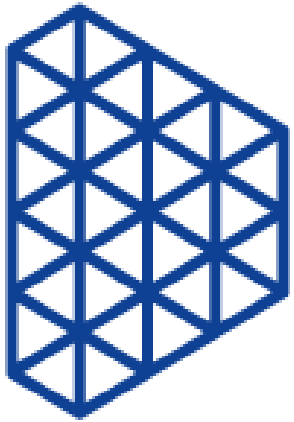

DISCOVERER

Sofia Petascale Supercomputer

What is DISCOVERER Sofia Petascale Supercomputer

EuroHPC Joint Undertaking

- **Discoverer** is a Petascale supercomputer can execute:
  - 4.5 PetaFlops Rmax
  - 6.0 PetaFlops Rpeak

  [1 PetaFlops= 10^{15} Flops = 10^{6} Flops x 10^{9} Flops]
- **Discoverer** is ranked at 91st place among the worlds top 500 supercomputers.
- Discoverer's infrastructure is **co-funded by EuroHPC JU (35%) and by PetaSC** and the Bulgarian government (65%).

91st TOP 500 The List.

- **PetaSC Bulgaria** is a legal consortium combining the knowledge and 15 years of expertise of the National Center for Supercomputing Applications, the Strategic Center for Artificial Intelligence, IICT-BAS and **Sofia Tech Park** (where it is hosted)

SCAI

- **Discoverer objectives are:** to foster better science for society; to facilitate cross-boarder collaborations between academic institutions and the business; and to help training the next generation Bulgarian and EU IT talent.

NCSA NATIONAL CENTRE FOR SUPERCOMPUTING APPLICATIONS

European High Performance Computing Joint Undertaking (EuroHPC JU)

DISCOVERER
Sofia Petascale Supercomputer

- Role: Developing a World Class Supercomputing Ecosystem by pooling the resources of the European Union, European countries and private partners
- Executive Director: Anders Dam Jensen
- Set up: 2020
- Number of staff: 15
- Location: Luxembourg, Luxembourg
- Website: European High Performance Computing Joint Undertaking: https://eurohpc-ju.europa.eu/index_en

EuroHPC JU

DISCOVERER
Sofia Petascale Supercomputer

EuroHPC JU aims to:

- develop, deploy, extend and maintain in the EU a world-leading federated, secure and hyper-connected supercomputing, quantum computing, service and data infrastructure ecosystem
- support the development of a demand-oriented, user-driven supercomputing system, based on a supply chain that will guarantee the provision of the related components, technologies and knowledge, thereby limiting the risk of disruption
- develop a wide range of applications optimised for these systems
- widen the use of this supercomputing infrastructure to a large number of public and private users
- support the development of key skills in high-performance computing in European science and industry.

System Architecture & Specs

DISCOVERER
Sofia Petascale Supercomputer

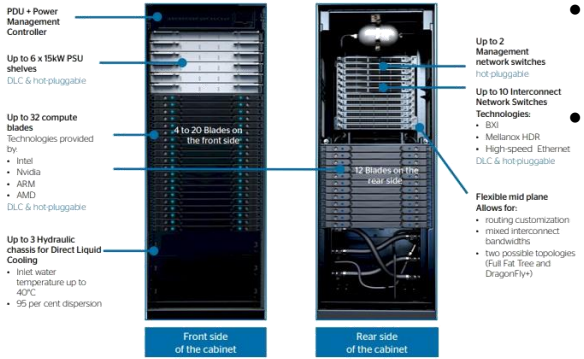

BullSequana X H2410 AMD blade

Design	1U blade comprising 3 compute nodes side-by-side
Processors	3x2 AMD® EPYC Rome® Processor
Architecture	3x1 motherboard
Memory	3x16 DDR4 memory slots (max 2048GB with 128 GB DIMMS)
I/O slots	InfiniBand HDR 1 port mezzanine board PCIe gen4 BXI 1 port mezzanine board
Storage	3x1 optional NVMe M2 format

- **12** Direct Liquid Cooling **BullSequana XH2000 Racks** with up to **32 blades per rack. i.e. Discoverer has 376 blades (12x32-8=376)**
- **376 blades** x **3 nodes/blade** = **1128 computing nodes**
  - **2** x **AMD EPYC processors per node i.e. Discoverer has in total of 2256 AMD EPYC CPUs (2.6 GHz normal freq)**
  - **256 GB shared memory per node with 18 fat nodes with 1024GB RAM (3200MT/s DR)**
- **2256 CPUs** x **64 cores/CPU** = **144 384 Nodes**
- Total size of the RAM reaches over 300 TB (1128x256+18x(1024-256)=316 416GB)
- **Fast disk storage DDN (>20GBps r/w IO) with total capacity of 2 PB**
- 2 racks with auxiliary (management, login and storage) infrastructure
- **Internode connection with IB 200Gbps HDR with DragonFly+ topology**
- The entire system is backed up against a power failure using an uninterruptible power supply with an output of **1 MW.**
- The whole system weight is over **30 tons & total power consumption of 1.3MW** (incl. supporting infrastructure).

