VideoFlow

Broadcast with Confidence

Impairments, Disruptions, and Outages

The limitations of SRT video transport in a cloud-connected world

VideoFlow White Paper

Introduction

The concept of a cloud datacenter has introduced great opportunities and great challenges. Migrating traditional linear television workflows has presented a challenge since the beginning of datacenter co-location.

The specific performance requirements of broadcast video over IP using industry standard formats such as MPEG2 Transport Streams can challenge even the most robust traditional cloud workflows that perform well for many other applications. Automated configuration and management of the cloud data center is not generally fine-tuned to the unique characteristics of low-latency video traffic.

In a cloud-integrated or fully cloud-based deployment, regardless of "direct-connects" or internet-based links to and from the cloud, at some point you are surrendering control and visibility of the underlying network to someone else, and the cloud essentially becomes a black box with few useful windows and control levers.

Moreover, the windows and levers that are available are often not understood by broadcast engineers nor corporate IT teams as the considerations for the special requirements of video are not well understood by many corporate network engineers. Tools that are available such as performance logging frequently do not have the sub-second resolution required to determine packet loss or other network issues are occurring.

While the cloud's strengths are vast, there are also many potential weaknesses when it comes to real-time video transport.

Grave Mistakes in Cloud Transitions

Planning for a migration to the cloud is often not as simple as advertised. There are many common pitfalls that might undermine the effort:

- Short term POC tests with predictable failures
- Believing the brand promise and "trusting the Cloud"
- Turning off your critical thinking under pressure from the C-Suite to move things to the cloud for non-technical reasons, such as to appear to be a more innovative company for the benefit of shareholder impression, without fully understanding the characteristics and risks associated with a Broadcast Video Cloud Deployment that must be mitigated to realize the unique benefits of the cloud.
- Not establishing an ongoing surveillance and proof of configuration testing scheme to insure background changes in the cloud fabric have not compromised your architecture.

Still, even if the big mistakes are covered with insightful planning, there are many particular details that must be ironed out for a successful migration to the cloud.

Latency Issues

Software-based ingest, playout, encoding, processing, and transcoding platforms tend to have higher latencies which result in glass to glass latency significantly higher than traditional satellite distribution architectures, which relied on systems with hardware based processing that had predictable and deterministic path latency.

More than ever, latency is a significant performance parameter which must be engineered at every point in the chain and consistently monitored for performance. This is not talking about synchronization in an audio video sense, but rather the total amount of delay that is introduced as a result of various steps in the broadcast chain.

How long is an acceptable glass-to-glass latency value?

Mean Time to Recovery

One of the most important realizations when considering a cloud deployment for video is a specification for MTTR, or Mean Time to Recovery. There are many excellent resources to study this concept in depth, but the most important question to answer with respect to cloud architecture of video is the following:

How long of a disruption or outage is acceptable?

Double Trouble

Latency and MTTR often combine in unfortunate ways by significantly amplifying the impact of various potential failures or degradations in the encoding and transmission chain to and from the cloud.

For example, if you were using a software-based VBR statmux encoder to optimize a multi program transport stream delivery, these frequently have multi-second processing delays due to the nature of running multiple VBR encoding passes in order to optimize the quality of each stream in relation to the maximum available bandwidth. In practice, we have seen cloud-based statmux encoders at 5-7 seconds of latency.

If an input source is lost on the front end of one of these systems, and if the system is set to failover to a backup input after 2 seconds of stream loss, then it will take the system 2 seconds to recognize the loss, switch to the backup, and then the feed will be restored perhaps 7-9 seconds later for viewers downstream as the statmux will need some time to fully integrate the new service. This is the best case scenario.

However, in more common scenarios of network impairment and disruptions, the stream is typically not totally lost for a period of time that exceeds the input switching detection interval, and so the input won't necessarily failover. While some equipment can detect ETR-290 stream errors like CC errors and make switching decisions based on a tuned set of parameters, this is surprisingly overlooked in many systems. Because of the processing latency, it takes 10-12 seconds to see the defects introduced by the impaired input, either visually by an operator (remember, the return path will also have latency to the human operators) or via an alarm generated in a QoE analyzer downstream, then many more seconds for an operator to assess the issue, determine the root cause, and manually reconfigure an element in the chain.

