Understanding IP Video Technologies

This white paper offers a candid look at the different technologies and use cases in the Audio-Visual (AV) sector. Aimed at both newcomers and more experienced specifiers, integrators and end users, this guide is a 10-minute read that helps to demystify the subject and aids the reader in understanding IP video technologies to help work out which is most suited for a range of use cases.

Exterity White Papers

Understanding IP Video Technologies

Mike Allan, CTO Exterity

ABSTRACT

The Audio-Visual (AV) sector has embraced the move from analogue to digital and is today heading towards an IP centric position which, according to analysts at Informa, is forecast to reach $325 billion in 2024, from $247 billion in 2019[1]. IP-based video offers many advantages over legacy technologies including reliability, flexibility, reach and cost – especially as Ethernet based networks are deployed into new buildings as standard. Today, more than 95% of video and digital signage installations use some version of IPTV.

However, IP-based video technologies vary in terms of their suitability for different use cases. In addition, the terms 'IPTV' and 'AV over IP' (AVoIP) have tended to blur, which has led to some confusion at both a technical and application level.

AV and the journey to IP

The first commercial AV installations emerged in the 1980s and were often as simple as a camera feed transmitting over coaxial cables to several screens situated within a location. Simple TV channel distribution and Video on Demand emerged in the 1990s using analogue RF transmission and were quickly adopted, for example, in hospitality for a fledgling in-room entertainment solution. Early digital offerings such as HDMI matrix switches addressed the need to deliver higher quality and lower latency video, but were expensive to implement with limited ability to scale beyond a small site.

The arrival of the internet and the growth of Internet Protocol during the early 2000s saw a major shift for the industry with the arrival of the first IPTV services – pioneered initially by Internet Service Providers as a way of offering TV alongside broadband services.

The original IPTV technology was designed for private networks using a set of widely understood and deployed technologies, including Real-time Transport Protocol (RTP), User Datagram Protocol (UDP), plus the most prevalent compression technologies based on the MPEG-2 family of Codecs.

Along with Transmission Control Protocol (TCP) for public internet connectivity, this quartet of technologies is widely used in related applications for everything from watching content on a TV or laptop, Netflix and YouTube on mobile devices and video chats via Apple Facetime, through to massive control rooms monitoring banks of CCTV cameras.

In the 2000s, the growing maturity of IPTV saw it start to systematically replace older analogue technologies based on RF and less flexible matrix switches. One of the biggest benefits is the simplicity and flexibility of using a single shared structured cable network and relatively low cost IP switches for internet connectivity, application data and AV – with the ability to move, add or change end points with ease.

The last decade has seen per port Ethernet switching costs decline while speeds have increased that have helped to propel IPTV into a dominant position for AV use cases. Today, more than 95% of video and digital signage installations use some version of IPTV – and in the vast majority of cases it still makes the most practical economic sense.

However, there have always been certain AV use cases such as medical imaging and live event simulcasting that require the highest levels of video fidelity along with extremely low latency. In the past, these have been the preserve of technologies such as HDMI switching, but new low compression codecs and higher speed lower cost networks have enabled IP to start to address these use cases. These use cases have spawned several AVoIP technologies designed to replace dedicated HDMI switching solutions, albeit with a few compromises, the most significant of which is the cost of upgrading network infrastructure.

How do IP-based video systems work?

IP-based video solutions, whether they are described as IPTV or AVoIP, work in a very similar way:

1. Content is ingested from a pre-compressed source. This could be a broadcast signal such as a TV channel from satellite, cable or DTV or a feed from a Video on Demand server. Alternatively, this ingested content may be a base band video source, e.g. HDMI or SDI, which is then compressed to reduce the amount of bandwidth.

2. This digital compressed video content is then turned into packets that are transmitted across an IP network in chunks. This transmission is normally multicast (one to many) or unicast (one to one) with packets flowing across the network through switches and along Ethernet cables or Wi-Fi networks to reach a destination.

3. At the destination device, the packets are received, assembled back into video, decompressed and displayed on a device such as a TV screen, video wall, projector or even smartphone.

