

Innovative Networking/Video Transport: Who's Checking?

Now that We've Moved the Encoder to the Cloud...
Where Does it Make the Most Sense
to Do the Monitoring/Alerting?

Sergio Ammirata, Ph.D.
CTO, DVEO



What We'll Present

- The Problem: Lack of Cloud Vendor Provided Stream Metrics/Analysis in Cloud Based Encoders
- The Solution: A Dedicated IP/TS Probe Image/VM to Place Next to Your Cloud Based Encoder, Wherever That May Be, Whichever Cloud Vendor, Whoever the Encoder Vendor
- Details:
 - What You Can Measure
 - Where You Can Measure and How to Monitor
 - Where You Can Test
 - CAPEX vs. OPEX Strategies



Everybody's Moving Their Encoder to the Cloud...


- Cloud based encoding is part of a group of applications growing so fast they now have their own category name and sub-categories—Cloud-Based Throughput-Oriented Applications, Public and Private. (Source: ABI Research).
- “the use of cloud services to store a large amount of data” is one of the “major factors driving the growth of the video encoder market.” (Source: Markets and Markets).



Cloud Vendors' Default/Built-in Encoders are *Basic*

- They usually aren't meant to do much more than simple rtp to hls
- Typically their analysis tools monitor virtual CPU load and network metrics...
not stream quality metrics, not content errors...

- Connection Count by Protocol
- Bandwidth In
- Bandwidth Out
- Total Connection Count
- Bandwidth In
- Bandwidth Out
- CPU Usage of the (vendor shall remain nameless) Streaming Engine server
- Heap Usage of the (vendor shall remain nameless) Streaming Engine server
- Total Memory Used
- Total Disk Used across all drives
- Cloud Vendor Metrics
- What Could That Miss?
- As Huge a Problem as No Audio!
- With a Probe Window as Below, the Two Green Bars Immediately Indicate Video and Audio are Present



Why Would That Be?

- The Major Cloud Vendors Sell Slices of Tangibles/Measurables: CPU, Storage and Bandwidth and/or Shared Versions of Already Produced Software (MS Office 365, for example)
- In Other Words, Amazon and Microsoft Aren't in the Encoder Business. It's an *Extra* For Them
- (Google Could Probably Own the Cloud Encoding Business Given its YouTube Base... but That Wouldn't Help Them Sell More Ads!)



So, No One-Stop Shopping?

- No, Not Really, But That May Not Be a Bad Thing
 - We Will Outline the Different Types of Vendor Services in a Non-Competitive Way
 - You Can Figure Out Your Company's Needs as a Content Provider *a la Carte*
- What We Can and Will Do Here Today is Show You How You Can *Measure and Monitor* Your Cloud Based Video Streams... in a Manner Consistent With Structured Data

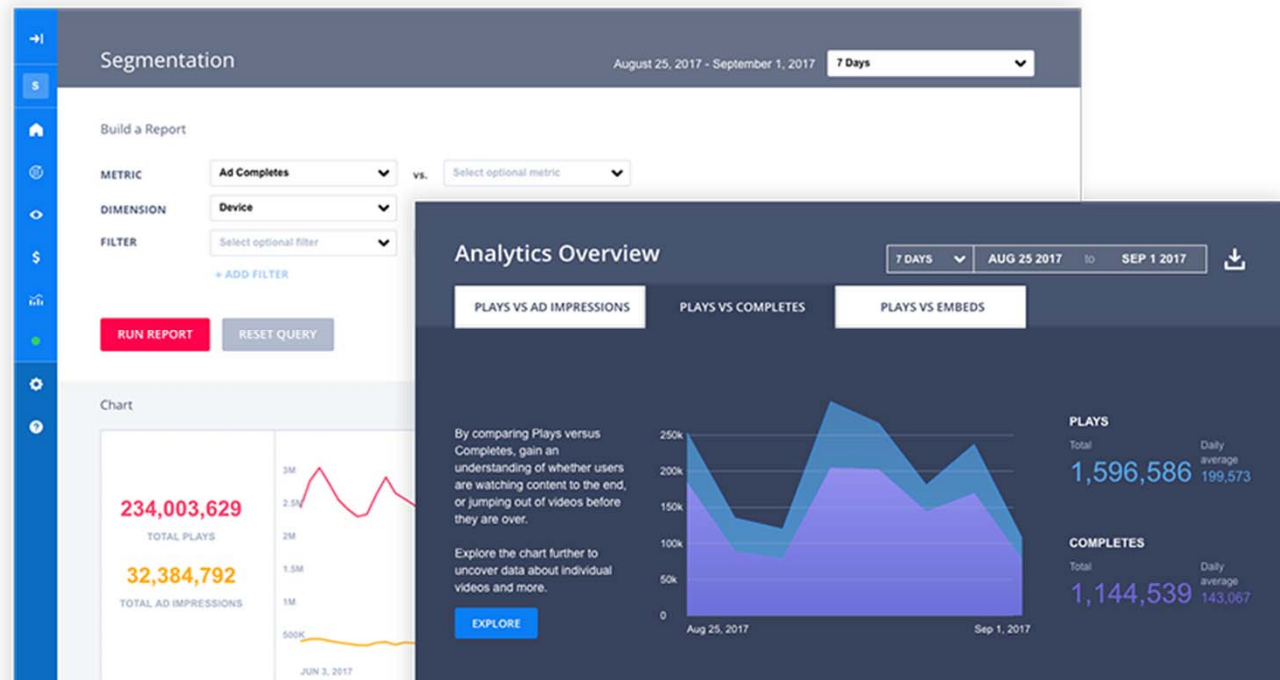


So, Who, If Not the Major Cloud Vendors?

