Vignette: Perceptual Compression for Video Storage and Processing Systems

Amrita Mazumdar, Brandon Haynes, Magda Balazinska, Luis Ceze, Alvin Cheung*, Mark Oskin

University of Washington
*UC Berkeley

Video storage engines use compression to trade visual redundancy for file size.

Baseline codec (HEVC) @ 20 Mbps
4 hours video playback

Source: Netflix Public Dataset

Video storage engines use compression to trade visual redundancy for file size.

fine details (noise, high frequencies)

color perception

fast motion

Baseline codec (HEVC) @ 20 Mbps
4 hours video playback

Source: Netflix Public Dataset

Vignette integrates new perceptual cues with video storage systems for reduced video sizes.

new perceptual cues

Baseline HEVC @ 20 Mbps
4 hours video playback

Vignette: 1 Mbps
6.5 hours video playback

Source: Netflix Public Dataset

Saliency is a powerful perceptual cue for compressed video workloads.

4K 360° video
300 MB

AI-generated saliency map
only 15% of pixels are important

Source: Lo et al., MMSys 2017

Saliency is a powerful perceptual cue for compressed video workloads.

Challenge:

Saliency prediction methods have improved in speed and quality, but video compression engines don’t leverage new perceptual cues.

Hypothesis:

We can integrate perceptual cues in cloud-scale video services for better storage and performance.

AI-generated saliency map
only 15% of pixels are important

Leveraging perceptual cues at scale presents design challenges.

Requires custom, outdated codecs

recent work:

- Applied to one outdated codec
- 1,500 lines of code per codec
- already worse than newer codecs without saliency compression

Leveraging perceptual cues at scale presents design challenges.

Requires custom, outdated codecs

No integration with storage manager

storage manager concerns :

- How can I switch between compression methods?
- Do I understand this format?
- How do I control quality?

Leveraging perceptual cues at scale presents design challenges.

Requires custom, outdated codecs

No integration with storage manager

No interface for applications

application concerns :

- Can users play this format?
- Can I use existing ASICs for playback?

Leveraging perceptual cues at scale presents design challenges.

Requires custom, outdated codecs

No integration with storage manager

No interface for applications

Goals:

✅ Modern codecs

✅ API for storage

✅ Application portable

Video processing pipeline through Vignette compression and storage system.

