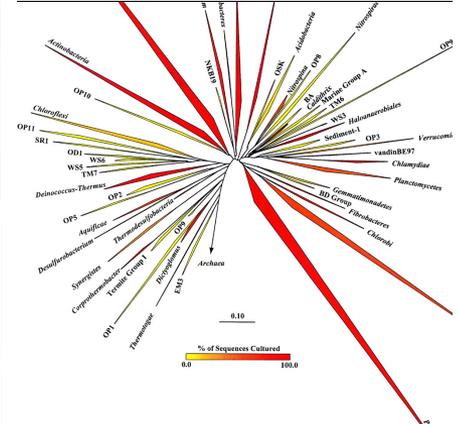
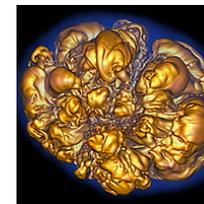
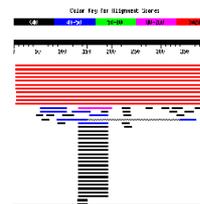
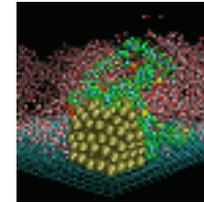
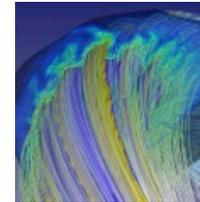
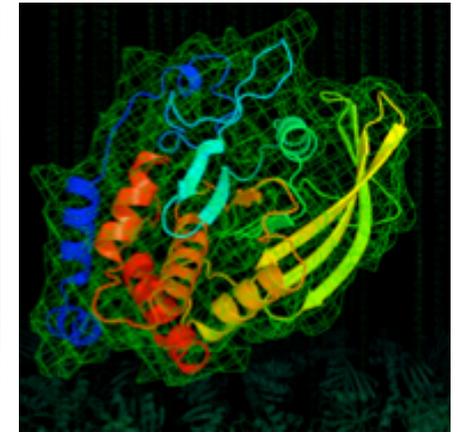


Lessons Learned – NERSC/JGI Partnership Kjiersten Fagnan, NERSC User Services/JGI



Joint Facilities User Forum on Data Intensive Computing

June 17, 2013

Outline

- **Overview of NERSC/JGI Partnership**
 - DOE JGI background
 - Team overview
 - Compute resources
- **Computing Strategic Plan**
 - JGI Goals
 - NERSC Goals
- **Lessons Learned**

