

IP TV distribution

A whitepaper discussing IP TV in broadband networks

TABLE OF CONTENTS

1	IP TV IN BROADBAND NETWORKS.....	3
2	BACKGROUND	4
3	THE TECHNICAL CHALLENGE.....	5
3.1	PRIORITY.....	5
3.2	THE IMPORTANCE OF HEADROOM	6
3.3	MULTICAST DISTRIBUTION	6
4	THE IMPORTANCE OF EASY ADMINISTRATION	8
5	THE CRITICAL CHOICE OF CPE	8
6	THE TECHNOLOGY.....	10
6.1	PRIORITY (QoS)	10
6.2	MULTICASTING	10
6.3	AUTHENTICATION AND SECURITY	12
7	SUMMARY.....	12

1 IP TV in broadband networks

The widespread adoption of IP technology in recent years has irreversibly changed the networking landscape. Once a technology purely intended to host data networks, IP now influences the design of other communication networks.

This document describes the technical concepts behind the distribution of television content in a common, IP-centric infrastructure. The document specifically describes the use of an environment denoted as “triple-play”, a network in which three different families of baseline services – data, voice and video – share the same underlying infrastructure.

This document also provides insight into the market logic and commercial consequences associated with different architectural approaches to TV distribution in IP networks. For many network operators, it has been a major dilemma to decide the extent to which an existing physical cable or network structure should be kept and enhanced. Some have already decided to implement a modern, IP-based network architecture based on twisted-pair cabling or, preferably, an all-fibre network. This document recommends this solution due to the ability it provides of easily adding services and in this way increasing network revenue.

The key success factors for distribution of TV services discussed in this document are:

- A guarantee that only the end users who subscribe to and pay for a particular channel receive this channel.
- The assignment of high priority to the TV services, to ensure high end-user-experienced quality.
- A cost-efficient approach to the configuration of access nodes and the distribution of TV services.
- A clear ownership of the customer premises equipment (CPE). This is important in order to decide where service functionality should reside; whether it should be embedded in the CPE or not, and who is responsible for this functionality.

Discussions concerning IP TV distribution often labour under confusion surrounding the Internet and its underlying technology platform (the TCP/IP family of networking protocols). It is important to notice the distinction between Internet connectivity and IP-based networking. Internet connectivity, IP telephony and IP TV distribution are all separate services. The only thing they have in common is that they all use the same, highly efficient underlying IP network.

Properly designed, TV distribution utilizing IP technology can be easily separated from an end user’s Internet access. Hence, licensing issues related to Internet distribution and measures to restrict uncontrolled re-distribution of content over the Internet do not need to become as major a concern as they are in situations in which all traffic is handled as a single service. This is the case when offering Internet access as a single service, supplemented by web casting or Internet streaming.

2 Background

Networks for the distribution of television have been based on analogue technology, merely relaying TV transmissions over a wireline distribution medium. These networks usually constitute separate network islands, each with its own terrestrial and satellite receivers, and a shared coaxial cable.

TV channel management was wide-meshed and often inadequate in such networks, since the individual mix of channel subscriptions was often defined in a set-top box, either hardwired or configured by means of the periodic distribution of decryption keys to end users. The outcome of this is common knowledge: it quickly attracted creators of counterfeit encryption engines, fake smart cards and similar fraudulent equipment.

A major obstacle has been the inability to create a truly tamper-proof mechanism of restricting access to premium subscription channels (such as movie channels). Creating this mechanism is difficult in a network environment in which many TV channels are distributed to large groups of households.

Star topologies with individual cables connected to a local hub or switch were introduced in the early years of cable television networks, enhancing the possibilities of offering user-specific services. In combination with smart set-top boxes, sometimes provided with additional functionality to allow Internet connectivity, these devices often did their job fairly well.

Several recent factors have created a need to reassess the situation. These factors include the internet revolution, a rapid increase in the performance of fibre infrastructure and a corresponding fall in its price, and the advent of IP telephony.

Network operators today are becoming increasingly aware of how crucial multi-service offerings will soon become. Other related issues have arisen, including the need to address the flexible configuration of end-user devices on a mass scale, and the possibility of hosting a number of different service providers in the networks. The amount of headroom available for expansion is another important topic, both in terms of the number of services, the capacity of each service (which determines, for example, the maximum number of television channels that can be included), priority issues, and the total number of users that can be handled by one network operations centre (NOC).

