Learn to Use Capped CRF with SVT-AV1 for Live Streaming

Jan Ozer
Streaming Learning Center
jan.ozer@streaminglearningcenter.com
September, 2023

Agenda

- Overview
- The command strings
- Performance
- Categories and CRF values
- CRF savings - Average
- CRF savings - Low Frame
  - *Analysis by media type*
- Conclusions

Online Course for New Streaming Professionals

Streaming Media 101:
Technical Onboarding for Streaming Media Professionals

Learn the technologies, techniques, and skills to succeed in a streaming media-related role, whether it's producing and distributing streaming video or creating the tools and services necessary to do so.

https://bit.ly/StreamingMedia101

What You Will Learn

In about 11 hours, this online course will teach you the terms, technologies, best practices, and skills needed to excel in a technical role in the streaming media industry. You will learn:

- How to encode and deploy streaming video using the H.264, HEVC, VP9, and AV1 codecs
- How to encode for single file and adaptive bitrate encoding and packaging for HLS, DASH, and CMAF
- About digital rights management (DRM) and distribution issues like choosing a CDN and how to measure and ensure Quality of Service and Quality of Experience
- Critical production-level decisions, like whether to encode on-premise or in the cloud, how to choose a per-title encoding technology and cloud encoder, and how to compute the breakeven on deploying an advanced codec like HEVC or AV1

You will learn to:

- Analyze files with MediaInfo, Bitrate Viewer, Apple's AVQT, and the Moscow State University Video Quality Measurement Tool
- Encode in FFmpeg and Handbrake
- Produce mezzanine files for upload to a streaming service
- Connect to YouTube Live and Facebook Live
- Embed a live or on-demand video into a web page

Overview

	CRF Value	Bitrate Savings	Average Quality Drop	Low-Frame Decrease	Recommendation
Media	33	-35%	-1.4%	-5.9%	Deploy but verify
60 fps sports	36	-18%	-1.6%	-29.2%	Use with caution
30 fps sports	41	-39%	-2.8%	-20.6%	Use with caution
Animation	42	-64%	-2.9%	-10.0%	Deploy but verify
Office	43	-80%	-2.5%	-5.7%	Deploy but verify
Average		**-44%**	**-2.1%**	**-13.5%**	

Capped CRF is an alternative encoding technique to VBR tested here for **live transcoding** with the SVT-AV1 codec using FFmpeg. We tested using preset 8 for reasons discussed on the next page.

During our comparison testing, we concluded the following:

- **Performance:** Capped CRF **delivered greater throughput** than VBR, though this varied by content type and logical processors deployed.
- **Bitrate savings:** Bitrate savings **averaged about 44%** over our five categories. For reference, the cap and VBR target for 60fps sports was 6 Mbps; otherwise, it was 4.5 Mbps.
- **Average quality:** The average VMAF score ***dropped*** 2.1% but still averaged 94.41, which is in the relevant target zone for-top rung quality for premium content producers.
- **Low-frame quality:** Low-frame quality is a predictor of ***transient*** quality problems. Capped CRF decreased the low-frame quality by 13.5%, though most of this was in sports-related footage.

Overall, all live streaming producers should consider capped CRF as an alternative to VBR.

Why Not Constant Bitrate Encoding (CBR)

The first version of this report compared capped CRF with CBR encoding, finding that capped CRF:

- Delivered ***much greater throughput*** than CBR (even more so than VBR)
- Delivered ***significant bandwidth savings*** over CBR (as with VBR)
- Delivered ***better*** low-frame scores than CBR (unlike VBR, which delivered better low-frame scores than capped CRF).

When I sent my report to a contact at AOM for review, he said “CBR is only implemented for the low-delay use case (video conferencing, very low latency live use cases), not for other use cases such as broadcasting and live streaming where latency isn't an issue.”

He attributed most of my negative CBR-related findings to this design intent. Since these tests were intended for general-purpose live event productions, I switched to VBR for this analysis.

If you’d like to see a copy of the report analyzing CBR, contact me at janozer@gmail.com.

Preset 8 Best for Live Transcoding

Key finding: Preset 8 was optimal for live transcoding.

Discussion: This chart compares all SVT-AV1 presets for version 1.7 by encoding time (blue), average VMAF quality (green) and low-frame (red). Low-frame is the lowest frame in the file, which is a predictor of transient quality issues.

In this test suite, we used preset 8, which delivered much faster than real-time encoding, offered good average quality, and nearly a 4-point boost in low-frame quality over preset 9. Preset 7 would decrease encoding speed/throughput by about 50%.

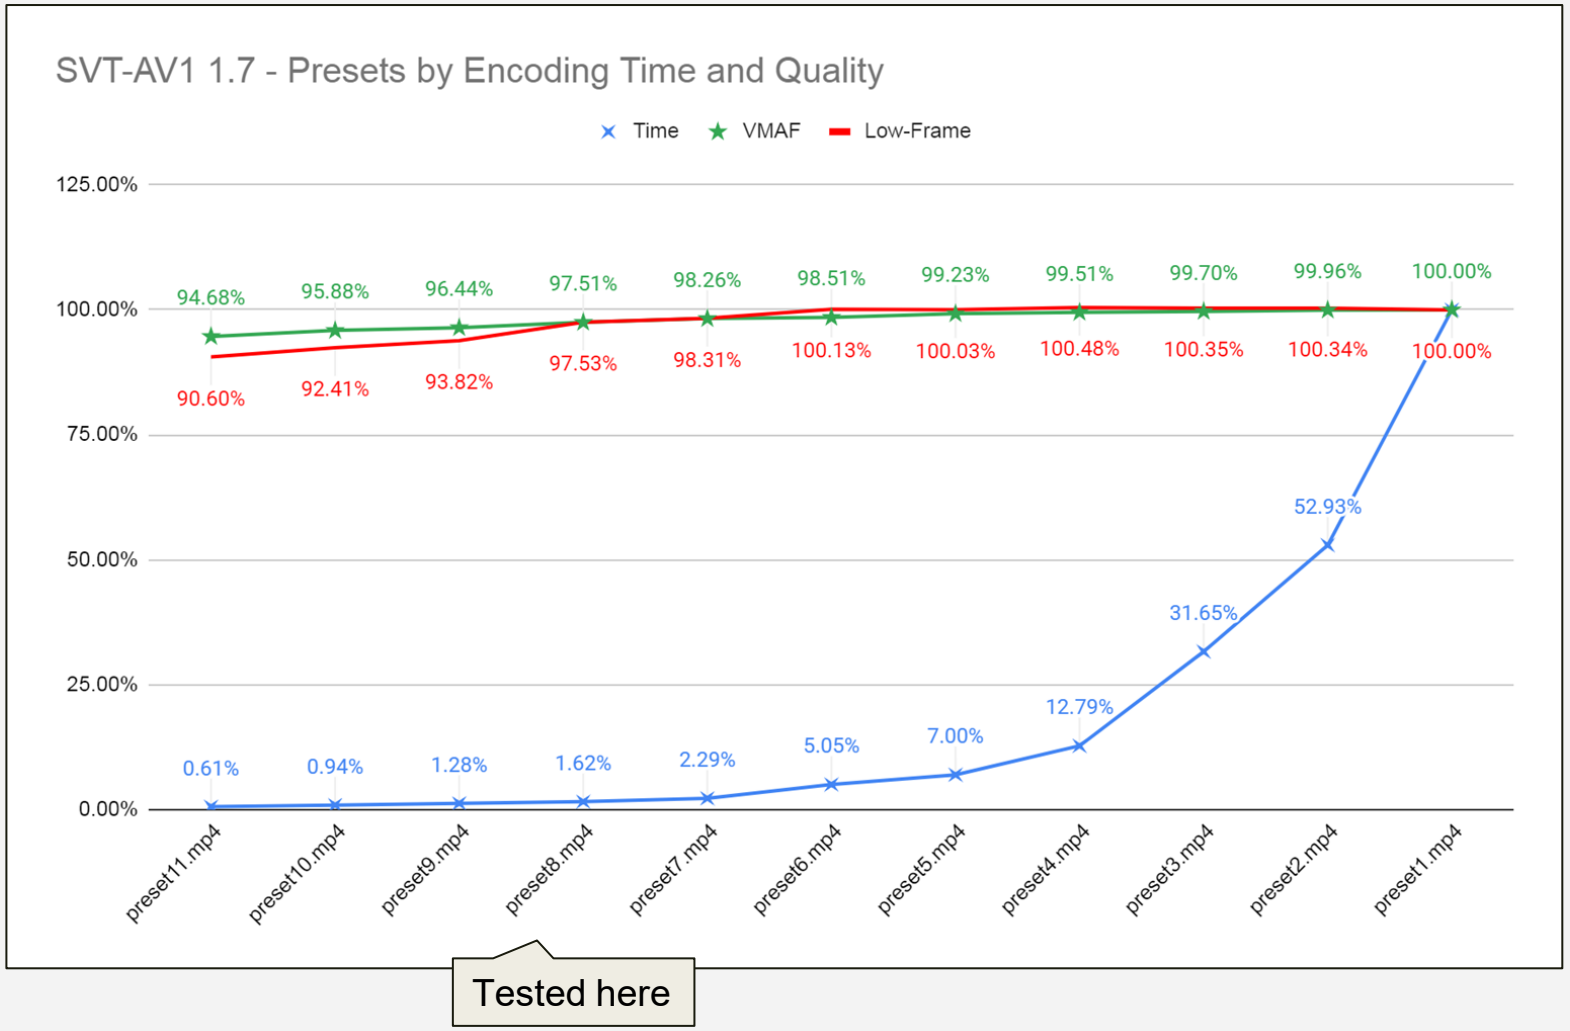

Live – Command String

- In a live application, capped CRF is an alternative to 1-pass VBR. Here are the two command strings, which are the simplest available to implement each encoding mode.

Capped CRF

```
ffmpeg -i input.mp4  -c:v libsvtav1  -g 60 -preset 8 -crf 42 -svtav1-params mbr=4500  output.mp4
```

1-Pass VBR

```
ffmpeg -i input.mp4 -c:v libsvtav1  -g 60 -preset 8  -svtav1-params rc=1:tbr=4500:enable-force-key-frames=0
output.mp4
```

With previous versions of SVT-AV1, you didn't need the force-key-frames expression, but without it, the command string crashed on 1.7. I asked a contact about this and he responded:

> It seems to be a bug in ffmpeg that was added recently. You can for now get around it by doing this:
> -svtav1-params rc=1:tbr=4500k:enable-force-key-frames=0
>
> Which turns off the enable-force-key-frames=0, which was added recently to support avif encoding, someone decided to turn it on by default.

Bitrate Profile - VBR vs. Capped CRF

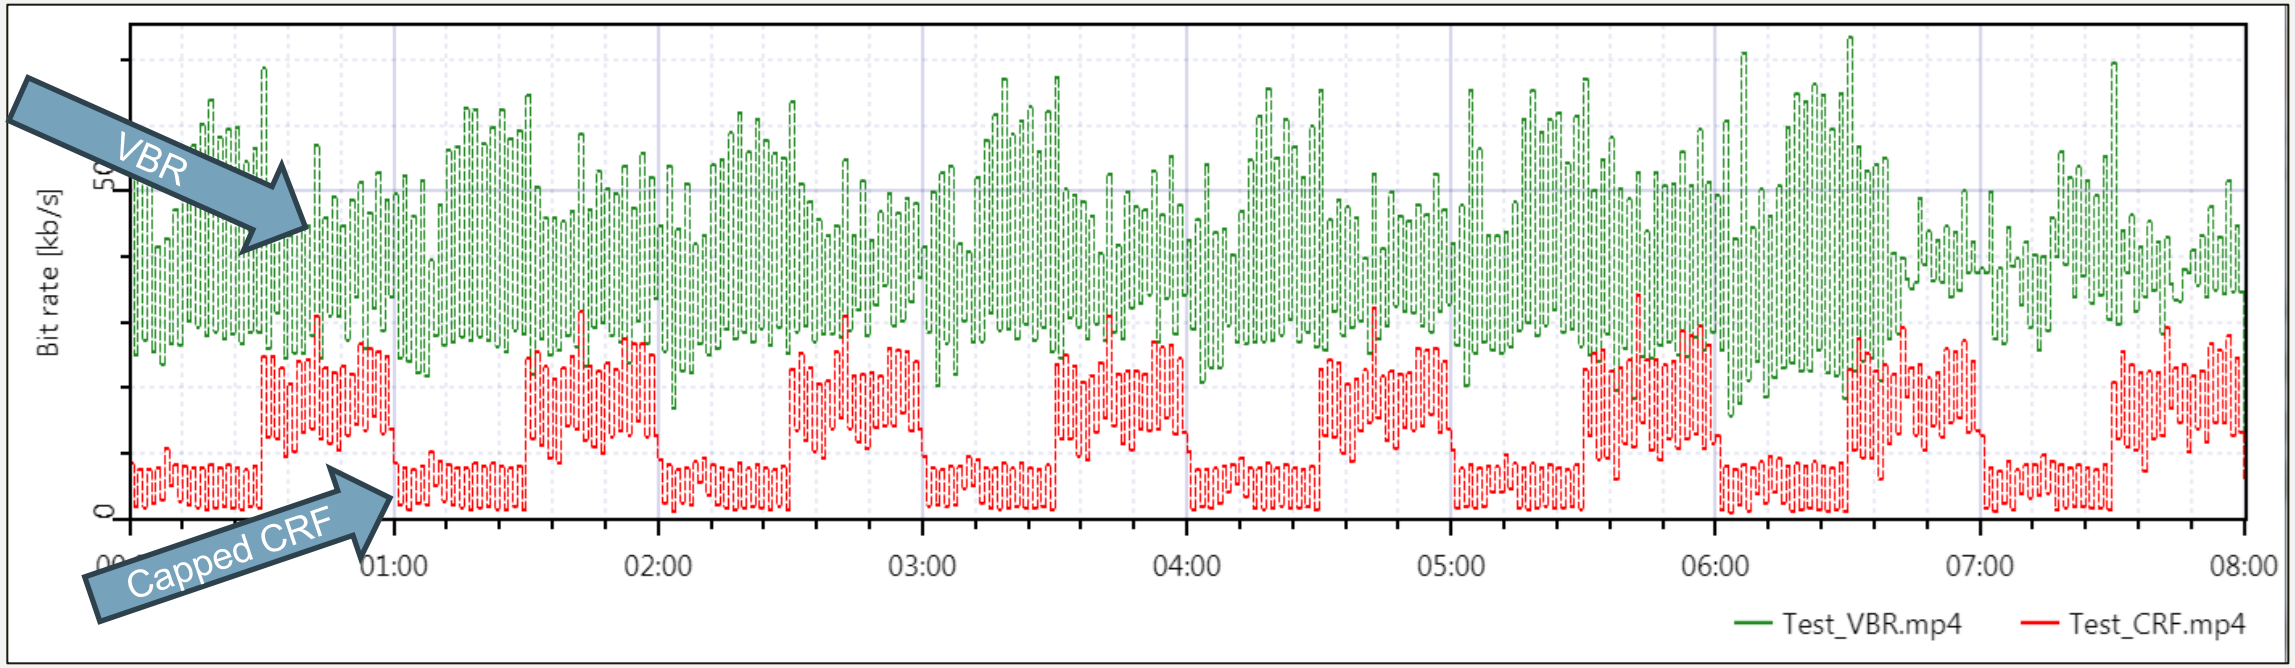

Key finding: VBR isn't as VBR-like as you would expect, which limits its effectiveness.

Discussion: I was curious about the bitrate profile produced by the two techniques, so I encoded a file that contains eight sequences of 30 seconds of talking head followed by 30 seconds of ballet.

Green shows the VBR file, which had an overall VMAF score of 97.23. Red shows capped CRF, which scored 94.59.

Looking at the graph, VBR appears more CBR-ish than expected, with a relatively consistent bitrate until the end. According to FFBitrateViewer, the average VBR bitrate was 4111, with a peak bitrate of 7361. In comparison, capped CRF averaged 1180 and peaked at 3,418.

As you see, capped CRF is very VBR-like, clearly adjusting to the changes in encoding complexity. This obviously contributes to its effectiveness.

Capped CRF vs VBR Performance

	40 LP (No Limit)		8 Logical Processors	
	Encoding Speed	CCRF increase	Encoding Speed	CCRF increase
Football				
Capped CRF	1.7		1.26	
VBR	1.38	23.2%	1.15	9.6%
Meridian				
Capped CRF	2.63		1.95	
VBR	2.03	29.6%	1.65	18.2%

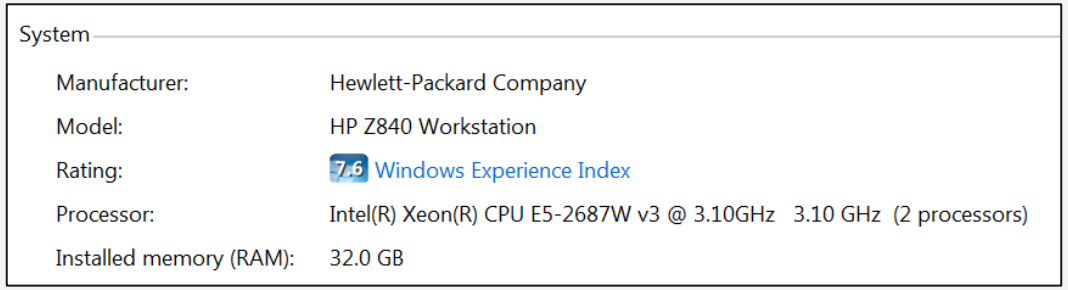

I tested the throughput of two files, Football and Meridian, on a 40-core workstation, first without limiting the logical processors (LP) in the command string, and then limiting the LP to 8. In both cases, capped CRF delivered higher throughput, though less so with 8 logical processors.

My contact on the SVT-AV1 dev team advised that at least part of capped CRF's speed advantage related to the lower bitrate. This is likely why capped CRF delivered a greater throughput advantage with Meridian (52% bitrate reduction over VBR) than the 30 fps Football clip (24% bitrate reduction).

Still, since capped CRF almost always delivers lower bandwidth than VBR, throughput should also be greater.

That said, this performance variability is concerning because CPU utilization will vary with the footage. If you're operating near 80% CPU utilization and clip complexity jumps, the system may not have sufficient CPU cycles to maintain realtime transcoding.

You need to leave plenty of performance headroom for any clip with mixed complexity or (gulp) consider ASIC-based hardware, which isn't impacted by content complexity.

Identifying the Optimal CRF Values

Obviously, you need to plug in a CRF value and a bitrate cap to encode using capped CRF. With most codecs, one value will work well for most content types. With SVT-AV1, this isn't the case.

Key finding; The optimal CRF value varies with content type. You should ascertain the optimal value(s) for your footage before deploying or even testing capped CRF.

Discussion: The CRF value sets the quality level of the video file and dramatically impacts the effectiveness of capped CRF. Lower CRF values increase quality, so CRF 19 quality is better than CRF 21. If you set CRF quality too high (like 15), you limit the bandwidth savings and exceed necessary quality levels. If you set it too low (like 45), quality may be inadequate.

The next section shows how we identified the optimal CRF value for each content type.

Recommended technique: Here's the recommended technique.

1. On a category-by-category basis, identify the CRF value that delivers a VMAF score of about 93-95 (see here). You do this by encoding at different CRF values without a cap.

2. Once you identify the optimal CRF value, encode with a bitrate cap, which usually equals the target VBR rate.

3. Then encode using VBR and measure and compare:

- Encoding time
- Bitrate savings
- Average quality differential
- Low-frame differential

The next few slides show the categories assessed, the values checked, and the recommended values used.

What CRF Value? Media

- Elektra - snippet from movie
- Freedom - concert video
- Haunted - snippet from faux horror video
- India - Harmonic test clip
- Meridian - snippet from movie
- Orchestra - snippet from concert
- Tears of Steel - snippet from movie
- Zoolander - opening sequence from movie

Media 30fps SVT-AV1	Bitrate	VMAF	Low Frame	1%
CRF 31	3305	95.73	85.15	89.98
CRF 33	2840	95.15	83.85	88.94
CRF 35	2150	94.55	80.81	86.48
CRF 37	2470	93.86	82.45	87.83

Tested using 33 to target VMAF 95

What CRF Value? 60fps Sports

- Football - from Harmonic test clip
- River Plate - from soccer match
- Soccer 1 - clip 1 from different soccer match
- Soccer 2 - clip 2

Total 60fps sports SVT-AV1	Bitrate	VMAF	Low Frame	1%
CRF 37	5247	94.41	83.72	88.49
CRF 39	**4595**	**93.69**	**81.60**	**87.35**
CRF 41	3566	92.06	76.88	84.89
CRF 43	4034	92.90	79.31	86.16

Tested using 36 to target VMAF 95

What CRF Value? 30fps Sports

- Basketball - snippet from basketball video
- Football - snippet from Harmonic test clip
- Hockey - snippet from hockey video
- Skateboard - snippet from skateboarding video
- Soccer - snippet from soccer match

30 fps Sports SVT-AV1	Bitrate	VMAF	Low Frame	1%
CRF 37	3975	97.11	81.22	87.94
CRF 39	3504	96.51	78.52	86.34
CRF 41	3093	95.77	76.42	84.67
CRF 43	2677	94.97	73.81	82.91

Tested using 42 to target VMAF 95

What CRF Values? Animation

- Big Buck Bunny- snippet from test clip
- El Ultimo - snippet from very simple 2D animation cartoon
- Sintel - snippet from test clip

Animations SVT-AV1	Bitrate	VMAF	Low Frame	1%
CRF 37	1935	95.96	83.36	90.18
CRF 39	1738	95.46	81.48	88.84
CRF 41	1567	94.89	78.98	87.44
CRF 43	1419	94.25	76.27	85.92

Tested using 41 to target VMAF 95

What CRF Value? - Office

- Epiphan - screencam
- Talkinghead - simple talking head
- Test - 30 seconds talkinghead, 30-seconds ballet
- Tutorial - PowerPoint with small talking head

Office SVT-AV1	Bitrate	VMAF	Low Frame	1%
CRF 37	822	95.72	88.58	93.02
CRF 39	753	95.46	89.04	92.50
CRF 41	696	95.19	88.35	91.90
CRF 43	641	94.88	87.61	91.27

Tested using 43 to target VMAF 95

What CRF Savings - Average VMAF

	Constant Bitrate			Capped CRF			Delta		
	Bitrate	VMAF	Low-Frame	Bitrate	VMAF	Low Frame	Bitrate	VMAF	Low Frame
Media	4058	96.14	82.88	2640	94.82	77.99	-34.9%	-1.4%	-5.9%
60 fps sports	5972	94.92	79.81	4872	93.39	56.48	-18.4%	-1.6%	-29.2%
30 fps sports	4283	97.04	78.28	2634	94.29	62.19	-38.5%	-2.8%	-20.6%
Animation	4187	97.63	87.77	1527	94.75	78.98	-63.5%	-2.9%	-10.0%
Office	4041	97.25	92.98	810	94.81	87.70	-80.0%	-2.5%	-5.7%
Average	**4,437**	**96.50**	**83.71**	**2,567**	**94.46**	**72.85**	**-44.2%**	**-2.1%**	**-13.5%**

Key Findings:

- **Overall, capped CRF delivered a bitrate savings of 44.2%** while reducing average VMAF from 96.50 to 94.46, which is still within the recommended 93 - 95 target, so this drop should not impact QoE.
- Low frame quality ***dropped*** by 13.5% on average, with sports clips showing the largest drop. We explore these findings below.

Discussion: As stated, the extent of the bitrate reduction relates to the bitrate cap and bitrate for the VBR comparison. At 4.5 Mbps, the benefits are probably overstated for animations and office footage, though about right for all other categories (60fps sports was 6 Mbps). You can perform the simple math to adjust these to whatever bitrate you think is appropriate.

Regarding low-frame scores, not all low-frame values translate to quality deficits that impact viewer quality of experience (QoE). Some are simply too transient to notice, some hidden by fast transitions or intentionally grainy source footage. For this reason, you have to actually examine the associated frames to gauge if they would impact QoE.

We do that next.

Low- Frame Analysis - Media

	Constant Bitrate			Capped CRF			Delta		
Media	**Bitrate**	**VMAF**	**Low-Frame**	**Bitrate**	**VMAF**	**Low Frame**	**Bitrate**	**VMAF**	**Low Frame**
Elektra	3,762	96.21	89.58	1,717	94.61	88.55	-54.4%	-1.7%	-1.2%
freedom.mp4	3,860	96.01	86.80	3,539	95.83	81.91	-8.3%	-0.2%	-5.6%
haunted.mp4	4,364	93.69	57.01	2,354	90.91	56.54	-46.1%	-3.0%	-0.8%
india.mp4	4,386	96.57	87.33	2,928	95.19	82.60	-33.2%	-1.4%	-5.4%
meridian.mp4	3,663	96.56	87.21	1,747	95.50	83.46	-52.3%	-1.1%	-4.3%
orchestra.mp4	4,018	94.46	89.20	3,176	93.65	83.28	-21.0%	-0.9%	-6.6%
tos.mp4	4,254	97.19	79.70	2,544	95.50	80.04	-40.2%	-1.7%	0.4%
zoo.mp4	4,158	98.42	86.24	3,116	97.34	67.56	-25.1%	-1.1%	-21.7%
Average	**4,058**	**96.14**	**82.88**	**2,640**	**94.82**	**77.99**	**-34.9%**	-1.4%	-5.9%

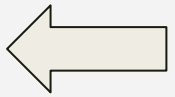

Here we analyze low frame performance for media files, as listed above.

