

Tiered Data Storage

Minnesota Supercomputing Institute



UNIVERSITY OF MINNESOTA

GlobusWorld 2017
Jeffrey McDonald, PhD
Assistant Director for Operations

© 2015 Regents of the University of Minnesota. All rights reserved.

Minnesota Supercomputing Institute
MSI



CASE STUDY

The Minnesota Supercomputing Institute upgrades to Spectra Logic BlackPearl, with LTO-7 tape drives and media, to its enterprise high-performance computing research data archive

“ At the University of Minnesota, Spectra Logic’s T950 and BlackPearl are important components in our strategic and comprehensive storage plan for hundreds of terabytes of critical research data. ”

*Jeffrey McDonald, Assistant Director for HPC Operations,
Minnesota Supercomputing Institute, University of Minnesota Twin Cities*



*Spectra T950
dual-frame Tape
Library and
BlackPearl Deep
Storage Gateway*

About Minnesota Supercomputing Institute

The Minnesota Supercomputing Institute (MSI) addresses the

The Minnesota Supercomputing Institute (MSI) is committed to expanding and developing the services it offers in order to continue to play a key support role across the growing spectrum of scientific innovation. The institute is dedicated to facilitating university and industry collaboration and to promoting technology transfer through the interchange of ideas in the field of supercomputing research, including the dissemination of research results accomplished with MSI resources. MSI utilizes Globus connect software so university data can be transferred easily.

<https://edge.spectralogic.com/index.cfm?fuseaction=home.displayFile&DocID=4674>

UNIVERSITY OF MINNESOTA

© 2015 Regents of the University of Minnesota. All rights reserved.

Minnesota Supercomputing Institute
MSI

Tape as a Archvie

- My collaborators have frequently made tape for a lot of jokes.
- IT folks prefer a lower-cost solution.
- IT folks have banked on cloud storage or other spinning solutions.
- But the costs are high and there is a cost/benefit trade-off.



The Empire has chosen Tape!



Minnesota Supercomputing Institute
MSI



HPC Resources



- **MSI Users**
 - PI Accounts: 700
 - Users: > 3000
- **Mesabi**
 - Cores: > 18,000
 - Memory: 67 TB
 - Accelerators: 80 K40 gpGPUs
 - Peak: 675 TF
 - 320 Gbps to Storage
- **New Technologies**
 - FPGAs
 - Nvidia GPUs
 - Intel PHI
 - Storage (Intel NVME)

UNIVERSITY OF MINNESOTA

© 2015 Regents of the University of Minnesota. All rights reserved.

Minnesota Supercomputing Institute
MSI



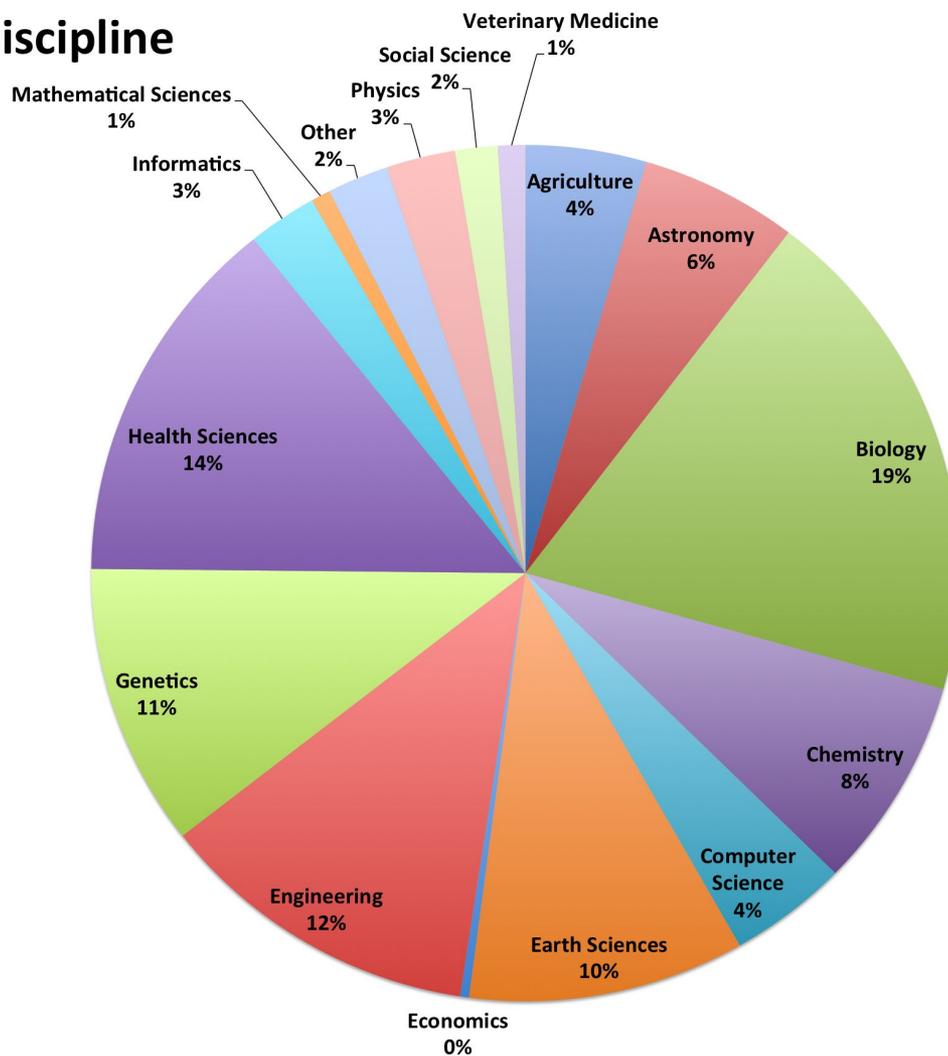
Storage

- High Performance Storage
 - 3.1 PB Usable
 - 48GB/s read/write
 - Available on HPC resources
- Tier-2 Storage
 - 3.1 PB Usable
 - Available via Amazon's S3 interface
 - Available anywhere in the world
- Archive Storage
 - > 3.5 PB tape-based storage
 - offline storage



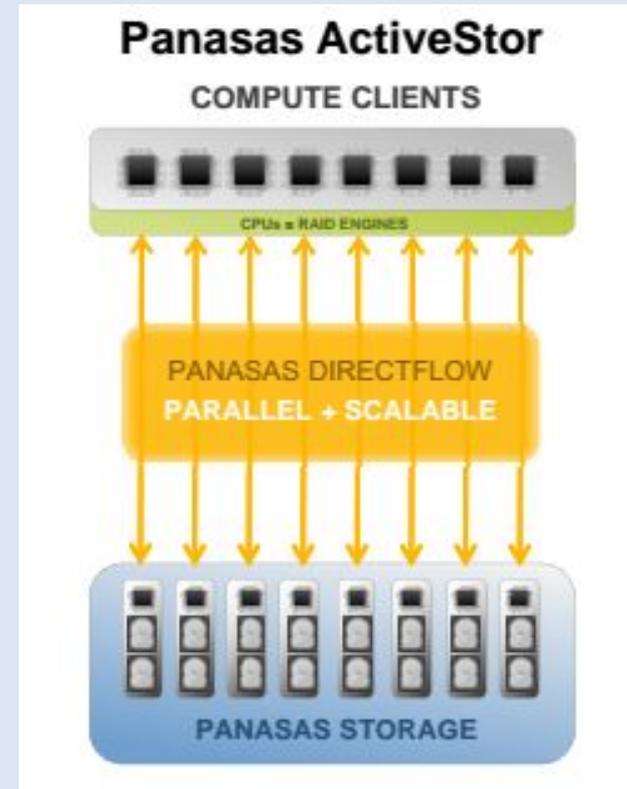
Who Uses MSI?

Storage Allocated by Discipline

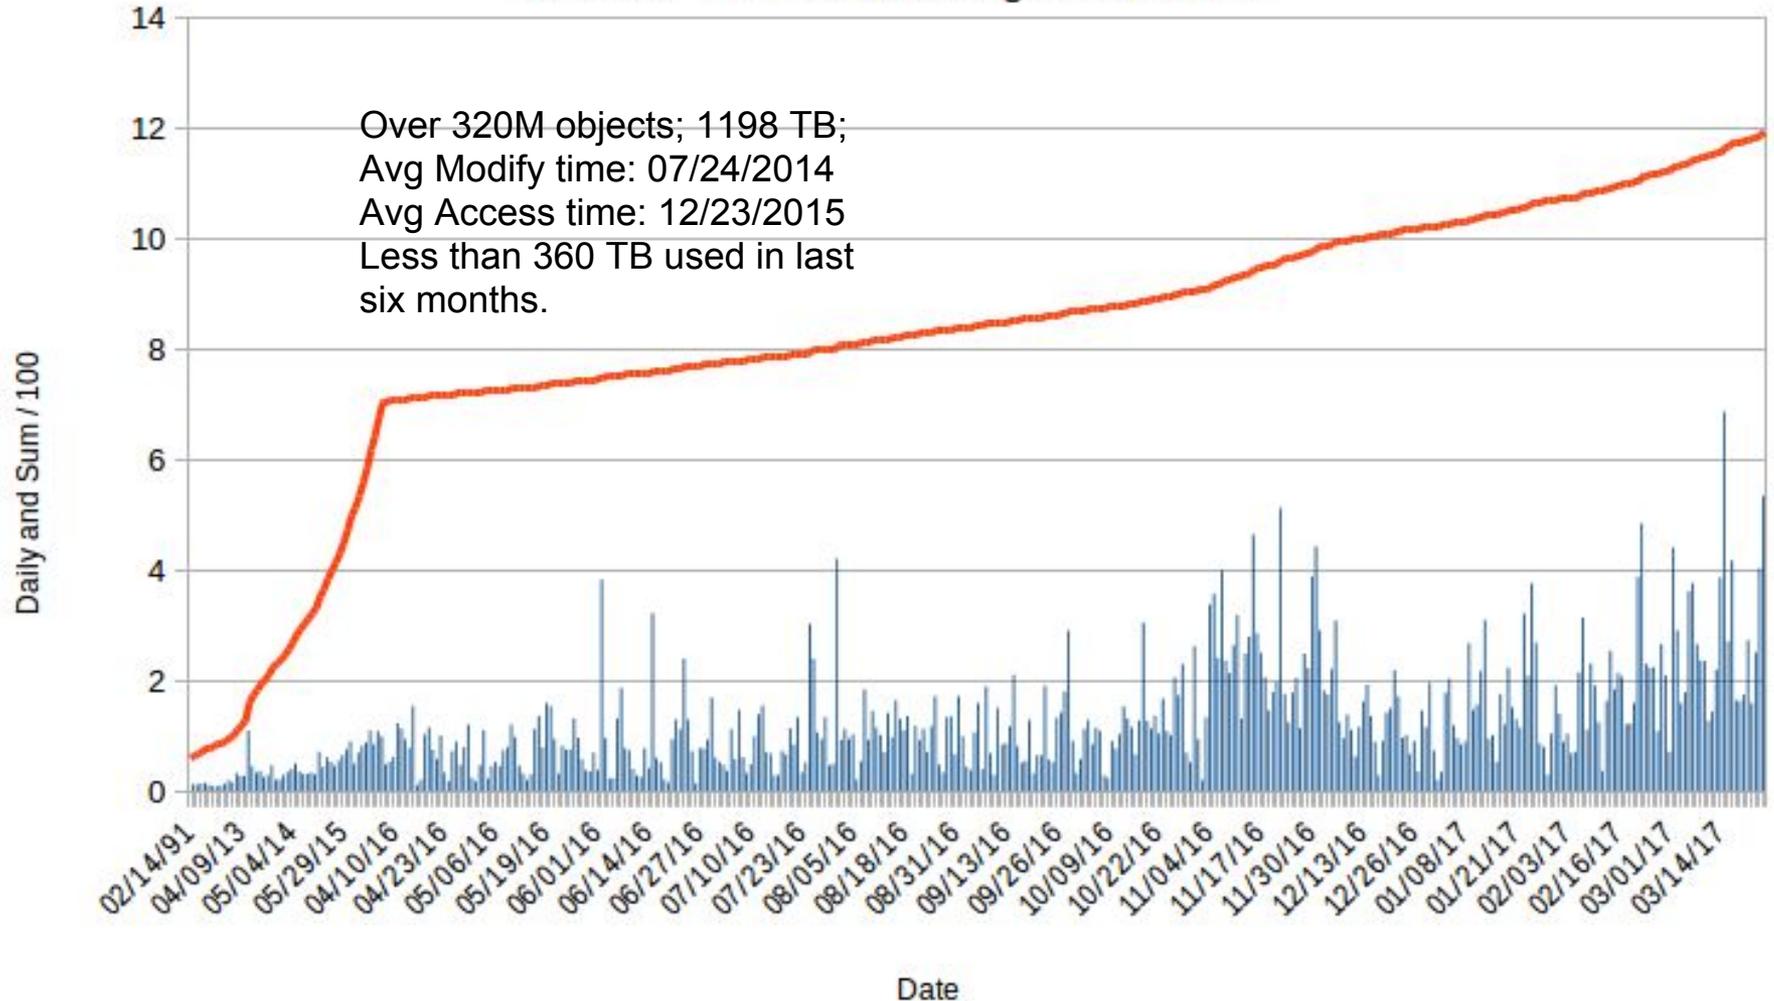


Global Panasas Storage

- 32 shelves of Panasas
- Each shelf has 10 storage + 1 director blade
- Each shelf is capable of 1.5 GB/sec and 50-70kIOPS
- System as a whole > 1 MIOPS, 48 GB/sec
- 4.5 PB; Globally available > 2000 clients



