

# SRMv2.2 in DPM



*Sophie Lemaître ([Sophie.Lemaître@cern.ch](mailto:Sophie.Lemaître@cern.ch))*  
*Jean-Philippe Baud ([Jean-Philippe.Baud@cern.ch](mailto:Jean-Philippe.Baud@cern.ch))*

*SRMv2.2 Workshop*  
*Edinburgh – 13th November 2007*

- **SRMv2.2**
  - No migration needed!
  - Just make sure the SRMv2.2 daemon is running
- **Space tokens vs. User space token descriptions**
- **SRMv2.2 useful tools**
  - dpm-reservespace
  - dpm-updatespace
  - dpm-releasespace

- **DPM architecture**
- **Authorization**
  - In the Namespace
  - On pools
- **SRMv2.2 in DPM**
- **Log files**
- **Admin corner**

- **DPM architecture**
- **Authorization**
  - In the Namespace
  - On pools
- **SRMv2.2 in DPM**
- **Log files**
- **Admin Corner**

`/dpm/cern.ch/home/dteam/my_file`



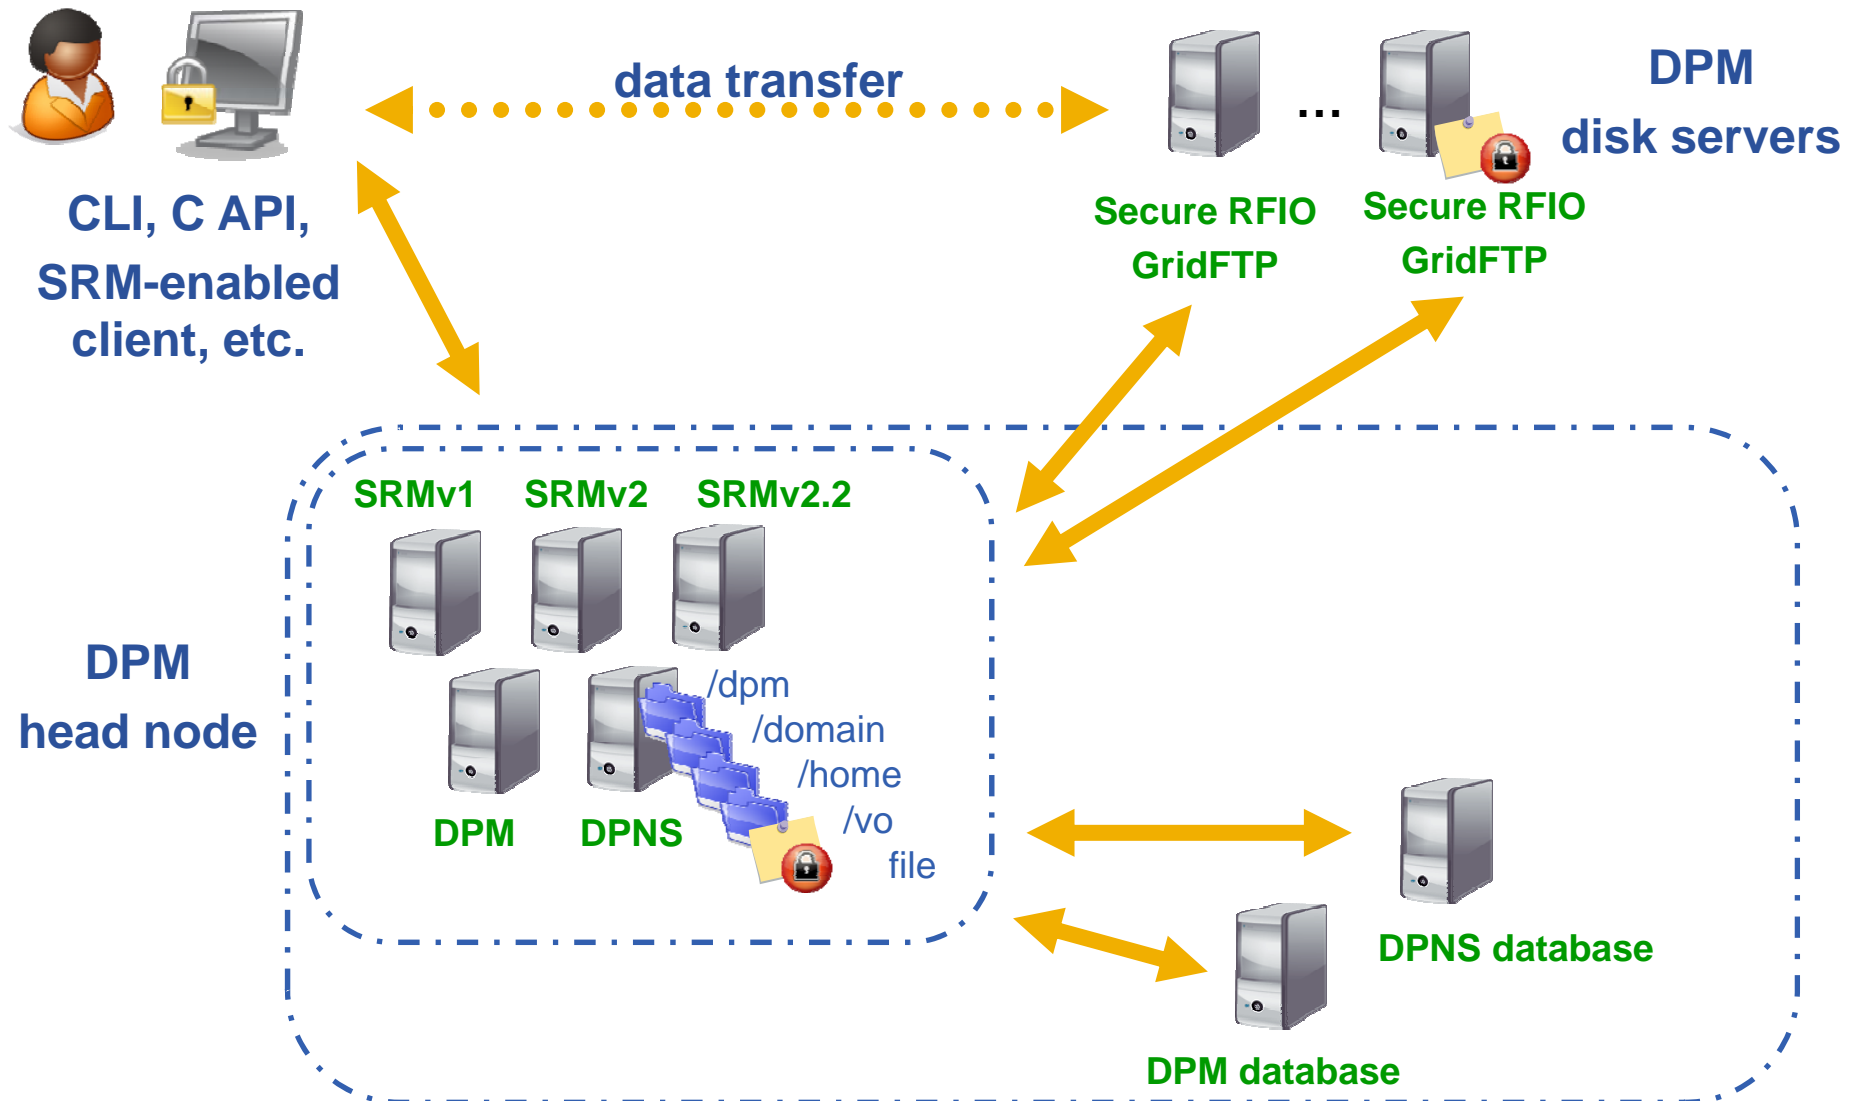
CLI, C API,  
SRM-enabled  
client, etc.



DPM

`disk01.cern.ch:/data01/dteam/2007-06-03/my_file.150.0`

# DPM architecture (details)



- **Ports**

- DPNS: 5010
- DPM: 5015
- SRMv1: 8443
- SRMv2: 8444
- SRMv2.2: 8446

- **Name Server database**
  - Namespace
    - Cns\_file\_metadata
  - Authorization
    - Cns\_groupinfo, Cns\_userinfo
  - Replicas
    - Cns\_file\_replica
  
- **DPM database**
  - DPM configuration
    - dpm\_pool, dpm\_fs
  - All requests (SRM, transfers, etc.)
    - dpm\_pending\_req, dpm\_req, dpm\_get\_filereq, dpm\_put\_filereq
  - Space reservation
    - dpm\_space\_reserv



