U.S. DEPARTMENT OF ENERGY | Office of Science

Next Generation Networking for Science Program Update

NGNS Program Managers
Richard Carlson
Thomas Ndousse
ASCAC meeting 11/21/2014

NGNS Program Mission

- **The Goal of the program is to research, develop, test and deploy advanced network and middleware technologies critical to DOE's science mission**
- **A unique focus is the end-to-end aspect of the high-performance, high-capacity network and middleware technologies needed to provide secure access to distributed science facilities, high-performance computing resources and large-scale scientific collaborations.**

NGNS Budget Trend

Funding in $K

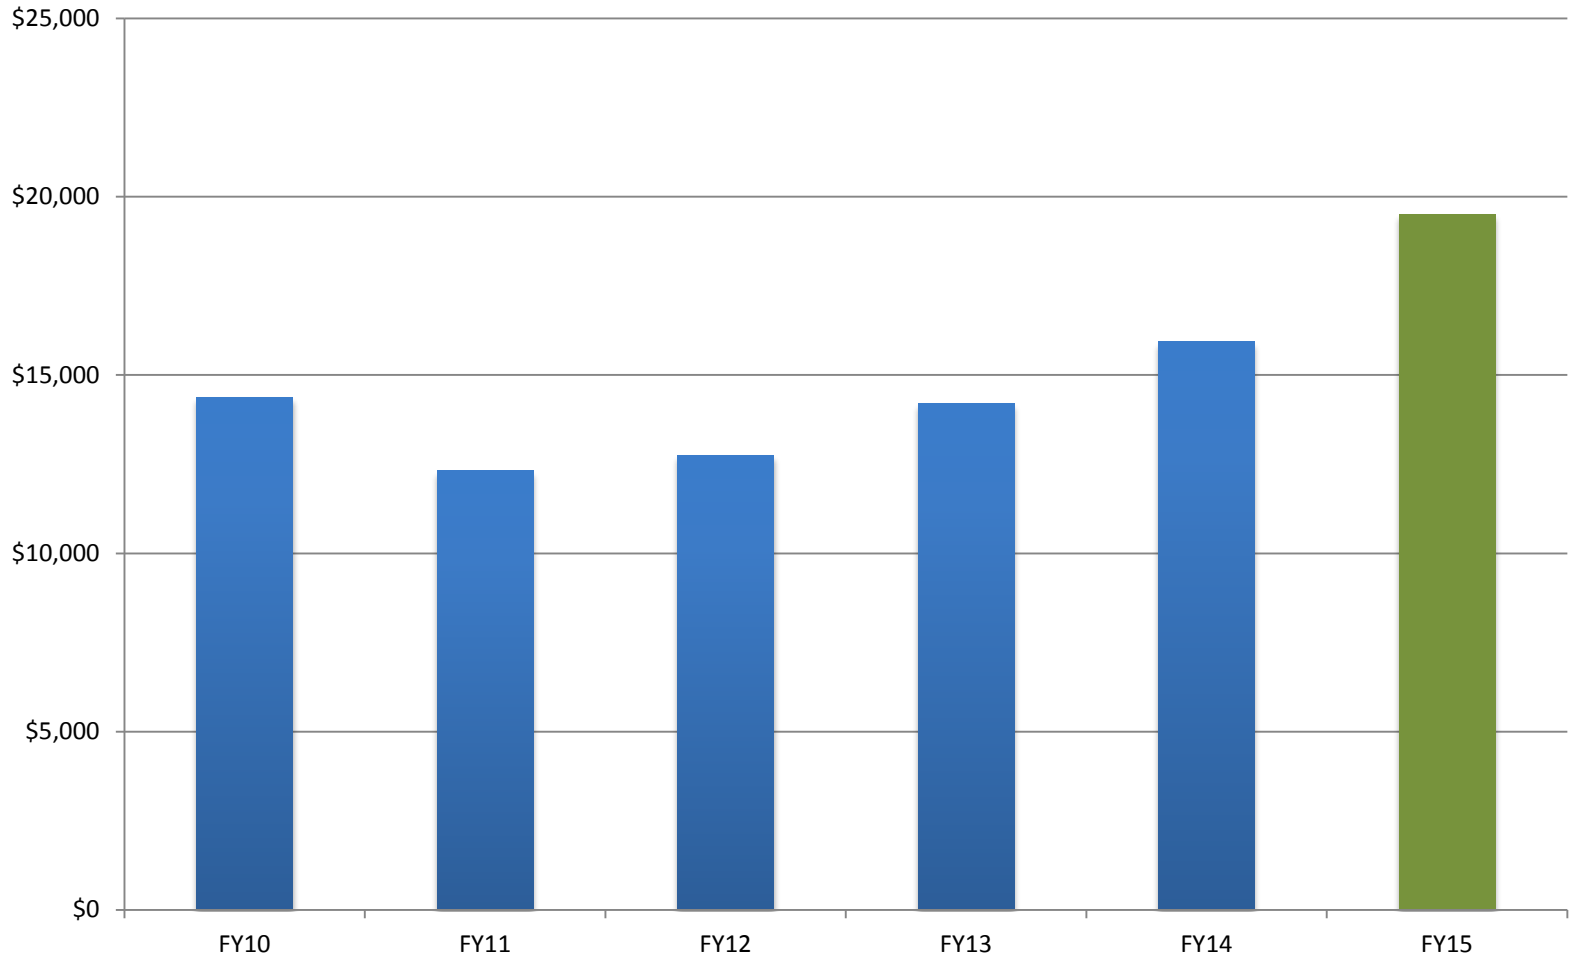

FY14 Funding Distribution

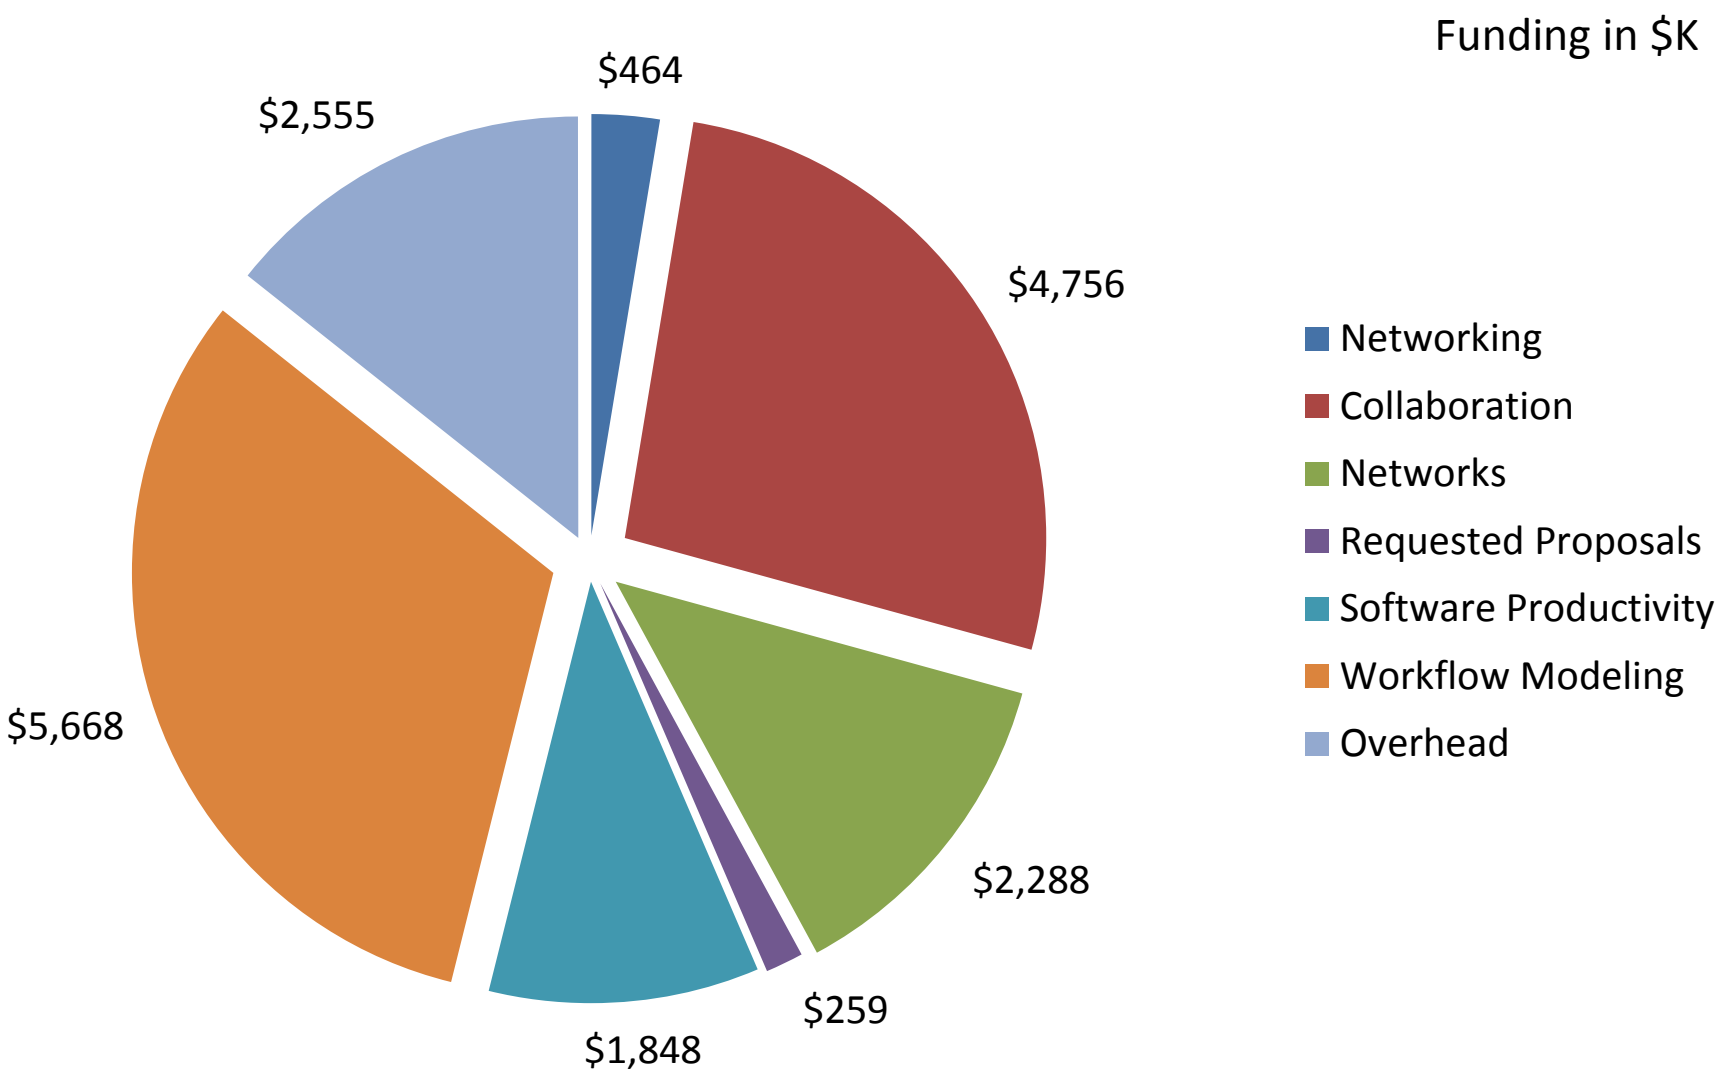

NGNS Program Management

- **Regular calls and interactions with PI's**
- **Workshops to gather new requirements**
- **Annual PI meetings**
  - URL's to last 3 PI meetings
    - https://indico.bnl.gov/conferenceDisplay.py?confId=489
    - https://indico.bnl.gov/conferenceDisplay.py?confId=566
    - https://www.orau.gov/ngnspi2014/agenda.htm
  - 2014 PI meeting highlights
    - 7 sessions for 24 projects
    - 3-4 speakers gave 15 minute talks per session
    - 45-60 minute discussion per session
    - Poster session

