IP NETWORKING IN HFC BROADBAND NETWORK

Oleh J. Sniezko, Vice President of Broadband Engineering
Xiaolin Lu, Vice President of Strategic Engineering and IP Networking
Luisa Murcia, Vice President, Data Services Technical Operations

AT&T Broadband, 188 Inverness Drive W.
Denver, CO 80112

Abstract

This paper describes the IP networking trends, the existing solutions, and possible evolutionary implementation of an integrated IP platform over hybrid fiber-coax (HFC) broadband telecommunication networks. Some of these implementations are analyzed based on the trials conducted and planned by AT&T Broadband.

In the first phase of these trials, the IP network is designed to enable AT&T Broadband Choice[SM] (ABC) for our customers. It provides them with a choice of multiple ISPs while at the same time enabling IP telephony supported by a broadband telephony interface (BTI).

The next phase of the planned trials will implement an end-to-end voice over IP (VoIP) solution. The ABC networking will be continuously optimized as more advanced routing equipment becomes available.

The bulk of this paper concentrates on those technical aspects of the field test that relate to networking, design and performance. The authors list and describe the existing equipment used in the trial and analyze the challenges to bringing the service reliability and quality to a level equal to or better than is being provided by traditional telecommunications networks.

INTRODUCTION

Today, HFC access networks deliver a multitude of services using several, very different multiplexing techniques for downstream signals and several different media access protocols in the upstream path. Similarly, on the network side, signals for different services are distributed over disparate transport systems, and interface with several different backbone networks. This situation exists mostly due to the historical evolution of telecommunications, data and entertainment service delivery systems.

The technological progress in the recent years allows for integration of most of these systems. Operators can now begin the transformation of their metropolitan and HFC networks towards packet-based service delivery networks for most of the interactive services. The metro network can be fully integrated and the multiple service platforms (MSPs) are becoming readily available from traditional transport vendors and from new vendors. These platforms are suitable for distribution of video, data and telephony signals. They often integrate layer 2 and layer 3 functionality. Together with the positive cost trends in video encoding technologies and in MPEG flow processing equipment, signals for all services, including broadcast analog video, can be distributed over these integrated platforms.

Similar integration on the HFC access side is technologically possible but not currently practical. The reasons for this are: the traditional consumer electronic equipment characteristics, some regulatory requirements and cost implications of the integrated approach. Most of the video services, including analog and digital broadcast video and VOD as well as analog and digital ad inserts are best distributed over a SCM network in their analog form or as MPEG signals. Most of the other services including telephony, data, iTV, signaling information, and the like can be cost-effectively transported over an integrated IP DOCSIS platform.

The platform integration can be accomplished in a one-time transformation or in a staged manner. In a staged transformation approach, the IP networking can be initially implemented in primary headends where different services, with the exception of broadcast video and MPEG VOD signals with local ads, could be converted for distribution in a packetized form over a DOCSIS-based platform to our customers. On the network side, the services could be initially separated to interface with different legacy transport platforms. From this initial point, the evolution could progress in two different directions. Pure packet-based systems could be driven deeper into the HFC access network with the concept of distributed CMTS. This approach would allow for simplification of networking in the access section of our systems. Only the last mile would remain optimized for RF SCM coaxial delivery. There, broadcast analog signals would be carried in their original format, and digital broadcast video and VOD signals would be carried in MPEG format, QAM modulated onto RF carriers (with digital ads spliced in). Most of the other signals would be carried in their packetized format over DOCSIS to the customer terminal devices.

In parallel, the packet-based networking would expand in metro networks to evolve into an end-to-end IP solution for all services, including traditional video services and telecommunications services. At the same time, some interfaces would be required to account for the fact that many of these services are carried on legacy systems such as PSTNs.

OBJECTIVES

Future IP networking focuses on the design and development of an AT&T owned and operated IP network, and the design and development of the client products that create value for our customers. The design provides our high-speed Internet access customers with a choice of multiple ISPs. The trial of this new infrastructure allows for testing how this network plugs into AT&T's current high-speed data backbone. It also supports IP telephony (IPT) on the access side. Moving forward, this IP network will support not only high-speed data (residential and commercial), but also iTV and other interactive services.

There are two primary differences between this new infrastructure and the conventional IP infrastructure:

1. Ability to efficiently provide AT&T Brodband's customers with convenient access to multiple ISPs and their services;
2. Ability to support multiple applications.

