NETINT

Logan
Video Server

Transcoding | Codensity ASIC G4

Executive Summary

Early adoption of ASICs slingshots streaming platforms ahead of their competition due to their low operational cost basis.

Designed as a quickstart solution for high density live video encoding and transcoding, the Logan Video Server comprises ten T408 video transcoders in a 1RU chassis that performs the equivalent work of roughly 10 dedicated servers running a typical open-source FFmpeg and x264 or x265 configuration. The server delivers the lowest TCO of any solution in the market, and is a drop-in replacement for existing CPU and GPU encoding stacks.

NETINT application specific integrated circuits (ASICs) are the secret to replacing software-driven video encoding for video platforms and delivery services wanting to decrease carbon emissions up to 40x, operating costs as much as 20x and reduce the number of servers needed by 90% for the equivalent video encoding workloads.

The results are profound and transformational.

ASICs will be the engine powering all future video streaming experiences

NETINT

HYPERSCALE PROFITABLY

Live streaming experiences are seeing rapid adoption

Applications:

- Live events
- Interactive video
- Cloud gaming
- Real-Time video
- Virtual worlds
- 360 / VR / AR

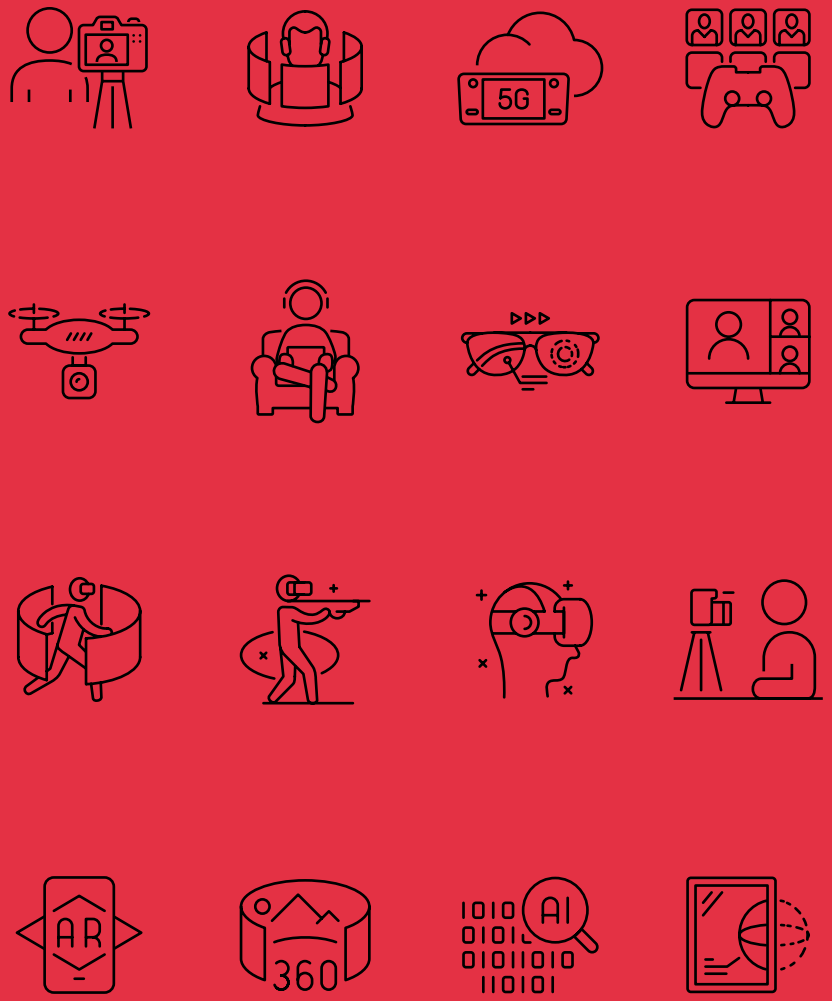

NETINT

The insatiable appetite of video consumers

They want nonstop, never-ending, high-resolution, non-buffering content accessible on any device. Now.

Viewers have developed an addiction to continuous content streaming. Video delivery and entertainment experiences are shifting from file-based to live where low latency and controlling operational costs are paramount.

