congatec

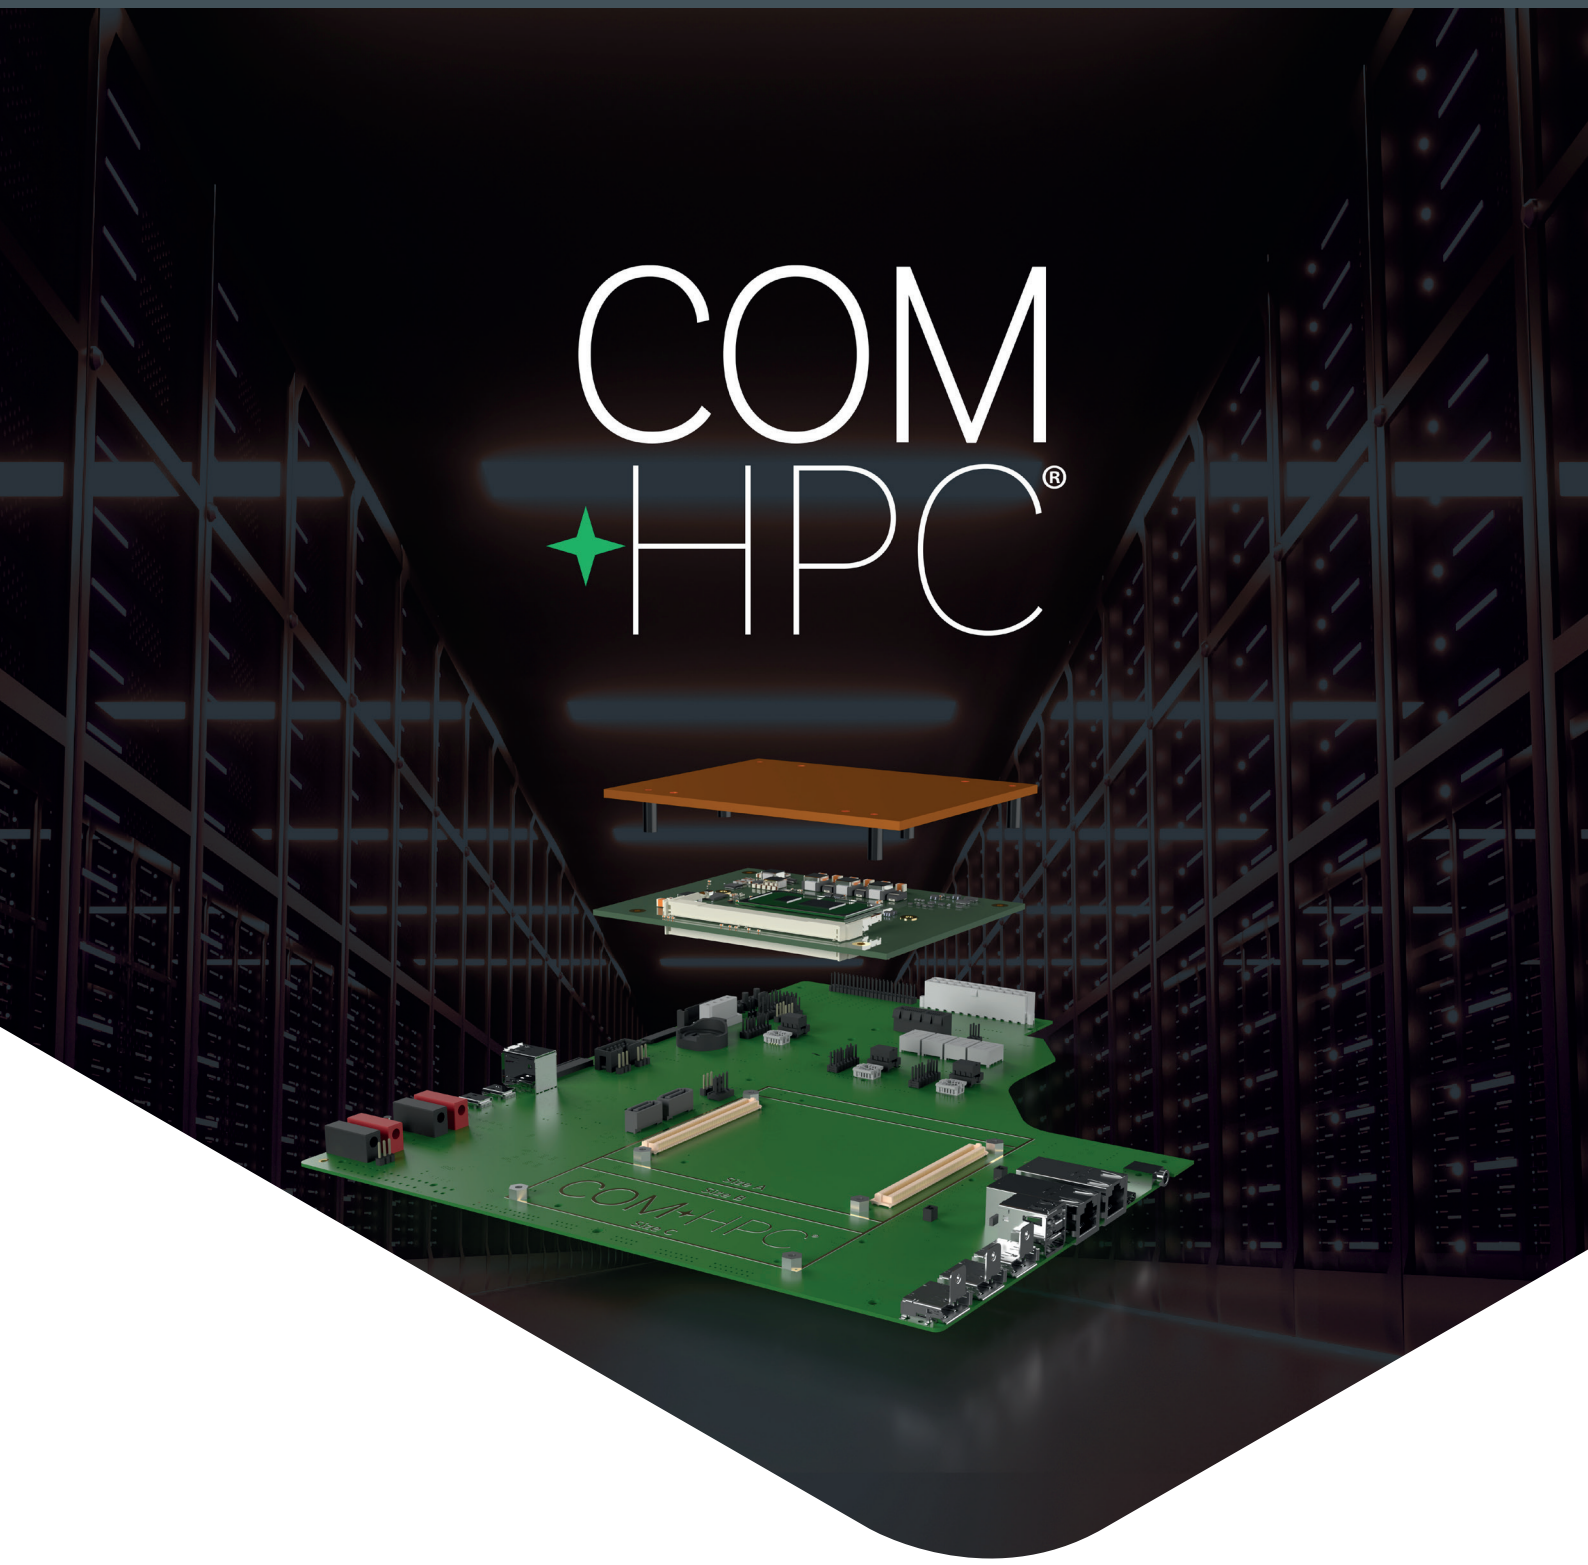

Whitepaper

A new standard for modular edge servers

COM HPC Server – the new Server-on-Module standard

A new standard for modular edge servers

The PICMG will soon officially adopt the COM HPC Computer-on-Module standard. Part of the new standard is the COM HPC Server specification. It was developed for high-end edge server computing, which eagerly awaits new high-speed interfaces such as PCI Express 4.0 and 5.0 as well as 25 Gbit Ethernet. Compared to established COM Express Type 7 based Server-on-Modules, the new specification has almost twice the interface bandwidth with its 800 signal lines, and therefore provides a significantly greater range of functions besides the performance boost. What are the most important markets and applications for the new high-end Server-on-Modules and what are the limits?

