Building Bridges: A System for New HPC Communities

HPC User Forum 59

LRZ, Garching

October 16, 2015

Presenter: Jim Kasdorf

Director, Special Projects

Pittsburgh Supercomputing Center

kasdorf@psc.edu

XSEDE
Extreme Science and Engineering Discovery Environment

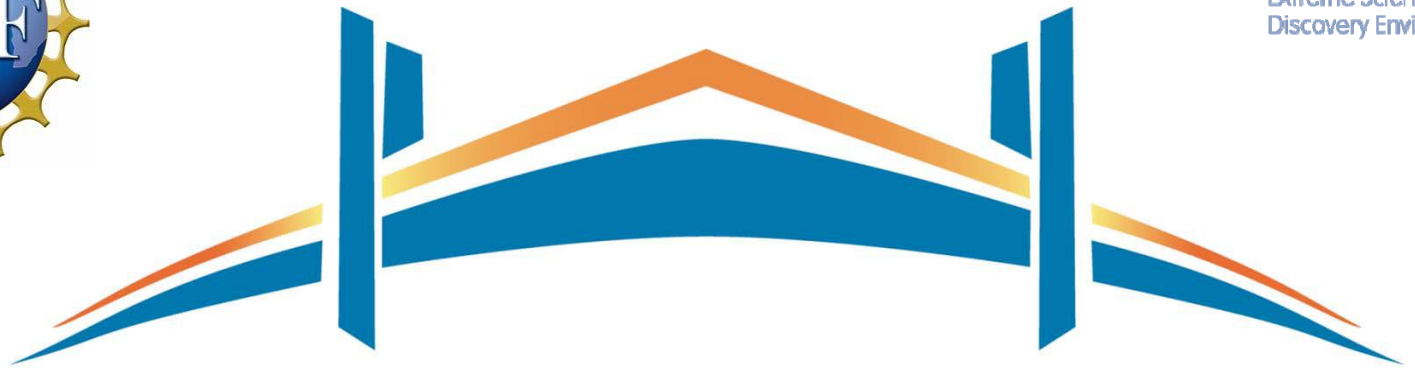

BRIDGES

A PITTSBURGH SUPERCOMPUTING CENTER RESOURCE

A Uniquely Flexible HPC Resource for New Communities and Data Analytics

The Shift to Big Data

Pan-STARRS telescope
http://pan-starrs.ifa.hawaii.edu/public/

Genome sequencers
(Wikipedia Commons)

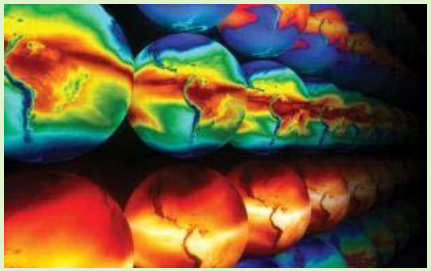

NOAA climate modeling
http://www.ornl.gov/info/ornlreview/v42_3_09/article02.shtml

