

Optimizing Video Transport in Your IP Triple Play Network

Service providers that compete for market share in next-generation consumer entertainment and communications services must move beyond “bundling” and offer personalized media and interactive IPTV services that blend entertainment, communications, and the Internet. With the deployment of these new IPTV services, existing network infrastructures will be pushed to their limits. To accommodate the needs of IPTV services, networks must be able to scale to millions of customers, maximize bandwidth resources, and provide quality of service (QoS) and security on an end-to-end basis. For these and other reasons, network intelligence is critical when deploying video over broadband. That is why many of the world’s most successful triple play service providers—including Comcast, Fastweb, Hong Kong Broadband, Neuf Telecom, and Surewest—have built their next-generation networks (NGNs) with IP NGN solutions from Cisco Systems®. Today, over 10 million subscribers around the world receive innovative video services over intelligent Cisco® IP NGNs.

The video services of the future will be delivered over next-generation IP networks. The requirements of these services are distinctly different from the requirements of high-speed Internet services. This paper examines the requirements for optimized video transport in an IP triple play network and shows how Cisco’s IP Next Generation Network (IP NGN) Architecture addresses the video challenges of today and tomorrow.

EXECUTIVE SUMMARY

Service providers offering high-speed Internet service typically define the service by the speed of upstream and downstream throughput, and thus the subscriber’s service-level agreement (SLA) is defined directly by transport parameters. With IPTV, in contrast, the service provider will define the service by the number of channels provided, the quality of the video, the size of the video-on-demand library, the user interface, the video recording capabilities, and the interactive features provided in the IPTV service. The SLA in this case is not defined by network transport parameters, but by the quality of the experience, which the network simply must support. Because of these distinct differences, Cisco has architected its IP NGN to enable optimum transport for each distinct service type. In this way, each distinct service can be protected by independent mechanisms, and SLAs can be enforced differently. For example, high-speed Internet is treated as a transport service, and Internet traffic passes through a broadband remote access server (BRAS), which provides per-subscriber QoS and authentication enforcement. In contrast, video traffic is treated as a managed application service. It runs natively over IP from the encoder to the set-top box and is not routed through a BRAS. In this case, QoS is delivered on a per-service basis and the application layer controls authentication and subscriber credentials. This per-service approach makes it possible for service providers to address the primary requirements of video over IP traffic. These requirements include:

- An optimized IP multicast design that takes full advantage of the IP architecture as video traffic moves from the video source to the set-top box
- QoS that is linked to the subscriber experience
- Network protection that goes beyond a basic link failure and provides meaningful failover mechanisms for coping with outages
- Admission control mechanisms that protect the service provider from video oversubscription, which can degrade the experience of many viewers simultaneously

TRANSPORT SERVICES VERSUS MANAGED APPLICATION SERVICES

When it comes to transport versus managed application services, subscriber requirements are very different. Bandwidth and availability are basically the only decision criteria subscribers have for high-speed Internet service selection. But prospective IPTV subscribers evaluate video services in many different ways. Subscribers want a wide selection of channels, a large on-demand movie library, an easy-to-use electronic program guide, and a large amount of high-definition content. When viewing this content, subscribers are looking for an enhanced experience that includes a continuous, high-quality TV picture and acceptable channel change times. Each of these subscriber requirements must be accommodated (with different SLAs that use different tools to apply and enforce different rules) for successful video over broadband delivery. Any provider that approaches video over broadband in the same way it approaches high-speed Internet service delivery will have a hard time differentiating services that attract and keep customers. The Cisco NGN is capable of addressing all requirements for the different types of voice, video, and data services that it carries. (See Table 1.)

Table 1. Comparison of Transport and Application Services Requirements

Service Type	Transport Service	Managed Application Service
SLA	Transport parameters <ul style="list-style-type: none"> Bandwidth, max drop, max latency, etc. 	Video application SLA <ul style="list-style-type: none"> Number of set-top boxes Basic vs. premium tier
Subscriber authentication/identification	Network Based <ul style="list-style-type: none"> PPPoE, 802.1X Per-subscriber VLANs, DHCP Option 82 	Application based <ul style="list-style-type: none"> Video middleware STB can be authenticated by the network when first connected
SLA enforcement	Network based <ul style="list-style-type: none"> Per-subscriber shaping/policing 	Application based <ul style="list-style-type: none"> Based on app signaling
QoS	Per subscriber <ul style="list-style-type: none"> Gold, silver, bronze Per-subscriber classification, queuing 	Aggregate <ul style="list-style-type: none"> Single queue for video

