# Video Streaming with the AV1 Video Codec in Mobile Devices

Joint white paper by Meta, Vodafone and YouTube

September 2025

# Contents

# Executive Summary

Video traffic on mobile networks is increasingly popular, and continuing to ensure efficiency of video delivery is important for both mobile network operators and content providers. Video codecs – the technology used to compress (encode) and decompress (decode) videos – play an important role in achieving better video efficiency. The AV1 video codec offers up to 30% better compression compared to previous generation codecs. Popular content providers (Meta, Google/YouTube, Netflix) already support AV1 streaming on their platforms by encoding videos in AV1.

AV1 decoding can be implemented in smartphones in both hardware and software. Compared to hardware implementation, software implementation has a modest overhead on CPU utilisation and battery life. AV1 hardware decode is supported on most modern, high end mobile SoCs (System-on-a-Chip – the main processor in a smartphone), but support in mid and low tier SoCs is limited. AV1 software decode provides a way for lower end and many older devices to decode AV1 streams where hardware support may not be available.

By making hardware-based AV1 decode available across all smartphone platform tiers on a timely basis, device SoCs would enable AV1 support across the full smartphone range. Support in mid and low tier chipsets will result in far greater support of AV1 hardware decode, increasing the benefits to mobile network operators, content providers and end users. We believe that, in the long term, the greatest benefits can be achieved by support for AV1 hardware decode.

For software decode, we recommend implementing a libdav1d based software approach. OEMs could investigate platform-specific optimisation for better performance. We also recommend development of an industry-wide model for benchmarking of SW video decoders to enable content providers to easily determine the SW decoder performance capabilities of a given device.

## Recommendations

**Device OEMs and SoC providers**

- SoC vendors should consider supporting AV1 in hardware, which allows for the highest resolution video and supports playing DRM-protected movies and TV shows
- In scenarios where HW AV1 is not an option, SOC vendors and device manufacturers should consider adopting a platform optimized SW AV1 decoder
- A SW AV1 decoder can help with the transition to AV1 in low/mid-tier devices
- A shift in low-end devices from A53 ARM Cores to A55 ARM cores will improve battery life and device performance for SW video decoding.

**Authors and Contributors**

| | **Name** | **Company** |
|---|---|---|
| **1** | José Almeida | **Vodafone** |
| **2** | Renuka Bhalerao | **Meta** |
| **3** | Carena Church | **Google** |
| **4** | Ben Hulse | **Google** |
| **5** | Des Keane | **Google** |
| **6** | David Ronca | **Ex Meta** |
| **7** | Neil Scully | **Vodafone** |
| **8** | Michael Smouha | **Google** |
| **9** | Hassene Tmar | **Meta** |
| **10** | Peter West | **Ex Vodafone** |

# 1. Video Traffic on Mobile Networks

In the past decade, mobile technologies have undergone significant advancements, becoming increasingly faster, more accessible, and versatile in their applications. As a result, smartphones have seamlessly integrated into our daily lives, serving as an extension of both work and entertainment. This shift has subsequently led to a substantial surge in demand for robust and reliable mobile networks.

Smartphones are used for a variety of services, including web-browsing, social media, e-mail, messaging etc. However, the largest contributor to data traffic is video. The percentage of data traffic that is video varies but is generally estimated to be around 70-80%[1] [2]. To put things into context, the average mobile webpage size is about 2MB[3] whereas watching a 720p video on your phone could require 2 MB in just 10 seconds. And video data traffic is increasing as new video experiences (short-form content, live streaming) become more popular. Figure 1 shows the trend of video traffic delivered from mobile networks to end-user devices. Whilst the year-on-year growth percentage of video traffic is reducing, volumes are still growing and will continue to grow over the next five years.

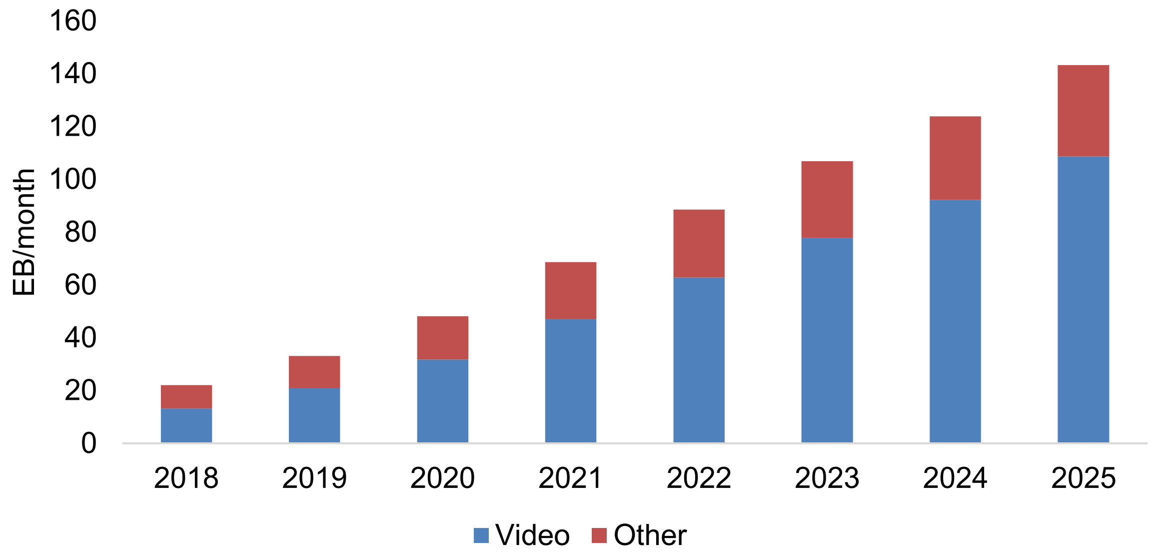


**Figure 1 Global Video Traffic Trend[4]**

---

[1] 50 Video Statistics You Can't Ignore In 2025 | Synthesia.io

[2] Ericsson Mobility Report, November 2024

[3] The Importance Of Page Size And How It Affects Your SEO - SEOptimer

[4] Data extracted from Ericsson Mobility Visualizer - Mobility Report - Ericsson

Video codecs are the technology used to **com**press and **decom**press video. The videos that creators upload to content providers like Meta and YouTube are typically 10x (sometimes even 100x) larger than the ones viewers ultimately watch, so content providers need to shrink videos down before viewers can watch them—this is the job of a codec.

Advanced video codecs like AOMedia Video 1 (AV1), the latest and most efficient codec developed by the Alliance for Open Media (AOM), present an opportunity to support greater efficiencies for video consumption while maintaining or even enhancing the viewer experience. This will benefit network operators, content providers, device manufacturers and end users.

## Network operator perspective: Vodafone

Customers are at the heart of everything that Vodafone does. Supporting video streaming on our networks enables customers to share their experiences on social media, catch up on the latest news reports or watch their favourite TV series while on the go. Our customers expect good quality video and do not want to experience video stalls.

Increasing demand for video is one of the drivers for traffic growth on our networks. To avoid congestion which would lead to a bad video experience, we continually upgrade the capacity of our networks. Due to the costs associated with network capacity upgrades, Vodafone is exploring ways to manage traffic growth. A key way to achieve that is by using more efficient video codecs.

Most video streaming on mobile networks is video hosted by content providers. From the network operator point of view, this video is 'over-the-top', i.e. it is not controlled by the network operators. Nevertheless, Vodafone does have a role to play in raising awareness of the benefits of more advanced video codecs, and is working with content providers to achieve better support of these codecs.

## Content provider perspective: Meta, YouTube

With the popularity of streaming video, it is important for Content and Application Providers (CAPs) to continue exploring new ways to support even greater video efficiencies.