New Emphases

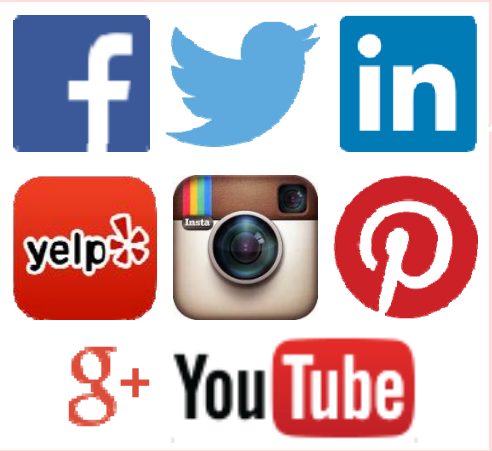

Social networks and the Internet

Video
Wikipedia Commons

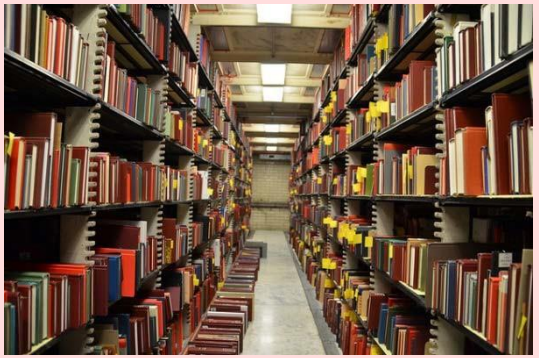

Library of Congress stacks
https://www.flickr.com/photos/danlem2001/6922113091/

Collections
Horniman museum: http://www.horniman.ac.uk/get_involved/blog/bioblitz-insects-reviewed

Legacy documents
Wikipedia Commons

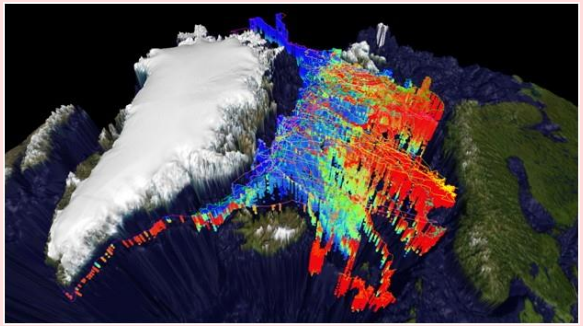

Environmental sensors: Water temperature profiles from tagged hooded seals
http://www.arctic.noaa.gov/report11/biodiv_whales_walrus.html

