TO FS OR NOT TO FS…

ROB ROSS, PHILIP CARNS, MATTHIEU DORIER, KEVIN HARMS, ROB LATHAM, AND SHANE SNYDER	Argonne National Laboratory
GARTH GIBSON, GEORGE AMVROSIADIS, CHUCK CRANOR, AND QING ZHENG	Carnegie Mellon University
JEROME SOUMAGNE	The HDF Group
GALEN SHIPMAN AND BRAD SETTLEMYER	Los Alamos National Laboratory

OUR ORGANIZERS HAVE IT RIGHT!

- How do the particular I/O **use cases** [**inform**] the way we manage data?
- How should we **present hierarchical storage** systems to user applications, …?
- How should we **manage data movement** through a storage hierarchy …?

SPECIALIZATION OF DATA SERVICES

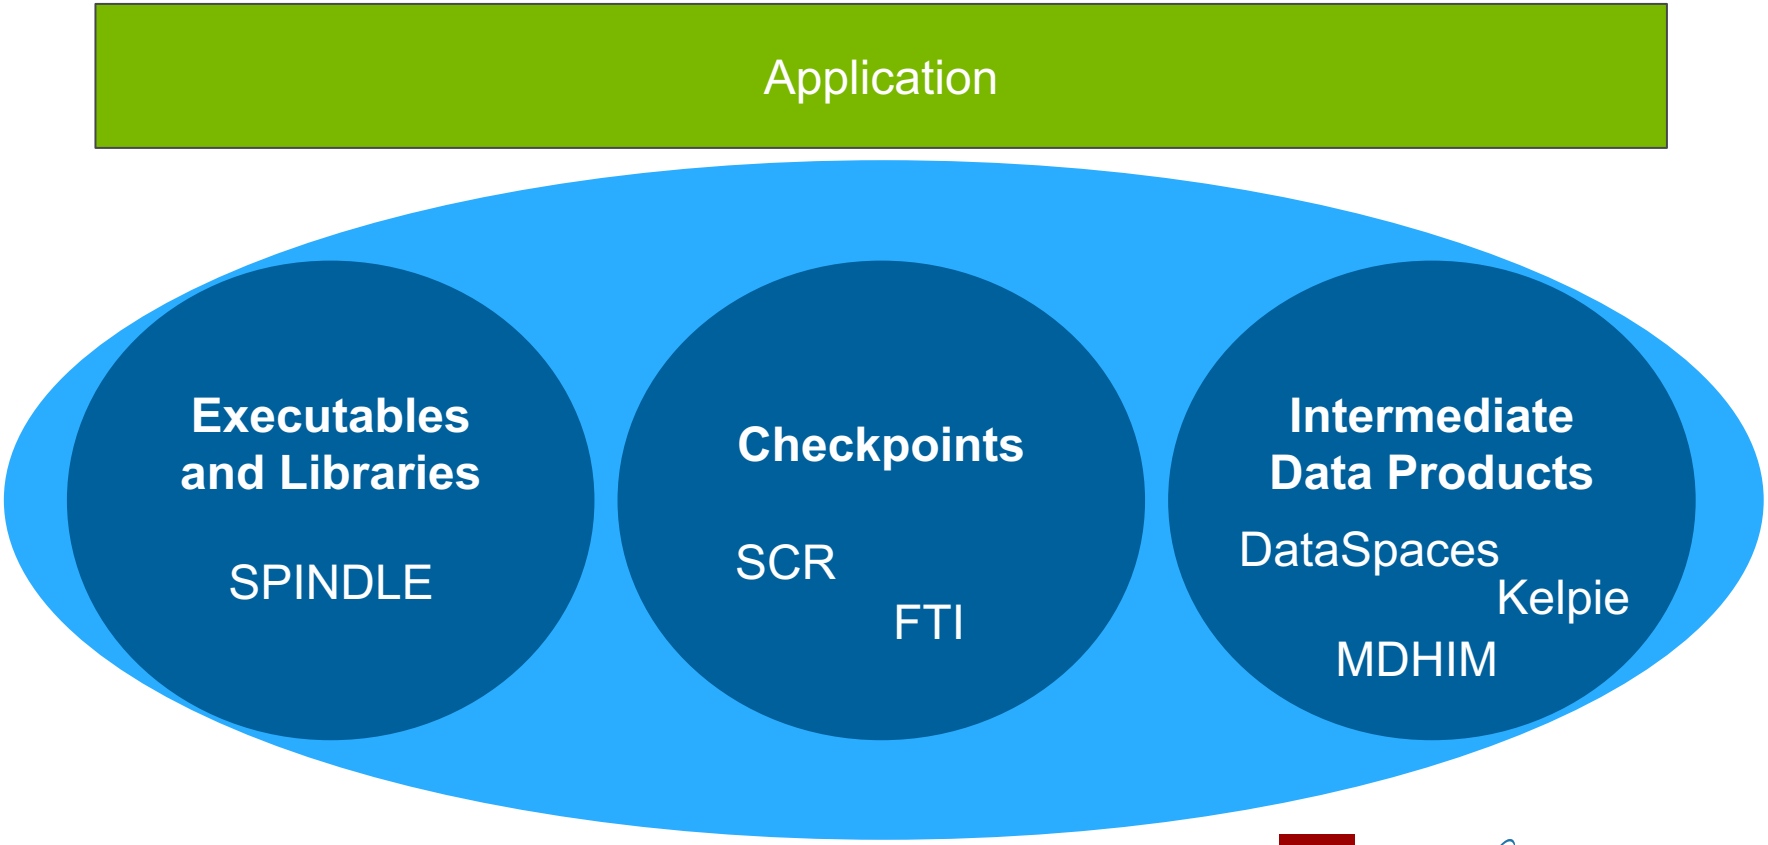

SPECIALIZATION OF DATA SERVICES

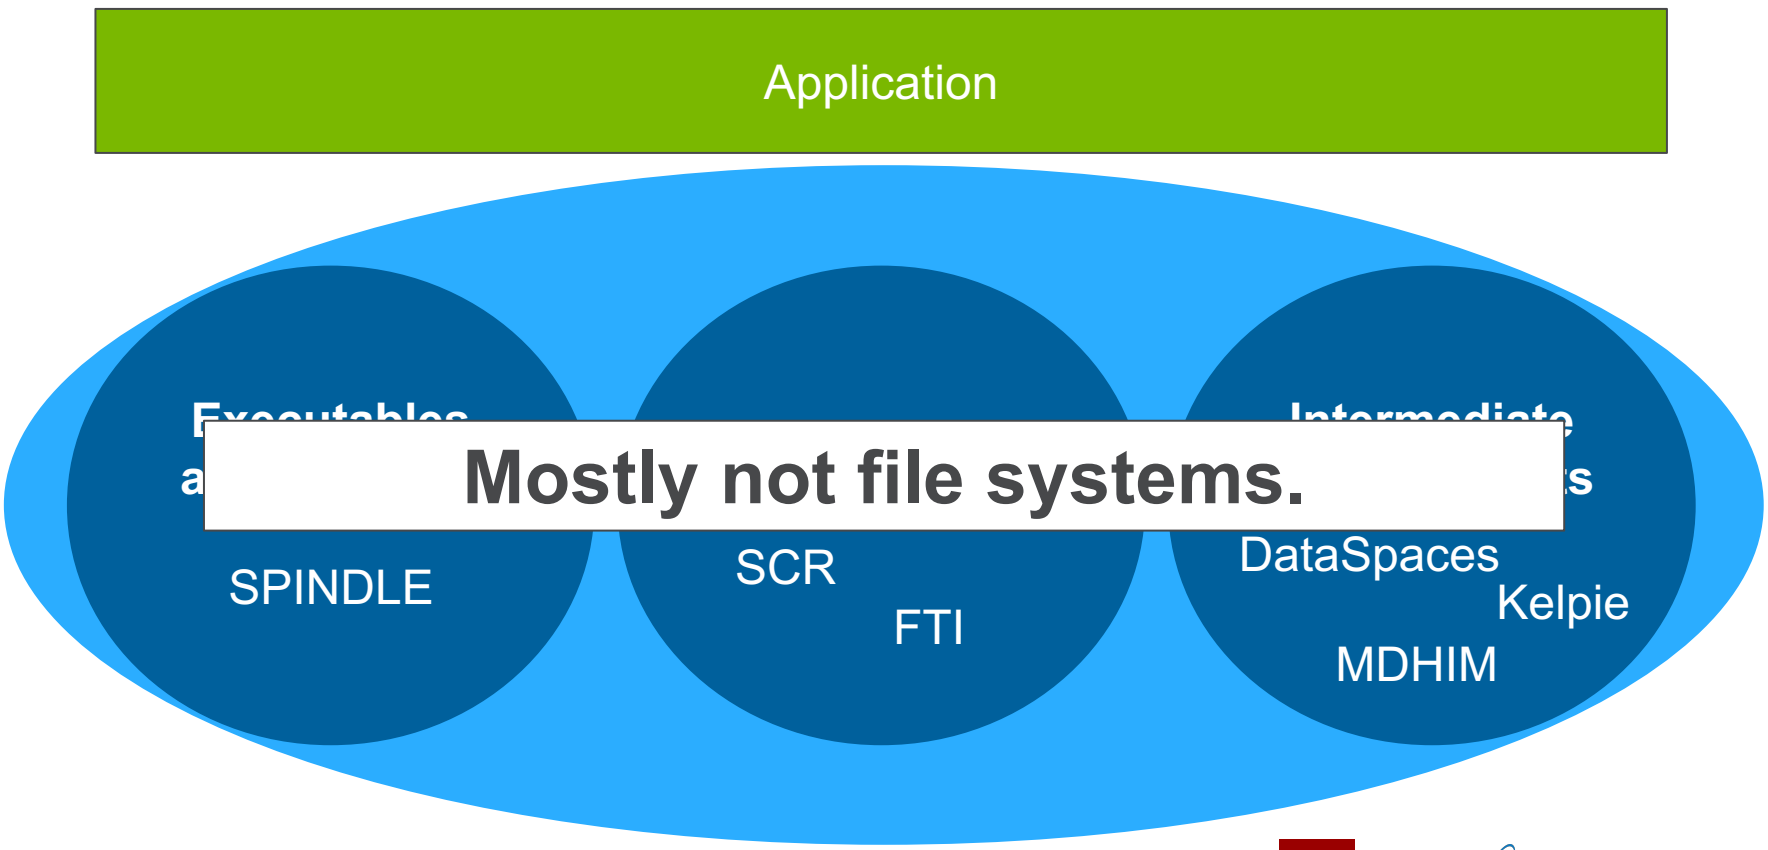

	Provisioning	Comm.	Local Storage	Fault Mgmt. and Group Membership	Security
ADLB *Data store and pub/sub.*	MPI ranks	MPI	RAM	N/A	N/A
DataSpaces *Data store and pub/sub.*	Indep. job	Dart	RAM (SSD)	Under devel.	N/A
DataWarp *Burst Buffer mgmt.*	Admin./ sched.	DVS/ Inet	XFS, SSD	Ext. monitor	Kernel, Inet
FTI *Checkpoint/restart mgmt.*	MPI ranks	MPI	RAM, SSD	N/A	N/A
Kelpie *Dist. in-mem. key/val store*	MPI ranks	Nessie	RAM (Object)	N/A	Obfusc. IDs
SPINDLE *Exec. and library mgmt.*	Launch MON	TCP	RAMdisk	N/A	Shared secret

ASSERTION: WE SHOULD BUILD AN ECOSYSTEM OF DATA SERVICES

- **Many components can be shared** across multiple services
- Some services will look like file systems, others not, that's ok
- **Need to tackle the hard problems:**
  - Group membership
  - Authentication/authorization
  - Good pub/sub (thanks Brad!)
  - Performance (we're HPC!)
- **Enable broader community to build better, more capable user-level data services than possible today.**

