Understanding the COM-HPC standard for modular system designs

The embedded computing industry is about to launch COM-HPC as the next-generation standard for modular system designs. Since COM-HPC is complex and sometimes misunderstood, there is a need for clear information.

Some see the PCI Industrial Computer Manufacturers Group (PICMG) COM-HPC standard as an entirely new platform addressing entirely new applications. The embedded edge server camp, which must manage massive workloads in harsh environments, thinks so. The second camp are the existing COM Express users. They are less interested in the server modules, and instead care more about the client modules of the new COM-HPC standard. They are a little more skeptical about COM-HPC. They want to protect existing COM Express investments, asking themselves: How long will COM Express be available and do I have to switch to COM-HPC now? What are the advantages for my customers? For them, it is most important to know what benefits the COM-HPC client modules offer and how they differ from COM Express. COM-HPC therefore addresses two separate target groups, each with different needs. So what potential do the two new sub-specifications have and how do they differ?

Open standard platform for embedded servers

A short history of high-performance Computer-on-Module standards	
Year	**Details**
2001	Foundation of the ETX-IG and launch of the first manufacturer-independent module standard
2005	PICMG publishes COM Express 1.0 specification
2010	COM Express 2.0 specification

2012	Sales of COM Express modules exceed those of ETX
2012	COM Express 2.1 specification
2018	Foundation of the PICMG COM-HPC Committee
2019	Release of the COM-HPC pinout
2020	Launch of the COM-HPC specification

Open and manufacturer-independent Computer-on-Module standards ensure that applications have a life cycle of several decades. OEMs can still purchase new ETX modules today, despite the fact that this form factor is based purely on legacy buses. Thanks to backward compatibility, standards based on PCIe will be around even longer.

COM-HPC Server is the first truly open standard for developing modular embedded rack server and box server designs for harsh environments. In today's classic server world, application-ready processor modules are still rarely leveraged, although this approach offers many advantages. For example, it makes it very easy to realize specific size and I/O requirements: Developers only need to design the appropriate application-specific carrier board; the complex core components such as processor, RAM and high-speed interfaces can be purchased in a standardized module.

Because carrier board design requires less effort than full-custom design, this approach can also be efficiently applied to smaller product series, where standard products were previously often an unsatisfactory but unavoidable compromise. What is more, the modular concept also reduces the cost of performance upgrades significantly. Compared to the full replacement of a 1U or 3U rack system, a modular server design can cut the cost of an upgrade by around 50% because only the module is replaced. This approach therefore improves the sustainability of the investment as well as the long-term availability and ROI of the solutions, as they can be used for longer.

More compute units in one standard

Next to the generic advantages of a modular concept, COM-HPC Server also offers some technical refinements that were previously not available in modules in this form. For instance, the COM-HPC standard is not limited to x86 processors but explicitly provides for the use of RISC processors, FPGAs and general-purpose

graphics processing units (GPGPUs). First samples with such alternative compute units were shown at the PICMG booth during Embedded World.

So, for the first time it has become possible to develop and implement heterogeneous server designs with a wide range of compute and accelerator units inside a single official specification and standardized ecosystem. To facilitate this, the new specification also supports slave modes for the modules for the first time. OEMs benefit not only from simplified and more efficient design-in but can also re-use their know-how more effectively.

More space for more performance

COM-HPC Server modules aim to provide edge and fog server applications with the high-performance computing power required by the new embedded edge server processors that semiconductor manufacturers are anticipated to launch very soon. The maximum specified power budget of 300 watts for COM-HPC Server modules gives an indication of the performance to be expected, at least in the medium term. For comparison: The most powerful COM Express Type 7 Server-on-Module today allows a maximum of 100 watts. Scale this, taking the predicted performance leap into account, and it is easy to see that COM-HPC will be able to cover immense server loads in the future.

Next to this high-performance potential, it is important how much space the modules provide for the processors or alternative compute units. This becomes clear when looking at the current high-performance CPUs from Intel and AMD with 16 or more cores or the powerful FPGAs, which can be palm-sized. For them, COM-HPC Server offers module footprints of 200 mm x 160 mm (Size E), or 160 mm x 160 mm (Size D) as shown in Figure 1. These comparatively large footprints also simplify heat dissipation, providing space for larger heat sinks that can distribute heat waste more efficiently.