Algorithms and Applications

Algorithms and Applications Have Also Changed

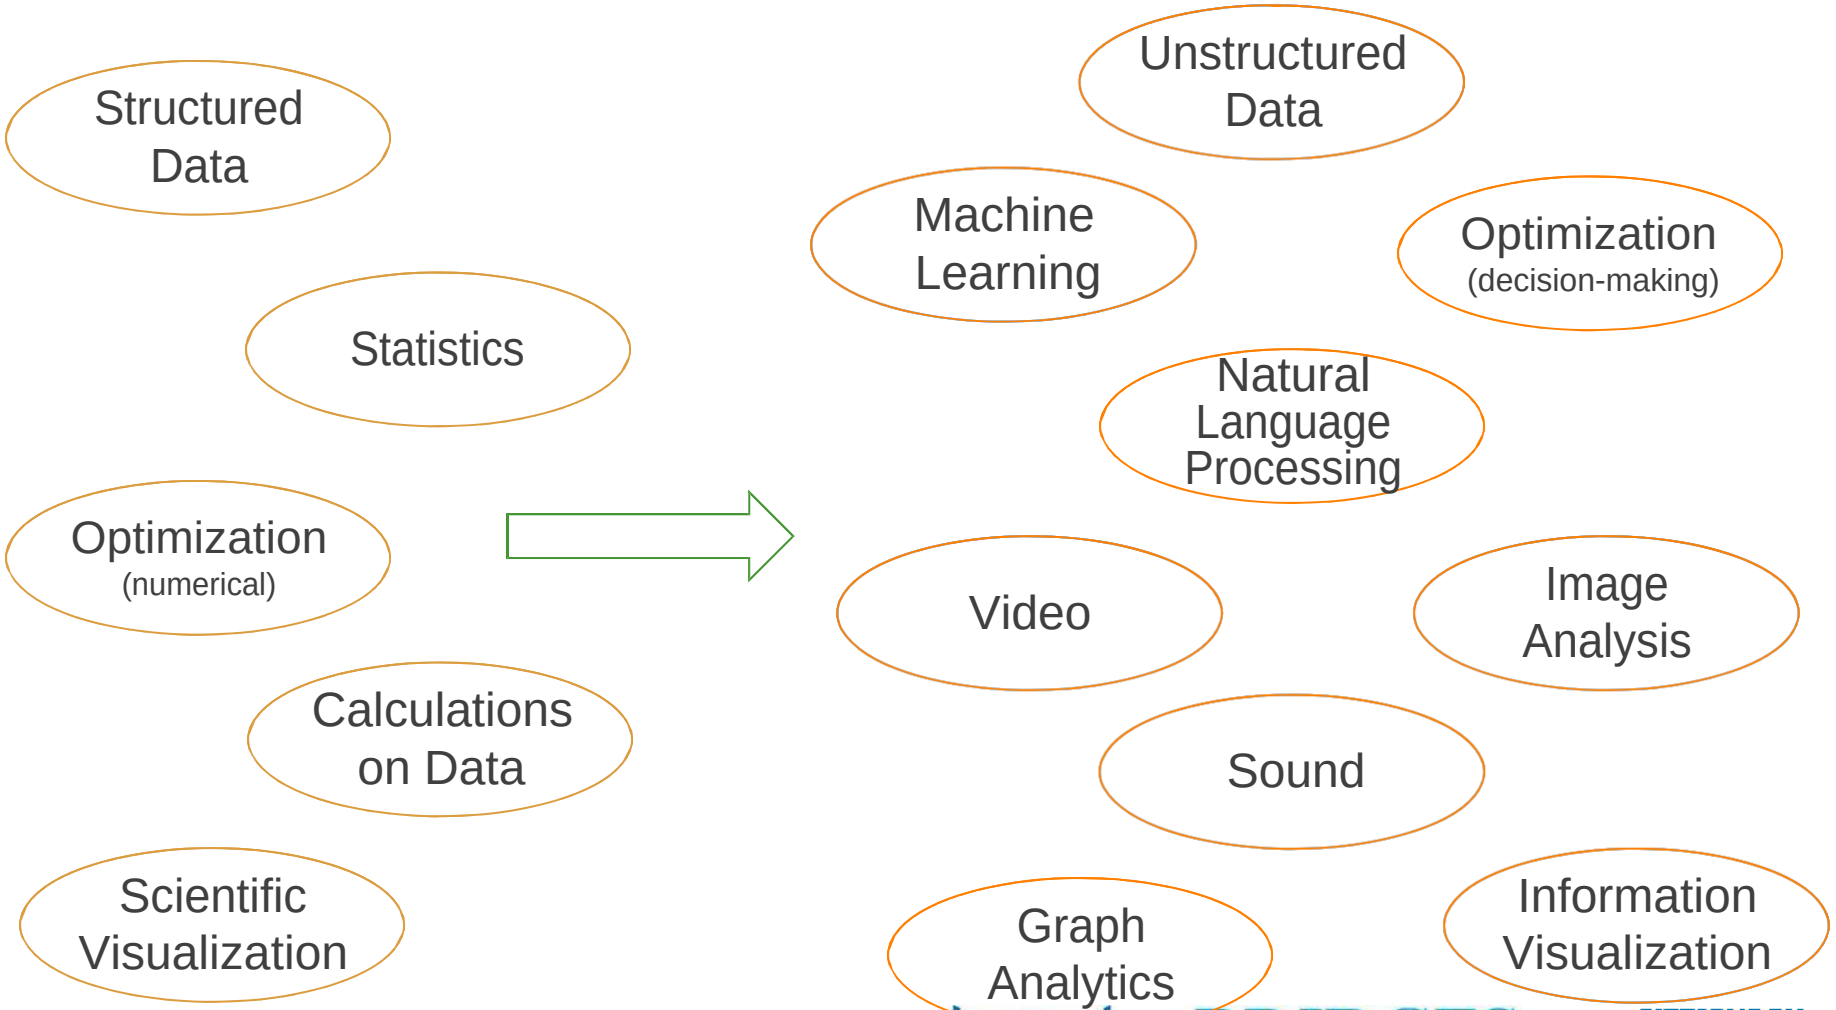

The $9.65M *Bridges* acquisition is made possible by National Science Foundation

Bridges: From Communities and Data to Workflows and Insight

Hewlett Packard
Enterprise

HPE is delivering *Bridges*

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

High-Level Use Cases

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

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 ease access and provide burst capability.
- Leveraging PSC's expertise with shared memory, *Bridges* will feature 3 tiers of large, coherent shared-memory nodes.
- *Bridges* will leverage its large memory for interactivity and seamlessly support applications through virtualization, gateways, familiar and productive programming environments and data-driven workflows.

