
NEXT-GENERATION TRIPLE-PLAY SERVICES

The IP networking world is buzzing with vendors releasing triple-play and multi-play ready products and platforms, service providers unleashing multimedia network extravaganzas on anticipating markets, and media pundits electrifying the buzz. In the beginning, triple-play referred to voice, video, and data services being delivered over an IP network. Multi-play extends this concept to refer to data services consisting of more than a single traffic type. These are premium data services. A gaming service, a high-speed download area, and a walled-garden network are examples of data services that can be treated differently to an Internet connection.

As explained in Chapter 1, “A History of Broadband Networks,” data services meant Internet access. Of course, there was access to corporate networks using narrowband services (X.25, Frame Relay, ATM, and dial-up), but by far the most common use of a data service in terms of numbers of connections is Internet access. Over the last several years, service providers have been offering voice over IP (VoIP) services. Such service providers include both fixed and nonfixed operators. The former covers operators enhancing their service offerings with a voice network more feature-rich than POTS. The latter includes VoIP providers that rely on customers providing their own connectivity over the Internet to access the provider network. Vonage and Skype are two examples.

Not until video over IP services were popularized did the concept of triple play really take off. Video over IP services include television channels transmitted to households over an IP network (IPTV) and Video on Demand (VoD), which allows users to request on-demand movies and TV shows streamed over their IP network connection. But video content is not restricted to these general categories. For example, YouTube and Network Private Video Recorders (NPVRs) are two examples of how the market dynamic is changing to encompass more than simple VoD and IPTV.

This chapter discusses in more detail what triple-play services are about and what subservices can be offered as part of the three broad topics under the triple-play umbrella. Also covered are technical aspects of service delivery. For example, a general technical overview of a video server architecture is provided for reference.

NETWORK TOPOLOGY

This section gives a high-level overview of a generalized network topology and any protocols and devices of interest. The major components are: server head-end, external Internet peers, IP core, Broadband Network Gateway (BNG) edge, Ethernet aggregation network, DSLAMs, local loop, and the customer household. Many of these components are explained in more detail later in this book, but Figure 2.1 shows a simple overview of the server infrastructure in relation to the core.

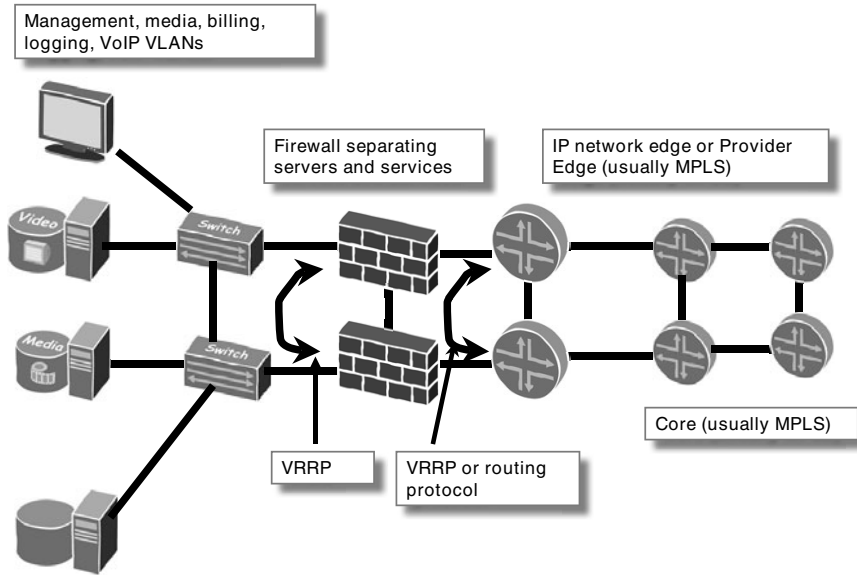


Figure 2.1 A high-level overview of servers, firewalls, IP edge, and core in a multiplay network.

One or two main data centers house the infrastructure servers, such as NMS, RADIUS, and log servers. Often colocated at these same data centers are servers for VoIP, IPTV, and VoD applications. These servers connect to switches using 100BASE-T and 1000BASE-T electrical connections, which are then fed into the core via optical GigabitEthernet or 10-GigabitEthernet links. To ensure that a failure of one IP router or switch does not affect connectivity between the core and server LANs, Virtual Router Redundancy Protocol (VRRP) provides a virtual IP and MAC address to the devices on the LAN. If one router fails, the other router assumes the role as the default gateway for the subnet.

If the network is running MPLS, edge routers are called Provider Edge (PE) routers, which are connected via point-to-point GigabitEthernet, 10-GigabitEthernet or SONET links to Provider (P) routers. Unicast and sometimes multicast IP packets are encapsulated into MPLS packets and are sent across the network of P routers toward a PE, at which point they are decapsulated back to IP packets. The core is a highly available, high-speed transport facility that switches MPLS packets to the destination label encoded in the packet header. To make the most of an

MPLS network, traffic engineering can be used to ensure that high-priority traffic is routed efficiently through the network using features such as fast reroute and link coloring. If packets are sent to a broadband subscriber, the PE should be the BNG, which performs service definition functions at the edge of the network.

Figure 2.2 shows the basic topology of the other side of the core network, including the BNG, historically called a Broadband Remote Access Server (BRAS). The BNG can be as simple as a router that forwards packets between the core and customer, providing little extra value than a high-density router. On the other hand, it can be a more complex router that implements dynamic per-subscriber IP policies, Quality of Service (QoS) profiles, rate limiters, packet manipulation, address assignment, session termination, and forwarding. The amount of complexity depends on the service provider’s requirements. At a minimum, the BNG performs some kind of session termination from subscribers, using either DHCP or PPP in tandem with an address pool or RADIUS.

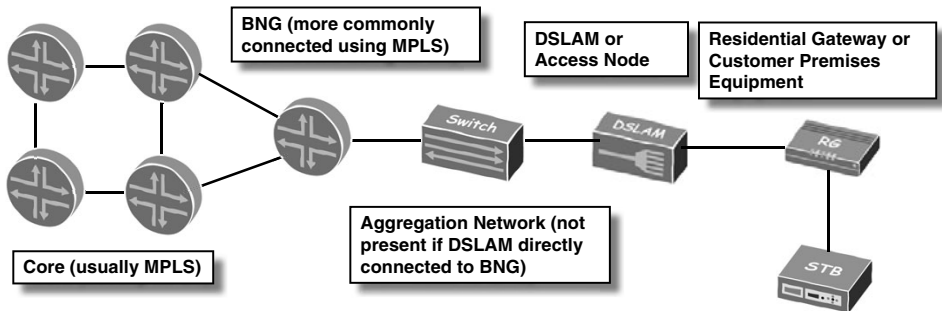


Figure 2.2 A view of the network from the core toward the subscriber.

The BNG connects to DSLAMs either using a switched aggregation network (also called a metro Ethernet network) or directly using dark fibre, CWDM, or DWDM. An aggregation network is useful when a switched layer between the DSLAM and the BNG provides cost-effective aggregation capacity. Such a scenario would arise if the bandwidth utilization per-DSL-port were not enough to justify connecting DSLAMs directly to the BNG. This benefit needs to be weighed against the expected traffic loads to and from DSLAMs. Also consider whether a switch can continue to offer enough statistical traffic multiplexing gains in the future.

DSLAMs terminate subscriber copper local loops and provide a DSL modulation service to the CPE. DSLAMs are also Layer 2-aware and take ATM AAL5 PDUs or Ethernet frames from the customer side and forward them to the BNG on the network side. If the DSLAM is Ethernet-capable, it forwards traffic toward the BNG as Ethernet frames with at least one level of VLAN tagging. Ethernet-capable DSLAMs are an important aspect of most of the designs in this book compared to the erstwhile DSLAMs, which use ATM as the interface toward the BNG. ATM-based DSLAMs also require the use of an ATM aggregation network, which is still common in many telcos. But the move is clearly toward Ethernet devices, which offer cost savings in the aggregation network and BNG due to the use of the cheaper Ethernet interfaces and the commoditized internal switching hardware.

