White Paper

IP Video Test in Transport Networks

By Reza Vaez-Ghaemi, PhD

Introduction

Internet Protocol (IP) Video service is delivered over a complex IP network. Impairment on this network can affect the video and/or audio component of a digital TV channel. The real-time nature of the service typically prevents the network from resending errored packets and thereby correcting the situation. As a result, the end users' perceived quality of experience (QoE) will be affected to varying degrees. Independent studies have shown that contrary to voice services, end users of IP television (IPTV) are not expected to compromise on quality, thus the signal quality across the IPTV network must be routinely tested or monitored to minimize and quickly resolve potential threats to service revenue.

This white paper focuses on the installation and maintenance test needs in Metro Core networks. Current test practices revolve around physical, link, and IP layer tests, which are insufficient for verifying and reproducing video impairments. As a result of these deficiencies, Tier 1 technicians cannot effectively resolve video service problems caused in Metro networks. This paper introduces IP video transport technology and provides an overview for testing and troubleshooting Metro Core field installation applications.

IP Video

IP Video is defined as multimedia services that can include television, video, audio, text, graphics, and data delivered over IP-based networks. These multimedia services are managed to provide the required level of quality of service (QoS) and QoE for security, interactivity, and reliability.

Figure 1 shows the protocol stacks used for both traditional digital broadcasting and IP Video. Motion Picture Experts Group 2 (MPEG-2) is a very common technology used to transport video streams in IP Video networks. Video content is carried in MPEG-2 transport streams.

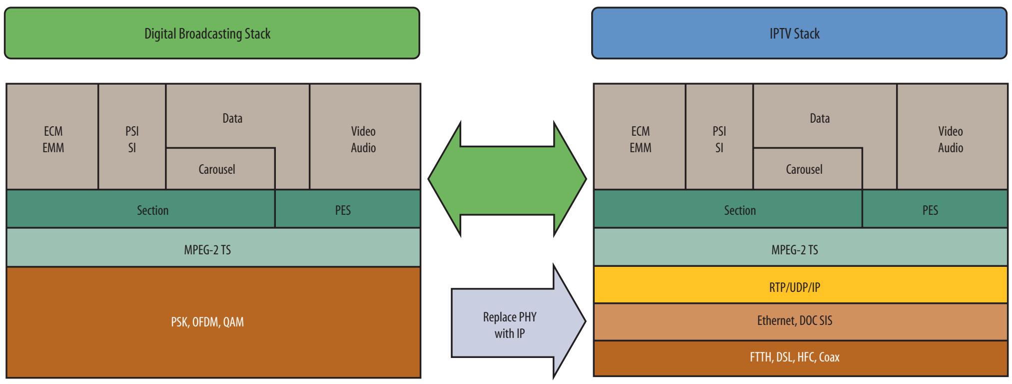

Figure 1 Comparison between Digital Broadcasting Stack and IPTV Stack

For digital broadcasting by terrestrial, satellite, and cable, the physical layer differs in the modulation schemes used. Phase shift keying (PSK), orthogonal frequency division multiplexing (OFDM), and quadrature amplitude modulation (QAM) are just three widely deployed modulation schemes. In a typical cable or satellite network, using broadcast video technology, all the contents constantly flow downstream to each customer, and the customer switches the content at the set top box (STB).

For IP Video, an Ethernet layer replaces the physical layer of the digital broadcasting. In other words, the MPEG-2 transport stream is delivered via IP multicast or unicast, depending upon whether the MPEG-2 transport stream carries live Broadcast TV or Video on Demand (VoD) contents. For a switched IP network, content remains in the network, and only the customer-selected content is sent into the customer's home, which frees up bandwidth, and the customer's selection is less restricted by the size of the "pipe" into the home.

MPEG-2 Transport Stream

Transport stream is a communications protocol for audio, video, and data as is specified in MPEG-2 Part 1, Systems (ISO/IEC standard 13818-1), which was designed to allow multiplexing of digital video and audio and to synchronize the output.

