IS THE DOCSIS CMTS SUFFICIENT FOR PACKETCABLE?

Burcak Beser *
Pacific Broadband Communications

Abstract

The DOCSIS 1.1 CMTS, as defined by the CableLabs, is almost ready for PacketCable. This paper discusses the additional functionality that will be required by a CMTS in order to satisfy the requirements of the PacketCable specifications. Equally important, this paper also identifies and discusses the external back office elements that are required for a successful PacketCable deployment.

This paper identifies the DOCSIS features that are used by PacketCable. Additional needs of telephony applications are identified, as well as an introduction to how these additional requirements are linked into DOCSIS features.

INTRODUCTION

PacketCable is a project instigated by Cable Television Laboratories, Inc. and its member companies. The PacketCable project is aimed at defining interface specifications that can be used to develop interoperable equipment capable of providing packet-based voice, video and other high-speed multimedia services over hybrid fiber coax (HFC) cable systems utilizing the DOCSIS 1.1 protocol. PacketCable defines a network superstructure that overlays the two-way data-ready broadband cable DOCSIS 1.1 access network. The initial phases of PacketCable cover only voice communications. This paper only addresses the issues regarding PacketCable Cable Modem Termination System (CMTS) support for carrier grade voice over IP (VoIP).

PacketCable assumes operation over DOCSIS 1.1, which adds features to the basic DOCSIS 1.0 release capabilities in the areas of managing and packaging of the data services. PacketCable augments the basic back-office elements such as Data Provisioning and Data Management servers of DOCSIS, with comparable equivalents for VoIP services. A Record Keeping Server, Call Management Server, and Gate Controller are shown in Figure 1.

PacketCable requires the following additional protocols and functionality from a DOCSIS capable CMTS:

- Theft of Service Prevention
- Quality of Service
- Legislative Support

THEFT OF SERVICE PREVENTION

The basic assumption regarding theft of service prevention is that the Multimedia Terminal Adapter (MTA) is not resistant to customer tampering, and that the incentive to illegally obtain free service will lead to some very sophisticated attempts to thwart any network controls placed on the MTA. This customer tampering includes, but is not limited to: opening the box and replacing ROMs; replacing integrated circuit chips;

