How Online Video Works

A technical overview

Senko Rašić
www.senkorasic.com

Table of contents

Introduction

Video is a big part of our online experience. It drives the majority of all internet traffic (73% in 2016, expected to rise to 82% by 2021), and people watch the equivalent of 3.6 billion DVDs every month (source: Cisco VNI 2016-2021). And that's just video on demand (VOD). Live events pull even more engagement, up to 10 times that of online video (source: Frost & Sullivan).

These trends make a compelling argument for including video as a key component of online presence for any brand, publisher, enterprise, or platform. Video can engage viewers and drive innovation in ways we never even imagined a decade ago.

Whether you decide to create or improve your online video strategy, you have a profusion of choice. The sheer volume and diversity of solutions, technology platforms, and approaches for producing, publishing, and delivering video can seem overwhelming.

This white paper outlines the important elements of a successful strategy and highlights key issues you should recognize.

Online video strategies

Broadly speaking, businesses use online video to support one of two activities: information sharing and monetization.

Information sharing encompasses marketing, enterprise communication, lead generation, knowledge dissemination, and other such activities. Monetization involves generating revenue through ads, subscriptions, or other means of charging for content access.

Successful online video use requires a well-defined strategy. For example, an enterprise seeking to publicize a global product launch will use a different approach and technology from a video blogger monetizing a YouTube channel. Social networks are popular platforms for individuals to create, publish, and stream their videos. However, businesses may want to exclusively brand their content and integrate it with their content management systems (CMSs), mobile apps, and other resources they publish.

Information sharing

Whether you want to market a new product or service, share knowledge with employees or clients, or increase brand visibility, video can be a powerful tool.

Inside an organization, video can align and unify workers, improve communication and knowledge transfer, and facilitate telecommuting. In a survey of 1300 global executives, 87% consider video beneficial to their organizations. (source: Cisco).

Video's dynamic and visual nature makes it more memorable and shared more often than other media types, such as text or images. More than 60% of professional marketers and small-to-medium-sized business owners believe in video enough to increase their current investment in it (source: Animoto).

Monetization through ads

Outside of Netflix, Amazon, and other streaming TV services, most companies monetize their online video through ads. Digital video advertising revenues reached $9.1 billion in 2016 and will likely become the fastest-growing online advertising segment in the next few years (source: iab).

Video ads generally come in three styles – in-stream, overlay, and companion – and an ad campaign usually mixes all three.

In-stream ads appear in the player before the content (pre-roll), interrupting the content (mid-roll), or following the content (post-roll). The skippable type presents a skip button after a short delay, and the non-skippable type forces the viewer to play the entire ad before showing the desired content. Overlay ads pop up on top of the video, obstructing part of the content. Companion ads run outside the video player but close to it, immediately above, next to, or below the video.

Video advertisers mostly use VAST, VPAID, and related standards, except on YouTube, which uses Google's AdWords and discourages third-party ad networks (it allows some VAST but no VPAID).

Keep in mind that the ads you see while watching videos differ from the marketing videos you'll find on web pages, Facebook feeds, or mobile apps. In the former, the video is the main content, supported or sponsored by ads. In the latter, the page, feed, or app is the main content, supported by the products or services marketed in the videos.

Video analytics

Online video technologies have a huge advantage over broadcast and cable TV: the ability to measure viewership and engagement numbers precisely.

Video metrics can range from simply counting the number of views or viewers to tracking every interaction the user has with the video, including rewinding or skipping. You can categorize viewers by location, type of device, demographics, and other parameters. This ability allows publishers and advertisers to fine-tune content to maximize impact and make better decisions about video and marketing strategies.

Integration

Online video makes up only one segment of overall brand and content strategy for enterprises. Well-implemented integration can maximize the return on investment both for video and other types of content.

Integration can range from simply embedding a video on a page to integrating the video content management with your own content management system, so you can organize and even automate the video workflow.

Video platforms usually distinguish between the video player API (application programming interface), which supports adding video players to the content and customizing their interfaces and behaviors, and the publisher API, which can control the production workflow.

When choosing an online video platform, you should determine your needs according to the type of online video you decide to provide, then find the platform APIs to support those needs.

