Proposal for obtaining installed capacity

Flavia Donno

CERN/IT

WLCG Management Board,
CERN 8 July 2008

The requirements

- **The goal**
  - Provide the management with information about ***installed capacity*** (per VO).
  - Provide the management with information about ***resource usage*** (per VO).
    - This info can also be used by VO operations people in order to "monitor" their usage of the resources
- **Main focus**
  - Storage resources
  - Computing resources
- **Dynamic view**
  - Information should be retrieved ***as dynamically as possible*** in order to be reliable
  - For the moment the ***source is the information system***
  - The information will be made ***available through the WLCG accounting system***

The documentation

- The technical specifications are available in the CCRC twiki in the SSWG section:

  https://twiki.cern.ch/twiki/bin/view/LCG/WLCGCommonComputingReadinessChallenges#Storage_Solution_Working_Group_S

- Storage only:
  - The documentation available reports on the conclusions reached during focused meetings with developers and information providers
    - Specific solutions found to cover dCache internal specialized buffers and avoid double counting

- The document has been agreed by *storage developers, storage information providers, data management developers,* …

The Storage Resources

- Current accounting harvests data from BDII
  - … but many errors in data.
  - combination of Glue1.2 and 1.3
- Need to improve information providers and configuration.
- Existing Reports
  - Current status (last hour)
  - Monthly Report (like manual Tier1 reports).

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	230.961	500	0	0	0	0	230.961	500	38.32%	25.63%
CH-CERN		multidisk	0	0.352	1.191	2.228	0	0.352	0	0.352	1.191	3.284	0.01%	0%
		unknownArch	1.431	4.292	501.907	1505.722	1000	3000	2500	7500	4003.338	12010.014	18.9%	18.26%
DE-KIT		disk	160.784	461.147	13.793	28.651	0.784	2.543	160.404	464.747	335.764	955.087	55.71%	48.76%
ES-PIC		multidisk	0	0	30.074	49.202	51.177	84.283	0	0	81.251	133.485	0.56%	0.19%
		tape	0.00	0.00	1000	1000	1000	1000	1000	1000	3000	3000	75.39%	13.76%
FR-CCIN2P3		disk	11.885	51.76	0.328	10.604	23.934	434.616	0.036	6.655	35.984	503.635	5.97%	25.71%
IT-INFN-CNAF		multidisk	0.076	0.102	12833.984	26500.109	1106.403	43099.876	0.076	1000.102	13940.519	70600.189	95.37%	98.79%
		tape	500	1500	50	150	0.00	0.00	0.00	0.00	550	1650	13.82%	7.57%
		unknownArch	50	150	50	150	50	150	100	300	250	750	1.18%	1.14%
NDGF		multidisk	0.002	0.072	0.002	0.072	0	0	0	0	0.004	0.144	0%	0%
NL-T1		multidisk	0	0	0	0	0.001	0.001	0	0.249	0.001	0.249	0%	0%
		tape	0	2.04	132.164	8261.704	0.00	0.00	44.752	8257.516	176.917	16521.26	4.45%	75.77%
		unknownArch	2.158	3.204	16900.257	52778.829	1.804	2.071	1.804	2.071	16905.824	52785.975	79.81%	80.26%
TW-ASGC		multidisk	0	0	152.25	207.109	152.25	207.109	0.00	0.00	304.499	414.218	2.08%	0.58%
		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	0	0	100.22	111.548	0.89%	0.16%
		tape	1.292	13.604	18.61	103.601	217.604	434.597	14.78	82.571	252.286	634.373	6.34%	2.91%
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	100	200	0	0	0	0	100	200	1.3%	0.28%
		unknownArch	0.00	0.00	0.00	0.00	23	220	0.00	0.00	23	220	0.11%	0.33%
Table Summaries														
	VO Totals (TB)	disk	172.469	512.907	246.082	537.255	24.718	437.158	160.44	471.402	602.709	1958.722	1.49%	1.22%
		multidisk	0.078	0.528	13298.226	27058.759	1319.304	43403.13	0.076	1000.703	14617.884	71463.118	36.2%	44.39%
		tape	501.292	1515.644	1200.775	9515.305	1217.604	1434.597	1059.532	9340.087	3979.203	21805.633	9.85%	13.54%
		unknownArch	53.589	157.498	17452.185	54434.351	1074.804	3372.071	2801.604	7802.071	21181.962	65765.988	52.45%	40.85%
		TOTAL	727.427	2186.572	32196.248	91545.669	3636.23	48646.957	3821.652	18614.262	40381.557	160993.461		
	Percentage	disk	28.62%	26.19%	40.66%	27.43%	4.1%	22.32%	26.62%	24.07%				
		multidisk	0%	0%	90.97%	37.86%	9.03%	60.74%	0%	1.4%				
		tape	12.6%	6.95%	30.18%	43.64%	30.6%	6.58%	26.63%	42.83%				
		unknownArch	0.25%	0.24%	82.39%	82.77%	5.07%	5.13%	12.28%	11.86%				
		TOTAL	1.8%	1.36%	79.73%	56.86%	9%	30.22%	9.46%	11.56%				