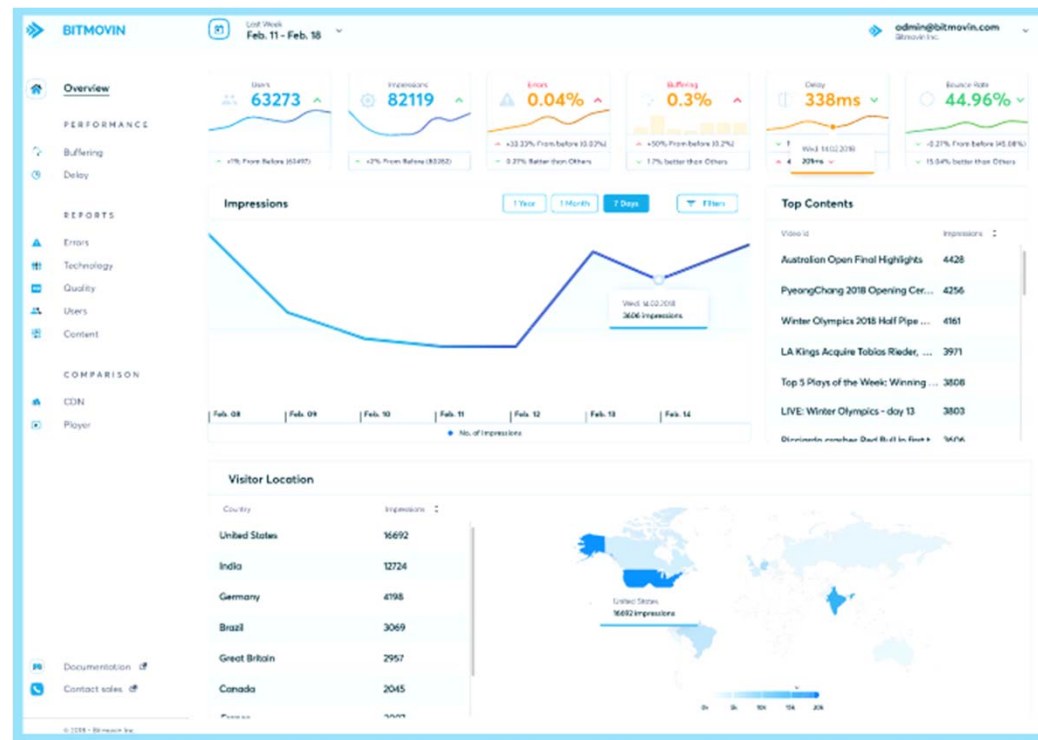
- Cloud Based Services e.g., JWPlayer
 - You won't find an image in the Azure or AWS markets, but they have their own cloud.
 - You contract them as a service not as your cloud based NOC replacement. (*Not that there's anything wrong with that*, as Jerry Seinfeld would say).
 - Interesting approach to analytics: demographic and economic oriented. (Note: we can identify other player based vendors: BitMovin is similar; TheoPlayer a little different)
 - They have marvelous tools that show feedback from the end user players.
 - Typically used for VOD content providers. You make a custom agreement providing for storage, CDN distribution, custom javascript coding for the user behaviors you want to know about, etc.
 - But that means if you're storing everything else in the Amazon or Azure clouds (*remember the quote about storage driving growth*), you now have two storage providers, two OpEx payments.



Demographic/Economic Analytics

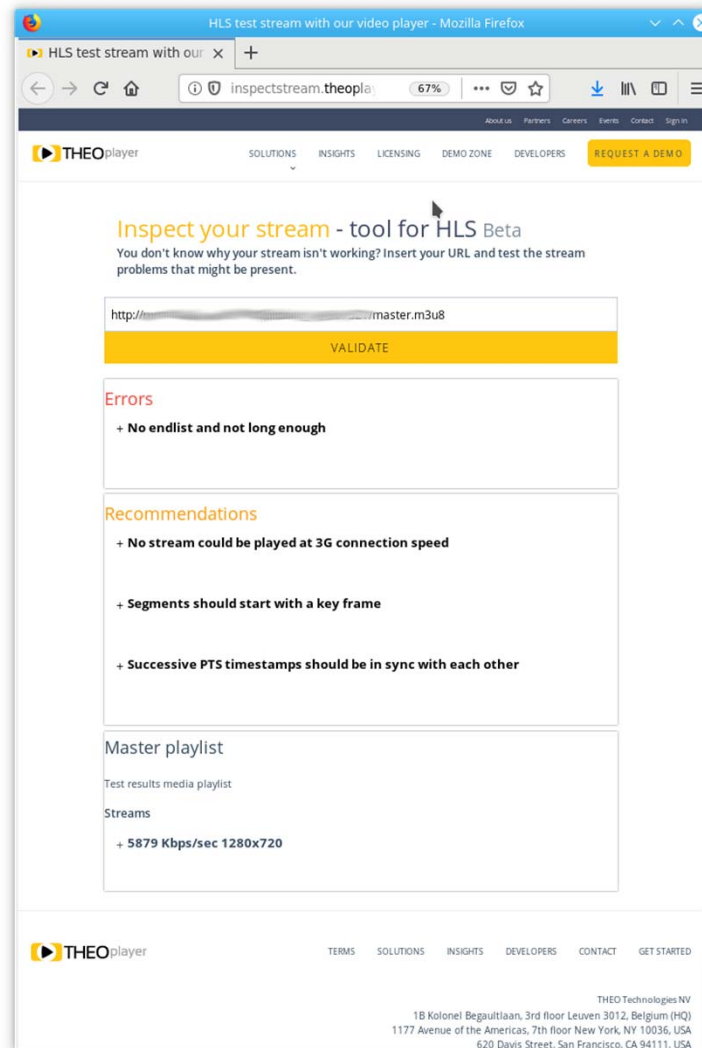


Demographic Analytics



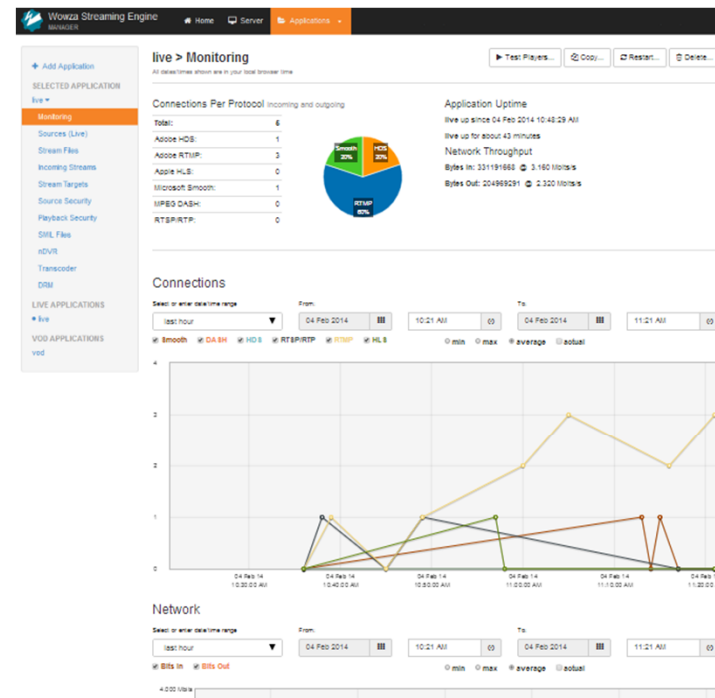
Non-Continuous Analysis Tool

- Perfunctory Stream Analysis, But it may be enough for many.
- For Stored Streams... Not for Live Streams.



There ARE Additional MarketPlace Images

- Example: Wowza via Azure
 - Easy to install.
 - But if you want stream quality metrics, you'll still probably have to buy an add-in. (Wowza does include an API to work with Google Analytics, for demographic analysis.)
 - Still, it's All on Azure; Though With Two Vendors Expect Your OpEx to Increase.

