

New User Communities

GridPP Storage Group

A perspective.

Things we have

- Things we can provide:
 - FTS
 - File Catalogues
 - Lots of Storage interfaces
 - Actual data storage
 - (Federated Authentication/Authorisation?)
 - *Expert knowledge*

WLCG Trends

(Things we may (not) have)

- SRM slowly going away in WLCG contexts.
 - informs how we provide quotas, storage protocol negotiation.
- Object Stores, HTTP, S3 becoming more popular.
 - provision of different classes of resources.
- "T2C"/"T2D" splits
 - determine where data might actually be placed.

Things VOs want?

- Key feedback:
 - Understanding the needs of each VO is important.
 - Initial interaction needs to establish needs of VO and their understanding of their needs
 - This is very important to get right.
 - We could consider organising VOs into taxa.
 - Allows us to "reuse" existing knowledge + solutions.

Taxonomies?

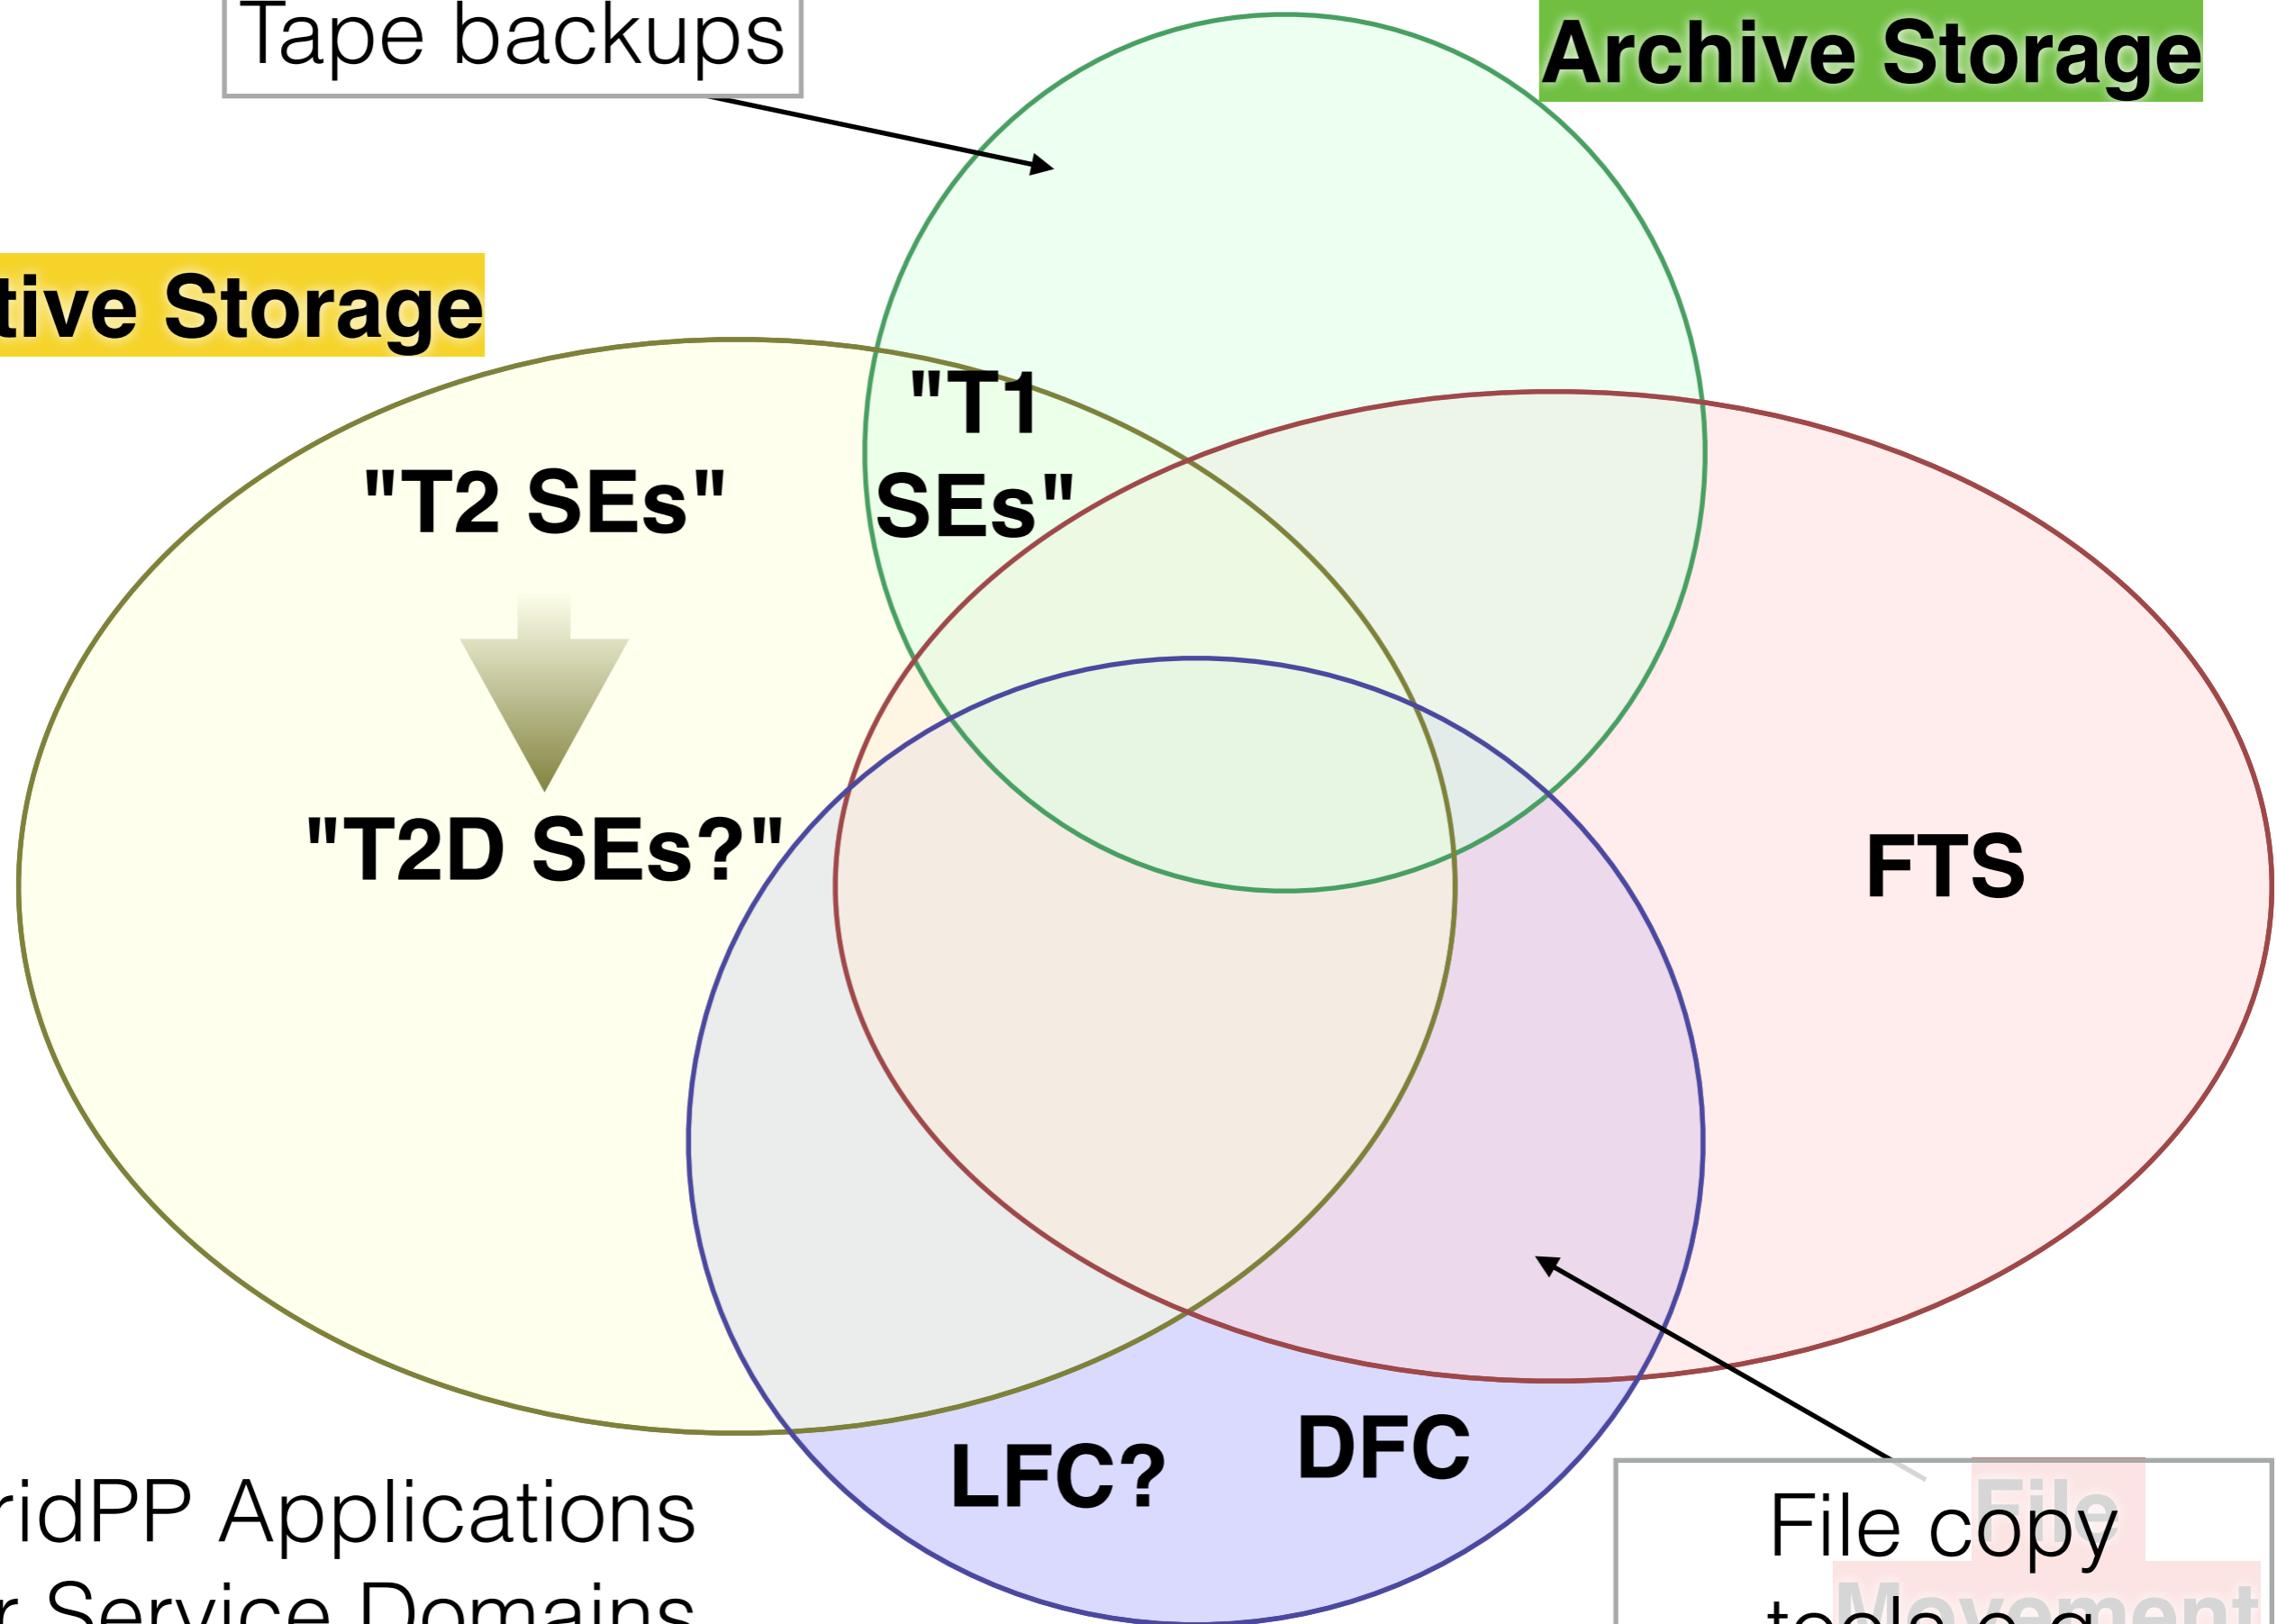
- By service types?
 - Data movement
 - Data archiving
 - Data storage
 - Data indexing
 - Data access

