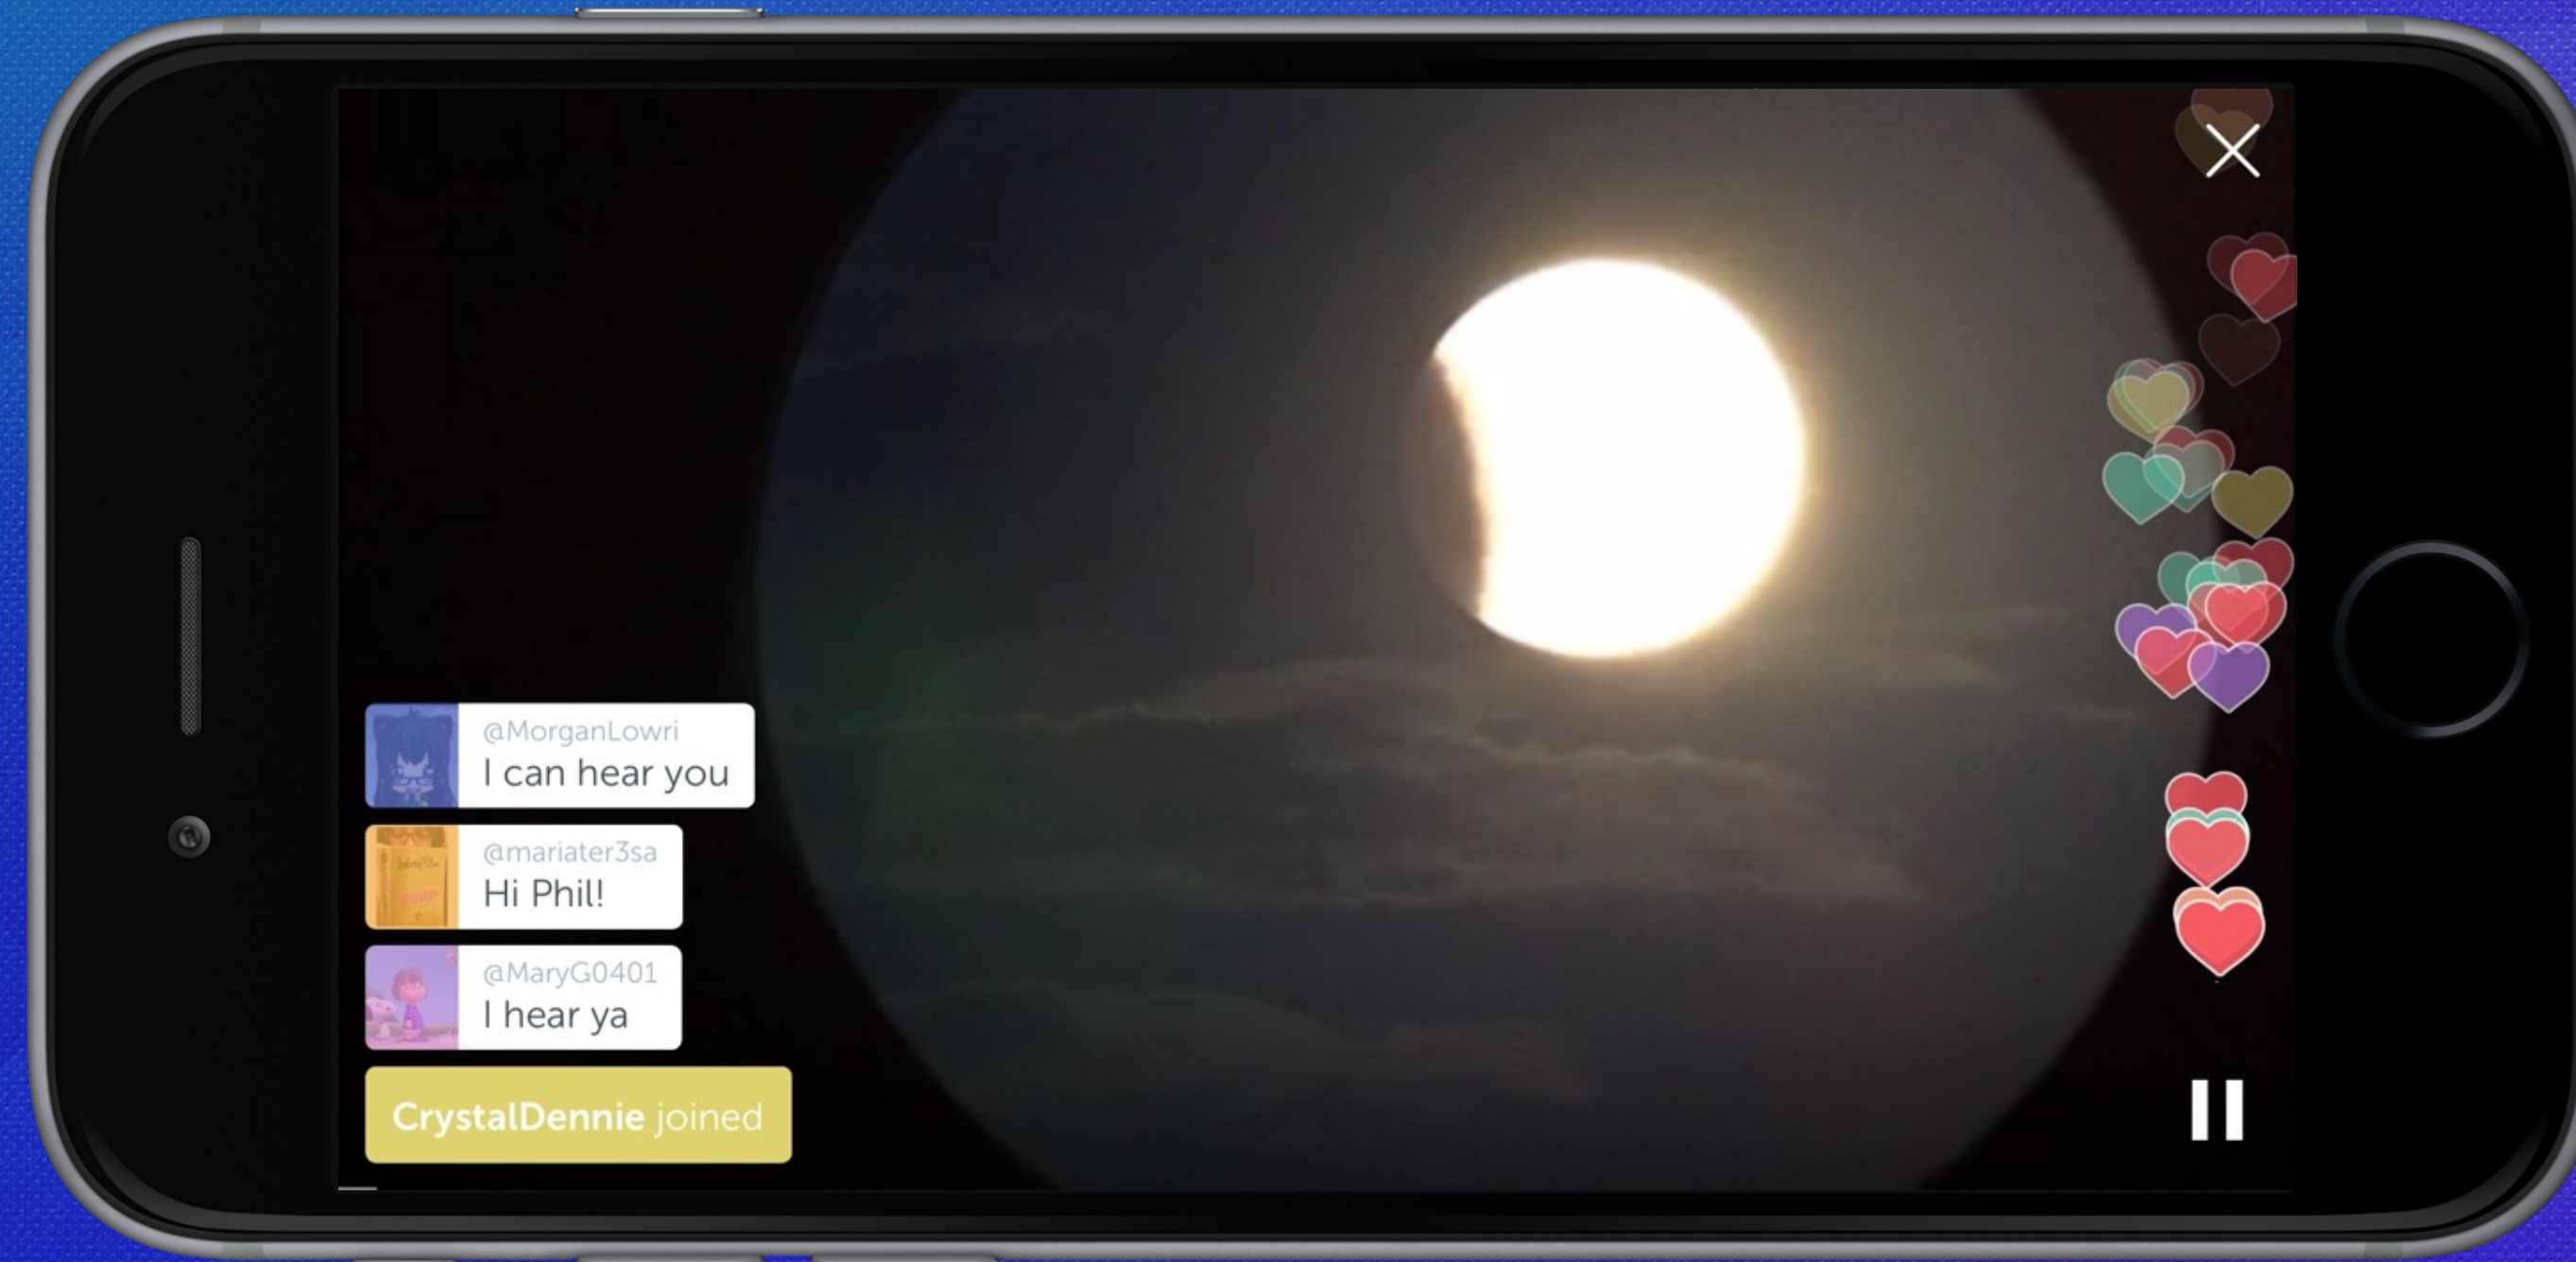


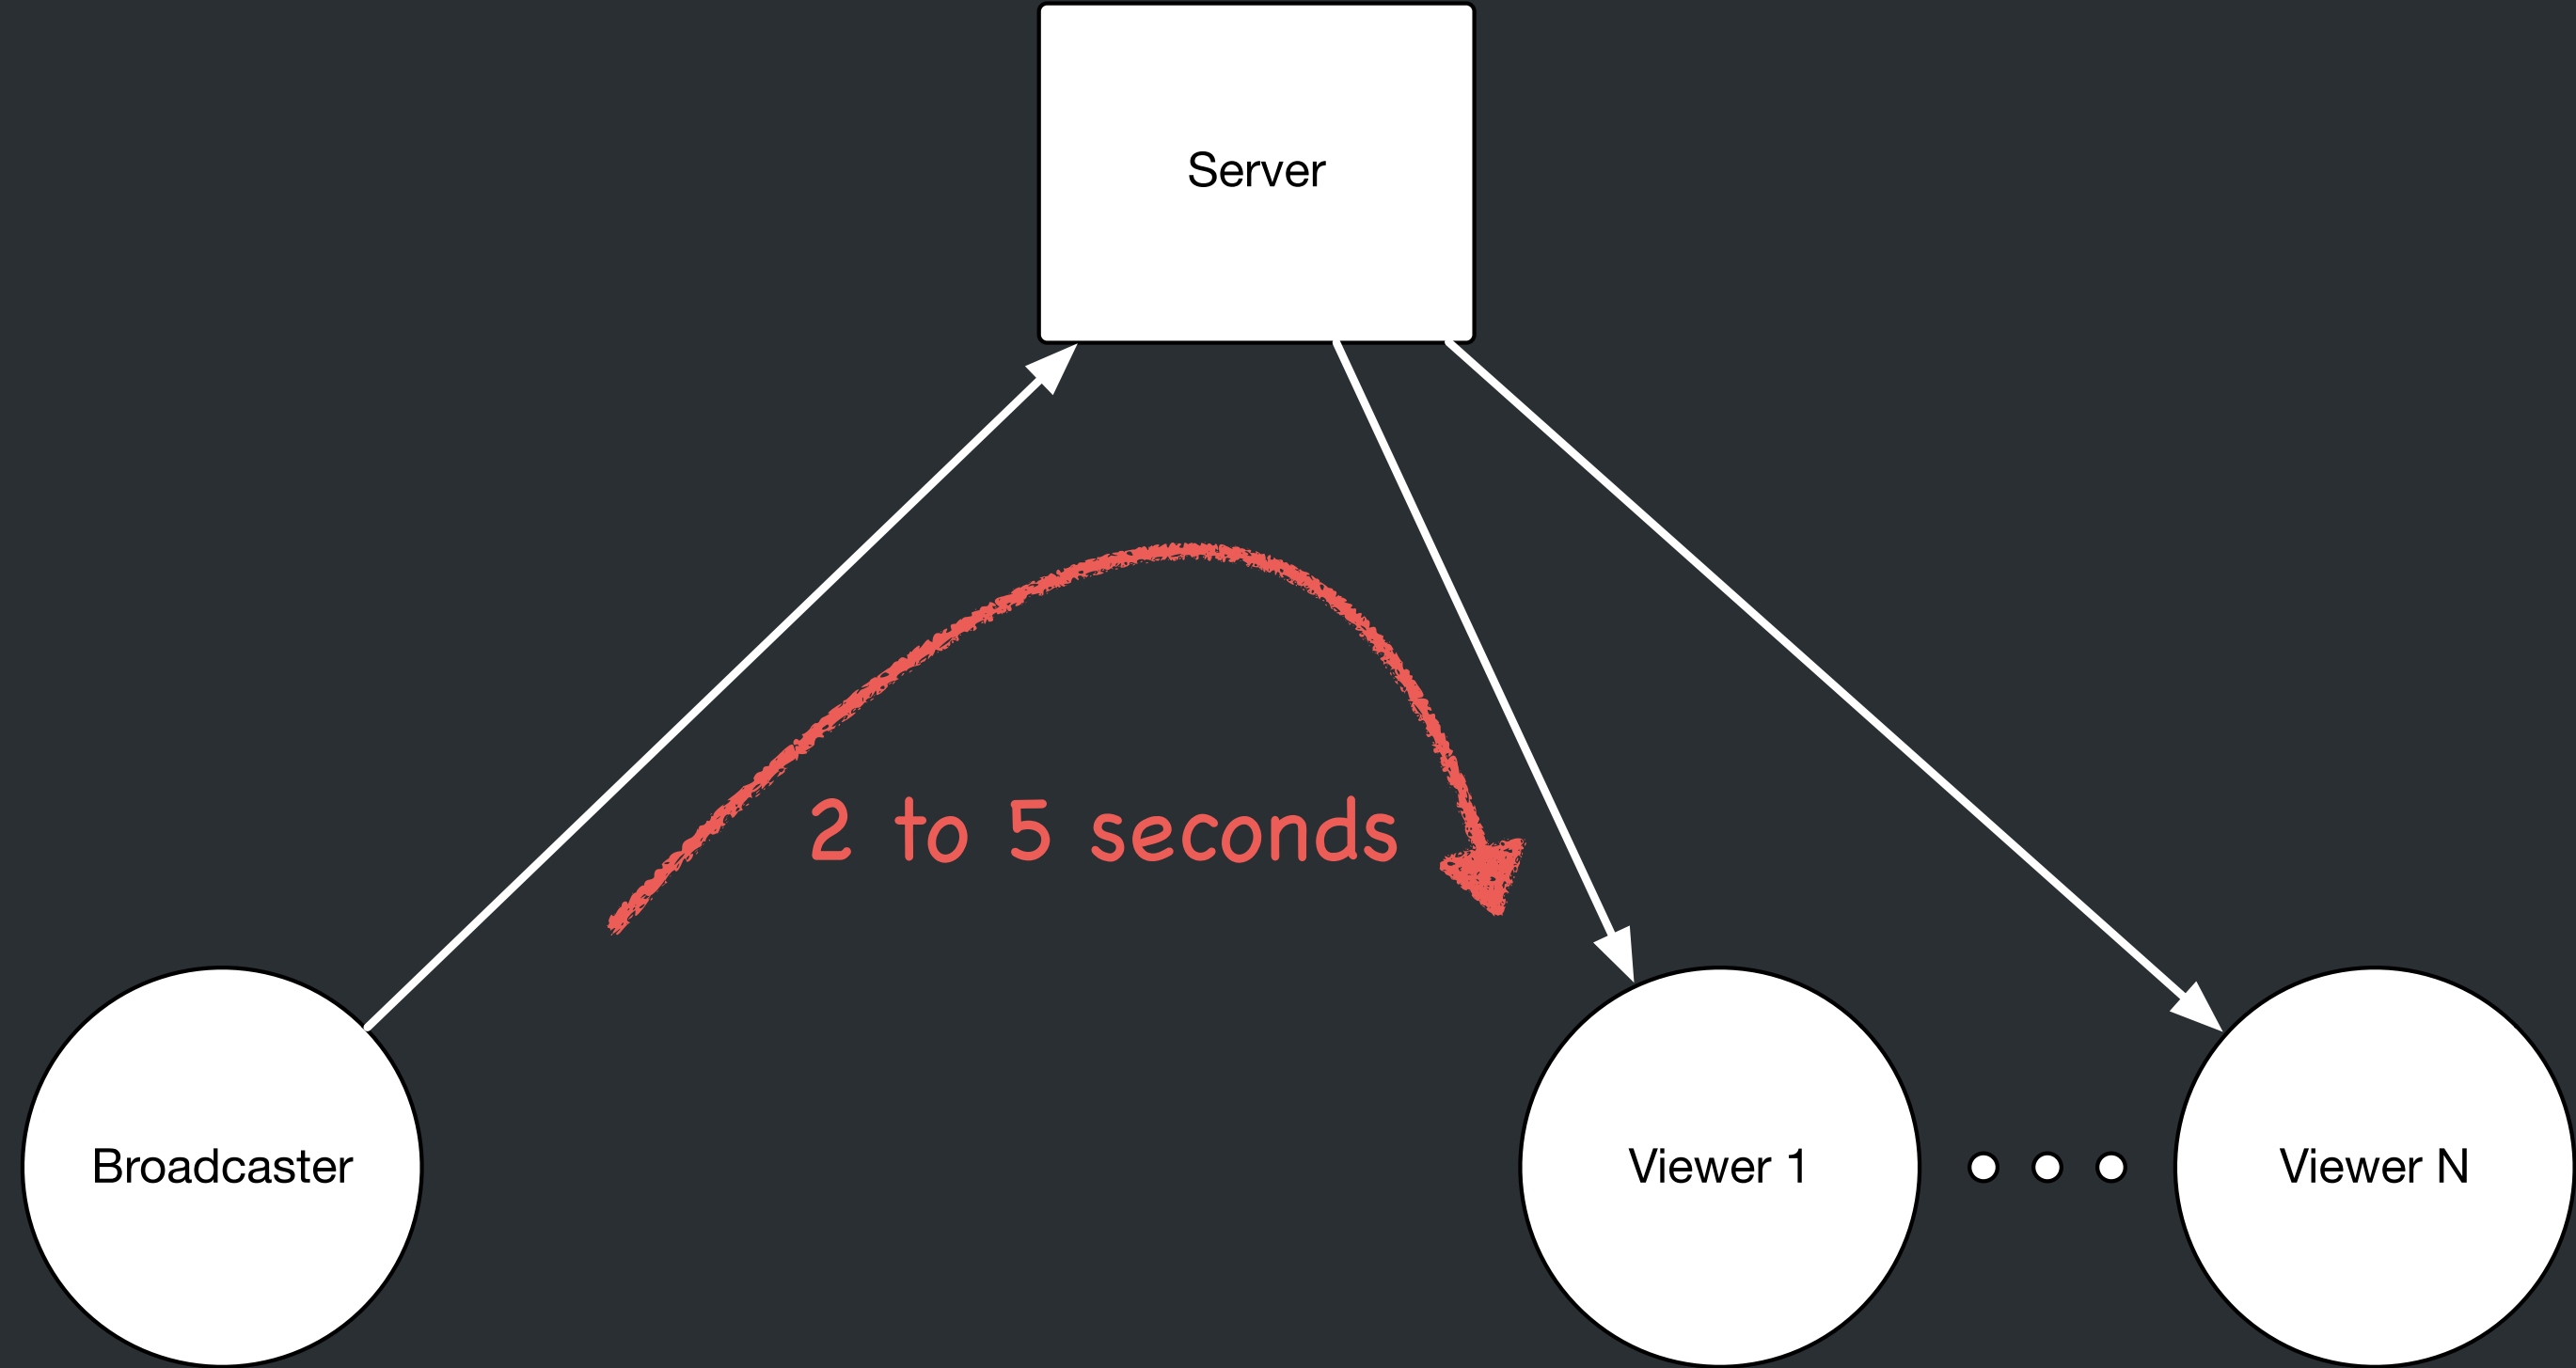
