CHAPTER 2

Video Application Components and Architecture

This chapter discusses the segmentation of the video application architecture into logical components that are required for broadcast video and video on demand (VoD) services. The function of each component is described, as well as the basic interfaces needed between each component and other components of the system.

This chapter describes possible video architectures and components only. For the actual tested implementation, see Chapter 4, "Implementing and Configuring the Solution."

Note Because there are currently few standards regarding application architectures for either broadcast or on-demand IPTV/video service over a DSL infrastructure, this solution makes no specific assumptions regarding the application architectures implemented by the vendors of specific video equipment. However, although there are few standards for video application architectures, the functionality implemented is fairly consistent from vendor to vendor.

For a list of the video components that were tested in this release, including product names and part numbers, see Table 3-1 on page 3-3.

This chapter presents the following major topics:

- Video Application Components, page 2-1
- IPTV/VoBB Transport Architecture and Issues, page 2-9

Video Application Components

Figure 2-1 on page 2-2 illustrates the logical relationship of the application-layer video components needed to deliver broadcast video and VoD services, as well as the basic interfaces between components. Components can be categorized as follows:

- Broadcast Video Components
- VoD Components
- Common Broadcast Video and VoD Components

Figure 2-1 Video Application Component Architecture

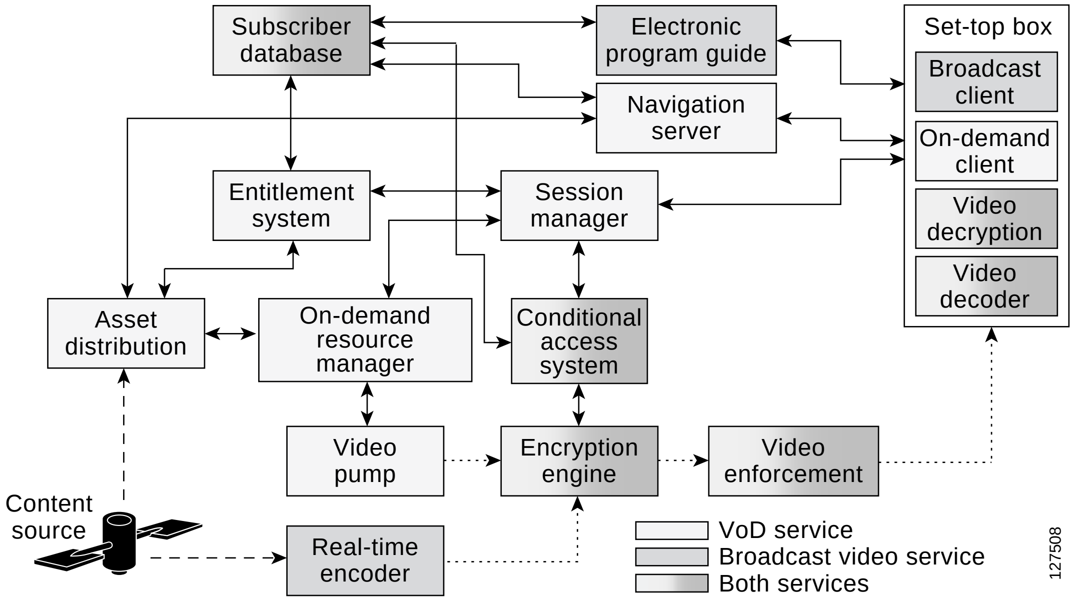

Broadcast Video Components

Broadcast video components (see Figure 2-1) include the following:

- Real-Time Encoder
- Electronic Program Guide
- Broadcast Client

These are described below.

Real-Time Encoder

The real-time encoder takes a live feed from a broadcaster in either analog or digital format and converts it into a compressed digital stream that is encapsulated in IP packets. The input to the encoder may be in a digital format that uses digital MPEG-2 over ASI format, or it may be in an NTSC, PAL, SECAM, or other analog format. The output of the encoder is a digitally compressed stream that is encapsulated in IP headers and sent to a multicast group address. The compression method used by the encoder may be either MPEG-2, WM9, or MPEG-4/AVC, while the IP-based transport encapsulation used is MPEG-2 transport over either UDP/IP or IP/UDP/RTP. Since the real-timer encoder is configured to encode a specific channel, no control interfaces are required between it and other video components.

Electronic Program Guide

The electronic program guide (EPG) provides information about available broadcast channels to the broadcast client application running on the IP set-top box (STB). The EPG is often implemented as an HTTP server and formats available channel listings as web pages. The EPG application authenticates and authorizes a subscriber for broadcast services. The EPG may also provide a customized view of channel listings that is based on the packages a particular subscriber has subscribed to. Both of these functions require an interface between the EPG application and the subscriber database. In addition to providing a graphical listing of available channels, the EPG provides the IP multicast address to which the channel is sent in the IP network. The broadcast client uses this address in Internet Group Management Protocol (IGMP) messages that are sent during the processing of a channel change.

Broadcast Client

The broadcast client is an application process running on the STB that is responsible for providing the user and control interface for broadcast video services. The broadcast client, in conjunction with the EPG, implements a subscriber authentication interface for set-top-based services. Authentication is typically done by means of an application layer authentication protocol such as HTTP in conjunction with a shared secret such as a username/PIN pair.

The broadcast client displays available broadcast-channel information using data from the EPG and implements the control interface for channel change by means of IGMP. Since the DSL line may be capable of supporting the bandwidth of only a single broadcast channel, the IGMP process for changing channels must ensure that only a single video broadcast stream is sent to the STB at a time. The broadcast client implements this by sending an IGMP leave for the current channel and then waiting for a configurable period of time for the broadcast stream to stop. After this timer expires, the broadcast client sends an IGMP join for the new channel. The full channel-change time, documented in Broadcast Video Channel-Change Time, page 2-14, includes these IGMP factors as well as other factors specific to video compression. (See also Effect of Multicast on Channel-Change Performance, page 3-19.)

VoD Components

VoD components (see Figure 2-1 on page 2-2) include the following:

- Asset Distribution System
- Navigation Server
- Session Manager
- Entitlement System
- Video Pump
- On-Demand Resource Manager
- On-Demand Client

These are described below.

Asset Distribution System

The asset distribution system (ADS) takes video content from content providers and uses business rules to distribute that content to different locations in the video service provider’s network. Video content may be provided to the ADS through a number of different methods. These methods include the use of pitcher/catcher systems, which receive video content from content providers over satellite links, and manual processes such as file copies from other network servers.

