White Paper

Assuring Quality of Experience for IPTV

Prepared by

HEAVY READING

www.heavyreading.com

TABLE OF CONTENTS

LIST OF FIGURES

I. Executive Summary

The year 2006 marks a time in which the IP television (IPTV) market is moving into the critical second phase of large-scale commercial deployments across many regions. This may vary between carriers and geographies, but as a whole, service assurance and quality of experience (QOE) largely define this evolution. Each phase is tightly coupled with ongoing underlying technology evolution, content acquisition, and scaling the overall number of IP video subscribers.

- **Phase I:** Prove technical viability of technology, architecture, and basic service delivery to match existing cable and satellite TV service offerings.
- **Phase II:** Deliver service assurance and QOE guarantees, increase personalization, and deliver any service, any time, in an effort to ensure and grow take rate.
- **Phase III:** Increase service differentiation and integration to achieve blended services and interactive TV on a large scale.

Assuring QOE for IPTV is rapidly becoming a top priority among vendors and service providers as the IPTV market evolves into Phase II, in which services are commercially deployed on a meaningful scale. IPTV is about the television business, and offering a differentiating and compelling value proposition is a key to success. In order for service providers to achieve target video take rates, the QOE of IPTV must meet and exceed the services cable and satellite providers are currently providing. Large and small service providers are either planning for or putting the technology in place to meet the requirements surrounding service assurance and QOE.

There are many factors influencing the challenge of achieving a high QOE in IPTV deployments. The mix of broadcast TV versus video on demand (VOD), network-based personal video recorder (nPVR), and other unicast video services is relatively unpredictable, which affects the ratio of multicast and unicast network traffic, and ultimately the design of the network architecture. The addition of high-definition (HD) content to the network drives greater bandwidth requirements and the need to assure the continuous availability of that bandwidth for the purpose of achieving flawless QOE. VOD usage patterns, as well as the proliferation of set-top boxes (STBs) and other appliances in the home, will dictate the optimal location of content placement in the network, which also has an effect on traffic patterns and bandwidth needs.

These issues should and will be transparent to the subscriber, who is simply expecting an always-on service offering. To that end, providing such QOE isn't necessarily simple, but it's a critical factor for the success of IPTV. It also is achievable with the right technology and network architecture as a foundation.

Intelligent network design, with an emphasis on avoiding congestion, is essential to address QOE issues before they arise. Peak service concurrency rates, based on a number of assumptions and real-world scenarios, will impact the network differently in the first mile than in the second, third, or fourth (e.g. as part of a congestion-avoidance and admission-control strategy). Blindly throwing bandwidth at the problem to overprovision the network is prohibitively expensive, so telecom equipment providers are proposing new and smarter means to right-size network capacity by gaining a better understanding of traffic patterns and usage trends.

The second phase of the IPTV market will drive new network requirements that must be addressed by industry-leading vendors in next-generation equipment. Robust service control and QOE assurance must be put in place to attract initial subscribers, ensure their comfort with the new IP-based video service, and further increase the average revenue per user (ARPU) over time. Optimizing the cost structure without compromising the service experience will enable them to maximize the revenue contribution of each paying customer.

This paper examines the critical challenges, required technology and features, and optimized network architecture necessary to achieve these goals.

II. Introduction

As IPTV evolves from its infancy into a large worldwide market, it's important to understand the factors that are shaping its evolution. This introduction explains the importance of IPTV/VOD services, explores how the nature of these services creates challenges for network architects – and then examines the network innovations proposed to overcome those challenges.

2.1 The Importance of IPTV/VOD

The time for debating whether or when service providers will roll out IPTV services has passed. It is happening now on a worldwide basis. Time to market and the ability to deliver a differentiated QOE are critical to seizing this window of opportunity. Regardless of the timing of any specific rollout, the drivers are the same and generally agreed upon industry-wide.

Service providers are experiencing intense competition to retain customers, attract new customers, and increase ARPU, which has led them to embark on plans to deliver comprehensive IPTV/VOD services. Next-generation broadband access, aggregation, and edge networks will be the cornerstones of IP video delivery, but robust service assurance and control solutions to ensure QOE will dictate the success or failure of these endeavors.

Figure 1 outlines the high-level inflection points guiding the evolution of the industry. As the diagram suggests, the addition of differentiated services largely defines the overall rollout of IPTV and translates into increased ARPU and customer loyalty. Cost-effectively scaling the number of subscribers and concurrent service usage over time is a primary goal that will result naturally, provided service assurance and QOE expectations are met for all services defined. Large-scale IPTV networks now measure on the order of hundreds of thousands of subscribers; by 2010, these numbers are expected to grow to tens of millions.

