# MOVING SET-TOP BOXES TO THE IP INFRASTRUCTURE

Doug Jones
YAS Broadband

*Abstract*

*Moving the STB to the IP infrastructure is a whole lot more than just delivering video over IP. In fact, of the five areas of consideration discussed in this paper, video over IP is probably the furthest from happening. When migrating STBs to the IP infrastructure, the first steps will include solving embedded CM provisioning and moving OOB transport and interactive applications to the IP infrastructure. Video over IP will follow when the economics prove viable.*

## INTRODUCTION

### Business Case

The near-term drivers for moving the STB to the IP infrastructure are neither Internet access nor delivering video over DOCSIS. Rather, moving the STB to the Cable IP infrastructure is about reducing both capital and operations costs by placing all services (data, voice, and the video platform) on the same network infrastructure.

The cable industry is migrating from proprietary technologies to using standard technologies based on the Internet Protocol (IP). It started when Cable operators began building an IP infrastructure to provide high-speed access to the Internet. Most recently this IP infrastructure will be used to provide telephone service using Voice over IP (VoIP) technology. Running a second service on top of a single infrastructure provides many business advantages. Specifically, the network is already there and does not have to be build from scratch. Equipment is available from multiple suppliers. Technicians are already trained for the network equipment.

What remains to happen is the migration of the video business to the common IP platform used by these other Cable services. Not only will additional cost savings be realized by making use of the existing infrastructure, there is more innovation and less operational expense associated with the existing IP infrastructure. The move does not include Internet access; rather, the move will have the STB business running on IP networks.

The IP infrastructure has proven to be both general purpose (supporting many different applications from web and email to VoIP and streaming media) and low cost. IP has stood up in the face of competition from other networking technologies such as AppleTalk and Novell Netware. IP technology just works, is widely deployed and available from multiple suppliers.

### Areas of Discussion

Providing either HSI or VoIP service through the STB is not discussed in this paper. Rather this paper discusses five areas when considering moving the STB to shared IP infrastructure. These are:

1. Embedded CM Provisioning
2. Out-of-Band Transport
3. Interactive Applications
4. Video Delivery
5. CMTS considerations

The section on CMTS considerations is include because items 2, 3 and 4 place traffic on the CMTS which could also be carrying HSI and Telephone service.

## OVERALL SYSTEM VIEW

A system view with the five areas of discussion is shown in Figure 1.

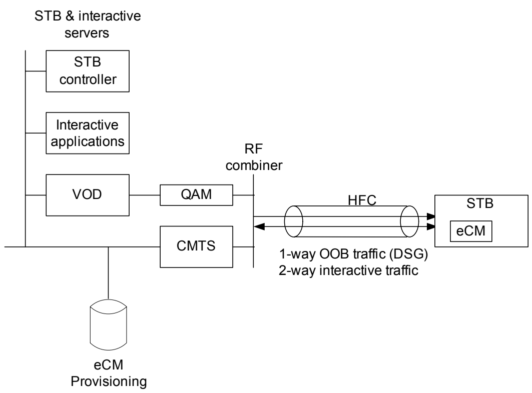


Figure 1 – IP STB system view

Note that in addition to the embedded CM (eCM) in the STB, a large portion of the diagram deals with integrating the Cable IP network with the headend STB systems.

The five areas of the paper are represented in this view. eCM provisioning is shown by the connection through the CMTS to the eCM provisioning servers. OOB transport is show by the connection from the STB controller, through the CMTS as a one-way stream of information to the STB. Interactive applications are shown in the headend via an IP connection to the STB through the CMTS. Video delivery is show through a QAM modulator and not the CMTS as this is the most cost effective means to deliver digital video at this time. There is no technical reason why digital video could not be delivered through the CMTS, however at this time the economic reasons are not there. Finally, with all the connections through the CMTS, a section is included to discuss the traffic engineering impacts of this device.

## 1. eCM PROVISIONING

### Introduction

Because the STB is being added to the IP network infrastructure, the STB will use a DOCSIS® embedded Cable Modem (eCM) for both one-way OOB traffic and interactive traffic over the HFC. The eCM does not have to be used for either Internet access or for video delivery, just for OOB and interactive traffic.

Having an eCM in a STB provides an interesting choice for operators. The STB is traditionally provisioned by the STB headend controller. DOCSIS CMs are provisioned by a set of HSI servers. These two worlds have not met before and as a result, the operator has the following choices when it comes to provisioning.

