WHITE PAPER

Powering 8K Video for Next-Generation IP Broadcasting

INTRODUCTION

Next generation high performance IP-based studios are revolutionizing the broadcast industry. This trend is even more apparent when it comes to 4K/Ultra High Definition Video (UHDV), 8K with/without High Dynamic Range (HDR), High Frame Rate (HFR), and other technologies.

For example, a single 4K (UHD-1) video stream requires more than 10Gb/s of network bandwidth; each stream delivers ~1M packets per second, 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 would not be able to instantiate the required number of cores to achieve it.

Even with the lightweight UDP communication protocol, kernel-based networking challenges can become the primary limiting factor for scale. 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, so does the required CPU utilization for processing the packets, resulting in higher latency (~100µs) and an increase in the number of cores.

Kernel Bypass to the Rescue

Video processing is a very CPU-intensive operation, while also having to deal with networking activities 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, VMA, and RDMA – the latter arguably one of the most popular. More recently, the BBC together with Mellanox developed an optimized Netmap implementation, leveraging Mellanox networking and kernel bypass technologies in advancing BBC's video networking technology, within their on-going IP Studio development project.

Mellanox VMA: Accelerating Your Application

Today leading companies in the Media & Entertainment (M&E) industry are using Mellanox VMA for application acceleration, and to streamline the costs and complexities of their broadcasting operations.

Mellanox VMA Messaging Accelerator is a dynamically-linked, user-space, socket-based, Linux 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 low latency.

VMA Technical Highlights

- VMA enhances the performance of socket-based networking-heavy applications over Ethernet networks. The standard BSD socket interface is the starting point for fast and easy integration with the applications, without requiring any application or OS changes the VMA allows for cutting through the entire IP stack, and offloading some of the non-networking application parts.
- Selective kernel bypass, where the kernel may handle traffic such as Address Resolution Protocol (ARP) and Internet Group Management Protocol (IGMP), eliminates the need for a full complex network stack in the user space. While VMA handles the latency and bandwidth critical data path.
- Designed and optimized for broadcast and streaming applications, VMA offloads sophisticated network processing to the adapter, instead of running it in the kernel. Extending socket API allowing zero copy, VMA dramatically reduces the CPU overhead (thereby improving application efficiency even more), while also delivering the highest bandwidth and lowest latency in the industry freeing CPU cycles for the application to do more pure video processing.

For particular networking requirements, Mellanox Connect-X® series adapters with VMA support, provide connectivity for up to 100GbE networks and all the popular kernel bypass solutions.

† For illustration only. Actual products may vary.

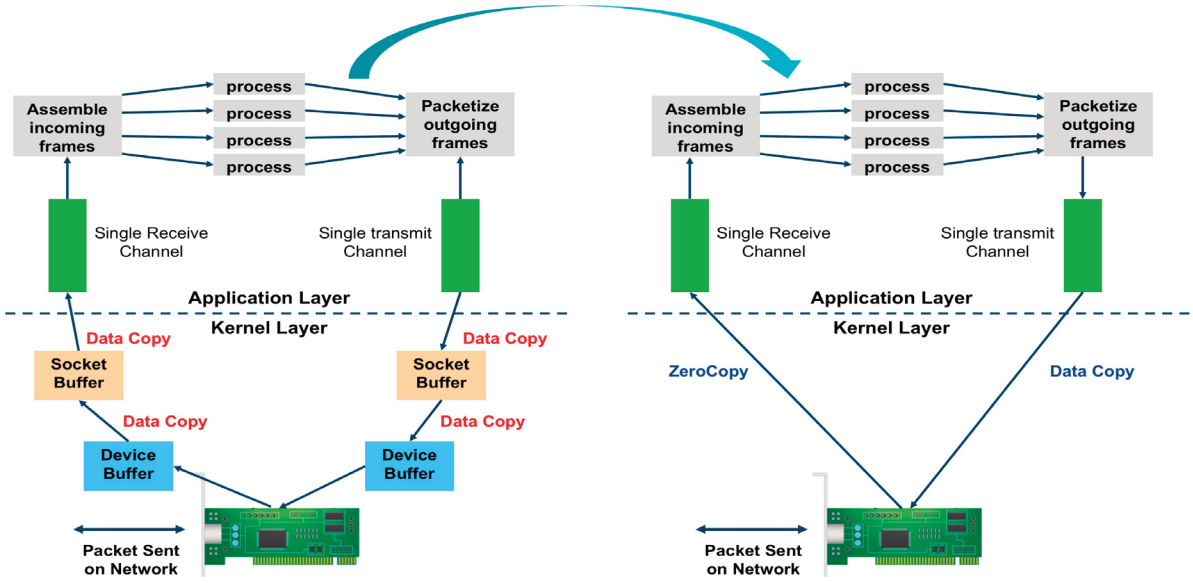

Figure 1: Zero-Copy with Mellanox VMA

The latest Mellanox VMA 8.3 release delivers several optimizations specifically designed for broadcasting and video streaming applications. These include:

Smart Buffer Management

VMA offers a unique mechanism through which applications compose video frames efficiently, using Mellanox Connect-X series adapters. With extended API support, 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 RTP, UDP, IP and MAC, or, alternatively storing these headers in another buffer. This way, the application can access the entire video segment from one contiguous buffer as shown in Figure 2.

