

Supercomputing resources for open, accessible, reproducible genomic analysis

Philip Blood

Senior Computational Scientist

Pittsburgh Supercomputing Center

blood@psc.edu


Pittsburgh Supercomputing Center



- Joint effort of **Carnegie Mellon University** and the **University of Pittsburgh**
- **National leadership since 1986** in:
 - High-performance and data-intensive computing
 - Data management technologies
 - Software architecture, implementation, and optimization
 - Networking and network optimization
 - Ground-breaking science and engineering
 - Enabling researchers nationwide
- **Supported by:** NSF, NIH, the Commonwealth of Pennsylvania, DOE, DoD, foundations, and industry
- PSC computing resources made available to researchers world-wide through **XSEDE** project

XSEDE: Making Advanced Infrastructure Freely Available to Researchers

- **\$230 million** invested by NSF over 10 years (5 years x 2 rounds)
- Currently in second 5-year funding going through 2021
- XSEDE does not fund major hardware
- Makes **resources** at NSF-funded centers **freely** available to researchers and educators involved in open scientific research
- Engage **research communities** that have not traditionally used national resources like XSEDE
- Major **investment in people** to support researchers and educators in using the hardware (~40 FTE)



XSEDE
Extreme Science and Engineering
Discovery Environment

www.xsede.org

Extended Collaborative Support Services (ECSS)

Provides collaboration between XSEDE staff and users. The objective of the program is to enhance effectiveness and productivity.

- **Extended Support for Research Teams (ESRT)**
 - In depth collaboration with a research group
- **Extended Support for Community Capabilities (ESCC)**
 - Support for community codes, initiatives
- **Extended Support for Science Gateways (ESSGW)**
 - Support for building, deploying science gateways




National Center for Genome Analysis Support

- Collaborative grant between PSC and Indiana University
- Funds bioinformaticians and computational scientists to work with biologists
- Specific focus on helping biologists do large-scale genomics analysis on advanced computational resources
- Address challenges with data, software, hardware in genomics pipelines
- Complements XSEDE collaborative support



Broadening access to large-scale genome analysis



- Providing computational expertise to help biologists take advantage of latest advances in genome analysis
- Optimizing genomics tools for high performance computing systems
 - 4x performance increase of Trinity RNA-Seq through collaboration between NCGAS and Trinity developers
 - Fixed ALLPATHS-LG to enable assembly of very large genomes (e.g. wheat genome)
- Creating and maintaining web-based portals for biologists to run advanced tools on HPC systems