The customer household is the network's demarcation point between service provider and customer. Services are delivered to the demarcation point, which may be a piece of equipment managed by the service provider. Or, in a simpler architecture, the customer may provide his or her own hardware. In both cases, they are called the Residential Gateway (RG), which is the Customer Premises Equipment (CPE). Most CPE are multi-play networks is complex pieces of equipment provided by the service provider that integrate a DSL modem and router. Sometimes only the DSL modem function is used; this means the CPE operates in bridging mode. In either case, it is useful for the service provider to manage the equipment for provisioning and remote diagnostics purposes, especially because the number of services a customer receives is more than just a basic Internet service. Connected to the CPE are one or more Set-Top Boxes (STBs), although most providers have only a single STB today. STBs take VoD and IPTV traffic coming from the network and send it to a TV connected via composite, SCART, HDMI, or RF outputs. STBs are controlled with a remote control, which is used to switch channels, browse videos, or, in some cases, send e-mail. Some STBs have a web browser built in to the STB User Interface (SUI). To let other devices connect to the broadband network, CPE have several Ethernet ports to connect PCs, gaming consoles, Macs, and other devices in the home. WiFi capability is commonly built into the CPE, which is useful for wireless phones with VoIP capability, PDAs, and any other nonwired device.

VIDEO OVER IP

As mentioned previously, video over IP is a broad term for a video service over an IP network. In this book, the focus of video over IP services is IPTV services delivered over multicast, and on-demand video services delivered using unicast. In general, the core and access networks can be built in a few different ways to support IPTV multicast. However, the multicast IPTV is a relatively straightforward service to categorize both at the network architecture level and at the service level. The service design is described in the following section. Microsoft has proposed a variation of the generalized multicast IPTV design, called MSTV. This is also discussed in the following section.

Unicast video services are more diverse. For networks with substantial network capacity, unicast video can in fact replace multicast IPTV services. This is called unicast IPTV. But to start with, the most identifiable service is basic VoD, similar to that found in hotels. Compared to hotel systems, not only is the range of content in multiplex VoD environments a lot wider than hotel systems, but the *concurrency rate* is a lot higher. Concurrency rate is the percentage of active video streams in the network versus the number of potential subscribers who can watch a video stream. These concepts and architecture are described in the section “Video on Demand.”

Some additional video services, such as Network PVR (NPVR) and specialized content agreements, are value-added services that can enhance a provider’s video value proposition. They are also discussed in “Video on Demand.”

To show how media content flows from a satellite or cable source to the network, that section shows an example of a video head-end, which is where most content is encrypted before being sent to subscribers. Encrypted content is usually channels that are received from the content distributor already protected. They must be decrypted and then re-encrypted, but with the provider’s certificate, before transmission to subscribers. This process is discussed in the section “Media Encoding, Security, and Encryption.”

To tie together all these services and platforms, the middleware acts as the control point. It handles video service provisioning, interfaces with most of the video elements, creates billing records, and interacts with customer STBs. Clearly this is an important piece of the network, and is described in the section “Middleware.”

IPTV

In most networks, IPTV services are generally characterized by efficient transport of television channels over a packet network using multicast. Multicast reduces utilization of network links between the video server and the customer by sending only a single copy of a media stream into the network. The network replicates the stream to individual subscribers closer to the edge, thus saving bandwidth in the core. The components involved before the core are the source video feed receiver and decoder, IP encoder, and encryption process. The media streams are then fed to the IP core. This process is described in the next section.

As mentioned previously, IPTV is not restricted to content via multicast. Networks with ample bandwidth to spare can get more control over the customer video stream if it is unicast. For example, it is easier to target advertising based on viewing habits to specific households if there is a uniquely addressable endpoint, that is, a unicast destination address. This is covered in the section “Unicast IPTV.” Another example of a service that makes use of unicast streams is Pause Live TV. One implementation by Bitband makes use of cache servers deployed throughout the network. If the user pauses an IPTV stream, the STB switches from the multicast stream to a unicast stream.

Video Head-End

Figure 2.3 shows a simple video head-end reference architecture that is used for the discussion in this section.

The content in Figure 2.3 is supplied by satellite, but there are a variety of transport mechanisms. Terrestrial microwave links or even IP feeds from content aggregators are other possibilities. Digital streams are received off the satellite using the Digital Video Broadcast, Satellite (DVB-S) standard. These streams are received by the satellite, are frequency-stepped down, and are fed to a receiver that handles the demodulation of the DVB-S signal. They are then fed to a suitable output format for the decoder to work with. If the channel is encrypted, it is the decoder’s job to decrypt the channel and forward it to the IP encoder. If the channel is not encrypted, it acts as a pass-through system and sends the stream directly to the encoder.




Figure 2.3 A simple head-end architecture fed by satellite.

The DVB-S and DVB-S2 systems transport the media streams in MPEG format, so the decoder’s job is to take the unencrypted MPEG stream and feed it to the IP encoder. The IP encoder does the hard work of taking the MPEG stream, trans-rating or trans-coding it to a lower bit rate. Whether it is trans-rated or trans-coded depends on the amount of bit-rate reduction required. The MPEG video feed is encapsulated into IP packets. Any streams that were originally encrypted after being received from the satellite need to be re-encrypted before being made available to subscribers. Such streams are sent to an encryption engine, which encrypts the data using the provider’s key or certificate. These streams can only be decrypted using an STB, which has received authorization from the certificate server. The certificate or key may have come via the middleware.

The encryption servers might be the last devices before the encrypted and unencrypted multicast streams are sent to the IP core. The types of devices (if any) that are between the encrypters and the IP network usually depend on the high-availability architecture. Dedicated servers with a heartbeat mechanism between

the two can provide a single multicast source to the network so that when one server fails, the other can resume the multicast forwarding duties. If the high availability (HA) mechanism is simple enough, the encrypters can also perform this feature.

This design uses a single IP source address for multicast group traffic. Each of the groups is distributed throughout the network using normal multicast forwarding replication, and the clients join the groups as needed.

Microsoft's MSTV

An alternative proposition by Microsoft is to extend the multicast distribution model with a hybrid unicast and multicast solution. One set of multicast *A-servers* (acquisition servers) sends the multicast groups into the network, similar to the architecture just described. In addition, a distributed layer of unicast *D-servers* (distribution servers) is located strategically throughout the network. The deployment architecture and location of the D-servers are intended to reduce the *channel zapping* latency for STBs.

From an architectural perspective, the primary use of D-servers is to compensate for the time interval between MPEG I-frames. An I-frame occurs every so often in the MPEG data flow and is a marker in the stream that can be used to build an instantaneous image on the screen. An MPEG decoder can only start decoding a stream starting with an I-frame. If an I-frame occurs more frequently in the stream, a receiver (such as an STB) can more quickly lock onto the stream and start the video output. But if the I-frame occurs too frequently in the stream, the video bit rate needs to increase, which is extra overhead. When an STB joins a new channel, a D server sends a unicast copy of the channel, beginning with an I-frame, at a much greater bandwidth (a burst in other words) than the multicast stream, to the STB. Simultaneously, the STB joins the multicast group in the network using IGMP. While the unicast stream is being used to display the channel to the television, the multicast group starts filling a video buffer. As soon as the buffer is sufficiently full, the unicast stream is terminated, and the STB switches to the multicast stream. Figure 2.4 shows a high-level architecture of the solution.

