

# The AliEn File Catalogue

**Jamboree on Evolution of WLCG Data &  
Storage Management**

17-June-2010


Pablo Saiz

- What is it?
- What the user sees
- What the user does not see
- Current developments
- Conclusions

- Global Unique name space
  - Mapping from LFN to PFN
- UNIX-like file system interface, ACLs and commands (+ some Grid-related goodies)
- Powerful metadata catalogue
- Automatic SE selection
- Integrated quota system
- Multiple storage protocols: xrootd, torrent, srm, file
- Physical file archival (helps with MSS)
- Used by several experiments



15 June 2010



CBM, MAGIC V

*Pablo Saiz - DAaM Jamboree*

- Same catalogue view independent of client location
- Everything is in the catalogue:
  - Input/output jobs
  - Application software packages with versioning and platform information
  - Metadata definition
  - File collection – search and filtering, based on filenames and/or metadata

- Several methods:
  - X509 certificates
  - SSH keys
  - User/password
  - Job tokens
- Different roles:
  - Admin
  - Any role defined by the VO:
    - Production, analysis, physics working groups, ...
- UNIX ACL
  - Rights to owner, group, rest of the world

```
psaiz@pcitgd26:~  
[pcitgd26] /home/psaiz > aliensh  
[ aliensh 1.0.14x (C) ARDA/Alice, Andreas.Joachim.Peters@cern.ch/Derek.Feichtinger@cern.ch]  
aliensh:[alice] [1] /alice/cern.ch/user/p/psaiz/ >ls -l  
drwxr-xr-x  psaiz  psaiz           0 Aug 04 08:39  bin  
-rwxr-xr-x  psaiz  psaiz           0 Mar 26 14:19  dasdas\@das  
-rwxr-xr-x  psaiz  cc             59 Apr 27 11:26  nuevo.jdl  
-rwxr-xr-x  psaiz  psaiz        187 May 26 09:30  nuevo2.jdl  
drwxr-xr-x  psaiz  psaiz           0 Mar 15 10:22  packages  
-rwxr-xr-x  psaiz  psaiz        187 Jun 08 10:12  quotafile  
drwxr-xr-x  psaiz  psaiz           0 Mar 15 10:13  test  
drwxr-xr-x  psaiz  psaiz           0 Mar 15 10:06  triggers  
drwxr-xr-x  psaiz  psaiz           0 Mar 15 15:07  uno  
aliensh:[alice] [2] /alice/cern.ch/user/p/psaiz/ >cd bin  
aliensh:[alice] [3] /alice/cern.ch/user/p/psaiz/bin/ >mkdir test  
aliensh:[alice] [4] /alice/cern.ch/user/p/psaiz/bin/ >ls -l  
drwxr-xr-x  psaiz  psaiz           0 Aug 04 08:39  .checkStartTime.sh  
-rwxr-xr-x  psaiz  psaiz        119 Aug 04 08:39  alienroot  
-rwxr-xr-x  psaiz  psaiz        850 Mar 11 13:55  analysis.sh  
-rwxr-xr-x  psaiz  z2             65 Aug 04 08:39  certificate  
-rwxr-xr-x  psaiz  psaiz        612 Aug 04 08:39  checkStartTime.sh  
-rwxr-xr-x  psaiz  psaiz           81 Aug 04 08:39  createFile  
-rwxr-xr-x  psaiz  psaiz        146 Aug 04 08:39  gsiProduction  
-rwxr-xr-x  psaiz  psaiz       2401 Aug 04 08:39  hello  
-rwxr-xr-x  psaiz  psaiz        956 Aug 04 08:39  peekaboo.csh  
-rwxr-xr-x  psaiz  psaiz       2519 Aug 04 08:39  peekaboo.csh~  
-rwxr-xr-x  psaiz  psaiz        767 Aug 04 08:39  pp_sim.sh  
-rwxr-xr-x  psaiz  psaiz        203 Aug 04 08:39  pwd.sh  
-rwxr-xr-x  psaiz  cc            1024 Aug 04 08:39  runtest.sh  
-rwxr-xr-x  psaiz  psaiz        622 Aug 04 08:39  Sh_AnaESD_MC_CheckGenCascade.sh  
  
drwxr-xr-x  psaiz  psaiz           0 Jun 15 18:34  test  
-rwxr-xr-x  psaiz  z2           2294 Aug 04 08:39  tpcRecAlien.sh  
aliensh:[alice] [5] /alice/cern.ch/user/p/psaiz/bin/ >
```

