



Percipient storage: A Big Data and Extreme Compute Storage Architecture



per·cip·i·ent (pr-sp-nt)

adj.

Having the power of perceiving, especially perceiving keenly and readily.


n.

One that perceives.

Objectives of the talk

- Discuss Requirements: Storage systems suited for Exascale & Data Intensive Computing
- Discuss A High Level Architecture

Future



Exascale Data Centric Computing
(**B**ig **D**ata **E**xtr**E**m**E** **C**omputing, or *BDEC*)

We cannot just build storage for Exaflop scaled processing we must accommodate the Ingest, Store, Process and Management of extreme amounts of data

Data Centric Computing


A Perspective

Science Paradigms

- Thousand years ago:
science was **empirical**
describing natural phenomena
- Last few hundred years:
theoretical branch
using models, generalizations
- Last few decades:
a **computational** branch
simulating complex phenomena
- Today: **data exploration** (eScience)
unify theory, experiment, and simulation
 - Data captured by instruments
or generated by simulator
 - Processed by software
 - Information/knowledge stored in computer
 - Scientist analyzes database/files
using data management and statistics



$$\left(\frac{\dot{a}}{a}\right)^2 = \frac{4\pi G\rho}{3} - K\frac{c^2}{a^2}$$




Data Centric Computing : Big Science



High Energy Physics
1 PB of raw data/ sec




Biology : 8PB for 100M Neurons



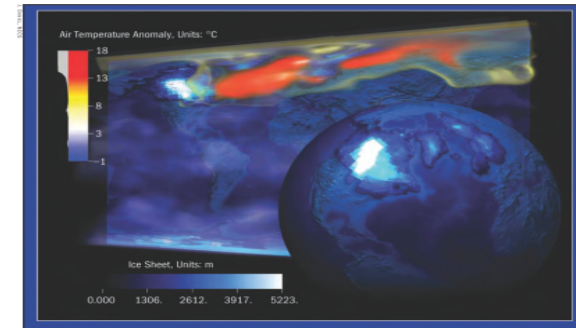
Medicine : Exascale to tackle mysteries of Protein folding



Astronomy
SKA:100PB - 3EB/yr



Industrial (eg:Oil and gas)
Visualizations and Volume rendering



Climate change effects

Data Centric Computing

HPDA(High Performance Data Analytics)