Meta and YouTube are committed to sustainable growth of the video ecosystem. We make significant resource investments in AV1 development and adoption. We are continuously optimizing adaptive streaming algorithms to adapt video quality to network and device conditions.

We regularly invest in modernizing our technical infrastructure, including specialized hardware to support the latest codecs.

We regularly optimize our content delivery networks, including sophisticated content caching to optimize network efficiency and reduce traffic load.

And we are focused on delivering the best possible user experiences through innovative technologies like AV1. We are already using AV1 to deliver high-watch content, and we plan to transition most of our content to AV1 in the coming years.

This is an opportunity to further collaborate among content providers and network operators to align with chipset manufacturers and device operating system developers to guarantee the best quality of experience to end-users, while ensuring optimal utilization of network resources with reduced congestion.

And for content providers, one key way to achieve this is by serving higher quality videos that use less data. When users can stream their favorite shows or movies in crisp, clear video without worrying about buffering or running out of data, they are more likely to be satisfied with their experience.

# 2. Benefits of More Efficient Video Codecs

A video codec is hardware or software that compresses and decompresses digital video. Video codecs play a crucial role in the delivery of video content, as file sizes are significantly reduced compared to uncompressed video. For the case of video streaming from content providers to smartphones, the encoding (compression) is done by the content providers, while the decoding (decompression) is done by the smartphone.

For videos that are recorded on a smartphone, the video will be encoded on the phone prior to being uploaded to the content provider. The content provider may transcode (decode followed by encode) the video into different codecs and quality levels. Given that the content provider transcodes videos and that far more data is streamed to phones than is uploaded from phones, encoding in the smartphone is significantly less important than decoding in the smartphone. On the receiving device, the client algorithm will select the appropriate resolution for optimal user experience based on factors such as the device screen type and available throughput . Adaptive Bitrate (ABR) will be used to adapt the quality and resolution as required. Figure 2 provides an overview of the infrastructure for encoding and streaming of videos.

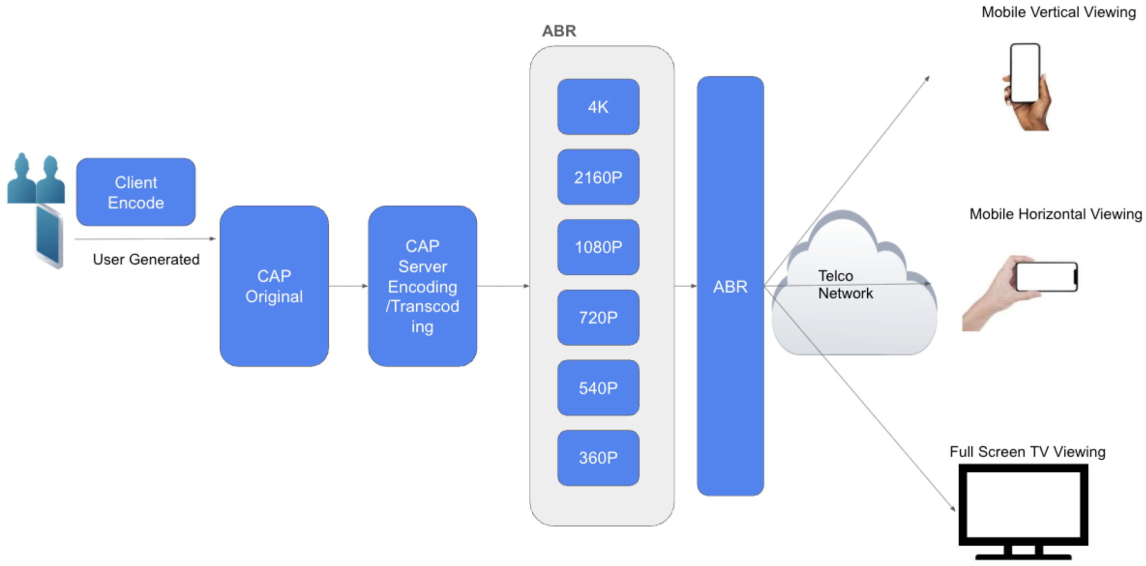


**Figure 2 Video Infrastructure**

There are several different video codecs that are relevant for smartphones. The oldest is H.264, also known as AVC, which was developed by ITU-T and MPEG. The H.264 specification was originally published in 2003, and it is now supported in close to 100% of smartphones. H.265, also known as HEVC, is the successor to H.264 codec and was also developed by ITU-T & MPEG. The first release of the specification was published in 2013.

VP9 is an alternate video codec developed by Google. The VP9 spec was originally published in 2013. AV1, developed by the Alliance for Open Media (AOM), is the successor to VP9. The first release of the specification was published in 2018.

H.266, also known as VVC, is the successor to H.265 and was also developed by ITU-T & MPEG. The first release of the specification was published in 2020. AV2 is a future codec from AOM and the first release of the specification is expected by the end of 2025. Figure 3 provides an overview of the evolution of video codecs.

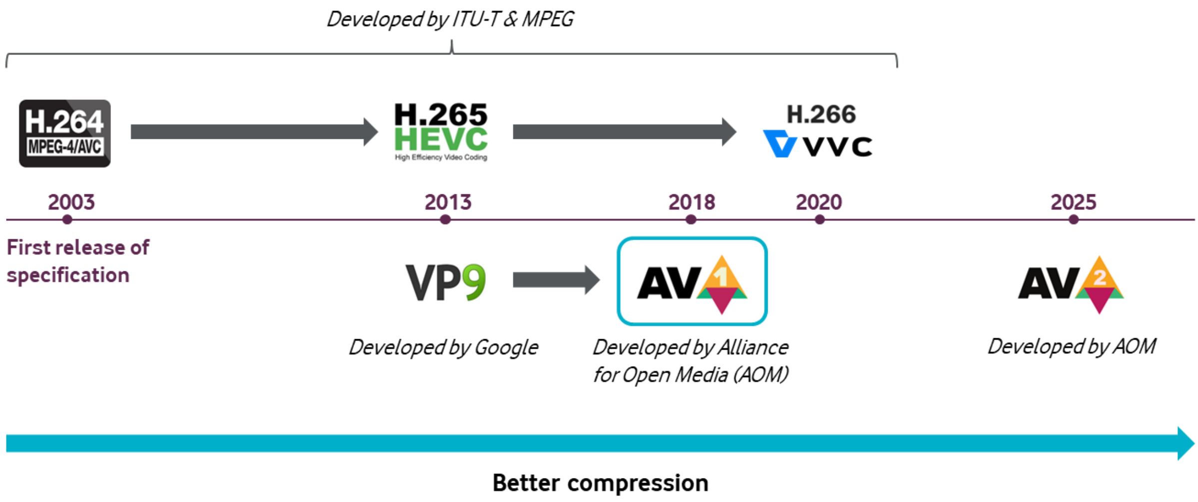


**Figure 3 Evolution of codecs**

The benefit of advanced codecs is that they are more efficient. This means that less data is needed to deliver video (of similar or better visual quality) from the content provider via the network operator and ultimately to the viewer. Less data means less demand on infrastructure for both content providers and network operators. More efficient codecs reduce the demand on cellular networks, improving overall performance. AV1 is one such codec.

AV1 not only benefits mobile network operators and content providers, but also provides a better user experience—faster loading times and smoother streaming experiences, even when network speeds are low. Network speeds can be low due to a weak signal strength or high network load.