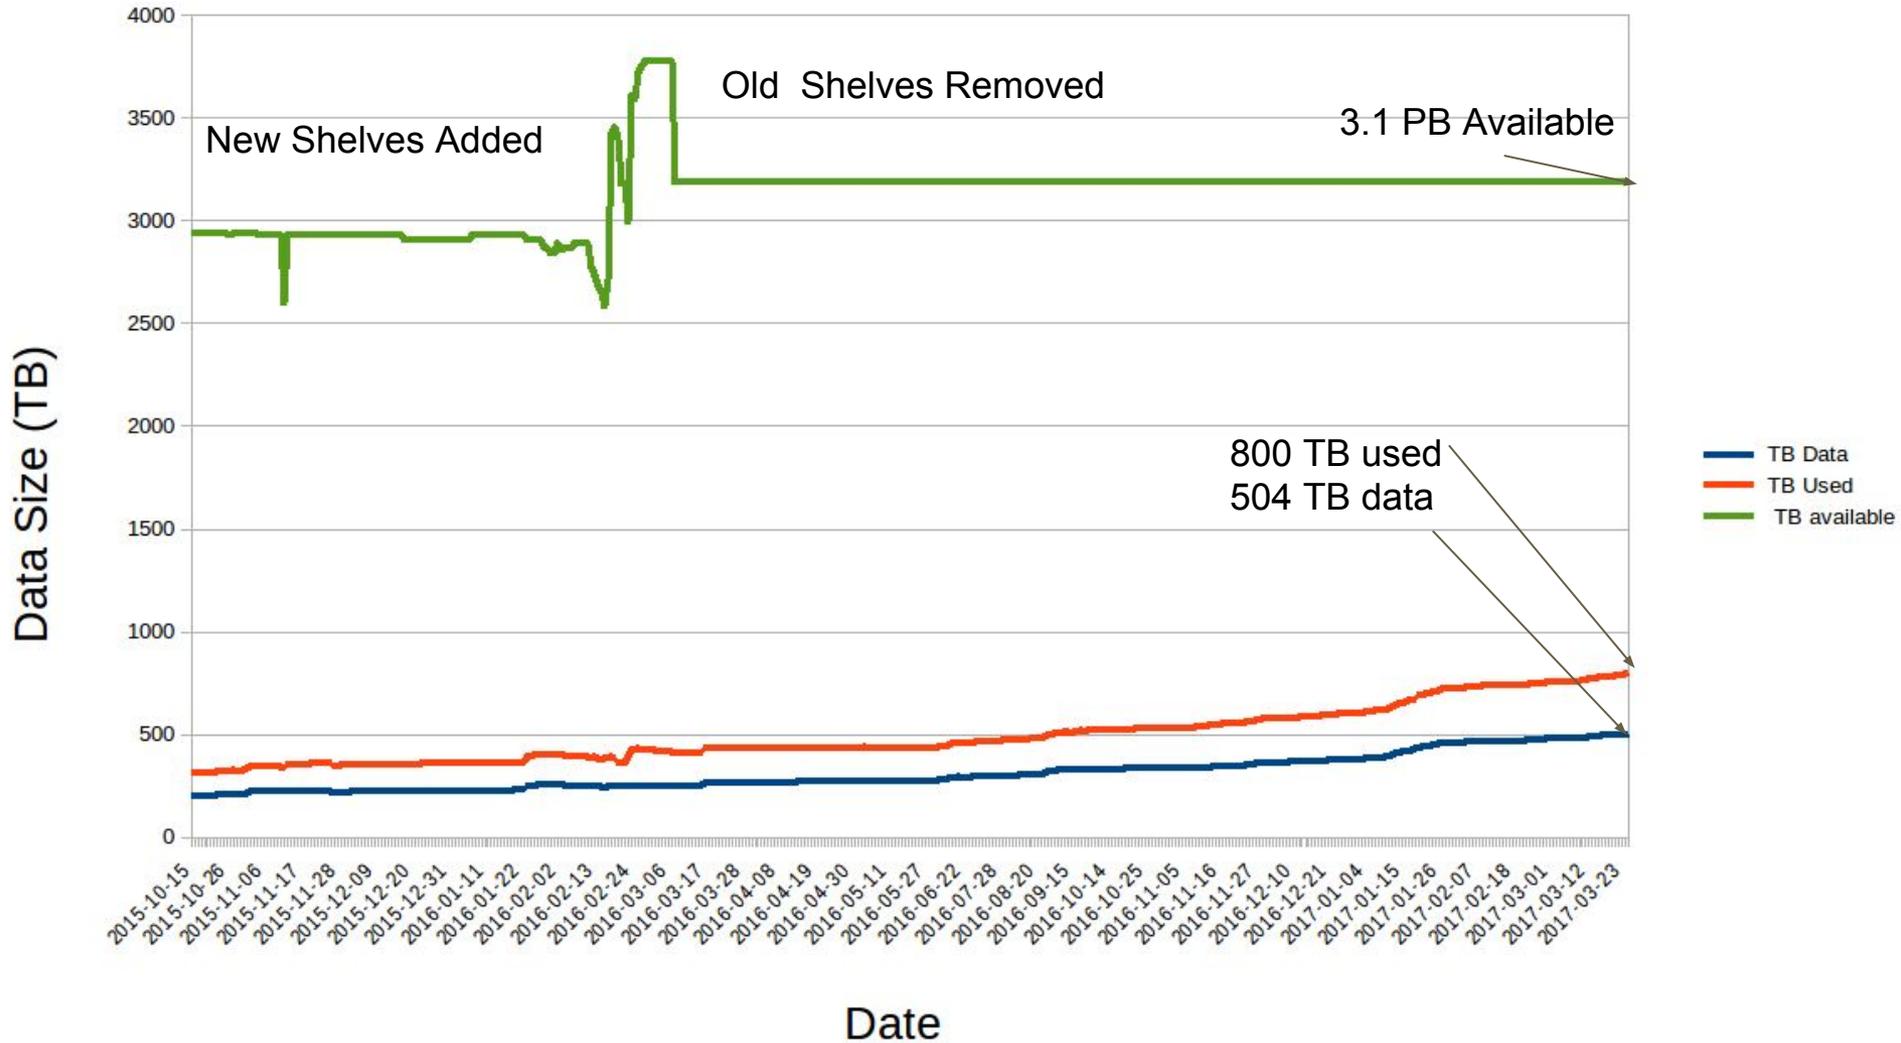
Panasas Tier 2 Data Storage March 2017



Tier Two Storage

- 2015 – Developed In-house expertise with CEPH (Object File Storage)
- 2016 – Deployed 3.1 PB of CEPH Storage available via an S3 interface.
- Deployed node “bricks” of 60 x drive systems with 12 SSDs System
- Multiple 10 GbE per “brick”
- Software-based Erasure Encoding 4+2 (RAID-6)
- Community is warming to CEPH (MSI presented BOF at Supercomputing 2015)

Ceph Storage at MSI



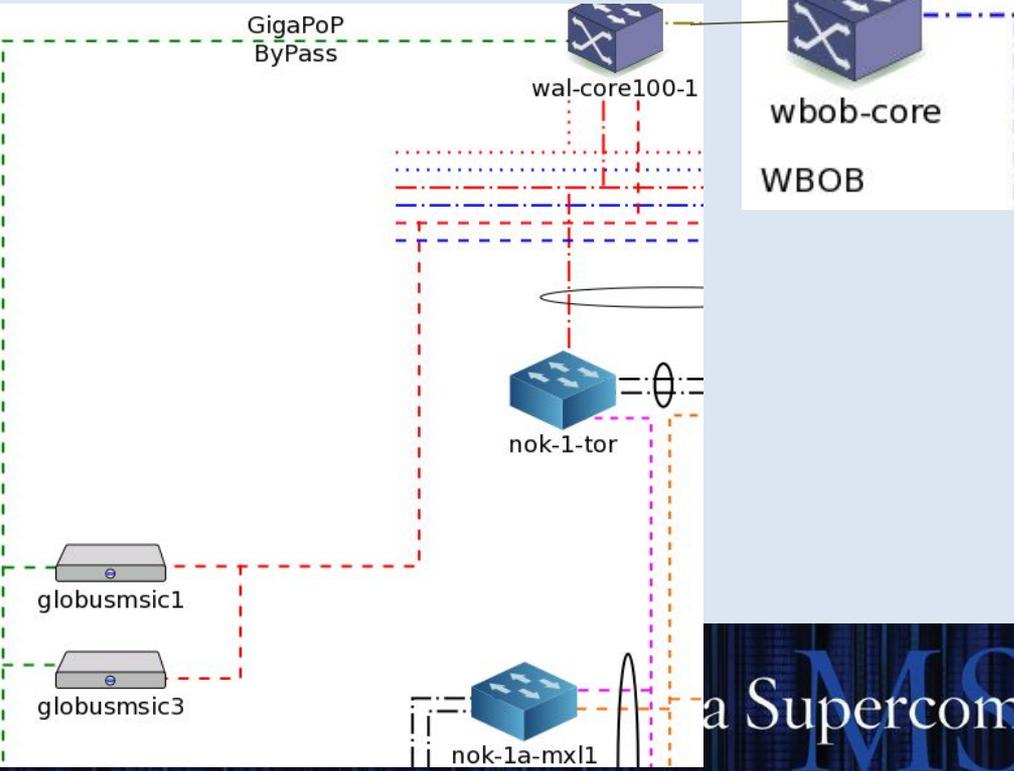
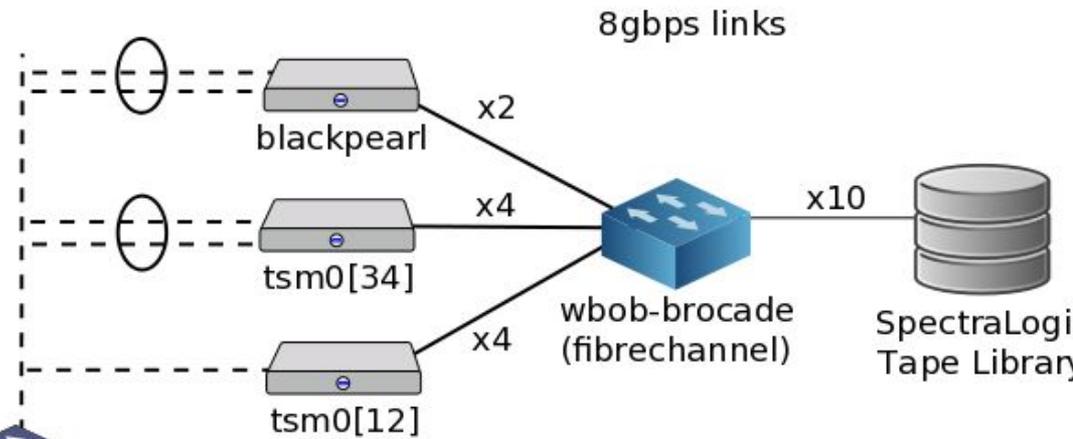
Tier Three Storage

- Blackpearl currently piloting program
- Media costs are as low as \$4/TB/year (5 year retention -- can be much lower based upon tape retention).
- Blackpearl device from Spectralogic to allow for self-service use of the tape archive. (Up to 1 GB/sec transfer rates.)
- Globus connects users to Blackpearl

Why use tape?

- 1000x more reliable than a USB drive
- 100x more reliable than an enterprise class drive
- Tape capacities are increasing and cost of maintaining storage is only 10% of initial acquisition (new media cost/TB)
- 15 years of lifetime or 5000 tape mounts or 200 re-writes
- Tapes don't need power or cooling.

GLOBUS Network Diagram



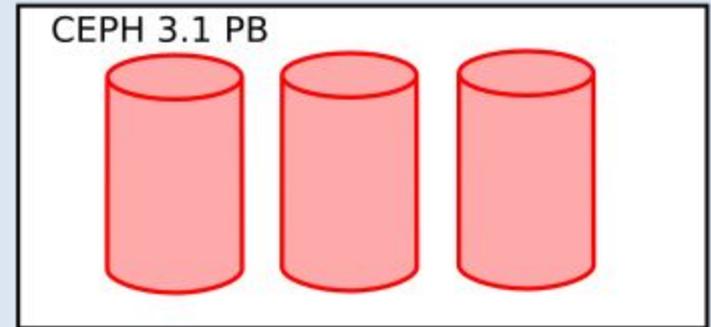
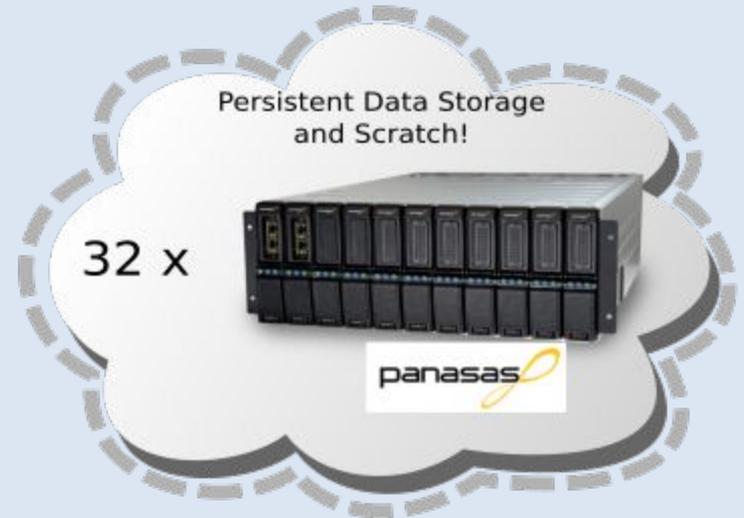
Globus Endpoints connected at 10 GBE, firewall bypass. Switch to switch on dedicated 40 GbE lane.

