

---

# Status of Data Management Services

Michael Ernst  
DESY

WLCG Service Coordination Meeting - CERN

---

# Disclaimer

- Information provided in this status report is obtained from developers responsible for the respective projects
- Primary topic is SRM
- SE Types covered in this status report
  - CASTOR (O. Barring, Shaun De Witt, Sebastien Ponce)
  - SRM/dCache (R. Kennedy, T. Perelmutov, J. Bakken, A. Rana)
  - DPM (J.-P. Baud)

# Enhancements required by Experiments

- Timescales – Experiment’s view
  - “Necessary to expose service well before start of SC4”
  - Five Priority items
    - (Explicit) Pin/Unpin
    - Relative Paths in SURLs
    - Permission functions – SRM Integration w/ VOMS
    - Directory functions
    - Global Space Reservation

# Explicit Pin/Unpin

- ❑ Issues associated with Explicit Pin/Unpin
  - ❑ Pins have a lifetime
    - Operational issue when used by average user who could interfere w/ the rest of community, especially in an HSM environment w/ Migration capabilities
  - ❑ Not part of the SRMv2.1 Specification
    - Method to support pin/unpin functionality though
      - ❑ srmPrepareToGet/Put has a lifetime field
      - ❑ Lifetime can be extended
      - ❑ Method to release pin (per job)
      - ❑ max. lifetime part of server configuration

# Explicit Pin/Unpin Implementation Status

- Castor
  - Integrated w/ Stager API
    - Currently being tested
    - Availability expected before Christmas
- dCache
  - Initial v2 release ready on 15 January
    - Includes data transfer functions
      - srmPrepareToGet/Put (used for pin/unpin)
      - srmCopy
- DPM
  - Implemented and tested since early 2005
  - Time interval can be specified (prepareToGet, separate command) or admin specified defaults (pool specific) can be used

# Relative Path

## ② Relative Path in SURLs

- Requirement as we understand it
  - (Re-)use the same file path beyond the SE access point  
srm://server:port/castor|pnfs|dpm/domain/VO/file-path
  - Only “file-path” provides valuable information to applications, the rest is concerned w/ site and implementation specifics the experiments are not interested in
  
- Developers have mixed feelings (raised by SRM/dCache team)
  - SRM URLs are supposed to be freely transferable between users and different roles assumed by the same user, which is not possible w / relative paths.
  - How can file locations have any meanings with relative paths?
  - Relative paths can only be interpreted knowing the detailed configuration of an SE as a function of the user and associated role

# Relative Path – Implementation Status

- Castor
  - Partial Implementation (w/o VOMS support) available before Christmas
  - Implementation incl. VOMS support by end January
- SRM/dCache
  - Since Relative Paths don't fit the concept - there are no plans to implement them:
  - SRM URLs are supposed to be freely transferable between users (cannot be done w/ relative paths)
- DPM
  - No concrete plans at the moment – will not be available before end January

# “VOMS Integration”

- ③ Permission functions: All experiments would like the permissions to be based on roles & DN's. SRM should be integrated with VOMS
  - ❑ Issues
    - Implementation of permission functions depends on the support from the underlying storage system (i.e. support for ACLs vs. UNIX permissions only)
  - ❑ Implementation Status
    - CASTOR
      - ❑ Availability announced for end of January
    - DPM
      - ❑ VOMS integration is done and tested
      - ❑ Code cleanup and compatibility testing (grid-proxy-init/voms-proxy-init)
      - ❑ ACLs based on virtual IDs associated w/ DNs and VOMS roles
      - ❑ Needs more work on Admin Interface

