grass valley

A BELDEN BRAND

Why Shaping Traffic at the Sender is Important

By Chuck Meyer, CTO, Production

December 2017

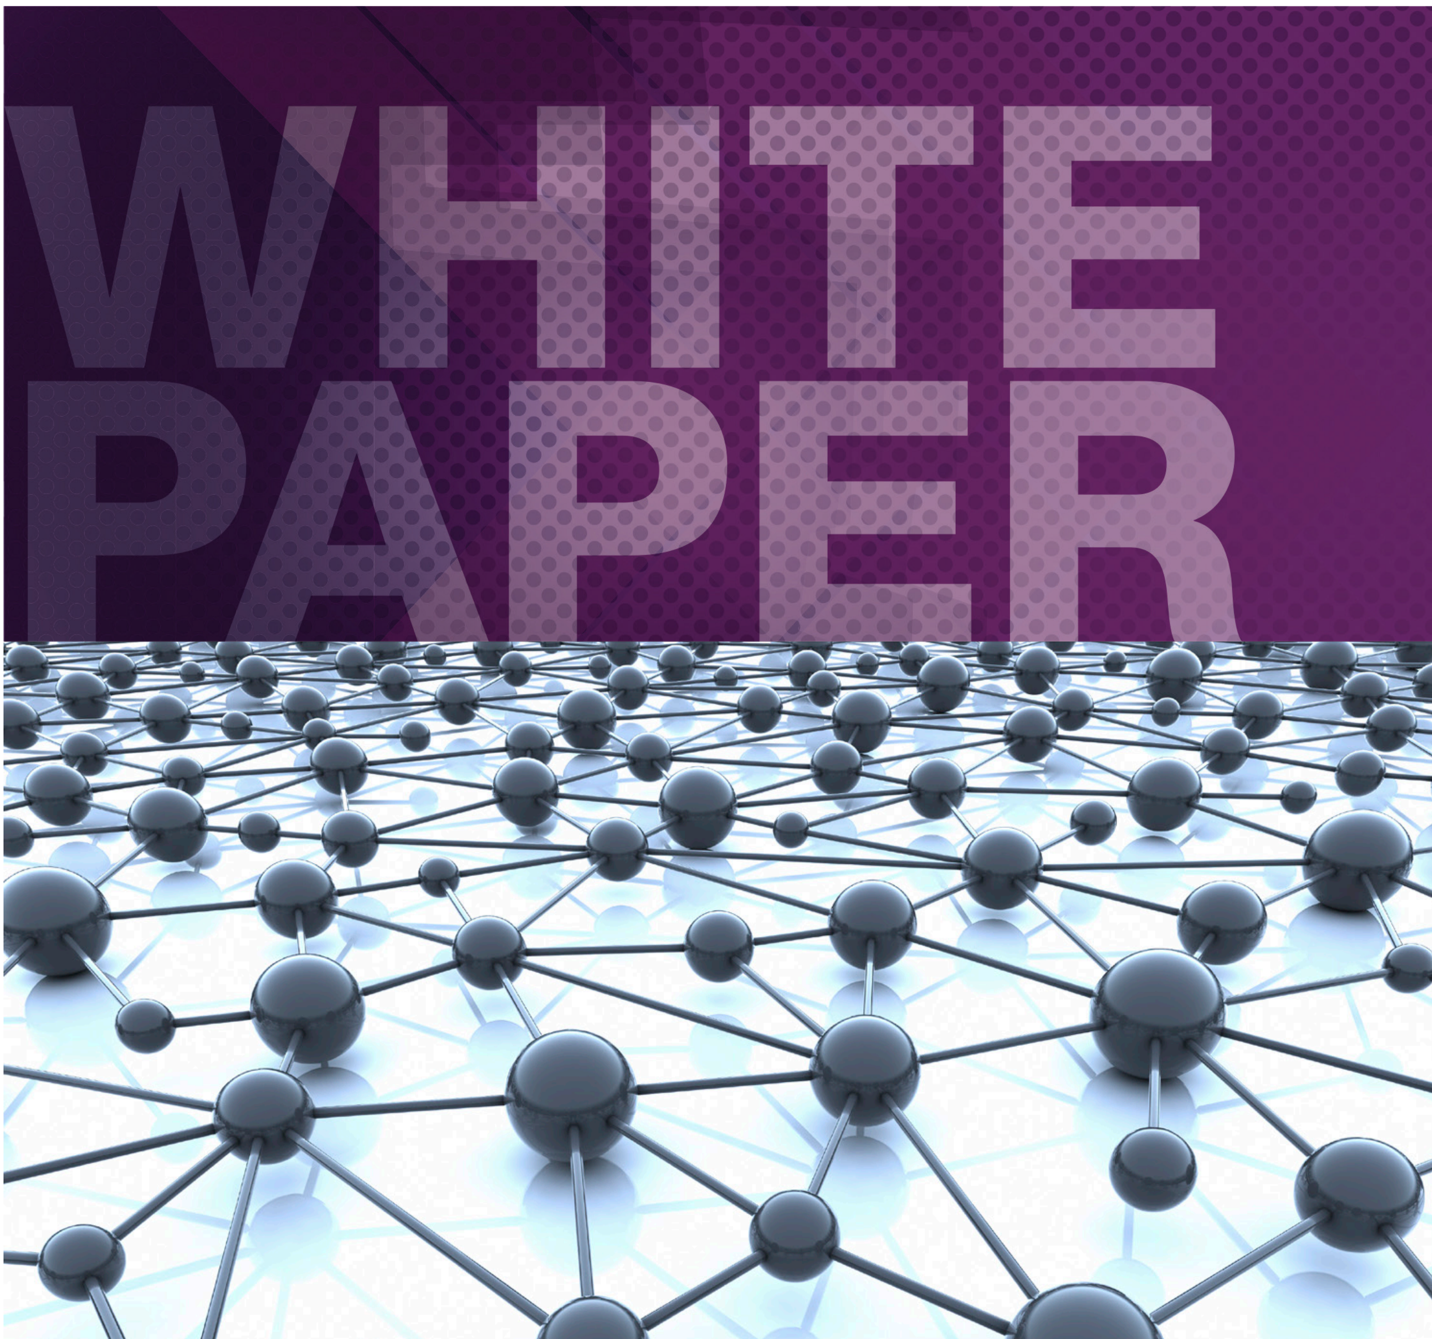

It's human nature to want to utilize a resource or tool to its fullest capacity. So, it stands to reason that broadcasters would expect to get the most out of their router, particularly in live video production. As the industry migrates to more flexible and scalable IP networks that leverage the advantages long enjoyed by the world of IT, maximizing utilization of IP routers requires broadcasters to look at contribution and distribution of signals in a new light.

Modern 10 GigE IP routers provide sufficient bandwidth and the very low latency needed for live content creation, but overutilization can cause problems as IP router design favors managing un-synchronized bursts of data from many data sources, called senders. However, the majority of video senders are typically synchronous, continuously sending streams of high-bandwidth data. This creates a challenge for those seeking to build or upgrade to more modern network architectures.

This challenge is easily addressed when an appropriate scheme of content contribution methods is applied and the data traffic is managed proactively at the source.

IP vs. SDI Bandwidth Management

The nature of IP routers, or switches, is to use memory to smooth out data bursts and manage the routing, or forwarding, of the data between ports. Traditional SDI routers, with only one stream per output, provide connections between inputs and outputs, which can sustain a continuous bitrate indefinitely. IP routers, however, have more than one steam flowing at every port and the time multiplexing process between signals assigned to a given data stream requires some buffering. As a result, some jitter is added to the transmitted packets.

An additional consideration of IP routers is oversubscription, which must be managed port-by-port and for all the ports combined within a switch. If there are too many signals assigned to an output, and their total bandwidth is too large, packets will be dropped. Similarly, if there are too many inputs to the router that require too much total bandwidth, packets will also be dropped.

Modern IP routers manage oversubscription effectively for traditional data, such as email and files, using Quality of Service (QoS) mechanisms to smooth data in the buffers and give priority to the most valuable traffic. Differentiated Services (DiffServ) mechanisms are then used to manage QoS automatically. IP routers also use admission control to manage input bandwidth. But, measuring traffic takes time, and by the time broadcasters discover there is too much traffic, the additional traffic has flowed to a memory buffer that may be full, or nearly full. While not necessarily a problem for traditional data, in high-bandwidth, high-speed video, the result of oversubscription can be unacceptable packet loss.