WP2 PanDA architecture for Titan

- **Pilot(s) executes on HPC interactive node**
- **Pilot interact with local job scheduler to manage job**
- **Data, produced on HPC automatically moves to external storage**

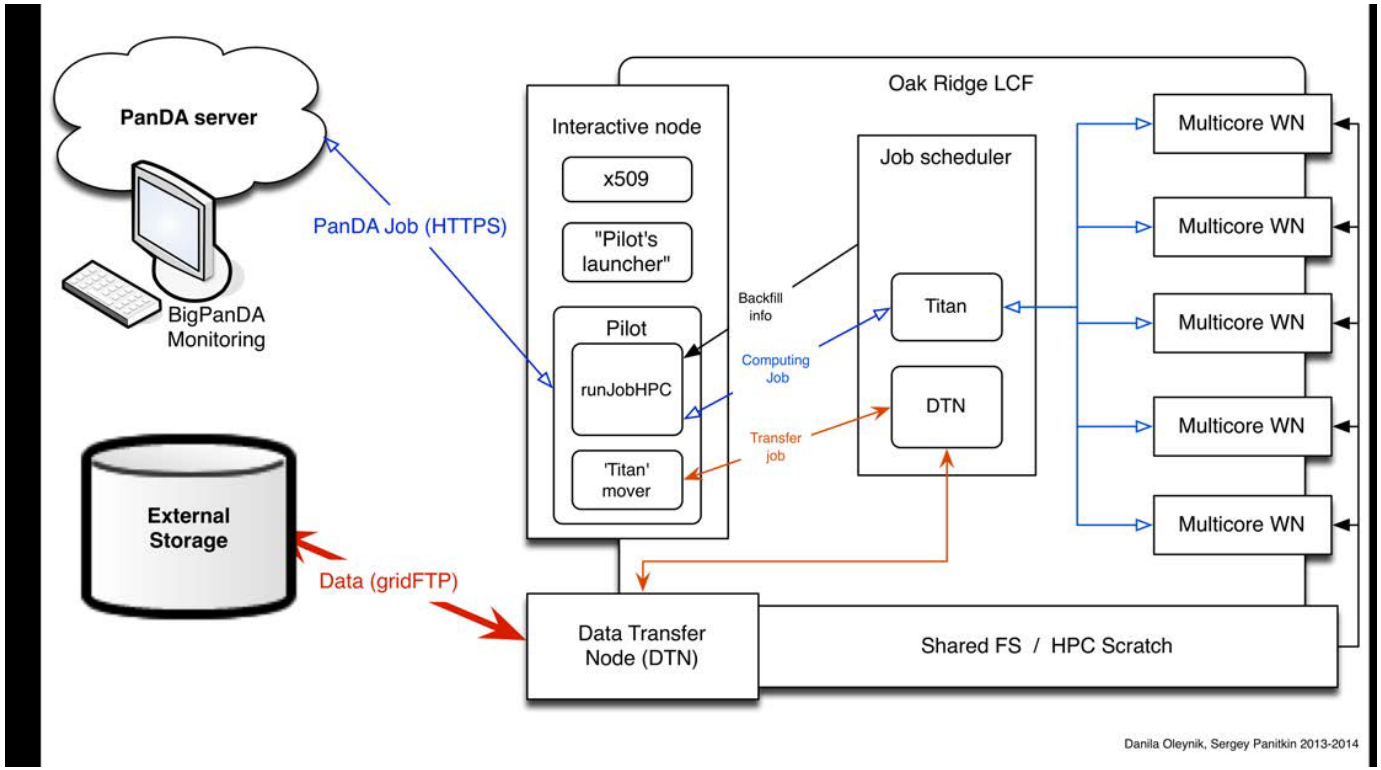

Started Integration work with NERSC using the same architecture

Multicore-aware Data Transfer Middleware

I/O-Centric architecture

Parallel data transfer

Data layout preprocessing

Data flow-centric scheduling

NUMA-awareness scheduling

I/O Locality optimization

Maximizing Parallelism

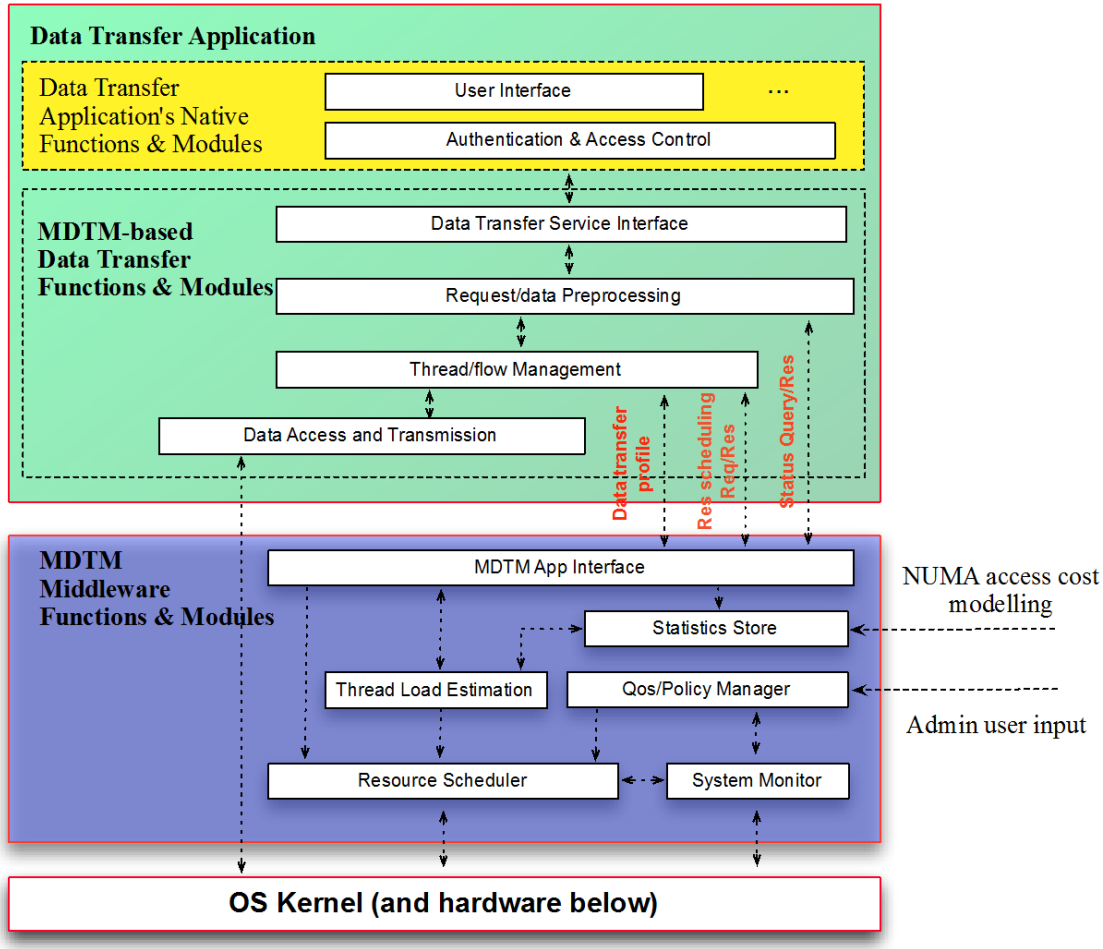

Automated Metadata, Provenance Cataloging and Navigable Interfaces