* e-mail: burcak@pbc.com

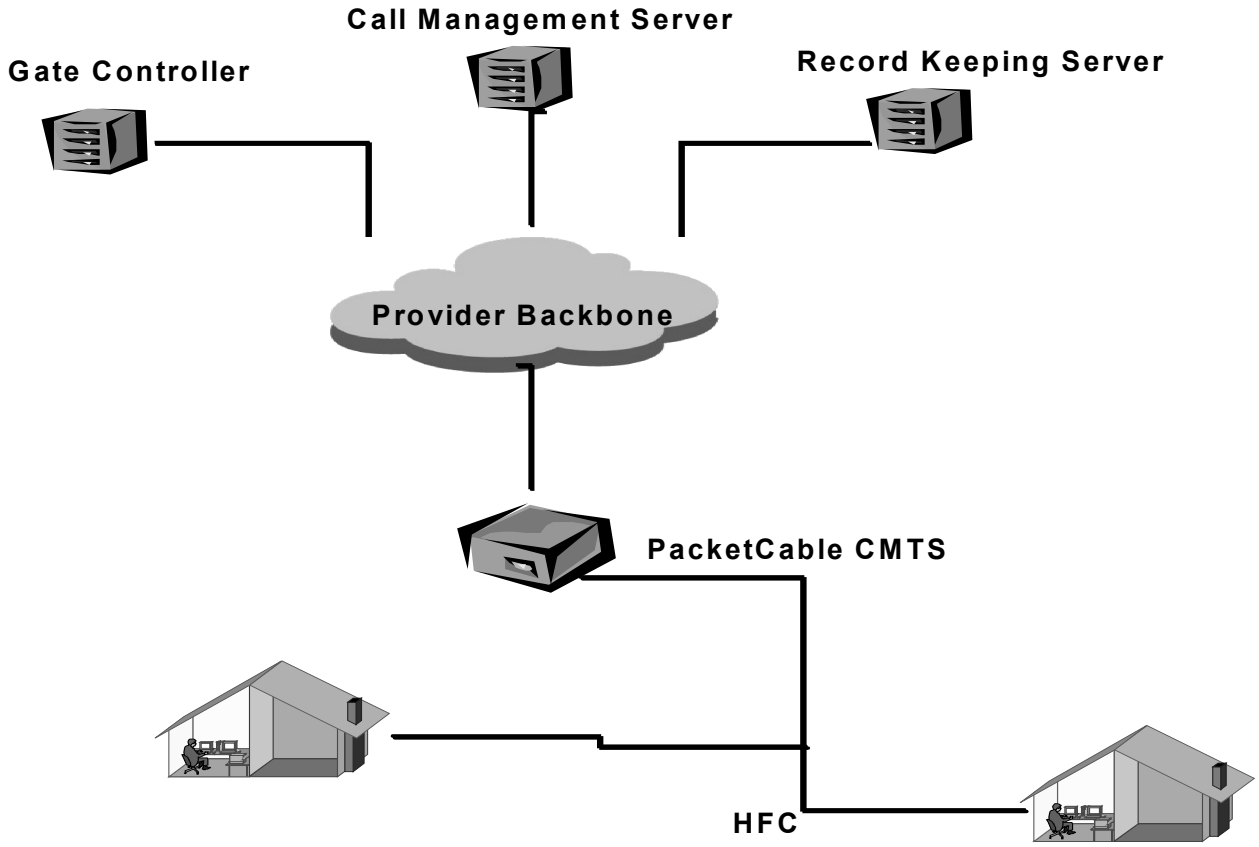

Figure 1 PacketCable CMTS and back office elements that are connected

probing and reverse engineering of the MTA design; and even total replacement of the MTA with a special black-market version. Examples of this degree of effort can be found in various other industries and technologies.

Since an individual MTA can be distinguished only by its communication over the RF network, it is possible, and quite likely that PC software may be written that will emulate the behavior of any MTA. In such a case the PC may be indistinguishable from a real MTA. In this case the software is under the total control of a customer.

Customers establishing high level QoS connections themselves

The MTA with sufficient intelligence can remember past destinations dialed as well as the destination address, or use some other mechanism to determine the IP address of the destination. It can then signal that destination itself (with some cooperation of the far-end client), and negotiate a high level quality-of-service connection via the RSVP mechanism or via the interface for an embedded client. Since no network agent is used in initiating the session, no billing record will be produced.

Even though the above scenario requires the cooperation of two altered MTAs, it is possible to achieve the same effect by manipulating/modifying only the originating MTA. If the originating MTA used the network agent to establish the session, thereby informing the destination in the standard manner of an incoming session, but again negotiated the high quality-of-service itself, there would be no billing record generated and the originator could obtain a free session.

Prevention of this scenario is accomplished by requiring per call authorization at the CMTS; without the proper authorization, any attempt to obtain carrier-grade quality-of-service to make the phone call will fail.

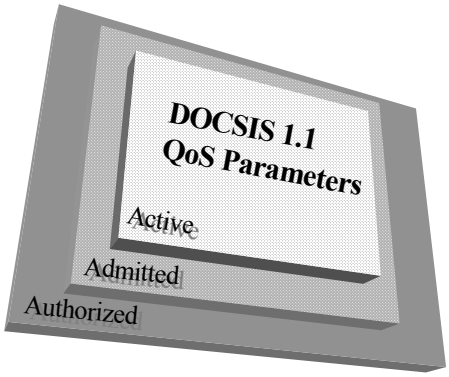

Figure 2 DOCSIS 1.1 QoS Envelopes

The DOCSIS 1.1 specification defines 3 envelopes that can be defined for an IP (service) flow: Authorized, Admitted, and Active. The "Authorized" envelope defines the boundaries of the QoS parameters that can be used by the initiator for specifying an IP flow. The "Admitted" envelope shows that a particular IP flow is admitted, and "Active" is the state in which the IP flow can be used with the QoS that was admitted. The relationship of these envelopes are shown in Figure 2.

This per-call authorization model requires the PacketCable CMTS to have an external means of setting the "Authorized" envelope by the Call Management Server (CMS). Since the Call Management Server knows when a call is in progress, the above-mentioned theft of service methods would be prevented.

The PacketCable DQoS specification includes the external interface that defines the "Authorized" envelope as Common Open Policy Service + (COPS+). The COPS protocol defined by the PacketCable has sufficient differences that a standard COPS client/server is not sufficient for the PacketCable needs.

The COPS+ protocol specification requires an extension to the objects and security mechanisms. The object that is manipulated by the gate controller via the COPS+ protocol is called a *PacketCable Gate*. The *PacketCable Gate* is a special construct that controls the "Authorized" envelope, and is at the heart of theft of service prevention in a PacketCable CMTS.

Customer alteration of the destination address of voice packets

Another theft of service scenario is shown in Figure 3. Two remote MTAs that are far apart, each make a local call. Once the bandwidth and connection for these local calls has been established, the MTAs change the destination addresses to cause their VoIP streams to point to each other. The billing system continues to bill each of them for a local call, while the customers are actually engaged in a long distance call.

