

SciDAC-2 Petascale Data Storage Institute

Presented by

Philip C. Roth

Future Technologies Group
Computer Science and Mathematics Division



The petascale storage problem

- Petascale computing makes petascale demands on storage.

- Performance
- Capacity
- Concurrency
- Reliability
- Availability
- Manageability



- Parallel file systems are barely keeping pace at terascale; the challenges will be much greater at petascale.

Petascale Data Storage Institute

- The PDSI is an institute in the Department of Energy (DOE) Office of Science's Scientific Discovery through Advanced Computing (SciDAC-2) program.
- Using diverse expertise with applications and file and storage systems, members will collaborate on requirements, standards, algorithms, and analysis tools.



Led by Dr. Garth Gibson, Carnegie Mellon University

<http://www.pdsi-scidac.org>

Participating institutions

Carnegie Mellon University

Lawrence Berkeley National Laboratory/NERSC

Los Alamos National Laboratory

Oak Ridge National Laboratory

Pacific Northwest National Laboratory

Sandia National Laboratories

University of California at Santa Cruz

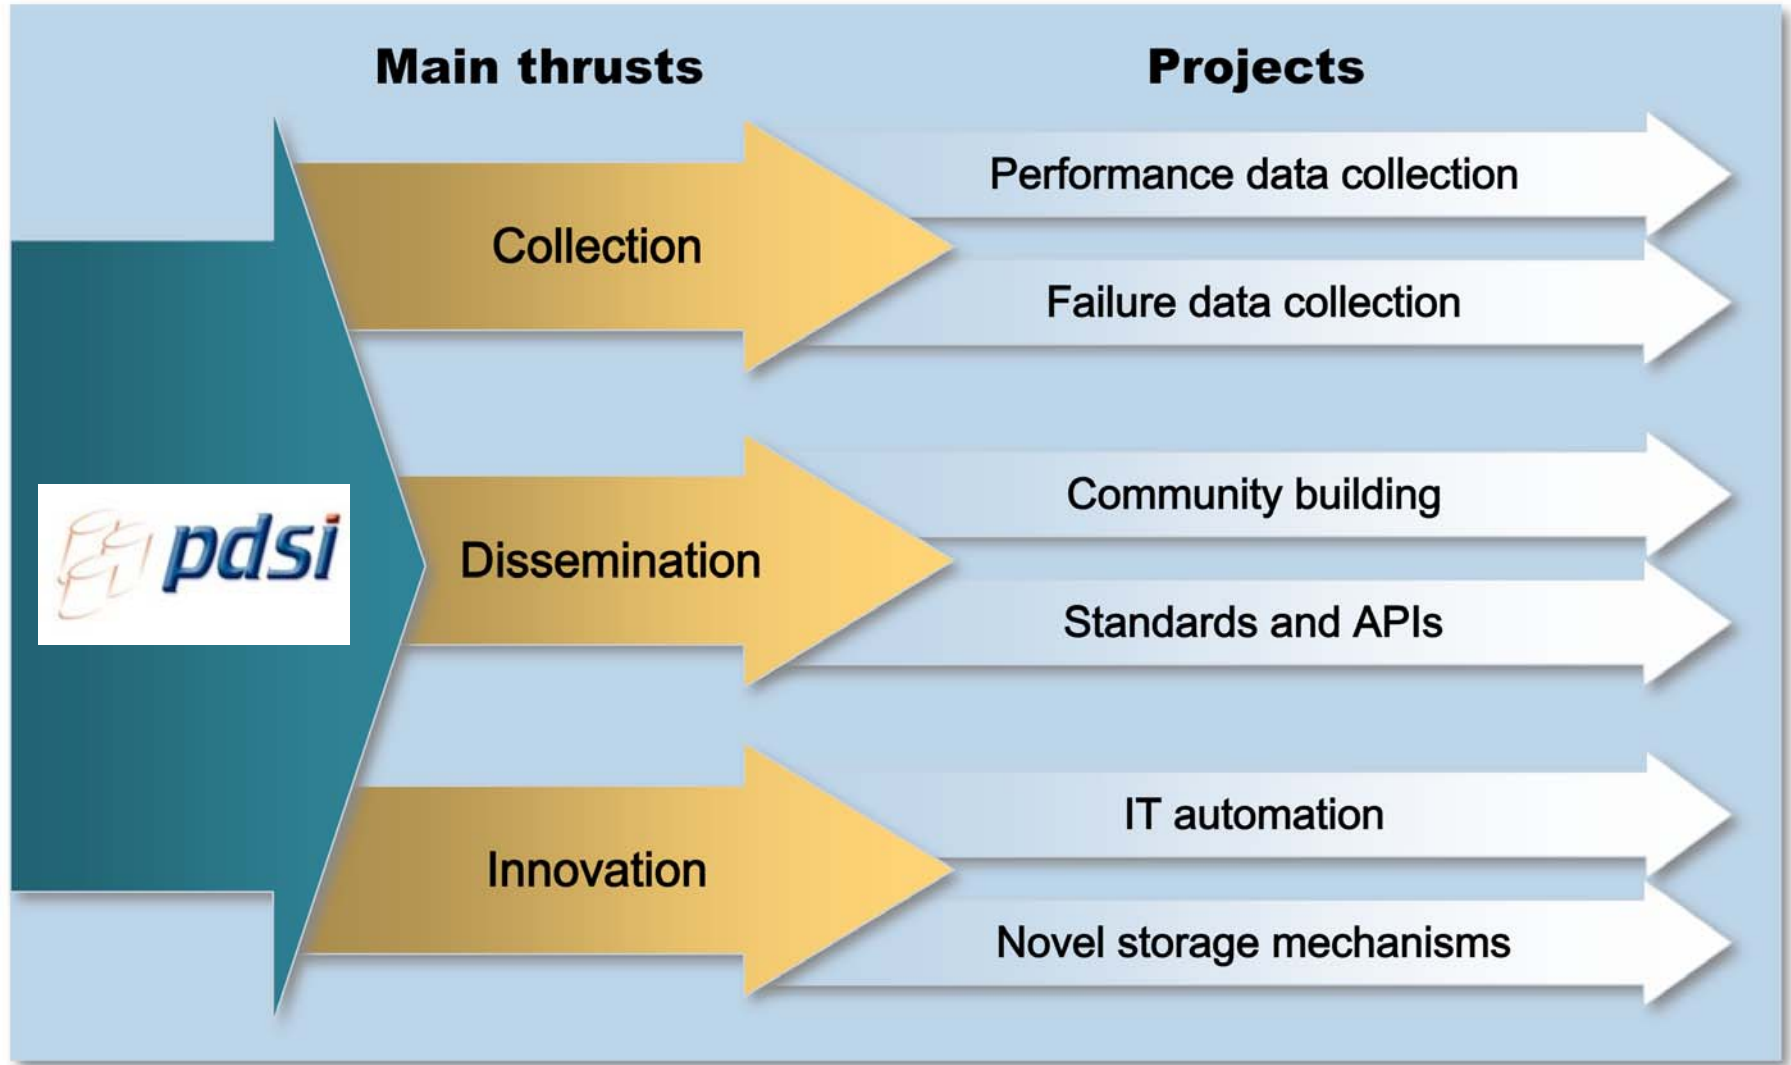
University of Michigan at Ann Arbor



Pacific Northwest
National Laboratory

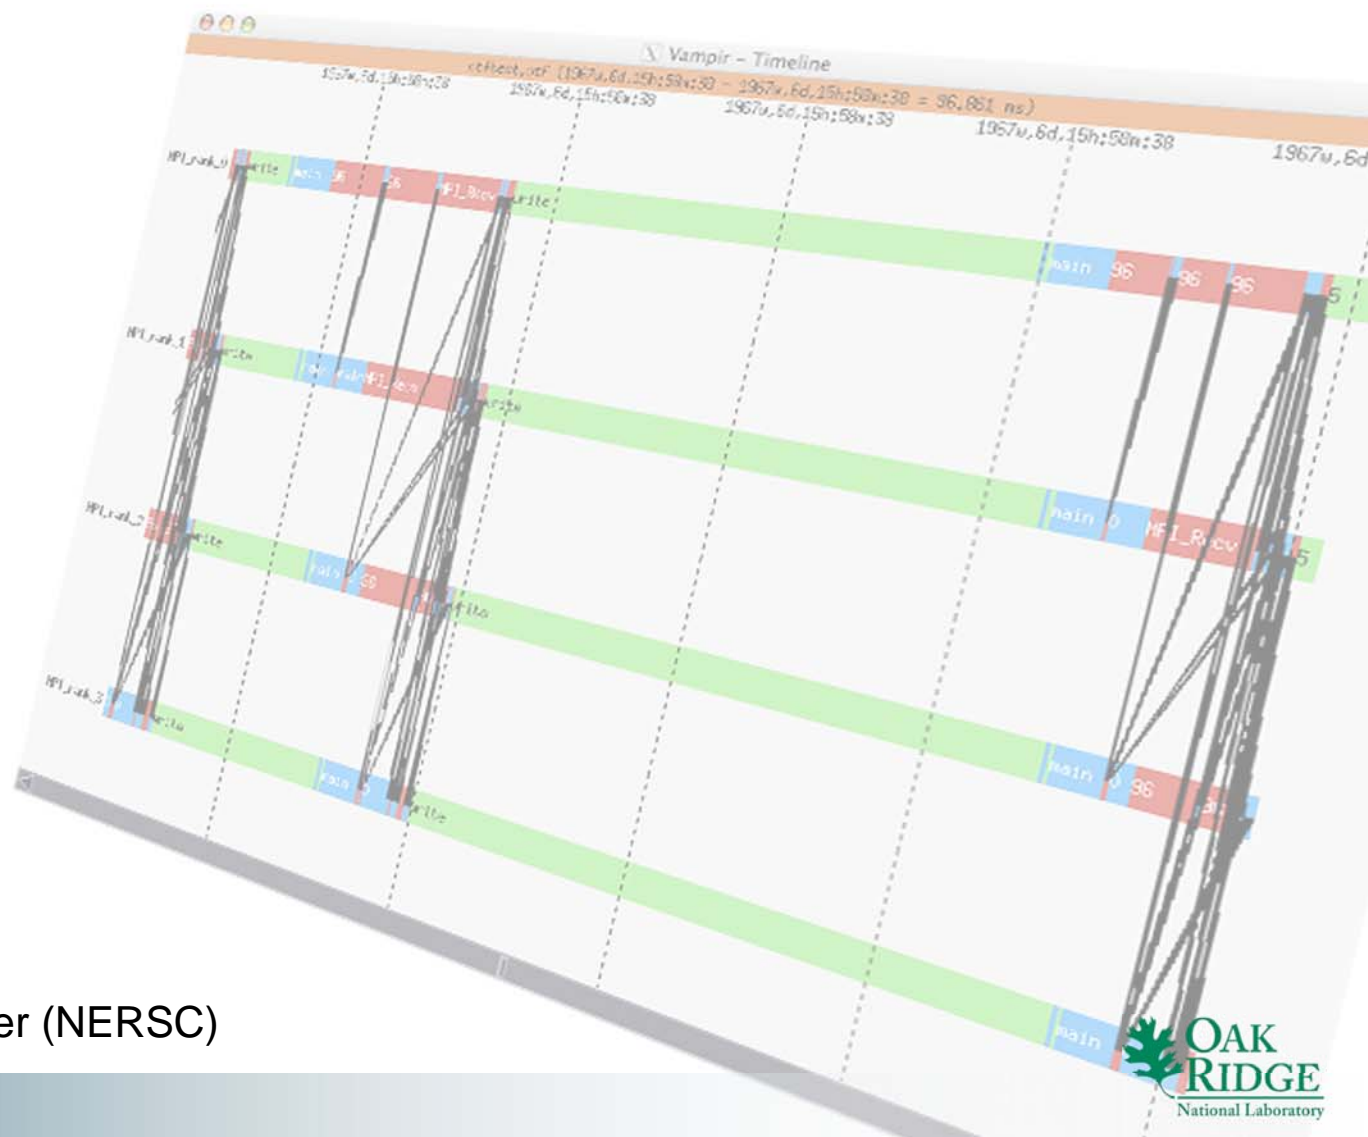


Petascale Data Storage Institute agenda



Collection: Performance analysis

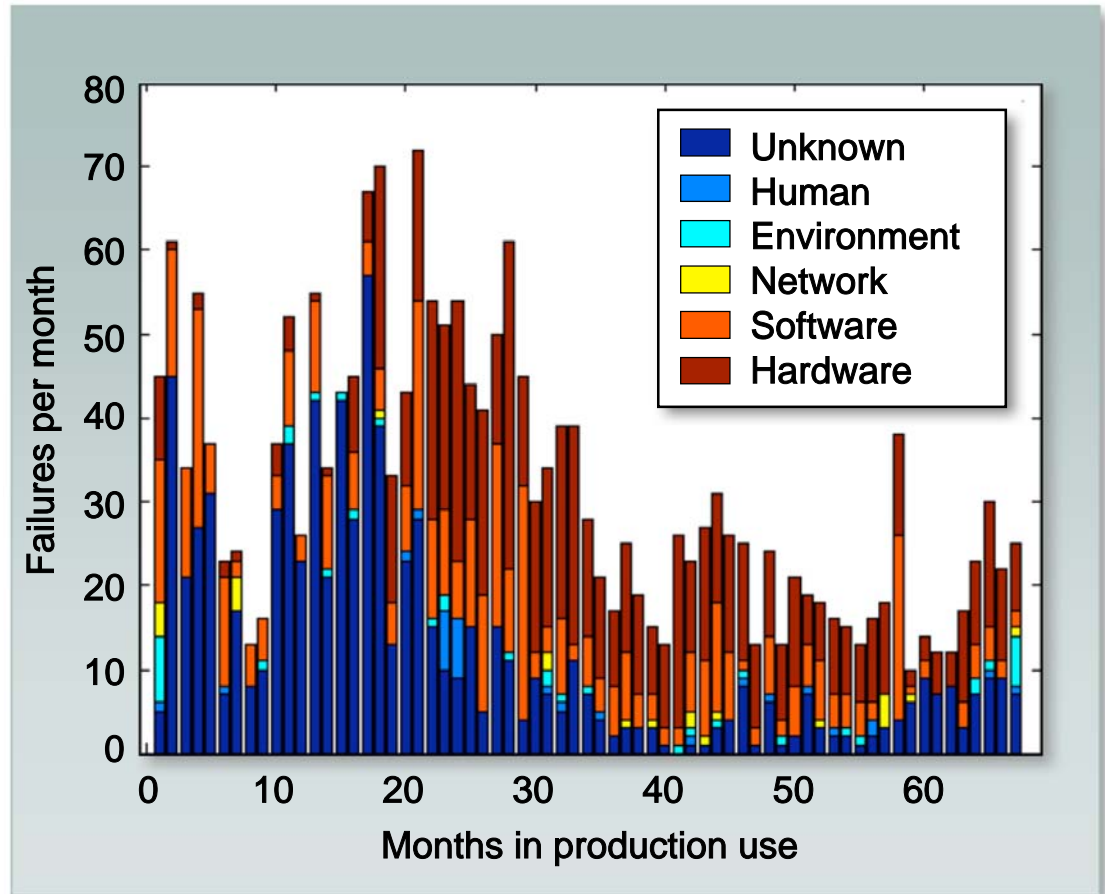
- Performance data collection and analysis
- Workload characterization
- Benchmark collection and publication



Led by William Kramer,
National Energy Research
Scientific Computing Center (NERSC)

Collection: Failure analysis

- Capture and analyze failure, error, and usage data from high-end computing systems
- Initial example: Los Alamos failure data available for 22 systems over 9 years with extensive analysis by Bianca Schroeder, Carnegie Mellon University



<http://institutes.lanl.gov/data>
<http://www.pdl.cmu.edu/FailureData>

Led by Gary Grider, Los Alamos National Laboratory

Dissemination: Outreach

