# Architectures for Broadband Residential IP Services Over CATV Networks

Enrique J. Hernandez-Valencia, Bell Laboratories, USA

## Abstract

The current state of the art in digital broadband access technologies to support emerging telecommunications services makes imminent the introduction of interactive broadband services — including data, video and the Internet — into the residential market. Over the last few years, much attention has been paid to the development of media access control protocols for cable TV networks that will allow the immediate support of broadband data services as the first step toward enhanced communications services for residential users. Here we review some of the architectural options that must be carefully considered in order to deliver IP services to such users in an efficient yet flexible manner.

Future residential cable data services are expected to deliver Internet access, work-at-home applications, small business access, local area network LAN-LAN interconnect, and LAN emulation services over cable TV (CATV) networks. These services are anticipated as a natural extension to the residential consumer market of the data networking capabilities in the business sector today [1]. Although related residential Internet Protocol (IP) services are already being trialed in the marketplace, substantive standardization efforts in this area did not materialize until quite recently. So far, the reference service model for such data services has been loosely fashioned from an initial request for proposals (RFP) by Cable Labs [2]. Over the last couple of years the IEEE Project 802.14 Working Group has been developing a set of specifications for physical and data link layer protocols [3, 4], including ADAPt+ [5], which will be applicable to supporting data service over cable networks and other broadband access infrastructures [6]. More recently, the Internet Engineering Task Force (IETF) has created the IP over Cable Data Networks (IPCDN) Working Group to address data service architectural issues similar, if not identical, to those applying to the initial Cable Lab's proposal. In order to ensure interoperability among vendors, a common data service framework that clearly specifies how IP-based services could be implemented over CATV and other similar access networks will be required. This service specification would be expected to identify required hardware/software components, signaling interfaces between them, and IP transport/routing services over the access network, as well as network provisioning, configuration, and management. Among the data networking issues, special attention should be paid to:

- Configuration of the physical and logical (IP) subnetworks
- Support for data forwarding/routing services, including IP Address Resolution Protocol (ARP) and the Internet Control Message Protocol (ICMP)
- Host address configuration
- Subscription
- Security

In addition, any proposed access architecture for broadband residential data services will be expected to support existing IP services such as the Dynamic Host Configuration Protocol [6], Domain Name System [8], IP Multicasting and Internet Group Management Protocol [9], Resource Reservation Protocol [10], Integrated Services [11], and IP Security [12], without much (if any) modification.

High-level system architectures for residential data services over CATV plants are depicted in Figs. 1a and 1b. On the customer premises side, a cable data modem (CM) with either an Ethernet or any other suitable PC interface provides the access interface to the CATV network. The CM connects the customer premises equipment to the standard CATV radio frequency (RF) coax drop cable. On the carrier's headend, (HE) the cable modem data termination system (CMTS) terminates the CATV RF link and implements the data link layer protocols in support of residential service. Given the broadcast characteristics of the RF link, multiple residential customers, and hence potentially many home-based LANs, will be serviced from the same CMTS interface.

The asymmetric nature of the CATV network, in terms of both access network topology and link rates, places special constraints on conventional point-to-point IP access models. Particularly for the so-called one-way CATV plant in Fig. 1a, the CATV access network would only deliver one-way (downstream) data transport services from the HE equipment in the cable data network (CDN) to the subscriber's home LAN. In

0890-8044/97/$10.00 © 1997 IEEE

such a scenario, a public switched telephone network (PSTN) link could be used to transport any (upstream) traffic from the subscriber to the CDN service provider. (Obviously, other residential access technologies may be used in place of the PSTN.) From the CDN service provider viewpoint, it is imperative that in this scenario most of the subscriber's downstream traffic is in fact relayed to the home LAN over the CATV access network, rather than the PSTN link (some two-way traffic may be required for connection initialization and control purposes). Additionally, as an enhanced service for one-way CATV plants, it may also be envisioned that the two-way PSTN link could be used as a CDN's access backup in the event of a CATV plant failure. In any event, it is further envisioned that the initial cable data service should be able not only to reuse exiting Internet technology, but also to provide a simple migration path to a two-way CATV infrastructure.

The traditional point-to-point dial-up access service model can more readily be adapted to the two-way CATV plant scenario. The main architectural difference from the conventional dial-up service is in the broadcast nature of the downstream RF link. Although the potential lack of privacy has never been of much concern when surfing the Internet, it is certainly troublesome when access to a company's private database is desired or when providing commercial transaction services over the Internet. Privacy has not been a showstopper with plain old telephone service (POTS) or integrated services digital network (ISDN) access to private networks since the link from the subscriber home to the corporate Intranet is over a dedicated circuit. For the existing CATV networks, enhanced data encryption capabilities will be required at either the media access control (MAC) layer or higher layers (e.g., IP). For the one-way CATV plant, however, in order to provide such an extra level of privacy, a mechanism is also required to identify which cable modems are attached to a given RF link. (This information would not be self-evident to the CMTS because the upstream traffic from the CMs will go to the homing CMTS over an interface other than the RF link.)

The next generation of IP (IPv6) is already being rolled out. An even more ambitious goal would be that the same broadband residential service architecture could be equally applicable to networks based on the current version of IP (IPv4) and to IPv6 networks. However, since the IP service models are somewhat different, it is natural to expect that their corresponding service framework over CDNs may also differ. Most of the discussion below is geared toward IPv4 residential services. Notice, too, that any IP service model over one type of CATV plant, say hybrid fiber coax (HFC) access networks, will share significant commonality with other asymmetric access infrastructures. Therefore, it can also be anticipated that the same service framework will be equally applicable to residential access technologies such as asymmetric digital subscriber link (ADSL), passive optical networks (PONs), wireless networks, and satellite networks. Similarly, it can also be envisioned that the PSTN fallback mode would make use of IP encapsulation and routing techniques akin to those under Mobile IP [13, 14].