Key finding: On average, capped CRF:

- ***Reduced bandwidth*** by 34.9%
- ***Lowered quality*** to 94.82 VMAF, still within the top-rung target range for most publishers
- ***Reduced low-frame quality*** by 5.9%

Discussion: One occasional by-product of capped CRF encoding is excessive low-frame scores. Before deploying capped CRF, you should measure low-frame scores and assess:

- How prevalent they are for each content type.
- How noticeable the quality drops would to the viewer. This involves duration and severity.

Let's explore Zoo (for Zoolander), the one file in this group that showed a significantly lower low-frame score.

Low-Frame Analysis - Zoo

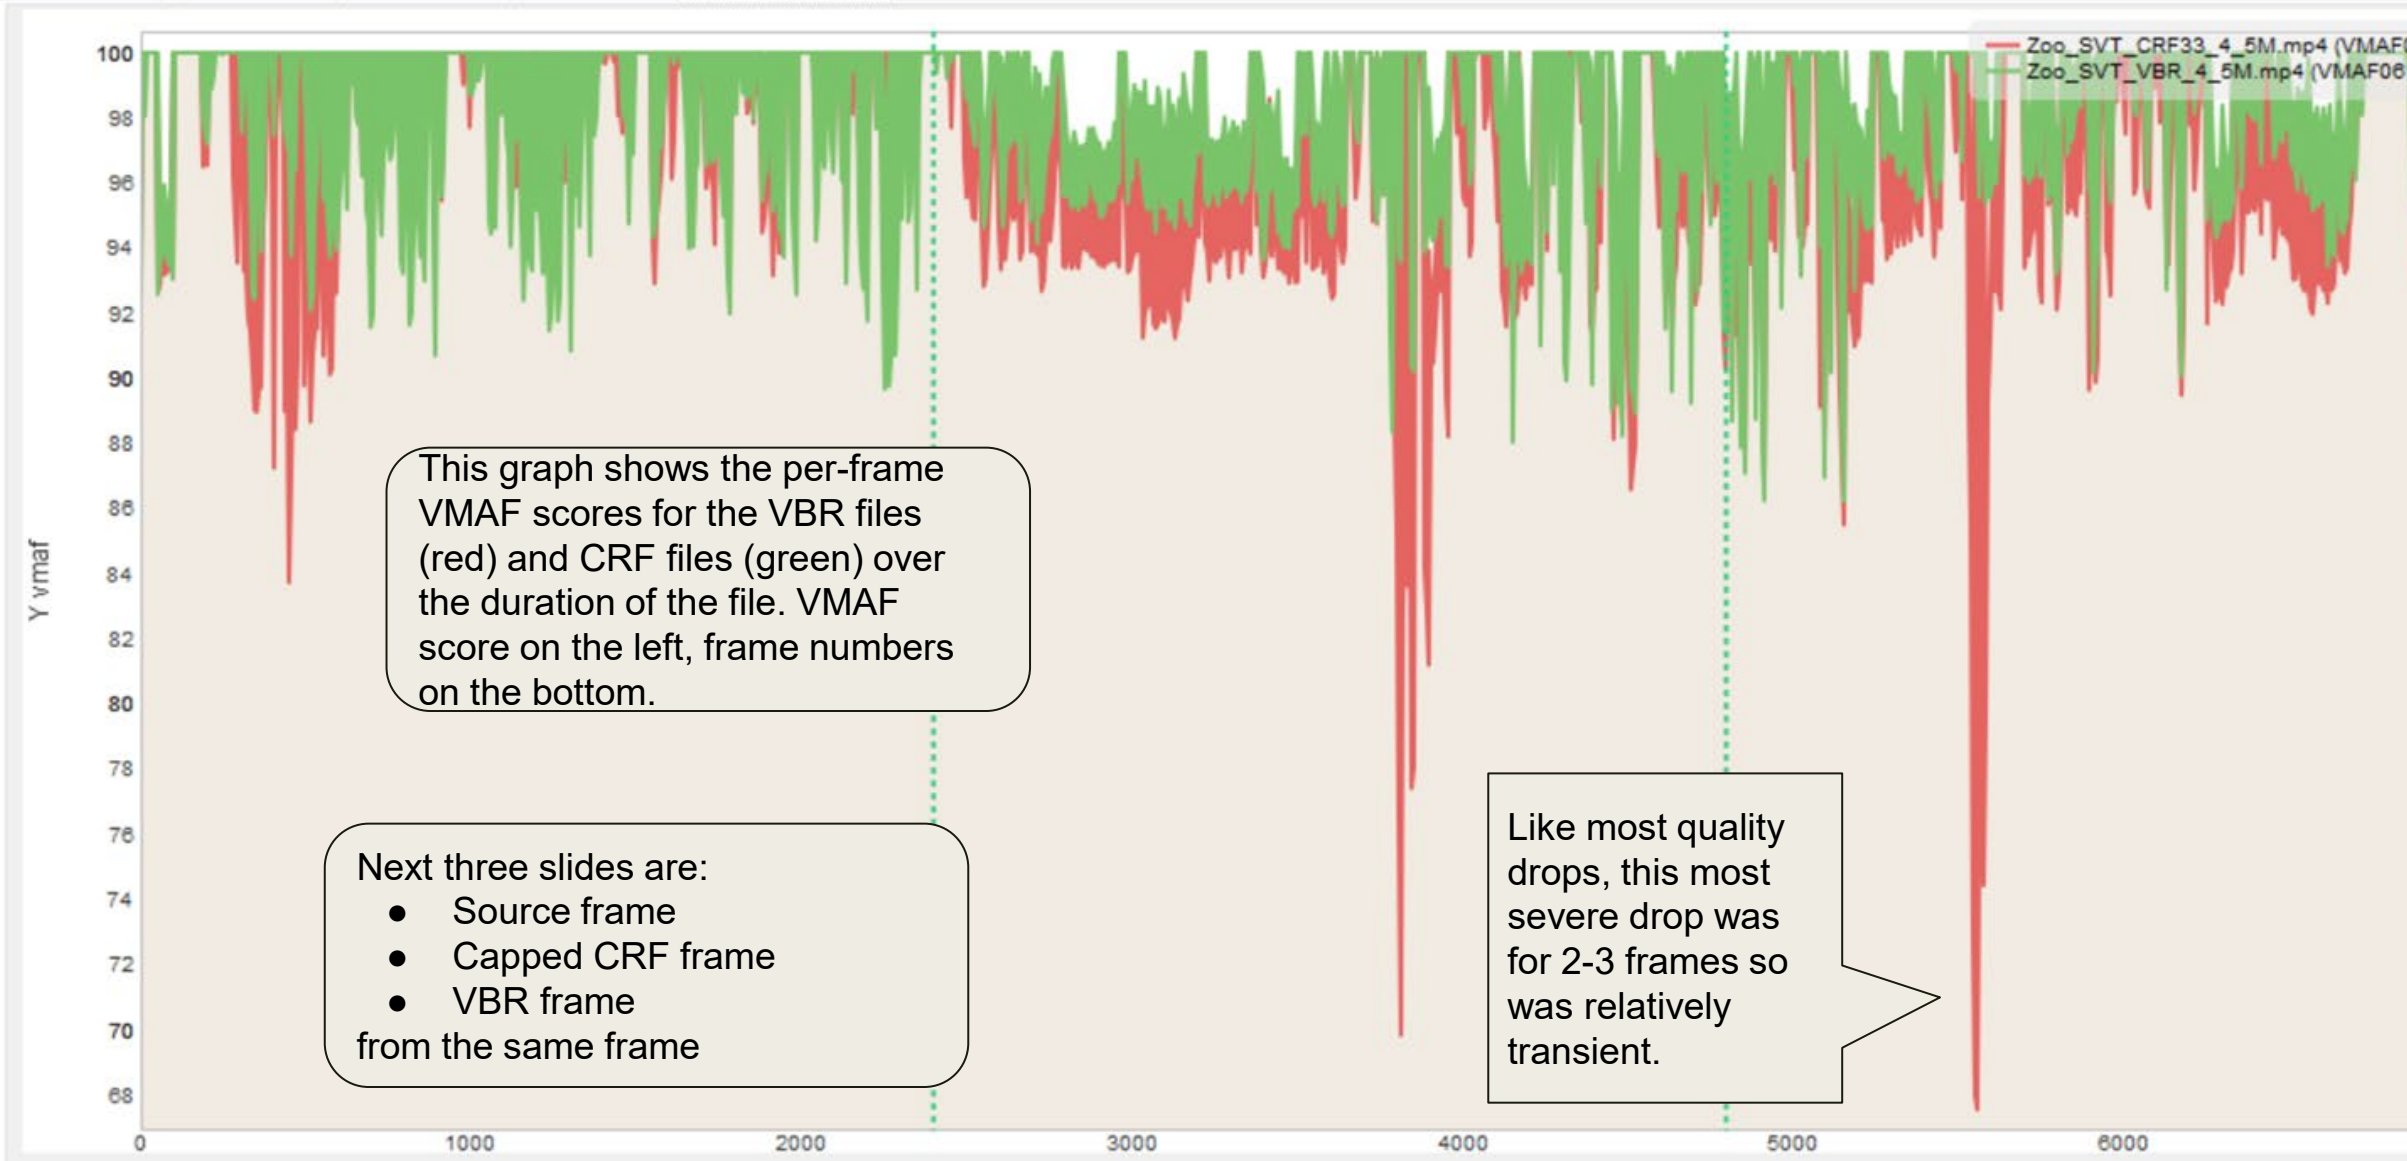

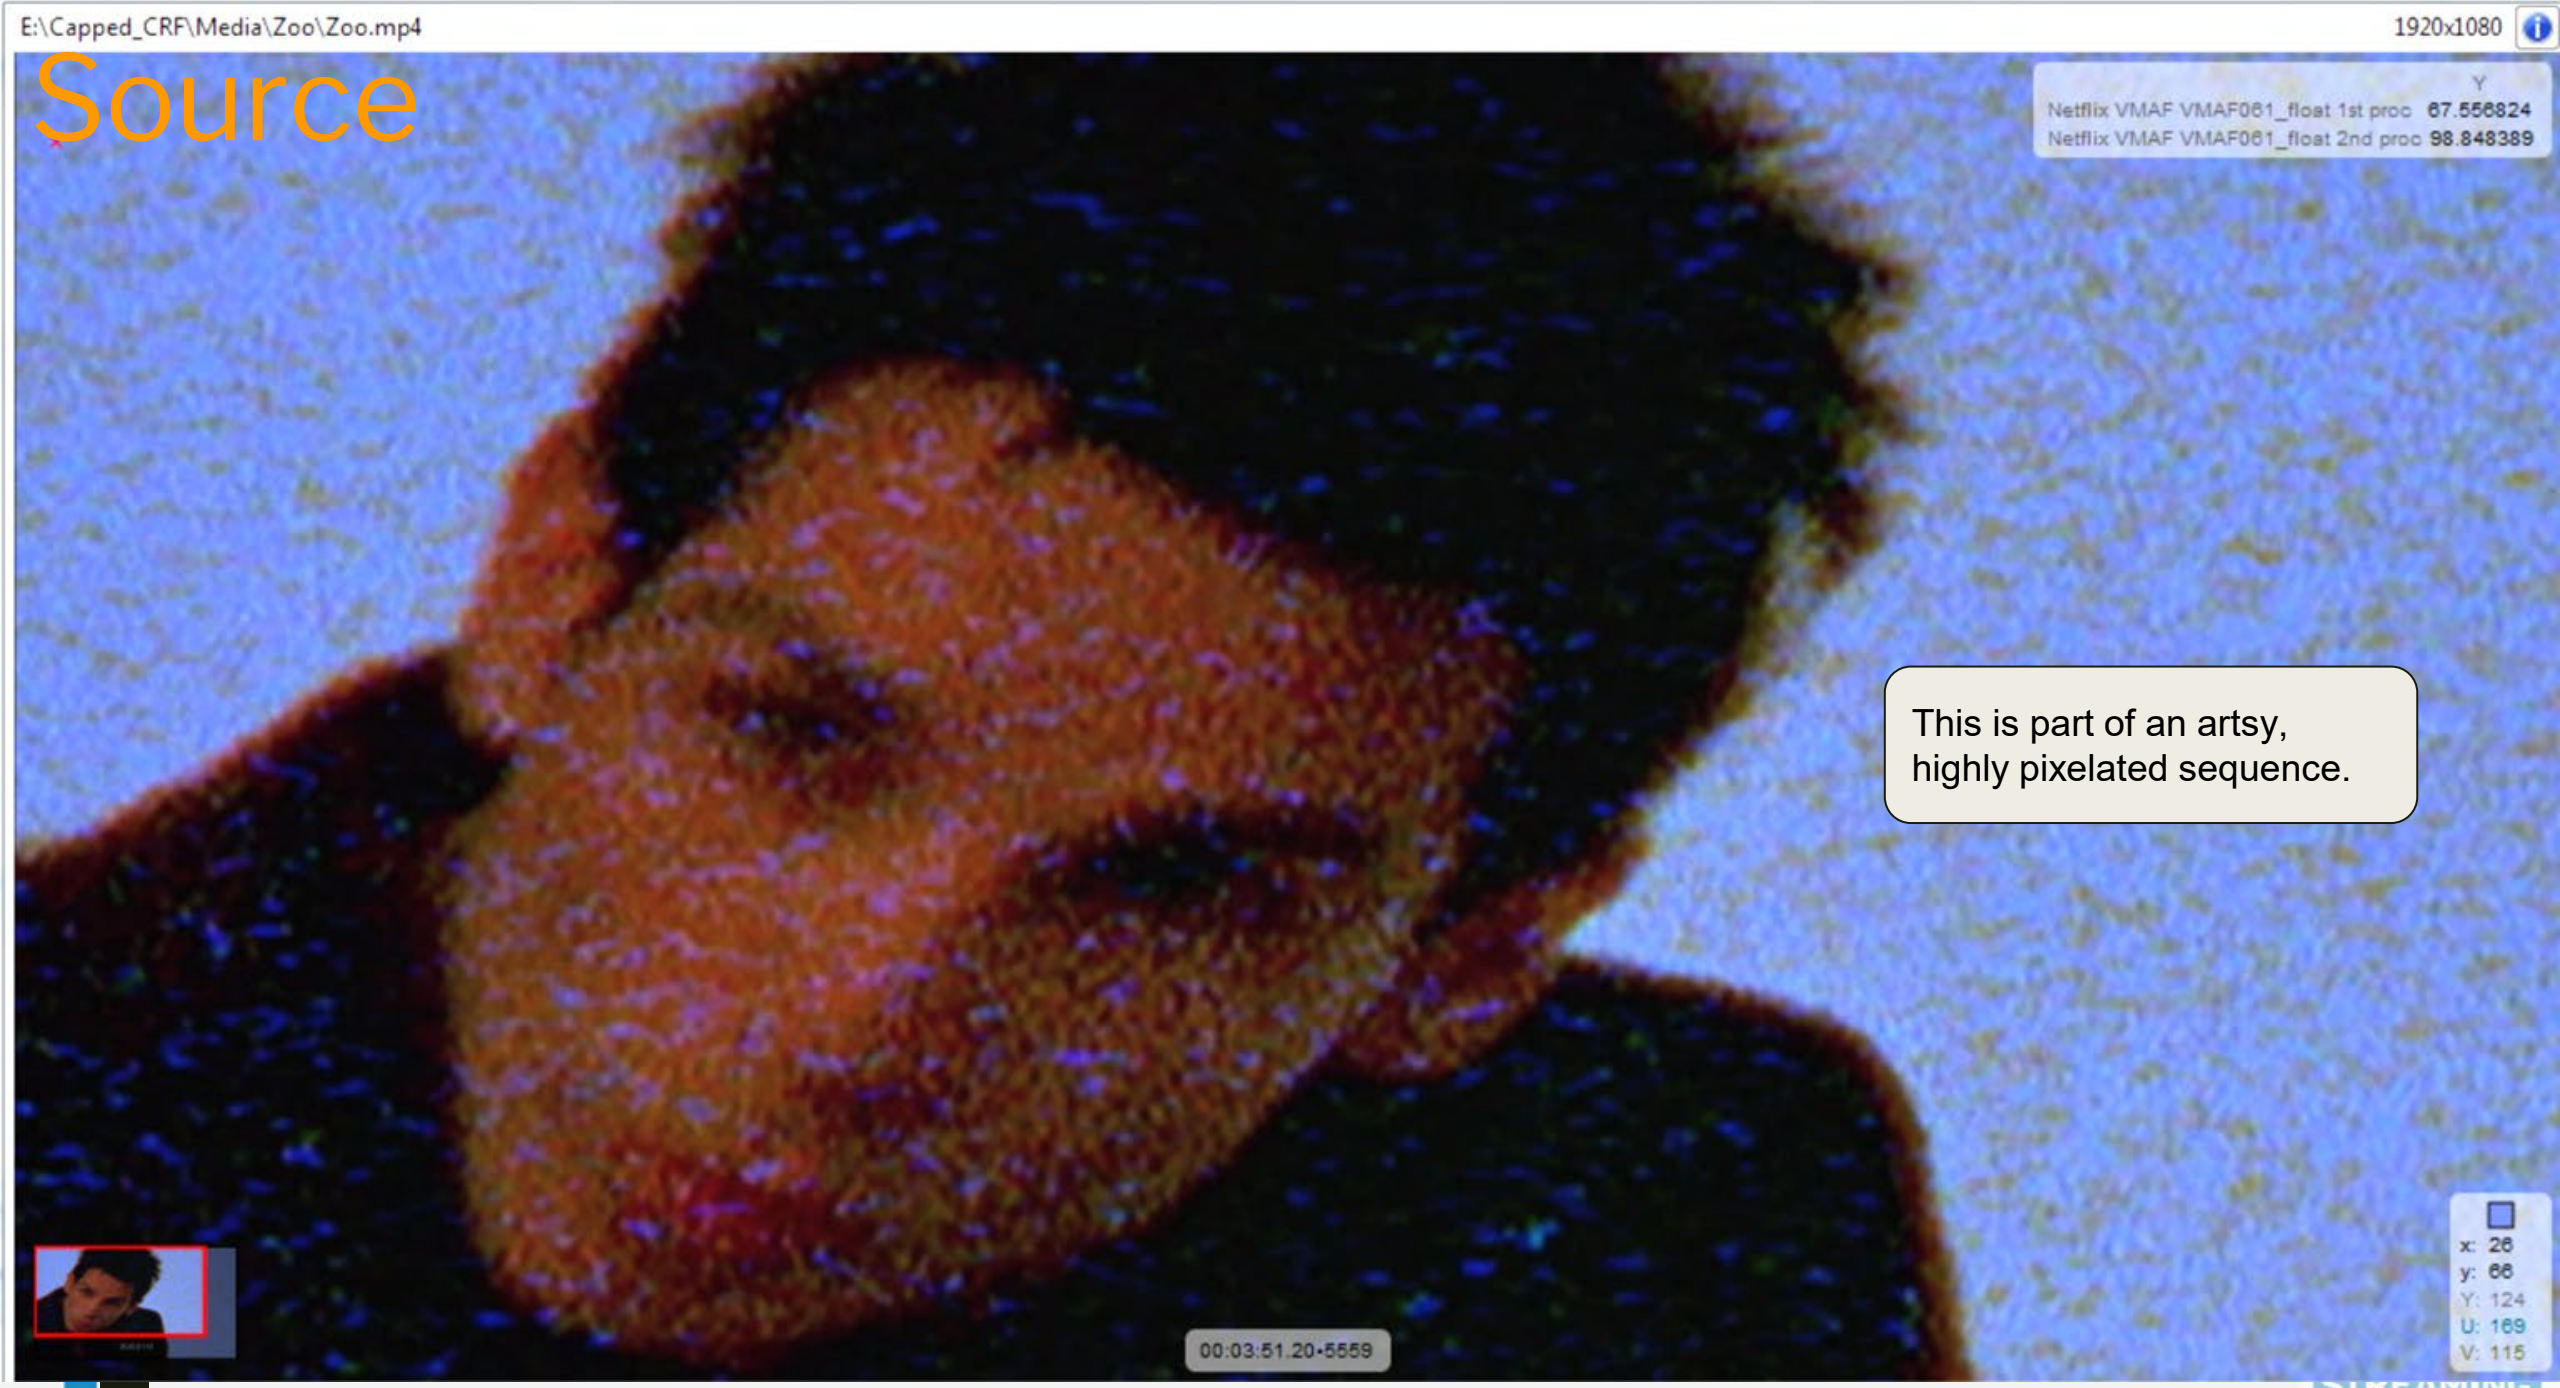
E:\Capped_CRF\Media\Zoo\Zoo.mp4
1920x1080
Source
Y
Netflix VMAF VMAF061_float 1st proc 67.556824
Netflix VMAF VMAF061_float 2nd proc 98.848389
This is part of an artsy, highly pixelated sequence.
00:03:51.20-5559
x: 26
y: 66
Y: 124
U: 169
V: 115

