

SharpMedia™ PCIE-8120 PCI Express Media Processing Accelerator

Data Sheet

The SharpMedia™ PCIE-8120 media processing accelerator enables high density voice and video processing to be integrated into a rack mounted server or other network appliances that feature standard PCI Express slots.


- Single slot full length full height PCI Express card with x4 interface
- High performance media processing core based on power-efficient DSPs
- Optional 2 x GbE ports (RJ45) with NAT function for direct network attachment providing server offload
- Comprehensive voice and video processing firmware and programmers interface included
- Support for 720p and 1080p video conferencing

The Artesyn Embedded Technologies SharpMedia PCIE-8120 is a PCI Express (PCIe) media processing acceleration engine to accelerate voice and video applications in standard server architectures or other appliances that feature full length PCIe slots.


Voice and video streaming is becoming more pervasive as customer demand for media consumption continues to rise, and the diversity of media sources, network conditions and individual consumption devices proliferate. Using traditional Intel® Xeon® server architectures, power and efficiency soon become a limitation when scaling to high channel density media stream transcoding. By employing highly efficient digital signal processors (DSP) to provide direct offload of the host CPU, the new SharpMedia PCIE-8120 acceleration engine can take the place of additional servers when adding high density voice and video processing to an application. Network equipment providers can either add or substantially increase the voice channel or video processing density within existing systems, so reducing overall power and space demands for their equipment as it scales to higher throughput. This board is particularly suited to the following applications:

- Session border controllers (adding voice or video transcode)
- Media gateways
- Media servers/media resource function
- Video/content optimization (transcode and transrating)
- Video communications servers
- Interactive voice and video response systems

SharpMedia PCIE-8120 brings high density DSPs with embedded voice and video firmware from Octasic together with Artesyn's strong embedded system heritage and thermal design skill, resulting in an industry-leading media processing density for the next generation of voice and video processing systems.



SharpMedia PCIe-8120 Block Diagram



SharpMedia™ PCIe-8120 is based around the Octasic OCT2224M multi-core DSP running Vocallo MGW firmware. The DSP array performs media processing acceleration for a host server, supporting both voice and video conferencing and transcode applications.

Media acceleration performance depends on both the codecs required and the number of DSPs available. A variety of board configurations allow for differences in application needs and server capabilities, and include a choice of 4, 8 and 12 DSPs, with maximum power consumption between 25W and 65W. For systems with limitations on PCIe slot power, an external power connection option is available by special arrangement.

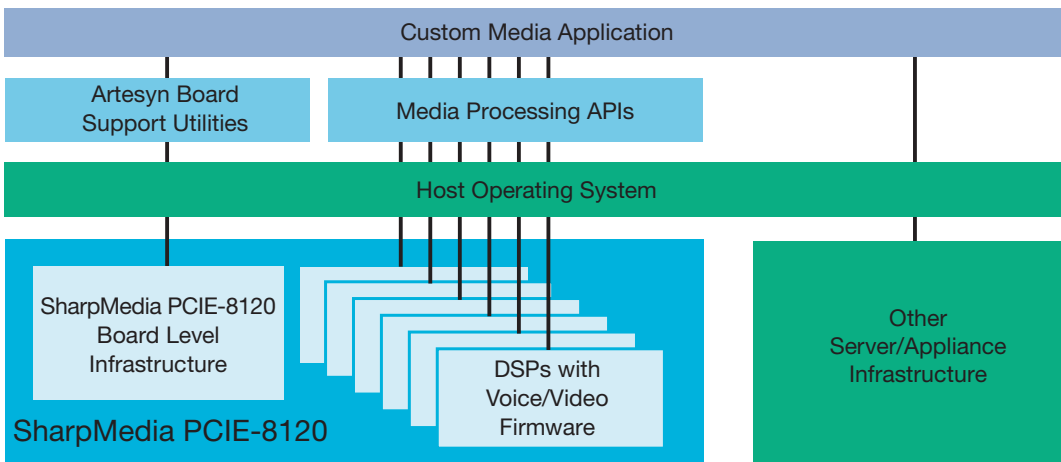
A comprehensive host-based media processing Application Programmers Interface (API) is provided. This is used to configure and execute voice and video stream processing functions. The API commands communicate directly with the DSP array based on an endpoint and stream resource model.

