OVERVIEW OF RCAC STORAGE RESOURCES

Lev Gorenstein, Senior Computational Scientist

PURDUE UNIVERSITY | Rosen Center for Advanced Computing

RCAC Storage Overview

Outline

What to Expect From This Talk

Objectives

- Overview of storage needs
- Storage solutions provided by RCAC (cluster or no cluster)
  - Features, capabilities, do's and dont's
- General strategies and approaches
- Demo – storage systems and access

RCAC Storage Overview

RCAC Storage Solutions

RCAC Storage Solutions

There is more to Research Computing than just computing!

- Sure, we are best known for our supercomputing clusters – but if it wasn't for storage and other cyberinfrastructure, where would you put all those nice things you've just calculated?
- See www.rcac.purdue.edu/storage for all our storage options
- An interactive storage solutions finder: www.rcac.purdue.edu/storage/solutions/
- Also check out www.rcac.purdue.edu/services for other services we provide

RCAC Storage Solutions

Storage is not created equal!

- **Capabilities**
  - ***Capacity:*** limited vs. unlimited
  - ***Performance:*** fast vs. slow
  - **Medium:** SSD vs HDD vs. tape
  - ***Longevity:*** short- vs. long-term
  - ***Access control and sharing:*** just me / my lab / group of collaborators / world?
  - ***Locality:*** local vs. cloud
  - ***Means of access:*** directly on filesystem vs. special tools
- **Cost**
  - ***Redundancy:*** "if disk crashes, am I safe?"
  - ***Recoverability:*** "if I *oopsed* a file, can I get it back?"

RCAC Storage Solutions

User needs are not created equal either!

- **Use cases**
  - I'm a cluster user, mostly ***compute-oriented***
  - I'm a cluster user, mostly ***data-oriented***
  - My lab needs ***shared space*** (for programs / data / documents) – ***even if non-cluster***!
  - I need ***long-term archive*** of my precious data – ***even if non-cluster***!
  - I need to ***share my data*** (or ***my code***) with collaborators (or make it accessible to the world per journal or funding agency rules)
  - I need to ***collaborate on a manuscript*** with local or remote collaborators
  - I have a ***scientific instrument*** and need a data pipeline (*"acquire – analyze – store"*)

RCAC Storage Solutions

RCAC provides multiple storage solutions for Purdue researchers

Local and "local" storage

- Are (or could be) mounted on clusters or lab/personal computers:
  - ***Home directory***
  - ***/tmp***
  - ***Scratch***
  - ***Data Depot***
- Indirect access:
  - ***Fortress*** tape archive

Cloud-based storage:

- ***Box.com*** folder
- ***REED folder*** for restricted data (based on Box.com, but with heightened security)
- **github.itap** (not quite storage, but kind of is)
- ***PURR*** (Libraries) - collaborative working space and data-sharing platform
  - Data management plans, DOI minting for generated datasets, citations tracking, etc.

www.rcac.purdue.edu/storage

RCAC Storage Solutions

Cluster Home directories

$HOME
/home/myusername

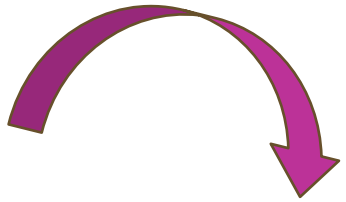

Features:

- Small (25 GB) and mildly performant
- Redundant hardware, never purged, protected by snapshots
- Cluster-specific (shared between nodes withing each individual cluster, but not across clusters)
- Everyone with RCAC account has it, but probably of little use outside of clusters
- *"My home is my castle"* – belongs to you. A known pain point for PIs (if student graduates and leaves important data behind)

Good for:

- **Your personal codes, programs, scripts, configuration files**
- **Your data and results** (if small)
- OK to run jobs off of it, but only if mostly number-crunching (little I/O)
- **Medium and long-term** storage

RCAC Storage Solutions

Cluster /tmp directories

/tmp
(the traditional temporary directory)

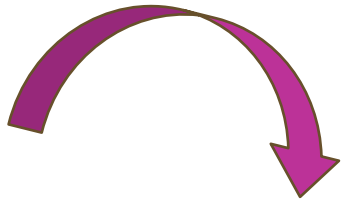

Features:

- What people used before mighty scratches appeared
- Moderate (200-400 GB), pretty performant
- Zero redundancy, regularly purged, and no snapshots (*"if it's gone, it's gone!"*)
- *Node-local* (each node has its own /tmp, only shared with other users on this same node)
- World-readable and writable by default

Good for:

- OS internal needs and almost nothing else
- **NOT** for valuable codes, programs, scripts, configuration files
- **NOT** for long-term precious final results
- **NOT** for something that needs to go between nodes

- Rarely needed (but priceless when it is). Unless you know what you are doing, **just use scratch or Depot instead**

RCAC Storage Solutions

Cluster Scratch directories

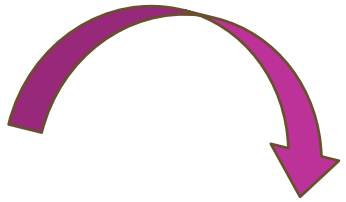

$RCAC_SCRATCH
/scratch/mycluster/myusername

Features:

- Huge (100+ TB) and very performant
- Internally redundant, but regularly purged (files inactive for 60 days are deleted), and no snapshots (*"if it's gone, it's gone!"*)
- Cluster-specific (shared between nodes withing each individual cluster, but not across clusters)
- Engineered for large sustained I/O (not so much for myriad of small files though – where /tmp might have an edge)
- Belongs to you (pain point for PIs)

Good for:

- Your massive generated intermediate data
- Perfect to run data-intensive jobs off of it.
- NOT for valuable codes, programs, scripts, configuration files
- **NOT for long-term precious final results**
- Short-term storage
- **Beware of weekly purges!** Watch email and use `purgelist` command regularly.
  Don't try to game the purge system
  (one day you'll forget and lose big)
- ***Just backup to Depot or Fortress!***

RCAC Storage Solutions

Data Depot space

/depot/mylab

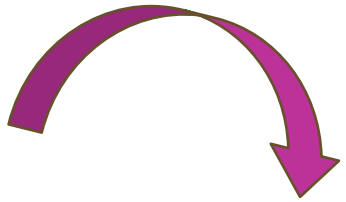

Features:

- Large (100 GB free, can grow as needed in 1 TB increments for $70/TB/year), reasonably performant
- Redundant hardware, never purged, protected by snapshots
- Visible on all RCAC clusters as well as on- and off-campus
- Access limited to members of your research group, very flexible controls for subfolders (per user, per project, common, etc) – collaborators too
- Basically, *"a home directory for labs"*
- **Don't have to buy cluster nodes! (60% don't)**
- *Belongs to the PI!*

Good for:

- **Your personal *and lab-shared* codes, programs, scripts, configuration files** (e.g., a common Anaconda environment everyone uses)
- **Your data, massive intermediate and final results**
- OK to run jobs off of if not too heavy of I/O (for heavy cases, use scratch)
- General lab documents
- Visible as a network drive on personal computers
- **Medium and long-term** storage

RCAC Storage Solutions

Fortress tape archive

Features:

- Huge (25 **PB**), free and practically unlimited
- Tape library with a robotic arm and a disk cache in front *(i.e. fast uptake, slow egress)*
- Redundant hardware, never purged, protected by multiple physical copies on separate tapes.
- Accessible from all RCAC clusters as well as on- and off-campus with additional tools
- Everyone with an RCAC account gets *personal* Fortress space. Additionally, labs with Depot also get *lab Fortress space* (access controlled by `mylab-data` Unix group membership)
- *Don't have to own cluster nodes!*
- Your space belongs to you, *group spaces belongs to PIs!*

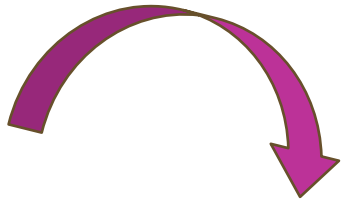

Personal and group spaces
/home/myusername and /group/mylab

Note: not your cluster /home or /group!

Good for:

- **BACKUP of your precious personal or lab-wide codes, data and results**
  - *"Vita brevis, ars longa"* – ancient wisdom that snapshots are transient and only true backups last!
- **NOT for running jobs off of it!** But perfectly ok to generate/process massive data in scratch, and then send to Fortress in one job
- Native tools `hsi` and `htar` available for Linux, SFTP for other command lines, and Globus endpoint makes life great for everyone
- **Tape prefers few large files over ton of small ones – bundle 'em!**
- **Long-term archival storage**

RCAC Storage Solutions

Box.com research lab folder

[L1FR] mylab

Features:

- Cloud-based sharing and collaboration system.
- Free. Login with an `@purdue.edu` address at purdue.box.com
- Individuals get a 1 TB personal spaces, for labs we set up unlimited Level 1 folder
- *Catch: no single file can exceed 50 GB*
- Performance good for day-to-day, inadequate for HPC (pretty kludgy to get working on clusters)
- Redundant, never purged, protected by versioning
- Very flexible document sharing with members of your research group and collaborators (internal and external)
- *Don't have to buy cluster nodes!*
- Lab folders belong to the PI!

