



# VR Streaming: Latency Matters

*Dutch Guild 50*

10 May 2022



tiled  
media



# Rob Koenen

CBO & Co-Founder, Tiledmedia  
rob@tiledmedia.com



# Tiledmedia

- 5 year young, 20 strong, Rotterdam-based
- Worldwide distribution of video with quite insane resolution
- Hardcore tech development
- B2B
- Products: ClearVR and Mosaic Multiview

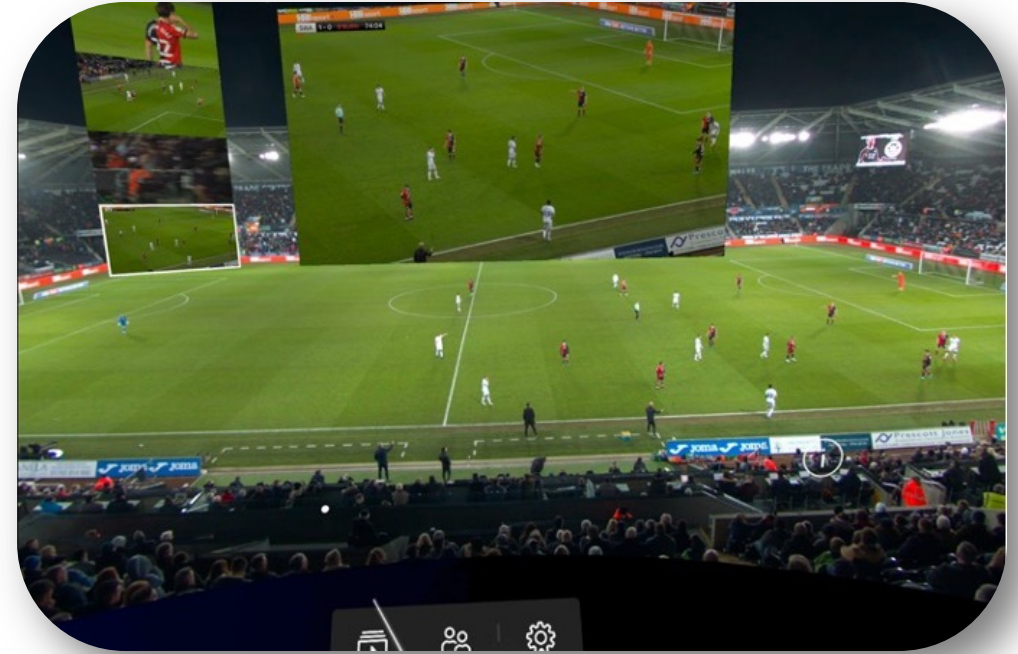


# Sky Worlds

Watching Football  
in VR, with Friends

# Sky Worlds

- 3 Premier League Matches / Week
  - One 4K fisheye cam
  - TV edit + 3 extra cameras
  - Watch with friends
- 
- Also: Films, cricket, netball, ..
  - <https://www.youtube.com/watch?v=hveGH9VfV2c>



# BT Sport

Seeing Premier League  
From All Angles on  
Mobile Devices



Winter 2022



**NBC**  
**OLYMPICS**  
By **xfinity**



# Winter 2022

- World-First 8K VR180
- Global distribution (China, US, Europe)
- Headset, web players





# How ClearVR works



ERP projection for illustration purposes only.  
Actual implementation uses cubemap



viewport






viewport




viewport





viewport







80% of tile switches completed in <math><50\text{ms}</math>  
98% of tile switches completed in <math><200\text{ms}</math>



