H3ABioNet Data Management workshop

June 6th, 2014

Radhika Khetani, Ph.D.

Technical Lead at HPCBio

University of Illinois, Urbana-Champaign

Data Management (Cluster/Server)

1. Data storage
2. Data security
3. Data transfer

Data Storage Outline

- File system organization
- RAID configuration
- Monitor disks for failure
- Data backup
- Data Archiving
- Network (distributed) file system

Data Storage Outline

File system organization

- For a shared system, it is important to start with an organization schema that will enable better storage, security and flexibility
- Divide users into several directories alphabetically
- Make distinct directories for storing databases and applications
- Make sure to accommodate groups working on shared data by giving them shared and non-shared user spaces
- A detailed hierarchical structure (intrinsic to most systems) enables granting hierarchical and stringent access permissions
- It also makes backups more manageable

RAID configuration

- Redundant Array of Inexpensive/Independent Disks
- Combines multiple disks into a logical component for data redundancy
- Data are distributed to several disks, and there are several schemas that can be used.
- Level of redundancy and performance (I/O) are the 2 major factors to be considered
- “Fault tolerance”, “Striping”, “Parity” and “Mirroring” are words commonly associated with RAID configurations

RAID configuration

- Fault tolerance – “the property that enables a RAIDed disk configuration to continue operating properly in the event of the failure of one or more disks”
- Parity – “If a drive in the array fails, remaining data on the other drives can be combined with the parity data (using the Boolean XOR function) to reconstruct the missing data”
- Striping – “segmenting logically sequential data, such as a file, and storing them on different disks”
- Mirroring – “replication of data onto separate physical hard disks in real time to ensure continuous availability”

RAID configuration

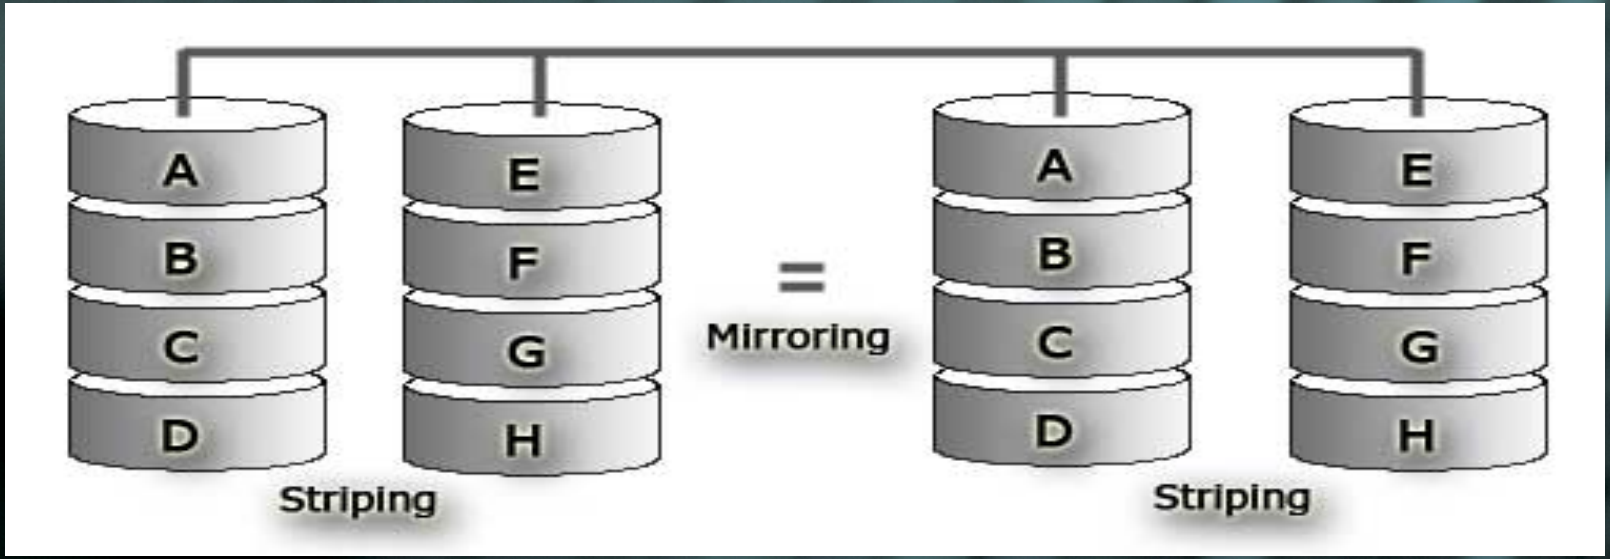

- Striping – “segmenting logically sequential data, such as a file, and storing them on different disks”
- Mirroring – “replication of data onto separate physical hard disks in real time to ensure continuous availability”