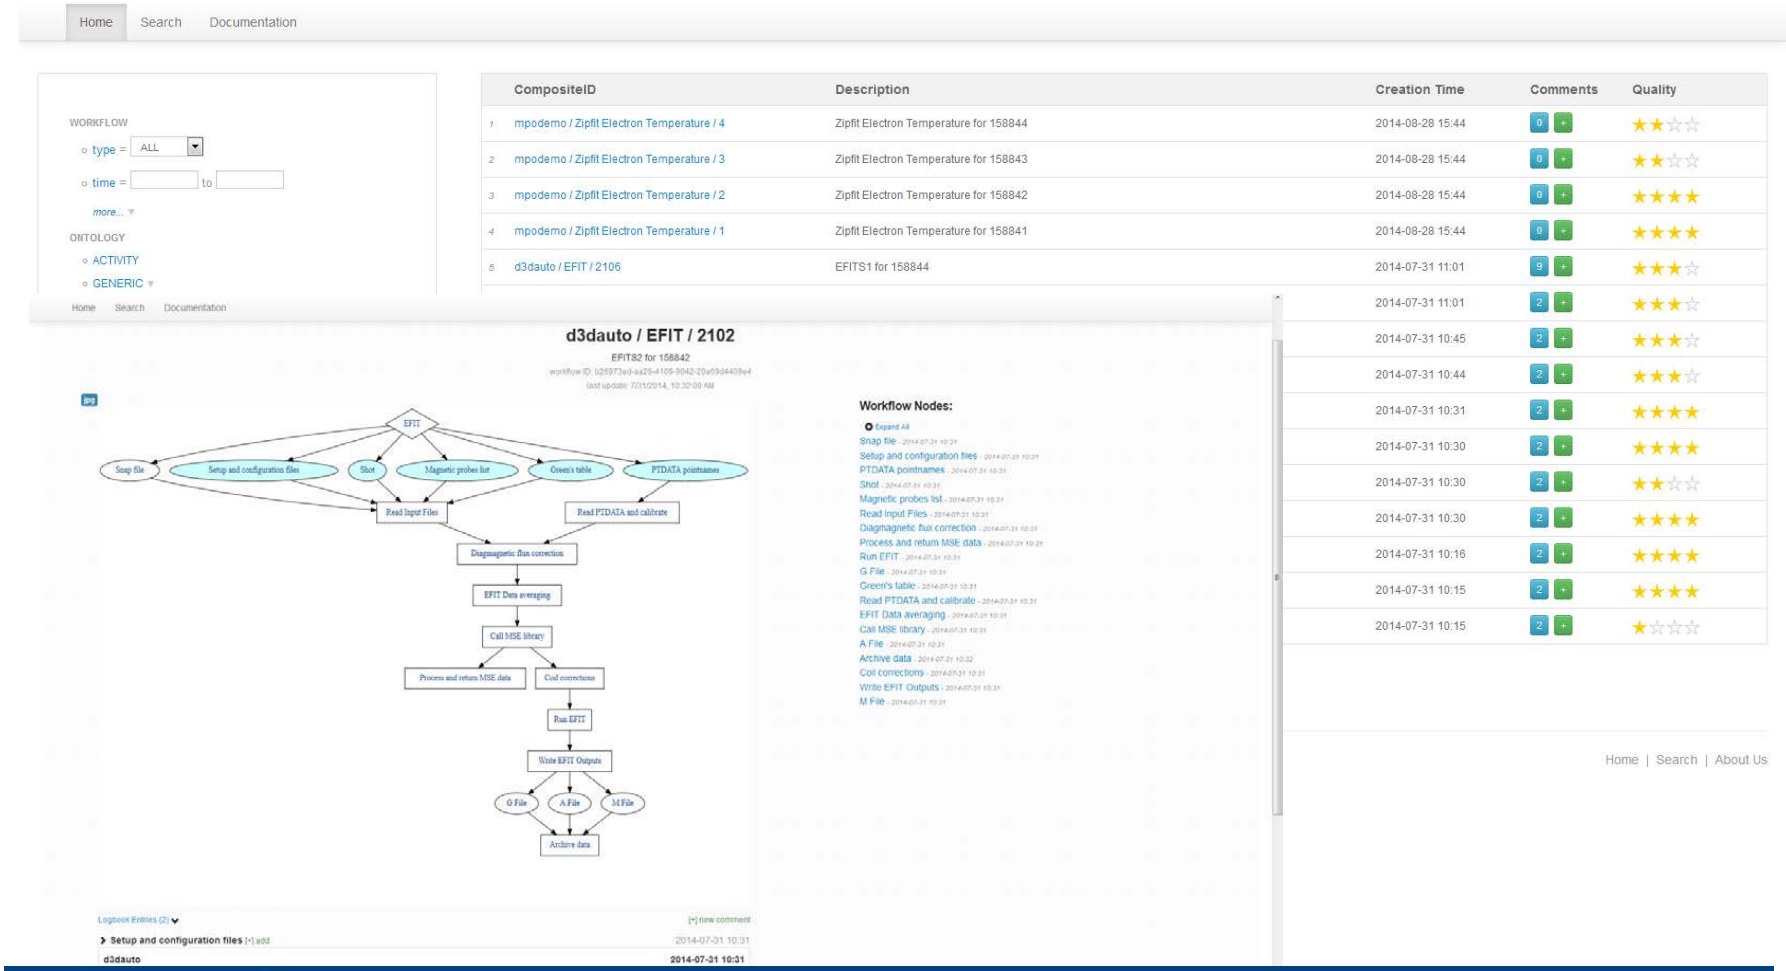

DIII-D NATIONAL FUSION FACILITY

BERKELEY LAB

PSFC

GENERAL ATOMICS

Distributed Computing Design Study

- **Two year 'Design Study' to create a roadmap for distributed computing**
  - Deep analysis of expected future for scientific distributed computing
  - Identify markers and signposts that indicate the roadmap is correct
  - Update Magellan cloud computing activity

Interoperable Design of Extreme-scale Application Software (IDEAS)

Motivation

Enable ***increased scientific productivity,*** realizing the potential of extreme scale computing, through ***a new interdisciplinary and agile approach to the scientific software ecosystem.***

