Technical Report 009

Peer-to-Peer (P2P) Technologies and Services

EBU **TECHNICAL**

Your reference in media technology and innovation

Geneva
May 2010

EBU•UER

° Issued by EBU TECHNICAL
L'Ancienne Route 17A
1218 Le Grand-Saconnex
Switzerland

tech@ebu.ch

A year in the Internet business is like a dog year, equivalent to seven years in a regular person's life. Vint Cerf, Internet pioneer

Contents

DISCLAIMER:

This document is written from the point-of-view of public broadcasters. It considers Peer-to-Peer (P2P) as a legal distribution mechanism for transport/distribution/delivery of legal TV/video content over the open Internet. In the past P2P protocol has been tainted by its association with copyright piracy. Many detractors still associate P2P with illegal file sharing. It is now possible to remove its bad name and consider P2P as a legal mechanism.

Acknowledgements

The present document is the result of collective work of Members of Project Group D/P2P (Peer-to-Peer) between February 2006 and March 2010. During these four years the group was able to review a large number of documents, perform some technical tests, organised several P2P demonstrations at IBC conventions, run a major trial on EBU Media Portal and evaluate a number of commercial propositions. The Group did not consider only technical issues but tried to address also some non-technical matters such as operational, business, legal and regulatory. The Group produced a number a valuable contributions to external bodies such as the DVB Consortium and the EC funded project *P2P-Next*.

The Group would like to express its gratitude to the EBU management for support and encouragement and the EBU Member organisations that were represented in the Group.

Particular thanks should go to those numerous EBU and non-EBU experts who actively contributed to this document, in particular:

Jari Ahola (VTT, Finland)
Anne-Catherine Berg (EBU Legal Dept.)
Stefan Clyck (VRT)
Frank Hoffsümmer (SVT)
Frans de Jong (EBU Technical)
Christian Klöckner (WDR)
Peter Lanigan (Philips)
Michael Moelants (VRT)
Johan Pouwelse (TUD Delft)
Mark Stuart (Pioneer)
Stephen Alstrup (Octoshape)
Thomas Buchholz (Deutsche Telekom)
Matthias Hammer (IRT)
Thomas Elton Jensen (DR)
Rainer Kirchknopf (ZDF)
Olli Karonen (Nokia)
Deon Noerskov Mathiesen (Octoshape)
Enrico Magli (Politecnico di Torino)
Heijo Ruijsenaars (EBU Legal Department)
Rafal Wiosna (TVP)

The document was compiled by Franc Kozamernik (EBU Technical). It was proofread and edited by Ed Wilson and Roger Miles (EBU Technical).

General Glossary

ADSL	Asymmetric Digital Subscriber Line
AS	Autonomous System (RFC 1930)
BCG	Broadband Content Guide
CDN	Content Distribution Network
CE	Consumer Electronics
CIF	Common Intermediate Format is a video format used to standardize the horizontal and vertical resolutions in pixels(typically 352 x 288) of YC_bC_r sequences in video signals.
DHCP	A protocol that allows networked devices to be assigned a unique IP address automatically from a pool of unused IP addresses.
DHT	Distributed Hash Table
DNG	Home Network Gateway
DNS	Domain Name System
DoS	Denial of Service
Downstream	In a direction towards the end-user
DRM	Digital Rights Management
DVB	Digital Video Broadcasting (Consortium)
DVR	Digital Video Recorder (usually hard-disk based)
EPG	Electronic Programme Guide
EPS	EBU P2P System
ESI	Edge Side Includes: is a small markup language for edge level dynamic web content assembly. The purpose of ESI is to tackle the problem of web infrastructure scaling
FLUTE	File Delivery over Unidirectional Transport
GL	Geolocation
HDTV	High Definition Television
HNED	Home Networked Device (DVB denomination)
HTML	Hypertext Mark-up Language
HTTP	HyperText Transfer Protocol
IANA	Internet Assigned Numbers Authority
ICAP	Internet Content Adaptation Protocol
IGMP	Internet Group Management Protocol (RFC 3376)
Internet-TV Device	An Internet-TV Device is defined as a consumer appliance (i.e. Set-Top Box, TV set, media-enabled PC, etc.), that is connected to the Internet via a broadband connection (such as ADSL, cable, fibre, etc) and which uses typically a remote control. The Internet-TV device shall be implementable with today’s chip sets for consumer digital TV devices including storage for the P2P system.
IPTV	Internet Protocol Television (managed service)
ISP	Internet Service Provider
LAN address	The private, internal IP address that locates a computer on a Local Area Network. A LAN IP address is not visible to users outside of the LAN. As described by RFC 1918, the following ranges are designated as reserved IP addresses for private LANs: 10.0.0.0 - 10.255.255.255 172.16.0.0 - 172.31.255.255 192.168.0.0 - 192.168.255.255

Legal Content	Content which is distributed and acquired according to the positive legal requirements (e.g. copyright)
MPEG	Moving Picture Expert Group
Multicast	This is a technique for one-to-many communication over an IP infrastructure in a network. It is often employed for streaming media and Internet television applications.
NAT	Network Address Translation: The changing of the source or destination IP address for a data packet. This usually occurs when one is behind a firewall or router, where it translates IP addresses so that multiple computers can exist on a LAN with while using the same WAN IP address.▫
NAT-PMP	NAT Port Mapping Protocol: An alternative to UPnP created by Apple, Inc. NAT-PMP is not as widely supported as UPnP is, and uptake of the protocol has been limited to Apple, Inc. products only thus far.
OPES	Open Pluggable Edge Service
P2P Flow	A logical data flow being exchanged between different elements of the P2P-Internet-TV Content Delivery system
P2P	Peer-to-Peer
PC	Personal Computer
PIP	Picture-In-Picture
QCIF	"Quarter CIF" - this implies that the height and width of the CIF frame are halved.
QoE	Quality of Experience (latency, download time, etc)
QoS	Quality of Service
RSS	Really Simple Syndication is a family of web feed formats used to publish frequently updated works such as blogs, news headlines, audio and video.
RSS feed	A file that is updated so that it delivers information and content in such a way that allows one to track updates quickly and easily.
RMS	Remote Management System (DVB)
RTP	Real-time Transport Protocol (RFC 3267)
RTSP	Real-time Streaming Protocol (RFC 2326)
SD&S	Service Discovering and Selection (DVB)
STB	Set-top box
Static IP Address	An IP address that does not change (remains static) across multiple sessions. A static IP address is necessary in port forwarding, as ports are usually forwarded to a specific IP address, where the rule does not change even if the computer's IP address does.
UDP	User Data Protocol (RFC 768)
Unicast	This is the sending of messages to a single network destination on a packet switching network. Unicast messaging is used for all network processes in which a private or unique resource is requested.
UPnP	A protocol that allows devices on a network to communicate with each other seamlessly. In the case of µTorrent, UPnP is used to forward a port on a router without the need to open the port manually.
Upstream	In a direction away from the end-user.
Users' context	Access network capability, storage capability
WAN	A computer network that covers a large geographical area. A WAN connects multiple LANs together. The Internet is an example of a WAN. ▫
xDSL	Any type of Digital Subscriber Line, including ADSL and VDSL.

Glossary of some specific BitTorrent related terms

Announce	The act of connecting to a tracker to update it on your status, and to obtain information from it as well, including (but not limited to) an updated peer list.
Availability	The number of complete copies of the torrent contents that are distributed in the part of the swarm you're connected to. The amount of the torrent contents you currently have is included in the availability count. A swarm with no seed and with availability below 1.0 will likely be unable to finish transferring the complete torrent contents.
Block	The units of data that comprise a piece. Because blocks do not directly affect whether torrent contents are considered to be finished transferring, it is not seen as an appreciable unit of data with regards to BitTorrent like the piece is.
Broadcatching	The act of downloading content from an RSS feed¤
Choked	This word describes the state of a BitTorrent connection. When a connection is choked, it means the person who is supposed to be doing the uploading on the connection does not want to send anything. This generally happens when the uploader's upload slots are full.
Client	The application a user is using when connected to a swarm. In this case, the application being used to connect to swarms is µTorrent, so it is the client.
DHT	Distributed hash Table: A distributed tracker that works similarly to a regular tracker in that you announce to it and get back a list of peers that are transferring the same .torrent file as you. Because DHT is distributed, there is no single point of failure, so even if a single node disconnects from DHT, the tracker will continue to work (unlike with normal trackers, where if the server goes down, it becomes unusable). DHT can be thought of as a backup tracker.
Disk cache	A feature that makes use of available memory to stores data for quicker access as well as ease disk thrashing. The use of a disk cache will cause an increase in memory usage in return for improved performance.
Disk thrashing	When a storage disk gets accessed very frequently. Extended disk thrashing may lead to hard drive wear and tear, shortening a drive's life.
Endgame mode	A change in the piece requesting strategy that occurs when a download is near completion during which the client requests pieces from all connected peers rather than requesting a piece from one peer at a time in the normal operating mode. Endgame mode is used because download rates often slow down considerably as a torrent job nears completion due to the tendency for the remaining pieces to be downloaded from peers with saturated connections. By requesting data from all peers rather than waiting for a single peer, such a bottleneck can be bypassed. This mode is not used during normal operating modes because of the large amount of overhead it potentially generates in sending requests to all peers.
Ephemeral port range	A range of port numbers automatically allocated by the operating system for use by any application on the system with network access. Ports in the ephemeral port range are typically used to make temporary outgoing connections. By default, the ephemeral port range is defined from ports 1024 to 5000, but this is configurable via the Windows Registry, and may vary from (operating) system to (operating) system.
Hash	A 'fingerprint' of data assumed to be unique to the data. Because of the assumed uniqueness of the data, it is used to verify that a piece of data is indeed uncorrupted (since the corrupted data's hash would not match its expected hash).
Hash check	The comparing of a piece of data's hash with a reference hash in order to verify the integrity of the piece of data.
Index	A site that lists .torrent files available for download.
Initial seeding	Also super-seeding: A method of seeding that attempts to decrease the bandwidth load for the initial seeder. With normal seeding methods, the initial seeder typically has to upload 150% to 200%, or even more, of the original data in before a full copy of the data has been distributed into the swarm. With initial seeding, the initial seed attempts to get the rarest pieces out instead of uploading identical pieces repeatedly, often

	lowering the initial upload requirement to 105%. Initial seeding does not necessarily improve upload speeds or decrease seeding time. It should be used only if you are the sole seeder on the swarm, and if there are at least 2 peers connected. Generally, initial seeding should not be used by people with high upload speeds.
Interested	This word describes the state of a BitTorrent connection. When a peer is interested, it means the peer is interested in the data that the peer on the other end of the connection has, and is willing to accept data from the other peer.
LPD	Local Peer Discovery: A method by which Torrent attempts to discover new peers local relative to your computer's network. Local Peer Discovery makes use of IP Multicast.
Leecher	A person who downloads, but fails to reciprocate the generosity of others by not sharing back. The word 'leecher' carries a strong negative connotation.
Micro transport Protocol	A UDP-based reliable transport protocol designed to minimize latency, but maximize bandwidth when latency is not excessive. This alleviates the bandwidth saturation that often occurs to BitTorrent users while they are transferring data and using the Internet for other purposes.
Optimistic unchoke	When a client tries to start a transfer on a previously choked connection in hopes that the connection becomes unchoked.
PE	Protocol Encryption: A specification designed jointly by Vuze and µTorrent developers, created as an attempt to bypass throttling and/or blocking of BitTorrent traffic by ISPs. There are different methods of encryption, ranging from full encryption of all of the data, to partial encryption (header encryption only, not unlike with PHE, although it is still not as easily detected as PHE).
PEX	Peer Exchange: A feature to exchange peer lists with other peers that support the same PEX implementation (generally limited to peers using the same BitTorrent client). By exchanging peer lists, it's possible to find peers not included in the peer list supplied by the tracker.
PHE	Protocol header Encryption: An old method of encryption created by the BitComet developer that encrypted only a part of the data (the header) in an attempt to bypass ISP throttling and/or blocking of BitTorrent traffic. Because its specification was designed in a relatively poor manner, ISPs were able to detect it with little trouble, rendering it useless.
Piece	The smallest appreciable unit of data in BitTorrent. The size of pieces can be different depending on the .torrent file in question.
Piece distribution	The general distribution of the pieces across the swarm. BitTorrent is generally most efficient when piece distribution is random, with minimal 'clumping' of pieces available in the swarm.
Poisoning	The act of intentionally feeding invalid data into the swarm, resulting in hashfails for peers receiving the invalid data. Outfits with (or hired by other entities with) anti-P2P agendas are the most common sources of swarm poisoning.
Port forwarding	The act of passing data on the forwarded port from one network device to another. In most cases regarding BitTorrent, port forwarding refers to the forwarding of connections from a router to a specific computer attempting to listen on that port.
Private flag	A piece of information stored in a .torrent file that tells any BitTorrent client that recognizes the flag to disable DHT, LPD, and PEX for that specific .torrent. The 'private' flag is typically used in .torrent files served by private trackers as a method of keeping a swarm isolated from people who aren't members of the private tracker.
Private tracker	A tracker that requires users to log in to use it. Private trackers typically enforce ratio requirements (by banning users whose ratios are too low) in order to prevent or minimize the leeching that is prevalent on many public trackers.
Public tracker	A tracker that is open for anyone to use (as opposed to private trackers, where only people who hold accounts can use the tracker).
Scrape	The grabbing of statistics (number of seeds and peers) from a tracker regarding a specific swarm.

Seed	A peer with 100% of the data in the torrent contents.
¤Seeding	The act of being connected to a swarm as a seed.
¤Snubbed	This word describes the state of a BitTorrent connection. A connection is marked as snubbed when the client has not received anything over the connection for an extended period of time.
Swarm	The collective group of peers (which include seeds) that are connected by a common .torrent file.
Throttle	A term used to refer to the intentional slowing down of transfer rates (download and/or upload), typically used in the context of ISP throttling. ¤
Torrent	A small file containing metadata from the files it is describing. In other contexts, it is sometimes used to refer to the swarm connected around that small file.
Tracker	Something that a client connects to in order to share its IP and port, as well as obtain information, including peer lists.
Upload	The act of transferring data from your computer onto another.
Wasted	Data that is tossed out either because it hashfailed, or because it was redundant data that the client had already downloaded.
Web seed	A seed that is basically a regular web server hosting the requested file. BitTorrent clients that support web seeds use them like any other seed, and can request data segments from the server much like requesting pieces from an ordinary seed. The use of web seeds ensures that a torrent swarm will never die as long as the file being seeded is left intact on the server and the server does not go down.

Peer-to-Peer (P2P) Technologies and Services

Keywords: P2P, Peer-to-Peer, Internet Radio, Internet TV.

Executive Summary

This document summarizes the studies carried out by EBU Project Group D/P2P (Peer-to-Peer). The Group was tasked to investigate the new Internet-based media delivery mechanism based on grid-type (distributed) technologies and assess its suitability for broadcasters to distribute large-scale high-quality audio-visual streams and files to the general public in a cost efficient, reliable, and user-friendly manner.

Background: The EBU's involvement in P2P goes back to February 2006 when the EBU organised a workshop 'From P2P to broadcasting'. This event was held at a time when more and more broadcasters were just starting to use regularly the Internet to distribute their live or on-demand programmes. At that time, P2P technology was still largely associated with illegal file sharing and tainted by its association with copyright piracy as a result of the illegal activities of Napster, Kazaa, Glockster, etc. The usage of Internet was exceedingly expensive for broadcasters, as they were charged by Internet Service Providers (ISPs) for each delivered stream (in gigabytes). Broadcasters therefore became victims of their own success: The more popular their content, the more they needed to pay for the distribution of their content. The situation became increasingly untenable.

The EBU workshop was a big success, with more than 100 participants, including representatives from more than 30 TV & Radio EBU Members and associate Member organisations. This landmark event showed that P2P could actually become a useful technical solution, potentially more cost efficient and robust than any other traditional Internet distribution method. The EBU members showed a significant interest in studying P2P technologies and were prepared to run experiments and trials in order to gain practical experience. Members felt it necessary to improve the public image of P2P and consider P2P a legal mechanism for the transport and distribution of legal TV/radio content. The workshop helped broadcasters to better understand what potential impact P2P could have on business, revenue, and distribution models.

EBU Project group D/P2P: As a result of this workshop, the EBU Delivery Management Committee (DMC) set up the D/P2P Project Group which held an exploratory meeting in April 2006. In the three years of its existence, the group held nine meetings, the last one on 6 and 7 May 2009 in Geneva. The Group counted as many as 96 members from all constituencies: public and commercial broadcasters, telecoms, ISPs, CDNs, P2P operators, technology providers, manufacturers, chip developers, universities, and legal experts. The group worked very closely with EC-funded projects and universities and contributed towards standardisation within DVB. In the course of its existence, it examined almost 150 input documents. The Group organised a real-life trial of EBU P2P media portal involving several EBU TV and radio stations. The Group provided support to the EBU's Eurovision Song Contest.

Document structure: This document is a comprehensive report on P2P, as seen by the EBU members. The first four chapters provide a tutorial on P2P technology and positions P2P amidst other distribution technologies such as Unicast, server-based content distribution networks, IP Multicasting and others.

Chapter 5 describes the P2P business environment and provides some P2P business models and cost considerations. The P2P services can be provided by content providers, service providers, or by special entities, P2P service providers. The latter ones play the same role as conventional CDN (Content Distribution Network) providers.

Chapter 6 offers some typical possible P2P use cases. These use cases are given as examples of possible usages and do not imply that P2P is the only possible technical solution for a given task – other solutions may also be used (and may often be even more efficient than P2P).

Chapter 7 is a key part of the report – it gives the technical and commercial requirements for a viable P2P system from a broadcast perspective. These requirements were submitted to the DVB Project and will represent a basis for a future P2P standard.

Chapter 8 provides a brief review of some existing commercial P2P services and systems that have been successfully deployed in the Internet market. These systems are all proprietary and are not interoperable. EBU members have evaluated many of them.

Chapter 9 is a collection of P2P trials, experiments, and experiences in which the Group (or its members) was involved. The experience running the EBU P2P Media portal available from the EBU web site is also described.

Chapter 10 describes P2P social networks in which peers share common activities, for instance exchange audio files, chat, socialise, etc. Social networks are an important mechanism for building communities around broadcast TV programmes. They may represent a way to attract new, particularly younger, audiences.

Chapters 11 and 12 were prepared by the manufacturer and telecom members of the group, respectively. These chapters significantly enrich the document by including some views and opinions from non-broadcasters. In many ways they provide a somewhat critical account of P2P.

Chapters 13, 14 15 and 16 address some ancillary topics, that are not linked directly to P2P technology but are important for P2P operations: geolocation, metadata, digital rights management (DRM), and the legal requirements. Geo-filtering and DRM may be important for copyrighted materials and premium services. It is recognised that neither DRM nor geo-filtering are able to ensure absolute security and protection, and can be hacked if sufficient effort is made. Nevertheless, they are seen by many broadcasters and content providers as indispensible 'best effort' protective measures that need to be applied. Metadata describe the content conveyed via P2P and the context in which the content is generated, transported and used.

Chapter 17 discusses the limitations of P2P networks if users utilise asymmetric lines with limited upstream capacity. The asymmetry of DSL lines is often a critical limiting factor determining the efficiency of the P2P system. The chapter gives some quantitative analysis of the problem and makes an estimate of how much capacity should be provided by a central server (in comparison to the capacity provided by peers).

Chapter 18 looks into the future: what are the future trends in P2P research and development. It is evident that P2P is a novel technology with a huge potential for future developments and refinements. This chapter lists the main projects involved in P2P R&D and how P2P will improve further in terms of its efficiency, robustness, availability and user friendliness. The fact that so many R&D projects are involved in P2P highlights the interest for this technology. It may be useful to single out a particular aspect: a conjunction of P2P with scalable video coding (SVC or MDC) enabling an adaptive P2P system. Developing network topology aware P2P systems is another interesting research topic.

Chapter 19 contains the SWOT analysis, which outlines the P2P system's strengths, weaknesses, opportunities and threats, as perceived by broadcasters.

Chapter 20, 'Main challenges and open questions', discusses the remaining issues that may still require further studies and careful considerations, should P2P become a viable distribution system to be used for broadcast purposes.

Key messages: Also to be found in Chapter 20. They can be summarised as follows:

'P2P is an attractive novel technology, which may potentially revolutionize the business of media distribution over the open Internet. P2P technologies are still being developed and refined and will definitely find their place in the media distribution landscape. P2P is highly disruptive but is likely to bring distribution costs down. Indeed, it has already brought the prices down by an order of magnitude. In order to exploit full potential of P2P, broadcasters need an open, standardised, license-free P2P solution that could be embedded in the range of consumer-electronic devices. It should be pointed out that P2P technologies will not replace conventional terrestrial, satellite and cable broadcast networks. For Internet distribution the trend now is to use a combination of centralised (e.g. CDN) and decentralised (e.g. P2P) approaches.'

This report addresses some particular issues related to P2P distribution of media over the Internet, including:

- **Internet distribution costs**: P2P is able to reduce significantly the server cost, bandwidth, and network load. Practically every individual user can now become a 'broadcaster'. The use of open source BitTorrent server by NRK and the BBC is an example of how broadcasters can reach the general public with their content almost for free.
- **Cost incurred by the user**: There is also a question of cost incurred by the user. The dominant view of the D/P2P group is that most people would be ready to pay for better quality. A business concept of tiered services may be one way forward. For example, users may be prepared to pay for high-quality services, whereas normal, baseline webcasts at, say, 350 kbit/s, would be for free.
- **New business models**: P2P has been instrumental in introducing new flexible business models for the end users, technology and service providers.
- **Social networks**: A P2P system can be regarded as a social network of 'friends' sharing the same interests and tastes – P2P may thus have an important social, community impact. This aspect should not be neglected, as broadcasters endeavour to build stronger links to their viewers and listeners.
- **Short and long tails**: P2P is particularly a suitable solution for distributing very popular content, as a large number of peers can participate in serving the content to other peers. However, P2P is also suitable to distribute 'long tail' content. Broadcasters may use it to target niche audiences. Such a P2P solution is undoubtedly cheaper than setting up new terrestrial broadcast networks to cover a few viewers and listeners. P2P delivery can potentially save significant amounts of broadcasting spectrum which is a very scarce resource.
- **Audience measuring**: One important advantage of the Internet is the possibility to measure audiences very accurately and not merely estimating them by using limited audience samples. This may radically impact the business models including personalised advertising used in the open Internet.
- **P2P embedded devices**: In 2008, Octoshape released an announcement that it launched in association with an American company CaptiveWorks a commercial P2P-enabled set-top box[1]. This product, along with the DVB efforts to standardise P2P-enabled Internet TV consumer electronic devices, may establish the Internet as one of the principal delivery

[1] Pioneer successfully demonstrated a TV STB using the Tribler P2P client at IBC 2008. This development was performed within the framework of the *P2P-Next* project. An improved version of the Pioneer STB was shown at IBC 2009

mechanisms of broadcast-quality television signals to the hybrid broadcast/broadband set-top box and integrated digital TV. DVB has launched a Study Mission in order to establish whether or not P2P compares favourably to other Internet distribution mechanisms.

- **Upload contribution of 'passive' peers**: It should be pointed out that P2P is not a panacea. This report provides a critical assessment of P2P weaknesses and identifies numerous open issues. There may be circumstances where P2P solutions may be less than optimal. Typically, such a case may occur in the countries where Internet consumption is limited to a certain amount of download and upload capacity. The concept of 'passive contributions' provided by stand-by peers to increase overall network upload capacity would help in such cases.
- **Centralisation**: One open question is whether the P2P system could be totally decentralised. The view of the Group is that content and network security could be better administered and controlled if a limited central network element managed by a content or service provider would be used.
- **Quality of service**: Some effort should be made to improve QoS (Quality of Service) and QoE (Quality of Experience) of Internet TV services. This report does not address the technically solutions to improve the all-important QoS-related issues, such as latency, channel-change time, buffering, resilience to packet losses and jitter, service availability, ability to cope with flash crowds, etc . Nevertheless, these issues will have to be addressed by an appropriate technical body, e.g. DVB.
- **Net neutrality**: The question of 'net neutrality' is still pending. This matter is highly political and should be resolved at the highest political level in Europe. It is clear that very tight rules will be required in order to prevent ISPs from shaping and filtering Internet (including P2P) traffic. Such traffic shaping may be potentially detrimental to broadcast content distribution. The distribution of professional broadcast content should not be adversely affected by any inappropriate ISPs activities.
- **Synergy of CDN and P2P**: The final observation is that P2P may be used as an add-on to conventional CDNs; rather than increasing the number of edge servers, one can use P2P to bring the signal to the general user through the existing infrastructure consisting of end-user devices (PCs, STBs, game consoles, mobile devices, etc).

It is recommended that EBU members continue considering P2P as one possible choice for distributing audio/video/data over the Internet. Members are encouraged to continue gathering operational experience by running P2P-driven services and sharing i he results with other Members. The future is uncertain and it is not possible to predict whether or not P2P will ultimately be successful in the media distribution market. The degree of success will probably vary across different European markets and will heavily depend on the regulatory arrangements adopted in the countries concerned. However, a significant impetus for the promotion of P2P may occur with the advent of large-scale affordable commercially-available P2P-enabled televisions and STBs.

1. Introduction

The Internet was designed to provide relatively short duration connections for delivery of web pages. It was never designed to carry huge video files and deliver them to a large number (potentially counting several millions) of concurrent users. It is relatively easy to carry a web page (some tens of kbyte), it is totally another matter to deliver a film which is equivalent to 100000 web pages (several Gbyte) [1]. If a web page fails to load properly, you click the refresh button and fix the problem easily. For video, however, long duration connections lasting several hours may be required. Clicking the refresh button in the middle of a film would re-start a film, and this would be a major inconvenience. Video delivery also requires larger and more consistent network bandwidth availability.

Today, the broadband Internet has become a consumer utility akin to the electrical power supply. More, it has become a useful complement to the traditional distribution channels of radio and television programmes, such as satellite, cable, terrestrial networks. During the forthcoming years, the Internet is likely to become a major means to bring radio, television and multimedia to the home and to users on the move[2]. Both back-bone Internet connections and access links to the home are getting faster and faster. Availability and deployment of broadband access is rapidly expanding worldwide. The ability of the network to provide reliability and low latency is improving and the cost of transmission is decreasing.

The key question is whether the Internet will be able to cope with the ever increasing traffic volumes or not. The major increase of traffic is due to all kinds of video content: live broadcasts, movies-on-demand, catch-up TV and user-generated content. Having initially accepted low-resolution small-screen video, the public now expects standard or even high-definition (HD) television quality content without buffering and delay. Better quality of digital media drives audiences and their viewing times constantly up. The growth in Internet media consumption presents significant economic and business challenges for content and network service providers.

Compared to other distribution mechanisms, the Internet still lags behind in terms of quality of service, scalability and cost [2]. The Internet channel generally is characterized by long latency, jitter and packet loss. This results in the need to provide retransmissions and long buffers at the player end. Internet can only provide a ‘best effort’ quality unless special provisions are taken. Furthermore, in order to reach a lot of members of the audience, it is often necessary to limit video resolution to QCIF or CIF. Also, Internet servers can often accommodate only a small number of simultaneous users. Scaling to large audiences is expensive for content providers and broadcasters, as they are charged by the Internet services providers (ISPs) per each stream.

This scaling problem means that today's Internet is a very costly medium. The more users there are, the higher the cost. Broadcasters are then penalized if their content is very popular, as they will have to pay more for distribution. The Internet business models are not yet clear and vary from country to country.

For Internet to become a successful medium, broadcasters need a reliable service delivery (e.g. no glitches, no interruptions, acceptably short delay, zapping time), high intrinsic quality (e.g. HDTV, multichannel audio, extensive metadata), scalability (large number - several millions - of concurrent users) and low transmission cost (ideally, cost should be independent of location, time, quality and number of users). Internet TV streaming and video downloading using the traditional server-client distribution mechanism is now well in place and a large number of commercial services are already deployed worldwide. Streaming traffic is among the fastest growing traffic on the Internet. In a recent white paper, Cisco predicts that by 2012, 90% of all Internet traffic will be

[2] The BBC iPlayer generated 94m views in December 2009. About 4m connected TVs are expected in UK homes by the end of 2010.

video. Compared to server-client distribution, using a peer-to-peer (P2P) distribution for downloading and particularly live streaming is becoming mature, however numerous technical, operational and legal challenges still need to be addressed.

There are two basic architectures for delivering streaming traffic on the global Internet: the client-server paradigm and the peer to peer (P2P) paradigm. A particular advantage of the P2P paradigm over the client-server paradigm is its scalability. As an example, PPLive, one of the largest P2P streaming vendors, was able to distribute large-scale, live streaming programmes such as the CCTV Spring Festival Gala to more than 2 million users with only a handful of servers. CNN reported that P2P streaming by Octoshape played a major role in its distribution of the historical inauguration address of President Obama. It is well demonstrated in practice that P2P streaming can deliver videos encoded at a rate of about 700 kbit/s, in the presence of rapid user joins/leaves, with positive user experiences.

P2P streaming is seeing rapid deployment. Large P2P streaming applications, such as PPLive, PPstream and UUSee, each has a user base exceeding 100 millions. P2P streaming traffic is becoming a major type of Internet traffic in some Internet networks. For example, according to the statistics of a major Chinese ISP, the traffic generated by P2P streaming applications exceeded 50% of the total backbone traffic during peak time in 2008. There were reports that major video distributors such as YouTube and Tudou are conducting trials of using P2P streaming as a component of their delivery infrastructures. Given the increasing integration of P2P streaming into the global content delivery infrastructure, the lacking of an open, standard P2P streaming protocol becomes a major missing component in the Internet protocol stack. Multiple, similar but proprietary P2P streaming protocols result in repetitious development efforts and lock-in effects. More importantly, this leads to substantial difficulties when integrating P2P streaming as an integral component of a global content delivery infrastructure. For example, proprietary P2P streaming protocols do not integrate well with existing cache and other edge infrastructures.

In the Internet traffic race, P2P enjoys mixed fortunes. For years, P2P traffic eclipsed HTTP traffic as broadband users slurped down music and movies. In June 2007, however, Ellacoya Networks reported that HTTP traffic accounts for 46 percent of all broadband traffic in North America, whereas P2P applications account only for 37 percent. The surge in HTTP traffic was largely due to huge popularity of YouTube HTTP video streaming services.

Broadcasters may wish to use P2P because it scales well to provide services to a large number of concurrent peers [3]. The server investment and maintenance cost as well as the server bandwidth required are much lower than with other content distribution mechanisms. Consequently, P2P may be more affordable and cost-effective than alternative content distribution systems. As the P2P systems are generally geographically distributed across the Internet, there is no single point of service failure. Therefore, P2P may ensure a reliable user service experience.

Nevertheless, broadcasters need to evaluate all pros and cons for P2P, investigate which P2P solutions are best suited for the applications required and how best to integrate the P2P solutions chosen with the traditional content delivery systems already in use.

These are the objectives of this EBU document.

The following extract from an article published in the Economist on 4 March 2010 (see box below) provides an excellent introduction to explaining the reasons why public broadcasters are interested in experimenting with P2P technology.

Video on the Internet: Why are public broadcasters experimenting with the 'peer-to-peer' technology beloved of online pirates?

Last year Norway's public broadcaster, NRK, filmed a stunning seven-hour train ride between Bergen and Oslo, shot entirely in high-definition video. Over one million Norwegians watched the film on television. But NRK faced a challenge in reaching a larger audience. How could it distribute the hard-drive-busting 246 gigabytes of raw footage to a global audience without bringing its servers to a grinding halt? The broadcaster made a somewhat surprising choice: it turned to BitTorrent, a peer-to-peer (P2P) Internet service best known as a means of sharing pirated movies and music.

Some at NRK worried that using a system associated with piracy would generate negative publicity. But BitTorrent itself is value neutral. It is a uniquely efficient distribution method that lets broadcasters 'seed' the Internet with one or two copies of their massive media files. It then relies on end users (called peers) who request the file and receive different pieces of it. To assemble a complete version of the file, these peers then share their pieces with each other (hence 'peer to peer'). It takes a while, but the broadcaster does not need expensive server farms or fat data pipes to deliver massive files to viewers anywhere in the world.

Several other public-service broadcasters have also been experimenting with P2P distribution, probably because they are relatively insulated from commercial pressures. In 2008, for instance, Canada's CBC used BitTorrent to distribute a high quality, unprotected version of a prime-time reality show called 'Canada's Next Great Prime Minister'.

The British experience, however, reveals that P2P distribution is no panacea. When the BBC rolled out iPlayer, its television catch-up service, in 2007, it initially relied on P2P techniques to offload the burden of supplying so many large files. At the time, the BBC's Anthony Rose believed that P2P was the only way to provide the service. Yet one year later, the BBC had switched to streaming content directly from its own servers. The reason? A 90% drop in the cost of bandwidth.

There were other problems with P2P. Many users did not realize that they had become video distributors by installing iPlayer, and complained about slower computers and upload speeds, or about exceeding a monthly data cap. From an ease-of-use standpoint, there is another drawback: most P2P systems work only with complete downloads, not with streaming video of the kind on YouTube. And if there is one thing that unites Internet users, it is impatience: multi-hour waits to download an episode of 'Top Gear' are intolerable.

None of this has stopped the European Union from investing in P2P as a route to the remaking of broadcasting. In 2008 it put €14m of funding into a four-year project called *P2P-Next*. It is an ambitious international undertaking backed by research institutes, companies and some broadcasters, including the BBC.

Television sets are already starting to stream Internet video from companies like Netflix in America. If *P2P-Next* has its way, televisions, computers and mobile phones will all support a standardized P2P network for streaming content distribution in the future. The idea is to create the world's best video service: from anywhere on earth, users can then use a standard protocol to pull up any video, at any time and on any screen.

Despite the hitches, P2P may yet be the right way to do this. Bandwidth may be cheap now, but many worry that the explosive growth of video, much of it in high-definition formats, could soon clog up the Internet. Traffic from legal online video sites like Hulu, iPlayer and YouTube has surged in recent years, increasing from 13% of all Internet traffic in 2008 to 27% in 2009. Furthermore, the current infrastructure of the Internet is not suited to the simultaneous transmission of live events to millions of viewers.

Johan Pouwelse, *P2P-Next*'s scientific director, imagines a brave new world for broadcasters in which interconnected television sets with P2P sharing can give any television station global reach. Barriers to market entry will be low, ensuring healthy market competition, he says. 'Satellite gave us hundreds of channels—Internet television can give us all the free-to-air channels of the globe.'

The business model that will support all this, however, is not yet clear. Advertising, the lifeblood of many broadcasters, is difficult to do globally. And distribution through P2P networks has costs. Many broadcasters make a great deal of money by selling international distribution rights. Though there may be little commercial demand for a seven-hour documentary of a train journey in Norway, and so no harm done by giving it away, the BBC makes lots of money selling programmes such as 'Top Gear'. It is unlikely that broadcasters, even in the public sector, would give away for nothing on the Internet what they might otherwise sell through traditional routes. Until a better reason for using P2P distribution emerges, it may only be of niche interest to the big broadcasters.

Courtesy of the Economist, 4 March 2010

2. Internet Media Distribution Technologies

This section provides high-level background information to the following Internet media delivery methods:

- Unicast
- IP Multicast streaming
- File Downloading
- Progressive Download
- Peer-to-Peer (P2P)
- Podcasting
- Content Distribution Networks (CDN) and
- Hybrid P2P-CDN systems.

2.1 Unicast streaming

Unicast describes a point-to-point network communication used for sending information packets individually from a server to a client. Server and client are identified by their individual IP address assigned at least for the duration of the communication session. For each additional client, requesting the same content, a new Unicast connection has to be established, requiring:

- a distinct copy of the information packets,
- establishing a connection on the IP layer,
- encapsulation of the information packet into a transport protocol and
- a broadband connection to the Internet (except for audio only applications).

The transfer of a separate copy of the content to each user leads to a massive multiplication of data transfer and a linear relation between cost and audience size.

Unicast streaming uses the Unicast protocol to transmit a continuous stream of data to a client, requiring:

- an adjustment of data volume to the channel - especially the downstream data rate of the client - by using audio and video codecs,
- use of flow control and optional error correction, either using TCP (as in HTTP streaming) or by streaming specific protocols on top of the UDP layer for example RTP/RTSP or proprietary protocols,
- the client to get information about the offer and especially the address of the streaming-server,
- the client to decode the stream using buffering, the selected codecs and a player to render and display video and/or audio.

As a result of potentially huge number of users typical for TV audience, a reliable and high-performance infrastructure has to be provided by the service provider (e.g. broadcast station), consisting mainly of streaming-servers and bandwidth. Due to the single point of failure, content delivery networks are used often, replicating the streams in different network segments providing redundancy. The use of content delivery networks also can provide economic and performance benefits, the latter by temporal and regional dispersion of traffic peaks.

2.2 IP Multicast streaming

IP Multicast represents a point-to-multipoint network communication used for sending information packets from one source to a defined group of receivers (hosts). Network routers are responsible for optimizing the paths between the source to the receivers and for replicating the traffic at appropriate points within the network. Because the source only needs to send each packet once, and it can be received by an arbitrary number of recipients, it has significant scaling advantages compared to Unicast streaming.

While Multicast would be the distribution mechanism of choice for broadcasters, today it works only in few network segments and autonomous systems (AS). Though inter-domain routing protocols exist and most Internet exchange points and routers are Multicast enabled, Multicast use in the Internet is not wide, due to the following reasons:

- around the year 2000, Telcos invested huge amounts for their Internet structure. As the Multicast protocols prevalent at the time had some problems, providers of content delivery networks opted for using caching solutions such as Akamai CDN (Content Distribution Networks),
- peering at no charge of Tier-1-providers among each other versus paid-up IP transit for Tier-2 providers is a permanent political and financial issue between Telcos, and Multicast traffic complicates the discussion further,
- ISPs are primarily interested in an efficient distribution in their AS and not necessarily in the whole Internet. To this end, large Telcos (like Telefonica and Deutsche Telecom) apply Multicast within their network but not to other domains.

In recent years the introduction of intra-domain Multicast within the networks of ISPs has increased as a result of the IPTV Triple Play service requirements.

The Multicast-enabled routers replicate the packets at OSI-layer 3. For every Multicast communication there must be

- one source IP address which is abbreviated (S) for the sender and
- one group address as destination abbreviated (G) within the IPv4 range of 224.0.0.0 and 239.255.255.255,

Both addresses are combined in the expression (S,G). Multicast packets are forwarded by a router to other routers on the way to receivers. The router knows the hosts within its subnet by a membership list. The membership list is defined by the Internet Group Management Protocol (IGMP), which exists in 3 versions. The actual IGMPv3 is defined by RFC 3376 and is backwards compatible. IGMPv3 administrates a *Membership-Query* using the methods *General-Query*, *Group-Specific Query* and *Group-and-Source-Specific-Query*. Receivers inform the router about their wish to participate in or to leave a group and the router stores this *source filtering* data in its *membership-report*.

