

# Delivering Multicast Video Over Asymmetric Digital Subscriber Line

*Laying the Foundation for Next-Generation Video Services*

The New World beckons with promises of unprecedented new revenues and profitability from exciting new services. But uncharted territory has its perils. Service providers must select a mix of new, differentiated services to attract and hold subscribers. The expanding availability of broadband access technologies such as digital subscriber line (DSL) offer high-speed access affordable to mass markets, giving service providers the ability to enable new types of service requiring ever-greater bandwidths. Providers who correctly gauge and sell to their markets need solutions that enable high-speed services. They must choose network architectures and technologies that enable new services at a reasonable cost. The underlying architecture and technologies must also provide flexibility for expansion and adding new capabilities as markets grow and mature.

It is no surprise, then, that many service providers are evaluating the viability of digital video services delivered over DSL connections to both residential and business markets. Cisco Systems, the worldwide leader in networking for the Internet, delivers multicast video solutions that service providers can deploy today. Commercial deployment of multicast video services is made possible through the maturation of enabling DSL technologies, the proliferation of DSL subscribers, and emergence of IP multicast video content providers.

To serve the requirements of multicast video services, the underlying DSL access network must offer enough downstream bandwidth. Some DSL deployments offer symmetric upstream and downstream speeds, and others

divide bandwidth asymmetrically. Streaming video requires high downstream bandwidth. Asymmetric DSL (ADSL) provides high downstream and low upstream bandwidth and currently is the most available DSL service in the market. It is also well-suited to deliver multicast video services. Therefore, this discussion focuses on ADSL as the foundation for video services because it is widely available from service providers and mature. When other DSL technologies such as symmetric DSL (SDSL) and very-high-speed DSL (VDSL) mature, they will enable new options for how video service is provisioned and delivered.

When planning for mass-market deployments, it makes sense to enable the most scalable, least expensive option. Multicast technology enables unique advantages for mass-market video services because it is far more scalable than unicast video streaming. Instead of delivering one stream for each viewer from the source, a multicast solution offers one stream through the network core that is locally replicated at the edge to all its subscribers. Signal replication in unicast networks occurs at the host, and the network core must support multiple streams to viewers, requiring massive bandwidth when multiplied over potentially millions of subscribers. This is a more expensive and less scalable delivery mechanism than multicast, where replication occurs at a router close to the viewer.

Both unicast and multicast video systems deliver the highest quality when quality-of-service (QoS) mechanisms are active from end to end (from host to client). Streaming video via the Internet cannot provide QoS controls to a

service provider, so the end-user experience often varies, depending upon changes in traffic load and the native intelligence and configurations of network devices along the route. Mass-market success will require high-quality delivery, so initial multicast services should be offered within a “walled-garden” environment. A “walled garden” is traffic that stays within a service provider’s own network infrastructure, enabling that provider to implement strong end-to-end controls over signal quality. Therefore, Cisco chose to begin its solutions development for video-over-ADSL services with multicast IP video in a walled-garden environment.

This discussion presents initial steps enabling multicast IP video services to be delivered over ADSL. It is a starting point for future growth of network capacity and capability as the industry explores emerging market for digital video services. We address the characteristics of multicast video-over-ADSL services, explore relevant enabling technologies, and compare the value and limitation of several access architectures for multicast delivery.

### **Characteristics of Multicast-over-ADSL Service**


In this section, we discuss the types of differentiated multicast video service that providers can deliver via ADSL, and some of the available service modeling options that can help providers attain profitability and growth. While multicast technology itself is relatively simple and well understood, it is important to clearly define the several necessary components for enabling multicast video services that generate revenue. These include how channels and subscribers are provisioned, subscription to services, content types, controlled content access, and billing.

Multicast IP video content delivered via ADSL enables more breadth of service than cable networks. Cable is a shared medium with fixed bandwidth, and the number of channels is limited. Because ADSL is a dedicated medium and multicast clients access multicast streams that are selectively routed through the network, it is possible to enable more channels with DSL than with cable. This opens the door to delivering more channels over DSL than is currently possible with either cable or broadcast television. By offering video-over-DSL, service providers can further differentiate themselves from cable operators by their content. Content can be either widely distributed or tightly controlled. With such differentiation in mind, providers can offer an array of tiered or bundled services that appeal to a wide range of audiences.

### **Video-over-ADSL Services Business Model**

A thorough video-over-ADSL service design recognizes that not all subscribers can achieve the same bandwidth rates. Most ADSL networks have heterogeneous bandwidths, a combination of line-length limitations and subscriber preference. Video quality is dramatically affected by available bandwidth, so it is desirable for providers to “simulcast” content at several speeds, rather than only broadcasting at the lowest-common-denominator rate. Users with higher-speed access can enjoy superior quality video. For example, the content provider CoolCast.com enables up to eight different simulcast speeds and provides both automatic and user-selectable options to select the optimal streaming rate allocated to particular channels (Figure 1).