DOE Joint Genome Institute

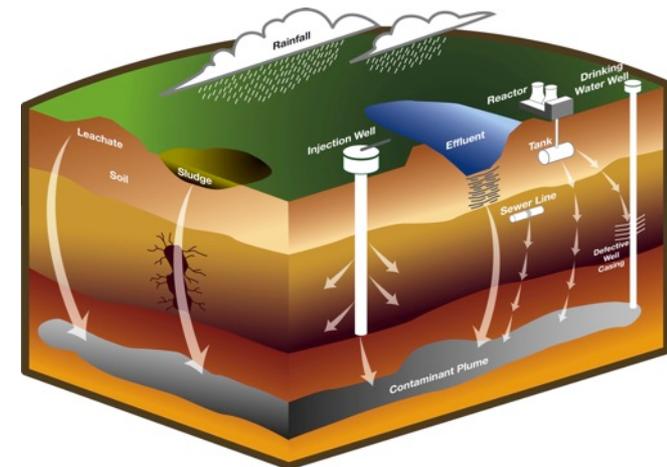
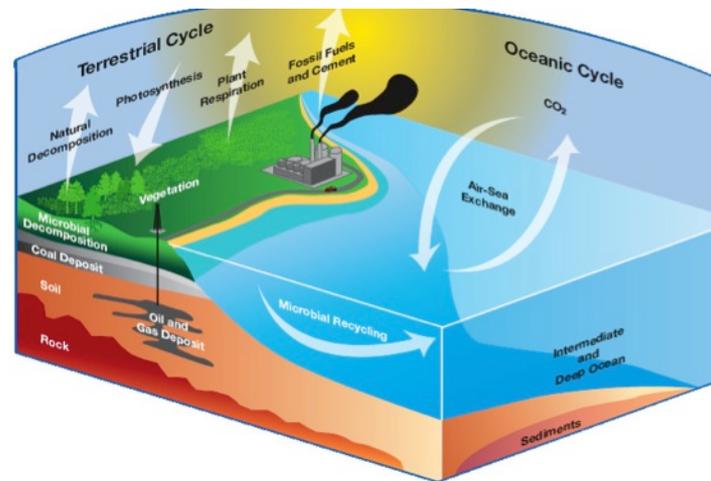
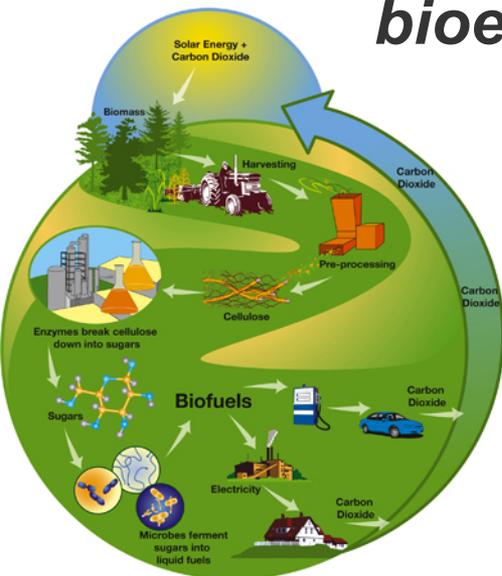


- Walnut Creek, CA facility opened in 1999
- 250 employees
- ~\$69M annual expenses

Mission:

DOE JGI, Serving as a genomic user facility in support of the DOE missions:

bioenergy, carbon cycling, & biogeochemistry



JGI Programs & Infrastructure

DOE
Mission
Areas



Bioenergy



Carbon Cycling



Biogeochemistry

NERSC/JGI Team

- User Services (2)
- Networking and Security (2)
- File Systems (1)
- Computational Systems Group (1)



Structure of the Genepool System

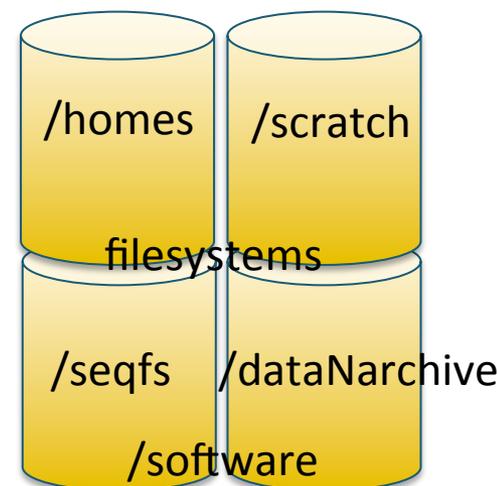
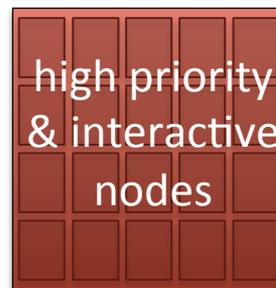
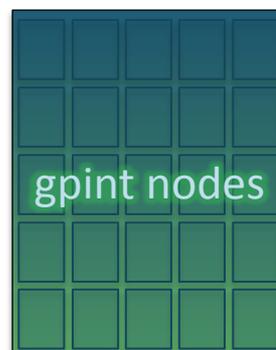
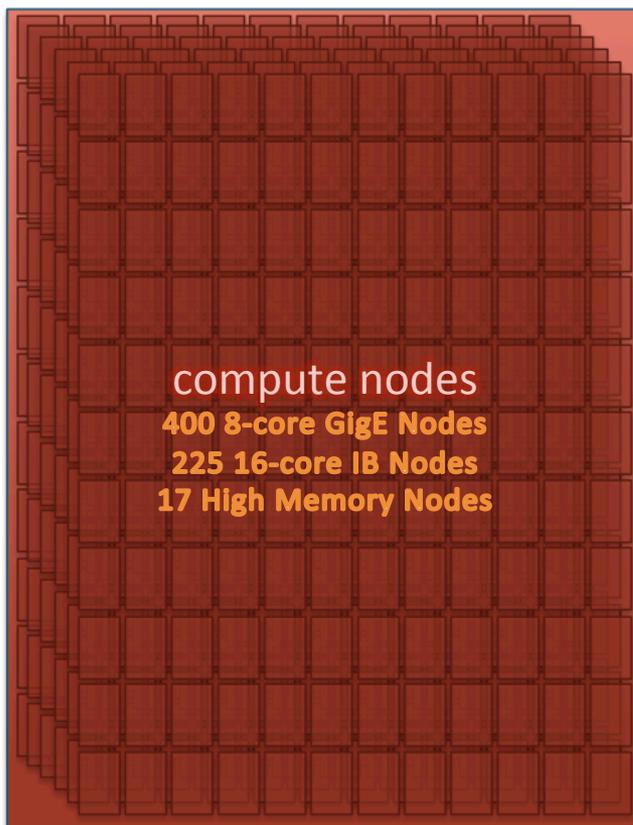
User Access

