



EUROHPC AND COMPANY: A ROADMAP TOWARDS EXASCALE

Dirk Pleiter | EXALAT Kick-off Workshop | 16.06.2020

Overview

- Joint Undertaking EuroHPC
- EuroHPC's research and innovation agenda
- Towards a European data infrastructure
- Summary and outlook

JOINT UNDERTAKING EUROHPC

Joint Undertaking EuroHPC



Joint Undertaking

- Legal entity involving European Commission, EuroHPC Member States, private members
 - 32 member states
 - Private members: ETP4HPC, BDVA

Mission

- Develop, deploy, extend and maintain in the European Union (EU) an integrated world-class supercomputing and data infrastructure
- Develop and support a highly competitive and innovative HPC ecosystem

Achieved milestones

- November 2018: Establishment of the Joint Undertaking
- July 2019: Publication of first calls for Research and Innovation actions
- November 2019: Call for tenders for 3 pre-exascale systems



Organisation



Governing Board

- Main decision body
- Composition: European Commission, member states

Research and Innovation Advisory Group (RIAG)

- Role: Create and regularly update the draft multi-annual strategic research and innovation agenda

Infrastructure Advisory Group (INFrag)

- Role: provides advice to the Governing Board for the acquisition and operation of the EuroHPC supercomputers, drawing up and regularly updating the draft multi-annual strategic agenda for such acquisition

EuroHPC office + Executive Director

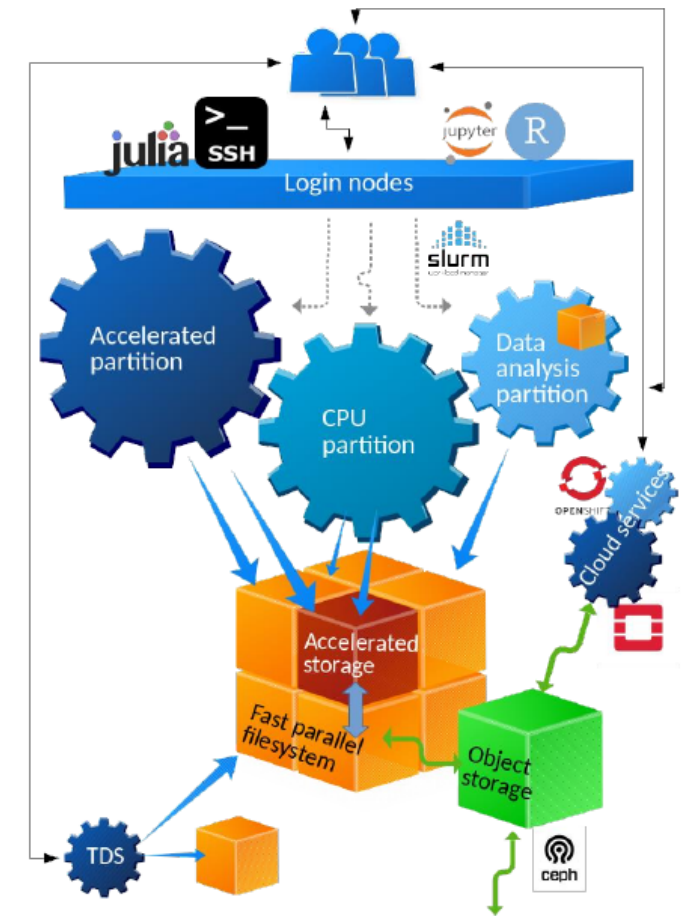
Upcoming EuroHPC Infrastructure: Plans

Pre-exascale systems

- Target HPL performance: >150 PFlop/s per system
- Acquisition budget: 120-150 MEUR
 - Up to 50% of TCO funded by EuroHPC, other by hosting entity
- Awarded hosting entities
 - Leonardo@CINECA (Italy)
 - LUMI@CSC (Finland)
 - Mare Nostrum 5@BSC (Spain)
- 10-20 MW power envelope → GPUs likely to play a key role
 - But not planned to be monolithic systems like Summit or Sierra

