SSIO and Programming Model/ System/Env Very Highly Related

Gary Grider

Recent SSIO Workshop Findings

- *Recent Workshop*
  - *day 1: apps people talking with SSIO and software Stack people*
  - *day 2: software stack and SSIO people using apps rqmnts to talk about cross cuts*
  - *day 3: SSIO researchers talk about needed R&D given previous 2 days input*

Programming Models/Systems/Environment related

- *In situ* data analysis is already an important component of many applications.
- New Solid state and disk storage layers are complicating the storage hierarchy.
- Scientists need a coherent view and mgmt methods of the storage resources.
- New requirements for results validation may change the role of SSIO systems
- New programming models/systems drive new persistence mechanisms
- Scientists desire increasingly complex data abstractions that improve productivity

- Current SSIO designs are hindered by their isolation from system-level resource management, monitoring, and workflow systems.
- Many important aspects of app/system SSIO behavior aren't well understood
- Community access to data on apps, systems, test environments for new technologies evaluation and bringing new talent into the community.

Near Term SSIO Research Priorities from Workshop Distillation

- SSIO architectures research for
  - managing deep and hetero storage hierarchies
  - explore alternative mgmt paradigms to the file system model
- In the area of metadata, name spaces, and provenance
  - new methods of mgmt of rich metadata
  - breaking away from the current file
- In the area of supporting science data
  - develop the next generation of I/O middleware and services to support new programming abstractions and workflows
- In the area of understanding SSIO
  - improve our ability to characterize storage activities to model and predict the behavior of SSIO activities on future systems.