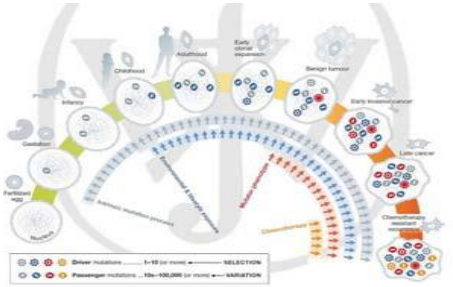

EMBO Mol Med (2013) DOI: 10.1002/emmm.201202388:
Proliferation of cancer-causing mutations throughout life

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

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.

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* will offer interactivity through a combination of virtualization for lighter-weight applications and dedicated nodes for more demanding ones.

Gateways and Tools for Building Them

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

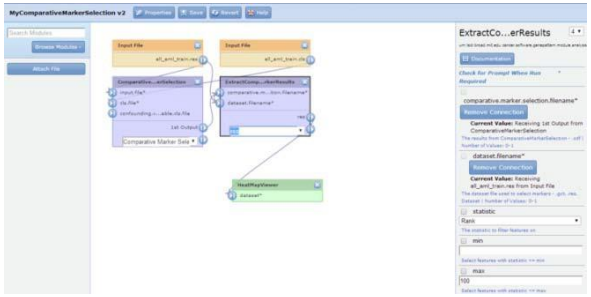

Interactive pipeline creation in GenePattern (Broad Institute)

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

High-Productivity Programming

Bridges will feature high-productivity programming languages and tools.

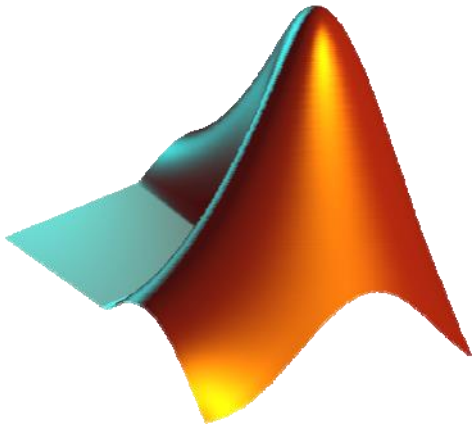

Hadoop Ecosystem

- *Bridges* will provide acceleration for Hadoop applications
- Large memory will be great for Spark.

Cassandra

APACHE HBASE

Campus Bridging

BRIDGES
A PITTSBURGH SUPERCOMPUTING CENTER RESOURCE

Federated identity management

Burst offload

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.