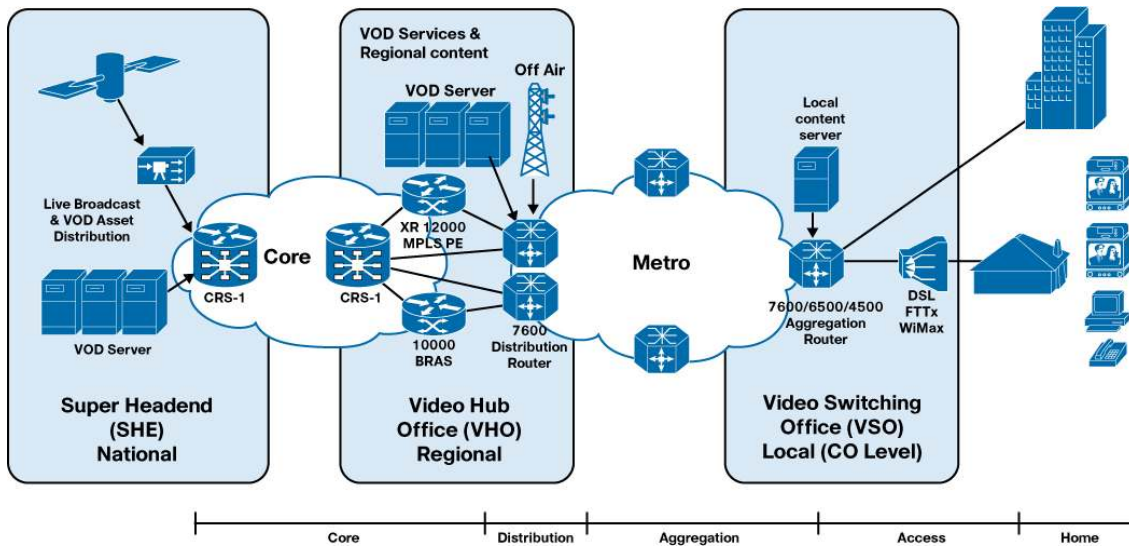
IPTV COMPONENTS OVERVIEW

The main components of an IPTV architecture include:

- **Super headend (SHE):** Live video feeds and real-time encoding of video broadcasts originate from the SHE, as do asset distribution systems for on-demand services. The SHE can also incorporate back-end systems such as the subscriber database. Video operators usually have one or two SHEs, which typically reside within the core of the transport network.
- **Video hub office (VHO):** VHOs often contain real-time encoders for local television stations and video on demand (VoD) servers as well as the network routers that connect the distribution network to the network core. Many VHOs house content servers for video-on-demand services. Many operators maintain a few dozen regional VHOs, predominantly in metropolitan areas. These VHOs typically service anywhere from 100,000 to 500,000 homes.
- **Central office and video switching office (VSO):** Central offices and VSOs house aggregation routers that aggregate traffic from subscriber homes. For DSL providers, the central office typically handles traffic from local digital subscriber line access multiplexers (DSLAMs).

A Cisco IP NGN can accommodate the full range of bandwidth and service requirements for IPTV because Cisco IP NGNs add scalability and intelligence throughout the network, but especially at the distribution and aggregation layers. (See Figure 1.)

Figure 1. Basic IP Video Network Architecture



VIDEO SERVICE REQUIREMENTS

When assessing an IPTV architecture it is useful to understand the overall demands video services place on the network.

- High bandwidth:** A home with an IPTV service subscription will require a substantially higher amount of bandwidth than a household that only receives high-speed Internet services. The increased traffic is because of the fact that video is delivered in constant, steady streams to the subscriber's set-top box. Image quality is controlled by the service provider, who determines the encoding rate. For example, the MPEG2 compression standard consumes approximately 3.75 Mbps. The newest compression standard, MPEG4¹, consumes only 2 Mbps while providing the same high-quality image. High-definition TV varies from 6 Mbps to 15 Mbps depending on the encoding rate. Another important consideration when assessing an IPTV architecture is the requirements of broadcast TV versus video on demand. Broadcast TV channels are delivered using IP multicast, which makes the bandwidth consumed dependent only on the number of channels offered and the encoding rate. For example, 200 channels of MPEG2 program content in standard definition will consume approximately 750 Mbps of bandwidth. Video on demand, however, is a unicast, per-viewer channel. One-thousand standard definition VoD viewers will consume approximately 3.75 Gbps. This makes managing VoD bandwidth complex.
- QoS:** QoS is extremely important when assessing an IPTV architecture because video over IP streams are sensitive to packet loss. And although the loss of one or several consecutive packets will not significantly affect a TV viewer's experience, if the event lasts more than a second it will degrade image quality. Set-top boxes have limited functionality for coping with frame losses. Many set-top boxes, for example, prevent visible artifacts from a network outage by utilizing Forward Error Correction² schemes that conceal missing information or by retransmitting lost information. Both of these methods are technically challenging³. Jitter is also an important parameter to consider, because set-top boxes have a limited capacity for jitter compensation (usually in the order of 150 ms). And although absolute delay is not that important for video delivery (if it is consistent over time), helping ensure consistent, end-to-end delay behavior is a primary attribute for video over IP network design. Finally, it is also important to consider the coexistence of video with voice over IP (VoIP) and other real-time traffic. When multiple services traverse the same network, different queuing needs emerge. Reliable scheduling and congestion avoidance mechanisms, combined with configurable queue sizes, should be part of the network solution.