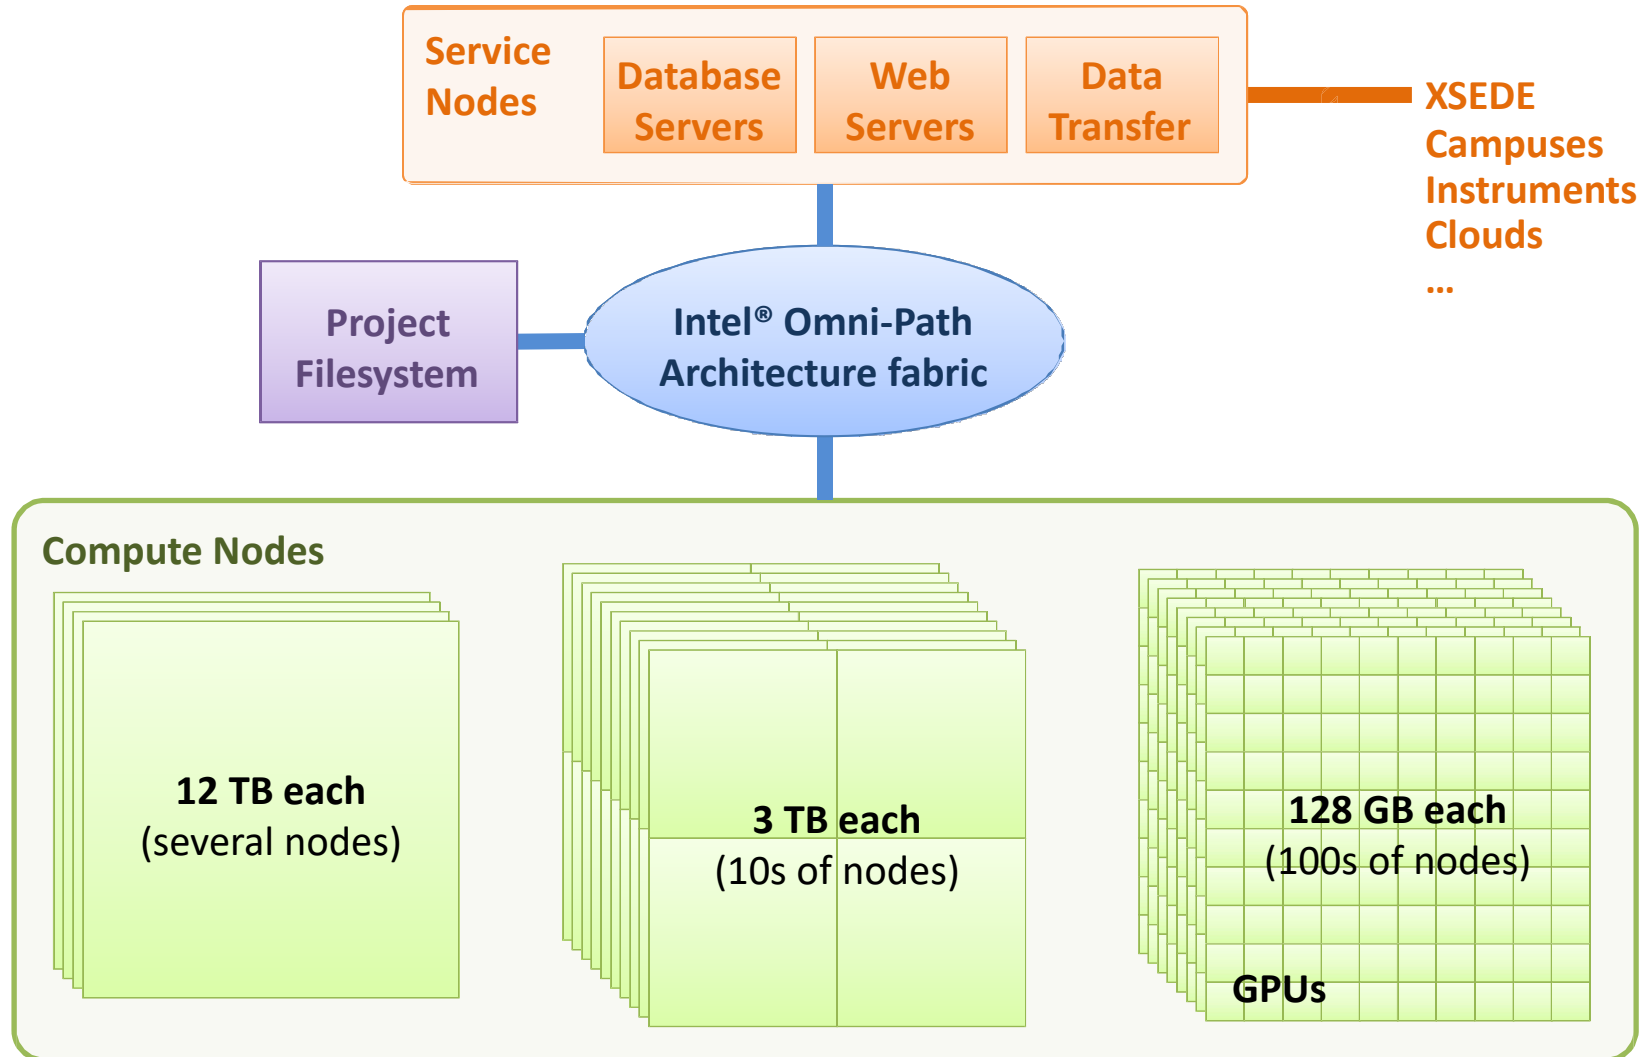


BRIDGES

A PITTSBURGH SUPERCOMPUTING CENTER RESOURCE

Connecting Researchers, Data & HPC

PSC's Bridges: High-Level Architecture

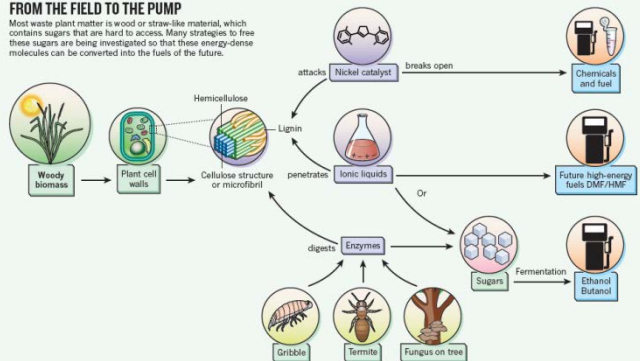


All nodes have both fast local disk and access to shared filesystems

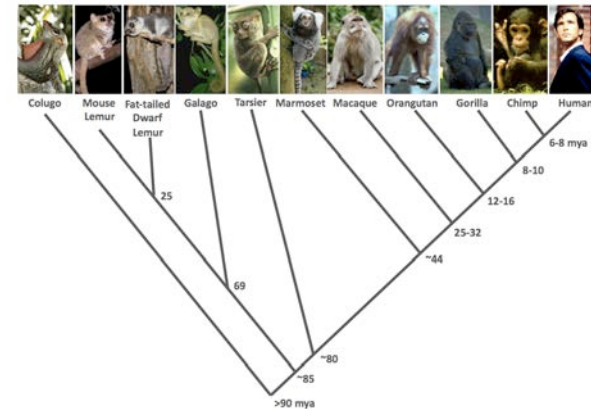
Flexible architecture for genomics applications

FROM THE FIELD TO THE PUMP

Most waste plant matter is wood or straw-like material, which contains sugars that are hard to access. Many strategies to free these sugars are being investigated so that these energy-dense molecules can be converted into the fuels of the future.



Reprinted by permission from Macmillan Publishers Ltd: [Nature 474, S12-S14](#), copyright 2011.



Metagenome/Plant Genome Assembly

4 x 12 TB RAM x
288+ cores:
Extreme-scale
Shared Memory



42 x 3 TB RAM x 64+
cores:
Very Large Memory
+ Throughput



Large Transcriptome Assembly

126 x 1 TB RAM x
20+ cores:
Large Memory + High
Throughput

...plus over 20,000 cores on regular memory nodes for other pipeline steps

Great for large-scale benchmarking projects!

Characterizing Diverse Microbial Ecosystems From Terabase-Scale Metagenomic Data

Brian Couger, Oklahoma State University

- ✓ Assembled 11 metagenomes sampled from diverse sources, comprising over 3 trillion bases of sequence data
 - Including a recent massive assembly of **1.6 Tbp of metagenomic data** from an oil sands tailings pond, a bioremediation target
 - Excellent performance of MPI-based Ray assembler on large memory nodes: **assemble 1.5 Tbp in 16 hours**
 - Analysis of assembled data in progress to characterize organisms present in these diverse environments and identify new microbial phyla

