



Case Study – Monitoring for Reliability

*Bill Reckwerdt, CTO
Video Clarity, Inc.*

Digital video is everywhere. While pristine audio & video can be stored and forwarded in non-real-time, the A/V must be compressed before sending them across a network to the end user. During this transformation, many problems can occur, which affect the video quality. In order to ensure reliable high-quality transmission, you first must assess your overall system performance. An effective A/V quality metric is needed to monitor reliability.

This paper is divided into 6 parts:

- 1) Monitoring Basics
- 2) General background on how to setup an effective monitoring solution
- 3) Using the Video Clarity Real Time Quality Monitoring solution
- 4) Defining real-life errors (video only for this paper)
- 5) Defining other reasons to monitor
- 6) Conclusions

Monitoring Basics

So why do we test? Reliability sells they say. While this may be true, the contra positive is definitely true – lack of reliability will result in no sale.

Reliability is an often misunderstood concept because people talk about it in too narrow a view. Each component in the chain can be reliable, and the entire system can be unreliable. Reliability must be built at the system-level across multiple equipment manufacturers using real-life data.

Monitoring Solutions

In order to monitor the quality, the first task is to read (capture) the data flowing across the network.

Once captured, many problems can be classified

- 1) The video is black
- 2) The audio is silent
- 3) The video and/or audio is distorted
- 4) The video and audio are out of sync with reference to each other
- 5) The ancillary data (closed captioning, subtitles, etc.) is not intact or timed properly

If you can determine the above by simply analyzing the data flowing across the network, then this is referred to as “no reference” quality monitoring. No reference is the simplest form of monitoring to understand, but it is easily fooled. If the original has silence or black frames or if the director applied special effects to distort the image, then false alerts will be generated.

A cleaner solution is to use the original video before it is transmitted as a reference and compare it against the captured data flowing across the network. This is referred to as “full reference” quality monitoring. Full reference will not generate false alerts, but sometimes access to the original content is not feasible.

Hybrid solutions that use some of the original video data exist called “reduced reference” quality monitoring, but these have a tendency of generating false alerts and suffer from the feasibility issue in accessing the original content.

Who Monitors


To classify and fix the errors, we must discuss who the user is. In general, 2 different groups use monitoring equipment:

- 1) Network Operations (TV broadcaster originators like BBC, NBC, Discovery)
- 2) Engineering & QA

Both groups send reference content through their processing chain and check the output. The content can either be special content that stresses the system or standard content. An automated system that alerts when the quality drops would greatly enhance the testing.

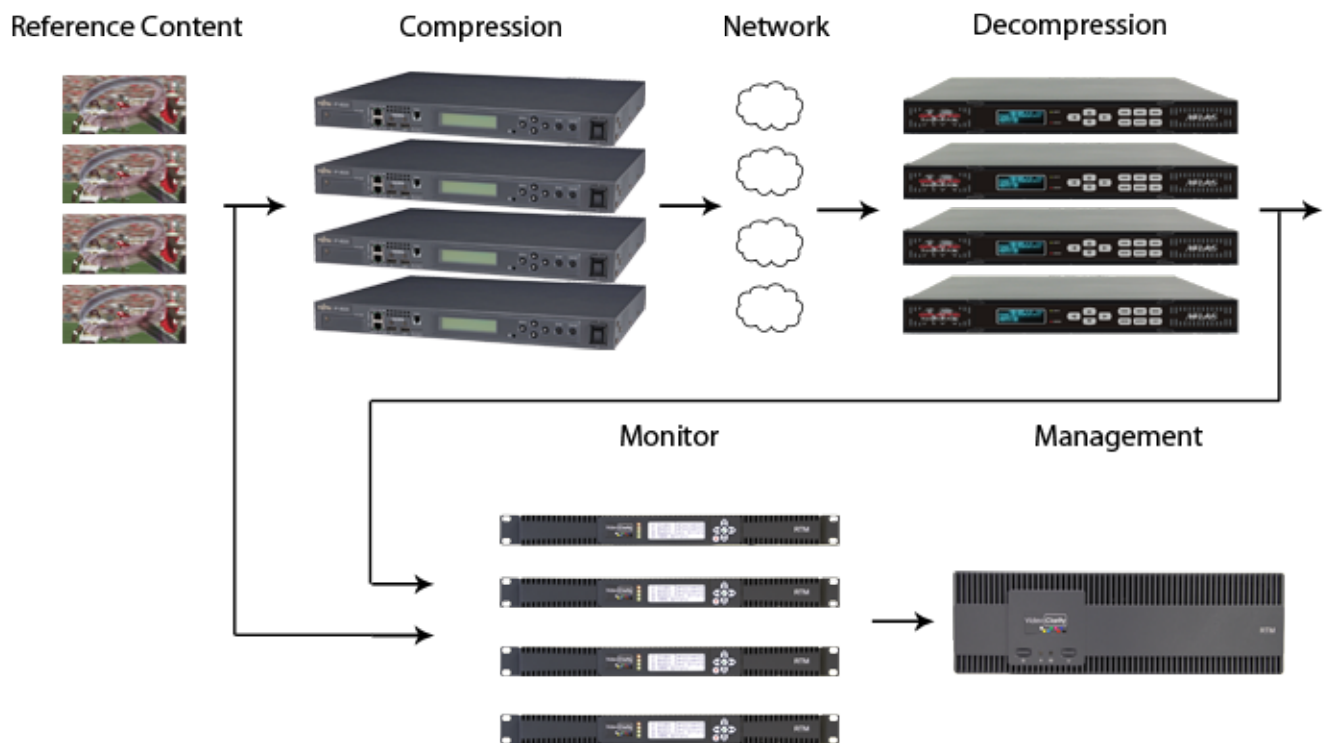
They have one key difference. The Engineering & QA groups can stop the monitoring at any point to review and classify the alerts. The standard workflow would look like the following.