RAID configuration

- RAID0 – fastest and efficient, but offers no fault-tolerance.
- RAID1 – fault-tolerant, and requires twice the number of disks

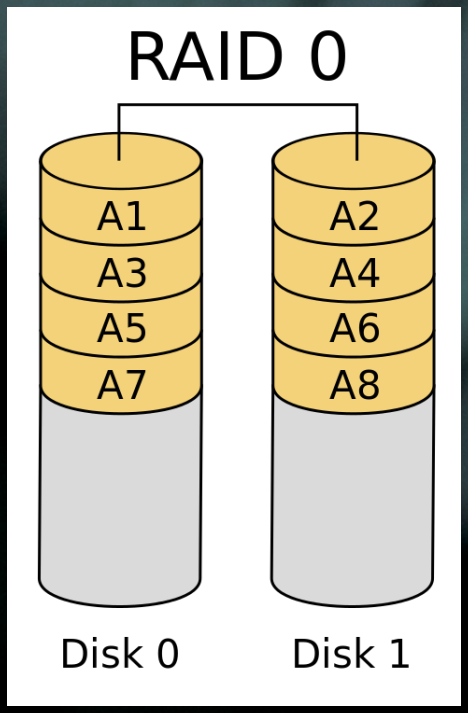

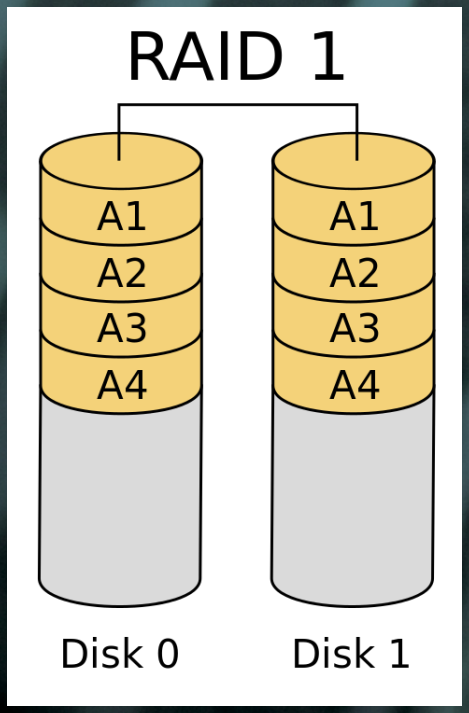

RAID configuration

- RAID5 – used in multi-user environments which are not I/O sensitive
  - needs a minimum of 3 disks
  - distributed parity
  - can allow for 1 disk failing
- RAID6 –similar to RAID5 however it allows extra fault tolerance
  - needs a minimum of 4 disks
  - 2 types of distributed parity
  - can allow for 2 disks failing

RAID configuration

- RAID5 – used in multi-user environments which are not I/O sensitive
  - needs a minimum of 3 disks
  - distributed parity
  - can allow for 1 disk failing
- RAID6 –similar to RAID5 however it allo
  - needs a minimum of 4 disks
  - 2 types of distributed parity
  - can allow for 2 disks failing

