ADLINK
LEADING EDGE COMPUTING

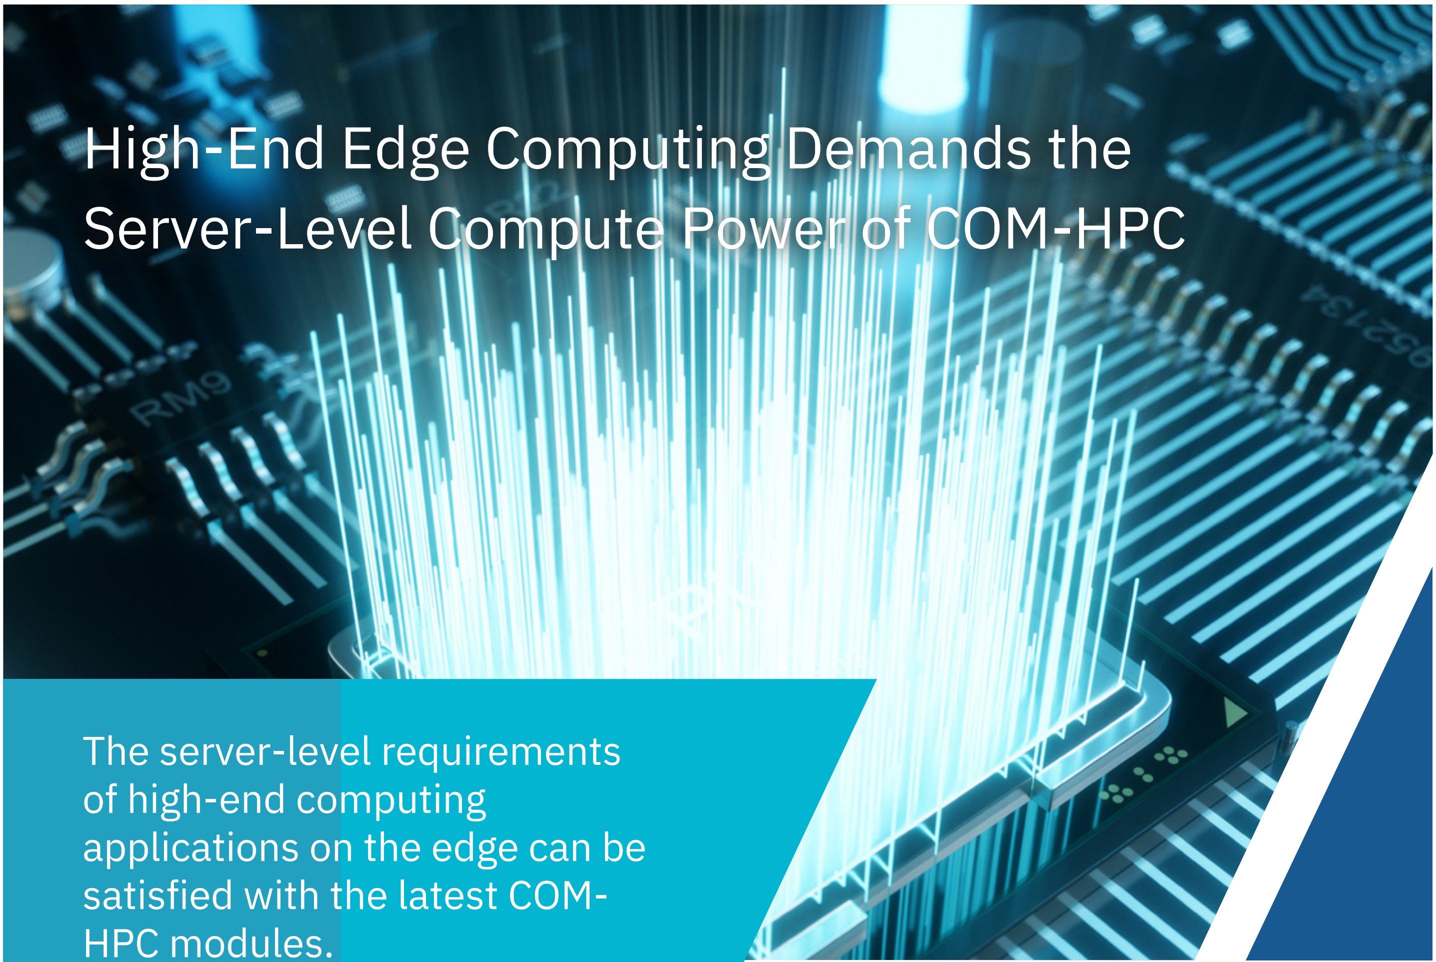

High-End Edge Computing Demands the Server-Level Compute Power of COM-HPC

The server-level requirements of high-end computing applications on the edge can be satisfied with the latest COM-HPC modules.

