



# Quadra Video Server

**Powered by NORSK**

High density video transcoding server with  
no-code/low-code workflow creation

# Norsk + NETINT

## The Perfect Synergy of Hardware and Software



### Low-cost, high-volume streaming hardware

NETINT's Quadra Video Server – Ampere Edition combines a 96-core Ampere ARM CPU and ten Quadra T1U video processing units (VPUs) in a Supermicro server chassis and costs \$19,000.

It's the ideal platform for low-cost, high-volume, high-quality streaming, with sufficient CPU cycles to host a range of transcoding and processing software.



### Simplified live video workflow creation

Norsk includes a low-code SDK and a no-code Studio interface that enable developers to easily create amazing, dynamic live video workflows and deploy at any scale. Norsk simplifies workflow creation, so you can concentrate on producing a differentiated experience for your customers and viewers—rather than the technical ins and outs of codecs or container formats.

- Norsk supports NETINT transcoders at the API level for efficient, fully featured transcoding
- Norsk runs efficiently on the Ampere Edition hardware, enabling a single 1RU solution for workflow orchestration, transcoding, and video processing
- If you're looking for the most cost-effective streaming solution, look no further—Norsk and NETINT represent the perfect synergy of technology and affordability

# Live streaming is hard



## Norsk makes it easy

Off-the-shelf tools are great, but they don't give you the freedom to create exactly what you and your customers need. And almost without exception, they're hard to customize and integrate.

But building a live video workflow from the ground up is extremely time-consuming and drains precious resources that would be better spent focusing on your customers.

Thankfully, there's a better way. Norsk by id3as lets you quickly and easily build amazing, dynamic viewer experiences, massively reducing development time and speeding time to market.

Norsk's battle-hardened technology has been used in the field for hundreds of thousands of live events per year for more than a decade, and now the low-code Norsk SDK and no-code Norsk Studio put that power in your hands. And Norsk can run anywhere—in the cloud, on-prem, or hybrid.

# Low code: Norsk SDK

The Norsk SDK lets you build any live streaming workflow—from single input, single output to complex multi-camera productions with source switching, browser overlays, and in-play highlights—in tens, or at most hundreds, of lines of simple-to-maintain code.



```
const source = await norsk.input.rtmpServer();
const browser = await norsk.input.browser({url: 'https://my.overlay.com...'});
const compose = await norsk.processor.transform.videoCompose(composeSettings);
const output = await norsk.output.whip();

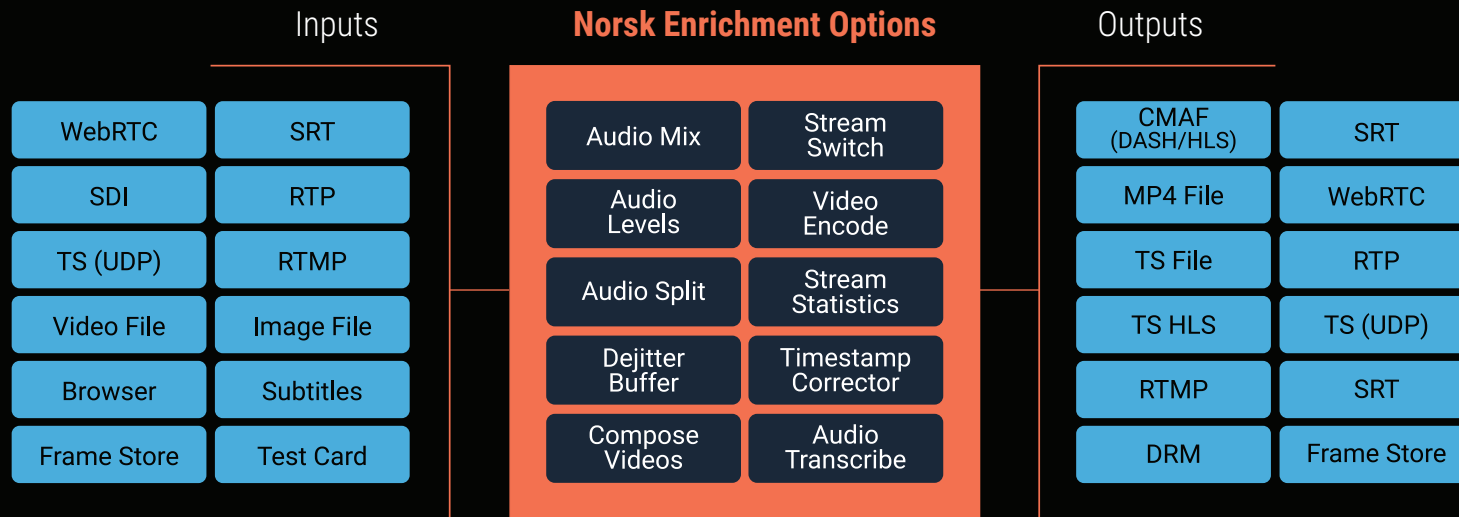
compose.subscribeToPins([
  { source: source, sourceSelector: videoToPin(background.pin) },
  { source: browser, sourceSelector: videoToPin(overlay.pin) },
]);

output.subscribe([
  { source: compose, sourceSelector: selectVideo },
  { source: source, sourceSelector: selectAudio },
]);
```