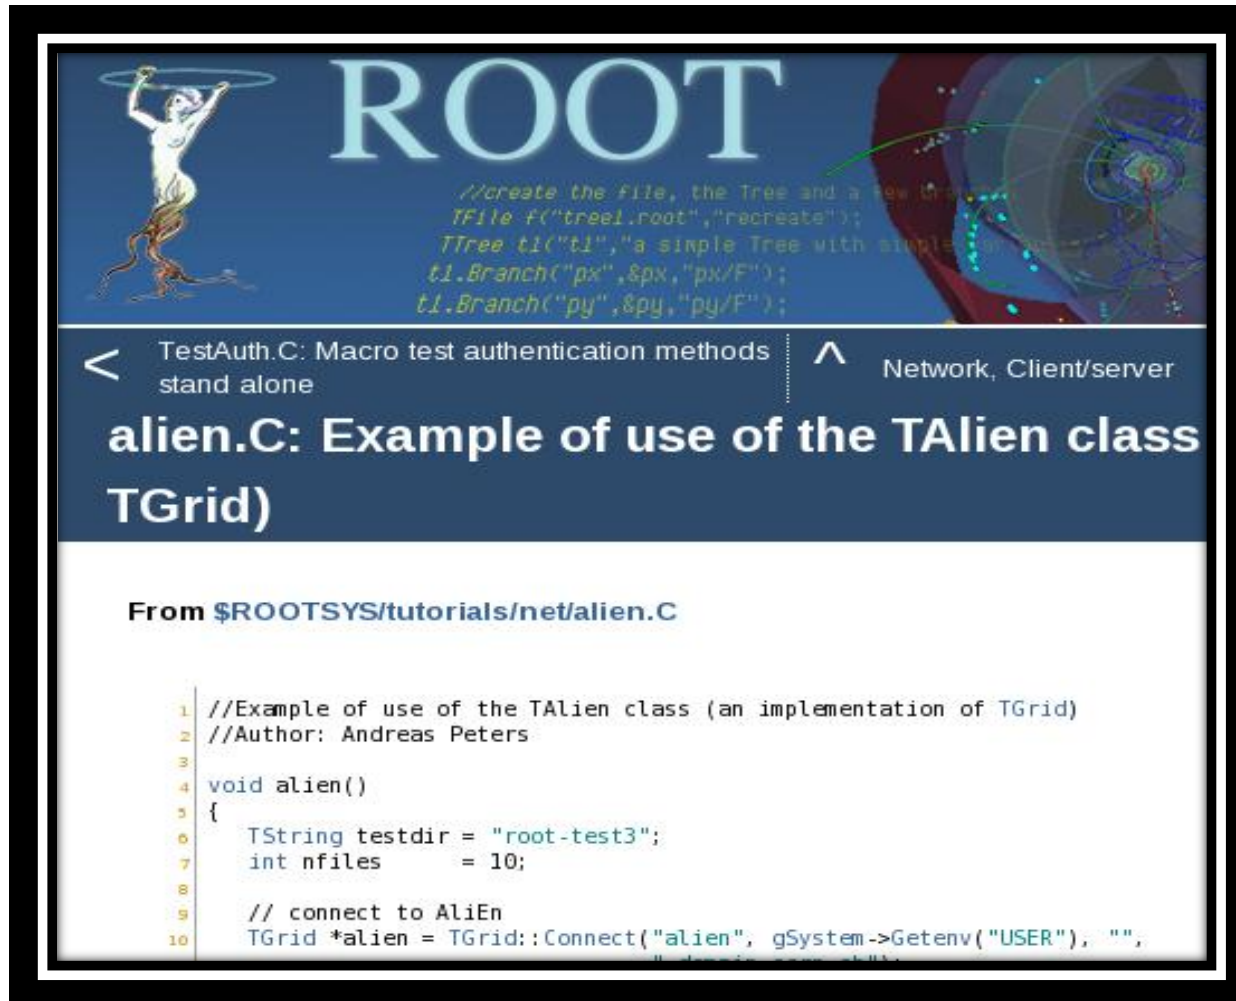
The screenshot shows a web-based file manager interface. On the left is a directory tree with folders like 'psaiz', 'bin', 'packages', 'test', 'triggers', 'uno', 'pscott', 'pshukla', 'psilvalo', 'pissot', 'pulvir', 'r', 's', 's', 't', 'u', 'v', 'w', 'x', 'y', 'z', 'data', 'jdt', 'macros', and 'packages'. The main area displays the path '/alice/cern.ch/user/p/psaiz' and a welcome message. Below this is a table of files and folders:

Permissions	Owner	Timestamp	Name
drwxr-xr-x	psaiz:psaiz	04 Aug 2009 08:39	bin
drwxr-xr-x	psaiz:psaiz	15 Mar 2010 10:22	packages
drwxr-xr-x	psaiz:psaiz	15 Mar 2010 10:13	test
drwxr-xr-x	psaiz:psaiz	15 Mar 2010 10:06	triggers
drwxr-xr-x	psaiz:psaiz	15 Mar 2010 15:07	uno

Below the first table is a section for creating a new folder, followed by another table of files:

Permissions	Owner	Timestamp	Size	Filename
-rwxr-xr-x	psaiz:psaiz	26 Mar 2010 14:19	0 B	dasdas@das
-rwxr-xr-x	psaiz:cc	27 Apr 2010 11:26	59 B	
-rwxr-xr-x	psaiz:psaiz	26 May 2010 09:30	187 B	
-rwxr-xr-x	psaiz:psaiz	08 Jun 2010 10:12	187 B	

At the bottom, there is an 'Upload a new file in this folder' section with a text input field and a 'Browse...' button. A context menu is open over the 'dasdas@das' file, showing options: View, Download, Delete, Edit, and Submit.



The screenshot shows the ROOT user interface. At the top left is a logo of a woman balancing a globe on her head. The word "ROOT" is displayed in large, light blue letters. Below it, there is a snippet of C++ code for creating a TTree and adding branches. The main window displays a breadcrumb trail: "< TestAuth.C: Macro test authentication methods stand alone" and "Network, Client/server". The title of the active window is "alien.C: Example of use of the TAlien class (TGrid)". Below the title, the source code for "alien.C" is shown, starting with a comment: "From \$ROOTSYS/tutorials/net/alien.C".

```
//create the file, the Tree and a few files
TFile f("tree1.root","recreate");
TTree t1("t1","a simple Tree with simple branches");
t1.Branch("px",&px,"px/F");
t1.Branch("py",&py,"py/F");
```

< TestAuth.C: Macro test authentication methods stand alone

Network, Client/server

### alien.C: Example of use of the TAlien class (TGrid)

From \$ROOTSYS/tutorials/net/alien.C

```
1 //Example of use of the TAlien class (an implementation of TGrid)
2 //Author: Andreas Peters
3
4 void alien()
5 {
6     TString testdir = "root-test3";
7     int nfiles      = 10;
8
9     // connect to AliEn
10    TGrid *alien = TGrid::Connect("alien", gSystem->Getenv("USER"), "",
```

And also a  
perl interface!