1. Both STB and eCM provision from the STB controller, which means the eCM may follow STB-supplier proprietary provisioning instead of defined HSI provisioning steps.

2. Provision the STB from the STB controller and the eCM from the HSI provisioning servers

3. Provision both the STB and the eCM provisioning from the HSI complex (which means the STB provisioning migrates to a single complex that is also provisioning DOCSIS, PacketCable and CableHome devices.

This paper will assume that option 2 is used because this allows continued use of the deployed STB controller while using standard eCM provisioning. In the longer term, since both HSI and VoIP are provisioned from the "data" side of the house, it is expected that STB provisioning will also migrate in this direction toward (option 3); however, this is a future topic.

## eDOCSIS™

There is an industry specification for Embedded DOCSIS, called eDOCSIS.

The first DOCSIS specifications were created for stand-alone cable modems that provide high-speed Internet services. The emergence of a class of devices that embeds additional functionality with a Cable Modem, such as packet-telephony, home networking and video, has necessitated the creation of this specification to define additional requirements such as interfaces, management and provisioning models. This is necessary to insure the eCM will function and interact properly with both the network and the device in which it is embedded.

The goals for eDOCSIS are described in the following list:

- To preserve functional separation of the DOCSIS cable modem entity from the device in which it is embedded. This ensures existing DOCSIS cable plant integrity, cable modem configuration, management and provisioning security are not compromised.
- To isolate DOCSIS cable modem functionality so that specification compliance can be tested for the eCM component independent of device in which it is embedded.
- To enable the service provider to enable or disable forwarding traffic between each application and the eCM within the eDOCSIS Device.
- To maximize compatibility with existing back-office management and provisioning infrastructure so that new services enabled by eDOCSIS devices can be deployed rapidly.
- To architect eDOCSIS devices in such a way as to scale to new services and applications, and to take advantage of technology innovations to achieve low cost and high functionalities.

An eCM has special requirements placed on it beyond a standalone CM (the kind used for HSI service). Specifically when the eCM is provisioned it includes information that identifies the device it is embedded in. eDOCSIS currently specifies how to embed a CM inside of two devices, a PacketCable Media Terminal Adapter (MTA) and a CableHome Residential Gateway. eDOCSIS does not specify how to embed a CM inside a STB; however, it seems reasonable to follow the blueprint already devised for these other devices.

For the purpose of provisioning, Figure 2 is one possible logical diagram of a STB with an embedded DOCSIS CM and it shows three separate elements within the STB that need to be provisioned:

1. eCM
2. eSTB
3. CableCARD™

The eCM will be provisioned by the HSI servers as a separate entity from the rest of the STB. The embedded STB (eSTB) is provisioned by the STB controller over the DSG connection (described in a later section). CableCARD provisioning depends on operator policies and is beyond the scope of this document.

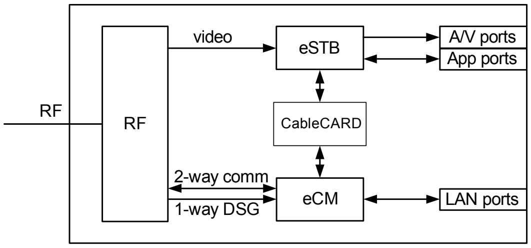


Figure 2

eDOCSIS provides guidelines for how the eCM should be provisioned, including DHCP options. In addition, eDOCSIS provides a blueprint for defining the logical connections within the STB to the CableCARD and the eSTB for consistent headend management practices.

## 2. OUT-OF-BAND TRANSPORT

### Introduction

OOB Transport enables a control-plane connection to the STB, even when there is an outage on the return path (and only the forward path is available to the STB). Traditionally this connection is from the STB controller over a QPSK modulator on the forward cable plant; a network completely separate from the IP infrastructure as shown in Figure 3.

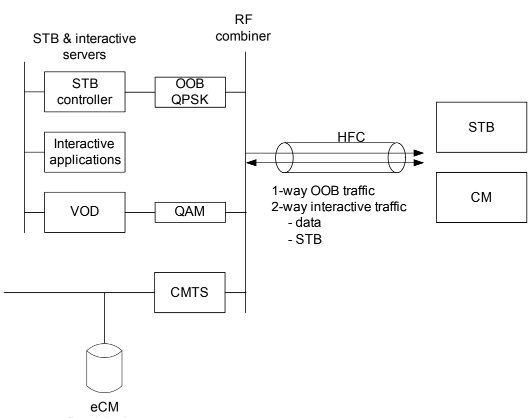


Figure 3

An industry specification, DOCSIS Set-Top Gateway (DSG), provides a method to migrate the legacy STB OOB channel to the common IP infrastructure. DSG supports retaining the legacy OOB transport for already deployed STBs but allows newly deployed boxes to operate on the infrastructure.