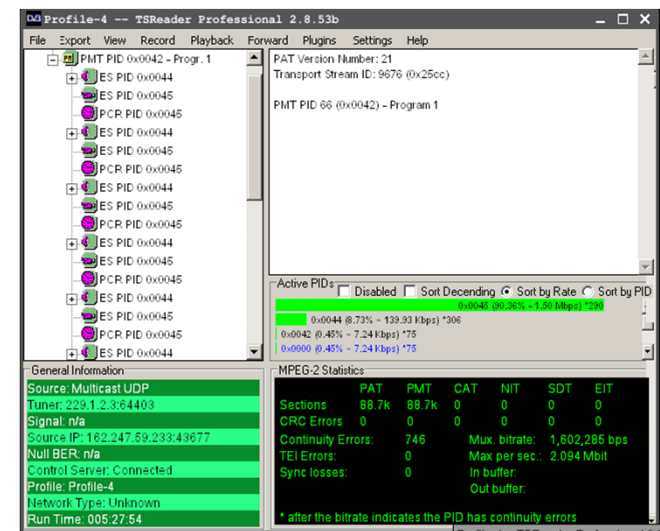


But Without Stream Metrics, What Could be Missed?

- As Important a Problem as No Audio!

With a Stream Probe Window as Below, the Two Green Bars Immediately Indicate Video and Audio are Present

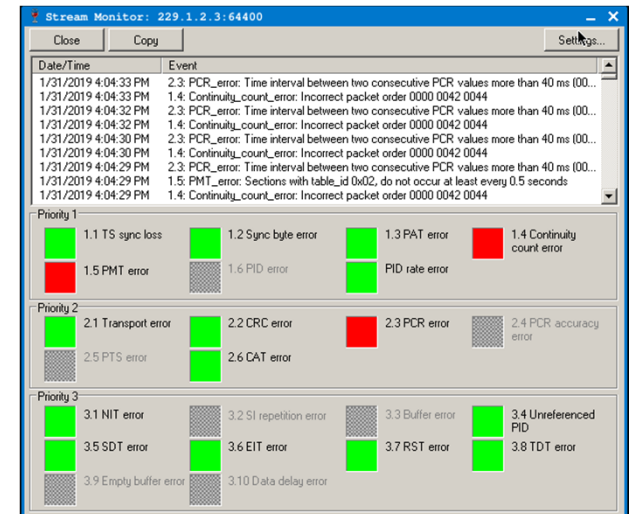
If Your Stream is a Live Stream, You Can't Be Sure From One Minute to Next That a Technical Problem Won't Be Present at the Source.



What Other Important Problems Could be Missed?


- As Nasty a Problem as Jitter in the Source

In this Monitor Window, Program Clock Register Errors Could Indicate Playback Errors Such as jitter



What Else Could be Missed?

- It Would Certainly Miss Non-fatal Source Discontinuities...



- Say, did I just blink or did the stream go down for half a second?

What Else Could Be Missed?

- *To sum up:* Almost any Video or Audio Problem that Might be Due to an Error in the Source Stream Itself—and Which Would Probably be:
 - Visible to Your Viewers
 - Visible in a Simple CPU or Network Measurement... *But Only Afterwards*
 - Visible in a Probe/Analysis of the Stream Itself
 - Visible in a Demographics/Commerce Oriented Analysis View... *But Only After You've Lost Many Viewers*



Is This an Innate Problem In Cloud Encoding?

- Not at All—But it May Indicate a Strategic Weakness: *Cloud Vendor Based Encoders Appear to ASSUME that the Source Content is PERFECT*
- Given the Usual (*basic price*) Cloud Based Encoder Options are ONLY PassThrough or One-Size-Fits-All Transcoding, Why Would They NOT Want to Assume The Source Stream is Perfect?
- You Can Find a Service... But it's a Service, Not Your Piece of the Cloud, So You're Probably Looking at Increased OPEX
- You Can Find Some Images...
But do they Have the Metrics You Want? Again Increased OPEX.
- So Do We Need More Options? YES.



First Let's Determine What We Need

- Live Streaming Is Different.
 - Need to Check Continuity
 - Need to Check Audio
 - Need to Alert on Minor Issues That Your Experience Shows Are Indicators of Imminent Failure
 - Visuals Are Good; in fact, a Video Wall is Essential.
 - Need to Ignore Certain Problems Unique to a Particular Stream But Which Experience Shows Can be Safely Ignored
 - Need to Know WHERE the Problem Is... Is it In the Source? Only After the Encoder?
 - Therefore Probe/Analysis Should Ideally be at a Place Where Encoder and Incoming Stream Meet, ideally, in a VM Image in the Same Subnet or Region in the Cloud



Cost Considerations

- Reasonable Cost
 - First We Need a Probe That Fits Inside a Linux VM Because That's The Cheapest Cloud Option
 - We Want to Limit It to Just One More VM to Provision, Though It May Need a Fair Amount of Processing Power to Provide a Video Wall.
 - We Want to Remember to use the Same Cloud Vendor Region as the Encoder for Lower Bandwidth Charges. (Assuming Your Cloud, Like Most Big Cloud Vendors Does Not Charge for Incoming Bytes, or Bytes Within the Region)



What We Did

- We Utilize a Probe That Fits Inside a Linux VM: in our case, we have a Windows Based One. We Run it Using "wine".
 - *Which believe it or not, gives us several advantages:*
 - We can start with a pristine environment, taken from a read only image, any time we start our daemon. As if we have a virgin copy of Windows every time!
 - Host VM can “take down” this environment like a normal daemon without affecting anything.
 - Completely passive monitoring that cannot affect the encoder process.
 - Easy to create multiple probes to monitor additional streams, or the same stream at different points in its “lifetime” or “grooming” stages.
- Then We Made It Highly Accessible, but Securely, via RDP, Which is a Fairly Lightweight Protocol (when not using vnc).



