


QUANTUM WHITEPAPER

4K, HDR, HFR: Calculating the Storage Impact in Media Workflows

Tom Coughlin, President
Coughlin Associates, Inc.
www.tomcoughlin.com



CONTENTS

Introduction	3
Inside the Camera	4
Calculating the Size of Digital Content	5
Digital Storage Technology Options for Modern Workflows	6
Display Technologies	9
Broadcast and Delivery	10
Distribution with Interoperable Master Format (IMF)	11
Technology Adoption for Various Video Markets	11
Summary	13
About the Author	14
Appendix	15

INTRODUCTION

Technology trends are driving an evolution of modern media workflows, including higher-resolution imagery, higher frame rates, higher dynamic range, wider color gamut, and simply more cameras per shoot. Their impact on the digital storage required to capture, process, distribute, and archive this richer digital video content is significant. Today, a single large Hollywood production may shoot 100 hours or more of footage. Productions embracing these advanced technologies generate more, requiring capacities like these:

- The David Fincher film, *Gone Girl*, generated about 500 hours of raw footage during its multi-camera 6K (4:1) Red Dragon production, the equivalent of about 261TB of storage.
- Back in 2009, James Cameron's *Avatar* was hailed as the first movie to break 1PB of total video production content. Five years later in 2014, *The Amazing Spider Man 2* required 2.4PB of storage to scan from film to 4K ¹.

Larger productions are coming and future workflows managing hundreds of petabytes of content are expected in the next ten years.

Leading the trend is higher-resolution content, with 4K becoming commonplace and early work beginning with 8K content. Another trend is higher frame rates with displays at 48 or 60 frames per second (fps) and 4K capture up to 120 fps or higher in specialty cameras designed for slow-motion capture. Cameras and display devices are sporting higher dynamic range and an expanded color space, providing content that looks even more lifelike.

Meanwhile, the TV and display world has moved to 4K, placing pressure on content delivery systems for higher-quality images, particularly with higher-dynamic-range content. Compression reduces the size of the original content so it is easier to distribute to consumers. The development of higher content compression technologies such as HEVC (requiring much more intense encoding processing) is a key enabler of delivering 4K content to consumers. Yet compression does nothing to relieve the pressure on production crews to capture as many high-quality pixels as possible.

With the rise of smaller professional cameras, including GoPro and similar devices, multi-camera projects are quite common. Capturing content from multiple cameras is a driving force for the creation of modern virtual reality content, and it also provides raw data that can be processed and used to create perspectives from anywhere within the field of view of a group of synchronized cameras focused on a scene (often called free viewpoint video).

All of these trends are driving the need for greater processing power, higher network and storage performance, and much more digital storage capacity to support the workflows that process this content.

In this whitepaper, we will explore the quality and size of original-captured video content for sports, film, episodic, and animation content. We will look at how this changes production, post-production, distribution, and display of finished work. We will provide a formula for calculating the size of frames with different resolution, dynamic range and

¹ *The Amazing Spiderman's 2.4 PB of 4k Scans*, <http://www.definitionmagazine.com/journal/2014/6/11/the-amazing-spidermans-24-petabytes-of-4k-scans.html>

color gamut, and frame rate so you will know how much digital storage is required for an hour of content. We will compare that calculation with actual camera recording formats. We will also talk about the storage devices and architectures that support various parts of the workflow to provide the required storage capacity, performance, and cost basis.

INSIDE THE CAMERA

Media professionals generally use the highest-quality cameras that they can afford for their content capture: the highest resolution, the best colors, and greatest dynamic range to support the rest of the video workflow. What is not captured by the camera cannot easily be added to content later. Directors may occasionally use less-expensive, lower-quality cameras to create useful artistic effects, to capture otherwise hard-to-get perspectives, or to provide inexpensive ways to create stereoscopic or free viewpoint video. *Generally, the captured content is the richest content in the entire workflow and therefore requires the highest storage capacity per hour of content.*

The richness of an image is determined by many characteristics.