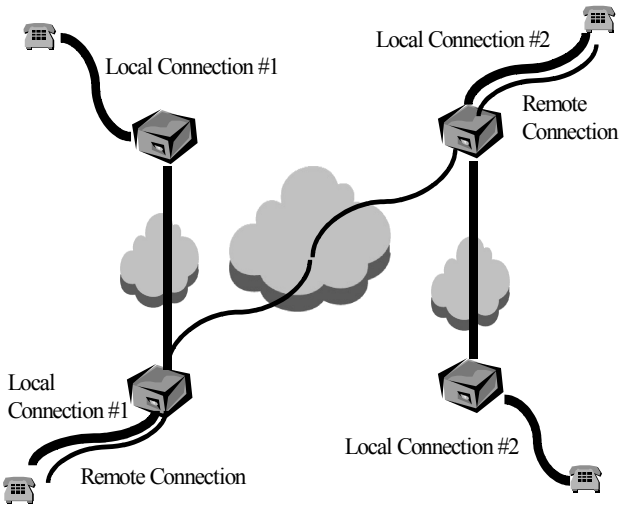

Figure 3 Using two local connections to achieve a remote connection

The solution to this theft of service mode has two parts. First, the previously defined authorization interface (PacketCable COPS) now should pass information regarding the IP addresses of the far end. Second, the CMTS should police the IP flow for compliance.

The DOCSIS specification does not mandate the use of egress policing of classifiers for the DOCSIS CMTS. A PacketCable-capable CMTS should implement the optional feature of policing egress IP flows.

MTA non-cooperation for billing

One can easily imagine what would happen if there was a message from the MTA on session establishment that said, "OK, called party has answered, start billing me now," or a message on hang-up that said, "Session has completed, stop billing now." However, there are more subtle ways that a user could have the same effect as tinkering with such messages if they existed.

With the current PSTN scheme, users are billed for the entire timeframe that they spend actually connected, but they are not billed, for example, for the 30 seconds that the far-end phone was ringing. As shown in Figure 4, the billing must be connected to actual VoIP QoS usage. Whenever the PacketCable CMTS activates a service flow, which is connected through a *PacketCable Gate*, it sends an event message to the Record Keeping Server, which is later used for billing purposes.

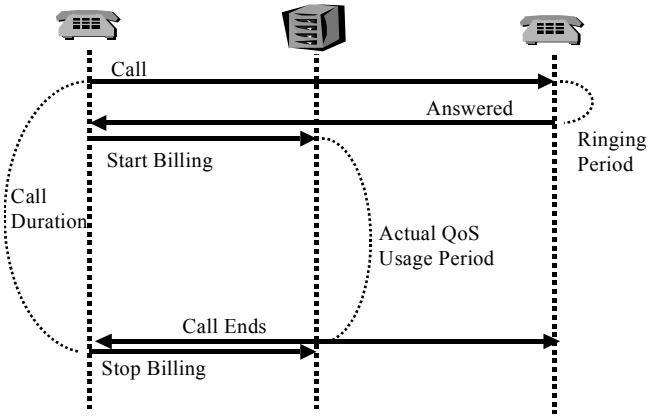

Figure 4 Call Billing

This state by which the connection is made and billing started is referred to as *PacketCable Gate Open*. The same event messaging occurs when the created service flow is deleted at the end of the call. This behavior is different from the DOCSIS standard where no messages are generated on the WAN interface when a service flow is activated or deleted.

The PacketCable specification defines the above-mentioned event-messaging protocol as Radius. The Radius protocol defined by PacketCable has sufficient differences from that of a standard Radius client/server and is not sufficient for PacketCable needs, and an additional content and security mechanism is specified by PacketCable.

Use of half-connections

In this theft of service scenario, one of the MTAs (originator) in the call signals the start of conversation but the MTA at other end (termination) does not signal the start of conversation at its end.

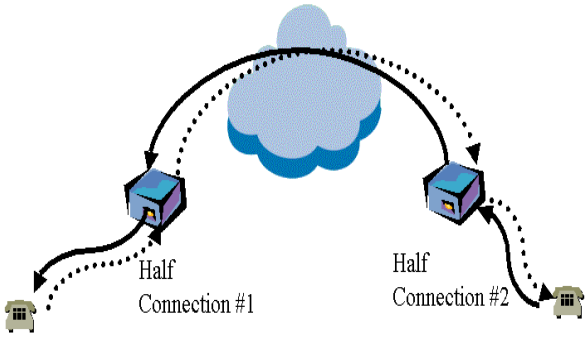

Figure 5 Using two half connections for bi-directional call

In this case, only one *PacketCable Gate* is opened, and the users and network are left with a "half-connection". Given that the terminating MTA did not send a "conversation-started" message, the network cannot legitimately bill the user for the half-connection. However, it is possible for two colluding clients to set up two half-connections, neither of which is billable, which can be combined to give a full connection between the two parties.