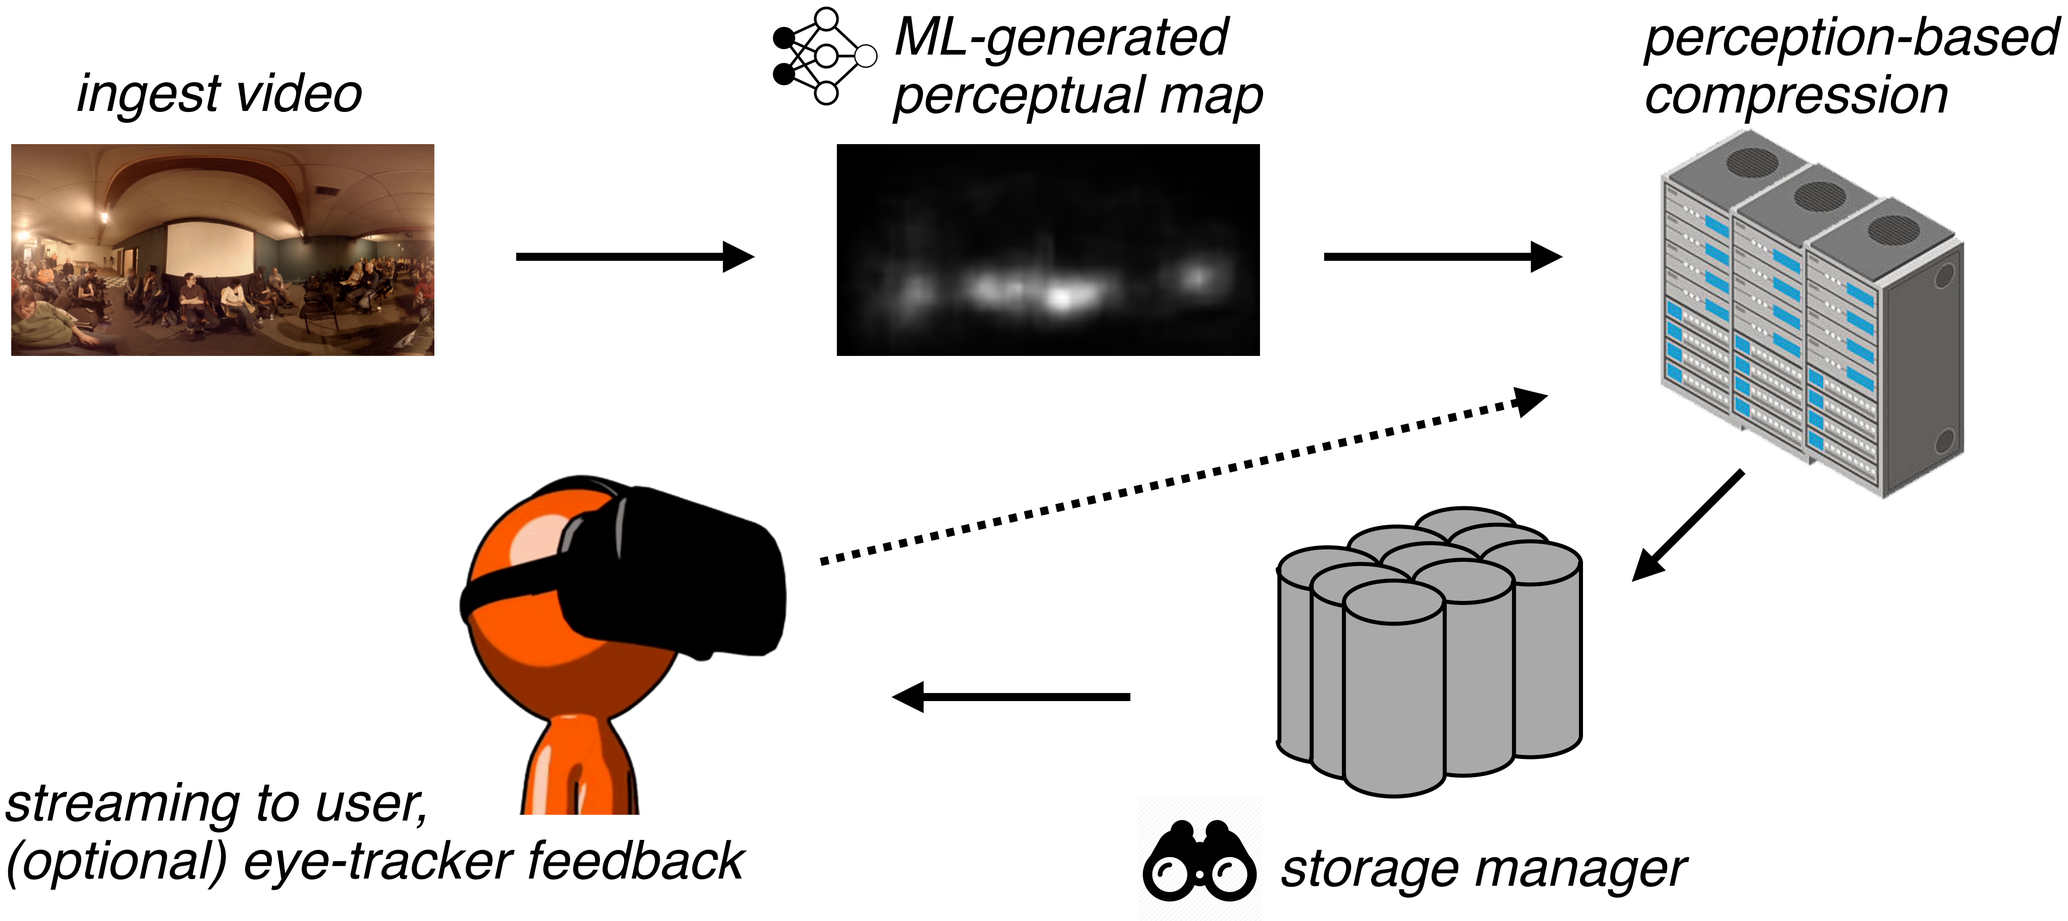

Video processing pipeline through Vignette compression and storage system.

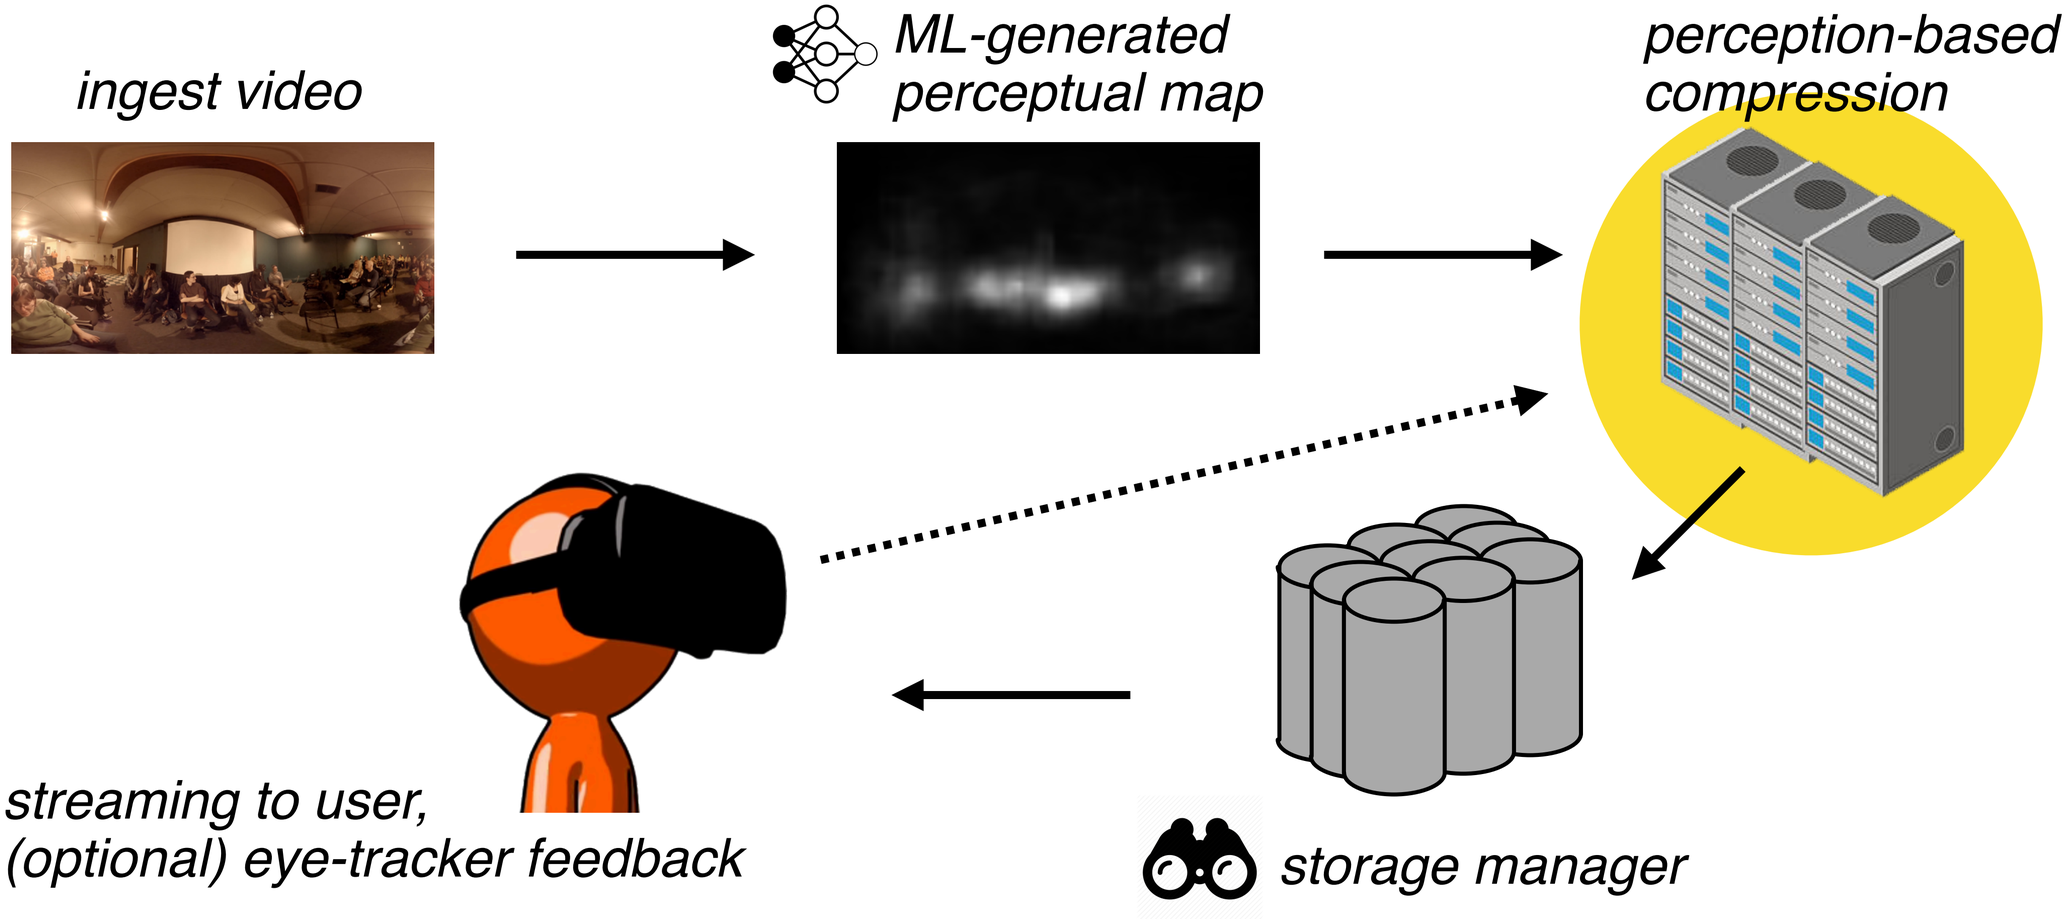

Vignette Compression uses tiles to convert saliency maps to video encoder parameters.

