HOW NEXT-GENERATION BROADBAND NETWORK ARCHITECTURES WILL DELIVER INTEGRATED SERVICES TO BUSINESS CUSTOMERS

Vikram Saksena
Narad Networks

Abstract

This paper describes a next-generation broadband network architecture for providing 100 Mbps/1 Gbps Ethernet services over an existing HFC infrastructure. By deploying Ethernet switching elements with integrated high speed modems at critical junction points, today's tree-and-branch cable network topology can be augmented with a symmetric, full duplex, data channel that provides orders of magnitude higher bandwidth than is currently available. With built in multi-service QoS capabilities, this next-generation broadband network becomes fully capable of addressing the large, highly lucrative commercial market segment.

Coupled with advanced network management and service delivery back-office systems, the entire process of deploying and managing broadband services can be fully automated. High levels of network reliability can be maintained by making the network elements remotely manageable and self-configurable. By radically compressing service deployment and provisioning intervals, cable companies can accelerate the process of new revenue generation. Customers get the benefit of being able to order services on demand from a broadband services portal without waiting for long service turn-up intervals and without worrying about bandwidth constraints. Flow-through provisioning and activation of back-office billing and customer care systems are accomplished through a robust Directory Enabled Networking (DEN) based service delivery platform.

INTRODUCTION

Recent years have seen the cable industry deploy a number of advanced services in the consumer marketplace. Driven by the need to generate new sources of revenue, the industry is adopting new technology and is rapidly transforming itself into one that provides multiple integrated services. This has resulted in increased cable modem penetration, expanding deployment of video-on-demand and digital TV, and cable telephony services. To accommodate these services, the bandwidth used by cable systems has steadily increased from 270 Mhz to 860 Mhz HFC over the last twenty years and the networks have been upgraded to support two-way traffic.

To continue on the revenue growth curve, the next major challenge for the cable industry is to address the needs of the commercial market. The small and medium business (SMB) segment of the commercial market in particular represents a lucrative growth opportunity for the cable companies and is a natural progression from the consumer markets addressed by the cable companies today. However, the existing cable infrastructure with its broadcast characteristics and asymmetric bandwidth capabilities is largely incapable of fully supporting the needs of the SMB market. This is supported by a recent Morgan Stanley study that reveals that the cable modem penetration in the commercial markets is insignificant when compared to DSL and other access technologies, and is not expected to increase in the foreseeable future. Recognizing these limitations, some cable companies have deployed separate physical networks based on ATM, SONET and other proprietary access technologies to address the commercial market. However, the high cost of deploying and operating a separate physical network has limited the addressable market to a few concentrated areas that are largely beyond the reach of the majority of SMB customers both in terms of availability and costs.

This paper proposes a new broadband network architecture that leverages a cable company's largest deployed asset - the HFC network - to reach a significant portion of SMB customers and serve them in a highly efficient and cost-effective manner. The proposed solution adapts switched Ethernet technology commonly deployed in today's enterprise networks to broadband HFC networks and enables a rich suite of IP services – high speed Internet access, virtual private networking (VPN), business grade telephony, web hosting, content distribution, and storage area networking. The technology foundation, based on high speed Ethernet switching, provides fiber-like reliability and performance characteristics over existing HFC networks at a fraction of the cost of alternative fiber-based solutions. This paper also introduces a service delivery platform that enables the rapid creation and provisioning of services in near real-time using a Directory Enabled Networking (DEN) architecture.

This next generation business solution gives the cable operators a true commercial-grade broadband network and service delivery platform that transparently coexists on their ubiquitously deployed HFC network. This new architecture provides a competitive advantage against the incumbent carriers who offer conventional narrowband services in a relatively expensive and inflexible fashion due to tariff restrictions and culture.

THE SMB OPPORTUNITY

Market Size

There are approximately 8 million small and medium businesses in the US that typically employ 5 to 100 employees. Approximately 60% of these businesses are within a few hundred feet of the local HFC network and can be cost-effectively reached by the cable operator. On average an SMB pays about $6000 in annual revenues to its service provider, which in most cases is the incumbent local exchange carrier (ILEC). Therefore, the addressable market opportunity for the cable industry is an annual $28.8 billion. Compared to the consumer market opportunity of $35 billion, the SMB market represents a significant growth opportunity for cable companies. Cable operators can significantly grow their revenues and EBITDA margins with modest penetration rates, given the higher margins associated with the SMB revenues.

