

Public Storage on the Open Science Grid

Ted Hesselroth
Fermilab

Purpose

- Provide managed storage for non-owner VOs on Storage Elements
- Public: “Maintained for or used by the people or community”
 - *Community = OSG VOs*
- Plan to develop and/or integrate software
 - Requirements have been reviewed.
 - Backend database and its operations have been defined.

Previous Users

- D0: Tevatron Experiment Monte Carlo production
- Fly's Eye: Highly Energetic Dispersed Radio Transients
- SCEC: Probabilistic Seismic Hazard Assessments
- Human Microbiome Project: Correlate microbiome/health
- LIGO: Einstein @ Home
- LIGO: Binary Inspiral
- LSST: Simulated images.

Current Procedure

- Contact OSG
- Discuss of Storage requirements
- Arrange space with sites
- Do test transfers
- Troubleshoot authorization or BDII info
- Upload files, notify by email
- Track uploaded files
- Run jobs

Current Procedure

- Contact OSG
- Discuss of Storage requirements
- Arrange space with sites
- Do test transfers
- Troubleshoot authorization or BDII info
- Upload files, notify by email
- Track uploaded files
- Run jobs

New Services

- Space Management needed for
 - Coordinate use of space in public storage areas by VOs
 - Track accessibility and use of public storage
- Transfer Management needed for
 - Enforcement of space management
 - Track files, cleanup
 - Storage resource selection

