Network structures — the internet, IPTV and QoE

Jeff Goldberg and Thomas Kernen
Cisco Systems

How would a broadcaster transmit TV transported over IP packets rather than using traditional broadcast methods?

This article introduces a view of a generic Service Provider IP distribution system including DVB's IP standard; a comparison of Internet and managed Service Provider IP video distribution; how a broadcaster can inject TV programming into the Internet and, finally, how to control the Quality of Experience of video in an IP network.

Transport of broadcast TV services over Service Provider managed IP networks

The architecture of IP networks for the delivery of linear broadcast TV services looks similar to some traditional delivery networks, being a type of secondary distribution network. The major components are:

- Super Head-End (SHE) – where feeds are acquired and ingested;
- Core transport network – where IP packets route from one place to another;
- Video Hub Office (VHO) – where the video servers reside;
- Video Serving Office (VSO) – where access network elements such as the DSLAMs are aggregated;
- Access network – which takes the data to the home – together with the home gateway and the user's set-top box (STB).

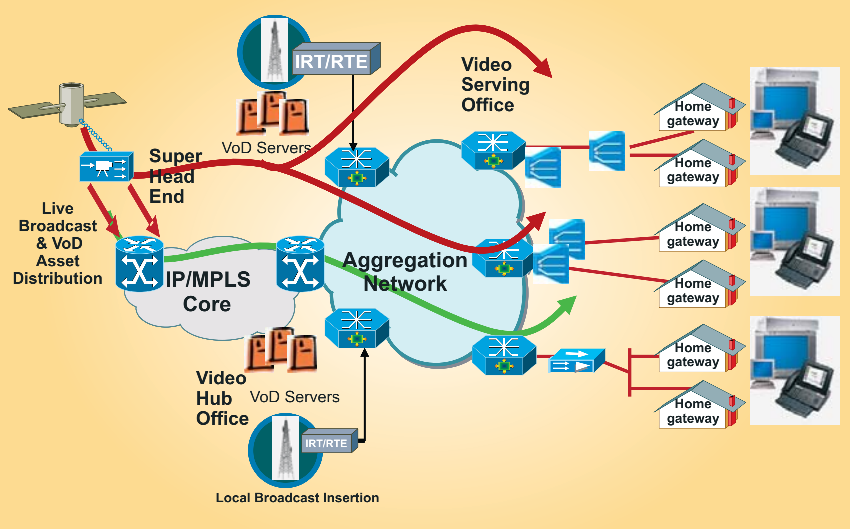

Figure 1
Broadcast TV over an SP-managed IP network

The whole network, however, is controlled, managed and maintained by a single Service

Provider (SP) which allows him to control all the requirements needed to deliver a reliable service to the end point. These requirements are, for example, IP Quality of Service (QoS), bandwidth provisioning, failover paths and routing management. It is this management and control of service that separates a managed Service Provider IP delivery of video streams transported over the public Internet.

The Service Provider acquires the video source in multiple ways, some of which are the same as in other markets, such as DVB-S. This results in significant overhead as the DVB-S/S2/T/C IRDs and SDI handoffs from the broadcasters form a large part of the acquisition setup. It is therefore preferable to acquire content directly from another managed network using IP to the head-end, something that is more efficient and becoming more common.

Once the content has been acquired, descrambled and re-encoded, it is then carried as MPEG-2 Transport Streams (TS) encapsulated into IP packets instead of the traditional ASI. The individual multicast groups act as sources for the services which are then routed over the infrastructure, though in some highly secure cases, these may go through IP-aware bulk scramblers to provide content protection. If security is important, then routers at the edge of the SHE will provide IP address and multicast group translation to help isolate the head-end from the IP/MPLS core transport network.

Abbreviations

AL Application Layer
ASI Asynchronous Serial Interface
ATIS Alliance for Telecommunications Industry Solutions (USA) http://www.atis.org/
AVC (MPEG-4) Advanced Video Coding
BER Bit Error Rate
BGD Broadband Gateway Device
CBR Constant Bit-Rate
CoP4 (Pro-MPEG) Code of Practice 4
DAVIC Digital Audio-Visual Council
DHCP Dynamic Host Configuration Protocol
DLNA Digital Living Network Alliance http://www.dlna.org/home
DNS Domain Name System
DSG (CableLabs) DOCSIS Set-top Gateway
DSL Digital Subscriber Line http://www.dslforum.org/
DVB Digital Video Broadcasting http://www.dvb.org/
DVB-C DVB - Cable
DVB-H DVB - Handheld
DVB-S DVB - Satellite
DVB-S2 DVB - Satellite, version 2
DVB-T DVB - Terrestrial
ETSI European Telecommunication Standards Institute http://pda.etsi.org/pda/queryform.asp
FC Fast Convergence
FEC Forward Error Correction
FRR Fast Re-Route
GUI Graphical User Interface
HGI Home Gateway Initiative http://www.homegatewayinitiative.org/
HNED Home Network End Device
HNN Home Network Node
IP Internet Protocol
IPI Internet Protocol Infrastructure
IPTV Internet Protocol Televison
IRD Integrated Receiver/Decoder
ISMA Internet Streaming Media Alliance http://www.isma.tv/
ITU International Telecommunication Union http://www.itu.int
IXP Internet eXchange Point
MDI Media Delivery Index
MLR Media Loss Rate
MPLS Multi Protocol Label Switching
MPTS Multi Programme Transport Stream
NGN Next Generation Network
NMS Network Management System
QAM Quadrature Amplitude Modulation
QoE Quality of Experience
QoS Quality of Service
QPSK Quadrature (Quaternary) Phase-Shift Keying
RF Radio-Frequency
RSVP ReSource reserVation Protocol
RTP Real-time Transport Protocol
RTSP Real-Time Streaming Protocol
SDI Serial Digital Interface
SDV Switched Digital Video
SHE Super Head End
SP Service Provider
SPTS Single Programme Transport Stream
STB Set-Top Box
TE Traffic Engineering
TS (MPEG) Transport Stream
UDP User Datagram Protocol
UGD Uni-directional Gateway Device
VBR Variable Bit-Rate
VHO Video Hub Office