Figure 1: IPTV Inflection Points

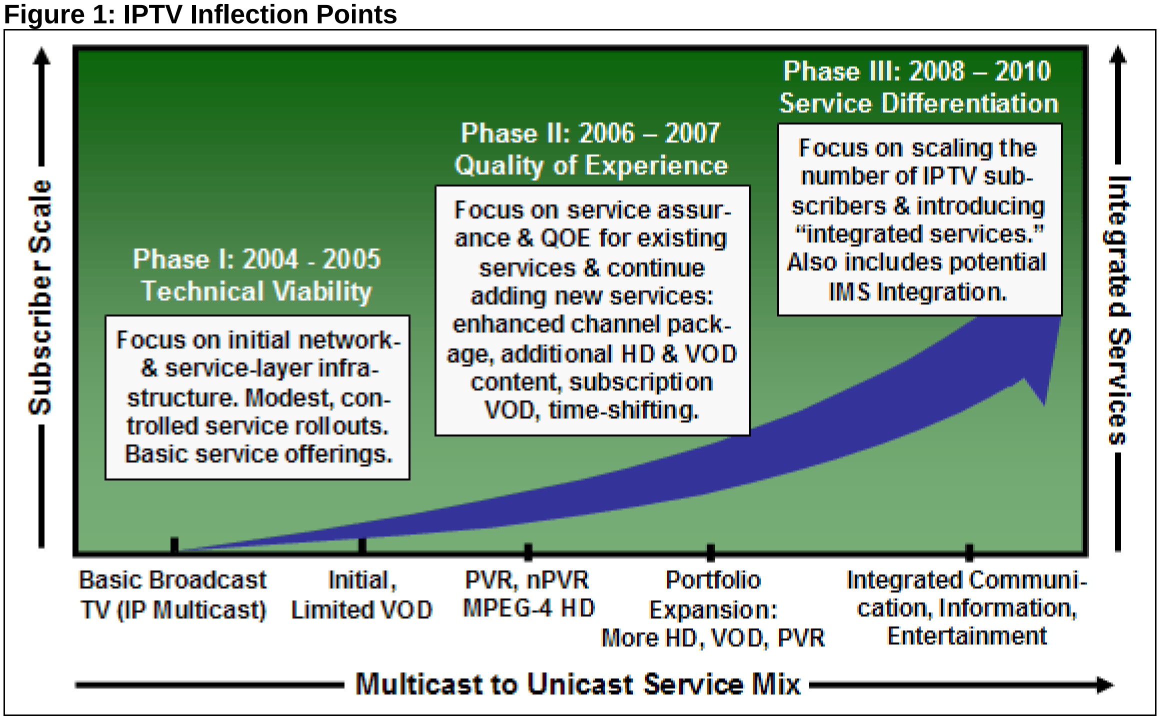

While initial take-rate targets of IP video services are modest – less than or equal to 10 percent of the service provider's overall subscriber base – a positive user experience for early service adopters is vital to establishing momentum and building brand equity via word of mouth. Longer-term take-rate targets are on the order of 30 percent, depending on the carrier. Provided carriers are successful in achieving such rates, their customer churn is expected to drop significantly. Assuming a carrier has 20 million broadband subscribers, a 30 percent take rate would equal 6 million subscribers. Assuming a modest ARPU increase of $50 per subscriber per month for these 6 million customers, this creates additional revenue of $3.6 billion per year for the service provider.

2.2 The Nature of IPTV/VOD Services

IPTV/VOD services are inherently resource-intensive, with unpredictable demand fluctuations. Although service providers and vendors are building networks in support of some concurrency assumptions, with additional solutions for dealing with fluctuating peak-level demands, the challenge remains paramount. The issues that define the nature of IPTV/VOD services and lead to unpredictability and resource-intensiveness include:

- **VOD/Unicast Concurrency:** VOD session concurrency has a direct effect on the amount of unicast network traffic, and is therefore a major variable in network design and reliable service delivery. Initial peak concurrency rates will likely be about 10 percent, although as VOD content and features such as nPVR are added, unicast video is expected to grow to more than 20 percent in the short term. This excludes unicast traffic as part of the channel-change mechanism, which can be on the same order of magnitude as VOD/nPVR during commercial blocks.
- **Broadcast Channel Concurrency:** The number of broadcast channels watched concurrently affects multicast replication throughout the network. Internet Group Management Protocol (IGMP) snooping can be deployed to block unwatched channels by pruning the related multicast tree branches, thereby optimizing bandwidth usage in the first, second, and third mile, based on actual viewing patterns.
- **HD Content Growth:** HD content grows in direct proportion to increased network bandwidth, whether broadcast or unicast. Even with advances in compression technologies, an MPEG-4 HD stream will consume a maximum of 8 Mbit/s. Many service providers believe the amount of HD will be a significant differentiator, making the addition of as much HD content, as quickly as possible, a top priority.
- **STB Proliferation:** The number of STBs per household – and features such as multi-channel viewing for picture-in-picture and multi-angle viewing – further compounds bandwidth requirements as more channels are accessed concurrently per household. North American service providers are projecting somewhere between two and three STBs per home on average.

Additional elements, such as trick-play commands and nPVR services, will add to the relative unpredictability of resource requirements. Network-based intelligence and quality-of-service (QOS) mechanisms such as Hierarchical QOS (H-QOS) need to provide the elasticity and efficiency to handle these dynamic real-time traffic changes based on policies defined (Committed Information Rate, Peak Information Rate, etc.).

The VOD server deployment architecture, be it centralized or distributed, will also play an important role in the availability of any particular VOD service. And the location of broadcast content insertion and replication points affects multicast distribution efficiency and cost. On top of all this, there are service-level agreements (SLAs) that will need to be maintained: For example, ordering a VOD asset from the network should not affect broadcast content on a separate TV in the same home, and changing from a standard-definition (SD) channel to an HD channel should not affect other services, either.

With all of these things to take into account, as well as the issues of serving a large subscriber base, accurate understanding of network resource availability and usage and performance details from IPTV middleware will be critical in achieving service assurance.

2.3 Service Delivery Challenges

Cable and satellite TV providers have set the bar extremely high in terms of quality and availability. Operators must consider many different peak-demand scenarios, which are identified and discussed below. But before getting into that level of detail, we must identify the primary service-delivery challenges that are common to them all. **Figure 2** outlines these challenges.