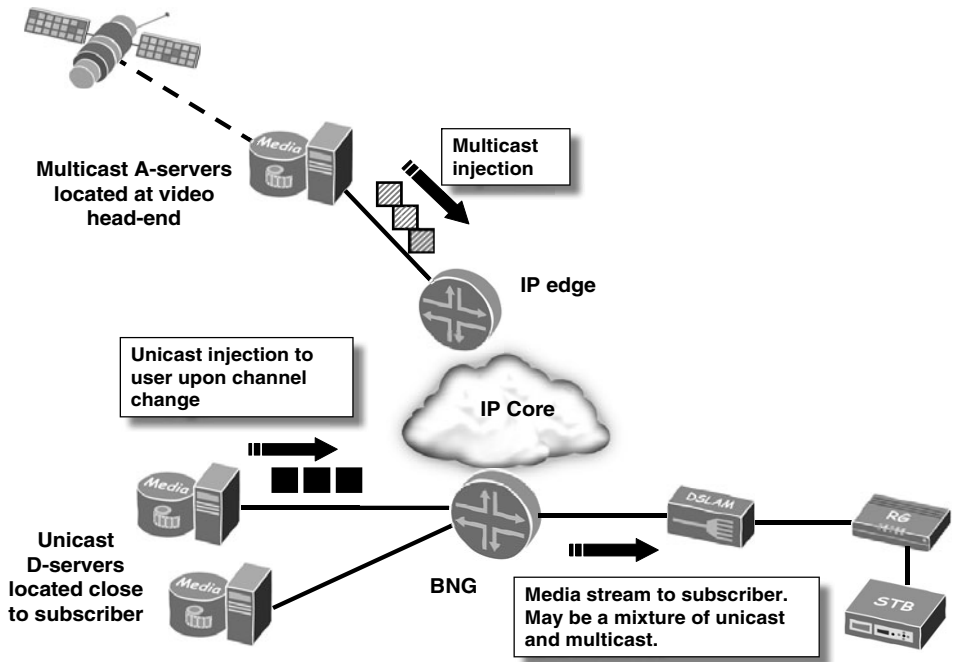


Figure 2.4 An MSTV architecture.

The solution requires servers to be distributed as close to subscribers as possible. This reduces the latency of the unicast stream between the STB and server. Also, the number of simultaneous connections that can be made between STBs and a server is limited by the server’s video processing capacity and interface throughput. Even though the STBs are connected to a server for only several seconds, the network needs to cope with many simultaneous channel changes during peak viewing times. The efficacy of this solution depends on the individual network. As with any architecture, it is advised that you test its scaling properties well in a lab environment beforehand.

Unicast IPTV

MSTV is a hybrid unicast and multicast model. A pure unicast model drastically increases the requirements of network transmission, forwarding capacity, and server cluster performance. Each server needs to generate a unique stream for each client, and at 4Mbps for a Standard Definition TV (SD-TV) channel, the requirements quickly add up. Of course, this offers the highest flexibility for the

service provider. A profile of each household's viewing habits can be built and targeted advertising created during the commercial breaks. Also, there are no issues with channel zapping latency due to I-frame interval and subsequent buffer fill. There are additional advantages for the service provider with a unicast stream. For instance, a service for rewinding, pausing, and fast-forwarding a program needs a unicast data stream. Multicast replication can be the standard delivery method until the user requests a pause, rewind, or fast-forward of the channel.

VIDEO ON DEMAND

Video on Demand as a concept is relatively well-known. Anyone who has been to a modern hotel has discovered VoD services available from the television in their room. There are two common models. The first is where popular movies run on a rotating schedule and, for a set period of time (24 hours, for example), you can watch the movie as many times as it comes up on the rotating schedule. This is not terribly flexible if you want to watch the movie right away. It is also difficult to implement video controls, such as pause and fast-forward, if the media stream does not have some kind of feedback loop to the server. This system is called Near Video on Demand (NVoD). The more preferred approach (at least from an end-user perspective) is a completely on-demand system in which the video stream is delivered to the user when requested. Usefully, a full on-demand approach is the way most systems are deployed nowadays; only a few older systems stream content on a loop system.

This section explains how VoD solutions can be deployed in a network service provider environment. At the press of a button, customers can watch the latest blockbuster movies or browse through hundreds of older cinematic releases from the comfort of their own home. One important aspect of a VoD over broadband service is, not surprisingly, a market differentiator over what traditional bricks-and-mortar movie shops offer. Breaking people's years-long habit of going to the local movie store and picking up a tangible, shiny disc needs more than just a strong marketing campaign. On-demand video over broadband has enormous potential over traditional media, but from the very beginning, clear advantages need to be pushed, such as previews of each movie or show, a large content archive, same-as or longer-than rental periods as DVDs and videos, and, most

importantly, a compelling price. These less-technical, more-service-oriented issues are covered later in this chapter.

Basic Video on Demand

Except in the most grandiose of VoD rollout plans, a server architecture normally starts out small, with a centralized video head-end. One or more GigabitEthernet links (either bonded or discrete paths) link the VoD server switches with the IP edge. For maximum flexibility and to reduce IP address and load-balancing complexity (but at the expense of interface cost), 10-GigabitEthernet links can be deployed from the outset.

There are varying models of how the servers interact with the IP edge and STBs. One example, proposed by BitBand, is to have one or more *streaming gateways* in addition to the video content servers. When an STB requests a media file from the middleware platform (described later in this section), the STB is sent a control file that lists the video servers to which it can connect. Each VoD server that is listed in the control file is ranked by preference so that the client can choose the best server. The preference metric is an inverse indicator of load, so the more preferred server is the least-loaded one. For redundancy reasons, the streaming gateway doesn't send the details of just the best video server to the client. If a server fails midstream, the client can reconnect to a different server automatically without having to reconnect to the streaming gateway.

Various high-availability models exist. These can cover various failure scenarios—server failure, switch failure, router failure, or a combination thereof. Because explaining these scenarios in low-level detail could consume several chapters, Figure 2.5 shows a sample architecture with the components involved.

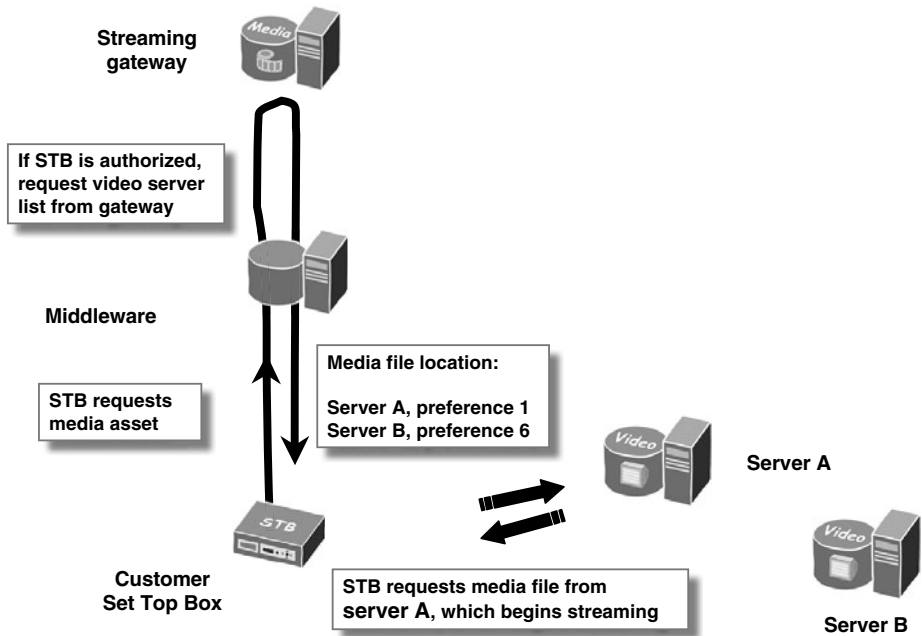


Figure 2.5 Video streaming setup.

Figure 2.6 shows a client re-requesting a media file to a different server after component failure. Less-impacting failures, such as a switch or router going down, cause a blip in the service until Layer 2 and Layer 3 protocols converge around the failure. Having a hot-standby server is difficult to achieve due to the amount of video session state that would have to be mirrored between the two servers.

