



The Internet can be a scary place for your live video

Don't let a bad stream cost you an audience





Don't Let a Bad Stream Cost You Your Audience

As a visual content creator, there are moments when you get that exhilarating, heart pounding, thrill of knowing you are creating compelling content to share with your audience. Capturing these shots is difficult enough – lighting, sound, focus, depth-of-field, angle, and so many other factors converge simultaneously with real-time subject matter that the thought of going live to a large viewer base on the web may seem daunting. The best shots are often from the most difficult places to be live from: court-side, front row, low and tight, high and wide and the creative necessity to be 100% mobile means you will need to rely on WiFi and cellular networks to deliver your creation reliably to your internet viewers at home.

But herein lies the challenge between reliability and ease of use. You have a clearly defined need to deliver your live video content as reliably as possible: 75% of today's online viewers will abandon a poor-quality video experience in less than four minutes*; but still, you can't spend all your time and energy ensuring that the "first mile" of video transmission, that is the delivery from the glass on the camera lens to the CDN or OVP, works flawlessly. We have all experienced issues with poor WiFi networks, and limited cellular bandwidth, and yet now you are going to need to trust these seemingly brittle networks for smooth and stable video delivery. It sounds difficult at best, but thanks to numerous advancements in technology such as connection bonding, adaptive bitrate encoding, and dynamic forward error correction, there can now be peace of mind that your content will be delivered at the highest quality, providing the best possible viewer experience.

The top three frustrations for online viewers:

- Buffering
- Pixelated or fuzzy content
- Slow startup time

Consumers are 62% more likely to have a negative impression of a brand that publishes a stuttering, buffering video

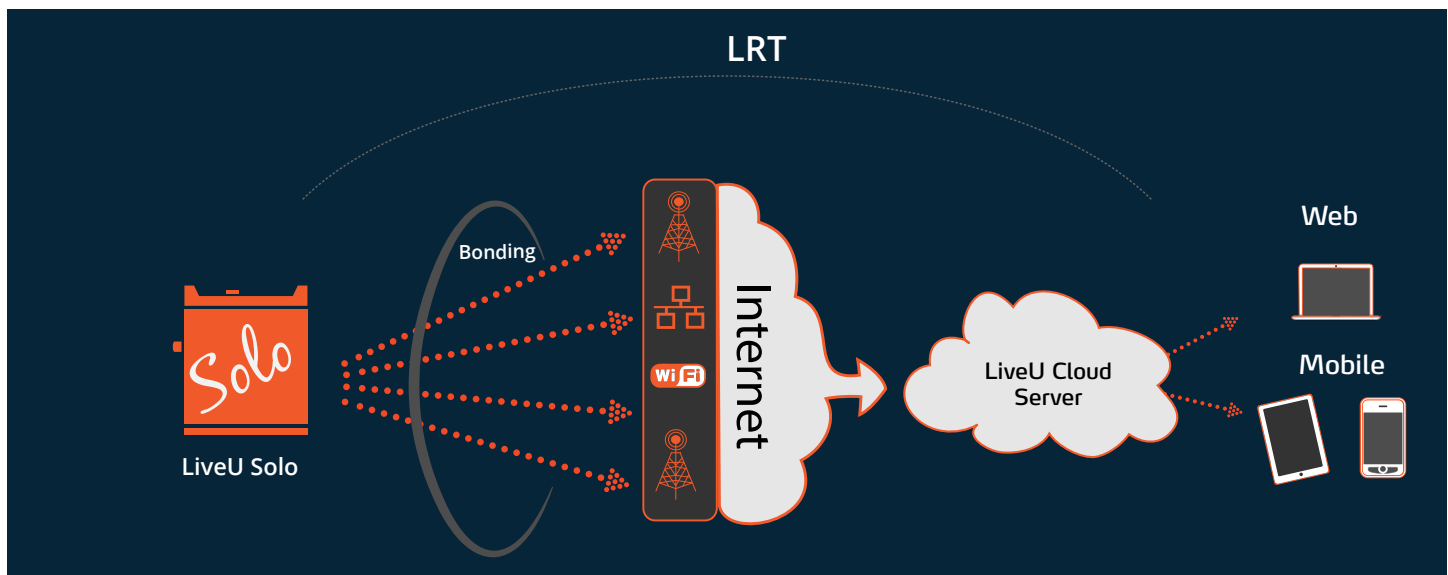
**Source: Conviva "How Customers Judge their Viewing Experience: The Business Implications of Sub-Par OTT Service" Report*

This white paper seeks to outline how the technologies described above can provide content creators a solid foundation on which to deliver their stunning live content from the edge to a distribution platform.

The Audience Experience

As our lives become ever more mobile, an increasing amount of live video content needs to be captured on the fly – and out in the real world. Be it from a drone, a body camera, on the move, or in the air – exciting video can happen anywhere and often in locations where it isn't easy to have a fixed network with reliable bandwidth ready for your use.

You also must consider the audience and how they will experience your content. Increasingly, content is being consumed on non-linear devices: smartphones, tablets, over-the-top set top boxes such as Apple TV, Amazon Fire, and Roku. In order to deliver a high quality viewing experience on these platforms, the “first mile”, or how your content is transported and encoded, must be pristine.