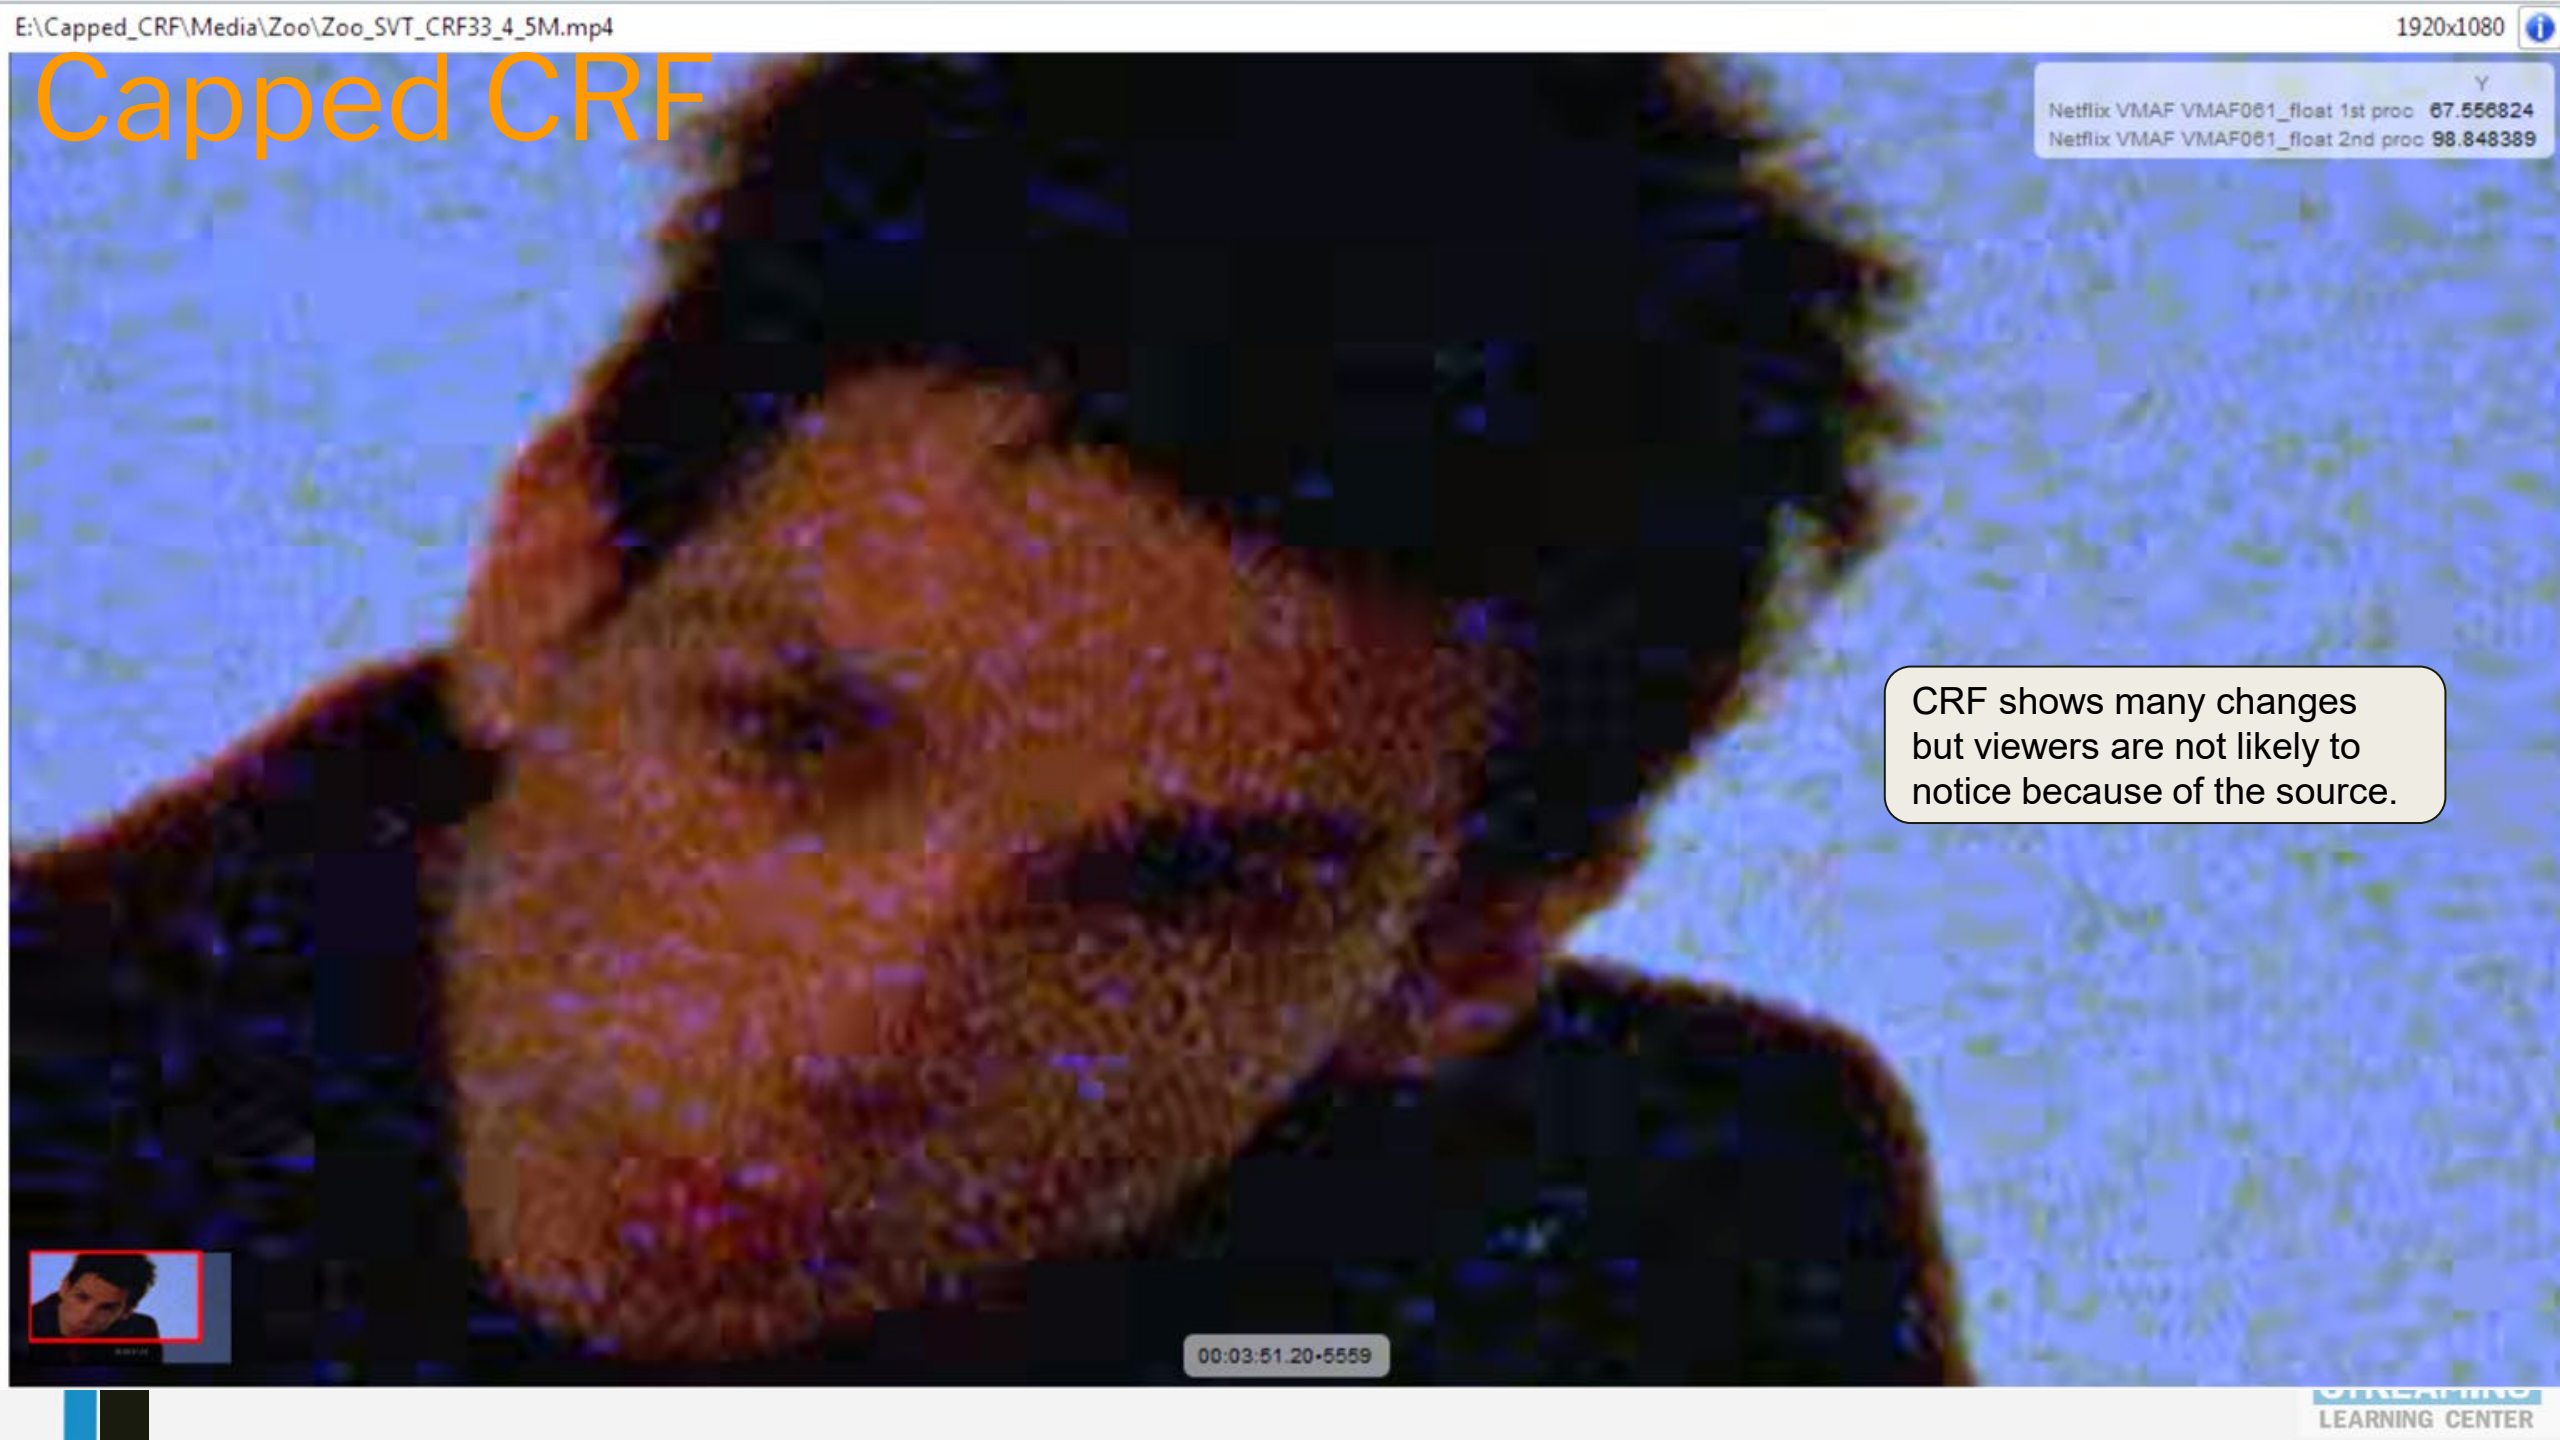
E:\Capped_CRF\Media\Zoo\Zoo_SVT_CRF33_4_5M.mp4
1920x1080
Capped CRF
Y
Netflix VMAF VMAF061_float 1st proc 67.556824
Netflix VMAF VMAF061_float 2nd proc 98.848389
CRF shows many changes but viewers are not likely to notice because of the source.
00:03:51.20+5559


E:\Capped_CRF\Media\Zoo\Zoo_SVT_VBR_4_5M.mp4
1920x1080
VBR
Y
Netflix VMAF VMAF061_float 1st proc 67.556824
Netflix VMAF VMAF061_float 2nd proc 98.848389
VBR looks close to the original. None of the major differences would have been noticeable to a viewer.
00:03:51.20+5559

Low- Frame Analysis - Media

	Constant Bitrate			Capped CRF			Delta		
Media	**Bitrate**	**VMAF**	**Low-Frame**	**Bitrate**	**VMAF**	**Low Frame**	**Bitrate**	**VMAF**	**Low Frame**
Elektra	3,762	96.21	89.58	1,717	94.61	88.55	-54.4%	-1.7%	-1.2%
freedom.mp4	3,860	96.01	86.80	3,539	95.83	81.91	-8.3%	-0.2%	-5.6%
haunted.mp4	4,364	93.69	57.01	2,354	90.91	56.54	-46.1%	-3.0%	-0.8%
india.mp4	4,386	96.57	87.33	2,928	95.19	82.60	-33.2%	-1.4%	-5.4%
meridian.mp4	3,663	96.56	87.21	1,747	95.50	83.46	-52.3%	-1.1%	-4.3%
orchestra.mp4	4,018	94.46	89.20	3,176	93.65	83.28	-21.0%	-0.9%	-6.6%
tos.mp4	4,254	97.19	79.70	2,544	95.50	80.04	-40.2%	-1.7%	0.4%
zoo.mp4	4,158	98.42	86.24	3,116	97.34	67.56	-25.1%	-1.1%	-21.7%
Average	**4,058**	**96.14**	**82.88**	**2,640**	**94.82**	**77.99**	**-34.9%**	-1.4%	-5.9%

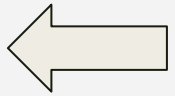

Discussion: Though we see confirmation of the VMAF score, the highly pixelated source footage makes it unlikely that the viewer would notice during real time playback.

Conclusion: With general purpose entertainment footage, capped CRF offers the potential for significant bitrate savings and relatively low risk of noticeable transient quality issues.

Recommendation: **Deploy but verify**. Seems low risk but check with your own footage.

Low- Frame Analysis - 60fps Sports

60 fps sports	Bitrate	VMAF	Low-Frame	Bitrate	VMAF	Low Frame	Bitrate	VMAF	Low Frame
football.mp4	5,919	95.79	81.95	5,125	93.57	39.07	-13.4%	-2.3%	-52.3%
riverplate.mp4	5,869	94.10	70.20	3,809	91.69	35.18	-35.1%	-2.6%	-49.9%
soccer1.mp4	6,049	94.72	85.41	5,585	94.18	81.28	-7.7%	-0.6%	-4.8%
soccer2.mp4	6,052	95.06	81.67	4,969	94.11	70.37	-17.9%	-1.0%	-13.8%
Average	**5,972**	**94.92**	**79.81**	**4,872**	**93.39**	**56.48**	**-18.4%**	**-1.6%**	**-29.2%**

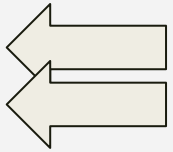

Overall, in this category, capped CRF:

- ***Reduced bandwidth*** slightly
- ***Reduced overall quality*** to the bottom end of the target zone
- Showed ***very low low-frame*** scores on two of four clips.

Let's look at the Football and Riverplate clips.

Low-Frame Analysis - Football

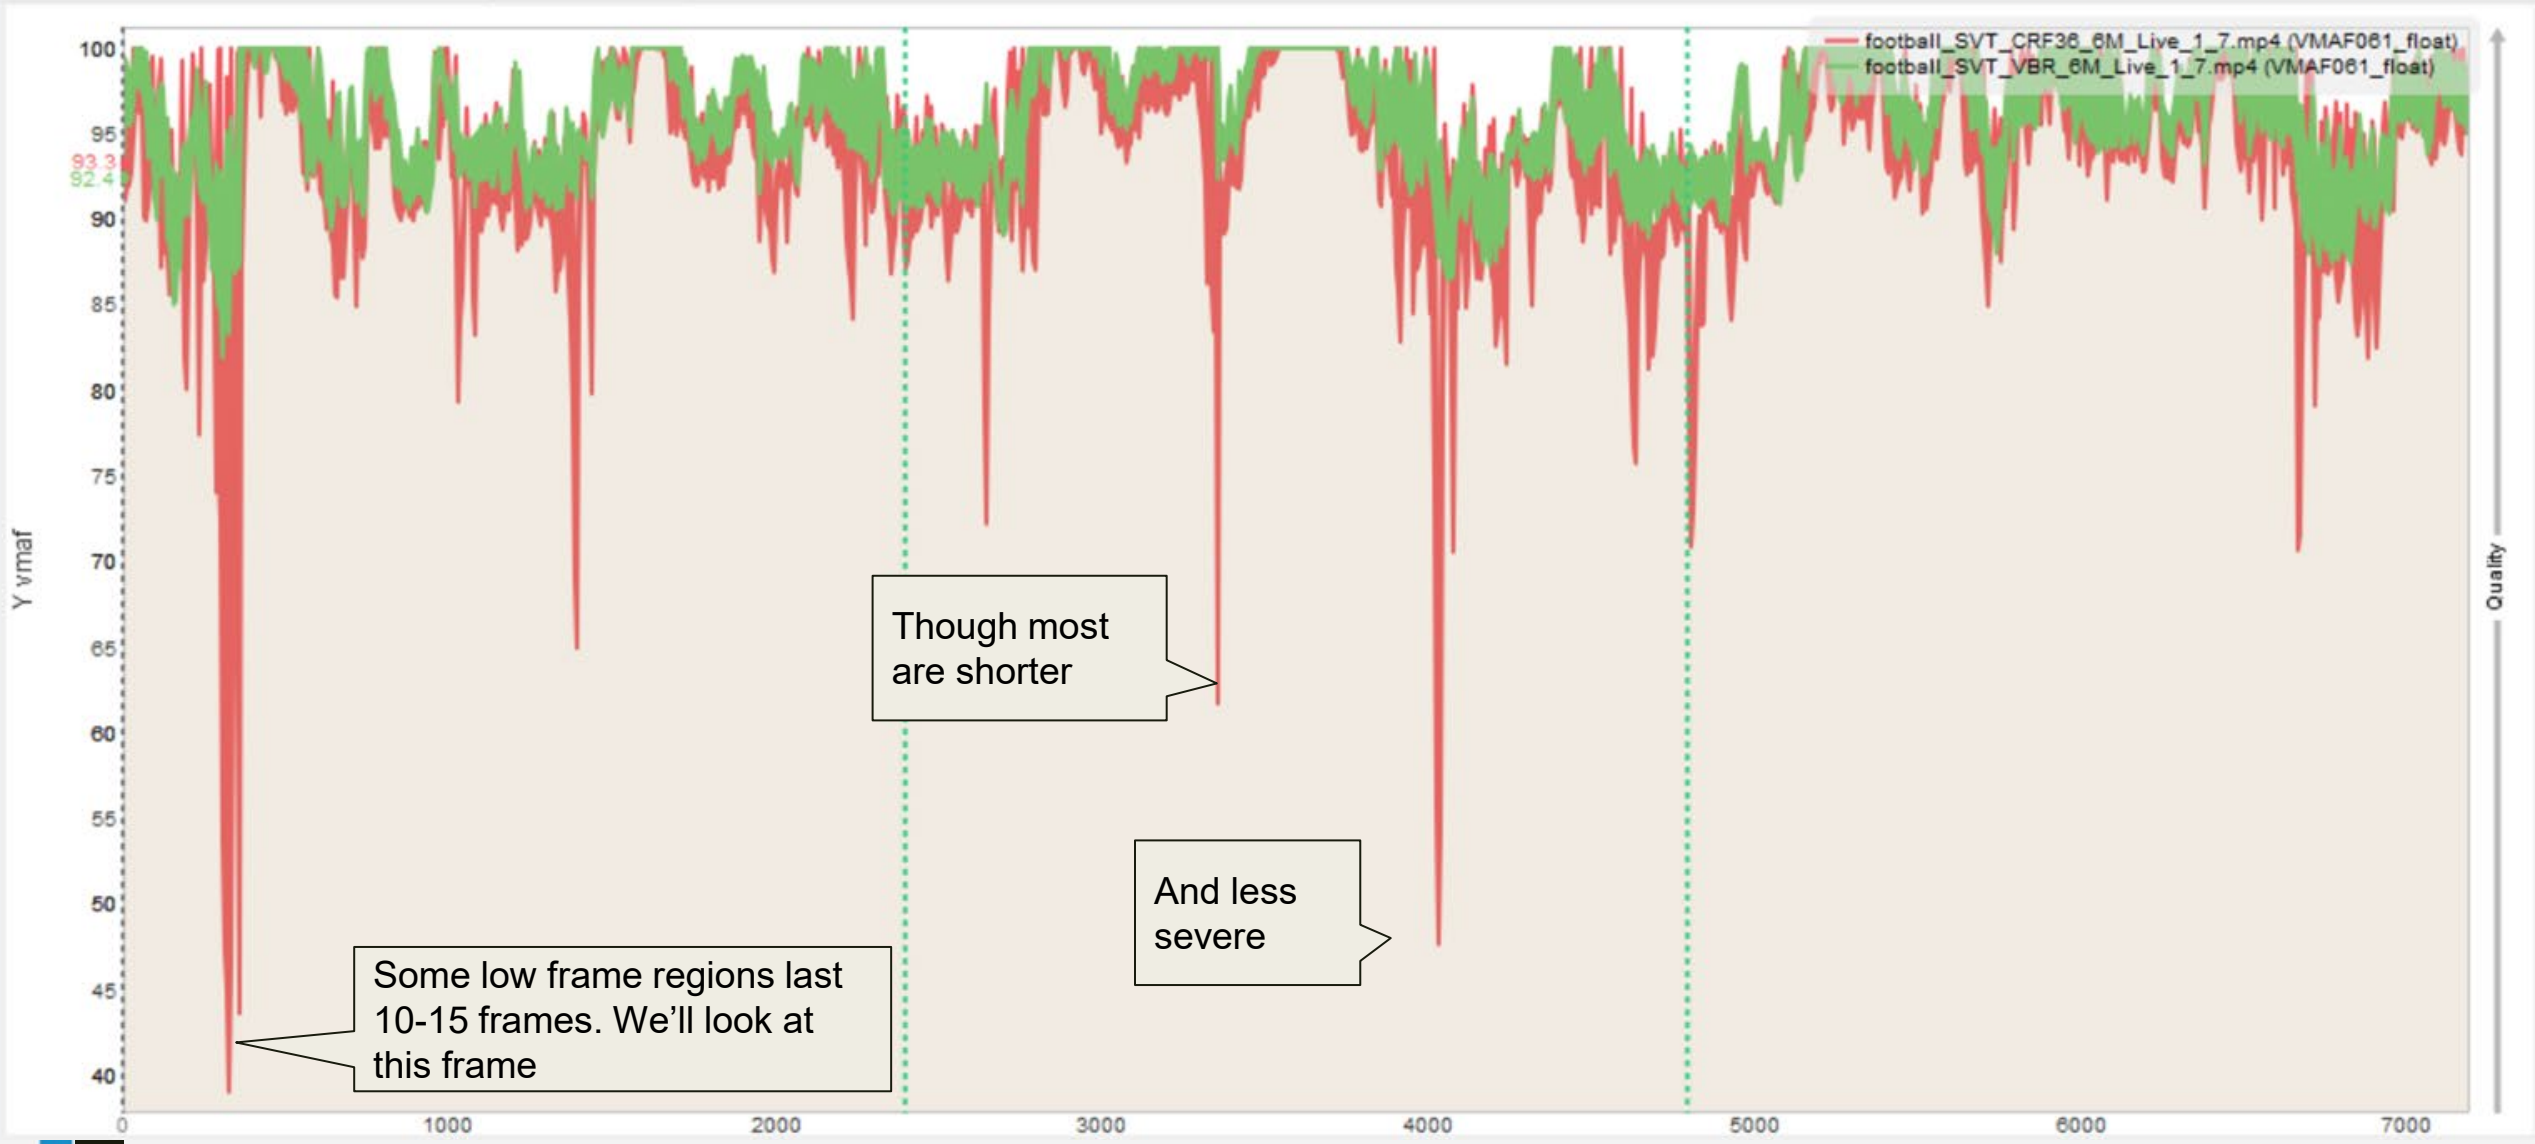


E:\Capped_CRF\60fps_sports\Football\football.mp4
1920x1080
Source
Netflix VMAF VMAF061_float 1st proc 39.071461
Netflix VMAF VMAF061_float 2nd proc 85.835419
00:00:05.27+327


E:\Capped_CRF\60fps_sports\Football\football_SVT_CRF36_6M_Live_1_7.mp4
1920x1080
Capped CRF
Anyone who doubts VMAF's accuracy should look at this frame, which clearly deserved its 39 score. You can check the file name on the top left to confirm this was the capped CRF frame. Still, it's only 3-4 frames, so it's tough to say if a typical viewer would perceive the quality drop.
Y
Netflix VMAF VMAF061_float 1st proc 39.071461
Netflix VMAF VMAF061_float 2nd proc 85.835419
00:00:05.27+327
x: 134
y: 324
Y: 94
U: 96
V: 111


E:\Capped_CRF\60fps_sports\Football\football_SVT_VBR_6M_Live_1_7.mp4
1920x1080
VBR
Netflix VMAF VMAF061_float 1st proc 39.071461
Netflix VMAF VMAF061_float 2nd proc 85.835419
VBR is not pristine, but clearly much better.
00:00:05.27•327