- **DPM architecture**
- **Authorization**
  - In the Namespace
  - On pools
- **Log files**
- **SRMv2.2 in DPM**
- **Admin corner**



**DPM  
Name Server**

- **Each user's DN**
  - Is mapped to a unique virtual uid
- **Each VOMS group, each VOMS role**
  - Is mapped to a unique virtual gid
- **Virtual uids / gids are created automatically**
  - the first time a given user / group contacts the DPM



(uid1, gid1)  
Ex: (102, 101)



**DPM  
Name Server**



**DB**

## Virtual uids mapping (example)

|  |     |
|--|-----|
| /C=CH/O=CERN/OU=GRID/CN=Sophie Lemaitre 2268 | 101 |
| /C=CH/O=CERN/OU=GRID/CN=Simone Campana 7461  | 102 |

## Virtual gids mapping (example)

|                       |     |
|-----------------------|-----|
| atlas                 | 101 |
| atlas/Role=lcgadmin   | 102 |
| atlas/Role=production | 103 |

```

$ grid-proxy-init
$ voms-proxy-init --vo atlas


=> Simone will be mapped
to (uid, gid) = (102, 101)
    
```



(uid1, gid1)  
Ex: (102, 103, 101)



**DPM  
Name Server**



**DB**

## Virtual uids mapping (example)

|  |     |
|--|-----|
| /C=CH/O=CERN/OU=GRID/CN=Sophie Lemaitre 2268 | 101 |
| /C=CH/O=CERN/OU=GRID/CN=Simone Campana 7461  | 102 |

## Virtual gids mapping (example)

|                       |     |
|-----------------------|-----|
| atlas                 | 101 |
| atlas/Role=lcgadmin   | 102 |
| atlas/Role=production | 103 |

```
$ voms-proxy-init -voms atlas:/atlas/Role=production
```

⇒ Simone will be mapped to  
(uid, gid, ...) = (102, 103, 101)

⇒ Simone still belongs to "atlas"

- **DPM architecture**
- **Authorization**
  - In the Namespace
  - On pools
- **SRMv2.2 in DPM**
- **Log files**
- **Admin corner**

- **DPM terminology**
  - A DPM pool is a set of dedicated filesystems on DPM disk servers
- **By default, pools are generic**
- **Possibility to dedicate a pool to one or several groups**
  - `dpm-addpool --poolname poolA --group alice`
  - `dpm-addpool --poolname poolB --group atlas,cms,lhcb`
- **Easy to add or remove groups**
  - `dpm-modifypool --poolname poolA --group +atlas,-alice`
- **Primary group only for writing**
  - For disk space accounting
  - Who actually uses the space gets the bill...

- **DPM architecture**
- **Authorization**
  - In the Namespace
  - On pools
- **SRMv2.2 in DPM**
- **Log files**
- **Admin corner**

- **No “migration to SRMv2.2” needed**
  - No need to update the DPM namespace
    - All files are accessible via SRMv1, v2 and v2.2
      - *Same logical file names*
      - *Same physical files*
    - The SRM daemons act as front end interfaces, they don't affect the namespace
- **SRMv2.2 daemon needs to run**
  - Automatically done by YAIM
  - You are probably already running SRMv2.2!



- **File Storage Types**
  - Volatile
  - Permanent
  - Any (default)
- **Some client tools use “volatile” type, others “permanent”**
  - Best to configure your pools as “-” (default)
    - `dpm-addpool --poolname MyPool --def_filesize 200M --s_type “-”`

- **Retention policies**
  - Given quality of disks, admin defines quality of service
  - Replica (default), Output, Custodial
- **Example**
  - `dpm-modifypool --poolname MyPool --ret_policy R`

- **Space token**
  - A UUID created internally by the DPM
    - Ex: `fe869590-b771-4002-b11a-8e7430d72911`
- **User space token description**
  - A case sensitive description
  - Provided by a user or the DPM admin
    - Ex: `myspace`, `atlas_prod`, `TutorialSpace`
- **No relationship between**
  - spaces, and
  - the DPM namespace

- **Static space reservation (admin)**

- **No limitation on duration and size of space reservation**

```
$ dpm-reservespace --gspace 20G --lifetime Inf --group atlas --token_desc Atlas_ESD
```

```
$ dpm-reservespace --gspace 100M --lifetime 1h --group dteam/Role=lcgadmin --token_desc LcgAd
```

```
$ dpm-updatespace --token_desc myspace --gspace 5G
```

```
$ dpm-releasespace --token_desc myspace
```

