SOLUTION BRIEF

# Can We Go Any Faster? Deploy a Sea of GPGPUs for HPC.

### Design Advantages

- Cooling and power for up to seven (7) double-wide or 14 single-wide GPGPU cards in a single system
- Supports up to four 100G intelligent network adaptor cards for optimal I/O accessibility and chaining
- Flexible system configuration with sub-division and SR-IOV support allows one or more servers controlling one or multiple cards – depends on card capabilities
- Low latency high speed PCI Express interconnect with support for multiple root complexes and virtual function awareness

*Unparalleled density and configurability for High Performance Computing (HPC) in a single system.*

## Introduction

High performance computing and neural networks are making their way into more and more applications that require learning algorithms and the need to parallelize calculations to enable timely and correct responses to a given input. In many cases, this functionality is implemented using General Purpose Graphics Processing Units (GPGPU), sometimes even using commercial graphics cards.

Many servers offer a fairly limited number of slots and cooling capabilities, typically limited to up to two (2) double-width or up to four (4) single-width graphics cards. This can be overcome using external cabled PCI Express extension kits, which creates a mixed environment while requiring additional cables, connections and power.

Artesyn's ***MaxCore™ platform***, with its unique architecture, allows implementing the "exactly right" configuration for a given environment with a high number of PCI Express slots and the cooling and power capabilities to enable all this. High bandwidth inside the system, and the ability to support up to 400G of external connectivity shared among multiple servers using SR-IOV if required, provides adequate and configurable connectivity for most applications. Paired with Artesyn's ***System Services Framework (SSF)*** software to configure the system using a GUI environment, the ***MaxCore platform*** presents an easy-to-use, flexible system for many HPC applications.

**MaxCore™ Platform**

INTEL® NETWORK BUILDERS PARTNER

PCI EXPRESS®

ARTESYN™ EMBEDDED TECHNOLOGIES

## How it's Done

The Artesyn MaxCore™ family of platforms offers a powerful standards-based combination of different technologies. Using PCI Express slots as its basis, the MaxCore platform supports a wealth of open-market I/O and computing modules that support most requirements for external connectivity, as well as functions such as GPGPUs.

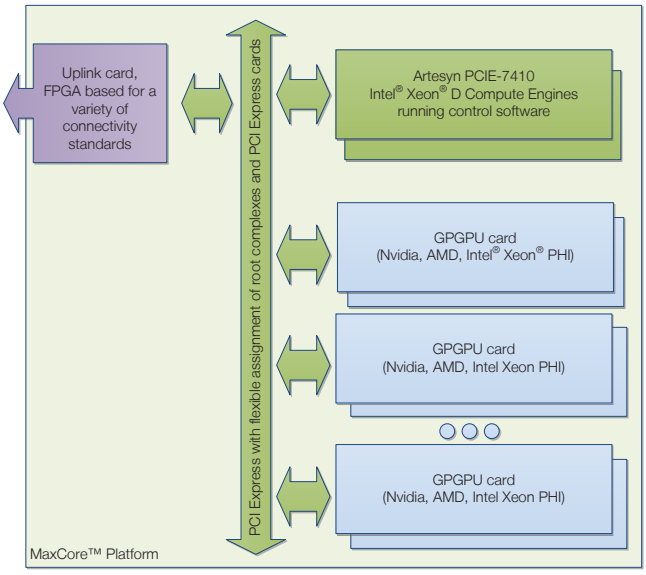


### *PCI Express Switch*

At the core of the MaxCore platform is a unique configurable PCI Express switch that supports multiple modes of operation, all of which are configured through the web GUI of Artesyn's ***System Services Framework (SSF)*** system management software.

When configured in basic mode, this switch functions as a generic PCI Express switch, allowing the use of all remaining 14 PCI Express slots for I/O and GPU cards. In this mode, the software running on the server will simply note a high number of PCI Express devices and will be able to support up to seven (7) dual-slot, 14 single-slot PCI Express cards, or a mix of these. The system is capable to cool and supply power for up to 150W per slot – so a double-width PCI Express card could support up to 300W of power!

