

File Transfer Services and Storage Management

ATLAS

Miguel Branco

Alessandro de Salvo



March 4, 2005

DQ



- Data Management for file-resident data
- “Grid-only”
 - Security
 - Transport protocols
 - Storage interfaces
 - Replica Catalogs
 - Information System
- Uniform Interface-layer
 - 3 GRIDS: LCG, NorduGrid, Open Science Grid
 - Client tools
 - For end-users
 - For production managers
- DQ implements VO (ATLAS) policies for handling file-resident data

File Transfer Service



- Reliable File Transfer Service:
 - Part of DQ
 - Backend MySQL
 - A set of services:
 - For agents
 - For clients
 - A set of agents to perform transfers
 - Auxiliary scripts to monitor transfers
- Current ATLAS policy:
 - Based on priorities
 - For production managers & end-users (different areas, statistical distribution for choosing requests, etc..)
- Accessible from end-user client tools and “power users”
- 4 TB/day, >15000/files

File Transfer Service

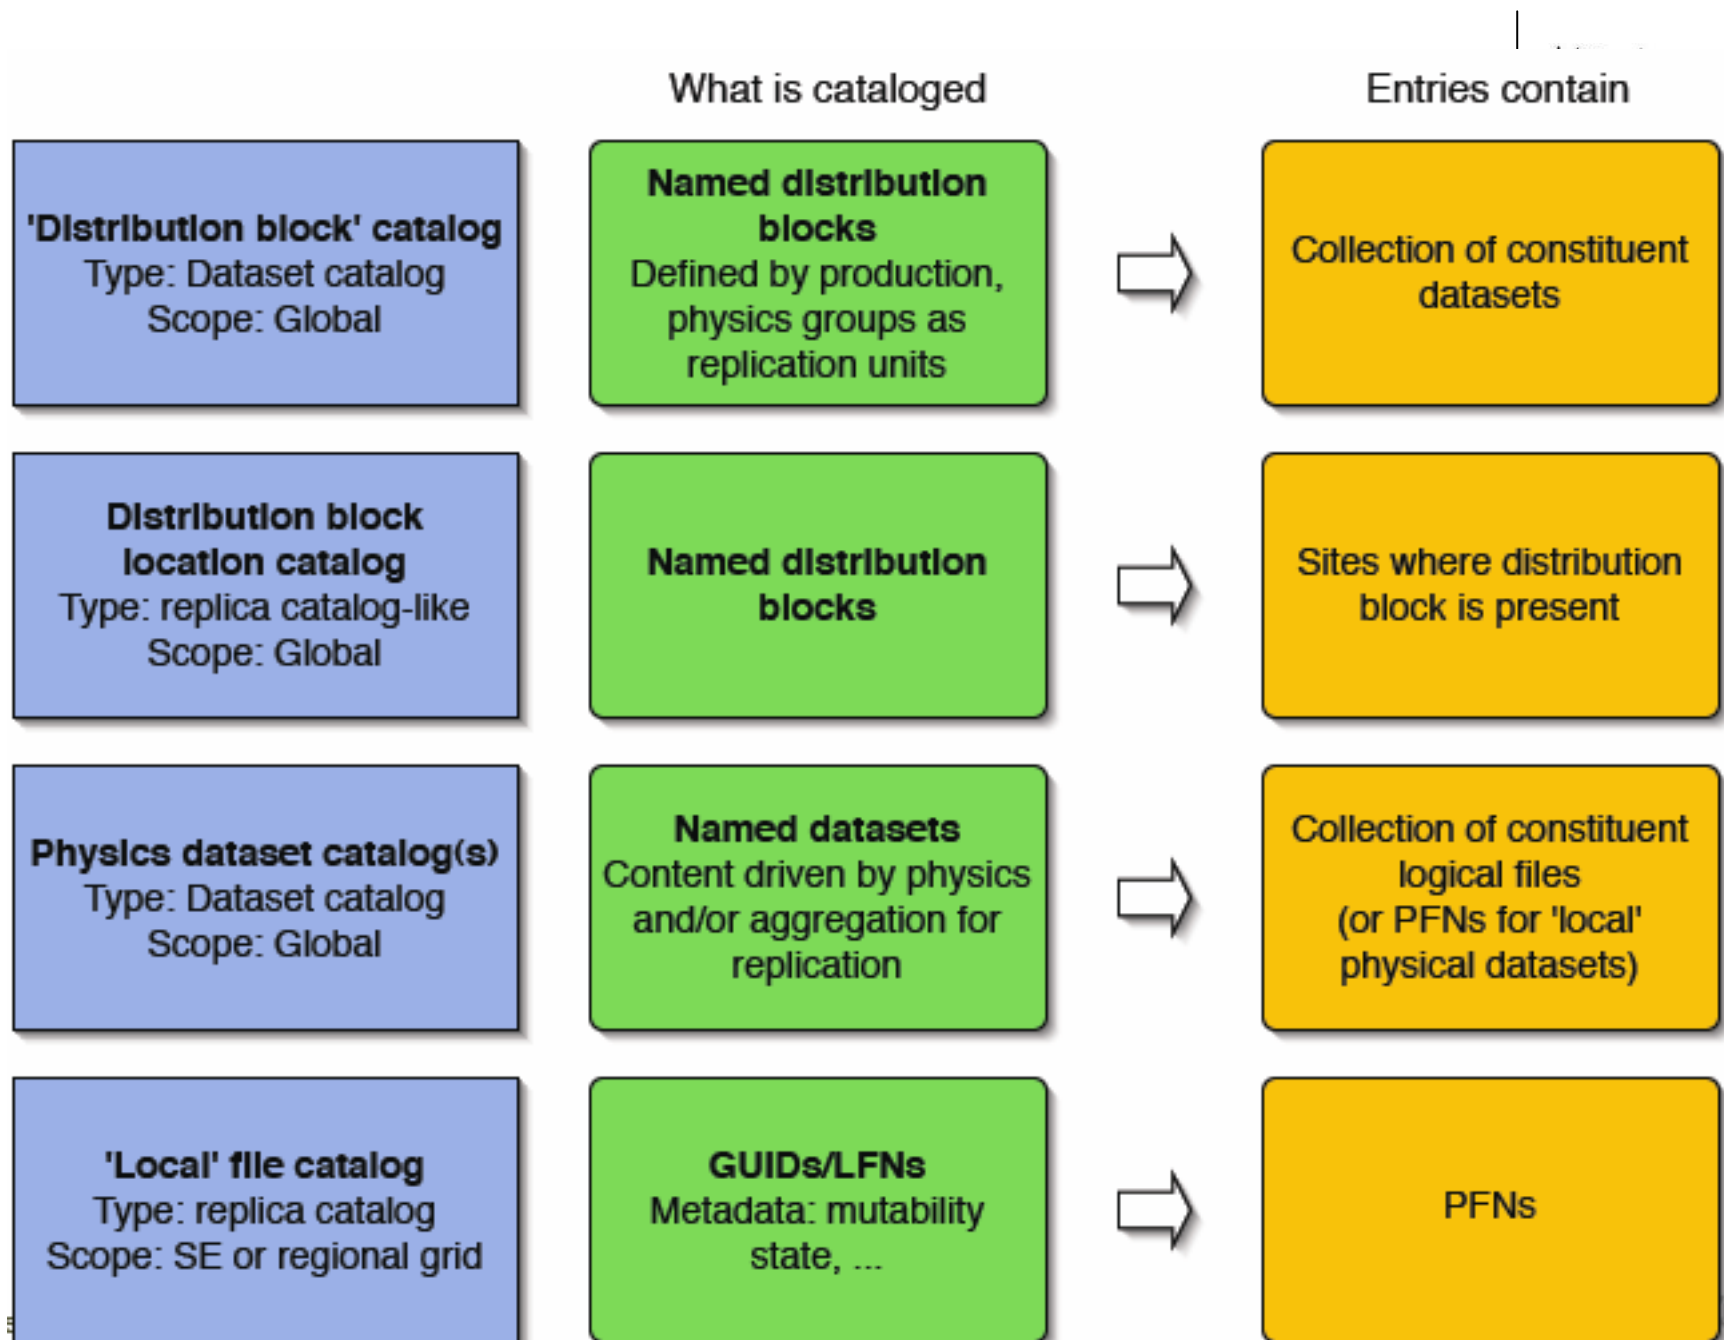


- Short term focus is adding Security:
 - We need to take advantage of identifying users (authentication and authorization)...
 - Permissions, user rights - how will ATLAS define these policies for accessing data? - ***not fully defined!***
 - User quotas & accounting - simple model using claim mechanism (shown next)
 - ATLAS needs roles in the VO to assign group/user quotas
 - Grid middleware must support permissions/user rights but DQ will always have to superimpose the restrictions defined by ATLAS policies

