



A Guide to Ensuring Digital Video Service Quality

Primer

Table of Contents

Introduction	4	Using QoE Scoring Methods to Develop a Service Benchmark	23
Section 1: Digital Video Content Readiness	5	Detecting Artifacts and Impairments before Subscribers.....	23
Containers, Mezzanines, and Codecs.....	5	Perceptual Video Quality: Compression and Encoding Artifacts	24
When Data Becomes Metadata.....	6	Regulatory Compliance	26
The Challenge of Quality Control.....	7	Closed Captioning.....	26
QC Methodology: Human vs. Automated Monitoring.....	7	Audio Loudness	27
Types of Errors: What Could Possibly Go Wrong?	7	Ad Verification.....	28
QC Tests, Checks, and Results	8	RF Monitoring.....	30
Examples of Content Readiness Testing in Cable Workflows.....	9	New Requirements for Monitoring OTT & Adaptive Bit Rate.....	31
Video on Demand.....	9	Summary.....	34
The QC “Pre-Flight” Check.....	10	Section 3: Network Troubleshooting & Diagnostics ...	35
Ad Insertion.....	11	Troubleshooting Background	35
Ad QC.....	12	Troubleshooting QoS.....	36
Adaptive Bit Rate.....	13	Troubleshooting QoE.....	40
“Segmenting” and QC Issues	14	How are QoS and QoE Interrelated?	42
Automating File-Based QC	15	Using QoS and QoE to Identify “Real” Problems: Troubleshooting Case Studies	43
Semi-Automated QC Workflows	15	Video QoE.....	43
Asset Management and the Fully Automated Workflow	16	Audio QoE.....	44
QC Reporting	17	Troubleshooting Over-Compression.....	44
How is a Workflow Like a Web Browser?.....	17	Troubleshooting Interoperability	44
When Systems Communicate	18	Troubleshooting Transport Streams with CaptureVu.....	45
Summary.....	18	Troubleshooting Ad Insertion	46
Section 2: Live Network Monitoring of Digital Video Services	19	Summary.....	47
Live Network Monitoring Background.....	19	References.....	47
Quality of Service (QoS) vs. Quality of Experience (QoE).....	20		
QoS Monitoring Tradeoffs.....	22		
Key Video Quality Monitoring Locations.....	22		

Introduction

Digital television has transformed the television viewing experience and has offered Video Service Providers the ability to deliver hundreds of channels of programming with better picture and sound quality. Since digital television was first launched, there have been a great deal of new technologies and capabilities to emerge including HDTV, digital ad insertion, tru2way, MHP, VOD and most recently—multiscreen or OTT (Over The Top) video. As networks converged and service offerings increased, Video Service Providers have dealt with the ongoing, arduous task of ensuring the quality of these services for their subscribers.

The first part of the digital video delivery workflow, **Content Readiness**, focuses on having all your On Demand content available in a form and condition suitable for distribution. It means content that's not only free from objectionable distortion but also equipped with the right markers and metadata and resolution to deliver the best possible service to subscribers, regardless of the receiving platform. This level of Content Readiness calls for a rigorous quality control (QC) program that inspects every individual file—from a 15-second ad to an HD movie—as it enters the cable plant and when it is stored for eventual play out. The sheer volume of content dictates some level of automation in the QC process.

The second part, **Live Network Monitoring**, involves ensuring both the Quality of Service (QoS) and the Quality of Experience (QoE) in real time across hundreds of digital services. Understanding the difference between QoS and QoE as well as how Perceptual Video Quality relates to QoE is crucial for any monitoring deployment. Knowing what, where and how to monitor leads to reduced trouble calls, a faster ability to detect and repair issues, reduced churn and a reduction in operational expenditures.

The final part of the workflow is **Network Troubleshooting**, which is important to help quickly identify and log the audio or video problem that occurred, then identify or pinpoint the equipment (or network link) that needs attention. To identify and isolate problems, it is critical to have access or test points throughout the facility. The minimum set of test points in any network should be at the point of ingest where the signal comes into the facility, the ASI or IP switch, and finally egress where the signal leaves as IP or RF. To begin testing a signal that may contain the suspected issue, both QoS and QoE methods are used. Both methods are useful in troubleshooting and analysis, but each of the two methods quantify issues using completely different metrics.

In this Guide to Ensuring Digital Video Service Quality, we will examine three parts of the digital video delivery workflow, and the tools and methods to ensure the QoS and QoE of your digital video services.

Section 1: Digital Video Content Readiness

Containers, Mezzanines, and Codecs

The file-based workflow begins with an ingest server whose job is to process received content into files that facilitate management, storage, and distribution. Content arrives in diverse formats and it is necessary to process everything into a uniform format for internal use.

The output of this process is digital files organized into “containers,” also known as wrappers. A container usually embodies a single file, though some containers may be made up of multiple files linked together. In any case containers encompass all the vital information about the files they contain: compressed video, audio and also importantly, metadata.

There is no single universal container format. In fact there is a variety of container formats including:

- MPEG Program Stream
- MPEG-2 Transport Stream
- MP4
- 3GP
- QuickTime File Format
- Material Exchange Format (MXF)
- General Exchange Format (GXF)
- Advanced Systems Format (ASF)

Codec	Container	Bit Rate
DV	MXF Op-Atom	100 Mb/s
MPEG-2 IMX	MXF OP1a	30–50 Mb/s
Apple ProRes	QuickTime	220 Mb/s
VC-3	MXF or QuickTime	220 Mb/s

Table 1.1. Commonly-used mezzanine file formats and codecs.

...and more. Some formats are optimized for playout, others for editing or capture or other points in the workflow. The common MXF container format alone has multiple variants. Containers are a cornerstone of file-based workflows.

In the cable domain, MPEG Transport Stream files are of course ideal for playout, ready to be packetized and sent out over the network. The .TS container format can be used for offline processing and QC operations as well. In addition the Quicktime and MXF formats, among others, are well suited to editing and transcoding (converting) by means of codecs such as those in Table 1.1.

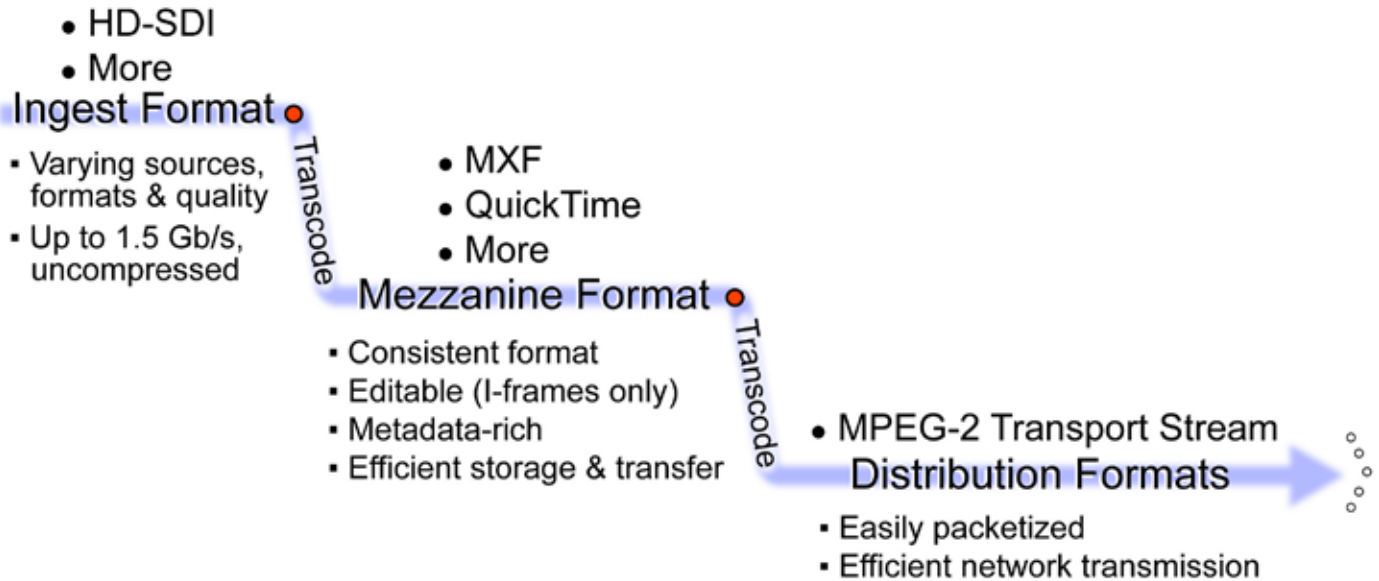


Figure 1.1. Mezzanine formats allow efficient handling and storage of files within the cable plant.

MXF and QuickTime formats are commonly used as “mezzanine files.” The term comes from the world of architecture, where it denotes an intermediate floor (often a lobby or a balcony) between two main floors. Similarly, mezzanine files are an intermediate step. They are working copies that are more expedient to use in the workflow than the original source material. Figure 1.1 is a symbolic illustration showing where the mezzanine’s transitional format fits in—midway between the high bit rates of the raw ingest and the lower distribution bit rates.

Though substantially compressed, the mezzanines suffer little noticeable loss in picture quality. They have sufficient resolution to minimize generation loss when transcoding. And being smaller files than the original source files (thanks to data rates of about 200 Mbits/sec compared to HD-SDI rates of 1485 Mbits/sec), they require less storage space and transfer time. Importantly, mezzanines are made up of I-frames only, which means that successive frames are complete and editable with no need for interpolation of I, P, and B-sequences, making editing and post-production tasks more efficient. For more details on the IPB hierarchy see the “Adaptive Bit Rate” section in this document.

When Data Becomes Metadata

Metadata is often stored in the file with the video and audio data, or it can be located in an auxiliary file in the same package. Metadata is overarching “data about the data”—a set of descriptors that can include the episode title, scene numbers, languages, ratings, and more. There may be information about the usage rights attached to the file. This specifies the number of layouts or the length of time in the licensing for the content. All these values are human-readable. In addition metadata may express attributes like frame size, frame rate, or aspect ratio; information that’s essential for correct playback. These entities are machine-readable.

Workflow operation is more efficient when metadata is written into the files. The MXF file format is being widely accepted in the broadcast industry specifically because it is metadata-rich. MXF is equally suitable for some cable applications.

The Challenge of Quality Control

QC Methodology: Human vs. Automated Monitoring

Almost like a law of physics, the challenge of quality control seems to grow with the square of content quantity. Perhaps there was a time when visual inspection of incoming programs was sufficient, but that era ended with the explosion of content required for services like Video On Demand (VOD). Human monitors are prone to overlooking subtle impairments, and of course can't detect metadata or embedded digital errors. Moreover, the task of scrutinizing thousands of old movies at two hours per selection can add up to decades' worth of man-hours. And with all the transcoded versions for diverse output platforms, the job just gets that much bigger.

A better idea is an automated QC strategy that operates 24/7 and pinpoints errors so human operators can spend their time fixing rather than finding errors. Automated QC is more thorough and consistent, whether it's evaluating one short commercial or an entire archive of old TV westerns. Equally important, capable automated QC systems can detect errors that are invisible to the human eye: metadata that doesn't match the measured file attributes, for instance, or syntax errors that might pass through the industrial-strength systems in the cable plant but could crash an ordinary set-top box.

Types of Errors: What Could Possibly Go Wrong?

File-based video is a technical blessing but like all other media, it is susceptible to flawed source material. Some errors originate in the baseband product. Such flaws range from improper camera or microphone levels to amateurish, out-of-gamut "homemade" graphics in a tire commercial. These errors are part of the file and must be rectified eventually.

Incorrectly encoded files are another problem area. These can arise from a faulty encoder that produces syntax errors due to buffer overflows or similar technical issues. Or the encoder may be misconfigured, as when a standard-definition profile is applied to an HD source.




Figure 1.2. Ingesting source material into a file-based workflow.

Transfer errors are yet another hazard. Problems can attach themselves during any transfer from system to system, including the ingest transfer itself. Even content that is sent correctly can be received with errors due to interruptions or faulty equipment.

An "instrumented decoder" (Figure 1.2) is the preferred tool to ensure detection of all these error types. Whereas a conventional decoder attempts to recover gracefully from errors, an instrumented unit reports them. In addition the latter tool finds errors at two different levels: in the encoded bit stream, and in the decoded baseband image raster and audio samples. A file can be syntactically correct but still contain block frames, muted audio, and so forth. Only an instrumented decoder can find these faults reliably.



Figure 1.3. “Super white” gamut violation.

QC Tests, Checks, and Results

Quality checks fall into three major categories:

■ Structural Checks

These QC checks focus on compliance with applicable industry standards and just as importantly, acceptance criteria for entities such as iTunes, Netflix and PBS. In addition CableLabs specifications for VOD content may apply.

“Compliance” typically encompasses format-related issues: codec and container type, the MPEG profile and level, GOP structure, bit rates, frame size, and frame rates. For example, there may be an incorrect number of streams, implying a missing audio portion. Or there may be a mismatch between the signalled bit rate and the actual bit rate.

■ Baseband Quality Checks

Errors can occur in the decoded image, the image raster or the decoded audio samples. These are the target of baseband checks. These measurements, when performed in the file-based domain, are very similar to observing a real-time live signal with a waveform monitor.

But file-based checks run faster than real playback time. The QC tool looks for video errors such as dropouts, frozen frames and unexpected letter- or pillar-boxes. Gamut problems such as super black or super white (Figure 1.3) also are revealed. In the audio domain, loudness violations such as clipping and CALM Act non-compliance can be not only detected (using the same ITU-BS.1770 algorithm used in a real-time loudness monitors) but also corrected.



Figure 1.4. Blockiness due to over-compression.

■ Encoded Content Checks

These tests watch for the low bit rates and over-compression that cause blockiness artifacts (Figure 1.4) that can be measured in the decoded image raster and reported. Field order problems in interlaced video are common as well, especially in Ad Insertion. These errors show up as motion artifacts. Similarly, MPEG errors such as incorrect slice order can cause large block distortions.

Files can get damaged during transfers. An interruption, for example, may cause a file to get truncated even though there are normally safeguards to ensure recovery from interruptions. An incomplete file like this would lack an End of Sequence marker.

All of these problems can be quickly exposed with a thorough syntax check and all lend themselves to an automated QC process that includes screening and reporting on every file.

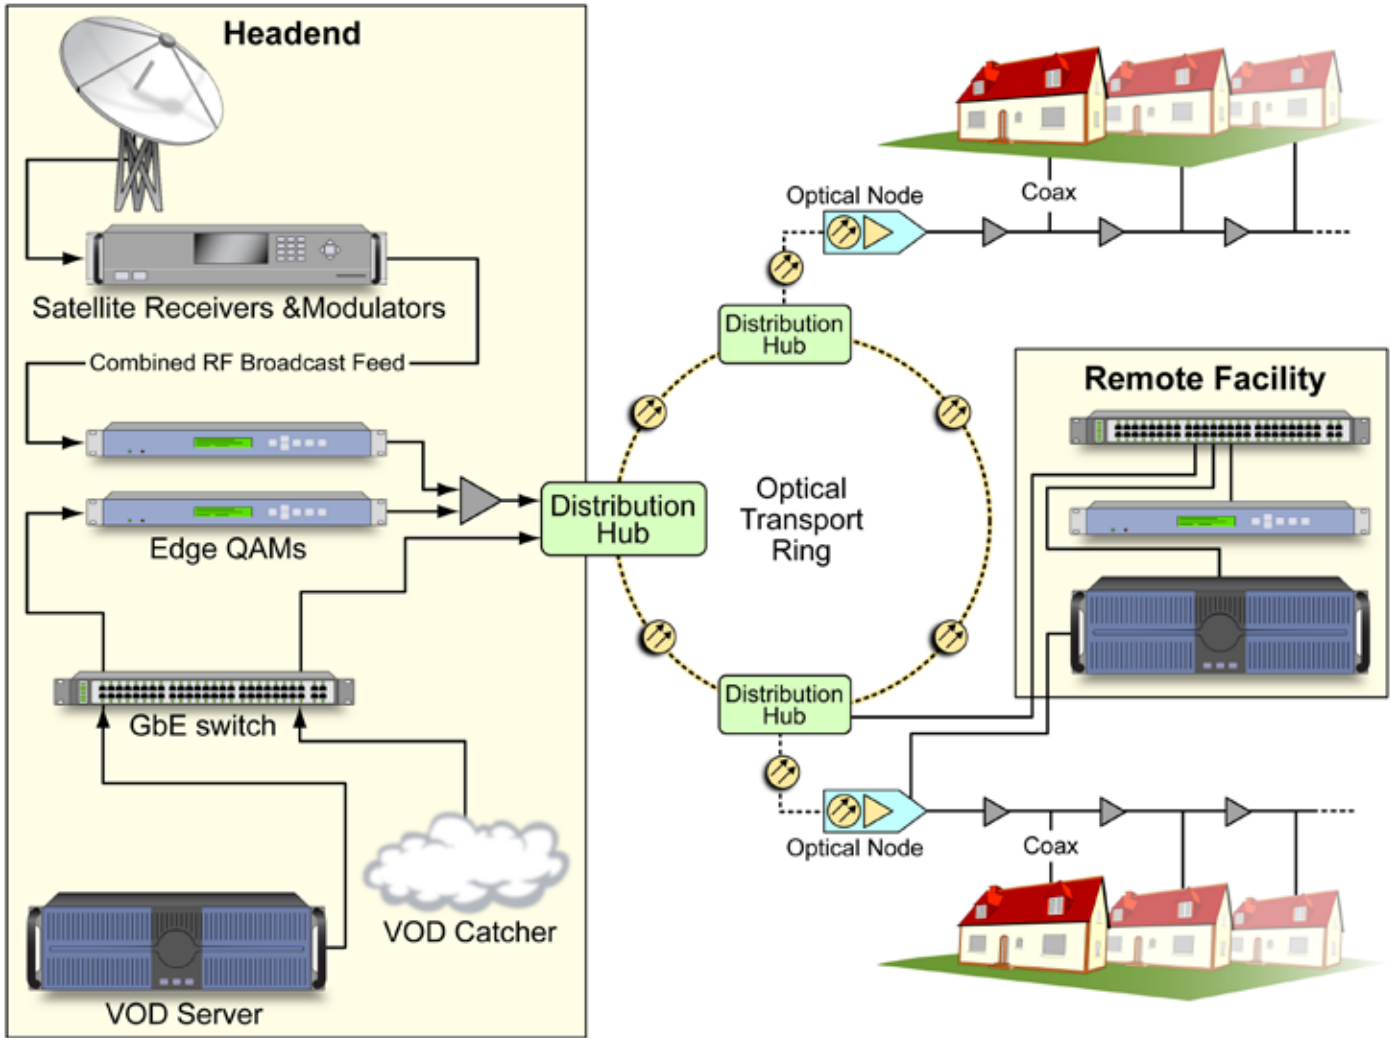


Figure 1.5. The Video On Demand (VOD) workflow.