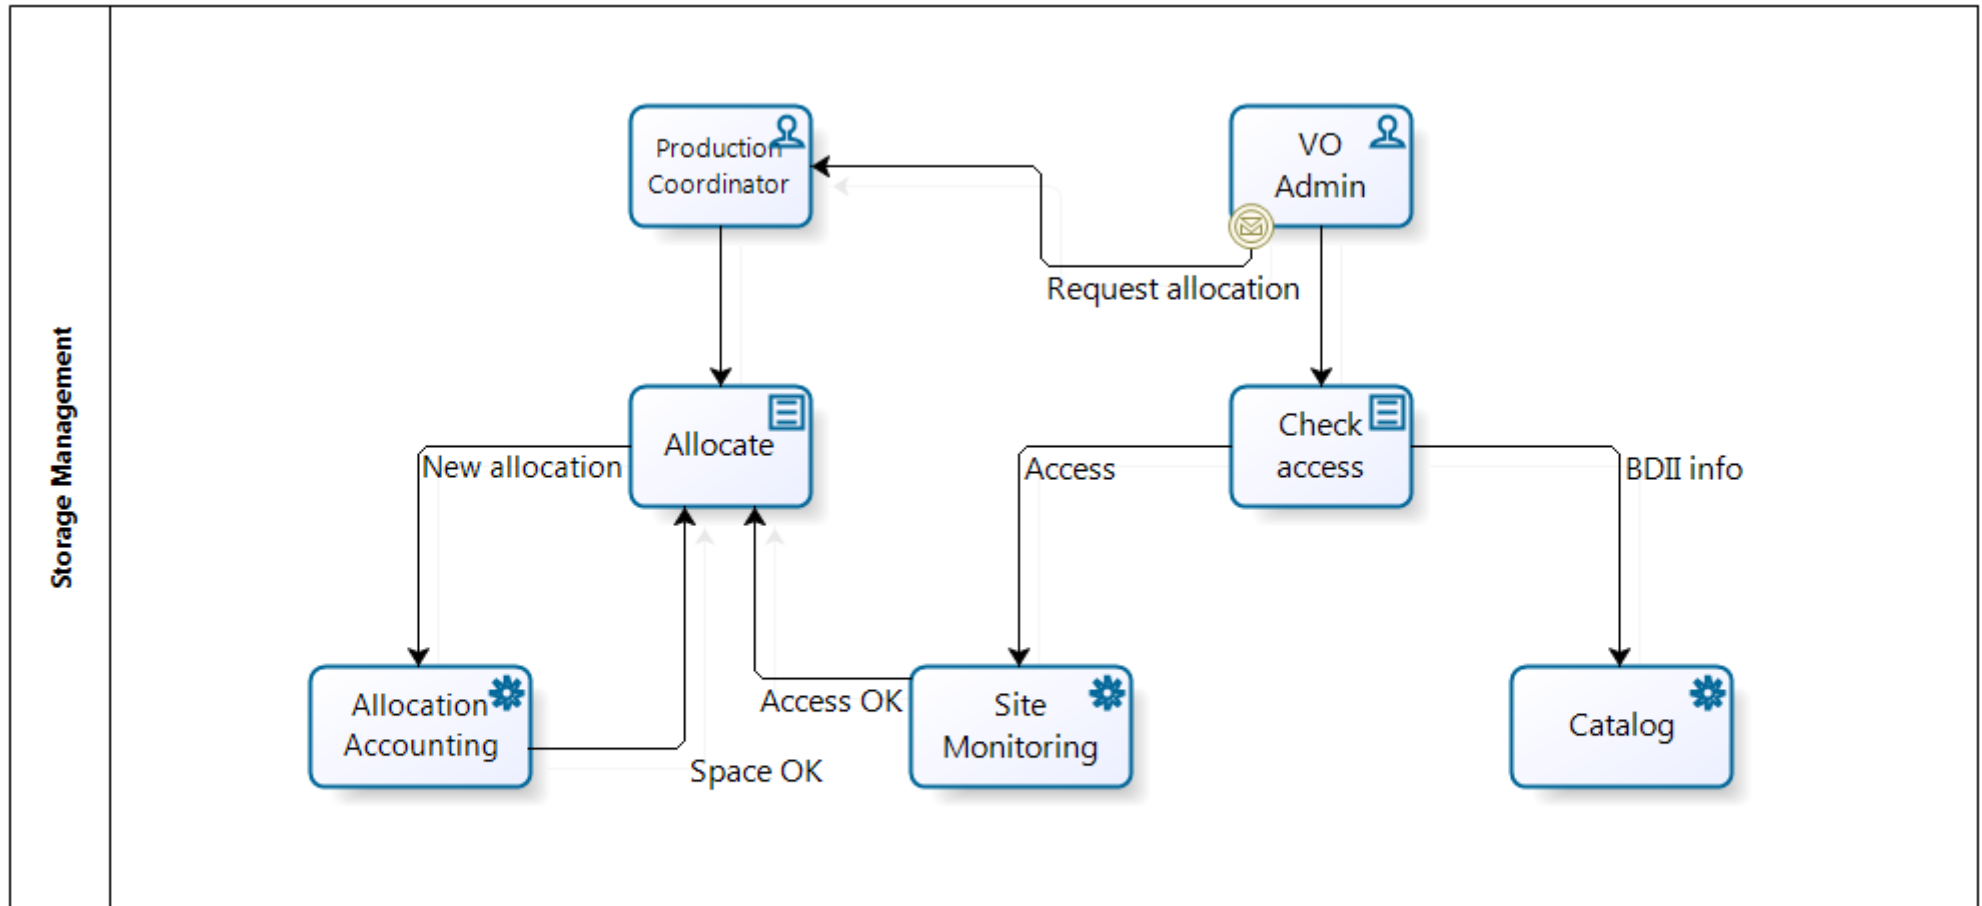
Constraints on the Design

- No alteration of Storage Elements
 - Access continues to be through current clients.
- No centralized OSG service
 - Software to be operated by VO
- Accommodate usage outside the service
 - Use of traditional means will not have an adverse effect

