



NVIDIA RIVERMAX

Unique IP-Based Video Streaming Library for Media and Entertainment

NVIDIA® Rivermax® is a unique IP-based solution that boosts video and data streaming performance, including high-dynamic range (HDR) and higher frame rates, for media and data streaming in 10, 25, 40, 50 and up to 100Gb/s networks. Powered by NVIDIA GPU-accelerated computing technologies and running on top of a commercial off the shelf (COTS)-based NVIDIA network interface card (NIC), Rivermax unlocks innovation for a wide range of applications in media and entertainment (M&E), broadcast, healthcare, smart cities, and more.

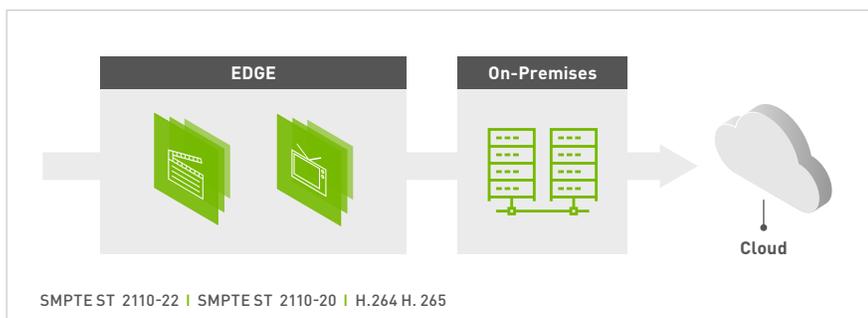
The Future of Media and Entertainment

With the increasing global demand for video content, the M&E market projects tremendous growth in 4K and 8K ultra-high-definition (UHD) video. Media production providers gearing up for the anticipated demand are moving toward next-generation IP-based infrastructures, which offer better scalability, flexibility, and performance over their proprietary software-defined infrastructure (SDI)-based video production solutions.

The Advanced Media Workflow Association (AMWA) Networked Media Open Specifications (NMOS) is used to define media node orchestration (discovery and registration, device connection management, and network control) of media flows.

Best-in-Class Throughput and Latency

Rivermax delivers best-in-class throughput and latency with minimal CPU utilization for streaming workloads, so applications can easily maximize the network usage without running out of system resources. Rivermax leverages NVIDIA Mellanox ConnectX® hardware-streaming acceleration technology that enables direct data transfers to and from the GPU, delivering best-in-class throughput and latency with minimal CPU utilization for streaming workloads. In addition, Rivermax implements a dedicated solution for IP-based video streaming of HD to UHD flows and is the only fully virtualized streaming solution to comply with the stringent timing and traffic flow requirements of the SMPTE ST 2110-21 specification, enabling future cloud-based broadcast solutions.



KEY APPLICATIONS

- > Broadcast and streaming—multi-viewers, playouts, mixers (compressed and uncompressed streams)
- > Video wall displays
- > Virtual production
- > Post-production
- > Monitoring and analysis equipment
- > AI and machine learning video analytics

KEY FEATURES AND BENEFITS

Flexible

- > Scalable, software-defined infrastructure adapts to evolving market needs

AI-Enablement

- > Allows fast integration with AI libraries
- > Reduces human workload/latency
- > Enhances software capabilities

Reduced CapEx

- > Reduces costs—allows moving to a single infrastructure with COTS hardware and dynamic, software-driven features

High Performance

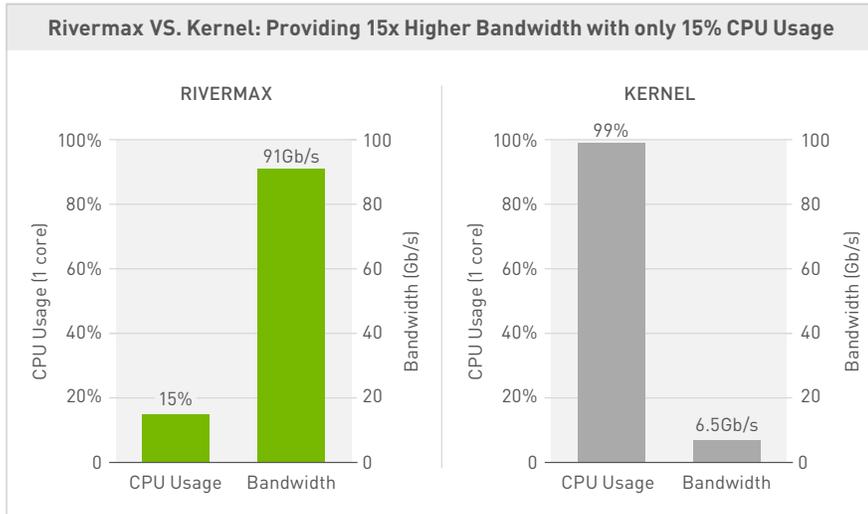
- > Provides industry-leading performance for complex SMPTE 2110 video workflows
- > Single-CPU core receives or sends 8 UHD 2160p 60 or 32 Full HD 1080p 60 streams
- > Scales performance linearly

ConnectX Off-the-Shelf Network Adapters

Rivermax runs on top of a ConnectX-5 or higher COTS adapter that's qualified and certified by industry-leading OEMs and enables applications to seamlessly run in both bare-metal and virtualized deployments.