WLCG Accounting Summary - 2008

			Jan-08	Feb-08	Mar-08	Apr-08	May-08	Jun-08	Jul-08	Aug-08	Sep-08	Oct-08	Nov-08	Dec-08	Total TB-months
Disk Space - TBytes	ALICE	Allocated	1340	861	4500	17178	12884	1105	0	0	0	0	0	0	37868
		Used	1043	573	1500	8589	6442	149	0	0	0	0	0	0	18296
	ATLAS	Allocated	1644	1133	5188	25768	25768	8124	0	0	0	0	0	0	67625
		Used	945	845	2119	12884	12884	3542	0	0	0	0	0	0	33219
	CMS	Allocated	1055	756	2764	21474	21474	5945	0	0	0	0	0	0	53468
		Used	649	608	1246	10737	10737	2371	0	0	0	0	0	0	26348
	LHCb	Allocated	608	397	3519	17178	17178	3061	0	0	0	0	0	0	41941
		Used	311	116	1509	8589	8589	649	0	0	0	0	0	0	19763
	TOTAL	**Allocated**	4647	3147	15971	81598	77304	18235	0	0	0	0	0	0	200902
		Used	5896	4284	12748	81598	77304	13422	0	0	0	0	0	0	195252

			Jan-08	Feb-08	Mar-08	Apr-08	May-08	Jun-08	Jul-08	Aug-08	Sep-08	Oct-08	Nov-08	Dec-08	Total TB-months
Tape Space - TBytes	ALICE	Used	0	0	0	2147	2147	500	0	0	0	0	0	0	4794
	ATLAS	Used	500	500	500	2147	2147	500	0	0	0	0	0	0	6294
	CMS	Used	500	500	50	2147	2147	500	0	0	0	0	0	0	5844
	LHCb	Used	500	500	500	2147	2147	500	0	0	0	0	0	0	6294
	TOTAL	**Used**	1500	1500	1050	8588	8588	2000	0	0	0	0	0	0	23226

The Storage Resources

- ***Use the GlueSA class***
  - It describes a Storage Area: a logical piece of space
    - which can include disks and tapes.
  - *Shared spaces are published as a Storage Area with multiple Access Control Rules*
  - *Express type of space through Capabilities (read-only, wan, lan, etc.)*
  - *GlueSAReservedOnlineSize*
    - *Space physically allocated to a VO (or a set of VOs) – Installed capacity*
  - *GlueSATotalOnlineSize (in GB = 10^9 bytes)*
    - *Total Online Space available at a given moment (it does not account for broken disk servers, draining pools, etc.)*
  - *GlueSAUsedOnlineSize (in GB = 10^9 bytes)*
    - Size occupied by files that are not candidates for garbage collection – *Resources usage*

Status and plans

- ***Storage***
  - Dynamic information providers for ***CASTOR*** already available by J. Jensen
    - Corrections needed to comply with what agreed
    - Packaging and distribution effort through the CASTOR CVS
    - First installation at CERN
    - Availability: end of July 2008 ?
  - Dynamic information providers for ***dCache*** available by R. Trompert
    - Need to verify with dCache developers that pinned files usage info is available, and other details
    - 2-3 weeks to implement the proposal
    - The changes will be reflected also in the new official dCache information providers by P. Millar
  - Dynamic information providers for ***DPM*** available by Michel Jouvin
    - The proposal already implemented
    - Installed at Edinburgh. – Testing phase started
  - Information already available for ***StoRM*** by the developers team
    - Probable minor additions for VOInfoPath
  - ***Work in progress with OSG and NorduGrid***

The Computing Resources

- Expressed in terms of KSI2000 per federation in current Megatable
- APEL provides already the information about *resources usage*
- Ongoing work to make information published about *installed capacity* more reliable (see presentation of Steve Traylen at GDB)
- *GridMap already provides the information per region. Ongoing work (in collaboration) to provide a view per federation and an interface to retrieve data.*
- ***Proposal:*** *Publish static pledges and the dynamic view in the current Megatable format (by federation) – provide APEL plug-ins*
  - *It allows for easily spotting discrepancies*
  - *Profit of the presence of a technical student who will be at CERN till the end of August.*