¹ H.264 Part10

² Like the Pro-MPEG Code of Practice #3

³ Approaches along these lines are under development in the industry

- **Admission control:** Successful video services lead to rapid increases in the number of subscribers who want to access the video resources of the network. Past VoD deployments, for example, have seen a tenfold increase in viewer volumes in just a few years. A network cannot be designed to oversubscribe video traffic because oversubscription results in random video packet drops, which deteriorate all downstream subscriber experiences. To prevent oversubscription, the network must have mechanisms that interact with video resources and grant “permission” to new video streams only if they will not cause congestion in the network.
- **Video broadcast channel change time:** Although subscribers will not base their decision to purchase TV service solely on channel change times, it does affect subscriber satisfaction with the service. Therefore it is important that the network solution is optimized to minimize channel change latency.
- **Comprehensive service availability:** Video on demand and broadcast TV have radically different availability requirements. Broadcast TV is made up of multicast video streams. If one multicast source is lost inside the network, it could affect hundreds of thousands of users. The network must be optimized for IP multicast and must provide ways to transparently restore the loss of a multicast source. Implementing geographically dispersed redundant IP multicast sources is a good way to increase availability because it allows the network to select the most efficient source and provides rapid failover switching from one source to another. In contrast, video on demand is a per-user service, so the loss of one stream is not catastrophic. However, the mismanagement of video on demand streams or lack of appropriate queuing could cause serious oversubscription problems. For example, if a network fails and the resulting backup path causes bandwidth shortage, random packet losses from different video sources will bring the whole service down. It is essential, therefore, when considering service availability, to utilize different QoS schemes for video on demand and for broadcast TV.
- **Service lifecycle:** When a new service such as video is implemented, the video take rate will vary depending on the density of the area served, the amount of time the service has been offered, the targeted marketing campaigns, promotions, and so on. In other words, the user base is not a static component of the service. The network must be able to easily accommodate these service lifecycle variables with no substantial upgrades. It must also scale easily to cover untapped subscriber populations. And because the service itself will change over time, the network must be flexible enough to address changes with a minimal effect on provisioning. So, for example, as subscriber interests and viewing habits change over time, channels can be easily removed and added. And when services become more popular because of seasonal demand, additional bandwidth can be easily provided for a few months.

VIDEO TRANSPORT ARCHITECTURE

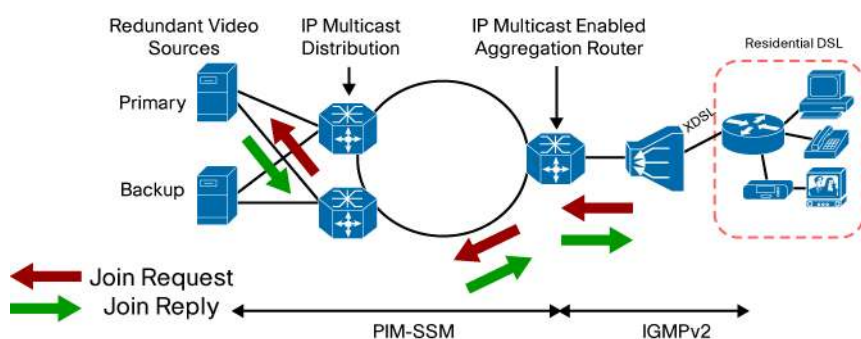
There are a variety of IPTV transport architectures, and each addresses the four logical network segments—core, distribution, aggregation, and access—differently, resulting in tradeoffs. In the reference model in Figure 1, distribution devices are directly connected to core routers, and the aggregation devices are the first devices crossed upstream of the subscriber that has redundant uplinks. In this model, the DSLAM is the access device, and the aggregation platforms are directly aggregating the DSLAMs.

The term *core* in this context refers to the network used to transport video from a few central locations (SHEs) to various metro areas (VHOs). Video over IP can run over a wide variety of transport technologies, and the Cisco IP NGN architecture can adapt to a wide array of transmission technologies. For example, the core of the Cisco IP NGN could be MPLS-based, VPLS-based, or native IP-based. Cisco offers several solutions for transporting multicast over MPLS. It should be noted that at this time, the largest broadcast video over IP networks use IP multicast in their core to distribute channels to the metro areas where they provide service. The reason for this is that IPTV broadcast is an IP multicast application. The QoS can be no better than the quality of the IP multicast implementation transporting it.

The distribution and aggregation segments of the network present unique challenges. These challenges can be addressed in a variety of ways, which require important choices, such as where the Layer 3 edge is placed, which failover capabilities are desired, and how bandwidth is optimized. The main architectural decision lies in the location of the Layer 3 edge, and there are several options, which include:

- **Placing Layer 3 in the distribution routers:** A distribution router can aggregate 150,000 users or more. Placing the Layer 3 edge close to the network core makes it difficult to scale MAC address and ARP cache in the distribution routers. It also prevents efficient protection from source failures.
- **Placing Layer 3 in the DSLAM:** Although this prevents the problems previously mentioned, deploying IP-enabled DSLAMs results in higher cost DSLAMs and more difficult IP address pool management.
- **Placing Layer 3 in the aggregation router:** There are compelling reasons for locating Layer 3 in the aggregation layer between the provider's network and the subscriber's set-top box. Figure 2 is an example of a network with IP multicast enabled in the aggregation routers.