Figure 1 End-to-End Multicast Video Services



A multicast video service delivers quality video to PC users. Users watch video content, perhaps within a Web page while simultaneously accessing other subscribed services. Service providers can design Web pages to control the “look and feel” of a service. For example, they could associate additional text and graphical information with video streams and display these in the Web pages displaying video. Or a user watching a sports event could also receive player or team statistics.

Let’s consider types of content, the channel access model, access control, and other service features.

**Types of Content**

Video content can be either widely available or tightly controlled, depending upon source and audience.

Anticipating multicast video service deployment over ADSL, content providers such as CoolCast have started to provide these channels at typical ADSL bandwidths to service providers as IP multicast streams.

Examples of such content could include, but are not limited to, the following:

- Broadcast television/cable channels commercially available to network service providers, such as:
  - Bundled, commercial channels (such as ESPN, Disney)
    - Basic network TV channels (such as ABC)
    - Premium channels (such as HBO)
  - Pay-per-view channels
  - Public channels
- Special interest group video channels targeted at niche audiences have the potential for rapid growth when basic multicast video service is available.
- Local channels spotlight local cultural events, sports, and other local activities that can be multicast and targeted within specific localities or audiences of any size
- WebCam content allows mobile users to visually monitor premises, such as homes or day care centers, and enables security agencies to enhance home security services via visual monitoring

- E-learning enables online training or education with video from training rooms that is multicast to online students
- Local Advertisements, which providers can insert to gain incremental revenue. Service providers can insert targeted advertisements into video channels received from content providers by temporarily switching the regular multicast stream with another carrying the advertisement.

#### Channel Access Model

To generate revenues, it is important to derive an appropriate channel access model that allows providers to bill for premium content. Some content (such as basic network television channels) will always be available to users free of charge, but content such as pay-per-view or premium channels must be controlled to retain its value as a source of potential revenue. Other services (such as WebCam, local, or special interest group channels) require authenticated access for privacy and security reasons. Users should have access only to those channels for which they are authorized via subscription. Service providers can deliver differentiated, community-focused services with specific channels as part of their multicast service offering. They can also employ content switching mechanisms to replace programming options. For instance, a local news program can preempt or replace national programming.

#### Access Control

Video services that offer multiple content channels can be organized into “packages” that collect related sets of channels and sell them as bundles to subscribers on a multitiered flat-rate basis. These correspond to familiar “basic” or “premium” packages common to cable television service. A service provider may offer one or more packages to subscribers. The network authenticates the user requests to access channels based on their subscriptions. Users can access all premium channels and free multicast channels in their subscription packages. Providers can also enable simultaneous access to multiple services such as Internet access and video game servers.

The initial end-user experience centers on standard PCs with a Web interface, though future applications and products may enable delivery to other types of system. Cisco Service Selection Gateway technology enables access control and enhances the end-user experience, allowing access to multiple destinations and services. In this startup model,

users log onto the network by connecting to a known URL of the service provider, then enter a login name and password on the login page. Successful log-ins display a dashboard menu of their subscribed service packages. Users select a multicast video package, which displays an associated Web page that lists available channels or channel categories, allowing the user to navigate and select a channel. When a user selects a channel, the channel is displayed in a viewer window in the Web page.

#### Operations Management and Billing

Service providers can gather information that assists them with operations management and billing. Service usage is logged as accounting records in standard RADIUS format to ensure wide-scale compatibility with existing billing systems.