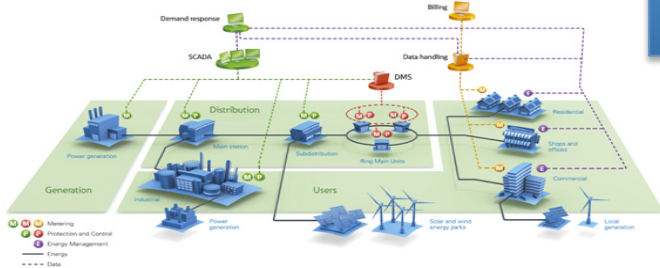


Automated Fraud Detection



Social Network Data Analysis

HPC for Big Data Analysis



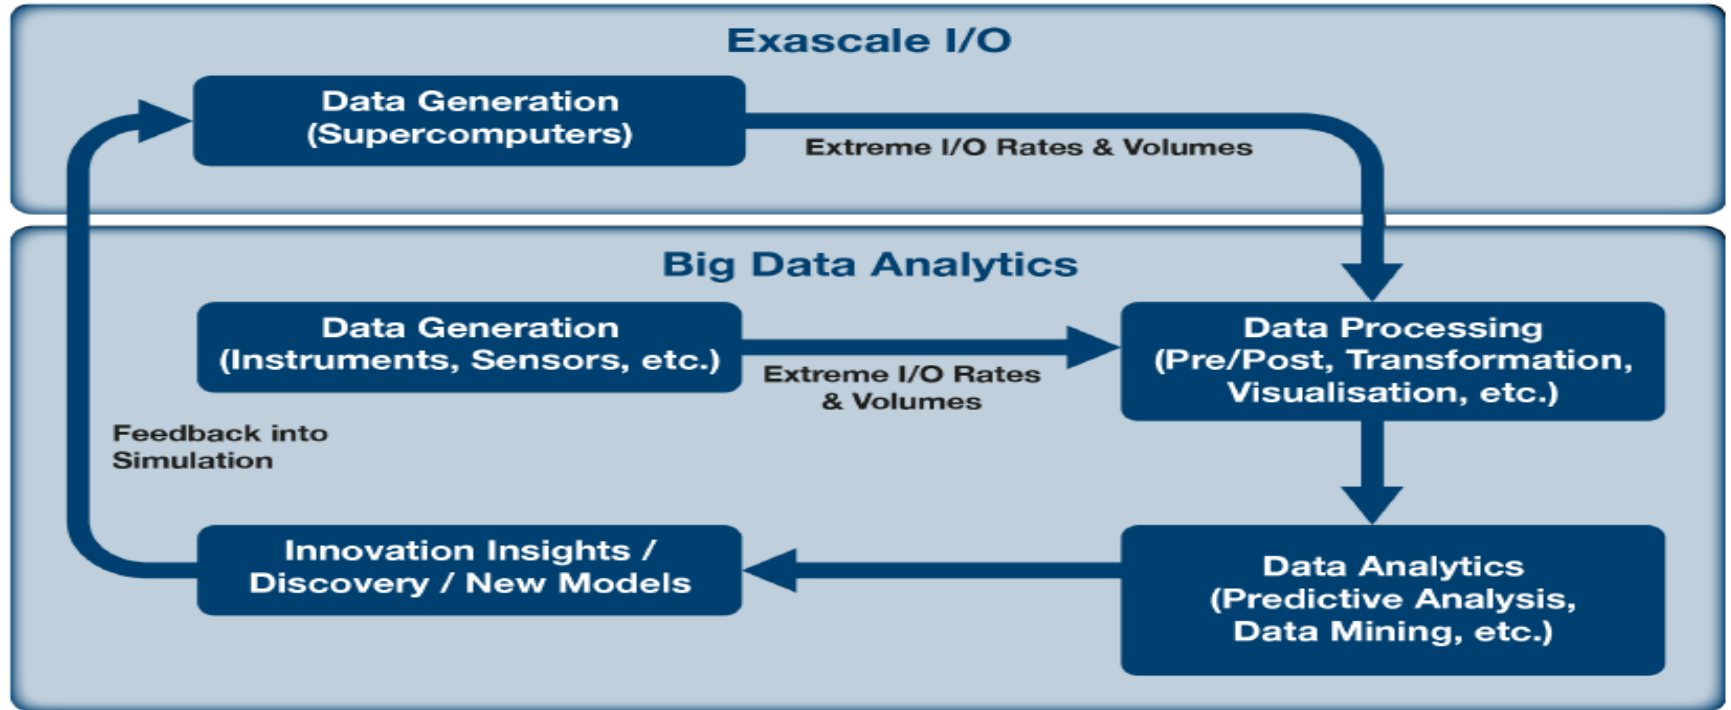
Smart Grid Data



Sensor Network Data

Knowledge Discovery in Data Centric Computing

The Links between Exascale I/O and Big Data Analytics



Key Solution Requirements

Dimensions of Needed Innovation*

Data Storage
Innovation



Data & System
Management
Innovation


Data Processing
and Analytics
Innovation

A “Percipient” Storage & I/O Oriented Solution


* *Based on Research Priorities for Storage & I/O identified by the Strategic Research Agenda(SRA) for Europe being continuously developed by the European Technology Platform for HPC (ETP4HPC)*

The Quality of “Percipience”




The Old Paradigm
'Isolated' Storage &
Computing

“PERCIPIENCE”