# Lots of ins and outs with options in between

Take inputs from almost any source and output them to almost any format while adding transformations, processing, and production functions in between. Norsk gives you access to all the most popular codecs and transport protocols.

Here are just a few of the inputs, outputs, and enrichment options available to you in both Norsk SDK and Norsk Studio. A comprehensive list can be found in our documentation [here](#).





**Inputs**

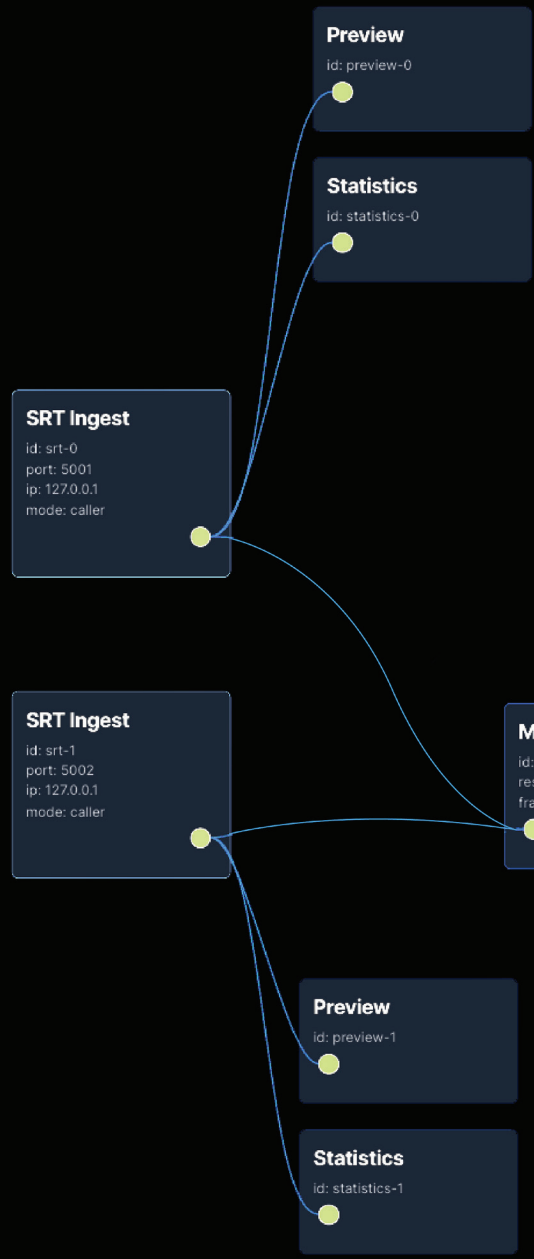
- Silence Generator
- SRT Ingest
- Video Test Card

**Processors**

- Browser Overlay
- Cascading Switch
- Encode Ladder
- MultiCamera

**Outputs**

- Auto CMAF
- Preview
- SRT Egest
- Statistics
- UDP TS Egest
- WHPE Egest



# No Code No Kidding

Norsk Studio is a no-code interface that makes using Norsk even simpler. It's ideal for both rapid prototyping and quickly building workflows from a predefined set of inputs, outputs, and processors.

Just drag and drop, connect the components, press play, and start your video sources.

# Norsk Studio is easy to extend

If the predefined nodes in Norsk Studio don't do exactly what you need, you can use the SDK to combine them with your own business logic to add to the no-code ecosystem.