# 3. Overview of the AV1 Video Codec

## 3.1 Introduction to AV1

AV1 provides improvements in video compression over earlier codecs such as VP9 and H.264, as highlighted in figure 4. It offers up to 30% better compression than VP9[5]. This provides the ability to deliver superior video quality at the same bitrate relative to earlier codecs or, conversely, to maintain equivalent video quality while reducing bitrate. The increase in compression efficiency comes at the cost of increased computation complexity. This increase in decode complexity creates challenges for video playback, especially on low and mid-tier phones devices with limited compute resources.

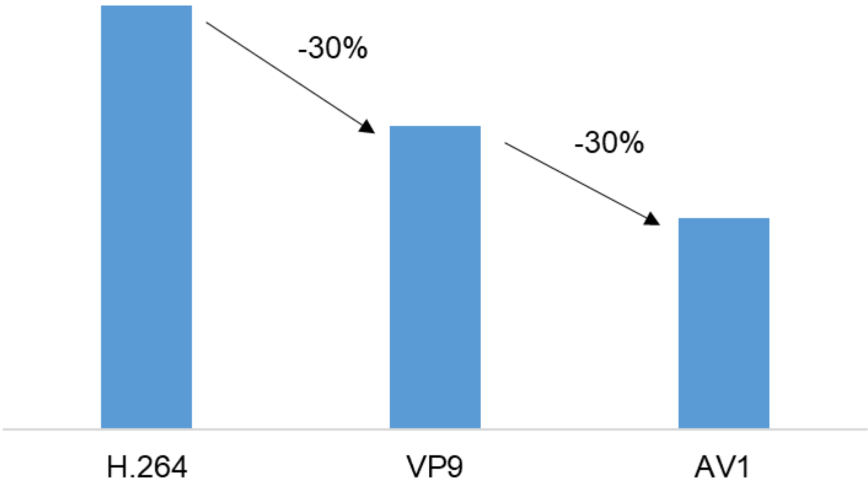


**Figure 4 Codec Compression Efficiency across the generations[7]**

AV1 decoding can be implemented in both hardware and software, which have different costs and benefits.

Hardware codecs are available on most modern, high end mobile SoCs and will ideally become available at lower price points in time. Hardware decode allows for the best network savings because it enables devices to take advantage of advanced AV1 coding features and deliver higher resolutions where the encoding costs are larger in absolute terms. Additionally a fully trusted execution environment is needed to play movies and TV shows on your phone (L1 DRM - Layer 1 Digital Rights Management), which can only be implemented with a hardware codec.

Software decoding using libraries such as the open source dav1d AV1 decoder originally developed by VideoLAN provides a way for lower end and many older devices to decode AV1 streams where hardware support may not be available. Dav1d is highly optimized making decode possible on cheaper (typically less than $300 USD) and older devices lacking a dedicated hardware decoder; however, this comes at the cost of modest increased

[5] Source: Indicative values based on AV1 beats x264 and libvpx-vp9 in practical use case - Engineering at Meta

CPU utilization and battery consumption and lack of support for L1 DRM.

Content providers and network operators generally prefer hardware decode and only fall back to software decoders when hardware decode is unavailable. AV1 provides the ability to optimize network resources and dav1d extends that reach far beyond the availability of hardware decoding. Despite this, the best experience for users would be AV1 video playback using a hardware decoder.

## 3.2 Ecosystem Support for AV1

AV1 was developed by the Alliance for Open Media (AOMedia), a non-profit consortium that develops standards and software for audio and video delivery. Standards contributors do not charge royalties to implementers of those standards. AV1 is open, offers transparency, customisation options, and benefits from a strong community that contributes to its improvement and support.

AOMedia membership spans the media and entertainment ecosystem—from content providers to mobile handset makers to SoC manufacturers and others. The Steering Committee Members of AOMedia are Amazon, Apple, Cisco, Google, Intel, Meta, Microsoft, Mozilla, Netflix, nVidia, Samsung, and Tencent.  In addition, there are currently over 3 dozen Promoter Members. See here for the complete list of AOMedia members.

Table 1 highlights the strong ecosystem support for AV1. The following sections provide more detail.

**Table 1 Overview of mobile ecosystem support for AV1**

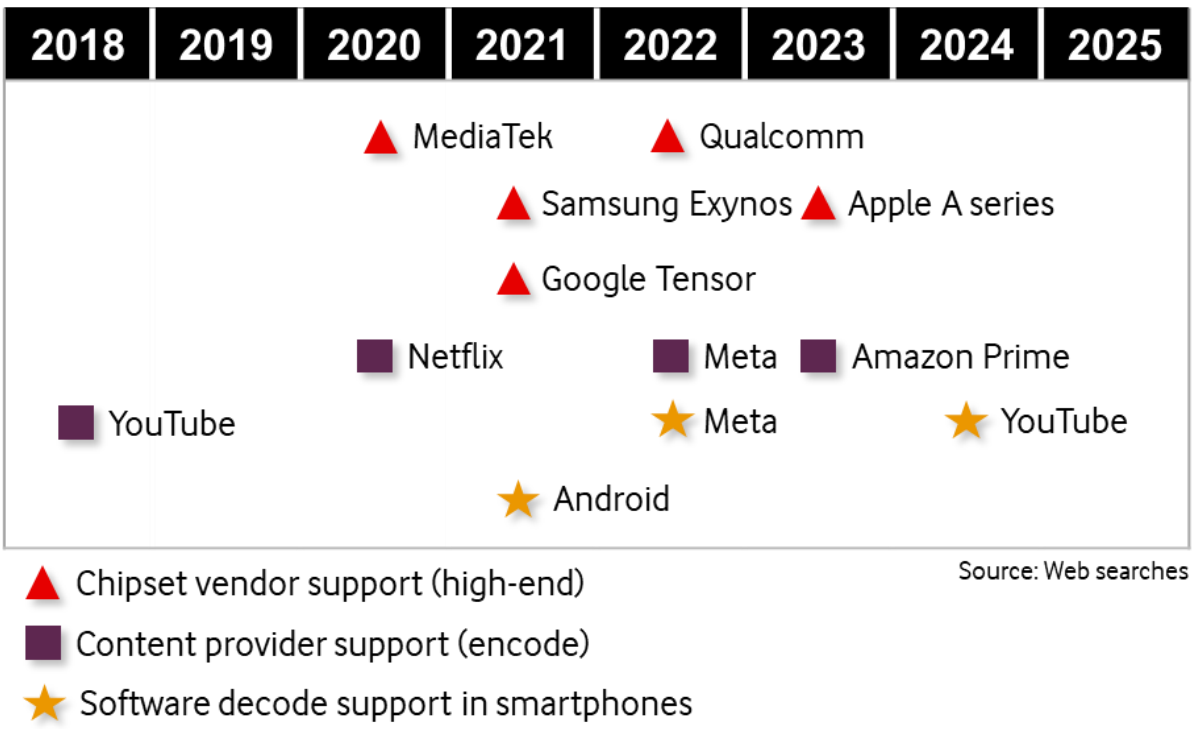

## Content provider support

YouTube launched AV1 for VOD (Video on Demand) playbacks on desktop web browsers in 2018 using software decode. YouTube then launched AV1 for televisions and mobile phones as AV1 hardware support became available, starting in 2023. In 2024, YouTube began deploying software AV1 decode for Android and iOS devices that do not yet have hardware AV1 support.

YouTube has also invested significantly in developing proprietary video processing hardware, engineered to manage AV1's high computational load and optimize AV1 deployment.

**Figure 5 YouTube Video Processing Hardware**

Today, over 75% of YouTube's catalog (weighted by watch time) is available in AV1.

In 2022, Meta launched AV1 for Reels (Facebook and Instagram) on iPhone with software decode and has since grown AV1 overall VOD playback to > 70% of global watch through a combination of hardware and software support.