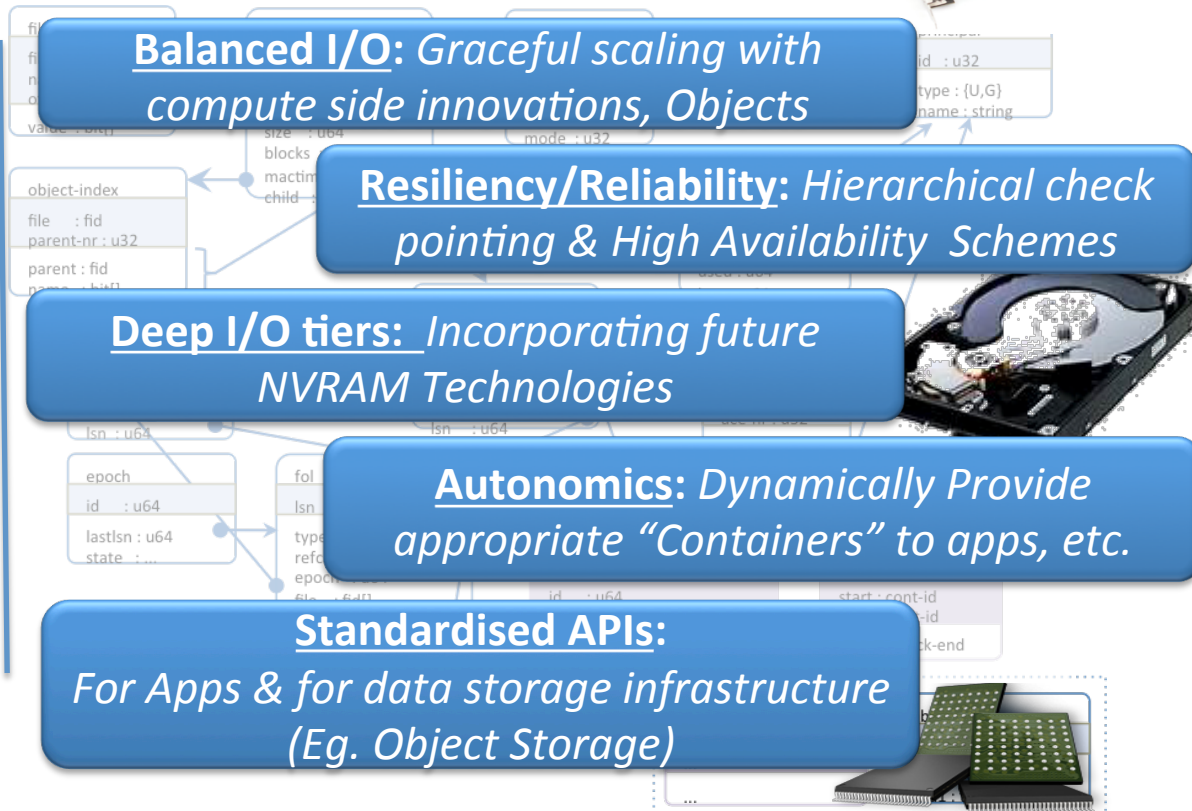


The New Paradigm

**Having a homogenous view of
data anywhere/everywhere**

Data Storage



Data Storage Innovation



Data Processing and Analytics

Data Processing
and Analytics
Innovation

Extreme Data Processing: *incl. Big Science Experiments, need advanced Big Data Analytics Methodologies*

Active Storage: *Storage Systems with compute capabilities*

On the Fly Data Manipulation:
*In-situ data processing in the I/O stack/
infrastructure (eg: Network)*