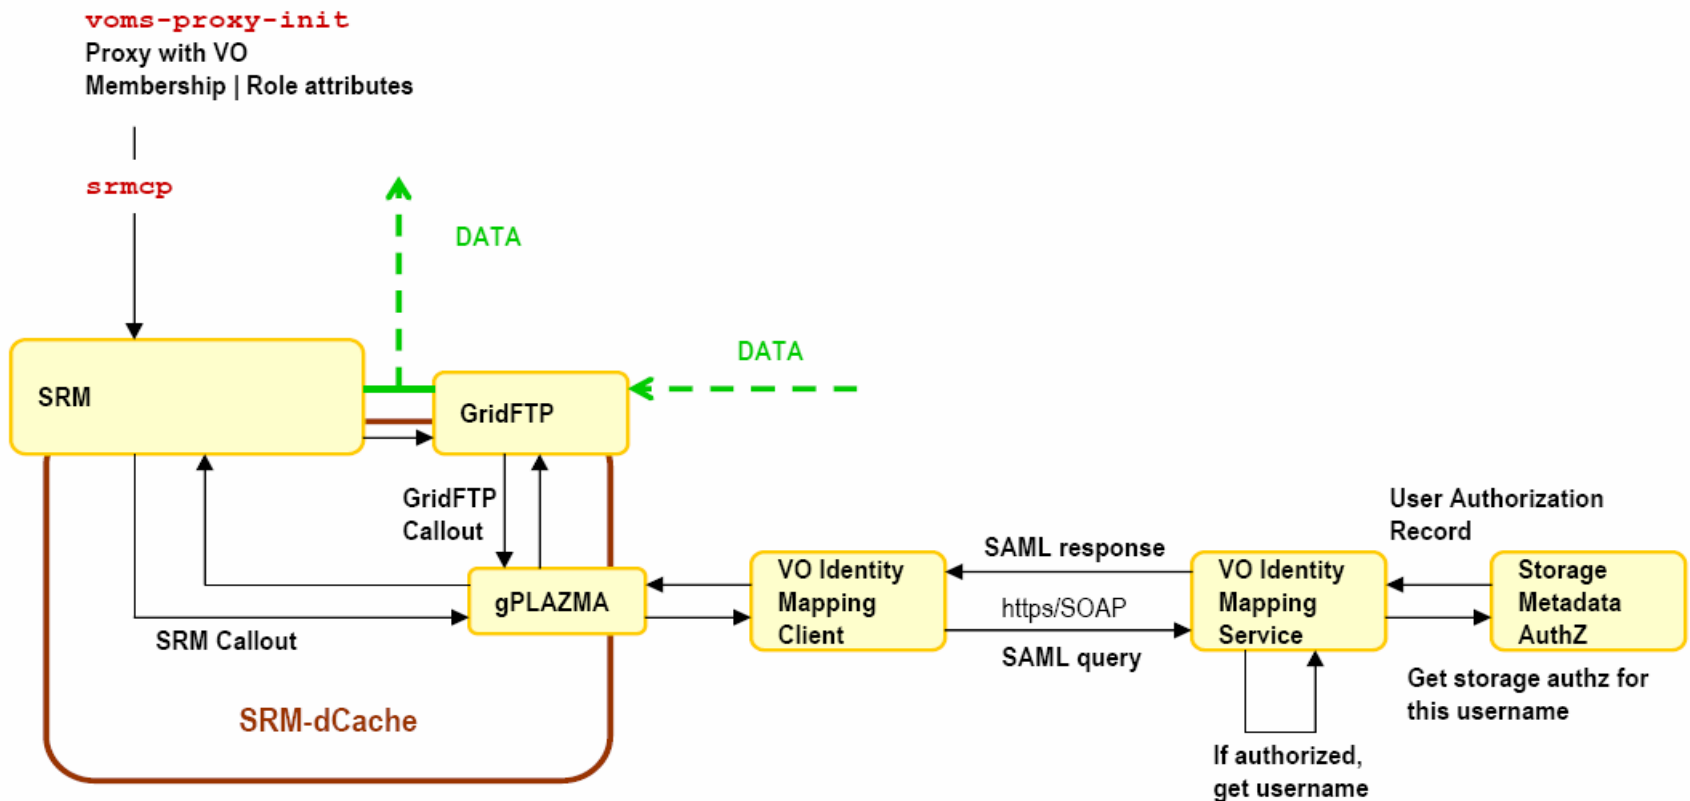


# Implementation Status SRM/dCache

Source: Abhishek Rana (UCSD), Timur Perelmutov (FNAL)

Development as part of the OSG development program

## An example use case: Fermilab VO-role mapping AuthZ



# SRM/dCache

## ■ Role based Authentication/Authorization Services with gPLAZMA

gPLAZMA authenticates.

Storage Authz Service contacts GUMS and gPLAZMA Storage Metadata Service.

GUMS translates {DN, Membership, Role} to Username.

System optionally translates Username to site-wide {UID, GID}.

gPLAZMA Storage Metadata Service translates Username to Storage-privilege Set.