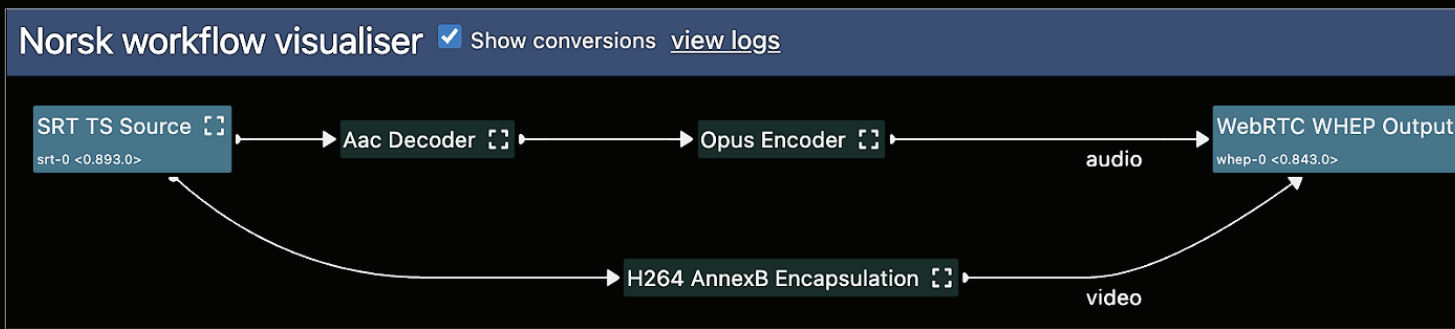

What's more, new open-source feature nodes are constantly being added to Norsk Studio.

MediaConnect Input	Fixed Ladder
MediaConnect Output	Cascading Switch
Browser Overlay	Stream Statistics
Monitor Wall	Instant Replay & Highlights
2-Camera Vision Director	Audio Mixer

# Visualize the flow

Norsk understands the intricacies of video and audio technology so you don't have to.

While you execute complex workflows using simple commands or the drag-and-drop interface, Norsk automatically executes conversions and transformations under the hood—but all of that detail is available for you to inspect in the included visualizer.

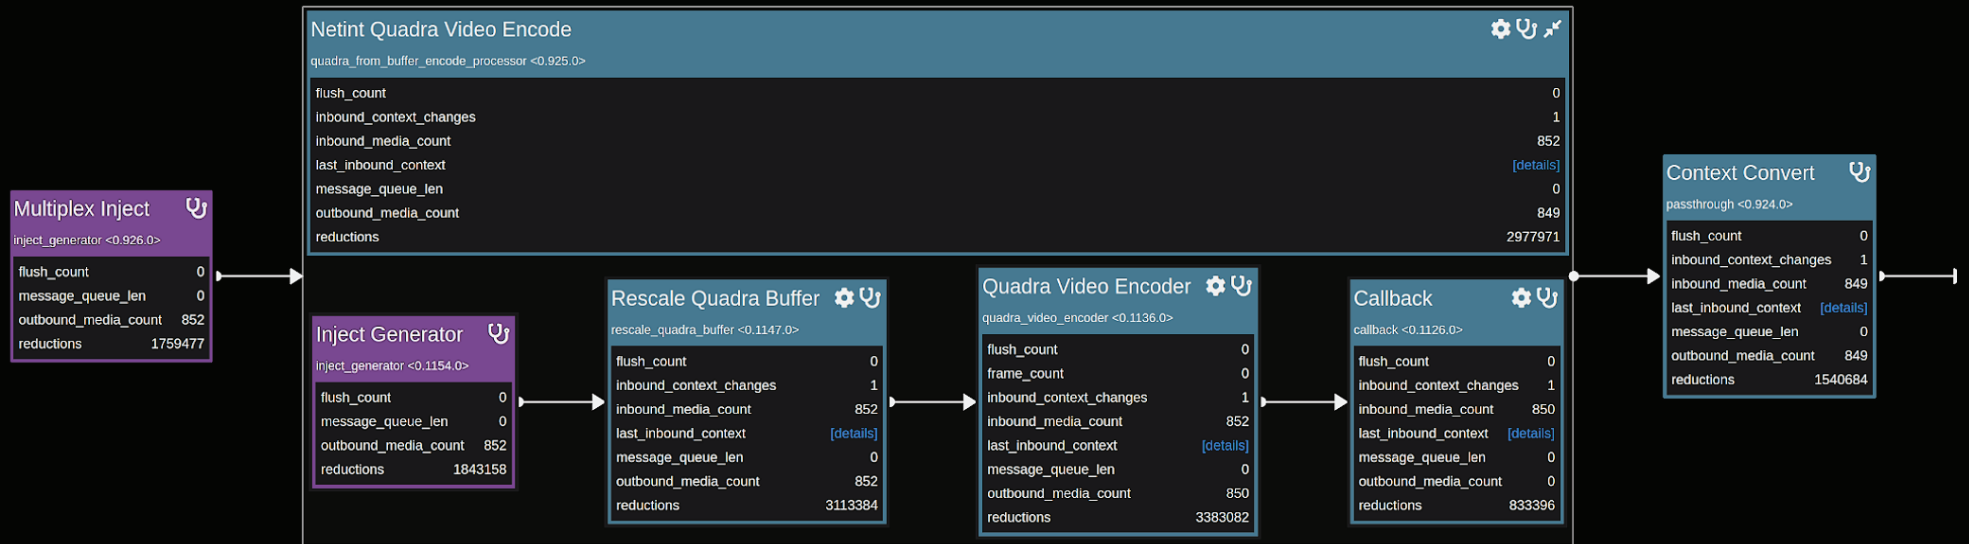




# Norsk + NETINT

When running on NETINT VPU hardware, Norsk automatically maximizes the VPU's capability by performing operations such as picture-in-picture, scaling, and overlays natively—on the hardware—not on the CPU. This leverages the efficiency of the hardware and optimizes overall system performance.