If your MTTR is 1-2 minutes, then this works just fine. But if you could lose significant revenue with an issue of that duration, then you must architect the systems to reduce the possibility of all stream degradation failure modes by all possible means so that there is a much smaller range of potential issues that can cause such an extended impact to image stability, and as a result, reduce your probable MTTR to the lowest possible acceptable values.

It is critical that the specific characteristics of how each part of the system responds to network issues be fully understood through thorough testing in order to establish a baseline MTTR for each potential failure mode throughout a broadcast chain to and from the cloud.

When a top of rack switch fails in the cloud, how quickly can the network heal itself? 100ms? 500ms? 5sec?

Are trunk links to server hardware optimized to eliminate disruptive switching between interfaces to load balance?

Network Impairments

After determining the acceptable latency and MTTR values for your requirement, and once you have determined the capabilities and limitations of the various software and hardware components within the processing chain, then the next step is to consider the possible network issues and limitations, and assess the suitability of a given transport protocol for your requirements.

It is of paramount importance to consider the distinctions between different types of network issues for which the stream transport protocol must be designed to overcome.

For ease of discussion, we will categorize network issues into three general categories:

Impairments

- Consistent Low Level Packet Loss
- Jitter
- RTT changes

Available Video Transport Solutions:

- Longer Receive Buffers
- FEC
- Basic ARQ

Disruptions

- Bursts of significant packet loss
- Full WAN traffic drops less than 500ms
- Packet reordering
- Short outages due to network fabric changes 500ms-2sec
- Interruption of bi-directional signaling traffic

Available Video Transport Solutions:

- Basic ARQ like SRT up to protocol limitations
- Longer Receive Buffers
- Selection of equipment specifically designed and tested to support very long buffers and can reorder RTP packets.

Outages

- Fiber Cuts
- Cloud, infrastructure configuration errors
- Cloud Outages
- Local Disasters

Available Video Transport Solutions:

- Redundant WAN Transmission Paths
    - Input Switching on Failover
        - Challenges: Detection, intelligent failover, reversion management
    - SMPTE 2022-7 Seamless Protection Switching Input
        - Challenges: stability while one path is down, visibility into the ongoing performance of each path
- ARQ (to minimize switching events and optimize disruptions after switching)
- Longer Receive Buffers (to optimized switching time)

It is important to recognize that while the various available solutions serve to solve 1 or perhaps or 2 categories of network issues, no one solution is suitable on its own.

For instance, 2022-7 Seamless Protection Input Switching will work great so long as there are two reasonably solid connections working, but often when you have a fiber cut on one leg, the other leg is of insufficient performance to carry the load without some kind of FEC or ARQ protocol. Furthermore, for many products, there is little visibility into the performance of each leg of the circuit so you may not know there is an impairment issue on the surviving leg during an outage event until the event actually occurs. In other words, you won't know the surviving path is underperforming until you need it the most.

For 2022-7 by itself, you must have a solution that employs at least 3 diverse redundant WAN links so that any two out of the three are likely to be online simultaneously under almost every conceivable scenario. With 2022-7 streams that are also delivered with ARQ, just one WAN link can provide a very high level of resilience and redundancy.

VideoFlow Key Advantages over Any Other Transport Solution

VideoFlow is truly distinct in the marketplace, with a range of features and technical sophistication designed to handle the most complex designs and technical challenges.

- Combines Multiple WAN Paths, 2022-2 FEC, 2022-7 Hitless Switching, and highly-tunable ARQ in one stream transmission.
- Retains the original packet RTP sequence numbers throughout the transmission chain.
- Supports up to 8 sources for Hitless inputs
- Supports connections with highly variable latency characteristics and can leverage them together to maximize reliability
- Provides visibility of each WAN path even in Hitless switching configurations for early warning of network issues.
- Supports any number of protocols and deployment methods including on-site hardware, virtual compute instances, bare metal compute instances, and docker containers for flexible and innovative architectures to mitigate cloud risks.
- Very Cost Effective Total Cost of Ownership compared to subscription-based service offerings

The unique combination of VideoFlow features provides near infinite deployment configurations to build a mesh ARQ transport deployment on top of the cloud infrastructure that is supremely redundant and resilient.