Tiling works better with a cubemap projection

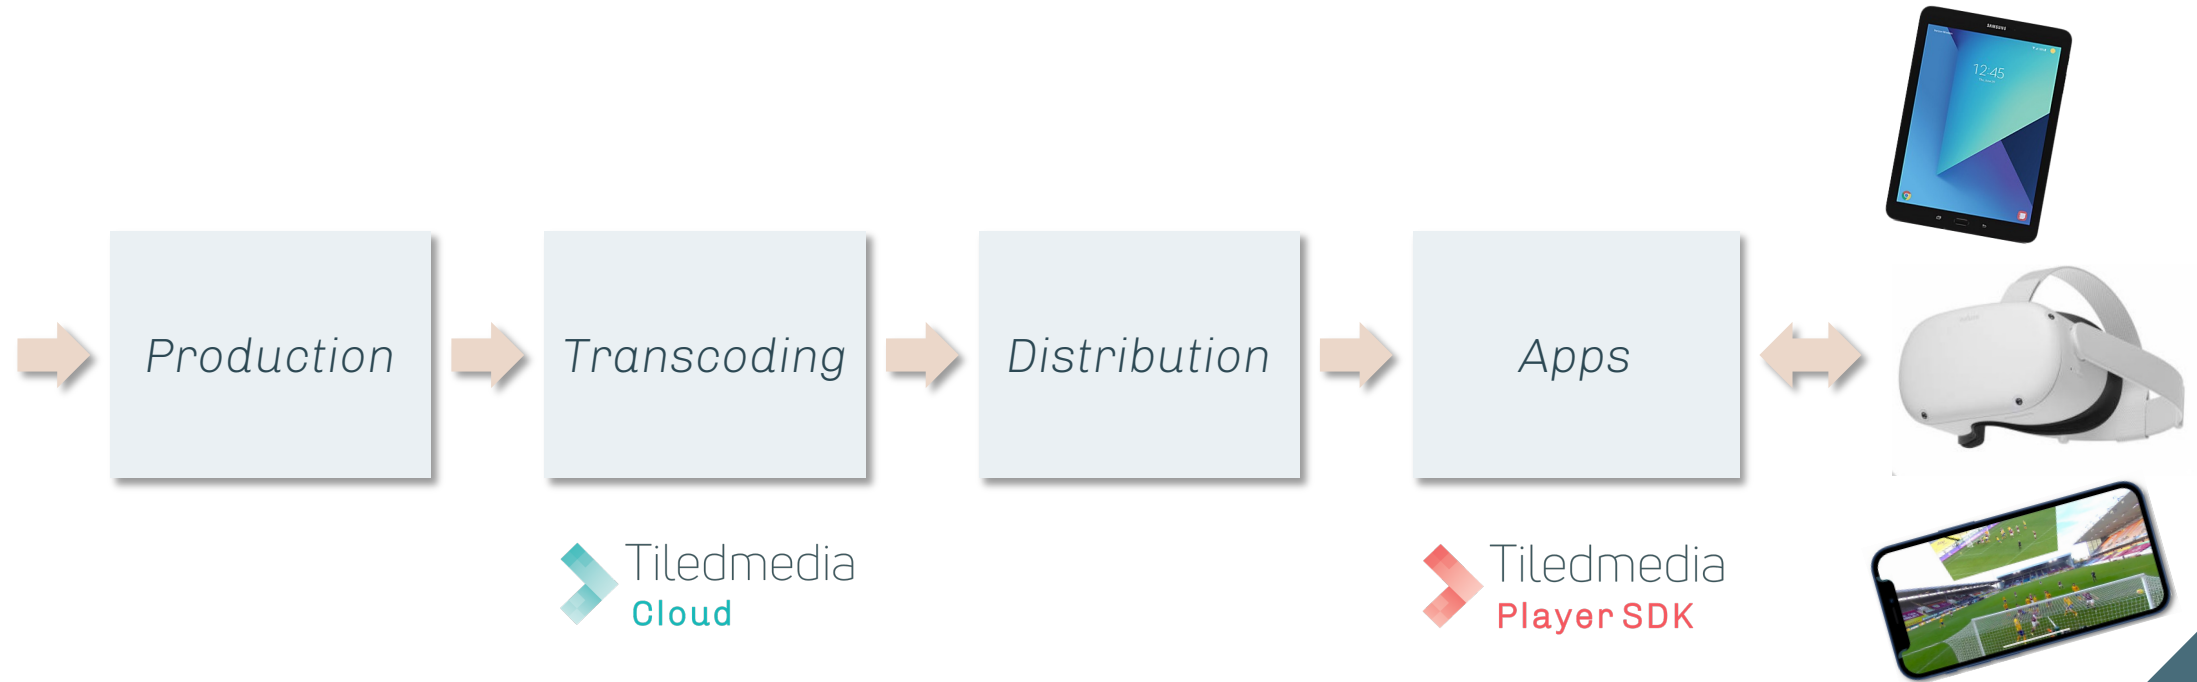


Raw decoder output (before rendering)