While Unicast routing is concerned about where the packet is going, Multicast routing is concerned about where the packet comes from. In order to achieve an efficient routing of the content to specially distributed receivers, so-called distribution trees are constructed, as Multicast paths to many receivers form a tree-topology. Possible methods are *Spanning Tree*, *Source Trees* and *Shared Trees*. The two latter are more flexible and thus preferred. The forwarding mechanism is based on routing protocols defining rules for forwarding and discarding Multicast packets. Thereby they enable routers to build a delivery tree between the sender(s) and receivers of a Multicast group. Within domains the Routing Protocol, called Protocol Independent Multicast (PIM) is mostly used in its Sparse Mode (PIM-SM) version. PIM-SM performs the following functions:

- builds and tears down distribution trees,
- requires an explicit join-message of a router willing to participate and
- uses externally-provided reachability table to build forwarding topology.

The central router within a PIM-SM is the Rendezvous Point (RP). A host joins a group by asking his router (First-Hop-Router) who in turn registers at the RP. The further development of PIM-SM is called PIM-Source Specific Multicast (PIM-SSM) defined in RFC 3569. PIM-SSM directly initiates a Source-Specific Multicast distribution tree to find the shortest path and does not necessitate an RP.

Source Specific Multicast thus is based on PIM-SSM and IGMPv3 and represents the Multicast state-of-the-art. The Internet Assigned Numbers Authority (IANA) has reserved the IP Address range 232.0.0.0/8 for SSM.

Implications within the OSI reference model that must be paid attention to are:

- Layer 2: the Multicast IP address range with 28 bits cannot be mapped unambiguously to Ethernet Multicast addresses containing just 23 bits,
- Layer 4: only UDP is supported.

Additionally, practical problems exist regarding insufficient Multicast support of inexpensive routers in the home network environment. There are open issues with wireless (WiFi) Multicast networks.

2.3 File Downloading

File transfer between hosts is the oldest form of computer networking. It accounts for the majority of traffic in today's Internet. As an application it is characterized by the ability of the recipient to consume the content multiple times, make copies and re-distribute. From a protocol viewpoint it can be implemented in many different ways, from traditional FTP and HTTP download, through first generation Peer-to-Peer protocols such as Napster, to the current leading technology of swarming protocols such as BitTorrent.

Multicast file transfer protocols also exist, e.g. File Delivery over Unidirectional Transport (FLUTE)[3], but are not widely deployed, for the same reasons that limit the use of Multicast in general.

2.4 Progressive Download

Progressive download is a simple (HTTP) file download using a standard web server. The media play-out starts immediately after enough data is stored (buffering) in the local RAM. Because the media file stays in the cache memory or other local storage space, it makes trick mode like pause, rewind or forward fast and easy to implement. Progressive download needs no large bandwidth. A slower download capacity will simply result in a longer time before the media file starts to play.

The biggest advantage of progressive download is cost efficiency. Progressive download uses no special streaming server and fits well in the existing Internet content delivery network. The downside of progressive download is poor scalability.

A new development called *Adaptive HTTP Streaming* is being considered by several companies including Microsoft (Smooth Streaming), Adobe, Apple and Move Networks. This system acts like streaming but is based on HTTP progressive download. It consists of a long series of very small progressive downloads. Adaptive HTTP Streaming has arguably become a fierce competitor to P2P.

[3] See RFC 3926

It has already received a prospect of significant commercial success, due to the following low-entry barrier features:

- Compatibility with existing HTTP network infrastructure such as CDNs
- No firewall, NAT, RTSP interoperability issues
- Smoothly adjusts the content quality to the changing network conditions and the device processing power
- The client does not need to download more than it actually consumes
- Fast start up – starts with low quality and increasing quality, as buffer fills up
- Provides better quality, as it uses the whole available bandwidth
- A non-adaptive streaming system forces the client to select a fixed bit rate which is below the available bandwidth; this is not the case with the adaptive HTTP streaming system.

The Adaptive HTTP Streaming system is in the process of being standardised within the Open IPTV Forum (OIPF). Currently, OIPF works on choosing the most suitable format of segments (e.g. MPEG-2 Transport Stream or MP4-F) and defining a client manifest.

2.5 Peer-to-Peer (P2P)

P2P systems have demonstrated their ability to provide large-scale content distribution in the Internet. The R&D work has now moved on from file sharing to multimedia streaming of live content such as live TV over P2P networks (P2P IPTV). There are already numbers of P2P IPTV applications deployed on the Internet, inspired by the P2P architecture of BitTorrent or eDonkey. Unlike a client-server based system, peers bring with them serving capacity. Therefore, as the demand of a P2P system grows, the capacity of the network grows too. This enables a P2P application to be cheap and scalable.

In a P2P network, the peers organise themselves into networks consisting of transient users (peers join and leave the network at any time). Peers upload files to other peers and download files from them. Thus, every client is a potential server (and vice versa). Since the nodes that participate in P2P networks also devote some computing resources, such systems scale with the number of peers in terms of hardware, bandwidth and disk space. As additional peers join a P2P network, they add to the download and upload capacity of the system.

Today, there are many P2P networks available in the market, each with their corresponding protocols. The protocols have the task to locate content and route query requests and query responses. All P2P protocols share the following three steps:

- Locate the content by issuing a search
- Based on responses received, choose peers from which content is downloaded
- Issue download requests and receive the response when download is completed.

In the peer-to-peer world there are different categorizations of P2P that range from completely centralised to completely decentralized. In the former camp are Napster and Seti@Home, which are centralised, and in the other camp, Freenet and Gnutella, which are decentralised. Napster incorporates a centralised indexing server that knows which files are on which clients; all clients send searches to this indexing server for locating resources. Freenet and Gnutella distribute the database of shared resources across the clients on the network, removing any need for a central server.

Gnutella uses a decentralised searching and downloading algorithm. An overlay of the various peers is dynamically constructed and queries are distributed (flooded) over portions of the overlay by neighbours forwarding the queries up to a certain radius. This leads to a linear growth of the load on the node as the number of queries increase, thus Gnutella does not scale well. To handle the scaling problem, distributed hash tables (DHT) have been proposed. The DTHs ease the query routing problem by organising the peers in a structured overlay. Due to P2P nodes joining and leaving the network, the hash data structure needs to be updated regularly. Some P2P systems (like Kazaa) rely on 'supernodes' that retain information about the data in the peers, and all queries are sent to them.

2.6 Podcasting

Podcasting is a variety of file download. It is typically (but not exclusively) intended for transferring the downloaded files to portable media players.

Podcasting allows the Internet audience to automatically receive the latest episode of a chosen programme as soon as it becomes available. One needs to *subscribe* to receive a podcast service, rather like one might subscribe to a magazine and get it delivered each week. The receiver's computer will periodically poll the set of subscriptions and automatically download all new content. The content may be automatically copied to the portable media device by the same application responsible for downloading the content.

Subscription uses the RSS (Really Simple Syndication) protocol whereas download may use any suitable file transfer protocol including P2P protocols. Some 'Podcatcher' software packages (e.g. Miro, Juice) accept RSS-Feeds containing Torrent Enclosures. This is also known as 'Torrentcasting', but it should be mentioned that these feeds are not compliant to the official definition of Podcasts. There is also the option to add the RSS functionality to BitTorrent Clients (e.g. µTorrent and Vuze). Websites such as www.mininova.org offer several torrents, which can be subscribed to using RSS.

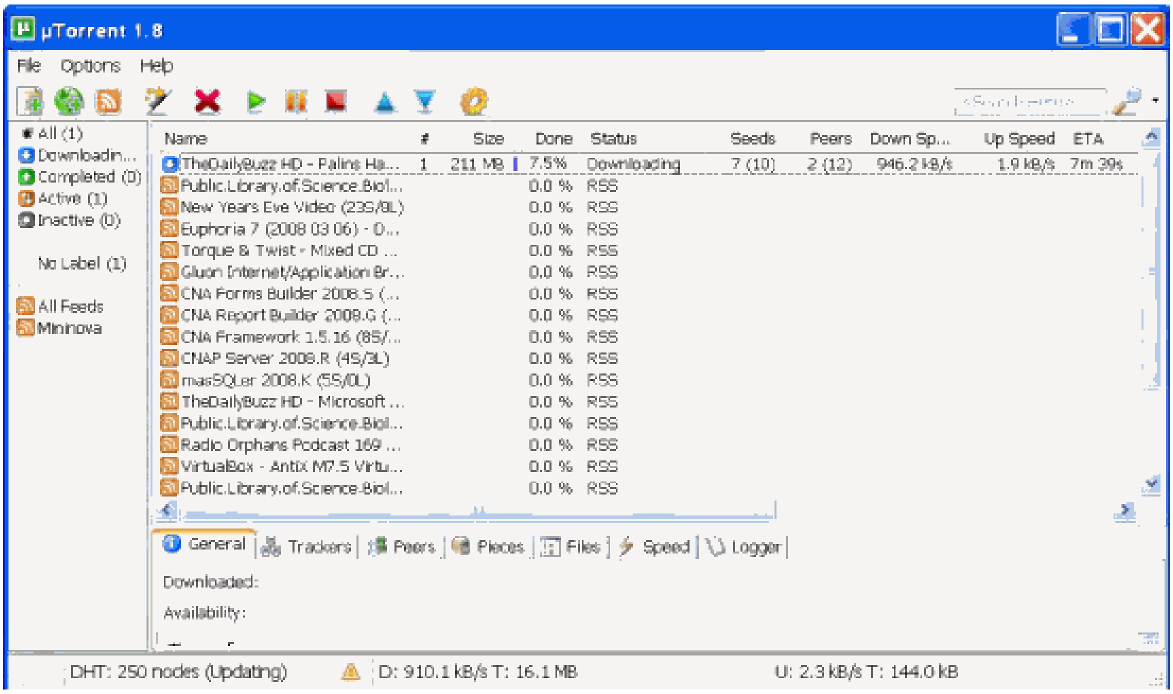

Figure 1: BT-Client µTorrent v1.8 downloading 'The daily Buzz HD...' announced by RSS

2.7 Content Distribution Networks (CDNs)

The Internet seems to be going through a similar phase of network evolution as the national television broadcast networks that were built out in the 1950s and 1960s. Regional broadcast transmitters were built in close proximity to viewers' homes to ensure they received high-quality

TV reception [4]. Today, many network service providers are building a new infrastructure of caching servers located at the edge of the service area. These networks are called Content Distribution Networks and their aim is to bring content closer to the end users.

A content distribution network (CDN) is a system of computers networked together across the Internet that cooperate transparently to deliver content to end users, most often for the purpose of improving performance, scalability, and cost-efficiency. The capacity sum of strategically placed servers can be higher than the network backbone capacity. This can result in an impressive increase in the number of concurrent users. For instance, when there is a 10 Gbit/s network backbone and 100 Gbit/s central server capacity, only 10 Gbit/s can be delivered. But when 10 servers are moved to 10 edge locations, total capacity can be 10*10 Gbit/s.

Strategically placed edge servers decrease the load on interconnects, public peers, private peers and backbones, freeing up capacity and lowering delivery costs. Instead of loading all traffic onto a backbone or peer link, a CDN can offload these by redirecting traffic to edge servers.

Since edge servers are usually placed near end users, assets are delivered via fewer hops, resulting in lower latency and increased delivery speed. End users will likely experience less jitter, fewer network peaks and surges, and improved stream quality - especially in remote areas. The increased reliability allows a CDN operator to deliver HD quality content with high QoS, low costs and low network load.

Modern CDNs can dynamically distribute assets to strategically placed redundant core fallback- and edge-servers. They can have automatic server availability sensing with instant user redirection. A CDN can offer 100% availability, even with large power, network or hardware outages.

Modern CDN technologies give more control of asset delivery and network load. They can optimize capacity per customer, provide views of real-time load and statistics, reveal which assets are popular, show active regions and report exact viewing details to the customers.

CDN nodes are deployed in multiple locations, often over multiple backbones. These nodes cooperate with each other to satisfy requests for content by end users, transparently moving content to optimize the delivery process. Optimization can take the form of reducing bandwidth costs, improving end-user performance, or increasing global availability of content.

The number of nodes and servers making up a CDN varies, depending on the architecture, some reaching thousands of nodes with tens of thousands of servers on many remote PoPs. Others build a global network and have a small number of geographical PoPs.

Requests for content are algorithmically directed to nodes that are optimal in some way. When optimizing for performance, locations that are best for serving content to the user are chosen. This may be measured by choosing locations that are the fewest hops or fewest number of network seconds away from the requesting client, so as to optimize delivery across local networks. When optimizing for cost, locations that are least expensive may be chosen instead.

In an optimal scenario, these two goals tend to align, as servers that are close to the end user may have an advantage in serving costs, perhaps because they are located within the same network as the end user.

Content Delivery Networks augment the end-to-end transport network by distributing on it a variety of intelligent applications employing techniques designed to optimize content delivery. The resulting tightly integrated overlay uses web caching, server-load balancing, request routing, and content services. These techniques are briefly described below.

- Because closer is better, web caches store popular content closer to the user. These shared network appliances reduce bandwidth requirements, reduce server load, and improve the client response times for content stored in the cache.

- Server-load balancing uses one or more layer 4–7 switches, also known as a web switch, content switch, or multilayer switch to share traffic among a number of servers or web caches. Here the switch is assigned a single virtual IP address. Traffic arriving at the switch is then directed to one of the real web servers attached to the switch. This has the advantages of balancing load, increasing total capacity, improving scalability, and providing increased reliability by redistributing the load of a failed web server and providing server health checks.
- A content cluster or service node can be formed using a layer 4–7 switch to balance load across a number of servers or a number of web caches within the network.
- Request routing directs client requests to the content source best able to serve the request. This may involve directing a client request to the service node that is closest to the client, or to the one with the most capacity. A variety of algorithms are used to route the request. These include Global Server Load Balancing, DNS-based request routing, Dynamic metafile generation, HTML rewriting, and 'anycasting'. Proximity—choosing the closest service node—is estimated using a variety of techniques including reactive probing, proactive probing, and connection monitoring.

Simple CDNs require manual asset copying. Earlier CDNs used active web caches and global hardware load balancers. Modern CDNs use cheap and simple edge servers and intelligent central CDN management technologies that distribute assets dynamically.

Several protocols suites are designed to provide access to a wide variety of content services distributed throughout a content network. The Internet Content Adaptation Protocol (ICAP) was developed in the late 1990s to provide an open standard for connecting application servers. A more recently defined and robust solution is provided by the Open Pluggable Edge Services (OPES) protocol. This architecture defines OPES service applications that can reside on the OPES processor itself or be executed remotely on a Callout Server. Edge Side Includes or ESI is a small markup language for edge level dynamic web content assembly. It is fairly common for websites to have generated content. It could be because of changing content like catalogs or forums, or because of personalization. This creates a problem for caching systems. To overcome this problem a group of companies created ESI.

2.8 Hybrid P2P – CDN systems

Peer-to-peer (P2P) is increasingly used in combination with CDNs to deliver content to end users. The Hybrid P2P-CDN system is considered an extension to conventional CDN content delivery/distribution services. The real strength of P2P shows when one has to distribute highly attractive data, like the latest episode of a soap opera or some sort of software patch/update, in a short period of time. Ironically, the more people download the (same) data, the more efficient P2P is, slashing the cost of the peering fees a CDN provider has to pay due to inter-peer delivery (in comparison to the same amount of data distributed using traditional techniques). On the other hand, 'long tail' type material does not benefit much from P2P delivery since to gain advantage over traditional distribution models a P2P-enabled CDN has to force storing (caching) of data on peers and this is usually not desired by users and rarely enabled.

Contrary to popular belief, P2P is not limited to low-bandwidth A/V signal distribution. There is no technical boundary, built-in inefficiency or by-design flaw in peer-to-peer technology preventing distribution of Full HD audio + video signal at, for example, 8 Mbit/s. It's just environmental factors such as low upload bandwidth or the lower computing power of CE devices that prevent HD material being widely available in P2P CDNs.

There have been several trials done by the Telewizja Polska SA, Poland's national broadcaster, involving live 800 kbit/s and 1.5 Mbit/s Windows Media simultaneous streams delivered over public Internet proving that the limiting factor in P2P delivery systems is the upload capacity of peers, which in DSL environments vary from 1/4 to even 1/16 of the download bandwidth.

There are some concerns about lack of QoS control over P2P distribution, but it is being developed and implemented by the *P2P-Next* consortium (see §17).

3. A short history of P2P

Peer-to-peer (P2P) technology evolved over the years from a simple 'share amongst friends' concept to complex implementations involving specially developed techniques such as 'ISP-aided Biased Query Search' and multilevel trust, reputation and the swarm concepts found for example in the Tribler social P2P implementation that was derived from BitTorrent

In fact, the root idea behind P2P concept can be traced back to the beginning of the Internet itself. ARPANet, a military project from which Internet has evolved, used the concept of peering, packet routing over interconnected network and traffic travelling from A to B via peers to overcome single point of failure in case of an atomic attack. It is the idea of multiple shared connections and utilizing both up- and down-links that brings P2P and ARPANet close together.

Up until 1999 the dominant data exchange model was the client-server approach. It proved efficient for small files or requests (cf. WWW/HTTP). For larger data sets, the old-style downloading method was just not enough and the idea of client-server architecture did not support the sharing needs Internet users started to demand. First of all, the clients did not utilize the 'reverse channel' (upload) for anything other than protocol-related acknowledgements or small requests to servers. The sharing concept was hard to achieve since client-server technique did not fit well and there was virtually no way to view who-had-what, apart from browsing flat-text files.

Even when a user was lucky enough to have a then not-that-uncommon symmetrical DSL or E1 Internet connection, the sharing rapidly saturated the upload capacity, even when most people downloaded the same file.

In addition, the data was scattered amongst users, with everyone having a different set, some pieces could be found in others' data sets and some were exclusive just to one user. A way had to be invented to make the data sets both trackable and queryable.

Note: the use of the term 'data set' might be somewhat misleading, and perhaps academic and artificial but it is done on purpose to move away from the real origin of P2P technology – the illegal distribution of MP3 files ripped from CDs. This early use of P2P caused a major blow to that technology's reputation and almost since the very start; P2P has had to prove it is 'not merely a technology for subverting intellectual property rights'.

Another thing is worth noting: the history of P2P shows that an advanced computer technology can be developed by its users from the ground up, with companies catching up only when it is mature and stable enough for commercial uses and when it has proved its superiority over traditional client-server architectures.

3.1 First generation P2P

One of the first applications of the 'applied P2P approach' was the Napster project conceived by Shawn Fanning who, motivated by his college roommate who was having difficulty accessing music files, wrote a tool to download and share music files amongst friends. It enabled users to share data, query the metadata describing the content itself and to make direct connections to parties willing to share.

This was also the concept for eDonkey 2000 network which was basically a star-based data indexing service with direct inter-peer file transfers but very much in a server-client fashion.

Applying the P2P concept to data sharing was revolutionary but both Napster's and eDonkey's downside was the centralisation of metadata which could easily lead to networks dieing instantly when the central server(s) were shut down, as was Napster's case. Without the central 'aid' provided by the server, no file-tracking or inventory of data was available. The user community learned that in order for a network to run freely without a single point of failure, metadata search has to be decentralised.

3.2 Second generation P2P

Even before Napster was shut down due to legal action there were attempts made to create fault-tolerant peer-to-peer file sharing networks (as in the original ARPAnet concept). The outcome was the Gnutella network, a completely decentralised file sharing framework in which metadata was distributed and had no single point of failure. This seemed like an antidote to a centralised approach but it turned out that treating all nodes as equal created its own bottlenecks and did not scale well. In fact it scaled very badly; finding the data was difficult since the query has to be sent to peers and this generates traffic in a pyramidal fashion, with far-distant nodes replying to the original query minutes after the data was sent into the network. The volume of query- and protocol-related data escalated, making the network slower to respond and making it difficult to obtain rare data. When there are a large number of peers in the network, most of the traffic consists of queries and metadata replies with no real file-transfer being performed. This 'query-flood' phenomenon was unacceptable.

This experience caused some enthusiasts from Sweden, Denmark and Estonia to establish FastTrack, a network consisting of normal peers and 'supernodes'. This effort, carried out in 2001, was so successful that it was later adapted for the Skype and Joost services, with the first of these itself making computing history.

The concept of 'supernodes' was to improve network response to metadata queries with caching and tracking of peers. Considered by some as a 'generation 2.5' P2P network, FastTrack spawned the infamous Kazaa client which, apart from being a malware and spyware vehicle, proved that some sort of server 'aid' is required for P2P networks to scale well.

The same concept of supernodes has been implemented and extended in numerous file sharing P2P networks. The main 2.5G P2P entity is the next generation eDonkey network (started in 2002) which introduced the concept of splitting large files into small chunks and increasing throughput by downloading different chunks of a file from a number of peers at the same time and thus making the network more resistant to peers leaving.

Another 2.5G network is Freenet, which has been in development since 2000 (release 0.7 of Freenet was launched in 2008). Freenet has an interesting approach since it employs numerous techniques that make second-generation P2P networks become more efficient, 'unbreakable' and 'uncensorable', such as key-based routing via distributed hash tables and nodes constantly attempting to swap locations (using the Metropolis–Hastings algorithm) in order to minimize their distance to their neighbours. Freenet's implementation of distributed storage and data caching make it a very advanced peer-to-peer distribution model. Similar techniques are being used to stream live video in a P2P network (for example, Octoshape).

An alternative approach to peer-to-peer and file sharing was begun by Bram Cohen in 2001. The idea was presented at the first CodeCon conference at the DNA Lounge in San Francisco, California. In this approach a small 'torrent'file named, for instance, 'MyFile.torrent' describes the files being shared and the means of peer tracking. Special servers called 'trackers' help the peers to find other peers hosting the file chunks they need. In the event of tracker failure, a different 'torrent' file could be used or an alternative tracker can be selected as this information is 'encoded' into the original 'torrent' file. Alternatively one might use the distributed tracker technique known as a distributed hash table (DHT) overlay network, which elects various nodes to index certain hashes, thus making the network fault-tolerant and the metadata dynamically decentralised.

One very important aspect of the BitTorrent protocol is that it applies the chunked file data model to achieve much lower cost to the content provider, much higher redundancy, and much greater resistance to abuse or to ‘flash crowds’ than a regular client-server distribution model, or tree-based P2P overlay networks

The introduction of ‘torrent’ files moved the finding of data away from the P2P network itself as it could conveniently be published, indexed and accessed over the traditional HTTP protocol. It is considered a ‘best of both worlds’ approach to file sharing and to P2P content distribution.

Second generation P2P networks addressed open issues of decentralisation, scalability and fault tolerance of the first generation. More or less all these were solved although no single network stands out as ‘THE network’ or ‘THE concept’ for the future P2P development. Second generation P2P networks still have their problems, mainly in the security area. When applied to distributing A/V material the current 2.5G P2P networks provide enough power, intelligence and technology to support either live (Freenet-like plus Unicast for fast channel switching) or VoD (BitTorrent-like) distribution models.

3.3 Third generation P2P

The new wave of P2P networks has anonymity and encrypting features built in. Yet they have not reached high levels of penetration because most current implementations incur too much overhead in their anonymity features, making them slow or hard to use. However, in countries such as Japan where very fast fibre-to-the-home Internet access is commonplace, a number of anonymous file-sharing clients have already reached high popularity.

Generally speaking, third generation P2P networks do address many relevant issues important for live/VoD audio-visual distribution. This report attempts to address those issues in the forthcoming sections.

4. A wealth of P2P networks and clients

Most BiTorrent clients are available from http://www.emule-mods.de/bittorrent/. The table below shows some of the most used P2P clients and the P2P networks on which these clients can be used for downloading audio and video files.

Table 1: Most popular P2P client software

LimeRunner	www.p2phood.com	Gnutella2, Gnutella, eDonkey, and BitTorrent
BitHost	http://p2p-boosters.com	Gnutella2, Gnutella, eDonkey, and BitTorrent
Shareghost	www.sharehost.com	Gnutella2, Gnutella, eDonkey, BitTorrent
Blubster	www.blubster.com	The MP2P Network
ZipTorrent	http://www.emule-mods.de/bittorrent/torrent-ZipTorrent.html	BitTorrent, MediaDefender 'decoy'!
Torrent Monster	www.goforsharing.com	Gnutella
G3 Torrent	http://sourceforge.net/projects/g3torrent	BitTorrent, current version 5.0.7
Vuze	www.vuze.com	BitTorrent, current version 4.4

Bit Pump	http://www.download3000.com/download_27555.html	BitTorrent, current version 1.00
Piolet	www.piolet.com	Manolito Network, current version 3.1.0
Sharetastic	http://www.goforsharing.com/index.php/home-mainmenu-1/sharetastic-mainmenu-126	Gnutella2, Gnutella, eDonkey, and BitTorrent
Yet ABC	http://www.goforsharing.com/index.php/home-mainmenu-1/sharetastic-mainmenu-126	BitTorrent, current version 3.1
Deluge	http://deluge-torrent.org	BitTorrent
WinMX	www.winmxworld.com	eDonkey/eMule, BitTorrent, Gnutella
Zultrax	http://zultrax.com	Gnutella, Zepp
Strong DC++	http://strongdc.sourceforge.net/download.php?lang=eng	Direct Connect, current version 2.41
Warez	www.warezclient.com	BitTorrent
iMesh	www.imesh.com	Gnutella and G2
BitLord	www.bitlord.com	BitTorrent
Pando	www.pando.com	Pando
Frostwire	www.frostwire.com	BitTorrent, Gnutella, current version 4.20.3
BitTyrant	http://bittyrant.cs.washington.edu/	BitTorrent
Lphant	http://lphant.com	BitTorrent, eDonkey2000, Kad
µTorrent	www.µTorrent .com	BitTorrent
eDonkey	http://www.softonic.fr/s/edonkey-2009	eDonkey2000, Overnet
eMule	www.emule-project.net	eDonkey2000, Kad
Ares	http://aresgalaxy.sourceforge.net	Ares, current version 2.1.5
KCeasy	www.kceasy.com	Gnutella, OpenFT, Ares
BitTornado	http://bittornado.com	BitTorrent
BitComet	www.bitcomet.com	BitTorrent
Shareaza	http://shareaza.sourceforge.net	eDonkey2000, BitTorrent, Gnutella, G2
LimeWire	www.limewire.com	Gnutella
BearShare	www.bearshare.com	Gnutella
Morpheus	www.morpheus.com	Gnutella2, FastTrack, eDonkey2K and Overnet P2P networks
Kazaa	http://www.infos-du-net.com/telecharger/Kazaa,0301-50.html	

Currently, the five most popular P2P clients for Windows are µTorrent, BitComet, Limewire, Vuze (formerly Azureus) and BitLord.

For MAC OS X the choice is much smaller. Amongst the most popular clients are µTorrent, Transmission, BitTorrent and Vuze (formerly Azureus).

The most popular P2P network is BitTorrent, with the following clients capable of working on it: ABC, Arctic Torrent, Vuze, BitComet, BitTornado, BitTorrent, BT++, G3 Torrent, Halite, TorrentFlux, µTorrent, Warez P2P and XBT Client.

According to the analysis available at www.p2pon.com, BitTorrent has the following strengths and weaknesses:

Table 2: BitTorrent strengths and weaknesses

BitTorrent Strengths	BitTorrent Weaknesses
Superior corruption managing	Trackers not very protected from legal attacks
Encryption provided	Supports only large files
High download speeds	Short-time availability for the files
Offers 'Hash'-like link functionality with torrents	Lacks global search option
It's open source	Your IP is known to trackers
Credit system to facilitate fair trading	
Well-known	

BitTorrent content can be searched for at a number of dedicated websites. Among most popular sites are the following: The Pirate Bay (contains some 600000 plus torrents, but it was taken down by Swedish law enforcers), Mininova.org (some 200000 indexes), BiteNova, BitSoup, BTjunkie, FileMP3.org, Isohunt.com (presently having some legal difficulties), Torrentscan.com (a 'meta-search' engine, which is a search engine that can find other search engines), LegalTorrents.com, Torrent-Finder.com, Torrentmatrix.com and many others.

Some examples of such decentralized, distributed systems are given below.

4.1 Freenet

As the name indicates, Freenet is a P2P-based overlay network application that allows for user anonymity. It permits the publication, replication, and retrieval of data while protecting the anonymity of both authors and readers. Freenet operates as an adaptive network of identical nodes that collectively pool their storage space to store data files and cooperate to route requests to the most likely physical location of data. No broadcast search or centralized location index is employed. Files are referred to in a location-independent manner, and are dynamically replicated in locations near requestors and deleted from locations where there is no interest. It is infeasible to discover the true origin or destination of a file passing through the network and difficult for a node operator to determine or be held responsible for the actual physical contents of her own node.

Freenet is free software available from http://freenetproject.org. Freenet lets you anonymously share files, browse and publish 'freesites' (web sites accessible only through Freenet) and chat on forums, without fear of censorship. Freenet is decentralised to make it less vulnerable to attack, and if used in 'darknet' mode, where users only connect to their friends, is very difficult to detect.

Communications by Freenet nodes are encrypted and are routed through other nodes to make it extremely difficult to determine who is requesting the information and what its content is.

Users contribute to the network by giving bandwidth and a portion of their hard drive (called the 'data store') for storing files. Files are automatically kept or deleted depending on how popular they are, with the least popular being discarded to make way for newer or more popular content. Files are encrypted, so generally the user cannot easily discover what is in his data store, and hopefully cannot be held accountable for it. Chat forums, websites, and search functionality, are all built on top of this distributed data store.

Ideas and concepts pioneered in Freenet have had a significant impact in the academic world. Freenet is based on a paper entitled 'Freenet: A Distributed Anonymous Information Storage and Retrieval System' [25].

An important recent development, which very few other networks have, is the 'darknet'. By only connecting to people they trust, users can greatly reduce their vulnerability, and yet still connect to a global network through their friends' friends' friends and so on. This enables people to use Freenet even in places where Freenet may be illegal, makes it very difficult for governments to block it, and does not rely on tunneling to the 'free world'.

The client that receives a packet is not able to readily establish from which client(s) this packet is coming from. Communication on Freenet occurs by sending a request to a client you are connected to, who in turn sends it on to another client, and so on. In Freenet, any client can insert content resources into the network as well as search for them and retrieve them.

In order to insert resources into the network, the resource is given a title identifier which is hashed using SHA-1 to generate a unique key to identify the resource. The resource is then associated with that key and stored locally. A request for insertion into the network may be made in which the resource with key attached is sent to other nodes to store. The 'direction' in which this packet will travel is based upon the 'nearness' of its key to other keys that other nodes are storing. This means that data with keys that are close, or near to other keys, reside close to each other in the network.

To search for a resource on the network it is necessary to know its title. The title is then hashed using SHA-1 and a request is made for that title by sending it to the most likely place to have the resource, based, once again, on key closeness. An initial search is then performed, searching for nodes that contain the required key. The search backtracks if the hop count reaches zero or if during the search a node is seen twice. When the required resource is found the search terminates and the client with the resource begins to send the matching resource back along the search route to the client who requested the resource. All clients along the way will cache the passing data which helps in the replication of popular resources and means that frequently requested data is cached and dispersed widely around the network increasing redundancy and reducing access times. No client participating in the resource retrieval knows if the resource it is downloading and passing on, is coming from the original publisher and going to the original requester, or whether it is just coming from or going to some other links in the chain.

4.2 Gnutella

Gnutella celebrated a decade of existence on March 14, 2010 and still has a user base of several millions. In March 2000 Nullsoft, a subsidiary of America Online, released a file sharing application called Gnutella that allowed file swapping without the need of a central indexing server and therefore no central point of failure, and no central point to sue for copyright infringements. On April 10, America Online declared Gnutella to be a rogue project and terminated it, but not before the program had been downloaded and replicated by thousands of users around the net. Over the next few weeks the protocol was reverse engineered and Gnutella clones began to appear. In 2007, it was the most popular file sharing network on the Internet with an estimated market share of more than 40%. In June 2005, Gnutella's population was 1.81 million computers increasing to over three million nodes by January of 2006.

Gnutella's architecture is similar to Freenet's in that it is completely decentralised and distributed, meaning that there are no central servers and that all computations and interactions happen between clients. All connections on the network are equal. When a client wishes to connect to the network, he runs through a list of nodes that are most likely to be up or takes a list from a website, and then connects to as many nodes as he needs. This produces a random unstructured network topology.

Routing in the network is accomplished through broadcasting. When a search request arrives at a client, that client searches itself for the file and broadcasts the request to all its other connections. Broadcasts are cut off using TTL (time-to-live) which specifies how many hops the packet passes before clients drop them rather than broadcast them.

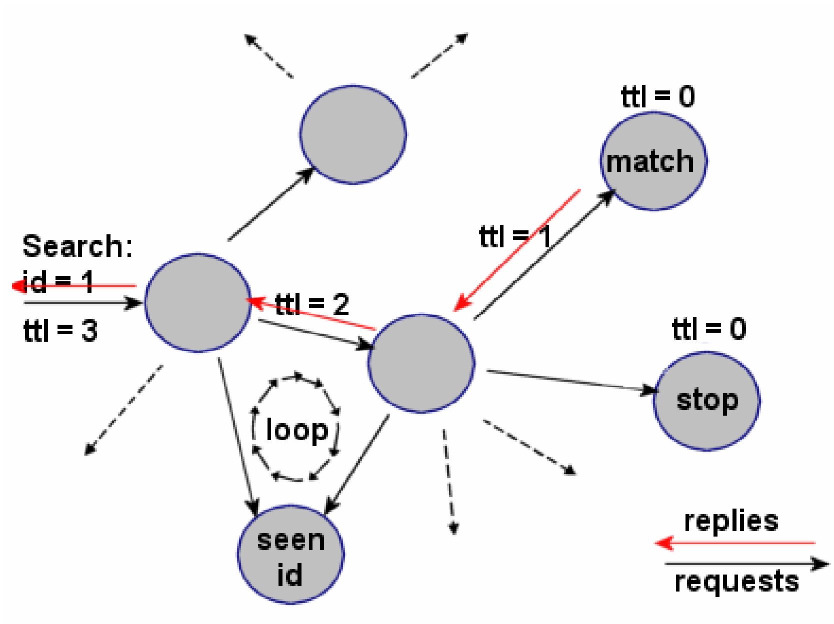

Figure 2: Example of a Gnutella broadcast search. Illustrates searching, replying, packet meeting TTL and redundant loops

There is a small degree of anonymity provided on Gnutella networks by this packet routing technique. A client that receives a packet does not know if the client it has received the packet from is the original sender or just another link in the chain.

Searching on Gnutella is accomplished by creating a keyword string that describes the file you want and broadcasting that string to all your connected neighbours. Your neighbours will then in turn broadcast that message to all their connected neighbours and so on until the packet's TTL has been reached.

Reaching a lot of nodes for search can be quite effective. With a TTL of 7 on the packet and an average of 8 neighbours per client, one million clients may be reached in theory. In reality this is not the case. Optimistically, one may have up to 40000 clients on Gnutella but only 2000 to 4000 of these clients may be reached with a search.

Upon a query match, clients create a packet that contains information on how to locate them and the file (a URL). To route the replies, all the clients send the query replies back along the same path they came. Following a certain time, all results will arrive back at the client who originally sent them out. At this point the client can decide which file, if any, he wants to download.

To download a file the client creates a direct connection to the client with the file he wants and sends a HTTP packet requesting the file. The client with the file interprets this and sends a standard HTTP response. This process, however, removes any anonymity in the system.

4.3 Kazaa

Kazaa uses the FastTrack architecture which follows a 2-tier system in which the first tier consists of fast connections to the network (Cable/DSL and higher) and the second tier consists of slower connections to the network (modem and slower). Clients on the first tier are known as SuperNodes and clients on the second tier are known as Nodes. Upon connection to the network, the client decides whether he qualifies as a SuperNode or not. If he qualifies, he connects to other SuperNodes and starts taking connections from ordinary Nodes. If he only qualifies as a Node, he finds a SuperNode that will allow him to connect to them and does so. In such a two-tier topology the nodes at the centre of the network are faster and therefore produce a more reliable and stable backbone.

Routing on FastTrack is accomplished by broadcasting messages between the SuperNodes. For example, when a Node issues a search request to the SuperNode it is connected to, the search request is taken by that SuperNode and then broadcast to all the SuperNodes it is currently connected to. The search continues in this way until its TTL has reached zero. Every SuperNode that it reaches searches an index that contains all the files of its connected Nodes. This means, that with a TTL of 7 and with an average number of Nodes per SuperNode of 10, a search request will search 11 times more nodes on a FastTrack network than on Gnutella. Unfortunately, since the searches are broadcast, the network will still produce enormous amounts of data that needs to be passed from SuperNode to SuperNode. However since the SuperNodes are guaranteed to be reasonably fast it does not produce as large a problem as on Gnutella.

Routing of replies follows the same lines as Gnutella. Replies are routed back along the path that they came from, until they reach the clients that originally issued them. A large problem with this type of routing in Gnutella was that clients making up its backbone were very transient and connected and disconnected to the network very sporadically which meant that packets being routed back along the path they came along could find the path gone because a link in the chain had disconnected. This problem occurs less on Kazaa as clients making up the backbone are guaranteed to be faster and more stable and therefore paths for return routing packets should be more reliable.

Downloading on Kazaa is the same as on Gnutella. Once the location of the file has been found, the client that wants the file connects to the client that hosts the file and sends a HTTP request for the file. The client hosting the file interprets the request and sends a HTTP response. The HTTP headers used in Kazaa have been modified to accommodate extra information such as metadata but standard HTTP/1.1 headers are supported. This means that files can be downloaded from Kazaa clients through a web-browser such as Internet Explorer or Mozilla.

Unfortunately, although the FastTrack topology is decentralised, one implementation of Kazaa requires that all clients register with a central server before being allowed to connect to the network which invalidates all the advantages of having a decentralised topology. Kazaa was undergoing legal proceedings with the RIAA (Recording Industry Association of America) but Kazaa's legal issues ended after a settlement of $100 million in reparations to the recording industry. Kazaa, including the domain name, was then sold off to Brilliant Digital Entertainment, Inc.

Today Kazaa operates as a monthly music subscription service allowing users to download unlimited songs.

