Automatic Resource & Usage Monitoring

Steve Traylen/Flavia Donno

CERN/IT

Outline

- Motivation
  - MOU Pledges and Reality
  - Better Understanding of Current State
- Measurement
  - Computing Resources
  - Storage Resources
- Validation
  - GridMap
  - GStat
  - Grid Configuration Monitoring
- Time Lines

Motivation

- Current Amount of Resources
  - More precise reliable measurements at any time of:
    - Storage allocation and use
    - Computing allocation and use
  - To be broken up by experiment, site, WLCG federation.
  - Information is needed by Grid and Experiment Operations
    - Individual sites must be able to see measurements of themselves.
    - Operators need to react to under/over productive sites.
- The WLCG resource pledges must be validated.
  - The WLCG MB needs to know these commitments are being met on a month by month basis.
  - Summary reports to be produced.

MOU Pledges

- Computing and Storage Pledges
  - Pledges are recorded in MOUs for T0,1 and 2 federations.
  - Today T0 and 1s submit their resources figures to WLCG MB for comparison.
    - Not possible or desirable for T2s given the number of them.
- Machine Readable Pledges
  - Computing now available. http://gridops.cern.ch/mou/
  - Storage will be added?
  - XML or JSON is available.

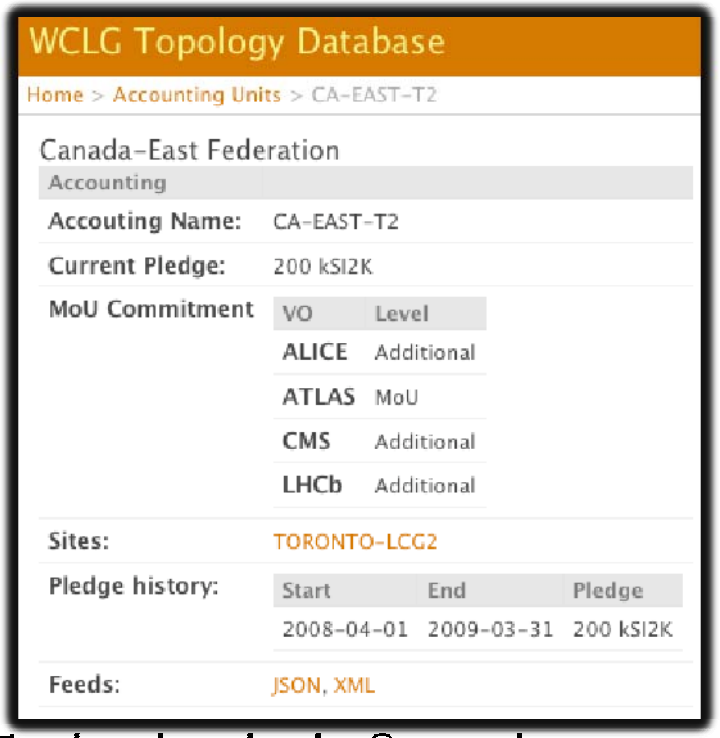

Publishing Compute Resource

- EGEE and OSG published a lot anyway.
  - We are just cleaning and disambiguating a few items.
- For Instantaneous Capacity
  - EGEE will publish required values.
  - OSG will publish required values
    - Extra requirement that endpoints are not necessarily published.

$$TotalCompute\ (SiteX, ExperimentY) = \frac{FS_Y}{100 * 1000} * \sum_{SC \in SubClusters_X} HostBenchMark_{SC} * NumberLogicalCPUs_{SC}$$

- Definitions and Clean Up.
  - FS_Y = % fair share of site to experiment Y
  - Benchmark is that of a LogicalCPU (a core).
  - An extra attribute added: ScalingReference
    - Allows actual and normalized SI00 to be split.

Publishing Storage Resource

- Storage reports static numbers used by management to collect info upon pledged resources (InstalledCapacity).
- Dynamic collection of resources used by the VOs (Total/Free/Used[Online|Nearline]Size).
- Measurements of disk and tape resources are made.
- In the Information Systems all storage areas used by the experiments (and not for internal optimization, such as replication pools) are described, including disk cache in front of tapes.
- Still discussions about how to make reports about storage.
  - What should be counted as pledged or not.