To test these features, AT&T Broadband asked several ISPs to participate in the first phase of the trials. The trial is taking place in Boulder, Colorado in the area with 25,000 homes passed. The following ISPs are currently participating:

- EarthLink,
- Juno,
- Worldnet, and
- Excite@Home.

This group may be joined by RMI.

AT&T Broadband has secured 400 volunteer customers of which 333 are already participating. Each of the multiple ISPs introduced tiered services. The table below lists the tiers of service for each ISP currently participating in the trial.

Table 1: HSD Service Tiering

ISP	Tier 1 # of Users @ 1.5 Mbps	Tier 2 # of Users @ 300 kbps	Tier 3 # of Users @ 128 kbps
EarthLink	21	1	3
Juno	47		
Worldnet	132	51	13
E@H	65		

The tiers presented in Table 1 reflect downstream capacity. The upstream capacity is held constant @ 128 kbps for all tiers.

PLATFORM DESCRIPTION

Generic Requirements

Supporting the infrastructure of the AT&T Broadband Choice[SM] (ABC) involves capabilities in three key areas of the infrastructure:

- a user-friendly interface that can run on the end user's PC to provide a convenient method of selecting services and ISPs;
- policy routing capabilities in the underlying routing network so that traffic destined for different ISPs can be sorted and routed accordingly; and
- a Service Activation System (SAS) for service provisioning and activation in the AT&T Broadband Choice[SM] (ABC) environment, capable of migration to an integrated provisioning platform for multiple-service offering.

GENERIC DESCRIPTION OF INFRASTRUCTURE

Integrated Platform in Access Network

The HFC access infrastructure was designed to meet DOCSIS 1.1 requirements, but the access platform was based on equipment compliant with DOCSIS 1.0. To support high-speed data delivery and VoIP over DOCSIS, the equipment was modified to meet the objectives of the DOCSIS 1.1 standard before the CableLabs certification commenced or to be capable of migrating to DOCSIS 1.1 standard. Both elements: CMTSs and BTIs, were deployed. Many other parameters related to VoIP in the HFC environment and covered by the evolving PacketCable standard had not been certified.

High-Speed Internet Access

On the network side of the CMTSs in primary hubs, the traffic for high-speed Internet access is separated from the telephony traffic. The HSD traffic is routed to different ISPs based on customer choice.

Voice over IP

After being separated from the IP data traffic, the IP voice traffic is routed via traditional gateway (IPDT for IP digital terminal) to Class 5 switches. This configuration on the network side does not differ from the existing approach; the telephony traffic is treated separately. A generic diagram of this network is presented in Figure 1. Additional expansion will be implemented to improve network reliability by creating a redundant point of connection (PoC) as presented in Figure 2.

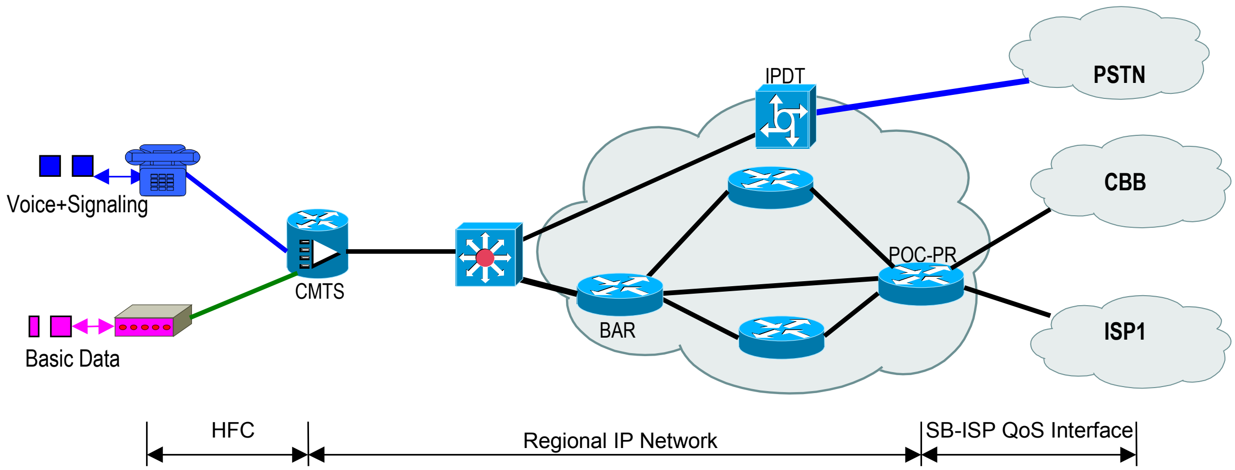

Figure 1: Generic Architecture for Integrated, Multi-ISP and Multi-Application, Broadband HFC, DOCSIS-Based Platform