★ User Support

- PSC staff helped user effectively download, prepare, and manage terabytes of input data for metagenome assemblies
- PSC staff installed and configured new metagenome analysis tools and provided initial job scripts to help the user run effectively on *Bridges*



Oil sands tailings pond. By NASA Earth Observatory - <http://earthobservatory.nasa.gov/IOTD/view.php?id=40997>, Public Domain, <https://commons.wikimedia.org/w/index.php?curid=8346449>

Optimization of RNA-Seq Analysis for Analyzing Disease Resistance in Shellfish

Bassem Allam and Sleiman Bassim, Stony Brook University

- ✓ Assembled ten 400 million read and five 1 billion read oyster transcriptomes in record times
 - *Bridges*-optimized Trinity assemblies reduced times from 3 days to 39 hours and from 5 days to **65 hours** for 400 million reads and **1 billion reads** respectively
 - Availability of many 3 TB nodes on *Bridges* allowed rapid testing of optimal methods for accuracy and speed
 - Research team will perform **hundreds of assemblies** on *Bridges* this year using optimized methods

★ User Support

- PSC staff provided initial Trinity script suitable for large assemblies
 - Using RAM disk for I/O
 - Using appropriate performance settings for each phase
- PSC staff helped user do initial assemblies on Greenfield last year and write a successful *Bridges* proposal, enabling the group to make rapid progress on *Bridges*



Pacific oysters (*Crassostrea gigas*), [CC BY-SA 2.0](https://commons.wikimedia.org/wiki/File:Pacific_oysters.jpg).
https://commons.wikimedia.org/wiki/File:Pacific_oysters.jpg