4.4 BitTorrent

Programmer Bram Cohen designed the protocol in April 2001 and released a first implementation on July 2, 2001. It is now maintained by Cohen's company BitTorrent, Inc. BitTorrent is a peer-to-peer file sharing protocol used to distribute large amounts of data. BitTorrent is one of the most common protocols for transferring large files, and by some estimates it accounts for about 35% of all traffic on the entire Internet. The initial distributor of the complete file or collection acts as the first seed. Each peer who downloads the data also uploads them to other peers. Because of this, BitTorrent is extremely efficient. A minimum of one seed is needed to begin spreading files between thousands of users (peers). The addition of more seeds increases the likelihood of a successful connection exponentially. Relative to standard Internet hosting, this provides a significant reduction in the original distributor's hardware and bandwidth resource costs. It also provides redundancy against system problems and reduces dependence on the original distributor.

Usage of the protocol accounts for significant Internet traffic, though the precise amount has proven difficult to measure. There are numerous BitTorrent clients available for a variety of computing platforms. According to isoHunt the total amount of content is currently more than 1.1 petabytes.

A BitTorrent client is any program that implements the BitTorrent protocol. Each client is capable of preparing, requesting, and transmitting any type of computer file over a network, using the protocol. A peer is any computer running an instance of a client.

To share a file or group of files, a peer first creates a small file called a 'torrent' (e.g. MyFile.torrent). This file contains metadata about the files to be shared and about the tracker, the computer that coordinates the file distribution. Peers that want to download the file must first obtain a torrent file for it[4], and connect to the specified tracker, which tells them from which other peers to download the pieces of the file.

Though both ultimately transfer files over a network, a BitTorrent download differs from a classic full-file HTTP request in several fundamental ways:

BitTorrent makes many small data requests over different TCP sockets, while web-browsers typically make a single HTTP GET request over a single TCP socket.

BitTorrent downloads in a random or in a 'rarest-first' approach that ensures high availability, while HTTP downloads in a sequential manner.

Taken together, these differences allow BitTorrent to achieve lower cost to the content provider, higher redundancy, greater resistance to abuse and can more easily handle 'flash crowds' than a regular HTTP server. However, this protection comes at a cost: downloads can take time to rise to full speed because it may take time for enough peer connections to be established, and it takes time for a node to receive sufficient data to become an effective uploader. As such, a typical BitTorrent download will gradually rise to very high speeds, and then slowly fall back down toward the end of the download. This contrasts with an HTTP server which rises to full speed very quickly and maintains this speed throughout (but is more vulnerable to overload and abuse).

The peer distributing a data file treats the file as a number of identically-sized pieces, typically between 64 kbyte and 4 Mbyte each. The peer creates a checksum for each piece, using the SHA-1 hashing algorithm, and records it in the torrent file. Pieces with sizes greater than 512 kB will reduce the size of a torrent file for a very large payload, but this reduces the efficiency of the protocol. When another peer later receives a particular piece, the checksum of the piece is compared to the recorded checksum to test that the piece is error-free. Peers that provide a complete file are called seeders, and the peer providing the initial copy is called the initial seeder.

The exact information contained in the torrent file depends on the version of the BitTorrent protocol. By convention, the name of a torrent file has the suffix .torrent. Torrent files have an 'announce' section, which specifies the URL of the tracker, and an 'info' section, containing (suggested) names for the files, their lengths, the piece length used, and a SHA-1 hash code for each piece, all of which is used by clients to verify the integrity of the data they receive.

Torrent files are typically published on websites or elsewhere, and registered with a tracker. The tracker maintains lists of the clients currently participating in the torrent. Alternatively, in a trackerless system (decentralized tracking) every peer acts as a tracker. This is implemented by the BitTorrent, µTorrent, rTorrent, KTorrent, BitComet, and Deluge clients through the distributed hash table (DHT) method. Vuze also supports a trackerless method but is incompatible (as of April 2007) with the DHT offered by all other supporting clients.

Users browse the web to find a torrent of interest, download it, and open it with a BitTorrent client. The client connects to the tracker(s) specified in the torrent file, from which it receives a list of peers currently transferring pieces of the file(s) specified in the torrent. The client connects

[4] Torrent files can be obtained legally from a number of web sites including www.vuze.com (movie trailers in HD, www.publicdomaintorrents.com and www.legaltorrents.com.

to those peers to obtain the various pieces. Such a group of peers connected to each other to share a torrent is called a *swarm*. If the swarm contains only the initial seeder, the client connects directly to it and begins to request pieces. As peers enter the swarm, they begin to trade pieces with one another, instead of downloading directly from the seeder.

Clients incorporate mechanisms to optimize their download and upload rates; for example, they download pieces in a random order to increase the opportunity to exchange data. This is only possible if two peers have different pieces of the file.

The effectiveness of this data exchange depends largely on the policies that clients use to determine to whom to send data. Clients may prefer to send data to peers who send data back to them (a tit for tat scheme), which encourages fair trading. But strict policies often result in suboptimal efficiency; for instance, when newly joined peers are unable to receive any data because they do not have any pieces yet to trade or when two peers with a good connection between them do not exchange data because neither of them wants to take the initiative. To counter these effects, the official BitTorrent client program uses a mechanism called 'optimistic unchoking', where the client reserves a portion of its available bandwidth for sending pieces to random peers (as opposed to the known, preferred peers), hoping to discover even better partners and to ensure that newcomers get an opportunity to join the swarm.

4.5 µTorrent

Codenamed Falcon[5], the µTorrent client has an easier, more secure and more complete web UI as well as support for streaming and remote downloading. Developed by BitTorrent Inc., µTorrent Falcon brings plenty of change to the BitTorrent client currently in use by more than 50 million people a month. Below we sum up some of the key features.

Allowing users to access their BitTorrent downloads from anywhere through a simple web-interface is one of the main goals of the Falcon project. Without having to configure µTorrent and home networks so that they can be accessed remotely, users can simply browse to the Falcon page http://falcon.uTorrent.com and connect to their client instantly.

The easy to use web interface is extremely secure and is a major improvement over the Web UI currently available. When logged in, it gives users all the controls they are familiar with in their PC client. Torrents can be added, paused and removed using an interface with a look and feel identical to that of the µTorrent application.

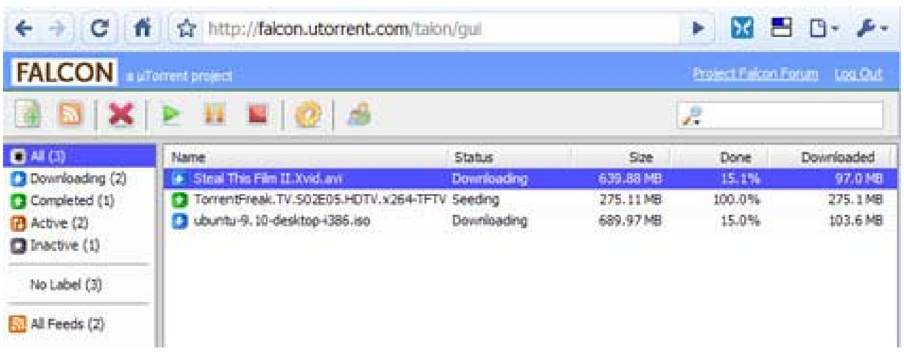

Figure 3: µTorrent Falcon web interface

Aside from the added security and easy setup, accessing your torrents via the Falcon web-interface offers another advantage – remote downloading. Once a file has finished downloading you can transfer a copy of the file to a remote computer via the web-interface.

[5] See TorrentFreak, 31 January 2010.

Another new feature of the Falcon project is the added option to stream video files while downloading. Instead of having to wait until a file has finished downloading, users can already start watching video provided that the download speed is sufficient. The Falcon intention is to transform getting media using µTorrent from a 'load-wait-watch-tomorrow' to more of a 'point-click-watch' experience.

Also new in the Falcon release is the 'Send Torrent' feature. This feature is particularly useful when you want to share torrents with people who do not have a BitTorrent client installed yet. Right clicking a torrent in µTorrent shows a 'Send Torrent' option which then brings up a URL which is a direct link to a download of the µTorrent client with the torrent file included.

The Falcon release makes it easier for users to find torrents. It provides a better ability for torrent sites to promote content or search within the client. The µTorrent client comes with a built in torrent search engine.

The latest alpha release also has a 'minify interface' option and the development team is also working on speed improvements, UI improvements and optional file security features.

4.6 Tribler

Tribler is a P2P-based application that enables its users to find, enjoy and share video, audio, pictures, and much more. Tribler is developed at the Delft University of Technology and the Vrije Universiteit Amsterdam. The Tribler software integrates state-of-the-art research into network technology, P2P networks, video streaming and user interaction. Tribler aims for an open source, fully decentralized peer-to-peer network, available for everybody, in every part of the world. Tribler supports UPnP to open a listening port on the firewall automatically.

When the Tribler application program is started, it will automatically start searching other users that have Tribler running on their computer. When a connection is established, it starts exchanging information. First it exchanges personal information (such as your avatar picture, your friends list, download history, etc.) and information about files that are available in the network. These files can be personal, shared files, but also files that one has received from another person.

The information about the discovered files and persons is available in the Tribler program. By browsing through the files and persons each user can find their preferred files and users. The Tribler program helps the user by giving extra information about each item (whether it is a file or a person) and also shows what other users think about it. When the user finds a person they like, he can add this person as a friend. An interesting file can be downloaded and will be available in the user's library. When the user presses download, his computer will make an inventory of which computers actually have this file (or a part of it) and then will download the parts from the different computers.

Tribler has three goals in helping the user: to find, consume and share content. Through improved search functionality the user searches the content of other Tribler users, and the content of large video web portals such as Youtube and Liveleak. Users can browse through different categories such as video, audio, pictures, etc. and also see what is most popular and what is made available recently. By making friends and getting in touch with users with similar tastes, the user can find content that he/she might find interesting. Users can also show to their friends what they like and what their friends should see.

Because of the integrated video and audio player the users can almost immediately start watching their favourite videos or listen to their favourite songs.

Tribler is a social application. This means that the user can make friends with other users and can show to others what he likes and dislikes. By sharing content the user helps other Tribler users to enjoy his favourite content.

4.7 P2P-enabled Adobe Flash player

The Adobe Flash player 10.1, called Stratus[6], will include P2P technology that could significantly reduce, and in some cases completely eliminate bandwidth costs incurred by the content providers. It is offered to developers free of charge and can support both live and on-demand video streaming. Besides video, Stratus can also be used for Flash based multi-player games and other forms of real-time communication.

Instead of serving the media from a central server, users will provide the necessary bandwidth. Adobe's Stratus system serves as an intermediary in this process, managing the communications between Flash players much like a BitTorrent tracker does for BitTorrent transfers.

Users first must agree to participate in a P2P-enabled Flash swarm, similar to how they are now asked to indicate whether Flash can use their webcam. If users do not want to share bandwidth, the broadcaster has the option to offer a regular stream, a degraded stream or no stream at all.

The impact of Adobe getting involved in P2P streaming could be of fundamental importance. One of the main advantages is that Adobe Flash is already installed on the majority of computers, which should facilitate the adoption rate among content providers.

5. P2P Business Environment

Delivering rich content, and in particular video and TV services, is a huge growth market opportunity. Since 1990 the amount of traffic on the Internet has been growing incredibly fast. There are two reasons for the growth: exponential increase in number of broadband users and the increase in traffic consumption per user. The latter is mainly due to the growth of the video type of content and services. The last five years, the traffic on the backbone has been growing at an aggregate rate of about 50 to 60% per year, reaching mid 2008 an equivalent of about 0.15 Gbyte per day per user. It is estimated that approximately one third of this traffic is delivered through a content delivery network [6]. This huge growth in IP traffic challenges operators to efficiently shape their network and to take advantage of new technologies to ensure quality of service for the end user. This growth presents both threats and opportunities for all players including broadcasters, content providers, Telcos, content delivery providers etc.

It is beyond any doubt that P2P is a very popular Internet protocol. The share of P2P protocol in the overall Internet traffic has been subject of numerous studies. As an example, the results of a study carried out in 2009 on the distribution of Internet backbone protocols across different parts of the world are shown below[7]. It is of no surprise that the P2P protocol takes typically more than half of the overall backbone traffic. According to [5], 65-70% of the Internet backbone traffic can be attributed to P2P. In the last mile, 50-65% of downstream and 75-90% of upstream traffic is P2P traffic. Moreover, P2P file sharing traffic continues to grow.

Ipoque (www.ipoque.com) performed an Internet study for 2007 assessing the impact of P2P file sharing, voice over IP, Skype, Instant Messaging, one-click hosting and media streaming. BitTorrent and eDonkey are the most popular P2P networks by a wide margin. The two protocols account for 70% and 97% of all P2P traffic, depending on the region. P2P is still producing more traffic in the Internet than all other applications combined. Its average proportion varies between 49% in the Middle East and 83% in Eastern Europe.

[6] See TorrentFreak, 19 May 2010.

[7] Source: EC-funded Project *P2P-Next* review, WP4, March 2009, Brussels

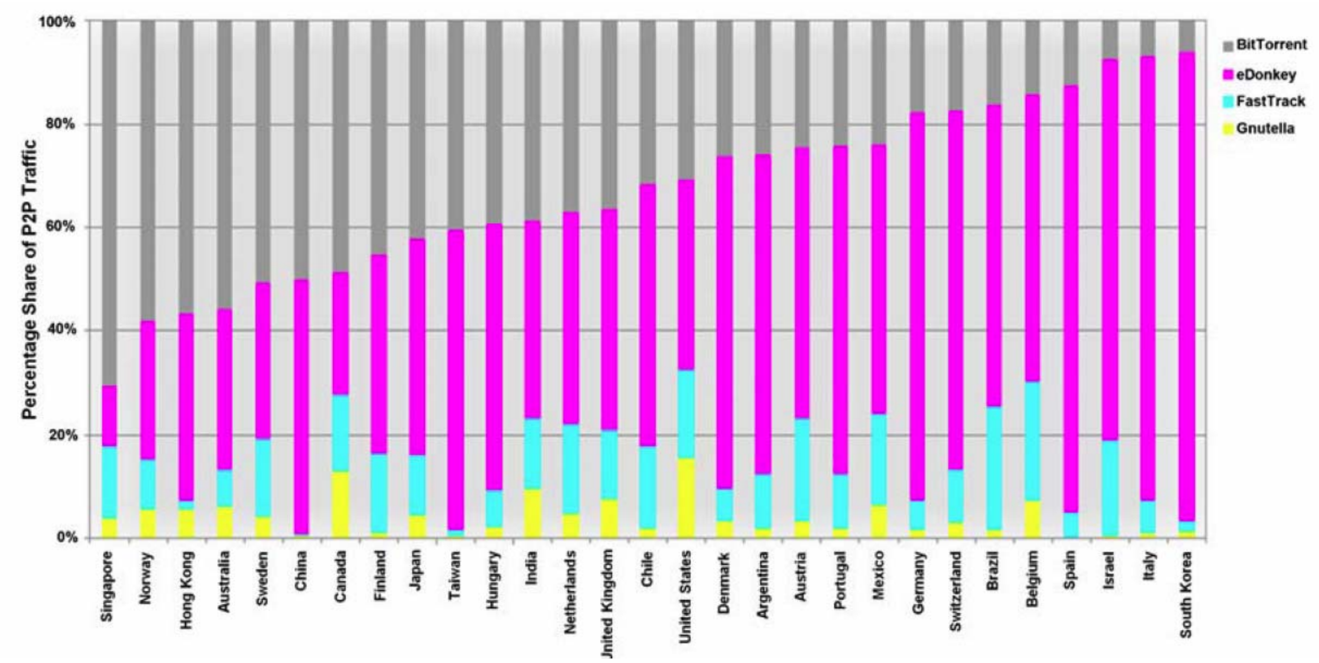

Figure 4: P2P traffic share in different countries

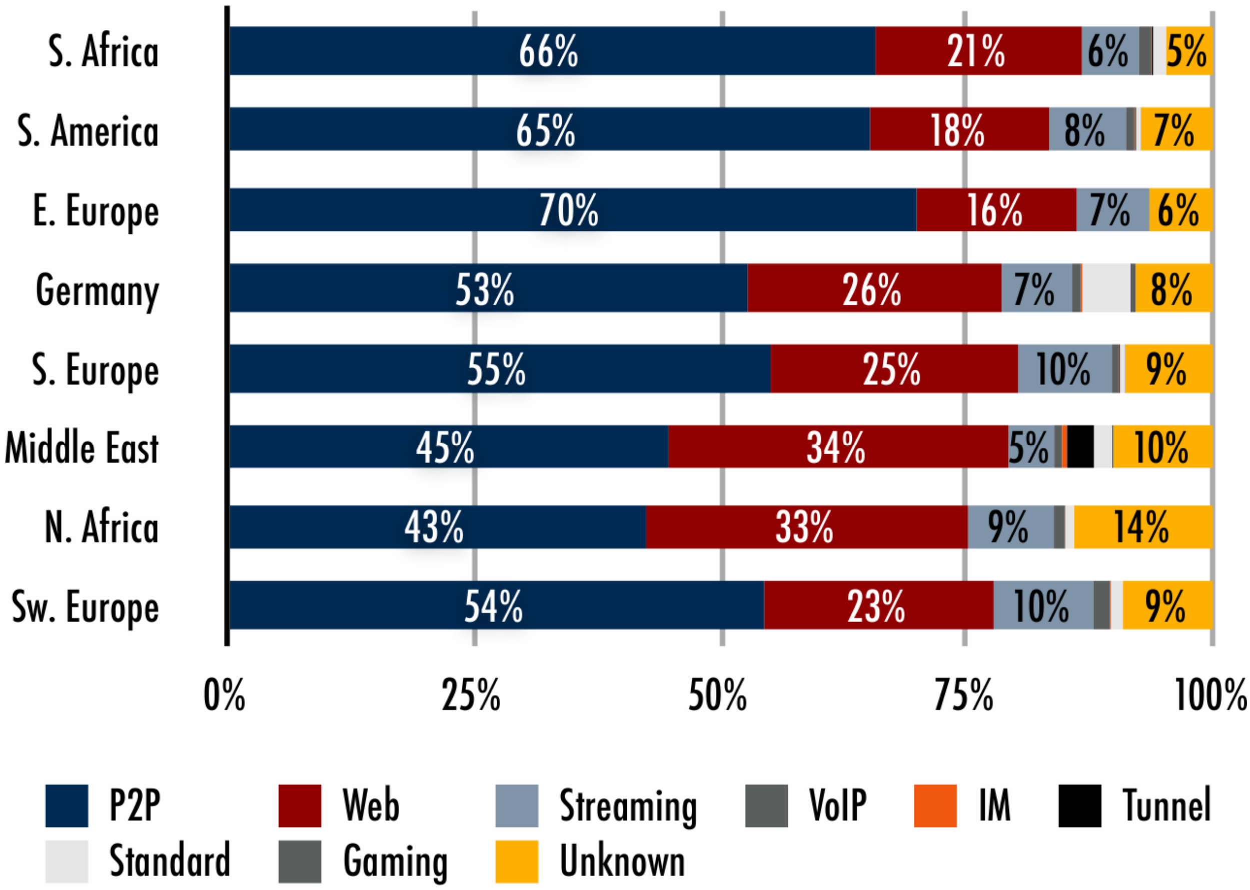

Figure 5: Usage of various Internet applications (Courtesy of Ipoque, 2007)

The complete protocol distribution for Germany is shown in Table 3:

Table 3: Protocol usage in Germany (Courtesy of Ipoque, 2007)

Protocol distribution in 2007	Germany [%]
P2P	69.25
HTTP	10.05
Media streaming	7.75
DDL	4.29
VoIP/Skype	0.92
IM	0.32
e-mail	0.37
FTP	0.50

The following pie chart provides an overview of BitTorrent traffic volume per content type for Germany. The leading BitTorrent content are video files with a share between 62% in Southern Europe and 79% in Germany.

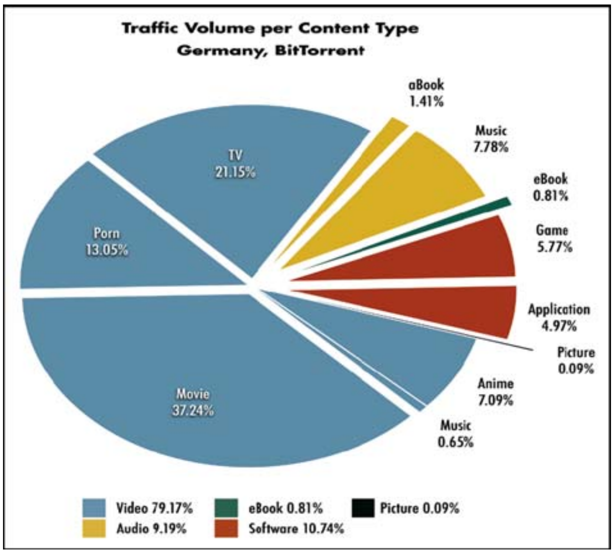

Figure 6: Traffic volume of BitTorrent based distribution in Germany

5.1 Business Model and Cost Considerations[8]

The P2P delivery system is an important ingredient of the digital convergence process which is about to drastically change the audio-visual media landscape. It will give rise to the following disruptive shifts:

- A shift from classical linear TV with a monopoly in audio-visual consumption towards non-linear, on-demand, anytime-anywhere, cross device patterns that will change the roles of publishers and consumers;
- A shift from passive collective, linear media consumption towards active personal behaviour at home and in mobility situations;
- A shift from highly popular audio-visual content to selective disclosure of 'long-tail' content,

[8] This section is based on the *P2P-Next* Work Package 2 results.

servicing communities with varying and diverse interests;

- A shift from awareness-building, non-interactive, expensive, interruptive TV advertisements to targeted, non-disruptive, sales-channel oriented forms of promotion.

The use of Audiovisual Media is moving from a collective and passive approach to a personal active approach, at home and in mobility situations outside the home. At the same time, use patterns are shifting towards non-linear usages and away from the classic models of linear broadcast TV. The TV set has no longer the monopoly of delivery of audiovisual content, the PC and related media centres, mobile phones, and potentially new devices such as game consoles, notebooks, tablets, e-books are all becoming increasingly important.

5.2 P2P Business Opportunities and Challenges

Because P2P services make use of the existing network infrastructure owned and controlled by the network service providers, they are often called the ‘over-the-top’ or ‘overlay’ services. In most cases, subscribers pay a flat-rate network access fee, independent of usage. This causes some problems for the network service providers that provide for transportation of the over-the-top video traffic but are not able to generate any associated revenues.

The P2P over-the-top content providers (such as Octoshape) may enter into direct relationship with end users, entirely by-passing the network service provider, thus entailing potential disintermediation of broadcasters, content packagers, service and network providers from subscriber base. P2P is thus most suitable for user-generated long-tail (niche) content distribution.

P2P content providers may become an ‘over-the-top’ content provider, competing with both the broadcaster and the network service provider for the same digital content services market. Network service providers are bypassed in terms of revenues, although they in fact carry a large part of the P2P traffic delivery cost. As a result of P2P services, the value of basic network access services may diminish.

In order to avoid content devaluation, broadcasters should embrace P2P technologies and enter the P2P content delivery market by themselves. Alternatively, they could outsource P2P delivery to an external specialized P2P-delivery company.

Since P2P entails bypassing the network service provider, it faces some inherent weaknesses which impact the P2P service providers:

- The P2P service is unable to provide quality guarantees for the video service provided. Such a best-effort service may prevent a lot of inherent value of the content to be monetized.
- Competitive differentiation: The over-the-top space is very crowded and there are already many players in the market. They all have access to the same type of content. To this end, content or quality are not effective differentiators. Revenues and profit margins are modest. Partnerships are emerging and collapsing every day, as the search for a viable business model continues.

5.3 P2P Business Models

As part of its business and regulatory studies, the *P2P-Next* project identified five P2P-related content delivery models:

- Business Model 1 – Free Content Distribution
- Business Model 2 – Advertisements Supported Distribution
- Business Model 3 – A la Carte Models
- Business Model 4 – Subscription Based Distribution

- Business Model 5 – Circular Content

5.3.1 Business Model 1: Free Content Distribution

A free content distribution model is generally performed by a public free-to-air broadcaster. This model encompasses business models such as: 'Give away', partly 'Attention grabbing', partly 'Use of large repositories' and 'Sales support channel'. Free content distribution may aim at revenues independent from the content itself, e.g. a free music video and revenue from resulting CD and Ticket sales.

'Give Away' Model

An example of this model is the offer of free music files embedded in mobile phones. Small record companies and individual artists who are seeking to make their name known often make their materials available on online networks for free, with the aim of attracting attention and interest for other revenue streams (concerts, merchandising).

'Attention Grabbing' Model

In the Internet environments, attention grabbing models with a large range of content as an attractive consumer offering can be promising, at least initially. These models can attract advertisers (e.g. using banners when consumers download software clients). If dubious, legality can be sued (suing can attract even more attention, e.g. Napster), bought up and then developed as a service accepted by major content owners, or simply closed down (MP3.com).

5.3.2 Business Model 2: Advertisements Supported Distribution

Typically, this model is used by a commercial broadcaster. This encompasses business models such as: 'Advertising', 'Sponsoring', partly 'Attention Grabbing', 'Customer Reach', 'Single-Click product buying'. It can be linked to advertising models (click to order) if not related to content.

Advertising Models

Advertising can help financing Internet audio-visual activities. A wide range of advertising models is available today:

- Banners when consumer downloads clients and use the service thereafter.
- Advertising prior to receiving content (as a condition).
- Regular advertising breaks (as in traditional TV with 30 second slots), or shorter breaks of a few seconds, making a fast forward unnecessary. Examples: Hulu (USA) for film and for music: Spiral Frog, Spotify, We7.com, Qtrax.
- Advertising can indirectly finance content.

Sponsoring Models

Example: Coca Cola product placement in the American Idol TV shows.

5.3.3 Business Model 3: Pay-Per-View Distribution

An example is a video rental store going electronic. This model can encompass business models as: 'A la carte'.

A la Carte Model

This model applies to a seller of audio-visual items. Example: Apple iTunes which is the first and still the leading supplier of music items. Apple is mainly interested in selling iPods and accepts low margin on sales of music. Other services in the music sector: Klicktrack, CDON, Amazon MP3, Media Markt, Tesco.

Generally, a range of choices available on iTunes cannot compete with what is available in file-sharing networks. Remuneration to artists/composers by iTunes is low, as opposed to very high remuneration offered to artists by record companies (as percentage of consumer cost). Selling music with a low margin (in order to market something else) may hinder development of a competitive retail market but can work well for ‘impulse’ purchases.

5.3.4 Business Model 4: Subscription Based Distribution

This model encompasses business models such as: ‘Flat rate’ and, partly, ‘Use of large repositories’. An example: a jazz music events producer.

Flat Rate

The Flat Rate model can be applied to both download and streaming services as well as a combination of the two (such as in Rhapsody). It can involve towards a variety of models, frequently combined with other variants (a la carte downloads as in the case of Rhapsody or Advertising as in Swedish Spotify). Many of these models often have the condition that what has been downloaded during the subscription period disappears if the subscription expires (e.g. Napster 2). Or if it is bundled with another subscription, e.g. broadband payment, the downloaded content disappears if one changes operator (example: TDC Denmark). At least one film distributor, Headweb (Sweden) considers offering a pure streaming flat rate service. Generally, consumers show limited favour of losing what they ‘have paid for’ – this tends to clash with the notion of ownership which is still governed by the physical world’s rules. Future solutions may involve blanket flat rate licences combined with broadband subscriptions.

Use of Large Repositories

Models based on accumulating large repositories of rights involve a) buying or selling and b) suing when such materials turn up in other business activities. Vertical integration in the media content industries has facilitated this model (and competition law has not been able to hinder the development).

Rights repositories and indirect incomes (e.g. from private copying levies, neighbouring rights payments) has also allowed major rights holders to earn more from such incomes than from traditional exploitation of rights (new creativity). Creativity is limited to finding new ways of generating money from rights, often without needing to remunerate the creators (again the example of the record companies buying shares in YouTube and generating an ‘equity profit’ which was not ‘royalty related profit’). This trend could lead to a withering of societal support for IPR regimes, both amongst creators as well as the general public. Artists’ managers are beginning to raise this as a legal issue with the major record companies/publishers.

5.3.5 Business Model 5: Circular Content

E.g. Super distribution, mash ups. Encompasses business models as: ‘Super distribution’, ‘Prosumer co-design’. Traditionally content is produced by professional content producers and solely consumed by users. Circular content describes an indefinite number of ‘re-produced’ content. This business model could for example reward users for super distributing contents (monetary or tokens)

Super-distribution of content

The expression ‘super-distribution’ means ‘(...) an approach to distribute digital products as software, videos, and recorded music, in which products are entirely free of copy protection and are made publicly available and distributed (...).’ In the meaning of a circular content model super-distribution takes place in controlled environments, for instance, in allowing certain explicitly marked videos or parts of videos to be distributed freely. To encourage super-distribution, e.g. for promotional contents, users could be rewarded by tokens, bandwidth or even monetary units.

Super-distribution of recommendations

This case of super distribution is very close to public playlist and involves individual consumers acting as agents, recommending and assisting further sales and receiving some form of reward for this. A number of such systems have emerged but have not survived the test of time. Altnet.com (linked closely to the Kazaa file sharing system) was one such. Another was Weedshare. They seem to suffer from allowing elements of ‘pyramid selling’ to infuse the system. A recent example in Sweden is Ameibo for film distribution. Selling on a widget which allows friends also to be part of a network is another alternative (e.g. akin to associations offering perks to those who get their friends to join).

6. P2P use cases

These cases show some typical examples where the P2P distribution mechanism can be used. However, they do not imply that P2P is the only technology which can be used, as other Internet distribution technologies (such as CDN, IP Multicasting) may be used to perform the same task with similar effect.

The P2P technologies are most useful in cases where the same content is distributed across the Internet to a large number of users at the same time. The same content is stored on a large number of devices at any given moment, and shared cooperatively between them.

6.1 Streaming media services

A user (we’ll call her Anna) is looking for a very popular scheduled programme to be streamed over broadband. She points her *P2P-enabled Internet device* to a content provider, selects a service and the TV or radio content is streamed in real time to her *P2P-enabled Internet device* via P2P.

6.2 P2P-enabled Internet-format catch-up TV services

Anna has booked to record a broadcast programme on her DVR which is connected to a P2P-enabled Internet network. At recording time there is a booking clash due to a late running schedule and the recording cannot be started. The device connects to the broadcaster’s P2P-enabled Internet-format catch-up TV service that delivers the programme (and associated adverts) to the P2P-enabled Internet device. The P2P-enabled Internet-format catch-up TV service uses some P2P-enabled Internet devices as sources among those devices that have been connected to the programme when it was broadcast.

6.3 P2P-enabled Internet device to record broadcast programmes

Anna has an interest in a specialist topic. She enters keywords into her P2P-enabled Internet device related to this topic. The device will then periodically check Internet format metadata provided by broadcasters on the open Internet. In the future Anna is informed whenever there are any programmes that are due to be broadcast that may match her interest, or the device may speculatively record the broadcast for her. The broadcaster may simultaneously push content (e.g. during the night when the network load is relatively low) via P2P to all users in possession of a P2P-enabled Internet device and that subscribe to this specialist topic.

6.4 Progressive download services

Anna points the P2P-enabled Internet device to a content provider, selects a service and the content starts downloading to her device, so that she is able to watch it after a short time and before the download is completed. This service is available via a P2P distribution mechanism.

6.5 Deferred download services

Anna points the P2P-enabled Internet device to a content provider, selects a service and the content starts downloading to her device. The download time can be significantly shortened by using a P2P approach compared to conventional Unicast.

6.6 Broadcast content related URI available from P2P-enabled Internet device

Anna is watching a film on a broadcast service. She presses the 'info' key which displays the related URI (or EIT) provided information. Her P2P-enabled Internet device also connects to the provided URI and downloads much richer information including reviews, cast lists, recommended films, etc via P2P connection.

6.7 Protected content with Portal

Some (multiple) Internet service providers require users to pay to access their content, and also control the use of the content by the user following acquisition. Anna buys a P2P-enabled Internet device supporting both free and paid for content. She is able to access both free and paid for content from all Internet service providers that operate P2P services. Anna registers with a portal (e.g. by registering using a web interface), and is then able to buy content from all providers associated with that portal without separately registering with each provider. She is subsequently able to use her P2P-enabled Internet device to access this content, subject to restrictions placed by the provider. Bob does not pay and is not able to access the content.

6.8 Protected content without Portal

Some (multiple) P2P-enabled Internet service providers require users to pay to access their content, and also control the use of the content by the user following acquisition. Anna buys a P2P-enabled Internet device supporting both free and paid for content. She is able to access both free and paid for content from all P2P-enabled Internet providers. In order to purchase content from a provider, Anna separately enters a contractual relationship (e.g. by registering using a web interface) with each provider and pays for access to some content. She is subsequently able to use her P2P-enabled Internet device to access this content, subject to restrictions placed by the provider. Bob does not pay and is not able to access the content.

6.9 PIP to display signing services

Anna's friend is hearing impaired and requires a signing service to be able to enjoy broadcast services. A broadcaster provides signing services for certain broadcasts and uses open Internet to stream these signing services synchronously to hybrid broadcast/broadband Internet-enabled devices. The signer (or its avatar) is displayed in a small PIP window which can be moved and resized on the main display. EPG (or BCG) informs of the programmes that include signing services.

If a large number of users subscribe to the signing services, it may be worthwhile for the service provider to use the P2P-enabled Internet distribution approach for these services.

6.10 P2P-enabled IPTV HNED accessing CoD services

Anna selects service discovery from her HNED and discovers that CoD (Content on Demand) services are available from managed network based service providers and Open Internet based service providers. She decides to configure her HNED to access the CoD guides of a managed network based IPTV provider and an open Internet based provider. She is able to browse both content guides and select to stream or download content from either provider. The CoD content which is available for download or streaming over the Internet can be distributed over a P2P-enabled Internet network to the general user.

6.11 Broadcast subtitles to be carried by P2P

Anna can speak a language which is not broadcast as subtitles. However the broadcaster has made DVB subtitle data for a specific programme accessible on the open Internet. She is able to select the language of the subtitle data. Her P2P-enabled Internet device is capable of receiving the subtitles of her choice. Her P2P-enabled Internet DVR records the broadcast programme and downloads the subtitle data. At playback time the recorded programme and subtitles are synchronised and displayed.

6.12 P2P-enabled Internet DVR to record broadcast content

Anna is using a P2P-enabled Internet device for previewing the content which is due to be broadcast shortly via a terrestrial channel. At a suitable point a dialogue box informs her of the option for her P2P-enabled Internet DVR to schedule a recording of the full quality broadcast, or if the broadcast is starting imminently then to switch to the service immediately.

6.13 Protected content on multiple devices

Anna owns several Internet devices which may be part of a home network. Content provider ContentCo is able to sell content to Anna that she or members of her family can consume on any of her Internet devices. Some of these devices may be P2P enabled, so that they can access TV services via P2P networks.

6.14 Picture-in-Picture (PIP) to display Internet content related to main broadcast

Anna's P2P-enabled Internet device supports picture-in-picture (PIP). She is watching a broadcast sporting event which signals that there are alternate viewing angles available. Anna selects this option. The main display continues to show the broadcast programme. The PIP window displays an alternate viewing angle streamed over the open Internet on a 'best endeavour' basis. Anna can cycle through a number of alternate angles in the PIP window.

6.15 Pre-positioned download services

Anna registers her interest in content with her service provider. The content provider downloads the selected content in accordance with her previously expressed preferences. The content can be pushed to the groups of the same taste or preferences via P2P.

6.16 Customized Programming Guides

Anna's P2P-enabled Internet device shows her a list of content to be broadcast. This Programme information is similar to an EPG and it can be customized based on her interests. She can be shown different lists such as her favourite channels, recommended channels from a portal, or channels of a specified genre in the EPG.

7. Broadcast Technical and Commercial Requirements for a Viable P2P System

This section provides the commercial requirements, as seen from the broadcasters' perspective, for the open Internet P2P distribution of audio-visual content (together with associated data) by means of the standardised CE (consumer electronics) devices (rather than PCs).

Rather than specifying the CE devices, their required functionalities are given. These CE device functionalities may be implemented either in software or hardware (or a combination of the two) and may be embedded in a set-top box, TV receiver, hand-held device, along with any other functionality.

The content can be delivered either in the form of the real-time television/radio streams ('channels') or/and on-demand files. The latter can be downloaded (and rendered after having been downloaded) or progressively downloaded (rendered almost immediately after downloading has started.

Any method of distribution via the Internet requires a means of performing the following tasks:

- Allowing consumers to discover and find content
- Allowing consumers to select content for consumption
- Transferring the content to the consumer
- Reducing the amount of data by using audio and video codecs in order to adjust data rates to the available channel bandwidth
- Provisioning of the infrastructure to enable users to access the content

Some methods provide additional functionality:

- Allowing consumers to subscribe to content, to be automatically provided with future editions, episodes, etc.
- Payment methods
- Supporting geo-restriction (geo-location), by restricting access to the content to geographic areas for copyright enforcement,
- Supporting rights management, by restricting access to paid subscribers, etc. or restricting the use of the content,
- Supporting a social model, allowing consumers to contribute reviews, ratings, referrals, etc.

In November 2007, EBU Project Group D/P2P was tasked with providing a basis for the standardisation of a system allowing for access to the content available to the general public across open Internet P2P networks.

Today, media delivery over the Internet implies the use of (relatively complex and expensive) computing devices (PCs). Our incentive is, however, to facilitate and speed up commercial introduction of Internet TV. To this end, it may be useful to replace PCs with *simpler and cheaper CE (consumer electronic) devices*. As the CE devices are simple, they can accommodate only a subset of all features available in PCs. These features need to be standardised[9].

7.1 Motivations and objectives

Broadcasters wish to deliver media over the open Internet to consumer electronics appliances such as set-top boxes and television sets which are connected to wired or wireless broadband access networks. TV sets seem to be the ideal target for broadcasting services delivered over the open Internet (example: ‘catch-up TV’ services using e.g. BBC iPlayer).

There is some evidence that the CE manufacturing industry is keen to enhance their product offering to the consumer in area of Internet connectivity. Examples: the DVB Project, Open IPTV Forum, DLNA, along with commercial projects such as the integration of NetFlix rental downloading directly to DVD players and iDTVs in the USA.

Ideally, horizontal markets for the media services and CE devices using the open Internet as a distribution mechanism of audio-visual content should be established.

Horizontal market for services: The horizontal service market means that any content and/or service provider may deliver their content or services to any CE device. Horizontal markets are very important as they enable FTA (free-to-air) broadcasters to fulfil their public remit by providing non-discriminatory announcement, transport and usage of their content offerings.

Horizontal market for CE devices: This implies that the CE devices used by the end user are not specific to only one service provider or Internet access technology (ISP). Any CE device can receive services from any provider in a given market. The horizontal market delivers many economic advantages for CE manufacturers, since products can be engineered once and sold across multiple regions with minimal regional variations.