From a data service perspective two important issues related to subscriber access are to be contemplated in any practical solution. First, some mechanism will be required to allocate IP addresses to the home appliances, particularly for situations where the CATV provider is not the Internet service provider. Second, mechanisms will be required to authenticate and secure data traffic to corporate Intranets. We will explore some of these issues below.

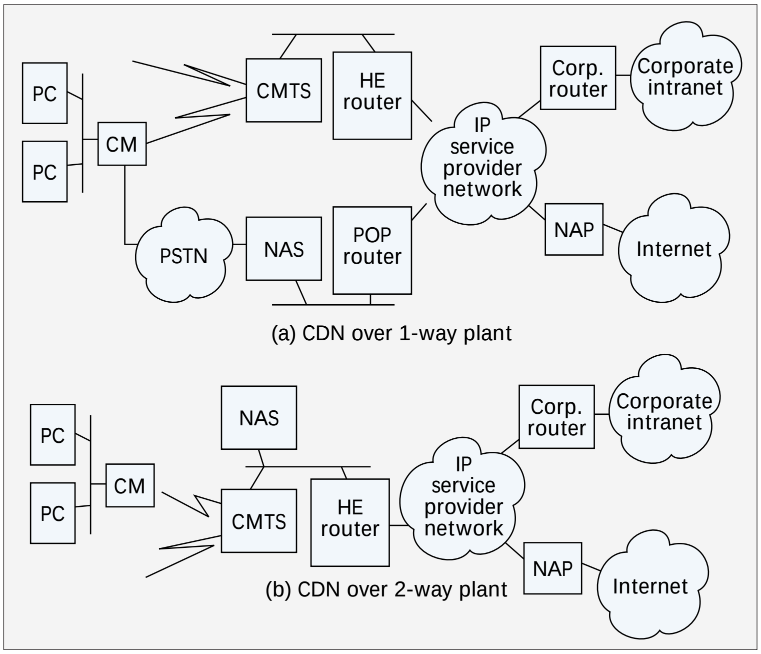

Figure 1. *a) A cable data network (CDN) over one-way plant; b) a CDN over two-way plant.*

## Generic IP Access Model

The generic access model envisions IP services over CATV networks as a one-to-one high-speed replacement of conventional IP dial-up services. For the two-way cable plant depicted in Fig. 2a that would certainly be the case, except for i) traffic filtering capabilities at the CMTS to eliminate broadcasting messages over all RF interfaces (this is certainly not a requirement, but it is highly desirable; otherwise, there could be significant throughput degradation), and ii) the potential support of data encryption at the MAC layer for improved data confidentiality. We will discuss security considerations later in this article.

In a one-way cable plant, however, the point of presence (POP) of the network access server (NAS, i.e., terminal server and related logical link control functions) controlling access to the networking facilities of an IP service provider (ISP) need not be collocated with the HE equipment controlling access to the CATV network. If only one IP address were given to a CM to support this service, an architectural choice must be made between placing the home LAN under the NAS' IP subnetwork or the HE/CMTS's IP subnetwork. Obviously, the choice would determine the path taken to forward data traffic to the home LAN.

### Cable Modem in the NAS Subnetwork

In this model, traffic to the home LAN would be forwarded first to the homing NAS in the ISP network. Since the CATV network is the desired last hop to the home LAN, downstream traffic would need to be redirected to the CMTS on the carrier's HE. Thus, some additional functionality would be required to implement this access model. One obvious approach can be readily borrowed from the data tunneling techniques used in Mobile IP [14]. Such an approach requires:

- An IP tunnel from the homing NAS in the ISP network to the homing CMTS/HE in the CDN on which downstream traffic to the home LAN is forwarded to the CMTS
- A (secure) mechanism to establish such tunnels

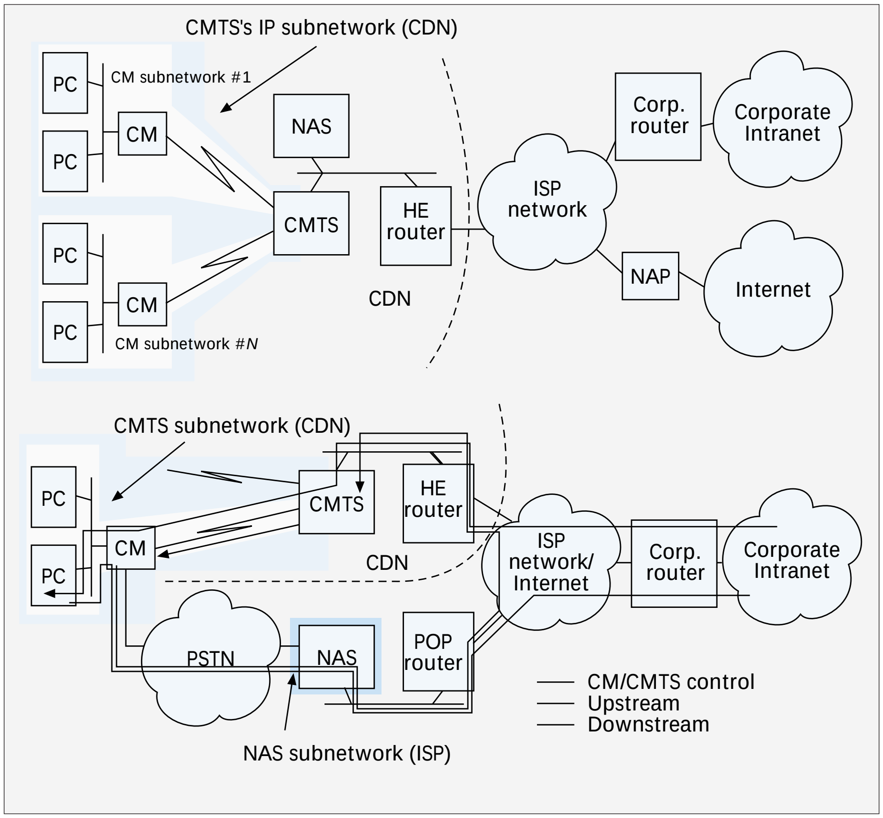

■ Figure 2. *a) A CDN architecture over a two-way CATV plant; b) a single IP subnetwork model for a 1-way CATV network with no PSTN fallback.*

