



Powering 8K Video for Next-Generation IP Broadcasting



INTRODUCTION

Next generation high performance IP-based studios are revolutionizing the broadcast industry with regards to 4K Ultra High Definition (UHD) Video, 8K UHD with/without High Dynamic Range (HDR), High Frame Rate (HFR), and other technologies. Consider the following: a single 4K UHD video stream requires more than 10Gb/s network bandwidth; each stream delivers ~1M packets per second (pps), which are processed sequentially by a single program. Given the strict requirement to process data packets at a rate of <1.5sec, the broadcasting application can't instantiate the required number of cores required to achieve it.

Even with the lightweight UDP communication protocol, kernel-based networking challenges can become a primary scale-limiting factor. The kernel-based networking software overhead often fails to deliver the required speed, due to the core's (~8Gb/s) performance ceiling. As bandwidth increases, the required CPU utilization for processing packets results in higher latency (~100us) and an increase in the number of cores.

Kernel Bypass to the Rescue

Video processing is a very CPU-intensive operation, especially when also having to deal with networking activities like sending and receiving packets from the network NIC. A key way to improve CPU efficiency and receive performance gains is to use "kernel bypass" – a technique that involves offloading the sophisticated I/O processing from the CPU to a smarter network adapter.

Over the years, several kernel bypass solutions have emerged including ZF_COPY, DPDK, and RDMA. More recently, Grass Valley, a Belden Brand, leveraged Mellanox networking and kernel bypass technologies to advance Grass Valley's iTX Integrated Playout Platform. The integrated solution delivers more than 10Gb/s playout throughput performance, packet-paced on Microsoft Windows with a single CPU core for 4K UHD streaming.

Mellanox's Rivermax™

Media & Entertainment (M&E) industry leaders are taking advantage of application acceleration to streamline the costs and complexities of their broadcasting operations. Mellanox's Rivermax is a user-space Linux/Windows library and kernel bypass solution that optimizes video processing by offloading network processing from the CPU. Bypassing the kernel and IP stack minimizes context switches, buffer copies and interrupts, resulting in extremely high bandwidth and low latency.

Rivermax Technical Highlights

- Enhanced video streaming performance of network-heavy applications over Ethernet networks – Rivermax provides a specially designed API for the application, tailored to today's IP-standards. Rivermax enables cutting through the entire IP stack and offloading some of the non-networking application parts.
- IP-Specification compliance – Rivermax enables SMPTE ST2110-21 standard compliance in NIC hardware, offloading all packet-pacing and network transmission handling.
- Selective kernel bypass – The kernel handles Address Resolution Protocol (ARP), Internet Group Management Protocol (IGMP) and other traffic, which eliminates the need for a full complex network stack in user space, while Rivermax handles the latency and bandwidth critical data path.
- Designed and optimized for broadcast and streaming applications – Rivermax offloads sophisticated network processing to the adapter, allowing applications to deal with lines/frames instead of individual packets. This results in reducing CPU overhead, thereby improving application efficiency even more; it also delivers the highest bandwidth and lowest latency in the industry, freeing up CPU cycles for the application to perform more video processing.

For particular networking requirements, Mellanox ConnectX-5® series adapters with Rivermax support provide connectivity for up to 100GbE networks in addition to all the popular kernel bypass solutions.

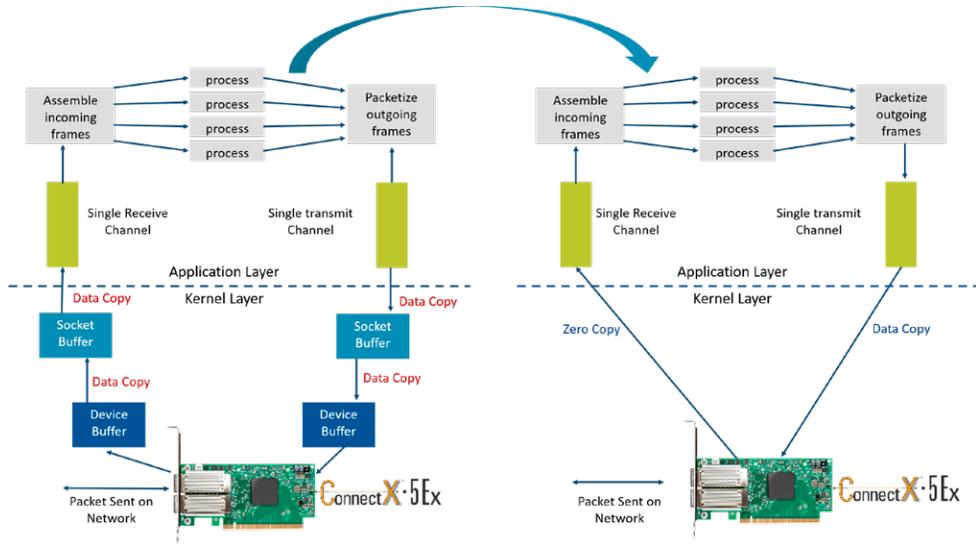


Figure 1: Zero-Copy with Mellanox Rivermax

Smart Buffer Management

Rivermax offers a unique mechanism through which applications compose video frames efficiently, using Mellanox ConnectX-5 series adapters. With Rivermax, applications can create a large cyclic buffer for each video flow, accommodating multiple video frames in sequence.

Mellanox’s adapters, in turn, store the actual video frames (packet payload) in the buffer while processing and striping the packet header information, such as UDP, IP and MAC, or, alternatively storing these headers in another buffer. By doing so, the application can access the entire video segment from one contiguous buffer as shown in Figure 2.

Furthermore, Rivermax optimizes buffer usage by enabling both the application and adapter to share and reuse the same buffer. With this capability, an application can receive video streams at 100Gb/s bandwidth using just a single Intel® Xeon® E5 server.

