

# *Research Storage 101*

Introduction to available RCAC Storage Systems

Geoffrey Lentner

*Lead Research Data Scientist*

*Rosen Center for Advanced Computing*



Rosen Center for  
Advanced Computing

Pre-Survey



# *Training Pre-Survey*

Please fill out beforehand to let us know where you are starting from

Scan this! →



# Overview of Covered Topics


## Available Systems

- What is Research Storage?
- What systems are available?
- Where are they available?



## Storage Use-Cases

- Why should I use a particular system over another?
- System details
- Strategies



## Transfers and Collaboration

- How to migrate data between systems?
- How to share data with collaborators?



# *More than just computing*

Cyberinfrastructure provided by RCAC

- RCAC is best known for the computing clusters, but where would you put those nice simulation results without storage?
- See [rcac.purdue.edu/storage](http://rcac.purdue.edu/storage) for details
- Interactive helper: [rcac.purdue.edu/storage/solutions](http://rcac.purdue.edu/storage/solutions)
- Other services: [rcac.purdue.edu/services](http://rcac.purdue.edu/services)

# ***TL;DR***

## Abbreviated Summary of Information

- Dedicated storage systems are specialized and balance **capacity, performance, availability, safety, and security**
- RCAC provides *multiple systems* that meet combinations of these features for both compute and non-compute users
- Consider your entire data management (life-cycle) and incorporate these different systems into your workflow
- Consider automation to ensure against data loss

# 1

## Overview

What do we mean by storage system and how is that different than my local machine?

# Research Storage

What do we mean by storage system?

intel.com



ibm.com




red-gate.com


- Persistence of digital information on physical media such as **tape**, spinning **disk**, or **solid-state** devices
- Even after losing power
- Available over network
- Often presented with standard file-system interface

# Features

Not all storage systems are created equal!




- How much space
- How scalable



- How fast (read/write)
- Latency




- Replication
- Recoverability



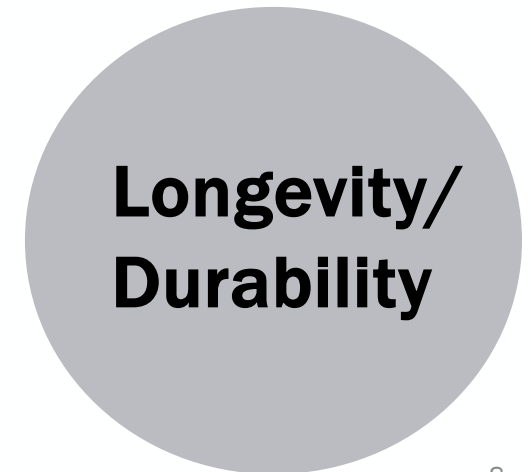
- Encryption
- Access controls



- Normal filesystem mount
- Or specialized tools
- Locality



- Price
- Maintainability






- Decay lifetime
- Fragility



# Use Cases

Not all applications have the same needs!

- *Simulations* are large-file write heavy, many writers
- *Genomics* is many parallel read/write on medium files
- *Deep Learning* is read-heavy on many tiny files
- *Optimization* doesn't require storage at all



# Use Cases


Not all users have the same requirements!

- Landing zone for scientific instrument data
- Shared space for research software development
- Long-term archival storage
- Collaborative documents

purdue.edu



ag.purdue.edu



nasa.gov

# 2

## Storage Systems

Specific details for each system.

