

Tier 1 Site Evolution In Response to Experiment Requirements

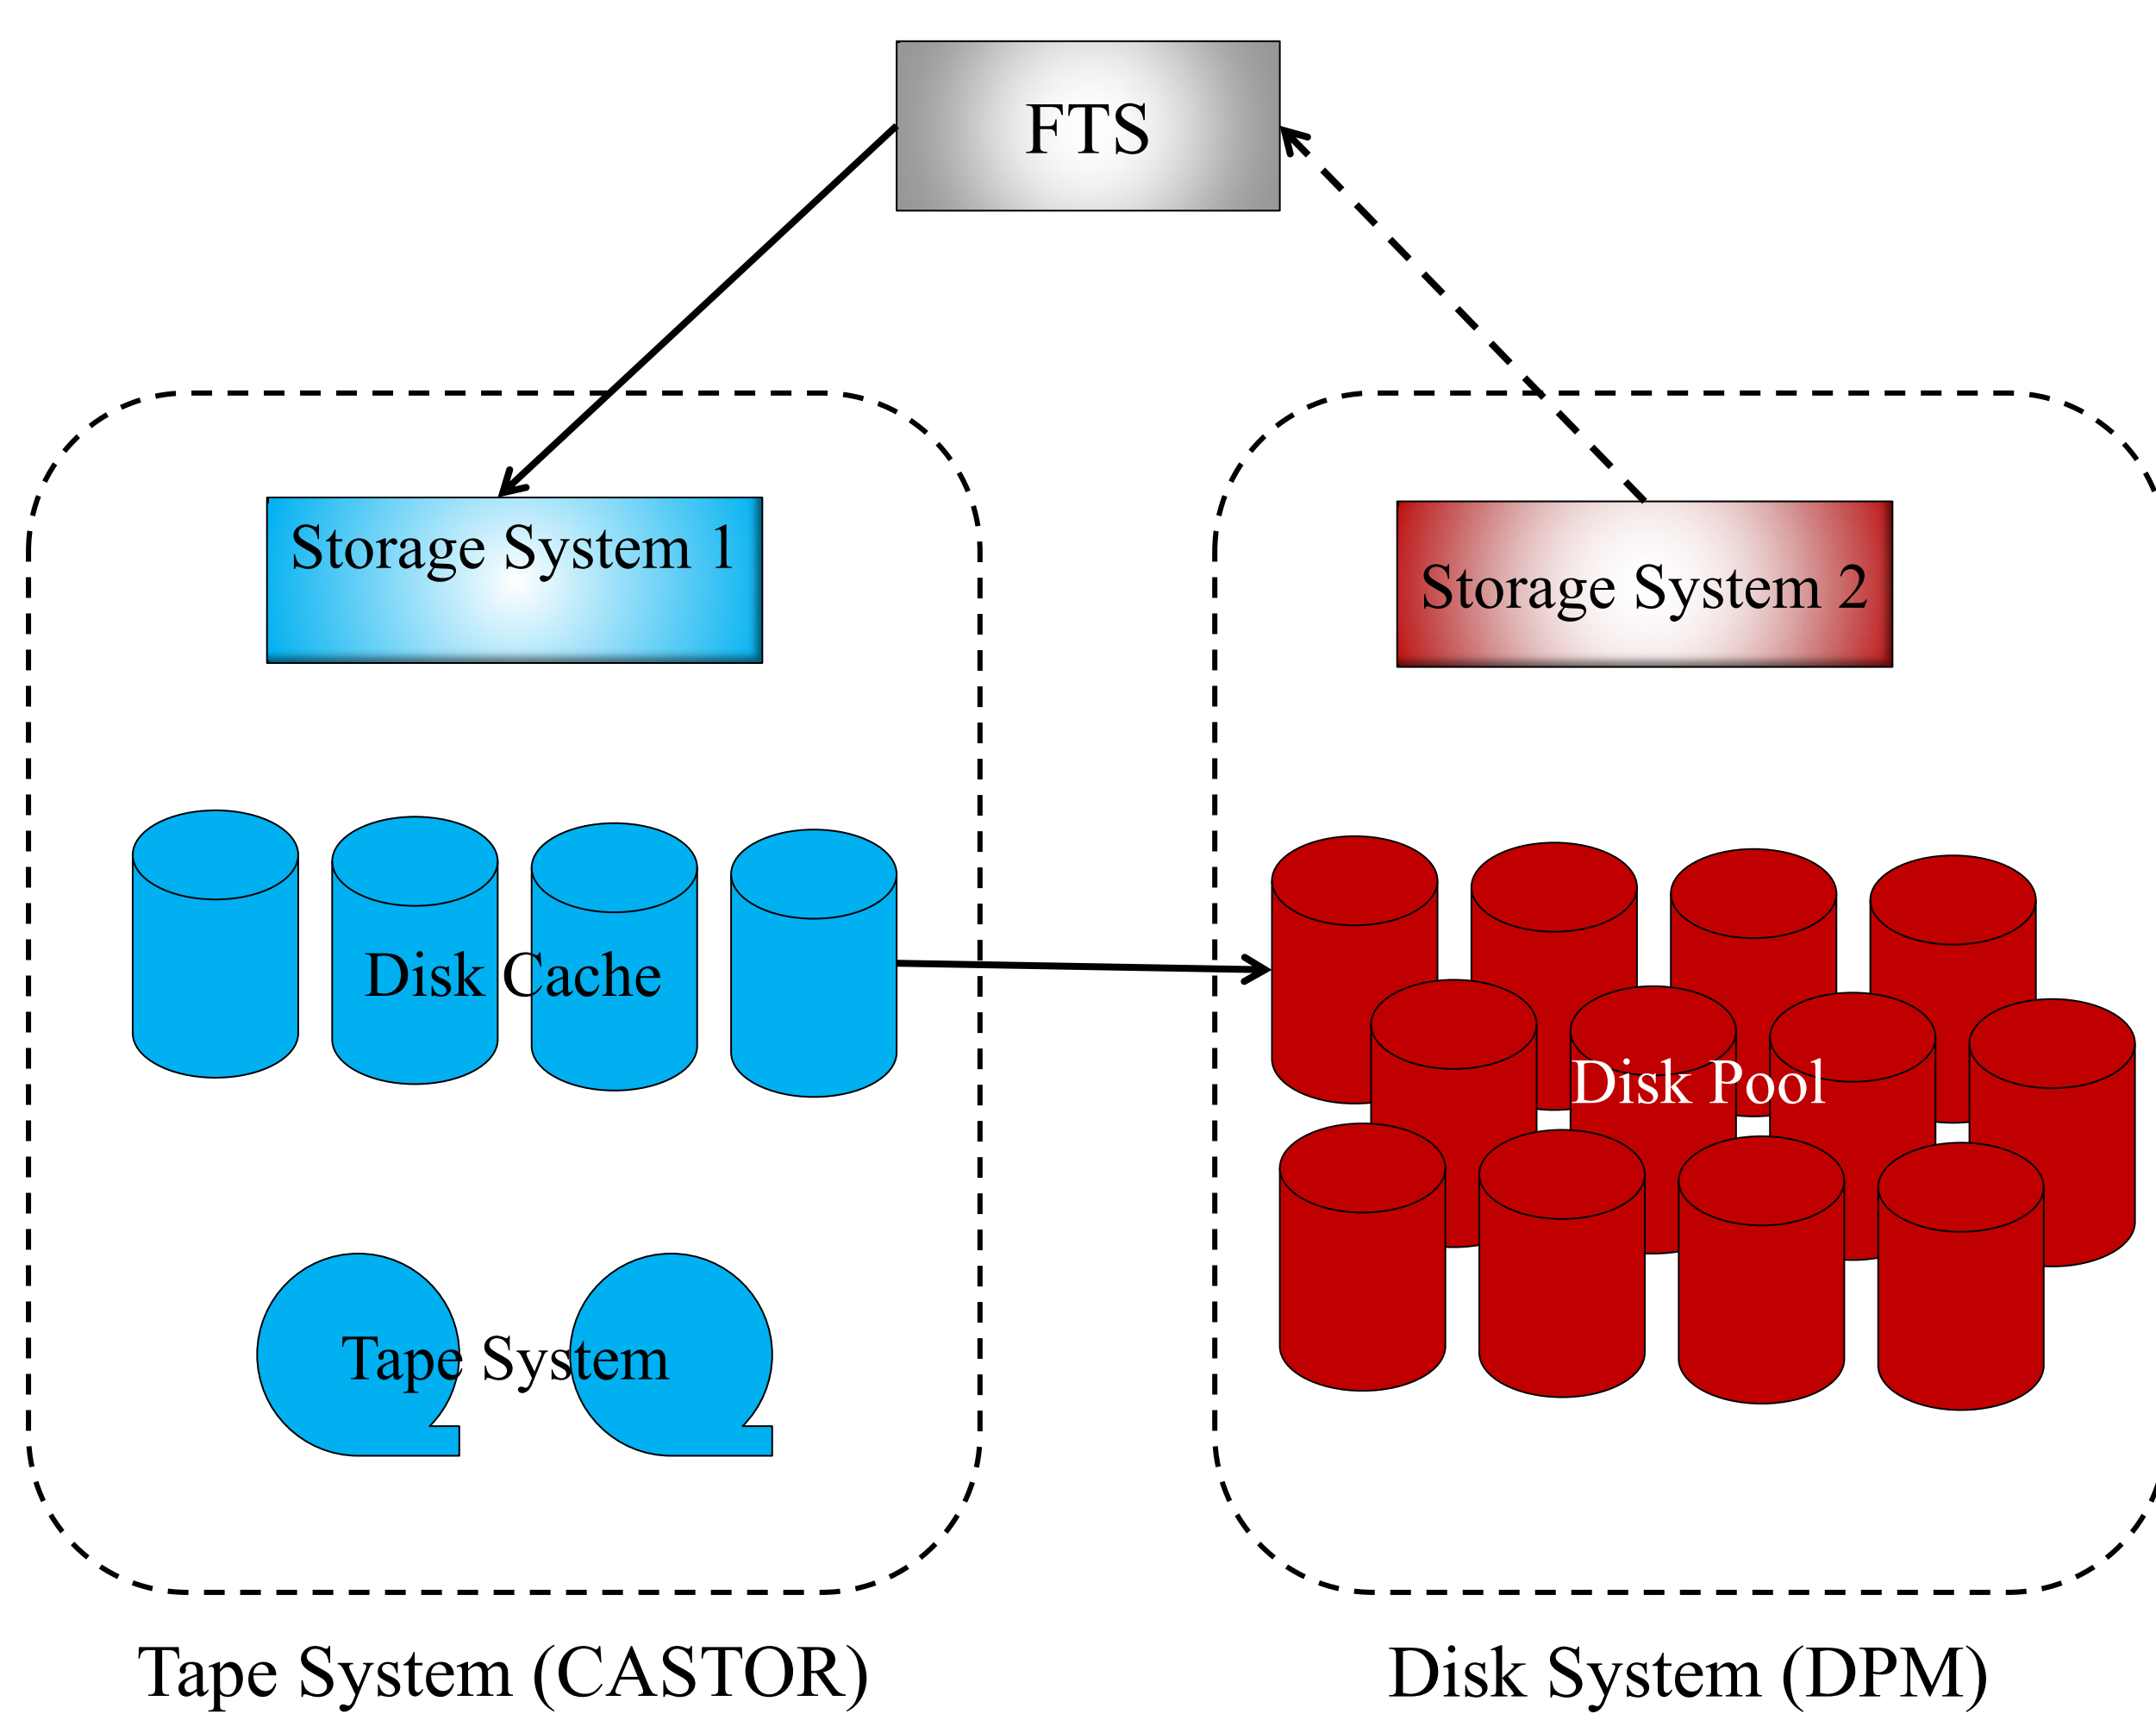
De Witt, Shaun (STFC), Huang, Jhen-Wei (ASGC)