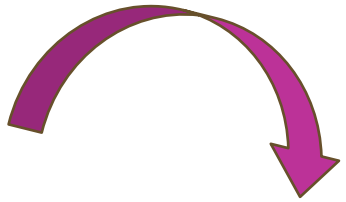

Good for:

- **Your personal and lab-shared documents, spreadsheets, manuscripts**
- Data (within 50 GB per file limitation)
- **Internal and external collaboration** is super easy (even data collection by link from non-tech-savvy participants!)
- **NOT** for storing sensitive or restricted data such as NIST SP-800-171, FERPA, HIPAA, etc. (but see REED folder)
- **Medium and long-term** storage

RCAC Storage Solutions

REED folder

[L3] mylab

Features:

- Cloud-based managed storage for projects requiring compliance with regulated or sensitive data
- Basically, a Box.com folder with heightened security
- Free, but requires training and execution of Data Use Agreement
- Unlimited, but same catch with 50 GB per individual file
- See Box folder description for more

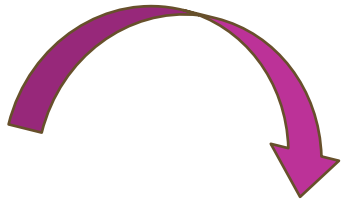

Good for:

- **Your personal and lab-shared documents, spreadsheets, manuscripts for research requiring sensitive or restricted data – such as FERPA, HIPAA, etc.**
- **Medium and long-term** storage

snapshots != backup!	$HOME	/tmp	$RCAC_SCRATCH	/depot/…	Fortress (HPSS)	Box folder	REED folder
Capacity	25 GB	150-400 GB	varies by cluster… 100 TB and up	100 GB free, then paid in 1 TB increments	unlimited	unlimited, but under 50 GB per file	unlimited, but under 50 GB per file
Resilience to hardware failures	yes	no	yes	yes	yes	yes	yes
Resilience to human errors	yes (snapshots)	no	no	yes (snapshots)	no	some (versioning)	some (versioning)
Subject to purging	no	yes	yes	no	no	no	no
Performance	medium	medium to high	high	medium	slow to very slow	slow to very slow	slow to very slow
Designed for HPC (running jobs off it)	no	no	yes	in moderation	• no (as main I/O) • yes (for staging and archiving)	no	no
Common access within cluster	yes	no	yes	yes	yes (hsi/htar)	possible but kludgy	no
Common access across clusters	no	no	no	yes	yes (hsi/htar)	possible but kludgy	no
Advanced sharing ACLs (past ugo/rwx)	no	no	no	yes	no	yes	yes
Globus endpoint and sharing	yes	no	yes	yes	yes	no	no
Overall good for	**Clusters**					**Non-clusters**	

RCAC Storage Overview

Strategies and Data Management

Strategies and Data Management

Types of storage and when to use them

- **"Capacity, speed, longevity – pick any two"**

- RCAC storage offerings are designed around these scenarios:
  - **Personal code, executables, scripts** – develop and store in ***$HOME*** (or ***/depot***)
  - **Lab-wide code, executables, scripts, settings** – in ***/depot***
  - Data and results (input/output):
    - **produce/analyze** – in ***$RCAC_SCRATCH*** (possibly ***/depot***)
    - **permanently store** – in ***Fortress archive*** and/or in ***/depot***
  - **Lab-wide documents, manuscripts, etc.** – in ***/depot*** and/or ***Box folder***
  - **Regulated data** – in ***REED folder***

Strategies and Data Management

"Hotness", or "ease of reach" analogy

Scratch.
Sizzling hot storage

Tools and data in your Home directory

Box.com

Home directory

As results cool off, they move from desk to binders to bookshelves and eventually to....

Fortress.
Cold storage

Data Depot.
Has more tools and a hot/warm storage.

Migrate between storage systems throughout data lifecycle

- **Amount of data never goes down**
- **Old data is precious**
  - Do you need 10 years worth of old data in hot immediate vicinity in your Depot space?
    - Very convenient
    - Risky (what if you accidentally `rm` in the wrong place?)
  - Maybe copy the ~~first 7~~ all 10 years out to Fortress, and delete the first 7 from Depot?
    - Peace of mind (full backup!) and $$ savings (fewer TB needed!)
    - If you need old data five years from now, get it from tape
    - Yes, it may take an hour to get it from tape. Negligible compared to 5 years.

- **Current data is precious, too!**
  - You finished a run, generated some raw data? Send them to Fortress or Depot first, then analyze.
  - Why not put an `htar` in almost every job script?