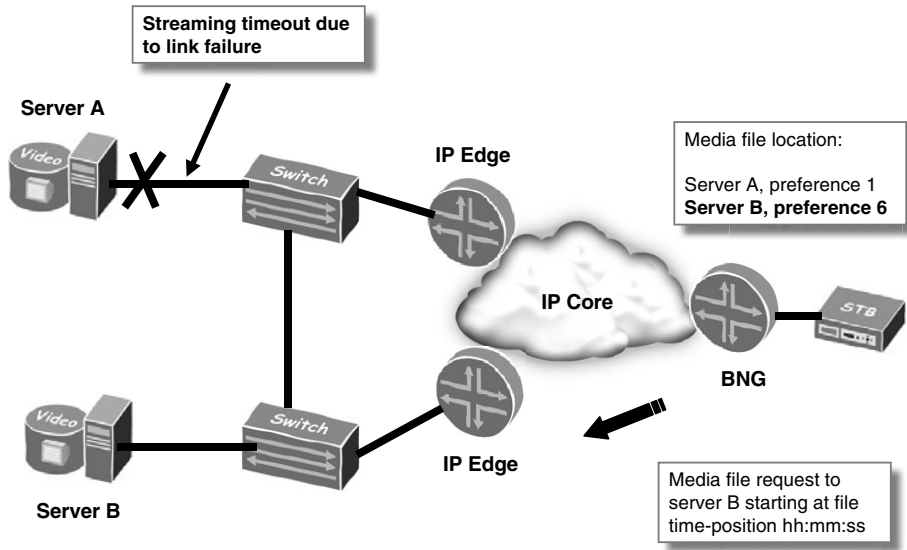


Figure 2.6 JSTB re-requests a file from the secondary server.

Distributed Video Clusters

To handle hundreds of concurrent video streams, the VoD server network needs to become more distributed. This has several benefits:

- It reduces latency between the server and the STB (although low latency is more important with an IPTV service).
- It creates Point of Presence (PoP) diversity to reduce the impact of a PoP going off the network.
- It improves bandwidth utilization in the core by moving VoD flows to the network edges.
- It reduces per-server bandwidth and CPU load.

When a new video PoP is built, it uses most of the components of the centralized video head-end, except on a smaller scale. Typically it has a central data center where all the TV shows and movies are stored. This data center also has certificate and encryption servers, middleware components, SUI content servers, and,

of course, some video streaming servers. If the servers do not share a central content store, popular content can be preseeded to the individual servers based on anticipated demand ahead of a popular movie release. A second option is to simply let viewer demand dictate how content is seeded to the servers. A third option is a mix of the two; this option is the most common.

Tier 2 video PoPs are geared more toward pure video content delivery rather than housing middleware platforms and streaming gateways. Instead, these lower-bandwidth services are centralized at the Tier 1 data center. In much the same way that the content can be replicated to the Tier 1 streaming servers from the content store, the Tier 2 streaming servers (or the content store in the PoP) are seeded with the media files beforehand. The most optimal server from which an STB should request a media file can be determined dynamically using heuristics such as lowest server round-trip time (RTT) or a trace route method. Or the intelligence can be shifted to the network, where the STB's physical location can be derived using methods such as IP address coupled with a DHCP option 82 stamp or a PPPoE Intermediate Agent (IA) string. DHCP option 82 and PPPoE IA are identifiers that the DSLAM can add to indicate from which physical port a subscriber DHCP-based IPoE session or PPPoE session comes. Then, an IP address can be linked to one of the two identifiers, depending on the access protocol used. The STB can then be directed at a subset of video servers from which it can choose the one with the lowest preference. Of course, if a requested media file is not available at a Tier 2 PoP closest to the STB, the most effective option is to direct the STB to a central location where the media asset is available. This is only one example of how content and server distribution can be deployed.

► **Caution:** One of the important aspects to consider in the early phases of the design is how the high availability of a particular architecture should work from an end-to-end perspective. If this is treated as an afterthought or left until late in the design stage, it can mean a lot of unnecessary redesign work.

Maintaining Video Quality with Call Admission Control

The principle of Call Admission Control (CAC) in this sense is much the same as it was in the PSTN. A session (a video stream) is allowed or disallowed to proceed based on a CAC algorithm. It is essentially a binary result. If a CAC request is denied, it is not possible for a session to proceed, even in a degraded manner.

This is different from a QoS mechanism on routers and switches, which can prioritize traffic based on embedded markings in packet or frame headers. If traffic cannot be transmitted at a point in time, it is queued for later transmission. Such a prioritization mechanism relies on IP traffic streams being able to be statistically multiplexed down a single path. Normally this is possible because often non-VoIP and non-video IP streams are not of a consistent bit rate and can tolerate some variable delay in the packet delivery process.

Most video over IP streams are not so forgiving of poor network conditions. These conditions occur for a variety of reasons, such as link failures that cause suboptimal traffic routing, poor network capacity planning, or a DoS attack. The end result is too much data offered to network links with insufficient capacity to transfer all required packets.

As a general principle, the network needs to be built in such a way that video traffic is high priority and is forwarded without delay. If it isn't possible to reliably send a video or VoIP call from end to end (such as due to lack of network resources), the session should not proceed. This requires the CAC mechanism to work during session setup. In other words, first a request is made to see if the user has the right credentials (he has paid his bill and has passed the authentication check, for example). Then the network links are checked for resource availability. Only then is the stream delivered to the client. Relying on a mechanism that takes periodic snapshots or averages of network utilization is not a reliable method, because the traffic constitution could easily change between polling periods.

Implementing a CAC function with a VoD network can work by extending the video server's functionality to interact with an external server. The server can model the network topology and have an end-to-end (or part thereof), real-time view of VoD traffic flows. Because the video streaming server is the last point of contact for an STB before the video stream is sent, it makes the most sense to have this device interface with the CAC engine.

This is a general overview of CAC. For more in-depth information on VoD and multicast IPTV CAC, see Chapters 8 and 11.

MEDIA ENCODING, SECURITY, AND ENCRYPTION

A mandatory feature of all video over IP networks is content security. Digital Rights Management (DRM) is a hot topic for anyone going into the media distribution business. Music and videos purchased from sites such as iTunes Music Store and Amazon.com have licenses restricting how the file can be used. The license is enforced through the player, which is also needed to decrypt the file. The element of trust ends at the player, so one principle is that as long as the player's integrity is intact, the content can be sent safely to the client. Of course, crackers always keep up with the efforts of the content distributors and encryption engines in trying to break encryption schemes, so some systems that stream content to the desktop incorporate some extra host checking software to detect devices or software that is trying to circumvent DRM.

Using a dedicated STB instead of a software player makes it harder for crackers to bypass DRM mechanisms, because the host environment is now unfamiliar to the cracker. There are other benefits of streaming video to a dedicated STB, which are not directly related to DRM and security, such as a consistent operating environment controlled by the service provider. This makes it easier to troubleshoot and maintain compared to a user's PC.

Before we get into too many specifics, an explanation of a broader security perspective is needed. There are many types of encryption systems on the market that use Certificate Authentication (CA), which authenticates the identity of the client. DRM with content protection encrypts the content and controls when the media file can be played. Figure 2.7 shows one example of a provider architecture with a centralized video distribution system, certificate authentication, and content encryption.

After encoding, the middleware system analyzes the media files, or assets, to get their correct bit rate, title, runtime, and any other file details. The asset is then encrypted using the provider's certificate. The servers used to encrypt the asset can be the same as those used for IPTV encryption. After the asset is added to the content store and entered into the middleware system, it is available for delivery by the video streaming servers.

