

ATPESC 2016



INTRODUCTION TO SESSION: DATA INTENSIVE COMPUTING AND I/O



ROB LATHAM

PHILP CARNS

Argonne National Laboratory

robl@mcs.anl.gov

carns@mcs.anl.gov

GERD HEBER

HDF Group

gheber@hdfgroup.org

STU MARTIN

Argonne National Laboratory

smartin@mcs.anl.gov

August 11, 2016

St. Charles IL

ABOUT US

- Rob Latham
 - Principal Software Development Specialist, MCS Division
Argonne National Laboratory
 - ROMIO MPI-IO implementation
 - Parallel netCDF high-level I/O library
 - Application outreach
- Phil Carns
 - Principal Software Development Specialist, MCS Division,
Argonne National Laboratory
 - Darshan characterization tool
 - PVFS file system

ABOUT US

- Gerd Heber
 - Applications Architect, The HDF Group
 - HDF5 on-ramp solutions
 - Application support
 - Data modeling


- Stu Martin
 - Software Development Manager, Computation Institute, University of Chicago
 - Globus Operations and Support Manager
 - Globus Toolkit
 - Used to play Ultimate at lunch with RobL

COMPUTATIONAL SCIENCE

- Use of computer simulation as a tool for greater understanding of the real world
 - Complements experimentation and theory
- Problems are increasingly computationally expensive
 - Large parallel machines needed to perform calculations
 - Critical to leverage parallelism in all phases
- **Data access is a huge challenge**
 - Using parallelism to obtain performance
 - Finding usable, efficient, and portable interfaces
 - Understanding and tuning I/O



IBM Blue Gene/Q system at Argonne National Laboratory.



Visualization of entropy in Terascale Supernova Initiative application. Image from Kwan-Liu Ma's visualization team at UC Davis.

GOALS AND OUTLINE

- Goals:
 - Share our view of HPC I/O hardware and software
 - Discuss interfaces that you can use to access I/O resources
 - Point to emerging and future trends in HPC I/O

- Outline (roughly)
 - Ways of thinking about I/O systems
 - How It Works: HPC I/O Systems
 - Using I/O systems
 - HDF5
 - Globus
 - Making sense of the whole mess
 - Emerging and future trends

- Notes
 - There will be slides that are hidden, don't be alarmed
 - After the morning break, we'll be looking through some of this code:
<https://xgitlab.cels.anl.gov/ATPESC-IO/hands-on-2016>