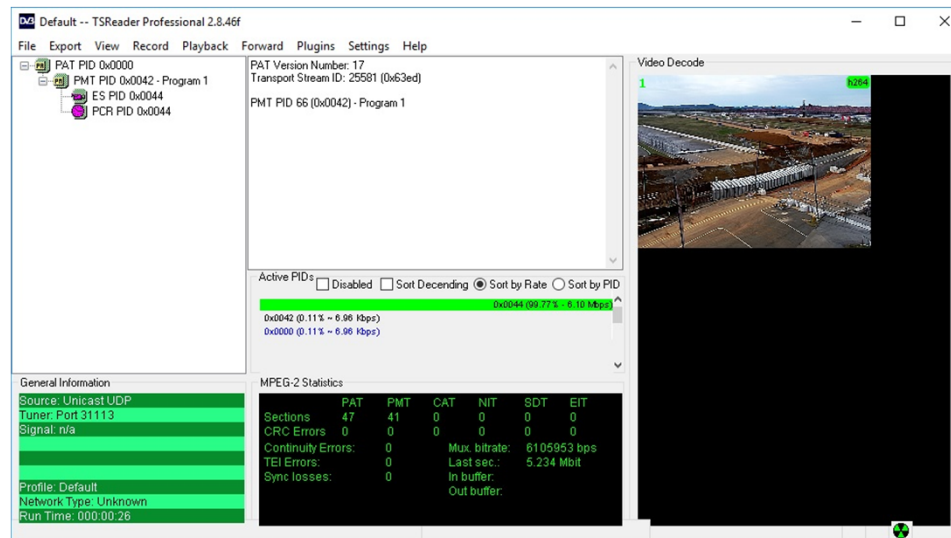
Portability

- Cloud Agnostic: We Made It Highly Portable, With Images for Azure and Amazon Marketplaces, and VMs for VMWare. If You Can't Figure Out How To Stick This Near Your Encoder, You're Not Thinking Very Hard.
- One Time License: The Only OpEx Costs May Be Outbound Bandwidth to Reach Your Monitoring Personnel



Probes Main Window Example

- There's a Lot of Data Behind That Window, Which We Will Explore



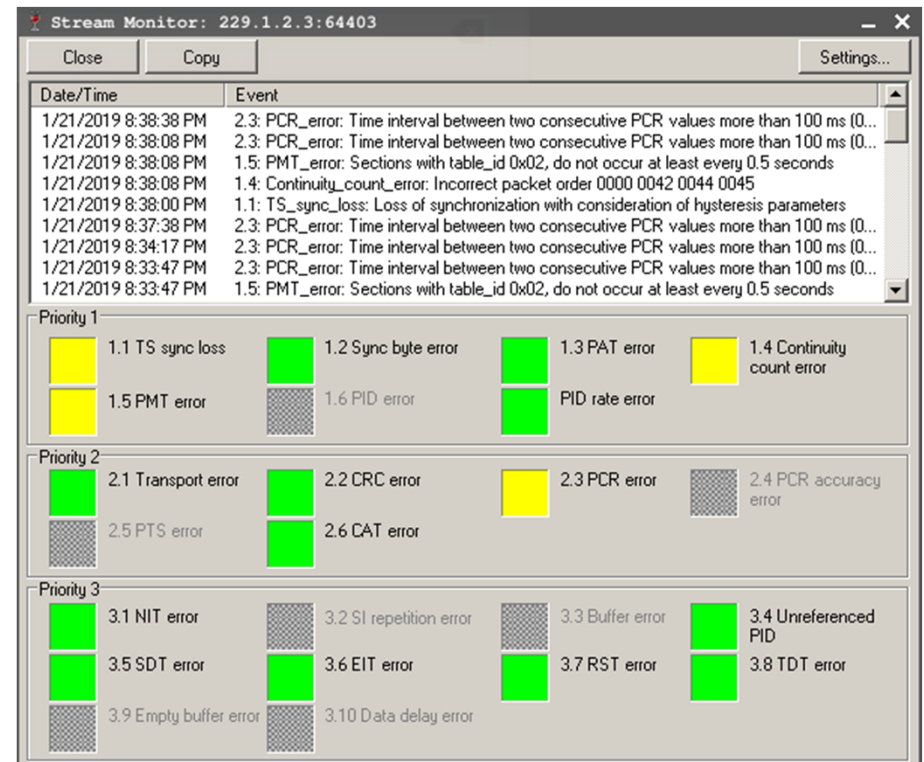
The What and Why of It

- What can we see in these probes that we can't see in the cloud vendors' tools view? *Metrics as a strategic approach to “preventative” stream health:*
 - Error counters as advanced early warning system
 - Drill down table analysis to identify fatal errors so that they can be corrected
 - Logging... to provide a clue as to when the problem started
 - Access to raw data in a structured (mysql) manner
 - For Instance...
 - Sync Issues
 - Continuity Count
 - PID Error
 - PCR Errors
 - Other Minor Errors
 - Graphs and visual aids for the boss



Sync Issues

- You can usually see a sync byte error or two prior to complete sync loss especially in a live stream
- Monitor (right) shows yellow (warn) for sync loss. It failed soon after. (Note: the PMT and PCR errors are related to each other, and not fatal).



Continuity Count

- You're not likely to see an indication of minor network errors such as out-of-order packets on the cloud vendor control panels... but this will usually provide such an indication (in this case, yellow; similar to previous, without the sync error)

The screenshot shows the 'Stream Monitor' application window. The title bar reads 'Stream Monitor: 229.1.2.3:64404'. Below the title bar are buttons for 'Close', 'Copy', and 'Settings...'. The main window is divided into two sections. The top section is a log table with columns 'Date/Time' and 'Event'. The bottom section is a grid of error status indicators organized by priority.

Date/Time	Event
1/21/2019 8:38:52 PM	2.3: PCR_error: Time interval between two consecutive PCR values more than 100 ms (0...
1/21/2019 8:38:35 PM	2.3: PCR_error: Time interval between two consecutive PCR values more than 100 ms (0...
1/21/2019 8:38:11 PM	2.3: PCR_error: Time interval between two consecutive PCR values more than 100 ms (0...
1/21/2019 8:36:18 PM	2.3: PCR_error: Time interval between two consecutive PCR values more than 100 ms (0...
1/21/2019 8:36:14 PM	2.3: PCR_error: Time interval between two consecutive PCR values more than 100 ms (0...
1/21/2019 8:36:14 PM	1.5: PMT_error: Sections with table_id 0x02, do not occur at least every 0.5 seconds
1/21/2019 8:36:10 PM	2.3: PCR_error: Time interval between two consecutive PCR values more than 100 ms (0...
1/21/2019 8:36:10 PM	1.4: Continuity_count_error: Incorrect packet order 0000 0042 0044 0045
1/21/2019 8:36:05 PM	2.3: PCR_error: Time interval between two consecutive PCR values more than 100 ms (0...

Priority 1			
1.1 TS sync loss	1.2 Sync byte error	1.3 PAT error	1.4 Continuity count error
1.5 PMT error	1.6 PID error	PID rate error	

Priority 2			
2.1 Transport error	2.2 CRC error	2.3 PCR error	2.4 PCR accuracy error
2.5 PTS error	2.6 CAT error		

Priority 3			
3.1 NIT error	3.2 SI repetition error	3.3 Buffer error	3.4 Unreferenced PID
3.5 SDT error	3.6 EIT error	3.7 RST error	3.8 TDT error
3.9 Empty buffer error	3.10 Data delay error		

PCR Errors

- Both streams at right are possibly losing sync... but some streams behave differently than others.
- One has to “learn” the character of the stream during a setup period and set an allowable range, stream by stream, for alerting.
- Top stream visibly bad; but the bottom (a real live stream we know intimately) can go on like that for weeks and weeks!