- Basic commands:
  - Identical to unix:
    - ls, cp, mv, rm, mkdir, cd, whoami, cat, vi, df, du,
  - With more options:
    - find (searching through name space and metadata)
- Advanced catalogue commands:
  - Quota based: fquota, jquota
  - File location: whereis (physical location of files), mirror (replication of files), listTransfers
  - File collections: createCollection
- Other AliEn commands:
  - Software packages: packman
  - Job interaction: submit, ps, top, kill, spy, trace

- Hierarchical metadata definition
  - Inherited by all entries under a directory
  - Users can define (and populate) their own metadata
  - The metadata definition is kept also as a plain file
- Most advanced use – the entire ALICE Conditions database is annotated in the catalogue
  - File container (root objects)
  - Metadata description and search
  - Used by every ALICE job!

- Metadata:
  - last\_run, first\_run, version, object\_classname, responsible, beam\_period, aliroot\_version, comment
- Each subdetector in a different directory
- Currently, 100K entries for 2010
- Access by every job

```
[aliendb06c.cern.ch:3307] /alice/data/2010/OCDB/ > find -y TRD *.root CDB:first_run<=114005 and CDB:last_run>=114005
Jun 15 18:56:15 info Filtering according to 'and' CDB first_run<=114005
Jun 15 18:56:15 info Filtering according to 'and' CDB last_run>=114005
/alice/data/2010/OCDB/TRD/Calib/PIDLQ/Run0_999999999_v1_s0.root
/alice/data/2010/OCDB/TRD/Align/Data/Run0_999999999_v2_s0.root
/alice/data/2010/OCDB/TRD/Calib/ChamberGainFactor/Run105160_999999999_v2_s0.root
/alice/data/2010/OCDB/TRD/Calib/ChamberStatus/Run106047_999999999_v3_s0.root
/alice/data/2010/OCDB/TRD/Calib/ChamberT0/Run0_999999999_v1_s0.root
/alice/data/2010/OCDB/TRD/Calib/ChamberVdrift/Run0_999999999_v1_s0.root
/alice/data/2010/OCDB/TRD/Calib/DCS/Run113994_999999999_v670_s0.root
/alice/data/2010/OCDB/TRD/Calib/DetNoise/Run113954_999999999_v16_s0.root
/alice/data/2010/OCDB/TRD/Calib/LocalGainFactor/Run0_999999999_v65_s0.root
/alice/data/2010/OCDB/TRD/Calib/LocalT0/Run0_999999999_v1_s0.root
/alice/data/2010/OCDB/TRD/Calib/LocalVdrift/Run0_999999999_v1_s0.root
/alice/data/2010/OCDB/TRD/Calib/MonitoringData/Run0_999999999_v1_s0.root
/alice/data/2010/OCDB/TRD/Calib/PadNoise/Run113954_999999999_v16_s0.root
/alice/data/2010/OCDB/TRD/Calib/PadStatus/Run113954_999999999_v16_s0.root
/alice/data/2010/OCDB/TRD/Calib/PIDNN/Run0_999999999_v1_s0.root
/alice/data/2010/OCDB/TRD/Calib/PIDThresholds/Run0_999999999_v1_s0.root
/alice/data/2010/OCDB/TRD/Calib/PRFwidth/Run0_999999999_v1_s0.root
/alice/data/2010/OCDB/TRD/Calib/RecoParam/Run0_999999999_v1_s0.root
/alice/data/2010/OCDB/TRD/Calib/HalfChamberStatus/Run106047_999999999_v1_s0.root
[aliendb06c.cern.ch:3307] /alice/data/2010/OCDB/ > []
```