Low-Frame Analysis - RiverPlate

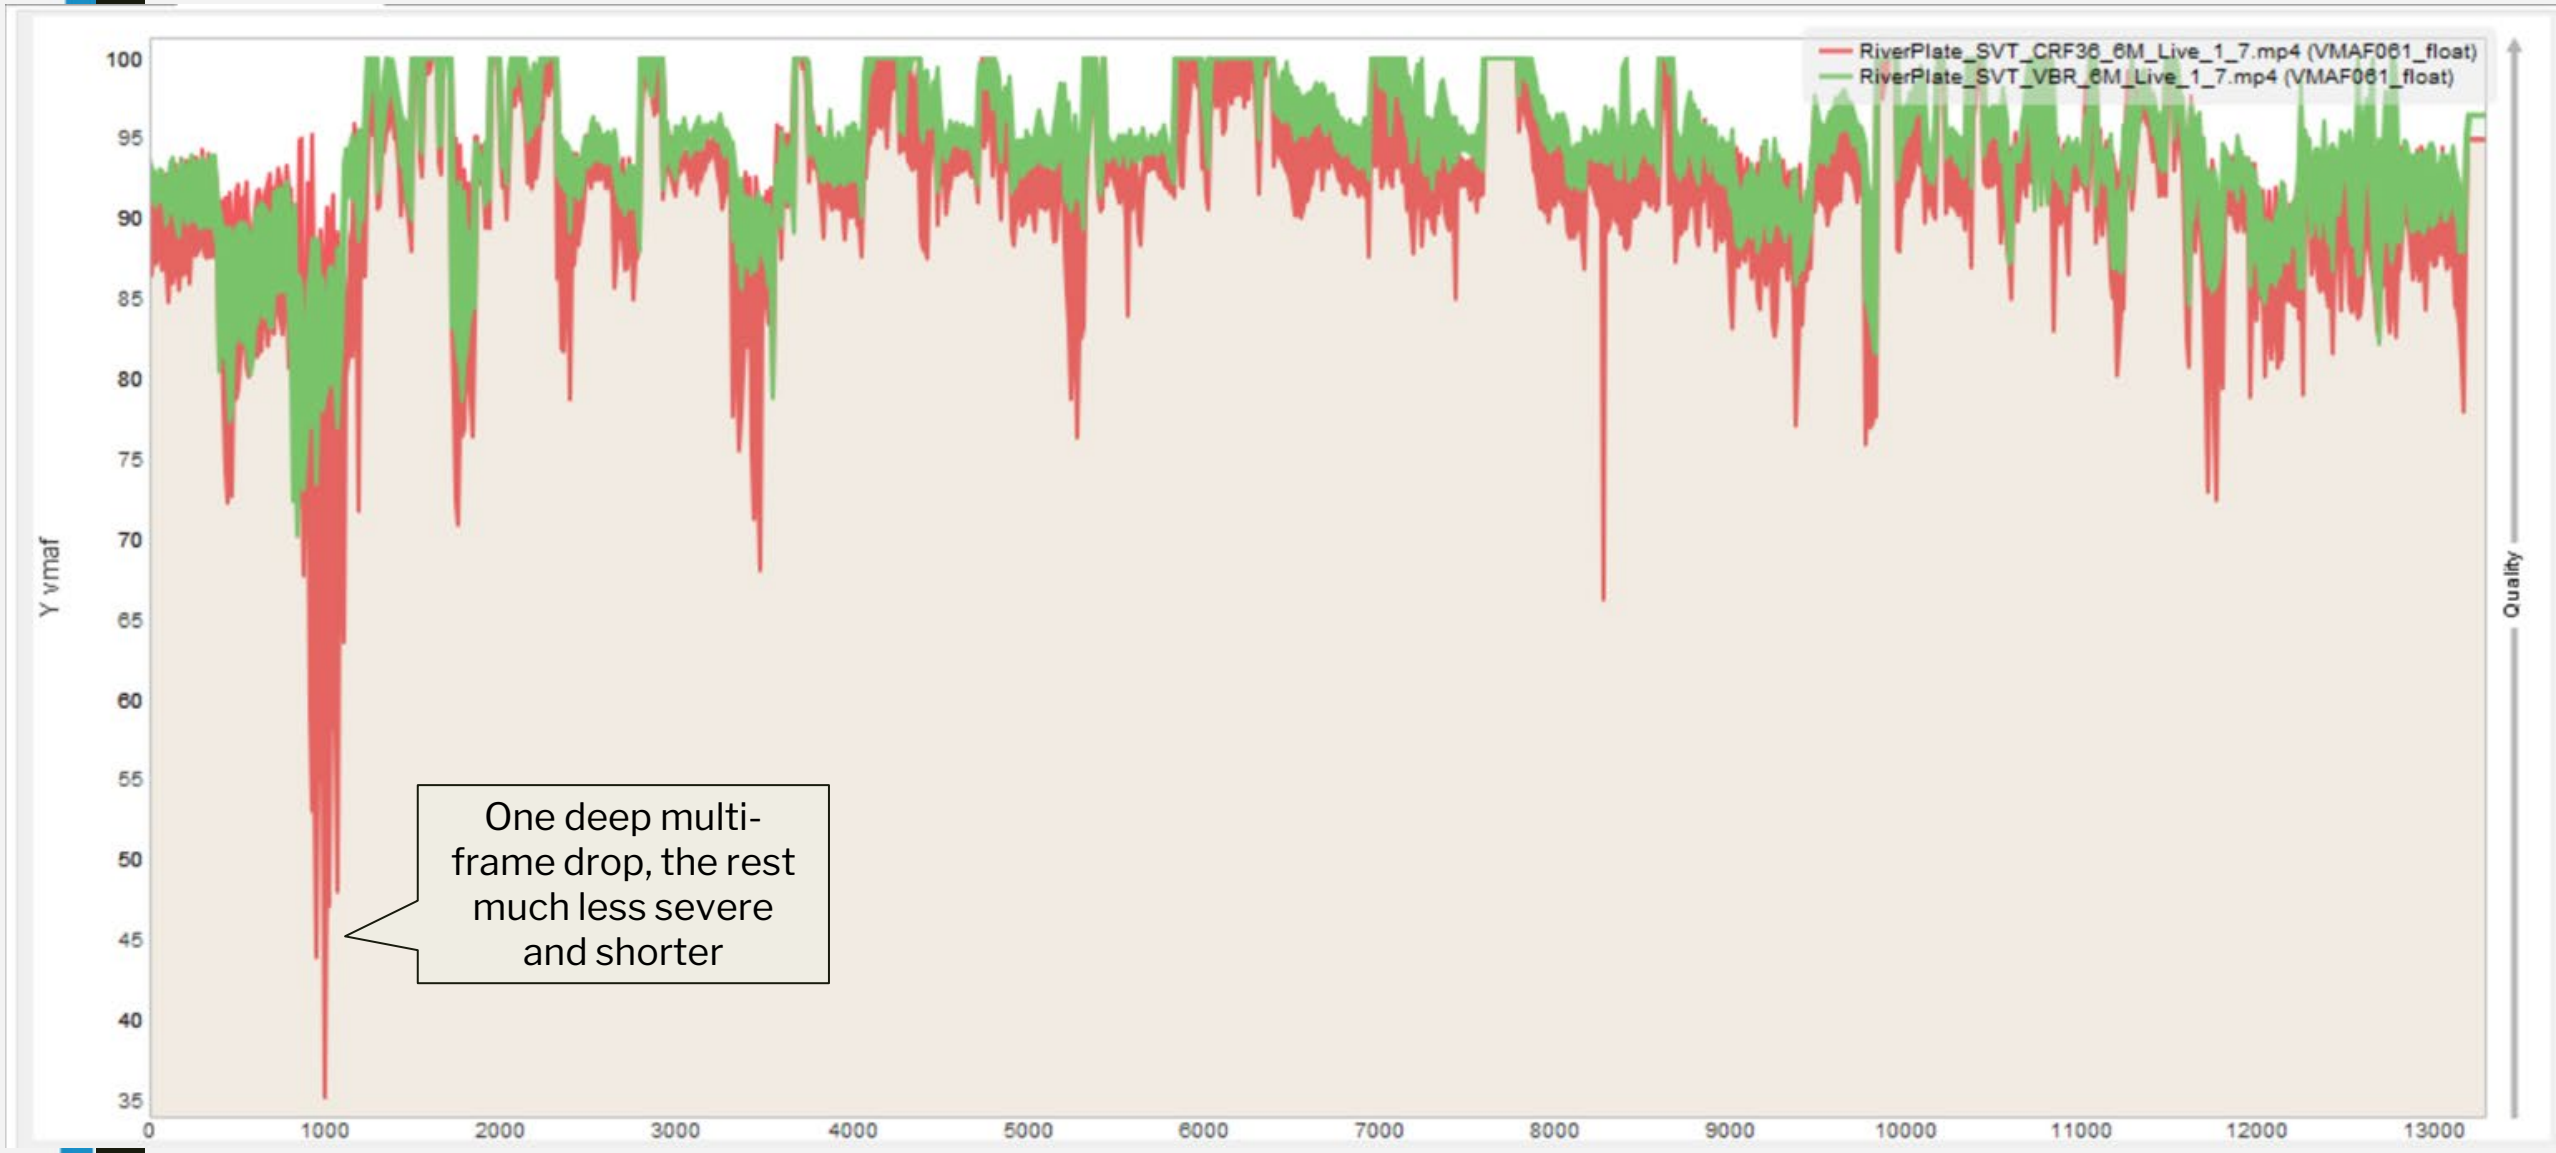


E:\Capped_CRF\60fps_sports\riverplate\RiverPlate.mp4
1920x1080
Source
Y
Netflix VMAF VMAF061_float 1st proc 35.181129
Netflix VMAF VMAF061_float 2nd proc 84.772751
Coca-Cola
Coca-Cola
00:00:16.48+1008


E:\Capped_CRF\60fps_sports\riverplate\RiverPlate_SVT_CRF36_6M_Live_1_7.mp4
1920x1080
Capped CRF
Y
Netflix VMAF VMAF061_float 1st proc 35.181129
Netflix VMAF VMAF061_float 2nd proc 84.772751
If you look closely enough, you can see the changes that produced a VMAF score of 35, though it feels unlikely that viewers watching in real time would notice.
Coca-Cola
00:00:16.48+1008


E:\Capped_CRF\60fps_sports\riverplate\RiverPlate_SVT_VBR_6M_Live_1_7.mp4
1920x1080
VBR
Netflix VMAF VMAF061_float 1st proc 35.181129
Netflix VMAF VMAF061_float 2nd proc 84.772751
Coca-Cola
Coca-Cola
00:00:16.48+1008

Low- Frame Analysis - 60fps Sports

60 fps sports	Bitrate	VMAF	Low-Frame	Bitrate	VMAF	Low Frame	Bitrate	VMAF	Low Frame
football.mp4	5,919	95.79	81.95	5,125	93.57	39.07	-13.4%	-2.3%	-52.3%
riverplate.mp4	5,869	94.10	70.20	3,809	91.69	35.18	-35.1%	-2.6%	-49.9%
soccer1.mp4	6,049	94.72	85.41	5,585	94.18	81.28	-7.7%	-0.6%	-4.8%
soccer2.mp4	6,052	95.06	81.67	4,969	94.11	70.37	-17.9%	-1.0%	-13.8%
Average	**5,972**	**94.92**	**79.81**	**4,872**	**93.39**	**56.48**	**-18.4%**	**-1.6%**	**-29.2%**

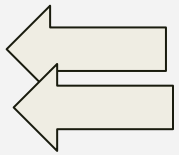

Discussion: Because this footage is fast moving and 60fps, it's challenging to compress, so capped CRF delivers ***reduced bandwidth savings***. Though perhaps not as noticeable as the frame graphs suggest, viewers ***may notice*** some of these quality issues.

So, less bandwidth benefit, more low-frame risk.

Recommendation: **Use with caution**. May not be appropriate for premium content but should be OK for more general distribution.

Low- Frame Analysis - 30fps Sports

30 fps sports	Bitrate	VMAF	Low-Frame	Bitrate	VMAF	Low Frame	Bitrate	VMAF	Low Frame
basketball.mp4	4,540	99.50	77.00	2,604	98.28	66.11	-42.6%	-1.2%	-14.1%
football.mp4	4,589	96.54	84.65	3,491	93.70	70.76	-23.9%	-2.9%	-16.4%
hockey.mp4	4,380	94.88	65.18	2,719	91.68	52.22	-37.9%	-3.4%	-19.9%
skateboard.mp4	4,033	96.95	85.68	1,313	91.22	66.90	-67.4%	-5.9%	-21.9%
soccer.mp4	3,873	97.33	78.89	3,041	96.58	54.96	-21.5%	-0.8%	-30.3%
Average	**4,283**	**97.04**	**78.28**	**2,634**	**94.29**	**62.19**	**-38.5%**	**-2.8%**	**-20.6%**

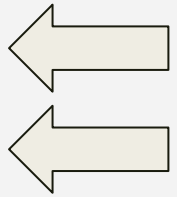

Discussion: With 30fps sports, capped CRF delivered:

- ***Significant bandwidth savings***
- ***Appropriate quality*** in the target range
- Some ***scary low-frame numbers***

Let's explore the frames in the hockey and soccer clips.

Low-Frame Analysis - Hockey

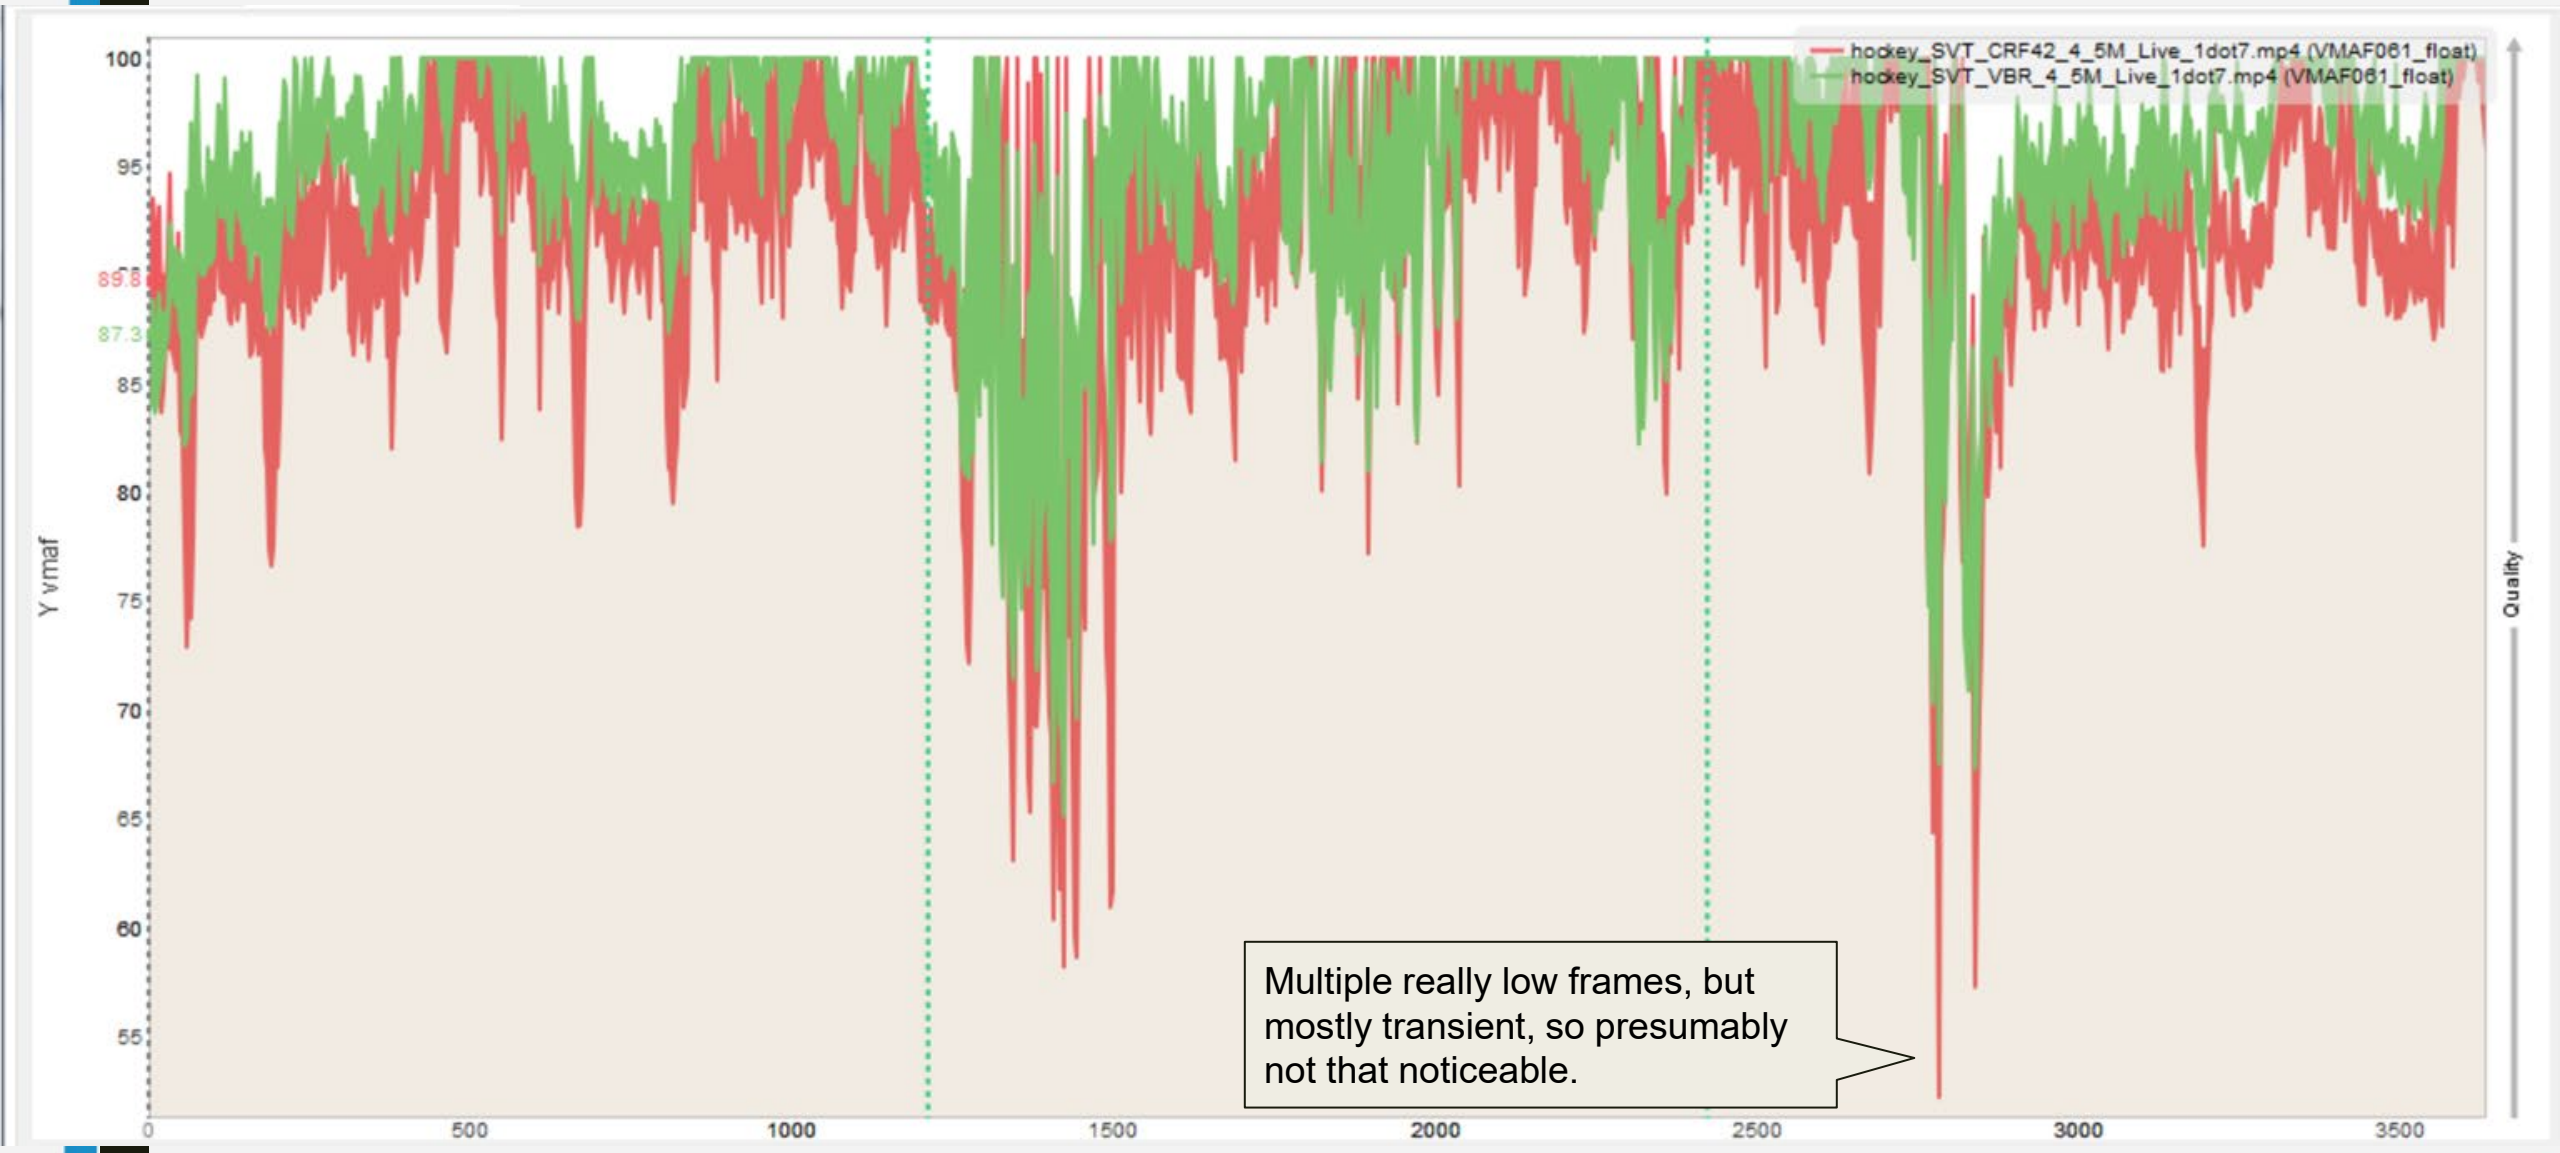


E:\Capped_CRF\30_sports\Hockey\Hockey.mp4
1920x1080
Source
Y
Netflix VMAF VMAF061_float 1st proc 52.220463
Netflix VMAF VMAF061_float 2nd proc 67.559456
x: 1391
y: 397
Y: 211
U: 139
V: 130
00:01:32.25•2783


E:\Capped_CRF\30_sports\Hockey\hockey_SVT_CRF42_4_5M_Live_1dot7.mp4
1920x1080
Capped CRF
Y
Netflix VMAF VMAF061_float 1st proc 52.220463
Netflix VMAF VMAF061_float 2nd proc 67.559456
Very fast game, frame is clearly degraded, but this is a very transient quality drop.
00:01:32.25+2783


E:\Capped_CRF\30_sports\Hockey\hockey_SVT_VBR_4_5M_Live_1dot7.mp4
1920x1080
VBR
Y
Netflix VMAF VMAF061_float 1st proc 52.220463
Netflix VMAF VMAF061_float 2nd proc 67.559456
Clearly better, though not perfect. Doubtful that viewers would notice in real time.
00:01:32.25-2783

Low-Frame Analysis - Soccer

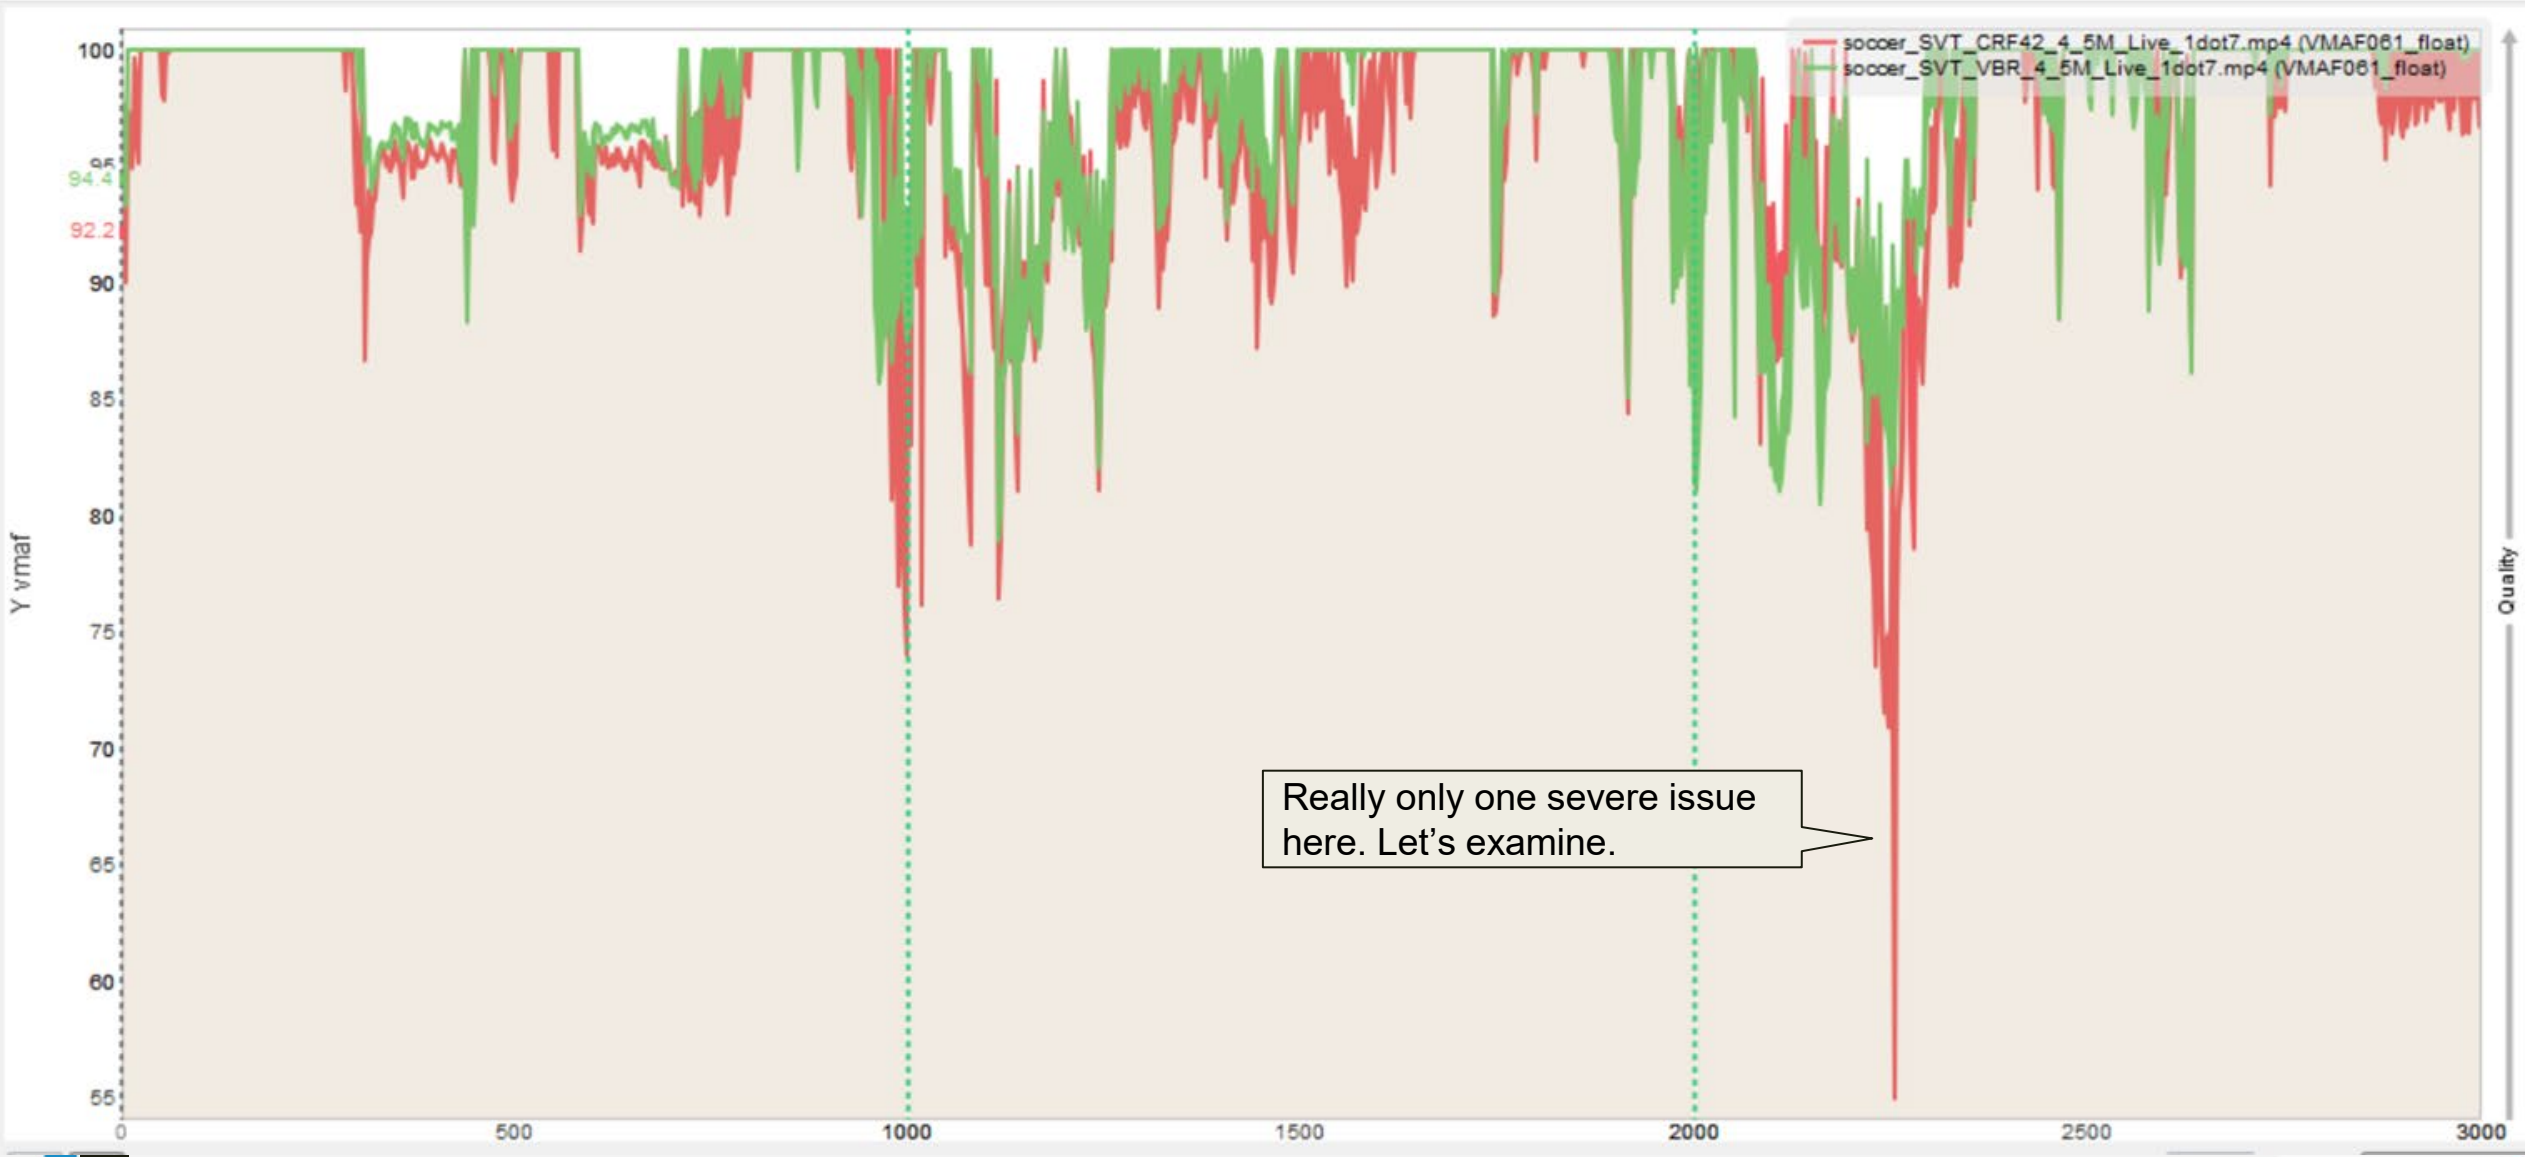


E:\Capped_CRF\30_sports\Soccer\Soccer.mp4
1920x1080
Source
Y
Netflix VMAF VMAF061_float 1st proc 54.962486
Netflix VMAF VMAF061_float 2nd proc 89.063103
Guga TV
00:01:30.4•2254