The core network lies at the centre of transporting the stream to its destination but it is the recent developments of high speed interfaces that have made it possible. The low cost and widely available Gigabit Ethernet, the more expensive 10 Gigabit Ethernet and the swift 40 Gigabit interface now provide the ability for the core to transport both contribution and distribution video streams. The modern optics used in these interfaces deliver Bit Error Rates (BERs) and latency that is lower than those of traditional transports such as satellite. These advantages, combined with an application layer Forward Error Correction (FEC) scheme – such as the Pro-MPEG Forum Pro-MPEG Code of Practice 4 (CoP4) and IP/MPLS Traffic Engineering (TE) – allow for redundant paths across the transport infrastructure. These paths can be designed in such a way that the data flows without ever crossing the same node or link between two end points, and delivers seamless failover between sources if the video equipment permits it. In addition, Fast Re-Route (FRR) and Fast Convergence (FC) reduce the network re-convergence time if a node or link fails to allow for swift recovery, should a path fail.

The transport stream can also use the characteristics of any IP network to optimize the path and bandwidth usage. One of these characteristics is the ability of an IP network to optimally send the same content to multiple nodes using IP Multicast, in a similar manner to a broadcast network. This characteristic has many applications and has proven itself over a long time in the financial industry, where real-time data feeds that are highly sensitive to propagation delays are built upon IP multicast. It also allows monitoring and supervision equipment to join any of the multicast groups and provide in-line analysis of the streams, both at the IP and Transport Stream level. These devices can be distributed across the network in order to provide multiple measurement points for enriched analysis of service performance.

The Video Hub Office (VHO) can act as a backup or a regional content insertion point but also may be used to source streams into the transport network. This sourcing can be done because of a novel multicast mechanism called IP Anycast, which enables multiple sources to be viewed by the STB as one single and unique source, using the network to determine source prioritization and allowing for source failover without the need of reconfiguration.

Primary and secondary distribution over IP

The bandwidth of individual or collective services in primary distribution between a studio or a playout centre and the secondary distribution hubs is traditionally limited by the availability and cost of bandwidth from circuits such as DS-3 (45 Mbit/s) or STM-1 (155 Mbit/s). This has restricted the delivery of higher bitrate services to such hubs that may benefit from a less compressed source.

The flexibility of IP and Ethernet removes these limitations and enables services to be delivered using lower compression and/or with added services. This means that delivery over an IP infrastructure is now possible:

- to earth stations for satellite (DVB-S/S2) based services;
- IPTV (DVB-IPI) or cable (DVB-C) head-ends;
- terrestrial (DVB-T) or handheld (DVB-H) transmitting stations.

We shall now look at two examples of this: firstly, Cable distribution and, secondly, IP distribution via DVB's IPI standard.

Example 1: Cable distribution

Cable distribution typically follows a similar pattern to primary and secondary distribution, with the major exception being the use of coaxial cable over the last mile. IP as a transport for secondary distribution in systems such as DVB-C has already been deployed on a large scale by different networks around the globe. Multiple Transport Streams (MPTS) are run as multicast groups to the edge of the aggregation network where edge "QAMs" receive the IP services and modulate them onto RF carriers for delivery to cable STBs.

The modulation onto RF carriers can be done in one of two ways: by translating a digital broadcast channel to the STB or by using a cable modem built into the STB to deliver it directly over IP. In the latter case, as it is a true IP system, the distribution could use DVB IPI described previously without any modification.