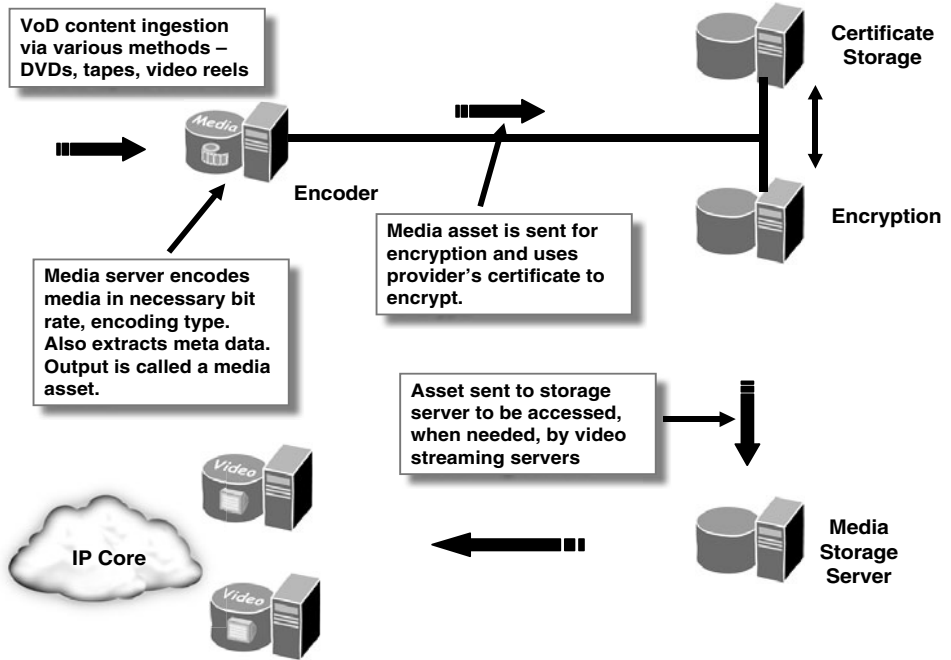


Figure 2.7 VoD asset encryption and delivery process.

Expanding on the example from earlier in this chapter that incorporates the Bit-Band video streaming servers, when an STB requests an asset, it contacts the middleware platform. The STB's identity is checked. It could be an IP address, an internal STB identifier, or a physical location gleaned earlier in the initialization process using DHCP option 82 or a PPPoE IA identifier. The identity is then used to check if the client has the right authorization for the VoD session. The authorization could be that the client's billing profile allows VoD streams, or that the requested movie fulfils a viewer age restriction, for example. One way of implementing parental control is for each person in the house to have his or her own PIN, which is entered when requesting a movie. Each PIN can be set up so that it can access all or only a subset of the content. TV shows and movies need to be rated as suitable for under-18s, for example. If the STB is allowed to receive the asset, it is sent a call entitlement, which includes information on how to decrypt the media file when it is streamed, how long the asset is valid (similar to a rental

period), and other control information. In most cases, the media file is streamed over UDP—ideally, with RTSP sequencing information.

MIDDLEWARE

The middleware platform can be considered the heart of the IPTV and VoD system. It controls the SUI, which is the interface that the customers use to navigate through the menus on their STB. The SUI can be customized by the service provider and normally is delivered using standard protocols, such as HTTP and SSL. Many unencrypted elements in user interface pages can be cached in an HTTP proxy to allow the SUI subsystem to scale well. The Electronic Programming Guide (EPG) is a list of current and future shows on an IPTV system. This data can be sent either as a multicast stream that all STBs can join or as a file that is periodically retrieved from an HTTP proxy via unicast.

The middleware also interfaces with the DRM platform, which stores the site certificates and encrypts content for distribution either via an on-demand system or through multicast IPTV. The actual encrypted unicast video content is stored for distribution by the video streaming servers. The middleware typically stores only metadata related to the media assets.

When a user purchases a media asset, the middleware also issues billing records to the service provider billing engine. One name for these records is Entitlement Data Records (EDRs). An EDR can be generated once after the asset has been authorized for delivery, regardless of how many times the user watches the asset within his or her rental period.

VIDEO SERVICES

Technically speaking, a full video solution is quite complex. All this complexity can be in vain unless a clear service framework drives the network functionality. Many technically superior products unfortunately have been relegated to footnotes in the history books because of poor marketing and vision. Although there is no substitute for a well-thought-out marketing plan, this section lists some of the services possible with a video over IP network.

Network Private Video Recorder

A Network Private Video Recorder (NPVR) lets the user select a show, which is normally streamed as part of an IPTV service, to be recorded by the network. Then it is sent at a time of the user's choosing. This works in much the same way that people have been recording television stations on their TiVo or VCR for years. An NPVR service offers additional functionality not possible with a device in the home that does all the storing and recording. A recorder in the home can record only one or two channels at once, but if the service is more network-centric, there is no limit to the number of channels a user can record at the same time.

The technical implementation of the service need not be reflected in the service to the customer. For example, a provider can record just a single copy of each of the most popular channels, which can be enumerated to users as required. This saves storage space and lets someone later request a TV show or movie that she may have forgotten to record. Any shows that are not caught under this policy—a viewer wants to record a foreign movie, for instance—can be specially recorded as long as the request comes in ahead of the program.

Targeted Advertising

Being able to tailor advertising to the home is a huge potential for marketers. And they are bound to pay extra for this privilege. In much the same way that Internet search engines return advertising related to what the user typed in his search query, targeted advertising can deliver commercials based on a profile of the channels and what time they are viewed. To work alongside multicast IPTV services, additional groups could be dedicated to semi-targeted advertising. One multicast group could be for a demographic of 18-to-24-year-olds. When the commercial break starts, the STB could switch from the general channel group to a group used for that particular demographic.

Extensive Media Catalog

One of the competitive advantages that online stores have over bricks-and-mortar stores is that they do not need large shops to display their wares. Amazon.com, for example, can keep less-popular titles in distribution warehouses, and only when the media are ordered are they sent from the inventory. This is called long-tail marketing. Amazon.com's *Booksurge* concept takes this

model one step further. If a book is out of print, no problem. An electronic copy of the book can be dispatched to a printing center, printed, bound, and sent to the customer as if it had been printed and sent directly from the publisher.

This is similar to the advantage that a service provider has with VoD over traditional bricks-and-mortar rental stores. Adding a few hundred movies to a program lineup involves adding an extra media storage array and is independent of how many users have *checked out* the movie. Physical space is not much of a concern.

An interesting value-add that IPTV services have over satellite TV providers is that the number of different channels that are offered does not increase network traffic load in a linear fashion—the usage is much better. The benefit comes from being able to have many niche channels added to the multicast group list. This increases the appeal of the television packages to a wider audience. And only when a customer wants to watch that channel is it added to the multicast tree. By contrast, a satellite provider consumes valuable bandwidth on its transponder for each channel it adds. The efficiency gains for a satellite provider occur when many households are connected to its service and only a few of the most popular channels are needed to fill its transponder capacity. For IPTV to have such a positive impact, the market dynamics need to be right. An ethnically diverse viewer base helps fuel the need for equally diverse content.

DATA SERVICES

Data services encompass anything that is neither video over IP nor VoIP—which is, in fact, a large basket of services. Where the triple-play concept has been expanded to cover multi-play is in data services. The data service began with Internet access and is the simplest and most understood service by service providers and customers alike. It is still the most common reason people get a broadband connection. However, Internet access is becoming commoditized through price; it is a war of attrition. Faster access through new technologies helps slow the commoditization process, but it is services and packages that differentiate one provider from the next. This section describes services that a provider can offer in addition to basic Internet access.

PREMIUM GAMING

Figure 2.8 shows a simple example of a multiplay environment. Individually, these services are not big technological advances but are branded as part of a service package. Using an end-to-end QoS and policy framework, a provider can differentiate itself from the competition. One of these services is premium gaming.

