
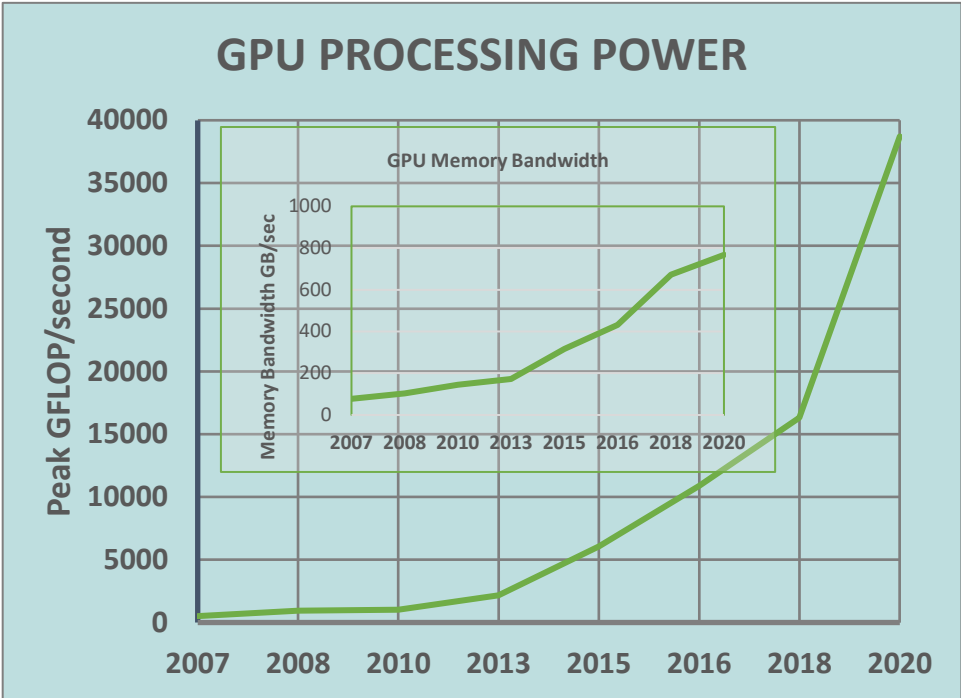


Tighter NIC/GPU Integration Yields Next Level Media Processing Performance

Thomas Kernen, NVIDIA
Thomas True, NVIDIA

INTRODUCTION

Media Production Workflows Transitioning from SDI to IP Meet COTS Hardware



ST 2110 & COTS network adapters

Modern offloading capabilities

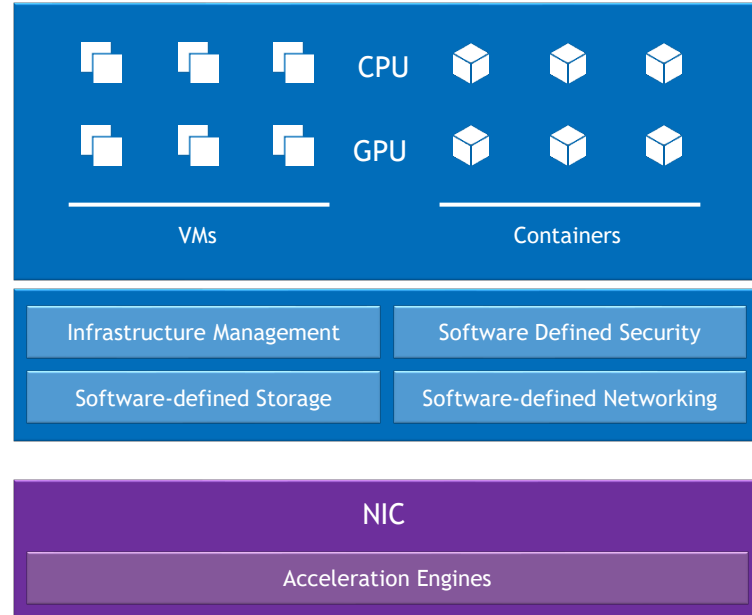
Maximizing I/O performance to meet SMPTE ST 2110-21's strict timing accuracy requirements:

- **Direct Memory Access (DMA):** Copy data to/from NIC to RAM
- **Packet checksum computation offload:** Send and/or receive messages to/from NIC
- **Segmentation Offload:** Move multipacket buffer from host to the NIC.
- **Receiver Side Scaling (RSS):** Enable distributing packets to separate queues to CPU cores
- **Large Send Offload (LSO):** TCP stack to build longer TCP messages & to NIC, re-segments the message into multiple TCP packets.
- **Large Receiver Offload (LRO)/Receive Segment Coalescing (RSC):** for received packets.
- **UDP Segmentation Offload (USO):** Similar capabilities to LSO for UDP packets.
- **Header-Data split:** Split received Ethernet frames headers and data into separate buffers.
- **Kernel bypass:** Reduce OS kernel overhead by enabling applications to directly access the NIC resources.

What is a Datacenter Processing Unit?