Figure 2: IPTV/VOD Service Delivery Challenges

CHALLENGE	DESCRIPTION & SOLUTION
"Always On" Service Expectations, Hard QOS Guarantees	The user expectation for "always on" services is based on current cable and satellite TV performance and drives end-to-end requirements from the IP STB to the end-to-end service delivery infrastructure and IPTV middleware. This entails everything from picture quality to channel-change performance to VOD availability at the time it's ordered, no matter what else is happening in the home or the network. Enhanced high-availability features in network equipment and many of the other items listed below will contribute to an always-on service experience.
Network and Service Capacity Planning	Network architecture and associated capacity planning are critical to optimize service-delivery cost, address multi-dimensional scaleability, and provide QOE guarantees. Assuring congestion-free video transport across each link and node in the IPTV network is key to service quality and performance. Thus, it is vital to understand trends in capacity utilization and engineer additional capacity in time for anticipated demand, while maintaining plenty of headroom for growth. Additional provisions, such as IGMP snooping, distributed policy enforcement, IP multicast replication, flexible content insertion at the most economical points, non-stop service capabilities, and video admission control (VAC), help to cost-optimize an IPTV network without having to overprovision bandwidth.
Network Congestion Avoidance	Network congestion avoidance and network capacity planning go hand in hand and are essential to avoid resource contention and minimize congestion. The first essentially depends on the second, to the extent that under-dimensioning the network is bound to result in congestion. While network oversubscription is the norm for traditional high-speed Internet and is acceptable for best-effort applications, support of deterministic H-QOS is needed to guarantee flawless QOE for triple-play, especially broadcast TV, VOD, and voice over IP (VOIP). This is not just a matter of putting in more capacity, because optimizing video content insertion by placing popular VOD content closer to end users helps avoid network congestion in the end-to-end video delivery path. It also encompasses tracking of capacity utilization and resource availability state changes (e.g., DSL training-rate variability in the first mile or potential fiber cuts in the metro aggregation network affecting second- or third-mile bandwidth) to anticipate and remedy potential resource contention before it occurs.
End-to-End QOE Measurement and Assurance	While proper capacity dimensioning and video admission control mitigate the risk of service quality or availability degradation as a result of network congestion, the only way to fully insure flawless perceptual QOE is by adding components to measure and verify IP video and audio quality, disturbance rate, and channel-change delay. Many vendors supply standalone solutions to measure and monitor video traffic, while others are integrating and/or partnering to provide an end-to-end solution taking advantage of network diagnostics and performance statistics, IPTV client measurements, middleware reporting, and MPEG-2/4 traffic analyzers and video monitors.

CHALLENGE	DESCRIPTION & SOLUTION
Video Admission Control	VAC is rapidly becoming an important new requirement for telecom equipment. It is most applicable to control dynamic admission of unicast VOD sessions, as VOD is a pay-per-use service and also the most bandwidth-intensive (i.e., the most likely to cause resource contention when unchecked). Multicast VAC is being proposed for cases where more broadcast channels are being offered than can be concurrently watched. Multicast VAC prevents bandwidth issues when changing channels, for example when bringing up a new channel (one that has not been previously watched). This may potentially result in some of the least-watched channels being unavailable, while ensuring that more popular channels are available with the expected quality. VAC is essentially a "safety valve" to ensure flawless QOE for video streams by preventing additional streams from entering the network when remaining capacity is insufficient to support them (e.g., extreme VOD, concurrency peaks with a new movie release, or reduced capacity due to failures). While VAC is a compromise between service quality and availability, the need to deny service requests due to insufficient network capacity should be an exception, not the rule. However, since the possibility of network congestion can never be ruled out, the implementation of an effective admission control strategy is an important issue involving network infrastructure, policy control systems, and IPTV middleware.
Additional Challenges	Although they lie outside the scope of this white paper, there are many additional service delivery challenges for IPTV service providers. Among them are regulatory issues, franchise requirements, working with local permitting agencies, vast and unique content rights, lack of industry standards, and the integration of multiple IPTV hardware and software service-delivery components.

III. Technical Requirements

Now that we've set up the discussion and outlined the primary challenges involved in achieving a high-quality IPTV experience, let's examine the technical requirements for congestion avoidance and service assurance in depth. The following two sections outline the guiding assumptions and the high-level requirements for achieving these goals. The final section of this white paper outlines best practices for implementation and network architecture.

3.1 Guiding Assumptions

Following are the critical guiding assumptions of this white paper. They are listed in no particular order of priority; they should be considered holistically as a set of interrelated guidelines for meeting service assurance expectations.

1. It is important to differentiate between IPTV and Internet TV. IPTV is a video service supplied by a telecom service provider that owns the network infrastructure and controls content ingestion and distribution over the broadband network for reliable delivery to the consumer (generally to the TV/IP STB). This is essentially a private/"walled-garden" network controlled by the service provider. Internet TV, which is rapidly emerging in parallel, consists of content sourced from anywhere on the Internet that can be streamed and/or downloaded by the user (generally on a PC). Both IPTV and Internet TV are delivered over a broadband connection, albeit with different levels of bandwidth, control, and QOS.

2. IPTV, when used as a standalone acronym, includes broadcast TV, VOD, and other forward-looking video-integrated data and voice services.

3. IPTV must meet and/or exceed the performance, availability, and QOE metrics currently being delivered by cable and satellite TV. Anything less will be a non-starter; hence the enormous importance of service assurance in IPTV networks.

4. Traditional best-effort IP networks are inadequate to support robust and scaleable IPTV services with an acceptable QOE. Many of the technical requirements discussed below are enhancements to best-effort IP networks for enabling reliable IPTV service delivery.

5. Formal SLAs must be defined by the service provider. To enforce policies appropriately, the supporting network architecture and policy-control solution needs to understand service priority under different usage scenarios. For example, broadcast TV will likely be assigned the highest priority, such that a VOD service competing for network resources does not disrupt the service contract associated with basic broadcast TV or other VOD sessions already in progress. The granularity of SLAs goes much deeper, and can differentiate between free and paid VOD, time shifting, and effect on other triple-play services.

6. Frequent "busy signals" and/or service denials are unacceptable and impair customer loyalty and recurring revenue opportunities for both multicast and unicast IPTV services. Service denial may be impossible to avoid in extreme peak-usage scenarios, but such cases should be rare and supported by a graceful message informing the subscriber why the service is unavailable or offering viable alternatives (e.g., movie content may still be available in SD format or downloadable for later viewing, if the STB has a hard drive). A tight coupling between the network layer and IPTV middleware will be important to clearly present this information and options – as defined by provider policies – to subscribers.

7. Grossly overprovisioning network bandwidth is not an acceptable solution to support congestion-free delivery of premium video services. Such a solution will not meet the cost targets of the service provider nor enable it to achieve a profitable business model. Instead, the operator must purpose-build a network that caters to video requirements.

