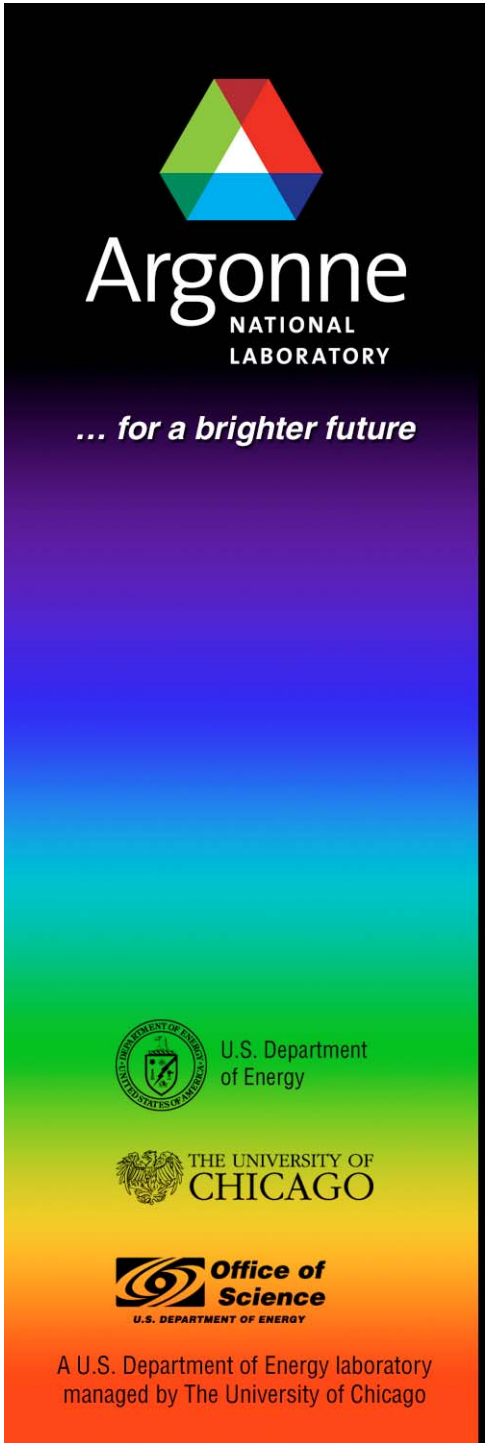

Spreading a Rumor: The SciDAC2 Scientific Data Management Center for Enabling Technologies

(Please just call it the SDM Center)

Rob Ross

Mathematics and Computer Science Division

Argonne National Laboratory

rross@mcs.anl.gov

SDM CENTER

What is SciDAC?

- Scientific Discovery through Advanced Computing
- Brings together application scientists and mathematicians and computer scientists to perform breakthrough science using the world's most powerful computing systems

  (at least, that's my take on it)

- Application groups might include
  - Astrophysics, Fusion, High-Energy Physics, Climate, Biology, etc.
- Computer scientists and mathematicians might provide
  - Solvers, software component infrastructure, I/O tools, etc.