# *Available Solutions*

Available solutions provided by RCAC (or Libraries)

## **Local (mountable)**

- /home
- /depot
- /scratch
- /tmp (cluster only)

## **Indirect Access**

- *Fortress* (HPSS tape archive)
- *Ceph* (S3 Object Storage)

## **Cloud**

- Box.com / REED
- GitHub ([github.itap.purdue.edu](https://github.itap.purdue.edu))

## **Libraries**

- PURR – research data repository providing data management, DOI, citation tracking,

# Home Directory

Personal directory on Community Clusters

`$HOME`  
`/home/<user>`

## Features

- Small (25G)
- Mild performance
- Cluster-specific (shared between nodes within cluster, but not across clusters), mountable as network drive
- Redundant hardware, never purged, protected by snapshots
- Private to user (cannot be accessed by other members of the research group)

## Good For

- Personal configuration files, software installation, scripts, etc.
- Personal data, job files (if small)
- Okay to run jobs against but not good for heavy I/O scenarios
- Medium – Long term storage



# *Depot Directory*

/depot/<group>

Group directory on Community Clusters

## **Features**

- Large (100GB free), can grow in 1TB increments, \$70 /TB /year
- Reasonable performance
- Mounted on all clusters but also as a network drive
- Redundant hardware, never purged, protected by snapshots
- Shared by group members with fine-grained access controls (belongs to PI)
- Can use without cluster purchase!

## **Good For**

- Shared configuration files, software installation, scripts, etc.
- Critical research data
- Okay to run jobs against but not good for heavy I/O scenarios
- Medium – Long-term storage

# Scratch Directory

High-performance directory on Community Clusters

`$RCAC_SCRATCH  
/scratch/<user>`

## Features

- Huge (100+ TB ... per user quota)
- High performance
- Cluster-specific (shared between nodes within cluster, but not across clusters), mountable as network drive
- Internally redundant, but **not recoverable**, regularly purged of older files
- Private to user (cannot be accessed by other members of the research group)

## Good For

- Massive / intermediate data I/O
- Perfect for data-intensive jobs
- Not for primary copy of data or software!
- ***NOT for long-term storage***

- Beware of regular purging of older files
  - Email notification, `purgelist` command
  - Don't try to game the system!
  - Just backup to *Fortress*!

# Temporary Directory

/tmp/...

Node-local /tmp directory on Community Clusters

## Features

- Moderate (200+ DB)
- Good performance
- Zero redundancy, files purged after job ends, no snapshots
- Node-local (each node has own /tmp)
- World-readable (and writable – sort of)
- What folks used before /scratch

## Good For

- Node-local caching of data and files
- **NOT** for valuable data or software
- **NOT** for long-term storage
- *Rarely needed* (but priceless when it is). Unless you understand the tradeoffs, **consider just using /scratch**



# Fortress (tape archive)

IBM HPSS system for long-term storage

/home/<user>  
/group/<group>

**NOT same as cluster /home or /group**

## Features

- Huge (25PB), free up to a point
- Tape system with robotic arm, disk cache makes *writes* fast, but *reads* slow
- Replicated, redundant, never purged
- Accessible from all cluster with specialized command-line tools, as well as network drive for external machines
- All Depot spaces get a /group Fortress directory

## Good For

- Backup (archive) of critical research data
- Large files, **NOT for many small files**
- **NOT** for running jobs against
- Long-term storage

# *Box.com (research folder)*

[ L1FR ] / <group >

Free cloud drive for research groups at Purdue

## Features

- Cloud-based, free
- Login with @purdue.edu credentials to [purdue.box.com](https://purdue.box.com)
- Individuals get up to 1TB space, lab groups get managed folder with unlimited space
- Okay performance for documents
- Redundant, versioning, never purged
- Flexible sharing and collaboration

## Good For

- Personal and group-shared documents, where collaboration is important
- Data (with 50GB per file limitation)
- **NOT** for sensitive or restricted data (such as NIST SP-800-171, FERPA, HIPAA, etc; see REED folder)
- **NOT** for HPC use (cannot compute against)
- Medium and long-term storage

# *REED (within Box.com)*

[L3]/<group>

Free cloud folder for sensitive or restricted data

## Features

- Cloud-based, free
- Login with @purdue.edu credentials to [purdue.box.com](https://purdue.box.com)
- Basically, same as Box.com but with heightened security controls
- Encrypted at rest
- See Box.com features