Challenge is to meet the requirements of Disk/Tape Separation:

- Tape is becoming more of a true archival medium, tending towards the ideal WORN (**W**rite **O**nce, **R**ead **N**ever)
- Tape is cheap, but prone to wear after 5-10k mounts
- Slow to read
- Disk at all Tiers is becoming more like a cache
 - If the file exists **on disk** somewhere else, use that copy, only go to tape as a last resort.
 - Xrootd fallback is the vehicle for this.
 - User jobs should only access disk pool; only production jobs should access the tape system

Two Methods to achieve this: **Physical** and **Logical** Separation

Physical Separation (ASGC)



Provide Two Separate End Points

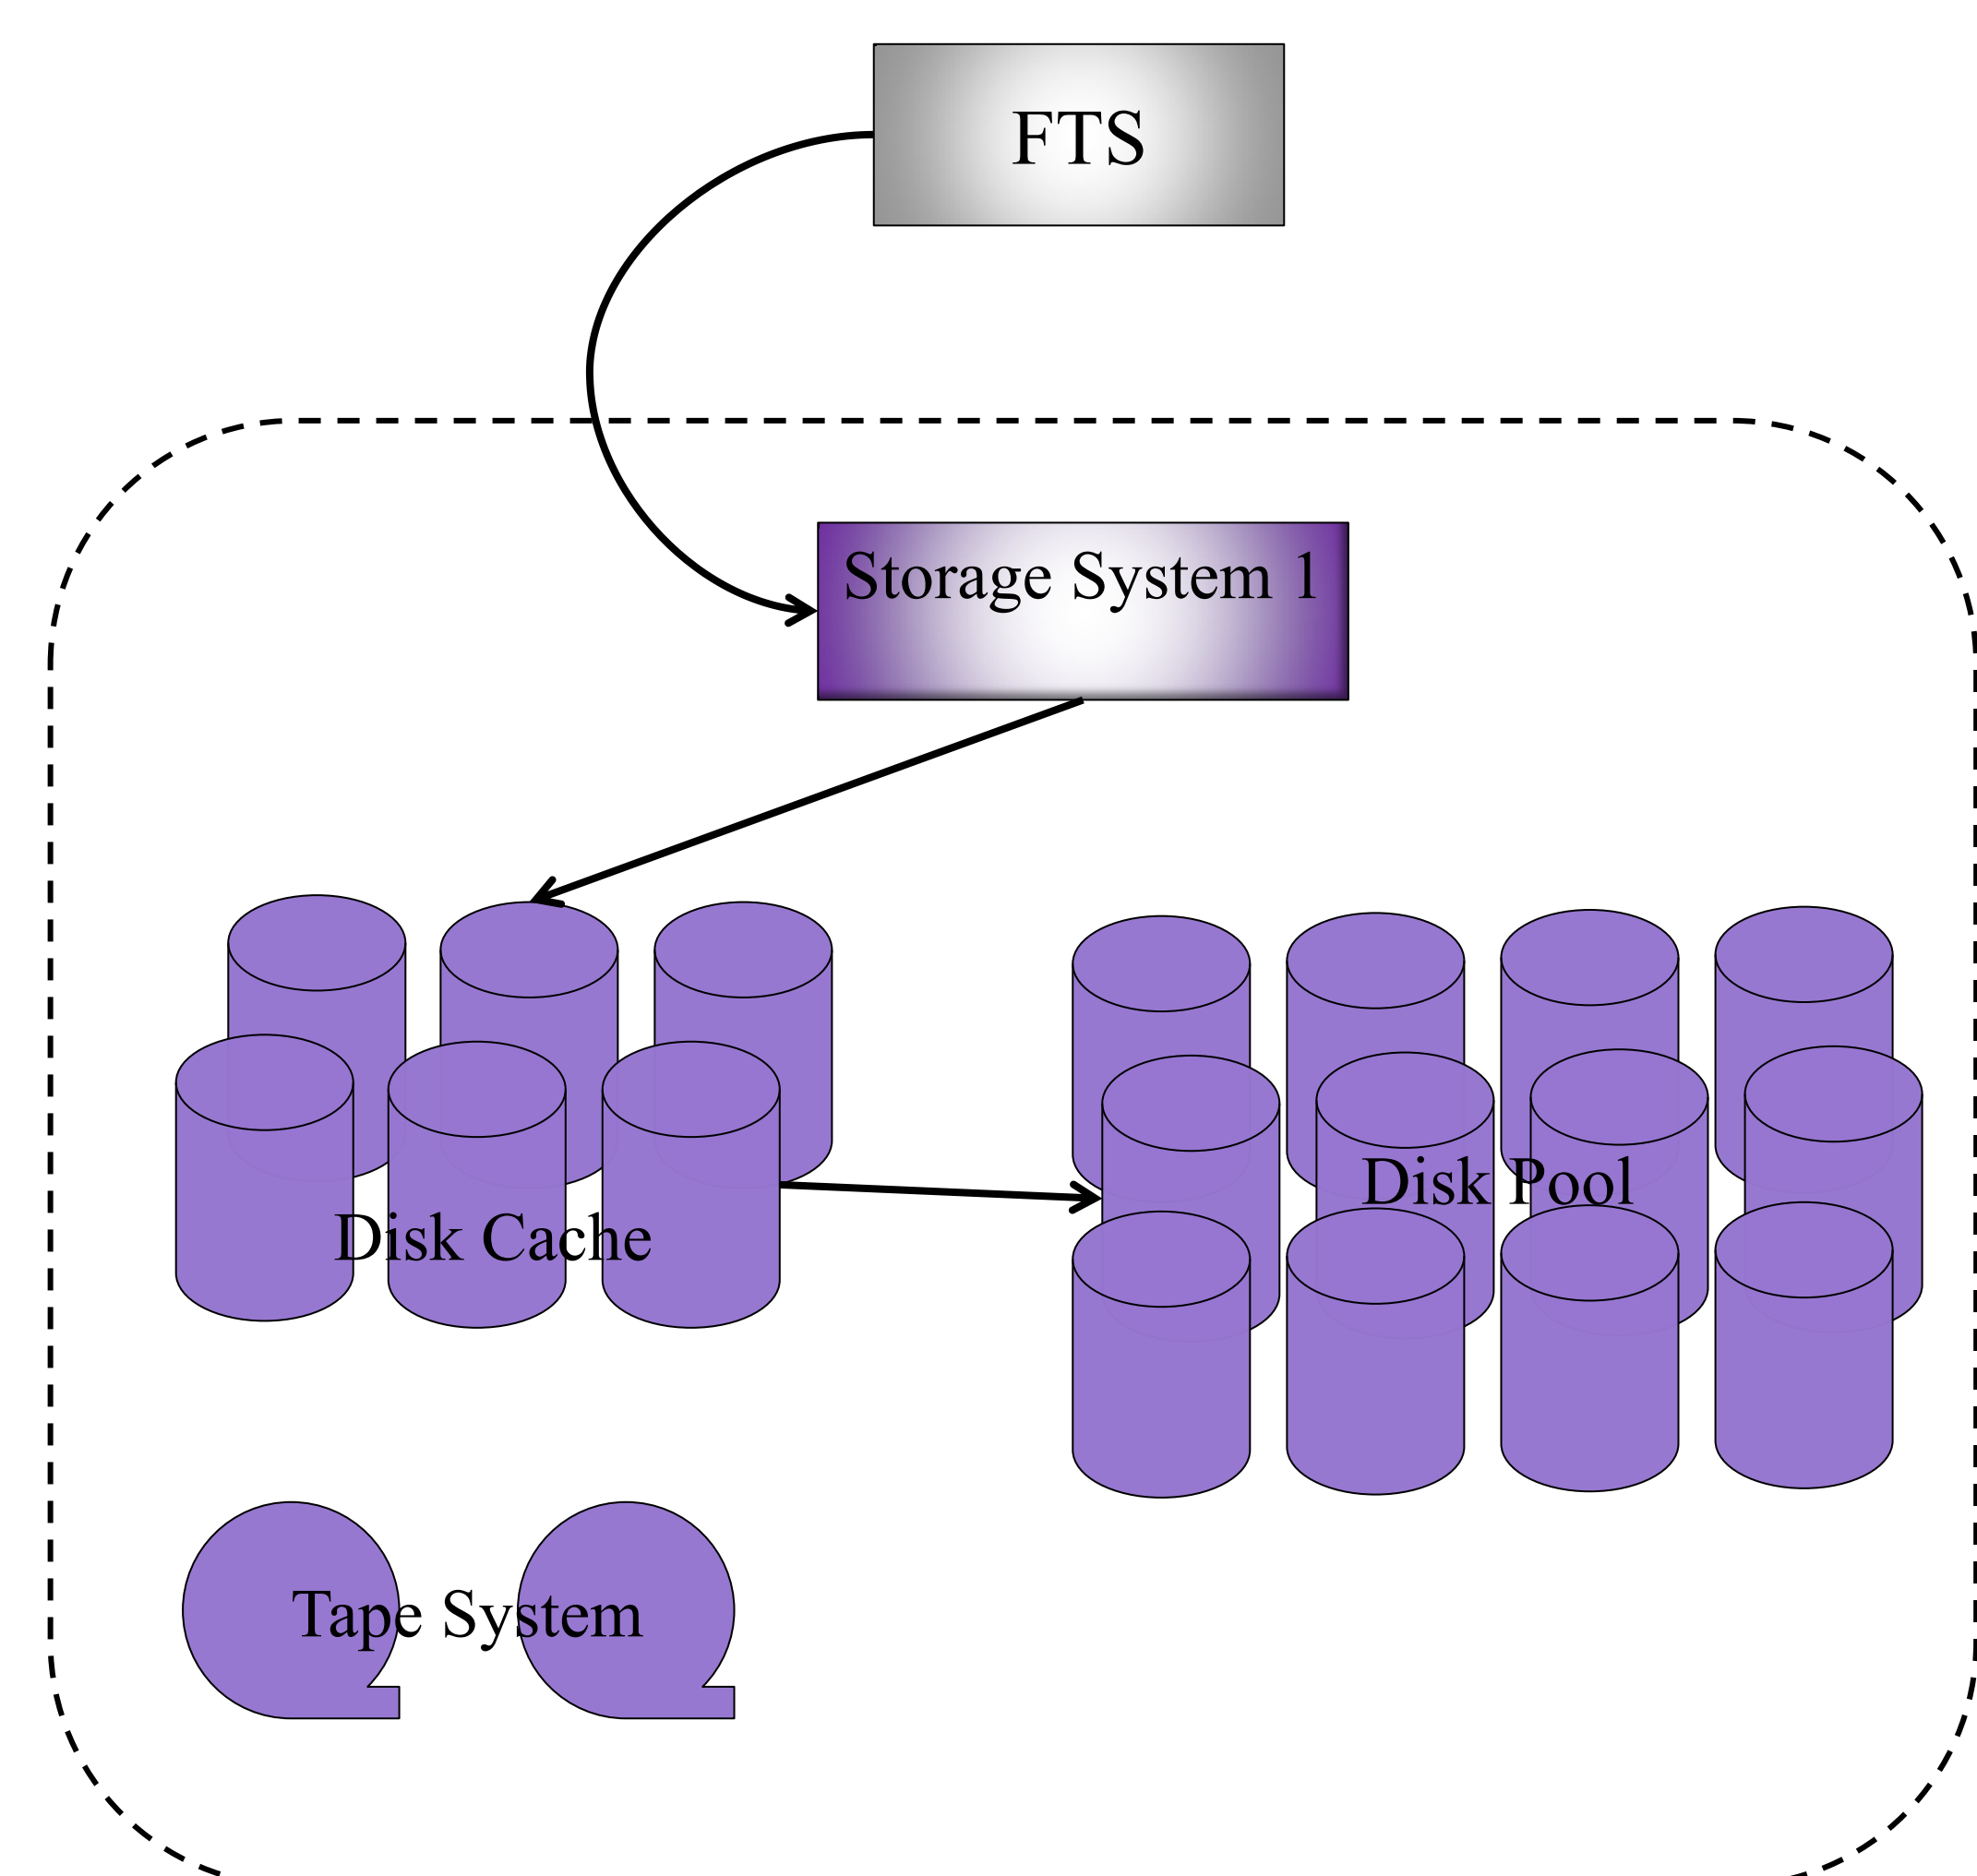
Advantages:

- True Disk/Tape Separation (no ‘backdoors’)
- Ability to select Best Solution for Each Task
- Can optimise performance separately
- Access to Tape System easily restricted to allowed users
 - Based on VOMS role
- Two systems can be upgraded independently
 - No blanket downtime

Disadvantages:

- Requires deployment of two (different) storage systems
- May require extra staff, particularly if both storage systems are different
- Not easy to move hardware between cache and pool

Logical Separation (STFC)



Provide Single End Point

Advantages:

- Fewer systems to manage (lower staff costs potentially)
- Relatively easy to move hardware between disk pool and tape cache
- Potentially lower licensing costs

Disadvantages:

- Difficult to restrict user access to disk cache
 - STFC use ACL’s and VOMS information
- Risk of ‘leakage’ if a user recalls a tape file directly onto the disk pool
- Storage system is a Single Point of Failure

FTS or *lcgutils* Mediates Transfer between tape cache and disk

- 3rd Party GridFTP is the usual protocol (but could use *xrdcp*)
- In case of logical separation, other protocols such as *file* are possible (if supported)

Compute Nodes **only** read from disk pool

- Access to tape is both **restricted** and **planned** such that typically many files are ‘staged’ in one hit to minimise tape mounts