The Fourth Industrial Revolution -- Industry 4.0 -- refers to the ongoing automation of traditional manufacturing and industrial practices using modern smart technologies. When coupled with the evolution of the Industrial Internet of Things (IIoT), the rapid growth of Industry 4.0 has resulted in an explosion of data emanating from sensors and actuators.

More recently, the computational demands of Industry 4.0 and the IIoT have started to undergo even steeper growth due to additional technology drivers, such as artificial intelligence (AI), machine learning (ML), and 5G applications.

The term "edge" refers to the boundary where the internet meets the real world. Embedded systems on the edge are currently experiencing exponential growth regarding their requirements for computational power and performance.

In order to minimize both latency and communication costs, this data needs to be processed at the source to provide actionable information in close to real-time. The term "edge computing" refers to the act of performing this processing at the edge.

Above and beyond their desire for increased computational capability, system designers and end users also wish to "future-proof" their systems. There are many aspects to this, such as having access to alternative vendors should an existing supplier exit the market. Also, it's necessary to have the ability to upgrade systems with enhanced capabilities to satisfy the requirements of future applications.

High-Performance Computing (HPC) COMs

Some COMs are proprietary, while others are based on industry standards. In the case of system designers and end users who wish to "future-proof" their systems, the best option is to use an industry standard that is supported by many COM vendors.

Founded in 1994, PICMG is the preeminent COM standards body. This is a nonprofit consortium of companies and organizations that collaboratively develop open standards for high performance telecommunications, industrial, test-and-measurement, medical, edge, and general-purpose embedded computing applications.

The design of COMs involves a wide range of technical disciplines, including very high-speed signaling design and analysis, power management, networking expertise, thermal management, mechanical design, high-availability software, and comprehensive system management. All of these disciplines are represented by PICMG, which boasts over 140 member companies.

HPC stands for "High-Performance Computing." PICMG is now entering the final stages in the development of its new COM-HPC computer-on-module specification, which will address the needs of the high end of the industrial edge computing market.

The exponentially growing edge computing market segment demands ever higher performance and ever greater quantities of system memory. The embedded processors used in high-end edge systems can have 20 cores or more, and the combination of multiple processing cores and virtualization means that some systems require 256 gigabytes (GB) of memory, or more. Another consideration is that many applications demand extremely high data bandwidths, which may range as high as 100 gigabits per second (Gbps). The COM-HPC specification addresses all of these high-end requirements.

Markets and Applications for COM-HPC

COM-HPC modules address the needs of a wide range of markets and application areas that require various combinations of high-performance and/or high memory and/or high-speed interfaces. These include, but are not limited to, autonomous vehicles, robotics, unmanned aerial vehicles, and 5G communications (small cells, base stations, and infrastructure).

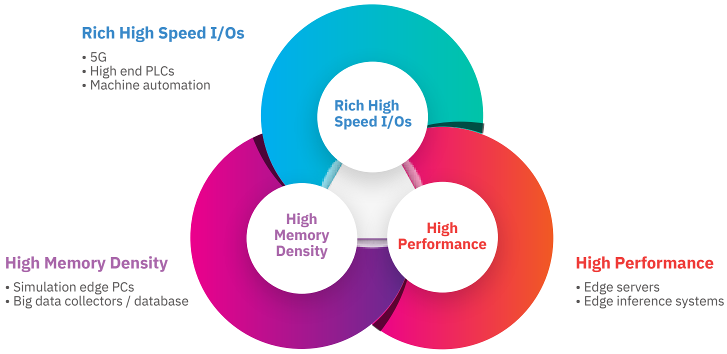

Figure 1: COM-HPC modules address a wide range of markets and applications

There are a number of big drivers for COM-HPC, not the least of which is task consolidation. Consider today's trains, for example, which may involve multiple IT-based services handling such tasks as reservation management, on-board Wi-Fi, and video surveillance. These services are typically provided by different companies in different "boxes" that are scattered throughout the train. Consolidating all of these tasks into a single COM-HPC "box" -- with each task running on its own virtual machine on its own processing core with access to tens of gigabytes of memory -- will dramatically reduce deployment and maintenance costs and resources.