Meta has made significant investments in dav1d SW AV1 decoder performance over the past few years. While the CPU and battery performance of dav1d is quite impressive on many SoC, the best results have been on high/premium SoC, where HW AV1 decoders are now available. Meta has new work in progress to expand AV1 SW into more low-tier devices.

Netflix has streamed in AV1 since 2020 and notes that it "improved members' viewing experience, particularly under challenging network conditions."[6]

## Smartphones and tablet support

The number of smartphones and tablets that decode AV1 in hardware is growing rapidly. Today, support is mainly concentrated in high end devices, but in the near future, we'll see more mid-tier devices able to decode AV1 in hardware. Table 2 lists some devices and Systems-on-a-Chip (SoCs) that support AV1 in hardware. In May 2025, Qualcomm announced support of AV1 in its Snapdragon 7 Gen 4 mobile platform[7], which will be used in

[6] Bringing AV1 Streaming to Netflix Members' TVs
[7] Snapdragon 7 Gen 4 Mobile Platform | Qualcomm

upcoming devices from Honor and vivo[8].

**Table 2 SoC and HW device support for AV1[9]**

| SoCs | Example Devices |
|---|---|
| Google Tensor (all generations) | • Pixel 6, 7, 8, 9 (all models) |
| Apple A17 Pro, A18 | • iPhone 15 Pro / 15 Pro Max<br>• iPhone 16 (all models)<br>• iPhone 17 (all models) |
| Snapdragon 8 series (all models from Gen 2 onwards) | • Samsung Galaxy S series since the S23 (regional variants with Qualcomm SoC)<br>• Samsung Galaxy Tab since the S9<br>• Oppo Find X6 Pro<br>• Xiaomi 15<br>• Honor Magic V3<br>• and many others |
| Samsung Exynos (2100, 2200, 2400) | • Samsung Galaxy S series since the S21 (regional variants with Exynos SoC) |
| MediaTek Dimensity (1100 1200, 1300, 8200, 9000 and subsequent Dimensity 9000 series) | • Oppo Find X8 Pro<br>• OnePlus Nord 3<br>• VIVO X80 |

Where hardware decode is not yet available, software decoders can be used to play AV1 video. As an example, Unisoc includes a SW AV1 decoder with the T606 SoC that powers the Moto e13, a sub-$100 Android phone (see case study below). And Android provides an optional software AV1 decoder in Android 10 and a more efficient dav1d based decoder in Android 12.

[8] Qualcomm Unveils the Snapdragon 7 Gen 4 Mobile Platform: A Multimedia Powerhouse Fueling the Next Generation of Mobile Experiences | Qualcomm

[9] Source: Web searches

## 3.3 Case Studies

### Case study: YouTube AV1 deployment

YouTube produces AV1 transcodes for long and short form user generated content and began serving AV1 transcodes to mobile devices as hardware codecs became available in 2021. Starting in 2024 YouTube began using software AV1 decode of short form content on Android for playback scenarios where battery impact would be minimal. In 2025, YouTube launched AV1 support for popular live streams. YouTube will continue to expand AV1 support, including deploying software AV1 to more iPhones and Android phones that do not have hardware AV1 support.

### Case study: Meta AV1 deployment

In 2022, Meta launched AV1 for Reels playback on the iPhone[10], and very quickly, AV1 Reels watch exceeded 70% on iPhones. Since that time, the focus has been on bringing AV1 to the Android ecosystem, which has been more challenging. Currently, > 70% of global video watch on the Meta Family of Apps is AV1. Key to Meta's success in deploying AV1 has been the work to understand the Android ecosystem, and increase reach in the mid and low tier populations.

10 "How Meta brought AV1 to Reels", Ryan Li, et al..

# 4. AV1 Video Encoding

## 4.1 Implementation of AV1 encoding

Efficient streaming of video at global scale requires careful allocation of finite compute and edge cache resources. Content providers need to use enough video codecs to maximize user reach while minimizing the number of transcodes to efficiently distribute video at scale. Social media platforms have additional challenges adopting new codecs due to the sheer volume of user generated content uploaded on a daily basis as well as the significant back catalog of content already available on the platform.

Effective adoption of AV1 requires the ability to transcode uploaded content quickly and efficiently, generally using a combination of hardware and software encoding. VOD content has some flexibility in the time taken to transcode but live streaming use cases require real-time transcoding, likely requiring dedicated encoding hardware. And as Live events increasingly move to OTT (“over-the-top”) delivery, real-time Live AV1 transcoding becomes a more important part of content provider video delivery strategies.

Transcode diversity places a strain on edge cache resources. The more codecs supported, the less efficient edge cache codes can be. AV1 has reached the point where device penetration and reduced transcode size has offset at least some of the inefficiencies created by introducing a new codec to the content delivery network. Ongoing investment will further improve efficiency and unlock AV1's full potential for the entire ecosystem.

## 4.2 Energy-Efficient Datacenter Encoding

When deploying video codecs at scale, it is essential to balance three critical factors: video quality, file size (contributing towards egress), and computational effort (which can affect energy consumption). Each of these factors significantly influences the overall performance and cost-effectiveness of the deployment, and when improving one factor, at least one of the other two will regress. For example, if improved bitrate efficiency is needed, one may choose to use a higher compression setting or move to a newer, and more complex codec. Either of these options will require more encoding compute. If additional compute is not available, and egress must be reduced, quality must be regressed. Put simply, better compression can be achieved by using more computational effort for the encoding, but this results in higher energy consumption. Traditionally, video codecs have focused primarily on enhancing video compression efficiency due to the historically high costs of bandwidth. However, in today's context it is equally important to prioritize energy efficiency. Note that tradeoffs are eased by investments in video processing hardware. YouTube’s custom hardware video encoder reduces by over 95% the amount of electrical power used to convert creators’ uploads into videos that can be watched by YouTube viewers.

With this in mind, it can be seen that considering video codecs solely on how many bits might be saved is inadequate. Rather, codecs should be considered from a cost/benefit perspective. The question is then shifted. For a given amount of energy, how does the bitrate efficiency of Codec A compare to Codec B? The convex-hull encoder analysis methodology was developed specifically to answer this question. This method aligns quality across codecs, measures egress efficiency, and evaluates it against the corresponding energy usage whether it's a software or a hardware-based implementation.

Figure 6[11] illustrates the application of this methodology, comparing different exemplary implementations of AVC, HEVC, VP9, and AV1 in terms of both compression and energy efficiency. Here, the x-axis represents, in logarithmic scale, the energy utilized by each implementation (with lower values being preferable), while the y-axis indicates the compression efficiency across the codecs as compared to SVT-AV1 M0. Again, lower is better.

