

# Multi-Tier Storage in Campus Environments

Brock Palen 2018 [brockp@umich.edu](mailto:brockp@umich.edu)



ARCTS

ADVANCED RESEARCH COMPUTING  
TECHNOLOGY SERVICES  
UNIVERSITY OF MICHIGAN

# Historic Campus Storage Offerings

- One size fits all
- Assumes the value of all data are the same
  - Hot vs. Warm vs. Archive/Publication

Otherwise build your own

# Non-Volatile Media Types

Solid State/NVMe ~1TB	Hard Drive ~8-12TB	Tape 12-15TB
\$1/GB	\$0.10/GB	\$0.01/GB
Seek Performance High	Medium	Very Low



# Justification - Flux 1.3 PB

File Size Range	Average Size MB	% of Total Files	% of Total Space
>1GB	5561.053	0.02%	58.14%
>512MB - 1GB	698.909	0.02%	8.73%
>256MB - 512MB	380.379	0.03%	5.96%
>128MB - 256MB	176.324	0.05%	4.90%
>32MB - 128MB	56.892	0.24%	8.12%
>1MB - 32MB	5.435	3.33%	10.98%
0-1MB	0.054	96.32%	3.16%

## FS Space Consumed vs File Size



# Proposal: Multi-Tier Offerings

## General Purpose

- Flash + HDD
- Access Type
  - Desktop & Cluster
  - NFS and CIFS
  - Globus
- 4-10GB/s
- Cost: \$115/TB/yr/Copy

## Large File

- HDD
- Access Type
  - Desktop & Cluster
  - NFS **or** CIFS
  - Globus
- 1-4GB/s
- Cost: \$40/TB/yr/Copy

## Archive

- HDD + Tape
- Access Type
  - Globus
- Cost: Not yet Available
- 1-4GB/s

Aligns with public Cloud: Eg. AWS S3, IIA, Glacier

# Turbo NAS Filer

- University Wide Access
- Simultaneous NFS/CIFS
- Hybrid NVMe HDD
- Thousands TB Capacity
- Sensitive Data Support
- Faster as system grows
  - Available Today
  - \$115.20/TB/copy/yr




ARC·TS

ADVANCED RESEARCH COMPUTING  
TECHNOLOGY SERVICES  
UNIVERSITY OF MICHIGAN

# Locker NAS Filer

- University Wide Access
- NFS or CIFS\* Access
- HDD System
- Thousands TB Capacity
- Sensitive Data Support\*
- Optimized for 1MB or Larger
- Low Cost
  - Available Today
  - \$40.09/TB/copy/yr\*



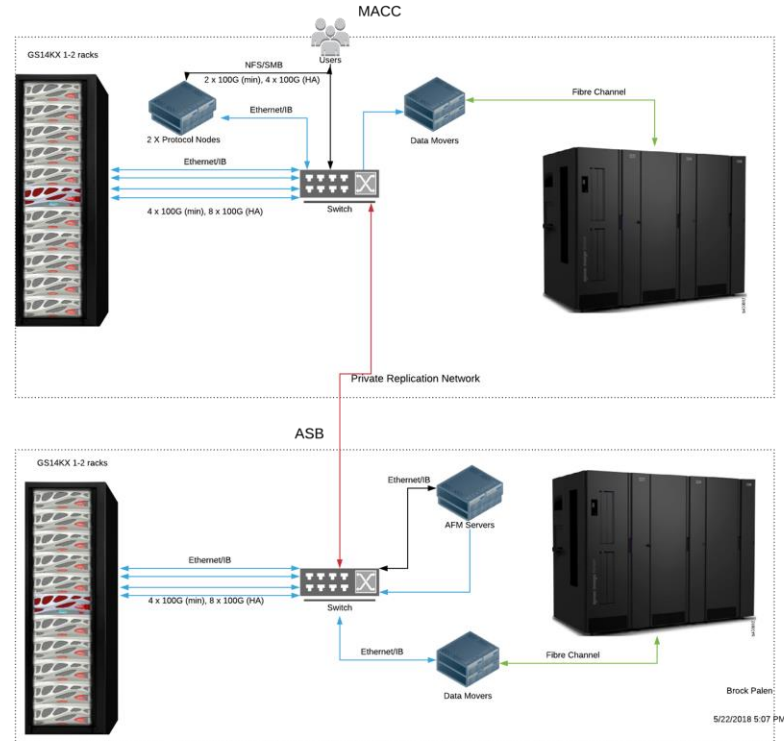
ARC·TS

ADVANCED RESEARCH COMPUTING  
TECHNOLOGY SERVICES  
UNIVERSITY OF MICHIGAN



# Data Den Archive

- World Wide Access and Sharing
- Globus / SCP
- Tape Based System
- 10's PB Capacity
- Sensitive Data Support\*
- Optimized for 100MB or Larger
- Two Copies
- Very Low One-Time Cost
  - Available Fall 2018
  - \$340-\$800 replicated / 12TB / 5years\*