Standard video content objects include the actual MPEG video, images for display during content navigation, trailers, and metadata files that provide information about the files contained in the object.

The ADS may be used to modify the metadata of a video asset to add business rules such as the price of the video, the distribution window, the VoD subscription package that the video is part of, whether the content needed to be encrypted, and so on. On the basis of these business rules, the ADS replicates the video asset to the on-demand resource management component of video servers in different locations.

Navigation Server

The navigation server provides information about available VoD content to the on-demand client application running on the STB. The navigation server is often implemented as an HTTP server and formats available VoD content as web pages. The navigation server uses information provided by the asset management system to determine which VoD content to display to the subscriber. For subscription-based VoD services, the navigation server may use the information in the subscriber database to customize the view of the video content presented to the subscriber, depending on the packages the subscriber has purchased.

Session Manager

The session manager is the central point of communication for VoD session requests that originate from the on-demand client on the STB. It manages the life cycle of a video session and is responsible for arbitrating the various resources required to deliver the video stream associated with the on-demand session request. Many vendors of VoD equipment and software provide a logical "session manager" function, though this function goes by a variety of different names.

When the session manager receives an on-demand session request from an on-demand client application, it first uses the services of the entitlement system to determine whether the subscriber is authorized to view the requested video content. If the request is authorized, the entitlement server includes additional information in the authorization response, such as the encryption format to be used for the content.

When the session manager receives the authorization response, it determines the best VoD server complex to use for the session request, based on the subscriber's IP subnet. The session manager then contacts the on-demand resource manager for that VoD server complex to request a video pump for the session. If the VoD content needs to be encrypted in real time, the session manager contacts the conditional access system (CAS) to request a real-time encryption engine for the session. The CAS responds with the decryption keys to be used by the STB to decrypt the video stream.

After all of the resources for a VoD session request are obtained, the session manager responds to the on-demand client with information about the IP/UDP/RTP transport parameters for the video stream to the STB. If the stream is to be encrypted, the session manager (or a key manager with which it coordinates) includes the decryption keys for encrypted video content in the response as well. Finally, the session manager includes the IP address of the video pump that was selected for the session. The IP address of the VoD pump is needed by the on-demand client in order to send stream control requests through Real Time Streaming Protocol (RTSP)—such as pause, fast forward, rewind—for the session.

Entitlement System

The entitlement system is responsible for determining whether the movie requested by an on-demand client is authorized for viewing by the subscriber associated with that client. The entitlement system uses information from the ADS to build a database indicating which videos are associated with different on-demand subscription packages. When the entitlement system receives an entitlement request from the session manager, it uses this database to determine with which orderable on-demand package the requested video is associated. The entitlement system then uses the services of the subscriber database to determine whether the subscriber associated with the entitlement request is entitled to view the requested video.

Video Pump

The video pump is the streaming storage component of a VoD system. The video pump contains locally or remotely connected storage that is organized in such as way that it can send any piece of stored media at a known constant rate. The streaming portion of the video pump is responsible for pulling requested files from the storage system and for pacing the output of each requested file to the network though the use of a shaper. Video pumps must be capable of implementing basic stream control functionality, such as fast-forward and rewind, to respond to requests from the on-demand client during the playout of a media file.

In addition to being able to stream media out, video pumps are also responsible for ingesting media for storage in the storage subsystem. While in general the functionality of a video pump is fairly independent of media format, the ingest function may have functionality that is specific to a particular media format. An example of this type of media-format dependence is the generation of trick files for use with stream control functionality such as fast-forward and rewind. Video pumps used in broadband environments are typically capable of storing and streaming both MPEG-2 and MPEG-4 content. (Only MPEG-2 was tested in this release.)

On-Demand Resource Manager

The on-demand resource manager (ODRM) is responsible for managing the streaming resources and storage of a group of video pumps. The ODRM is responsible for locating and replicating content, as well as for allocating video pumps for the on-demand session requests it receives from the session manager.

On the ingest side, the ODRM is responsible for taking content received from an asset management system and replicating it to one or more of the video pumps it controls. The ODRM makes decisions on when and where to replicate content on the basis of such information as asset metadata and the demand for each title (as indicated by on-demand session requests).

On the streaming side, the ODRM responds to on-demand session requests from a session manager by locating a video pump that contains the requested title, has the capacity to stream the title, and is capable of reaching over the transport network the subscriber that generated the session request.

On-Demand Client

The on-demand client (ODC), an application process running on the STB, is responsible for providing the user and control interface for on-demand services. The ODC provides the user interface for browsing on-demand content using the services of the navigation server. The browsing interface of the ODC is typically implemented by means of an embedded HTTP-based browsing application.

The ODC interfaces to the session manager to make requests to stream on-demand content. It also interfaces to video pumps to make stream-control requests for movies that are actively being streamed.

Common Broadcast Video and VoD Components

Common broadcast video and VoD components (see Figure 2-1 on page 2-2) include the following:

- Conditional Access System and Encryption Engine
- Broadcast Video Bandwidth Enforcement
- Set-Top-Based Video Decryption and Video Decoder
- Set-Top Box

- Subscriber Database

These are described below.

Conditional Access System and Encryption Engine

The conditional access system (CAS) is responsible for the key management and distribution infrastructure associated with the encryption of video content. Video encryption is used as the second tier of protection against theft of content. The first tier of protection for both broadcast and on-demand services is performed as part of the on-demand and broadcast client applications running on the STB. These applications use the services of the EPG and navigation server to authenticate the subscriber and provide a customized view of available channels and content based on the services the subscriber has purchased. For on-demand services, the entitlement system also checks whether the subscriber is authorized to view requested titles, with the result that the ODC does not allow the subscriber to view unauthorized content. While application-layer authorization protects against the theft of content from authorized STBs, it does not protect the video stream itself. Video encryption using CAS provides this second layer of protection against theft and unauthorized viewing of video content.

Note CAS was not implemented as part of the first release of the solution.

Because conditional access adds an additional level of complexity and cost to a video delivery system, service providers typically use CAS-based encryption only on premium-tier broadcast channels and on-demand titles. For broadcast services, encryption must be done in real time as the video stream is delivered. For on-demand services, encryption may be done either in real time as the content is streamed or as part of the process of replicating content to video pumps. The process of encrypting video content as part of replication is called pre-encryption.