Energy Reduction Strategies:
*Computation takes less energy than data
movements*

System & Data Management

System & Data
Management
Innovation

SYSTEM

Infrastructure Telemetry, Analytics

Infrastructure Simulation

HSM/ILM: Storage Management with
many Layers

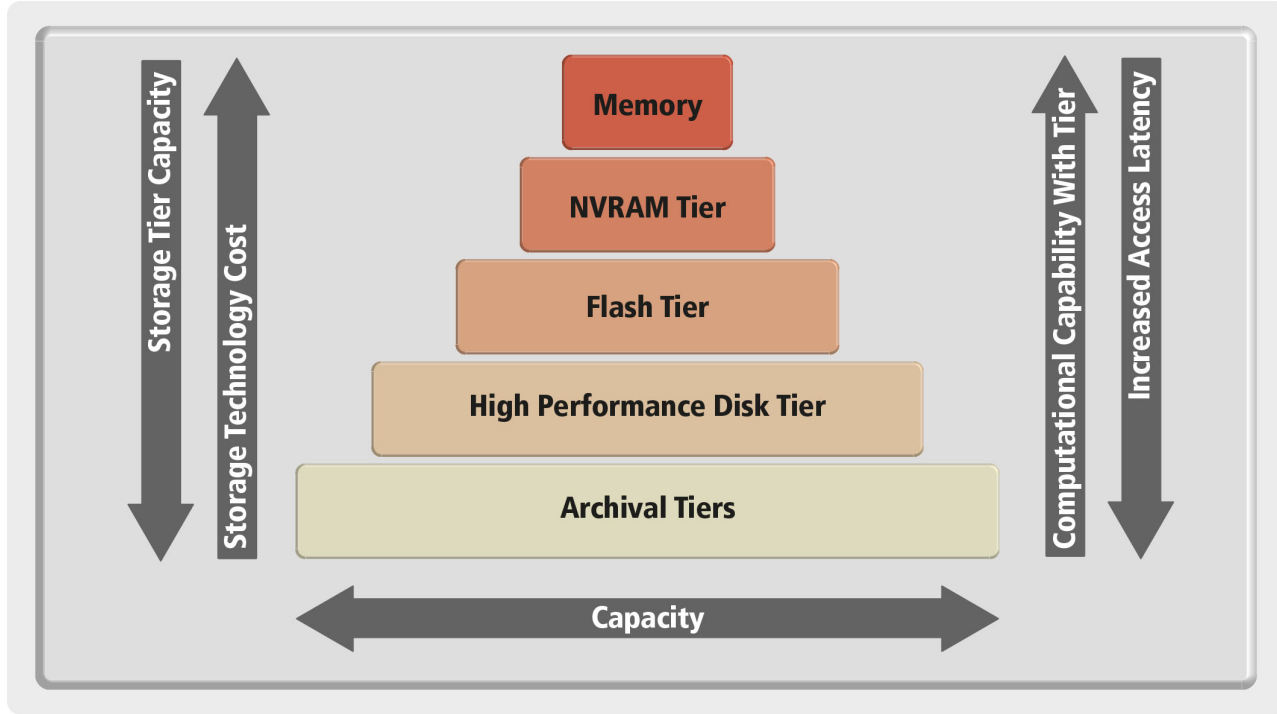
DATA

Data Layouts/Transformation:
*Different "Views" to the same data
(eg: POSIX, HDF5, NetCDF)*