- A (secure) mechanism to indicate to a NAS which users have access to particular CMTS's such that their traffic can be forwarded over the appropriate NAS/CMTS IP tunnel

Irrespective of the additional software complexity, there are clear performance drawbacks to this approach: the potential for a large number of extra forwarding hops from the ISP to the cable data provider, double encapsulation, and so on. The extra hops will not only introduce more latency into the data forwarding path, but also, given the high throughput anticipated for the downstream path, will significantly increase the amount of network bandwidth and processing resources required at the NAS and intermediate routers to cope with this service. One potential advantage of such an approach is that a fallback to dial-up access would be easier to implement for the case of RF link failure since most of the functionality to forward the user data to the homing NAS may already be in place.

### Cable Modem in the CMTS Subnetwork

In this model, traffic to the home LAN would first be forwarded to the homing CMTS on the HE LAN, which is already the preferred data forwarding path. Notice, as depicted in Fig. 2b, that this simple approach can support the minimum data forwarding functionality required to deliver IP services over the one-way CATV plant; but more functionality will be required if a fallback path through the PSTN is desired in the event of cable plant failures. In that case one of the following is needed:

- CMTS has to be able to identify homing NAS in the ISP network to open an IP tunnel to redirect downstream traffic.
- CMTS could source route IP packets to the CM (but source routing is not widely supported by router vendors).
- NAS in the ISP network needs a mechanism to map the incoming PSTN port to the CM's IP address for data forwarding purposes

Hence, by the time fallback access over the PSTN is incorporated into the picture, the hardware/software complexity of both approaches seems about even. However, the latter approach gives the most direct data forwarding path to the CATV network. One potential concern, though, is that upstream traffic may need to be redirected to the HE LAN under conventional (lack of) security practices.

## A Dual IP Subnetwork Model for a One-Way CATV Plant

A simpler CDN architecture that supports communications over either the access network or a PSTN can be achieved by associating different IP subnetworks with each of the different CM interfaces outside the home LAN. Specifically, the CM's interface over the CATV network can be logically associated with the IP subnetwork of the homing CMTS in the carrier's HE. This assignment would establish the logical point of attachment of the CM to the CDN. Similarly, the CM's interface over the dial-up network can be logically associated with the IP subnetwork of the ISP's NAS. This assignment establishes the logical point of attachment of the cable modem to the PSTN. Given that the CATV access network is the desired data forwarding path to the home LAN, rather than the PSTN link, the IP addresses to be used by the home appliances to communicate with their target destinations should have the IP subnetwork address of the homing CMTS as a prefix.

Hence, under normal operation conditions, CDN bound traffic from the home LANs would be relayed upstream by the CMs to the ISP network via the PSTN link. Traffic bound to the home appliances would always be forwarded through the homing CMTS since the home LAN would appear to the external world as an IP subnetwork of the homing CMTS/HE router, assuming that classless interdomain routing [15] is deployed in the CDN (Fig. 3a). The CMTS would then "relay" the traffic to the CM in the home LAN over the appropriate RF port, from where the CM would next relay this traffic to the target home appliance. Implicit in this model is the need for mechanisms that allow the CM to "discriminate" local traffic from CDN bound traffic, and match an RF port on the CMTS to home appliances. Also, one must define the data relay function provided by the CM and the CMTS.

Under an RF link failure scenario, the CMTS would not be able to forward IP traffic directly to the CM. However, if the CMTS knows the IP address of the CM on the PSTN, the CMTS can establish an IP tunnel to the CM to forward the home LAN bound traffic. The CM would then extract the traffic from the tunnel and "relay" it to the home appliances (Fig. 3b). The CMTS can learn the IP address a CM is using over the PSTN when the CM initiates the CDN connection request to the CMTS since, at that point, it must already have an IP address assigned by the ISP. (Note that in a one-way CDN the connection to the ISP network must exist prior to any possible exchange of user data between a CM and its homing CMTS.) Hence, besides the capability to establish a secure tunnel to the CM, there is no other major functional requirement placed by this approach.

