



GDB

# Storage Accounting

John Gordon  
GDB, June 2008



GDB

# Status

- A Storage Accounting schema has been defined
- A single storage accounting sensor at Edinburgh queries the BDII once per day to gather information on current storage usage from all UK sites
- This sensor populates an R-GMA archiver with the gathered information
  - StorageRecords table
- The information is aggregated per VO/Site/Month as in APEL
  - Value stored is average over the number of measurements
  - $\text{Sum}(\text{measurements})/\text{N}(\text{measurements})$
- APEL-like visualisation interface has been developed
  - Grids, countries, regions
  - Only tables for country and region implemented so far
  - <http://goc02.grid-support.ac.uk/accountingDisplay/>



GDB

# Recent Progress

- Page now displays about sites or other points in the EGEE Tree
- Lists types of storage, allocated and used for LHC VOs.

Resource Tree

- EGEE
- OSG
- Tier1
  - CA-TRIUMF
  - CH-CERN
  - DE-KIT
  - ES-PIC**
  - FR-CCIN2P3
  - IT-INFN-CNAF
  - NDGF
  - NL-T1
  - TW-ASGC
  - UK-T1-RAL
  - US-FNAL-CMS
  - US-T1-BNL
- Tier2
- Unassigned

Storage Accounting Display [v2.1]

Select Interval:

**VO Groups**

LHC
  non-LHC
  ALL
  Custom

**VOs:**

atlas
  biomed
  cms
  des
  dteam  
 gear
  Ifae
  lhcb
  magic
  ops  
 pps
  swetest

**SEArchitecture:**

All  
 castor(permanent)
  disk(permanent)
  disk(unknown)  
 disk(volatile)
  multi-disk(permanent)
  multidisk(permanent)  
 multidisk(unknown)
  multidisk(volatile)
  srm(permanent)  
 srm\_v1(permanent)
  tape(permanent)
  tape(unknown)  
 unknownArch(permanent)
  unknownArch(unknown)
  unknownArch(volatile)

Refresh

Disk Usage Reported in the last Hour (TB)

- [GridPP Wiki Storage Accounting](#)
- The term Allocated is defined as `GlueSAStateSpaceUsed + GlueSAStateSpaceAvailable`
- The table reports what is actually published into the information system from sites.
- Storage areas that are shared by VOs are not treated in any special way.




### Disk Usage Reported in the last Hour (TB)

- [GridPP Wiki Storage Accounting](#)
- The term Allocated is defined as GlueSAStateSpaceUsed + GlueSAStateSpaceAvailable
- The table reports what is actually published into the information system from sites.
- Storage areas that are shared by VOs are not treated in any special way.
- [Savannah Bug 25139](#) There is a problem in adding together the allocated storage for each VO at a particular site as it often leads to double (triple...) counts of resources.
- **Navigation:** Click on a heading in the results table, to see the information underneath it.

Disk Used and <b>Allocated</b> per VO (Units are TB) for ES-PIC in the last Hour												
Entity	Reports	Architecture	VO						Total (TB)		Percentage	
			atlas		cms		lhcb		Used	Alloc	Used	Alloc
			Used	Alloc	Used	Alloc	Used	Alloc				
pic		multidisk	86.332	111.787	86.245	130.189	35.465	35.533	208.041	277.51	100%	100%
		tape	1000	1000	1000	1000	1000	1000	3000	3000	100%	100%
<b>Table Summaries</b>												
<b>VO Totals (TB)</b>		multidisk	86.332	111.787	86.245	130.189	35.465	35.533	208.041	277.51	6.48%	8.47%
		tape	1000	1000	1000	1000	1000	1000	3000	3000	93.52%	91.53%
		TOTAL	1086.332	1111.787	1086.245	1130.189	1035.465	1035.533	3208.041	3277.51		
<b>Percentage</b>		multidisk	41.5%	40.28%	41.46%	46.91%	17.05%	12.8%				
		tape	33.33%	33.33%	33.33%	33.33%	33.33%	33.33%				
		TOTAL	33.86%	33.92%	33.86%	34.48%	32.28%	31.6%				

Available resources may be shared between multiple VOs depending on the disk pool configuration of the SE at each site.



