WHITE PAPER

January 2017

# Powering Next-Gen IP Broadcasting with Mellanox Efficient Network

## Contents



## Introduction

Next generation high performance IP-based studios are currently revolutionizing the broadcast industry specifically in the broadcast production arena. This is due to the complexity and bespoken nature of the less efficient Serial Digital Interface (SDI) that's used today. Migrating to an IP-based infrastructure empowers the broadcasters to innovate in all the areas of content creation & distribution, multi-platform support and future video formats. With the emergence of open networking and commercial off the shelf (COTS) solutions, the same technology that powers today's datacenters and cloud, provides the underlying foundation for this innovation.

As the shift to an all IP-based studios takes hold, larger studios are looking to deploy on IP instead of SDI. The joint collaboration across several organizations, led by the Advanced Media Workflows Association (AMWA) and the Joint Task Force on Networked Media (JT-NM), has accelerated both the innovation and adoption of IP-based studios. This has paved the way to multiple new standards from organizations such as the Society for Motion Picture and Television Engineers (SMPTE), Video Service Forum (VSF) and Audio Engineering Society (AES) like:

- Wrapping SDI in its entirety within an IP packet (SMPTE 2022-6)
- Discovery & Registration Protocol (AMWA IS-04)
- Separating out the essence flows1[1] as multiple IP streams (TR03 (ST2110), TR04)
- Managing successful time synchronized transport (PTP – RFC1588, SMPTE2059)

However, as with any technological innovation, these new standards bring multiple levels of complexity that broadcast engineers will need to understand and adopt.

1 In broadcasting, all media content consists of two parts: essence and metadata. Essence data is the actual video and audio media. This data can be one or more of a variety of essence data types including compressed or uncompressed video and any digital audio source

## Broadcast Challenges With IP Network:

There are fundamental differences in packet (IP) vs. circuit (SDI) switched solutions. Hence delivering a solution with this paradigm shift means careful adoption of broadcast workflow on IP networks. Historically, IP networking fabrics have been designed to carry IT traffic, that have a different set of needs than with broadcast media. Quality of Service (QoS), buffers and pause frames all create variencies in network latency. Both QoS and pause frames with PFC (Priority Flow Control) are designed to address contention and allows flows to be prioritized accordingly, while buffers hold packets whilst the network is congested.

SDI is a circuit based switching solution, so switching sources in SDI is both immediate and instantaneous. However, the newer generation video formats like 4K and 8K and broadcast requirements such as High Frame Rate (HFR) and High Dynamic Range (HDR) demand a new architectural approach on a congestion free network. As video flows (IP packets) traverse the IP fabric, they are interspersed with other flows that have the same precedence. In a dedicated purpose built broadcast IP plant, the environment can be controlled to ensure contgestion is managed and video delivery is guaranteed. In more publicly shared networks such as those found at large sporting events, controlling who gets access to which resources and when can be a challenge.

The successful transport of audio and video requires predictability, reliability and determinism not seen and largely not required in the traditional IT based environments. Being prescriptive on how the IP resources are used; mitigating the effect on switch buffers by rate limiting multiple synchronized senders; ensuring evenly distributed traffic across the fabric are some of the technical recommendations for the IP broadcast plant.

Packets need to be evenly dispersed or distributed over the available network resources to ensure correct load distribution, addressing both contention and fairness. With a plethora of variables available with video such as compressed, uncompressed, HD-SDI, 3G-SDI, 4K, 8K, HFR and HDR high bandwidth and a range of port speeds is essential. Building out a scalable fabric that is non-blocking and where traffic for all the viewers are fairly distributed is the key to successful adoption of an IP broadcast infrastructure.

## Mellanox's Solution to the Next-Gen Broadcasting Infrastructure

For the past two years, Mellanox has been leading the way with JT-NM, AMWA and SMPTE interoperability testing and technical discussions and working alongside the major broadcasters such as Fox, BBC and NBC, helping define and deliver the next generation IP Studio.

Mellanox Technologies is a leading supplier of end-to-end Ethernet and intelligent interconnect solutions. Mellanox's industry-leading technologies improve IP Studio efficiency by providing the highest throughput and lowest latency, delivering data faster to endpoints and unlocking system performance. Mellanox offers a choice of high performance solutions: network and multicore processors, network adapters, switches, cables, software and silicon that can accelerate and balance throughput and maximize results.

Not all technology is made equal. Mellanox have the technologies to provide the end-to-end IP-based studio, as follows:

### Mellanox Low Latency Switches

Mellanox provide a number of fixed form factor switching platforms based on the same silicon, so the performance is known, tested and proven from one platform to another.