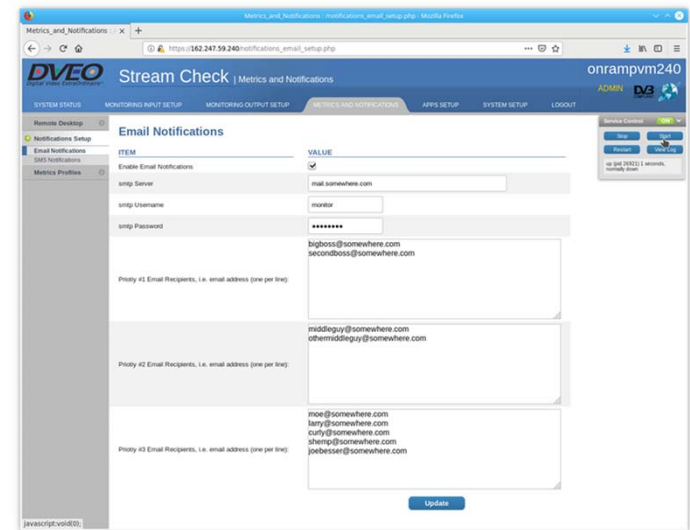
Logging

- Time stamps can lead to identification of problems quickly if you know what time the product started
- Access to this level of details is essential; any sysadmin will tell you that the least sexy but most important things are logs.
- We'll show you what we did with this later.




Strategic Alerting

- You Need a Hierarchy of Alerts and Alerted Personnel
 - Big Bosses Get Alerted at Complete Failure Situations:
“The Barn is On Fire”
 - Middle Managers Get Alerted at the Start of the Emergency:
“I Smell Smoke in this Barn”
 - NOC Folk Get Alerted at Everything:
“That Cow’s Getting Close to That Kerosene Lamp”
 - Note: Screen Shot of Stream Check from a KVM VM



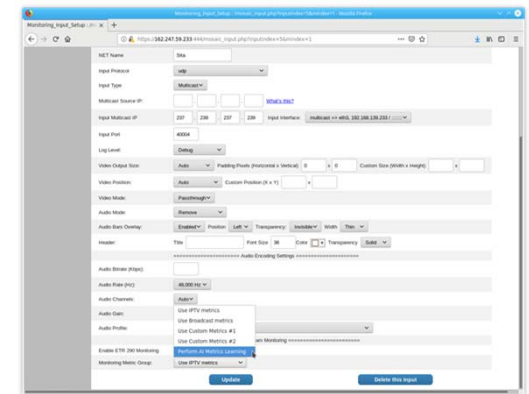
Notes: Machine Learning as it Relates to Monitoring

- How much can AI be depended on to learn?
- As mentioned, streams vary to a great degree. Certain errors as measured by metrics in one source may be tolerated to a much greater degree than in another.
- What machine learning can do is to establish base minimums and maximums during a setup period, establishing a “recommendation” so that alerting can be triggered below or above those levels.



How To Implement Machine Aided Monitoring

- Since alerts have minimum/maximum triggers, if we keep a log with every single error for a period of say, 24 hours, and if we say that the stream quality was acceptable in that period, we've got minimum and maximum values for as many metrics as we've logged. It's a matter of query the log, and output to a config file.



How Far Can We Go With It for QA?

- Since the probe knows when the stream goes down, we could conceivably grep backwards from any stop event and change a max or min level (though we haven't implemented this).
- Likewise, if monitoring personnel “approve” performance for say, the last 24 hours, the machine could “loosen” the min and max levels.
- Since we already have scripting present to restart the stream when the monitor in the DVEO encoder detects a major problem, we could implement an API in the cloud monitoring VM to similarly research and create a new rule based on the moment the DVEO encoder detected the error.
- *Note: we're also providing the DVEO encoder for Amazon and Azure clouds.*

