EuroHPC

Andrej Filipcic
Jozef Stefan Institute, Ljubljana, SI

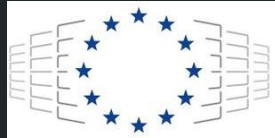

EuroHPC
Joint Undertaking

EuroHPC Joint Undertaking

- Phase-I : 2018-2021
  - 3 pre-exascale (250-350 PFlops)
  - 5 peta-scale (5-20 PFlops)
- Phase-II: 2021-2027, currently approved:
  - High-range: 1 exa, Jupiter@Julich
  - Mid-range: 4 HPCs
  - More coming in the next few years
- Quantum Computers:
  - Collocated to some HPCs
  - Likely 3 approved next month, 3 different technologies

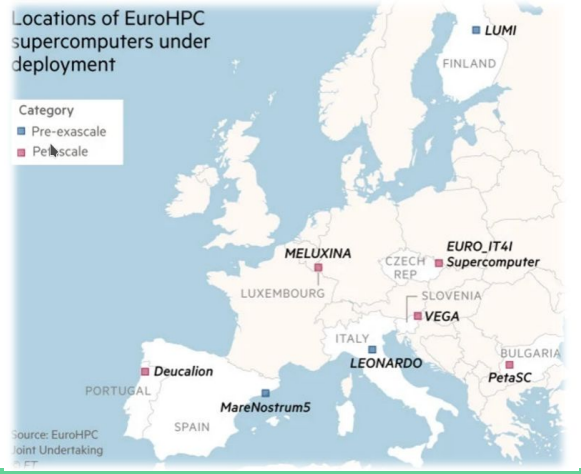

Mission

- 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 and uptake of demand-oriented and user-driven innovative and competitive supercomputing system based on a supply chain that will ensure components, technologies and knowledge limiting the risk of disruptions and the development of a wide range of applications optimised for these systems;
- widen the use of that supercomputing infrastructure to a large number of public and private users and support the development of key HPC skills for European science and industry

Budget

- 50% European Commission, 50% hosting states (65% for peta/mid-range)
- Phase-I: ~1B€
- Phase-II: ~7B€
  - 200M€ for hyperconnectivity
  - ½ for infrastructure

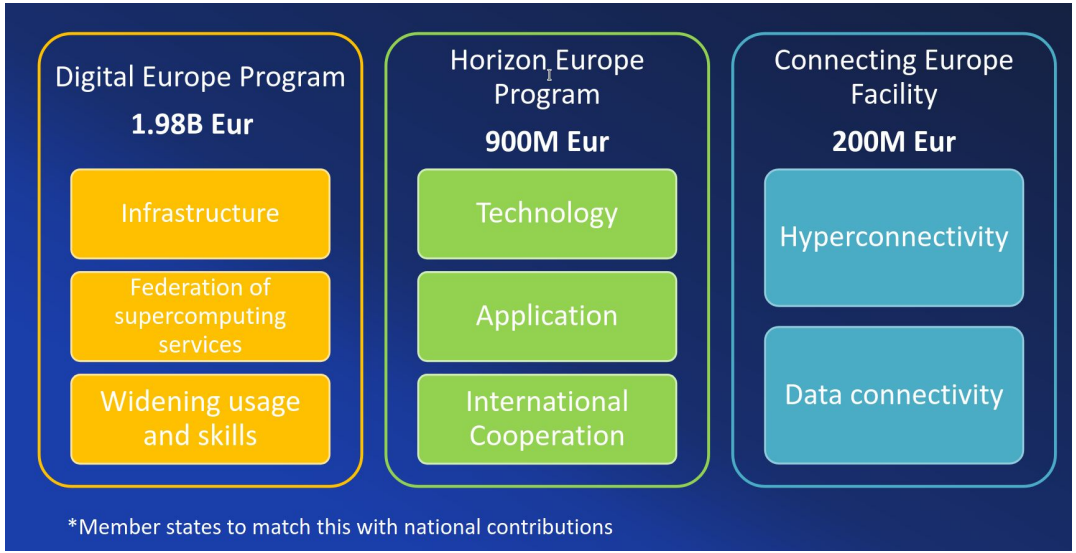

Projects and activities

- EuroCC, EuroCC2 - competence centers:
  - In each member state, coordinated by Castiel project, budget ~1M€/country/year
  - Goals:
    - Develop and display a comprehensive and transparent map of HPC competences and institutions in their country
    - Act as a gateway for industry and academia to providers with suitable expertise or relevant projects, may that be national or international
    - Collect HPC training offers in their country and display them on a central place together with international training offers collected by other NCCs
    - Foster the industrial uptake of HPC
- Centers of Excellence:
  - Support sw development, scalability etc... ~10, 6-8M€/project
- RIAG (research and innovation), INFRAG (infrastructure) advisory groups
  - Recommendations for evolution and development
- EPI - European processor initiative
- EUMaster4HPC - common university program ... many other projects... (~30 ongoing)

Resource allocations

- Access for EU entitled users
- EC resource share - PRACE-like calls for applications
  - development & benchmarking (immediate)
  - regular (peer reviewed)
  - Industry & public sector
- Hosting entity share:
  - Depends on the country policies and decisions
  - Eg WLCG pledges

Design

- Some HPCs are “classical”, but not all
- Vega (SI):
  - Scaled up for extreme networking, including WAN (200Gb/s atm)
  - Data processing (440GB/s actual input file reading for ATLAS), ceph, lustre…
  - Outbound connectivity (ipv4 works, ipv6 not supported by IB-ETH GW)
- LUMI (FI):
  - Dedicated partition for virtualization and services
  - Ceph Object Storage
- MareNostrum5:
  - Not clear yet, but they will support LHC (government agreement) with pledges

.....