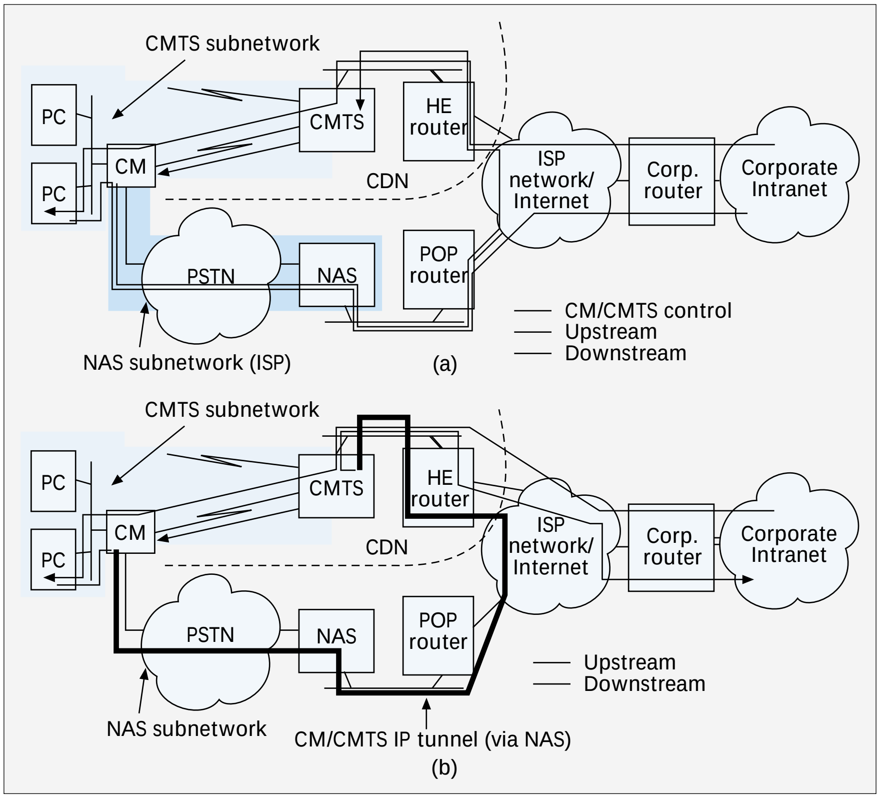

■ Figure 3. *a) Dual IP subnetworks for 1-way CDN — upstream/downstream data flows; b) dual IP subnetworks for 1-way CDN — fall back upstream/downstream data flows.*

## Bridge or Route

Both the single and dual IP subnetwork models for the CM on the home LAN allow the CM and the CMTS to operate as either intelligent repeaters/bridges or routers. Clearly, the particular choice would determine the routing and address resolution capabilities to be expected from the CM and CMTS, as well as the scaling properties of the CDN service with respect to the number of subscribers. We briefly consider some of the architectural issues below.

### CM and CMTS as Routers

With the CM and CMTS as routers, the CMTS need not be aware of the home appliance's MAC address. IP address resolution must be supported by the CM or some other device at the home LAN. The CMTS can treat its connection to the CM as a virtual point-to-point link. Home LANs could be configured as stub IP subnetworks sharing the IP subnetwork prefix of the tending RF link. This approach significantly simplifies operations, administration, maintenance, and provisioning (OAM&P) requirements for the CDN service provider since supporting data services over the CATV network, including IP multicast, would not be any different than supporting IP services over a conventional router network. Another benefit from this approach is that the IP routing table size for the CMTS can be made to grow in proportion to the number of RF links rather than to the number of CMs or home appliances, by making all home LANs share the homing CMTS's IP subnetwork prefix. Also note that, similarly, the IP routing table entries at the HE routers can be made to grow in proportion to the number of supported CMTSs. However, the total number of routing entries in a CMTS may still grow with the number of active CMs (which may not be much different from the number of active home appliances) if MAC-layer encryption is implemented to deliver some level of data confidentiality for broadcast tree-and-branch access networks such as HFC and PON-based access networks.

Here, the CM would be expected to support MAC address resolution on behalf of the home appliances. Thus, it should be able to either use Address Resolution Protocol (ARP) for the MAC address of the target home appliance, or provide any other mechanism to associate MAC addresses with home appliances' IP addresses. Although ARP support is common on most routers and UNIX-based systems, it is rather uncommon in most entry-level PCs for the residential market, which could become an important architectural consideration. For routing purposes, the CMTS would only be required to provide a mechanism to map the cable modem's IP subnetwork prefix to an RF port and exchange routing information with other routers in the same administrative domain, but not the CMs themselves.

Figure 4 shows the various communications stacks that could be involved in a plausible communications scenario for CDN services over a one-way CATV plant. Clearly, the router-based approach places substantial data networking functionality, and hence implementation complexity, on the CM. With the addition of support for IP tunnels to the CMTS, when PSTN fallback in the event of RF link failures is desirable, almost any conventional data service can easily be delivered to the home LAN. The cost of the CM and the scaling properties of the CDN with respect to IP unicast/multicast routing become the most relevant design and implementation concerns.

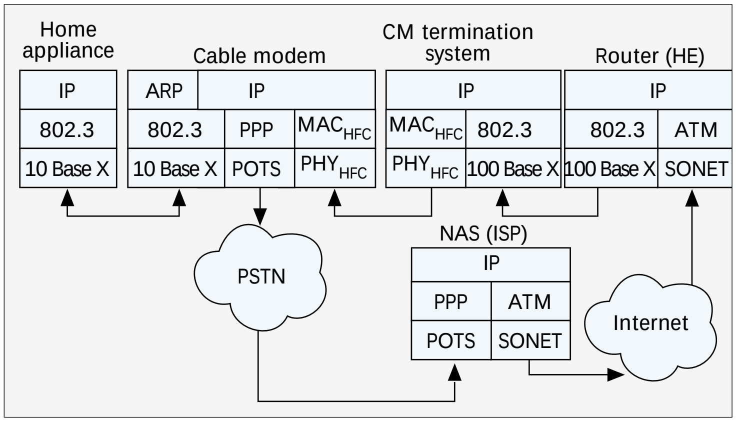

■ Figure 4. *A plausible communications model for the CDN over a 1-way CATV plant.*

### CM and CMTS as Bridges / Repeaters

The idea of bridging the home LAN to a carrier-based LAN in the HE is intrinsically appealing since it makes supporting data service to the home LAN no different than supporting those services over 802.3 LANs. However, with the CM and CMTS as a repeater/bridge, the routers on the carrier's HE must be able to resolve the home appliance's MAC address from its IP address. IP routing table requirements for HE routers would grow proportionally to the number of home appliances to be supported, which could be rather large (anywhere between a few hundred subscribers to tens of thousands per HE). The CMTSs would at least be required to filter traffic on an RF link basis; otherwise, they would needlessly broadcast traffic over all RF links, reducing the useful access network capacity. Forwarding cache tables on the CMTS would also grow proportionally to the number of active home appliances supported over its RF links. Additionally, the CATV link is not likely to be a true two-way share link. (The more likely scenario, according to most IEEE 802.14 proposals, is that it will be shared downstream, but point-to-point upstream.) Thus, the shared nature of the 802.3 MAC would need to be emulated by the CMTS.