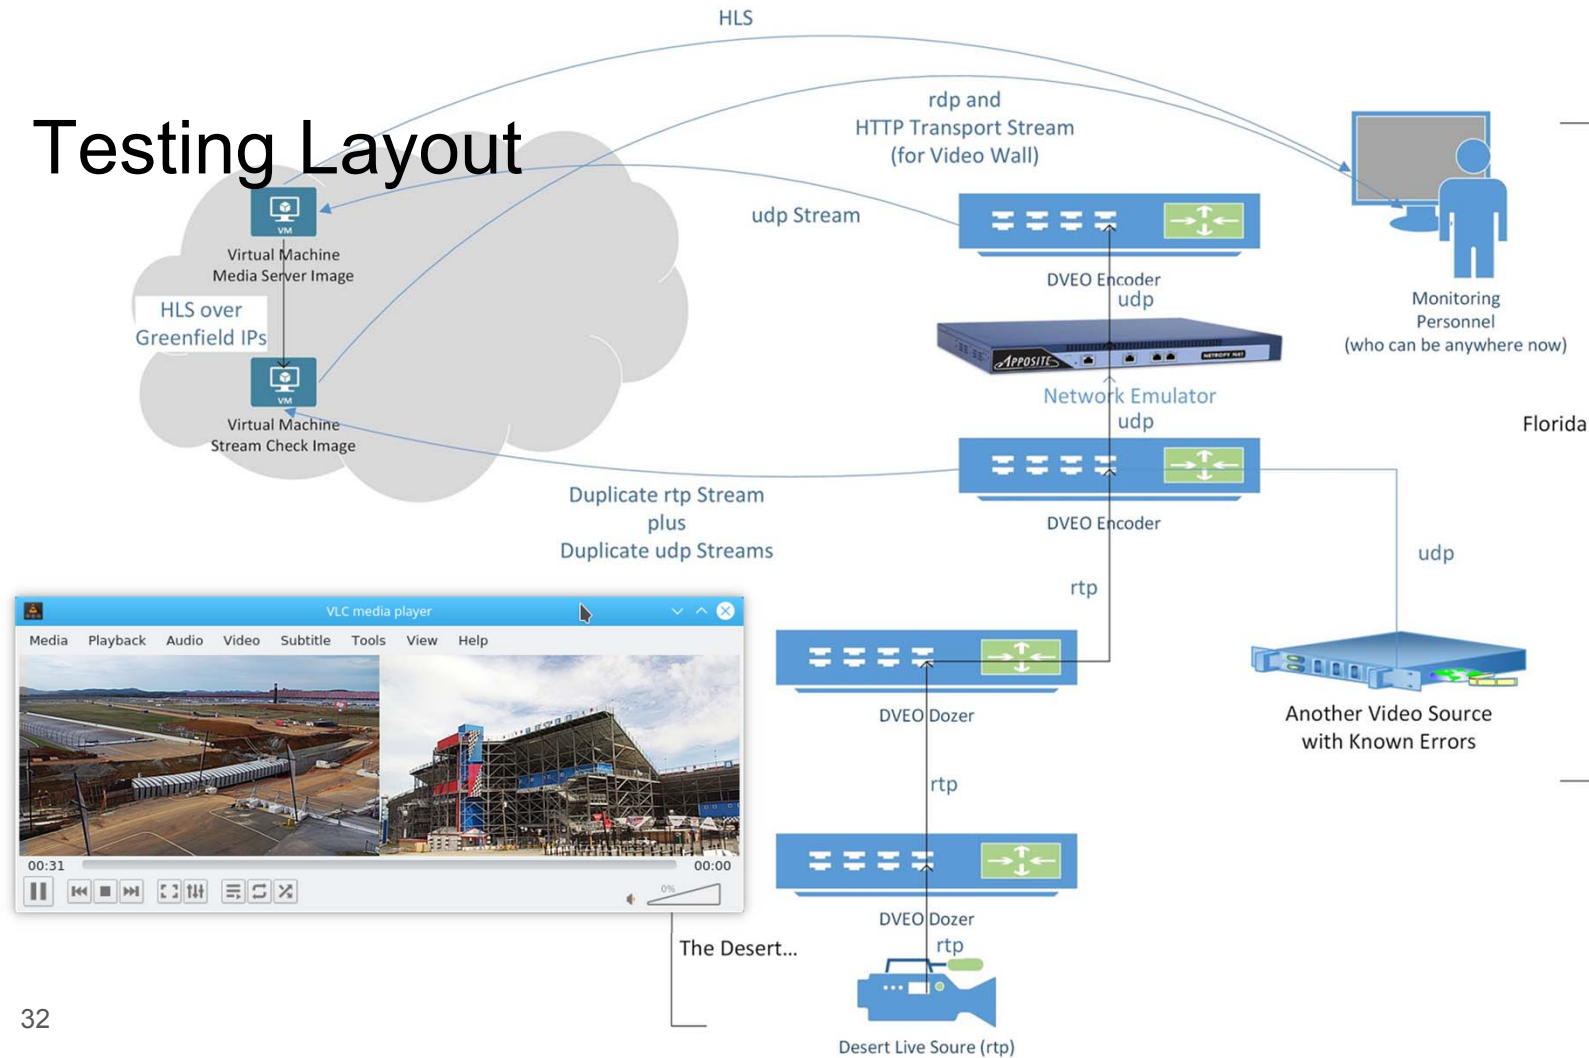


Real World Testing

- We sent a (1) our “fickle” desert camera stream, plus (2) a recorded stream with known errors to a media encoder in our own cloud, from which HTTP/TS and udp were monitored through our cloud Stream Check
- We sent them through our Network Emulator in the lab, out to the cloud and through our Stream Check
- See Layout Next Page





Testing Layout



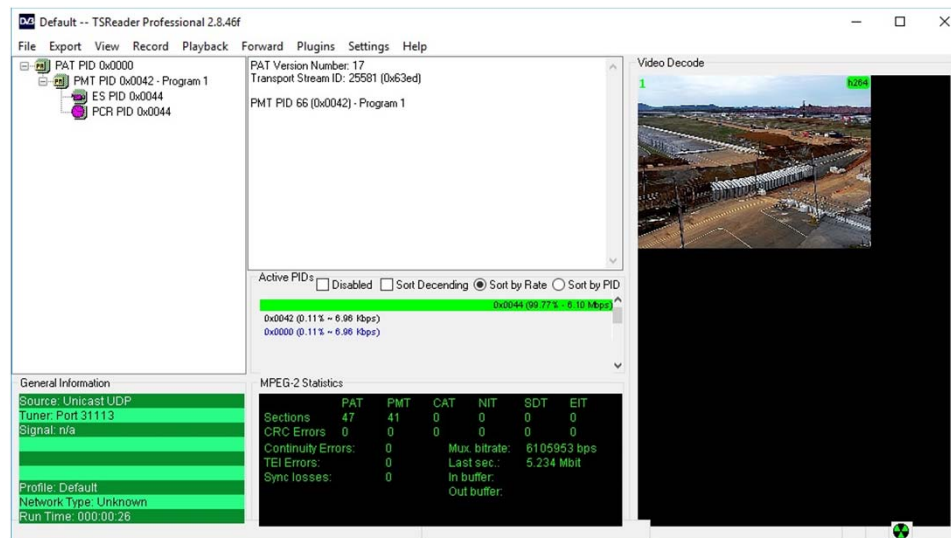
Network Emulator View

- Left: bypass mode; Right: dropped random packets 2%
(to help you can visualize what you'll be seeing in the probes and data).



Probe Window

- Main Window, Before Introducing Errors
- We'll be Using the Log Windows to See the Actual Errors Once We Introduce Them With the Network Emulator



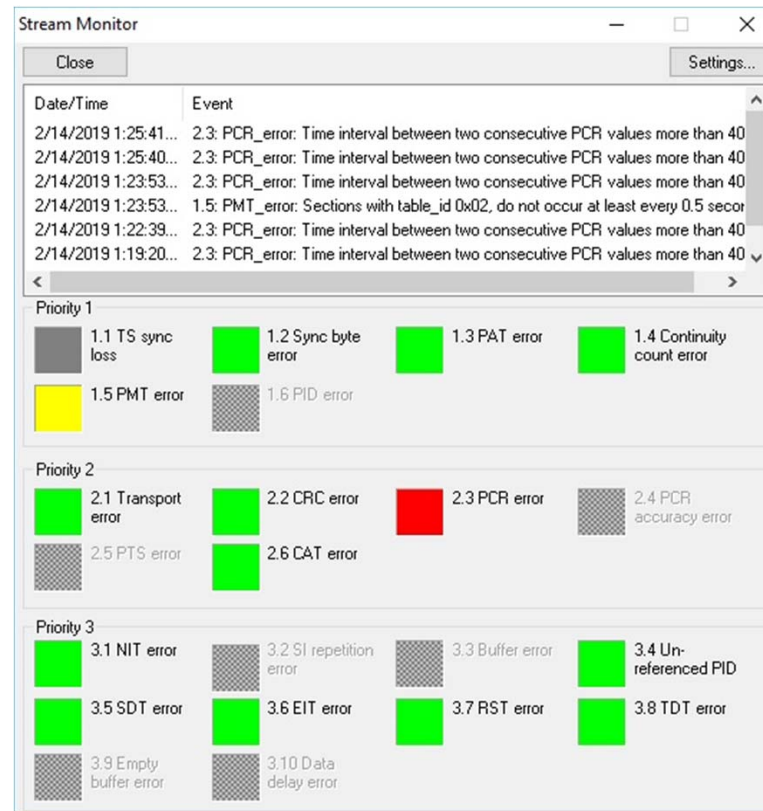
Probe Lights/LogViews

- Corrupt a few Packets...