Goal: To disseminate information about techniques, mechanisms, best practices, and available tools

- **Our approach**
 - Workshops (SC07 Petascale Data Storage Workshop, November 11)
 - Tutorials and course materials
 - Online, open repository with documents, tools, and performance and failure data
- **Target audience**
 - Computational scientists
 - Academia (professors and students)
 - Industry (storage researchers and developers)

Led by Dr. Garth Gibson, Carnegie Mellon University

Dissemination: Standards and APIs

Goals: To facilitate standards development and deployment and to validate and demonstrate new extensions and protocols

Some work under way

- **POSIX extensions**
 - e.g., support for weak data and metadata consistency
 - <http://www.pdl.cmu.edu/posix>
- **Parallel Network File System (pNFS)**
 - In IETF NFSv4.1 standard draft
 - University of Michigan Center for Information Technology Integration producing reference implementation
 - <http://www.pdl.cmu.edu/pNFS>

Led by Gary Grider, Los Alamos National Laboratory

Innovation

IT automation applied to high-end computing systems and problems

- Storage system instrumentation for machine learning
- Data layout and access planning
- Automated diagnosis, tuning, and failure recovery

Led by Dr. Garth Gibson,
Carnegie Mellon University

Novel mechanisms for core high-end computing storage problems

- WAN/global storage access
- High-performance collective operations
- Rich metadata at scale
- Integration with system virtualization technology

Led by Darrell Long,
University of California at Santa Cruz

Summary

- The Petascale Data Storage Institute brings together individuals with expertise in file and storage systems, applications, and performance analysis.
- PDSI is a focal point for computational scientists, academia, and industry for storage-related information and tools, both within and outside SciDAC-2.



<http://www.pdsi-scidac.org>

Contact

Philip C. Roth

Future Technologies Group

Computer Science and Mathematics Division

(865) 241-1543

rothpc@ornl.gov