3.2 High-Level Requirements

The following sub-sections identify the most significant high-level technical requirements the network must support so that an IPTV service provides a consistently high-quality user experience.

Differentiated & Deterministic QOS

Best-effort IP networking is simply not good enough for IPTV/VOD services. Differentiated QOS support is needed because a variety of services with conflicting QOS needs share the same network. Deterministic QOS is a clear requirement not only for the variety of IP video services, but also other triple-play services available to the subscriber, such as VOIP. That means oversubscription for these classes of services is strictly controlled. A high take rate of IP video subscribers will not allow for the traditional oversubscription models used for basic high-speed Internet.

Support of H-QOS is particularly important so that each service in the network and each subscriber to each service is guaranteed the bandwidth required. If provisioned bandwidth is temporarily not needed for a particular service, H-QOS enables the unused bandwidth to become available for other services (e.g., unused VOD bandwidth can be used by high-speed Internet), which results in bandwidth efficiencies. One major benefit of H-QOS is the resilience it provides against the peak-demand conditions of a given service. Another advantage is that it provides a reasonable margin of error for applying VAC, as video traffic is extremely bursty.

Intelligent, End-to-End Service Control & Policy Enforcement

Policy control and enforcement remains a critical requirement for IPTV and other triple-play services. There are several architectural choices for implementing policy control. In each case, the control is centralized, but the enforcement may be more distributed, depending on the network architecture. *Heavy Reading* has done extensive research on centralized and decentralized policy control and network architectures, including the report entitled **IP Video and the New Broadband Edge** (Vol. 3, No. 20, December 2005).

Regardless of the network model, intelligent policy-control capabilities are growing more important relative to other requirements, such as multicast/unicast VAC. Policy control is also becoming more wide-ranging in scope, with devices such as residential gateways and application middleware being integrated more tightly into the end-to-end solution. Requests for proposals have been published positing tighter integration between policy-control solutions and customer premises equipment, pushing policy onto next-generation multimedia home network devices. The DSL Forum specification TR-069 and additional specifications and working texts, including TR-111 and WT-135, extend even further into the multimedia home network of the future.

Non-Stop Service Delivery

The access, aggregation, and edge networks have all undergone significant high-availability improvements over the past several years in the interest of supporting today's real-time services. Further availability enhancements continue to hit the market, including IGMP stateful switchover, which protects IGMP multicast groups and joins/leaves throughout the network in the event of a failure scenario. Likewise, numerous non-stop routing and service-operation techniques have been delivered to support a variety of mission-critical residential and business IP services.

Applying non-stop routing intelligence to multicast protocols is an obvious next step benefiting multicast IPTV services. This is critically important when one considers the investments being made to get exclusive broadcast rights for live events such as the World Cup, the revenues from ads during commercial breaks, and the enormous number of viewers such content will attract to a given IP multicast tree instance – not to mention the fury of viewers should there be a glitch or outage at a decisive moment. Clearly, these are critical feature enhancements for broadcast IPTV services that will form the baseline of carrier service offerings.

Optimizing Broadcast TV Delivery

Support for IGMP and IGMP snooping are table stakes for multicast IPTV service delivery, as it is part of the IPTV channel-selection mechanism. By means of IGMP join and leave messages, the STB indicates what IP multicast group it wishes to join in order to receive broadcast content. Channel popularity will dictate the number and location of IP multicast streams that are replicated and forwarded at various points throughout the network. IGMP snooping, and its location of implementation, is important to optimize network bandwidth for broadcast TV services, because it enables the network to forward only those channels subscribers have requested.

This capability will continue to grow in importance as more channels and HD streams are added to the network. The most popular channels will usually be "always on" and multicast to a point closest to the subscriber in the access network. For the less popular channels, or channels not being watched, IGMP snooping can block them from joining a particular multicast group, thereby saving bandwidth in the access and aggregation network. Also, by recording IGMP snooping events, a service provider can track channel viewer density to a fine-grain level. This is not only invaluable for marketing and advertisement purposes, but also enables them to plan the amount of bandwidth that must be reserved for broadcast video delivery. Features such as IGMP snooping and others ultimately impact the capital costs of network equipment and available bandwidth resources to better guarantee service assurance.

Optimizing Video Content Insertion

Supporting a network architecture and solution that will allow for flexibility in content insertion, be it broadcast TV or VOD, will be important in optimizing network bandwidth and cost, and ultimately in improving service assurance. Broadcast content from different sources will undoubtedly be inserted at different points in the network, affecting channel availability to specific subscriber groups and multicast efficiency.

For example, broadcast content from regional and local stations is likely to be inserted and viewed in a video hub office within a metro area; whereas national channels such as ABC will be inserted at the super head-end office. Local insertion of ads is another possible way to cut costs. As VOD grows more prevalent and user patterns are better understood, video asset placement will be an internal operation that will enable the carrier to optimize cost and service delivery. The network infrastructure, the VOD server software, and the IPTV middleware will be tightly integrated and centrally managed to achieve successful distributed VOD architectures.

Accurate Intelligence on Resource Availability

Service assurance and QOE are dependent on bandwidth availability in the network. Providing the tools in the management infrastructure and network elements to accurately report available bandwidth will be critical to support IPTV/VOD service requests. This entails consistently keeping up with bandwidth changes at various links and potential contention points in the network and keeping detailed records so that services can be admitted and policies can be enforced appropriately based on SLAs.

To complete the solution, network infrastructure must be tightly integrated with IPTV middleware so that actual service requests can be granted or gracefully denied. The latter is the purpose of VAC (see below), an emerging requirement that is shaping differentiated vendor equipment as we speak. Some vendors support this today, while others are planning to deliver it in the near term.

Video-Quality Measurement & Monitoring

Video-quality monitoring is another requirement heating up the industry. The service provider needs to measure video traffic and monitor the performance of various network-layer and service-layer elements, including those in the home network, such as the remote gateway and IP STB, as well as video head-end equipment such as MPEG encoders. These solutions are being delivered by niche vendors with a specific focus in this area and by large incumbents as part of an end-to-end IPTV solution. Solutions could be made up of integrated third-party products and/or products developed fully in-house.

