



Cyberinfrastructure of a Multi-Petabyte Microscopy Resource for Neuroscience Research

K. Benninger

G. Hood

D. Simmel

L. Tuite

A. Wetzel

A. Ropelewski

ropelews@psc.edu

Pittsburgh Supercomputing Center,

Carnegie Mellon University

Pittsburgh, Pennsylvania

S. Watkins

A. Watson

swatkins@pitt.edu

Center for Biologic Imaging,

University of Pittsburgh

Pittsburgh, Pennsylvania

M. Bruchez

bruchez@andrew.cmu.edu

Molecular Biosensor and Imaging

Center, Carnegie Mellon University

Pittsburgh, Pennsylvania

ABSTRACT

Advances in microscopy technologies and techniques are driving development of an integrated cyberinfrastructure to manage the vast amounts of image data being generated. To address this need, a team of neuroscientists and cyberinfrastructure engineers has integrated mass storage, networking, and high-performance computing components to create the Brain Image Library (BIL). BIL serves the neuroscience community by providing a persistent repository for the massive amounts of data being generated (multi-petabyte scale) and by offering brain researchers supercomputing-class processing and visualization capabilities for working with high-resolution brain image datasets.

CCS CONCEPTS

• **Applied computing** → **Bioinformatics**; • **Information systems** → Storage area networks; Network attached storage; **Digital libraries and archives**; **Computing platforms**.

KEYWORDS

image archive, neuroscience, storage architecture

ACM Reference Format:

K. Benninger, G. Hood, D. Simmel, L. Tuite, A. Wetzel, A. Ropelewski, S. Watkins, A. Watson, and M. Bruchez. 2020. Cyberinfrastructure of a Multi-Petabyte Microscopy Resource for Neuroscience Research. In *Practice and Experience in Advanced Research Computing (PEARC '20)*, July 26–30, 2020, Portland, OR, USA. ACM, New York, NY, USA, 7 pages. <https://doi.org/10.1145/3311790.3396653>

1 BRAIN IMAGE LIBRARY (BIL) PROJECT INTRODUCTION

Advancements in the field of microscopy have pushed the boundaries of what was once thought possible. Fluorescent microscope

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for third-party components of this work must be honored. For all other uses, contact the owner/author(s).

PEARC '20, July 26–30, 2020, Portland, OR, USA

© 2020 Copyright held by the owner/author(s).

ACM ISBN 978-1-4503-6689-2/20/07.

<https://doi.org/10.1145/3311790.3396653>

systems such as very fast resonant scanning confocal (ribbon scanning), multiphoton, and fluorescence micro-optical sectioning tomography (fMOST) are all capable of generating very large multi-terabyte whole mouse brain data sets. These technologies and techniques allow researchers to probe with greater specificity of observation and increased sensitivity in both resolution and frequency illuminating cellular level detail that was previously not possible - at a cost of increased data storage. These whole brain datasets pose a model problem for data archiving, access, and analysis: they are large; they are highly structured; they possess much information that may be useful beyond initial analyses (e.g. by other investigators, with distinct questions); multiple specimens must be analyzed for each experiment; they can be interrogated at many levels of complexity: elucidating properties of single synapses, single cells, brain regions, or the whole brain; diverse datasets can be rapidly generated with a range of different cell-types using readily available genetic resources (e.g. the library of Cre-driver mice used in neuroscience); and importantly, no single “average” value is going to capture the data represented within the collection. Finally, the archival library (composed of image collections showing connectivity between many distinct cell types and across many individuals) has significantly greater value than the individual images as a resource for understanding principles of neural connectivity across brain regions, circuit architectures, developmental stages and individual variations. Providing such a collection of datasets is non-trivial. It is not unusual for individual mouse datasets to exceed 10TB of storage and full-brain marmoset datasets to exceed 800TB. Due to limitations of data storage, individual labs frequently must decide the value of data immediately after collection, and allow discretionary deletion regardless of the reuse value of the dataset to the general scientific community. Furthermore, acquiring terabyte data sets from remote sources for reuse is generally infeasible for individual labs due to barriers such as the amount of local storage required, the amount of data that must be transferred, limitations of the bandwidth and the computational power required to harmonize the datasets to a uniform analysis standard.

The Brain Image Library (BIL)[6] is a National Institutes of Health funded public resource that is solving this challenging problem by creating a collection of high-resolution high-quality brain microscopy datasets. BIL provides access to those datasets

along with high performance computing and visualization services. These storage and compute resources, which are located and managed at Pittsburgh Supercomputing Center (PSC), enable BRAIN researchers to deposit, analyze, mine, share and interact with large brain image datasets. BIL encompasses the deposition of datasets, the integration of datasets into a searchable web-accessible system, the redistribution of datasets, and a computational enclave to allow researchers to process datasets in-place and share restricted and pre-release datasets.

2 DATASETS AND METADATA

BIL is the designated repository to accept and make microscopy data publically available for the BRAIN Initiative[16]. BIL currently contains over 1000 whole brain mouse datasets, including high-resolution volumetric microscopy data from the Brain Initiative Cell Census Network[2]. Datasets from a wide variety of experimental study purposes including connectivity, cell counting, cell morphology, expression characterization and others are held in the collection. The technologies used to create the datasets include serial two-photon tomography (STPT), fluorescence micro-optical sectioning tomography (fMOST), oblique light-sheet tomography (OLST) and others. The typical whole-brain image datasets that the library receives are slice-wise multi-channel image data. A categorization of the data included in the current BIL data release is shown in Figures 1 and 2.