A DATA SERVICE FOR MULTI-SCALE SIMULATIONS
A FAR OUT EXAMPLE

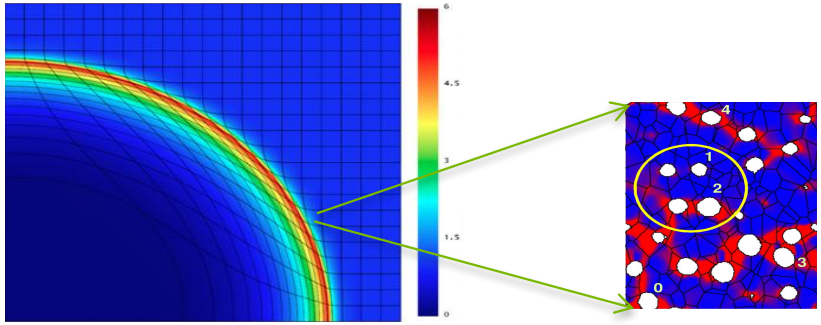

Lulesh continuum model:
- Lagrangian hydro dynamics
- Unstructured mesh

Viscoplasticity model [1]:
- FFT based PDE solver
- Structured sub-mesh

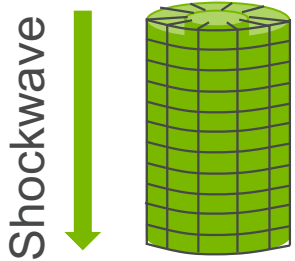

- Future applications are exploring the use of multi-scale modeling
- As an example: Loosely coupling continuum scale models with more realistic constitutive/response properties
  - e.g., Lulesh from ExMatEx
- Fine scale model results can be cached and new values interpolated from similar prior model calculations

R. Lebensohn et al, Modeling void growth in polycrystalline materials, Acta Materialia, http://dx.doi.org/10.1016/j.actamat.2013.08.004.

Carnegie Mellon University · HDF · Los Alamos National Laboratory · Argonne National Laboratory

A DATA SERVICE FOR MULTI-SCALE SIMULATIONS
A FAR OUT EXAMPLE

- Goals
  - Minimize fine scale model executions
  - Minimize query/response time
  - Load balance DB distribution
- Approach
  - Start with a key/value store
  - Distribute approx. nearest-neighbor query
  - Distribute data to co-locate values for interpolation
  - Import/export to persistent store

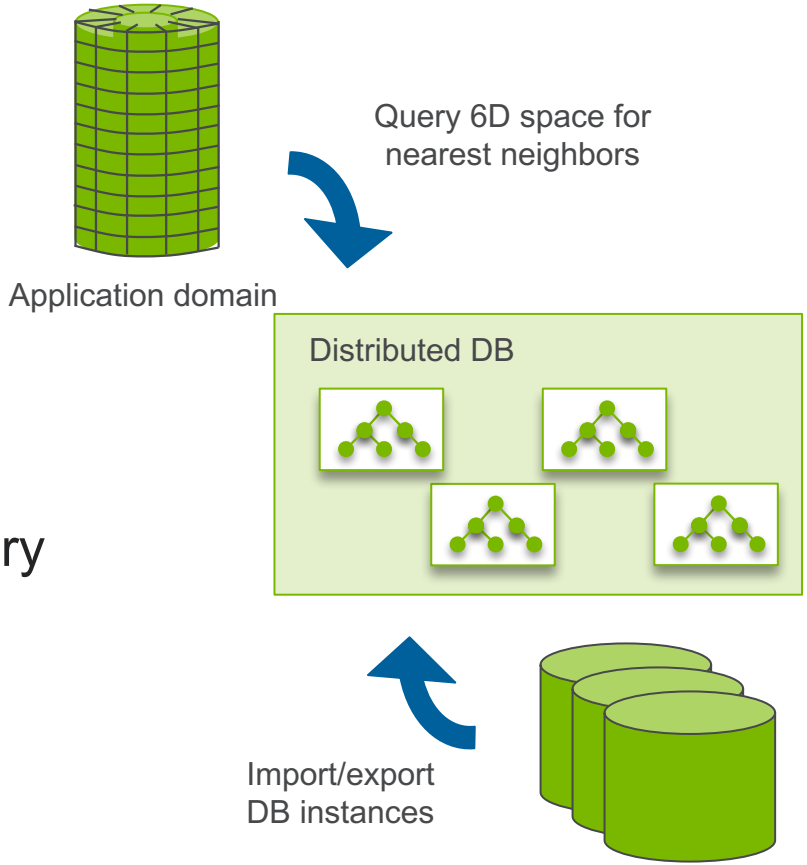

THANKS!

THIS WORK IS SUPPORTED BY THE DIRECTOR, OFFICE OF ADVANCED SCIENTIFIC COMPUTING RESEARCH, OFFICE OF SCIENCE, OF THE U.S. DEPARTMENT OF ENERGY UNDER CONTRACT NO. DE-AC02-06CH11357.

Argonne NATIONAL LABORATORY