GGF-Grid Interoperability Now (GIN)
SRM Interoperability
Testing Results

September 10, 2006

Alex Sim, Vijaya Natarajan
Scientific Data Management Research Group
Lawrence Berkeley National Laboratory

http://sdm.lbl.gov/srm-tester/ggf18.html

Abstract:

Storage Resource Managers (SRMs) are based on a common interface specification (http://sdm.lbl.gov/srm-wg). SRMs can have different implementations and services for the underlying storage systems. The purpose of the GIN SRM testing is to check the compatibility and interoperability of SRMs according to the specification.

The tests include the following:

- read access to a file in a remote grid storage managed by SRM
- write access to a registered user account in a remote grid storage managed by SRM
- a file replication for a registered user between two independent storage systems
- space reservation and write access to the reserved space for a registered user in a remote grid storage by managed SRM (for SRM v2.2)

Approach:

Participating storage sites that support SRM access will be tested with the SRM-Tester for various operations, according to the SRM specification. In case of failures, other client tools may be used to test the SRM servers further.

SRM-Tester:

The SRM-Tester tests Storage Sites managed by SRMs. It is an independent implementation of SRM specification for the purpose of SRM managed storage testing. It is designed to:

- test the conformity of the SRM server interface according to the SRM spec v1.1 and v2.2
- test the compatibility and interoperability of the SRM servers according to the SRM specifications
- test the following operations for SRM v1.1:
  - ping, put, get, advisory delete, copy
- test the following operations For SRM v2.2:
  - ping, put, get, bring online, release file, copy, reserve space, release space, remove
- test the following file transfer protocols:
  - gsiftp, ftp, http, https
- publish the testing results to a web site.

The test results for GGF 18 GIN-Data SRM testing are published at http://sdm.lbl.gov/srm-tester/ggf18.html.

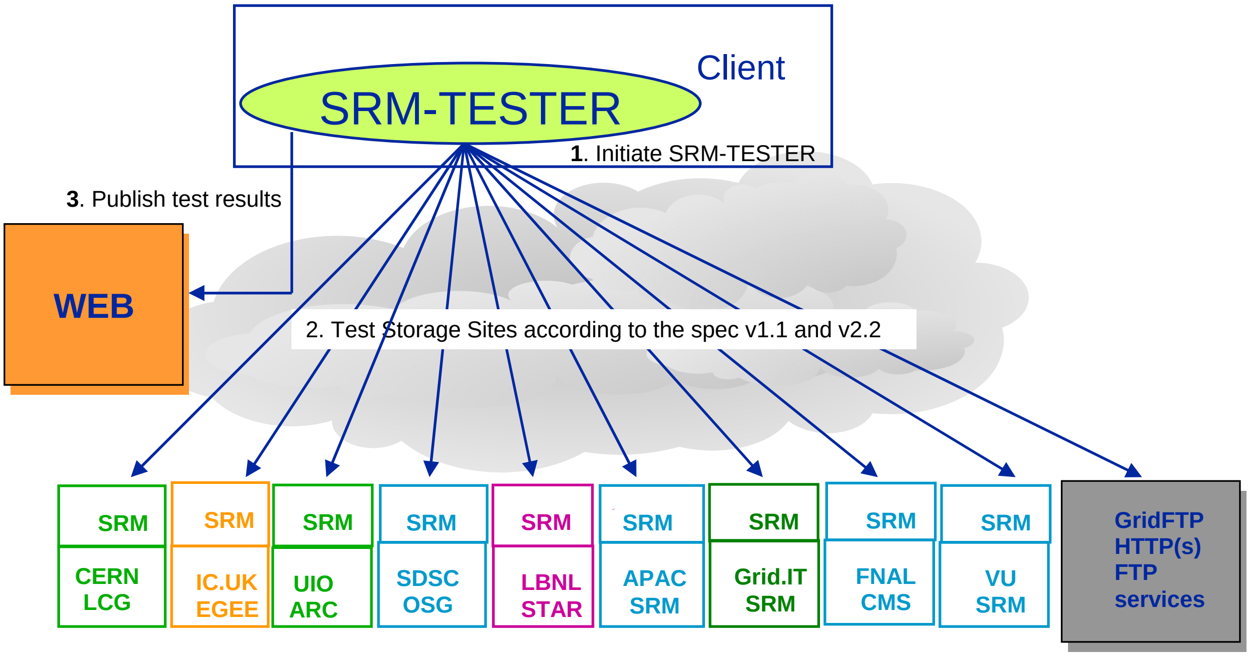

Participating Sites:

There are 9 sites that participated in the GIN-data SRM island testing. 6 SRM V1.1 and 6 SRM V2.2 servers were tested. SRMs at 2 sites were not ready at the time of the testing. Compared to last GGF-17, two additional site and 5 more SRM servers have participated.