Figure 2. Native IP Multicast in the Aggregation Network

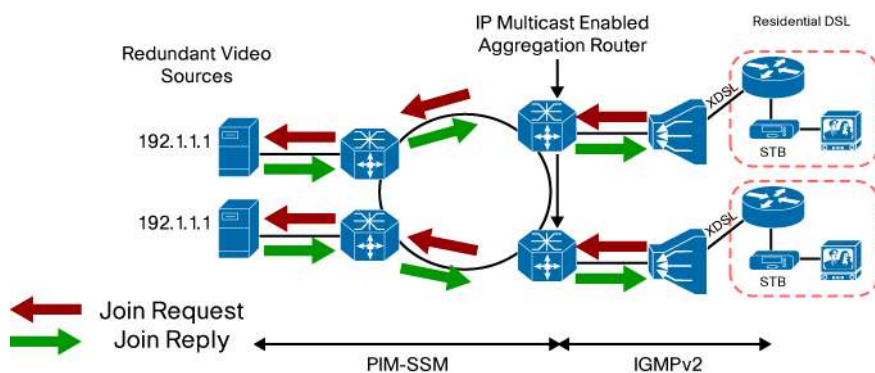


The optimal design (Layer 3 in the aggregation router) places IP multicast capabilities in the aggregation layer, which provides many advantages for service providers, including:

- **Local content insertion:** Deploying Cisco IP multicast capabilities in the aggregation and distribution layers gives service providers greater control so they can determine the video source from which a set-top box receives its streams. IP multicast also makes it easy to add and remove video sources without changing the configuration of the set-top box. For example, the service provider can deploy local ads that are targeted to a specific subscriber community. Local ad insertion is one of the best ways to generate additional revenues from an IPTV service. However, local ad insertion also leads to an increase in the number of IP multicast channels: for example, one channel with 10 unique ad zones suddenly becomes 10 channels of multicast.
- **Admission control:** An IP-aware aggregation router can support integrated admission control to prevent the admission of new video streams that would cause congestion in the network and prevent video packet drops that would lead to an unacceptable user experience.
- **Optimal bandwidth allocation:** Bandwidth is at a premium in this part of the network. IP multicast video streams in Layer 3 are only sent to VSOs and DSLAMs, which must receive the streams. In a Layer 2 aggregation network, all multicast—regardless of destination—is replicated throughout the aggregation network. This is particularly problematic in the case of local ad insertion, which will result in multiple video streams inserted per channel being replicated throughout the aggregation network (from the preceding example, all 10 channels will be replicated throughout a Layer 2 aggregation network even though they are destined for specific neighborhoods). This replication results in potentially gigabits of wasted bandwidth.
- **Asymmetric networking:** Another advantage of having a Layer 3 device to aggregate DSLAMs is the capability of using unidirectional components. In essence, IPTV traffic is overwhelmingly unidirectional. With this in mind, Cisco has developed unidirectional optical components, which allow service providers to deploy an asymmetrical network that provides significant savings on the transport equipment. For the return path, Cisco utilizes generic routing encapsulation (GRE), a Layer 3 encapsulation technique that enables this capability in the aggregation router.


- **Efficient video source selection and protection:** In a Cisco IP NGN, set-top boxes can specify from which source they want to receive a video stream. This more efficient use of bandwidth is enabled by anycast, which allows multicast sources to provide the same content and to share the same IP address. Source redundancy and fast source failover are other advantages of using anycast. Should a multicast source fail, the Cisco NGN automatically and efficiently selects the next closest source available. (See Figure 3.)

Figure 3. Anycast and Closest Source Selection



Having a network that recovers quickly when a failure occurs is an essential part of a successful video over broadband service. Of course, there are many types of failures. A common misconception is to assume that if network elements can protect against a link failure in less than 50 ms, the subscriber is guaranteed good service. For example, consider the typical design in Figure 4.

Figure 4. Typical Design



Assume every link in this network is protected by a sub-50-ms recovery mechanism. If no other protection mechanisms are put in place, the subscriber will suffer seconds of outages from any of the following equipment failures:

- The video source from which the set-top box receives its stream
- The routers used by the video source to reach the STB (R1 through R5)
- The unidirectional link⁴
- The DSLAM or the home gateway

To prevent service disruption, the network must be optimized to allow rapid switching from a server or encoder to another server or encoder. IP multicast mechanisms such as anycast enable this video component redundancy. Application levels implemented in the servers are also available. And some IP video streaming vendors have implemented redundant output Ethernet interfaces. The Cisco IP NGN architecture has been tested and optimized to help ensure that any type of failure will not result in more than 1 second of outage *in any scenario*. A 1-second outage will not cause the subscriber to place phone calls into the support center to complain about the shortage.

⁴ Recovery behavior is dependent on the protection technology used.

Cisco enables *network* redundancy with a wide array of mechanisms that protect against a variety of failures. (See Table 2.)