- Increased public cloud provider costs are stressing businesses
- Live experiences are growing in resolution, color depth, and quality expectations
- Playback is expected on every device using its full capability
- More data centers are needed to handle capacity increases

2021

Social video viewing surpassed Google search traffic

> 1 billion active montly users on short-form video apps.

65%

Percentage of ALL internet traffic is video streaming, increasing 24% year over year.

40%

Percentage of people 18 to 24 turning to visual-based social media platforms for internet searches.

NETINT

Why ASICs are needed.

Density is a dirty expensive problem

Global corporations spend 20% of their annual OPEX powering data centers.

Data centers operate 24/7, massively consume energy, and are depleting our planet's resources at an accelerated and unsustainable rate. Today, there are 8,000 data centers globally and their collective consumption is expected to double by 2025.

NETINT

Our solution.

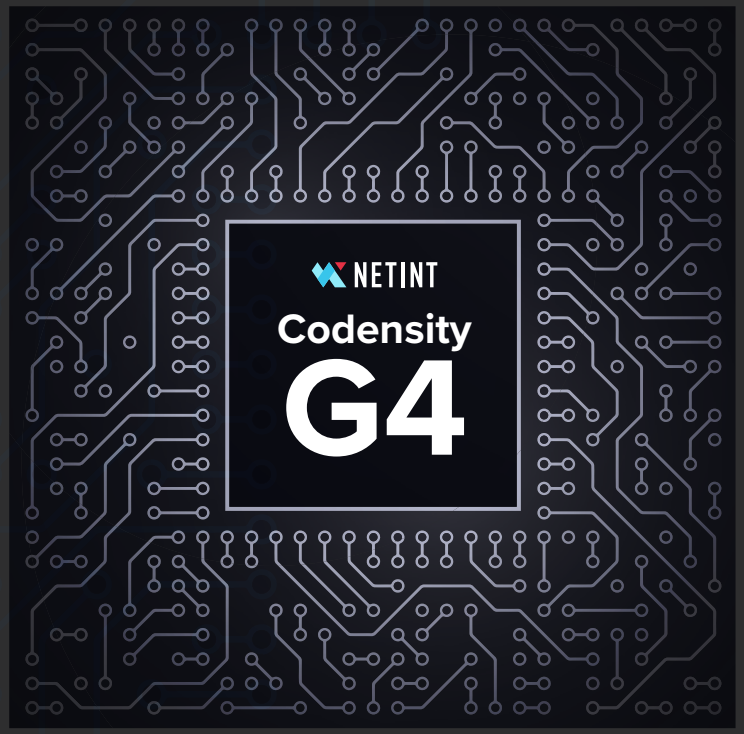

We designed an ASIC to slash the encoding footprint up to 40X

By replacing video encoding software with ASIC video transcoders, you deliver these benefits:

1. Massively increased capacity
2. Exponentially reduced power consumption

This saves your bottom line and the planet.
That's a win-win.

NETINT

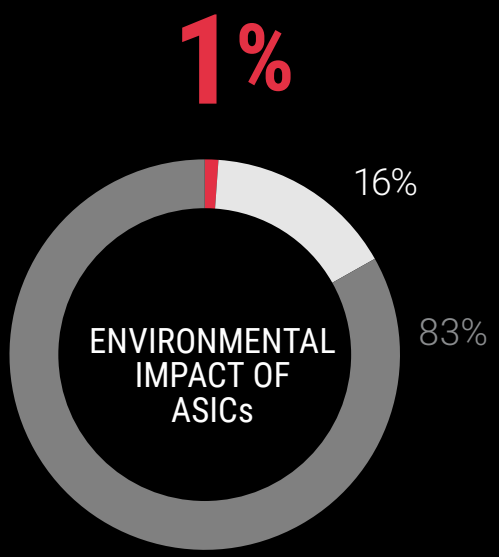

This is why Google built a custom ASIC for YouTube

For everyone else who isn't Google, we did the heavy lifting for you.

Google's ASIC called 'Argos,' helps YouTube encode and process videos much more efficiently. Argos chips provide "up to 20-33x improvements in compute efficiency compared to its previous traditional server set up," according to a Google executive. Another report suggests that Argos replaced over 10 million Intel CPUs in YouTube.

