



sage 1 ¹ (sā)

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*.

Percipient **StorAGE** for Exascale Data Centric Computing

Project Summary and Overall Strategy

Sai Narasimhamurthy

Seagate Systems UK

EXDCI Prague 2016 – 10/05/2016

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



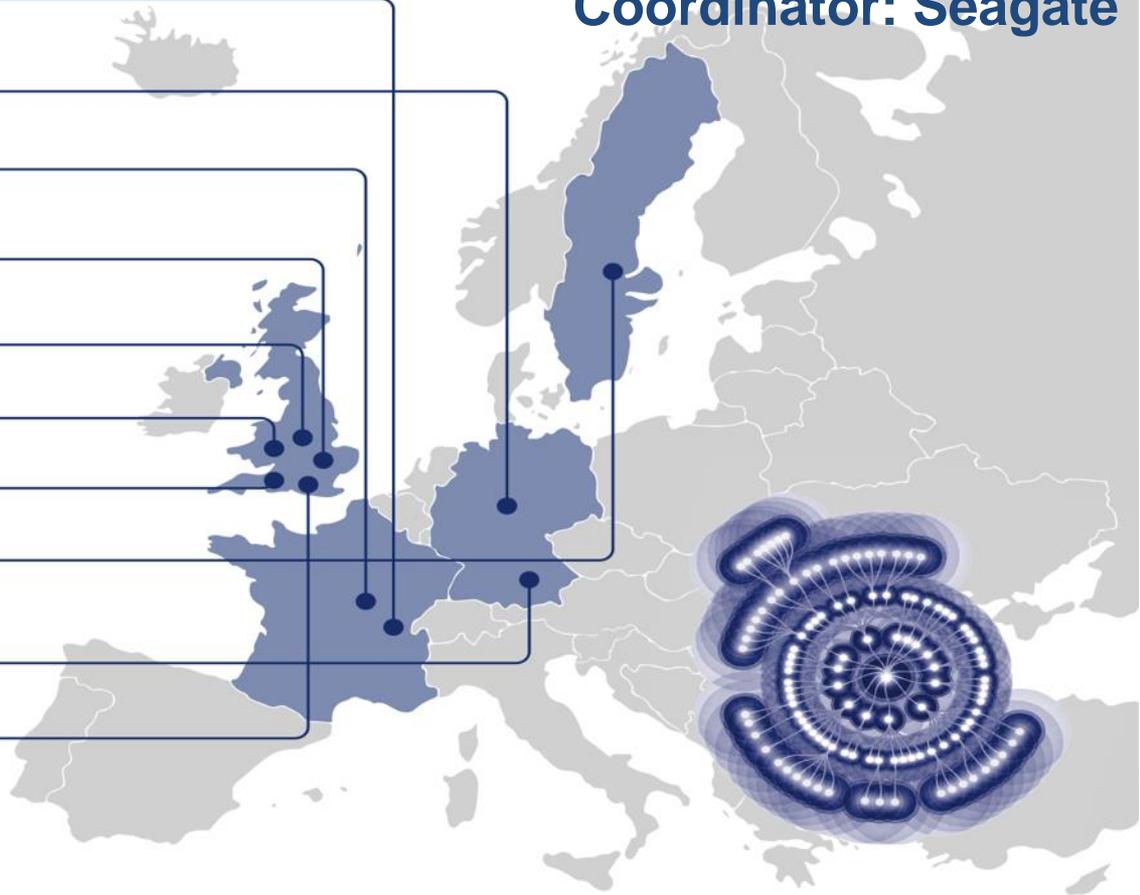


Need Exascale Data Centric Computing Systems
Big Data Extreme Computing (BDEC Systems)

