



The Data eXaCell – DXC

J. Ray Scott

DXC PI

May 17, 2016

DXC Leadership

- Mike Levine Co-Scientific Director Co-PI
- Nick Nystrom Senior Director of Research Co-PI
- Ralph Roskies Co-Scientific Director Co-PI
- Robin Scibek Project Manager PM
- J. Ray Scott Senior Director of Facilities Technology PI

Pittsburgh Supercomputing Center

- The Pittsburgh Supercomputing Center:
 - Joint effort of Carnegie Mellon University and the University of Pittsburgh
 - 30 years national leadership in:
 - High-performance and data-intensive computing
 - Data management technologies
 - Software architecture, implementation, and optimization
 - Enabling researchers nationwide
 - Networking and network optimization
 - Supported by: NSF, NIH, the Commonwealth of Pennsylvania, DOE, DoD, foundations, and industry

DXC/DIBBs

- Project in a nutshell:
 - DXC/DIBBs:
 - Accelerated, development pilot project
 - Creating, deploying and testing relevant software and hardware building blocks
 - Functionalities designed to support data-analytic capabilities for data intensive scientific research
- Guided by selected collaborating research groups
 - Diverse set of emerging and existing data-intensive & data-analytic applications
 - Not well served by local resources or existing HPC systems
 - Learn what they (and presumably others) need.

Radio Astronomy at Green Bank (NRAO)