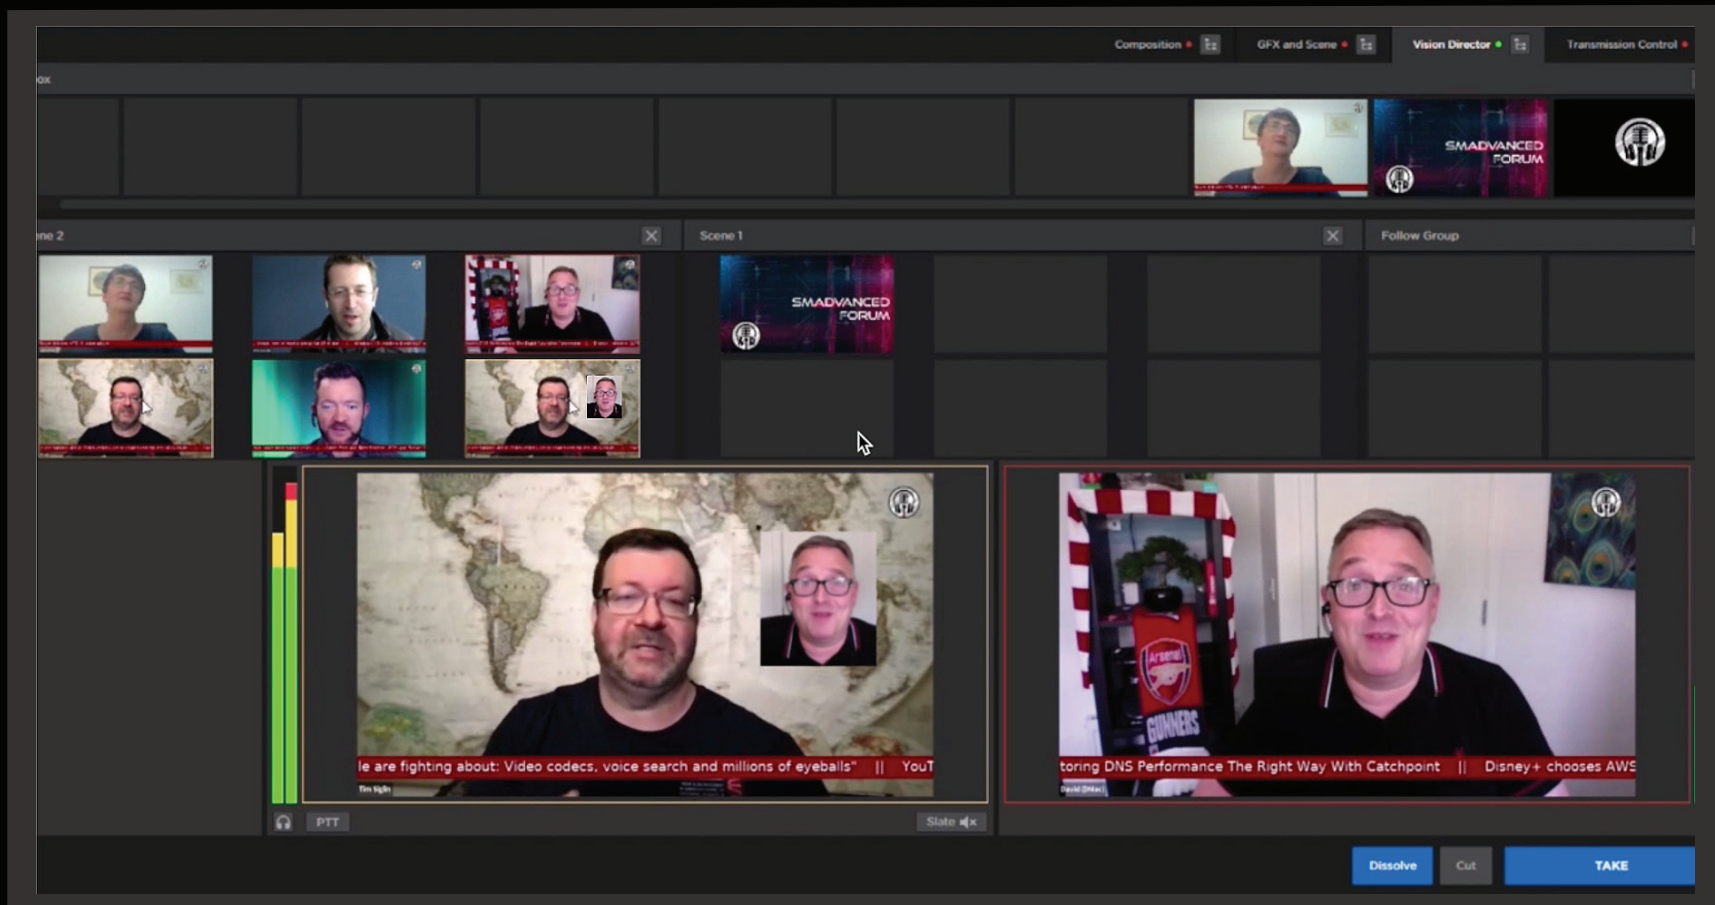
You can deliver hundreds of channels—each with complex processing, transformation, and encoding—**using a single 1RU rack server**. The VPU hardware also frees up the CPU for other pertinent tasks such as speech-to-text or AI processing.