Collaborative R&D PM/E, RT, SSIO

Fast Forward I/O Architecture

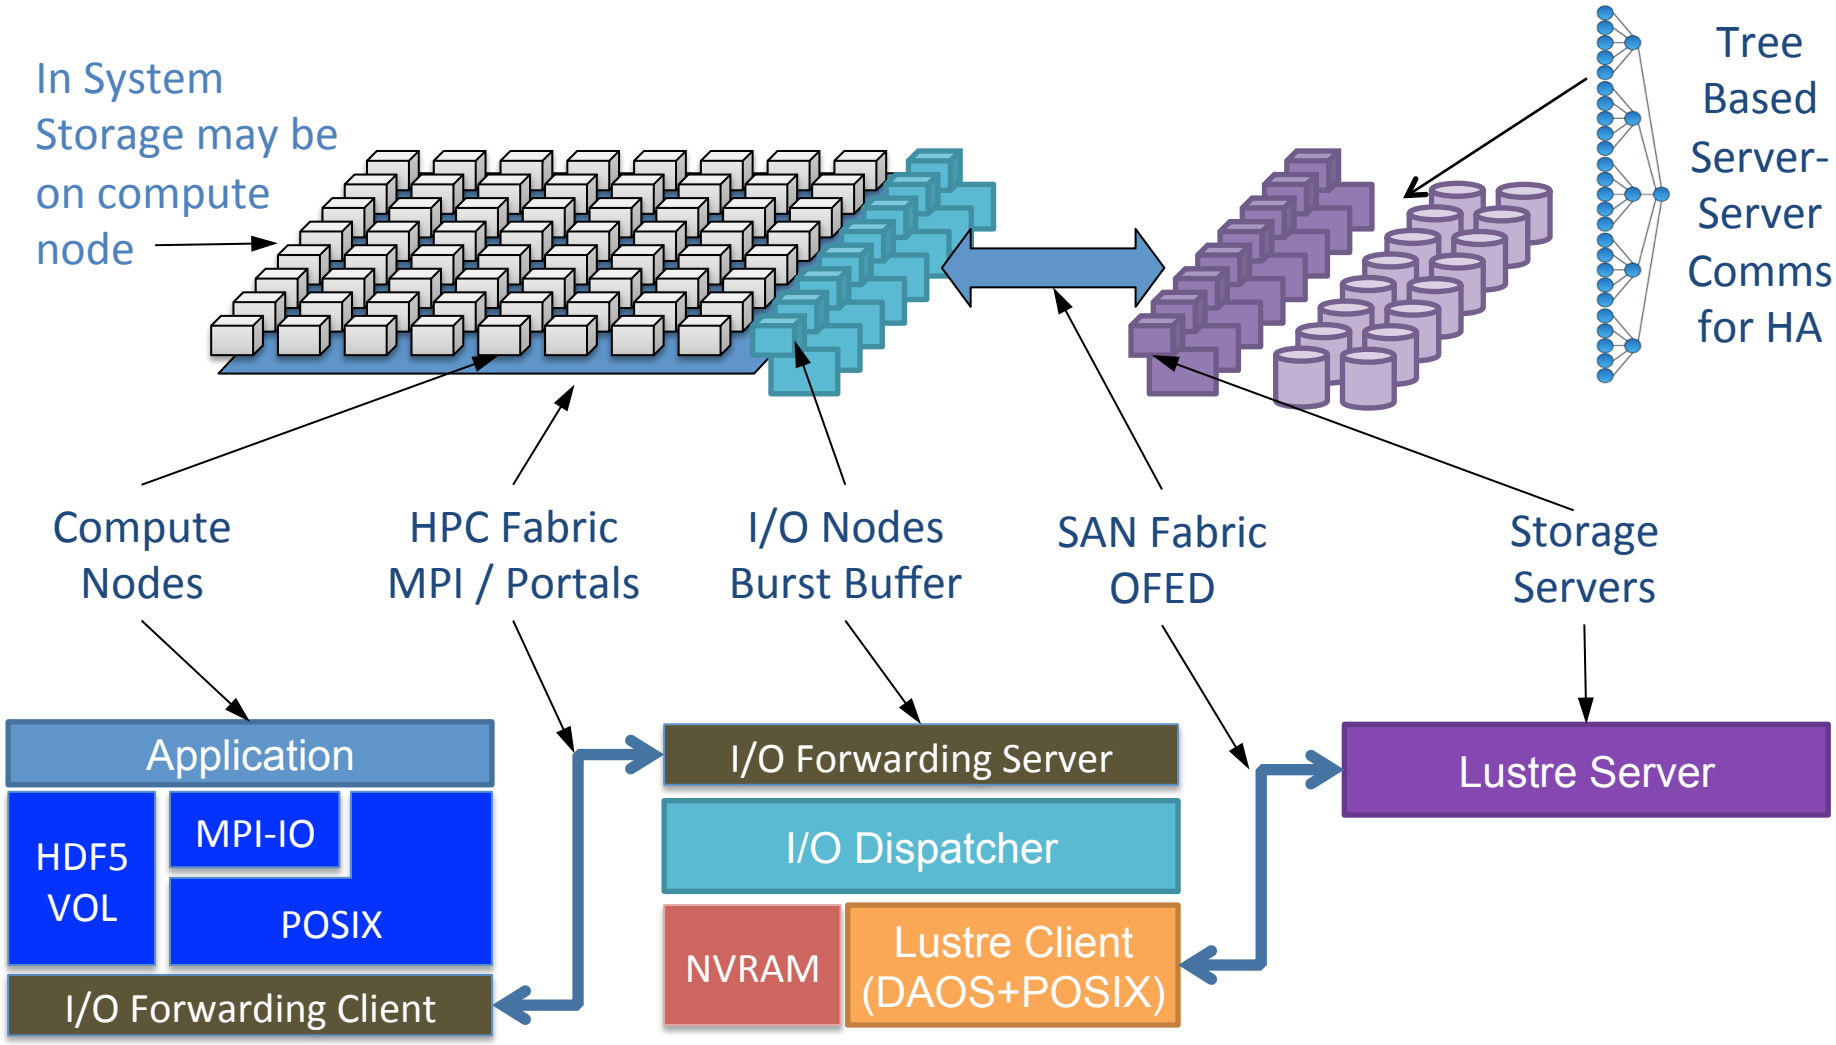

DOE Storage FFwd Accomplishments

- Tiered storage with coordination across tiers (client libraries, burst buffer mechanisms and disk mechanisms.)
- Enabling versioning capabilities at one or more tiers of storage
- Enabling multiple IO object like KVS, arrays, unstructured scientific data, etc. to be stored/understood by the data management stack
- Transactional APIs from to to bottom to provide clean interaction with failure and prompt recovery with defined consistency
- Server/burst buffer side collectives for operating on/organizing data
- Server side scalable failure detection and remediation
- Active burst buffer/object storage capabilities (shipping analysis functions to be run by parts of the data management stack
- Rudimentary serial organization/indexing of data by the DM stack
- End-to-end data integrity capabilities in the data management stack.

