

Understanding and Using High-Definition Video

Introduction

This document is written for digital media professionals interested in moving to high-definition production, post-production, and distribution. This document focuses on the powerful Adobe® desktop applications, Adobe After Effects® and Adobe Premiere® Pro, but it also includes information about many other tools and products as they integrate into this workflow.

The two goals of this document are to help you understand what's involved in making the transition to authoring and distributing high-definition content and to help you get the best results out of that transition.

TABLE OF CONTENTS

1	Introduction
2	What is high definition?
4	How HD is used today
6	HD in production
7	HD tape formats
8	HD storage
9	Using RAID systems for HD storage
10	HD and connectivity
11	Capturing HD content
12	Adobe Premiere Pro-compatible HD solutions
13	Post-production for HD content
15	SD color space versus HD color space
15	Deciding on the distribution media for HD content
17	Distribution formats for HD content
19	Display devices for HD
20	Computer playback configurations
21	Playing HD audio
22	Conclusion

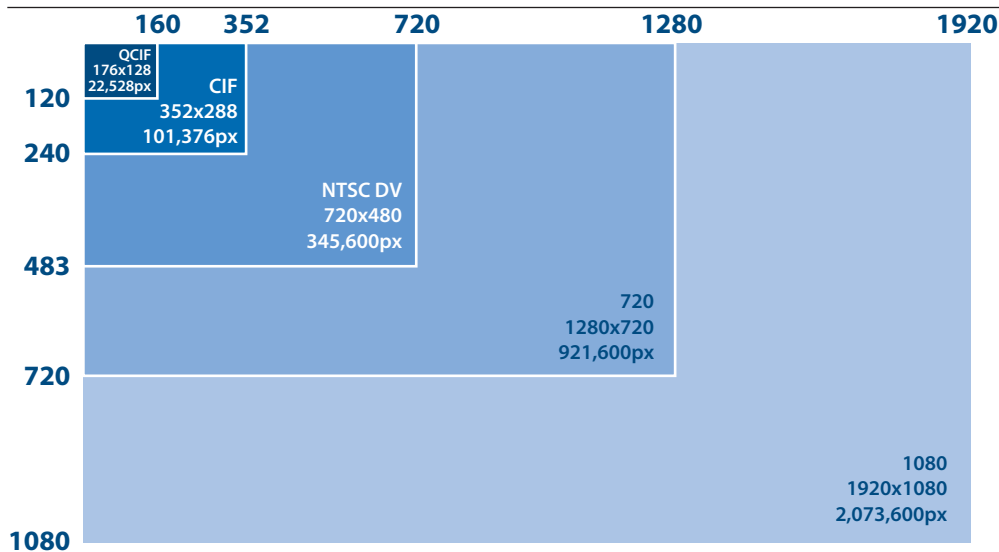
What is high definition?

After years of anticipation, high definition (HD) has finally, if not stealthily, gone mainstream in professional video. Over-the-air HD is rapidly building an audience, and HD DVD development is proceeding apace. People are also using HD during production and post-production in a wide variety of ways. For example, the pioneering DVD company, Criterion, uses HD on D5 tape for digital intermediates from their telecine captures. Suddenly, companies are finding that they have HD content to distribute and are looking for ways to distribute it.

The myth is that HD is very expensive and difficult to author. The reality is that HD has become easier and cheaper than standard definition (SD) was only a few years ago. Today, you can build a complete entry-level HD editing system, including storage, software, and camera, for under US \$10,000. Higher-end and uncompressed HD editing systems can still be much more expensive, but they are also declining in price as more options enter the market.

Defining HD

High definition simply means “more than standard definition.” The highest resolution SD format is PAL, with 576 lines. Thus, almost any video with a frame size greater than 576 lines tall is a type of HD. HD video is generally either 1920 x 1080 or 1280 x 720, with a 16:9 aspect ratio. However, there are many situations where other sizes may be appropriate, such as when you plan to play back the video on a computer.



Types of HD

Type	Size	Frames per Second
720 24p	1280 x 720	23.976 fps progressive
720 25p	1280 x 720	25 fps progressive
720 30p	1280 x 720	29.97 fps progressive
720 50p	1280 x 720	50 fps progressive
720 60p	1280 x 720	59.94 fps progressive
1080 24p	1920 x 1080	23.976 fps progressive
1080 25p	1920 x 1080	25 fps progressive
1080 30p	1920 x 1080	29.97 fps progressive
1080 50i	1920 x 1080	50 fields per second/25 fps interlaced
1080 60i	1920 x 1080	59.94 fields per second/29.97 fps interlaced

HD terminology

There is a lot more variability within HD than within SD. The terminology can be confusing, and it changes depending on the source used. This document uses the following format when discussing HD content: <frame height> <frame/field rate> <interlaced/progressive>. For example, 720 60p means a frame that is 1280 x 720 pixels and 60 progressive frames per second. Conversely, 1080 60i means a frame that is 1920 x 1080 pixels and 60 interlaced fields per second. Of course, for interlaced formats, each frame is made of two fields, so the frame rate is half the field rate.

In NTSC formats, the frame/field rate is actually 0.1% lower than listed, so 24p is really 23.976 frames per second, and 60i is really 59.94 fields per second. PAL formats use the listed frame rate—25p is really 25 frames per second.

Often times, 720 60p is called 720p and 1080 60i is called 1080i. However, those labels are ambiguous as to frame rate. 720 can run at 24p, and 1080 can run at 50i in its European version.

Video formats and sizes

This chart illustrates the relative sizes of different video formats' frames, with a count for the number of total pixels. These aren't aspect ratio corrected. For the actual shapes of video frames in different formats, see the next image.

The most common types of HD

Note that the 720 sizes don't have interlaced options, and the 1080 sizes don't support progressive at the highest rates.

NTSC normally reduces the frame rate by 0.1%; thus, for 30 frames per second you would get 29.97 frames per second to match the frame rate of NTSC color broadcasts. While this reduction is optional in most HD products, most people use the lower rates.

Frame Size Title	Width	Height	Area	Aspect Width	Aspect Height	Corrected Width	Corrected Height
NTSC VCD	352 px	240 px	84480 px	4	3	336 px	252 px
PAL VCD	352 px	288 px	101376 px	4	3	368 px	276 px
NTSC SD 4:3	720 px	480 px	345600 px	4	3	679 px	509 px
NTSC SD 16:9	720 px	480 px	345600 px	16	9	784 px	441 px
PAL SD 4:3	720 px	576 px	414720 px	4	3	744 px	558 px
PAL SD 16:9	720 px	576 px	414720 px	16	9	859 px	483 px
1280 x 720	1280 px	720 px	921600 px	16	9	1280 px	720 px
1920 x 1080	1920 px	1080 px	2073600 px	16	9	1920 px	1080 px

Relative dimensions of the frames for aspect-corrected video

Each dimension is aspect-ratio corrected with the correct number of pixels for the frame, so 720 x 480 16:9 is drawn wider and shorter than 720 x 480 4:3.

A (very) short history of HD

HD video for consumers has been around for about 20 years, even though it has only recently developed a sizable audience. The history of broadcast HD is fascinating and has been ably told by the New York Times' Joel Brinkley in his book *Defining Vision: How Broadcasters Lured the Government into Inciting a Revolution in Television*.

The original HD-broadcasting standard for consumers was Muse, created by Japan's NHK. Muse was an analog HD system intended for satellite broadcast. While it produced very good image quality, market realities kept Muse from getting much traction.