The answer to these requirements in the context of the currently available technology has been rather unclear. Networks based on TCP/IP have traditionally lagged behind dedicated distribution networks for television in terms of real-time performance. However, the situation has now changed. With correct network design, containing priority mechanisms and multicast functionality (intelligent handling of bandwidth-consuming real-time bitstreams), an IP networking infrastructure is now the most attractive alternative for realizing the vision of true triple-play networking. The most demanding challenge has undoubtedly been to host multi-channel television distribution. The discussion below will explain why this is the case, and it will show how the challenge can be met.

3 The technical challenge

Television is a highly demanding network service. TCP/IP was not designed to cater for such time-critical and bandwidth-hungry bitflows. TCP, the most common transport mechanism on the Internet, was primarily designed for reliability and for a situation in which 10 ms or even 100 ms of extra delay did not affect the overall impression of service performance (such as the transmission of an e-mail message or a large data file). The prime concern was to correctly transfer each and every bit of information.

Extra delay is, of course, unacceptable when dealing with real-time traffic such as telephony and television. Real-time traffic protocols for the Internet have therefore been developed, with less demanding transport mechanisms, where a continuous bitflow has been a more important concern than receiving every bit correctly. Consequently, Internet streaming has been designed to accept a fairly large level of packet loss (dropped packets).

3.1 Priority

Premium voice and video services cannot operate acceptably when subject to the level of dropped packets that is usually experienced when streaming traffic on the public Internet. It is therefore necessary to employ other solutions. One important step has been the introduction of priority mechanisms. In modern IP-based triple-play networks, the dominating technology that determines priority is called “Diffserv”. Time-critical traffic is given a priority labelling when it enters the network, and the traffic is then handled accordingly by routers and other equipment along the path to the end user.

Telephony (voice) traffic is commonly labelled with the highest priority in the network, followed by video and audio services. Video frames are generally not as time-critical as voice synchronisation, which is a factor worth mentioning here.

