

WHITEPAPER

AVC-Intra for HD Editing and Production

Why AVC-Intra is now a good choice for high-end HD production

Charlie Dunn, *Grass Valley Servers and Digital News Production*

September 2010

Until recently, content creators have had to juggle multiple different codecs in HD production because, in general, acquisition codecs (DVCPRO HD and XDCAM HD) were not capable of higher end productions, and production codecs (ProRes 422 and DNxHD) were not practically available for acquisition. Although AVC-Intra offered some promise of being the first 10-bit codec with a broader appeal, it suffered from being too computationally complex for most editors of the day and at its introduction was not widely supported beyond Panasonic camcorders.

Today, with the advent of powerful multi-core CPUs and the widespread acceptance of technology based on open standards, this is no longer true. This paper explains why AVC-Intra 100 has some compelling reasons to make it an emerging favorite for producing HD content today and into the future.

Contents

Introduction	1
What Are the Main Codec Choices Today?	1
What Makes a Good HD Production Codec?	2
Is it Fast?	2
Looks	2
Support	2
Usage Costs	2
How AVC-Intra Measures Up	3
AVC-Intra 100 is Fast	3
AVC-Intra 100 Looks Good	4
Full HD	4
Advanced Algorithms	5
Comparisons Between Codecs	6
AVC-Intra 100 is Widely Supported	8
AVC-Intra 100 is Efficient	8
Conclusion	8
Bibliography	8

Introduction

Today, one of the most important decisions that producers of HD content make is to decide which video codec is best suited to what they want to accomplish. This is because each of the many codecs that are available come with a specific set of tradeoffs. You may be trading full HD resolution to avoid a few transcode steps; or a tight budget may be driving you to use the lowest bit rates; or you may just be using the codec that comes with your editor by default without being aware of the set of compromises placed on you by that choice.

This paper will compare the main choices of codecs available today including DVCPRO HD, XDCAM HD, ProRes, DNxHD, and AVC-Intra, explain the tradeoffs represented by each of them, and illustrate the practical impacts these have on your business. It will also present information on two important trends in today's production environment that, taken together, make the 100 Mb profile of AVC-Intra the best all-around choice for production today.

What Are the Main Codec Choices Today?

A summary of the available codecs and their key characteristics can be found in Table 1. It includes uncompressed SMPTE 292 video on the top line to show both why compression is needed (1,500 Mb/s is a lot of data) and how good a match each codec makes to the main video characteristics of raster, chroma sampling, and bit-depth.

Name	Raster	Chroma	Bit depth	Bit rate, Mb/s
Uncompressed SMPTE 292	Full	4:2:2	10-bit	~ 1,500
ProRes 422 (HQ)	Full	4:2:2	10-bit	220
ProRes 422	Full	4:2:2	10-bit	147
ProRes 422 (LT)	Full	4:2:2	10-bit	102
DNxHD 220x	Full	4:2:2	10-bit	220
DNxHD 220	Full	4:2:2	8-bit	220
DNxHD 145	Full	4:2:2	8-bit	145
HQX	Full	4:2:2	10-bit	120-220
DVCPRO HD	Scaled	4:2:2	8-bit	100
XDCAM 50	Full	4:2:2	8-bit	50
AVC-Intra 50	Scaled	4:2:0	10-bit	50
AVC-Intra 100	Full	4:2:2	10-bit	100

Table 1 – Main HD production codec choices.

Two popular choices today are DVCPRO HD and XDCAM HD 50, but they have some inherent drawbacks that limit their long-term potential as a production format:

- Both are only 8-bit, which increases the potential for noticeable contouring and posterization artifacts
- DVCPRO HD only captures two-thirds of the horizontal resolution that is available in the uncompressed signal, so images could appear softer
- XDCAM HD 50 uses Long GOP coding which introduces concatenation (repeated encoding, decoding, and re-encoding), noise even in the simplest editing scenarios, and also causes changes to frames that have nothing to do with those being edited

Because of these issues with both DVCPRO HD and XDCAM HD, some large editing software companies developed their own compression solutions that were higher quality and more friendly to the CPUs of the time. Avid pioneered this with its DNx line of codecs, and was followed by Grass Valley™ with the HQ codec, and then Apple followed suit with the ProRes 422 family of codecs. All of these solutions require an initial transcode from the acquisition format that takes time, but they effectively eliminate the image quality problems we've just mentioned. However some of them have important drawbacks that we'll discuss later that are not shared by the codec that is the focus of this paper: AVC-Intra 100.