On the other hand, the CM would only need to keep local home LAN traffic away from the access network. Thus, the CM forwarding cache table need only support a few tens of entries per home. IP address resolution need not reside on the CM. Since the bridged network could be rather large, ARP traffic could also become a significant fraction of CATV access network traffic. This issue can partially be managed by reducing the size of the logical IP subnetworks as seen by the routers at the HE, at additional OAM&P complexity for the CDN service provider.

> *Although the lower cost and complexity requirements on cable modes for 2-way CATV plants are clearly appealing, it is unclear that those savings also translate to cable modems for one-way cable plants.*

For one-way cable plants the situation is more complex. First, the upstream and downstream paths are not the same; hence, some external mechanism is required to populate forwarding cache tables at the CMTS. Second, ARP responses to an HE router's ARP request, if supported, cannot be forwarded blindly over the PSTN link as they are most likely to be filtered out by the ISP network (ARP is not a routable protocol). Bridging increases the amount of overhead, which decreases end-to-end throughput. Additionally, the slow PSTN link and the additional forwarding hops from the ISP network to the homing CMTS/HE could make responsiveness of ARP rather sluggish. Supporting ARP efficiently may require that the CMTS either i) acts as a proxy ARP server on behalf of its downstream home appliances, or ii) require the CM to forward ARP responses (when supported by the home appliance) over the CM/CMTS IP tunnel such that the CMTS can relay them to routers in the carrier's HE. ARP proxies may so be used in two-way CATV plants to keep ARP traffic off the cable plant. Notice that any information required to match an RF port on the CMTS to an IP subnetwork prefix for a CM could be either dynamically acquired from the CM as part of the CDN connection setup phase or statically provisioned.

The bridging approach allows for somewhat simpler implementation requirements for the CM at the expense of higher implementation and operation complexity at the CMTS and HE routers. It also provides the least desirable scaling of routing tables at the CMTS and the HE routers with respect to the number of subscribers. Although the lower cost and complexity requirements on CMs for two-way CATV plants are clearly appealing, it is unclear that those savings also translate to CMs for one-way cable plants.

## Static vs. Dynamic Address Allocation

The unassigned pool of public IPv4 addresses is quickly being exhausted given the explosive growth of the Internet over the last few years. Supporting millions of potential residential users would only exacerbate this problem. This design constraint applies to dial-up and cable data services alike.

### Statically Allocated Addresses

Draft proposals of CDN services still envision allocating IP addresses to home appliances, and even CMs, statically. That may be adequate for service demonstration purposes or initial service introduction, but it would not scale to large user populations using IPv4 since these addresses are already scarce. IP service providers are already required to justify any additional IP address allocation requests to the Internet Assigned Numbers Authority (IANA), and any static address allocation is strongly deprecated by the Internet address registers. Because of this practice, very few ISPs could afford to statically allocate public IPv4 addresses to their customers for free, particularly considering that the number of subscribed customers to most dial-up services is about 100 to 1000 times the number of simultaneously active customers during peak business hours. Thus, any static allocation of IPv4 addresses from the public[1] address space would most likely be deemed a liability.

A separate but relevant observation about the allocation of IP addresses is that neither user traffic constrained within the boundaries of a given ISP nor intermediate networking devices such as routers, CMTSs and CMs within a given CDN require public IP addresses. Private IP address allocation is an appealing approach to alleviate the reliance of data service providers on public IP addresses, even for relatively large subscriber populations, and extends the lifetime of static address allocation. Such implementations necessitate the deployment of network address translators (NATs) that translate private IP addresses into public IP addresses for traffic flowing into the Internet, as well as private address filters in the edge routers of the CDN to avoid advertising such private IP addresses on the Internet. Clearly, this approach would be highly intrusive on any ISP network not already supporting private IP addressing since it would require a significant overhaul of their transport network to provide the required access gateways for address translation and route filtering. This approach may be deemed by some as shortsighted. NATs do not enjoy much support within the IETF community as a long-term solution to the depletion of IP address spaces, but its popularity has been growing steadily given the lack of other mature IP addressing options. Note, however, that implicit to the usage of private IP addressing is dynamic allocation of public IP addresses somewhere at the edges of the CDN. Hence, the combination of statically allocated IP addresses to end users and NATs could be used as a convenient mechanism to transition to more dynamic address allocation, that is, from the network access gateways directly to the end user.

Of course, address space concerns are moot for any data services relying on IPv6. This approach will require even more

[1] *The term "public IP address" is used to refer to those network addresses that can be globally routed using dynamic IP routing from anywhere on the Internet. A "private IP addess" is one that is only valid within a given Intranet but cannot be advertised outside the intranet, thus, they cannot be accessed from the global Internet.*

significant infrastructure upgrades to data service providers, given the obvious need to interwork with IPv4-based networks. Although this functionality can be assumed to be available in most ISPs in the not too distant future, there are still multiple interoperability concerns that must be ironed out before such an approach may be considered commercially mature.

### Dynamically Allocated Addresses

Most ISPs are already architected for dynamic allocation of either public or private IP addresses. With an IP address from the public block, no further addressing/routing functionality is required from a CDN service provider to support access to the Internet. However, under IPv4, a CDN must still face the problem of address space depletion. For private Intranet access services, such as work-at-home, allocated IP address to home users may as well come from the private address block since most Intranet firewalls will perform address translation to satisfy some of their internal security concerns [16]. Thus, CDN service providers can leave to the corporate network policy the decision on whether such users would require a public/private IP address, or whether such an address should come directly from the corporate's IP address block (public or private), not from the ISP. For instance, CDN service providers could allocate IP addresses from the public block just for Internet access services. Such addresses would more likely be taken from the CDN's allocated public block. For Intranet access services, such as work-at-home, the allocated IP address could come from the Intranet's private block, assuming that Intranet-specific services are only allowed. If the corporate employee were also to be allowed to demand concurrent Internet and work-at-home access services, an IP address from the Intranet's public block could be allocated such that the employee can access the Internet directly *through* the company's Intranet. Note that the use of IP addresses from the private address space as discussed above does not necessitate the deployment of network address translators (but the use of route filters is still required).