© 2015 Regents of the University of Minnesota. All rights reserved.

Tier 3 policies

- One 'bucket' per group
- Tools to chunk data to adequate sizes
 - Duplicity
 - Tar
- Must purchase 5 years of storage
- May manipulate archive
- Globus provides data movement channel

MSI Storage Access 2017



Blackpearl Storage System
2 x LTO7 tape drives

© 2017 Regents of the University of Minnesota. All rights reserved.

UNIVERSITY OF MINNESOTA

Minnesota Supercomputing Center
MSI

Conclusions/Future

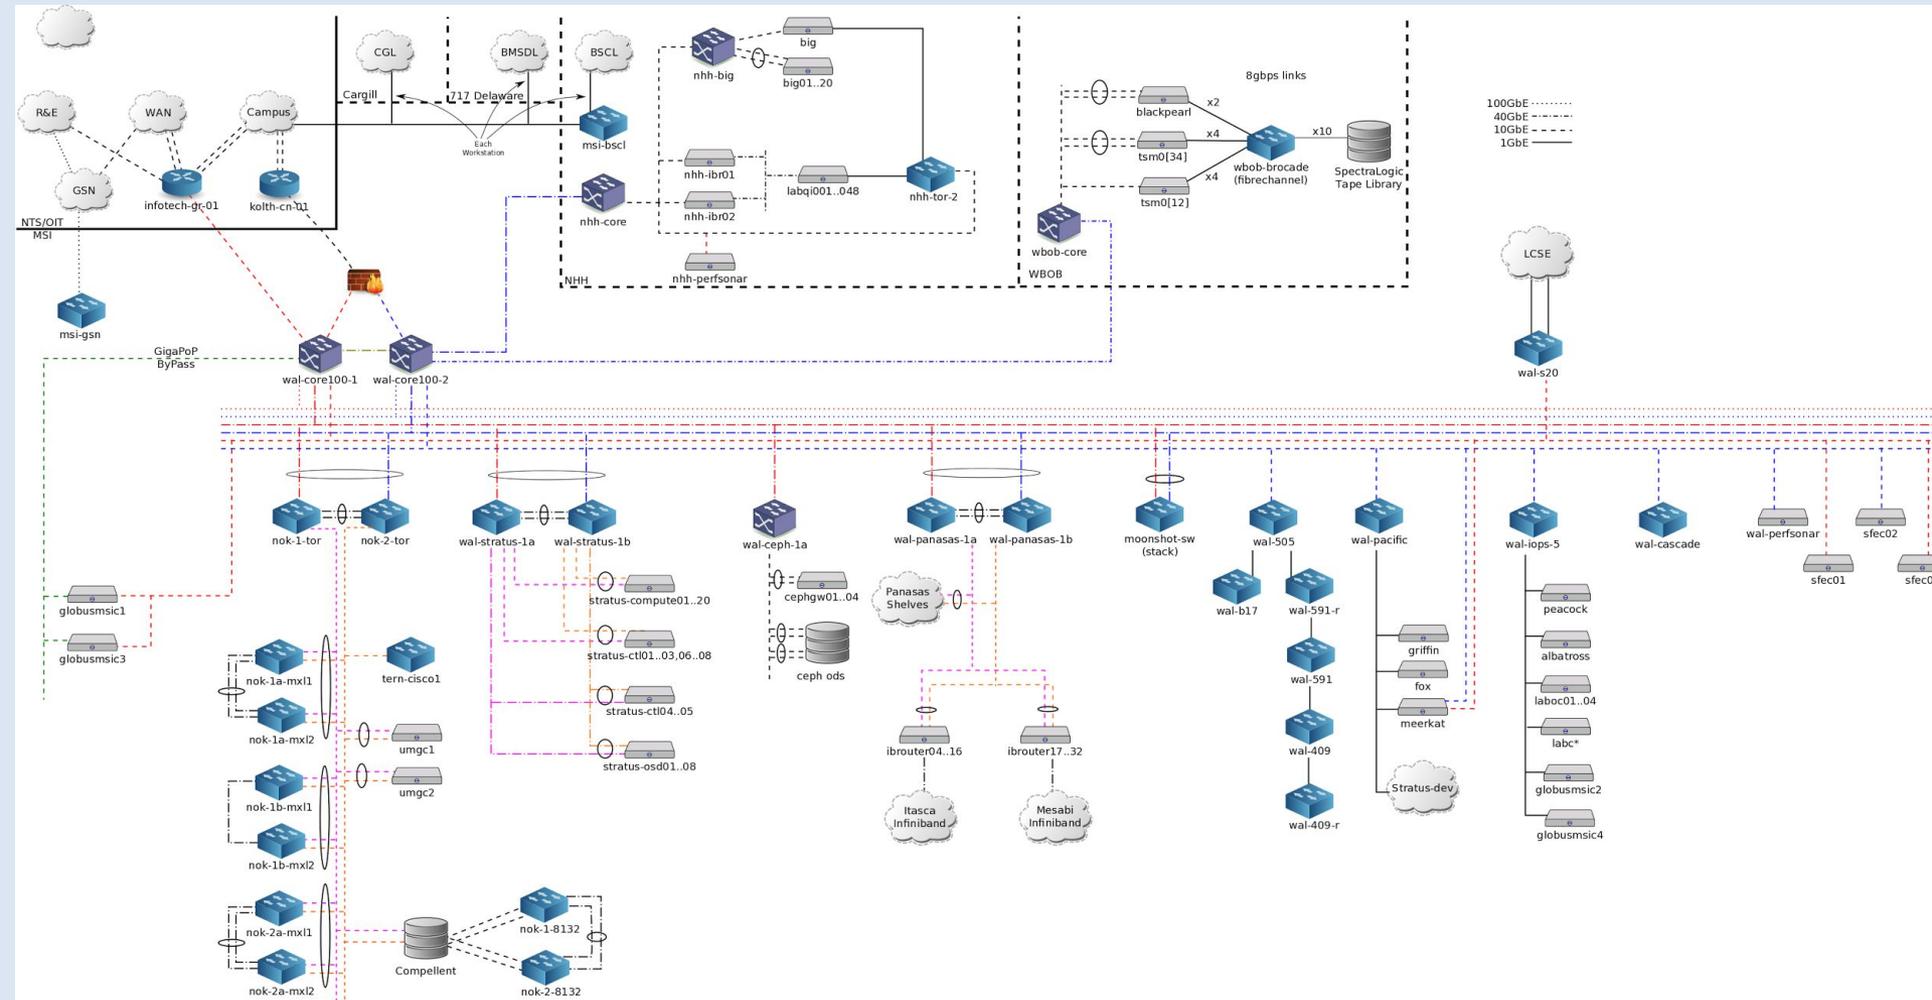
- MSI will launch provider service for Tier-3 archive storage in July 2017.
- Keys will be managing user expectation and MSI is developing policies to do just that.
- Discussing possibly offer service University-wide.
- Stay tuned.

Backup Slides

UNIVERSITY OF MINNESOTA

© 2015 Regents of the University of Minnesota. All rights reserved.

Minnesota Supercomputing Institute
MSI



UNIVERSITY OF MINNESOTA

© 2015 Regents of the University of Minnesota. All rights reserved.

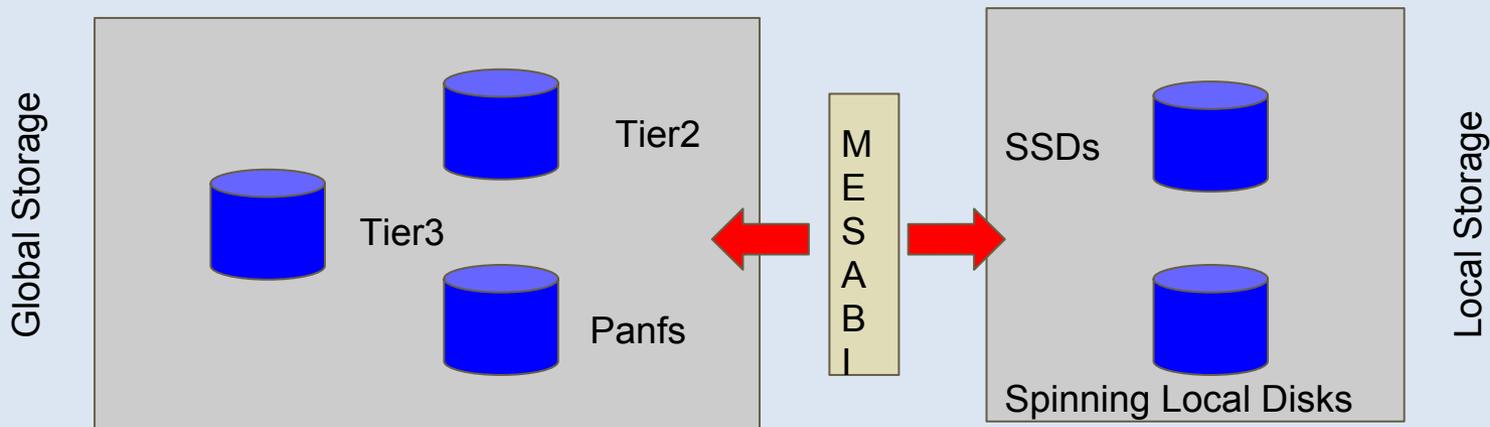
Minnesota Supercomputing Institute
MSI

HPC Center Differences

- MSI provides high performance computing resources
- Less emphasis on reliability (1 or 2-9s of uptime) typically 97% or better.
- Our users require high performance global storage shared with > 2k clients.
- MSI tolerates a high degree of risk in the deployment of systems.

Big Data What does it mean?

- Lines are blurred at MSI.
- Each Mesabi node has access to multiple storage platforms.

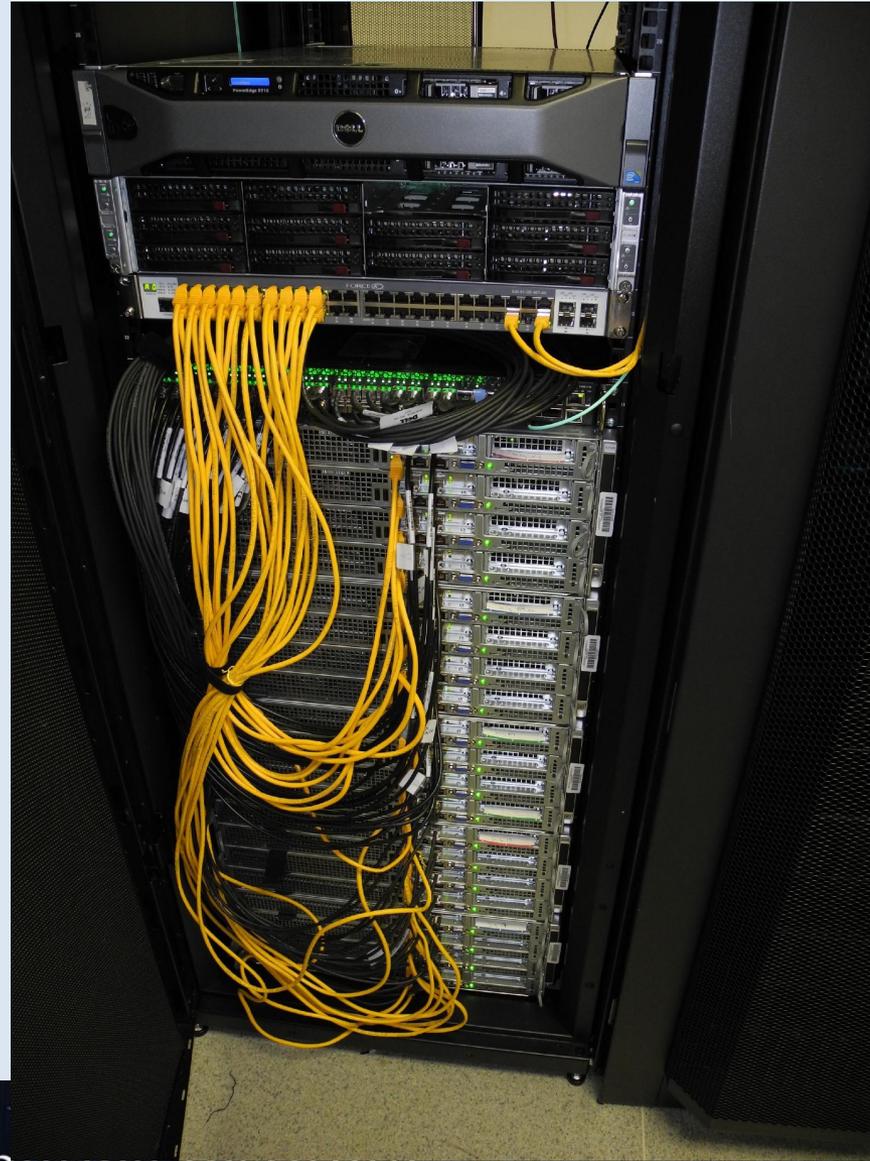


Big Data Requirements

- Sustainable Storage Model to provide the PBs needed.
 - Data Retention Policies
 - Multiple Storage Options
- Compute Resources to analyze
- Bandwidth to Move Data Around
- Share the results

Persistent Big Data Resources

- 20 SuperMicro nodes
E5-2680v2 10 core
processors; 128 GB
ram.
- Each node has 12 x 6
TB SATA3 drives, 1.2
PB of aggregate
space.
- 10 GbE network.