At the time, there was a drive in the U.S. Congress to reallocate some unused UHF channels from broadcasters to emergency communications and other uses. As an excuse to keep the channels, the National Association of Broadcasters (NAB) proclaimed it needed the channels for HD broadcasts in the future. This proclamation succeeded in the short term, but it committed the NAB to broadcasting in HD and (in theory) eventually to giving up analog channels for other uses.

The Federal Communications Commission (FCC) set up the Advanced Television Systems Committee (ATSC) to define the specification for HD broadcasts. The process took years more than expected, but the standard was finally created, and broadcasters began HD broadcasts. In 2004, most prime time shows are broadcast in HD, and both cable and satellite systems are introducing HD channels.

Even today, only a minority of consumers in the U.S., and a much smaller minority in other industrialized nations, actually have access to HDTV systems. But HD capable displays are selling very well now, even if many displays are only used to view standard definition content, like DVD.

The next big advancement for HD is likely to be HD DVD for which the standard is currently in development.

Moore's Law and HD

Intel co-founder Gordon Moore made an audacious prediction about computing decades ago: the number of transistors on an integrated circuit doubles at regular intervals.

The modern interpretation of Moore's Law is that computing power, at a given price point, doubles about every 18 months. Consequently, a new U.S. \$3000 editing system is twice as fast as a U.S. \$3000 system from a year and a half ago. A system three years ago was a quarter of the speed, and a system four and half years ago was one-eighth of the speed. Similar predictions hold true for other measurements, such as the amount of RAM available in a computer, and the speed and size of hard drives.

A pixel size of 1920 x 1080 60i is only 6.5 times greater than the NTSC SD standard of 720 x 480 60i. The idea of a computer having to do more than 6.5 times more work may seem daunting, but think about Moore's Law—if current trends continue, we can expect that type of improvement in about four years!

And Moore's Law is just a measurement of performance. In the last four years, we've seen great gains in terms of usability of video and audio-authoring tools.

How HD is used today

Currently, HD is used in a variety of ways including delivery to users. HD is also becoming popular as a way of distributing intermediates and proofs in the film industry, and it is increasingly used for video-editing installations, kiosks, and digital signage.

HD and the ATSC

HD for most consumers means over-the-air ATSC content. ATSC supports 480, 720 and 1080 at all of the NTSC frame rates. It uses MPEG-2 for video and Dolby Digital (AC-3) for audio; the same codecs used on DVD.

The experience of watching HD at home is still not as refined as it is for SD. For example, the amount of available content is still relatively small, and HD video recorders and personal video recorders (PVRs) are much more expensive than their SD counterparts. But vendors view HD as a growth area and are rapidly enhancing products for use with HD. In some ways, HD is easier and cheaper to work with than analog SD because you can avoid analog-to-digital conversion or compression.

While much of the rest of the world is moving toward digital broadcasting, only the U.S. and Japan have seen the broad adoption of HD as a broadcast technology. Developments in HD lag somewhat behind digital broadcasting, but some areas of Europe have begun HD broadcasting; Japan has announced it, and it should continue to be adopted worldwide.

HD and digital broadcast systems (DBS)

Since the ATSC began, the U.S. market has become dominated by cable and satellite, with only a small minority still watching over-the-air television. Analog cable is quickly going away and is being replaced by digital cable and satellite (collectively DBS). An odd twist in the digital broadcast revolution is that many viewers are now getting less than SD at home. Typical digital broadcasts are now 352 x 480 (often called half D1) instead of the full-resolution 720 x 480, and often contain very noticeable compression artifacts.

Many DBS providers are adopting HD with vigor. Their HD offerings are quite different than those for SD. For example, there are only a few dozen channels because HD bandwidth is much higher. They also look better—the HD channels are getting significantly more pixels and more bits per pixel, which results in a dramatic leap in quality.

Most DBS vendors use ATSC data to create their HD broadcasts. However, Cablevision's satellite service will use MPEG-4. Because new equipment (a set-top box) is required for HD, and because bandwidth is so high, there is interest in the possible gains from a more modern compression system than ATSC's MPEG-2. The two most discussed compression systems are MPEG-4 AVC and Microsoft's VC-9.

HD in film production and the post industry

Film has been moving to digital post production for a while. Originally, people used digital video for a few special effects shots; now, digital color correction and other techniques are becoming standard. And the new digital cinema cameras provide excellent quality. These cameras go beyond the video-oriented HD formats into the RGB color space without subsampling and with more than 10 bits of precision.

HD in post production often uses the 2K and 4K standards—respectively 2048 and 4096 pixels wide. Most people feel that 2K is more than good enough for typical productions. An increasing number of companies are using 1080 24p HD video cameras and equipment.

There is hope at the low end of the market as well. JVC's HDV-based cameras provide great results at a miniscule price. Similar cameras with higher resolution and native support for 1080 24p are coming. These new cameras will achieve better image quality than 16mm film, will cost less, and will be available and from a variety of vendors.

In both cases, it's easier to go back to film with a film recorder from HD 24p content than from any type of NTSC content. The 6.5 times as many pixels makes a huge difference, and authoring in film's native 24p frame rate makes motion much smoother.

In the post industry, HD video has proven to be an effective method for distributing intermediates, proofs, and clips for review. A downloadable HD file looks substantially better and can be available much sooner than a Beta SP tape sent by mail or even overnight express.

HD and digital cinema

After films can be digitally made, why go back to film? Digital projectors are rapidly dropping in price and improving in quality. Film projectors are expensive as is developing and distributing film prints. In addition, a digital print doesn't lose quality each time it's played. While it's not economically necessary for theater owners to make a complete transition to digital projection today, it probably will be necessary within this decade.

The definition of digital cinema is still evolving. The quality requirements are much higher than they are for consumer HD, and Hollywood likes it this way—they want the theatrical experience to have an advantage over the home experience that goes beyond popcorn. Expect whatever standard is adopted for theatrical HD projection to be well beyond the standard that will be affordable for the mass market.

The good news is that it is already possible for digital projection to equal or exceed the quality of film projection. It's just a question of getting those projectors cheap enough to put in the theaters.

Even though this point hasn't been reached yet, many theaters are putting in video projectors for showing pre-movie loops, slide shows, and other content. These projectors are typically cheaper 1280 x 1024 projectors. Although the resolution is far from 1920 x 1080, the most viewers don't notice the difference between this resolution and film resolution.

One interesting possibility, after digital cinema is standard, is a transition to higher frame rates. For many decades, 24p has been the film standard because updating cameras and projectors is very expensive, as is the cost in film stock. But with digital projection, 60p isn't that much more expensive to create, distribute, or project than 24p. And 60p can provide a vivid immediacy that is currently impossible with film.

HD and computer playback

Modern personal computers are very viable HD playback systems. Mainstream computers of the last year or two can handle 720 24p, and today's fastest computers can display 1080 frame sizes as well. One benefit of computer playback is that the frame size of the movie can precisely match that of the display, resulting in a crystal-clear, pixel accurate image.

Microsoft has been pioneering this HD-playback market. They are working with content companies to create a line of two-disc sets: a conventional DVD version of the movie and a DVD-ROM that contains an HD Windows Media Series 9 (WM9) version of the movie. The first title releases had some problems with playback performance, Digital Rights Management (DRM), and software players, but content vendors have clearly learned from those problems and the second generation of titles provided a more seamless experience. The technology is not exclusive to Microsoft; any content creator can make similar discs.