Vital to a robust video-quality measurement and monitoring solution are tools including detailed network-layer diagnostics and reporting; network performance and fault management; MPEG-2/4 analyzers and video monitors; TR-069-based home network-management tools, and client QOE measurement capabilities.

Multicast VAC

Multicast VAC is applicable for broadcast IPTV services to ensure that all channels requested in the broadcast channel lineup (SD or HD) can be viewed without congestion within the limits of the available bandwidth, provided it is part of the user's service package. First- and last-mile bandwidth limitations may exist in many IPTV deployments, and service providers typically limit the amount of SD and/or HD channels a multi-STB user can access. This can be controlled by the IP STB version supported in a given subscriber's home and by the use of stream tokens that give an STB the right to receive video streams up to the total number that may be concurrently used by a home. For the master TV in the home – assuming that's the location of the HD-capable IP STB – seamless channel zapping between SD and HD content should not be an issue and is not subject to further admission control decisions for the STB itself.

Multicast VAC is not related to these first/last-mile admission-control mechanisms, but is tightly coupled with how and where IGMP and IGMP snooping are supported in the network and where content is inserted. As a result, it can control the maximum amount of broadcast channels that can be brought up to a broadband access node simultaneously. Intelligent multicast VAC capabilities in the network are useful to maximize the number of channels that can be offered to end users while assuring that concurrent use of channels stays within pre-defined capacity limits, thus safeguarding QOE. However, blocking of multicast channels should occur only in exceptional situations, as broadcast TV is a pre-paid service for which there are delivery obligations, and the users most likely affected would be those with the largest channel lineup (i.e., the top percentile of ARPU contributors).

Unicast VAC

Whereas multicast VAC applies to broadcast IPTV services, unicast VAC applies to VOD, nPVR, and unicast multimedia services such as video telephony. Unicast VAC is responsible for checking potential resource-contention issues in the end-to-end network, including the VOD service overlay, for available resources and interworking with the IPTV middleware to grant or deny service. Since denying a service request is a bad user experience and a lost revenue opportunity, the latter should be kept to a minimum with a robust and cost-optimized network design.

The availability of a particular VOD service could be affected by end-to-end network bandwidth availability or by resources specific to the VOD infrastructure itself, including server placement and storage/streaming performance. Unicast VAC is tightly coupled with SLAs because unicast traffic will be generated in many forms, including traditional paid VOD, subscription VOD, free VOD, nPVR playback, and trick modes. Granting access to a variety of different VOD services will depend on how the service provider prioritizes the availability of those particular services.

Both multicast and unicast VAC must be coupled with network resource intelligence to obtain an accurate view of resources in the service delivery path in order to make an informed admission-control decision, as well as with middleware service delivery platforms in order to effectively enforce and communicate this decision.

IV. Implementation & Network Architecture

Now that the overarching goals, primary service challenges, and associated network requirements are understood, it's important to examine best practices for implementation and network architecture design. Depending on the footprint and specific service offerings provided by the carrier, there may be several network architecture options available. Also, since the service offering and traffic patterns will evolve over time, it is important that the network architecture and implementation be able to adapt to changes in a cost-effective manner.

To that end, this section provides insight into a competitive network implementation approach being deployed in the industry to address many of the usage scenarios that will potentially stress IPTV service delivery, while maintaining a competitive cost structure and providing flawless QOE.

Figure 3 provides a high-level view of an end-to-end IPTV network architecture in support of these goals. It should be understood that the details of such a network and its implementation are far more extensive, and beyond scope of this white paper.

Figure 3: End-to-End IPTV Network Architecture Example

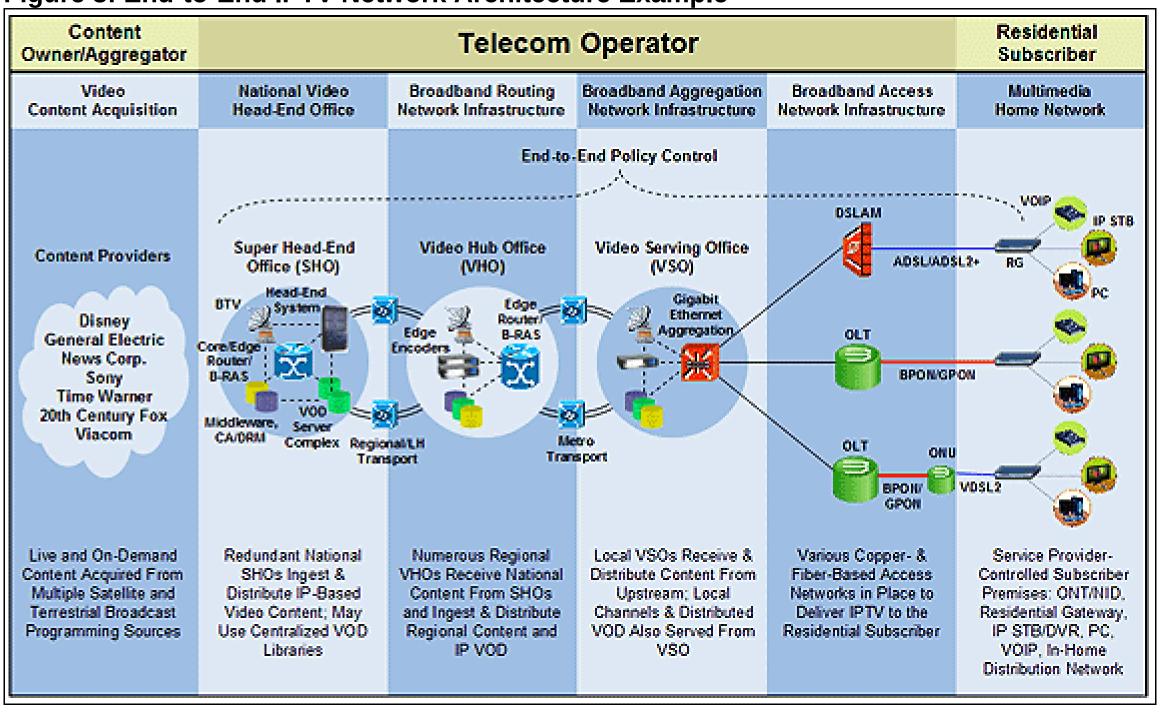

4.1 Peak Usage Scenarios Affecting Network Design

Although the number of resource-intensive usage scenarios is vast, there are several that guide network implementation and design. Understanding and planning for peak-usage scenarios will go a long way toward achieving service assurance and a high QOE for the user. Many of the technical requirements previously discussed work together in designing for such scenarios.