SMB Requirements

The core set of services demanded by SMB customers includes high speed Internet access and telephony. In many cases, SMBs also require additional value-added services such as Centrex, virtual private networking, web hosting, and firewall services. Most of these services require symmetric bandwidth. A common service configuration includes an integrated access device (IAD) on the customer premises that aggregates voice and data traffic over a symmetric, full-duplex access line such as T1 or fractional T1 to the ILEC's central office (CO). At the CO, traffic is demultiplexed; voice traffic is sent to a Class 5 switch that connects to the PSTN and data traffic is sent to an IP network that provides Internet and data services.

While the 1.5 Mbps bandwidth provided by a T1 line is generally adequate for basic voice and data services, bandwidth requirements increase when customers begin demanding additional value-added services. Although NxT1 services are available from the ILECs in a few areas, by and large the next step up in access bandwidth is DS3 (45 Mbps). However, DS3 costs are prohibitively expensive and beyond the range of affordability of most SMB customers. Therefore, in many cases, as traffic demand increases customers are forced to use techniques such as voice compression to squeeze more traffic into the same T1 line. This not only leads to degradation in voice quality but also causes higher delays and packet loss for data traffic. In this scenario, lack of affordable bandwidth from the ILECs becomes a barrier to the growth and proliferation of broadband services.

It has become a common practice for the commercial customer to enter into service level agreements (SLAs) with its service provider that binds the service provider to contractual commitments related to service performance metrics such as delay, throughput, packet loss, and availability. In order to effectively deliver on these SLAs, the service provider needs a QoS-enabled, multi-service network infrastructure that is highly reliable and scalable. In addition, the service provider also requires a highly sophisticated back-office operations and management infrastructure that is fully automated to reduce provisioning cycle times and mean time to repair activities.

The incumbent service providers are at a competitive disadvantage because of their high cost legacy networks and outdated back-office operations systems that create unduly long provisioning intervals leading to customer dissatisfaction. This scenario presents a significant opportunity for the cable companies to leverage their HFC network to offer broadband services to commercial customers at the lowest cost per bit. By adopting state-of-the-art back-office support systems, cable companies can deliver on the promise of customer SLAs in a fashion far superior to the incumbents.

CURRENT HFC PLANT CHARACTERISTICS

In current HFC networks, the coaxial portion of the plant is a tree-and-branch topology with a shared LAN-like structure. In the downstream direction, RF signals are broadcast over the 54-860 Mhz spectrum. These signals often traverse several miles of the coaxial plant while crossing multiple amplification and splitting stages before reaching the customer endpoint. To maintain acceptable signal-to-noise ratio up to the very last subscriber on the cable plant, bandwidth beyond 860 Mhz cannot be effectively utilized in the current design.

The broadcast nature of the downstream plant also raises privacy concerns, especially in commercial applications, because a given customer's data is visible at every endpoint on the cable plant. To mitigate privacy concerns, HFC encryption techniques such as BPI have been added. These techniques introduce additional overhead and complexity.

In the upstream direction, bandwidth between 5 and 42 Mhz is shared across all subscribers. Due to ingress noise issues and interference with other signals, the actual usable portion of the upstream spectrum is further limited to a fraction of this spectrum. One of the methods to access this shared media channel is through a media access protocol defined in the DOCSIS standard. Endpoint cable modems send requests for transmission to a CMTS at the headend which grants permission for accessing unused time-slots in the upstream frame. If a request is denied, as might happen during periods of heavy load, the cable modem continues to retry until a request is granted. Such a centrally controlled reservation based access scheme leads to large and unpredictable delays making SLA commitments very difficult to meet. While DOCSIS 1.1 allows for prioritized access to time-sensitive traffic such as voice, the delays could still be on the order of tens of milliseconds.

The nature of the initial DOCSIS technology (1.0 and 1.1) is that symmetrical services are nearly impossible and are not possible at all as demand increases. Data flow asymmetry makes it very difficult for DOCSIS-based access systems to penetrate the commercial market.

In summary, the current cable plant architecture has bandwidth and performance limitations that make it difficult for use in commercial applications resulting in low cable modem penetration.

NEXT-GENERATION ARCHITECTURE

The architecture of broadcast cable networks is very similar to that of Ethernet local area networks (LANs) that have been successfully deployed in the enterprise environment for many years. Ethernet LANs have evolved over time to meet the growing needs of the enterprise in terms of increasing bandwidth and applications requirements. The next-generation cable network architecture proposed here for cable networks draws heavily from the lessons learned from the evolution of Ethernet.

Evolution of Enterprise Ethernet LANs

First-generation Ethernet LANs are based on a shared bus architecture where all stations contend for access to the bus using the CSMA/CD access protocol. As the number of stations and traffic increased, excessive packet collisions cause significant retransmissions that lead to severe delays and throughput degradation.