Examples of Content Readiness Testing in Cable Workflows

Video on Demand

Video On Demand (VOD) is now widespread in the consumer market thanks to aggressive campaigns by content producers. VOD has become a substantial revenue source for cable providers, and continues to increase its market share at the expense of traditional consumer media. Figure 1.5 is an overview of the VOD flow from content capture to end user.

VOD implementation begins at the Headend with the acquisition of content. This may be via satellite receivers or by file transfers. Industry insiders use the baseball metaphor of “pitching” and “catching” to describe this process. A satellite antenna may act as the VOD catcher, though increasingly the task is simply one of transferring files from the Internet

cloud. Large (and successful) enterprises have been built up to provide tools that ensure the fastest possible transfer of massive data files such as multi-gigabyte movies. Increasingly, both the content provider and the cable operator maintain a cloud connection through a proprietary file transfer vendor to speed the exchange from pitcher to catcher. This method is likely to supplant the older FTP solution eventually, but conventional FTP is also in common use at this time.

Watermarking is added by the VOD Catcher. The purpose of watermarking is to build in an “invisible” undetectable means of tracing illegal copies. Conceivably a subscriber with the right tools could record a VOD movie and resell it, but a watermark identifies the material’s source and provides recourse against this kind of piracy. Metadata, added at ingest and stored in an XML format compliant with the [CableLabs Asset Distribution Interface specification](#), plays an important role in automating VOD playback.

The QC “Pre-Flight” Check

Subscribers pay for the VOD product, often included as part of their subscription package, and they expect a level of quality that lives up to the vendor’s claims. Therefore the quality of the content stored in the VOD servers is crucial to the success of the whole process. The moment when the requested content exits the facility is clearly not the right time to evaluate its quality. Like a pre-flight check, it is important to confirm the quality of the product before it leaves the plant.

In Figure 1.5 the headend contains the broadcast feed which is modulated and inserted into the optical transport ring and onto the network. Of course this is also the distribution point for VOD. It is not necessary to run everything from a single centralized headend and in Figure 1.5 a remote facility (which may be one of several) shares the load. These remote sites have their own VOD servers which store separate copies of the VOD assets, increasing the efficiency of distribution in their locales. A high-speed network is required for streaming video to remote edge QAMs.

Incoming VOD assets can undergo QC checks when they are received at the Ingest cache, or while they are stored on the VOD servers. What kind of problems do we look for in the VOD workflow?

- File integrity problems: These can occur during the automated pitching-catching process at ingest. Files can become corrupted or truncated. Does the syntax check detect an EOS (End of Sequence) flag, and is the measured play time accurate?
- Format compliance issues: Files must be checked for compliance with CableLabs specification CEP 3.0 (Content Encoding Profile). Are the bit rates within the prescribed range; are the PID numbers correct (e.g. 481 for video or 482 for audio); are the GOP lengths correct?
- Does the VOD material comply with governmental regulations? Are the mandated CEA 608/708 captions present? Is the audio loudness within the limits set by the CALM Act? In every case these characteristics must be not only checked but also corrected when errors are found.
- Does the machine-readable metadata agree with measured values such as frame rate, displayed picture size, and play time?

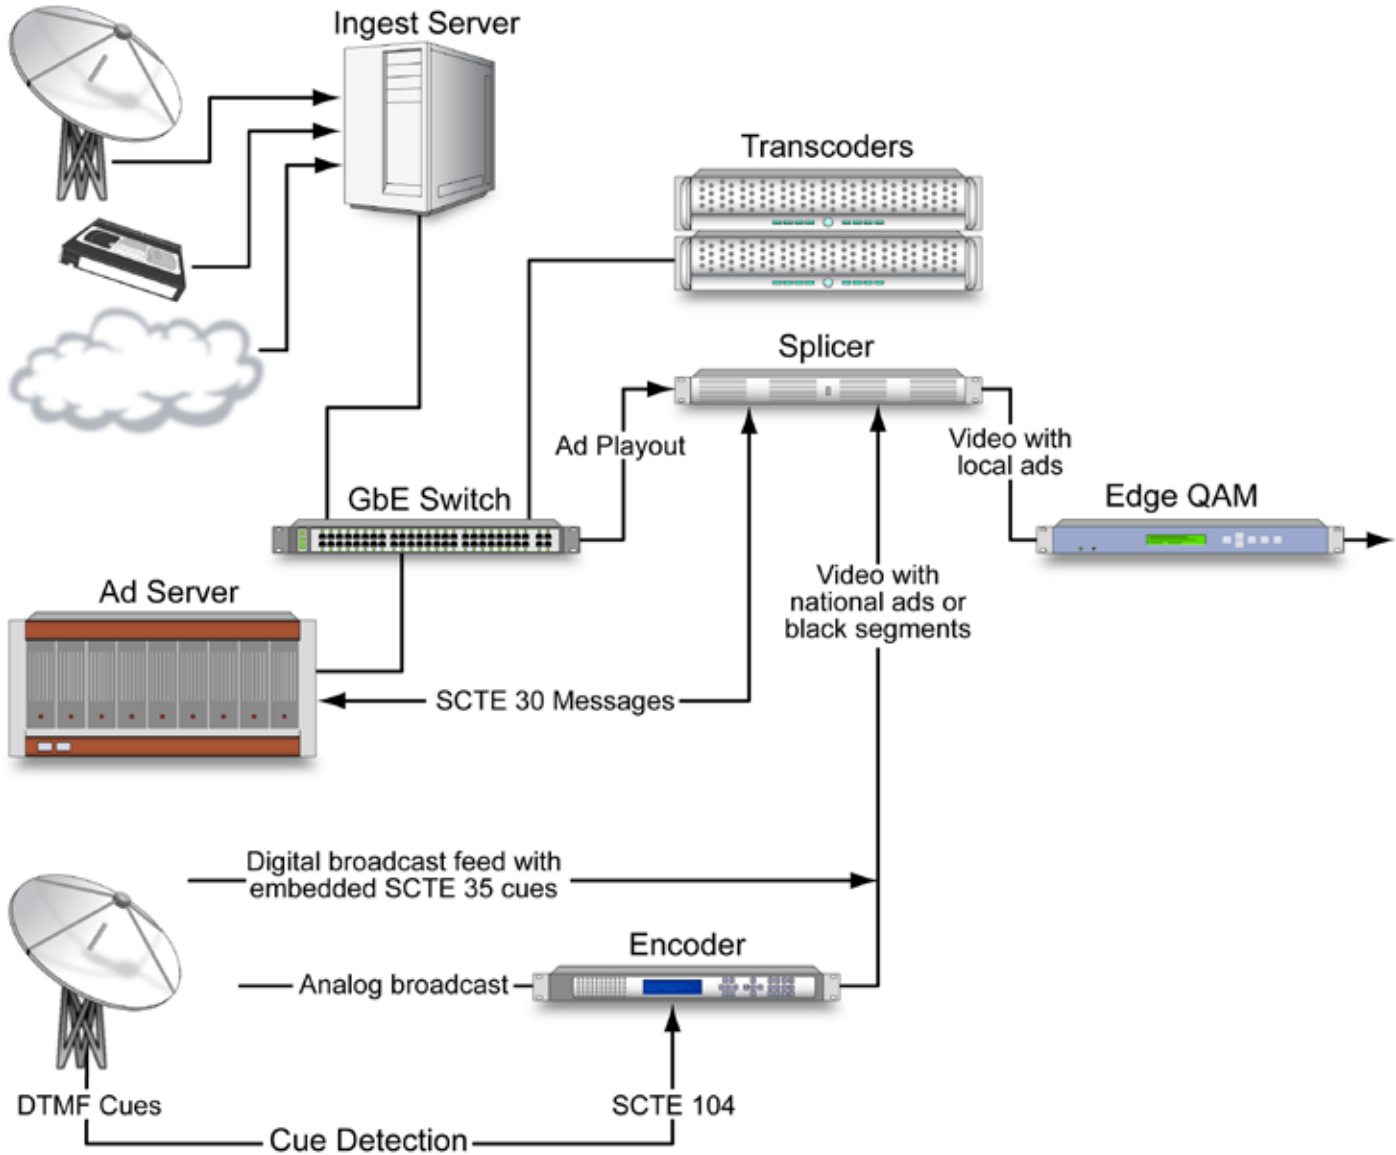


Figure 1.6. Ad insertion workflow.

Ad Insertion

Advertising is of course the lifeblood of scheduled cable and broadcast networks. A content stream isn't complete until advertising is embedded. In the ad insertion workflow (Figure 1.6), the challenge is one of managing a myriad of codec, container, and media formats. Local commercials may

be delivered in the form of tapes, DVDs, memory cards, or digital files. National ads typically arrive as files from a content delivery network. The ingest process must ensure that all of this is stored in a consistent file format. In every case, vigilant QC is required to ensure continuity of visible quality as programming switches from entertainment to ad content and back.

Rigorous incoming inspection is the ideal, and theoretically content that fails to meet acceptance standards will be rejected. But the reality is that cable operators often must groom local ads before they are suitable for distribution. The mattress outlet at the mall and the car dealership in town sometimes produce ads on a scant budget with too little attention to normal video standards such as gamut compliance and loudness.

In the early part of the ad insertion workflow, the QC focus is on content quality rather than absolute format correctness. Ads that don't meet minimum standards, for example having a very low bit rate and consequently poor picture quality, are candidates for rejection. Gross errors in play time (for example a 32-second spot for a 30-second timeslot) are another disqualifier. Other flaws such as loudness violations may be accepted with the understanding that the operator will correct them—sometimes with additional charges.

SCTE standards specify various types of ad cues and messaging encoded in diverse ways. Incoming content may include national spots that can be used “as is” or replaced by local spots. There may be black segments meant to be populated with local ads drawn from an in-plant ad server.

As ingest the content gets transcoded to the format the splicer needs in order to seamlessly insert it into the broadcast feed. Because the splicer can only switch ads in and out on GOP (Group Of Pictures) boundaries, one of the quality checks must ensure that the clip contains an integral number of GOPs; in other words it is “closed” so that there are no references to GOPs preceding or following the clip.

Ultimately the ad server in Figure 1.6 contains only ads that are in the correct format and ready for insertion. The active broadcast feed passes through the splicer, which inserts local ads at the correct time.

Ad QC

The majority of errors in the Ad Insertion workflow are quality-related. A local merchant simply can't afford the lavish production values of a big national campaign, so compromises are made. A small video house might not have a loudness meter, for example, which means that audio levels may be too high—even to the point of clipping. Or perhaps an inexperienced hand produced graphics with garish, attention-grabbing colors that just happen to be out of gamut. Or there may be over-compression, causing blockiness and artifacts. Clearly there are many pitfalls.

It is the cable operator's prerogative to reject this content or accept it and repair it. In either case it's crucial to have a QC regime that can detect these flaws.

Another job for the QC system is to confirm that the ad content meets submission guidelines. Many operators constrain the delivery formats they will accept, with the intent of reducing complexity. Limitations may include codec type, container type, audio channel assignments, picture size, bit rate, frame rate, and more.

One of the most bothersome errors in ad content is the format mismatch. For example, ads may be submitted in 4:3 standard definition (again the result of cost-cutting production) even though the broadcast will be in HD. Unfortunately this causes letterboxing as the subscriber's equipment tries to make the best of a 4:3 aspect ratio on a 16:9 screen. It is even possible to have letter-boxing and pillar-boxing occurring at the same time; a small picture appears in the center of the screen, surrounded by black on all sides. This is a distracting effect! One way to avoid it is to set up separate ingest paths for SD and HD deliveries. In addition, the QC tool should routinely verify that the video fills the active image and that letter-boxing/pillar-boxing effects are not a permanent part of the file. And it is becoming common practice to check the Active Format Descriptor (AFD) to confirm that the playout code is correct.

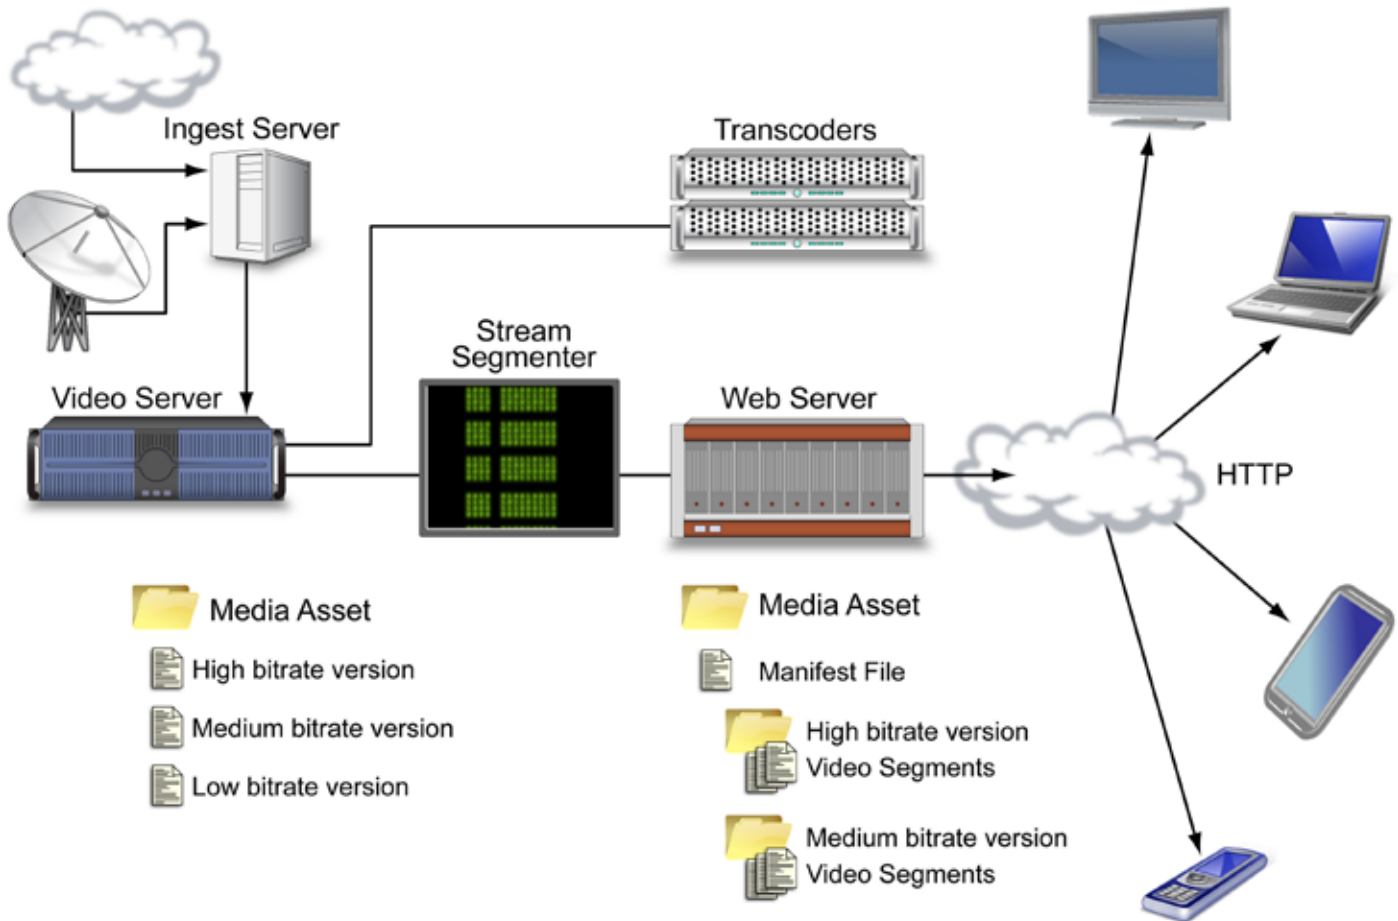


Figure 1.7. Adaptive Bit Rate (ABR) workflow.