Figure 2 illustrates the basic multiplexing approach for single video and audio elementary streams. The resulting compressed elementary streams are packetized to produce Packetized Elementary Stream (PES) packets. Each PES packet is broken into fixed-sized transport packets to be carried inside an MPEG-2 transport stream.

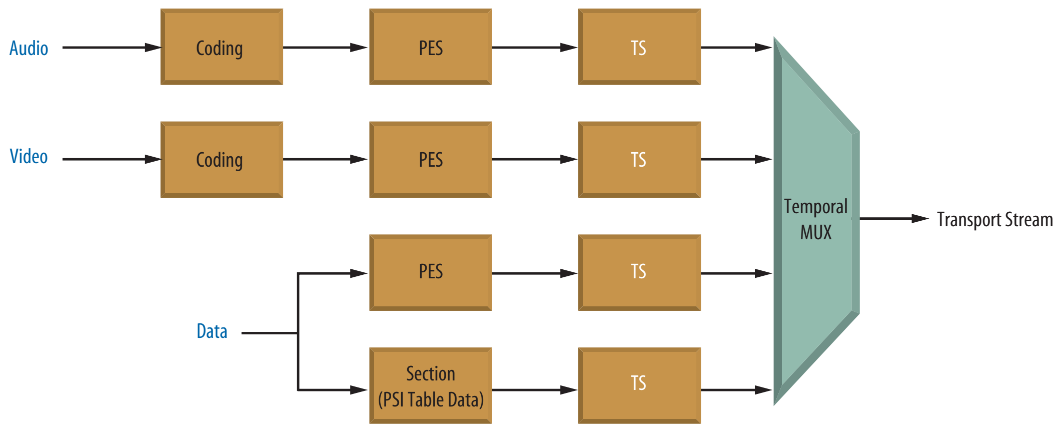

Figure 2 Encoding of IP Video in transport streams

A transport streams packet is the basic unit of data within transport streams. The transport streams packet consists of a sync byte with a value of 0x47 followed by three 1-bit flags and a 13-bit Program Identifier (PID) and then followed by a 4-bit continuity counter. Additional optional transport fields may follow, as signaled in the optional adaptation field. The rest of the packet consists of payload. Packets are most often 188 bytes in length. Figure 3 shows a detailed MPEG-2 transport stream header structure.

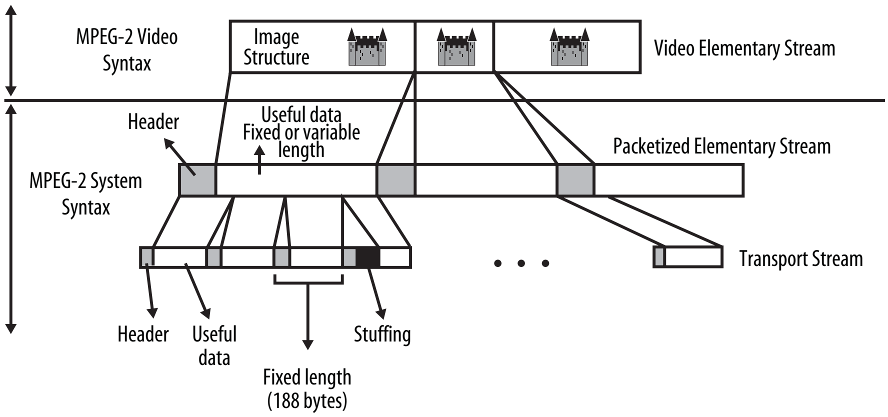

Figure 3 MPEG-2 detailed transport stream packet

Packetized Elementary Stream

Figure 4 illustrates a detailed PES header structure, which starts with a 3-byte start code, followed by a 1-byte stream ID, and a 2-byte length field. Additional optional fields may follow. The PES payload includes the elementary streams data.