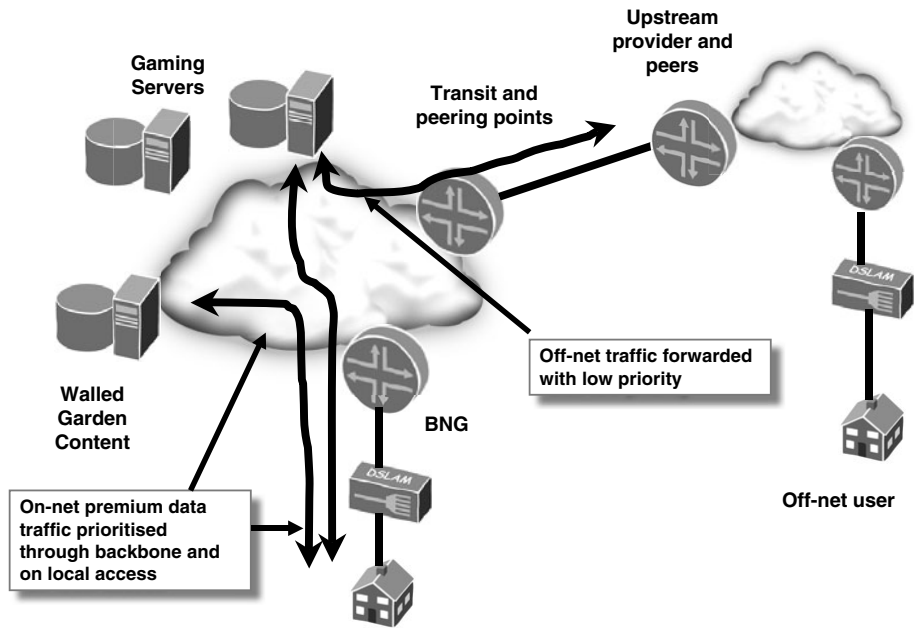


Figure 2.8 Prioritized gaming service.

To entice subscribers to a provider’s network, locally hosted gaming servers are a popular draw card because of the enticement of lower ping times to the servers. Whether this is enough of an enticement for many subscribers depends on a variety of factors, not just the latency through the network to the servers. The branding image among the gaming community of a provider running hosted servers is a strong one. And so is network latency. If a gamer can shave a few tens of milliseconds off his ping by being with a provider, many enthusiasts make the move just for this reason. An even more attractive proposition for service providers and customers is to ensure that gaming traffic is treated with high priority when there is a real chance of network traffic congestion. Traffic congestion causes latency increases and, worse, jitter. The latency is surely a negative factor,

but jitter has a worse impact on online games, which need fast, consistent feedback of players' movements to the servers. A well-engineered core should not be congested. The aggregation network and customer local loop are most susceptible to congestion. A customer doing heavy downloads at the time of online game-playing can disrupt her gaming session.

A solution is to mark all IP traffic going from gaming servers to on-net subscribers with a DiffServ marking that gives this traffic higher priority over less important best-effort traffic. Internet traffic can be marked with a DiffServ value of 0 and on-net gaming traffic as slightly higher than this—DiffServ AF11, for example.

Such a scheme could be applied to the network in a variety of ways. First, all packets coming from gaming servers and going to on-net subscribers can be marked with a nonzero DiffServ value. On the BNG, dynamic policies can be applied to subscribers who have purchased the gaming service, to modify their QoS policies, which would match the previously marked packets and transmit them ahead of Internet packets. Of course, this means that such data would be prioritized before reaching the BNG too, regardless of whether the destination is a paying premium customer. If this is viewed as a significant problem, a network-wide, BGP-based QoS signaling mechanism could be used. When a subscriber connects to the network, his /32 host route can be announced to the network with a special BGP community ID. When all BGP-speaking routers receive prefixes matching the special community, they can update policies in their forwarding tables, which gives traffic going to these destinations a slightly higher priority than best-effort. The scaling properties of this approach are poor. For all customers who subscribe to this service, there is a /32 route in the network, which consumes extra router control- and forwarding-plane memory. Despite there being services where subscribers can have a static address, which can trigger a /32 route in the provider network, these do not affect any QoS or filtering policy in router's forwarding databases.

A more scalable solution is to have dedicated address pools for subscribers who have a premium gaming service. These pools can be announced as aggregated networks from BNGs with BGP communities in a similar, but more scalable way than the per-subscriber /32 method previously mentioned.

WALLED-GARDEN SERVICES

Walled-garden services have been on the market for many years. CompuServ and AOL are two shining examples of walled-garden environments. Throughout the mid-1980s and early 1990s, CompuServ “was one of the largest information and networking services companies in existence.”¹ Its model provided content that was available only if you connected via the CompuServ network. At the time, the World Wide Web (WWW) was not yet conceived, so this was an ideal way to access graphical content over the CompuServe network using CompuServ’s Win-CIM client. But as the WWW gained popularity, providers who were hosting sites on CompuServ’s network shifted their sites to the Internet to have access to a wider audience. Therefore, the network’s popularity dropped sharply. AOL’s walled-garden content was its primary focus until the popularity of the Internet started pushing AOL to another model.

Other examples of popular walled gardens still exist. Cellular handsets accessing the Internet often start inside a walled garden. Here, the service provider sends a portal to the customer. News aggregation, content subscription control pages, and e-mail portals are a few examples. If users want to access the Internet, the data volume tariff might be higher (or not free) to access off-net or non-walled-garden content, but the Internet is still possible. Sirius and XM satellite radio are walled gardens, because users need a special handset to receive the broadcasts, and some of the content is exclusive to these providers.

Some lessons can be learned from walled-garden content. Even if the access method is unique—that is, if users need a special device to access the content—a walled garden can still be popular. To access satellite radio, users need a special receiver. Nonetheless, users are prepared to buy hardware for the additional benefits of satellite radio—a large coverage area, a digital signal, no retuning to a different frequency when moving between areas, and content. For Internet access on cellular handsets, walled gardens serve a useful purpose to reduce the number of pages that are needed to navigate to commonly accessed pages.

It could be tempting to have a special area with premium content on a service provider’s network for more traditional devices—PCs, Macs, gaming consoles. This has to be carefully thought through, because the danger is to restrict a user’s

¹ Wikipedia page on CompuServ: <http://en.wikipedia.org/wiki/Compuserve>

access to the wider Internet, as happened with CompuServ during the late '90s. Because the Internet is currently the network with the most people online, users want Internet access to chat with their friends and visit their home pages, blogs, and discussion forums. A simple model is to have a section of on-net content that is accessible with a special tariff and to have the Internet accessible by a more expensive tariff. If the provider offers flat rate (it does not bill according to volume), on-net content could be accessible with a faster speed than Internet content. One good example of having content available to off-net customers is gaming servers. If gaming servers are not available to off-net players, there might not be enough people continually active on the servers to make the service popular enough—a critical-mass problem.

One alternative is if a service provider is willing to pay the necessary license fees to a content provider so that it can rehost or rebrand the media and sell it as part of a package. One example of such a service discussed throughout this book is IPTV. But there needs to be a good reason that a user can't go to the site with his or her web browser and access the same content.

BUSINESS CONNECTIVITY

DSL lines are no longer just for residential access. Business customers are connecting to the Internet and their other offices using DSL access. Frame-Relay, ATM, and ISDN circuits cost a lot more to run than DSL-based ones, and DSL is proving to be a suitable alternative to these older modulations and encapsulations. Broadly speaking, three types of business services in the marketplace use DSL access: Layer 3 VPNs, Layer 2 VPNs, and enhanced Internet access.