Capacity planning

When calculating the cost of online video publishing, the major metrics to watch are number of views per month, total size of recorded video on demand, and number of simultaneous live streams.

Recorded video is usually measured by storage size in gigabytes or video duration in hours. One GB roughly equals one hour of high-definition video. You can measure traffic either by view count (how many times the video was watched) or duration (how many hours have been watched). You can estimate duration by multiplying the average video length with the average number of views per video.

Most online video platforms define their pricing plans and product limits using a combination of these metrics. If using an in-house video solution, the storage size and bandwidth usage will also directly determine the operating cost of the solution.

To stay on top of these costs, you have to plan your usage according to your online video strategy, estimate these key metrics, take into account your selected platform, and determine their combined costs.

Another key aspect of video strategy is live and on-demand video.

Live and on demand

Online video generally fits in two categories, live and on demand. Viewers watch live content as it happens – competitions, concerts, news, and such. On-demand video is produced beforehand and watched at the viewer's convenience – movies, TV shows, or YouTube clips. Other types of video content, such as conferencing, monitoring, and CCTV, usually don't fall into the online video category.

Video producers can record live content and publish it later as on-demand video. In some cases, publishers allow viewers to pause live video or move backwards in it, achieving a DVR-like effect and further blurring the distinction between live and on demand.

Live video production workflow

Here's a simplified overview of the production workflow for live video:

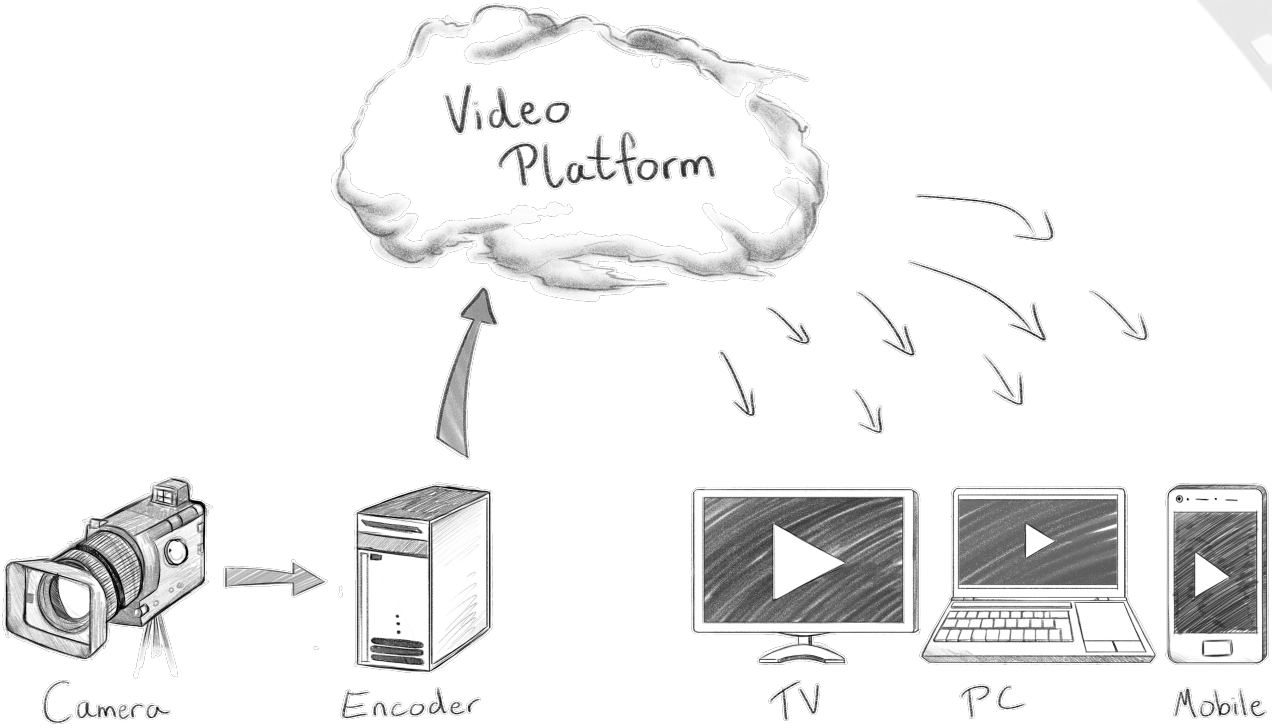

It starts with capturing the video and sound at the event. This can be as simple as a webcam or mobile phone or as sophisticated as multiple high-quality cameras and microphones, supplementary video sources such as presentation displays or B-rolls, and video and audio mixers to combine all the sources into one raw video stream.