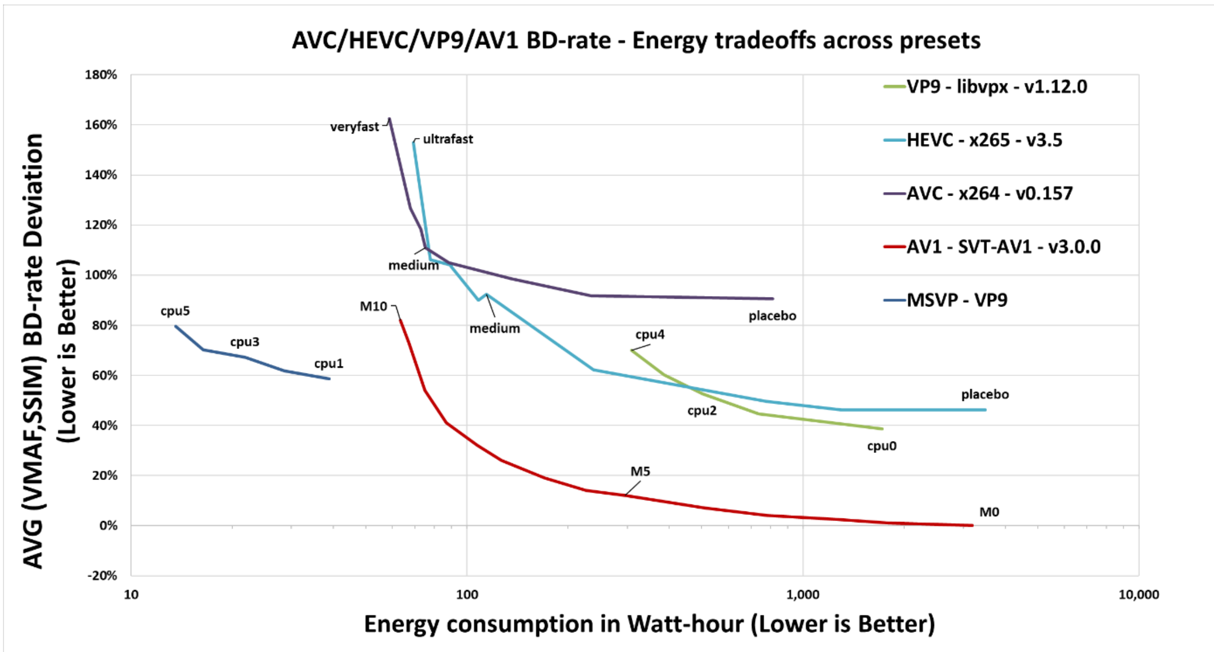


**Figure 6 Compression and energy efficiency for various video codecs**

As illustrated, the red curve representing the AV1 software implementation "SVT-AV1" demonstrates a substantial potential for energy-efficient compression efficiency gains when comparing these exemplary encoders and versions. In fact, SVT-AV1 has the convex hull over 2.5 orders of magnitude, delivering ~40% fewer bits at the same power cost as x264 'veryfast' for the same energy expenditure.  Further, SVT-AV1 can deliver ~30% fewer bits than x265 at the same energy level.

[11] https://engineering.fb.com/2023/02/21/video-engineering/av1-codec-facebook-instagram-reels/

# 5.0 AV1 Decoding in the Mobile Ecosystem

## 5.1 Mobile Device Tiers

Devices in the mobile ecosystem are thought of in tiers, based on capabilities, performance, and price. High-tier devices such as Google Pixel 10 Pro, iPhone 17 Pro, and Samsung Galaxy S25 Ultra have cutting edge technologies, fast processors, advanced batteries, and high-quality displays typically 1440 pixels wide. Further, the SoC in such devices will have a modern ARM CPU and use the most advanced process technologies, typically ≤ 4nm (discussed below in Appendix 2). Finally, high-tier devices will be the first devices to include HW decoders for the latest video codecs such as AV1. High-tier devices represent ~25% of new mobile devices globally[12].

Mid-tier devices offer a balance between price, performance, and features, and are typically in the $300-$600 price range in the USA. While these devices do not match the performance of high-tier devices, they still have modern ARM cores and will typically use manufacturing technologies between 4nm - 6nm.

Low-tier devices are budget phones, and typically range from $30 - $300 in price in the USA. These devices have large variations in performance and capabilities. The cheapest low-tier devices might have ARM 7 and/or 28nm process technology (discussed below), and run 32-bit Android OS. Further, they will have a smaller amount of slower memory, less power-efficient CPUs, and less capable batteries.

Low-tier and mid-tier devices represent ~75% of new mobile devices.

## 5.2 AV1 Decoder Levels

The AV1 specification defines levels as a means for HW decoders to signal their capabilities, and for encoders to ensure decoder compatibility. The fundamental metrics of AV1 decoder levels are pixels/second - which in turn, determines the resolution and framerate limits - and bitrate. A simplified way to think of levels is in terms of resolution and framerate as shown in table 3. Note that the max bitrates defined by the spec for a given level are well above the practical bitrate needed for a given video at the specific resolution.

[12] Source: Global Premium Smartphone Share Climbs to 25% in 2024 as Premiumization Continues

**Table 3 AV1 level examples**

| AV1 Level | Main Profile Max Bitrate | Example |
|---|---|---|
| 3.1 | 10 Mbps | 720x1440 @30fps, 540x1080 @60fps |
| 4.0 | 12 Mbps | 1080x2160 @30fps, 720x1440 @60fps |
| 4.1 | 20 Mbps | 1080x2160 @60fps, 1440x2880 @30fps |
| 5.0 | 30 Mbps | 1440x2880 @60fps, 2160x4320 @30fps |
| 5.1 | 40 Mbps | 2160x4320 @60fps |
| 6.0 | 60 Mbps | 7680×4320 @30fps |
| 6.1 | 100 Mbps | 7680×4320 @60fps |

## 5.3 AV1 Decoding in Hardware

Decoding video in hardware requires less power than SW, and thus HW video decoding is preferred. Further, HW decoding is necessary for premium video use-cases where L1 DRM[13] is required. Over the past few years, most high/premium tier Android SoC have included HW AV1 decoder support, and the iPhone 15 Pro and iPhone 16 onwards all have HW AV1 decoding.  Since most low and mid-tier Android devices implement lower cost SoCs that do not include a HW AV1 decoder, those phones are not able to play AV1 using HW decode. It is expected that AV1 HW may reach some - but not all - mid-tier SoCs by 2026, but HW AV1 decoder reach in the Android ecosystem is on track to grow slowly.

### Way forward on AV1 decode in hardware

We are seeing increased AV1 adoption at the top of the middle tier price range both on iOS and Android through recently released phones such as the iPhone 16e, OnePlus 12R, etc. These phones are generally based around higher end SoCs such as the Apple A18 or Qualcomm Snapdragon 8. In future, hardware decoder availability within middle tier SoCs such as the MediaTek Dimensity 6000 and 7000 series, Qualcomm Snapdragon 6, or Samsung Exynos 1500 line would be a positive development.

We note that SoC revisions are slowing down and it is common to see newer phones based around SoCs from prior years. Additionally, SoC vendors are focusing more on increased AI capabilities through more powerful NPUs, which leaves less space on the silicon die for decoders. Video playback is still an important mobile phone use case and increased hardware AV1 availability allows users to benefit from an improved video playback quality, lower data usage and increased battery life.

One key to increasing the reach of AV1 HW decoding at devices with lower price points is to

[13] L1 DRM requires a “Trusted Execution Environment" (TEE), and SW decoding can not meet the TEE requirements.

level-limit the decoder to the devices. For example, an SoC that will sell into low-tier devices that will typically have a 720p display could be level-limited to 3.1 (720p @30fps). The following tier-specific level limits are recommended.