A non-blocking command/ response protocol aids multi-channel programming efficiency. Additional board support utilities can set the board's internal switching infrastructure into various modes, and provide diagnostic information.

The internal data flows of the board are all based on Gigabit Ethernet connections with all DSPs accessible via a local Ethernet switching subsystem. Host access to all the DSPs is via a 2 x 1Gb/s PCI Ethernet controller. The Ethernet switching subsystem also supports two 1Gb/s links to each DSP to support special 1080p video conferencing modes.

Media streams can be routed to the DSPs either via the host CPU or optionally via two external Gigabit Ethernet ports provided for direct traffic termination. In this configuration, packets for transcode can bypass the host computer entirely while a special Network Address Translation (NAT) device makes the hardware and software architecture.

Software Architecture



Media Processing Functions

TRANSPORT AND ENDPOINT SUPPORT

- Voice and video over IP endpoints
 - RTP/UDP/IP endpoint
 - IPv4 with DHCP & IPv6
 - Supports secure RTP
 - Adaptive jitter buffer
 - Supports RTCP (RTCP-XR in future)
- Line echo cancellation and voice quality enhancements
 - G.168 (2004) compliant line echo cancellation, up to 128ms tail, with HLC and music protection
 - Manual and automatic level control (G.169)
 - Acoustic Echo Control (AEC)
 - Adaptive Noise Reduction (ANR)
 - Natural Level Enhancement (NLE)
- Endpoint statistics
 - RTP/UDP/IP per channel and per port packets and errors
 - Per channel/port states, terminations used, media stream events

AUDIO STREAM PROCESSING

- 3GPP/3GPP2 wireless voice codecs
 - GSM Full Rate (FR) and Enhanced Full Rate (EFR)
 - GSM Adaptive Multi-Rate (AMR)
 - GSM wideband AMR-WB (G.722.2)
 - EVRC and EVRC-B (subject to additional upgrade license fee)
- ITU-T, IETF, and other voice codecs
 - Clear channel operation
 - G.711 μ -law/A-law with appendices I and II silence suppression with spectral comfort noise generation
 - Narrowband: G.729AB, G.723.1, G.726
 - Wideband: G.722, G.722.1
 - T.38 fax relay
 - V.152 voice band data
 - iLBC
 - SILK (Skype codec) – subject to additional upgrade license fee
 - OPUS – *subject to additional upgrade license fee*
- In-band announcements and signaling support
 - Announcements and message playback
 - Tone detection, generation and removal
 - SS5, MFR1, MFR2, DTMF
 - Tone relay (RFC 2833/4733)
 - Caller ID generation and detection (ETSI, BellCore and NTT)

- N-way audio conferencing
 - All channels on a DSP can participate in one or many conferences
 - Noise reduction for conferencing
 - Prime speaker detection
 - Conference supervisor function

VIDEO STREAM PROCESSING

- Video encoder/decoder support
 - H.263 profile 0 (baseline) and 3 level 10, 20, 30, 40, 45 and 50 (CIF/QCIF) RFC4629
 - MPEG-4 simple profile, single object level 0, 1, 2 and 3 (CIF/QCIF) RFC3016
 - H.264 Baseline Profile and High Profile level 1, 1b, 1.1, 1.2, 1.3 (CIF/QCIF) level 2.2 (VGA), 3.2 (720p) and 4.2 (1080p) RFC3984
- Scaling and adaptation
 - Advanced-CBR, CBR and VBR rate controllers
 - Frame rate adaptation up to 30 FPS/channel
 - Scaling between resolutions, from SQCIF to 1080p
 - Packet loss concealment
 - A/V synchronization (lip sync) via delay compensation
- Video conferencing/mixing support
 - Video MCU function
 - Voice activity based source selection
 - Graphics overlay with key color
 - Video mix: up to 16 layers per stream for video or graphics