BIL receives the majority of its original data in the TIFF[1] or JPEG 2000[12] formats, but can receive data in any microscope format that can be converted by bioformats[13] into TIFF format. Higher level data (e.g. reconstruction of axons and dendrites) is accepted in the swc format[7]. Researchers can choose to download BIL data via anonymous ftp or interact with the data in-place on

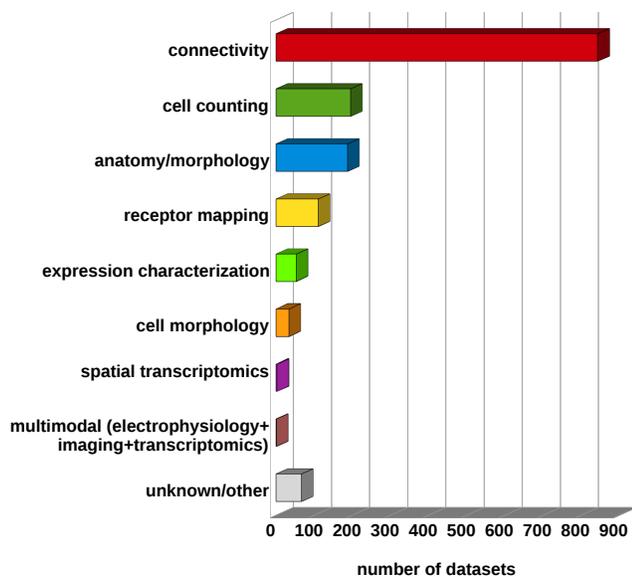


Figure 1: Modalities of datasets included in the current BIL release

high performance computing and visualization resources dedicated to the BIL project or offered on PSC’s Bridges[17] supercomputer. BIL collects several fields of descriptive metadata alongside the submitted microscopy images. Metadata fields include type of organism, its taxonomy identification number, transgenic line information, imaging method/technique/modality, and other key fields. Metadata is provided to BIL both through an ingestion portal and through template metadata spreadsheets.

3 BIL ARCHITECTURE

A primary architectural goal of BIL is to offer high-capacity, scalable storage infrastructure to house an ever-growing collection of image datasets. Thus, the overall architecture is oriented around a robust set of storage servers. Offering computational resources to locally operate on these datasets is another important consideration. We couple our available computation resources to the storage by a high-speed network. Figure 3 presents an overview of the BIL architecture, and includes resources available within the coupled Bridges supercomputer.

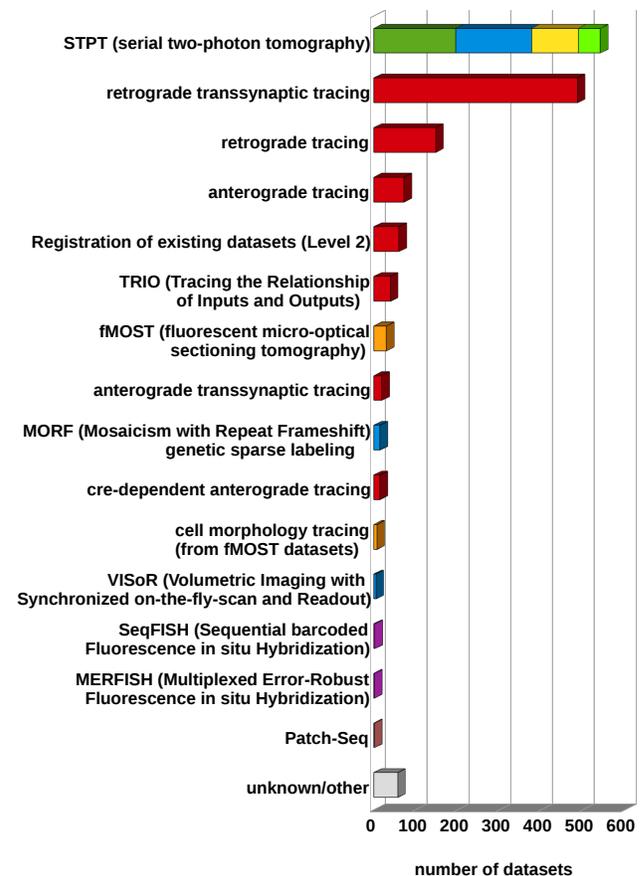


Figure 2: Techniques used to generate datasets included in the current BIL release. The colors are keyed to the modalities in Figure 1.