SDM Center: Reducing Data Management Overhead

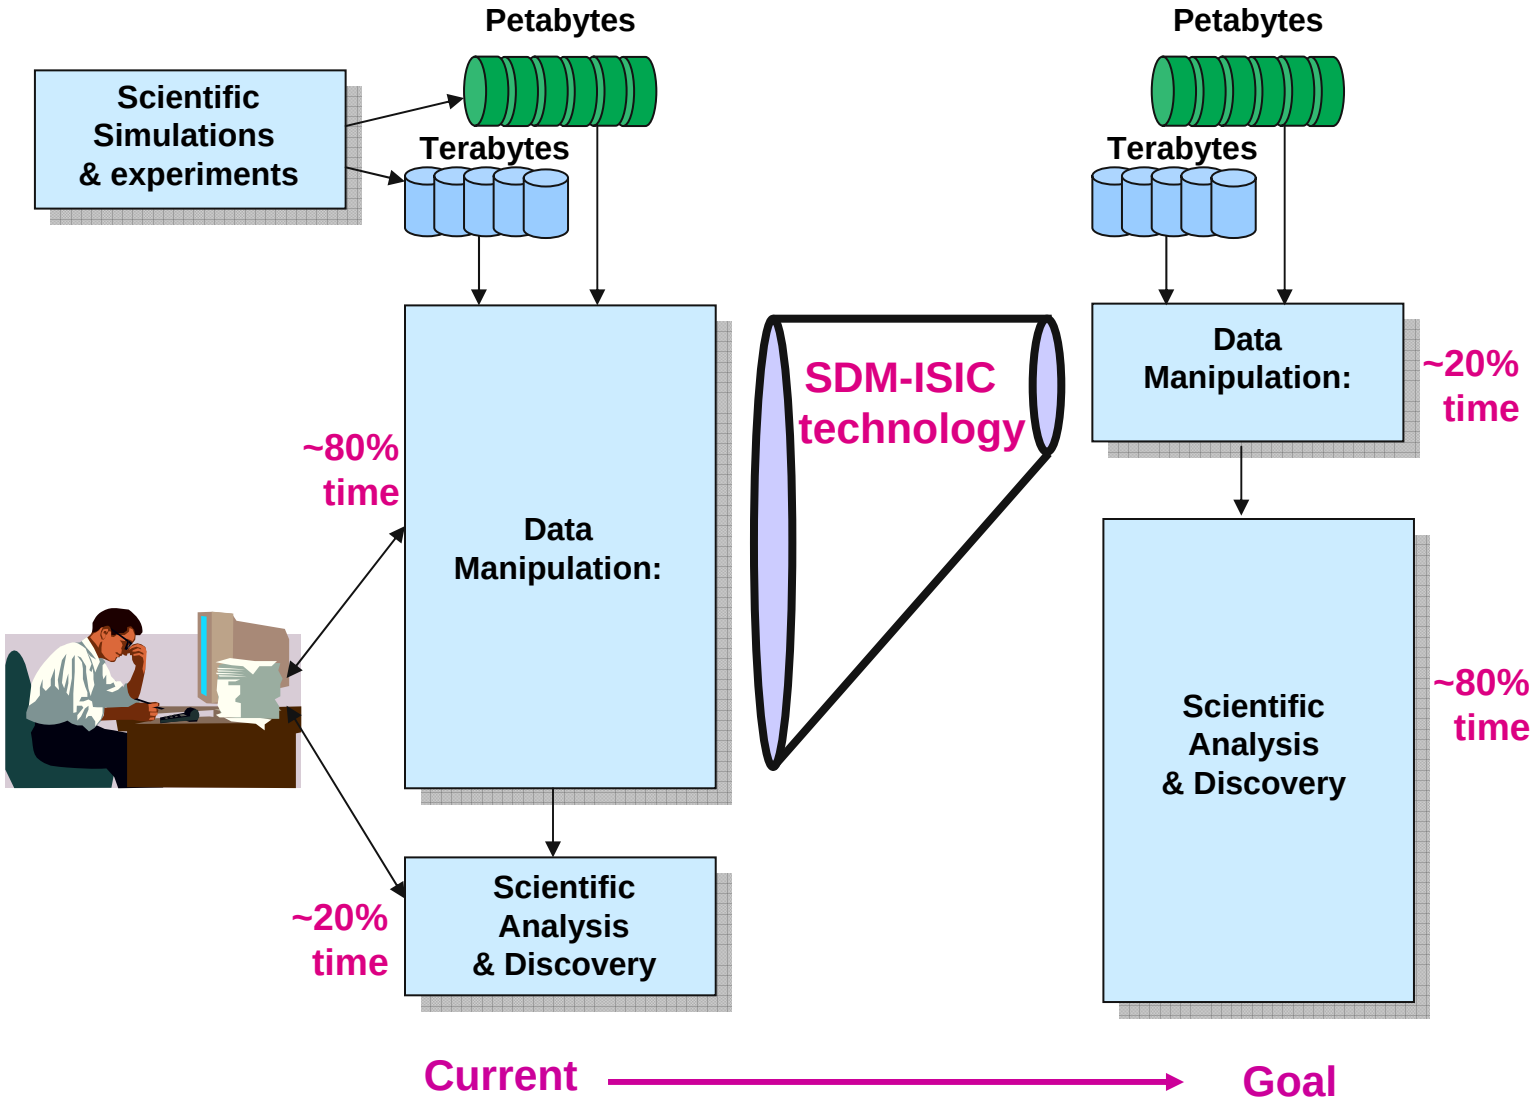

Participating Institutions

- Lawrence Berkeley National Laboratory
  - Arie Shoshani (PI), Doron Rotem
- Argonne National Laboratory
  - Rob Ross, Bill Gropp, Rajeev Thakur
- Lawrence Livermore National Laboratory
  - Terence Critchlow, Chandrika Kamath
- Oak Ridge National Laboratory
  - Nagiza Samatova, Jeff Vetter
- Pacific Northwest National Laboratory
  - Jarek Nieplocha
- Northwestern University
  - Alok Choudhary
- North Carolina State University
  - Mladen Vouk
- University of Utah
  - Steve Parker
- University of California at Davis
  - Bertram Ludaescher
- San Diego Supercomputer Center
  - Ilkay Altinas

- Expertise in a wide range of I/O and data management technologies, from wide-area data movement, to data analytics, to automation of workflows, to I/O middleware and parallel file systems

Proposed Approach

- Use an integrated framework that:
  - Provides scientific workflow capabilities
  - Supports data analytics tools
  - Accelerates data storage and access
- Simplify data management tasks for the scientist
  - Hide details of underlying I/O and indexing technologies
  - Permit convenient assembly of workflows
- **Focus on leadership class machines and associated applications**