De Novo Assembly of the Sumatran Rhino Genome

James Denvir and Swanthana Rekulapa, Marshall University

- ✓ First assembled the 1 gigabase Narcissus flycatcher (*Ficedula narcissina*) genome
 - On the users' local resources, they could only assemble $\frac{1}{3}$ of the data due to memory limitations, taking 16 hours
 - Assembly of *all* data on a 3TB node of *Bridges* required 1.5 TB of memory and only **6.6 hours**, which was at least 3-4× faster than they had anticipated
- ✓ Then assembled the 3 gigabase Sumatran rhino (*Dicerorhinus sumatrensis*) genome
 - Required **1.9 TB** of memory and completed in only **11 hours**, again 3-4× faster than the users had expected
- ★ User Support
 - PSC staff developed a job script for DiscoverDeNovo to use *Bridges*' features for optimal performance
 - Using RAM disk for I/O
 - Increasing DiscoverDeNovo memory settings to take advantage of large memory
 - PSC staff also helped diagnose and fix an issue with input data causing initial job failure



Narcissus Flycatcher (*Ficedula narcissina*) in Osaka, Japan, Kuribo. [CC BY-SA 2.0.
https://commons.wikimedia.org/wiki/File:Narcissus_Flycatcher-cropped.jpg](https://commons.wikimedia.org/wiki/File:Narcissus_Flycatcher-cropped.jpg)




Sumatran Rhinoceroses at the Cincinnati Zoo & Botanical Garden, Charles W. Hardin. [CC BY 2.0.
https://upload.wikimedia.org/wikipedia/commons/3/33/Sumatran_Rhino_2.jpg](https://upload.wikimedia.org/wikipedia/commons/3/33/Sumatran_Rhino_2.jpg)

Creation of World's Largest K-mer Database

Rachid Ounit and Chris Mason, Cornell University

- ✓ Created a database of 153 billion species-specific nucleotide sequences (k-mers)
 - Analyzed entire **NCBI Reference Sequence (RefSeq) archive** containing over 15K species to create the k-mer database
 - Required massive in memory hash table
 - **Computation took 24 days and 4.8 TB of RAM on *Bridges* 12 TB node**
 - Allows rapid identification and classification of species in metagenomics samples
- ★ **User Support**
 - Calculation performed by user who had not previously used XSEDE resources
 - PSC staff guided user through successful calculation
 - Discussed application and helped develop best strategy
 - Helped write and debug job script
 - Created reservation and helped monitor progress of long-running job
 - Helped group write proposal for more time on *Bridges*



By Madprime - Own work, CC BY-SA 3.0,
<https://commons.wikimedia.org/w/index.php?curid=2121945>

Supercomputing

Open

Accessible

Reproducible

User-Friendly HPC & Data Analytics

- **Interactivity** is the feature most frequently requested by nontraditional HPC communities and for doing data analytics and testing hypotheses.
- **Gateways and tools for gateway building** will provide easy-to-use access to Bridges' HPC and data resources.
- **Database and web server nodes** will provide persistent databases to enable data management, workflows, and distributed applications.
- **Popular programming languages & software environments** will let users scale applications and workflows.
- **Virtualization** will allow users to bring their particular environments and provide interoperability with clouds.