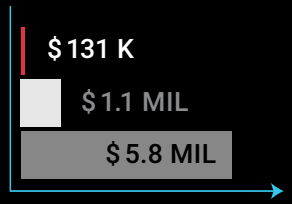

ANNUAL OPEX

Required to deliver 10K concurrent live HD streams

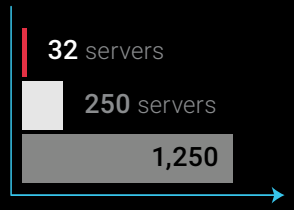

SERVER DENSITY

Servers required to deliver 10K concurrent live HD streams

ASIC *NETINT ASIC video processing units (VPU)*
GPU *NVIDIA T4 graphic processing units*
CPU *Intel SVT with CPU-based encoding*

NETINT

The real cost of live streaming

CAPEX and OPEX comparison using 3 video process/encoding scenarios.

Test assumptions:

- Servers run 100 concurrent five-rung encoding ladders
- x264 very fast preset used for CPU-only processing

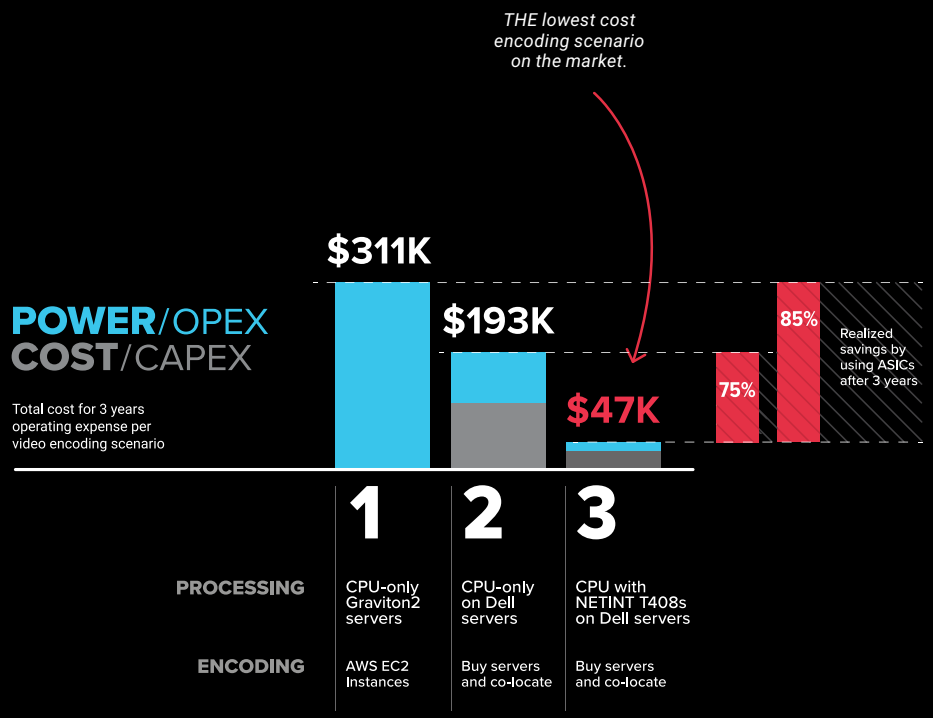

SOURCE: https://info.netint.com/hubfs/Whitepaper/Slash_CAPEX_OPEX_Carbon-Emissions-extended.pdf