- APAC
  - SRM is not ready due to undergoing SFT tests
- ARC
  - srm://grid.uio.no:58000
- CMS
  - srm://fledgling05.fnal.gov:8443
  - srm://fledgling06.fnal.gov:8443 (SRM v2.2 dCache)
- EGEE
  - srm://gfe02.hep.ph.ic.ac.uk:8443
- Grid.it
  - SRM is not installed publicly and not ready for testing
- LCG/EGEE
  - srm://lxdpm01.cern.ch:8443
  - srm://lxdpm01.cern.ch:8446 (SRM v2.2 DPM)
  - srm://lxb1389.cern.ch:8442 (SRM v2.2 CASTOR)
  - srm://castor300.ads.rl.ac.uk:8443 (SRM v2.2 CASTOR)
- OSG
  - srm://t2data2.t2.ucsd.edu:8443
- STAR
  - srm://dmx09.lbl.gov:6175
  - srm://dmx09.lbl.gov:8443 (SRM v2.2 DRM)
- VU
  - srm://vanderbilt.edu:8443 (SRM v2.2 L-Store)

Information requested for SRMs:

1. SRM end point:
   e.g. in case of an SRM at SDSC, it looks like
   srm://t2data2.t2.sdsc.edu:8443//srm/managerv1
2. Version of SRM running on the above end point:
   e.g. v1.1 or v2.2
3. writable data directory for GIN-VO
   e.g. /data2/gin
4. if "delete" on the file in SRM managed cache is available
   e.g. true, false
5. if "overwrite" on the existing file in SRM managed cache is available
   e.g. true, false
6. if SRM managed cache is "read-only" storage. In other words, no writing is possible.
   in case of read-only cache, the URL of the file that we can pull out
   e.g. srm://t2data2.t2.sdsc.edu:8443//srm/managerv1?SFN=/data2/gin/mytest.file
7. supported transfer protocol list
   e.g. http, gsiftp, dcap

From the above information, we could compose an URL and proceed with tests:
srm://t2data2.t2.sdsc.edu:8443//srm/managerv1?SFN=/data2/gin/my_put_test.file

Testing Description for SRM v1.1

1. "ping" the target SRM server to check if the server is running
   a. GSI security may or may not be enforced.
2. "put" (write) a file into the SRM managed, VO specific directory from client's local disk
   a. checks the transfer protocol during the file transfers
3. "get" (read) a file from SRM managed, VO specific directory to the client's local disk
   a. Requests a file from the SRM managed, VO specific directory
   b. When the file is ready, SRM server returns Transfer URL (TURL) with transfer protocol
   c. File transfers from the TURL to the client's local disk is tested
4. "unpin" (release) files
5. "advisory delete" (remove) files
6. "copy" a file from one SRM to another SRM
   a. SRMs are in two separate independent data grids
7. "copy" a file from one remote gridftp server to an SRM

Testing Description for SRM v2.2

SRM v2.2 adds a few functionalities on the previous SRM versions. SRM v2.2 is a new interface, and its implementation and deployment efforts are very active. All current implementations are still in-progress, and the test results are only showing the current status.

1. "ping" the target SRM by "srmPing"
2. "put" (write) a file into SRM managed, VO specific directory from client's local disk
   a. Checks the transfer protocol during the file transfers
   b. Tests srmPutDone() after file transfer is completed
3. "get" (read) a file from SRM managed, VO specific directory to client's local disk
   a. Requests a file from the SRM managed, VO specific directory

    b. When the file is ready, SRM server returns Transfer URL (TURL) with transfer protocol
    c. File transfers from the TURL to the client's local disk is tested