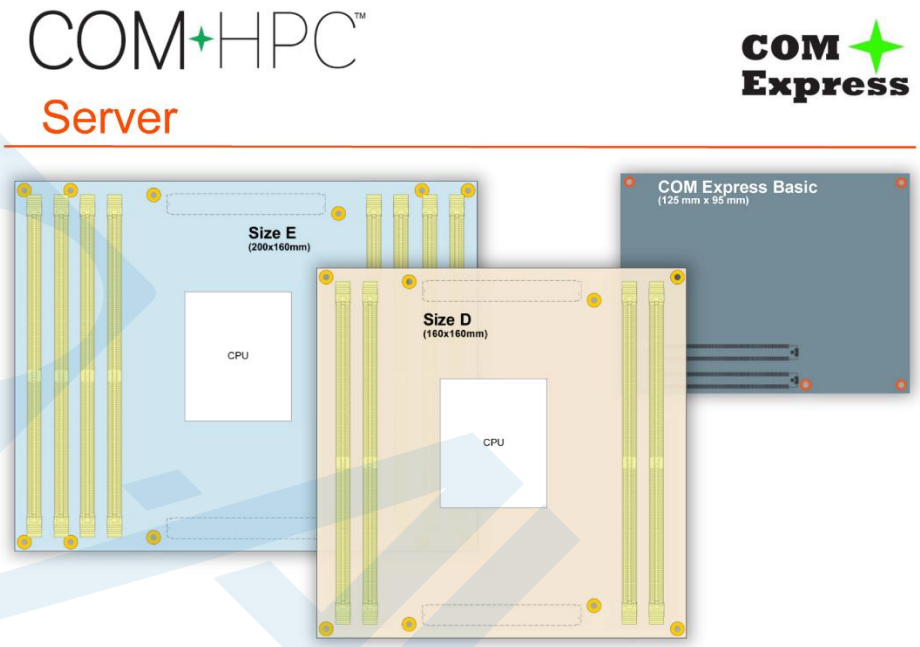

Figure 1: COM-HPC Server specifies two different footprints: Size E with space for up to 8 DIMM sockets for currently 1 terabyte of RAM, and the 20% smaller Size D footprint for 4 DIMM sockets. While COM-HPC Server and Client use the same connectors with 2x 400 pins, they are positioned at different distances from each other. This prevents damage from accidentally mounting the wrong type of module (Source: congatec).

More memory performance

With these large footprints, the COM-HPC Server modules also provide more memory performance. They have enough space for full DIMM memory modules that meet the high memory bandwidth and size requirements of micro, edge and fog servers. In the Size E footprint, they can host up to 8 DIMM sockets for currently up to 1.0 terabyte of memory. In the Size D footprint, they can host a maximum of 4 DIMM sockets for currently up to 512 gigabytes of memory.

More I/O performance

For carrier board connection, COM-HPC Server defines 8x 25 GbE, as well as 65 PCIe lanes for PCI Express Gen 4.0 and Gen 5.0 (Figure 2). One of these lanes is reserved for communication with an optional board management controller (BMC) on the carrier board, while the remaining 64 PCIe lanes can be used to connect peripherals.

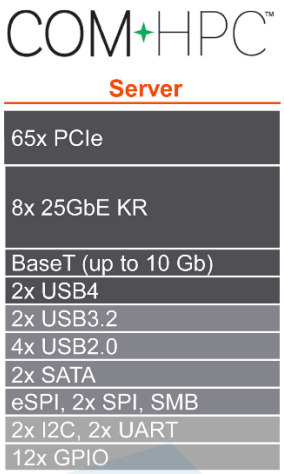

Figure 2: The key features of COM-HPC Server modules: Extraordinarily large number of high-speed interfaces, outstanding network bandwidth and headless server performance (Source: congatec).

COM-HPC Server therefore offers extremely broad and powerful connectivity, for example to connect additional computing accelerators such as GPGPUs, FPGAS and ASICS – for instance in the form of matching COM-HPC modules – or NVMe-based storage media. In total, COM-HPC Server designs benefit from an I/O performance of up to 256 Gigabyte/s via PCIe. A further 2x 40 Gigabit/s can be added via the two USB 4.0 interfaces on Thunderbolt 3.0 versions, as well as 2x 20 Gigabit/s via the two specified USB 3.2 interfaces. Four additional USB 2.0 interfaces round off the USB offer on COM-HPC Server modules. Besides 2x native SATA, they also support eSPI, 2xSPI, SMB, 2x I2C, 2xUART and 12 GPIOs to integrate simple peripherals and standard communication interfaces, for instance for service purposes. An additional 10 Gb Ethernet port provides a dedicated communication channel that can be used for remote and out-of-band management.