Resolution

Resolution refers to the number of pixels per image. The greater the number of pixels, the finer the details available in the image and the more immersive the resulting content. The image resolution required for that sense of immersion is a function of the size of the display screen and the distance of the viewer from that screen. *In order to future-proof their content, many modern video professionals shoot their content at the highest resolution that they can afford.*

Computer-Generated Content

Like camera footage, animation and CGI content is also rendered at the highest resolution possible. As the price of processing comes down, some effects can be rendered in real time during the editing process. *Regardless of whether content is captured or created, it will always be at the highest quality possible.*

High Frame Rate

Most professional cameras today have frame rates of up to 120 fps for 4K, with some specialized cameras designed for slow-motion video capturing content at thousands of frames per second. Frame rate's impact on storage is so great that some cameras trade off the resolution (number of pixels per frame) with frame rate. *The additional storage capacity required for higher frame rates is linear. For example, doubling the frame rate from 60 fps to 120 fps will double the storage.*

High Dynamic Range

Professional video cameras generally have a High Dynamic Range (HDR) and thus can show great levels of detail across very dim and very bright areas of a frame. Cameras with the greatest dynamic range have the most sensitive light sensors or may combine multiple frames with different exposure settings of a scene to enable more details in the dark and light areas of the image. These high-end cameras have high-quality optics that are designed to capture as much light energy as possible, even from dark imaged regions. *Higher dynamic range can increase the captured content frame capacity in the camera by up to 100%. For instance, Red Camera's HDRx setting doubles the storage capacity per hour of shooting.*

Color Depth

The color depth or bit depth is the number of bits used to indicate the color of a given pixel in a video frame buffer, or the number of bits used for each color component of a single pixel. Color depth represents how finely levels of color can be expressed, or the color precision. Generally, the greater the bit depth of pixels, the more smoothly the color of one pixel can shade to the next pixel. *Greater bit depth also adds to the overall storage capacity required for each video frame. For example, increasing the bit depth from 8 bits to the 10-bit minimum standard for Ultra HD increases the capacity required by 25%.*

Color Gamut

The color depth is only one aspect of color representation. The other important aspect of color representation is how broad a range of colors is being expressed. This is called the color gamut. The color space recorded in a camera is usually a subset of all the colors (light frequencies) that are sensitive to the human eye (although some cameras have sensitivities at higher or lower light levels than those of the human eye).

In general, professional video cameras are sensitive to a very large color space. This is referred to as a Wide Color Gamut (WCG). The color encoding specification for the camera defines both color precision and the gamut by assigning a digital code value to a location in a color space. *Unlike color depth, which increases the number of bits per pixel, the ability of cameras to capture a wider color gamut does not inherently increase storage requirements.*

Compression

Many professional video cameras also offer RAW recording options that don't compress the content using logarithmic encoding. These RAW recordings are very different between manufacturers because of the different imaging sensors in the cameras. Modern workflows can handle encoded as well as RAW video images. RAW image output requires considerably greater capacity per hour of shooting than the logarithmic encoded content used in non-RAW formats. *The size of RAW images can be 2-3 times greater than logarithmic encoded images.*

CALCULATING THE SIZE OF DIGITAL CONTENT

A few simple equations can be used to calculate the overall digital storage requirements for an hour of content with a given level of resolution. Given the Ultra HD-1 spec at 4K and 10-bits/pixel, these are:

Equation 4.4.1. Frame size = Width x Height x Bits/pixel x Colors

$3,840 \text{ pixels} \times 2,160 \text{ pixels} \times 10 \text{ bits/pixel} \times 3 \text{ colors} = 248\text{Mb/frame} / 8 \text{ bits/byte} = 31.1\text{MB/frame}$

Equation 4.4.2. Data Rate = Frame size x Frames/second

$31.1\text{MB/frame} \times 60 \text{ fps} = 1,866\text{MB/sec}$

Equation 4.4.3. Data in 1 Hour = Data Rate x 3600 seconds

$1,866\text{MB/sec} \times 3600 \text{ sec} \times 1024\text{MB/GB} = 6,718\text{GB}$

Table 1 shows various representative characteristics for professional media standard formats used in content capture. Note that the actual data rate and storage capacity of working or archived content can depend upon additional parameters not discussed in detail here, such as specialized compression and color coding. A general feature of digital content is that, as the content resolution increases, the data rate and resulting capacity demand increases as well.