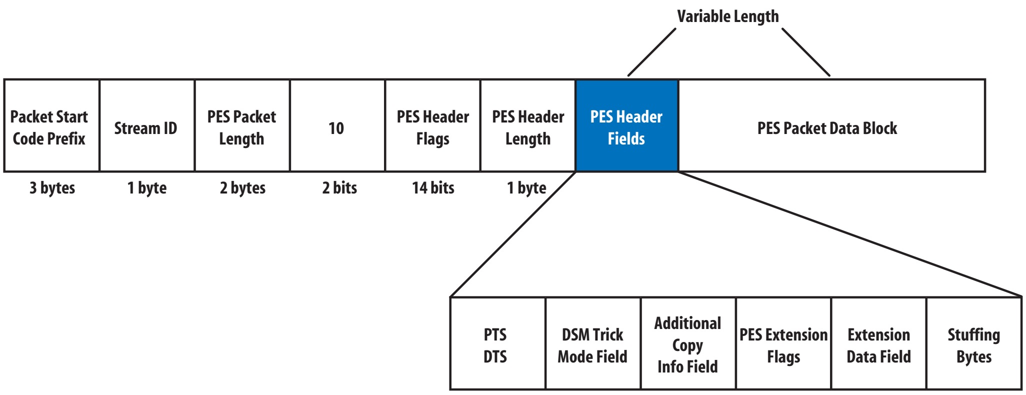

Figure 4 Detailed PES packet

To receive a particular transport stream, the user must first determine the PID being used and then filter packets with a matching PID value. Identifying which PID corresponds to which program requires the transmission of a special set of streams, known as signaling tables, with a description of each program carried within the MPEG-2 transport streams. Signaling tables are sent separately to PES and are not synchronized with the elementary streams, because they are an independent control channel.

The tables (called Program Specific Information [PSI] in MPEG-2) consist of a description of the elementary streams that must be combined to build programs and a description of the programs themselves. Each PSI table is carried in a sequence of PSI sections. Although they may vary in length, they are usually small in comparison to PES packets. Each section is protected by a CRC (checksum) which is a simple way to protect the integrity of data by detecting errors in data transmitted. Checksum works by adding up the basic components of a message and storing the resulting value. At a later time, someone can perform the same operation, compare the result to the previous value and determine if the was or wasn't corrupted. To verify the integrity of the table being carried, the length of a section allows a decoder to identify the next section in a packet. A PSI section may also be used for downloading data to a remote site. Tables are sent periodically by including them in the transmitted transport multiplex.

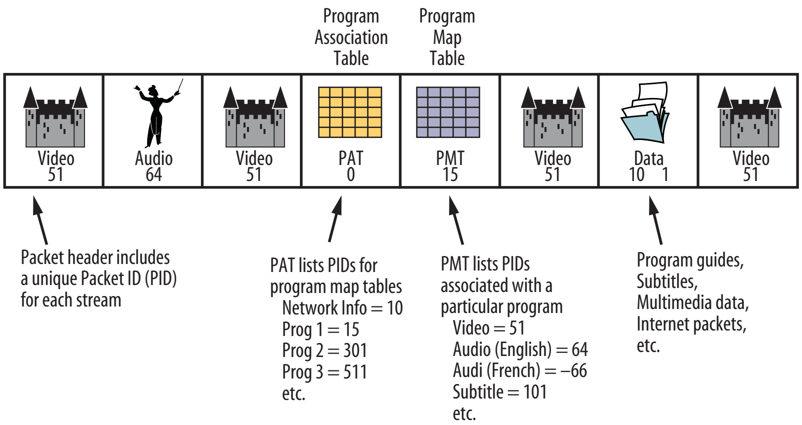

Figure 5 Signaling tables and transport layer PIDs

MPEG-2 signaling tables (Figure 5) include the following:

– Program Association Table (PAT) lists the PIDs of tables describing each program. The PAT is sent with the well-known PID value of 0x00.

