

How av1 encoding quality is changing the way we watch video

Why av1 encoding quality matters more than ever

For years, video compression has been a quiet battle between file size and visual fidelity. Every streaming platform faces the same challenge: deliver crisp, detailed video without breaking the bandwidth budget. av1 encoding quality has emerged as the most promising solution to that trade-off. It allows creators and platforms to push higher resolution and better clarity over the same internet connection. That is a big deal for anyone who uploads or streams video regularly.

I have been working with video codecs for over a decade, and I have seen the shift from h.264 to h.265 and now to AV1. Each step brought improvements, but none felt as fundamental as what AV1 offers. The difference is not just incremental. It is structural. av1 encoding quality, when properly tuned, can deliver the same visual result as h.265 at roughly 30 to 40 percent lower bitrate. That means better looking video for viewers with slower connections, and lower hosting costs for creators.

The technology behind better compression

AV1 was developed by the Alliance for Open Media, a consortium that includes major tech companies. The codec uses a set of advanced tools that previous standards lacked. For example, it can partition frames into much smaller blocks, adaptively choose between different prediction modes, and apply more sophisticated loop filters. These features allow the encoder to remove redundant information more aggressively while preserving the details that matter to the human eye.

One of the most noticeable effects of good [av1 encoding quality](#) is how it handles high-motion scenes. Fast action sports, explosions in movies, or quick camera pans often break older codecs into visible artifacts. Blockiness and mosquito noise creep in. AV1, on the other hand, maintains cleaner edges and smoother motion at the same bitrate. That is not just a theoretical advantage. It is something you can see when you compare side-by-side clips on a large screen.

But there is a catch

AV1 encoding is computationally expensive. Encoding a video in AV1 can take several times longer than encoding the same video in h.264 or even h.265. That is because the encoder has to evaluate many more options to find the most efficient way to compress each frame.

Hardware encoding is helping, but software encoders still dominate for high-quality production work. If you are running a small channel or a boutique production studio, the extra render time can be a real cost.

That said, the payoff is real. Once the video is encoded, playback is efficient. Modern GPUs and mobile chips include hardware decoders for AV1, which means the viewer does not bear the computational burden. The trade-off is upfront processing time for downstream savings in bandwidth and storage. For many workflows, that is a worthwhile exchange.

Practical advice for content creators

If you are producing video for a platform like YouTube, you have probably noticed that your uploads get transcoded into multiple formats anyway. YouTube re-encodes everything you upload, and they have been using AV1 for their own delivery for years. But the starting quality of your source matters. If you upload a highly compressed h.264 file, the re-encode into AV1 can only work with what is there. Garbage in, garbage out.

To get the best av1 encoding quality from your own pipeline, start with a high-bitrate master. Export your edit as a ProRes or DNxHD intermediate if possible. Then feed that into an AV1 encoder. For software encoding, SVT-AV1 from Netflix is a solid choice. It offers a good balance of speed and quality. Use the preset that matches your patience: slower presets produce better quality but take longer.

Here are a few settings I have found useful:

- Use CRF mode instead of fixed bitrate for variable quality needs. A CRF value of 18 to 22 works well for most content.
- Enable film grain synthesis. It preserves natural noise in film-like content without wasting bits.
- Set the keyframe interval to match your scene changes. Two to four seconds is typical.
- If your source is already noisy, apply light denoising before encoding. AV1 handles clean sources better.
- Test a short clip before encoding the full video. Adjust settings based on what you see.

These tweaks are not magic, but they help you get closer to the theoretical efficiency of the codec. Every bit counts when you are trying to deliver high quality at low bitrates.

Measuring quality is not simple

There are objective metrics like PSNR, SSIM, and VMAF. They give you numbers, but numbers do not always match human perception. I have seen clips that score well on VMAF but look flat and lifeless. And I have seen clips with mediocre scores that look great because the encoder preserved texture and detail in the right places. The best judge is still your own eyes, especially on a calibrated monitor.

That is why I recommend building a small test set of challenging scenes: a night shot with city lights, a fast-moving sports sequence, and a slow pan across a landscape. Encode each at several bitrates and compare them side by side. You will quickly learn where your encoder struggles and where it excels. That hands-on knowledge is more valuable than any benchmark.

The future of video compression

AV1 is not the last codec we will see. VVC, also known as h.266, is already being developed. But AV1 has a head start in adoption. Every major browser supports it. Most new smartphones have hardware decoders. Streaming services use it for their premium content. The infrastructure is in place.

What matters for creators is that av1 encoding quality keeps improving. The encoders get faster every year. The decoders get more efficient. And the visual results keep getting better. If you have not tried encoding in AV1 yet, now is a good time to start. The learning curve is real, but the payoff in quality per bit is hard to ignore.

I have personally switched my primary encoding workflow to AV1 for all new projects. The extra encoding time is a small price to pay for the confidence that my videos will look their best on any screen, at any connection speed. And I have seen viewers notice the difference in comments and feedback. That kind of real-world validation is the best metric of all.

Video compression is not glamorous, but it is the invisible engine that makes modern streaming possible. AV1 is the best engine we have right now. Understanding how to use it well is a skill that will serve any creator for years to come.