Raw HD video requires about one gigabit per second (Gbps) of bandwidth, making raw video too big to stream directly to viewers (according to Statista, the high-speed Internet connection in the U.S. averages around 18.75). Therefore you must send it to an encoder to compress it to a manageable size while retaining its quality. Quality preferences and bandwidth requirements can vary, so encoders are highly configurable. That means you can tweak them for the optimal quality and bandwidth balance you need.

A common setup is to use one or more cameras and microphones with a single computer that acts as a digital mixer and encoder. Popular encoding solutions such as Telestream Wirecast, Adobe Media Encoder, and OBS Studio all provide features that support such use.

The encoder sends the content to the streaming service, either a single server within your network or an online video platform (OVP), a cloud-based service provided by a third party. The streaming service may have to present the content to hundreds or millions of viewers, so it must be scalable. For example, the 2017 U.S. presidential inauguration drew 4.6 million viewers, using 8.7 Tbps of bandwidth at peak times (source: Akamai). The streaming service can also record live video, collect viewer analytics, deliver ads, implement pay-per-view models, and many other tasks related to publishing video content.

The last component in the workflow is the video player. This may be an app on a mobile phone, a computer with a video player connected to a monitor, a media center connected to a TV, or a sophisticated digital projection system in a theater. The player component connects to the streaming service, receives the content, and presents it to the viewers, rather than allowing viewers to download the content to their own machines.

Video on demand (VOD) workflow

The VOD workflow resembles the live workflow but consists of two parts, production and publishing.