Compared to PCs, CE devices are characterised by limited processing power and storage capability, but benefit from lower power consumption and the ability to engender higher levels of trust and reliability than PCs. Furthermore, economies of scale will drive CE device costs down. Consequently, they are cheaper and simpler, and might be more widely accepted by consumer markets worldwide[10]. Table 4 summarizes these differences.

Table 4: Principal differences between the PC and CE Devices

PC / high range functionality	CE device / low range functionality
Intel Core 2 CPU~30,000 MIPS	Embedded CPU Core ~300MIPS
≥ 1 Gbyte RAM	16 -128 Mbyte RAM
S/W Decode Audio/Video	‘On Chip’ H/W Audio/Video Decode
Lean forward UI: Keyboard and mouse	Lean backward UI: Remote control from 3m
PC Powerful, User Extendable	Device Appliance, designed for purpose
Relatively expensive	Low cost
Relatively high energy consumption	Low energy consumption (<5W)

(source: John Adam, Samsung)

[9] The EBU D/P2P work is parallel (and complementary) to the DVB-CM IPTV group work (see for example doc. CM-IPTV0379 and its revisions).

[10] An open Internet P2P device can be combined with traditional digital TV set-top boxes or integrated in TV receivers.

7.2 Working assumptions

The P2P network is assumed to be hybrid or decentralised (i.e. does not need a central distribution server) and consists of (distributed) peers which up-stream and down-stream media streams and files.

Not all peers are created equal. Super peers may be provisioned by content providers to fill the bandwidth gap due to asymmetries in the Internet access networks of their users.

There are two options for locating content in P2P networks:

- A central metadata repository (e.g. Napster or tracker-based approaches) or
- distributed metadata repository (e.g. eDonkey and other systems that leverage DHT for torrent acquisition).

No controlling organisation is required for the system to function properly.

We define a CE device (or appliance) as a simple and affordable consumer electronics device with a limited memory footprint and processing power. Such CE devices are peers participating in the P2P network. It is also assumed that all CE devices are able to participate actively or passively[11] in the P2P network. The P2P distribution technologies embedded in the CE device should be standardised according to the 'tools approach', as used in DVB. This means that more than one P2P technology or policies for a specific aspect of system operation (e.g. peer selection) can be defined. One or more standardised P2P clients may be embedded in the CE device.

As far as possible, generic solutions based on open standards and open source implementations should be used.

Backwards compatibility of future versions of the standard should be mandated, though the legacy installed base shall experience lesser features and lower performance.

7.3 System architecture

A simplified system stack is shown in Figure 7. It consists of the following layers (from the bottom up):

Physical, Network and Transport Layer: This layer represents the existing Internet infrastructure network, upon which P2P transport is overlaid.

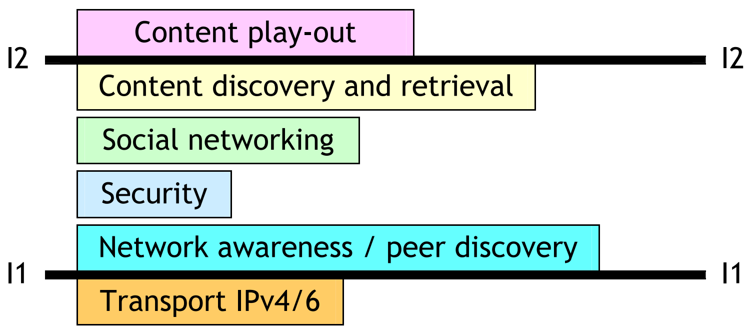

Figure 7: Definition of I1 and I2 Interfaces within the P2P system architecture

[11] Passive participation implies a stand-by status of the CE device and it does not require active user involvement. Passive participation may help improve the efficiency of the P2P network.

Peer discovery layer: This layer represents mechanisms for search and discovery of peer CE devices that share the same P2P standard client and are able to share/exchange content.

Interface I1 between bearer and P2P system functionalities is derived from the OSI reference model layers. An important task of I1 is to provide bidirectional transport and measurements of connection to peers in order to optimise the transport.

Interface I2 interfaces the P2P system which may host a metadata repository and the functions of a P2P-enabled device which renders the content and processes metadata.

Security layer: This layer prevents any malicious peers wreaking havoc in the P2P network by spreading viruses or by trying to make resources unavailable. This layer is designed to achieve availability of files/streams and peers, file/streams authenticity, peer anonymity, payments and access control. The latter may restrict access to media resources only to those peers that have the rights to access those resources.

Social networking layer: This layer specifies the social networking characteristics of the P2P system such as communities, user profile, friends, rating, tagging, recommendations and personalisation, reputations, message boards, chatting, etc.

Content discovery and retrieval: content-related metadata, content search, rights management (DRM), watermarking and fingerprinting, programme guides (EPGs), etc.

Content play-out: Audio/video decoding, content rendering, data display, presentation facilities, content storage, PVR.

Figure 8 is another perspective of the system architecture.

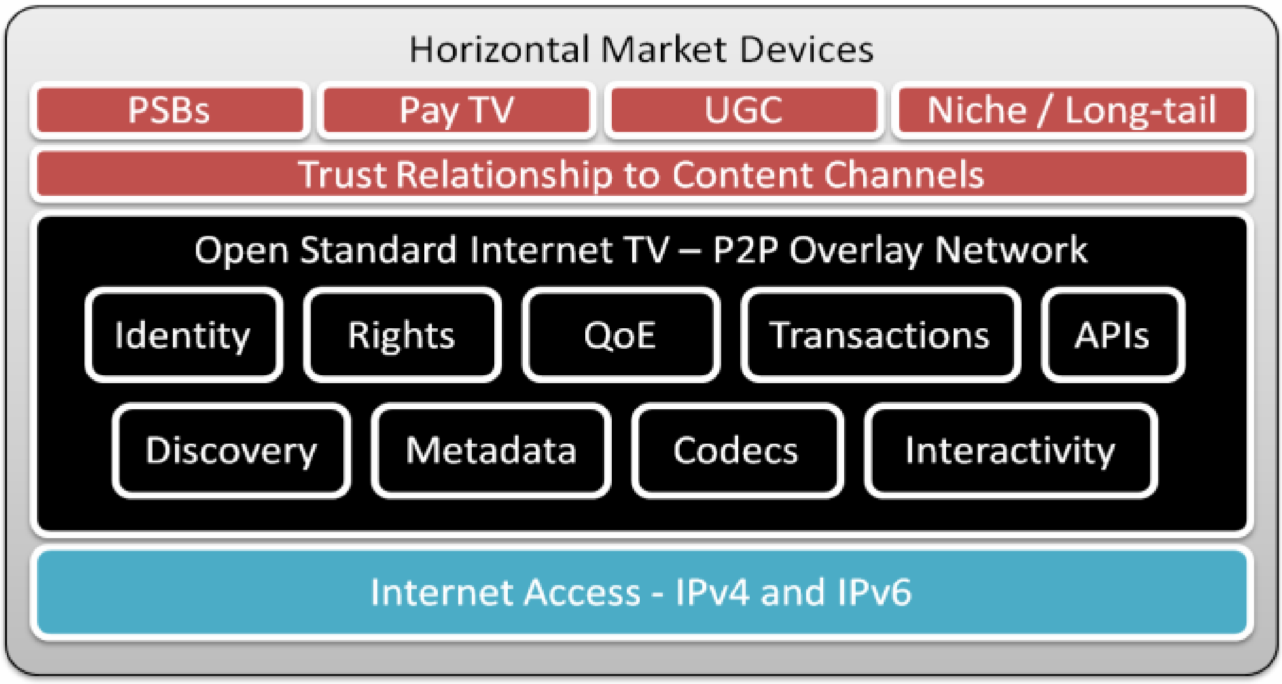

Figure 8: Classification of market devices

Figure 9 illustrates the concept of a P2P infrastructure and positions a standardised P2P distribution platform with the over-the-top services that exploit it and the value-chain. Importantly, it also emphasises the need for a trust platform to enable implementation of the business models identified earlier.

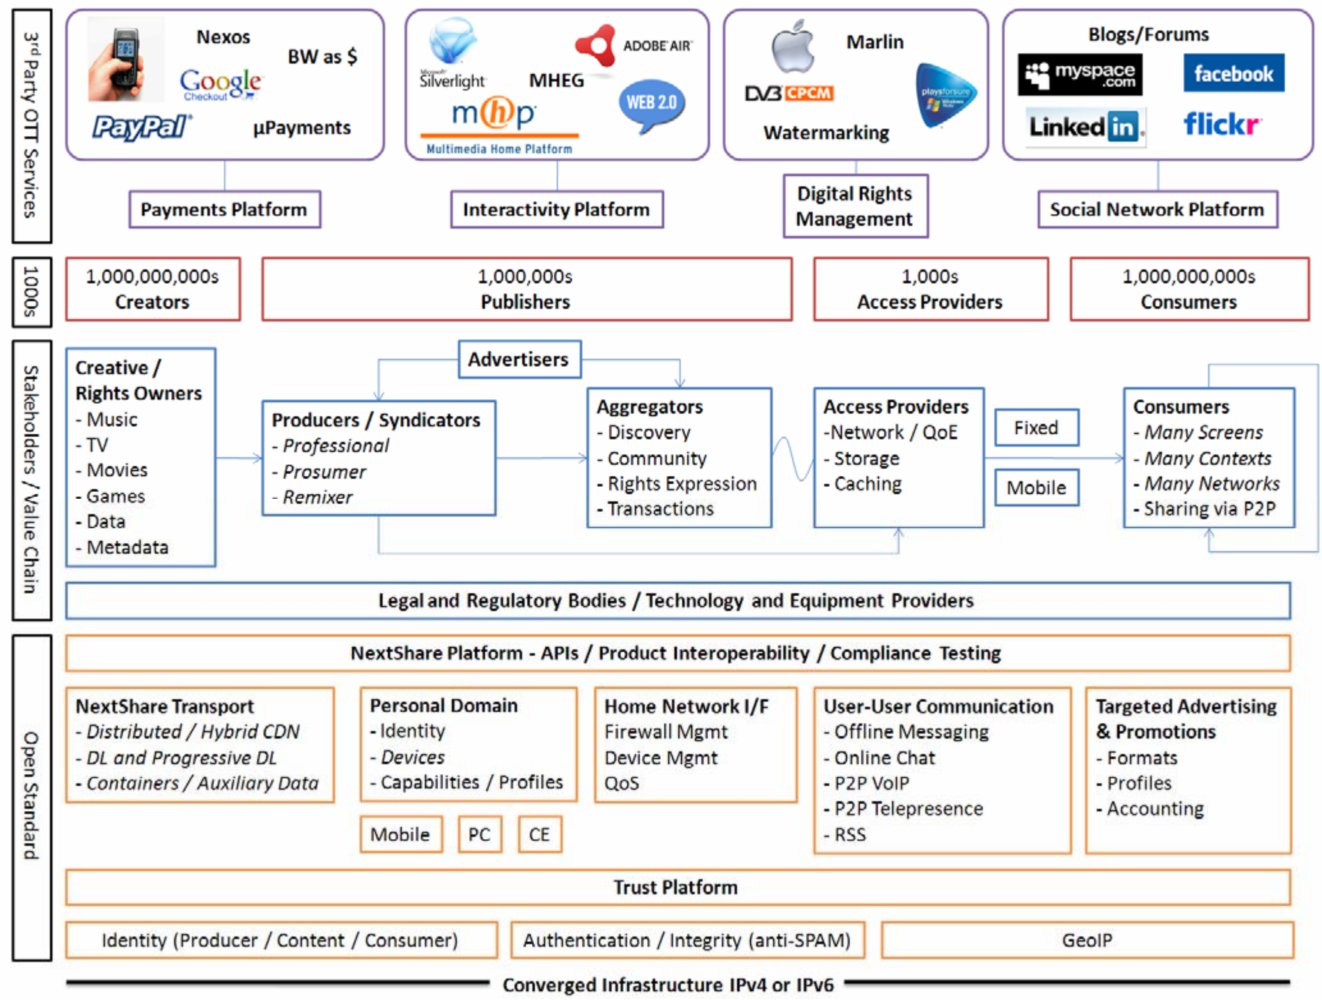

Figure 9: Concept of a P2P-infrastructure

7.4 General P2P-related requirements

7.4.1 A CE device should enable open Internet access via a broadband connection (via cable or xDSL, wireless), in order to locate and consume open Internet content and services.

7.4.2 CE devices should be able to connect to both server-based (Unicast/CDN, Multicast) services as well as P2P-based services via either bridged or routed access.

7.4.3 The P2P clients should be standardised by competent bodies and should be embedded in the CE device.

7.4.4 The CE device functionality is either a part of an IPTV managed system or a standalone open Internet consumer appliance (example: Internet radio appliance).

7.4.5 Any technical specification prepared by the international standardisation bodies should provide a common profile of the necessary protocols and data formats required to deploy a DVB service over the open Internet.

7.4.6 The CE device allows for a variety of business models in a horizontal market.

7.4.7 The CE device allows for collecting statistical audience data and capturing user behaviour, while fully complying with international regulations concerning user anonymity and user data confidentiality. Such data should not be divulged to third parties without the user's consent and should not be used for commercial purposes.

7.5 P2P delivery-related requirements

7.5.1 The P2P based delivery and contribution of content should allow for different implementations to enable live-streaming, progressive-download and deferred download. However, these implementations should be able to cooperate with each other in non-discriminating way in terms of quality, performance/quantity and bandwidth sharing, although peers must inevitably discriminate based on traffic localisation and the priority of reducing access latencies.

7.5.2 A P2P architecture should allow for scalability in terms of the number of clients. This should be true for a typical network case where the upstream capability of some peers may be limited as well as there being adverse issues such as heavy usage, high network packet loss rate and jitter or NAT/firewall problems.

7.5.3 No upper limit for the number of concurrent users should exist.

7.5.4 The system should be able to cope with flash-crowds and may be improved by additional super-peers provided by content or service provider.

7.5.5 In the case of considerable network asymmetry, suitable compensation should be provided by caching or super-peers.

7.5.6 Appropriate caching should be provided in case of content tails, e.g. in order to offer infrequent content pieces or partial streams.

7.5.7 Client-site caching of P2P packets on a designated area of the device hard disk should be considered to complement usefully the ISP-level cache.

7.5.8 Continuity of service should be ensured in case of situational breakdowns, peer failures or opt-outs, change of IP numbers and network mobility as well as temporary packet losses.

7.5.9 In order to maintain adequate service continuity, it may be occasionally necessary to switch to a reduced bitrate stream (for the same service). This should be performed automatically. In such a case the service providers should generate a range of different stream qualities (bitrates) or suitably encode the video streams in some form of scalable video coding and/or multiple description coding.

7.5.10 Under extremely severe circumstances a fallback to Unicast streaming may be required.

7.5.11 A fast service build-up is required especially for live streaming.

7.5.12 Latency/switching time of P2P-based live streaming should be comparable to the DVB-IPTV standards.

7.5.13 P2P may be complemented or combined in some cases by IP Multicasting or CDN approaches. In the former case, if no IP Multicast is available, the client switches to P2P based distribution[12].

7.5.14 Ability to seek and resume playback at any position within a P2P-based VoD stream should be provided.

[12] See www.delco.cs.tut.fi

7.6 Social networks and user preferences

7.6.1 The P2P system should enable the implementation of social network functionalities such as

- on-line communities, friends, friends-of-friends
- rating the content and usage preferences of the user
- user personalisation
- user comments and tagging
- reputations and recommendations

7.6.2 If not desirable (e.g. for some specific TV channels), it should be possible to restrict or deactivate (wholly or partially) the social network functionalities, either temporarily or permanently.

7.7 Content formats

7.7.1 Content providers should support audio and video streaming and file as well as associated container formats that are customary and predominantly used in the Internet environment. If possible, these formats should be compatible with those used in conventional digital broadcasting such as DVB.

7.7.2 Content formats used by content providers and implemented in the P2P CE devices should be aligned in terms of codecs, levels and profiles used.

7.7.3 Carriage of large files (e.g. HDTV) should be enabled for on-demand services.

7.8 Content delivery

7.8.1 Content should be delivered in the form of

- live streaming of TV/radio channels,
- on-demand file downloading (incl. subscription based content as podcasts),
- on-demand progressive downloading.

7.8.2 Content may originate from any content sources including traditional broadcasting sources, industrial content providers as well as general users (user-generated content).

7.8.3 Effort should be made to preserve high end-to-end content quality, with minimal impairments and degradations of the original content.

7.8.4 Every means should be taken in order to minimise computational capacity in the CE P2P device and to allow simple recomposition of partial P2P streams, regardless of the transport system used (either conventional MPEG TS or direct A/V streaming over RTP).

7.8.5 In case of P2P-based forwarding, the FTA-content must be reconstructed after reception in an unchanged manner and completely in terms of quality and transported advertisement.

7.8.6 Copy management such as defined by DVB CPCM should cover usage restrictions after acquisition of the content, in response to either simple signalling or Digital Rights Management (DRM) information delivered alongside the content. A minimum requirement would be central administration of content usage via P2P based unique identification of this content (e.g. allocation to other peers' upstream capacity by communicating their current IP addresses).

7.8.7 Access Services including subtitling, spoken subtitling, audio description, deaf signing may be added to the core audio and video content, in order to provide value-added services and satisfy requirements of different interest groups.

7.8.8 Watermarking services may be imbedded in the audio content, video content (or both), in order to help differentiate our broadcast traffic from any illegal P2P traffic. In addition, watermarking could help detect copyright information, authenticate users and monitor content usage. Watermarking should not be perceptible to the user or otherwise degrade end-to-end quality. It should be rugged enough to resist any functional or hostile processing (e.g. attacks). Watermarking data protocol needs to be specified in order to contain the required information.

7.9 Content acquisition

7.9.1 Content discovery and acquisition mechanisms in the P2P environments should be similar or identical to those used in IPTV or open Internet TV services.

7.9.2 The P2P system should use efficient and easy-to-implement ('lite') content acquisition mechanisms, suitable for discovery of both live and on-demand content. The system should also be used to display service information. Announcements of Content on Demand (CoD) files should be possible.

7.9.3 The P2P system may allow for a suitable display of RSS (Really Simple Syndication) feeds. RSS allows publishing and frequently updating new content. This should be enabled on clients for a browser-based display of feeds, user's combination of different feeds and the announcement of available content. If a browser doesn't support subscription of a feed, the final device should comprise an additional feed aggregator. Only versions RSS 2.0 and higher should be taken into account. A broadcaster can choose RSS to transfer URLs (e.g. of his own websites or torrents) to customers. RSS typically offer the options to subscribe, which is valuable for broadcasters' customer retention[13].

7.9.4 A central server should be available for each service to check whether the content is the same as labelled originally and thus provide authentication of the content and its provider[14].

7.10 Service management and service monitoring

7.10.1 The P2P system should be designed to allow efficient service management performed by the service/P2P provider.

7.10.2 The overall service quality and service experience should be controlled and monitored constantly by the service and content providers. The network and traffic load characteristics should be taken into account.

7.10.3 There should be a mechanism allowing the bandwidth usage shared by different applications to be suitably prioritized by the user.

7.10.4 Statistics about audience tracking are to be available to content providers and service

[13] The announcement of torrents by RSS is already known as 'Broadcatching' and realised by µTorrent, Miro player and Vuze plug-in for RSS Feeds.

[14] For example, each broadcaster could provide a Tracker-Server at an URI, such as 'tracker.broadcaster.eu' or otherwise broadcasters could install a joint architecture.

operators at any juncture (dynamically).

7.10.5 The P2P service should be interoperable with geolocation, DRM and other applications required for fulfilling copyright and right management.

7.11 Management of P2P functions in a CE device

7.11.1 The P2P clients should be embedded in the CE device. They should be upgradeable remotely by the service/network provider/CE manufacturer

7.11.2 The P2P client should require only a ‘moderate’ hardware footprint in order to comply with the low cost nature of CE devices.

7.11.3 The P2P client should be transparent to different OSs, browsers and players

7.11.4 The P2P clients should be of modular design and future proof

7.11.5 Excessive usage of user’s upload capacity and user’s storage and processing power by the P2P client should be prevented. Other applications should always be able to work in conjunction with the P2P application.

The maximum upload capacity that is used should be adaptable (i.e. limited, automatically and possibly also manually) as a function of the download capacity and the maximum available upload-bandwidth.

7.11.6 The P2P clients should be capable of being uniquely identified and authenticated

7.11.7 The P2P CE device should be able to work in passive (idle) mode comprising a wake-up and a sleep function triggered by a P2P announcement and P2P-based transfer at minimum CPU workload

7.11.8 The P2P CE device should have a simple user interface to install, operate and upgrade

7.11.9 The P2P CE device can be stand-alone or else can be integrated with IPTV, DTT, cable and/or satellite receiving terminals.

7.12 Security issues

7.12.1 The P2P CE device provides a secure network connection to a known person (rather than CE device)

7.12.2 The P2P system is able to combat eavesdropping and man-in-the-middle attacks

7.12.3 The P2P system contains the necessary content and network security mechanisms including authentication, authorisation and accounting (AAA) mechanisms

7.12.4 The integrity and authenticity of content should be guaranteed by using the certificates and trusted public key infrastructure (PKI). This can be achieved for the whole file and the announcement (as a torrent). Each piece of a file can be checked using a lower level of security comparable to hash sums. Note that hash sums cannot be calculated for live streams in advance of broadcast, but pieces can be signed to verify their integrity and authenticity at the time they are seeded to the P2P system.

7.13 Geolocation, DRM and storage-related issues

7.13.1 Broadcasters and content providers should be able to use the geolocation (GL, geo-restriction, geographical limitation) services in order to adhere to contractual obligations and target the contractually agreed footprint. The use of GL has no bearing on the design of a CE device.

7.13.2 The delivery of content is most efficiently controlled at the first ‘server’ the P2P client contacts, which can be either a tracker server for download or a resource-allocating server for live streaming. At this central entity access to individual swarms/live streaming topologies can be controlled and changed on-the-fly. The central metadata repository should provide the possibility for broadcasters to adjust geolocation management according to their transmission and copyright requirements.

7.13.3 CE devices with or without copy management or DRM will be available on the market to support different business models. These devices should adhere to non-discriminatory standardised solutions such as defined e.g. by DVB CPCM (but not limited to this solution).

7.13.4 The content provider or a broadcaster should be able to inject live streams and on-demand files via P2P networks and users should be able to record these. Access to these streams and files should be controlled in order to allow or disallow redistribution or remote access from/to the PVR over the Internet. Another control mechanism may be required in order to provide services to the P2P network under additional conditions (for example, up until, say, 7 days after live transmission).

7.13.5 Transmission of files and live streams across P2P networks to the general public should be in the clear. After reception of such transmission, the content should remain in the clear when stored for further redistribution via the P2P network, unless encryption is required by the FTA content provider using e.g. the appropriate DVB signalling for FTA.

7.14 Network neutrality

7.14.1 Internet Service Providers (ISPs) should provide access to any legal content, including one based on P2P delivery, indiscriminately.

7.14.2 Any traffic shaping leading to a degradation of service quality (including bandwidth, packet-loss or jitter) compared to an average data transfer Quality of Service (QoS), should not be tolerated.

7.14.3 The use of port filtering to deteriorate service quality of P2P-based FTA reception also is not acceptable.

8. Existing commercial P2P services

Across the Internet, content can either be delivered as live (real-time) streaming or on-demand content (files). Content can be of any type: audio, video, data or any combination of these. Specifically, the most popular media applications and services in the Internet are the following:

- Linear and on-demand radio over the Internet
- Linear and on-demand television over the Internet
- Interactive and personalized applications
- Music downloads
- Movie downloads
- Media podcasting

These services may lead to an improved relationship with our audiences and offer new revenue opportunities

The following is a collation of information provided by the most commonly encountered commercial P2P systems.

[Editor's Note: Statements made in the following section are contributions from the companies' perspectives and may be of a promotional nature.]

8.1 Octoshape

Octoshape Streaming Solutions (www.octoshape.com), [7],[8],[9] delivers video and audio content from studio to online audience. Octoshape's principal mission is to lower the costs and increase the quality of live and on-demand streaming. Having created a P2P-based technology for streaming audio and video, Octoshape wants to better the user experience of streaming content, while cutting costs for the broadcasters. Using Octoshape grid-cast technology broadcasters can scale to accommodate millions of concurrent users, and yet provide a quality media experience. Traditional architectures transmit the whole stream directly from one server hosted by radio or TV stations. With Octoshape, the end-user receives fragments of the live programme from many peers in the grid. The result is a quality media experience, and a 97% reduction of bandwidth consumption.

Octoshape has been successfully providing live streaming services for the EBU's Eurovision Song Contest since 2005. In 2008 the EBU conducted a P2P Media Portal trial using Octoshape technology. The company has recently partnered with content delivery networks (CDNs) Highwinds Network Group Inc. and CDNetworks Co. Ltd. Its technology was chosen, along with the CDNetworks CDN, to help deliver 2 Mbit/s video streams over the Seoul Broadcasting System (SBS) Website during the Korean broadcaster's Olympics coverage. Turner's CNN.com is using Octoshape's P2P for Live Feeds.

The main characteristics of the Octoshape P2P system are:

- Cost-Effectiveness: Reducing price making live streaming profitable
- No 'server busy': Having unlimited number of simultaneous viewers
- No stream fall out: Using a multi fallback system
- Compatibility: Using the leading media codecs and players
- Statistics: Follow on-line number of users
- Instant access: The stream starts in seconds
- Reliability: Always a quality broadcast because the stream never stops
- Freedom: Users can choose their preferred players
- Availability: No more 'server busy' messages
- Speed: 'Plug-and-play' using transparent integration technology

8.2 RawFlow

RawFlow, www.rawflow.com, is a provider of live P2P streaming technology that enables Internet broadcasting of audio and video. RawFlow was incorporated in 2002. Its main office is in London, UK. A P2P computer network relies on the computing power and bandwidth of the participants in the network rather than concentrating it in a relatively low number of servers. When using this technology, the bandwidth requirement of the broadcast is intelligently distributed over the entire network of participants, instead of being centralized at the broadcast's origin; as the audience grows so do the network resources available to distribute that broadcast without adding any additional bandwidth costs.

The RawFlow ICD (Intelligent Content Distribution) server provides the first contact point for clients in the network. When initially installed, the ICD Server connects to the broadcaster's existing media server and begins receiving the stream, maintaining a buffer of stream data in memory at all times. It then begins accepting connections from clients and serving them with the stream. When launched by a user, the ICD Client first contacts the ICD Server and begins receiving the stream from it. The media player plays the stream as it is received by the ICD Client. If it has available resources, the ICD Client also accepts connections from other clients in the grid to which it may relay a portion or the whole stream it receives as requested.

The ICD Client monitors the quality of the stream it is receiving and upon any reduction of quality or loss of connection it again searches the grid for available resources while continuing to serve the media player from its buffer. The buffer prevents interruption to the playback and ensures that the end-user experience is not affected.

The ICD Server is always available as a last resort for clients that cannot find sufficient available resources in the grid. This guarantees a constant seed of the stream.

RawFlow's user-generated broadcast (UGB) technology offers a media platform which enables individuals, brands and communities to easily produce live and on-demand digital broadcasts, distribute them to multiple destinations simultaneously, engage with an audience through multimedia communications, and monetize their brand. RawFlow provides the P2P distribution technology for the Selfcast platform which went live in early 2007 and it has since then, attracted thousands of broadcasters and millions of live viewers.

Selfcast is a free and easy-to-use service for anyone, and since its launch it has been used by very diverse groups of broadcasters such as politicians, music events, sports clubs and Churches.

Selfcast utilises the basic functionalities of the UGB platform which include:

- One-click broadcast
- One-click recording
- Chat
- Scheduling
- Widget technology
- Embed functionality
- Channel pages
- Personalisation
- High Quality encoding
- Alerts

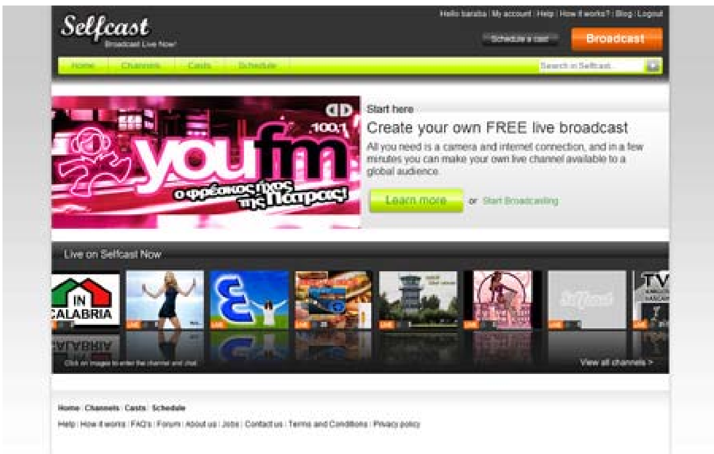

8.3 Abacast

Abacast, www.abacast.com, is a commercial, quality Hybrid Content Distribution Network (CDN), offering the content industry the ability to distribute and monetize live video, online radio broadcasts, video-on-demand (VOD), games and software. It was established in 2000 and it is based in the US.

It helps business models with revenue-generating features such as its Ad Injection System, Subscription Systems or Synchronized Ads. It has efficient live and on-demand delivery options including standard Unicast, P2P or a Hybrid combination of both. It has the ability to monitor the objectives with Understandable Analytics that are accurate, real-time and designed for business and media people. Brand promotion and a superior end-user experience is provided with premier and custom media players

Abacast provides the following services

- **Internet Television**: End-to-end solutions for live broadcasts, video-on demand, advertising integration, brand promotion and custom players.
- **On-line Radio**: Abacast provides the solutions required to make online radio streaming effective, legal, reportable and profitable.
- **Content Monetization**: Video and radio monetization services including pre-roll management, ad injection, banner ad sync and subscription management.
- **File delivery**: Complete solutions for high-volume games, software, patch, and video delivery.
- **Education/Distance Learning**: Archived or live course delivery services to anywhere on the globe.
- **Corporate Communications**: Enterprise solutions for corporate communications, sales training, compliance and certification.

8.4 BitTorrent DNA

BitTorrent DNA™, www.bittorrent.com, is a content delivery service that uses a secure, private, managed peer network to power faster, more reliable, more efficient delivery of richer content. BitTorrent DNA works with your existing CDN or origin servers, seamlessly accelerating your downloads or HTTP media streams.

With over 160 million clients downloaded, BitTorrent is the consumer standard for software and content distribution on the Internet. BitTorrent DNA extends the open BitTorrent protocol into a managed platform for commercial-grade content delivery.

Founded in 2004, BitTorrent is a privately held company backed by venture capital firms, Accel and DCM. The company is headquartered in San Francisco, California.

Peer Accelerated Content Delivery

BitTorrent DNA™ is the next step in the evolution of digital content delivery; it combines the efficiency and organic scalability of peer networking with the control and reliability of a traditional content delivery network (CDN). BitTorrent DNA™ uses one or more existing origin servers or CDNs to seed a managed peer network. Use of the peer network is tightly controlled by a specialized tracker operated by BitTorrent, Inc. and accessible to BitTorrent DNA™ customers through a web-based dashboard that provides control and reporting tools.

Assured Delivery

BitTorrent DNA™ is designed to complement existing delivery mechanisms, making the best use of both peer and infrastructure resources. Customers using BitTorrent DNA™ for downloads may associate with each object a quality of service (QoS) parameter that defines a required minimum bitrate. When BitTorrent DNA™ is used for streaming media, QoS is set automatically to ensure smooth playback with no buffering interruptions, while still making the most use of the peer network. Throughout the download process, BitTorrent DNA™ carefully balances its use of peer and CDN or server resources, downloading from all, in parallel, to meet per-object or streaming media QoS requirements.

Organic Scalability

By unobtrusively harnessing end-users' unused network capacity, BitTorrent DNA™ scales organically with demand, providing capacity exactly where and when you need it. Need delivery capacity in some particular corner of the world? BitTorrent DNA™ will give you capacity there. Do you occasionally experience unexpected spikes in demand that strain your delivery infrastructure? BitTorrent DNA™ automatically scales its delivery capacity with demand to ensure a consistently high-quality user experience.

Advanced Bandwidth Management

BitTorrent DNA™ runs quietly in the background with minimal impact to the end-user experience. Our proprietary transport technology leverages the full available network capacity of all paths without disrupting other applications. By detecting the presence of other applications, computers and devices sharing the consumer's broadband connection, BitTorrent DNA™ automatically moderates its use of the network to ensure that web browsing, voice over IP (VoIP), Internet gaming and other applications are not disrupted.

Friendly to Service Provider Networks

BitTorrent DNA™ contains a number of enhancements to mitigate the impact of peer networking on service provider networks. These enhancements include: BitTorrent's sophisticated congestion-avoiding transport technology; an intelligent peer selection algorithm that prefers peers on the same LAN, network or AS; and work with vendors of BitTorrent caching products to support local cache discovery. By keeping traffic local and non-congestive, BitTorrent DNA™ reduces long-haul and peering traffic for service providers, while improving the end-user experience.

Multi-CDN Acceleration

BitTorrent DNA™ is designed to complement existing delivery mechanisms, including content delivery networks (CDNs) and traditional web servers. To provide maximum flexibility and robustness, BitTorrent DNA™ can seed its managed peer network from multiple CDNs in parallel.

Multi-CDN Analytics

BitTorrent DNA™'s client-side telemetry and web-based dashboard provide performance visibility across all deployed content delivery solutions, including third-party CDNs. With BitTorrent DNA™, you see accurate reports of the actual performance experienced by your end users.

8.5 PPlive

PPLive, www.pplive.com, is arguably the largest P2P streaming video network worldwide. It was created in December 2004 by Huazhong University of Science and Technology, People's Republic of China. PPLive programmes are targeted at Chinese audiences. A majority of them are categorized as movie, music, TV series or live TV streaming. Also available are some specialties covering sports, news, game shows, etc. Most available programmes are in Mandarin, Cantonese or Korean. There is also an increasing amount of programmes in English, such as Hollywood blockbuster movies and popular American TV shows. All these English-speaking shows are hard-coded with Chinese subtitles.

In addition to PPLive, many commercial P2P TV services have been developed recently in China: *TVants*, *TVUPlayer*, *PPLive*, *QQLive*, *Feidian*, *PPStream* and *SopCast*). The majority of available applications broadcast mainly Asian TV stations, with the exception of TVUPlayer, which carries a number of North American stations including CBS, Spike TV and Fox News. Some applications distribute TV channels without a legal license; this utilization of P2P technology is particularly popular to view channels that are either not available locally, or are only available by paid subscription, as is the case for some sports channels.

Since the PPLive video stream depends on the network connection and numbers of peers, occasional glitches such as short pauses during the viewing or re-buffering are not unusual. In some circumstances, the stream may stop completely if the source video file crashes or not enough peers are available to establish a smooth streaming.

A recent report by Jin Li of Microsoft [10] states that:

'... current PPLive platform do[es] experience playback freeze for as long as 1 minute, and as frequent as 4 freezing incidents in a 7 minute interval...PPLive also incurs a relatively long playback lag. It is observed that PPLive can incur startup delay for about 20s to 30s for popular channels, and up to 2 minutes delay for unpopular channels, and some peers may watch frames in a channel minutes behind others'.

8.6 SwarmPlayer

SwarmPlayer is being developed within the EC-funded *P2P-Next* project (see §18). It is based on the BitTorrent protocol and allows a player to download movies, watch video-on-demand, and watch live video streams using one technology, while taking advantage of the popularity and maturity of existing BitTorrent clients.

The Project has completed the SwarmPlayer software development to support the above streaming modes, but require an audience to test it on. After all, P2P technology is designed to support thousands of users, and to properly test this, many users have to watch the same video at the same time.

In centralised video streaming systems, such as YouTube, a single set of computers provides the video to all viewers. Such a solution requires a massive number of computers (as YouTube has) to serve all of the videos to a large set of users.

Peer-to-peer technology takes a different approach. The video stream is served to a few users, after which users exchange and forward the video stream among each other. The users thus help serving the video, reducing or even removing the need for a central server park. Starting a YouTube-like system becomes orders of magnitude cheaper when P2P technology is used. The downside of peer-to-peer is that the quality is harder to control since the responsibility of forwarding the video is shifted from the central server park to the users themselves. If the users cannot or will not forward the video among each other, the quality of the system will suffer.

The goal of this trial is to examine how well SwarmPlugin scales when serving thousands of users on the Internet. It is not possible to reproduce thousands of users around the globe in a lab, so the *P2P-Next* project uses a trial network involving a large number of participants. It is hoped that in return for their participation, participants will receive a glimpse of what the future of online video might look like.

To understand how well streaming works with SwarmPlugin, statistics are being collected to understand what is going on in the network. Also, when the video is finished, several experiments are run, including a new NAT traversal mechanism, an ISP-friendly congestion control algorithm and a novel UDP-based light P2P swarming core.

P2P-Next is developing a platform that takes Open Source development, open standards, and future-proof iterative design as key design principles. By using P2P technology we aim to provide an efficient and low-cost delivery platform for professional and user-created content.

9. P2P Trials, Experiments and Experiences

9.1 Eurovision Song Contest

On 20 May 2006 the EBU performed a technical trial of live-streaming the Eurovision Song Contest (ESC) using the Octoshape P2P system. The following are some of the statistical and operational data obtained:

- A maximum of 15000 concurrent streams were logged
- There were more than 70000 unique users.
- No congestion was experienced through downloading the Octoshape media player plug-in.
- The MS Windows Media player was used.
- 3 quality levels were provided: 200 kbit/s (Q2CIF), 450 kbit/s (QCIF), 700 kbit/s (CIF)
- There were no major technical problems and almost no complaints from users.
- Network dependent quality – asymmetry was an issue.
- Neither rights management (DRM) nor geolocation (geographical restrictions) were applied.

The EBU repeated the experiment with Octoshape for the Junior ESC’06 and ESC’07. For ESC'07 the observed peak of concurrent streams was above 25000.

In May 2008, the Eurovision Song Contest was streamed live from Belgrade (Serbia) with the following statistics:

- More than 155000 unique visitors.
- 45000 concurrent streams.
- 338000 sessions.
- Total 171000 viewing hours.

9.2 EBU Member Services

RTVE has been offering two TV channels in high quality with Octoshape since the summer of 2006.

RTVSLO has been offering a number of services provided by Octoshape for some time. Today RTVSLO has 3 TV and 8 radio channels. All channels are in high quality.

RTP has been offering 1 TV channel in high quality with Octoshape since the beginning of 2007.

DW has been offering a high quality TV channel with Octoshape since 2006.

RNE has been offering three radio channels in high quality with Octoshape since 2006.