Technical Specifications

HARDWARE

- Form factor: PCIe long card
 - Full height (106 mm) x full length (311 mm), single-slot width
 - PCIe X4 Gen 2 electrical connection
- DSP core
 - Multicore Octasic OCT2224M DSPs running Vocallo
 - MGW firmware
 - Standard variants: 4, 8 or 12 DSPs
- External Ethernet port option
 - 2 x Gigabit Ethernet (GbE) 1000BASE-T RJ45
 - Network address translation function makes DSP array look like single IP address (up to 1Gbit/s)

- Internal Ethernet infrastructure capacity
 - 2 x 1Gbit/s to host via PCI Ethernet controller
 - 2 x 1Gbit/s to each DSP to support 1080p video conference function
 - 2 x 1Gbit/s to external ports
- Board control and management functions
 - On-board temperature monitoring and reporting
 - Individual DSP and other function resets

DEPLOYMENT ENVIRONMENT

- Board power consumption (estimated worst case)
 - 4 DSP variant: 25W
 - 8 DSP variant: 46W
 - 12 DSP variant: 65W
- Board operating temperature
 - Normal operation: 0 °C to 40 °C

- Cooling requirement
 - Passive heatsinks (requires forced air flow)
 - Approx 4CFM for 12 DSP card at 40 °C ambient
 - Approx 5CFM for 12 DSP card at 55 °C ambient

HOST SOFTWARE ENVIRONMENT

- Host operating system
 - Red Hat Linux 6.3/CentOS
- Programmers environment
 - Octasic Vocallo MGW SDK and documentation
 - Octasic debug tools
 - Artesyn SharpMedia™ PCIE-8120 board support and configuration

Standard Board Configurations

Part number	Number of DSPs	External GBE ports + NAT function	Ext Power	Audio processing	Video processing
PCIE-8120-A12-N-PP	12	No	Slot	Yes	No
PCIE-8120-V12-N-PP	12	No	Slot	Yes	Yes
PCIE-8120-A12-N	12	No	Yes	Yes	No
PCIE-8120-V12-N	12	No	Yes	Yes	Yes
PCIE-8120-A04	4	Yes	Slot	Yes	No
PCIE-8120-V04	4	Yes	Slot	Yes	Yes
PCIE-8120-A08	8	Yes	Slot	Yes	No
PCIE-8120-V08	8	Yes	Slot	Yes	Yes
PCIE-8120-A12	12	Yes	Slot	Yes	No
PCIE-8120-V12	12	Yes	Slot	Yes	Yes

Representative Performance*

Media processing application	4 DSPs	8 DSPs	12 DSPs
Standard voice gateway/transcode: G.711 (20ms) <=> G.729AB (20ms)	2560 channels	5120 channels	7680 channels
Wireless voice gateway/transcode: G.711(20ms) <=> AMR NB (20ms)	1716 channels	3432 channels	5148 channels
Mobile video transcode MPEG4/CIF => H.264/CIF 15fps	128 channels	256 channels	384 channels
HD video conference H.264/720p 30fps x 4 participant conferences	2 bridges	4 bridges	6 bridges

*Note: exact performance depends on many factors including exact codec mix and use patterns. These figures are provided as guidance to potential performance.