Video encryption may be done on either a tier or session basis. In tier-based video encryption, a single set of encryption/decryption keys is used for all of the video content associated with a particular service offering. Subscribers that are authorized to view the content associated with the service are delivered the decryption keys needed for that service ahead of time. Conditional access for broadcast video services is always implemented by means of tier-based encryption, because a single video stream may be viewed by many subscribers simultaneously. Decryption keys for broadcast video services are delivered in a secure manner to the STB through the EPG. In session-based video encryption, decryption keys for a piece of content are generated and delivered to the subscriber on a per-session basis. Session-based encryption may be used with VoD content. Because decryption keys are generated only on a per-session basis for session-based encryption, they may be used with either real-time or pre-encryption techniques.

In a typical CAS, the encryption of digital services can be achieved by using entitlement control messages (ECMs) and entitlement management messages (EMMs). In order to generate the final keys needed to decrypt a particular video stream, the STB must receive and decrypt the correct ECMs and EMMs. EMMs provide keys that can be decrypted only by a specific subscriber, while ECMs provide keys that are specific to a particular video stream. Because EMMs are specific to a subscriber, they are always generated ahead of time. Because ECMs are specific to a particular video stream, they may be generated ahead of time when pre-encryption is used, or they may be generated in real time when real-time encryption is used. ECMs are typically delivered in band as a component of the video stream.

Whether the content must be encrypted may be determined by a number of factors. For on-demand services, content providers can require content to be encrypted by enabling the “Encryption” field in the metadata file associated with the video asset. For broadcast services, the video service provider statically configures the video stream from real-timer encoders to be sent either directly to a multicast group or to a real-time encryption engine, depending on whether that channel is to be encrypted.

The encryption engine takes MPEG streams in and encrypts them in real time using encryption keys received from the CAS. Encryption engines typically use a DES algorithm for encryption.

Broadcast Video Bandwidth Enforcement

Broadcast video bandwidth enforcement in Release 1.0 is implemented as part of the functionality of the aggregation router (AR). The AR enforces a maximum broadcast bandwidth limit by limiting the number of IGMP joins on the ranges of multicast addresses associated with broadcast video to a configured maximum on the aggregation links it controls. The mapping of video channels to multicast addresses can be done in such a way that the AR can associate the bandwidth for different classes of video (for example, standard definition or high definition) with different ranges of multicast addresses. IGMP join limits can then be set for each range of multicast addresses.

For more details on the network enforcement for video broadcast, refer to Multicast Admission Control, page 3-17.

Set-Top-Based Video Decryption and Video Decoder

The set-top box, or STB (see below), includes two components that are responsible for turning the incoming video stream, delivered as IP packets, into an uncompressed digital stream that can be directly turned into an analog TV signal ready for display by a television set. These components are the video decoder and the video decryptor.

The video decoder is responsible for decompressing the incoming video stream. It uses a decompression algorithm that is matched to the compression algorithm used by the real-time encoders for broadcast services. The video decoder may also support additional decompression algorithms for VoD services if VoD assets are compressed by a different algorithm than broadcast channels use.

The video decryptor is responsible for performing decryption on the video stream if the stream was encrypted by the encryption engine when real-time encryption is used, or by an offline encryptor when pre-encryption is used for on-demand assets.

Set-Top Box

The set-top box (STB) is the hardware and common software infrastructure component that is used by the on-demand and broadcast clients as well as by the video decryptor and the video decoder. The STB hardware typically consists of a general-purpose processor and video subsystem that produces an analog television output. The hardware may also include a hardware-based decoder and decryption subsystem. The STB software typically includes an embedded operating system, and may also include application infrastructure components such as a web browser.

Subscriber Database

The subscriber database contains service-level information about each subscriber, for example, which services the subscriber is authorized to use, information used for billing, and so on. The subscriber database may also contain information that can be used for subscriber authentication. An example of this type of information is the username/password information that is used by the EPG to identify and authenticate a subscriber for broadcast services.

IPTV/VoBB Product Architecture

This section describes how the logical components described previously are commonly combined into products that are supplied by current video-component vendors. Figure 2-2 illustrates how the video components presented in Figure 2-1 on page 2-2 are "bundled" into application products. This bundling reduces the number of products and vendors that must be integrated to build a complete video system. It also reduces the number of interfaces that must be agreed upon by vendors.

Note The logical components described in this section represent functional blocks that are common to most video application architectures, and do not necessarily reflect how these functions are bundled into products by video equipment vendors.

The following classes of video products are needed to build a complete broadband video solution:

- Middleware
- VoD Server
- Conditional Access System
- Real-Time Encoder and Set-Top Box

These are described below.

Figure 2-2 Common Video Components

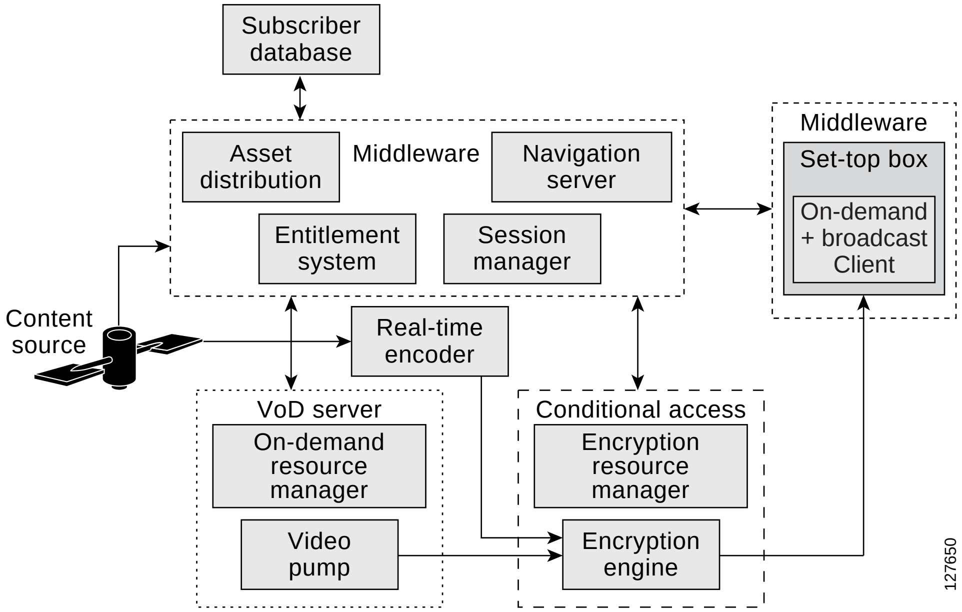

Middleware

Middleware, as defined, has the role of gluing a number of logical components together into a more comprehensive IPTV/video software system. (Note that there are several different middleware implementations. Thus, the following description is a typical example for illustrative purposes.) Middleware implements the user interface for both broadcast and on-demand services. It is also used as the glue software that integrates products from other vendors into an application level solution. Middleware products are often used to integrate multiple VoD servers, conditional access systems, and set-top boxes from different vendors into the same deployment.