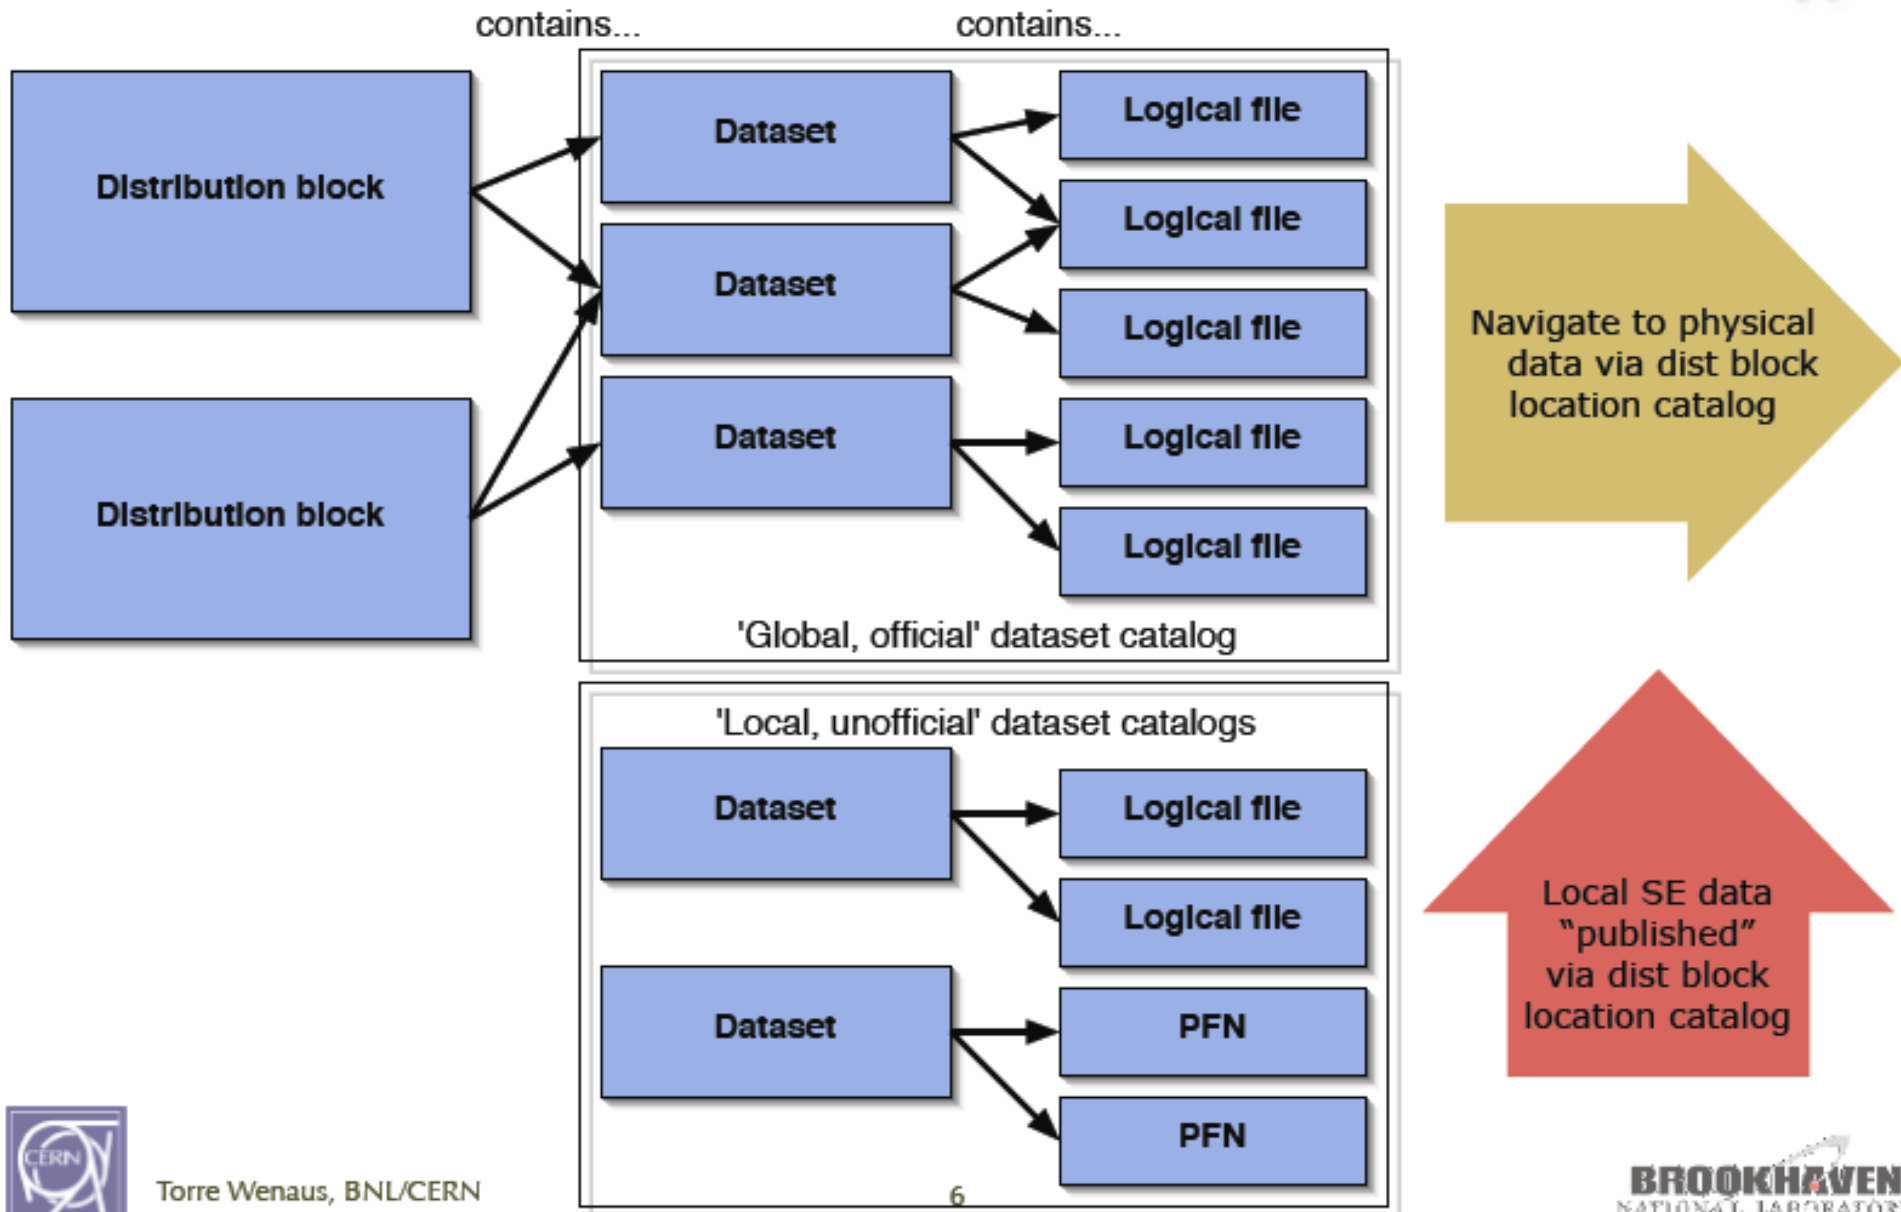
Future Plans for DQ



- □□ *Redefinition of DQ as part of ATLAS Data Management effort*
- **Automatic** filling of transfer requests
 - Users insert high-level transfer definitions:
 - “I want to move this block of data/dataset/set of partitions to this site”
 - .. Or ..
 - “I want this dataset to be available, replicated twice only, on any site which is Tier1”
 - Example above could in principle be done now with reliable file transfer but given 10 000 files per dataset, it’s not manageable.. And doesn’t take into account current status of data distribution per dataset; doesn’t know about Tiers, user roles, ...
 - **High-level transfer definitions are meant to be a “representation” of the ATLAS computing model for data flow:**
 - “ a) the derived AOD is archived via the CERN analysis facility and an instance is shipped to **each** of the external Tier-1s (a full copy at each Tier-1); “
 - Reliable infrastructure must be in place to move data, but monitoring will eventually become more important