Many other Members are currently offering P2P live streaming in high quality. As such, high quality P2P streaming may no longer be considered an experiment, but a ‘normal’ 24/7 service to the public.

9.3 EBU Test at IBC 2006

As part of the 'EBU Village' stand at IBC 2006 (International Broadcast Convention, in Amsterdam) the EBU presented a live P2P streaming demonstration. Figure 10 shows the equipment setup.

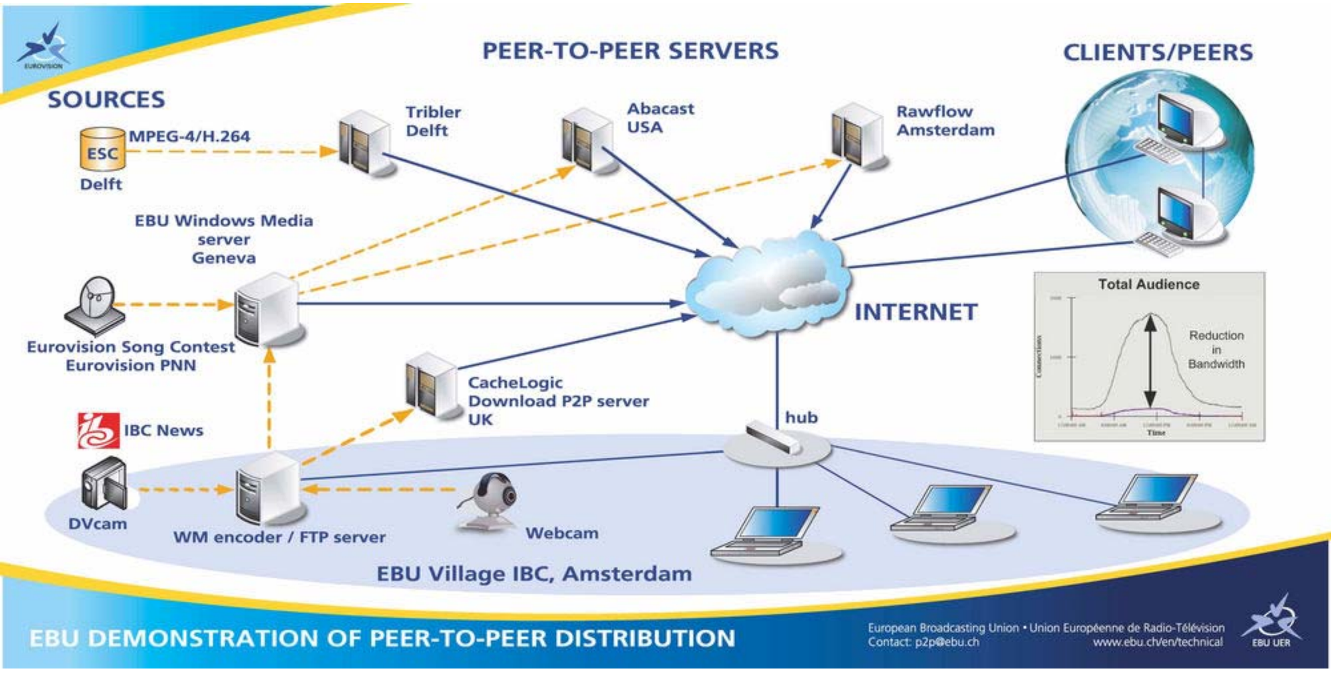

Figure 10: EBU demonstration of P2P distribution at IBC 2006

Three providers participated in this demonstration, RawFlow (in association with CacheLogic), Abacast and Tribler (TUD). Octoshape had been invited to participate, but it decided not to do so.

Each provider was fed by content served from the EBU Stream Farm in Geneva. Three stream channels and one file repository were set up:

IBC Daily News: a daily news programme of about 15 minutes length, produced by the IBC. A compressed version was sent to the Geneva servers over a secure connection and the programme was streamed in a loop all day at 800 kbit/s.

EBU HQ: a palette of programmes produced by the EBU at high bitrate (800 kbit/s).

EBU: a live stream showing the Eurovision PNN contribution news channel.

Velocix: the IBC daily file was sent to Cachelogic so that it was available for fast download in DVD-like quality using a P2P BitTorrent client such as Vuze. An attempt to produce an iPOD-compatible version in MPEG-4 failed because no software could be identified that ensured lip-sync after the conversion (several seconds of delay were observed with Imtoo, Cucusoft and Videora).

Five laptops were available to receive the streams and files on the P2P booth from the different providers. A sniffing software (NetLimiter) was used to monitor IP packets being received or sent by the individual laptops (or from a central server).

9.4 Prix Europa 2006 concert

The Prix Europa 2006 opening concert was given on 14 October 2006 in Berlin by a Portuguese World Music group called Gaiteros de Lisboa. On the occasion of this one-hour long concert, the EBU organized a technical experiment to distribute multichannel 5.1 audio. The main purpose of the experiment was to integrate a multichannel HE AAC real-time encoded sound programme with an Octoshape P2P streaming server and then provide it as a webcast across the Internet. Many who listened to the webcast were impressed by the high quality and smoothness (no interruptions) of the sound delivered. On the downside, however, the number of users of the webcast was unimpressive.

- Number of unique users: 52;
- Number of sessions: 914;
- Number of different countries: 13;
- Average session duration: 10 min 44 sec;
- Peak number of users: 11.

The experiment was significant because for the first time an event was ‘broadcast’ live in 5.1 multichannel audio across the Internet, proving its potential to address large audiences with high-quality surround sound.

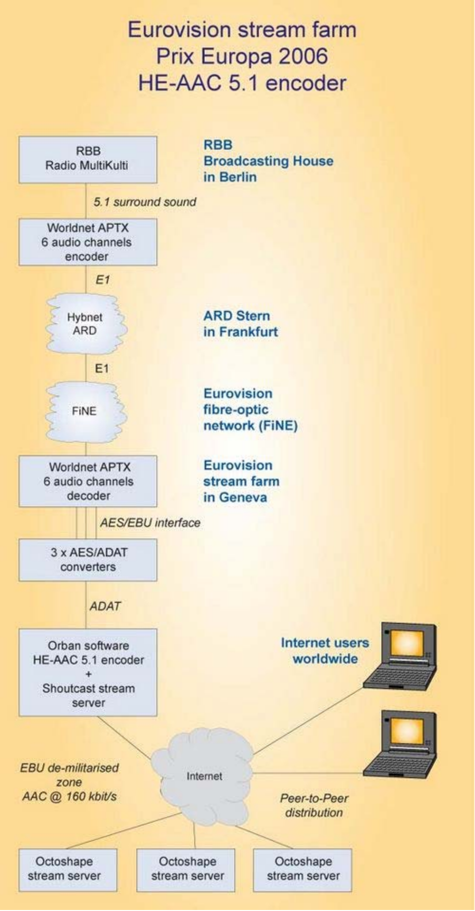

Figure 11: P2P experiment at the Prix Europa 2006

9.5 WDR evaluations of PPLive and TVants

The following outlines the results obtained when testing the functionality of the *PPLive* and *TVants* P2P live streaming systems at WDR, Cologne in 2006. PPLive and TVants are Chinese IPTV applications based on P2P. In their native markets, the video content that they distribute stems from channels such as CCTV 5 and Shanghai Sports, but also the BBC and in part, premium content such as the English Premier League football. Further special sport events such as the FIFA World Cup or the Olympic Games are also streamed live by PPLive and TVants. According to its own account PPLive had up to 500000 simultaneous users and the software was downloaded over 20 million times. Whilst the RealVideo codec only works with PPLive, both systems support Windows Media encoded streams using bandwidths of typically 200 – 400 kbit/s, which is commensurate with the typical upstream capacities of broadband users.

Users of PPLive receive an update channel list from the central PPLive channel list server. After selecting a channel the user asks root servers to retrieve online peers for this channel. Peers augment this list by sending their own actual lists of peers, so the content chunks can be shared with each other. The TV engine is responsible for downloading the content chunks and streaming the video to the local video player. Two buffers exist within local player memory; the PPlive TV engine buffer and the media player buffer. While a media player buffer is commonplace for streaming, the PPlive engine buffer provides an efficient forwarding of content-chunks to peers.

TVants is similar to PPlive but contains more features for analysis such as monitoring network activity, a graphical display of the buffer level and an event list. A search function allows finding specific channels. Furthermore the numbers of trackers and seeders, the shutdown and build-up of connections can be displayed. Figure 12 shows the TVants stream of the FIFA World Cup game between Italy and the Czech Republic.

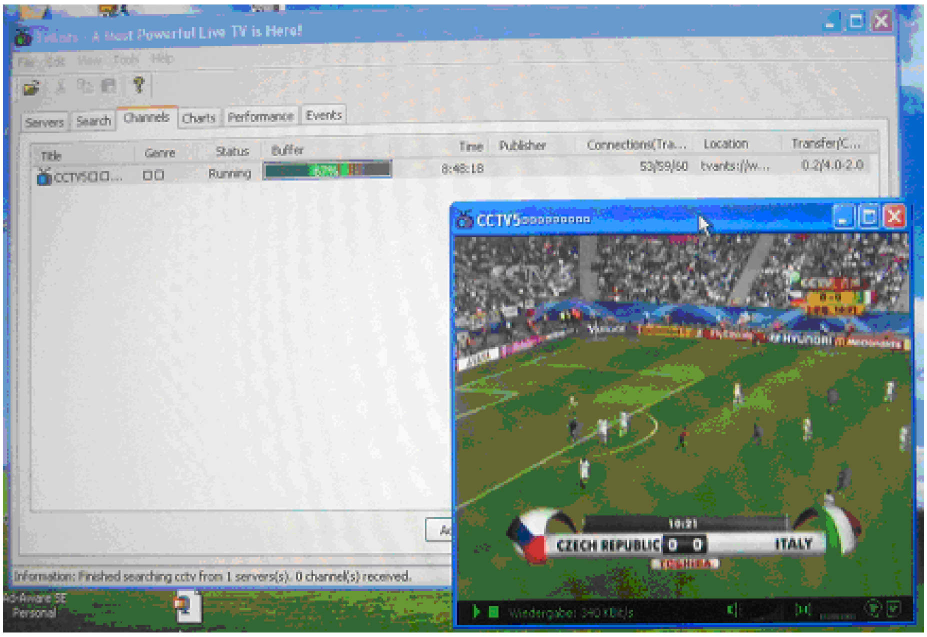

Figure 12: Screenshot of the TVants live stream of the FIFA World Cup

9.6 *Danish Broadcasting corporation P2P trials using RawFlow and Octoshape*

Trial with RawFlow during 2005-2007

From March 2005 to December 2007, DR performed a live trial with 3 Internet radio services from www.dr.dk. RawFlow was selected mainly because of better user-experience and easy integration with DR's HTML player. The trial was a parallel test where users had the option to use P2P or choose the normal stream.

There were three radio streams available; music and journalism, jazz (music only) and soft (music only). The format used was MS Windows Media Audio at 32, 64 and 96 kbit/s

Figure 13: Player integration, install RawFlow ActiveX plugin

Results of the trial:

- Simple player integration
- Simple server setup
- 100% server uptime
- Easy install and transparent for the users
- 90% end-user acceptance
- 5000 users peak (programme 3, 150000 unique/month)
- Close to 1 million client downloads (2007)
- DR saved 70-80% bandwidth from P2P

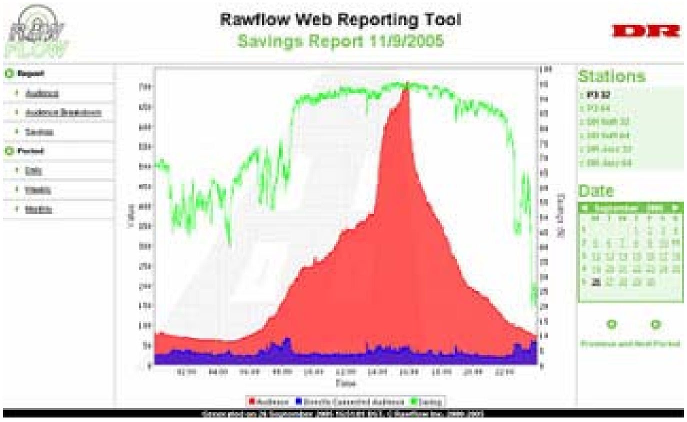

Figure 14: RawFlow reporting tool, 1 day view, red=P2P traffic, Blue=normal traffic, green=savings in%

Trial with Octoshape during UEFA 2007

The trial consisted of a live TV simulcast of the UEFA Cup in May 2007. Two of the semi-finals were streamed on the Internet only, whilst the final was simulcast on TV and the Internet. The format used was MS Windows Media Video at 1.3 Mbit/s.

Results of the trial:

- Easy to setup encoder
- Total 1700 users
- Peak ~ 1000 users
- Fast start and stable streaming
- DR saved about 60% bandwidth

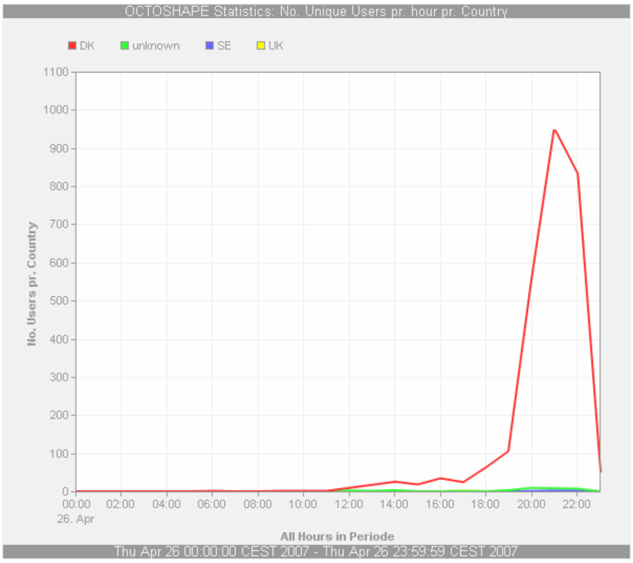

Figure 15: DR's Octoshape UEFA 2007 test

9.7 EBU P2P Media Portal

The trial on the 'EBU P2P Media Portal' ('EBUP2P' for short) was set up by EBU Project Group D/P2P (Peer-to-Peer) in autumn 2007 in order to perform technical evaluations of the P2P (Peer-to-Peer) technology provided by Octoshape. The EBU may choose to set up a one-stop Internet platform on the EBU home page, www.ebu.ch, through which any Members' television and radio channels may be made available to the general public worldwide using a broadband (WiFi) connection and a personal computer.

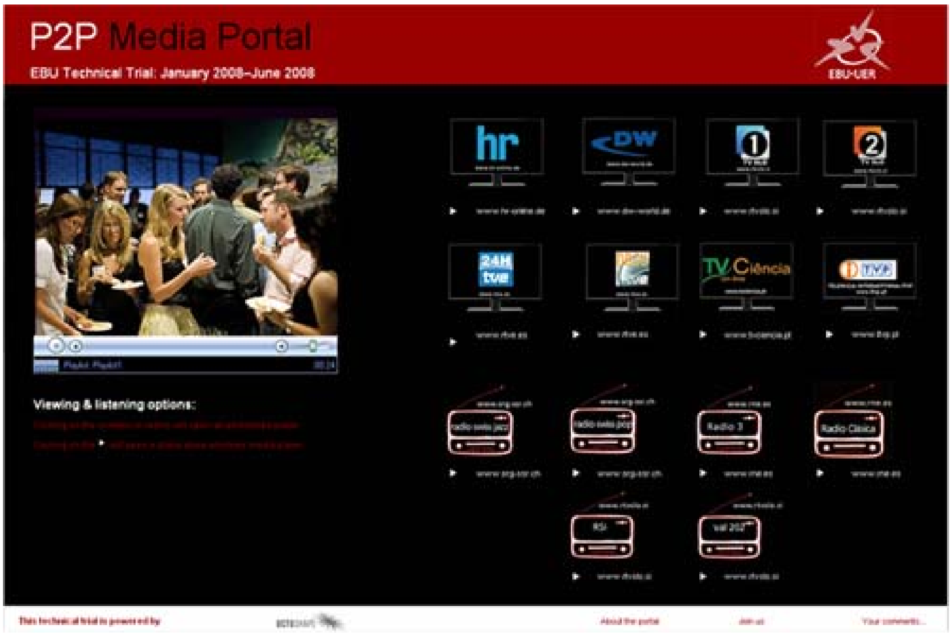

Figure 16: EBU P2P Media Portal interface GUI

The EBUP2P was designed to become a unique shop window for EBU Member organisations, showing their creative efforts and potentially forging their public broadcasting mission identity internationally. All active (and associate) EBU Member organisations could join the EBUP2P with their national, regional and local channels, both radio and television.

The principal conclusions of this trial can be summarized as follows:

- EBUP2P represented the state-of-the-art technical solution and fulfilled all technical and operational requirements in terms of the service quality, scalability, video and audio quality, accessibility, security and user-friendliness,
- EBUP2P had no technical limitations regarding the number of channels to be accommodated in the Portal, Members could flexibly join in and opt out at any time,
- EBUP2P fulfilled our requirements concerning copyright, by applying territorial filtering (geolocation) and watermarking,
- EBUP2P enabled a number of business models,
- EBUP2P was at the time the most efficient commercial proposition for media distribution over the Internet.

Table 5: EBU P2P channels (2007)

		CHANNEL	ORGANISATION	URL
TELEVISION	1	HR	Hessischer Rundfunk (HR)	www.hr-online.de
	2	DW - TV	Deutsche Welle (DW)	www.dw-world.de
	3	TV SLO 1	RTV Slovenia (RTVSLO)	www.rtvslo.si
	4	TV SLO 2	RTV Slovenia (RTVSLO)	www.rtvslo.si
	5	24H tve	RTV Spain (RTVE)	www.rtve.es
	6	DOCU tve	RTV Spain (RTVE)	www.rtve.es
	7	TV Ciencia on-line	TV Portugal	www.tvciencia.pt
	8	iTVP	Polish TV (TVP)	www.itvp.pl
	9	Taiwan TV		
RADIO	10	radio-suisse jazz	SRG-SSR	www.srg-ssr.ch
	11	radio-suisse pop	SRG-SSR	www.srg-ssr.ch
	12	radio 3	RNE	www.rne.es
	13	radio classica	RNE	www.rne.es
	14	youfm	HR	www.hr-online.de
	15	RSi	RTVSLO	www.rtvslo.si
	16	Val 202	RTVSLO	www.rtvslo.si

9.8 *EBU's Tribler experiences*

Several EBU Member organizations have been testing the open-source Tribler P2P system developed by numerous universities and coordinated by the Technical University Delft. Tribler is a social community that facilitates sharing through a P2P network.

When the Tribler application program is started it will automatically start searching other users that have Tribler running on their computer. When a connection is established it starts exchanging information. First it exchanges personal information (such as your avatar picture, your friends list, download history, etc.) and information about files that are available in the network. These files can be personal, shared files, but also files that have been received from another person.

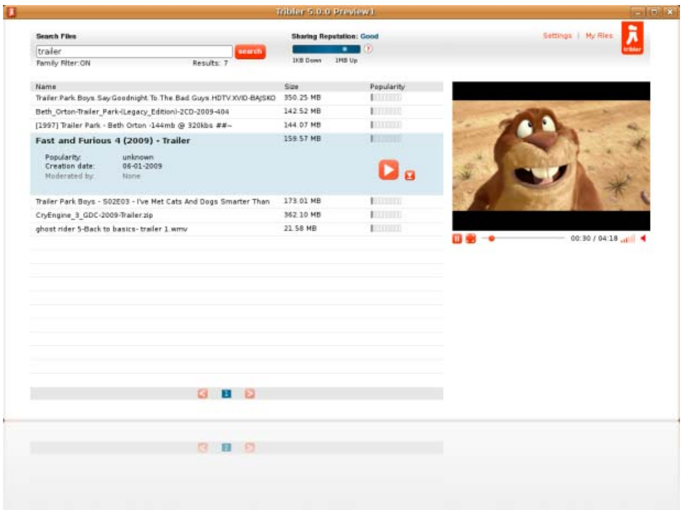

Figure 18: Tribler screenshot

The information about the discovered files and persons is available in the Tribler program. By browsing through the files and persons each user can find their preferred files and users. The Tribler program helps you by giving extra information about each item (whether it is a file or a person) and also shows what other users think about it. When you find a person you like, you can add him/her as a friend. Any file which you find interesting can be downloaded and will be available in your library.

In early 2009, Tribler was the only Open Source environment that supported all three modes of operation:

- Download the video and watch it afterwards (typical BitTorrent behaviour)
- Video-on-demand
- Live real-time streaming (web-cams, live TV broadcasts, etc)

Two EU Framework Programme 7 projects, with a combined budget of 26.26€ million are actively extending the Tribler foundation in the 2008-2013 timeframe. The first project, called *P2P-Next*, is developing a platform that takes as key design principles Open Source development, open standards and future-proof iterative design. By using P2P technology, *P2P-Next* aims to provide an efficient and low-cost delivery platform for professional and user-created content. It is focused on crafting a production-level Open Source reference implementation of 'next-generation P2P technology'. During Q3 of 2008 *P2P-Next* held a large public trial with 85,000 participants. This successfully tested the first unified P2P algorithm supporting live streaming, on-demand viewing, and HDTV download. Quality of Experience was measured to be very high with nearly no frame drops and video initialisation delays of the order of seconds.

9.9 IBC 2008 and IBC 2009 – P2P-Next trial using a Pioneer set-top-box

During the IBC 2008 in Amsterdam, the *P2P-Next* Project released its first version of NextShare, an Open Source P2P video delivery platform. NextShareTV is a STB (Set-Top-Box) implementation of the NextShare content delivery platform running on low-cost embedded hardware.

The *P2P-Next* project successfully released and tested the first Beta version of their P2P live streaming technology. This new technology allows broadcasting a live stream, such as a TV channel or webcam feed, to potentially millions of Internet users. Key is the bandwidth efficiency of this technology, by expanding the proven BitTorrent protocol you can stream to thousands of people using roughly the same amount of bandwidth as for a single user. This platform may enable large audiences to stream and interact with live and on-demand (VoD) content via a set top box or a TV receiver. In addition, it may allow audiences to build communities around their favourite content via a fully personalized system.

Figure 17: EBU *P2P-Next* demonstration team at IBC 2009

The harnessing of open source P2P video streaming technology like NextShare by low-cost consumer electronic devices represents a change and an exciting opportunity for the market. The next step is to agree Open Standards for interoperability between CE peer devices across Europe.

The P2P trial at the IBC 2008 was the world's first end-to-end streaming of live TV via a P2P network to a set top box using professional content at professional quality.

10. P2P Social Networks

One of the underlying principles of any P2P system is the participation of several peers in a common activity (for example, sharing a video clip). This topic is currently the subject of many studies and considerations. P2P social networks could potentially facilitate the introduction of several social features in content distribution, as follows:

- communication primitives
- strong peer authentication
- strong content integrity checks
- permanent storage of context
- semantic clustering
- recommendations
- reputations
- micro-payments
- Intuitive usage/insertion of own content, no matter which source type, codec etc (ingest should be profiled to align with standards – thereby enabling horizontal market integration)
- feedback and rating to myself as content creator, as well as feedback and rating to the community
- support for time-shifted-viewing, realized via a buffer or the 'swarm' itself as a buffer
- metadata, like ratings, tags, recommendation lists, social preferences
- community tagged browsing accessible and changeable/expansible by the users.

Social networking web sites, which allow users to create identities and link them to friends who have also created identities, are highly popular. Systems such as Facebook and MySpace utilize a traditional client-server approach to achieve this, which means that all identities and their social links (the entire social network) are stored and administered on central servers. Although this approach supports highly mobile user access - users can log-in from any computer - it also poses high dependence on predefined central server(s), which results in possible exploitation of private data. An alternative approach is presented in a paper produced by the Delft Technical University (TUD) [11]. This TUD protocol is based on a gossip (or epidemic) protocol, which uses a completely decentralized peer-to-peer system to create and store the social network information. The system is self-administered and works in a highly transient environment of peer availability. The design and implementation of a distributed social networking system is scalable and robust, allowing users to perform core social networking functions of establishing and removing social links without any requirement for centralized servers or administration.

10.1 Free-riding issue

The problem of free-riding arises if some peers in the P2P network do not forward the video content to other users [12]. A free-rider in a P2P network is a peer that consumes more resources (or downloads data) than it contributes (or it uploads little or nothing). The burden of uploading falls on a server and/or on some altruistic peers (but the latter may be too few to provide all peers with an acceptable quality of service). Consequently, the whole P2P network may suffer. Several solutions to motivate peers to contribute upload capacity have been proposed. In Give-to-Get (G2G), peers have to forward (give) the chunks of content received from a peer to others in order to get more chunks from that peer. By preferring to serve good forwarders, free-riders are excluded in favour of well-behaved peers. When bandwidth in the P2P system becomes scarce, the free-riders will experience a drop in their quality of service. Free-riders will be able to obtain video data only if there is spare capacity in the P2P system. G2G has been implemented by the TUD in Delft for VoD services using any video codec. The system splits the video stream into chunks of fixed size, which need to be played at a constant bit rate.

An important indicator of the peer's 'reputation' is the sharing ratio, i.e. the ratio of upload and download capacity over a given time period. Popular closed BitTorrent trackers such as www.TVTorrents.com require a set minimum sharing ratio for a peer to stay in the system and ban peers that upload content of low quality or spam. However, these trackers make use of a centralised management authority which arbitrates among peers. A fully distributed system for reputation management has been developed and deployed in the open source file sharing network Tribler by the Delft University (TUD). This decentralised protocol is called BarterCast [13]; the real-time upload and download statistics are broadcast to all peers. Each peer calculates its own reputation based on their local traffic information. The spreading of false and misleading information is minimised by applying the maxflow algorithm in computing indirect contributions of one peer to another peer.

Incentive mechanisms are essential components of P2P systems as they enforce peers to share their resources and participate. Recent P2P systems that distribute live media streams take their inspirations from BitTorrent which is more focused on sharing of files. Studies have shown that BT incentive mechanisms which may be suitable for file sharing are not well suited to streaming live media. To this end, new incentive mechanisms specifically designed for continuous media such as live streaming are required [14].

10.2 User data and privacy issues

In considering P2P-driven social networks it is critical to preserve user privacy versus sustainable advertisement-driven business models. On the one hand, many consumers would not like their past history of downloads being recorded and/or disseminated to whole of the network; on the other hand, many would like more personalisation of search results and recommendations. The latter case is similar and relevant to the scenario of targeted advertisements, where sponsors target a particular audience based on their past preferences (such as Amazon) and/or search keywords (Google's 'Sponsored Links').

With most online social networks, users are invited, at the time of registration or first usage, to fill out social features such as 'My Favourites'. Typical fields (in addition to optional personal data such as age, gender) are: movie-actors, music bands, brand names, books, colours, dream holiday destinations, movies, TV series and football teams.

This provides more dense connections among users, and using this information, similarity and correlation among them can be established despite having little or no history of download preferences. This thereby helps alleviate the above two sub-problems. A topic requiring further research is whether the free-to-air/Internet business model can work when users fill in their profiles, even in the absence of recording their history?

11. Consumer Electronics

11.1 P2P-enabled Devices

A number of CE products capable of receiving television services via the Internet have recently appeared on the market, and many more have been announced by the major CE vendors. Products include televisions, games consoles, set top boxes and PVRs. It is likely that many more types of device will become connected in the coming years.

It is assumed that Internet TV services will be adapted from the current versions intended for PCs to better suit consumer electronics devices. This will include both adaptation of the service front end to provide a suitable '3 metre' user interface; and possibly technical changes such as different codecs and security schemes to match the capabilities of CE products.

11.2 Implications of P2P for CE Devices

Compared to an Internet TV-enabled CE product implementing 'conventional' delivery protocols, the following additions or changes are required to enable P2P protocols:

- An implementation of the P2P client, including protocols and control
- Integration with browser and/or embedded applications (such as EPG)
- Ability to execute a global search for content to be consumed, that spans multiple content providers and service provider networks, and in addition spans all the content on cooperating device types (PC, STB, Mobile, Console)
- Sufficient storage internal or external storage for the P2P cache. This may already be present if the device is a PVR
- Sufficient CPU and memory resources to run the P2P client in parallel with other television functions, without impairing the user experience
- Possibly a graphical user interface to allow users to configure any necessary P2P-related settings and provide information such as uplink bandwidth usage, share ratios, etc.

In general, the larger the number of peers available, the more efficiently the P2P network will operate. Users may also gain some 'status' within an online community based on the amount of bandwidth they contribute to the P2P network. This could imply that the device should continue to function as a peer in the network while in 'standby' mode. However, this is in conflict with the trend in the CE industry to reduce power consumption of devices for environmental reasons. For example, Philips 2008 range of large LCD televisions has a standby power consumption of 0.15W [15]. The power required to run a P2P client in the device, including CPU, storage and network interface would far exceed this figure. If such a feature was present on millions of devices, the increase in overall power consumption would be significant, and may well attract the attention of regulators.

11.3 Motivation for CE vendors to build P2P Enabled Devices

The business model behind the manufacture of consumer electronics devices can very roughly be divided into two categories.

- 'Horizontal' market products, not tied to any particular service provider and sold in retail. These products typically implement open standards and in some limited cases proprietary components that are de-facto standards. This enables them to receive services from many providers.
- 'Vertical' market products, distributed by a service provider; typically subsidised by the

provider and only capable of receiving services from that provider.

The motivation to include a P2P client in a CE product will depend on the business model. The following sections discuss the conditions for P2P clients being added to Internet TV products in horizontal and vertical markets. It should be noted that many variations of these business models are possible and are likely to emerge. Time will tell which will be successful.

11.3.1 Horizontal Market Products

In a ‘horizontal’ market, the manufacturer’s interest is to make attractive products that users will want to buy, and to maximize the profit realised from selling these products.

In order to make a product attractive, it should be able to access a wide range of services including the most popular services for the market in which the product is sold. It should be easy to use the product to access the services and the quality of the pictures and sound should be as high as possible.

An important factor in maximising the profit for the manufacturer is that the product must be sold in as many markets as possible. Variations between models for individual countries are very expensive for manufacturers. Indeed, it is often not viable to develop and sell products that are specific to a single country, meaning that such products may simply not emerge.

The likely conditions for P2P clients appearing in horizontal markets are as follows:

- The P2P client must provide access to a wide range of services that would not otherwise be available, or at least access to the most important, popular services in particular markets.
- The P2P client must be applicable across multiple markets. In Europe, this means the same P2P client should be capable of receiving relevant services in many countries.
- The client side P2P protocols must be stable, so that the product does not require long term after-sales support. In general, the business models behind horizontal market CE products do not allow ongoing upgrade of products in the field. Note that this may change if there is some revenue stream to the CE vendor after the product has been sold, for example from services provided by the CE vendor or revenue sharing of paid-for services. Of course, only software upgrades are possible in the field – any requirements for new hardware cannot be met.
- It is strongly preferred that a single open standard emerges for the client-side interface to the P2P system. This greatly increases the chances of the above conditions being met, and in addition, openly standardised interfaces prevent lock-in by particular companies to their specific client (and indeed server) implementations. Such lock-in to a single source of components may not be acceptable to CE vendors.

11.3.2 Vertical Market Products

In a ‘vertical’ market, products are a part of a specific provider’s service offering, and are usually manufactured to that provider’s specification. The use of proprietary components or components not widely used by other providers is feasible. Indeed, features not available from other service providers may be an important differentiator for a vertical operator compared to their competitors.

However, vertical markets can also benefit from, and even be enabled by, standardisation. This is especially true for components that are not differentiators between vertical operators, as has been the case for content delivery protocols in the past. Where an open standard exists and is widely used, implementations become commodities available from multiple sources, or even as open source software. This reduces costs for the operator, avoids proprietary lock-in and ultimately leads to greater levels of penetration for the technology concerned, which benefits all stakeholders.

11.4 Internet TV Products

Below are some examples of existing and announced CE products with Internet TV features.

- Game consoles including Sony Playstation 3, Nintendo Wii and Microsoft X-Box 360. These products are all capable of receiving Internet TV services. Currently all are to a large degree 'vertical' in their nature, as the services available are controlled by the manufacturer and are made available with that manufacturer's branding. Software updates are regularly provided by the manufacturer, often adding new features. These are funded through sales of games (which include a licensing fee) and from revenue derived by the manufacturer from services offered. It is not believed that the consoles include P2P clients at present, but it should be possible to add them in the future via a software update, as the devices are very powerful and include storage.
- Philips Net TV – televisions including Internet TV features. A Services Portal will be provided by Philips, offering the user facilities for finding suitable services and content. Services and content will be available from multiple providers. Once the user has navigated away from the Portal, they interact directly with the service provider and content is distributed using that service provider's infrastructure. Services are browser based so branding is under the control of each service provider.
- AcTVila – a standard for Internet TV from CE products in Japan, based around a common Portal run by several CE vendors.

12. Telecom Views on P2P Distribution Systems and Services

12.1 Business Opportunities

Beside conventional IP Multicast and Unicast systems, P2P systems may offer new options for enhancing content delivery systems provided by Telcos[15].

Whereas IP Multicast provide high bandwidth efficiency and scalability within a single Multicast domain and CDNs are well established systems with a predictable and high QoS up to a distinct number of customers, P2P systems are simple to deploy, they do not need specific network provisioning and, in principle, scale gracefully with number of viewers.

It is necessary to point out that the P2P service providers have generally adopted a different business model than the traditional Telcos. In contrast to the traditional client-server applications, where communication is to and from a central server, P2P applications comprise peers that serve simultaneously as clients and servers.

P2P technology enables new traffic-intensive services with minimal infrastructure and seamless scalability and may helping transform Internet users from information-consumers to producers and broadcasters.

The dominance of P2P applications in the Internet traffic is stimulating broadband demand and creating a huge market potential for Telcos for the next years, but straining the current business model of Telcos on the other hand. Telcos want to increase operational efficiency and service quality while reducing costs, improve customer satisfaction and introduce new innovative services

[15] Telco is an abbreviation for Telecommunications Company

beyond pure bit-transport.

12.2 Challenges for Telcos

The emergence of P2P is posing significant new challenges to achieving efficient and fair utilization of network resources.

- P2P systems build independent overlay topologies on the existing network infrastructure whose traffic cannot be easily controlled by Telco's traffic engineering
- Increase in P2P related traffic does not translate to additional revenues for Telcos due to 'all-you-can-eat' flat pricing
- Potential congestion due to unexpected traffic peaks in some network areas leads to customer dissatisfaction and forces Telcos to invest in significant infrastructure updates without getting additional revenue streams
- 50-70% of P2P traffic flows to destinations outside the single ISP network, leading to high inter-ISP peering costs.

Presently, little or no direct communication and coordination between Telcos as network providers and P2P service providers takes place. Consequently, network resources are used inefficiently and P2P applications do not take into account the existing network topology. In particular:

- P2P nodes connect to neighbours randomly, irrespective of underlying network topology
- P2P causes a large amount of cross-ISP traffic and bottlenecks in traffic routing, leading to inefficiencies, and consequently both ISP and customers suffer.

12.3 ISP- P2P Collaboration

An improved collaboration between ISPs and P2P systems will certainly increase the overall system performance for ISPs as well as P2P systems aiming to decrease backbone traffic and bring down network operation costs by enabling service providers to communicate information about network conditions to client applications for the purpose of facilitating improved P2P traffic performance. Instead of selecting peers at random without exact knowledge of the underlying network topology, ISP- P2P collaboration can align P2P system topology with the Telco's network topology, reducing cross-ISP traffic and ISP peering costs, as well as improving P2P application performance.

Telcos can use P2P systems for optimized content-delivery networks (CDNs).

Different systems for ISP-P2P collaboration are under development such as:

- P4P activity conducted by DCIA (Distributed Computing Industry Association), see http://cs-www.cs.yale.edu/homes/yong/p4p.html.
- OnO activity by AquaLab Project, see http://www.aqualab.cs.northwestern.edu/projects/Ono.html.
- Peer Mapping Service activity provided by TU-Berlin/Deutsche Telekom Laboratories, see http://www.net.t-labs.tu-berlin.de/research/isp-p2p/.
- SmoothIT FP7 project, see §18 of this report.

The IETF started to standardizes an ISP-P2P collaboration solution in a new ALTO (Application-Layer Traffic Optimization) group. See
http://www.ietf.org/html.charters/alto-charter.html.

12.4 Integration of P2P Client at Customer Premises Equipment

In addition to the network related components (content ingest, content management and control, network caches), P2P clients located at the customer side (home network) are an essential part of P2P based content delivery systems provided by Telcos.

There are three main options to integrate a P2P client into existing in-home components, each having its advantages and disadvantages:

12.4.1 Option 1: P2P-TV Client on STB

- P2P TV Client is located at STB which is connected directly via Home Gateway to P2P network
- All CDN/P2P related traffic (management & control, streaming content) should be routed to the STB via the Home Gateway and In-home Network
- Extended effort to assure Quality of Experience (QoE) necessary
- STB should be equipped with sufficient computing power and storage to host and run the P2P TV Client
- P2P Client and Home Gateway have to deal with Firewall/NAT issues of CDN/P2P solutions, especially inbound connections are an issue
- In-home Network shall provide direct access to P2P-Client at the STB for Management & Control by the service provider (not in the case of certain horizontal market deployments)
- Collaboration between service and network provider can enhance QoE

12.4.2 Option 2: P2P-TV Client on PC

- P2P TV Client is located at PC which is connected directly via Home Gateway to P2P network
    - This is current state-of-the-art solution, nearly all available CDN/P2P solution support this approach
    - Several solutions offer P2P client as Browser Plug-in but more commonly P2P Client as separate software application
    - Customer needs to install P2P client software at her/his PC
- All CDN/P2P related traffic (management & control, streaming content) should be routed to the STB via the Home Gateway and In-home Network
- Extended effort to assure Quality of Experience (QoE).
- P2P Client and Home Gateway have to deal with Firewall/NAT issues of CDN/P2P solutions, especially in-bound connections are an issue
- In-home Network shall provide direct access to P2P-Client at the STB for Management & Control by the service provider

12.4.3 Option 3: P2P-TV Client on Gateway and STB

- P2P TV Client functionality is distributed, between Home Gateway and STB
    - All CDN/P2P related traffic (management & control, streaming content) is terminated on the Home Gateway
    - Home Gateway also hosts storage for content buffering and a streaming server (UPnP media server) to deliver a content stream in the home network
    - Based on UPnP-AV algorithm a media renderer running on the STB can access the content

provided by the media server decode and replay it

- All CDN/P2P related traffic from/to open Internet is decoupled from Home network
  - No P2P management & control traffic is transmitted in the home network
  - Firewall/NAT issues of CDN/P2P solutions can be avoided
- Service provider can access and manage the P2P-Client directly without user interaction
- No specific software is needed for STB – current DLNA enabled STBs can be used in principle
- Home Gateway needs additional computing power and storage to host P2P – client

Table 5 is a comparison of the three optional approaches described above.

