sage [1] (sāj)
n.
One venerated for experience, judgment, and wisdom.
adj. **sag·er, sag·est**
1. Having or exhibiting wisdom and calm judgment.
2. Proceeding from or marked by wisdom and calm judgment: *sage advice.*

On Building a storage system for HPC + Big Data Use Cases

Sai Narasimhamurthy (Seagate)

Presented at the European Big Data Value Forum, 2017

Per-cip-i-ent (pr-sp-nt)
Adj.
Having the power of perceiving, especially perceiving keenly and readily.
n.
One that perceives.

This project has received funding from the European Union's Horizon 2020 research and innovation programme under grant agreement No 671500

Coordinator: Seagate

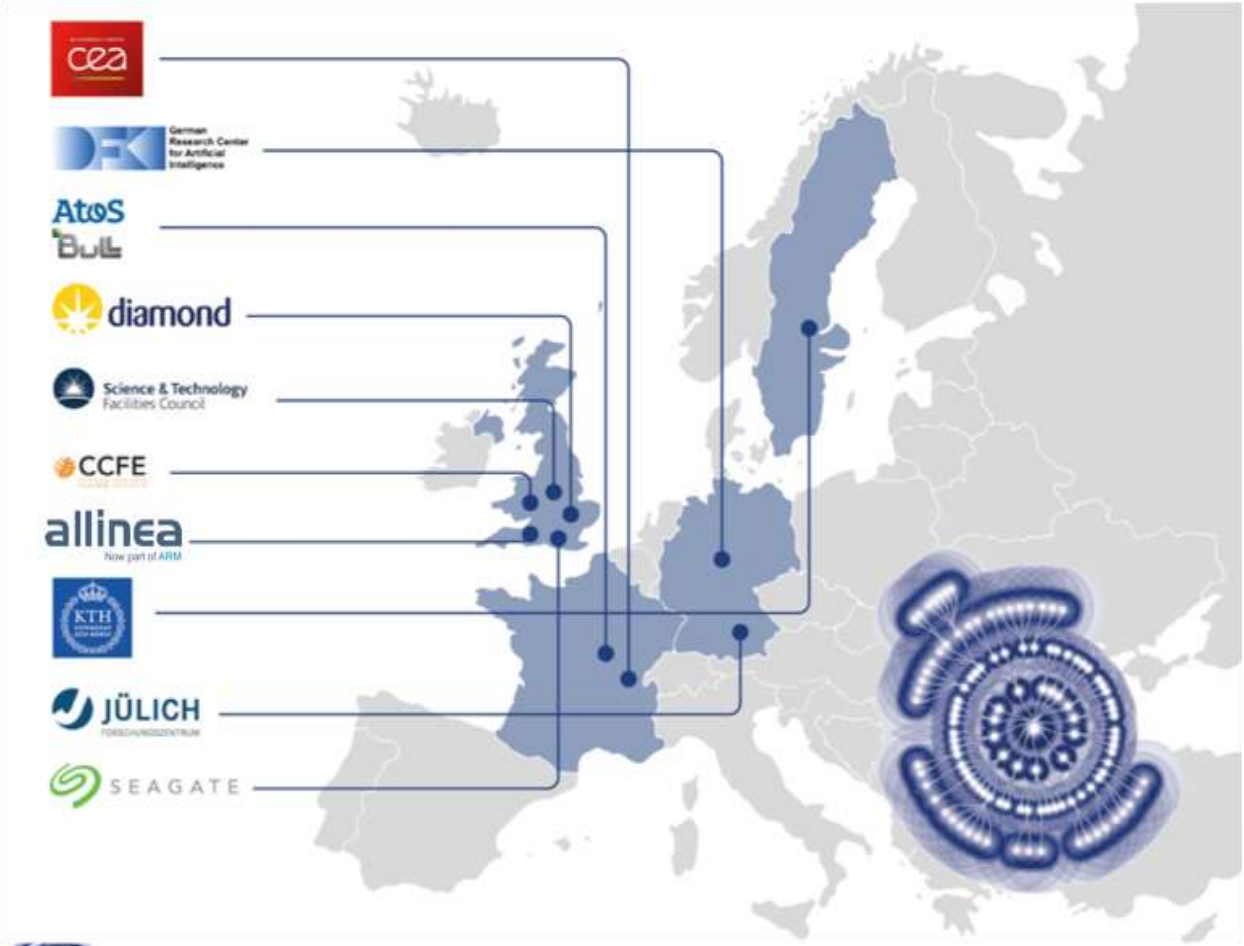

"SAGE"
Percipient
Stor**AG**e
for **E**xascale data
Centric computing

FETHPC Project
~€7.9M