To alleviate this situation, second-generation architectures segment the shared bus using Ethernet hubs to reduce the serving group size on each segment. However, as bandwidth intensive, multi-service applications continued to emerge, even this architecture can be inadequate. While the hub based architectures allowed some degree of control on managing the overall utilization of a shared Ethernet segment, enforcing Quality-of-Service (QoS) for multi-service applications such as voice, data, and video is extremely difficult as packet collisions cannot be completely eliminated.

The only way to enforce true QoS is to eliminate the shared nature of the network and transition to a point-to-point switched network where every end-user gets a dedicated network link. This recognition leads to the creation of third-generation Ethernet technology - Switched Ethernet - that uses a store-and-forward switching paradigm for forwarding packets.

Switched Ethernet enables several capabilities besides allowing dedicated bandwidth per user. Switching allows packets to be directed only to those endpoints to which they are addressed rather than broadcast to everyone on the network. This eliminates privacy concerns of shared media networks. Through the 802.1q virtual LAN (VLAN) header extension, virtual private network services can be enabled at Layer 2 and the network can be logically segregated into multiple VLAN domains. Enterprises are using this capability effectively to create VLANs for each administrative domain within the corporation on a common switched Ethernet network. The VLAN header also has a priority field that can be used to classify packets and to service them at different priority levels. Modern Ethernet switches offer priority queuing at each port so that end-to-end delays for critical traffic can be managed more effectively. With adequate buffering at each switching point, the store-and-forward paradigm also allows the network to absorb traffic bursts much more efficiently without losing packets and degrading throughput. Thus switching enables QoS and multi-service capabilities in the network.

Ethernet Switching over Cable Networks

Current cable networks are analogous to first-generation Ethernet LANs – both use a shared bus architecture. Node splitting and the mini fiber-node concept for segmenting the cable network to support smaller serving group sizes is analogous to second-generation hub-based Ethernet LANs.

Due to the inherent deficiencies of the shared nature of existing HFC architectures, the existing network designs cannot scale to meet the needs of the commercial market segment. In order to enable QoS, true multi-service capabilities, and to ensure privacy for commercial customers, a new architecture is needed. Similar to the latest Ethernet technology, the new architecture must have a switching element that is capable of providing secure data to the end user. This new architecture must also deliver all of the services that the commercial customers demand. Clearly, unless the cable operator is willing to give up existing spectrum for the commercial customer, new bandwidth must be activated on the HFC network. The next natural evolution in HFC is to go beyond the existing 860 Mhz spectrum limit. Since continuing to grow a shared network is not feasible, a new data channel along with a switched network overlay must be created on the cable plant. The coaxial cable medium has usable spectrum beyond the 860 Mhz spectrum used by existing systems. The proposed design makes this previously unused spectrum usable by breaking the tree-and-branch coaxial network into a number of point-to-point coaxial segments interlinked by switching elements. Since each point-to-point coaxial segment is only a few hundred feet long, additional bandwidth beyond 860 Mhz can be effectively utilized for transmitting and receiving high data rate signals while maintaining acceptable signal characteristics. Each coaxial segment can be viewed as a point-to-point link over which the new data channel operates. Data is recovered and relayed at each endpoint. Synchronization is required only link-to-link, eliminating the need to synchronize with a central headend control device and the corresponding end-to-end ranging operations typical of DOCSIS systems. The existing services below 860 Mhz coexist on the same physical cable plant and continue to operate as before.

The best choice for the switching technology to be used in this new HFC architecture is Ethernet. Using Ethernet switching in the access network allows for a seamless interconnection of enterprise LANs to metro-optical networks, both of which use Ethernet switching. Additionally, with the use of Ethernet, the headend interconnection architecture is greatly simplified. Since routers support standard Ethernet interfaces, the fiber facilities carrying Ethernet signals from the HFC optical node can be directly terminated on headend routers without requiring any specialized equipment. Although specifically tailored for HFC, the Ethernet segment between the headend and the customer premises functions like a standard Ethernet LAN segment to the headend access router. Since IP over Ethernet services and applications are ubiquitously deployed, this architecture allows the cable companies to readily deliver a rich suite of IP services to their business customers using proven techniques.

Figure 1 shows the next-generation broadband architecture that uses Ethernet switching over the HFC network.

The Broadband Interface Unit (BIU) is the customer premises device that functions as an intelligent IAD. It supports 10/100BaseT interfaces to connect Ethernet attached switches or routers, a coaxial interface for TV and cable modem services operating below the 860 Mhz spectrum and a T1 or POTS interfaces for telephony applications.

available from Frame Relay and ATM networks. Outgoing Ethernet packets are modulated and transmitted over a 100 Mbps upstream channel on the coaxial drop. Ethernet packets are received at the BIU from the coaxial drop over a separate 100 Mbps downstream channel.