COM HPC Server addresses the entire range of edge servers as the new Server-on-Modules have a maximum power consumption of up to 300 watts and can therefore use the powerful Intel Xeon and AMD EPYC server processors. However, first modules are not expected to become available until the launch of the next Intel Embedded server processors.

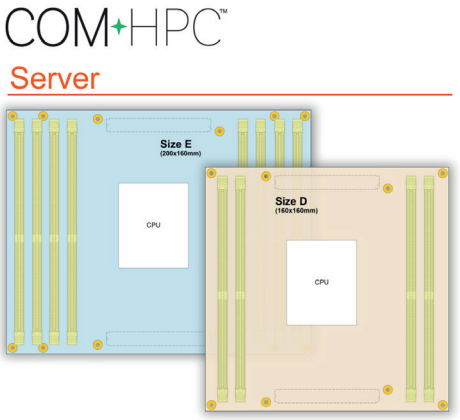

There are two different COM HPC Server form factors that can host either 4 or 8 SO-DIMM sockets.

COM HPC modules can currently support up to 1 Terabyte RAM on up to 8 SO-DIMM sockets. Ultimately, that's the limiting factor. The memory hosted on the modules cannot be made directly available to other parallel processors via Intel Ultra Path Interconnect or AMD Infinity Fabric Inter-Socket. Instead, it is exclusively available to each individual COM HPC module. Dual-die/single-socket Cascade Lake processor technology with up to 12 socket RAM is also not feasible with COM HPC.

Master/slave function for artificial intelligence

The roadmap nevertheless includes multi-module installations on carrier boards since COM HPC allows master/slave configurations. However, this feature is not so much designed for dual-processor carrier boards, where two or more COM HPC Server modules could be operated in master/slave mode. Rather, it aims to enable communication between COM HPC Server modules and COM HPC FPGA and/or COM HPC GPGPU modules – a third application area of the COM HPC standard besides COM HPC Server and COM HPC Client modules. The master/slave function, available for the first time with Computer-on-Modules, was primarily implemented to make flat GPGPU or FPGA modules available in a form factor that is significantly more powerful than MXM. Hence the COM HPC standard also stands for the integration of functions such as artificial intelligence, deep learning and digital signal processing (DSP), which is important for telecommunications. With COM HPC, all these functions can be made available in a robust, space-saving and scalable manner. Since COM HPC Server-on-Modules are generally not implemented in dual-processor boards, the standard positions itself in the mid-performance segment of server technology, which is required in many diverse configurations at the edge.

5G infrastructure servers and connected edge systems

Much of the demand for robust server-class performance comes from edge applications that reside directly in and at the network and communication infrastructures. Examples are outdoor small cells on 5G radio masts for tactile broadband Internet, or carrier-grade edge server designs with 20 inch installation depth, which are currently found in 3G/4G edge server rooms where the installation depth of the racks is just 600 mm. In both cases, the systems are often enriched with GPU or FPGA boards for future 5G designs, for example to enable real-time scheduling of data transmission in the microsecond range , which again underlines the need for the master/slave function in COM HPC. COM HPC Server modules are, of course, also attractive for the related 5G measurement technology.