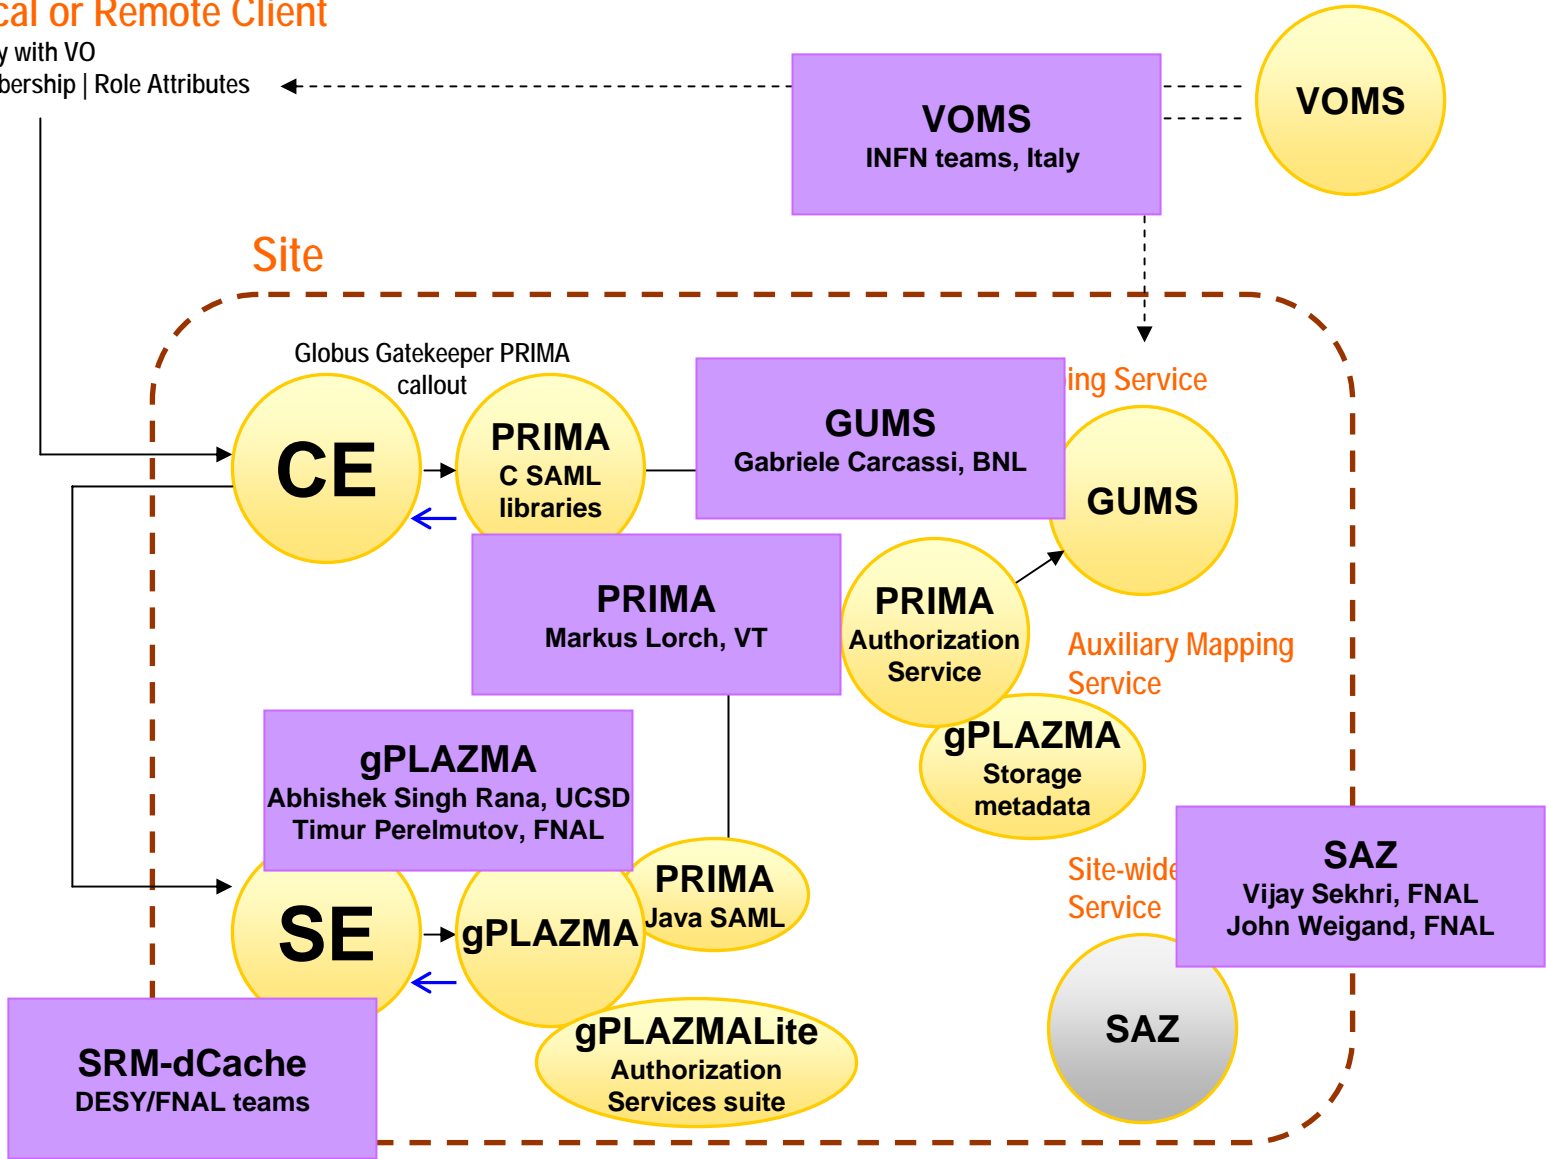
Storage-privilege Set is {UID, GID, permitted storage area, R/W permissions}.