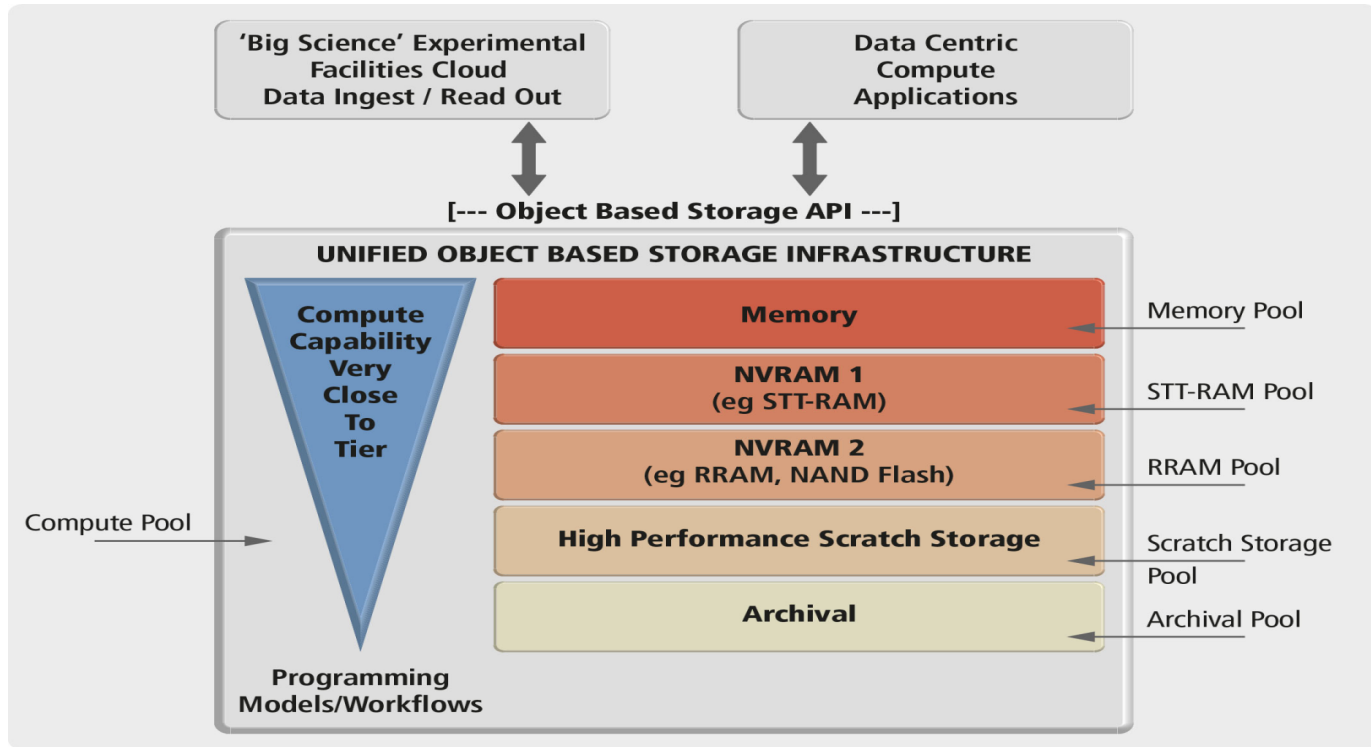
Data Indexing/Provenance Tracking/ILM

Plug-in Frameworks for adding data
management "apps"