Automatically generate a saliency map

Split the video segment into tiles

Map saliency values to tiles

Source: Wong 2000

Vignette Compression uses tiles to convert saliency maps to video encoder parameters.

Automatically generate a saliency map

Split the video segment into tiles

Map saliency values to tiles

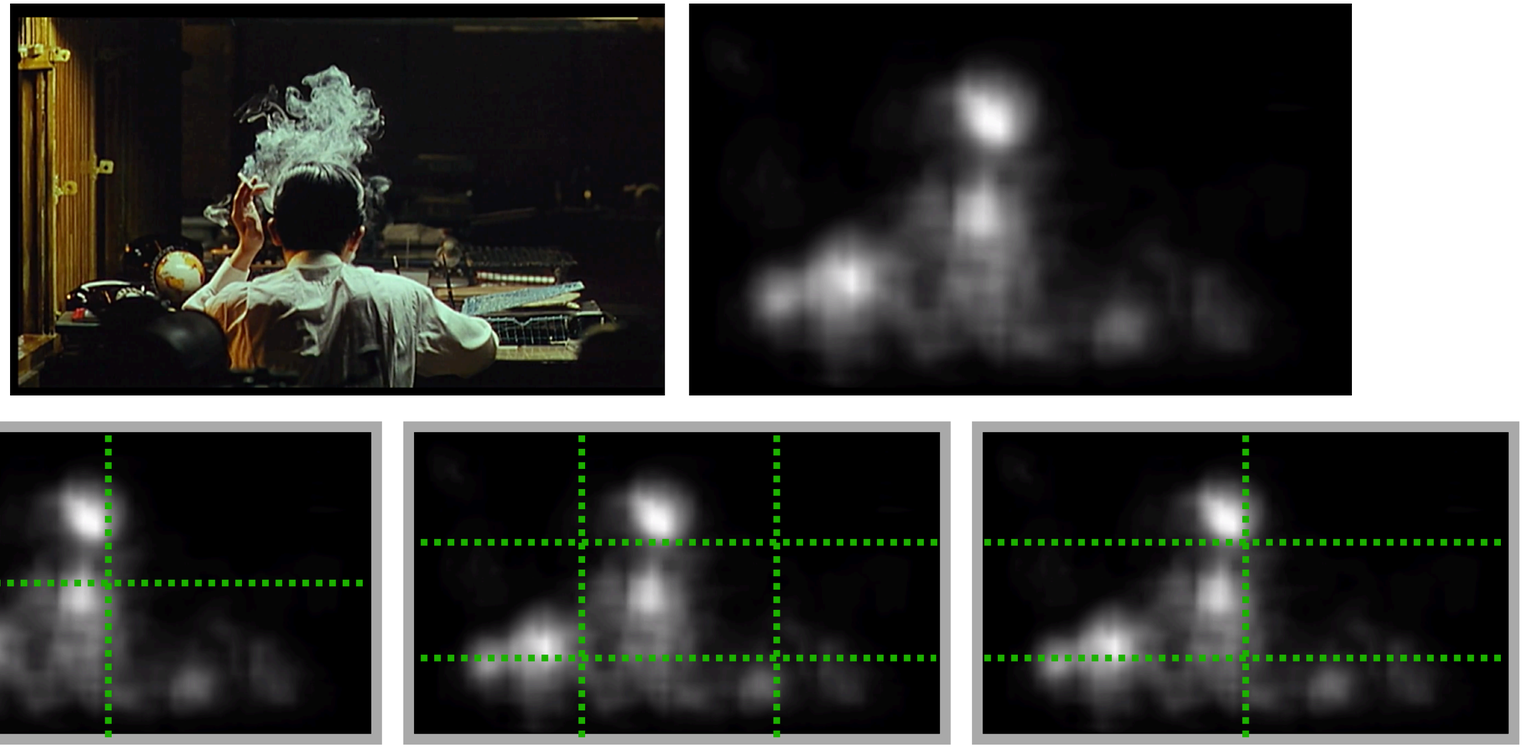

Source: Wong 2000

Vignette Compression uses tiles to convert saliency maps to video encoder parameters.