Passing Resources to the MB

- EGEE
  - The information systems are probed a few times a day.
  - Results stored at the CESGA accounting portal.
  - Summary reports are generated and presented to MB.
- OSG
  - OSG will calculate summary results.
    - Exact method to be determined.
  - Summary data will be presented to the MB.
- Monthly report will contain:
  - Maximum Resource Provision.
  - Average Resource Provision.

Map

CERN-PROD

	Total	Physical	Logical	
Installed Capacity (WLCG):			**1856**	**(3567 kSI2k)**
ce127.cern.ch:	3608			(1922 SI2k/CPU)
ce106.cern.ch:	3608			(1922 SI2k/CPU)
ce128.cern.ch:	464			(1922 SI2k/CPU)
ce103.cern.ch:	3608			(1922 SI2k/CPU)
ce126.cern.ch:	3608			(1922 SI2k/CPU)
ce111.cern.ch:	3608			(1922 SI2k/CPU)
ce105.cern.ch:	3608			(1922 SI2k/CPU)
ce109.cern.ch:	3608			(1922 SI2k/CPU)
ce110.cern.ch:	464			(1922 SI2k/CPU)
ce104.cern.ch:	3608			(1922 SI2k/CPU)
ce108.cern.ch:	3608			(1922 SI2k/CPU)
ce107.cern.ch:	3608			(1922 SI2k/CPU)
ce115.cern.ch:	3608			(1922 SI2k/CPU)
ce114.cern.ch:	3608			(1922 SI2k/CPU)
ce112.cern.ch:	3608			(1922 SI2k/CPU)
ce125.cern.ch:	3608			(1922 SI2k/CPU)
ce124.cern.ch:	3608			(1922 SI2k/CPU)
ce113.cern.ch:	3608	1856	1856	(1922 SI2k/CPU)
ce129.cern.ch:	464			(1922 SI2k/CPU)