Phase 2 of Fast Forward – Hook Up Apps

- Mostly concentrate on trying object/array/kvs/blob, async, transactional, versioning API to allow apps to utilize tiered storage (in and off system)

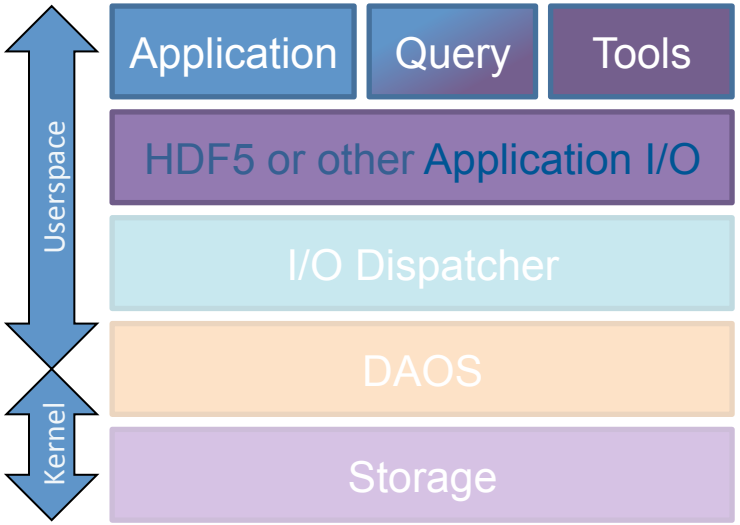

- Want to try different apps/levels
  - Legion persistent pool with triggers for scheduled "make persistent" tasks
  - Maybe an async AMR mini-app/code
  - Maybe Argonne GVR (global view resilience)

Collaborative R&D PM/E, RT, SSIO

Please Talk to Your SSIO Folks When Prototyping (example at LANL)