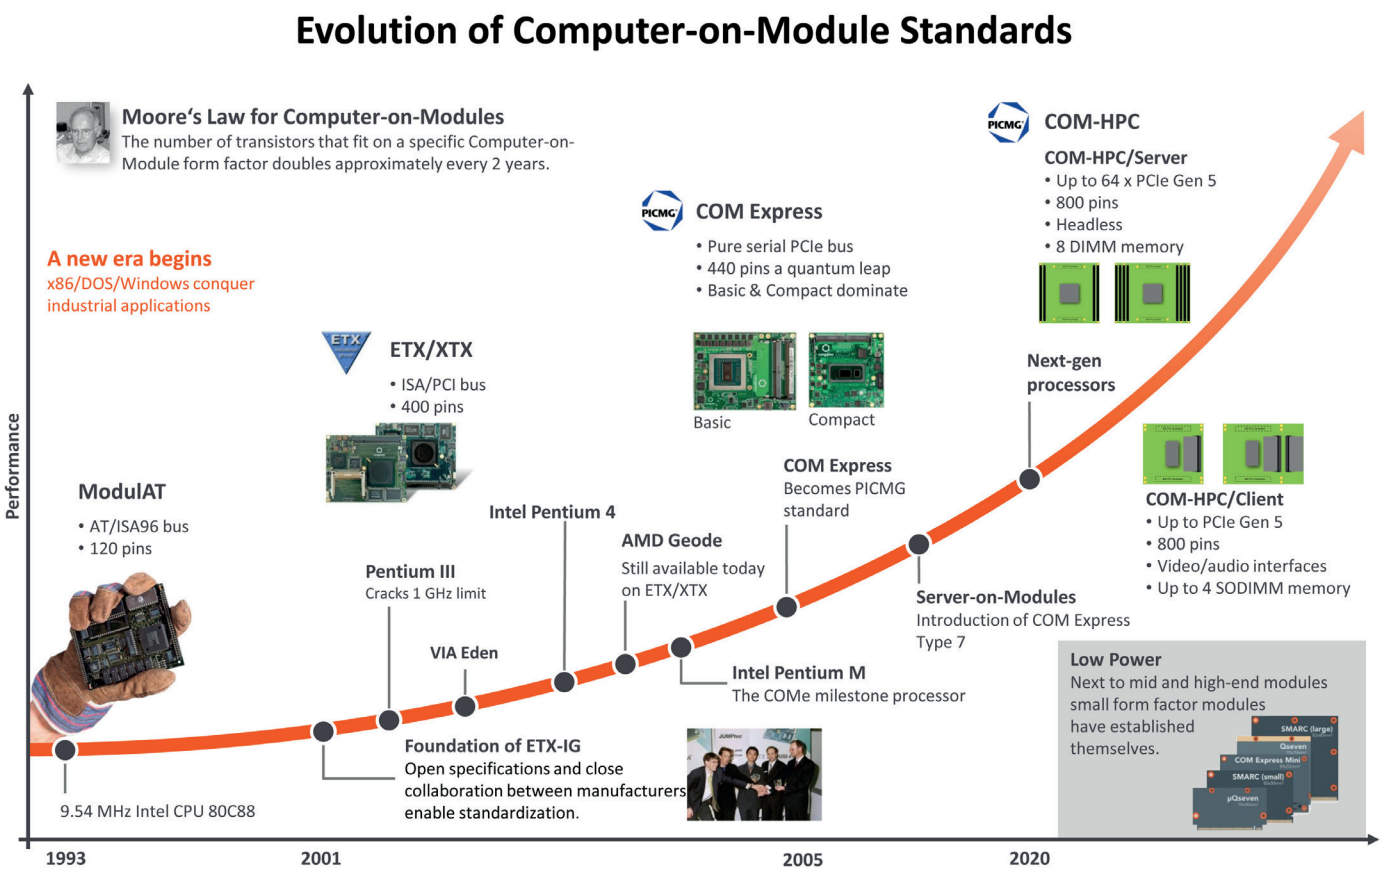

COM-HPC is a consistent step in the progression of the Computer-on-Module market. However, it will likely be years before COM-HPC reaches market shares similar to COM Express, since COM Express also needed about 5 years to outstrip ETX in terms of quantities. And with ETX modules still being sold today, existing COM Express customers can also expect to be able to purchase COM Express modules for many years to come.