4. "release" files
    a. releases the file after file transfer is completed (following "get" operation)
5. "bring online" a file from SRM managed, VO specific directory to SRM managed online space
    a. Requests a file from the SRM managed, VO specific directory
    b. When the file is ready online, SRM server returns the proper status information
6. "non-recursive directory listing" (ls)
    a. browses a file or a directory in the SRM managed, VO specific space
7. "copy" a file
    a. Request to "replicate" a file from one SRM to another SRM. SRMs are in two separate independent data grids.
    b. Request to "replicate" a file from a remote gsiftp server to an SRM.
8. "reserve a space"
    a. requests to reserve a volatile or durable space under the SRM for an amount of byte size and gets a corresponding space token
9. "status" of the reserved space
    a. checks the status of the space with the space token received from the above space reservation
10. "put" (write) a file into the above reserved space with the space token
    a. Similar to "put" operation, but a specific target space token is used.
11. "get" (read) a file into the above reserved space with the space token
    a. Similar to "get" operation, but a specific target space token is used.
12. "bring online" a file into the above reserved space with the space token
    a. Similar to "bring online" operation, but a specific target space token is used.
13. "release a space"
    a. releases the above reserved space with the space token

Test results for SRM v1.1

Table 1 : Test results for SRM v1.1

	ping	Get	put	advisory delete
ARC (UIO.NO)	pass	Pass	fail	fail
LCG/EGEE (CERN.CH)	pass	Pass	pass	pass
CMS (FNAL.GOV)	pass	Pass	pass	pass
STAR (LBL.GOV)	pass	Pass	pass	pass
OSG (SDSC.EDU)	pass	Pass	pass	pass
EGEE (IC.UK)	pass	Pass	pass	pass

- SRM at ARC fails for operations put
    - SRM-Tester supports HTTP POST method to put the file, but ARC uses HTTP PUT method.
    - Another SRM client, SRMCP also failed with the same reason.
    - NGCP (NorduGrid ARC standalone v0.5.46) works for operations except copy
        - ngcp does not handle different formats of certificates yet, such as service certificates

COPY between SRMs in PULL mode

- SRM-Tester sends a request to the target SRM to "pull" a file from the source SRM
- After a copy request is submitted to the target SRM, communication thereafter is between the two SRMs.

- SRM-Tester checks the status of the copy request.

Table 2 : Test results for COPY between SRMs in PULL mode

Target SRM → Source SRM ↓	ARC (UIO.NO)	LCG/EGEE (CERN.CH)	CMS (FNAL.GOV)	STAR (LBL.GOV)	OSG (SDSC.EDU)	EGEE (IC.UK)
ARC (UIO.NO)		N.A.	fail	pass	N.A.	fail
LCG/EGEE (CERN.CH)	fail		pass	pass	N.A.	pass
CMS (FNAL.GOV)	fail	N.A.		pass	N.A.	pass
STAR (LBL.GOV)	fail	N.A.	pass		N.A.	pass
OSG (SDSC.EDU)	fail	N.A.	pass	pass		pass
EGEE (IC.UK)	fail	N.A.	pass	pass	N.A.	

- SRM at CERN (LCG/EGEE)
  - The request goes through the interface, and the status shows request queued and in progress.
  - We were informed that actual copy is not implemented yet
- SRM at SDSC (OSG)
  - The request goes through the interface, and the status shows request queued and in progress, with no further explanation
  - We were informed that only “push” mode into the SRM at SDSC works because the underlying storage resides inside the private network.
- SRM at UIO (ARC)
  - The request failed and we have reported the problem and they are working on it.

COPY between SRMs in PUSH mode

- SRM-Tester sends a request to the source SRM to “push” a file from the source SRM to the target SRM
- After a copy request is submitted to the source SRM, communication thereafter is between the two SRMs.
- SRM-Tester checks the status of the copy request.

Table 3 : Test results for COPY between SRMs in PUSH mode

Target SRM → Source SRM ↓	ARC (UIO.NO)	LCG/EGEE (CERN.CH)	CMS (FNAL.GOV)	STAR (LBL.GOV)	OSG (SDSC.EDU)	EGEE (IC.UK)
ARC (UIO.NO)		fail	fail	fail	fail	fail
LCG/EGEE (CERN.CH)	N.A		N.A	N.A	N.A	N.A
CMS (FNAL.GOV)	fail	pass		pass	fail	pass
STAR (LBL.GOV)	fail	pass	pass		pass	pass
OSG (SDSC.EDU)	N.A	N.A	N.A	N.A		N.A
EGEE (IC.UK)	fail	pass	pass	fail	pass	

- SRM at CERN (LCG/EGEE)
  - The request goes through the interface, and the status shows request queued and in progress.
  - We were informed that actual copy is not implemented yet
- SRM at UIO (ARC)
  - The request failed because ARC supports HTTP POST method to put files, but clients use HTTP PUT method.