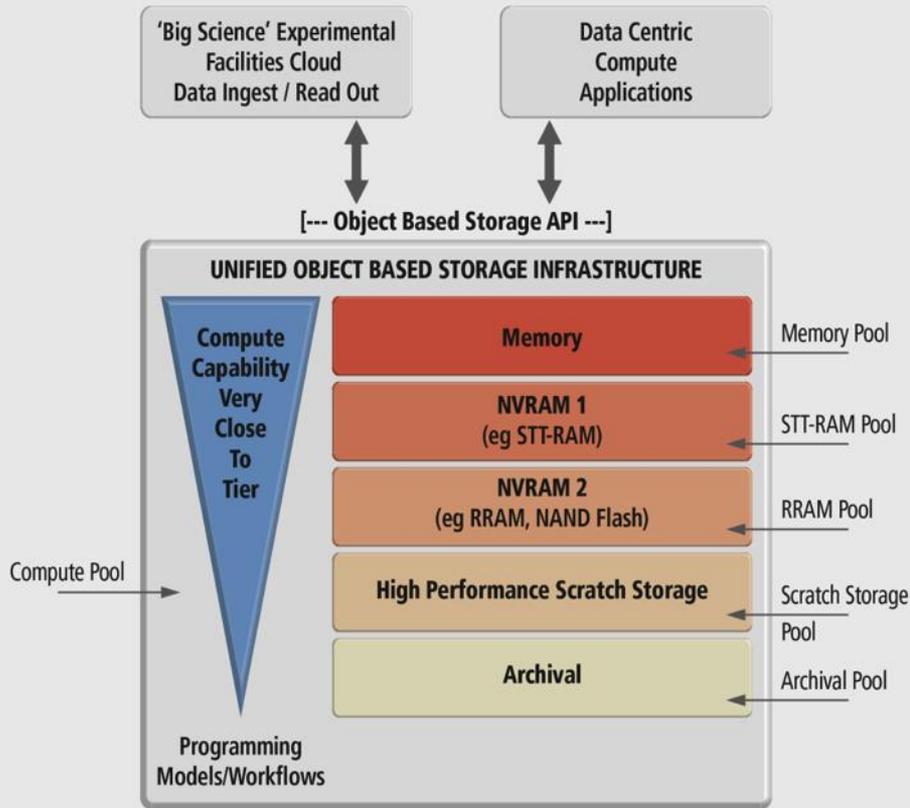
SAGE Validates a BDEC System which can Ingest,
Store, Process and Manage extreme amounts of data