E:\Capped_CRF\30_sports\Soccer\soccer_SVT_CRF42_4_5M_Live_1dot7.mp4
1920x1080
Capped CRF
Y
Netflix VMAF VMAF061_float 1st proc 54.962486
Netflix VMAF VMAF061_float 2nd proc 89.063103
This is the middle frame of a 3-second transition. Quality confirms 43.9 VMAF score but very transient.
Guga TV
00:01:30.4•2254


E:\Capped_CRF\30_sports\Soccer\soccer_SVT_VBR_4_5M_Live_1dot7.mp4
1920x1080
VBR
Y
Netflix VMAF VMAF061_float 1st proc 54.962486
Netflix VMAF VMAF061_float 2nd proc 89.063103
VBR is clearly better, though not perfect.
Guga TV
x: 845
y: 486
Y: 133
U: 74
V: 113
00:01:30.4•2254

Low-Frame Analysis - 30fps Sports

30 fps sports	Bitrate	VMAF	Low-Frame	Bitrate	VMAF	Low Frame	Bitrate	VMAF	Low Frame
basketball.mp4	4,540	99.50	77.00	2,604	98.28	66.11	-42.6%	-1.2%	-14.1%
football.mp4	4,589	96.54	84.65	3,491	93.70	70.76	-23.9%	-2.9%	-16.4%
hockey.mp4	4,380	94.88	65.18	2,719	91.68	52.22	-37.9%	-3.4%	-19.9%
skateboard.mp4	4,033	96.95	85.68	1,313	91.22	66.90	-67.4%	-5.9%	-21.9%
soccer.mp4	3,873	97.33	78.89	3,041	96.58	54.96	-21.5%	-0.8%	-30.3%
Average	**4,283**	**97.04**	**78.28**	**2,634**	**94.29**	**62.19**	**-38.5%**	**-2.8%**	**-20.6%**

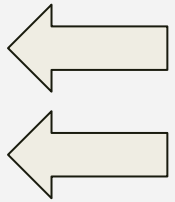

Discussion: At 30fps, the bandwidth savings were significant, and the low-frame issues were very transient and contained within generally fast-moving footage where quality deficits are hard to perceive.

Recommendation: **Use with caution**. May not be appropriate for premium content but should be OK for more general distribution.

Low- Frame Analysis - Animation

	Constant Bitrate			Capped CRF			Delta		
Animation	**Bitrate**	**VMAF**	**Low-Frame**	**Bitrate**	**VMAF**	**Low Frame**	**Bitrate**	**VMAF**	**Low Frame**
bbb.mp4	4,327	97.72	87.22	1,926	95.68	81.26	-55.5%	-2.1%	-6.8%
el_ultimo.mov	3,887	97.72	95.75	717	95.09	91.31	-81.6%	-2.7%	-4.6%
sintel.mp4	4,348	97.45	80.35	1,937	93.49	64.36	-55.5%	-4.1%	-19.9%
Average	**4,187**	**97.63**	**87.77**	**1,527**	**94.75**	**78.98**	**-63.5%**	**-2.9%**	**-10.0%**

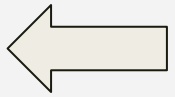

With these animated clips:

- The ***bandwidth savings were very significant*** compared to 4.5 mbps VBR (which probably is a bit too high for AV1 animated footage).
- ***Overall quality*** was a more than adequate ***94.75 VMAF***
- ***Low-frame deltas aren't significant***

Let's have a look at the issues in the Sintel clip.

Low-Frame Analysis - Sintel

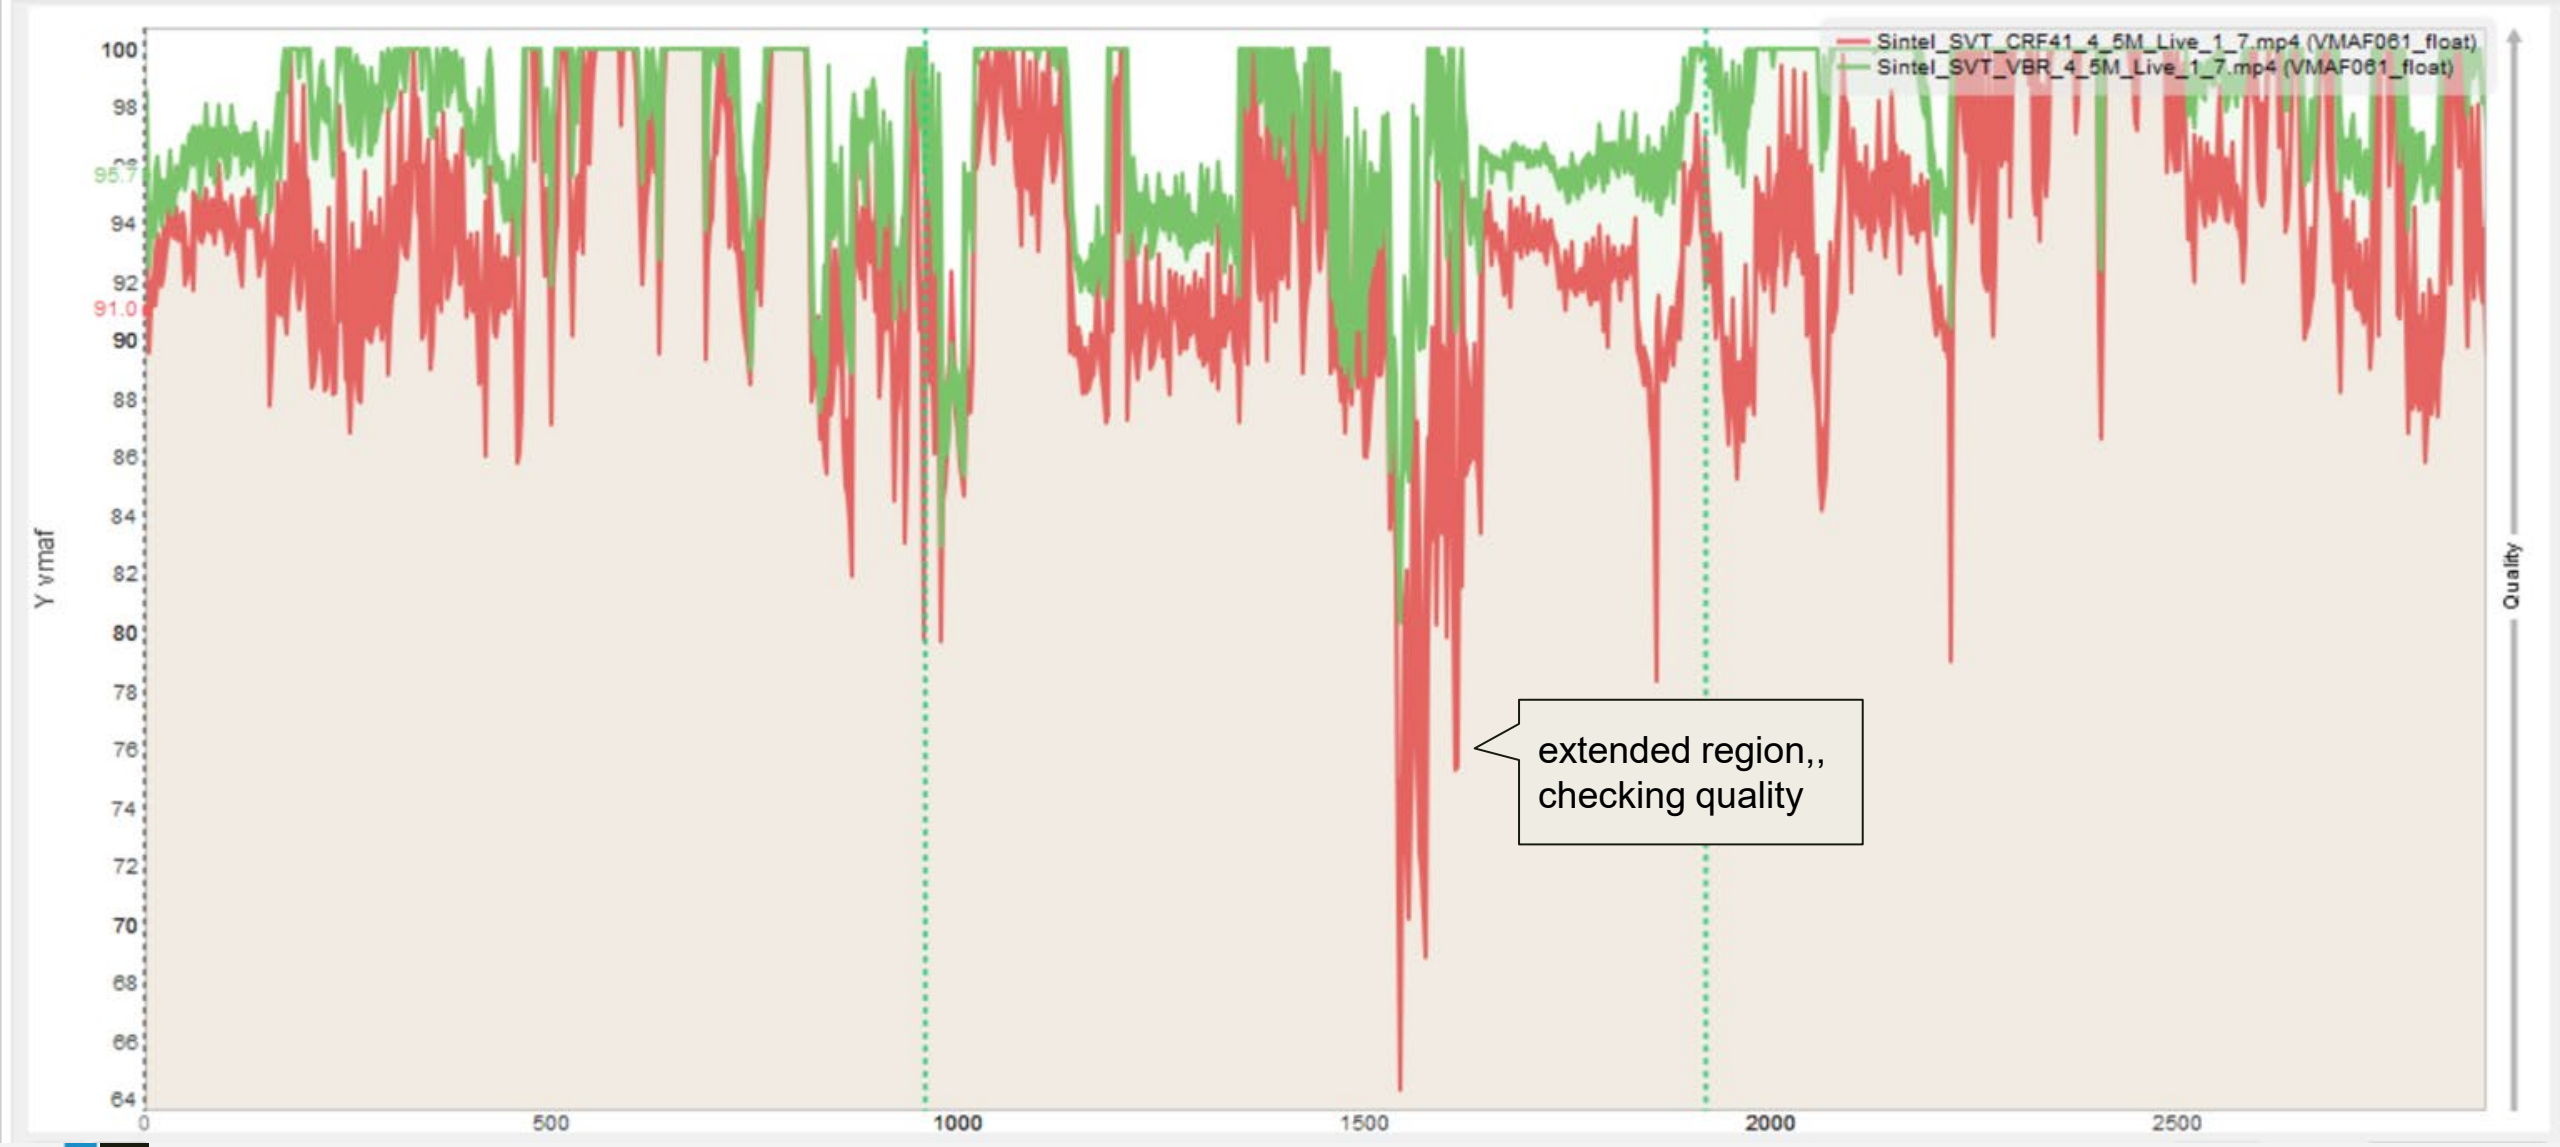

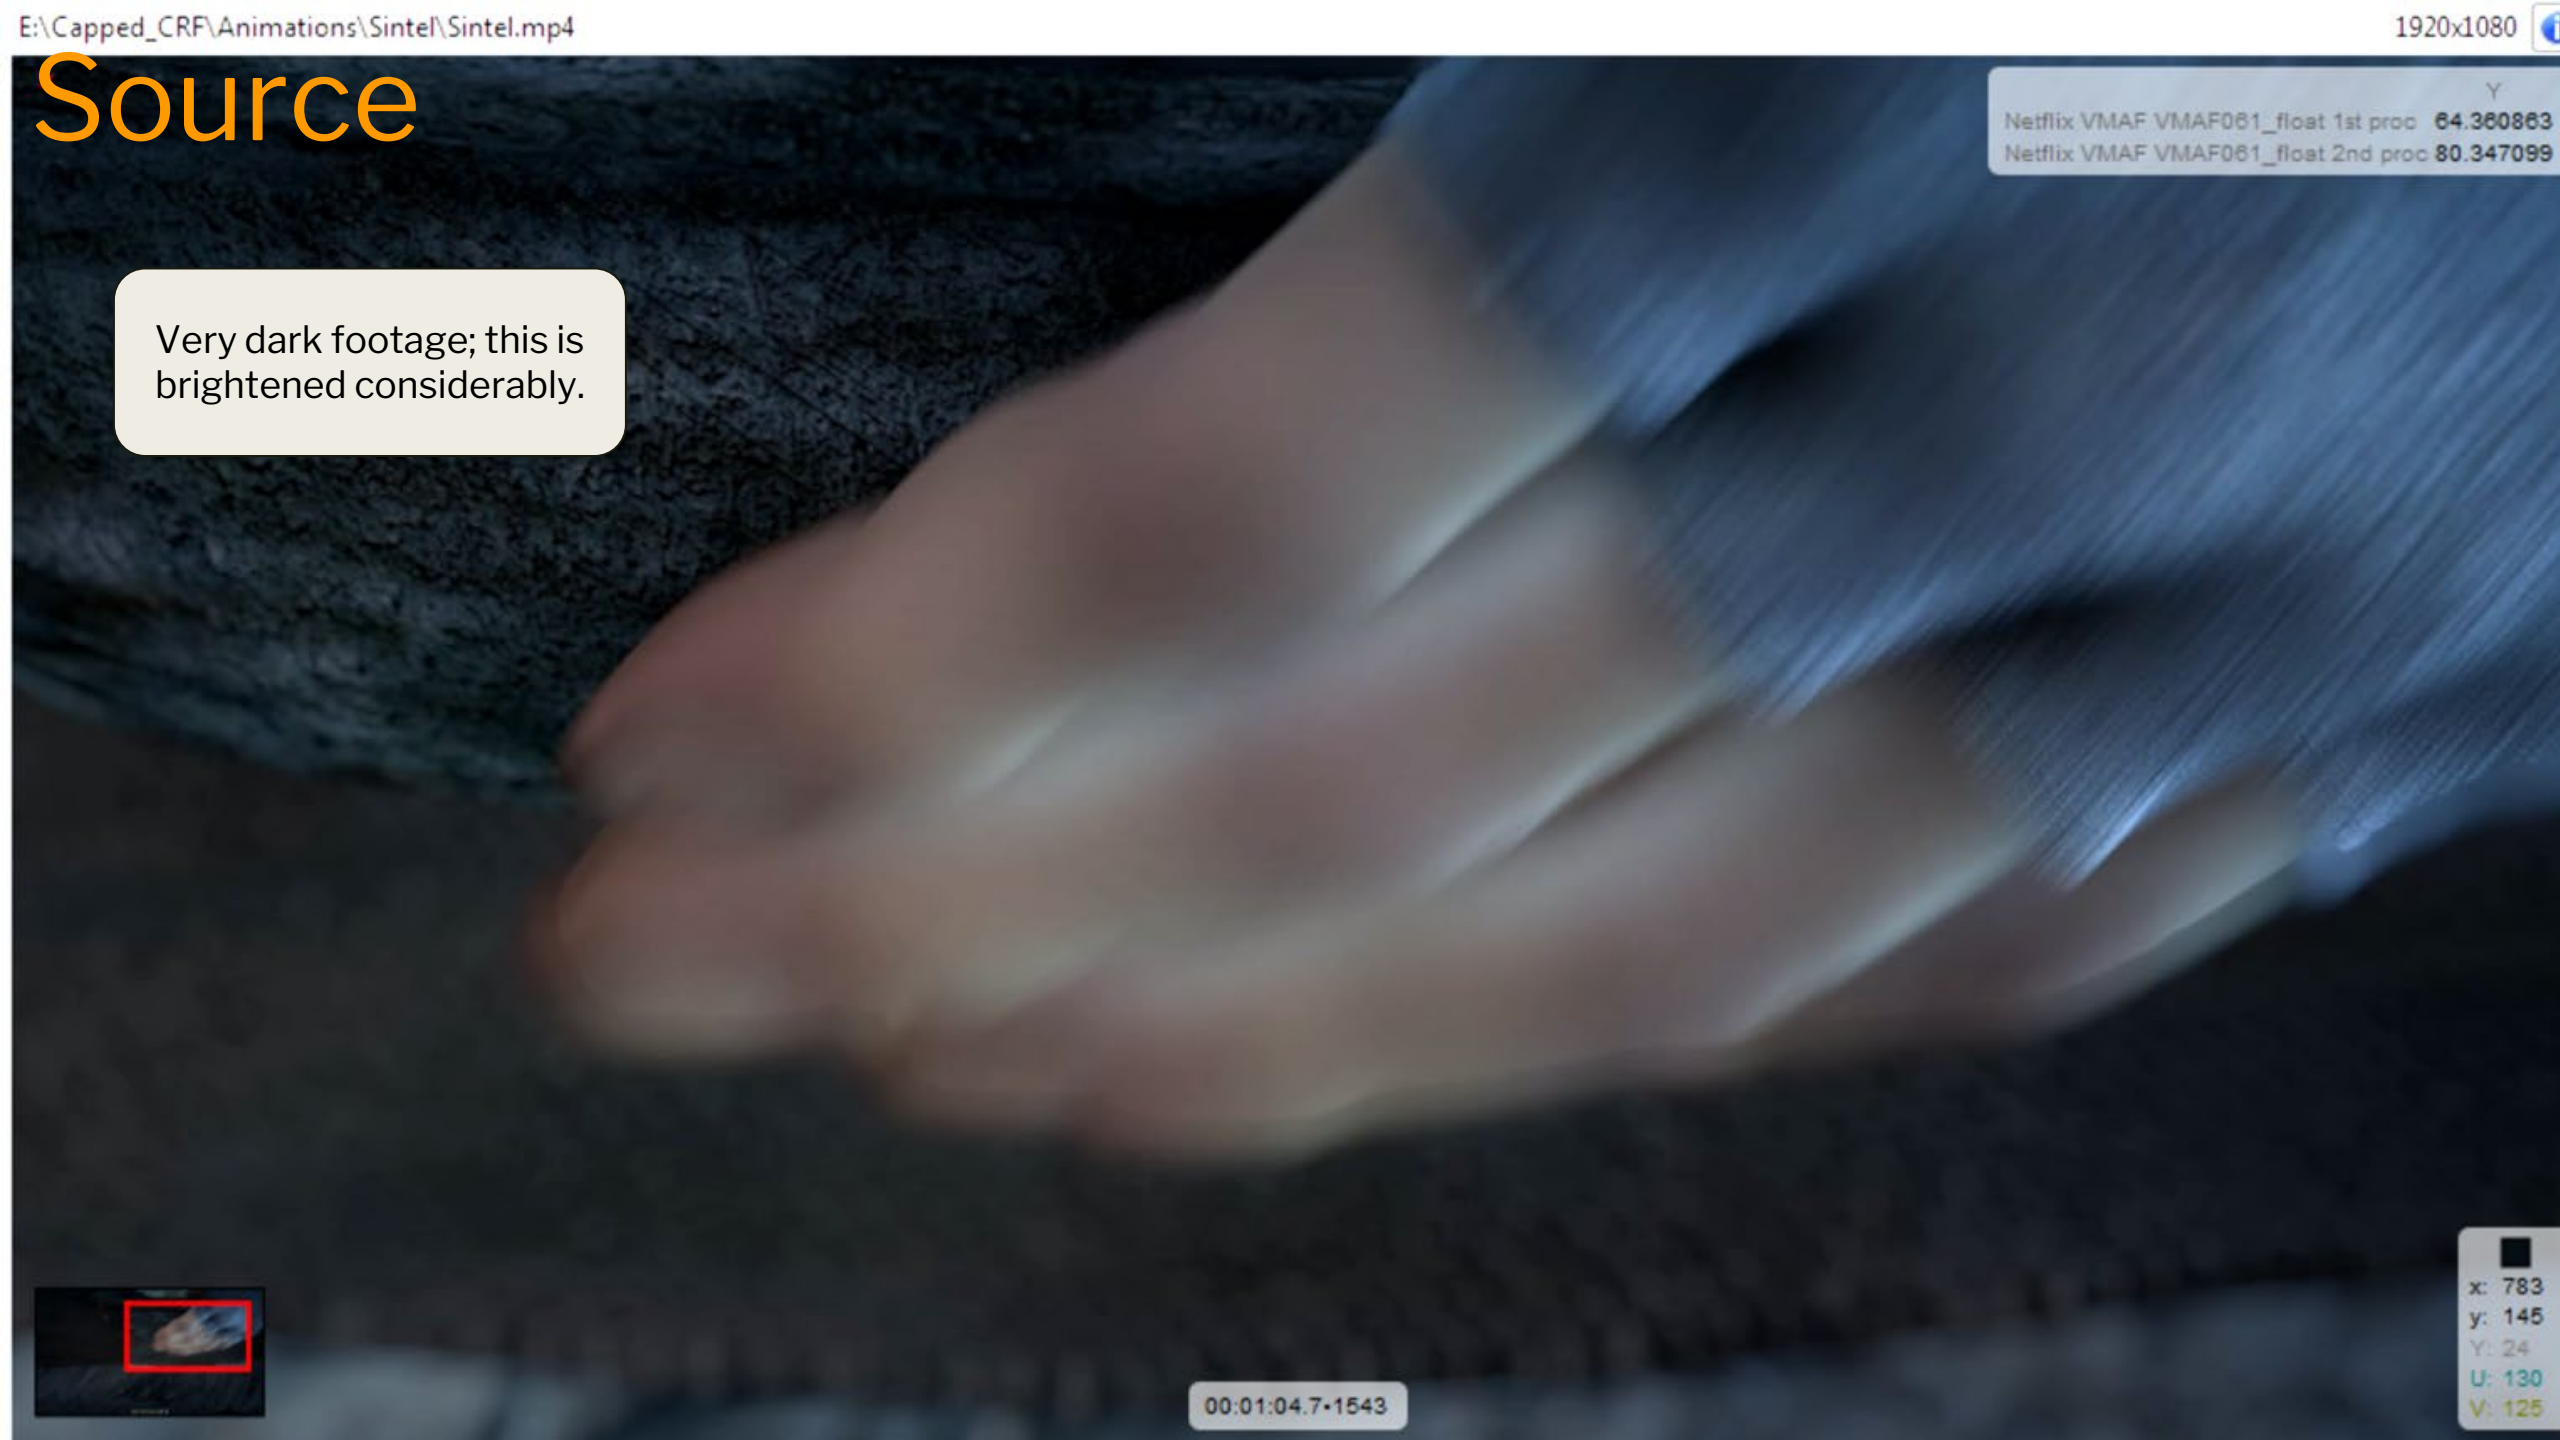
E:\Capped_CRF\Animations\Sintel\Sintel.mp4
1920x1080
Source
Y
Netflix VMAF VMAF061_float 1st proc 64.360863
Netflix VMAF VMAF061_float 2nd proc 80.347099
Very dark footage; this is brightened considerably.
00:01:04.7•1543
x: 783
y: 145
Y: 24
U: 130
V: 125