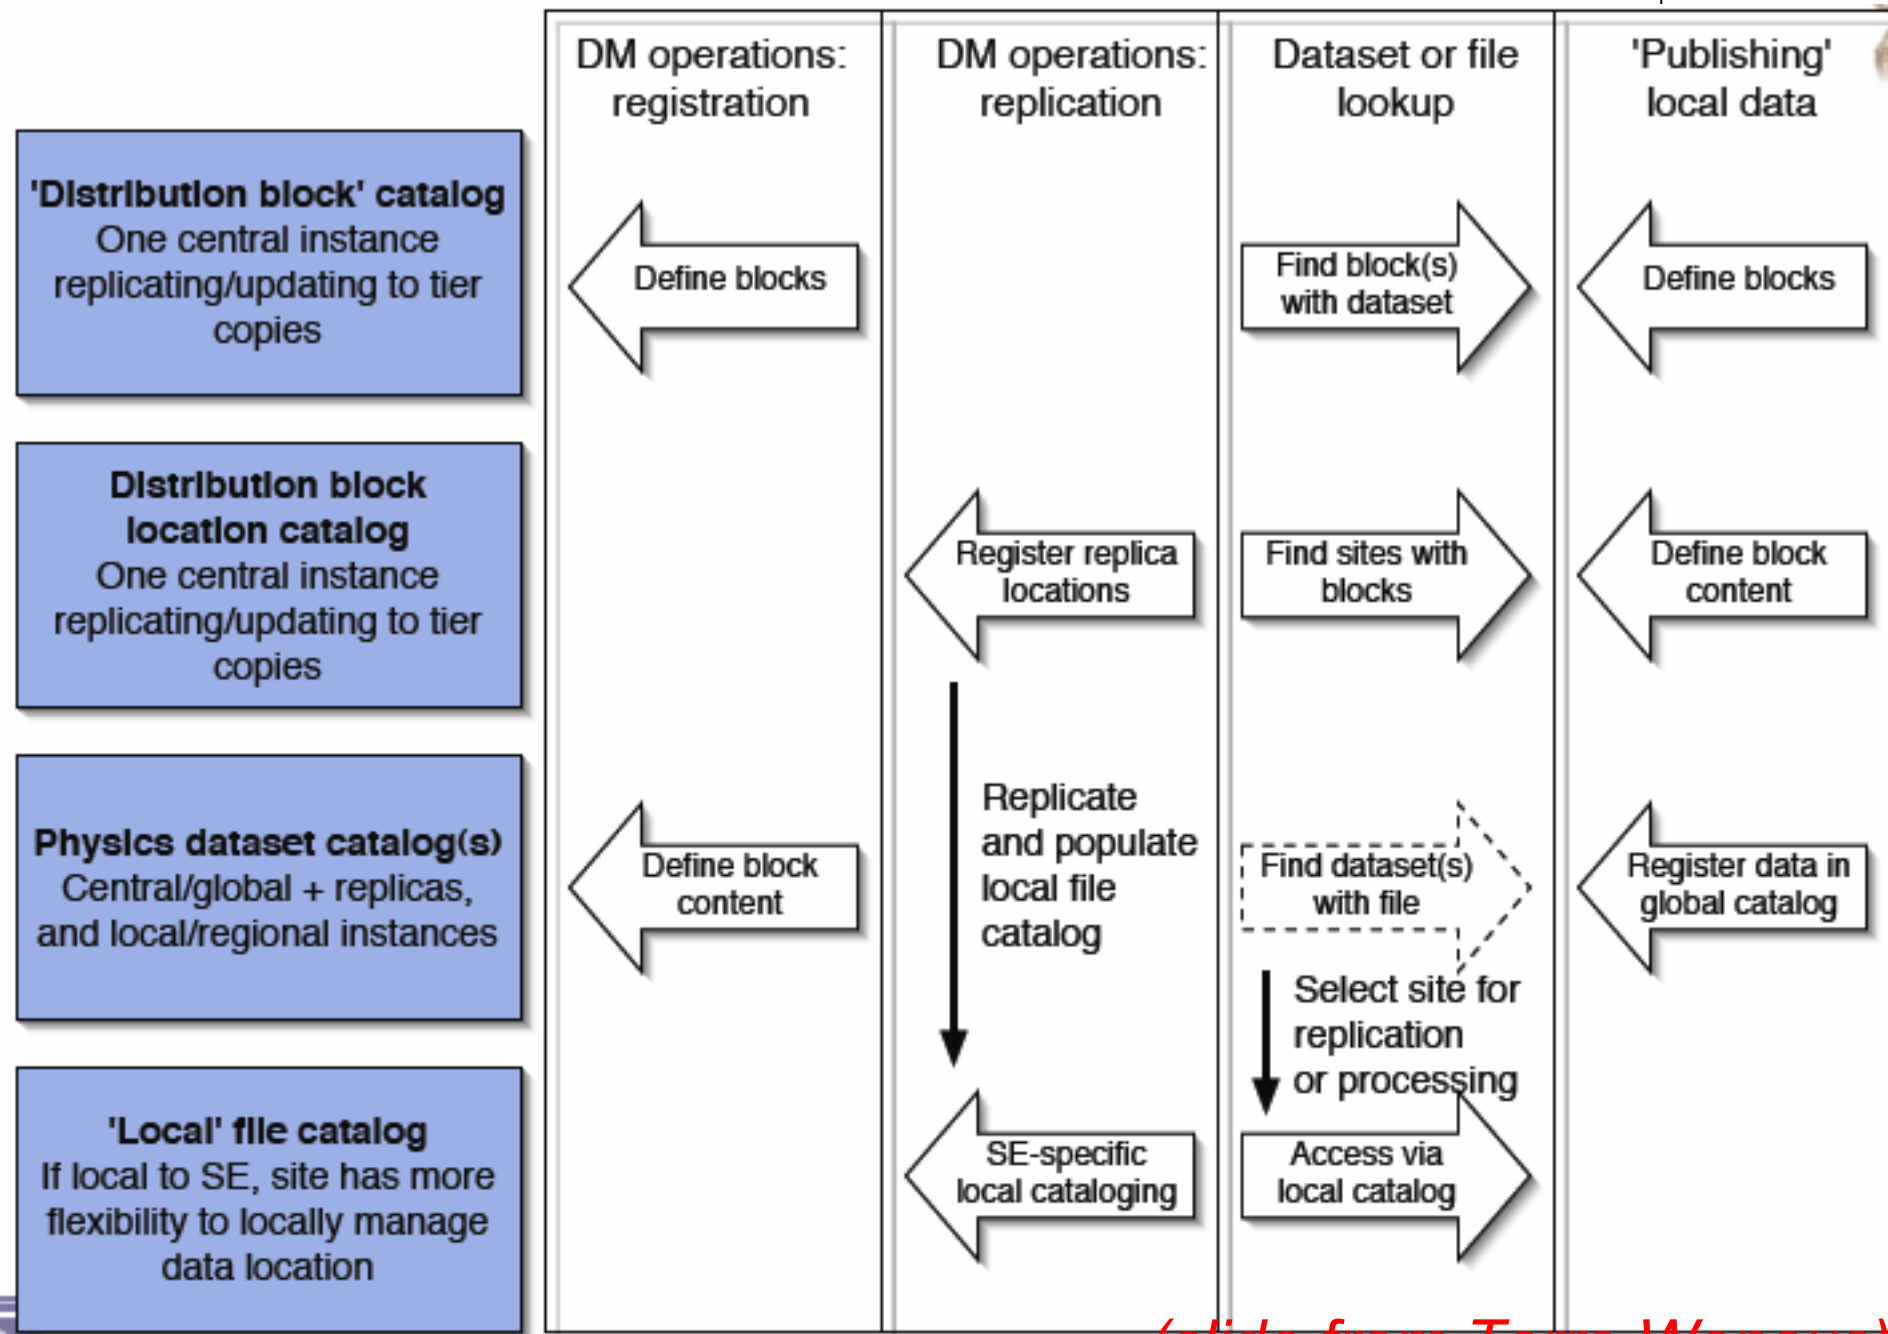
(slide from Torre Wenaus)