10 partners

SAGE

SAGE Consortium

- *Storage cannot keep up w/ Compute!*
- *Way too much data*
- *Way too much energy to move data*
- *New Storage devices use unclear*
- *Opportunity: Big Data Analytics and Extreme Computing Overlaps*

SAGE

Storage Problems at Extreme Scale

Need Exascale Data Centric Computing Systems

Big **D**ata **E**xtreme **C**omputing (*BDEC Systems*)

SAGE Validates a BDEC System which can Ingest, Store, Process and Manage extreme amounts of data

SAGE

SAGE Project Goal

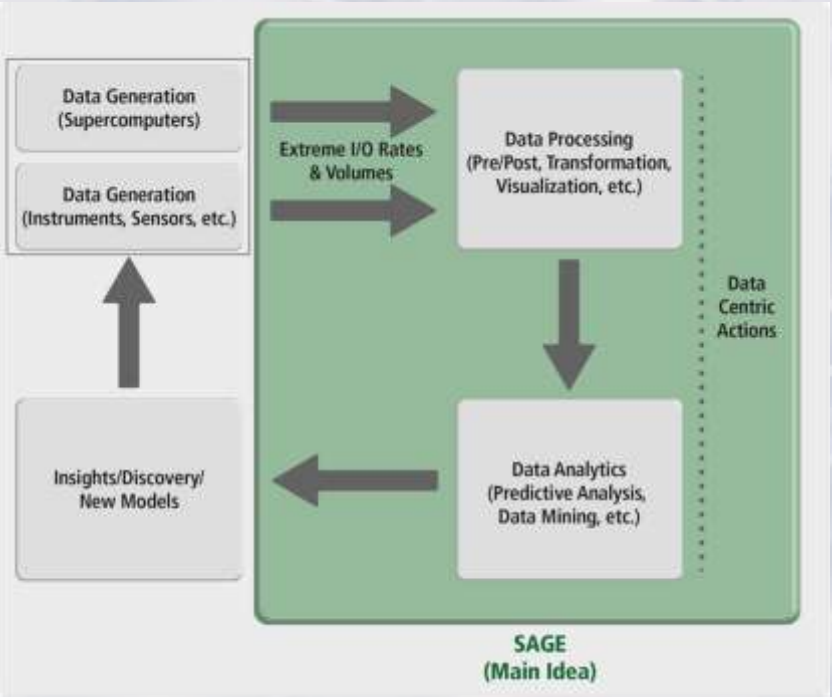

- Deal with high I/O rates
  - Instruments
  - Running Simulations
- Performing data Processing
  - In parallel with ongoing simulation and Data collection
- Insights could be fed back into
  - Running Simulations
  - Data collection decisions