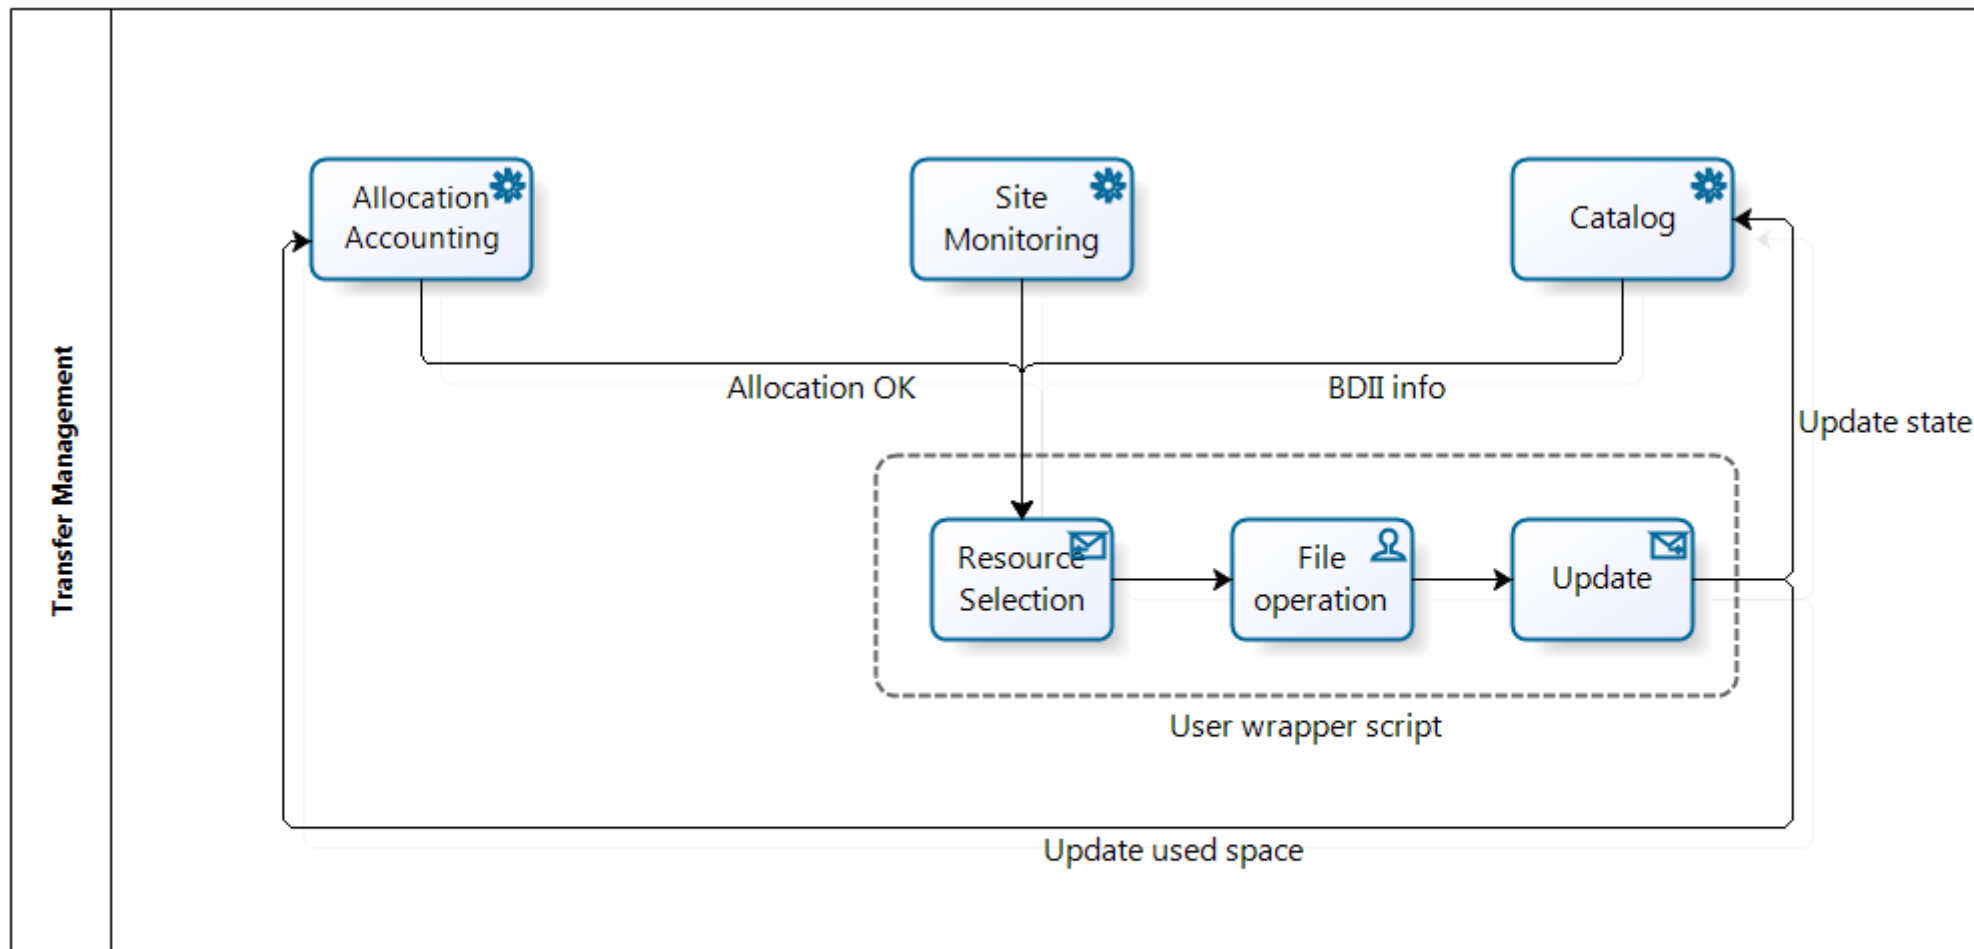
Public Storage Space Management

- Site administrator deploys a storage area and authorizes VOs
 - Puts storage area discovery info on OSG-central BDII.
- VO Administrator discovers and checks access to storage areas.
 - Requests an allocation
- OSG Production Coordinator allocates space on a storage area.
 - Allocation = Storage Area, Size, Max number of files, Expiration.
- User accesses a storage resource through its allocation.
 - Allocation parameters are enforced and updated.
- VO Administrator cleans up expired allocations.

Storage Management Process



Transfer Management Process



Components of the Public Storage Service

- Allocation Accounting Service
 - Used to grant an allocation to a VO.
 - Tracks space used in allocations.
- Site Monitoring Service
 - Saves records of access check tools.
- Public Storage Catalog
 - Maps logical file name to location on storage resource.
 - Tracks replicas for access and cleanup.
 - Saves dataset definitions.