Compute node design

- **CPU model: AMD EPYC 7H12** , **64core**, 2.6GHz, 280W ; Next generation x86 “Zen2”
- **CPU sockets per node : 2** ;
- **CPU Cores per node : 128**;
- **Main memory per node : 256GB** (Each of the 18x Fat nodes has 1024GB Memory)
- Memory type and frequency : 16GB DDR4 RDIMM 3200MT/s DR; (The fat nodes are equipped with 64GB DDR4 RDIMM 3200MT/s DR)
- **Node DP TeraFlop/s peak : 5.325TFlops;**
- **TFlop/s sustained Linpack : 3.940TFlops**;
- DP TeraFlop/s peak vs Linpack : 0,74 ;
- **Linpack node power consumption : 665.1 W per 256 GB compute node**; 747.0 W per Fat compute node (Cooling subsystem power consumption excluded);
- Number and bandwidth of network interfaces : 1x 200Gbps HDR

High performance Network

- Interconnect family : IB HDR;
- Interconnect bandwidth per link : 200Gbps (IB HDR) ;
- Expected latency (worst case for a 1 kB message): 520ns;
- Interconnect topology : DragonFly+ ;
- Number of compute nodes per isle (2 Racks) : 192 ;
- Blocking factor within isle : 2:1.

High-speed Interconnect network topology

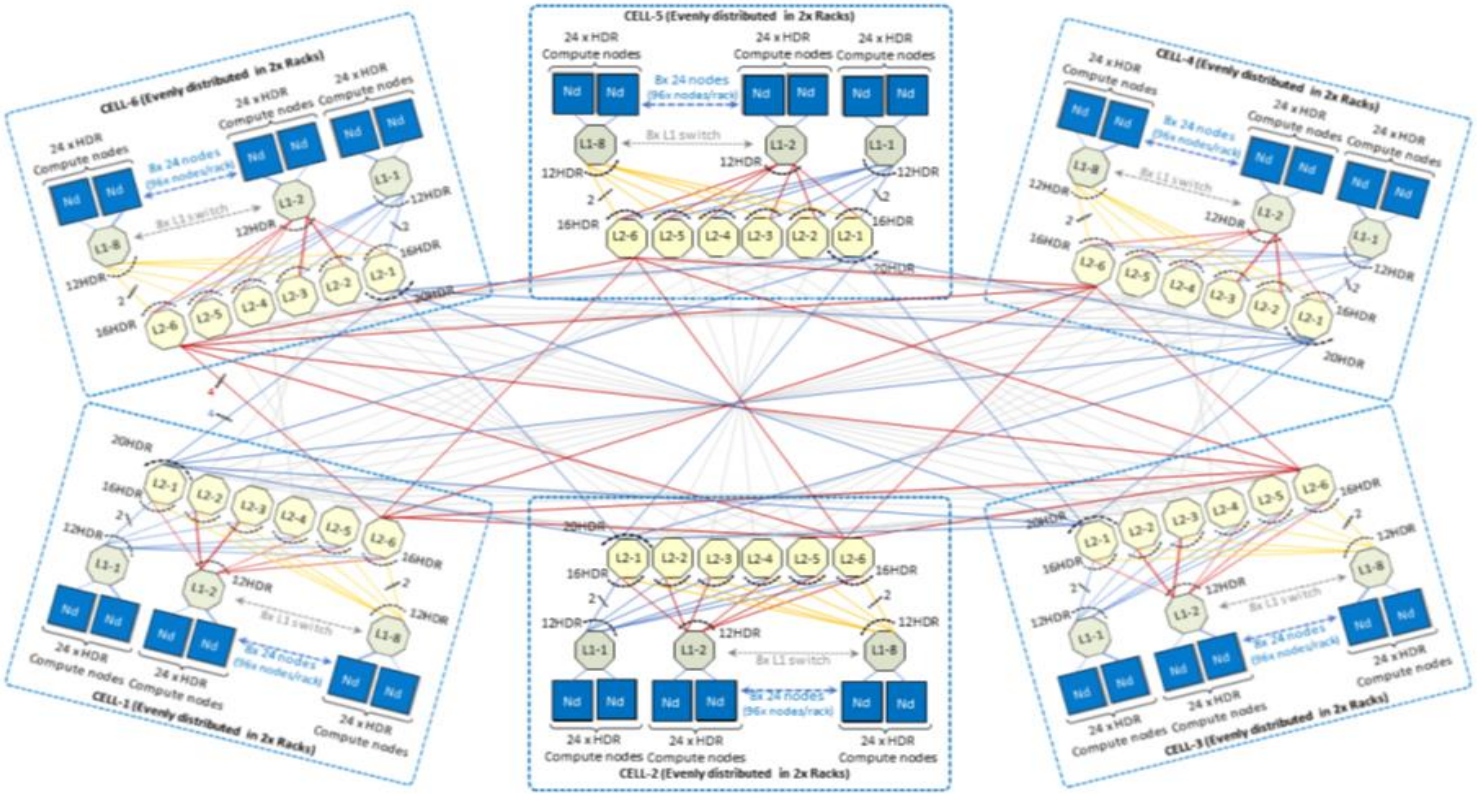

Software Environment

System Software

- Operating system : Red Hat Enterprise Linux 8;
- Compiler Suite(s) : AMD Optimizing C/C++ Compiler ; AMD Open64 SDK; x86 Open64 Compiler System; GNU Compiler Collection ;Python; R;

Numerical libraries

- AMD Optimizing CPU Libraries (AOCL); FFTW; Portable Extensible Toolkit for Scientific Computing (PETCs Library); Trilinos; BLAS; LAPACK/ScaLAPACK; EigenSolver

Debugging/ profiler tools

- AMD µProf ; IO Instrumentation (IOI) ;Lightweight Profiler (LWP) ;Modular End-of-Job Report (MEJR); Scalasca; BSC Performance Tools, Extra-P , SCORE-P

Resource and workload manager

- SLURM ; Singularity;

Software & Target Application Areas

DISCOVERER
Sofia Petascale Supercomputer

Open-source Software (short list)

- Bioinformatics / Genomics (BLAST/RAY/EXCALATE/(HAD)DOCK/ROSETTA)
- Computational & Quantum Chemistry (CP2K/CPMD/Quantum Expresso/ GAMES)
- Molecular Dynamics & Mesoscale Modelling, Monte Carlo (GROMACS/NAMD/LAMMPS/DL POLY)
- Computational Fluid Dynamic / Finite Elements Methods (Open FOAM/Alya/SALOME)
- AI / Big Data Analytics (Tensor Flow/Python ML Libraries /NEURON)

Application Areas:

- *In-silico* Drug Discovery
- Structure-Property Predictions & Molecular Discovery
- Digital Product Formulation & Optimization
- Climate & Whether Forecasting / Environmental Modelling
- Simulated Environments in Automotive & Civil Engineering