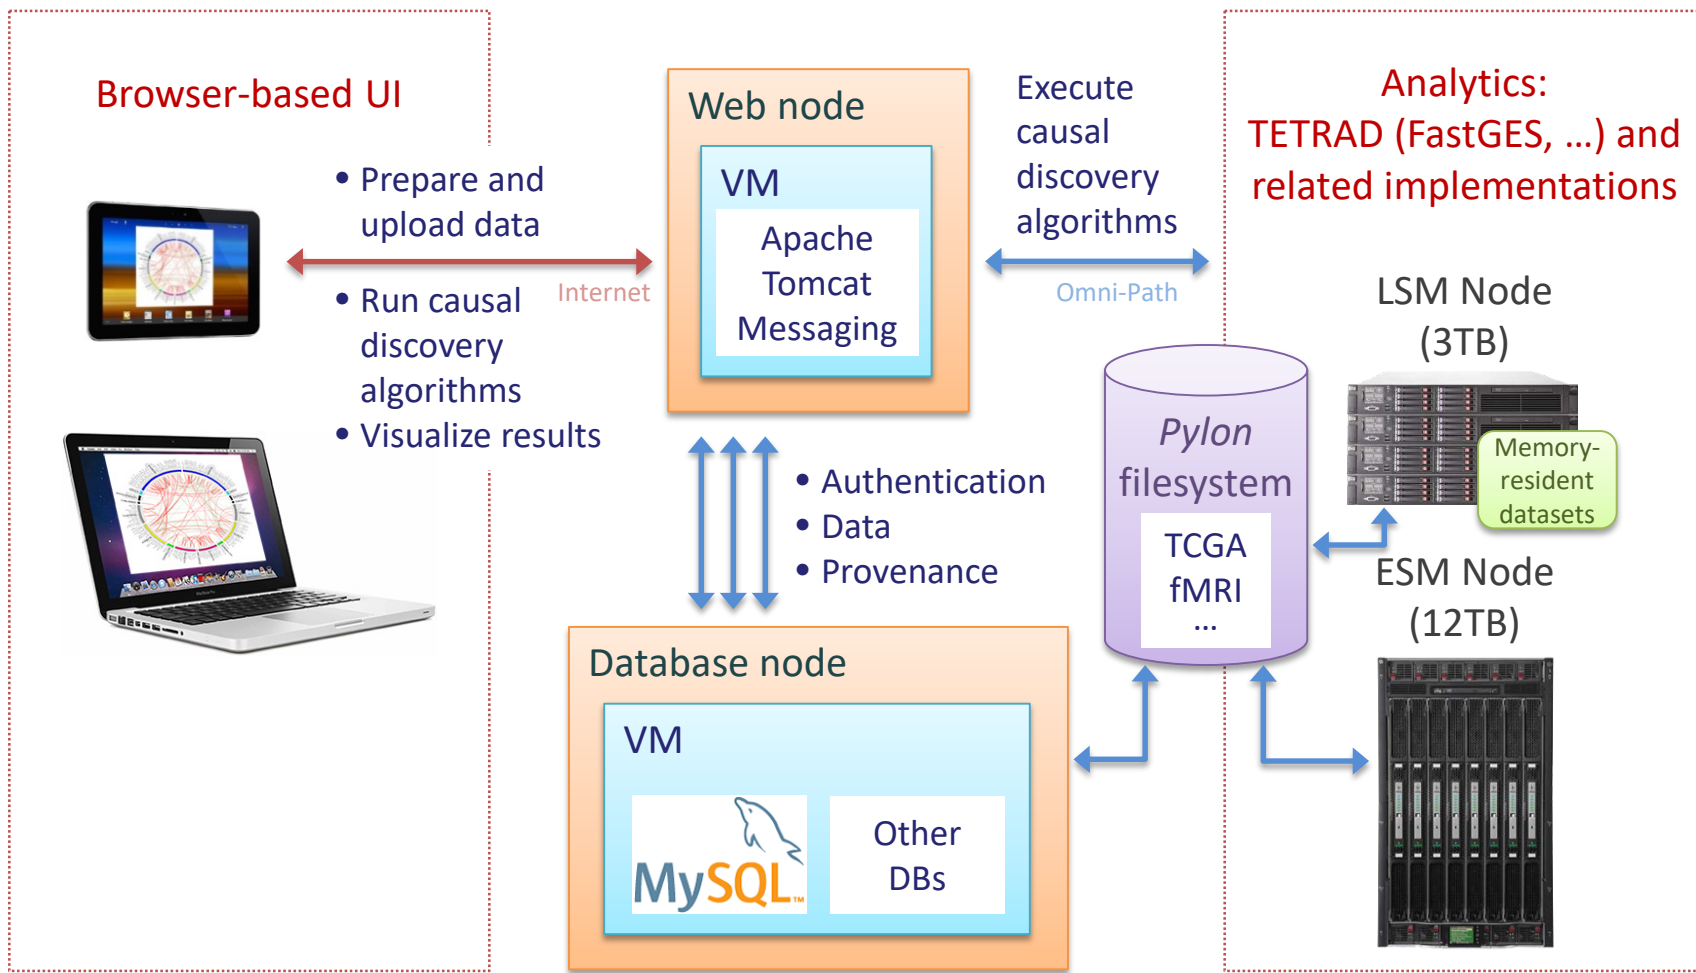
Database and Web Server Nodes

- Dedicated database nodes will power persistent relational and NoSQL databases
 - Support data management and data-driven workflows
 - SSDs for high IOPs; RAIDed HDDs for high capacity



- Dedicated web server nodes
 - Enable distributed, service-oriented architectures
 - High-bandwidth connections to XSEDE and the Internet

Example: Causal Discovery



XSEDE “Science Gateways” for Bioinformatics

- Science Gateways enable communities to access XSEDE resources through web interfaces, providing ease-of-use and community-specific tools & resources
- The CIPRES gateway lets biologists run parallel versions of **BEAST, GARLI, MrBayes, RAxML, & MAFFT** via a browser interface on SDSC’s Trestles supercomputer, an XSEDE resource
- Hosts optimized versions of RAxML and MrBayes (superior parallel scaling)