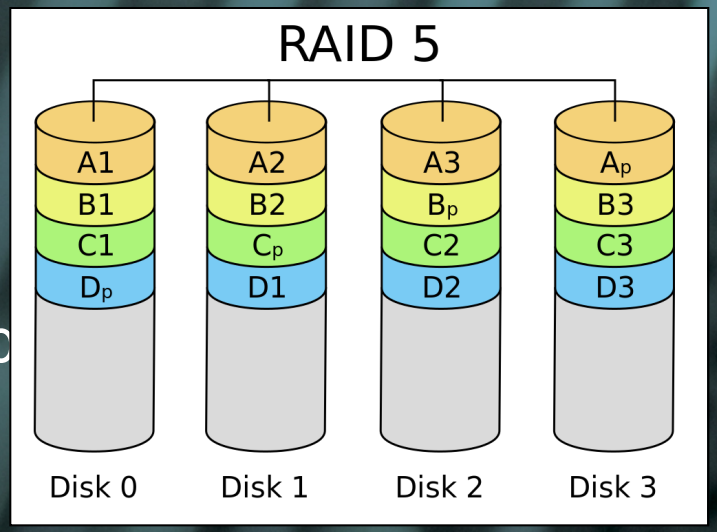

RAID configuration

- RAID5 – used in multi-user environments which are not I/O sensitive
  - needs a minimum of 3 disks
  - distributed parity
  - can allow for 1 disk failing
- RAID6 –similar to RAID5 however it allows extra fault tolerance
  - needs a minimum of 4 disks
  - 2 types of distributed parity
  - can allow for 2 disks failing

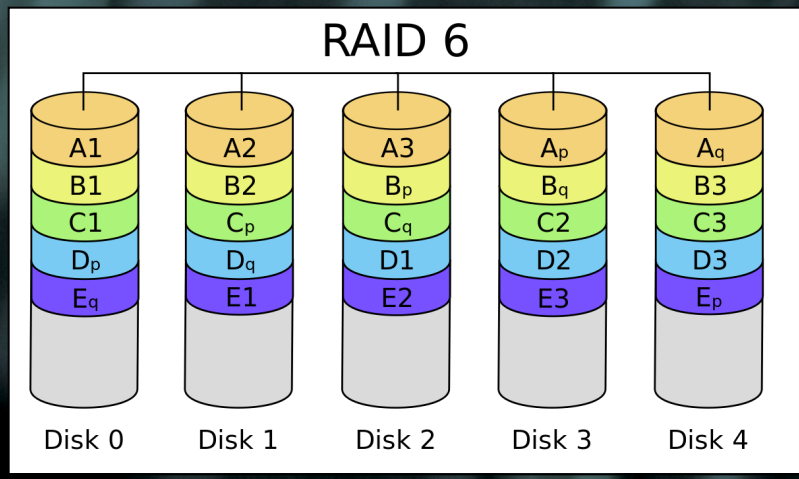

RAID configuration

- RAID6 –similar to RAID5 however it allows extra fault tolerance
  - needs a minimum of 4 disks
  - 2 types of distributed parity
  - can allow for 2 disks failing
  - recommended!!
  - always keep 2 extra disks handy
  - with constant monitoring, this can provide a relatively stable storage environment

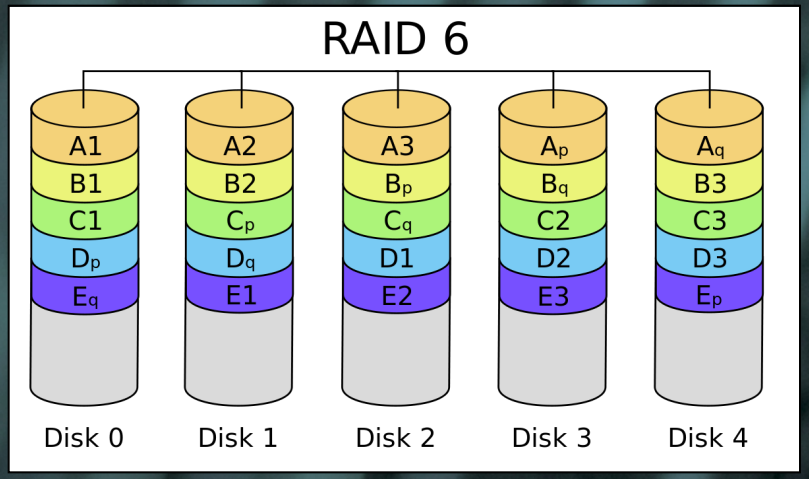

Monitor disks for failure

- No matter which RAID set up you have, set up a system for daily disk (RAID array) monitoring
- Any script that needs to be run daily can be set up in /etc/cron.daily/ for Linux systems
  - For example, a script using mdadm to test the disks can be added to the directory
  - If set up correctly, the results of “mdadm –monitor” will be delivered to your inbox daily or weekly or monthly