Business Layer 3 VPNs

This business service is a provider-hosted Layer 3 VPN service and uses a DSL line as the access technology. The PVC or VLAN from the DSLAM is terminated on an edge router, providing connectivity to the business VPN cloud. Various encapsulations could be used on the local access. Also, the edge router could be different from that used for the residential BNG. The simplest approach is to reuse the residential broadband infrastructure and use the same PPPoE or DHCP-based IP over Ethernet (IPoE) addressing and encapsulation between the CPE and BNG. The BNG has Virtual Routing and Forwarding (VRF) configured

for each Layer 3 VPN. Reusing PPPoE encapsulation makes it easier to dynamically provision a customer session to the correct VRF by using RADIUS to map a PPP username (such as username@realm) to a VRF. If the local loop addressing and encapsulation are DHCP and Ethernet, respectively, a more static approach is needed to map a customer VLAN (Ethernet backhaul) or PVC (ATM backhaul) to a VRF. For more information on various backhaul architectures, see Chapter 4, “Designing a Triple-Play Access Network.” Figure 2.9 shows how the same access and aggregation network that is used for residential session transport can also be used for business VPN access. The remote user on the bottom logs in using PPPoE and the username sam@acewidgets.com and is placed in the business Layer 3 VPN. The Internet user, bob@isp.com, is placed in the global routing table.

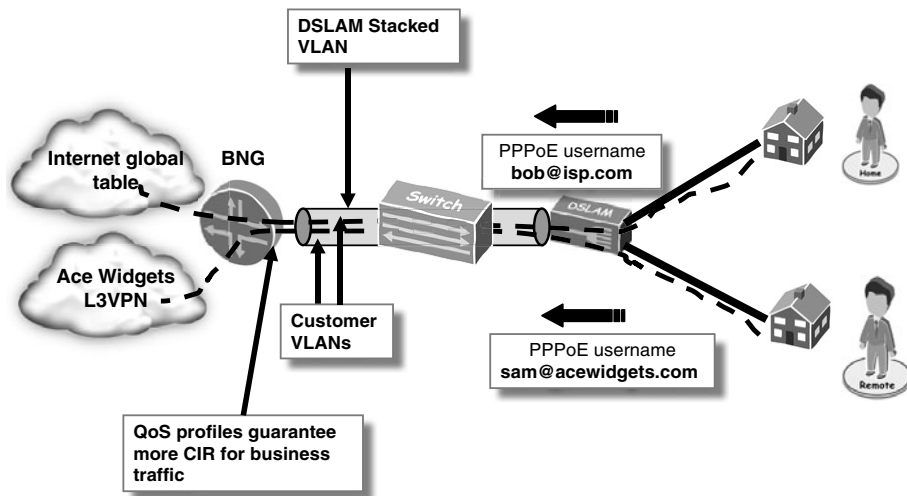


Figure 2.9 Business Layer 3 VPN access using residential broadband infrastructure.

Despite being a residential broadband infrastructure, many of the same QoS and SLA guarantees needed for business can still be achieved in the model shown in the figure. For instance, one of the most common requirements in a business Layer 3 VPN connection is to ensure a minimum bit rate from the IP edge to the CPE. This is known by various terms, such as Committed Information Rate (CIR), assured rate (Juniper JUNOSE terminology), bandwidth (Cisco IOS terminology), and transmit rate (Juniper JUNOS terminology). The circuit’s peak rate may be much higher, but this is not guaranteed throughput if the same links

that the circuit traverses are being heavily utilized by other customers. If the aggregation network between the DSLAM and BNG is oversubscribed, CIR rates can be managed on the BNG using QoS mechanisms. This makes the BNG *the* point of control to ensure that business connections receive their committed rates in the event of circuit or interface congestion. A cruder alternative to managing oversubscription in the aggregation network is to set the Ethernet priority in the Ethernet header and let switches and DSLAMs handle any congestion with simple scheduling. The disadvantage is that without some very complex traffic modeling, it is hard to ensure that an accurate assured rate is delivered to the end subscriber.

However, it is common for providers to opt for dedicated edge routers for business connections. One of the reasons is that bigger providers often have one group manage VPN services and another manage residential broadband. Using separate hardware helps streamline operations, design, and testing. Even when one group manages everything, there might be a purely architectural reason to have different routers. It might be for stability reasons, to reflect the importance of business connections by using dedicated hardware, or additional routing policy features available on a different platform. Because usually there are less business connections than residential, there are less changes occurring on a dedicated business router. These changes and *fingers in the network* are things that can also cause service outages. Figure 2.10 shows how a customer VLAN can be dedicated to a DSL port and is switched through a metro Ethernet network.

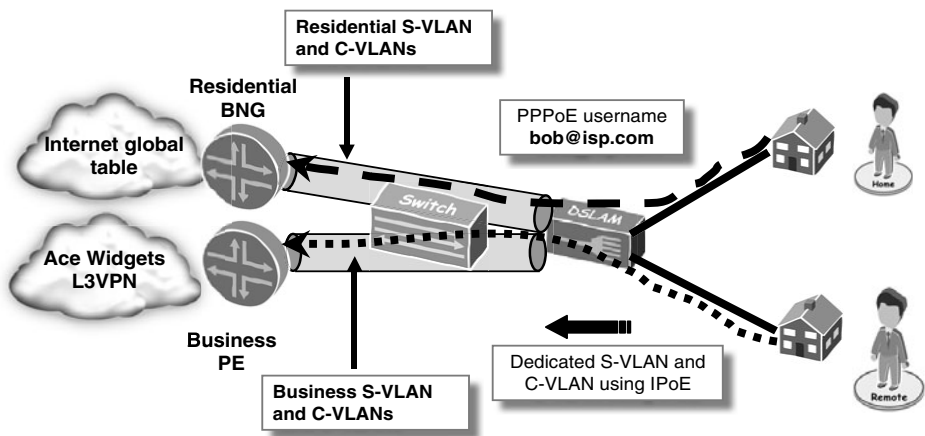


Figure 2.10 Business Layer 3 VPN access using dedicated circuits.

In the figure, a stacked VLAN (S-VLAN) tag identifies the DSLAM and a customer VLAN tag (C-VLAN) for the DSL port. There are several ways of switching the customer through the metro Ethernet network. The architecture shown in Figure 2.10 is the most recommended for several reasons. An outer S-VLAN tag for the DSLAM lets the architecture scale well in a large metro Ethernet network with many DSLAMs. A dedicated C-VLAN tag for the DSL port as opposed to a service-based VLAN for all business connections attached to a DSLAM ensures Layer 2 separation between customers. This provides a good level of security between DSL ports. It also suits more traditional routers that enforce QoS on a per-VLAN basis rather than per-subscriber. In the latter scenario, this would be the case with a service VLAN environment. More details about service VLANs, C-VLANs, and S-VLANs are covered in Chapter 4.

Business Layer 2 VPNs

In the access and aggregation networks, the architecture is much the same, except with Layer 2 VPNs, it is almost mandatory to have a dedicated C-VLAN. Ideally an S-VLAN between the DSLAM and the IP edge is used for scaling reasons. A Layer 2 VPN is a customer-managed VPN service. The service provider (SP) delivers a Layer 2 connection between two of the customers' CPE. A simple approach is for a CPE to have a DSL connection over which is a bridged Ethernet connection to the provider's PE. The SP runs an MPLS network and uses Layer 2 VPN circuits between two PE routers. At the other PE, a similar architecture is employed, with a VLAN to the other CPE. Bridged Ethernet (or pure Ethernet if using VDSL) is also used on the access link. All devices on the two LANs are on the same broadcast domain. This contrasts with the Layer 3 case, in which the two CPEs need to be running in routed mode.

Enhanced Internet Access

Internet access is a very important service for business customers, especially if their income is derived solely from an Internet presence. An Internet service can be delivered either using a stand-alone router at the provider, as described in the two preceding sections, or on the same BNG as residential broadband services. If a provider uses separate hardware to deliver services to business clients, business-grade Internet services are usually delivered on the same routers as those for Layer 3 and Layer 3 VPNs. What is business-grade Internet access? On a technical

level, it can be supposed benefits from the hardware separation from residential access. Even if there is no physical service separation, a business connection usually has a higher CIR than a residential customer.