Table 1. Example Resolution, Data Rates and Storage Capacity Required for Professional Media Content (assumes no chroma subsampling)²

Format	Resolution (width x height)	Frame Rate (fps)	Data Rates (MB/s)	Storage Capacity (GB/Hour)
SDTV (NTSC), (8-bit)	720 x 480	~30	31	112
HDTV (1080p, 8-bit) RGB	1920 x 1080	24	149	537
UHD-1 4K (10-bit) RGB	3840 x 2160	60	1,866	6,718
UHD-2 8K (12-bit) RGB	7680 x 4320	120	17,916	64,497
Digital Cinema 2K (10-bit) YUV	2048 x 1080	24	199	717
Digital Cinema 4K (12-bit) YUV	4096 x 2160	48	1,910	6,880
Digital Cinema 8K (16 bit) YUV	8192 x 4320	120	25,480	91,729

Note that 8K content, whether UHD-2 or Digital Cinema format, produces significantly larger files than 4K or 2K (HD). Digital Cinema 8K generates 13 times larger capacity per hour than 4K and 127 times larger capacity than 2K. Despite the much greater requirements for 8K compared to 4K, there are parts of the video industry that are pioneering 8K workflows. As stated before, image capture will always take advantage of the most advanced technology to capture the best image possible.


Table 2 in the appendix shows examples of real 4K codecs being used in professional cameras. Because of different color depth, color subsampling, frame rate, compression, HDR, and many other factors, the actual 4K data rate and storage capacity vary considerably and aren't the same as in Table 1.

DIGITAL STORAGE TECHNOLOGY OPTIONS FOR MODERN WORKFLOWS

Today's digital media workflow has many requirements and constraints. High-performance storage is needed to support capturing and editing high-resolution content, and at the same time IT budgets aren't necessarily increasing at the same rate. That makes economical storage an important element in most video workflows. As a consequence of these often-conflicting needs, many different storage technologies are often used together, including flash memory, hard disk drives, and magnetic tape.

Figure 1 represents a Digital Entertainment Content Workflow showing various aspects of the role of digital storage in the creation, distribution, and archiving of digital entertainment content.

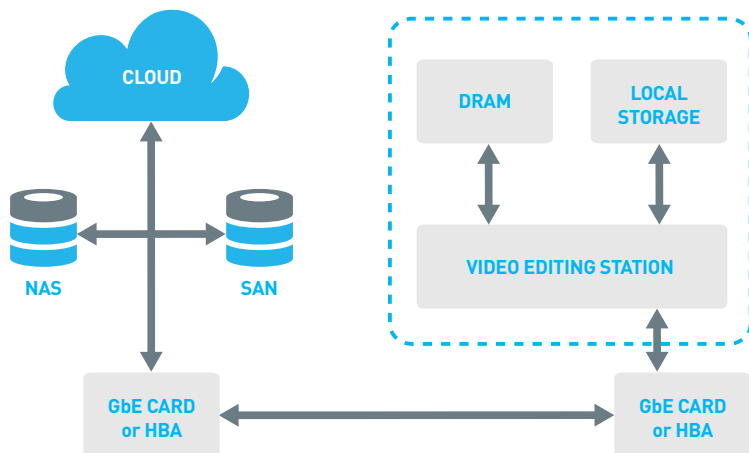
Figure 1: Digital Entertainment Content Workflow



Digital storage that supports older workflows with many streams of HD content will often choke on just a few streams of 4K content. The higher frame rate of slow-motion cameras and many 4K and 8K cameras can chew up storage even faster. In addition to higher resolution, HDR, the need to ingest many content delivery formats adds to the demand on storage systems that serve these content delivery systems.

At the same time, modern workflows can be very collaborative activities, making network storage important in local facilities. Collaboration over a distance requires Internet-accessible storage systems (either private or public cloud storage). Network storage and especially cloud storage are growing parts of the total storage in media and entertainment. Figure 2 illustrates a modern video editing station showing local direct attached, network as well as cloud storage.