- Multiple Bandwidth connectivity (1G, 10G, 25G, 40G, 50G and 100G):
  Supporting the throughput required for all video requirement including 4K, 8K, HFR and HDR, whether compressed or not. 25Gb Ethernet delivers 4K uncompressed to the client while 100GbE will be used for a high-speed switch interconnect.
- Predictable buffer allocation:
  Flexible large buffer pool that is available to all ports on the switch rather than being divided into groups of ports. Flexible switch buffers with zero-packet loss allows storing of bursty traffic which is prevalent in broadcast applications and ensures that all packets are forwarded as required.

- QoS with DSCP marking:
  Allows the marking of important flows to assist with ensuring a non-blocking fabric.
- Predictable Network Performance:
  Both previous generation SwitchX-2-based and current generation Spectrum-based switches consistently deliver predictable performance for all packet size. This has been described and well documented in the Tolly test reports.

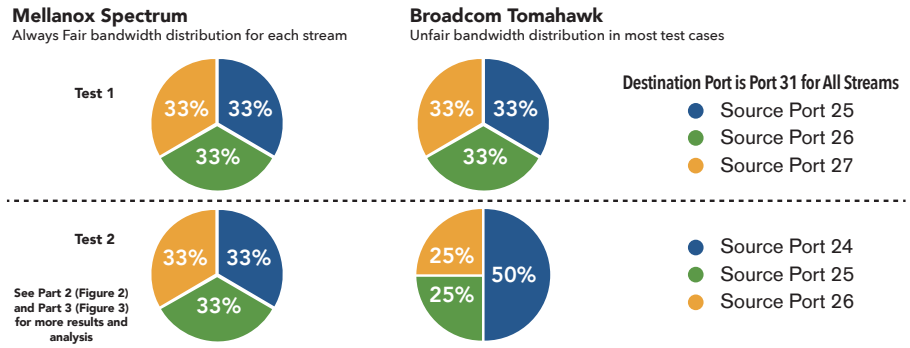


***Figure 1.*** *Predictable Network with Mellanox Spectrum*

- Consistent and very low port-to-port latency and jitter:
  Mellanox switches have lowest port-to-port latency and jitter in the industry. Alongside other IP switch vendors, this has been tested by both FOX Networks Engineering and Operations and confirmed by Snell Advanced Media (S-A-M) . During these tests, Fox concluded that the Mellanox platforms proved to be the clear leader in switch silicon performance. Both port-to-port latency and Packet Delay Variation (PDV) was tested and Mellanox proved to have the lowest and most deterministic latency out of all the vendors under test. As the IP based broadcast studio scales with a Mellanox fabric it does so with known and well understood constants. Fox and Aperi Corporation had earlier presented their test results across other switch vendors at 2014 SMPTE Annual Technical Conference and Exhibition.