It is common for business customers to have provider-hosted Layer 3 VPNs and Internet access from a single provider. There are several ways of delivering these services to customers from a perspective of PE, VLAN, and CPE configuration. For instance, a single VLAN between the CPE and PE could be used to deliver all services, and the PE does the routing between the Layer 3 VPN and the Internet. Alternatively, one VLAN can be assigned per service, and the CPE handles the routing between the Layer 3 VPN and the Internet, and, if applicable, NAT between any private address and public Internet addresses. More details on these connectivity solutions are covered in Chapter 3, “Designing a Triple-Play Backbone.”

VOICE SERVICES

Voice over IP (VoIP) is a hugely popular service and is being deployed by traditional wire line telcos, competing ISPs, and Internet-based VoIP providers such as Vonage and Skype. The competitive landscape is pushing traditional carriers toward migrating their traditional TDM-based switches to more cost-effective packet-based ones. Additional services can be delivered using online self-provisioning tools. A customer can log on to his or her portal, enable the second VoIP line into his or her home, and select a range of voice mail and unified messaging options. Many of the more powerful features of use to businesses can be ordered online, too. An extra ten outgoing lines from a business PBX can be enabled at the press of a button, for example.

This section explains some of the VoIP services that providers can offer to customers. Even though VoIP services can be offered over any IP connection, the target audience for this section is wire line providers delivering voice services. Purely Internet-based providers such as Vonage and Skype share some similar attributes, but their services are more Internet-based and are not covered here.

POTS ACCESS TO DSLAM

One of the simpler approaches to delivering a VoIP service is to deploy it in stages. Voice services that coexist with DSL connections have used a frequency splitter at the customer premises to separate the lower-frequency voice band from the DSL signals. A VoIP architecture does not need to be IP end-to-end. Instead, a voice switch at the central office (CO) or exchange can provide the interworking between the POTS and IP voice signaling. This avoids the extra complexity of having to deploy SIP-capable CPE. Many DSLAM vendors are integrating voice switch capability into their products, which enables the local loop at the CO to be terminated on the DSLAM for both voice and data services. This eliminates the need for external splitters and voice switch hardware if a single provider has exclusive use of the copper to the customer.

In terms of the service to the customer, the features are almost the same as POTS or ISDN, which have been used for many years. On the provider side, the voice switch in the CO takes any analog calls from the customer and sends them down a VoIP trunk to one of the soft switches. These soft switches perform call routing and billing. A sample end-to-end architecture is described in the next section. This is a good stepping-stone before a full-blown end-to-end VoIP architecture. Such an architecture might also be used in environments where a competitor wants to supply both a voice and DSL service to the customer, rather than allow the incumbent to have the voice-band service.

END-TO-END VOIP

This section explains the components involved in a simple VoIP architecture. Designing and implementing a VoIP network is complex business, so the intent of this section is to provide a casual introduction to how the signaling and media flow both within and beyond the network. Figure 2.11 shows a VoIP server network located on the left at one of the provider's PoPs. The server network consists of an SDH interconnect with a local voice provider. A VoIP VPN also is used as a private interconnect with another VoIP provider, either domestic or international.

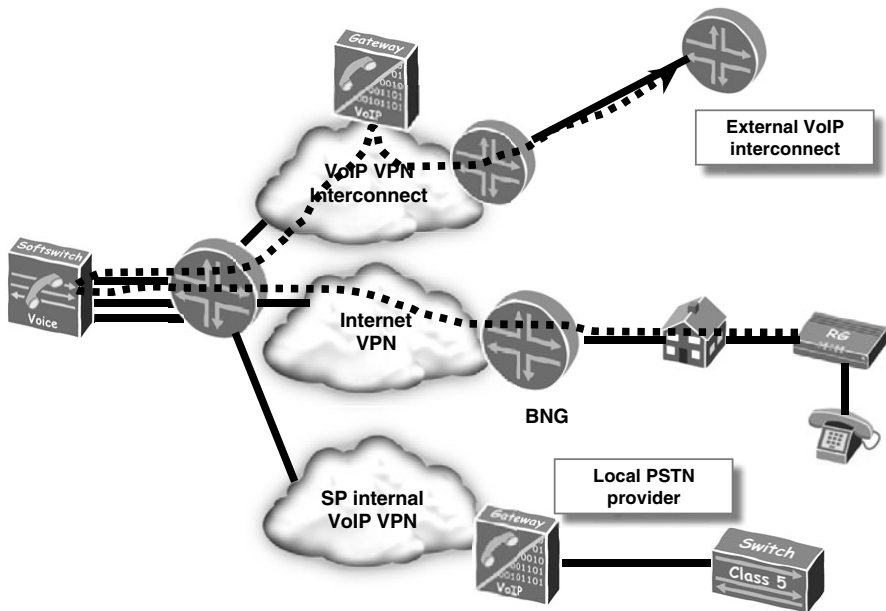


Figure 2.11 An end-to-end simple VoIP network architecture.

In Figure 2.11, a SIP CPE has a POTS phone connected. The SIP CPE presents a POTS interface to the phone and, at the same time, registers with the soft switch. After the CPE has registered itself with the soft switch, usually using a special username and password as an identifier, the soft switch then knows at which IP address a number can be reached.

As soon as the SIP user agent (UA)—the SIP CPE in this case—has registered with the voice switch, it can place calls with and receive calls from the network. In Figure 2.11, the CPE makes an international call via the external VoIP interconnect. The UA signals its call intentions to the soft switch, which establishes a signaling session with the international provider’s voice switch. If the call can proceed, the UA is directed to send its media session (the voice part of the call) to the soft switch or Session Border Controller (SBC) (described in more detail in Chapter 12, “Security in Broadband Networks”). This media session is represented by the dashed line, which flows to the switch/SBC, and then via the VoIP gateway, shown at the top of the diagram. To add some complexity, for the purposes of illustration, the international provider requires that calls be received in a lower bit rate than the G.711 encoding that the provider uses internally. In this

case, the encoding is G.729, so the media gateway needs to *trans-code* the call before sending it on to the international VoIP provider's voice switch.

This is just one example of a VoIP architecture. There are many variations, such as more integrated units performing media transcoding and media signaling.

SUMMARY

This chapter has briefly illustrated the network components that are described in more detail in the rest of the book. These range from the video server head-end to the MPLS core and from the aggregation network to the BNG. Approaching a multiplay network from a service perspective rather than being driven purely by technical capability is paramount for a successful deployment. This chapter discussed some ideas for multiplay data services, such as a premium gaming service or walled-garden content. From the video perspective, adding extra value that traditional bricks-and-mortar video shops cannot easily provide is key to enticing users to switching to video over IP. Some features are solved in a more routine manner, such as having a large range of content. Other services are more complex to implement, such as an NPVR service.

Finally, a VoIP service is almost as important as, if not more important than, an Internet service to deliver to customers. Some customers will want just a telephone connection without any data or video services, so, at a minimum, a stand-alone VoIP service needs to be available. This could be delivered using a method with only incremental complexity by using DSLAM-based VoIP trunks (implemented using H.248 trunking, for example) and using a low-pass frequency splitter at the customer premises. Or you might not bother with a splitter if only a PSTN service is provided. On a more complex level, providers can deliver VoIP services directly to the customer. For the more residential customers, the ideal option is to have one or more analog POTS ports on the CPE, to which they connect analog phones. The CPE then acts as a SIP UA on the phone's behalf. This is a VoIP service directly to the home. For more-complex business requirements, the SIP UAs are devices on the LAN and generate the necessary SIP or H.242 signaling and RTP data streams themselves.

QoS needs to be carefully managed for all these services. Chapter 8, "Deploying Quality of Service," covers quality-of-service implementations for effective triple-play networks.