Adaptive Bit Rate

Cable content distribution used to be straightforward, with scheduled programming going out to subscribers via operator-provided set-top boxes. But that market has changed forever. The set-tops are still there, of course, but so are tablets, phablets, phones, laptops and PCs. To be competitive, operators must offer streaming services for all these devices.

The solution for this multi-faceted challenge is Adaptive Bit Rate (ABR) streaming, also known as Over-the-Top (OTT) delivery. HTTP Live Streaming (HLS) and Smooth Streaming are two of the leading streaming architectures in use today. Both HLS and Smooth Streaming rely on transcoding each asset to multiple bit rates. Thus there are several coexisting versions of each asset. Standard HTTP network transport protocols are used for client (subscriber) access.

The multiple bit rates make it possible to optimize transmissions for moment-by-moment network capacity and to tailor the content to the receiving device. An ABR

implementation must be able to repeatedly change the bit rate, switching it higher or lower depending on network conditions. And on the client (receiving) side, a computer with a 21-inch LCD screen requires much more data to support a credible picture than does a phone with a four-inch screen. The client-side player determines the available bit rate and requests the best content available to match that rate.

ABR demands are complex, and optimizing the bit rates is only part of the story. A file must be broken into short segments rather than being transmitted in one full-length delivery to the subscriber. As shown in Figure 1.7 the “stream segmenter” tool is dedicated to this step. The segments are just a few seconds long, usually ten seconds or so. The reason for this layout is to provide boundary points at which the bit rate can be switched. As bandwidth availability changes, the client player requests the optimum bit rate and the switch occurs at an appropriate Group-of-Pictures (GOP) boundary.




Figure 1.8. H.264 compressed video content with Groups of Pictures (GOP) made up of I, P, and B frames. An Instantaneous Decoder Refresh (IDR) frame follows each GOP boundary. Frames in segment $n+1$ cannot reference previous segments.

Figure 1.8 illustrates the frame types and GOP boundary crossings that make up a segmented file. H.264 compressed video provides the basis for this simplified view. The content includes I, P, and B frames. The I frame is the ongoing reference frame and the P and B frames contain incremental changes relative to the I frame. This scheme enables very high compression ratios but it puts constraints on when and where bit rates can be switched.

The H.264 video file in Figure 1.8 also includes an Instantaneous Decoder Refresh (IDR) frame. In effect this clears the buffers that refer to preceding frames and restarts the I, P, and B sequence to define a new segment. The IDR marks the boundary between any two segments. A frame within the $n+1$ segment in Figure 1.8 cannot reference a frame from the preceding segment.

“Segmenting” and QC Issues

A key part of the QC regime in an ABR workflow is to verify the content’s readiness for segmenting. It is important to establish early in the workflow that the file can be correctly divided into usable segments. Note that this is not a check on the individual pieces; that is the job of a different tool designed to monitor and measure the ABR stream in real time. At QC time, the task is one of ensuring that smooth segmentation is possible.

In this context “content readiness” implies that IDR frames are embedded at regular timing intervals. These form the segment boundaries. The QC checks should confirm that stored content—whether movies or commercials—has the IDRs in place and timed correctly.

Another quality check relates to picture quality in an environment where content is stored in numerous versions, each with a separate bit rate. Are the lowest bit rates still delivering acceptable picture quality? For that matter, are *all* of the rates providing the expected image quality? This is a test that is not practical to perform on every piece of content in real-time; instead it is best used to guide the design of an effective ABR workflow in the cable plant. Using a picture quality analysis system, it is possible to fine-tune encoder performance to get the most out of each bit rate. In H.264 compressed video, for example, there are numerous settings that can be adjusted to maximize the final picture quality on the receiving device.

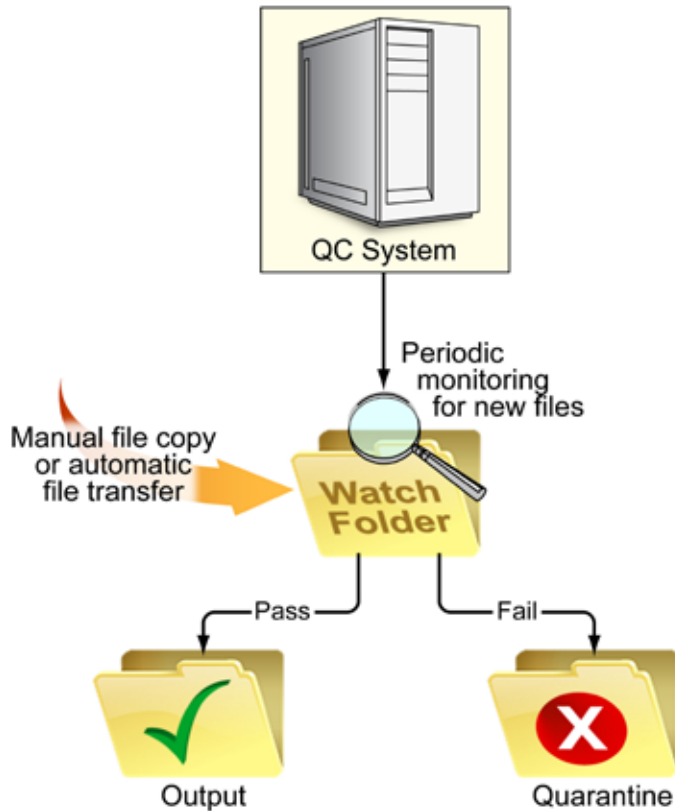


Figure 1.9. Semi-automated QC flow.

Automating File-Based QC

Content readiness means consistent high quality every day, in every asset that is stored for distribution. And that may add up to many thousands of hours of material, particularly in a VOD service. Manpower is a costly and admittedly fallible resource for checking content quality on this scale of quantity.

Semi-Automated QC Workflows

A far better solution to the problem is some degree of automation in the QC effort. The more automation the better, of course, but a partially automated workflow can ensure content quality very effectively. Figure 1.9 illustrates this scheme. A semi-automated work flow is optimal for small-to-medium operations and tasks like ad insertion.

In a semi-automated workflow, the QC System is in charge of passing files from acquisition to storage. On the input side, the QC System periodically polls one or more Watch folders and creates a directory listing to detect content that needs to be ingested. Both manual file copy operations and automated transfers via the “catcher” can go into the Watch Folder(s) anytime during the day or night.

When a new asset arrives, the QC System runs tests to measure the quality and compliance of the material. Based on the Pass/Fail results, the file goes to either an Output folder or a Quarantine folder. In the latter case an e-mail alert can be sent to an operator who can repair or reject the content as he/she sees fit.

Asset Management and the Fully Automated Workflow

A fully automated QC workflow is the solution of choice for enterprises that handle VOD libraries or other material in great volume. The amount of content that must be reviewed for aberrations and errors makes manual evaluation impracticable. The overarching term for this content is media assets.

The heart of a media asset is known as the “essence,” which consists of video and/or audio material. When metadata is added to this, the sum is “content,” but the package is still not complete. With the addition of “usage rights” the viewable media asset is complete. Figure 1.10 is a symbolic view of the hierarchy.

A media asset has value because, like any inventory item in any company, it costs money to acquire it. The cost is embodied in the usage rights. Cable providers pay for the right to show a particular program a limited number of times, and usually for a limited time span.

The need to manage media assets is pervasive across the workflow. Asset management tasks include searching for and retrieving assets and importantly, keeping track of the usage rights and the playouts accrued for each asset. A media Asset Management System (AMS) is the clearing house for all such transactions. The key to making an AMS architecture work is ensuring the quality of the metadata stored with the assets.

With an AMS and effective metadata, it is feasible to automate the operation and move beyond a semi-automated model where the QC tools must look for files and then test them. Under the control of an Asset Management System, QC is based on where an asset is in the workflow.

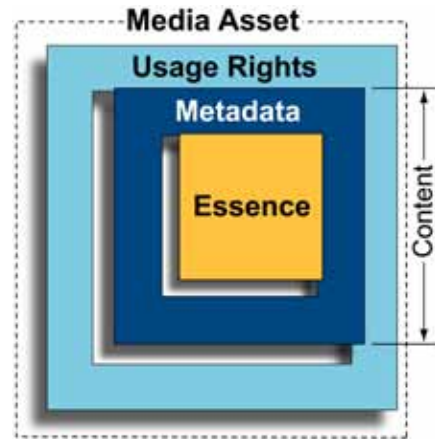


Figure 1.10. A hierarchy of elements makes up the media asset.

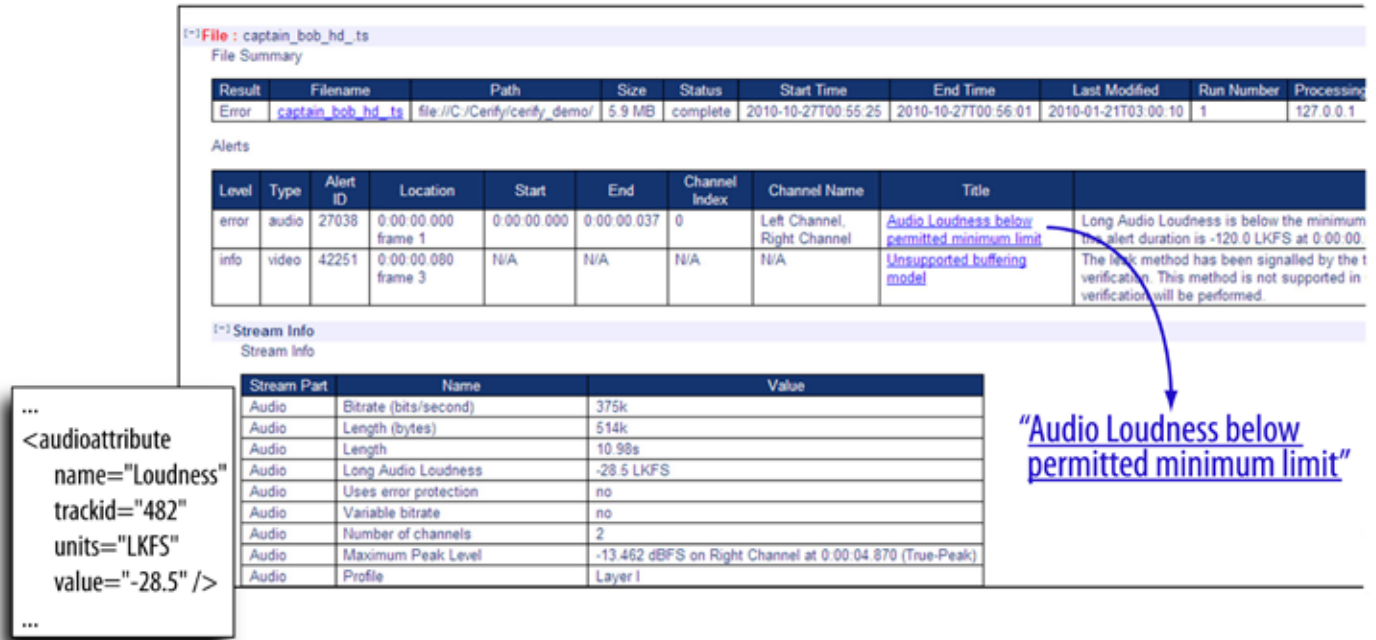


Figure 1.11. A human-readable QC system report.

QC Reporting

A QC system produces a report upon completion of its tests on each file or set of files. That report can be saved with the asset and stored in the Asset Management System's database. It includes both errors and measured values such as play duration, peak levels, and more. The human-readable components of the report enable operators to respond efficiently to any detected errors that were found in the file. To implement a fully automated QC process, it is of course necessary to include machine-readable data in the report.

The technology that supports these exchanges between machines and also facilitates human intervention is the XML text format, a solution borrowed from the IT world. The box on the left of Figure 1.11 depicts a short segment of XML code as seen by the AMS. Adapted to a style sheet as shown in the figure, this same information is human-readable.

How is a Workflow Like a Web Browser?

How is all this data exchanged among the various systems in a workflow? The reality is that many workflow elements are simply software tools residing within computers. All this software needs communicate without adding a lot of complexity to the flow. Increasingly, conventional web services are being used to integrate software tools from diverse sources (vendors) into coherent systems.

These web services are actually the same proven protocols commonly used for web browsing. The functional needs of the video workflow are very similar, with local clients (say, the AMS) sending requests to remote servers and receiving responses. The Simple Object Access Protocol (SOAP) submits requests and receives responses over HTTP in a server-client architecture. The Web Services Description Language (WSDL) is, as its name implies, a machine-readable description of the operations available by web services. WSDL can automatically generate library code and user documentation for the services.

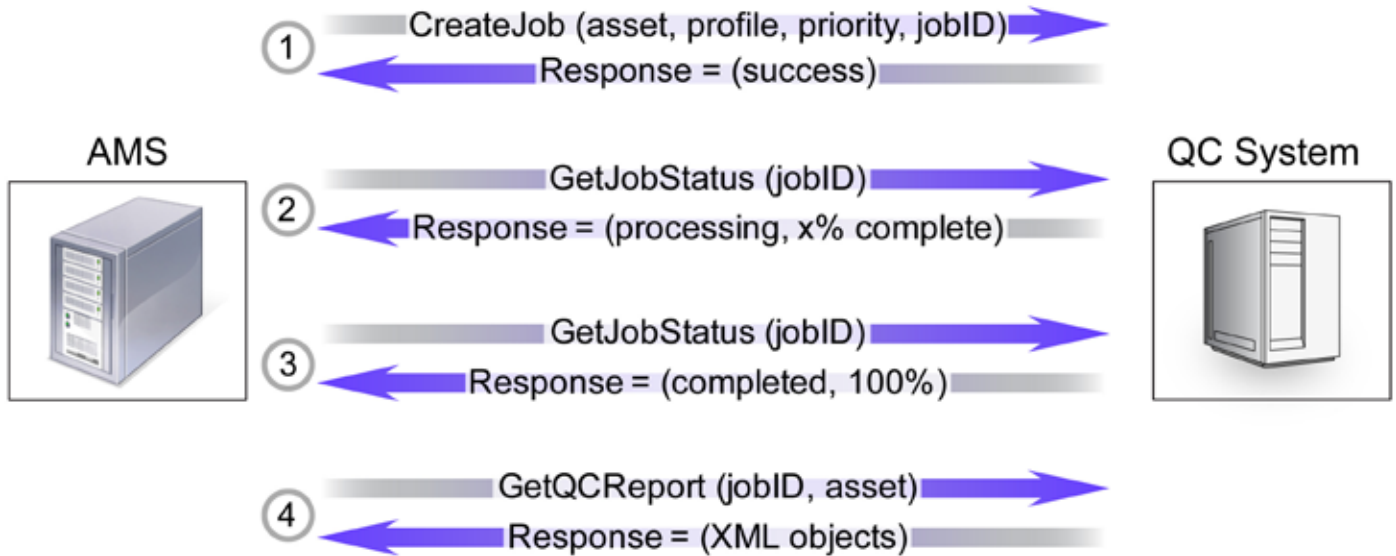


Figure 1.12. Steps in a transaction between the AMS and the QC system.

When Systems Communicate

Again, the interaction among systems is a process of sending requests and getting responses. An example of a transaction between the AMS and the QC system might proceed as depicted in Figure 1.12, with the AMS acting as the client:

1. When the AMS receives a new asset from the VOD catcher or other transfer source, it sends the QC System a new job request to test the file. Crucial parameters about the file are included in the request: where the file is located, priority (high, medium, or low), and a profile that describes the type of asset and the checks to be performed on it. The Job ID is a token that accompanies the file through the remaining steps in the process.
2. Next, query the QC System. The Asset Management System needs to poll the QC system periodically. What is the status of the job with this ID? In a web services model, it is always necessary to send a request to get a response. Without a request there would be no notification when the job gets completed, so frequent polling ensures a timely response. If the job is not finished, the QC system will respond with its current percent of completion.

3. Repeat the status queries until the response indicates that the job is 100% complete.
4. Lastly, get the QC report for the asset. The response comes back in the XML format explained earlier. Now the AMS can interpret the results and make decisions (guided by predetermined failure codes) about what to do if the file failed. For example a file that's missing captions will get treated differently than one that has failed the loudness test.

Summary

The spread of services like Video on Demand is a challenge to cable providers who must not only manage a much larger volume of content than ever before, but also deliver that content with consistently high quality. Many cable providers are joining their broadcast industry colleagues in adopting automated QC inspection tools to speed their file-based workflows. In a field where profits depend on vast queues of vibrant, instantly deliverable entertainment, content readiness is a top priority. Modern integrated solutions like the Tektronix Cerify automated video content verification system are the shortest path toward an efficient content readiness strategy.