- Command Line
- Scheduler
- Service

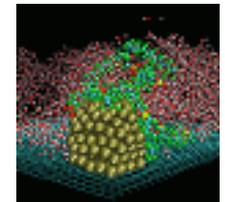
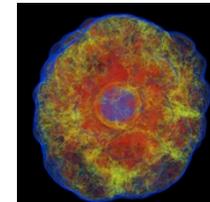
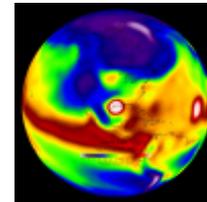
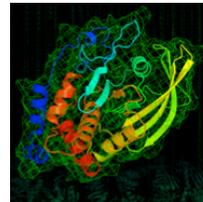
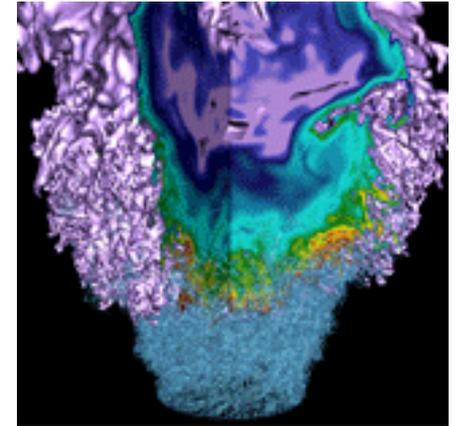
```
> ssh genepool.nersc.gov
```



 <http://...jgi-psf.org>



JGI Strategic Computing Plan



Computing Strategic Plan

1. Leveraging Computing Resources

- move computations to NERSC HPC resources

NERSC/JGI
partnership

NERSC Strategic Plan: Partner with DOE facilities to better understand data-intensive workloads

6. Training in Programming for the HPC Environment

NERSC/JGI
partnership

Computing Initiative GOALS

- **Learn more about the data-intensive workloads used for production and R&D at the JGI**
- **Recommend hardware and software needed to support these workloads**
- **Work with JGI staff to analyze pipeline/software efficiency**
- **Support JGI in making progress on strategic computing initiatives**

JGI/NERSC HPC Initiatives

Optimization of the RQC/Jigsaw pipeline for HPC environments

Leads: James Han, Bryce Foster, Kjersten Fagnan

Goal:

Port RQC and Jigsaw pipelines to NERSC Cray Systems (FY14)

JGI Strategic
Goals 1,2,3,6

Annotation and Comparative Analysis of Metagenomes

Leads: Amrita Pati, Kjersten Fagnan

Goals:

1. Port metagenome annotation pipeline in production-mode to NERSC Cray Systems (FY14)
2. Deploy framework for comparison of metagenomes on NERSC Cray Systems(Prototype, benchmarks FY14; Production (FY15)