Today, almost all of the STBs have no cable modem internally so the IP stream terminates in the hub-site closest to the STB and even if they did, the data infrastructure is often separate from the video infrastructure. This separation is beginning to change as cable data modems become much cheaper and the data infrastructure costs become lower. An in-between stage is emerging where most of the broadcast channels are as before, but some of the little-used channels are sent via IP, known as "Switched Digital Video" (SDV). The consumer notices little difference between a Switched Digital Video channel and a standard digital cable channel since the servers and QAMs in the hub and/or regional head-ends do all the work. The SDV servers respond to channel-change requests from subscriber STBs, command QAM devices to join the required IP multicast groups to access the content, and provide the STBs with tuning information to satisfy the requests. The control path for SDV requests from the STB is over DOCSIS (DSG), or alternatively over the DAVIC/ QPSK path. In some designs, encryption for SDV can also take place at the hub in a bulk-encryptor, so minimizing edge-QAM encryption-key processing and thus speeding up the channel-change process.

Example 2: IP distribution to the STB via DVB IPI

DVB has had a technical ad-hoc committee (TM-IPI) devoted to IP distribution to the STB since 2000 with a remit to provide a standard for the IP interface connected to the STB. In contrast to other standards bodies and traditional broadcast methodology, it is starting at the STB and working outwards.

In the time since TM-IPI started, many groups around the world have discovered IP and decided to standardize it *(see Fig. 2)*. The standards bodies shown are:

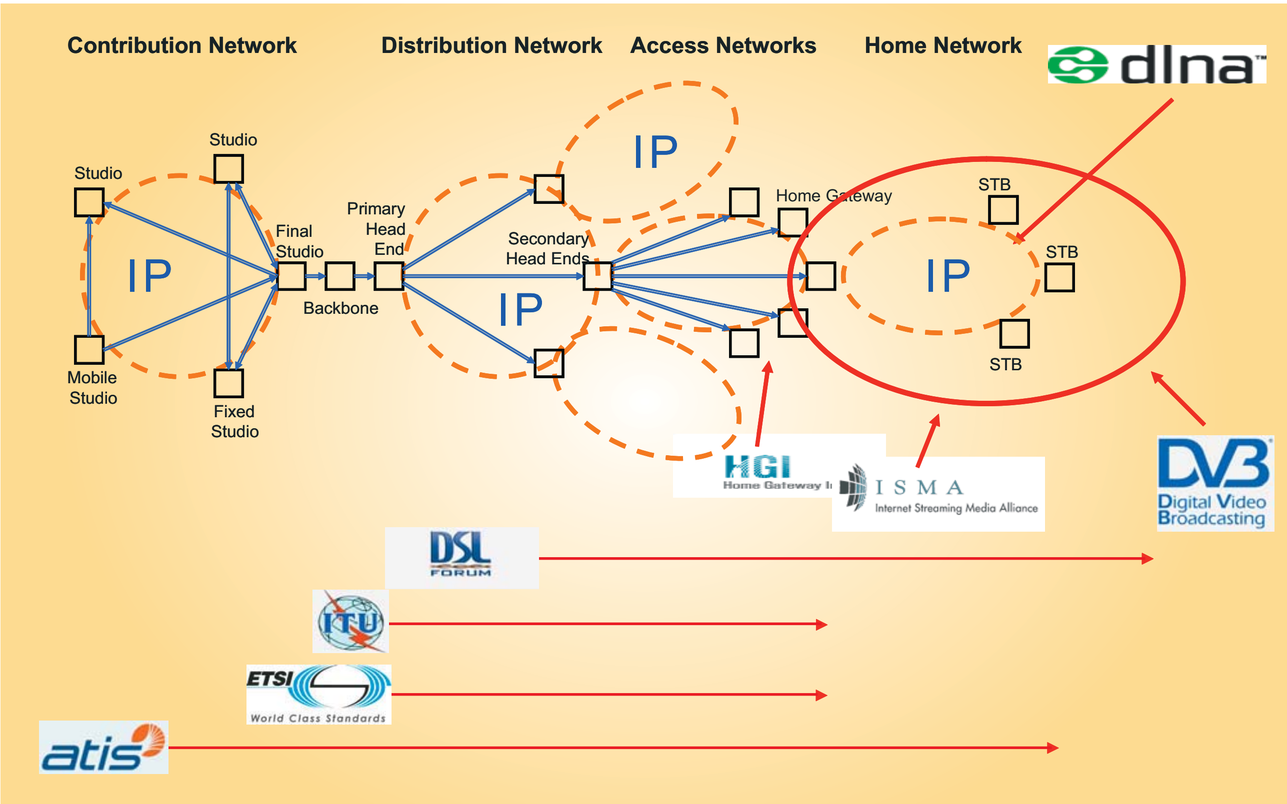

Figure 2
IPTV-related activities of selected standardization bodies

- **DLNA** (Digital Living Network Alliance) for the home network – see also the section “The Home Network and IP Video”;
- **HGI** (The Home Gateway Initiative) for the standards surrounding the residential gateway between the broadband connection and the in-home network;
- **ISMA** (The Internet Streaming Media Alliance) for the transmission of AVC video over IP;
- **DSL Forum** for the standards surrounding DSL and remote management of in-home devices including STBs and residential gateways;
- **ITU** which, via the IPTV Focus Group, is standardizing the distribution and access network architecture;
- **ETSI** which, via the NGN initiative, is standardizing the IP network carrying the IPTV;
- **ATIS** which, via the ATIS IPTV Interoperability Forum (ATIS-IIF), is standardizing the end-to-end IPTV architecture including contribution and distribution.

