SuGI
SUSTAINABLE GRID INFRASTRUCTURE

dCache.ORG

dCache.ORG

# *SRM 2.2 in dCache*

Patrick Fuhrmann

support and funding by






GEFÖRDERT VOM
Bundesministerium für Bildung und Forschung
D-GRID

## Assignment of directory (Tags) to pools

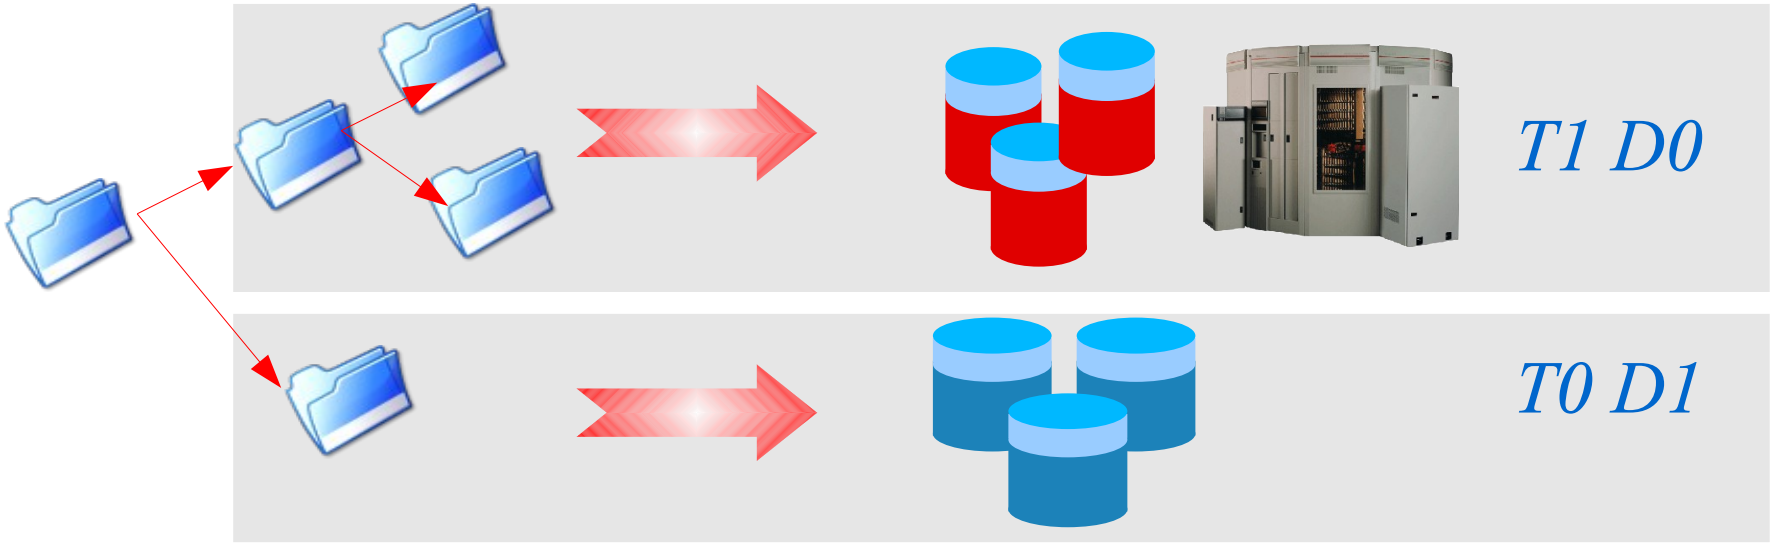


dCache.ORG

## Assigning 'space tokens' to pools.

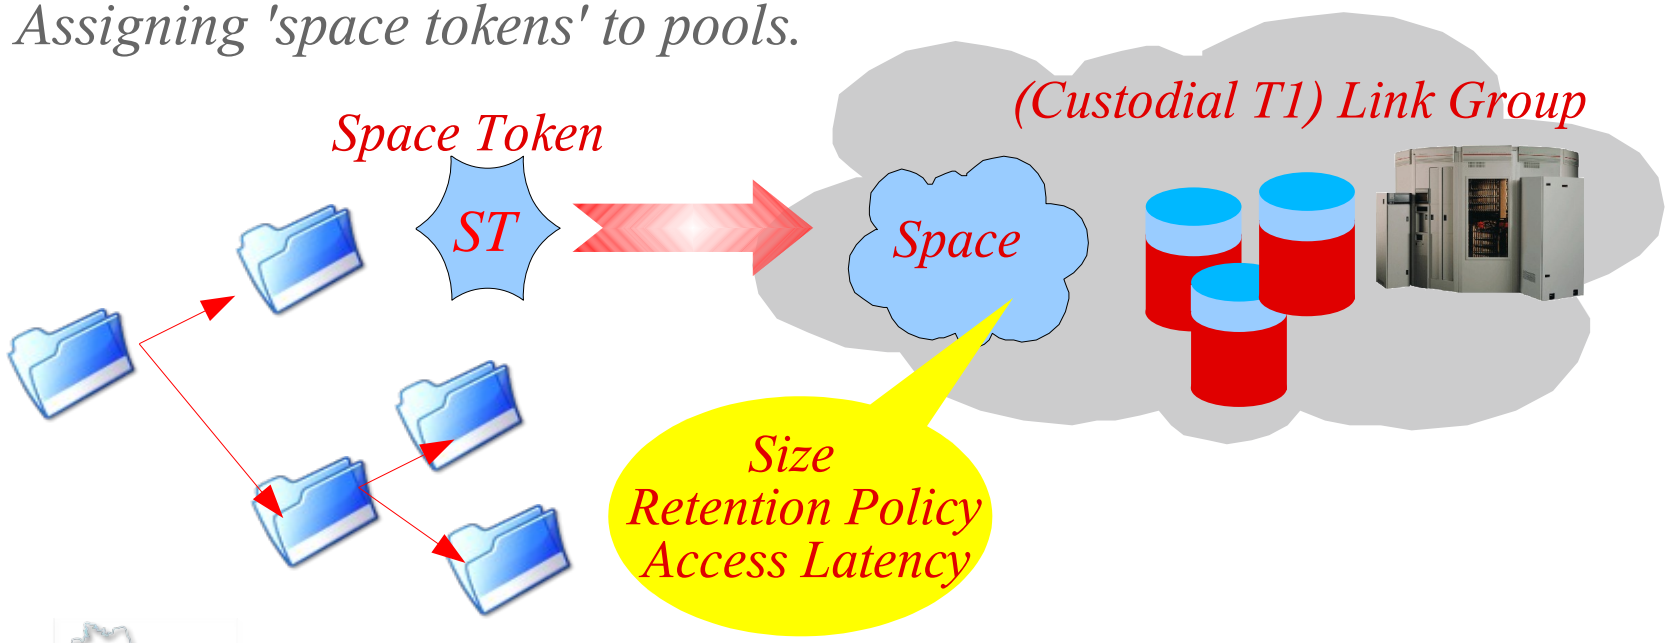

# Request Flow

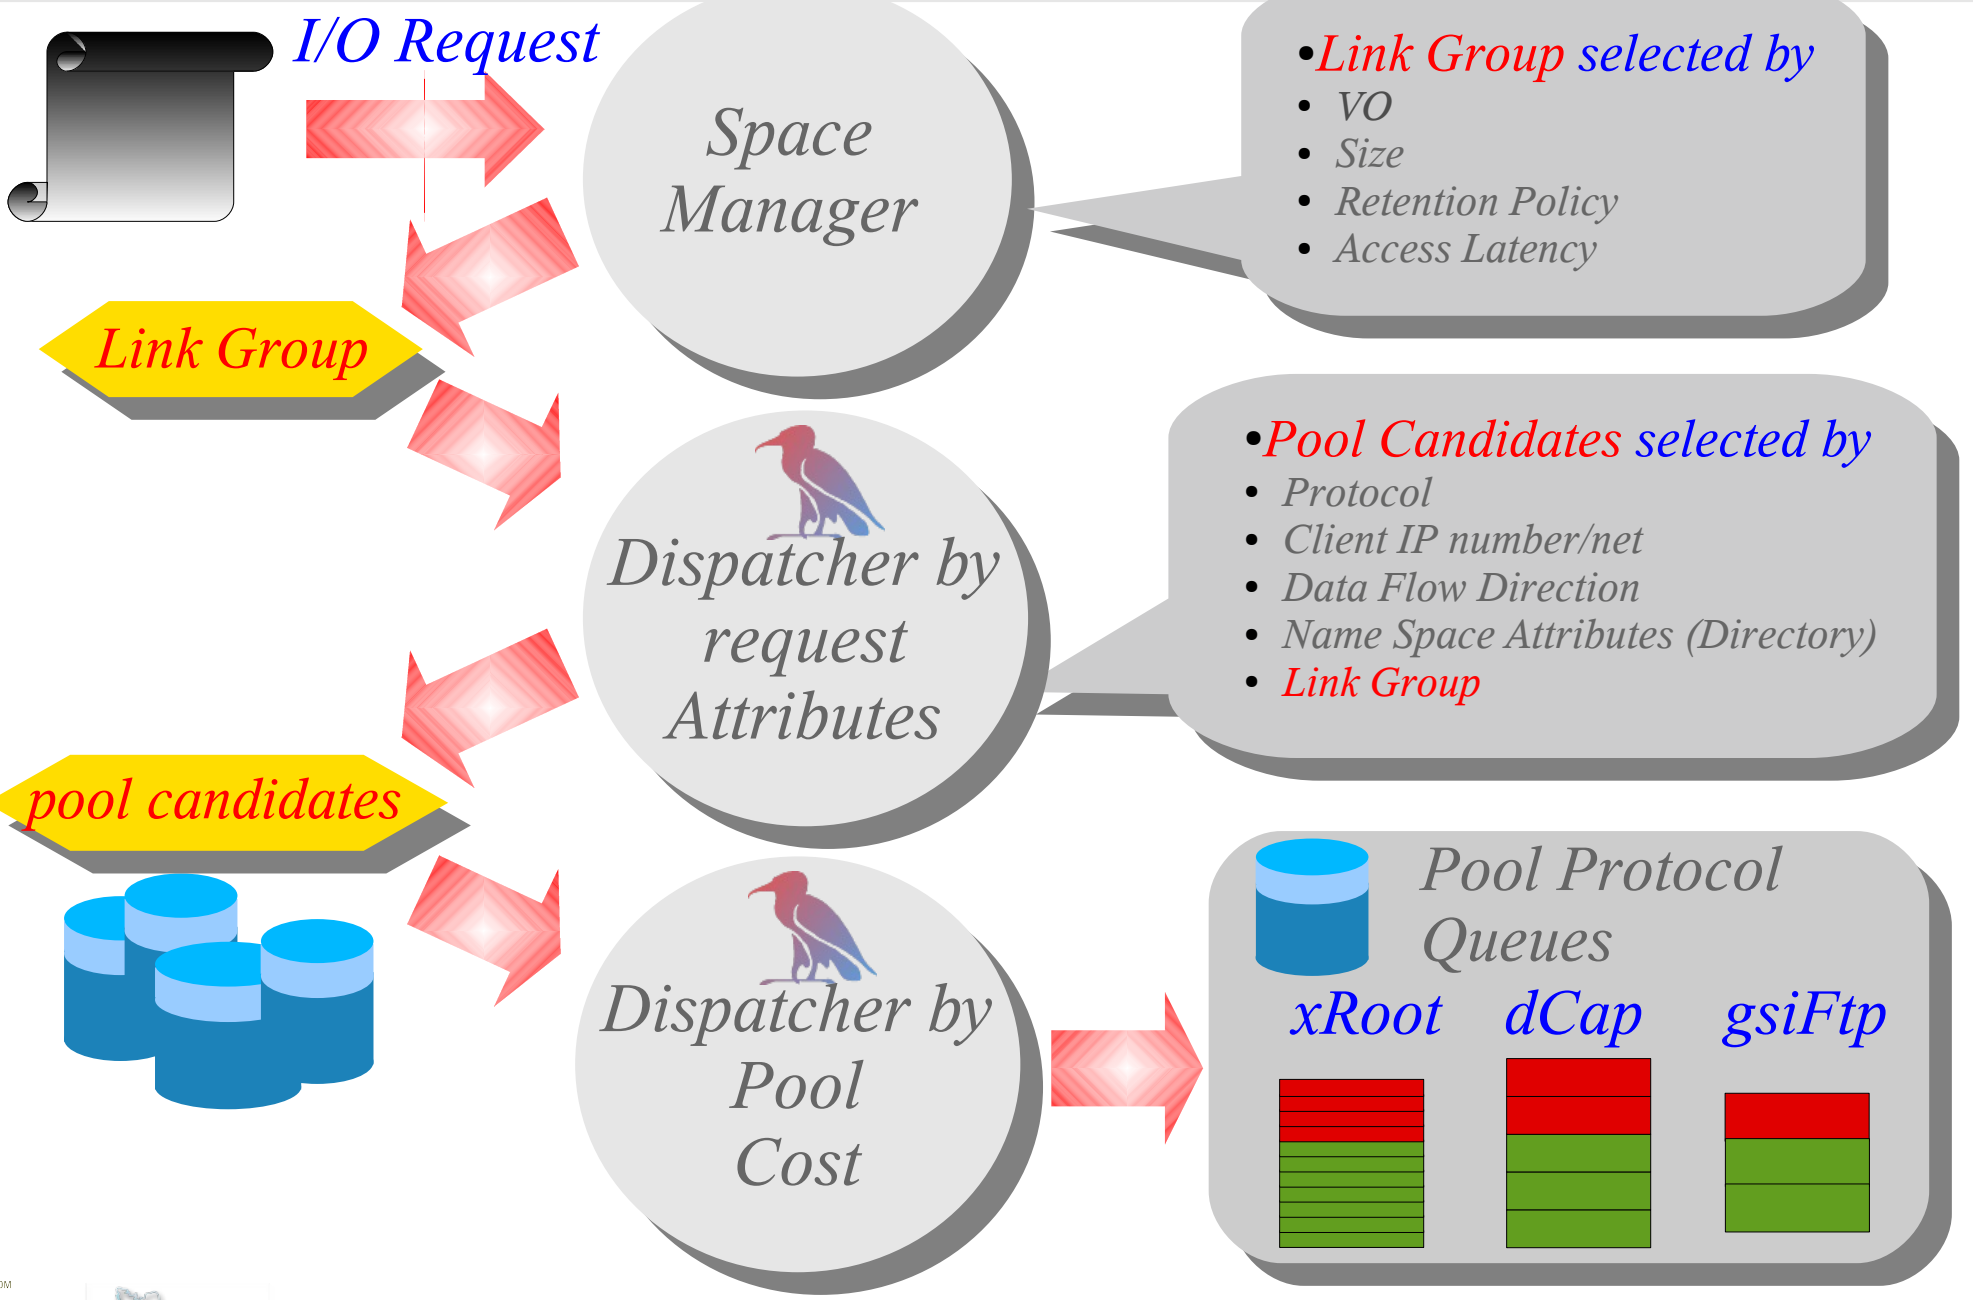


dCache.ORG

dCache.ORG

# Recall of **pool selection** w/o Space Management

*A link associates a set of input parameter with a set of pool groups or pools.*

*Request :* { IP Number , Storage Class , Protocol, I/O Direction }

*Link :*

{ IP Subnet A, IP Subnet B, IP Subnet C } { Storage Class A, Storage Class B, Storage Class C } { Protocol A, Protocol B, Protocol C }

**Link** *foo*

readpref=x
writepref=y
***

Pool Group A

Pool Group B

*e.g. :* {131.169.3.3 , cms:generated ; gsiFtp , PUT}

{ 0.0.0.0/0.0.0.0 } { cms:generated, cms:mc } { gsiFtp, gsiDCap }

WRITE_LNK

readpref=0
writepref=10

Write-Pools

- Without Space Management, a *Request* is compared with *all Links*, permitting the requested I/O direction (PUT/GET/STAGE)
- Links which match the request attributes are resolved into pools, which subsequently are considered as candidates for the final Cost Calculation.