The Subscriber Access Switch (SAS) on the cable network combines the operation of a conventional Tap for existing services below 860 Mhz, and HFC Ethernet switching

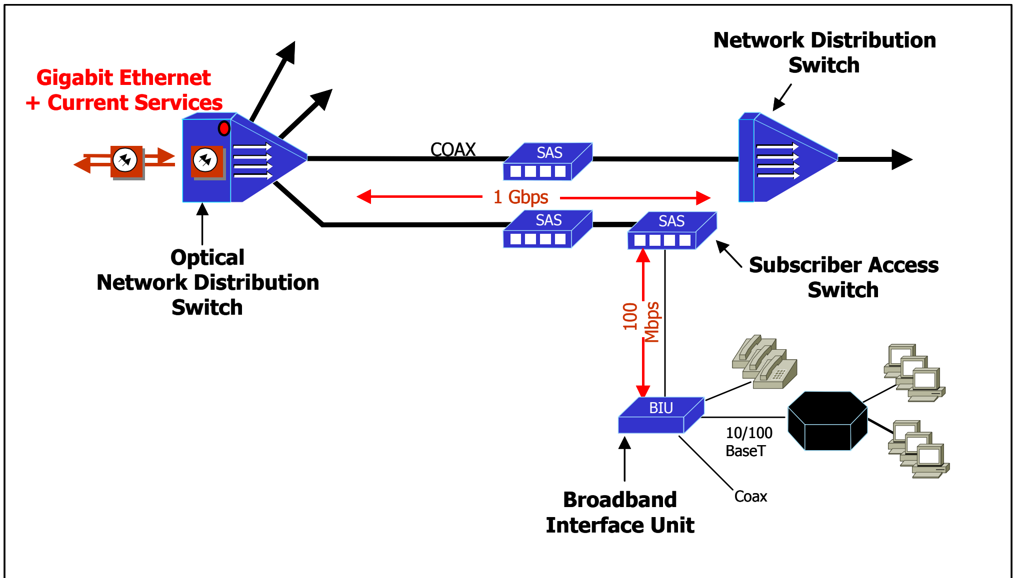

Figure 1: Next-Generation Broadband HFC Architecture

Traffic is classified and prioritized according to the services provisioned at the BIU. For example, voice over IP (VoIP) traffic is given the highest priority to minimize queuing delays. Mission-critical VPN traffic can be given the next level of priority while Internet bound traffic can be treated at an even lower priority level. The priority fields are appropriately set on an HFC header as well as in the IP header Type-of-Service (TOS) field so that the priority treatment can be extended through the access, metro, and wide-area networks. Traffic policing and shaping can also be enabled at the BIU to support Committed Information Rate style data services similar to those

functions for the high-speed data channels. It supports 100 Mbps full duplex subscriber data channels from customer premises and 1 Gbps full duplex trunk channels that connect to other upstream and downstream SASs over coaxial trunk cables. The Ethernet switching function supports priority queuing and weighted round-robin queuing so that multi-service traffic can be treated appropriately and all subscribers get a fair delay treatment independent of their distance from the headend. The SAS also supports a flow control scheme to avoid packet loss in periods of traffic congestion. The SAS also supports a bypass capability so that when local failures happen, all signals from the upstream trunk cable are directly coupled to the downstream trunk cable thereby maintaining service to downstream subscribers.

The Network Distribution Switch (NDS) provides the functions of a distribution amplifier for signals below 860 Mhz as well as trunk-to-trunk switching function between the 1 Gbps trunk channels. The optical NDS replaces a conventional optical node and combines the signals below 860 Mhz from the headend with the Gigabit Ethernet signals from the headend router. Switching functions, QoS, and flow control capabilities available at the NDS and the optical NDS are similar to those at the SAS.

By augmenting the existing HFC infrastructure with 100 Mbps and 1 Gbps Ethernet switching functions, and multi-service QoS capabilities, the proposed architecture provides all the benefits of an all fiber solution at a significantly lower cost. By adopting this architecture cable companies can leverage the ubiquitous nature of the HFC plant effectively.

NETWORK MANAGEMENT

Network management plays a key role in maintaining smooth operation of the network and in ensuring that the network is able to achieve the high standards of reliability demanded by commercial customers. The goal is to be able to remotely manage the network from a local or regional Network Operations Center (NOC) and to fully automate the functions of fault, configuration, and performance management so that truck rolls and other manual operations can be eliminated to the extent possible.

Each element in the new switched Ethernet HFC network is an IP addressable entity supporting an SNMP management agent and protocol. The SNMP management information base (MIB) supports all the operations necessary to manage the network element.