Ordering Information	
<i>PCIE Boards</i>	<i>Description</i>
PCIE-8120-A04	PCIE card, 4X OCT-2224M, audio license
PCIE-8120-A08	PCIE card, 8X OCT-2224M, audio license
PCIE-8120-A12	PCIE card, 12X OCT-2224M, audio license
PCIE-8120-A12-N	PCIE card, 12X OCT-2224M, audio license, no RJ45 GBE, external power
PCIE-8120-A12-N-PP	PCIE card, 12X OCT-2224M, audio license, no RJ45 GBE, no external power
PCIE-8120-V04	PCIE card, 4X OCT-2224M, audio and video license
PCIE-8120-V08	PCIE card, 8X OCT-2224M, audio and video license
PCIE-8120-V12	PCIE card, 12X OCT-2224M, audio and video license
PCIE-8120-V12-N	PCIE card, 12X OCT-2224M, audio and video license, no RJ45 GBE, external power
PCIE-8120-V12-N-PP	PCIE card, 12X OCT-2224M, audio and video license, no RJ45 GBE, no external power
Software Upgrades	
<i>Marketing Number</i>	<i>Description</i>
8XXX-UPGR-V04 / -V08 / -V12	Video license upgrade for audio only boards - 4 DSPs / 8 DSPs / 12 DSPs
8XXX-UPGR-EVRC-X04 / -X08 / -X12	PCIE card, 8X OCT-2224M, audio license
Development Support Packages	
<i>Marketing Number</i>	<i>Description</i>
OEM-DEV-JMP	Jump Start Package - 3 month period - Includes 1 day training session, regular expert meetings
OEM-DEV-ADP	Advanced Developer Package - 3 month period - includes 1 day on-site consultation, regular expert meetings + 10 service requests

IMPORTANT NOTICE - STANDARDS ESSENTIAL PATENTS AND THE USE OF CODECS

The Artesyn End User License Agreement covering software related to the ATCA-8120 does not represent or warrant that the codec software is free of infringement of any third party patents, copyrights, or trade secrets. Many codecs and other recognized standards may require licensing arrangements involving the execution of license agreements or payment of fees to an intellectual property rights (IPR) holder or an IPR agent acting on behalf of the IPR holder.

It is the user's responsibility to determine, for any codecs or other standards they intend to use, whether any additional IPR licenses are required, including the payment of royalties or license fees. The availability of implementations including codecs packaged in products acquired from Artesyn does not imply the right to practice these standards nor does Artesyn grant a license or the right to use or practice some or all of these standards. Depending on the country involved, the end user may be legally obliged to contact an IPR holder or agent and conform to their patents licensing requirements.

STANDARDS ESSENTIAL PATENTS

Standards Essential Patents (SEPs) are an unavoidable consequence of complex standards developed by consortia. Patents are essential when the technology covered by the patent must be practiced in order to comply with the Standard. A patent is typically defined as essential if a standard cannot be practiced without infringing the patent. The contributor of the IP is usually a company involved in the standards process and almost always retains ownership. Companies that own SEPs that are often members of standards setting organizations (such as ETSI or IEEE) and may be required to declare that they will license their patents on Reasonable And Non-Discriminatory (RAND) terms. Most standards organizations do not review patents alleged to be essential to determine that they in fact are essential. This is a determination that may not occur until decided in a contested legal matter.

COVERAGE FOR STANDARDS-ESSENTIAL PATENTS (See Table 1)

Artesyn does not generally provide indemnification against infringement of SEPs related to codecs, but there are some exceptions where Artesyn does offer limited indemnification. Table 1 summarizes Artesyn's understanding of licensing requirements with respect to selected codecs. Except where specifically stated that Artesyn does offer indemnification, it should be understood that Artesyn does not offer indemnification.

COVERAGE FOR NON-ESSENTIAL PATENTS

Non-Essential patents are by definition patents that are not necessarily infringed in order to practice the standard implemented by a codec. As part of a custom license, Artesyn may offer indemnification against infringement of non-essential patents with respect to the codecs embedded in the ATCA-8330 software. Contact your Artesyn sales representative for further details.

