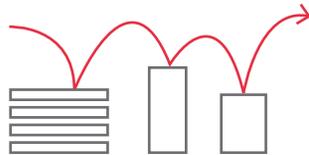




# Quadra Video Server

Ampere® Edition

# The problems with video transcoding started way back:



## Transcoding functions

historically, have been siloed in standalone servers with ABR packaging and other production tasks performed elsewhere.



Over time, this schema become inefficient and inflated both CAPEX and OPEX.

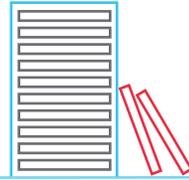


Ultimately, the pandemic isolated everyone forcing the need to stream:

- Video content
- Social media
- Video conferencing
- Remote working
- Education
- E-commerce shopping
- Gaming
- Health and wellness



In a rush to meet demand, companies pushed their platforms to the lowest hanging clouds.



Today, most data centers are full of bulky and power hungry hardware.

The costs are unsustainable.



## Now what?

The streaming industry has reached a maturity that no longer tolerates gross power consumption and profit erosion previous accepted with CPU-based video processing. This inflection point is forcing video streaming providers to do more with less in order to remain competitive in this explosive-growth industry.

# Re-Introducing ASIC technology

ASIC technology has been at the forefront of innovation since traditional TV broadcasting. Here's how the renaissance of ASICs have revolutionized today's video streaming workflow:

- Efficient encoding and decoding
- Ultra-low power consumption
- Reduced latency in high-quality real-time
- CDN optimization in acceleration and scalability
- Robust security and user privacy

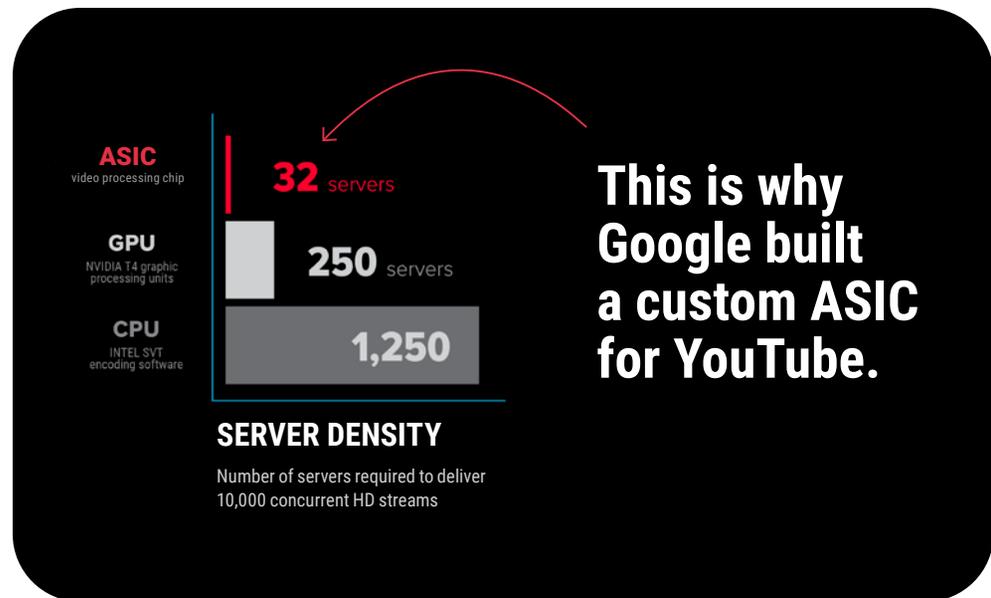
Industry migration is underway to capitalize on ASICs.

“

“There are two types of companies in the video business.

Those that are using video processing ASICs ...and those that will.”

*David Ronca  
Video Encoding Expert, Meta  
(Formerly from Netflix)*



This is why Google built a custom ASIC for YouTube.

## Industry titans already pivoted into ASICs

Both Google and Meta built proprietary transcoding chips to reduce costs, reduce footprint and increase density for their streaming platforms only.

For everyone else, we built one for you.

# We built one better: Smart VPUs with AI

We call ours Video Processing Units because they only process video. They're smart because the added AI will supercharge your productivity 10X. Lastly, our VPUs are based on ASIC technology and available worldwide.



**NETINT Quadra Smart VPU**  
ASIC chip, 2nd Gen G5

## Built for Speed

Tasking VPUs with only video transcoding, they process, scale, and overlay streams faster than any other transcoder.

## Built for Volume

VPUs shift substantial computational loads to the cloud by specializing in a singular processing task so it pushes 10X more streams to the cloud.