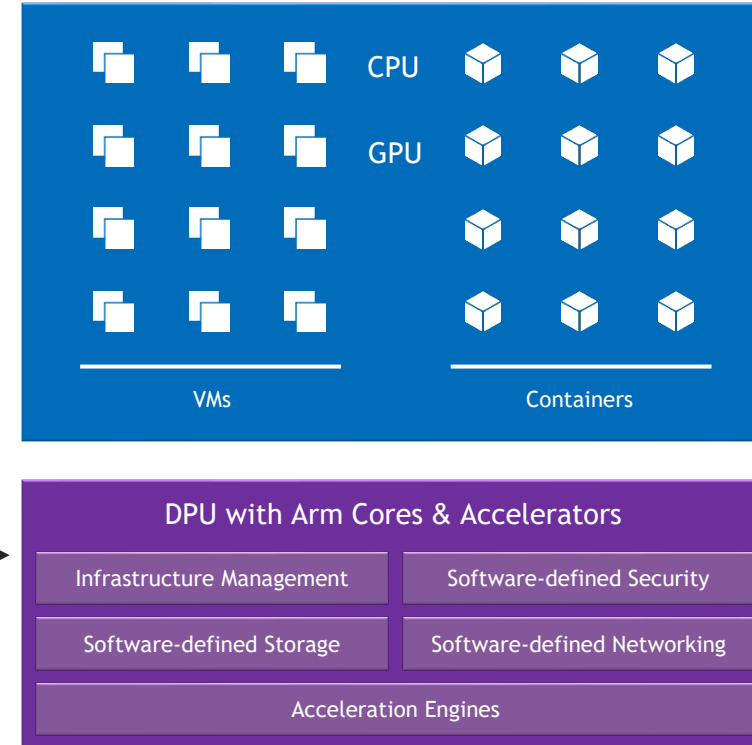
3rd pillar in Datacenter architecture

TRADITIONAL SERVER



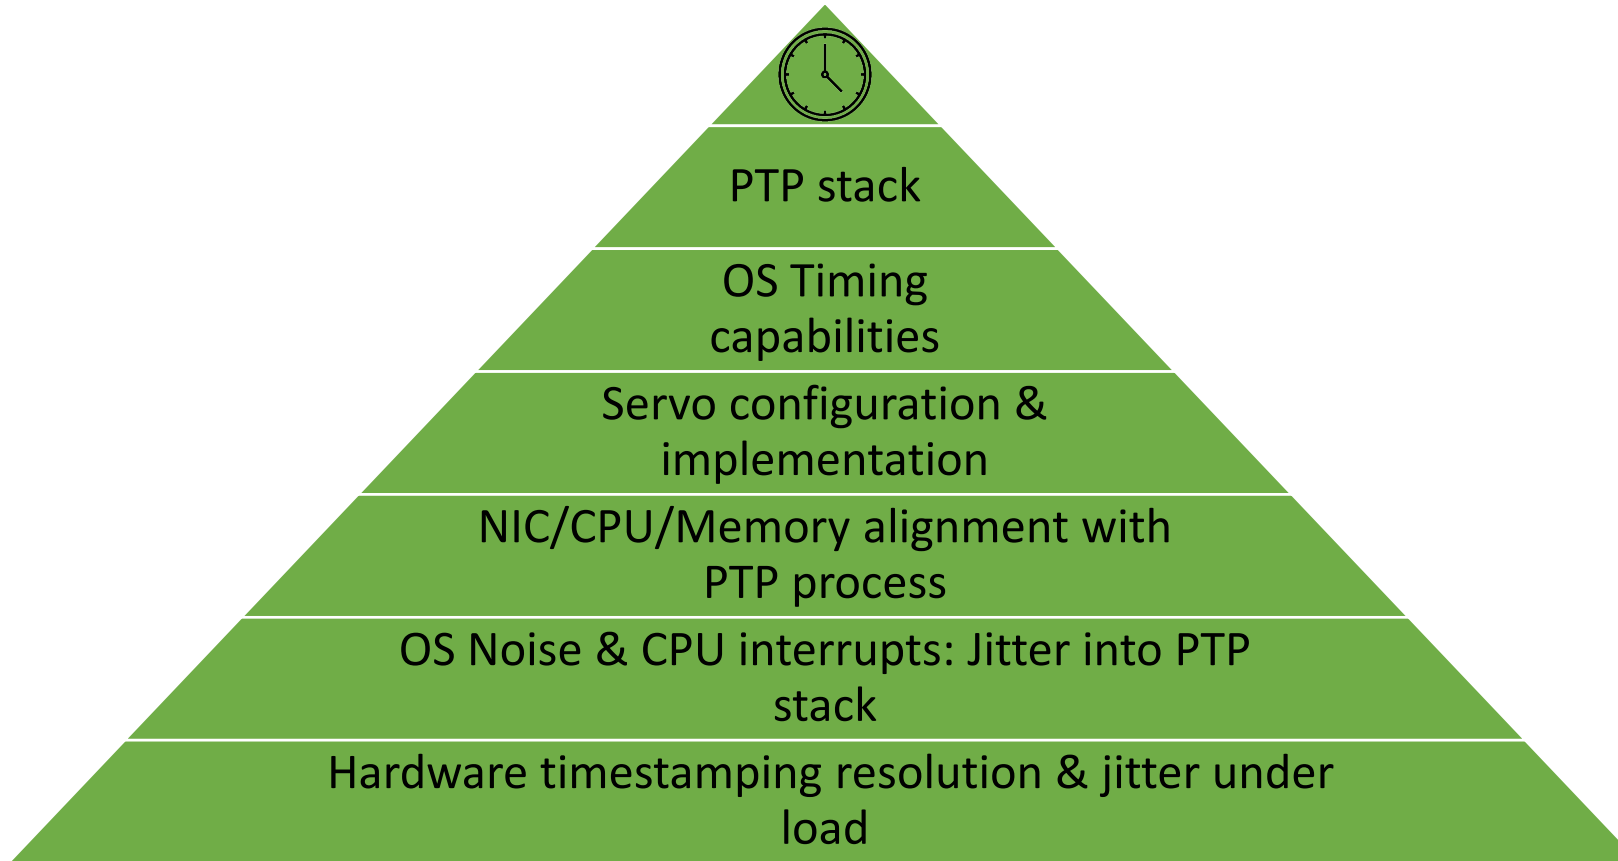
Manual Infrastructure Management | Security Appliances |
Storage Systems | Static Networks | Microservices |
East-West Traffic | Storage Access | Zero Trust Security