Storage-privilege Set is User-level ACL governed by {DN, Membership, Role} .



# Local or Remote Client

Proxy with VO  
Membership | Role Attributes



---

# Implementation Status

- Packaging / Distribution available by end Jan
- Need to integrate gPLAZMA into dCap door
  - No firm date w.r.t. availability given

---

# Directory Functions

- Implementation Status

- CASTOR

- All implemented except srmLs and srmCp that will be made available by mid January

- SRM/dCache

- Initial v2 release (beta on 15 Jan) will support ls, rm, mv etc.

- DPM

- All implemented since first release except srmLs does not support recursive listings nor offsets

# Global Space Reservation - Implementation Status

- **Castor**
  - Fully available
  
- **SRM/dCache**
  - Tentative date for alpha release is 15 March
  
- **DPM**
  - Current implementation just returns token, which at the moment is not used in PrepareToGet/Put methods
  - A full implementation will not be available before March 2006

# Interoperability Castor-dCache-DPM-?

- ❑ Issues include
  - Specifications leave room for interpretation
  - srmCp interoperability problems in v1.1
- ❑ SRMv2 test suite by Gilbert Grosdidier
  - LBL suite for DRM part of VDT but not suitable
  - Flavia Donno will create RPMs
- ❑ Instances for Interoperability tests available for Castor and DPM (?)



# Interoperability issues dCache ⇔ DPM

- ❑ Based on current/recent distributions
- ❑ Problems observed on transfers from dCache => DPM
  - srmCp push goes at good rate
  - Direct gridFTP to gridFTP door goes at good rate
  - 3<sup>rd</sup> party gridFTP get/put choked at 250 kB/s
  - FTS currently uses gridFTP get/put (basic version of srmCp client exist)
    - Invoke srmGet and srmPut, receive 2 TURLs and make 3<sup>rd</sup> party gridFTP between source and destination host
  - Problem appears to be independent from FTS code
    - Manual invocation of commands leading to same result

# SRMv2 / FTS

- ❑ Main Issue w.r.t. SC-4 is Timeline
  - FTS developers need to integrate after SRMv2 implementations are available to them
- ❑ Other worries
  - Differences in Implementations (Castor, dCache, DPM,?)
  - srmCp is important for Performance
  - Cross SRMv2 ⇔ SRMv1 support
  - Need to understand behavior of new features (i.e. Space Reservation etc.)

# Data Integrity for Data in Transit

- Timur (FNAL) has made proposal to enable integrity verification in SRM servers when clients are writing into Storage Systems
  - 2 optional fields in PrepareToPut
    - Checksum type and checksum value
  - Parameters could be optionally set by the client
  - Server calculates checksum on data transferred into SRM managed storage via any of the supported protocol
  - Server can reject transfer in case locally calculated checksum does not match the transferred checksum
  - Include checksum type and value in srmPrepareToPut file status structure
  - In FNAL SRM implementation implement integrity verification functionality as part of the gridFTP v2 specification (CKSM, SCKSM) as it is defined in

<https://forge.gridforum.org/projects/gridftp-wg/document/draft-ggf-gridftp-v2protocol-2.pdf/en/1>

# Summary

- Good progress made w/ all three implementations
  - Castor
    - Transfer and directory functions ready by end January
    - Global Space Reservation fully available today
  - SRM/dCache
    - Transfer and directory functions ready by end January (Beta)
    - Global Space Reservation by mid March (alpha)
  - DPM
    - Transfer and directory functions ready by end January
    - Space Reservation not before March (optimistic date)
  - SRMv2 test suite will help sort out Interop issues
  - See no showstopper but schedule is tight