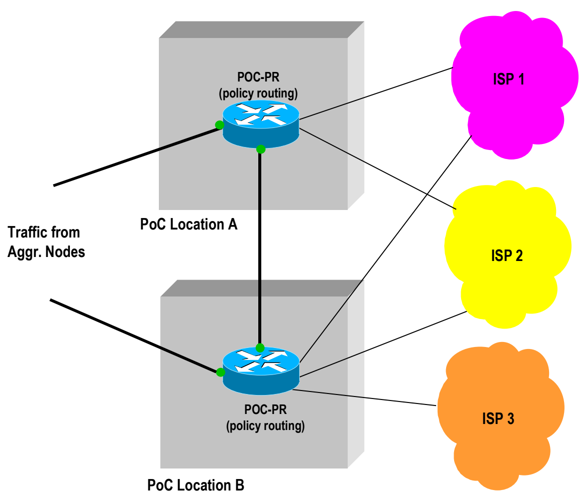

Figure 2: Redundant Configuration of PoCs for Higher Reliability

MULTIPLE ISP NETWORKING

Layered Broadband IP Infrastructure

The multiple-ISP network comprises two main sub-networks: regional broadband access infrastructure and intra-regional backbone. ISPs can interface directly with the regional network (local ISPs) or via inter-regional backbone. A simplified configuration of this layered approach is presented in Figure 3.