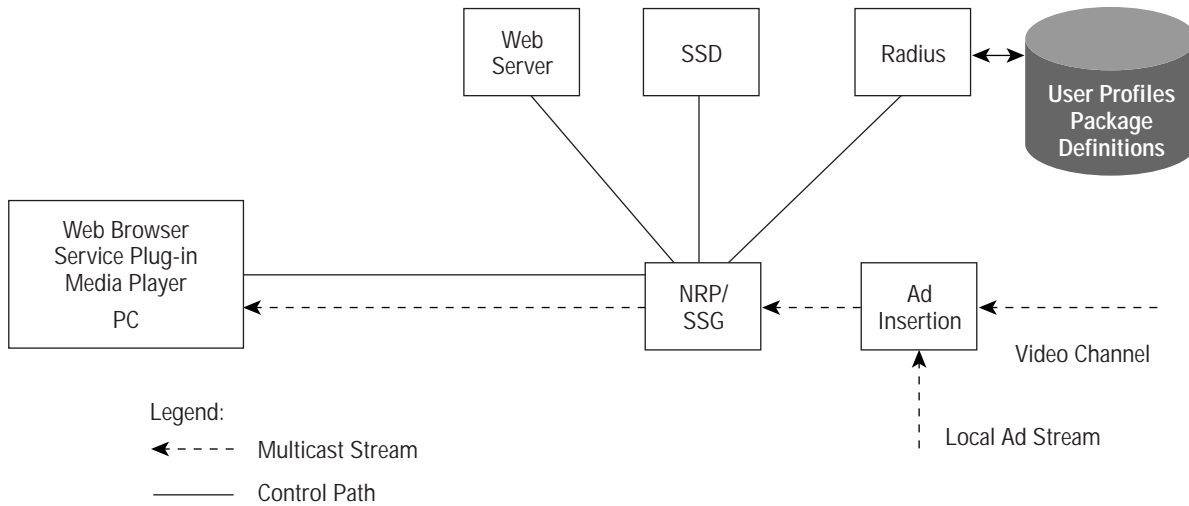
#### Browser Plug-in

Service-specific client software (a Web browser plug-in) controls the look and feel of the service. Client software can either be explicitly downloaded from the service provider’s Web server or automatically downloaded when a channel is accessed. Automatic download is easier for users because they do not need to explicitly download software or keep track of its updated versions. For providers delivering simulcast content at several ADSL speeds, another client software feature could be automatic detection of available bandwidth to the user. Based on this information, the plug-in chooses the most suitable channel size for optimal video quality. The CoolCast plug-in and Web server support this feature. The browser plug-in can also serve as a strong statistics gathering tool to facilitate advertising sales. For example, the CoolCast plug-in periodically provides demographic data (such as time, channels watched, and Web sites visited) to the CoolCast Web server.

#### Video-over-ADSL Services Architecture

The higher-layer services architecture for multicast video services includes several servers that enable its various elements, including user authentication, access control, service selection, service provisioning, and billing. A typical service architecture includes RADIUS servers, one or more Web servers, service selection gateways (SSG), and service selection dashboards (SSD). There may be optional servers for content processing and advertisement insertion (Figure 2).

Figure 2 Multicast Video Service Components



**User Authentication**

An authentication, authorization, and accounting (AAA) server (a RADIUS server) maintains user and service profiles, and provides standard RADIUS-based functionality. User profiles correlate user identity with the service packages that they've subscribed to. RADIUS service profiles contain a URL for the corresponding services Web page. The format of each Web page can be service specific.

To access a video service, users log into the service provider network, select a multicast video service, and select a video channel, as detailed below:

- **User Login**—users connect via a Web browser to the known address of the SSD server and are presented with the user logon page. Users enter a username and password, and the page is forwarded to the SSD, which communicates with the RADIUS server (via SSG) to authenticate the user. The RADIUS reply contains the list of services that the user has subscribed to. SSD displays this list of services to users via their browsers in a dashboard menu format.
- **Service Login**—when users select a service from the dashboard, the selection is forwarded to the SSD, which retrieves the service profile from the RADIUS server. The SSD visually indicates to the user that the user is logged on to the service. The SSD then redirects the browser to the proper URL. This displays the service Web page that typically contains a list of channels or channel categories.

- **Channel Access**—when users select a channel, the user's PC issues an Internet Group Management Protocol (IGMP)-join request for the channel via interaction between plug-in and media player. The plug-in learns channel-to-IP-address mapping by interacting with the Web server associated with the service or by listening to a well-known multicast stream that continuously multicasts mapping information. When users join a multicast group, the Cisco 6400 Universal Access Concentrator (UAC) acts as an access router, forwards multicast packets to each user for display on the user PC screen.
- **Multicast authentication**—the SSG provides subscription-based access to multicast video streams. It intercepts user channel access requests (IGMP-join messages) and allows the join to succeed if the channel is included in one of the services to which the user has currently logged in (or if the channel is a free channel). Because users can only log into subscribed packages, this limits user access to only those channels included in their subscribed packages, in addition to free channels. SSG intercepts IGMP-join requests, not multicast data packets, so authenticated access to video streams does not degrade multicast data throughput of the Cisco 6400 platform.

Multicast Authentication is demonstrable in lab, although it is currently not available on Cisco 6400 NRP. However, even without this feature, service providers can deploy video multicast service on a flat rate billing basis.

### Provisioning and Billing

Service and subscriber provisioning involves a RADIUS server and a Web server that provides browser plug-ins to the PC, channel-to-IP address mapping, and interaction with the plug-in. Typically, Web servers belong to service providers and are used to maintain Web pages associated with service packages from one or multiple content providers.

For provisioning a new multicast service using SSG functionality, an administrator creates a new service profile in the RADIUS server for each package included in a service. Package profiles include information such as the name, description, and URL of the package home page.