For the two-way cable plant, it would be expected that, in most situations, the CDN service provider would also be the subscriber's ISP. For a one-way cable plant, the case is not that obvious, since the ISP may want to establish agreements with local CATV providers to deliver data services to the home. Either way, when the CATV provider and ISP are not the same, there will be a more pressing need to allocate IP addresses dynamically to simplify CDN operations. Otherwise, a significant amount of provisioning may be required to configure CMs, CMTSs, routers, and network access servers to add new subscribers into service.

From the above discussion, the following IP addressing/routing functionality in the provider network can be assumed as required to support dynamic allocation of public/private IP addresses:

- Dynamic host address configuration server for CDN attached hosts
- Network address translators at the edge of the CDN and/or the edge of a private Intranet
- Private IP address filters at the access points to the Internet (to be implemented by the service provider)
- Suitable security procedures and associated hardware/software such that the CDN user can interact with the NAS, the CMTS, and the Intranet's security server, if any

From the end user's perspective, dynamically allocating public/private IP addresses to CDN users (home appliances) requires:

- Dynamic host address configuration client to obtain IP addressees from the CDN. The CMTS or a configuration server/adjunct could interact with a DHCP server in the ISP network to acquire the IP addresses to be allocated to the CDN subscribers during the duration of a connection to the CATV access network.
- A yet to be defined mechanism to either pass IP addresses from the CM to the home appliances, if the addresses are initially allocated to the CM, or a mechanism to administer IP addresses directly between the CDN and the home appliances via the CM.

Other enhancements may also be needed to the POP software to provide differential treatment to Internet and work-at-home users.

By default, there is no intrinsic security afforded by the Internet data transport service. Such security must be provided as an add-on service by the CDN.

## Secure Corporate Intranet Access

Traditionally, home users have gained access to terminal servers on corporate Intranets by dialing up their remote access servers over point-to-point links (POTS or ISDN). This approach has allowed most corporations to reduce security concerns to an appropriate user authentication procedure as long as all data exchange occurred over a dedicated circuit-switched connection. Security in both user authentication and data confidentiality was assumed from the dedicated circuit-switched nature of the phone service. When corporate Intranet access is to take place over the Internet, the data exchanges take place over a distributed (shared) and open data networking infrastructure. By default, there is no intrinsic security afforded by the Internet data transport service. Such security must be provided as an add-on service by the CDN.

As before, it is useful to draw a distinction between public and private IP networks. When the corporate Intranet has public IP visibility, that is, when any IP host on the corporate network can be readily reached over the Internet, it will always be up to the corporate network to support strong authentication/encryption and related security mechanism at each host. Security in this case would be provided above the data transport layer, and hence can be deemed corporate, or application-specific. Protocols such as Keberos [17], a user-to-host authentication protocol, would fall within this category. A more sophisticated implementation of the same principle would direct incoming traffic to the corporate network to a security gateway at the edge of its private Intranet, where network address translation may also be performed. Both types of services would be functionally indistinguishable to the end user. Established secure Web access approaches such as Secure Socket Layer [18] and Secure Hyper-Text Transfer Protocol [19] would also fall within the general category of host-based security services. However, the Secure Socket Layer protocol only addresses the issue of authenticating end systems (host-to-host authentication) rather than users. Secured Hyper-Text Transfer Protocol, on the other hand, is only applicable to Web-based applications.

### Network-Layer Security

A more comprehensive approach to provide flexible data services to residential customers, which is applicable to public and private IP networks alike, is to provide the desired security services at the data transport layer. This solution can coexist with any of the host-/user-based authentication procedures mentioned above. Although the Internet community has been

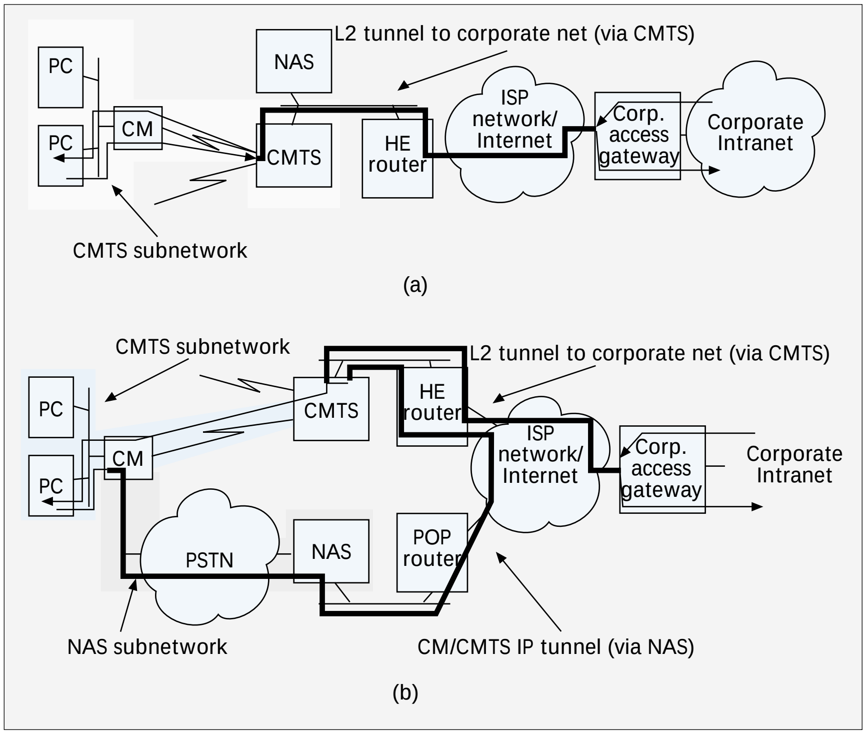

■ Figure 5. *a) L2 tunnel for remote user access over two-way CATV plant (L2 tunnel over two-way plant); b) one-way CATV plant with L2 tunnel for remote user access (dual IP subnetwork model).*

working in this direction for quite a while, no mature standard to deliver such transport-layer authentication and confidentiality services over the Internet currently exists. IPSEC, an IETF draft proposal currently under development, is the logical candidate to fulfill this role. Unfortunately, IPSEC is missing key components such as a public key management system. IPSEC will also require software upgrades to the communication stacks of any home-based appliance or corporate host, as well as the deployment of a public key management infrastructure. Of course, proprietary solutions are also viable; a strategy to migrate to IETF standards as they emerge would be needed.