Objectives

Address confluence of trends in hardware and increasing demands for predictive multiscale, multiphysics simulations.

Respond to trend of continuous refactoring with efficient agile software engineering methodologies and improved software design.

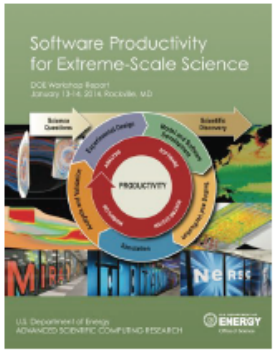

Impact on Applications & Programs

Terrestrial ecosystem ***use cases tie IDEAS to modeling and simulation goals*** in two Science Focus Area (SFA) programs and both Next Generation Ecosystem Experiment (NGEE) programs in DOE Biologic and Environmental Research (BER).

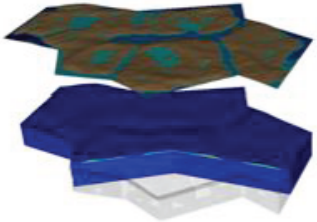

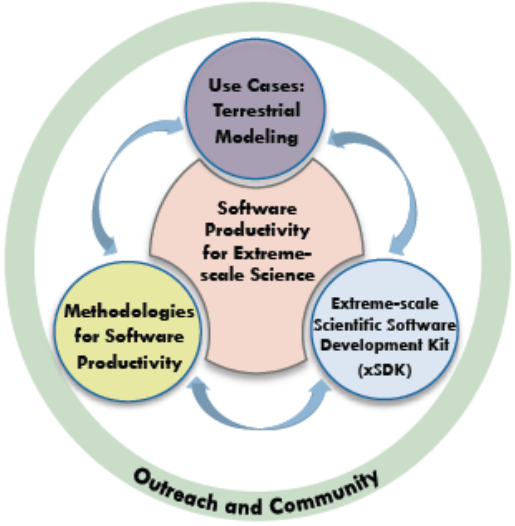

Approach

ASCR/BER partnership ensures delivery of both cross-cutting methodologies and metrics with impact on real application and programs.

Interdisciplinary multi-lab team (ANL, LANL, LBNL, LLNL, ORNL, PNNL, SNL)

ASCR Co-Leads: Mike Heroux (SNL) and Lois Curfman McInnes (ANL)

BER Lead: David Moulton (LANL)

Topic Leads: David Bernholdt (ORNL) and Hans Johansen (LBNL)

Integration and synergistic advances in three communities deliver scientific productivity; outreach establishes a new holistic perspective for the broader scientific community.