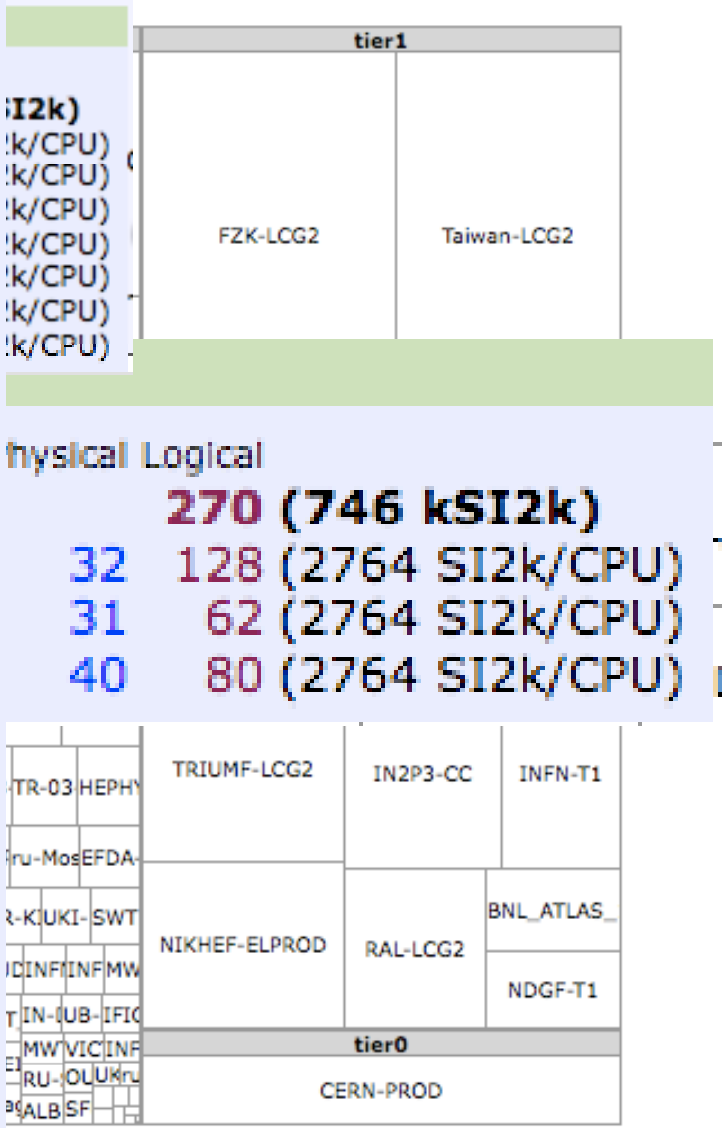

GridMap – Strict View

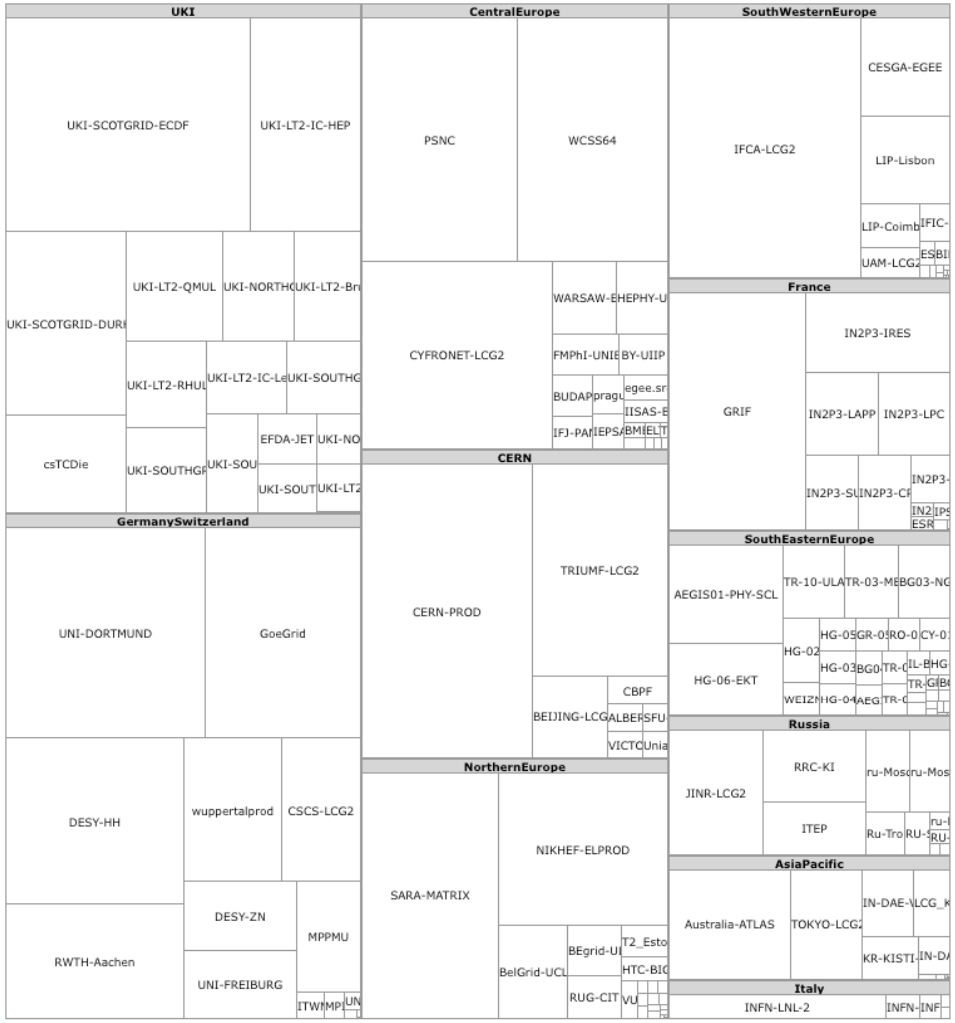

- GridView has a strict to WLCG Capacity mode since 2 days ago.
- Work to be done:
- Italy is tiny.
  - Sites publish 0 logical CPUs.
- Scotgrid-EDCF huge.
  - Double publishing.