Figure 2: Professional Non-Linear Editing System Schematic²



² 2015 Digital Storage in Media and Entertainment Report, Coughlin Associates, <http://www.tomcoughlin.com/techpapers.htm>

Solid-State Disks (SSDs)

In video workflows, flash memory—often SSDs—seems to be growing in popularity as a means of providing higher-performance access for work in progress. Most workflows still provide a working archive of content on HDDs that can be drawn up into the SSDs as needed to support a fast workflow. SSDs are also growing in use to provide fast content delivery to consumers, often backed by a HDD-based content library.

That said, SSDs are still very expensive. Otherwise, every facility would build their entire infrastructure from flash. Workflows have different requirements for data performance (bandwidth), access (latency), and cost (\$ per capacity unit).


Archive

Archiving content is an important element in the long-term preservation and monetization of that content. Archiving also serves to reduce aging content from primary storage to improve capacity utilization and increase performance. Today's digital archives, whether in the cloud or in a data center, generally use hard disk drives, magnetic tape, and sometimes optical discs to store data.

Object-Based Content Libraries

Object-based content libraries are the basis for most private and public cloud storage infrastructures. As Figure 3 shows, archiving in the cloud is expected to be the biggest application—in terms of total storage capacity—for cloud storage in the next few years.

Figure 3: Media and Entertainment Cloud Storage Capacity Projections ¹



Digital Tape and Vaulting

Magnetic tape is the least expensive digital archive media. The cost of raw storage on magnetic tape is less than a penny per GB. Furthermore, magnetic tape (like optical discs) does not use any energy when sitting in a library. Thus, both the capital expense (CapEx) and operating expense (OpEx) for magnetic tape archives is very low.

While tape is inexpensive, it has a longer latency (time to retrieve the first byte) than some other storage technologies. For this reason, a magnetic tape library is often coupled with some hard disk or flash storage where archived content can be cached when it is more frequently accessed. A sophisticated asset management system can help decide which data to copy to the HDD arrays for faster access.

As a media archive grows, it must ultimately move beyond removable external drives or small tape libraries to a secure digital archive. That archive may exist at the production house or at the content owner facilities, or it may be kept in Internet-accessible Cloud Storage.

Asset Management

Valuable video created in a workflow should also be managed by an asset management system that captures any analytic data gathered during the production and couples it to the content. Metadata on all aspects of the digital content makes content more searchable, helping content producers find what is needed, when it is needed. This helps provide insight when the content is delivered and helps monetize this content in the future. Archiving in particular benefits most from media asset management to help production teams find and use assets in the digital library.

Tiered Storage Management

To help manage content throughout its lifecycle, storage management software can analyze usage patterns and use workflow-specific policies to determine the optimal location for assets. Then, assets can be automatically migrated between storage types, ideally keeping the assets transparently accessible to users and applications regardless of location.

In a modern digital media workflow, flash memory, hard disk drives, and magnetic tape provide an optimal trade-off between performance and cost. As the price of flash decreases, the total storage capacity of flash memory should increase since many parts of the media and entertainment industry value the low latency and high performance that flash-based storage provides. Although the total mix of storage technologies—whether local or in the cloud—will change with time, the total storage demand and bandwidth requirements will increase with the rapid pace of media and entertainment technology development.

DISPLAY TECHNOLOGIES

In today's world, people watch content on many different size displays, from large theatre screens to home displays to tablets and smart phones. All these different types of displays require content that is tailored to them. Real-time transcoding of content to match the needs of a particular communication channel and different screen sizes provides important efficiencies, as these different content formats may not need to be stored long for use, but rather created more or less as needed.

As the TV and display world moves to 4K, greater storage capacity for content delivery systems will be required, particularly with higher dynamic range. Higher definition is generally most noticeable for very large screens, or where the viewer is likely to be very close to a screen. Because the eyes of computer and tablet users are generally close to their screens, some computers and tablets have 4K-resolution screens.

While resolution perception is a function of the size of the screen and the distance of the viewer, perceptions of color and high dynamic range seem to make a difference to viewers for all size screens and reasonable distances from a screen. Let's look at the way modern display devices provide colors and dynamic range.