What Makes a Good HD Production Codec?

So what makes a good production format in practice? Too often we see white papers on compression focus on the technical aspects of the compression itself, but in our view, this is just one of the factors that should be considered. In most applications it can be reduced down to four primary considerations that have to be weighed to make an informed decision:

- **Speed:** Is it fast?
- **Looks:** Does it yield the picture quality I need?
- **Support:** Can I put together the workflow that I desire?
- **Usage costs:** Does the total cost represent good value?

Is it Fast?

A format is fast when you can:

- Decode and encode multiple streams in real time on common computer platforms so that the editor can work without interruption or intermediate rendering
- Acquire and edit in the same format on the same media to avoid a baseband ingest, a file transfer, or a file transcode before the editing process can begin
- Hand off the finished content to a playout server without having to take the time to flatten or transcode to yet another format

This is where a codec like DVCPRO HD really shines. It's friendly to any generation of PC/Mac platform and you can create workflows that stay in the same format all the way through. XDCAM HD shares many of these characteristics, but there's often a penalty in using a Long GOP codec because the finishing process can take longer when the entire project has to be leveled (i.e., re-encoded) as a final step to re-build the GOP structure. If you just look at the editing part of the project, ProRes, DNX, and HQ are extremely fast because they are friendly to the PC in terms of the number of real-time streams they can support, but they can be slowed down by transcodes either into or out of the editor (HQ does not work on the Mac).

Looks

A format looks its best when:

- It encodes in full resolution (with no scaling artifacts introduced)
- It has full 10 bits per pixel and full 4:2:2 chroma sampling, just like uncompressed HD-SDI
- Multiple generations can be encoded and decoded without noticeable loss in quality (no concatenation)

DNxHD, ProRes422, and AVC-Intra 100 all look really good by these criteria. By comparison, what makes DVCPRO fast is what limits its potential for higher quality applications. A codec like DNX, ProRes, and HQ overcomes this with more bits, but this again is another tradeoff when it comes to what those extra bits ultimately cost. The best solution is the combination of the best looking video with the lowest overhead.

Support

A fast, good-looking format is not of much value unless it is also widely supported by the production tools that are available in the market. Formats based on open standards achieve this more readily than more proprietary solutions, but it's not the only reason for adoption, as we've seen with JPEG 2000. Because Sony and Panasonic have experience and invest a great deal in making standards, they have done an extremely good job at making sure that DVCPRO HD and XDCAM HD have wide support in our industry. While DNX and ProRes have been made available for third-party integrations, and DNX has been made a standard, the number of solutions available is decidedly fewer in comparison with the kind of workflows you can put together with DVCPRO and XDCAM. However, AVC-Intra is now approaching this same adoption and with recent updates and announcement from major suppliers, it is now the only 10-bit format can be used natively on editors from all the "Three A's": Avid, Apple, and Adobe.

Usage Costs

Last, but definitely not least, the costs of using a format weighs on how good it is for production overall. The factors that contribute to the costs of a given codec choice can be some or all of the following:

- **Costs of the bits.** A given bit rate is a direct multiplier to the cost of your infrastructure like an editing SAN, an archive, or network backbone. As an example, the SAN you need to support six editing systems at 100 Mb might easily cost you 2 to 3 times less than the same system using a 225 Mb format, and that could add \$200-\$400k to a system and when money could be better spent elsewhere.
- **Cost of acquisition,** either directly through a camcorder or through a baseband signal to a server connected to a SAN. Do I need to hang another device off my camera (i.e., a KiPro)? Are the prices for servers that record DNX competitive relative to other choices? Do I need to use a secondary tape machine and then ingest directly through the edit suite via baseband?
- **Costs of extra steps.** If you can't build a full workflow around the codec you've chosen there's an inherent cost in buying transcoding products or in the time people spend executing or waiting for the steps to complete.