UNIVERSITY OF MINNESOTA

Minnesota Supercomputing Institute



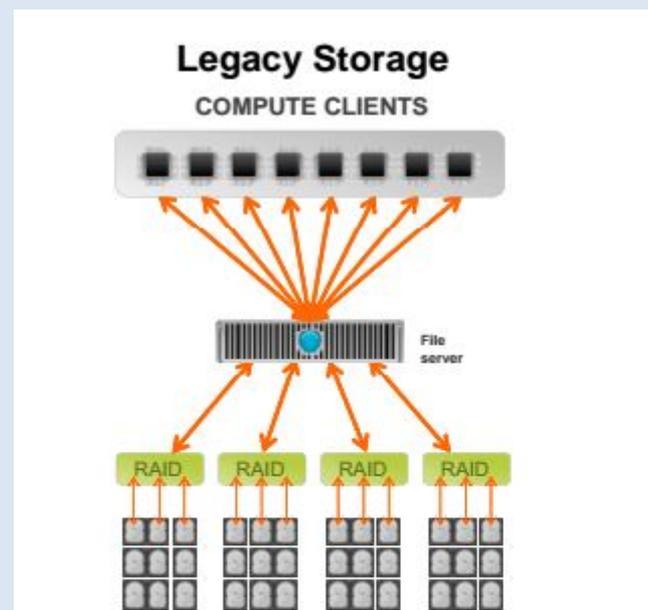
Storage Soap Box

- RAID is challenged by large disk units, rebuild time and filesystem metadata limits.
- Object based storage is scalable and uses “RAO”= Redundant Array of Objects.
- Users with data want to share that data.



Storage Soap Box

- NAS storage is always limited by a metadata server which is a single point of failure.
- RAID units scale as the disk size, all blocks must be rebuilt. No intelligence in most systems.
- Distributed Model needed.



Analysis of Global Storage

- Monthly home directory snapshots of MSI storage (stat of each file), 250M objects per snapshot
- These are aggregated on a hadoop file system and Apache PIG is used to perform the analysis.
- 20 months of snapshots used in analysis up to April 2016. Over 5 B records (2.6 TB of data).

Hadoop System Applications

- Use Hadoop version 2.7.1
- Pig script version 0.15.0
- These Applications Run on our 1.2 PB/20 node cluster
- Spark available
- Possible to run on an HPC system with a non-persistent HDFS

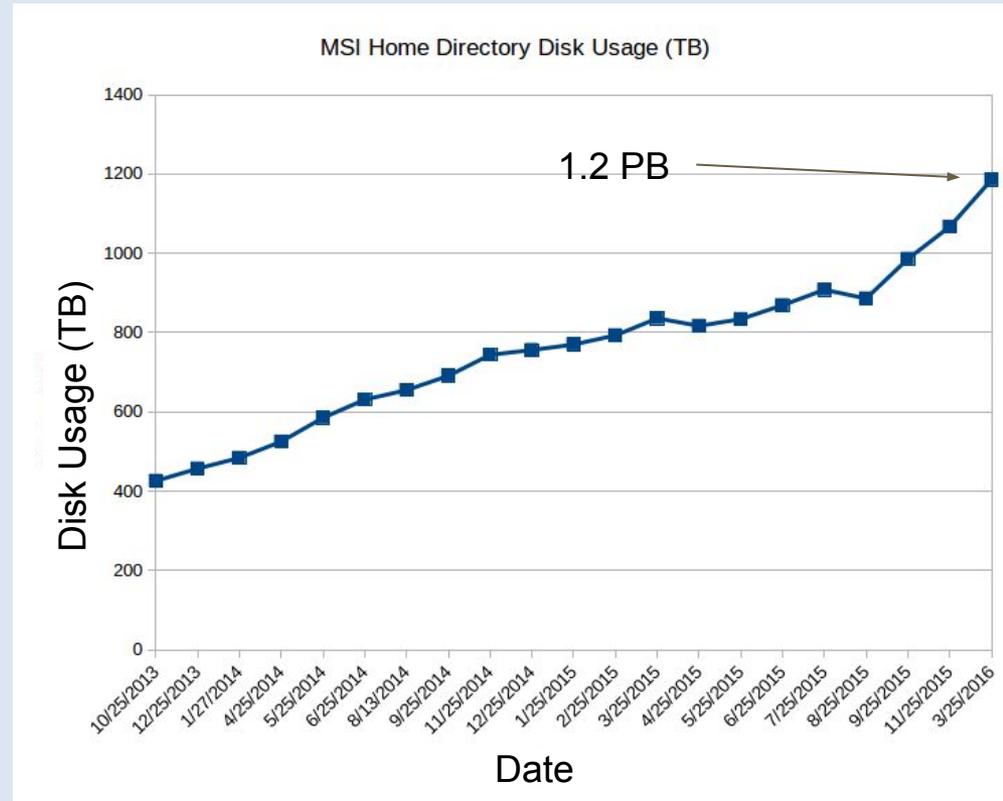
Example Pig Script

```
filestat= LOAD '...' using org.apache.pig.piggybank.storage.CSVExcelStorage(',',  
YES_MULTILINE', 'UNIX', 'SKIP_INPUT_HEADER') AS (fields...);  
fstat_all = FILTER filestat by gid > 0;  
fstat_all2 = FILTER fstat_all by mtime > 0L and mtime < 1461579014L;  
fstat_all3 = GROUP fstat_all2 by gid;  
my_table_distinct = foreach fstat_all3 GENERATE group, COUNT(fstat_all2) as numfiles,  
SUM(fstat_all2.size) as rawsize, SUM(fstat_all2.blocks) as sizeblocks, AVG(fstat_all2.mtime)  
as mtime1, AVG(fstat_all2.atime) as atime1;  
ordered_table = ORDER my_table_distinct by mtime1 DESC;  
dump ordered_table;
```

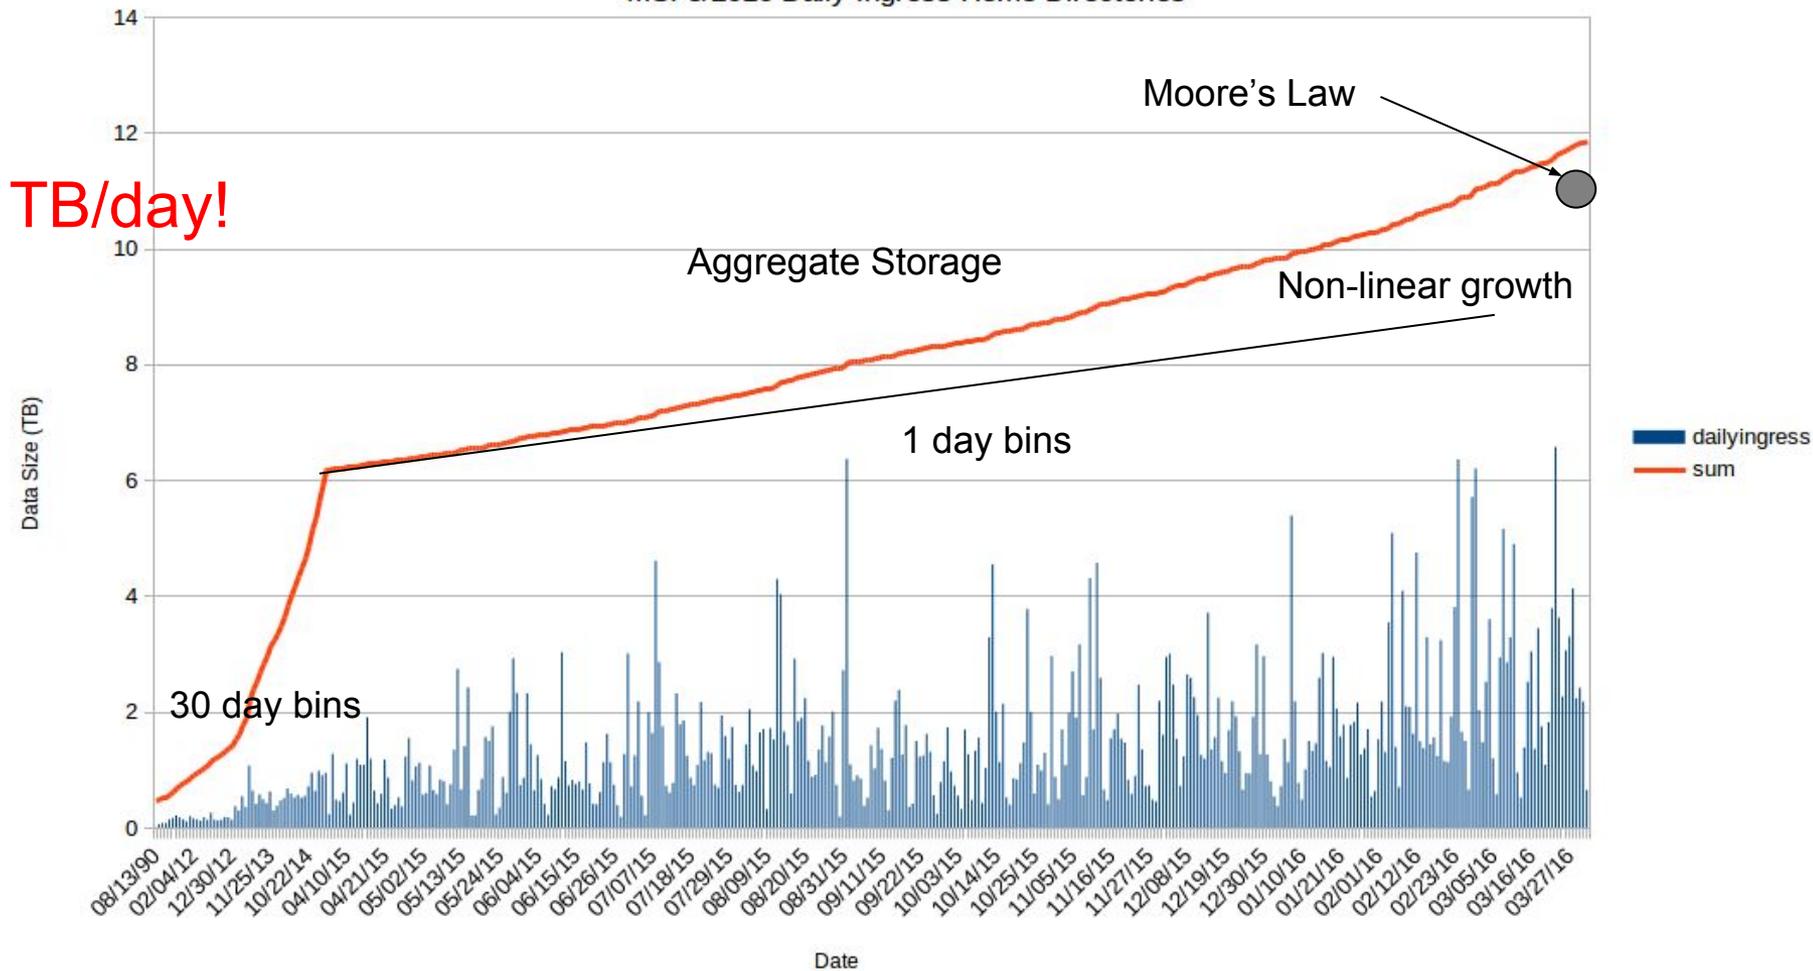
No Parallel Programming Needed!
Experience Needed!