Transport and encoding are two separate issues but both impact how the audience experiences your content. Transporting deals with how bytes get from one place to another; encoding deals with how sounds and images are converted to bytes and back. When cutting the cord to capture content over wireless, your transport engine can't rely on just one connection. Whether using cellular or WiFi, issues with congestion or interruption can ruin your stream. Bonding multiple signals allows for consistency in bandwidth even when one source becomes unstable. To ensure high quality, the encoder needs to be part of this process. When the transmission engine informs the encoder, the encoder can adapt dynamically when bandwidth inevitably increases and decreases to compensate for any potential losses and to automatically help maintain a high quality viewer experience.

Mezzanine Limitations: The stream can never be better than the mezzanine format – it has to be encoded at a higher bitrate than the distribution stream. For example, if you capture a mezzanine stream at 480i, but your distribution is 1080i, you will be unable to distribute at HD resolution.

Streaming directly to an Online Video Platform (OVP) or Content Delivery Network (CDN) introduces further considerations versus streaming to a studio, where the video is made available as a baseband signal. A common practice today is to send a single, high quality bit rate stream to the cloud, often called a mezzanine stream, and then transcode it again into multiple formats and bit rates to facilitate delivery to different devices that have their own delivery requirements. Luckily, there are a multitude of vendors that enable you to reach every viewer

on every device such as YouTube Live, UStream, Brightcove, Wowza, and others. The delivery across these platforms to the viewer is frequently referred to as “the middle mile” and as you might have realized, you, as a content provider, have little control over this part of the live media distribution.

In order to ensure that your content comes out the other end as high quality as it went in, and with as little delay as possible, you need to solve the first-mile problem of getting your video into the cloud over IP-based technologies and mitigate the issues that make the internet such a challenging environment for video streaming. Using advanced integrated encoding technology allows you to match the output bitrate of the encoder to the available bandwidth dynamically, while you are live streaming, and automatically, without any interaction needed from you.

Transporting Your Live Stream

There are two commonly used methods of sending data as packets over IP networks – TCP (Transmission Control Protocol) and UDP (User Datagram Protocol). These two layer 3 protocols differ in how they insure a packet of data made it to its destination. Specifically, TCP sends an acknowledgement every time a packet arrives (called an ack), whereas UDP does not. This makes TCP good for data that must arrive entirely, but often adds latency and additional overhead to accomplish that. Examples of such data include web pages or email. UDP, which does not use acks, is inherently low in latency; since real-time streaming puts a premium on timely delivery, UDP is often used.

But, while UDP has no-overhead transmission, it does not ensure all data arrives. What payloads would be appropriate for a situation where not all data arrives? Audio and video are good examples. This is because there is so much data in audio and video, that losing a few bytes is often inconsequential to actually seeing images and hearing sound.

However, there are limits to this rule of thumb: a little bit of data loss can be acceptable for consumer video and audio, but as data loss increases, it leads to all the things a content streamer hates: broken images, poor audio, buffering, and dropped viewers. Thus there is an inherent conflict in the methods commonly used to stream audio and video: on unpredictable networks, UDP can have big enough data loss to suffer these problems and with TCP, you still have the problems of overhead and latency.



Video suffering from packet loss results in broken images, buffering and dropped viewers

This has led to some additions to the UDP base that close some of the gaps. Some protocols use packet numbering (such as Real-time Transport Protocol – RTP) to be able to put packets back into a proper order if they arrive out of order. Other protocols use mathematical schemes to add some redundancy to the packets (very similar to how RAID works with disk drives). Lastly, some protocols add a form of acknowledgement and resends that don't acknowledge every single packet (like TCP) but acknowledge groups of packets and can resend data that never arrived, if it's critical.

So, while there are different benefits to each method available, nothing seems truly appropriate for high quality online streaming. What you really need is something that incorporates all of these techniques into a single, integrated protocol: TCP feedback, UDP speed, and RTP reliability – all without the delay. You need reliable transport over the most unreliable of networks including cellular networks where the transmitter is moving at high speed. You need all this and one other element that many of the transport level solutions on the market today leave out: tight integration with the encoder. Seems like something unattainable or too complex to exist?

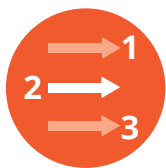
Integrated Reliable Transport Protocol of the Future, Available Today

LiveU has been at the forefront of IP-based live video services for 10 years. As such, LiveU has seen these inherent issues for the online streaming community and developed a Reliable Transport protocol. LiveU Reliable Transport (LRT™) brings together the best of all these techniques into one integrated protocol that works with the encoder so you don't have to do anything. How?



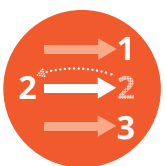
Packet Ordering

LRT uses numbered packets so that the packets can be re-ordered when they arrive out of order. This is a common practice since data often arrives in a different order than it was sent, but is an absolute requirement with connection bonding, where data will almost always arrive in a different order than intended!