*Adapted from information provided by Wayne Pfeiffer, SDSC

Tools

search tools

Fetch Alignments/Sequences

NGS: QC and manipulation

NGS: DeepTools

NGS: Mapping

NGS: RNA Analysis

NGS: SAMtools

NGS: BamTools

NGS: Picard

NGS: VCF Manipulation

NGS: Peak Calling

NGS: Variant Analysis

NGS: RNA Structure

NGS: Du Novo

NGS: Gemini

NGS: Assembly

Trinity (Beta) De novo assembly of RNA-Seq data Using Trinity on PSC's Bridges

Operate on Genomic Intervals

Statistics

Graph/Display Data

Phenotype Association

BEDTools

Genome Diversity

EMBOSS

Trinity (Beta) De novo assembly of RNA-Seq data Using Trinity on PSC's Bridges (Galaxy Version 0.0.1) Options

Paired or Single-end data?

Paired

Left/Forward strand reads

No fasta or fastq dataset available.

Right/Reverse strand reads

No fasta or fastq dataset available.

Strand-specific Library Type

Not set

Group pairs distance

500

Maximum length expected between fragment pairs

Path reinforcement distance

75

Minimum read overlap required for path extension in the graph

Use Additional Params?

No

Job Resource Parameters

Use default job resource parameters

Execute

History

search datasets

Unnamed history

(empty)

This history is empty. You can [load your own data](#) or [get data from an external source](#)

Snapshot from Galaxy Main, <https://usegalaxy.org/>

galaxy.bridges.psc.edu

Galaxy Analyze Data Workflow Shared Data Visualization Help User Using 37.8 GB

Tools search tools

- Trinity: NGS RNA Seq
- Get Data
- Send Data
- Lift-Over
- Text Manipulation
- Filter and Sort
- Join, Subtract and Group
- Convert Formats
- Extract Features
- Fetch Sequences
- Fetch Alignments
- Get Genomic Scores
- Statistics
- Graph/Display Data
- Phenotype Association

NGS TOOLBOX BETA

NGS: QC and manipulation

Workflows

- All workflows

✓ Welcome to Galaxy on Bridges!

BRIDGES

A PITTSBURGH SUPERCOMPUTING CENTER RESOURCE

Galaxy is an open, web-based platform for data intensive biomedical research. The [Galaxy team](#) is a part of [BX at Penn State](#), and the [Biology](#) and [Mathematics and Computer Science](#) departments at [Emory University](#). The [Galaxy Project](#) is supported in part by [NHGRI](#), [NSF](#), [The Huck Institutes of the Life Sciences](#), [The Institute for CyberScience at Penn State](#), and [Emory University](#).