Research Data Experience

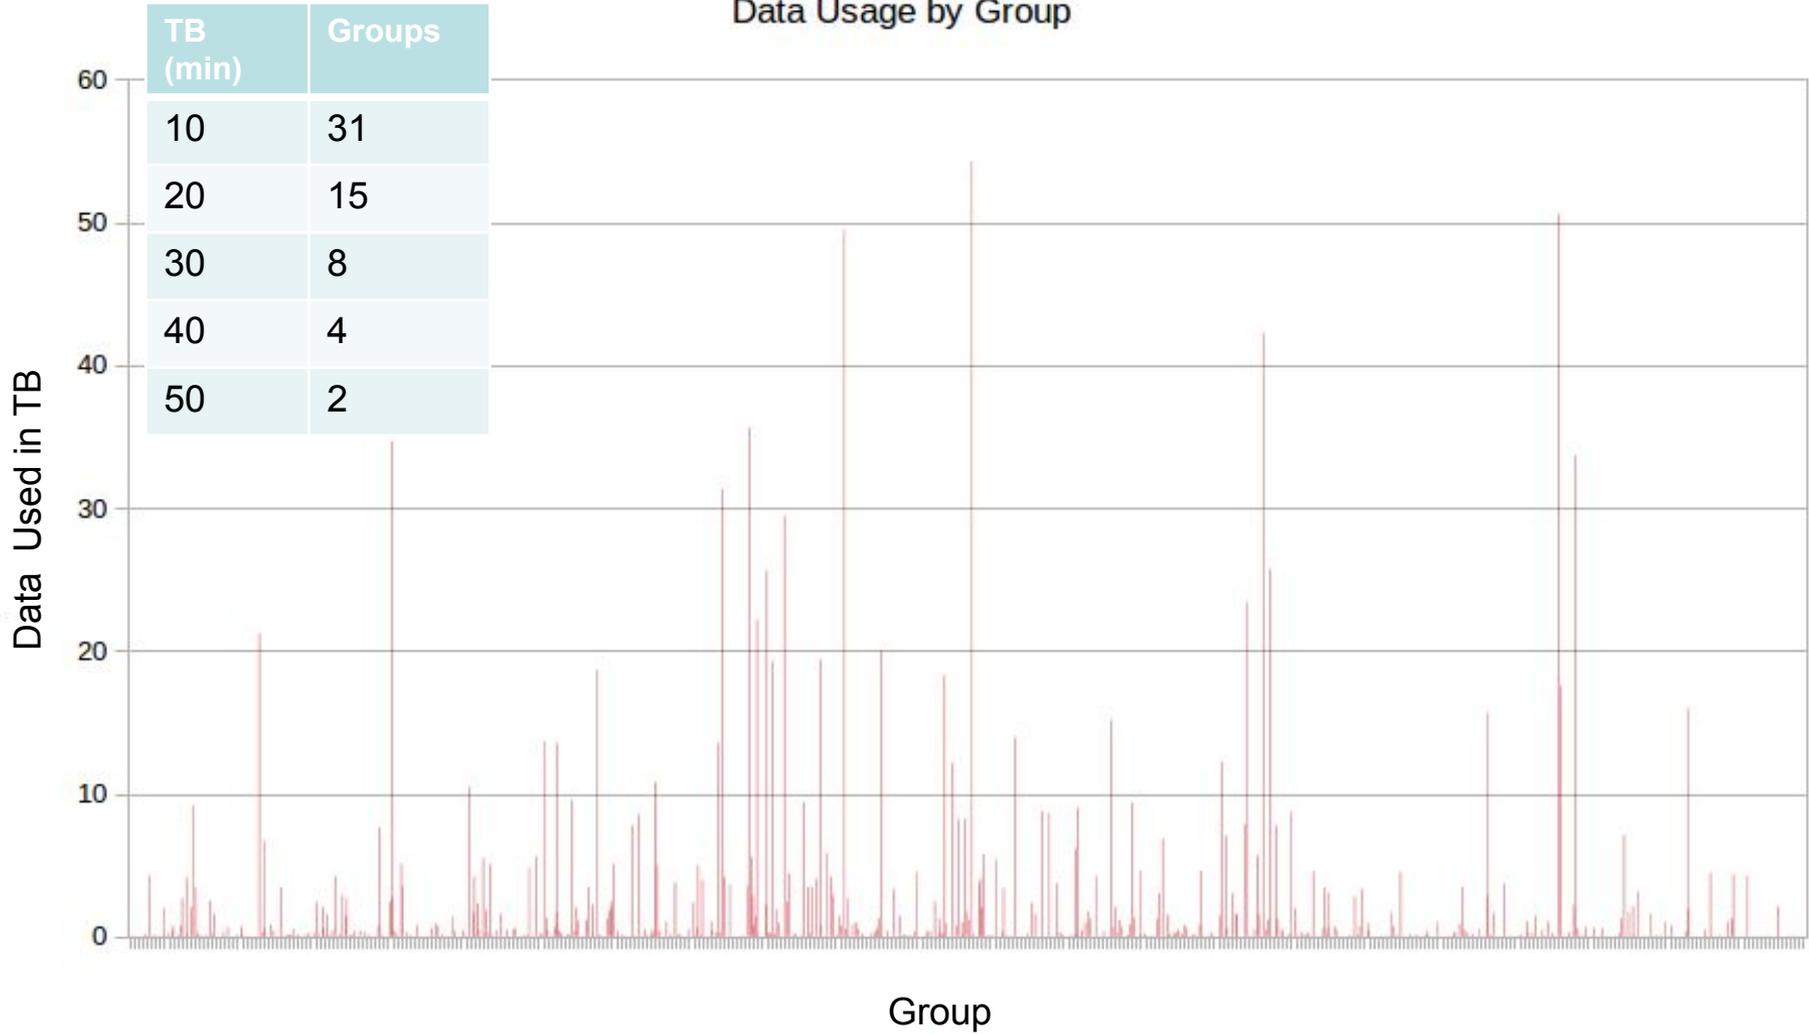
- Huge amount of churn: 86 TB deleted + 109 TB added/monthly.
- Users are quota-bound and alternative storage is available.



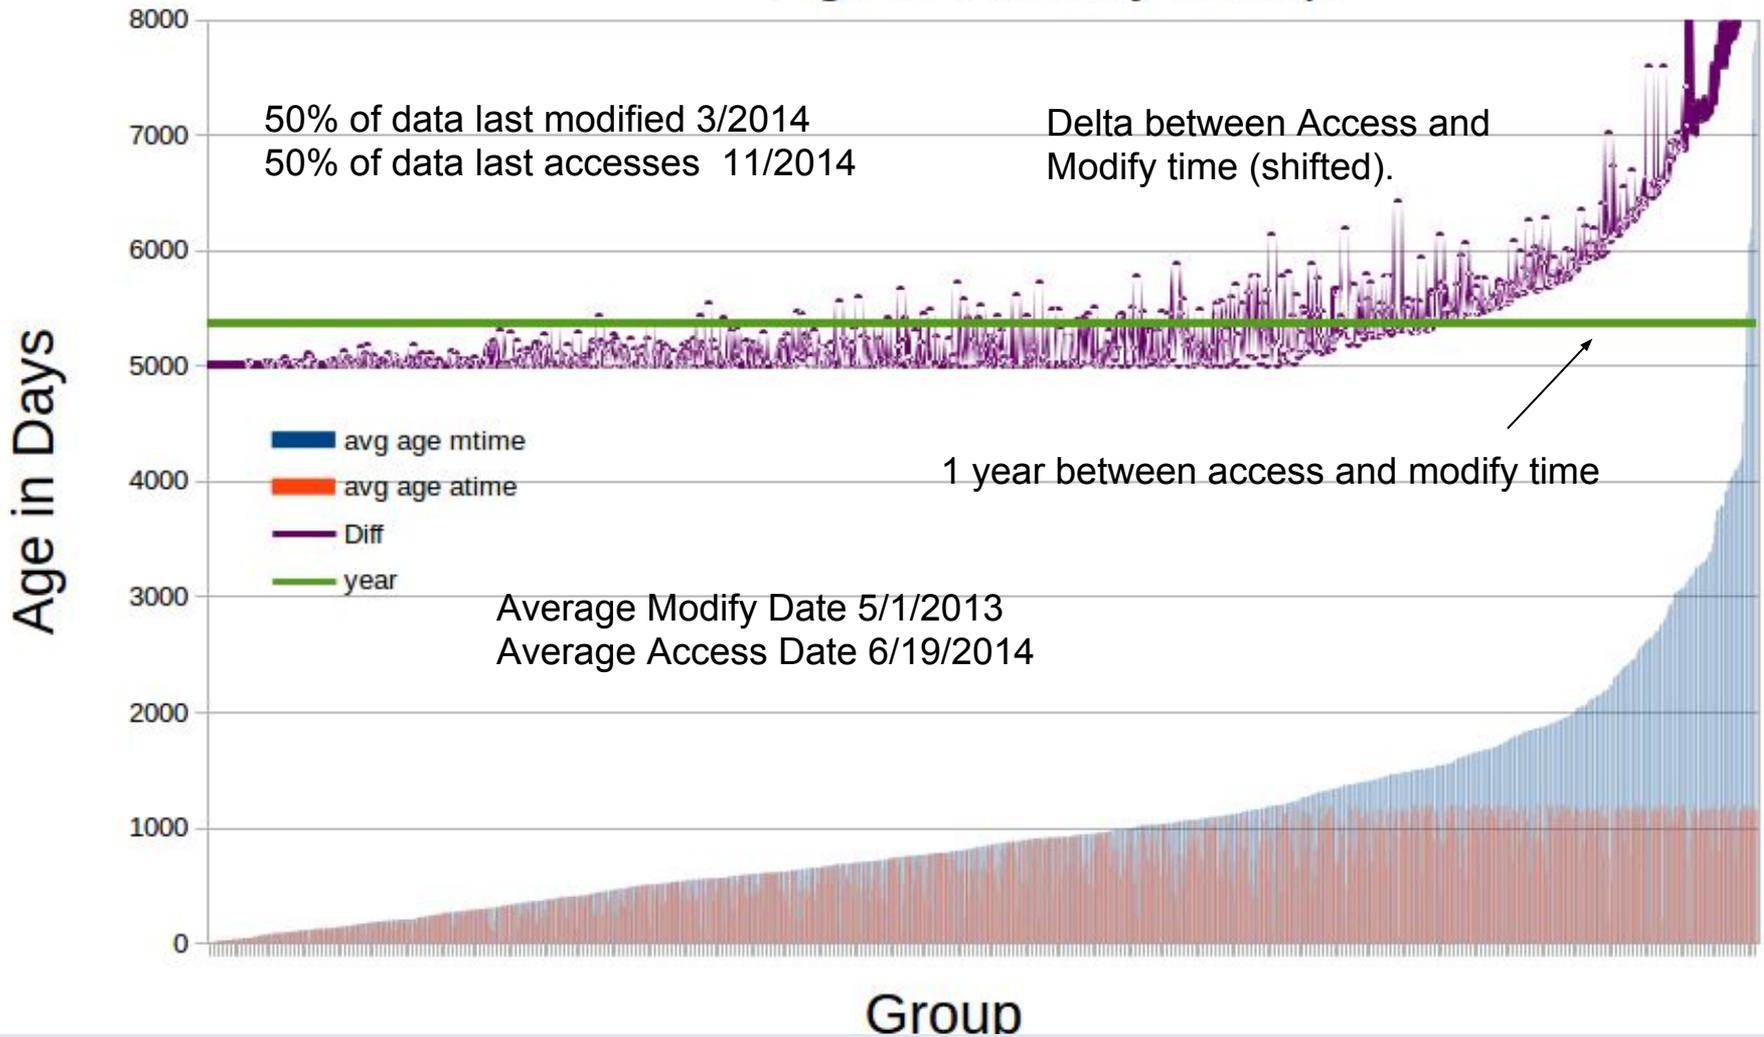
MSI 3/2016 Daily Ingress Home Directories



Data Usage by Group



Age of Files by Group



Toward Sustainable Storage

- MSI is shared resource; demand for storage is high and drives MSI to look for alternatives relatively small amount of “hot” data.
- Developing tiers of storage.
- Spinning disk tier for short term data (CEPH, Tier 2)
- Tape, yes, Tape, for long term storage—at the least expensive, for a spinning disk system per TB/year, the electrical+cooling costs start to equal the capital costs.

More tips

- Work product should go into a cloud storage and/or use RCS (github,svn,etc)
- Backup user's systems with a backup product (if you don't they will and will naturally pick the least efficient model)
- Offer a tiered storage model for hot, medium and cold data.

10.31.0.209



Search

/home/tech/jmcdonal

Name	Size	Type	Date Modified	Name	Owner	Size	Kind	Created
▶ db_be...		Folder	9/19/13 12:26 PM	▶ archived_users	jmcdonal	--	Bucket	December 30
▶ debug		Folder	5/6/14 1:30 PM	▶ example	jmcdonal	--	Bucket	October 20, 2
▶ delaw...		Folder	3/27/16 1:43 PM	▼ jmcdonal1	jmcdonal	--	Bucket	October 19, 2
▶ dell		Folder	9/19/13 12:26 PM	1gbfile-4	jmcdonal	1.0 GB	Object	December 30
▶ Desktop		Folder	2/2/15 8:55 AM	ANSYSLICM...	jmcdonal	155.5 MB	Object	February 4, 2
▶ Docum...		Folder	5/18/16 2:42 PM	ANSYS_160...	jmcdonal	4.3 GB	Object	February 4, 2
▶ Downl...		Folder	9/19/13 12:26 PM	ANSYS_160...	jmcdonal	4.1 GB	Object	February 4, 2
▶ fluent...		Folder	4/4/16 3:07 PM	ANSYS_160...	jmcdonal	4.4 GB	Object	February 4, 2
▶ fsstats		Folder	5/24/16 1:14 PM	ANSYS_160...	jmcdonal	3.0 GB	Object	February 4, 2
▶ git-tuto...		Folder	9/11/15 1:55 PM	DellCompell...	jmcdonal	296 KB	Object	February 4, 2
▶ gpus		Folder	9/19/13 12:27 PM	DellCompell...	jmcdonal	375 KB	Object	February 4, 2
▶ hopsp...		Folder	2/18/15 2:58 PM	DellCompell...	jmcdonal	344 KB	Object	February 4, 2
▶ Install...		Folder	4/21/16 2:35 PM	Front USB-Vi...	jmcdonal	19 KB	Object	February 4, 2
▶ lib		Folder	2/17/15 3:15 PM	Panasas-Scr...	jmcdonal	343.0 MB	Object	February 4, 2
▶ mail		Folder	5/3/16 11:35 AM	Procmon.exe	jmcdonal	2.4 MB	Object	February 4, 2
▶ mesab...		Folder	2/17/15 3:58 PM	RESULTS.txt	jmcdonal	4 KB	Object	February 4, 2
▶ minne...		Folder	9/19/13 12:27 PM	csd5621.txt	jmcdonal	2 KB	Object	February 4, 2
▶ moab-i...		Folder	9/19/13 12:27 PM	dmf_logfile.txt	jmcdonal	84 Bytes	Object	February 4, 2
▶ modul...		Folder	4/13/16 3:21 PM	dsmcron.pan...	jmcdonal	11.8 MB	Object	February 4, 2
▶ mpi		Folder	9/19/13 12:27 PM	dsmcron.pan...	jmcdonal	20.1 MB	Object	February 4, 2
▶ MSlim...		Folder	12/2/15 9:26 PM	dsmcron.pan...	jmcdonal	14.9 MB	Object	February 4, 2
▶ Music		Folder	9/19/13 12:27 PM	flexlmlogs.tar	jmcdonal	61.2 MB	Object	February 4, 2
▶ NoMac...		Folder	8/3/15 9:32 AM	http_https_1...	jmcdonal	12.2 MB	Object	February 5, 2
▶ nvidia		Folder	4/27/15 9:58 AM	jeff-script.sh	jmcdonal	298 Bytes	Object	February 5, 2
▶ OpenCL		Folder	9/19/13 12:27 PM	jeff-script2.sh	jmcdonal	304 Bytes	Object	February 5, 2
▶ openmpi		Folder	4/8/15 3:41 PM	labqueueava...	jmcdonal	51.2 MB	Object	February 5, 2
▶ Pictures		Folder	9/19/13 12:27 PM	mbaynton-te...	jmcdonal	118.2 MB	Object	February 5, 2
▶ prnsl		Folder	9/19/13 12:27 PM					

Jobs Log

Log

```

STARTING DS3 Browser Session
AUTHENTICATING Session
AUTHENTICATED Session
BULK GET BUCKETS http://10.31.0.209
BULK GET OBJECTS http://10.31.0.209/jmcdonal1&delimiter=/

```

Data Transfers

- MSI has a Globus subscription.
- Dedicated Servers for endpoint (transfer-optimized)
- Transfers occur out of band/in batch system.
- Endpoints for Tier 2 and 3 storage systems available.
- Transfer Rates exceeding 2.5 Gbps between MSI and NCSA, others.