### DOCSIS Set-Top Gateway (DSG)

The DOCSIS Set-top Gateway function resides in the CMTS and is called the DSG Agent. The DSG Agent provides reliable and transparent transport for Out-Of-Band messaging that has traditionally been carried on dedicated legacy OOB channels, specifically those defined in SCTE 55.1 and SCTE 55.2, over a DOCSIS forward channel.

A DSG configuration is shown in Figure 4 reflecting how as part of DSG the STB controller is networked to the CMTS.

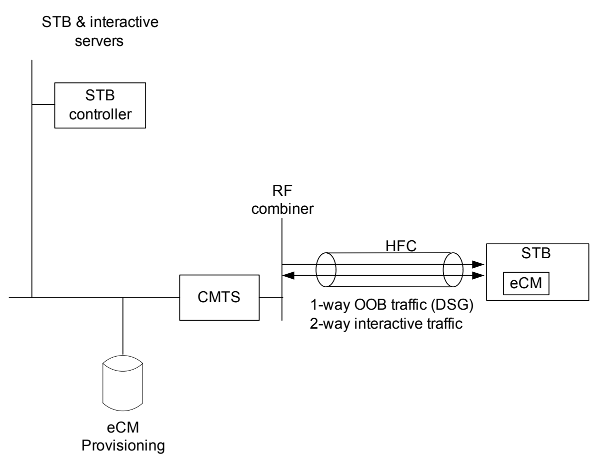


Figure 4

To support operation during return path outages, DSG uniquely modifies the eCM to continue receiving downstream information in the event of a return path outage. Receiving the OOB information through

DSG takes precedence over registering as a two-way CM.

Legacy OOB messages from the STB controller arrive at the CMTS where the DSG agent encapsulates them into Ethernet frames which are delivered over the DOCSIS forward path. DSG further describes how to address these Ethernet frame to allow an eCM in a STB to find them on the DOCSIS channel. The eCM passes the DSG information to the CableCARD or eSTB depending on how the system architecture is defined.

DSG supports a minimum of 8 tunnels per STB Conditional Access system (CAS). In this way, information can be segmented onto individual tunnels for easier management. For instance, System Information (SI) and Emergency Access System (EAS) information could be placed on one tunnel while Electronic Program Guide (EPG) information could be placed on a second tunnel and CAS information such as Entitlement Management Messages (EMMs) can be placed on a third tunnel. The operator has the flexibility to configure the DSG tunnels based on their needs and preferences.

### DSG Basic Mode

In DSG Basic Mode the destination Ethernet address of the DSG Tunnel may be either a multicast (group) MAC address or a unicast (individual) MAC address. The DSG Client in the STB recognizes a DSG Tunnel solely by the uniqueness of the DSG Tunnel Address.

The tunnel addresses are known by the STB ahead of time. The tunnel addresses are either included in the software image of the STB or they are learned from a CableCARD™ when it is inserted into the STB.

### DSG Advanced Mode DSG

DSG Advanced Mode provides several benefits over Basic Mode.

First, DSG Tunnel Addresses are learned dynamically from the network; in this case the DOCSIS downstream. This provides the operator with more flexibility when associating STB controller information to DSG Tunnels.

Second, Advanced Mode provides additional options for identifying DSG tunnels. In addition to Ethernet addresses, both CA_system_ID and Application_ID are included. These two IDs still map to Ethernet MAC addresses on the DOCSIS forward path, but using the IDs allows a more seamless integration with existing STB operations.

Third, Advanced Mode provides a keep-alive message on the DOCSIS forward path. In the event of a backend network outage when STB controller messages may stop appearing on the DOCSIS downstream, this keep-alive message provides a mechanism to keep the embedded CM on that downstream channel so when the network is restored the STB will begin receiving OOB information as quickly as possible.

Fourth, Advanced Mode provides additional security features achieved through a combination of techniques. First, the destination MAC address of the DSG Tunnel may be replaced dynamically. If the MAC address of a tunnel ever becomes widely known, it may provide the opportunity for a PC to spoof that MAC address and snoop the DSG Tunnel. DSG Advanced Mode also provides a downstream filter which will further qualify the DSG Tunnel based upon destination IP address, source IP address, and destination UDP port.