Middleware provides the client and server functionality that implements the user interface for both broadcast and on-demand services. The components that provide the client-side functionality are the broadcast and on-demand client applications on the STB, while the components that provide the server-side functionality are the electronic program guide and the navigation server.

Middleware uses the entitlement system and session manager components to integrate the VoD servers used in an on-demand service. The entitlement system integrates the asset ingest function of a VoD server, while the session manager integrates the session plane of the VoD server into an on-demand service.

Middleware uses the session manager and on-demand client to integrate CAS into an on-demand service. These components may be used to pass decryption keys from the conditional access system to the video decryption component in the STB. These components also determine when to use the services of the CAS based on the encryption requirements of the service and each asset associated with the service. Middleware uses the EPG and the broadcast client to integrate CAS into a broadcast service. The broadcast client determines when to use the services of the CAS based on information it obtains from the EPG on each broadcast channel.

VoD Server

The VoD server (one or several) implements storage and real-time streaming functionality for on-demand services. The VoD server consists of a set of video pumps that are managed by an on-demand resource manager. The VoD server integrates with middleware and may also be integrated with the CAS when preencryption is used.

Conditional Access System

The conditional access system (CAS) provides encryption and decryption services, as well as key generation and distribution functionality, for both broadcast and on-demand services. The CAS consists of the encryption resource manager, the encryption engine, and the video decryption process in the STB.

The CAS interfaces to middleware when session-based encryption is used for on-demand services. The CAS may also interface to middleware for encryption key distribution between the encryption resource manager and the decryption process on the STB. Finally, the CAS interfaces to VoD servers where preencryption is used for on-demand content.

Real-Time Encoder and Set-Top Box

The real-time encoder and STB components described in Real-Time Encoder, page 2-2, and Set-Top Box, page 2-7, respectively, are identical to product classes of the same name shown in Figure 2-2 on page 2-8.

IPTV/VoBB Transport Architecture and Issues

To meet the end-to-end transport requirements for broadcast VoD services, the IPTV/VoBB transport architecture provides functional requirements and configuration recommendations for each switching node in the path from the VoD servers to the STBs. This section presents the following topics:

- Video Sites

- Video Service Requirements
- Potential Video Service Architectures
- Service Separation in a Triple-Play Architecture

Note Although this solution is focused on video service, it must work within the context of a triple-play solution. Because VoBB services are fairly new, vendors and service providers do not user the same terminology to describe the major sites. This section describes terminology commonly used for triple-play solutions.

Video Sites

The video sites described in this section are the super headend (SHE), the video headend office (VHO), and the video switching office (VSO). Figure 1-1 on page 1-2 shows the location and roles of the sites and components in a typical IPTV/VoBB deployment.

Super Headend

The SHE is where live feeds for the broadcast video service are located. This site contains the real-time encoders used for the broadcast video service, along with the asset distribution systems for on-demand services. This site may also contain back-office systems such as the subscriber database. Most IPTV/VoBB deployments have a single SHE site; this is the source of most of the multicast streams for the broadcast video service. The SHE typically resides in the core of the transport network.

Video Headend Office

The manned VHO is where the video server complex resides (as well as where optional local/PEG content may be inserted). The VHO is where the majority of the video pumps used for on-demand services are typically located. It is also where the real-time encoders for local television stations reside. A VHO typically serves a metropolitan area of between 100,000 and 1,000,000 homes. The VHO is equivalent to (and may be contained in) the same facility as the point of presence (POP) for Internet access services. Transport for video traffic between the VHO and IP/MPLS core network is provided by a distribution edge router (DER). The DER interconnects the core network and the local video sources to a high-bandwidth distribution network that carries both broadcast and on-demand video to VSOs.

Video Switching Office

The VSOs house the aggregation routers that aggregate local or remotely attached GE DSLAMs. The VSO is typically located in the central switching office. The central switching office is the physical termination point for the majority of the copper loops for the residences it serves. Because ADSL and ADSL2+ rely on short loop lengths to obtain maximum training rates and throughput, the copper loops used for DSL service are often terminated at a location closer to the subscriber than the VSO. This means that the DSLAMs that the VSO aggregates may or may not be colocated in the VSO. The switching equipment in the VSO interconnects the aggregation and distribution networks. Traffic to and from the DSLAMs is aggregated by the aggregation router (AR). The AR resides in intermediate and terminal VSOs.

In order to minimize the bandwidth requirements between the VSO and the VHO, a VSO may include local video pumps that are used to cache popular on-demand content. While the Release 1.0 transport architecture does not preclude the use of video pumps in the VSO, this configuration is not tested as part of the solution test effort.

Figure 2-3 presents a more detailed view of the distribution, aggregation, and access layers of the IPTV/VoBB transport architecture, with video streams acquired, multiplexed, and distributed from the SHE through the VSO (not shown). Both unidirectional and bidirectional interfaces are shown.

