


COMPLEMENTARY SYSTEMS



In addition to the Karolina and Barbora supercomputers, IT4Innovations also operates smaller complementary systems. These systems represent emerging, non-traditional, and highly specialised hardware architectures. At the same time, these technologies are promising or widely accepted. New programming models, libraries, and software development tools are also deployed in the complementary systems. They provide users of IT4Innovations computational resources, specifically software developers, an opportunity to test their applications on hardware that they may encounter in other computing centres (e.g. the EuroHPC LUMI and Deucalion systems) and upcoming systems, both national and international (e.g. JUPITER).



Fujitsu A64FX processor




The compute nodes are built on **Arm A64FX processors** with integrated fast HBM2 memory. It is a fragment of *Fugaku*, one of the world's most powerful supercomputers, installed at the RIKEN Center of Computational Science in Japan. This first HPC Arm processor is used as a platform for developing a European HPC processor, primarily for the identically used instruction set, including the **SVE** (Scalable Vector Extension) vector instructions. The configuration comprises eight HPE Apollo 80 compute nodes interconnected by a 100 Gb/s Infiniband network.

Edge server




This is an HPE EL1000 server with an **Intel Xeon D-1587** processor (TDP **65 W**), which is designed to process artificial intelligence tasks directly at the data source, often outside the data centre. The server has high computing power for AI inference thanks to the **NVIDIA Tesla T4** GPGPU accelerator (TDP **70 W**), several communication technologies (10 Gb Ethernet, Wifi, LTE) and low power consumption. The server's power consumption is low compared to high-performance but power-intensive hardware in production systems.

Virtual Desktop Infrastructure (VDI)



It consists of two HPE ProLiant DL 385 Gen10 servers, each equipped with two **AMD EPYC 7413** and two **NVIDIA A40** 48 GB GPUs. VDI is primarily designed for efficient pre-processing and post-processing of tasks and models for large-scale simulations and computations. It provides users with a high-performance MS Windows OS virtual environment with a graphical interface and a focus on 3D OpenGL and RayTracing applications for convenient remote desktop work.

AMD Zen3 processors, AMD accelerators, Xilinx FPGA




This partition is built on AMD technologies. The servers are equipped with third-generation AMD EPYC processors (**AMD EPYC 7513**), four **AMD Instinct MI100** GPGPU cards, which are interconnected by AMD Infinity Fabric, and two different **Xilinx Alveo** FPGA cards (U250 and U280). AMD MI100 is the predecessor of the graphics cards installed in the *LUMI supercomputer* (AMD MI250X).

COMPLEMENTARY SYSTEMS


FPGA Synthesis Server

FPGA synthesis tools typically run for several hours up to one day to generate the final bitstream (logic design) of large FPGA chips. These tools are usually sequential; therefore, the system includes a dedicated server for this task. The FPGA Synthesis Server is used by the tools required for the FPGA cards that are installed in complementary systems equipped with Intel Xeon Ice Lake Gold 6338 and AMD EPYC 7513 processors.




NVIDIA Grace CPU Superchip

The **NVIDIA Grace CPU Superchip** server is built on two processors, each with 72 **Arm Neoverse V2** cores. The Arm architecture and its particular implementation with fast RAM (LPDDR5X) represents a new trend in HPC. This processor, combined with NVIDIA GPUs, will also be installed in **Europe's first exascale supercomputer – JUPITER**, at the German Jülich Supercomputing Centre.



AMD Zen3 processor with large L3 cache

This part is based on an HPE Proliant DL 385 Gen10 server with **AMD EPYC 7773X Milan-X** with **768 MB L3 cache**. This unusually large L3 cache allows for significant performance improvements for tasks limited by the random access rate of RAM.




Intel Ice Lake, Intel Optane, Intel FPGA

The compute nodes are built on Intel technologies. The servers are equipped with third-generation Intel Xeon processors (**Intel Xeon Ice Lake Gold 6338**), **persistent (non-volatile) Intel Optane memory** with a total capacity of 2 TB or 8 TB per server, each accelerated by two **Intel Stratix 10** FPGA cards (Bittware 520N-MX).



Intel Sapphire Rapids + HBM

This partition of the complementary systems is equipped with two 48-core **Intel Xeon CPU Max 9468** processors. The server is equipped with a combination of HBM and DDR5 memory to take advantage of both memory types. This allows developers to adapt applications to such a memory system, which will be used, among other things, in the **future European HPC SiPearl Rhea processor**. Users can also compare the impact of HBM memory with that of the large L3 cache included in the AMD Milan-X processor.



IBM Power10

It consists of a single server with two **IBM Power10** processors, DDR4 RAM and fast NVMe storage with a total capacity of 16 TB. The servers are suitable both for computations, the power of which is limited by the performance of the memory system, and for porting applications to the POWER platform, which is specific, for example, by the number of eight logical cores per physical one. The IBM Power10 processors are successors to processors installed, for example, in the US **Summit and Sierra** supercomputers (**Oak Ridge National Laboratory and Lawrence Livermore National Laboratory**) and **Marconi** at the Italian CINECA supercomputing centre.



Arm Ampere Altra + NVIDIA A30 + NVIDIA BlueField-2 E-Series DPU

This partition consists of two Gigabyte G242-P36 servers, with **Arm Ampere Altra Q80-30** processors augmented with NVIDIA A30 programmable GPGPU accelerators and **NVIDIA BlueField-2 E-Series** DPU processors. This complementary system offers users access to hardware that enables the development of applications supporting “in network computing,” i.e. data processing using programmable network elements.