Table 5: Comparison of the different approaches

	P2P-TV Client on STB	P2P-TV Client on PC	P2P TV Client on Home Gateway / STB
Pros	Compact solution Hardware and Software requirements nearly met with current STB already - effort for integration is limited Approach already addressed by several solution providers Standardization activities started	Compact and well established solution Hardware and Software requirements also met with less powerful PCs Standardization activities started Streaming to additional devices based on known technologies possible	Direct access to Gateway including P2P-Client facilitate complete service management for ISPs Decoupling Internet traffic and In-home Network traffic Better Support for QoE only media streaming traffic in the home network; no P2P control traffic Firewall/NAT issues easier to handle In principle all Streaming client enabled devices can be addressed
Cons	Complete P2P control traffic should be routed through home network (QoS issues) Control of P2P client in the home network more difficult No complete solution for Firewall/NAT issues available No support for streaming to additional in-home devices	PC is not target Device for the living room Complete P2P control traffic should be routed through home network (QoS issues) Control of P2P client in the home network more difficult No complete solution for Firewall/NAT issues No support for streaming to additional in-home streaming client devices	New Home Gateways with enhanced functionality necessary (enhanced Hard- and Software Requirements) Currently approach not addressed by solution and service provider

13. P2P-related Geolocation Systems

A geo-restriction (geolocation) system allows delimiting geographically multimedia delivery over the Internet, so that end users could access the media only within pre-determined areas, such as national territory or the EBU Broadcasting Zone. Geo-restriction services are achieved by making use of the assignment of IP addresses to a geographic location. The IP address ranges are assigned by the Internet Assigned Numbers Authority (IANA) and its regional assignment agencies. This identification can be improved by means of checking the browser time, browser language, version, etc.

The use of GL services may facilitate the successful resolution of webcasting rights issues and may be helpful in EBU negotiations with the various sports federations. Using P2P based Live-Streaming and (Progressive) Downloading, access to content offered by the broadcaster can be administrated by identifying the geographic position of the requesting user. This control can obviously be best achieved by a central entity.

The GL system chosen must satisfy some minimal technical and operational requirements. Those agreed by the EBU members are given below

13.1 Accuracy

While the accuracy of a decision based only on a single criterion as the IP address is often not accurate enough, the probability of a correct determination of location is increased by combining it with the evaluation of the language version of the accessing browser and media player, the time zone etc. Generally accuracy thus is a property of the GL system that qualifies how successful the determination of the end user's location is. For national and supranational level accuracies of 99% or more are offered, while values of 95% are often accepted by right owners (except for premium content).

The technical notion of accuracy includes detection of both false positives and wrong negatives. It includes reproducibility and consistency of the results. Unfortunately the numbers given for accuracy usually do not specify the number of wrong negatives, which could be problematic for the public remit of broadcasters.

13.2 Robustness

A GL system should be able to withstand stress such as when heavy traffic is experienced as a result of concurrent requests.

13.3 Security

A GL system should disallow any illegal or illicit usage (e.g. hackers). Visitors with unknown addresses (anonymous visitors) should be rejected, redirected and given a pop-up caption. Visitors who may attempt to access the media on the Member's site via known or unknown proxies should be disallowed access. To this end, the GL system should be able to identify all known and emerging proxies as well as anonymizers and put them in the « blocking area ».

13.4 Secure connection

A certified secure mechanism has to be included in the connection to the stream server. The security solution proposed should be standardized as an open standard and be proven.

13.5 Server side components

There should be no executable or server components that would need to be installed on each server. The script needed should be part of the HTML page.

13.6 Interoperability

The GL system shall be interoperable with any operational platforms used by EBU Members including Windows, Solaris and Linux and including the different types of streaming servers (e.g. Windows Media, RealNetworks, QuickTime). Furthermore the GL system architecture should allow one or several instances of the GL database running on one or several servers, as required. Upon registering these databases, the load should be spread in a uniform way among these instances.

13.7 Flexibility of coverage

The coverage of the allowable GL area should be flexibly modifiable by each EBU Member in order to suit its specific requirements on an event-by-event basis:

- National borders
- Part of the country (e.g. German Länder)
- Group of countries (e.g. common-language countries)
- Pan-European coverage
- European Broadcasting Zone (including, optionally, some additional countries)

13.8 End User Access

The end user should be able to access the streaming media published by an EBU Member only via the Member's web site by clicking on a corresponding hyperlink. Direct access through the media player by cutting-and-pasting the URL should not be possible.

14. P2P-related Metadata

Content distributed in a P2P environment needs to be described by a suitable metadata specification. At the time of writing this document, there is no internationally agreed P2P-related metadata standard. The closest to such a standard is a proposed *P2P-Next* Rich Metadata Specification [16].

The *P2P-Next* Rich Metadata specification provides a minimum set of attributes which are necessary to describe a content item in the Next-Share system.

This Specification consists of the mandatory Core and Optional metadata. Splitting the metadata in Core and Optional metadata provides flexibility to adjust the metadata system to any business model that may be utilized for the content item.

The Core metadata contain only those attributes that are essential for all business models and that do not change over the course of time. The Core metadata is packetized together with the actual media content and the hash values for the .torrent-file are created based on this package.

In addition to Core metadata, several types of Optional metadata are specified, like payment or advertising metadata.

To ensure compatibility with as many applications and existing metadata collections as possible, the attributes of the core and optional specifications of the *P2P-Next* Rich Metadata specification are mapped to three state-of-the-art metadata standards: TV-Anytime, MPEG-7, and URIPlay. Additionally, the Channel Metadata extensions provide a solution to describe the programming guide of a TV channel such as BBC 1.

The two-step metadata creation process is illustrated in Figure 18.

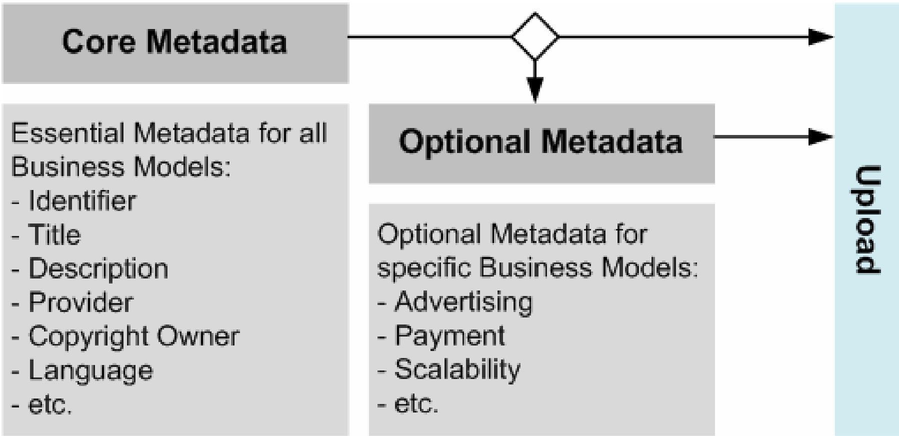

Figure 18: A two-step *P2P-Next* metadata creation (Courtesy: *P2P-Next* project)

The *P2P-Next* Metadata Specification also include a 'Signpost Specification' called Kendra, which provides a concept for content providers to ingest their content into the Next-Share system supporting various metadata schemes The Signpost implementation enables the translation of the metadata of content providers to a scheme compatible to the *P2P-Next* Rich Metadata specification and thus provides a convenient way for content providers to ingest all their content in a Next-Share compatible way.

Another part of the *P2P-Next* Metadata Specification represents the Content Collections Specification which describes how the content can be announced to interested parties utilizing so-called 'Atom feeds'. The specification describes the discovery feeds, which provide an overview of the available live programmes and on-demand collections of a single content provider, and the content feeds, which provide more details for specific content collections.

Table 6: *P2P-Next* Core Attributes

Attribute	Description
GUID/CRID	A unique identification of the file.
Title	The title of the file.
Description	A description of the file's content.
Provider	The user that ingested the file into the P2P system.
Duration	The duration of the content in a suitable time unit, if the content has a timeline, or none if the content has no timeline.
Chapters	An indication of the chapters of the content, or none if the content has no chapters.
Creation Time	The time when the content was produced.
Language	The language of the content.
Subtitles	The language of the subtitles of the content, or none if there are no subtitles.
Captions	The language of the captions of the content, or none if there are no captions.
Age Appropriateness	The rating in respect of the age appropriateness of the file.
AV	The AV attributes of the file, like the file-format, the bit-rate or the frame-rate of the content, etc.
Originator	The original producer of the content.
Genre	The genre of the content, e.g., a documentary.
Series/Episode	Describes the series and the episode of the content, if applicable.
European Content	Describes if the content has been produced in Europe.

Table 7: ***P2P-Next*** **Optional Attributes**

Cat.	Attribute	Description
Advertising	Business Model	Describes the business model applied for consuming the content.
	Live Content	Describes if the content is distributed as live transmission.
	Target Group	Provides information about the target group. Such information include the age, gender, country, etc.
	Allow Advertising	Specifies if advertising may be provided together with the content.
	Circular Content	Describes if the users are allowed to redistribute the commercial content.
	Ad Type	Describes the type of advertisement.
	Streaming Type	Describes how the advertisement provided, i.e., in- or out-stream.
	Ad Format	Describes how the advertisement is integrated into the content, e.g., as overlay.
	Content Format	Describes the content format of the advertisement, including its resolution, etc.
Payment	Payment Identifier	A second identifier of the content, as a payment system should not be aware of the content for privacy reasons.
	Price	The price of the content.
	Currency	The currency for the price.
	Payment Recipient	The recipient of the payment, if the content is not offered for free.
	Accept Donations	Describes if voluntary donations for the consumption of the content are accepted.
	Advanced Information	A link to a website with further information on the payment system.
Scala-bility	Layer Information	Information about the layers of a scalable bit stream.
	SPS Information	Information about the Sequence Parameter Sets (SPS) of the bit stream.
	PPS Information	Information about the Picture Parameter Sets (PPS) of the bit stream.

In the table above, optional attributes for three categories are provided. The attributes provided for advertising are used to describe properties of the commercial content and how they can be displayed together with the actual content. The payment attributes provide the information needed for acquiring a specific content. The attributes provided for scalability are only required for scalable bitstreams and provide the scalable header information to the packetizers and de-packetizers of the P2P system.

These optional attributes provide an initial metadata support for the *P2P-Next* business models, but each of the categories might be enhanced or additional categories might still be added in the course of the project. Please note that also additional metadata that are not described in this specification might very well be used within the *P2P-Next* system.

15. P2P-related Digital Rights Management Systems

Digital Rights Management (DRM) is a system for protecting the copyright of data by enabling secure distribution, access or usage to authorized users. There are two methods of preventing consumers using content illegally: a) hardware-based and b) software-based methods. These methods are not specific to P2P but generally apply to any media distribution across Internet. Software-based solutions are more flexible than hardware ones, and involve encrypting or marking the content with a digital watermark. Once the content is encrypted, a license is required allowing (or not) the user to use the media. The license consists of a key to unlock the content and a set of media usage rights (e.g. play only, play and record, single play, play for 30 days, etc.).

16. Legal and Regulatory Matters pertaining to P2P[16]

This section addresses some specific P2P-related legal issues that may be relevant for Internet delivery of TV and radio. For a long while, P2P has been synonymous with illegal file sharing activity akin to those of Napster and Kazaa. The music industry has been taking legal action against file-sharers and against companies that have made P2P file-sharing possible, but has largely failed to stop the massive scale of copyright infringement that happens over P2P networks.

Recently, however, the situation has changed. According to Out Law News, http://www.out-law.com/, research published in July 2009, illegal file-sharing activity amongst 14 to 18 year olds has gone down from 42% to 26% over the past two years. The research demonstrated that the music industry could best battle piracy by creating attractive legal routes to the acquisition of music, whether this is an unlimited streaming service such as Spotify or a service such as Virgin's which aims to offer unlimited MP3 downloads as well as unlimited audio streams. Young people find services such as YouTube much more convenient for checking out new music than file sharing.

The former European Commissioner for the Information Society, Vivian Reding, launched a Digital Europe Strategy in 2009, in which she called on the content industries of Europe to create services which compete with piracy. She declared that *'...Internet piracy is a vote of no-confidence in existing business models and legal solutions – it should be a wake-up call for policy makers'*.

16.1 Distinction between Scheduled versus On Demand content

It has been shown in this document that audiovisual (AV) services delivered over P2P networks can carry a variety of different types of service, including live streaming, progressive download, on-demand streaming, on-demand download, catch-up and podcasting services. The Audiovisual Media Services Directive (AVMS Directive), formerly called the 'Television Without Frontiers' Directive, is based on the assumption that a clear distinction can be made between non-linear services - Video on Demand (VoD) and linear services (scheduled programmes).

The question is whether, and more importantly, how, the AVMS Directive applies to these AV services. This question is important because the rules applied for linear services are somewhat tougher that those applied for non-linear services.

Under the definitions in the AVMS Directive it is clear that catch-up and podcasting services come, in principle, under the rules for non-linear services. However, it should also be noted that the definitions in the AVMS apply for the purposes of this Directive only.

It is important to realize that, despite the distinction, both linear and non-linear services are covered by the new AVMS Directive, which is already recognition of their convergence. The Directive maintains the principle of technological neutrality, so that the rules applicable to linear services apply in a horizontal manner, independently of the technical modes of signal delivery (over-the-air, by satellite, cable, broadband, microwave, telephone line, format, resolution, screen-size, protocol, etc.). Transmissions via the Internet or mobile phones which are not triggered by the consumer himself thereby remain a broadcasting act, whether the technology used for the act is based on P2P delivery, via point-to-multipoint or otherwise. The new notion of linear services has therefore safeguarded broadcasting from becoming a mere technology-dependent activity and is therefore a crucial achievement also vis-à-vis the use of P2P technology by broadcasters.

[16] This section was produced with the assistance of the EBU Legal Department

16.2 European content

Article 3(i) of the AVMS Directive states that on-demand audiovisual services should promote European works but does not impose any quotas. The quota provisions only apply for television services or linear services but have been unchanged. Only certain categories of programmes, such as feature films, television films and series, are subject to the quota system. The following are excluded from the system: news, sports events, games, advertising, teletext services and teleshopping. The promotion of European works should not apply to user-generated content. However, Member States are free to adopt stricter or more detailed rules for media services provided under their jurisdiction. They had until the end of 2009 to implement the provisions of the AVMS Directive in their national law.

16.3 Advertising rules

In the AVMS Directive only quantitative rules based on insertion and hourly limits for advertising exist, and they have been relaxed for television advertising. No quantitative rules exist for non-linear services. For example, the quantitative limit on advertising spots, which is the most relevant in practice, is 12 minutes per hour. Moreover, the 20-minute rule (which required that at least 20 minutes elapse between each successive advertising break within a programme) has been abolished. In addition, the principle of separation of advertising and editorial content has been maintained but has been reformulated in a way which opens up more possibilities for the use of new advertising techniques.

Thus, Members States are free to adopt more liberalized rules for non-linear services, bearing in mind that it is ultimately the user who decides. Too much advertising kills advertising.

However, all media audiovisual services (linear and non-linear) are subject to ‘qualitative’ rules (on the protection of minors, on human dignity and on public health).

16.4 The (future?) role of ISPs

European countries seem to have different approaches towards Internet users who share illegal content and exceed certain limits of file sharing. France, for example, seems in favour of the so-called graduated approach which allows, upon the decision of a special body, Internet service operators to remove a customer’s Internet connection after repeated offences following on from a warning. Sweden has not adopted this graduated approach, but allows content owners, as in other countries, to sue infringers before the courts. It remains to be seen what the developments in other Member States are. At the moment, there is no specific legislation on this issue envisaged at the EU level.

It must be recalled that the ‘graduated approach’ favoured by France is not officially linked with the EU Enforcement Directive, but is an initiative on the basis of a Government-commissioned report (the ‘Commission Olivennes’) for which most media stakeholders in France - as well as the French association of ISPs - have provided their express support. On the other hand, consumers - and recently the European Parliament in the current revision of the European telecom rules - have opposed this approach, arguing that a costumer's Internet connection should not be cut off unless there is a prior decision by the judiciary. Whether such a gentleman's agreement - and broad support - could be obtained in other countries remains to be seen, given that, so far, the proposal has not been openly embraced in Europe. Clearly, several categories of right-holders consider the ISPs the best-placed to act against unlawful file-sharing (given that the ISPs profit from their subscriptions). It could be argued that, although the graduated approach seems to include a practical remedy to reduce, at least over time, the number of ‘innocent’ copyright infringers (those uploading protected content without authorization but without knowledge that the act is unlawful), it could well be that to achieve that effect a warning system as such would be sufficient. After all, the main difficulties with the proposed final measure (the annulment of the ISP subscription) are

1) its appropriateness, as such annulment would cut off complete access to a network which is used for many other purposes too (and possibly by other persons within the same household),

2) the replacement of judicial relief by administrative proceedings (and therefore shifting litigation costs to society as a whole), and

3) its presumably limited effectiveness with regard to professional infringers who would probably find ways of continuing to keep their activities anonymous.

For example, the United Kingdom Government recently threatened the UK ISP and broadband providers with a compulsory levy if they did not come forward voluntarily with a warning system.

This matter is likely to be debated again later, as it could be one of the major issues when the e-Commerce Directive (which includes the most relevant liability provisions for ISPs) goes through its revision process.

16.5 Global Internet Licence

The idea of a 'global Internet licence', as originally proposed some years ago by consumer organizations and musicians' societies in France, during the debate on the revision of new French copyright law, has not yet achieved support from other stakeholders in the copyright community.

This is not surprising because it is thought that such 'culture flat rate' should then apply to all works and protected matter on the Internet, without bearing in mind that music is used and consumed differently (in particular, repetitively over a long period of time) from, for example, films and television series. Given the peculiarities of music usage, such a flat fee concept would thus not be appropriate for all types of works. Moreover, the conditions of such permitted use do not need to be the same for all works, as it could be liable to undermine legal (and paid-for) offers via the Internet.

For example, at first sight the nature and modes of exploitation of audiovisual works (such as feature films) would be unlikely to fit easily into a collective fee scheme. However, the underlying concept of a collective 'flat fee' payment has certain similarities with existing (collective) licensing practices as regards massive use of music (e.g. by broadcasters) and is therefore certainly worth analysing and exploring in more detail. This is particularly the case since any other legal remedies to Internet piracy of music have so far proven ineffective or inadequate. Such a concept, if thoroughly developed, could also be more favourable to the less well-known musicians and would then be a far better approach than, for example, the idea of extending the term of protection for performers and record producers.

16.6 Must carry

Must-carry rules seek to ensure that consumers have access to a wide, varied range of radio and television channels and services. The rules are thus an important instrument whereby Member States can guarantee media pluralism, cultural diversity and consumer protection. They remain essential in the digital environment.

In the context of the review of the Telecom Package, the European Parliament followed, in its second reading (6 May 2009), the Council position not to extend the potential scope of must-carry rules to non-linear audiovisual media services. Nevertheless, the text now at least ensures that all complementary services, and not just accessibility services (as in the Commission proposal), can be covered by must-carry rules, and that there is no mandatory review every three years.

Must-carry rules are of potential benefit with regard to not only traditional cable television networks but also other 'closed' networks, where all media services are provided/controlled by the service/network operator (such as IPTV services offered by broadband access providers). On the other hand, must-carry rules are not an appropriate instrument with respect to 'open' networks, where there is open access and where media services are provided/controlled by the media service providers themselves (e.g. on the open Internet). For the latter, the principle of net neutrality is more relevant.

16.7 Net neutrality

Must-carry rules are not an appropriate instrument with respect to 'open' networks. As the Internet is an essential platform for broadcasters and other audiovisual media service providers for the distribution of non-linear audiovisual media services to the general public, the open architecture of the Internet, underpinned by new legal concepts of 'net neutrality' or 'openness of the Internet', is of crucial interest. Must-carry rules and net neutrality principles can thus be seen as complementary instruments.

Broadcasters are expected to use all available platforms including terrestrial, satellite, cable, broadband and the Internet for the distribution of their non-linear audiovisual media services. The Internet is an indispensable platform. One problem that has already emerged in several European countries relates to the recurring demands by network operators that broadcasters (and other audiovisual content providers) should pay if they do not want their content to be slowed down on the Internet. This could lead to a dramatic increase in distribution costs for broadcasters and other content providers.

The concept of net neutrality is based on principles such as transparency of network management, minimum requirements as regards 'quality of service', and non-discrimination with respect to content, applications, services, providers and users. Net neutrality is thus likely to play an important role in facilitating the distribution of (linear and non-linear) audiovisual media services to the general public, irrespective of whether these are provided by public service media, by commercial media or by individual users'.

The discussion on net neutrality has developed in Europe only recently, in contrast to the United States. This may be explained by the different business models used by Internet access providers in the US and in Europe (with the effect that problems with net neutrality first arose in the United States), and by the different regulatory environment for the telecom sector in Europe (which has been harmonized, as far as EU and EEA Member States are concerned, by a bundle of EU Directives which are generally referred to as the 'Telecom Package').

Under the current EU Regulatory Framework for electronic communications services and networks, it is possible or even mandatory for national regulatory authorities to impose, on operators designated as having significant market power, certain obligations regarding access or interconnection (see Articles 8-13 of the Access Directive). However, such measures primarily serve to ensure effective competition among telecom operators, and do not directly protect end-users.

The new telecoms regulatory framework adopted in 2009 ensures that European consumers have an ever greater choice of competing broadband service providers available to them. Internet service providers have powerful tools at their disposal that allow them to differentiate between the various data transmissions on the Internet, such as voice or 'peer-to peer' communication. Even though traffic management can allow premium high quality services (such as IPTV) to develop and can help ensure secure communications, the same techniques may also be used to degrade the quality of other services to unacceptably low levels. This is why, under the new EU regulatory framework, National Regulatory Authorities are required to promote '*the ability of end-users to access and distribute information or run applications and services of their choice.*' This sets a very important principle for net neutrality, as it recognises and safeguards the basic freedoms of Internet users [17].

The new framework explicitly foresees the possibility for National Regulatory Authorities, after consulting the Commission, to set *minimum quality levels* for network transmission services so as to promote 'net neutrality' and 'net freedoms' for European citizens. This should ensure that traffic management and possible prioritisation does not lead to degradation of content and services provided by non-commercial actors or by new entrants.

In addition, thanks to the new *transparency requirements*, consumers will be informed – even before signing a contract – about the nature of the service to which they are subscribing, including traffic management techniques and their impact on service quality, as well as any other limitations (such as bandwidth caps or available connection speed).

16.8 Country-of-origin-of-initial-transmission

Another legal issue that would need to be considered in the context of Internet TV services is the question which law should apply to those services, given that, if the public part of the Internet is involved, the initial transmission of the content may cross national borders. For media law purposes, the AVMS Directive upheld the rule under the former Directive that the pertaining law is that in the country in which the originator of the broadcast resides (with certain exceptions for services specifically targeting foreign countries). A similar rule exists under the 1993 Satellite and Cable Directive for satellite services (the act of communication is determined by the law of the country of the satellite uplink). Much can thus be said for applying the same principle to Internet transmissions.

17. Discussion on xDSL Network Asymmetricity

The vast majority of Internet use today is for high speed download from a service provider to the consumer. Where business users required a symmetrical service, the initial solution was the Integrated Digital Services Network (ISDN). ISDN offered guaranteed quality of service (QoS) and an equal data rate in both down (from server to user) and up (from user towards server) directions. As user content has increased in both directions, ISDN was no longer sufficient.

To this end, DSL (Digital Subscriber Line) broadband connections came into operation. However, DSL local loops are generally highly asymmetric and may provide a limited upstream data rate; therefore they are less suitable to support video-conferencing and P2P traffic[17]. To provide a higher upstream data rate, the DSL infrastructure would require a significant re-design and re-engineering. As this is unlikely to generate new revenues and provide a viable business case, telecom companies seem not interested in making their networks more symmetric.

In February 2008 the EBU Technical Committee posed the following question:

'What are the limitations of P2P networks for video streaming if users utilize asymmetric lines with limited upstream capacity?'

A collective response from Project Group D/P2P is given below.

17.1 Members experiences

Several EBU members have evaluated the live TV distribution over P2P networks with a high degree of asymmetry.

The VRT carried out a test with the Octoshape P2P system for the Tour De France in 2007, where they streamed at 1.5 Mbit/s. ADSL lines in Belgium are mostly highly asymmetric, using upstream

[17] Services such as Slingbox could also be enabled by making DSL more symmetric. Slingbox relays a home television receiver viewing via the Internet to any other location.

capacity of 256 or 512 kbit/s. Some larger network providers have symmetric lines but some of them block P2P traffic.

The VRT's experience is that, because of the limited upstream capacity, the P2P grid did not contribute much to the overall bitrate. Most capacity needed therefore came from the central servers. These servers can be spread over the different ISPs (as in a simple CDN).

The WDR's experience from the Rock-am-Ring festival is that the P2P/Unicast ratio of live streaming using Octoshape's P2P system was 40% - 60% at a bitrate of 700 kbit/s.

Hence the VRT believe that P2P is today most useful for radio/audio streaming and low quality video.

For high quality video the gain for the broadcaster compared to Unicast is very limited.

17.2 Technical discussion on Asymmetry

The *Network Asymmetry* (NA) is the ratio of *upstream capacity* (UC) to *downstream capacity* (DC).

i.e. NA = UC / DC

The following is a real-life example with nominal numbers:

'Swisscom': NA = 0.6 Mbit/s / 5 Mbit/s = 0.12

and

'ADSL 16': NA = 1 Mbit/s / 16 Mbit/s = 0.065

NA tends to increase with more complex technologies such as ADSL 2 and ADSL 2+ in some markets. Quite in contrast the VDSL-Flatrate offerings of Deutsche Telekom comprise

'VDSL 25': NA = UC/DC = 3 Mbit/s / 25 Mbit/s = 0.12

and

'VDSL 50': NA = UC/DC = 5 Mbit/s / 50 Mbit/s = 0.10

Generally, the bigger the ratio of the access, the better it can be used by P2P-applications. In the examples above some networks which provide potentially a better P2P performance (due to the higher upload-capacity) have a lower ratio.

Using nominal numbers of the connection capacity makes it difficult to prove the individual P2P usability by this ratio. More significant information can be achieved by using values of the bandwidth used by the *received material* (SB) and the *upstream capacity* (UC) (nominal or better which is available for the applications upstream).

The efficiency of the P2P grid depends on the ratio (R) between the (available) upstream capacity (UC) and the media stream bandwidth (SB).

i.e. R = UC / SB

This results, for the examples mentioned above, (3.5 Mbit/s streaming material and about 80% of available upstream bandwidth):

'VDSL 50': R = UC/SB = 4 Mbit/s / 3.5 Mbit/s = 1.14 (in contrast to the NA = 0.10)

and

'ADSL 16': R = NA = 0.8 Mbit/s / 3.5 Mbit/s = 0.23 (in contrast to the NA = 0.065)

Three use cases apply:

CASE 1: R >= 1

In this case the media streams are of lower bitrate than the usable upload bandwidth, typically audio only or low-definition video. The efficiency of P2P grid is high and no additional caching or CDN is required.

Typical cases: Assuming UC = 600 kbit/s

- SB = 100 kbit/s audio stream (e.g. AAC stereo) (-> R = 6)
- SB = 200 kbit/s video stream (e.g. for QVGA screens - mobile phones) (-> R = 3)
- SB = 450 kbit/s video stream (e.g. for VGA screens - portable devices) (-> R = 1.33)

CASE 2: 1 > R >= 0.75

In this case, the media streams are larger than the upstream capacity of the network, albeit of the same order of magnitude. The P2P network is no longer able to perform all the distribution only by itself - additional help from the server-based technologies is required (either in the form of 'super-peers', caching servers or small-scale CDNs)[18].

Typical case: Assuming UC = 600 kbit/s

- SB = 700 kbit/s video stream (-> R = 0.86)

CASE 3: R < 0.75

In this case, a standalone (pure) P2P grid is no longer able to perform the distribution of the media streams. The distribution predominantly relies on a server-based distribution such as CDN. A P2P grid, superimposed to CDN, can complement CDN to improve its scalability (but cannot replace it).

Typical cases: Assuming UC = 600 kbit/s

- SB = > 1 Mbit/s video stream ('standard Internet TV') (-> R = 0.6)
- SB = > 2 Mbit/s video stream (sometimes called 'HDTV Internet TV') (-> R = 0.3)

In addition it should be mentioned that available upstream capacity can be partially reduced by firewalls, parallel usage and data volume-based pricing of ISP.

Comments

Media Quality is steadily being improved. There are three reasons for that:

- Networks: upstream capacity of DSL lines is improving, as DSL line capacities are generally going up (both in terms of downstream and upstream, although in different proportions). In addition, links become more and more symmetric (consumer SDSL, HDSL, VDSL2+) in some markets.

[18] It should be pointed out that there are only a very limited number of P2P solutions that are technically capable of using upload capacity in CASE 2.

- Audio and video compression technologies: have recently made significant improvements and are more and more efficient, both in encoding and decoding. For example, the state-of-the-art Advanced Audio Coding (AAC) is about 3 times more efficient than MPEG-2 Layer 2 coding developed about 10 years ago (AAC requires 3 times less bit rate to achieve the same quality).
- The third reason is the progress in P2P technologies (such as Octoshape) which are able to benefit from many new approaches (for example, exploit even small upload contributions from 'stand-by' users).

The above technology advances may allow the standalone P2P networks in the next five to ten years to become a viable distribution mechanism even for 'standard TV quality'.

Service continuity: As the number of peers in the network varies over the time (and may eventually increase to several tens of thousands), it is important to keep the service running continuously. P2P has a unique capability (compared to CDNs) of avoiding a 'single point of failure' (which may occur in case peers come and go). Also, P2P is capable of scaling down quality (lowering media stream bitrate) smoothly, if the number of peers suddenly increases (thus increasing the need for aggregate upload capacity).

Scalability: The P2P systems are highly scalable both in terms of media quality and numbers of simultaneous users. This is particularly important for covering live events with unpredictable size of audience.

Cost aspects: For CASE 1 systems, commercial P2P systems are currently significantly cheaper than CDNs (approximately by an order of magnitude), measured as cost per gigabyte, although this cost gap is getting less.

Progress in P2P systems: use of standby users: Octoshape is able to benefit from using upstream capacity contributions from standby (idle) peers who are not using the upload capacity by themselves. Such contributions are generally very small, however since the number of such peers may be large, such contribution may become significant to help with the upload deficit to the peering users.

17.3 Some conclusions on asymmetry

The upstream capacity limitations of today's access networks (including DSL, HSPA, WiMAX, etv) severely affect the functioning of P2P systems. Nevertheless, it is really the ratio between the upstream network capacity and the media stream bitrate that matters, as identified in the three use-cases above. In CASE 1, where media streams are of lower bandwidth, P2P generally should work fine as a standalone system, and should be scalable and cheap. In CASE 2, P2P requires a limited support from the server-based technologies. CASE 3 requires a significant contribution from the server-based technologies.

In other words, currently no true HDTV quality is possible with the pure (no caching) P2P distribution systems in regions with highly asymmetric networks (such as existing ADSL networks in Europe). However, due to advances of the DSL networks, audio and video compression technologies as well as P2P system, P2P is gradually expanding into a better quality media delivery zone.

For the future it is likely that a combination of CDN and P2P will prevail, whereby the share of CDNs will gradually diminish, as the media technologies and networks advance.

Some benefits of SVC and MDC redundant coding methods are being considered to gain maximal value from the use of asymmetric connections.

18. Trends in P2P Research and Development

18.1 P2P-Next

The Next Generation Peer-to-Peer Content Delivery Platform (*P2P-Next* for short) http://www.*P2P-Next*.org/ is an EU FP7 Collaborative Research (Large Scale Integrated Project) which started in January 2008 and is due to end in 2011. The objective of *P2P-Next* is to build an open peer-to-peer-based content delivery platform, supporting broadband video portals and delivering content via P2P to TV screens via STBs, and PCs. The *P2P-Next* system is designed to include payment and DRM functions, as well as functions to help content producers turn linear video content into interactive content. The goals are also to define use cases and sustainable business models for the *P2P-Next* system and the various actors and users.

At the time of writing, the *P2P-Next* project team had reported a number of technical achievements for the first 24 months. Among those, the highlights are the streaming video-to-STB solution (exhibited at IBC 2008), the Internet-HD quality end-to-end streaming distribution of professional content to low-cost STB hardware (exhibited at IBC 2009). A number of trials have been conducted in the Living Lab and a large number of users (unique IP addresses) have successfully installed and executed the software.

The main strength of the project is in the large scale of integration it pursues. In addition to the P2P core functionality the project promotes a practical approach to content nonlinearity and interactivity (the LIMO concept), as well as several other fundamental elements of the media delivery and consumption chain.

Following the audit in March 2010, the Commission stated that most of deliverables have been delivered in time, except for the mass production of the NextShareTV CE device. This milestone has been delayed because of technical problems which are beyond the control of the *P2P-Next* team.

The *P2P-Next* platform approach allows modular development and modular applications, enables knowledge sharing and facilitates technology integration, code- and skill re-use. This translates to fast development of new content delivery applications that build value for service and content providers.

P2P-Next will advance the state-of-the-art in important areas, including evolutionary content distribution, easy access to vast amount of content with metadata federation, social networking, and innovative business models for advertising. The sum of these advances is a large step towards moving the information access from the hands of a producer to the hands of the consumer, and allowing consumers to enjoy and utilize content resources in a mobile and pervasive manner, across the great online space.

Distribution of radio and television programmes, movies, music, ring tones, games, and various data applications to the general public is today possible via a variety of dedicated networks and special end user terminals. As broadband Internet becomes ubiquitous, all content distribution services will be combined (bundled) and conveyed to the general public via a common pipeline – the Internet. Today several technologies are used for the media distribution across the Internet: Unicast, IP Multicast, content distribution networks and most recently – Peer-to-Peer (P2P).

P2P-Next will be developed through collaboration with European and national initiatives, as well as some of the largest and most important actors in the media and telecommunications sector, ensuring industrial relevance and worldwide application reach. *P2P-Next* involves 21 partners in 12 different countries, including large European players to ensure the future project's sustainability, SMEs, universities and R&D institutes to manage highly-focused technology components[19].

Broadcasters and content providers consider *P2P-Next* as a promising approach towards a future-proof, universal, and ubiquitous two-way (interactive) distribution mechanism. *P2P-Next* seems to be well positioned to complement the existing distribution mechanisms such as satellite, cable and terrestrial networks.

The P2P-Next Project extends the notion of a conventional media distribution network. It introduces a concept of on-demand, personalised and social network.

In the following some details of *P2P-Next* are given.

P2P-Next is an overlay to the existing infrastructure

The *P2P-Next* system is an application-layer media delivery system, which can in principle be overlaid to any two-way communication system. Media delivery does not require a dedicated network such as the DVB-H or 3G networks. Therefore,

P2P-Next requires much lower infrastructure investments (no 'streaming farms' are required), management costs and maintenance costs compared to dedicated distribution networks.

P2P-Next may dramatically improve the network economics

P2P-Next changes the conventional business model for media distribution over the Internet. Using conventional technologies (e.g. Unicast, IP Multicasting, content distribution networks), distribution costs are proportional to the number of users, the bandwidth required and service quality required (SLA).

P2P-Next cost is not directly proportional to the number of users; in fact, the cost per user diminishes with the number of users. Therefore, *P2P-Next* significantly reduces distribution costs.

The number of services that can be accommodated is practically limitless. Note that conventional systems do require frequency spectrum, which is a scarce resource and has a limited capacity of channels.

To this end, P2P-Next may considerably ease regulatory and frequency management problems.

P2P-Next approach is NOT limited to computers

The concept of P2P distribution can be ported to virtually any consumer-electronics terminal devices such as DVB set-top boxes and home gateways. For reasons of interoperability, such P2P solutions need to be standardised and validated. Embedded *P2P-Next* devices could be used for open Internet access as well as closely controlled and managed IPTV systems.

Extension of P2P to connected non-PC devices can dramatically change the business model of media distribution. In this manner, *P2P-Next* could become platform- and device agnostic.

The P2P-Next Project plans to contribute to worldwide standardisation of an open-source, scalable, and modular P2P plug-in, which can potentially be embedded in any CE media device.

P2P-Next is a win-win technology as it helps all actors in the value chain

By rolling out *P2P-Next* services, all actors of the media distribution value chain could benefit

[19] P2P-Next includes several public broadcasting organisations including, RTVSLO, BBC and EBU

significantly:

- Content owners and providers can enlarge their markets and can make more profits.
- For broadcasters, *P2P-Next* can represent yet another outlet for distributing their programmes and contents.
- Consumer manufacturers and IT manufacturers could sell more hybrid[20] STBs and other devices. As the market grows, prices go down.
- Network providers will benefit as the overall network load will be reduced – P2P packet paths are much shorter than conventional in traditional server-client networks.

P2P-Next aims to support all actors in the media distribution value chain.

P2P-Next provides live steaming, downloading, and progressive downloading

It has been often stated that the audio-visual media are moving away from linear channels and are becoming more and more on-demand, thus available to the end user when they want them, where they want them, and on any terminal.

The P2P-Next system will help end users to live-stream a TV channel, download a song or video clip, or progressively download a file (and watch it while it is being downloaded).

P2P-Next enables horizontal market solutions

As *P2P-Next* is a non-proprietary, open source, open standard solution and will be put forward for international standardisation. It will enable the deployment of a horizontal market. A not-for-profit foundation will be established to roll out *P2P-Next* technology and perform the required product compliance and interoperability tests.

P2P-Next will facilitate the opening of horizontal markets for any kind of business.

P2P-Next facilitates social networking

One of the underlying principles of any P2P system is the participation of several peers in a common activity (for example, sharing a video clip). *P2P-Next* will facilitate the introduction of several social features in content distribution: communication primitives, strong peer authentication, strong content integrity checks, permanent storage of context, semantic clustering, recommendations, reputations, micropayments, etc.

P2P-Next will enhance P2P distribution with social networking features to support user and business communities.

P2P-Next development is driven by actual users

There are many examples of promising technical solutions that dramatically failed, because users' perception did not match with marketers' vision. Much of the 2001 Internet bubble collapse was caused by this phenomenon.

In P2P-Next an incremental approach will be followed by early and short-cyclic releases involving actual user communities (living lab approach).