- By size?
 - number of sites?
 - amount of data?
 - rate of data?

Tape backups

Archive Storage

Active Storage



"T2 SEs"

"T1 SEs"

"T2D SEs?"

FTS

LFC?

DFC

File Catalogues

File copy tools e.g. Globus Connect

GridPP Applications for Service Domains

Case Study: DiRAC

- Data Movement to Archive.
- Services used: FTS3, (Castor Tape)
- Issues:
 - Authn/z (Proxy lifetime, mappings)
 - Transfer efficiency (small/zero sized files, v large files)
 - Metadata retention (need to retain Posix file ownership, paths)
 - "Archive" is actually in production use! (need changesets?)
 - FTS max queued transfer limit hit!

Case Study: DiRAC

- Support provided via "semiclosed" list.
- Resolutions:
 - Script renews proxies (make proxy long lived)
 - Archive filesystem to TAR before transfer
 - retains metadata
 - regularises file size
 - reduces # transfers to manageable value

Case Study: DiRAC

- Lessons:
 - Authn/z is still difficult, even for intelligent experts.
 - Need to be *flexible* adapting "our" solutions to different requirements.
 - Thinking outside "Grid" can provide solutions to existing problems (*tar* for archiving!)

Case Study: SNO+

- Pure Data movement
- Issues:
 - Source data from a secure administrative subdomain
- Leveraged existing experience with DiRAC VOs similar needs.
- Support via the *gridpp-support* mailing list
- Solution is very "traditional"

Case Study: LSST

- Full Stack Data/Compute
- Services used: Dirac File Catalogue, Dirac File Transfer, SEs (DPM), (FTS)
- Issues:
 - file migration into DFC from early work
 - file migration from NERSC (bare gridftp?)
 - file migration from BNL (https w/ passwd)
 - Effort available for Data Management?

Case Study: LSST

- Methods:
 - File migration into DFC via "copying" into the same storage (but via DFC)
 - Manual copying from passwd auth'd endpoints.
 - DFC automation of file movement for internal file handling (no copies outside of Dirac!)
 - Ganga for automation + orchestration

Case Study: LSST

- Lessons:
 - If you might need to use a file catalogue, use one from start (esp. with DFC)
 - Authn/Authz interactions with external repositories need to be improved / add complexity.
 - gfal2 and other frameworks should help with this. (?)