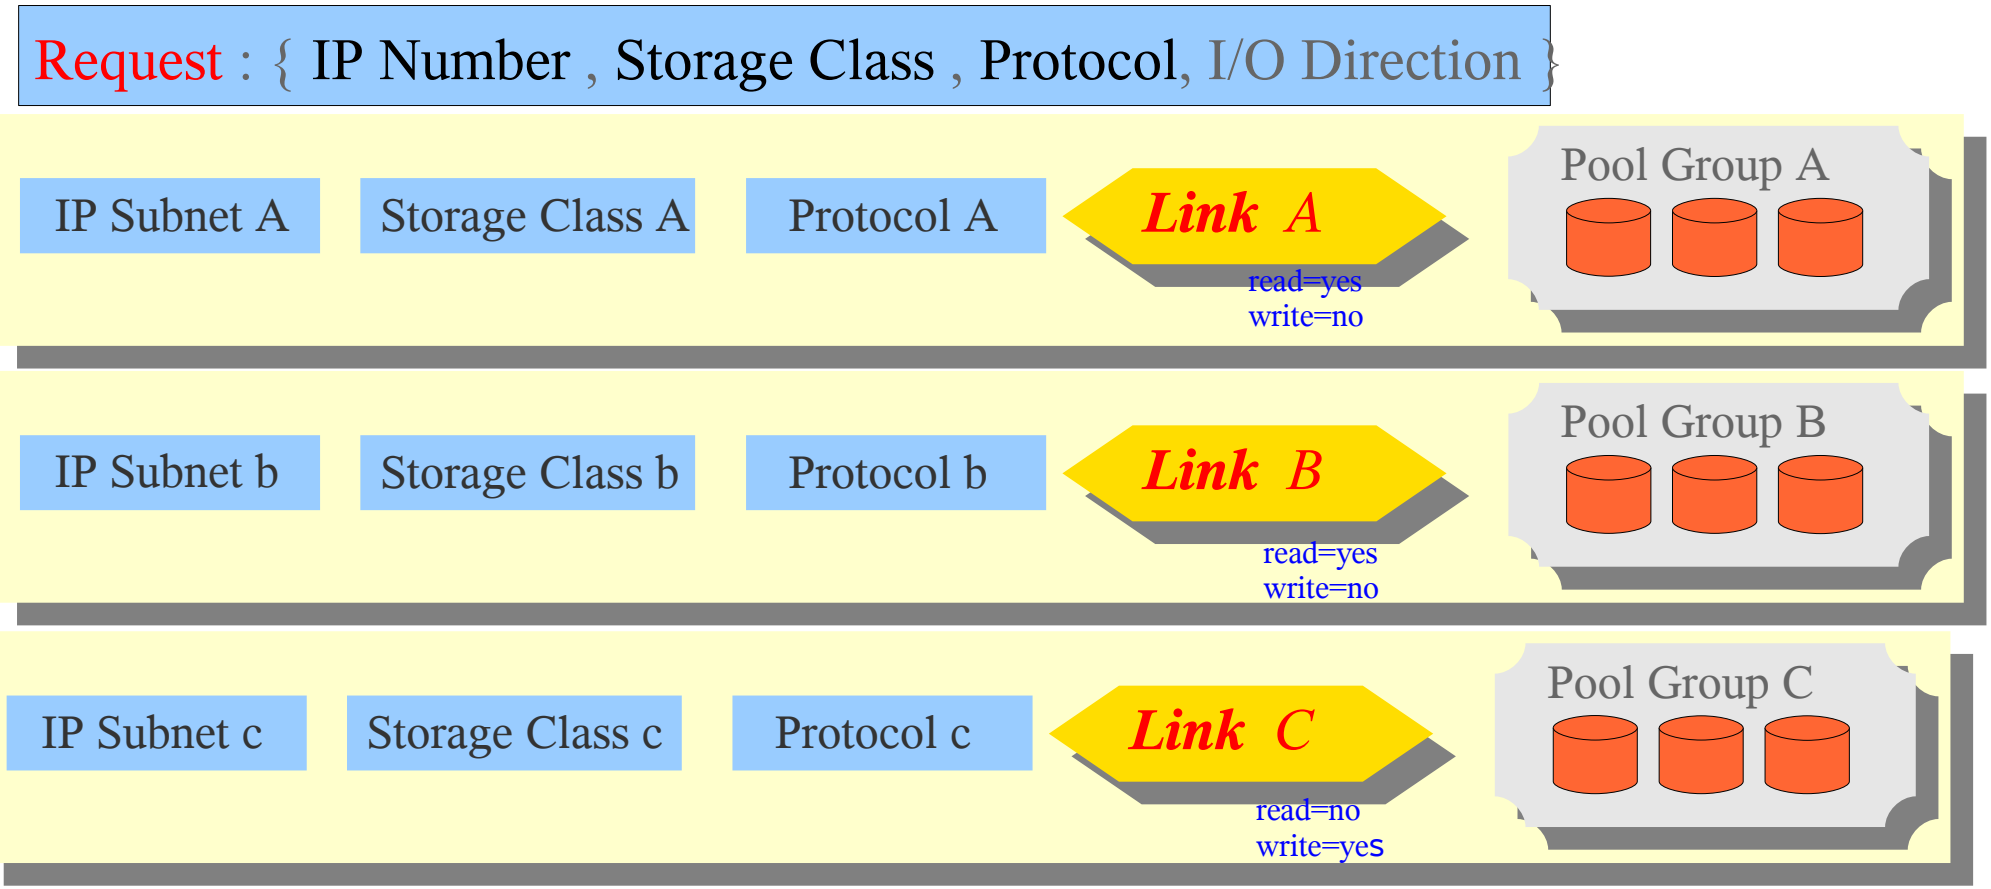


dCache.ORG

# What is a dCache *LinkGroup*

- A *Link Group* is a collection of *Links*.
- Space Reservation can only be performed on Link Groups
- Space Reservation can either be
  - **explicit**, using Space Tokens.
  - **implicit** for individual files if no Space Token is specified.
- A *Link Group* may only allow a subset of Storage Attributes for Space Reservation (Online, Nearline, Custodial, Replica)
- A Link Group supports a configurable set of VO's.
  - Only those VO's (FQANs) are permitted to reserve space on that Link Group.
- The total size of a *Link Group* is the sum of the sizes of all *Pools* in that *Link Group*.
  - Resolving : Link Group -> Links -> Pool Groups -> Pools
  - A pool is only counted once
- A Pool must not be part of more than one Link Group.

dCache.ORG

# What is a Space Token ?

- A Space Token is a reserved storage area on a Link Group. Space Token properties are :
  - Total size, used space, available space.
  - Pair of Storage Attributes. (Access Latency, Retention Policy)
  - Owned by a VO