Automatically generate a saliency map

Split the video segment into tiles

Map saliency values to tiles

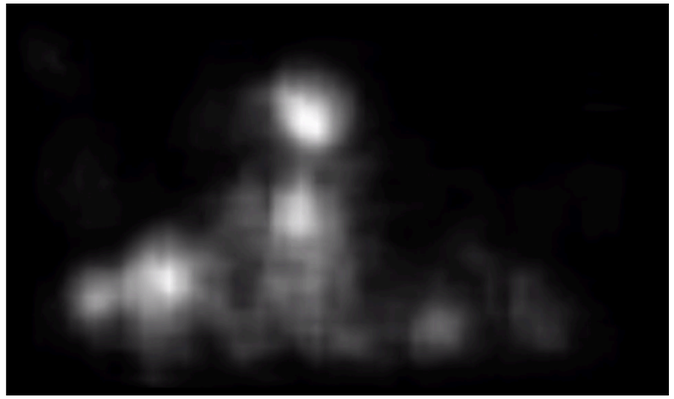

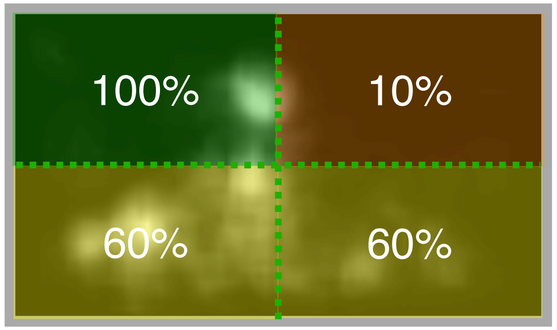

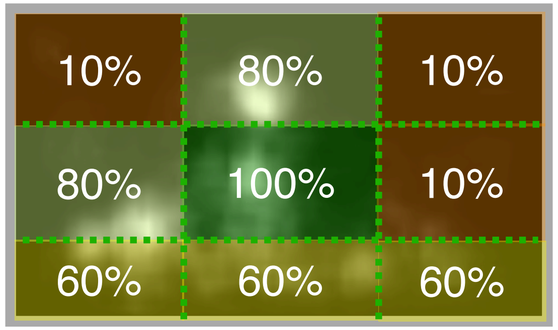

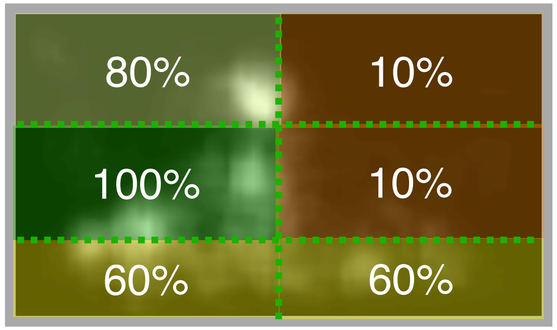

Source: Wong 2000

Video processing pipeline through Vignette compression and storage system.

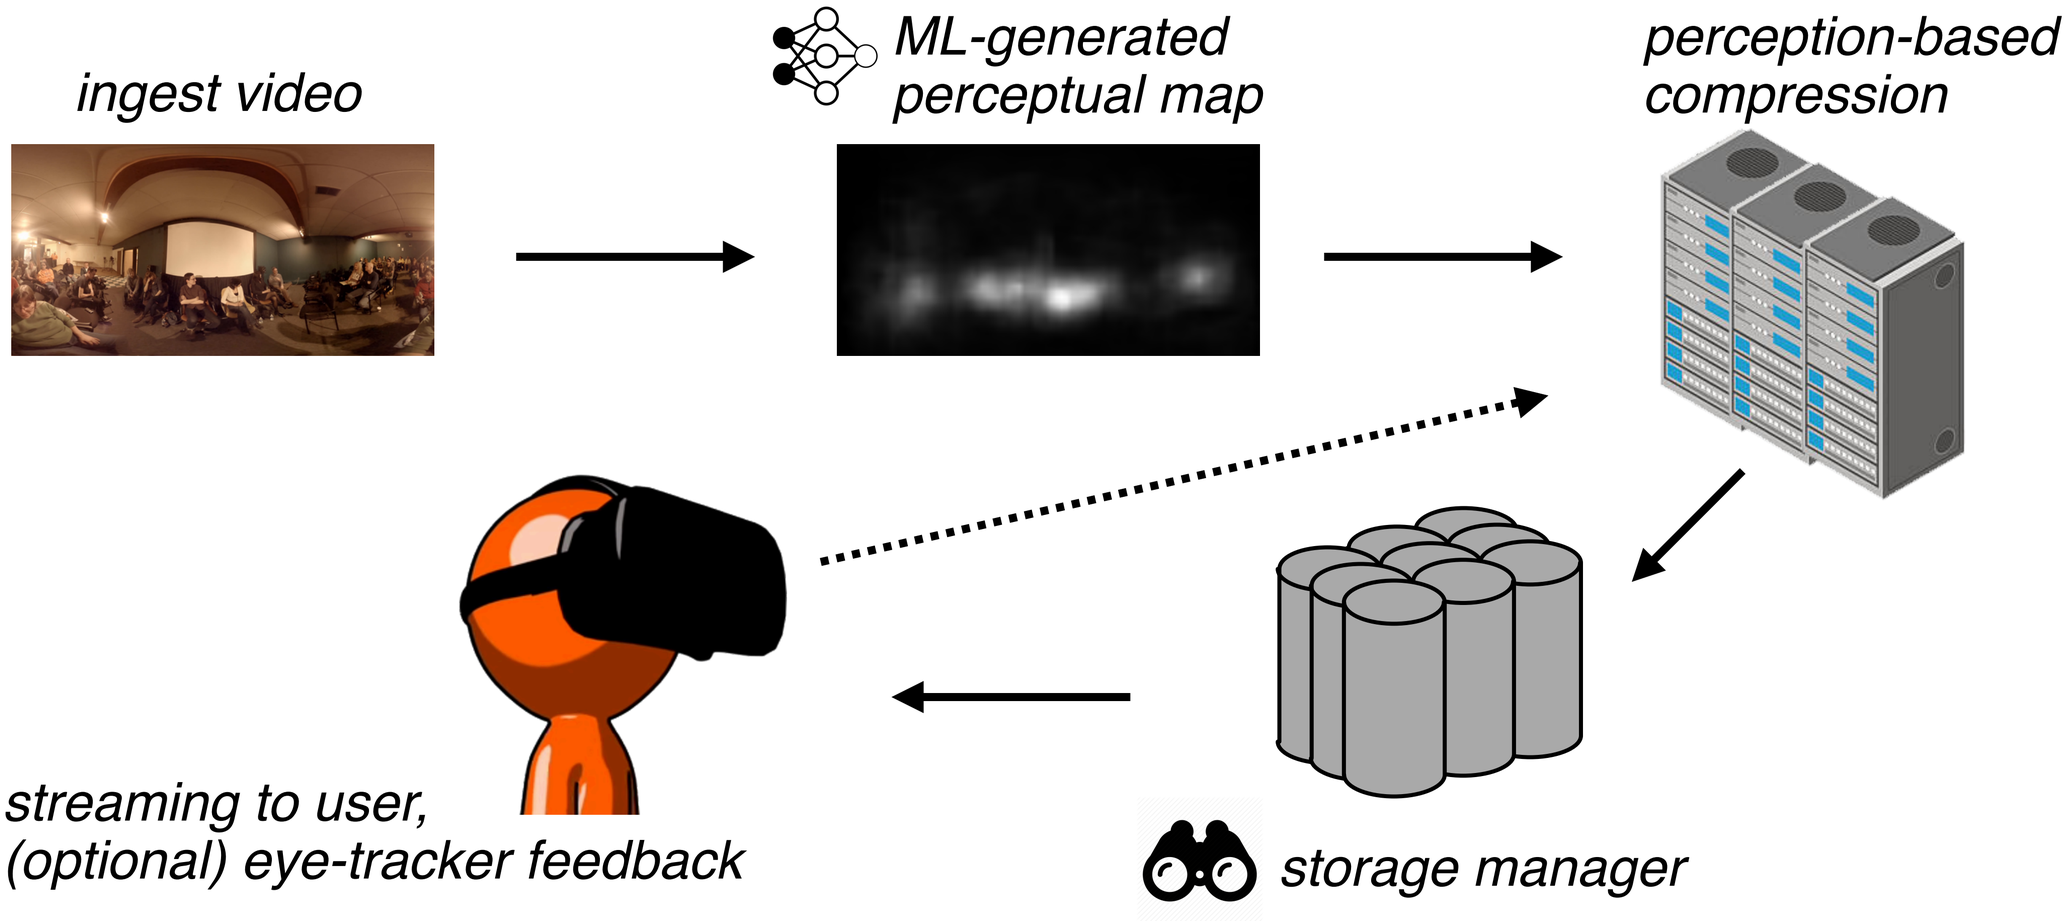

Vignette Results

Baseline HEVC @ 20 Mbps
4 hours video playback

Vignette @ 1 Mbps
6.5 hours video playback

Full Study Results: https://homes.cs.washington.edu/~amrita/vignette_socc19.html

Vignette Results

Participants either preferred Vignette or perceived no difference for 75% smaller videos.