All of these factors have practical consequences that are oftentimes the result of choosing a format based on either the camera, the editor, or what organizations have been comfortable in using versus looking at the bigger picture of what's needed and what implications those choices will have on the system. Some of those practical considerations may yield the following overall improvements:

- I get double the storage capacity for a given amount of disk space
- It takes half the time for me to move files around
- I can afford to build a SAN that really allows the sharing they need

How AVC-Intra Measures Up

The table below summarizes how well each codec measures up against four key criteria. What we believe comes to light is the overall strength of AVC-Intra 100 as choice for a production format. As we've pointed out in the table, this wasn't true when AVC-Intra was introduced, but we believe those factors have changed.

Codec	Fast	Looks Good	Widely Supported	Efficient
AVC-Intra 100	Yes, though it was complex to decode when introduced	Yes	Yes, though it was limited to Panasonic camcorders at introduction	Yes
DVCPRO 100	Yes	Only 8-bit and horizontal resolution scaled by two-thirds	Yes	Yes
XDCAM HD 50	Long GOP structure can introduce delays in editing	Only 8 bit	Yes	The most efficient
ProRes 422 Family	Yes	Yes	Comparatively fewer options available	45% to 120% less efficient than AVC-Intra
DNxHD Family	Yes	Yes	Comparatively less options available	45% to 120% less efficient than AVC-Intra

Table 2 – Comparison of codec key features.

We aren't making the claim that AVC-Intra is the best for any application, but we are noting that when all the factors outlined are considered, AVC-Intra stands out to us as making the fewest tradeoffs and thus providing a more secure choice for the long run. In this we'll take a closer look at how and why AVC-Intra 100 now meets all four of the key characteristics that a modern HD codec needs.

AVC-Intra 100 is Fast

The class of "intermediate" codecs, including DNxHD, ProRes 422, Grass Valley's HQ, and others that arose between three to five years ago, did so to address the problem that acquisition formats were either too compromised to make good pictures when used in editing and production workflows, or too complex to support enough channels on practical editing platforms. The complexity argument is no longer valid because multi-core, 64-bit processors along with optimized chipset instruction sets specifically engineered to optimize encoder/decoder performance have become available at reasonable prices.