Network File System (NFS)

- "NFS allows one computer (a client) attached to a network to access the file systems present on the hard disk of another computer (an NFS server) over the network."
- For a system with several computers connected over a local network, the file system can be distributed across them using this set up, e.g. compute clusters
- Each disk should be RAIDed appropriately
- User should not be able to differentiate between a distributed system and a local system, both from the standpoint of directory structure as well as speed of access (internal network speed notwithstanding)

Data backup versus Data Archiving

- Backing up is the act of making sure that all the data are copied to a completely separate disk array, ideally at a different location, regularly
- Archiving is the act of backing up compressed data for the long-term, and is done when a project completes or reaches a breaking point

Data backup


GLASBERGEN
"We back up our data on sticky notes because sticky notes never crash."

Data backup

- Ensure there is enough disk space available for backup
- Backup everyday (night)
- Depending on the type of data and the amount of data you have, you might want to consider different solutions
  - For ≤ 50TB on a single server, rsync works very well
  - If you have multiple servers with many large files, Amanda is an open source solution
  - For >100TB, you might want to consider a commercial solution like Symantec's NetBackup, Bacula (open source), etc.
- The local network connection of 1Gbit is recommended when backing up large datasets

Data archiving

Data archiving

- Make 2 copies and store in 2 different locations
- Magnetic tape data storage
  - Linear Tape-Open (LTO)
  - Stores ≥2.5TB, but much cheaper than regular hard drives
  - Ultra reliable for an extended period of time - “50 years from now you can tape the tape together with tape” – D. Slater
  - Built in Encryption
  - Cheaper!
  - Requires special set up to read and write
- Amazon Glacier, and other commercial solutions

Data backup and Data Archiving

- md5sums – a digital fingerprint for a file
- Always compare the md5sum before and after transfer to ensure data integrity

Costs

- Storage is cheaper than it was 5 years ago, but if you consider the RAIDed set up along with backup facility, storage is not cheap
- Depending on the users and type of data, some facilities choose to have quotas
- Usually these quotas are associated with an additional cost for the additional storage
- Costs for archiving and long-term storage of tapes should also be considered for maintaining standards

Data Security Outline

- Permissions and access
- Firewalls
- Monitor system
  - illicit activity
  - vulnerabilities