Table 2. Network Redundancy Protection Mechanisms

Layer	Mechanism	Link	Node	Path
Layer 3	Cisco OSPF enhancements ⁵	X	X	X
	Cisco IS-IS enhancements	X	X	X
	Cisco PIM enhancements	X	X	X
	Bidirectional forwarding detection (BFD)	X	X	
MPLS	Fast Reroute (FRR)	X	X	X
	PW redundancy	X	X	X
Ethernet (Layer 2)	Rapid Spanning Tree Protocol	X	X	
	Resilient Packet Ring (RPR)	X	X	
	802.3ad LACP	X		
	Flexlink	X	X	
Layer 1	DWDM and SONET	X	X	

MULTICAST AND CHANNEL CHANGE

A common misconception is that an IP multicast network causes slow channel changes. Although it is undeniable that stopping the reception of a multicast channel and starting the reception of another one is not instantaneous, the multicast operation typically takes about 50 ms⁶. The reason why this delay is so low is because an aggregation router normally serves thousands of subscribers simultaneously. Because this number exceeds the number of broadcast channels that are offered, the probability of a multicast join request going all the way to the video source (as shown in Table 3) is next to null. In addition, IP multicast is not the main contributor to channel change latency. A video stream is a succession of pictures the set-top box has to display. Pictures are sent to the set-top box sequentially and then grouped in a group of pictures (GoP), which contains a variable amount of pictures (usually 15). The first picture of the group is called the I-Frame. It is the foundation of the GoP in the sense that all the other pictures after it are deduced from it using picture change information in the video coding. If the GoP is made of 15 pictures, there is an I-Frame every 0.5 seconds. When a user changes channel, the set-top box sends a message to the network to request a channel change. As seen later, IP multicast represents less than 10 percent of the overall channel change time. There are many other factors than IP multicast that contribute to the channel change time, and the Cisco multicast network architecture is optimized to minimize the network contribution to channel change.

Table 3. Channel Latency Delay

Channel Change Latency Factor	Typical Latency
Multicast leave for old channel	50 ms
Delay for multicast stream to stop	150 ms ⁷
Multicast join for new channel	50 ms
Jitter buffer fill	150-200 ms
Conditional access delay ⁸	0 msec–2 sec
I-Frame delay	500 ms

⁵ Cisco proprietary enhancements

⁶ Figure obtained on a Cisco 7600 with 100 joins/sec

⁷ DSLAM dependent.

⁸ The conditional access delay is applicable to broadcast channels that are encrypted using a conditional access system that modifies decryption keys periodically and carries updated decryption keys in band in the video stream. The set-top box must wait for the latest set of decryption keys to be delivered in the video stream before it can perform any decoding. The amount of time associated with this delay depends on how often the conditional access system sends updates decryption information in the video stream.

SOURCE SPECIFIC MULTICAST IN DISTRIBUTION NETWORK

All Cisco routers and Layer 3 switches support IP Multicast and Protocol Independent Multicast (PIM) by default. PIM is a routing protocol that makes it possible to establish a multicast distribution tree, which propagates broadcast TV channels from their source throughout the network. Cisco offers PIM-Source Specific Multicast (SSM), which is an enhancement to Internet Standard Multicast (ISM). It allows a receiver (set-top box) to specify which source it wants to receive a channel from. PIM-SSM provides many advantage, including:

- **Optimum distribution tree:** Because the network has a multicast distribution tree that is directly rooted at the video source, the video stream takes the shortest path to reach the set-top box. This is especially important for those service providers who have large video bandwidth requirements.
- **Enhanced security:** The network is more secure because it prevents rogue sources from transmitting content to any group and performs a denial-of-service (DoS) attack against the current viewers of a given channel. The set-top boxes could learn the source IP address of the channels from the electronic program guide (EPG) and transmit it when they join a multicast group. If the set-top box does not support this capability—because of lack of IGMPv3 support—Cisco routers such Cisco 7600 Series Routers are capable of gathering source information from a DNS server on behalf of the set-top box. This allows traditional IP set-top boxes to enjoy the advantages of PIM-SSM without knowing which source the channel should come from.
- **Easy installation and management:** Because the network does not need to keep track of the active sources, it is very easy to add and remove video servers. This capability is very attractive to growing networks, which must accommodate rapid addition and removal of broadcast channels. Another benefit is that video servers or encoders can easily be serviced, removed, and replaced.

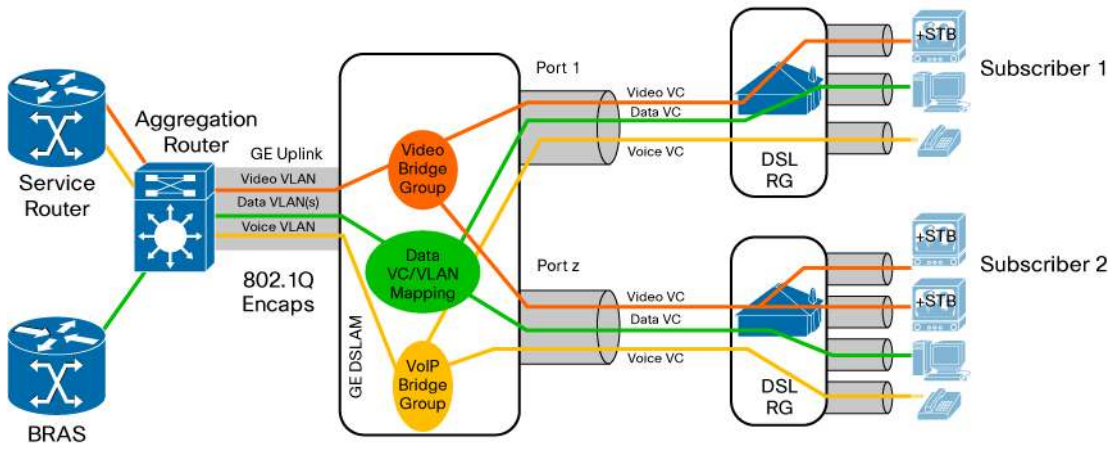
AGGREGATION AND ACCESS LAYER

There are many ways to configure the access layer. Configuration options include:

- Single virtual connection (VC) versus multiple VCs on the DSL link
- Video virtual LAN (VLAN) per service or per user
- High-speed Internet VLAN per service or per user
- Layer 2 versus Layer 3 DSL home gateway
- Using VLANs or Layer 2 priority bits for QoS

There is no one clear best solution, and each combination of the above mentioned options will have advantages and drawbacks. Cisco feature-rich equipment can accommodate virtually all of the design options a provider might choose. However, Cisco has found that the deployment model shown in Figure 5 provides numerous advantages in terms of scaling, ease of management, and security. In this straightforward configuration, one VLAN is used per service. This approach improves scalability because each VLAN can accommodate as many subscribers as the bandwidth allows. Each is independent to the number of users subscribed to the service. Therefore, a DSLAM and its attached subscribers will receive three VLANs.

Figure 5. Aggregation and Access Example



In this model the home gateway separates the traffic and maps each VLAN to its own port(s) inside the house. The aggregation router can be configured with one set of service VLANs or can use a unique set of service VLANs for each DSLAM, an option that provides more security. A single IP address pool can be used for all the subscribers connected to all the DSLAMs aggregated on that router. For example, an aggregation router would typically be connected to about 10,000 users. By using unnumbered IP switch virtual interfaces (SVIs), all the users can be placed in the same subnet, making the management of the IP addresses scalable.

With video, the subscriber is authenticated at the application layer. However, the service provider might want to identify the set-top box. The DSL Forum identifies ways to do this in its WT101 specification, which uses Point-to-Point (PPP) tag or Dynamic Host Configuration Protocol (DHCP) option 82. DHCP option 82 works as follows: when a set-top box is connected to the network, it sends a DHCP request to the DSLAM. The DSLAM snoops the DHCP request and inserts the line ID on which the request has been received into a text field called option 82. This information is placed into the DHCP request and forwarded to the DHCP server, which can match the set-top box with its line ID.

SERVICE SEPARATION

An important aspect of the video transport network for the triple play architecture is how much support the network provides for the different performance and functions required by each type of service. Minimally, the network must provide the ability to meet the delay and drop requirements for each type of service when multiple service types share the same physical link. This capability is inherent in the QoS architecture of the Cisco solution. In addition, the Cisco IP NGN can be configured to provide separate forwarding/routing domains for each service type. This level of separation is very useful when a service provider wants to manage the address space, topology, and IP infrastructure associated with each service type separately. This transport architecture allows traffic associated with different types of service to be aggregated and/or terminated at different sites using different infrastructure components. This architecture also allows traffic associated with Internet access services to be aggregated at a broadband remote access server while traffic associated with video services is terminated using the video infrastructure components.

QOS

Video over broadband is based on IP. Therefore, IP Differentiated Services Code Point (DSCP) bits⁹ become a de facto reference when defining QoS policies. In this architecture, a per-hop behavior (PHB) based on IP DSCP bits is the most scalable way of handling QoS.

⁹ DSCP is a six-bit field in an IP packet header from which as many as 64 classes of service can be created. For more information, please see **Error! Reference source not found.**

Voice, video, high-speed Internet, and voice/video signaling are placed respectively in their own queue. This allows a deterministic end-to-end jitter/delay/loss behavior. Table 4 shows an example of settings a user could choose to configure the access and aggregation elements.

Table 4. QoS Access and Aggregation Configuration

Traffic Class	DiffServ PHB	DiffServ DSCP Value	Queuing Method	Queue Weight
Video Broadcast	Assured Forwarding (AF)	AF 41	Weighted (1)	<ul style="list-style-type: none"> • 80% downstream¹⁰ • 20% upstream¹¹
Video on Demand		AF 42		
Voice + Video Signaling		CS3		
Voice	Expedited Forwarding (EF)	EF	Priority	N/A
Internet Access	Default	0	Weighted (2)	<ul style="list-style-type: none"> • 20% downstream • 80% upstream

When a network element such as a DSLAM is not capable of looking at the IP DSCP bits, the aggregation router is configured to map those settings to the Layer 2 Ethernet 802.1p bits. The DSLAM can place the frames on the appropriate ATM VC, and the ATM SAR function will use proper scheduling to honor the different priorities.

ADMISSION CONTROL

The combination of the traffic generated by video on demand and video broadcast can easily exceed the maximum capacity of any given link in the network, especially in failure scenarios. If this happens, packets from the video queue will start to be dropped randomly. This will affect the viewing experience of the entire subscriber community located downstream of the problem. Providers must have ways to prevent such catastrophic scenarios from happening. Some providers massively overprovision the network so it can accommodate all possible VoD scenarios. However, based on peak concurrency estimates, the bandwidth consumed by all subscribers receiving a VoD stream makes the cost of this approach unpalatable. Cisco offers a wide variety of admission control tools that help ensure a good user experience even when the network reaches its capacity. So providers can better balance the quality of experience delivered with an optimal cost model.