When in Express mode, SSF allows slots to be configured as Root nodes (dividing the one PCI Express tree into multiple trees with private nodes) to run multiple servers inside a single 3U chassis. Additionally, the PCI Express switch understands virtual functions, so devices that support SR-IOV can be shared among multiple servers, for instance, enabling them to share a single 100G network interface controller (NIC) among themselves or even the virtual machines they are running, removing the need for running virtual switches on every CPU.

### *Server*

Artesyn's SharpServer™ PCIE-7410 is a server card based on the Intel® Xeon® D SoC with up to 16 cores. Every server card features two CPU complexes, so the system may provide up to 30 server CPUs in a single enclosure, depending on connectivity needs.

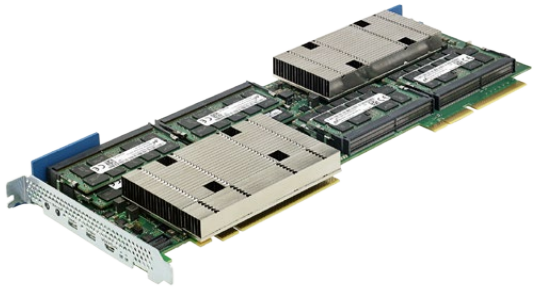

### *Networking*

Artesyn's ***SharpSwitch™ PCIE-9205*** intelligent network adaptor card features a high-bandwidth Intel® FM10840 network interface SoC with switch characteristics, which allows offloading many functions that would normally be executed on a CPU onto the interface silicon. The card allows redundant uplinks and connects into the system at 100Gbps. Additionally, for further offload functions or other server functions, the card features an Intel® Xeon® D processor complex similar to the one on the SharpServer PCIE-7410. The system is designed to host up to four (4) of these cards in certain slots.

ARTESYN™
EMBEDDED TECHNOLOGIES

## GPGPUs

The system is prepared and ready to host and cool multiple Nvidia Tesla, AMD or Intel® Xeon® PHI cards. With up to 150W per slot of power and enough cooling and connectivity capacity, the MaxCore family of platforms is ready to support massive applications that can also be daisy-chained using the network adaptors.

Nvidia Tesla Card

## Get Started!

If you are interested in using and deploying, or just testing, our MaxCore solutions, please contact the nearest Artesyn sales office or request a demo online — look for the "Request a Demo" button on the *MaxCore platform web pages*. We will be happy to discuss your application in more detail. Contact us to get started today!

REQUEST A DEMO

## System Configuration

Rounding off the technical capabilities, Artesyn's ***System Services Framework software*** offers a web GUI-based, easy-to-use user interface into the system configuration and status. It allows storage of multiple configurations, control and status review of all sensors available in the system. Agents running on the Artesyn boards enable even deeper insights into status and loads of every single component in the system.

As a separate package, Artesyn also offers support for OpenStack software and container architectures either in a pre-configured manner or through partner software packages.

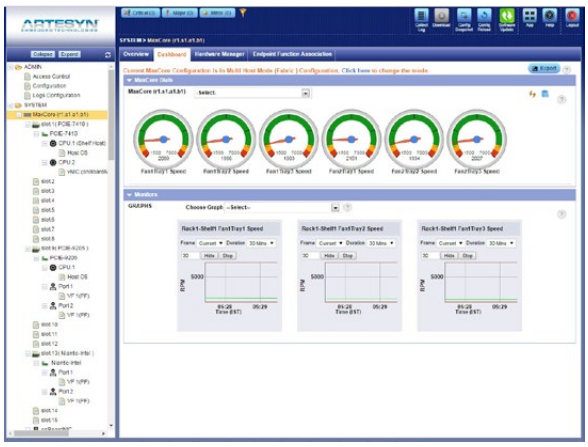

**www.artesyn.com**

**+1 888 412 7832 or +1 602 438 5720**



HighPerformanceComputing-SolutionBrief-15Jul2016

ARTESYN™
EMBEDDED TECHNOLOGIES