Nevertheless, the DVB-IPI standard does mandate some requirements on the end-to-end system *(see Fig. 3)*, including:

- The transmission of an MPEG-2 Transport Stream over either RTP/UDP or over direct UDP. The method of direct UDP was introduced in the 1.3.1 version of the handbook. Previous versions only used RTP, and the use of AL-FEC requires the use of RTP.
- Service Discovery and Selection either using existing DVB System Information, or an all-IP method such as the Broadband Content Guide.

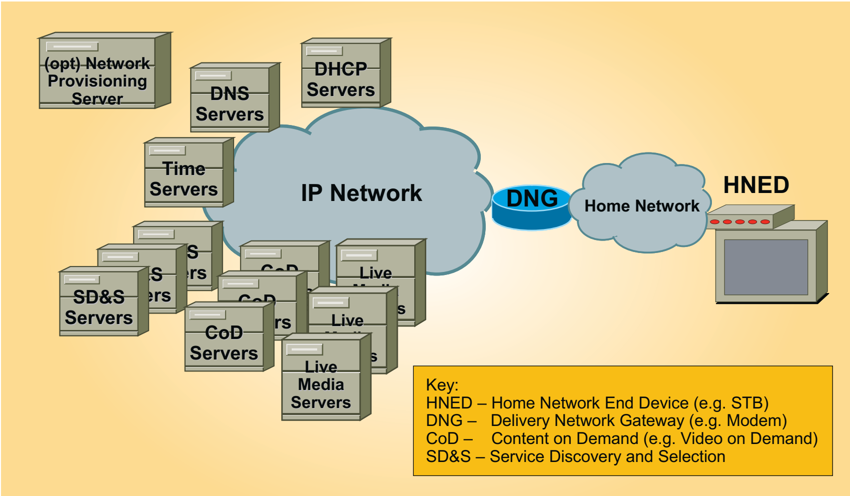

Figure 1
DVB-IP version 1.3 Architecture

- Control of content on demand using the RTSP protocol.
- The use of DHCP to communicate some parameters such as network time, DNS servers etc. to the STB.

It is normal in IPI to use single-programme transport streams (SPTS) as the content are normally individually encoded and not multiplexed into MPTS as they would be for other distribution networks. This provides the added flexibility of only sending the specifically-requested channel to the end user, which is important when the access network is a 4 Mbit/s DSL network as it reduces bandwidth usage.

IPTV and Internet TV convergence

The two worlds of managed STB and unmanaged Internet TV are coming together with sites like YouTube or MySpace showing user-generated content and excerpts from existing TV programming. Internet TV demonstrates what can be done with an unmanaged network across a diversity of different networks, including one in the home. In this section we’ll cover what the home network will look like, compare IPTV to Internet TV, and show how a broadcaster can place content on the Internet via an Internet Exchange.

The Home Network and IP Video

Improving technologies of wireless networks, increases in hard-disk-drive sizes and the increasing number of flat-screen TVs in European households, makes the home network inevitable in the near future. Unfortunately the home network still remains more of promise than reality for high-quality broadcast TV transmission, mainly because the standards and interoperability are some way behind.

