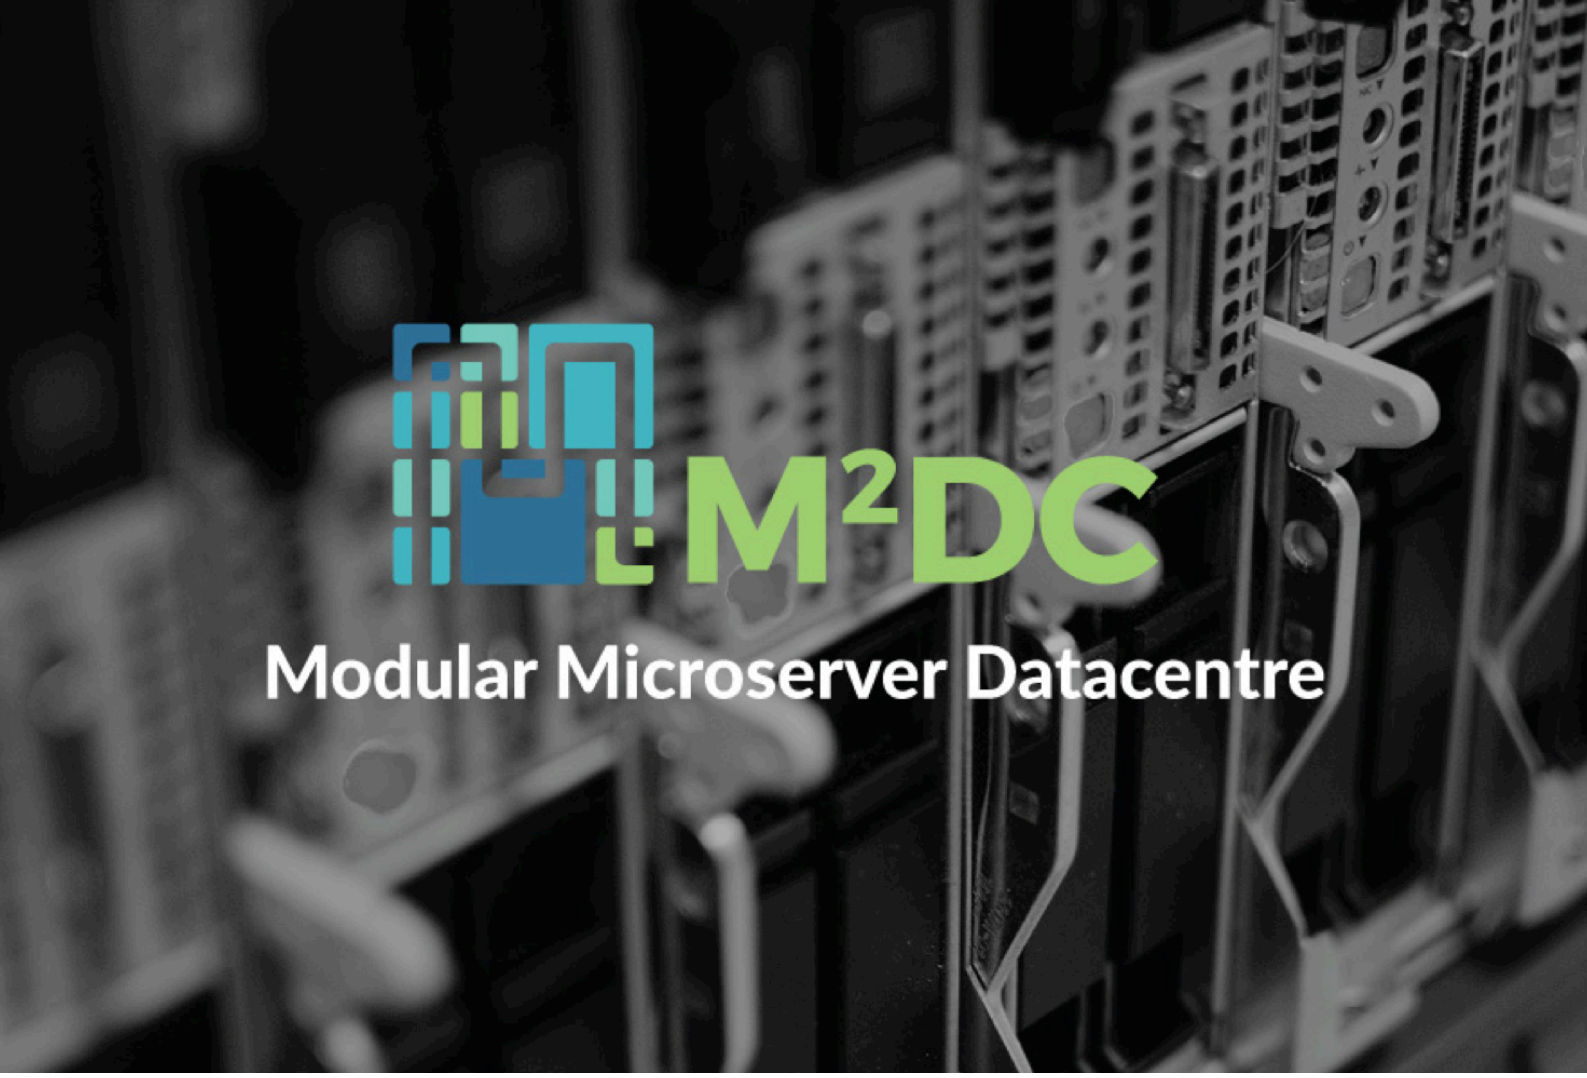

The new class of low-power appliances with built-in efficiency and dependability enhancements

The M2DC White Paper

Partners

Vodafone Automotive · PSNC · ARM · CHRISTMANN · CEA · HUAWEI · Universität Bielefeld