Furthermore, VMA 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 core of a server.

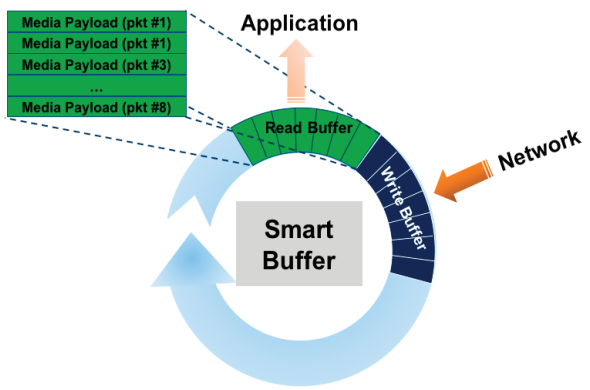

Figure 2: Smart Buffer Management with Mellanox VMA

Advanced Steering Technology

With advanced flow steering technology, VMA enables applications to handle each video flow in a different buffer, thereby dramatically reducing the CPU overhead for sorting/identifying the packet origin and destination. For example, a multi-viewer running an Intel® Xeon® E5 processor can handle 64 x ST2022-6 or 64 x ST2110 video streams with a sustained bandwidth of 100Gb/s. In addition, the Connect-X Network Adapter hardware is configured to check the IP header including checksum verification, as well as the UDP checksum.

Smart Protocol Offload

VMA offloads video protocol as video packet headers are striped and stored in a single contiguous virtual buffer, the software logic that checks if a packet is missing is simplified and works as described below.

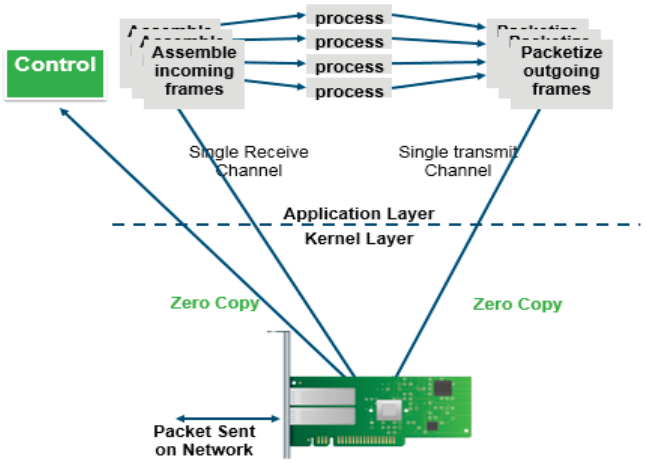

Figure 3: Smart Protocol Offload Enables Kernel Bypass

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

1. For every packet:
    - 1.1 If last packet in a frame:
        - 1.1.1 Go to next packet
        - 1.1.2 NextPSN = currentPacket.PSN+1
        - 1.1.3 Jump to 2.3
2. For every video frame:
    - 2.1 Increment the pointer by N*Header size
    - 2.2 Expected PSN = NextPSN + N
    - 2.3 If (currentPacket.PSN != NextPSN)
        - 2.3.1 Declare packet loss event, frame is incomplete
        - 2.3.2 Go to 1
    - 2.4 Declared complete video frame received
    - 2.5 Go to 2.1

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 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 TCP window for how much to send and the rates.

Mellanox Future Developments

The Mellanox roadmap will bring continuous development of products and solutions designed to improve and accelerate video application and flow. Examples will include a Tx side offload as well as new silicon features in the upcoming Mellanox ConnectX® products.

MELLANOX VMA KEY BENEFITS

- Simple API set for seamless application integration with minor application changes
- Lowest CPU overhead with highest bandwidth: 1.5% CPU core utilization per 1.56Gb/s
- Designed specifically for broadcasting and video streaming applications
- Delivers video processing acceleration
- VMA Smart Buffer: Offloads application handling of packets to the NIC hardware providing the application with a contiguous multi-frame video buffer, while separating data from headers
- Stream video at 100Gb/s with a single core of a Xeon® E5 core processor

CONCLUSION

With the advent of broadcasters migrating to IP infrastructures handling next-generation video workloads, kernel bypass is playing a key role in overcoming native Linux 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 VMA is enabling these broadcasting studios to enrich the video experience for their users, without requiring any application changes, and even allowing application code removal. By implementing open source VMA technology, video processing at the required throughput is achievable, while freeing more processing power to the primary application and shortening ROI times.

Take advantage today of Mellanox ConnectX-5 100G InfiniBand and Ethernet interconnect adapter, and VMA technology, to future-proof your broadcast networking platforms for tomorrow's demands. Visit www.Mellanox.com.

350 Oakmead Parkway, Suite 100, Sunnyvale, CA 94085
Tel: 408-970-3400 • Fax: 408-970-3403
www.mellanox.com

53248WP
Rev 1.1