Figure 2 shows the architecture of the Network Management System (NMS). It includes an NMS applications server that hosts management applications and services, a topology server that discovers and maintains

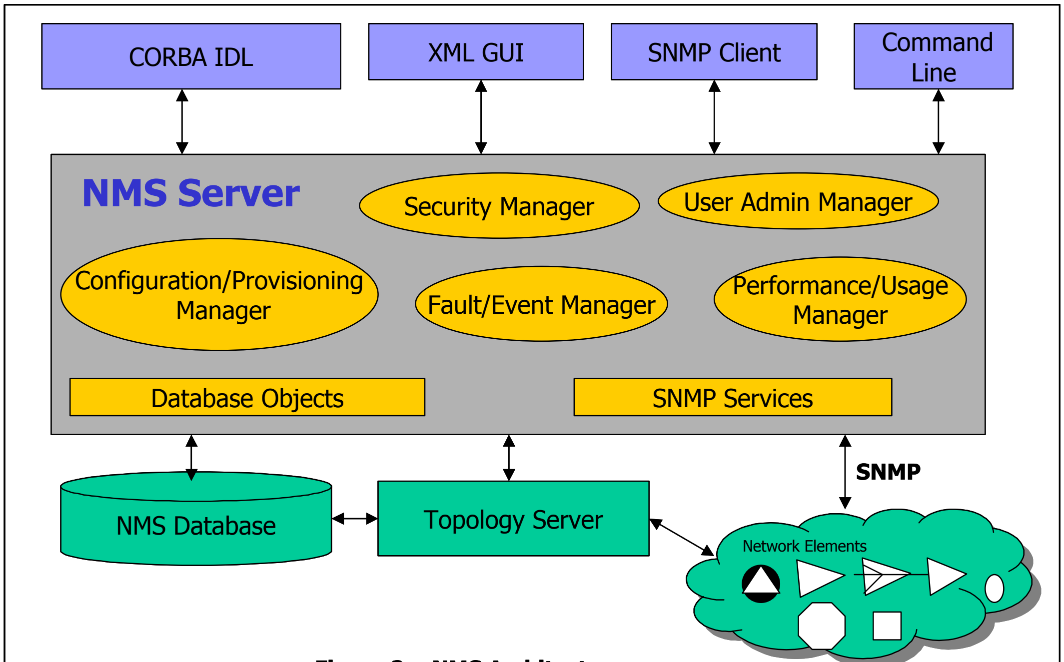

Figure 2: NMS Architecture

the network topology, and a database server that serves as the repository of network management information. The system supports a number of client interfaces including XML, SNMP, command line, and a Corba IDL for ease of OSS integration. This architecture provides maximum flexibility in distributing and scaling the management functions across multiple hardware platforms and work centers. The goal of the NMS is to integrate as many traditional back office processes as possible. It is only with full integration that the new HFC network will be able to serve the commercial market with timely and reliable services with availabilities at par with alternative fiber based delivery systems. The key management applications supported by the NMS platform are described below.

Configuration Management

When a new element is inserted into the network, it announces its presence and is acknowledged. The appropriate configuration files are then downloaded to the new element. The path taken by the announcement is captured by the topology server which allows it to discover the exact location of the network element. The information about a new element discovery and its location on the network topology is made available to various management applications. At this point, the network element becomes an actively managed element and is subject to all management operations. Thus the initial discovery and configuration of network elements is a fully automated process requiring no manual steps other than physical installation.

Fault Management

The network elements support a broad range of SNMP traps that communicate status changes of various types. The network administrator has the ability to turn these traps on or off and to direct them to specific operators. Alarms can be classified as minor, warning, or critical and can be used to generate and assign trouble tickets to various technicians. The fault management application also allows the network administrators to run a full suite of diagnostics and to perform various trouble isolation tests remotely to quickly diagnose fault conditions. The network elements themselves support a built-in self-test (BIST) capability for carrying out successful initialization tests. The network administrator can initiate BIST capability for fault diagnosis when required. The network administrator has the ability to remotely reset devices to clear fault conditions. These operations ensure that truck rolls are minimized and the mean time to detect and repair faults is as short as possible thereby improving network availability.

Performance Management

The NMS is capable of collecting a number of statistics related to the physical layer and data link layer. At the physical layer the operator can remotely collect and display power levels of the RF signals. In this system, gain and slope characteristics of the amplifiers can be adjusted either automatically or under operator control based on RF measurements. At the data link level, statistics on frame and byte counts and dropped packets can be collected by service classes. In addition to validating SLA performance, the various statistics can be used for accounting purposes and fed to the billing systems after appropriate processing. The network administrator has the ability to set what statistics are to be collected and at what intervals. As with other state of the art management platforms, the NMS system described is able to filter and present custom views of the network and its performance to individual users responsible for specific deliverables.