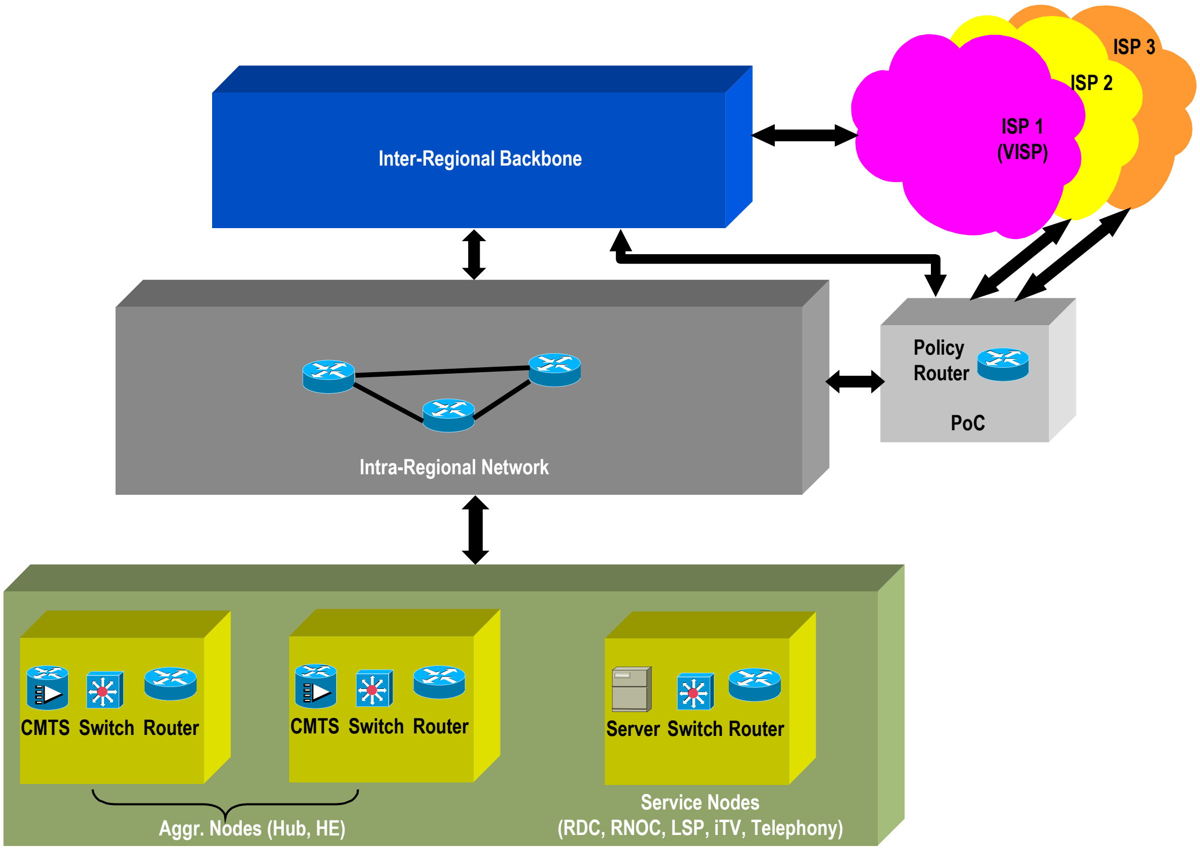

Figure 3: Regional and Inter-Regional IP Infrastructure

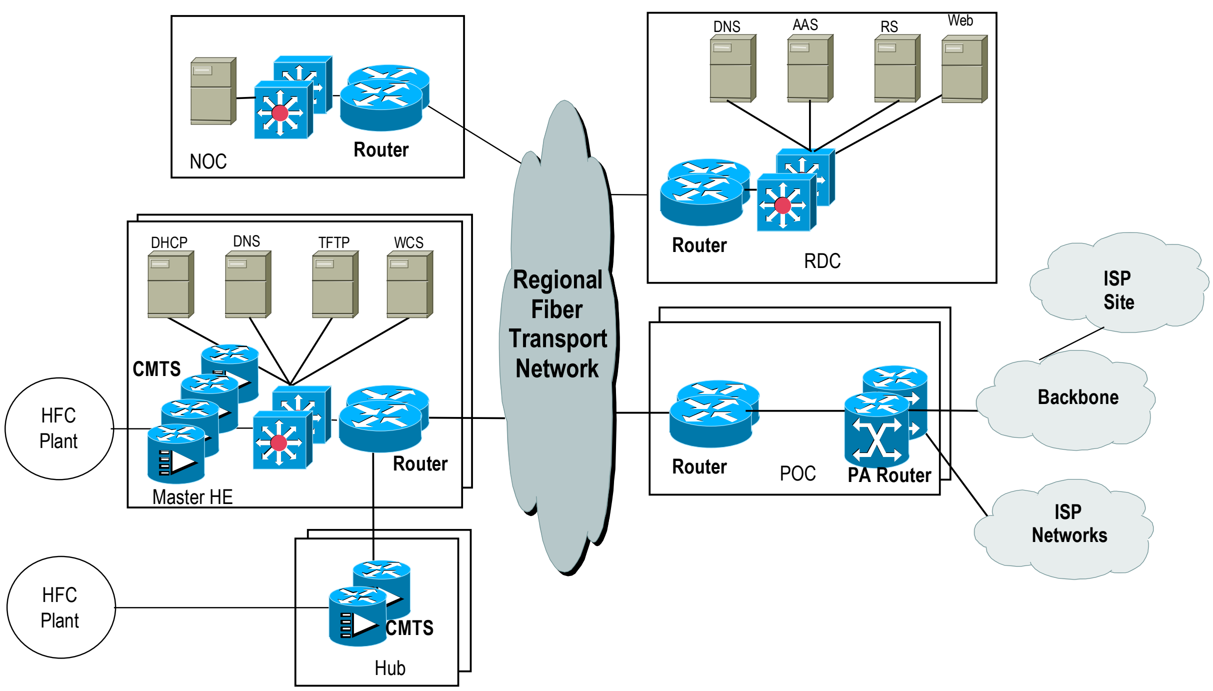

Figure 4: Regional Broadband IP Infrastructure

Regional Broadband Access

Figure 4 depicts the components of the regional infrastructure. The regional network:

- aggregates end-user traffic in aggregation nodes located in hubs and headend; and
- provides local services such as local ISP connections, regional registration, and regional NOC through service nodes.

Customer base projection and network/traffic engineering determine geographical size and boundary of a region.

Inter-Regional Backbone

The inter-regional backbone allows for connecting remote ISPs (virtual ISPs) to the regional network. It also provides a means of distributing content from a central location to the regions. It is used by applications/services requiring inter-regional connectivity. It also provides connectivity to a centralized NOC

Policy Routing

In the ABC environment, one of the major tasks is to route traffic to the ISP of choice in an optimal way to:

- accurately determine the allocation of resources and billing per an ISP; and
- to avoid congestion in the transport systems (including inter-regional backbone).