GridView – Strict View

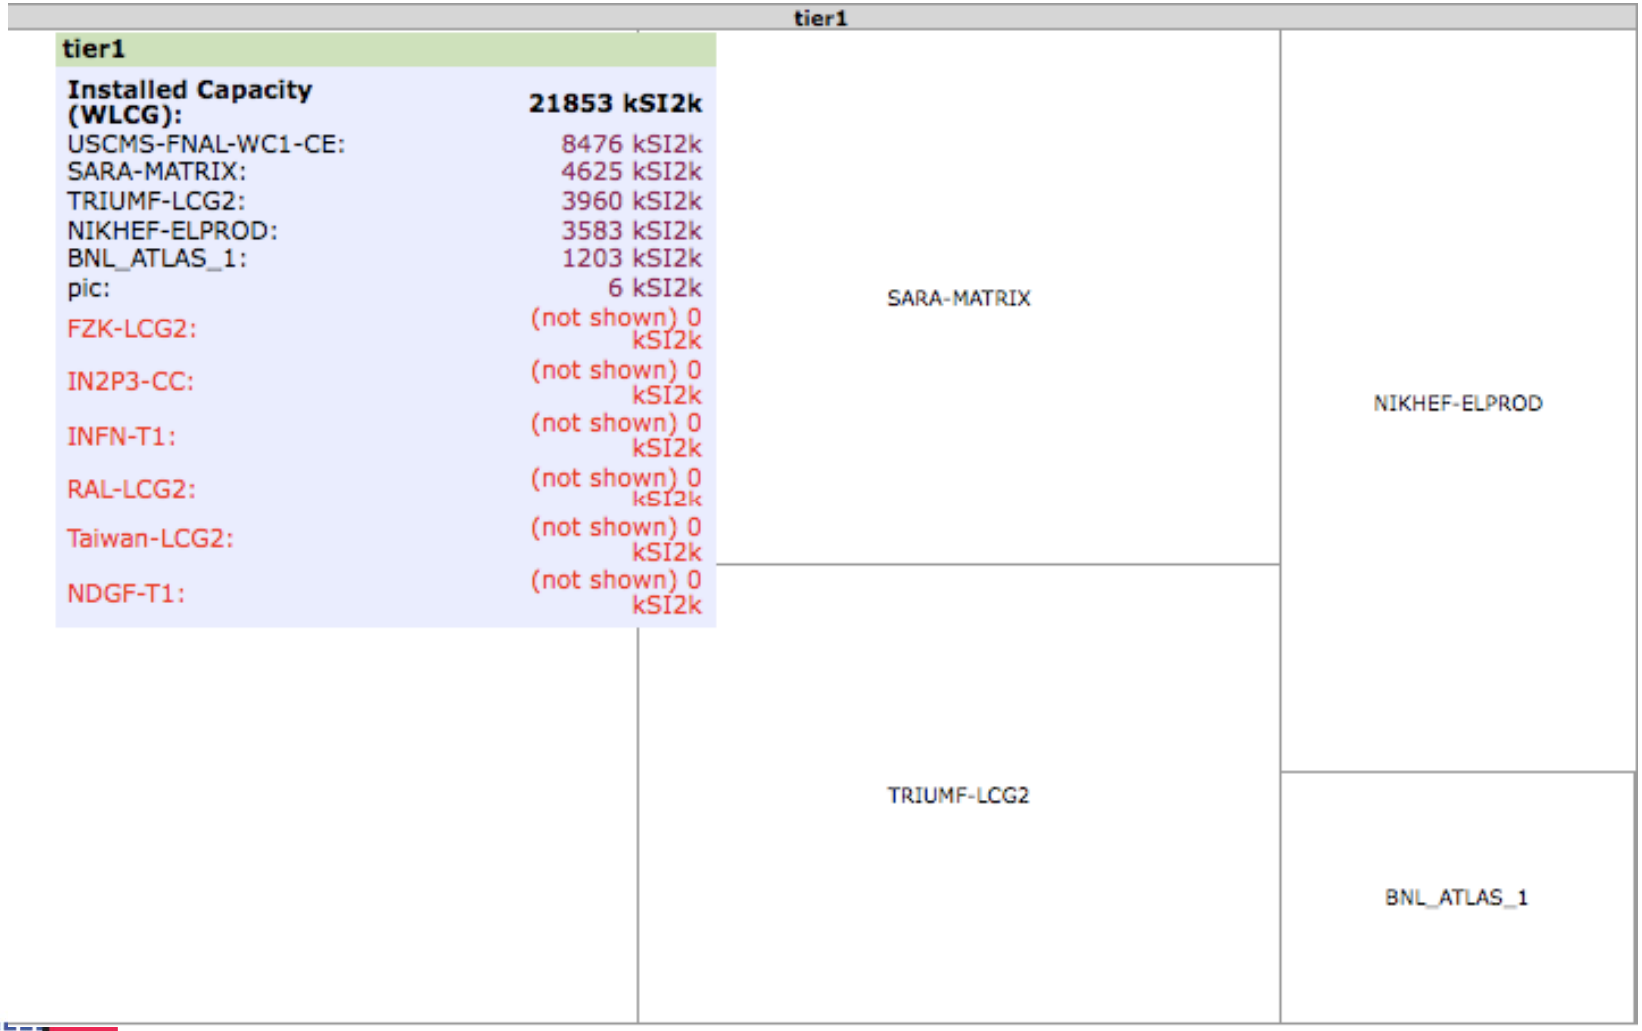

Validation – Gstat

- Gstat as always performs sanity checks.
  - Is there > 1 Logical CPU
  - Physical <= Logical
  - …..
- Like all gstat checks these are now nagios probes.
  - Will be deployed with EGEE SA1 Nagios.
  - Will be evaluated by OSG to run in their framework.
- A dozen or so GGUS tickets already submitted based on results.
  - Most sites were able to correct their information.
  - Some require more information.
- Subsequently gstat needs correcting, free > total?.
  - Sites please bare with us, there will be a transition phase.

Validation - GCM

- Grid Configuration Monitoring - real measurements.
  - Per WN measurements of:
    - SubCluster Membership of WN, Processor Vendor and Model, Number of cores, Memory, local disk space, memory, OS, ..
  - First time the SubCluster Information can be validated
    - Logical CPUs per SubCluster can be compared against published values.

- GCM data would be used for evaluating BDII published data in the first instance.
- No plans now as to when this might be.

GFAL Version in Site UKI-LT2-RHUL

no	hostname	status	result	latest update
1	node020.cvos.cluster	successful	1.10.18	27-Jan-2009 16:45:28
2	node023.cvos.cluster	successful	1.10.18	22-Jan-2009 15:25:14
3	node025.cvos.cluster	successful	1.10.18	22-Jan-2009 15:42:15
4	node042.cvos.cluster	successful	1.10.18	28-Jan-2009 14:12:44