# Example: Remote production

That's everything you need to create a vision mixer. Now any director can control a multi-camera production with stings, browser overlays, and picture-in-picture compositions.

This is just an example—Norsk's capabilities are limited only by your imagination.




# NETINT Hardware

The Ideal Platform for Norsk

## We designed an ASIC to slash the encoding footprint up to 80%

By replacing CPU-software encoding with Smart VPUs you get:

1. Increased encoding capacity using fewer VPU chips
2. Fewer chips require smaller hardware footprint
3. Less hardware consumes less power



# This is why Google built a custom chip for YouTube

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

We developed commercial-ready Smart VPU cards for easy drop-in replacement and immediate deployment.

## SMART VPU


NETINT video processing units powered by AI

## GPU

NVIDIA T4 graphic processing units

## CPU

INTEL SVT encoding software




## SERVER DENSITY

Servers required to deliver 10,000 concurrent HD streams

## SMART VPU

## GPU

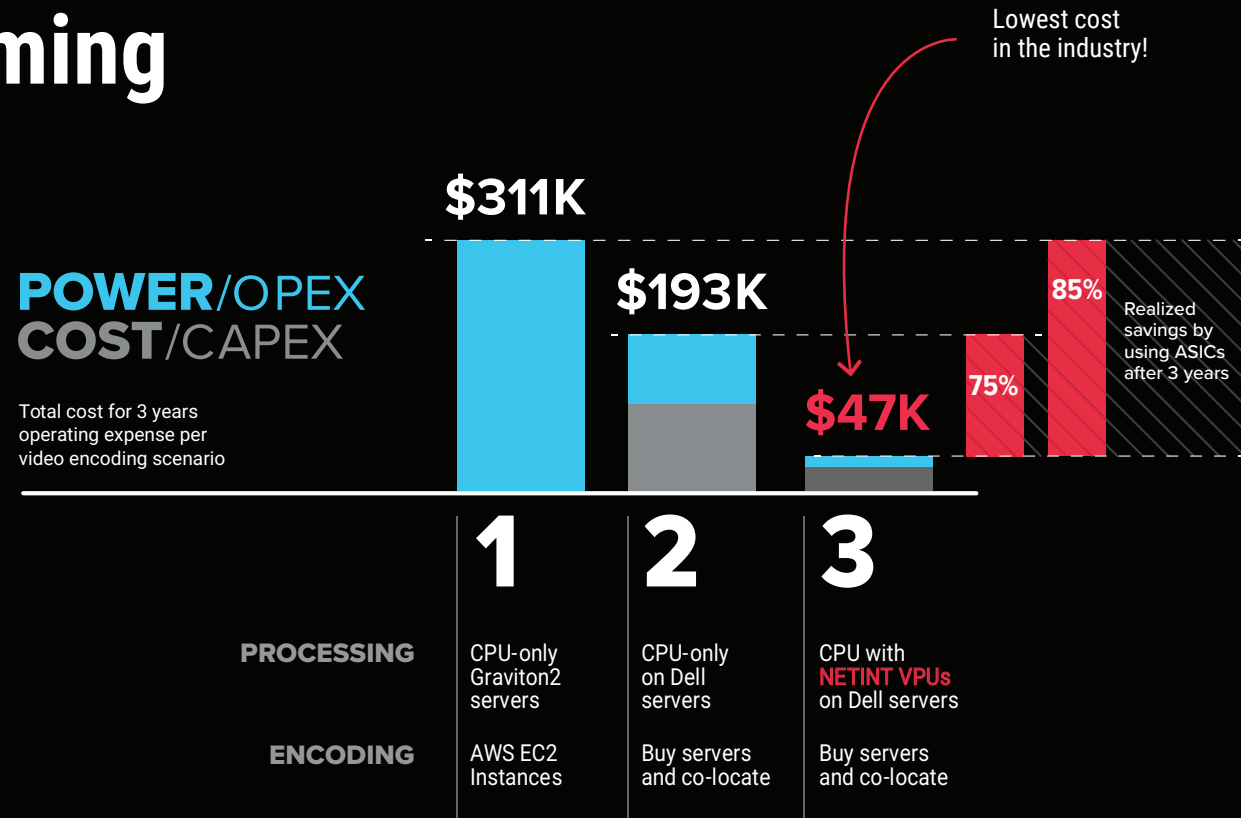
## CPU



## ANNUAL OPERATING COST

Cost running servers required to deliver 10,000 concurrent HD stream

# The real cost of live streaming




Test assumptions:

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




SOURCE: Slash CAPEX & OPEX

# Quadra Specifications



# Quadra Video Server

Ampere Edition



Ultra-high density, low cost and low power

- HEVC, H.264 and AV1 encoding
- HEVC, H.264, and VP9 decoding
- Up to 8K resolution
- 10-bit HDR

Ultra-low latency encoding of up to 320 broadcast quality 1080p30 streams in a compact 1RU form factor. Massive transcoding capacity enables breakthrough reductions of up to 90-95% in OPEX and CAPEX costs compared to software-based encoding systems.