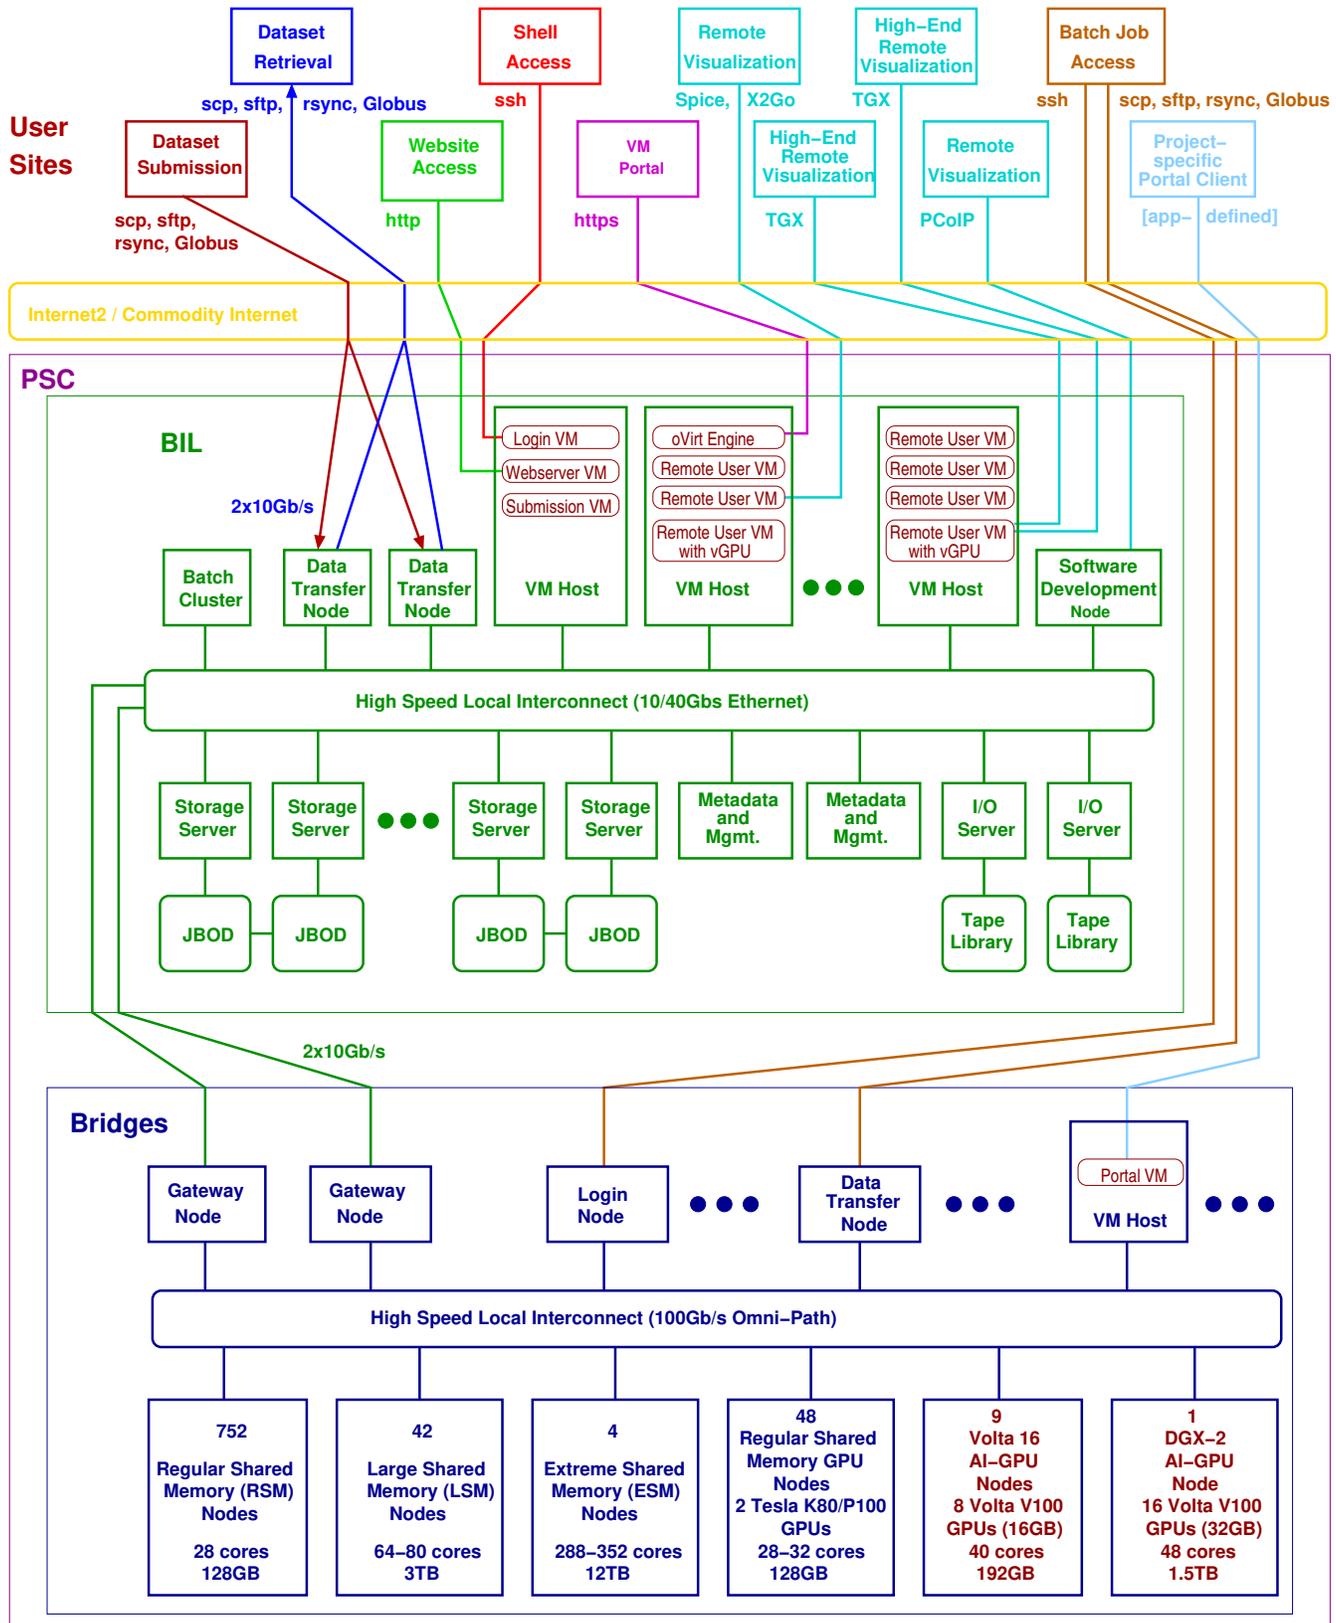


Figure 3: BIL/Bridges Architecture: These systems together support a variety of services to user sites. There is tight coupling between the two systems so that the /bil filesystem is directly mounted on the Bridges login and compute nodes.

While using the “Cloud” for storage might seem an attractive alternative to managing the storage hardware in-house, recurring costs become substantial for a multi-PB repository. While the advertised storage costs of \$0.0125/GB/month for Amazon S3™ infrequent-access storage[3] appear low at first glance, this would scale to \$25,000/month for 2PB, the current total size of files stored in the primary BIL filesystem. The other disadvantage of cloud storage is the Internet bandwidth costs and retrieval fees for accessing stored datasets. Downloads of 100TB/month from Amazon S3™ infrequent-access storage would cost an additional \$8000/month (\$1000 retrieval fee + \$7000 Internet transfer fee). BIL currently has a 20Gb/s link to the Bridges supercomputer at PSC for operating on large datasets, and will soon have a 200Gb/s link to the upcoming Bridges-2 supercomputer. This allows for much faster access than would be realistic for the Cloud, and at almost no retrieval cost.

3.1 Storage

A primary goal in the BIL project is to create a scalable filesystem with a target capacity of 10PB or larger by 2022. Since there was already extensive experience with the Lustre[14] filesystem at PSC, we wanted an architecture well-suited to supporting a filesystem of that type. For individual Lustre Object Storage Targets (OSTs) we used ZFS[5][25] with six drives in each zpool in a RAID-Z2 configuration. Thus, it would be possible to tolerate the failure of two drives and still not lose any files. Additionally, we arranged each server (running an Object Storage Server (OSS)) in a pairwise configuration with