Petascale systems

- IT4Innovations (Czech Republic), LuxConnect (Luxembourg), University of Minho (Portugal), Maribor (Slovenia) and Sofia (Bulgaria)



Upcoming EuroHPC Infrastructure: Status



Pre-exascale systems

- Call for tenders launched in November 2019
 - Competitive dialogue → likely >6 months until awarding contracts
- Installation of (at least phase 1 of) the systems by end of 2020
- Operations starting early 2021

Petascale systems

- All tenders launched (last in May 2020)
- Direct procedure → fast awarding of contracts possible

EUROHPC: R&I AGENDA

ETP4HPC's Strategic Research Agenda

ETP4HPC = European Technology Platform for HPC

- Private, industry-led organisation
- Main mission: promote European HPC research and innovation

Regularly updates a Strategic Research Agenda

- Goal: Outlines a roadmap for the achievement of exascale capabilities by the European High-Performance Computing (HPC) ecosystem

Work results from open working groups

- System Architecture
- System Hardware Components
- System Software & Management
- Programming Environment
- I/O & Storage
- Mathematics & Algorithms
- Application co-design
- Centre-to-Edge Framework



<https://www.etp4hpc.eu/sra.html>

Topics Relevant to the LQCD Community

System architectures

- Integration of heterogeneous resources
- Network

System Hardware Components

- Integration and packaging

I/O and storage

- Integration of non-volatile memory in the I/O stack
- Memory-style addressing of persistent storage
- Automatic placement of data

Mathematical methods and algorithms

- Robust methods and algorithms enabling extreme scalability
 - Increase of parallelism, reduction of synchronisation
 - Exploitation of reduce precision and new floating-point formats
- Vertical integration and validation of mathematical methods and algorithms

European Processor Initiative (EPI)



Mission

- European independence in High Performance Computing Processor Technologies
 - EU Exascale machine based on EU processor by 2023
- Based on solid, long-term economic model, Go beyond HPC market
- Address the needs of European industry
 - Example: car manufacturing market

Organisation

- 27 partners from 10 European countries
- Including company founded to produce EPI processors: SiPearl