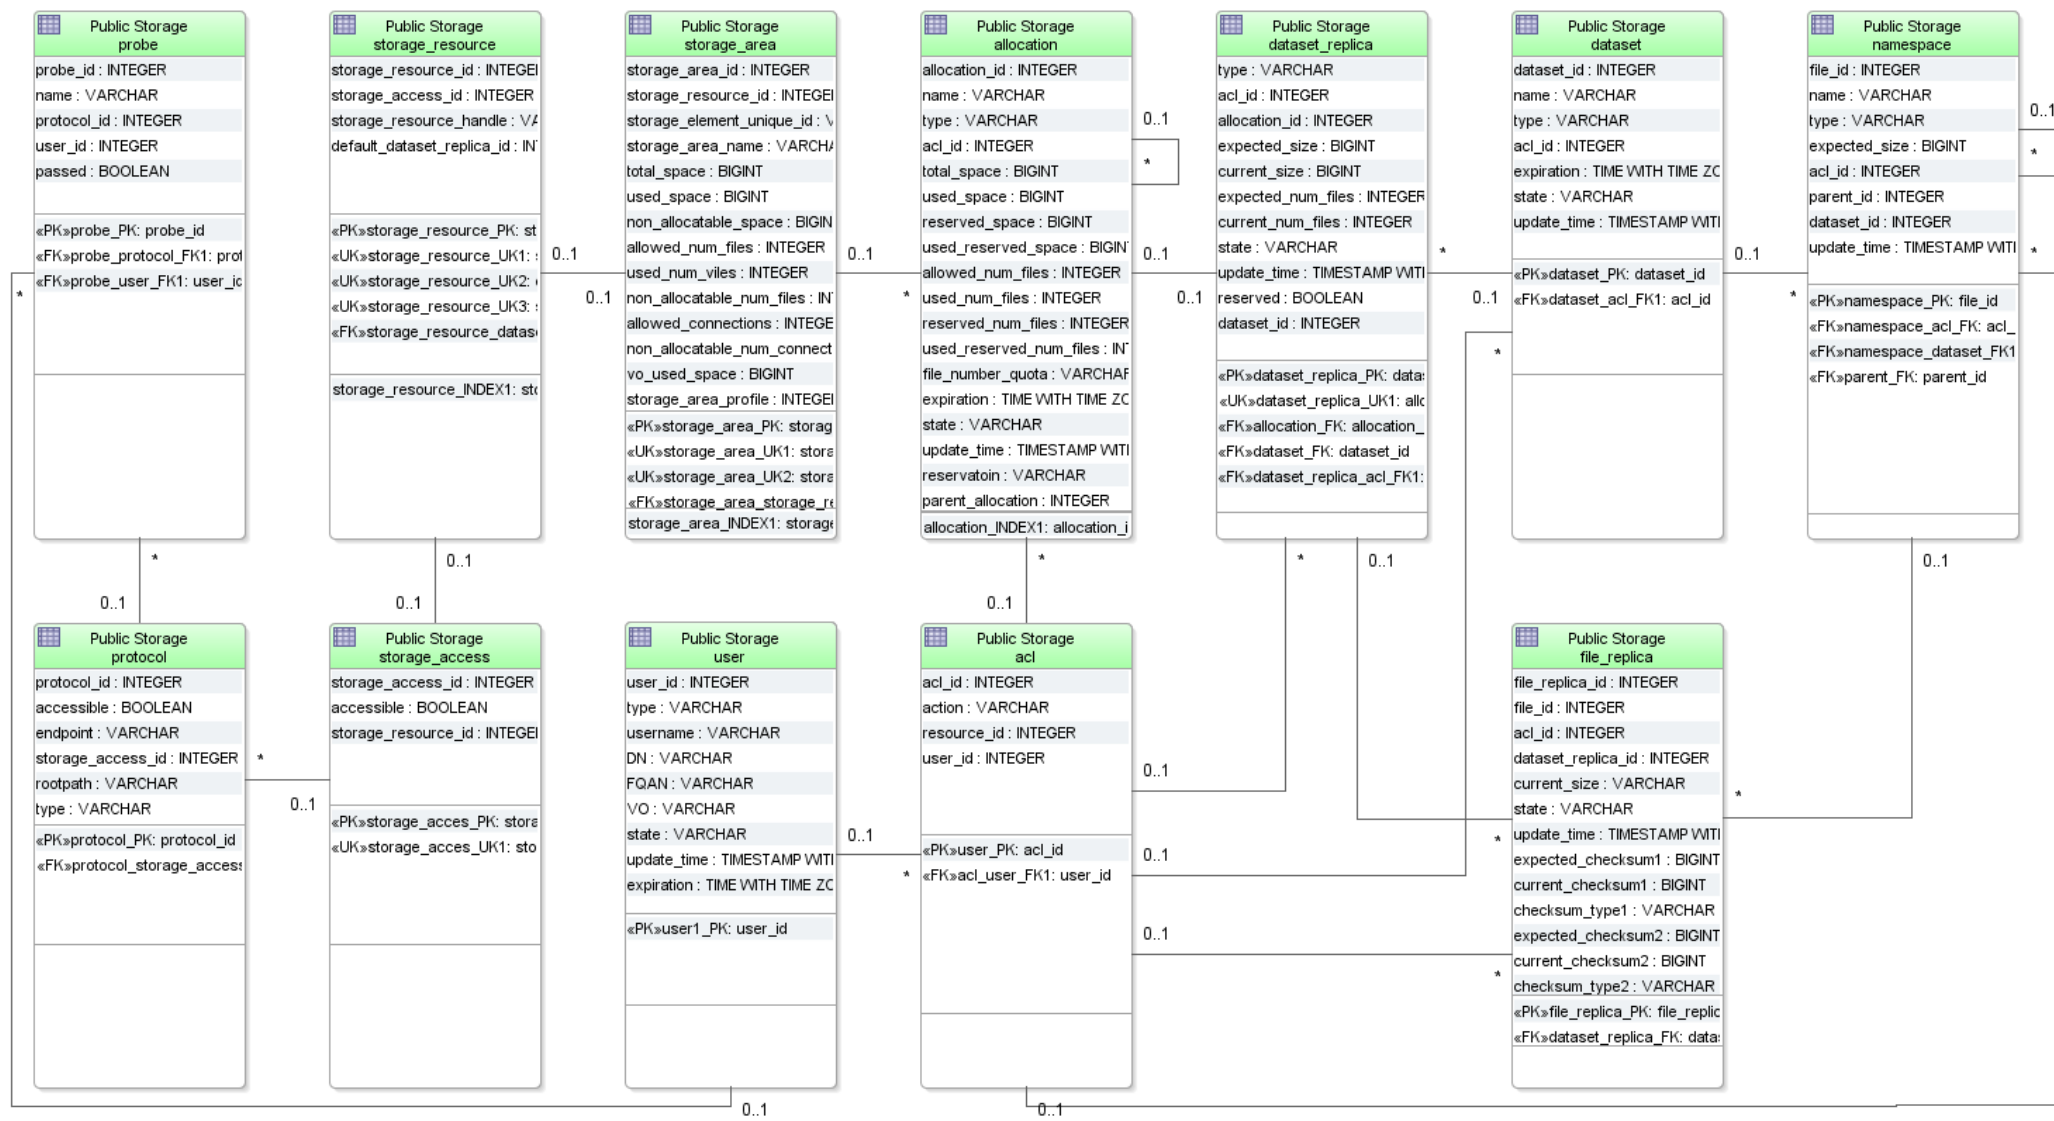
Support for Dataset Operations

- User defines a dataset.
 - By ancestor directory or metadata entry
- VO Administrator optionally makes suballocations for the dataset.
 - Same parameter types as for allocations.
- User uploads the dataset files.
 - All files to the same destination storage resource.
- User replicates the dataset files to another resource.
- User runs jobs on the dataset.
- User deletes expired datasets.

Cleanup Policy

- Ensure availability of unused space over the long term.
- Expiration times for allocations and datasets.
- In order of preference
 - Users clean up datasets (delete files in them)
 - VO Administrators clean up allocations
 - Site administrators clean up allocations
 - At the request of the OSG Production Coordinator
 - Using a provided tool and input list

The Actual Schema



The Pigeon Access Checker

- Discover Servers
 - The test “generator” calls Discovery Tools on a daily basis to get a list of servers from BDII that advertise support for the VO.
- Create Tests
 - Using this list, create a custom suite of tests to verify the SRM/gridftp functionality of each site.
- Execute Tests
 - Execute the tests to verify the site's storage functionality.
- Show Results
 - Publish the results on a web page and in a XML.

Pigeon Screenshot

Mozilla Firefox

File Edit View History Bookmarks Tools Help

[←](#)
[→](#)
[↻](#)
[✕](#)
[🏠](#)
fnal.gov
<https://gw014k0.fnal.gov:8443/gip/>
☆
Google

<https://gw014k0....l.gov:8443/gip/>
+

GLOW

[Create Link](#)

1ping GLOW	2010-08-27 09:18:01.767522	Success	Command finished.	Archive Create Link