Higher dynamic range provides a stronger contrast between the bright parts and the dark parts of the image. In an HDR display, bright white clouds can have texture like a real cloud, rather than looking like a washed out white patch. The peak brightness of an HDR television is about 1,000 nits while the peak brightness of most non-HDR TVs is about 100 nits. Thus, there is a ten-fold increase in the brightness of the highlights on the screen. As a result, HDR in video images shows more details in both bright and dark areas of the display, creating a better sense of reality and thus the illusion of being there. *HDR requires more storage capacity and thus requires more bandwidth for content distribution.*

The color gamut available for display depends upon the display device and the encoding of the color space. *Although color depth will impact the capacity requirement and bandwidth of the transmitted video, a wider color gamut generally will not.*

For these reasons, The UHD Alliance (made up of studios, consumer electronics manufacturers, content distributors, and technology companies) is incorporating HDR (SMPTE ST2084 EOTF) into the Ultra HD Premium standard for TVs with 4K resolution and an expanded color gamut (Rec 2020). TVs branded with this logo are expected to provide a much more immersive consumer experience than 4K TVs that don't have HDR and a wider color gamut (WCG).

BROADCAST AND DELIVERY

The development of higher content compression technologies such as HEVC is a key enabler of delivering 4K content to consumers. Compression reduces the size of the original content so it is easier to distribute to consumers. Without it, very few locations could enjoy 4K without a considerable investment in communication bandwidth. However, HEVC and other distribution compression technologies in use today are lossy, and if the content is compressed too much, the effective resolution of the resulting images suffers.

However, with up-rezzing technology available in today's 4K TV displays, 2K resolution content with higher dynamic range and wide color gamut can rival the image quality of non-HDR 4K. Netflix says that delivering an HDR picture requires about 20% more bits than the equivalent non-HDR resolution. Thus, while 4K is normally delivered by about 15 Mbps, 4K HDR requires 18 Mbps. Likewise, normal 2K is delivered by 5-6 Mbps, while 2K HDR requires 8 Mbps. Compared to the original content data rates in Tables 1 and 2, this content is obviously significantly compressed.

³ A nit is a unit of visible-light intensity. One nit is equivalent to one candela per square meter. One candela is approximately the amount of light emitted by a common tallow candle.

⁴ Netflix: *High Dynamic Range is 'more important' than 4K*, Sophia Curtis, *The Telegraph*, Jan. 12, 2015.

Some content creators are starting to capture and display 8K content for special broadcasts. There are plans to broadcast 8K TV in Japan by 2020. 8K distribution outside of sports events and theatres will likely remain limited until bandwidth increases to homes and through future wireless (5G) networks and compression technology becomes even more sophisticated than it is today. It remains to be seen whether 8K content will become popular in homes in the next decade. Most consumers don't have a large enough wall to mount a display that can effectively show 8K at comfortable viewing distances. *Perhaps 8K content will be used in VR environments where the viewer is very close to the display.*

DISTRIBUTION WITH INTEROPERABLE MASTER FORMAT (IMF)

Video media can be sent to people all over the world with many different cultures and languages. In addition to the technical developments that have enabled the amazing videos that are available today, there are developments that are making it faster to repurpose content for an international audience. An SMPTE media standard (SMPTE ST 2067-2) called Interoperable Master Format (IMF) for the file-based interchange of multi-version finished audio-visual works is an important element in this capability.

IMF allows multi-language versions with subtitles and closed captions, video inserts, and after-the-fact delivery of content with supplemental packages. Using IMF, versions of content can be tailored to the viewing audience. Companies such as Disney have used this capability to rapidly monetize content tailored to many different markets and cultures. *IMF eliminates the need for long-term storage of multiple complete versions of video for different markets since the versions can be generated as needed.*

TECHNOLOGY ADOPTION FOR VARIOUS VIDEO MARKETS

Cinema

Movie video production covers a wide variety of practitioners, from small-scale-budget independent productions to large studios with the latest state-of-the-art equipment. With movie producers using a variety of cameras with different resolutions, frame rates, color gamuts, and dynamic ranges, the amount of video generated is getting larger.

Some of the largest productions in recent years and in the near future will involve many PBs of storage. If feature films break into shooting 8K video with multiple cameras with high frame rates, high color depth, and high dynamic range, this number could rise to hundreds of PBs, perhaps approaching an Exabyte (EB) or 1,000PB of data generated in a single movie production.