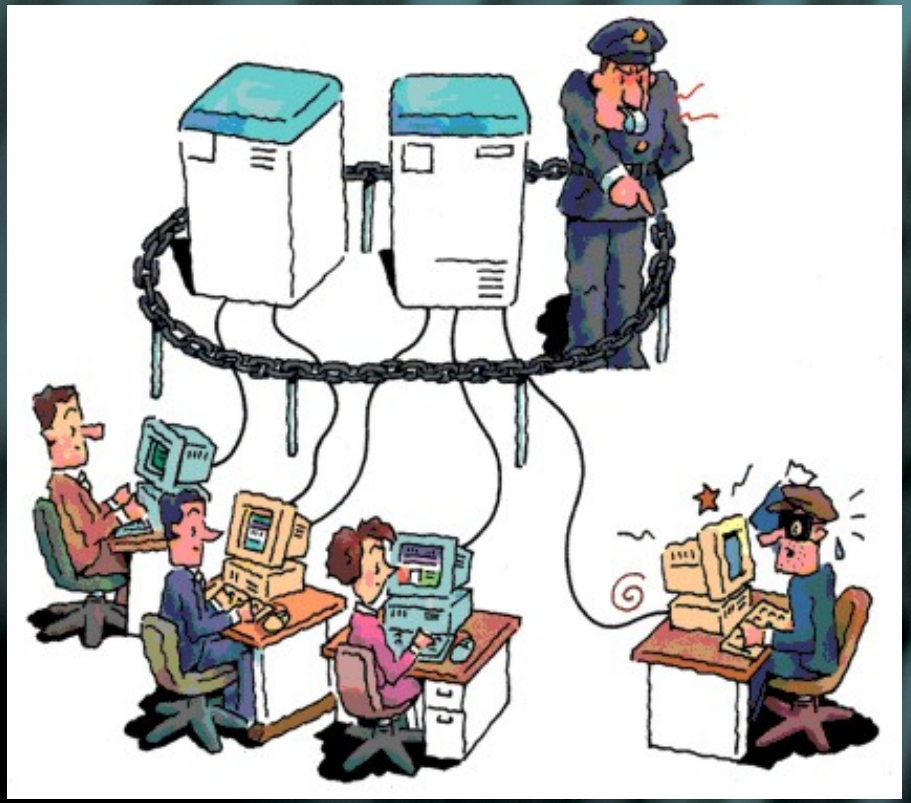

Permissions

- In multi-user systems, access and access restrictions are key
- Typically, you are the owner of every file/directory you create or bring into a system
- What other files and directories you can read, write or execute will depend on how the system is set up

```
-rw-rw-r-- 1 rkhetani hpcbio  888 Sep 26  2013 R1_files.list
-rw-rw-r-- 1 rkhetani hpcbio  888 Sep 26  2013 R2_files.list
-rw-rw-r-- 1 rkhetani hpcbio  654 Sep 26  2013 normalization.sh
drwxrwsr-x 2 rkhetani hpcbio  32K Sep 27  2013 normalized_data
-rw-rw-r-- 1 rkhetani hpcbio 3.3G Sep 27  2013 R2_files.list.normalized_K25_C40_pctSD100.fq
-rw-rw-r-- 1 rkhetani hpcbio 3.3G Sep 27  2013 R1_files.list.normalized_K25_C40_pctSD100.fq
-rw------- 1 rkhetani hpcbio 6.5K Sep 27  2013 Normalize-Leafy.o466519
-rw-rw-r-- 1 rkhetani hpcbio  565 Sep 27  2013 inchworm_only.sh
-rw-rw-r-- 1 rkhetani hpcbio  564 Sep 27  2013 trinity-post-inchworm.sh
-rw-rw-r-- 1 rkhetani hpcbio 6.2M Sep 27  2013 runTrinity.stdout
```

Permissions

- In multi-user systems, access and access restrictions are key
- Typically, you are the owner of every file/directory you create or bring into a system
- What other files and directories you can read, write or execute will depend on how the system is set up

drwxrwxrwx – owner (u) group (g) others (o)

- A “sticky bit” is applied to shared directories to protect files such that only the owner has the ability to change permissions
- “chown” and “chgrp” are commands that let you change owner and groups respectively

Permissions and access

- Structured permissions across the file system are essential for a multi-user, multi-group system
  - More permissive at the top levels
  - Less permissive at the bottom levels
  - set-group-ID bit is used to set up directories so that any file created in the directory will retain the same group as the parent directory (setgid)
- Access Control Lists (ACL)
  - Extension of the standard UNIX permissions to give system administrators more fine-grained control
  - Easier to set up permissions for pre-determined groups