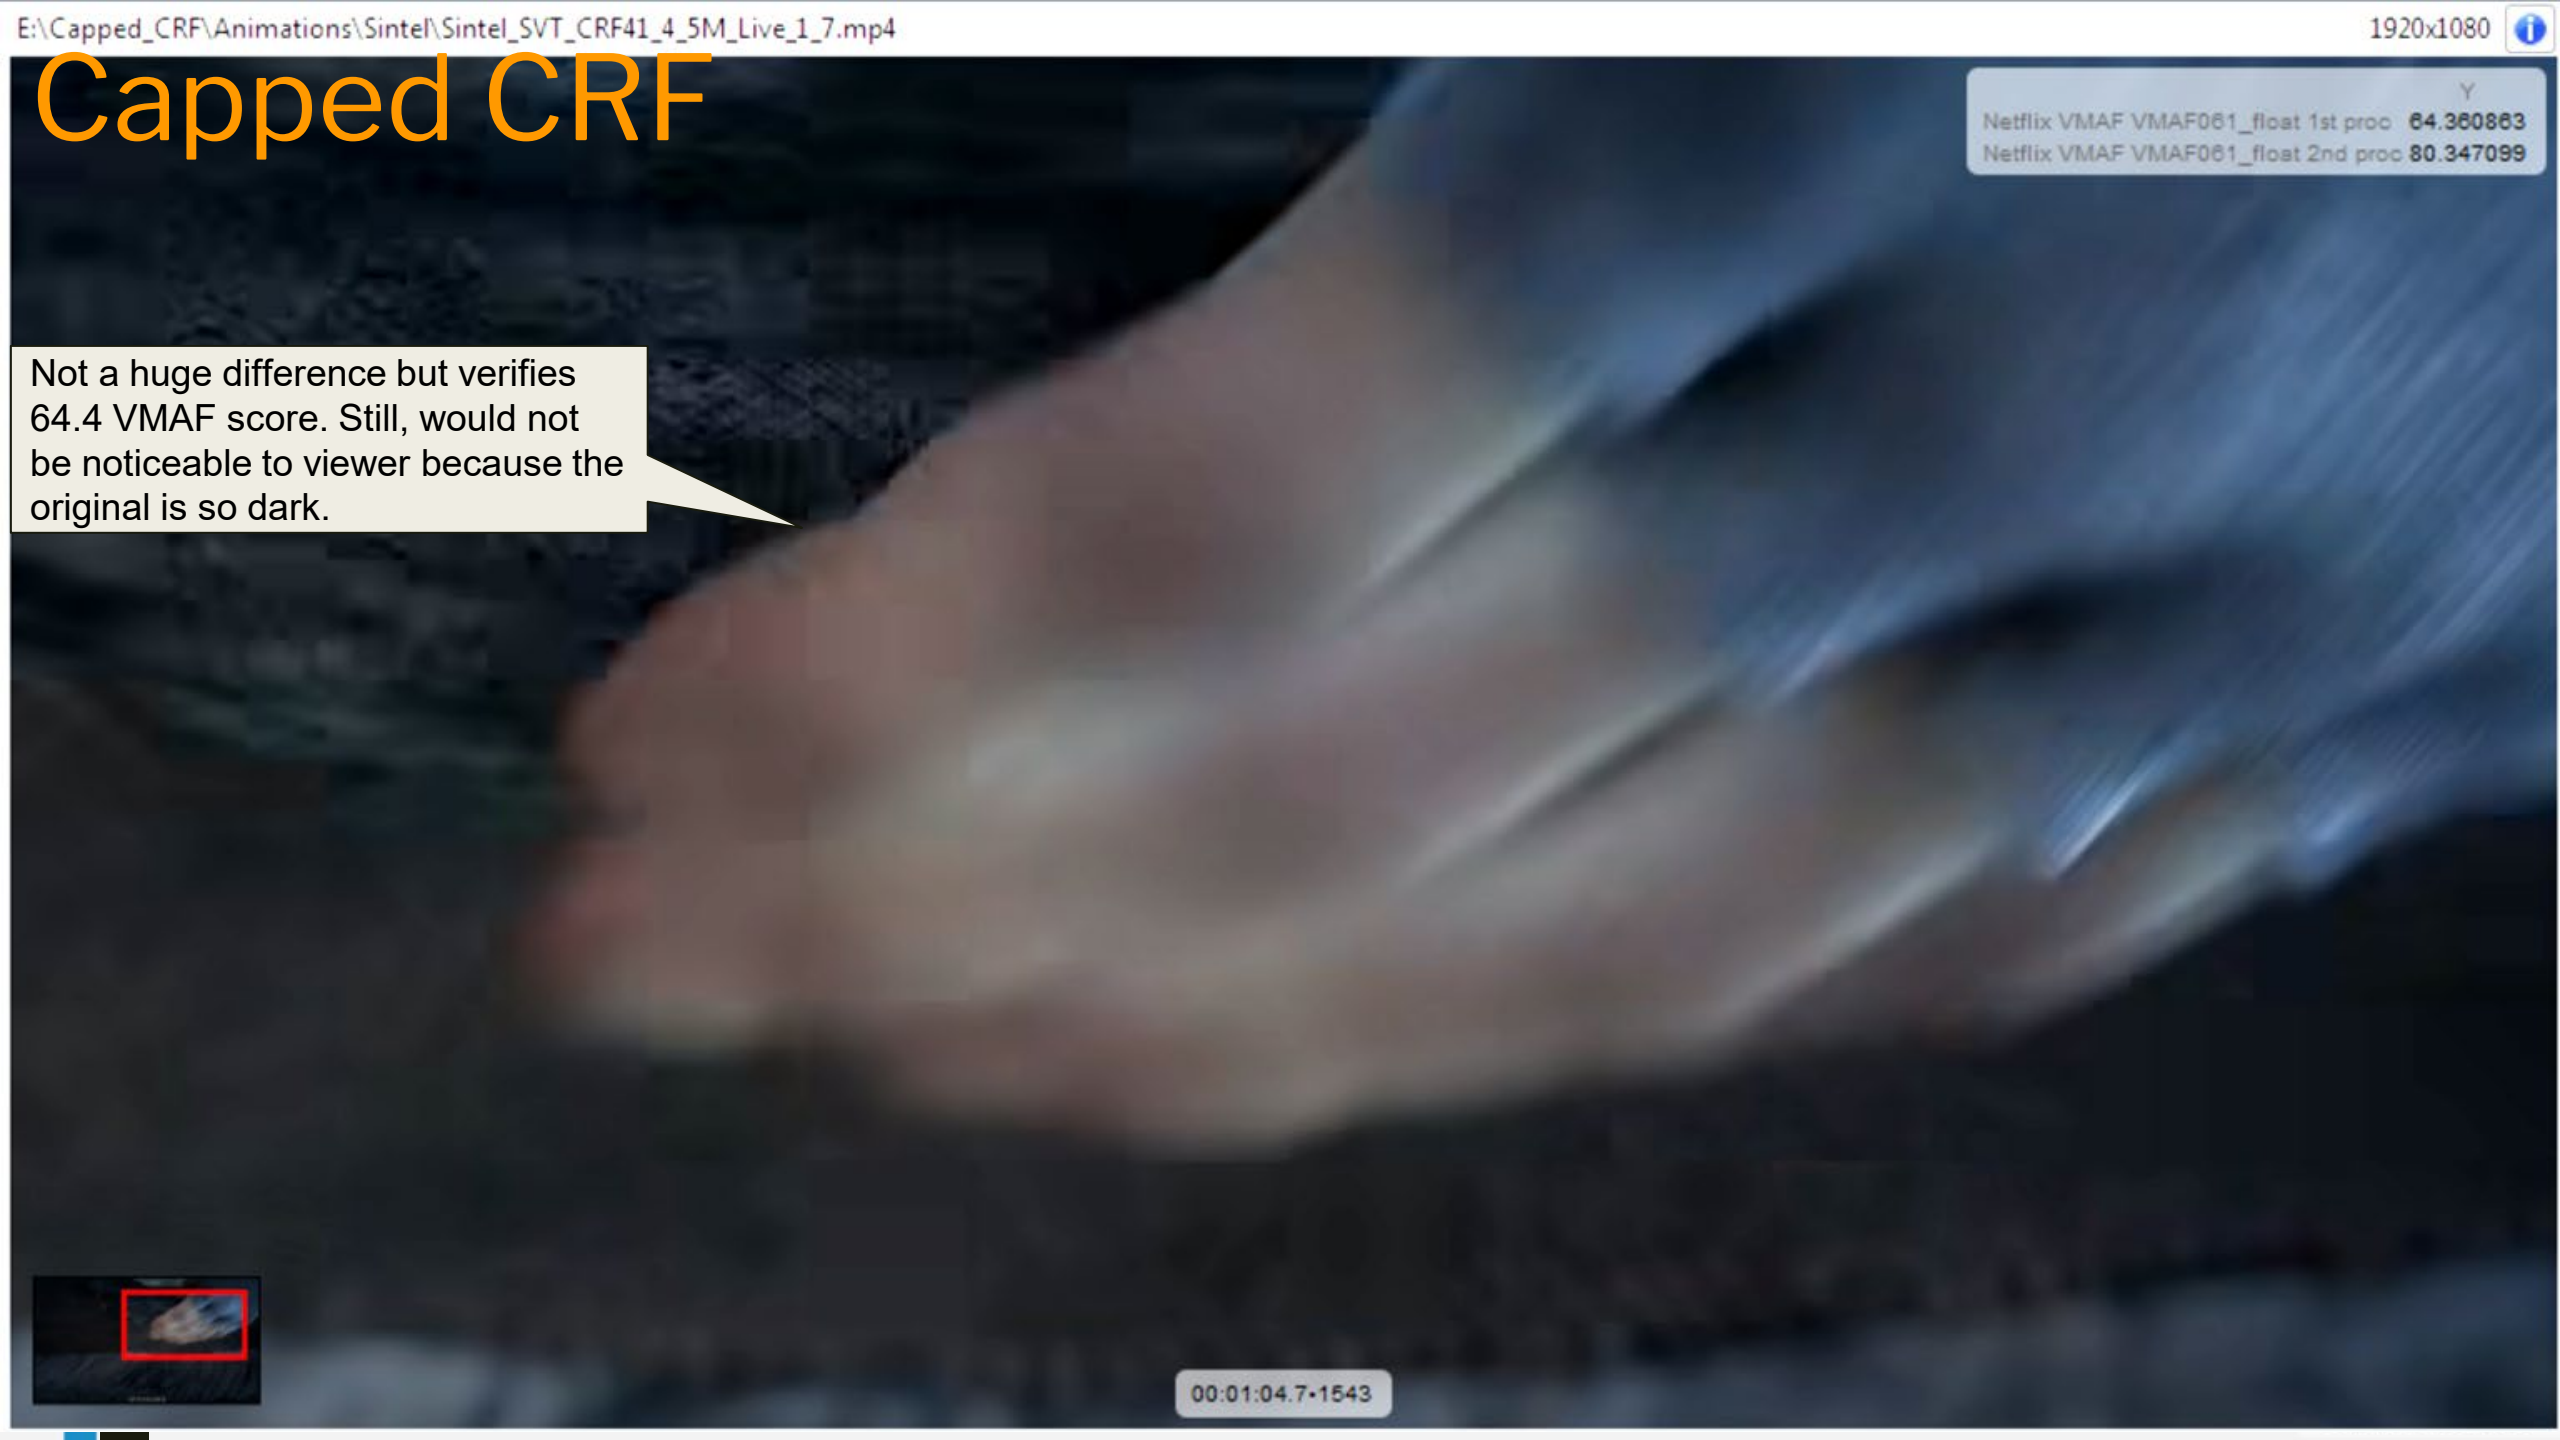
E:\Capped_CRF\Animations\Sintel\Sintel_SVT_CRF41_4_5M_Live_1_7.mp4
1920x1080
Capped CRF
Y
Netflix VMAF VMAF061_float 1st proc 64.360863
Netflix VMAF VMAF061_float 2nd proc 80.347099
Not a huge difference but verifies 64.4 VMAF score. Still, would not be noticeable to viewer because the original is so dark.
00:01:04.7•1543

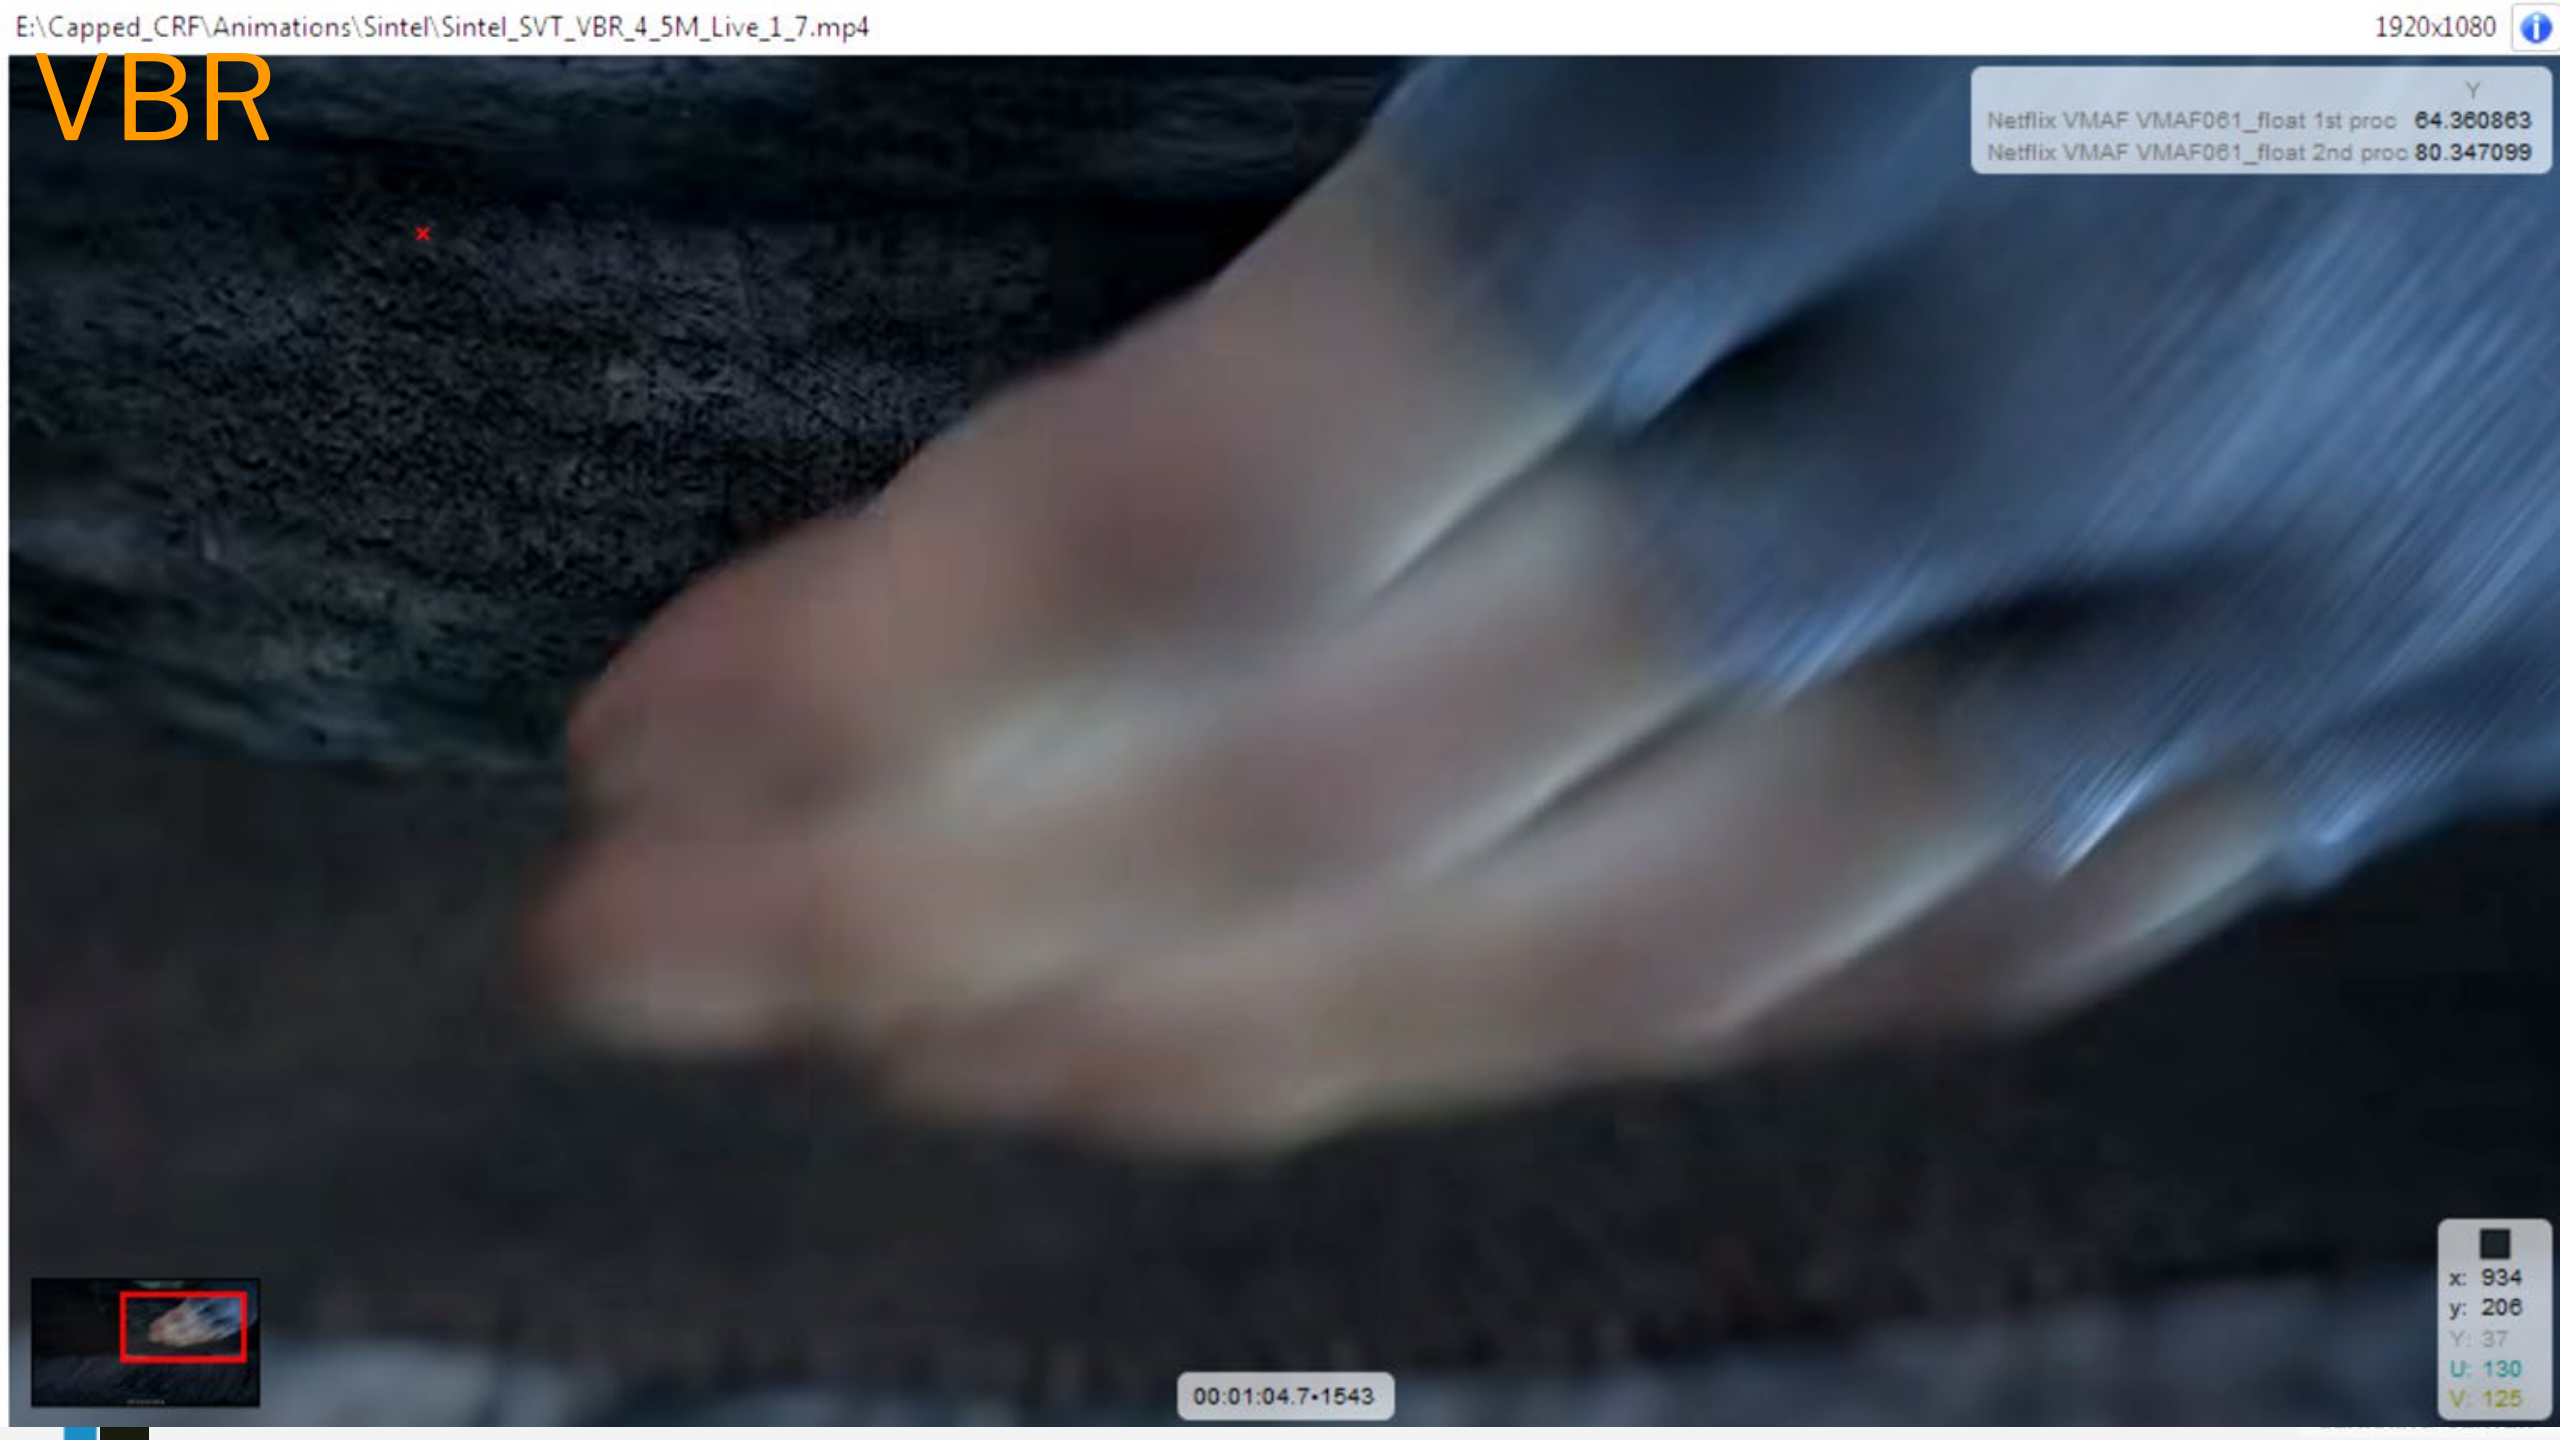
E:\Capped_CRF\Animations\Sintel\Sintel_SVT_VBR_4_5M_Live_1_7.mp4
1920x1080
VBR
Y
Netflix VMAF VMAF061_float 1st proc 64.360863
Netflix VMAF VMAF061_float 2nd proc 80.347099
00:01:04.7•1543
x: 934
y: 206
Y: 37
U: 130
V: 125

Low- Frame Analysis - Animation

	Constant Bitrate			Capped CRF			Delta		
Animation	**Bitrate**	**VMAF**	**Low-Frame**	**Bitrate**	**VMAF**	**Low Frame**	**Bitrate**	**VMAF**	**Low Frame**
bbb.mp4	4,327	97.72	87.22	1,926	95.68	81.26	-55.5%	-2.1%	-6.8%
el_ultimo.mov	3,887	97.72	95.75	717	95.09	91.31	-81.6%	-2.7%	-4.6%
sintel.mp4	4,348	97.45	80.35	1,937	93.49	64.36	-55.5%	-4.1%	-19.9%
Average	**4,187**	**97.63**	**87.77**	**1,527**	**94.75**	**78.98**	**-63.5%**	**-2.9%**	**-10.0%**

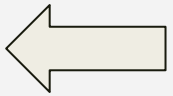

Discussion: Capped CRF delivered significant bitrate savings, minimal average quality decreases, and nothing scary in the low-frame department.