ENABLING COMMERICAL SERVICES

Enabling services for SMB customers on the proposed broadband network infrastructure is key to new revenue generation. The BIU platform becomes an intelligent service edge capable of delivering value-added functionality to the end user. Some of these services are described below.

IP Telephony

Telephony service is part of the core set of services demanded by commercial customers. While some cable companies have been offering residential telephony services using circuit-switched technology, this technology is expensive and difficult to scale due to its inefficient use of the frequency spectrum. SMB customers require upwards of four telephone lines on average and therefore the aggregate number of lines could easily become large with sufficient penetration rates. Cable modem technology using DOCSIS 1.1 standards would allow more telephone lines to be supported on the cable network by packetizing and compressing voice traffic. However, while adequate for residential customers, voice compression degrades performance and may not be acceptable to business customers.

The QoS enabled switched Ethernet architecture described in this paper presents an ideal platform for supporting IP telephony. The 100 Mbps / 1 Gbps link bandwidth is adequate to support a large number of telephone lines without requiring voice compression. Furthermore, the priority treatment at each hop ensures that end-to-end delays experienced by voice packets over the cable access network are in the sub-millisecond range from the premises to headend. Therefore, next-generation broadband technology has the potential to accelerate the deployment of toll-quality IP telephony.

From a telephony perspective, the intelligent BIU may function as an IAD by providing T1 or POTS interfaces to support on site PBX or analog phones. Voice traffic is packetized into VoIP format and transmitted over the Ethernet interface. Signaling information is transmitted using standard signaling protocols such a H.323, MGCP or SIP. On the network side, the BIU communicates with a VoIP gateway and a softswitch which shunts traffic to and from the PSTN. Value-added services such as Centrex can be offered by a feature server connected to the softswitch. Recent advances in the softswitch architecture are making it possible to support key business services without requiring an expensive Class 5 switch.

Internet Access

Many SMBs require a high speed Internet access service. This service can also be readily enabled by the intelligent BIU. Most cable companies have shown a preference for using policy based routing (PBR) as the mechanism to support multiple ISP access. The BIU can be provisioned with the appropriate routing and QoS policies for different source IP subnet and MAC address combinations based on the levels of service offered by different ISPs.

In addition to PBR, the BIU can support other options for enabling multiple ISP access. For example Layer 2 VLAN technology can also be used for this purpose. Each ISP subnet can be represented by its own VLAN domain and the BIU can be provisioned to map specific source MAC addresses to VLAN domains based on the ISP selected by the subscriber. Different QoS policies can be applied within each VLAN domain if the ISP wishes to offer different grades of service to its subscribers.

The BIU's ability to uniquely classify packet flows using any combination of IP address, TCP/UDP port, and MAC address, can be exploited to support fine-grained QoS features linked to the various applications and services offered by the ISP.

Virtual Private Networking

Small and medium size businesses require VPN services to link branch offices and to support remote access for telecommuters and work-at-home employees. VPN services can be enabled at the BIU by supporting various forms of tunneling mechanisms. Site-to-site VPN connectivity can be provided either by Layer 2 VLANs or by Layer 3 tunneling mechanisms such as IPSec. BGP/MPLS VPNs can also be enabled at the BIU. For telecommuter access, L2TP tunneling can be supported from the BIU without requiring any specialized client software. All forms of tunnels can be provisioned with QoS attributes to support multiple grades of service.

Storage Area Networking

The high bandwidth and low latency characteristics of a Gigabit Ethernet infrastructure can be exploited by the cable companies to offer new services such as storage area networking (SAN) to small and medium businesses. The incumbent service providers cannot offer such a service today because they lack a true broadband infrastructure.

SAN technology has emerged as a mechanism to provide servers with shared access to storage resources over a high-speed network. Currently businesses spend a lot of money to build dedicated SANs using highly specialized fiber channel networking protocols. By outsourcing this function, businesses can save significant amounts of money in both SAN hardware and software as well as in the IT staff that would manage the SAN infrastructure.

With the emergence of Gigabit Ethernet, there is a move to consolidate SAN applications over an IP network so that specialized networks can be avoided. The recent progress in SCSI-over-IP standards, and the development of iSCSI host adapters for servers and routers promises to bring SAN and IP networking together. With a ubiquitous broadband network, cable companies can exploit this opportunity and position themselves as the SAN service provider of choice to small and medium businesses in major metropolitan areas.

SERVICE DELIVERY PLATFORM

With the creation of a high-speed, symmetrical broadband access network, broadband services can be delivered from a metro service delivery center that attaches to the cable company's metro-optical network. This architecture is shown in Figure 3.

The metro service delivery center hosts a number of servers and gateways that support the kinds of services discussed earlier. For IP telephony, media gateways, softswitches, and feature servers provide value-added services and connectivity to the PSTN. Gateway routers provide connections to multiple ISPs. Content and storage servers are used to provide hosting and storage services.