NETINT

Logan
Video Server

Transcoding | Codensity ASIC G4

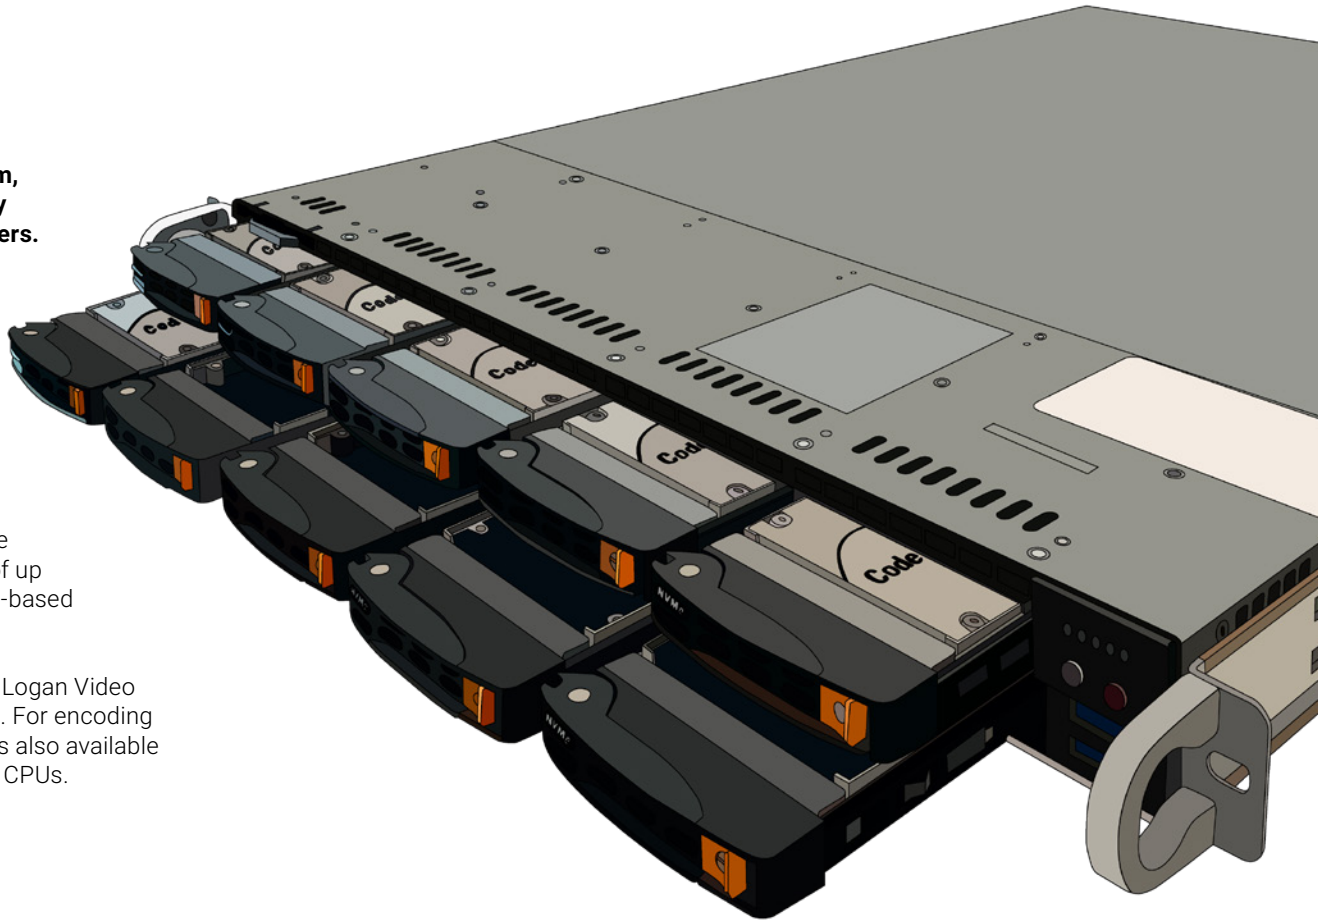

Built on the Supermicro 1114S-WN10RT server platform, NETINT's Logan Video Server boasts ultra-high density encoding capacity enabled by ten T408 video transcoders.

Supports:

- **HEVC and H.264 video encoding**
- **Up to 4K resolution**
- **10-bit HDR**

Ultra-low latency encoding of up to 80 broadcast quality 1080p30 streams in a compact 1RU form factor. Massive transcoding capacity enables breakthrough reductions of up to 80% in OPEX and CAPEX costs compared to software-based encoding systems.

Performance results in this brochure are for the NETINT Logan Video Server powered by an AMD EPYC™ 7543P (32-core) CPU. For encoding workflows with different encoding demands, the server is also available with the AMD EPYC 7232P (8-core) and 7713P (64-core) CPUs.

NETINT

T408
Video Transcoder

Codensity ASIC G4

The video engine inside.