Baseline HEVC @ 20 Mbps
4 hours video playback

Vignette @ 1 Mbps
6.5 hours video playback

Full Study Results: https://homes.cs.washington.edu/~amrita/vignette_socc19.html

Vignette videos reduce bitrate in non-salient regions, maintaining visual quality at lower storage

YouTube

NETFLIX

VR-360

blender

Vignette videos reduce bitrate in non-salient regions, maintaining visual quality at lower storage

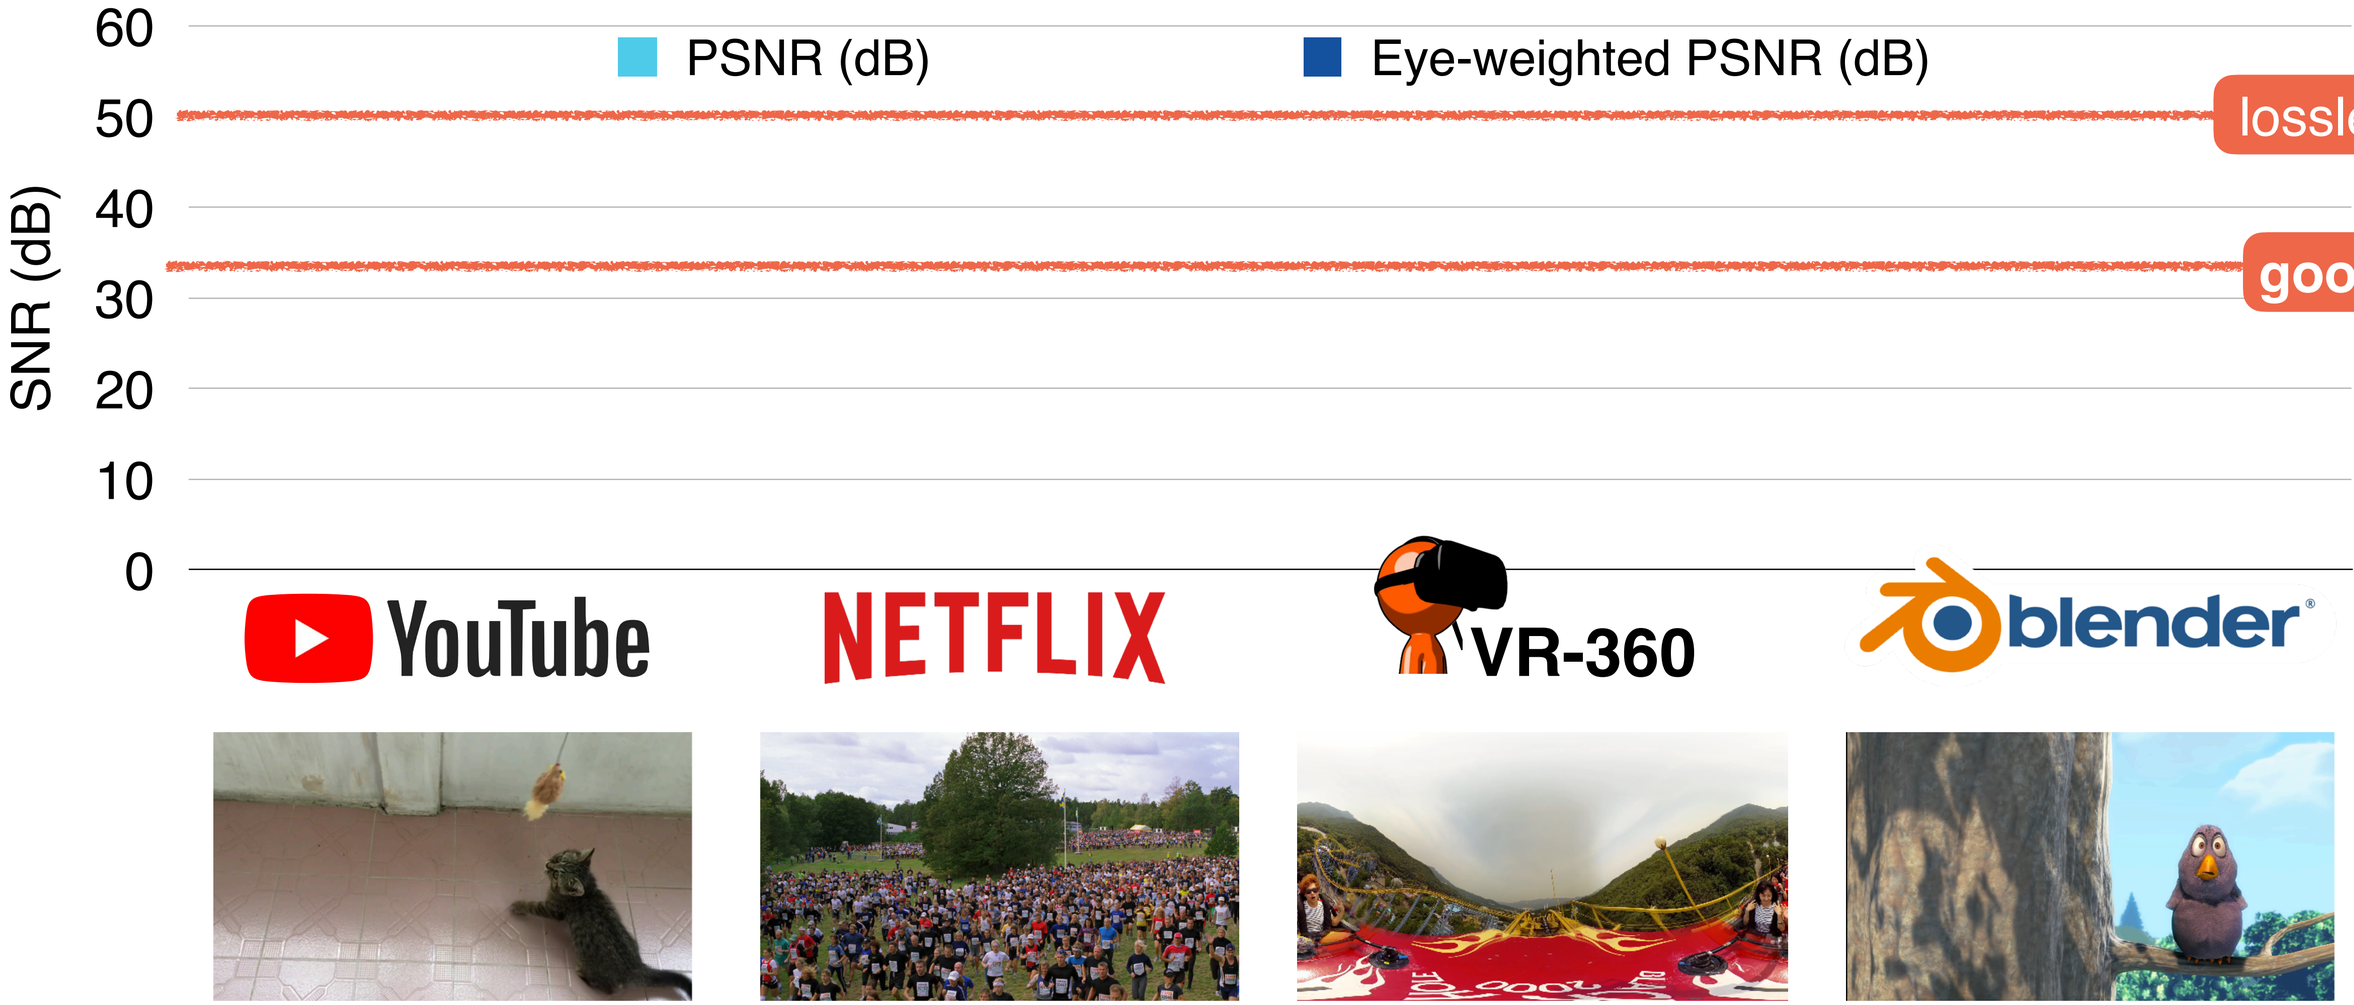

Vignette videos reduce bitrate in non-salient regions, maintaining visual quality at lower storage

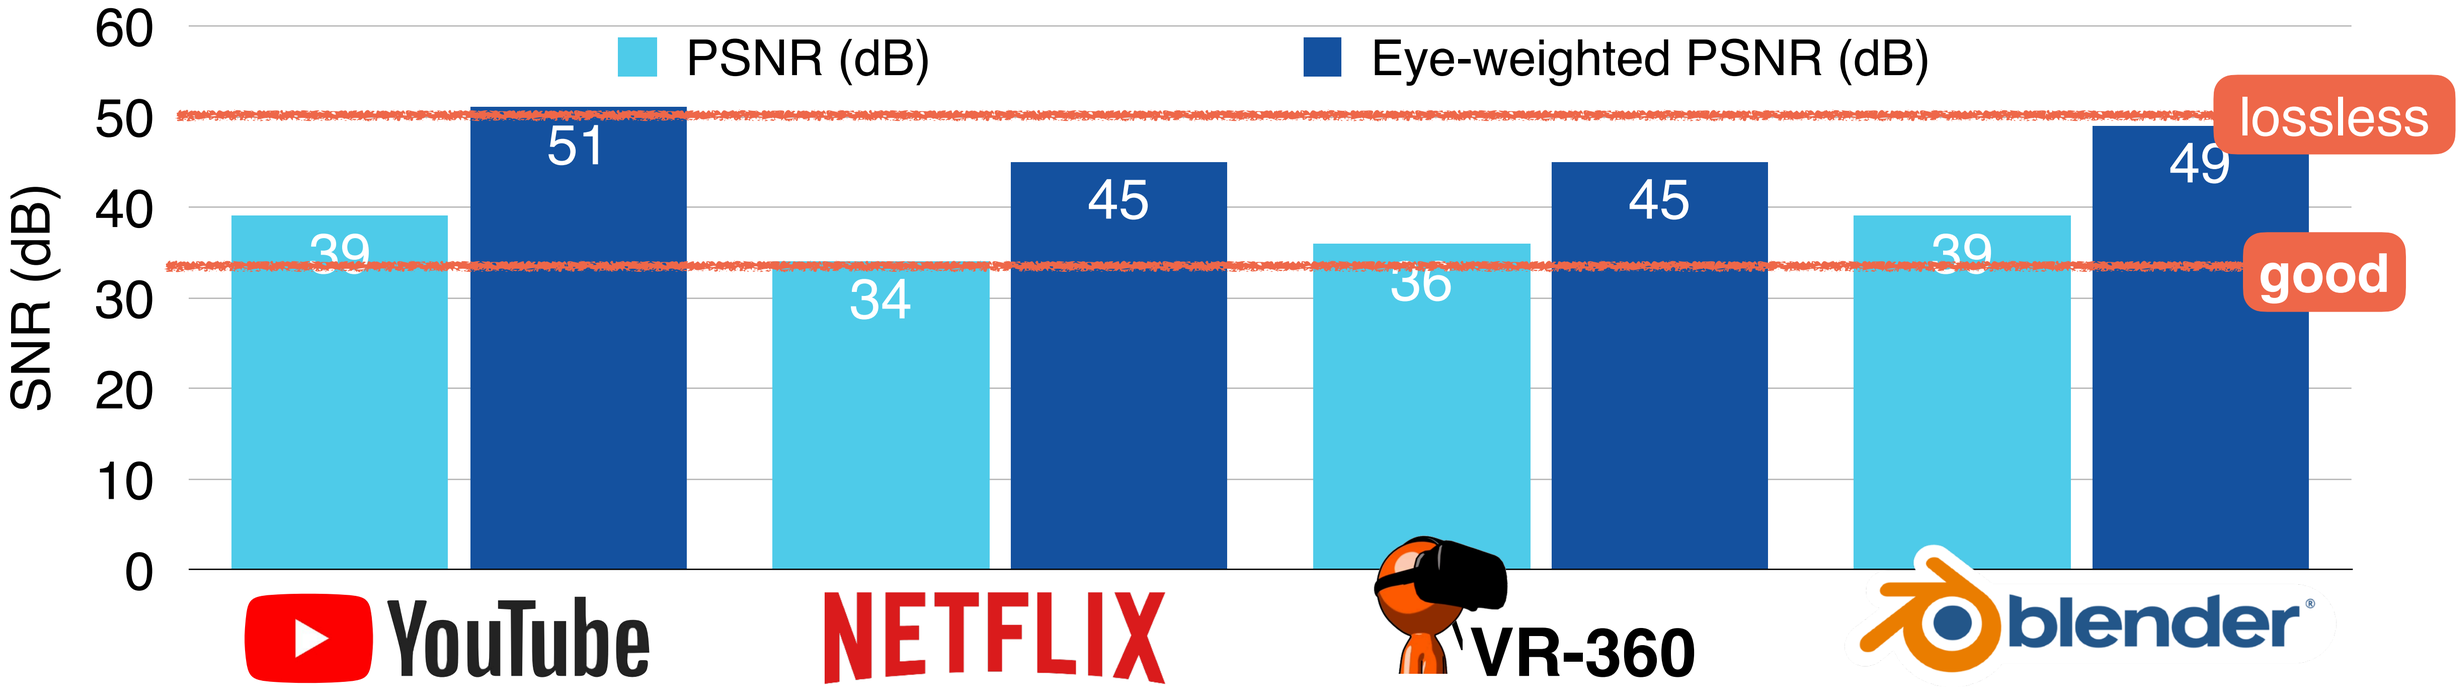

TCO analysis of Vignette in an AWS data center, storing and streaming a 1-million video library.

