



BRIDGES

A PITTSBURGH SUPERCOMPUTING CENTER RESOURCE

A Uniquely Flexible HPC Resource for
New Communities and Data Analytics

Nick Nystrom • Director, Strategic Applications & Bridges PI • nystrom@psc.edu
XSEDE ECSS Symposium • December 15, 2015



The \$9.65M *Bridges* acquisition is made possible by National Science Foundation (NSF) award #ACI-1445606:

Bridges: From Communities and Data to Workflows and Insight



**Hewlett Packard
Enterprise**

is delivering *Bridges*

Disclaimer: The following presentation conveys the plan for *Bridges*. Certain details are subject to change.




BRIDGES
A PITTSBURGH SUPERCOMPUTING CENTER RESOURCE



An Important Addition to the National Advanced Cyberinfrastructure Ecosystem


Bridges will be a new resource on XSEDE and will interoperate with other XSEDE resources, Advanced Cyberinfrastructure (ACI) projects, campuses, and instruments nationwide.



Examples:



High-throughput genome sequencers



Reconstructing brain circuits from high-resolution electron microscopy



Social networks and the Internet

Data Infrastructure Building Blocks (DIBBs)

- Data Exacell (DXC)
- Integrating Geospatial Capabilities into HUBzero
- Building a Scalable Infrastructure for Data-Driven Discovery & Innovation in Education
- Other DIBBs projects

Other ACI projects



Carnegie Mellon University's Gates Center for Computer Science



Temple University's new Science, Education, and Research Center





Motivating Use Cases

(examples)

Data-intensive applications & workflows

Gateways – the power of HPC without the programming

Shared data collections & related analysis tools

Cross-domain analytics

Graph analytics, machine learning, genome sequence assembly, and other large-memory applications

Scaling research questions beyond the laptop

Scaling research from individuals to teams and collaborations

Very large in-memory databases

Optimization & parameter sweeps

Distributed & service-oriented architectures

Data assimilation from large instruments and Internet data

Leveraging an extensive collection of interoperating software

Research areas that haven't used HPC

Nontraditional HPC approaches to fields such as the physical sciences

Coupling applications in novel ways

Leveraging large memory and high island bandwidth




Potential Applications (*Examples*)

- Finding causal relationships in cancer genomics, lung disease, and brain dysfunction
- Analysis of financial markets and policies
- Improving the effectiveness of organ donation networks
- Assembling large genomes and metagenomes
- Recognizing events and enabling search for videos
- Understanding how the brain is connected from EM data
- Addressing societal issues from social media data
- Analyzing large corpora in the digital humanities
- Cross-observational analyses in astronomy & other sciences
- Data integration and fusion for history and related fields



Objectives and Approach

- Bring HPC to nontraditional users and research communities.
- Allow high-performance computing to be applied effectively to big data.
- *Bridge to campuses* to streamline access and provide cloud-like burst capability.
- Leveraging PSC's expertise with shared memory, *Bridges* will feature 3 tiers of large, coherent shared-memory nodes: **12TB**, **3TB**, and **128GB**.
- *Bridges* implements a uniquely flexible environment featuring interactivity, gateways, databases, distributed (web) services, high-productivity programming languages and frameworks, and virtualization, and campus bridging.



EMBO Mol Med (2013) DOI: 10.1002/emmm.201202388:
Proliferation of cancer-causing mutations throughout life



Alex Hauptmann et. al.: *Efficient large-scale content-based multimedia event detection*

Interactivity

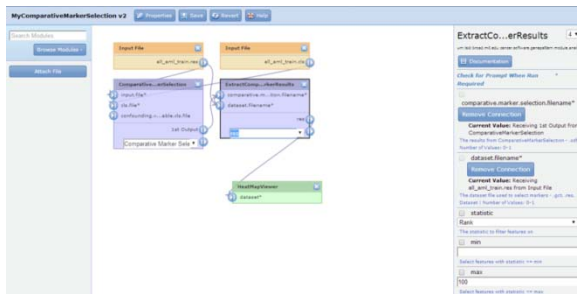
- *Interactivity is the feature most frequently requested by nontraditional HPC communities.*
- Interactivity provides immediate feedback for doing exploratory data analytics and testing hypotheses.
- *Bridges* offers interactivity through a combination of virtualization for lighter-weight applications and dedicated nodes for more demanding ones.