Data Storage (Local)

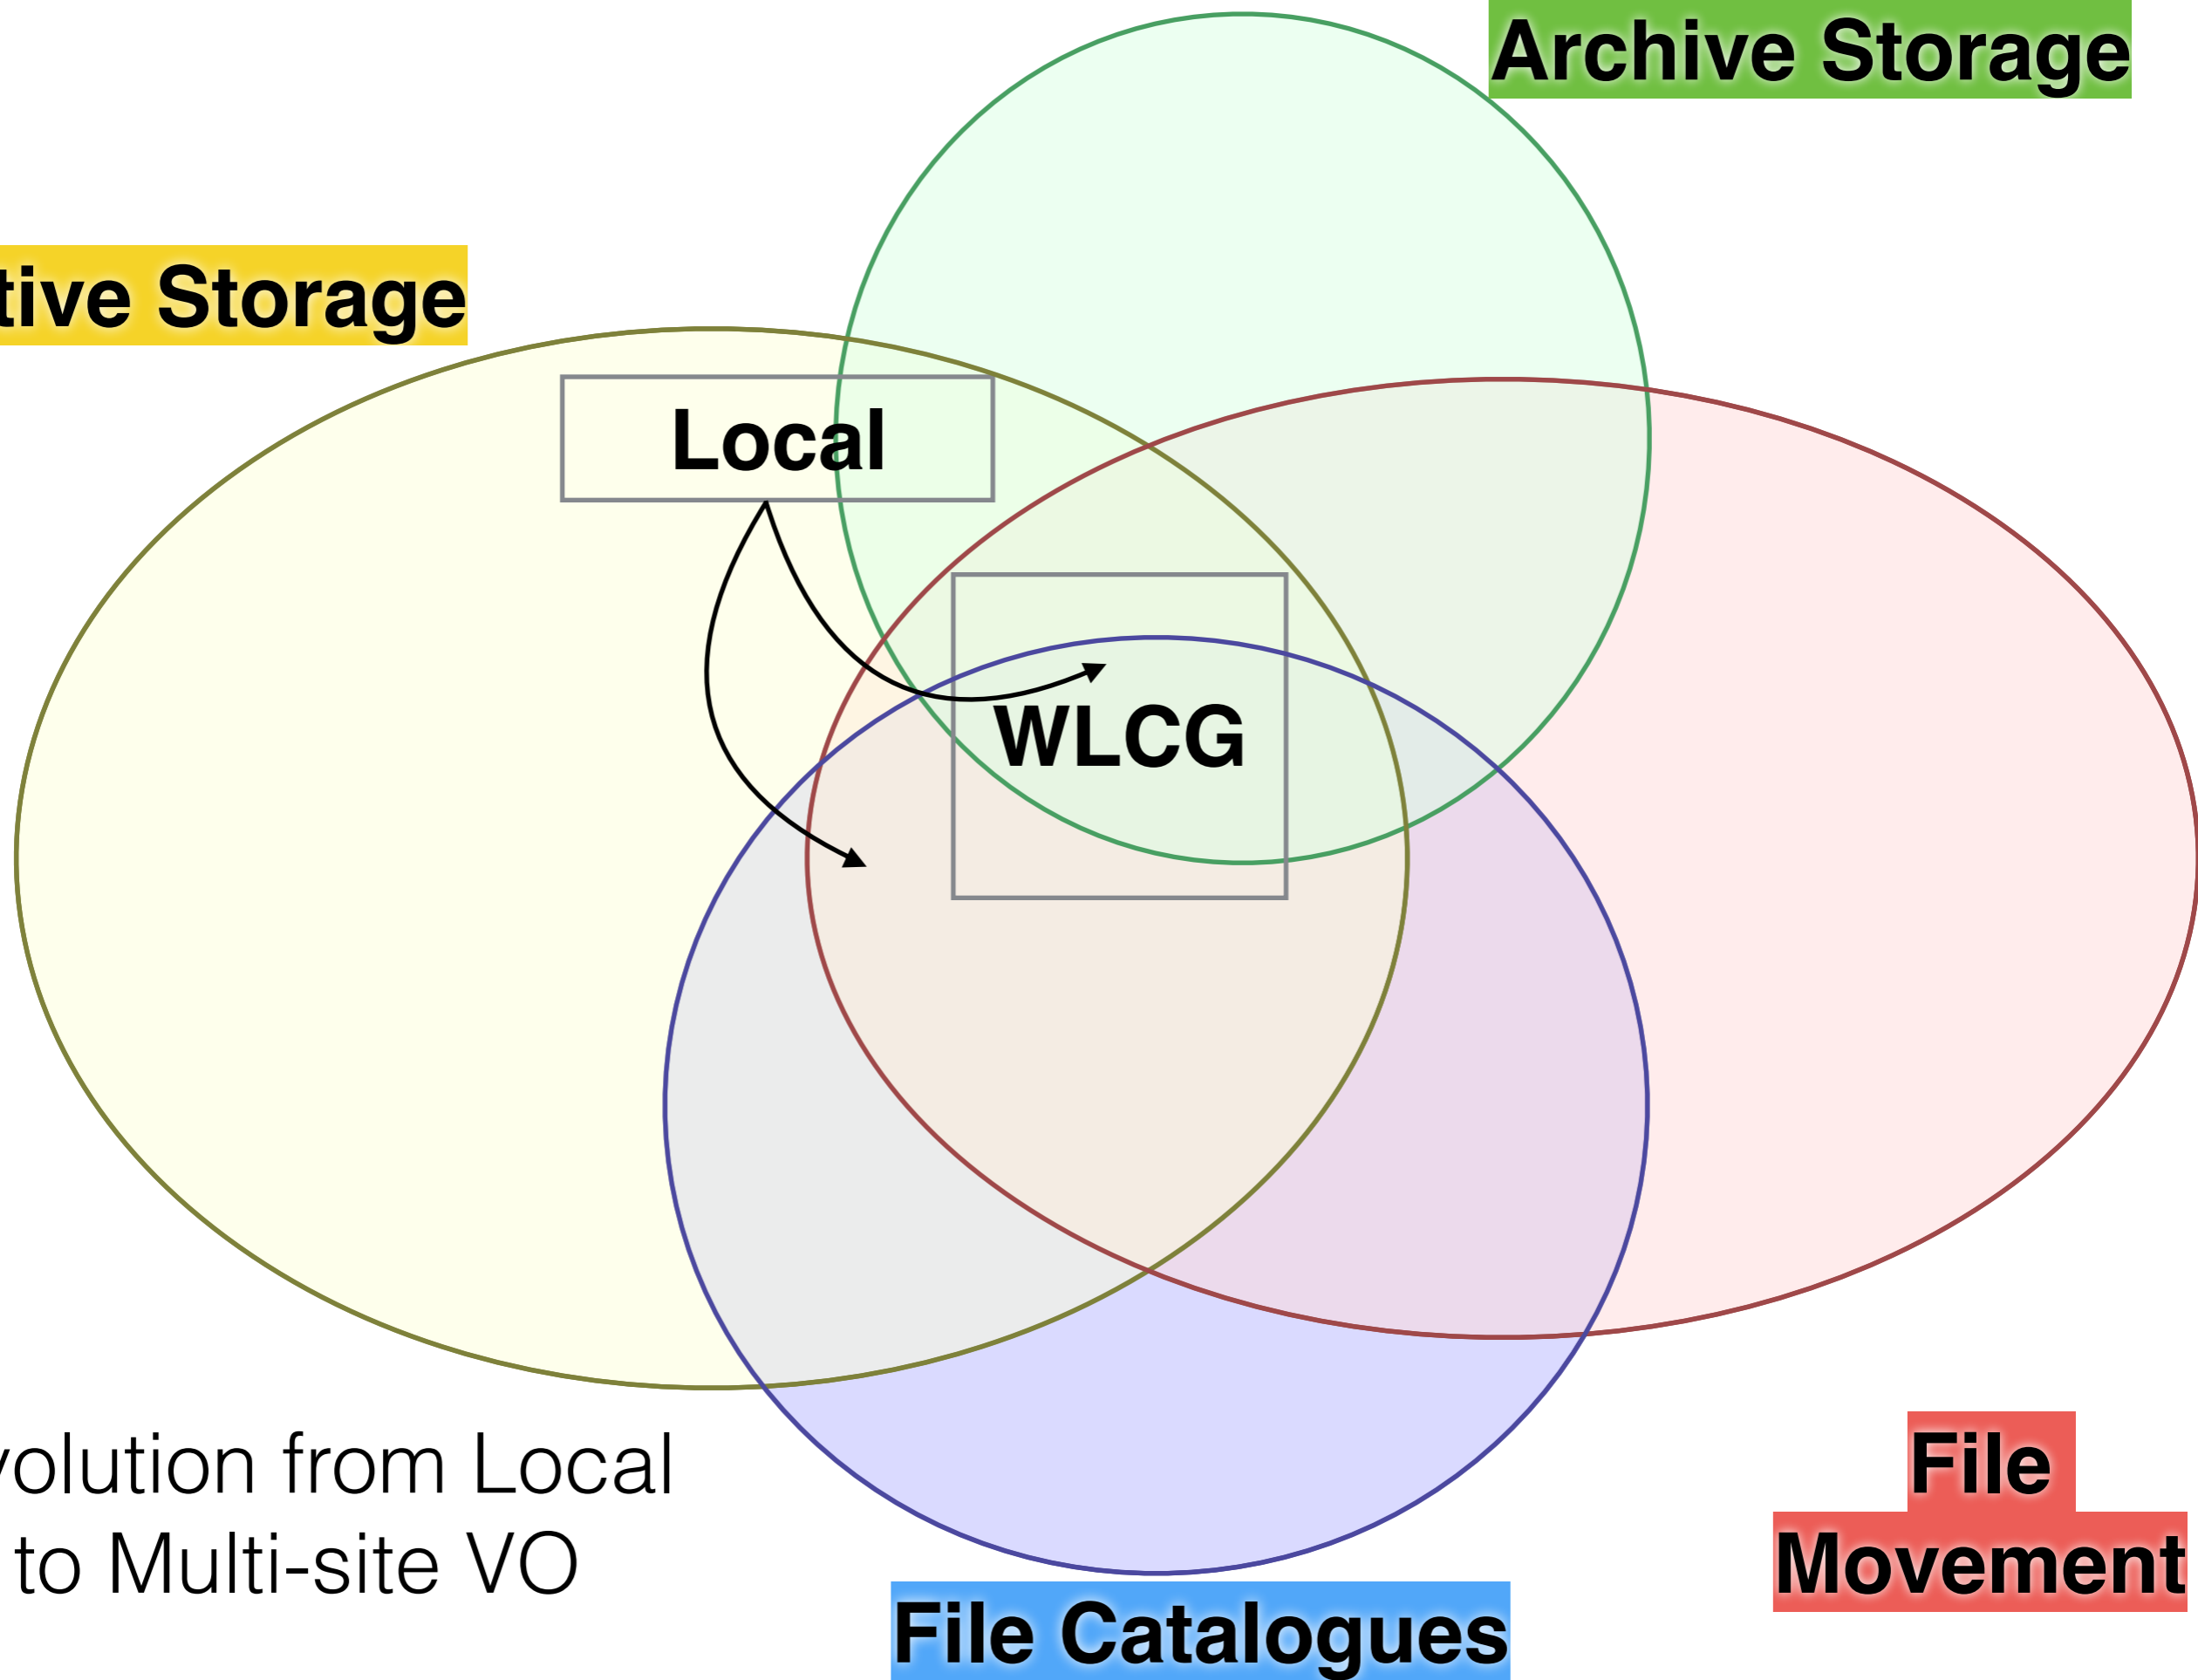
- Do interfaces matter?
 - "Modern": S3, HTTP
 - GridFTP - still default
 - local posix?
- Local users don't even need "Grid" Storage models, if never planning on moving to "distributed".
- Landslides uses CouchDB!
- FTS support, DFC/LFC support limit our agility.
- With move to T2Ds, where is "local"?