As channels are added or removed from service packages, the RADIUS service profile stays the same. The administrator may have to modify the package Web page accordingly.

Subscribers are provisioned by creating a new user profile in the RADIUS server that includes information such as subscriber username and password, names of subscribed services, and network data such as subnet mask and framed

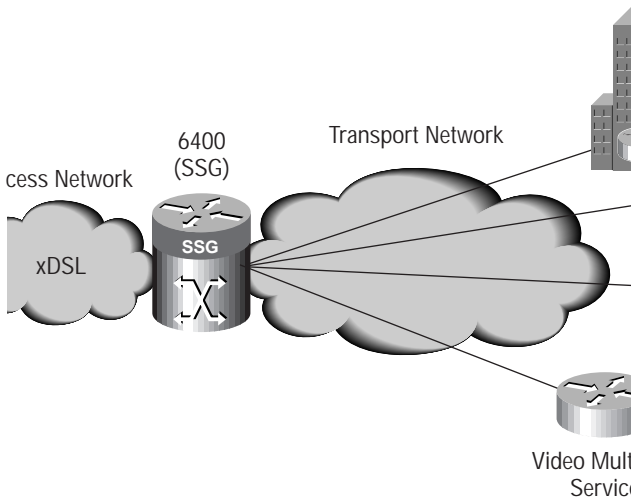
IP address. When users update their subscriptions, profiles are modified to reflect these changes. Billing multicast video services could be usage based for pay-per-view channels, or offered on a multitiered, flat-rate subscription basis for premium channels. Billing can be based on statistics gathered by the RADIUS server as the system records user login/logout and service login/logout.

### Service Selection

The service provider provides an SSD as a Web-server-based application. The SSG is an integral component of Cisco IOS® software on the Cisco 6400 node route processor (NRP). The SSG interacts with the SSD for user authentication and connection control as previously described. SSG internally maintains user-service connection states and accordingly controls a user's access to services. It also enables subscribers to simultaneously connect to multiple services from one or multiple content providers (Figure 3). These services could be

any combination of data, voice, and video services offered by the provider, such as Internet access, telecommuting VPN services, and multicast video.

Figure 3 Simultaneous Access to Services from Multiple Providers



### Enabling Technologies

Cisco IOS software contains many features that service providers need to deliver high-quality multicast video services. Among these are multicast routing protocols and QoS mechanisms. In addition to Cisco IOS software, standard video-encoding algorithms are also required. The technologies referenced in this section are explicitly and thoroughly described in other documentation, so readers are presumed to have a working knowledge of these features.

#### Multicast Routing Protocols

IP multicast uses multicast routing protocols, several of which have been proposed in standards bodies, including Distance Vector Multicast Routing Protocol (DVMRP), Multicast Open Shortest Path First (MOSPF), and Core-Based Tree (CBT). DVMRP and MOSPF do not scale well. CBT is new and has not achieved wide-scale acceptance. Protocol-Independent Multicast (PIM) has emerged as the *de facto* standard.

IGMP—Used by IP hosts to report multicast group memberships to an adjacent multicast router. Note that the standard multicast protocol IGMP does not support controlled access. Cisco addresses this need via specific Cisco IOS enhancements to provide authenticated access to video channels. For example, the SSG has a feature that intercepts each IGMP-join request and checks to see whether the requester has logged into a service that includes the requested

channel. If so, the request is accepted and, if not, an error message is returned to the user (currently not available in IOS, but demonstrable in lab).

PIM—A multicast routing protocol that enables multicast routing on existing IP networks. It is the most widely deployed multicast routing protocol because it is more scalable than either DVMRP or MOSPF, and unlike them will interoperate with any existing unicast routing protocol in the core such as OSPF and EIGRP. PIM can be operated in either dense or sparse mode:

- PIM dense mode is data driven, and resembles typical multicast routing protocols. Packets are forwarded on all outgoing interfaces until pruning occurs. Receivers are densely populated and downstream networks want to receive and will probably use the datagrams forwarded to them. The cost of using dense mode is its default flooding behavior.
- PIM sparse mode tries to constrain data distribution so that a minimal number of routers in the network receive it. Packets are sent only if they are explicitly requested to a designated router called a “rendezvous point” (RP). Receivers are widely distributed and downstream networks have simultaneous demand for few multicast streams. The cost of using sparse mode is its reliance upon periodic refreshing of explicit join messages and its need for rendezvous points. Sparse mode is more scalable than dense mode.

PIM can be operated in a hybrid sparse-dense mode that offers flexibility to offer different video channels in either mode depending on channel viewing patterns. This is the recommended PIM mode.