A NOTE ON WIRELESS VOICE CODECS

Although the patent situation for wireless voice codecs is very complex, it is not standard industry practice for embedded board-level product manufacturers like Artesyn to offer IPR licensing cover for AMR and EVRC codecs because they are not well placed in the IPR value chain. Therefore, Artesyn does not typically offer indemnification for AMR and EVRC codecs. The IPR holders usually approach end-product manufacturers because many end-products include additional standards compliant technologies that may infringe patents. Many IPR holders prefer to offer portfolio licenses that cover much more than just the codecs. Additionally, a starting point for licensing is based on the number of channels used in a product, and the end-product manufacturer is much better placed to quantify and control usage than the embedded technology provider. Artesyn strongly recommends that customers undertake a full product mapping exercise to determine the feature sets that need to be covered by patent portfolios.

Table 1

01 October 2014

CODEC FAMILIES AND STANDARDS ESSENTIAL PATENTS

Class	CODEC Family	Standards Essential Patent Situation
Wireline Voice	G.711 PCM	Artesyn believes that these codecs are currently unencumbered by standards essential patents.
	G.722 Wideband	
	G.726 ADPCM	
	G.723.1	Indemnification for infringement of SEPs for these codecs is included within the Artesyn custom license described here.
	G.729AB	
	G.722.1	G.722.1 (also known as "Siren") is licensed royalty free by Polycom provided that the end-product manufacturer executes the license at http://www.polycom.com/company/about-us/technology/siren.html . Customers interested in this codec should contact Polycom directly.
iLBC	iLBC is an open source royalty-free codec available directly to Customers under the "revised BSD" license. The full text of the Revised BSD License can be found at: http://opensource.org/licenses/BSD-3-Clause .	
Wireless Voice	GSM FR	GSM and GSM AMR codecs may incorporate SEPs held by Ericsson, Voiceage, Nokia, NTT, and France Telecom. Please see "A NOTE ON WIRELESS VOICE CODECS".
	GSM AMR	
	GSM AMR WB	
	EVRC-A	EVRC codecs may incorporate SEPs held by Qualcomm, Ericsson, NTT, France Telecom and others. Please see "A NOTE ON WIRELESS VOICE CODECS".
	EVRC-B	
Internet Voice	SILK	SILK can be used royalty free under a patent license at http://developer.skype.com/silk/license . Customers interested in SILK should contact Skype directly.
	SPEEX	Speex is an open source royalty-free codec available directly to Customers under the "revised BSD" license (see iLBC above for reference).
	Opus	Opus is an open source royalty-free codec available directly to Customers under the "revised BSD" license (see iLBC above for reference).
Video	H.263	MPEG LA (www.mpegla.com) offers coordinated patent licenses for video codecs to end-product manufacturers. Customers interested in these codecs should contact MPEG-LA directly.
	MPEG-4	
	H.264	

SOLUTION SERVICES

Artesyn Embedded Technologies provides a portfolio of solution services optimized to meet your needs throughout the product lifecycle. Design services help speed time-to-market. Deployment services include global 24x7 technical support. Renewal services enable product longevity and technology refresh.

WORLDWIDE OFFICES

United States	+1 888 412 7832	Japan	+81 3 5403 2730
Hong Kong	+852 2176 3540	Korea	+82 2 6004 3268
China	+86 400 8888 183		

Artesyn Embedded Technologies, Artesyn and the Artesyn Embedded Technologies logo are trademarks and service marks of Artesyn Embedded Technologies, Inc. Intel and Xeon are trademarks of Intel Corporation or its subsidiaries in the United States and other countries. All other names and logos referred to are trade names, trademarks, or registered trademarks of their respective owners. Specifications are subject to change without notice. © 2018 Artesyn Embedded Technologies, Inc. All rights reserved. For full legal terms and conditions, please visit www.artesyn.com/legal.



www.artesyn.com

PCIE8120-D2 29Jan2018