HD and consumer electronics

Currently, we lack mass-market consumer devices for storing HD. The only options available are the new line of HD PVRs, which are not intended for long-term storage, and Digital-VHS (DVHS) video tape recorders (VTRs), which put MPEG-2 on a VHS-sized tape. The DVHS technology works well, is reasonably easy to edit, and can produce good quality video. Its market share isn't high; although, this may change as tapes and decks continue to drop in price. There are a few dozen prerecorded movies in the DVHS format as well—largely action titles.

The goal is an HD DVD format, and it's coming. However, there are at least four different competing technologies in development. The leading proponents are the DVD Forum and Blu-Ray. The DVD Forum's efforts are the furthest along. They've already announced a plan for video codecs supporting MPEG-2, MPEG-4 AVC, and Microsoft's VC-9.

A writable format is in development for the DVD-Forum format, and it should become a viable target for one-off media in a product-introduction cycle that is much faster than DVD-Video was.

HD in production

Options for HD production have exploded in the last few years. Like SD, there is a broad range of products with wildly different price points and features. The HD experience is functionally very similar to SD, except for the additional bits and pixels. Equivalent cameras, monitors, and other workflows are available for SD and HD, but they have somewhat higher price points for HD. Even the formats are similar with derivatives of DVCAM and D5 as the dominant high-end production formats.

HD cameras

There are a wide variety of HD cameras, which range from 720 HDV models for U.S. \$3000 to more than U.S. \$100,000 for uncompressed RGB cameras used in motion picture production. Cameras range vastly in price, and they support all of the production HD tape formats. The high-end cameras use hard drive arrays instead of tape and store high-precision RGB data for film production. Because these cameras are changing so rapidly, this document doesn't mention any particular models.

HD techniques and issues

The main difference between HD and SD is the significant increase in the number of pixels. Experienced SD video professionals are often startled by the details HD reveals that aren't normally visible in SD video. These professionals often need to tune sets, makeup, and framing to make them work well in HD. The level of detail used in film is a good starting benchmark for tuning.

The targeted 16:9 aspect ratio of HD is also different. The combination of 16:9 and HD can be a boon for sports because much less panning is required to see the details. Well-shot HD hockey and basketball can be a true revelation—HD sports are one of the major drivers for consumers upgrading to HD.

HD tape formats

While there are some live broadcasts of HD, most HD is edited nonlinearly on personal computers and workstations, similarly to SD.

Currently, there are a variety of digital tape formats in use for professional HD production, but the industry seems to lean toward the Sony HDCAM and Panasonic D5 formats. All formats use the existing physical tape formats originally designed to record SD video but with new compressed bitstream techniques to store the additional data HD requires. Fortunately, there isn't a significant archive of analog HD tape content.

Sony HDCAM format

Sony's HDCAM supports 1080 resolutions at frame rates of 24p, 25p, 50i, and 60i. HDCAM internally stores the video at 1440 x 1080, a 33 percent reduction horizontally from 1920. It also uses unique 3:1:1 color sampling. HDCAM is 4:4:1 compressed and has only half of the color detail of other HD formats. It is internally 8-bit but supports 10-bit input and output and four 20-bit audio channels.

HDCAM is a time-tested format that produces visually impressive results—the compression issues have not proved to be much of a practical limitation. And, HDCAM has excellent price performance. Sony is working on higher-end HD video formats that won't have these compression issues.

Panasonic D5-HD format

Panasonic's D5-HD format uses the D5 tape shell. In addition to the HDCAM formats, D5-HD can handle 720 60p, 1080 24p, 1080 60i, and even 1080 30p. Higher-end decks also handle 1080 24p, which is the standard for HD in Hollywood. Note that the lower-end decks can't play back the 24p format—a source of confusion and frustration. D5-HD supports eight channels of audio.

Panasonic DVCPRO-HD/DV100 format

Panasonic's other HD format is DVCPRO HD (sometimes called D7-HD or DV100), and it is based on the same tape shell used for DVCAM and DVCPRO. It supports 720 60p with 1080i24p in development. It is one of the most lightly compressed formats; it uses 1.67:1 compression. It supports 10-bit per channel and eight channels of audio.

JVC HDV format

JVC uses HDV technology for their ground-breaking mid-range HD camera, the JY-HD10. HDV is an MPEG-2 transport stream (which includes a lot of error correction) on a miniDV tape. JVC video uses interframe compressed MPEG-2 at 19 Mbps for 720p or 25 Mbps for 1080i. Audio is encoded in 384 Kbps MPEG-1 Layer II stereo. Other vendors are expected to ship HDV-based solutions in the near future.

The interframe encoding allows HDV to achieve high quality at lower bit rates, which means there is more content per tape, but that content is more difficult to edit. Solutions like CineForm's Aspect HD are required to convert from HDV to an editable format.

Note that HDV uses the same bitstream as DVHS, but on a smaller tape.

HD storage

Getting uncompressed HD content onto a computer for editing or processing is a straightforward process—the hardest single element being storage. HD requires fast storage and a lot of it.

Compressed options are also getting onto the market, which makes things simpler and cheaper. These options are discussed in more detail in the following sections.

Big numbers to store

How big? How fast? Let's look at some numbers. These show DV, HDV, SD, and HD formats storage requirements in their common formats (the SD and HD numbers assume the video is the native uncompressed format). There are substantial differences among different types of HD. Going from 720 24p 8-bit to 1080 60i 10-bit is an increase of about 3.5 more pixels per second to process and store.

Width	Height	Frame Rate	Type	Mbps	GiB per hour
720	480	29.97	DV25	25	11
720	480	29.97	DV50	50	22
720	486	29.97	8-bit 4:2:2	168	74
720	486	29.97	10-bit 4:2:2	210	92
1280	720	29.97	HDV 720 30p	19.2	8
1280	720	23.976	8-bit 4:2:2	354	155
1280	720	23.976	10-bit 4:2:2	442	194
1280	720	25	8-bit 4:2:2	369	162
1280	720	25	10-bit 4:2:2	461	203
1280	720	29.97	8-bit 4:2:2	442	194
1280	720	29.97	10-bit 4:2:2	552	243
1280	720	50	8-bit 4:2:2	737	324
1280	720	50	10-bit 4:2:2	922	405
1280	720	59.94	8-bit 4:2:2	884	388
1280	720	59.94	10-bit 4:2:2	1105	486
1920	1080	23.976	8-bit 4:2:2	795	350
1920	1080	23.976	10-bit 4:2:2	994	437
1920	1080	25	8-bit 4:2:2	829	365
1920	1080	25	10-bit 4:2:2	1037	456
1920	1080	29.97	8-bit 4:2:2	994	437
1920	1080	29.97	10-bit 4:2:2	1243	546

Understanding storage capacity

Storage capacity can be quite difficult to understand. Here are the terms used in this document:

Mbps Megabits per second (1,000,000 bits per second). Many drives measure in Megabytes per second (MBps), which is eight times greater than Mbps.

GB Gigabyte. A GB is properly 1,000,000,000 bytes, although many applications assume it is really 2^{30} , which is about 7% larger and properly called a Gibibyte, or GiB. The same issue is true of MB, which is really 1,000,000, but is often considered to be 2^{20} , which is about 4% larger.

Storage capacity table

This table shows the ranges of file sizes and the required storage capacities.

Required bandwidth for storage

The first issue in storage is sustained bandwidth—can the storage system reliably transfer all of the bits of video onto the disc in real time? Many storage systems measure “peak” bandwidth, which is a best-case performance during a short burst of activity. For video it is essential to know the sustained, or guaranteed, bandwidth that can be maintained long term. When the transfer can’t keep up, frames are dropped from the capture, which means the footage stutters.