Each of these steps adds a small time delay, or latency, between transmission from the source to being displayed at the destination. In addition, each time the source video is compressed and decompressed, there is a degradation in visual quality, which is a trade-off for the video content being mathematically altered to use less bandwidth. Although a very simplistic description, this concept is broadly the same for IPTV, AVoIP, Software Defined Video Over Ethernet (SDVoE) and several other specialist technologies. The major difference between the different technologies is the degree of compression and the resulting latency.

What are the differences between the various IP-based video technologies?

The varying IP-based video technologies tend to focus on different elements of this delivery chain to better serve different use cases. Comparing each technology is challenging as it depends on the specific use case. But let's consider a couple of benchmarks.

If you are considering carrying uncompressed video over your network, you need to be aware of the bandwidth it will consume. Uncompressed video at 1080p, 60 Hz, YUV 4:2:2, 10-bit, consumes around 2.49 Gbps. The bit rate (Gbps) required is dependent upon many factors, including resolution, frame rate, colour subsampling, bit depth per channel and HDR format. This means that a 4K video can require bit rates of up to 18 Gbps, but typically requires around 9-10 Gbps.

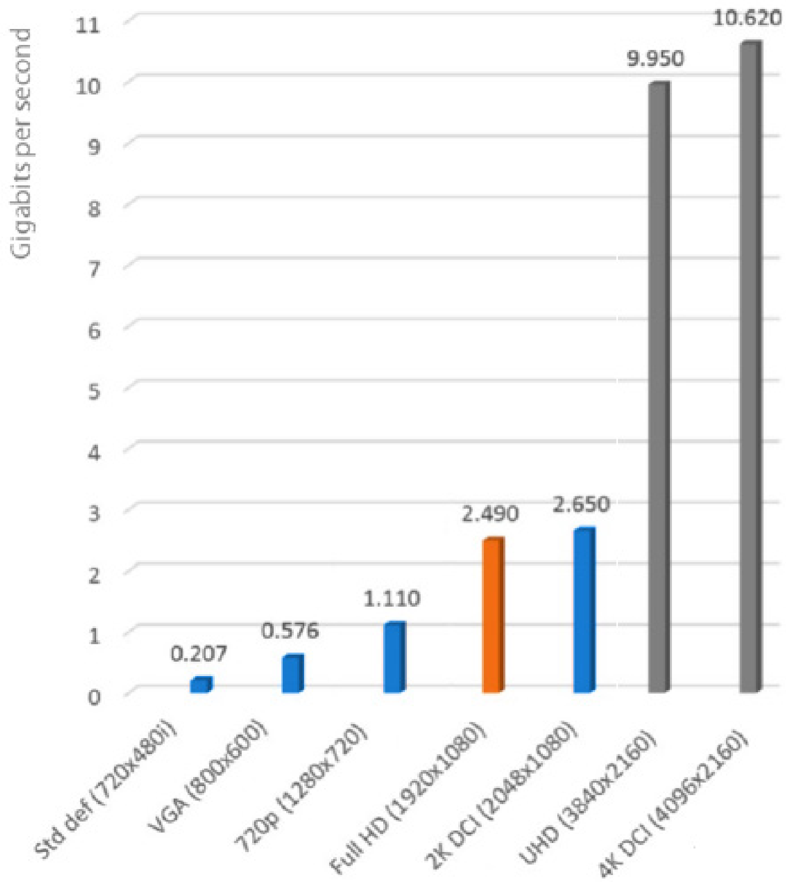

Bandwidth comparison of different video formats [2]

Most end points installed within the last 5 years are connected via 1000Base-T (IEEE 802.3ab Gigabit Ethernet that uses Category 5 UTP wiring, and there are still many legacy end points which are older 100Base-T. Although 10GBase-T (802.3.an) ports are becoming more common as cost per port drops, these are typically for core networks or server to server connectivity as they offer little advantage for a desktop or endpoint connection. It is clear that delivery of uncompressed video is simply not possible to most Ethernet end points.