The T408 video transcoder is a real-time, low-latency video transcoder. Available in U.2 form factor, the T408 enables hyperscale video platforms to easily transition from software to hardware encoding.

The T408 is based on the NETINT Codensity G4 ASIC that supports H.264 and HEVC video encoding at up to 4K resolution with 10-bit HDR. The high throughput of the T408 enables ultra-low latency encoding of up to 80 broadcast quality 1080p30 streams in a 1RU server using ten T408's.

NETINT

Codensity G4
ASIC

Application Specific Integrated Circuit

ASIC Video Transcoder

The Codensity G4 ASIC combines on-chip H.264 and HEVC video encoding, decoding, and processing engines which deliver scalability for video-intensive live streaming applications. The core of NETINT's Codensity technology is an in-house built ASIC that increases encoding density compared to CPU-based software encoding solutions.

This increase in encoding density expands the number of channels that can be encoded without increasing the rack footprint. Reduced power and HVAC cost means a lower TCO without sacrificing video quality or latency.

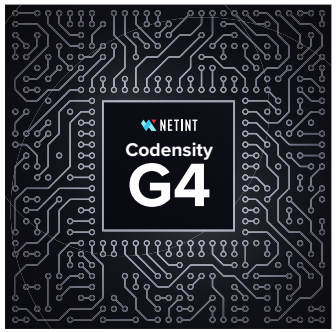

4K UHD Video Transcoding

On-chip H.264 and HEVC encoders and decoders deliver 4K live streaming scalability. Today, video is streamed using the ubiquitous H.264 standard while HEVC is a more complex codec. This limits the scalability of CPU and GPU-based encoders, which precipitously drop in throughput when encoding HEVC.

Not so for the Codensity G4 ASIC, which produces nearly identical throughput for both H.264 and HEVC. For both codecs, the Codensity G4 delivers the flexibility and quality of software with the performance of hardware for 4K live transcoding.

Flexible Architecture

The Codensity G4 is built on a programmable microprocessor architecture that optimizes firmware and pipeline processing and enables continual performance and quality improvements. This counters a criticism that silicon-based encoders lack upgrade flexibility.

Designed for the Cloud

High-density live UHD Transcoding

The T408 video transcoder takes full advantage of the video processing capability inside the Codensity G4 ASIC to support H.264 and HEVC live encode and transcode functionality of 4K UHD video in SDR and HDR with HDR10 and other popular high dynamic range standards. By offloading complex encode and decode processing to the Codensity G4 ASIC, the T408 video transcoder minimizes host CPU utilization. The result is a significant improvement in real-time transcoding density compared to any software or GPU-based transcoding solution.

As many as ten software-based video encoding servers may be replaced for every NETINT Logan Video Server that is installed in the data center.

High power efficiency

Each NETINT T408 U.2 module consumes only 7W of power at full load. This makes the Logan Video Server the most energy efficient video transcoder available.

Enterprise NVMe Integration

Available in the U.2 form factor, the T408 offers a simple upgrade path from CPU-based software to ASIC-powered video transcoding on any x86 or Arm-based server.

NETINT's Logan Video Server hosts ten T408 video transcoders supporting up to 80 simultaneous live 1080p30 transcoding sessions.

NETINT

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 T408 includes highly efficient FFmpeg and GStreamer compatible SDKs, allowing operators to apply a patch to complete the integration.

The libavcodec patch on the host server functions between the T408 NVMe interface and the FFmpeg and GStreamer software layers, simplifying integration and enabling fast and efficient performance and capacity upgrades.