## Good For

- **Sensitive or restricted** data (such as NIST SP-800-171, FERPA, HIPAA, etc)
- **NOT** for HPC use (cannot compute against)
- Medium and long-term storage

| snapshots != backup!                          | \$HOME             | /tmp           | \$RCAC_SCRATCH                        | /depot/...                                   | Fortress (HPSS)   | Box folder                             | REED folder                            |
|---|--------------------|----------------|---------------------------------------|--|---|--|--|
| <b>Capacity</b>                               | 25 GB              | 150-400 GB     | varies by cluster...<br>100 TB and up | 100 GB free,<br>then paid in 1 TB increments | unlimited   | unlimited, but<br>under 50 GB per file | unlimited, but<br>under 50 GB per file |
| <b>Resilience to hardware failures</b>        | yes                | no             | yes                                   | yes  | yes   | yes                                    | yes                                    |
| <b>Resilience to human errors</b>             | yes<br>(snapshots) | no             | no                                    | yes<br>(snapshots)                           | no  | some<br>(versioning)                   | some<br>(versioning)                   |
| <b>Subject to purging</b>                     | no                 | yes            | yes                                   | no   | no  | no                                     | no                                     |
| <b>Performance</b>                            | medium             | medium to high | high                                  | medium                                       | slow to very slow   | slow to very slow                      | slow to very slow                      |
| <b>Designed for HPC (running jobs off it)</b> | no                 | no             | yes                                   | in moderation                                | <ul style="list-style-type: none"> <li>no (as main I/O)</li> <li>yes (for staging and archiving)</li> </ul> | no                                     | no                                     |
| <b>Common access within cluster</b>           | yes                | no             | yes                                   | yes  | yes<br>(hsi/htar)   | possible but kludgy                    | no                                     |
| <b>Common access across clusters</b>          | no                 | no             | no                                    | yes  | yes<br>(hsi/htar)   | possible but kludgy                    | no                                     |
| <b>Advanced sharing ACLs (past ugo/rwx)</b>   | no                 | no             | no                                    | yes  | no  | yes                                    | yes                                    |
| <b>Globus endpoint and sharing</b>            | yes                | no             | yes                                   | yes  | yes   | no                                     | no                                     |
| <b>Overall good for</b>                       | <b>Clusters</b>    |                |                                       |  | <b>Non-clusters</b>   |  |  |