JGI Strategic
Goals 1,2,4,6

Port metagenome assembly pipeline to Hopper/Edison

Leads: Alex Copeland, Kjersten Fagnan

Goal:

1. Investigate alternative approaches to metagenome assembly using distributed (MPI-based) computing on commodity hardware (benchmarks, prototype FY14)
2. Deploy production pipeline (FY15)

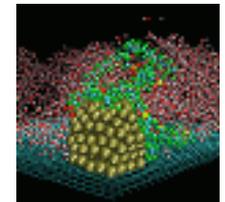
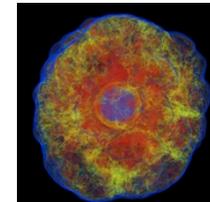
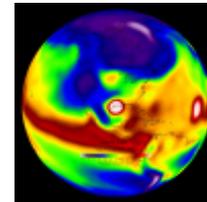
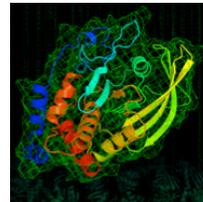
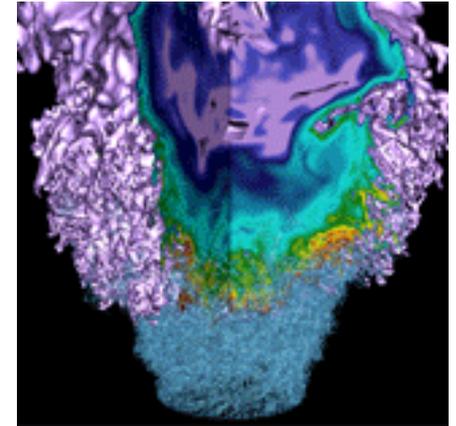
JGI Strategic
Goals 1,2,4,6

Lessons Learned: Migrate Pipelines to Single Framework

- **Functional Annotation and ReadQC pipelines have been (mostly) migrated to FireWorks (still testing)**
 - facilitates running software on multiple clusters through queue adapters
 - software is currently used heavily by Materials Genome Project
- **For JGI-specific implementation, integrate data ingest and output with JAMO through JAT**