Firewalls

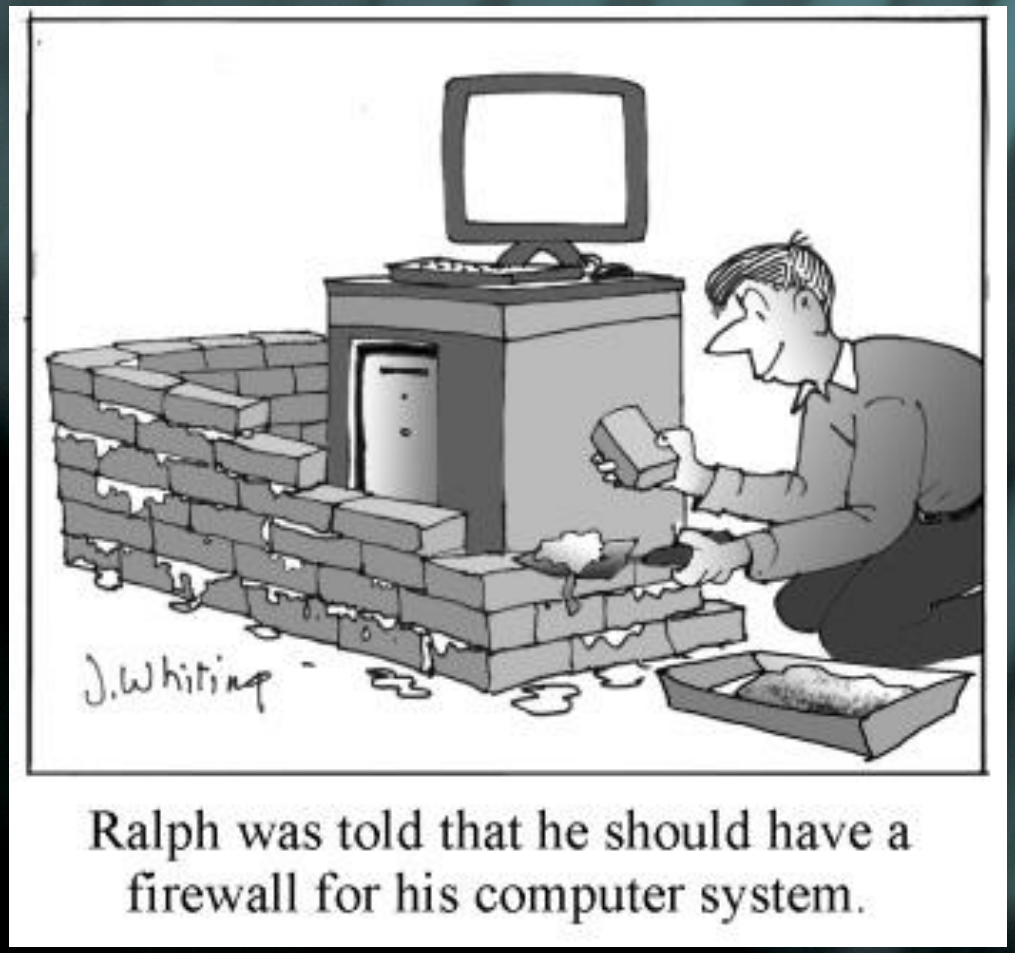

Ralph was told that he should have a firewall for his computer system.

Firewalls

- It is important to set up a firewall to protect data on a given system from hackers
- They can filter network traffic by content or user (IP addresses)
- Public areas of servers should be more heavily protected (web servers etc.)
- Private or restricted-access areas can be less heavily protected

Monitor system security

- Illicit activity
  - Check logs for anyone trying to unsuccessfully log on multiple times, e.g. ≥ 4 attempts in 1 minute means that IP address cannot log on again
  - Scan records regularly by setting up cron jobs (hourly or daily, Linux)
  - It is possible to also add the offending IP address to your firewall's list of untrustworthy addresses
- Vulnerabilities
  - Check computer systems, networks or applications for any security holes
  - Programs like openVAS offer network vulnerability checks and suggestions on how to manage an issue. Another commonly used paid program is Nessus

Data Transfer Outline

- FTP versus sFTP
- rsync
- GridFTP (Globus)
- Shipping data
- md5sums

FTP

- File Transfer Protocol
- FTP is a very commonly used network protocol to transfer files over the internet
- Files can be accessed anonymously
- Easily implemented and simple to use

FTP

Cons:

- There is no encryption
- Third parties can easily access the data moving through the network, and can even “hijack” the transfer
- Data can be edited *en route* by malicious third parties
- Login credentials are transferred in clear text and no authentication
- It cannot perform md5sum comparisons to ensure proper transfer

sFTP

- Secure File Transfer Protocol
- Uses ssh or secure shell (a cryptographic network protocol used widely)
- Data transfer is over a secure channel
- Both data and user information is encrypted
- A variety of authentication methods available

FTP versus sFTP?

- Use sFTP when possible!