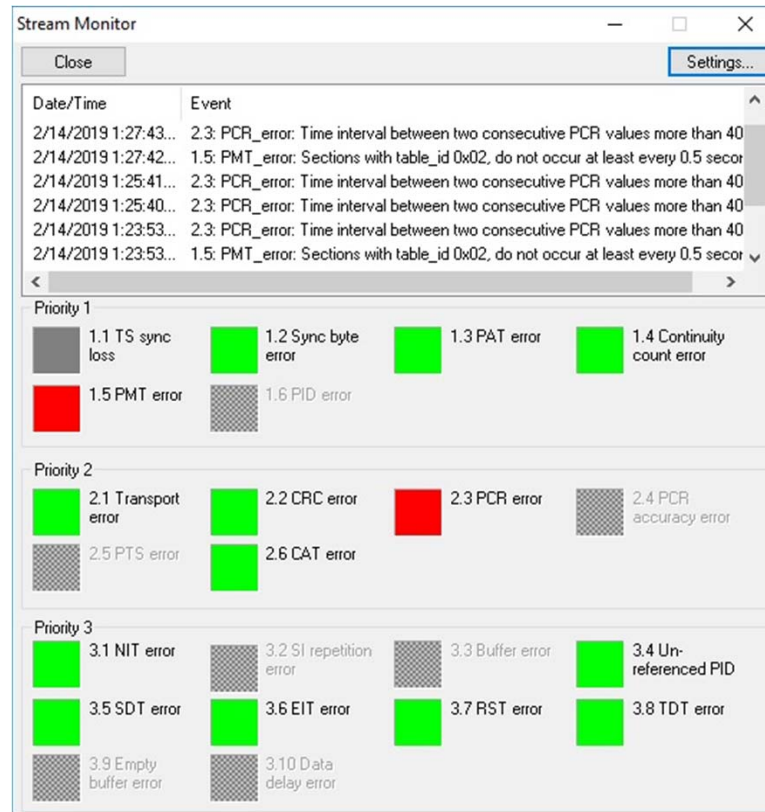
Probe Lights/LogViews

- Remove Corruption, Add a Touch of Jitter...



Probe Lights/LogViews

- Add More Jitter and a Few Dropped Packets...




Probe Lights/LogViews

- Jitter, Corruption, but no Dropped Packets...



Learning So Far

- Different Patterns of Warning Lights for Different Error Conditions
- Though Different Mixes of Error Conditions Could Be Confusing (let's face it, in the real world you'll *never* have *just* jitter, or *just* corruption), Certain Conditions (such as corrupt packets) Appear to Trip the Same *Types* of Error Categories
- We Have TimeStamps in the Log (the text area at the tops of those windows)
- We Have Specific Error Types Organized by Category of Severity (Priorities 1,2,3)
- *Therefore...* We Have Structured Data We Can Capture to a Database



Database Table Views

- If We Let mysql TimeStamp Each Error, Then Summarize by Minute...

The screenshot displays a MySQL database interface with two windows. The left window shows the 'Prod Database/streamcheck/ProbeErrors1' table with columns: Timestamp, ProbeType, ProbeName, and ProbeData. The right window shows the 'Prod Database/streamcheck/ProbeSummary1' table with columns: MinuteMax, 1_1, 1_2, 1_3, 1_4, 1_5, 1_6, 1_7, 2_1, 2_2, 2_3, 2_4, 2_5, 2_6, 3_1, 3_2.

The query result for the 'ProbeSummary1' table is shown below:

MinuteMax	1_1	1_2	1_3	1_4	1_5	1_6	1_7	2_1	2_2	2_3	2_4	2_5	2_6	3_1	3_2
2019-02-10 20:51:00	0	0	0	2	2	-1	0	0	0	2	-1	-1	0	0	0
2019-02-10 20:52:00	0	0	0	2	2	-1	0	0	0	2	-1	-1	0	0	0
2019-02-10 20:58:00	0	0	0	2	2	-1	0	0	0	2	-1	-1	0	0	0
2019-02-10 20:59:00	0	0	0	2	2	-1	0	0	0	2	-1	-1	0	0	0
2019-02-10 21:00:00	0	0	0	2	2	-1	0	0	0	2	-1	-1	0	0	0
2019-02-10 21:01:00	0	0	0	2	2	-1	0	0	0	2	-1	-1	0	0	0
2019-02-10 21:02:00	0	0	0	2	2	-1	0	0	0	2	-1	-1	0	0	0
2019-02-10 21:03:00	0	0	0	2	2	-1	0	0	0	2	-1	-1	0	0	0
2019-02-10 21:04:00	0	0	0	2	2	-1	0	0	0	2	-1	-1	0	0	0
2019-02-10 21:05:00	0	0	0	2	2	-1	0	0	0	2	-1	-1	0	0	0
2019-02-10 21:06:00	0	0	0	2	2	-1	0	0	0	2	-1	-1	0	0	0
2019-02-10 21:07:00	0	0	0	2	2	-1	0	0	0	2	-1	-1	0	0	0