– Condition Access Table (CAT) defines the type of scrambling used and PID values of transport streams that contain the conditional access management and entitlement management message (EMM). The PAT is sent with the well-known PID value of 0x01.

– Program Map Table (PMT) defines the set of PIDs associated with a program, such as audio, video, and data

The tables described above are illustrated as PSI sections in Figure 6.

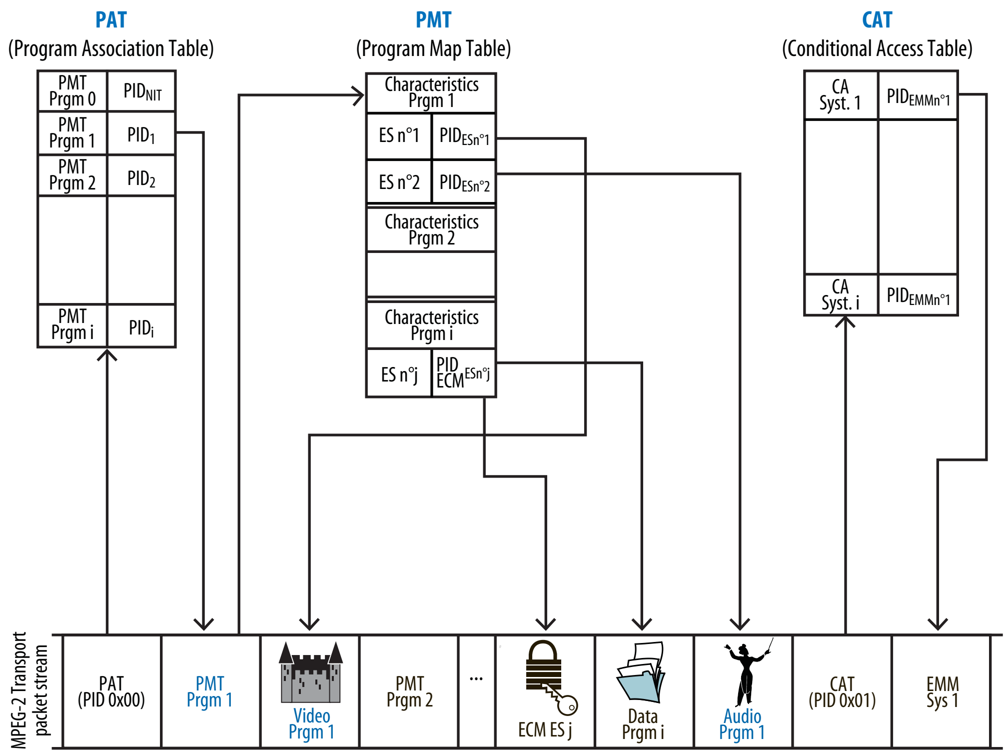

Figure 6 Mapping of PAT/PMT/CAT in MPEG-2 transport streams

IP Video Network Architecture

A typical IPTV network consists of the following components as illustrated in Figure 7.

- Super headend (SHE)
- Core network
- Video hub office (VHO)
- Metro network
- Video serving office (VSO)
- Access network
- Residential gateway (RG)
- Home network
- STB, PC, phone

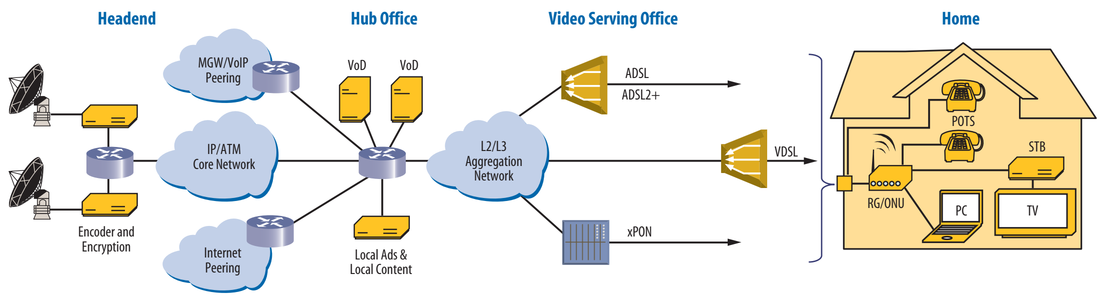

Figure 7 Generic IP Video network architecture

A typical IPTV network is composed of the following functional blocks (see Figure 7):