2xncp GLOW	2010-08-27 09:25:02.252111	Failure	Authorization Error	Archive GOC TICKET 8799
3xmkfile GLOW		Not Run	Dependent test 2xncp_GLOW has failed.	Archive Create Link
4xmkmdir GLOW		Not Run	Dependent test 2xncp_GLOW has failed.	Archive Create Link
5xmkmdir GLOW		Not Run	Dependent test 4xmkmdir_GLOW has not run yet due to dependencies.	Archive Create Link
6xmkmdir GLOW		Not Run	Dependent test 4xmkmdir_GLOW has not run yet due to dependencies.	Archive Create Link
7xmkrm GLOW		Not Run	Dependent test 2xncp_GLOW has failed.	Archive Create Link
globus url copy GLOW	2010-08-27 09:15:01.485599	Success	Command finished.	Archive Create Link
globus url run GLOW	2010-08-27 09:20:02.205427	Success	Command finished.	Archive Create Link

GRASE-CSE-MAGIC

globus url copy GRASE-CSE-MAGIC	2010-08-27 09:34:01.856264	Success	Command finished.	Archive Create Link
globus url run GRASE-CSE-MAGIC	2010-08-27 09:39:04.892010	Success	Command finished.	Archive Create Link

GridUNESB CENTRAL

The OSG Discovery Tool

- Reads BDII information
- Can search using XPath and XQuery
 - XML representation more future-proof than LDAP
 - Composition of complex queries is easier than with Idapsearch
- Prewritten Compute Element Searches
 - OS Version, Runtime Version, SE mount point
- Prewritten Storage Element Searches
 - SRM and gridftp URLs. Client command templates.
 - Storage Area available space and authorization info.

get_runtime_versions

- Finds software version with matching string

```
[discovery-1.0.6]$ get_runtime_versions --vo engage --match "OSG 1.2."
```