Lessons Learned



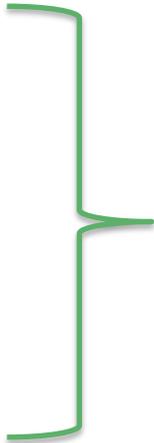
Lessons Learned - Scheduling

Heterogeneous jobs are difficult to schedule efficiently

UGE Logs

Procmon

Utilization



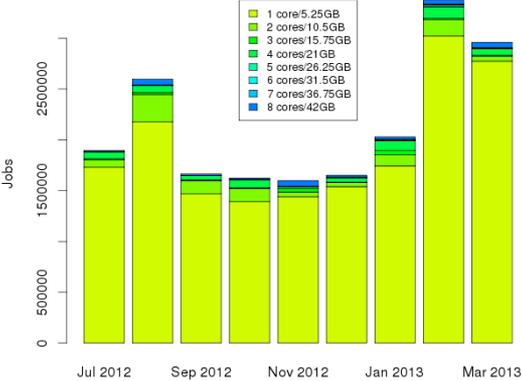
Analysis



Action: Used to inform NERSC scheduler procurement

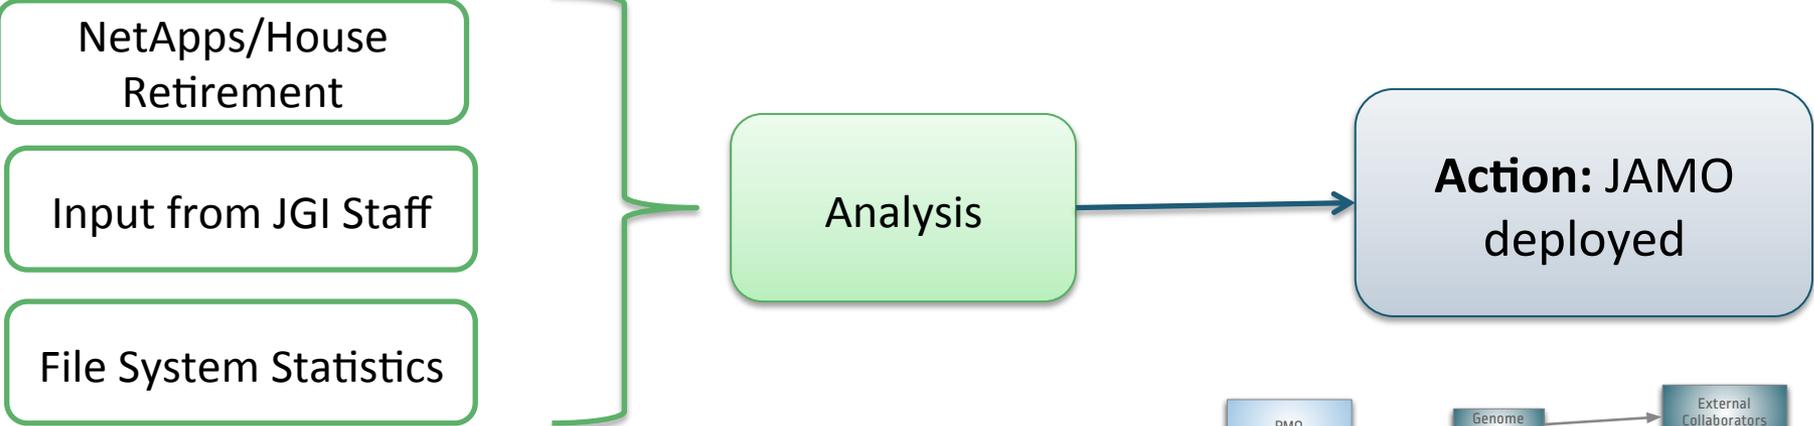
Contributors: D. Jacobsen, L. Pezzaglia, J. Srinivasan