## Eats Less Power

VPUs have a tiny appetite for power, drawing <7W, and uses 80% less energy than CPU-processing.

## Boosted with AI

AI inferencing accelerates accuracy and productivity with key object filtering and facial recognition.

**NETINT**  
**2024 Winner**

**TECH**  
**EMMYS®**

Design & Deployment  
of Efficient Hardware Video  
Accelerators for Cloud

We feel it's worth mentioning,  
co-winners for this same category  
were Google, Meta and AMD.

<https://theemmys.tv/tech-75th-award-recipients/>



# Finally, a server with monster throughput has a tiny appetite for power

## The Quadra Video Server Ampere Edition

Industry leaders united to consolidate your video transcoding process, accelerate core functionality and integrate multiple production processes into this high-performance server.

- HEVC, H.264 & AV1 encoding
- HEVC, H.264, & VP9 decoding
- Up to 8K resolution, 10-bit HDR
- AI enhanced productivity

**NETINT**  
**Smart VPU**

Quadra  
T1Us

**AMPERE**  
**CPU**

2.8 GHz 96-core  
Altra Max  
M96-28

**SUPERMICRO**  
**Server**

MegaDC  
ARS-110M-NR  
1RU



# Built for pushing streams to the cloud

When Quadra Video Processing Units (VPUs) handle complex video encoding and decoding tasks, it leads to significant advantages. Firstly, this approach offloads the workload from the Ampere 96-core CPU, enabling it to focus on other operations such as deinterlacing, and AV1/MPEG-2 decoding. Secondly, leveraging Quadra VPUs for these tasks results in a synergy that boosts throughput by tenfold while reducing costs -80%. This dual benefit represents a highly efficient and cost-effective solution.

## Enterprise application integration

The Quadra Video Server Ampere Edition empowers video engineers to run additional publishing-related applications such as:

- Real-time captioning and translating with OpenAI/Whisper
- Dynamic ABR packaging
- Streaming orchestration and
- Content management

320

simultaneous live 1080p30 encoding sessions at ultra-low latency in a compact 1RU server



# Ampere® Altra® Max Processor



## Predictable High Performance

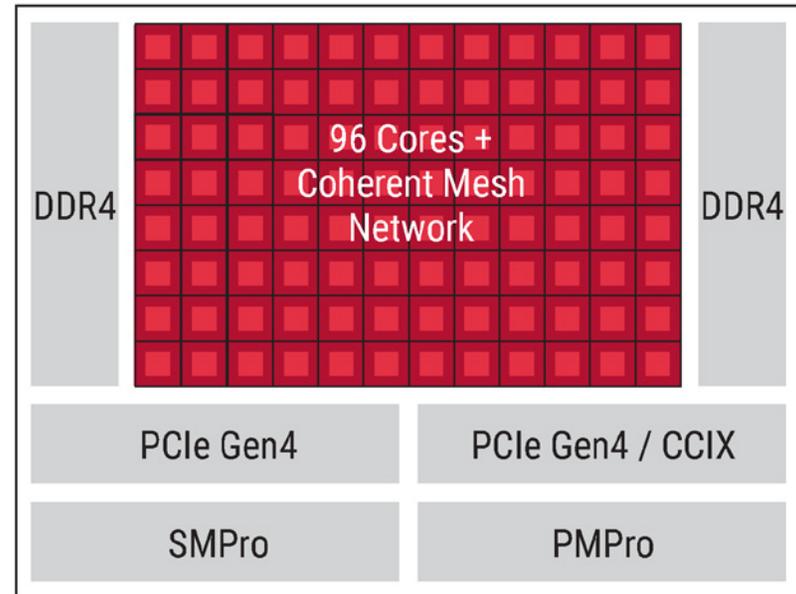
- 96 cores | 128 cores
- Coherent mesh-based interconnect
- High-memory bandwidth and density

## High Scalability

- Industry leading power/core
- Cache-coherent multi-socket support
- Flexible I/O connectivity

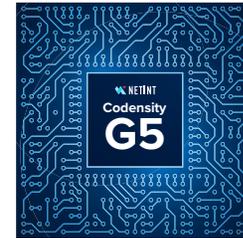
## Power Efficiency

- Advanced system, security and power management
- Monolithic die on leading 7nm process
- Leading power/core