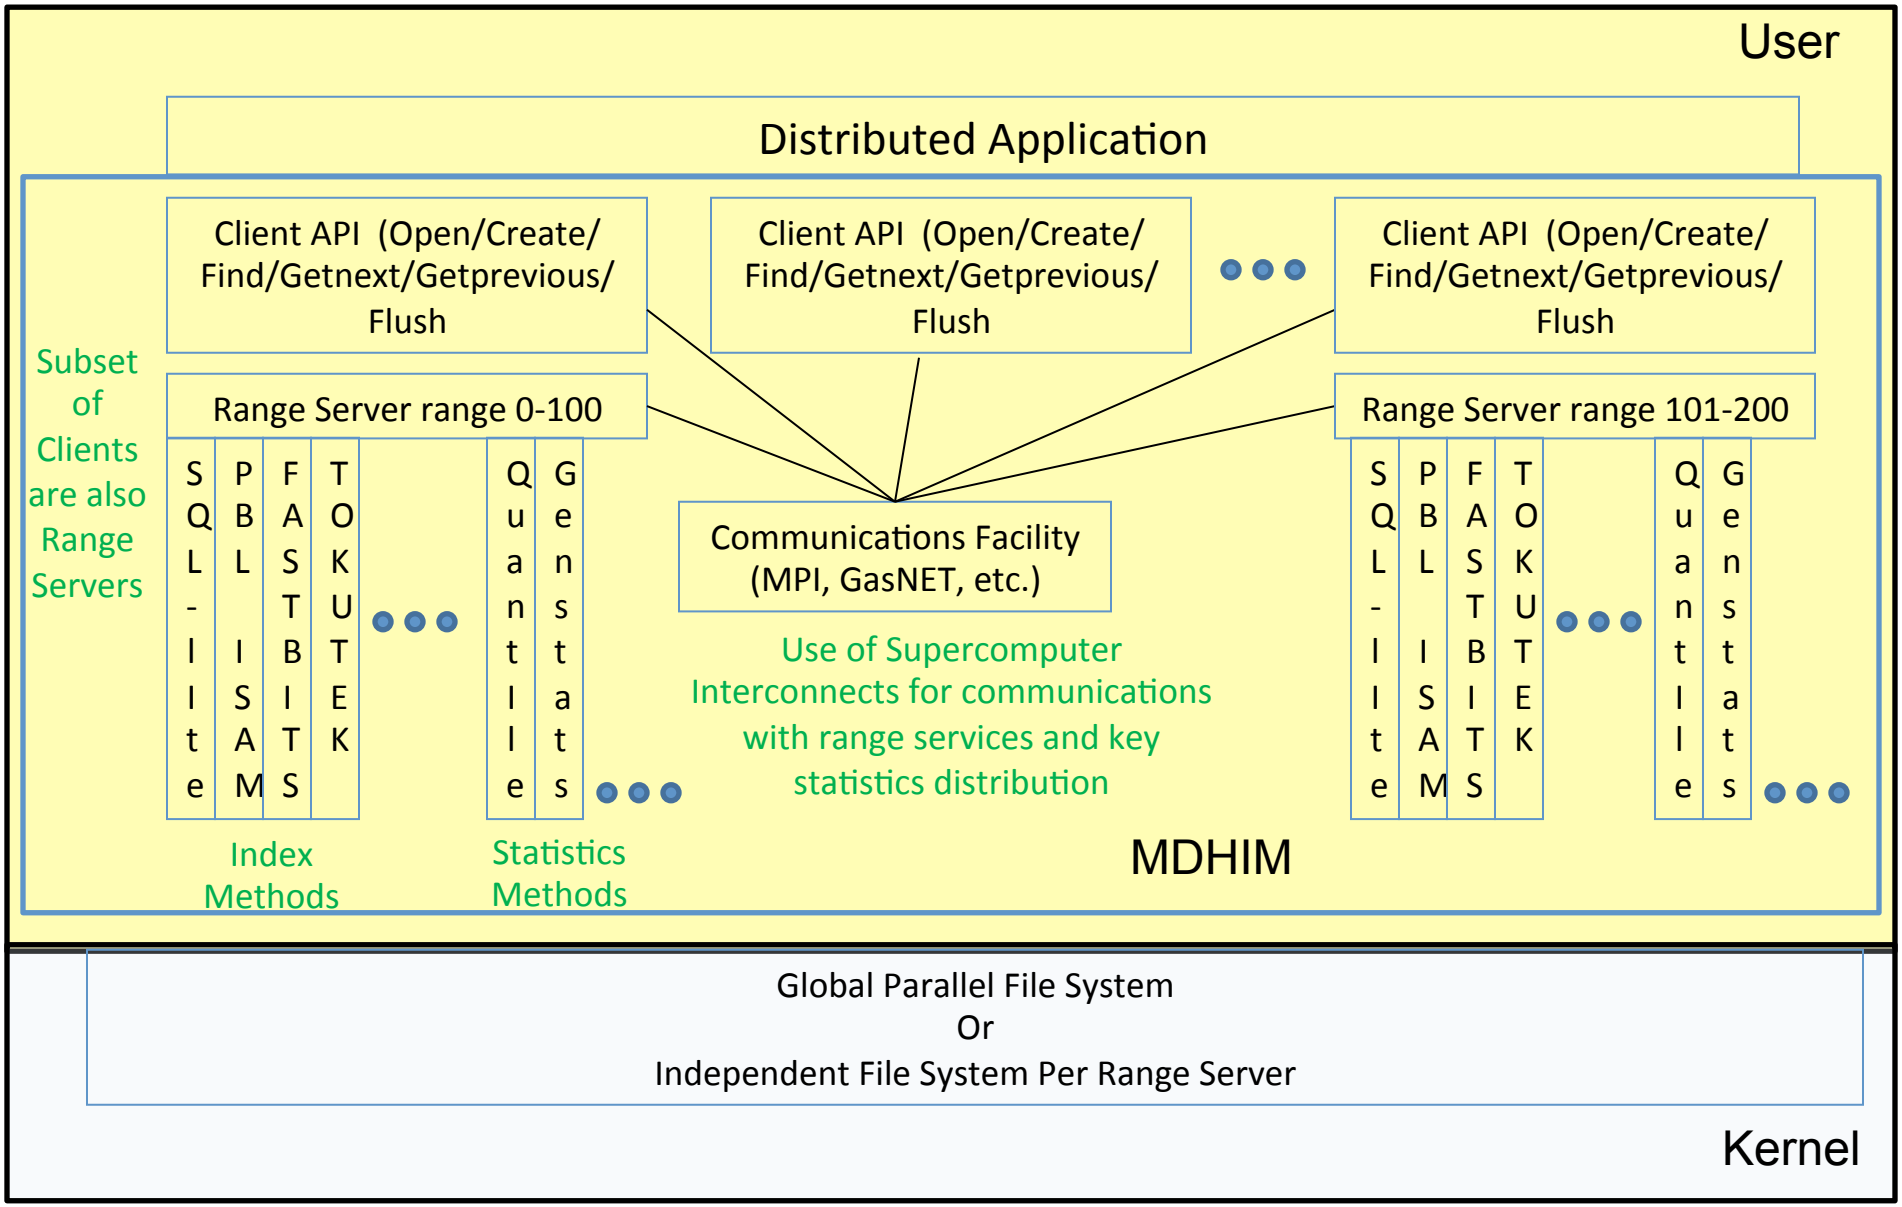

Backup

When we talk about complicated IO stacks, this is what we mean!