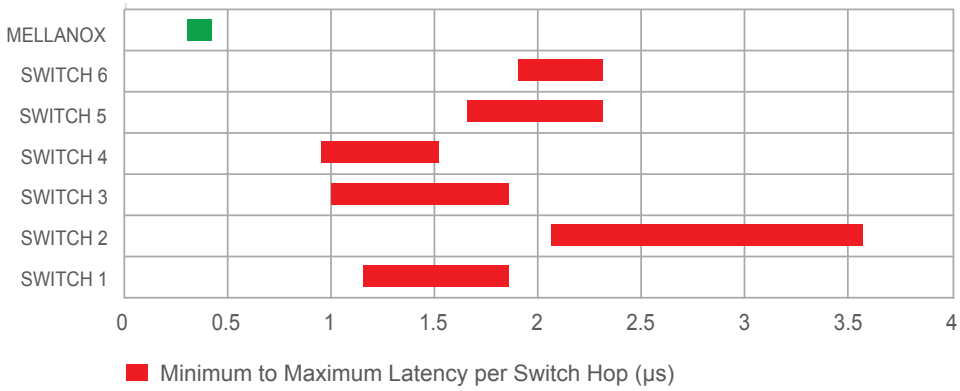


***Figure 2.*** *Port-to-Port (Switch-to-Switch) Latency*

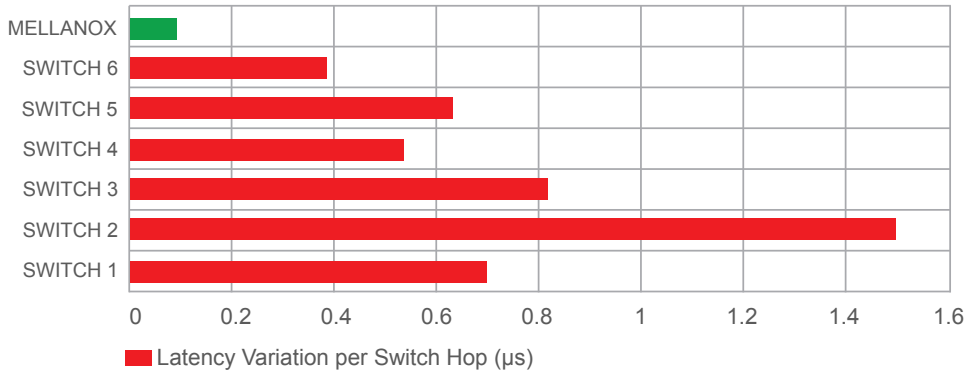


***Figure 3.*** *Packet Delay Variation (Jitter)*

- In fabric containerized broadcast services:
  By containerizing IP studio services and running them on the switch allows broadcast engineers can focus on building ideal IP media fabric for their studio without utilizing additional servers and virtual machines. For example, Mellanox can build an NMOS aware switch with both the NMOS Registration and Discovery service and the NMOS aware Connection Manager running within it.

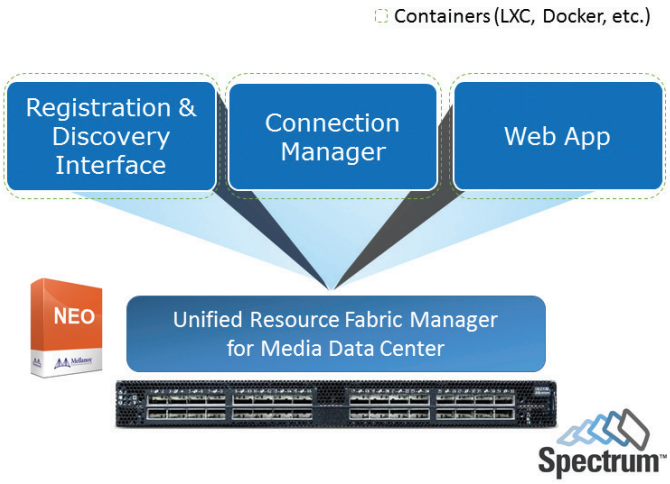


***Figure 4.*** *Containerizing IP Studio services on Spectrum Switches*

- OpenFlow Integration:
  With OpenFlow, Spectrum switches can allow seamless configuration for Multicast-to-Unicast and Unicast-to-Multicast translation, stream duplication, and NAT-based stream switching. Spectrum supports the latest OpenFlow 1.3 standard which allows the manipulation of flows within a switch and across the fabric. With it's RESTful API, OpenFlow provides dynamic control of how broadcast flows are forwarded throughout the networking fabric. Further, Spectrum allows support for multiple configurations: OpenFlow only, a port-based hybrid and a protocol-based traditional mode of operation. It is also integrated with Open Daylight (ODL), Open Network Operating System (ONOS) and other controllers.

## Mellanox Network Adapters

The Network Interface Adapter (NIC) plays a critical part in delivering predictable and high performance broadcast flows onto the network.

- Mutiple Bandwidth Connectivitiy: 1, 10, 25, 40, 50, 100GbE
- Kernel bypass:
  Mellanox ConnectX-based NICs support multiple kernel bypass solutions including RoCE (RDMA over Converged Ethernet), Netmap, VMA and DPDK (Data Plane Development Kit). This provides faster packet processing, enabling lower latency and higher throughput. Kernel bypass technologies reduce the overhead of context switches and memory copy for I/O processing. This returns more CPU cycles to the application, lowering jitter and improving throughput.

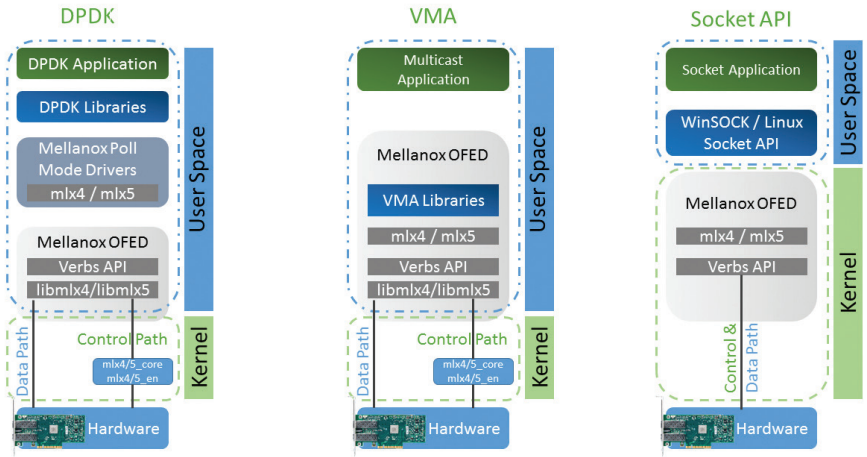


***Figure 5.*** Mellanox Kernel Bypass Solutions

**Netmap**: Line-rate raw packet I/O from user space, developed by the University of Pisa

**Mellanox VMA Messaging Accelerator:** A dynamically-linked user-space Linux library for accelerating UDP based multicast traffic to offload network processing from the CPU. Bypassing the kernel and IP stack minimizes context switches, buffer copies and interrupts, resulting in extremely low latency.