- **Dynamic space reservation (user)**

- **Defined by user on request**

- `dpm-reservespace`, `dpm-updatespace`, `dpm-releasespace`

- `srmReserveSpace`

- **Limitation on duration and size of space reservation**



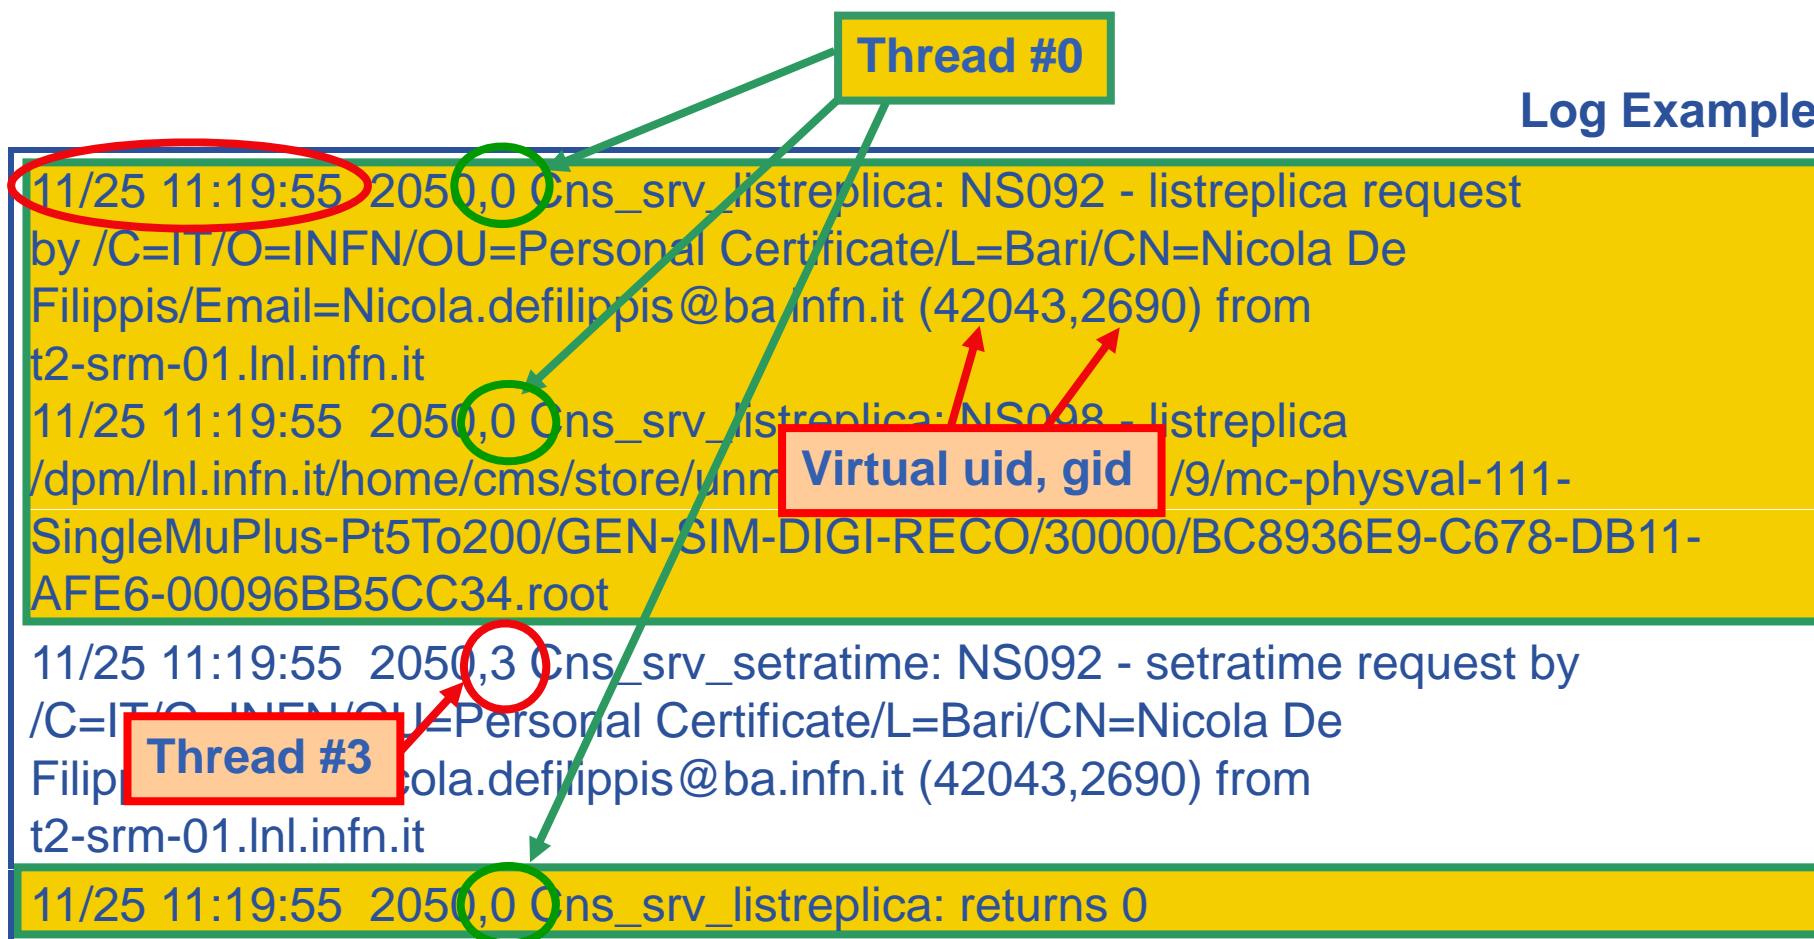
**Features released since DPM 1.6.3**

- **DPM architecture**
- **Authorization**
  - In the Namespace
  - On pools
- **SRMv2.2 in DPM**
- **Log files**
- **Admin corner**

- **Example: DPNS log**

- /var/log/dpns/log

## Log Example



The diagram illustrates a sequence of log entries from a DPNS log file. A yellow box highlights the first three entries, which are associated with Thread #0. A red box highlights the fourth entry, associated with Thread #3. A red box labeled 'Virtual uid, gid' points to the '2050,0' values in the second and third entries. Green circles and arrows connect the thread labels to their respective log entries. Red circles highlight the timestamps '11/25 11:19:55' in the first and fourth entries.

```

11/25 11:19:55 2050,0 Cns_srv_listreplica: NS092 - listreplica request
by /C=IT/O=INFN/OU=Personal Certificate/L=Bari/CN=Nicola De
Filippis/Email=Nicola.defilippis@ba.infn.it (42043,2690) from
t2-srm-01.lnl.infn.it
11/25 11:19:55 2050,0 Cns_srv_listreplica: NS098 - listreplica
/dpm/lnl.infn.it/home/cms/store/unm Virtual uid, gid /9/mc-physval-111-
SingleMuPlus-Pt5To200/GEN-SIM-DIGI-RECO/30000/BC8936E9-C678-DB11-
AFE6-00096BB5CC34.root
11/25 11:19:55 2050,3 Cns_srv_setratime: NS092 - setratime request by
/C=IT/O=INFN/OU=Personal Certificate/L=Bari/CN=Nicola De
Filippis/Email=Nicola.defilippis@ba.infn.it (42043,2690) from
t2-srm-01.lnl.infn.it
11/25 11:19:55 2050,0 Cns_srv_listreplica: returns 0
    
```