- National headend: Where most of the IPTV channels enter the network from national broadcasters
- Core network: Usually an IP/MPLS network transporting traffic to the access network
- Access network: Distributes the IPTV streams to the DSLAMs
- Regional headend: Where local content is added to the network
- Customer premises: Where the IPTV stream is terminated and viewed

A variety of content sources are typically found with the SHE, such as satellite and terrestrial feeds. Contents acquired at the SHE are transported over a backbone core network to a VHO. The key roles of the VHO are content localization and protection. The function of the VSO is to provide digital distribution of content to remote terminals from which homes are served via various types of physical access networks such as copper, optical fiber, or, in some cases, coaxial cable.

IP Video Test in Transport Networks

Network-specific variables that affect the QoS of a network include content quality, program clock reference (PCR) jitter, error indicator, network quality, packet loss, packet jitter, and Internet Group Management Protocol (IGMP) latency (see Figure 8).

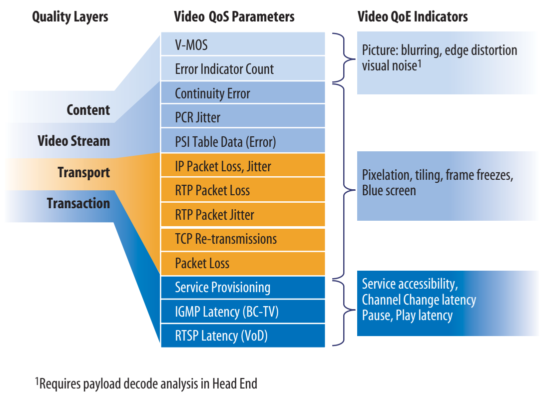

Figure 8 Video quality parts

The objective of IP video test in transport networks is to:

- Validate service flows by emulating a service end point
- Verify performance of the transport network for video streams by measuring critical transport parameters (bandwidth, packet loss, jitter, and delay)
- Isolate between physical, transport, and video elementary stream problems by evaluating and classifying measurements to various protocol/video quality layers
- Isolate between access and metro problems by simultaneous testing of multiple streams destined for different end customers/locations
- Localize transport problems by conducting and comparing tests at various locations from the SHE to the remote Digital Subscriber Line Access Multiplexer/optical line termination (DSLAM/OLT)

Video Stream Quality

The quality of the content begins at the starting point. Decisions made in the video headend, where the content is acquired, determine variations in quality. The content sources used establish the initial quality of the video stream, such as the compression algorithms implemented, the encoders employed, and the source quality monitoring system used. The data output from the encoders start the video packet flow. Two critical source quality parameters can be measured in MPEG- 2 transport stream video flows at the customer premises and/or in the last-mile access network: PCR jitter and the video transport packet error indicator count.