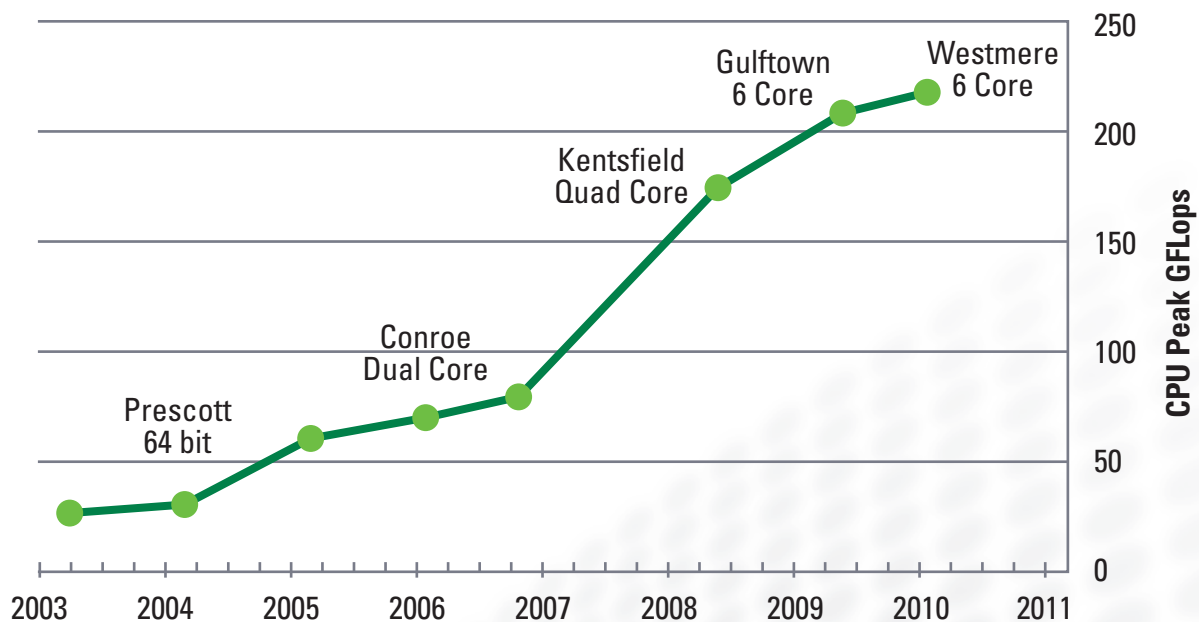


Figure 1 – Recent evolution of processor performance.

How AVC-Intra Measures Up (cont.)

This point is illustrated in Figure 1, which shows that the computational power available today is three to four times greater than it was when DNxHD and ProRes 422 were introduced. In other words, the evolution of computer platforms has eroded the complexity that initially disadvantaged AVC-Intra as a production codec.

But being “fast” isn’t all about computer performance and how quickly one can encode or decode video. Sometimes it’s about avoiding that step altogether. A typical production workflow is sequenced into three stages: acquire, edit, and playout. Intermediate codecs

are only widely supported in the edit phase. This means that to use them is to require transcoding operations as content is ingested to the editor and rendered out of it. There are two problems with transcoding:

- It takes time proportional to the duration of the material being ingested or rendered
- It creates a new generation of coding, which adversely affects picture quality

AVC-Intra can be used in all three phases of production, which eliminates the need for transcoding and avoids these issues altogether.

As well as editing, ingesting, and rendering, the other thing you need to do with your content is to move it around over computer networks which is where the high efficiency of AVC-Intra makes it a faster proposition than the other intermediate codecs covered here. Table 1 illustrates the real time savings that AVC-Intra provides when moving a 20-minute clip across at Gigabit Ethernet network¹.

Codec	Transfer Time	Time Penalty vs. AVC-Intra
AVC-Intra	2:51	N/A
ProRes 422	4:08	1:17
DNxHD 220	6:17	3:25

Table 3 – Transfer time of 20-minute clip over a Gigabit Ethernet network.

So, whether you are moving material from your acquisition device, restoring from an archive, or pushing it to your playout server, it’s going to move faster by 40-100%.

Of course these benefits of speed only matter if AVC-Intra provides image quality comparable to the alternatives.

AVC-Intra 100 Looks Good

So why and how does AVC-Intra provides image quality that’s comparable to the other high-end HD codec choices?

Full HD

AVC-Intra gets off to a good start by encoding the full HD raster with 4:2:2 chroma sampling and 10 bits per pixel. If you’re working with multiple generations of your content or with graphics, using a 10-bit format with 4:2:2 chroma sampling is a must because:

- 8-bit formats will show banding in shaded graphics
- 4:2:0 will cause colors to bleed over subsequent generations

¹Assumes Gigabit Ethernet with effective throughput of 70% wire speed.

How AVC-Intra Measures Up (cont.)

Advanced Algorithms

AVC-Intra manages to compress the full HD signal with best-in-class efficiency by taking advantage of modern compression techniques such as Intra Frame coding. Such techniques have had the focus of two major international standards bodies over the past three decades as illustrated by Figure 2.