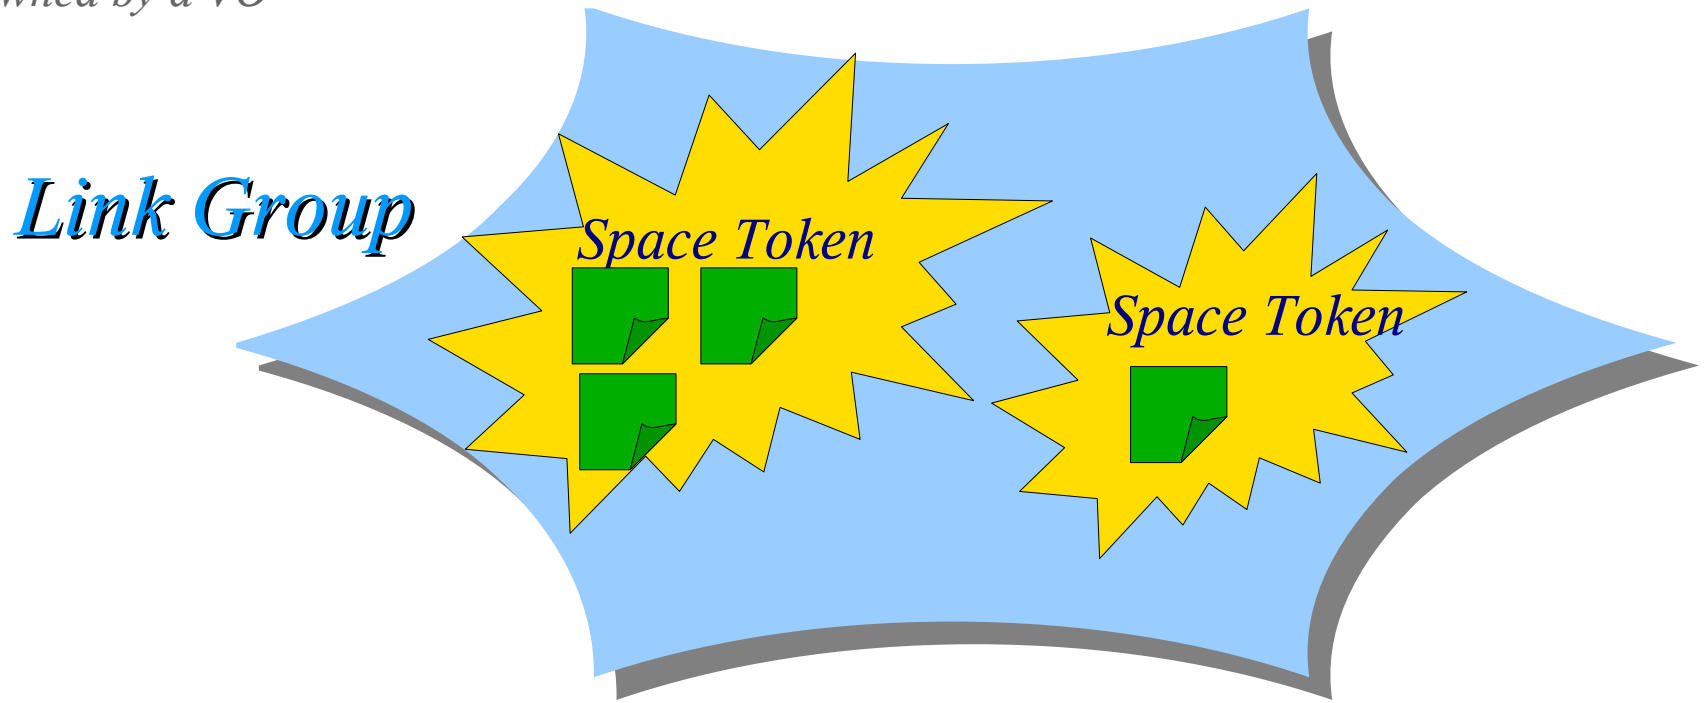

# Creating a Space Token



dCache.ORG

- Using the SrmSpaceManager admin command '*reserve*'.
  - Can only be issued by local 'sysadmin'.
  - The Link Group, where the Space Token should be created on, can be explicitly specified.
- Using the SRM client command '*reserve space*'.
  - Because the concept of a Link Group is not part of the SRM specification, the Link Group, where the Space Token should be created on, can't be specified directly.
  - The Link Group is selected by dCache, considering the following constrains:
    - FQAN of the requester must match the configured VO's of the LinkGroup.
    - The Link Group must allow the requested Access Latency and Retention Policy.
    - The requested size of the space reservation must fit into the available space of the Link Group

# How to use a token.

SuGI SUSTAINABLE GRID INFRASTRUCTURE



dCache.ORG

*Space Tokens may be used by any SRM 2.2 clients (prepare-to-put only)*

- *dCache SRM 2.2 client (srmcp)*
- *FTS (gLite File Transfer Service)*
- *LCG Tools e.g. lcg-cr, lcg-cp ****