**Table 4 Tier-specific minimum level limit**

| Tier | Level Limit | Example |
|---|---|---|
| Low | 3.1 | 720x1440 @30fps, 540x1080 @60fps |
| Mid | 4.0 | 1080x2160 @30fps, 720x1440 @60fps |
| High | 4.1 | 1080x2160 @60fps |

## 5.4 AV1 Decoding in Software

### AV1 SW decoding across the mobile device tiers

As discussed above, devices in the mobile ecosystem are thought of in tiers based on capabilities, performance, and price. High-tier devices have cutting edge technologies, fast processors, SoCs built using advanced process technology, and advanced batteries. In addition, nearly all new high-tier devices have HW AV1 decoding. For the small cohort that do not include HW AV1 decoder, SW decoding is not challenging, even at high resolutions and framerates (i.e. 1440p @60fps). HW AV1 decoders are expected to arrive in more mid-tier devices in the next few years. In the meantime, modern mid-tier devices can decode AV1 in SW up to 1080p @60fps but with suboptimal user experience including increased battery drain and less responsive application performance. Many low tier devices might not see HW AV1 decoders for several years, if ever, and also present the most challenges for SW AV1 decoding. A successful AV1 strategy requires AV1 delivery to as many low-tier devices as possible, which in turn, means that AV1 SW decoding is necessary on these devices. Fortunately, many low-tier devices are capable of SW decoding AV1 so long as streaming applications can determine the performance capabilities of the device, and determine whether or not that device can successfully SW decode AV1, and if so, at what level. These considerations are discussed in Appendix 2 below. The quick summary for device makers to support AV1 SW decoding is to move off of older and less-capable ARM CPU and on to modern A55-class ARM V8 cores. As shown by the Moto e13 (discussed below), even sub-$100 devices can support the newer ARM architecture.

### Way forward on AV1 decode in software

While the ultimate goal is AV1 HW decoders in all mobile devices, the timeline for AV1 HW decoders in low-tier devices is likely very long. Therefore, SW AV1 decoding is necessary. As previously discussed, SW AV1 decoding works well if the mobile device meets the baseline criteria for ARM version & core types, process technology, and 64-bit OS. However, it is not possible, for example, for an application to probe the device to learn the process technology. Thus, to get good AV1 performance, content providers may need to implement complex models to determine device suitability. Further, they may need to bundle the dav1d decoder with their apps in order to tune performance or support older devices. Adding such complexity to content providers puts significant friction on using AV1, and will hinder the adoption of AV1.

One approach is to follow the example of Unisoc (see case study below), and have SoC vendors provide an AV1 SW decoder that is optimised for their architecture, and which has level limits appropriate for their platform. This will allow CAPs to deliver AV1 to devices with built-in SW decoder that is level-limited, and have confidence that the device can deliver a good AV1 experience within the level limits.

For cases where the SW decoder is included in the app or the OS, having a single industry-wide model for benchmarking SW video decoders, and specifying AV1 support capability will enable a better experience for users as content providers will have a clear understanding of the SW decoder performance capabilities of a given device. Meta's Video Codec Acid Test[14] (VCAT) is one such model.

## Case study: Moto e13 with Unisoc T606 SoC

In 2023, Moto introduced the e13, a sub-$100 Android phone based on the Unisoc T606[15] SoC with 2GB RAM and a 720x1440 display. When the e13 shipped, it was observed that it included an AV1 decoder, 'c2.unisoc.av1.decoder'. While appearing as HW, the decoder is actually a Unisoc-supplied variant of the dav1d SW AV1 decoder. Seeing a $90 device with AV1 support gave motivation to run benchmarks to determine how well the device could support AV1, and it was decided to also test the SVT-AV1 'fast decode' recipes on this device. The test conditions were radio off, display brightness at 25%. The test video was 720p @30fps, encoded in (reference) and SVT-AV1 Fast Decode 0/1/2, and was played from a local file. Playback done using VLC, and was looped for 16 hours, using the default VLC SW decoders[16]. The results are shown in figure 8. The VP9 reference[17] had a battery drain of ~2.6%/hour. AV1-FD2 had drain of ~3.2%/hour, a difference of 0.6% hour. While the test was only run for 16 hours, it can be seen that in a full battery drain test using AV1-FD2, the Moto e13 would run for about 30 hours. This is quite impressive for a $90 phone. Of note is that the T606 SoC has a modern arm core, and is manufactured with 12nm process technology (see Appendix 2 below), which is the baseline for good AV1 SW decoder support.

[14] https://www.streamingmediaglobal.com/Articles/News/Featured-News/Review-Metas-Video-Codec-Acid-Test-(VCAT)-170508.aspx

[15] The T606 uses 12nm process technology, and has ARM V8 Octa-core with 2x1.6 GHz Cortex-A75 & 6x1.6 GHz Cortex-A55.

[16] VLC was using dav1d 1.4.2.

[17] Starting in 2021, VP9 SW decoding represented a majority of Meta VOD delivery.

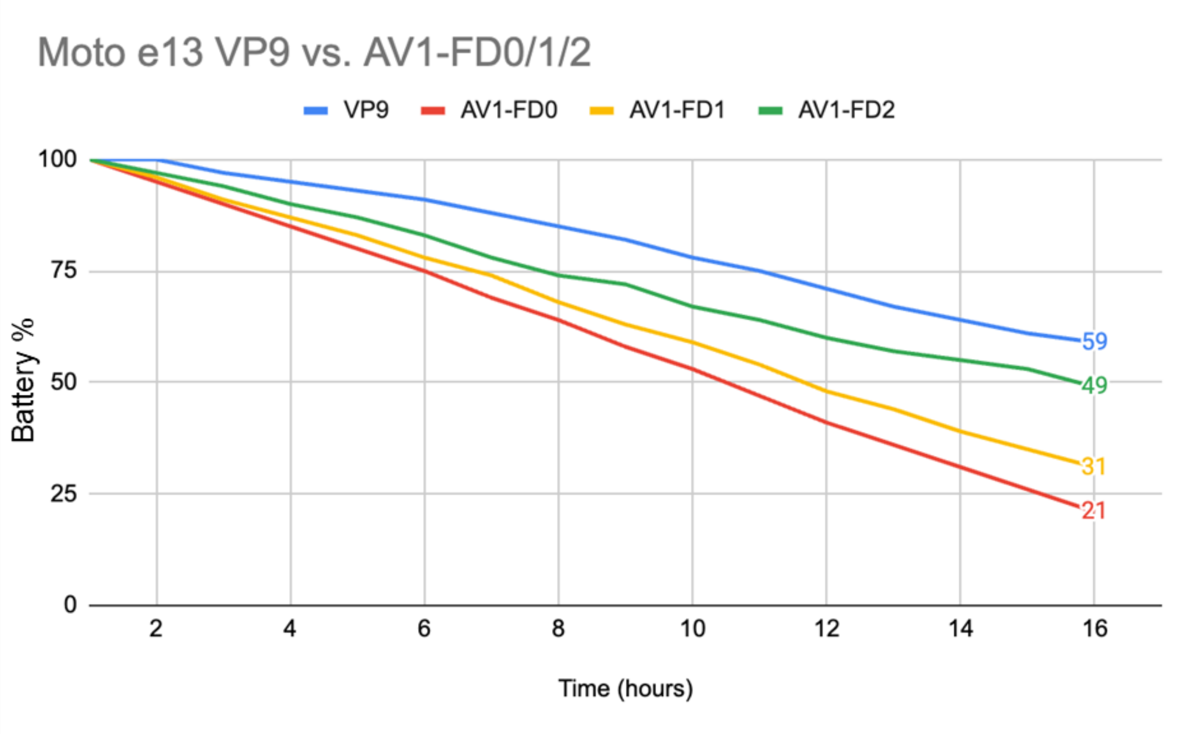


**Figure 7 Moto e13 battery drain during video playback**

## 5.5 Leveraging Both Hardware and Software Decode to Maximise AV1 Benefits

The benefits of AV1 are clear, and most content providers are moving to AV1 as their flagship codec. The CAPs such as Meta, Netflix, and YouTube have embraced AV1 and all have roadmaps to further shift traffic to AV1. Further, since Mobile World Congress 2024, many Communication Service Providers (CSPs) are also seeing AV1 as an opportunity to reduce video bits.

Android's Compatibility Definition Document[18] lists AV1 as one of multiple video encoders and decoders that are required on Android-compatible handheld devices, whether that support is in hardware or software.