PI: David Halstead, National Radio Astronomy Observatory



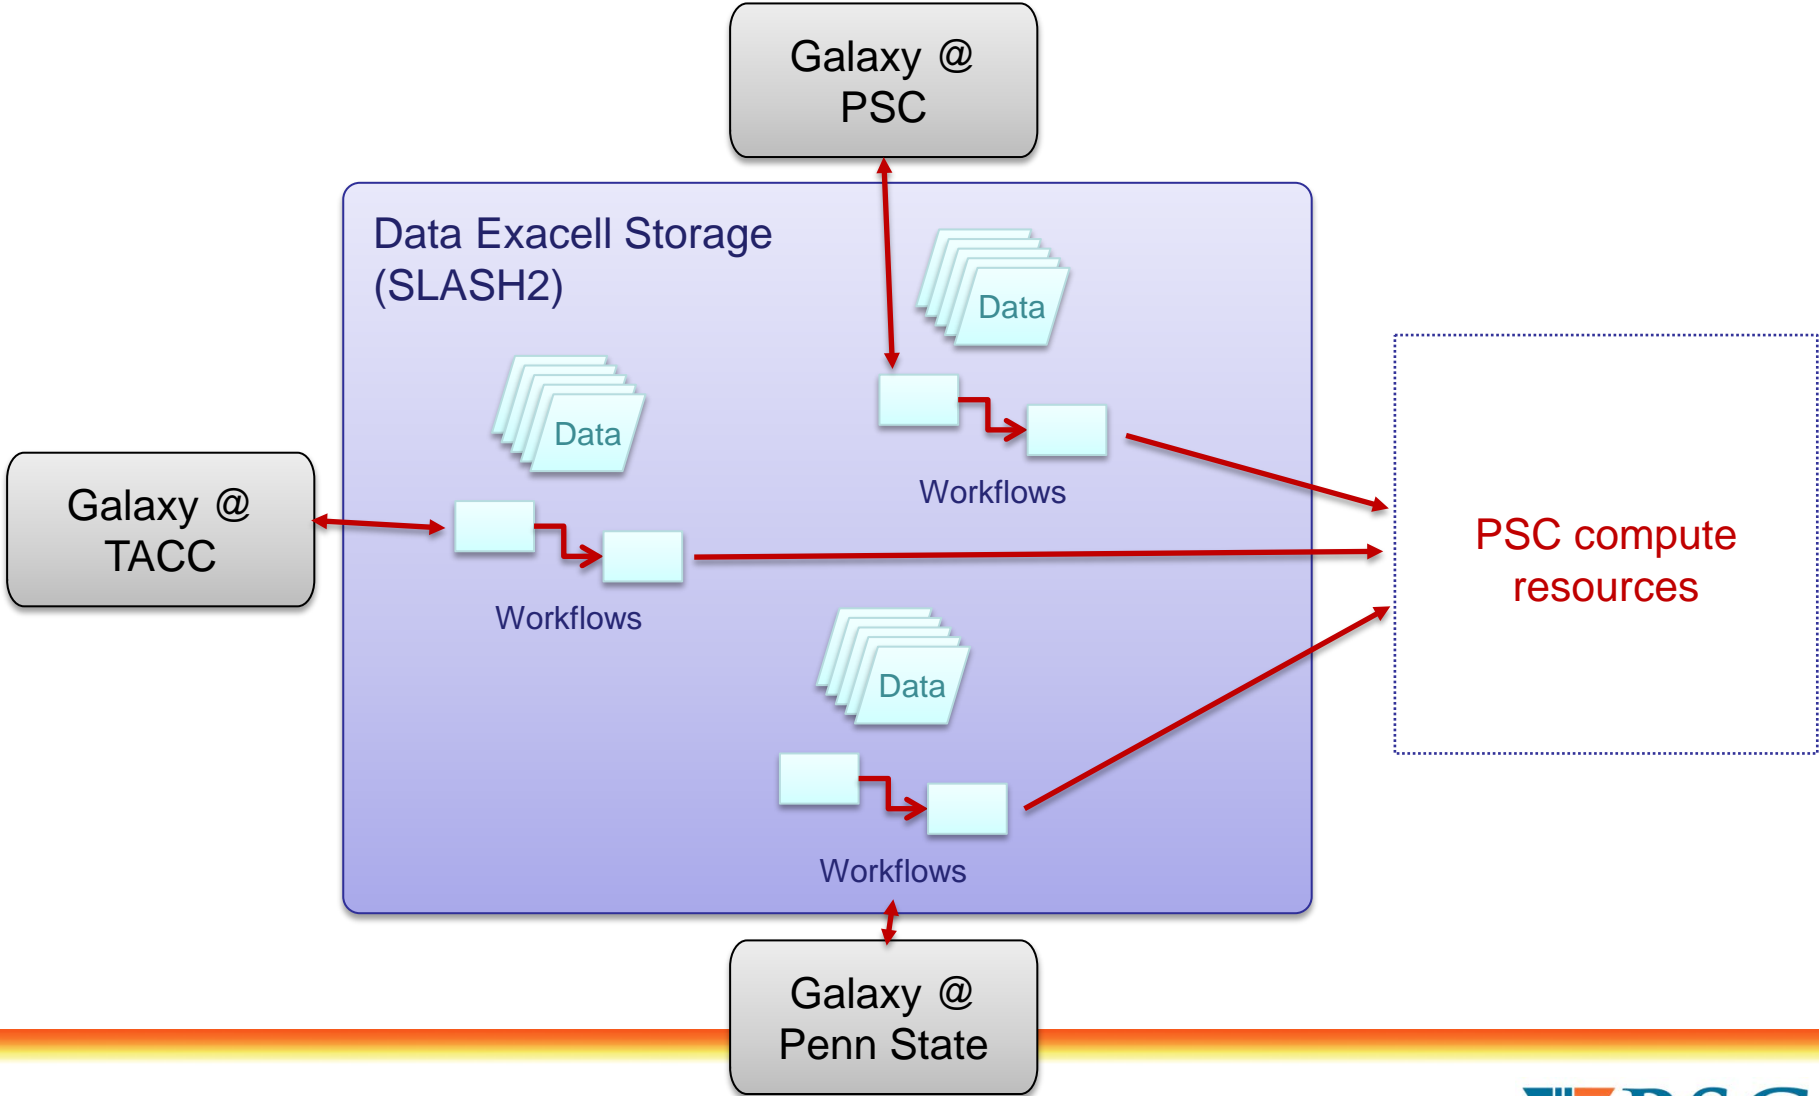
The Robert C. Byrd Green Bank Telescope (GBT) has a dish diameter of 100 meters and wavelength sensitivity from 3m down to 2.6mm.

Thanks to new focal plane receivers and back-end equipment, the volume of data produced by the GBT is rising rapidly.

The GBT Mapping Pipeline is a new software tool intended to ease the production of sky maps from this massive data stream. Mapping of large patches of sky is one of the main uses of the GBT, and is complementary to the highly focused studies from facilities like the EVLA.

NRAO and PSC are collaborating to leverage coupled storage and analytics on the DXC (and later, Bridges) for the Mapping Pipeline.