In addition to servers and gateways, the metro service delivery center also hosts a service delivery platform (SDP) that supports provisioning, billing, customer care, and other back-office functions necessary to ensure smooth delivery of services. The SDP interfaces with a broadband services portal that is used by customers to order and provision services online and provides the interfaces to provision and control the gateways and servers that provide service functionality.

The main challenges in delivering services to end-users are:

- Ensuring that services can be ordered and provisioned in near real-time,
- Ensuring that new services can be rapidly created and deployed without requiring major software upgrades on network elements,
- Providing customers with a level of control in ordering and self-provisioning services and in monitoring their service status.

and the workflow, thereby allowing existing databases to be used and preventing unnecessary duplication of databases. By adopting a Java based framework similar to the NMS, the DEN applications and databases can be fully distributed leading to operational flexibility and scalability. Figure 4 shows the proposed DEN-based SDP architecture.

The main platform components of the SDP architecture are described below.

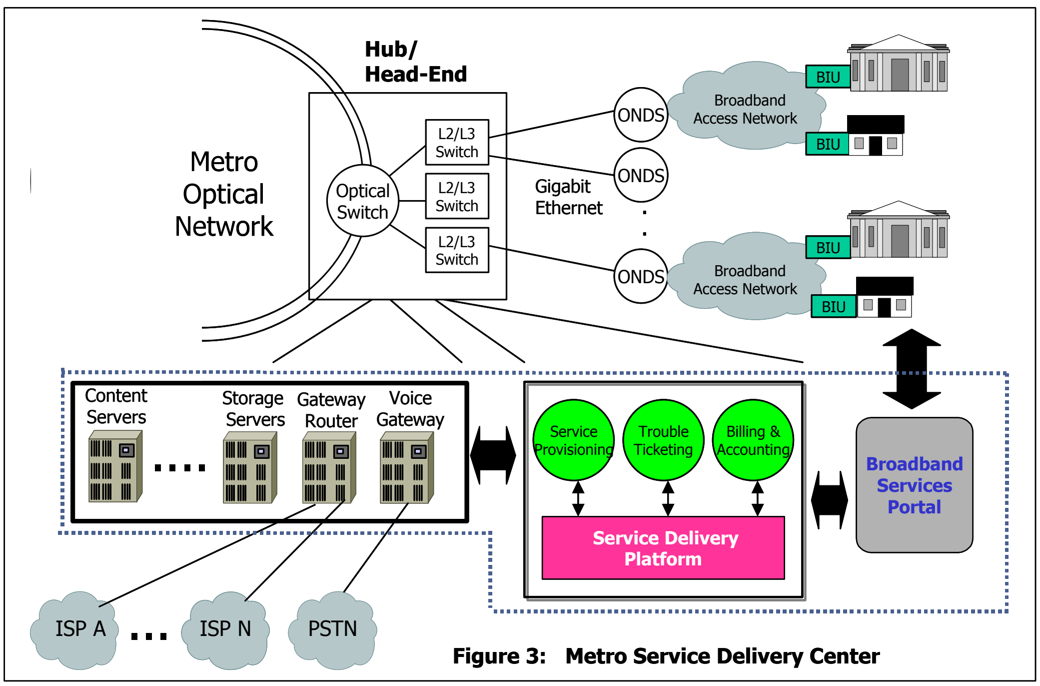

Figure 3: Metro Service Delivery Center

A Directory Enabled Networking (DEN) architecture for the SDP can properly address these challenges. DEN has been used successfully in enterprise networks to manage networks, service agents, and the client interaction with services in a policy-driven framework. By properly associating subscribers with network elements and services, the DEN model allows services to be rapidly created and provisioned. DEN decouples service data from the service logic

Service Creation Platform

The service creation platform (SCP) allows service providers to define the service logic associated with offering new services. The definition of service logic includes the description of a service template and a description of the provisioning workflow associated with the service. A service template describes all the characteristics of the service such as the service elements used by the service and a list of service parameters that need to be configured on each service element. The provisioning workflow describes all the tasks and the sequence in which they need to be executed to fully enable the service for the subscriber. Service creation can be a complex task requiring highly skilled personnel. A well-developed SCP is able to hide this complexity by providing the cable operator with a GUI and service creation object classes that makes service creation a simple and easy exercise.

LDAP Directory

The directory is simply a description of an information model and a data model that captures

based LDAP interface. Applications can retrieve any information related to subscribers, their services and their attached network elements through a unified LDAP query to the directory without knowing where the data is physically stored. This capability eliminates significant complexity in application design.