another server so that the individual disks in any zpool were distributed over 4 chassis, and thus the filesystem could continue operation even with the complete loss of a chassis. This configuration is illustrated in Figure 4. Each server can potentially access all the drives in its partner’s zpools over the SAS3 links, and so it would be possible to run both OSS’s on a single server should one server fail. As of the writing of this paper, we have not experienced any server failures, so have not had to fall back on this capability.

We initially started with two Supermicro® 5049P-E1CTR36L servers and two Supermicro® 847E1C-R1K28JBOD JBODs. These have 160 3.5in disk slots. We use 120 of these for 20 OSTs, each consisting of six 10TB Seagate® ST10000NM0016 drives. We now have a total of eight servers and eight JBODs, with the most recent pair of servers and JBODs having 16TB disk drives. Our current effective filesystem capacity is ~2.7PB with roughly 85% utilization, though we have many idle disks (>3PB raw) in reserve that can be quickly added into new OSTs to increase effective capacity to >4PB. Spinning up drives roughly a month before they are actually needed increases their longevity and reduces power consumption.

One decision we made early on was to have a single very large primary filesystem rather than multiple smaller filesystems. The reason is that it is easier to manage a single filesystem, and it allows large collections of data to be moved via Linux “mv” from one location to another (e.g. from the location of file ingest to the publicly-accessible file distribution area) virtually instantaneously.

For running VMs on the BIL local compute nodes we needed a filesystem to hold the OS images. The access patterns for this are quite different than for doing I/O on datasets, and we wanted to select a filesystem type that had been well-tested with our choice