Data Storage Shares

- The Tragedy of the Commons is a classic issue for growing VOs.
- "SRM solution": Spacetokens?
 - site administered (not VO!)
 - hard to have "default" for some SEs
 - increasingly deprecated for WLCG VOs
- This has never really been properly addressed.
- WebDAV "Quotas" coming... (but orthogonal to spacetokens)

Active Storage

Archive Storage



Local

WLCG

Evolution from Local
to Multi-site VO

File Catalogues

**File
Movement**

Authentication (aside)

- Moving from local -> "Grid" operation requires federation of Auth.
- Trust model / Identity sources
 - WLCG Grid Certificates [local to us]
 - Edugain - see *AARC* etc
 - *INDIGO* (IdS ~~-translation-~~ -> OpenID Connect/OAuth2)
- Usability?
 - X509 is not bad, but our tools for using it are.
 - Consider what everyone else is using.

File Catalogues

- Core element added in transition from local -> distributed
- GridPP DIRAC '**DFC**' (Imperial)
 - Need opt-in at start
 - LSST's "local" EDCF data needed to be migrated into namespaces compatible with DFC
 - Need to move files with DIRAC tools (which talk to FTS)
 - Maintained by Imperial

File Catalogues

- Core element added in transition from local -> Grid
- **LFC**
 - Deprecated for WLCG
 - (But adopted by Belle 2!)
 - Does not need initial opt-in (supports post-registration in any path)
 - Supported by FTS directly
 - How many instances now exist?

"Best Practices"?

- For New VOs:
 - Local VOs *who will remain local* - can have any solution suitable to them?
 - VOs which will go Distributed:
 - should start using an integrated solution(DIRAC?) from the start, if possible
 - should be offered "modern" protocols, access in preference to legacy (despite experience gap)
- *We should try to be more flexible in the domain of Authn/Authz*
- *We should try to be more flexible in decomposing our solutions to solve only the problems each VO needs solved.*

Parting Thoughts

- GridPP has a range of tools to solve different problems in Data.
- We also have a range of requirements from different customers.
- Flexibility is important in supporting these different requirements.
- Building up a knowledge base of effective solutions will make us more effective as time goes on.
- There is a tension between providing *fast* solutions (with existing models) and *future-proof* solutions (within projected evolution).