COPY from gsiftp server to SRM

- SRM-Tester sends a request to the target SRM to “pull” a file from the source gsiftp server. Gsiftp server is based on GT 4.0.1.
- After a copy request is submitted to the target SRM, communication thereafter is between the SRM and the source gsiftp server.
- SRM-Tester checks the status of the copy request.
- When the first copy test failed, gsiftp server based on GT 2.4.3 at LBNL (STAR) was used as the source for additional testing, for possible certificate format handling issues.

Table 4 : Test results for COPY from GSIFTP server to SRM

Target SRM → Source GSIFTP ↓	ARC (UIO.NO)	LCG/EGEE (CERN.CH)	CMS (FNAL.GOV)	STAR (LBL.GOV)	OGS (SDSC,EDU)	EGEE (IC.UK)
Gsiftp server from GT4.0.1	fail	N.A.	pass	pass	N.A.	pass

- SRM at CERN (LCG/EGEE)
  - The request goes through the interface, and the status shows request queued and in progress.
  - We were informed that actual copy is not implemented yet.
- SRM at SDSC (OSG)
  - The status returned failure with no further explanation
  - We were informed that only “push” mode into the SRM at SDSC works because the underlying storage resides inside the private network.
- SRM at UIO (ARC)
  - The request failed and we have reported the error and they are working on it.

COPY between SRMs with 3rd party file transfer

- SRM-Tester sends a request to the source SRM to “get” a file from the source SRM server. Source SRM returns the TURL back to SRM-Tester when the file is ready.
- SRM-Tester sends a request to the target SRM to “put” a file into the target SRM server. Targer SRM returns the TURL back to SRM-Tester when space is ready.
- SRM-Tester makes a 3rd party file transfer from the TURL from the source SRM to the TURL from the target SRM.
- SRM-Tester makes a putDone to the target SRM after the transfer is complete, makes a “release” to the source SRM, and checks the status.

Table 5 : Test results for COPY between two SRMs with 3rd party file transfer

Target SRM → Source SRM ↓	ARC (UIO.NO)	LCG/EGEE (CERN.CH)	CMS (FNAL.GOV)	STAR (LBL.GOV)	OSG (SDSC.EDU)	EGEE (IC.UK)
ARC (UIO.NO)		pass	pass	pass	fail	pass
LCG/EGGE (CERN.CH)	fail		pass	pass	fail	pass
CMS (FNAL.GOV)	fail	pass		pass	fail	pass
STAR (LBL.GOV)	fail	pass	pass		fail	pass
OSG (SDSC.EDU)	fail	pass	pass	pass		pass
EGEE (IC.UK)	fail	pass	pass	pass	fail	

- SRM at SDSC (OSG)
  - The status returned failure.
  - This did not work because of their underneath storage policy.

- SRM at UIO (ARC)
    - The request didn't go through, because their put is using HTTP PUT and SRM-Tester and another srm client, SRMCP use HTTP POST.

Test results for SRM v2.2

Table 6 : Test results for SRM v2.2 (as of 9/4/2006)

Methods	CERN Castor	CERN DPM	FNAL dCache	RAL Castor	VU L-Store	LBNL SRM
Ping	Ok	Ok	Exception	Ok	Ok	Ok
PrepareToPut, Status, PutDone	N.A.	Ok	Ok	Ok	Ok	Ok
PrepareToGet, Status	Exception	Ok	Ok	Failed	Ok	Ok
Release	N.A.	Ok	Ok	N.A.	Ok	Ok
Mkdir, Rmdir	Ok	Ok	Ok	Ok	Ok	Ok
Mv	N.A.	Ok	Ok	Ok	Ok	Ok
Rm	N.A.	Ok	Ok	Ok	Ok	Ok
Ls	Exception	Exception	Ok	Exception	Ok	Ok
ReserveSpace	Failed	Ok	Ok	Failed	Ok	Ok
ReleaseSpace	N.A.	Ok	Ok	N.A.	Ok	Ok
BringOnline, Status	N.A.	Ok	N.A.	N.A.	Ok	Ok

Table 7 : Test results for SRM v2.2 with Space Reservation (as of 9/4/2006)

Methods	CERN Castor	CERN DPM	FNAL dCache	RAL Castor	VU L-Store	LBNL SRM
PrepareToPut, Status, PutDone	N.A.	Ok	Ok	N.A.	Ok	Ok
PrepareToGet, Status	N.A.	Ok	N.A.	N.A.	Ok	Ok
BringOnline, Status	N.A.	Ok	N.A.	N.A.	Ok	Ok

COPY between SRMs in PULL mode

- SRM-Tester sends a copy request to the target SRM to “pull” a file from the source SRM
- After a copy request is submitted to the target SRM, communication thereafter is between the two SRMs.
- SRM-Tester checks the status of the copy request.