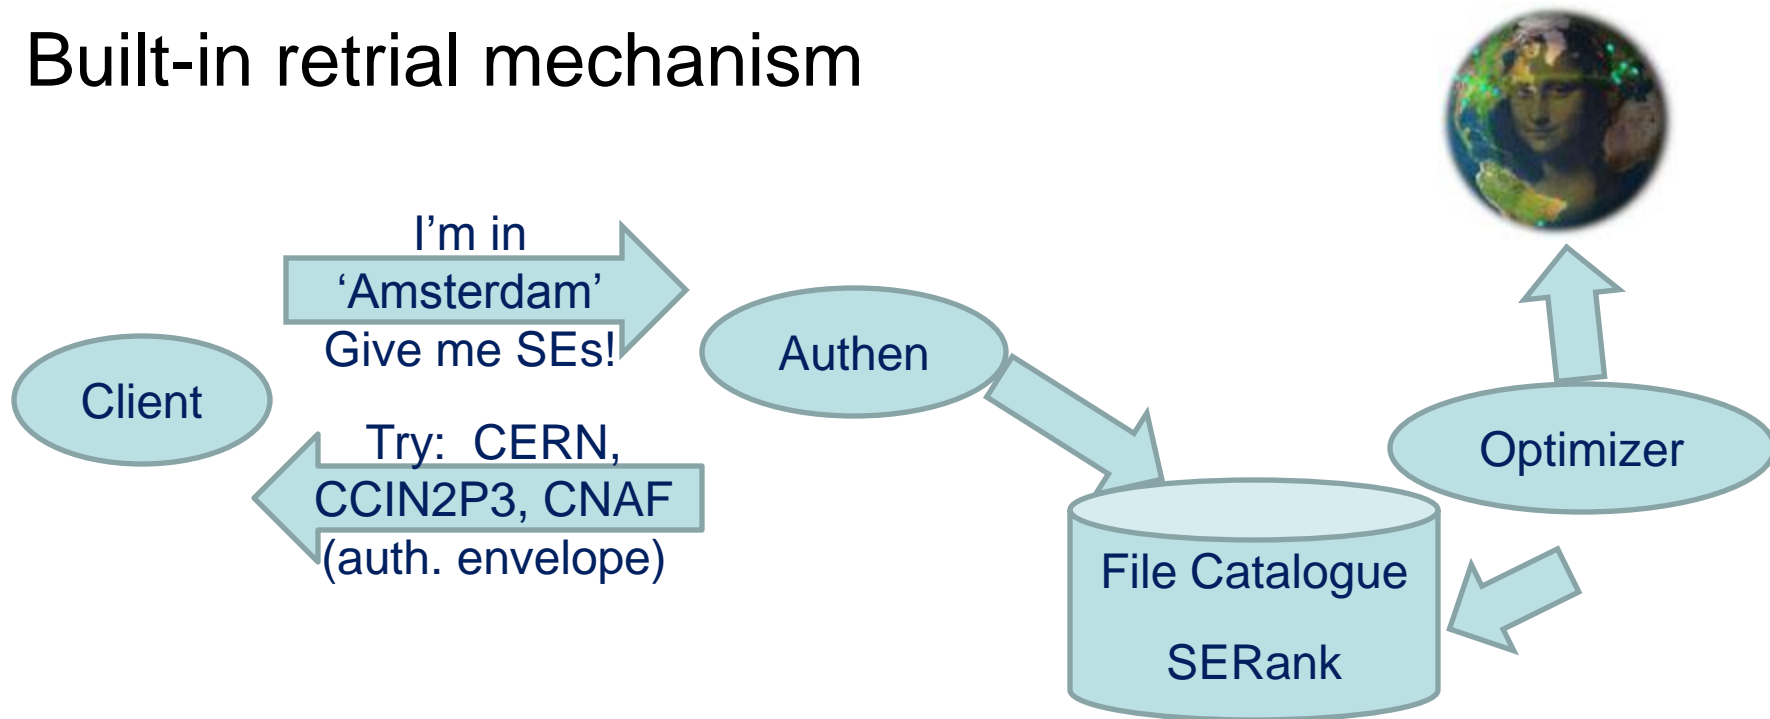
- Users usually need many files
- A collection is a list of LFNs
- Appears as a single entry in the catalogue
  - And can be split according to the location
- Facilitate actions on multiple files: cp, mirror, analyze...