At the time when the first phase of the trial commenced, source address-based policy routing was developed to meet the requirements of the ABC environment. In this approach, a PC is assigned an address from the address block of the ISP of choice. The packets going out of the PC are forwarded to the ISP based on the source address (policy routing). The incoming packets are routed in a conventional way to the client PC based on its address. This routing method is reflected in Figure 5.

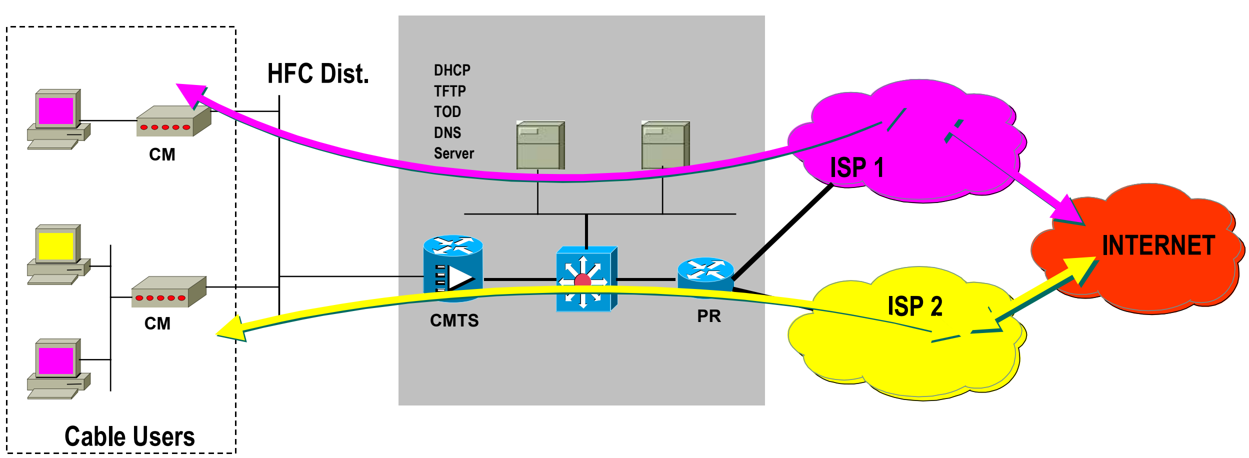

Figure 5: Operation of Source Address-Based Policy Routing

This method of policy routing was supported during the trial by a single policy router (Cisco 7513). The choice of this policy routing was based on technology availability and test results. Some limitations existing at the beginning of the trial were later eliminated by improvements in the policy router. The trial results confirmed that this method of policy routing is effective. However, it revealed some shortcomings such as:

- limited number of ports available on a single router;
- no redundancy provided for ISP connections into the region;
- no load balancing provided for traffic to an ISP;
- policy decisions need to be provisioned (and re-configured) centrally on the single policy router; and
- requires a very high performance policy router to support the large number of policy filtering decisions that may be required in a region.

Many of these shortcomings can be addressed in a redundant architecture of PoCs (refer to Figure 2). In this configuration, policy routing is applied to incoming packets on all the marked interfaces of PoC-PRs. These PoC-PRs may be collocated. The incoming traffic from the region is load balanced into the PoC-PRs. ISPs with connectivity to both PoC-PRs benefit from load balancing automatically.

With the progress made in routing protocols, a different method of policy routing can be adopted. One method that is being considered for the next phase of the trials is MPLS-based policy routing. The application of MPLS (multiprotocol label switching) and LDP (label distribution protocol) will allow for many other features and will provide sufficient improvement in provision of QoS and CoS. An MPLS-based routing table can be constructed for each ISP. This table would span all PoC locations to which the ISP is connected. Under this scenario, ISPs can connect to as many of the PoC locations as they choose.