Total Job Counts for Genepool High Throughput Nodes

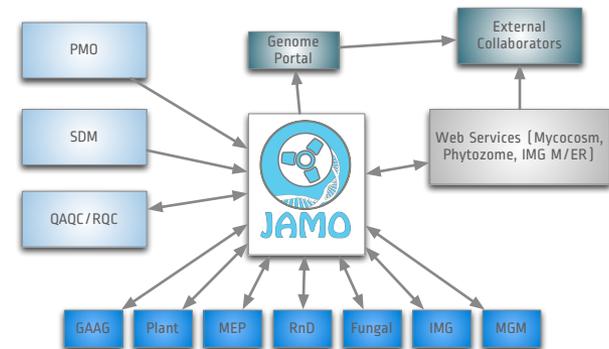


Lessons Learned – Data Management

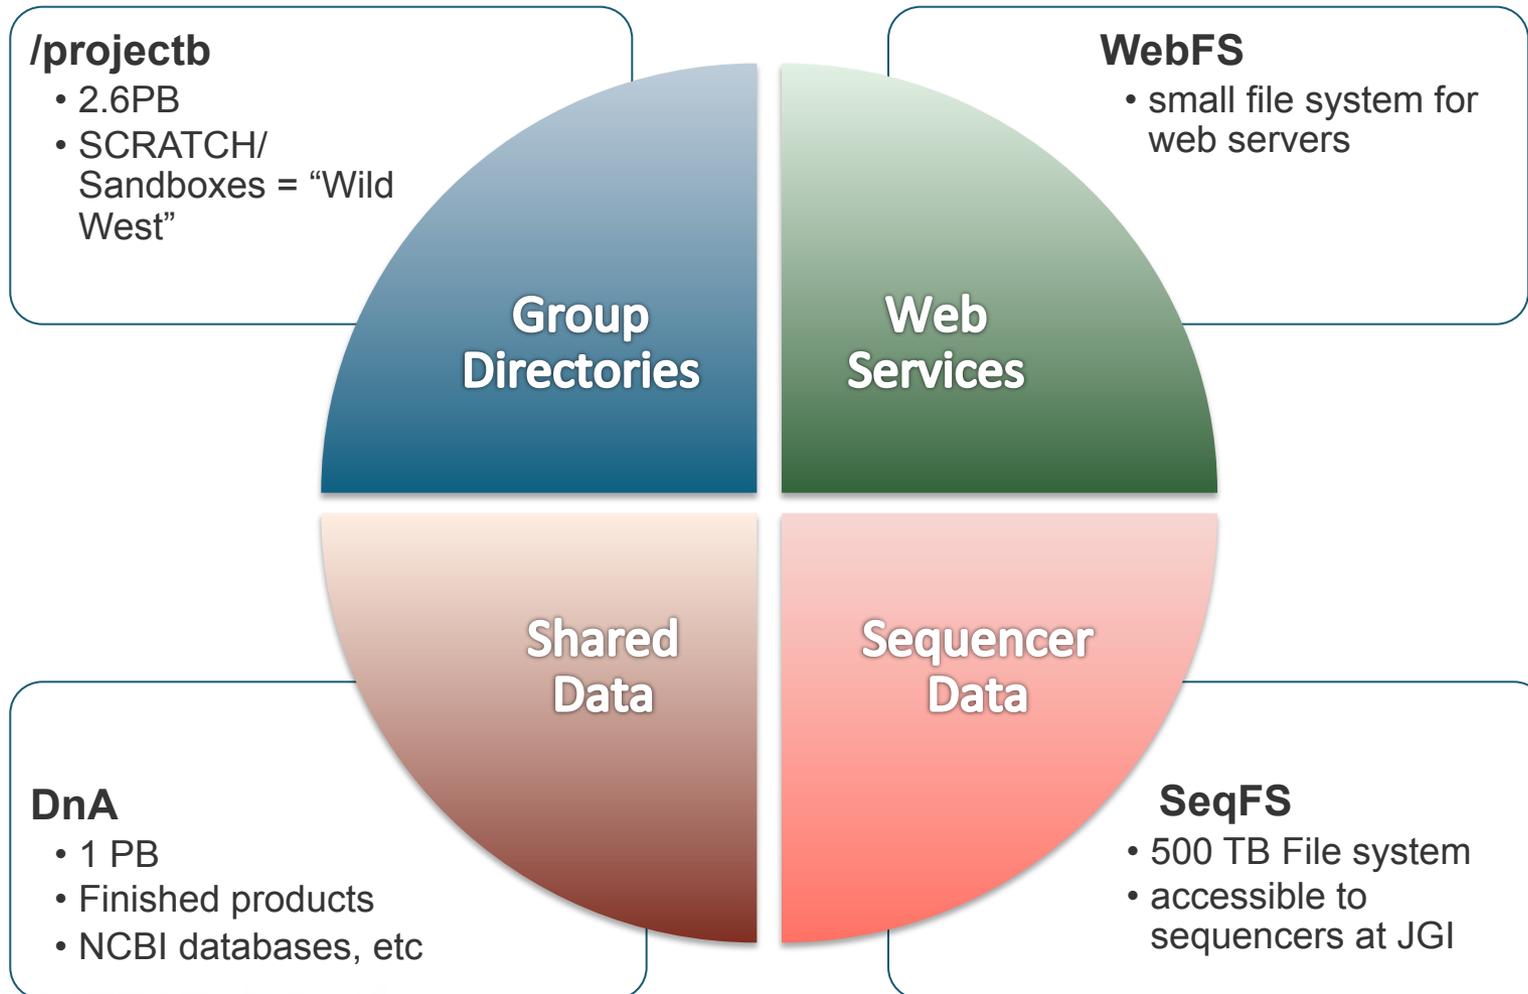
Volumes of JGI data benefit from hierarchical data management system



Contributors: C. Beecroft, A. Boyd, A. Copeland, K. Fagnan, J. Hick, N. Balthesar, W. Hurlburt