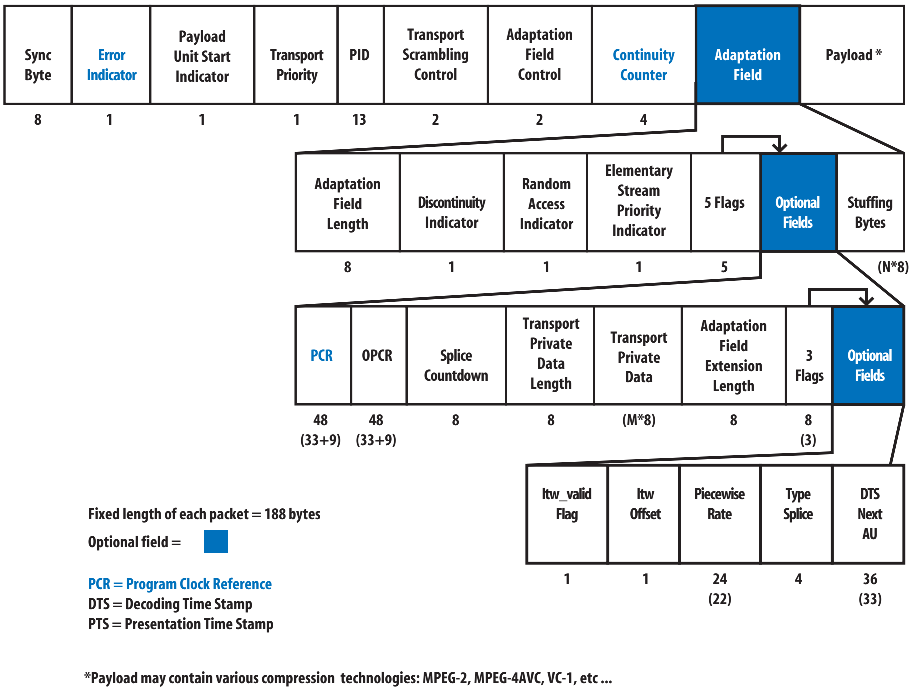

Figure 9 MPEG-2 frame format

PCR Jitter

Timing in the transport stream is based on the 27 MHz system time clock (STC) of the encoder. To ensure proper synchronization during the decoding process, the decoder's clock must be locked to the encoder's STC. In order to achieve this lock, the encoder inserts a 27 MHz time stamp, or PCR, into the transport stream for each program. Video decoders use the timing signal to synchronize to the encoded data stream so that they can derive two timing parameters embedded within each audio and video program. These timing parameters, the decode time stamp (DTS) and the presentation time stamp (PTS), are used in the decoding process to properly present the decoded video to the display unit, for example, the television.

If excessive PCR jitter is present, the decoder cannot synchronize itself correctly to the data stream, resulting in visual impairments, such as pixelization, frame freezes, and loss of color. The amount of PCR jitter that is considered excessive is not a constant but rather various parameters determine it, including the input buffer sizes of the decoder and the design of the STB software. However, in today's typical packet video networks, PCR jitter should be less than 10 ms, or preferably less than 5 ms, depending on the specific decoder/STB design.

Several factors can cause PCR jitter, however, most likely causes include:

- Overall network packet jitter
- Transcoding problems in the encoder
- Local ad insertion issues

If packet jitter is not excessive when PCR jitter is present, then the cause is specific to the particular program flow. For example, if an encoder is not performing up to specifications, PCR jitter will be constantly excessive. If PCR jitter is not constant, then a momentary problem from inserting local programming may be the cause.

Error Indicator

The encoders will set a bit called an error indicator in any transmitted video packet where they detect corrupted source content. The presence of packets with this indication is strictly an issue related to content quality and not to the performance of the distribution network. Monitoring of video encoder output streams in the headend can detect this condition and provide an early opportunity for problem resolution. Error indicator counts seen at the customer premises reveal a source quality problem.

Continuity Errors

Receipt of continuity errors indicates that MPEG-2 transport stream packets were lost, out of sequence, or duplicate packets were detected. If continuity errors are received, observe physical and transport quality layer results to see if errors are present. Not having errors present at these layers typically indicates that the continuity errors are due to congestion on the network.

Program-Specific Errors

Each transport stream is composed of audio and video packet flows. Each flow is identified with a PID and includes program specific information (carried in the PAT and PMT tables), which is required to identify and organize packet flows when the stream is received by a decoder. The PSI information must be received at regular intervals; when monitoring, transport stream the module will declare a PMT error or PAT error if the data in the corresponding table is not received within the minimum required interval.