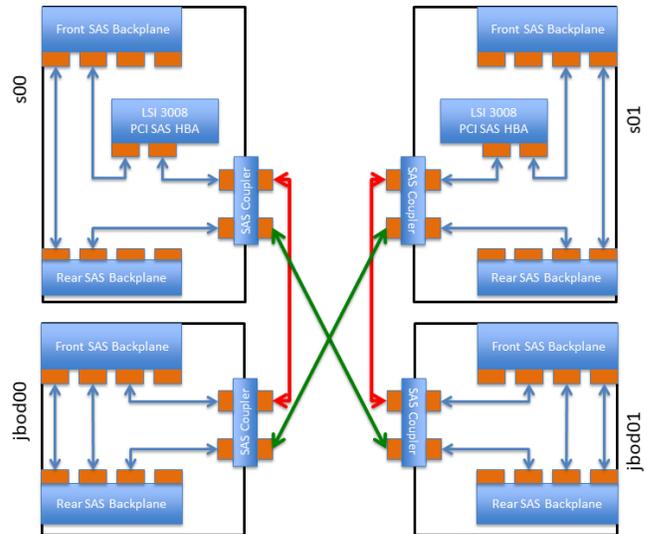


Figure 4: SAS Connections Among a Pair of Storage Nodes and JBODs

of VM software, oVirt (see below). We decided to set up a GlusterFS instance to support that need. Lustre and GlusterFS are both parallel filesystems so increasing the number of storage nodes will increase performance. In light of this, we overlaid a GlusterFS filesystem onto the set of server nodes that support Lustre, after confirming that these packages could coexist without interference on a single node, and that a single node could accommodate the CPU/memory requirements of both filesystems. In order to tolerate loss of a single server, we created a distributed replicated volume with replica 3 redundancy. Currently we use six slots holding 1TB SSDs on each of the storage nodes for the GlusterFS filesystem.

3.2 Backup and Tape Import

We chose LTO-7 tape (9TB capacity when written on LTO-8 drives) as a medium both for making archival backups of the primary BIL filesystem and for bulk exchange of data with other sites. The advantages to tape are its lower cost (~\$7/TB) compared to disk (~\$24/TB), zero power consumption during storage, and longer projected storage life. We have two Spectra Logic Spectra® Stack libraries, one located with the majority of the BIL equipment in our datacenter and one in the PSC office building for ease of dealing with tapes for export/import to and from other sites. In particular we have recently been receiving large fMOST datasets from Huazhong University of Science and Technology in China on LTO-7 tapes, which has proved more convenient than dealing with 6-12TB disks, their former transport medium. The tapes’ lighter weight in comparison with disks has also reduced shipping costs.

Archival backup tapes of all validated BIL data collections are written and stored at a secondary location to provide for data recovery. Backups and tape management are implemented with BETSOL®’s Amanda Enterprise (AE). Several factors led to the selection of AE: it is open source software; backups are written in

standard GNU tar format and can be restored without the AE software if necessary; it scales to meet the expected archive needs of BIL; AE has native support for the Spectra Logic tape library; and the software license with phone support was within the budget constraints of the project.

3.3 Computation and Visualization Resources

The BIL filesystem is mounted on PSC's Bridges supercomputer, which provides significant computational ability to process data in-place. These computing resources comprise nodes with 12 terabytes of memory, GPU nodes with NVIDIA® Tesla® K80, P100, and V100 GPUs as well as an NVIDIA® DGX-2™ to support high-bandwidth multi-GPU applications.

In addition, BIL has several servers dedicated to performing computations on datasets. Some are for internal use only (e.g. dataset validation and metadata harvesting), and include an HPE ProLiant® DL580 large-memory server with 1TB RAM and four 16-core Intel® Xeon® E7-8860 CPUs, and a cluster composed of 11 Supermicro® 1U 1042G-TF servers with 512MB RAM and four 16-core AMD Opteron™ 6376 CPUs each. Others are for remote BIL users to do processing locally. We wanted to avoid duplication of resources available on Bridges, and one capability that was not present on Bridges was the ability to do high-quality remote visualization. So, these nodes are focused on providing remote users with the resources to run interactive graphical programs on high-end (V100) GPUs, sufficient memory to accommodate large datasets (typically 500GB), and high-speed access to the BIL filesystem. Figure 5 shows an example session using Vaa3D[18–20] to view the morphology of a neuron from marmoset prefrontal cortex.

Using VMs allows us to flexibly partition these resources and insulate users from each other so that problems in one VM (e.g. overflowing memory) do not impact the other VMs. While we considered VMware and Citrix as a VM platform, we ended up going with oVirt because of systems staff familiarity with it and Red Hat®/CentOS products in general. For users needing high-quality remote desktops with fast, smooth updating, we acquired Mechdyne TGX[15] licenses. Mechdyne TGX uses the GPU to encode the video stream with low latency. We also support X2Go[23] remote desktops (software encoding) for users who prefer clients of that type.

3.4 Data Transfer Infrastructure

3.4.1 Transfer Nodes (DTNs). A fundamental component of BIL is the capability to efficiently support large file transfers with other institutions. BIL's data transfer infrastructure comprises two data transfer nodes (DTNs) integrated with wide area (WAN) and local area (LAN) network connections. The architecture is based on the Energy Sciences Network Science DMZ[8] model. Having two DTNs allows for load-sharing and redundancy to mitigate a single server failure. Each DTN has a 10G externally-facing WAN interface along with a 10G LAN interface, which provides connectivity to the other BIL components.

3.4.2 Wide Area Network. Both WAN and LAN design choices have significant impact on data throughput. The differences between each collaborator's WAN connectivity have presented challenges to data transmission between sites. In the best case scenario, the research university sites connect to Internet2's national research and