# Periscope LHLS Media Streaming

**Mark Kalman**  
Twitter | Periscope

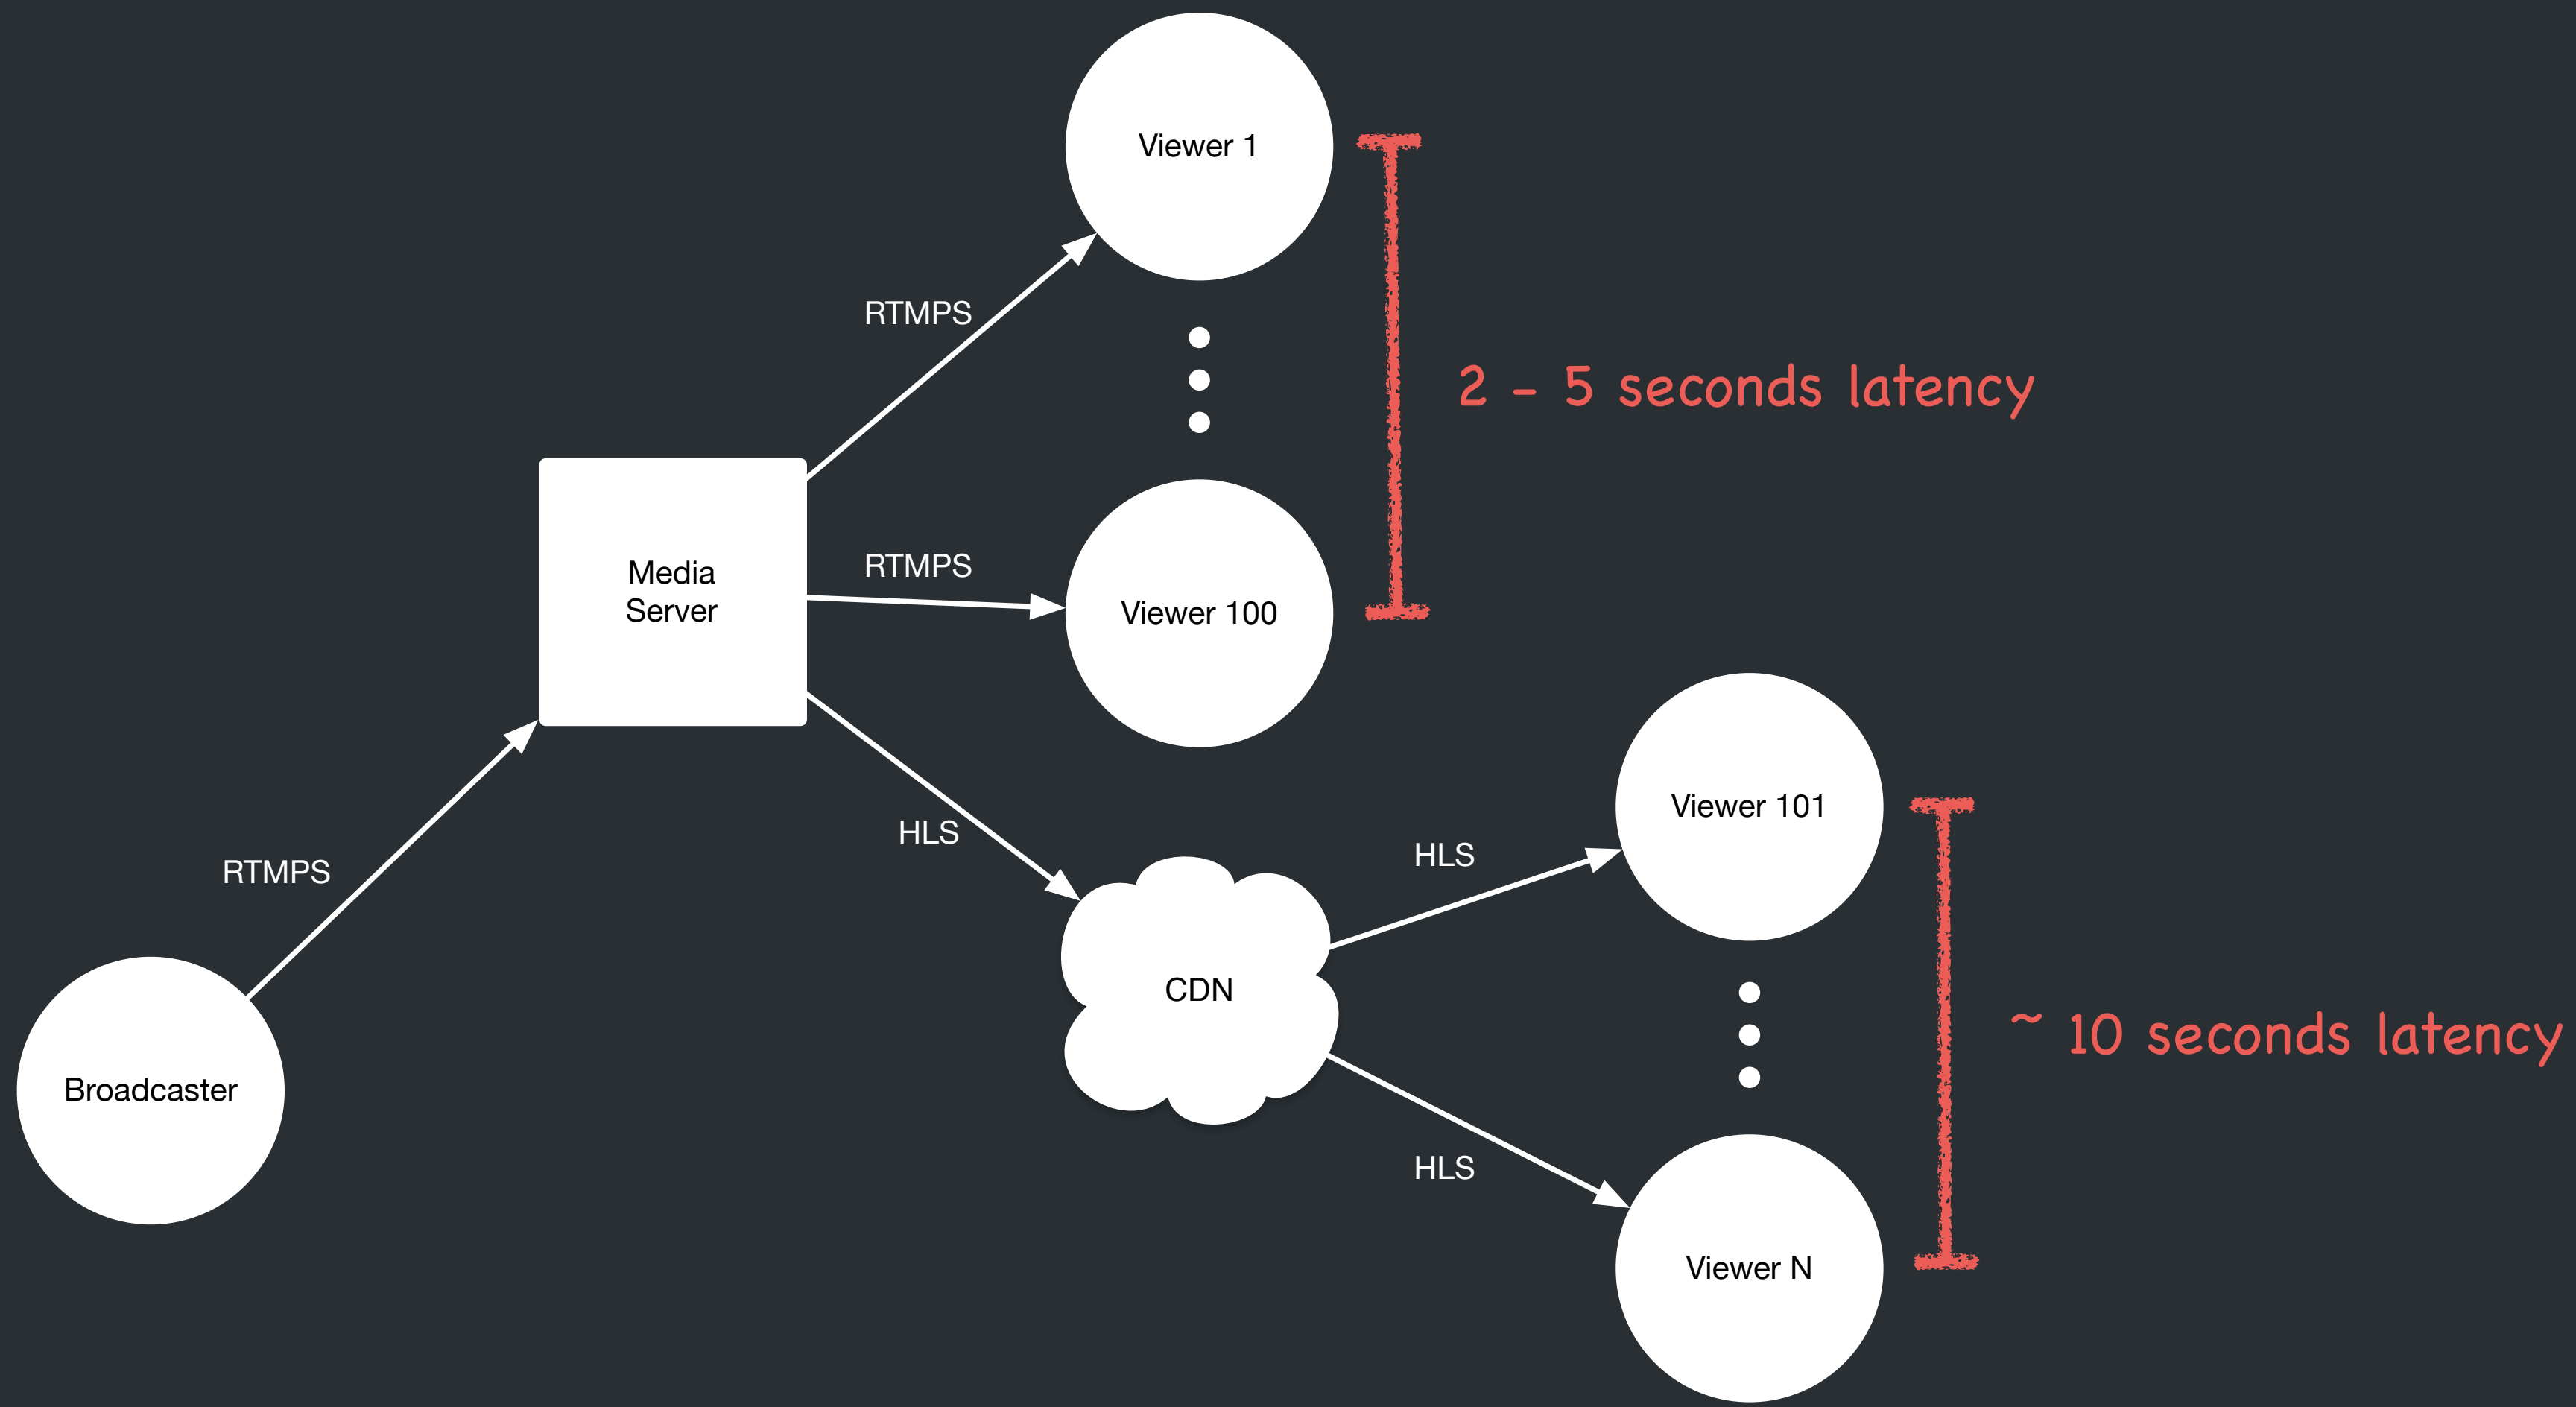
# Interactive Streaming from Mobile Devices