Leading Performance

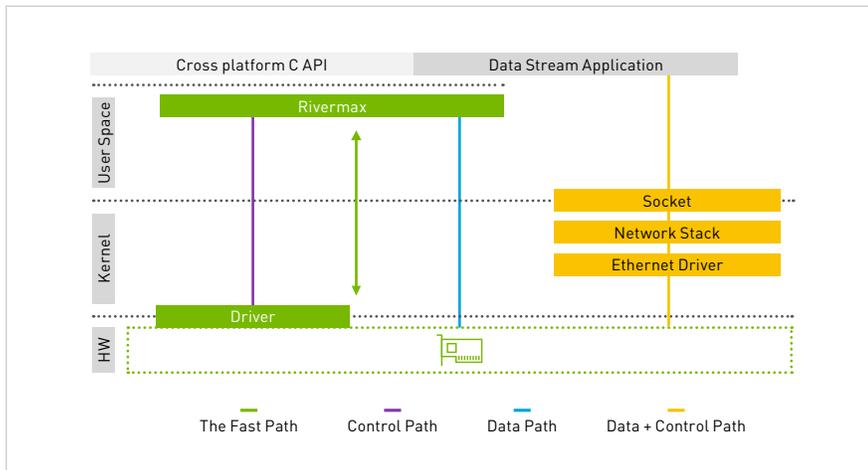
Rivermax delivers from full HD to 8K UHD (7680x4320 50fps 4:4:4 12bpp), equivalent to 82.2Gb/s video bandwidth, using a single CPU core with low CPU usage.



Video Analytics Enablement

Rivermax improves GPU density by enabling UDP/IP NVIDIA Mellanox GPUDirect®, a highly optimized engine that can move data in and out from the GPU without loading the CPU cores.

For video analytics, AI libraries can be integrated with Rivermax to enhance the generation of metadata from live video streams.



Ease of Use

- > Data abstraction—API can be used to represent data in terms of frames/lines rather than individual network packets
- > Simple, straightforward API meets streaming application needs
- > Easy integration with PTP stack
- > Supports both Linux and Windows
- > Short integration time
- > Cloud ready
- > Easily configurable through NMOS management software

Specification Compliance

- > SMPTE ST 2110-21-compliant hardware packet pacing at line rate (up to 100Gb/s)
- > SMPTE ST 2022-6 Transport, ST 2110-30 Audio, and ST 2110-40 Ancillary for live production
- > SMPTE ST 2110-22 with JPEG-XS*
- > Hardware-based seamless reconstruction of SMPTE ST 2022-7 redundant streams

* Requires third-party software

SYSTEM REQUIREMENTS

Library	Rivermax
CPU Architectures	x86_64 and aarch64
RAM	1.5GB to 4.5GB, depending on the amount of video bandwidth sent by the application
Networking	ConnectX-5 or higher adapter
Software*	10/25/50/100Gb/s Ethernet Windows 10 Pro Windows Server 2016 Linux Red Hat CentOS Ubuntu VMware ESXi

* See release notes for full matrix support

Feature Summary

Cloud Ready

- > Supports IP streaming over VMWare ESXi or OpenStack virtual machines (VMs)
- > Each VM owns an adapter instance
- > Enables streaming on both Linux and Windows guest OSs
- > Maintains SMPTE 2110-20/21/30/40 spec compliance on VMs

Software Library Key Characteristics

- > A single standardized cross-platform SDK for Linux and Windows, x86, and Arm CPU architectures
- > Slim set of APIs to transmit and receive flows
- > Straightforward API that meets the needs of any data streaming application
- > Reference code for fast integration
- > Data abstraction—API can be used to represent data in terms of frames/lines rather than individual network packets
- > Cloud ready—supports containers and virtualization
- > Easy integration with any PTP stack
- > Reference examples include NMOS integration for management orchestration
- > Compliance with SMPTE ST 2022-6, 2022-7, 2110-21, 2110-22, 2110-30, 2110-31, and 2110-40 (video, audio, and ancillary)
- > Seamless reconstruction in hardware of SMPTE ST 2022-7 redundant streams

Packet Pacing in Hardware

- > Complies with SMPTE ST 2110-21N (narrow sender), with minimum dependency on the CPU strength, interrupt level, or application
- > Provides packet pacing for any resolution bit rate in a standard network card

Offloading Packet Handling to Network Adapter

- > Kernel bypass technology—hardware-based direct copy between an application's user space memory and the network interface
- > Hardware-based packet handling offloads
- > GPUDirect—zero-memory copy to GPU, reducing PCIe transactions, CPU usage, jitter, and latency
- > Performance scales linearly both in terms of CPU loading and network bandwidth
- > All network adapter hardware-packet pacing is in accordance with the SMPTE ST 2110-21 specification
- > Joint Task Force on Networked Media (JT-NM) tested**

** For more details on the JT-NM Tested program, please visit: jt-nm.org/jt-nm_tested

Learn more

Learn more at developer.nvidia.com/networking/rivermax