Figure 2-3 IPTV/Video over Broadband Transport Architecture: Detailed View

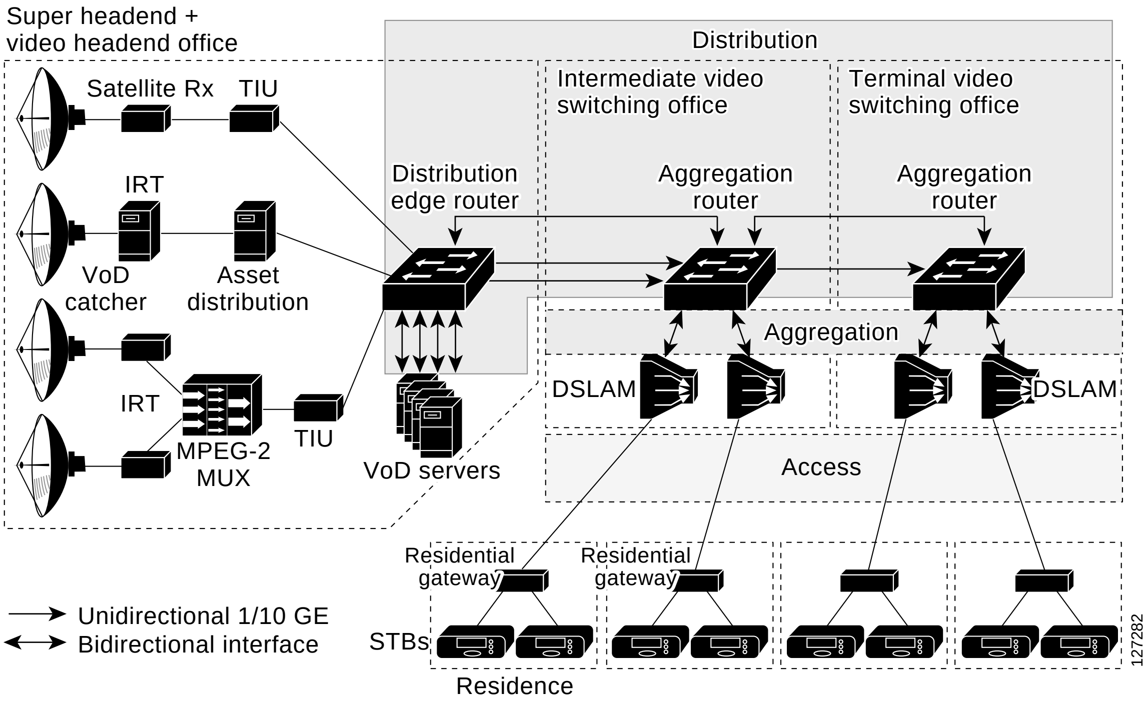

Video Service Requirements

In order to understand better some of the design tradeoffs associated with the transport architecture, it is important to understand common requirements for a video service and how an IP network can be optimized to meet these requirements. This section outlines some common requirements for broadcast video and VoD services. It also describes what design factors in the transport network are relevant to these requirements.

High Bandwidth

The amount of bandwidth that a network must be capable of transporting to support video services is typically an order of magnitude more than what is required to support voice and Internet access services. A standard-definition IP video stream that is carried as an MPEG-2 SPTS stream over RTP uses about 3.75 Mbps of bandwidth. A high-definition IP video stream using the same type of compression and transport uses about 16 to 18 Mbps of bandwidth.

These bandwidth requirements mean that a DSL access infrastructure that is designed for real-time video transport must be capable of carrying significantly more bandwidth than what is needed for VoIP and Internet access services. It also means that the DSL line itself is typically constrained to carrying only one or two video streams simultaneously. The result is that video over DSL service offerings must limit the service to one or two simultaneous broadcast channels or on-demand sessions to a household.

Because the video streams associated with on-demand services are unicast while the video streams associated with broadcast services are multicast, the amount of bandwidth required in the aggregation and distribution networks to carry on-demand streams is much greater than what is required for broadcast services. Also, because broadcast video services use multicast, the amount of bandwidth required in the access and distribution networks scales with the number of channels offered. As an example, a broadcast video service that uses MPEG-2 compression and offers 300 channels of standard-definition content requires more than 1 Gbps/sec of capacity in the distribution network to handle worst-case usage patterns. On-demand services use unicast, so the amount of bandwidth required in the distribution network scales with the number of subscribers and peak on-demand utilization rates the network is designed to carry. For example, a distribution network that is designed to deliver MPEG-2 compressed standard-definition content to 50,000 on-demand subscribers at a 10% peak take rate requires about 19 Gbps of capacity.

Asymmetric Bandwidth

Video traffic is inherently asymmetric, as both broadcast video and VoD flows are unidirectional. The only traffic that is sent in the upstream direction of either service is control traffic that is used to instantiate the video flow. For on-demand services, this control traffic is the session and resource signaling that is described as part of the component descriptions in Video Application Components, page 2-1. For broadcast services, the control traffic is IGMP and PIM signaling that is used to instantiate the multicast flow for the broadcast channel.

Because of this asymmetry, the cost of the distribution network can be reduced by incorporating unidirectional links in the transport path. One of the transport alternatives tested in this release of the solution includes unidirectional transport in the distribution network.

Quality of Service

When broadcast and on-demand video is carried over an IP network, there is an assumption that the video quality experienced by the subscriber is comparable to that experienced by those watching MPEG-2 video carried by cable and satellite networks today. To ensure that any degradation in video quality resulting from the IP transport network is negligible from a subscriber's point of view, most providers allow only one visible degradation in video quality roughly every two hours.

While this allowance is similar to what is allowed for VoIP services, the resulting allowed packet-drop requirement for an IP transport network designed for video services is much more stringent. There are two reasons for this:

- Video is much more highly compressed, so losing a packet may result in the loss of more-valuable encoded video information. If the network drops a single video packet, there is a visible degradation of video quality of anywhere from a single frame up to a loss of one second of video, depending on the kind of encoded information that is lost.
- The receiving decoders, such as the STBs, generally do not have loss-concealment algorithms, whereas VoIP phones and gateways typically support algorithms that conceal dropouts in the voice signal caused by lost packets. In the VoIP case, the network can drop a single voice packet without the listener perceiving any degradation in voice quality—unlike the case for video.

The DiffServ (Differentiated Services) architecture defines packet marking and scheduling behaviors that can be used to ensure that video flows meet the required 10^{-6} drop rate when links are congested. (For details on the QoS architecture for Release 1.0, see QoS Architecture, page 3-46.) Packet drops due to bit errors on physical links must be addressed on a link-by-link basis.

The link-layer technologies used in video networks employ cyclic redundancy check (CRC) algorithms to ensure that packets with errors are not delivered. This means that a single bit error in a video packet results in that packet being dropped when a CRC is performed. Video over IP is typically carried in packets that are approximately 1400 bytes. If bit errors are assumed to be distributed randomly, the resulting requirement for transport links is to ensure a bit-error rate (BER) less than 10^{-10}.

The BER on optical links can be engineered to 10^{-14} or less by ensuring a high signal-to-noise ratio (SNR) on those links. Because Release 1.0 uses optical connectivity in the access and distribution networks, degradation in video quality resulting from bit errors on these links should not be an issue.

However, packet drops due to bit errors on the DSL line can have a significant effect on video quality. The SNR on a DSL line varies as a result of many factors, including loop length, proximity to noise sources, and other factors. In addition, the SNR may vary over time because of factors such as corrosion at connection points, moisture, and so on. Consequently, it may be difficult to qualify a DSL line at the time of installation to ensure a BER of less than 10^{-10} over the life of the video service.