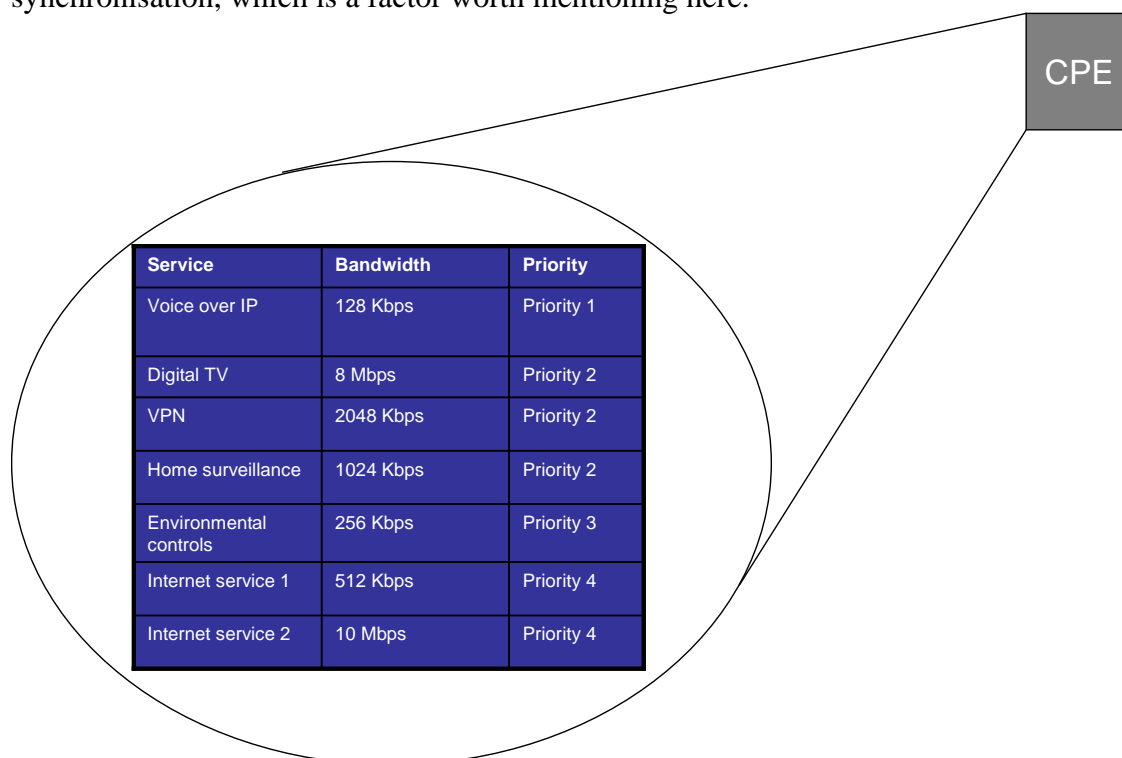


Figure 1: The priority of the different types of services offered to the end user.

3.2 The importance of headroom

Other factors such as latency (the total, accumulated network delay) and jitter (a variable packet reception rate) also affect quality. These quality issues are generally solved by adding large capacity headroom in the network (additional data bandwidth). This is fairly easy to achieve today with respect to IP telephony, since telephony traffic only requires data capacity in the region of 10-100 kbps. The accumulated bandwidth requirement for TV distribution, however, rapidly becomes unmanageable when the number of channels increases. Each digital video bitstream requires 1.5-15 Mbps, depending on the content and the predetermined quality level. It is simply impossible to implement more than a few additional TV channels in the network solely by adding more extra headroom. We may add here that this is not strictly a question of technology, rather one of operational economy.

3.3 Multicast distribution

The solution to the accumulated capacity problem is a procedure known as “multicasting”, in which bandwidth-hungry traffic such as TV channels are sent once and directed only to the receiving party specifically requesting them, thereby minimizing unnecessary distribution. This can be compared with traditional traffic in an IP network, which is of unicast type, where every user requests his or her own bitstream from the source.

Initially, multicast was intended to revolutionize the Internet, but it did not achieve widespread acceptance in the interwoven structure of independent networks on the Internet, mainly because multicast traffic did not fit the volume-based business model of most IP network operators. (The setting of priority was another “failed revolution”, by the way. It proved to be difficult to set priority when “high priority” traffic crossed administrative borders in the public Internet.)

In a private network, such as a triple-play access infrastructure, the situation is different. Priority labelling and traffic optimization by means of multicasting have become key factors, and are crucial to offering real-time services and keeping capacity development within reach.

Another important feature of multicast distribution is the increased security that it makes possible. Multicasting also implicitly brings forward the need to control traffic at the network layer (also referred to as “layer 3” or simply “L3”) in the access segment of the network – but it will not lead to the same complexity in the network terminals (set-top boxes, etc.) installed at each end user as a layer 2 approach would do. Figure 2 and 3 illustrate the differences between a layer 2 and a layer 3 approach.

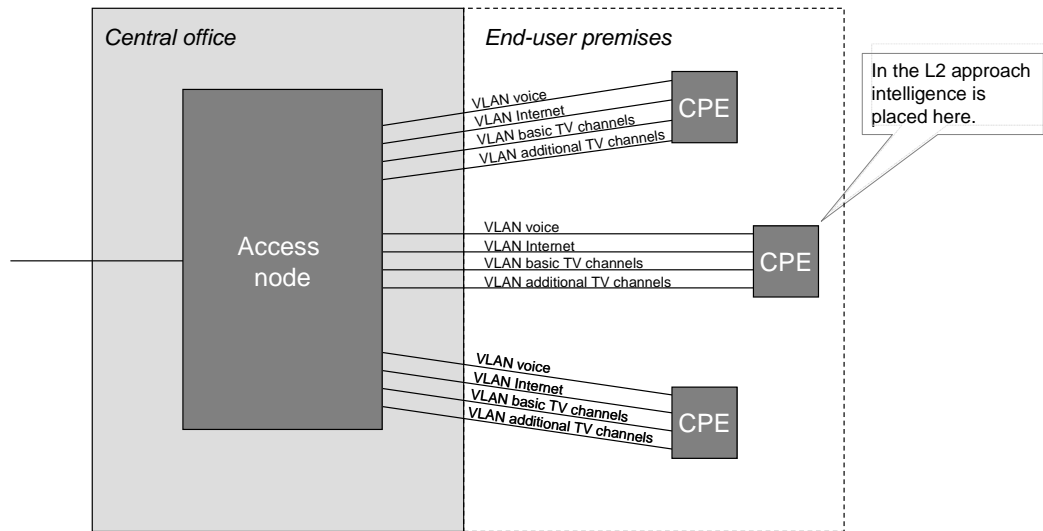


Figure 2: The layer 2 approach. In this approach advanced intelligence is required in the customer premises equipment (CPE). Each service offered to the end user requires configuration of the CPE to establish a new VLAN. Complexity at the end-user premises will increase as more services are added.