# 3

## Strategies

Best practices for typical HPC workloads

# *Strategies and Data Management*

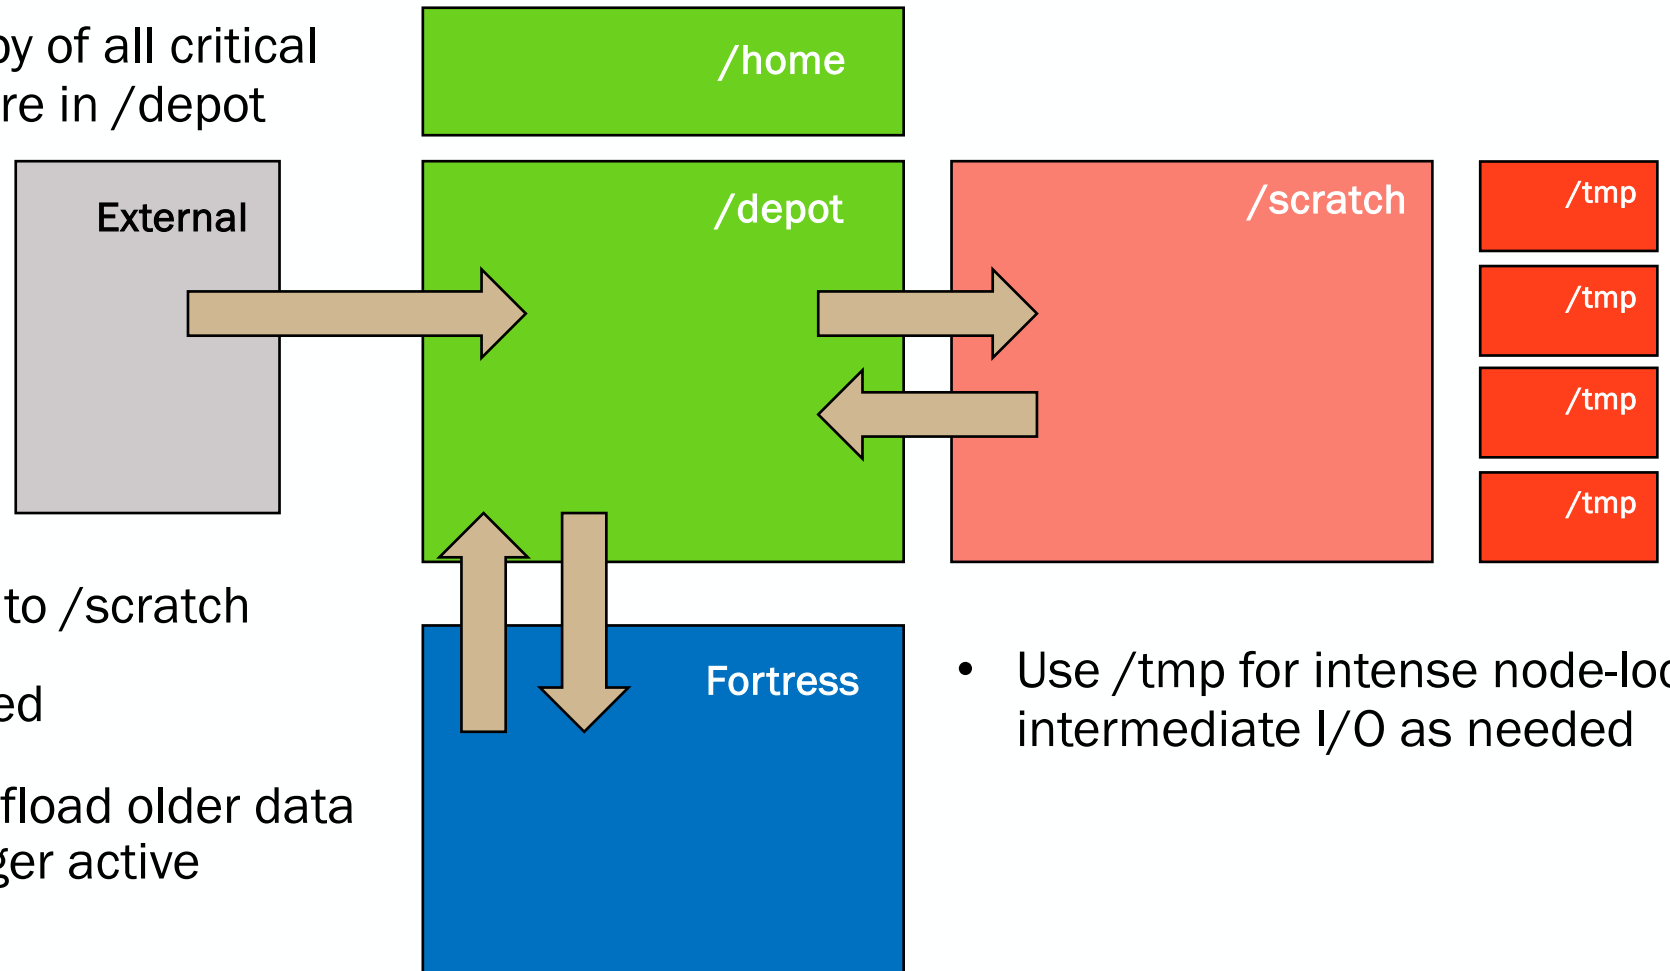
RCAC Storage offerings are designed around these scenarios