Multiple technologies are available to deal with bit errors on the DSL line. Some common technologies are forward error correction (FEC) and real-time retransmission (RTR). While Release 1.0 does not include the testing of these technologies, future releases of the solution will include technologies to help deal with bit errors on the DSL line.

Service Availability

Service providers deploying video services often have different availability requirements for VoD and broadcast video services, as contrasted below.

Broadcast video services are inherently real time. A subscriber who experiences an outage in the broadcast service cannot come back and continue watching at that point when the outage is over. Because of this and the higher usage rates associated with broadcast services, the availability associated with broadcast services must be very high.

In contrast, the customer disruptions associated with an outage in VoD services are typically much less problematic. A subscriber who experiences an outage in a VoD service can come back at a later time and replay the content—either from the point of disruption or from the beginning. In addition, the peak usage rates associated with VoD are typically between 10 and 20% of the subscriber population. This is much lower than the peak usage rates for broadcast services.

Because of the above factors, service providers have much higher availability requirements for broadcast services than for on-demand services. Consequently, the differing availability requirements between the two services may result in differing transport requirements for each service. For example, the high-availability requirement for broadcast video typically results in the requirement that there be redundant transport paths between the DER and AR nodes of the distribution network. (See Figure 2-3 on page 2-11.) Because of the higher bandwidth and lower availability requirements associated with VoD services, the topologies used for these services may not necessarily require redundant transport paths.

The test topologies for Release 1.0 include a distribution network design that provides path redundancy for both services, as well as a cost-optimized distribution design that provides path redundancy only for broadcast services. In addition, the quality of service (QoS) architecture includes DiffServ marking for broadcast and on-demand services, allowing the network to drop VoD traffic preferentially over broadcast traffic in the event of a network outage. Finally, Release 1.0 supports redundant broadcast video encoders, as well as a method to fail over in a timely manner from one encoder to another.

Broadcast Video Channel-Change Time

An important aspect of a broadcast video service is the amount of time it takes for the system to respond to a channel-change request from a subscriber. While the channel-change time for current analog broadcast services is perceived by the subscriber to be instantaneous, the channel-change time for digital broadcast services is between one and one-and-a-half seconds. The majority of this time is due to the differential encoding and decoding methods used to compress digital video streams.

To reduce the amount of bandwidth required for digital video transmission, compression methods such as MPEG compress the video frames of a digital video stream into three different types of frames. These frames are called I-frames, B-frames, and P-frames. An I-frame is a compressed version of all of the information in one frame of a video stream. An MPEG decompressor can recreate the original frame using just the information in the I-frame. A P-frame is an incrementally encoded video frame that can be decoded with the information in the preceding anchor frame (I-frame or P-frame). A B-frame is an incrementally encoded video frame that can be decoded with the information in the preceding and following anchor frames (I-frame or P-frame).

Because of incremental coding, an important factor in how long it takes to change a channel for a digital video service is the I-frame gap. The I-frame gap defines how often I-frames are included in the MPEG stream. Shorter I-frame gaps result in shorter channel-change times, while longer I-frame gaps result in longer channel-change times.

When a digital broadcast service is run over a DSL access infrastructure, the following additional factors must be added to the delay caused by the I-frame gap:

- STB performance in processing a channel-change request
- Multicast latency in terminating the IP video feed associated with the “tuned from” channel
- Multicast latency in joining the IP video feed associated with the “tuned to” channel
- Whether or not the “tuned to” channel is encrypted by means of a CAS
- Delay to the next cryptoperiod and the time needed to acquire CAS/DRM (digital rights management) decryption keys before the decryption of the “tuned to” channel begins
- Delay in refilling the jitter buffer for the decoder in the STB

The goal of this solution is to provide subscribers with a channel-change experience similar to that currently experienced for digital broadcast services. Most of the additional channel-change delay associated with a DSL access infrastructure is due to the amount of time it takes for the network to stop sending the multicast stream for the “tuned from” channel and to begin sending the multicast stream for the “tuned to” channel. Distribution and Aggregation Transport Architecture, page 3-4, provides a recommendation for a scalable multicast architecture that best meets the channel-change requirements for broadcast video services.

Potential Video Service Architectures

One aspect of a transport architecture for video that must be considered initially is how the service provider sells the video service to the subscriber. This section examines how two potential video service-level agreement (SLA) models affect the requirements of a transport network implemented to deliver the service to the subscriber.

- The SLA for a video transport service is based on transport parameters. A typical transport-based SLA includes factors such as maximum bandwidth, packet-loss rate guarantees, and jitter and latency guarantees.
- The SLA for an application service is based on service-level parameters. A typical video application-based SLA includes the following:

- – The number of simultaneous video channels (live or on-demand) a subscriber a is authorized to view
- – The broadcast channel line-up (basic or premium tier) that the subscriber has signed up for
- – Any subscription VoD content that the subscriber has signed up for

The services the network provides to deliver a transport-based SLA as opposed to an application-based SLA are very different. Table 2-1 provides an overview of the technologies used to deliver the basic functionality of a transport service as opposed to an application service.

Table 2-1 Service-Delivery Technologies: Transport vs. Application

Service Type	Transport Service	Managed Application Service
SLA	Transport parameters: • Bandwidth • Max. drop • Max. latency • Etc.	Video application SLA: • Number of STBs • Basic vs. premium tier
Subscriber authentication/ identification	Network based (examples): • PPPoE • 802.1x • Per-subscriber VLANs • DHCP option 82	Application based: • Video middleware
SLA enforcement	Network based: • Per-subscriber shaping/policing	Application based: • Based on application signaling
QoS	Per subscriber: • Gold, silver, bronze • Classification • Queueing	Aggregate: • Single queue for video service[1]

[1] Assumes all network devices can support DiffServ-type congestion management

The two SLA models are examined in detail in the following sections:

- Transport-Based SLA
- Managed Application-Based SLA

Transport-Based SLA

Subscriber authentication and identification for a transport service is done at the transport layer. Subscriber authentication technologies rely on shared secrets such as passwords or private/public key pairs to establish a trust relationship between the subscriber and the network. Subscriber identification technologies use a well-known property of a subscriber (such as the DSL line to which the subscriber is attached) to identify all packets coming from or to the subscriber. Transport SLA enforcement requires a subscriber identification technology and may also include a subscriber authentication technology.