For anything other than an HDV capture, a single drive does not provide sufficient bandwidth for capture, and you will need a Redundant Array of Independent Disks (RAID) system (described in the next section).

Beyond capture, real-time effects require enough read performance to play back two streams off of the drive in real time.

Required capacity for storage

Once the drives are fast enough, there needs to be enough storage for the intended use. Acquisition is easier than editing. Because HD requires a large amount of storage, editing an offline version at standard definition and then capturing only the required sections in HD for finishing, is a great option.

You can now purchase a single 300-GB drive that doesn’t provide many minutes of storage for anything other than HDV, which again suggests a RAID solution is required for working with larger HD formats.

Using RAID systems for HD storage

The key to making storage for HD work is RAID. The basic idea is to tie multiple drives together into a single volume. Because this process makes multiple drives appear as one drive, storage is increased, and because it enables writing to multiple physical disks at the same time, throughput is increased. Duplicating data on more than one drive provides redundancy, which allows the volume to survive with no loss of data, even during the failure of one of its members. When a drive fails, the array must rebuild the redundancy by using this new drive. This rebuilding can often reduce the available throughput or can require the RAID array to be taken out of service during the rebuild. RAID storage vendors build different capabilities into their systems; consequently, it is important to know the level of redundancy and rebuild performance you require when you choose a RAID to store your video.

Types of RAID Systems

There are many different types of RAID that combine drives in different ways. The ones relevant to video work are:

RAID 0 uses two or more drives and splits the data between them. RAID 0 has no redundancy and offers no protection from data loss—if one drive fails, all of the data in the set is lost. However, all drives are used for storage, and performance is blazing fast. If the source is backed up, RAID 0’s performance and price can be worth its fragility. For editing with CineForm’s Aspect HD, a two-drive RAID 0 is sufficient for 720 30p. More drives are needed as speed requirements go up, and the more drives you add, the greater the chance of failure.

RAID 3 uses three or more drives combined and dedicates one to be used for redundancy. The volume can survive any one drive failing. RAID 3 doesn’t perform quite as well as RAID 0 on total throughput, but the fault tolerance is well worth it. Because the parity is dedicated to one drive in RAID 3, there is no effective limit to how you size your blocks for transfer.

RAID 5 uses three or more drives in a set, and the redundancy is shared across all of the drives by dedicating equal space on each drive to parity, or redundant, data. RAID 5 provides aggregate transfer rates that are better than RAID 3.

RAID 5+0 uses two or more RAID 5 sets combined into one RAID 0. Typically, each of the RAID 5 sets is on its own drive controller, allowing bandwidth to be split over two controllers. This division of bandwidth provides better performance than RAID 5 and still provides redundancy (one drive per RAID 5 set can fail without data loss). RAID 5+0 is the optimal mode for most HD use.

Comparing SCSI drives and ATA drives

The debate between SCSI drives and ATA drives (a class of IDE drives) for high-performance storage has been going on for over a decade. Historically, SCSI drives were faster than ATA drives, and ATA required more CPU power to run. Currently, Serial ATA has, for the most part, caught up with SCSI. Even uncompressed 1080 60i 10-bit video plays well off of a very fast ATA RAID. And because the cost of fast IDE drives is so much less than SCSI and cabling is easier with Serial ATA, the industry should continue the rapid transition to ATA, even in the HD world.

Either way, HD performance requires a stand-alone controller card (potentially several for throughput). The RAID controllers on many motherboards aren't up to the task of 1080 uncompressed video yet. For the purposes of reliably handling video, it is not recommended to use software RAID controllers. Hardware RAID controllers for video work often contain onboard cache's that are essential for efficient management of the data transfers.

HD and connectivity

RAID systems work well for storing files locally. But when the work needs to be split among multiple workstations, for example, between an Adobe Premiere Pro editor and an After Effects compositor, you need a system that can handle high bandwidth and high speed.

Transferring HD content via Ethernet

The simplest way to move files is by means of the Ethernet. Transferring between two or more Windows® computers is easy by means of Server Message Block (SMB), as the transfer is from Mac OS X version 10.3 or higher to a Windows system.

Ethernet speeds vary widely. For HD work, you should use a gigabit Ethernet, which can typically support several hundred Mbps of bandwidth with good cables and a fast switch.

Transferring HD content via SAN

When a workgroup wants to share the content for editing, not just transfer it, a storage area network (SAN) is the solution of choice. SANs allow multiple computers to access the same storage as if they were on local drives on each computer. You can directly capture HD content to stored media without having to transfer it after the capture. It costs more per workstation to connect, but storage isn't much more expensive, and the workflow advantages are well worth it.

Transferring HD content via sneakernet

The final solution for transport is the venerable "sneakernet." This solution involves simply transporting physical media, be it a FireWire drive, a tape, or a whole RAID array, from one computer to another. Overnight mail can take 300 GB across the country in 24 hours, but it would take two weeks on a T1 line.

Capturing HD content

How does the content get from tape to storage? That's the role of capture, which is ably handled in Adobe Premiere Pro by using one of several third-party products.

You can use one of two dominant protocols to capture HD content. The high-end formats use High Definition Serial Digital Interface (HDSDI), and HDV uses FireWire.

Capturing using the HDSDI Protocol

HDSDI is a straightforward extension of the Standard Definition Serial Digital Interface (SDSDI) protocol into HD. It supports all of the standard HD frame sizes and rates with 4:2:2 10-bit sampling. Delivery codecs, like Windows Media Video 9 (WMV9) and MPEG-2, are almost always 4:2:0. The horizontal color detail is cut in half when exporting to those formats. This color reduction isn't a big deal for most content, especially at HD resolutions.

The drawback when using HDSDI is that HDSDI is always uncompressed. Consequently, the bit rate of a relatively compressed format, like HDCAM or DV100, balloons. Ideally, all vendors would provide a way for you to capture their compressed bitstreams and natively edit them instead of forcing a conversion to uncompressed. This conversion increases storage requirements and requires another lossy generation of compression every time the signal goes back to tape.

In addition, HDSDI doesn't directly support device control; it requires traditional serial device control.

You can embed audio in the HDSDI bitstream; however, not all capture cards support embedded audio. It's more common to use an Audio Engineering Society/European Broadcast Union (AES/EBU) connection for transporting audio.

Capturing using AES/EBU

You can embed audio into an HDSDI bitstream or externally transport it over an AES/EBU connection. AES/EBU is an uncompressed audio bitstream that can transport the full quality of the information on the HD tape.

The disadvantage of using AES/EBU instead of embedded HDSDI audio is that using AES/EBU requires more cables.

Capturing with FireWire

IEEE-1394 (also known as FireWire and i.Link) took the video world by storm as a transport connection for the popular DV video format. However, FireWire can do more than DV. DV is 25 Mbps, but first generation FireWire has a theoretical peak of 400 Mbps. The 800 Mbps 1394b standard already appears in computers.

Currently, the only major HD format supporting native FireWire transport is HDV, which also maximizes at 25 Mbps. However, the compressed bitstreams of many of the professional formats could easily fit in the original 400 Mbps FireWire, and all formats could fit in 800 Mbps. Hopefully, vendors are rising to the challenge and providing a professional HD workflow with the ease of use of DV.