Figure 1: Engineering / QA Workflow



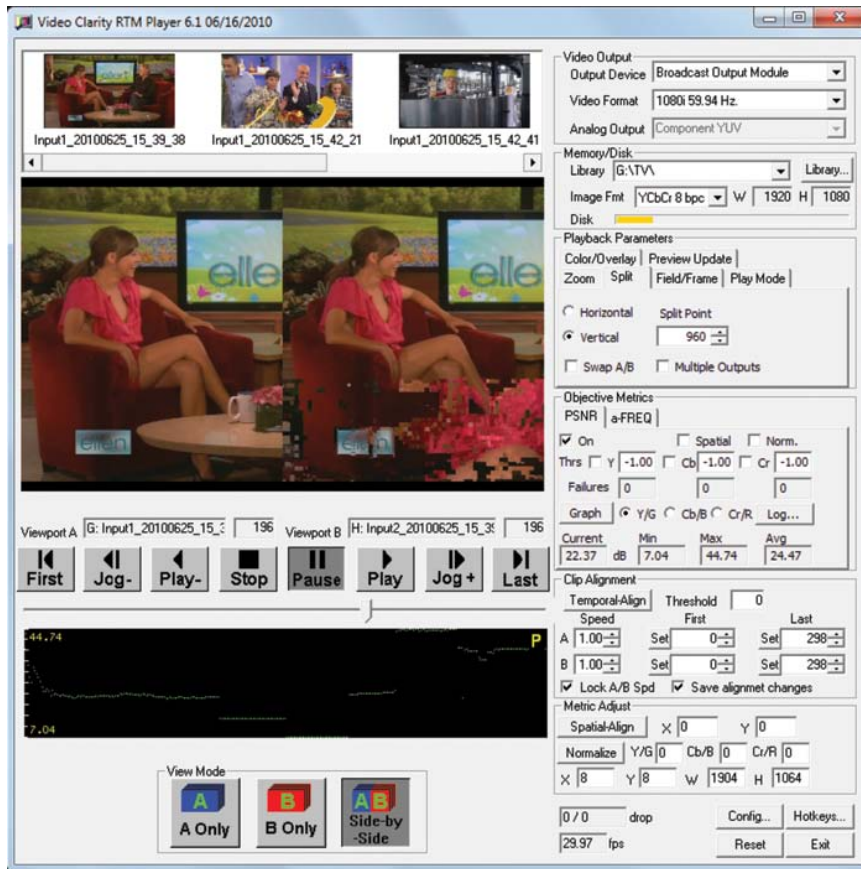
The operation group runs multiple channels, uninterrupted 24x7. To review and classify the alert, they must offload the alerts to an external network management device. The standard workflow would like the following.

Figure 2: Operation Workflow



In either case, the review process includes looking at the reference and processed content in many viewing modes with or without perceptual, objective metrics. An example of such a viewing tool is the Video Clarity RTM Player shown below.

Figure 3: RTM Player showing side-by-side viewing



Video Clarity's RTM

Video Clarity, using the technology insight from its ClearView line of video quality analyzers, built RTM as a full reference quality monitor.

The RTM operation is straight forward. The original reference and processed content are read through SDI (HD/SD with embedded audio), the video data is temporal and spatially aligned, and the audio data is separately aligned. After the alignment phase, RTM

- 1) measures the video quality,
- 2) measures the audio quality,
- 3) checks the ancillary data integrity, and
- 4) calculates the audio and video synchronization offset.