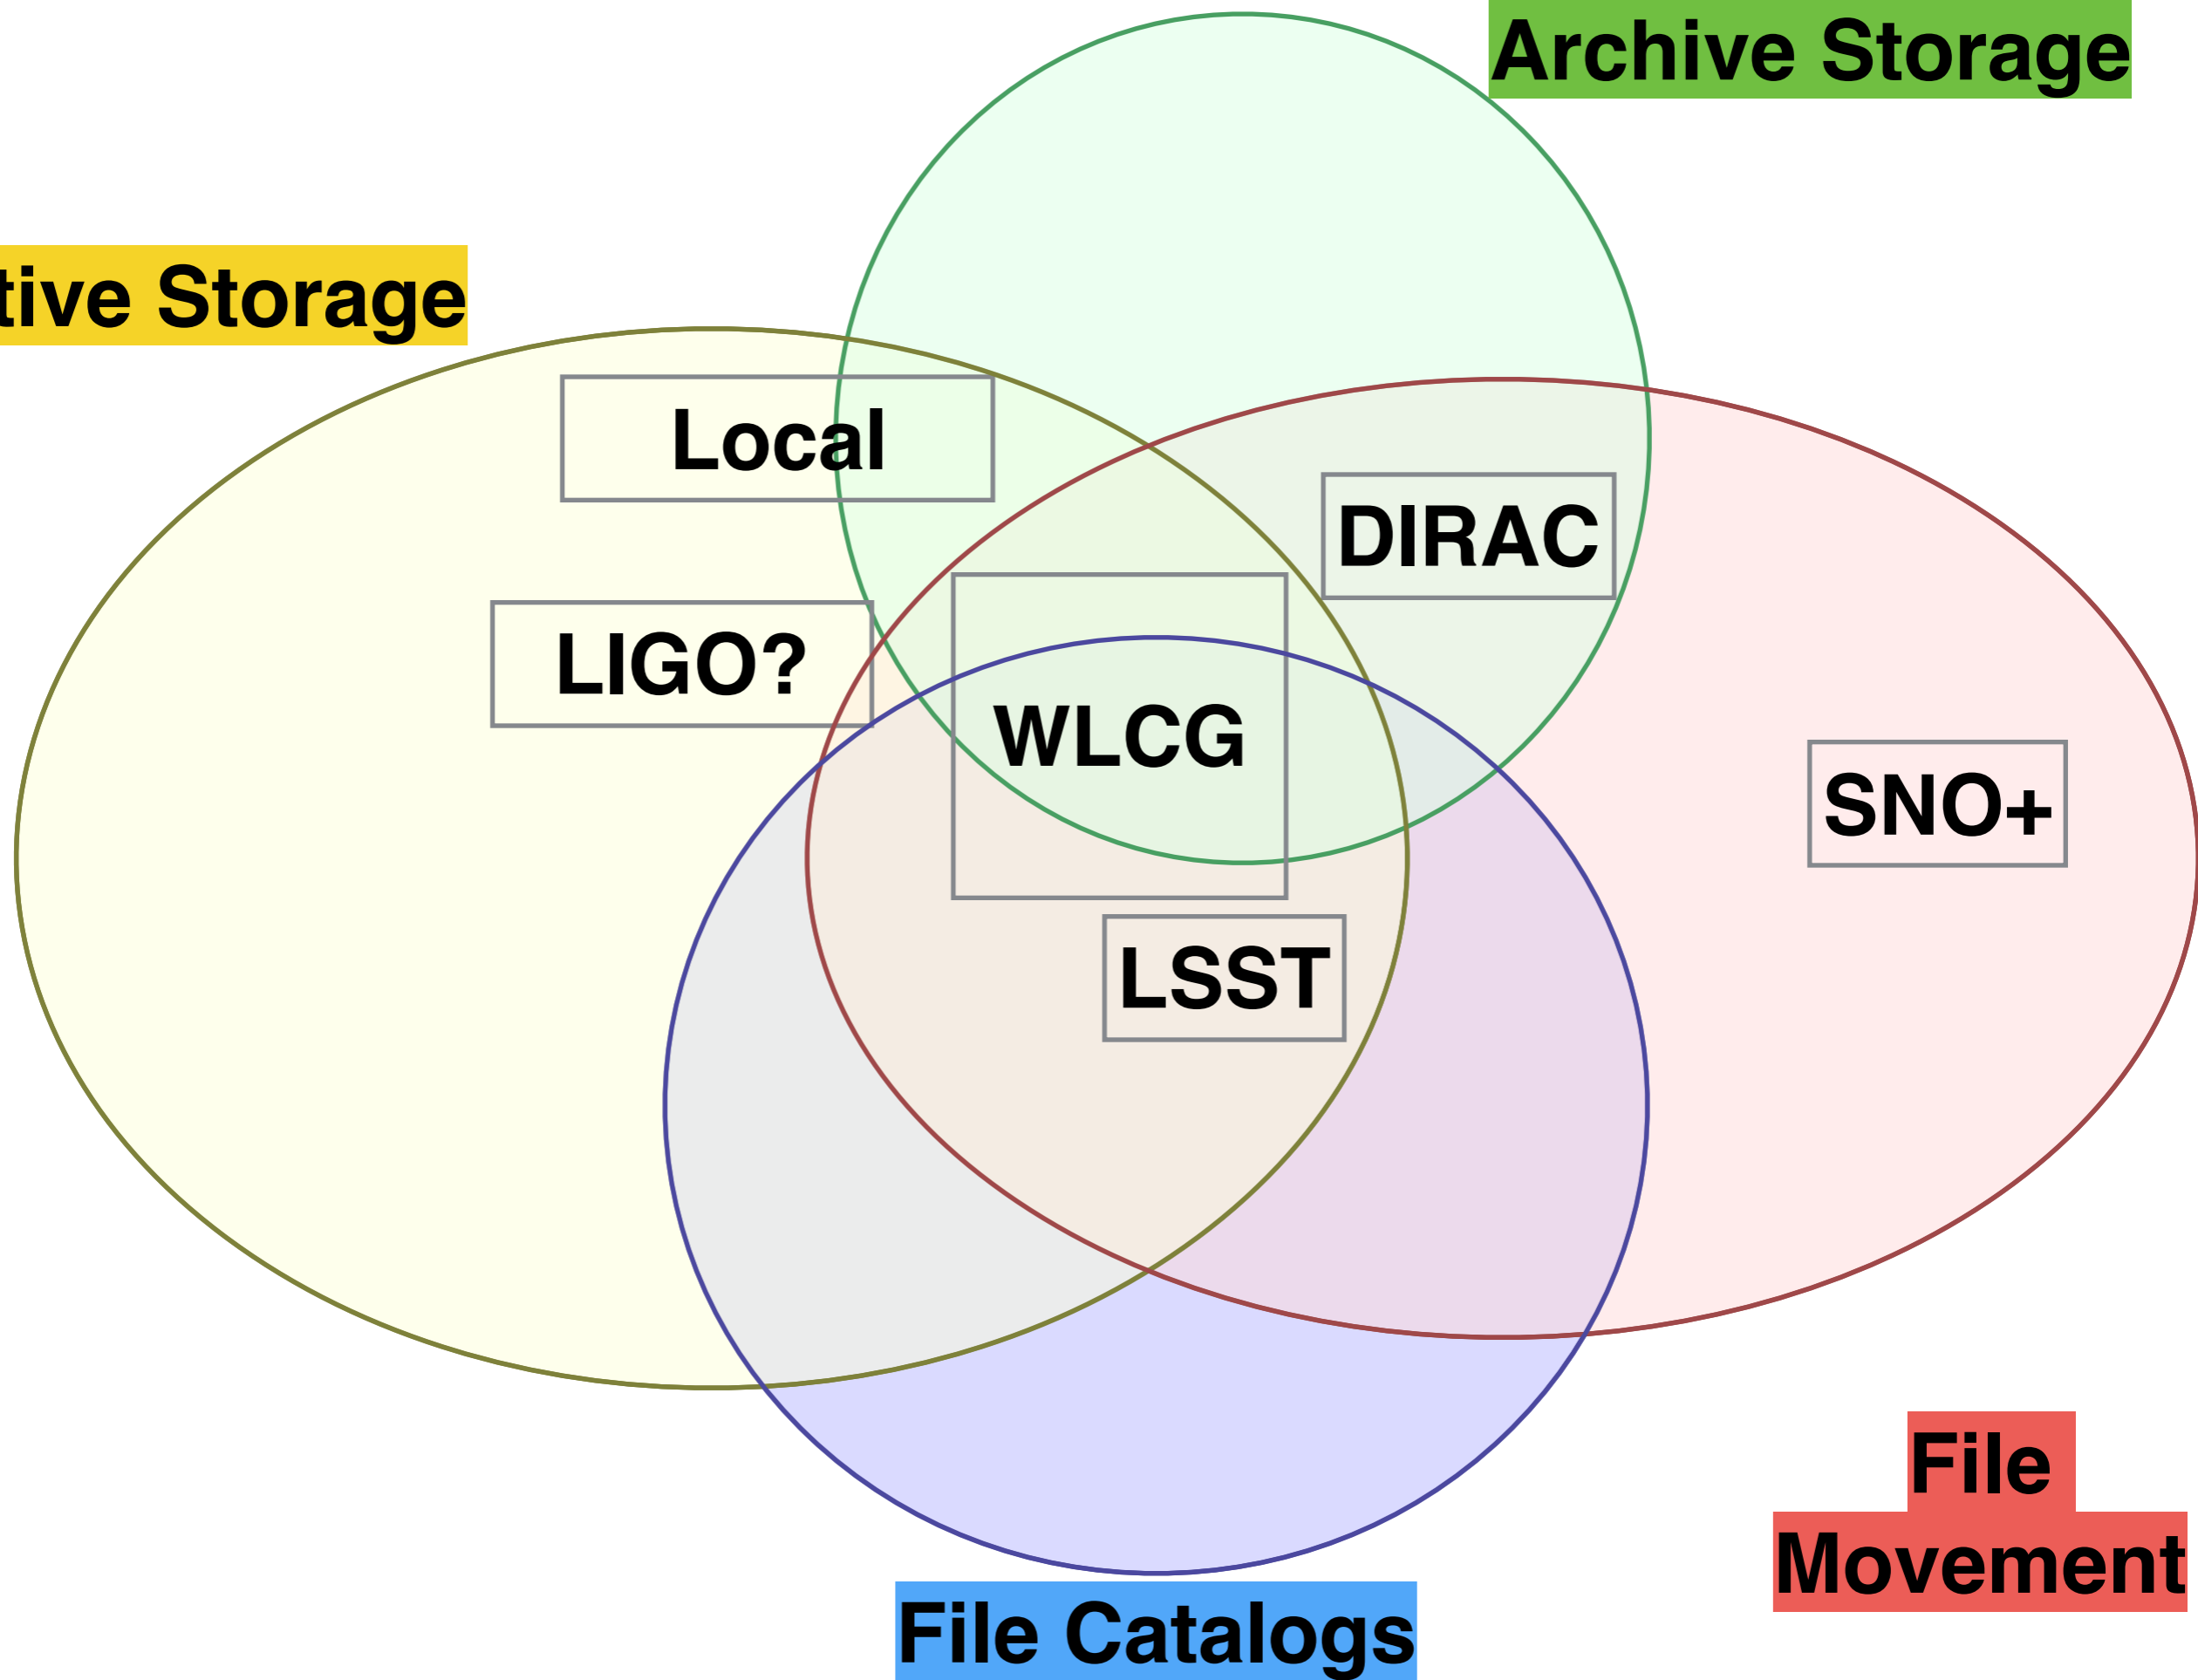
Backup

When does a new user/VO need a spacetoken?
(spacetokens are the correct way?)

DIRAC - subVO per DIRAC site.

Active Storage

Archive Storage



File Catalogs

File Movement