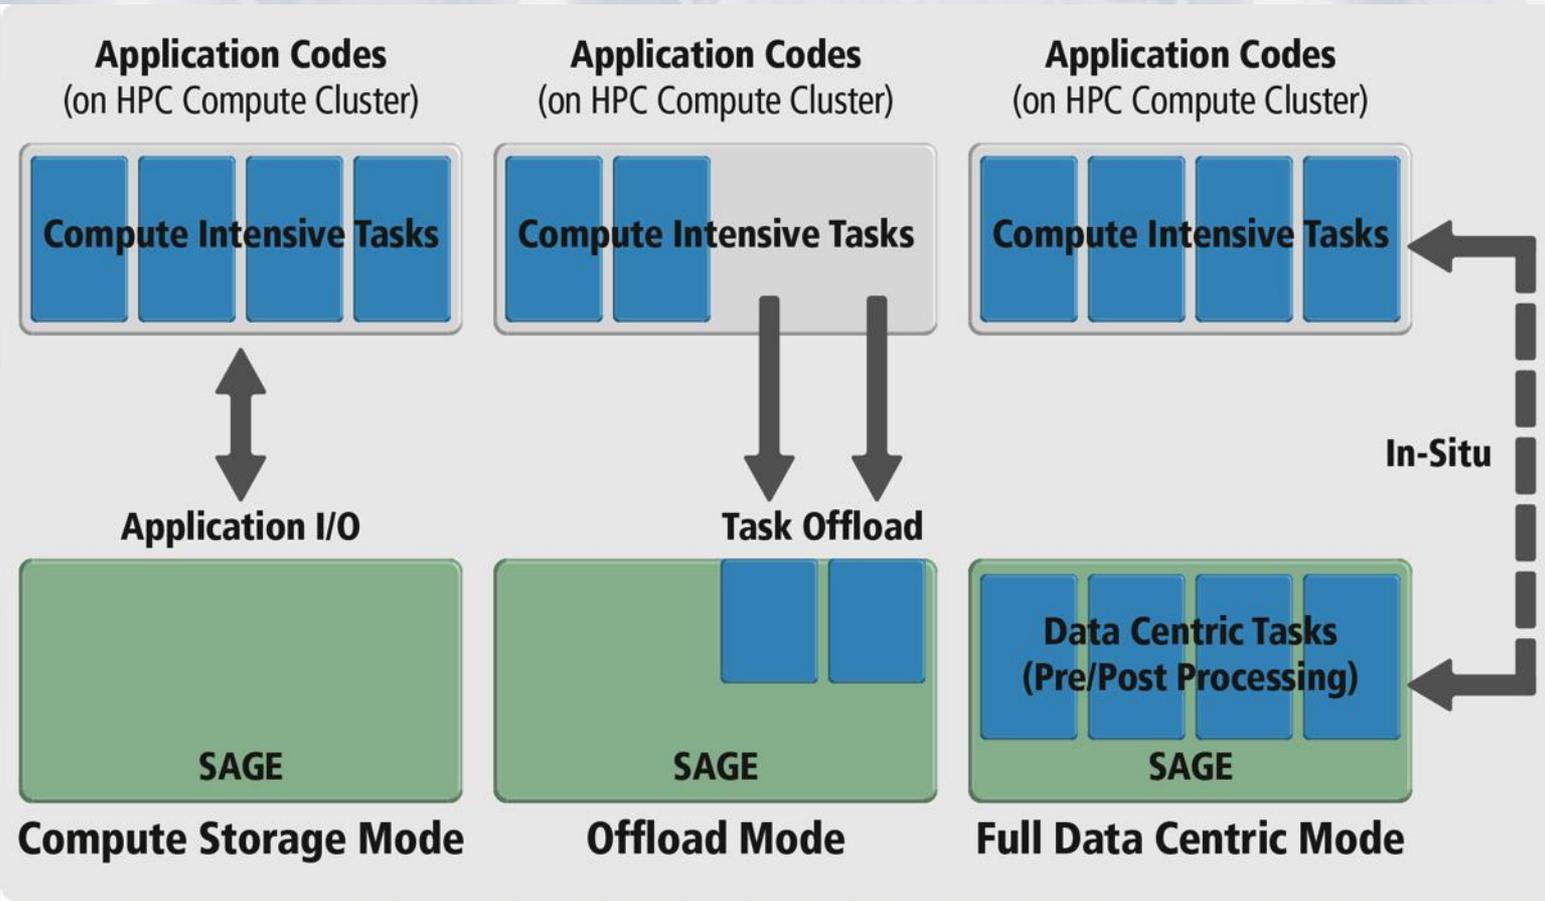
Coordinator: Seagate



Consortium

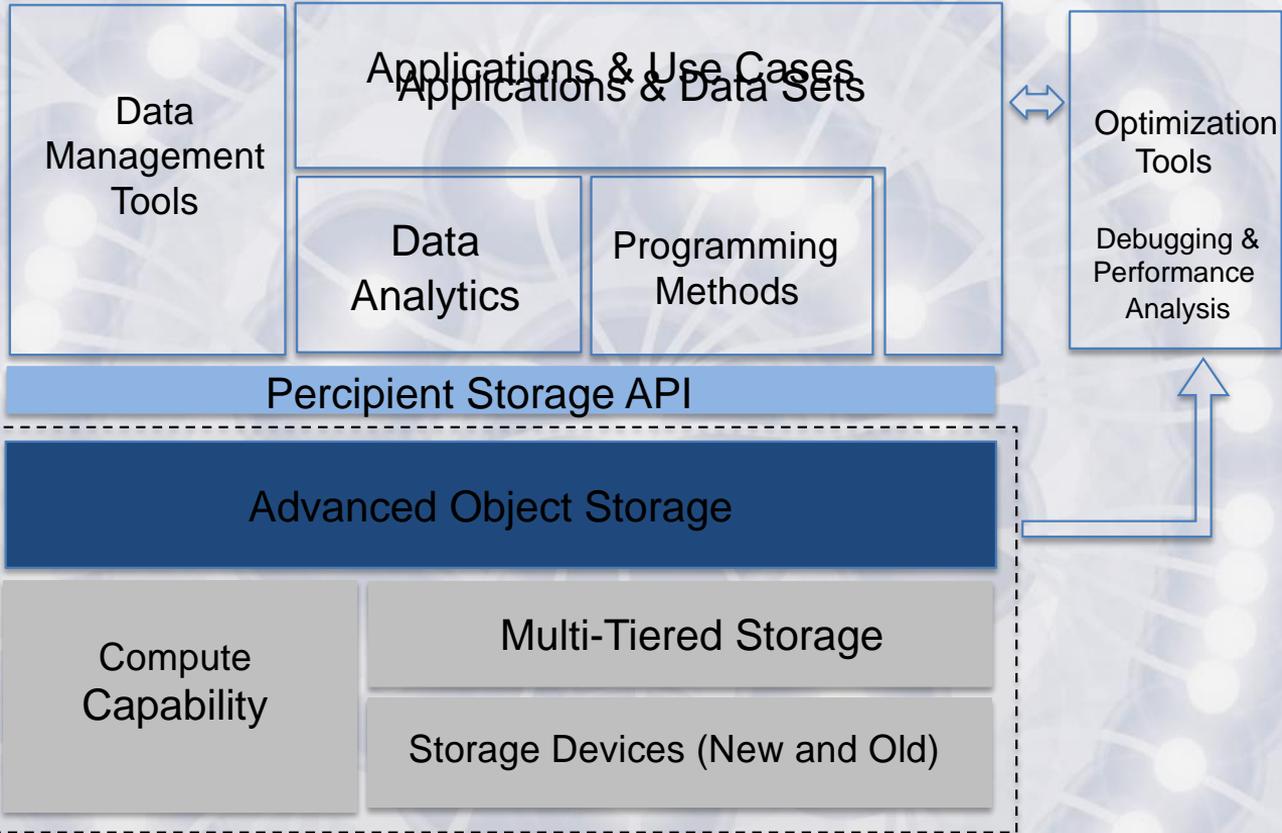


- **Goal**
 - Build the data centric computing platform
- **Methodology**
 - Commodity Server & Computing Components in I/O stack
 - **New NVRAM Technologies** in I/O stack
 - **Ability for I/O to Accept Computations**
 - Incl. Memory as part of storage tiers