OCCERA server with full NFV and SDN support

Another application area for COM HPC Server modules are robust edge server designs that are based on the Open Compute Carrier-grade Edge Reference Architecture (OCCERA) and built in rack depths of 1000 or 1200 mm. Due to increasing virtualization and support for NFV and SDN, dedicated infrastructure equipment is being implemented on such open standard platforms in growing numbers. This includes firewalls, intrusion detection systems, traffic shaping, content filters, deep packet inspection, unified threat management, and anti-virus applications. Here too, users benefit from on-demand scalability and cost-effective performance upgrades for the next generation of servers.

From Industry 4.0 to workload consolidation

Ultimately, edge equipment also includes all the fog computers close to the applications, i.e. redundant clouds at the edge. In the harsh industrial environment, they are required for Industry 4.0 communication and/or to implement new business models and comprehensive predictive maintenance using big data analytics at the edge. Another application area is opening up from the possibility to consolidate the computing performance on a single edge server instead of distributed computers – for example, to control a manufacturing cell with several robots.

Energy grids, autonomous vehicles and rail transport

Similar edge servers are also needed in smart energy networks to reconcile energy generation and consumption in microgrids and virtual energy networks in real-time. In such distributed applications real-time 5G will play a role as well, since the required high bandwidths along the infrastructure paths can be provisioned comparatively easily. Autonomous vehicles will benefit from 5G, too. They also need COM HPC Server performance in the cockpit as they have to process high bandwidths onboard to analyze data from hundreds of sensors in fractions of a second. Further applications can be found in transportation, namely in the rail sector to provide a full range of services – both onboard the trains and within the infrastructure ¬– from security technology to streaming services for infotainment. Mobile broadcasting systems and streaming servers as well as equipment for non-public mobile terrestrial radio networks for use by security authorities and organizations (BOS Digital Radio Network) are further application areas for COM HPC Server modules.

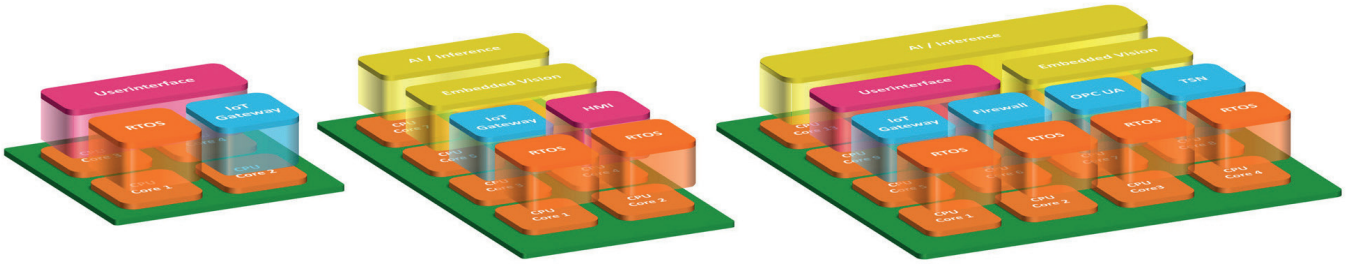

COM HPC Server designs enable load balancing at the edge not only through virtualization but also thanks to the ability to scale cores on demand.

Smart vision systems

A considerable number of COM HPC Server applications can also be found in medical backend systems, industrial inspection systems, and video surveillance solutions at the edge for increased public safety, where – as with autonomous vehicles – massive video and imaging data have to be processed in parallel. However, a significant share of these applications will not rely on COM HPC GPGPU, but be looking to use even more powerful FPGA extensions or PEG graphics.

Modular motherboard designs

In addition to the often dedicated system designs for the application areas described so far, COM HPC modules will also be deployed on robust standard motherboards, which are long-term available and ideal for the design of smaller servers. Used mostly in rackmount, workstation or tower designs, these boards range from the µ ATX format (244 × 244 mm) to the extended ATX form factor (305 x 244 mm). EEB mainboards (305 x 330 mm) will also be able to use COM HPC. As some of the latter feature two processors, they will only be able to conquer part of this

form factor market. However, they will bring previously unimaginable scalability to this market, coupled with the option to customize the I/O interfaces extremely cost-effectively and efficiently. This is an enormous advantage for edge applications as the industrial communication landscape is extremely heterogeneous in the direction of the process and field level. So in summary, COM HPC Server modules are suitable for an extremely wide range of edge and fog server applications in harsh environments, whereby the focus tends to be clearly on the processor.