	Category 2	Category 5	Category 5a	Category 6	Category 6a	Category 7
Cable Type	UTP	UTP	UTP	UTP or STP	STP	SSTP
Max Data Transmission Speed (Mbps)	10	10/100	1000	1000	10 000	10 000
Max Bandwidth	16 MHz	100 MHz	100 MHz	250 MHz	500 MHz	600 MHz

Different Ethernet categories [3]

Compression technologies are used to reduce the bandwidth consumption of video transmission and range from lossless formats such as MJPEG that offer a 2:1 ratio to newer HEVC variants that can reach as high as 1000:1. However, as compression becomes more aggressive, picture quality can degrade and latency is introduced due to the time it takes to compress and decompress each frame. The impact of compression on a viewer's perception of quality is highly dependent on the use case, video source and output device.

In the UK, content transmitted in HD by the likes of BBC HD or Sky and delivered for UK HD TV transmission must be 1920 x 1080 pixels in an aspect ratio of 16:9, 25 frames per second (50 fields) interlaced — now known as 1080i/25 with a colour sub-sampled at a ratio of 4:2:0. This HD format is fully specified in ITU-R BT.709-5 Part 2.

However, in the case of BBC HD, its channels are encoded in H.264/MPEG-4 AVC for satellite and terrestrial broadcasts and in MPEG-2 for cable transmissions. This compression results in an HD channel that consumes the equivalent of 8-12Mbps of bandwidth, with a negligible impact on quality when watched on a typical 55" TV.

The possible permutations of resolutions, frame rates, colour depths and compression schemas make IP-based video distribution incredibly powerful as different combinations can be used to serve different use cases. For example, a typical hotel room infotainment system where the end point is 100Mbps may use 720p / 25 / 4:2:0 with MPEG2 compression resulting in less than 4Mbps of bandwidth consumption per channel. In contrast, a sports bar with 4 large screens connected via 1000mb (Gigabit Ethernet) may instead run 2160p (UHD), 4:4:2 with HEVC compression that consumes around 20MBps, while an IPTV-based digital signage solution at a large stadium could be connected via a Wi-Fi link with each static image or simple animation.

IPTV is used to live stream matches at TSG Hoffenheim, Germany

Understanding IPTV

IPTV is the most mature IP-based video technology and, in many ways, the 'TV' moniker is a misnomer in the modern age. Many of the most prevalent use cases for IPTV are for hybrid services that mix on-demand services with TV. For example, most of the world's large hotel chains have deployed IPTV for in-room entertainment where the same system is used for TV, on-demand films, casting, booking of services and checking of bills. The technology is also deployed widely in common settings such as the electronic menus in fast food outlets, the large screens in sports bars and across shopping malls the world over.

The primary strength of the technology is its maturity, which has led to a level of ubiquity within the market. This does not mean stagnation, as IPTV has progressed to support higher resolutions, such as 4K, along with integration with other IP-based systems. Flexibility is another strength. As the technology enters its third decade, a growing ecosystem of technologies support the combination of protocols and codecs that make up IPTV. As such, modern displays from the likes of Samsung, LG, Philips and Sony can directly display IPTV streams without the need for additional local equipment. This flexibility also extends to newer categories of devices such as tablets and smartphones, which can display IPTV streams via a wide array of apps.

Compromises of IPTV emerge in the more specialist use cases that need to display uncompressed video with extremely low latency. Imagine a medical procedure such as an endoscopy; in such a scenario, image degradation through compression or delay between what the surgeon is doing on the screen versus the audio commentary can have critical consequences. Working with very large video walls measured in the tens of metres where the video source is beyond 4K is another area where highly compressed video is less suitable. In this scenario, the quality of the video may be impacted if the source is compressed and then upscaled. Codecs used for IPTV are much more suited to natural scenes; however less compressed video is more applicable to small text or detailed diagrams. Another benefit is where there is a need for interactivity, for example a remote control application, where the reduction in latency can provide smoother control.

Understanding AVoIP and SDVoE

AVoIP, including SDVoE at the top end, has emerged predominately as a technology that serves use cases in which higher quality and lower latency are the key requirements. AVoIP is functionally like IPTV in many ways, although its ability to deliver lightly compressed video at the typical 1Gb Ethernet network level is its most compelling advantage. However, for any end point with a typical consumer grade 55" LED display, this advantage does not translate into a tangible visual benefit for the end viewer when looking at content such as TV, movies or sport, as the screen is not big enough to reflect the improvement in resolution. Yet where detail such as smaller text or precision diagrams is important, the higher visual quality can be of benefit.