Section 2: Live Network Monitoring of Digital Video Services

Live Network Monitoring Background

Today, digital video services do not readily show transmission impairments or video artifacts until the signal is extremely corrupt (also known as the “cliff effect”). This is due to RF symbol redundancy and error protection in the digital modulation schemes (e.g., PSK, VSB, QAM, COFDM, etc.). Once the signal becomes extremely corrupt, the digital video service normally freezes or becomes quite useless. To make monitoring today more challenging, the previously used analog test signals in the line just above active video are now stripped out during the compression process. Only the active video lines get compressed and sent. Monitoring equipment must now rely on new parameters for assessing the quality of the digital video service.

To maintain high quality using minimal bit rates, digital video services are segmented into several layers and associated measurements. They include the following:

- RF/IP layer with frequency, power level, modulation formats, etc., or IP headers, checksums, payloads, packet timing (jitter), etc.
- Transport Stream (TS) layer with headers, payloads, continuity counters, Program Clock Reference (PCR) timing, and Program Specific Information (PSI) tables (basic electronic program guide)
- Packetized Elementary Streams (PES) including headers, payloads, audio/video decode and presentation timing (also known as Access Units)

- Elementary Stream Sequence headers including codec format, frame size and rate
- Picture frame slice headers, Macroblocks (16x16 set of pixels), and finally Blocks (8x8)
- Audio Frames or Access units in small blocks of time (e.g., 32 ms for Dolby AC-3 at 448 kbps using 5.1 surround)

Monitoring equipment must be able to validate each of these many layers in order to achieve high confidence that the digital video service can be received and decoded on any compliant digital TV or set top box. Anything less would question the mandated interoperability between the service provider and the TV.

Initially, MPEG test equipment vendors submitted a group of important transport stream monitoring requirements for broadcasters to ensure encoder/multiplexer completeness (ETSI TR 101 290). This is an excellent standard for testing digital services, but it only covers one of the many different layers. For example, to say that the TS headers have been measured and comply to the ETSI TR 101 290 requirements has nothing to say about the audio levels or the picture quality being delivered. One must be able to traverse from the highest layer of RF/IP all the way down to the pixel or audio level before one can be confident that the digital video service is acceptable. There are two different approaches to help maintain high quality. One approach is called quality of service (QoS) testing and monitoring which looks for errors at the physical layer and TS layer. The second approach is called quality of experience (QoE) testing and monitoring that focuses more on the video and audio aspects of the decoded program. Both methods are very important, but approach the issue of testing and monitoring in very different ways.

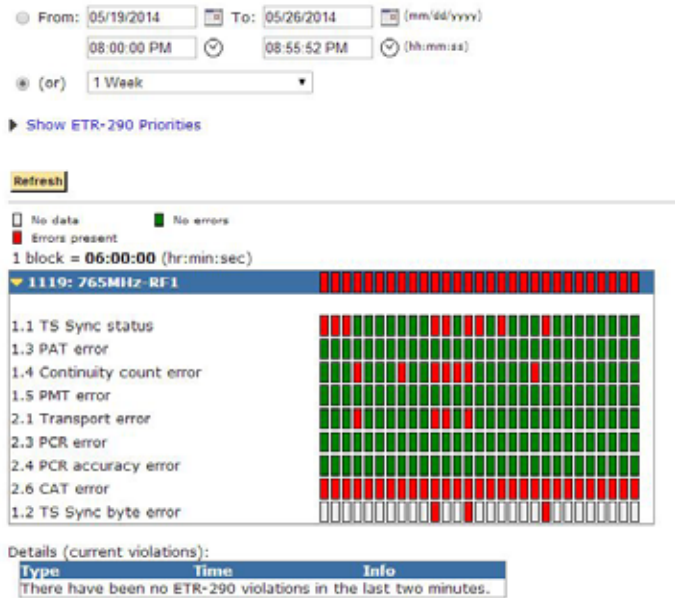


Figure 2.1. ETR290 history over seven days on channel 765 MHz (EIA-119).

Quality of Service (QoS) vs. Quality of Experience (QoE)

Using ETSI TR 101 290 (ETR290) to measure the TS is a good way to measure QoS. The ETR290 test includes three levels of testing:

1. **Priority 1** for highly important tests such as Sync error/loss and missing tables or packets
2. **Priority 2** for demodulation failures and PCR/DTS/PTS timing failures
3. **Priority 3** for Electronic Program Guide failures and buffer failures

Testing and monitoring with ETR290 is a good way to understand the health of the network and TS. But, it will be

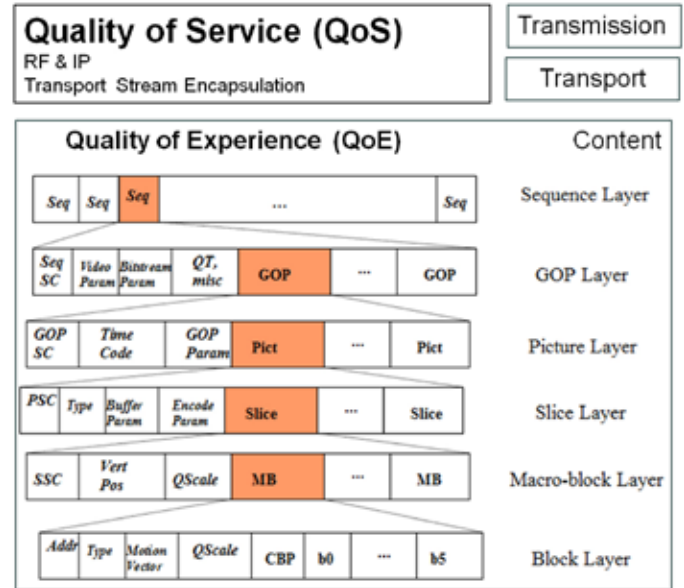


Figure 2.2. QoS and QoE layers.

almost impossible to determine how badly the audio and video service will be affected by an ETR290 failure. Figure 2.1 summarizes one of many RF channels on a local cable TV network. This specific channel is being reviewed here because it was known to have significant RF impairments. As we can see, there have been several TS Sync and dropped packet issues over the last seven days.

Due to the problematic nature of this specific RF channel, it becomes imperative to also look at RF characteristics that might be leading to TS problems. RF testing will be discussed in more detail later in this Technical Brief.

A good graphical example of QoS and QoE layers is shown in Figure 2.2, which shows the different layers that will be evaluated during QoS and QoE testing and monitoring.

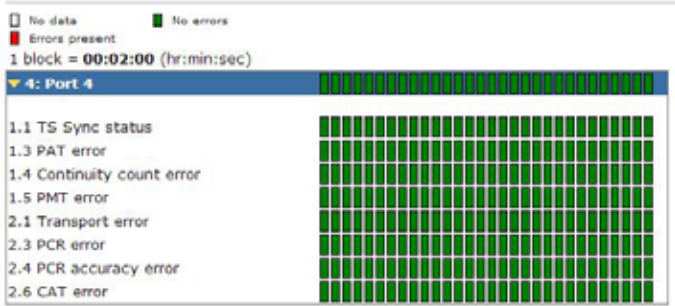


Figure 2.3. ETR290 with no TS errors.

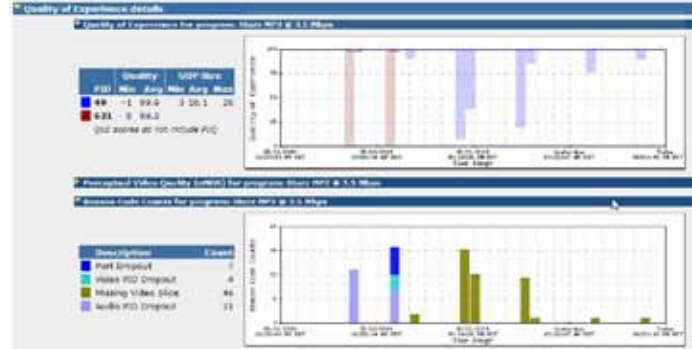


Figure 2.4. QoE plot with audio/video missing packets and slice errors.

To contrast QoS and QoE, think of QoS as a means of rating the quality of the signal which when error free, should create a good digital video service at the TV or set top box. Think of QoE as a means of watching the video and listening to the audio and rating the quality independent from the physical or TS layer quality. We can easily have a bad QoS with a bad QoE, but there are times when we can have a good QoS and still measure a bad QoE due to video or audio coding problems somewhere upstream. QoS will tell us when the TS is broken, and by default, the QoE should degrade also, but it may be more important to know when the QoE drops no matter if it is due to poor RF conditions, or upstream codec failures.

An example of a good QoS and bad QoE can be seen in Figures 2.3 and 2.4 where we show perfect ETR290 with bad audio and video QoE issues. These are good case examples where we know the transmission was perfect, and the error **must** be coming from the content provider.

The conclusion of Figures 2.3 and 2.4 tell us that the encoder or multiplexer from the content provider is failing, but the transmission is perfectly fine. The solution here would be to make a call to the content provide to correct equipment failures.

A good monitor should go deep into the QoE testing and monitoring of every layer of every video and audio service in every TS. Whenever the monitor detects an audio or video codec command that is in error, it denotes a drop in QoE. The monitor increases the relative weight of each video protocol impairment (i.e., syntax or semantic) depending upon the type of video frame (i.e., I, B, or P), as well as where the impairment landed in the frame (i.e., center vs. corner). See Figure 2.4 for the drop in QoE ratings related to audio and video errors.

This is an excellent example of QoE ratings where bad audio and video protocol drop the scores from a perfect 100 down to something significantly lower depending upon the error in the video or audio frame, and the repetition of the error.

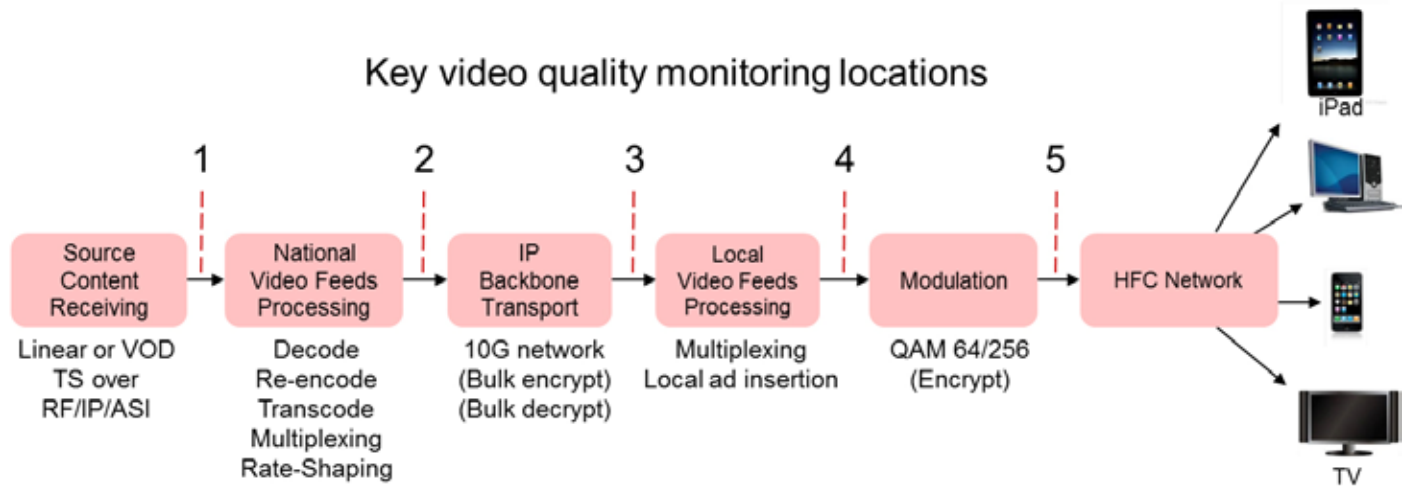


Figure 2.5. Five suggested monitoring points from ingest to egress in a typical digital video cable network.

QoS Monitoring Tradeoffs

Measuring QoS can help us to understand the relative health of the transmission path, but it can mislead an operator to thinking that a video service/program channel is error free when it might have a codec error several layers below the TS. Think of QoS as a good indicator of the transmission path and multiplexer completeness, but QoE is still needed to determine the health of each digital video service. Another minor point with ETR290 is that it is for the entire transport which often includes several digital video services. If you see a continuity counter error (i.e., dropped TS packet), you will not readily know which video service/program channel was affected, or if it was an audio packet or video packet.

Key Video Quality Monitoring Locations

Choosing the right monitoring locations is just as important as which parameters to measure. See Figure 2.5 for an example of typical monitoring locations in a cable TV architecture.

We need to monitor at the ingest (test-point #1) for two reasons:

1. To make sure that we have a clean error-free link from the content provider to the ingest receiver. If we have dropped or erred TS packets, then we already know that the digital video service will suffer.
2. We need to monitor the encoded service to make sure that there are no audio or video syntax or semantic errors. This is because TV's and set top boxes are not designed to handle services with errors. Receiving a clean ETR290 transport with embedded audio/video errors means that the QoE is going to be less than perfect. Therefore, it is critical to always measure the source.

In order to localize the service, it might be transcoded into a different format, or simply modified slightly by inserting a commercial at various locations in the stream (test point #4). With today's highly compressed services, problems can occur in the video layer as well as the audio layer. Therefore, it is also critical to monitor the output of any equipment that modifies the stream.

Now, what if an error was detected after the QAM modulator? It would now be possible to know which piece of equipment introduced the error based upon the results of the monitoring equipment just upstream from the modulator.



Figure 2.6. Program Dashboard Reporting.

Using QoE Scoring Methods to Develop a Service Benchmark

With video services often originating from a wide variety of sources, there tends to be a wide difference in QoE impairments when all are compared against each other. In a large collection of ingest sources, a QoE report can be helpful to focus work on the worst or bottom 10% rather than treat all sources equally. Another idea is to congratulate the top 10% as high achievers. QoE ratings and reports can also be generated on a daily, weekly, or monthly basis depending upon your needs and urgency of reports.

One example of scoring programs is to create a dashboard of a set of measurement categories. Within each category, a summary of all of the programs is shown. Figure 2.6 shows eight categories of measurements with a rating summary for each of the many programs. In most cases, green is a good sign, and red is a bad sign. The dashboard makes it easy to quickly see the health of the entire network.

Summary | Templates | **Transport** | BFS | OCAP
 Summary | Programs | PIDs | Tables | PMTs | DSM-CC | IP Stats

Editing Program Alerts

- Select alert type: Perceptual Video Quality (eMOS)

Out of a possible score of 5, generate an alert when the perceptual video quality (eMOS) score for any program goes
- Use program group:
 - 1 program selected:

Port#	Port	TSID	Pgm	Name
1	Sports Files	Any	2	No Name
- When alert is generated:
 - Save in [Alert History](#)
 - Send SNMP trap (configure trap host in the [System Settings](#))
 - Send email Always (or)
 - At most email(s) in

<input type="checkbox"/>	Name	Email
<input checked="" type="checkbox"/>	Administrator	dennis@tek.com

Figure 2.7. Configuring an Alert condition.

Detecting Artifacts and Impairments before Subscribers

The last thing that a network operator wants is to get a call from a frustrated subscriber explaining audio or video problems in a service. The more frustrated that the customer becomes, the more likely they are to cancel their subscription. Therefore, it is imperative to catch audio and video impairments before customers call in with complaints.

In order to track such impairments, the monitoring system should allow for triggered alerts to send out an SNMP trap, or an email message to one or more operators or administrators. Alert definitions should be available in a wide array of choices from no audio or video over a defined window of time, to audio DialNorm/Loudness deviations or video over-compression. Once the alerts have been defined and then applied to each or all of the video services, the operators can rest assured that all defined alerts will raise immediate SNMP traps and or email messages when triggered. For example, Figure 2.7 shows how an alert message gets triggered via email whenever the video quality drops below 3.75 (i.e., 25% drop).




Figure 2.8. PVQ showing blockiness impairments due to running MPEG-2 HD and only 3.5 Mbps (over-compression related to motion).



Figure 2.9. Mezzanine-level compression at 50 Mbps without blockiness or over-compression.



Figure 2.10. Transcoding from 50 down to 3.5 Mbps resulting in blockiness and over-compression.

Perceptual Video Quality: Compression and Encoding Artifacts

Within HD encoding systems, two choices are available related to bandwidth (bandwidth). We can choose variable bit rate (VBR) where the encoder uses less bits as the scene becomes stable or easy to compress, and then we have constant bit rate (CBR) which is excellent if the video is to be remultiplexed downstream with other programs. The problem with CBR is that we often get a variable quality depending on the complexity of the video scene. One moment there is enough bandwidth to preserve all of the original video attributes, and then when the scene becomes much more complex, and the encoder does not have enough bandwidth to keep up with the complexity, the encoder must over-compress parts of the video causing a reduction in the video quality (i.e., blockiness). The perceptual video quality (PVQ) test performs a quality rating based upon over-compression. If a video program always has enough bandwidth to maintain high quality, then

the PVQ score will be at 5.0. Figure 2.8 started with an HD program in MPEG-2 at 50 Mbps with a PVQ rating of 5.0 for the entire 30-minute clip. After running the clip through a transcoder and recompressing the clip to 3.5 Mbps in 1080i HD, the scenes with low motion stayed at 5.0, but then drop significantly lower as motion increases.