Optimized server-grade board management

Another industry first to be introduced by COM-HPC is a dedicated system management interface. This interface is currently being developed in the PICMG Remote Management Subcommittee. The aim is to make parts of the feature set specified in the intelligent platform management interface (IPMI) available for remote edge server module management. Similar to the slave function, COM-HPC will therefore also provide extended communication functions for remote management. Thanks to this feature, OEMs and users will be able to ensure reliability, availability, maintainability, and safety (RAMS), a common set of

requirements for servers. For individual needs, this function can be expanded via the optional board management controller on the carrier board. This provides OEMs with a uniform basis for remote management that can be adapted to specific requirements.

COM-HPC Client – bigger, faster, more

While the COM-HPC Server specification concentrates on completely new embedded edge server designs, there are of course also the “classic” high-performance embedded systems, which have been leveraging COM Express Type 6 up to now. OEMs are wondering whether COM-HPC will make their existing COM Express designs obsolete, when would be the best time to switch to COM-HPC, and what advantages COM-HPC has for them and their customers. To answer these questions, it is important to know in detail which functions COM-HPC Client modules offer and to compare them with the COM Express functions.

Three sizes

In many respects, the two standards have more similarities than differences. Like COM Express, COM-HPC Client specifies three module sizes: 120 mm x 160 mm (Size C), 120 mm x 120 mm (Size B), and 120 mm x 95 mm (Size A). This means the smallest COM-HPC Client footprint is virtually identical with COM Express Basic, which measures 125 mm x 95 mm. This alone shows that COM-HPC Client sits above COM Express, addressing applications that cannot be reached with COM Express. (Figure 3).

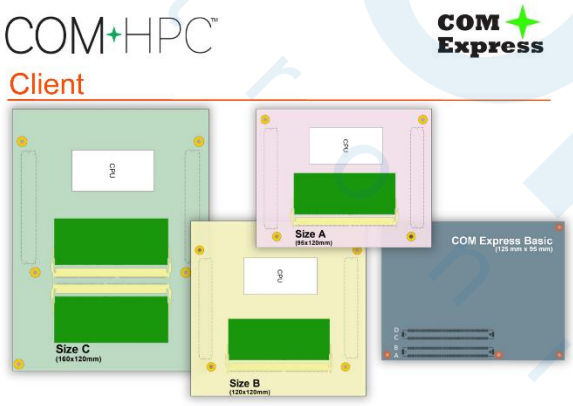

Figure 3: COM Express and COM-HPC Client both define three different footprints. However, with the smallest COM-HPC Size A almost identical with COM Express Basic, it is immediately evident that COM-HPC is positioned above COM Express (Source: congatec).

More power

This is also reflected by the supported power budget of 200 watts, which is about three times as much as today's most powerful COM Express Type 6 modules. As far as memory is concerned, while COM-HPC Client and COM Express both use SODIMMS or soldered memory, COM-HPC can accommodate more memory with up to 4 SODIMM sockets. However, with COM Express already capable of supporting 96 GBytes today, it meets high memory requirements as well.

More and faster interfaces

From a layout point of view, the most important difference between COM Express and COM-HPC modules concerns the connector, and the number of signal pins connecting the module to the application-specific carrier board as Figure 4 shows.

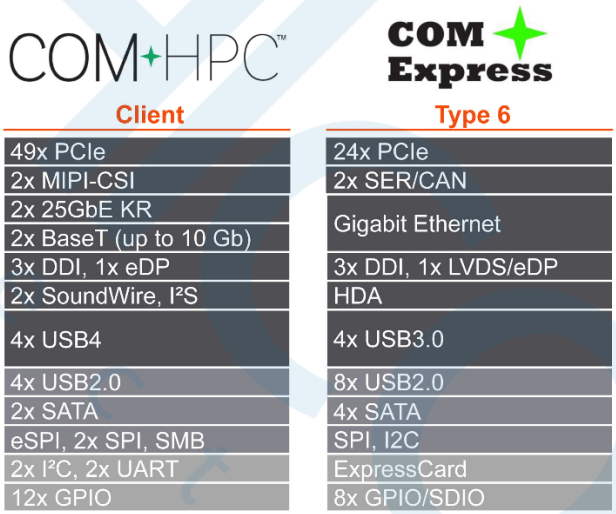

Figure 4: COM-HPC Client and COM Express Type 6 interfaces differ mainly in the number of PCIe lanes and the bandwidth, the Ethernet interfaces and USB ports, and the extended remote management support still to be specified (Source: congatec).