Transport Quality

IP, Real-time Transport Protocol (RTP), User Datagram Protocol (UDP), and Transmission Control Protocol (TCP) issues occur in the IP, RTP, UDP, or TCP packets, and are usually a result of synchronization issues, congestion in the distribution network, or noise hits on the access network. Symptoms include excessive packet loss and packet jitter on the network. These symptoms are typically evaluated using packet loss, packet jitter, and Media Delivery Index (MDI) analysis.

Packet Loss

Packet loss is measured by analyzing video packet flows and determining the presence of a continuity error event. Standards define the process. Because each video packet carries a sequence number, continuity errors can be determined with certainty. For example, because an MPEG-2 transport packet is 188 bytes in length, an IP frame carries seven MPEG-2 transport packets within it; thus, losing one IP frame results in the loss of seven MPEG-2 transport packets. Conversely, if the MPEG-2 continuity counter jumps by seven MPEG-2 transport packets between two consecutive IP frames, one can be fairly certain that an IP frame has been lost.

Missing packets, out of sequence packets, and duplicate packets are all counted as errors. Each of these events can cause decoding errors. Depending on the temporal or spatial components contained in the frames within an MPEG-2 transport packet, a single packet error event may or may not be seen on the TV screen. However, actual network performance is measured by the packet loss parameters regardless of whether or not the decoder can hide the problem.

Packet Jitter

If the overall packet flow experiences excessive jitter due to congestion problems and resulting Class of Service (CoS) mechanism performance issues, packet jitter can be the cause of PCR jitter. If it is excessive enough, packet jitter can cause decode buffers to deplete, which in turn, causes gaps in the decoder output. Gaps may appear as freeze frame or pixelization events seen on the TV screen.

Due to the structure of Ethernet and IP networks, the quality of the video/audio traffic is primarily influenced by network jitter and packet loss. Due to the type of video encoding that is used in MPEG or other similar compression algorithms, the actual impact to the user perception depends on the packet type that is lost in the network. In MPEG-2, the transported packets that are used to form an image are divided into I-frames, P-frames, and B-frames. In simple terms, I-frames contain a complete image, while P-frames and B-frames contain predicted information from the other frames. Figure 10 provides a sample of the relationships between the various types of frames included in a group of pictures (GOP). As shown, I-frames are independent and provide input to support the other frames; which means that an error in the I-frames will have more repercussions to the image being viewed than losing P-frames or B-frames.

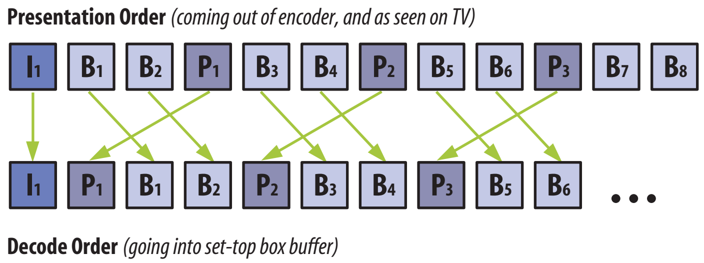

Figure 10 Group of Pictures (GOP)

Several metrics exist to quantify the impact of the network on the quality of the channel that the end user receives. The most popular parameters are MDI as well as PCR jitter for MPEG-2 transport streams. Other parameters are also used in the IPTV network, but they typically require additional packet inspection to compile the information required for deeper analysis.

IPTV is an evolving technology, and it is not completely driven by specific standards for testing and monitoring. However, the aforementioned parameters must be measured as a first alert to help qualify the user's QoE of the service delivered by the network.