Sender Traffic Shaping

To manage this risk in advance, broadcasters should consider how the buffers in the switch work and ensure that data senders behave in such a way that the buffers never become too full. This practice is known as sender traffic shaping. To ensure pre-compensating for this adverse effect is efficiently applied to IP, broadcasters must follow one rule: ***the network must be designed and managed for the desired traffic load.*** Logical router control systems, such as GV Convergent, manage IGMP connections between senders and receivers. The system is aware of the bandwidth required for a connection, and simply does not allow the connection to take place if a port will be oversubscribed.

Keeping this in mind, the IP router must be modeled for three different boundary or capacity levels and the following rules must be met:

1. ***The main memory for the switch must not overflow.*** This is a maximum data requirement on total traffic into the switch and memory size of the switch.
2. ***The output port must not overflow.*** This is a maximum data requirement on signal bandwidth at a port, the port's buffer size.
3. ***A data sender, or device, must send packets at a nearly uniform pace.*** If a burst of data has too much bandwidth for too long, buffers may overflow.

The following sequence of figures provides a simple example.

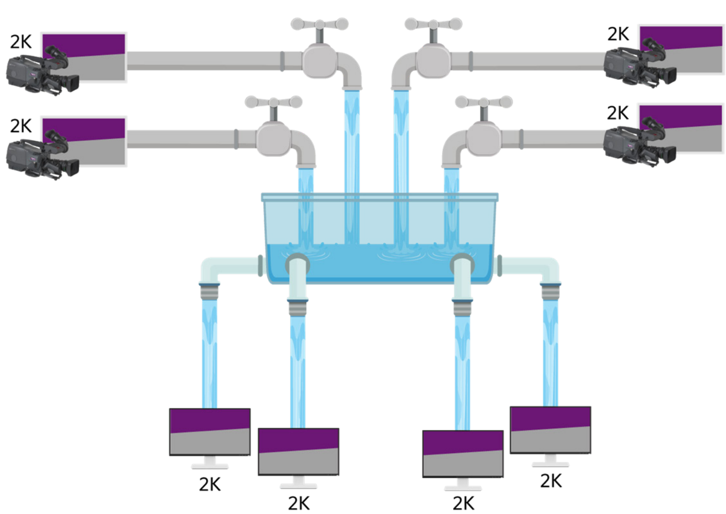

1) The network is designed for a known traffic flow. All the data in is equal to, or less than, the data out. Ideally, this bucket could be nearly empty, and latency is then nearly zero. The video for every source carried error free, to every output.

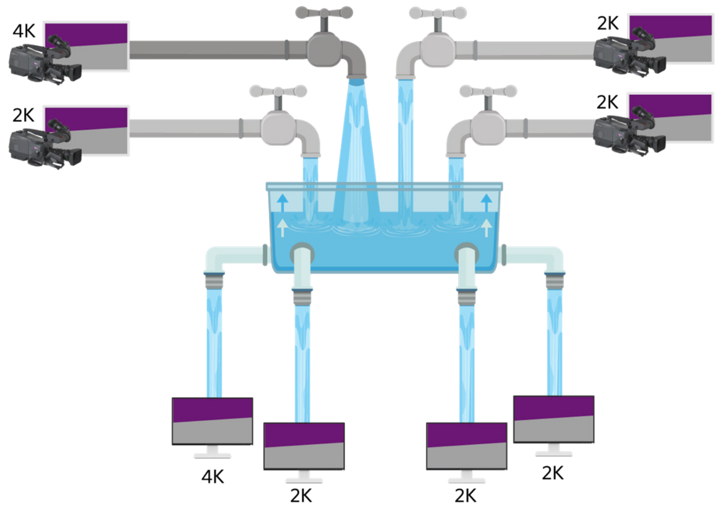

2) The switch has capacity for the maximum data load anticipated. The addition of a 4K source (upper left) should not cause the bucket or the output streams to overflow. Based on design, the bucket may need to hold more water, so latency will increase, but data is not lost. Note the received video is still error free.