COM-HPC leverages a new connector that's designed for the latest high-speed interfaces and already specified for the high clock rates of PCIe 5.0 and 25 Gb/s. COM Express supports PCIe Gen 3.0 and PCIe 4.0 in compatibility mode. But of course, embedded processors with PCIe Gen 4.0 must first become available. Like COM Express, COM-HPC supports two connectors, but with 400 pins each. So with 800 signal pins in total, COM-HPC has almost twice as many pins as COM Express Type 6 modules with 440 pins. Needless to say, this also provides space for many more interfaces.

COM-HPC Client modules use these for 49 PCIe lanes to the carrier board, one of which is again intended for communication with the carrier board's BMC. That's twice as many lanes as COM Express Type 6 provides with a maximum of 24 lanes. Two 25 GbE KR Ethernet and up to two 10 Gb BaseT interfaces are also provided directly on the module. COM Express Type 6 supports 1x 1 GbE, with the option to implement further network interfaces via the carrier board.

4x graphics

The graphics support is identical in both standards; at the same time, it is what distinguishes these modules from the headless COM Express Server-on-Modules and COM-HPC Server modules. Both standards support up to four displays via three digital display interfaces (DDI) and 1x embedded DisplayPort (eDP). In terms of multimedia interfaces, COM-HPC uses SoundWire in place of the HDA interface specified for COM Express. SoundWire is a new MIPI standard that requires only two lanes: clock and data, with a clock rate of up to 12.288 MHz. Up to 4 audio codecs can be connected in parallel over these two lanes, with each codec receiving its own ID for analysis.

More USB bandwidth plus MIPI-CSI

Geared for the future also in terms of the supported USB standards, COM-HPC specifies four USB 4.0 interfaces, supplemented by 4x USB 2.0. While this means that COM-HPC Client modules offer four USB ports less than COM Express Type 6 modules, which execute up to 4x USB 3.1 and 8x USB 2.0, this is compensated by more bandwidth, since USB 4.0 is designed for transfer rates of up to 40 Gbit/s. Another attractive feature of COM-HPC Client modules is that they provide two MIPI-CSI interfaces, enabling cost-effective camera connections for situational awareness and collaborative robotics.

COM-HPC Client further offers 2x SATA interfaces for connecting traditional SSDs and HDDs, which are almost legacy devices today, plus industrial interfaces such as 2x UART and 12x GPIO. 2x I2C, SPI and eSPI complete the feature set. All these features are comparable to COM Express Type 6 modules, which in contrast to COM-HPC offer optional CAN bus via the connector as a unique feature.

Judging by the differences, OEMs with COM Express based designs can rest assured that they will be well served with COM Express for many years to come. This is also because COM-HPC does not introduce a new system bus – unlike the switches from ISA to PCI and from PCI to PCI Express. It's also worth remembering that COM Express modules did not replace ETX as the best-selling modules until 2012 – a

good 11 years after the introduction of ETX. And ETX modules are still sold today. Since PCIe generations are backward compatible with their predecessors, designs with PCIe Gen 3.0 will remain in service for a long time, even after PCIe Gen 4.0 has been introduced across all processor levels. So, as long as the given interface specification is adequate, there is absolutely no need to change.

However, if you need more than 32 PCIe lanes, or you require PCIe 4.0 in full bandwidth, USB 4.0 multiple 25 Gbit/s Ethernet and/or advanced remote management features, it is worth making the switch. Otherwise stick to the motto ‘never change a running system’.

By ***Christian Eder*** *is director of marketing for EMEA at congatec and chairman of the PICMG COM-HPC Subcommittee. He holds a degree in Electrical Engineering, University of Applied Sciences Regensburg, Germany. Eder has over 30 years of experience in embedded computing and is one of the founders of congatec.*