[20] 'Hybrid' implies both broadcast and broadband front ends.

P2P-Next Content generation and distribution

The *P2P-Next* project designs the mechanisms for professional content distribution, which is particularly interesting for public broadcasters. The project has adopted metadata standards based on MPEG-21, MPEG-7 and TV Anytime, designed a lightweight interactivity mechanism (LIMO) based on the emerging HTML 5 standard, and defined how content will be packaged and delivered within *P2P-Next*. The adoption of the Kendra Signposts service for the metadata transformation has the potential to attract a larger base of professional content providers. Temporal addressing of content segments and its recommendation to friends is highly valued by the users and hence should be provided on the STB.

A basic toolbox was developed that covers the control, management and ingestion of VoD content. The *P2P-Next* Rich Metadata specification has been extended with the specification of optional metadata, i.e., advertisement, payment and scalability metadata. ATOM Feed based content discovery has been implemented on the STB and was demonstrated at IBC and NEM Summit in 2009.

P2P-Next PC and CE device platform

The PC platform is designed and prototyped and is addressing several advanced features such as SVC and MDC coding, zoomable UI, adaptive media playout and forward error correction techniques, etc. The current SVC decoder implementation is not real-time, which might turn out as a risk to the end-to-end demonstration of the benefits of content scalability.

Developments of the low-cost CE platform, supporting the NextShare applications, have been pursued. An EPG and an ATOM-based Feed Navigator, allowing content discovery by end-users have already been implemented. Initially, the P2P network stack was based on Python but exhibited performance limitations (in terms of computational complexity and resource requirements) when running on the P2PNextshare CE hardware platform. Effort has then been dedicated to the development of an optimized version of the NextShare content transport to allow for higher bit rates and lower memory resource utilization.

The mass production of the units to be used in the Living lab trials has been slightly delayed due to technical problems in the manufacturing of the SoC (system on a chip), which are beyond the control of the P2PNext consortium. The support for SATA storage has been lost and the consortium is compensating for it by relying on USB flash media. It has been agreed that 300 units will be manufactured, instead of the original 500 budgeted for in the Living lab trials.

SWIFT protocol

One of the focus points of the *P2P-Next* project[21] is to develop a generic content-centric multiparty transport Internet protocol called SWIFT. This protocol is designed to distribute content very efficiently among a swarm of peers using UDP with LEDBAT (Low Extra Delay Background Transport) congestion control. Once developed and tested, SWIFT will be capable of integrating smoothly into current and future browsers and operating systems accommodated in practically all consumer-electronic devices including set-top boxes, televisions, mobile and portable devices, so that it will be able to serve most of the Internet traffic.

[21] As of March 2010

18.2 SmoothIT

The project has the following major objectives to be pursued in support of the overall aim of radically advancing technology:

- SmoothIT will structure overlays in a way that is efficient or optimal, both for user communities and for ISPs. This is to be attained by means of economic and incentive mechanisms
- SmoothIT will study and define key requirements for a commercial application of economic traffic management schemes for ISPs and Telcos.
- In order to advance traffic management beyond traditional limits, specialized economic theory will be applied for building, in a fully decentralized way, network efficient Internet-based overlay services in multi-domain scenarios, solving the information asymmetry problem.
- SmoothIT will design, prototype, and validate the necessary networking infrastructure and their components for an efficient implementation of such economic traffic management mechanisms in an IP test-bed and trial network.
- SmoothIT will develop an optimized incentive-driven signalling approach for defining (theory) and delivering (technology) economic signals across domain boundaries in support of co-operating and competing providers in an interconnected heterogeneous network environment.
- SmoothIT will integrate concepts from previous work, such as M3I and MMAPPS, leveraging existing knowledge and applying it to future overlay services.
- SmoothIT will stress operator-orientation by verifying key results of the work through ISP and Telco requirements as well as its supporting technology.

18.3 NAPA-WINE

TV services over the Internet can be provided either by exploiting IP Multicast functionalities or relying on a pure end-to-end (P2P) approach. The first technique unfortunately will only work on a network infrastructure controlled by a single broadband operator due to limitations of IP Multicast facilities. The P2P approach, on the other hand, has been successfully exploited to overcome these limits and can potentially offer a scalable infrastructure. Recently, several P2P-TV systems started up, with the last generation offering High Quality TV (P2P-HQTV) systems, providing a ubiquitous access to the service. These same potentialities of P2P-TV systems constitute a worry for network carriers since the traffic they generate may potentially grow without control, causing a degradation of quality of service perceived by Internet users or even the network collapse (and the consequent failure of the P2P-HQTV service itself!).

Starting from these considerations the NAPA-WINE project, funded by the European Commission within the seventh framework programme, aims at:

- providing a **careful analysis** of the impact that a large deployment of both general P2P-TV and P2P-HQTV services may have on the Internet, through an in detailed characterization of the traffic they generate;
- providing **guidelines for P2P-TV developers** regarding the design of systems that minimize the impact on the underlying transport network while optimizing the user perceived quality;
- providing a **road map for Internet Service Providers** to better exploit the network bandwidth by showing simple and minimum cost actions that can be taken in presence of P2P-TV traffic.

18.4 NaDa

NaDa (Nanodatacentres), http://www.nanodatacenters.eu/, is an EC funded project which investigates into the next step in data hosting and in the content distribution paradigm. By enabling a distributed hosting edge infrastructure, NaDa can enable the next generation of interactive services and applications to flourish, complementing existing data centres and reaching a massive number of users in a much more efficient manner.

Increased computational power, combined with advances in data storage and global networking, has made Internet services a critical resource in our everyday life. Data centres (buildings that host large numbers of networked computer servers and power supplies) are often critical enablers of such services. Data centres are known to be a major source of cost and complexity for operators, while they are inherently not scalable due to their centralised nature. As a result, router companies, server manufacturers, and hosting facilities hasten to produce more efficient hardware and software for data centres. They also try to improve the efficiency of operation of such components. For instance, operators may dynamically shut down some processes in machines or even entire machines, depending upon the current load. They may also redirect surplus load to other idle machines in the same data centre. While this effort improves efficiency, it is bound to produce rather short-term remedies. Indeed, **the entire paradigm of monolithic data centres seems to be challenged, not the specifics of their numerous possible realizations**.

The changes in costs, combined with the observed changes in highly interactive demand profiles, illustrate the need for a paradigmatic shift towards highly distributed data centres. NaDa is tailored towards servicing interactive applications to a massive number of clients. This solution requires a large number of geographically dispersed nano data centres (instead of a few large data centres). In addition, it will materialise from the composition of pre-existing, but underutilised resources, and thus does not require heavy capital expenditures. Indeed, there are large amounts of untapped resources at the edges of the network today that, if integrated intelligently, could provide a substantial complement to existing data centres, if not a complete substitute. Such resources include: next generation home gateways, set-top boxes, wireless access points, etc. Most of these devices are nearly as powerful as standard PCs, with a great deal of processing power and reasonable storage, but unlike PCs, they are often idle and moreover controlled by a single service provider. This idleness is largely due to 'always-on' user habits, which result in systems that are being powered most of the time, despite most of their computing and storage resources remaining inactive. Similarly, the (broadband) link that connects such boxes to the Internet stays idle for long periods of time. The NaDa objective is to tap into these underutilised resources at the Internet edge and use them as a substitute/aid to expensive monolithic data centres.

The NaDa approach is classic in one respect, and revolutionary in others. It moves content and complexity to the edge, which is perfectly in line with the Internet's original philosophy and offers the maximum guarantee of network performance and availability (no additional complexity in the network). However, it is a revolutionary approach in next generation Internet research, especially when compared to the approach currently taken in the US which is to re-design the architecture of the network core in order to better handle content instead of using existing resources at the edge. Still, NaDa does not ignore the current CDN or cache-based content delivery architecture. In fact, NaDa will use existing caches or CDNs to improve the quality of service experienced by users.

In order to combine all unused edge resources, NaDa will use a new, managed peer-to-peer (P2P) communication architecture. The P2P paradigm allows the deployment of new services such as file sharing or telephony quite easily without having to scale servers for peak capacity. However, most of the currently deployed P2P systems have focused on simple file sharing or streaming applications (and often for illegal content). Thus, several fundamental issues must be addressed in order to invent a new P2P paradigm for the NaDa system.

18.5 VICTORY

The IST VICTORY project (Audio-VIsual ConTent search retrieval in a distributed p2p repositORY), www.victory-eu.org, which started in January 2007, aims to build a distributed repository of heterogeneous visual objects, accompanied with related textual documents, videos and images (MultiPedia objects). In light of three dimensional (3D) object retrieval evolving from text annotation to content- and context-based and from standalone applications to web-based search engines, VICTORY aims at the creation of **a search engine for 3D and associated multimedia distributed content** into Peer-to-Peer (P2P) and mobile P2P (m-P2P) networks.

Driven by the very successful concept of Wikipedia (www.wikipedia.org), the main goal of the project VICTORY is to create the first distributed **MultiPedia** Object Repository. A Multipedia Object is a 3D object along with its accompanied information i.e. 2D views, text, audio, video to which any peer can contribute.

The main objective of VICTORY is the development of novel search and retrieval framework that allows easy integration of different search methodologies. This will result in an integrated platform which will allow processing and accessing data and knowledge by using ontology-based management and semantic-based retrieval mechanisms.

The challenge within VICTORY is to bridge the gap between content-based and knowledge-based search and to apply this really innovative technology to **MultiPedia** content, especially to 3D objects.

Content-based search will be based on a) content, which will be extracted taking into account low-level geometric characteristics and b) context, which will be high-level features (semantic concepts) mapped to low-level features. VICTORY aims at introducing a solution so as to bridge the gap between low and high-level information through automated knowledge discovery and extraction mechanisms. High level features will be a) appropriate annotation options provided by the system or generated by the user dynamically (active learning) and b) relevance feedback where the user will mark which retrieved objects he thinks are relevant to the query (user's subjectivity). These high level features are expected to improve significantly the retrieved results.

For supporting sophisticated 3D object search and retrieval, an ontological framework is needed that will allow for combining text/metadata-based searching with 3D content searching. This ontology will be used either as organizing principle to navigate through the 3D objects, or for restricting/guiding the search through the objects' classes. As a result, VICTORY will introduce advanced search methodologies, in addition to the filename-based method prevalent in P2P applications.

The search engine will feature novel multimodal personalized interfaces in order to take into account the users' interests and to offer capabilities of matching between 2D/3D objects (image to 3D), sketches and text (mixed-media queries). The search engine will be able to retrieve the appropriate 3D object(s) along with its accompanied **MultiPedia** information.

VICTORY will also develop a P2P scheme so as to utilize not only the distributed data storage but also the computational power of each peer for the pre-processing, interpreting, indexing, searching, retrieving and representing of **MultiPedia** data. Through the VICTORY framework, users will be able to handle, share and retrieve 3D and audio-visual data among peers around the world. Moreover, every peer will be responsible for extracting and indexing the features of the shared 3D data, thus the efficient manipulation of the 3D data will be accomplished.

The key driver for the design of the VICTORY P2P scheme is the enhanced Quality of Experience (QoE) of the user. QoE can be realized as the combination of a multitude of Quality of Services (ranging from bitrate, processing speed, and graphics quality, to power consumption). Existing P2P-based standards and products are oriented towards content and services discovery while QoS (and

QoE) issues for real-time demanding applications (such as 3D content manipulation and rendering) have not been adequately addressed. The envisaged scheme will incorporate the protocols, semantics, and intelligence allowing the negotiation of resources (computational, communications, data storage, etc) sharing and the optimization of the P2P framework for enhancing the user quality of experience. Moreover, the P2P collaborative framework will be capable of utilizing new techniques to perform visualization and rendering of the retrieved 3D objects into low bandwidth and processing power mobile devices. Finally, the search and retrieval efficiency of the system will be improved by restricting the search procedures to the neighbouring peers. Copyright protection and 3D objects' ownership will be provided developing/deploying highly innovative 3D watermarking algorithms and DRM.

Technical approach

The work in VICTORY is divided into four phases covering the following topics: Project management and evaluation, research and development, prototyping and integration, and user testing and validation. Research and development in VICTORY will mainly focus on the development of the following components:

- **P2Pnetwork**: involves research in P2P networking of standard PCs and mobile devices and interoperability between standard and mobile P2P networks. The main objective of this discipline is to develop a QoE (Quality of Experience)-aware P2P-based framework, which will provide the advanced middleware infrastructure needed for the effective realization of the VICTORY applications for 3D content searching and manipulation. This middleware framework will allow standard and mobile peers to advertise, access, and negotiate the sharing of available computational, communications, services, and content resources in a user centric, flexible and secure manner.
- **3D Search engine**: involves research into 3D content and context based search and retrieval techniques, and their integration to a distributed P2P environment. The search engine will consist of the following distinct parts :
  - 3D low-level feature extraction algorithms
  - Annotation mechanisms
  - Relevance Feedback algorithms
  - Ontology Based Retrieval
  - 2D/3D and sketch/3D matching algorithms
- **Multimodal personalized user interface**: involves the creation of a user interface where the user will be able to input information about himself, his interests and so on. Additionally, the retrieved data will be visualized through the user interface where the user will be able to interact with them (roto-translate the objects, change the rendering mode (wire frame, solid, etc.), add/remove textures, add annotations, mark relevant results and so on). The use of semantic concepts, interfacing with end-users' interests, along with the different modalities and the interactive capabilities are expected to provide great flexibility, high efficiency of the retrieval process and ease of use.
- **Visualization on handheld devices**: involves research so as to bridge the gap between high quality 3D graphics and mobile devices. VICTORY will allow mobile users to search, retrieve and finally display, on their own devices, extremely complexity scenes. The infrastructure will be 'flexible' enough in order to support the widest spectrum of devices independently of their hardware and operating system.
- **Security and Copyright protection mechanisms**: involves research into security and copyright protection techniques to guarantee secure delivery of 3D content across standard and mobile P2P networks. This will be achieved firstly, using a DRM architecture that will allow a super-distribution model of content distribution and secondly, protecting the

MultiPedia content using a combined model consisting of two methods, 3D digital watermarking and public-key encryption.

18.6 P4P

P4P [18] is a framework that can be used to enable Internet service providers (ISPs) and peer-to-peer (P2P) software distributors to work jointly and cooperatively. P4P stands for 'Proactive network Provider Participation for P2P'. P4P is a *set of business practices and integrated network topology-awareness models designed to optimize ISP network resources and enable P2P based content payload acceleration*.

The P4P framework can be used to ascertain appropriate and voluntary best practices to accelerate distribution of content and optimize utilization of ISP network resources in order to provide the best possible performance to end-user customers.

P4P allows for more effective cooperative traffic control between applications and network providers. The Yale University conducted extensive simulations and real-life experiments on the Internet to demonstrate the feasibility and effectiveness of P4P. Their experiments demonstrated that P4P either improves or maintains the same level of application performance of native P2P applications, while, at the same time, it substantially reduces network provider cost compared with either native or latency-based localized P2P applications.

18.7 Digital Video Broadcasting (DVB) Project

The Digital Video Broadcasting (DVB) Project is an industry-led consortium of around 300 broadcasters, manufacturers, network operators, software developers, regulatory bodies and others in over 35 countries that are committed to designing open technical standards for the global delivery of digital television and data services. To date, DVB has produced and published a number of DVB IP television specifications, focused on 'managed IPTV track'. During 2009, the DVB IPTV group developed a document entitled 'Commercial Requirements for Peer-to-Peer Internet-TV Content Delivery' which represented a first brick of the unmanaged 'Internet-TV track'[22].

18.7.1 DVB Commercial Requirements on P2P

The DVB Commercial Requirements on P2P document covers the deployment of DVB services over unmanaged networks, such as the Internet, where at least one segment does not provide guaranteed QoS and where advanced network features (e.g. Multicast) cannot be fully assured. This means that no assumption can be made on availability or reservation of resources in the core and access network. The system shall assume that only best effort data delivery can be used, i.e. functionalities that the typical broadband Internet access provides today. Managed content delivery networks are not considered here, although P2P could also be used in the managed, QoS-controlled environments.

The DVB describes the complete set of commercial requirements from the Internet delivery side. It represents a basis for further development work which may result in a complete technical specification for carrying DVB content across P2P networks. Such a specification may be produced in due course (although currently there is no plan to do so).

Extracts from the DVB P2P Commercial Requirements document are reproduced in **Appendix 2** of this document.

[22] In producing this document, the EBU representatives played a significant role.

18.7.2 DVB Study Mission on Internet TV Content Delivery

In March 2009, DVB decided to launch a major technical Study Mission on Internet TV Content Delivery. The rationale for this study mission was mainly to investigate technology options to deliver DVB-type content over the Internet to a large number of CE devices (including game consoles), PCs or mobile devices. The primary purpose of the study was to identify if DVB should start any specification activities in the area of Internet-TV content delivery, complementary to the previous DVB work on managed and QoS-guaranteed IPTV delivery. Another key question was whether or not DVB members consider P2P as an appropriate technology for media delivery over the Internet, so that they could start a standardisation process.

In the light of a large number of responses received, the Study Mission showed that many DVB members, including broadcasters, CE manufacturers, technology integrators, network providers and others, are already actively involved in the open Internet-related commercial activities and are therefore highly interested in furthering the specification activities.

Based on the extensive information collected by the Study Mission report [26], many important conclusions were taken, however for the purpose of this report we extracted the following items:

- DVB should continue considering specifying of certain important Internet-TV content delivery interfaces, formats and protocols. Most importantly, there is considerable scope for improving technologies for the reliable distribution of *high-quality commercial AV content* over the Internet to a *large number of consumer end devices*, which requires considerations that are not already sufficiently addressed by any standardization organization. Proprietary solutions exist but are generally not targeted or adapted to typical DVB services, content and end devices.
- DVB should focus on components that have clear and well-specified interfaces, and can be integrated and deployed in different end-to-end specifications and deployments and in combinations with different technology components already specified by DVB or elsewhere, such as service discovery, middleware, content protection, etc.
- DVB should reuse technologies within the existing DVB specifications in areas of, for example, DVB-IPTV, DVB-AVC (Audio Video Coding), DVB-FF (File format) and DVB-GBS (Generic Data Broadcasting and Service Information Protocols), but only to the extent that those technologies are well adapted and may potentially improve performance or/and add functionalities, if deployed.
- DVB should adopt an evolutionary concept of refining Internet-TV specifications in different phases. Specifications of Internet-TV content delivery can be refined and extended more easily than traditional broadcast specifications.
- Internet-TV Content Delivery in mixed Broadcast/Internet deployments where the main service is still distributed over DVB-S/T/C or IPTV. Specifications for the delivery via the Internet of supplementary services such as content download, content on-demand or auxiliary services should be considered initially.
- No technical indication has been given during the Study Mission why the MPEG-2 TS and the DVB AVC codecs could not be used in the Internet-TV content delivery. However, due to the lack of QoS support, typically resulting in varying bitrates, adaptive content delivery for streaming and download should be considered, along with suitable application and service signalling.
- In Internet-TV content delivery HTTP is considered as the primary protocol from the network to the consumer end device. DVB should investigate if such an approach can fulfil the commercial and technical requirements for the use cases and services as initially considered by DVB.
- The Study Mission has found evidence that a classical client-server approach can efficiently be augmented by adding caching servers - Content Distribution networks (CDNs), distributed

across different Internet domains and located at the coverage edges. DVB should consider this option in its initial release and attempt to standardise the interfaces and protocols between the CDNs and the end user CE devices connected to the Internet.

- Extensibility of the initial versions must obviously be taken into account from day one and for this purpose relevant use cases should be considered and checked for possible integration in later releases. *Optional extensions should be considered, in particular augmenting the specification with P2P-based delivery and/or combined CDN/P2P solutions*.
- DVB should focus on technical specification for the interfaces to the consumer end device, as it has done successfully for all other broadcast systems. DVB should not duplicate its efforts with other organisations such as OIPF, HbbTV or MHEG-5 IC, but should attempt to provide sufficiently clear interfaces such that the existing specifications can be integrated in emerging solutions. For example, an interface solution for a DVB Internet-TV technology may be part of a Interactive TV offering just as existing DVB-T or DVB-S is.
- Due to lack of time, the Study Mission report could not provide in-depth considerations of some recent Internet-TV content delivery developments, such as those specified by OIPF, DVB CDS Unicast, Apple and Microsoft IIS-Smooth Streaming. Nevertheless, these specifications could provide a sound starting point for preparing DVB technical specifications for Internet-TV content delivery.
- *As shown by the numerous inputs to the Study Mission from NextShare, Samsung P2P-TV and emundoo as well as the work in the IETF on PPSP, the DVB Internet-TV specification could be usefully enhanced by adopting an optional standardised P2P interface*.

For our purposes, the statements referring to P2P delivery are highlighted in italics. The message given is that the DVB Consortium does not exclude considering the P2P delivery, either independent or in combination with CDN, and may develop the appropriate P2P protocol and interface standards, provided that the market requires them, as optional additions to the DVB Internet-TV specifications.

18.8 Use of SVC/MDC (Scalable Video Coding/Multiple Description Coding) in P2P Systems

The advantages in terms of robustness, reconfigurability and scalability make P2P a promising technology. The P2P approach permits to serve a larger audience without the need of proportionally increased resources; in fact, only a small fraction of the audience should be served directly by a content provider. Users may then connect to each other and exchange their content. From the user point of view, the P2P should allow to experience high quality video in a cost effective fashion.

A P2P client (node) usually receives data from many peers simultaneously, only partially relying on centralized servers. All peers provide the community with resources in terms of upload bandwidth and storage capability. The distributed nature of P2P overlay networks increases robustness in case of failures, as data are usually replicated over multiple peers. On the other hand, the use of P2P for video distribution is not devoid of problems. Real time streaming applications require the overlay network to guarantee a constant flow of data and low start-up latency, as well as on-the-fly content adaptation and inherited resilience. The P2P video technology still experiences problems of long **start-time** and **churn-induced** instability. The aggregated peer uplink **bandwidth** is still typically insufficient to support large scale distribution, due to the asymmetry of residential broadband connections. **Packets losses** happen for several reasons: router congestion and transmission errors on the physical network, node departure from the P2P overlay, strict timing out due to real time visualization (which makes data packets lose relevance very quickly). The lack of guarantee about the actual delivery of the data may cause unacceptable drops in the reproduction quality, and frequent service outages. For example, due to predictive encoding, a packet loss not only corrupts the current frame in a video transmission, but the error also propagates to subsequent frames. Another relevant aspect that is affecting research in multimedia is **unreliability and heterogeneity**. Internet is a best effort network, with no guarantee about the actual delivery of the transmitted data. Typically data get lost because of router congestion or transmission errors (especially over wireless channels). In a P2P system, the vulnerability is also due to node departure from the overlay network. In applications such as file download or web browsing, specific transport mechanisms request the retransmission of the lost data; however, these cannot be used in a streaming context, due to time out and network flooding problems. Needless to say, data losses often cause an unacceptable drop in the reproduction quality perceived by the user, and frequent service outage. Moreover, we expect for some years to come a communication scenario where a user can access information in a seamless way, using terminals with very different characteristics in terms of resolution and bandwidth, and different access technologies, ranging from DSL, WiMax, GPRS/UMTS and so on. The demand for mobility and ubiquitous access poses stringent challenges, because of the use of low power terminals and unreliable access networks.

Even though these problems are intrinsic in the P2P paradigm, some countermeasures can be adopted at the video coding level. To this end, **scalable video coding (SVC)** can be an interesting tool. In fact, different levels (layers) can be recovered from a scalable video stream, provided that some basic information (the so-called base layer) is received. This helps avoiding abrupt service breakdown, and also facilitates the exchange of contents among users equipped with different display resolution. The following SVC modalities are possible:

Temporal scalability refers to the fact that one can decode a video streaming at different frame rates according to decoding capabilities. The scalable stream can offer a base layer coded at low frame rate and one or more enhancement layers with increasing frame rates.

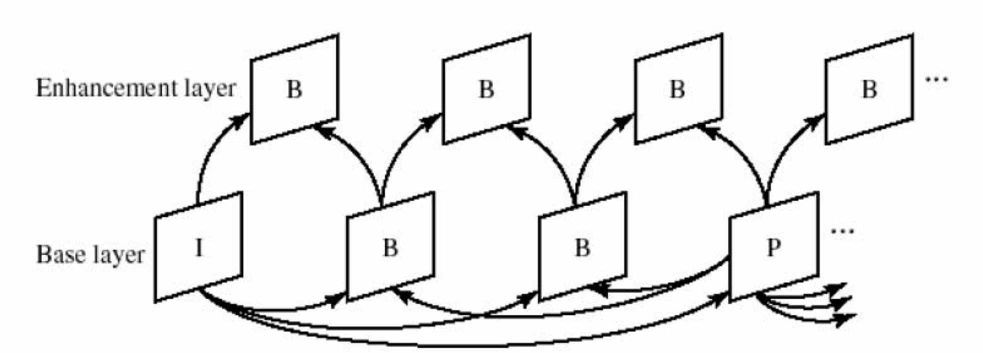

Figure 19: Scalable Video Coding - Base and Enhancement layers

Spatial scalability is obtained when a video is coded at multiple spatial resolutions. The embedded stream can offer a base layer at QVGA resolution with the possibility to scale up to full HD according to the number of decodable layers. The data and decoded samples of lower resolutions are used to predict data of higher resolutions in order to reduce the bit rate to code the higher resolutions.

SNR/Quality/Fidelity scalability can also be exploited in order to increase the visual quality of a video streaming. A coarse-quality version of a multimedia stream could represent the base layer for preview and free-of-charge distribution with the capability to shift to high-quality resolution for editing, post-processing and pay-per-quality services.

The different video layers can be transmitted in different bit streams called sub-streams or they can also be transmitted in the same bit stream, which is called an embedded bit stream.

Figure 20: Spatial and SNR scalability

A major boost to scalability is given by the MPEG/ITU-T JVT H.264/SVC co-decoder [19]. It is conceived as an extension of H.264/MPEG-4 Part 10 AVC [20], [21] and it implements temporal, resolution, and quality scalability and any combination of them, with limited compression penalty with respect to AVC.

As for P2P, SVC is a very promising technique given that it is possible to guarantee the correct and timely delivery of the base layer, through centralized servers or a reliable network infrastructure. The peers would then be able to improve the visual quality by decoding enhancement layers received from other peers. SVC can be used to serve clients with different requirements (e.g., low quality and low bandwidth, high quality and high bandwidth), e.g. through application layer Multicast, with every layer being served to a different Multicast group, and users joining as many groups as desired. SVC requires that the base layer be delivered with very high reliability, and this calls for suitable error control or scheduling techniques.

At the moment, P2P clients do not have the capabilities to store scalable content nor to exchange different layers of a given content. Moreover, another open issue is how to define a scalable stream, in terms of the number and type of layers, and the rate devoted to each of them. H.264/SVC, for instance, only standardizes the bit stream syntax in order to allow a decoder to successfully decode a scalable stream. It is thus possible to develop non-normative tools at the encoder-side to generate such a bit stream.

The number and type of layers, as previously explained, should be carefully tuned in order to avoid inefficiency of the system together with a proper and jointly rate-distortion optimization. Given two resolution layers (i.e. QCIF and CIF), it is still an open question how to distribute rate among these layers. The rate-distortion optimization may become even more sophisticated in case of many layers.

Multiple Description coding is another coding tool that is gaining popularity in the video coding community. In the MDC approach two or more representations of the same data, or descriptions, are generated, which yield mutually refinable information and can be independently decoded. Descriptions are created so that the quality of the recovered signal is only a function of the number of received descriptions.

Similarly to SVC, MDC allows for graceful degradation, in that a basic quality level is obtained from a single description, and quality enhancement stems from the decoding of further descriptions. Different from SVC, which requires the reception of the base layer, no hierarchy exists among descriptions, so any description is useful to improve quality.

The benefits of MDC features are paid for in terms of some impairment of the compression efficiency, due to the inserted redundancy among the descriptions. This redundancy is generally measured in terms of the extra rate required by the MDC scheme, compared to a single description reference system achieving the same performance.

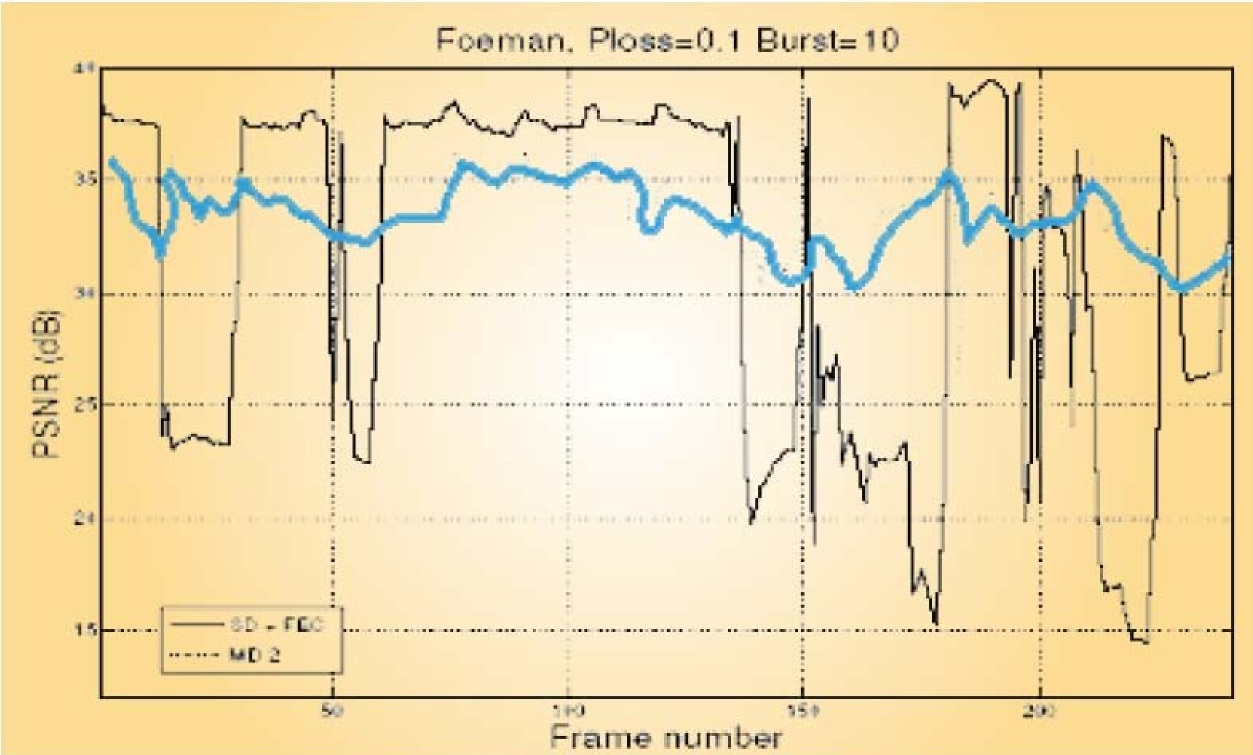

Figure 21: MDC versus standard encoding, quality frame-by-frame with 10% packet loss.

The black line corresponds to standard coding protected by Reed-Solomon forward error correction (all-or-nothing behaviour), the blue line corresponds to two descriptions (slightly lower average quality, but much lower variance).

Descriptions must be somewhat correlated to each other, so that the lost information of the missing descriptions can be estimated from the received descriptions. As the redundancy is detrimental in the case of error free transmission, it is clear that the number of descriptions and the related redundancy should be carefully tuned in order to match the actual network conditions.

Several MDC techniques have been proposed, such as MDC quantization, correlating transforms, lapped orthogonal transforms, correlating filter banks, video sub-sampling and Reed-Solomon forward error correction. The interested reader may refer to [22] for a detailed tutorial on this topic. Recently, MDC schemes have been proposed that are compatible with coding standards such as JPEG2000 or H.264/AVC. H.264 options such as redundant slices and FMO, are exploited, for example, in [23], to create two H.264 streams, which can be either separately decoded, or pre-processed prior to be input to a standard decoder. This latter algorithm which makes use of the redundant slice option present in H.264, generate two descriptions, each of which contains some primary and some redundant slice representations. At the decoder side, the finest representation of each slice is retained. Even though it is subject to the drift effect, the algorithm has proven to yield good performance, with little extra complexity with respect to H.264 co-decoding.

One of the objections to the use of MDC in P2P is that peers receiving a subset of the descriptions generally recover only an estimated version of the missing ones; consequently, peers may be called to share non equal contents. This drawback has a few possible solutions, and is also relevant to SVC. However, it is worth noticing that MD schemes allowing for exact reconstruction of the signal do exist. In MDC using redundant bases (frames), the descriptions, obtained as the output of an oversampled filter bank, exhibit deterministic redundancy and enable exact recover of missing information. As such descriptions are generally visually meaningful, they can be individually compressed using standard tools; this way, the description generation and decoding can be conceived as pre-/post-processing tools at the transmitter/receiver. Moreover, as MDC exploits diversity in the network path from the server to a client, such diversity must be provided. One way to do so is to employ a multiple-tree architecture. Mesh-based overlays should also employ some form of path diversity, e.g., multiple meshes.

Interesting research topics regard the comparison/merging of MDC and SVC features, and the adaptation of MDC tools so as to create scalable descriptions. Another issue is the relationship between MDC, SVC and distributed video streaming, which makes use of the distributed source coding principles for efficient encoding of correlated sources (e.g. different resolution versions of the same video).

19. SWOT Analysis of P2P Systems

The following section gives a bulleted SWOT analysis *from the broadcasters' perspective*:

19.1 Strengths

- very effective for broadcast (one-to-many) mass distribution – compared to Unicast
- cost effective – processing and storage resources contributed by the users
- low infrastructure requirements for 'broadcasters'
- low barrier to entry by new CPs leading to an enrichment of the content industry and greater consumer choice
- 'network edge solution' - no need to upgrade the Internet to embrace new technologies – a problem faced by technologies such as IPv6, Multicast, etc
- software-only solution - no specific hardware required at the clients' end
- distributed solution - less sensitive to congestion, node failures,...
- global reach, equivalent to the reach of the Internet
- supports content download, progressive download and live streaming (ideally the same client software)
- scalability – so long as stream bandwidths are < average (upload capacity)
- P2P client could be embedded in the CE devices and mobile/portable devices
- can be integrated in hybrid (P2P/CDN) systems

19.2 Weaknesses

- relies on a 'neutral', open Internet that does not filter out specific types of content
- relies on the 'generosity' of users to share their resources as part of the commons
- leeching behavior is possible
- an open door to content piracy if not administered
- audience measurement difficult
- content distribution control difficult (for fully decentralized P2P systems)

- content search and discovery difficult (for fully decentralized P2P systems)
- tends to require a symmetric access to the end-user, which is not the case for broadband Internet today (uplink 1/5th to 1/10th of downlink capacity)
- relies on firewall puncturing strategies for a wide variety of devices (in order to optimize sharing and efficiency)
- special software has to be installed on client platforms if no plug-ins are used
- best effort delivery (Internet property)
- latencies of access to VoD and Live streaming content are very high (often in the order of 10-20 seconds) as compared to the near instant access properties of modern Unicast solutions

19.3 Opportunities

- new distribution channel (either complementary to traditional satellite, cable and terrestrial) or standalone, for professional content
- distribution opportunity for user-generated content – everyone can become a content provider
- extend the reach of content to new communities such as younger target audiences and geographical locations
- distribution costs lowered substantially
- accessibility to new market players – new business opportunities, increased competition
- network neutrality trend can help protect the political and legal context for P2P
- In line with the participation/conversation trend on the Internet through lowered barrier to entry
- the public already relies heavily on P2P technologies. P2P content represents 10% to 50% of the overall Internet traffic today
- P2P may be a means to force ISPs to enable Inter-domain Multicast
- with an increase in untrustworthy content providers the strategic advantage arising from credibility of public service providers becomes even more important.
- growth in uplink capacity tends to be more sustainable than growth in video bitrate due to HD
- ISPs may migrate customers to symmetric access products (by tiering services with the promise of HD on-demand) leading to higher ARPUs and other downstream opportunities
- an open public platform which no one company can control or switch off

19.4 Threats

- More competition on the content market since almost any individual or organization can afford delivery for modest cost.
- regulation could ban peer-to-peer technologies due to bad reputation in relation to content piracy
- ISPs could slow down or block P2P traffic to prioritize their own solutions (like Multicast) or reduce traffic travelling over their infrastructures
- the media industry has a natural tendency to expect precise audience metrics for Internet applications, although it was never the case for traditional TV and Radio
- longtail trend: the attention goes towards niche content and less towards mass media.
- gatekeeper role of broadcasters and media organizations could be weakened

- CDN providers and ISPs may integrate P2P technologies which may become opaque to broadcasters, growth in efficiency may benefit ISP/CDNs only
- Greater piracy leads to crisis in intellectual property conservation and recession of creative industries in general.

20. Main Challenges and Open Questions

This section gives a short list of some important study and research topics which need to be considered before P2P services come into regular operation. The list is neither complete nor is it prioritised.

20.1 Efficiency

What is the benefit of P2P compared to Unicast in terms of bandwidth saving?

Some P2P solutions use all available upload capacity that a user disposes. For example, if a peer gets a radio stream at 200 kbit/s and have a 20 Mbit/s upload capacity line (say, at work), a P2P solution could use the whole upload capacity of the 20 Mbit/s, thus sending streams to 100 other peers. This approach is typically taken by traditional P2P solutions, where P2P is seen more like a simple distributed server setup. Many newer P2P solutions however have default behaviour by which upload and download for each individual peer is more balanced and takes into account only a portion of upload capacity that is really needed.

Some P2P solutions use capacity from 'passive peers', namely the peers which do not actively render (consume) content but could contribute to aggregate upload capacity. Such a contribution could be fairly small; however, as the numbers of passive peers can be large, the total contribution could be important.

In summary, efficiency of a P2P system depends on two elements: a) whether or not uses a totality of upload capacity available or just a minimal portion of it, and b) whether or not it makes use of upload capacity provided by passive peers.

This topic requires further study.

20.2 Upper bound/scalability

Is there an upper bound (i.e. maximum number) of concurrent users beyond which the P2P network degrades or even fails to work properly?

P2P networks could be expected to scale to several millions of concurrent users. In principle, this should be possible since each peer that joins the network increases the capacity of the network. However, this theory ignores the real-life of available uplink and downlink bandwidth for each peer. In practical xDSL networks downlink and uplink may be heavily unbalanced; their ratio could be between 10 and even 20. This means that uplinks could get saturated 10 times earlier than the downlinks. Consequently, the number of uploading peers should be at least 10 times greater that the number of downloading peers. If the number of downloaders is the same as the number of uploaders, then the average download capacity is limited by the average uplink capacity.