Extreme Single-Channel Concurrency – Live Events, Breaking News

This case can best be characterized as the "Super Bowl Effect" – when the majority of subscribers are watching the same broadcast channel. Considering this an extreme condition may in fact be false, given that the demand for the expected channel may be known ahead of time and the need to replicate other channels at the network edge during this time will likely be limited. In this scenario, architectures using IGMP snooping with distributed multicast replication extending into the access and aggregation networks will realize very high bandwidth efficiency, as few different channels will be watched and forwarded. Highly available multicast routing protocols will be critical, as many viewers are subscribing to few channels, which results in large multicast trees. The need to reconverge multicast trees in case of failure conditions would take relatively long, affect many viewers, and is best avoided.

Total Broadcast-Channel Concurrency – Weeknight Prime-Time

This represents a weeknight prime-time viewing scenario, in which many users ask for many broadcast channels concurrently. This may be a peak-demand scenario for more broadcast TV channels, even as the majority of subscribers will be accessing the most popular channels. Highly available multicast routing will be critical to support and preserve the higher frequency of channel changes in these conditions. Multicast VAC may apply in extreme conditions to set an upper limit, in case too many channels are requested concurrently. As an alternative to denying access to less-frequently watched channels, it is possible to preempt some bandwidth allocated for high-speed Internet traffic up to a certain level, as Internet traffic has a lower best-effort priority.

Peak Channel-Change Demands – Commercial Blocks

Peak channel-change demands occur during commercial blocks and will be most excessive when large numbers of subscribers are accessing the same broadcast channels, driving up demand for unicast traffic towards broadcast TV caching servers, if present. Broadcast TV caching servers will help reduce the amount of actual multicast channel changes, since many channel changes are short-lived; picture-in-picture capabilities also help. Since broadcast TV caching servers also have finite capacity and will likely serve STB requests on a first-come, first-served basis, extreme load conditions may result in the STB issuing the IGMP join request directly, in case the caching server's response times fall under a certain threshold.

The important point is that even in extreme load situations, no viewers are denied access to broadcast TV content. Supporting reliable and responsive channel-zapping capabilities during potential peak scenarios, and capabilities such as IGMP stateful switchover and highly available unicast and multicast routing, will be critical in ensuring a high QOE to the user.

Peak VOD Concurrency – New Movie Releases

Peak VOD concurrency is one of the most critical traffic patterns, yet it remains largely unknown, given the early stage of IPTV deployments. Getting a handle on peak VOD concurrency versus average VOD concurrency will be important. Peak VOD concurrency will increase unicast traffic in the network significantly, since for each subscriber requesting VOD content, a separate unicast stream must be supported between the STB and the VOD server. The type of content being requested, and the priority it is assigned, will further differentiate VOD content and service assurance. For example, free VOD content may have a lower admission priority than new releases, and resuming a paused movie may have a higher priority than requesting a new movie.

As peak VOD concurrency enters double-digit percentages, unicast VAC and VOD server network architecture will play critical roles. In addition to appropriate network dimensioning to avoid peak conditions, they will lead resource contention and enable carriers to maximize revenues from available capacity by making policy- and priority-based admission control decisions.

Resource Failure Detection & Recovery

If not planned for appropriately Resource failures can and will have a negative impact on service assurance and QOE. First-mile access bandwidth is a critical asset that may be affected by environmental conditions. Using loop analyzers and conditioning tools, the ongoing line quality and bandwidth capacity can be assured at stable levels. Unless redundancy is provided, resource failures in the network may create any one of the previously mentioned peak-usage scenarios, given that network resources once planned for are no longer available, even if only for a short time. A loss of 100 Mbit/s in the second mile, between the DSL access multiplexer and the aggregation node, may affect a small number of subscribers, but the impact per subscriber could be significant. A similar bandwidth loss in the third mile, between the aggregation node and the edge router, would affect a larger number of subscribers, but the impact per subscriber would be less, as capacity loss is distributed over a larger subscriber base.

The first element is to provide for contingencies in the network to prevent a single point of failure resulting in capacity loss. Redundancy provisions within a node, such as non-stop routing and forwarding and in-service upgrades, can protect against aspects such as line-card, control-card, power-supply, and software failure. Preprovisioning backup capacity with Sonet/SDH protection or MPLS fast-reroute options can provide a mean time to respond of less than 50 milliseconds, and are typical solutions to increase capacity availability.

Regardless of the network-failure scenario and recovery scheme, real-time updates to available bandwidth in the network will be critical to guaranteeing service assurance. The VAC mechanism thus may require the ability to revoke previously admitted sessions, in case the allocated capacity for user sessions falls under the available capacity left in the end-to-end service delivery path. By doing so, the network may be restored to congestion-free operation for remaining services. Policy-based decisions would apply to determine priorities between services in progress, such as which ones to maintain and which to revoke. Since these decisions may take seconds to complete, faster preemption and mitigation mechanisms can assure that voice and video services are affected much less (if at all) than best-effort services such as high-speed Internet.