SAGE

A Vision: BDEC Storage System

Improving the time to solution

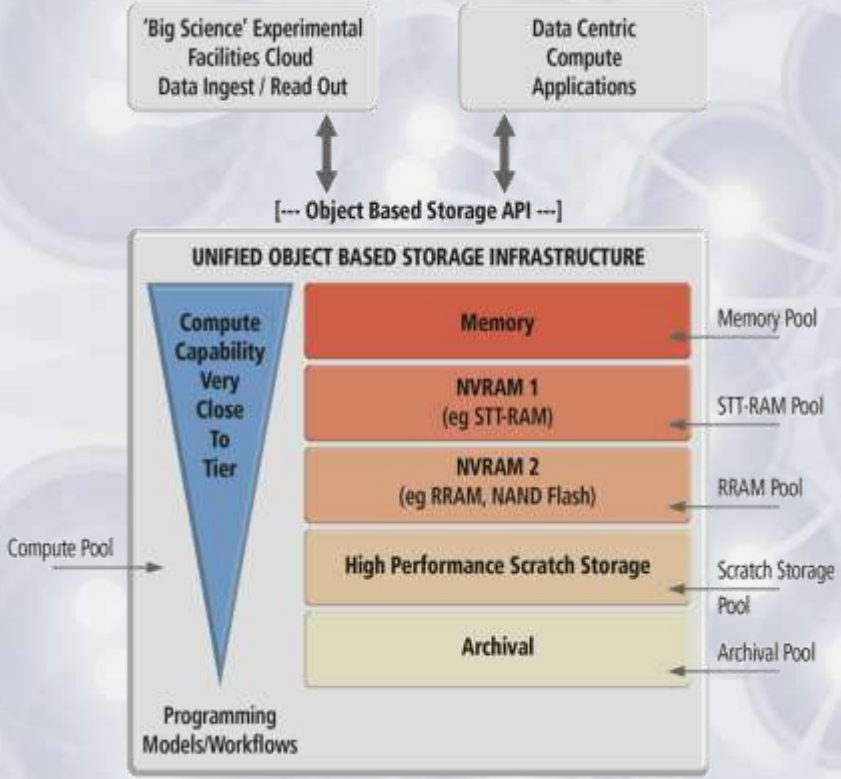

- **Goal**
  - Build the data centric computing platform
- **Methodology**
  - Advanced Object Storage
  - New NVRAM Technologies in I/O stack
  - Ability for I/O to Accept computation
    - Incl. Memory as part of storage tiers
  - API for massive data ingest and extreme I/O
  - Commodity Server & Computing Components in I/O stack