This is an example of theft of service that could occur in the absence of resource use coordination at both ends of the call. Coordinating the operation of the two *PacketCable Gates* can prevent fraud of this type.

Whenever a *PacketCable Gate* is opened (thereby granting carrier grade QoS on the Cable segment) a message is sent to the far end PacketCable CMTS that the half connection is started, and if it does not receive a message that indicates the other end has committed to the call by some certain time then the call will be terminated.

The PacketCable specification defines this *gate-coordination* protocol as Radius with PacketCable specific content and security mechanisms.

The theft of service attack defined above can also be executed by early termination of the connection by one end only. In this case the Records would show very short call duration. This mode of attack can easily be prevented by coordinating flow deletion.

PacketCable defines that the same *gate-coordination* protocol used above will be used for flow deletion as well.

Fraud directed against unwanted callers

Due to details of the call setup sequence, it is possible that bandwidth authorization at the destination will be more generous than that at the source. Given this, it is then possible for a called party to reserve and allocate bandwidth far in excess of the final negotiated amount, resulting in the calling party being charged more than expected. If available, this technique would likely be used against telemarketers by individuals in response to unwanted calls during dinner.

For this reason the *gate coordination* defined by PacketCable includes the examination of the bandwidth that is being used.

QUALITY OF SERVICE

The use of DOCSIS 1.1 enables PacketCable systems to provide carrier grade QoS for VoIP communications. Some of the properties of DOCSIS 1.1 that are applicable to MTAs can be summarized as:

- Multiple service flows, each with its own class of upstream traffic
- Both single and multiple voice connections per DOCSIS service flow
- Prioritized classification of traffic streams to service flows.

For a DOCSIS 1.1 CMTS the following are the QoS settings that play a crucial role in providing carrier grade QoS:

- Guaranteed minimum/constant bit rate scheduling services
- Constant bit rate scheduling with traffic activity detection service (slow down, speed up, stop, and restart scheduling)
- DOCSIS packet header suppression for increased call density
- DOCSIS classification of voice flows to service flows
- TOS packet marking at the network layer
- Guarantees on latency and jitter
- Reclamation of QoS resources for dead/stale sessions
- Two stage QoS setup

The primary mechanism for providing low-latency quality of service for media streams in the access network is the DOCSIS 1.1 flow classification service. This service classifies packets into specific flows based upon packet fields such as IP source and destination addresses and UDP port number parameters. In the upstream, such classified packets are transported via an appropriate constant bit rate service (for current codecs), as dynamically scheduled by the CMTS. In the downstream the packets are transported via an appropriate high priority queuing and scheduling mechanism. DQoS (between CMS and CMTS) and DOCSIS (between CMTS and CM) signaling mechanisms are used to dynamically set up the media stream flow classification rules and service flow QoS traffic parameters.

Providing Timely Call Setup

During call setup a number of messages are exchanged between various entities. Of these messages, the QoS setup messages are handled by the CMTS. Since the QoS messaging takes place in real time while callers wait for services to be activated, the protocol that is used for QoS setup must not impose unnecessary delays. The number of messages, which traverse end-to-end, should be minimized. For this reason PacketCable partitions the resource management into two distinct segments: access and backbone. Segmented resource assignment is beneficial for two reasons:

- It allows for different bandwidth provisioning and signaling mechanisms for the originating network, the far end network, and the backbone network.
- It allows for resource-poor segments to maintain per-flow reservations and to carefully manage resource usage. At the same time, when backbone segments have sufficient resources to manage resources more coarsely, it allows the backbone to avoid keeping per-flow state, and thus enhances scalability.

DOCSIS only specifies QoS control for the cable segment through DOCSIS MAC messages. DOCSIS MAC messaging is useful only if the QoS requesting entity is directly connected to the RF link. If a client is connected through an IP network behind a cable modem, then that device cannot request QoS from a DOCSIS CMTS.

For this reason, the PacketCable specification uses RSVP+, which is a modified version of the RSVP protocol. Using this modified RSVP protocol it is possible for all client devices to request high quality IP links from the PacketCable CMTS.