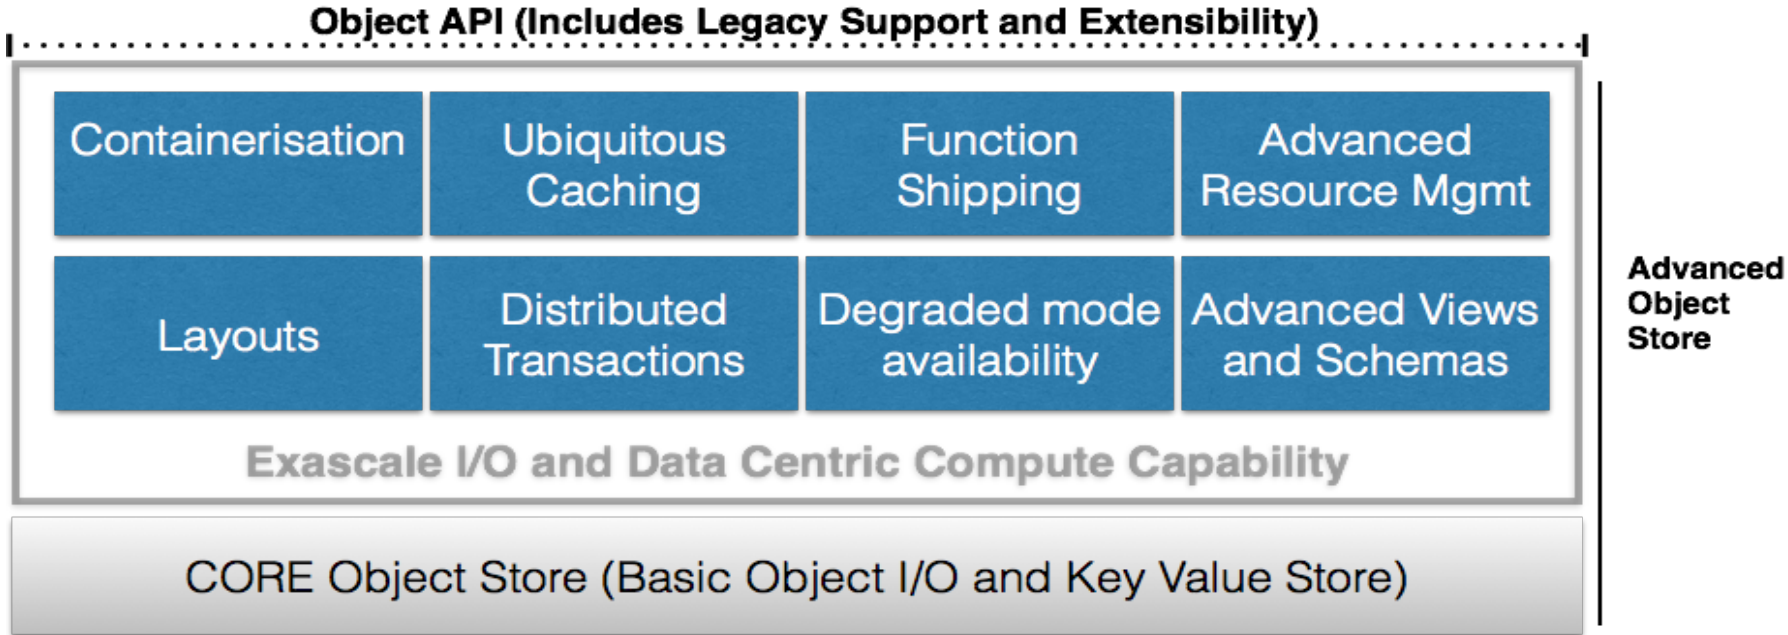
A Percipient Storage Architecture



A Percipient Storage Architecture

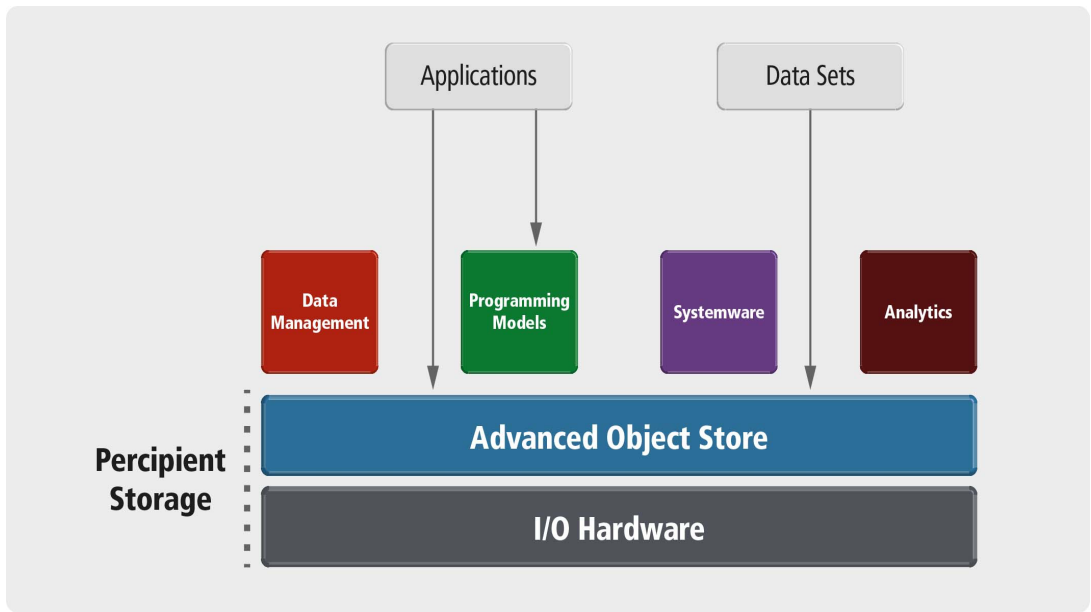


Advanced Object Store in Percipient Storage



Percipient Storage Ecosystem (examples)