DPU-ACCELERATED SERVER

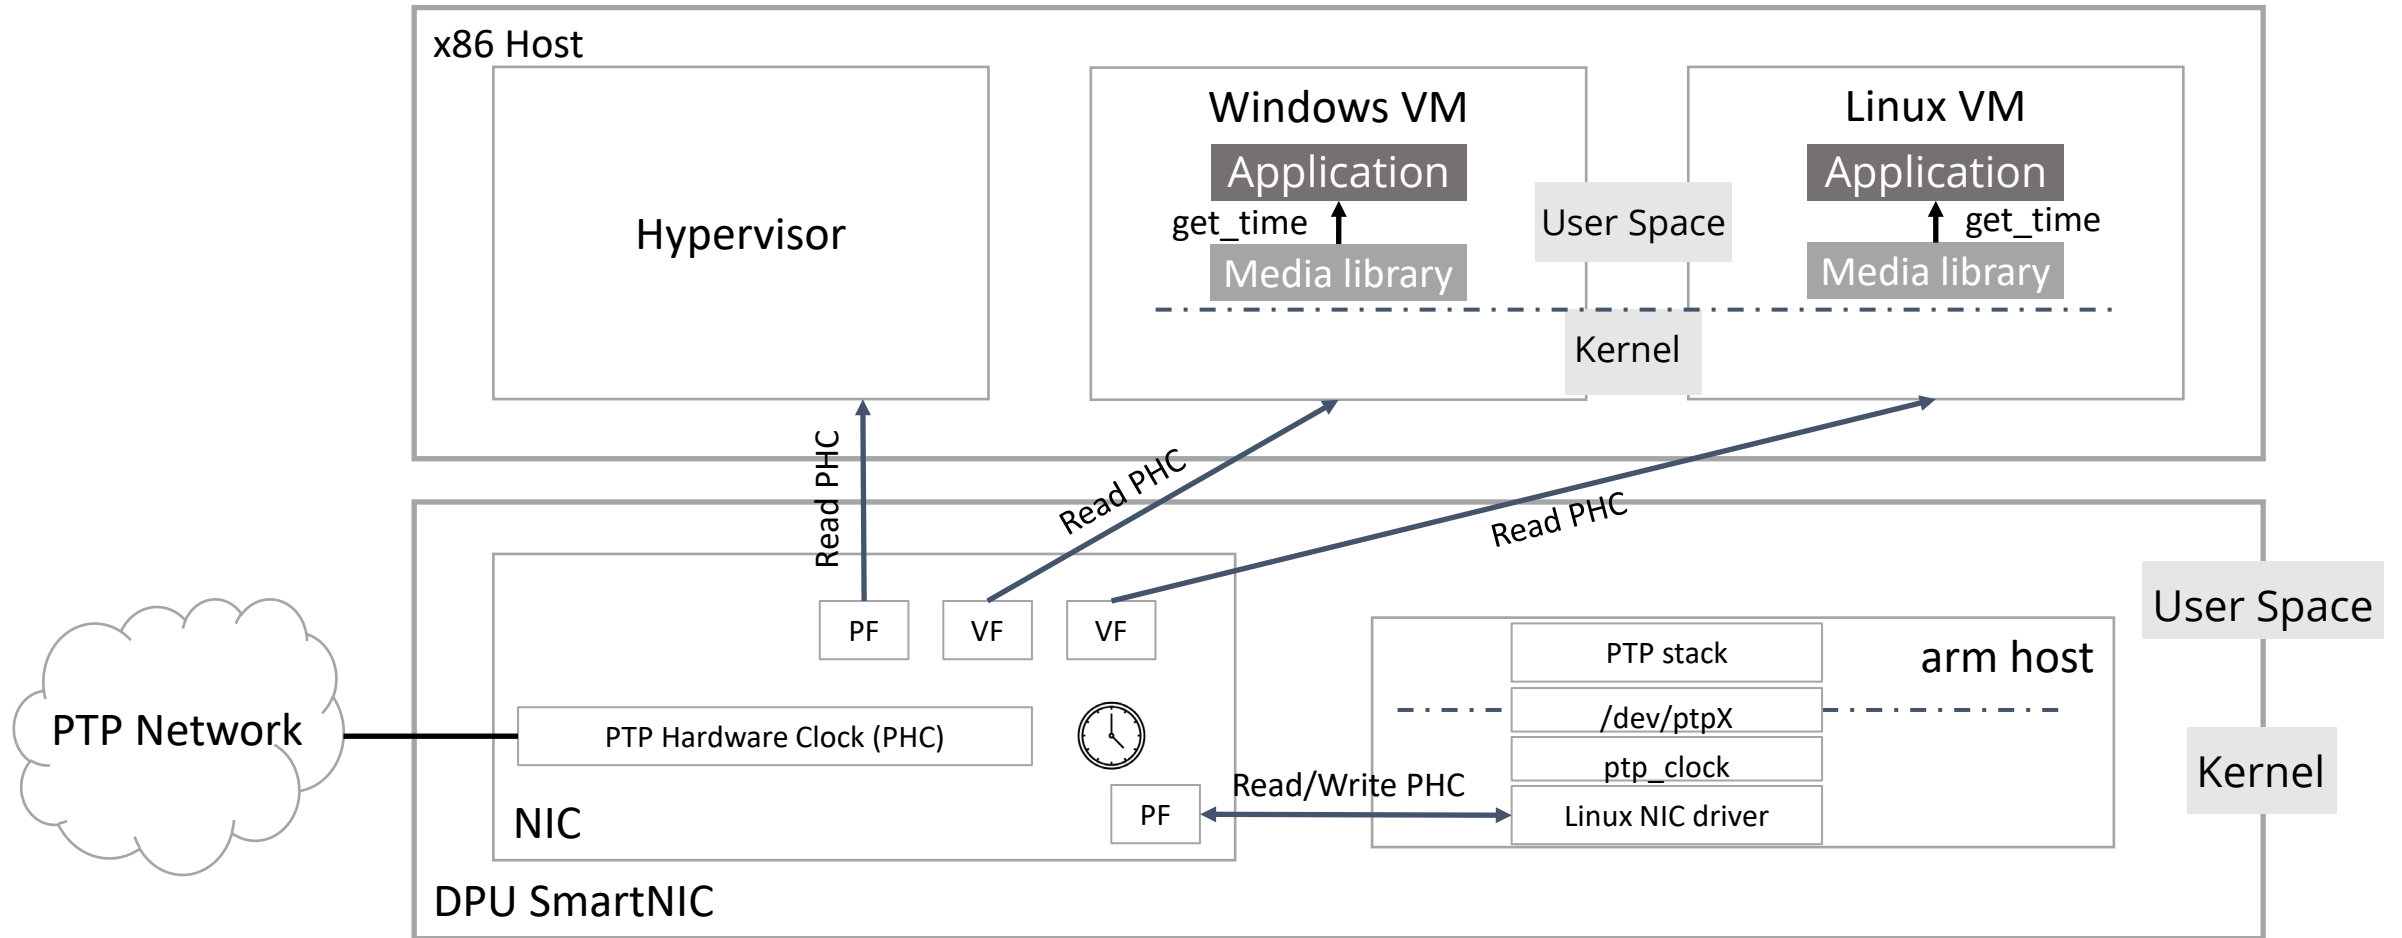


Delivering Consistent Time Across Hosts

Challenges to be overcome



DPU Timing Architecture

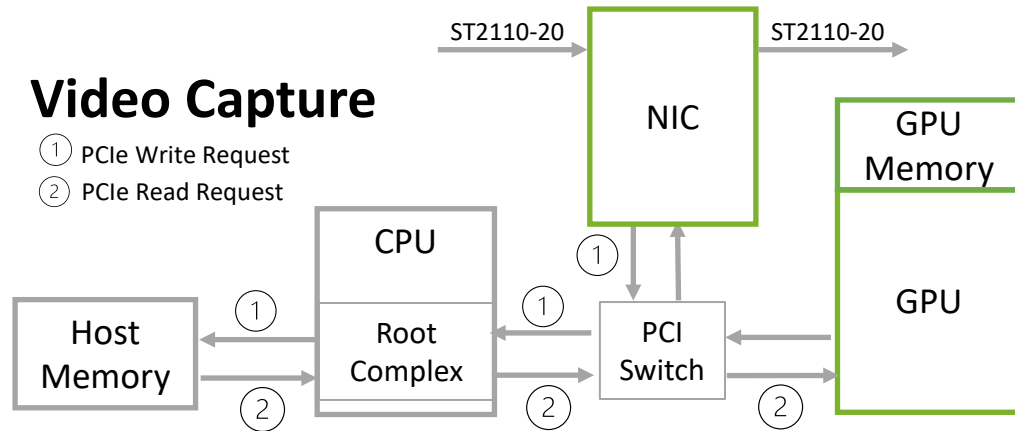


CLASSIC DATA TRANSFER BETWEEN GPU AND NIC