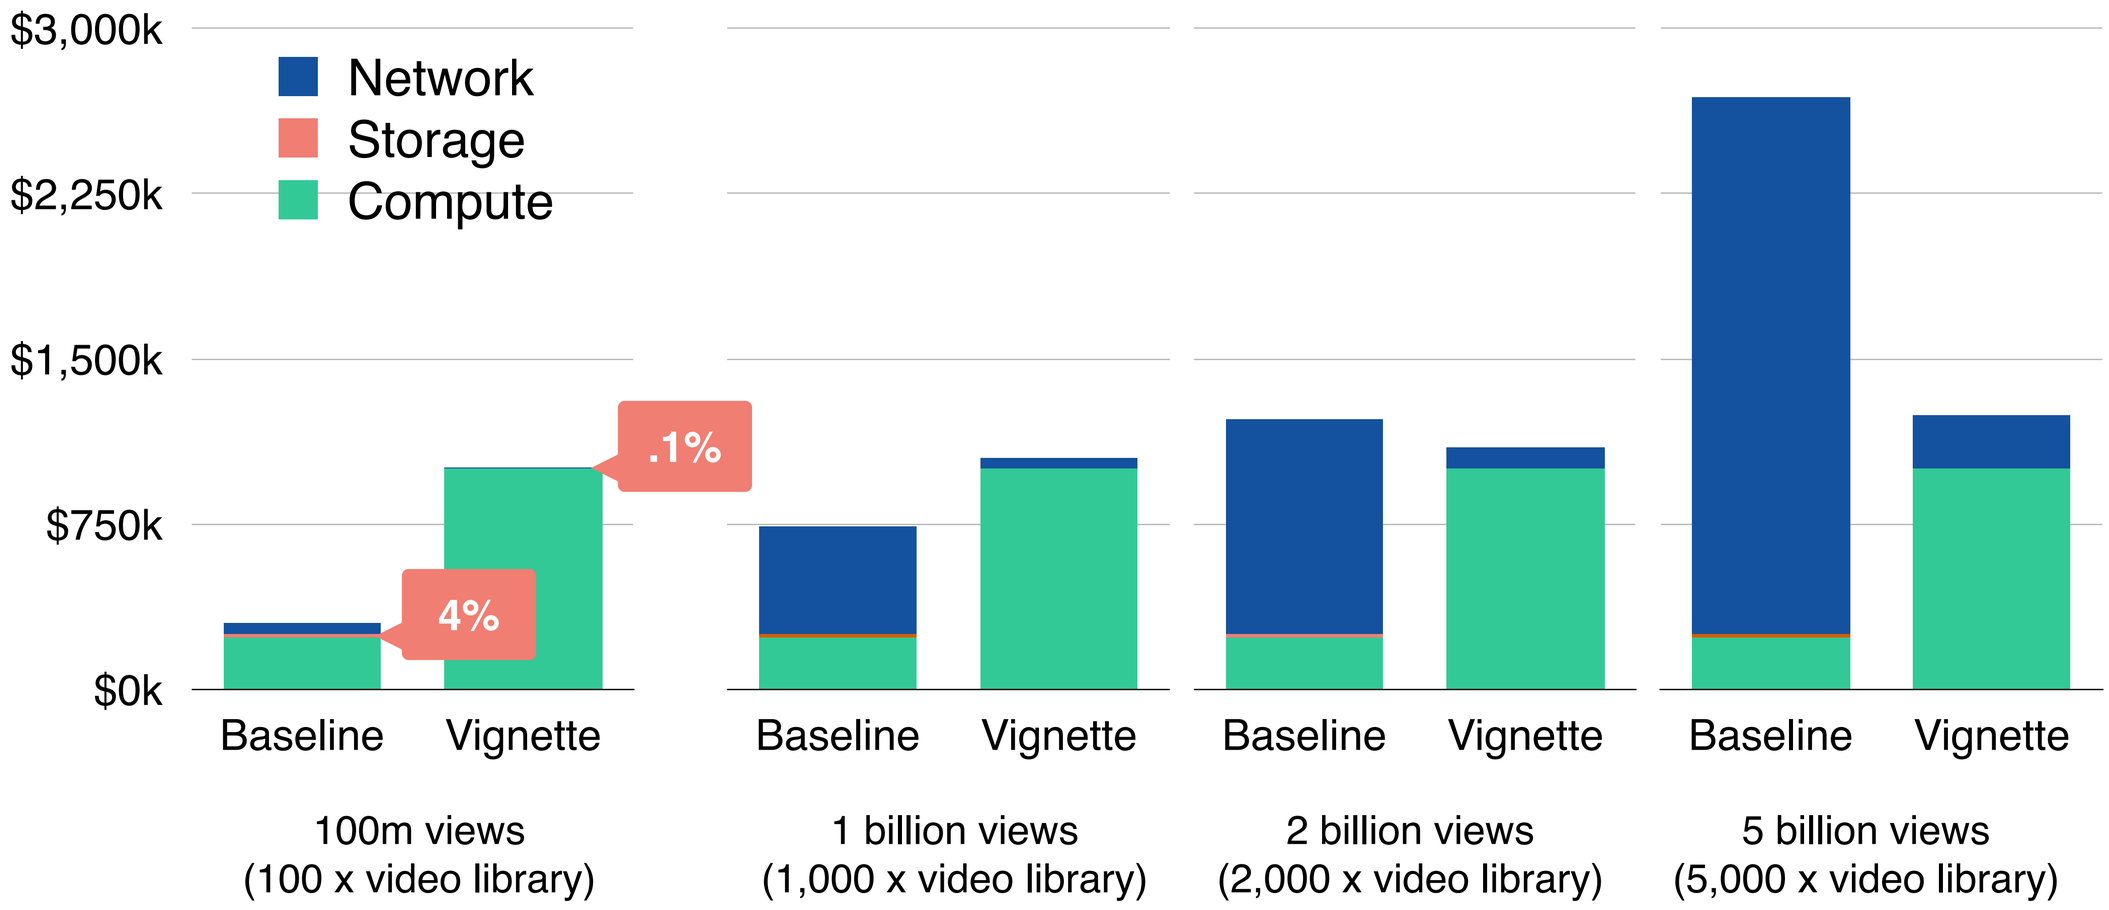

TCO analysis of Vignette in an AWS data center, storing and streaming a 1-million video library.