dCache.ORG

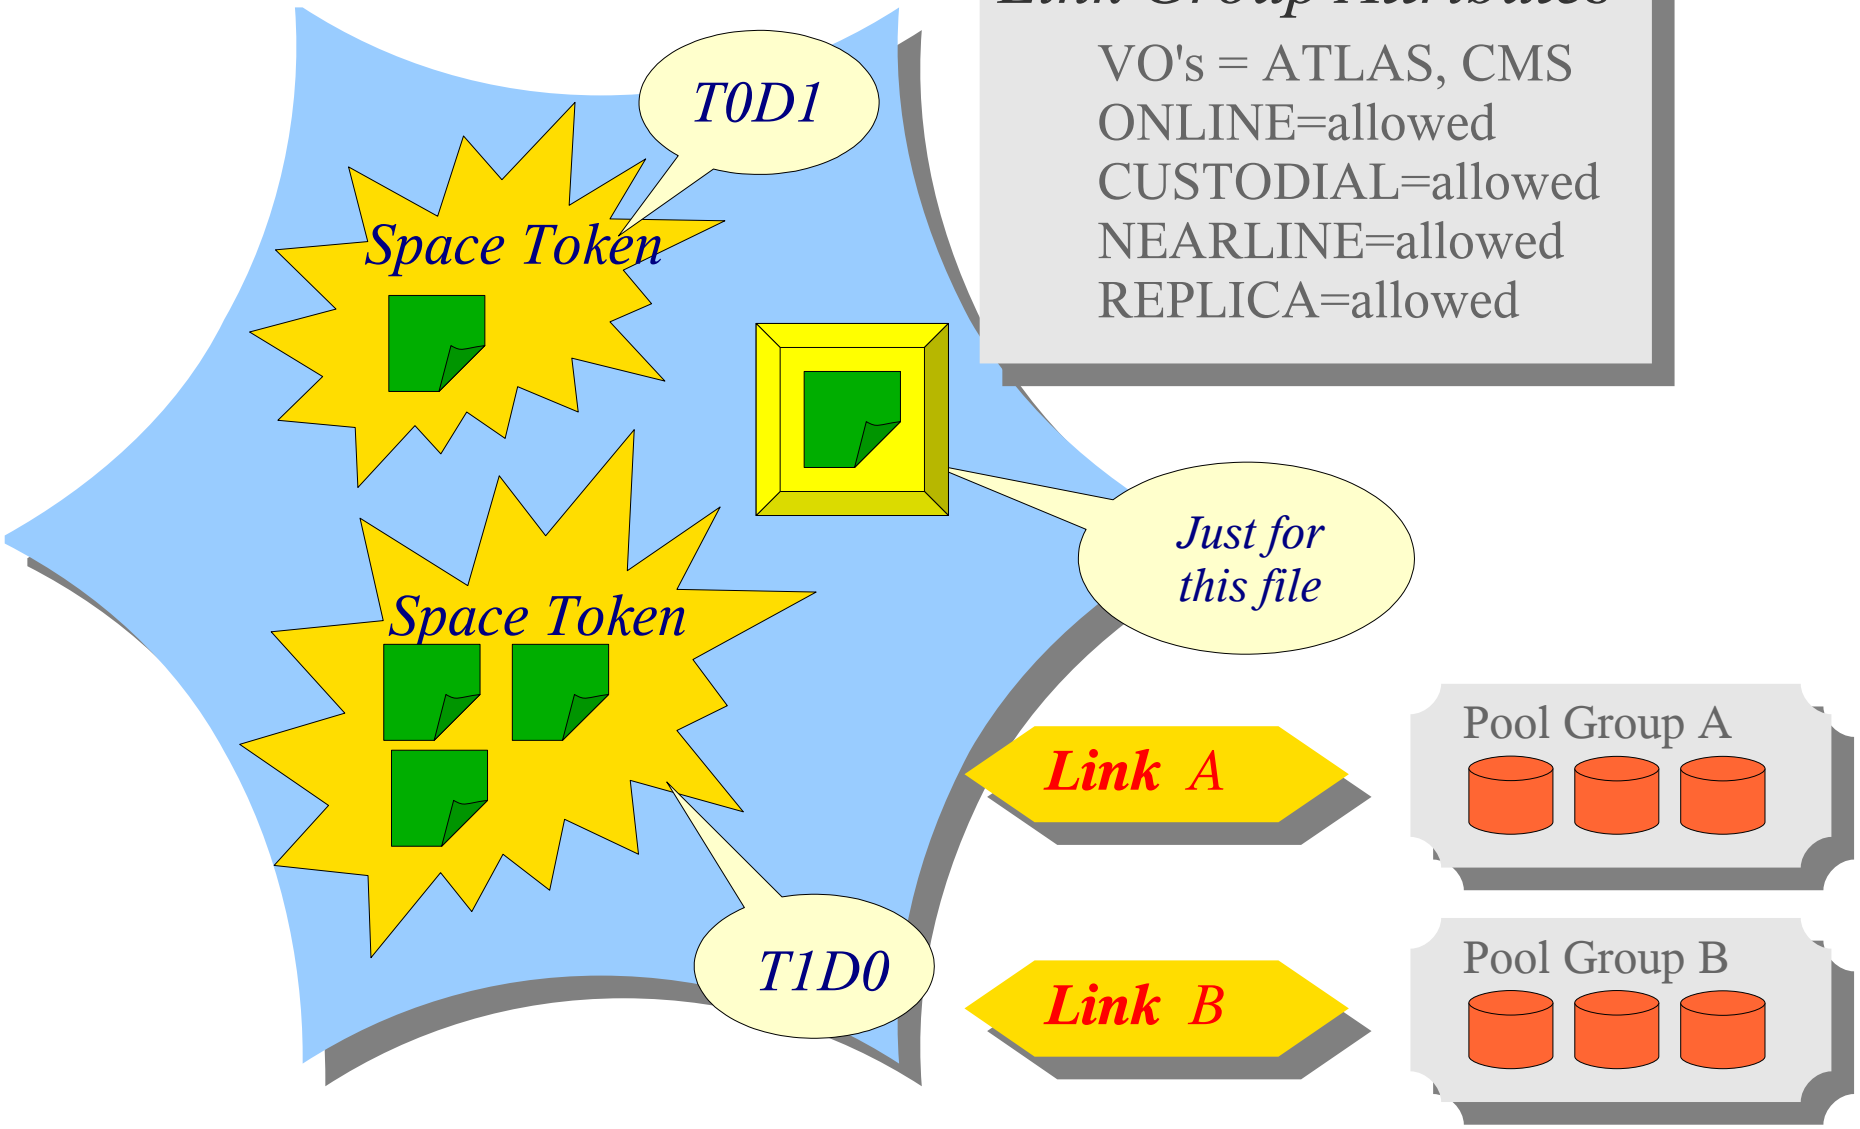

# Link Selection, using a Token

- Writing into a Space Token limits the subsequent selection to the associated Links (A and B).
- When not writing into a Space Token, the behavior depends on the system configuration. See next slides.