Another major driver is artificial intelligence. It won't be long before every robotic device, including stationary robots, mobile robots, and collaborative robots, or cobots -- that is, robots that interact collaboratively and in close proximity with their human colleagues -- will be equipped with sophisticated AI capabilities. Similarly, it won't be long before everything that can drive, fly, or float is going to have an AI co-pilot or be fully autonomous. Even hard-to-come-by skills that have historically been dominated by humans (surgery, for example) will soon be augmented by AI.

Yet another driver is rugged and/or special purpose communications devices that are oftentimes deployed in harsh environments. Examples include 5G base stations, packet inspection systems, and front running trading systems. These are all applications that cannot use off-the-shelf solutions and for which COM-HPC modules are ideally suited.

Dimensions and High-Speed Connectors

The COM-HPC specification currently defines five different options in the form of two server modules and three client modules. The largest of the server modules offers eight on-board DIMM sockets supporting up to 512 GB of memory (these are the vertically orientated sockets in the images below).

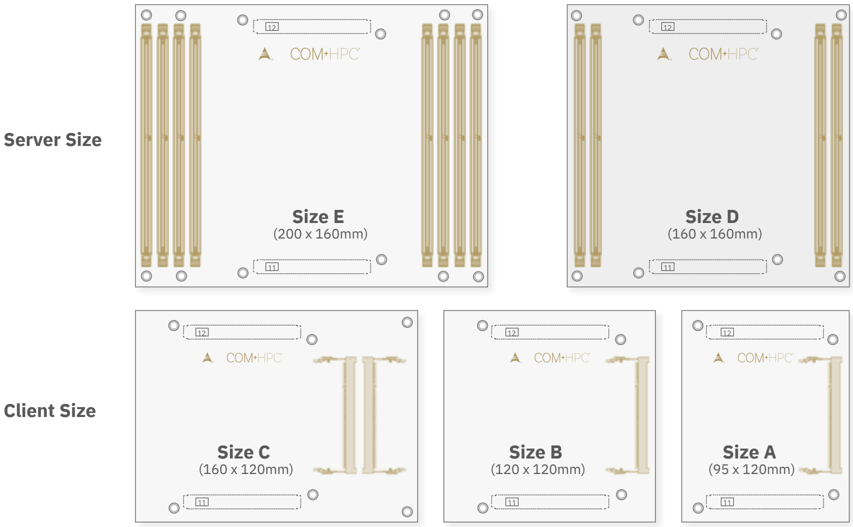

Figure 2: New COM-HPC module sizes

In addition to their DIMM sockets, each COM-HPC module is equipped with two new high-speed, high-density 400-pin connectors (the horizontally orientated connectors in the image above), resulting in 800 pins per module.

Figure 3: COM-HPC employs a new kind of connector.

Pinout Overview for COM-HPC Server

There are several points that stand out regarding the new high-speed, high-density COM-HPC connectors discussed above. For example, the power input is a single 12-volt supply that can deliver up to 358 watts of power, which is more than sufficient to drive high-end CPUs with 20 or more cores, 512 GB of memory, and any other components.

With regard to the PCIe (GEN5) interfaces, there are 64 lanes in total, with 1x16 = 16 lanes on the J1 connector and 3x16 = 48x lanes on the J2 connector. An additional 1x PCIe lane is provided on the J1 connector for board management control (more on this below). Furthermore, there are 8x USB 2.0, 2x USB 3.X, and 2x USB 4.0; along with 1x 1GbE and 8x 10GbE connectors.

The developers of many high-performance systems will also be interested to hear that COM-HPC server modules support the Intelligent Platform Management Bus (IPMB) and the Intelligent Platform Management Interface (IPMI). These are a set of computer interface specifications for an autonomous computer subsystem that provides management and monitoring capabilities independently of the host system's CPU, firmware, and operating system.