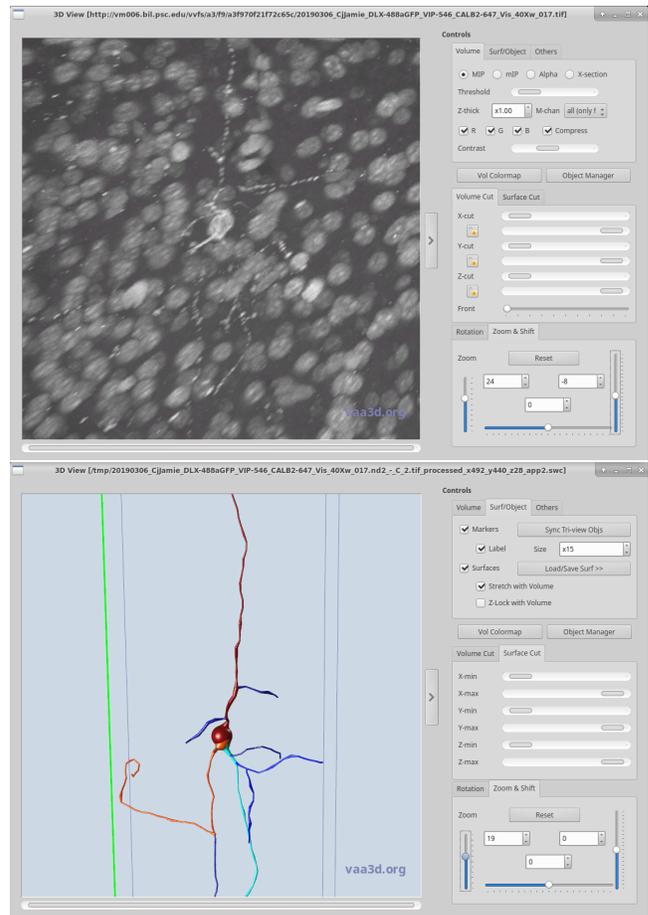


Figure 5: Using Vaa3D[18–20] to examine a neuron's morphology as imaged in a 3-D stack. The upper window shows a maximum intensity projection through the data slab. The lower window shows a depiction of the neuron's morphology obtained by manually tracing the neuron's processes through the image stack.

education (R&E) network at speeds of 10Gbps or greater with minimal traffic congestion. The Internet2 backbone itself is designed, instrumented, and closely monitored to support high-performance connectivity between sites. In contrast to the research university environment, the independent research institutions tend to be connected to the WAN at significantly lower data rates, e.g. 1Gbps via a "commodity" Internet connection. With modern servers and laptops connected at a minimum of 1Gbps, external data transfers by a single user can readily consume an organization's entire external bandwidth. The problem is compounded as multiple labs within a facility share a single 1Gbps external interface, resulting in serious contention for the limited bandwidth available.

3.4.3 Local Area Network. BIL LAN connections are Ethernet-based, with 40G interfaces to the compute nodes, two channel-bonded 10G interfaces on each of the eight storage nodes, and 10G interfaces on the DTNs, the tape server, and the remaining BIL hosts.

A data-center-class 10G/40G Juniper switch provides the internal interconnect for the BIL servers. A valuable lesson learned was that switch buffer tuning can be a critical factor in LAN performance. We discovered that the default Juniper switch configuration could not adequately support BIL traffic when multiple servers were writing to a single server (e.g., six storage servers writing backups to the one tape server). Reconfiguring the switch shared buffer settings to optimize performance for LAN traffic successfully mitigated the problem.

Since a Lustre filesystem can be run over InfiniBand™ and Intel® Omni-Path networks, we also considered those options; however, the support for these in VM contexts is much less developed (especially with oVirt), so we went with more well-tested Ethernet infrastructure. The slightly lower latency of InfiniBand™ and Omni-Path has little impact on the types of applications that BIL users run, since these typically do not rely on fine-grained synchronization across nodes.

3.4.4 Data Transfer Testing. To set realistic expectations for data transfer speeds, the BIL network team used traceroute and iperf3, commonly available network test tools, to check the path and the maximum achievable data throughput between each site. Lessons learned during this testing were: 1) the paths between sites were not all as expected and the routing needed to be corrected; and 2) sustained data transfer speeds from 1Gbps-connected sites to the 10G BIL DTNs at PSC ranged between 150-600Mbps, requiring days to transfer multi-TB data collections. We are still investigating ways to increase the data transfer rates for sites that are constrained by 1G commodity Internet connections.

3.4.5 BrainBall. Sites with limited Internet bandwidth and no LTO-7/8 tape facilities have the option to load data onto a portable storage device for shipment back to the BIL storage platform at PSC. These “BrainBall” storage devices offer 30TB of usable storage capacity in a 4x10TB RAID1 configuration, and, when augmented with an alternate set of four disks, allow 60TB to be transferred in a single shipment. BIL has also transferred processed datasets back to labs in this fashion. Most of our collaborators have attached the BrainBalls at their site via the USB3 interface. In practice, the transfer rate onto and off the BrainBalls has been somewhat mediocre (~100MB/s), so we limit their use to where Internet connections are substantially slower than this.

3.5 Account Administration & Authentication

PSC’s NSF-funded, XSEDE-allocated[22] Bridges supercomputer has direct access to the BIL Lustre filesystem. This facilitates broader computational transformations and analyses using BIL data with the extensive resources of the Bridges system. To streamline access to both BIL and Bridges systems, new users are required to provide their XSEDE User Portal username (creating one if they do not already have one[24]), which BIL administrators validate and then add to the XSEDE allocated project for BIL on the Bridges system. Once added to the XSEDE project for BIL, the XSEDE accounting infrastructure initiates account creation at PSC for the new users on the Bridges system. PSC then adds these users also to the BIL platform to enable data submission and access to BIL’s dedicated storage and computing resources. BIL users may then apply

independently for computing time and resources on the Bridges supercomputer for their computational science projects. Startup allocations may be reviewed and awarded immediately. Larger allocation requests are reviewed quarterly by the XSEDE Resource Allocations Committee[10].