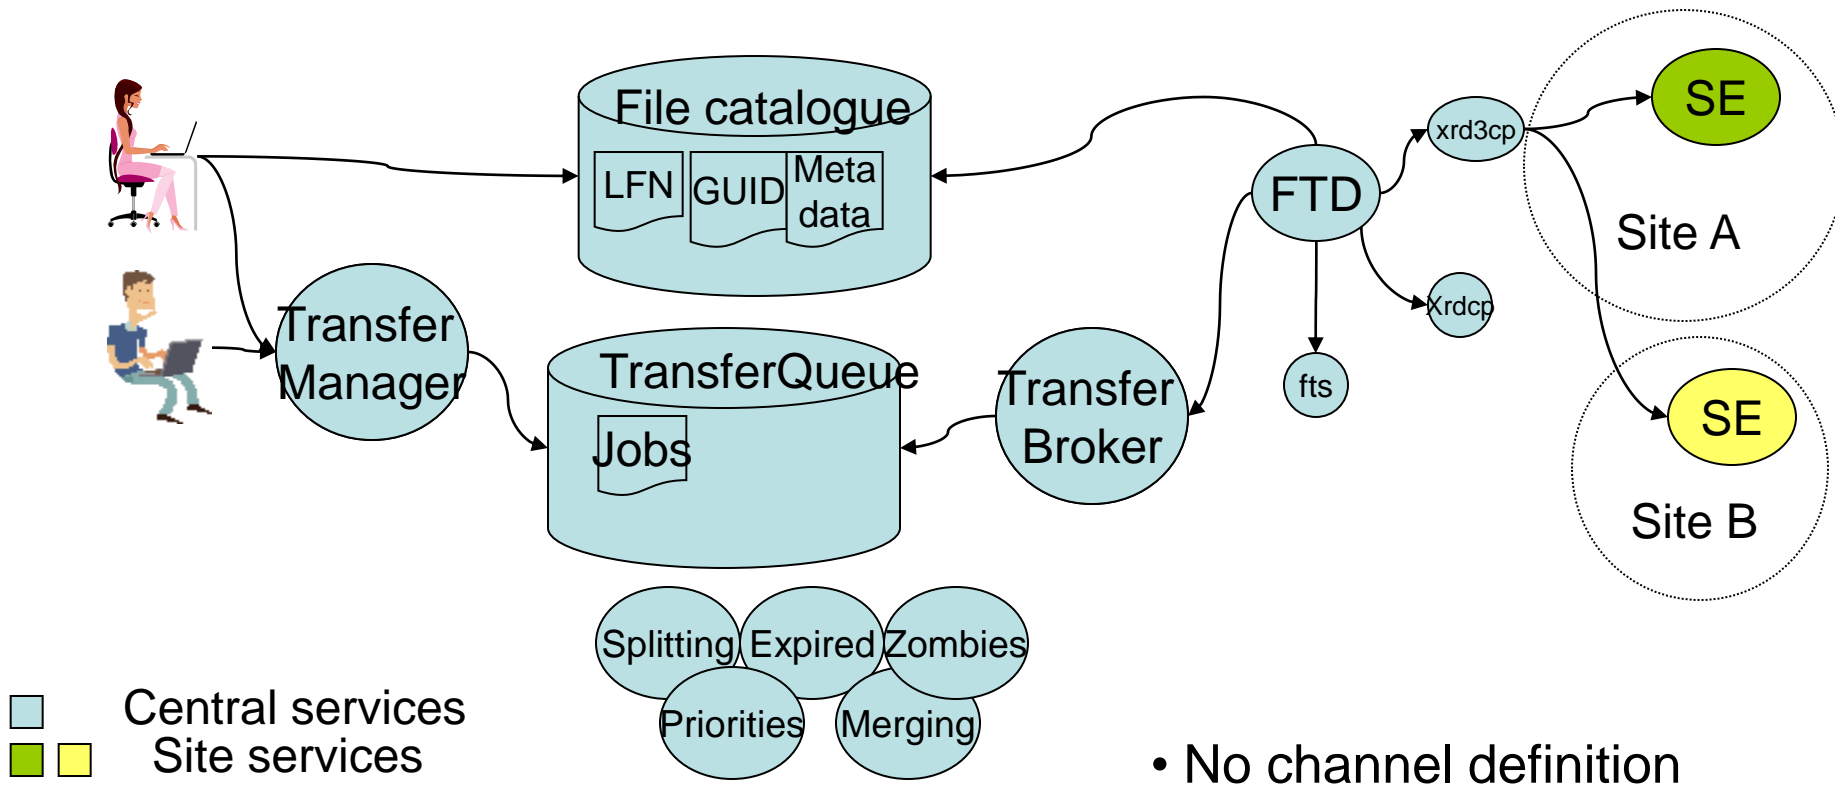
```
[aliendb06c.cern.ch:3307] /alice/data/2010/ > ls -al LHC10c/000120677/collection
crwxr-xr--  alidaq  alidaq   9564798201 May 24 00:03      collection
[aliendb06c.cern.ch:3307] /alice/data/2010/ > listFilesFromCollection LHC10c/000120677/collection
3C67BA10-66A9-11DF-8001-5523A10ABEEF (from the file /alice/data/2010/LHC10c/000120677/raw/10000120677048.10.root)
61AEF5FE-66A9-11DF-8001-5523A10ABEEF (from the file /alice/data/2010/LHC10c/000120677/raw/10000120677048.100.root)
6538478E-66A9-11DF-8001-5523A10ABEEF (from the file /alice/data/2010/LHC10c/000120677/raw/10000120677048.110.root)
68E76D24-66A9-11DF-8001-5523A10ABEEF (from the file /alice/data/2010/LHC10c/000120677/raw/10000120677048.120.root)
6C99E104-66A9-11DF-8001-5523A10ABEEF (from the file /alice/data/2010/LHC10c/000120677/raw/10000120677048.130.root)
70478C8E-66A9-11DF-8001-5523A10ABEEF (from the file /alice/data/2010/LHC10c/000120677/raw/10000120677048.140.root)
73F4D1E8-66A9-11DF-8001-5523A10ABEEF (from the file /alice/data/2010/LHC10c/000120677/raw/10000120677048.150.root)
77A60474-66A9-11DF-8001-5523A10ABEEF (from the file /alice/data/2010/LHC10c/000120677/raw/10000120677048.160.root)
7BBA4386-66A9-11DF-8001-5523A10ABEEF (from the file /alice/data/2010/LHC10c/000120677/raw/10000120677048.170.root)
7EC525F0-66A9-11DF-8001-5523A10ABEEF (from the file /alice/data/2010/LHC10c/000120677/raw/10000120677048.180.root)
821E1BC6-66A9-11DF-8001-5523A10ABEEF (from the file /alice/data/2010/LHC10c/000120677/raw/10000120677048.190.root)
43DCC43E-66A9-11DF-8001-5523A10ABEEF (from the file /alice/data/2010/LHC10c/000120677/raw/10000120677048.20.root)
857E2892-66A9-11DF-8001-5523A10ABEEF (from the file /alice/data/2010/LHC10c/000120677/raw/10000120677048.200.root)
89464BBC-66A9-11DF-8001-5523A10ABEEF (from the file /alice/data/2010/LHC10c/000120677/raw/10000120677048.210.root)
8CBF8E84-66A9-11DF-8001-5523A10ABEEF (from the file /alice/data/2010/LHC10c/000120677/raw/10000120677048.220.root)
```