Common subscriber-authentication technologies used for a transport service include Point-to-Point Protocol over Ethernet (PPPoE) and IEEE 802.1x. These technologies are used to authenticate a subscriber transport session. To enforce a subscriber's transport SLA at the transport layer in PPPoE environments, every packet associated with a subscriber's transport session can be identified with a PPPoE session ID that is specified as part of the PPPoE tunnel encapsulation. To associate packets with an 802.1x transport session, one either incorporates SLA enforcement functionality in the switching node attached to the physical link terminating the 802.1x session, or switches the packets from the physical link terminating the 802.1x session into a VLAN that is terminated on the node that enforces the SLA. In the absence of an authentication protocol, either VLAN tags or DHCP option 82 could be used to identify the DSL line that every packet is coming from or going to. These technologies could be used to enforce a subscriber transport SLA without the use of an explicit subscriber-authentication protocol.

Note PPPoE-based architectures could also use a VLAN tag as opposed to a PPPoE session ID to identify the traffic associated with a particular PPPoE.

SLA enforcement used for transport services relies on per-subscriber shaping or policing. The resulting QoS architecture relies on per-service-classification queueing and scheduling. SLA enforcement is typically implemented in the same node that terminates the transport session (PPPoE or 802.1x). Packets are classified per subscriber according to the transport session identifiers described above. The downstream traffic for each subscriber is typically shaped to a maximum rate based on the parameters of the transport SLA. If the transport SLA sold to the subscriber includes more than one class of service (gold, silver, or bronze), then additional classification, queuing, and scheduling are done to guarantee the transport parameters of the SLA associated with each class. For transport services, the node that terminates the transport session and enforces the subscriber SLA is typically the broadband remote-access server (BRAS). Figure 2-4 illustrates the per subscriber control and data plane functionality used by the network to implement a transport service.

Figure 2-4 Per-Subscriber Control and Data-Plane Functionality Used to Implement a Transport Service

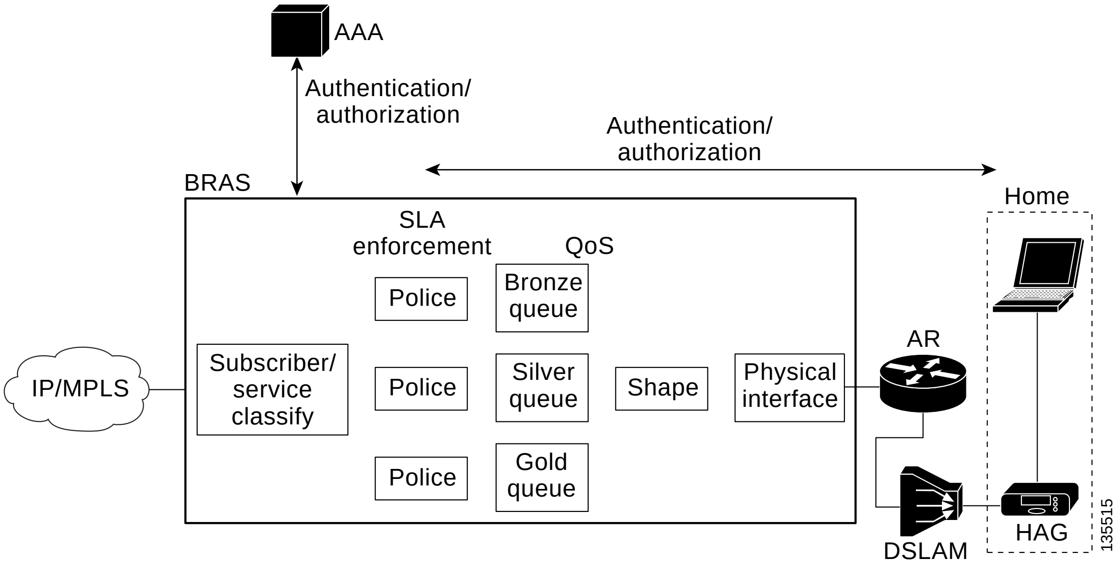

Managed Application-Based SLA

Subscriber authentication for an application service is implemented by means of application-aware components. For example, Electronic Program Guide, page 2-2, describes how subscriber authentication for a video service is typically implemented as part of the electronic program guide (EPG) function. The EPG, a component of video middleware, often authenticates a subscriber's video STB by means of an application-layer challenge such as an HTTP authentication challenge. If the EPG is not able to authenticate the STB, the subscriber cannot use the STB for broadcast video services. SLA enforcement for a managed application service is also performed by application-aware components. As an example, the number of simultaneous video streams that a subscriber may have active for a video application service is limited by a combination of (1) the number of authorized STBs the subscriber has in the home, and (2) the video session limits enforced by the video middleware.

Because SLA enforcement for a managed application service is performed by application-aware components, the QoS architecture required to support an application-aware service can be greatly simplified. Instead of having shapers and queues per subscriber, QoS architectures that use class-based classification and scheduling, such as the DiffServ architecture, can be used for QoS. Figure 2-5 on page 2-17 illustrates the application and transport architecture used to implement a video application service.

Figure 2-5 Application and Transport Architecture Used to Implement a Video Application Service

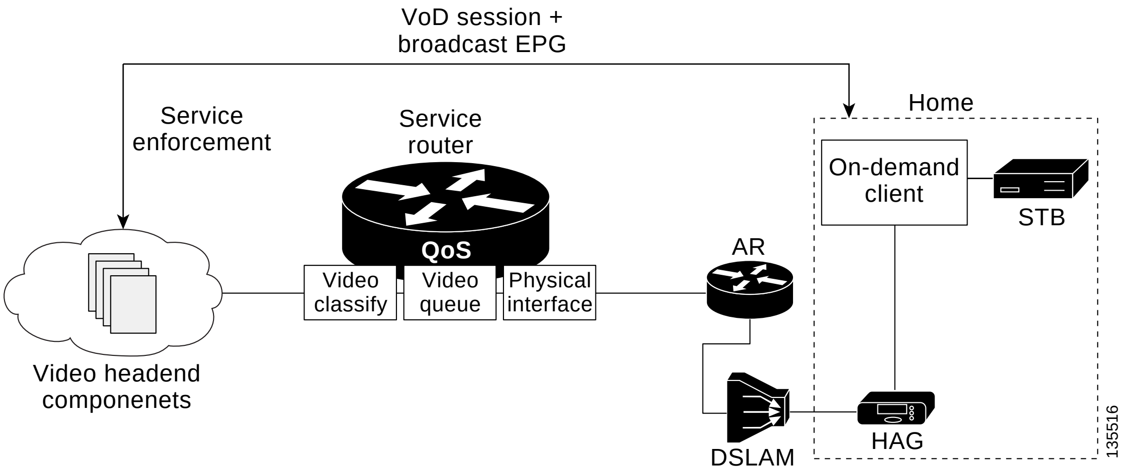

While an Internet access service is typically sold as a transport service, a video service may be sold to a subscriber as ether a transport or an application service. From the discussion above, the transport architecture needed for an application service is significantly different from that needed for an application service.