4.2 Network Dimensioning for Optimized QOE

Compared to services such as traditional high-speed Internet and VOIP, IPTV/VOD service demands will result in vastly different bandwidth and QOS requirements at various segments of the end-to-end broadband network. Before going further, it is important to define the high-level network segments that are assumed to exist:

1. **First Mile:** The link between the subscriber's home and the first access network element.
2. **Second Mile:** The link between the broadband access network element and the broadband aggregation node.
3. **Third Mile:** Link between the broadband aggregation node and broadband edge router.
4. **Fourth Mile:** The link between the broadband edge router and the video source. It should be noted that in distributed-VOD server architectures, and in scenarios where regional and/or local content is injected closer to the subscriber, the "fourth mile" could be a direct link between a video source and a broadband aggregation node.

Although not the focus of this paper, it is worth mentioning that the "last 100 feet" (the in-home network link connecting the IP STB and the residential gateway) plays a role in QOE, although VAC does not need to take this link into consideration. The in-home network link will, in many cases, have higher bandwidth than the first-mile link, obviating the need for VAC involvement. Characteristics other than bandwidth availability, such as in-home coax quality, are more likely to create adverse QOE effects. Solutions are available in the market to address these challenges.

Figure 4 provides a view of how various IPTV and other triple-play services affect the network links defined above. Provisioning appropriate network bandwidth to handle peak-usage scenarios while maintaining an optimized cost structure starts with an understanding of the link requirements throughout the network. Many of the requirements discussed in this paper work in parallel with such a network design to assure acceptable QOE.

As the diagram depicts, bandwidth requirements for the first mile remain relatively constant and fixed, based on the aggregate access bandwidth needed to support the subscribed services. The access loop is provisioned and conditioned to ensure it can support these commitments regardless of potential loop-rate variations. Service providers looking to deploy very-high-bit-rate DSL access networks are targeting a minimum of 25 Mbit/s per subscriber, which would support a minimum of one HD channel, two SD channels, variable-rate high-speed Internet, and multiple VOIP lines, among other services.

The second and third miles will be most affected by the number of concurrent VOD sessions and broadcast TV channels, and demand-driven bandwidth requirements may fluctuate dynamically. In some Ethernet-centric network architecture designs, separate virtual private LAN service instances may be deployed to ensure bandwidth separation on a per-service basis.

Fourth-mile connections, as defined in this paper, will match the physical connection of the video source, be it unicast (VOD server) or multicast (head-end system), while streaming capacity limits of the VOD server are enforced by the IPTV middleware to avoid server overload conditions.

Figure 4: Service Impact on Capacity Needs at Each Network Link

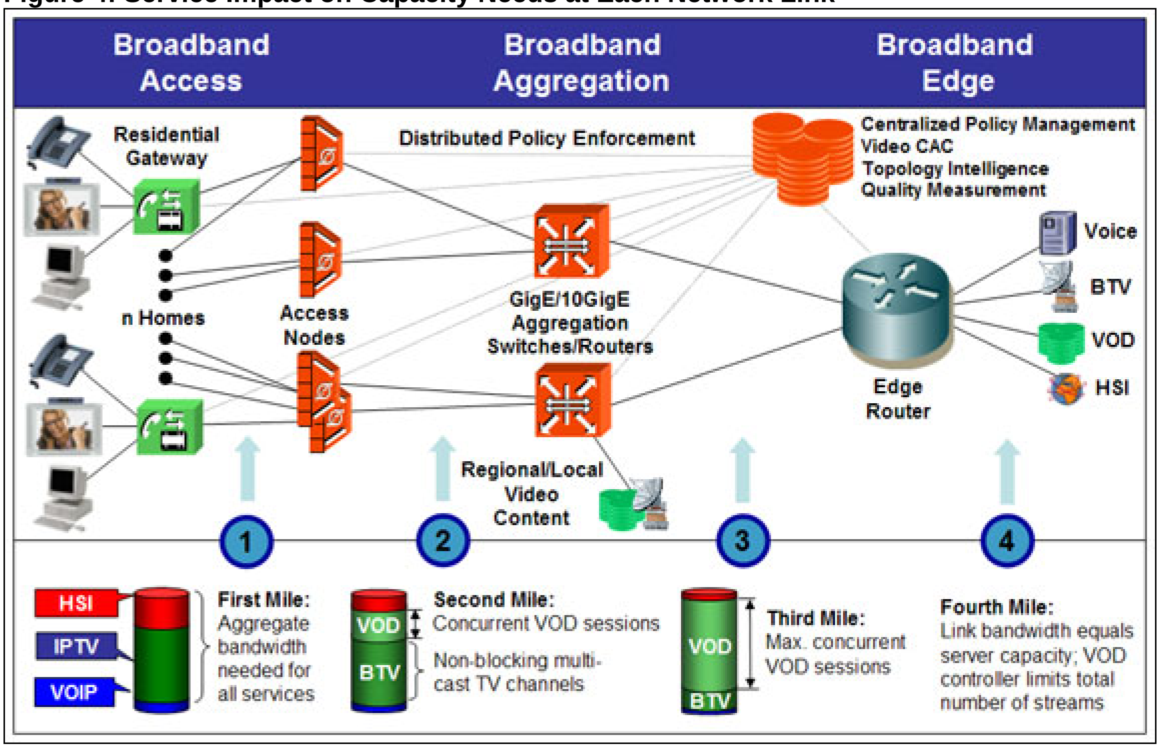

The architecture shown above assumes a centralized policy control model, whereby policies are enforced in a distributed manner throughout the broadband network. A key challenge to VAC is that video traffic is very dynamic and bursty, which makes per-session bandwidth reservation approaches inherently ineffective and inaccurate. Historically, per-session bandwidth-reservation techniques never succeeded in wide-area data networking environments.

The solution to this problem is to apply and operate VAC at link levels where aggregate bandwidth-utilization levels and thresholds can be monitored more accurately and the statistical multiplexing effects of concurrent video streams are taken into account. The VAC management and control function is logically centralized to ensure that it operates with an updated near-real-time view of available bandwidth resources throughout the end-to-end network, even as VAC servers may be geographically assigned to different service regions to provide control scaleability.