#### Encoding and Compression

For all services, video is typically encoded in MPEG-2, though alternate encoding schemes such as MPEG-1 and MPEG-4 can also be used. Encoding happens at the content source, and decoding occurs in a user’s PC. The network can support any encoding scheme as long as the PC can decode it using a plug-in.

The ClearBand server uses a PC with software-based encoding. It is an example of a new encoding scheme that generates higher-quality video at lower bit rates than those of standard MPEG-2. It can generate multicast video streams from video sources such as video cameras, video cassette players, and stored video files. The software encodes video

streams into an MPEG-2-compliant format at user-specified rates (300 kbps and upward), and transmits data as an IP multicast stream. Users need a thin-client (a plug-in to a standard Web browser) to decode and view video streams. In the case of ClearBand, the Web server delivers the plug in as part of the multicast stream to simplify user setup. Depending on available ADSL bandwidth and video encoding technology, users can expect quality video at full screen or smaller screen sizes (half or quarter screen). For example, ClearBand encoding allows full-screen video of acceptable quality at about 500 kbps.

#### Quality of Service

Multicast video service is sensitive to packet loss and delay because they adversely affect the video quality; hence the expectation that most initial offerings will be restricted to walled-garden environments where service providers can implement tight QoS controls from end to end. QoS mechanisms are configured to control how specific traffic classes are handled. Providers can minimize loss and delay by assigning higher QoS priority to video traffic. The network then drops or delays lower-priority, time-independent data in favor of video traffic, thus mitigating or avoiding traffic congestion.

A clearly defined QoS policy establishes traffic classifications, balancing business priority with application sensitivity to loss and delay. For example, mission-critical applications are typically time and error sensitive; voice and video are very time sensitive but can tolerate occasional packet loss; noncritical traffic (such as file transfers), have no stringent delay requirements but are error sensitive. Typically, multicast video streams are classified at or close to the highest priority.

Traffic can be classified and controlled in several ways (including source, destination, protocol, port). Layer 2 classification as found in ATM networks may be insufficient for traffic that must cross several networks before reaching users. For end-to-end service, Layer 3 IP precedence (three type-of-service bits in IP headers) provides a convenient way to assign priority to traffic. Each packet is classified at the edge of the network, its IP precedence bits are set, and then packets are forwarded to the network core. Otherwise, packets may get delayed before reaching the core. Committed Access Rate (CAR) incorporates features that set IP precedence and provides rate limiting to deliver a high degree of control at the edge. Providers may use its rate-limiting capability to limit video streams from a content provider to a maximum agreed data rate.

Congestion may also occur in the network core. Congestion avoidance QoS mechanisms selectively drop lower-priority packets to “throttle back” the rates of low-priority flows to assure delivery of higher-priority ones. If congestion avoidance in the core is required, the service provider may consider deploying Weighted Random Early Detection (WRED). The effectiveness of WRED will depend

on the traffic mix and volume and the service provider's QoS traffic classification. WRED is useful with protocols that intelligently slow down transmission when packet drops are detected (such as TCP). Note that if multicast video streams are directly received from satellites at the PoPs (bypassing the core), any existing QoS policy of the core will not be influenced by multicast video service deployment.

IP precedence-based queuing techniques such as Weighted Fair Queuing (WFQ) are appropriate for the access network, where lower last-mile bandwidth can cause congestion. WFQ allocates usable bandwidth based on the precedence level of packets in queuing buffers.

### **Multicast-over-ADSL Network Architectures**

The foundational network architecture significantly determines the inherent scalability, QoS, security, and provisioning characteristics of a multicast video service.

This section discusses how network architecture influences the quality of multicast video service and identifies those aspects that help or hinder multicast video service deployment. There are three logical elements of the network architecture: content acquisition, transport, and access.

### **Content Acquisition Network**

Content can be locally generated by a service provider or received from a content provider. As discussed earlier, video streams are usually compressed in MPEG format and transported as IP multicast streams.

There are several ways to obtain content for delivery, such as satellite, stored content from video servers, or from live camera feeds. Input streams are aggregated by a central site switch, such as the Catalyst 8500 series, and transported to the Points of Presence (POPs) via the transport network. Content providers could also transport video streams via a private satellite network directly to the POPs, bypassing the core. This would conserve transport network bandwidth for other services. This also makes video quality independent of QoS configurations in the core network. On the downside, using satellites calls for additional equipment and increases management overhead. Coolcast is an example of a content provider that uses a private satellite network to distribute content and provides satellite receivers for the POP or central office.

### **Transport Network**

For providers who do not receive content at POPs via satellite, the transport (core or backbone) network carries video streams from the source to the POP. As previously explained, only one multicast video stream travels the

transport network to an access POP, independently of the number of subscribers. The router nearest the subscriber dynamically (on demand from subscribers) replicates the multicast streams and forwards them into the access network to subscribers using the IGMP protocol. To support video service, routers in both the transport and access networks must be multicast enabled.