OFFIS · XLAB NOT IDLE · Politecnico Milano · cewe · beyond · REFLEX · EOLEN

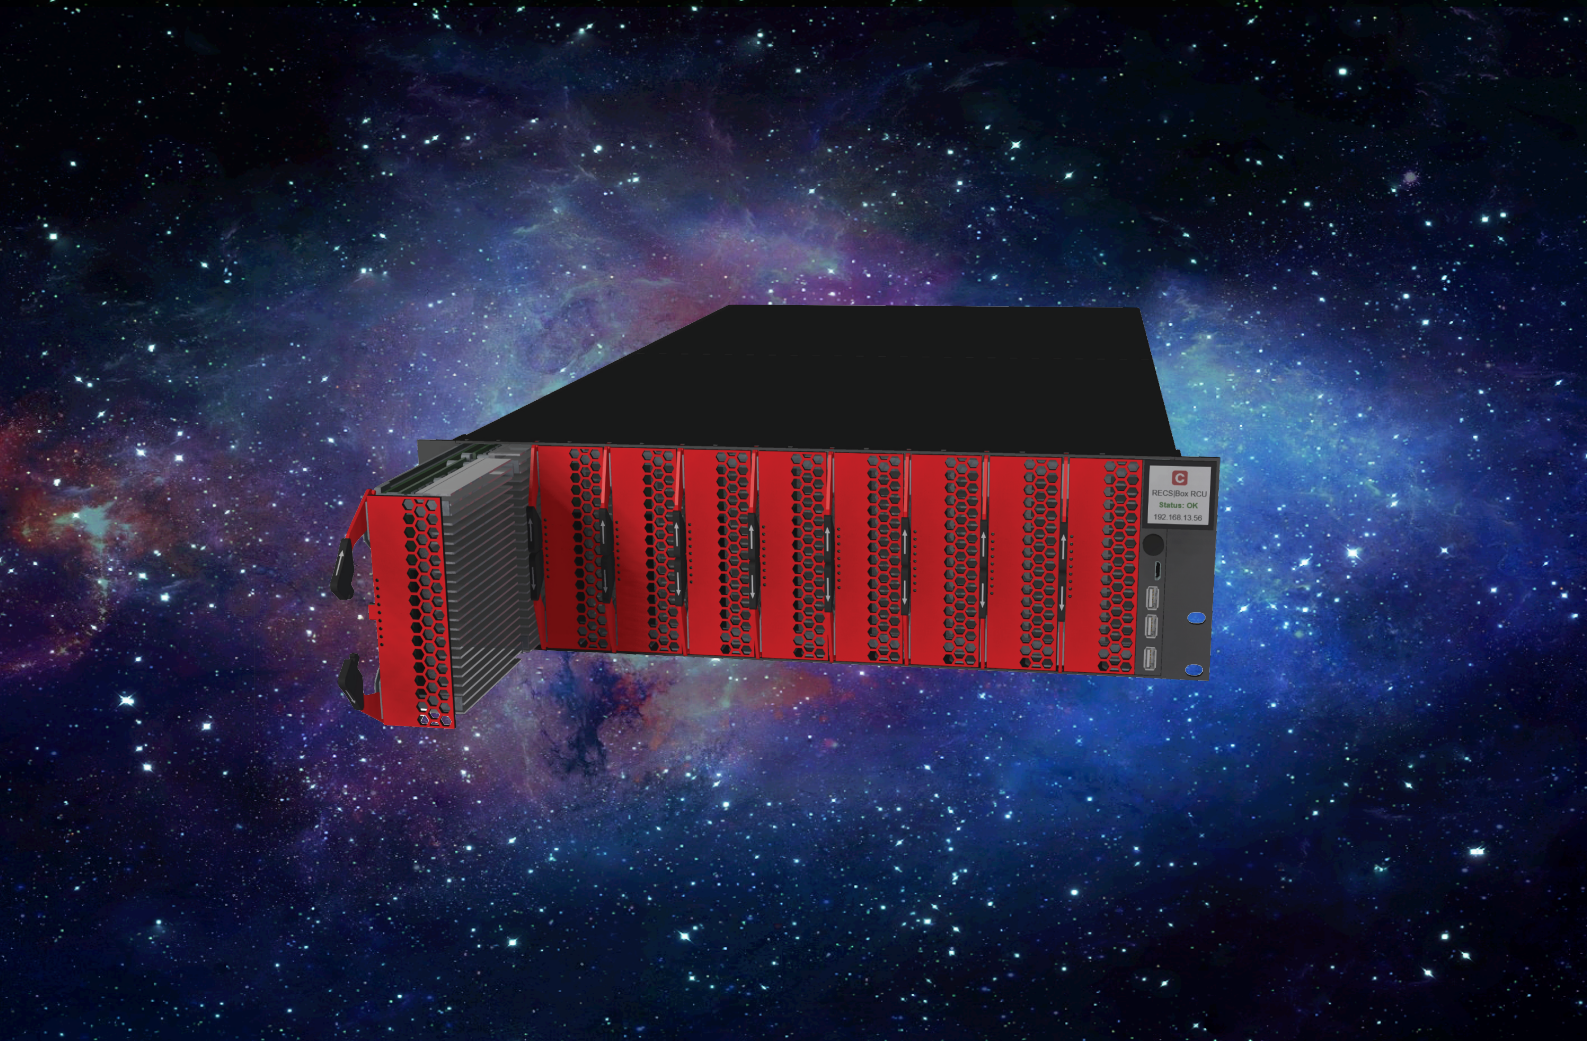

THE VISION

New technologies including advanced mobile devices, Internet of Things (IoT), 5G and machine learning, pose steadily increasing demands on the performance and energy efficiency of server platforms and data centres. **Heterogeneous hyperscale data centres** target these challenges with a combination of highly scalable server platforms and integrated hardware accelerators, e.g., based on GPUs and reconfigurable hardware. Against this background, 14 companies, research institutes and Universities are collaborating in the H2020 project M2DC (Modular Microserver Data Centre) on the development of **turnkey appliances, which can be easily configured, produced, installed and maintained**. The basis for these appliances is the M2DC server, a **modular, highly efficient, cost-optimised server architecture, seamlessly integrating heterogeneous microservers and hardware accelerators**. Built-in efficiency and dependability enhancements are combined with an intelligent power management for continuous optimisation of power and performance.