 - **API** for massive data ingest and extreme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



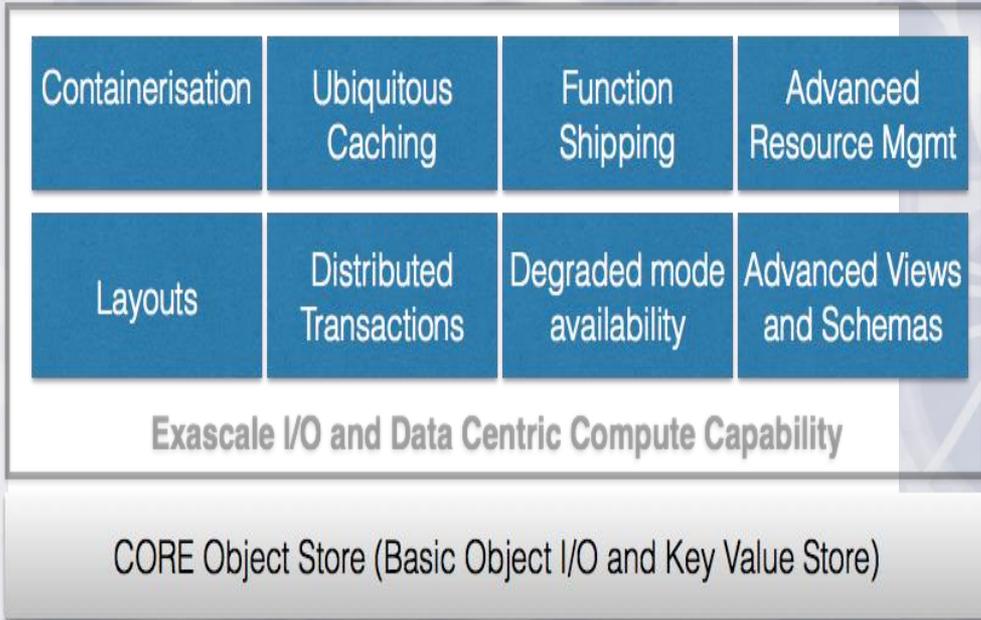
Percipient Storage Usage: Vision

SAGE Ecosystem
& Applications



Percipient Storage System Stack

Object API (Includes Legacy Support and Extensibility)

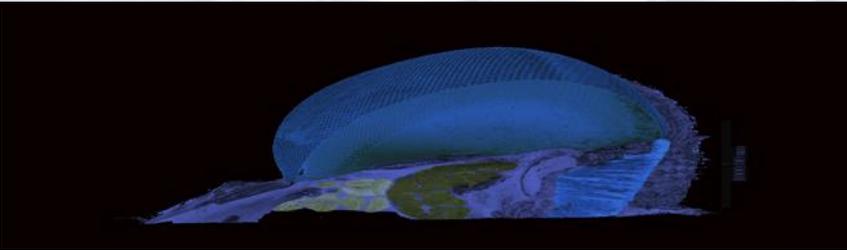
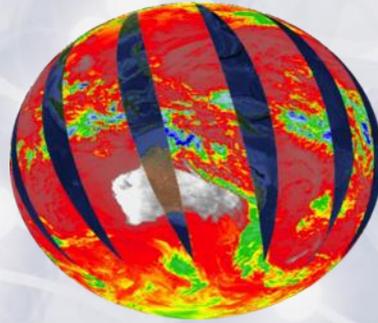
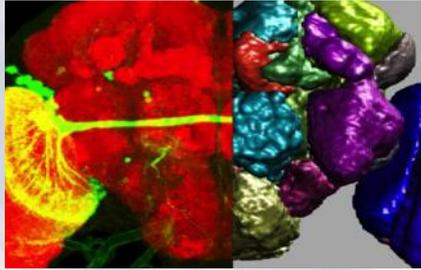