Frame size and HD capture

Some capture cards can support on-the-fly upsampling and downsampling to convert to and from 720 and 1080. On-the-fly sampling can be very useful when you're creating a project from mixed HD sources or a 720 project from 1080 sources. Other capture cards can only capture at the native frame size, which may be all that some people require. A few solutions support only 720 or 1080.

Working with progressive and interlaced content

For HD terminology, we indicate progressive with “p” and interlaced with “i”, putting the letter between the frame size and frame rate. There are various ways to express this distinction, for example 1080 60i. For progressive content, the frame rate indicates frames per second. For interlaced content, it indicates fields per second. So, 60i is 30 frames per second, and each frame has two fields.

The 720 formats are all progressive, but 1080 has a mixture of progressive and interlaced frame types. Computers and computer monitors are inherently progressive, but television broadcasting is based on interlaced techniques and standards. For computer playback, progressive offers faster decoding and better compression than interlaced and should be used if possible.

Capturing 8-bit versus 10-bit video

Video formats are either 8-bits per channel or 10-bits per channel. Video with 8-bits per channel has 256 steps from black to white, and 10-bit video has 1024 steps. The extra detail in 10-bit video can improve quality. Compressed delivery formats, like MPEG-2 and Windows Media, are 8 bit, so there isn't any need to capture in 10 bit if the video is going straight to compression.

Most HD authoring formats are 10 bit, but most HD delivery formats are 8 bit. Converting from one to another can be tricky. Some tools might truncate the least significant 2 bits. This process normally works fine, but in content with very smooth gradients, some banding might occur. Some capture cards can dither on-the-fly to 8 bit while capturing, which reduces storage requirements and provides a better dither than many software exporters.

Audio capture support

Audio for HD can be virtually any bit depth and sampling rate. Professional video editors frequently use the old standards of 48 kHz 16 bit. Higher bit depths, such as 20 bit and 24 bit, and higher sampling rates, such as 96 KHz, are common.

Audio for HD is normally mastered in a multichannel format, such as 5.1 surround sound (five speakers plus a subwoofer) or 7.1 surround sound (seven speakers plus a subwoofer). Most HD tape formats support at least four channels of audio and many support eight. When delivering HD content, AC-3 can support 5.1 48 KHz, and Windows Media Audio 9 Professional technology can deliver 5.1 and 7.1 at 96 kHz.

Adobe Premiere Pro compatible HD solutions

Adobe Premiere Pro currently offers two great HD solutions. Others are coming soon as well. As of NAB 2004, the following companies have introduced or committed to providing hardware, software, or both to support HD video editing with Adobe Premiere Pro:

- 1 Beyond (www.1beyond.com)
- BlackMagic Design (www.blackmagic-design.com)
- Bluefish 444 (www.bluefish444.com)
- Boxx Technologies (www.boxxtech.com)
- Canopus (www.canopus.com)
- CineForm (www.cineform.com)
- Matrox (www.matrox.com/video/home.cfm)

For a current list of all certified third-party vendors for Adobe Premiere Pro, visit the Adobe Premiere Pro third-party resource page at www.adobe.com/products/premiere/6cards.html.

Post-production for HD content

After you've acquired the content, post-production begins. Post-production is the domain of Adobe Premiere Pro for editing and Adobe After Effects for compositing and finishing.

Post-production for HD is similar to that for SD. One difference is that with HD you're dealing with significantly more data and consequently increasing the load on the CPU and video card. But if you worked in Adobe Premiere Pro and After Effects for SD, you can use the same tools and enjoy the same workflow in HD today.

Choosing an HD monitor

HD video monitors can be quite expensive, especially models with HDSDI input. Someone authoring video for SD broadcast always uses a professional grade video monitor; similarly, anyone making HD for broadcast should have a true HD video monitor.

For some tasks, a computer monitor is a less expensive way to view the actual HD output, especially for progressive projects. Computer monitors use a different color temperature from true video monitors, so it's beneficial to use a video monitor to ensure that you see the true representation of the video as it will be delivered. A number of vendors provide HDSDI to Digital Video Interface (DVI) converters to drive computer monitors. Because high-quality 1920 x 1200 LCD monitors are now available for around U.S. \$2000, this converter is cost effective. Because computer monitors are natively progressive, these converters are not a good solution for natively interlaced projects in Adobe Premiere Pro (the preview in After Effects is progressive, so this isn't an issue for that application).

Another advantage for using broadcast monitors is calibration. These monitors provide features that easily set them to the canonical 709 color space of HD. While this setting is possible with a computer monitor as well, correctly setting the calibration can be difficult. Even systems with a colorimeter for automatic color configuration don't always include 709 presets (although, this standard should be rapidly changing).

Adobe Premiere Pro, with the proper hardware, can also directly send previews out to a video or computer monitor. Sending previews directly is the optimum way to work if hardware that supports this feature is available.

HD can be slow!

Moore's Law has many formulations, one of which is: a computer at a given price point doubles in performance every 18 months or so. Because 1080 60i is 6.5 times as many pixels as 480 60i, rendering time will be similar to how it was four to five years ago, all things being equal. HD is much slower than SD, but on a modern computer, it's about as fast as a turn of the millennium system.

The good news is that modern tools give us a number of techniques to get around these problems.

Specifying rendering quality and compression

After Effects has several rendering modes. As with any project meant for final delivery, you should use the highest quality mode for the final output (and as of version 6.5, new layers automatically default to Best Quality mode now).

High-quality rendering is especially important for content that is going to be compressed for final delivery. When you use lower-quality rendering modes, the errors introduced in the video are difficult to compress and result in more artifacts in the output.

Using 16-bit rendering in After Effects

16-bit per channel rendering was introduced in the After Effects Production Bundle and is now available in both After Effects 6.0 Standard and After Effects 6.0 Pro. The 16-bit rendering process doubles the bits of each channel, which increases the precision of video rendering by 256 times. That extra precision is quite useful, especially in reducing the appearance of banding in an image with subtle gradients. The drawback of 16-bit rendering is that it takes twice the time and twice the RAM as 8-bit.

Rendering in 16-bit isn't necessary when encoding to an 8-bit delivery format like Windows Media or MPEG-2. It does provide a very welcome quality boost when going back out to 10-bit formats like D5.

Support for real-time effects

Many content creators have become accustomed to real-time effects in SD, which were made possible by advancements in nonlinear editing software tools. Currently, achieving similar functionality in HD is not as easy. Common HD workflows include uncompressed video streams where the amount of data required to be transferred to and from the storage sub-system and the load on the CPUs make real-time performance nearly impossible without the use of specialized hardware. These hardware-based systems are still being perfected and, as of NAB 2004, the most common HD workflow requires even simple effects to be rendered.

If uncompressed video streams are not a firm requirement, the use of advanced compression techniques in MPEG-2 and Wavelet technologies can produce visually lossless media at data rates that allow various levels of real-time editing, including multiple streams, effects, filters and graphics. Both hardware- and software-based solutions are available to configure these real-time HD solutions.

Working offline in SD

Due to the intense storage and rendering requirements of HD, it is a good idea to capture and edit a rough cut in SD, and then conform to HD. Adobe Premiere Pro makes this process quite easy. When you use this technique, the standard real-time effects of Adobe Premiere Pro are available when you're editing the offline version. This allows work to proceed quickly with real time visual feedback.

Adobe Premiere Pro 1.5 includes a project trimmer that you can use to edit in a lower-resolution format, such as SD or DV, using any assets with little concern for storage space. You can then save the project offline including the critical data about the assets, edit points, effects, and filters used. Then you can recapture only the HD footage that is necessary to reassemble the project for quick finishing and delivery.