# NETINT Codensity G5

Smart VPU with AI



## High Performance Encoding

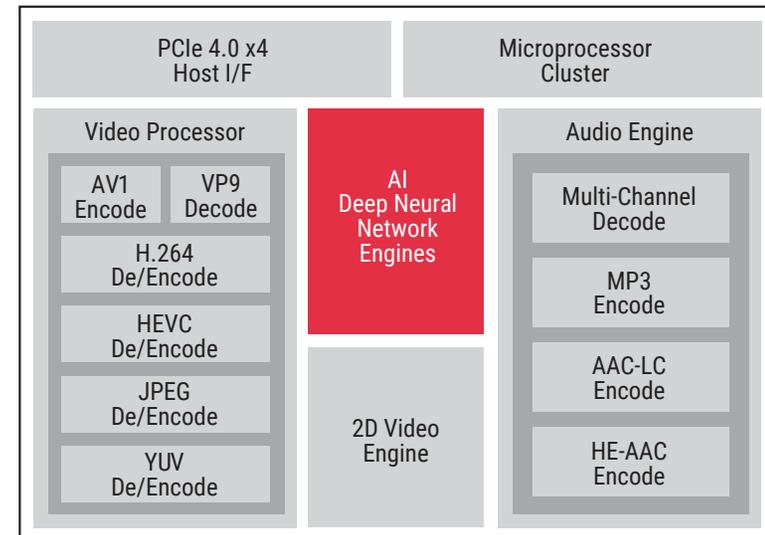
- 8K UHD Video Encoding
- On-chip AV1, H.264 and HEVC
- Up to 60% reduction in bitrate

## Flexible Architecture

- Programmable micro-processor enables firmware optimization
- Pipeline processing improves performance and increases video quality

## AI Processing Engines

- Two Deep Neural Network DNN engines onboard at 18 TOPS
- Enables object detection, classification and segmentation for image quality improvement and content-adaptive rate control
- Advanced performance and seamless integration for ROI (region-of-interest) encoding and background replacement

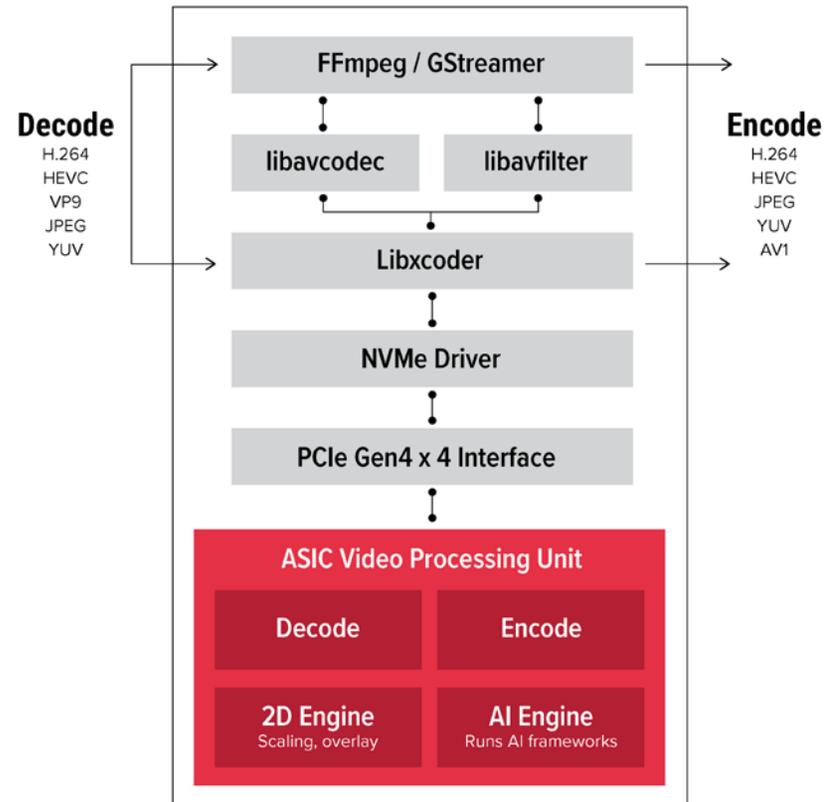


# Simple Integration

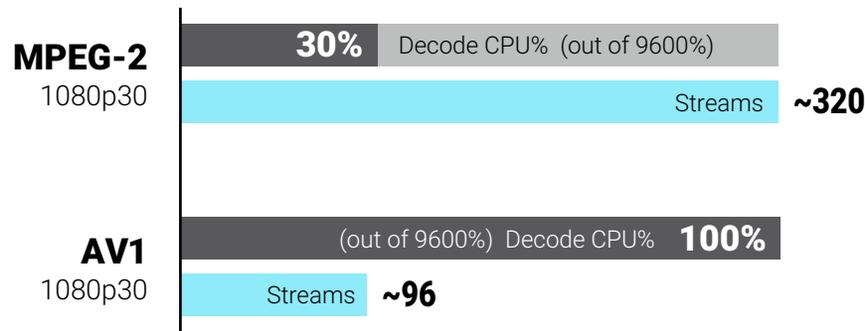
## Open-source suite of processing tools