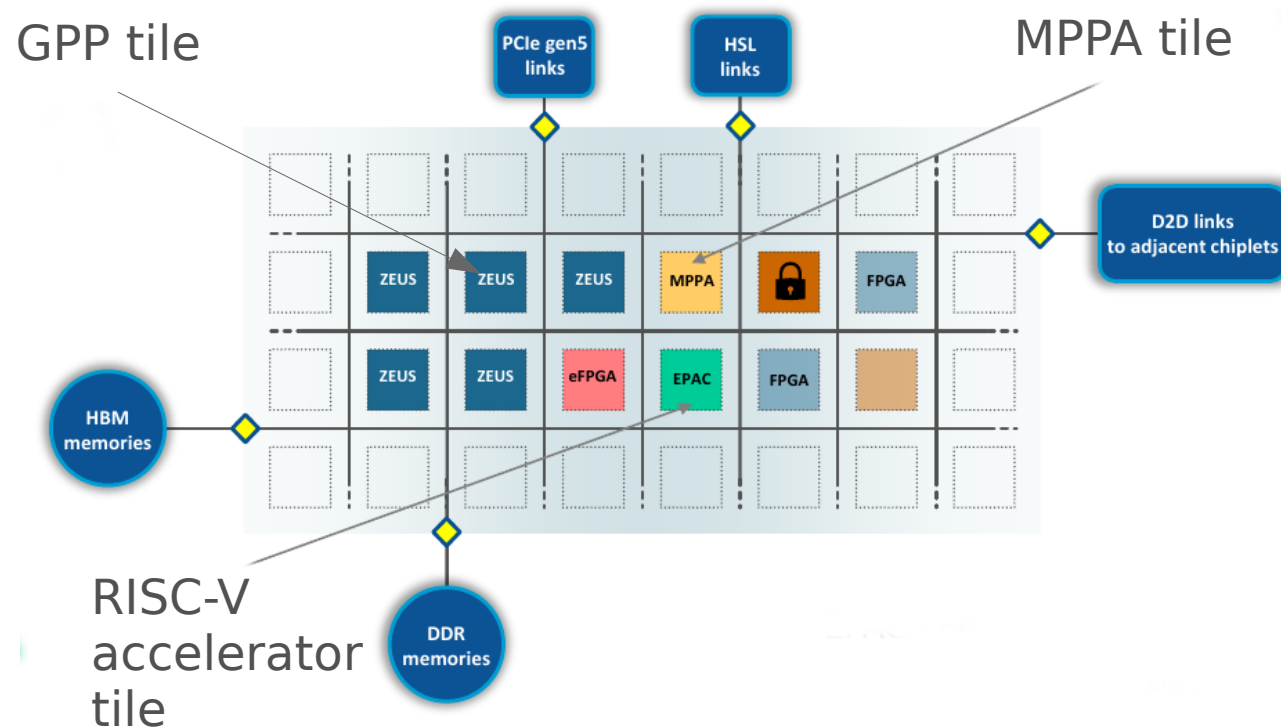
EPI: Common Platform Approach

Integration of different technological components

- E.g. Tiles with Arm cores, RISC-V accelerators, MPPA, eFPGA, ...

Global approach for power management and security

Modular approach allows (in principle) flexible customization



EPI: GPP Stream

Based on established processor core technology

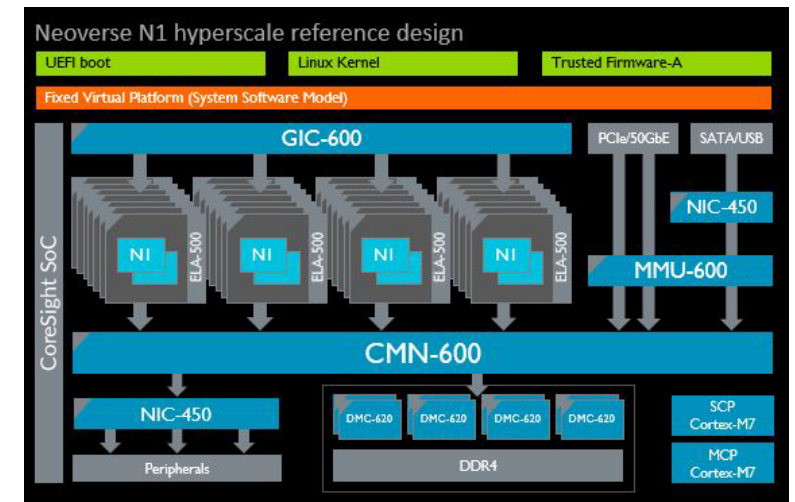
- First generation based on Arm's Zeus core

Zeus is part of Arm's Neoverse processor family

- High-performance processors for AI, Cloud, HPC and edge
- Aim for processor designs with large number of cores
- Armv8-A ISA

Support for Scalable Vector Extension

- Key feature: Vector length agnostic
 - Vector length not defined at compile time
- Required hardware support for VLA
 - Update of predication registers
 - Update of loop counters