FUNDED BY
THE EUROPEAN UNION
No. 688201

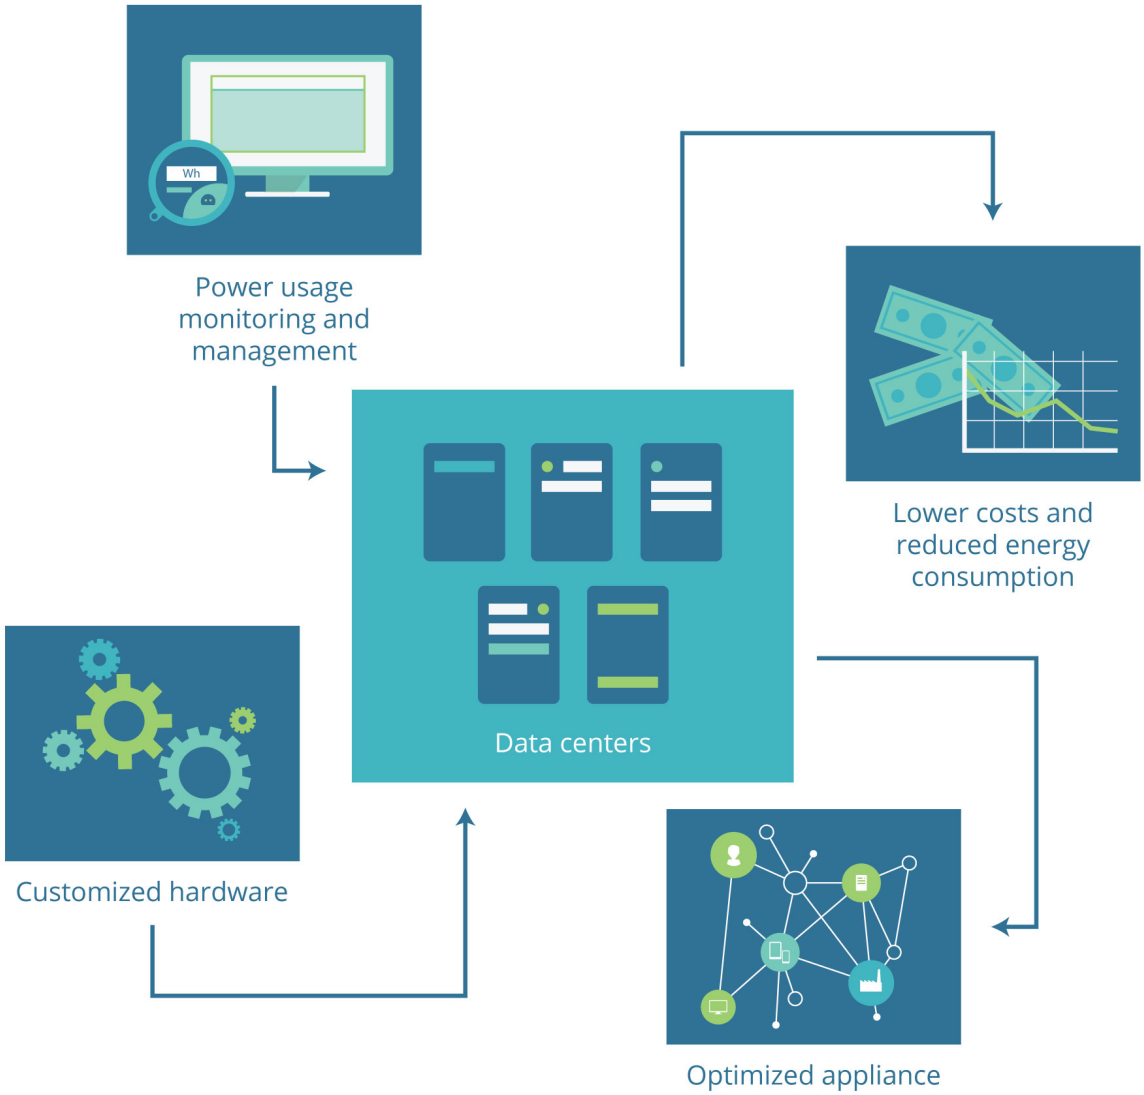

WIDE IMPACT

M2DC enables customization and smooth adaptation to various types of applications, while **advanced management strategies** and **system efficiency enhancements** (SEE) will be used to achieve high levels of energy efficiency, performance, security and reliability. The M2DC middleware provides a data centre capable abstraction of the underlying heterogeneity of the server. This **includes new microserver designs based on ARM64 and Intel Stratix 10**. Baseline benchmarks show the high potential of accelerators for the targeted applications including **photo finishing systems, IoT data processing, cloud computing, CNN and HPC**. In the next months we will finish the system design including an optimized middleware for deployment of the optimized appliances.