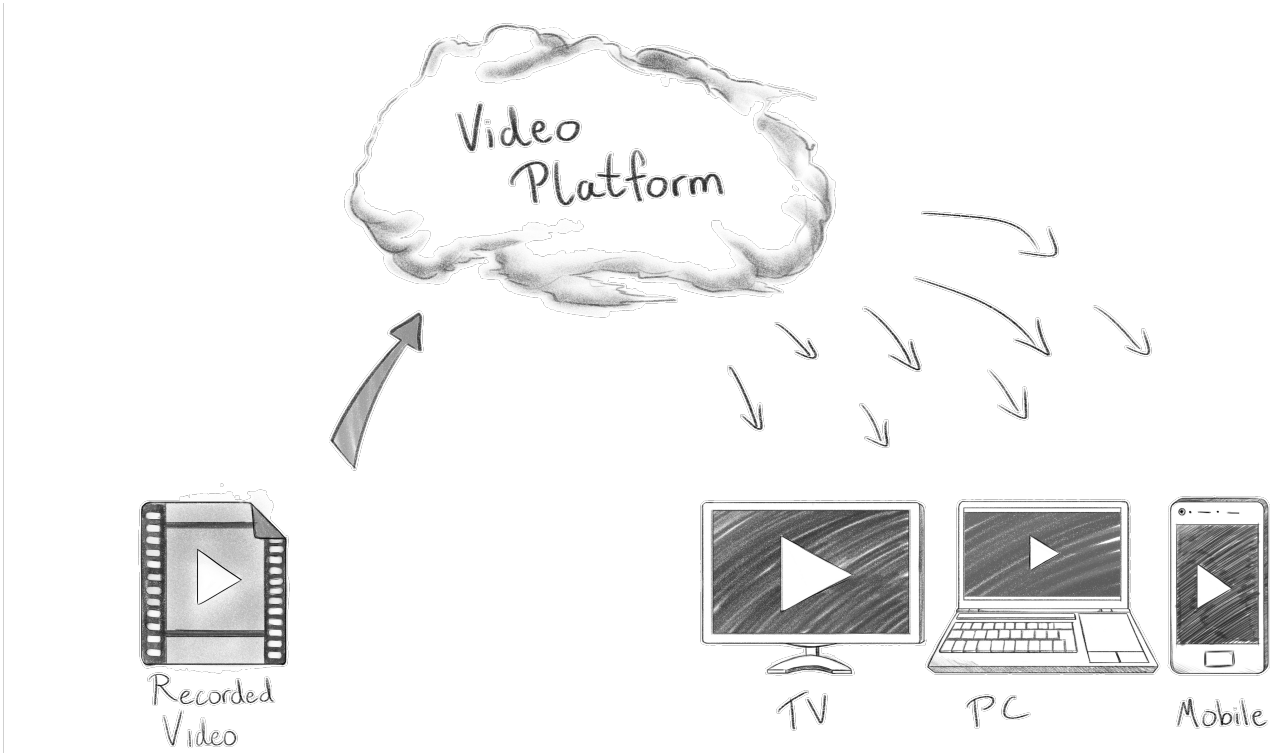

Video production encompasses video and sound capture, mixing, editing, and encoding into suitable output formats. Because it doesn't happen in real time, production can involve complex editing and post-processing, which can take more time than recording itself. Once ready, the results go to the streaming service, where the viewers can connect and watch at their convenience.

In both live and video on demand, viewers stream content from the service instead of downloading content locally and watching it later. Some providers allow their customers to download content for offline use, like Netflix, but most use streaming.

Adaptive streaming

While publishers want to produce and deliver content of the highest possible quality, they have to work within the constraints of connection quality, available bandwidth, and video-player performance.

To achieve the highest quality under these constraints, publishers use a technique called adaptive streaming, where the streaming service provides the content in multiple versions of varying quality, like high definition (HD), standard definition (SD), and low resolution. That way the video player can switch between them based on network and device conditions. For example, if you're on a train watching a movie on your tablet and the train enters a tunnel, your player will switch to a lower-quality stream to improve speed. Once the train exits the tunnel, the player switches back to the better-quality stream.

Codecs, formats, and protocols

If you've never worked with video before, you'll need to learn a few technical terms.