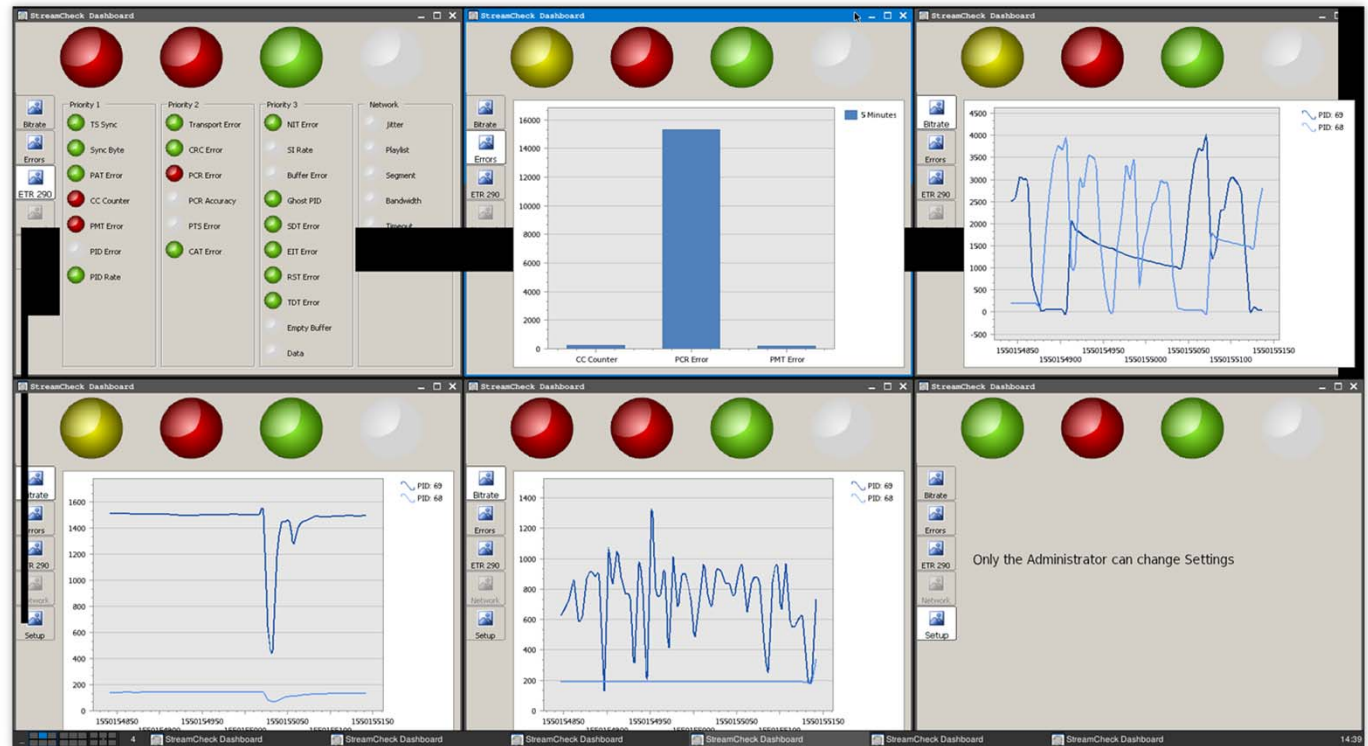
The query used to generate this result is:

```

23 MAX(ProbeStatus1.3_9) AS 3_9,
24 MAX(ProbeStatus1.3_10) AS 3_10,
25 FROM ProbeStatus1
26 WHERE ProbeStatus1.Timestamp < STR_TO_DATE(CONCAT(YEAR(NOW()), '-', MONTH(NOW()), '- ', DAY(NOW()), ' '), '%Y-%m-%d %H:%i:%s')
27 GROUP BY MinuteMax
    
```

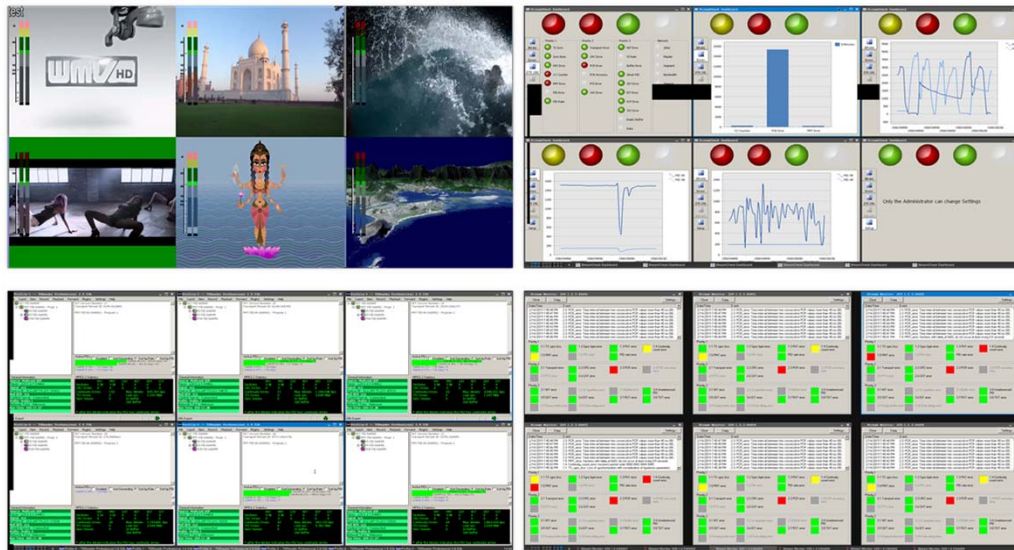

Dashboards

- We Now Have the Data to Provide Color Coded Real Time Analysis




Putting It All Together...

- 4x 1920x1080 Screens Monitoring All Aspects of Stream Life, From the Cloud. (below is a different installation with six streams, rather than two).
- If Your Monitoring Personnel Miss An Error, They Won't Have Much Excuse!



Could We Do More? Location, Location, Location

- Cloud-based Monitoring Provides Additional Options:
 - Check stream health at entry point. For example, since we also have our Dozer Error Correcting Protocol available for the cloud, the receiver could multicast the stream to both the monitoring probe and the encoder at the same time; real-time analyzing the source at the *same time* it enters the encoder. (we didn't do it in this testing because our network emulator doesn't do multicast)
 - Check stream health at encoder exit, even each bitrate/size for ABR
 - Check HLS packaged uplink to CDN
 - Check CDN "received" stream. (though if it's Azure to Azure CDN or Amazon to CloudFront this might not be necessary; but it would be worth a real world test!)



Additional Learning

- If you manage a stream right with a setup period to learn how it “behaves” (or “misbehaves”) you can configure the alert levels just so.
- If something is seriously wrong, you can’t beat having a huge amount of information. Between live measurements, and especially the log, with its timestamps, you don’t have to be Sherlock Holmes to find the problem. Just patient.
- A video wall as a first “line of defense” is excellent. Just make sure to put labels and the audio level indicators in each!



So It's Just Common Sense?

- Yes.
- If your encoder and distribution system are in the cloud, put your monitoring alongside it. Besides an immediate reduction in CAPEX, in the end it may save OPEX – because you've eliminated transport problems between the cloud encoder and the monitoring system in your NOC which may create false positives.
- If you can expose all stages of a stream to your probes, it's more likely you'll identify any problem. You just have to learn how to identify the time and place where the problem begins. Which is sometimes quite easy.



Do I Still Need A NOC for Probe/Monitors?

- Maybe Not. Did You Plan to Get Rid of It?
This Brings Up the CAPEX vs OPEX Issue.
If You Put a Probe in the Cloud:
 - Inbound Bandwidth Will Probably be Free, so That's a Win
 - It's Another VM, so That's an OPEX cost
 - It's Going to do a Video Wall, Which Requires Processing Power, so that's an OPEX cost
 - That Video Wall Will Also Require Outbound Bandwidth (from the cloud) for Those Who Monitor It, so that's an OPEX cost
 - Then There's The Question of Seeing its Analysis... That's what the RDP protocol is for.



The Case For Cloud-Embedded Monitoring/Alerting

- Potentially Big Benefits:

- You've placed a probe in your IP stream to count every little error and to make every internal table transparent. That's debug Information that can help you solve the toughest problems.
- Since your probe is in your cloud, your 24x7 monitoring personnel aren't tied down to a specific location or NOC... they can be *anywhere* with a good dependable connection, which answers the NOC question.
- If you make intelligent use of cloud vendors' regions, you can place the monitor practically on top of your encoder and monitor multiple points in the stream's lifetime.



End

- Thank you.