FUNDED BY THE EUROPEAN UNION No. 688201

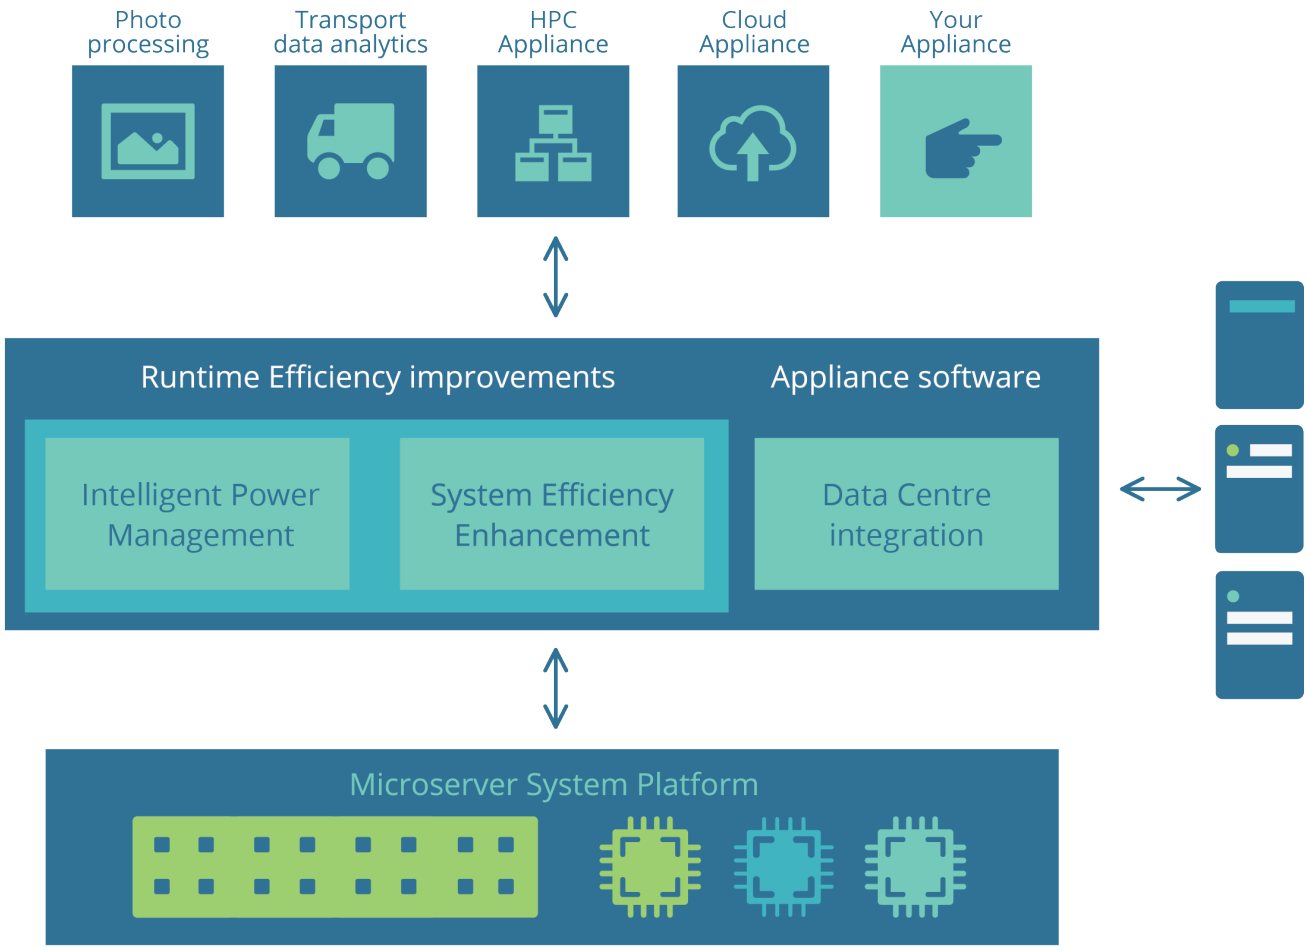

GLOBAL VIEW

M2DC targets the development of a **new class of energy-efficient appliances** with built-in efficiency and dependability enhancements. The appliances will be easy to integrate with a broad ecosystem of management software and fully software-defined to enable optimisation for a variety of future demanding applications in a cost-effective way. The highly flexible M2DC server platform will **enable customisation and smooth adaptation to various types of applications**, while advanced management strategies and system efficiency enhancements (SEE) will be used to improve **energy efficiency, performance, security and reliability**. Data center capable abstraction of the underlying heterogeneity of the server is provided by an OpenStack-based middleware.