BIL user accounts are maintained in PSC’s site-wide account management systems, with authentication via LDAP and Kerberos. BIL users accessing BIL and Bridges login and data transfer services authenticate using their PSC username and password. Users may also employ their XSEDE username, password and DUO credentials to access Bridges systems directly, or to authenticate for Globus GridFTP file transfer services to/from BIL and Bridges filesystems. Users at InCommon-member institutions subscribing to InCommon credential services may also employ their local institutional login credentials via CILogon[4] to authenticate for Globus GridFTP file transfer services. Published BIL datasets are available publicly for download online without user authentication.

3.6 Web Access Portals

BIL maintains web access portals for project-specific purposes and can host user-created access portals for researchers who want to provide unique ways of accessing BIL data. A key feature of the BIL architecture is that it is designed to host web portals in two ways: through VMs on PSC’s Bridges system and through BIL VMs. Bridges VMs are ideal for user-created portals and provide the flexibility to take advantage of Bridges high-performance computational resources. BIL VMs are useful for applications that require access to GPU resources to serve image data.

4 FUTURE PLANS

In the near future, BIL plans to enhance the remote desktop environment, make BIL data available on PSC’s new Bridges-2 high performance computing platform and implement several data usability enhancements, including the accommodation of additional data types, issuing DOIs for datasets, and integrating social media.

4.1 Remote Desktop Environment

Currently, we have statically created VMs for external users that have expressed an interest in a graphical desktop environment with high-speed access to BIL datasets. We are planning on moving to a scheme that requires less manual intervention for each new user. Thus, we are exploring the use of a connection broker for dynamic VM assignment. This would also be able to reclaim resources that are idle for long periods.

4.2 Bridges-2

PSC received an award from NSF in 2019 to develop and build the Bridges-2 platform, a \$10-million dollar system designed for data and compute-intensive research, AI, and non-traditional HPC communities. This system, expected to be deployed in the summer of 2020 will replace the current Bridges system and provide for significantly more computational power and ease of use over the current Bridges system. When this system is deployed, it will include direct access to the BIL filesystem in a manner similar to the current Bridges system.

4.3 Data Enhancements

At present, we keep the datasets in the identical format and with the same directory and file structure in which they were uploaded by the generating lab. For example, some datasets consist of .tiff files and others consist of .jp2 files. Whatever file naming conventions were used by that lab are also visible to people looking through the dataset directories, and these conventions can vary considerably from lab to lab. We plan on offering ways for dataset consumers to download datasets in a more customizable way, for instance, to retrieve a 3-D volumetric dataset as a collection of LZW-compressed .tiff images with sections numbered from 0000-1499, with the x, y, z axes oriented in a standard way.

In the past month BIL has begun to receive data from other imaging modalities including Spatial Transcriptomics and Patch-Seq. We expect an increasing number of submissions of this type, and will be developing better procedures for indexing and working with these datasets. Of these data modalities Spatial Transcriptomics presents the larger challenge due to the increasing combination of microscopes and rapidly developing methods available (such as MERFISH[26] and seqFISH[21]).

The BIL project is working on providing all BIL datasets with digital object identifiers (DOIs)[11] to enhance the FAIRness[9] of BIL data, by providing a globally universal persistent identifier. DOIs will provide additional advantages including allowing cleaner citations of datasets in peer-reviewed papers and providing the archive with the ability to use citation-tracking methods to measure dataset citations and dataset engagement (such as twitter posts) about the dataset.

Finally, we intend to integrate social media into the data release process. Platforms such as Twitter provide an easy way to follow and tag items as well as give information on the reach and activity of access to the datasets. Twitter's API gives developers helpful information on traffic and clicks to the data provided through the portal.

ACKNOWLEDGMENTS

The Brain Image Library is funded by a grant from the National Institute of Mental Health of the National Institutes of Health under award number R24-MH-114793 and is operated as a partnership between the Biomedical Applications Group at the Pittsburgh Supercomputing Center, the Center for Biological Imaging at the University of Pittsburgh and the Molecular Biosensor and Imaging Center at Carnegie Mellon University.

This work used the Extreme Science and Engineering Discovery Environment (XSEDE), which is supported by National Science Foundation grant number ACI-1548562. Specifically, we used the Bridges supercomputer at the Pittsburgh Supercomputing Center through allocation TG-IBN180004.