Custom PSC topology
for data-intensive HPC

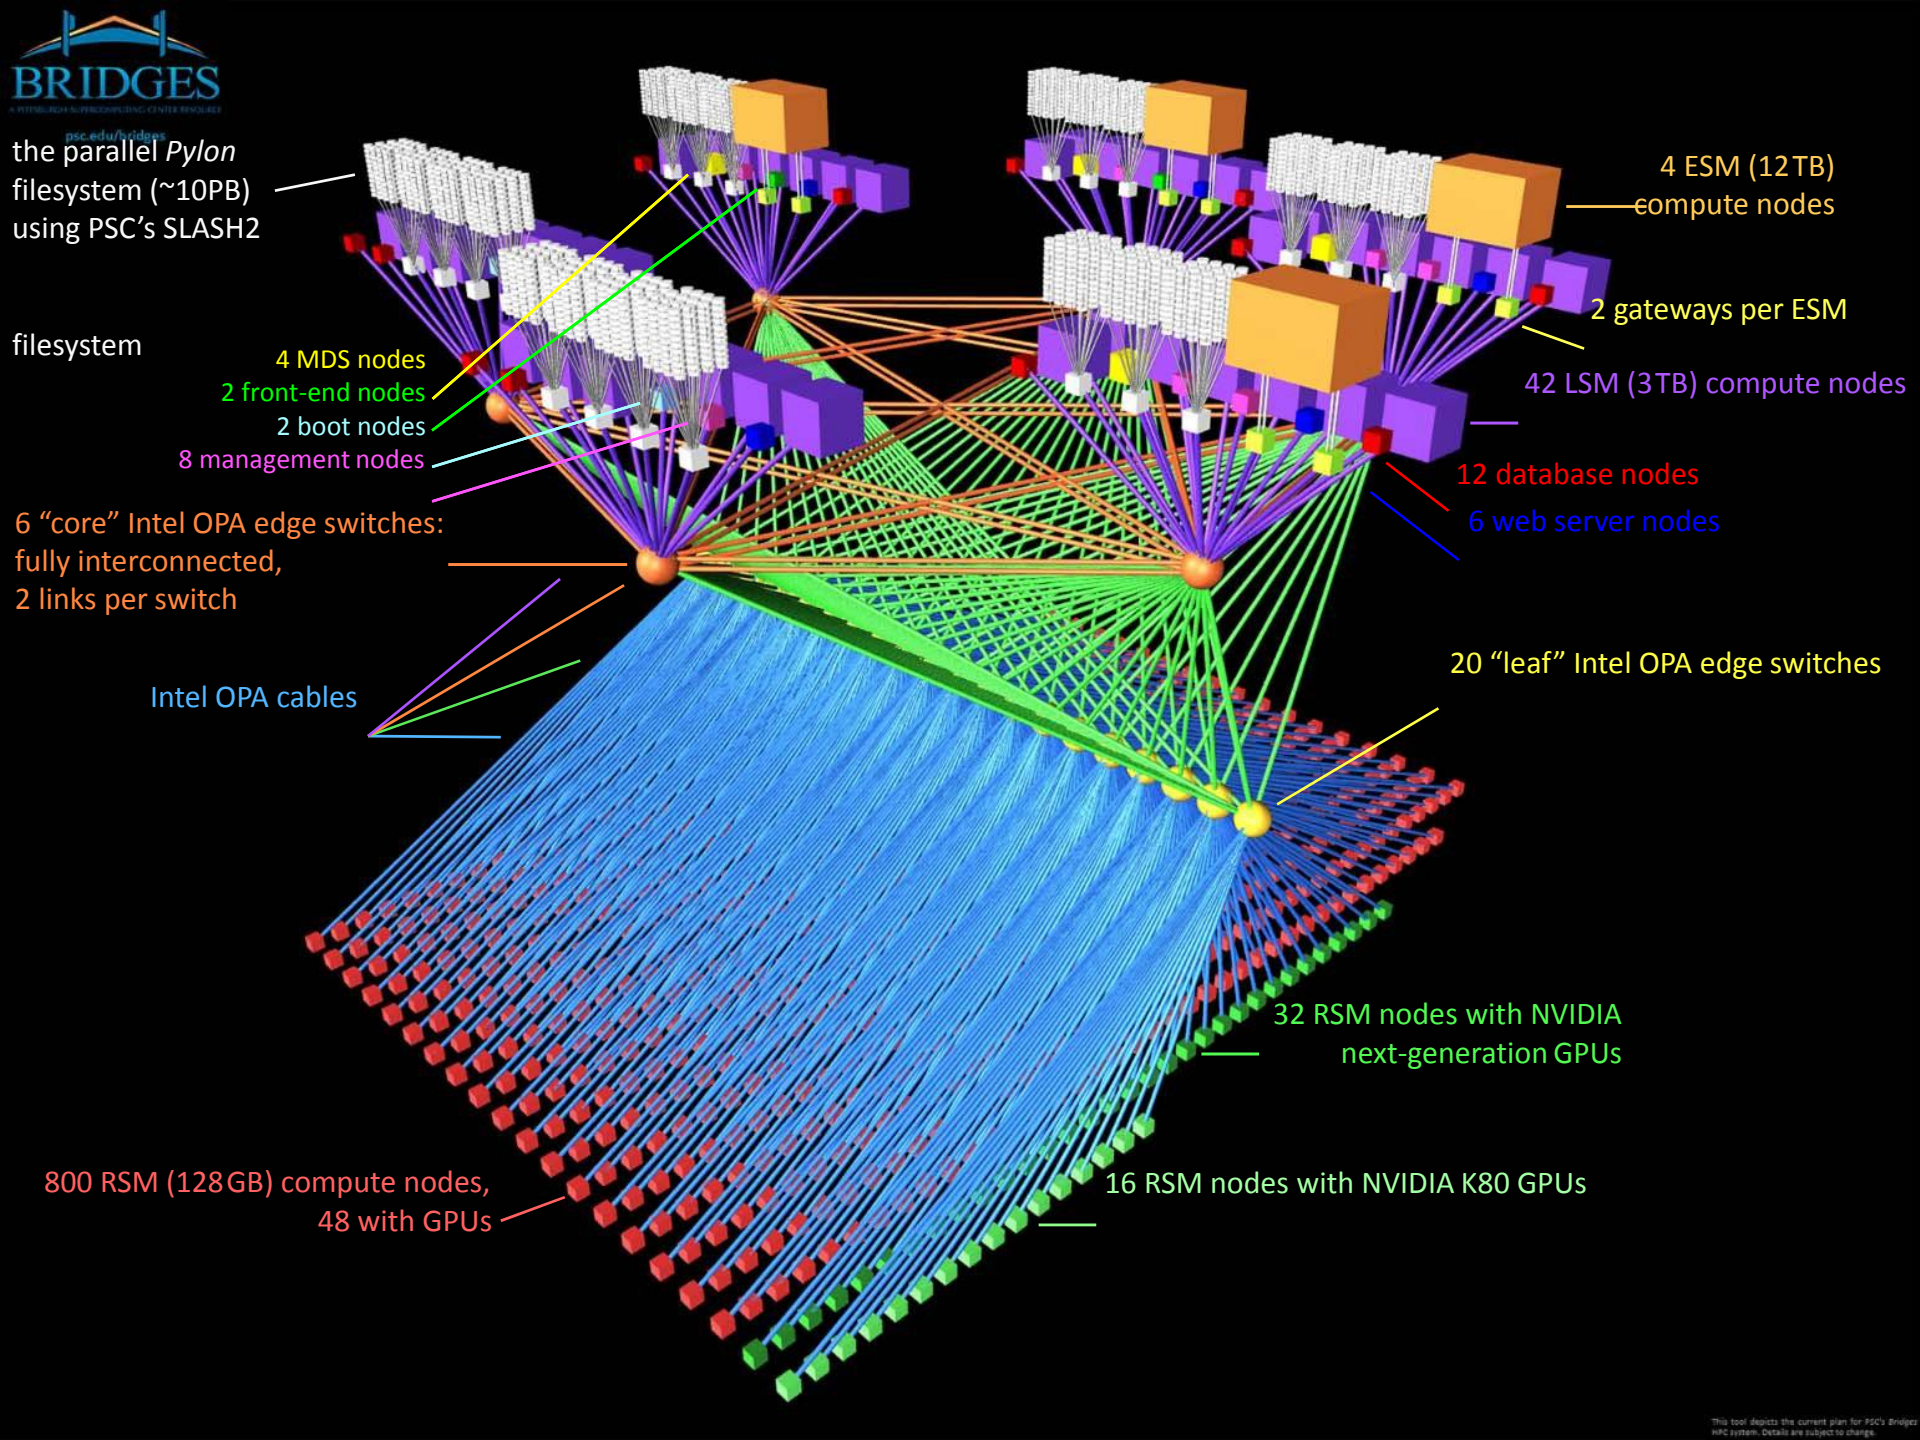
BRIDGES
psc.edu/bridges
the parallel *Pylon* filesystem (~10PB) using PSC's SLASH2
filesystem
4 MDS nodes
2 front-end nodes
2 boot nodes
8 management nodes
6 "core" Intel OPA edge switches: fully interconnected, 2 links per switch
Intel OPA cables
4 ESM (12TB) compute nodes
2 gateways per ESM
42 LSM (3TB) compute nodes
12 database nodes
6 web server nodes
20 "leaf" Intel OPA edge switches
32 RSM nodes with NVIDIA next-generation GPUs
16 RSM nodes with NVIDIA K80 GPUs
800 RSM (128GB) compute nodes, 48 with GPUs
This tool depicts the current plan for PSC's Bridges HPC system. Details are subject to change.