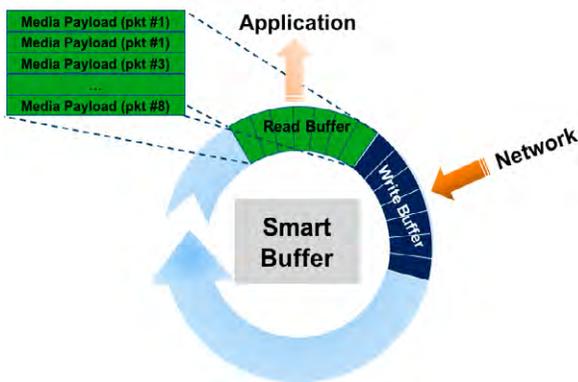


Figure 2: Smart Buffer Management with Mellanox Rivermax

Advanced Steering Technology

With advanced flow steering technology, Rivermax enables applications to handle each video flow in a different buffer. In the case of multiple incoming low-bit-rate streams, Rivermax enables applications to use a single mixed buffer for multiple media streams. Rivermax does this by assigning each packet a unique flow ID tag that identifies which flow it belongs to. In addition, Rivermax performs packet sorting/dispatching by means of the packet’s unique flow ID, offloading the application from having to parse packet headers to identify the packet’s stream, dramatically reducing the CPU overhead (Figure 3).

Example: a multi-viewer running an Intel® Xeon® E5 processor can receive 64 x ST2022-6 or 64 x ST2110 video streams with a sustained bandwidth of 100Gb/s. In addition, the ConnectX-5 Network Adapter hardware is configured to check the IP header including checksum verification, as well as the UDP checksum.

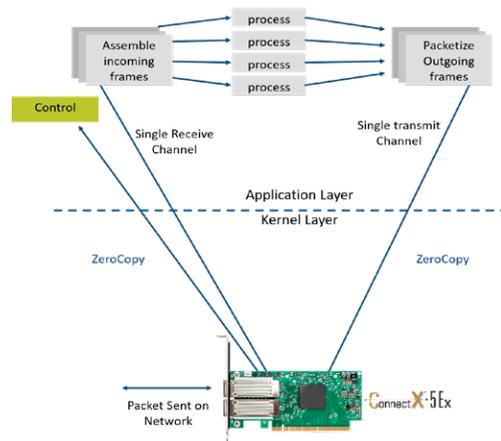


Figure 3: Smart Protocol Offload Enables Kernel Bypass

Smart Protocol Offload

Rivermax offloads video protocol as video packet headers are stripped and stored in a single contiguous virtual buffer. As described below, the simplified software logic checks if a packet is missing.

Sync for frame start, where you know that the video frame contains <N> packets for every frame:

1. InSync = false
2. For every packet:
 - 2.1 If last packet in a frame:
 - 2.1.1 NextPSN = currentPacket.PSN+1
 - 2.1.2 Go to next packet
 - 2.1.3 Jump to 3.3
- 3 For every video frame:
 - 3.1 NextPSN = NextPSN + N
 - 3.2 Increment the pointer by N*Header size
 - 3.3 If (currentPacket.PSN != NextPSN)
 - 3.3.1 Declare packet loss event, frame is incomplete
 - 3.3.2 Go to 1
 - 3.4 If InSync == True
 - 3.4.1 Declared complete video frame received
 - 3.5 Else
 - 3.5.1 InSync=true
 - 3.6 Go to 3.1

This software implementation is format-agnostic, e.g., UHD-1 and UHD-2, as the complexity is dependent on the number of frames per seconds, and not on the content.

Congestion-Free Streaming with Packet Pacing

Video broadcasting and streaming applications are bursty and bandwidth-intensive in nature, and thus can easily cause network congestion. Packet pacing, and specifically SMPTE ST 2110-21, overcomes this challenge by limiting the bandwidth for each flow from the video origin server. With packet pacing, traffic is evenly spaced out; so there is minimal queuing until the load matches the bandwidth. Instead of transmitting packets immediately upon receipt of an acknowledgement, the sender spreads packet transmission, defining both the max burst size to send and the rates. The SMPTE ST 2110-21-compliant ConnectX-5 implements packet pacing in its hardware, relieving the software to deal with application-level tasks.

Take advantage today of Mellanox ConnectX-5 10/25/40/50/100 Gb Ethernet interconnect adapters and Rivermax technology to future-proof your broadcast networking platforms for tomorrow's demands. Visit www.Mellanox.com.

RIVERMAX KEY BENEFITS

- Simple API set for easy application integration
- Lowest CPU overhead with highest bandwidth: 1.5% CPU core utilization per 1.56Gb/s
- Support for Linux and Microsoft Windows
- Designed specifically for broadcasting and video streaming applications
- Delivers video processing acceleration
- Offloads application handling of packets to the NIC hardware keeping application at line(s)/frames level
- Stream video at 100Gb/s with a single core of a Xeon® E5 core processor
- SMPTE 2110-21 compliance with NIC packet pacing - even at speeds of 100GbE
- Cloud ready (VM/Containers/SRIOV)

CONCLUSION

With the advent of broadcasters migrating to IP infrastructures handling next-generation video workloads, kernel bypass and hardware offloads are playing a key role in overcoming native Linux/Windows kernel bottlenecks. Even greater benefits are achieved with deeper application integration.

In recent years, Mellanox Technologies has been working alongside major broadcasters to help define and deliver the next generation IP studio. Mellanox Rivermax is enabling these broadcasting studios to enrich the video experience for their users.

By implementing Rivermax technology, video processing at the required throughput is achievable, while freeing more processing power to the primary application and shortening ROI times.