Multicast video service can be deployed over any existing transport architecture that can support the required bandwidth and supports appropriate multicast routing protocol and QoS. The transport network must have enough capacity to support aggregate bandwidth of each central office or POP. An OC-48 or OC-12 core may be suitable, depending on aggregate traffic volume. Packet over SONET (POS) provides 25- to 30-percent gain in efficiency over comparable IP over ATM over SONET architectures. One option is the Cisco 12000 gigabit switch router (GSR) connected via OC-48 POS. GSRs would reside in each major central office. Each location without a GSR is connected to one that does via an ATM virtual circuit (VC) (Figure 4).

### **Access Network**

Access network architecture plays a crucial role in multicast service deployment in terms of scalability, security, QoS, and subscriber provisioning. In general, the access network consists of a Cisco 600 series DSL modem at the subscriber premises, with digital subscriber line access multiplexers (DSLAMs) and access routers (Cisco 6400 UACs) in the central office. At the subscriber premises, one or more PCs can be connected via a 10/100 Ethernet LAN to the customer premise equipment (CPE). Video channel bandwidth cannot exceed the "last-mile bandwidth" to the CPE. Typical subscriber bandwidths range from 256 kbps to 1.5 Mbps in many current Asymmetric Digital Subscriber Line (ADSL) access networks, although higher speeds are available in certain areas from certain providers.

Multicast replication takes place in the access router, and streams are forwarded to subscriber CPE (Figure 5). The SSG feature of the Cisco 6400 access router allows users to connect simultaneously to multiple destinations, and is also required for supporting authenticated access to multicast streams.

Figure 4 Transport Network Architecture

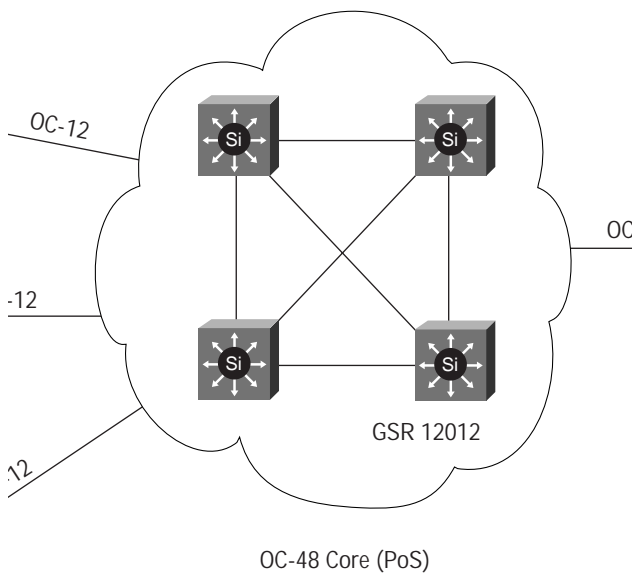
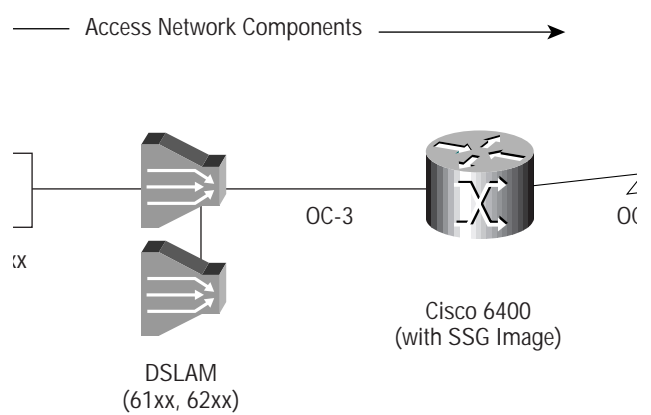


Figure 5 Access Network Architecture



Scalability is an important consideration because video requires high bandwidth and some access architectures scale better than others. Because packets are replicated by the access router, downstream bandwidth from the access router is an important scalability factor. As with any consumer service, security is an important feature. Certain access network architectures that are susceptible to Address Resolution Protocol (ARP) spoofing and IP hijacking are not suitable for video service deployment. Another important

factor in an access architecture is proper QoS for video traffic. Reducing the complexity of subscriber provisioning is an important component enabling profitable, large-scale service deployment.