Galaxy: DXC Pilot



SLASH2




- SLASH2 is designed from the ground up to be:
 - wide-area
 - portable
 - Scalable
- Features
 - files are managed as chunks
 - system managed replication
 - error checking

DXC SLASH2 Schematic



A DSC hardware building block


DXC SBB
512 TB (u)





4*44 4 TB (r)

'hJB' = half of
a JBOD

DXC SLASH2
Command and Control



File Systems Development Support


- DXC will involve development of advanced file system support
- Initial effort to revamp support tools used in SLASH2
- Portable File system Libraries (PFL) 
- Weldable Overlay Knack File System (WOKFS) 

ADAPT-FS: Active Data Processing and Transformation File System

- On-the-fly CPU/GPU computation
- Replaces explicit storage of processed images
- Enables collaborative processing and sharing of large image data sets with minimal data duplication
 - 3D electron microscopy data of brain tissue
 - currently in the 100TB range,
 - petabyte scales forthcoming)
- Portable File system Library module with a flexible interface
- per-dataset specification of data interpretation, preparation, and transform as submodule drivers.



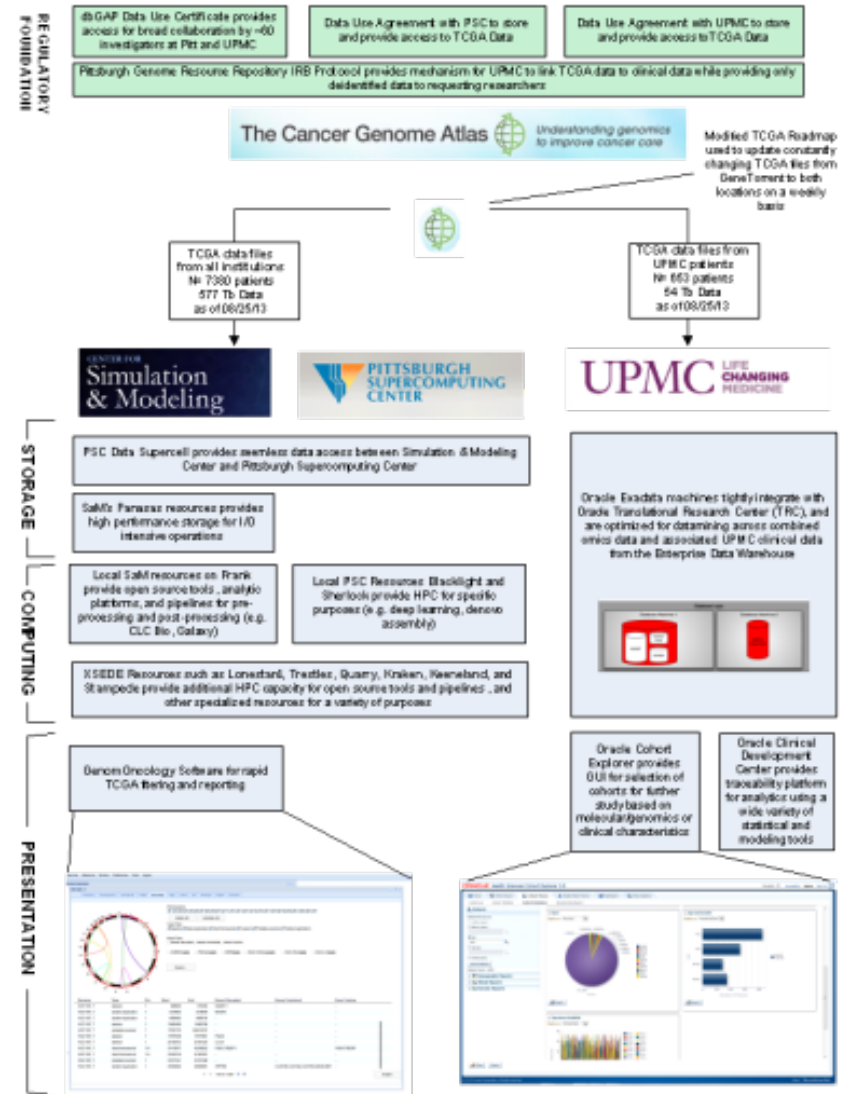
Multi-site Support

- Multiple metadata servers (MDS)
 - global mount support
 - foundation for further multi-MDS development
- SLASH2 ↔ local file multi-site file import/export
- Workflow integration
 - XSEDE Extended Support for Science Gateways
- Cross site UID mapping
 - security
 - federated authentication
- Enhanced access controls 
 - e.g. read/delete only file access
 - building block: SCAMPI file system
- Public cloud support