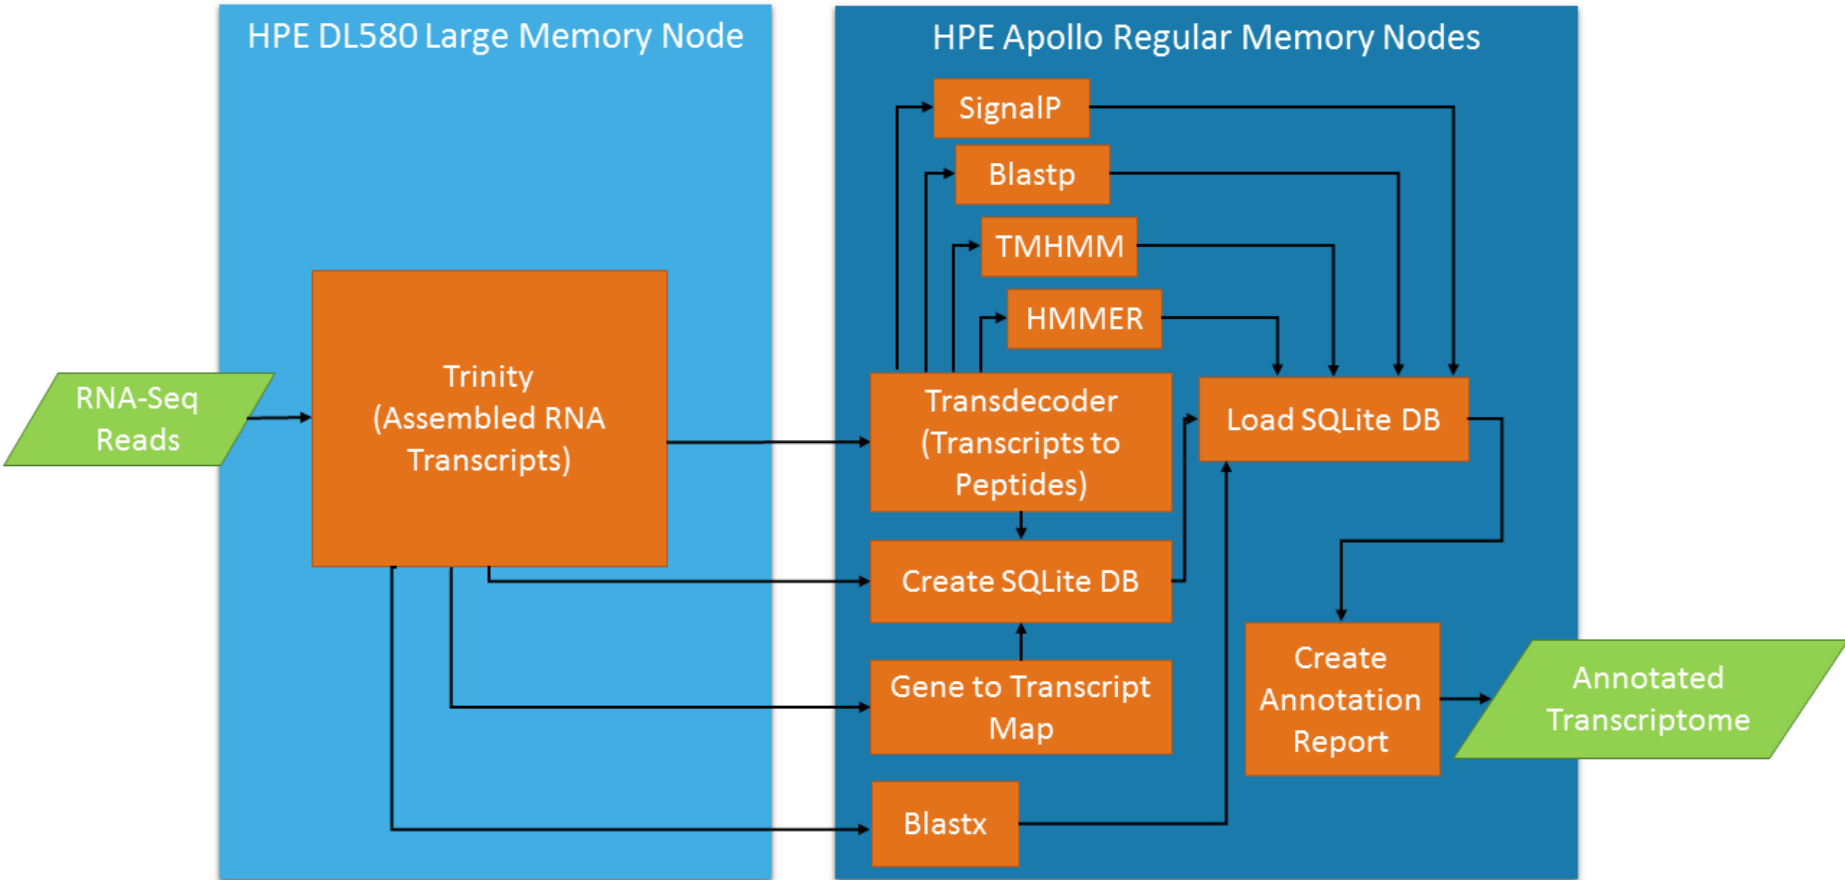
History search datasets

Bridges RNA-Seq Workflow
26 shown
11.5 GB

- 26: TRINOTATE ANNOTATION REPORT on data 25: Trinotate Annotation Report
- 25: LOAD SQLITE DB on data 21, data 20, and others: Trinotate sqlite database
- 24: LOAD SQLITE DB on data 21, data 20, and others: stdout.txt
- 23: CREATE SQLITE DB on data 8, data 13, and data 4: Trinotate sqlite database
- 22: CREATE SQLITE DB on data 8, data 13, and data 4: stdout.txt
- 21: TMHM on data