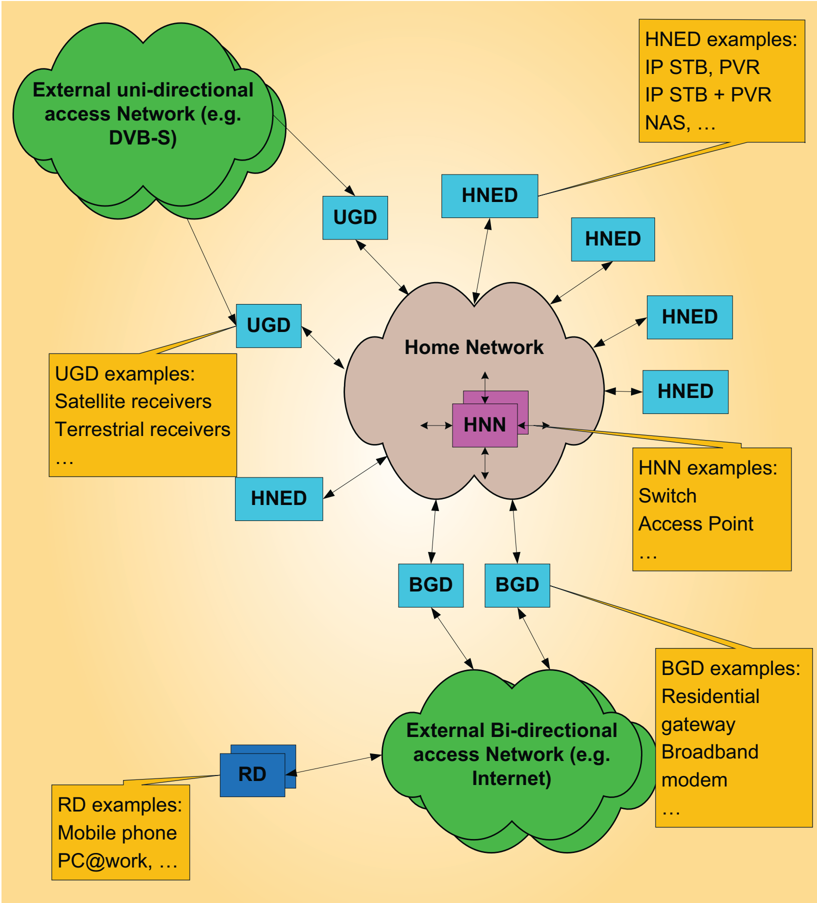

Figure 4
DVB IPI Home Network Reference Architecture

DVB has just released a Home Network reference model which is the first part of a comprehensive specification which will be completed in 2008. The home network consists of several devices *(See Fig. 4)*:

- **Broadband Gateway Device** (BGD) – The residential gateway or modem connected to the IP Service Provider, usually via either cable or DSL.
- **Uni-directional Gateway Device** (UGD) – A one way device that converts broadcast TV to a stream on the home network. For example a DVB-T tuner that converts the stream to IP and sends it wirelessly over the home network.
- **Home Network End Device** (HNED) – The display, controlling and/or storage device for the streams received either via the BGD or UGD.
- **Home Network Node** (HNN) – The device, for example a switch or Wireless Access Point, that connects the home network together.

The Home Network Reference Model, available as a separate DVB Blue Book, is based on work done by the DLNA (Digital Living Network Alliance). DLNA already has existing devices that do stream video over the home network but from sources within the network. The DVB Home Network is the first that integrates both programming from broadcast TV and in-home generated video.

Comparison of Internet video and IPTV

Although IPTV and Internet-based video services share the same underlying protocol (IP), don't let that deceive you: distribution and management of those services are very different.

In an IPTV environment, the SP has a full control over the components that are used to deliver the services to the consumer. This includes the ability to engineer the network's quality and reliability; the bitrate and codec used by the encoder to work best with the limited number of individually managed STBs; the ability to simplify and test the home network components for reliability and quality; and prevention of unnecessary wastage of bandwidth, for example by enabling end-to-end IP Multicast.

Control over the delivery model doesn't exist with Internet video services. For example, IP Multicast deployments on the Internet are still very limited, mostly to research and academic networks. This means that Internet-streamed content services use either simple unicast-based streams between a given source and destination or a Peer-to-Peer (P2P) model which will send and receive data from multiple sources at the same time.

One of the other main differences is the control of the required bandwidth for the delivery of the service. A Service Provider controls the bitrate and manages the QoS required to deliver the service, which allows it to control the buffering needed in an STB to ensure the audio and video decoders don't overrun or underrun, resulting in artefacts being shown to the end user. Internet video cannot control the bitrate so it must compensate by implementing deeper buffers in the receiver or attempting to request data from the closest and least congested servers or nodes, to reduce latency and packet loss. In the peer-to-peer model, lack of available bandwidth from the different nodes, due to limited upstream bandwidth to the Internet, enforces the need for larger and more distant "supernodes" to compensate which, overall, makes the possibility of packet loss higher so increasing the chance of a video artefact.

The decoding devices in the uncontrolled environment of Internet TV also limit encoding efficiency. The extremely diverse hardware and software in use to receive Internet video services tend to limit the commonalities between them. H.264, which is a highly efficient codec but does require appropriate hardware and/or software resources for decoding, is not ubiquitous in today's deployed environment. MPEG-2 video and Adobe Flash tend to be the main video players that are in use, neither being able to provide the same picture quality at the equivalent bitrates to H.264.

Challenges of integration with Internet Video services

Internet Video services are growing very fast. The diversity of the content on offer, the ease of adding new content and the speed with which new services can be added is quite a challenge for managed IPTV services. This leads to the managed IPTV service providers wanting to combine the two types of IP services on the same STB.

The most natural combination is the "Hybrid" model which has both types of services, probably by integrating the peer-to-peer client within the SP's STB. This would allow for collaboration between the two services and would benefit the users by allowing them to view the Internet video content on a TV rather than a computer. The Service Provider would then make sure that the Internet video streams obtain the required bandwidth within the network, perhaps even hosting nodes or caching content within the Service Provider network to improve delivery. They may even transcode the Internet TV content to provide a higher quality service that differentiates itself from the Internet version.