Many video processing and transcoding applications developers use FFmpeg and GStreamer, two open-source software libraries offering a vast suite of video processing functions. The Quadra video server includes highly efficient FFmpeg and GStreamer compatible SDKs, allowing operators to apply an FFmpeg or GStreamer patch to complete the integration.

The libavcodec patch on the host server functions between the Quadra NVMe interface and the FFmpeg and GStreamer software layers simplifying integration and enabling fast and efficient performance and capacity upgrades.



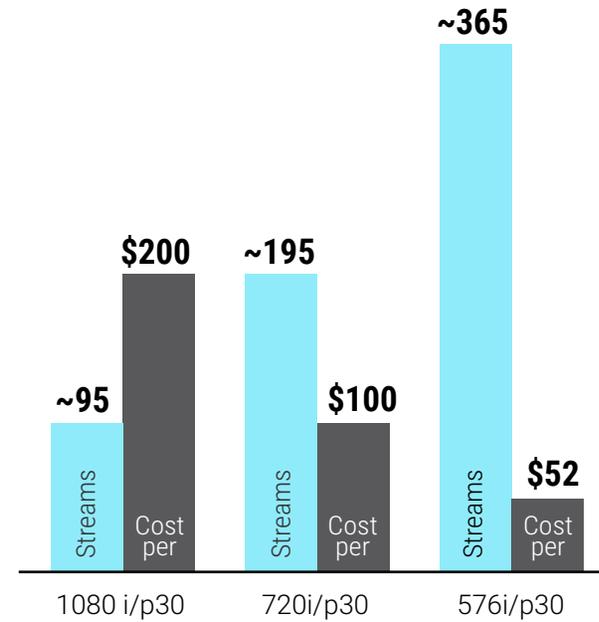
# What 96 cores can do for you



## Decode MPEG-2 and AV1

At 1080p30, the Ampere CPU can decode incoming streams:

- over 320 MPEG-2 streams, utilizing only 30% CPU%
- around 100 AV1 streams, maxing out the CPU%



## Deinterlace with Ease

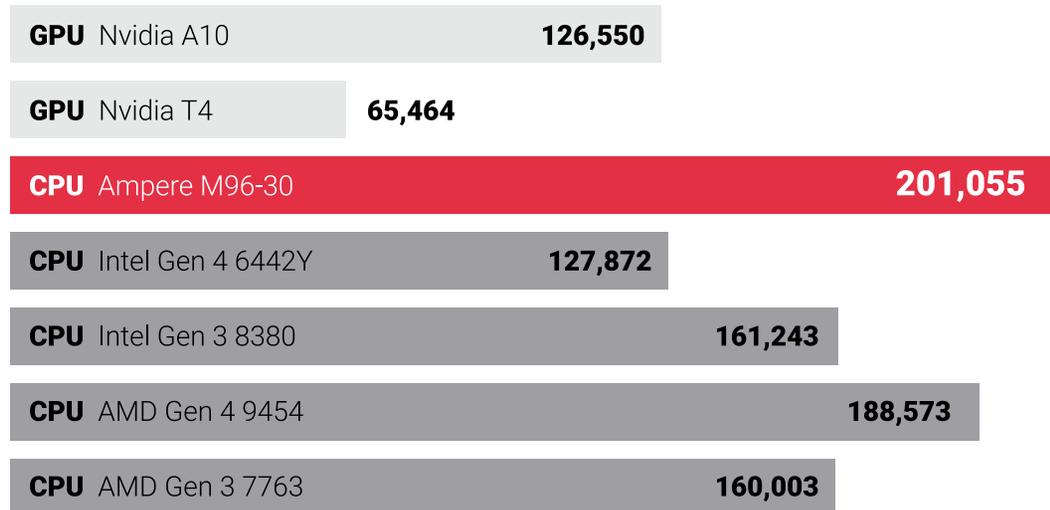
The Ampere CPU is ideal for video engineers working with interlaced source content.

# Add real-time captioning and translating to your live streams with AI

## International accessibility with OpenAI/Whisper

The 96-core Ampere CPU is ideal for video engineers to run additional publishing-related applications like Whisper AI for live captioning and translating.