# Latency Target: 2 to 5 s



# Streaming Architecture at Launch



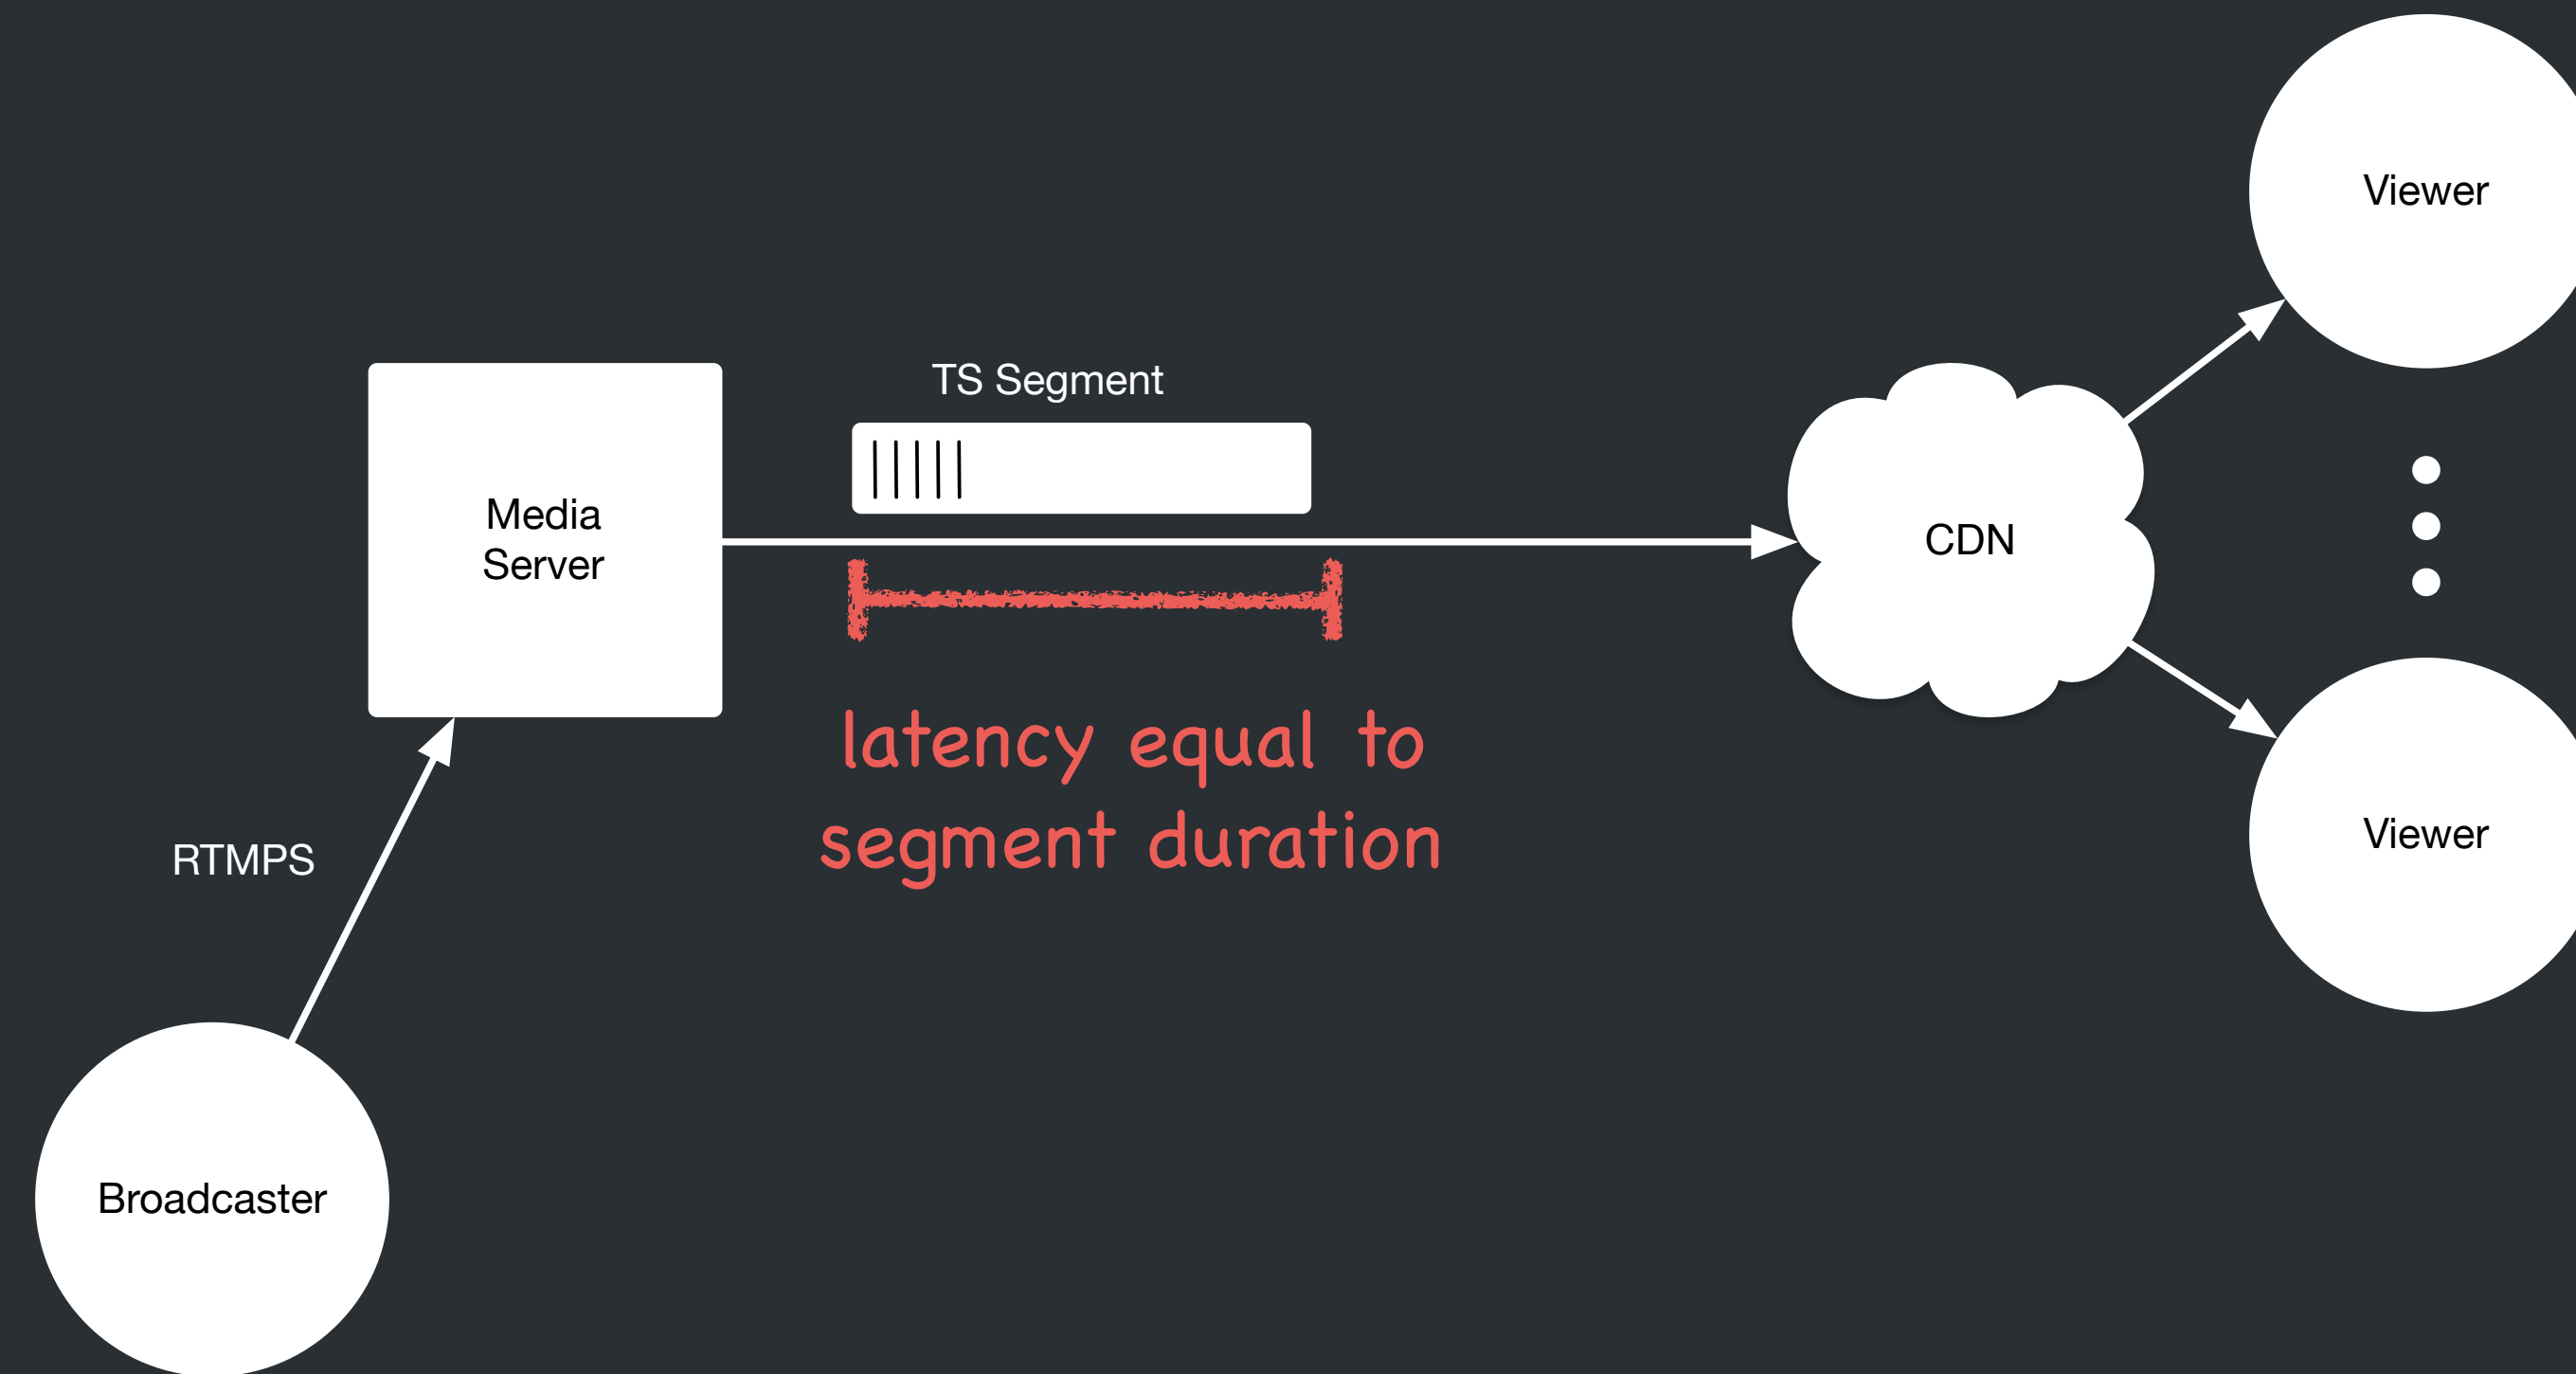
# Downsides Of Hybrid Architecture

- Not all viewers received a low-latency RTMP stream
- Egress bandwidth from AWS for RTMP viewers #1 operational cost
- Long TCP connections between ingress server and viewers