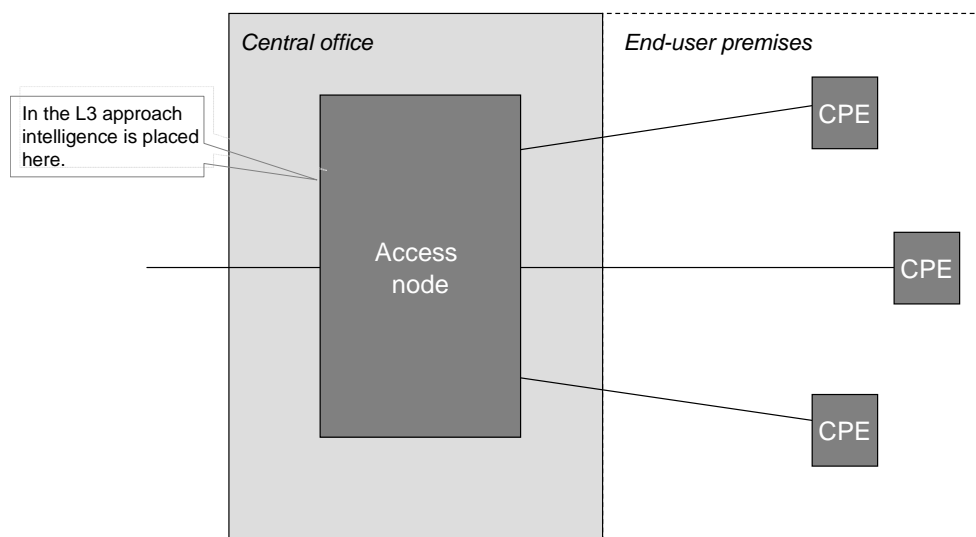


Figure 3: The layer 3 approach. In this approach the intelligence is placed in the access node, where service delivery is handled as routing decisions. The connection between the access node and the CPE carries a customised mix of triple-play services forwarded from the L3 access node. This approach reduces complexity in the end-user equipment.

4 The importance of easy administration

Layer 3 is the recommended architecture when designing a triple-play network that should cater for flexibility and allow for several service providers sharing the same network. A major advantage of a layer 3 approach over a traditional link layer (layer 2, L2) approach is the increased level of control that it offers, while avoiding much of the headache traditionally associated with the configuration of end-user equipment.

In the layer 2 architecture, the end user receives services based on the definition of the VLANs. The end-user terminal will usually have one VLAN defined for voice (telephony) and another VLAN for Internet access. More VLANs must be assigned when video services such as TV channels are added. This approach increases the complexity of the end-user equipment, and the task of simultaneously upgrading services (reconfiguring) 10,000 – or even more – end-user boxes is highly demanding.

In contrast, the task of configuring new services becomes much easier when a) using a priority labelling mechanism, b) using the multicast technology, and consequently c), controlling the traffic at the network level all the way out to the access node. This is true for the configuration of new services, both individual configuration for each end user and for configuration on a mass scale, such as that necessary during major network upgrades or permanent changes in the content offerings.

The difference is, in fact, so significant that it will result in a network in which each individual service can be controlled in a second-by-second manner. No TV channel will be distributed to an end user unless it has been subscribed to, and a number of new business cases can be foreseen, given that a channel can both be configured and deactivated in a matter of seconds.

Another advantage of the layer 3 approach is the fact that it can easily cope with many different service providers. In a layer 2 environment, the end-user equipment must be reconfigured for each and every change in subscription status. Administrative problems may therefore arise when several service providers use such a network, if every provider is to be allowed access to individual end-user devices.

5 The critical choice of CPE

The most critical question for any access network operator today is undoubtedly whether service configuration should reside in the end-user device (CPE, customer premises equipment) or not. This will affect not only which individual services will have the potential to become profitable, it will also affect the associated price-floor for new services (which is defined by the cost for setting up and maintaining the service). Furthermore, the decision will influence the level of control over services that is possible, and it will influence security and reliability issues.