No Direct Data Transfer

Video Capture

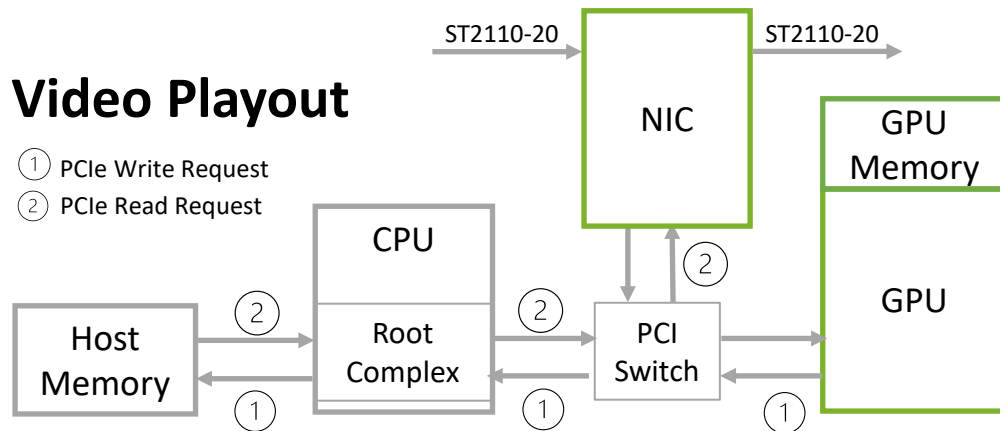
- ① PCIe Write Request
- ② PCIe Read Request



- Host memory buffers required
- Two-step memory transfer process between NIC and GPU
- Peer-to-host-to-peer DMA transfers