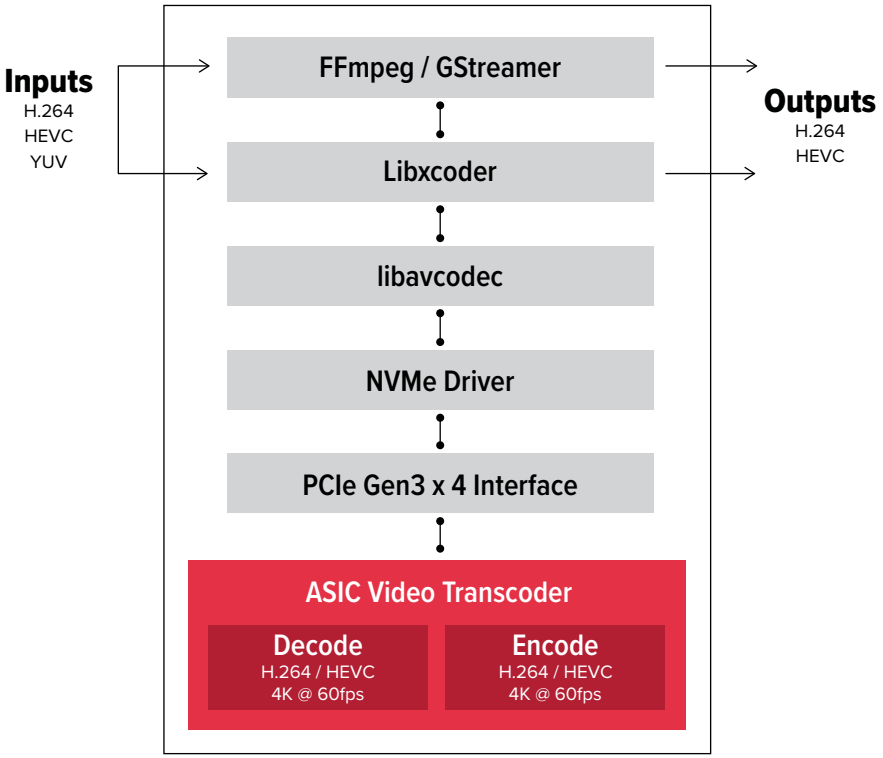

NETINT

Logan
Video Server

Transcoding | Codensity ASIC G4

CPU Options	AMD EPYC™ 7232P Server Processor (8-core)
	AMD EPYC 7543P Server Processor (32-core)
	AMD EPYC 7713P Server Processor (64-core)
Operating System	Ubuntu 20.04.05 LTS *(as of May 2023)*
Memory	8x 16GB DDR4-3200
Storage	400GB M.2 SSD
NVMe Support	10x
PCIe Expansion	Up to 3x PCIe slots
Network Options	Dual 10GBase-T LAN
Power Supply	700W: 100 - 140Vac
	750W: 200 - 240Vac
	750W: 200 - 240Vdc (CCC only)
Transcoders	10x NETINT T408
Encoding Capacity	Up to 10x 4Kp60 or 80x 1080p30 (HEVC and H.264)
Codec Support	H.264 - Encode/Decode
	HEVC - Encode/Decode
Transcoder Software	FFmpeg, GStreamer

Physical Dimensions	W: 17.2” (437mm), H: 1.7” (43mm), D: 23.5” (597mm)
Rack Size	1U
Weight	39 lbs (17.69 kg) *(includes 10 processors)*
Environmental	50 degrees F to 95 degrees F Operating Temperature, 8% to 90% Operating Relative Humidity
Power Inputs	100 - 140Vac / 8 - 6V / 50-60Hz
	200 - 240Vac / 4.5 - 3.8A / 50-60Hz
	200 - 240Vdc / 4.5 - 3.8A (CCC Only)
Certifications	RoHS Compliant, UL Approved

Specifications

T408
Video Transcoder

Codensity ASIC G4

Form Factor	U.2 (SFF-8639)
Interface	PCIe 3.0 x4
Protocol	NVMe
Power Consumption (Typ)	7W
Usage	24/7 Operation
Operation Temperature	0 degrees C to 70 degrees C
RoHS Compliance	Meets requirements of European Union (EU) ROHS Compliance Directives
Product Health Monitoring	Self-Monitoring, Analysis, and Reporting Technology (SMART) commands Temperature Monitoring and Logging
Hardware Interface	Available U.2 slot

	H.264 AVC Encode/Decode	**H.265 HEVC Encode/Decode**
Profile	CBP / BP / XP / MP / HiP / Hi10P	Main / Main 10
Level	1 to 6.2	1 to 6.2 Main Tier
Min / Max Resolution	Min: 32 x 32 / Max: 8192 x 5120	
Scan Type	Progressive	
Bitrate	64kbit/s to 700Mbit/s	
Software Integration	FFmpeg and GStreamer SDKs and direct integration with LibXcoder API	
Capacity	4K @ 60 fps | 1080p @ 240 fps	