Torre Wenaus, BNL/CERN



(slide from Torre Wenaus)



(slide from Torre Wenaus)

Lifetimes



- Files should be assigned lifetimes
 - To avoid cluttering catalogs and filling disks
 - Files at all scopes
 - Logical (VO) file
 - Archived file replica
 - Staged file
 - Multiple users will reserve a given file
 - So a single lifetime will not work
 - Introduce claims...

(slide from David Adams)

Claims



- Usage
 - User request with lifetime is assigned a claim
 - Put, get, copy, ...
 - User may add claims on existing files
 - Logical, archived or site
 - Claim owner may (should) release claim when done
 - Claim owner may extend lifetime of claim
- Behavior
 - Each claim has an expiration time (now plus lifetime)
 - Claim is active until released or expired
 - File may have multiple active claims
 - File should not be deleted while claim is active
- Accounting
 - Claims provide mechanism for accounting

(slide from David Adams)

Other Issues



- Storage and Space management
 - Security accessing catalogs, disks and mapping to user accounts:
 - 1st phase: ATLAS-wide (granted or refused)
 - 2nd: Few user account (per PWG and/or per Production role (??))
 - 3rd: Fine-grained down to user level? Maybe not.. Few user accounts (group accounts) might be enough
 - Access to the File Transfer service: using same roles/groups/users mapping
 - Implement Claims mechanism
 - Initial (*short term*) implementation: on clients only
 - From DC-2 experience: **important use case missing**: not easy to produce overview of catalogs; they are too slow
 - Think of alternatives.. Every option being considered
- Usage of GridFTP and SRM for file movement
 - Current infrastructure in place (must be improved)
 - Grid m/w should provide metrics on network topology
 - Upgrade file movement mechanism to use service challenges infrastructure?