Figures 2.9 and 2.10 show the results of transcoding from 50 Mbps to 3.5 Mbps over a 30 minute clip. To provide a visual understanding as to how these two video services compare, let’s look at the same frame in the two clips where a fast action scene is being played out. The frame below correlates to about 20 minutes into the above graph where the PVQ plot drops from 5.0 to about 2.5. The example is where a pirate is being catapulted over the deck of a ship. Notice that in Figure 2.9 at 50 Mbps, the pirate is not blocky, although he may look a little blurry due to the film and camera shutter-speed. The same pirate in Figure 2.10 at 3.5 Mbps is heavily over-compressed due to high motion and limited bandwidth.

Choose report type(s):

Select all types

<p>CGMS</p> <p><input type="checkbox"/> CGMS - No restriction</p> <p><input type="checkbox"/> CGMS - No copy</p> <p><input type="checkbox"/> CGMS - 1-copy</p>	<p>Closed Caption Data</p> <p><input checked="" type="checkbox"/> 708</p> <p><input checked="" type="checkbox"/> 608</p> <p><input checked="" type="checkbox"/> SCTE-20</p> <p><input checked="" type="checkbox"/> Combined</p>	<p>Other Data</p> <p><input type="checkbox"/> Scrambled</p> <p><input type="checkbox"/> Discontinuity</p> <p><input type="checkbox"/> Teletext</p> <p><input type="checkbox"/> Subtitles</p>
---	--	---

Generate Report

No data available Data present for some time
 Data present Data present but in error for some time
 Data present but in error

1 block = 00:02:00 (hr:min:sec)

Port #:	Port Name	Program	Name	Current Status
1	Sports Files	2	No Name	UP

Closed Caption: 608
 Closed Caption: 708
 Closed Caption: SCTE-20
 Closed Caption: Combined

Figure 2.12. Closed Captioning Report Setup and availability bars.

Regulatory Compliance

Many government agencies require digital video service providers to also provide additional services beyond audio and video. One such requirement is for closed captioning (or subtitling/teletext) services. Another requirement is for audio loudness.

Closed Captioning

In the case of Closed Captioning (CC), or Subtitling/Teletext, many governments require this service for the hearing impaired. If required, there can be a hefty fine for failing to carry such services. A monitor should have the ability to detect the presence and absence of such services. These results, like all other details, should continuously be written to a 60-day database available to anyone with a browser. Figure 2.12 shows the supported formats, as well as an example of 708 CC in a digital video service.

An example of compliance testing is to know if 85 percent of all content is including closed captions. Figure 2.13 shows a summary of such a test where about 150 digital video services are in compliance, yet over 300 are still failing to run closed captioning over 85 percent of the time.

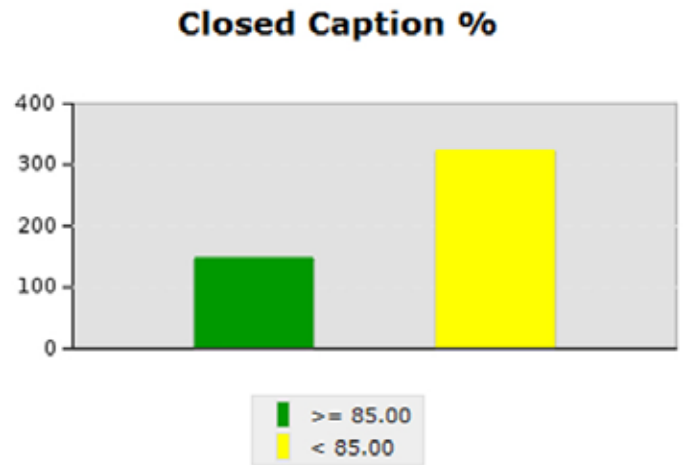


Figure 2.13. Percent of digital video services with closed caption support over 85% of the time.




Figure 2.14. Dynamic DialNorm and Audio Loudness.

Figure 2.15. Configuring Audio Loudness Program Alerts.

Audio Loudness

Another important regulatory compliance requirement around the world today is for audio to be within a small range of the DialNorm value. The ITU BS 1770 standard describes the formulas used to calculate short term and long term audio loudness levels. The ATSC A/85 (2013) Recommended Practice document further restricts the limits to a window from 3 to 10 seconds for a short term test, whereas the EBU R 128 document recommends 3 seconds for the short term measurement window.

These standards are designed to address the difference in loudness values between programs, as well as between advertisements. Failure to comply can also carry a hefty fine when not maintained. A Video Quality Monitor should have

the ability to monitor loudness on every audio element on every program in every TS. Figure 2.14 shows a dynamically changing DialNorm (dotted line) value and its associated Loudness measurements over time.

One use-case example for Audio Loudness testing was to configure a trigger to alert anytime that the loudness values deviated more than 5 LKFS (Loudness, K-weighted, relative to Full Scale) for over 15 seconds. See Figure 2.15.

This setup was configured for a specific group of channels and then ran 24x7. As seen in Figure 2.15, the DialNorm may change from program to program, or during commercial insertion. As the DialNorm changes, the monitor uses that new reference value to measure the delta to the measured loudness over that same time period.

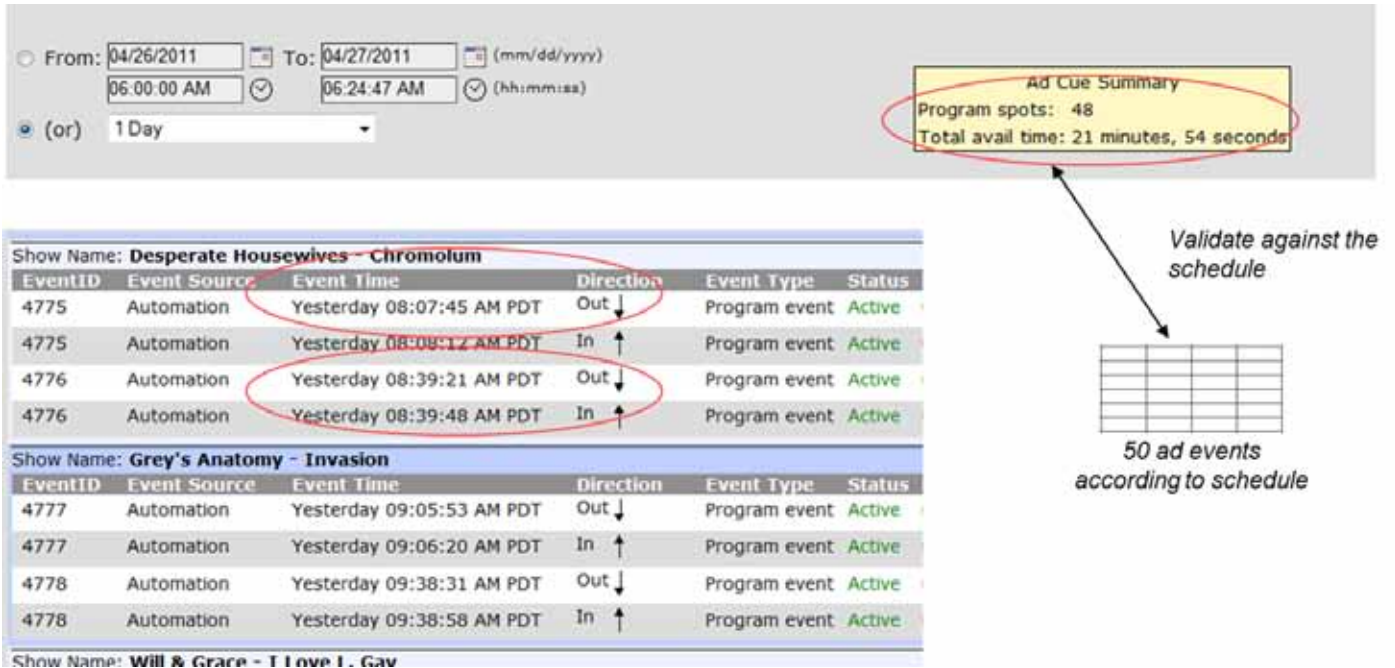


Figure 2.16. SCTE 35 Ad Cue summary.

Ad Verification

Today, an average of 2% of digital advertisements fail to air or air incorrectly due to scheduling, insertion, and other errors. In addition, advertisements often suffer the same audio and video quality issues that plague regular programming. As a result, monitoring and auditing capabilities are critical to successful ad delivery.

A monitor should provide the most complete digital ad insertion monitoring capabilities by combining real-time monitoring and alerting with historical auditing across the entire channel lineup in all advertising zones. The monitor should also deliver extensive data including historical thumbnails which improve digital ad insertion on any platform, allowing engineering teams to ensure proper function of insertion technology by identifying and correcting system errors when they occur.

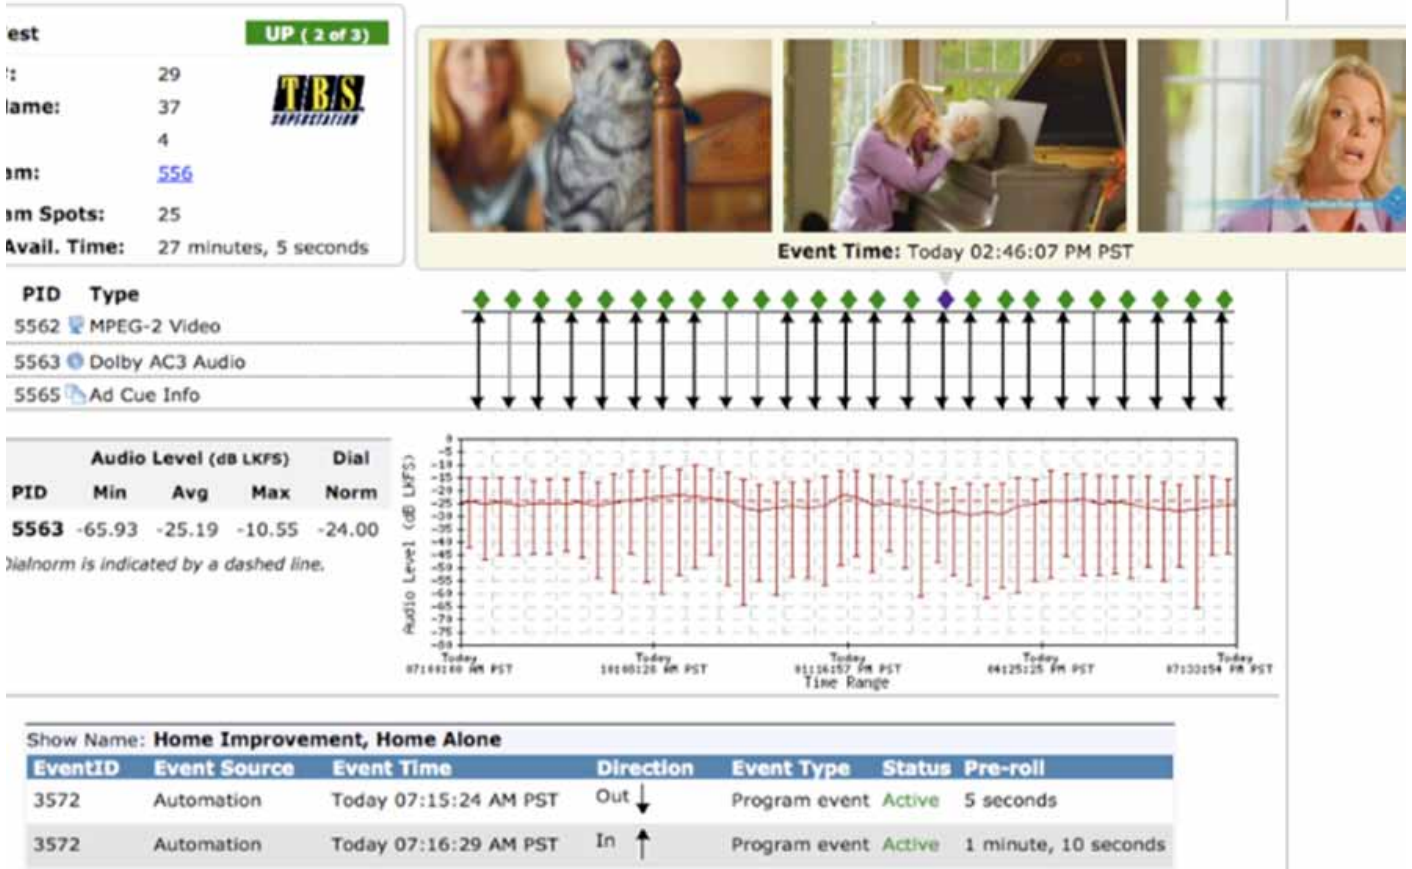


Figure 2.17. Video thumbnail around the SCTE 35 Ad Cue event.

In addition, by performing ad insertion verification, ad sales groups are able to provide higher levels of customer service, resulting in greater revenue potential. By strategically placing monitoring probes in each of your ad zones, you should be able to monitor and be alerted on all insertion opportunities network-wide, as well as issues that arise from problems. Figure 2.16 shows an example of the SCTE 35 commands in a digital video service showing when to trigger an ad-insert. The summary also shows that 48 ads were detected even though the expectation was for 50 ads.

An example that helps in Ad Verification is to show a video thumbnail just prior, during, and after the Ad Cue trigger from SCTE 35. Figure 2.17 shows the SCTE 35 triggers and audio details, but also includes three video thumbnails surrounding the selected key event.

RF Stats									
Port Name	Freq	Current Activity	Const.	Pre-RS BER			Post-FEC erred TS packets		
				Min	Avg	Max	Min	Avg ↕	Max
675MHz-RF1	675 MHz	UP		0.00E+0	4.23E-5	2.91E-4	0.0/sec	43.4/sec	475.0/sec
753MHz-RF1	753 MHz	UP		0.00E+0	1.35E-5	2.64E-4	0.0/sec	0.6/sec	16.0/sec
669MHz-RF1	669 MHz	UP		0.00E+0	2.49E-6	5.58E-5	0.0/sec	0.1/sec	5.0/sec
57MHz-RF1	57 MHz	UP		0.00E+0	2.16E-6	4.15E-6	0.0/sec	0.0/sec	0.0/sec
63MHz-RF1	63 MHz	UP		0.00E+0	0.00E+0	0.00E+0	0.0/sec	0.0/sec	0.0/sec
69MHz-RF1	69 MHz	UP		0.00E+0	0.00E+0	0.00E+0	0.0/sec	0.0/sec	0.0/sec
177MHz-RF1	177 MHz	UP		0.00E+0	0.00E+0	0.00E+0	0.0/sec	0.0/sec	0.0/sec
183MHz-RF1	183 MHz	UP		0.00E+0	0.00E+0	0.00E+0	0.0/sec	0.0/sec	0.0/sec
189MHz-RF1	189 MHz	UP		0.00E+0	0.00E+0	0.00E+0	0.0/sec	0.0/sec	0.0/sec
195MHz-RF1	195 MHz	UP		0.00E+0	0.00E+0	0.00E+0	0.0/sec	0.0/sec	0.0/sec

Figure 2.18. RF Monitor measuring BER on more than 100 RF QAM channels from a local cable TV service provider.

RF Monitoring

QoS starts at the physical layer and measures parameters such as signal level, noise, and forward error correction (FEC). All digital RF transmissions (e.g., 8PSK, QPSK, 8VSB, COFDM, QAM, etc.) assume a lossy environment and add additional data for redundancy and error correction (i.e., inner-FEC). These additional overhead rates range from 9/10ths (very little) up to and beyond 1/2 (almost two for one, or full redundancy) depending upon the RF standard and how much noise or interference is expected to be received or recovered from. Interleaving is another good means to prevent loss of data (or to aid in error recovery), but does not add any additional overhead to the stream. Interleaving works best by taking a short burst of errors and spreading them out over longer amounts of time making the FEC more effective. If the bit error ratio (BER) for the received transport stream (after demodulation and inner-FEC) is less than about 5×10^{-3} , then it is assumed to be quasi-error free (QEF) due to the Reed/Solomon (R/S) portion of the transport being

able to correct up to eight errors per TS packet (five errors for 8VSB). The reason for calling it quasi-error free is due to the statistical nature of random errors where we might rarely find a short burst of nine or more errors in a single TS packet causing the R/S algorithm to pass a TS packet to the video/audio decoders with one or more bit errors. Figure 2.18 shows different RF ingest points (usually satellite and terrestrial) as well as the common egress QAM RF feed going out to the many set top boxes. Each of these RF ingest and egress points need to be monitored for QoS and preferably a post-FEC BER of 0.0. Figure 2.18 is from a QAM receiver measuring over 100 RF channels in parallel from 57 MHz up to 1000 MHz. It is important to note that the “Post-FEC Avg” column was selected and sorted showing the worst offenders at the top of the list. In this specific case, channel EIA-104 at 675 MHz is receiving a huge amount of RF impairments leading to a very high BER in both Pre-RS and Post-FEC. In this case, it appears that the TS from 675 MHz is getting about 43 TS packet errors every second.