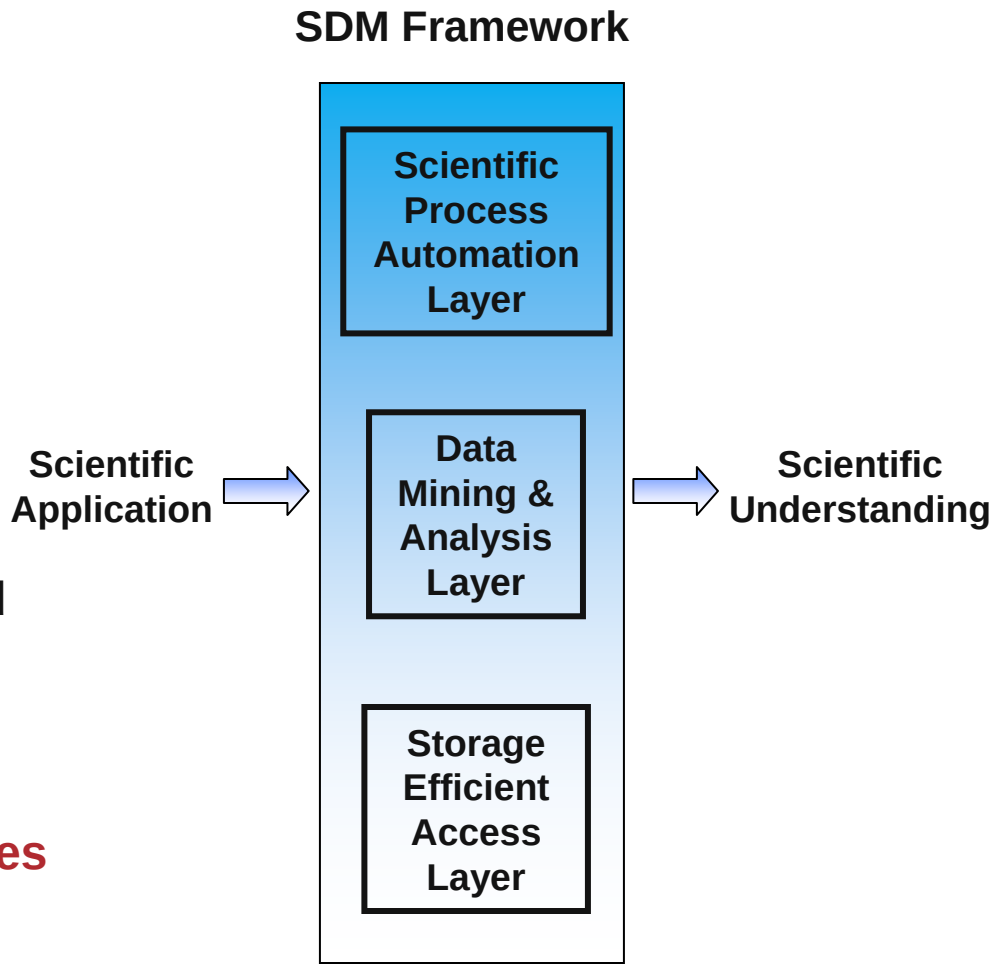

SDM Center Impact

Astrophysics:

High speed storage technology, parallel NetCDF, parallel VTK, and ASPECT integration software used for Terascale Supernova Initiative (TSI) and FLASH simulations

Tony Mezzacappa – ORNL, Mike Zingale – U of Chicago, Mike Papka – ANL

Scientific Workflow

John Blondin – NCSU

Doug Swesty, Eric Myra – Stony Brook

ASCI FLASH – parallel NetCDF

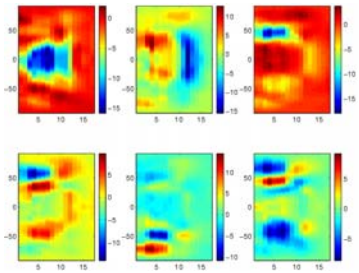

Dimensionality reduction

Climate:

High speed storage technology, Parallel NetCDF, and ICA technology used for Climate Modeling projects

Ben Santer – LLNL, John Drake – ORNL, John Michalakes – NCAR

Combustion:

Compressed Bitmap Indexing used for fast generation of flame regions and tracking their progress over time

Wendy Koegler, Jacqueline Chen – Sandia Lab

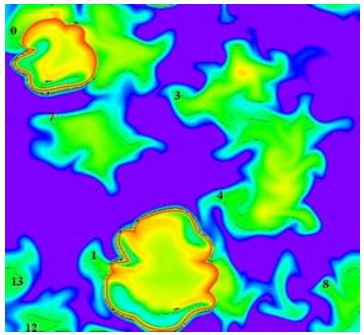

Region growing

SDM Center Impact (cont'd)

Biology:

Kepler workflow system and web-wrapping technology used for executing complex highly repetitive workflow tasks for processing microarray data

Matt Coleman - LLNL

High Energy Physics:

Compressed Bitmap Indexing and Storage Resource Managers used for locating desired subsets of data (events) and automatically retrieving data from HPSS

Doug Olson - LBNL, Eric Hjort – LBNL, Jerome Lauret - BNL

Fusion:

A combination of PCA and ICA technology used to identify the key parameters that are relevant to the presence of edge harmonic oscillations in a Tokomak