## /var/log/srmv2.2/log example

```

11/05 05:06:16 25696,0 PrepareToPut: request by /C=IT/O=INFN/OU=Personal Certificate/L=Pisa/CN=Flavia
Donno/Email=flavia.donno@pi.infn.it from lxdev25.cern.ch
11/05 05:06:16 25696,0 PrepareToPut: SRM98 - PrepareToPut b2c7ba01-ca72-4c9a-9714-312e70db7f22
11/05 05:06:16 25696,0 PrepareToPut: SRM98 - PrepareToPut 0
srm://lxdpm102.cern.ch:8446/srm/managerv2?SFN=/dpm/cern.ch/home/dteam/20071105-050615-1865/f0
11/05 05:06:16 25696,0 PrepareToPut: SRM98 - PrepareToPut 1
srm://lxdpm102.cern.ch:8446/srm/managerv2?SFN=/dpm/cern.ch/home/dteam/20071105-050615-1865/f1
11/05 05:06:16 25696,0 PrepareToPut: returns 0, statusCode=SRM_REQUEST_QUEUED
11/05 05:06:16 25696,0 StatusOfPutRequest: request by /C=IT/O=INFN/OU=Personal Certificate/L=Pisa/CN=Flavia
Donno/Email=flavia.donno@pi.infn.it from lxdev25.cern.ch
11/05 05:06:16 25696,0 StatusOfPutRequest: SRM98 - StatusOfPutRequest b2c7ba01-ca72-4c9a-9714-312e70db7f22
11/05 05:06:16 25696,0 StatusOfPutRequest: SRM98 - StatusOfPutRequest
srm://lxdpm102.cern.ch:8446/srm/managerv2?SFN=/dpm/cern.ch/home/dteam/20071105-050615-1865/f0
11/05 05:06:16 25696,0 StatusOfPutRequest: SRM98 - StatusOfPutRequest
srm://lxdpm102.cern.ch:8446/srm/managerv2?SFN=/dpm/cern.ch/home/dteam/20071105-050615-1865/f1
11/05 05:06:16 25696,0 StatusOfPutRequest: returns 0, statusCode=SRM_REQUEST_INPROGRESS
11/05 05:09:17 25696,0 StatusOfPutRequest: request by /C=IT/O=INFN/OU=Personal Certificate/L=Pisa/CN=Flavia
Donno/Email=flavia.donno@pi.infn.it from lxdev25.cern.ch
11/05 05:09:17 25696,0 StatusOfPutRequest: SRM98 - StatusOfPutRequest b2c7ba01-ca72-4c9a-9714-312e70db7f22
11/05 05:09:17 25696,0 StatusOfPutRequest: SRM98 - StatusOfPutRequest
srm://lxdpm102.cern.ch:8446/srm/managerv2?SFN=/dpm/cern.ch/home/dteam/20071105-050615-1865/f0
11/05 05:09:17 25696,0 StatusOfPutRequest: SRM98 - StatusOfPutRequest
srm://lxdpm102.cern.ch:8446/srm/managerv2?SFN=/dpm/cern.ch/home/dteam/20071105-050615-1865/f1
11/05 05:09:17 25696,0 StatusOfPutRequest: returns 0, statusCode=SRM_SUCCESS
  
```

**SRMv2.2 calls received**

**SRMv2.2 status codes**

- **Useful tools**
  - DPM check consistency tool
  - DPM request database cleanup tool
  
- **No need to restart the daemons when the host certificate changes**
  
- **SLC3 to SLC4 migration**
  - OS reinstallation, not an update
  - Backup the MySQL “cns\_db” and “dpm\_db” databases
  - Need to backup:
    - Log files
    - Configuration files (NSCONFIG, DPMCONFIG)
    - etc.
  - Follow this procedure:
    - <https://twiki.cern.ch/twiki/bin/view/LCG/DpmMigratingFromSL3TOSL4>



- **On the TO DO list**
  - Blacklisting via LCAS
  - ChangeSpaceForFiles
    - To be able to move a file from one space to another
  - Quotas

- **Main DPM entry point**
  - <https://twiki.cern.ch/twiki/bin/view/LCG/DataManagementTop>
- **DPM Admin Guide**
  - <https://twiki.cern.ch/twiki/bin/view/LCG/DpmAdminGuide>
- **Troubleshooting**
  - <https://twiki.cern.ch/twiki/bin/view/LCG/LfcTroubleshooting>

- **Global Grid User Support (GGUS)**
  - [helpdesk@ggus.org](mailto:helpdesk@ggus.org)
    - Your ROC will first help
    - If needed, DPM experts
- **Users and administrators forum**
  - [dpm-users-forum@cern.ch](mailto:dpm-users-forum@cern.ch)
    - Place to ask questions
    - Place to share your experience with other users/admins
    - To subscribe, go to <http://simba.cern.ch>

- **SRMv2.2**
  - No migration needed!
  - Just make sure the SRMv2.2 daemon is running
- **Space tokens vs. User space token descriptions**
- **SRMv2.2 useful tools**
  - dpm-reservespace
  - dpm-updatespace
  - dpm-releasespace

# Questions?

*Sophie Lemaitre ([Sophie.Lemaitre@cern.ch](mailto:Sophie.Lemaitre@cern.ch))*  
*Jean-Philippe Baud ([Jean-Philippe.Baud@cern.ch](mailto:Jean-Philippe.Baud@cern.ch))*