Savings would obviously be less if lower cap/VBR bitrate was applied. Still, for live streaming of animated footage, capped CRF is definitely worth testing.

Recommendation: **Deploy but verify**. Seems low risk but check with your own footage.

Low- Frame Analysis - Office

	Constant Bitrate			Capped CRF			Delta		
Office	Bitrate	VMAF	Low-Frame	Bitrate	VMAF	Low Frame	Bitrate	VMAF	Low Frame
epiphan.mp4	3,572	97.10	90.62	520	94.00	78.52	-85.4%	-3.2%	-13.4%
talkinghead.mp4	3,862	97.22	95.13	646	93.65	89.64	-83.3%	-3.7%	-5.8%
test.mp4	4,111	97.23	90.89	1,317	94.59	87.94	-68.0%	-2.7%	-3.2%
tutorial.mp4	4,619	97.44	95.28	757	97.00	94.69	-83.6%	-0.5%	-0.6%
Average	**4,041**	**97.25**	**92.98**	**810**	**94.81**	**87.70**	**-80.0%**	**-2.5%**	**-5.7%**

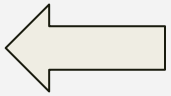

Discussion: Capped CRF delivered significant:

- Exceptional bandwidth savings
- Very good average quality at selected CRF value
- Very similar low-frame quality to VBR.

Let's check the issues in the Epiphan clip. This is a mixed source clip with:

- Real world video of talking heads and shots of an Epiphan encoder
- Some PowerPoint slides
- Some screencam

Low-Frame Analysis - Epiphan

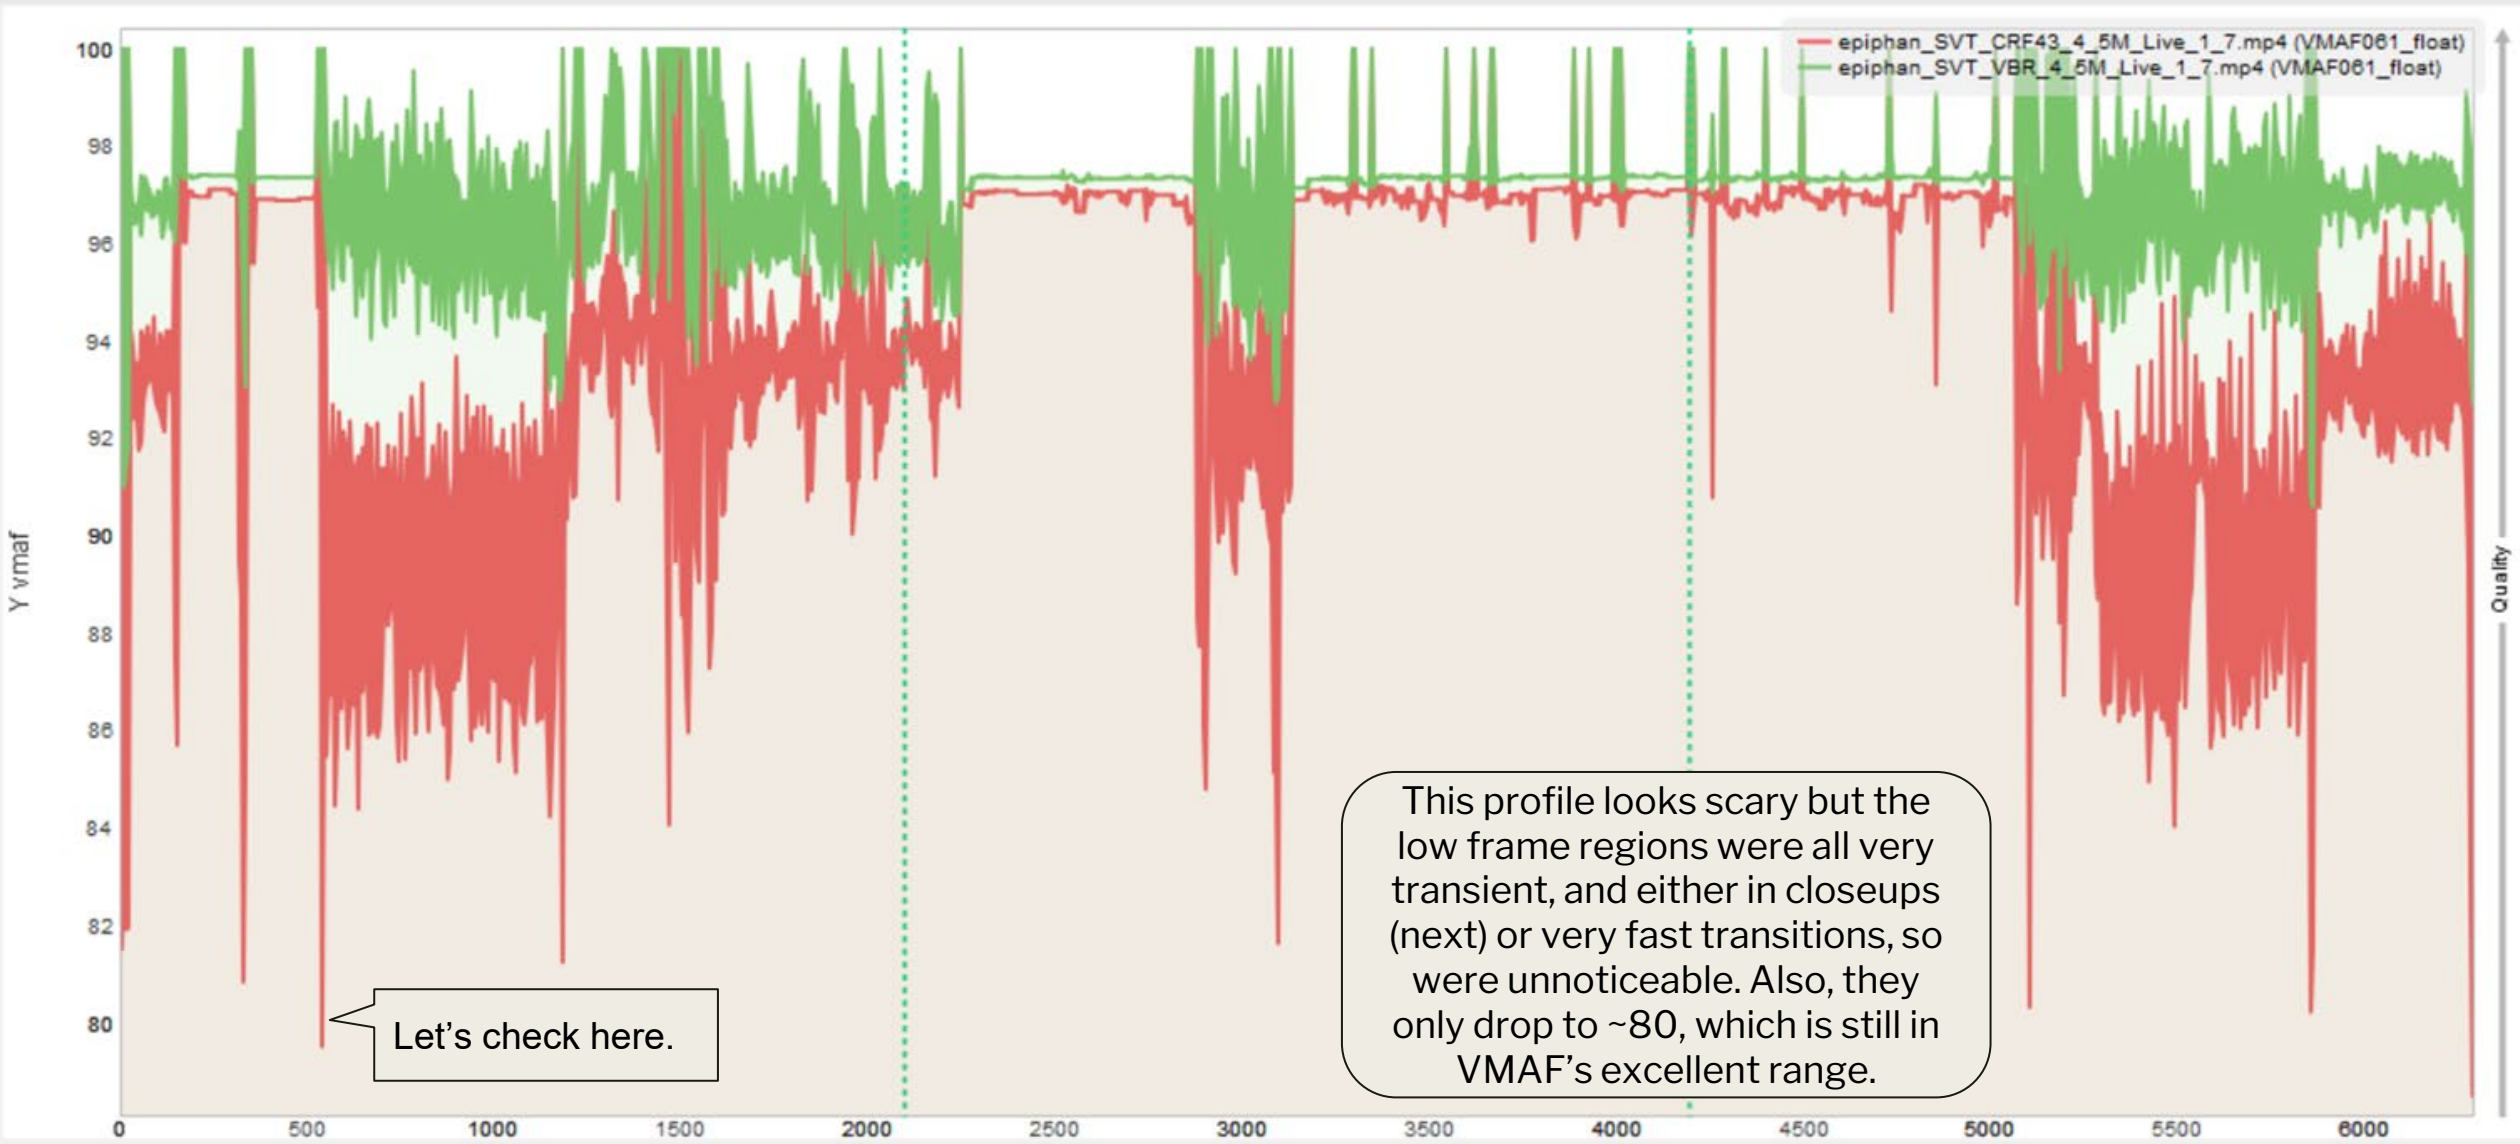

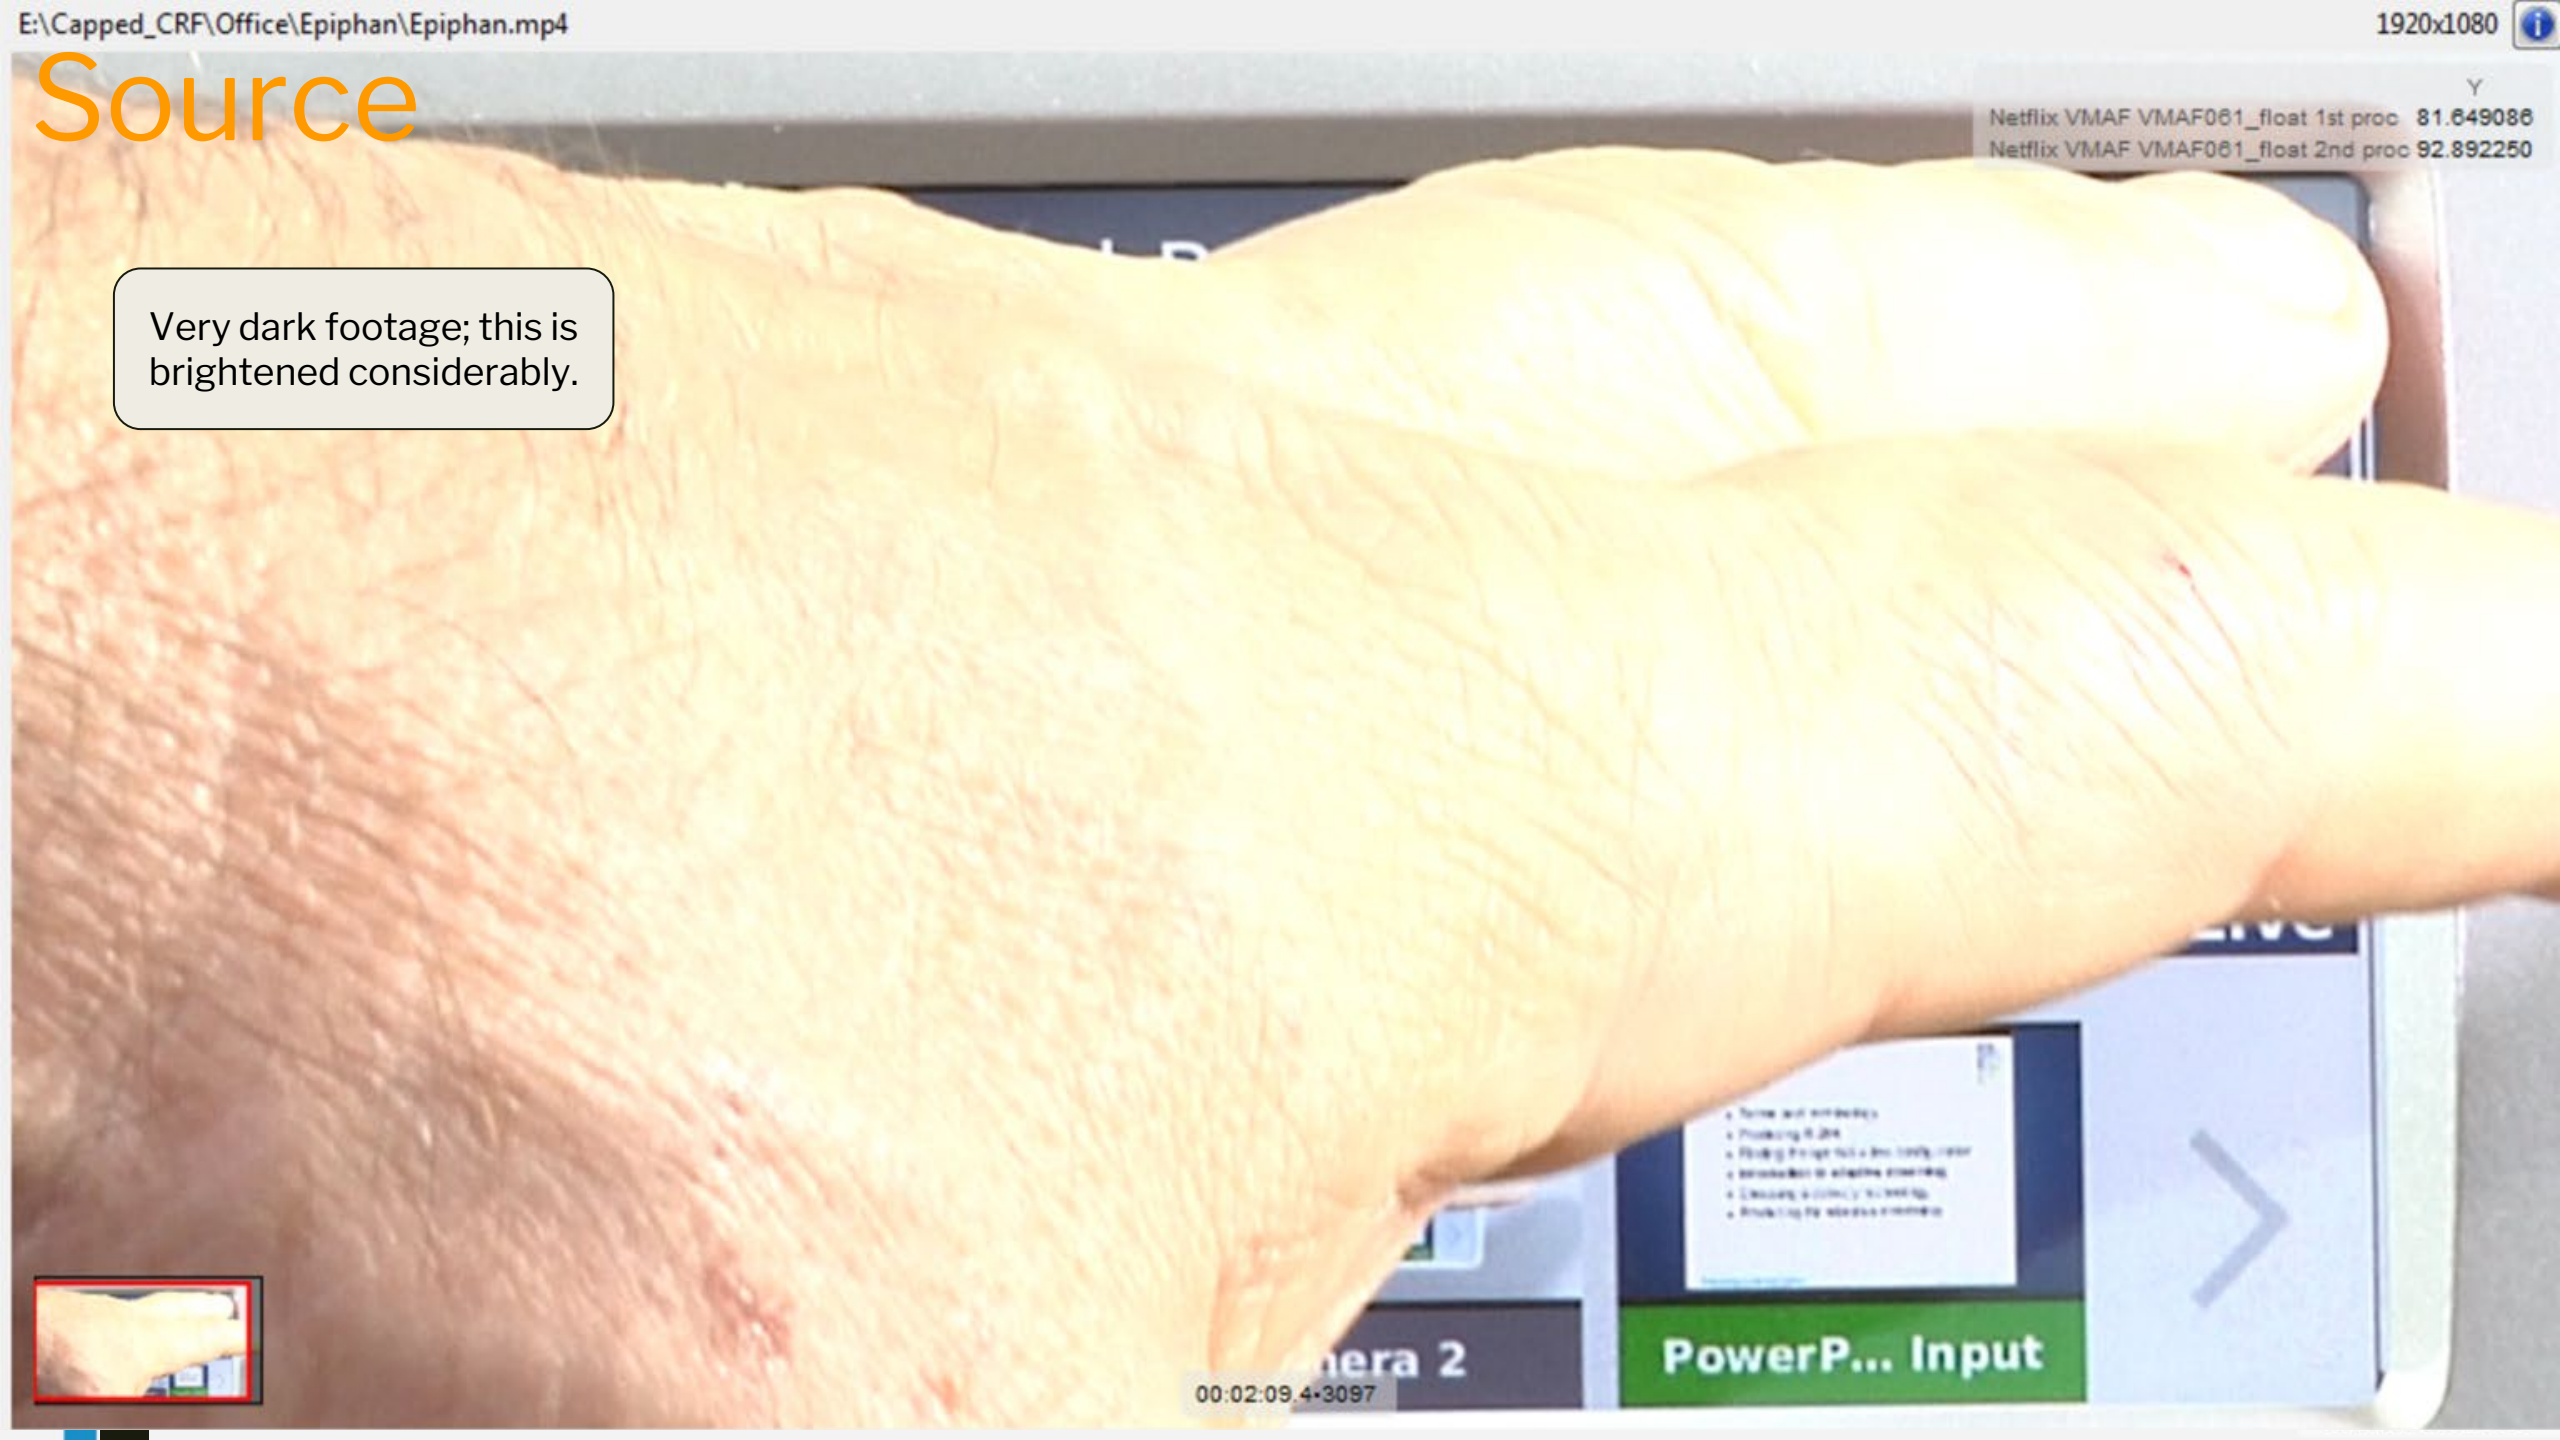
E:\Capped_CRF\Office\Epiphan\Epiphan.mp4
1920x1080
Source
Y
Netflix VMAF VMAF061_float 1st proc 81.649086
Netflix VMAF VMAF061_float 2nd proc 92.892250
Very dark footage; this is brightened considerably.
PowerP... Input
00:02:09.4+3097