Media Delivery Index

IPTV services have inherent characteristics that are the primary drivers affecting the quality of the image being viewed, such as bandwidth availability, packet loss, and jitter. The use of MDI as a testing metric can provide users with the tools to measure and diagnose network-induced impairments for IPTV streaming media. MDI is the only standards-based (RFC 4445) video-quality metric available today, which is endorsed by the IP Video Quality Alliance. MDI consists of two distinct measurements: delay factor (DF) and media loss rate (MLR). Together, they provide a QoS measurement of a delivered media stream, which can be directly correlated to ultimate QoE for the end user.

One of the key benefits of using MDI is that it does not perform any type of stream decoding to achieve its metrics; therefore, it does not require significant real-time processing power. It can also be used with encrypted media payloads. Additionally, it is not dependent on any one type of video encoding technique, so MDI can easily be scaled to monitor video quality on hundreds of simultaneous channels. MDI is typically sampled at multiple points throughout the stream path with the measurements serving as indicators of problems in the network that can be proactively addressed before it affects service. Because MDI relies on transport-layer metrics (DF and MLR), it can be used to set network margins, and it directly correlates to impending network problems with respect to video quality. Moreover, because uses packet-level metrics, it plays a key role in validating network equipment such as switches and routers since these network elements are important in determining whether a packet is delayed or dropped.

Delay Factor

The DF is the time difference between the arrival and the drain of the media packets. It takes into account the amount of jitter present in the media stream and provides the necessary buffer required for error-free transmission at the next downstream point. Very large DF values indicate severe jitter in the network, which in turn indicates that the network requires more latency (larger buffers) to compensate for the time needed to fill the buffers before the packets can be sent to the receiver. Networks experiencing high DF and insufficient buffering will eventually experience packet loss due to buffer underflow or overflow conditions further exasperating the poor video quality.

Media Loss Rate

The MLR is the number of lost or out-of-order flow packets counted over a period of time. It is important to include out-of-order packets in the MLR metric, as many stream consumer-type devices do not rearrange the order of packets received out of order. Therefore, any lost or out-of-order packets will introduce errors and visible distortions to the media stream, which may be perceptible to the end viewer. This fact makes the MLR component of MDI a popular measure for service level agreements (SLAs), as it is a much better indicator of network and video quality issues than a simple mean opinion score (MOS).

Transaction Quality

IGMP is the signaling protocol used to access broadcast video services that use a multicast network design to efficiently manage network bandwidth. IGMP lets each STB obtain only the programming that the viewer is interested in watching, conserving bandwidth in the access network. In this implementation, a join message is sent from the STB to the network. The join message asks the network to send the requested program/channel to the STB by joining a multicast group carrying the desired broadcast channel.

IGMP Latency

IGMP latency is the period between the time the join message is sent and the time the first video packet is received by the STB. Thus, IGMP latency is a measure of both service provisioning and network response performance. These messages travel upstream into the network to the first device that can add (join) the requestor to an existing broadcast channel flow. This parameter measures network performance, but not the end user's experience, with regard to channel changing time. The IGMP latency plus the time it takes to fill the decode buffer and to decode and display the content is the total user experience time. However, the buffer fill time and the decode time are functions of the network architecture and are not variables. Thus, the measurement of the variable network performance aspect of IGMP latency is the critical parameter for measuring actual network performance.

Summary

Current test practices in IP video transport networks focus on physical, link, and IP layer tests, which are insufficient for verifying and reproducing video impairments. These deficiencies result in the inability for Tier 1 technicians to effectively resolve video service problems caused in Metro networks. Technicians can now conduct measurements on TCP, UDP, RTP, and MPEG Transport Stream layers to verify the transport problems that can impact video quality. When errors are found, the proposed measurements help identify the source of the problem.

Test & Measurement Regional Sales

NORTH AMERICA	LATIN AMERICA	ASIA PACIFIC	EMEA	www.jdsu.com/test
TEL: 1 866 228 3762	TEL: +55 11 5503 3800	TEL: +852 2892 0990	TEL: +49 7121 86 2222	
FAX: +1 301 353 9216	FAX: +55 11 5505 1598	FAX: +852 2892 0770	FAX: +49 7121 86 1222	