The high-performance 96-core CPU boosts performance during transcoding-related operations not supported on the Quadra hardware, like deinterlacing and AV1 and MPEG-2 decode. Video engineers can also run additional publishing-related applications on the server, like Whisper real-time audio transcription, dynamic ABR packaging, streaming orchestration, or content management. By consolidating these operations on a single server, the Ampere Edition delivers unprecedented transcoding and publishing density, saving CAPEX, OPEX, and rack space.



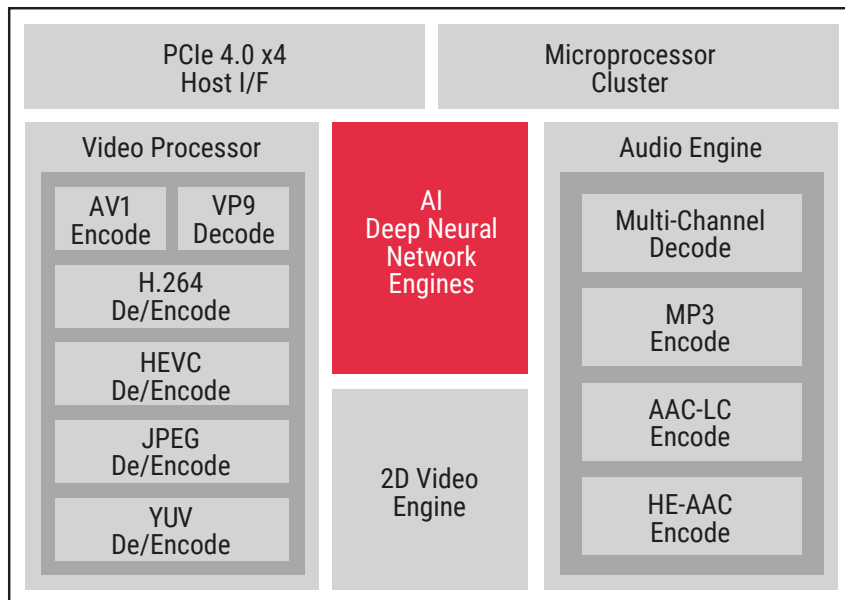
# Codensity G5

Smart VPU

ASIC chip, powered by AI



The Codensity G5 architecture uniquely combines on-chip AV1, H.264 and HEVC video encoding and AI processing engines to deliver encoding density and scalability for metaverse, live streaming, and interactive applications.



## 8K UHD Video Encoding

Our Codensity G5 ASIC enables up to 8K video transcoding using HEVC and H.264 codecs (AV1 is limited to 4K). Advanced codecs like AV1 and HEVC deliver superior quality with up to a 60% reduction in bitrate.

## Flexible Architecture

The Codensity G5 is built on a programmable micro-processor architecture to optimize the firmware and pipeline processing for improved performance and increased video quality.

## AI Engine

Two Deep Neural Network engines capable of up to 18 TOPS (trillion operations per second) enable object detection, classification and segmentation for image quality improvement and content-adaptive rate control. Advanced performance and seamless integration also for region-of-interest (ROI) encoding and background replacement.