In After Effects, using low-resolution compositions while roughing out a piece provides similar advantages in speed. You can apply the finished low-resolution work to the identical media in HD.

Taking advantage of network rendering

After Effects 6.5 includes grid rendering, which can be a massive workflow improvement for HD authoring, especially with complex projects. After Effects has supported network rendering for a number of years, but the new grid model makes the process much easier. It also enables network rendering for previews, not just final rendering in the render queue—a great enhancement.

You can set up a four-node render farm for less than the cost of a single high-end HD capture card, and it can pay off enormously in a more efficient workflow.

Choosing progressive or interlaced output

Unlike broadcast TV, which is natively interlaced, HD is natively progressive in many cases. This was a contentious issue during ATSC's development because the computer industry was lobbying very strongly to drop interlaced entirely from the specification, and the traditional video engineers were fighting hard for interlaced to remain. In the end, the 720 frame sizes are always progressive, and the 1080 frame sizes can be either progressive or interlaced.

It is always best to perform the post-production in the same format that was used for delivery. Consistency in formats prevents the need for transcoding, which can add artifacts and require long rendering sessions.

If you're creating content to be used in a variety of environments, working in 24p is a great universal mastering format. The 24p format easily converts to both NTSC and PAL, as well as back to film. And its lower frame rate reduces hardware requirements.

SD color space versus HD color space

SD and HD define colors differently. Specifically, SD uses International Telecommunication Union (ITU) 601, and HD normally uses ITU 709. To complicate this matter further, JVC's HDV cameras are HD, but still use the 601 color space.

The good news is that these differences are transparent in most cases. For example, when After Effects converts from its internal RGB to Y`CbCr in a video codec, the codec internally handles the transformation. Therefore, exporting from the same After Effects project to an SD and an HD output should work fine.

Complexity is added when you work in a native color space tool, such as Adobe Premiere Pro. If you capture, edit, and deliver in the same color space (for example, 709), color accuracy and fidelity are maintained. The native filters in Adobe Premiere Pro assume that the 601 color space is used, but this assumption rarely has much effect on the math used in video processing. Still, it's a good idea to use a 709-calibrated monitor for previews when using Adobe Premiere Pro.

Deciding on the distribution media for HD content

All video is created with an audience in mind, and HD is no exception. There are a variety of ways that you can distribute HD content to customers and clients.

SD DVD-video distribution

Most HD content is distributed in SD at some point in its lifecycle. DVD-video is still the de facto standard for mass-distribution of video content, so it's important to understand how you can achieve this.

The MPEG-2 engine in both Adobe Premiere Pro and After Effects can output DVD-compliant MPEG-2 content from HD content. It accomplishes this output by scaling and, in some cases, some frame-rate trickery.

By using the Adobe Media Encoder, you can specify a target frame size for the video. Before the HD video frames are compressed to MPEG-2, the encoder scales them down to the target frame size. Because this scaling occurs at the end of the production process, it can achieve the highest quality compression. In the case of 23.976 fps content, the Adobe Media Encoder can insert data in the MPEG-2 stream so that it plays back interlaced data on a standard TV and DVD player and progressive data on a progressive scan DVD player and TV. This is the same technique that is used in most Hollywood DVDs, which are produced from a 24-fps film master.

After the MPEG-2 content is created, you can output to DVD by using Adobe Encore® DVD software. You can use the video in the DVD-authoring application in the same way as any other video.

DVHS and HDV delivery

Video has been delivered to clients and customers on tape for years. The professional tape formats for HD are too expensive for this use, but there is now DVHS tape format. DVHS adds an MPEG-2 bitstream onto a tape by using the familiar VHS shell (although, the tapes themselves aren't compatible with common VHS players). The good news is that the prices of decks keeps falling rapidly—some good models are now below U.S. \$400.

The MPEG-2 engine built into both Adobe Premiere Pro and After Effects is capable of making DVHS compatible MPEG-2 bitstreams. The CineForm software previously described includes profiles for DVHS authoring and a utility for copying the MPEG-2 data back to the tape.

As mentioned earlier, HDV uses the same bitstream as DVHS, just on a smaller cassette. While HDV decks aren't yet available, they should be available on the market soon and might provide an alternative to DVHS. Still, the physically bigger tape of DVHS should translate into greater durability, so kiosks and other environments where multiple viewings are intended should stick with DVHS.

DVD-ROM delivery

Another major method of distributing HD content is on DVD-ROM, be it high-volume commercial titles, like the Windows Media titles released so far, or one-off DVD-Rs. The installed base of computers that can play HD video is a lot higher than any other HD devices, and you can build an HD compliant system for only an estimated U.S. \$1000.

MPEG-2 and Windows Media are the most popular formats for DVD-ROM in use today, and both can work well. Windows Media format is the first to be used in commercial products, and it has the compression efficiency to get a 2.5 hour movie on a dual-layer DVD in high quality. MPEG-2 requires much higher bit rates for equivalent quality.

One drawback to Windows Media is that it doesn't work at HD frame sizes on Macs. If you want playback on anything other than a computer that runs Windows, another format, such as MPEG-2, is appropriate. Note that there are issues with the availability of MPEG-2 playback software—this topic is discussed later.

Because you can play back the files on a computer, you can build various levels of complexity into the user interface. However, it's a good idea to allow users to pick their own player for the raw files because it is common for users to have their own preferences for players.

Hard drive delivery

For kiosk and laptop presentations, playback from the hard drive is the simplest and most reliable way to go. Space limitations do not exist, and a hard drive is more durable for continual use than a DVD-ROM drive. And because peak data rates are so fast from a hard drive, the real limitation is the system's decode performance.

For delivery of compressed content, any modern hard drive is more than good enough.

Progressive download delivery

Progressive download is the original form of web video. Unlike real-time streaming, progressive download transfers the file using the same protocols as web pages. There are no buffering errors or image glitches due to dropped packets, but there also is no guarantee of real-time performance.

With progressive download, you can encode the content at higher data rates than the viewer's connection speed. Only the portion of the file that has been transferred plays, but the viewer can watch that portion and decide whether or not to continue the download.

Real-time streaming delivery

The most difficult aspect of HD is real-time streaming. While you can serve and play back HD in real time, bandwidth remains a barrier. A visually pleasing 720 24p video requires at least 4 Mbps, even in high-end formats. This data rate is far beyond what is commonly used in homes today and can be more than a business intranet can sustain.

Still, there have been impressive demonstrations of HD streaming over the academic Internet2 network. The fastest consumer connections are only 3 Mbps at best today, but the combination of improving codecs and increasing bandwidth should make HD streaming to users a feasible market later this decade.

Broadcast delivery

Last, but not least, you can broadcast HD as ATSC. In most cases, you provide broadcasters with master tapes of the programs, and then they handle the encoding themselves. However, local broadcasters may be able to accept programming as MPEG-2 transport streams.