Broadcast

Traditional broadcast captured content is somewhat constrained by the channel bandwidth available to distribute that content through the air, in cable, using satellite, or through the Internet. This has resulted in a somewhat slower adoption of 4K and higher resolution, particularly in local and even international news.

To some extent, this is due to the use of any available technology (even cell phones) to capture current events, but it is also due to a somewhat slower adoption of this technology compared to episodic TV (traditional and OTT) content as well as sports content. However, as 4K TVs become more standardized and common, traditional broadcasters will need to switch to higher-resolution content capture and workflow.

Episodic TV

Episodic TV content has always been shot at a higher resolution than regular television programs (besides sports and other special interest programming) in order to attract regular viewers. So naturally, with the introduction of 4K HEVC and other compression technologies that allow distribution of higher-resolution content through existing channels, combined with the rise of high-resolution original episodic content from OTT providers like Netflix and Amazon, episodic TV shows are embracing 4K UHD as well as HDR workflows.

For instance, Amazon's *Mozart in the Jungle* is shot with HDR 4K and released in high dynamic range and standard dynamic range (SDR) versions. The content is shot in 4K using Sony F55 cameras and has a full 4K post-production. The producers worked in RAW content through editorial and then rendered out 10-bit 4K DPX for review. This is done in SDR through approval, and then they take the non-color-corrected 10-bit DPX files, add the P3 color space, and trim for HDR. *This level of quality has historically only been seen on the highest-budget feature films. So while new video technologies are driving higher requirements, new storage, network, and compute technologies are bringing higher-quality video within reach of smaller-budget projects.*

Animation

Animation and special effects use a rendering process to generate video in computers and servers. High-performance computing, network, and storage are characteristics of rendering environments. As the movie industry moves to 4K workflows, it is driving the development of 4K and even higher-resolution rendering for special effects, which can also include HDR rendering if required.

Likewise, animation can be more lifelike with higher resolution, although this requires a lot more processing for long-form video and more lifelike content that looks less like a cartoon. Because of the cost of a render farm, many smaller production houses rent rather than buy these services, often using online rendering services.

Sports

While the rest of the broadcast and over-the-top (OTT) online TV industry is making the conversion to 4K video, many sports producers are starting to produce 8K content. In July 2015, Major League Baseball (MLB) recorded its first game ever in 8K, streaming the content to displays scattered around Yankee stadium. During the 2015 Wimbledon matches, attendees watched tennis in 8K on 105-inch displays on the roof of the Broadcast Center. During the FIFA Women's World Cup, matches were broadcast in 8K for live public viewings at two venues in Japan.

In addition to the higher resolution of 8K content, sports producers are also using higher frame rates and greater dynamic range. During the 2015 Kentucky Derby, slow-motion cameras running up to 8X progressive high frame rates in 4K allowed viewing the races in great detail. MLB used HDR cameras during the 2015 All-Star game.

In addition to the higher resolution, frame rates, and high dynamic range, many sports events use a large number of cameras to capture the action. ESPN used 32 cameras on Monday Night Football in 2014, and it used 40 cameras for the 2015 NBA finals. Showtime and HBO used 18 cameras to cover a ring that is only 22 feet on a side for the Mayweather-Pacquiao fight.

With so much rich content to manage and with so many ways that this content is consumed—from large displays to smart phone screens—there are constantly increasing demands on digital workflows, in content capture, production, and distribution of content. These demands for the content increase the demand for storage capacity as well as bandwidth to move around that capacity.

SUMMARY

Many factors are driving the growth of digital storage and bandwidth requirements, both for modern media and entertainment workflows, and the content delivered to consumers. Cameras and rendering are increasingly creating 4K content with some markets, such as sports, pushing the envelope even further with 8K content.

In addition to higher-resolution modern content creation, technologies are providing greater color depth, increased color gamut, and higher dynamic range. These factors increase the storage capacity required per video frame and, combined with higher frame rates, increase the required data rates and storage capacity for an hour's worth of raw content.