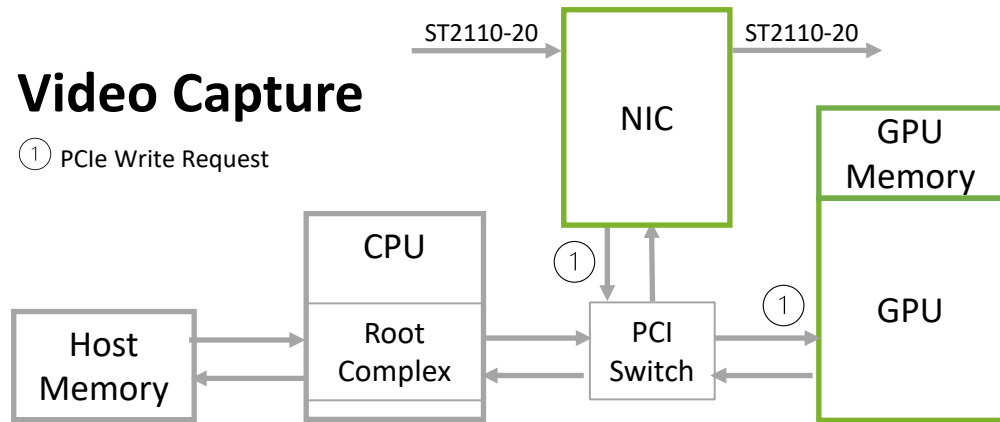
Video Playback

- ① PCIe Write Request
- ② PCIe Read Request

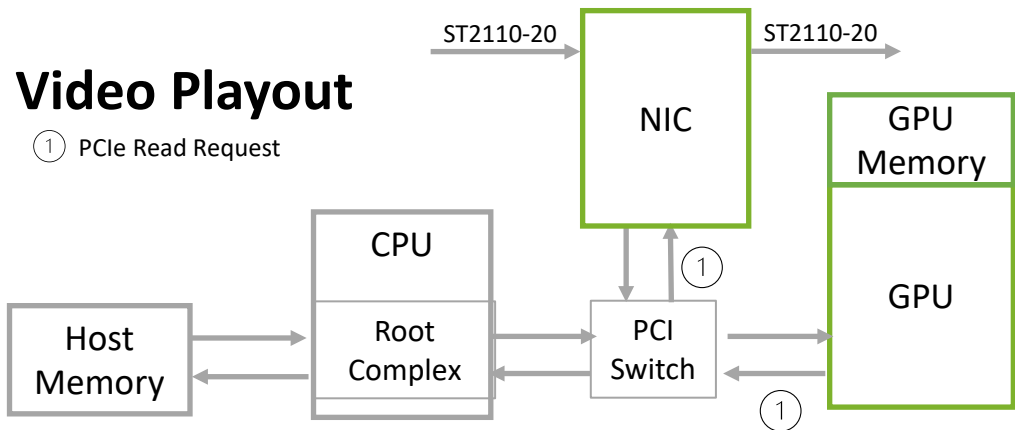


DIRECT DATA TRANSFER BETWEEN GPU AND NIC

Direct Data Transfer