Why use COM HPC Server modules?

The COM HPC Server specification offers edge servers the full advantages of standardized Server-on-Modules, which we are familiar with from the Computer-on-Module market, i.e. the performance can be tailored to the specific needs of the application. In a rack with identical server slots it is therefore possible to adapt the physical performance through the design of each individual module, thereby optimizing the price/performance ratio for the application. This is a particular advantage for real-time applications, where it's not an option to simply scale the load balancing regardless of the real physical resources. Rather, it must be possible to allocate real resources in a dedicated way. When performance upgrades are implemented on the basis of an existing system design, the costs for moving to the next generation can be reduced to about 50% of the initial investment because, in the majority of cases, it is enough to simply replace the processor module. This significantly reduces the TCO and accelerates the ROI. Manufacturer-independent standardization also ensures second-source strategies, competitive prices and high long-term availability. Finally, Server-on-Modules are also more sustainable, ultimately contributing to a lower CO2 footprint of server technology in production, which – despite all the energy consumed in continuous operation – significantly lowers the environmental impact.

Advanced server management interfaces

Another advantage of the new specification is the availability of a dedicated remote management interface. This interface is currently being defined by the PICMG Remote Management Subcommittee. The aim is to make part of the function 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 remote management functions to communicate with the modules. This feature set guarantees OEMs and users the same reliability, availability, maintainability, and security (RAMS) as is common in servers. For individual needs, this function can be expanded via the board management controller to be implemented on the carrier board. This provides OEMs with a uniform basis for remote management that can be adapted to specific requirements.

Unbeatable custom interface design and performance upgrades

All in all, COM HPC Server offers many advantages. Whenever edge servers need a very specific mix of interfaces that must also fit on a small footprint, Server-on-Modules are simply unbeatable. In such cases a carrier board design can be implemented much faster and more cost-effectively than a full-custom design. Up to 80% of the NRE costs, sometimes even more, can be saved this way. What is more, COM HPC Server modules are also extremely likely to find widespread use in standard industrial servers as the cost advantages of performance scaling for the next generation are unmatched. The use of modules will further be strengthened by the fact that server performance will increasingly be provided as-a-service, so that the investment will be borne by the service provider instead of the user.

Development Concepts in Comparison

COM & Carrier Designs	Full Custom Designs
One supercomponent	Complex BOM
Open standard	Proprietary processor implementation
Application ready BSPs	Complex implementation of hardware-related software
Comprehensive design-in support	Limited support options
Large ecosystem	No community
Efficient re-use of existing building blocks	Wheel reinvented every time
Long term availability	Complex lifecycle management
High design security	Greater risk of design errors
Low development costs	High NRE costs
Short time-to-market	Long development cycles
High scalability	Each variant a new product
Easy upgrades	Always a new design
Ideal for small to medium series	More complex than COM & carrier fusion
COM & carrier fusion for large series	

An edge server can be provided with highly customized I/Os much more cost-efficiently when using Server-on-Modules since individual carrier boards incur significantly lower NRE costs.

COM HPC: One standard, three application fields

The PICMG COM HPC standard specifies overall three variants: First, the so-called COM HPC Server modules, which are positioned above the COM Express Type 7 headless specification, and second the COM HPC Client modules, which are traded as successors of the COM Express Type 6 specification. A third application field for the specification are FPGA/GPGPU modules, for which suitable master/slave functions have been reserved in the standard. While the COM HPC Server and FPGA/GPGPU modules are mainly perceived as innovations, COM HPC Client modules give rise to concerns about the disruptive effect on the popular COM Express Type 6. Users of these modules are therefore a little more skeptical

about the COM HPC standard. They want to protect existing COM Express investments, and ask themselves: How long will COM Express be offered and do I have to switch to COM HPC now? What are the advantages for my customers? They primarily want to know what the COM HPC client modules have to offer and how they differ from COM Express.