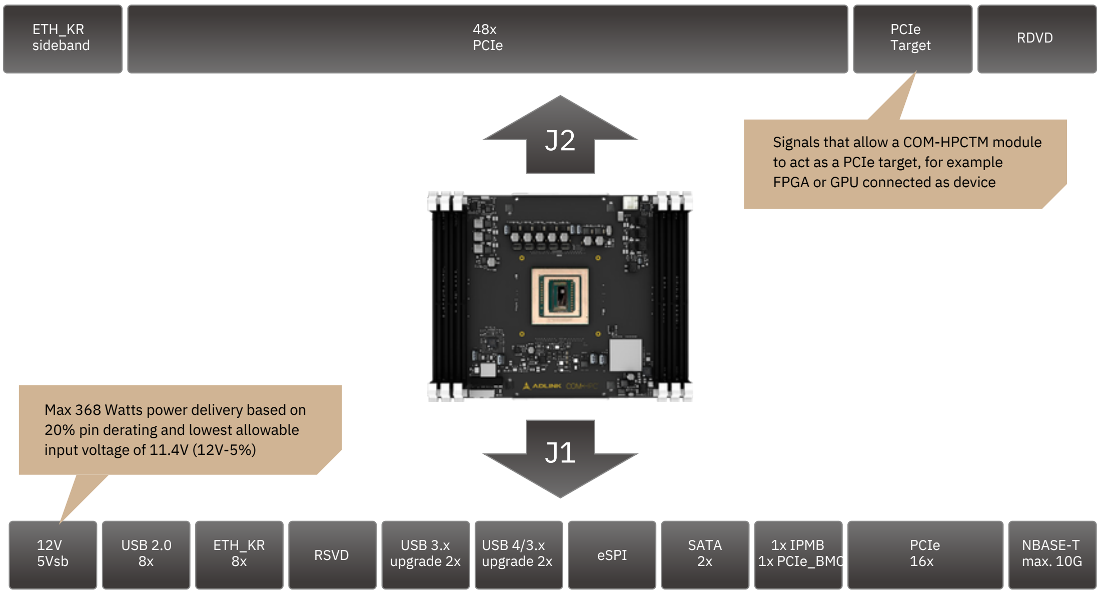

Figure 4: COM-HPC server pinout

The Price-Capacity Sweet Spot for Memory

Not surprisingly, when it comes to the memory-hungry markets for which COM-HPC is targeted, a major consideration for embedded systems architects and developers is the cost of the memory.

By employing DIMM technology, the COM-HPC server type not only enlarges the total possible memory capacity, but also facilitates the best price per GB. At the time of this writing, the sweet spot for price-capacity is 32 GB, which means users can enjoy 8x32 GB = 256 GB at the best price-capacity on the market.

In the case of applications that demand even more memory, developers can achieve this by employing 8x64 GB = 512 GB. Although this currently comes at a 5% price disadvantage, it is expected that the sweet spot for price-capacity will migrate to 64 GB in the not-so-distant future.

Sophisticated Remote Management Capabilities

Many of the early COM modules supported a limited form of management control by means of an application programming interface (API) called EAPI. However, this typically focused only on lower-level embedded management functions that are not abstracted through drivers, such as watchdogs, display brightness and backlight controls, logistics, EEPROM storage, and I2C/SMBus control.

EAPI will continue to be supported by both client and server COM-HPC modules. Additionally, the COM-HPC server type will also support full IPMI or its open source variant, Open BMC, because of its network centric nature.