Keith Burrell - General Atomics

Scott Klasky - ORNL

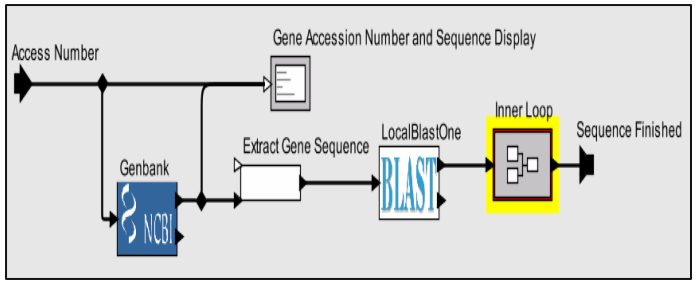

Building a scientific workflow

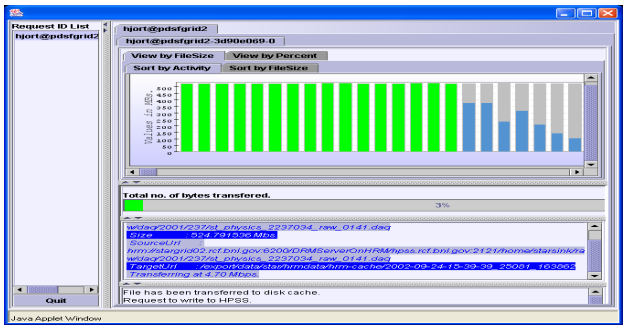

Dynamic monitoring of HPSS file transfers

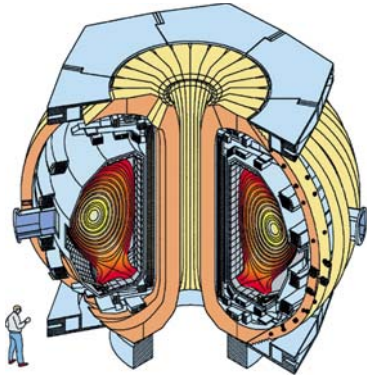

Identifying key parameters for the DIII-D Tokamak

Technology Details by Layer

Scientific Process Automation (SPA) Layer

Data Mining and Analysis (DMA) Layer

Storage Efficient Access (SEA) Layer

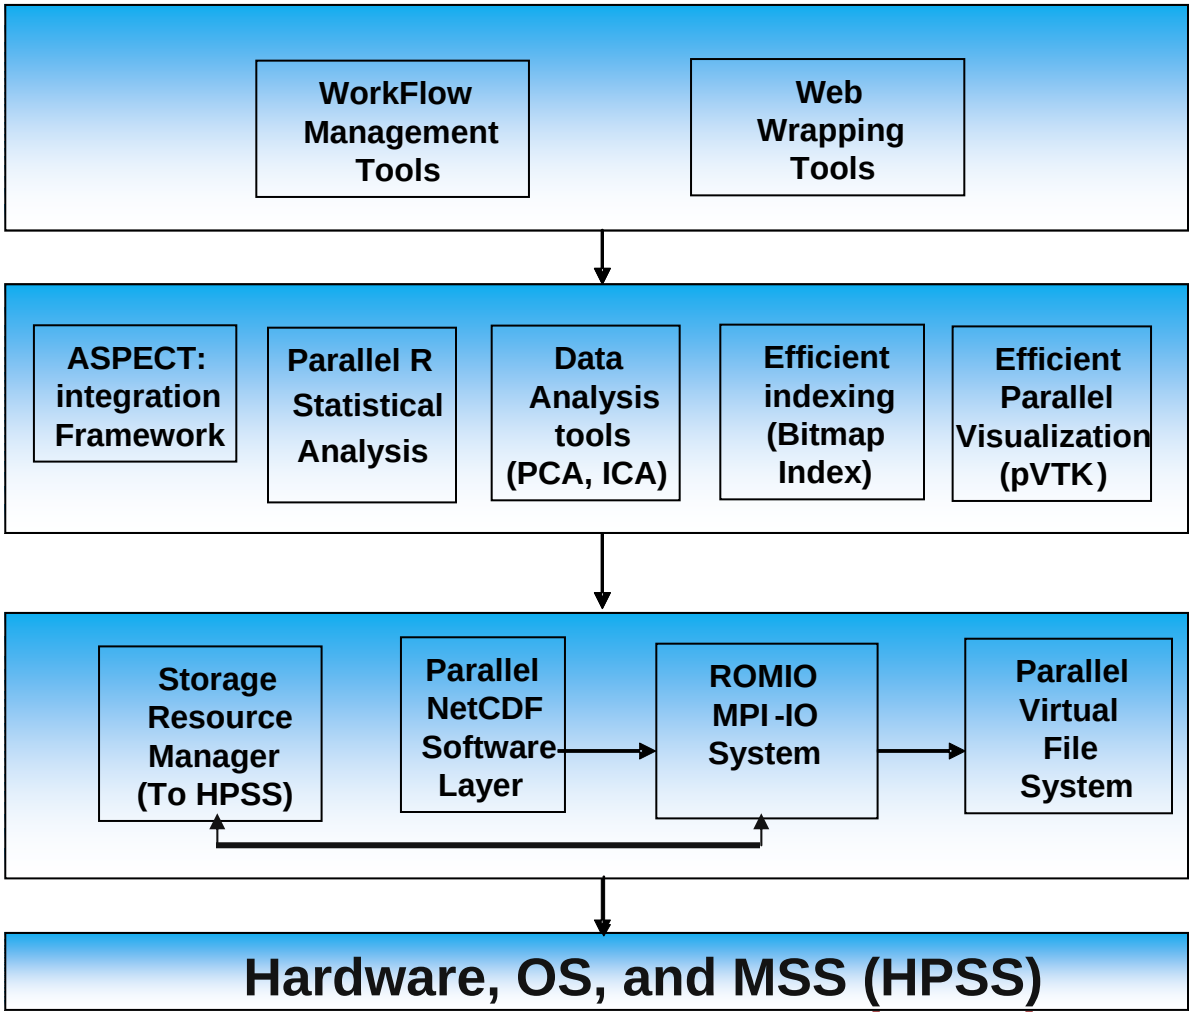

SDM Storage Efficient Access (SEA) Layer

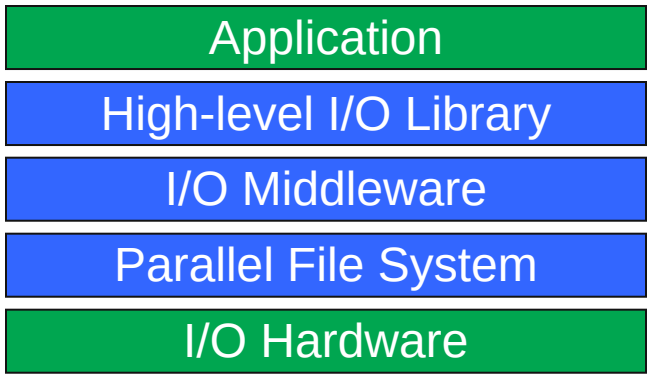

- This layer is the most relevant to the HECIWG IO and FS
- Break up I/O support into multiple layers with distinct roles:
  - High level I/O library maps app. abstractions to a structured, portable file format (e.g. HDF5, Parallel netCDF)
  - Middleware layer deals with organizing access by many processes (e.g. MPI-IO, UPC-IO)
  - Parallel file system maintains logical space, provides efficient access to data (e.g. PVFS, GPFS, Lustre)
- Argonne, Northwestern, LBNL, ORNL, and PNNL active in this layer

The SDM Center and You

- Stages in technology development
  Research ➔ Prototype ➔ Product ➔ Infrastructure
  - Technology starts off as research ideas
  - Technology becomes "infrastructure" when sites start installing it by default
    - *Happens when use in applications hits a certain critical mass*

- Primary role of SDM center: Prototype ➔ Product
  - Apply technologies that have been prototyped
  - Fill gaps in I/O software stack on leadership class machines
  - Make application groups aware of these products and their uses
- SDM center might be helpful to HECURA I/O researchers by
  - Identifying critical needs on leadership class machines
    - *Possibly low-hanging fruit?*
  - Connecting you with SciDAC application groups with relevant I/O needs
    - *Your research might provide revolutionary capabilities*
  - Aiding in integration of prototypes into products
    - *Wider distribution of software often uncovers new challenges…*

SDM Center Contacts

- SDM Center web site
  http://sdmcenter.lbl.gov

- Arie Shoshani (Center PI)
  ashoshani@lbl.gov

- Rob Ross (SEA area leader)
  rross@mcs.anl.gov