

CRAY[®]
EMEA RESEARCH LAB



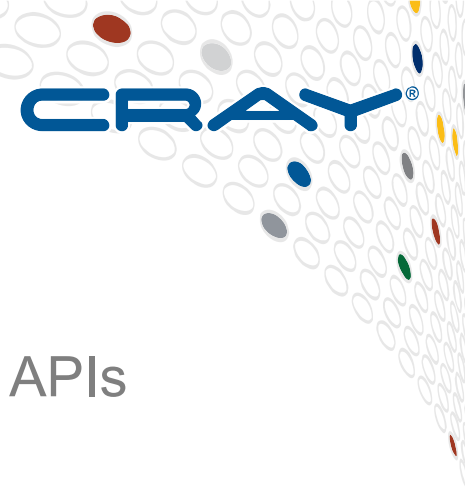
Human Brain Project



Cray's HBP-PCP project

Adrian Tate, Director Cray EMEA Research Lab
adrian@cray.com

Project Overview



- **Objectives**

- Evaluate non-volatile memories, storage class memory and their access APIs
- Study non-standard workloads and their execution / software support

- **Next-gen use-cases**

- In-situ visualisation + neurosimulation
- In-transit visualisation
- Machine learning + simulation

- **Pilot System “JULIA” : installed in Juelich**

- Variety of compute nodes (KNL, data-nodes, viz nodes)
- Intel Omnipath 100Gbit/s network
- Memory / storage system : SSD-backed Ceph, incorporating new memories
- Data-intensive software stack
- Define abstractions and/or software tools to assist with memory usage and transit



Anticipated technology suggested for inclusion in an EsD project



- *Tiers of varied memory and storage*
- *Object store*
- *Memory abstraction layer*

- *All are immature, but highly relevant to Exascale*

How should this technology be used?

- *Do not burden the user more than already!*
- *Software abstraction is necessary*

Are there any pre-or co-requisite items?

- *Desirable : A simple non-posix open API for IO*
- *non-volatile memories are not yet in general release*

Any extra work/interaction?

- *Large amounts of specific technical integration*
- *Need inter-connected projects exploring various memories / APIs*
- *Questions remain*
 - *performance implications*
 - *the precise requirements for and constraints on middleware development*
 - *relevance of open frameworks such as openStack,*
 - *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”***