3) If an input provides too much data, even if only for a short period of time, the bucket can overflow and data is lost. Data loss can occur in unrelated streams if the data loss occurs in the main switch memory, rather than just the forwarding port buffer. It only takes one sender flooding the switch to cause errors in every received video signal.

Standardizing the Practice

The proposed SMPTE 2110-21 standard for timing sets forth criteria for the sender profile. There are three profiles: Narrow (N), Narrow Linear (NL) and Wide (Wide).

N and NL provide for a nearly smooth transmission of packets with the ideal spacing. The difference between N and NL is that N is used to transmit the video packets as soon as possible from a rasterized source where there is blanking, while NL is used to transmit video packets as soon as possible from a source that is not rasterized. Almost all sources of video today are rasterized, so N facilitates the transition from SDI to IP. NL will be used more and more in the future as ancillary data moves out of the video raster, and video formats no longer use today's raster structure. N and NL are essentially identical from a system router design standpoint because they both require a very fast transmitter capable of sustained throughput. An N transmitter needs to manage packets to within a few microseconds, and N senders are typically implemented in FGPAs, but software-based transmitters which meet the N criteria have been demonstrated (see below).

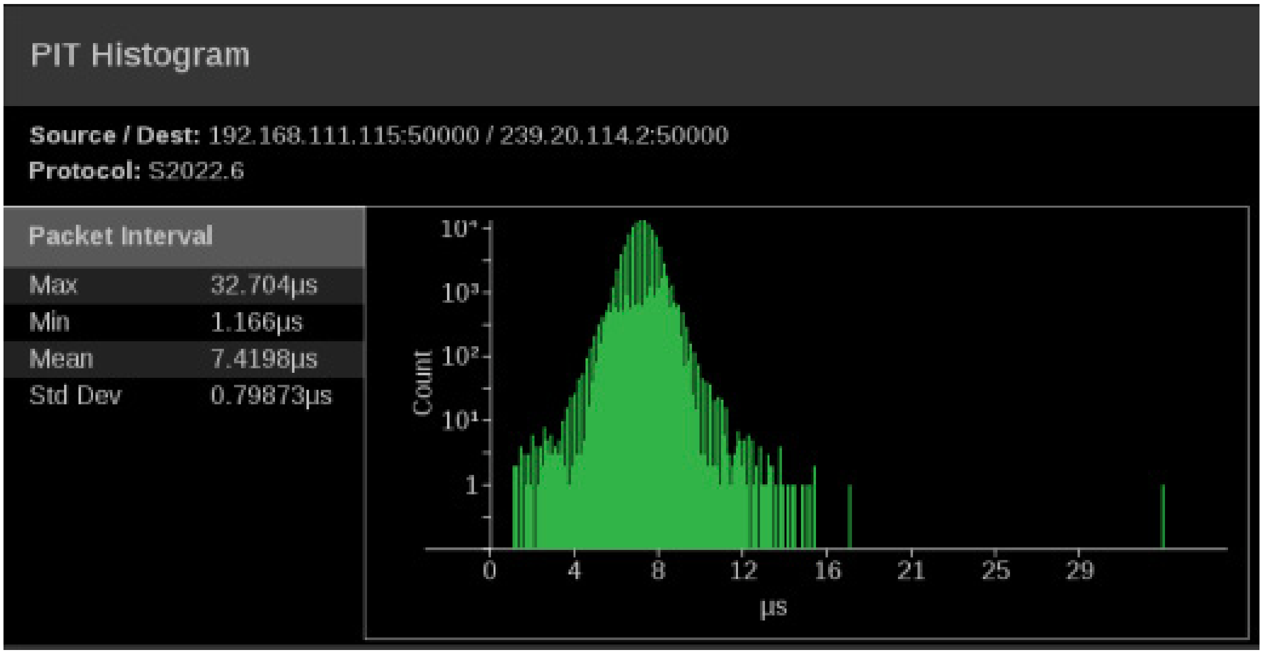

A Grass Valley software implementation of an N sender shown at the Houston VSF Interop, held in February 2017.