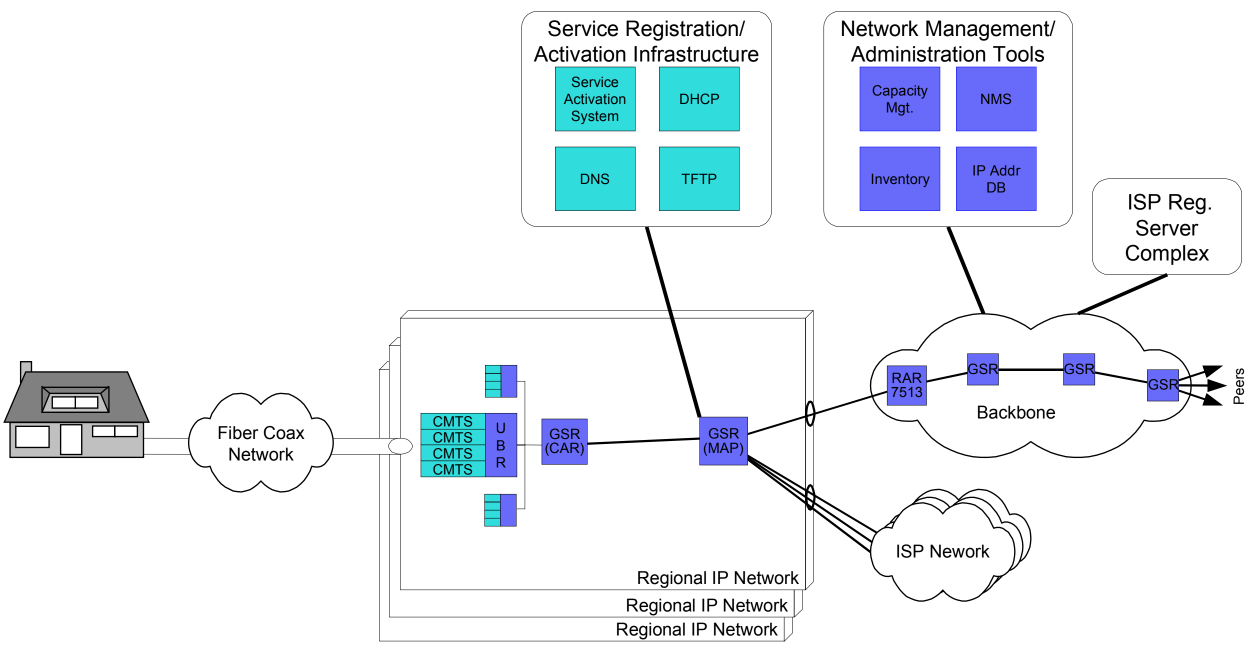

Figure 6: Network Management and SAS Infrastructure

<u>Network Management</u>

The service provisioning and activation as well as network management in the multiple ISP environment becomes critical. The objective of system scalability and adaptability to the new services required that a significant effort be dedicated to the development of a suitable platform. The configuration of the network management system is presented in Figure 6

In addition to these requirements, the issue of management of the IP subscriber addresses allocated to ISPs had to be resolved. This issue was especially critical for source address-based policy routing. A detailed address allocation plan was developed. It allows for re-use of Net 10-address space between regions but this feature is not necessary. In the newly built networks, the address allocation does not present a major challenge. However, during migration of the existing networks to the ABC service, a careful evaluation of Net 10 usage is essential. Some address re-assignment may be necessary but can be avoided by careful planning. Address management tools, SAS (DHCP) address usage reporting and automation of address management are essential components of successful address management in the ABC environment.

VOICE OVER IP OVER DOCSIS (IP TELEPHONY)

Generic

The same integrated broadband IP platform was used to provide telephony services. On the access side, BTIs (broadband telephony terminals) were used at customer premises. The calls were terminated into CMTSs located at the aggregation nodes. At the time the trial commenced, DOCSIS 1.1 and PacketCable compliant equipment was not available. Moreover, call management servers (CMSs) and soft switches were either not available or could not provide calling features commonly considered as basic. To provide these features, the calls originating in IP DOCSIS network were routed to Class 5 switch via IP digital terminals (IPDTs), known also as smart gateways, developed for this application to provide fully compliant GR-303 interface. This configuration is presented in Figure 1.

Bandwidth Efficiency and Traffic Models

AT&T Broadband elected to mix data and telephony traffic on shared downstream and upstream channels for two reasons:

1. more efficient bandwidth use, and
2. possibility of simultaneous use of BTIs and CMTSs for digital telephony and data communication (efficient use of equipment).

Statistical data about today's traffic parameters and user behavior were collected to determine the capacity (i.e., a number of users) supported by a single downstream and six upstream (3.2 MHz QPSK) channels (a typical configuration of a CMTS). The methodology used converted a telephony user into an equivalent number of data users and calculated a cumulative number of composite users based on an assumed total number of telephony and data users. With an assumption of low initial telephony penetration and heavy usage of data in the downstream path, it was determined that a single CMTS can support approximately 2,500 HHP and the limiting capacity is the capacity of the downstream channel (highly asymmetrical data traffic at low telephony penetration).