JGI File Systems



Lessons Learned – File Systems

Projectb is not configured properly for JGI workloads

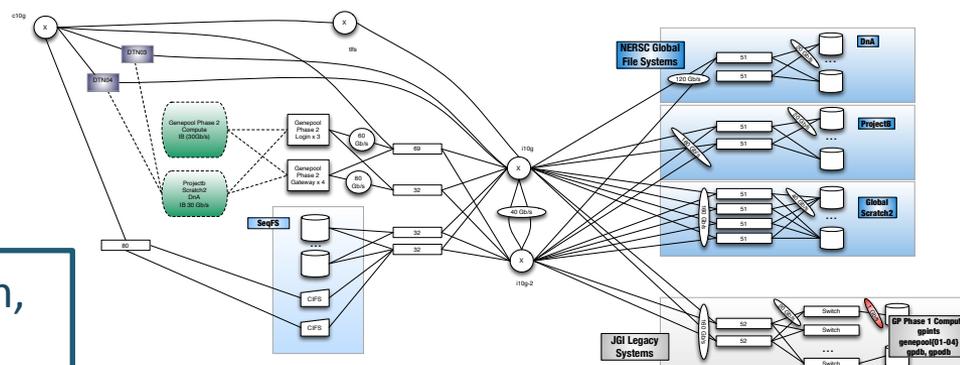
Canary Jobs

HPC Initiatives

File system
Monitoring

Analysis

Action: changing configuration during CRT move



Contributors: R. Egan, A. Copeland, K. Fagnan, D. Jacobsen, J. Hick, R. Cheema, S. Canon

Lessons Learned – Pipelines & Services

Pipelines and services have different, complex hierarchical dependencies; hard to guarantee uptime

Procmon - monitor dependencies

File system dependencies

Database dependencies

Analysis

Action: Develop services dashboard for improved monitoring

Contributors: D. Jacobsen, B. Yumae, L. Pezzaglia

Lessons Learned – Pipelines & Services

Pipelines and services have different, complex hierarchical dependencies; hard to guarantee uptime

Procmon - monitor dependencies

File system dependencies

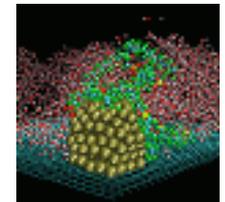
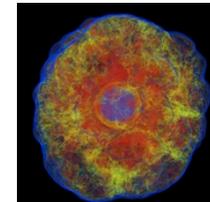
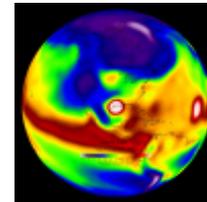
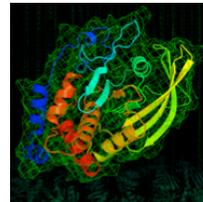
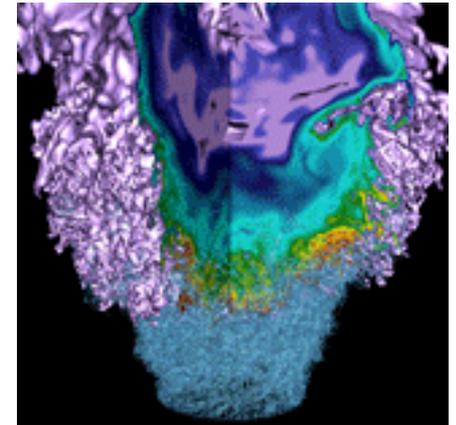
Database dependencies

Action: Develop

Tools and pipelines from years ago will not continue to work efficiently and we need to invest human effort to
REFACTOR

Contributors: D. J. Pezzaglia

Summary



JGI/NERSC Partnership & Lessons Learned

- **Significant time and effort needed to understand their compute and data needs**
- **Collaboration is the key to success**
 - NERSC staff sit at the JGI, JGI staff sit at NERSC
 - HPC Initiatives
 - JGI/NERSC Coordination Committee
- **Heavy dependence on third party tools means we should focus on data movement and workflows for optimizing their computations**
- **JGI has a relatively small, but increasing data footprint**
 - data management, sharing data with collaborators, reproducible research
- **Parallelism and new data formats can be exploited to improve I/O performance, but community needs to buy into the changes**