- **Broadcast Video Admission Control:** It is not uncommon for service providers to plan on offering 200 or more broadcast channels. On heavily utilized Gigabit Ethernet DSLAM uplinks, this could mean that broadcast video consumes 75 percent of the available bandwidth by itself.¹² The leftover bandwidth might not be enough to provide high revenue-generating VoD service. In this case, the provider can choose to reduce the amount of bandwidth consumed by the broadcast video (without degrading the user experience) by using a connection admission algorithm for broadcast service. Since the aggregation router runs IP multicast, it has the capability of limiting the number of channels simultaneously sent to the DSLAM. The criteria for doing this could be general popularity or bandwidth used by a specific channel (for example, high definition versus standard definition). It could also guarantee that selected core channels are never blocked.
- **VoD Admission Control:** Video on-demand can be even more challenging. If a VoD service is more successful than forecasted, then the bandwidth consumed between the subscribers and the VoD server complex can become oversubscribed. For example, assume a peak concurrency of 20 percent of all set-top boxes watching on-demand content at the same time. In this case, a central office supporting 10K subscribers would need 2K streams at peak, or roughly 4G to 7Gbps (depending on which codec is used) to support standard definition. Add high-definition on-demand content, and the bandwidth demand can be quite large. What happens when demand rises above the peak concurrency assumed when designing the network and/or VoD service capacity? If a particular

¹⁰ The downstream queue weight for video is a recommendation assuming that all video traffic will consume no more than 70 percent of the physical link bandwidth for the link being configured. If the expected ratio of video traffic to total link bandwidth is significantly less, then a lower queue weight can be used.

¹¹ The upstream queue weight for video takes into account only voice + video signaling since video broadcast and video on demand traffic is unidirectional. The actual value used for the queue weight can vary depending on the expected ratio of signaling traffic compared to total link bandwidth.

¹² Considering 200 MPEG2 channels encoded at 3.75 Mbps

VoD server complex cannot service a particular request, it can get rerouted to another VoD server complex that has available capacity. If the network capacity over a set of link(s) is exceeded, then allowing too many VoD sessions to be set up could cause a high packet drop rate for all the video streams. This is known as the 1001st stream problem: the admission of one stream results in a large number of subscribers experiencing degradation of the video and audio signals. Furthermore, in the event of a network failure that occurs simultaneously with the peak concurrency window, it is possible that all users will perceive an outage. One way of avoiding a widespread perceived outage is through the use of integrated admission control. If a video session cannot be supported due to oversubscription anywhere in the network or service, this integrated control would deliver a “could not be serviced at that time” signal to the requesting set-top box. Although no one likes a busy signal, the possibility of mass degradation of the VoD service is much worse. It is easy to imagine that the application knows that there are streams that will be ending soon and thus could provide more sophisticated busy messages. This would give the subscriber more choices for delayed start of the VoD stream, alternative video service offerings, commerce opportunities, and so on. In the event of a failure, the admission control system could be configured to drop all free VoD streams *before* any pay VoD streams.

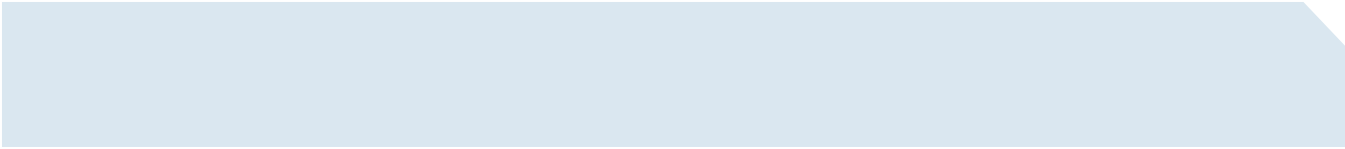
- **Integrated On-Path and Off-Path Admission Control:** The VoD admission control solution must be able to take into account complex network topologies. This would include topologies that have redundant and load-sharing paths in the transport network as well as access link utilization and/or business policies that can be enforcing other types of constraints on the subscriber’s service. To do this, the network’s routers, in coordination with policy managers and on-demand server/managers, need to collectively perform an admission control function, called Integrated VoD Admission Control. Cisco has developed an in-path method to perform admission control for the complex core and distribution network topologies found in NGN networks designs. It utilizes RSVP for in-path signaling, sent by the VoD server or a component prior to starting the VoD session. The RSVP message traverses the same exact path the VoD session will use, thus tracking—in real-time—any changes in the complex network topologies in the core and distribution layers. Along the path, Cisco IP routers perform a bandwidth accounting function, either allowing the session or denying it if bandwidth is not available for that VoD stream. What makes this in-path admission control possible is having IP or Layer 3 routing, present on every network element from the video on-demand server complex to the aggregation router in central office. Unlike previous use of RSVP in voice networks, the number of video streams in the Metro will likely only be in the thousands, so scaling is not an issue.