It is not wise to leave these critical devices in the hands of the end-user. CPE units can be dropped on the floor or disassembled, and may be costly to replace.

Another important question is how much should be hardwired in the CPE. One of the worst cases for IP telephony may prove to be voice-over-IP (voip) hardwired or awkwardly implemented directly into the CPE. This is a good example of the unfortunate situation in which the network owner's equipment will dictate how service providers can implement their services.

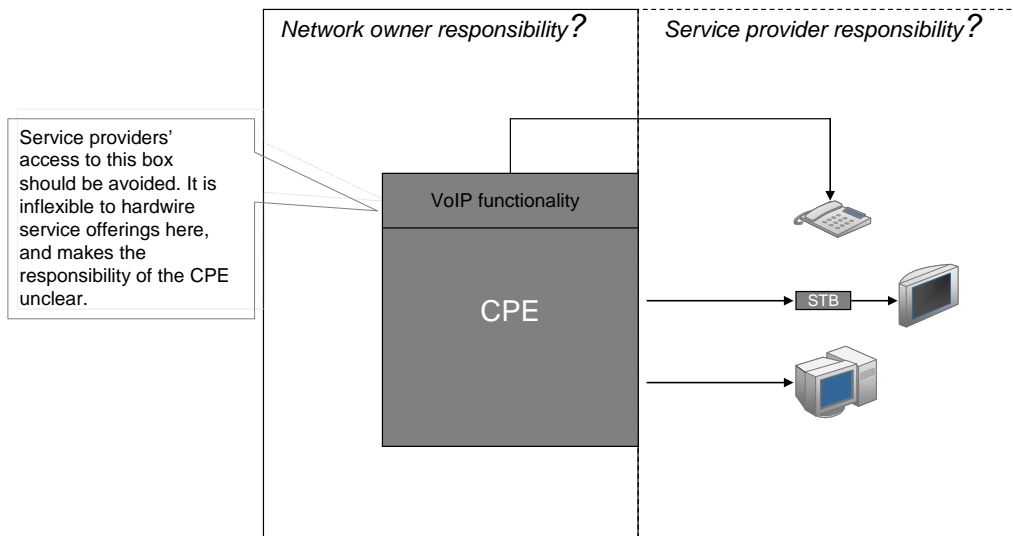


Figure 4: Any functionality embedded in the CPE will decrease the flexibility and make administration of this box more complex. It becomes even more complex in an open-access environment in which several service providers need access to the configuration. The issue of responsibility is also cumbersome, and such questions may arise as whether the CPE is the responsibility of the network owner or the service provider.

The best solution is to establish a clear boundary between the responsibility of the network owner and the responsibility of each service provider at the outgoing ports of the CPE. The complexity of the hardware is reduced and the network investment will be secured for any future changes in the service offering.

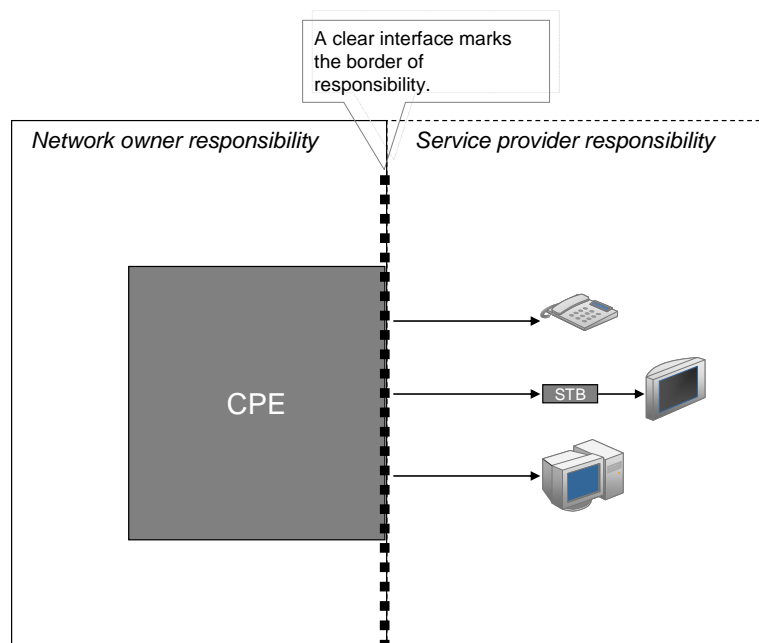


Figure 5: The CPE should not have embedded functionality for any service, due to the unclear issue of responsibility. Here, the CPE is the responsibility of the network owner, and a clear boundary is drawn between the network owner and the service provider.