SAGE

SAGE "Percipient" Storage Overview

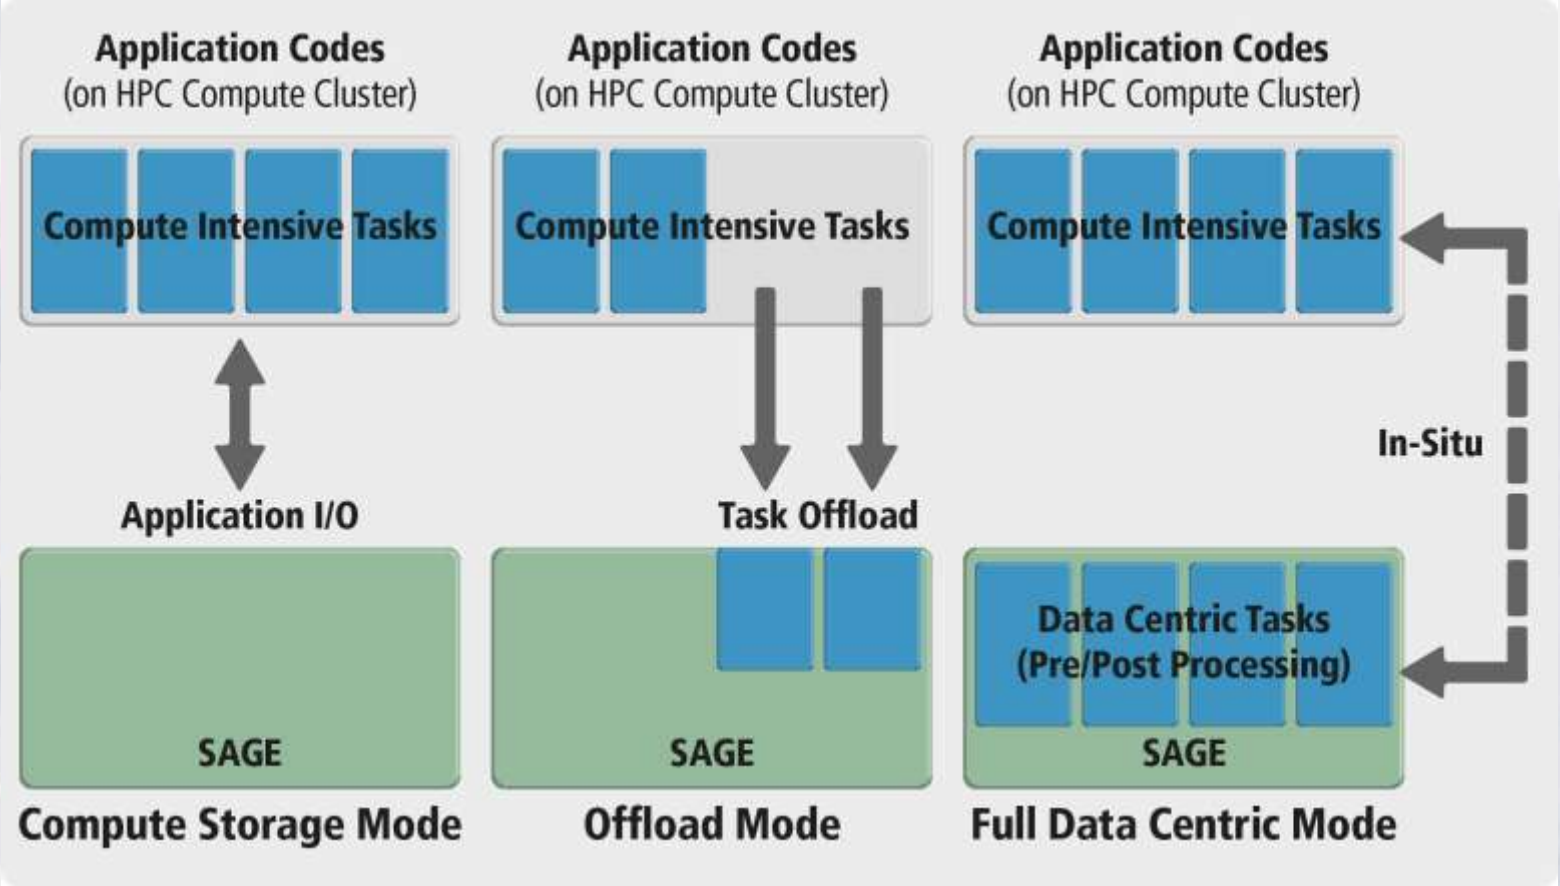

SAGE