- Integrated with a job management system
  - Using the same interface
  - Input/Output/Executable from the catalogue
    - Automatically split according to file co-location
  - Packages from the catalogue


```
[
Executable = "/alice/cern.ch/user/p/pgonzale/bin/analysis.sh";
Packages = { "VO_ALICE@AliRoot::v4-19-13-AN",
            "VO_ALICE@ROOT::v5-26-00b-5",
            "VO_ALICE@APISCONFIG::V1.1x" };
OutputFile = { "Diffractive.ESD.1.root@disk=2",
              "EventStat_temp.root@disk=2" };
Split = "se";
OutputDir = "/alice/cern.ch/user/p/pgonzale/PedroDiff_LHC10b_New_1/output/014/#alien_counter_03i#";
InputFile =
{
    "LF:/alice/cern.ch/user/p/pgonzale/PedroDiff_LHC10b_New_1/KOAnalysis.C",
    "LF:/alice/cern.ch/user/p/pgonzale/PedroDiff_LHC10b_New_1/ESD.par",
    "LF:/alice/cern.ch/user/p/pgonzale/PedroDiff_LHC10b_New_1/AOD.par",
};
User = "pgonzale";
TTL = 30000;
InputDataCollection = {"LF:/alice/cern.ch/user/p/pgonzale/PedroDiff_LHC10b_New_1/000117112.xml,nodownload" };
Validationcommand = "/alice/cern.ch/user/p/pgonzale/PedroDiff_LHC10b_New_1/analysis_validation.sh"
]
```