We encourage SoC vendors to invest in HW AV1 across all device tiers to meet the AV1 requirements to deliver an improved viewer experience, device battery savings and enhanced network operator infrastructure efficiency.

[18] https://source.android.com/docs/compatibility/16/android-16-cdd

# 6. Conclusions and Next Steps

Continuing to ensure efficiency of video delivery is important for both mobile network operators and content providers. Video streaming can be made more efficient by using improved video codecs. The AV1 video codec offers improved compression and better video quality at constant bitrate vs the HEVC and VP9 codecs — expect further improvements with AV2. Popular content providers (Meta, Google, Netflix) already support AV1 streaming on their platforms by encoding videos in AV1.

AV1 hardware support in high-tier smartphones is already good and is expected to expand to more mid-tier smartphones starting in 2026. By making hardware based AV1 decode available across all smartphone platform tiers on a timely basis, device SoCs will enable AV1 support across the full Android smartphone range. Support in mid and low tier chipsets will result in far greater support of AV1 hardware decode, increasing the benefits to mobile network operators, content providers and end users.

Software AV1 support is available but has a modest negative impact on phone battery life relative to hardware decode. However, content providers only use software decode as a fallback when hardware decode is unavailable. It is important to ensure that these software implementations are efficient. Our recommendations for software decode are:

- For software decode, we recommend implementing a libdav1d based software approach. OEMs should investigate platform specific optimisation for better performance.
- For cases where the SW decoder is included in the app or the OS, a single industry-wide model for benchmarking SW video decoders and declaring level-support will enable content providers to easily determine the SW decoder performance capabilities of a given device.

We believe that, in the long term, the greatest benefits can be achieved by support for AV1 hardware decode.

## Recommendations

**Device OEMs and SoC providers**

- SoC vendors should consider supporting AV1 in hardware, which allows for the highest resolution video and supports playing DRM-protected movies and TV shows
- In scenarios where HW AV1 is not an option, SOC vendors and device manufacturers should consider adopting a platform-optimized SW AV1 decoder
- A SW AV1 decoder can help with the transition to AV1 in low/mid-tier devices
- A shift in low-end devices from A53 ARM Cores to A55 ARM cores will improve battery life and device performance for SW video decoding.

# Appendix 1: Impact of Accelerating AV1 Adoption in Devices

Video codecs reside inside end user equipment and Content and Application Provider (CAP) service platforms. They determine the quality and compression ratio of IP video signals, and thus the IP data requirements of a given quality video stream.

Accelerating the adoption of AV1 is expected to have a significant positive effect on telecommunications operators' investments. AV1 can also help manage data traffic growth. A video encoded with AV1 has an average bit-rate reduction of approximately 30% over that same video encoded in VP9 and HEVC, and around a 70% reduction over the video encoded in H.264, with actual performance varying based on video complexity.

Given that video traffic is projected to represent 70% to 80% of total Internet traffic by 2030, the widespread use of AV1 could significantly reduce the amount of traffic transmitted across telecommunications networks.

However, the actual cost savings or investment impact from AV1 adoption will vary depending on each operator's network characteristics, operational strategies, investment plans, levels of cooperation, and potential for synergies or efficiencies. As a result, the benefits of AV1 adoption may differ significantly across regions and countries.

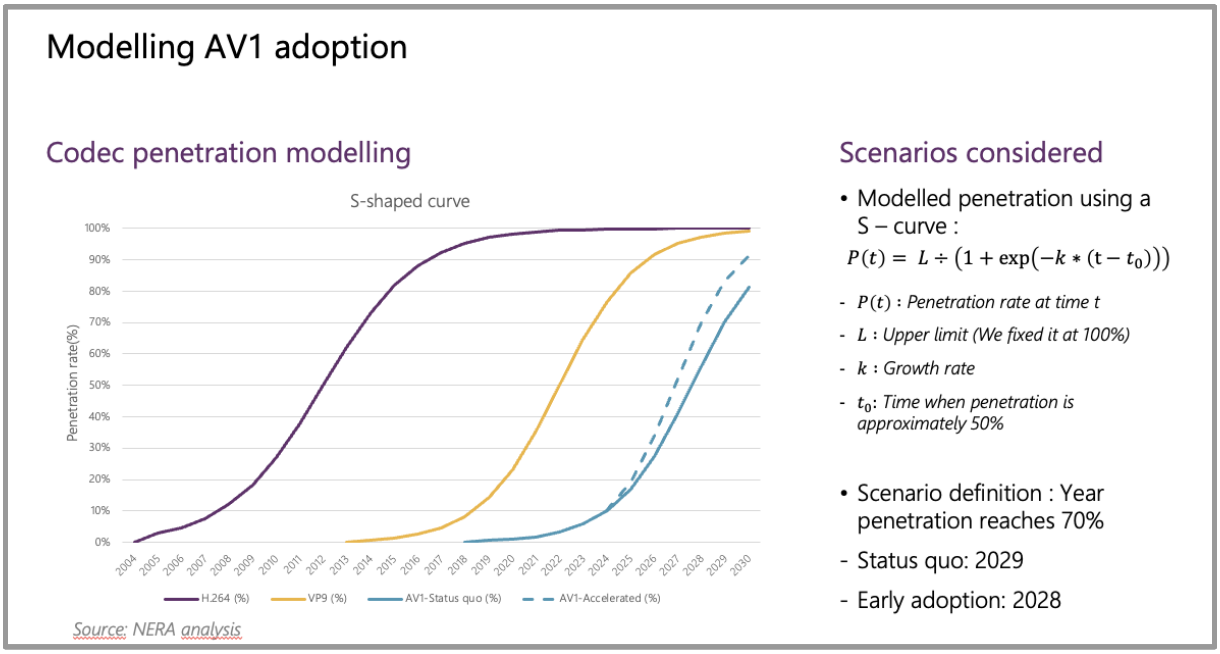


**Figure 8 Early Adoption of AV1**

**Table 5 Hypothetical calculations of traffic potentially avoided in Europe assuming early and full adoption of AV1**

| Country | Avoided traffic (PB/year) | | | | | | Total avoided traffic 2025-2030 (PB) |
|---|---|---|---|---|---|---|---|
| | 2025 | 2026 | 2027 | 2028 | 2029 | 2030 | |
| EU-27 | 5,968 | 21,624 | 47,093 | 72,008 | 82,629 | 76,225 | 305,548 |
| UK, Switzerland and Norway | 1,567 | 5,490 | 11,588 | 17,201 | 19,187 | 17,226 | 72,258 |
| Total Europe | 7,534 | 27,114 | 58,682 | 89,209 | 101,816 | 93,451 | 377,806 |

Those pushing for an accelerated adoption of AV1 stand to benefit all parties across the digital ecosystem—from network operators and OTTs to end users—by enabling more efficient, sustainable, and cost-effective video delivery.

Source: Nera Consulting Report on AV1 codecs

# Appendix 2: Key Factors in Software Decoder Performance

The capability for a smartphone SoC to handle software decode of AV1 depends on the technology and architecture used for the CPU. Most CPUs in smartphones use an ARM architecture. A number of factors will determine whether the CPU has the required processing power and power efficiency to decode AV1, and these are highlighted in table 6.