Percipient Storage Modes of Operation

SAGE

Percipient Storage: Research Areas

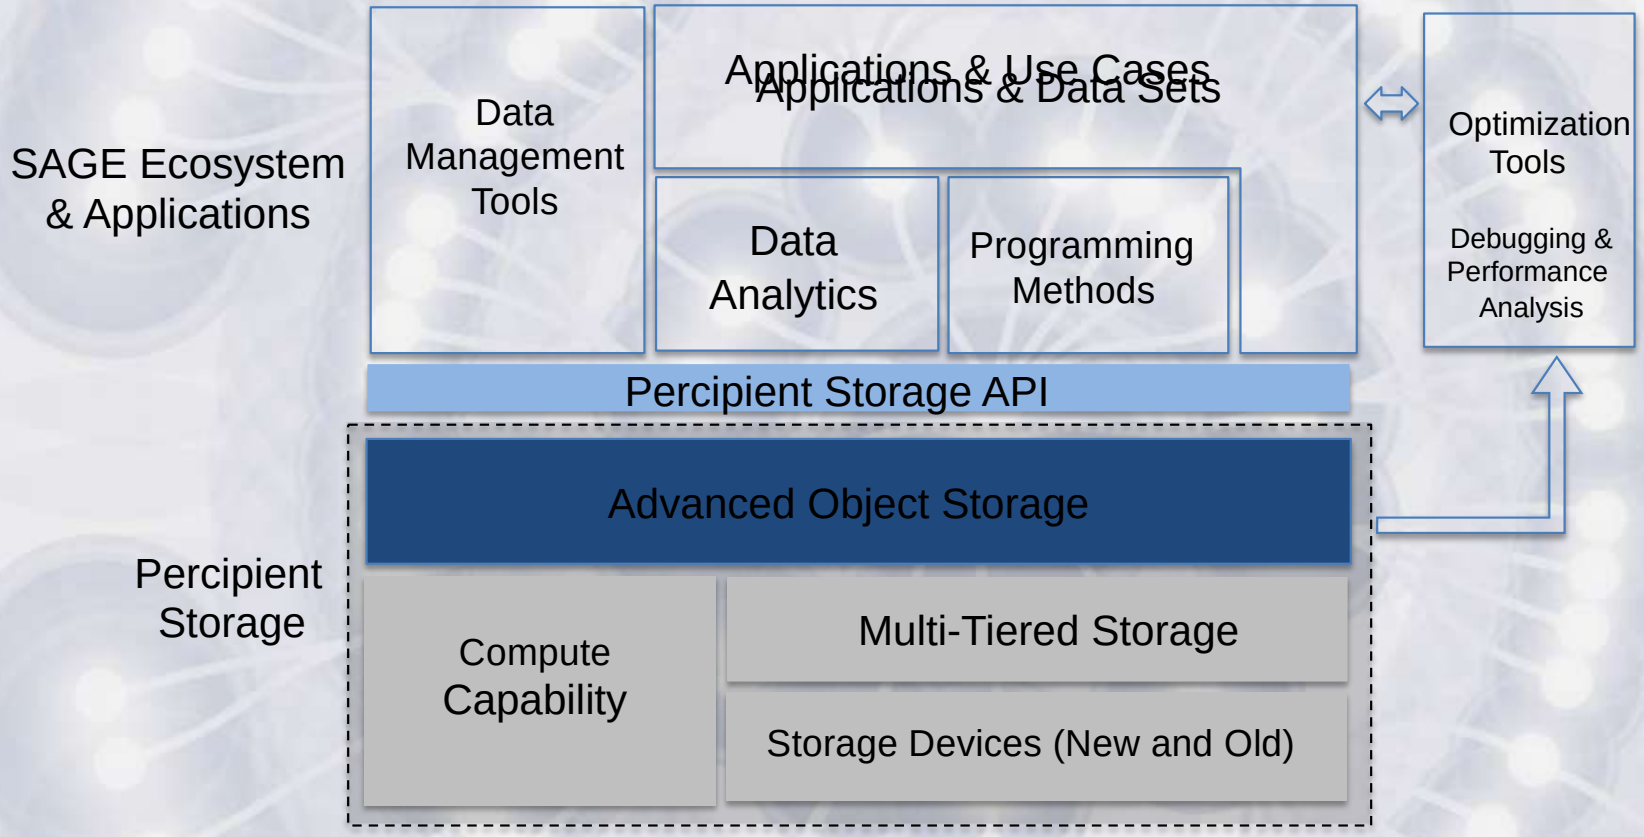