# Globus - Transfer and Sharing

- World Wide Access and Sharing
- Laptop Friendly
- Provided to Campus no Cost
- Support Lab Devices
- Available on ARC-TS Systems
- Fire and Forget
- PHI/HIPAA Support Coming
- <http://arc-ts.umich.edu/systems-services/globus/>
- <http://globus.org/>

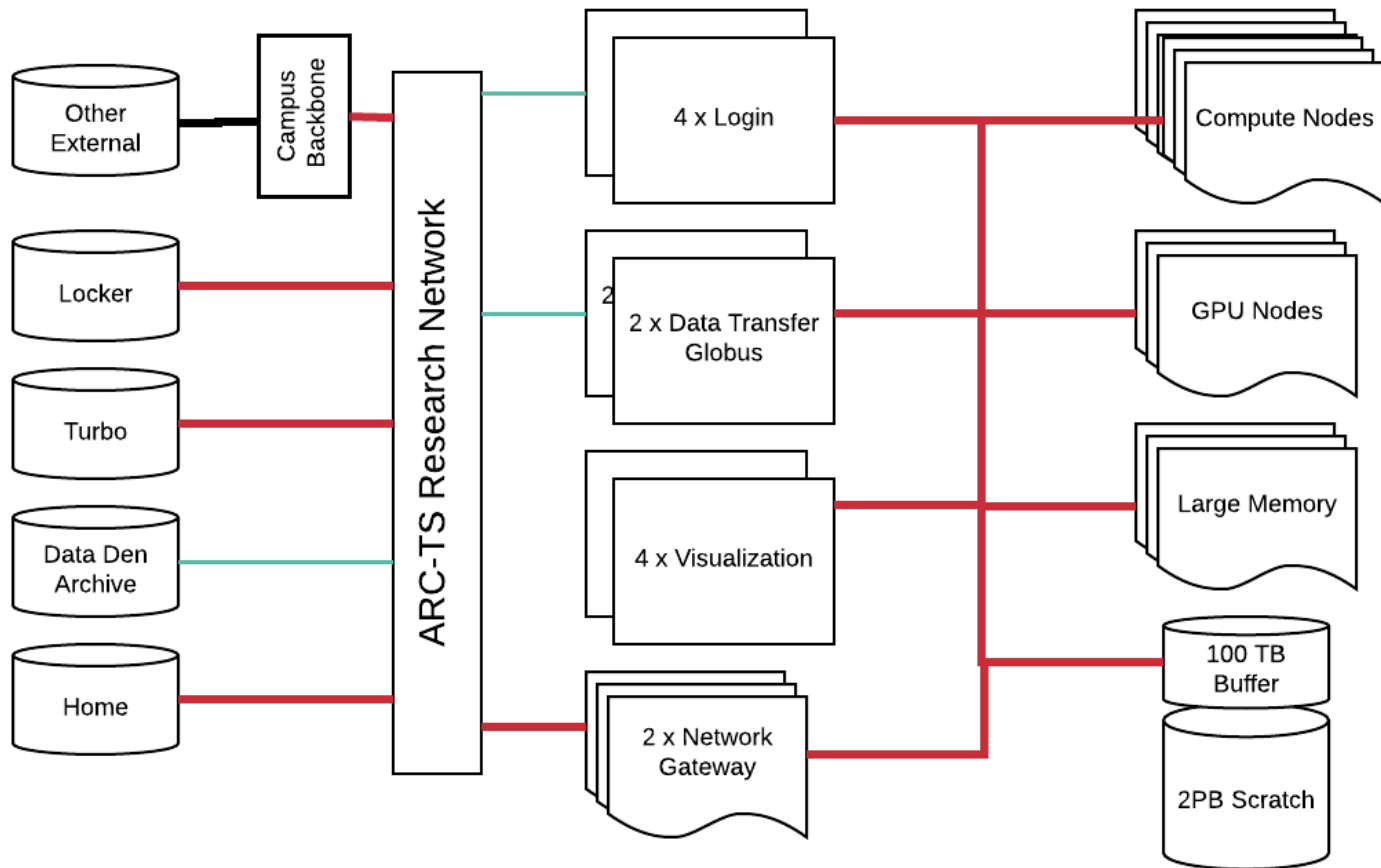


# Available Now

- Turbo recent upgrade for better performance and encryption
- Locker 2 month test allocations
- Data Den pilot feedback period (ie free for now)

Contact: [hpc-support@umich.edu](mailto:hpc-support@umich.edu)

— 100 Gig IB/IP  
— 25 Gig IP



ARC-TS

ADVANCED RESEARCH COMPUTING  
TECHNOLOGY SERVICES  
UNIVERSITY OF MICHIGAN

# Data Den Archive Costing

Goal: USB drive competitive cost durable option for data storage in the post award period with self service access for all data types (PHI, ITAR, etc)

Currently 5PB replicated, 20PB capacity with no additional university investment

Example Costing 5 years:

\$850k 5 PByte Replicated: \$34/TB/yr

\$993k 10 PByte Replicated: \$19/TB/yr

\$1,279k 20 PByte Replicated: \$12/TB/yr



ARC·TS

ADVANCED RESEARCH COMPUTING  
TECHNOLOGY SERVICES  
UNIVERSITY OF MICHIGAN