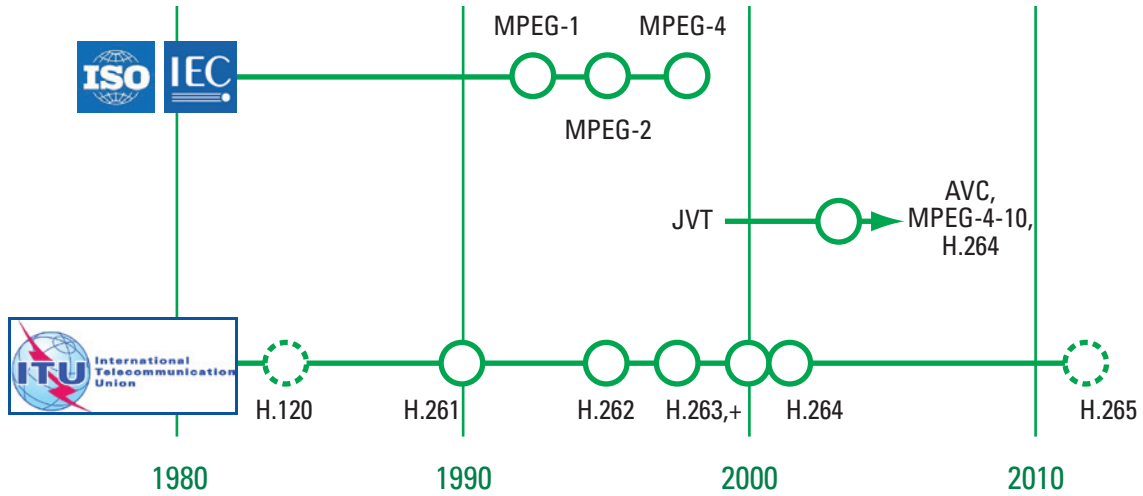


Figure 2 – Development history of H.264 family of codecs, including AVC-Intra.

In contrast to the codecs developed for DNX, ProRes, or HQ that were created in private labs by a few key people, the AVC-Intra compression was developed from a much larger effort and hence has the benefit of more advanced compression tools and techniques developed by a worldwide effort. One particularly powerful tool is called “intra-frame prediction,” and is illustrated in Figure 3. Intra-frame prediction exploits the fact that each part of a picture can look similar to other parts.



Figure 3 – Intra-frame prediction technique. (Images used with permission of Panasonic Corporation)

How AVC-Intra Measures Up (cont.)

Intra-frame prediction finds the best match between similar blocks of pixels in the picture, and uses copies of these as a replacement for the original pixels. Doing this achieves, on average, a huge 60% reduction in the data needed to represent the picture, but does introduce errors that are captured in the so called “residual” image which is typically 30% the size of the original picture. The original image can be perfectly reconstructed by adding the residual back to the intra-frame predicted picture which means that the process nets a 30% compression before any of the more normal compression schemes are applied.

Comparisons Between Codecs

The normal way of making objective comparisons between different video codecs is via their PSNR performance for some well known reference sequences. Although the PSNR does not model the way in which the human visual system perceives noise created by image compression, it is both objective and straightforward to compute (see sidebar). Differences between codecs are only significant if the difference in PSNR exceeds approximately 2 dB. It's important that the same reference content is used when comparing codecs because codec performance is a function of the uncompressed image.

The European Broadcasting Union (EBU) makes available a set of uncompressed reference sequences that are commonly used to evaluate codec performance. When they evaluated AVC-Intra 100 in 2007 with a panel of experts viewing content at a normal viewing distance of three screen heights, they judged the first generation to be visually identical with the reference (Visca, 2008).

The EBU sequences, comprising 1250 frames of video, were used to evaluate two key performance measures for a variety of codecs used in HD production.

The acronym “PSNR” stands for “Peak Signal to Noise Ratio”

It's a handy, single-figure measurement of how different a picture that's been through a codec is compared to the uncompressed original.

Its calculation requires a few simple steps:

- 1) Compress and decompress a reference picture to produce a resultant picture.
- 2) Subtract the resultant from the original; the non-zero values in the resultant are errors that were introduced by the codec.
- 3) Square all the resultant pixels and add them together, then divide by the total number of pixels to get a figure for the Mean Square Error.
- 4) Form the Ratio by dividing the Mean Square Error into the Peak value a pixel can represent (SMPTE 292 defines “peak white” as code value 940, but provides headroom up to 1019).
- 5) Express the ratio in dB.

$$MSE = \frac{1}{N \cdot M} \cdot \sum_{i=1}^M \sum_{j=1}^N (ref_{i,j} - res_{i,j})^2$$

$$PSNR = 10 \cdot \log_{10} \left(\frac{Peak^2}{MSE} \right)$$

How AVC-Intra Measures Up (cont.)

The results are illustrated in Figure 4, and we've used D5 as a reference.