This "Hybrid" model offers collaboration but may still incur some limitations. The Internet TV services might be able to be delivered to the STB but the amount of memory, processing and increased software complexity might make it too difficult within the existing STB designs. This would increase the cost of the unit and therefore impact the business models, whilst competition between such services may lock out specific players from this market due to exclusive deals.

How can a broadcaster get content into an Internet Video service?

First some Internet history: Today, the Internet is known worldwide as a "magical" way to send e-mails, videos and other critical data to anywhere in the world. This "magic" is not really magic at all, but some brilliant engineering based on a network of individual networks, so allowing the Internet to scale over a period of time to cover the entire world, and continue to grow. This network of networks is actually a mesh of administratively independent networks that are interconnected directly or indirectly across a packet switching network based on a protocol (IP) that was invented for this purpose.

The Internet model of a network of networks with everyone connected to everyone individually was fine until the cost and size of bandwidth became too high, and the management of individual links became too difficult. This started the movement towards Internet Exchange Points (IXP) which minimized connections and traffic going across multiple points by allowing the Service Providers to connect to a central point rather than individually connecting to each other. One of the first was at MAE-East in Tyson's Corner in Virginia, USA, but today they exist across Europe with LINX in London, AMS-IX in Amsterdam and DE-CIX in Frankfurt being among the largest and most established ones.

The Internet Exchange Point, by interconnecting directly with other networks, means that data between those networks has no need to transit via their upstream SPs. Depending on the volume and destinations, this results in reduced latency and jitter between two end points, reducing the cost of the transit traffic, and ensuring that traffic stays as local as possible. It also establishes a direct administrative and mutual support relationship between the parties, which can have better control over the traffic being exchanged.

Being at the centre of the exchange traffic means that IXPs can allow delivery of other services directly over the IXP or across private back-to-back connections between the networks. Today, this is how many Voice-over-IP and private IP-based data feeds are exchanged.

This also makes the IXP an ideal place for Broadcasters to use such facilities to establish relationships with SPs to deliver linear or non-linear broadcast services to their end users. The independence of the IXP from the Service Provider also allows content aggregation, wholesale or white-labelled services, to be developed and delivered via the IXP. For example, the BBC in collaboration with ITV is delivering a broadcast TV channel line-up to the main broadband SPs in the UK. They also provide such a service for radio in collaboration with Virgin Radio, EMAP and GCA. This service has been running for a couple of years and has been shortlisted for an IBC 2007 Award within the "Innovative application of technology in content delivery" category.

Quality of Experience

The Quality of Experience (QoE), as defined by ETSI TISPAN TR 102 479, is the user-perceived experience of what is being presented by a communication service or application user interface. This is highly subjective and takes into accounts many different factors beyond the quality of the service, such as service pricing, viewing environment, stress level and so on. In an IP network, given the diversity and multiplicity of the network, this is more difficult and therefore more critical to success than in other transports *(see Fig. 5)*.

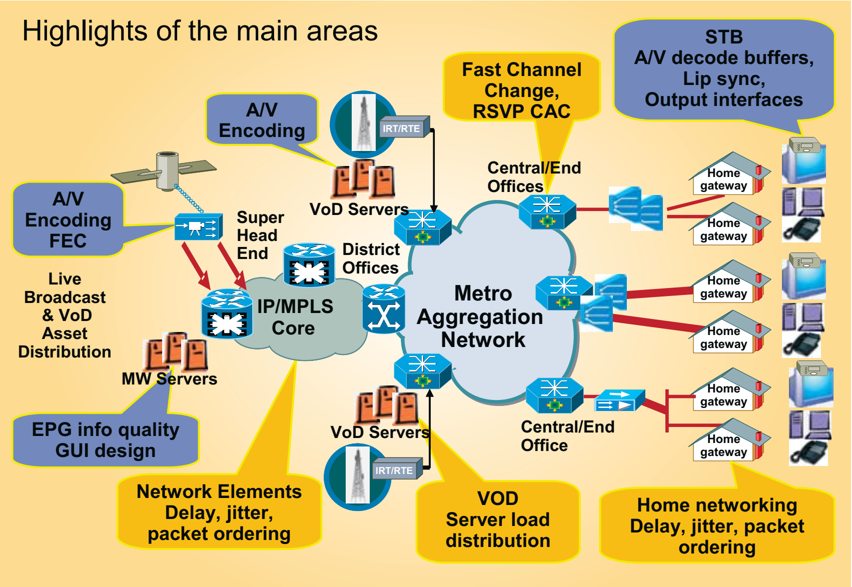

Figure 5
IPTV QoE in the end-to-end model

Subjective and Objective requirements

Subjective measurement systems, such as ITU-R BT.500-11, provide a detailed model for picture-quality assessment by getting a panel of non-expert viewers to compare video sequences and rate them on a given scale. This requires considerable resources to set up and perform the testing, so it tends to be used for comparing video codecs, bitrates, resolutions and encoder performances.