For live P2P streaming of high quality streams, the imbalance and saturation represent a significant challenge in delivering full quality to all participants. The instantaneous upload capacity is the constraint. It is still possible to serve all peers but with lower quality. A possible solution is to use a hybrid P2P/caching approach, where distributed caches are used to seed streams into the P2P network.

Analysis of optimum relationship between P2P and caching for a given network is required.

20.3 Network Topology

The P2P networks can have either a tree or mesh (grid) topology or indeed a mixture of the two. A study of pro’s and con’s of different P2P network topologies is required in terms of efficiency, latency, reliability and other technical parameters.

20.4 (De-)centralization balance

A P2P network can have different degrees of centralized and decentralized (distributed) functionalities for media (video, audio, data) distribution and control/management data. Considering BitTorrent-type tracker control/management functionalities, they can be either centralized in a central server, or distributed across peers or indeed a combination of the two.

Is there an optimum balance of centralization or decentralization of the P2P network?

20.5 Latency

As the signal propagates through the network, there may be some significant delay in accessing the streams. Latency is the time that elapses from injection of input signal in the encoder to rendering at the end user device. P2P delivery often occurs over several hops; however, as ping times on the Internet are down to a few milliseconds such extra hops do not significantly degrade latency. The main source of latency is the need for extra buffering which may amount to several tens of seconds (typically between 20 to 30 s) up to several minutes.

Latency is a general Internet distribution problem and not specifically P2P-related problem.

20.6 Service reliability (QoS)

Internet is a best effort, no-QoS environment. There is no guarantee for service quality in terms of service interruptions, discontinuities, dropouts. Typically these occur due to buffering, traffic congestion, jitter, packet loss, etc.

Specifically, P2P quality may suffer due to clients opting out (the tree-type of P2P is more fragile than the mesh-type P2P). A possible remedy to poor quality of P2P system is to use an ‘intelligent’ plug-in/application to be installed at the end user PC. A plug-in could control whether or not content has been received and how well it has been played out (and not just sent). If there is any service problem, the plug-in could try to access other sources on-the-fly and pick the best one.

The question is what other means to improve the quality of P2P services could be used? Would it be possible to use FEC, packet retransmission and other measures often used in the Internet distribution?

20.7 Traffic Peaks

The ability to accommodate large traffic peaks arising, for example, at the start of major broadcast events (such as the Eurovision Song Contest), is very important performance characteristic. It is largely expected that a P2P system should be able to cope with sudden expected or unexpected surges of traffic.

There are two problems.

a) Peak requests just prior to the event, and

b) Build-up rate capability (i.e. how fast the number of connected peers receiving certain quality can grow after the start of the event)

a) In traditional solutions using a central server, bottleneck problems can arise when many client requests are arriving in a short timer interval. A P2P system can be designed in such a way as to spread the instantaneous burden and find the best way to get the stream. This can be done by distributing the ability to capture requests from the newly joined peers by the nearby peers that are already members of the P2P network. In other words, a 'plug-in solution' embedded in each client can help to spread a burden more evenly.

b) The build-up capability of the P2P network can be significantly improved by injecting several streams at a number of locations at the same time. A hybrid solution consisting of special caching servers or super-peers (along with P2P clients) could be helpful.

This topic requires further quantitative study.

20.8 Critical mass of peers

Ideally, P2P systems should be able to operate with a small number of peers. The question applies what is the critical mass (i.e. a minimum of users) to make P2P viable?

20.9 Switching time

Zapping between TV channels in P2P environment can take several seconds, as only one TV channel is transmitted at the same time. The problem of switching time may be quite severe, as most P2P systems use long buffering provisions.

A study is required to shorten the channel switching time. Would it be possible to use the FCC (fast channel change) approach, as used in IPTV?

20.10 DSL asymmetry

Clearly, P2P efficiency would improve if networks on which P2P is overlaid were more symmetric. However, the reality of DSL networks is that upstream bandwidth is very low (for ADSL2+ asymmetry may be as high as 24:1). This problem may be partly resolved by using the principle that the upload bandwidth would be proportional to the download bandwidth. In other words, the more you download, the more capacity of your computer you should allow others for upload.

20.11 Large bitrate (HDTV) streaming

The question is how P2P networks could be able to handle real-time streaming of very high bitrate streams, e.g. HDTV (typically, HDTV may require a total of 8 to 12 Mbit/s).

20.12 Geolocation

A geolocation provision is often required to satisfy copyright requirements. How can a decentralized P2P system handle geolocation to cope with copyright?

20.13 Virus contamination

A P2P-type network is 'ideal' for fast propagation of viruses and spyware. What measures shall be taken to prevent virus propagation? In cases where the prevention is not effective, how can a virus be removed promptly and efficiently?

20.14 Possible Blockage of P2P packets

This question is closely related to net neutrality. The Internet Service Providers (ISPs) are confronted with the estimate that P2P traffic may consume up to three quarters of all Internet bandwidth. This leaves very little room for other applications. In order to preserve subscriber access to these applications, especially in peak times, some ISPs try to simply block (filter out) P2P packets, thus disabling all P2P applications.

Fortunately, most of ISPs endeavour to identify more constructive solutions aiming at preserving subscriber access to P2P services. 'While ISPs are desperately seeking to reduce their spiraling bandwidth costs relating to increased P2P traffic, they also recognize file sharing as a significant driver of new broadband subscribers to their network services'.

(Michael Hoche, Aberdeen Group). The measures taken to save P2P may be multiple, ranging from imposing download quotas, traffic shaping, 'charge more' and 'least-cost routing' to caching. Additional caching can significantly reduce transit and network costs while maintaining the user experience. However, there is some cost involved in installing and maintaining P2P cache servers. Therefore the question applies as to which entities should share the caching costs, so that the expected quality level of P2P services could be maintained and indeed increased.

21. Summary and Conclusions

Access to audio-visual content via Internet is an increasingly valuable resource to both citizens and broadcasters. Over the last several years, traffic on the Internet has grown by more than 50% a year. Distribution of audio-visual materials across the Internet has become an important activity involving extensive investments in the Internet network infrastructure, end user devices and new media-encoding technologies. As the creation and consumption continue to grow, the traffic volumes will grow at an even faster pace. As the video files and streams are very large, the danger is that the network load becomes unsustainable and the Internet may eventually collapse. To ease the network load, novel distribution technologies are required. It has now become clear that the centralised server-client approaches practiced by many network providers are no longer the only viable proposition and that distributed models are rapidly developing to support more Internet traffic.

To this end, the Peer-to-Peer (P2P) system - as a representative of a decentralized, distributed network architecture - continues to be an attractive solution for carrying media across the Internet. P2P seems to be an attractive long term solution for the following main reasons.

- P2P is an overlay to the existing broadband network and does not change the existing network infrastructure
- Relatively low service cost (per GB delivered) compared to Content Delivery Networks (CDN)
- Low investment and maintenance cost
- Scalability to millions of concurrent users
- High service reliability
- No single point of failure
- Lower network load (compared to Unicast).

Today, a large number of broadcasters and portals are carrying experiments and operational P2P services for downloading, VoD and streaming using several commercially available P2P systems. However, these services are not compatible and are only available via PC. We need to develop a standardised P2P solution which will enable access to all the content, services, applications and end user devices.

P2P is not the only solution for delivering content across the Internet. There may be situations, where P2P is not the most suitable choice, and other delivery approaches may be more appropriate. There are many challenges still to be resolved with P2P. The market reality is that P2P is increasingly being used in conjunction with other delivery mechanisms such as CDNs, as opposed to stand-alone P2P systems. These hybrid approaches are increasingly popular in the Internet distribution market (examples: Octoshape going with CDNetworks, Joost moving to CDN, etc.)

It should be made very clear that the adaptive HTTP streaming approach is a very strong candidate for content delivery, potentially displacing P2P from the no.1 position. However, it remains to be seen how the promising technology will be converted into a viable market proposition.

Today, pure P2P systems are most successful at low bitrates (such as live radio/audio-only streaming) but – due to asymmetry of today's broadband networks – high-quality video broadcasts (1 Mbit/s or more) are more challenging and require additional mechanisms (super peers, edge caching, CDNs, etc).

P2P is most efficient for local, close area services but seems to be less suitable for inter-domain services involving two or more different ISPs, especially if the broadband network is star based (examples: the Netherlands, the UK).

Several P2P approaches have recently been refined (or are in the process of becoming more 'intelligent') in order to understand the broadband network topology, so that media flows can be optimised (e.g. smallest number of hops, shortest latency, lowest transport cost, etc) – examples: Octoshape, Tribler (NextShare). Such 'intelligent' P2P systems are potentially able to substantially ease the network load and improve service quality, compared to traditional P2P systems (such as BitTorrent, Kazaa, etc).

Intelligent P2P such as *P2PNext* libswift should be able to radically change the Internet distribution market. Libswift is a promising P2P solution, which is lightweight, inexpensive and simple, so that it can be smoothly integrated in any 'thin client' devices including mobile and portable devices. It should ideally enable all players in the value chain to improve their services (win-win situation), however, all players will need to adjust themselves to the new P2P paradigm (if P2P is commercially successful). For example, in the light of commercial pressure and vigorous competition, ISPs may wish to make their networks more symmetrical, offer higher download and upload capacities, introduce flat rates and remove service caps (ceilings).

It is important to develop, test, standardize and market as soon as possible an 'intelligent' P2P algorithm which is network topology aware. Such algorithms are being developed by the EC-funded project P2P-Next, NAPA/WINE and others. These research projects include all constituencies including broadcasters, operators, manufacturers, etc. and could provide a sound basis for the system stack standardisation of advanced P2P approaches. The possible use of scalable video codecs based on H.264 and specifically optimised for the use in P2P networks has been alluded to.

It has been shown that P2P can be implemented not only in PCs but also on low-cost consumer devices. For example, Pioneer has already developed a first low-cost prototype P2P-enabled STB which is able to perform live TV streaming. There are indications of significant interest among CE manufacturers to develop prototype P2P-based devices –examples: Pioneer NextShare, Tribbox (ST Microelectronics), AHT, etc. However, our expectations should be realistic – there is no standardised P2P solution allowing the mass market deployment, as yet.

P2P has been implemented on an Integrated Circuit designed by ST Microelectronics (Tribler algorithm). Any STB/TV set manufacturer is able to use this chip to market their products in the open market.

In order to accommodate P2P-enabled CE devices in the market, there is a requirement to standardise a P2P algorithm and other layers of the stack (e.g. video coding, middleware, metadata, signalling, etc). This document provides some reflections on how P2P could be successfully introduced in the non-expensive consumer electronic devices and gateways.

The future P2P standard should be able to accommodate any business model (examples: flat rate, various solidarity approaches, bandwidth as currency, etc.). Ideally, the same P2P standard should be able to perform streaming, downloading, VoD, catch-up TV, etc.

The Digital Video Broadcasting (DVB) Project (in cooperation with the IETF) is undoubtedly the right standardisation development body to agree the P2P standard for TV sets and CE devices. However, the standard should be adopted at the appropriate moment, not too early and not too late. Research into P2P has been underway for over 10 years and the R&D work is now mature enough for the foundations of a standard to be created. The standard should be modular, expandable and evolvable. Other international bodies involved in Internet media standardisation (such as Open IPTV Forum) are expected to adopt the DVB approach on P2P, rather than develop their own approach.

It is important to understand the difference between PC-based P2P and CE device-based P2P services. The latter may be constrained by a limited processing power and memory footprint, upgradeability of software and introduction of new services. However, it should be pointed out that the both PCs and CE products are becoming more and more powerful in time but the difference will probably remain to keep the cost affordable.

P2P requires efficient audio and video codecs in order to minimise the bitrate of the stream. Ideally, these codecs should be standardised, open, licence-free, and may become part of a browser.

In order to overcome high asymmetry of DSL-type access networks, possible contributions of ‘passive’ or ‘stand-by’ P2P-enabled devices to the aggregate upload capacity would be highly advantageous. It would be important to identify mechanisms which encourage users not to switch off their devices when they do no longer watch TV programmes. Overall energy consumption of potentially millions of such stand-by devices should therefore be minimized.

All Internet TV players should endeavour to establish a suitable legal framework which would disallow unfair blocking or degradation of content, and prevent arbitrary discrimination against particular providers of content, applications, services or devices. Internet service providers (ISPs) play an important role in this effort and should be able to develop their businesses and obtain the fair revenue streams for them (example: subscription, multi-tier service levels, etc)

P2P services should include accurate and timely audience statistics, information about the share between CDN and P2P bandwidth (bandwidth saving due to P2P use), and the individual contributions to aggregate bandwidth.

Building Communities: P2P is a ‘natural’ environment for building groups of users interested in the same programme and willing to chat, share, communicate, invite friends (and friends of friends) and help each other. Incentive mechanisms in the P2P environment are very important – people should be encouraged to share content and contribute to the aggregate upstream capacity. Each of them can become a broadcaster.

P2P can be used not only in totally unmanaged open Internet environment but also in a closely managed environment by using a transaction server controlling the database of all users, all transactions, all content and metadata. In fact, there are indications that commercial system would always require some degree of system management.

A vision of an ‘Open Service Sphere’ platform that can be globally used to develop any application of user interest is now being promoted by the European Commission [24]. Such an open service platform will integrate innovative web-based services, sensor networks, mobile networks and cloud computing platforms to develop new business applications in the Internet. It goes without saying that P2P-based network technologies may provide a suitable technology framework for such a platform.

In the past P2P protocol has been tainted by its association with copyright piracy. Many detractors still associate P2P with illegal file sharing. It is now high time to remove this bad name and consider P2P a suitable legal mechanism for transport/distribution of legal TV/video content. All content-related copyrights should be cleared in advance to P2P broadcasts. A suitable Digital Rights Management system should be used to enable the legal use of content.

We should be aware that P2P is not a universal distribution panacea. It will not replace conventional satellite, terrestrial and cable delivery mechanisms. It will rather complement them when and if appropriate. In the most likely future scenario, P2P technologies will be used in combination with the CDN, IP Multicast and adaptive HTTP streaming solutions for efficient, reliable and high-quality distribution of audio-visual content across the Internet.

In concluding this report, we should state that it is presently not possible to take a clear-cut position in favour of using P2P for all purposes. It is not possible to recommend P2P as a single best solution without any reservations. There are many excellent solutions now available in the Internet media distribution market, and there is a good prospect of more solutions to come in the near future.

22. References

[1] Discovering P2P : Michael Miller, Sybex, San Francisco, 2001

[2] From P2P to Web services and Grids – Peers in a Client/Server World : Ian J. Taylor, Springer-Verlag London, 2005

[3] Eric Klinker: When should you choose P2P?, Streaming media Magazine, September 2008

[4] Enabling Digital Media Content Delivery, Emerging Opportunities for Network Service Providers, Strategic White Paper, Velocix – An Alcatel-Lucent Company, March 2010

[5] On peer-to-peer (P2P) content delivery, Jin Li, Peer-to-Peer Networking and Applications, Vol. 1, No. 1. (7 March 2008), pp. 45-63.

[6] Open Content Aware Networks (OCEAN), FP7 proposal, Networked Media & 3D Internet, Annex 1, January 2010

[7] Solving Open Internet Peer-To-Peer Content Distribution Using CE Home Devices (IBC 2008, 12. September, 2008)

[8] How P2P Technology Can Change Broadcasting As We Know It (11. September, IBC 2006)

[9] EBU Technical Review, July 2005: EBU introducing Octoshape — a new technology for large scale streaming over the Internet

[10] Hei X, Liang C, Liang J, Liu Y, Ross KW (2006) Insights into PPLive: a measurement study of a large-scale P2P IPTV system. In: Workshop on Internet Protocol TV (IPTV) services over World Wide Web in conjunction with WWW2006, Edinburgh, Scotland, May

[11] S.M.A. Abbas, J.A. Pouwelse, D.H.J. Epema, H.J. Sips, 'A Gossip-Based Distributed Social Networking System,' Enabling Technologies, IEEE International Workshops on, pp. 93-98, 2009 18th IEEE International Workshops on Enabling Technologies: Infrastructures for Collaborative Enterprises, 2009.

[12] Give-to-Get (G2G): Free-riding-resilient Video-on-Demand in P2P Systems: J.J.D. Mol, J.A. Pouwelse, M.Meulpolder, D.H.J. Epema, and H.J Sips, Department of Computer Science, Delft

[13] M. Meulpolder, J.A. Pouwelse, D.H.J. Epema, H.J. Sips, 'BarterCast: A practical approach to prevent lazy

freeriding in P2P networks,' ipdps, pp.1-8, 2009 IEEE International Symposium on Parallel&Distributed Processing, 2009

[14] Towards an Incentive Mechanism for Peer-to-peer Multimedia Live Streaming systems: Thomas Silverstone, Olivier Fourmaux, Jon Crowfort, International Conference on Peer-to-Peer Computing 2008, Aachen, 8-11 September 2008

[15] Standby Power Consumption for 2008 Philips TVs, http://www.philips.com/shared/assets/global/sustainability/philips_tv2008_stand_by_overview.pdf

[16] P2P-Next Deliverable 5.3.1 – Tools for Rich Metadata and Signposts, Michael Eberhard (UNIKLU) et al, December 2009

[17] Net Neutrality in Europe, Neelie Kroes, Vice president of the European Commission Commissioner for the Digital Agenda, Address at the ARCERP Conference (L'Autorité de Régulation des Communications Electroniques et des Postes), Paris, 13 April 2010

[18] P4P Provider Portal for Applications, Haiyong Xie et al, SIGCOMM 08, August 17-21, Seattle, Washington, USA

[19] ITU-T Rec. H.264 Annex G/AVC Amendment 3, Scalable Video Coding (SVC), January 2007

[20] Joint Video Team JVT of ISO/IEC MPEG and ITU-T VCEG, International Standard of Joint Video Specification (ITU-T Rec. H.264, ISO/IEC 14496-10 AVC), March 2003

[21] T. Wiegand, G.J. Sullivan, G. Bjntegaard, and A. Luthra, 'Overview of the H.264/AVC video coding standard,' IEEE Trans. on Circuits and Systems for Video Technology, vol. 13, no. 7, pp. 560–576, July 2003.

[22] V. K. Goyal, 'Multiple description coding: compression meets the network,' IEEE Signal Processing Magazine, Vol. 18, Issue 5, pp. 74 – 93, Sept. 2001

[23] T.Tillo, M. Grangetto, G. Olmo, 'Redundant slice optimal allocation for H. 264 multiple description coding,' IEEE Transactions on Circuits and Systems for Video Technology, Vol. 18, No. 1, pp. 59-70, Jan. 2008

[24] Future Internet - the way forward, Neelie Kroes, Vice president of the European Commission Commissioner for the Digital Agenda, Future Internet assembly, Valencia, 14 April 2010

[25] Freenet – A Distributed Anonymous Information Storage and retrieval System: Ian Clarke, Oskar Sandberg, Brandon Wiley, Theodore W. Hong, published on July 1, 2000, available from http://freenetproject.org/

[26] Digital Video Broadcasting (DVB): Internet TV Content Delivery Study Mission Report, DVB Document A145, December 2009

Appendix 1: Functional and Operational Requirements for an EBU P2P System

(Guidance for P2P System Developers)

The 'EBU P2P system', hereafter termed as 'EPS', will be used by EBU members for distributing their streams and files across open Internet.

Purpose

EBU Members can use EPS for the following purposes:

- live streaming
- on-demand archive files distribution to individuals
- progressive downloading and on-demand streaming
- narrowcasting/podcasting (pushing content to selected groups of users/subscribers)

Example Usage Scenarios

This section gives some examples of usage scenarios, as follows:

- streaming of live radio or television events (simulcasting)
- downloading archive files after live events
- viewing/listening only, no storage possible
- progressive downloading of high-quality content
- downloading and re-distributing of files to all users over the public Internet
- downloading and re-distributing of files to other members in a private Internet network (tunnelling).
- after a file has been acquired, usage can be restricted in terms of time, location and number of devices
- applying rights management to files and streams
- applying various billing mechanisms
- applying social networking mechanisms

Basic System Requirements

The following are some basic requirements for EPS.

The EPS system refers to the transport mechanism only. Data about the programmes (e.g. EPG type data) will be available in the same way it is available now (e.g. Media channels, websites, etc.) The transport mechanism should be Codec agnostic.

EPS:

- shall be based on an open source approach
- shall support all platforms, i.e. operating systems such as Windows, Mac and Linux and browsers such as IE, Firefox and Opera

- should not hamper or disrupt users' PC or any existing applications (Java, etc)
- should be capable of distributing both files and streams
- should be easy to install on user's machine and should not require administrative rights
- servers should be made available to any EBU member who requires it
- should be highly scalable, resistant to peaks and 'tails'.
- should be able to cope with virus contamination
- should allow for quasi real-time delivery (i.e. latency should be reasonable (below 1 min))

End User License agreement

Not necessary. However, for tracking purposes EUL would be a useful mechanism for a broadcaster to be able to authenticate and track its users.

System plug-in requirements

When turning the computer on, the EPS plug-in should be automatically launched.

- EBU EPS icon should appear in the task panel, as soon as EPS application is launched
- EPS can be started from the 'start' menu.
- EPS should support a range of video/audio players: WM (optionally also Real, QT and Flash)
- EPS MMI (Management Interface) should be a media player.
- Plug-in settings should be simple and available on 'right-click' of EPS icon
- Any EPS updates should be performed on in a user-transparent manner

DRM

EPS should be able to apply a DRM system to control media access

Geolocation

For copyright, a geolocation system should be used on the fly to contain access to a given territory. This system should be transparent to the user.

Geolocation systems may be distributed across the EBU zone in order not to delay propagation of signal.

Other Broadcasters' Requirements

- Broadcasters should be able to track the usage of streams/files (who is connected, how long, which channel)
- Downloading of EPS plug-ins should be 'easy' and scalable (the latter is important for live events)
- Broadcasters should be able to apply GEO or not - depending on content rights issues.

Appendix 2: DVB requirements of P2P

The following extracts have been taken with permission from the DVB IPTV 0442 rev15 document (June 2009).

Internet-TV Content Delivery Architecture

As depicted in the Figure below, a P2P Internet-TV Content Delivery system is composed of several elements:

- **P2P Clients** are network elements which share AV files and/or streams among themselves rather than receive them from a central server. They are typically embedded in Internet-TV devices and work in best effort mode. Typical products are STBs, TV sets, media-enabled PCs, etc. Primary goal is to have the P2P client in the Internet-TV device, but it is not excluded that network equipment can host that function. It could also be part of a managed device connected to DVB-IPTV managed system. In this case the P2P Client is still assumed to operate in best effort mode.
- A **P2P Tracker**[23] is a central element hosting dedicated features and is under the control of the Internet-TV service provider. The P2P tracker coordinates transfer of content (being files or streams) among the P2P Clients. For instance the P2P tracker maintains information about all P2P clients either uploading or downloading data. It may also track which segments of content each P2P client stores, in order to assist in efficient data sharing between clients. A P2P tracker can be centralised or geographically distributed across the network.
- A **P2P Auditing Server** is a logical element under the control of the Internet-TV service provider and hosts system and network monitoring functions. The role of the auditing server is to provide data that reflects the operational status of the system. Its role is similar to the RMS in an IPTV system.
- A **P2P Portal** is a portal which allows the user to connect to many trackers using a single logon facility. A user registers his/her P2P Internet-TV device with the supported Portal – providing information, possibly including user-ID, user name, address and billing details. The P2P Internet-TV device automatically downloads a list of channels provided by different trackers from the portal without logging into all of the trackers. Users select a service from a provider that has an agreement with the portal and are able to buy and consume services from that provider without having to separately register, logon and provide personal details to that provider-related tracker. The portal provides APIs to both P2P Trackers and P2P Clients. A P2P Client does not need to modify anything when a new tracker is added.

[23] Some P2P systems are trackerless, all metadata information is distributed among P2P clients.

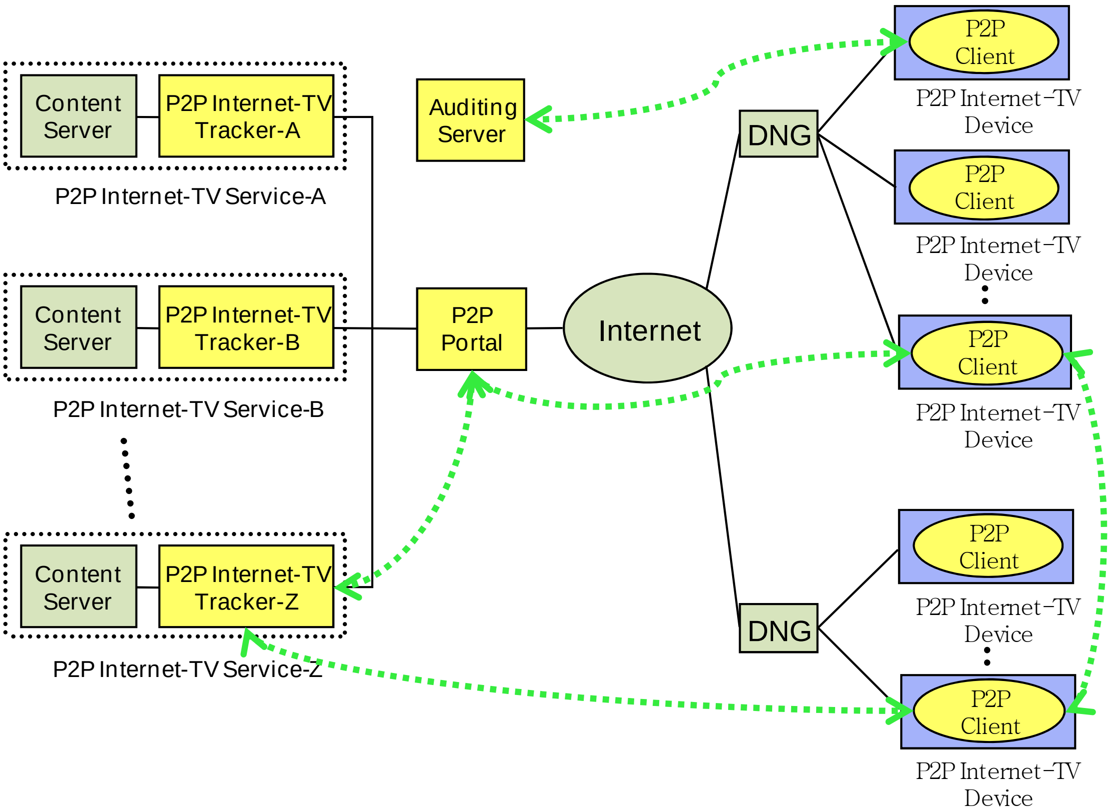

Figure 22: Notional P2P Internet-TV Content Delivery system architecture (Courtesy: DVB IPTV)

The P2P Internet-TV Content Delivery system includes 5 types of interfaces which are proposed to be standardized, as follows:

- P2P Client / P2P Client
- P2P Client / P2P Tracker
- P2P Client / P2P Auditing Server
- P2P Client / P2P Portal
- P2P Portal / P2P Tracker

The P2P Internet-TV device may be owned either by the user or by the Internet-TV service provider who rents it to the user.

The P2P Internet-TV Content Delivery system deals with content delivery and is widely agnostic to other technologies such as e.g. Content Protection, audio/video codecs, content related metadata, interactive middleware, service signalling, etc.

The technical group should give feedback on potential interdependencies, so that work in the commercial group can be prioritized accordingly.

Commercial Requirements for P2P Internet-TV Content Delivery

This chapter provides commercial requirements for P2P Internet-TV Content Delivery.

General requirements

1000	The P2P Internet-TV Content Delivery system shall not assume any specific requirements from the DNG for a successful operation.
1010	The P2P Internet-TV Content Delivery system shall be compatible with both bridged and routed Internet access.
1020	The P2P-Internet-TV Content Delivery system shall allow various types of applications such as • CoD (streaming, progressive downloading and deferred downloading) • Live streaming of TV / Radio events and channels It shall allow for both traditional content and lawful user generated content. ⓘ: For the user the services may look similar to the related IPTV services, but QoE might be lower.
1030	The P2P-Internet-TV Content Delivery system shall not impose any limitation on the content format (coding, resolution, frame rate, etc). Subject to P2P network capabilities, carriage of large streams and files (e.g. HDTV) should be enabled. ⓘ: Internet-TV services might provide lower QoE than managed IPTV services.
1040	It shall be possible to distribute content (or pieces of content) using typical content delivery networks based on servers only (e.g. CDN). In other words, the P2P Internet-TV Content Delivery system shall be able to interface with CDNs. If the P2P system faces extremely severe circumstances (e.g. lack of suitable peers), it shall be able to fallback to CDN mode.
1060	The P2P Internet-TV Content Delivery system shall scale in terms of number of users. It shall not impose any limitation in terms of maximum number of concurrent users. ⓘ: Please consider commercial TV services deployed over the Internet as examples for the typical number of users.
1070	The P2P Internet-TV Content Delivery system should be able to support abrupt increases in the number of peers that may occur at the start of live TV streaming events when potentially large numbers of users may wish to connect and start watching the event.
1080	The P2P Internet-TV Content Delivery system shall support a fast response to user requests for content delivery and content presentation, especially for live streaming. ⓘ: QoE with Internet-TV is expected to be lower than with IPTV.
1100	The P2P Internet-TV Content Delivery system shall not prevent the delivery of complementary streams such as subtitling, spoken subtitling, audio descriptions and signing, in order to provide value-added services and satisfy requirements of different interest groups.
1110	The P2P Internet-TV Content Delivery system shall support high QoE for the user.
1130	The P2P-Internet-TV Content Delivery system shall not preclude the usage of geolocation data. Note: Geolocation means that service availability depends also on the geographic position of the user.
1140	The technical group is requested to consider the existing P2P protocols that are currently widely used on the Internet as a technical basis for P2P Internet-TV Content Delivery system.
1150	Effort should be made to align with other IPTV security related DVB requirements and specifications, and build on them.
1160	The P2P Internet-TV Content Delivery system shall not prevent the implementation of a secure transaction mechanism allowing pay TV or free-to-view content to be provided according to usage guidelines defined in the DVB CPCM commercial requirements.
1170	Within the P2P Internet-TV Content Delivery System channel zapping time shall be minimized.
1180	The P2P Internet-TV Content Delivery System shall support seek operation based on time, allowing the user to move to a relative or absolute time of playing position.

1190	The P2P Internet-TV Content Delivery System shall not prevent content providers from delivering user created contents.
1200	The P2P Internet-TV Content Delivery System shall support using DVB EPG (or BCG) to locate the P2P-delivered channels.
1210	The P2P Internet-TV Content Delivery System shall allow content providers to update their services (including title, description and content) without terminating the service.
1220	The P2P Internet-TV Content Delivery System shall not prevent the user from accessing multiple services from different portals at the same time.

P2P Internet-TV Content Delivery Management

This section captures the requirements on the P2P Internet-TV Content Delivery system from the Service provider point of view.

Storage usage and management

One of the key points in a P2P Internet-TV Content Delivery is the level of replication for the various pieces of content. When a content item is highly replicated, QoE may be satisfactory. When a content item is rarely replicated on the network, QoE may be poor. Therefore, one of the key success factors of a P2P content delivery system is the capability to guarantee a certain response level of the system. A P2P content delivery system will behave well, once a given content is sufficiently replicated within its peers.

For instance, when a new content is introduced in a P2P content delivery system, it is available at only a few peers. Therefore, as long as it is not sufficiently replicated, the response level of the P2P content delivery system may be quite weak. One way to solve this problem is to allow a service provider to manage the content replication, in order to be able to reach the required objectives in terms of the response levels.

In order to achieve this, the storage available in the P2P client is modeled in two logical distinct areas, one under the control of the user where he can manage content (e.g. save/delete), called the user storage area; one under the control of the P2P-Internet-TV service provider, called the service provider storage area.

Different business models are envisaged and it is up to the service provider to decide whether he requires two distinct storage areas or not. So, the support of two storage areas is optional, but may be required by some service providers for implementing their services.

This section provides requirements, so that a service provider can use and manage the storage that is available in a P2P system. These requirements apply to both the user storage area and the service provider storage area.

2000	The storage in a P2P Internet-TV Content delivery system may be divided into two distinct partitions, *user storage area*, where the user can download, save and delete content; and a *service provider storage area*, where the P2P Internet TV Content Service Provider can download and delete content from. Business models shall be supported that have either one or both storage area present in the ITD. ⓘ: It is envisaged that for live streaming services, a separate storage area in the ITD will be used temporarily during the streaming session, the size of which shall be defined at the P2P installation time.
2005	An ITD supporting P2P content delivery as well as other download services such as CDS shall be able to support separate storage areas for each download service.
2010	Content on the user and service provider storage areas shall be available to the P2P Internet-TV Content Delivery system, i.e. other peers from the system shall be able to get content from both

	storage areas subject to access restrictions.
2015	Service provider shall be able to manage the level of the content replication (or the replication of content segments) in the service provider storage area to reach the required objectives in terms of content presentation latency.
2026	It shall be possible for the P2P Internet-TV Content Delivery system to find out the size of the total storage capacity of the P2P ITD (independent of user and service provider storage areas). The storage size information may be used to assess the suitability of the P2P ITD to participate in the P2P Internet-TV Content Delivery system. The P2P Internet-TV Content Delivery system may reject joining if the P2P ITD does not meet a minimum storage requirement.
2027	The P2P system shall be able to query, create and resize P2P storage areas (user and service provider) subject to access restrictions imposed by the user and/or ITD. This may be carried out via a remote management system or via another mechanism.
2028	It shall be possible to create and resize P2P storage area (user and service provider) via a user interface subject to access restrictions imposed the service provider.
2030	There shall be a mechanism allowing a P2P Client to receive updates to a content item, so that the latest version of that content item is automatically available on the client.
2070	It shall be possible to restrict access to content stored in both the user storage area and the service provider storage area to authorized P2P clients..
2075	It shall be possible to detect the tampering of content stored in both the user storage area and the service provider storage area.
2080	It shall be possible to restrict the ability to identify content stored in the both the user storage area and the service provider storage area to authorized P2P clients.

Auditing

An auditing system is necessary in order to monitor the performances of the P2P Internet-TV Content Delivery system. This will require features to be supported in the P2P Clients as well as on an Auditing server (see Figure).

2200	It shall be possible to monitor and report relevant data of the P2P Internet-TV Content Delivery system in order to get information on the operational status of the system. ⓘ: Possible data may include statistics on uploaded and downloaded data, up & down bandwidth, play out errors, etc.
2210	The P2P Internet-TV Content Delivery system shall be able to routinely monitor non-user-specific information (e.g. global usage, statistics, network traffic, etc).
2220	The P2P-Internet-TV Content Delivery system shall not preclude the use of audience measurements mechanisms, such as measuring the size of the audience of a given event, possibly on the fly.

P2P Client management

2300	The P2P Client shall be able to operate behind a NAT without a manual reconfiguration. The P2P client shall be able to operate without a manual reconfiguration of a firewall and/or DNG. The design should consider the need to conserve resources (e.g. ports) in the firewall, NAT and/or DNG.
2310	The specification shall not preclude using several P2P Internet-TV devices on the same Home Network.
2315	The specification shall not preclude running several instances of the P2P client on the same P2P Internet-TV device.
2320	It shall be possible to define priorities between different P2P Flows. For instance, user requested P2P Flows (CoD or Live) may have higher priority than background content placement & management P2P Flows.

2330	The P2P Client shall make its context information available to other P2P Clients or/and Tracker/Portal (see Figure 1). ⓘ: Possible context information includes its IP address & port, its maximum uplink bandwidth, the content locally available, etc.
2350	It shall be possible to embed the P2P Client in a CE device, based on chip sets and storage, typically used in high volume consumer digital TV devices.
2360	It shall be possible to minimise free-riding by allowing P2P Clients to contribute some upstream capacity when they are not actively used (e.g. a kind of active stand-by mode). ⓘ: The technical group is expected to make proposals which are compatible with low power consumption modes.

Content acquisition and consumption

This section captures the requirements on the P2P Internet-TV Content Delivery system from the user point of view.

3000	When the user asks for a specific content on the P2P Internet-TV Content Delivery system, it shall be possible to calculate the estimated start-up delay or download time.
3020	The specification shall not preclude that manufacturers build devices allowing users to start several Internet-TV services in parallel and that priorities for QoE can be handled.

Security requirements

P2P content security

The following content security requirements shall apply to the P2P Internet-TV system:

4000	The P2P Internet-TV Content Delivery system shall allow detecting modifications of the content that happened during its transport over the Internet and Home Network.
4010	It shall be possible that only authorized P2P Clients can download certain content items.
4020	The P2P Internet-TV Content Delivery system shall allow verifying consistency between the downloading content and the metadata content title.
4030	It shall be possible to prevent an eavesdropper from accessing content transmitted over P2P Internet-TV Content Delivery system.
4040	The P2P Internet-TV Content Delivery system shall support means to enable Service Providers to detect illegal content in the system.
4050	The P2P Internet-TV Content Delivery system shall not preclude the distribution of protected content (including CA-protected, DRM-protected and CPCM-protected content).

P2P network security

As there is no work on network security in CM-IPTV so far, we ask the technical group not to preclude inclusion of such technology in the system. DVB may decide later to start certain work in this area.

4100	The specification shall not preclude that measures can be taken, so that the P2P Internet-TV Content Delivery system resists to Sybil[24] and pollution[25] attacks. In particular, this would mean the unique identification and authentication of P2P Clients during a P2P session using secure mechanisms.
4120	The specification shall not preclude that measures can be taken, so that the P2P Internet-TV Content Delivery system protects the communications between P2P Clients and between a P2P Client and infrastructure elements against unauthorized modifications. This includes both signaling and content.
4130	The specification shall not preclude that measures can be taken to detect overload of P2P Flows or overload of Peers in case of malicious attacks (e.g. DoS[26]).
4140	The specification shall not preclude that P2P Clients authenticate infrastructure elements from which content is downloaded. ⓘ: It is open how authentication can be made in horizontal devices, not initially related to a service provider.

User privacy

The overall P2P-Internet-TV Content Delivery system shall fairly preserve the user's privacy as further described below:

4200	The P2P Internet-TV Content Delivery system shall not disclose the user's credentials (if any) to other users, externals and operators unless explicitly allowed by the concerned user.
4210	It shall be possible to prevent an eavesdropper from identifying content transmitted over P2P Internet-TV Content Delivery system through metadata.
4220	The P2P Internet-TV Content Delivery system shall protect the communications between P2P Clients and between a P2P Client and infrastructure elements against eavesdroppers. This includes both signalling and content.

[24] Attack wherein a reputation system is subverted by creating a large number of pseudonymous identities in P2P networks.

[25] Attacker-initiated insertion or alteration of passive data in a system which degrades the quality of the service provided by this system.

[26] Attack whose purpose is to prohibit a legitimate user the use of a service or of a system.