We've plotted the PSNR for each frame in the sequence for four different codecs. Remember bigger numbers are better, so the bright green line at the top of the graph shows the codec with the best picture quality by this method.

You can see that D5 is clearly superior, though that's to be expected with a 250 Mb/s codec. ProRes422 and DNxHD

perform about the same as each other, and slightly better than AVC-Intra, but worse than D5. Again this reflects the bit rate of around 145 Mb/s for these codecs.

Over the long run, AVC-Intra comes in between 1 to 1.5 dB worse than ProRes 422 and DNxHD, which isn't much to sacrifice to gain around 50% efficiency savings. This may be justification enough to edit with AVC-Intra, but there is another strong reason to consider: multi-generation performance of the various codecs.

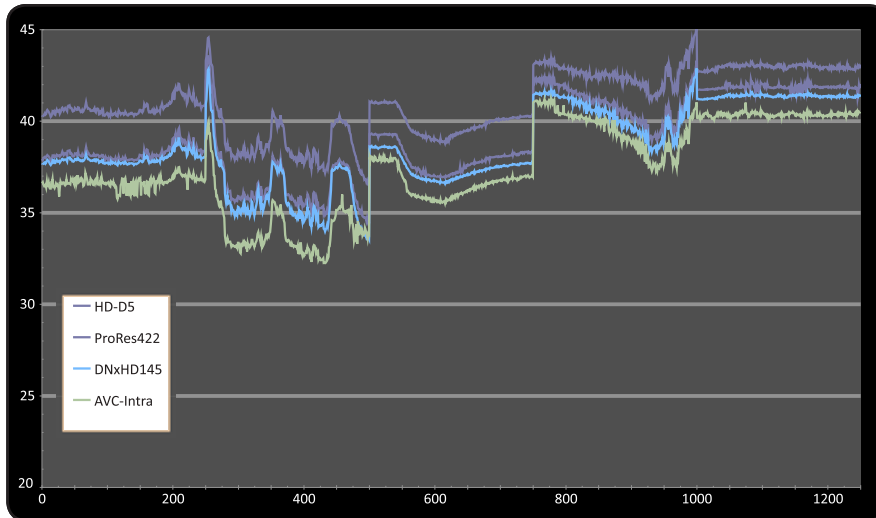


Figure 4 – PSNR comparison using EBU reference sequences.

The impact of repeated encode/decode cycles on the PSNR of a reference image is graphed in Figure 5 with each point on the line representing another generation of image.

You can see that both ProRes 422 and DNxHD suffer from a steep degradation in the early generations, whereas AVC-Intra maintains picture quality across multiple generations much better.

Unless you're doing anything but the simplest editing, it's easy to rack up three, five, or seven generations of a clip during the process. The multi-generational behavior of ProRes 422 and DNxHD all but eliminates their apparent advantages that showed up in the previous Figure. This is especially true when you consider the penalty paid for the initial generation working with DNxHD or ProRes 422 requires, which can be avoided when AVC-Intra is also used as the acquisition codec.

So for practical purposes, AVC-Intra is as good for production as ProRes 422 and DNxHD.

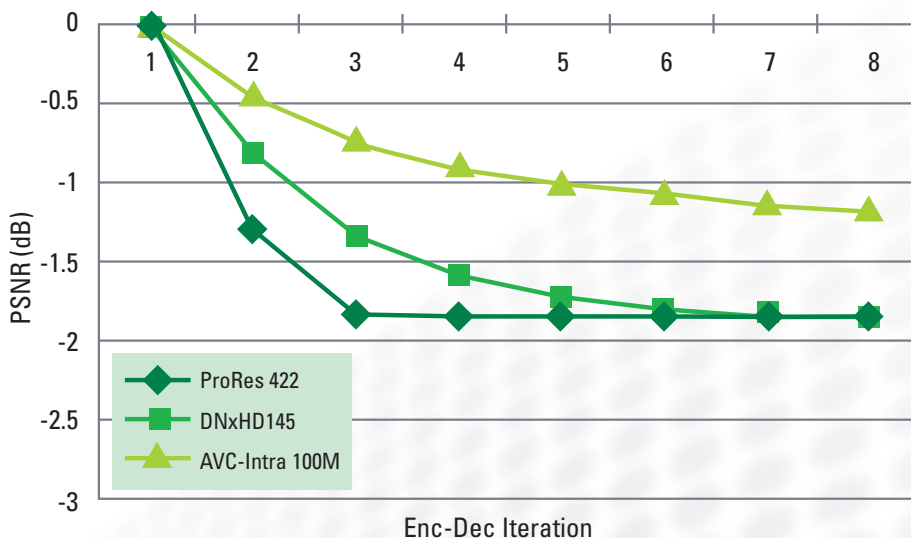


Figure 5 – Multi-generation performance.