FUNDED BY THE EUROPEAN UNION
No. 688201

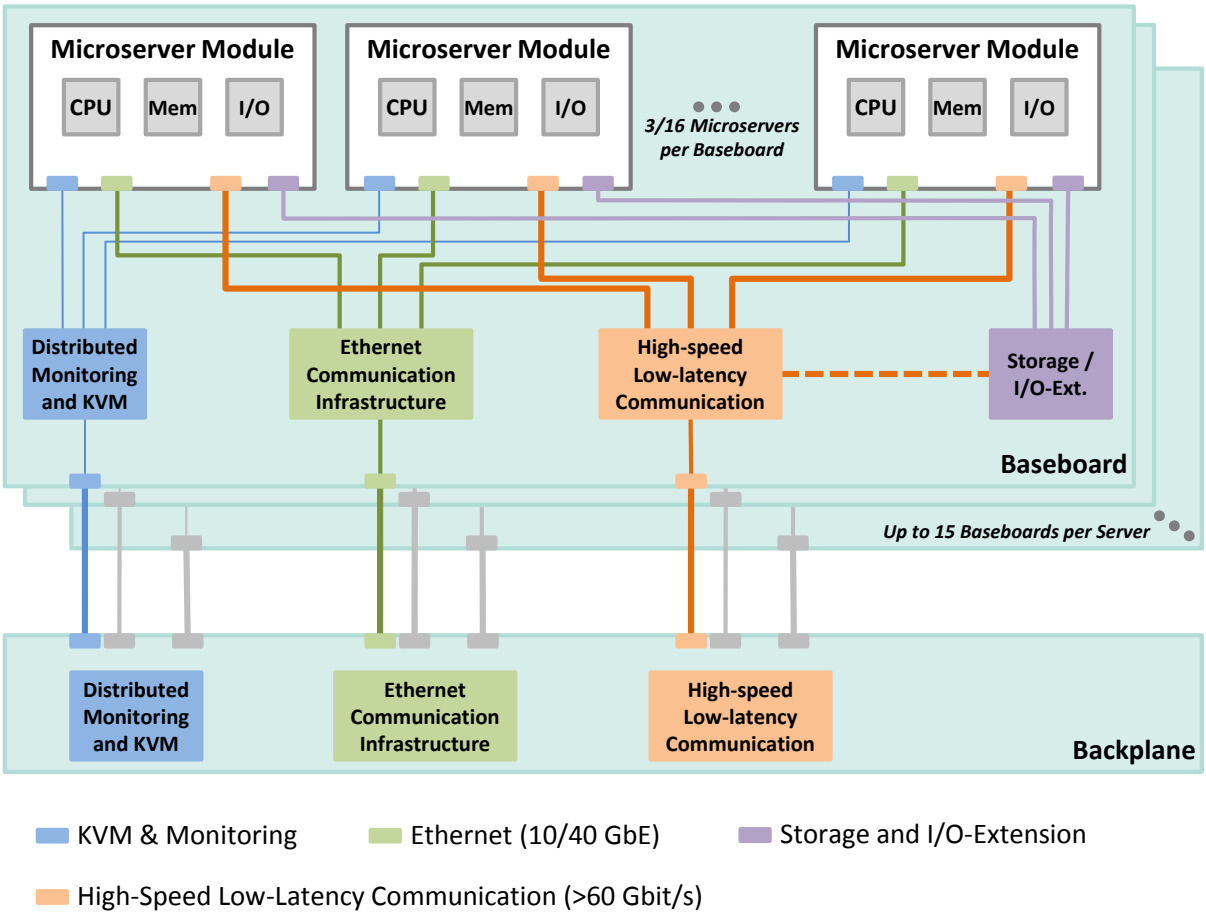

POWERFUL HARDWARE

The M2DC server is a resource-efficient, highly scalable heterogeneous microserver platform integrating: high-performance ARMv8 server processors; low-power ARMv8 embedded/mobile processors; x86 processors; and accelerators (MPSoCs, GPUs, and FPGAs). It's goals are: (i) **resource efficiency**; (ii) **high reliability**; and (iii) **high maintainability**. With a blade-style design, the platform integrates up to 240 low-power microservers or 45 high-performance microservers per chassis. The system has three levels of communication for flexibility and performance: monitoring and control; management and compute; and flexible high-bandwidth low-latency communication between the microservers. Microservers are based on established Computer on Module standards (COM Express, Nvidia Jetson, Toradex Apalis) with **novel microservers developed for M2DC** including the ARM64 Microserver (32 Core Cortex-A72 @ 2.1 GHz) and the Intel Stratix 10 Microserver (Quad-core ARM Cortex-A53 and high density FPGA fabric).