Table 8 : Test results for COPY between SRMs in PULL mode (as of 9/4/2006)

Target SRM ➔ Source SRM ↓	CERN Castor	CERN DPM	FNAL dCache	RAL Castor	VU L-Store	LBNL SRM
CERN Castor		N.A.	Failed	N.A.	Failed	Failed
CERN DPM	N.A.		Ok	N.A.	Ok	Ok
FNAL dCache	N.A.	N.A.		N.A.	Ok	Ok
RAL Castor	N.A.	N.A.	Failed		Failed	Failed
VU L-Store	N.A.	N.A.	Failed	N.A.		Ok
LBNL SRM	N.A.	N.A.	Failed	N.A.	Ok	

- SRMs at CERN and RAL

    - The request goes through the interface, and the status shows request queued and in progress.
    - We were informed that actual copy is not implemented yet

COPY between SRMs in PUSH mode

- SRM-Tester sends a request to the source SRM to “push” a file from the source SRM to the target SRM
- After a copy request is submitted to the source SRM, communication thereafter is between the two SRMs.
- SRM-Tester checks the status of the copy request.

Table 9 : Test results for COPY between SRMs in PUSH mode (as of 9/4/2006)

Target SRM → Source SRM ↓	CERN Castor	CERN DPM	FNAL dCache	RAL Castor	VU L-Store	LBNL SRM
CERN Castor		N.A.	N.A.	N.A.	N.A.	N.A.
CERN DPM	N.A.		N.A.	N.A.	N.A.	N.A.
FNAL dCache	Failed	Ok		Ok	Failed	Failed
RAL Castor	N.A.	N.A.	N.A.		N.A.	N.A.
VU L-Store	Failed	Ok	Ok	Ok		Ok
LBNL SRM	Failed	Ok	Ok	Ok	Ok	

- SRMs at CERN and RAL
    - The request goes through the interface, and the status shows request queued and in progress.
    - We were informed that actual copy is not implemented yet

COPY from gsiftp server to SRM

- SRM-Tester sends a request to the target SRM to “pull” a file from the source gsiftp server. Gsiftp server is based on GT 4.0.1.
- After a copy request is submitted to the target SRM, communication thereafter is between the SRM and the source gsiftp server.
- SRM-Tester checks the status of the copy request.
- When the first copy test failed, gsiftp server based on GT 2.4.3 at LBNL (STAR) was used as the source for additional testing, for possible certificate format handling issues.

Table 10 : Test results for COPY from GSIFTP server to SRM (as of 9/4/2006)

Target SRM → Source GSIFTP ↓	CERN Castor	CERN DPM	FNAL dCache	RAL Castor	VU L-Store	LBNL SRM
Gsiftp server from GT4.0.1	Failed (work in progress)	*N.A.*	Ok	Failed (work in progress)	Ok	Ok

- SRMs at CERN and RAL
    - The request goes through the interface, and the status shows request queued and in progress.
    - We were informed that actual copy is not implemented yet.

Table 11 : Test results for COPY between two SRMs with 3rd party file transfer (as of 9/4/2006)

Target SRM ➔	CERN	CERN	FNAL	RAL	VU	LBNL
Source SRM 🡇	Castor	DPM	dCache	Castor	L-Store	SRM
CERN Castor		*N.A.*	*N.A.*	*N.A.*	*N.A.*	*N.A.*
CERN DPM	*N.A.*		Ok	*Ok*	*Ok*	*Ok*
FNAL dCache	*N.A.*	*Ok*		*Ok*	Ok	Ok
RAL Castor	*N.A.*	*N.A.*	*N.A.*		*N.A.*	*N.A.*
VU L-Store	*N.A.*	*Ok*	*Ok*	*Ok*		*OK*
LBNL SRM	*N.A.*	*Ok*	*Ok*	*Ok*	OK	

