



# Spaces, tokens and descriptions

- Storage classes
- What is a space?
- Space reservation
- Spaces and quotas
- Spaces and paths
- Space usage



# Storage classes

- A storage class is a quality of storage defined by the Retention Policy and Access Latency
- The WLCG SRM v2.2 [MoU](#) defines 3 cases:
  - Custodial x Nearline → "Tape1Disk0"
  - Custodial x Online → "Tape1Disk1"
  - Replica x Online → "Tape0Disk1"
- TapeN → N copies guaranteed on tape
  - Or other high-quality media
    - Tape1/Custodial → "Do not lose this data!"
    - Tape0/Replica → "No disaster if this data is lost." (a custodial copy may be elsewhere)
- DiskM → M copies guaranteed on disk
  - Disk0 managed by system, Disk1 managed by VO
- WLCG SURLs are permanent, can only be removed by the user
  - Volatile files do not imply Tape0Disk0 (!)



# What is a space?

- A space is a chunk of disk storage with given:
  - Retention Policy
    - Custodial  $\leftrightarrow$  Replica
  - Access Latency
    - Nearline  $\leftrightarrow$  Online
  - Size
  - Lifetime
    - Infinite for static spaces
  - Access Control List
  - Use case(s)
- Nearline access latency applies when the disk is a system-managed cache in front of a tape system back-end
  - A file need not always have a copy on disk for immediate use
    - It may need to be temporarily brought online first
- Tape storage (when applicable) considered infinite as far as space reservation is concerned



# Space reservation

- srmReserveSpace allows space to be reserved for use by a DN or a VOMS role/group
  - Space is taken from the unnamed “basin” of resources available to the VO
  - An opaque, site-dependent space token is returned
  - A space token user description can be associated with the space
    - Allows for site-independent identification of spaces
- TapeNDiskM storage classes only require static reservations by admins
  - For a certain VOMS role/group per space
  - Space token user description indicates intended use
    - LHCb\_RAW, LHCb\_RDST, ...
    - Can be published in information system
- Dynamic (temporary) reservations by users can be used to guarantee room for job output files
  - Requires significant portion of “basin” not to be allocated statically
  - Not supported by CASTOR



# Spaces and quotas

- In SRM v2.2 one cannot reserve a dynamic space within another space
  - May become possible in future extension
- Static spaces are for well-behaved activities
  - Typically controlled by production managers of big VOs
  - Not accessible (for writing) by ordinary users
  - Overbooking normally avoided
- Dynamic spaces allow for some protection of user jobs
  - But CASTOR does not support it
  - Overbooking can only be prevented with quotas
- SRM v2.2 does not support quotas
  - Implementations may support them through proprietary APIs



# Spaces and paths

- In SRM v1.1 the quality of storage is defined by the path
- In SRM v2.2 spaces and paths are orthogonal
  - A directory may contain files of different storage classes
  - The logical name space need not be driven by storage considerations
    - Name space indicates type of data, run number etc.
    - A site will have custodial responsibility for a subset of the VO's files
    - Files may have been replicated from other sites into the same directories
  - Some implementations can still offer better quality of storage when the name space is used as before
    - E.g. fine-grained tape set selection
- A file may be moved into a different space
  - The SURL remains the same
  - srmChangeSpaceForFiles can be used to change the storage class
    - Only Tape1Disk0 ↔ Tape1Disk1 is supported



# Space usage

- WLCG MoU: token can only be specified when file is written  

```
lcg-cr --vo lhcb file:/... -d se.domain --st LHCb_RDST  
glite-transfer-submit -t LHCb_RAW ...
```
- The room in an online space shrinks by the file size
  - The room can only be recovered by deleting the file
- The room in a nearline space temporarily shrinks by the file size
  - The room will be recovered when the file has been written to tape and the copy on disk is not in use
- Some implementations allow a token to be specified on reads
  - Get a copy of the file into a preferred space
  - By default the SE will decide where to stage the copy
    - Depending on the storage class, user DN, VOMS attributes, client IP address, pool load, ...