Galaxy RNA-Seq Workflow on Bridges

RNA-Seq De Novo Assembly and Annotation Workflow



Singularity on *Bridges*

- PSC is partnering with the CAMI project for reproducible evaluation of metagenomics tools
- CAMI and Joint Genome Institute defined bioboxes standard for Docker containers to encapsulate bioinformatics tools
- Problem: Docker security issues on shared systems (requires running root-level process)
- **Now Docker images (including bioboxes) can be converted to Singularity images and run on *Bridges* at PSC**



github.com/bioboxes



CAMI

<http://www.cami-challenge.org/>



Running Velvet in Singularity on *Bridges*

```
[blood@br006 velvet]$ ls
bioboxes_velvet.img input_data output_data run_singularity.sh
[blood@br006 velvet]$ ls input_data/ output_data/
input_data/
biobox.yaml reads.fq.gz

output_data/
[blood@br006 velvet]$ cat run_singularity.sh
#!/bin/bash
image="$1"
command="$2"
imageinput="$3"
imageoutput="$4"
input="$5"
output="$6"

module load singularity

singularity exec --bind "$input:$imageinput" --bind "$output:$imageoutput" $image $command
[blood@br006 velvet]$ sbatch -p RM-small -n 1 ./run_singularity.sh bioboxes_velvet.img "assemble default" /bbx/input /
bbx/output input_data output_data
Submitted batch job 1136824
[blood@br006 velvet]$ squeue -j 1136824
      JOBID PARTITION    NAME    USER  ST       TIME  NODES NODELIST(REASON)
      1136824  RM-small  run_sing  blood CG        0:00     1  r005
[blood@br006 velvet]$ ls
bioboxes_velvet.img input_data output_data run_singularity.sh slurm-1136824.out
[blood@br006 velvet]$ tail slurm-1136824.out
[0.024444] Concatenation...
[0.024444] Renumbering nodes
[0.024444] Initial node count 1
[0.024445] Removed 0 null nodes
[0.024445] Concatenation over!
[0.024460] Writing contigs into /tmp/tmp.01cIwTVBIO/contigs.fa...
[0.024609] Writing into stats file /tmp/tmp.01cIwTVBIO/stats.txt...
[0.024633] Writing into graph file /tmp/tmp.01cIwTVBIO/LastGraph...
[0.024834] Estimated Coverage cutoff = 3.619238
Final graph has 1 nodes and n50 of 2703, max 2703, total 2703, using 0/228 reads
[blood@br006 velvet]$ ls output_data/
biobox.yaml contigs.fa
```

Next steps: Facilitating open, accessible, reproducible, metagenomic analysis (on big data)

- Deploy CAMI bioboxes as Singularity containers on *Bridges*
- Benchmark analysis pipelines on CAMI + other metagenomics datasets on PSC *Bridges* - estimate resource requirements
- Establish a web gateway for running metagenomics pipelines in Singularity containers on *Bridges*
- Renew CAMI community allocation on *Bridges* for CAMI 2
- Enable free use of *Bridges* for CAMI 2 (through web or command line)

Getting Started on *Bridges*

- **Starter Allocation** <https://www.xsede.org/allocations>
 - Can request *anytime*... including *now!*
 - 1-year effective duration
 - Can request XSEDE ECSS (Extended Collaborative Support Service) or help from NCGAS (National Center for Genome Analysis Support)
- **Research Allocation (XRAC)** <https://www.xsede.org/allocations>
 - Appropriate for larger requests
 - Quarterly submission windows; *Next: Jun 15 - Jul 15 2017*
 - Can request ECSS

Contact me: blood@psc.edu

<https://psc.edu/bridges>

<https://xsede.org>

<https://ncgas.org>

Custom PSC topology for data-intensive HPC

20 Storage Building Blocks, implementing the parallel *Pylon* filesystem (~10PB) using PSC's SLASH2 filesystem

4 MDS nodes
2 front-end nodes
2 boot nodes
8 management nodes

6 "core" Intel OPA edge switches: fully interconnected, 2 links per switch

Intel OPA cables

4 ESM (12TB) compute nodes

2 gateways per ESM

42 LSM (3TB) compute nodes

12 database nodes

6 web server nodes

20 "leaf" Intel OPA edge switches

32 RSM nodes with NVIDIA next-generation GPUs

800 RSM (128GB) compute nodes, 48 with GPUs

16 RSM nodes with NVIDIA K80 GPUs

blood@psc.edu

<https://www.psc.edu/index.php/bridges-virtual-tour>