**Data Plane Development Kit (DPDK):** A set of data plane libraries and network interface controller drivers for fast packet processing

By utilsing implementation features withinthe adapter such as above, we can choose which data uses the kernel stack and which bypasses the kernel.

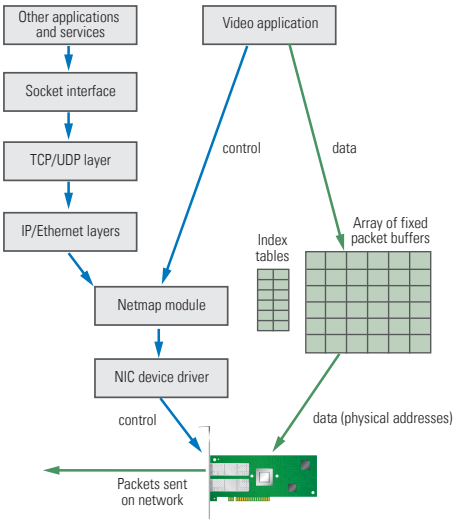


***Figure 6.*** *Optimized Netmap solution, jointly developed with BBC*

- Packet Pacing to prevent network congestion:
  Packet Pacing overcomes the challenge where multiple synchronized streams all send data at the same time thereby clashing and overflowing switch buffers. Envisage an IP network with a multitude of bursty senders- as the video streams to the network synchronizes, the volumes of bursty packet also synchronizes which may cause port blocking through buffer exhaustion. Within a non-blocking fabric this is something that needs to be addressed and eliminated in both the server and the switch. By pacing the flows out from all of the senders on the server, this issue can be mitigated.

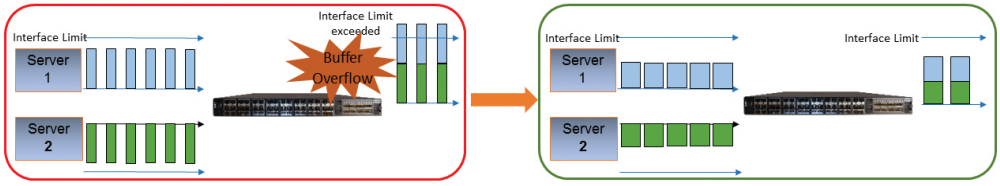


***Figure 7.*** *Prevent Network Congestion with Packet Pacing*

- Hardware timestamping:
  The accuracy of IEEE1588 time based synchronization is dependent on the accuracy of the available clock and where suitable timestamps are terminated. These time stamps in turn are used to calculate the clock skew/offset from real time. Mellanox hardware time stamping overcomes the variances introduce by the software stack on server hardware. By terminating the timestamps at the hardware NIC Mellanoc adapters can eliminate the highly unpredictable jitter seen in software based solutions such.

## Mellanox Optical and Copper Cables

Transmitting video over any fabric requires a very high degree of performance and accuracy. All Mellanox Interconnects are built to a very high standard to support and maintain our company mantra of high speed and low latency. The passive copper and active fibre (VCSEL [Vertical-Cavity Surface-Emitting Laser] Silicon Photonics) cables as well as the optical transceivers are all built with lowest power and lowest cost in mind in both SFP and QSFP form factor. Tested to an industry's lowest Bit Error Rate (BER) of 10e-15 means fewer transmission errors and retries compared to competing products.

## Conclusion

The move to next-gen IP based broadcast world is inevitable – thanks to the inefficient proprietary SDI fabric. This is even more evident with companies like BBC and FOX embracing their infrastructure using IP network. Mellanox end-to-end efficient broadcast network is radically changing the economics of broadcast applications including live sports coverage, production studios, data transport, content distribution and storage applications for broadcast professionals. With a proven and scalable solution, Mellanox Spectrum switches, ConnectX series adapters and LinkX cables allows broadcasters to save time and money, and deliver extremely reliable HDR video to their viewers.

Mellanox is an active participant of IEEE, SMPTEE, AMWA and JT-NM standards to help further advance IP network requirements needed for the Media and Entertainment Industry.

## Glossary

1. AMWA – Advanced Media Workflows Association - http://www.amwa.tv
2. JT-NM – Joint Task Force on Networked Media - http://www.jt-nm.org
3. SMPTE – Society for Motion Picture and Television Engineers - https://www.smpte.org
4. NMOS
5. VSF – Video services Forum - http://www.videoservicesforum.org
6. AES – Advanced Engineering Society - http://www.aes.org
7. VCSEL – Vertical-Cavity Surface-Emitting Laser

Mellanox TECHNOLOGIES

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