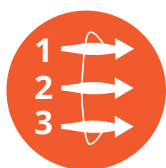
Dynamic Forward Error Correction

That mathematical technique mentioned previously is called Forward Error Correction – it adds some overhead to the stream, with the idea that the small amount of additional overhead can be used to recover lost data faster than a resend. For example, 20% of additional stream bandwidth can result in enough redundancy that entire groups of lost packets can be recovered without ever requesting or waiting for a resend. LRT not only uses FEC, but it uses a dynamic version of FEC. This means that it automatically varies how much FEC to use based on monitored network conditions.



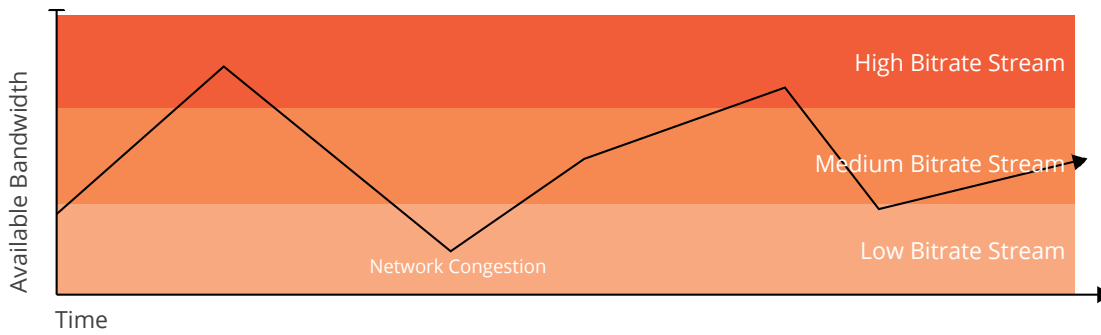
Acknowledge and Resend

LRT uses a form of acknowledge and resend that is appropriate to streaming video and audio. It can acknowledge large groups of packets if they all arrived. If some did not arrive, it can inform the streaming engine to resend needed data. By acknowledging large groups of packets at a time, the overhead and latency of TCP is not re-introduced. Only the packet numbers are used to let the system know what was delivered (or not) so that only the data that is absolutely needed is requested and resent. You also never encounter the main drawback of UDP where you can only hope that your data makes it to its destination. With LRT, you get a complete feedback loop so you know it is consistently arriving.



Tight Encoder Integration or Adaptive Bit Rate Encoding

The last but perhaps most important piece of the LRT protocol is its tight integration with the encoder. As the bandwidth condition changes, LRT automatically recognizes this and informs the video encoder to allow it to adapt the bit rate of video it is delivering and keep the best possible stream within the available bandwidth at any given moment. This ensures you are not trying to push more bits than will fit on the pipe and minimize buffering.



As bandwidth conditions improve, LRT signals the encoder to modify its parameters and increase the bandwidth it is using to stream. In this way, you continue to push the best possible video stream at all times. This backward

flow of information insures that video continues to be sent, and sent at the highest quality and lowest latency possible no matter how unpredictable the network conditions are.

How Scared are You Now?

Freedom and mobility are necessary for creating engaging live content. High quality and reliable streams are essential to the viewers' experience. But erratic networks and unpredictable Internet bandwidth can be a fear of the past with the right reliable transport protocol and encoding system. Especially if it is something that happens automatically without needing your interaction. Bonding greatly improves available bandwidth in the field and reliable transport protocols enhance bonding and quality anywhere. When the transport and encoder engines are integrated, all you have to focus on is getting great content.

All LiveU online streaming solutions allow you to acquire exciting live content in even the most remote locations for your viewers and deliver this content reliably without having to worry. With an LRT integrated solution, the highest quality content streaming is ensured and you can expect the best possible online video experience for your audience – differentiating your content from everyone else.



LiveU Reliable Transport Benefits:

The best of all technologies in one:

- Dynamic Forward Error Correction, Adaptive Bit Rate, Acknowledge and Resend
- Automatic – no interaction is needed
- Complete feedback loop
- Tight integration with the encoder
- Consistency even in the most challenged bandwidth situations
- Robust continuous video over any internet condition
- Broadcast-quality IP video stream delivery over challenging network environments
- Low latency

LRT is at the heart of all LiveU products to deliver optimal video performance.

To learn more about LiveU's unique Reliable Transport protocol or to schedule a demo, contact us here: get.liveu.tv/lrt/



About LiveU

LiveU sets the standard for high-quality and reliable live video acquisition, management and distribution over IP. LiveU's award-winning technology enables live video transmission from any location around the world. With top-tier customers in 60+ countries, LiveU's solutions are being utilized for breaking and developing news and high-profile events, such as the FIFA World Cup™, Winter and Summer Olympic Games, Presidential Campaigns, Super Bowls, US Collegiate Championships and red-carpet events.

LiveU Inc.
2 University Plaza, Suite 505, Hackensack,
New Jersey 07601, USA
Tel: 1-(201)-742-5228

www.liveu.tv/contact-us

www.liveu.tv | www.facebook.com/LiveU.Fans | www.twitter.com/liveu