6 The technology

Three factors are critical in order to make access networks designed for triple play and TV distribution truly efficient. These factors are: quality-of-service, the possibility of multicasting, and a reliable authentication mechanism. Quality-of-service (QoS) offers a reasonable way of ensuring that there is enough capacity available for every bitstream in the network at any given moment. Multicasting is a highly efficient method of avoiding unnecessary traffic. Authentication is crucial in order to control and restrict access to network resources such as premium television channels for individual end users.

6.1 Priority (QoS)

The technical standardization body behind TCP/IP and the Internet, IETF, has been discussing efficient mechanisms for setting priority for more than 10 years. These discussions have led to a specification known as “Diffserv”. Diffserv allows a network administrator to manage network traffic in several different categories. Unlike earlier solutions to the problem of priority, which allowed only a basic traffic-class management, Diffserv allows the administrator to add a certain kind of metric value to every packet in the network, in a way somewhat similar to modern routing metrics.

The value added to the packet is known as the “DSCP” (Diffserv Code Point), and it is added to the TOS field in the IP packet header when a packet enters the network, either from the boundary of another network or from an access connection (an individual end user).

A well-known problem with priority is that of administrative control. This is also the reason to why QoS mechanisms have not been very successful in network environments in which traffic must pass between many administrative domains, such as networks run by different network operators. The Internet is a good example, formed by a large number of independent domains called “Autonomous Systems” (ASs).

It is important to verify all DSCP values in order to maintain a reasonable level of control. A “web of trust” is created in this way between the routers in the network. Any packet not originating from a trusted source will be examined and, if necessary, re-labelled. This applies, for example, to traffic from individual end users. This trust relationship ensures that an end user cannot gain a higher priority for his or her own traffic in the network simply by labelling it with higher priority. This avoids the sorry situation that has arisen with spam e-mails, in which every message is marked with “highest priority”.

It should be noted that it is much harder to reach this level of control in a network in which control is placed in the link layer (a layer 2 network).

6.2 Multicasting

In a traditional cable TV (CATV) network every television channel is distributed over the complete physical cable infrastructure, and the reception in the individual end user’s home is controlled by a simple filter or encryption device. This design is simple, but inefficient. It is also limited, in the sense that it is sometimes necessary to group several TV channels in order to match a certain filter, and access is commonly based on shared encryption key rather than an individual encryption key.

The optimal method of distributing a large number of television channels over a network is, obviously, to distribute a video channel only to the end users who explicitly request it. This is not possible in a traditional TCP/IP network.

Multicast is a method in which the source sends the packets once, and the packets are then distributed in a tree-like fashion only to parties that explicitly subscribe to a particular service. Multicast was specially designed to optimize traffic distribution, and it thus solves the problem with congested networks that is caused by inefficient distribution of TV content. Network routers along the route must be multicast-enabled in order for multicast to function, and a number of rendezvous points (RPs) for the traffic must be established.

In order to receive a certain television channel, an end user sends a request to the nearest RP in the network to “join” the multicast stream that contains this channel, and the distribution tree is then instantly extended to include that end user. Similarly, a user may request to “leave” the channel. The association is then immediately removed and the tree no longer contains this end user. These requests are sent with a routing protocol called IGMP, and the multicast-enabled routers use a router-to-router protocol called PIM to exchange information about changes in the multicast tree structure. See below illustration of multicast.