Leveraging the API



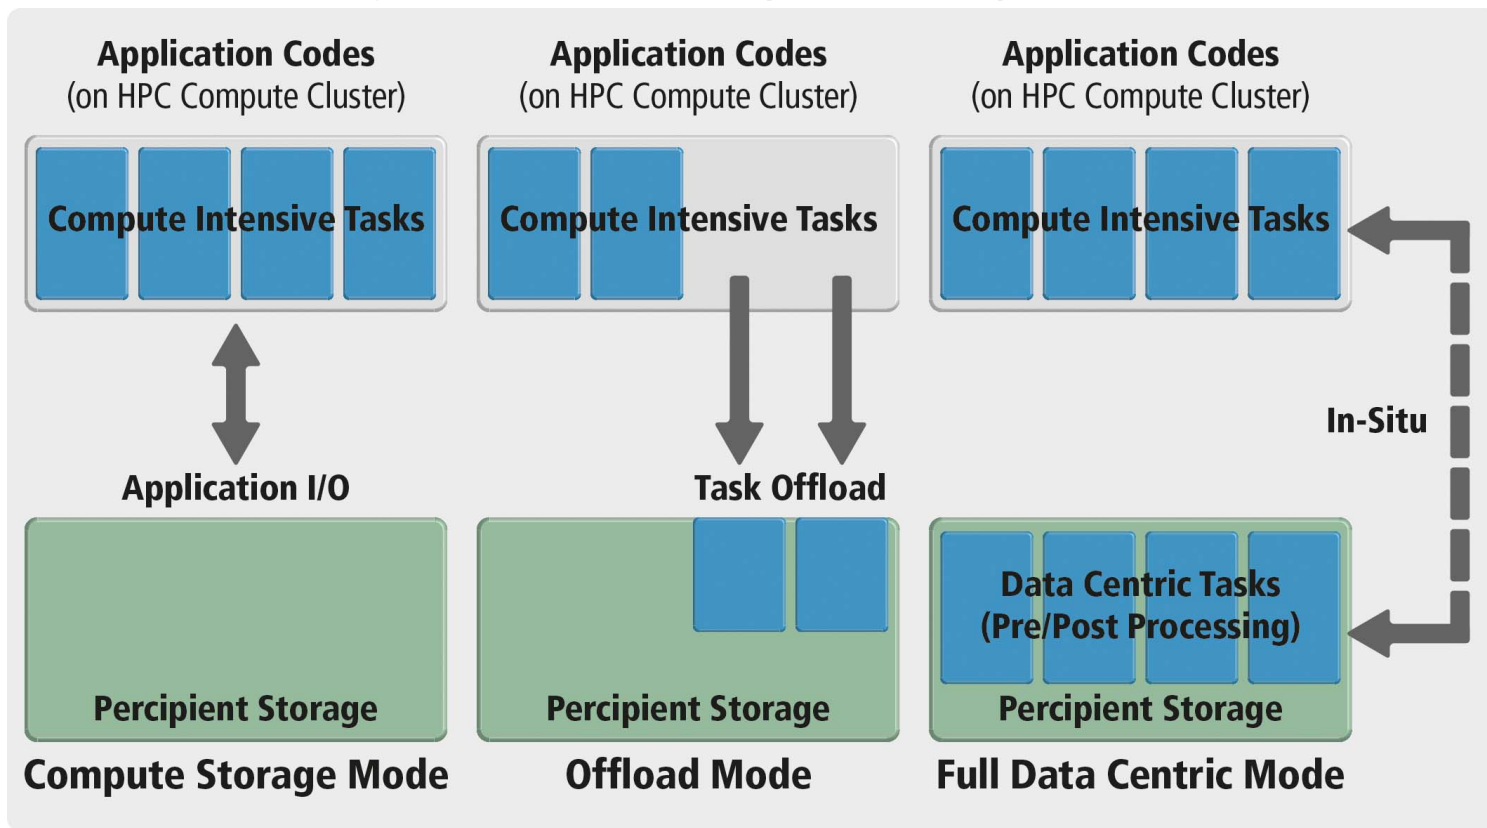
Deep HSM, Data Integrity Check, pNFS interface, etc