## 3. INTERACTIVE APPLICATIONS

### Introduction

Interactive applications require a method for 2-way communication between the STB and the application server on the back-office network. Since the network has migrated to IP, supporting the end-to-end connection over IP is not unreasonable. And since the method for Cable to offer IP over the HFC network is DOCSIS, having an embedded CM in the STB provides a means for communications between the STB and application server to be over common infrastructure.

Three areas relating to interactive applications are discussed:

- OCAP
- network topology and services
- security

All relate to support for IP networking and connections.

### OCAP

OCAP, the OpenCable™ Application Platform, provides Application Programming Interfaces (APIs) that support two-way IP networking. Specifically OCAP 1.0 provides support for the basic forms of IP transport (unicast and multicast) as well as obtaining an IP address using the Dynamic Host Configuration Protocol (DHCP).

While the existing APIs are sufficient for basic IP communication, additional APIs could be used for more advanced forms of IP communication.

### Network Topology

The back-office IP network is generally very specific to the operator. There are numerous protocols and connectivity choices available that can be used to meet business objectives. Important considerations include:

- address plans that incorporate STB application servers in the network with HSI and VoIP servers.
- prioritizing STB interactive application traffic relative to HSI and VoIP traffic.
- designing and sizing network connections
- routing protocols for a resilient network.

Interconnection of the Video and HSI networks should use common protocols and links designed from a systems perspective.

### Security

Figure 5 shows an additional connection not included in Figure 1.

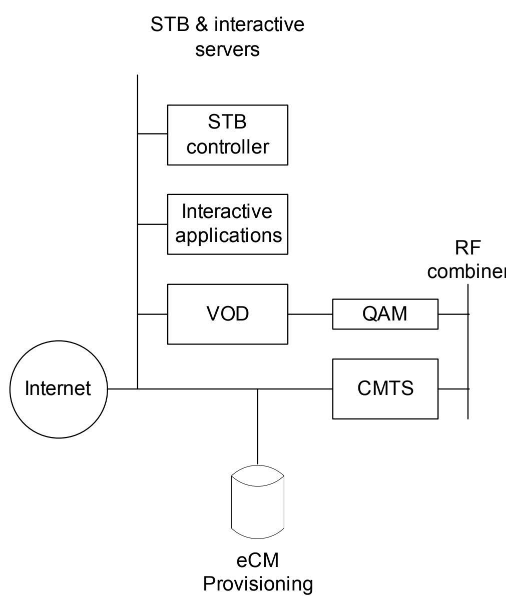


Figure 5

Specifically the CMTS connection to the Internet is shown for HSI and VoIP services, providing an indirect connection between the Internet and the STB platform.

This is not a cause for concern for a well-engineered network. The HSI and VoIP service elements also have this indirect connection and are kept safe through prudent networking practices. The operator needs to be aware of the threats and take appropriate measures to secure the STB platform.

## 4. VIDEO DELIVERY

### Introduction

Video delivery over IP is not a question of when but how far to take it. Video is delivered from the video server to the edge QAM modulator over Ethernet connections (typically Gigabit Ethernet, GigE) as IP packets today. Therefore arguably video is being delivered over IP today.

But it is also technically possible to deliver video over IP all the way to the STB by delivering the video to a CMTS. Both scenarios are show in Figure 6.

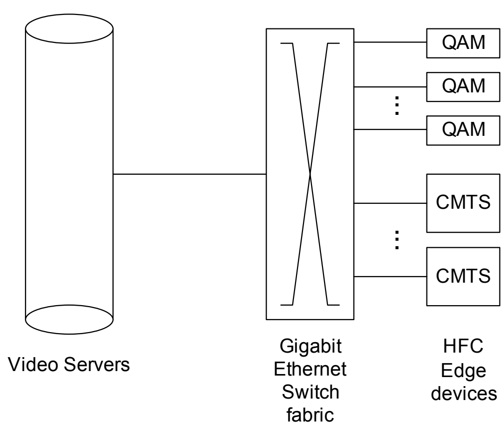


Figure 6

Video Delivery will be discussed in two segments: from the video server to the HFC edge and from the HFC edge to the STB. The HFC edge can either be a QAM modulator that delivers video over native MPEG transport or the HFC edge can be a CMTS that delivers video over DOCSIS. Technically both work, economically there are differences.