Percipient Storage Stack

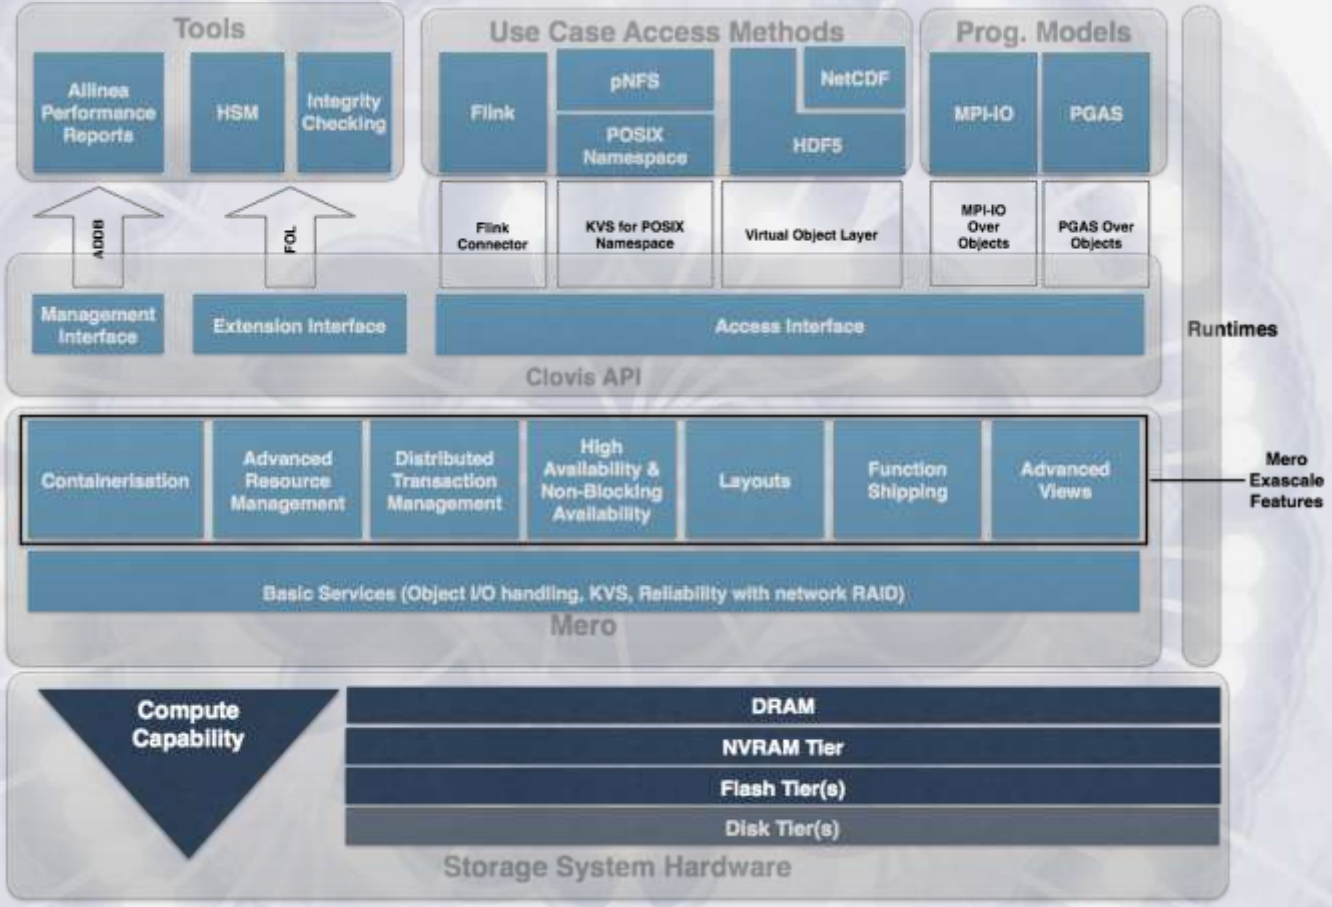

SAGE

Percipient Storage Stack..