MPI for Data Centric Compute
PGAS

Debugging, Performance Analysis

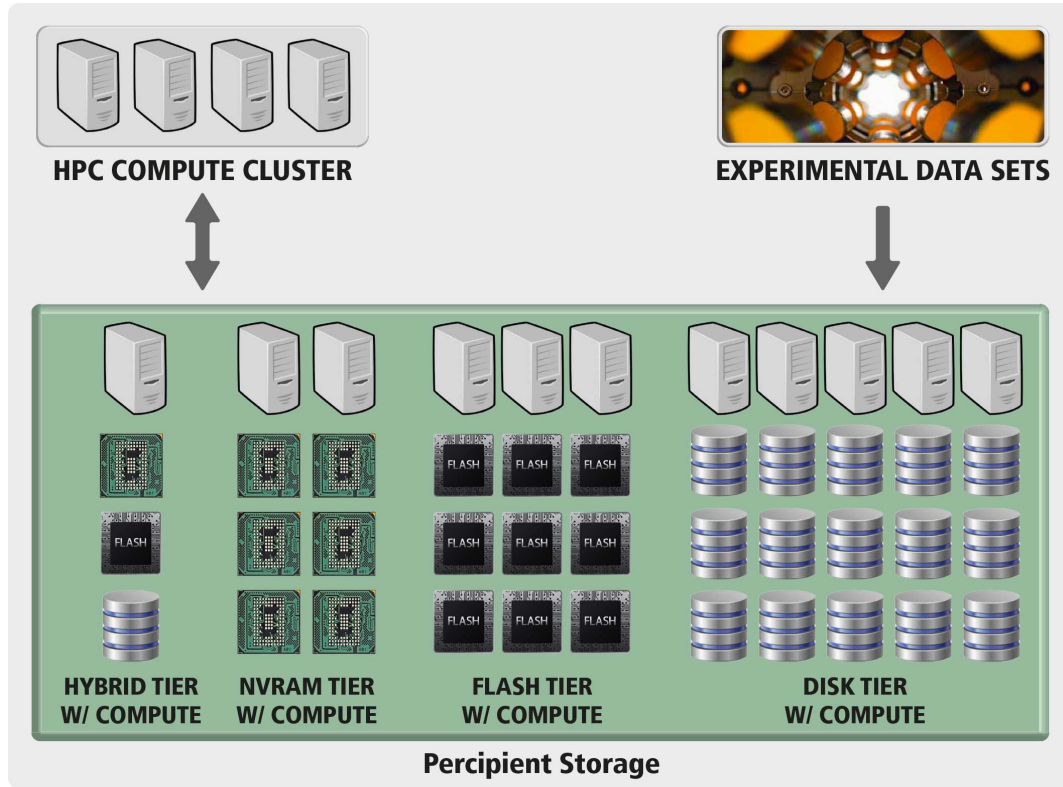
Apache Flink

Percipient Storage Usage Modes



SEAGATE

Percipient Storage Implementation



Commodity Hardware for front end servers
(eg: Embedded Server Modules)

Easily Available Disk Drive/Flash Technology

NVRAM Techs as and when they become available

Easy interfacing with HPC clusters



sai.narasimhamurthy@seagate.com