When referring to Computer-on-Modules instead of Client modules, users quickly grasp the new concept. They are familiar with the advantages of these products since the ETX Computer-on-Module standard was published at the beginning of the millennium. And it is precisely this expertise that provides the answer to almost all the questions: A leading COM standard has always been the guarantee for massive long-term availability and this also applies to all existing COM Express designs. COM Express Computer-on-Modules will never disappear quickly as thousands of OEMs use them in their industrial applications, relying on the long term availability promises of embedded processor and embedded form factor standards. And rightly so! So you can safely stick with your COM Express designs. There is definitely no need to change as long as the given interface specification suits your requirements.

On the other side, moving from COM Express to COM HPC designs is definitely recommended for those developers who wish to benefit from these new features: PCIe Gen 4.0, more than 32 PCIe lanes for massive I/Os, 4 video interfaces, USB 4.0, 25 Gbit Ethernet, and advanced remote management. But unless you fall into this group, there is no reason for a rushed change. So, relax and continue to rely on the established open standards.

Author:

Christian Eder,

Director Marketing of congatec and Chairman of the PICMG COM HPC Subcommittee

About congatec

ccongatec is a rapidly growing technology company focusing on embedded computing products. The high-performance computer modules are used in a wide range of applications and devices in industrial automation, medical technology, transportation, telecommunications and many other verticals. Backed by controlling shareholder DBAG Fund VIII, a German midmarket fund focusing on growing industrial businesses, congatec has the financing and M&A experience to take advantage of these expanding market opportunities. congatec is the global market leader in the computer-on-modules segment with an excellent customer base from start-ups to international blue chip companies. Headquartered in Deggendorf, Germany, congatec currently has entities in USA, Taiwan, China, Japan and Australia as well as United Kingdom, France, and the Czech Republic. More information is available on our website at www.congatec.com or via LinkedIn, Twitter and YouTube.

Headquarters

congatec GmbH
Auwiesenstraße 5
94469 Deggendorf
Germany
Phone +49 (991) 2700-0
info@congatec.com
www.congatec.com

Subsidiaries

congatec Asia Ltd.
2F., No.186, Sec. 3,
Chengde Rd.
10366 Taipei, Taiwan
Phone +886 (2) 2597-8577
sales-asia@congatec.com
www.congatec.tw

congatec, Inc.
6262 Ferris Square
San Diego
CA 92121 USA
Phone +1 (858) 457-2600
sales-us@congatec.com
www.congatec.us

congatec Japan K.K.
Ingresso Shiodome building 301,
Minato-ku Hamamatsucho 1-2-7,
105-0013 Tokyo-to, Japan
Phone +81 (3) 6435-9250
sales-jp@congatec.com
www.congatec.jp

congatec China Technology Ltd.
Sunyoung Center, 901 Building B,
No. 28 Xuanhua Road, Changning District,
Shanghai 200050, China
Phone +86 (21) 6025-5862
sales-asia@congatec.com
www.congatec.cn

congatec Australia Pty Ltd.
Unit 2, 62 Township Drive
West Burleigh
Queensland 4219, Australia
Phone +61 (7) 5520-0841
sales-au@congatec.com
www.congatec.com

congatec Korea Ltd.
Mr. Yoonsun Kim
(Yatap-Dong), #707,
42, Jangmi-ro, Bundang-gu,
Seongnam-si, Gyeonggi-do,
South Korea
Phone: +82 (10) 2715-6418
yoonsun.kim@congatec.com

Real-Time Systems GmbH
Gartenstrasse 33
88212 Ravensburg
Germany
Phone +49 (751) 359558-0
info@real-time-systems.com
www.real-time-systems.com

Sales Offices

France
Mr. Luc Beugin

Phone: +33 6 44 32 70 88
cfr-sales@congatec.com

United Kingdom / Ireland
Mr. Darren Larter

Phone: +44 784 581 4792
cuk-sales@congatec.com

Nordics & Baltics
Mr. Anders Rasmussen

Phone: +45 285 649 92
cdk-sales@congatec.com