*(Available Q2 2024)*



*Number of inferences per second*

# Quadra T1U

Smart VPU

Video Processing Unit with AI | Codensity G5



<b>Form Factor</b>	U.2
<b>ASIC</b>	1x Codensity G5
<b>Interface</b>	PCIe 4.0 x4
<b>Power Consumption (Typ)</b>	17W
<b>Usage</b>	24/7 Operation
<b>Operation Temperature</b>	0 - 50°C
<b>RoHS Compliance</b>	European Union (EU) ROHS Compliance Directives
<b>Product Health Monitoring</b>	Self-Monitoring, Analysis, and Reporting Technology (SMART) commands Temperature Monitoring and Logging
<b>Video Encoding Standards/Formats</b>	AVC/H.264 Baseline, Main, High, High 10 HEVC/H.264 Main, Main 10 JPG YUV 420 8 bit/10 bit encoding AV1 Main
<b>Video Decoding Standards/Formats</b>	AVC/H.264 Baseline, Main, High, High 10 HEVC/H.265 Main, Main 10 VP9 Profile 0, 2 JPEG YUV 420 8 bit/10 bit decoding
<b>Throughput Capacity</b>	Up to 32x 1080p30, 8x 4Kp30, 2x 8Kp30
<b>Audio Codecs</b>	MP3, AAC-LC, HE-AAC
<b>Level</b>	1 to 6.2 Main Tier
<b>Resolution</b>	32 x 32 to 8192 x 5120
<b>Scan Type</b>	Progressive
<b>Bitrate</b>	64kbit/s to 700Mbit/s
<b>Software Integration</b>	FFmpeg SDKs, GStreamer, LibXcoder API integration
<b>AI Deep Neural Network Engines</b>	15 TOPS AI Assisted Encoding
<b>Region of Interest (ROI)</b>	ROI enables the quality of some regions to be improved at the expense of other regions
<b>Closed Captioning</b>	EIA CEA-708 for H.264 and HEVC encode/decode
<b>High Dynamic Range (HDR)</b>	HDR10, HDR10+, HLG for H.264 & HEVC encode/decode
<b>Low Latency</b>	Sub-frame latency
<b>IDR Insert</b>	Forced IDR frame inserts at any location
<b>Flexible GOP Structure</b>	8 presets plus customizable GOP structure
<b>Video 2D Processing Engine</b>	Crop & Padding/Scaling/Overlay/YUV & RGB Conversion

# Quadra Video Server

## Ampere Edition

Supermicro | MegaDC ARS-110M-NR



<b>CPU Options</b>	Ampere Altra Max M96-28, 96-cores
<b>Operating System</b>	Ubuntu 22.04.3 LTS
<b>Memory</b>	256GB of DDR4-3200 RDIMM
<b>Storage</b>	400GB M.2 SSD
<b>NVMe Support</b>	10x
<b>PCIe Expansion</b>	Three PCIe 4.0 x16 LP slots, one PCIe 4.0 x16 AIOM slot
<b>Network Options</b>	1 RJ45 Dedicated IPMI LAN port 2x 25Gb SFP28 Ethernet LAN Ports
<b>Power Consumption</b>	~500W
<b>Power Supply</b>	800W Redundant Platinum Level power supplies
<b>Transcoders</b>	10x NETINT Quadra T1Us
<b>Encoding Capacity</b>	Up to 20x 8Kp30, 80 4Kp30 or 320x 1080p30
<b>Codec Support</b>	H.264 - Encode/Decode
	HEVC - Encode/Decode
	JPG - Encode/Decode
	VP9 - Decode
	AV1 - Encode
<b>Software Integration</b>	FFmpeg, GStreamer, NETINT SDK

<b>Physical Dimensions</b>	W: 54.5 mm x H: 40.25 mm x D: 220 mm
<b>Rack Size</b>	1U
<b>Weight</b>	39 lbs (17.69 kg) <i>(fully loaded with 10 T1U VPUs)</i>
<b>Environmental</b>	50 degrees F to 95 degrees F Operating Temperature, 8% to 90% Operating Relative Humidity
<b>Power Inputs</b>	750W: 100-127Vac / 50-60Hz
	800W: 200-240Vac / 50-60Hz
	800W: 230-240Vdc / 50-60Hz
<b>Certifications</b>	RoHS Compliant, UL Approved



For more information on NETINT encoding solutions, contact us.

[sales@netint.com](mailto:sales@netint.com)  
[netint.com](https://netint.com)