

Cray's HBP-PCP project

Highlights of your project

Cray, Juelich, EPFL, CSCS and other partners are currently engaged in the co-design of a platform for the exploration of next-generation mixed workload use-cases, such as in-situ and in-transit visualisation of neuro-simulation, and usage of machine learning for brain image classification. A pilot system with Intel Knights Landing compute nodes, NVIDIA visualisation nodes and the Intel Omnipath interconnect is installed at Juelich. An experimental storage/memory system consists of high capacity SSD storage, Ceph object store / filesystem, high bandwidth memory and (later) multiple types of non-volatile memory. The project will develop abstractions and frameworks to enable the complex data management, transit and storage needed for seamless execution of the use-cases on the pilot system and elsewhere, and will assess performance and usability implications of doing so.

What are anticipated technology (hw/sw/methodology) suggested for inclusion in an EsD project and describe the current maturity?

- 1) *Tiers of varied memory and storage*
- 2) *Object store*
- 3) *Memory abstraction layer*

All of these technologies are immature, but each seeks to solve a piece of the Exascale IO/data problem and hence is relevant to EsD projects.

How should this technology be used / integrated (I/F, APIs)

The number and variety of access methods for memories mean that Applications should not directly access new memories. We believe that this is the role of some missing middleware. Our project will not produce that middleware, but it will define some requirements to assist with that definition and assess any existing partial solutions.

Are there any pre-or co-requisite items?

A simple non-posix open API for IO would help this investigation and all Exascale IO projects. Also most interesting non-volatile memories are not yet in general release, there is a pre-requisite stabilising and maturing of those technologies required before serious consideration in EsDs.

Any extra work/interaction (on top of current project roadmap) needed to make them ready?

Large amounts of specific technical integration is required, and a small number of inter-connected projects exploring various memories and approaches should be performed. Questions remain around performance implications of the usage of new memories in applications, the precise requirements for and constraints on middleware development, relevance of open frameworks such as openStack, and the implications on resource managers.

What information / actions are needed to best prepare for EsD projects?

Clear definition of some driving EsD use-cases including the data-intensive and mixed workloads. Understanding the data roadmap of target applications (how data sizes will grow over time). Clear definition of next-generation usage models.