Region of Interest (ROI)	ROI enables the quality of some regions to be improved at the expense of other regions
Closed Captioning	T408 supports EIA CEA-708 closed captions for H.264 and H.265 encode and decode
High Dynamic Range (HDR)	T408 supports HDR10 & HDR10+ for H.264 & H.265 encode and decode
Low Latency	T408 supports sub-frame latency
IDR Insert	Forced IDR frame inserts at any location
Flexible GOP Structure	8 presets plus customizable GOP structure

NETINT

Transcoding with Scaling

This table details the H.264 and HEVC output at the specified resolutions and frame rates; and the associated cost per stream. All inputs are scaled to the designated targets. Though the host CPU performs the scaling in these tests, **Note CPU utilization remains exceptionally low, reducing power costs and carbon emissions.**

Input	Output	Codec		FFmpeg	FFmpeg Low Delay	GStreamer	GStreamer Low Delay
4Kp30	1080p30	AVC > AVC	Instances	20	20	20	20
			CPU Usage	25.7	24.1	2.4	2.5
		AVC > HEVC	Instances	20	20	20	20
			CPU Usage	25.6	23.7	2.5	2.4
		HEVC > AVC	Instances	20	20	20	20
			CPU Usage	24.8	25.6	2.4	2.5
		HEVC > HEVC	Instances	20	20	20	20
			CPU Usage	24.2	25.2	2.4	2.6
1080p30	720p30	AVC > AVC	Instances	80	80	80	80
			CPU Usage	30.2	30.6	11.2	11.7
		AVC > HEVC	Instances	80	80	80	80
			CPU Usage	30.6	30.8	11.3	11.6
		HEVC > AVC	Instances	90	90	90	90
			CPU Usage	34.5	34.4	12.8	13.2
		HEVC > HEVC	Instances	90	90	90	90
			CPU Usage	35.4	35.1	12.8	13.1

Ladders

This table shows the number of full encoding ladders produced by the server and the cost per ladder for that output. Note the low CPU usage, despite all lower resolution rungs being scaled by the host CPU.

Input	Output	Codec		FFmpeg	FFmpeg Low Delay	GStreamer	GStreamer Low Delay
1080p30 - 5 Ladders	1080p30 @ 5Mbps	AVC > AVC	Instances	30	9	30	30
	1080p30 @ 3.5Mbps						
	720p30 @ 2Mbps						
	540p30 @ 1Mbps		CPU Usage	31.6	8.7	7.8	8.8
	360p30 @ 600kbps						
1080p30 - 4 Ladders	1080p30 @ 3.5Mbps	AVC > HEVC	Instances	26	14	28	28
	1080p30 @ 1.8Mbps						
	720p30 @ 1Mbps		CPU Usage	20.9	10.4	6.4	7.0
	360p @ 500kbps						
4Kp30 - 6 ladders	4Kp30 @ 12Mbps	AVC > HEVC	Instances	3	NA	7	7
	2Kp30 @ 7Mbps						
	1080p30 @ 3.5Mbps						
	1080p30 @ 1.8Mbps		CPU Usage	13.2		6.7	6.8
	720p30 @ 1Mbps						
	360p30 @ 500kbps						

Power Consumption

One of the key strengths of ASIC-based transcoders is ultra-low power consumption, which reduces OPEX and carbon emissions. You see this in the power figures, **particularly the Watts/Output, which are orders of magnitude lower than comparable figures for CPU-based transcoding.**

Operation	# Streams	Watts / Stream
Transcode 720p HEVC > HEVC, low delay	150	**2.1**
Scale 1080p > 720p HEVC to AVC, low delay	90	**3.4**
Five-rung AVC ladder, low delay	30	**10.8**

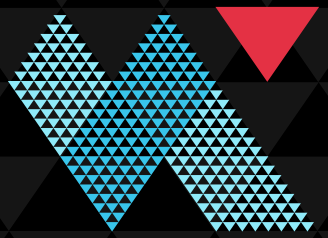

NETINT

HYPERSCALE PROFITABLY

For more information on NETINT encoding solutions, contact us.

go@netint.com

netint.com