Status

1. Preparation
   a. Setting up a VO (gin.ggf.org) access to the host and the VO designated account was error-prone.
      i. Email communication was needed with system people
      ii. However, once it's set up, no problems
   b. The necessary information to perform testing can be summarized as the following five items.
      i. SRM end point: e.g. srm://t2data2.t2.sdsc.edu:8443//srm/manager
      ii. Version of SRM running on the above end point: e.g. 1.1, 2.2
      iii. working directory for GIN-VO: e.g. /data/gin
      iv. Available "copy" operational type: e.g. PULL, PUSH, ALL
      v. Optional supported transfer protocol list and its specific parameters: e.g. gsiftp, http, ftp, https, http:put, http:post.
2. SRM v1.1 operations
   a. Most implementations are compliant with the specification.
   b. Interoperability between SRMs was mostly tested successful.
   c. Incompatibility mostly comes from the transfer protocols and the underlying storage configurations, not from the interface incompatibility. Information service to advertise capabilities of individual SRMs would help.
3. SRM v2.2 operations
   a. Implementation is in-progress.
   b. Interoperability testing will be continued until deployment

Summary

Most of the SRM testing operations were successful. Most problems come from the mismatch of the file transfer protocols and its parameters. Information service or discovery mechanism to advertise capabilities of the SRM may help reduce such failures.

References

1. Grid Interoperability Now (GIN) Working Group, https://forge.gridforum.org/projects/mgi
2. GIN VOMS, https://kuiken.nikhef.nl:8443/voms/gin.ggf.org/
3. SRM Collaboration, http://sdm.lbl.gov/srm-wg
4. SRM Specification v1.1, http://sdm.lbl.gov/srm-wg/doc/srm.v1.0.pdf
5. SRM Specification v2.2, http://sdm.lbl.gov/srm-wg/doc/SRM.v2.2.html
6. SRM-Tester, http://sdm.lbl.gov/srm-dist
7. SRMCP, http://srm.fnal.gov
8. NGCP, ftp://ftp.nordugrid.org/nordugrid/tags/0.5.46/

9. GGF-18, GIN SRM island testing results, http://sdm.lbl.gov/srm-tester/ggf18.html
10. GGF=17, GIN SRM island testing results, http://sdm.lbl.gov/srm-tester/ggf17.html

Acknowledgement

- APAC
  - Univ. of Melbourne : Glenn Moloney <glenn@physics.unimelb.edu.au>
- ARC
  - Lund Univ. : Oxana Smirnova <oxana.smirnova@hep.lu.se>
    Aleksandr Konstantinov <const@takas.lt>
- EGEE
  - CERN : Erwin Laure <Erwin.Laure@cern.ch>
  - Imperial College : Mona Aggarwal <m.aggarwal@imperial.ac.uk>
    Olivier van der Aa <o.van-der-aa@imperial.ac.uk>
    David Colling <d.colling@imperial.ac.uk>
- FNAL
  - Timur Perelmutov <timur@fnal.gov>
- Grid.IT
  - INFN : Riccardo Zappi <riccardo.zappi@cnaf.infn.it>
    Luca Magnoni <luca.magnoni@cnaf.infn.it>
- OSG
  - UCSD : Frank Wuerthwein <fkw@ucsd.edu>
    Abhishek Singh Rana <rana@ucsd.edu>
- STAR
  - BNL: Jerome Lauret <jeromel@bnl.gov>
  - LBNL : Vijaya Natarajan <vnatarajan@lbl.gov>
    Junmin Gu <jgu@lbl.gov>
    Arie Shoshani <shoshani@lbl.gov>
    Alex Sim <asim@lbl.gov>
- VU
  - Vanderbilt Univ. : Surya Pathak <surya.pathak@accre.vanderbilt.edu>
    Paul Sheldon <paul.sheldon@vanderbilt.edu>
- WLCG/EGEE
  - CERN : James Casey <James.Casey@cern.ch>
    Jean-Philippe Baud <Jean-Philippe.Baud@cern.ch>
  - RAL: Shaun De Witt <S.De.Witt@rl.ac.uk>
    Jiri Mencak <j.mencak@rl.ac.uk>
- GIN VOMS VO
  - NIKHEF : Oscar Koeroo <okoeroo@nikhef.nl>
  - INFN : Vincenzo Ciaschini <vincenzo.ciaschini@cnaf.infn.it>