- Users/jobs should be automatically directed to the best SE to store/retrieve files
  - Find the closest working SE of a given QoS
- Built-in retrial mechanism



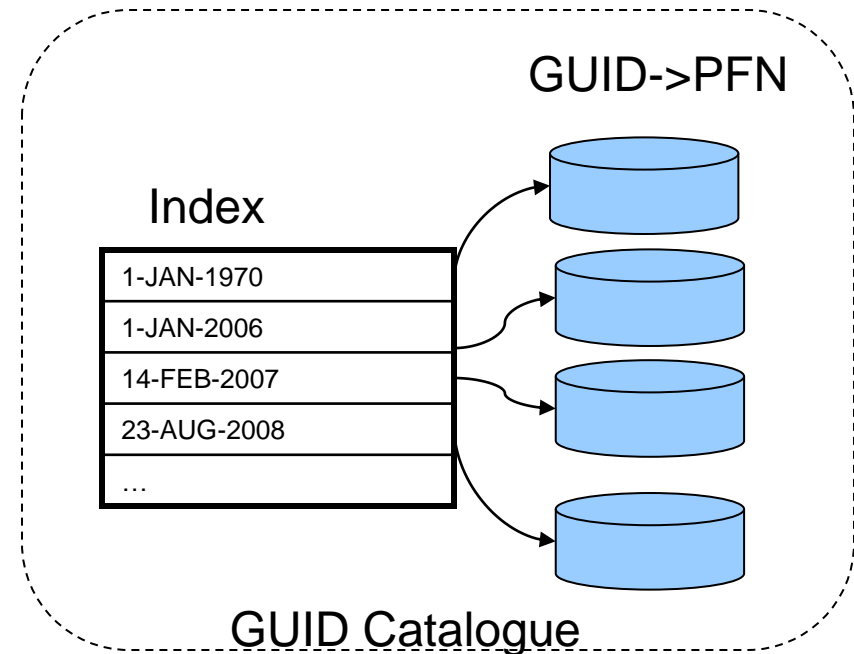



Central services  
Site services






- No channel definition
- Limit on transfers per SE
- Multiple protocols

- Entries can be retrieved by LFN or GUID
- Flexible table structure, allowing for easy scalability
- ALICE: 150M entries (100GB DB size)





- Improving SE layout
  - Currently, guid based, moving towards lfn based
  - Global namespace with human readable files also on storage (replicas and xrootd global namespace fully supported)
- File deletion
  - Efficient physical removal of files (with their replicas)
  - Currently the ‘orphaned’ files (SE malfunctions) are a problem
- AliEn over ORACLE
- Multiple copies of the catalogue
  - Already replicated (used for backups)
  - Enable slaves as ‘read-only’ (if the load is too high for one single machine)

- Provides a Global catalogue for all users and jobs
- UNIX-like file system
  - Same set of commands, with Grid-related extensions
  - Easy to use (for UNIX users)
- Included metadata catalogue:
  - Powerful search algorithms for all types of jobs, conditions data storage and annotation
- And more: file collections, file transfers, quotas, automatic SE selection, software packages
- Used by several experiments   

<http://alien.cern.ch>