### Disk Used and Allocated per VO (Units are TB) for Tier1 in the last Hour

Entity	Reports	Architecture	VO								Total (TB)		Percentage	
			alice		atlas		cms		lhcb		Used	Alloc	Used	Alloc
			Used	Alloc	Used	Alloc	Used	Alloc	Used	Alloc				
CA-TRIUMF		disk	0	0	0	0	0	0	0	0	0	0	0%	0%
CH-CERN		multidisk	0	0.436	1.191	2.311	0	0.436	0	0.436	1.191	3.62	0.08%	0.2%
		unknownArch	1.431	4.292	2.384	7.153	0.954	2.861	2.384	7.153	7.153	21.458	8.12%	7.09%
DE-KIT		disk	0.771	31.265	142.885	148.062	7.764	16.501	45.039	55.749	196.459	251.578	56.43%	52.84%
ES-PIC		multidisk	0.00	0.00	86.332	111.787	86.245	130.189	35.465	35.533	208.041	277.51	14.5%	15.3%
		tape	0	0	1000	1000	1000	1000	1000	1000	3000	3000	95.02%	85.87%
FR-CCIN2P3		disk	3.836	5.661	39.91	50.144	36.336	78.589	71.593	90.164	151.675	224.559	43.57%	47.16%
IT-INFN-CNAF	Tier1	multidisk	500	500	501.267	516.322	50.132	191.163	1.44	5.222	1052.839	1212.707	73.38%	66.88%
		tape	0.00	0.00	50	150	0.00	0.00	0.00	0.00	50	150	1.58%	4.29%
		unknownArch	0.048	0.143	0.048	0.143	0.048	0.143	0.095	0.286	0.238	0.715	0.27%	0.24%
NDGF		multidisk	0.012	0.022	0	0	0	0	0	0	0.012	0.022	0%	0%
NL-T1		multidisk	0	0	0	0	0.001	0.001	0	1.85	0.001	1.851	0%	0.1%
		tape	0	0.419	1.57	1.573	0	0	1.048	1.049	2.617	3.041	0.08%	0.09%
		unknownArch	2.111	3.204	52.419	52.967	1.595	2.071	1.595	2.071	57.721	60.313	65.51%	19.94%
TW-ASGC		multidisk	0	0	0.013	60.059	0.013	60.059	0.00	0.00	0.026	120.118	0%	6.62%
		unknownArch	0	0	0	0	0	0	0	0	0	0	0%	0%
UK-T1-RAL		multidisk	0	0	90.746	100.039	9.474	11.509	72.373	85.823	172.594	197.372	12.03%	10.89%
		tape	1.292	13.604	9.827	79.759	50.767	182.808	42.889	64.58	104.774	340.752	3.32%	9.75%
US-FNAL-CMS		disk	0	0	0	0	0	0	0	0	0	0	0%	0%
US-T1-BNL		disk	0.00	0.00	0	0	0.00	0.00	0.00	0.00	0	0	0%	0%
		multidisk	0	0	0	0	0	0	0	0	0	0	0%	0%
		unknownArch	0.00	0.00	0.00	0.00	23	220	0.00	0.00	23	220	26.1%	72.73%

**Table Summaries**

	disk	4.608	36.927	182.795	198.206	44.1	95.09	116.632	145.914	348.135	476.137	6.92%	7.82%
	multidisk	500.012500.458	679.549	790.519	145.865	393.357	109.278	128.864	1434.704	1813.198	28.53%	29.79%	
<b>VO Totals (TB)</b>	tape	1.292	14.024	1061.396	1231.332	1050.767	1182.808	1043.937	1065.629	3157.392	3493.793	62.79%	57.41%

CONDOR

sar

ics

# OSG

a\_Storage

S

\_Tier3

OSG-Prod-SE

srm\_v1

S

-WC1-CE2

-WC1-SE

e

-WC1-CE

ier2

:

## Disk Used and Allocated per VO (Units are TB) for OSG in the last Hour

Entity	Reports	Architecture	VO				Total (TB)		Percentage		
			atlas		cms		Used	Alloc	Used	Alloc	
			Used	Alloc	Used	Alloc					
			multidisk	0	0	0	0	0	0%	0%	
			unknownArch	0.00	0.00	23	220	23	220	100%	100%
			unknownArch	0	0	0	0	0	0	0%	0%
			disk	1.305	3.909	1.305	3.909	2.611	7.817	10%	8.59%
			unknownArch	0	0	0	0	0	0	0%	0%
			disk	0	0	0	0	0	0	0%	0%
			unknownArch	0	0	0	0	0	0	0%	0%
			disk	0	0	0	0	0	0	0%	0%
			unknownArch	0	0	0	0	0	0	0%	0%
			disk	0	0	0	0	0	0	0%	0%
			unknownArch	0	0	0	0	0	0	0%	0%
			disk	0.00	0.00	0	0	0	0	0%	0%
			unknownArch	0	0	0	0	0	0	0%	0%
			disk	1.075	1.74	1.075	1.74	2.149	3.48	8.23%	3.82%
			unknownArch	0	0	0	0	0	0	0%	0%
			disk	7.031	20.484	7.031	20.484	14.063	40.969	53.86%	45.03%
			unknownArch	0	0	0	0	0	0	0%	0%
			disk	0.022	6.56	0.022	6.56	0.044	13.12	0.17%	14.42%
			unknownArch	0	0	0	0	0	0	0%	0%
			disk	0.739	2.106	0.739	2.106	1.479	4.211	5.66%	4.63%
			unknownArch	0	0	0	0	0	0	0%	0%
			disk	0.37	1.053	0.37	1.053	0.739	2.106	2.83%	2.31%
			unknownArch	0	0	0	0	0	0	0%	0%
			disk	0.083	0.225	0.083	0.225	0.165	0.45	0.63%	0.49%
			unknownArch	0	0	0	0	0	0	0%	0%
			unknownArch	0	0	0	0	0	0	0%	0%
			unknownArch	0	0	0	0	0	0	0%	0%
			disk	0	0	0	0	0	0	0%	0%
			disk	0	0	0	0	0	0	0%	0%
			disk	2.429	9.371	2.429	9.371	4.859	18.741	18.61%	20.6%