In addition, the ease and lower cost of digital production are encouraging moviemakers and others to use more digital cameras and to shoot many more hours of content per camera than are eventually needed, leading to even higher-capacity workflow requirements.

Also driving these bandwidth and storage requirements are consumer demand for 4K content for their new 4K UHD TVs and the means to economically provide this content to them using media asset management combined with on-the-fly transcoding and the version flexibility that an IMF workload provides. Whether on a mobile device, on a home display, or in a UHD theatre, consumers want higher definition, wide color gamut, and high dynamic range video.

In order to meet these needs, media and entertainment companies are turning to networked storage and storage in the cloud to enable fast and low-cost worldwide collaborative workflows. The trend to put content in the cloud, particularly archived content (to enable future monetization), is driven by the changing economics of storage.

Storage in M&E workflows uses a blend of storage devices to provide high performance where it is needed and economical storage of vast and growing video-stored assets. For this reason, flash memory, hard disk drives, magnetic tape, and even some optical storage are used separately or together in various parts of the industry.

Digital storage is a key technology to enable the blockbusters of today and tomorrow. All the factors that make today's media compelling to consumers require smart management and application of multiple storage technologies. Today's M&E information technology managers are enabling these smart tools for modern artists. Continuing to develop that art will depend on continued dedication to more immersive technologies.

ABOUT THE AUTHOR



Thomas M. Coughlin, President, Coughlin Associates is a widely respected storage analyst and consultant. He has over 30 years in the data storage industry. Dr. Coughlin has many publications and six patents to his credit. Tom is also the author of *Digital Storage in Consumer Electronics: The Essential Guide*, published by Newnes Press. Tom publishes the *Digital Storage Technology Newsletter*, the *Digital Storage in Media and Entertainment Report*, and other reports.

Tom is active with SNIA, SMPTE, IEEE, and other professional organizations. He is Education Chair for the SNIA Solid State Storage Initiative. He is Chair of Future Directions for the IEEE Consumer Electronics Society as well as Director for IEEE Region 6. He is a long-standing member of the CE Society BoG and was Vice President of Operations for three years. Tom is the founder and organizer of the Annual Storage Visions Conference as well as the Creative Storage Conference. He is the general chairman of the annual Flash Memory Summit.

APPENDIX

Many factors are driving the growth of digital storage and bandwidth requirements, both for modern media and entertainment workflows, and the content delivered to consumers. Cameras and rendering are increasingly creating 4K content with some markets, such as sports, pushing the envelope even further with 8K content.

Table 2. *Some 4K and Beyond Camera Codecs*⁵

Camera/Codec	Resolution (width x height)	Frame Rate (fps)	Data Rates (Mbps)	Storage Capacity/Hour (GB/Hr)
Panasonic GH4 4K	4096 x 2160	24	100	45
Red 4K (6:1)	3840 x 2160	24	432	194
XAVC 4K	4096 x 2160	30	300	135
KineMINI 4K CinemaDNG	4096 x 2160	24	332	150
AVC-Ultra 4K	4096 x 2160	24	400	180
Canon 1DC MJPEG 4K	4096 x 2160	24	500	225
ProRes 422 4K	3840 x 2160	60	489	220
KineMAX 6K CinemaDNG	5760 x 3240	24	672	302
ProRes 422 HQ 4K	3840 x 2160	60	734	330
Sony F5/55 RAW 4K	4096 x 2160	25	1,000	450
ProRes 4444 4K	3840 x 2160	60	1,100	495
RED 6K WS (4:1)	6144 x 3160	24	1,160	522
ProRes 4444 XQ 4K	3840 x 2160	30	1,650	742
Sony F65 RAW SQ 4K	4096 x 2160	24	2,000	900
BlackMagic 4K PL RAW	4000 x 2160	24	2,120	954
RED 6K WS (4:1) with HDRx	6144 x 3160	24	2,320	1,044
Canon RAW 4K (12-bit)	4096 x 2160	30	2,664	1,199
Phantom Flex 4K RAW	4096 x 2160	938	102,400	46,080

⁵ Some data from *4K and Beyond—Video Data Rates*, VashiVisuals, <http://vashivisuals.com/4k-beyond-video-data-rates/>

Coughlin Associates

Data Storage Consulting