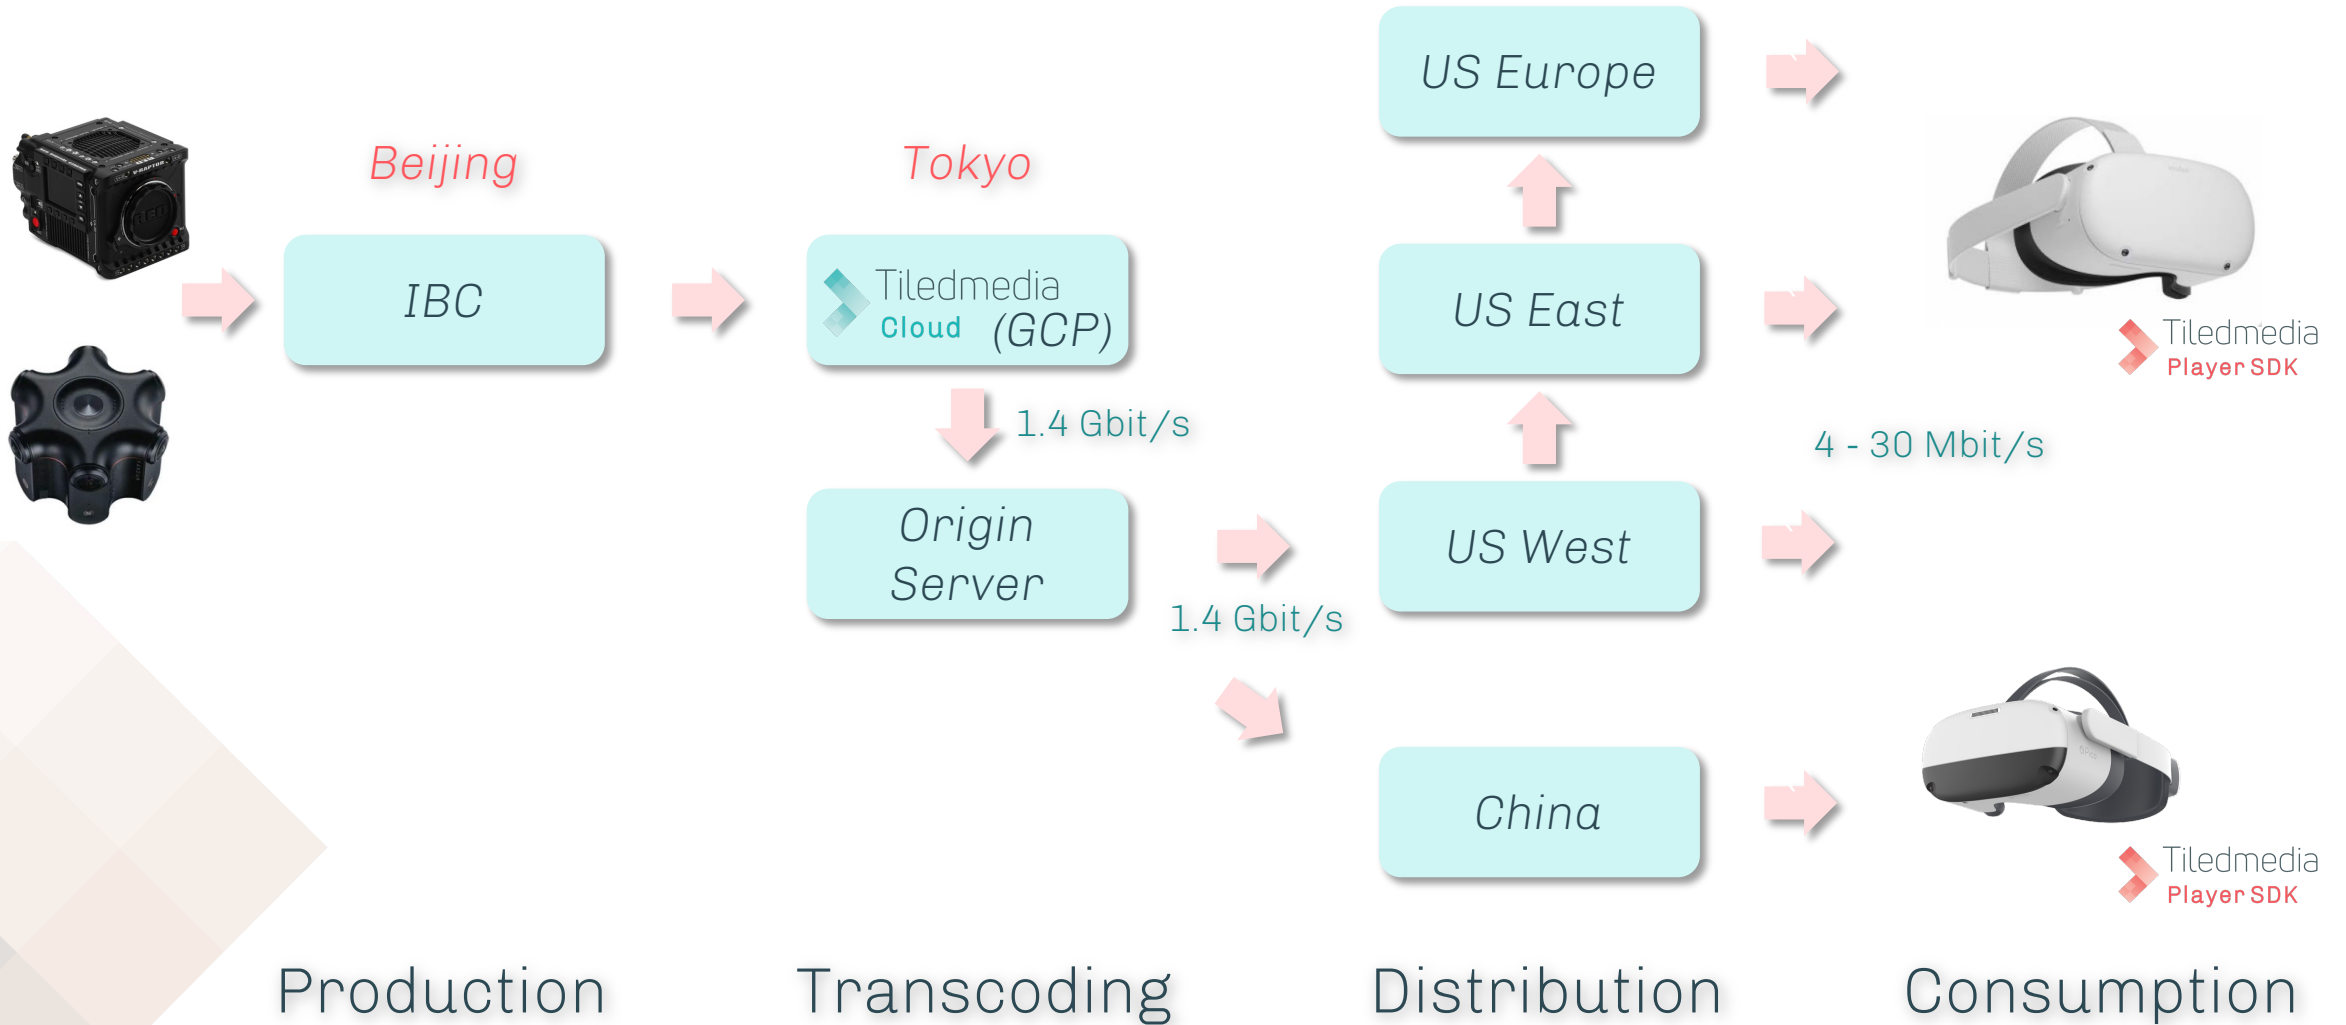


3 ABR Layers  
3 Random Access Intervals (GOPs) per Layer

# Integration into processing chain



# This Winter



# This Winter



- 1 x 8K 360
- 1 x 8K VR180 (5-cam edit)

Production

Beijing

IBC

~900 parallel tile encodes + 4 HLS levels per feed

Tokyo

Tiledmedia Cloud (GCP)