The calculation was performed for the particular trial application. The model, however, is valid for generic applications with adjusted traffic statistics dependent on:

- user behavior for different service types,
- service type,
- compression ratios for IP telephony,
- penetration levels for different services,
- downstream and upstream channel capacity (modulation level on

downstream channel, symbol rate and modulation level on upstream channel),
- CMTS configuration,
- rate limiting policy, and others.

Performance

Nine of the IP telephony loops were tested for 10 loop and data throughput (dial up modems) parameters. The test results were compared against analog loop performance and DT HFC loop performance. The results confirmed that the CMTS/BTI based IP telephony quality below the quality provided by HFC digital telephony. This is understandable in light of the fact that neither BTIs nor CMTSs were tested for compliance with DOCSIS 1.1 and PacketCable standards and BTIs were among the first of this type of terminal.

The major deficiencies were experienced in round trip delays, impulse hits and phase hits. These deficiencies affected a cumulative measure of GoS (grade of service). Moreover, current BTIs were not compatible with most of the dial-up modems. The effective data throughput for the dial-up modem compatible with BTIs was much lower than that for HFC digital telephony NIUs. The test also showed lower than required ringing voltage.

The test results were communicated to the equipment vendors. They were also used in preparation of an RFI for BTIs and IP telephony CMTSs.

QUALITY OF SERVICE (QOS)

The trial IP broadband network was used to provide two different services over shared channels. Moreover, ISP services were offered at three different levels of guaranteed downstream throughput. Therefore, at the minimum, QoS functions had to enable:

1. Rate limiting, and
2. Priority handling.

Implementation of these two functions involves control and configuration of HFC and IP parameters. These configurations are presented in Figure 7. DOCSIS 1.0 allows for configuration of class of service (CoS) parameters:

- Max downstream rate: enforced per CM by CMTS on a packet basis,
- Max upstream rate: enforced per SID by CMTS via time-slot scheduling,
- Min guaranteed upstream rate: enforced per SID by CMTS via CIR type mechanism,
- Upstream channel priority: used by CMTS when scheduling time-slot grants per SID.

An enforcement of the control mechanisms for these parameters was assigned to the vendors participating in the trials. To allow prioritization of telephony traffic, DOCSIS 1.0+ with some additional expansions was used. It provided:

- unsolicited grant service (UGS), and
- ToS overwrite

In the IP part of the network, a DiffServ model, based on ToS (type of service bits) was applied to enable different priority queues.

Three different tiers of high-speed Internet access service were provided, all of them under Best Effort. The CoS was defined in the CM configuration file. Voice service was provisioned under UGS with a minimum upstream for voice signaling traffic. The ToS bits were reset by CMTS to DiffServ Code Point EF (Express Forwarding) signaling.

This implementation allowed to provide all QoS functionality required for the trial. The task was simplified by the fact that the IP voice traffic was terminated at the IPDTs.