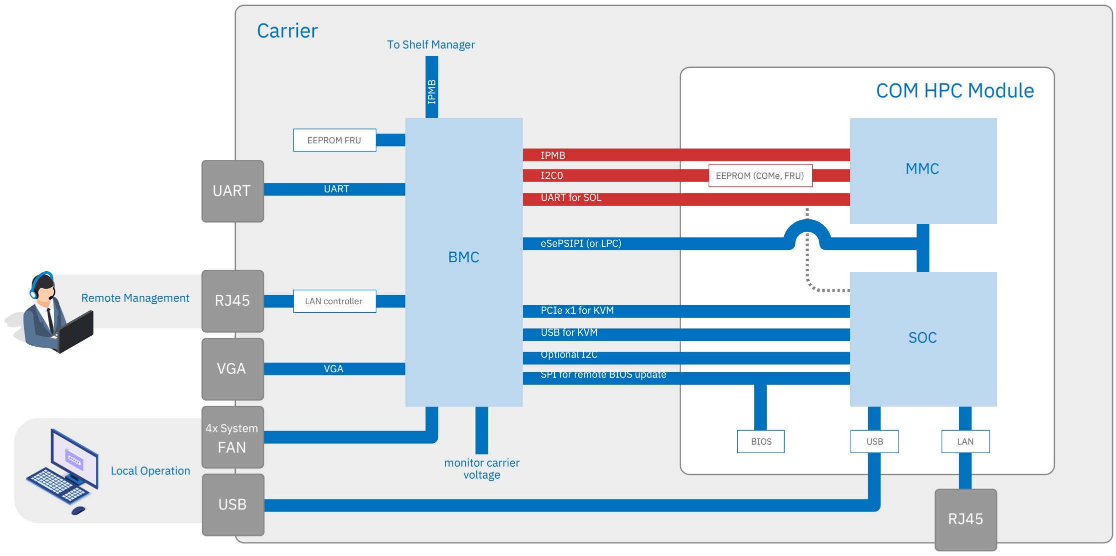

Figure 5: Typical out-of-band management control scenario

A typical implementation of IPMI in combination with a modular formfactor is shown in the diagram above. Although this illustration reflects only a single COM-HPC module, multiple modules can be controlled by a single BMC on the carrier.

The Future is Closer Than You Think!

ADLINK is a leading member of the PICMG consortium. As part of this, the company plays a large part in defining the COM-HPC standard and is at the forefront of COM-HPC development.

As shown below, ADLINK's proof-of-concept COM-HPC module is Size E (160mm x 200mm), has a server type feature set, and provides up to 16 computing cores in a 110W platform with multiple DIMMs.

Observe that, in order to extract heat to an external cooling device, a heat-spreader is mounted on top of the module. For companies wishing to develop their own COM-HPC modules and carrier boards, ADLINK can provide design verification or design services, as required. Also, ADLINK can provide software, firmware, and BIOS adaptation services to suit customers' specific requirements.

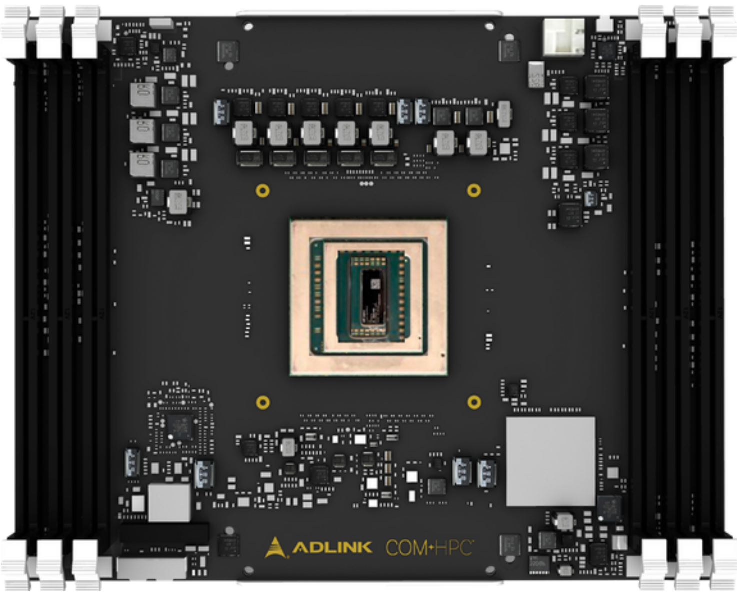

Figure 6: A proof-of-concept COM-HPC Server module

Furthermore, ADLINK's customers can take advantage of local and global networks of ADLINK support centers located in all major geographic regions.

Backed by top-class production and manufacturing logistics coupled with lifetime product stability, ADLINK's COM-HPC solutions provide long-term availability. The combination of extended temperature operation, mechanical ruggedness, and highly accelerated life testing (HALT) means that ADLINK's COM-HPC solutions provide high reliability. Also, the ability to provide hardware and firmware customization services regarding specific customer applications means that ADLINK's COM-HPC solutions provide a high degree of flexibility.

COM-HPC server modules will soon be appearing in myriad diverse locations and environments to perform a vast array of applications—all designed to make our lives better—and ADLINK will continue to be at the forefront of this exciting technology.