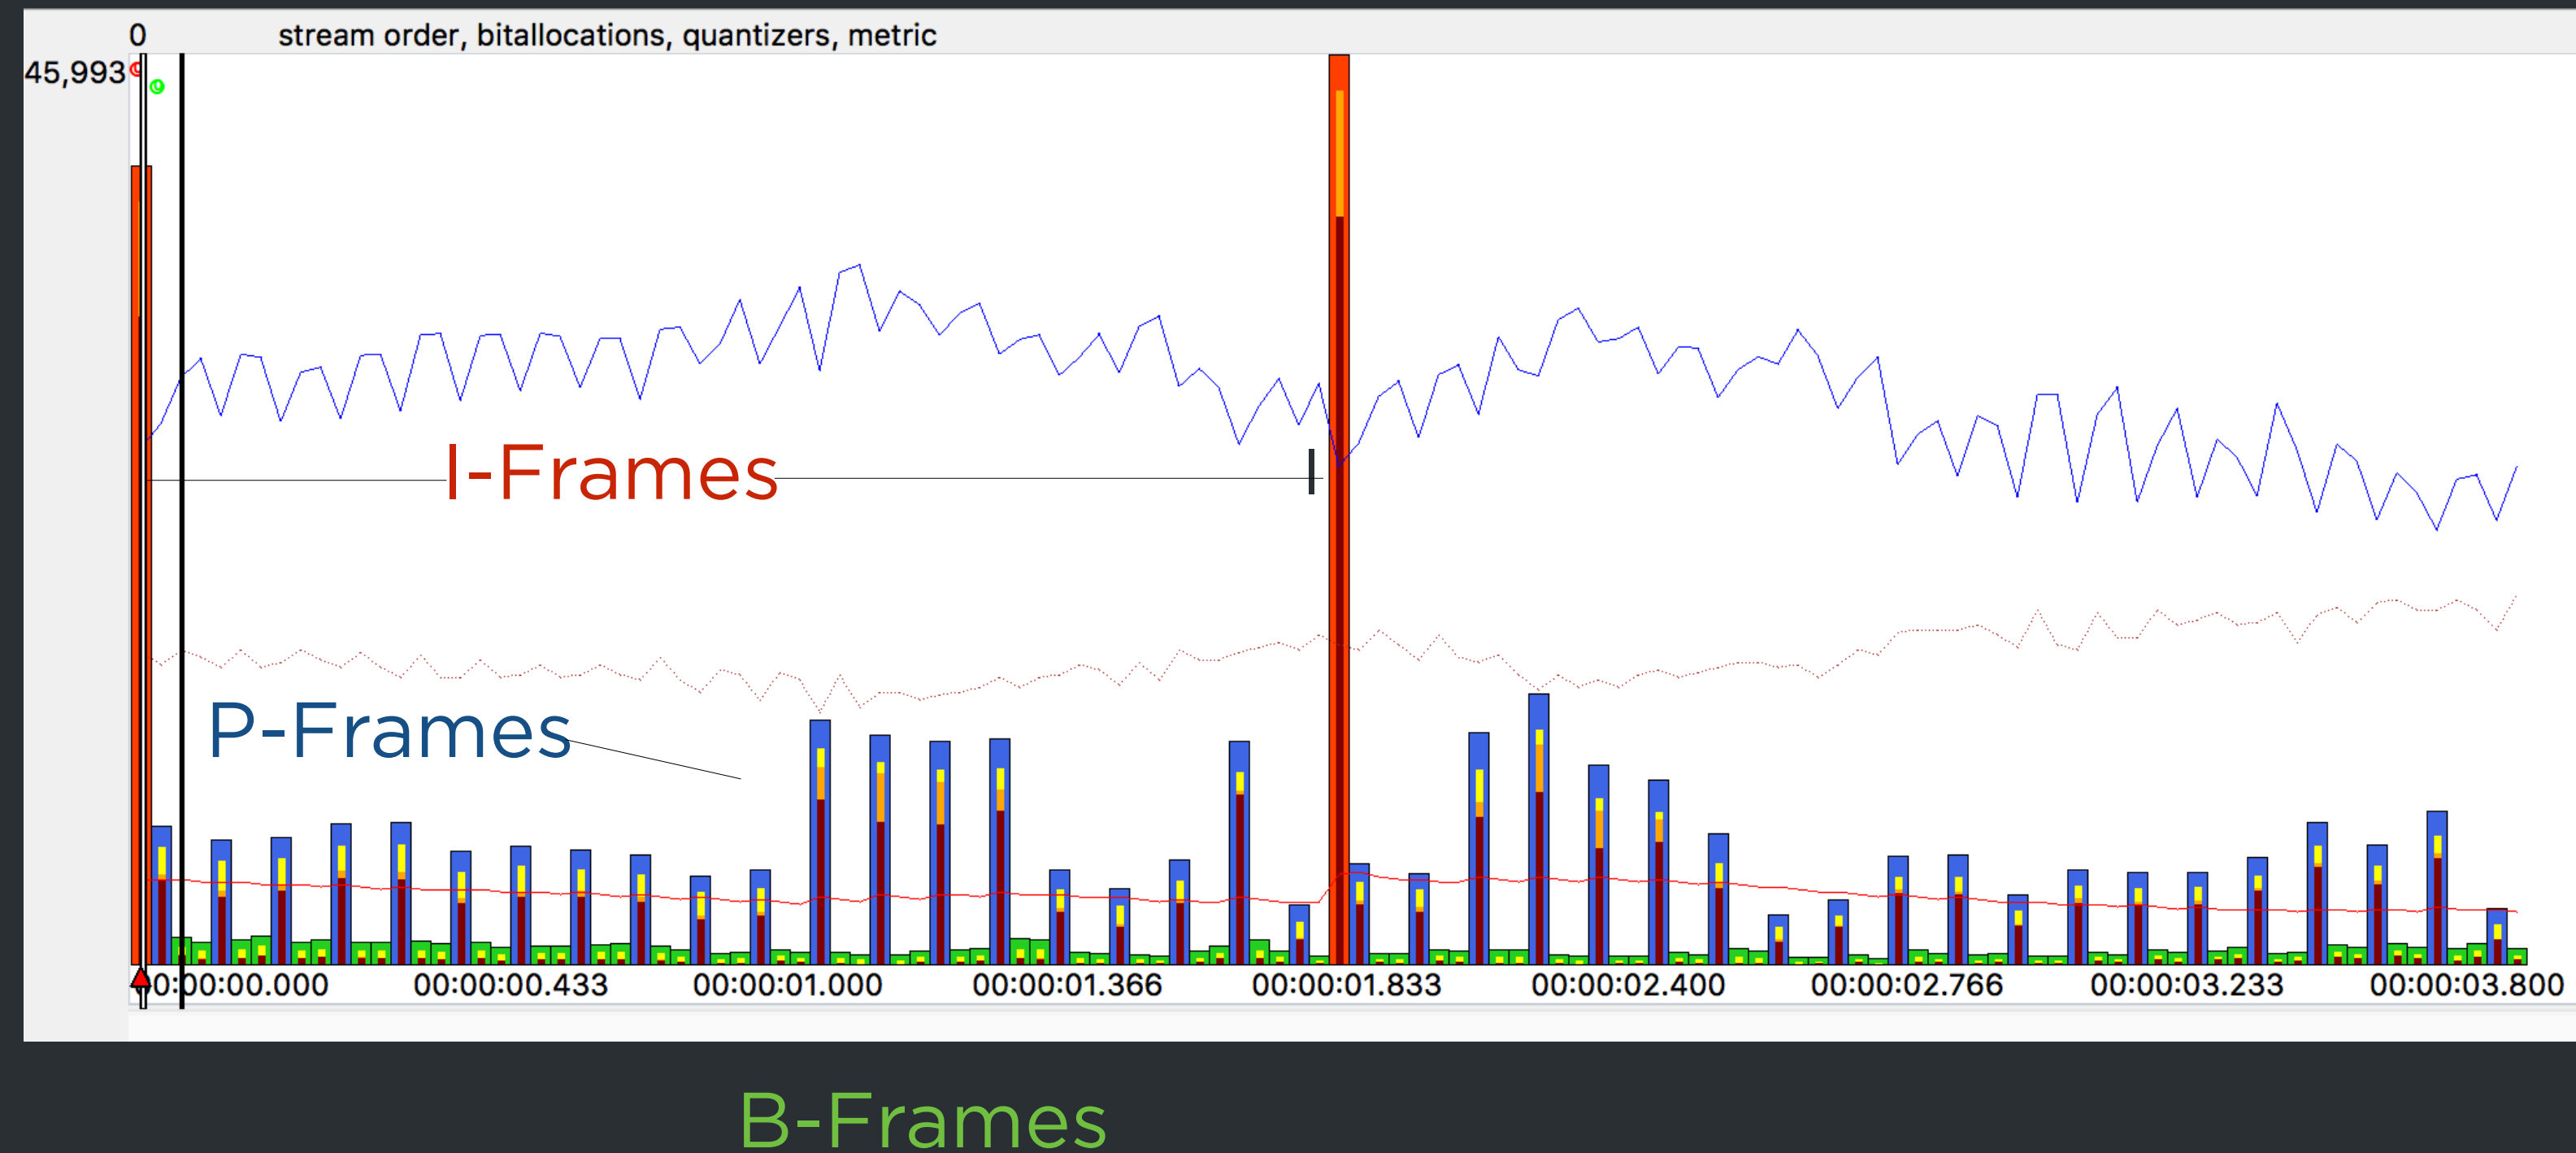


# Sources Of Latency In HTTP-Based Streaming (1)

## Segmentation Latency

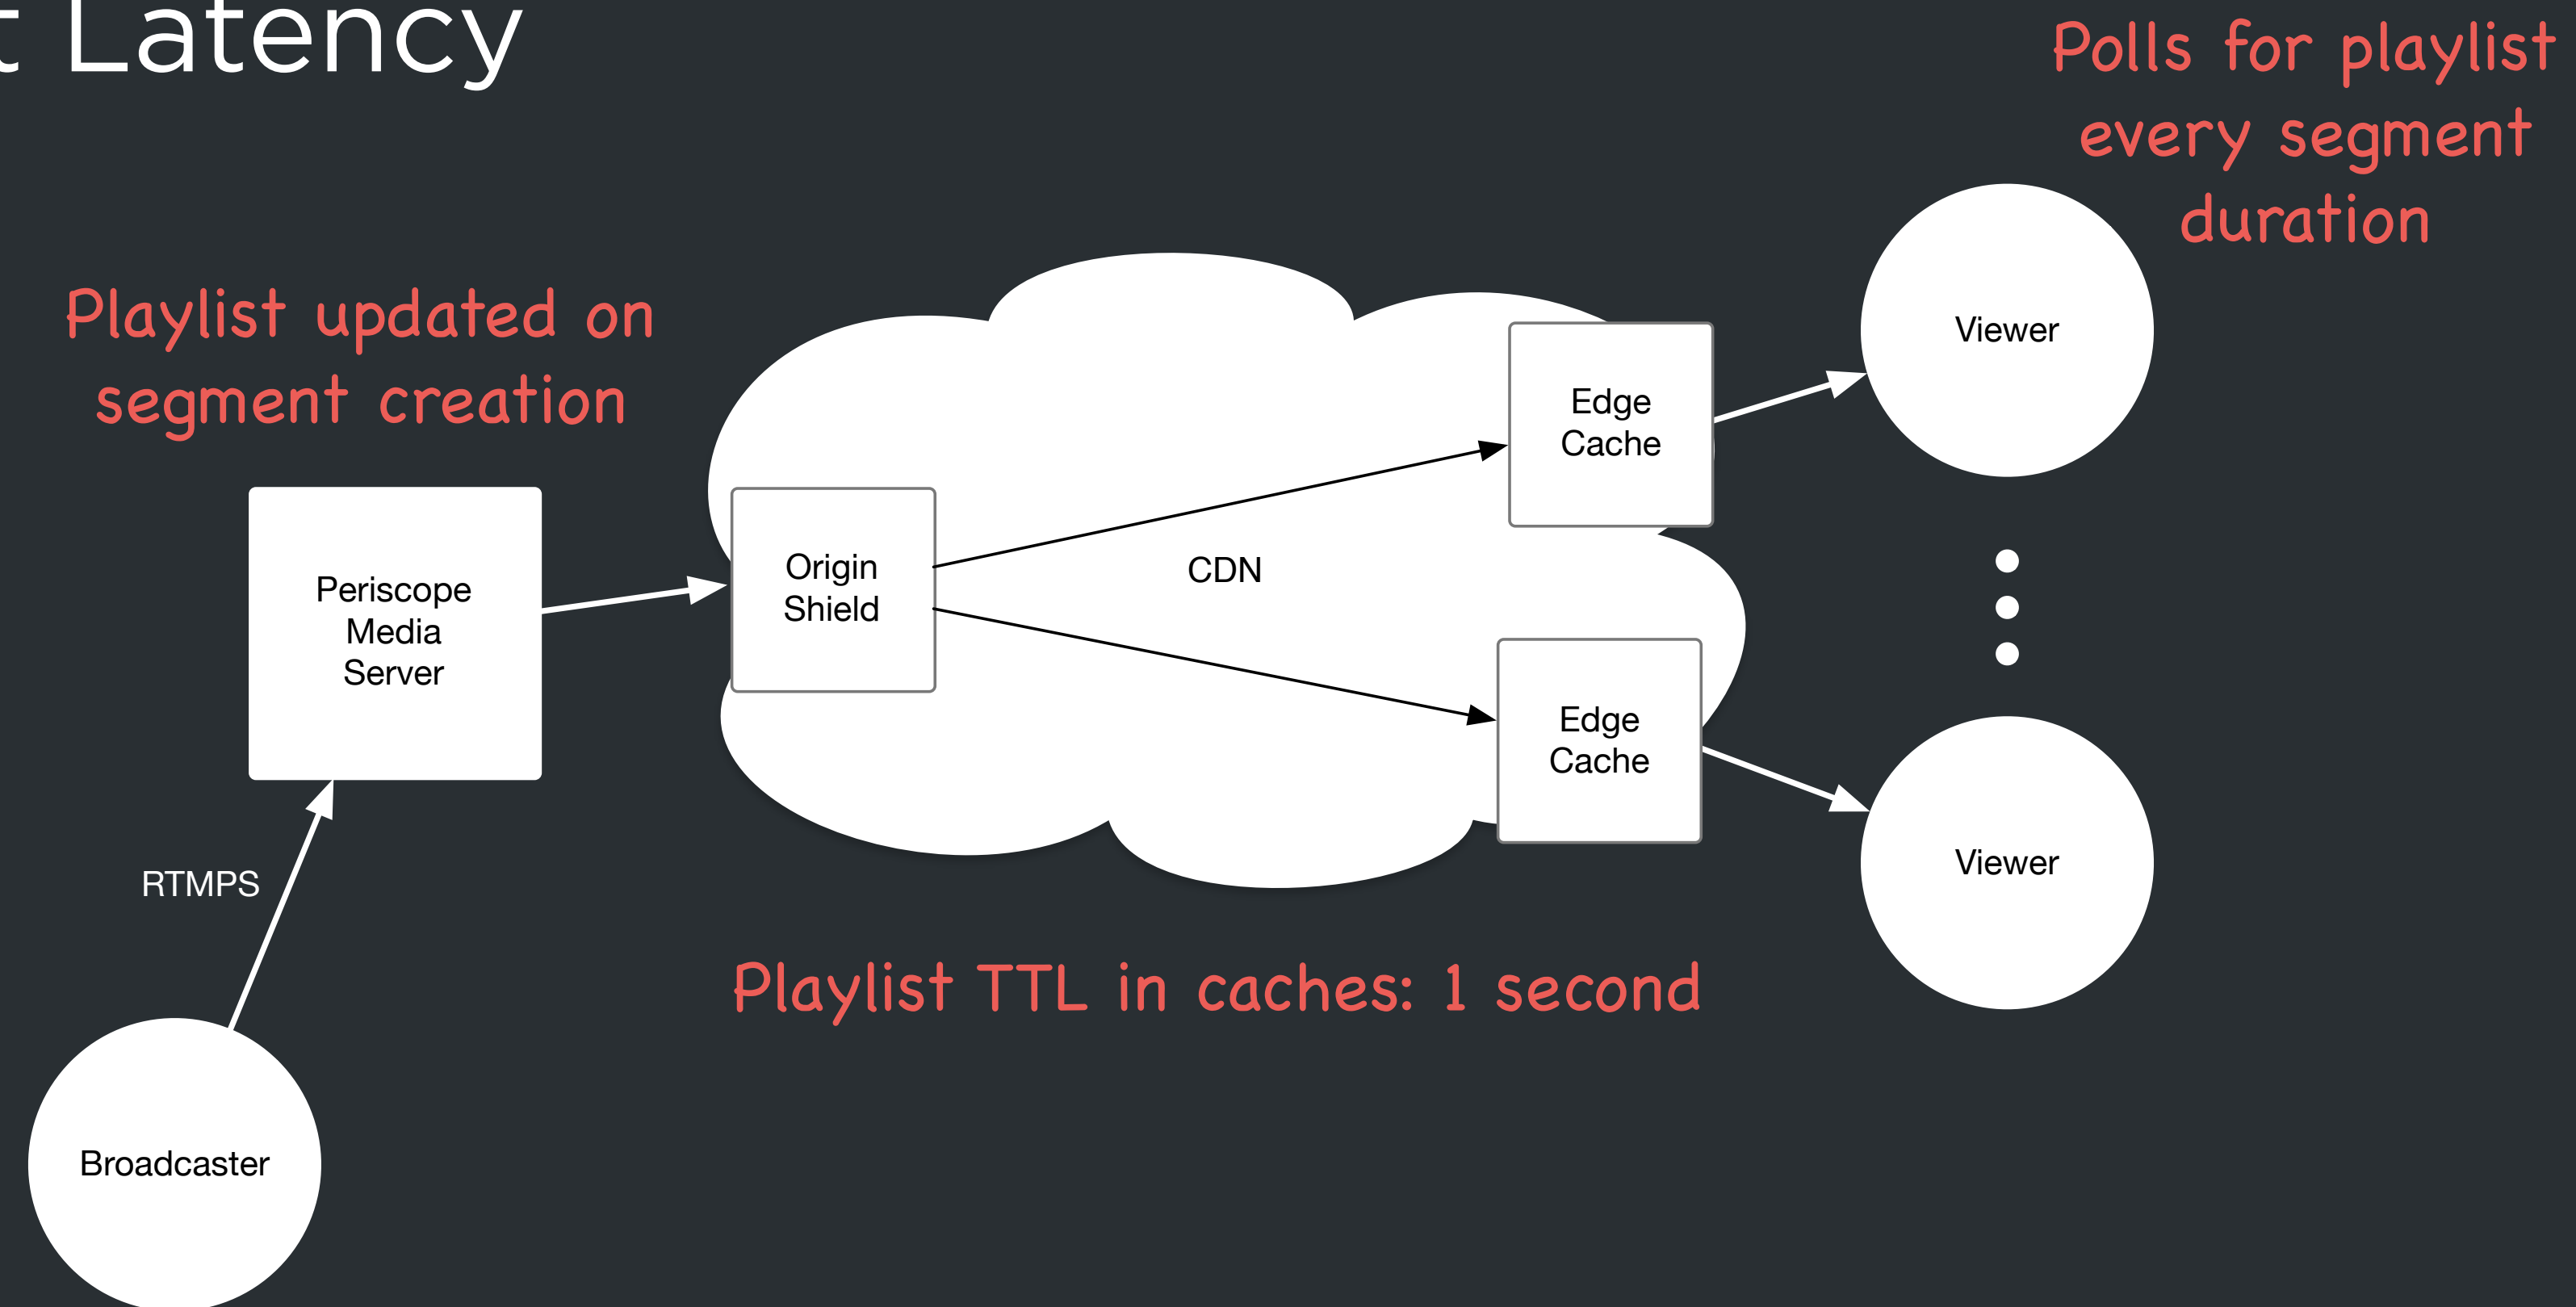


# Why You Can't Just Make Segments Really Short



# Sources Of Latency In HTTP-Based Streaming (2)

## Manifest Latency

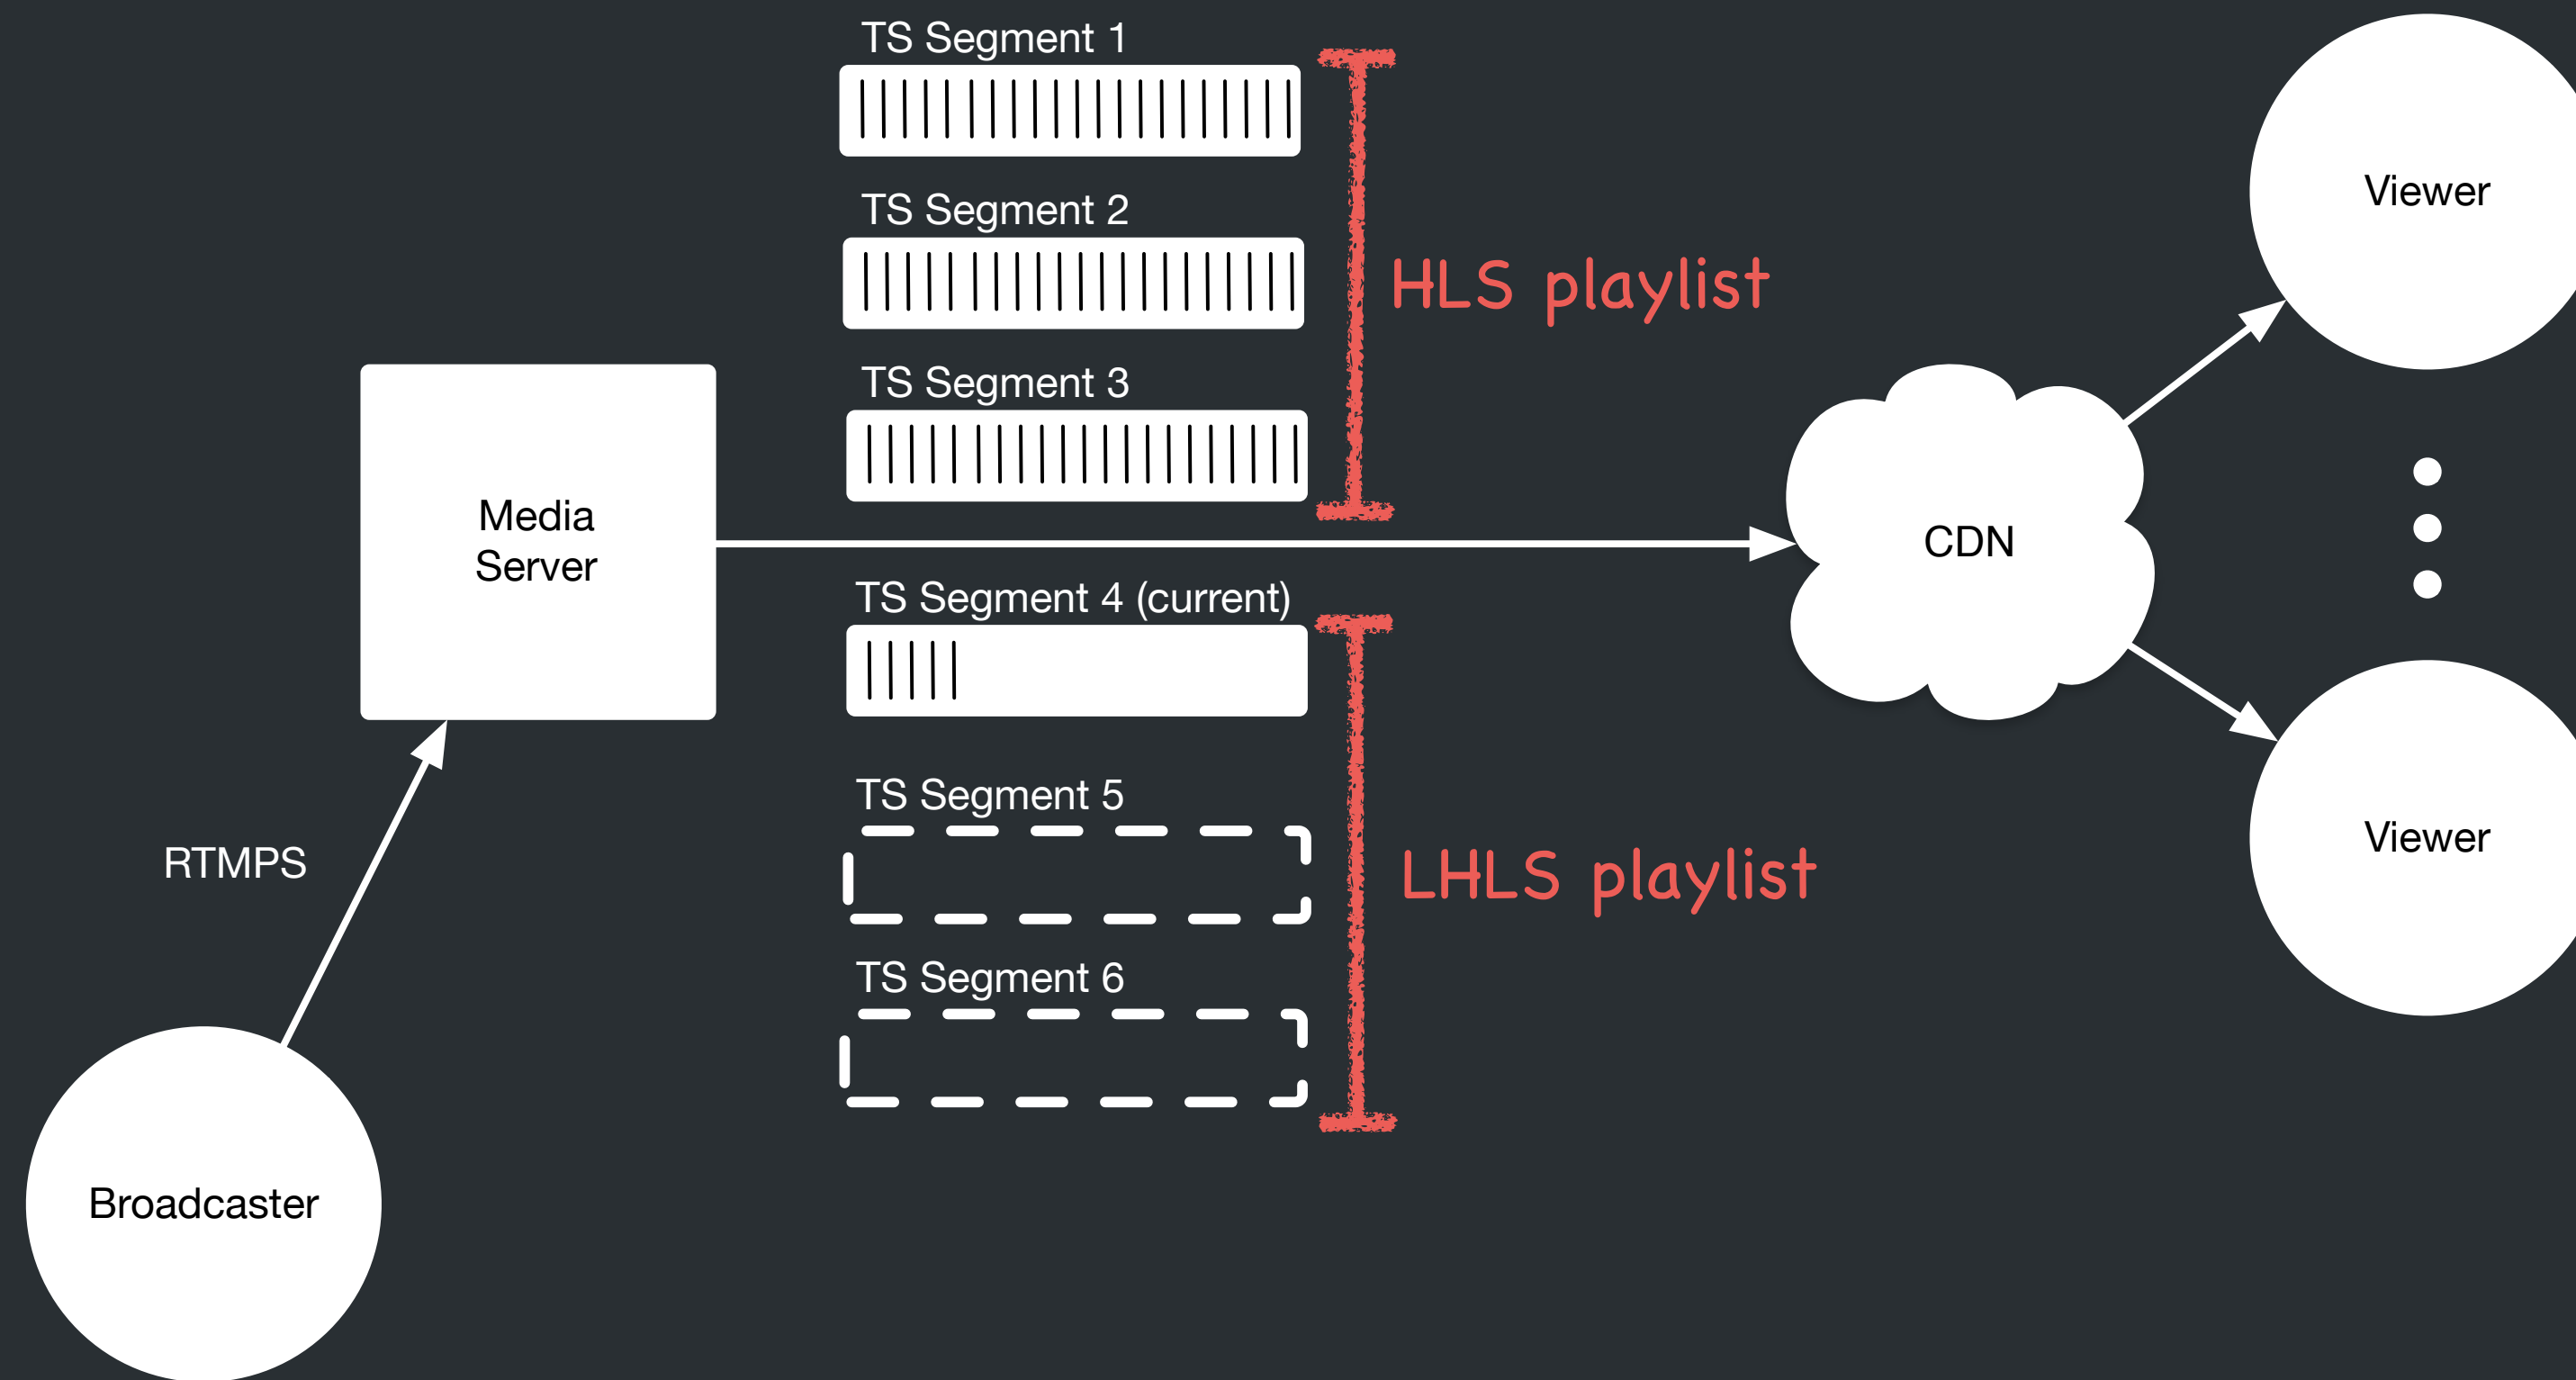