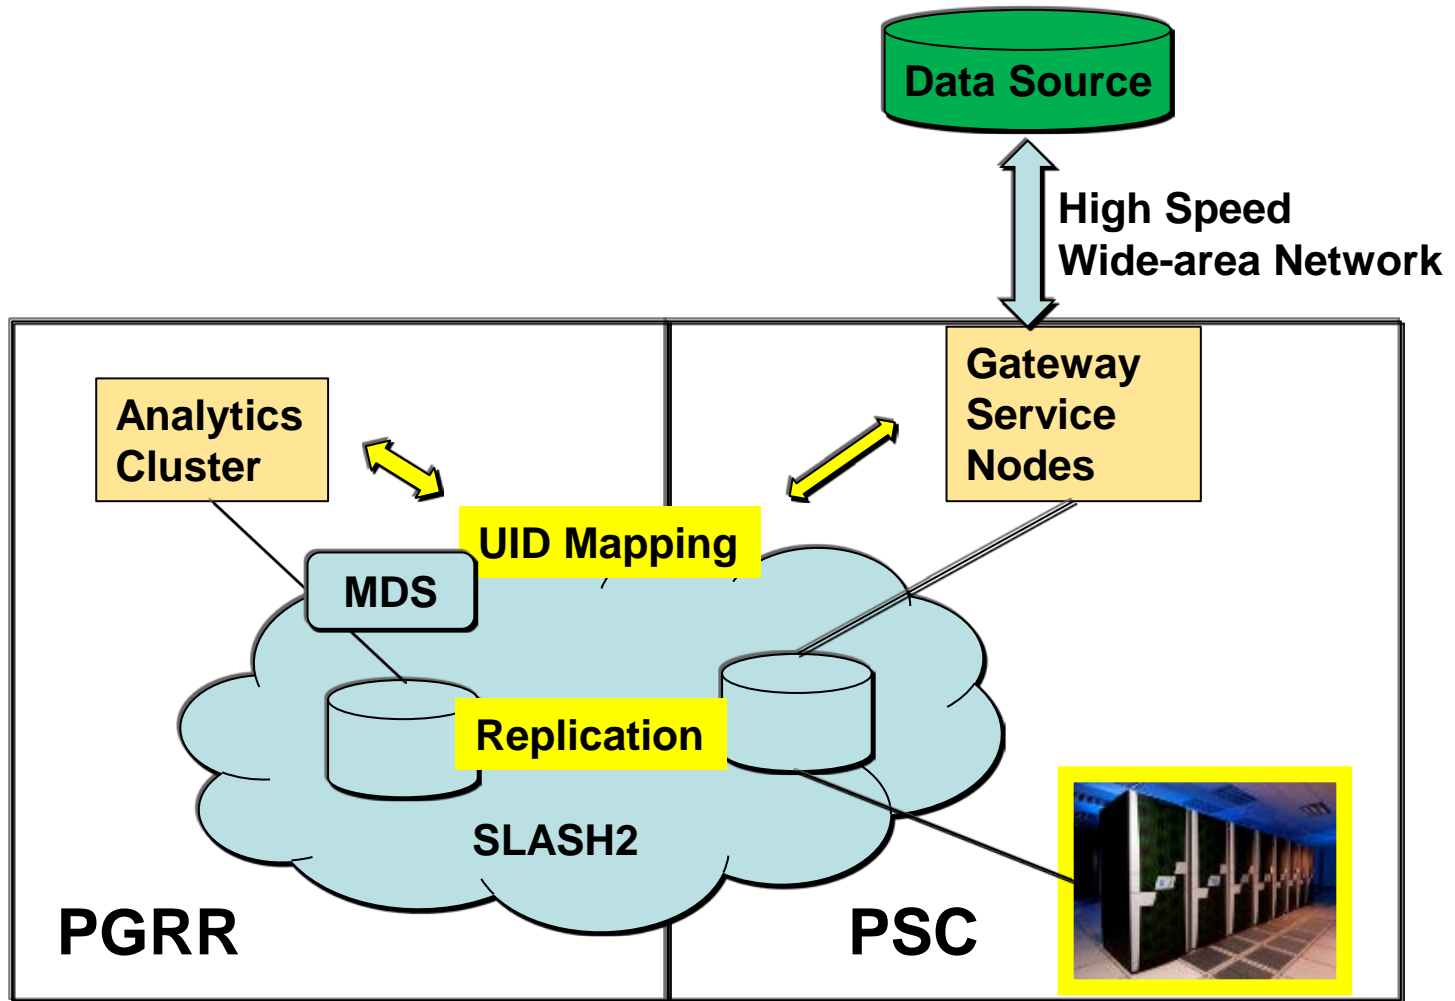
Pittsburgh Genome Resource Repository (PGRR) pgrr.pitt.edu