An outcome of the current lack of a complete security framework for IP has been that, without significant hardware/software upgrades, most private networks today can only afford marginal support to external corporate traffic from unrecognizable IP addresses. This is prompted by the large variety of attacks to which they would be exposed if they were not to filter such traffic. Also, conventional networking practices strongly encourage large private companies to allocate their own IP addresses and to keep the internal details of their IP addressing plan away from third parties, including any external service provider. Hence, the ability to allocate IP addresses to end users from the corporate's address space, as well as the ability to manage such address allocation directly rather than differing IP address administration to the IP service provider can be anticipated to be critical to the deployment of secure private corporate network access over the Internet.

## Link Layer Security

Most authentication services (as well as IP address allocation services) for residential users are implemented at the data link layer via point-to-point protocols such as PPP. Any equivalent approach that preserves the current dial-up user authentication interface over other access technologies, such as CATV access networks, would be inherently appealing because it would require little if any modification to the existing base of customer premises hardware/software. Two very similar IETF proposals, the Point-to-Point Tunneling Protocol (PPTP) from Ascend, Microsoft, and US Robotics among others, and the Layer 2 Forwarding (L2F) from Cisco Systems, attempt to do so by defining a secure mechanism to tunnel PPP traffic over IP networks [20], hence extending the current dial-up access framework to allow virtual dial-up services to be supported over the Internet. The proposals provide mechanisms to establish a secure link between an NAS in an ISP network and a foreign access gateway, as well as the framework under which all standard services available under dedicated access configuration, such as authentication, authorization, address allocation and accounting, could be provided by the foreign network, separate from the equivalent services provided by the ISP.

Approaches that provide virtual dial-up access to remote users, such as PPTP/L2F, could be extended to the two-way CATV plan (Fig. 5a). There, it is assumed that the conventional NAS functions are split into two separate components: one still provided by the NAS on the ISP network and performing physical and media access control, configuration, and management functions; the other, provided by a remote access gateway on the foreign network, performing logical link control and network layer configuration and management functions, including PPP termination. Note that the CMTS could either only provide physical and MAC layer functionality to support data services over a CATV plant, for which a collocated NAS will be required, or be subsumed as one of the line interfaces of an NAS. The later approach may be more convenient since an ISP will prefer that managing the CATV RF link would not be significantly different from managing a PSTN or ISDN link. An IP tunnel would be established between the CMTS and the corporate access gateway, over which PPP traffic from the home LAN is forwarded. Note that minimal software upgrade would be required to current NASs and remote access gateways, as well as to the home appliances, since, it does not rely on a public key management system as does IPSEC. Hence, in principle, remote user access over a two-way CATV plant could easily reuse virtual dial-up access technology currently being developed by IETF.

It is reasonable to anticipate that any remote user access solution for a one-way CATV plant will also provide a simple migration path to a two-way CATV plant. This can be done trivially by allowing the CMTS to support the layer 2 tunneling protocol illustrated in Fig. 5b. Current IETF draft proposals for layer 2 tunneling expect all home traffic to flow over the CMTS-to-access gateway IP tunnel. Hence, a mechanism is required to forward corporate-bound traffic back to the CMTS, as opposed to sending it directly to the corporate access gateway from the NAS. Since the CM may have established an IP tunnel with the CMTS for backup purposes, the CMTS-to-CM communications protocol for a one-way CATV plant could easily be enhanced to support this additional functionality. Namely, the CMTS/CM must be able to multiplex/demultiplex traffic from/to the home appliance and traffic to/from the CMTS/CM itself. If necessary, two separate communication channels could be established for data relay and connection control functions. Although this seems rather convenient, because it reuses most of the functionality already anticipated for the CMTS, it does have some performance drawbacks; for instance, the most direct path between the home LAN and the foreign network would not be employed. Using such a direct path would require either a mechanism to authenticate the CMs with the foreign access gateways, which would deviate significantly from the current draft proposals for layer 2 tunneling, or a mechanism for the CM to act as a layer 2 tunneling proxy for the CMTS, which would also add additional complexity to existing remote access server proposals.

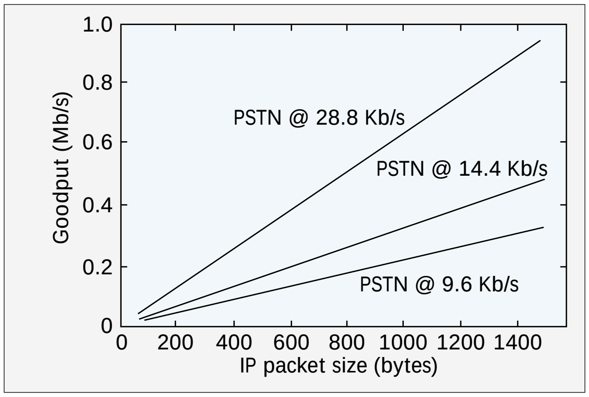

■ Figure 6. *Downstream end-to-end goodput as a function of packet size.*

## Practical Performance Considerations

In any reliable data transmission, end-to-end throughput is a function of the bandwidth-delay product, the error characteristics of the communications path, and the sender/receiver buffer sizes. The TCP transmitters regularly probe the data path and adjust their transmission window sizes to match the available network bandwidth. Most TCP receivers are expected to acknowledge every IP segment, or every other segment if implementing delayed acknowledgments [21]. Although the raw link rate for cable data services is anticipated to be about 30 Mb/s for downstream traffic, and about 1.5–3 Mb/s for upstream traffic, actual data rates can be constrained by configuration options for the various TCP/IP stacks.

