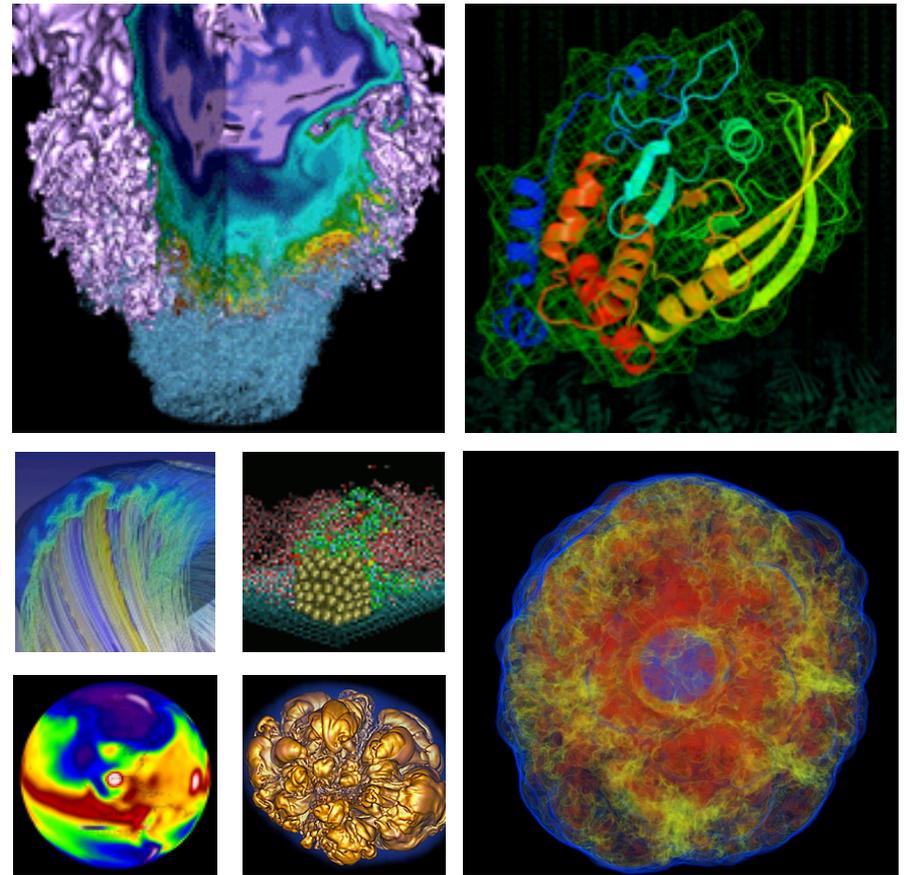


# Biosciences Computing and Storage for JGI



**Shane Canon**  
**Group Leader for Technology  
Integration**

February 12, 2013

- 1 -

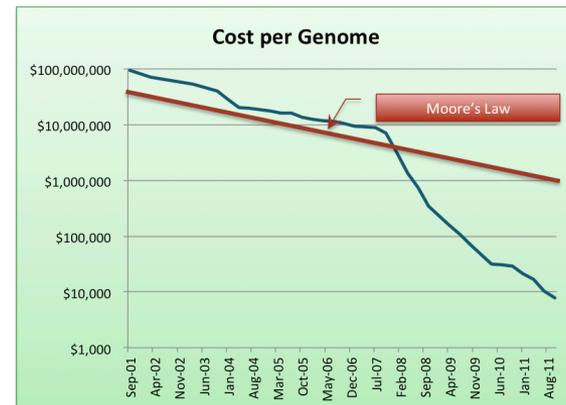
- **Biofuels**
  - Engineering better plants for biofuels
  - Engineering microbes to convert cellulose into ethanol
- **Bioremediation and Restoration**
  - Finding microbes to aid in cleanup
- **Restoration and Carbon Cycle**
  - Role of Microbes in restoration and impact on carbon sequestration



# History and Motivation



- NERSC and JGI formed partnership in 2010
- JGI was struggling to keep pace with the growing demand from NGS and NERSC wanted to grow its “Big Data” capabilities



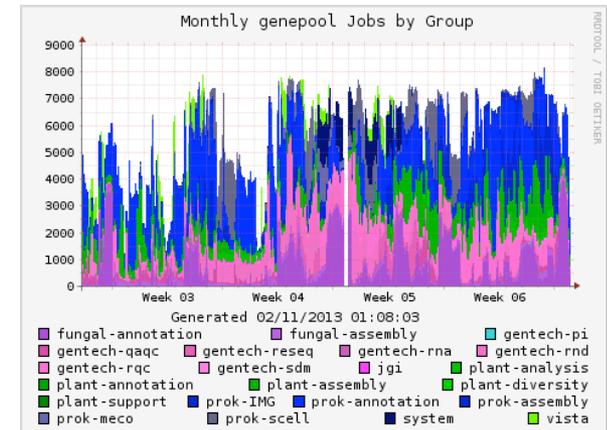
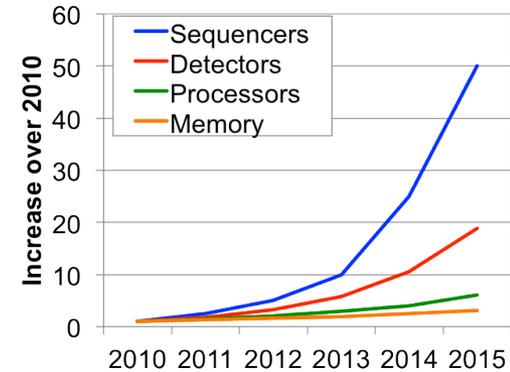
Source: National Human Genome Research Institute