Another place to see when BER negatively affects a TS is in ETR290 Priority 2.1. This test looks at the Transport Error Indicator bit (one bit after the TS sync byte) which is set by every RF digital demodulator. If the bit is a “1” value, then we know for a fact that the TS packet has one more bit errors inside the 188-byte TS packet. This is a bad sign, but it could be even worse if it happens to land inside an anchor video frame (i.e., I-frame) being used to display many other video frames. If an error rarely occurs in an audio TS packet, or a bi-directional video frame, then the error will probably be missed by the viewer, or at worst, only occur for a small fraction of a second. See Figure 2.18 for an example of a demodulated TS with errors (ETR290 2.1 failures).

There are usually two key points in the network where RF monitoring is important to any broadcaster of Cable, Terrestrial, or Satellite digital video services. Those two points are at the ingest where source content is brought into the system (by dish or aerial antenna), as well as the egress where the signal leaves the system. The most important point to RF testing is signal or transport integrity. This is achieved by monitoring the RF Level, noise, and most importantly, the Post-R/S BER (should be always 0.0). It is usually OK if the pre-R/S BER is worse than 5×10^{-3} as long as the post-R/S BER is zero, and the ETR290 Priority 2.1 test is always green/good.

Another good reason to measure the RF signal at ingest and egress is to make sure that we stay as far away from the digital cliff as possible. This means maintaining high levels of signal power, low levels of noise, and minimal pre-R/S BER. As long as these tests stay within predefined values, then we can be sure that the stream is transmitted and received error-free.

Width	Height	Bit rate
1280	720	3 Mbps
960	540	1.5 Mbps
864	486	1.25 Mbps
640	360	1.0 Mbps
640	360	750 Kbps
416	240	500 Kbps
320	180	350 Kbps
320	180	150 Kbps

Figure 2.19. An example of eight different ABR formats and rates available from within a manifest file.

Even though we have concluded that the stream was transmitted error-free, there may have been audio, video, or multiplexing errors embedded inside the stream that are hidden to any RF or ETR290 testing. In this case, it would be advisable to also measure the audio and video protocol for correctness and decodability. The additional test (QoE) is only required on the RF-ingest side because the egress side should have already been tested for compliance prior to modulation and transmission.

New Requirements for Monitoring OTT & Adaptive Bit Rate

Over the top (OTT) technology is new and growing quickly. Where the previous sections discussed digital video services being broadcast from one to many, OTT is more of a one to one service, usually over the HTTP broadband network. It also includes dynamic bandwidth (bandwidth) allocation (slightly decreasing or increasing every few seconds) depending upon bandwidth availability at the users' PC, tablet, or mobile phone. This dynamic change in bandwidth is called Adaptive Bit Rate or ABR. Figure 2.19 shows an example of the many possible video formats and rates.

System Status

Engine Status: UP

NTP Status: UP

Database Used: 4%

Device Status

LAN 1 (web interface) UP

LAN 3 (ABR) UP

Data Activity Status

Current total input rate: 112.301 Mbps

Origin Servers

[Edit](#)

Origin Server	Current Bitrate	Bandwidth Limit
10.0.3.46	113.548 Mbps	<input style="width: 80px;" type="text" value="300.000000"/> Mbps

Current Input Bitrate								
Port #	Name	Device	Current Bitrate	Format	Encrypted	VOD/Live	Rep. Count	Status
0	Show Time -1	ABR	9.813 Mbps	HLS	No	Live	7	UP
1	Show Time -2	ABR	9.745 Mbps	HLS	No	Live	7	UP
2	Show Time -3	ABR	9.761 Mbps	HLS	No	Live	7	UP
3	first bipbop	ABR	826.536 Kbps	HLS	No	Live	4	UP
5	NTSC Test - 1	ABR	826.689 Kbps	HLS	No	Live	4	UP
6	NTSC Test - 2	ABR	826.310 Kbps	HLS	No	Live	4	UP
7	NTSC Test - 3	ABR	826.724 Kbps	HLS	No	Live	4	UP
10	Impaired 404 & 302	ABR	19.524 Mbps	HLS	Yes	Live	6	UP
11	Impaired 404	ABR	20.092 Mbps	HLS	Yes	Live	6	UP
12	Encrypt. -1	ABR	20.031 Mbps	HLS	Yes	Live	6	UP
13	Encrypt. -2	ABR	20.029 Mbps	HLS	Yes	Live	6	UP

Figure 2.20. ABR monitor with 13 unique ABR services.

In order to make this shift quickly and seamlessly, the content must be compressed to up to 16 different rates, and then divided into small chunks of time (e.g., 2-second chunks). The link between the service provider and end-user-equipment negotiates for the best estimate of bandwidth for the next few seconds. This scenario relies on a manifest file describing the content and available rates. The end user does not see any of the negotiations, but instead should simply be able to view the highest quality video within the browser or application.

To make sure that all of this goes off without any problems, the ABR monitor will check each manifest file as well as the many supporting rates over time. This process can be performed actively or passively. The active approach will attempt to watch every defined video over every rate and create a report showing the results of the video over rates and time. Figure 2.20 shows the ABR monitor actively monitoring several video services. In each case, the service will include the same content or video program over several different bit rates.

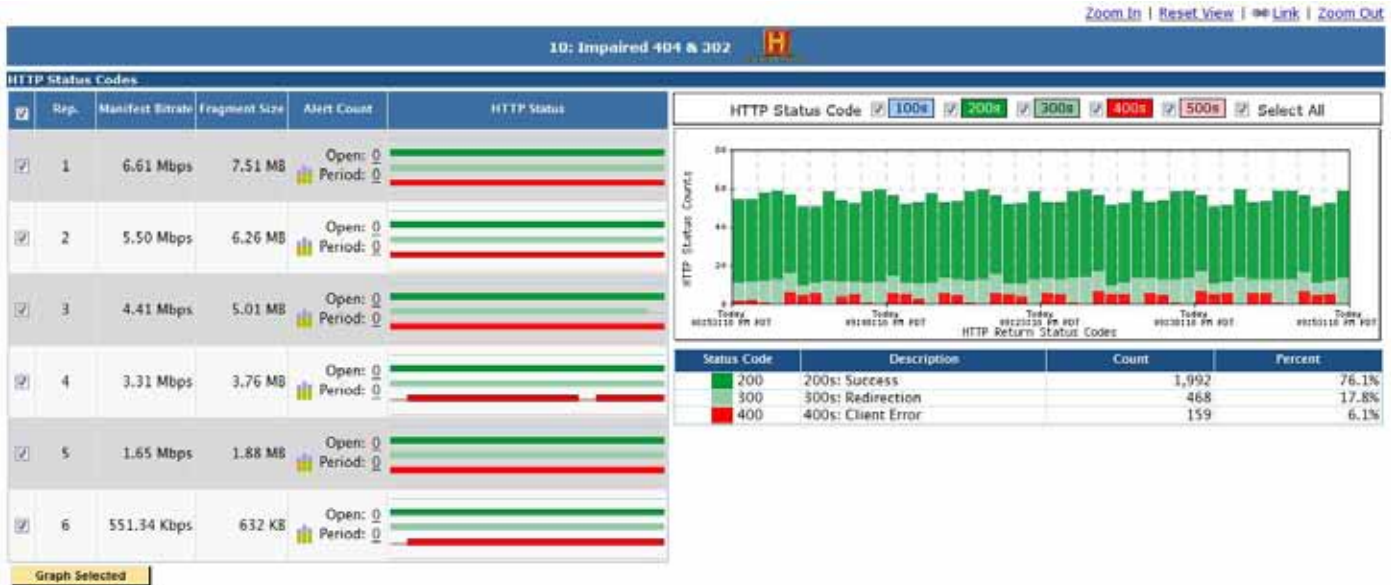


Figure 2.21. Multiple formats and rates for a single service with occasional HTTP failures.

For the passive process, the ABR monitor will only monitor the selected stream based upon the end user request (only one of the many rates listed in the manifest file).

As mentioned earlier, each service is divided into small chunks of time allowing viewers to dynamically switch to a higher or lower bandwidth depending upon how much bandwidth is currently available. Figure 2.21 is an example of one specific service that offers six different video formats and bit rates from

6.6 Mbps down to 550 kbps. The plot is over several hours showing that most of the HTTP handshakes were successful, but a few were redirected to another server, and also a few 400-series client errors (e.g., Page Not Found).

Since the quality of any ABR service is heavily dependent upon the source material, it is wise to also measure QoS, QoE, and PVQ on each upstream digital video service.

Summary

To summarize digital video service testing and monitoring, we need to place high importance on testing QoS to ensure that the transmission link is working well, but not to rely solely on this one metric. Next, we must rely on QoE because we know that it is possible to deliver a bad video or audio element in a TS and get a perfect QoS score. Therefore, QoE testing will verify that the audio and video elements are decodable on any compliant TV or set top box. Lastly, we must monitor audio loudness for differences in audio levels, as well as video over compression which is legal, but negatively affects to quality of the video.

These testing and monitoring requirements will help keep viewers happy and will reduce the amount of calls from subscribers complaining about adverse video services.

According to MRG Cable & IPTV Operator Surveys, the four top issues causing people to call in and complain were Macroblocking, Blackout, Freeze, and Audio Silence - see Figure 2.22.

Top Video Error*	Operators
Macroblocking	89.5%
Blackout	88.4%
Freeze	84.2%
Audio Silence	52.6%

Figure 2.22. Viewer Reported Errors – Poor QoE.

These top four errors account for 54% of the total number of view complaints. Of these errors, 60% were caused by the operator, and 40% were caused in the home. Therefore, proactively monitoring and measuring each service from RF/IP to individual pixels will help to keep the customers happy. This will reduce subscriber loss in today's highly competitive markets and will also reduce operational expenses.

Section 3: Network Troubleshooting & Diagnostics

Troubleshooting Background

Video Service Providers deliver TV programs using a variety of different network architectures. Most of these networks include satellite for distribution (ingest), ASI or IP throughout the facility, and often RF to the home or customer premise (egress). The quality of today's digital video and audio is usually quite good, but when audio or video issues appear at random, it is usually quite difficult to pinpoint the root cause of the problem. The issue might be as simple as an encoder over-compressing a few pictures during a scene with high motion. Or, the problem might be from a random weather event (e.g., heavy wind, rain, snow, etc.). In some cases, it is as simple as adding too many 3 dB RF splitters in the home.

No matter where the problem comes from, it is important to be able to quickly identify and log the audio or video problem that occurred, then identify or pinpoint the equipment (or network link) that needs attention. To identify and isolate problems, it is critical to have access or test points throughout the facility. The minimum set of test points in any network should be at the point of ingest where the signal comes into the facility, the ASI or IP switch, and finally egress where the signal leaves as IP or RF. With a minimum of these three access points, it is now possible to isolate the issue to have originated at either: ingest, facility, or egress.

To begin testing a signal that may contain the suspected issue, two related methods are often used, Quality of Service (QoS), and Quality of Experience (QoE). Both methods are useful in troubleshooting and analysis, but each of the two methods quantify issues using completely different metrics. Network Troubleshooting utilizes an MPEG Analyzer for both types of testing methods, and the benefits and faults of each type will be illustrated.




Figure 3.1. As the amount of noise per symbol increases, the MER drops, bringing it closer to the digital cliff.

Troubleshooting QoS

One simple explanation of QoS is simply the ratio of good bits to total bits. This is also related to the bit error ratio (BER). The optimum method for testing a transmission link is to take it out of service in order to make use of special test patterns. These patterns are repetitive, making it easy to quantify when a bit or byte is in error. The problem with testing live TV programs is that there is very little repetitive data to test. The use of BER is critical in monitoring the satellite signals because it does carry a small amount of redundancy. To make any additional tests beyond BER, we need to look to the MPEG^{1,2} and RF^{3,4,5,6}/IP standards for QoS as it relates to digital TV. Each RF standard has its own modulation requirements and characteristics. The ETSI TR 101 290⁷ document focuses on a variety of modulation types and defines measurement algorithms for modulation error ratio (MER) as well as many other RF and transport measurements.

Before we begin with the more complex measurements, let's agree to start with one basic, but extremely important RF test - signal power. TV receivers are designed to work with a reasonable amount of signal power. Too much, or too little causes the receiver to fail. Next, the receivers are designed to work with reasonable amounts of noise. With these two measurements, we get the signal to noise ratio (SNR). The higher the SNR, the easier it is for the receiver to recover the transmitted bits or symbols. The lower the SNR, the more probable the bit will be received in error. For digital TV transmission measurements, we refer to MER rather than SNR. As long as the receiver has a high MER, then the QoS is assumed to be very good. The problem begins when the receiver is near the fringe area of reception, or when low power and high noise corrupt the signal. This is often referred to as the "digital cliff" area. When approaching this

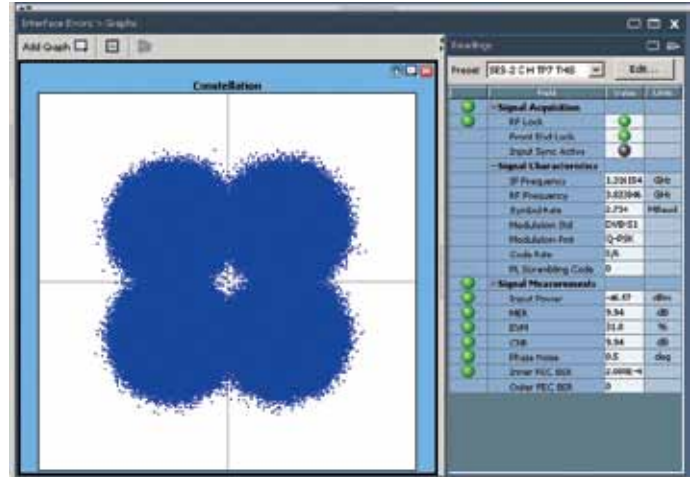


Figure 3.2. Even though this DVB-S QPSK signal includes high amounts of noise (poor MER) and some symbols landing in the wrong quadrants, FEC and R/S have effectively corrected all errors (QEF).

point, some of the symbols are incorrectly received. Figure 3.1 shows MER gradually decreasing as it approaches the cliff. After the MER decreases enough to cross the cliff, the TV picture and sound go from great to terrible. The first two important RF QoS measurements are signal level and MER.

Transmitting digital TV over RF is assumed to occur in relatively hostile domains, so redundancy is always added to overcome the loss or corruption of good symbols. Some digital RF modulation methods provide varying degrees of redundancy to account for different weather or interference conditions. The term used to describe this redundancy is called Forward Error Correct (FEC, or inner FEC). One additional layer for digital TV is called Reed/Solomon (R/S, or outer FEC). This inner code is the Viterbi convolutional coding and the use of both together is often called "concatenated" FEC coding. R/S can correct up to eight bad bytes in each transport packet of 188 bytes for DVB broadcasts (satellite, cable, and terrestrial). ATSC⁶ Terrestrial and SCTE⁸ Cable use slightly different amounts of R/S protection, but are also in less harsh environments and travel much shorter distances than satellite signals (ten bytes per 187 for ATSC 8VSB, three bytes per 122 for SCTE cable). For most RF transmissions, 100% of the digital transport can be received "effectively error-free" or "quasi-error-free" (QEF) as long as the RF bit error ratio (inner FEC) is better than 5×10^{-3} (i.e., 995 good bits for every 1000). R/S will take care of the rest resulting in zero received errors. In this case, the QoS of the demodulated signal could be considered perfect because 100% of all bits sent have been received error-free. Figure 3.2 shows a signal with a BER better than 5×10^{-3} , and therefore, the received transport is considered to be quasi-error-free.

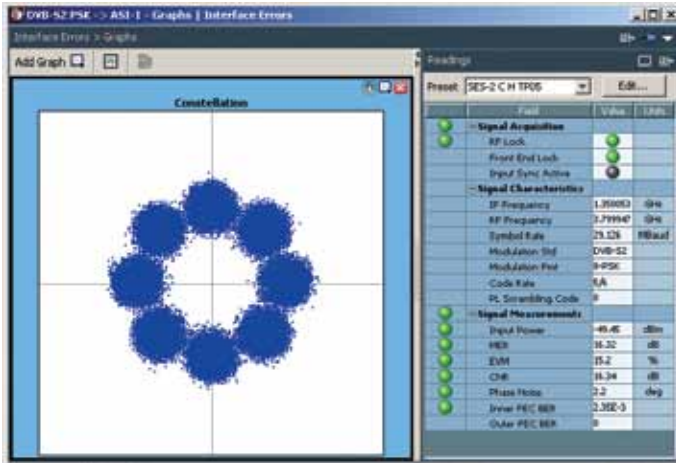


Figure 3.3. DVB-S2 8PSK signal with high amounts of noise, but enough Forward Error Correction to provide a quasi-error-free transport stream.