Collaborative effort to address challenges with TCGA data:

- **University of Pittsburgh:** Institute for Personalized Medicine (IPM), U. Pitt. Cancer Institute (UPCI), Department of Biomedical Informatics (DBMI), Center for Simulation and Modeling (SaM)
- **University of Pittsburgh Medical Center (UPMC)**
- **Pittsburgh Supercomputing Center (PSC)**



PGRR DXC Architecture



DXC Hardware To Support Research Collaborators

- Equipment in place
 - 41 servers
 - 128GB
 - GPU
 - 3TB
 - 12TB
 - 5 PB of SLASH2 managed shared storage
- Being used by both developers and collaborators

DXC Shared File System

- Available for DXC applications and development efforts
- Mixed capability components to facilitate optimization testing
- Subsets withheld to allow invasive testing without harm to collaborators' data
- DXC shared file system is a Building Block deliverable
 - commodity components
 - testing and optimization in progress
 - a prototype for research groups requiring inexpensive, large-scale storage

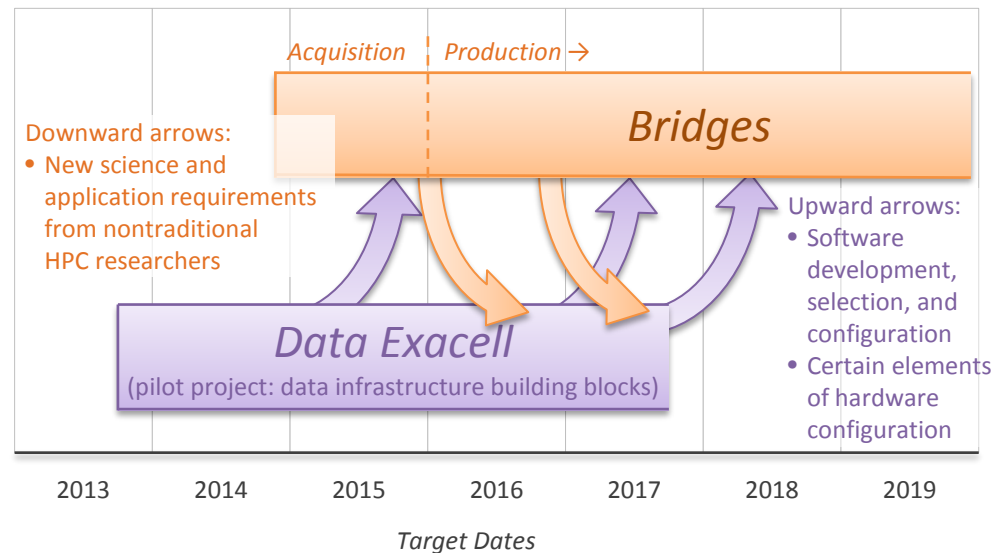


XSEDE Service Provider Support

- Blacklight
 - SGI UV 1000
 - 16 TB * 2 coherent shared memory
- Greenfield
 - HP Superdome X
 - 12 TB coherent shared memory
 - HP DL580
 - 3 TB coherent shared memory
 - Shared file system
 - 800 TB useable
 - SLASH2

Bridges and the Data Exacell : A Valuable Engineering Lifecycle

- Hardware and software “building blocks” developed through the *Data Exacell* (DXC) pilot project enables new application architectures on *Bridges* and convenient, high-performance data movement between *Bridges* and users, campuses, and instruments.
- *Bridges* and *DXC* will provide complementary roles for production and application prototyping.



Summary

- First phase
 - Building new storage and analytic facility
 - Gathering users
 - Understanding needs
 - Prototype solutions in place
- Next phase
 - Larger user experience
 - Multi-site support
 - Authentication
 - Metadata services
 - Distributed MDS
 - Data tagging