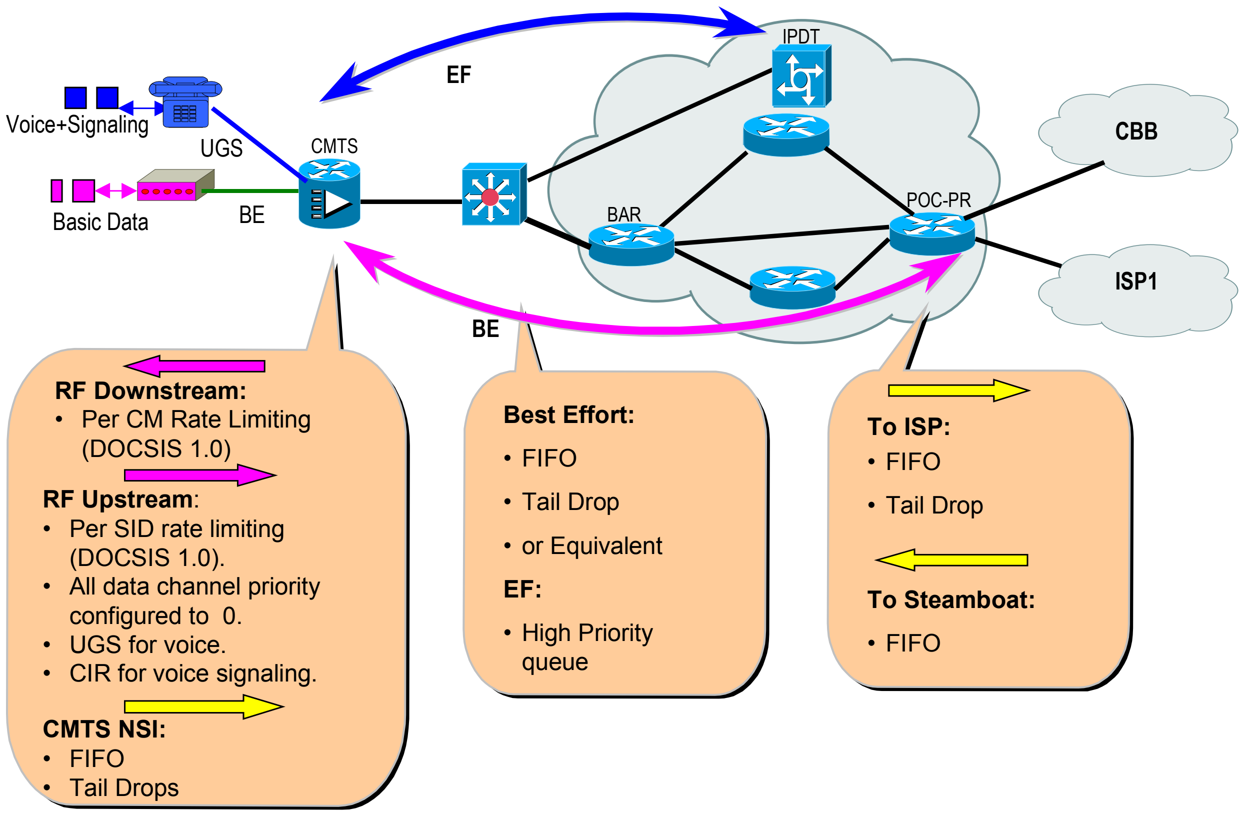

Figure 7: Quality of Service Functionality in IP/DOCSIS HFC and IP Networks

CONCLUSIONS

The first phase of the ABC trial has met its objectives. It provided a unique experience to our high-speed Internet access customers. The new network allowed for implementation of a service model where:

1. Customers can:
   - have a relationship with a selected ISP or ISPs;
   - upgrade the service, change an ISP or any other feature;
2. ISPs can:
   - offer significantly improved service;
   - provide regional or national service;
   - develop applications based on the demand;
3. MSOs can:
   - respond to the marketplace;
   - create a specific value proposition to a particular customer;
   - maximize value delivered to their customers.

The trial also showed that an additional effort has to be expended to match requirements for IP digital loop performance and make them comparable to those achieved in HFC CBR digital telephone loops. Standard compliant equipment should provide significant improvement in these parameters.

The IP networking required more engineering effort than the setting up of IP telephony. This is mostly due to the fact that IP telephony was contain to the HFC network and the section of the IP network between the

CMTSs and IPDTs. An end-to-end VoIP solution would require a significant amount of effort to implement, even if it complies with the industry standards.

Next phases of the IP network trials will most likely be conducted with end-to-end IP telephony systems and possibly include other IP-centric services. The trial area will pass approximately 100,000 homes. This phase is scheduled in late 2002 but it will depend on the equipment availability.

ACKNOWLEDGMENT

The authors want to acknowledge many of their colleagues from AT&T Broadband and other AT&T organizations, who contributed significantly to the entire project. Specifically, we want to thank Samir Saad, Han Nguyen, Tayfun Cataltepe, Steve Bulka, Pete Zarrelli, Brian Field, Julie Regalado, Cameron Gough, and Allen Stone, but we are certain that many more deserve the acknowledgment.

AUTHORS' CONTACT INFORMATION

Xiaolin Lu
Vice President of Strategic Engineering and IP Networking
Tel: (303) 858-5575
Fax: (303) 858-3181
lu.xiaolin@broadband.att.com

Oleh J. Sniezko
Vice President of Broadband Engineering
Tel: (303) 858-5027
Fax: (303) 858-3181
sniezko.oleh@broadband.att.com

Luisa Murcia
Vice President, Data Services Technical Operations
Tel: (303) 858-3553
Fax: (303) 858-3181
murcia.luisa@broadband.att.com

All authors are located at AT&T Broadband,
188 Inverness
Englewood, CO 80112