© Arm

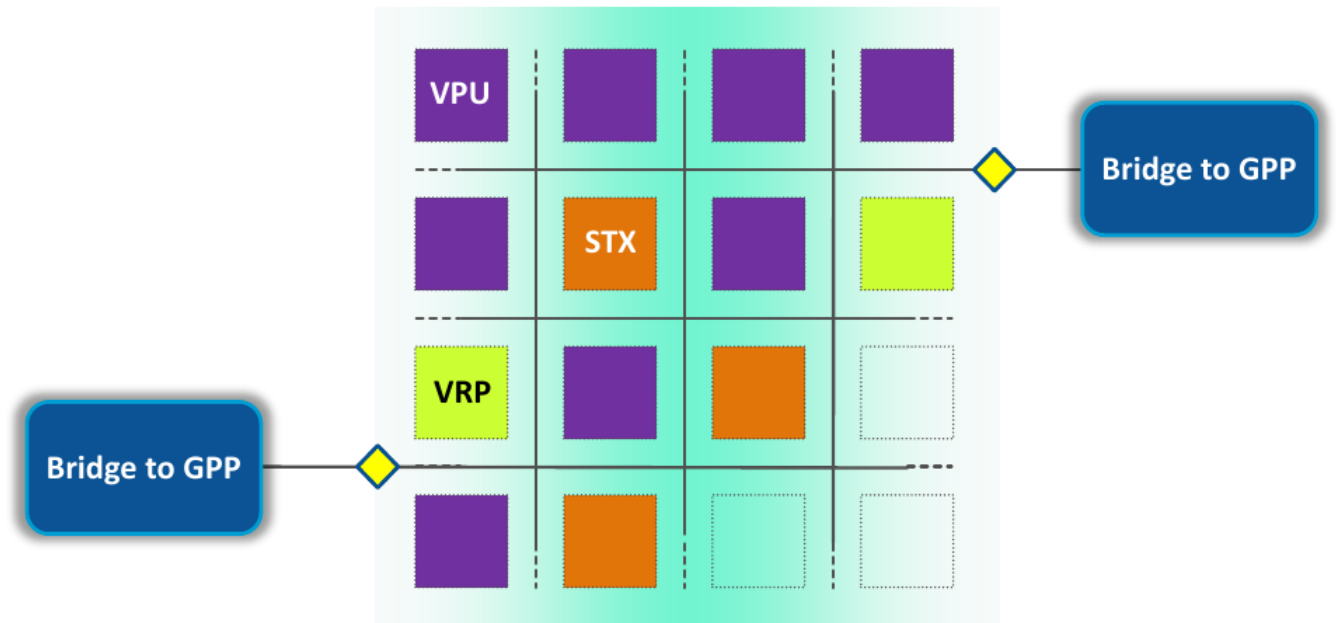
EPI: Accelerator Stream

Vector processing unit

- Based on RISC-V vector ISA

Domain specific accelerators

- Stencil/tensor accelerator (STX)
- Variable precision processor



TOWARDS A EUROPEAN DATA INFRASTRUCTURE

European Open Science Cloud (EOSC)

Initial mission statement

- The main goal of the EOSC initiative was to offer European researchers a virtual environment with free, open, and seamless services for the storage, management, analysis and re-use of research publications, data and software that are linked to their research activities across borders and disciplines

Focus on FAIR data policy

- Principles of find-ability, accessibility, interoperability, and re-usability

Emerging cloud and data infrastructure

- Common AAI
- Efforts towards solutions for long-term archiving of large quantities of open data
- Cloud services for, e.g., data analysis workflows
- Opportunities for better integration with HPC infrastructures

From ILDG to EOSC?

International Lattice Data Grid	EOSC-based Solution
X.509 certificate-based authentication challenging	Industry-standard token-based solutions (OIDC, SAML)
Proprietary, difficult to install clients	Clients based on software-components support in typical Linux distributions/Python package managers
Lacking integration of integration of HPC systems and federated data infrastructure	Projects working towards better integration of HPC and Cloud-based storage resources (e.g. Fenix)

SUMMARY AND OUTLOOK

Summary and Outlook

EuroHPC's future infrastructure

- Successful effort for putting 3 pre-exascale systems in place
- Precursor for exascale systems

Upcoming R&D efforts

- LQCD applications not a key focus, but opportunities for LQCD applications to benefit
 - Higher visibility through Centres of Excellence would be beneficial

European Processor Initiative

- Arm SVE + use of high-bandwidth memory technologies interesting for LQCD
- Opportunities related to accelerators to be explored

European data infrastructure

- New opportunities for sharing LQCD data