FUNDED BY THE EUROPEAN UNION No. 688201

	CPU Microserver	GPGPU	FPGA Microserver	PCIe-Extensions
Low Power Microserver	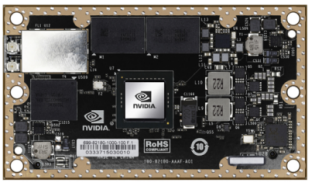	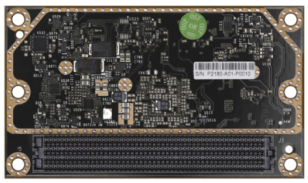	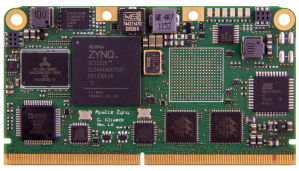	RAPTOR FPGA acc. 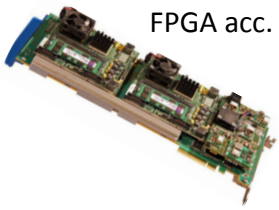
	NVIDIA Jetson TX1/TX2 4 Core A57@1.73 GHz + Maxwell GPGPU@1.5 GHz 2 Core Denver + 4 Core A57 + Pascal GPGPU@1.5 GHz		Xilinx Zynq 7020 85 kLC	Xeon Phi 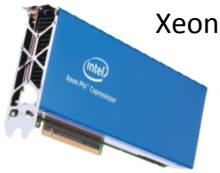
High Performance Microserver		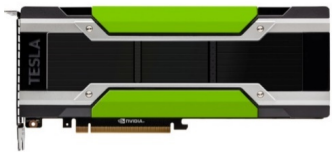	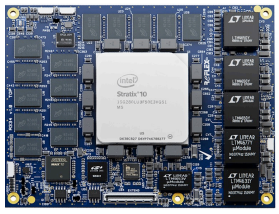	PCIe SSD 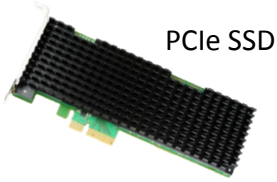
	ARMv8 Server SoC 32 Core A72@2.1 GHz	NVIDIA Tesla P100 Pascal GPGPU@1.3 GHz	Intel Stratix 10 SoC 2,800 kLC	

ENERGY EFFICIENCY

M2DC's energy-aware workload management combines the advantages of several referenced approaches which each focus only on limited aspects. By applying a combined proactive and reactive allocation algorithm, the M2DC workload management is able to **optimize aggressively while also providing emergency measures for sudden spikes in workload**. However, the most appealing advantage of M2DC's approach is the strict focus on its applicability: turn-key ready, easy-to use framework for appliances with **embedded fine-grained system monitoring with distributed preprocessing of sensor data** (e.g., power, voltage, temperature). While the usage of OpenStack as base platform should help spreading the solution due to its broad community and popularity, the consideration of alternative compute nodes in the modelling and management process guarantees a future relevance when GPUs and FPGAs become more popular in general-purpose computing. The M2DC middleware stack is based on OpenStack Ironic. Provides **bare metal (micro)server software deployment and lifecycle management**, with integration of extensions for handling the dynamic and heterogeneous nature of the microservers and hardware accelerators. Considers OpenStack extensions for dynamic node composition, combining CPU and accelerator nodes as required by the application. Includes a management layer for energy and thermal aware distribution of workload on server and rack level.