An IP network operator cannot have a team of humans sitting looking at pictures to assess picture quality, particularly with the number of channels these days. They therefore test quality with automated measurement systems which provide real-time monitoring and reporting within the network and services infrastructure. The measurement systems usually use some subjective human input to correlate a baseline that objective measurement methods can be mapped to. An operator usually deploys probes at critical points in the network which report back to the Network Management System (NMS) a set of metrics that will trigger alarms based on predefined thresholds.

When compared to a traditional broadcast environment, video services transported over an IP infrastructure introduce extra monitoring requirements. The two main categories of requirements are:

- **IP transport network**
  Whilst transporting the services, IP packets will cross multiple nodes in the network(s) – possibly subjected to packet delay, jitter, reordering and loss.
- **Video transport stream** (MPEG-2 TS)
  Traditional TS-monitoring solutions must also be used to ensure the TS packets are free of errors.

The two categories are also usually in different departments: the IP transport monitoring is within the Network Operations Centre, and the video transport stream monitoring within the TV distribution centre. One of the keys to a good Quality of Experience in IP is sometimes just good communication and troubleshooting across the different departments.

Finally, although this is beyond the scope of network-based management, additional measurements should be taken into account in a full system, such as the following:

- **Transactional** – GUI and channel change response time, service reliability.
- **Payload (A/V compression)** – Compression standards compliance, coding artefacts.
- **Display (A/V decoding)** – Colour space conversion, de-blocking, de-interlacing, scaling.

Measurement methods

The main measurement methodology for the IP transport network is the Media Delivery Index (MDI) as defined in IETF RFC 4445. MDI is broken down into two sub-components: Delay Factor (DF) and Media Loss Rate (MLR) which are both measured over a sample period of one second. The notation for the index is DF:MLR.

DF determines the jitter introduced by the inter-arrival time between packets. This shouldn't be viewed as an absolute value but is relative to a measurement at a given point in the network. Jitter can be introduced at different points by encoders, multiplexers, bulk scramblers, network nodes or other devices. It is important is to know what the expected DF value should be, which can be determined by a baseline measurement in ideal operating conditions. The value can change dependent on the stream type: Constant Bitrate (CBR) streams should have a fixed inter-arrival time whilst Variable Bitrate (VBR) streams will have a varying value. Once a baseline value has been determined, you normally set a trigger significantly above this value before alerting via an alarm.

MLR provides the number of TS packets lost within a sample period. This is achieved by monitoring the Continuity Counters within the TS.

If the stream contains an RTP header, the sequence number can be used for identifying out-of-sequence or missing packets without the need to examine the IP packet payload. This will reduce the computational requirements and speed up the monitoring process. It is normal therefore to distribute MDI probes across the IP forwarding path to allow supervision on a hop-per-hop basis. This helps troubleshoot potential issues introduced by a specific network element.

To complement the IP packet metrics, DVB-M ETSI TR 101 290 (ETR 290) is used to provide insight within the transport stream itself. This operates in the same way as in a traditional ASI-based infrastructure.

The combination of MDI and ETR 290 delivers a scalable and cost-effective method for identifying transport-related issues. By triggering alarms at the IP and TS level, these can be aggregated and correlated within the NMS to produce a precise reporting tool between different events and their insertion point within the network infrastructure.

Improving QoE with FEC and retransmission

DVB has considerable experience in error-correction and concealment schemes for various environments, so it was natural – given the difficulty of delivering video over DSL – that the IPI ad-hoc group should work in this area. They spent a significant time considering all aspects of error protection, including detailed simulations of various forward error correction (FEC) schemes and quality of experience (QoE) requirements.

The result is an optional layered protocol, based on a combination of two FEC codes – a base layer and one or more optional enhancement layers. The base layer is a simple packet-based interleaved XOR parity code based on Pro-MPEG COP3 (otherwise known as SMPTE standard 2022-1 via the Video Services Forum, see **http://www.videoservicesforum.org/activities.shtml**) and the enhancement layer is based on Digital Fountain's Raptor FEC code (**http://www.digitalfountain.com**). It allows for simultaneous support of the two FEC codes which are combined at the receiver to achieve error correction performance better than a single code alone.

FEC has been used successfully in many instances; however, another technique in IP can also be used to repair errors: *RTP retransmission*. This works via the sequence counter that is in every RTP header that is added to each IP packet of the video stream. The STB counts the sequence counter and if it finds one or more missing then it sends a message to the retransmission server which replies with the missing packets. If it is a multicast stream that needs to be retransmitted then the retransmission server must cache a few seconds of the stream in order to send the retransmitted packets *(see Fig. 6)*.