### Server To HFC Edge

Backbone and Metropolitan Area Network digital video transport (broadcast and on-demand) is already using MPEG-2 framing over User Datagram Protocol (UDP)/IP carried over Gigabit Ethernet. Both Single Program Transport Streams (SPTS) and Multi-Program Transport Streams (MPTS) are carried this way. An example of SPTS is a VoD stream at the streaming server's Gigabit Ethernet output. An example of MPTS is the multiplexed broadcast streams from a multiplexer.

Current industry practice is to encapsulate seven MPEG-2 transport frames into a single UDP/IP packet. Both constant bit rate (CBR) and variable bit rate (VBR) encodings can be supported. The IP-based video transport is directed to an edge termination that may be either a QAM modulator or a CMTS.

### HFC Edge To The STB

The debate about video over MPEG or video over IP should be about video over the most cost effective HFC transport. MPEG and IP are both digital technologies that deliver content in the form of 1's and 0's; however the current costs of the technologies are different. To determine which is the most cost-effective, the cost of MPEG QAM modulators should be compared to the cost of IP QAM modulators. An IP QAM modulator is the QAM modulator in a CMTS. Count the number of downstream ports in a CMTS (each is a QAM modulator) and divide that number into the cost of the CMTS to get a rough estimate of the cost of an IP QAM modulator. Compare this to the cost of the QAM modulator being used for VOD. While there are other considerations, this provides a rough estimate. The decision is a matter of economics, not technology.

## 5. CMTS CONSIDERATIONS

### Introduction

Two considerations will be discussed that illustrate issues when carrying multiple services over a DOCSIS consideration. The first is ensuring each class of traffic receives the proper delivery guarantees. The second is ensuring overall that the CMTS does not run out of other resources before becoming bandwidth limited.

These issues need to be solved when a second service is added to the CMTS, such as adding VoIP when already offering HSI. With DSG and possibly other traffic types on the horizon, solving these issues becomes even more important.

### Traffic Classes

The CMTS may be carrying several types of traffic including:

- DSG traffic
- STB interactive application traffic
- HSI traffic
- voice traffic
- video traffic
- on-line game traffic, etc.

All of these traffic types can be carried over IP; however, each traffic type may demand different resources from the CMTS.

Simply mixing these traffic types on the DOCSIS channel may result in the jitter or delay of all packets or in the worst case dropping some packets. DOCSIS 1.1 provides tools that can guarantee the delivery of traffic types, not all traffic, but some types of traffic. The operator will need to apply these DOCSIS Quality of Service (QoS) tools, along with proper network engineering, to meet service guarantees.

The PacketCable™ Multimedia project (PCMM) describes one method to apply and manage DOCSIS QoS tools. A discussion of PCMM is beyond the scope of this paper; however, it will provide a valuable method to manage QoS on the DOCSIS network.

### CMTS Resources

The DOCSIS protocol has certain design limits that should be respected when engineering a network. With HSI service the CMTS usually runs out of bandwidth before it runs out of other DOCSIS resources; however, as different types of services are added to the DOCSIS network, it is possible that other DOCSIS resources will come into play before available bandwidth is exhausted.

One such limit within the DOCSIS protocol is the pool of available Service Identifiers (SIDs). Each upstream connection between a CM and CMTS requires a SID. Within a DOCSIS domain, defined as the group of upstream channels associated with a downstream channel, there are approximately 8,000 available SIDs.

Field experience shows that with HSI a DOCSIS domain will run out of available bandwidth before 8,000 CMs are present; however, the addition of DSG to the network will increase the number of low bandwidth CM connections. DSG is assumed to be low bandwidth at the start because current two-way STB connections are much more concentrated than HSI connections.

With the addition of low bandwidth connections to the already existing HSI connections, it is possible the CMTS SID limit could be approached. Theoretical calculations show that even with multiple low-bandwidth connections the industry has a low probability of reaching the SID limit; however, further field experience is needed to fully understand the issues.

## SUMMARY

Moving the STB to the IP infrastructure involves more than just delivering video over IP. Nor is Internet access from the STB a prime concern. Rather, there are fundamental provisioning and traffic engineering issues to address as more services are added to the IP infrastructure.

## AUTHOR CONTACT

Doug Jones
YAS Broadband Ventures
1877 Broadway
Boulder, CO 80302
voice: +1 303.415.9400 x303
doug@yas.com