- FinTech/MarkeTech & Big Data (DL/ML/AI)

Discoverer Access Policy

DISCOVERER
Sofia Petascale Supercomputer

- **EuroHPC has 35% resource share from the Discoverer's resources, which will be allocated via regular calls for:**
  - Regular HPC projects (~80-90%) (for projects that require significant amount of HPC resources)
  - HPC project benchmarking (few %) (to determine project suitability & scalability, which is prerequisite to apply for a regular project)
  - HPC Software Development & benchmarking (few %)
  - Fast Track applications (~10%) (for example to get extra data quested during the peer-review process of previous project)
- EuroHPC calls will be organised via **PRACE** and are targeting pan-EU HPC projects including academic & industrial applications (with 80:20 split – see bellow)

- **PetaSC's resource share is 65% and it will be allocated using similar to EuroHPC access policy**
- PetaSC calls will be organised by PetaSC's Scientific & Applicant Board and the applications will be assessed by national & international scientific panel.

- For both EuroHPC (35%) and PetaSC (65%) share, there an additional **(min)80:(max)20** split between **free access** for purely scientific/non-profit and **paid access** for industrial/for-profit applications.
  - the distinction between free & paid access will be made on the base of results dissemination, where free access projects have obligation to openly share the all results from the project, while paid access calls can retain (part of) their results and use them for-profit /business proposes.
  - i.e. consortium from academic & industrial partners can qualify for free access to the Discoverer, provided they prepare suitable project proposal, apply for a call and get selected by the panel and after project end to openly share all the data and results from the project.

Applications

DISCOVERER
Sofia Petascale Supercomputer

Examples of Successful Applications for the Discoverer PS supercomputer:

- Emergency situations forecast
- Quantum Computing Applications
- Computational Chemistry
- Preventing Asthmatic People from their problems
- Industrial Applications: Rendering, Digital Twins, …

Thank you for your Attention!

- „There is nothing more useful for the practice than the good theory.“

  Kurt Lewin

  (1890-1947)

- „In theory, the theory and practice are one and the same. However, in practice they are not.“

  Albert Einstein

  (1879-1955)