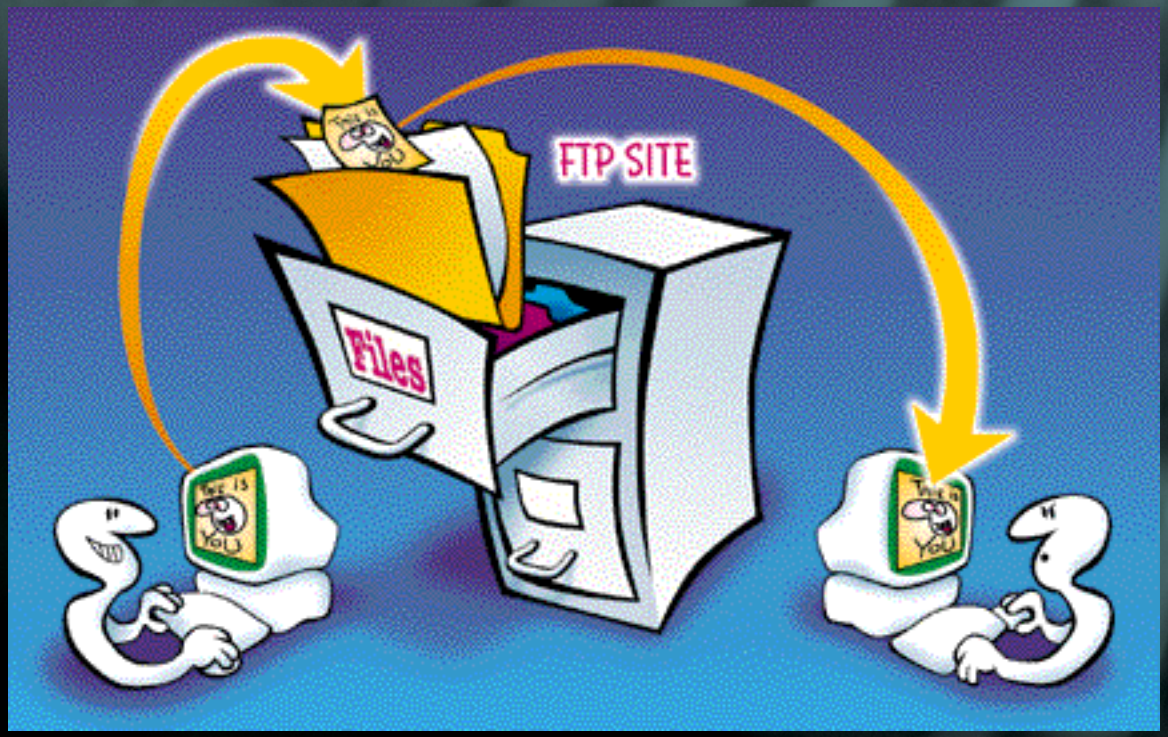

rsync

- Network protocol for UNIX-like systems
- It synchronizes files in the 2 locations by checking the modification time and size of each file in the destination directory
- It is also able to perform md5sum comparisons and modify the destination directory to match the one in the start location
- Makes for very fast and efficient transfers (especially for regular backups)
- Can be encrypted using ssh

GridFTP (Globus)

- Reliable, faster and secure File Transfer Protocol
- Developed to meet the needs of the grid computing community
- Data can be moved around to predetermined endpoints using an easy-to-use web interface
- Transfer can be set up and Globus takes care of making sure the data gets there intact
- Enables sharing large data files in a secure environment and over a secure network

Shipping Data

- Physical shipment of data via a secure courier is an alternative
- Encryption of the information is important, lots of commercial options available
- Shipping hard drives can damage data – too many moving parts that can be “roughed up”
- Shipping LTO tapes is an alternative
  - Built-in encryption
  - Not damaged or tampered with easily
  - Need hardware to read and write LTO tapes at both end points

md5sum

- A program that generates a digital fingerprint for a file
- Used to verify file integrity after transfer
- Example:

```
$ md5sum filetohashA.txt

595f44fec1e92a71d3e9e77456ba80d1  filetohashA.txt
```

- Files can be edited in a way that keeps the md5sum unchanged, but it requires a lot of work and is therefore rare

Conclusions and final thoughts

- Data security, storage and transfer are intertwined and in many ways and share concepts
- Data provenance is an important aspect of data storage as well as data transfer
- Collaborative research and large file sizes have made these concepts an important aspect of biology education

Acknowledgements

Dr. C. Victor Jongeneel

David Slater

Dr. Christopher J. Fields

Daniel Davidson

Wikipedia