Transfer Files

RECENT ACTIVITY ○ 0 ▾ 0 ○ 0

Endpoint ☆

Path

Endpoint ☆

Path

select all up one folder refresh list

- Folder Ansoft
- Folder Desktop
- Folder Documents
- Folder Downloads
- Folder InstallShield
- Folder MSIimages
- Folder Music
- Folder NoMachine
- Folder OpenCL
- Folder Pictures
- Folder Public
- Folder SVI
- Folder Templates
- Folder Videos
- Folder VirtualBox VMs
- Folder analytics
- Folder archived_users
- Folder blackpearl
- Folder db_bench
- Folder debug

select all up one folder refresh list

- Folder archived_users
- Folder example
- Folder jmcdonal1
- Folder Ito7newbucket
- Folder put_sample2
- Folder test3

Label This Transfer

This will be displayed in your transfer activity.

Transfer Settings

- sync - only transfer new or changed files ?
- delete files on destination that do not exist on source ?
- preserve source file modification times ?
- verify file integrity after transfer ?
- encrypt transfer ?

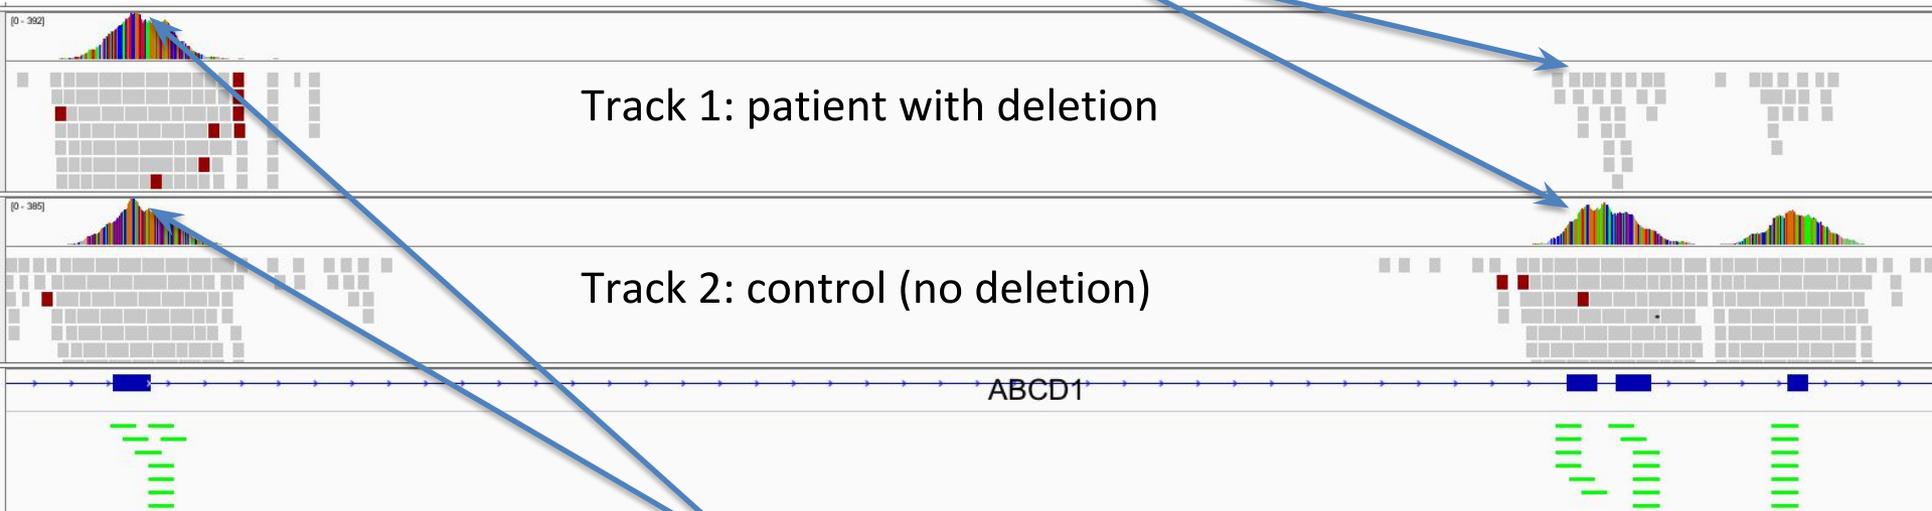
[Get Globus Connect Personal](#)
 Turn your computer into an endpoint.

Hadoop RIS Projects

- MSI RIS group (Research Informatics Solutions) provides solutions in the biological/medical space.
- RIS used hadoop on two projects
 - Speeding up a program for detecting deletions and duplications on genes (CNV program)
 - Programmatically searching pubmed

Copy Number Variation Personalized Medicine

Exon missing in patient



Track 1: patient with deletion

Track 2: control (no deletion)

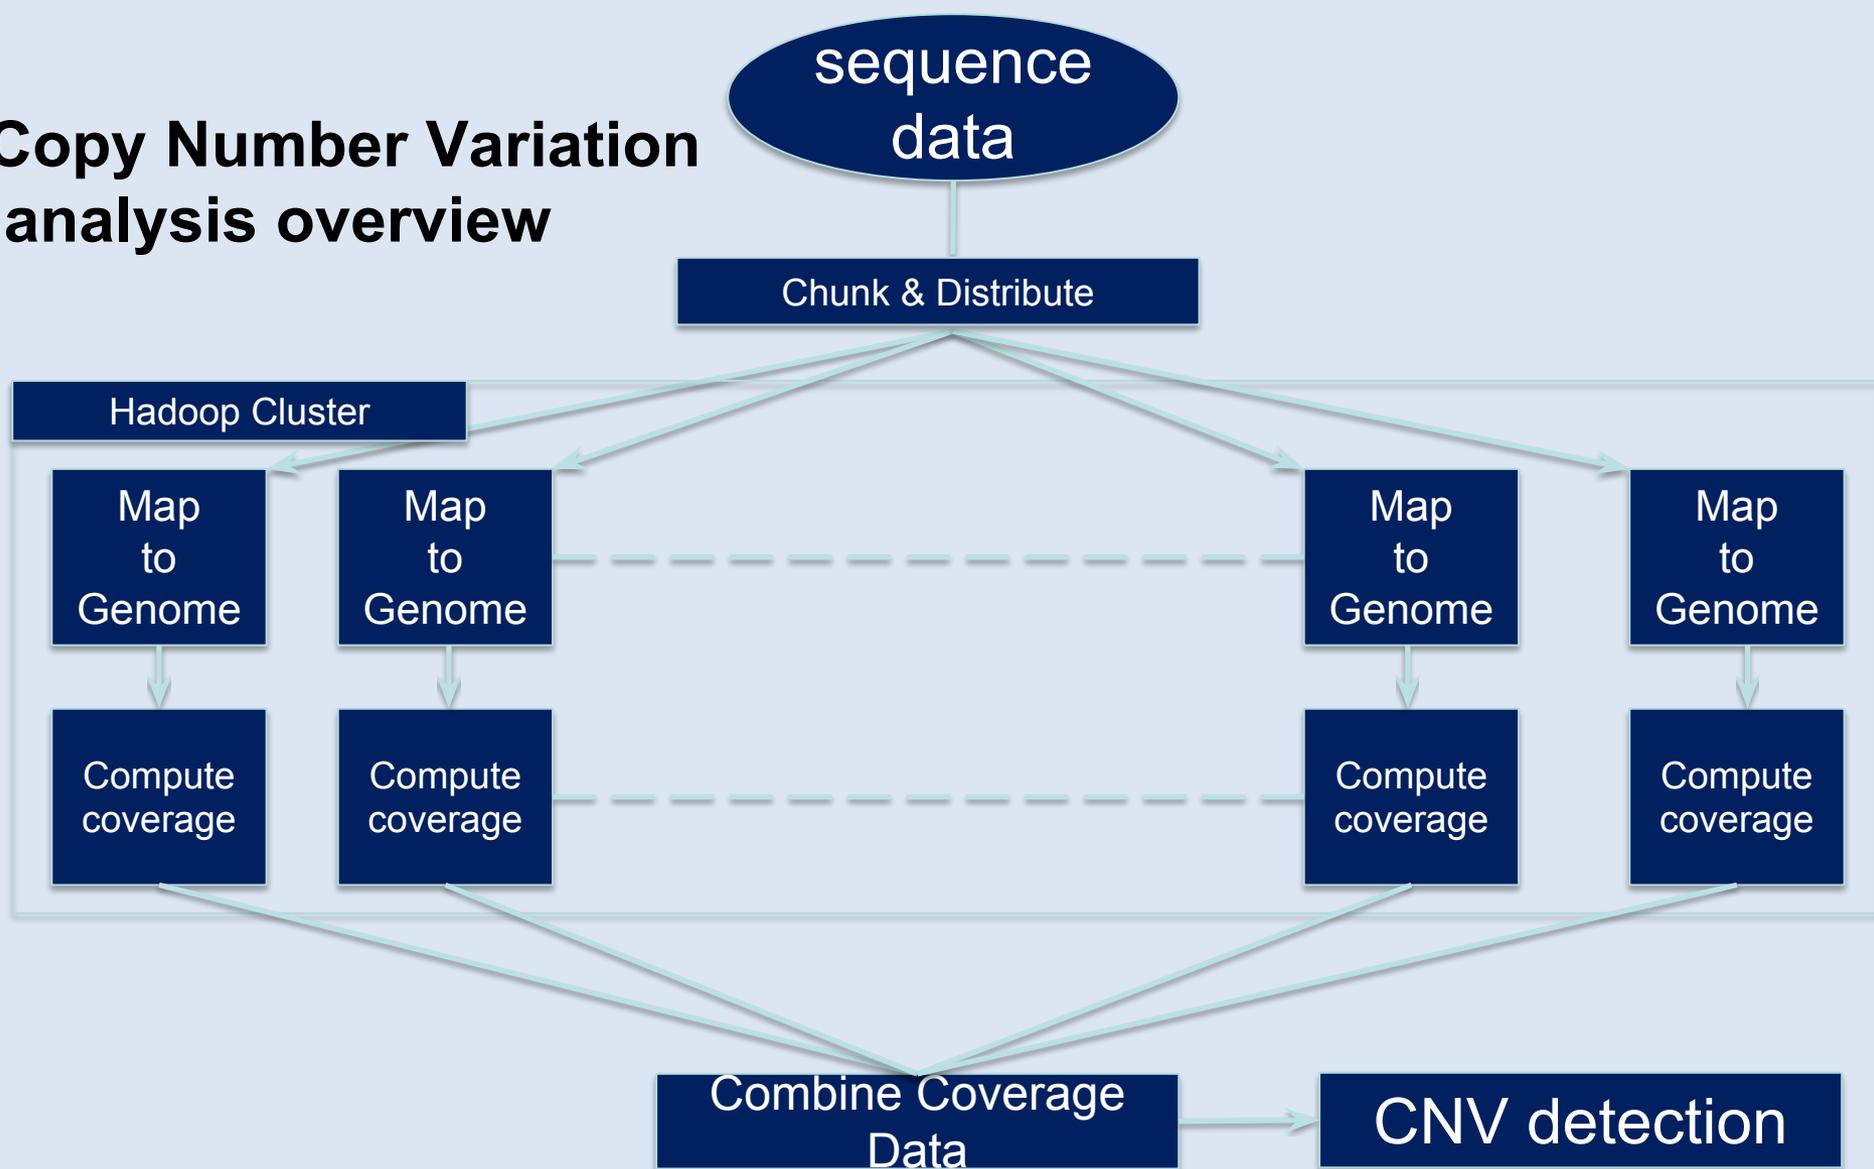
exon present in both patient and control

Exon = part of a gene

CNV = Copy Number Variation = different in number of gene copies (typically we have two, one from each parent)

The above slide shows an instance where a patient is missing part of a gene. Top right, missing coverage graph indicates loss of that portion of the gene. Stop by if you have any questions about this figure.