Besides being aware of the end-to-end network state, it is also important that the VAC function be aware of subscriber entitlements and account status, SLAs, and relative admission and preemption priorities, as well as application behavior. By means of policy-based admission control rules, the service provider can enable VAC to make informed decisions to help optimize revenues from available network resources. For example, different link-utilization levels may be defined and policy-based admission control rules may dictate admission of free VOD content up to certain utilization levels, beyond which only requests for premium VOD content or resume requests of paused VOD sessions would be granted.

Time-based policies may be applied to influence these utilization thresholds over the course of the week or day to maximize network utilization in relation to historical or anticipated demand patterns (e.g., upcoming VOD releases). Through integration with subscriber management, the VAC function can evaluate potential subscriber privileges that may apply to the admission decision. Through the integration with IPTV middleware, a potential service request denial may result in a proper "busy tone" message to the end user, rather than providing a "black screen." In its operation, VAC is complemented by network-based QOS control mechanisms, such as H-QOS, which ensure that potential video traffic bursts are properly handled without packet loss.

Monitoring and auditing of the VAC function itself provides insight into network utilization and the frequency and duration of periods in which network utilization reaches levels that result in service denials. This information may be used as input to provisioning decisions for additional capacity.

Finally, as mentioned above, end-to-end QOE measurements of video and audio quality are important to verify that both the network and VAC system are operating as they should and that users receive the quality to which they are entitled. Video quality monitoring and reporting tools may also be centrally resident, taking advantage of in-network diagnostics that are collected at various points in the broadband network and correlated with diagnostics obtained from STBs, video headend equipment, and servers to obtain an accurate end-to-end understanding of the user's QOE.

Policy-based admission control has more applications than just IPTV/VOD; IP multimedia services such as video conferencing can benefit equally from VAC. Therefore, the VAC component needs to be designed in a way that provides openness to easily integrate with additional applications and interfaces.

V. Summary & Forward-Looking Thoughts

As the industry moves forward with IPTV deployments, the understanding of the technical pitfalls and solutions required to address them becomes more apparent with each passing day. Phase II of the IPTV market, driven by the need for service assurance and an emphasis on QOE in order to build an initial and substantial customer base, will likely turn out to be the most important phase in establishing industry leaders among both vendors and service providers.

Initial deployments and early trials can be supported to some extent without all the advanced features outlined in this white paper. Transitioning IPTV into a mass-market service, however, will undoubtedly require more refined and robust solutions, as discussed above. In other words, there may be no Phase III – or at least, it will take a long time to get there – if Phases I and II are not completed successfully. Without support for many of the features identified in this paper, accomplishing this may be difficult, if not impossible. These critical features are summarized in **Figure 5**.

Figure 5: Key Challenges & Requirements

CHALLENGE	KEY REQUIREMENT	RATIONALE
"Always On" User Expectations, Flawless Quality	Differentiated and deterministic QOS	Video has stringent SLA needs
	End-to-end service and policy control	Effective SLA enforcement
	Non-stop service delivery/operation	Uninterrupted viewing, 24/7
Capacity Planning, Cost Control, Right-Sizing	Optimized broadcast TV delivery (distributed multicast, IGMP snooping)	Minimize BTV transport cost; maximize channel bouquet
	Optimized video content insertion (flexibility to insert video content at centralized and decentralized locations)	Minimize VOD transport cost; regional content, ad insertion; scale higher VOD concurrency
Congestion Avoidance, Service Assurance	Accurate intelligence on resources	Detect and prevent contention
	Video quality measurement/monitoring	Verify perceptual QOE for users
	BTV/multicast video admission control	Limit BTV channel overbooking
	VOD/unicast video admission control	Limit VOD session overbooking

With that said, it is always important to continue looking forward to the next challenges, so that the network infrastructure and technology deployed today anticipates tomorrow's potential problems and opportunities. Some future challenges may be network-centric, while others may be more service- and user-centric.

One future service-centric problem will be the transition from channel surfing to intelligent content navigation. Much emphasis has been (rightly) placed on broadcast TV support, high-performance channel zapping, and a robust underlying broadband network architecture in support of today's "surfing" paradigm. But as additional content is added to IP-based networks and IPTV services become more personalized and user-centric, the ability to shift from traditional channel surfing to intelligent content navigation and a personalized, on-demand TV experience will be the key competitive differentiator. The ability to get any content, at any time, in a user-friendly presentation via next-generation electronic programming guides will shape the next wave of video services. The resulting consequences of growing on-demand video content and concurrent usage on network traffic patterns, capacity projections, and content distribution approaches are dramatic and potentially disruptive. Therefore, a number of precautionary measures must be taken with respect to network design, capacity dimensioning, traffic engineering and optimization, congestion avoidance, and finally service assurance mechanisms, to handle the unpredictable nature and demand peaks of on-demand video services – in combination with other services – in a cost- and performance-optimized manner.

Furthermore, as integrated services – data and voice blended with IPTV services – become better defined and delivered, the importance of sharing budgeted bandwidth among triple-play services may become more critical. Increased QOS intelligence and bandwidth-sharing priorities will need to be defined on an end-to-end basis to support potentially unexpected blends of services.

Also, we cannot overlook Internet TV: The business model and optimal service delivery model for this service is not yet fully understood, but various Internet service providers have expressed a strategic interest in providing it. The fact of the matter is that the importance and bandwidth demands of Internet-sourced video will grow, and providing assured QOS/QOE will become a factor if and when such services evolve from the current download model to an instant-viewing model. Internet TV could emerge as either complementary to or competitive with walled-garden IPTV, depending on how the value chain of video content delivery evolves. Regardless, network- and service-infrastructure layer solutions, in addition to newly emerging solutions, will need to take this trend into account over the next few years.

But let's not get too far ahead of ourselves: The highest priority should be placed on solving problems for short- to medium-term market requirements, and those problems lie squarely in the domain of service assurance and QOE. These will be the critical elements to watch for the remainder of 2006 and into 2007.