High-Performance, Data-Intensive Computing

- Three tiers of large, coherent shared memory nodes

Memory per node	Number of nodes	Example applications
12 TB	Several	Genomics, machine learning, graph analytics, other extreme-memory applications
3 TB	Tens	Virtualization and interactivity including large-scale visualization and analytics; mid-spectrum memory-intensive jobs
128 GB	Hundreds	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
  - 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

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

- Omni-Path will connect 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
  - OFA-compliant applications supported without modification
  - Early access to this new, important, forward-looking technology
- The Intel Omni-Path Architecture will deployed in *Bridges* using a novel topology developed by PSC to enable data-intensive HPC

Example: Causal Discovery

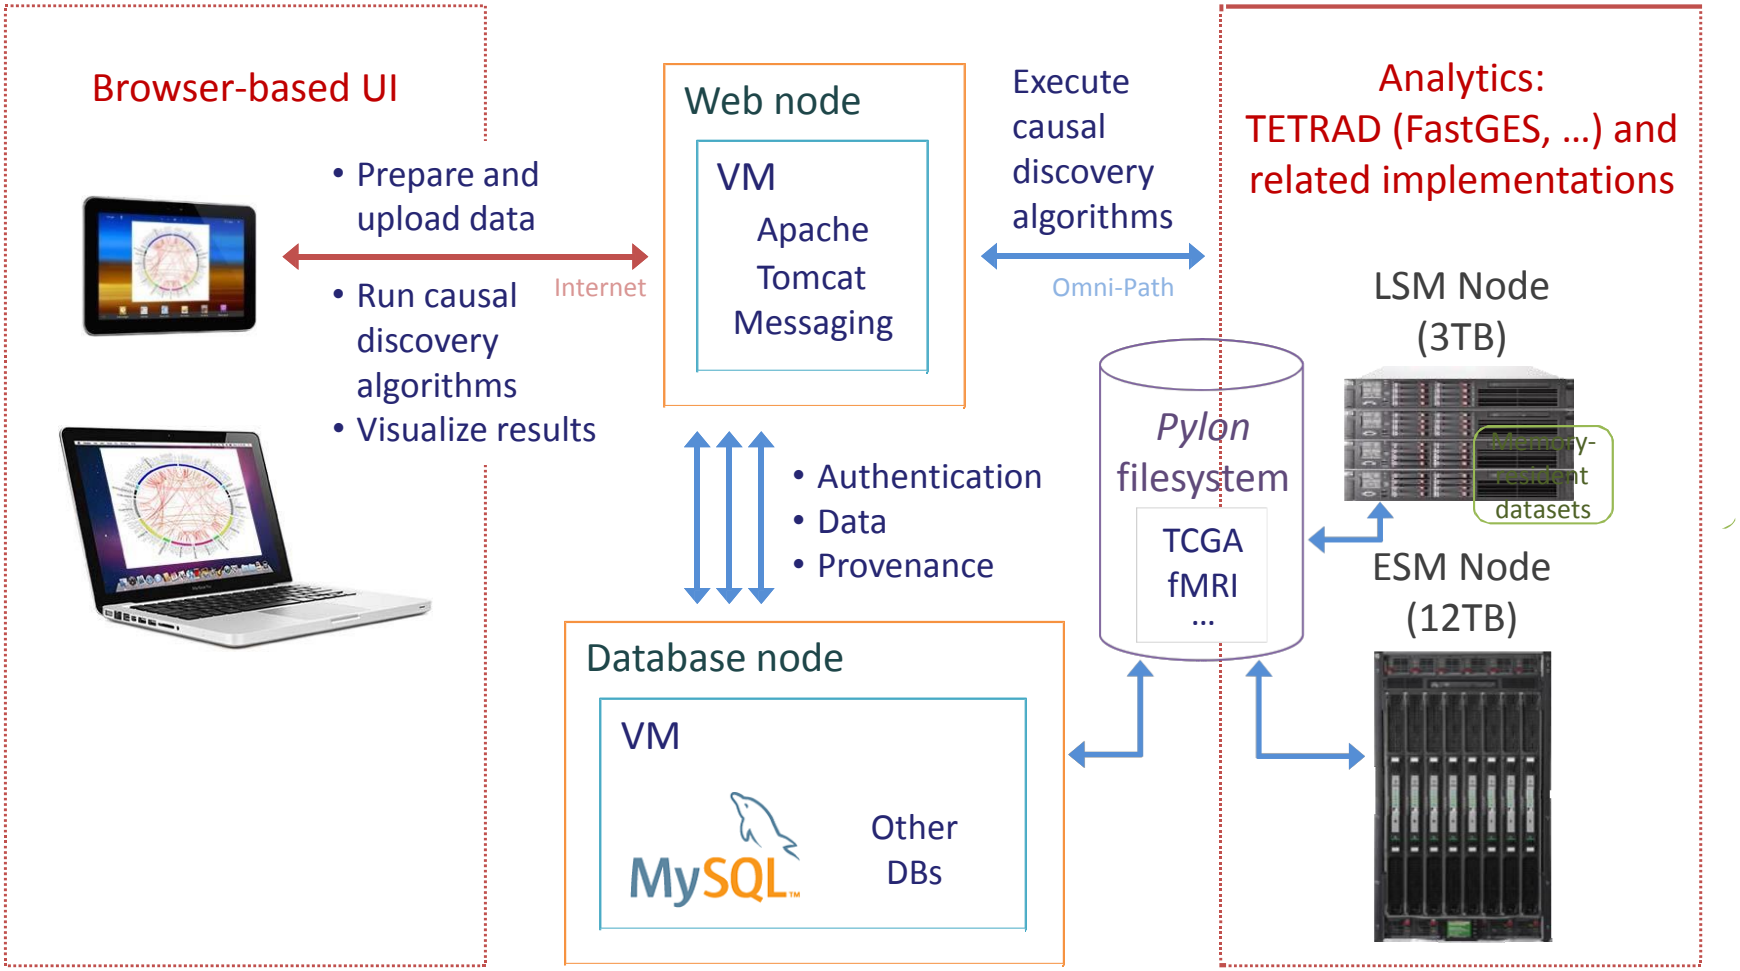

Getting Started on *Bridges*

- Starter Allocation https://www.xsede.org/allocations
  - Can request anytime... including *now!*
  - 1-year effective duration, to begin when *Bridges* comes online; can begin running on *Greenfield*
  - Can request XSEDE ECSS (Extended Collaborative Support Service)
- Research Allocation (XRAC) https://www.xsede.org/allocations
  - Appropriate for larger requests
  - Quarterly submission windows; *Next: Sept. 15 – Oct. 15, 2015*
  - Can request ECSS
- Early User Period
  - Users with starter or research proposals may be eligible for *Bridges'* Early User Period (late fall '15)

Bridges Target Schedule

- Acquisition
  - Construction planned to begin October, 2015
  - Early User Period starting in late 2015
- XRAC Allocated Use
  - 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

For Additional Information

Project website: www.psc.edu/bridges

Bridges PI: Nick Nystrom
nystrom@psc.edu

Co-PIs: Michael J. Levine
Ralph Roskies
J Ray Scott

Project Manager: Janet Brown