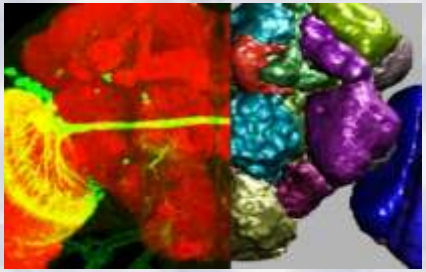

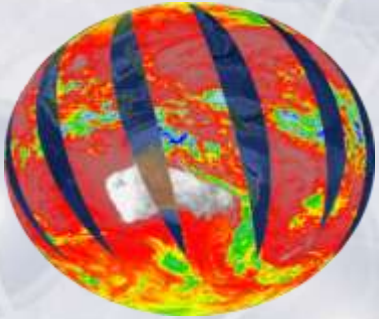

Space weather
iPIC3D

Satellite Data Processing
JURASSIC

Bio-informatics
RAY

Nuclear Fusion
Spectre, ALF, PARAFEM

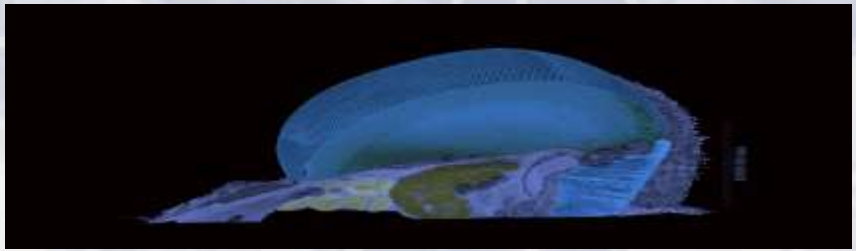

Big Data Analytics
Apache Flink Workflows

Synchrotrons
SAVU

Human Brain Science
NEST

SAGE

Use Cases

- **Goal**
  - Explore tools and services on top of Mero
- **Methodology**
  - “HSM” Methods to automatically move data across tiers
  - “pNFS” parallel file system access on Mero
  - Scale out Object Storage Integrity checking service provision
  - Allinea Performance Analysis Tools provision

- **Goal**
  - Explore usage of SAGE by programming models, runtimes and data analytics solutions
- **Methodology**
  - Usage of SAGE through MPI and PGAS
    - Adapt MPI-IO for SAGE
    - Adapt PGAS for SAGE
  - Runtimes for SAGE
    - Pre/Post Processing
  - Volume Rendering
    - Exploit Caching hierarchy
  - Data Analytics methods on top of Clovis
    - Apache Flink over Clovis, looking beyond Hadoop
    - Exploit NVRAM as extension of Memory

SAGE

WP4: Programming Models and Analytics

- **Goal**
  - Hardware definition, integration and demonstration
- **Methodology**
  - Design and Bring-up of SAGE hardware
    - Seagate Hardware
    - Atos Hardware
  - Integration of all the software components
  - Integration in Juelich Supercomputer Center(JSC)
  - Demonstrate use cases
    - Extrapolate performance to Exascale
    - Study other Object stores vis-à-vis Mero

SAGE

Integration, Demonstration

- M1 (Sept 2015) – Project Start
- M3 – Draft System architecture
- M9 – Co-Design Inputs from Apps Complete
- M12 – Design of Key Software Components (1st Draft)
- M18 – Prototype system available (Juelich)
- M27 – Design of All Software Components (2nd Draft) – In Progress
- M18 – M28 – Prototype system testing in Juelich – In Progress
- M28 – M36 – Application use cases and performance tests of the SAGE Platform in Juelich

SAGE

Overall Project Timeline

Sage is extremely well aligned to the broader goals for Europe in the area of Storage, I/O and Energy Efficiency

- M-BIO-1: Tightly coupled **Storage class memory io systems** demo
- M-BIO-3: **Multi-tiered** heterogeneous storage system demo
- M-BIO-5: **Big data analytics tools** developed for hpc use
- M-BIO-6: '**Active Storage**' capability demonstrated
- M-BIO-8: Extreme scale **multi-tier data management** tools available
- M-ARCH-3: New compute nodes and **storage architecture use nvram**
- M-ENER –X: Addresses **Energy goals** by avoiding data movements
  - 100x more energy to move data compared to compute!!
- M-ENER-FT-10: Application survival on unreliable hardware

Alignment with European Goals [ETP4HPC SRA2]

- **We will be opening up the SAGE prototype to select external users in the coming few months**
  - **Please talk to us!**
  - **info@sagestorage.eu**
  - **Twitter: @sagestorage**
- Please visit www.sagestorage.eu and download the white paper

SAGE

Additional Information