Audio Engineering Society

Convention Express Paper 15

Presented at the 153rd Convention
2022 October

AES67 Wide Area Network Transport Utilizing the Cloud

Bill Rounopoulos[1] and Andreas Hildebrand[2]

[1] *Ross Video, Ottawa, Canada*
[2] *ALC NetworX, Munich, Germany*

Correspondence should be addressed to Bill Rounopoulos (brounopoulos@rossvideo.com) or Andreas Hildebrand (andreas.hildebrand@alcnetworx.de)

ABSTRACT

This paper highlights the challenges with transporting AES67 through a public cloud infrastructure and how these can be solved along with topics for further study. This will be done through the lens of results from recent live demonstrations that prove it is feasible including one where audio was transported between North America and several locations in Europe. Along the way, the learnings, as well as future considerations, to enable the practical use of AES67 cloud transport in real-world applications, will be reviewed.

1 Introduction

Remote production is clearly an important topic today as companies around the world race to maintain their existing workflows with their talent dispersed in many off-site locations. Remote production is the new normal and - now that we have experienced its potential – is likely to continue to be a focus in the future. While many solutions have been hastily cobbled together, there is a need for higher quality productions with lower latency that integrate easily with existing equipment.

2 Starting Point

We started with a basic question: can we send AES67 traffic over the public infrastructure and over long distances? Then we wondered would we be able to listen to something resembling audio? Would it be good quality? It is one thing for a single company with their own equipment to do it, but could we also interoperate with equipment from other companies? After all, this is the whole point behind AES67. Finally, we also wondered how we would do it and what challenges would we face.

Before digging into the setup and issues that needed to be overcome, it is important to understand that AES67, even though it uses IP, is designed to be used in local area networks (LANs) and well-managed corporate networks. Despite this, AES67 has been proven and is being used commercially in wide area network (WAN) applications across private networks, even though its use in WANs was never contemplated by the standards. It is important to note that private dedicated networks, whether 'owned or leased' are well-architected, support the necessary protocols, have predictable behaviour, and come with performance and service level guarantees. Public networks, on the other hand, are the equivalent of the 'wild west'. You can't control them. They have limited protocol support and are congested and unpredictable. Public networks suffer from packet loss due to link failures and have large, sometimes dramatic, latency due to packet re-transmissions. This makes the public environment very inhospitable for AES67!

3 Challenges

There are three main challenges: latency and packet jitter; packet loss; timing and synchronization.

Fortunately, the increased latency and packet jitter of the public network can be handled by design through the use of large receiver buffers. While AES67 only requires 3 msec, it strongly recommends supporting the equivalent of 20 msec of audio content for receiver buffers. Most well-designed AES67 solutions, like all the equipment used in this experiment, have even bigger buffers that can compensate for the added delay. The AES Standard Committee working group SC-02-12-M has issued its guidelines for AES67 over WAN applications and a key recommendation is to increase the buffer size within devices.

Packet loss is another important challenge. AES67 is not designed to cope with dropped packets. Fortunately, there are other methods and transport protocols that are architected to deal with dropped packets without introducing a lot of extra latency. These include redundant stream transport (SMPTE ST 2022-7), Secure Reliable Transport (SRT), Zixi, and Reliable Internet Stream Transport (RIST) but there are many others. We solved the challenge of packet loss by encapsulating AES67 traffic within SRT.

The final but significant challenge is timing and synchronization. We start by having a separate Precision Time Protocol (PTP) Grandmaster (GM) at each site that is synchronized to GPS. PTP is then only used locally at each location to achieve synchronization among all participating devices across all sites. No PTP packets are sent across the WAN or through the cloud, which would simply not be practical as packet jitter is too high to achieve adequate synchronization precision.

4 The Demo Setup

These musings resulted in a very ambitious proof-of-concept demo involving AES67 equipment from 3 companies (Ross Video, Merging and DirectOut) across 4 sites over two continents (North America and Europe) that leverages the public cloud infrastructure from Amazon Web Services (AWS). The picture below gives a generalized view of the demo setup. Ross equipment (Ottawa, Canada) interfaced with AWS Virginia, while the Merging and DirectOut setups (Grenoble, France; Lausanne, Switzerland and Mittweida, Germany) communicated with AWS Frankfurt (as shown in Figure 1). On-site in Ottawa, Mittweida, Lausanne and Grenoble, various AES67 gear was used to create and receive standard AES67 streams. Gateways on the local networks were used to wrap these AES67 streams into SRT flows which in turn were handed off to the cloud provider's access points using the public Internet. The flows were then transported within the public cloud between the access points, from where they were handed off (secured by SRT) to the local SRT gateways via the public Internet again. The gateways unwrapped the AES67 streams so that they appeared unchanged in the local destination networks and could be received by the AES67 devices.

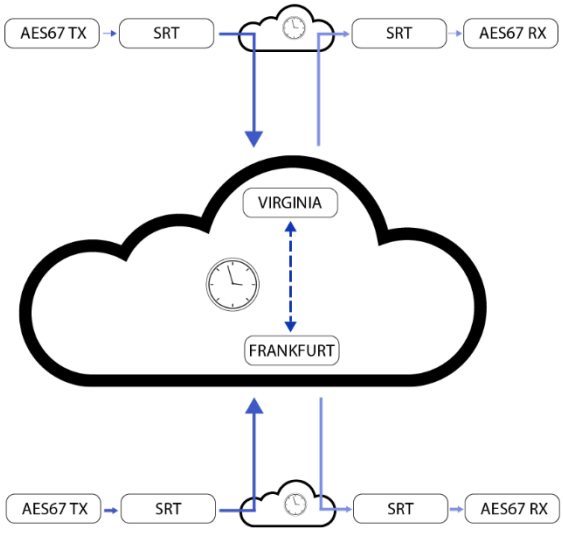

Figure 1. Generalized view of demo setup.