Advanced Object Store

- Extreme I/O handling
- Infrastructure fault handling
- Enable Application Resiliency
- Handle Computations from apps
- Different views to same data
- Complex data layouts & caches in I/O hierarchy



Advanced Object Storage (w/ Clovis API)



- **Overall Goal**
 - Demonstrate value with SAGE at Juelich Supercomputing Center
- **Methodology**
 - Obtain Requirements from:
 - Visualization
 - Satellite Data Processing
 - Bio-Informatics
 - Space Weather
 - Nuclear Fusion (ITER)
 - Synchrotron Experiments
 - Extrapolate needs to Exascale

ComputeStorage:
Offload

Classic HPC Apps
IPIC3D, JURASSIC,
NEST, RAY

Scientific Workflows
CCFE(IMAS, MASS,
EFITT++),Savu

Data Analytics Use Case
Flink Benchmark

Full Data Centric:

Hardware Assumptions
APIs (POSIX, Object Storage)
Function Shipping Semantics

Co-Design

Detailed Workload Description
I/O Interfaces/Libraries
Data/Metadata/Fault Tolerance
In-Storage Processing
Scalability
Data Retention

SAGE Hardware

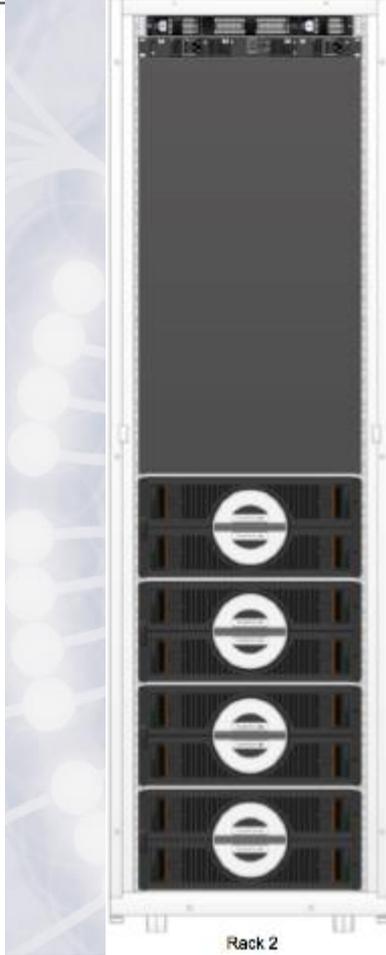
SAGE Software



Co-Design in SAGE



Rack 1



Rack 2

Tier	H/W Tier	Storage Technologies	Storage Enclosure
Local compute/GPU/ visualization (Not part of the core SAGE storage system)	0	N/A	 Bull RX
Memory	1a	RAM	 Bull Bullion S
NVRAM 1	1b	NVRAM (NVDIMM / new tech)	
NVRAM 2	1c	NVRAM - FLASH NVMe	 Seagate OneStor 2U24
	2	NVRAM - FLASH 2.5" SSD	
Scratch or Archive (optional)	3 / 4	Tier 3: Fast 3.5" SAS disk Tier 4: Cost optimised, 3.5" SATA disk	 Seagate OneStor 5U84



Validation System

- Application Analysis, Co-Design Inputs & Platform definition Complete
- Function Shipping Extensions Design Complete
- HSM Tools Design Complete
- Integrity Checking tools Design Complete
- Data Analytics Tools Interface Design Complete
- Legacy Parallel File System Access Design Complete

- 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

- **Co-operation**
 - Links with CoE: EsiWACE
 - FP7 Project(s): EpiGRAM
 - Other International: Nuclear Fusion Community - ITER
 - Planning to explore bringing Analytics & HPC community together
- **Leveraging EXDCI**
 - Better homogenization of European HPC players towards common solutions
- **Role in Extreme scale demonstrators**
 - We expect SAGE to be the foundation for at least 1 ESD – where Data Centric Operational characteristics are critical!

Questions ?

sai.narasimhamurthy@seagate.com