The maximum window (WIN) and segment (MSS) sizes applying to a specific connection are negotiated during connection setup. TCP/IP hosts exchange the MSS and WIN values they are willing to support and choose the lowest MSS as the one applying to the connection. Since the MSS is typically derived from the maximum transmission unit (MTU) on the outgoing network interface, this process could lead to very inefficient use of the bandwidth when the upstream interface has a much smaller MSS value than the downstream interface. Also, since the data forwarding path changes after a reroute, there is a further potential for IP segment fragmentation. At a minimum, path MTU discovery [22] must be supported to avoid unnecessary IP segment fragmentation. Ideally, MSS renegotiation should also be supported (but most TCP/IP implementations do not).

Additionally, because of the asymmetric nature of the data forwarding paths over a one-way CATV plant, the reduced network bandwidth over the upstream PSTN link will throttle the rate of TCP acknowledgments to the TCP transmitter and further constrain the end-to-end goodput over the downstream link. Figure 6 depicts the maximum end-to-end goodput as a function of the packet size for common PSTN access speeds. End-to-end throughput, or goodput, can be improved by using standard techniques such as TCP/IP header compression [23]. Further goodput gain can be obtained by also taking advantage of PPP's various packet compression mechanisms.

## Closing Remarks

In this article we have reviewed high-level architecture design options available to deliver cable data services over both one- and two-way CATV plants. In the process, some of the performance trade-offs among design options were also identified. At this time, it seems feasible to deliver residential data service for Internet access, work at home, small business access, LAN-LAN interconnect, and LAN emulation services over two-way CATV networks with little modification to existing data networking protocols for conventional dial-up services. Although slightly more work is required to extend such services over one-way plants, which constitute a significant fraction of the deployed CATV distribution systems, it is anticipated that cable data networks will become more prevalent over the next few years as the by-products from the recent standardization efforts reach the market.

### Acknowledgments

The author would like to thank Donald Blahut, James Bobeck, Robert Dalton, Roger Levy, Harrison Miles, Robert Rance, and Ronald Sharp for their contributions to many of the views expressed in this article.

### References

[1] S. A. Grzelak *et al.*, "Residential Data Services via Broadband Local Access Network," *Bell Labs Tech. J.,* vol. 1, no. 1. Aug. 1996.
[2] Cable Television Laboratories, "High Speed Cable Data Services Request for Proposal," Apr. 1995.
[3] IEEE P802.14 WG, "Cable-TV Functional Requirements and Evaluation Criteria," IEEE 802 Committee, Feb. 1995.
[4] IEEE P802.14 WG, "MAC Convergence Proposal," IEEE 802 Committee, work in progress. Sept. 1996.
[5] B. T. Doshi *et al.*, "A Broadband Multiple Access Protocol for STM, ATM and Variable Length Packet Services on Hybrid Fiber-Coax Networks," *Bell Labs Tech. J.*, vol. 1, no. 1. Aug. 1996.
[6] J. W. Eng, "Standards for HFC-based Residential Broadband: IEEE Project 802.14, Its Mission, Charter and Status," *SPIE Proc.*, vol. 2609, Oct. 1995, pp. 2–9.
[7] R. Droms, "Dynamic Host Configuration Protocol," RFC 1531. Oct. 1993.
[8] P. V. Mockapetris, "Domain Name: Concepts and Facilities." RFC 1034, Nov. 1987.
[9] S. E. Deering, "Host Extensions for IP Multicasting," RFC 1112, Aug., 1989.
[10] L. Zhang *et al.*, "RSVP: A New Resource ReSerVation Protocol," *IEEE Network*, 1994.
[11] R. Braden, D. Clark, and S. Shenker, "Integrated Services in the Internet: an Overview," RFC 1633. June 1994.
[12] R. Atkinson, "Security Architecture for the Internet Protocol," RFC 1825. Aug. 1995.
[13] S. Hanks *et al.*, "Generic Routing Encapsulation," RFC 1701. Oct., 1994.
[14] C. Perkins (ed.), "IP Mobility Support," IETF Mobile IP WG, work in progress, 1996.
[15] V. Fuller *et al.*, "Classless Interdomain Routing (CIDR): An Address Assignment and Aggregation Strategy," RFC 1518. Sept. 1993.
[16] W. R. Cheswick and S. M. Bellovin, *Firewalls and Internet Security*, Reading, MA: Addison-Wesley, 1994.
[17] J. Kohl and C. Newman, "The Kerberos Network Authentication Server (V5)," RFC 1510, Sept. 1993.
[18] Netscape, "Secure Socket Layer Protocol," Mar. 1996.
[19] IETF Web Transaction Security WG, work in progress, 1996.
[20] IETF PPP Extensions WG, work in progress, 1996.
[21] W. R. Stevens, *TCP/IP Illustrated*, Vol. 1, Reading, MA: Addison-Wesley, Nov. 1994.
[22] J. C. Mogul and S. E. Deering, "Path MTU Discover," RFC 1191, Apr.1990.
[23] V. Jacobson, "Compressing TCP/IP Headers," RFC 1144. Jan. 1990.

### Biography

ENRIQUE J. HERNANDEZ-VALENCIA received his Electrical Engineer degree from the Universidad Simon Bolivar, Caracas, Venezuela, in 1982, and his M.Sc. and Ph.D. degrees in electrical engineering from the California Institute of Technology, Pasadena, California, in 1985 and 1987, respectively. He is a member of the Performance Analysis department at Bell Laboratories, Holmdel, New Jersey, where he works on issues related to the design and engineering of data communications networks.