All SRT gateways were built from Haivision's open source SRT implementation. While some of the companies used separate host machines to run the SRT gateways, others were able to include the gateway functionality into their audio equipment.

Since all AES67 devices at each site were synchronized to the same time source via local PTP GMs referenced to GPS, the generated streams received exact RTP timestamps which were transparently transported through the cloud, so that a deterministic and stable play-out latency and inter-stream alignment could be configured at the receiving ends.

Since streams were not processed or altered in the cloud or by the SRT gateways, the audio data was bit-transparently passed through with full quality. Since any packet loss was coped with by the SRT protocol, a higher latency setting needed to be configured to accommodate the larger packet delay variation (PDV) due to occasional packet retransmission (as shown in Figure 2). All the AES67 receiver devices used in this demo provided ample buffering capacity to allow adequate configuration. In practice, buffer settings (which equate to the overall latency setting) ranged from 200 to 600 ms, depending on quality and bandwidth of the local Internet connection.

Typical AES67 audio and stream settings were employed for the demo to guarantee interoperability. Multicast AES67 streams supporting a 1 ms packet time were used to transport 24-bit audio sampled at 48kHz.

A monitoring web page, connected to a local loopback server hosted in the cloud, enabled listening to the live streams via http from within any browser, including display of live VU metering and accumulated (unrecoverable) packet loss per stream.

5 Lessons Learned

The proof-of-concept demo worked well, and we are very pleased with the results. It required some expertise and fiddling with manual settings, which included adjusting items like buffer settings and SRT timeouts, to get it to work. Many lessons were learned from the proof-of-concept; here are a few:

i. "Local only" PTP synchronization locked to GPS works fine.
ii. Even though there is occasional packet loss, it can be managed via SRT and the extra processing penalty wasn't a significant delay factor.
iii. Latency, at significantly less than 1 second, is lower than what we expected, but still substantial. Note that the latency could be further reduced if the local sites were directly connected to the cloud provider's network. In this demo, the streams had to travel hundreds of kilometers over the public Internet, in some cases, to get to the cloud provider's access point.

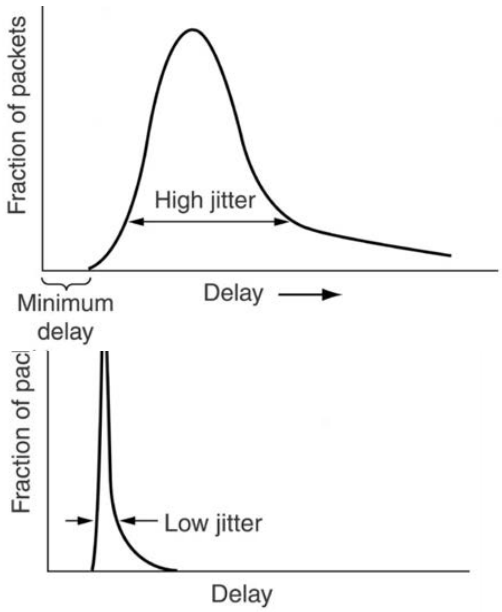

Figure 2. PDV effect on overall latency

iv. To manage increased network delay, manual tuning of the link offset at each location was required, as expected, but the deep buffers of the receivers were able to compensate for it.

6 Future Considerations

There are few items that require further study to make the demo a more practical and usable solution:

i. An important one is how to transport PTP timing through the cloud.
ii. We consciously decided on manual connections using Session Description Protocol (SDP) files to keep things simple. It would be very valuable to be able to use a registration and discovery solution, such as RAVENNA or NMOS, over the cloud to automate the connection process.
iii. Ease of use would be greatly enhanced if the link-offset could be handled automatically to compensate for network delay.
iv. To manage packet loss, it would be interesting to learn if SMPTE ST 2022-7 redundancy would work in similar setups.
v. Although SRT worked great, it would be good to experiment with RIST to understand if there are any performance or reliability benefits.

7 Concluding Remarks

The proof-of-concept (POC) demo clearly showed that it is feasible to transport AES67 across the WAN utilizing the public cloud. The POC also demonstrated that it was possible, using well-designed AES67 devices, to tolerate a fairly big but reasonable amount of latency. Finally, we also confirmed that synchronization via GPS could be maintained across all participating sites despite the lack of PTP within the public cloud. Overall, the POC showed there is a lot of promise for AES67 transport using the public cloud and we are excited and motivated to continue exploring what's possible.

8 Acknowledgments

There were many people who contributed to this ambitious "global" demo and we are grateful for their efforts. In particular, we would like to thank Angelo Santos (Ross) for not only installing and configuring the equipment in Canada but also putting together a diagram of the setup and for his insights; Nicolas Sturmel (Merging) for the installations in France and Switzerland and for programming a monitoring setup in the cloud that was open to the public to view the live results; and Claudio Becker-Foss (DirectOut) for equipment setup in Germany and his thoughts on gateway programming.