The Release 1.0 transport architecture is optimized for service providers that sell video as a managed application service, as opposed to a transport service, with the assumption that the access network devices provide congestion management or DiffServ QoS. In the solution transport network design, architectural tradeoffs have been made with this assumption in mind.

Note Alternatively, for networks where the access network devices do not provide the minimal congestion management described, distributed per-subscriber QoS control for all applications, including video, may be required. This transport SLA-like approach is not the design direction taken for the first release of the solution.

Service Separation in a Triple-Play Architecture

An important aspect of the transport network for a triple-play architecture is how much support the network provides in isolating each service.

Minimally, the network must provide the ability to meet the delay and drop requirements for each service when multiple services share the same physical link. This capability is inherent in the QoS architecture of the solution. (See QoS Architecture, page 3-46.)

In addition, the network may be configured to provide separate forwarding and routing domains for each service. This level of service separation is very useful when a service provider wants to manage separately the address space, topology, and IP infrastructure associated with each service. The following section explains why a service provider may want to have different transport attributes for different sets of services.

Forwarding Architectures

The transport architecture associated with different services may require the use of different encapsulations and therefore different types of packet forwarding. If one creates separate logical topologies for different services, these services can be forwarded by means of different forwarding techniques. The paragraphs below illustrate how the different transport architectures of Internet access and video services require that there be separate logical forwarding planes for the two service categories.

As explained in Potential Video Service Architectures, page 2-14, Internet access service is typically sold as a transport service. In a DSL environment, this typically results in a transport architecture that uses a PPPoE session from a CPE device to a BRAS that authenticates subscriber sessions and enforces the SLA associated with that session. Because PPPoE encapsulation requires an 802.3 header, PPPoE packets must be forwarded by means of Layer 2 switching between the PPPoE client (the CPE device) and the PPPoE server (the BRAS).

Also from Potential Video Service Architectures, page 2-14, the transport architecture for Release 1.0 assumes that the SLA for video services is an application SLA. Because authentication and enforcement are application services implemented in application components, there is no need to use a Layer 2 tunneling protocol such as PPPoE or a transport-layer authentication and enforcement component such as a BRAS for video services. Instead, video services can use IP encapsulation between the STB and the video infrastructure components described in Video Application Components, page 2-1. Since IP encapsulation is used, there is no need to forward packets between STBs and the video infrastructure components in the VHO using only Layer 2 switching. The solution transport architecture described in this document uses a combination of Layer 2 and Layer 3 forwarding for broadcast video and VoD services.

Note that the Internet access transport architecture described above requires that the access, aggregation, and distribution networks switch Internet access packets at Layer 2, while the video transport architecture allows these networks to switch video packets at either Layer 2 or Layer 3. To allow Layer 3 switching for video and Layer 2 switching for Internet access, the network must be configured into separate logical topologies that are switched by means of different encapsulations and packet switching functions (Layer 2 vs. Layer 3). The transport architecture separates these logical topologies in the aggregation and distribution networks by using different 802.1q VLANs for the different services.

Service-Availability and Bandwidth Requirements

Because different services have different service-availability and bandwidth requirements, a service provider could potentially reduce the cost of the network while maintaining the requirements for each service by creating separate logical topologies for different services.

As an example of different service availability requirements, Service Availability, page 2-13, describes the different availability and bandwidth requirements of broadcast video and VoD services. A service provider could optimize the network for both services by creating separate logical topologies for each service. These topologies could be created by using VRF-based technologies such as MPLS VPN or VRF-lite. [VRF stands for virtual private network (VPN) routing and forwarding, as well as a VRF instance.] In addition, the separate logical topologies could be created by populating the routing table with multiple instances of routing processes running on the different topologies and not exchanging routes between these processes. The differing availability requirements for broadcast video and VoD may lead to a transport requirement that the network must provide redundant paths for broadcast video but not for VoD. To meet this requirement cost-effectively, separate logical topologies can be created for the two services. The logical topology for broadcast video maps the address space associated with real-time encoders and STBs into a topology with redundant physical paths, while the address space associated with VoD servers and STBs maps into a VRF with nonredundant physical paths.

Note Release 1.0 test configurations did not include the use of VRF technologies to map services to different VRFs.

Organizational Structure

A service provider may have an organizational structure in which different services are managed by different organizations. The ability to map different services to different logical topologies allows each organization to manage and debug the transport as well as the IP infrastructure components separately.

IP Infrastructure Components

When different services are managed by different organizations within a service provider, it may be operationally simpler to have separate IP infrastructure components such as Dynamic Host Configuration Protocol (DHCP) servers for different services. Using different DHCP servers for different services allows the IP address spaces for these services to be managed separately. It also allows the DHCP servers to be configured separately for different services without having to use static configuration on the DHCP server to associate different CPE devices with different services.

Service Separation in the Release 1.0 Architecture

Because of the transport architecture issues described in Forwarding Architectures, page 2-18, it is likely that early IPTV/VoBB deployments will not use a unified transport architecture for all services. Because of this, Release 1.0 uses a service separation architecture in which traffic associated with each service is forwarded to or received from a separate logical access topology at the CPE device. This service-based logical topology separation is continued through the aggregation and distribution networks.

This transport architecture allows traffic associated with different services to be aggregated or terminated at different sites by means of different infrastructure components. This architecture allows traffic associated with Internet access services to be aggregated at a BRAS, while traffic associated with video services (specifically the managed video application service types) is terminated by means of the video infrastructure components described in IPTV/VoBB Transport Architecture and Issues, page 2-9.

Video Forwarding Architecture, page 3-11, describes how service separation is implemented in the aggregation and distribution networks in Release 1.0, while Edge Transport Architecture, page 3-39, describes how service separation is implemented at the CPE device and in the access network.