A key application of the directory is to interface to other existing back-office systems used for customer care and billing. This is done through the use of appropriate "adapter" technologies that allow data to flow between the systems. Through these adapters, existing customer databases can be accessed for customer account information, and the appropriate billing data can be passed to existing billing systems.

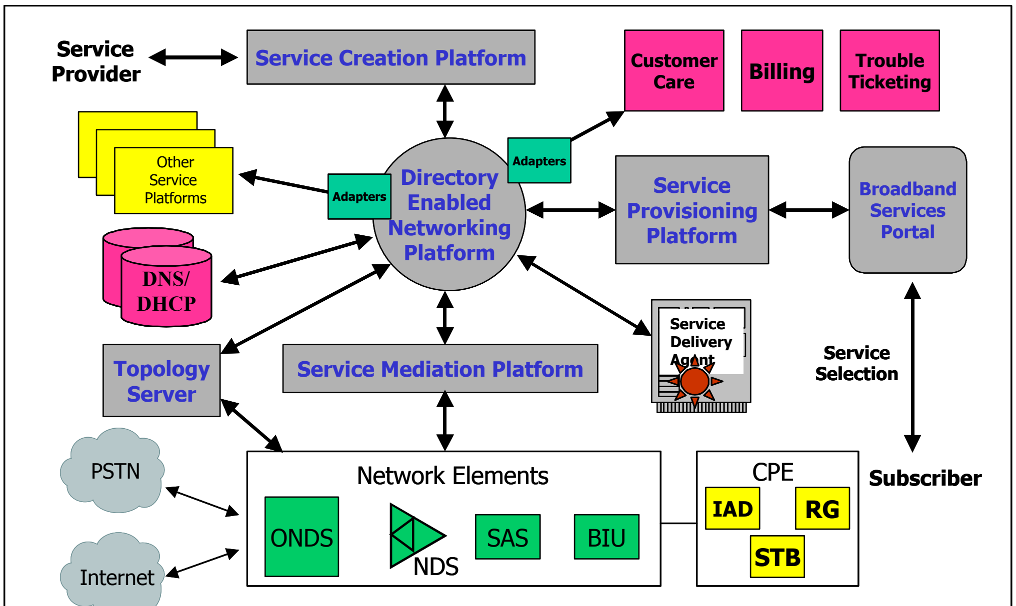

Figure 4: Service Delivery Platform Architecture

the relationship between subscribers, the network elements they are attached to, and the services to which they subscribe. The service templates and the provisioning workflows are also captured in the directory. The data model captures the data associated with a subscriber's service. The actual data can reside anywhere on the network in physical databases. The information model is exposed to external applications through a standards

Service Provisioning Platform

The service provisioning platform (SPP) is the major workhorse for executing all flow-through provisioning tasks. It includes a powerful workflow engine that can execute the tasks defined in the provisioning workflow, undo tasks when some of the provisioning steps fail, and update the appropriate back-office systems that support other service management functions when the provisioning steps are successfully executed. The SPP is a high performance platform capable of executing several provisioning requests simultaneously.

The SPP gathers subscriber provisioning data from the broadband services portal and builds the association with network and service elements by accessing the directory information model. It parses out the provisioning data according to the impacted service elements and then proceeds to carry out the provisioning tasks.

Service Mediation Platform

The service mediation platform (SMP) is used to manage the network elements during service provisioning. It hosts the service MIB that captures the management semantics of the particular service being provisioned. It takes provisioning commands from the SPP and translates them into the appropriate SNMP set requests to provision the device MIB on the BIU. By separating the service MIB from the device MIB, this architecture allows multiple services to coexist on the device without hard-wiring service level semantics into the device. By supporting a Java Virtual Machine (JVM) based service execution environment on the BIU, service logic can be downloaded at provisioning time and executed on the fly. This architecture embodies a truly dynamic service delivery environment.

A major function of the SMP is to perform admission control on both real-time services such as IP telephony and provisioned services such as VPNs. This ensures that adequate capacity is available in the network to support the service at the requested QoS levels so that customer SLA commitments can be honored.

SUMMARY

The next-generation broadband HFC architecture described in this paper unlocks the full potential of the most widely deployed asset of a cable company – the HFC network. By adopting a point-to-point switched Ethernet architecture using spectrum beyond 860 Mhz, previously unused bandwidth can be effectively utilized to create a high capacity two-way packet data channel. This multi-service QoS enabled channel can be used to deliver a wide range of IP based broadband services that were not possible before. In this new broadband network, services can be created on the fly and made available to the subscribers in near real-time through a sophisticated service delivery platform. Armed with such a state-of-the-art broadband infrastructure, cable companies can position themselves strongly in the marketplace, giving end-users affordable access to true broadband services, for the first time.