Architecture

- Most are general purpose machines
- CPU (AMD, Intel, ARM)+ GPU (NVidia, AMD) partitions
  - Most nodes have GPUs (60-80% on large machines)
- Storage: Lustre, Ceph, some GPFS - throughputs from 200-2000 GB/s
- Future: ARM (near) + RISC-V (medium)
- Plans for Tbps interconnects, calls next year

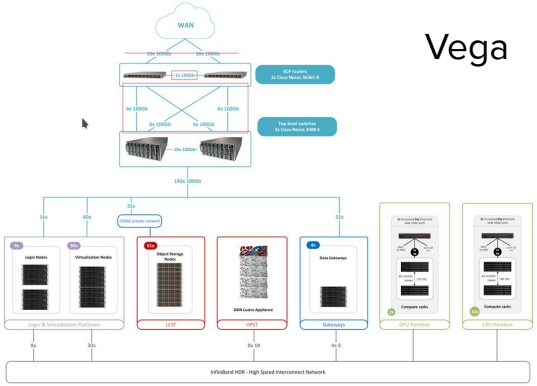

Vega

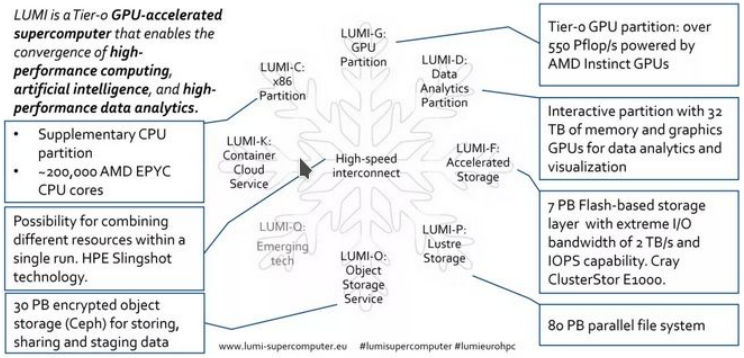

After 1.5 years of operation

- ~500 users per HPC
- Many different kinds of workloads:
  - Many-node compute (cpu, gpu), mostly chemistry, material science
  - Small node/core parameter scans - independent jobs
  - Machine learning - rising rapidly (tensorflow)
  - Extreme data processing (not only LHC), eg Cryomicroscopy
- Some HPCs allocate full nodes only
- Many (non-LHC) users are unhappy with “default” data copy over ssh - something more organized needed
- Many users require containers or even virtualization

Impact

- Much more users than in PRACE, a lot of newcomers due to simplicity of access
- Usage for industry is rising, mostly small or medium enterprises
- Several countries decided to provide resources through (Euro)HPCs to LHC (Slovenia, Spain, Germany...)
- Network connectivity to Geant will likely be boosted in the next 2 years
- Long term data storage is not part of the plans yet - needs cooperation with EOSC, EGI and WLCG...
- Some leadership EU project, eg DestinationEarth - resource pledges approved for production by EuroHPC JU
    - More to follow, (LHC?)

Sites: ~7k CPU nodes, ~800k cores

- Vega:
  - **960** CPU nodes, 128-core AMD Rome
  - 60 GPU nodes, 128-core AMD Rome + 4x A100
- LUMI:
  - **1536** CPU nodes, 128-core AMD Treno
  - ~4000 GPU nodes, 64-core AMD Trento, 4x AMD MI250X
- Leonardo:
  - **1536** CPU nodes, 52-core Intel Sapphire Rapids
  - 3456 GPU nodes, 32-core Intel Ice Lake, 4x A100-64
- Deucalion:
  - 1632 Fujitsu PRIMEHPC FX700
  - **400** CPU nodes, 128-core AMD Rome
  - 33 GPU nodes, 128-core AMD Rome + 4x A100
- Discoverer:
  - **1128** CPU nodes, 128-core AMD Rome
- Karolina:
  - **720** CPU nodes, 128-core AMD Rome
  - 72 GPU nodes, 128-core AMD Rome, 8x A100
- MeluXina
  - **573** CPU nodes, 128-core AMD Rome
  - 200 GPU nodes, 128-core AMD, 4x A100
- Mare Nostrum 5: no details yet (procurement)

EuroHPC Roadmap

	2019 & 2020	2021	2022	2023	2024	2025	2026	2027
HPC Infrastructure	3 pre-exascale and 5 petascale HPC systems	Several pre-exascale systems and 2 exascale HPC systems				One or more exascale and post-exascale HPC systems		
Quantum Infrastructure	Quantum simulators interfacing with HPC systems	First generation of quantum computers		Quantum simulators interfacing with HPC systems		Second generation of quantum computers		

Details on EuroHPC Supercomputers

- https://eurohpc-ju.europa.eu/about/our-supercomputers_en