The advantages for N and NL are that system latency is reduced, packet loss is essentially eliminated and routers without deep memory buffers may be used. Reduced latency makes N ideal for live production, where there are many hundreds of sources, all of which must be in perfect phase and timing to ensure that human operators have the best possible working environment. An additional benefit for N is that it makes router fabric easy to design. The first step is to count the signals in the plant, and then, almost like SDI, a common value of bandwidth, 1.5 Gb/s for example, can be used. In fact, N was chosen based on simple arithmetic and low-cost 10 GigE wire speed switch silicon used widely in the market today.

The third type, W transmitters, allow for a longer burst of data. This profile is desired by software product developers to enable a wider range of SMPTE 2110-20 senders. W transmitters need only be accurate to about 100 µsec and can be useful in playout where the server needs to run at real time, although some latency is tolerable.

Server NIC cards often have their own memory buffers used to offload data transfers, thereby simplifying the process of streaming output signals. This memory buffer results in longer bursts of higher bandwidth data. In the case of video, senders are often synchronized to a house reference, which means that they emit data at very similar times. If too many W senders transmit at once, and their bursts are too long, then packets could be dropped. Therefore, the IP router and overall system must be more carefully designed. One option is to use a router with additional internal memory. This will stop packet loss, but add more packet jitter attributed to the simultaneous alignment of sender burst and more signal latency.

On average, both senders are running in real time. The difference is that a W system will typically exhibit longer delay and higher Packet Delay Variation (PDV). An N receiver may not accommodate the burst of a W sender, while a W receiver could accommodate an N sender; but a W receiver must be designed to operate with the extra precision required for N senders to operate with the lowest latency. Unfortunately, this may not be possible given the generalized nature of a W receiver design.

Another aspect to consider is that if a system is composed with the intent of N latency and jitter, an unexpected W sender, if not managed well, can cause problems. In the case of back pressure, if a W sender is routed to an output port along with N senders, the W data burst will impact the N traffic profile by introducing excess packet jitter.

If a common switch is used with some W output ports, or output ports that effectively do not care, then the number of W senders, when taken together, must not consume too much of the switch memory, thereby causing the other N output ports to fall out of specification.

Here is a simple example using the proposed SMPTE 2110-21 standard. The key design points from the standard are:

1. A CMAX value of 4 for the N and NL profiles
2. A CMAX value of 16 for the W profile
3. Beta of 1.1

CMAX is analogous to the depth of the bucket, as allocated to each stream of water coming in. It is a number of packets. Beta is essentially a cap on the maximum data for a given physical network link to provide assurance against over-subscription.

Simple arithmetic yields basic design parameters (an exact analysis is included in SMPTE 2110-21). One hypothetical router chip might contain 128 ports and 10 bytes of memory. This 128 port, 10 GigE switch chip can accommodate 384 streams each at 3 Gb/s. This is a total data rate of 1.152 Tb/s, or 144 GB/s. The switch will buffer data to demultiplex, forward and remultiplex the streams.

Data input storage time
= (CMAX * 1540 bytes per/packet * 8 bits/byte) / 3.0 Gb/s
= 16.43 µsec

Total Data Storage per time needed
= Total Input Rate * Beta
= 1.267 Tb/s

Memory size
= (1.267 Tb/s * 16.43 µsec) / 8 bits per byte
= 2.6 MB

The memory utilization is 26 percent, well under full capacity. However, the router algorithms and operations often use this same memory, so using 50 percent as a safe window means that this chip will work.

However, if CMAX = 16 for a W sender, the required memory is 10.4 MB. The channel count will need to be reduced below 384 to ensure no packets are dropped. At this point, it is essential that the router vendor is consulted to determine the best approach. As stated earlier, it may be acceptable to simply add memory to the router, if latency and PDV tolerance can be increased.

Conclusion

Shaping sender traffic is important for system reliability, and it is equally important to consider its impact for low latency and jitter in live production systems. To ensure the optimal workflow and performance, use the proposed SMPTE 2110-21 standard as a way to design the system, and then verify this, as well as system performance requirements, with your router vendor.

GVB-1-0657B-EN-WP