# Designed for the Cloud

## High-density live UHD transcoding

The NETINT Quadra VPU takes full advantage of the video processing capability inside the Codensity G5 ASIC to support H.264, HEVC, and AV1 live encode functionality of up to 8K UHD video. By offloading complex en/decode processing to the Codensity G5 ASIC, the Quadra VPU minimizes host CPU utilization. The result is a significant improvement in real-time transcoding density compared to any software or GPU-based processing solution.

Every NETINT Quadra Server installed in a data center would replace as many as 25 software-based video encoding servers.

## High power efficiency

Each NETINT Quadra U.2 module consumes only 17W power at full load. This makes the Quadra Video Server, the most energy efficient video transcoder available.

## Enterprise NVMe integration

Deployed in a U.2 and HHHL AIC form factors, Quadra offers a simple upgrade path from CPU-based software to ASIC video encoding on any enterprise-class server.



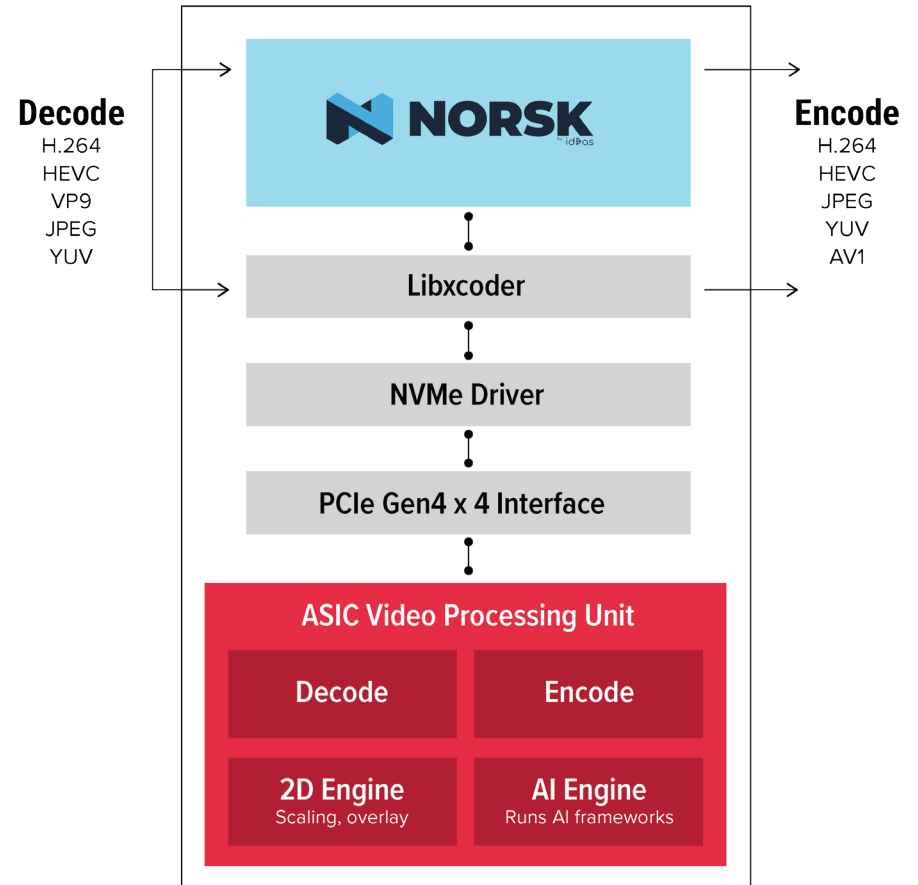
**NETINT's  
Quadra Video Server  
hosts ten Quadra VPUs supporting  
up to 320 simultaneous live 1080p30  
encoding sessions.**

# Norsk / NETINT Integration

## Open-source processing tools

Norsk integrates with Quadra at the API level, ensuring feature-rich, low-latency, and responsive control over Quadra's hardware-accelerated decode, scaling, overlay, and encode functions.

In operation, Norsk automatically performs operations like picture-in-picture and scaling and overlays natively on the hardware rather than the CPU, which leverages Quadra's hardware efficiency and optimizes overall system performance.



# Quadra T2A

## Smart VPU

Video Processing Unit with AI | Codensity G5



<b>Form Factor</b>	AIC (HH HL)
<b>ASIC</b>	2x Codensity G5
<b>Interface</b>	PCIe 4.0 x4x4
<b>Power Consumption (Typ)</b>	40W
<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 64x 1080p30, 16x 4Kp30, 4x 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>	36 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 T1A

## Smart VPU

Video Processing Unit with AI | Codensity G5




<b>Form Factor</b>	AIC (HH HL)
<b>ASIC</b>	1x Codensity G5
<b>Interface</b>	PCIe 4.0 x4
<b>Power Consumption (Typ)</b>	20W
<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>	18 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 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