Copy Number Variation analysis overview



Searching Pubmed

- Why hadoop for searching pubmed
 - NCBI API offer limited type of queries
 - Only useful with a small dataset
 - Need ability to executes 1000s of queries for a single dataset
 - We used PIG (on top of hadoop)
 - Next two slides illustrate use of PIG to search pubmed and validation we are getting correct results

Using hadoop to search pubmed

```
onson001@labh01 [~/hadoop/pig] %  
onson001@labh01 [~/hadoop/pig] % pig-0.10.0/bin/pig -x local  
Picked up _JAVA_OPTIONS: -Xmx22528000k  
Picked up _JAVA_OPTIONS: -Xmx22528000k  
2014-11-04 16:51:33,209 [main] INFO org.apache.pig.Main - Apache Pig version 0.10.0 (r1328203) compiled Apr 19 2012, 22:54:12  
2014-11-04 16:51:33,210 [main] INFO org.apache.pig.Main - Logging error messages to: /panfs/roc/groups/14/msistaff/onson001/hadoop/pig/pig_1415141493203.log  
2014-11-04 16:51:33,405 [main] INFO org.apache.pig.backend.hadoop.executionengine.HExecutionEngine - Connecting to hadoop file system at: file:///br>  
grunt>  
grunt> REGISTER /home/msistaff/onson001/hadoop/pig/trunk/contrib/piggybank/java/piggybank.jar;  
grunt>  
grunt> pigdata = LOAD '/home/msistaff/onson001/hadoop/pig/data/medline/' USING org.apache.pig.piggybank.storage.XMLLoader('article') as (article:chararray);  
grunt>  
grunt> medline_subset = foreach pigdata GENERATE FLATTEN(REGEX_EXTRACT_ALL(article, '.*pmid">(\d+)<.*<article-title>(.*?)</article-title>.*<abstract>\s(.*?)</abstract>.*')) AS (pmid:int, title:chararray, abstract:chararray);  
grunt>  
grunt> filter_one = FILTER medline_subset BY ((LOWER(title) MATCHES '.*proteomic.*' OR LOWER(title) MATCHES '.*microscopy.*') AND (LOWER(abstract) MATCHES '.*neuropeptides.*' OR LOWER(abstract) MATCHES '.*nanopharmacology.*'));  
grunt>  
grunt> pmid = FOREACH filter_one GENERATE pmid;  
grunt>  
grunt> dump pmid;  
2014-11-04 16:52:42,818 [main] INFO org.apache.pig.tools.pigstats.ScriptState - Pig features used in the script: FILTER  
2014-11-04 16:52:43,049 [main] INFO org.apache.pig.backend.hadoop.executionengine.mapReduceLayer.MRCompiler - File concatenation is disabled  
2014-11-04 16:52:43,087 [main] INFO org.apache.pig.backend.hadoop.executionengine.mapReduceLayer.MultiQueryOptimizer - MR plan size after optimization: 1  
2014-11-04 16:52:43,087 [main] INFO org.apache.pig.backend.hadoop.executionengine.mapReduceLayer.MultiQueryOptimizer - MR plan size after optimization: 1  
2014-11-04 16:52:43,122 [main] INFO org.apache.pig.tools.pigstats.ScriptState - Pig script settings are added to the job  
2014-11-04 16:52:43,137 [main] INFO org.apache.pig.backend.hadoop.executionengine.mapReduceLayer.JobControlCompiler - mapred.job.reduce.markreset.buffer.percent is not set, set to default 0.3  
2014-11-04 16:52:43,163 [main] INFO org.apache.pig.backend.hadoop.executionengine.mapReduceLayer.JobControlCompiler - Setting up single store job  
2014-11-04 16:52:43,225 [main] INFO org.apache.pig.backend.hadoop.executionengine.mapReduceLayer.MapReduceLauncher - 1 map-reduce job(s) waiting for submission.  
OpenJDK 64-Bit Server VM warning: You have loaded library /nfs/soft-el6/hadoop/0.20.205.0/lib/libhadoop.so which might have disabled stack guard. The VM will try to fix the stack guard now.  
It's highly recommended that you fix the library with 'execstack -c <libfile>', or link it with '-z noexecstack'.  
2014-11-04 16:52:43,258 [Thread-4] WARN org.apache.hadoop.util.NativeCodeLoader - Unable to load native-hadoop library for your platform... using builtin-java classes where applicable  
2014-11-04 16:52:43,268 [Thread-4] WARN org.apache.hadoop.mapred.JobClient - No job jar file set. User classes may not be found. See JobConf(Class) or JobConf#setJar(String).  
2014-11-04 16:52:43,365 [Thread-4] INFO org.apache.hadoop.mapreduce.lib.input.FileInputFormat - Total input paths to process : 8  
2014-11-04 16:52:43,365 [Thread-4] INFO org.apache.pig.backend.hadoop.executionengine.util.MapRedUtil - Total input paths to process : 8  
2014-11-04 16:52:43,378 [Thread-4] INFO org.apache.pig.backend.hadoop.executionengine.util.MapRedUtil - Total input paths (combined) to process : 1  
2014-11-04 16:52:43,729 [main] INFO org.apache.pig.backend.hadoop.executionengine.mapReduceLayer.MapReduceLauncher - HadoopJobId: job_local_0001  
2014-11-04 16:52:43,730 [main] INFO org.apache.pig.backend.hadoop.executionengine.mapReduceLayer.MapReduceLauncher - 0% complete  
2014-11-04 16:52:43,745 [Thread-5] INFO org.apache.hadoop.util.ProcessTree - setsid exited with exit code 0  
2014-11-04 16:52:43,754 [Thread-5] INFO org.apache.hadoop.mapred.Task - Using ResourceCalculatorPlugin : org.apache.hadoop.util.LinuxResourceCalculatorPlugin@2d94c3f0  
2014-11-04 16:52:43,776 [Thread-5] INFO org.apache.pig.backend.hadoop.executionengine.mapReduceLayer.PigRecordReader - Current split being processed file:/home/msistaff/onson001/hadoop/pig/data/medline/AAPS_J_2010_Oct_19_12(4)_716-728.nlm  
:0+119422  
2014-11-04 16:52:49,711 [communication thread] INFO org.apache.hadoop.mapred.LocalJobRunner -  
2014-11-04 16:52:50,881 [Thread-5] INFO org.apache.pig.backend.hadoop.executionengine.mapReduceLayer.PigRecordReader - Current split being processed file:/home/msistaff/onson001/hadoop/pig/data/medline/AAPS_J_2010_Aug_24_12(4)_635-645.nlm  
:0+119422
```

Total of 5
commands

```
2014-11-04 16:53:18,756 [main] INFO org.apache.pig.backend.hadoop.executionengine.mapReduceLayer.MapReduceLauncher - Success!  
2014-11-04 16:53:18,762 [main] INFO org.apache.hadoop.mapreduce.lib.input.FileInputFormat - Total input paths to process : 1  
2014-11-04 16:53:18,762 [main] INFO org.apache.pig.backend.hadoop.executionengine.util.MapRedUtil - Total input paths to process : 1  
(20957528)  
(20734175)  
grunt>  
grunt>
```

two articles returned: 20957528 ,
20734175

UNIVERSITY OF MINNESOTA

© 2016 Regents of the University of Minnesota. All rights reserved.

Find any article with *proteomic* or *microscopy* in the title and *neuropeptides* or *nanopharmacology* in the abstract

Minnesota Supercomputing Institute

Same search on pubmed site : to validate we get the same results

2014-11-04 16:53:18,756 [main] INFO
2014-11-04 16:53:18,762 [main] INFO
2014-11-04 16:53:18,762 [main] INFO
(20957528)
(20734175)
grunt>
grunt>

[Display Settings:](#) Abstract

[Send to:](#)

AAPS J. 2010 Dec;12(4):716-28. doi: 10.1208/s12248-010-9232-y. Epub 2010 Oct 19.

Multidimensional atomic force microscopy: a versatile novel technology for nanopharmacology research.

Lal R¹, Ramachandran S, Arnsdorf MF.

[+ Author information](#)

Abstract

Nanotechnology is giving us a glimpse into a nascent field of nanopharmacology that deals with pharmacological phenomena at molecular scale. This review presents our perspective on the use of scanning probe microscopy techniques with special emphasis to multidimensional atomic force microscopy (m-AFM) to explore this new field with a particular emphasis to define targets, design therapeutics, and track outcomes of molecular-scale pharmacological interactions. The approach will be to first discuss operating principles of m-AFM and provide representative examples of studies to understand human health and disease at the molecular level and then to address different strategies in defining target macromolecules, screening potential drug candidates, developing and characterizing of drug delivery systems, and monitoring target-drug interactions. Finally, we will discuss some future directions including AFM tip-based parallel sensors integrated with other high-throughput technologies which could be a powerful platform for drug discovery.

PMID: 20957528 [PubMed - indexed for MEDLINE] PMID: PMC2976997 [Free PMC Article](#)

Conclusions

- HPC is becoming more data-driven
- Important to understand data lifecycles and plan for them
- Input-Output can continue to be optimized (processors are limited here)
- Big Data means a place for storage of data, processing of data and mechanisms to share that data.

MSI is an External Service Organization

Storage resources,
Consulting, and System time
available for external entities

Current and Former Customers

- US Golf Association
- PepsiCo
- Cargill
- Syngenta
- Third Wave Systems

<https://www.msi.umn.edu/content/service-catalog>

or

email: help@msi.umn.edu

Questions

UNIVERSITY OF MINNESOTA

© 2015 Regents of the University of Minnesota. All rights reserved.

Minnesota Supercomputing Institute
MSI

Organization

Office of the Vice President for Research

User Gateway Group

- User Support Lead
- User Training
- On Boarding
- Communications
- Outreach

Scientific Computing Solutions

- Optimization
- Benchmarking
- HPC Research
- Workflow & pipeline Development

Research Informatics Solutions

- Informatics education
- Informatics research
- Informatics services
- Life Science Computing