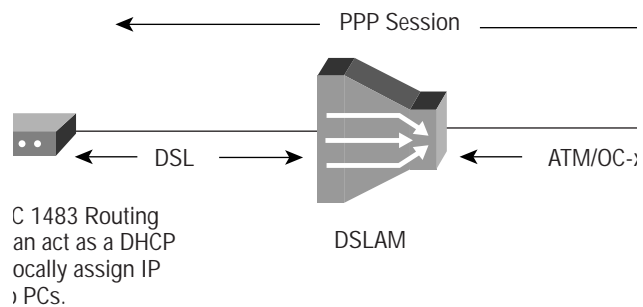
Four DSL access architectures offer varying suitability for multicast video services. Of these, three are viable for multicast video, Point-to-Point over ATM (PPPoA), Point-to-Point over Ethernet (PPPoE), and Route Bridge Encapsulation (RBE), so they are discussed in detail below.

- *Integrated Routing and Bridging (IRB)*—allows a router to act as both bridge and router on the same interface. However, this architecture is not suitable for multicast video services. It does not scale well because it floods downstream multicast packets to multiple subscribers in the same bridge group, overloading ADSL links. It also presents security concerns with possible ARP spoofing or IP hijacking. While this is usually not a problem in enterprise deployments, it is a serious issue for services provided by a service provider.
- *PPPoA*—an architecture endorsed by the ATM Forum
- *PPPoE*—a relatively new, IETF-standard access architecture with limited deployments that requires each PC to have preinstalled PPPoE client software.
- *RBE*—an enhancement of IRB that enables scalable multicast service deployment. This new Cisco technology addresses the multicast scalability and security issues of bridged networks.

#### PPP over ATM

The RFC 2364 PPPoA access architecture is suitable for deploying scalable multicast video service. CPE configured for PPPoA works as a router, establishing a PPP session with the Cisco 6400-NRP and routing subscriber data from PCs to the NRP through the PPP session. CPE obtains an IP address via PPP/IPCP negotiation and can serve as a DHCP server to assign IP addresses to its connected PCs (Figure 6).

Figure 6 PPPoA Architecture



booting, establishes PPP session with 6400 UAC. It gets IP address (unless statically assigned).

PPPoA often requires per-subscriber CPE configuration with user name and password for PPP authentication and an IP address pool. To avoid per-subscriber CPE configuration, service providers can configure all CPEs with identical user names and passwords, and live with the fact that no real PPPoA CPE (RADIUS) authentication is possible. Cisco provides features in Cisco IOS release 12.0(5)DC to address this. These features enable mass CPE configuration in PPPoA environments while allowing individual CPE authentication. They are:

- VPI/VCI-based authentication, which allows RADIUS authentication of individual CPEs by identifying CPE via its associated VC rather than a username, enabling individual authentication of CPE provisioned with a default user name/password.
- The IPCP subnet feature, which bypasses manual configuration of CPE DHCP pools. This feature enables CPE to automatically configure its local IP pool based on the subnet mask it receives during PPP negotiation.

#### Advantages of PPPoA:

- PPP allows feature negotiation between peers; in the future, CPE implementing automatic feature negotiation will be able to negotiate common supported features with the Cisco 6400 access router; for example, if CPE is upgraded to enable data encryption, it can automatically activate encryption if the Cisco 6400-NRP supports it
- CPE mass configuration is possible with VPI/VCI-based CPE authentication and IPCP subnet features
- Scalability is comparable to RBE
- Does not have IP hijacking issue of IRB

#### Limitations of PPPoA:

- CPE configured for PPPoA needs an IP address

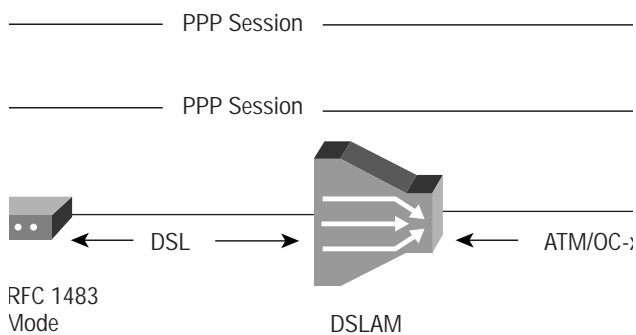
- VPI/VCI-based authentication needs an authentication, authorization, and accounting (AAA) server that can be customized to support the feature.

#### PPP over Ethernet

Like RBE and PPPoA, PPPoE is suitable for scalable multicast video service. Like RBE and unlike PPPoA, CPE runs in 1483-bridged mode and does not need an IP address. A PC behind a CPE typically gets its IP address via PPP/IPCP. Many feel that the requirement for preinstalled client software and ensuing support issues hinder the widespread adoption of PPPoE. Current Windows operating system software does not have this capability as built-in software, but third-party client software is available.