The user sets the quality threshold levels based on their experience, but standard threshold levels could be

- 1) if the quality drops below 30 for 5 consecutive frames
- 2) if the quality drops below 20 for 10 frames over the span of 100 frames

If the threshold levels are exceeded,

- 1) an alert is generated and logged,
- 2) the associated audio or video data is saved (both the reference and processed), and
- 3) RTM continues to look for the next threshold level failure.

By saving the original data around the error, the user can

- 1) review the alerts and classify them (operation, engineering, and QA), and
- 2) add the reference data to the set of regression tests (engineering or QA)

RTM can be operated from a script, SNMP, or the GUI shown below.

When using the GUI, the following items are displayed:

- 1) the reference and processed videos,
- 2) the reference and processed audio levels,
- 3) the audio and video quality scores over the last 30 seconds,
- 4) the number of alerts,
- 5) the video and audio offset (lip-sync),
- 6) the A/V delay between reference and processed,
- 7) the current video quality score and the threshold values, and
- 8) the current audio quality score and the threshold values

All of this data is also present in the log files with a few additions:

- 1) the periodic average audio and video quality scores are since the test began
- 2) the periodic average audio and video synchronization (lip-sync) values are since the test began, and
- 3) pointers to the audio and video sequences when an error occurs so you can drag & drop to see the errors using either RTM Player or ClearView.

Figure 4: RTM's User Interface and Product Photo



RTM - 1RU



Monitoring Real-Life Errors

Errors will occur. Simple errors are easily found and corrected, but some happen infrequently and/or in the presence of special conditions. For example we will outline 3 types of errors.

Reference Content Caused the Error

At times the reference content can have break-ups in it, and this can cause the processed content to further break-up or completely freeze.

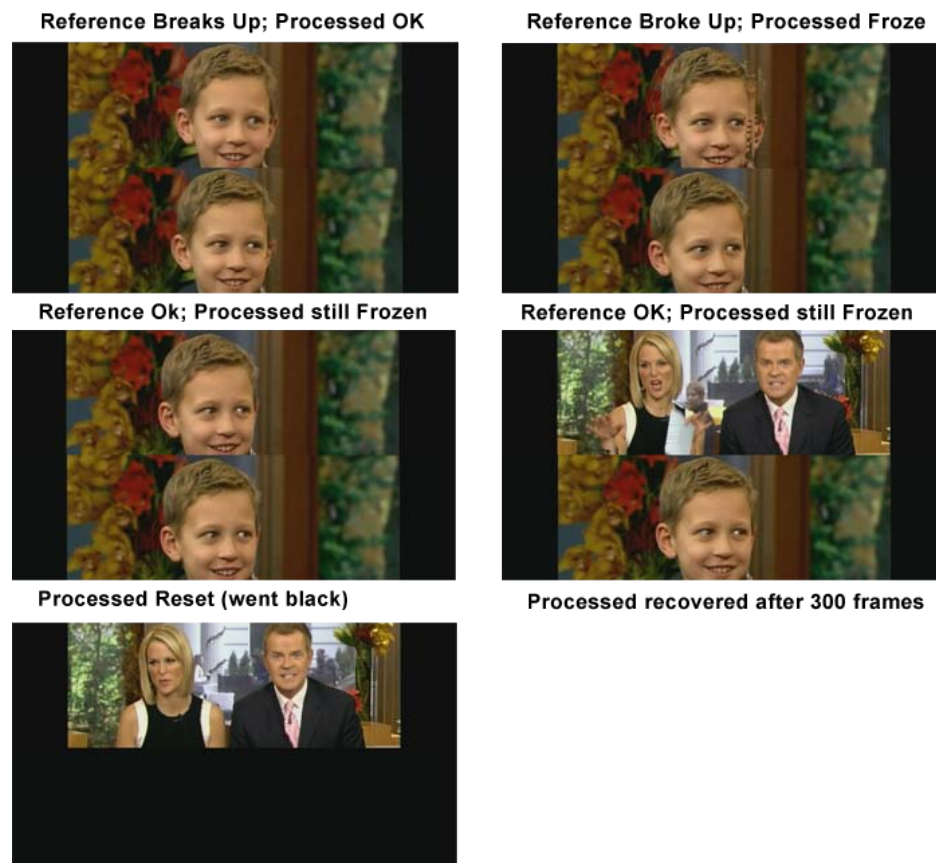
In the following examples, the reference content had a momentary breakup for approximately 10 frames, and this caused the processing units downstream to completely freeze. If you were not watching carefully, you may have thought that the processing units created the error do to network error, but this was not the case. The reference content caused the error.