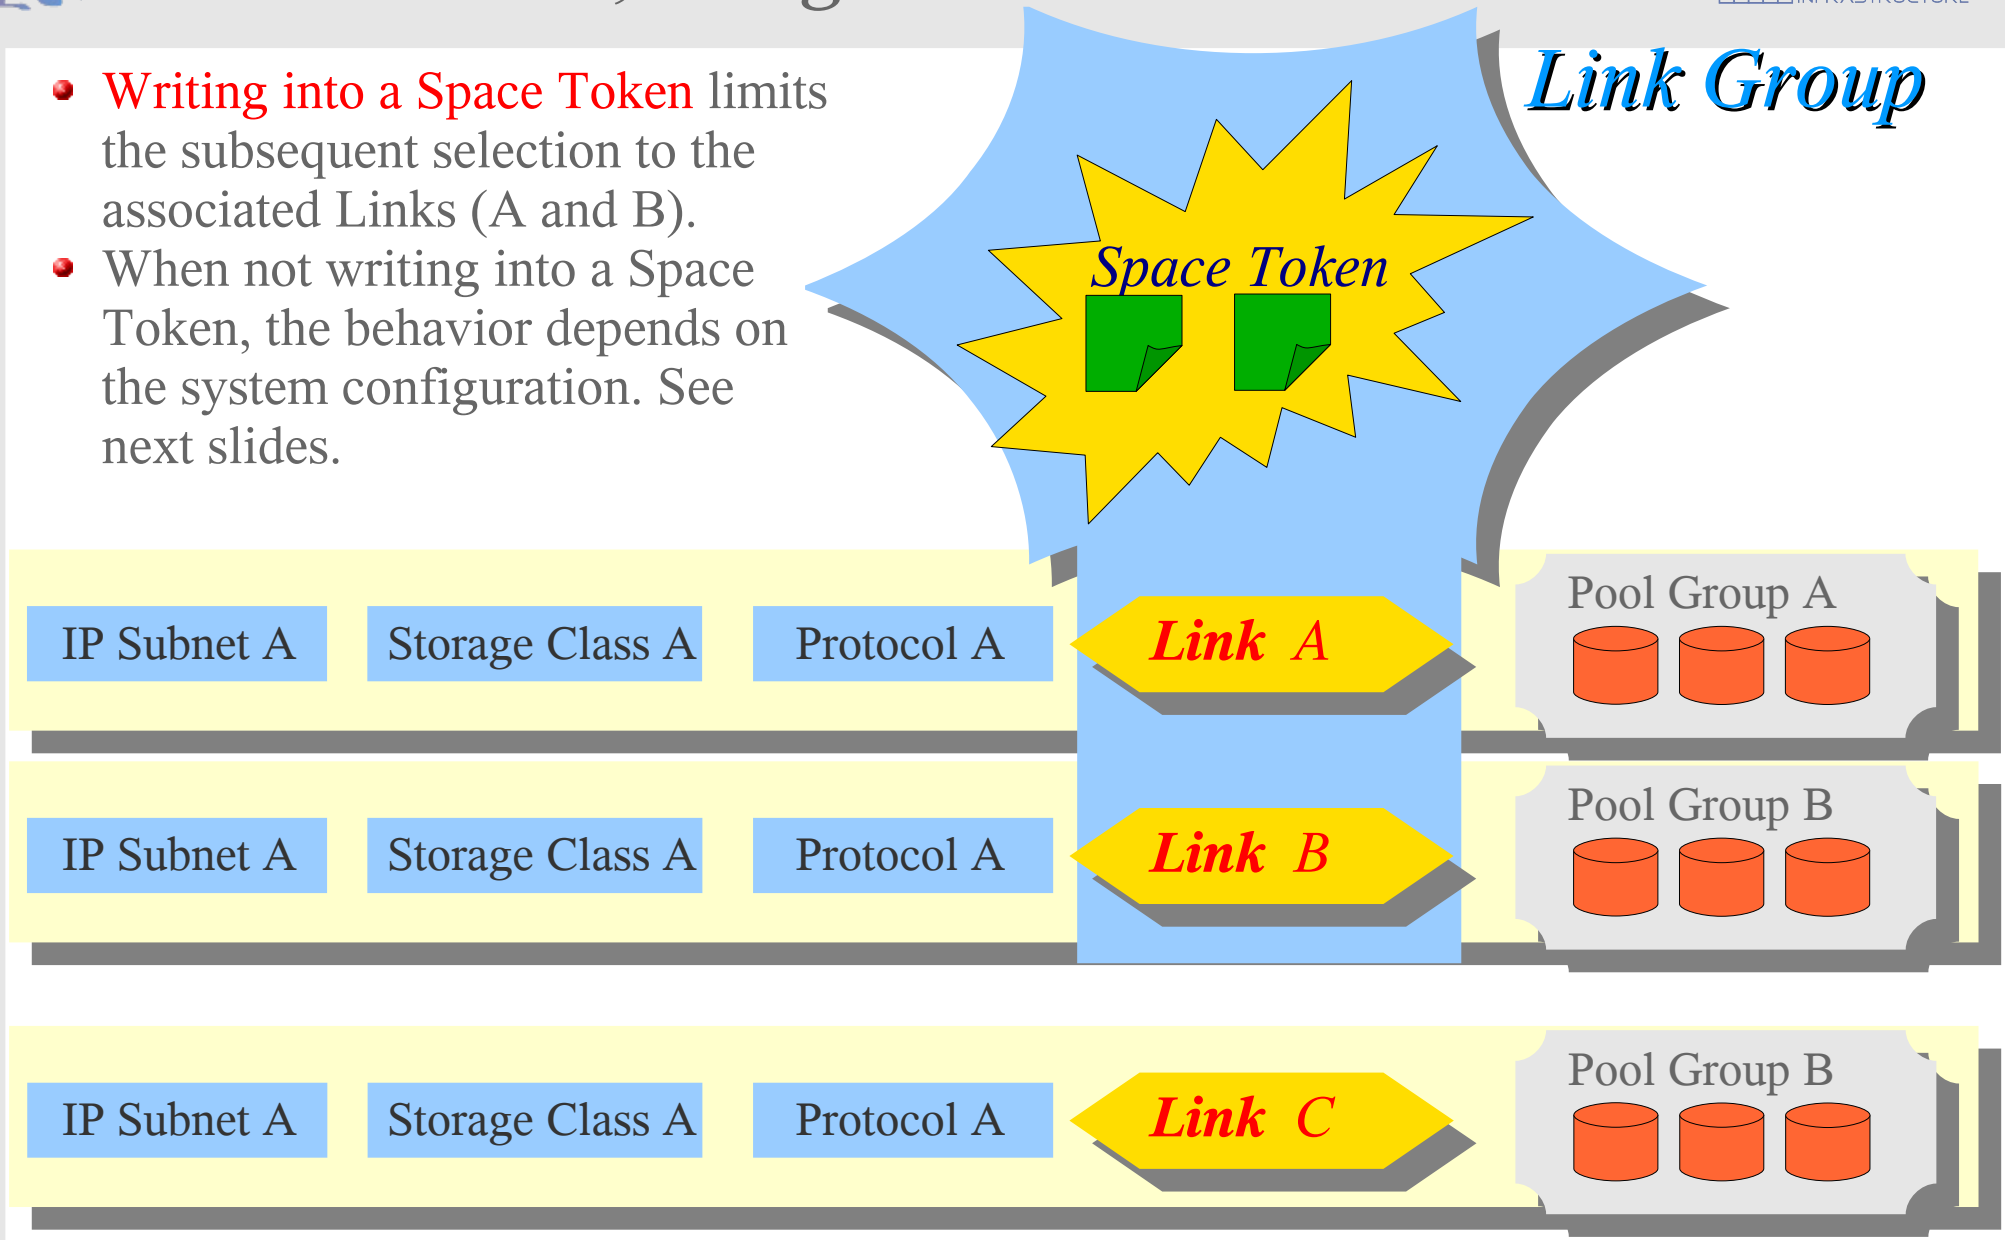

# Implicit Space Reservation

SuGI SUSTAINABLE GRID INFRASTRUCTURE



dCache.ORG

- What happens if you write into dCache with a protocol that does not support space management and only have managed space configured ?
  - SRM 2.2 without Space Token
  - SRM 1.1, gsiFtp
  - we are working gsidCap and gsiXroot
- Implicit Space Reservation resolves this issue with a 'virtual Space Token'.
- The rules for creating an implicit Space Reservation are similar to the rules of creating of Space Tokens.
- Virtual Space Tokens are deleted after file transfer.
- File or copies of the file inherit the Storage Attributes (AL/RP) of the 'virtual Space Token'.

# Implicit versus Explicit Space Reservation

SuGI SUSTAINABLE GRID INFRASTRUCTURE



## Explicit Space Reservation

- SRM 2.2 transfers, specifying a space token.
- SRM 2.2 transfers, where the destination directory has a Token Tag.
- SRM 1.1 transfers, where the destination directory has a Token Tag.

## Implicit Space Reservation (implicit SR = on )

- SRM 1.1 & 2.2 transfers. Access Latency and Retention Policy from
  1. Command line
  2. Directory Tag