PPPoE offers a familiar dial interface to subscribers over a bridged CPE. It allows PCs connected to CPE to directly establish separate PPP sessions with the Cisco 6400 NRP. The CPE operates in RFC-1483 bridged-mode. PPPoE includes a discovery protocol (RFC 2516) to learn the Ethernet address of its remote peer (the Cisco 6400 NRP) and receive a unique session identifier to establish a PPP connection over Ethernet. (Figure 7).

Figure 7 PPPoE Architecture



' session with 6400 UAC. It gets IP address via IPCP (unless static

#### Advantages of PPPoE:

- Suitable for offering the familiar PPP interface to subscribers in an RFC 1483-bridged CPE environment
- Does not have the IP hijacking issue of IRB
- Scalability comparable to PPPoA


#### Limitations of PPPoE:

- Windows 95, 98, and NT operating systems do not support the PPPoE protocol stack, so subscriber PCs require third-party software installation.

#### Route Bridged Encapsulation (RBE)

RBE is particularly well-suited for providing multicast video service over ADSL networks. It is a new feature introduced in Cisco IOS release 12.0(5) DC and is specifically designed to address the multicast scalability and security issues associated with IRB. RBE forwards a downstream multicast packet to only subscribers currently accessing the stream and does not require bridge groups, thus eliminating the scalability problem. Because ARP packets are not broadcast but sent to a subscriber's CPE, other subscribers can't access the packet for the purpose of ARP spoofing. RBE uses separate subnets per CPE. Because each subinterface has a unique subnet, the Cisco 6400 NRP detects a fake IP address from another subscriber's subnet in an ARP reply, drops the packet, and generates a "wrong cable" error (Figure 8).

Figure 8 RBE Architecture



nd forwards a downstream multicast stream to only to those CPEs

In RBE, CPE operates in bridged mode and does not need an IP address (unlike PPPoA). PCs are typically configured to obtain IP addresses from a DHCP server.

#### Advantages of RBE:

- Scales well for multicast service
- Offers protection against ARP spoofing/IP hijacking among subscribers
- Offers an easy migration path for existing IRB-based networks to add multicast video service
- Simple CPE configuration for new subscribers (simpler than PPPoA)
- CPE need only support RFC 1483 bridging
- CPE can be mass configured

#### Access Architecture Comparison

While all four architectures discussed here can be used with



Corporate Headquarters  
Cisco Systems, Inc.  
170 West Tasman Drive  
San Jose, CA 95134-1706  
USA  
<http://www.cisco.com>  
Tel: 408 526-4000  
800 553-NETS (6387)  
Fax: 408 526-4100

European Headquarters  
Cisco Systems Europe s.a.r.l.  
Parc Evolic, Batiment L1/L2  
16 Avenue du Quebec  
Villebon, BP 706  
91961 Courtaboeuf Cedex  
France  
<http://www-europe.cisco.com>  
Tel: 33 1 69 18 61 00  
Fax: 33 1 69 28 83 26

Americas  
Headquarters  
Cisco Systems, Inc.  
170 West Tasman Drive  
San Jose, CA 95134-1706  
USA  
<http://www.cisco.com>  
Tel: 408 526-7660  
Fax: 408 527-0883

Asia Headquarters  
Nihon Cisco Systems K.K.  
Fuji Building, 9th Floor  
3-2-3 Marunouchi  
Chiyoda-ku, Tokyo 100  
Japan  
<http://www.cisco.com>  
Tel: 81 3 5219 6250  
Fax: 81 3 5219 6001

Cisco Systems has more than 200 offices in the following countries. Addresses, phone numbers, and fax numbers are listed on the

**Cisco Connection Online Web site at <http://www.cisco.com/offices>.**

Argentina • Australia • Austria • Belgium • Brazil • Canada • Chile • China • Colombia • Costa Rica • Croatia • Czech Republic • Denmark • Dubai, UAE  
Finland • France • Germany • Greece • Hong Kong • Hungary • India • Indonesia • Ireland • Israel • Italy • Japan • Korea • Luxembourg • Malaysia  
Mexico • The Netherlands • New Zealand • Norway • Peru • Philippines • Poland • Portugal • Puerto Rico • Romania • Russia • Saudi Arabia • Singapore  
Slovakia • Slovenia • South Africa • Spain • Sweden • Switzerland • Taiwan • Thailand • Turkey • Ukraine • United Kingdom • United States • Venezuela

Copyright © 1999 Cisco Systems, Inc. All rights reserved. Printed in the USA. Packet is a trademark; Cisco, Cisco Systems, Cisco IOS, and the Cisco Systems logo are registered trademarks of Cisco Systems, Inc. or its affiliates in the U.S. and certain other countries. All other trademarks mentioned in this document are the property of their respective owners. The use of the word partner does not imply a partnership relationship between Cisco and any of its resellers. (9910R) 12/99 BW5680