Analytical Modeling for Distributed Computing

- **Significantly enhance our ability to predict how scientific workflows perform and to explain workflow behavior on distributed extreme-scale infrastructures**
- **Specific tasks include**
  - Workflow component analytical models.
  - Integrated end-to-end models.
  - Experimental validation and verification.
  - Simple to use tools.
- **Start of a larger program in modeling and simulation for DOE science**

Analytical Modeling for Extreme-Scale Computing Environments

- **Four large projects started in FY14**
  - IPPD: Integrated End-to-End Performance Prediction and Diagnosis for Extreme Scientific Workflows – Kerbson, et.al. (PNNL)
  - Panorama: Predictive Modeling and Diagnostic Monitoring of Extreme Science Workflow – Deelman, et.al. (USC/ISI)
  - Ramses: Robust analytical models for science at extreme scales – Foster, et.al. (ANL)
  - X-Swap: Extreme-Scale Scientific Workflow Analysis and Prediction – Strohmaier, et.al. (LBNL)

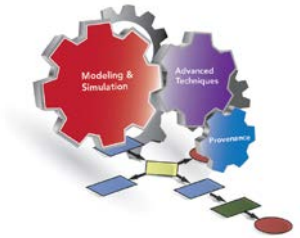

PANORAMA

ramses
robust analytical models for science at extreme scale

X-SWAP PROJECT

http://science.energy.gov/ascr/research/next-generation-networking/am-projects/

Proposed Activities

- **Workshops**
  - Joint NGNS/CS on workflow science (Science Automation)
  - Explore optical networking from computers to networks
  - What else
- **Potential areas**
  - SDN-Enabled Terabit Optical Network for Extreme-scale Science
    - Explore the potential of network programmability in the design of a new generation of federated intelligent terabit optical networks
  - Dynamic Distributed Resource Management
    - A mixture of applied and basic research to support complex computationally intensive and data intensive science experiments

Conclusions

- **NGNS program continues to support both applied and basic research in networking and scientific collaborations**
- **Existing programs augmented by new areas of research**
  - Software Productivity
  - Analytical Modeling of scientific workflows
- **Future programs being developed to address challenges emerging from current work**
  - What is the top question your current research motivates you to now ask?

Supplemental Slides

Existing NGNS Programs

- **FY11 – Networking (8 funded)**
  - Innovative approaches to develop federated terabit network tools and services that address DOE's emerging network challenges
- **FY12 – Collaboration (10 funded)**
  - Transformative approaches to understanding and/or enabling scientific collaborations at a scale not possible with today's knowledge or using current Internet-based services and tools
  - http://science.energy.gov/ascr/research/next-generation-networking/extremescalecollaboration/
- **FY13 – Networks (7 funded)**
  - Conduct research on terabit data-aware network technologies and storage system management to address scientific Big Data infrastructure challenges
  - Explore innovative theoretical frameworks, network-aware data management concepts, and storage system performance optimization models

FY14 New Starts

- **Design Study**
- **Software Productivity**
- **Analytical Modeling of Science Workflows**

SDN-Enabled Terabit Optical Network For Extreme-scale Science

Focus Ares

Current Approach

SDN-Enabled Internet

Internet Apps

Internet Function | Internet Function

Open Interfaces

Network Operating System Control Plane

Open Interfaces

Data Plane

SDN-Enable Science Networking

Science Apps | Science Apps | Science Apps

Workflow interface

Terabits HPC-Networking Function | Terabits HPC-Networking Function

Open Interfaces

Network Operating System Control Plane

Open Interfaces

Multi-layers Data Plane

Network Intelligence

Current Networking trends: Network Control and traffic (data) flows are hardwired into a proprietary hardware

Emerging SDN paradigm; Separation of network Control and traffic (data) forwarding

SDN-Enabled Science Terabit Networking Adaptation and extension of SDN to distributed data-intensive science application