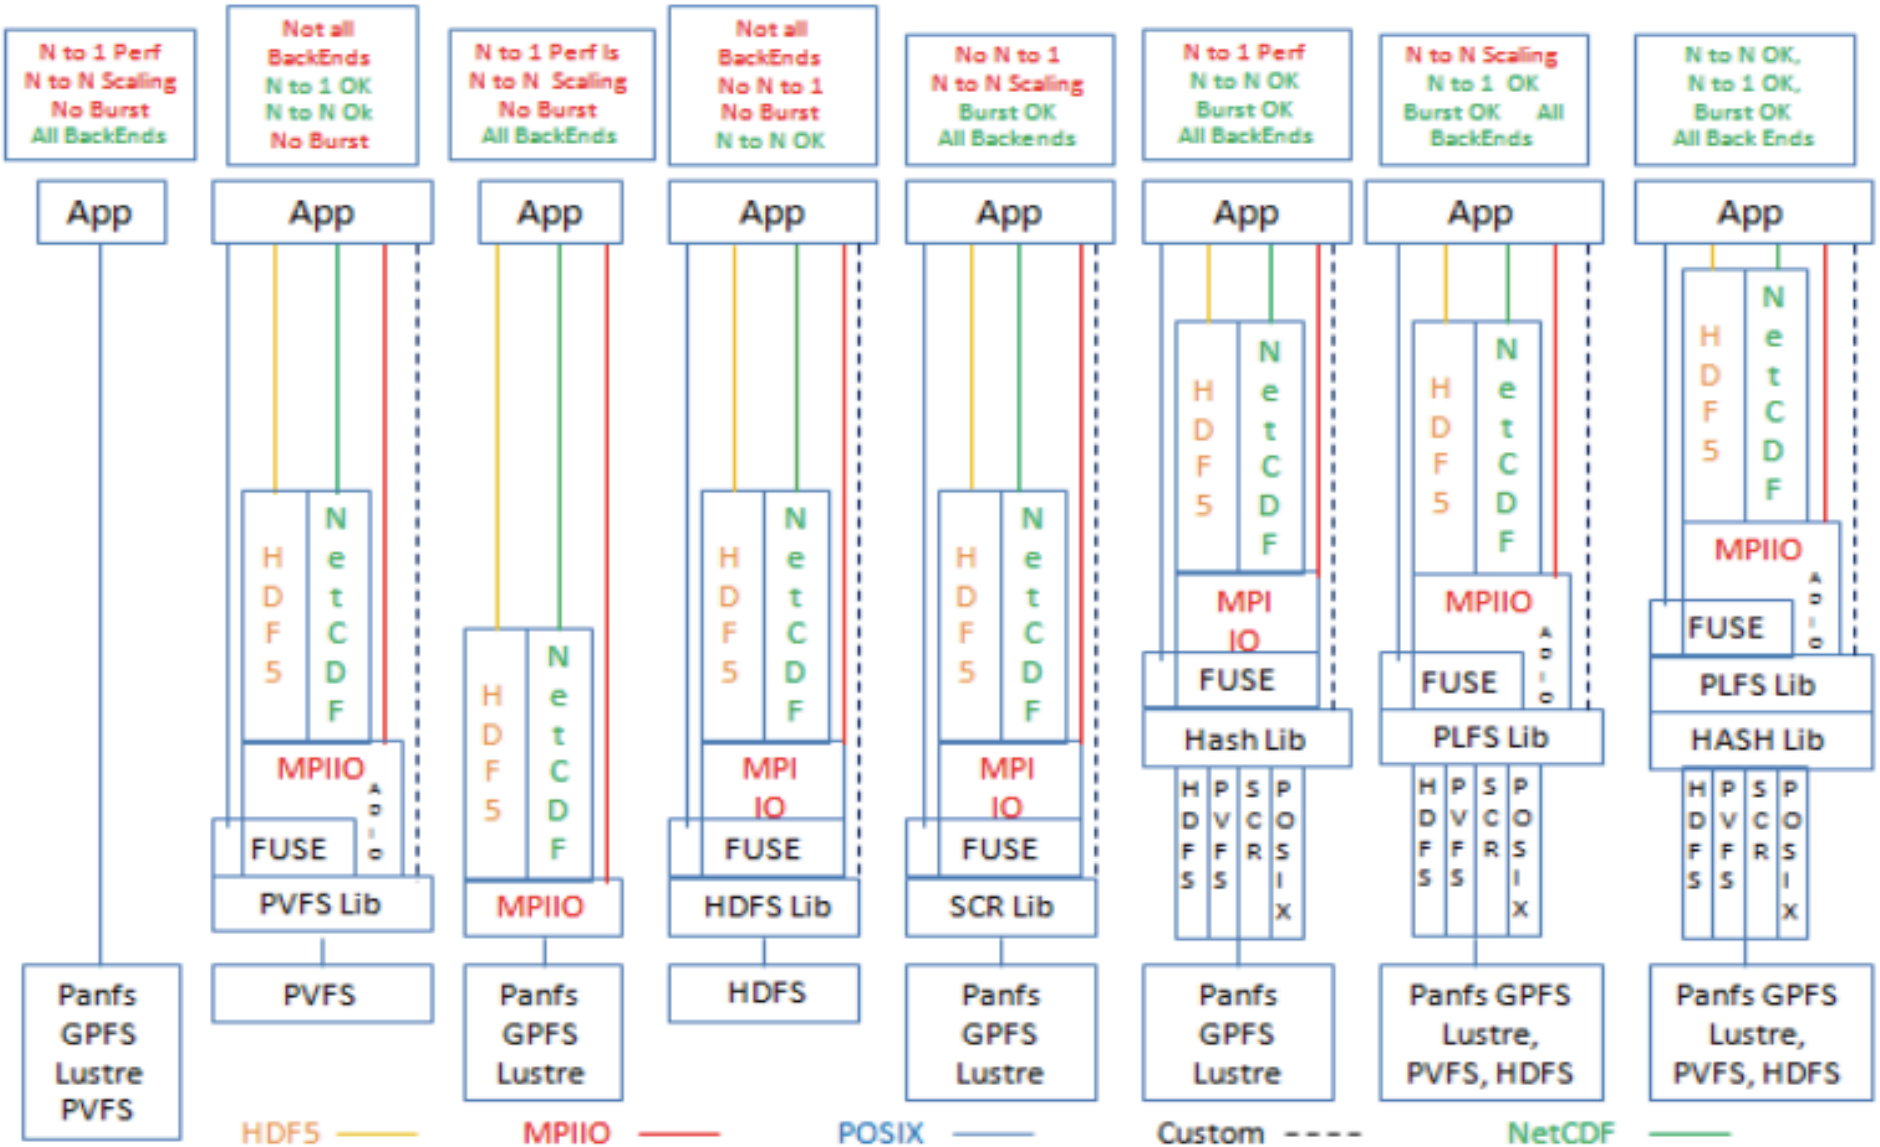

SSIO Exa Nexus Scope

The Data Management (DM) Nexus is focused on the infrastructure necessary to manage, understand, and share data in support of DOE computational science activities at Exascale:

- **Scalable and reliable data management software infrastructure**
  - check pointing, data generation, and analysis
- **Workflow, provenance, and data curation**
  - facilitating execution of complex computational science processes
  - management of information necessary to interpret and reproduce results
- **Data collection, reduction, and transformation**
  - transformation and analysis of scientific data where it resides and as part of data movement, to reduce the cost (in time and energy) to solution
- **Data organization and discovery**
  - technologies for indexing and reorganizion of data to enable time- and power-efficient scientific data