FUNDED BY THE EUROPEAN UNION No. 688201

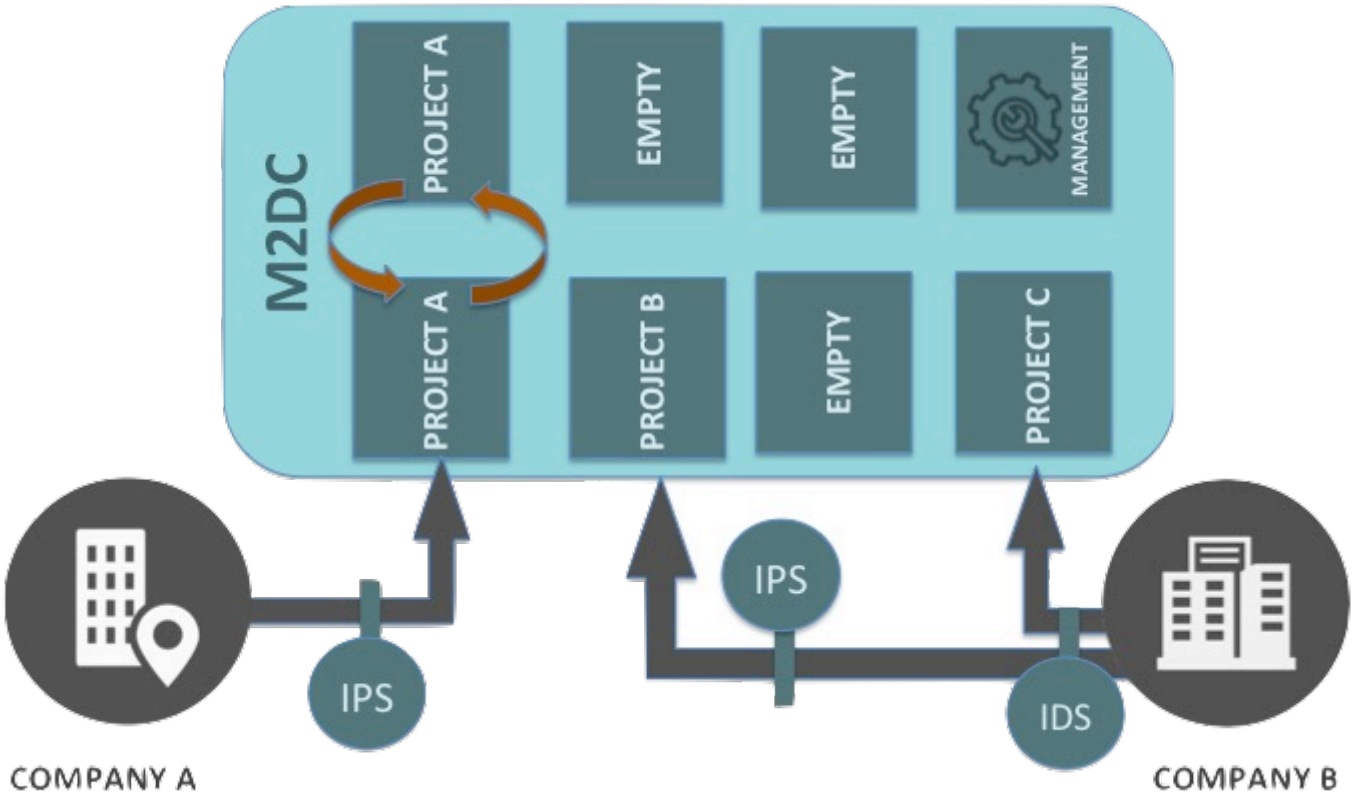

UNBREAKABLE SECURITY

Network security is of the utmost importance, everywhere. The knowledge of what is going on inside your private network is invaluable. Due to the fast development and introduction of new technologies, the pool of potentially vulnerable applications and protocols is quickly increasing. Firewalls alone can't provide network protection. The M2DC platform includes **Intrusion Detection and Prevention systems** (IDS/IPS) that can detect anomalies in internal networks, for incoming (before the fact) and outgoing traffic (after the fact). When IPS is enabled the alarms permit to actively block the unwelcome traffic. **At M2DC we take security very seriously**. Our security appliance is based on the next-generation firewall solution Suricata. This means cost-optimised server architecture enabling seamless integration of IDS/IPS systems with the microservers. The multi threaded approach of Suricata allows a significantly high throughput and is capable of **real time 10 Gbit and faster** network traffic inspection.

FUNDED BY THE EUROPEAN UNION
No. 688201

M²DC

HIGH-PERFORMANCE COMPUTING

Demanding complex distributed applications focusing cost-effectiveness and energy efficiency, with optimised reconfigurable communication and hardware acceleration

TRANSPORT DATA ANALYTICS

High-volume, high-velocity computation of real-time analysis focusing IoT sensors and other data sources

PHOTO PROCESSING

Efficient and scalable image and signal processing services for ow-cost image processing

CLOUD COMPUTING

Cloud vendors with focus on Platform as a Service based on M2DC micro-server platform equipped with ARM processors

Ariel Oleksiak

Head of Energy Efficient ICT Department at PSNC
Coordinator of the M2DC Project

"M2DC appliances will fit market needs for optimised microserver solutions with low maintenance costs and wide options of reconfiguration. This meets energy efficiency requirements imposed by both cost minimisation objectives and emerging regulations."

Andrew Donoghue

Former European Research Manager
Advisory Board of the M2DC Project

"M2DC rightly recognises that homogenous, commodity architectures will not be sufficient to meet the compute requirements of the next generation of IT workloads. The goals of M2DC fit well with the current push to develop new, highly efficient, heterogeneous application-specific architectures for emerging workloads such as IoT, machine and deep learning, and big data analytics."

Contact

Project manager: Ariel Oleksiak
ariel@man.poznan.pl

Scientific manager: Mario Porrmann
mporrmann@cit-ec.uni-bielefeld.de

Dissemination manager: Joao Pita Costa
joao.pitacosta@xlab.si

Innovation manager: Stefan Krupop
stefan.krupop@christmann.info