Much of today's video satellite transmission uses QPSK (e.g., DVB-S, DigiCipher II, DSNG, etc.) to send a transport stream from point to point, or point to multipoint. Recent technology improvements such as 8PSK from DVB-S2⁴ allow more bandwidth while using the same frequency spectrum as a QPSK signal. Figure 3.3 shows an example of a DVB-S2 signal carrying a 70 Mbps transport stream with 31 H.264² SD programs. Similar to the QPSK example, there is a significant amount of noise in each of the eight unique symbols, but enough FEC and R/S has been provided to effectively correct the misplaced symbols (QEF).

Similar digital modulation schemes can be used on the egress side for sending a transport stream over cable. Figure 3.4 shows a 256 QAM signal with well-behaved symbols in its constellation display.

Continuously measuring BER is a good thing, but when the error ratio exceeds $5 \times 10E-3$ and one or more errors make it into a transport packet (ISO/IEC 13818-1 – MPEG-2), it is almost impossible to tell what will happen to the video and audio. It might make no difference at all. If the error ratio is high enough ($>5 \times 10E-3$), then the errors will be landing in many of the audio and video frames. With this level of errors, it can be assumed that the problems will be noticed by the all viewers. The more difficult issue in monitoring a digital network is determining the extent of the problem on less frequent errors.

When monitoring ASI⁹ or video over Ethernet¹⁰, FEC is usually not performed as the delivery does not take place in a hostile environment. Therefore, it is assumed that the BER is always

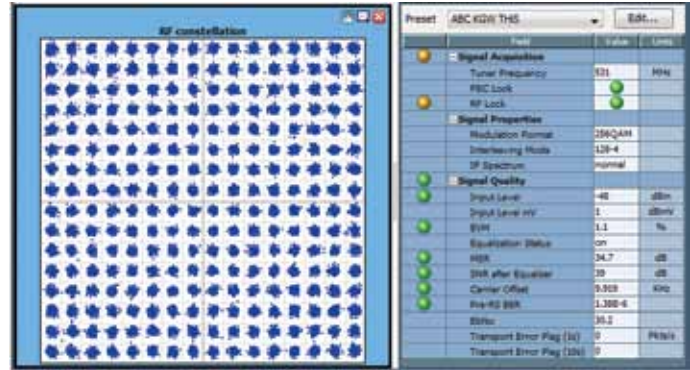


Figure 3.4. 256 QAM cable TV signal.

0.0. It would be nice if this were always the case, but given the complexity of IP switches, occasional failures or overloads do occur and cause IP packets to get dropped. And to make matters worse, the low-latency requirement mandates UDP (send and forget), which is unlike TCP/IP which uses a handshake for every IP packet. Video over IP commonly puts seven 13818-1 transport packets into each IP packet, so the occasional loss of a single IP packet translates into seven transport packets being lost. If video over IP used RTP, then the additional counters in the IP header make it trivial to know when an IP packet was lost. But, since most video over IP uses UDP, then we need another method to know when an IP packet has been lost. With MPEG-2 transport packets, each TV program includes separate 4-bit counters for each audio and video element, so it is possible to determine when a packet might have been dropped. Calculating the ratio of missing packets to total packets is possible, but not too helpful as each missing packet will usually generate audio or video degradation. The indication of a single or occasional missing transport packet is a very bad thing, so once the event occurs, troubleshooting usually begins immediately. The magnitude of the degradation will depend upon where the error occurred within each video frame.

Beyond RF and BER testing, the TR 101 290 document includes a recommended set of transport stream measurements for broadcast operators (section 5 of the standard).

- Priority 1 errors mean that the receiver will not be able to lock the signal.
- Priority 2 errors imply that quality of video and audio may be impaired.
- Priority 3 errors imply a problem in the electronic program guide.



Figure 3.5. TR 101 290 table tabulating several corrupt TS packets as well as lost or missing packets.

This parameter set was chosen in an attempt to quickly determine the QoS of a live TV signal. From a troubleshooting and diagnostics point of view, the most critical of these tests are Sync Byte Error, Continuity Counter Error, and Transport Error Indicator Flag. Any errors found in these three categories usually means that something very bad is happening in the transmission of the stream, or possibly in the building or multiplexing of the stream. The TR 101 290 test parameters are a great way to quickly get an idea of the health of the transport stream and its audio and video elements, but some of the other tests are often misleading. Most important are the timing tests of the MPEG-2 Program Association and Program Map Tables (PAT and PMT) tests. Without these two tables, a normal digital TV or set top box would fail to decode a program. Figure 3.5 shows an example table from the MTS4000 MPEG Analyzer with the TR 101 290 Priority 1 and 2 tests. The green LEDs show a perfect error-free history, and the amber LEDs and their counters show a series of impairments in the past. The 2.1 Transport error means the RF demodulator failed to recover a transport packet (357 times), and therefore, introduced one or more bit-errors in each of the 357 transport packets. The 1.4 Continuity Counter errors (453 times) means that 453 transport packets were lost from the stream, or possibly the counter was corrupt during a 2.1 Transport error. Errors of this type and magnitude would normally cause video and audio artifacts or impairments.

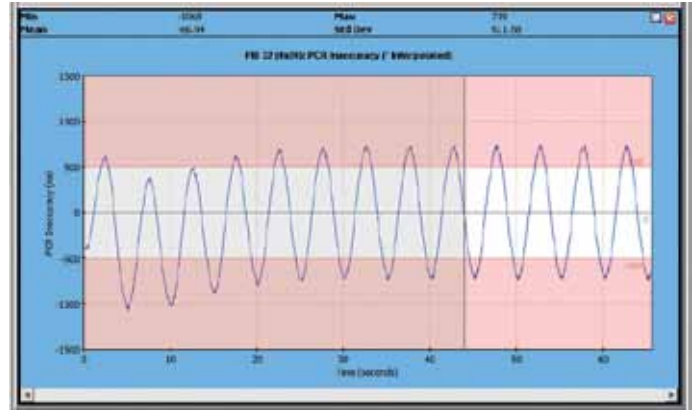


Figure 3.6. TR 101 290 PCR Inaccuracy test.

It is true that PAT and PMT are needed (Priority 1), and a minimal arrival time interval requirement is good, but if the tables arrive just one millisecond late, then the TR 101 290 display goes from green to red. This is considered to be a critical error, even though the extra millisecond of latency in the table is never noticed by the TV, set top box, or the TV viewer. For this reason, some interpretation of the TR 101 290 results are needed. Relying solely on TR 101 290 can get you into trouble. It is common practice for many network operators to increase the threshold of the arrival time intervals to allow for some deviation, while still testing that the tables are arriving at some rate or interval. Keeping the interval times tight will ensure end users can change channels quickly, but TR 101 290 might falsely flag a problem when a stream or program has a slight delay in parts of its electronic program guide. The same goes for many of the program clock reference (PCR) measurements in priority 2. A deviation of a few percent in the interval will not make a difference to virtually every TV or set top box, but this deviation will cause TR 101 290 to change from green to red. Figure 3.6 shows a graph where the PCR Accuracy is out of specification by a few nanoseconds (deviations outside of the white band).



Figure 3.7. Tracking IP packet gap, packet loss, packet rate, etc.




Figure 3.8. High resolution histogram of selected IP stream.




Figure 3.9. Plotting jitter trend up to 7 days.

Given that most facilities today route MPEG over IP, it may become necessary to diagnose transports in an IP network as well. All of the same TR 101 290 tests can be applied to transports over IP, but it is critical to pay attention to massive latency issues between IP packets. RF modulators that use IP as an input will require consistent arrival times in order to keep the buffers from overflowing or underflowing. Therefore, measuring the gap between any two IP packets can be critical. Figure 3.7 shows an MPEG Analyzer measuring the Packet Interval Gap on several MPEG transports, all in parallel on a GigE link.

A high resolution histogram plot is also available for detailed jitter measurements on IP signals. Figure 3.8 shows a mean gap of about 543 μ s, along with a small grouping around 300 μ s and 900 μ s.

Tracking jitter or packet gaps over time can be very powerful in proving the maximum gap found over a 7-day period. Figure 3.9 shows the same 543 μ s average gap (dark blue line), but there are rare occasions where the gap between two IP packets is up to 16.7 ms (gray envelope).

To measure the QoS of the MPEG-2 transports, TR 101 290 will suffice. Determining the quality of the audio and video (QoE) will require much more processing than simply looking at MPEG transport packet headers and packet gaps.

Given that compressed TV is not linear like uncompressed audio, a bad bit landing in the middle of an MPEG video I-frame is much worse than a bad bit landing in the corner of a B-frame. The reason that this issue is not considered linear is because an error in an I-frame may reside on the TV screen for an entire Group of Pictures (about 500 ms). With a bad B-frame picture, it only stays on the TV screen for one frame (e.g., 1/30th or 1/25th of a second). Also, a bad bit could cause an entire slice or rows of pixels to turn green. Or maybe the bad bit only causes a signal pixel to change color. Therefore, unlike in telecom testing where BER was an adequate measurement, in digital TV, we need to measure BER, and also keep track of impairments within each audio and video element.



Figure 3.10. Low blockiness from satellite distribution feed.

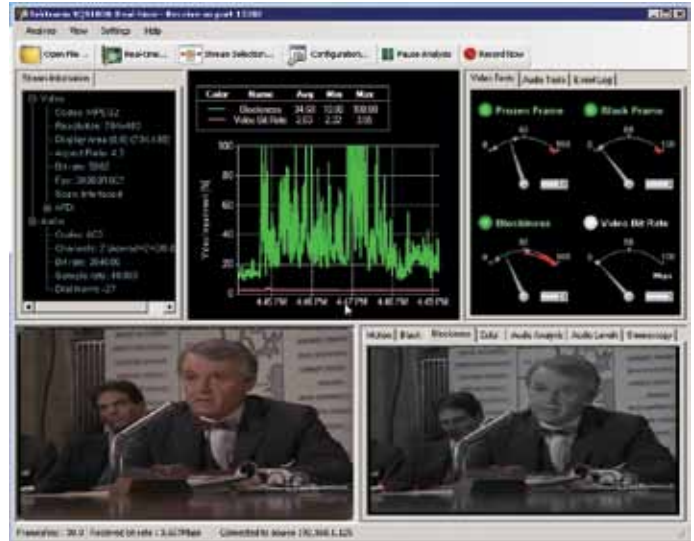


Figure 3.11. Transcoded content with significantly higher blockiness (green plot) over the same timeframe.

Troubleshooting QoE

QoE is a metric that relies much more upon our ability to notice a problem in the midst of a TV program rather than a single bit error. For instance, if a single bit error affected the far corner of a B-frame, then that bit error would have a much lower impact to QoE than if the error had landed in the middle of an I-frame. Subsequently, if the single error in the I-frame ended up corrupting an entire slice of 16x16 pixel blocks, then the impact to QoE would be huge. Therefore, to be able to measure the impact of an error upon a TV program, it is critical to know exactly where the error occurred. Another issue is that the transport may be perfectly error free (no syntax or semantic errors), but the video is objectionable because there is not enough bandwidth to clearly portray the high-motion video scene. In this case, there are no protocol errors, but the picture frame or video sequence is made up

of large 16x16 solid squares rather than a clean picture. Viewers often refer to this as a blocky video problem. Figure 3.10 shows a low blockiness measurement on a TV program from satellite distribution at 4.1 Mbps. Figure 3.11 shows the same TV program over cable (256QAM), but only after it has been transcoded (or rate-shaped, clamped, etc.) to a lower rate of 3.6 Mbps. The resulting changes cause significantly more blockiness artifacts over the same period of time. In this case, the QoE has deteriorated enough that viewers would recognize the blockiness issues in the video.

Another video QoE issue occurs when the same video frame occurs repeatedly for a long period of time (frozen or black frames). Obviously, video frames repeat when there is no activity, but at some point (e.g., 2 minutes), an alarm is needed to alert the operator if the frozen frames are on purpose (for example an all black scene for artistic effect), or from a broken link or piece of equipment.

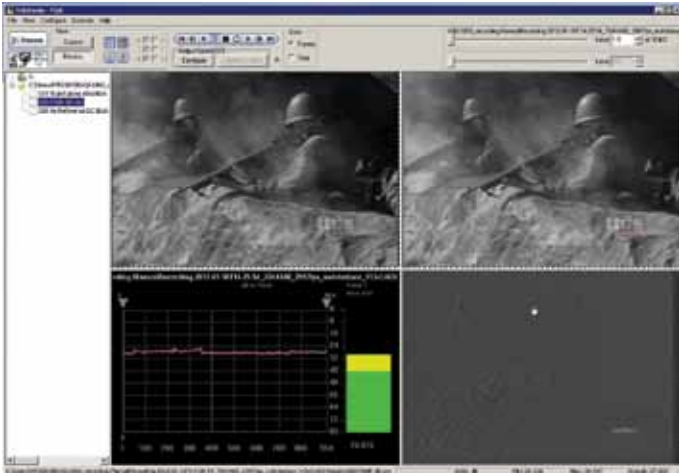


Figure 3.12. The Picture Quality Analyzer performs direct comparisons of video frames before and after a live transcode.

When a reference signal is available for direct comparison as in the case shown above from satellite and cable, an objective picture quality rating can be made using a Picture Quality Analyzer. In this case, a stream is recorded from both satellite and cable, and then the Picture Quality Analyzer will convert to YUV, perform spatial and temporal alignments, and then measure the difference between the two video clips. The basic Peak Signal to Noise Ratio (PSNR) measurement is shown, but the Picture Quality Analyzer also performs many more important tests such as Picture Quality Rating, Attention model, etc. Figure 3.12 shows the Picture Quality Analyzer making a frame by frame PSNR comparison between every frame of the satellite and cable video clips.

The Picture Quality Analyzer has been pitted directly against human view trials in a variety of studies. The IRT study¹¹ (Institut fuer Rundfunktechnik in Munich, Germany) showed extremely high correlation between the PQR results and human viewer results using ITU-R_BT.500-11¹² testing.

Audio is another element of QoE as it occasionally increases to an unbearable level, or maybe drops out altogether over a long period of time. To accurately track these problems, loudness measurements have been defined for instantaneous peaks (ITU-R BS.1770/1771¹³), as well as short term and long term filters for ATSC¹⁴, ITU, and EBU¹⁵. These measurement algorithms allow test equipment to track audio over long periods of time and trigger on deviations outside allowable levels. One method of tracking this is to use a DialNorm level as a target for normal dialog levels. Averaged or filtered levels are allowed to rise or fall within a few dB of this set level. Many governments have adopted this requirement and now impose



Figure 3.13. Satellite ingest with audio loudness at 2 dB quieter than DialNorm.



Figure 3.14. Cable broadcast with audio at 2 dB louder than DialNorm. This is a 4 dB increase over the satellite feed.

fines on broadcast operators who deviate beyond the agreed limits. With this agreement in place, TV viewers should now be able to switch from channel to channel without having to adjust the audio levels. The same goes for the commercials or ad-inserts between programming. The loudness of commercials must also stay within the agreed upon levels. As an example of audio levels being altered, Figure 3.13 shows the satellite distribution signal at one audio level (-29 LKFS¹³) while its cable-broadcast counterpart (Figure 3.14) is at a different audio level (-25 LKFS). Both signals use a DialNorm reference level of -27 LKFS.



Figure 3.15. Significant loss of RF or IP packets cause blockiness and slice errors.

How are QoS and QoE Interrelated?

Both QoS and QoE are interrelated, but their relationships are not equal. Case in point:

- 1) A good QoE usually means a good QoS.
- 2) But a good QoS does not always mean you will have a good QoE. For example, Encoder bandwidth starvation or over-compression.
- 3) A bad QoS, or at least frequently occurring QoS issues usually lend toward a poor QoE. With enough bad bits randomly distributed into all video and audio elements, poor quality of video and audio is bound to occur (see Figure 3.15).
- 4) A bad QoE is not always related to a bad QoS, but it is always worth validating the QoS performance. Poor QoE can come from misbehaving ad-inserted and any other equipment (not related to TR 101 290). Poor QoE can come from over compression (this is perfectly legal). Poor QoE can come from dropped IP packet, corrupt RF packets, or other transmission problems.

Therefore, it is common to measure both QoE and QoS in order to quickly identify the root cause of a problem and in turn correct the issue.

One problem with QoS issues is that they can be easily hidden or masked when the program is handed off from one network to another (e.g., dropped video packets). Once the program is decoded back to baseband for manipulation and re-encoding, all previous QoS errors are lost forever. After encoding a second time, the transport and video protocol is now error free, but visually, there may be large slice errors seen in the middle of the video frames. At this point, then only thing available to catch these artifacts is to decode the video and

audio content and look for anomalies. A common problem in sporting events is when the remote transmission link is lost for a fraction of a second. The video breaks apart, but the receiver is able to catch this QoS issue due to Transport Error Indicator flags (TR 101 290 Priority 2.1 – Also bit #9 in each transport packet). Although, if the decoded and blocky or broken video is passed on to the network as a re-encoded program, then the program will often have a perfect QoS, as well as its compressed video syntax being perfect due to its decode and then second encode process. The video may still look like garbage (as seen in Figure 3.15), but according to QoS and video syntax, the results are error free. To find an artifact, we must actually look at the decoded picture content and determine if the frames contain objectionable artifacts (blockiness, slice errors, frozen frames, etc.).