Site Name	Compute Element ID	Software Version
BNL-ATLAS	gridgk01.racf.bnl.gov:2119/jobmanager-condor-default	OSG 1.2.9
BNL-ATLAS	gridgk02.racf.bnl.gov:2119/jobmanager-condor-default	OSG 1.2.9
CIT_CMS_T2	cit-gatekeeper.ultralight.org:2119/jobmanager-condor-default	OSG 1.2.8
CIT_CMS_T2	cit-gatekeeper.ultralight.org:2119/jobmanager-condor-osg_user	OSG 1.2.8
CIT_CMS_T2	cit-gatekeeper2.ultralight.org:2119/jobmanager-condor-default	OSG 1.2.8
CIT_CMS_T2	cit-gatekeeper2.ultralight.org:2119/jobmanager-condor-osg_user	OSG 1.2.8
GridUNESP_CENTRAL	ce.grid.unesp.br:2119/jobmanager-condor-default	OSG 1.2.11
MIT_CMS	ce01.cmsaf.mit.edu:2119/jobmanager-condor-default	OSG 1.2.3
USCMS-FNAL-WC1	cmsosgce3.fnal.gov:2119/jobmanager-condor-default	OSG 1.2.8
FNAL_FERMIGRID	d0cabosg2.fnal.gov:2119/jobmanager-pbs-osg	OSG 1.2.11
FNAL_FERMIGRID	fermigridosg1.fnal.gov:2119/jobmanager-condor-default	OSG 1.2.11
Firefly	ff-grid.unl.edu:2119/jobmanager-pbs-default	OSG 1.2.4
Firefly	ff-grid2.unl.edu:2119/jobmanager-pbs-default	OSG 1.2.4
Firefly	ff-grid3.unl.edu:2119/jobmanager-pbs-default	OSG 1.2.4

...

get_mount_path

- Finds the mount point of the Storage Element on worker nodes

```
[discovery-1.0.6]$ get_mount_path --vo ligo
```

COMPUTE ELEMENT ID	MOUNT POINT
osg-gw-2.t2.ucsd.edu:2119/jobmanager-condor-default	/hadoop/ligo
osg-gw-2.t2.ucsd.edu:2119/jobmanager-condor-group_ligo	/hadoop/ligo
cit-gatekeeper.ultralight.org:2119/jobmanager-condor-cms_monitor	/mnt/hadoop/osg/LIGO
cit-gatekeeper.ultralight.org:2119/jobmanager-condor-default	/mnt/hadoop/osg/LIGO
cit-gatekeeper2.ultralight.org:2119/jobmanager-condor-default	/mnt/hadoop/osg/LIGO
cit-gatekeeper2.ultralight.org:2119/jobmanager-condor-osg_user	/mnt/hadoop/osg/LIGO
ce.grid.unesp.br:2119/jobmanager-condor-default	/store/ligo

```
...
```

get_storage_area

- Finds total and free space, authorized VOs
- Can be used to find public storage areas

```
[discovery-1.0.6]$ get_storage_area
```

STORAGE ELEMENT ID	STORAGE AREA NAME	TOTAL	FREE	VOs
atlas.bu.edu	Default Storage Area	122984	26264	atlas,cdf,mis,ops,osg
atlas07.cs.wisc.edu	Default Storage Area	0	0	atlas
bsrm-1.t2.ucsd.edu	Default Storage Area	1010185	100463	accelerator,astro,atlas,...
charm.ucr.edu	Default Storage Area	6730	1388	atlas,cdf,cms,engage,...
cit-se.ultralight.org	Default Storage Area	526869	109617	DOSAR,LIGO,alice,atlas,...
cmssrm.fnal.gov	flushPools:replica:online	6984011	1870180	cms,mis
cmssrm.fnal.gov	LFSOnlyPools:replica:online	529440	377902	cms,mis
cmssrm.fnal.gov	ResilientPools:replica:online	428079	195002	cms,mis
cmssrm.fnal.gov	stagePools:custodial:nearline	161695	37560	cms,mis
cmssrm.hep.wisc.edu	black-pools:replica:online	20186	11618	LIGO,atlas,cdf,cms,...

...

Further Info

- OSG Twiki
 - OSG Public Storage
 - Opportunistic Storage
 - OSG Storage Discovery Tool
- Pigeon: <http://home.fnal.gov/~dstrain/pigeon>
- OSG docdb
 - OSG Public Storage Requirements
 - OSG Public Storage Design