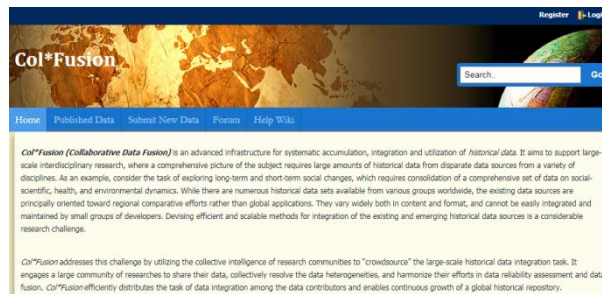
Gateways and Tools for Building Them

Gateways provide easy-to-use access to *Bridges'* HPC and data resources, allowing users to launch jobs, orchestrate complex workflows, and manage data from their browsers.

- *Extensive leveraging of databases and polystore systems*
- *Great attention to HCI is needed to get these right*



Interactive pipeline creation in GenePattern (Broad Institute)



Col*Fusion portal for the systematic accumulation, integration, and utilization of historical data, from <http://colfusion.exp.sis.pitt.edu/colfusion/>



Download sites for MEGA-6 (Molecular Evolutionary Genetic Analysis), from www.megasoftware.net




Virtualization and Containers

- Virtual Machines (VMs) enable flexibility, customization, security, reproducibility, ease of use, and interoperability with other services.
- Early user demand on PSC's Data Exacell research pilot project has centered on VMs for custom database and web server installations to develop data-intensive, distributed applications and containers for reproducibility.
- Bridges leverages OpenStack to provision resources, between interactive, batch, Hadoop, and VM uses.



High-Productivity Programming

Supporting languages that communities already use is vital for them to apply HPC to their research questions.




Spark + Hadoop + ... Ecosystem

Bridges' large memory is great for Spark!

Bridges enables workflows that integrate Spark/Hadoop, HPC, and/or shared-memory components.

hadoop



Cassandra




**APACHE
HBASE**



BRIDGES
A PITTSBURGH SUPERCOMPUTING CENTER RESOURCE



Campus Bridging



http://www.temple.edu/medicine/research/RESEARCH_TUSM/

- Through a pilot project with Temple University, the *Bridges* project will explore new ways to transition data and computing seamlessly between campus and XSEDE resources.
- **Federated identity management** will allow users to use their local credentials for single sign-on to remote resources, facilitating data transfers between *Bridges* and Temple's local storage systems.
- **Burst offload** will enable cloud-like offloading of jobs from Temple to *Bridges* and vice versa during periods of unusually heavy load.



BRIDGES
A PITTSBURGH SUPERCOMPUTING CENTER RESOURCE



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 HPE Integrity Superdome X (12TB) compute nodes

42 HPE ProLiant DL580 (3TB) compute nodes

12 HPE ProLiant DL380 database nodes

6 HPE ProLiant DL360 web server nodes

20 "leaf" Intel® OPA edge switches

32 RSM nodes with NVIDIA next-generation GPUs

16 RSM nodes with NVIDIA K80 GPUs

800 HPE Apollo 2000 (128GB) compute nodes

Purpose-built Intel® Omni-Path topology for data-intensive HPC


<http://psc.edu/bridges>

High-Performance, Data-Intensive Computing

- 3 tiers of large, coherent shared memory nodes

Memory per node	Number of nodes	Example applications
12 TB HPE Integrity Superdome X	4	Genomics, machine learning, graph analytics, other extreme-memory applications
3 TB HPE ProLiant DL580	42	Virtualization and interactivity including large-scale visualization and analytics; mid-spectrum memory-intensive jobs
128 GB HPE Apollo 2000	800	Execution of most components of workflows, interactivity, Hadoop, and capacity computing