RTM was able to find this problem. By saving the reference and processed content and alerting, the problem was further diagnosed by playing the content in over-under mode using RTM player. RTM player displayed the reference content on top and the processed on the bottom and some of the frames are illustrated below.

A no reference monitor would have found this problem, but the cause would not have been determined. If a person watched this error, they may or may not have noticed the 10-frame breakup in the reference.

Since the reference material is saved, it can be used for regression testing after engineering has worked on possible fix for the processing chain. Further, if the reference content came from a content originator (NBC, Discovery, etc.), it can be sent back to them so that they can fix the problem on their end.

Figure 5: Reference Caused Breakup (Reference on Top)



Processing Content Caused the Error

In the following sequence, processing the content caused a problem. By putting traffic on the network some of the compressed packets were dropped before they made it to the decoder (set-top box). The set-top box went into error recovery mode, but could not fix the problem. Finally, the network needed to be reset. This was the case of a MSO (re-broadcaster) having problems with its internal network.

RTM was able to find this problem. By saving the reference and processed content and alerting, the problem was further diagnosed by playing the content in side-by-side mode using RTM player. RTM player displayed the reference content on left and the processed on the right and some of the frames are illustrated below.

A no reference monitor would have found this problem, but the cause would not have been determined. If a person watched this error, they probably would have correctly assumed that this was a network problem while processing the data.

Since the reference material is saved, it can be used for regression testing after engineering has worked on possible fix for the processing chain.

Figure 6: Network Processing Breakup (Reference on left)



Dropping a Frame

The following is a simple commercial. For commercial TV, the broadcaster is paid when most of the commercial is played until completion (something around 95%). In this example, a car dealer showed the same vehicle (a 2010 mustang) in 2 different colors. Each vehicle was shown for 10 seconds (20-second commercial).

When the blue car was supposed to be displayed, the black car continued to be displayed. It took approximately 5 seconds for the correct frame to be displayed. In this instance, the broadcaster would not have been paid since the commercial was only correctly displayed for 15 seconds.

RTM was able to detect this error and alert. The operator could have switched to their backup feed and would have been paid.

A no reference monitor would not have alerted on this error since the video looks correct. A person watching this would have detected the error if they were viewing the reference and processed during the 5 seconds that they were different.

Figure 7: Lost Frame (Reference on the left)



Monitoring Other Uses

RTM measures the audio and video quality and writes the average quality in a log file. A long duration test could be setup where you change the compression and/or network parameters periodically during the test. The quality will be measured continuously.

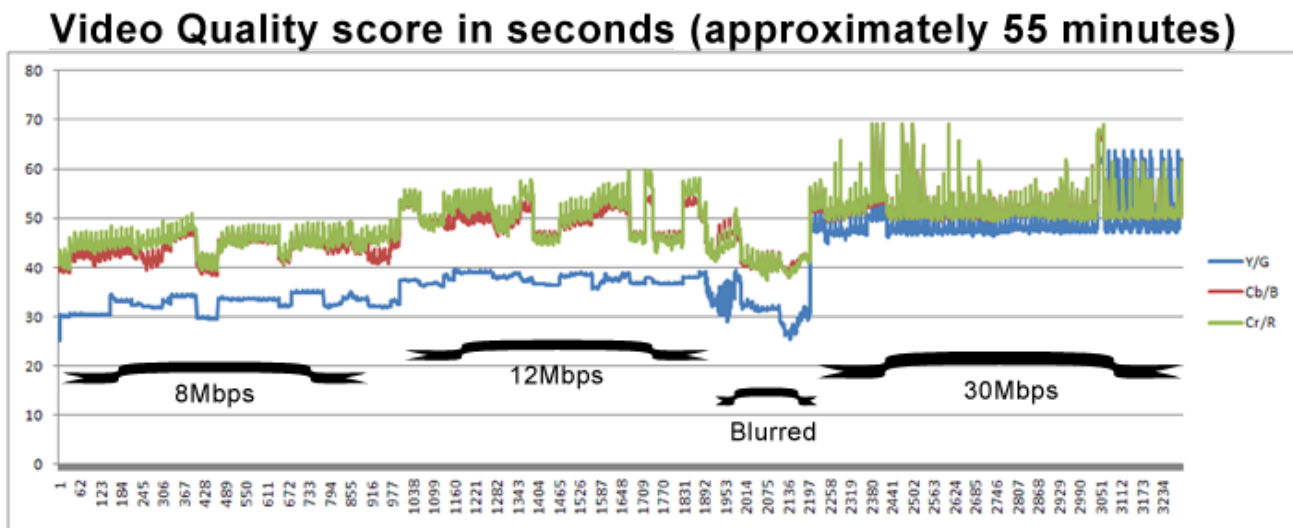
As an example, a 55 minute series of commercials were sent across the Video Clarity network and it was monitored by RTM. The goal was to judge the quality under different conditions.