# Eliminating Segmentation And Manifest Latency

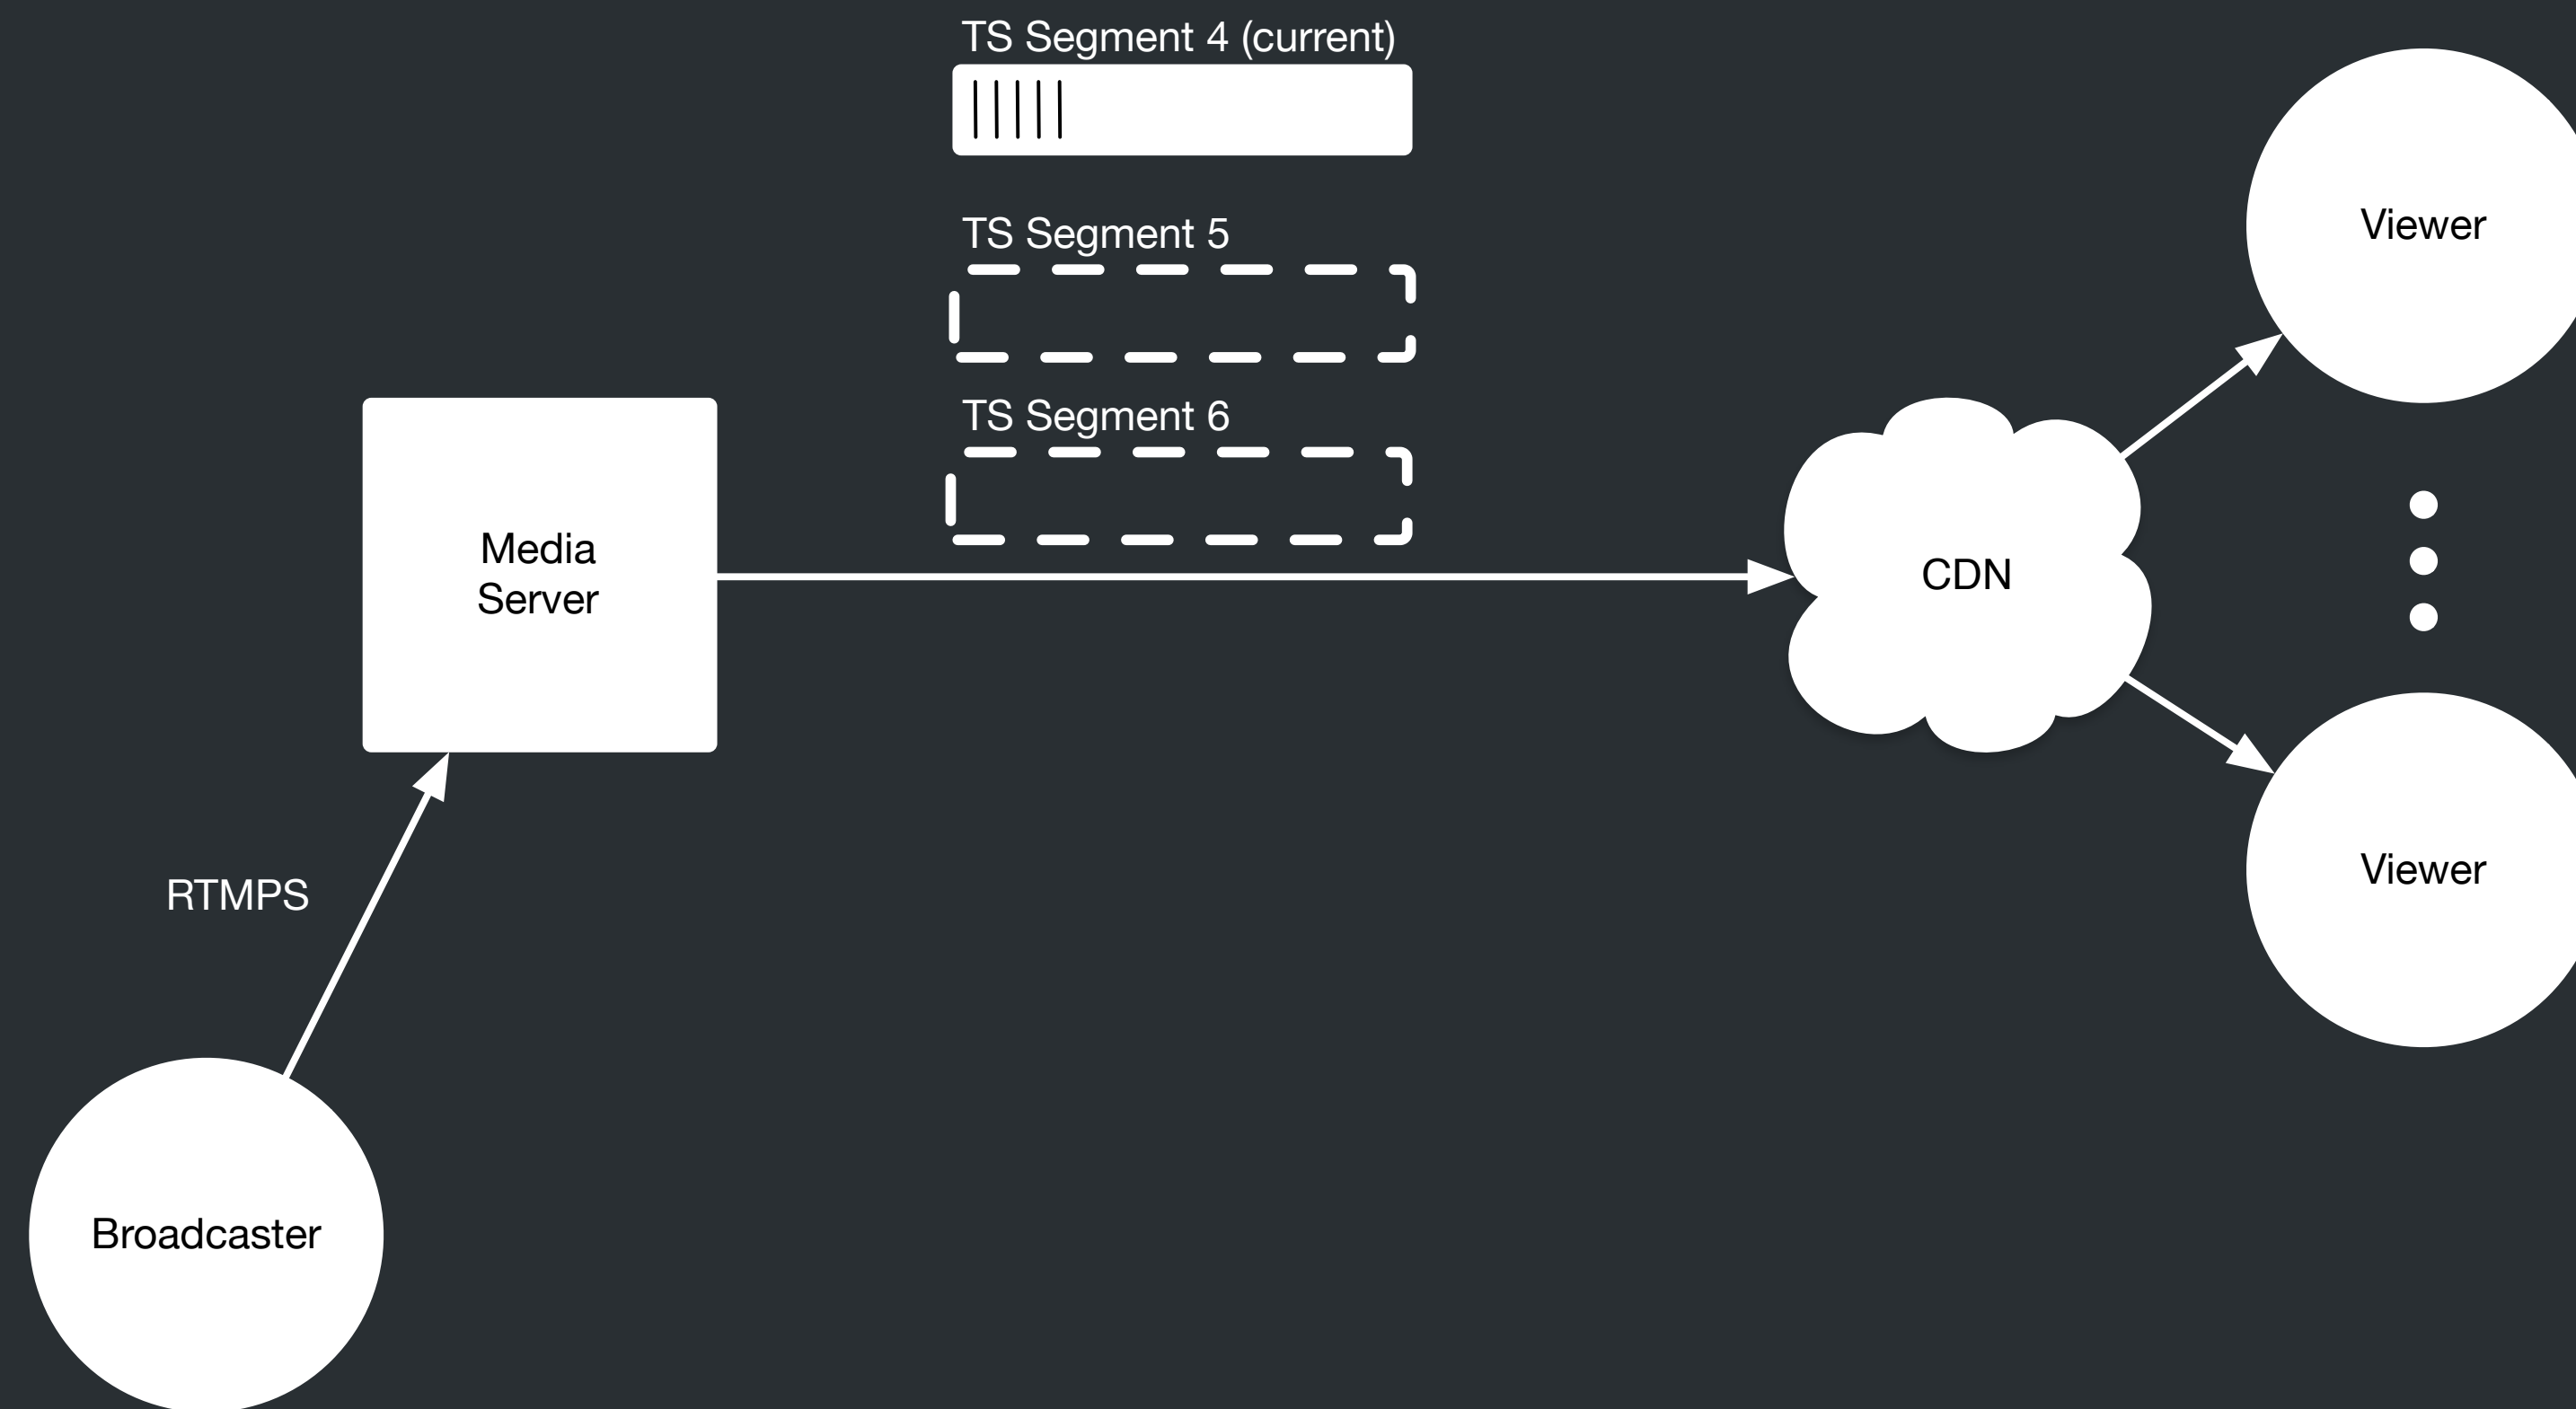
- Use HTTP/1.1 Chunk Transfer Coding
- Add segments to the manifest early




# LHLS And HLS Compared




# CDN Caching Behavior



Video Frames Received at the Playback Client



Playback Buffer Occupancy



# Bandwidth Estimation For Bit-Rate Adaption

- Chunked transfers make it hard to estimate
- Practical solutions
  - Non chunked downloads for probing
  - Heuristics based on observing playback buffer
- Recent work and small chunks
  - Arrival timing accuracy limitations?



# User Generated Content Considerations

- Segments advertised before properties are known
  - Duration, frame rate, and resolution are uncertain
  - Discontinuities unknown
- Why not transcode to remove irregularities?



# Summary

- RTMP meets latency requirements but doesn't scale
- HLS scales but doesn't meet latency requirements
- LHLS scales like HLS with latency of RTMP
- It's been working for us!