- **Wide diversity in application characteristics**
  - Long running, large memory for assembly
  - high-throughput for annotation
  - Growing number of MPI apps for assembly and metagenome analysis
  - Wall times vary from seconds to months
- **Complex workflows**
  - Deeply nested dependencies
  - Large task arrays with serial sections
- **Heavy I/O**
  - Reading and comparing against large references
  - Reading large raw sequence files
  - Writing and reorganizing large outputs

# Resulting Challenges



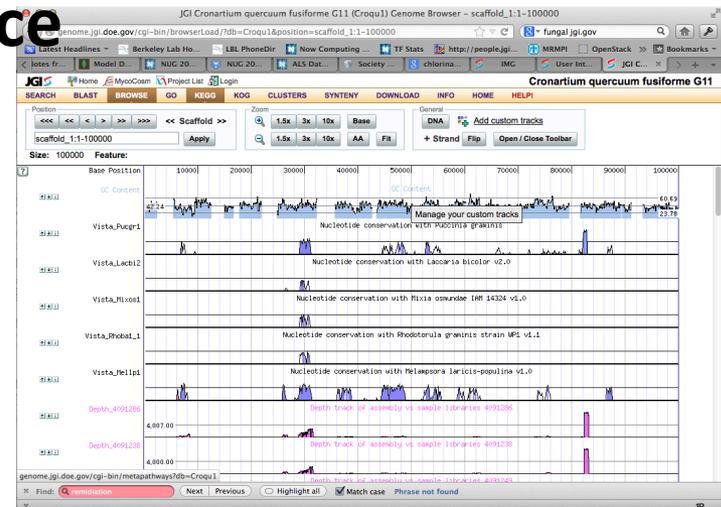
- **Resource Requirements**
  - Needs outpacing Moore's Law
  - Need to take full advantage of the available resources and improve efficiency
- **Scheduling**
  - Insuring fairness across groups
  - Effectively scheduling small and big, short and long jobs
- **File Systems**
  - Lots of data
  - Lots of files
  - Lots of problems



# Other Aspects



- Heavy reliance on database systems (postgres, MySQL, and Oracle)
- Growing use of Hadoop Ecosystem (annotation pipelines and Biopig extensions to Pig)
- Heavy use of Web-based Science Gateways by end users
  - IMG collection
  - Phytozome
  - Portal



# Compute and Storage Resources

- **Genepool**
  - Heterogeneous Mixture of “standard memory” and large memory nodes
  - Newest hardware is Intel Sandy-Bridge, 128 GB per node
  - Total of ~8500 cores (776 nodes)
- **Storage**
  - Nearly 5 PBs of storage
  - Mixture of legacy NFS and GPFS
  - Growing use of HPSS



# Accomplishments

---

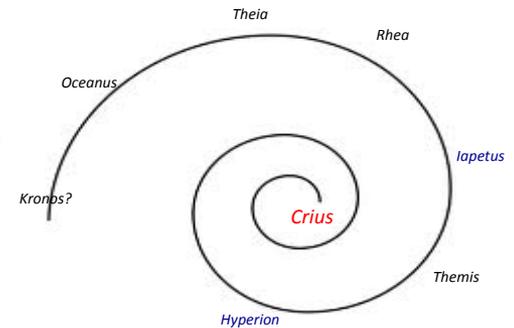


- Stabilize Initial Hardware
- Gain Understanding of workload
- Integrate Compute Resources
- Improve Utilization of Resources
- Build Up User Services and Support
- Train Users
- Improve Application Efficiency
- Improve I/O Performance and Reliability

# Accomplishments



- Stabilize Initial Hardware
- Gain Understanding of workload
- Integrate Compute Resources
- Improve Utilization of Resources
- Build Up User Services and Support
- Train Users
- Improve Application Efficiency
- Improve I/O Performance and Reliability



# State of the Partnership



- **JGI's workload (and similar communities) are challenging. Different set of requirements compared to HPC and PDSF (HEP/NP)**
- **NERSC continues to work with JGI to hone the scheduling configuration**
- **Continuing to modify file system configuration to better meet JGI's requirements and improve robustness**
- **We're not there yet. But we are getting close.**



**National Energy Research Scientific Computing Center**