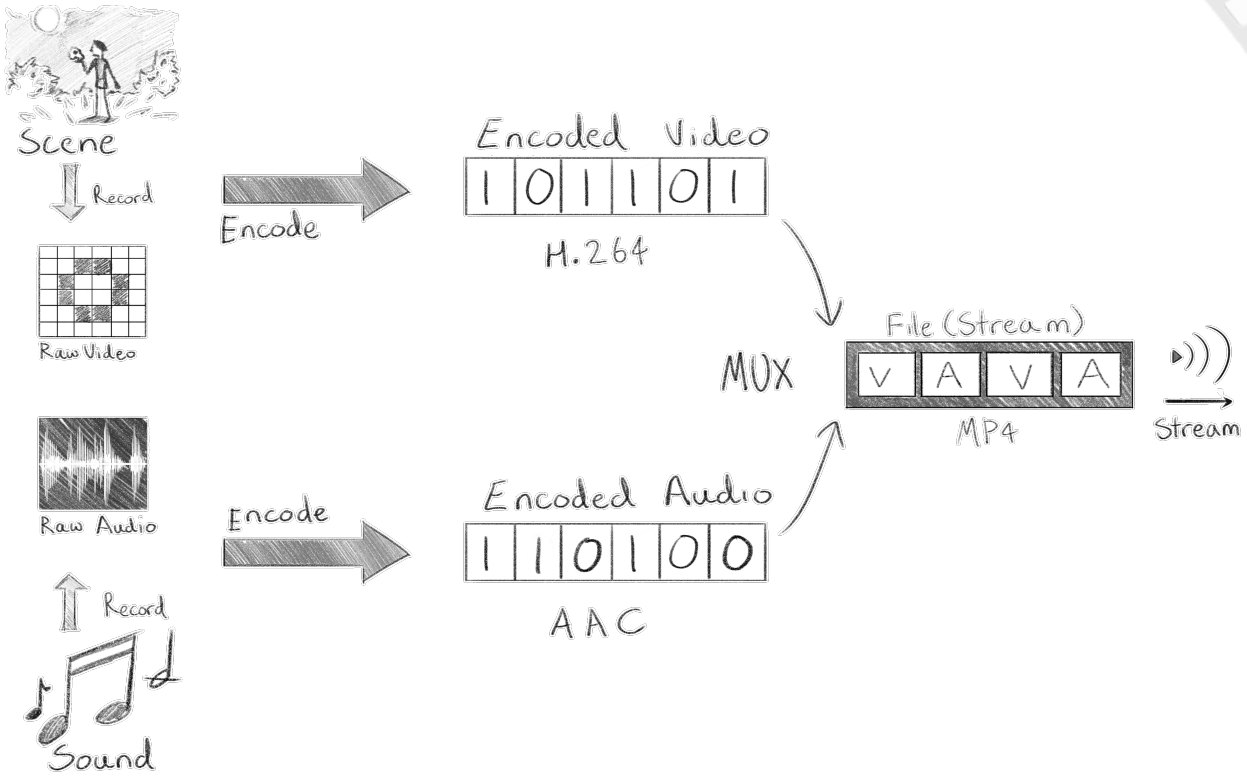

A codec is a video compression technique that converts the huge raw video into a file that will stream on an average internet connection. Currently the industry standard, H.264, can reduce a video's bandwidth by 100 to 300 times. Several next-generation video codecs have also gained ground in the industry, including HEVC/H.265 and VP9. The most popular audio codec today, AAC, can reduce audio bandwidth 10 to 20 times.

A format (also known as container format or bitstream) packs video, audio, and other information into a single file. In a given file, you may find a video stream, an English audio track, a French audio track, several subtitle tracks, and so on. Try not to confuse stream with audio channels; one audio stream may have several audio channels, like two for stereo and five or more for surround sound.

Although MP4 remains the most widely used format for online video files, streaming often uses the more suitable MPEG-TS, and you'll find YouTube videos in Google's WEBM format.

A protocol (sometimes confusingly referred to as a streaming format) transports content from the encoder to the streaming service and then to the user. The protocols for streaming to and from the service usually aren't the same.

The industry standard for streaming from encoder to service is RTMP, originally developed by Macromedia (now Adobe) for Flash. Although Flash use has declined, RTMP remains the standard; online video providers (OVPs) such as Bitmovin, Ustream (now IBM Cloud Video), LiveStream, YouTube, Facebook, and Twitch use it almost exclusively.

Streaming from the service to user devices usually involves Apple HLS, with MPEG-DASH becoming more widespread. Both protocols allow seeking (for video on demand and DVR stream) and adaptive bitrate functionality.

Latency

A low-quality internet connection full of glitches and freezes can make viewing video frustrating, even if the user has a "high-speed connection," which most providers measure purely by download speed. To avoid these problems, video players usually employ buffers to smooth over the irregularities. The streaming service also uses buffers to reduce upload variances, temporarily store content while transcoding into various output formats, and so on.

For live video, these buffers affect the latency, the time delay between the recording of the video and its display on a video player. Many online video platforms consider a 60-second latency normal but provide low (10 seconds) or ultra-low (a few seconds) latency for videos requiring more immediacy and interactivity.

However, for real-time purposes such as video-conferencing, most users consider a two-second latency unacceptable. As a result, such video often employs different technical solutions, such as UDP-based transfer protocols and lower compression and quality levels in exchange for quicker error recovery.

How to choose

When designing and implementing your video strategy, the choices you make in regards to the video platform and workflow will directly affect the user experience, user engagement, publishing cost, and ultimate success of your video strategy.

With so many options available today, choosing the right one for you can be daunting. This white paper highlights the key issues involved in such choices, but no single resource can provide all the answers needed.

If you're faced with such a challenge, partnering with an industry expert can help you make more informed choices, maximizing potential and minimizing miscalculation. I've helped clients design and implement a variety of online video strategies as well as choose and integrate with video platform providers.

Contact me today and find out how I can help with your online video strategy.