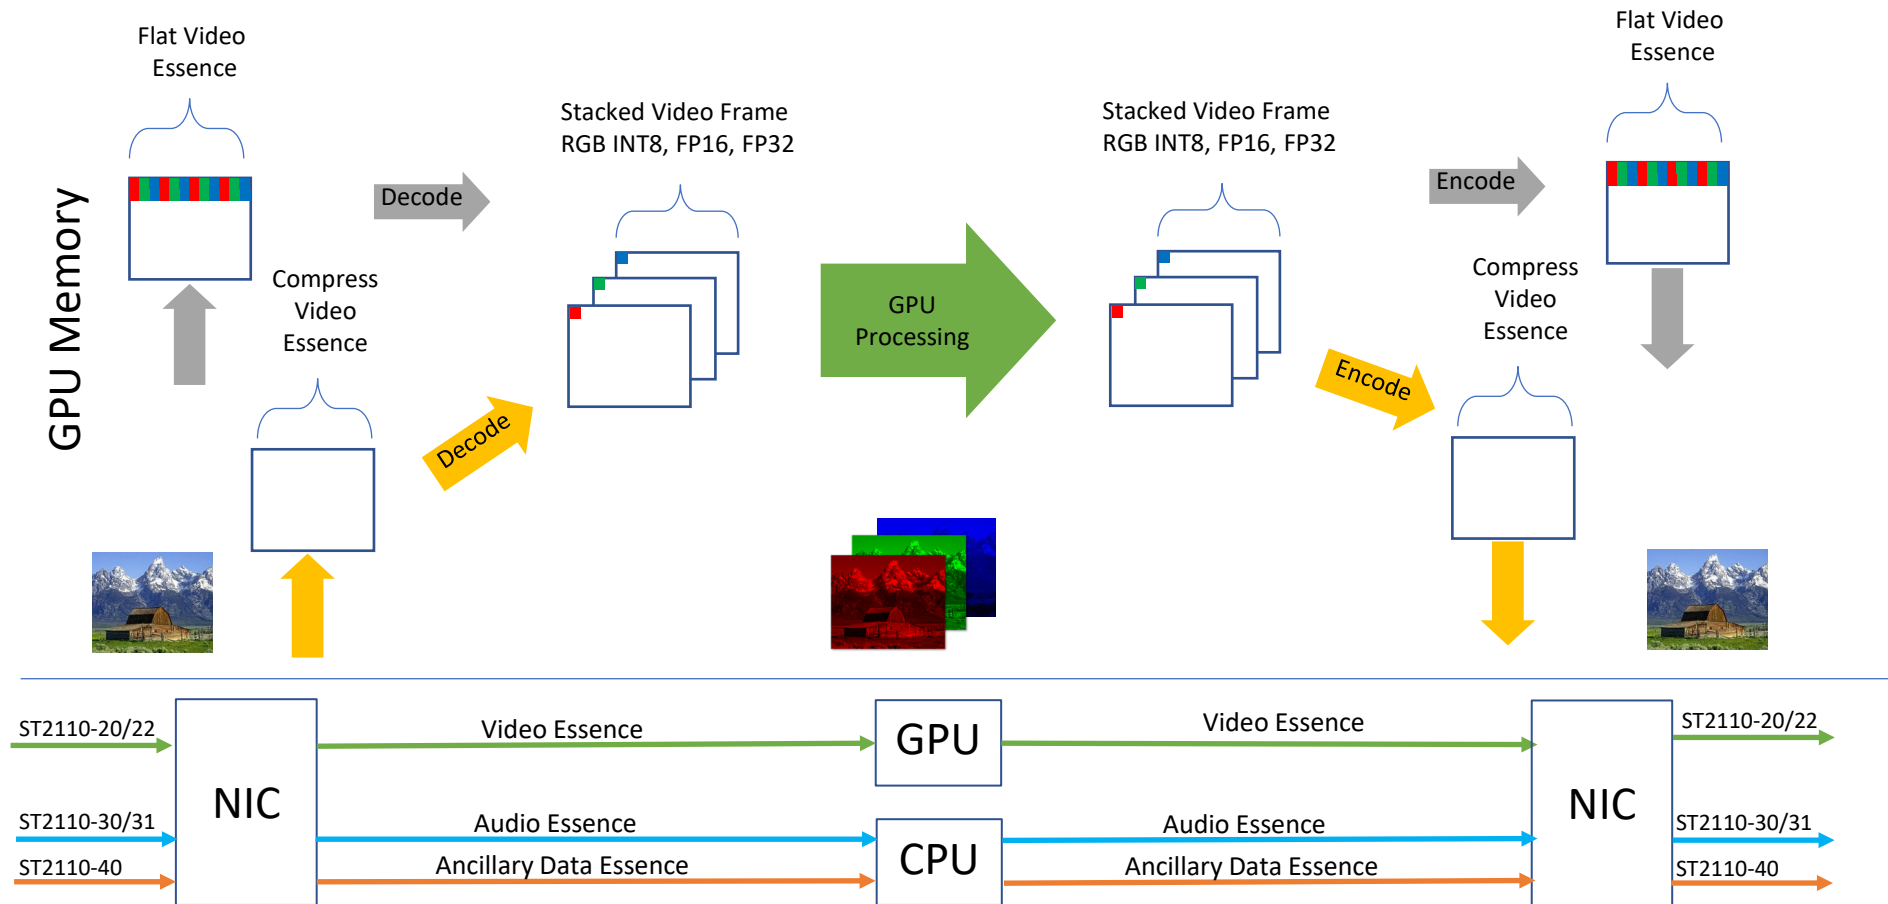


- Utilizes only GPU device buffers
- Single PCIe read or write request
- Peer-to-peer DMA transactions between GPU and NIC
- Optimum performance achieved when NIC and GPU reside on same PCIe switch