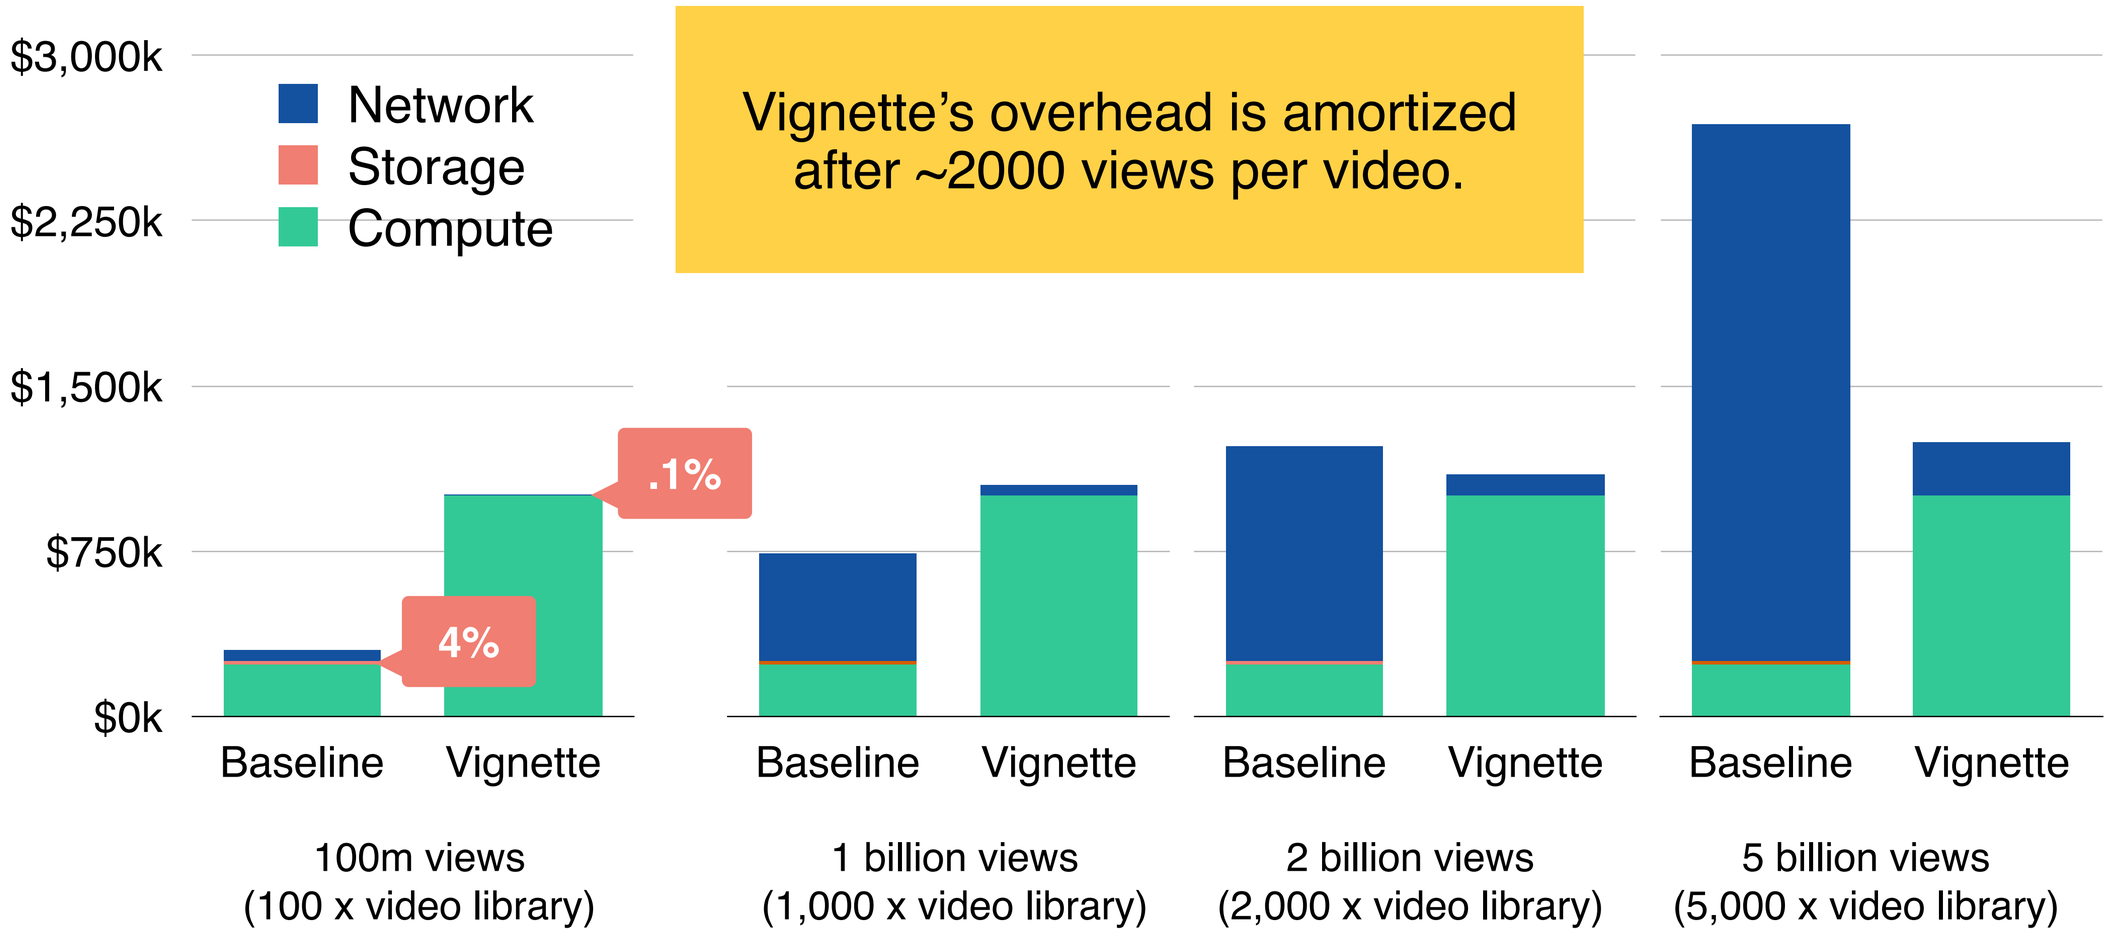

Vignette is a system for perceptual compression and storage.

Vignette Compression
codec-agnostic perceptual video compression

Vignette Storage
storage manager for perceptually-compressed videos

Reduces storage by up to 75% with little quality loss