REFERENCES

- [1] Adobe. 2020. *TIFF*. <https://www.adobe.io/open/standards/TIFF.html>
- [2] Allen Institute. 2020. *Brain Cell Data Center (BCDC)*. <https://biccn.org/>
- [3] Amazon Web Services. 2020. *Amazon S3 Simple Storage Service pricing*. <https://aws.amazon.com/s3/pricing/>
- [4] Jim Basney, Heather Flanagan, Terry Fleury, Jeff Gaynor, Scott Koranda, and Benn Oshrin. 2019. *CILogon: Enabling Federated Identity and Access Management for Scientific Collaborations*. In *Proceedings of the International Symposium on Grids and Clouds (ISGC) (PoS(ISGC2019)031)*. <https://doi.org/10.22323/1.351.0031>
- [5] Brian Behlendorf. 2011. *ZFS on Linux for Lustre*. http://cdn.opensfs.org/wp-content/uploads/2012/12/1130-1200_Brian_Behlendorf_LUG11_ZFS_on_Linux_for_Lustre.pdf
- [6] Brain Image Library. 2020. *Brain Image Library: A public resource enabling researchers to deposit, analyze, mine, share and interact with large brain image datasets*. <http://www.brainimaginglibrary.org/>
- [7] R. C. Cannon, D. A. Turner, G. K. Pyapali, and H. V. Wheal. 1998. An on-line archive of reconstructed hippocampal neurons. *J. Neurosci. Methods* 84, 1-2 (Oct 1998), 49–54.
- [8] Energy Sciences Network (ESnet). 2020. *Science DMZ*. <https://fasterdata.es.net/science-dmz/>
- [9] GO FAIR. 2020. *FAIR Principles*. <https://www.go-fair.org/fair-principles/>
- [10] David Hart and Ken Hackworth. 2019. *XSEDE Resource Allocation Committee (XRAC) Reviewer Manual*. <https://www.ideals.illinois.edu/handle/2142/104685>
- [11] International DOI Foundation. 2020. *Digital Object Identifier System*. <https://www.doi.org/>
- [12] Joint Photographic Experts Group. 2020. *JPEG - JPEG 2000*. <https://jpeg.org/jpeg2000/>
- [13] M. Linkert, C. T. Rueden, C. Allan, J. M. Burel, W. Moore, A. Patterson, B. Loranger, J. Moore, C. Neves, D. Macdonald, A. Tarkowska, C. Sticco, E. Hill, M. Rossner, K. W. Eliceiri, and J. R. Swedlow. 2010. Metadata matters: access to image data in the real world. *J. Cell Biol.* 189, 5 (May 2010), 777–782. <https://doi.org/10.1083/jcb.201004104>
- [14] lustre.org. 2020. *Lustre*. <http://lustre.org>
- [15] Mechdyne Corporation. 2020. *TGX - High Performance Remote Desktop*. <https://tgxremotedesktop.com/>
- [16] National Institutes of Health. 2020. *Brain Initiative*. <https://braininitiative.nih.gov/>
- [17] Nicholas A. Nystrom, Michael J. Levine, Ralph Z. Roskies, and J. Ray Scott. 2015. Bridges: A Uniquely Flexible HPC Resource for New Communities and Data Analytics. In *Proceedings of the 2015 XSEDE Conference: Scientific Advancements Enabled by Enhanced Cyberinfrastructure* (St. Louis, Missouri) (XSEDE '15). ACM, New York, NY, USA, Article 30, 8 pages. <https://doi.org/10.1145/2792745.2792775>
- [18] Hanchuan Peng, Alessandro Bria, Zhi Zhou, Giulio Iannello, and Fuhui Long. 2014. Extensible visualization and analysis for multidimensional images using Vaa3D. *Nature Protocols* 9, 1 (01 Jan 2014), 193–208. <https://doi.org/10.1038/nprot.2014.011>
- [19] Hanchuan Peng, Zongcai Ruan, Fuhui Long, Julie H. Simpson, and Eugene W. Myers. 2010. V3D enables real-time 3D visualization and quantitative analysis of large-scale biological image data sets. *Nature Biotechnology* 28, 4 (01 Apr 2010), 348–353. <https://doi.org/10.1038/nbt.1612>
- [20] Hanchuan Peng, Jianyong Tang, Hang Xiao, Alessandro Bria, Jianlong Zhou, Victoria Butler, Zhi Zhou, Paloma T. Gonzalez-Bellido, Seung W. Oh, Jichao Chen, Ananya Mitra, Richard W. Tsien, Hongkui Zeng, Giorgio A. Ascoli, Giulio Iannello, Michael Hawrylycz, Eugene Myers, and Fuhui Long. 2014. Virtual finger boosts three-dimensional imaging and microscopy as well as terabyte volume image visualization and analysis. *Nature Communications* 5, 1 (11 Jul 2014), 4342. <https://doi.org/10.1038/ncomms5342>
- [21] seqFISH. 2020. *Technology - seqFISH*. <https://www.seqfish.com/technology>
- [22] J. Towns, T. Cockerill, M. Dahan, I. Foster, K. Gaither, A. Grimshaw, V. Hazlewood, S. Lathrop, D. Lifka, G. D. Peterson, R. Roskies, J. R. Scott, and N. Wilkins-Diehr. 2014. XSEDE: Accelerating Scientific Discovery. *Computing in Science & Engineering* 16, 5 (Sept.-Oct. 2014), 62–74. <https://doi.org/10.1109/MCSE.2014.80>
- [23] x2go.org. 2020. *X2Go - everywhere@home*. <https://wiki.x2go.org/doku.php>
- [24] XSEDE. 2018. *Getting Started - XSEDE*. <https://www.xsede.org/for-users/getting-started>
- [25] ZFS on Linux. 2020. *FAQ - zfs/linux/zfs Wiki*. <https://github.com/zfs/linux/zfs/wiki/FAQ>
- [26] Xiaowei Zhuang. 2020. *Zhuang Research Lab*. <http://zhuang.harvard.edu/merfish.html>