GPU PROCESSING

Flexible Data Conversion Between NIC and GPU

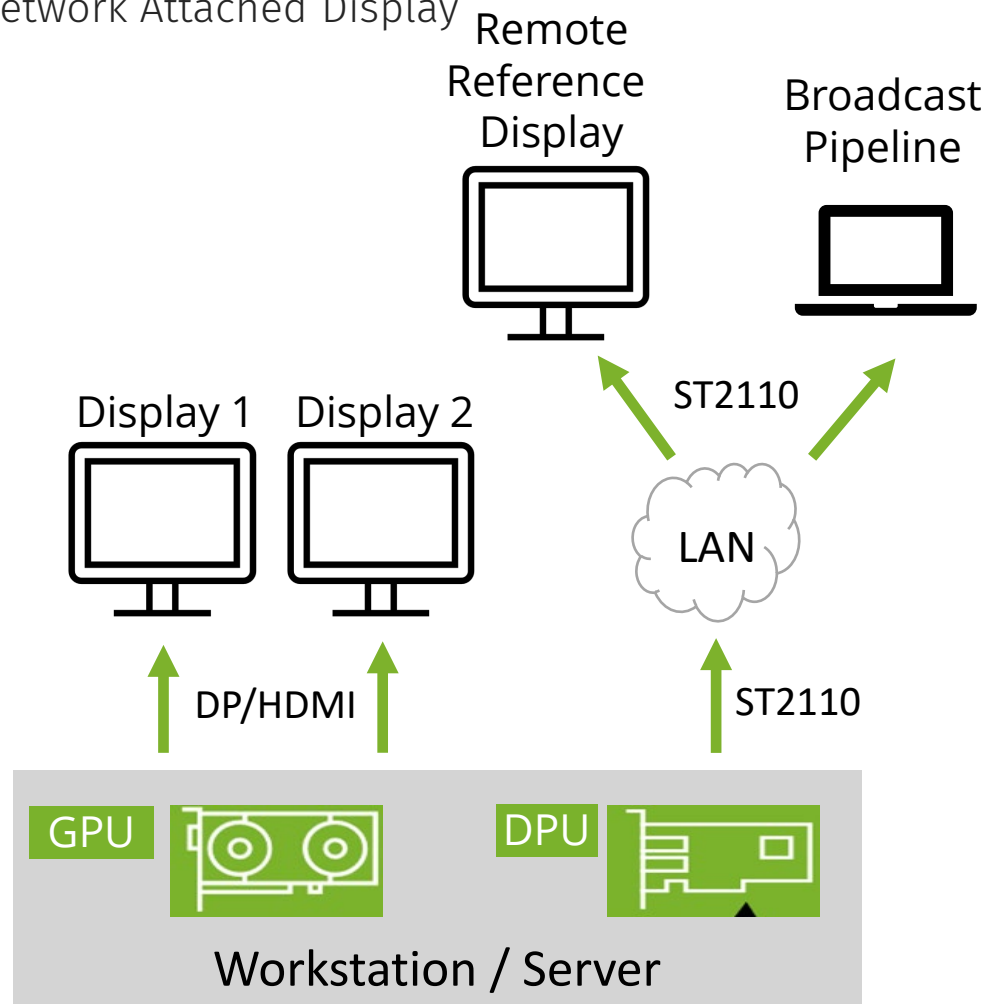


https://commons.m.wikimedia.org/wiki/File:Barns_grand_tetons.jpg

BRINGING IT ALL TOGETHER

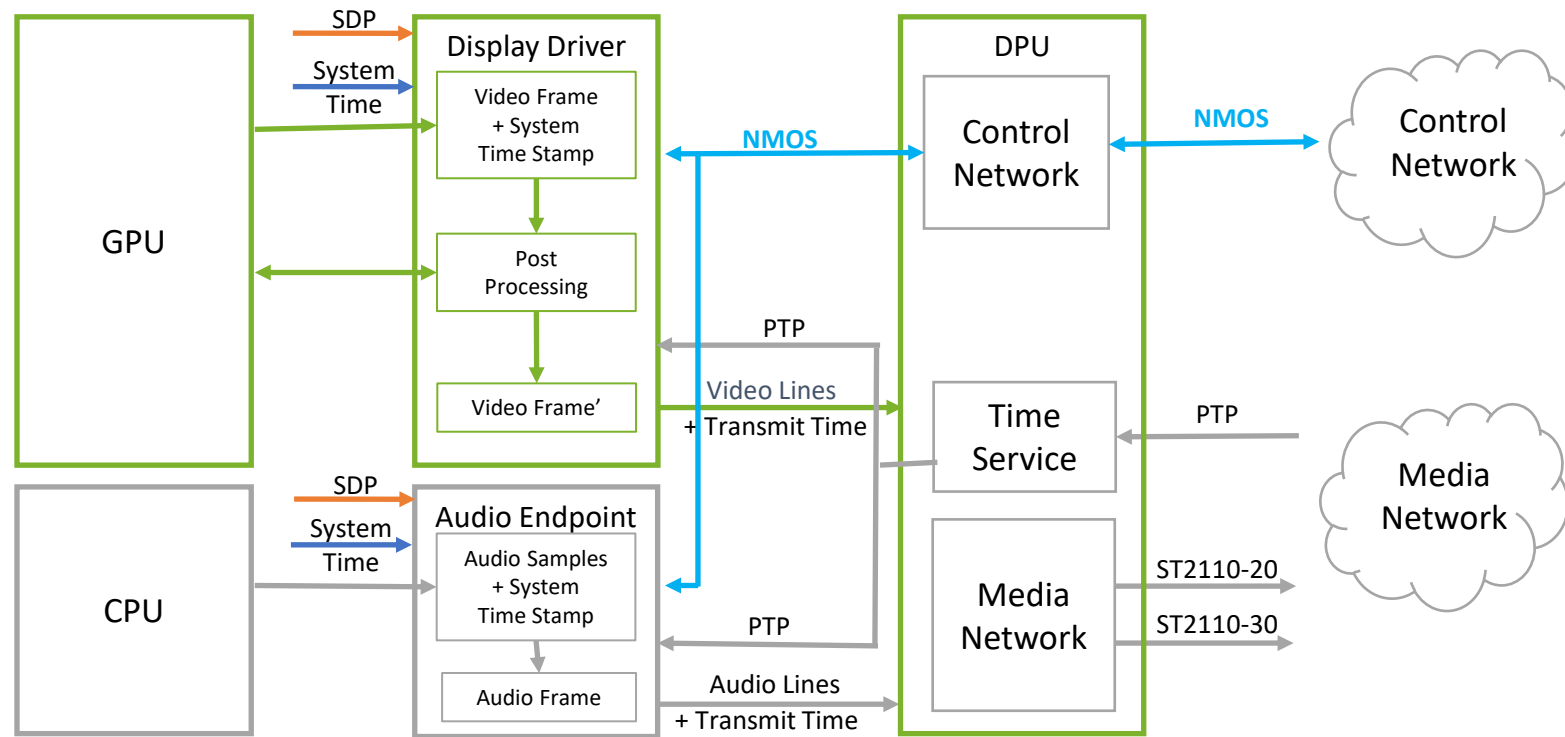
GPU + DPU = Network Attached Display

- Combines GPU and NIC technologies
- Virtual display on the Windows desktop
- GPU rendered frames transmitted by NIC as SMPTE 2110-20 stream
- OS audio stream transmitted as SMPTE 2110-30 stream
- Transparent output to high-quality reference display or broadcast pipeline
- Adds ST2110 output to any desktop application
- Works on bare metal system or virtualized in the data center



BRINGING IT ALL TOGETHER