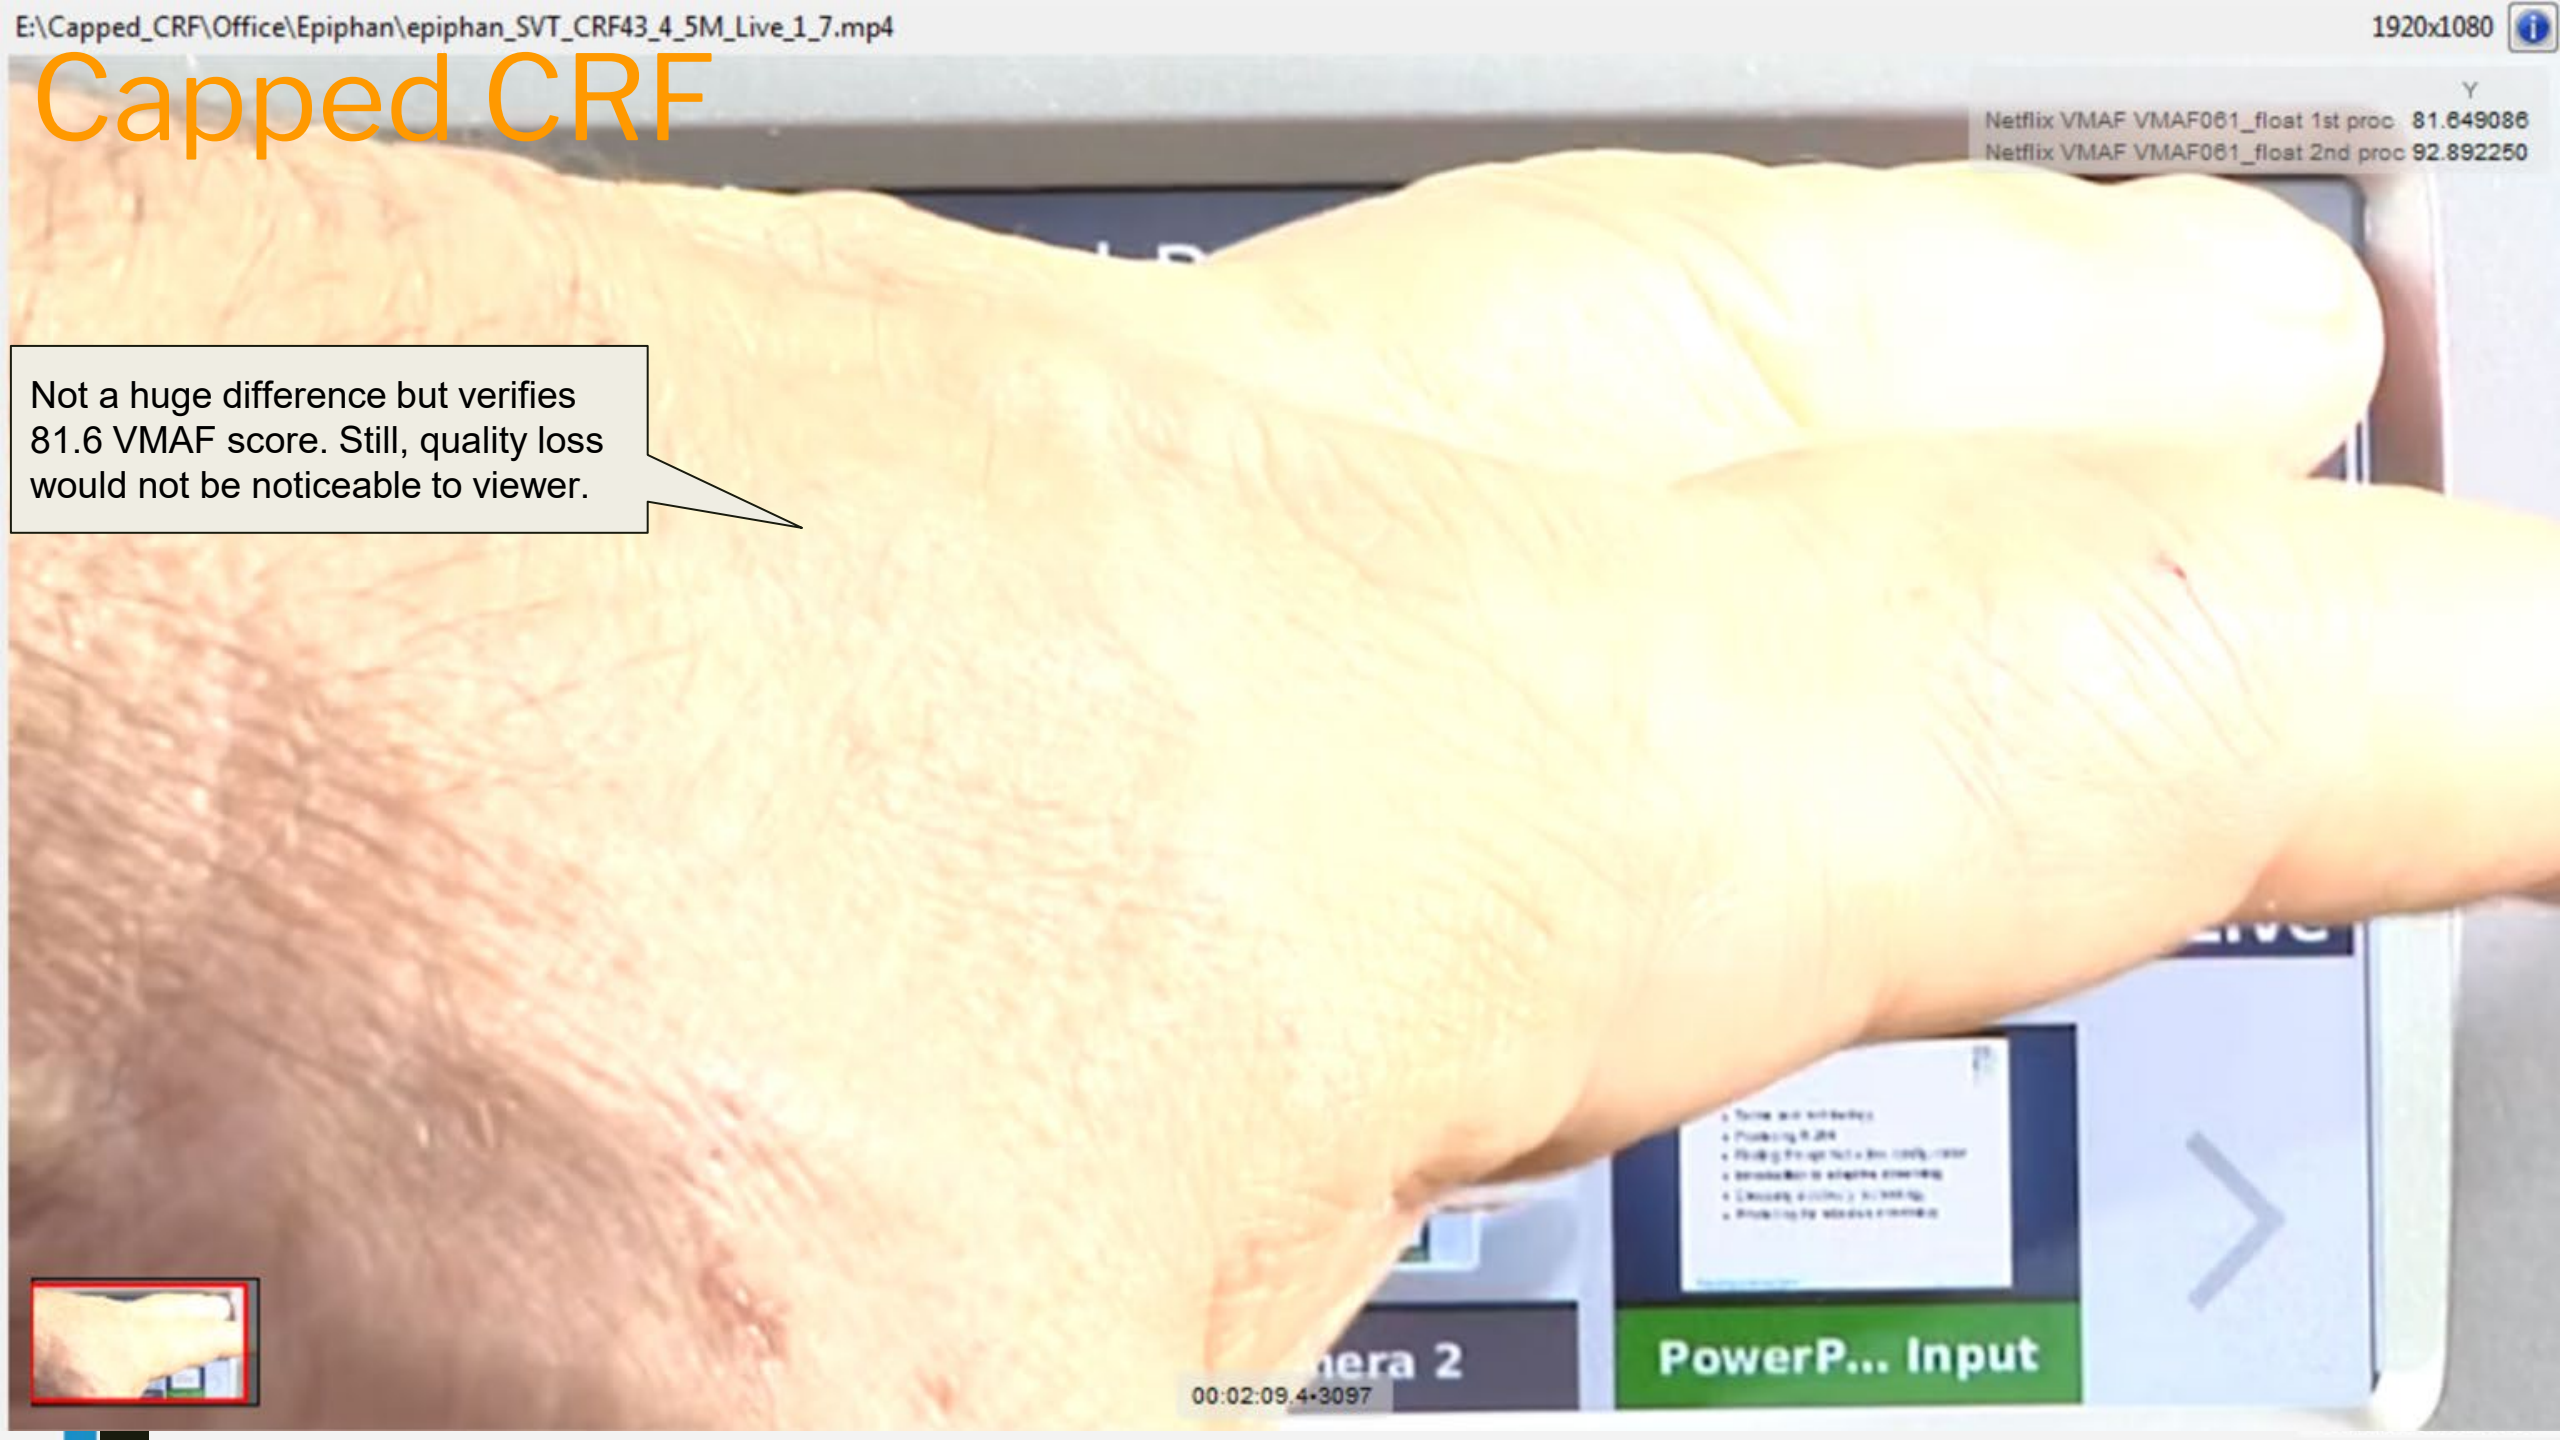
E:\Capped_CRF\Office\Epiphan\epiphan_SVT_CRF43_4_5M_Live_1_7.mp4
1920x1080
Capped CRF
Y
Netflix VMAF VMAF061_float 1st proc 81.649086
Netflix VMAF VMAF061_float 2nd proc 92.892250
Not a huge difference but verifies 81.6 VMAF score. Still, quality loss would not be noticeable to viewer.
PowerP... Input
00:02:09.4•3097

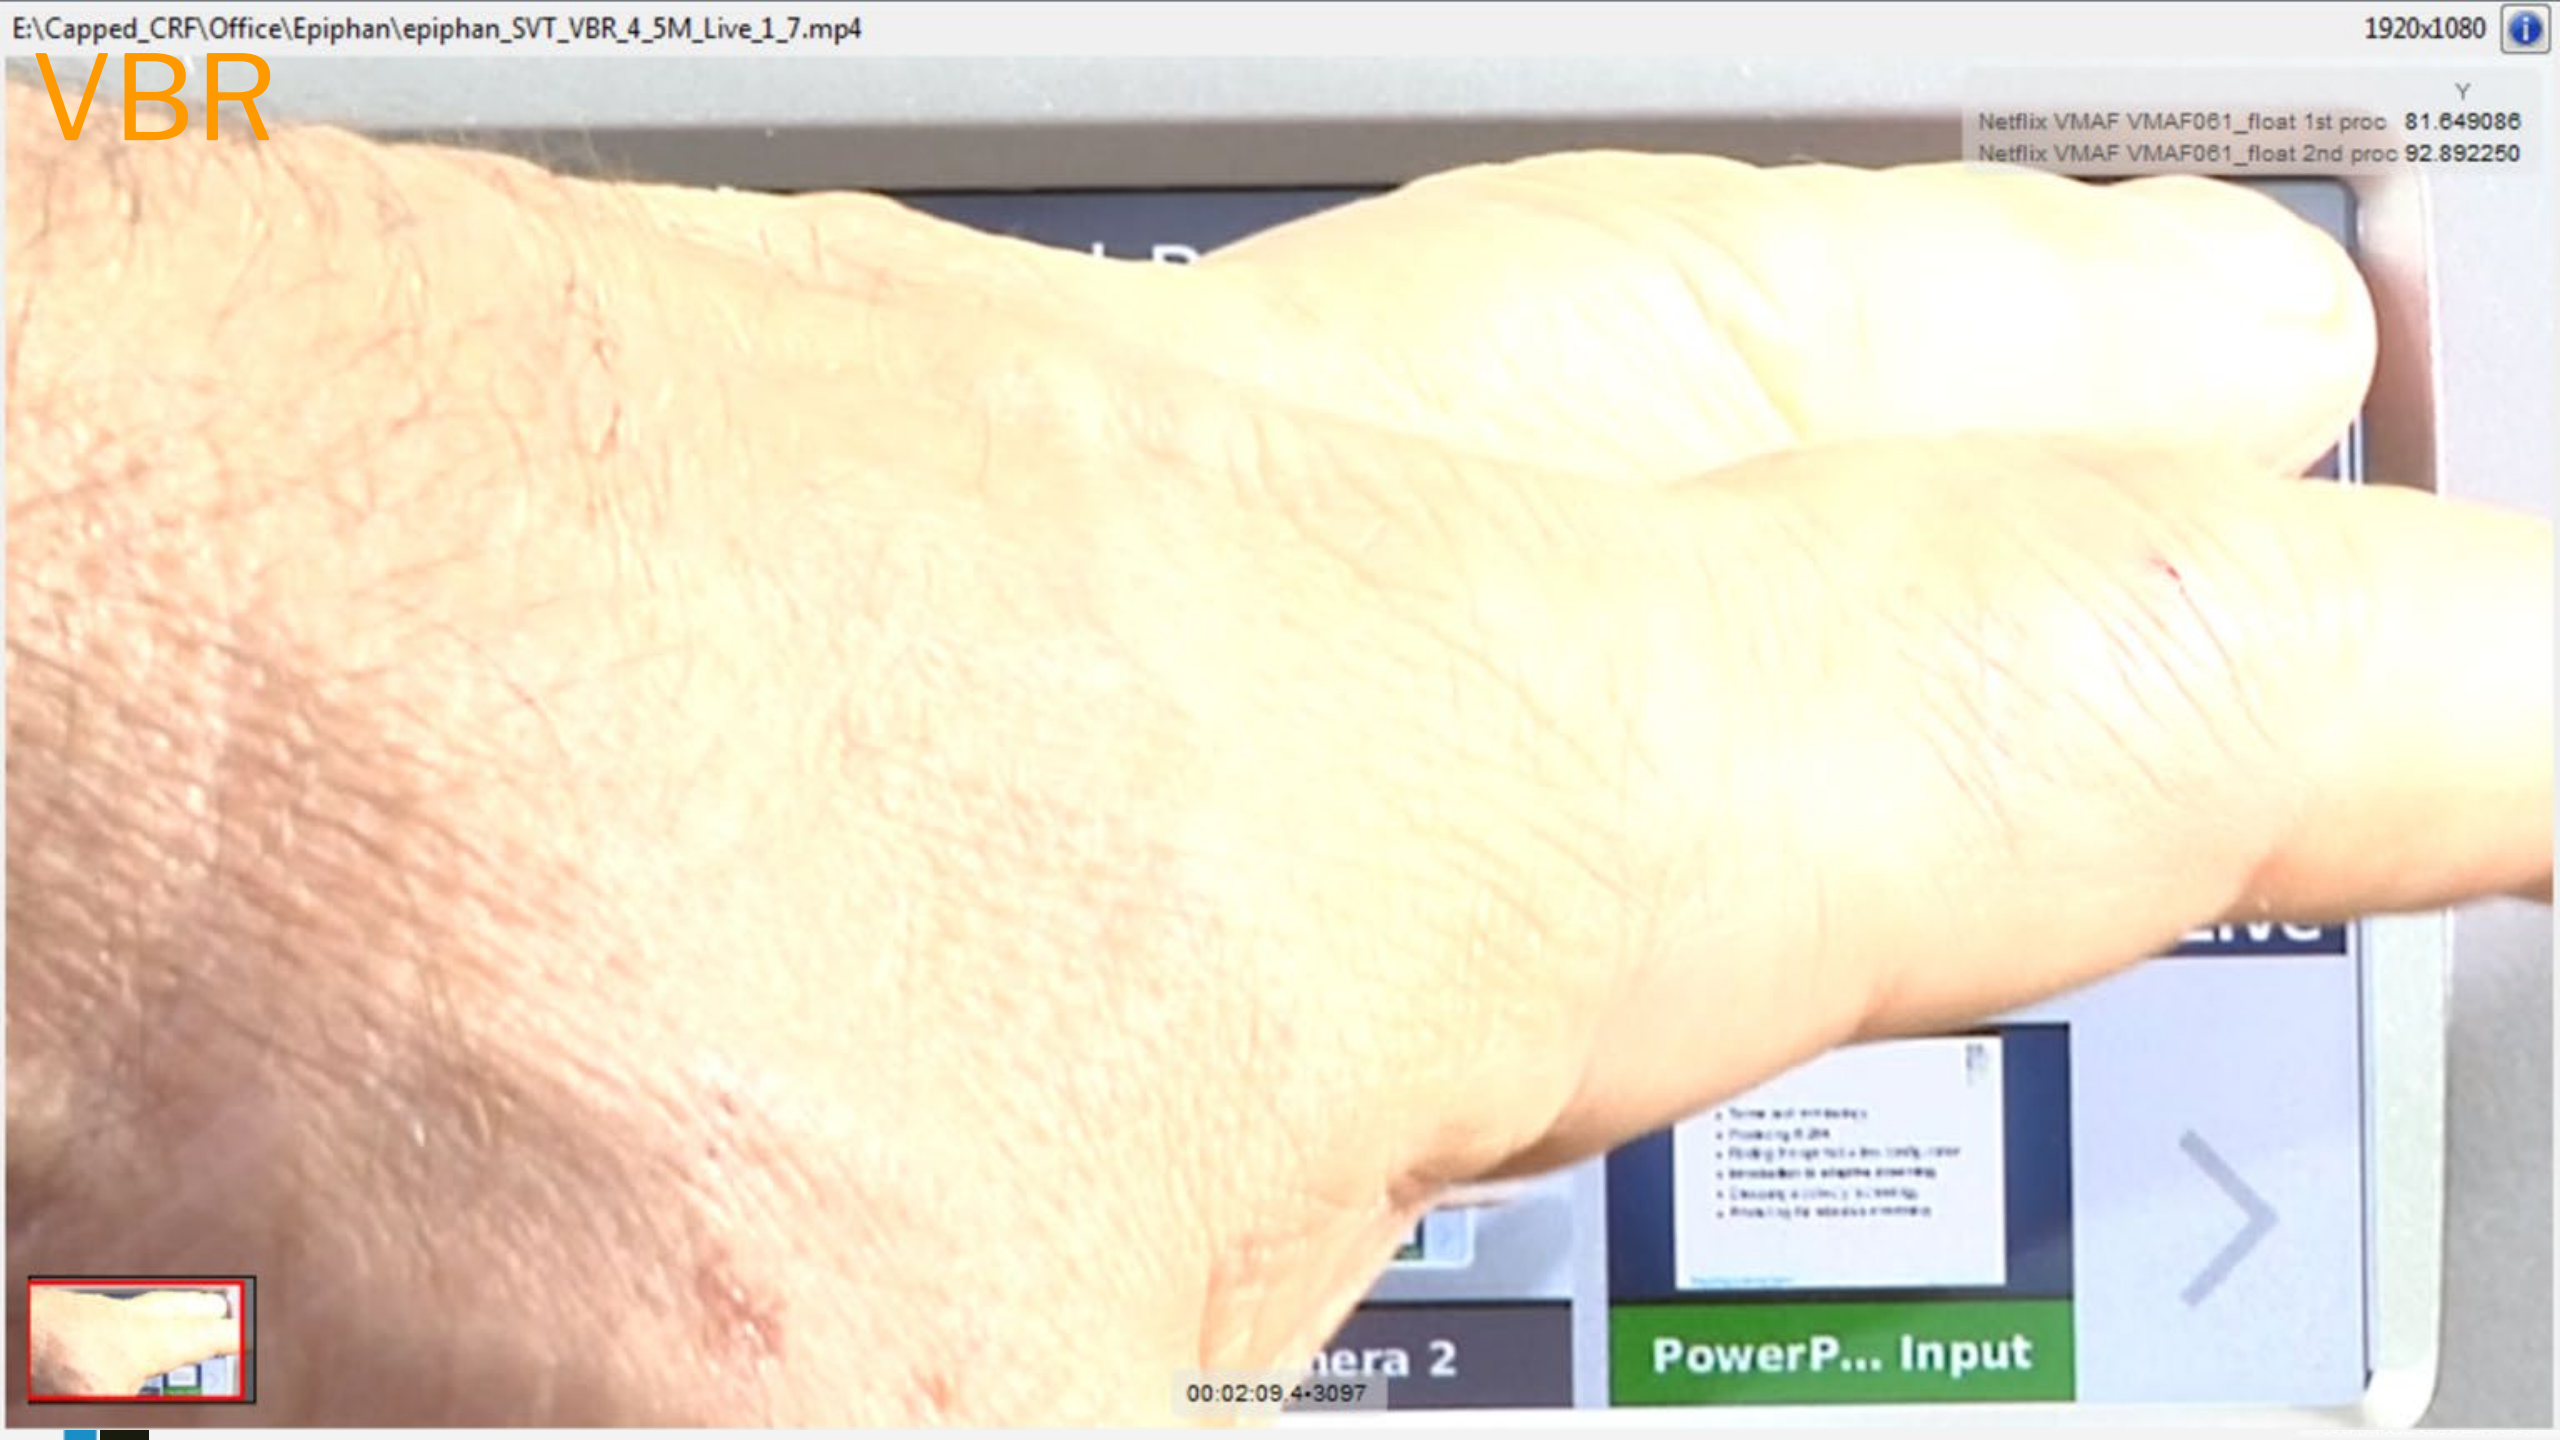
E:\Capped_CRF\Office\Epiphan\epiphan_SVT_VBR_4_5M_Live_1_7.mp4
1920x1080
VBR
Y
Netflix VMAF VMAF061_float 1st proc 81.649086
Netflix VMAF VMAF061_float 2nd proc 92.892250
PowerP... Input
00:02:09.4•3097

Low- Frame Analysis - Office

	Constant Bitrate			Capped CRF			Delta		
Office	**Bitrate**	**VMAF**	**Low-Frame**	**Bitrate**	**VMAF**	**Low Frame**	**Bitrate**	**VMAF**	**Low Frame**
epiphan.mp4	3,572	97.10	90.62	520	94.00	78.52	-85.4%	-3.2%	-13.4%
talkinghead.mp4	3,862	97.22	95.13	646	93.65	89.64	-83.3%	-3.7%	-5.8%
test.mp4	4,111	97.23	90.89	1,317	94.59	87.94	-68.0%	-2.7%	-3.2%
tutorial.mp4	4,619	97.44	95.28	757	97.00	94.69	-83.6%	-0.5%	-0.6%
Average	**4,041**	**97.25**	**92.98**	**810**	**94.81**	**87.70**	**-80.0%**	**-2.5%**	**-5.7%**

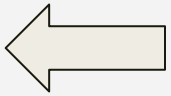

Discussion: Low frame issues don't appear particularly scary.

Recommendation: **Deploy but verify**. Seems low risk but check with your own footage.

Overall Findings

	CRF Value	Bitrate Savings	Average Quality Drop	Low-Frame Decrease	Recommendation
Media	33	-35%	-1.4%	-5.9%	**Deploy but verify**
60 fps sports	36	-18%	-1.6%	-29.2%	**Use with caution**
30 fps sports	41	-39%	-2.8%	-20.6%	**Use with caution**
Animation	42	-64%	-2.9%	-10.0%	**Deploy but verify**
Office	43	-80%	-2.5%	-5.7%	**Deploy but verify**
Average		**-44%**	**-2.1%**	**-13.5%**	

Overall, as compared to VBR, capped CRF:

- Should produce substantial bandwidth savings (overall 44%)
- Should, in all cases, produce average quality scores in the relevant range
- May increase the likelihood of transient quality problems in sports footage
- Should deliver greater throughput than VBR

- High-volume producers (or those transcoding) should consider a hardware solution because CPU requirements will vary significantly for SVT-AV1 based upon the complexity of the source footage

Procedure for Deploying Capped CRF

1. Create separate test files for each genre. Find between 3-6 files between 1-2 minutes long
2. Encode at various CRF values without a cap to identify value that delivers a VMAF score from 93-95
3. Once you have that CRF value, encode test files with a cap and check overall and low-frame quality.
    a. *Choose a cap that equals the bitrate you would use for VBR transcoding*
4. Apply to production files and reassess quality and bitrate savings