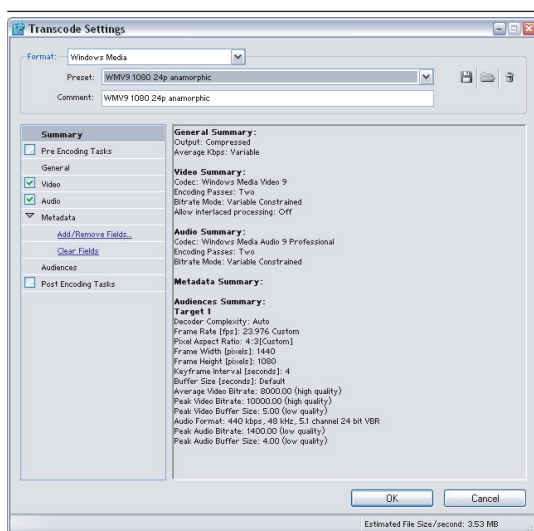
Most HD broadcasting that occurs during the day is in the form of loops of compressed MPEG-2 content—only prime-time programs are broadcast in HD, in most cases. Enterprising HD producers might be able to get some airtime by providing free, high-quality HD MPEG-2 content for local broadcasters to include as part of their loop.

Distribution formats for HD content

A variety of distribution file formats are available for delivering HD—either for final distribution to consumers or for distribution of proofs of work in progress for review.

Windows Media format

The leading format for computer-based distribution of HD content so far is Microsoft's Windows Media format. Multiple titles have been released in the format, which demonstrates excellent quality and compression efficiency, and rock-solid DRM to prevent piracy. Currently, no other format offers the ability to put a two hour movie on a single DVD in an HD format.

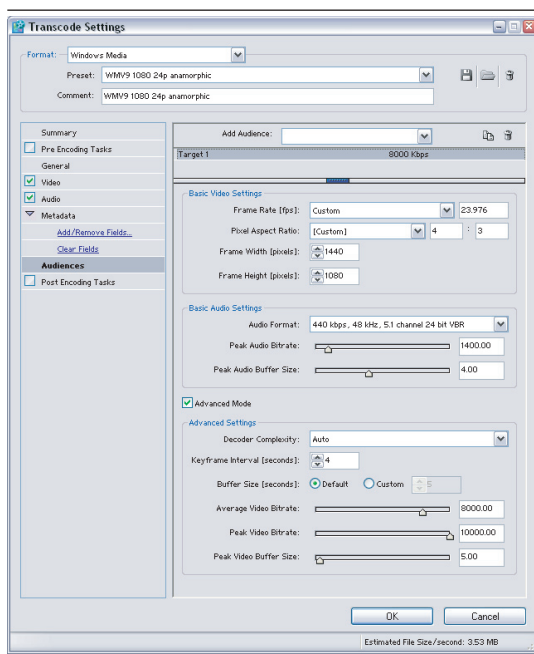


WMV9 1080p video
Detailed video settings for the anamorphic 1080 24p WMV9 setting.

Currently, HD WMV9 works well only on computers that run Windows 2000 or Windows XP®. There is a Mac player from Microsoft, but it doesn't have sufficient performance capabilities for HD playback.

The 1080 WMV format normally uses a frame size of 1440 x 1080 with a 4:3 horizontal compression, which makes the files somewhat easier to decode and improves compression efficiency. Currently, very few HD displays can resolve more than 1440 lines horizontally. WMV does 720 at its normal 1280 x 720. Windows Media has no requirement for frame size, so it is preferable to crop out letterbox bars instead of encoding them—these elements waste bits and CPU cycles. Thus, you might encode a 2.35:1 movie from 1080 source at 1440 x 816.

Microsoft's HD efforts go beyond their own products. They've submitted the WMV9 codec specification (as VC-9) to the Society of Motion Picture and Television Engineers (SMPTE), and they're one of the three proposed mandatory codecs for HD DVD.



WMV9 codec
Audience settings for the Anamorphic preset in the WMV9 codec.

MPEG-2 format

The standard MPEG-2 used for HD is called Main Profile @ High Level or MP@HL. This mode uses the standard techniques for MPEG-2 that are used on DVD and other media but with much higher limits for frame size and data rate.

One drawback to MPEG-2 is that HD compatible decoders don't ship by default with most computers. These decoders are commercially available for U.S. \$30 or less, but you have to either distribute the HD MPEG-2 content with a licensed player or ask the end-users to acquire a player themselves. There are a number of freeware HD-compatible MPEG-2 players. One of the most successful is the VideoLAN Client (VLC). However, it is unclear if it is legal to distribute those players on a DVD-ROM without paying an MPEG-2 licensing fee.

The MPEG-2 specification has Layer II as the standard audio codec, but most HD MPEG-2 products, such as ATSC, use the Dolby Digital (AC-3) audio codec. An advantage to using AC-3 is that it natively supports multichannel audio, but most Layer II implementations are stereo only.

MPEG-2 is another of the three proposed mandatory codecs for HD DVD.

MPEG-4 Part 2 format

The original release of MPEG-4 standard had the video codec defined in Part 2; consequently, that's what the original type of MPEG-4 video (as opposed to AVC) is normally called.

Part 2 is a relatively linear enhancement of the time-tested H.263 video-conferencing codec. Part 2 performs reasonably well at web data rates, but it hasn't proven to have much of an advantage over standard MPEG-2 for HD use. Compression efficiency is somewhat better, but the advantages are not enough to justify the costs of switching from MPEG-2.

MPEG-2 defines the functionality available to a format by means of profiles. The two main profiles available for MPEG-4 for HD are Simple and Advanced Simple. Simple uses a smaller subset of features, which reduces its compression efficiency. Advanced Simple produces smaller files at the same quality but with higher processor requirements. It isn't used much for HD.

MPEG-4 HD almost always uses the AAC-LC audio codec. Most implementations are two-channel, but 5.1 does exist.

The DivX format, which is being touted as an HD format, uses MPEG-4 part 2 video in an AVI file. While the profiles for lower-frame-size DivX allow Advanced Simple, the HD profiles are only Simple profile for decoder performance reasons.

Other than DivX, there isn't much momentum behind MPEG-4 Part 2 HD implementations.

MPEG-4 Part 10/AVC/H.264

Many vendors skipped MPEG-4 Part 2 implementations entirely in favor of the new MPEG-4 Part 10 codec, better known as AVC or H.264. This is a new from-the-ground-up codec that primarily focuses on optimum compression efficiency. The first products to use the new standard are rolling out in 2004. The focus on compression efficiency means that AVC takes more computational power to decode than competing formats. But the goal of much improved compression efficiency was realized.

One open question with AVC is how well quality holds up at HD and high data rates. The initial implementations have some difficulty preserving fine details, like film grain. There aren't yet any real-time HD decoders for AVC, software or hardware, so it's difficult to judge the real-world quality of mature products. Initial indications suggest that it doesn't do as well as WMV9 at HD, largely due to the detail issue, but a definitive answer to that question needs to be determined later.

AVC is one of the three proposed mandatory codecs for HD DVD.

Display devices for HD

You can use a variety of display devices for computer playback. Each device has its own strengths and weaknesses. You can create content that works well on most displays, but if you're authoring for a particular setup, you can optimize for the particular display.

Using CRT computer monitors

The most common type of computer playback is the venerable CRT monitor that uses a VGA connector. CRTs don't have a native resolution, so they can handle a variety of resolutions.

Because CRTs use the analog VGA connector, they can be sensitive to analog noise, especially at high resolutions like 1920 x 1440. By using short runs of high-quality cable, you can avoid this problem. Look for cables with ferrites (barrels near each end). Running through a K Virtual Machine (KVM) switch or using long runs of cables can result in soft images, ghosting, or other

artifacts. If you need long runs or switching, you can use digital DVI connections and a DVI-to-VGA converter at the display to keep the image digital as long as possible. The cheap DVI-to-VGA adaptors won't help—they simply put the analog VGA signal on the DVI cable instead of converting the signal into true DVI.