Computing Resources Example

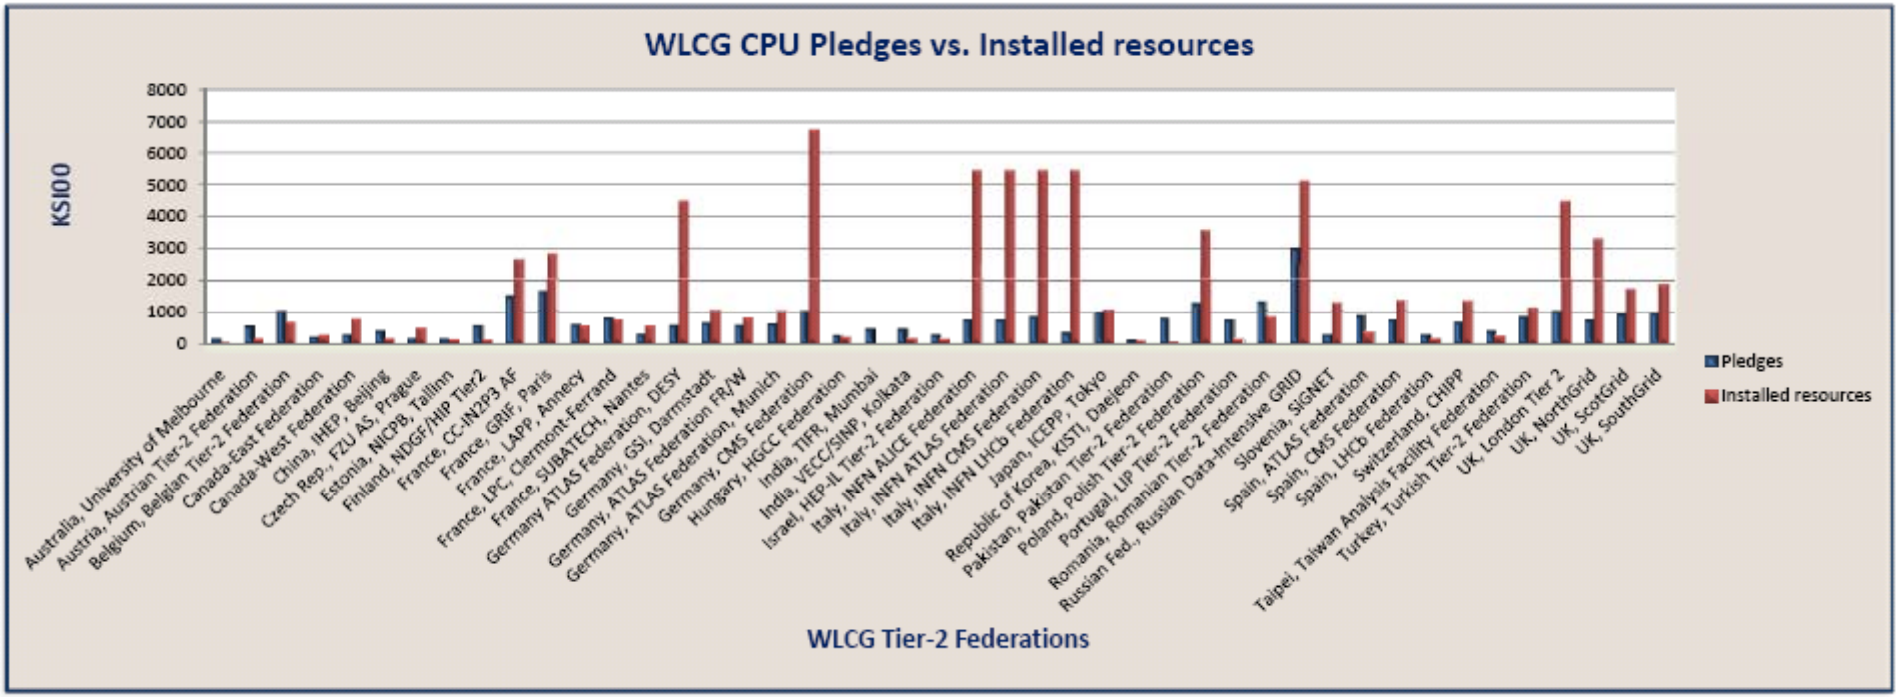

- Need to interpret results
- Refinement and cleaning (HEP shares, cpus vs cores, job slots, etc.) – Follow up with sites
- Improve information providers (Glue 2.0 helps)

Status and plans

- ***Computing resources***
  - First procedures in place. Looking for better technologies.
  - We are currently in the process of understanding and refining current information
  - Work with sites to make information more accurate
  - Check that what is needed will be covered by Glue 2.0
  - Help automate the publishing process (work with Steve Traylen)
  - Work with OSG and NorduGrid
  - Provide a working prototype for APEL by the end of August 2008
  - Give feedback to the MB and help producing monthly reports

Side effects

- Static information
  - Pledges
  - Federation – Site names association
- Following recommendations from Monitoring Group
  - Experiment django project
    - ***http://www.djangoproject.com/***

Thank You

WLCG Management Board,
CERN 8 July 2008