- Personal code, executables, scripts – store in **/home** or **/depot**
- Lab-wide code, executables, scripts – store in **/depot**
- Data and results (input / output):
  - produce/analyze – store in **/scratch**
  - retain/archive – store in **/depot** or **Fortress**
- Lab-wide documents, manuscripts, etc. – store in **/depot** or **Box.com**
- Regulated data – store in **REED** folder (**Box.com**)

# Strategies and Data Management

Basic recommendations for typical use-cases

- Store primary working copy of all critical research data and software in /depot



- Personal configuration in /home

- All read/write during jobs to /scratch
- Stage in and out as needed
- Archive to Fortress and offload older data from /depot when no longer active

- Use /tmp for intense node-local intermediate I/O as needed

# *Strategies and Data Management*

Miscellaneous recommendations for Depot vs Fortress archiving

- **Don't wait to send to Depot or Fortress,**
  - risky (what if you forget to migrate something)
  - why not put an *htar* in your job script?
- **Fortress is cheap**
  - keep things in both
  - offload older data from /depot when no longer active
  - stage in and out as needed



# 3

## Transfer Tools

How to move data between systems and your own machine.

# *Migrating Data*

How to move data between systems/clusters and external machines?

## **Command-Line Tools**

- Standard UNIX tools: `scp`, `sftp`, `rsync`
- To Fortress: `hsi`, `htar` (available for installation)

## **GUI Tools**

- Network drive (SMB)
- *WinSCP*, *MobaXterm* on Windows, *CyberDuck*
- Countless others!

# *Migrating Data*

How to move data between systems/clusters and external machines?

## **Demo**

- Standard UNIX tools: `scp`, `sftp`, `rsync`

Using SCP to migrate folder to RCAC Cluster (even on Windows)

```
C:\Users\Me> scp -r Data\MyData username@negishi.rcac.purdue.edu:/depot/mylab/data/  
...
```

Using RSYNC to migrate folder to RCAC Cluster (even on Windows)

```
C:\Users\Me> rsync -avh Data\MyData\ username@negishi.rcac.purdue.edu:/depot/mylab/data/mydata/  
...
```

# *Migrating Data*

How to move data between systems/clusters and external machines?

## **User Guides**

- Every resource has a user-guide
- *Education > User Guides* ([rcac.purdue.edu/knowledge](http://rcac.purdue.edu/knowledge))

## **Other Platforms**

- Globus ([transfer.rcac.purdue.edu](http://transfer.rcac.purdue.edu) | [globus.org](http://globus.org))

# *Fortress Tools*

Moving data to and from the HPSS archive

## **Command-line (on clusters)**

- `hsi` is a remote shell-like interface similar to `sftp`
- `htar` is like UNIX `tar` but sends to Fortress

## **Globus**

- Web interface (choose ***Purdue Fortress HPSS Archive***)
- Command-line interface with `gLOBUS` command

# *Globus*

Moving data with Globus

## **Command-line (on clusters)**

- In every user-guide under “File Storage and Transfer”
- Go to [transfer.rcac.purdue.edu](http://transfer.rcac.purdue.edu)
- In the “File Manager” choose source and destination
- Can be scheduled (repeated on a timer!)
- Getting started: [docs.globus.org/how-to/get-started](http://docs.globus.org/how-to/get-started)
- Can also be used for sharing (Shared Endpoints)

# 4

## What now?

Visit [rcac.purdue.edu/training](https://rcac.purdue.edu/training) to see other topics related to research storage

# *Thank You*

Slides and recording: [rcac.purdue.edu/training/storage](https://rcac.purdue.edu/training/storage)

Send an email to [rcac-help@purdue.edu](mailto:rcac-help@purdue.edu) for help and consulting

Coffee hour consultations: [rcac.purdue.edu/coffee](https://rcac.purdue.edu/coffee)





# *Training Post-Survey*

Please fill out afterward to give us helpful feedback

Scan this! →