Most CRTs are 4:3, so 16:9 HD content displays in letterbox format.

Using flat panel displays

There are some important differences among CRT, LCD, and other flat panel displays. The flat panels have a single native resolution, but CRTs can handle a variety of resolutions. And modern flat panels use the digital DVI connector instead of VGA. Avoid using any flat panel by means of VGA—the analog noise of the VGA connector easily results in a degraded image.

When authoring content for a particular flat panel, you should try to encode content for that panel's native resolution to avoid scaling and achieve the crispest image.

Using laptop displays

For demonstration work, using a laptop's built-in screen works well. Due to lower system performance, laptop playback is generally weaker than a high-end desktop computer. You can generally play back 720 24p video on today's fastest laptops, but it can be problematic and difficult to achieve 1080 laptop playback.

When authoring content specifically for playback on a laptop, try to encode at the native resolution of the screen.

Using HDTV displays

HDTV devices include televisions and projectors that are meant to play HD video content as opposed to computer content.

One issue with DVI-based televisions, as opposed to DVI-based computer monitors, is that many DVI-based televisions have DVI ports that don't properly communicate its supported operating modes through the DVI connector. In some cases, directly connecting a computer to an HDTV by means of DVI can confuse the computer and lead to deactivating the computer's video display altogether. DVI dongles that solve these issues, such as the DVI Detective, are available, although in-the-field success with these devices has yet to be quantified.

If a dongle isn't an option, you can use a utility like PowerStrip to tightly control all of the elements of the VGA signal going out to the projector. Using a utility can be a lengthy and frustrating process and doesn't always produce the best results. If possible, use a card with DVI output and a display with proper DVI support.

Computer playback configurations

Even though delivering HD content on computers reaches the largest audience, there are a number of system configuration issues that need to be faced.

Optimizing system performance

The most fuzzy requirement of all is system performance. Traditionally, MHz have been used as the measurement of performance, but that's similar to measuring a car's performance in RPM—it's related to performance, but it's not analogous. There are many other elements that are also important, such as memory bandwidth.

Microsoft's Windows Media Series 9 HD DVD-ROM titles recommend a minimum of a 3 GHz CPU, such as a Pentium 4. That recommendation is made because that processor comes with a minimum memory bus speed that is required. Also, note that frame size and frame rate matter a lot. Playing back 720 24p is less than half as challenging as playing back 1080 60i.

Currently, the fastest Intel and AMD-based systems are more than capable of great HD playback, and many users report successful playback of Microsoft's 1080p content with as little as an Athlon XP2100. However, given the many other variables in an HD playback system, be liberal with your requirement estimates.

For laptop applications, note that the GHz rating of Pentium® M computers may understate their actual performance. Usable 720 playback results have been reported with 1.7 GHz Pentium M systems. The faster the system, the better the performance experience. If you intend to work with HD content of any type, get the most powerful system your budget will allow.

Using high-bandwidth video cards

A fast video card is critical for good HD playback performance. Ideally, you should use an Accelerated Graphics Port (AGP) 8x bus with a 128 MB card or better. Later this year, PCI Express will start to sell high-end video cards that provide excellent performance. Three-dimensional performance doesn't matter for video playback—it's really a matter of the bus bandwidth and the memory bandwidth in the card. High-end 3D cards also come with those features.

Modern laptops, especially those using the ATI Radeon Mobility 9700 chipset, are also capable of quite good HD performance. A few laptops with DVI outputs are also now entering the market.

Playing HD audio

Playing back the audio content from an HD source isn't particularly different than the audio playback associated with SD video. Stereo audio files tend to be the same bandwidth no matter where they are used. One complication with HD is that it is often accompanied by multichannel audio, such as 5.1 surround sound. This accompaniment is a great feature, but it makes playback more complicated.

Playing analog stereo audio

For straight stereo encoded audio, playing back audio is very simple because it works with any stereo equipment. The biggest issues are sound-card quality and analog noise. Cheap sound cards with mini jacks can introduce line noise. It is better to use an external audio breakout box, which leaves the digital-to-analog conversion outside of the computer chassis where there is a lot less noise. Of course, the ideal output is digital, as described in the next section.

Playing analog matrix surround audio

The older Dolby Pro Logic system provided matrix surround audio where two channels of audio used phase and frequency tricks to encode surround information for four channels: left, right, center, and surround. This system works reasonably well, although it doesn't offer the precise spatial acuity that discreet channel encoding does.

There is a massive library of matrix encoded titles available, and they work on the oldest computers. Unfortunately, there is no good way to detect whether or not content is matrixed without playing it through a Pro Logic decoder and listening for directional effects. Stereo played in Pro Logic mode can sound audibly altered, although rarely annoying to the casual listener.

Some systems might offer the ability to down-convert multichannel audio codecs into Pro Logic for playback. This process does work, but it produces lower quality audio and more separation than a discreet solution as described earlier.

Playing analog multichannel audio

With analog multichannel, there is one analog output per channel of audio—typically 5.1, with 7.1 available in the high end. Like analog stereo, an external breakout box can provide better quality than mini jacks on a consumer-grade sound card inside a computer.

Depending on the solution, output is either in stereo pairs, or one per channel, that can be connected to an amplifier, or directly connected to powered speakers. These days, 5.1 speaker setups for gaming are becoming popular. And although 7.1 isn't too significant for computers yet, solutions do exist on the high end.

Playing digital stereo and multichannel audio

The ideal solution for playing back audio from an HD source is an all-digital system. Digital allows perfect quality, so you don't have to worry about noise or make adjustments.

The three significant formats for digital audio are the professional AES/EBU, the consumer-focused (but still used by professionals) SPDIF, and Dolby Digital. All of these formats run over either Toslink optical cable or coax. For playback, both AES/EBU and SPDIF offer perfect multichannel playback.

Normally, the AC-3 solutions only work with AC-3 encoded audio tracks. One feature of some of Nvidia's Nforce motherboards is that they support real-time transcoding from any multichannel audio source into AC-3 for easy integration with existing AC-3 decoders. This option is not as ideal as having multichannel audio out, but it works well in practice.

Conclusion

By understanding the basics of the HD formats and the equipment that supports them, you can make a comprehensive decision about which HD solution that meets your post-production needs—whether you're a hobbyist or a feature film editor.

The advances in technology and the growth in the number of manufacturers providing HD solutions continue to drive prices down to the point where choosing to edit in HD is a real option for video professionals at almost any level.

For a comprehensive overview of Adobe Premiere Pro, please visit <http://www.adobe.com/products/premiere/main.html>

For a comprehensive overview of Adobe After Effects, please visit <http://www.adobe.com/products/aftereffects/main.html>

Adobe Systems Incorporated • 345 Park Avenue, San Jose, CA 95110-2704 USA • www.adobe.com

Adobe, the Adobe logo, Adobe Premiere, Adobe Encore DVD, and After Effects are either registered trademarks or trademarks of Adobe Systems Incorporated in the United States and/or other countries. Mac and Macintosh are trademarks of Apple Computer, Inc., registered in the United States and other countries. Intel and Pentium are trademarks or registered trademarks of Intel Corporation or its subsidiaries in the United States and other countries. Microsoft and Windows are either registered trademarks or trademarks of Microsoft Corporation in the United States and/or other countries. All other trademarks are the property of their respective owners.

© 2004 Adobe Systems Incorporated. All rights reserved. Printed in the USA.



Tools for the New Work™