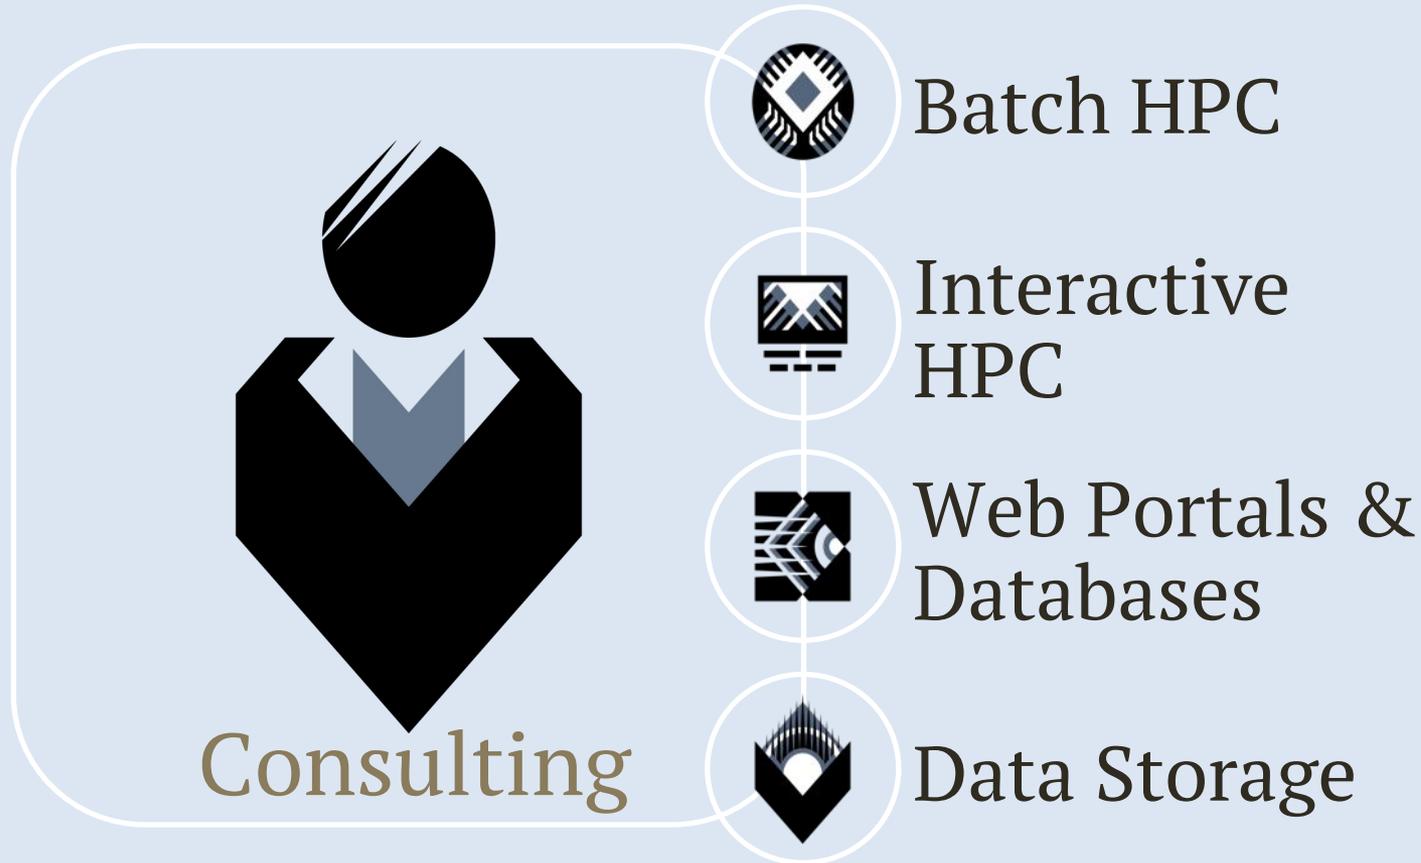
Application Development Solutions

- Custom App Dev
- System Programming

Advanced Systems Operations

- Common Services
- HPC Systems
- Storage Systems
- Hosted Services

Services



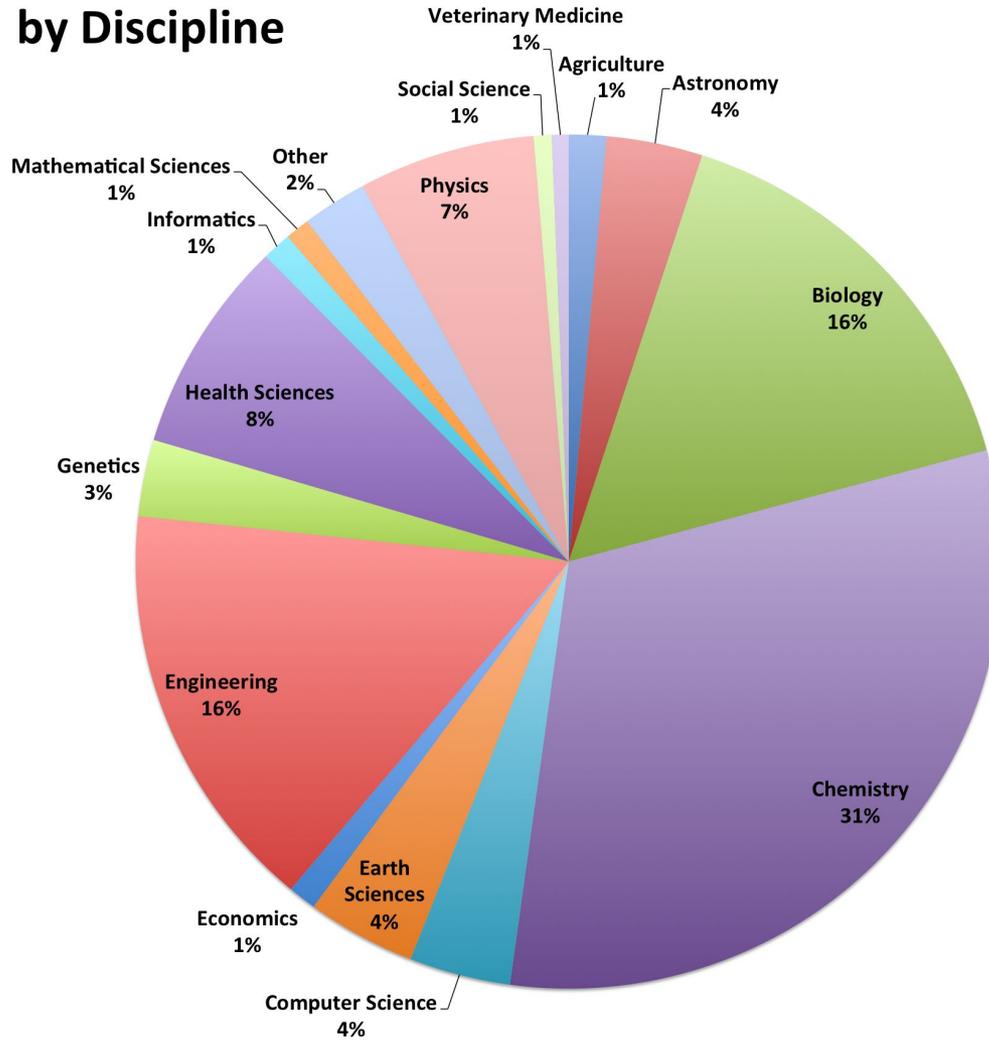
UNIVERSITY OF MINNESOTA

© 2015 Regents of the University of Minnesota. All rights reserved.

Minnesota Supercomputing Institute
MSI

Who Uses MSI?

Services Units Allocated by Discipline

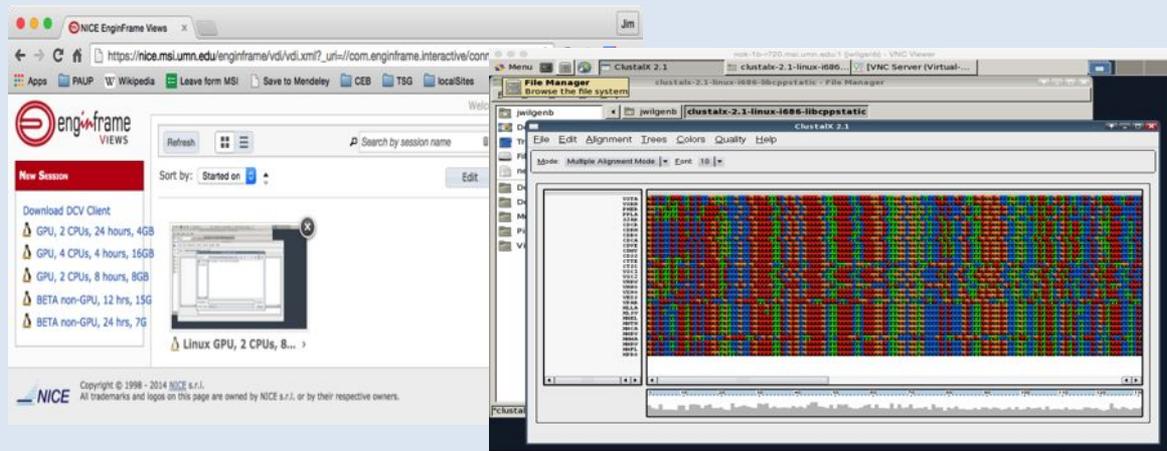




Interactive HPC

Environments

- iSub:
- NICE:
- Citrix for Windows



- **Itasca**

- Cores: 8,744
- Memory: 31.3 TB
- Peak: 100 TF

- **Mesabi**

- Cores: 16,848
- Memory: 67 TB
- Accelerators: 80 K40 gpGPUs
- Peak: 675 TF

- **Cascade**

- 32 Tesla gpGPUs
- 8 Kepler gpGPUs
- 2 Intel Phi
- 192 CPU Cores



Portals & Databases

Multiple “omics” platform
Emphasis on Mass Spectrometry based Proteomics

The screenshot shows the GalaxyP web portal at <https://usegalaxy.org>. The interface includes a navigation bar with 'Galaxy / GalaxyP' and menu items like 'Analyze Data', 'Workflow', 'Shared Data', 'Visualization', 'Help', and 'User'. A left sidebar lists various tools under categories like 'CORE TOOLS', 'PROTEOMICS', 'GENOMICS', and 'EMBOSS'. The main content area features a 'Welcome to GalaxyP' message, a 'GalaxyP' logo, an 'Updates' section with recent activity, a 'Tweets' section with a tweet from @usegalaxy, and a 'Links' section with a link to a paper on multi-omic data integration. A right sidebar shows 'History' and 'Unnamed history' sections.



Portals & Databases

Minnesota Biodiversity Atlas

Working with the Bell Museum and Zooniverse to organize and annotate images of specimens.

The screenshot shows the search interface of the Minnesota Biodiversity Atlas. The header includes the Bell Museum logo and the text "Minnesota Biodiversity Atlas". Below the header is a navigation bar with links for Home, Search Collections, Map Search, Image Search, Browse Images, Inventories, Interactive Tools, and Log In / New Account / Sitemap. The main content area is titled "Enter Search Parameters" and contains several sections for filtering search results:

- Taxonomic Criteria:** Includes a checkbox for "Include Synonyms from Taxonomic Thesaurus" and a text input field for "Family or Scientific Name".
- Locality Criteria:** Includes input fields for Country, State/Province, County, Locality, and Elevation (with a "to" field).
- Latitude and Longitude:** Includes two search methods: "Bounding box coordinates in decimal degrees" (with fields for Northern, Southern, Western, and Eastern Longitude) and "Point-Radius Search" (with fields for Latitude, Longitude, and Radius in Kilometers).
- Collector Criteria:** Includes input fields for Collector's Last Name, Collector's Number, and Collection Date.
- Specimen Criteria:** Includes input fields for Catalog Number and Other Catalog Number, and checkboxes for "Limit to Type Specimens Only" and "Limit to Specimens with Images Only".

Search buttons are located to the right of the Taxonomic and Specimen Criteria sections.

UNIVERSITY OF MINNESOTA

© 2015 Regents of the University of Minnesota. All rights reserved.

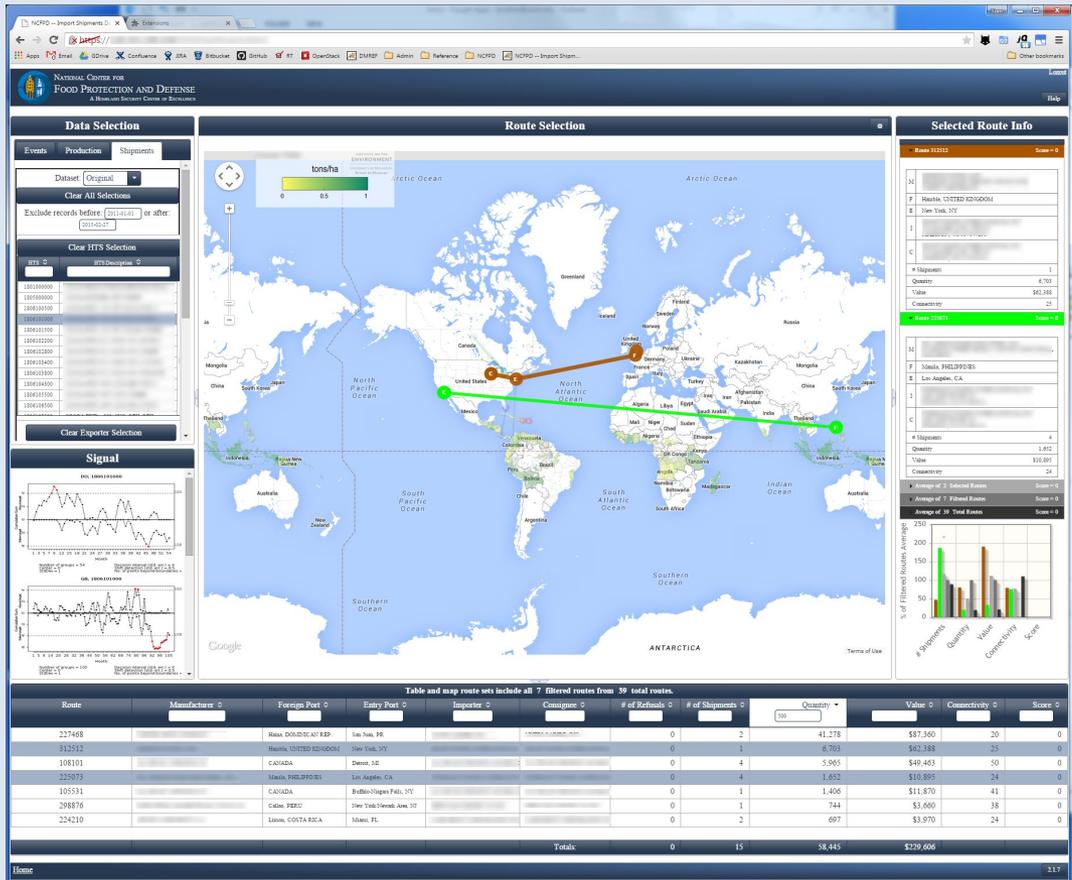
Minnesota Supercomputing Institute
MSI



Portals & Databases

Food Protection and Defense Institute

- FIDES
 - ID potential disruptions to food supply
 - ID food system disruptions using open news media
- FIDES extension to water supply data



UNIVERSITY OF MINNESOTA

© 2015 Regents of the University of Minnesota. All rights reserved.

Minnesota Supercomputing Institute



Collaborations

International AgroInformatics Alliance



Reimagine the role of informatics to improve agricultural, food and nutritional outcomes worldwide

- **Clusters:**

- **Corn Cluster:** Creating and Testing Spatially Integrated Database Design Options and Prototypes
- **Wheat Cluster:** Scoping Alliance Partnerships
- **Rosaceae and Solanaceae Cluster:** Characterization of the Molecular Resistance Repertoire
- **Soybean Cluster:** Rescue, Redeployment and Re-envisioning Breeding Related Data
- **Engineering Access:** Application & DBs in Real World Environments
- **IAA-Innovation Partnerships:** IPR, Data Privacy and Data Sharing Practices
- **Potato Cluster:** Rescue, Redeployment and Re-envisioning Breeding Related
- **Oats Cluster:** Rescue, Redeployment and Re-envisioning Breeding Related Data



Consulting



- 24 -- Consulting Staff
- 17 -- Staff with advanced degrees
- 10 to 20 -- Workshops and Tutorials annually

UNIVERSITY OF MINNESOTA

© 2015 Regents of the University of Minnesota. All rights reserved.

Minnesota Supercomputing Institute
MSI

Research Outcomes

- 1,831 MSI PI publications
- 935 Unique journal titles
- 165 MSI PI pubs in high impact journals
- \$427 Million C&G funding by MSI PIs
- \$66 Million by top 25 MSI users

Questions

UNIVERSITY OF MINNESOTA

© 2015 Regents of the University of Minnesota. All rights reserved.

Minnesota Supercomputing Institute
MSI