- The latest Intel® Xeon® CPUs
- NVIDIA® Tesla® dual-GPU accelerators



Database and Web Server Nodes

- Dedicated database nodes will power persistent relational and NoSQL databases **HPE ProLiant DL380**
 - Support data management and data-driven workflows
 - SSDs for high IOPs; RAIDed HDDs for high capacity



- Dedicated web server nodes **HPE ProLiant DL360**
 - Enable distributed, service-oriented architectures
 - High-bandwidth connections to XSEDE and the Internet

Data Management

- *Pylon: A large, central, high-performance filesystem*
 - Visible to all nodes
 - Large datasets, community repositories (10 PB usable)
- **Distributed (node-local) storage**
 - Enhance application portability
 - Improve overall system performance
 - Improve performance consistency to the shared filesystem
- **Acceleration for Hadoop-based applications**



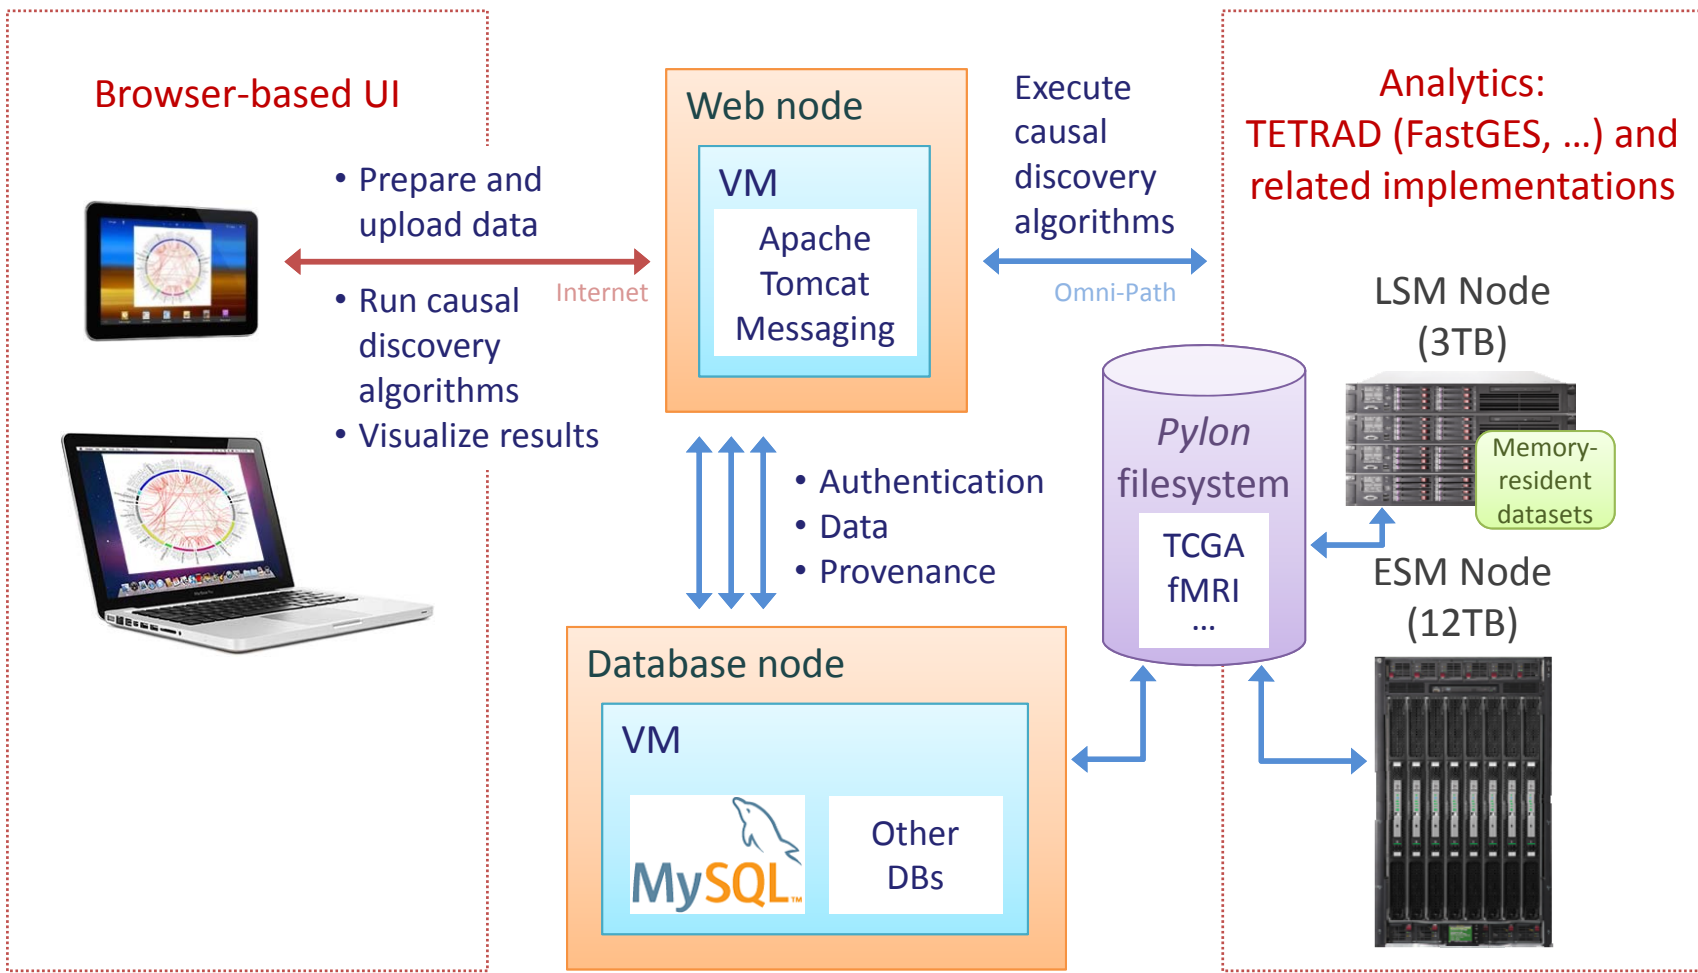
Intel® Omni-Path Architecture (OPA)