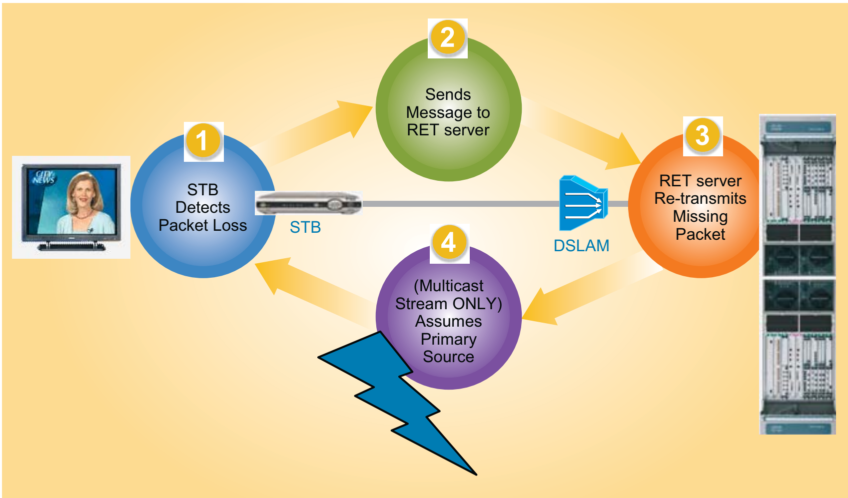

Figure 6
IPTV QoE in the end to end model

Bandwidth reservation per session

One of the advantages of IP is the ability to offer content on demand, for example Video on Demand (VoD). This is resulting in a change in consumer behaviour: from watching linear broadcasts to viewing unscheduled content, thus forcing a change in network traffic. This makes corresponding demands on the IP infrastructure as the number of concurrent streams across the managed IPTV infrastructure can vary from thousands to hundreds of thousands of concurrent streams. These streams will have different bandwidth requirements and lifetime, dependent on the nature of the content which is being transported between the source streamers playing out the session, across the network infrastructure to the STB.

The largest requirement is to prevent packet loss due to congestion, which can be prevented if the network is made aware of these sessions and makes sure enough bandwidth is available whenever setting up a new stream. If there isn't enough bandwidth, then the network must prevent the creation of new streams – otherwise all the connected users along that path will have a degraded viewing experience *(Fig. 7)*.

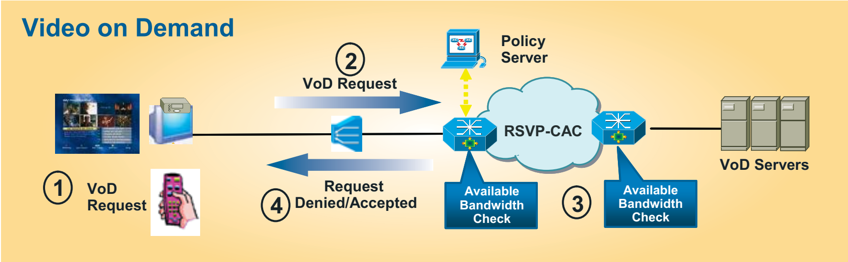

Figure 7
Connection Admission Control

RSVP CAC (based on RFC2205, updated by RFC2750, RFC3936 and RFC4495) allows for per-session bandwidth reservation to be established across the data path that will carry a given session. Step 1 & 2 in *Fig. 7* show the VoD session starting between the STB and the middleware. The authorization credentials will be checked to make sure that the customer can play the content, based on a set of criteria such as credit, content rating, geography and release dates. Once these operations are authorized by the middleware and billing system, the middleware or VoD system manager identifies the VoD streaming server for this session. In step 3, the server initiates a request for an RSVP reservation path between the two end points across the RSVP-aware network infrastructure. Finally, in step 4, if the bandwidth is available then the session can be initiated; otherwise a negative response will be sent to the middleware to provide a customized response to the customer.

Conclusions

Delivery by IP of broadcast-quality video is here today and is being implemented by many broadcasters around the world. The nature of IP as a connectionless and non-deterministic transport mechanism makes planning, architecting and managing the network appropriately, which can be done with careful application of well-known IP engineering. When the IP network is the wider Internet, the lack of overall control makes guaranteed broadcast-level quality difficult to obtain, whereas on a managed IP network, Quality of Service techniques, monitoring and redundancy can be used to ensure broadcast-level quality and reliability.

The techniques to monitor video are similar to the ones used for any MPEG-2 transport stream. However, these need to be related to the IP layer, for example using MDI, as debugging the problem will often require both network and video diagnostics.

Thomas Kernen is a Consulting Engineer working for Cisco Systems' *Central Consulting Group* in Europe. He works on Video-over-IP with broadcasters, telecoms operators and IPTV Service Providers, defining the architectures and video transmission solutions. Before working for Cisco, he spent ten years with different telecoms operators, including three years with an FTTH triple play operator, for which he developed their IPTV architecture.

Mr Kernen is a member of the IEEE, SMPTE and active in the AVC group within the DVB Forum.

Jeff Goldberg is a Technical Leader working for a Chief Technology Officer within Cisco. He has been working on IPTV, IP STB design and home networking since 1999, and has been working for Cisco since 1994. He was part of the founding group of DVB-IPI and has been working on it ever since, particularly on the home networking, reliability, Quality of Service and remote management parts.

Before working for Cisco he designed handheld devices and PC software.