Where screens start to exceed 100" and up to multi-metre video walls, this improvement in video quality also starts to become more noticeable. To date, the majority of AVoIP deployments have focused on specialist use cases. At the higher end dominated by SDVoE, the use cases are almost exclusively working with almost uncompressed 4K content where latency is the biggest factor. A half second delay may well be critical for a stadium's main video wall or medical imaging application, but makes no difference to a hotel guest consuming entertainment content. In remote control applications, for example drone piloting, SDVoE quality and latency benefits are also a major plus point. At the same time as SDVoE has emerged, we have seen the development of a range of SMPTE standards designed to transport video over IP. Initially aimed at the broadcast industry, as a standard developed by a recognised standards body, this could represent the way forward for distribution of uncompressed or lightly compressed video over an IP network.

The biggest compromise inherent in both technologies, and more so with SDVoE, is the requirement for a more advanced network to deliver uncompressed video. This means that customers who need to carry multiple streams have to invest significantly in a move to 10Gb Ethernet and potentially upgrade not just switches, but structured cabling across a facility.

AVoIP at work in a sports competition arena

IP Technology	Interoperability	Latency	Quality	Bandwidth consumption	Network requirements	Cost / complexity	Typical use case
IPTV	Uses standards-based technologies with hundreds of vendors / thousands of products alternatives. Also suitable for digital signage	Typically around 400ms	Supports SD, HD using MPEG and up to 4K/ UHD using a modern codec such as HEVC	Typically scales from under 1Mbps for SD video to 16Mbps for highly compressed 4K video. Suitable for low bandwidth links such as Wi-Fi and cellular networks	Suited to 100Mb to 1Gb Cat 5 Ethernet – any off the shelf switchers. Can support 10Gb to effectively deliver hundreds of simultaneous HD streams	Lowest cost due to mass market adoption, simpler network architecture including shared deployments plus a wider ecosystem of suppliers	All content, everywhere with slightly higher latency
AVoIP	New standards along with a number of proprietary versions from tens of vendors / products	Between 10ms to 200ms depending on transmission protocol and type of compression	Typically up to 4K/UHD with support for 8K. Also supports lossless J2000 full frame	Typically not used for SD but HD starts at 100Mbps (e.g., NDI) and close to 1Gbps for several AVoIP solutions	Suited to 1Gb to 10Gb networks – any off the shelf switches. Typically 6/3 HD/4K channels over a 1000Base-T network	Slightly higher cost of system elements and often requires upgraded network architecture – smaller ecosystem reduces choice	A lot of content to selected users or groups with lower latency
SDVoE	Single standard adopted by 10-20 vendors within a closed ecosystem	Typically under a single frame of latency	Typically up to 4K/UHD with support for 8K Fully uncompressed transmission	Typically each stream is native at 2.5Gbps for HD, 10Gbps for 4K	Requires 10Gb network across all endpoints. 4/1 HD/4K channels over 10GBase-T (802.3an) network segment	Requires 10Gb network from source to destination along with more expensive system components	Limited range of uncompressed content to specific locations with lowest latency

Summary

The IP-based video market has a range of options, but the crux of the matter is that the use case should always dictate the technology choice, reflecting costs, genuine needs and future flexibility.

IPTV, in all its guises, including TV services delivered alongside broadband, hospitality, digital signage, live TV re-broadcast and video-based training is currently the industry standard. In most use cases, IPTV offers more flexibility, lower cost, simpler deployment and ultimately more freedom to adapt to new requirements.

The remaining situations are cases in which AVoIP/SDVoE can provide significant advantages that are needed for uncompressed video, extremely low latency or an unusual use case. A 500ms second delay can be critical for a sports arena video wall, remote control drones, or laparoscopic surgery, but makes no difference to most organisations which are consuming a general mix of video and graphics, such as digital signage.

At present, many of the AVoIP solutions are still relatively immature, which makes the number of suppliers and available alternatives rather limited. This means a higher sticker price for solution elements and almost certainly a network upgrade to contend with. That may change as AV technologies and use cases evolve, but for the foreseeable future, IPTV meets most organisations' needs in the majority of use cases.

References: [1] AVIXA, [2] biamp, [3] Overblog