  3. Global Defaults in dCacheSetup

dCache.ORG

## Implicit Space Reservation for Non SRM Transfers = on

- gsiFtp. Access Latency and Retention Policy from
  1. Directory Tag
  2. Global Defaults in dCacheSetup

## NO Space Reservation

(gsi)dCap, xroot

GEFÖRDERT VOM
Bundesministerium für Bildung und Forschung
D-GRID

# Where do files go (summary)

SuGI SUSTAINABLE GRID INFRASTRUCTURE

dCache.ORG

dCache.ORG

| | implicitSpaceReservation | |
|---|---|---|
| **SRM 2.2** | OFF | ON |
| Token specified in SRM command | uses specified token | |
| AL/RP specified in SRM command | Uses a Link, not member of a Link-Group | Uses Link-Group which matches the specified AL/RP for implicit space reservation. |
| **SRM 1.1 & 2.2** | | |
| Token as Tag in destination directory *NEW* | uses token in directory | |
| AL/RP as Tags in destination directory | Uses a Link, not member of a Link-Group | Uses Link-Group which matches the specified AL/RP for implicit space reservation. |
| Nothing specified, no tags in directory | Uses a Link, not member of a Link-Group | Uses Link-Group which matches the specified AL/RP for implicit space reservation. (AL/RP from dCacheSetup) |