Redundant Global Origin Replication

1.4 Gbit/s

Origin Server

1.4 Gbit/s

US Europe

US East

US West

China

4 - 30 Mbit/s



Tiledmedia Player SDK



Tiledmedia Player SDK

Transcoding

Distribution

Consumption

# Latencies in VR

1. Glass-to-Glass
2. Motion-to-photon
3. Motion-to-high resolution



# Latencies in VR

1. Glass-to-Glass

2. Motion-to-photon

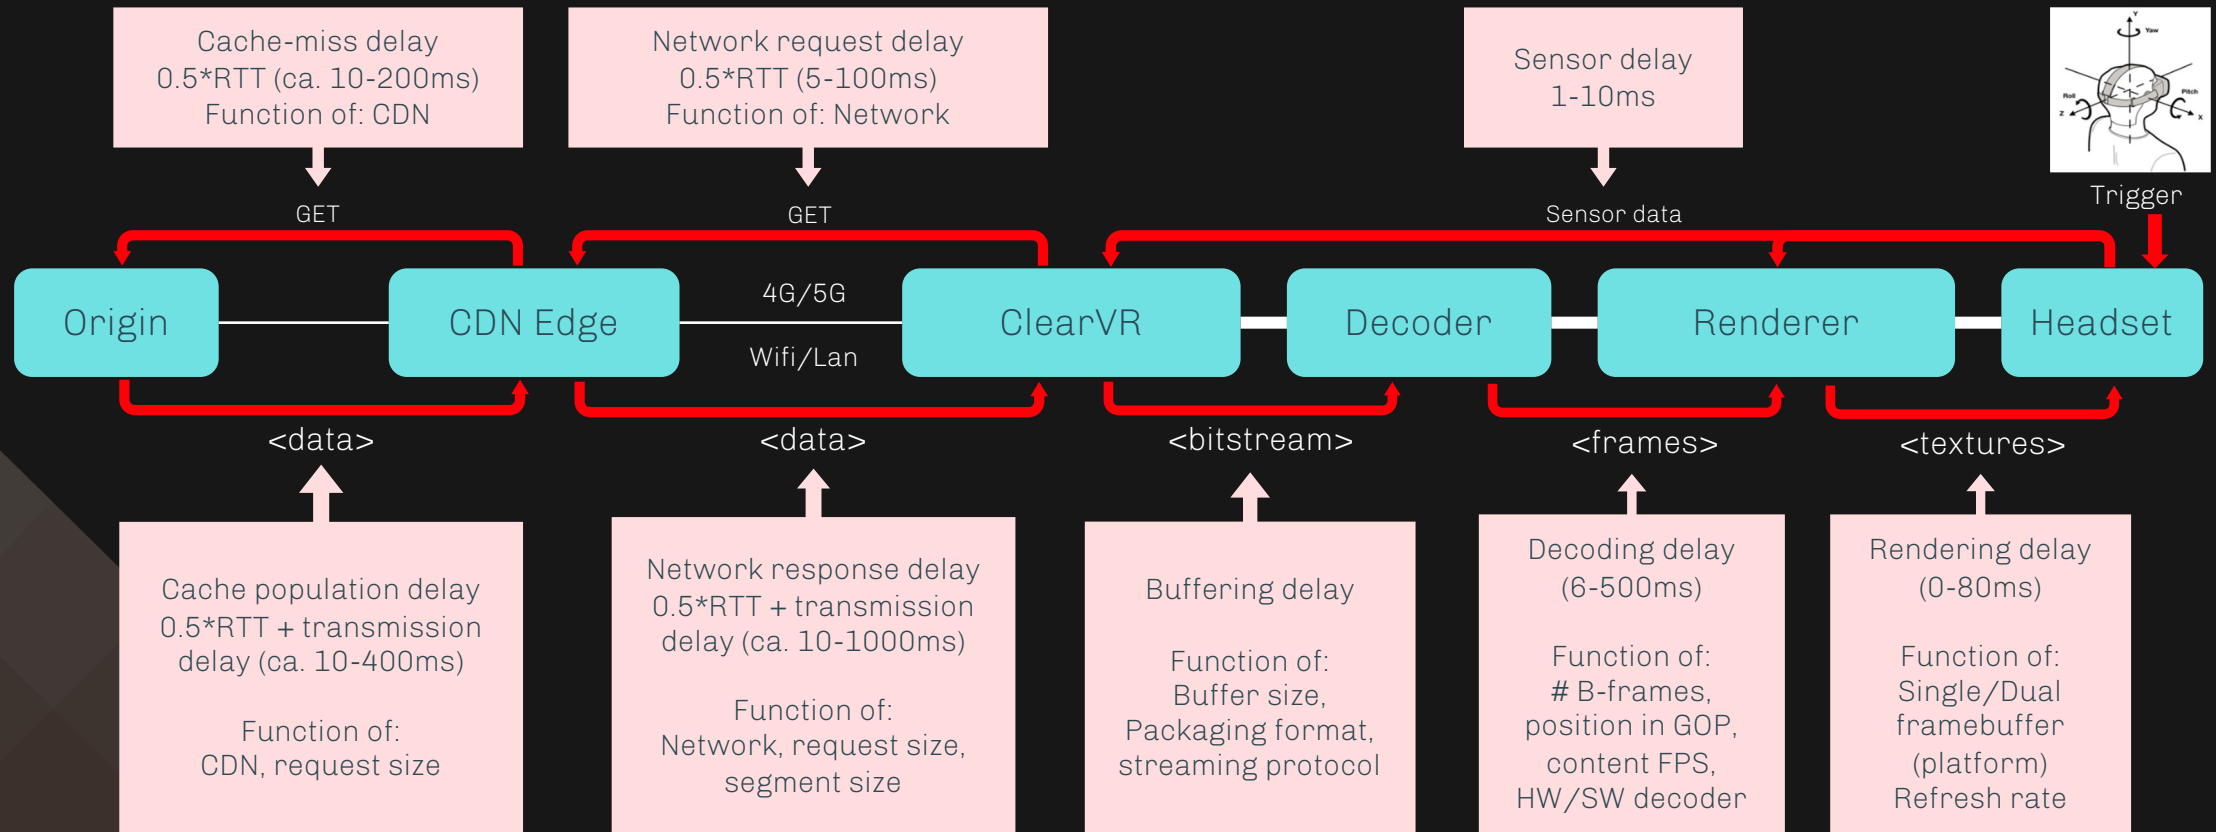
3. Motion-to-high resolution

1. As in HLS streaming  
(3-60 secs)

2. Depends only on local device;  
(virtually zero)

3. The focus of our efforts  
(invisible)

# Motion-to-High-Resolution Latency



# Conclusions

- Head and eye speed  $\sim 500^\circ/\text{s}$  ( $\sim 30$  msec per tile)
- Viewport covers  $\sim 4$  tiles horizontally
- You always watch in the middle of the VR display
- Vestibulo-ocular reflex: first head motion, then eye motion, then refocus
  
- All taken together, there is some 200 msec to work with
- Virtually always sufficient to retrieve high-res tiles (98% of tile switches)

En dan nu tijd voor ...