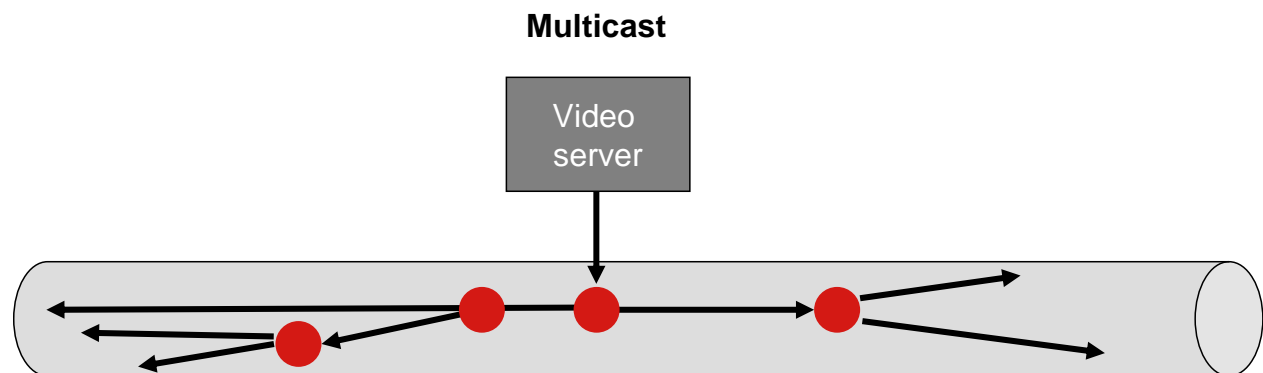


Figure 6. Multicast uses a logic-tree-structure. The data is sent only once, and then branched off at certain router nodes in the network called rendezvous points. Individual end users can quickly join or leave a multicast channel by sending a request to such a rendezvous point.

The use of multicast has been somewhat limited until recently, due to its technical (routing) complexity and a lack of interest among public IP network operators for the function.

The problem of routing complexity, however, can be greatly reduced within a well-structured administrative domain such as a triple-play network run by an access network operator. This means that multicast has become an attractive choice in order to keep traffic volumes low and reach a high level of cost efficiency in the network.

The ability to keep a tight control over who may join the different multicast channels in the network is, of course, critical. A reliable authentication mechanism that is hard to bypass must for this reason be added.

6.3 Authentication and security

Access to TV channels must be restricted in almost every network that distributes such channels. Early solutions included simple filters that scrambled certain frequency ranges or brought unsubscribed premium channels out-of-sync for the receiving party. These are, however, rather basic solutions to the access problem, and have been replaced to a certain extent by various encryption devices and regularly distributed shared encryption keys. This has, to some extent, increased the level of security.

The ultimate way of controlling television content would be to use a method in which only the exact numbers of subscribed channels are distributed from the access node to the end user – and only after a robust means of authentication.

A multicast-enabled level 3 network that passes routing information all the way out to the access nodes provides exactly this capability. Here, the network operator is offered full control over which packets are allowed over the cable to each individual end user. Moreover, the control is carried out at the access node rather than at the user device (CPE), simplifying the CPE.

A very efficient method of authenticating end user equipment – and ensuring that the connected device is an authorized one – is to use its MAC number as identification. The device issues a DHCP request to the network to obtain an IP address, and the access node adds some information about the physical connection (cable) over which the request was made. This combined information can then be checked against centrally stored information about the end user and the related individual mixture of services. This information is sent back to the access node, which enables the services. In this way, the user is not involved at all in the authentication process. No manually entered encryption keys are used, and it is almost impossible to manipulate the network to deliver unsubscribed services.

7 Summary

Distribution of television channels is an activity that consumes large quantities of bandwidth, despite modern video compression formats. It is, therefore, essential to optimize the way television signals are handled in an IP-based triple-play network.

This paper outlines an architectural approach to triple-play networks based on level 3 (network level) switching from the core of the network out to each access node connecting individual end users. This design makes administration easier than it is in level 2 (link level) networks and it reduces the need to invest in complex end-user devices. It also paves the way for the use of multicast technology to minimize network load, at the same time giving the network operator unprecedented control of service delivery to each end user.

With the multicast method, television channels are neither shared by all and decrypted by some, nor requested by and sent to many end users in parallel, which is a bandwidth-consuming strategy. The television channels are sent once from one point in the network, and then branched off repeatedly, forming an efficient tree-structure that distributes the signals only to requesting and authorized end users. In a well-tuned triple-play network, individual end users can join or leave multicast television channels in the matter of a second.

Combined with Diffserv, an advanced method of ensuring Quality-of-Service in IP networks, a multicast-enabled triple-play network performs excellently with respect to such key economic factors as cost of administration, level of access control and optimal use of network resources.