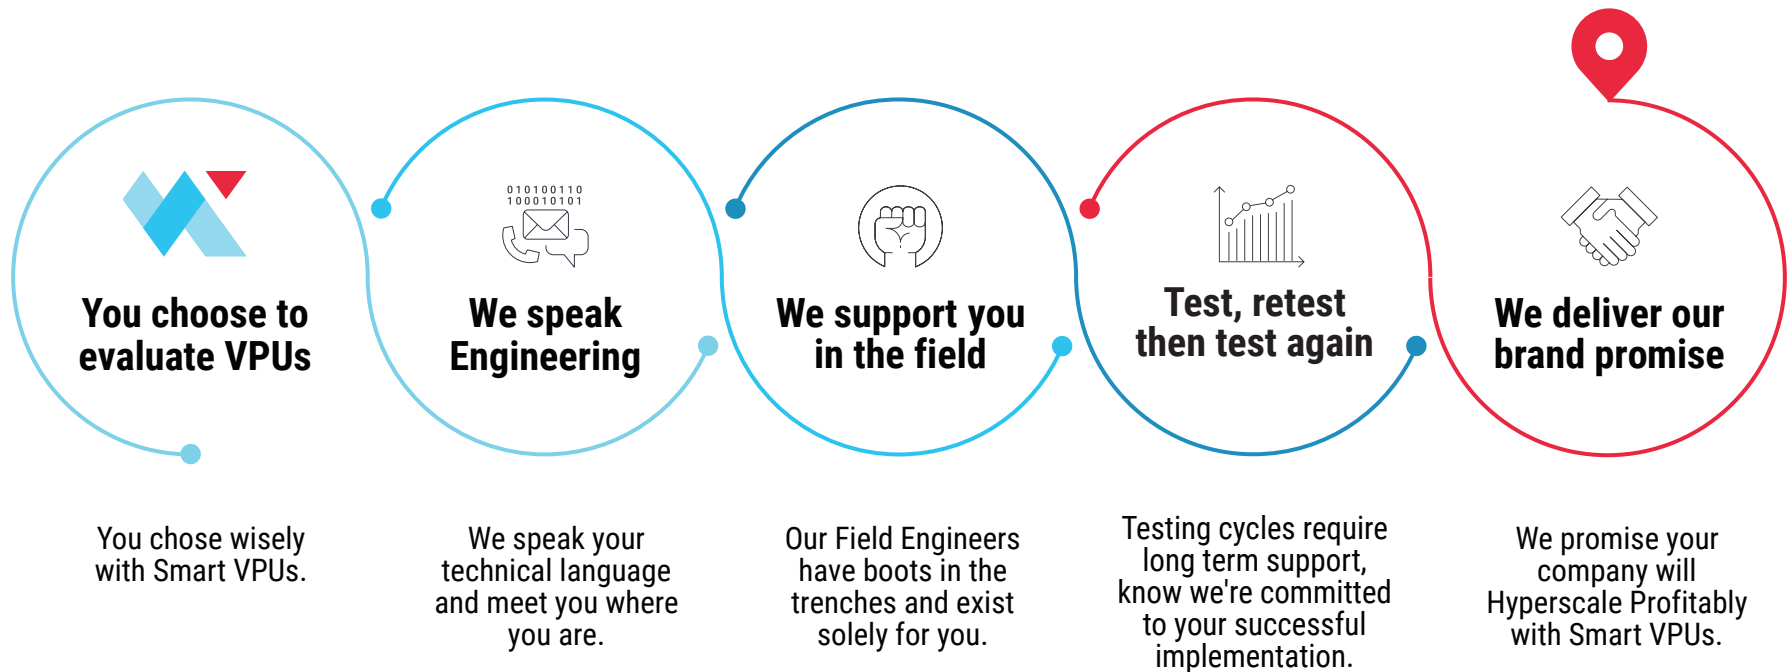
<b>CPU Options</b>	Ampere Altra CPU M96-28 CPU
<b>Operating System</b>	Ubuntu 22.04.3 LTS
<b>Memory</b>	256GB of DDR4-3200 RDIMM
<b>Storage</b>	400GB M.2 SSD
<b>NVMe Support</b>	10x
<b>PCIe Expansion</b>	Three PCIe 4.0 x16 LP slots, one PCIe 4.0 x16 AIOM slot
<b>Network Options</b>	1 RJ45 Dedicated IPMI 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

# Your Buying Journey

## What to expect when evaluating NETINT

We know the typical sales cycle prospective buyers endure is a 12-18 month process and we're prepared to stand beside you and navigate you through. We're demonstrating our commitment to supporting you by heavily investing in this process so you can realize the value in our product and in working with us.





sales@netint.com  
netint.com



sales@norsk.video  
norsk.video