For the IP backbone, the PacketCable specification is more relaxed and allows the use of a number of protocols including RSVP, DiffServ and Aggregated RSVP. Due to limitations on the scalability of both RSVP and of DiffServ, it is expected that aggregated RSVP will be used for signaling IP backbone high quality links.

Providing low delay voice transport on a Cable Network

For voice services, the end-to-end packet delay needs to be small enough that it does not interfere with normal person-to-person interactions. For normal telephony services using the PSTN, the ITU recommends no greater than 300 ms roundtrip delay. Given that the end-to-end backbone propagation delay may comprise a significant percentage of this delay budget, it is important to control delay on the access channel, at least for long-distance calls.

DOCSIS 1.1 specifies a new scheduling service called Unsolicited Grant Service (UGS) in order to reduce the delay that is introduced by the upstream cable segment. Unfortunately just the existence of UGS is not sufficient to minimize upstream cable network delay. The generation of the VoIP packets has to be synchronized to the UGS schedule, which is itself synchronized to the DOCSIS timestamps.

On the far end the VoIP client will play out the incoming data stream. The far end can be another VoIP client or can be a PSTN gateway, which converts the VoIP packets into a PSTN data stream (a DS-0).

In PSTN networks the voice samples are transmitted using common reference timing to which all the PSTN gateways should be synchronized. Due to the sampling synchronization, the incoming frames must be synchronized to UTC as well. If the synchronization does not take place, the playout buffer would underflow or overflow over time, which would result in pops and clicks and would effect fax/modem communications carried as VoIP data.

The DOCSIS specification does not mandate that the DOCSIS CMTS be synchronized to any external clock source, and due to frame synchronization issues, the PacketCable CMTS should be synchronized to the PSTN common timing reference.

LEGISLATIVE

The PacketCable system should be ready to adopt the same legislative requirements that are in place for PSTN systems. Two of these requirements, electronic surveillance and privacy, are specific to a PacketCable CMTS.

Even though it is possible to achieve compliance outside of the CMTS, the PacketCable CMTS is at the right place to achieve compliance without unnecessary overhead.

Electronic Surveillance

Electronic surveillance includes both interception of communications and the acquisition of call-identifying information. Since the call-identifying information is not retained in the CMTS, only communication interception should be handled by the PacketCable CMTS.

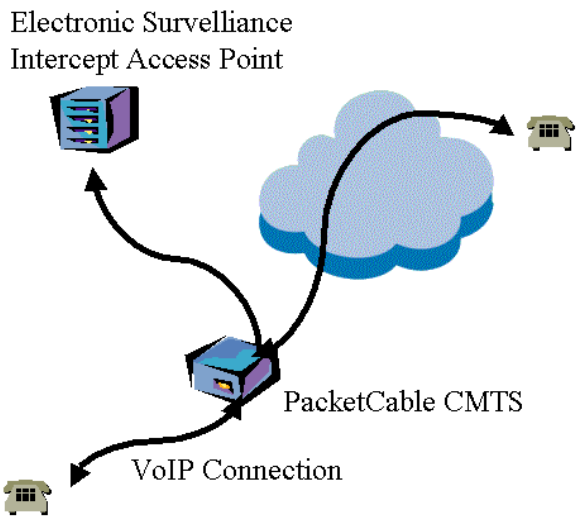

Figure 6 Electronic Surveillance using PacketCable CMTS

A PacketCable CMTS that supports electronic surveillance is responsible for duplicating the VoIP packets of a specific IP flow and sending them as defined in the PacketCable Electronic Surveillance Specification to a predetermined end-point, such as to a local Electronic Surveillance Intercept Point at a local federal law enforcement office as shown in Figure 6.

Privacy

The privacy issue on VoIP systems has multiple dimensions: One dimension is the telephony representation of caller ID, which is achieved by careful design and implementation of the PacketCable call signaling protocol.

The second dimension is IP address privacy. This issue arises from the fact that the IP address that is contained within IP packets can be used to determine the location of the caller. Today there are multiple systems that can accurately pinpoint the location of any given IP address.

IP address privacy can be achieved if the PacketCable CMTS performs Network Address Translation on VoIP packets.

There is a beneficial side effect that comes from double NAT: The telephony devices do not have to be globally routable. Having telephony devices that are in the private address space would help greatly to alleviate the depletion of routable IP addresses, especially if one thinks that there would be millions of such devices.

IN SUMMARY

Even though DOCSIS 1.1 is an essential baseline for PacketCable support, a DOCSIS 1.1 CMTS alone is not sufficient for supporting carrier grade voice transport over the Internet.

The PacketCable CMTS needs to support additional protocols and functionalities for the areas of theft of service prevention, carrier grade quality and legislative issues.