Table Summary








But the data is not good.

GDB



# WLCG Accounting Summary

April 2008  
Centre PIC

# GDB

Please report accounting data in the shaded cells and return the report to [lcg\\_office@cern.ch](mailto:lcg_office@cern.ch)

			Jan-08	Feb-08	Mar-08	Apr-08	May-08	Jun-08	Jul-08
Disk Space - TBytes	ALICE	allocated							
		used							
	ATLAS	allocated	113	221	224	224			
		used	105	136	140	143			
	CMS	allocated	126	167	176	176			
		used	66	66	62	62			
	LHCb	allocated	29	156	149	149			
		used	27	31	38	41			
	TOTAL	allocated	268	694	549	549	0	0	
		used	196	263	260	245	0	0	
installed capacity <sup>***</sup>			268	694	597	577			
disk usage as % installed <sup>***</sup>			100%	54%	62%	59%	n/a	n/a	n/a
MoU pledge <sup>a</sup>	allocated		220	220	220	970	970	970	970
	used		154	154	154	679	679	679	679

			Jan-08	Feb-08	Mar-08	Apr-08	May-08	Jun-08	Jul-08
Tape Space Used - TBytes	ALICE								
	ATLAS		40	46	56	75			
	CMS		142	191	338	280			
	LHCb		24	23	26	25			
	TOTAL		208	260	419	380	0	0	
	installed capacity <sup>***</sup>		240	293	427	520			
tape usage as % installed <sup>***</sup>			88%	89%	98%	73%	n/a	n/a	n/a
MoU pledge <sup>a</sup>	allocated		240	240	240	950	950	950	950
	used		240	240	240	650	650	650	650

## Notes - disk and tape tables

Space used/allocated at end of the calendar month




### Disk Usage Reported in the last Hour (TB)

- [GridPP Wiki Storage Accounting](#)
- The term Allocated is defined as GlueSAStateSpaceUsed + GlueSAStateSpaceAvailable
- The table reports what is actually published into the information system from sites.
- Storage areas that are shared by VOs are not treated in any special way.
- [Savannah Bug 25139](#) There is a problem in adding together the allocated storage for each VO at a particular site as it often leads to double (triple...) counts of resources.
- **Navigation:** Click on a heading in the results table, to see the information underneath it.

Disk Used and <b>Allocated</b> per VO (Units are TB) for ES-PIC in the last Hour													
Entity	Reports	Architecture	VO						Total (TB)		Percentage		
			atlas		cms		lhcb		Used	Alloc	Used	Alloc	
			Used	Alloc	Used	Alloc	Used	Alloc					
pic		multidisk	86.332	111.787	86.245	130.189	35.465	35.533	208.041	277.51	100%	100%	
		tape	1000	1000	1000	1000	1000	1000	3000	3000	100%	100%	
<b>Table Summaries</b>													
<b>VO Totals (TB)</b>			multidisk	86.332	111.787	86.245	130.189	35.465	35.533	208.041	277.51	6.48%	8.47%
			tape	1000	1000	1000	1000	1000	1000	3000	3000	93.52%	91.53%
			TOTAL	1086.332	1111.787	1086.245	1130.189	1035.465	1035.533	3208.041	3277.51		
<b>Percentage</b>			multidisk	41.5%	40.28%	41.46%	46.91%	17.05%	12.8%				
			tape	33.33%	33.33%	33.33%	33.33%	33.33%	33.33%				
			TOTAL	33.86%	33.92%	33.86%	34.48%	32.28%	31.6%				

Available resources may be shared between multiple VOs depending on the disk pool configuration of the SE at each site.





GDB

## But the data is not good.

- Published values, don't agree with monthly accounts.
- Needs work to reconcile with Tier1s
- Information Providers need work
- Storage Tokens/Spaces need recording
  
- This work has started
- There will be a proposal soon to MB on reporting.