If we have a high BER or poor QoS, then it is inevitable that the QoE will deteriorate also (as in Figure 3.15). In this case, we do not care which frames the errors are landing in because the rate is high enough to land in every frame. Therefore, the QoE rating becomes highly objectionable. In this case, the QoS measurement is the key to isolating the problem.

When the QoS is quite good (no transmission errors), then there is a very low correlation between QoS and QoE. In this case, it is more important to focus on the QoE.

There is an important case where the QoS can be perfect, but the QoE is highly objectionable. This is due to the common case with networks running constant bit rate (CBR) video. Normal video with low motion may look decent, or even very good. But the scene occasionally changes from slow moving objects (e.g., talking head) to something with a high rate of change (e.g., sporting event, quick successive scene cuts, etc.). If the encoder does not have enough bandwidth to accommodate the transition to the high-motion events, then the only recourse for the encoder is to start decimating the 8x8 blocks of video. At its worst, the 8x8 blocks may only contain a single value representing the video luminance and color. This scenario is important for QoS and QoE testing because the end-user notices a very objectionable TV program segment, although the difference between over compression and near total loss of signal are almost indistinguishable. Figures 10 and 11 show a program at both 4.1 and 3.6 Mbps with the lower rate having much more blockiness due to high motion in a bandwidth starved stream. Figure 3.15 shows blockiness too, but due to a completely different reason (missing packets). Therefore, QoE testing is extremely important, but testing QoS helps to determine the difference between malfunctioning equipment causing a problem verses a poor transmission link.

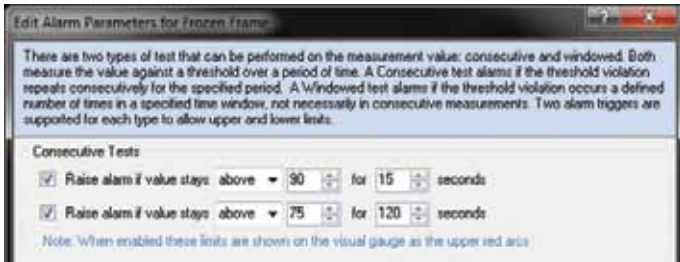


Figure 3.16. Dual trigger levels for Frozen and Black video.

Using QoS and QoE to Identify “Real” Problems: Troubleshooting Case Studies

Most facilities use an ingest, manipulate, and egress architecture. Cable headends use a wide variety of video equipment to capture, manipulate, and ultimately, broadcast. Most all test points would be in the form of QPSK/8PSK, 8VSB, QAM, IP (GigE or 10G), or ASI.

Monitoring QoE aspects on live programming falls into several different audio and video categories.

Video QoE

Video QoE includes monitoring video for frozen frames, black, and blockiness. In the case of frozen frames, a long series of nearly identical frames should trigger a problem since live TV programming is usually made up from moving video. We know that there are many occasions where it is acceptable for a frame to be repeated many times, but at some point, the operator should be alerted to excessive repetitive frames. It could be the sign of a broken encoder or multiplexer, or any of a number of pieces of hardware. Figure 3.16 shows the QoE software setup for Frozen Frames that will be triggered if

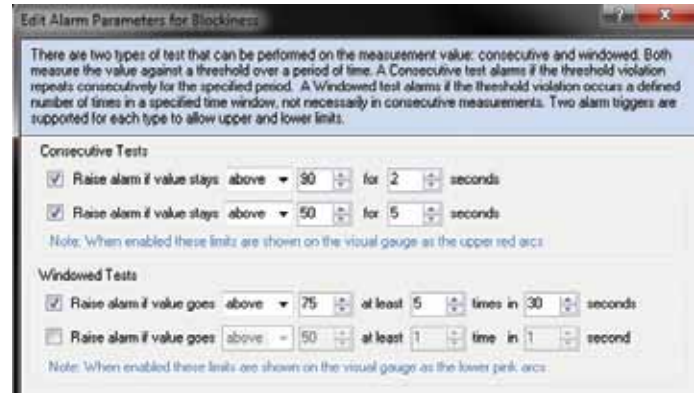


Figure 3.17. Dual trigger levels and gating for Blockiness.

90% of the frame is frozen for more than 15 seconds, or when 50% of the frame is frozen for more than 120 seconds. The requirements are similar for excessively long periods of black. Blockiness is a little different in that it can occur in a small portion of the picture, or over the entire picture. Blockiness will often occur during short scenes where the video content is moving too quickly for the encoder to faithfully compress all of the details. Therefore, the encoder tends to throw away high frequency details rather than failing completely. In this case, the monitor must be set to trigger once the blockiness level has crossed a threshold, and then maintain that level for a significant amount of time. Otherwise, the monitor would trigger an alert every time the scene was overly blocky, even for one single frame. Figure 3.17 shows the VQS1000 QoE software setup for Blockiness triggering when

- level of 90% is reached for over 2 seconds, or
- level of 50% is reached for over 5 seconds, or
- level of 75% is reached at least 5 times within 30 seconds.

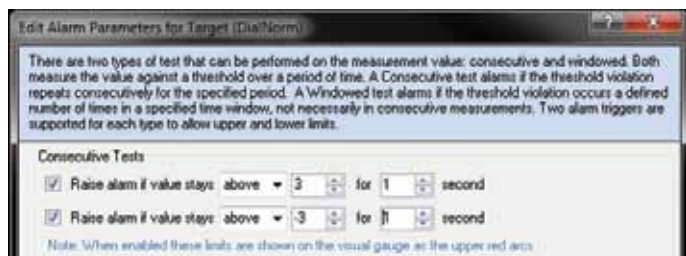


Figure 3.18. Target DialNorm settings.

Audio QoE

Audio QoE includes monitoring audio for loudness levels. With such a high emphasis on audio levels today, it has become critical to measure overall loudness levels according to the new guidelines. Just as with video QoE monitoring, triggers for levels, deviations, and durations are important to minimize the many false positives that can occur. Figure 3.18 shows audio loudness limits must be within 3 dB of DialNorm.

Troubleshooting Over-Compression

The QoE software allows for real-time monitoring of multiple RF and IP signals in both the transport layer as well as the video and audio layer. When blockiness is found to be excessive on in a TV program, the MPEG Analyzer can measure the broadcast program as well as the ingest distribution feed.

When the QoE of a TV program is called into question, the MPEG Analyzer with its QoE software can quantify the magnitude of the audio and video issues. Once the levels are proved to be unacceptable (as seen in Figure 3.11), the MPEG Analyzer can quantify the ingest program (as seen in Figure 3.10), and then a comparison can be made between the two programs. In the example here, the cable broadcast program has a stable QoE or blockiness as long as the content does not contain too much action. Once the scene changes quickly, the blockiness rating spikes dramatically. In comparison, the ingest content is much less blocky and only varies slightly when the scenes change quickly. It can be noted that the ingest content is coming into the facility at about 4.1 Mbps whereas the broadcast program is leaving at only 3.6 Mbps. Therefore, the blockiness issue can be identified as a result of decoding and re-encoding without providing sufficient bandwidth to maintain the quality of the original picture.




Figure 3.19. Terrestrial encoder generating negative DTS-PCR delta values (should be impossible).

Troubleshooting Interoperability

The MPEG Analyzer can also be used to identify non-compliance issues when televisions and set top boxes begin to react abnormally to specific programming. In a local broadcast case, several viewers complained about inaudible audio on the SAP channel (audio sounding like motorboats). Other viewers had no problem at all on the SAP channel. It sounded like an interoperability issue, so the MPEG Analyzer was tuned to the local off-air channel to look at its TR 101 290 results. The TR 101 290 tests, and virtually every other test came up clean except for one. The MPEG Analyzer can graph the audio or video Decode Time Stamp (DTS) value as it arrives, against its current PCR value. This delay is always a positive number (representing buffer delay time) and by definition it is required to be between zero and one second. The newly added SAP audio had its DTS-PCR delta ranging above and below zero, which is by definition, never allowed to happen. Figure 3.19 shows the local broadcaster SAP audio delay issue.

This same terrestrial broadcast feed was also sent to both local cable companies for redistribution.




Figure 3.20. Cable headend encoder generating negative DTS-PCR delta values (should be impossible).

Another similar audio example was found at the cable headend where the audio buffer delay was a negative value, which is impossible for a set top box to support. The MPEG Analyzer once again showed the values going below zero. Figure 3.20 shows the same audio problem, but at a different facility.

In both cases, reconfiguring the encoder fixed the problem.

The MPEG Analyzer with its unique DTS-PCR measurement was able to quickly pinpoint the problem to the recent configuration change to the encoder, and thus take the blame off the small set of failing decoders throughout the city.

Troubleshooting Transport Streams with CaptureVu

A key part of delivering a quality experience is finding the root cause of problems in the transport stream. It's no longer good enough just to spot a problem, reboot the set-top box and hope that it goes away. The MPEG Analyzer can uniquely go as deep as is needed into the transport and elementary streams to track down sources of picture anomalies, like discovering that a closed caption stream contained too much information, causing set-top box buffers to overflow and precipitate automatic reboots.




Figure 3.21. MTS4000 Transport Stream Compliance Analyzer: Trigger and capture enabled for missing encryption key packet.

While the MPEG Analyzer provides real-time analysis, the instrument's ability to capture events for deep analysis is critical to identifying the root cause of problems. Especially useful are the pre-triggers in CaptureVu that not only capture an error, but also provide information leading up to the error.

This last troubleshooting example occurred with a DVR from a major consumer electronics manufacturer that was misbehaving whenever it was tuned to a Switched Digital Video (SDV) channel. The MPEG Analyzer discovered the device was dropping the encryption key whenever a service was added to or removed from the multiplex. Figure 3.21 shows the MPEG Analyzer TSCA setup screen for enabling a trigger on a specific test or missing PID. By documenting the scenario with screen shots and test reports, the DVR manufacturer was contacted, and they provided a software update that solved the problem and eliminated many expensive truck rolls for the Cable Operator.




Figure 3.22. MTS4000 Transport Stream Compliance Analyzer enabling the CaptureVu and Recording for SCTE35 DPI packets.

Troubleshooting Ad Insertion

All broadcast companies fund their businesses based upon selling advertising time between programs. The ads need to be inserted at very specific time, and often timed by SCTE 35¹⁶ Digital Program Insertion (DPI) cueing. If these cue tones do not make it through, then the national ads go through and local revenue is lost. The MPEG Analyzer can trigger on the STCE 35 packets as well as record programming to disk before and after the event (pre-trigger). Figure 3.22 shows the Transport Stream Compliance Analyzer enabling the CaptureVu and Recording feature based upon SCTE 35 packets. The trigger/record function allows the recorded file to be as small as possible.

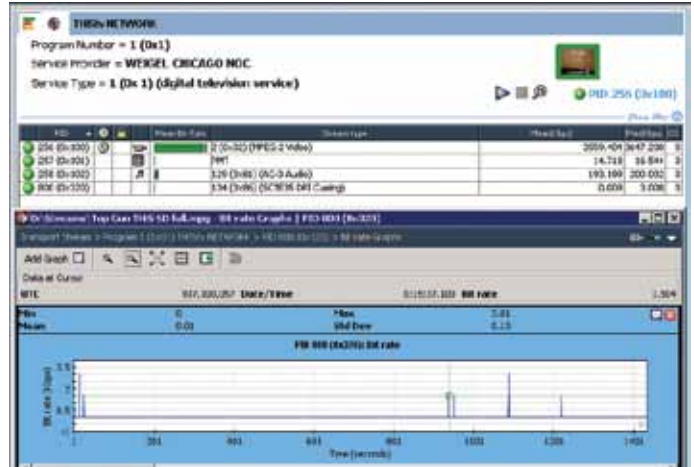


Figure 3.23. PID 800 DPI packets arrive about every 15 minutes.

In the case of much larger recordings, the MPEG Analyzer can record for many hours. In the example shown in Figure 3.23 running at 4 Mbps, the MPEG Analyzer could hold over 11 days of continuous programming. Figure 3.23 also shows the PID 800 DPI packets arriving around 15 minutes apart, as you would expect for local ad-inserts.

With these SCTE 35 DPI tools, troubleshooting and debugging missing DPI packets becomes relatively simple.

Summary

Troubleshooting video signals in a broadcast or cable facility requires the use of an MPEG Analyzer that provides multiple input signals all running in parallel. The minimum set of test points should be at ingress, ASI/IP switch, and egress. It is critical to measure QoE at both egress and ingest, but in the case of audio and video problems, QoS testing at multiple points in the facility may be needed to pinpoint the source of the impairments. Critical capabilities in an MPEG Analyzer should include:

- Transport stream generation, modification and analysis
 - Real and deferred time
 - Including MPEG-2, H.264 & MPEG-4 AAC , H.265/HEVC
 - Multiplexer
 - Automatic error triggered recording and capture
- A range of physical interfaces
 - Multi-port ASI
 - RF (DVB-S2, 8-VSB & QAM-B)
 - Video over IP analysis and generation (1Gbps and 10Gbps)
- Comprehensive suite of software tools for analyzing all layers of video
 - High accuracy RF layer analysis
 - Transport, Program & Elementary Streams
 - Video & Audio Quality of Experience (QoE)
 - Picture Quality

A facility equipped with an analyzer having the feature set listed above should be able to resolve problems in minutes rather than hours or days using alternative methods.

References

1. ISO/IEC 13818 MPEG-2 System, Video, and Audio
2. ISO/IEC 144961-10 MPEG-4 Part 10 Video, also known as H.264
3. EN 300 421 DVB-S QPSK Satellite
4. EN 302 307 DVB-S2 Second Generation Satellite
5. ITU-T J.83B QAM-B, also known as SCTE 07
6. ATSC A/53 Digital Television Standard
7. ETSI TR 101 290 DVB Measurement Guidelines
8. SCTE 07 Digital Transmission for Cable Television
9. EN 50083-9 DVB-PI ASI and parallel
10. RFC 0791 Internet Protocol
11. IRT / Tektronix Comparison of Subjective and Objective Picture Quality Rating of MPEG-2 Video Report
12. ITU-R_BT.500-11 Methodology for the subjective assessment of the quality of television pictures
13. ITU-R BS.1770/1771 Algorithms to measure audio program loudness and true-peak audio level
14. ATSC A/85 Techniques for Establishing and Maintaining Audio Loudness for Digital Television
15. European Broadcast Union R 128 Loudness normalization and permitted maximum level of audio signals
16. SCTE 35 Digital Program Insertion Cueing Message for Cable

Contact Tektronix:

- ASEAN / Australia** (65) 6356 3900
- Austria*** 00800 2255 4835
- Balkans, Israel, South Africa and other ISE Countries** +41 52 675 3777
- Belgium*** 00800 2255 4835
- Brazil** +55 (11) 3759 7627
- Canada** 1 (800) 833-9200
- Central East Europe and the Baltics** +41 52 675 3777
- Central Europe & Greece** +41 52 675 3777
- Denmark** +45 80 88 1401
- Finland** +41 52 675 3777
- France*** 00800 2255 4835
- Germany*** 00800 2255 4835
- Hong Kong** 400-820-5835
- Ireland*** 00800 2255 4835
- India** +91-80-30792600
- Italy*** 00800 2255 4835
- Japan** 0120-441-046
- Luxembourg** +41 52 675 3777
- Macau** 400-820-5835
- Mongolia** 400-820-5835
- Mexico, Central/South America & Caribbean** 52 (55) 56 04 50 90
- Middle East, Asia and North Africa** +41 52 675 3777
- The Netherlands*** 00800 2255 4835
- Norway** 800 16098
- People's Republic of China** 400-820-5835
- Poland** +41 52 675 3777
- Portugal** 80 08 12370
- Puerto Rico** 1 (800) 833-9200
- Republic of Korea** +822-6917-5000
- Russia** +7 495 664 75 64
- Singapore** +65 6356-3900
- South Africa** +27 11 206 8360
- Spain*** 00800 2255 4835
- Sweden*** 00800 2255 4835
- Switzerland*** 00800 2255 4835
- Taiwan** 886-2-2656-6688
- United Kingdom*** 00800 2255 4835
- USA** 1 (800) 833-9200

* If the European phone number above is not accessible,
please call +41 52 675 3777

Contact List Updated June 2013

For Further Information

Tektronix maintains a comprehensive, constantly expanding collection of application notes, technical briefs and other resources to help engineers working on the cutting edge of technology. Please visit www.tektronix.com



Copyright © 2014, Tektronix. All rights reserved. Tektronix products are covered by U.S. and foreign patents, issued and pending. Information in this publication supersedes that in all previously published material. Specification and price change privileges reserved. TEKTRONIX and TEK are registered trademarks of Tektronix, Inc. All other trade names referenced are the service marks, trademarks or registered trademarks of their respective companies.

10/14 EA/WWW

25W-30904-0