To prevent a video stream from being sent to a set-top box (if the access link to a subscriber’s home doesn’t have enough capacity to carry the stream), the VoD server or a network component in the path mechanism will send a request to an off-path component. This might be a policy server that is keeping track of the access network via bandwidth accounting. The policy server can check to see if the access link has enough unused bandwidth or has any business policies that might or might not allow the stream to be supported. The policy server will then allow the session or deny it at that time. The combination of in-path admission control with an off-path policy server at the edge is the most reliable and efficient way for an admission control solution to decide on whether or not a new VoD stream should be allowed to reach a specific subscriber. Streams will be denied if the business rules in either approach are triggered.

The VoD admission control and related bandwidth accounting added to a DiffServ QoS allowed to be set up can meet the required 10^{-6} packet drop rate. Combining the different admission control technologies with DiffServ QoS allows Cisco to offer a simplified queuing strategy and have all the video traffic shared in the same queue thus avoiding complex and very costly centralized, hierarchical queuing strategies.

CONCLUSION

Tomorrow’s video services will be delivered over IP. Cisco IP leadership and IP NGN innovations raise the performance of the entire IP video and triple play system by optimizing video transport with unique features developed by Cisco. Cisco offers the industry’s first IPTV network solution that uses Service Separation to optimize QoS and forwarding based on service type. Service Separation uses different



approaches for transport services versus managed application services such as video. This architecture drives Layer 3 intelligence into the aggregation network and delivers a number of key advantages:

- **Optimal source selection and bandwidth efficient delivery of IPTV services:** Source Specific Multicast (SSM) in Cisco aggregation routers allows for optimal source selection and bandwidth efficient delivery of IPTV services while providing a level of simplicity and inherent security that is greater than delivering multicast routing over Layer 2 aggregation networks.
- **A high-quality viewing experience:** by delivering integrated admission control and combining the benefits of on-path and off-path solutions, Cisco provides the control required to prevent the oversubscription of video. This helps ensure a consistent, high-quality viewing experience.
- **Comprehensive network protection:** advanced resiliency mechanisms (such as Anycast) with failover of mission-critical broadcast sources provides comprehensive consistent recovery times for any type of failure including links, nodes, video servers and encoders.

With innovative, intelligent networking capabilities like these, it's not surprising that more than 10 million subscribers in North America alone now access their entertainment services over Cisco IP NGN infrastructures.

REFERENCES

1. Multicast Virtual Private Networks Concepts
http://www.cisco.com/en/US/tech/tk828/technologies_white_paper09186a00800a3db6.shtml.
2. Internet Protocol Multicast <http://www.cisco.com/en/US/docs/internetworking/technology/handbook/IP-Multi.html>.



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

European Headquarters
Cisco Systems International BV
Haarlerbergpark
Haarlerbergweg 13-19
1101 CH Amsterdam
The Netherlands
www-europe.cisco.com
Tel: 31 0 20 357 1000
Fax: 31 0 20 357 1100

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

Asia Pacific Headquarters
Cisco Systems, Inc.
168 Robinson Road
#28-01 Capital Tower
Singapore 068912
www.cisco.com
Tel: +65 6317 7777
Fax: +65 6317 7799

Cisco Systems has more than 200 offices in the following countries and regions. Addresses, phone numbers, and fax numbers are listed on the **Cisco.com Website at www.cisco.com/go/offices.**

Argentina • Australia • Austria • Belgium • Brazil • Bulgaria • Canada • Chile • China PRC • Colombia • Costa Rica • Croatia • Cyprus • Czech Republic
Denmark • Dubai, UAE • Finland • France • Germany • Greece • Hong Kong SAR • 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 • Scotland • Singapore • Slovakia • Slovenia • South Africa • Spain • Sweden
Switzerland • Taiwan • Thailand • Turkey • Ukraine • United Kingdom • United States • Venezuela • Vietnam • Zimbabwe

Copyright © 2006 Cisco Systems, Inc. All rights reserved. CCSP, CCVP, the Cisco Square Bridge logo, Follow Me Browsing, and StackWise are trademarks of Cisco Systems, Inc.; Changing the Way We Work, Live, Play, and Learn, and iQuick Study are service marks of Cisco Systems, Inc.; and Access Registrar, Aironet, BPX, Catalyst, CCDA, CCDP, CCIE, CCIP, CCNA, CCNP, Cisco, the Cisco Certified Internetwork Expert logo, Cisco IOS, Cisco Press, Cisco Systems, Cisco Systems Capital, the Cisco Systems logo, Cisco Unity, Enterprise/Solver, EtherChannel, EtherFast, EtherSwitch, Fast Step, FormShare, GigaDrive, GigaStack, HomeLink, Internet Quotient, IOS, IP/TV, iQ Expertise, the iQ logo, iQ Net Readiness Scorecard, LightStream, Linksys, MeetingPlace, MGX, the Networkers logo, Networking Academy, Network Registrar, Packet, PIX, Post-Routing, Pre-Routing, ProConnect, RateMUX, ScriptShare, ScriptShare, SlideCast, SMARTnet, The Fastest Way to Increase Your Internet Quotient, and TransPath are registered trademarks of Cisco Systems, Inc. and/or its affiliates in the United States and certain other countries.

All other trademarks mentioned in this document or Website are the property of their respective owners. The use of the word partner does not imply a partnership relationship between Cisco and any other company. (0601R)

Printed in USA

C11-343474-00 03/06