**Table 6 Key factors in software decoder performance**

| Factor | Description |
|---|---|
| Process Technology | Process technology (also called node size) is measured in nanometers (nm), and refers to the transistor fabrication process in a chip. Smaller nodes offer higher transistor density, lower power consumption, and better performance per watt. For example, when compared to 28nm process technology, 12nm will use ~40% less power, 6nm ~65% less power, and 3nm ~75% less power. In addition, compared to 28nm, heat generation is 25%, 50%, and 65% less for 12nm, 6nm, and 3nm respectively. |
| ARM Version & core types | ARM 8 & 9 are 64-bit big-little architecture. Video decoding can typically run on little cores, requiring less power. ARM7 are 32-bit and many ARM7 CPU’s lack NEON. |
| NEON | NEON is ARM’s SIMD (single-instruction, multiple-data) extension, and reduces CPU load and power consumption. NEON can reduce the computational load of AV1 decoding by up to 5-6X. SIMD is essential for power-efficient video decoding. |
| ARM8 Core types | ARM8 CPUs have either A53 or A55 little cores and A73 or A75+ large cores. A55-class SoC are 15%-20% more power efficient than A53-class devices. Further, SoC with A53 are typically much older, and thus use older process technology. |
| 64-bit vs. 32-bit Android OS | When 32-bit Android is installed on a 64-bit device, it limits performance and efficiency by preventing apps from utilizing 64-bit processing optimizations, such as improved memory addressing and SIMD extensions like SIMD/NEON. |

## Level-based classification of devices for AV1 SW decoding

The AV1 specification defines levels as a means for HW decoders to signal their capabilities, and for encoders to ensure decoder compatibility. The fundamental metrics of AV1 decoder levels are pixels/second - which in turn, determines the resolution and framerate limits - and bitrate. A simplified way to think of levels is in terms of resolution and framerate as shown in table 7. Note that we put max bitrate limits on the levels that are well below those specified by the AV1 specification[19], but well above the practical bitrate needed for a given video at the specific resolution.

**Table 7 AV1 level examples**

| AV1 Level | SW Max Bitrate | Example |
|---|---|---|
| 3.1 | 4Mbps | 720x1440 @30fps, 540x1080 @60fps |
| 4.0 | 8Mbps | 1080x2160 @30fps, 720x1440 @60fps |
| 4.1 | 12Mbps | 1080x2160 @60fps, 1440x2880 @30fps |
| 5.0 | 16Mbps | 1440x2880 @60fps, 2160x4320 @30fps |
| 5.1 | 20Mbps | 2160x4320 @60fps |

The concept of HW levels can be applied to SW decoding. Using level-limited playback for SW AV1 decoding significantly simplifies the problem of delivering AV1 only to devices that can successfully decode AV1[20]. This is illustrated in the figure below. By classifying mobile devices in terms of max level for SW AV1 decode, and limiting streams delivered to each device to its maximum level capability, the problem of AV1 SW decoding is greatly simplified. Of note is that this is exactly how a content provider should approach HW decoding on a device. In lab benchmarking, many A53-class devices appear to be capable of decoding AV1 to level 3.1 reasonably. However, there are many with similar specifications that can not.  Thus, as a group, A53-class mobile devices can not reliably support AV1, and new low-tier devices should move to ARM A55 CPU, and modern process technology.

[19] Higher bitrates result in significantly higher decode complexity due to entropy coding.

[20] In this white paper, “successful decoding” of AV1 means being able to decode without frame drops or thermal events, and with the battery able to sustain 14+ hours of continuous decoding on a single charge.

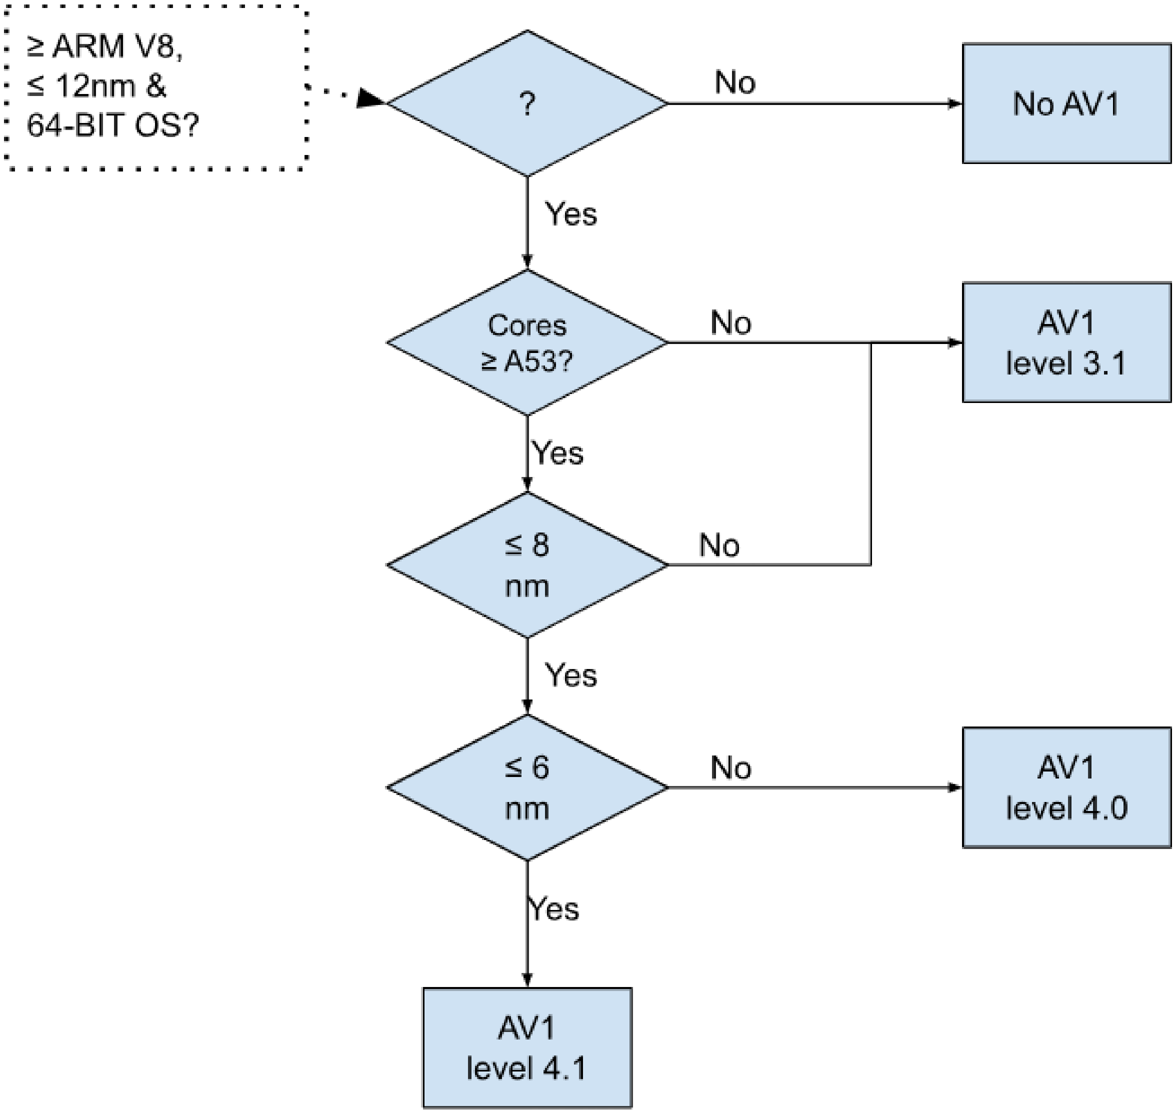


**Figure 7 Determining a mobile device's AV1 SW decoding capability**[21]

[21] In benchmarking and field tests, there is a wide variation in the ability of A53-class devices to support SW AV1 decoding. There are many possible reasons such as bus, memory, cache, higher leakage, and lower efficiency curves. Generally the older the SoC, the more troublesome AV1 SW decoding will be.