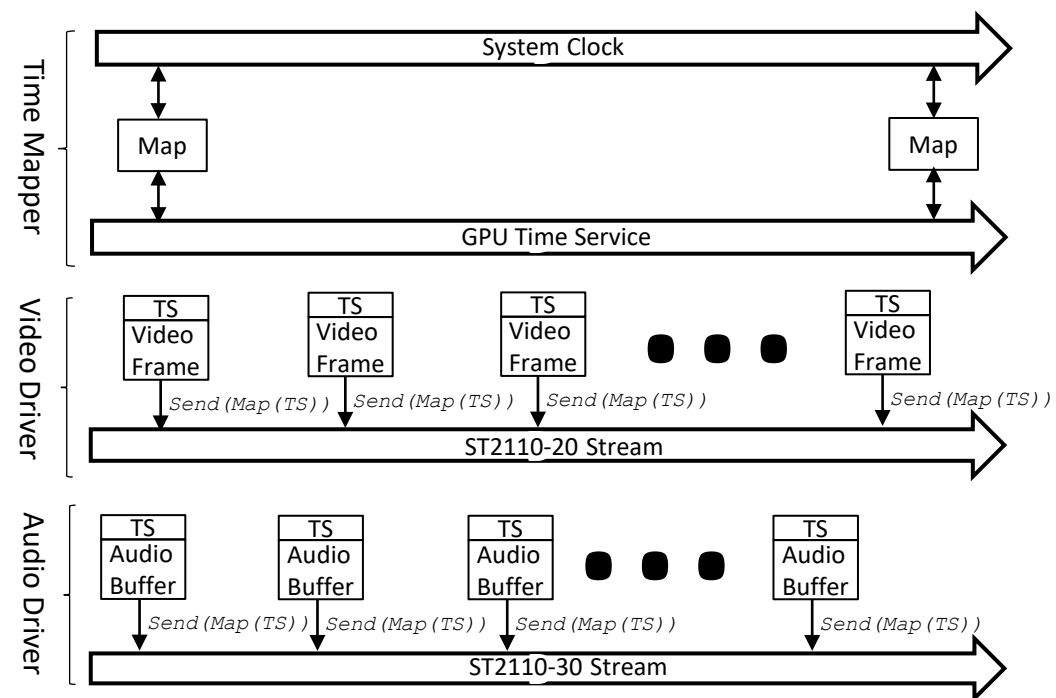
GPU + DPU = Network Attached Display



AUDIO-VIDEO SYNCHRONIZATION

GPU + DPU = Network Attached Display

- Query PTP time from DPU
- Map Windows OS timestamps to PTP time
- DPU uses mapped PTP times to schedule transmission



CONCLUSION

Tight Coupling Between GPU and NIC Reduces System Overhead

- GPU/NIC performance increases unlock COTS based workflows
- Tighter integration required to maximize performance increase
- Network protocol termination offload from CPU to DPU
- OS/system agnostic accurate timing delivered via DPU
- Bypass CPU for efficient data transfer between DPU/GPU
- Network attached display as a technology demonstrator

Thank You!