# Where do files go (summary)



dCache.ORG

| | implicitSpaceReservation = ON | |
|---|---|---|
| **gsiFtp** | **implicitSpaceReservationForNonSrmTransfers** | |
| | OFF | ON |
| Token as Tag in destination directory | uses token in directory | |
| AL/RP as Tags in destination directory | Uses a Link, not member of a Link-Group | Uses Link-Group which matches the specified AL/RP for implicit space reservation. |
| Nothing specified, no tags in directory | Uses a Link, not member of a Link-Group | Uses Link-Group which matches the specified AL/RP for implicit space reservation. (AL/RP from dCacheSetup) |
| **dCap, gsiDCap, xroot** | | |
| unconditional | Uses a Link, not member of a Link-Group | |

# File Status changes and Token Size calculation

SuGI SUSTAINABLE GRID INFRASTRUCTURE



dCache.ORG

- File is written into a T1D1 Space Token.
- File becomes 'precious' on arrival.
- Token Space shrinks by the size of the file.
- File is flushed to Tape, becoming 'cached,sticky'.
  - File can not be removed from disk by the system.
- If the File is removed from the namespace, the file is removed from Disk and Tape and the size of the file is added to the available Space of the Token again.

# File Status changes and Token Size calculation

SuGI SUSTAINABLE GRID INFRASTRUCTURE



dCache.ORG

- File is written into a T1D0 Space Token.
- File becomes 'precious' on arrival.
- Available Token Space shrinks by the size of the file.
- File is flushed to Tape, becoming 'cached'.
- Available Token Space grows again by the size of the file.
- File can be removed from disk by the system if space is needed.
- After the File has been migrated to Tape, there is no association any more between the file and the original Space Token. The File is no longer part of a *Space Token* Space.
- Removing the file from the name-space will not affect the available Token space.

dCache.ORG

*LinkGroups and Space Tokens Authorization is work in progress and is likely to change in the near future.*

- To create a Space Token within a Link Group the requester FQAN needs to match the FQAN specified in the Link Group Authorization File.
- All other operations on Space Token are not checked or protected.
- With the appendum to the MoU, Space Tokens will be protected by ACL's.
- Everybody, who know a space token, may use it for writing.

SuGI
SUSTAINABLE GRID INFRASTRUCTURE

dCache.ORG

dCache.ORG

# *gPlazma Authentication, Authorization*

dCache.ORG

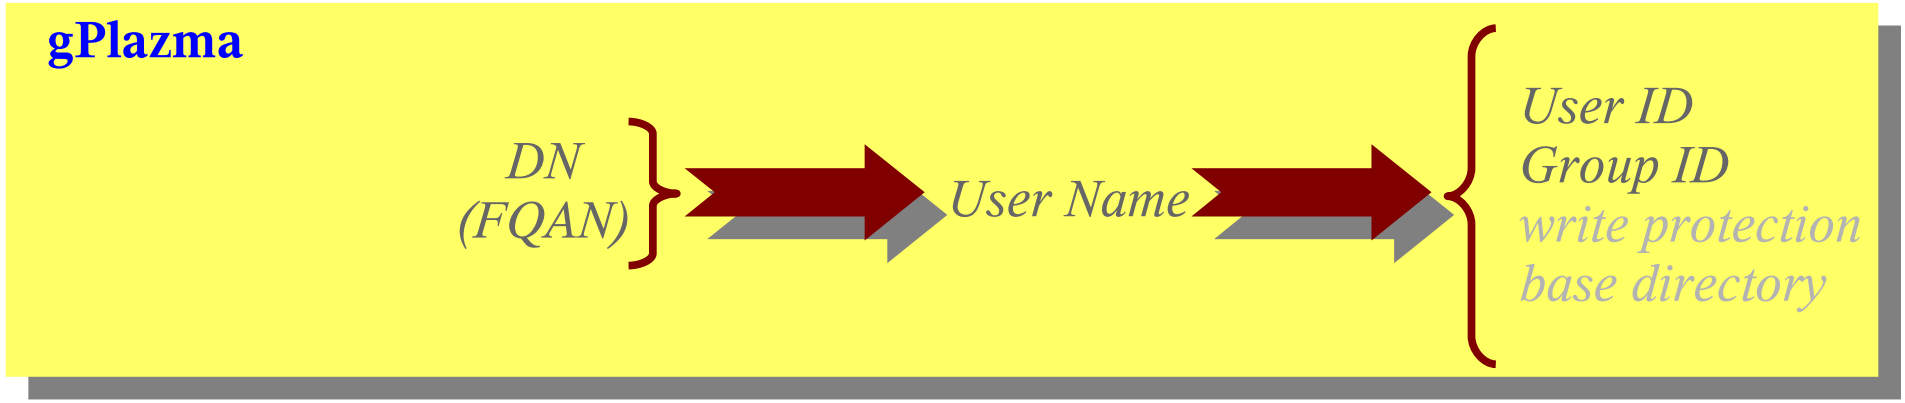


| Plugin Name | | File | | File | |
|---|---|---|---|---|---|
| kpwd | DN | dcache.kpwd | user | dcache.kpwd | uid,gid,... |
| grid-mapfile | DN | grid-mapfile | user | storage-authzdb | uid,gid,... |
| gPlazma-lite-vorole | DN<br>FQAN | grid-vorolemap | user | storage-authzdb | uid,gid,... |

# gPlazma Authentication, Authorization

SuGI SUSTAINABLE GRID INFRASTRUCTURE



- *kpwd and grid-mapfile methods can only map DN's to uid,gid pairs.*
- *A single DN can not get different uid,gid for different FQAN's*

- *The gPlazma lite vorolemap method maps different DN, FQAN pairs to different user names.*
- *Different user names can be mapped to different uid,gid pairs.*

dCache.ORG

# e.g: gPlazma setup for this tutorial

SuGI SUSTAINABLE GRID INFRASTRUCTURE

dCache.ORG

dCache.ORG

***Example for gPlazma lite vorolemap***

*/C=DE/O=DESY/CN=Kermit the frog*
***/desy/Role=NULL*** } ***desy01*** { *UID-1* *GID*

*/C=DE/O=DESY/CN=Kermit the frog*
***/desy/Role=production*** } ***desy02*** { *UID-2* *GID*

SuGI
SUSTAINABLE GRID INFRASTRUCTURE

dCache.ORG

dCache.ORG

# *Further reading*

*www.dCache.ORG*