How AVC-Intra Measures Up (cont.)

AVC-Intra 100 is Widely Supported

The growth in support for AVC-Intra across the industry is a market trend that's as important and powerful as the dramatic improvement in CPU power touched on earlier. In a way, it's one of the best-kept secrets of the content production industry, but when you look, you start to find support everywhere:

- If you're editing, grading, or compositing, you'll find timeline-level support for AVC-Intra in Grass Valley EDIUS® 5 and 6, as well as Adobe Creative Suite 5; Apple Final Cut Studio 3; Autodesk Inferno, Smoke and Flame, Flint, and Lustre; and all Avid editors, such as Media Composer 5, DS 10, Symphony 5, and NewsCutter 9
- You'll find across the board support for AVC-Intra from Grass Valley (K2 media server family) and other video servers including Avid AirSpeed Multi Stream, DVS VENICE, EVS XT2 (scheduled for Q4 2010), Harris Nexio, Omneon Spectrum, and Quantel sQ
- Transcoding is supported by Telestream and Rhozet Transcoding Services

And there's a wealth of available technology from providers including MainConcept should you wish to develop your own solutions. This is a software codec that will work in both PC and Mac environments. Part of the benefit of using an open standard is that MainConcept can develop its own codecs and use them for a wider range of applications by making them fully MPEG-4 compliant.

In addition to Grass Valley, a number of companies have also developed their own AVC-Intra decoders including Apple and Avid.

There are also many viewing and ingest solutions available for AVC-Intra, including those from Calibrated Software, MOG Technology, MXF4Mac, and OpenCube.

Notice that at no point in the full-HD workflow—from editing to serving to transcoding—are you restricted to a single choice of vendor. AVC-Intra is the only format that can make this claim today.

AVC-Intra 100 is Efficient

Not including audio, AVC-Intra 100 is a 113 Mb/s format. ProRes 422 varies from 145 to 220 Mb/s with Avid DNxHD occupying a similar range. What does this 40% to 120% efficiency advantage gain for you?

- On a single editing system, the savings may not be that significant, though they might mean that you can use that rugged solid-state drive you'd like instead of a spinning disk
- But multiplied over thousands of hours of content or over a large editing SAN, the potential efficiency savings mount up significantly. Potentially you can halve the storage requirement from your capital budget, and halve the archive data tape requirement from your operating budget

Production system costs are driven both by the storage costs and bandwidth. Choosing a higher bit rate will increase them both.

Conclusion

We've shown that AVC-Intra 100 is the best all-round solution for high-quality HD production today when evaluated on the four codec features that make a difference to the production values and economic efficiency of your workflow:

- AVC-Intra is fast because you can avoid time-consuming trans-codes and eliminate their impact on image quality, and computers today are powerful enough to support multiple channels on a practical workstation
- AVC-Intra looks good because of its superior multi-generation performance, and the fact that it's possible to avoid two encode/decode generations by deploying AVC-Intra from acquisition to payout...
- ... which is possible because AVC-Intra is widely deployed—meaning you can choose the point solutions that are best for you from among multiple vendors when defining your workflow
- As the most efficient full HD codec available, AVC-Intra can deliver important cost savings relative to the alternatives. You can halve your storage costs and double your transfer rates

Bibliography

Visca. (2008). HDTV Production Codecs. *EBU Technical Review*, 1-15.

SALES

Local and regional sales contacts can be found by visiting
www.grassvalley.com/sales

SUPPORT

Local and regional support contacts can be found by visiting
www.grassvalley.com/support