

Grid File Name (GFN) and Site File Name (SFN)

Arie Shoshani, Alexander Sim - LBNL
Asad Samar - Caltech
Heinz Stockinger, Wolfgang Hoeschek - CERN
Shahzad Muzaffar, Rich Wellner - FNAL
29 June 2001

It is generally accepted that in a grid environment every file accessible to clients on the grid should have a grid-wide globally unique file name, and that every replica should have a site-specific file name. These names should be stored in the Replica Catalog (RC) in order for clients to find all replica site locations and file names used by that site. The terminology used for these file names are: `logical_file_name` (LFN) and `physical_file_name` (PFN), respectively.

In other words, the LFN uniquely identifies a file name for a set of identical replicas (files) which in turn are uniquely identified by their PFNs. One of the aims of this document is also to provide a terminology that is understood by both grid projects, PPDG and European DataGrid .

Note that by “file name”, we mean a “`directory_path/file_name`”, such as “`ppdg/fermi/projectX/foo`”. This applies to both LFNs and PFNs.

In the European DataGrid project similar terms are defined in the Architecture Task Force (ATF) document [1]. In detail, three different file names are defined in order to distinguish between file names and names for actually transferring a file. The following three names are proposed:

LFN (logical file name), e.g. `lfn://cms.org/run1/file1`

PFN (physical file name), e.g. `pfn://host1.cern.ch/data/any_other_dir/run1/file1.dbf`

TFN (transfer file name: contains the transfer protocol),
e.g. `gridftp://host1.cern.ch/run1/file1.dbf`

Recently, we have considered what these names mean in the context of using a Hierarchical Resource Managers (HRMs). Since HRMs manage the access to both files on tape as well as files on disk, an HRM’s PFN can have an ambiguous meaning. For example, the Fermi-HRM (that uses SAM/Enstore) considers the PFN an independent name (which they refer to as the “logical file name”) that they internally map to the actual `storage_file_name`, which could be either a disk address or a tape address. In contrast, the LBNL-HRM considers the PFN the actual storage address.

Furthermore, different projects may choose the PFN to be either a tape address or a disk address. For example, GDMP currently considers the `disk_file_name` as the PFN, thus its HRM would be required to support the mapping from the `disk_file_name` to the `tape_file_name` in case that the file was migrated to tape, and removed from disk. A similar approach is used in the AMS-HPSS interface in Objectivity where the default location for a PFN is a file name on disk.

To avoid this ambiguity for a PFN, we propose next terminology and procedures for assigning file names. Our goal is to provide the following features:

- 1) We wish to accommodate file access through HRMs and DRMs (if they exist), or directly through GridFTP, or directly with other protocols (such as plain FTP).
- 2) We wish to allow “site transparency”, i.e. to let a site manage internally the file location on storage devices as it see fit.

First, to avoid the ambiguity of the meaning of LFN and PFN, we introduce the following terms:

- A `grid_file_name` (GFN): a globally unique file name for the grid. This name is assigned by data collection creator, and registered in the RC. In DataGrid terminology, the GFN is equivalent to the LFN.
- A `site_file_name` (SFN): a site specific file name for a replica.. It is the choice of the site administrator what these file names reflect; they can be a disk address, a tape address, or a name independent of both. The site name for a replica is registered in the RC as part of a site URL that includes the protocol and site location. In DataGrid terminology, the SFN is equivalent to the file path of the TFN. The DataGrid TFN is similar to a URL and thus contains the protocol name, the host and the file path. Note that in DataGrid the term `StorageElement` is used for indicating the smallest granularity of data access. In detail, a site can have several `StorageElements`.

An example of a GFN is: “`ppdg/fermi/projectX/foo`”. Note that the collection name is: “`ppdg/fermi/projectX`”, and the file name is “`foo`”.

An example of a SFN URL at LBNL is: “`hrm://sleepy.lbl.gov:1000/hpss_ppdg/foo`”. Note that “`hrm`” is the protocol, “`sleepy.lbl.gov`” is the site address, “`1000`” is the port, and “`hpss_ppdg/foo`” is the SFN. In this case, the SFN may mean a storage location on HPSS tape. When a file is requested, the HRM will check if it is on its disk. If so, the disk URL is returned. If not, the file is staged to disk, and the disk URL is returned.

An example for the same file at Fermi is: “`hrm://sam.fermi.gov/projectX/john/foo`”. Here, the SFN is “`projectX/john/foo`”, and it may be an independent name that the FERMI-HRM maps internally (using the SAM directory) to a disk or tape location. As above, it will stage the file to disk if necessary and return the disk URL.

The above solution achieves the two goals we mentioned above. The first goal is satisfied by using the protocol in the URL. Specifically, if the system accessed has no HRM or DRM, but supports GridFTP, the site URL will reflect the storage location. For example: “`gridftp://dm.lbl.gov/disk10/temp_dir/foo`”. Similarly, if only FTP is supported, the site URL could be: `ftp://dm.lbl.gov/disk10/temp_dir/foo`.

The second goal is satisfied by permitting the site URL to be anything the site chooses to use. For example, if the site uses an HRM, the HRM could do internal mappings of site addresses to storage addresses. Another HRM may choose to have no internal mappings, and will use tape addresses. Similarly, a site that manages a disk cache may choose the

site location to be a disk address. This will permit access to the file directly by gridFTP or FTP without involving a DRM.

Reference:

[1] ATF homepage: <http://www.cern.ch/grid-atf/>, Version 2 of the architecture document June 2001.