- *Bridges* is the first production deployment of Omni-Path
- Omni-Path connects all nodes and the shared filesystem, providing *Bridges* and its users with:
 - 100 Gbps line speed per port;
25 GB/s bidirectional bandwidth per port
 - 160M MPI messages per second
 - 48-port edge switch reduces interconnect complexity and cost
 - HPC performance, reliability, and QoS
 - OPA-compliant applications supported without modification
 - Early access to this new, important, forward-looking technology
- *Bridges* deploys OPA in a two-tier island topology developed by PSC for cost-effective, data-intensive HPC



Example: Causal Discovery Portal

Center for Causal Discovery, an NIH Big Data to Knowledge Center of Excellence



Getting Started on *Bridges*

- **Starter Allocation** <https://www.xsede.org/allocations>
 - Can request *anytime*... including *now*
 - Can request XSEDE ECSS (Extended Collaborative Support Service)
- **Research Allocation (XRAC)** <https://www.xsede.org/allocations>
 - Appropriate for larger requests; can request ECSS
 - Quarterly submission windows; *Next: Dec. 15, 2015–Jan. 15, 2016*
- **Early User Period**
 - Users with starter or research proposals may be eligible for *Bridges'* Early User Period
- **Questions?**
 - See <http://psc.edu/bridges>
or email Nick Nystrom at nystrom@psc.edu



Bridges Target Schedule

- Acquisition
 - Construction planned to begin October, 2015
 - Early User Period starting in late 2015
 - Base-Level System: installed, early users very soon
 - Phase 1 underway: completion ~March 2015
 - Phase 2 Technical Update: summer 2016
- XRAC-awarded projects
 - Planned to begin in January, 2016
- Related resources
 - *Greenfield* to transition from *Blacklight* and to provide data for developing XRAC proposals for *Bridges*
 - *Data Exacell*: a research pilot project to explore the coupling of data analytics with novel storage