5	node046.cvos.cluster	successful	1.10.18	28-Jan-2009 11:40:26
6	node050.cvos.cluster	successful	1.10.18	27-Jan-2009 16:46:40
7	node13.beowulf.cluster	failed		20-Jan-2009 10:21:35
8	node64.beowulf.cluster	failed		29-Jan-2009 09:38:11
9	node68.beowulf.cluster	failed		22-Jan-2009 15:26:51
10	node70.beowulf.cluster	failed		20-Jan-2009 16:19:20
11	node73.beowulf.cluster	failed		26-Jan-2009 10:37:42

Timeline

- YAIM Release – In certification now.
  - Note, sites can configure non-zero capacity now.
- OSG Release – On Integration TB now.
  - Wider deployment during May.
- Sites and Regions to check “strict” view of GridMap.
  - General support and site specific help will be given.
  - Ask GGUS for help.
- First preliminary EGEE and OSG reports
  - July , April. To be considered with WLCG federations first.
- At some point during this:
  - New info providers for Xroot, CASTOR, dCache.

Summary

- Strict view in GridMap allows regions and sites to get a good view of their measured capacity now.
  - Sites should correct as far as possible.
- Reference is the Install Capacity Document.
  - Main Editor – Flavia Donno.
  - https://twiki.cern.ch/twiki/pub/LCG/WLCGCommonComputingReadinessChallenges/WLCG_GlueSchemaUsage-1.8.pdf
- Compute resource is well defined, info providers done, validation well under way, reporting understood.
- Storage most providers are complete, and are being deployed. Data being collected.
  - More visualization will be needed , e.g Storage GridMap.