- Also true for large data-generating scientific instruments (sequencers, microscopes, etc.)

RCAC Storage Overview

Demos

Demo – access by SCP, SFTP, Windows Network Drive

- Every one of our resources (compute and storage) has a User Guide.
- On our web site, *Education -> User Guides* (or just www.rcac.purdue.edu/knowledge)
  - pick your resource
  - see *"File Storage and Transfer"* section for all supported access methods
- Many command-line utilities or GUI tools (WinSCP, CyberDuck, MobaXterm)

Demo – Fortress access by `hsi` and `htar`

- All supported methods: www.rcac.purdue.edu/knowledge/fortress/storage/transfer
  - `hsi` and `htar` are special command-line utilities for HPSS tape archive (installed on all clusters, available for download for other Linux computers)
  - Other OS can use SFTP (via fortress-sftp.rcac.purdue.edu) or Globus
- `hsi` is remote shell-like interface to Fortress
  - navigate, list, manipulate files, transfer things in or out
  - `ls`, `cd`, `mkdir`, `cp`, `mv`, `put`, `get` – great for interactive work (but can also be batched)
  - www.rcac.purdue.edu/knowledge/fortress/storage/transfer/hsi
- `htar` closely mimics regular tar
  - create, list, extract tarballs stored on Fortress (also makes a matching index file for faster searches)

```
htar -cPvf /path/on/fortress/archive.tar datadir(s)      # store
htar -xvf /path/on/fortress/archive.tar  [file(s)]       # extract
htar -tvf /path/on/fortress/archive.tar                  # list
```

  - www.rcac.purdue.edu/knowledge/fortress/storage/transfer/htar

Demo – Globus

- In every RCAC resource's User Guide under *"File Storage and Transfer"* section
- Go to transfer.rcac.purdue.edu or globus.org and login using *"Purdue University Main Campus"* as organization from drop-down menu. Use BoilerKey 2FA.

- Globus transfers:
  - In the *"File Manager"* tool, search for source collection in one panel, destination collection in another panel… navigate paths, highlight files, tweak options, hit *"Start!"*
  - Can be scheduled (repeated on a timer!)
  - Globus getting started guide: docs.globus.org/how-to/get-started/

- Globus can also be used for sharing – even when recipient does not have account on our system!
  - *"European colleague needs to get (or put) a terabyte of data in my scratch space"*
  - Globus sharing guide: docs.globus.org/how-to/share-files

Summary

TL;DR summary

- One size storage does not fit all
- RCAC provides multiple storage options for individual researchers and labs
- You do not have to be a cluster user!

- Think about a comprehensive data-management strategy for you and your lab based on your needs, and we will help implementing it
- "Hot / warm / cold" storage is a useful paradigm. Saves money and nerve cells, too.

- Data pipelines (e.g., "acquire on an instrument – send to cluster for processing – sent to storage for work and archival") are on the rise. Talk to us if you have one!
- Globus is awesome

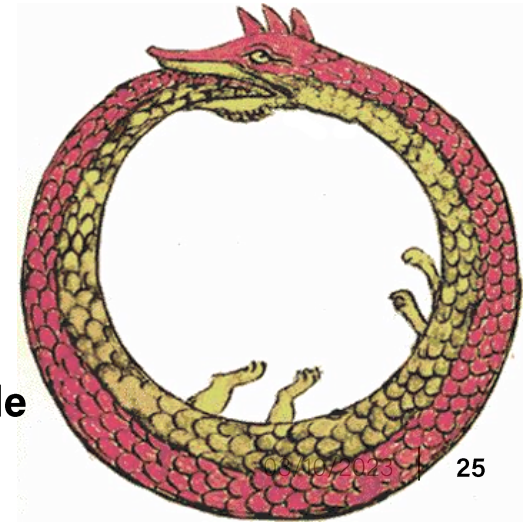

Research Data Lifecycle

BACKUP!!!

RCAC Storage Overview

What Comes Next?

What Comes Next?

Upcoming Seminars

- Open OnDemand 101: March 24
- Containerized Bioinformatics Applications for HPC: March 29
- Workflow Automation Tools for Many-Task Computing: March 30
- NLP 101: March 31
- Time Series Forecasting 101: April 7
- https://www.rcac.purdue.edu/news/events

THANK YOU

Feel free to reach out to lev@purdue.edu with questions.

Slides and recording are posted at:
https://www.rcac.purdue.edu/training/storage

General help: rcac-help@purdue.edu

Coffee Hour: https://www.rcac.purdue.edu/coffee