The processed video quality was changed 4 times during this period:

- 1) 8Mbps for the first 1000 seconds
- 2) 12Mbps for the next 1000 seconds
- 3) 12Mbps with enhanced blur for the next 300 seconds
- 4) 30Mbps for the final 1000 seconds

As suspected, the video quality at 12Mbps is better than 8Mbps, and 30Mbps far exceed either 8Mbps or 12Mbps. When 12Mbps was blurred, the quality dropped below 8Mbps. In some cases, it could occur that the quality difference between 8 and 12Mbps is not discernable. In these cases, setting the compression to 8Mbps will save bandwidth without an perceived video quality degradation.

Figure 8: Video Quality Over Time



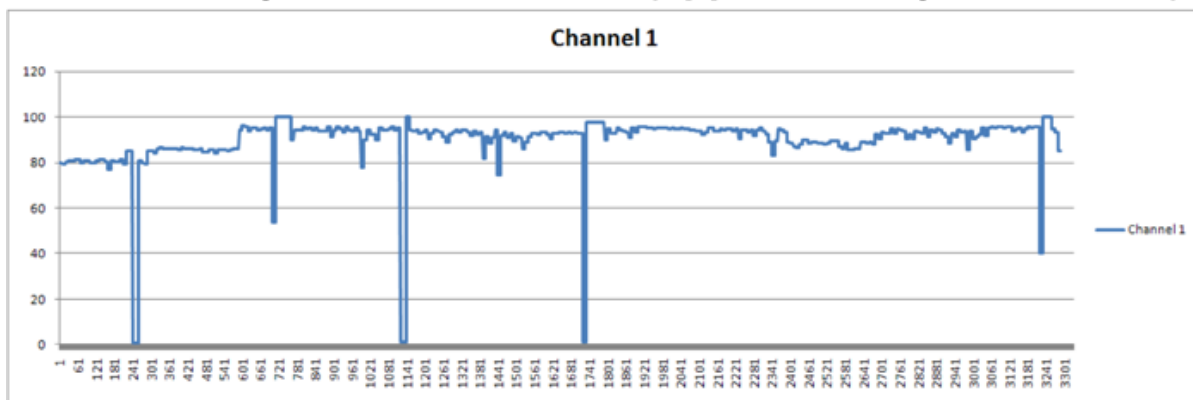
The processed audio quality was kept constant at 224Kbps, and silence was injected 5 times duration the audio processing. The duration and position of the audio silence is listed below:

- 1) 3 seconds around 240 seconds
- 2) 0.5 seconds around 720 seconds
- 3) 3 seconds around 1140 seconds
- 4) 2 seconds around 1740 seconds
- 5) 1 second around 3240 seconds

As expect the audio quality stayed pretty consistent throughout the test. The forced silence created 5 spikes. The spikes reached zero when the duration was greater than 1 second (the periodic average under which the test results were logged – RTM measures audio quality enveloped to the video frame).

Figure 9: Audio Quality with 5 Forced Silences

Audio Quality score in seconds (approximately 55 minutes)



Conclusions

Reliability sells and the only way to ensure reliability is to monitor the end-to-end system's performance.

Video Clarity created RTM as a full reference monitoring solution which alerts on drops in audio or video quality, reports the audio and video quality over time, measures the audio and video offset, measures the lip-sync, checks the VANC data integrity, and saves the audio and video sequences when the quality drops below the defined thresholds.

This allows the user, regardless of whether they are in operations, engineering, or QA, to review the logged errors, classify them, and design and test fixes.

RTM Availability

Video Clarity's products are currently being used by many [broadcasters and equipment manufacturers](#). To get a demonstration, please contact Video Clarity or one of its channel partners (www.videoclarity.com/sales.html) or visit us at one of the shows listed at www.videoclarity.com.

The Author

Bill Reckwerdt has been involved in digital video since the early 90's from digital compression, video on demand, to streaming servers. He received his MS specializing in Behavioral Modeling and Design Automation from the University of Illinois Urbana-Champaign.

He is the VP of Marketing and the CTO for Video Clarity, which makes quantitative, repeatable video quality testing and monitoring tools.